ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOAG001

# STATUS OF THE APS ACCELERATOR CONTROLS FOLLOWING ITS MAJOR UPGRADE*

G. Shen† on behalf of APS Controls Group, Argonne National Laboratory, Lemont, IL, USA

*Abstract*

The APS Upgrade replaced the original storage ring with a fourth-generation multi-bend achromat lattice and modernized its accelerator controls while retaining the injector systems. The upgraded controls integrate a high-speed network, enhanced timing and event distribution, advanced beam diagnostics, and time-correlated data acquisition within the EPICS framework. These improvements enabled a rapid transition from decommissioning to beam commissioning and the restoration of user operations within one year. This paper summarizes the current status of the APS accelerator controls, highlights key system upgrades, and presents lessons learned that may serve as a reference for future control system modernization.

## INTRODUCTION

The Advanced Photon Source (APS) [1] at Argonne National Laboratory recently completed a major upgrade, transforming its third-generation 7 GeV storage ring into a fourth-generation 6 GeV multi-bend achromat light source [2]. This upgrade doubles the beam current to 200 mA and reduces the emittance from 3100 pm rad to 32–42 pm rad, while preserving the original 20 double-sector configuration and 1.1 km circumference.

As part of this upgrade, the accelerator control system was modernized for the new storage ring, while the injector controls were largely retained for compatibility. A new software and hardware infrastructure, integrated with the EPICS [3, 4], was deployed to support high-speed data acquisition, precision timing, enhanced machine protection, and many other subsystems.

This paper presents the current status of the APS control system, outlines major hardware and software upgrades, and discusses lessons learned during the transition from decommissioning to beam commissioning [5].

## CONTROLS SYSTEM OVERVIEW

The upgraded APS control system builds upon the EPICS framework, originally developed at APS in collaboration with LANL [6, 7], and preserves its well-established three-layer architecture: front-end electronics, EPICS IOCs, and user interfaces, as shown in Fig. 1. This structure has been enhanced to provide a highly integrated platform that spans all technical subsystems within the facility, ensuring seamless operation across both the upgraded storage ring and the retained injector controls.

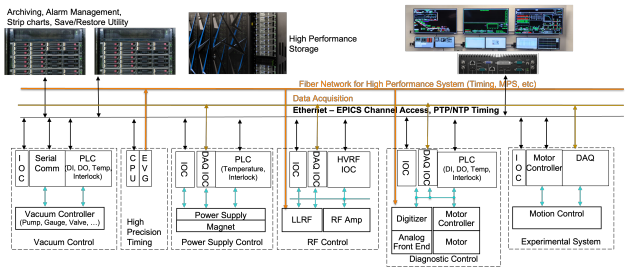



Figure 1: APS control system: three-layer architecture.

In the original APS configuration, most equipment was controlled through a variety of VME-based systems, utilizing standard I/O components such as VME IP carriers, IP modules, and transition cards, supplemented by custom-built VME modules. Over time, a phased approach transitioned many of these controls to more modern hardware platform like Programmable Logic Controller (PLC)-based solutions, particularly for systems like vacuum control, where improved reliability and maintainability were required.

As a pioneer in the EPICS community, APS developed numerous client tools, including MEDM (Motif Editor and Display Manager) [8] and ALH (Alarm Handler) [9], which remain actively used today at APS to maintain backward compatibility with unchanged subsystems such as the injector and storage ring RF systems. The upgraded control system balances modernization with continuity, ensuring that legacy tools and interfaces continue to function alongside newly implemented high-performance hardware and software components.

## CONTROL SYSTEM REMOVAL

During the APS Upgrade, portions of the original control system were carefully decommissioned while preserving critical infrastructure, particularly global timing fibers and other critical cables. A thorough pre-shutdown assessment was conducted to segment, label, and document all necessary hardware and connections. This preparation enabled the safe removal of obsolete components and facilitated the replacement of aging cables, many of which had been in service since the facility's initial installation.

As a result, no damage occurred to critical systems during the removal phase. More than 400 hardware components were harvested, with roughly 280 refurbished and reinstalled to support the upgraded accelerator.

## HARDWARE SYSTEM MODERNIZATION

### *Network Structure*

The APS-U control system incorporates a redesigned three-tier network architecture to meet the demands of high-

* The work is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

† gshen@anl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOAG002

# DESIGN AND DEVELOPMENT OF DIAMOND-II ACCELERATOR CONTROL SYSTEM

S. Singh*, M. Heron, P. Hamadyk, M. Abbott, A. Wells, R. Mercado, L. Hudson, K. Baker, G. Christian, M. Gaughran, M. Stubbings, A. Alexander, A. Cousins, C. Colborne
Diamond Light Source, Didcot, United Kingdom

*Abstract*

Diamond Light Source is a 3rd-generation synchrotron light source that has been operating since 2007. The existing accelerator control system is based on EPICS version 3, and uses a mixture of VME hardware, PCs and embedded devices. An upgrade of Diamond to Diamond-II is now in the construction phase with installation set to begin in January 2028, followed by storage ring commissioning in Oct 2028. A new control system is currently under development, using existing infrastructure while modernizing key components. The updated control system will be built on EPICS 7 with software deployed via Kubernetes clusters. This paper outlines the system requirements, development activities, planning, and deployment strategy for the Diamond-II accelerator control system.

## INTRODUCTION

The Diamond-II project is a major upgrade of the UK's national synchrotron light source, introducing a new machine, advanced beamlines, and improvements to optics, detectors, sample handling, delivery, and computing. This initiative will boost Diamond Light Source's scientific capabilities and strengthen its status as a leading synchrotron research facility [1, 2].

Diamond-II key operational parameters:

- Beam energy: 3.5 GeV.
- Beam current: 300 mA.
- Horizontal emittance: $\leq$ 120 pm rad.
- Vertical emittance: 8 pm rad.
- Beam stability: $\leq$ 3% (1–1000 Hz).
- Top-up transparency: $\leq$ 1% flux drop over 100 ms.
- In-vacuum ID gap (min): 4 mm.

The Diamond-II control system handles supervisory control, synchronization and timing, machine protection, safety, and archiving of machine variables for diagnostics and improvement purposes. A well-structured and dependable control system contributes to beam availability and the enhancement of machine performance.

A lot of the current Diamond control system [3] is based on VME based systems operating on VxWorks and EPICS version 3. Diamond-II will retain a small number of the existing VME systems, but VxWorks will be replaced with RTEMS on all of them. Most systems will be upgraded with new hardware and software system. Most of the existing network infrastructure will be used in Diamond-II. Kubernetes will be deployed for the management of most control systems and IOCs.

* sudheer.singh@diamond.ac.uk

## REQUIREMENTS

The Diamond-II control system is designed to support a modular, scalable, and interoperable architecture that meets the evolving needs of accelerator development, commissioning, and operation. It must accommodate legacy systems like VME with RTEMS while enabling new infrastructure across the machine. Key functional requirements include extendibility, portability, and scalability, ensuring control modules can operate across diverse platforms and scale with demand. The system must support multiple user groups including operators, physicists, system experts, beamline users, control engineers, and facility managers, each with distinct needs such as intuitive GUIs, detailed diagnostics, remote access, synchronisation between machine and beamline controls, and safe operational protocols. These requirements ensure robust, flexible operation of Diamond-II Functional Requirements. The control system for Diamond-II must support:

- Extendibility and Interoperability: Modules across diverse infrastructures must interact seamlessly.
- Portability: Modules should run on varied platforms.
- Scalability: Systems must scale with operational demands.
- Timeline Compliance: Development must align with the Diamond-II project schedule, see Fig. 1.

The functional requirement has been grouped in to three views User View, Physical View and Deployment View of the system.

### *Use Case Requirements*

The use case view of the Diamond-II control system is a user-oriented design approach [4] that prioritises the needs and interactions of various user groups over the physical deployment or software architecture. This ensures that the control system is not only technically robust but also functionally aligned with the workflows and expectations of its users. From this perspective, the system must support:

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOAG003

# UPGRADE AND MODERNIZATION OF CSNS ACCELERATOR CONTROL SYSTEM TOWARDS CSNS-II*

Y. L. Zhang[†,1,2], P. Zhu[1,2], X. H. Lu[1,2], X. Wu[1,2], Y. C. He[1,2], L. Wang[1,2], K. J. Xue[1,2], M. T. Li[1,2], S. N. Cheng[1,2]
[1]Institute of High Energy Physics, Beijing, China
[2]Spallation Neutron Source Science Center, Dongguan, China

*Abstract*

The CSNS-II project, launched on 1st January 2024, aims to significantly enhance the beam power from 100 kW to 500 kW. The current accelerator control system, commissioned in 2018, was designed based on hardware and software platforms finalized in 2012. Over time, these systems have begun to exhibit obsolescence issues. To meet the advanced requirements of CSNS-II, a comprehensive upgrade and modernization of the control system is essential. This presentation outlines the overall upgrade plan and design considerations for the control system. The key upgrades include: transitioning the EPICS framework from version 3 to the modern and feature-rich version 7, migrating the hardware platform from VME to MTCA, adopting the latest Phoebus as part of the Control System Studio suite, incorporating support for big data analytics and artificial intelligence capabilities to enhance system performance and diagnostics. These enhancements will ensure the control system meets the demanding operational requirements of CSNS-II while improving reliability, scalability, and future readiness.

## BRIEF INTRODUCTION OF CSNS AND UPGRADE PROJECT CSNS-II

The China Spallation Neutron Source (CSNS) is designed and constructed by the Institute of High Energy Physics, Chinese Academy of Sciences. The construction of CSNS includes an 80-MeV Linac, a 1.6-GeV Rapid Cycling Synchrotron (RCS), two beam transport lines, a solid target station of 100 kW, three initial neutron instruments and other utility facilities[1].The CSNS-II's plan for future improvements to the existing accelerator complex aims to deliver a beam with a minimum power of 500 kW for the target station. The schematic layout of CSNS-II upgrade project is shown in Fig. 1

The CSNS-II Linac will utilize multiple 324 MHz pulsed double spoke cavities and 648 MHz elliptical 6-cell pulse-capable cavities to achieve a peak beam current of up to 40 mA within the energy range of 80–300 MeV. The double spoke cavity portion should accelerate from 80 MeV to 165 MeV using 20 cavities in 10 cryomodules. The elliptical cavity portion of the Linac should accelerate from 165 to 300 MeV using 24 elliptical dressed cavities in 8 cryomodules [2]. The schematic layout of upgraded Linac in CSNS-II is shown in Fig. 2.

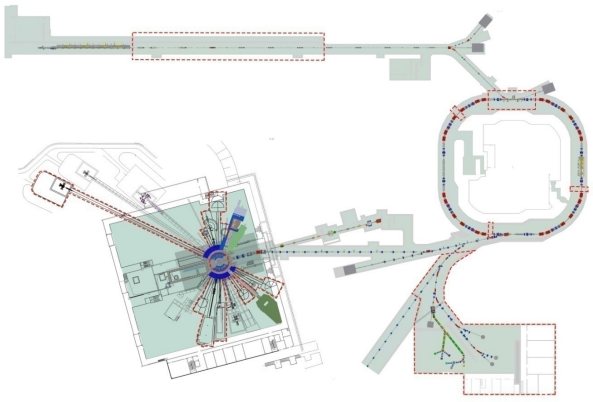


Figure 1: Schematic layout of CSNS-II.

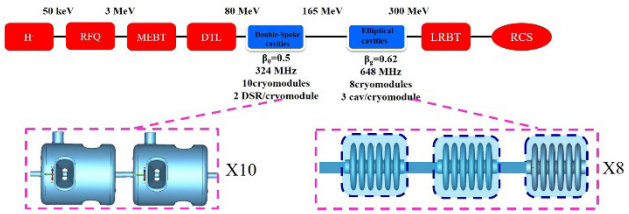



Figure 2: Schematic layout of CSNS-II Linac.

## OVERALL DESIGN OF CSNS-II ACCELERATOR CONTROL SYSTEM

### *Main Tasks and Requirements of Control System*

The primary task of the CSNS-II accelerator control system is to monitor, control, and protect the various equipment distributed along Linac, RCS and the beam transport lines. It provides a hardware and software platform for global control and information interaction, enabling operators, physicists, and system experts to stay informed in a timely manner and to monitor and operate the equipment remotely from the control room via human-machine interfaces (HMIs). Furthermore, the accelerator control system must provide global machine mode, operation mode, and beam operation management functions. Fig. 3 depicts the main functions of accelerator control system.

The main tasks required for the CSNS-II accelerator control system upgrading and modernization include the following aspects:

- Expand the timing subsystem and implement a globally high-precision synchronous timestamp based on the timing system.
- Optimize and update the machine protection system logic.

[†]zhangyl@ihep.ac.cn

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOAG004

# SLS 2.0 BEAMLINE UPGRADE EXPERIENCE: NAVIGATING MODERNIZATION, LEGACY, AND COMMISSIONING CONSTRAINTS

T. Celcer†
Paul Scherrer Institute, Villigen, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

After successfully reaching the key milestones of the SLS 2.0 machine upgrade, focus and prioritization have shifted to the beamline upgrades. These have been structured into three distinct phases. In the "pre-dark time" Phase 0, core technical solutions and controls hardware portfolio were validated on selected beamlines. We are currently finalizing the upgrades for Phase 1 beamlines, while preparations for Phase 2 — scheduled for 2026 — are gradually ramping up. In this contribution, we first reflect on the Phase 1 commissioning deliverables and assess how far we could implement the originally proposed control system upgrade strategy, particularly regarding hardware modernization and the coexistence with legacy components. Second, we analyze the impact of resource and time constraints on our commissioning activities. Delays in the prerequisite steps for control system commissioning (device list provision, schematic design, hardware assembly, testing, installation and cabling) largely due to limited capacity in infrastructure groups, ultimately resulted in a significantly compressed commissioning window during a critical project phase. We discuss adopted mitigation strategies - including pre-commissioning using test systems, solution standardization efforts, task prioritization driven by the minimum viable product (MVP) delivery, and strong cross-team coordination activities. These insights offer practical lessons for managing the coming SLS 2.0 Phase 2 beamline upgrades.

## INTRODUCTION

The Swiss Light Source (SLS) upgrade began in October 2023, following more than two decades of successful operation since first light in 2001. At its core, the project focused on replacing the storage ring to deliver fourth-generation performance [1]. Remarkably, the first electrons circulated in the new ring only 15 months later, in January 2025. Since then, attention has shifted to the upgrade of the beamlines.

Prior to the project, SLS operated 18 user beamlines. Within SLS 2.0 upgrade, two additional beamlines are being constructed (Debye and I-Tomcat), one beamline (IR) was decommissioned, while three existing ones (PEARL, microXAS, QUEST) will be relocated to new sectors.

The beamline upgrades were structured into three phases. Phase 0, completed with the start of the dark time in October 2023, included the partial commissioning of the new Debye beamline, the upgrade of the PX III crystallography beamline, and selected system upgrades elsewhere. This phase also served to benchmark the new controls hardware portfolio, particularly the motion platform, thereby validating strategic choices and establishing the basis for subsequent phases. Phase 1, now nearing completion, covers the commissioning of nine existing beamlines and the new I-Tomcat beamline. In total, nine beamlines have already received first light in their end stations, with three more expected to follow shortly. Phase 2, scheduled for 2026, will include a three-month shutdown dedicated to commissioning the remaining beamlines, which are set to receive first light in the second half of 2026. Due to budget constraints, two additional beamlines will be upgraded only after the project concludes, from 2027 onward.

## REVIEW OF CONTROL SYSTEM UPGRADE STRATEGY

After more than two decades of operation, the SLS faced significant technical debt, including reliance on aging VME hardware and fragmented higher-level experiment control, alongside growing user demands for complex workflows such as fly scanning. Consequently, upgrading the control system was identified as a key deliverable to fully exploit the improved beam characteristics at the experimental stations, enhance the user experience, and support scientific productivity. To address this, a comprehensive control system upgrade strategy was developed [2]. In this paper, we will focus primarily on aspects related to the hardware portfolio upgrade, its migration away from VME, and selected control system software developments.

### *Modernizing Hardware: Moving Beyond VME*

The central part of the control system upgrade strategy was the migration away from the aging VME platform, which no longer met modern requirements. Continued reliance on VME was also impractical due to several other factors, including uncertain licence support and unpredictable licensing costs for VxWorks 6, the inability to migrate EPICS to versions 7.1 and beyond because of VxWorks 6 limitations, and the incompatibility of existing MVME CPU cards with VxWorks 7. Continuing to depend on the VME platform at the upgraded SLS was therefore not an option.

Table 1 presents the curated hardware portfolio selected to replace the VME platform. Since motion control constitutes the majority of the beamline system, selecting the next-generation motion platform was critical. We chose an EtherCAT bus system with Beckhoff hardware modules, controlled via the EtherCAT Motion Controller (ECMC) software platform [3, 4]. While ECMC is primarily a motion control solution, leveraging industry-standard Beckhoff modules also allows support for lower- and mid-tier ADC/DAC applications within the EtherCAT bus limits.

† tine.celcer@psi.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOAG005

# UPGRADING FERMILAB'S ACCELERATOR CONTROL SYSTEM WITH ACORN*

C. Roehrig†, P. Dobbin, J. Case, A. Tiradani, B. Harrison, Fermilab, Batavia, IL, USA

*Abstract*

The Fermilab Accelerator Complex is the largest national user facility in the Office of High Energy Physics (DOE/HEP) program and the only national user facility operating at Fermilab. Fermilab serves as the host to the Long Baseline Neutrino Facility/Deep Underground Neutrino Experiment (LBNF/DUNE), the laboratory's flagship project for neutrino science that is under construction. LBNF/DUNE will be powered by megawatt beams from an upgraded accelerator, the Proton Improvement Plan II (PIP-II) that will replace the laboratory's aging linear accelerator with a new one based on superconducting radio-frequency cavities. The Accelerator Controls Operations Research Network (ACORN) Project will support LBNF/DUNE and PIP-II by modernizing the accelerator control system. The project is at the conceptual design phase and looking to achieve Critical Decision 1 (CD-1) later this year. The scope and structure of the project will be presented, along with an overview of how that has changed in the past year. Current design and technology choices will be shared. Specific challenges facing the project will be addressed, along with current thinking on solutions.

## INTRODUCTION

The Fermilab accelerator complex consists of several accelerators, all connected by beam transfer lines. Beam is extracted from the ion source and delivered to the Linac. The Linac accelerates the beam to 400 MeV before sending it onto the Booster. In the Booster, proton energy is increased to 8 GeV. From there, beam can be sent to either the Recylcer ring or the Main Injector. The Main Injector accelerates beam to 120 GeV before delivering protons to the neutrino experiment, NuMI or other fixed target areas. ALong the way, protons can be sent to the 400 MeV Test Area, the Booster Neutrino Beam experiments, and the Muon Campus.

The Fermilab accelerator control system, generally known as ACNET [1], was originally developed to support the start of the Tevatron Collider operations in 1983. At the time, it was a large improvement, providing a single control system for what had been separate subsystems. The design implemented a three tiered architecture, with Front Ends interfacing to hardware, User Applications providing users with access to the control system, and a Central Services tier in the middle, providing support to both applications and front ends.

The Department of Energy is funding two large projects at Fermilab. The Proton Improvement Plan Two (PIP-II) [2] will replace the existing linac with a new 800 MeV, 215 meter linac, capable of delivering over 1 megawatt of continuous beam power to the Booster. At the other end of the accelerator complex, the Long Baseline Neutrino Facility (LBNF) [3] will use the 120 Gev protons from the Main Injector to generate a beam of neutrinos. That beam will travel through the ground to the Deep Underground Neutrino Experiment (DUNE), located in Lead, South Dakota.

To support both of these projects, the DOE is also funding the ACORN project. This project will provide needed modernization of the accelerator control system. This modernization will focus on four main areas.

- Develop and test new control electronics for magnet and RF bias power supplies.
- Design and install control electronics to replace existing CAMAC and rack monitor systems.
- Implement a controls software infrastructure that will enable real-time data acquisition, modern cybersecurity practices, and AI/ML capabilities.
- Develop and deploy modern user applications and an application framework to support future development.

## ACCELERATOR POWER SYSTEMS

As is true for any accelerator, magnets are used to define the path of the particle beam. Many of the electronics used to control the delivery of power to Fermilab's magnets use components that are past end of life and, in some cases, can no longer be built. In other cases, the performance of existing solutions does not meet the requirements of the LBNF/DUNE era. To address these problems, the ACORN project is designing new electronics.

### *Voltage Regulation*

The large magnets of the accelerator complex have their power delivered by power supplies that use a 12-Pulse rectifier topology (Fig. 1) to convert the 60 Hz 3-phase AC input power into a DC output with 720 Hz ripple (Fig. 2). The electronics used to regulate the voltage output of these power supplies are 40 years old. The components are unavailable, making building new voltage regulation chassis using the original design impossible.

The existing design uses 10 cards, shown in Fig. 3, each with a different function.

- System Monitor - provide diagnostics
- Voltage Monitor - analog readback
- Digital Program - converts 16-bit reference voltage to analog

* This work was produced by FermiForward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. Publisher acknowledges the U.S. Government license to provide public access under the DOE Public Access Plan.

† croehrig@fnal.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBG003

# NEW MACHINE PROTECTION SYSTEM AT THE SPALLATION NEUTRON SOURCE – DESIGN PROCESS AND PERFORMANCE ANALYSIS

M. Bobrek, H. Hao, K.-U. Kasemir, K. Vodopivec, T.A. Justice
Oak Ridge National Laboratory, Oak Ridge, TN, United States

*Abstract*

A New Machine Protection System (MPS) at the Spallation Neutron Source (SNS) was developed and implemented on µTCA-based hardware platforms. The system can monitor more than 2500 field inputs and shuts off the beam within 10 µs if adverse events occur. We will present system level design process of various firmware and software components as well as the system integration into EPICS environment. The performance analysis of the MPS after two SNS run cycles will also be presented.

## INTRODUCTION

The recent proton power upgrade at the Spallation Neutron Source elevated potential hazards of uncontrolled energy release and subsequent physical damage and radioactivation of the critical accelerator equipment [1]. This created a need to upgrade the existing Machine Protection System (MPS) based on obsolete VMA architecture. The new MPS is using µTCA hardware platforms including VT815 and VT819 from VadaTech® (Fig. 1). The new MPS provides increased channel capacity for additional monitoring points at the First Target Station as well as for the future Second Target Station. In addition, it offers shorter response time, a comprehensive health monitoring capabilities, improved controllability and observability, and a seamless integration in EPICS environment [2].

In the subsequent sections we will cover the MPS top-level architecture, communication links and messages, low latency solutions, and software integration. At the end, we will present the MPS performance assessment during the two run cycles since the system was deployed.

## MPS ARCHITECTURE

The main components of the new MPS include the central processor unit called MPS Master, the peripheral units called Field Nodes, and the fiber-optical interfaces connecting the MPS Master to the Filed Nodes (Fig. 2). The primary function of the MPS Master is to shut down the beam within 20 µs when a fault occurs. The qualifying faults include channel faults from the Field Node field inputs as well as the system faults coming from any of the hardware, firmware, or software components. The faults are reported via fault messages generated by the affected node. The messages are propagated from Node Processors to the MPS Master Trigger Controller module.

The MPS Master Trigger Controller module controls the beam delivery via three gate control signals connected to the ion source plasma RF generator, the RF generator for the first accelerating structure in the accelerator (RFQ), and the gating pulse for the LEBT chopper.

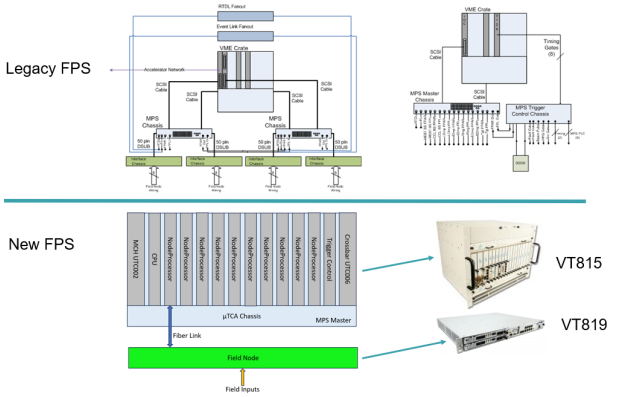



Figure 1: New vs legacy MPS.

To implement timestamping of critical events in MPS, the Master is receiving timing events from the Accelerator Timing Master and distributes the timing packets across the system.

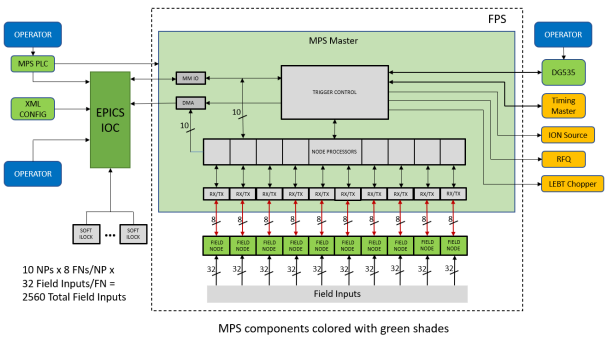



Figure 2: MPS architecture.

The MPS is interfacing with the EPICS IOC via PCIe bus on the VT815 backplane. The Memory Mapped IO (MMIO) and DMA channels exist for every Node Processor and the Trigger Controller. A separate crate module called µTCA Carrier Hub serves as the PCIe switch between the CPU and other crate modules. In addition, the IOC is connected to a number of soft interlocks that monitor various power supply health parameters. The response time for the soft interlocks is in order of few ms.

DG535 shown in Fig. 2 provides a hard upper limit for the current pulse duration and repetition rate in case the Trigger Controller sets those outside the safe limits. DG535 setup is done by the operator every time the accelerator running parameters change.

MPS PLC control from Fig. 2 is a safe, key-controlled way to shut down the beam via MPS Master.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBG004

# CURRENT STATE AND IMPROVEMENTS OF THE LUMINESCENT SCREEN PROTECTION SYSTEM IN THE KARABO CONTROL SYSTEM AT THE EUROPEAN XFEL

I. Karpics, A. Koch, A. Garcia-Tabares
European XFEL, Schenefeld, Germany

*Abstract*

During startup, maintenance, and the user program, stable beam conditions are crucial for smooth operation of beamlines and instruments at many light sources, including the European XFEL. For visual verification of the beam position and jitter, invasive luminescent screens are mostly used. These screens made of various materials are located at several locations along the beam path and can withstand a particular photon intensity and flux. Excessive intensity within a few x-ray pulses may damage the screen resulting in increased maintenance costs and operational down time. At the European XFEL, to protect these screens, a set of Karabo control system services are deployed. The main functionality of these services includes calculation of screen damage thresholds and the estimation of the actual photon flux. In the case of an overdose the user is warned, the screen is retracted and moved to a predefined safe position. In this contribution the current protection system and upcoming improvements of the optical luminescent screen integration in the Karabo control system at the European XFEL are presented.

## INTRODUCTION

The European Free Electron Laser Facility (EuXFEL), located in Schenefeld, Germany, provides the scientific community with experimental soft and hard x-ray experimental capabilities [1]. High-brilliancy x-ray radiation is generated in the self-amplified spontaneous emission (SASE) undulator systems and guided through the beam transport system to seven experimental stations.

In beam transport tunnels and instruments, more than 30 imaging units for beam diagnostics are installed [2]. These x-ray beam imaging units (further called imagers) are custom build vacuum chambers consisting of one or several luminescent screens, motorized stages to move them into the x-ray beam and a camera. Screens are typicaly made from Ce:YAG, polycrystalline diamond or pyrolytic-hexagonal-boron-nitrite ceramics (BN). Motorized stages are based on Beckhoff stepper motors and for image acquisition CMOS GigE cameras with mostly fixed focus are used. Imagers differ in geometry, resolution and most importantly their location, which define the limitations and functions of them. The following types of imager are available:

- Transmissive (TR) imagers containing the thinnest scintillators are closest to the Free Electron Laser (FEL) source, and, in correlation with other imagers, can be used to measure the beam trajectory;
- Spontaneous radiation (SR) imagers are used to detect spontaneous radiation from the undulators;
- FEL imagers are used to characterize the FEL beam by measuring beam position, size and shape with particular high spatial resolution;
- Pop-in (PopIn) imagers are available for all seven instruments and are used to locate the beam position and, in several cases, to provide feedback for steering the beam to a predefined accepted position;
- Exit slit (ES) imagers are mounted on the three exit slits downstream of the SASE3 monochromator;
- Other type of imagers for special operation modes like HXRSS (Hard X-Ray Self Seeding) and XFELO (X-ray Free Electron Laser Oscillator demonstrator).

Most of the installed imagers are invasive, i.e., the luminescent screen is moved into the beam and partially or fully absorbs by the x-rays. Based on the physical properties of the screen various levels of damage within a single x-ray pulse may occur resulting in partial or full destruction of the screen. To avoid potential damage of the screens, a real time monitoring system in the Karabo control system has been developed and is actively used.

## CURRENT LUMINESCENT SCREEN PROTECTION WITH KARABO CONTROL SYSTEM

At the European XFEL, Karabo, a distributed control system is used [3]. It provides a unified ecosystem for hardware integration, monitoring, control, and data acquisition. Communication between beamline and instrument components is achieved via a centralized broker and organized into Karabo services and devices. Written in Python or C++, a Karabo device is a self-contained software component that might interact directly with the hardware or perform higher level procedures. For hardware integration various Karabo devices allowing the operator to control actuators, motors or valves, or orchestrate several hardware components, like high speed cameras, spectrometers and detectors, are available. For higher level automation tasks, various devices like x-ray feedback [4], the scantool [5] and others [6] have been developed and are actively used.

Karabo's graphical user interface (GUI) is based on a single desktop application written in Python using the Qt library [7]. It allows direct access to hardware components and higher-level automation tasks via Karabo devices. Control system components are organized in Karabo projects containing Karabo backbone services, devices and graphical

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBG005

# ARCHITECTURE OF THE PERSONNEL PROTECTION SYSTEMS FOR SPALLATION NEUTRON SOURCE SECOND TARGET STATION*

T. Michaelides†, S. M. Hartman
Oak Ridge National Laboratory, Oak Ridge, TN, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Oak Ridge National Laboratory (ORNL) is implementing a major upgrade to the Spallation Neutron Source (SNS) facility, encompassing the addition of the Second Target Station (STS). Preliminary design reviews have been conducted on several STS Personnel Protection Systems (PPS). The reviews focused primarily on the integration with the existing SNS PPS, the new proton transport tunnel, and the target areas. Development of the PPS is ongoing, to ensure a coherent safety system with the mission of protecting users and workers from prompt radiation hazards while providing high beam availability to operations. The STS PPS element in the Integrated Control System (ICS) is a facility-wide system composed of multiple safety subsystems, including the Ring to Second Target (RTST) beam transport tunnel, Target, Bunker and Instruments. Personnel working in all these geographic areas are protected by modular reliable PPS solutions. The safety system enforces access controls, radiation monitoring, beam destination control, and application of critical device inhibit upon detection of abnormal condition. It uses well-documented processes, Common Industrial Protocol (CIP) safety, pulsed test, and redundancy to achieve the desired Safety Integrity Level (SIL) [1, 2]. This paper gives an architectural overview of the STS PPS and a detailed safety plan for the SNS facility, addressing safety solutions and human factors.

## PPS SEGMENTS

The safety design of the PPS segments is to facilitate the monitoring, operations and prompt radiation safety of the facility. Each segment has its group of safety hardware, treated as a unit, that monitors and controls access to its segment. The PPS access controls, and beam containment functions treat each segment as a single unit, including its associated exclusion area(s). A segment is accessible to personnel if it is in an operating mode that allows personnel to enter an exclusion area monitored by that segment. Six segments sharing critical information make up the SNS PPS, and the STS adds another three as shown in Fig. 1.

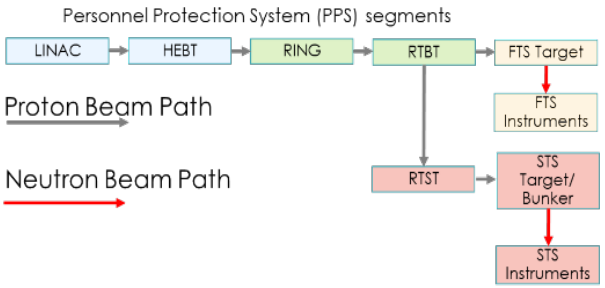



Figure 1: SNS & STS PPS segments block diagram.

### *SNS PPS*

The SNS PPS segments consist of the Linac, HEBT (High Energy Beam Transport), Accumulator Ring, RTBT (Ring to Beam Transport), Target and Instruments. The STS PPS segments include the RTST, Target/Bunker and Instruments. Each segment using safety rated Programable Logic Controllers (PLC) and well documented Operating Procedure Manuals (OPM) enforce:

- Personnel access controls
- Radiation monitoring
- Search procedures
- Trapped key interlock
- Inhibiting radiation generating devices
- Inhibiting critical bending magnets

The beamline that transports protons from the extraction point in the Ring to the Target Beam Transport (RTBT) line to the Second Target Station (STS) is indicated in Fig. 2.

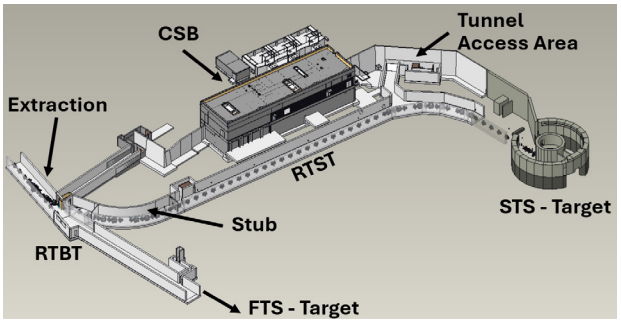



Figure 2: Extraction point to the Ring to Second Target (RTST).

### *STS PPS Critical Devices*

The STS PPS uses two sets of critical magnets to keep the beam from being extracted into STS-controlled areas shown in Fig. 3. The magnets are selected for the desired beam-control state, reliability, determinism and verifiability. Four pulsed dipoles and two transport magnets effectively provide beam attenuation independently.

* This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).
† michaelidest@ornl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBR001

# DEVELOPMENT OF A DIGITAL TWIN TOWARDS ENHANCED MONITORING OF THE $SF_6$ SUBSYSTEM ON THE Z MACHINE *

I. McCabe†, C. Pyles‡
Sandia National Laboratories, Albuquerque, NM, USA

*Abstract*

In this work, we develop a digital twin (DTW) of the Sulfur Hexafluoride ($SF_6$) subsystem on the Z Machine at Sandia National Laboratories. The Z Machine is a premier pulsed power research facility for studying high energy density science, to which $SF_6$ serves as an insulating gas. A DTW can dynamically update alongside its physical counterpart using real-time data, providing a precise estimate of the system's internal state which can augment monitoring processes. Thus, partially motivating the development of this DTW is a desire for online leak detection, as $SF_6$ emissions are regulated and reportable. The DTW discussed in this work supports real-time, online IO and interaction with the $SF_6$ control system; it has been incorporated as the plant model in a state observer, using online telemetry to synchronize the DTW's state with that of the physical system in a software in the loop framework. We discuss future enhancements to the digital twin aimed at minimizing $SF_6$ loss, including leak detection, leak localization, and leak isolation algorithms, as well as potential predictive maintenance strategies.

## INTRODUCTION

The Z Machine is a high precision instrument for studying high energy density science across three broad domains: dynamic material properties, inertial confinement fusion, and radiation science. The Z Machine generates extreme conditions found within celestial bodies or nuclear reactions, for example, by storing up to 22 mega joules of electrical energy in its Marx capacitor banks and releasing this energy in a 100-nanosecond pulse to the target at the center of the machine. This process can produce an output pulse of over 29 million amperes and 80-100 trillion watts of electrical power, as well as an x-ray output of 2 megajoules.

$SF_6$ gas is an exceptional electrical insulator and arc quencher. Accordingly, Z's $SF_6$ subsystem ($SF_6$SS) provides centralized $SF_6$ distribution to high-voltage components, providing insulation to the current running through the machine and enabling nanosecond switching operations. The $SF_6$SS consists of many control points, valves, pressure regulators, and pressure vessels. Due to varying experimental requirements, it is highly dynamic, heavily automated, and equipped with numerous monitoring systems. The complexity of the subsystem, as well as the gravity of potential gas leaks, necessitates a robust solution –one that provides fault detection and offline testing capabilities.

A digital twin is a highly customized virtual model of a physical system, which leverages real-time data supplied from sensors on the physical system to continuously simulate its behavior and track its operations [1]. It can be interfaced with control algorithms, provide enhanced insight into the behavior of the physical system in real-time, and support proactive maintenance. Additionally, a DTW can be easily modified and used to perform automated testing under various conditions of interest. [2]. Digital twins are especially useful in the gas industry where sustainability is a serious concern, increasing efficiency, reducing downtime, and improving safety [3]. In the context of the $SF_6$SS, a DTW can not only reduce the risk of $SF_6$ loss but can also reduce risks associated with directly modifying the subsystem by providing an environment for designing controllers and evaluating software and hardware alterations.

This DTW was designed to take the same set of inputs as the physical $SF_6$SS, namely pressure setpoints and valve configurations, and replicate the behavior of the $SF_6$SS's individual components. The DTW should also be capable of serving as a state observer, synchronizing itself with the $SF_6$SS when connected to its live control infrastructure and producing state estimates. To enable reliable real-time fault detection, the DTW must maintain high accuracy in its predicted outputs and estimated states.

## METHODOLOGY

### *Development Environment*

The simulation environment chosen for this project was MATLAB Simulink and Simscape as it contains many of the gas domain components found on the $SF_6$SS. The DTW interface software was developed in LabVIEW, enabling real-time IO interaction with the DTW.

### *Digital Twin Structure*

**Simulink Simulation** The simulation consists primarily of physically realistic gas chambers, valves, pipes, reservoirs, and pressure regulators, as well as data routing components. It contains 45,120 blocks, 5,526 discrete states, and 110 continuous states. Eighty inputs and correction signals are supplied to the simulation, resulting in 83 outputs along with 79 estimated $SF_6$ masses, one for each major subsystem. The simulation uses a variable-step ODE45 solver and runs approximately ten times real-time speed. Structurally, the simulation is composed of four major layers. Figure 1 depicts the main individual components (ovals) as well as the major subsystems (rectangles) within

* Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. SAND#2025-11130C
† inmccab@sandia.gov
‡ cspyles@sandia.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBR002

# ACCELERATOR COMMISSIONING TOOLS REVIEW, DESIGN PATTERNS, AND ARCHITECTURAL APPROACHES

W. Sulaiman Khail*,P. Schnizer
Helmholtz-Zentrum Berlin, BESSY, Berlin, Germany

*Abstract*

Particle accelerators are costly, complex machines. Modern light sources in particular cannot reach design performance without a dedicated software stack for commissioning and tuning, validated during design and construction using Monte-Carlo studies and digital twins. As both the software and the twins grow in scope, architecture matters.

We review representative tool stacks, extract recurring design patterns, and show how treating the system as interactions over a large shared state leads to simpler, more consistent implementations of commissioning tools and digital-twin workflows.

## INTRODUCTION

Complex experiments or accelerators require a sophisticated design, build, assembly, and commissioning procedure. We use accelerators as our main example, as we are working within an accelerator environment.

### *Designing and Building Complex Projects*

Accelerators are typically first designed using computer simulations. During the design process, technology experts join the design teams to ensure that the required devices comply with proven technology. If not, appropriate research and development programs are initiated.

Devices are designed and built to provide the intended functionality of lattice elements. This process improves the prediction of parameter variations that these devices will have. Along with this process, the lattice design is revisited and tested to ensure that the accelerator complex will achieve the desired performance parameters. Typically, this check is performed during the device production process as soon as measurements of manufactured devices become available. Finally, after the entire complex has been assembled, it is commissioned: first, its devices, then functional units, and finally with beam. To summarize:

- An accelerator lattice is specified as a set of lattice elements.
- Devices are designed, built, and tested to realize the elements' intended functions within tolerances.
- Devices are assembled, installed, and aligned into the machine.
- The accelerator is commissioned and tuned to achieve design performance.

* waheedullah.sulaiman_khail@helmholtz-berlin.de

### *Commissioning Tools Becoming Part of Operations Infrastructure*

Since MAX IV – the first 4$^{th}$-generation synchrotron light source – has been successfully put into operation [1], subsequent machines that are pushing the limits of this design [2,3] depend heavily on commissioning carried out by appropriate computer-based tools, i.e., self-adaptive systems. The devices of these machines cannot be built and operated with sufficient accuracy. Typically, these machines cannot simply be turned on using the parameter set derived from a dedicated commissioning period; instead, all devices must be carefully tuned to operation. As a consequence, these commissioning steps must be executed periodically – often every time the machine is started up, which is typically every week. Therefore, these tools must also be executed periodically, and machine operation relies on their functionality.

The development of commissioning tools begins concurrently with the design of the machine. The design process not only verifies the performance of the ideal lattice but also ensures that it is robust enough to work with variations in device performance. This means that the design must be either robust enough to work out of the box, or that appropriate commissioning tools must be available to measure the actual machine parameters with sufficient accuracy, so that mitigation actions can be applied to achieve the desired performance.

As modern machines depend heavily on these commissioning tools, they have become critical components that are designed and tested alongside the machine using Monte-Carlo simulations.

## REVIEWED TOOL STACKS

We examined several commissioning toolboxes and identified recurring patterns that these toolboxes employ. In this paper we present these patterns and demonstrate how they can be used to structure different aspects of commissioning tools and digital twins.

Our ideas build on our previously published and ongoing work [4–7] together with tools and tool stacks used to design or commission accelerators:

- Johan Bengtsson's tracking code Tracy II and associated tool set [8, 9].
- Online models

  **NSLS II** This online model allows simulation of the whole accelerator complex [10]. It uses Tracy II as its simulation engine.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBR003

# VIRTUAL BEAMLINE++: ACCELERATING THE DESIGN OF HIGH-ENERGY, NEXT GENERATION LASER ARCHITECTURES

J. Penner†, R. Aden, K. McCandless, S. McLaren, L. Waxer
Lawrence Livermore National Laboratory, Livermore, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Virtual Beamline++ (VBL++) is an advanced nonlinear optical laser beam propagation code used for the design and operation of high-energy and high-power laser systems. It is the latest tool employed to precisely setup the National Ignition Facility's (NIF) laser beamlines and support a broad range of missions from fusion ignition to high energy density science.

VBL++ also plays a critical role in supporting a growing base of users with emerging and diverse laser designs, technologies, and applications. This tool actively contributes to the design of novel architectures, including some for Inertial Fusion Energy applications, proposed by NIF researchers and collaborators. VBL++ offers an intuitive user interface for building optical chains and supports high-resolution runs using High Performance Computing (HPC) resources. This talk will focus on a subset of features in VBL++ that help users to elevate their modeling above single simulations. These features include:

- Parameterization and optimization of inputs
- An interface for smooth desktop-to-HPC integration.
- An inverse solver used to determine the input low-power pulse shape that will achieve the high-power request on target.
- An interface written in Python, MATLAB, and IDL that enables users to build chains of optical components and run simulations with greater configurability.

This application is still under active development: new models and enhanced features are added regularly to address evolving user requirements at the cutting edge of fusion class laser design and operations.

## INTRODUCTION

As laser physicists continuously seek to enhance the capabilities of their systems, it is essential to ensure that experimental configurations yield the desired outcomes without risking damage to costly equipment. To address this requirement, scientists and engineers at Lawrence Livermore National Laboratory have developed specialized software to model laser behavior and simulate energy gain as well as potential sources of damage. At Lawrence Livermore National Laboratory, precursors to the NIF facility used MALAPROP as early as the 70's [1] to validate their setups. In the mid 90's, its successor PROP92 and later PROP [2] were used in the design of the NIF facility itself. The Virtual Beamline (VBL) [3] code was developed in the mid 2000's to support the daily operations of the NIF laser, with an architecture that allowed for new optical models to be added regularly to increase the code's fidelity.

As NIF's operating range was pushed even further in the hunt for ignition, users of the VBL code required higher resolution simulations to capture all the fine-grained spatial and temporal effects in their laser setup calculations. Developers and users both concluded that a major upgrade to the code was required to allow simulations to be performed in parallel execution environments and capture more computationally intensive laser physics effects

VBL++ – named as an homage to the C++ upgrade from the C programming language – is an application written in C++ that provides an intuitive interface for users to build simulations that could represent anything from a tabletop setup to the NIF laser chain. It can then run the simulations anywhere from a personal laptop to an HPC server. This paper will explain the core functionality of the VBL++ application as well as highlight features that have been added to support the wide range of use cases that scientists require.

## OVERVIEW

### *Modes of Propagation*

VBL++ operates using a paraxial approximation of the classical wave equation [4] which allows for a split-step algorithm that separates propagation of an electric field array and nearfield interactions into two separate steps. The propagation step can be implemented by performing a Fast-Fourier Transform (FFT), applying a phase term, and performing the inverse FFT. This approximation significantly decreases the computational complexity of the propagation algorithm and minimizes interactions between data points within the grid, thereby facilitating the implementation of data-level parallelism.

This default mode of propagation – with the electric field values stored in a 3-dimensional complex array representing the electric field's shape in 2-dimensional space (with an x and y axis) and a pulse length in time – is called the "Monochromatic Fourier" mode. VBL++ also offers a further approximation named the "Hankel" mode, which assumes the laser beam has a revolution symmetry over space and therefore can be represented in a single radial dimension. The "Broadband" mode is named as such because it aims to model broadband laser effects –where a Fourier transform can be performed along the time dimension to manipulate the frequency domain– such as those captured in the unidirectional pulse propagation equation (UPPE) [5]. Because the Broadband mode often requires a highly resolved computational grid, a "Chirped" approximation mode is included that assigns instantaneous frequencies to each time slice [6].

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOBR004

# LEVERAGING SOFTWARE SIMULATORS IN SKA DISH LMC DEVELOPMENT AND TESTING

S. N. Twum*, SKA Observatory, Jodrell Bank, United Kingdom
A. Lethole, M. Nzama, B. Ojur, N. Padavattan, D. Petrie, Y. Reddi, U. C. Silere, A. Venter
South African Radio Astronomy, Cape Town, South Africa

*Abstract*

The Square Kilometre Array (SKA) project employs the TANGO Controls framework to manage its telescopes. Within SKA MID, the Dish Local Monitoring and Control (LMC) software integrates dish sub-components into the overall control system. Because hardware availability is often limited, Dish LMC development and validation are heavily based on software simulators derived from Interface Control Documents (ICDs) and system-level requirements. These simulators emulate the behaviour of the respective devices for requirement verification. Their use has enabled iterative testing in both staging and integration environments, accelerated development, and reduced risks ahead of hardware integration. This paper presents the design and implementation of Dish LMC simulators, their role in supporting system-level validation and advancing software maturity.

## THE SKA CONTROL SYSTEM - DISH LMC'S ROLE

The SKA project is a global endeavor to construct the world's largest radio telescope: SKA MID in South Africa (operating between 350 MHz – 15.4 GHz) and SKA LOW in Australia (operating between 50 – 350 MHz). SKA MID will comprise 197 Gregorian offset dishes, including 64 MeerKAT dishes and 133 SKA dishes. In parallel, SKA LOW will employ 131,072 dipole antennas grouped into 512 stations. Together, these telescopes will cover a collecting area of one square kilometre [1].

A distributed Tango control system provides remote access to the telescope subsystems [2], as well as other hardware on site. The dish subsystem implements a Tango interface that aggregates sensor data from its subcomponents to communicate the overall dish state. Most monitored subcomponents expose a Tango interface, while a few implement either OPCUA, SNMP, or other proprietary communication protocols.

Dish LMC provides master control and aggregated monitoring of the dish subsystem shown in Fig. 1. Dish LMC's architecture, shown in Fig. 2, includes a power manager, single-pixel feed controller (SPF), single-pixel feed receiver (SPFRx) controller, Band 5 down converter (B5DC), dish structure manager (DS), and dish manager [3] [4]. The development of Dish LMC is guided by ICDs and system requirements. However, the limited availability of hardware has necessitated the use of software simulators to accelerate development, mitigate risks, and validate requirements early.

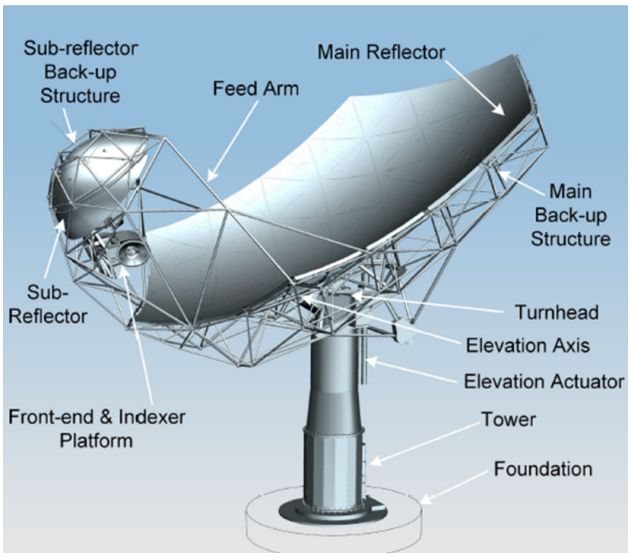



Figure 1: SKA MID Dish elements.

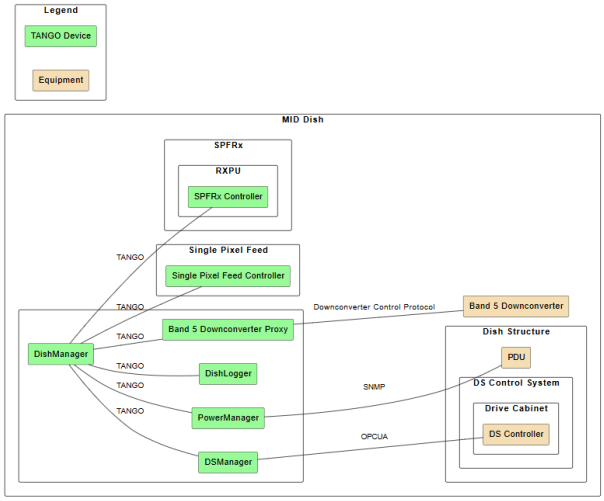



Figure 2: Dish LMC architecture overview.

## DISH LMC SIMULATION APPROACH

The various software components comprising the Dish LMC are developed by different teams who are not contracted to provide simulators for their product. Consequently, significant effort is invested in building a simulator that represents the hardware. The Dish LMC simulators are designed to reproduce the behaviour of hardware subcomponents based on ICDs, ensuring that software integration aligns

* samuel.twum@skao.int

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCG002

# NEURAL NETWORK-BASED MODEL PREDICTIVE CONTROL FOR ENERGY EFFICIENT HVAC SYSTEMS FOR CERN ACCELERATORS

N. Bunijevac*, D. Monteiro, CERN, Geneva, Switzerland


## *Abstract*

In times of concern over the environmental impact of high-energy physics organizations, our research in CERN's Cooling and Ventilation group in the Engineering department (EN/CV) investigates energy-saving strategies for heating, ventilation, and air conditioning (HVAC) systems. Widely used in both the residential and industrial settings, HVAC systems contribute up to 40 % of the residential and 70 % of industrial consumption, making their optimization a global concern. At CERN, cooling and ventilation systems are used to ensure appropriate temperature and humidity conditions in the accelerator complex. Together with general water-cooling systems, these systems account for up to 12 % of the total electricity consumption of the Large Hadron Collider (LHC), CERN's flagship accelerator. Despite their energy intensity, these systems are typically managed by classical controllers, which are reliable for temperature and humidity regulation, but not optimal in terms of energy efficiency. This study in the EN/CV group aims to quantify the potential energy savings in HVAC systems using model predictive control (MPC), an advanced control strategy with two main elements: 1. a system predictive model (digital twin), that leverages external data such as weather forecast, and 2. an optimization process. The optimization algorithm task is to provide an energy-optimal solution over a predefined time in future, while also keeping the controlled process variables, such as zone temperature, within required limits. In a quest for a reproducible solution that can be easily adapted to different HVAC plants, the digital twin is built using physics-informed neural networks, following recent advances in academic research, while the optimization algorithm is the genetic algorithm. For the pilot project, a specific plant at CERN is chosen - HVAC311, where the integration of MPC strategy is currently undergoing, as a cloud high-level controller that provides optimal setpoints to the existing low-level system controller (PID). An example of the resulting optimization of the MPC is given, as well as discussion on future work and suggested improvements.


## INTRODUCTION

### *Problem: High Energy Consumption of HVAC Systems*

HVAC systems (heating, ventilation, and air conditioning) are the crucial elements to provide comfort in residential and public buildings, as well as the proper operation of the many industrial complexes. Air conditioning systems maintain required air conditions in indoor spaces, in terms of temperature, humidity and air quality. However, they are among the largest contributors to global energy consumption and carbon emissions. The impact of HVAC systems on the environment is only expected to increase in the future [1]. The building sector takes 40 % share of total global primary energy use and about 30 % of global greenhouse gas emissions [2] while HVAC systems contribute to around 38 % of residential building consumption [3]. Therefore, improving HVAC systems efficiency plays a crucial role in energy and emissions reduction.

### *Efforts of CERN and Cooling and Ventilation Group Towards Higher Energy Efficiency*

In response to increasing awareness regarding growing environmental consciousness, together with energy rising costs, many governments, companies and organizations have been taking actions to reduce their energy footprint, including Organization for Nuclear Research (CERN). In a report published in 2021, CERN management set a goal "to establish itself as the model for transparent and environmentally responsible research organizations," which includes implementing "actions and technologies aiming at energy saving and reuse" [4]. Among many initiatives, Cooling and Ventilation group (CV), a part of engineering department (EN), focuses on improving efficiency of cooling and ventilation systems. Cooling and ventilation systems maintain appropriate temperature and humidity conditions for accelerator machinery and experiments equipment, but are, like other systems of this type, very energy intensive. During the experiments run phase, the electricity consumption dedicated to cooling and ventilation makes 12 % 75 GWh [5] of the total electricity consumed by the Large Hadron Collider (LHC), CERN's flagship accelerator. Innovation in control strategies is important since it results in significant savings with low-cost changes. A key approach in CV is the "controls optimization method", the approach in which several control strategies are simulated and analyzed in a safe digital environment (digital twin), and the most optimal one is chosen and deployed [5]. More recently, the team decided to develop an advanced control strategy Model Predictive Controller (MPC) with application on HVAC systems. The core part of MPC is the model that is capable of predicting system behavior in the future, based on the initial state, control sequence being tested, and weather forecast. This predictive power is what differentiates MPC from standard control strategies and is a key element in the optimization process that identifies the energy-optimal set of controls. This method can be applied on many different complex systems, while for the HVAC systems in particular, the literature declares energy savings of up to 20 % [6].

In the study presented in this paper, the authors propose

* nikolina.bunijevac@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCG004

# OPTIMIZATION OF THE TRANSFER LINE AGAINST COLLECTIVE EFFECTS USING PHYSICS-CONSTRAINED GENERATIVE PHASE SPACE RECONSTRUCTION

S. Kim*, G. Kim, H. Yang, PAL, POSTECH, Pohang, Republic of Korea
R. Roussel, A. Edelen, SLAC, Menlo Park, CA, United States
J. P. Gonzalez-Aguilera, University of Chicago, Chicago, IL, United States

*Abstract*

Optimization of the transfer line against collective effects such as space charge and coherent synchrotron radiation (CSR) effects is crucial to preserve the beam quality. While simple conventional diagnostic methods provide ensemble averaged beam parameters or limited information of phase space, they are still limited in obtaining precise, complete 6-dimensional phase space with all the correlations due to hardware and dedicated time requirements. A generative phase space reconstruction method (GPSR) has been developed as a robust diagnostic framework that reconstructs complete 6-dimensional phase space. Here we show a physics-constrained GPSR model that incorporates known physical parameters, such as RMS beam sizes and emittances, as constraints. We performed simulated demonstrations at the Pohang Accelerator Laboratory X-ray Free Electron Laser facility that the Physics-informed GPSR can be performed for complete 6-dimensional phase space. Furthermore, by using the reconstructed phase space, we performed non-differentiable particle tracking simulations to investigate the phase space evolution against the space charge and CSR along the bunch compressor. We present the trend of predicted CSR-induced emittance growth, which closely matches the ground truth.

## INTRODUCTION

As the particle accelerator becomes an essential scientific tool to explore not only fundamental physics [1] but also characteristics of materials [2] or protein [3], understanding the distribution of particles beams is one of the essential tasks of accelerator control. Conventional beam diagnostics techniques, such as quadrupole scan [4], slit-scan [5], or tomography [6–8] are commonly used in accelerator facilities to measure the ensemble-averaged beam parameters (e.g., Twiss parameters or emittances) and measure the phase space distribution of the beam. While the previous demonstrations on phase space diagnostics using slit-scan [7] or tomography [8] were successfully performed, yet they are practically limited to be used as typical diagnostic methods due to the requirement of dedicated beam time (> 24 hours) or special diagnostic equipment [9].

Recently, a phase space reconstruction method [10] (known as generative phase space reconstruction; GPSR) has been developed using differentiable simulations and neural networks, which basically utilizes the measured screen images along the beamline. The main algorithm of the GPSR method is as follows: 6-dimensional particle distribution $(x, x', y, y', z, \delta)$ is generated by the neural networks (parameterized transform). The initially generated particle distribution is used as an input of the differentiable simulation where the beamline lattice is the same as the experimental setup (e.g., quadrupole scan that consists of single quadrupole magnet, drift space, and YAG scintillating screen). Then, the simulated beam images at the beamline lattice are compared with the experimentally measured data (ground truth) to perform the back-propagation using the gradient descent method, enabling updates of weights of the neural networks to refine the initial phase space. This process is repeated until the reconstructed phase space predicts the experimental images.

The first experimental demonstration has been successfully performed to reconstruct 4-dimensional phase space of typical round beam using the conventional quadrupole scan method [10]. In addition, characteristics of special beams such as magnetized beam and flat beam have been investigated [11]. Moreover, complete 6-dimensional phase space reconstruction was achieved by introducing transverse deflecting cavity and spectrometer magnet [12] during the experimental demonstration; in such a system, studies on the reconstruction for CSR-dominated beams [13] were also carried out [14]. In case of the complete 6-dimensional phase space reconstruction, recent study shows that the reconstruction can also be performed using the accelerating cavity (instead of deflecting cavity) [15] for capturing $(z, \delta)$ correlations, making the AI/ML-based beam diagnostics feasible with readily available components without dedicated diagnostic equipment.

In this study, we extended the study on reconstruction of accelerating cavity-based 6-dimensional phase space of the beam for optimization of the beamline together with a non-differentiable simulation to investigate the CSR effect on the beam emittance along the Pohang Accelerator Laboratory X-ray Free Electron Laser (PAL-XFEL) beamline. We show simulated demonstrations to discuss the main approach on how to consider the physical constraint for the GPSR in our setup, and present the discussion that the reconstructed 6-dimensional phase space successfully predicts the trend of the CSR-induced emittance growth and optimal beamline setting to preserve the emittance against the effect.

* sykim12@postech.ac.kr

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCG005

# AI-DRIVEN AUTONOMOUS TUNING OF RADIOACTIVE ION BEAMS*

S. Lopez-Caceres†, D. Santiago-Gonzalez, Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Facilities that produce radioactive ion beams (RIBs) enable unique advances in nuclear structure and astrophysics, but their operation demands frequent, expert-guided tuning across a high-dimensional, multi-device parameter space. We describe an AI-driven online optimization framework for the Californium Rare Isotope Breeder Upgrade (CARIBU) at Argonne National Laboratory that couples Bayesian Optimization with the accelerator control system. The system features a plugin-based interface for devices and diagnostics, a human-in-the-loop workflow, and the use of archived configurations to ensure robust, safe, and repeatable tuning. We demonstrate closed-loop optimization of electrostatic beam transport using beta-decay detector feedback, achieving performance comparable to expert tuning while reducing operator cognitive load and enabling automation. The approach is portable to other RIB facilities and extensible to additional complex beamline subsystems.

## INTRODUCTION

Radioactive ion beam (RIB) facilities provide access to rare isotopes for forefront research in nuclear physics and nuclear astrophysics. Achieving reliable beam delivery requires coordinated tuning of many devices (steerers, electrostatic quadrupoles, deflectors, dipoles, and diagnostics), traditionally performed manually by domain experts. This practice is time-consuming, non-trivial to reproduce, and sensitive to slow drifts and configuration changes.

Artificial Intelligence (AI) and Machine Learning (ML) methods have shown promise in automating complex beam-based tuning tasks with minimal prior data [1, 2]. In particular, Bayesian Optimization (BO) provides an efficient, data-driven framework for optimizing expensive-to-evaluate systems with unknown response surfaces [3]. BO uses a probabilistic surrogate—typically a Gaussian Process (GP)—and an acquisition function to balance exploration and exploitation.

Building on these developments, we report an online BO framework integrated with the CARIBU [4] beamline controls at the Argonne Tandem-Linac Accelerator System (ATLAS). We emphasize control-system integration, operation workflow, and operator experience, aligning the presentation with ICALEPCS interests.

* Work supported by the U.S. Department of Energy, Office of Nuclear Physics, under contract DE-AC02-06CH11357, and DE-FOA-0002875 funds.

† Corresponding author: slopezcaceres@anl.gov slopezcaceres1@gmail.com

## FACILITY CONTEXT AND OPERATIONAL MOTIVATION

CARIBU produces RIBs from the spontaneous fission of $^{252}$Cf and feeds experimental stations for research in nuclear structure and astrophysics. A typical RIB transport involves in-tandem sections with steerers and electrostatic quadrupoles, mass selection with a dipole-based isobar separator early in the line, and diagnostic stations that monitor beam radioactivity (Fig. 1).

Operationally, each beamline section presents 5–15 control variables constrained by preset limits and operational procedures. The full RIB transport path involves on the order of 200 parameters. Manual tuning must be repeated upon ion species change, beam interruptions, or drifts, increasing the burden on limited operations staff. Automating routine tuning with a human-in-the-loop reduces overhead while maintaining safety and quality.

## SYSTEM INTEGRATION

The optimizer was implemented within the Badger [5] optimization environment and connected to the CARIBU control system. Integration required:

- Reliable communication between Badger and device control interfaces,
- Safety checks to limit parameter ranges and protect equipment from harmful settings,
- Real-time monitoring of beam rate and device status,
- Beam operator access through GUI to allow human-in-the-loop intervention.

Our system integrates an online BO engine with the control system via a plugin architecture: - Software stack: Badger [5] (optimization GUI and plugin system) with Xopt [6] (optimization backend) built upon BoTorch [7], GPyTorch [8], and PyTorch [9] for GP/UCB models and acquisition optimization. - Device Interface plugin: encapsulates read/write operations and custom procedural steps. - Environment plugin: defines optimization variables, initial "good-tune" seeds, and constraints. - Diagnostics: silicon detectors count beta electrons as a proxy for RIB intensity. - Controls integration: remote access is brokered by an operator-enabled server within the ATLAS Web Access and Control Server (AWACS). - Data and reproducibility: per-run configurations, model state, and results are archived to enable rollback and ad-hoc analysis. The details of the underlying algorithm and theoretical justification are discussed extensively in reference [10].

Figure 2 sketches the workflow: network operators enable the control gateway, beam operators initially connect via

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCR001

# REIMAGINING THE MEERKAT ARCHIVE: MORE THAN MEETS THE UI

Z. Smith*, A. Gangat, C. Schollar
South African Radio Astronomy Observatory (SARAO), Cape Town, South Africa

*Abstract*

We recently upgraded the interface for the *MeerKAT Telescope* data archive. Scoped as a straightforward rewrite of an existing list-based search interface to address longstanding user feedback and bugs, we used the opportunity to rethink the role the archive can play within the South African Radio Astronomy Observatory (SARAO) and the broader astronomy research community.

As we introduced a new architecture, added capabilities, and improved the user interface (UI), surprising user feedback highlighted aspects of the user experience (UX) that we had overlooked. This study shows how reimagining existing legacy systems provides opportunities to reveal assumptions, identify non-obvious user workflows, and improve the human/system interface.

With over 700 active users searching and downloading data, the new interface has been well received. We are able to respond to support requests quickly and have improved both usability and access to MeerKAT data. Comprehensive usage metrics underpin quick, effective software support as we continue to develop a nuanced understanding of MeerKAT data access and management.

## INTRODUCTION

The MeerKAT, located in the Northern Cape (South Africa), is one of the world's largest and most sensitive radio telescopes. As of late 2025, the MeerKAT data provides open access to over 5,000 high-quality datasets, spanning 6 years of operation. This represents more than twenty thousand hours of observations stored across several petabytes of files. As the telescope continues to explore the cosmos and embargo periods on observations expire (to give proposing teams time for initial publications), interest in MeerKAT data will continue to grow. The archive provides an ever-increasing accumulation of observations; some are repeats of the same sources over multiple epochs. Many potential discoveries lie in wait as this data becomes available retrospectively.

Scientific publications in astronomy are increasingly reliant on archived data rather than exclusively on new observations (e.g., [1, 2]). Accompanying this growing reliance on archives is the challenge of ensuring that the data remain accessible. MeerKAT data is catalogued, searchable, and quality-controlled, but volume doesn't guarantee usability. Making use of these observations still requires a thoughtful archive interface, reliable web-based tooling, and realistic workflows for accessing and analysing large amounts of data. We recently relaunched SARAO's data archive—reconfigured as the *MeerKAT Archive*—to support long-term access to MeerKAT data [3].

The MeerKAT is the successor to the Karoo Array Telescope (KAT-7) [4]. The previous UI implementation, which was initially designed to support data from the KAT-7 telescope, contained anachronisms obscuring the logic of the underlying archive. Much of the logic in the old archive interface was tied to fields and data products from this precursor instrument, which are now defunct. The code was also old—developed in a deprecated framework—and otherwise hard to maintain or evolve due to earlier architectural decisions and coding patterns. Lack of observability in the form of usage metrics, event logs, and access information for the old interface meant that much of how users engaged with the interface was hidden. Redeveloping the new interface therefore became not just a redesign, but an exercise in reimagining the archive itself.

With the initial interface as a reference, the new work builds on the previous UI effort to deliver an interface that is more maintainable, facilitates quicker response to bugs and feedback, and otherwise more cleanly expresses systems, in source code, that support and communicate MeerKAT data. Guided by user-feedback, event logging, and usage metrics, the project has become a deliberate, open-ended exploration of the role the MeerKAT Archive plays both today and into the future.

As users slowly migrated from the old interface to the new interface, we discovered from system metrics that the archive is used by over 700 active users. Despite its apparent simplicity as a search tool, we frequently misunderstood what these users—engineers and scientists—wanted to search for, and how. Incorrect expectations of usage patterns shaped early design decisions in ways that only became clear in hindsight.

This paper introduces the new archive interface by providing an overview of the system architecture and our UX journey as we gained insights into our users. It also highlights how our reframed approach to authentication and authorization enabled us to better define the role our software could play within SARAO as an organization. The MeerKAT Archive is accessible at `archive.sarao.ac.za` (Fig. 1).

## A LAYERED ARCHITECTURE

Building the new portal as a static Single Page Application (SPA) paired with an Application Programming Interface (API)—a public interface exposing internal systems over HTTP—ensures that programmatic access to the archive evolves in parallel with the React-based UI. Every

* zsmith@sarao.ac.za

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCR002

# PHOEBUS: AN OPEN ECOSYSTEM FOR CONTROL SYSTEM APPLICATIONS AND SERVICES

K. Shroff, Brookhaven National Laboratory, Upton, NY, USA
R. Lange, ITER Organization, St. Paul lez Durance, France
T. Ford, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
G. Weiss, S. Brewer, European Spallation Source ERIC, Lund, Sweden
K. Kasemir, Oak Ridge National Laboratory, Oak Ridge, TN, USA
I. Finch, Science and Technology Facilities Council, UK

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Phoebus toolkit continues to evolve as a modern, extensible platform for control system user applications and middle-layer services. As the next generation of the Control System Studio, Phoebus retains the familiar, integrated toolset experience while replacing the Eclipse RCP framework with a modular architecture built on standard Java technologies and JavaFX. This transition simplifies maintenance and extensibility while providing foundational building blocks for scalable services and modern applications.

Recent work has focused on strengthening infrastructure and streamlining deployment. These efforts include support for recent Java LTS releases, modernization of microservices with updated Kafka and Spring Boot, and expanded documentation. Core services—covering alarm management, save/restore, name resolution, and logbooks—continue to mature, with emphasis on easier configuration, improved scalability, and alignment with modern web standards and containerized workflows. The Phoebus collaboration now includes contributors from dozens of facilities worldwide, many of them first-time participants. Alongside technical progress, the project prioritizes a sustainable and inclusive collaboration model to support future developers and users.

## INTRODUCTION AND MOTIVATION

Phoebus is the successor of Eclipse as the framework for Control System Studio and a growing number of control system services [1, 2]. Its development was motivated by the need to simplify the architecture, remove the heavy dependency on the Eclipse RCP framework, and modernize the technology stack. The result is an ecosystem that is both extensible and sustainable, as seen in Fig. 1, and which has been adopted by an increasing number of facilities worldwide.

At the heart of Phoebus is the integrated user experience. The toolkit provides a cohesive set of tools and applications that are easy to navigate between, where data can be seamlessly transferred from one tool to another, ensuring a consistent experience that supports the user workflows within the same environment. For example, an operator can investigate the details of an alarm, retrieve historical data for the PVs that triggered it, open related OPI/HMI displays to interact with devices, and finally record the investigation in the logbook. Each of these steps uses different applications—Alarm UI, Data Browser, Display Builder, and Logbook—but Phoebus enables seamless transitions between them. Selection services and adapter services automatically transfer context including PV names, current values, timestamps, alarm states, archive sources for historic data, OPI screen locations, log entries, and more. This rich context propagation reduces user effort and improves efficiency.

For developers, Phoebus provides a framework for creating applications that integrate smoothly into this ecosystem without requiring tight dependencies on other applications. Shared modules offer access to optimized resources such as connections to EPICS Process Variables (PVs), REST clients, and other common services, reducing duplication and lowering the barrier to entry for new contributors.

Community adoption has also been a central motivation. The project emphasizes open collaboration and inclusivity, encouraging contributions from new developers while maintaining a consistent architecture. This ensures that Phoebus continues to evolve in response to user needs while remaining sustainable in the long term.

To understand the motivation for this architecture, it is helpful to situate Phoebus within the broader controls ecosystem used at accelerator and large-scale scientific facilities. As illustrated in Fig. 1, instrumentation and Input/Output Controllers (IOCs) at the lowest level communicate via EPICS protocols such as Channel Access and PVAccess [3]. Middle-layer services—including the archiver, logbook, save-and-restore, and alarm system—provide persistence, workflow integration, and operational infrastructure. On top of these, multiple client environments coexist: physics toolkits, scripting interfaces, web clients, legacy CS-Studio, and increasingly, Phoebus. This context highlights the dual role of Phoebus: providing operator-facing applications and frameworks for new tools, while also offering middle-layer services that integrate tightly with existing site infrastructure.

## ARCHITECTURE OF THE PHOEBUS FRAMEWORK

Phoebus is organized as a modular framework with a small set of reusable core modules and a flexible extension mechanism. The design balances stability and adaptability: stable libraries and APIs provide common behavior, while

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCR004

# IMPLEMENTING DESIGN PRINCIPLES IN THE ROCK-IT GUIs: A UI/UX DEVELOPMENT CASE*

Z. I. Isik Dursun†, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Effective UI/UX design is essential for scientific GUIs, especially when complex workflows and time constraints challenge usability. By grounding GUI development in UI/UX design principles, accessibility, and user feedback; one can implement user-centred improvements that will lead to higher user satisfaction and increased productivity in experimental physics control systems. This paper explores the UI/UX design phases within the demonstrator case of the ROCK-IT project, shaped by theory and in-person user interviews. The key points in each phase are also highlighted; such as principles of screen design, visual accessibility, and examples of user feedback. The paper ends with how this process evolved in response to tight timelines, workforce availability, and shifting priorities in ROCK-IT - steering the route toward an efficient, yet user-centred and accessible solution.

## INTRODUCTION

ROCK-IT [1] (Remote, Operando-controlled, Knowledge-driven, and IT-based) is a project that aims to develop all necessary tools for the automation and remote access of in-situ and operando synchrotron experiments, funded by the Helmholtz Association in Germany. Aiming at users from various experience levels, the project enables remote access for experimental users, experiment automation including the use of robots, data management following FAIR data principles, and nearly real-time analysis and optimization using machine learning (ML) and artificial intelligence (AI). The ROCK-IT project is being implemented in four partners (and their beamlines) as a pilot development case, mainly for catalsysis operando experiments: Deutsches Elektronen-Synchrotron DESY (P65 beamline), Helmholtz Zentrum Berlin (HZB) (mySpot beamline), Helmholtz Zentrum Dresden-Rossendorf (HZDR) (HAMSTER beamline) and Karlsruhe Institute of Technology (KIT) (CATACT beamline). Sharing a considerable amount of mutual users, one of the project's aims is to maintain a consistent look-and-feel of user interfaces between these Helmholtz institutions.

Within this context, a task was to revise and enhance the existing GUIs among the partner institutions, which varied from command-line interfaces (CLI) to patchwork graphical user interface (GUI) solutions developed by various beamline scientists throughout a wide timeframe. This process consisted of 5 phases (see Fig. 1): Phase 1 - Field Research I; Phase 2 - Field Research II; Phase 3 – Evaluation; Phase 4 – Drafting UI/UX concepts; Phase 5 – User testing; Phase 6 – Final outcome. The following part will explore these phases within the UI/UX development case of ROCK-IT.

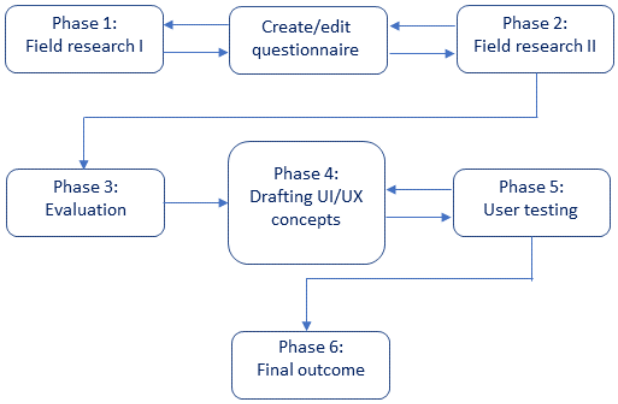



Figure 1: UI/UX Design Process Workflow at ROCK-IT.

### *Phase 1: Field Research I*

In order to outline the necessities of the task at the beginning, it was necessary to learn about the configuration and working process of the beamlines. This process was supported by in-person project meetings including control room and beamline visits at the partner institutes. Discussions among developers and beamline scientists were very valuable to gather information on the system, the needs, and possibilities of what can be done. These meetings also gave a chance to examine the current GUIs and observe some user interactions, creating a base for the user interviews.

### *Phase 2: Field Research II*

The second phase was to understand the users' needs with a qualitative research by in-person user interviews. The user interviews had two steps; i) Answering the questionnaire, ii) Collecting feedback through the existing GUI visuals. In this context, a questionnaire [2] was created to collect information on users' experiences and their needs. The questionnaire[1] consisted of 3 parts: i) Use cases; ii) User experience questions; and iii) Experiment processes[2]. These questions were first addressed to users at the ROCK-IT Outreach Meeting [3], which was held as a hybrid satellite meeting within the DESY Photon Science Users' Meeting in 2024 [4]. Following some revisions on the content, a preliminary question on the users' experience level was added to get an insight of different user types, and the questionnaire evolved into its current state with 18 questions. The question "Since what year have you been

*The ROCK-IT project is funded by the Helmholtz Association in Germany.

†zeynep.isik.dursun@desy.de

[1]Special thanks to Alexander Schökel (DESY) and Mehdi Kazemi (DESY) for their support and feedback on the questionnaire.

[2]Special thanks to Diana Rueda (DESY), Marc-Olivier Andrez (DESY), and Tim Schoof (DESY) for creating the "Experiment Processes" questions.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCR005

# BEYOND LIKE-FOR-LIKE: A USER-CENTERED APPROACH TO MODERNIZING LEGACY APPLICATIONS

M. Polzin*, M.Guzman†, Fermi National Accelerator Laboratory, Batavia, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

When modernizing a legacy application, it is easy to fall back on a like-for-like replica with new tools and updated design stylings, but this is an opportunity to explore making a more intuitive application that supports user tasks and efficiency. Rather than having a blank canvas–unburdened by legacy tech debt–to create a new application, you are working with an existing application that is integral to accelerator operations and one that expert users are already familiar with. Due to this, you might assume people will prefer the like-for-like, but you could be carrying forward the pain points, processes that are inefficient, and ultimately wind up with an application that no one wants to use because it doesn't solve existing problems.

Getting users involved can make all the difference in your approach to modernizing a legacy application that caters to both newer and expert users. It also can bridge the gap between like-for-like and introducing new GUI design. Having a legacy application doesn't have to make the modernized one difficult to develop, as the existing application is a tool in how you move forward with the new application. It provides insight into areas that an application with a clean slate doesn't give you.

## INTRODUCTION

At Fermilab, the Accelerator Controls Operations Research Network (ACORN) project is working to modernize Fermilab's legacy particle accelerator control system [**?**]. Within this larger project, our team is tasked with modernizing the existing control system applications and increase usability by creating intuitive and effective user interfaces. We are emphasizing user experience efforts by incorporating user input throughout our process. Although we could take a like-for-like approach and rebuild existing applications with newer technology, we have decided to seek opportunities for improvements to better meet the needs of our users.

## MODERNIZING LEGACY APPLICATIONS

Fermilab has over 600 legacy applications that are a combination of text based C/C++ applications and Java applications. A goal of the ACORN project is to modernize existing applications by providing a web browser-based interface to new applications that provide the functionality of existing legacy applications. In order for ACORN to modernize all existing applications, a full assessment of legacy applications must be done. The initial assessments were to inventory the lines of code for each application and look at usage statistics over a span of two years. Having the lines of code allowed to have a high level understanding of the complexity of each application and having usage statistics allowed to prioritize higher used applications.

### *Like-for-Like Pitfalls*

When modernizing applications one strategy could be to develop a like-for-like replacement. This strategy would take an existing application and replicate all functionality as well as the interface. This could be short term success as the development cycle could be quick. However, this strategy could be filled with pitfalls. A like-for-like replacement would carry over the user interface of legacy applications–which includes structural design choices and user workflows. This could mean that existing pain points from users will continue to exists. This strategy also does not take into account that the design decisions were based on the limits and strengths of technologies that are no longer being used. While the application will be modernized, the experience would not be.

### *Moving Beyond Like-for-Like*

Another strategy would be moving away from a like-for-like application. This allows for a reevaluation of the needs of users and purpose of the application, as those may have changed since the development of the initial application. A reevaluation of the application allows the development to leverage the strengths of current technologies and address inefficiencies in the user workflows. For example, moving to a browser-based application allows for leveraging the built in usability features of browsers, cross platform usage, and responsive experience. This can introduce an increase in usability and a reduction in learning curves. By moving to modern technologies you can also move away from technical debt and allow for adaptability and future improvements.

## LEGACY APPLICATIONS AND USERS

Modernizing legacy applications may feel constraining in comparison to applications that do not yet exist where you have more freedom, but there are ways to utilize the legacy applications and find opportunities for improvements in efficiency and usability.

Learning from users how they use the legacy applications helps define tasks and goals as well as what they do with the information they engage with. You can focus on user workflows within the application to ensure you are meeting their needs that may not be addressed in a one-for-one replica. You may also find certain tasks are completed by interacting with multiple applications. By investigating the relationships between applications and tasking, you can uncover redundancies and organize information in a way that better fits the user needs.

* mpolzin@fnal.gov
† mguzman@fnal.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCR006

# DEVELOPMENTS OF PAnTHer: 2D AND 3D WEB MAPS

L. Zambon*, A. Apollonio, Elettra Sincrotrone Trieste, Trieste, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

PAnTHer (Particle Accelerator on THree.js) is an innovative visualization tool designed to create interactive 2D and 3D maps for particle accelerators, leveraging advanced web and touch technologies. These dynamic maps are seamlessly integrated with real-time data sourced from accelerator control systems, simulation outputs, and an external component database.

The mapping process utilizes a JSON-formatted lattice file, which can be generated on-the-fly by extracting essential parameters from a simulator device server. This allows for instantaneous presentation to remote users, complete with visualizations of particle trajectories at refresh rates of up to 25 frames per second.

To ensure efficient and secure data delivery, we have implemented several techniques that facilitate access for external users while maintaining data integrity.

The 3D visualization component was specifically developed to strike an optimal balance between visual fidelity and performance, enhancing the user experience in exploring complex accelerator environments.

This presentation will detail the architecture, functionality, and performance metrics of PAnTHer, highlighting its potential applications in accelerator physics and education.

## INTRODUCTION

PAnTHer is an advanced web-based mapping tool designed specifically for control room personnel, machine physicists, and plant specialists, while also serving as an educational resource for the general public. This tool strikes a balance between fidelity to real-world representations and clarity of information (Fig. 1).

To enhance usability and reduce processing time, unnecessary components and details that could cause confusion have been intentionally omitted or simplified.

The architecture of PAnTHer consists of two primary JavaScript files — one dedicated to 3D visualization and the other to 2D mapping. Additionally, it utilizes "JSON lattice" files for machine data, along with up to five component bundle files for each machine. The system includes a configuration file, core modules, and additional loadable modules as specified in the lattice file. Furthermore, PAnTHer incorporates utilities for data collection from the control system, automated inspection tools, and preprocessing scripts to streamline operations..

## LATTICE FILE

Machine-specific information within PAnTHer is encapsulated in a "JSON lattice" file, a concept inspired by the lattice structures used in particle accelerator simulators. To facilitate this, we employed automated scripts to convert data from Matlab format to JSON [1].

However, this initial conversion served merely as a foundation; we subsequently enriched the dataset with extensive additional information sourced primarily from the mechanical department, as well as various other contributors The nodes situated between the straight sections form the core of the JSON file structure. Typically, the center of these nodes corresponds to the optical center of the bending magnet, while mirrors are also represented as nodes within optical circuits. The edges of the structure denote the walls.

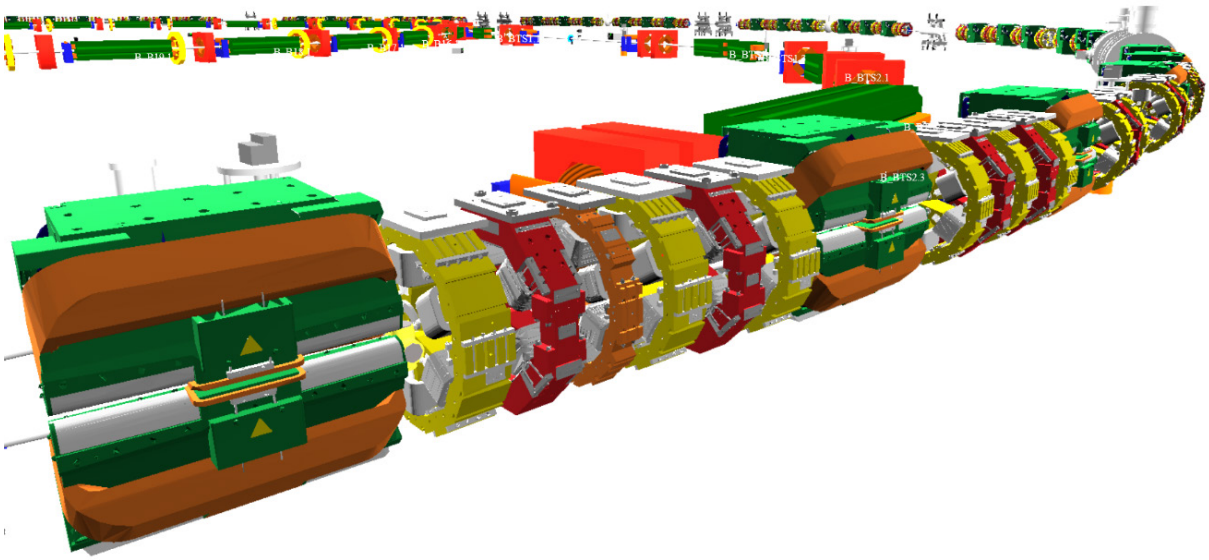


Figure 1: 3D high resolution map of Elettra 2.0.

†lucio.zambon@elettra.eu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MOCR007

# MODERN CONTROL SYSTEM APPLICATIONS AT FERMILAB WITH DART AND FLUTTER

J. S. Diamond[†,1]
[1]Fermi National Accelerator Lab, Batavia, United States of America

*Abstract*

The control system at Fermilab is undergoing an unprecedented modernization effort. Hundreds of legacy applications originally developed with technology from the early nineties will be replaced with a suite of modern web applications. The selection of a user interface framework is a key technology decision that will impact on all the applications developed over the next decade. The Controls department at Fermilab has decided to use Google's Dart language and Flutter framework for future application development. In this paper we will discuss the decision process that selected Dart/Flutter and the development of early general purpose control system applications with the framework.

## INTRODUCTION

Fermilab is the premiere particle accelerator facility in the United States hosting five particle accelerators spread out over 6800 acres. A sophisticated, distributed control system known as ACNET has been at the core of operations for this facility since the era of PDP-11s. Over the decades the control system computing hardware has migrated from mini-computers to modern day multi-core Intel-based servers, but an alarming amount of the original software remained unchanged. Of particular note, and the subject of our work, the user-facing graphical user interfaces used by operators in the Fermilab Main Control Room are still built on top of a custom technology stack from the early 90s written in Fortran or C and using Xlib at its core (Fig. 1).

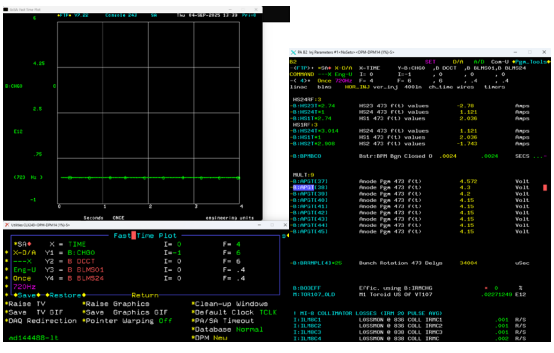


Figure 1: ACNET controls applications based on Xlib.

An attempt was made in the early 2000s to modernize the ACNET control system using Java and Swing. A few applications were developed and adopted by users, most notably the Synoptic display builder but this effort stopped short of replacing over 400 legacy applications (although significant progress was made in migrating software in the middle-tier of the control system).

In 2016, the PIP-II project (a super-conducting linear accelerator expected to turn on and replace Fermilab's original Linac in 2029) selected EPICS as it's control system. This makes the Fermilab accelerator control system hybrid – supporting EPICS for the new linear accelerator and ACNET for the rest of the complex.

To address the legacy ACNET control system infrastructure and to position Fermilab for the next generation of scientific discoveries the ACORN project has been established to modernize the control system. This modernization includes re-imagining a select subset of core and critical ACNET applications (approximately 50-60) as web applications.

## MODERN APPLICATION ARCHITECTURE

The ACORN conceptual design envisions web applications as the future interface to the accelerator control system (Fig. 2). To communicate with the control system, web applications call an API gateway using GraphQL This API gateway aggregates interfaces to a collection of micro-services that form the central tier of the control system. Most of these micro-services expose an API using gRPCs.

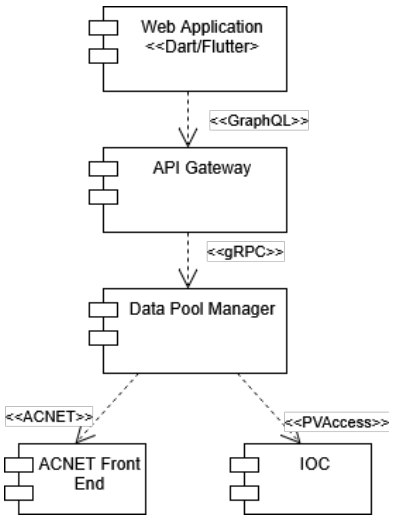



Figure 2: A simplified high-level component diagram of Fermilab's modern control system architecture.

Access to the data-tier is provided by a collection of gateway services, most notably the Data Pool Manager

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODG001

# THE CBXFEL CONTROL SYSTEM: BRINGING NEW TECHNOLOGIES INTO THE SLAC LCLS ACCELERATOR CONTROL SYSTEM*

M.A. Montironi†,[1], W. Lewis[2], S.J. Stein[2], N. Balakrishnan[1], T. Thayer[1], J. Mock[1], C. Curtis[1], M. Dunning[1], A. Le[1], S. Karimian[1], S. Saraf[1], D. Sanchez[1], A. Shetty[1], X. Permanyer[1], M. Balcazar[1], D. Zhu[1], A. Halavanau[1], A. Lutman[1], H.-D. Nuhn[1], Z. Huang[1], Y. Shvydko[3], R. Lindberg[3], D. Shu[3], M. White[3], G. Lanza[1]
[1]SLAC National Accelerator Laboratory, Menlo Park, CA, USA
[2]Osprey DCS, MD, USA
[3]Argonne National Laboratory, Lemont, IL, USA

## *Abstract*

The Cavity-Based Free Electron Laser (CBXFEL) project is proposing to produce a recirculating X-ray cavity and deploy it to the SLAC LCLS (Linac Coherent Light Source) Hard X-ray (HXR) undulator line. The objective of the project is to demonstrate multi-pass gain and to eventually increase peak brightness of the produced X-ray beam. To support this novel application, new electronics and control system architectures have been adopted and integrated into the LCLS accelerator control system. This includes nanoscale precision multi-axis motion control, high-precision in-vacuum temperature control, high-speed and high-resolution USB cameras, and high-speed digitizers. Moreover, the existing accelerator network, vacuum, timing, and machine protection system (MPS) control systems were expanded to support the new devices. This paper describes the main requirements to be met, how the technologies were integrated into the accelerator control system, and the main lessons learned.

## INTRODUCTION

The Cavity-Based Free Electron Laser (CBXFEL) project is proposing to produce a recirculating X-ray cavity and deploy it to the SLAC LCLS (Linac Coherent Light Source) Hard X-ray (HXR) undulator line [1, 2]. Figure 1 shows the experimental layout where the optical cavity, defined by four Bragg-reflecting crystal mirrors (C1, C2, C3, C4), encloses the first seven HXR undulator segments. Dipole magnets in a chicane configuration divert the incoming electron beam around C4 and re-insert it in the first undulator segment to generate x-rays. Another chicane diverts the electron beam exiting the seventh undulator segment, around C1, and re-inserts it in the eighth undulator segment just downstream of C1.

Beyond the C1-C4 crystals, the CBXFEL system is composed of several in-vacuum and in-air beam diagnostics and movers organized in six stations: A, B, C, D, E, and F as shown in Figure 1. Design for an additional station, G, planned for the experimental hutches, will be finalized once Station B installation is completed.

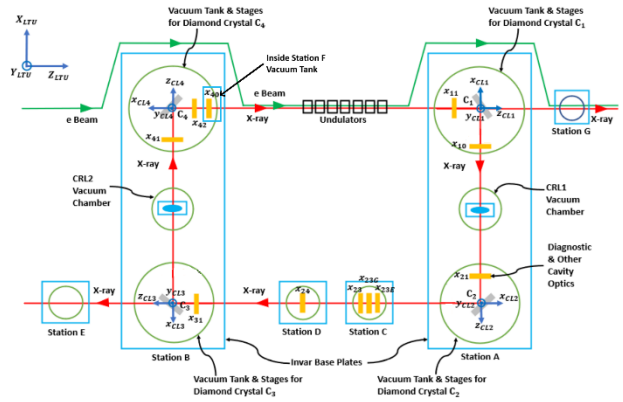



Figure 1: CBXFEL experimental layout.

### *Station A*

Station A, located in HXR cell 21, hosts the C1 and C2 diamond crystals, three X-ray Beam Position/Profile Monitors (X11, X10, and X21), and one Compound Reflective Lens (CRL1)

### *Station B*

Station B, located in HXR cell 13, hosts the C3 and C4 diamond crystals, three X-ray Beam Position/Profile Monitors (X31, X41, and X42), and one Compound Reflective Lens (CRL2)

### *Station C*

Station C, located along the return line, hosts three X-ray Beam Position and Profile Monitors (X23E, X23G, and X23) as well as a beam sampling diffraction grating.

### *Station D*

Station D, located along the return line, hosts one X-ray Beam Profile Monitor (X24).

### *Station E*

Station E, located along the return line, hosts a Beam Position/Profile monitor and two photodiodes.

### *Station F*

Station F, located downstream of Station B in HXR cell 13, hosts a Beam Overlap Diagnostic [3].

* Work supported by US Department of Energy Office of Science Under Contracts DE-AC01-06CH11357 and DE-AC02-76FO0515
† alexmon@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODG002

# OPEN SOURCE EtherCAT MOTION CONTROL, ECMC, AT PSI

A. Sandström†, A. Acerbo, E. Divall, F. Maier, I. Mohacsi, T. Celcer, X. Wang
Paul Scherrer Institute, Villigen, Switzerland

*Abstract*

The open-source EtherCAT motion control EPICS module, ECMC, is now being extensively deployed at the Paul Scherrer Institute (PSI), with the number of motion axes currently exceeding 600. ECMC is used across PSI facilities, particularly for the ongoing upgrade of the Swiss Light Source (SLS 2.0), as well as in SwissFEL and HIPA. The scale and diversity of motion control and data acquisition (DAQ) applications have driven the development of many new features in ECMC. Recent developments include support for PVT (Position-Velocity-Time) motion profiles for advanced fly scanning applications, along with beam-synchronous data acquisition capabilities. A concept for flexible and mobile motion control solutions has also been introduced, aiming to support beamline common-pool equipment and facilitate the integration of user-provided motors. Efforts have also focused on improving user-friendliness, defining best practices, streamlining commissioning, and simplifying troubleshooting processes. This contribution highlights the ongoing development and deployment of ECMC within the complex motion control environments at PSI.

## INTRODUCTION

The Paul Scherrer Institute (PSI) operates four accelerator facilities: the HIPA and PROSCAN proton facilities, and the SLS and SwissFEL electron facilities. The ongoing upgrade of the Swiss Light Source (SLS 2.0) [1] aims to transform it into a fourth-generation synchrotron by renewing the storage ring and replacing large parts of the facility electronics. For this project, ECMC, an open-source EtherCAT-based motion control module [2], has been adopted as the standard motion control platform. The total scope of the upgrade includes approximately 1500–2000 motion axes, most of them located in the beamlines.

ECMC is fully integrated into the EPICS (Experimental Physics and Industrial Control System) [3] environment, and the preferred way of configuration is by the ecmccfg [4] configuration module. ECMC, maintained and developed at PSI, builds upon the open-source EtherLab EtherCAT master [5] to provide real-time control of EtherCAT [6, 7] devices. Its core functionality includes motion control, soft-PLC capability, and data acquisition, expandable with additional features via a modular plug-in concept.

## IMPLEMENTATION ACROSS FACILITIES

Although the SLS 2.0 project is the main driver for ECMC adoption, the framework has also been deployed at SwissFEL and HIPA, where different technical requirements dominate.

### *SLS*

The high number of motion axes and diversity of components have driven efforts to simplify configuration, expand hardware support, and streamline commissioning workflow. Many SLS components require synchronized motion, prompting the development of new features such as Position Velocity Time (PVT) motion. In addition, certain systems are shared between beamlines, leading to the development of a portable motion control concept.

Due to tight schedules and late hardware availability, virtual commissioning techniques have been tested to configure and validate complex systems before physical access.

### *SwissFEL*

While SwissFEL shares many requirements with SLS, its pulsed operation has driven the implementation of beam-synchronous data acquisition features that can capture both motion-related signals as well as other data.

### *HIPA*

At HIPA, applications require interfacing with legacy and radiation-hard components including potentiometers, resolvers, DC motors, and hydraulics that are not typically found in electron accelerators.

## SYSTEM ENVIRONMENT

Common across all facilities, ECMC deployment requires a robust, standardized hardware and software environment to ensure uniform operation.

ECMC systems are deployed on standardized hardware at all facilities. Two controller types are used:

- HP DL20 small-form-factor server (preferred), equipped with four additional network ports supporting up to four ECMC IOCs (EtherCAT masters), Fig. 1.
- Beckhoff C6025-0010 industrial PC, a fan less controller (82 × 127 × 47 mm) mountable on a DIN rail, supporting up to two ECMC IOCs, Fig. 2.

All controllers run Debian 12 with a real-time kernel and the PSI EPICS environment. Systems boot diskless via network provisioning using Warewulf [8], which loads the operating system image entirely into RAM at startup. This ensures uniform environments, rapid updates, and improved reliability by eliminating local storage. Real-time thread jitter below 10 µs has been achieved with this platform.



Figure 1: Preferred controller, HP DL20 server.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODG004

# A SCALABLE APPROACH TO CAMERA INTEGRATION AT NIF

B. Hackel, A. Barnes, C. Estes, D. Tucker, J. Vaher, V. Gopalan
Lawrence Livermore National Laboratory, Livermore, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Integrated Computer Control System (ICCS) at the National Ignition Facility (NIF) manages devices essential to fusion ignition experiments, including 494 specialized cameras supporting beam alignment, target alignment, cryogenic layering, optics damage inspection, collision avoidance, and other key processes. These heterogeneous camera systems, sourced from multiple vendors, use varied bus technologies and Windows-based proprietary SDKs, posing challenges to integrate into a unified control architecture. To address this, ICCS adopted a diskless Front End Processor (FEP) platform using Linux and open-source camera control libraries [1]. Leveraging a network-booted Linux environment and community-supported video drivers, ICCS integrated diverse vendor hardware while enabling rapid adoption of new camera models. Consolidating multiple cameras onto single FEPs improved maintainability and reliability through strategic grouping. Open-source libraries also position NIF as a contributor to the broader controls community. This modernization demonstrates a systematic approach to integrating dissimilar hardware components across a high-stakes facility. By bridging various proprietary protocols, consolidating scattered device controls, and leveraging community-supported software, ICCS achieves a maintainable, expandable architecture for key imaging and streaming functions.

## INTRODUCTION

The Digital Video (DVID) subsystem is an integral part of the NIF control system, underpinning nearly every aspect of the ignition workflow. From laser alignment and wavefront correction to target chamber situational awareness, deuterium-tritium (DT) cryogenic fuel layering and optics damage inspection, many diagnostics require precise, near-real-time imagery. The Integrated Computer Control System (ICCS) serves as the unified framework for camera control, still image acquisition and video streaming. However, the facility's diverse hardware fleet – 24 disparate camera models from ten separate vendors – has historically necessitated unique Windows-based SDKs and dedicated FEP operating system images. As Windows XP approached end-of-life in April 2014, cybersecurity concerns mounted, and sustaining dozens of diverging images became untenable. In response, the ICCS FEP team embarked on a modernization effort to consolidate controls, reduce maintenance burden, and streamline new camera integration.

## THE INTEGRATION CHALLENGE

The sheer diversity of camera make and model throughout NIF (outlined in Table 1) means that each specific DVID implementation could require a specific SDK variant, or a particular model-dependent quirk on top of an existing implementation for another camera. The lack of a common protocol typically resulted in integration times of over two months for new vendor devices, while adding a model from an already-supported vendor required up to a three-week development and testing period. Since each FEP was booting from its own local hard drive, patches to Windows images had to be applied individually, exposing the system to divergent configurations. As modern datacenter metrics demonstrate, multiplying hundreds of individual drives collapses system-level Mean Time To Failure (MTTF) – leading to frequent replacements and escalating maintenance overhead [2]. Maintaining backward compatibility with obsolete camera variants (some 15–20 years old) also ballooned the ICCS FEP codebase. Every incompatible SDK update effectively doubled development and validation effort, as both old and new branches needed to remain live. Vendor heterogeneity, in short, translated into a labyrinth of source code decision branches, mounting technical debt and risk.

Table 1: NIF Cameras by Make and Model

| Manufacturer | Models | Cameras |
|---|---|---|
| Basler | 9 | 345 |
| Cohu (analog) | 1 | 106 |
| PixeLINK | 7 | 20 |
| Spectral Instruments | 3 | 16 |
| Illunis | 1 | 2 |
| Dage (analog) | 1 | 2 |
| Imperx | 1 | 2 |
| Phasics | 1 | 1 |
| **TOTAL** | **24** | **494** |

## A UNIFIED, SCALABLE, DISKLESS LINUX SOLUTION

To break this cycle, NIF's ICCS team architected a single, version-controlled operating environment based on Gentoo Linux [3]. By leveraging PXE network boot [4], all FEPs load one of only two images[1] hosted on a centralized server, eliminating on-disk configuration drift and enabling atomic updates. Camera control is implemented through two open-source libraries: Aravis for GigE Vision [5] devices and libdc1394 [6] for FireWire cameras. These libraries utilize industry-standardized APIs, allowing a single code branchpoint to support multiple camera types.

[1] Due to kernel driver constraints and other external dependences in the respective open-source camera libraries, separate operating system images were required to support FireWire and GigE functionality.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODG005

# DESIGN AND COMMISSIONING OF POSITION CONTROL FOR THE HIGH-DYNAMIC CRYOGENIC SAMPLE STAGE OF THE SAPOTI NANOPROBE AT THE CARNAÚBA BEAMLINE (SIRIUS)

G. O. Brunheira*, V. B. Falchetto, R. R. Geraldes, J. P. S. Furtado, T. R. S. Soares, A. C. Piccino-Neto, Brazilian Synchrotron Light Laboratory, Campinas, Brazil

*Abstract*

Scanning Analysis by Ptycho for Tomographic Imaging (SAPOTI) is the second cryogenic nanoprobe station recently installed and under commissioning at the Coherent X-Ray Nanoprobe Beamline (CARNAÚBA) beamline at Sirius, designed for multi-analytical X-ray techniques, including 2D and 3D ptychographic imaging. A high-dynamic mechatronic system aimed to provide sample positioning at single-nanometer resolution was developed using a decoupled architecture, force actuators, high-speed and high-accuracy metrology and a dynamic filter for reaction forces. This paper presents the design and implementation of a XYZ positioning controller, the fly-scan operation for imaging and results from initial technical commissioning. First, the sample stage mechatronic architecture is overviewed. Next, control strategy based on feedback and trajectory feedforward is detailed, followed by the system identification procedure and design of controllers using loop-shaping and model inversion techniques. Implementation in a real-time digital controller, along with fly-scan trajectories and triggering scheme for detectors are discussed. Finally, in-situ performance in stability, trajectory tracking and imaging reconstruction are demonstrated with experimental results.

## INTRODUCTION

CARNAÚBA is a tender-to-hard X-ray beamline at Sirius, the 4th-generation synchrotron light source of the Brazilian Synchrotron Light Laboratory (LNLS), designed to host two end-stations for imaging and spectroscopy [1]: the operational TARUMÃ [2] and the novel SAPOTI. SAPOTI is a cryogenic nanoprobe that uses Kirkpatrick-Baez (KB) mirrors to deliver diffraction-limited beam sizes from 140 to 35 nm for multi-techniques experiments in 2D and 3D, including ptychography, fluorescence, absorption, diffraction, and luminescence [3].

Following offline integration and characterization [4], the station was installed at CARNAÚBA. The technical commissioning, performed at cryogenic (110–120 K) and non-cryogenic (280 K) sample temperatures, was mostly completed. Scientific commissioning was conducted over three periods totaling four weeks, since beamtime was shared with users from the TARUMÃ. An overview of the system integration and beamline commissioning results, focusing on the thermo-mechanical performance of the sample stage and KB mirrors, will be presented in [5]. This work details the design, implementation and commissioning results of the digital controllers for the XYZ 3D scanning module of the cryogenic sample stage [6].

## SAMPLE STAGE

For a detailed description of the cryogenic sample stage's mechatronic architecture, see [6, 7]. Figure 1 illustrates the key components relevant to this work.

A goniometer stage (GON) with 220° rotation range is located in the base level to select angular positions in $R_y$ direction for 2D mappings and 3D tomography. It is operated by a stepper motor and a rotary optical encoder with a Omron Delta Tau Power Brick controller [8]. As seen in Fig. 1(a), above it is mounted the 3D scanning stage module, known as short-stroke (ShS), which provides positioning control by means of tip-tilt-piston motion $(Rx', Rz', y')$ from a flexure-based parallel kinematic mechanism - a solution successfully adopted in other high-precision systems at Sirius [9]. Figure 1(b) shows the three vertically oriented voice-coils (VCs) in triangular symmetry to provide force actuation between the sample support (SS) and a balance mass (BMS). These two are fixed to the same force frame (FFR) via folder leaf-springs (FLS) and tuned in different compliant suspension frequencies to filter reaction force dynamics.

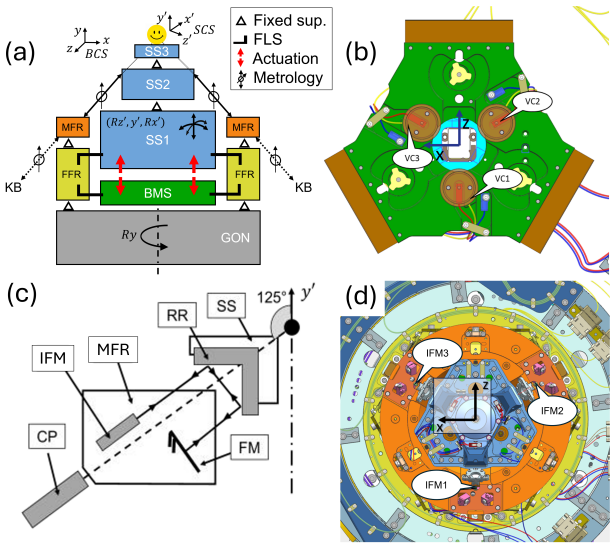



Figure 1: Sample positioning stage from SAPOTI: a) simplified schematic drawing; b) top-down view of balance mass (BMS) with voice-coils triangular distribution; c) metrology schematic; d) top-down view of sample stage, detailing interferometer (IFM) heads triangular arrangement.

* gabriel.brunheira@lnls.br

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODG007

# STANDARDIZED LOW-VOLTAGE DC POWER DISTRIBUTION FOR CONTROL SYSTEM INFRASTRUCTURE AT LCLS

D. Kameswaran †, N. Waters, A. Wallace, R. Nazir
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

Modern accelerator facilities like the Linac Coherent Light Source (LCLS) demand highly reliable and maintainable control system infrastructure to ensure uninterrupted operation and rapid fault recovery. To meet these requirements, LCLS has implemented a centralized, standardized low-voltage DC (LVDC) power distribution architecture that supports mission-critical subsystems, including vacuum controls, system programmable logic controllers (PLCs), and other essential control components.

The architecture features dual-voltage distribution—24 V DC for PLCs and select vacuum systems, and 48 V DC for powering motors. At its core are intelligent DC circuit breakers that offer integrated protection, real-time monitoring, and remote control. These breakers are integrated with input/output controllers (IOCs) and are fully manageable via the Experimental Physics and Industrial Control System (EPICS) software. This integration enables remote power cycling, fault isolation, and system recovery without requiring physical access, significantly improving uptime and maintainability.

The standardized LVDC infrastructure reduces system complexity and supports modular scalability. Integrating with the control system enables advanced diagnostics and continuous monitoring. This work reviews the current LVDC implementation at LCLS, explores the challenges involved in designing DC systems for new applications, and outlines planned enhancements and future developments.

## INTRODUCTION

Modern accelerator facilities such as the Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory operate continuously, supporting experiments where downtime is costly and must be minimized. In these environments, the reliability and maintainability of the control system infrastructure directly influence overall facility performance. Traditionally, industrial control systems have relied on AC/DC converter power modules that are mounted close to the devices they supply. While this distributed approach is common in industrial settings, it poses challenges in the accelerator environment.

Device-level AC/DC converters are often difficult to access, especially when installed deep within equipment racks or close to sensitive beamline components. Troubleshooting or replacing such modules typically requires physical intervention in restricted areas, delaying fault recovery and introducing unnecessary disruption. Moreover, as the number of devices increases, this approach becomes bulky and complex, with multiple converters generating additional heat and vibration near the beamline. These conditions are undesirable in high-precision experimental facilities. To address these limitations, LCLS has developed and deployed a centralized low-voltage DC (LVDC) power distribution system. This architecture consolidates power conversion and distribution, providing a standardized, maintainable, and remotely manageable alternative to conventional distributed power modules.

## MOTIVATION AND CURRENT CHALLENGES

The shift toward a centralized LVDC system was motivated by several recurring challenges with legacy distributed power schemes. First, physical access to AC/DC power supplies placed near the beamline often required interrupting experiments, resulting in significant downtime. Second, the accumulation of power supplies across multiple devices introduced additional points of failure, creating unnecessary maintenance overhead. Third, the heat and vibration generated by localized AC-to-DC conversion were detrimental to experimental stability. Finally, power recovery in the event of faults required on-site intervention, limiting the ability to respond quickly and efficiently.

A standardized LVDC approach addresses each of these challenges by relocating power conversion into modular rectifier chassis, distributing DC power through intelligent panels, and integrating monitoring and control into the Experimental Physics and Industrial Control System (EPICS) [1]. This enables not only higher availability but also system-wide visibility into the power infrastructure, improving fault response and predictive maintenance capabilities.

## FUNCTIONAL REQUIREMENTS

The LVDC system was designed with clear functional requirements tailored to the needs of a 24/7 accelerator facility. At its core, the system must guarantee uninterrupted delivery of reliable DC power to mission-critical subsystems. These include programmable logic controllers (PLCs) that form the backbone of automation, vacuum systems that ensure beamline integrity, and motors used in essential mechanical operations. Equally important is the ability to recover quickly from power-related faults.

In legacy setups, restoring power to a failed subsystem often required technicians/engineers to visit the equipment rack, identify the problem manually, and reset or replace hardware. Such interventions increase downtime and limit

†divyak@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODG008

# ADVANCED POLARIZATION AND ENERGY CONTROL FOR APPLE-II TYPE UNDULATOR BEAMLINES AT MAX-IV

Á. Freitas*, B. Bertrand, M. Holz, M. Lindberg, J. Lindón-Simón, C. Takahashi, H. Tarawneh, L. Zhu
MAX IV Laboratory, Lund University, Lund, Sweden

*Abstract*

Precise control of photon beam properties is essential for modern synchrotron beamlines, particularly those utilizing APPLE-II type undulators. This paper presents the control system architecture developed at MAX-IV using IcePAP drivers and TANGO control system to achieve advanced polarization and energy manipulation. The system implements the BLUES (Beamline Universal Polarization Mode) framework, allowing dynamic control of both helical and inclined polarization states through synchronized phase motor settings. Central to this approach is the use of parametric lookup tables to define nonlinear motion trajectories for the undulator gap and phase axes. This system enables linear energy ramps, supporting constant electron volts per second (eV/s) scans crucial for high-resolution spectroscopy and imaging techniques, taking full advantage of the high flux provided by fourth-generation light sources and improving data collection efficiency without compromising the stability or quality of the photon beam. Integration between the beamline and accelerator control systems allows for the complex coordination required to manage polarization settings. To ensure electron beam stability during undulator motion, the control system integrates feedforward correction loops that compensate for orbit and optics distortions induced by gap and phase changes. This approach offers a scalable and precise method for enhancing beamline capabilities, tailored specifically for the challenges posed by APPLE-II undulators.

## INTRODUCTION

The MAX-IV Laboratory in Lund, Sweden, is the first diffraction-limited storage ring in the world. It is composed by two storage rings at 1.5 GeV and 3 GeV with a full energy injector (linac). The linac serves as both an injector for the storage rings and a driver for the Short Pulse Facility (SPF), enabling femtosecond-scale time-resolved experiments [1].

Soft X-ray beamlines at MAX-IV are primarily located on the 1.5 GeV ring and selected sections of the 3 GeV ring. These beamlines rely on APPLE-II type undulators, which are elliptically polarizing devices capable of producing variable polarization states linear (horizontal and vertical), circular, and elliptical by adjusting the relative phase of magnet arrays. This capability is essential for experiments requiring precise control over photon polarization, such as X-ray magnetic circular dichroism (XMCD), X-ray linear dichroism (XLD), resonant inelastic X-ray scattering (RIXS), X-ray absorption spectroscopy (XAS), and ambient-pressure photoelectron spectroscopy (APXPS). Polarization control is a fundamental requirement in synchrotron-based soft X-ray spectroscopy. The polarization state of the incident photons—linear, circular, or elliptical—directly influences the interaction mechanisms between X-rays and matter, thereby determining the sensitivity and specificity of various spectroscopic techniques. Figure 1 illustrate the polarizations schema.

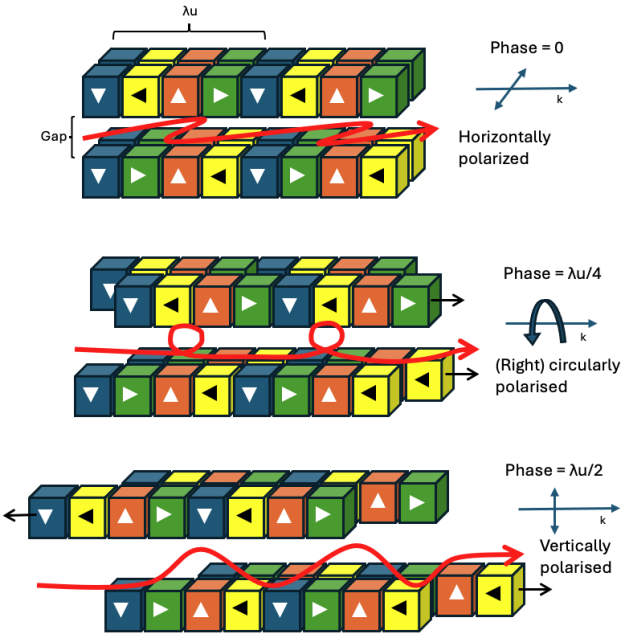



Figure 1: Polarization schema example of APPLE-II.

Currently, MAX-IV operates several APPLE-II undulators built in-house [2], including those that serve the HIPPIE, Veritas, and SoftiMAX beamlines on the 3 GeV ring. These devices offer high brilliance and polarization flexibility across the soft X-ray spectrum. Additionally, the BLOCH beamline on the 1.5 GeV ring utilizes a quasi-APPLE-II undulator, extending the facility's polarization control capabilities at lower photon energies.

## MAX-IV APPLE-II CONTROL SYSTEM

The MAX-IV APPLE-II undulator is driven by eight independent servomotors: $x_1, x_2, x_3, x_4$ control the longitudinal phase (phase motors), and $z_1, z_2, z_3, z_4$ control the magnetic gap. The system uses IcePAP [3] controllers to interface with the servo drives, allowing the servomotors to be controlled in a stepper-like fashion. Position feedback is provided by absolute encoders, enabling closed-loop control via an integral controller. Additional PLC-connected encoders are

* aureo.freitas@maxiv.lu.se

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODR001

# READOUT OF THE FPGA-BASED GLOBAL TRIGGER SYSTEM OF THE ATLAS EXPERIMENT

T. Xu* on behalf of the ATLAS TDAQ Collaboration
Argonne National Laboratory, Lemont, IL, USA

*Abstract*

A new FPGA-based Global Trigger system is intended for the Phase-II Upgrade of the ATLAS experiment at the High-Luminosity Large Hadron Collider (HL-LHC). The system will process data from the experiment with fixed latency to allow the selection of individual collisions of proton bunches with physical potential. Intermediate data from the Global Trigger system are read out for the collisions of interest for trigger decision verification with commodity computing. The readout of the Global Trigger is handled by the readout firmware, which interfaces with the ATLAS readout system. The firmware receives trigger decisions, timing and control signals via the 9.6 Gb/s ATLAS Local Trigger Interface (LTI) link and outputs data via the 25 Gb/s Interlaken link. A prototype of this readout firmware has been designed and tested on the Global Common Module (GCM) boards equipped with AMD Versal Premium FPGAs.

## INTRODUCTION

ATLAS is the largest general-purpose particle detector experiment at the Large Hadron Collider (LHC) at CERN [1]. It is designed to study a wide range of physics, including the Higgs boson and physics beyond the Standard Model. Using advanced tracking, calorimetry, and muon systems, it provides precision measurements of particles and interactions. The high-luminosity phase of the LHC (HL-LHC) will provide an order of magnitude higher collision rate [2]. To cope with this challenge, ATLAS is being upgraded to handle higher collision rates and improve sensitivity, enabling discoveries of rare processes and new phenomena. The LHC collides proton bunches every 25 ns, producing 40 million bunch crossings per second. An event is defined as the set of detector-sampled data from a single crossing.

As part of the HL-LHC upgrade, ATLAS will introduce a new Level-0 Trigger [3] operating at an increased maximum output rate of 1 MHz with 10 µs latency. The enormous output data rate from the ATLAS detector makes it impossible to record every event. To solve this, ATLAS employs a trigger system to select events of interest [4]. Its core component, the Global Trigger, runs complex algorithms on full-granularity calorimeter data on the Global Common Module (GCM), a common hardware platform for the Global Trigger [5]. This design enables the implementation of algorithms such as topological clustering, refined particle identification, pile-up suppression, jet finding, exotic-object selection [4], while improving scalability and maintainability.

As illustrated in Fig. 1, the Global Trigger system consists of three major firmware layers: Data Aggregation and Multiplexing (MUX), Global Event Processor (GEP), and Global to Central Trigger Interface (gCTPi). Each layer consists of multiple firmware nodes, with each node targeted to an FPGA on a GCM board. The Global Trigger system receives various raw detector data from the source systems and processes them through the nodes in each layer. Within each node, a common readout firmware is integrated to support the readout of data of interest as selected by the Central Trigger Processor (CTP). The CTP sends trigger decisions to the Local Trigger Interface (LTI) card, which forwards them to a Front-End Link eXchange (FELIX) system [6] for distribution to all nodes in the Global Trigger system. Within the Global Trigger system, the term event data refers to the processed outputs of the MUX, GEP, and gCTPi nodes. Each node's readout firmware decodes the LTI messages received from FELIX, uses them to select the corresponding event data, and passes the selected data back to FELIX, which in turn delivers them to the data handler (a host-based software package) for storage and further processing.

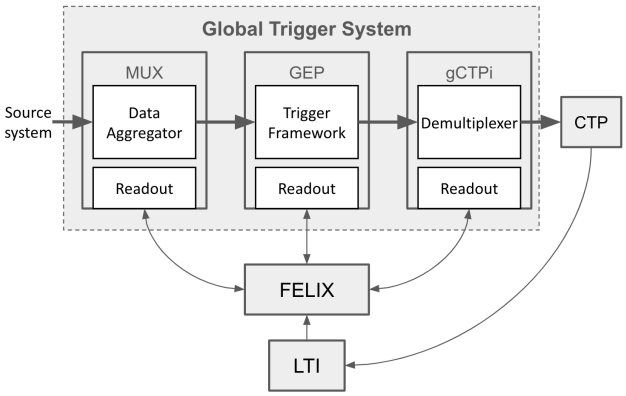



Figure 1: Global Trigger system data flow.

## INTERFACES

The readout firmware on each node provides the following major interfaces: 9.6 Gb/s LTI link and 25 Gb/s data link to FELIX, application and event data input buses as shown in Fig. 2. As the readout firmware is common to MUX, GEP, and gCTPi nodes, it provides a common application and event data interface to each node. Currently, the readout firmware has been integrated with the GEP Trigger Framework firmware.

### *Interface to FELIX*

Each GCM node is connected to FELIX through two optical links, including the 9.6 Gb/s LTI link from FELIX to

* tong.xu@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODR003

# MODERNIZING EMBEDDED CONTROLLERS AT THE NIF

A. Barnes, V. Gopalan, B. Hackel, T. Hillard, W. Massey, A. Pao, P. Singh, K. Yu
Lawrence Livermore National Laboratory, Livermore, CA, United States

*Abstract*

As the world's most energetic laser, the National Ignition Facility (NIF) plays a critical role in advancing high energy density physics and inertial confinement fusion research. The NIF relies on a distributed control system to automate setup and execution of experiments. This includes over 1,000 embedded controllers split between 17 distinct types. Most of these controllers were designed in the late 1990s to early 2000s, using platforms ranging from Echelon Neuron microcontrollers to STD Bus computers.

Over the decades, many of our chosen software technologies, hardware components, and platforms have reached end of life. Therefore, we have begun a modernization effort for the NIF embedded controllers. Modern hardware/software will allow us to procure additional units in order to maintain adequate spares and support upcoming upgrades to the NIF. It will also allow us to tackle incompatibilities developing between our existing firmware and modern IT infrastructures. This paper will provide an overview of our embedded controller strategy. It will focus on component selection, minimizing risk during the transition using an automated test bench, and firmware/hardware upgrades to minimize on-going maintenance costs.

## WHY MODERNIZE?

As anyone who's works on a long running project knows, modernization efforts come with significant costs and risks. Therefore, projects normally try and maximize the lifetime of their existing hardware. For the NIF embedded controllers, a number of factors have come together driving our need to modernize.

At this time, our primary driver stems from the fact that we cannot purchase new units of our Power Conditioning Embedded Controllers (PCECs). These were originally developed in the late 1990s and built using the STD32 Bus. Unfortunately, since the turn of the century, industry has rapidly moved away from the STD Bus. As a result, the number of manufacturers willing to produce STD Bus components has plummeted, and the costs have skyrocketed. For example, in 2018 we only found a single manufacturer willing to produce an STD Bus CPU board. Since production would require retooling a manufacturing line, they required a minimum order of 1,000 boards. For reference, NIF has less than 250 units in total (production, offline labs, and spares). Given that we'd run into similar issues with the IO boards, we decided to focus on refurbishing components on our existing boards. While this helped maintain our pool of spares, unfixable failures still reduce the number of spares each year. In addition, NIF is now looking to upgrade our power conditioning system as part of an Enhanced Yield Campaign (EYC). The number of new units required exceeds the number of spares we have on hand.

Our next driver comes from incompatibilities between embedded controllers and our Information Technology (IT) infrastructure. Our complex ECs use operating systems that were modern at the time of their design (such as older versions of VxWorks and Embedded Windows XP). These legacy operating systems have led to a number of issues in recent years. First, both vendors and the open-source community are actively dropping support for legacy file system protocols (such as SMBv1 used by Windows XP). Second, both Microsoft Visual Studio and gcc compilers have dropped support for Windows XP. While we can still use older versions of the compilers, they require special considerations to address known security vulnerabilities. Consequently, the Controls and IT teams collaborated on deploying mitigations for these cybersecurity concerns. Finally, as the age between legacy components and more modern systems increase, compatibility issues are beginning to arise.

Another driver, and headache for IT, comes from the need to support legacy build machines maintained by software developers. In some cases, the original Software Development Kit (SDK) came with a node locked license. Since then, the vendor's IP was purchased by another company unable to support transferring the old licenses to new machines. In other cases, libraries used by our ECs, such as OpenCV, have undergone major API revisions. Upgrading the build machine to a modern OS would require upgrading the external library. Doing so would necessitate a major rewrite of the embedded controller. Developer maintained build machines also introduce the risk of knowledge loss in the event an employee leaves the organization. Frequently, the next developer to work on the EC must rediscover how to build it.

In addition to the issues with the legacy hardware, we hope to improve our ECs during the modernization process. Over the years, we've identified a number of issues directly caused by design decisions made in the original implementation. For example, one of our ECs responds to queries with only the numeric or string value requested. The response provides no information to link the EC response with the initial query. This has resulted in numerous cases where communications desynchronized due to temporary issues (network or EC performance).

We also hope to make improvements to minimize ongoing maintenance costs for the ECs. The largest potential reduction comes from consolidating the number of hardware platforms. Various different teams independently designed NIF's current ECs. As a consequence, NIF ended up with a unique hardware platform from each team. Also, as time progressed, some of the teams changed their

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODR004

# DISTRIBUTED I/O TIER AS A REFERENCE PLATFORM FOR HARNESSING SYSTEM-ON-CHIPS IN CERN'S CONTROL SYSTEM: GATEWARE DESIGN, BUILD SYSTEM, AND SOFTWARE SERVICES

A. Arias Vázquez*, CERN, Geneva, Switzerland
V. Amoiridis, G. Daniluk, A. Pinho, B. Guncic, M. Lettrich, F. Vaga, CERN, Geneva, Switzerland

*Abstract*

The Distributed I/O Tier (DI/OT) project was initially launched to develop a common, modular hardware platform for custom electronics at the lowest layer of the CERN control system. With the adoption of the AMD Zynq UltraScale+ MPSoC for the high-performance System Board, DI/OT has also become a reference platform for integrating System-on-Chip (SoC) technology into CERN's control system. This paper presents two key aspects of DI/OT's role as a SoC reference platform: (1) tools and methodologies to streamline end-application development and (2) integration of DI/OT into CERN's control system as a Front-End platform. The first aspect includes a user-friendly build system and reference design that enable seamless integration of custom FPGA IP cores and Linux device tree entries while providing the reference design with essential DI/OT functionality and monitoring interfaces for local crate peripherals. This build system also automates synthesis and low-level software compilation to generate a complete bootable binary. The second aspect covers a fail-safe and reliable SoC boot mechanism, network booting of FECOS (a Debian-based CERN Linux image for Front-End Computers), and integration with standard monitoring services.

## INTRODUCTION

The use of System-on-Chip (SoC) platforms has grown rapidly within the Accelerator and Technology Sector (ATS) at CERN, driven by their ability to combine real-time logic, embedded processing, and flexible I/O in a compact form factor. This trend is expected to intensify in the lead-up to Long Shutdown 3 (LS3) of CERN's accelerator complex, as new systems demand tighter integration, higher performance, and lower latency at the front end.

The DI/OT project [1] was the first SoC-based platform developed specifically for integration into the CERN control system. Its early adoption of the Zynq MPSoC, combined with a modular hardware design and a complete software stack, established it as the reference platform for subsequent SoC developments in ATS.

However, as more SoC platforms emerge, common challenges are being faced independently: managing gateware complexity, ensuring safe boot and configuration, integrating with control system services, and supporting standard workflows. This fragmentation highlights the need for a unified SoC framework where the DI/OT project becomes the reference.

* alen.arias.vazquez@cern.ch

## DI/OT PLATFORM OVERVIEW

The Distributed Input/Output Tier (DI/OT) project provides a centrally supported modular hardware kit for building custom electronics systems in the lowest tier of the control stack, directly interfacing to accelerator equipment. The basic kit is illustrated in Fig. 1. The modules are housed in a 3U Europa crate featuring an in-house designed, passive backplane compliant with the CompactPCI Serial (CPCI-S.0) standard. Having a fully passive backplane prevents potential radiation-related issues in this critical component [2].

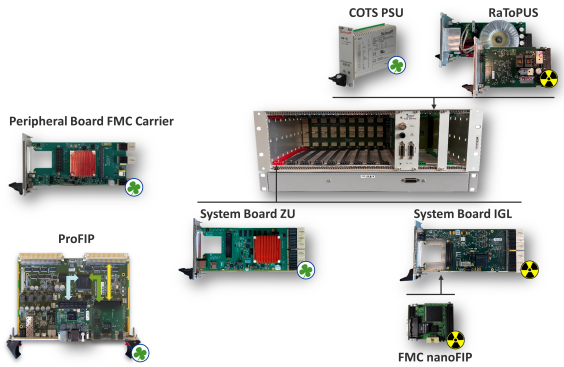



Figure 1: Distributed I/O tier hardware kit [1].

As shown in Figure 1, the leftmost slot of the crate hosts the system board, while a star topology in the backplane allows high-speed point-to-point communication between the system board and the peripheral boards.

There are two variants of the platform: radiation-tolerant and non-radiation-tolerant [3]. While they share exactly the same 3U crate, they differ by the hardware modules hosted inside that DI/OT crate:

- **Radiation-tolerant variant:**
  - Custom-designed power supply (RaToPUS) developed in-house.
  - System board featuring a flash-based Microsemi IGLOO2 FPGA.
- **Non-radiation-tolerant variant:**
  - Off-the-shelf, CPCI-S power supply.
  - System board based on an AMD Xilinx Zynq UltraScale+ MPSoC.

Rather than replacing established platforms such as PLCs or VME, DI/OT is designed to address two specific but essential use cases:

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODR005

# NEXT GENERATION DIRECT RF SAMPLING LLRF CONTROL AND MONITORING SYSTEM FOR LINEAR ACCELERATORS *

C. Liu†, E. Snively, R. Herbst, K. Kim, E. A. Nanni
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

The low-level RF (LLRF) systems for linear accelerating structures are typically based on heterodyne architectures. The linear accelerators normally have many RF stations and multiple RF inputs and outputs for each station, so the complexity and size of the LLRF system grows rapidly when scaling up. To meet the design goals of being compact and affordable for future accelerators, or upgrading existing ones, we have developed and characterized the next generation LLRF (NG-LLRF) platform based on the RF system-on-chip (RFSoC) for S-band and C-band accelerating structures. The integrated RF data converters in RFSoC sample and generate the RF signals directly without any analogue mixing circuits, which significantly simplified the architecture compared with the conventional LLRF systems. We have performed high-power tests for the NG-LLRF with the S-band accelerating structure in the Next Linear Collider Test Accelerator (NLCTA) test facility at SLAC National Accelerator Laboratory and a C-band structure prototyped for Cool Cooper Collider (CCC). The NG-LLRF platform demonstrated pulse-to-pulse fluctuation levels considerably better than the requirements of the targeted applications and high precision and flexibility in generating and measuring the RF pulses. In this paper, the characterization results of the platform with different system architectures will be summarized and a selection of high-power test results of the NG-LLRF will be presented and analyzed.

## INTRODUCTION

The conventional low-level RF (LLRF) systems of particle accelerators are typically implemented with heterodyne-based architectures. The RF signals are up and down mixed from the baseband with analog RF mixers. There are many RF input and output channels in large accelerators, so the size and cost of analog circuits of the LLRF system could become substantial. With the latest direct RF sampling technology of the RF system-on-chip (RFSoC), we have designed and implemented the next-generation LLRF (NG-LLRF) system for linear accelerators (LINACs) [1], including the future collider concept, the Cool Cooper Collider ($C^3$) [2] and the upgrade solution for the S-band accelerator test station at the Next Linear Collider Test Accelerator (NLCTA) [3], which have stringent RF performance requirements for LLRF systems.

RFSoC technology has been applied to a range of high-energy physics and astrophysics instruments with RF frequencies up to 6 GHz and we have performed RF performance characterizations with different RFSoC devices specifically configured for different applications [4–9]. The development of NG-LLRF at SLAC begins with a C-band linear accelerator(LINAC). The prototype has demonstrated low amplitude and phase fluctuation levels for C-band applications [10–12] and a selection of the evaluation results will be summarized in this paper. There are many S-band RF stations around SLAC, including the S-band test stations in test laboratories, the S-band station in NLCTA, and the RF guns and approximately 200 klystrons of LCLS [13] and FACET. The NG-LLRF could be an excellent candidate for upgrading the LLRF systems for those RF station. To start the exercise, we have configured the NG-LLRF for the S-band test station at NLCTA and performed some initial performance evaluation [3]. In this paper, the test configuration used for the initial high-power test results of the NG-LLRF with the S-band station at NLCTA will be introduced and some test results will be discussed.

## NG-LLRF FOR C-BAND LINACS

### *NG-LLRF System Architecture*

Figure 1 shows the block diagram of the NG-LLRF with two rf input channels and a single rf output channel for a single set of accelerating structure cavities. However, the NG-LLRF system can have up to 16 rf input channels and 16 rf output channels, which can be customized per application.

The center frequency of the C-band NG-LLRF is around 5.712 GHz and the bandwidth is 245.76 MHz. The ADCs direct sample the RF signals with sampling rate of 2.4576 giga samples per second (GSPS) at higher order Nyqusit zone. The digital samples are digitally down-mixed in a harden IP implemented in RFSoC. The baseband output of the mixer in in-phase (I) and quadrature (Q) format is then decimated and filtered. The feedback control algorithm block in the firmware reads IQ components of the down-converted cavity signal and calculates a new set of IQ values based on user input parameters. The updated IQ is modulated with the user-defined baseband pulse. The modulated pulse is then interpolated and digitally up-converted with a hardened IP implemented in the RFSoC. The updated RF pulse is generated by the integrated DAC in RFSoC at sampling rate of 5.89824 GSPS and the RF frequency is in the second Nyquist zone.

### *Pulse-to-Pulse Stability with High-Power Test*

Pulse-to-pulse fluctuation level is critical for an LLRF system. We measured phase jitter within 60 consecutive RF pulses at four different stages with appropriate attenuation,

* This work was supported by the U.S. DOE, Office of Science contract DE-AC02-76SF00515.

† chaoliu@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-MODR007

# DESIGN OF A STANDARDIZED FPGA ARCHITECTURE FOR EIC COMMON PLATFORM DAUGHTERCARDS*

P. Bachek†, K. Mernick, S. Mai, M. Costanzo, G. Narayan
EIC, Brookhaven National Laboratory, Upton, NY, USA

*Abstract*

The EIC Common Platform is a modular system architecture which will serve as the basis for the EIC Controls Systems. It consists of a SoC based Carrier with up to two independent pluggable FPGA based Daughtercards. Different types of Daughtercards have custom electronics catering to the specific needs of an application. All types of Daughtercards will have FPGA logic to support a common protocol for communication with the Carrier as well as a basic set of features for programming and telemetry. Logic to support Daughtercard specific functionality will be implemented in the same FPGA. Daughtercard FPGA projects will be organized with a common modular structure to facilitate reuse of IP cores while allowing for flexibility within the Daughtercard specific logic design. The FPGA Firmware Framework (FWK) developed at DESY will be leveraged for managing the generation and building of FPGA projects. The basic functionality and organizational structure of EIC Common Platform Daughtercard FPGA projects is presented.

## INTRODUCTION

To support the needs of the EIC controls system, the EIC Common Platform has been developed. It constitutes a standardized hardware platform that will be shared by many subsystems including RF, Instrumentation, and Controls. The Carrier supports two pluggable Daughtercards for flexibility with different application specific functionality requirements.

Several different types of Daughtercards are being developed with features such as Digital I/O, SFP networking, ADC, DAC, RF ADC, RF DAC, and a BPM processor. Although each Daughtercard will have unique electronics catering to its specific features, they will all share much of the same hardware design. This decreases the development time required for designing new Daughtercards and helps to reduce issues introduced by completely unique custom approaches.

Like the physical hardware, the FPGA logic design of each Daughtercard will differ depending upon the application. But all Daughtercards will include a common set of FPGA logic to support a standard set of features. Code reuse enables much faster bring-up of new Daughtercards and again reduces bugs introduced by different implementations of the same thing.

* Work supported by Contract Number DE-AC02-98CH10886 with the auspices of the US Department of Energy
† pbachek@bnl.gov

## HARDWARE DESIGN

All Common Platform Daughtercards will use the same FPGA, or a pin compatible one from the same family. They will also share the same power regulator circuitry for the necessary FPGA power rails. Every Daughtercard will have 64 MB of NOR flash for nonvolatile configuration memory, and a 256 Byte EEPROM which includes a preprogrammed 48-bit unique id.

The connector interface to the Carrier is standardized for all Daughtercards. The Carrier directly controls the configuration pins of the Daughtercard FPGA. The Carrier can select one of two configuration modes: QSPI flash or Slave Serial. In QSPI flash mode the Daughtercard configures itself from the local flash configuration memory. In Slave Serial mode, the Carrier sends the bitstream to the Daughtercard over serial.

The Daughtercard EEPROM is directly accessible by the Carrier via I2C. This allows the Carrier to configure a Daughtercard on powerup by reading an identifier code from the Daughtercard EEPROM. The Carrier can then identify the proper bitstream to use for configuring the Daughtercard. Once the Daughtercard is programmed the Carrier can also access the Daughtercard flash memory for reading or writing a new nonvolatile image.

## FPGA PROJECT HEIRARCHY

The FPGA logic is designed in such a way as to facilitate as much code reuse as possible. It is broken up into three main hierarchical blocks Daughtercard Common, Daughtercard BSP, and Daughtercard Application. The Daughtercard Common section is shared logic between all Daughtercards. The Daughtercard BSP is unique to each type of Daughtercard. The Daughtercard Application is specific to each unique application of a particular type of Daughtercard.

A simplified block diagram of the project structure is shown in Fig. 1. The Daughtercard BSP is the top-level entity and must instantiate the Daughtercard Common block. It may or may not include the Microblaze Processor depending upon the needs of that particular Daughtercard, but this diagram shows an example of what a typical project might look like. Similarly, for the Daughtercard Application the internal structure will change but the port map will remain the same.

### *Daughtercard Common*

Every Daughtercard will support high bandwidth communication with the Carrier via a high-speed serial link. This is implemented using the AXI Chip2Chip IP coupled with the Aurora 64b66b IP from AMD for 10 Gbps serial

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUAG001

# ELECTRON ION COLLIDER PROJECT AND CONTROL SYSTEM UPDATES FOR 2025*

J. P. Jamilkowski†, M. Costanzo, T. D'Ottavio. C. De Silva, W. Halbig, M. L. Kabir, K. Kulmatycski, K. Mernick, C. Montag, S. Nagaitsev, S. Nemesure, L. Nguyen, V. Ptitsyn, V. Ranjbar, K. Shroff, K. Smith
Brookhaven National Laboratory, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

As the Relativistic Heavy Ion Collider (RHIC) completes its final physics run in 2025, design activities for the future Electron Ion Collider (EIC) that will probe the building blocks of nuclear physics for decades to come have made critical advances. Recent improvements including a new Electron Injector System and updated Low Energy Cooler for electron-based cooling of hadron beams will be described. Advancements in the planning and demonstration activities for accelerator controls elements such as the Common Platform front-end computer designs and related infrastructure, EPICS-based software infrastructure, bridging tools to controls for legacy systems required to support hadron injector systems, and networking and computing infrastructure designs will be highlighted. Our analysis of potential scope of AI/ML integration with controls for the EIC accelerator and detector systems will be introduced.

## COMPLETING THE RHIC ERA

The RHIC has been the workhorse since 2001 for Nuclear Physics research supported by the US Department of Energy (DOE) involving colliding beams of polarized protons or a range of heavy ions species between energies 3.85 GeV [1] and 255 GeV [2] for experiments that have themselves significantly improved in terms of capabilities over the last two decades [3]. The current expectation is that RHIC beam operations will permanently cease at the end of 2025, which amounts to a short delay of several months from what was predicted in our previous report.

Plans to reuse RHIC facilities and equipment have advanced in recent years. We still intend to adapt the RHIC Tunnel to host the Hadron Storage Ring (HSR) and Electron Storage Ring (ESR), where the HSR will primarily be comprised of magnetic elements from the existing Yellow Ring with modifications. A limited portion of the Blue Ring will be adapted to serve as the 41 GeV Bypass Line, though the location of that segment has moved to Arc Sector 1 to mitigate space concerns in the Collider Tunnel.

## A REVISED INTRODUCTION TO THE EIC

The basic concepts of the EIC design remain close to those described in our 2023 report [4]. In addition to the previously described HSR and ESR machines that provide up to two Interaction Regions where physics experiments can be performed, EIC will require a photo-injector, a LINAC, and a Rapid Cycling Synchrotron (RCS) to accelerate polarized electrons that can be utilized in the ESR. Between 2024 and 2025, it was decided that physics and logistical issues involving tunnel installation of the new accelerator rings in the existing tunnel used by RHIC would require a revised set of designs for the Electron Injector Systems (EIS). This culminated in an updated design for the EIS that is comprised of an updated 750 MeV LINAC with a revised RCS Ring operating between 5 and 18 GeV and a newly added Beam Accumulator Ring (BAR) that are planned to be constructed in a new facility adjoining the Collider Tunnel in the existing 4 o'clock region. Figure 1 shows the updated cartoon layout for the entire EIC Project.

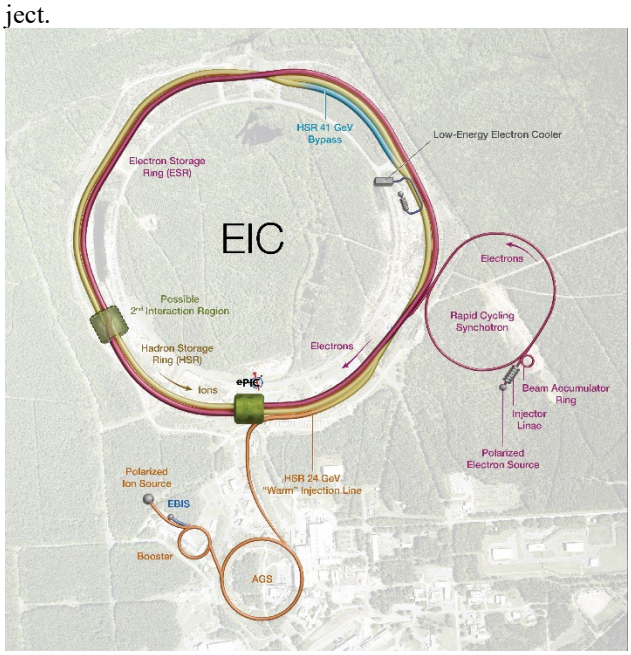



Figure 1: An updated cartoon layout of EIC as of September 2025.

Another set of strategic decisions that have been made recently is the removal of the Strong Hadron Cooler (SHC) systems from the Project and the concentration on providing a Low Energy Cooler (LEC) that uses a co-propagated electron beam to cool hadron beam in the HSR only at the injection energy for that Ring, which would still be installed in the 2 o'clock region of the Collider Tunnel. The motivation for these decisions is to reduce overall Project costs and has been supported by recent assessments of the limited projected benefit to beam luminosity yielded by the SHC design. The overall impact to Controls scope due to

* Work supported by Contract Number De-AC02-98CH10886 with the auspices of the US Department of Energy.
† jpjamilk@bnl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUAG002

# THE SMALL SIZE TELESCOPE CONTROL SYSTEM DESIGN OF THE CHERENKOV TELESCOPE ARRAY OBSERVATORY

V. Conforti[,1], G. Tosti[2], V. Pastore[1], P. Bruno[3], S. Germani[2], G. Giavitto[4], S. Iovenitti[5], A. Marchetti[5], E. Rol[4], A. Sulich[7], A. Trois[6], V. Voitsekhovskyi [4], J. Watson[4], R. White[4]
[1]INAF Osservatorio di Astrofisica e Scienza dello Spazio di Bologna, Bologna, Italy
[2]Università di Perugia, Dipartimento di Fisica e Geologia, Perugia, Italy
[3]INAF Osservatorio Astrofisico di Catania, Catania, Italy
[4]Deutsches Elektronen-Synchrotron, Berlin, Germany
[5]INAF Osservatorio Astronomico di Brera, Milano, Italy
[6]INAF Osservatorio Astronomico di Cagliari, Selargius, Italy
[7]INAF Osservatorio Astronomico di Trieste, Trieste, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Cherenkov Telescope Array Observatory (CTAO) will include telescopes of three different sizes, the smallest of which are the Small-Sized Telescopes (SSTs). In particular, the SSTs will be installed at the southern site of CTAO, on the Chilean Andes, and will cover the highest energy range of CTAO (up to ~300 TeV). The SSTs are developed by an international consortium of institutes that will provide them as an in-kind contribution to CTAO. The optical design of the SSTs is based on a Schwarzschild-Couder-like dual-mirror configuration. They are equipped with a focal plane camera based on SiPM detectors. The Telescope Control System (TCS) is the system responsible for the control and supervision of each telescope. The TCS includes several supervisor components that interface with the telescope local control systems, the hardware and software that control the telescopes hardware devices such as the telescope mount drive systems and the Cherenkov camera. The TCS is also the interface between the telescope and the CTAO Array Control and Data Acquisition system (ACADA). As far as the mechanical structure is concerned, the TCS is also derived from what has already been developed within the ASTRI project. The design of the SST telescopes was evaluated and approved during the Critical Design and Manufacturing Readiness review (CDMR) organized with CTAO. In this contribution we will present the design of the Telescope Control System, including the results of the CDMR.

## INTRODUCTION

The Cherenkov Telescope Array (CTA) project [1] represents a major advancement in ground-based gamma-ray astronomy, aiming to provide unprecedented sensitivity and resolution through an extensive array of telescopes of different sizes. Within this ambitious framework, the Small-Sized Telescope (SST) sub-array [2] plays a crucial role in detecting gamma rays at the highest energies, thus significantly contributing to CTA's overall scientific capabilities. Central to the operational success of the SST is its Telescope Control System (TCS), responsible for coordinating and managing all telescope subsystems, ensuring precise control and seamless integration within the array. The TCS provides essential functionalities such as telescope pointing, tracking, camera synchronization, thereby directly influencing the scientific output and reliability of the observations. The design of the SST structure is almost the same of the ASTRI Mini-Array telescopes currently under construction at the Teide Observatory [3]. Consequently, parts of the SST TCS software design, particularly those interfacing with telescope hardware, closely follow the ASTRI Mini-Array model. This shared design approach significantly enhances the probability of successful implementation and reduces development time. Additionally, some SST telescopes are financed through Italy's National Recovery and Resilience Plan (PNRR) [4], imposing stricter deadlines compared to the other SST telescopes. This factor necessitates efficient development practices and rigorous schedule adherence. This paper presents the current development status of the SST-TCS. We outline the ongoing design refinements, which are guided by the feedback received from the recent Critical Design & Manufacturing Review (CDMR). We then detail our implementation strategy, which follows a dual approach: adapting core components from the proven ASTRI software framework while concurrently developing the new, high-level manager applications. Finally, instead of presenting formal results, we discuss the initial and crucial phase of integration testing, focusing on the preliminary hardware-in-the-loop activities to verify the interface with our hardware simulators.

## OVERVIEW OF THE SST PROJECT

The Small-Sized Telescopes (SSTs) [5] are one of three types of telescopes that will make up the Cherenkov Telescope Array Observatory (CTAO). The SST sub-array will be installed at the CTAO southern site in the Chilean Andes and is designed to cover the highest part of the observatory's energy range, up to approximately 300 TeV. The SSTs are provided as an in-kind contribution to CTAO by an international consortium of over 20 institutes from countries including Italy, France, Germany, the UK, Switzerland, and Japan. The current plan for the observatory's "alpha-configuration" includes the construction and installation of 37 SSTs. The additional SSTs will be delivered earlier and have been found by PNRR. The SSTs' optical design is a dual-mirror, modified Schwarzschild-Couder

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUAG005

# IFMIF-DONES CONTROL SYSTEMS: GENERAL ARCHITECTURE

C. Carvajal[†,1], J. Aguilar[1], M.J. Gutierrez[1], R. Lorenzo[1], L. Maindive[1], A. Moreno-Cortes[1], I. Podadera[1], C. Torregrosa[1], F. Martín-Fuertes[1,2], H. Coronel[3], E. Botta[4], J. Cruz[5], M. Cappelli[6]
[1]IFMIF-DONES Spain Consortium, Granada, Spain
[2]CIEMAT Madrid, Spain
[3]ATG Science and Engineering S.L., Barcelona, Spain
[4]Ansaldo Nucleare, Genova, Italy
[5]University of Granada, Spain
[6]ENEA Frascati Research Center, Rome, Italy

*Abstract*

IFMIF-DONES (International Fusion Materials Irradiation Facility DEMO-Oriented NEutron Source) Facility Project is a novel research infrastructure for testing and qualifying structural materials for future fusion reactors under relevant irradiation conditions. The project is implemented through in-kind contributions within the DONES Programme. IFMIF-DONES Control Systems are organized in two levels: Central and Local Instrumentation and Control Systems (CICS and LICS), connected through a complex network infrastructure. The CICS comprises three systems: Control, Data Access and Communication which is responsible for normal operation via central supervision and control, timing and synchronization, data management, alarm handling, system administration, and software management; Machine Protection System that ensures protection of investment against failures or operational errors; and Safety Control System, that implements the safety functions to protect personnel and the environment, including radiological protection functions. Several of the key technologies have been validated through the Linear IFMIF Prototype Accelerator (LIPAc), enabling risk mitigation and design consolidation. This contribution provides an overview of the IFMIF-DONES Control Systems architecture, focusing on strategic decisions for availability, maintainability, and long-term operability, as well as coordination frameworks addressing complex challenges throughout the project lifecycle.

## INTRODUCTION

### *Context and Background*

The IFMIF-DONES (International Fusion Materials Irradiation Facility DEMO-Oriented NEutron Source) was conceived within the Fusion Programme to provide a fusion-like neutron source for qualifying materials, with the primary objective of generating irradiation data essential for the design, licensing, construction, and safe operation of the DEMOnstration fusion power plant (DEMO).

IFMIF-DONES will accelerate a continuous 125 mA deuteron beam up to 40 MeV, delivering 5 MW per beamline (two planned in a future upgrade). The beam, shaped to a cross section of 10 or 20 cm width and 5 cm height, is directed onto a 25 mm-thick liquid lithium target flowing at 15 m/s, producing fusion-like neutrons that irradiate material samples in the Test Module. Lithium Systems include the largest-capacity liquid lithium loop ever built, with auxiliary oil and water loops, and impurity control system. Plant infrastructure, auxiliary systems, and buildings support the entire operation, all coordinated through the Central Instrumentation and Control Systems (CICS).

The project is based on an In-Kind (IK) contribution strategy. The parties are represented in the Steering Commitee, which is the board to take top-level agreements, while the consortium "IFMIF-DONES España" is serving as Technical Authority, facility owner, operator, and responsible for Spanish IK contributions mainly concentrated on parts of accelerator, buildings, plant systems, and control systems. This approach allows for broad international involvement but requires rigorous coordination, standardization, and integration frameworks. Since 2015, the project has progressed throughout preparatory and consolidation phases within the EUROfusion WPENS framework. Since 2023 we have started the construction phase, with several systems under final design and manufacturing. Integrated commissioning is scheduled to start in 2029, with first experimental operations expected in 2034, and full operation planned during 20 years.

### *Prototype Accelerator*

A key technical reference for IFMIF-DONES is the Linear IFMIF Prototype Accelerator (LIPAc), developed under the International Fusion Materials Irradiation Facility / Engineering Validation and Engineering Design Activities (IFMIF/EVEDA) project and located in Rokkasho, Japan. LIPAc, currently under commissioning, is designed to generate a 125 mA, 9 MeV continuous-wave deuteron beam with the objective of validating the IFMIF accelerator design [1-3]. It shares similar accelerator technology with IFMIF-DONES and has been instrumental in validating critical systems such as the Medium-Voltage Radio-Frequency (RF) systems, the Experimental Physics and Industrial Control System (EPICS)-based CODAC architecture, and the Machine Protection System (MPS) logic, reducing technical risk and consolidating the system architecture.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

† celia.carvajal@ifmif-dones.es

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBG002

# LEVERAGING IT-INSPIRED WORKFLOWS FOR PLC SOFTWARE

J. M. de Paco*, W. Niinimaeki, M. Marin, O. Minaiev, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

PLCs play a crucial role in operating, controlling, and interlocking high-power distributed systems at CERN, including main magnet power supplies and static VAR compensators. The increasing number, complexity, scale, and specialization of these critical applications make development and maintenance particularly challenging for a small team. To address this, we have introduced IT-inspired engineering workflows and technologies that wrap around the UNICOS framework, balancing customization with standardization. This approach enhances engineering efficiency and software quality. This paper highlights the benefits of using continuous integration pipelines, automated functional tests, reviewed merge requests, and the necessary virtual computing infrastructure, demonstrating how these innovations streamline development and maintenance in high-power system automation.

## INTRODUCTION

At CERN, the Electrical Power Converters (EPC) group is responsible for power converters and Static VAR Compensators (SVCs) [1] across all accelerators and facilities. Its Converters Controls Software (CCS) section develops and maintains PLC-based applications using CERN's standard stack—Siemens S7-1500 PLCs, *UNICOS-CPC* [2], and Siemens WinCC OA SCADA [3]—for the control, interlocking, and operation of high-power systems.

Our PLCs play a crucial role in the control, operations and diagnostics of high power-systems that are critical for the availability and performance of the accelerators. They protect expensive equipment, orchestrate large distributed control systems, interface with other services of the accelerators (power converter regulation, electrical network, cooling plants, magnet protection, access systems, etc.) and provide human machine interface (HMIs) both for machine experts and for CERN's Control Center. The group has adopted in-house PLC-based control systems as the standard for integrating and controlling large distributed power converter systems, which were previously outsourced to industry. This includes the upgrades of the main power converters of the PSB [4], PS [5], and SPS [6], as well as the SVCs.

### *Challenges and limitations of traditional workflows*

With the growing number and complexity of new projects and systems under our responsibility, the team faces significant technical and resource-related challenges. While maintaining and supporting the systems currently in operation, we work in projects to be commissioned during the CERN's Long Shutdown 3 period (2026-2029). These are the challenges we are facing and the limitations we would have if we keep traditional workflows:

- Productivity: Manual builds, uploads, and testing in UNICOS-CPC and TIA Portal are slow, repetitive, and error-prone.
- Quality assurance: Updates to UNICOS-CPC framework and PLC logic modifications risk breaking existing functionality. Limited time during technical stops prevents full manual regression tests.
- Teamwork: Large projects require version control for shared development, code reviews, and traceability. However, traditional automation engineering environments lack the built-in version control and collaborative features that are standard in modern IT software development.
- Commissioning phase: Debugging during commissioning is costly and risky, as engineers must work under pressure within short technical stops to commission multiple machines simultaneously. Any issue at this stage can jeopardize accelerator restart schedules and significantly increase costs.
- Knowledge transfer and resourcing: Expertise is typically concentrated in a few senior engineers. Traditional methods and technologies discourage knowledge transfer and make it harder to retain young engineers with modern IT skills and expectations.

### *Strategic plan towards DevOps in industrial automation*

To address the listed challenges, the team has strategically adopted IT-inspired technologies and workflows, transitioning from traditional PLC development practices to a modern industrial DevOps framework. DevOps unites software development with IT operations, leveraging automation, collaboration and fast iterative cycles. The core of this methodology are version control and CI/CD. Applying these principles into the context of Operational Technology (OT) aims to solve long-standing challenges related to software quality, efficiency and machine commissioning.

This is a gradual process involving changes in tools, processes, and team culture. In this paper we are presenting this transition in four adoption levels, according to our experience in the previous years. In any case, the process should be adapted to the needs and context of each organization. The four adoption levels are: 1. Version control and code editing; 2. Continuous Integration and Continuous Delivery (CI/CD); 3. Automated testing: virtualization and simulation; 4. Cultural adoption and advanced techniques.

* josemanuel.depaco@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBG003

# ACCELERATING CONTROL SYSTEMS WITH GITOPS: A PATH TO AUTOMATION AND RELIABILITY

M. Gonzalez*, M. Acosta†, Fermi National Accelerator Laboratory, Batavia, Illinois, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

GitOps is a foundational approach for modernizing infrastructure by leveraging Git as the single source of truth for declarative configurations. The poster explores how GitOps transforms traditional control system infrastructure, services and applications by enabling fully automated, auditable, and version-controlled infrastructure management. Cloud-native and containerized environments are shifting the ecosystem not only in the IT industry but also within the computational science field, as is the case of CERN and Diamond Light Source among other Accelerator/Science facilities which are slowly shifting towards modern software and infrastructure paradigms. The ACORN project, which aims to modernize Fermilab's control system infrastructure and software is implementing proven best-practices and cutting-edge technology standards including GitOps, containerization, infrastructure as code and modern data pipelines for control system data acquisition and the inclusion of AI/ML in our accelerator complex.

## INTRODUCTION

Fermilab's control system, ACNET [1], has supported accelerator operations for over four decades and continues to be the backbone of daily accelerator activities. Today, it comprises more than 600 active applications and services, developed primarily in FORTRAN, C++, and Java. These applications and services run on a mix of bare-metal servers and virtual machines, and most are accessed remotely through X11 over SSH. Source code has been maintained for decades in internal Concurrent Version System (CVS) repositories, reflecting the long-term continuity and stability of the system.

These practices were well suited to the technological context and operational needs of their time, providing robustness and reliability for accelerator operations. However, evolving requirements in scalability, collaboration, and security, along with advances in software engineering and infrastructure, now motivate a transition toward more modern approaches as seen at CERN [2] and Diamond Light Source [3] that are shifting towards modern software and infrastructure paradigms.

Several key shifts are being implemented by the Accelerator Control Operations Research Network(ACORN) [4] project as part of the modernization of Fermilab's control system:

- Architectural evolution: The system is moving away from monolithic applications that are tightly bound to specific hardware, toward microservices-based designs. This shift enables modularity, independent scaling of components, and faster development cycles.
- Ownership and collaboration: Instead of relying on a single expert or owner for each application, development is transitioning to team-based ownership. This model fosters peer reviews, improves resiliency of knowledge, and distributes responsibilities across multiple contributors.
- Programming paradigms: Legacy applications written in FORTRAN, C/C++, and Java are gradually giving way to more modern, web-based languages and frameworks.
- Security practices: The traditional approach of relying primarily on perimeter firewalls for protection is being replaced by systematic, standards-based security models. By following CISA and NIST hardening guidelines, the system adopts auditable, consistent, and adaptive practices that extend security beyond the network boundary.
- Infrastructure modernization: Applications and services are shifting from bare-metal and virtual machine deployments to containerized workloads orchestrated on Kubernetes [2] clusters. This provides declarative deployments, elasticity, automated recovery, and greater operational efficiency.

These changes create an environment where GitOps [3] emerges as a natural enabler, providing the necessary framework for automation, traceability, and reliability. By managing both infrastructure and applications as code, GitOps aligns with these modernization efforts and accelerates the integration of control systems into scalable and secure scientific environments.

## BACKGROUND

The ACNET control system comprises two main categories of applications and services: Console (CLIB framework) and Java [1]. Console applications are primarily written in C++ and are tied to the CLIB library; a monolithic framework that supports the integration of code and new functionality to control system software. The CLIB code base is the foundation for control system central services, data acquisition, data manipulation, network messaging and user-facing applications for several control system user groups. Applications that use CLIB as base are primarily written in C++ and use a tightly coupled monolithic system of shared libraries that enable machine experts, physicists and engineers to produce software that operates the Fermilab accelerator complex [1]. As a result, modifications and deployments to this framework and the applications that rely on it require careful peer-reviews and involve many manual

* mariana@fnal.gov
† macosta@fnal.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBG005

# TEAM RED: EXPERIENCES OF AN AGILE TASK FORCE SUPPORTING MULTIPLE SOFTWARE PRODUCTS

A. Asko*, S. Deghaye, C. Roderick, CERN, Geneva, Switzerland

*Abstract*

Controlling CERN's accelerator complex requires a significant number of domain specific applications. Typically these are developed by small teams comprised of a staff member with essential domain knowledge, and 1 or 2 students or graduates staying for 1-to-3 years. Developments may span over months or even years, according to the complexity and requirements, which may evolve significantly over time to follow the accelerator programme.

After some years, applications need to be consolidated or re-written to follow the inevitable technology evolution and remain secure, hardware compatible, and based on technologies for which skills exist in the marketplace. This development approach has produced results, but is inefficient due to the isolated developments sometimes leading to black-box solutions that eventually need reverse-engineering. Furthermore, long-term staff members can become siloed within specific domains, increasing the risk of single points of failure.

This paper shares the results and experience gained over 18 months in applying a more agile, task-force team-based approach to address the aforementioned limitations and inefficiencies. Together with the application of micro-frontend architectures, the team has been able to develop and maintain multiple applications with enhanced distribution of both domain and technical knowledge, as well as predictable delivery times. The team's workflows will be described, together with challenges faced and lessons learned.

## INTRODUCTION

Over the years, numerous applications have been developed to meet the diverse requirements of CERN's accelerator complex. For more than a decade, web applications have played a central role in the accelerator controls at CERN. Due to the complexity of an environment like CERN, many of these applications (also known as products) are developed by small teams, typically consisting of a staff member and one or two interns, over extended periods that often span several years.

Given that web development is defined by rapid evolution and continuous technological advancements, with new solutions regularly emerging, many of these applications ultimately come to rely on outdated technologies shortly after being considered feature-complete.

In turn, this triggers significant migration efforts to transition to the latest web technology stacks, a process that could also span several years. Meanwhile, new requirements emerge from the accelerator complex, occasionally requiring the development of entirely new applications that utilize the latest available web frameworks.

This eventually led to an environment where different teams are responsible for different applications, each built on distinct technology stacks according to the preferences and knowledge of the teams concerned. While the web served as a common platform, the diversity of stacks and limited collaboration across teams ultimately resulted in siloed groups focused on different domains with varying resources.

To address the aforementioned technical limitations, we have adopted a new migration approach over the past three years: Micro-Frontends [1]. This strategy has allowed us to gradually replace outdated, monolithic applications at a much faster pace. While this has also been effective in enhancing cross-team collaboration, this technological strategy does not address challenges related to domain separation, resource allocation, and delivery timelines.

It was recognized that the current workflow needed improvement, particularly the practice of fixed assignments of one team to one application product. This approach often led to one-person teams, as resources were frequently redirected to higher priority initiatives, leaving products understaffed and progress slow. To address this challenge, an agile task force approach has been adopted. In practice, a CERN organizational unit (known as a Section) within Controls, has been structured into several teams. The remainder of this paper describes the new agile task force approach from the perspective of "Team Red".

## OBJECTIVES

With an inventory of more than 30 web applications, varying in complexity and built on different technology stacks, it was required to find a more optimal way to continue evolving them while moving towards a more maintainable future. Before proceeding with changes in the approach and potentially in workflows, it was important to define the objectives that would guide all future decisions.

### *Shorten Delivery Times*

Distributing resources across different parallel product developments can lead to long delivery times, depending on each product's complexity, required competencies, and other factors. One of the key objectives was to determine how delivery times could be shortened as much as possible across the various products.

### *Increase Knowledge Transfer*

Improving knowledge transfer between Product Owners and eliminating single points of failure is highly desirable.

* anti.asko@cern.ch, stephane.deghaye@cern.ch, chris.roderick@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBG006

# HUNTING FOR HIDDEN BUGS: DEALING WITH TEST FLAKINESS IN SKA CONTROL SOFTWARE

G. Marotta*, C. Baffa, E. Giani, INAF-OAA, Firenze, Italy
S. Di Frischia, INAF-OAAb, Teramo, Italy
I. Novak, M. Colciago, Cosylab Switzerland, Brugg, Switzerland
G. Brajnik[1], E. Lena, Interactive Design Solutions Srl, Udine, Italy
[1]also at University of Udine, Udine, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Test flakiness—when a test intermittently passes or fails without changes to the code—poses a significant challenge in the validation of distributed control systems. This paper presents an investigation into test flakiness in CSP.LMC (Local Monitoring and Control for the Central Signal Processor), a key subsystem of the SKA (Square Kilometre Array) telescope. CSP.LMC is a Python application built on the TANGO framework, that is tested using a multi-level testing approach combining unit, component, and integration tests. To achieve scalable and reproducible deployment, the entire SKA control software runs within a Kubernetes environment. We systematically collect test outcomes and execution benchmarks to monitor system stability over time. A data mining approach is applied to uncover correlations and hidden patterns associated with test instability. Our analysis aims to uncover subtle software issues that are not easily detected through standard test evaluation. Furthermore, we aim to explore how the complexity of both the software architecture and its deployment may introduce sources of non-determinism that can lead to flaky tests. We discuss the impact of flakiness on the reliability of SKA control software and propose practical strategies to benchmark, detect, and mitigate flaky tests in complex distributed environments.

## INTRODUCTION

The reliability of control systems is a fundamental requirement in large-scale scientific projects, where even minor software faults can have severe consequences for operations and data integrity. This is particularly true for the Square Kilometre Array Observatory (SKAO), an international effort to build the world's largest radio telescope [1]. The SKAO control system is responsible for orchestrating and monitoring a vast and heterogeneous set of distributed components. Ensuring the robustness of this system is essential to guaranteeing uninterrupted scientific operations and safeguarding the quality of astronomical data.

The complexity of the SKAO control system makes reliability an especially difficult goal. The telescope encompasses thousands of devices and services, each with intricate interactions and dependencies. This scale demands advanced strategies for testing, monitoring, and maintaining software reliability. At the same time, the system's architecture introduces challenges such as nondeterministic behaviour arising from the interaction of heterogeneous hardware and software components.

Within this context, test flakiness has emerged as a significant obstacle. A flaky test produces inconsistent results—passing or failing without corresponding changes to the system under test. Flakiness reduces trust in the testing process, obscures genuine defects, disrupts continuous integration and deployment (CI/CD) workflows, and ultimately leads to frustration among developers and testers. In projects such as SKAO, mitigating flaky tests is of paramount importance. Furthermore, the design and execution of the test suite itself may amplify system complexity, creating additional opportunities for flaky behaviour through its interactions with the System Under Test (SUT).

This paper presents an investigation of test flakiness in the Central Signal Processor Local Monitoring and Control (CSP.LMC) subsystem, a critical element of the SKAO control software. Our methodology combines the systematic collection of test outcomes from CI/CD pipelines, database-driven storage, and the analysis of trends through dashboards and exploratory sessions. With this approach, we aim to identify the causes and impacts of flaky tests in distributed environments and propose strategies for their detection and mitigation in large-scale scientific software systems.

## SKAO CONTROL SYSTEM INFRASTRUCTURE

The SKAO control system (CS) is built around a modular, distributed software architecture designed to manage and coordinate a large number of heterogeneous devices. The CS software components are implemented in Python and based on the TANGO Controls framework [2], an open software middleware widely adopted in large-scale scientific facilities. Within this framework, each device is represented as a TANGO Device, encapsulating the logic for controlling hardware or providing higher-level coordination functionality. This model provides a uniform interface for monitoring and commanding system elements, ensuring interoperability across the diverse landscape of SKAO subsystems.

To support scalability, reproducibility, and resilience, all TANGO devices in the SKAO control system are containerized and deployed as pods within a Kubernetes cluster [3]. This design enables uniform deployment practices and facilitates automated orchestration, scaling, and recovery of software components. Kubernetes, originally developed for cloud-native applications, also offers capabilities that align

* gianluca.marotta@inaf.it

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBR001

# SQUEEZING EVERY LAST DROP: LIFE EXTENSION OF THE LHC AND PS ACCESS SAFETY SYSTEMS

T. Hakulinen, D. Chapuis, P. Gaviglio, G. Godineau, F. Havart, T. Ladzinski,
T. Lang, P. Ninin, R. Nunes, T. Riesco, V. Rios
CERN, Geneva, Switzerland

*Abstract*

The LHC Access Safety System (LASS) was commissioned in 2008 for the initial launch of the LHC, while the current PS Access Safety System (PASS) entered in operation in 2015 after a full multi-year revision. By the end of the upcoming LHC Long Shutdown 3 in 2030, LASS will have accumulated over 20 years of service and PASS 15 years. Both systems are based on Siemens series 400 safety PLCs, a product line that is being progressively phased out by the vendor. After evaluating several consolidation strategies to ensure the smoothest possible continuity of operation of both LHC and PS, we have opted for a life extension of both current access safety systems. This strategy retains the existing systems but introduces a number of architectural enhancements to mitigate current limitations that could impact accelerator availability. The objective is to secure reliable long-term operation of both safety systems, while also improving operational autonomy, maintenance capabilities, and resilience to vendor product life cycles. This paper outlines the rationale behind the selected approach, the external factors influencing our decision, and the key steps planned to implement this consolidation effectively.

## INTRODUCTION

The Large Hadron Collider (LHC) [1] is the world's largest and CERN's current flagship particle accelerator with a circumference of 27 km, 8 major underground sites, and with 4 large and 5 small physics experiments (Fig. 1).

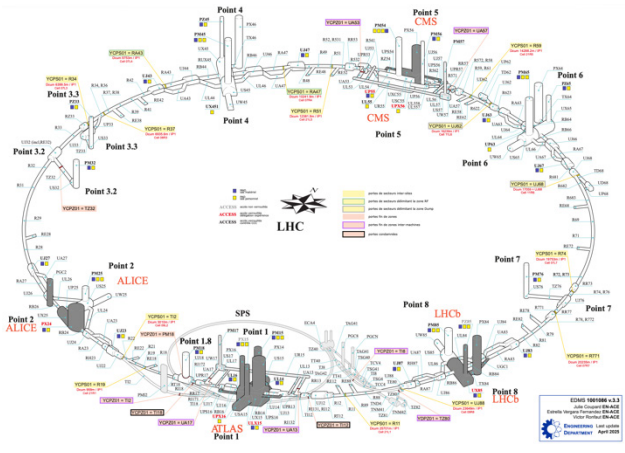



Figure 1: Schematic of the LHC.

The PS Complex comprises Proton Synchrotron (PS – CERN's first circular accelerator, in operation since 1959), PS Booster, LINAC 3, LINAC 4, LEIR, Antiproton Decelerator (AD), and the various experiments and transfer lines associated with these accelerators (Fig. 2).

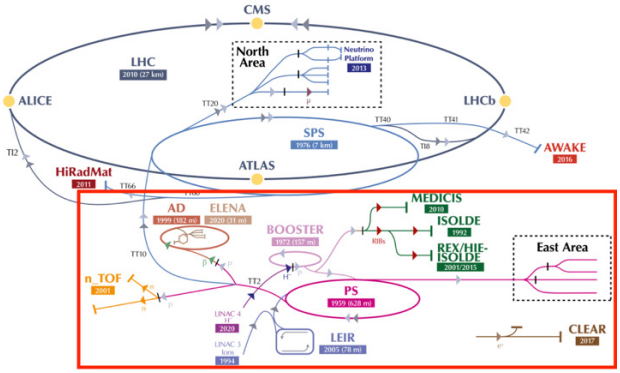



Figure 2: Schematic of CERN accelerators and experiments. The PS Complex within the red borders.

The LHC and PS Personnel Protection Systems (PPS) consist respectively of the *LHC and PS Access Control Systems* (LACS and PACS), as well as the *LHC and PS Access Safety Systems* (LASS and PASS). The Access Control Systems manage user access to the LHC and PS beam, service, and experimental areas by enforcing strict conditions on user identification, authentication, authorization, and people counting. Access to the controlled areas for personnel is via specific Personnel Access Devices (PAD) and for material via Material Access Devices (MAD). As of this writing, the LHC PPS contains 10 sites, 34 access points, 43 PADs, 29 MADs, over 350 access doors of different kinds, and over 10000 wired signals. The PS PPS contains 14 zones, 21 access points with the same number of PADs, 13 MADs, 236 access doors, and 2242 wired signals.

Accelerator Access Safety Systems are critical to the successful running of the accelerators, as they enforce the safety of all personnel entering the accelerators by ensuring that if the beam is on, there can be no people present, and if people are present, there can be no beam. These systems run completely isolated from other systems and networks, and only operate within the strict parameters of monitoring the absence of beam or people depending on the operation mode.

The LASS and PASS manage the safety conditions of all zones by monitoring various "Elements Important for Safety" (EIS), which include access doors (EIS-access), beam safety elements (EIS-beam), and "machine elements" such as various radio frequency (RF) equipment (EIS-machine). The LASS and PASS will allow passing the facility into access or beam modes according to the detected conditions and interlocks the accelerator, if these conditions are not met.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBR002

# MANAGING TECHNICAL DEBT ACROSS LARGE-SCALE CONTROL SYSTEMS *

Adam Watts†, Beau Harrison‡,Finley Novak§, Fermilab, Batavia, IL, USA

*Abstract*

This paper provides a comprehensive overview of technical debt in the context of control systems for large-scale physics facilities. We explore various forms, common causes, and potential consequences on system reliability, maintainability, and extensibility. Drawing on experiences from multiple projects, including the ACORN control system modernization at Fermilab, we present a range of strategies for proactively managing technical debt, including best practices in design, development, testing, and documentation, as well as reactive approaches for identifying and mitigating existing issues. We emphasize the importance of a collaborative approach and the use of modern tools, such as fault tracking systems, to ensure the long-term health and success of control systems in our field.

## INTRODUCTION

Particle accelerator laboratories operate complex control systems that must function reliably for decades while adapting to evolving scientific requirements. These systems often inherit legacy components, emergency fixes implemented under time pressure, and undocumented solutions developed by personnel no longer available to explain their design rationale. The cumulative effect of these compromises creates what software engineering terms "technical debt" – a metaphor describing the implied cost of additional rework caused by choosing expedient solutions over comprehensive approaches [1].

Technical debt, analogous to financial debt, represents a conscious trade-off between immediate delivery and long-term maintainability. Just as financial debt can be strategically beneficial when managed properly, technical debt becomes problematic only when it accumulates without systematic repayment or when the interest payments (ongoing maintenance costs) become unsustainable.

Accelerator control systems often fall out of sync with current industry standards and technology due to the decades-long operational life-cycles of accelerator facilities coupled with high experimental demand for constant up-time. This inherited debt includes dependencies on obsolete hardware and programming paradigms that are no longer supported or understood by the broader engineering community. Paying down this type of technical debt requires extensive, focused, and well-funded efforts.

This paper examines technical debt specifically within the context of large-scale accelerator control systems, where the consequences of system failures extend beyond inconvenience to potential safety hazards, experimental data loss, and facility downtime cost measured in millions of dollars. Drawing on experiences from the ACORN [2] control system modernization project at Fermilab and other facility experiences, we present both proactive strategies for preventing excessive debt accumulation and reactive approaches for identifying and mitigating existing technical debt in operational systems.

## TECHNICAL DEBT IN CONTROL SYSTEMS CONTEXT

### *Unique Characteristics of Control Systems*

While similar to technical debt as traditionally defined in software development, technical debt in particle accelerator control systems can differ in ways that necessitate specific management methods. These control systems operate in environments where safety considerations are critical, yet response time and reliability requirements often necessitate stringent and unwavering performance. The operational lifetime of major accelerator components can span decades, far exceeding typical software product life-cycles.

Integration complexity represents another distinguishing factor. Modern accelerator facilities incorporate thousands of devices spanning multiple technologies, communication protocols, and vendor platforms. Control systems must interface with legacy equipment that cannot be easily replaced while simultaneously supporting cutting-edge instrumentation with rapidly evolving interfaces.

The staffing model for accelerator facilities also influences technical debt accumulation. Control systems are maintained by mixed teams of engineers, technicians, operators, and even scientists, each with different backgrounds, priorities, and constraints. Solutions that appear elegant to software engineers may prove incomprehensible to technicians responsible for troubleshooting at 3 AM.

### *Types and Manifestations of Technical Debt*

Technical debt in accelerator control systems manifests in various forms, each with distinct characteristics and implications for facility operations.

**Design Debt** occurs when expedient solutions are implemented without consideration of broader system architecture. A prominent example at Fermilab is the hard-coded assumption of 15 Hz operation embedded throughout legacy control systems. As future accelerators and experiments such as PIP-II and LBNF/DUNE require 20 Hz operation rates,

* This manuscript has been authored by Fermi Forward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.

† awatts@fnal.gov
‡ beau@fnal.gov
§ fnovak@fnal.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBR003

# MANAGING OBSOLESCENCE OF ALIGNMENT CONTROL SYSTEM ON THE LASER MEGAJOULE FACILITY

J. Hannier, J-P. Airiau, R. Doazan, A. Gavello
CEA CESTA, Le Barp, France

*Abstract*

The Laser MegaJoule (LMJ), a 176 beam laser facility developed by CEA, is located near Bordeaux. It is part of the French Simulation Program, which combines improvement of theoretical models used in various domains of physics and high performance numerical simulation. It is designed to deliver about 1.4 MJ of energy on targets, for high energy density physics experiments, including fusion experiments [1].

The LMJ technological choices were validated on the LIL, a scale-1 prototype composed of 1 bundle of 4 beams. The first bundle of 8 beams was commissioned in October 2014 with the implementation of the first experiment on the LMJ facility. The operational capabilities are increasing gradually every year until the full completion by 2026. By the end of 2025, 22 bundles of 8 beams will be assembled (full scale) and 19 bundles are expected to be fully operational.

The alignment control system is an essential part of the LMJ facility that allows the alignment of all the beams through the bundles to the target. It is built to control several hundred of devices, mainly drives, but also cameras and lasers.

The architecture of the alignment control system is based on physical and software PLCs that were chosen 15 years ago. These PLCs are now obsolete and we have to replace them. This article discusses our strategy to manage the obsolescence of alignment control system by minimizing the impact on the operation of the LMJ.

**Key words:** Laser facility, LMJ, IT, Obsolescence, PLC, Control System

## PRESENTATION

### *Hardware Architecture*

The hardware architecture of the alignment control system consists of computers and PLCs. It can be divided into three subsystems (Fig. 1).

The supervision subsystem includes the computers necessary for the operation of the control system's SCADA Panorama software: a storage server, servers for acquiring and processing data from the PLCs, and HMI stations.

The common subsystem includes the PLCs and computers used to perform operations that are common to all laser bundles. It consists of an OPC server, a Siemens S7-400 PLC, and an image processing PC (IP PC), all of which are duplicated to provide redundancy in case of failure. This subsystem controls the lasers and shutters via remote Siemens ET200 input/output modules, as well as the common reference camera in the center of the chamber via a WCF server.

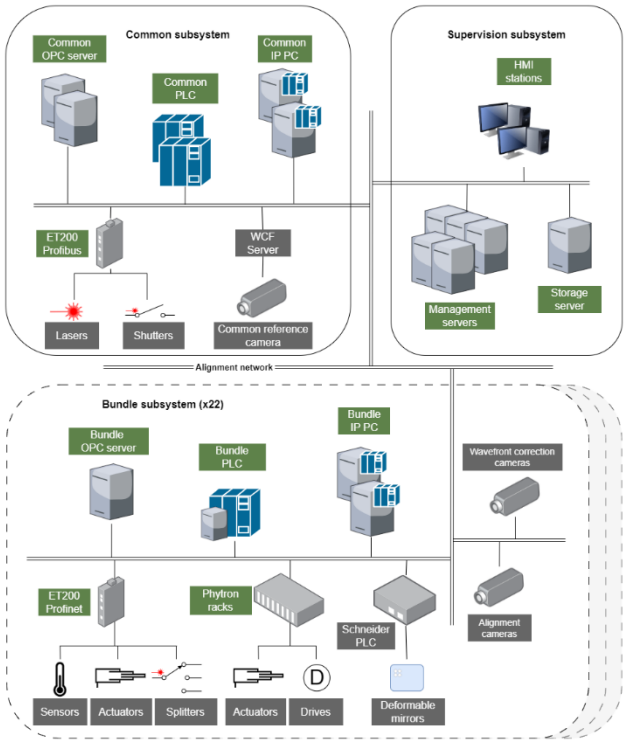



Figure 1: Hardware architecture.

The bundle subsystem includes the PLCs and computers used to perform the operations on a laser bundle. There are 22 of these subsystems (one per laser bundle). It consists of an OPC server, a Siemens S7-mEC PLC, and two image processing PCs (IP PCs) that operate in parallel with half of the cameras when both are operational, or in redundancy with all cameras when one of them fails. The Siemens S7-mEC PLC is a PC packaged as a PLC that runs with a software PLC on a real-time system (Siemens WinAC RTX). This PLC has an operating system whose features can be used via a Siemens SDK. The bundle subsystem controls sensors and actuators via remote Siemens ET200 input/output modules, drives using Phytron racks, and deformable mirrors at the bottom of the cavity using a Schneider PLC.

Common or bundle image processing PCs also run with a software PLC on a real-time system (Siemens WinAC RTX). This software PLC acts as an interface between the common or bundle PLCs and the PC operating system via a Siemens SDK.

Numerous communication mechanisms are used in alignment control system. OPC Classic is used between the supervision and the PLCs, and Siemens's S7 Communication protocol is used between the PLCs. The Profibus and Profinet protocols are used to control most of the equipment (ET200, Phytron). The Modbus TCP protocol is used

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBR004

# OPERATIONAL EXPERIENCE AND OPTIMIZATION TECHNIQUES FOR LARGE EPICS CONTROL SYSTEM ENVIRONMENTS *

K. Vodopivec†, D. Cashion‡, Oak Ridge National Laboratory, Oak Ridge, TN, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Spallation Neutron Source (SNS) utilizes the Experimental Physics and Industrial Control System (EPICS) control system to manage the main accelerator, multiple test stands, and most neutron scattering instruments. The control system environment for the main accelerator alone includes approximately 300 Input/Output Controllers (IOCs), 200 soft IOCs, and 100 Operator Interface (OPI) workstations, all communicating through over 1,000,000 process variables. The EPICS distributed control system relies heavily on UDP broadcast communication to locate process variables on the network, which can significantly impact the operation of connected devices and, in turn, affect accelerator uptime. In this paper, we present operational challenges and discuss optimization techniques applied to the EPICS environment at SNS.

## INTRODUCTION

EPICS [1] provides a distributed control system architecture that links sensors and actuators with client tools and high-level services over Ethernet. At facilities such as the Spallation Neutron Source (SNS), this scales to thousands of devices and more than one million Process Variables (PVs). In such a distributed system, nodes must locate PVs that may reside on other nodes. The traditional mechanism is a UDP broadcast search, followed by establishing a TCP connection once the PV is found. UDP broadcast packets are sent to every host on the subnet, which is efficient for discovery and scales surprisingly well. However, problems emerge when clients continually search for PVs that do not exist. At facilities operating for decades, this situation is unavoidable. Examples of how non-existing PVs arise:

- Temporary IOCs for experiments. For example, a beam-diagnostics experiment may deploy a temporary IOC with new PVs and operator interfaces. When the experiment ends, the IOC is removed but the operator interface remains in a shared location. When opened, it generates repeated searches for now-missing PVs.
- Upgrades and obsolescence. Hardware upgrades replace aging systems with new equipment that provides similar functionality but different PVs. Old PVs remain embedded in client tools such as data archivers, alarm handlers, GUIs, diagnostics, high-level IOCs, or AI models.

As a fully distributed system, EPICS does not rely on any central facility to track PVs. As a corollary, EPICS does not provide a native way to trace stale PV references, allowing them to silently accumulate. Over time, the buildup of non-existing PVs leads to excessive network traffic, wasted resources, and in some cases degraded control system performance. Repeated broadcast searches for missing PVs generate unnecessary network load, and directly ties to the scalability and performance concerns.

## POTENTIAL PROBLEMS

At SNS, occasional EPICS broadcast floods have hindered operations by rendering parts of the control system invisible to operators. In many cases, individual IOCs locked up and rebooted, leading to unplanned shutdowns—most often VME IOCs running vxWorks 6.x. In more severe incidents, intense broadcast traffic caused (mostly VME) IOCs to stop responding to EPICS communication while still executing local control tasks. In some cases, the thread responsible for Channel Access communication crashed, requiring a manual reboot, while in other instances the IOC rebooted automatically. Because this created uncertainty in machine safety, operators chose to manually shut down the accelerator to prevent uncontrolled operation. Other devices with limited resources, such as PLCs and terminal servers, also experienced failures such as dropped communication or locked-up network stacks, resuming normal operation only after the traffic subsided.

Each of these events resulted in several hours of facility downtime and triggered post-mortem investigations. Cisco network equipment provides basic monitoring of broadcast traffic, but it cannot resolve EPICS protocol details such as which PVs are being searched for, or where the traffic originated from. The open-source tool *caSnooper* [2] can decode Channel Access search requests and is valuable for troubleshooting, but it does not support real-time reporting or system-wide analysis. To address this, a graphical analysis tool was developed as shown in Fig. 1. It subscribes to EPICS broadcast traffic like any IOC, parses each packet, and extracts key fields (origin host, PV name, protocol version). All requests are stored in an internal database for retrospective analysis. The tool provides two modes of operation:

- Continuous mode: plots a live histogram of search requests per second.
- Introspection mode: enables zooming into a selected time window and filtering by hostname or PV name.

* This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paidup, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

† vodopiveck@ornl.gov

‡ cashionmd@ornl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBR005

# LEAVING AGILE ISLAND: USING FLIGHT LEVELS AT THE ALBA SYNCHROTRON

O. Matilla*, J. Avila, J.L. Fernandez, T. Fernández, B. Lorenzo, Z. Reszela, N. Soler, S. Vicente
ALBA Synchrotron, Barcelona, Spain

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Integrating agile methodologies within large, non-agile organizations often results in isolated agile islands and inefficient Water-Agile-Fall models. This work outlines the practical approach adopted by ALBA Synchrotron's Computing Division to embed agile practices effectively within a traditional, hierarchical structure. Instead of complex, large-scale frameworks (SAFe, SoS or LeSS), we employed a flow-based agile integration strategy based on Klaus Leopold's Flight Levels.

Teams utilize agile development at the operational level (Flight Level 1). The critical integration occurs at Flight Level 2 (Coordination), focusing on optimizing end-to-end value streams across teams. Key FL2 mechanisms at this level include visualizing the entire workflow, enforcing strict Work-in-Progress (WIP) limits to enhance predictability and throughput, fostering targeted communication, and ensuring reliable commitments. In practice, we have developed a custom app drawing data using the Jira API (Atlassian ecosystem) as a key stone, which ensures the crucial aspects of workflow visualization and enhanced communication. To bridge the gap with non-agile stakeholders, we track progress via time investment versus estimates and utilize milestones (e.g., MVPs, system validation) as governance and delivery checkpoints.

Additionally, we will discuss how these practices illustrate reliability and transparency, cultivating organizational trust between agile teams and the wider hierarchical structure.

## OVERVIEW OF A DEMANDING ORGANIZATIONAL CONTEXT

The Computing Division at the ALBA Synchrotron is divided into five sections and comprises nearly 70 staff members and 10 students. The Division is responsible for delivering software and hardware solutions, maintaining operational excellence in Controls Systems, Data Acquisition, and Data Management, operating the central IT infrastructure, and providing other corporate services.

The Division serves numerous scientific groups and other key clients; therefore most team members work with multiple clients simultaneously.

In scientific research infrastructures, project specifications are often open-ended and evolve during execution; yet clients require precise delivery times and commitments. This pressure can sometimes lead to overlooking the inherent complexity of technical developments and undervaluing the need for specialized expertise.

* omatilla@cells.es

At the same time, teams must consistently ensure the operational excellence of existing systems. Leveraging this dual role creates workload imbalances, limiting the available resources for operational maintenance and technological readiness. As deadlines approach, client pressure naturally intensifies. Prioritization then becomes highly complex, since direct client-driven activities compete with essential maintenance and readiness tasks, where the cost of delay should ideally guide decision-making.

Most sections within the Computing Division use agile methodologies, which provide flexibility to evolving needs. However, aligning these methods with the institute's traditional waterfall project management culture remains a challenge. At times, agile practices have been misinterpreted as a lack of ability to define schedules or estimate resources. To address this, the Division has implemented a fully transparent policy covering all activities in which it is engaged.

Nevertheless, given the large number of ongoing complex activities, this transparency has occasionally resulted in information overload for our clients, making it difficult for them to extract key information from the technical intricacies. To mitigate this and preserve the Division's reputation, multiple levels of meetings have been established to ensure alignment across the organization.

This dynamic strains the trust between the Division and its clients. Unpredictable timelines can be read as weak performance, while reprioritization can undermine the integrity of our commitments.

## ON THE IMPORTANCE OF TRUST

Multiple studies have tried to identify the key factors that lead organizations to achieve the highest level of productivity. Historically, factors such as clearly defined roles, detailed workflows, risk management frameworks or business alignment assurance (KPIs, SLAs) have been considered essential. Yet, one distinct aspect stands out above all the previous ones: **Trust**.

The concept of trust is familiar in everybody's daily life: trusted people are expected to fulfil their commitments, act with good intentions, and deliver results. Yet, there is a fundamental psychological component: the close relationship between trust and vulnerability. Psychologically, trust is the willingness to be vulnerable, relying on others to act as expected without needing to be monitored or controlled.

Trust enables progress in situations where people and groups cannot fully control the outcomes. It serves as the glue that holds an organization together and the lubricant that allows processes to run smoothly. In summary, high levels of trust foster agility.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUBR006

# MAINTAINING CONTROL SYSTEM VIABILITY AMIDST DECLINING BUDGETS, RESOURCES, AND SANITY (HOW I LEARNED TO STOP WORRYING AND LOVE SAYING NO)

G. Croke
Jefferson Lab, Newport News, VA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The controls systems group at Jefferson Lab has gradually diminished in size over 20 years as the scope of their responsibilities has continued to increase. Following a recent RIF, the group contains only 7 software developers responsible for maintaining hundreds of applications and EPICS IOCs, as well as tackling numerous pressing challenges. The size of the group is currently inadequate for the task. I examine how this situation arose, potential effects, and plans to remedy.

## JEFFERSON LAB

Jefferson Lab — formally the Thomas Jefferson National Accelerator Facility — is a U.S. Department of Energy national laboratory in Newport News, Virginia. It's dedicated primarily to nuclear physics research, especially the study of how quarks and gluons (the fundamental building blocks of matter) form protons, neutrons, and nuclei.

### *Accelerator Facilities*

Jefferson Lab's main facility is the Continuous Electron Beam Accelerator Facility (CEBAF), a superconducting radiofrequency (SRF) electron accelerator [1]. CEBAF was commissioned at a 4 GeV machine in the 1990s. Design energy upgrades to 6 GeV and then 12 GeV followed [2].

Jefferson Lab operates several secondary accelerator facilities, including the Upgraded Injector Test Facility (UITF), Low-Energy Recirculating Facility (LERF), Gun Test Stand, Cryomodule Test Facility (CMTF), and Vertical Test Area (VTA).

### *Control Systems at Jefferson Lab*

CEBAF initially operated on the CEBAF Control System (CSS), a fully in-house designed system. JLab migrated to using EPICS in the mid-1990s [3]. For the next decade, Jefferson Lab's control systems teams maintained a significant presence in the EPICS collaboration and made many contributions to the community. Today, all accelerator facilities at Jefferson Lab have EPICS control systems.

## ACCELERATOR COMPUTING GROUP

The Accelerator Computing Group (ACG) at Jefferson Lab comprises the Accelerator Computing Software (ACS) team and the Accelerator Computing Environment (ACE) team.

### *ACS Responsibilities*

**IOCs** ACS staff are responsible for maintaining the operation of hundreds of EPICS IOCs. These are primarily VME, PC-104, and softIocs.

Table 1: EPICS IOCs at Jefferson Lab

| IOC Type | OS | Number |
|---|---|---|
| VME | vxWorks | 248 |
| PC-104 | RTEMS | 407 |
| SoftIOC | RHEL | 275 |

**Software** ACS maintains an extensive suite of software applications developed over decades of operations.

Table 2: Software Maintained by ACS

| Application Type | Number |
|---|---|
| Low-level (EPICS layer) | 1,211 |
| High-level (Linux) | 405 |
| Web applications | 50+ |
| C++ libraries | 25+ |
| Scripts | 100s |

**ACE Systems** The Accelerator Computing Environment team manage an environment and network comprising many hundreds of hosts:

- 400+ RHEL 9/7/6 (~230 Bare Metal and 170 VMs)
- 450+ Windows 10 / 2008R2 / 2012R2
- ~200 Virtual Machines (Windows, Linux)
- 100+ Console Servers (TSLinux)
- 180+ APC UPSs & reboot switches
- Various Ethernet to DAQ devices
- 8 VMware ESXi HyperVisor Hosts (~100 VMs)
- 5 Routers (Cisco)
- 160+ Switches (Cisco)
- 60 Defined Subnets/VLANS

The EPICS control systems at Jefferson Lab contain approximately 1,000,000 PVs, about half of which are archived. The EDM display manager has been used for operator panels at Jefferson Lab for two decades. ACS maintains a library of about 15,000 static EDM screens.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCG001

# EPICS AREADETECTOR OVERVIEW AND UPDATE

M. L. Rivers, University of Chicago, Chicago, USA

*Abstract*

areaDetector is an EPICS framework to support 1-D and 2-D detectors. It is modular C++ code that greatly simplifies the task of writing drivers for new detectors. It supports plugins, which receive detector data from the driver and process it in some way. For example, plugins perform statistics calculations, image processing, Region-Of-Interest extraction, file saving, fast Fourier transforms, and export to EPICS Channel Access or pvAccess for image display in clients like ImageJ. Plugins can each run in their own threads, permitting parallel processing on multi-core machines. Drivers have been written for many of the detectors commonly used at synchrotron and neutron sources, including CMOS imaging detectors, X-ray pixel array detectors, X-ray and visible light spectroscopy detectors, and X-ray beam position monitors.

## INTRODUCTION

Most experiments at synchrotron, FEL, and neutron scattering facilities use 1-D or 2-D detectors. These include X-ray detectors such as direct-detect pixel array detectors and scintillation-based detectors with CMOS, CCD, and amorphous silicon sensors. It also includes visible light CCD and CMOS cameras for imaging and optical spectroscopy, energy dispersive Si and Ge detectors for X-ray spectroscopy, and various types of X-ray beam position monitors. These detectors need to be integrated with the control system used for other components of the experiment. Many beamlines around the world use the EPICS control system [1], which is a collaborative open-source control system toolkit. areaDetector is an EPICS framework for controlling detectors [2, 3]. An early (2010) version of areaDetector was described in [4], and a more mature (2017) version was described in [5]. This paper will very briefly summarize the framework architecture and then describe in greater detail the features that have been added since 2017.

The goals of the areaDetector module are to:

- Provide a standard interface defining the functions and parameters that a detector driver should support.
- Provide a set of base EPICS records that will be present for every detector using this module. This allows the use of generic EPICS clients for displaying images and controlling cameras and detectors.
- Minimize the amount of code that needs to be written to implement a new detector.
- Allow easy extensibility to take advantage of detector-specific features beyond the standard parameters.
- Have high-performance.
- Provide a mechanism for device-independent real-time data analysis such as regions-of-interest, statistics, image processing.
- Save files in community standard formats.
- Provide drivers for commonly used detectors at facilities running the EPICS control system.

## ARCHITECTURE

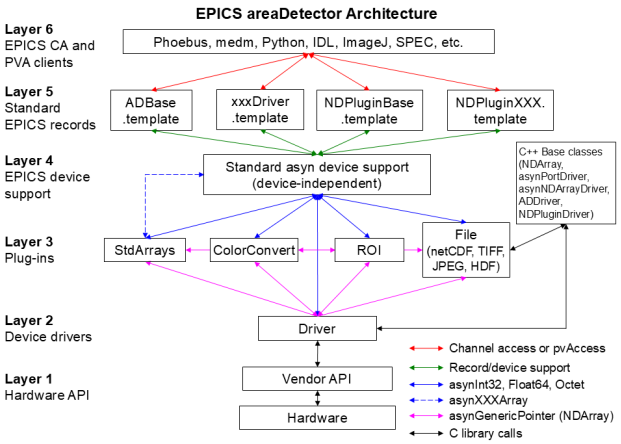



Figure1: areaDetector architecture.

The architecture of the areaDetector module is shown in Fig. 1. From the bottom to the top this architecture consists of the following six layers:

1. This is the layer that allows user written code to communicate with the hardware. It is usually provided by the detector vendor. It may consist of a library, a socket protocol definition, or other type of API.
2. This is the driver that is written for the areaDetector application to control a particular detector. It is written in C++ and inherits from the ADDriver class. It uses the standard asyn interfaces for control and status information. Each time it receives a new data array it passes it as an NDArray object to all Layer 3 clients that have registered for callbacks. This is the only code that needs to be written to implement a new detector.
3. Code running at this level is called a plugin. This code registers with a driver (or another plugin) for a callback whenever there is a new NDArray.
4. This is standard asyn device support that comes with the EPICS asyn module [6,7,8].
5. These are standard EPICS records and EPICS database (template) files that define records to communicate with drivers at Layer 2 and plugins at Layer 3.
6. These are EPICS channel access clients that communicate with the records at Layer 5. areaDetector includes ImageJ client applications that can display images using EPICS Channel Access or pvAccess. It also includes detector and plugin screens for the EPICS display managers CSS/Phoebus, CSS/BOY, MEDM, edm, and caQtDM.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCG002

# DATA ACQUISITION AND ON-THE-FLY PROCESSING FROM HIGH RATE DETECTORS AT MAX IV

A. Lilius, P. Bell, M. Cascella, Z. Matej, A. Salnikov, MAX IV Laboratory, Lund, Sweden

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

At MAX IV, we have developed a high-performance data acquisition (DAQ) system to handle the high rate detectors exploiting the brightness of the fourth generation source. This system integrates multiple detector types, including photon counting and charge integrating detectors as well as sCMOS cameras, into a unified DAQ framework. Data are streamed to a central Kubernetes cluster which mounts an IBM Storage Scale (GPFS) storage, with control provided via Tango. The system provides live feedback from the detectors/cameras and is furthermore extended to provide on-the-fly data reduction via the "Dranspose" framework, a horizontally scalable, distributed data analysis pipeline. We describe the components of our DAQ and processing framework, highlighting its performance for live data streaming and on-the-fly reduction.

## INTRODUCTION

The MAX IV synchrotron laboratory in Lund, Sweden, currently operates 16 beamlines for users, providing X-ray techniques in diffraction and scattering, imaging, and spectroscopy covering a wide scientific program. The 3 GeV ring [1] was the first fourth generation light source when it began operation in 2016, offering low emittance, high brilliance and a high degree of coherence. The ability to fully exploit these beam properties is dependent on the performance of the detectors at the beamline experimental end-stations. Although each technique has its own specific requirements, there is a general need for large-area pixelated detectors to run at high frame rates. Such devices generate large volumes of data at high rate which must be captured and stored without losses, in experiments where the detectors are operated continuously for extended periods. Moreover, it is not sufficient to simply record this data to disk for later analysis; to guide the experiments, on-the-fly analysis and visualisation must be provided. These requirements motivate a data acquisition (DAQ) and online analysis framework based on *data streaming* as opposed to file-based operations. In accordance with this, a unified streaming format has been implemented to ensure more consistent and reliable streaming across the facility network. The system standardisation is necessary to cover all detector types for ease of support and operation, but also enough flexibility for different applications.

## MAX IV DETECTOR-DAQ SYSTEM

An overview of the detector-DAQ scheme including a simplified view of the network connectivity is shown in Fig. 1. It can be divided into a control and data streamer layer, which is detector-specific and runs close to the detector at the beamline, and a data stream receiver and processing layer, which is detector-agnostic and takes place in a central Kubernetes-managed DAQ cluster. The DAQ cluster mounts the IBM spectrum scale GPFS storage where data are saved in hdf5 files, directly as raw detector images and optionally after some processing step(s). The two layers are described in detail in the subsequent sections, followed by a description of the operation of the combined system.

### *Data Streaming From the Beamline*

Hybrid photon counting from commercial suppliers are usually delivered with a custom server referred to as the Detector Control Unit (DCU). The communication protocol between the detector head and the DCU is typically a closed system that we do not have access to, using for example UDP over twisted pair or optical fibre. In case of CMOS cameras, a server fulfilling the role of the DCU is usually provided by MAX IV. This may hold a "frame grabber" card supplied with the camera, for example a CoaXpress PCI Express card or a Camera Link PCI Express card. Additionally, some CMOS cameras use a USB3 connection. The distance between the DCU located in the beamline electronics rack and the camera in the experimental hutch means that for these types of connections some method of range extension is often required, for example by conversion of the Camera Link to optical fibre and back.

For genuinely on-the-fly analyses and visualisation it is necessary to operate on a data stream rather than on files. This data stream utilises a dedicated "purple" network separate from the "blue" network of the control system (which at MAX IV is Tango [2]). The Eiger detectors from Dectris [3] are unique in natively providing a stream of data from the DCU over ZeroMQ, typically over a single 40 Gbit fibre network interface. For all other detectors and cameras, we have developed a *streamer* application that runs on the DCU, providing a data stream from a ZeroMQ push socket similar to that of the Eiger. The streaming format most commonly used and the one that has been chosen as the standard uses multipart JSON and binary encoding for the message headers and frames, respectively. On the DCUs the ZeroMQ stream is made available on a 40 Gbit fibre network interface (or 10 Gbit if sufficient for the data rate of the detector). Data are then streamed to the DAQ cluster over dedicated 100 Gbit capable Ethernet (purple line in Fig. 1).

The implementation of the streamer application is specific to the particular detector or camera. In some cases the streamer application can be developed and deployed on the DCU thanks to a high level API or SDK from the detector or camera supplier, which provides methods to get the image frames into memory. Implementation in Python is preferred, though sometimes with bindings to code supplied in another language, such as C via CFFI. The Dectris Pila-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCG003

# A MOTION CONTROL SYSTEM FOR FLY SCANS SYNCHRONIZED TO X-RAY PULSES*

N. S. Lentz†, Z. L. Lentz, U. Farooq, B. Fishler,
SLAC National Accelerator Laboratory, Menlo Park, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

With the High Energy Upgrade project in the installation phase at the SLAC National Accelerator Laboratory Linac Coherent Light Source (LCLS), soon the repetition rate of the X-ray beam will exceed the capability of the present motion control system to reliably deliver samples. We describe the design of a motion control system which aims to deliver solid crystal samples within 5 um of the focal point of the X-ray beam at a rate of 1000 Hz using a fly scan. Throughout the fly scan, a triggered timestamping device is used with an EtherCAT based distributed clock to predict the position error at the time of each X-ray pulse. This position error is used in a feedback loop to bias the velocity command to each motion axis and restore synchronism. The design is flexible enough to expand to any number of synchronized axes up to the limits of the computing hardware. Additionally, with most of the software written in Structured Text language defined in IEC 61131-3, it is transferable to other EtherCAT based real-time systems with few modifications. Finally, we describe how the TwinCAT PLC programming environment can be used to develop and test almost all of the functionality of the software without hardware present, through a combination of unit testing and simulation axes.

## DESIGN GOALS

This document describes the design and implementation of a synchronous scan feature for the LCLS Experiment Control Systems (ECS) Motion Control System (MCS). This feature enables high-speed deterministic motion synchronized to X-ray pulses from the X-ray Free-Electron Laser (XFEL).

### *Requirements*

The synchronous scan feature must meet or exceed the following requirements:

1. Deliver solid samples within 5 µm of the X-ray beam focal point at 1000 Hz.
2. Provide a general interface usable in any experiment hutch without significant customization.
3. Maintain all existing ECS MCS features.

## DESIGN AND IMPLEMENTATION

Figure 1 shows the synchronous scanning architecture.

### *Hardware Used*

The synchronous scan feature is designed to work with EtherCAT Distributed Clock (EDC) hardware [1]. Typical MCS hardware includes a Beckhoff CX5240-0185 Programmable Logic Controller (PLC), EL1252-0050 timestamp terminal, and EDC-compatible drives: EL7062 for stepper motors and AX5206 for servos.

Due to demanding performance requirements, drives use Cyclic Setpoint Positioning (CSP) mode, which utilizes faster internal positioning control loops rather than PLC scan time [2]. Encoders are selected for protocols directly supported by the drives to enable CSP operation.

### *Clocks and Triggering*

Synchronous scan uses the EDC as a timing reference and a "Trajectory Clock" (TC) to track elapsed time since a specific EDC time. All synchronous scan times reference either EDC or TC time.

The LCLS Timing System outputs TTL triggers synchronized to within microseconds of XFEL X-ray pulses [3]. These triggers connect to an EL1252-0050 timestamp terminal, which detects rising edges and records timestamps in EDC time. By measuring intervals between triggers, the XFEL firing frequency can be accurately determined, enabling prediction of upcoming shots based on previous timing.

A trigger is output for every XFEL pulse. The first trigger starts the scanning sequence, and subsequent triggers indicate when the system should reach the next target point. Each new trigger timestamp applies a timing bias to synchronize setpoint generation with the XFEL, as detailed in the pulse trigger synchronization section.

### *EPICS Interface*

The EPICS interface is designed for easy transfer to any LCLS experiment setup. The real-time controls require three EPICS interface categories: sequencer commands, status, and target points.

1. Sequencer commands are boolean controls for program flow: "start", "stop", "pause", "resume", "reset", and "give permission to move to first target".
2. Sequencer status is summarized as one enumeration representing the sequencer state variable value and an error string.
3. Target points are 64-bit float pairs: position (in axis engineering units) and TC time (in seconds).

This completes the EPICS interface. It intentionally uses target points to meet the design requirement for general compatibility across all LCLS hutches, assuming any experiment requiring synchronous scan ultimately requires reaching target positions at target times. A GUI will be created for future system operation.

* Work supported by DOE contract number DE-AC02-76SF00515
† nlentz@slac.stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCR001

# INTEGRATING EPICS, OPC UA, AND TSN: A UNIFIED, AI-READY CONTROL ARCHITECTURE FOR PARTICLE ACCELERATORS AND LARGE RESEARCH FACILITIES

J. Cruz-Miranda[1], E.N. Al-Soliman[1], C. Megías[1], M. Damas[1], M. Cappelli[2], J. Díaz[1]
[1] ICAR Department, University of Granada, Granada, Spain
[2] Nuclear Department, Italian National Agency for New Technologies, Energy and Sustainable Economic Development (ENEA), Frascati (Rome), Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The increasing complexity and data demands of modern particle accelerator and large research facilities necessitate a paradigm shift towards unified, intelligent control system architectures. This paper proposes a study to enhance a particle accelerator control framework architecture in order to develop an AI-ready control system, drawing upon the strengths of both open-source and commercial control frameworks, on the basis of the IFMIF-DONES control design experience. By integrating the widely adopted Experimental Physics and Industrial Control System (EPICS) with industrial Supervisory Control and Data Acquisition (SCADA) systems via a key component –an OPC UA server for EPICS pvAccess– a seamless and standardized communication layer is established. This hybrid approach enhances flexibility, scalability, and long-term maintainability while leveraging the benefits of both ecosystems. Furthermore, the proposed architecture explores the potential of unifying traditionally separate networks, such as Timing, Control/Monitoring, and Interlock, through Time-Sensitive Networking (TSN) to simplify infrastructure and improve bandwidth utilization. This paper outlines the architectural design, discusses the advantages of this integrated approach in achieving AI readiness and improved interoperability, and highlights the role of a EPICS - OPC UA Gateway as a cornerstone for future advancements in accelerator control.

## INTRODUCTION

The control system is a critical component in particle accelerator plants and large research facilities, ensuring the precise coordination of complex subsystems to achieve stable and reliable operation. Its role is fundamental for safety, performance optimization, and enabling advanced experimental capabilities in high-energy physics and fusion research.

This work has been conducted within the context of the Work Package Early Neutron Source (WPENS) [1] project, taking the International Fusion Materials Irradiation Facility - Demo Oriented Neutron Source (IFMIF-DONES) facility [2, 3]. as a base example, but the scope of the proposal is broader, and may be applicable not only for the DONES and WPENS projects, but also for other accelerators and large research facilities, and the EPICS community.

The design of the control system for the IFMIF-DONES project has been proposed during a preliminary phase of WPENS, encompassing both the Central and Local Integrated Control Systems (CICS and LICS). All that work is documented in a series of publications [4-–10]. Looking ahead, it is expected that future WPENS efforts will shift toward research, feasibility studies, and validations, with the esDONES consortium taking over the design and construction responsibilities. The overarching goal remains a collaboration to deliver a unified and highly optimized control system that meets both established plant standards and the latest scientific advancements.

This study, builds upon the foundational WPENS design, proposes modifications to meet distinct requirements and specially unites some of the last research performed by ICAR department of the University of Granada in collaboration with the CISC group of the WPENS project.

The IFMIF-DONES facility's operation is contingent on the Central Instrumentation and Control System (CICS), a sophisticated system responsible for overarching supervision. The CICS oversees and manages all instrument parameters, handles data storage and visualization, and maintains continuous, bidirectional communication with the Local Instrumentation and Control Systems (LICSs). The LICSs, in turn, provide local monitoring and control of individual subsystems to ensure process variables remain within specified ranges.

A primary function of the CICS is to guarantee the safe and reliable irradiation of samples and to prevent damage to the equipment. Its architecture is composed of three main systems:

- Control, Data Access, and Communication (CODAC): This system manages supervision, timing, data, and the human-machine interface, and includes the alarm system.
- Machine Protection System (MPS): It implements strategies to protect against component, system, or operational failures.
- Safety Control System (SCS): This system ensures compliance with safety standards through its plant safety, personnel access safety, and radiation monitoring subsystems.

EPICS [11] is the control framework proposed for the CODAC. Particularly, the piece of software of the EPICS ecosystem that will offer the CODACs main services is CSS/Phoebus [12] which provides graphical operator interface, data archiving, and alarm management.

On the other hand, in the WPENS scope, an industrial SCADA has been proposed for the Control Framework of the other two CICS Systems (MPS and SCS); specifically,

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCR002

# CONTROL SYSTEM UPGRADES AT THE NATIONAL IGNITION FACILITY FOR HIGHER LASER ENERGY AND HIGHER FUSION YIELDS

M. Fedorov, A.I. Barnes, L. Beaulac, A.D. Casey, J.R. Castro Morales, J. Dixon, C.M. Estes, S. Heerey, R. Lacuata, B.P. Patel, M. Paul
Lawrence Livermore National Laboratory, Livermore, CA, USA

*Abstract*

Following the landmark achievement of fusion ignition in December 2022, the National Ignition Facility (NIF) has established ignition as a repeatable, high-yield operational regime. To access new experimental opportunities, NIF is advancing the Enhanced Yield Capability (EYC) upgrade, which will raise delivered laser energy from 2.2 MJ toward 2.6 MJ. This expansion enables projected target yields above 30 MJ while supporting ongoing high energy density and stockpile stewardship programs. The central challenge is to integrate hardware enhancements with the long-term sustainment of control systems, now approaching two decades of continuous operation. The NIF Integrated Computer Control System (ICCS) was originally designed with data-driven abstractions and modular upgrades in mind, allowing most changes to be implemented through incremental device models and configuration updates. However, select subsystems, notably embedded controllers, require architectural renewal to meet both reliability and performance demands at higher energies. We present the technical strategies and patterns that reconcile capability upgrades with sustainment in a live facility, highlighting the roles of phased deployment, end-to-end test automation, and the integration of applied AI and machine learning as tools for reliability and cadence. This experience demonstrates how tightly coupled enhancement and renewal provide a practical path for extending NIF's experimental reach and operational continuity.

## INTRODUCTION

NIF has demonstrated and repeatedly achieved fusion ignition, establishing a new operational baseline for high-yield experiments. The next step is the Enhanced Yield Capability, which raises delivered laser energy from 2.2 MJ toward 2.6 MJ and opens access to fusion yields projected above 30 MJ [1, 2]. In parallel, the core control systems, now approaching two decades in service, must continue to support higher energies and higher shot rates under evolving security and reliability requirements.

We frame the NIF controls interplay of challenges and opportunities as a "triangle": first, ignition is now a reliable regime, enabled by precision pulse shaping, robust diagnostics, and automated shot execution [3, 4]. Second, EYC extends this capability at the same precision standard, requiring coordinated changes in power conditioning, beam conditioning, and optics-damage mitigation. Third, long-term sustainment of the ICCS and hardware underpins continuous operations, blending legacy and modern components across the facility, and supporting ongoing modernization of both hardware and software [5].

Central to our approach is the deliberate pursuit of design and operational synergies: planning upgrades and sustainment actions together to minimize conflicts and maximize efficiency. Rather than treating enhancement and renewal as competing efforts, we integrate them, leveraging shared design elements, modular upgrades, and phased deployment to maintain experimental cadence. This approach extends beyond hardware, now encompassing applied AI and machine learning as tools for reliability, anomaly detection, and automated response. These AI/ML initiatives are embedded in daily operations rather than separated as research pilots, tying directly to facility reliability, uptime, and shot rates.

The following sections demonstrate how this integrated approach delivers practical results in the modernization of power conditioning, the architectural pivot in beam-based PSS, and other sustainment-driven upgrades. In each case, we show how blending enhancement with renewal and navigating the yin and yang of capability and constraint, enables ICCS to advance NIF's mission at the frontier of experimental science.

## IGNITION ACHIEVED AND REPEATED

NIF achieved the first laboratory fusion ignition on December 5, 2022, producing 3.15 MJ from 2.05 MJ of laser energy [6-8]. As of April 2025, NIF has achieved ignition nine times, setting a new record with an 8.6 MJ yield and a target gain of 4.1 using 2.08 MJ at 456 terawatts peak power. Multiple shots now deliver more than twice the input energy, and fusion yields continue to rise as symmetry control, capsule robustness, and energy delivery are further optimized (see Fig. 1).

This performance rests on the ICCS, which coordinates more than 66,000 devices across 2,300 front-end processors and embedded controllers, executing over two million device operations per shot. ICCS delivers beam pointing within tens of micrometers and shot synchronization to tens of picoseconds. The High-Fidelity Pulse Shaping system reduced shaping error from 5–20% to below 0.5% and improved shot-to-shot reproducibility by a factor of two to three, supporting repeatable delivery at ignition conditions.

Beam flattening using the Programmable Spatial Shaper and grey edge blockers suppresses high-fluence hot spots, enabling reliable operation above 2 MJ, while automated cryogenic layering forms uniform deuterium-tritium ice with millikelvin control using x-ray feedback [4].

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCR003

# PARALLEL CONTROL SYSTEMS: AN EFFICIENT AND LOW RISK APPROACH FOR A MIGRATION FROM VSYSTEM TO EPICS

K. L. Winter*, O. J. Kuiken†, M. Niemand, A. Gerbershagen
Particle Therapy Research Center (PARTREC), University Medical Center Groningen
Groningen, The Netherlands

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

For over 30 years, the AGOR cyclotron control system at UMCG PARTREC has relied on Vsystem. However, the limitations of Vsystem's aging technology stack hinder efforts to improve reliability. To address this, we have decided to migrate to EPICS. Given the limited IT resources at PARTREC, a cost-effective migration strategy is essential. Additionally, planned and unplanned accelerator downtime must be kept to an absolute minimum. Instead of the conventional approach of a gradual transition, we have opted for a different method: running both Vsystem and EPICS concurrently as fully configured control systems. During migration, all controllers will communicate with both systems simultaneously, ensuring continuity and minimizing downtime. This paper outlines the feasibility of this approach, its cost-effectiveness, and the proofs of concept conducted to validate its implementation.

## INTRODUCTION

In 2020 PARTREC, then KVI-CART, was transferred from the University of Groningen (RUG) to the University Medical Center Groningen (UMCG) [1]. This transition created opportunities for investments in the reliability and longevity of the AGOR accelerator, leading to numerous modernization projects [2, 3].

These modernisation efforts have also renewed the focus on the software behind AGOR's control system. The current control system is based on Vsystem by Vista Control Systems [4], but our implementation has several issues and limitations that impede reliable operation. The system exhibits stability problems requiring frequent restarts, and Vsystem's limited communication protocol support constrains our options for equipment integration into the control system. Furthermore, the fact that only a limited number of institutes have reported the use of Vsystem in publications for the ICALEPCS conferences over the past decade raises concerns about whether continued support can be guaranteed.

These issues have prompted us to look for a replacement, and an investigation into the available options has found EPICS to best fit PARTREC's requirements. However, several factors complicate the migration away from Vsystem. Control system migrations require substantial effort, yet our IT team is small. Additionally, the available time for a traditional subsystem-by-subsystem migration is severely constrained by operational demands. The accelerator produces beam approximately 26 weeks per year in 2-3 week periods, with maintenance periods typically lasting 1-2 weeks (except for a longer summer period). Control system modifications are not allowed during beam production, restricting subsystem migration to these brief maintenance periods. Furthermore, these maintenance periods are often allocated to other control system upgrades, as work is also underway to replace our Bitbus based controllers with Modbus TCP based controllers.

Given these constraints, a migration strategy has been developed prioritizing operational continuity and implementation efficiency. Our approach centers on running both Vsystem and EPICS in parallel, allowing for a transition with almost no operational downtime, and enabling implementation, testing, and verification during beam time. The EPICS design has been optimized to support this parallel operation and to enable an efficient implementation. This paper presents the EPICS design and migration strategy and demonstrates how they enable a cost-effective transition with minimal impact on AGOR operations.

## BACKGROUND INFORMATION

The superconducting AGOR cyclotron (see Fig. 1) began development in the late 1980s and was commissioned in 1997. AGOR's subsystems were controlled using I/O controllers that used the Bitbus fieldbus for communicating with the Vsystem based control system, which ran on the VAX/ELN architecture.



Figure 1: The AGOR cyclotron.

During commissioning, both Vsystem (investigated as early as 1992) and EPICS (investigated as early as 1995) were

* k.l.winter@umcg.nl
† o.j.kuiken@umcg.nl

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUCR004

# CONTROL SYSTEM CONSIDERATIONS FOR LANSCE MODERNIZATION AND INTEGRATION OF THE LAMP FRONT-END*

H. Watkins†, D. A. Dimitrov, C. D. Hatch, E. E. Westbrook
Los Alamos National Laboratory, Los Alamos, NM, United States

*Abstract*

The Los Alamos Neutron Science Center (LANSCE) has embarked on a major modernization effort through the LANSCE Modernization Project (LAMP), which has recently received approval for mission need. LAMP proposes a new 100 MeV front end to replace aging components, including the proton sources, Cockcroft-Walton generators, and 100 MeV drift tube linac. This new design will integrate with the existing cavity-coupled linac to reach the facility's full design energy of 800 MeV, with full deployment anticipated by 2030. The modernization project introduces significant control system challenges, particularly in integrating new high-performance subsystems while maintaining full operability of the legacy infrastructure, which will remain responsible for approximately 85 percent of the accelerator complex. This paper discusses control system strategies for timing synchronization, high-speed data acquisition, and software integration. Key topics include compatibility between legacy and modern control protocols, deployment of real-time data systems, and software development to ensure operational continuity. The LANSCE control system must provide seamless support for both existing and modernized hardware, enabling efficient operation and long-term sustainability.

## INTRODUCTION

LANSCE is a National Nuclear Security Administration (NNSA) facility that has delivered high-current proton beams for fundamental science, national security, and industrial applications for over five decades. Originally commissioned in 1972 as the Los Alamos Meson Physics Facility (LAMPF), the accelerator complex has since evolved into a multi-beam, multi-user facility, providing mission-critical capabilities to the U.S. scientific and defense communities. Today, LANSCE operates a high-power 800-MeV proton linear accelerator feeding five experimental areas, including the Lujan Neutron Scattering Center, the Ultra-Cold Neutron facility, the Weapons Neutron Research facility, Proton Radiography, and an Isotope Production Facility [1].

While LANSCE's beam delivery capabilities have remained robust and highly utilized, many of its systems date back to the facility's early operational era. Legacy hardware, aging custom electronics, and software environments built on now-obsolete platforms present increasing challenges for system maintainability, integration, and performance scalability. These limitations have become more acute as the operational demands on the accelerator grow in complexity, and as expectations for data integration, system diagnostics, and automation align more closely with modern scientific facilities.

To ensure long-term reliability, improve operability, and support future mission requirements, LANSCE has embarked on a comprehensive modernization effort. This modernization effort is not merely a technological refresh; it is a strategic investment in the resilience and agility of a facility whose experimental output continues to have national impact.

This paper outlines the architecture, objectives, and phased deployment strategy of the LANSCE control system in response to the modernization program. It highlights technical constraints specific to a continuously operating legacy facility and future directions toward aligning LANSCE's controls infrastructure with best practices in accelerator and large-scale experimental physics environments.

## LAMP

The current front-end of the LANSCE accelerator has exceeded its expected operational lifespan by more than three decades and suffers from increasing unreliability and diminishing maintainability. Its obsolete design and aging hardware have become significant sources of downtime and present a high risk of single-point failure that threatens all five mission-critical capabilities. The facility has experienced significant downtime due to legacy equipment such as the Cockcroft-Walton 750kV high voltage column as well as experiencing cracks in the copper cladding internal to the DTL [2].

To address these challenges, the LANSCE Accelerator Modernization Project (LAMP) aims to replace the aging front-end with a robust, maintainable system that reliably supports the facility's five principal capabilities: dynamic multi-frame proton radiography, advanced material characterization, ultra-cold neutron production, isotope production, and fundamental nuclear physics research. LAMP reached Critical Decision-0 in November 2024 and is an official DOE 413.3B project.

## LAMP DESIGN AND COMPONENTS

The LAMP design replaces LANSCE's aging front-end systems—originally commissioned in the 1970s—with modern, high-reliability technology from the ion sources all the way through the end of the 100MeV Drift-Tube Linac (DTL) [3].

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). LA-UR-25-29226

† hwatkins@lanl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG001

# UPGRADE OF THE LOS ALAMOS NEUTRON SCIENCE CENTER (LANSCE) BEAM CHOPPER PATTERN GENERATOR*

A. Braido[1, †], T. Ramakrishnan[1], S. A. Baily[1], L. Kennel[1], K. Griego[1], P. Torrez[1], C. Boyle[1,2]
[1]Los Alamos National Laboratory, Los Alamos, NM, USA
[2]Texas A&M University, College Station, TX, USA

*Abstract*

LANSCE delivers macropulses of beam, hundreds of microseconds in duration and at a nominal repetition rate of 120 Hz, to five experiment areas. These macropulses are distributed to four $H^-$ areas and one $H^+$ area. Each of the $H^-$ experiment areas require a unique beam time structure within the macropulse. This time structure is imposed on the beam by a traveling wave chopper located in the $H^-$ Low Energy Beam Transport (LEBT) section of LANSCE. The chopper is driven by pulsed power systems which receive digital signals generated by the LANSCE chopper pattern generator. This chopper pattern generator system must maintain tight synchronization with multiple LANSCE RF reference signals and is triggered by the LANSCE master timer system. This paper describes a recent upgrade to the LANSCE chopper pattern generator from its original NIM/CAMAC/VXI form factor, including details in software and hardware, test results, and future plans.

## LANSCE'S $H^-$ TIMING STRUCTURE

The chopper is a pulsed traveling wave structure at the front end of LANSCE, which rapidly deflects and undeflects 750keV $H^-$ beam during LANSCE's 120Hz macropulses, as beam is sent to different areas [1]. The chopper's time structure is unique to each area where $H^-$ beam is delivered within a given machine cycle. The time structures are as follows:

### *Weapons Neutron Research (WNR)*

WNR is a user facility based around a spallation neutron target which is fed by beam directly. The beam gate for WNR is known as the "Micropulse Enable Gate" (MPEG).

To support typical beam delivery to WNR, the Chopper produces beam pulses synchronous with a 16.77MHz structure in the LEBT known as the Low Frequency Buncher (LFB). Typically, one pulse is produced by the chopper and is bunched by the LFB, followed by 29 LFB RF cycles of off time to allow for time-of-flight measurements to be made at WNR.

### *Proton Radiography (pRad)*

pRad uses beam pulses within a macropulse to produce images of either static or dynamic (explosively driven) experiments. For pRad's operation, the chopper's timing is arbitrary and is set by the pRad experimenters to tailor to the specific experiment. Generally, pRad's "shot pattern" produced by the chopper is a 10's of ns precursor pulse followed by a train of 20-40 "radiography pulses" 100's of μs later.

### *Lujan Center/Proton Storage Ring (PSR)*

The Lujan Center at LANSCE is an experimental area, fed by a neutron spallation target, which is in-turn fed by the LANSCE PSR. Beam is accumulated in the PSR over many turns and is extracted onto the Lujan Center Spallation target at a 20 Hz nominal rate. The beam gate for PSR is known as the "Long Bunch Enable Gate" (LBEG).

For proper beam extraction from the PSR, the chopper must produce a pulse and a gap, synchronous with the PSR's circulation frequency of 2.8MHz to allow for some time for the PSR extraction kickers to fire during the beam gap. The timing for PSR beam is shown in Fig. 1.

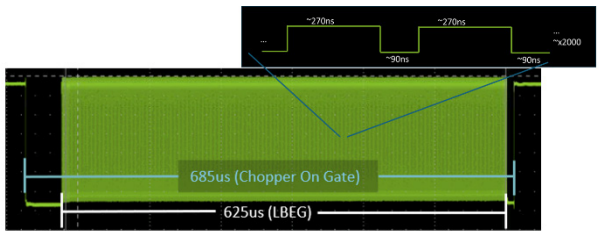



Figure 1: Typical time structure of PSR beam, ~290 ns on and ~90 ns off, within a macropulse of typically 625 us. Note that the chopper begins deflecting briefly before the beam gate and continues to deflect briefly after, within an envelope known as the "Chopper On Gate" (CHOG).

### *Ultra Cold Neutrons (UCN)*

UCN's typical time structure at the macropulse scale is 20 Hz burst of 9-11 macropulses with a ~5 second interval. However, inside these macropuplses is a chopper pattern approximately like PSR beam but with no requirement to be synchronous with or circulate in the PSR. The control room can modulate UCN's average macropulse current by changing the amount of "off time" they receive every 2.8 MHz cycle. UCN and pRad both are downstream of line X at LANSCE, and both use a beam gate known as "$H^-$ Go down line X" (H-GX) [2].

## THE LEGACY CHOPPER PATTERN GENERATOR

At a high level, the chopper pattern generator produces pulses synchronized with RF and with variable pulse width, phase, and decimation. The pattern also changes

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001. LA-UR-25-29065

† anthony.braido@lanl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG002

# ACCELERATOR PROCESS WATER UPGRADE AT ANL/APS

J. Stevens†, G. Shen, G. Swetin, R. Wright, J. Dench
Argonne National Laboratory, Lemont IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The APS accelerator process water control system was upgraded with a network of Allen-Bradley ControlLogix programmable logic controllers (PLCs) for pump stations and integrated with the Experimental Physics and Industrial Control System (EPICS) for remote control and monitoring. A Python tool, plcEpics, generates EPICS databases from structured Microsoft Excel worksheets, enforcing naming consistency and reducing integration errors. This paper describes the control-system architecture and commissioning experience, highlighting key integration issues and lessons learned.

## INTRODUCTION

Since user operations began in 1996, the Advanced Photon Source (APS) has used scheduled maintenance shutdowns to replace outmoded components. As EPICS software and PLC technology advanced, segments of the accelerator process-water system were upgraded incrementally. The APS Upgrade (APS-U) dark year starting in 2023 enabled comprehensive modernization of the entire process-water control system.

The APS accelerator process-water system comprises approximately 36 pump stations and 72 pumps (primary plus backups). This paper focuses on the 20 stations on the Storage Ring Mechanical Mezzanine (SRMM) located above the storage-ring experimental floor (Fig. 1).



Figure 1: Typical SRMM copper pump station.

A primary cooling circuit circulates nearly 10,000 gpm (≈ 630 L/s) to heat exchangers that serve closed-loop secondaries; individual secondary loops operate at 200 gpm or less (≈ 12.6 L/s) with tight temperature control.

* Work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No, DE-AC0206CH11357
† jbs@anl.gov

Two secondary subsystems serve the facility: the copper system supplies deionized water to copper-based loads, including magnets, power supplies, and absorbers; the aluminum system circulates deionized water to control vacuum-chamber temperature [1].

Each station includes pumps, heat exchangers, control valves, and instruments to measure flow, pressure, and supply and return temperatures. Deionized loops monitor conductivity (resistivity) to protect the aluminum vacuum chambers. Copper-loop water is degassed to reduce dissolved oxygen to below 10 ppb. The low oxygen content and high resistivity reduce erosion of components. Control objectives are to hold supply temperature and flow within limits, protect equipment through interlocks and alarms, and minimize thermal transients during accelerator operations.

The control software is organized by location—ITS (Injection Test Stand), Linac, PAR (Particle Accumulator Ring), Booster, RF Wing East/West, Storage Ring, and SRMM—with consistent EPICS process-variable names to support operator displays, alarming, and trending.

## PHASE 1 (1996 – 2007)

At the start of Phase 1, the SRMM process-water (PW) control system used 20 Johnson Controls local control panels (LCPs), one per water station.

In 1999, Allen-Bradley PLC-5 controllers were added to supplement the Johnson LCPs, enabling circulating-water temperature control within ±0.1 °F. The PLCs were integrated into EPICS for monitoring and supervisory control [2].

EPICS integration used a VME-based input/output controller (IOC) running VxWorks and iocCore. The IOC software was built on a separate host, booted over the network via FTP, and ran as the sole application on the single-board computer in the VME crate.

The PW IOC, installed in the accelerator Main Control Room (MCR), hosted two Allen-Bradley 6008-SV2R VME scanner cards (each dual port). The four scanner ports were wired via “blue hose” twisted-pair copper to four dedicated Allen-Bradley 1771-AF fiber-optic modules. Multimode fiber connected these 1771-AF modules in the MCR to corresponding 1771-AF modules on the mechanical mezzanine, with each pair serving one storage-ring quadrant (Fig. 2).

From each mezzanine quadrant 1771-AF fiber-optic module, a blue-hose Remote I/O (RIO) link ran in series (daisy-chained) to five Allen-Bradley PLC chassis. The fiber link between the MCR and mezzanine provided electrical isolation. With five chassis per quadrant, the system comprised 20 PLCs in total.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG003

# THE IRRAD PROTON IRRADIATION FACILITY DATA MANAGEMENT, ANALYTICS, CONTROL AND BEAM DIAGNOSTIC SYSTEMS: CURRENT PERFORMANCE AND OUTLOOK BEYOND THE CERN LS3*

F. Ravotti†,1, B. Gkotse[1,2,3], P. Jouvelot[3], M. Malkin[1], E. Matli[1], J. McCarthy[1], N. Minafra[1], G. Pezzullo[1], C. Scheuerlein[1], J. Szumega[1,3]
[1]Experimental Physics, Beams and Technology Department, CERN, Geneva, Switzerland
[2]University of Wisconsin-Madison, Madison, USA
[3]Mines Paris, PSL University, Paris, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The recent Proton IRRADiation Facility (IRRAD) at the CERN East Area was built in 2014 during the Long Shutdown 1 (LS1), and later improved during LS2 (2019), to address the needs of the HL-LHC accelerator and detector upgrade projects. IRRAD, together with the CHARM facility on the same beamline, exploits the 24 GeV/c proton beam of the Proton Synchrotron (PS), providing an essential service at CERN while showcasing more than 4400 samples irradiated during the last decade. IRRAD is operated with precise custom-made irradiation systems, instrumentation for beam monitoring (IRRAD-BPM), monitoring GUIs (OPWT) and a dedicated data management tool (IDM) for experiments follow-up and sample traceability. Moreover, performance tracking generated by custom-made analytics tools (Jupyter, etc.) guarantees operational feedback, thus maximizing beam availability for IRRAD. While the HL-LHC components qualification is coming to an end with LS3 (2026-2028), new challenges arise for future detector, electronics components and material irradiations: reaching extremely high fluence levels, operating low-momenta or heavy-ion beams, being some of those. In this context, we first describe the last improvements implemented at IRRAD after LS2 and then present the challenges ahead that will drive future upgrades such as, for example, applying ML techniques to the beam data, aiming to achieve real-time automatic beam steering and control.

## INTRODUCTION

The Proton IRRADiation Facility at CERN (IRRAD), located in the East Area of the Proton Synchrotron (PS) accelerator, is an experimental area primary tailored for the qualification of inner-detector components for High-Energy Physics (HEP) experiments. The IRRAD facility was first built in 1999 [1], originally on the T7 beam line, and underwent major upgrades in 2014 and 2019 during the CERN Long Shutdown 1 (LS1) and 2 (LS2) respectively.

During LS1, IRRAD was extended and moved to the T8 beam line, to cope with the increasing demand of irradiation experiments of the HEP community working for the High-Luminosity Large Hadron Collider (HL-LHC) [2]. Downstream IRRAD on T8, a twin experimental area named CERN High-energy AcceleRator Mixed field facility (CHARM) and equipped with a massive 50cm-long cylindrical target was also built during LS1. The secondary radiation field from the protons on target (e.g., the portion of beam that did not interact with the upstream material in IRRAD) reproduce accurately a broad mixed field (hadrons, photons, electrons, etc.) up to GeV energies, as expected in the CERN accelerator infrastructure, in the service areas of the HEP experiments but also mimicking ground- and atmospheric-neutron spectra as well as space environments [3].

From the very beginning of IRRAD operation post LS1 (see Fig. 1), several hundred samples are exposed yearly to the proton beam, serving today an enlarged community extending to electronics and material testing, accelerator components qualification and R&D users. As a result, the registration, planning and follow-up of these experiments require the management of a considerable amount of data. Therefore, during LS2, beside the commissioning of a custom-made IRRAD Data Manager (IDM) application, several upgrades were also made to the IRRAD control system, beam diagnostics, dosimetry and facility infrastructure required by the users to safely handle, store and perform post-irradiation measurements [4].

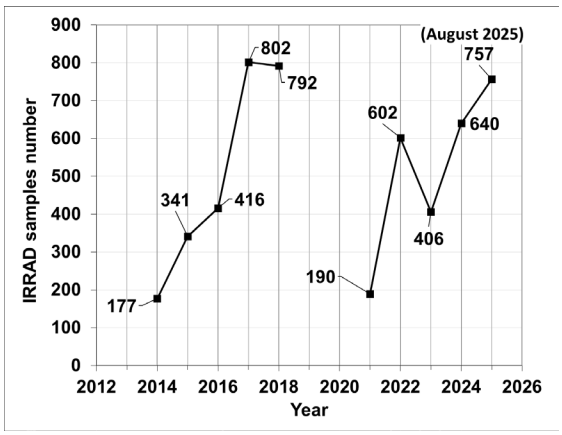



Figure 1: IRRAD statistics towards the end of the second run of operation after LS1 (2014). Data for the 2025 run (ending mid-November) refer to the month of August.

Looking beyond HL-LHC, CERN has now completed the FCC Feasibility Study [5], showing the technical feasibility of a future hadron collider (FCC-hh) with a centre-of-mass energy of approximately 100 TeV. This accelerator

* Work supported in part by the European Union's 2020 research and innovation program under grant agreement no. 101004761 (AIDAinnova), and no. 101057511 (EURO-LABS).
† Federico.Ravotti@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG004

# DEVELOPMENT OF AN EtherCAT-BASED CONTROL SYSTEM FOR AN IN-VACUUM UNDULATOR FOR SPring-8-II

K. Yamakawa[1, †], T. Fukui[2], M. Ishii[1, 2], M. Yamaga[1, 2]
[1]Japan Synchrotron Research Institute, Sayo, Japan
[2]RIKEN SPring-8 Center, Sayo, Japan

*Abstract*

SPring-8, a third-generation light source, has been operating for almost three decades. To transition to a fourth-generation light source, SPring-8 will undergo a major upgrade to its storage ring, which will be renamed SPring-8-II, with the upgrade scheduled for 2027-2028. This upgrade includes the installation of new In-Vacuum Undulators for SPring-8-II (IVU-IIs) and a corresponding modernisation of the accelerator control systems. As part of this, we are replacing the legacy VME-based systems with new EtherCAT-based slow control systems. The installation plan calls for maximum seven IVU-IIs to be deployed annually between 2023 and 2026, each with an EtherCAT control system. In 2024, we successfully implemented the first fully EtherCAT-based control system. This system meets the required 50 um synchronisation accuracy and ensures stable operation of the steering magnets and reliable temperature monitoring, confirming its suitability for the SPring-8-II upgrade program.

## INTRODUCTION

SPring-8 has been operating for approximately thirty years as a third-generation light source. The mainstream of storage ring light sources has shifted towards generating higher brilliance beams with low-emittance, a characteristic of fourth-generation light sources. Several such facilities have recently come into operation. SPring-8 will renew its storage ring, to be called SPring-8-II, to achieve a beam 100 times brighter than the current one by decreasing the beam emittance from 2.4 nmrad to 100 pmrad [1]. The scheduled shutdown will last approximately one year, with SPring-8 shutting down in 2027 and the new SPring-8-II commencing operation in 2028.

The upgrade of Insertion Devices (IDs) is one of the key components of the SPring-8-II upgrade programme. IDs, especially In-Vacuum Undulators (IVUs), are crucial instruments for generating light at SPring-8. In this programme, an upgraded IVU is referred to as the IVU for SPring-8-II (IVU-II) [1, 2]. The IVU-II features the cancellation of the attractive force between magnetic arrays. Installations of these devices began in 2023.

Furthermore, the accelerator control systems are also being upgraded prior to the SPring-8-II project. One of these control system upgrades involves the implementation of MADOCA 4.0 [3]. We have also decided to replace legacy VME modules for slow controls with EtherCAT [4] devices, which employ a master-slave relationship, due to the shrinking market for VME boards [5, 6].

† kosei.yamakawa@spring8.or.jp

## REQUIREMENTS

For an IVU-II, we are required to use EtherCAT to control the following elements:

- Gap control and monitoring
- Power supplies for steering magnets
- Magnet temperature monitoring

The gap is defined as the distance between the upper and lower magnetic arrays. For the IVU-II, we must control these arrays independently. To ensure beam stability and prevent the generation of unwanted spectra, both arrays are required to drive in sync. This synchronisation is critical to prevent the electron beam from being kicked off the center of the gap. The required accuracy is a difference of less than 50 um, a level of synchronisation that is sufficient to neglect any detrimental effects on the beam.

Steering magnets are essential for generating and correcting magnetic fields that cancel out erroneous components from the magnetic arrays. This precise correction is vital for maintaining the high-quality beam needed for advanced experiments.

The magnets' temperatures must be continuously monitored to ensure the operational stability and safety of the IVU-II. This is because temperature fluctuations can directly affect the magnetic field properties, which may lead to deviations in the beam's trajectory and stability. The typical magnet temperature range is 24 to 32 degrees Celsius.

Figure 1 shows the schematic view of the IVU-II magnetic arrays, driving mechanism, and control devices. Two steering magnets are positioned both upstream and downstream of an IVU-II.

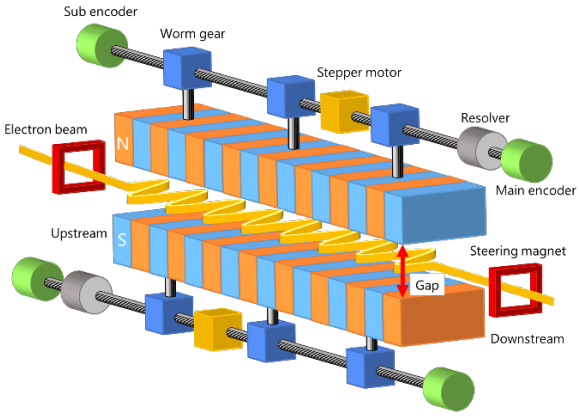



Figure 1: Schematic view of IVU-II magnet arrays, driving mechanism, measuring devices, and steering magnets.

The new system is VME-free and employs EtherCAT, because a lot of VME boards, used for IVUs at SPring-8, were discontinued and EtherCAT is widely used in industry,

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG005

# CONTINUOUS INTEGRATION ON TOP OF THE EXISTING CONTROL SYSTEMS OF LIPAc, FOR A RF CONDITIONING TEST BENCH*

L. Maindive†,1, J.P. Adam[3,4], A. De Franco[2], J. Franco Campos[2], L. González-Gallego[1], J. Hyun[2], I. Moya[3], F. Scantamburlo[3], and the IFMIF/EVEDA integrated project team
[1]IFMIF-DONES España Consortium, Granada, Spain
[2]QST-Rokkasho, Aomori, Japan
[3]Fusion for Energy (F4E), Garching, Germany
[4]CEA Paris-Saclay, Gif-sur-Yvette, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Under the Broader Approach agreement between Japan and Europe, the Linear IFMIF Prototype Accelerator (LIPAc) aims the validation of the International Fusion Materials Irradiation Facility (IFMIF) accelerator design, to produce a deuteron beam of 125 mA at 9 MeV in continuous wave (CW). While the superconducting linear acceleration stage is being installed, a high-power test bench (HPTB) was set up for the testing and conditioning of four pairs of radio-frequency (RF) couplers for LIPAc RF quadrupole (RFQ). The RFQ RF couplers are a key component to operate LIPAc toward CW.

Accordingly, the HPTB control systems (CS) were implemented in parallel to the existing CS of the LIPAc, benefiting from the tools available while avoiding their modification, and additional functionalities and devices were integrated to tackle the specificities of the test bench. This work was continuously performed during the test bench operations, identifying and answering further needs, such as deploying an automated conditioning tool or enabling slow feedback loops for automatic parameter tuning. Furthermore, this test bench became a testing environment for LIPAc CS refurbishments, such as upgrading the CS framework to EPICS v7, switching to CS-Studio Phoebus and its applications for the operator interfaces, or using Debian for the operating systems and ProxMox for the virtual environment. "The experience gained will prove valuable for the DEMO-oriented neutron source implementation of IFMIF (IFMIF-DONES) [1]."

## INTRODUCTION

In the scope of the engineering design and validation of the International Fusion Materials Irradiation Facility (IFMIF), the Linear IFMIF Prototype Accelerator (LIPAc) is reaching its final commissioning phases [2-4]. It is currently in a long maintenance period for the integration of its last acceleration stage, a cryomodule nesting eight superconducting radio-frequency (SRF) cavities.

After a first version for testing the temporary O-ring radio-frequency (RF) couplers of the RF quadrupole (RFQ) cavity before beam operations in 2023 [5-7], a new high-power test bench (HPTB) was developed for the testing and conditioning of the original RFQ RF couplers [8], having brazed ceramic windows instead of O-ring joints. Since summer 2024, the conditioning is processed by pairs of couplers, to reach thermomechanical validation up to 200 kW and continuous wave (CW).

The new configuration of the test bench is exploiting an RF module of the LIPAc dedicated to power injection into the SRF cavities, with additional hardware that is not part of the LIPAc. As such, it is the union of systems that do not belong together in the LIPAc design nor in beam operations [8] (Fig. 1).

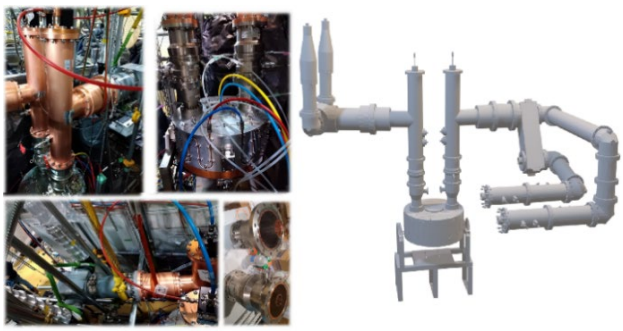


Figure 1: Latest configuration of the HPTB.

Since its first version, the HPTB operations have greatly improved with the addition and continuous implementation and integration of dedicated controls. These include not only a new abstraction layer that enables the test bench to be operated from a unique operator interface (OPI) (Fig. 2), but also the controls for software and hardware specific to the test bench, along with the integration of all these systems with LIPAc control systems (CS). These aspects will be detailed in this document, regardless of their chronological implementation.

## METHODOLOGY

The work presented in this document was performed "on-the-go", while the test bench was being set up and then during its operations. Because the main motivations were to provide tools that would facilitate the operations, a special care was put on streamlining the interfaces between the CS members and the test bench experts and operators.

### *Close to the Experts*

To fully comprehend the functional needs and the systems behaviours, the CS members made good use of the

* This work was undertaken under the Broader Approach Agreement between the European Atomic Energy Community and the Government of Japan with the support of EUROfusion (Grant Agreement 101052200). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the institutions listed.
† lucas.maindive@ifmif-dones.es

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG007

# DEVELOPMENT OF GigE VISION CAMERA CONTROL SYSTEM AND APPLIED TO BEAM DIAGNOSTICS FOR SPring-8 AND NanoTerasu

A. Kiyomichi[†,1], H. Dewa[1], T. Fukui[2], K. Ishii[3], E. Iwai[1,2], T. Maruyama[2], S. Matsubara[1], H. Oikawa[4], H. Sumitomo[3], T. Tsuchiyama[4], K. Ueshima[4], R. Yamamoto[3], K. Yanagida[5]
[1]JASRI, Sayo, Hyogo, Japan
[2]RIKEN SPring-8 Center, Sayo, Hyogo, Japan
[3]SPring-8 Service, Sayo, Hyogo, Japan
[4]QST, Sendai, Miyagi, Japan
[5]BREXA Technology Inc., Kawasaki, Kanagawa, Japan

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

As an imaging system supporting beam diagnostics using screen monitors (SCMs) at the SPring-8 site, we have continuously developed and improved a GigE Vision camera control system and expanded its adoption. By adopting the versatile open-source library Aravis, we eliminated vendor dependency and built an image acquisition system integrated into the SPring-8 control framework, MADOCA 4.0. Key features include the ability to control up to eight GigE cameras per computer with centralized management of camera power, trigger distribution, and screen operations. Its long-distance cabling enables flexible and simple deployment. Operation is achieved by writing the configuration file without programming, significantly reducing development costs and time.

As part of the SPring-8 upgrade, this system was successfully implemented for the SCMs of the beam transport line (XSBT) that uses the SACLA linac as the injector for the SPring-8 storage ring. We expanded the application of this system to the SCMs of the SACLA linac and the SACLA-BL1 linac (SCSS+), replacing the complex and costly Camera Link cameras. We also newly applied it to NewSUBARU injector linac and NanoTerasu in Sendai.

## INTRODUCTION

SPring-8, a large-scale synchrotron radiation facility in Japan, is a global leader in its field, driving innovation alongside the next-generation synchrotron radiation facility NanoTerasu. As a part of the SPring-8 upgrade project, the SACLA linac was integrated as a new injector for the SPring-8 storage ring. This project required a significant upgrade of the beam monitor system [1], which includes approximately 30 screen monitors (SCMs) along the 600-meter beam transport line.

Historically, beam diagnostic imaging at SPring-8 had been updated from analog cameras to Camera Link [2] standard cameras. While this standard offered high-bandwidth data transfer and trigger synchronization, its major drawback was a limited transmission distance of only 10 meters. This limitation forced the use of expensive optical converters and complex wiring to extend transmission distance using optical fiber, which complicated system architecture and increased maintenance costs. To address these challenges, we started the development of a new GigE Vision-based [3] control system, which enables long-distance data transmission over standard Ethernet cables, thereby simplifying the system and reducing costs.

This paper provides a detailed account of our efforts to overcome the technical hurdles of this system and its subsequent successful deployment in various facilities. We will elaborate on the fundamental design philosophy, the intricate details of hardware and software development, and the valuable insights gained from extensive operation, including the key improvements made to enhance system stability for cutting-edge applications.

## DESIGN CONCEPT

The core design philosophy was to build a flexible, simple, and integrated SCM control system that could seamlessly replace the complex, costly, and difficult-to-maintain Camera Link infrastructure. Our strategy was multifaceted, focusing on three key principles: vendor independence, system simplification, and efficiency.

First, to achieve vendor independence, we adopted the open-source library Aravis [4], which is a glib/gobject-based library conforming to the GenICam (Generic Interface for Camera) [5] standard. GenICam provides a unified interface for cameras from various vendors, allowing us to support a wide range of hardware without being tied to a single supplier. This approach ensures long-term maintainability and flexibility in hardware selection.

Next, we aimed for system simplification. The GigE Vision standard's capability for long-distance transmission (up to 100 meters) eliminated the need for complex cabling and expensive signal extenders. This simplified architecture not only reduced installation costs but also made the system easier to troubleshoot and maintain.

Finally, we prioritized efficiency and convenience in both system operation and management. The system was designed to control a maximum of eight cameras from a single computer, centralizing the management of camera power, triggers, and screen operations. The control software is seamlessly integrated into the SPring-8 control framework, MADOCA 4.0 [6]. This integration enables a "programming-less" operation, as users can start and manage the system by simply editing a configuration file, which drastically reduces development time and costs.

† kiyomichi@spring8.or.jp

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG009

# UPDATE ON MIGRATION TO EPICS AT THE ISIS ACCELERATORS

I. D. Finch*, A. Al Shafei, K. R. L. Baker, A. Kurup[1],
J. L. Fernandez-Hernando, G. Howells, M. Leputa, P. Ray, M. Romanovschi
ISIS, STFC, Rutherford Appleton Laboratory, Chilton, Oxon, UK
[1]Also at Imperial College London, London, UK

*Abstract*

The ISIS Neutron and Muon Facility accelerators are migrating to an EPICS control system. The tools developed to run two control systems in parallel and to automate the migration of hardware and user interfaces to EPICS have been previously presented. We now detail our emerging EPICS setup. Hardware interfaces are implemented as a mixture of conventional EPICS IOCs, in-house developed equivalents in Python, and bridged through our old control system. Our user interfaces are primarily the Phoebus stack but web interfaces in Python are being explored, particularly to support machine learning purposes such as automated optimisation and anomaly detection. We present issues which may arise at any site in transition, such as handling continuity of data archiving.

## INTRODUCTION

The ISIS Neutron and Muon Sources [1, 2] accelerators began the migration from the Vsystem [3, 4] commercial control system to the EPICS [5,6] control system in 2019 [7]. ISIS experiment operations have completed their transition to EPICS [8, 9] so, in this paper, we refer exclusively to the transition in progress in the accelerators control system.

The key technology in our transition remains pvecho [10, 11], software developed in-house in Python to allow the two control systems to remain synchronised and used simultaneously. Pvecho has two components:

- **v2e** (Vsystem-to-EPICS), creates a PV (with the same name by default) for each Vsystem channel, and keeps this PV synchronised with the Vsystem channel.
- **e2v** (EPICS-to-Vsystem, which monitors or sets an EPICS PV from another source, e.g. an IOC, and keeps a mapped Vsystem channel in synch.

The migration to EPICS has progressed to a stage where a significant number of hardware interfaces have been deployed, and operator interactions are performed through the EPICS control system. Further details are given regarding hardware interfaces, our services and middleware stack, user interfaces, and finally considerations other than hardware and software.

## INFRASTRUCTURE

Our server and network infrastructure remains largely unchanged since last reported [12]. However, further details of the server infrastructure and notes on changes to network monitoring are given here.

### *Servers and Docker Swarms*

Vsystem runs on a cluster of four HP Itanium servers running OpenVMS. The number of active scripts and databases has fallen as the transition to EPICS has continued.

EPICS related software is deployed into Docker Swarms managed using Portainer [13]. Operationally important EPICS hardware interfaces and CS-Studio Phoebus [14, 15] (hereafter Phoebus) services run on a production Docker Swarm [16].

The development and machine learning (ML) swarms have seen revisions. The development swarm is now a three node swarm to better replicate the production environment. The ML server has been separated from the development swarm.

The ML server is a Proxmox [17] based virtual machine host hosting three virtual machines (VMs). One VM is a single-node Docker Swarm cluster, hosting ML deployments as well as the experiment management instance MLFlow [18]. It is configured to be identical to the production swarm nodes and has full access to the ISIS controls network, except that its PVA Gateway restricts access to PVs prefixed ML: or DEV:. The MLOps workflow has been previously outlined [19]. The other two VMs are user machines, each with access to an Nvidia A100 GPU and 38 vCPU cores, as well as 200GB of RAM. These are multi-user terminal machines used for ML development [20] that have restricted access to the control network.

A common configuration [21] of the external PVA Gateways and internal UDP broadcast relays on the production, development, and ML swarms is required for management of pvAccess protocol EPICS traffic.

### *Networks and Network Device Monitoring*

ISIS has a controls network over two primary subnets and a third specialist isolated subnet [12].

Monitoring of the status of the network components has been moved to Libre NMS [22]. It is currently being used to monitor endpoints via ICMP, SNMP, SSH, TCP, DNS, DHCP, and/or HTTP status as relevant for operationally important network-attached devices. This has aided in identifying device failure, particularly in the inner synchrotron, before they impact operations.

The inner synchrotron's radiation environment causes frequent switch failures and CPS [23] boot issues which may usually be resolved through reboot or power cycling. During ISIS operational periods, access to the inner synchrotron requires that beam is off for a least 30 minutes to allow the radiation environment to become safe for personnel. This

* ivan.finch@stfc.ac.uk

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG013

# CONTROL SOFTWARE AND TECHNOLOGY CHOICES FOR THE ELECTRON-ION COLLIDER*

T. D'Ottavio, C. De Silva, J. Jamilkowski, M. L. Kabir†,
J. Laster, S. Nemesure, R. Olsen, K. Shroff, Z. Yin
Brookhaven National Laboratory, Upton, USA

*Abstract*

The Electron-Ion Collider (EIC) will succeed the current Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. For over two decades, RHIC and its injectors have relied on a homegrown Accelerator Device Object (ADO)-based control system, which has provided a reliable and efficient operational framework. However, the EIC's requirements - such as a greater number of subsystems, higher uptime, increased data rates, and other factors - demand significant enhancements. Advances in both hardware and software technologies since the RHIC era have expanded the range of available options, each with its own set of benefits and challenges. In response, the EIC plans to deploy state-of-the-art technologies to meet these elevated demands, favoring open-source and community-driven solutions wherever feasible. This paper will focus on the control software and the technology choices under consideration and the strategies being adopted for the EIC.

## INTRODUCTION

Brookhaven National Laboratory (BNL), located on Long Island, New York, is home to the Relativistic Heavy Ion Collider (RHIC), the world's only polarized proton–proton collider [1, 2]. In addition to polarized protons, RHIC is capable of accelerating and colliding heavy ions of various species. Its operations are supported by a chain of injectors—the Linac, Booster, and Alternating Gradient Synchrotron (AGS)—along with ion sources such as the Electron Beam Ion Source (EBIS) and the Tandem Van de Graaff (TANDEM). Operational since 2000, RHIC has a circumference of 2.4 miles. At present, two major experiments, STAR and sPHENIX, are collecting collision data at the interaction points [3]. Figure 1 presents an aerial view of RHIC, including its injectors and ion sources. The two hadron storage rings, known as the Blue and Yellow rings, are also indicated in the figure.

The RHIC will be succeeded by the future Electron-Ion Collider (EIC), designed to investigate the quantum chromodynamics (QCD) structure of matter. Construction approval will be sought in 2026 for the initial stage of the project. The EIC will provide both hadron and electron beams and is planned to include a Hadron Storage Ring (HSR), an Electron Storage Ring (ESR), a Rapid Cycling Synchrotron (RCS), a Low Energy Cooler (LEC), a Bunch Accumulator Ring (BAR), and Electron Injector Systems (EIS). Figure 2 presents a schematic of the proposed EIC and its major components. The injectors and ion sources from RHIC are expected to be reused. At present, a single interaction point has been approved.

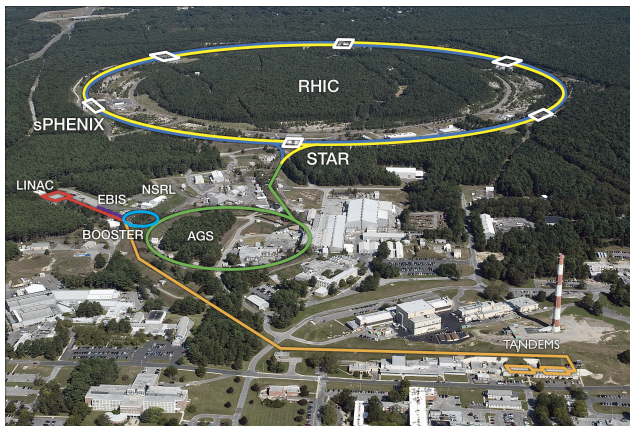



Figure 1: Areal view of the Relativistic Heavy Ion Collider at BNL.

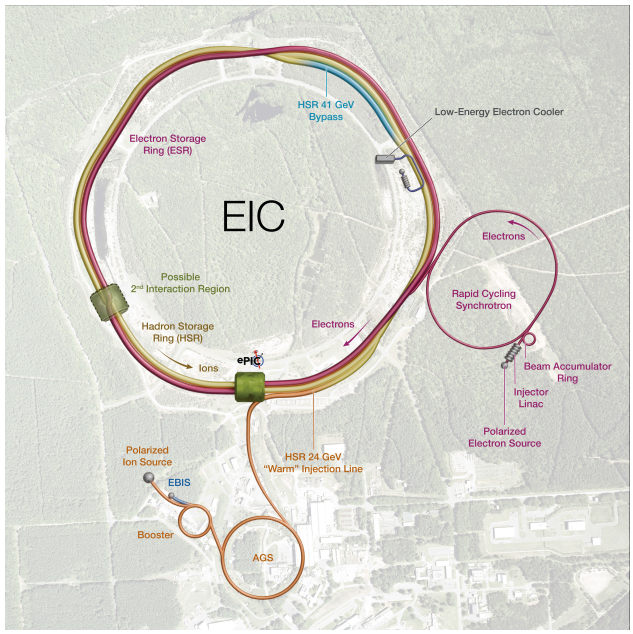



Figure 2: The proposed layout of the Electron Ion Collider at BNL.

## EIC CONTROL SYSTEM REQUIREMENTS

To establish the requirements for the EIC control system, it is useful to first review the RHIC control system for context. The RHIC and its associated injectors and ion sources

* Work supported by US DOE Contract Number DE-SC0012704
† kabir@bnl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG014

# NEW L2SI DYNAMIC REACTION MICROSCOPE ENDSTATION IN TMO: CONTROL SYSTEM DESIGN, INSTALLATION AND INTEGRATION*

J. Yin†, D. Kameswaran, Y. Mandza, B. Aljamal, D. Schafer
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

To harness the unprecedented capabilities of the world's most powerful X-ray beam delivered by the LCLS-II project, the former Atomic, Molecular, and Optical Science (AMO) instrument at the SLAC Linac Coherent Light Source (LCLS) has been upgraded to the Time-resolved AMO (TMO) instrument through the L2SI (LCLS-II Science Instruments) project. The new Dynamic Reaction Microscope (DREAM) endstation, also covered by the L2SI project and located at the second interaction point of the TMO, will offer unique capabilities to support cutting-edge research in the fundamental science of matter and energy. This paper provides an in-depth overview of the control systems for the DREAM endstation, detailing the architecture, design methodology, implementation, and seamless integration with the broader LCLS control infrastructure. It will also address key challenges, including integrating SmarACT motion control systems with the X-ray Machine Protection System (MPS) across different platforms, developing a robust and flexible equipment protection system, and implementing automated vacuum controls to meet stringent reliability and operational requirements.

## INTRODUCTION

The DREAM endstation requires precise control to overlap the X-ray beam and pump laser with extremely tight tolerances. To enable controlled gas-phase target delivery while preserving ultra-high vacuum conditions, the system employs a multistage differential pumping scheme. This approach allows localized gas injection into the interaction region without compromising the vacuum integrity of adjacent chambers. Additionally, a pair of Helmholtz coils is installed around the COLTRIMS (Cold Target Recoil Ion Momentum Spectroscopy) spectrometer to actively cancel the ambient geomagnetic field. This magnetic field compensation is critical for maintaining the momentum resolution required for coincident ion and electron detection. In addition, the system must maintain an ultra-high vacuum of $3 \times 10^{-11}$ Torr in the main chamber during gas injection [1]. An overview of the DREAM endstation and its major components is shown in Fig. 1.

The DREAM endstation is supported by three key control subsystems: A Beckhoff Programmable Logic Controller (PLC) system manages motion control for critical components such as the main chamber, gas jet assembly, and Helmholtz coils. This PLC-based setup allows for high-precision positioning and reliable operation under experimental conditions. A SmarAct piezo-based motion control system is deployed for the fine alignment of scatter slits, diagnostic paddles, and laser in-coupling optics, enabling submicron resolution and high repeatability. Another Beckhoff PLC oversees the vacuum infrastructure, including the differential pumping systems for the gas jet and catcher chambers, as well as the central experiment chamber. This centralized control approach ensures coordinated vacuum management across high- and ultra-high vacuum regions.All control subsystems are fully integrated with the LCLS Supervisory Control and Data Acquisition (SCADA) platform [2], notably the Experimental Physics and Industrial Control System (EPICS) [3]. While these control systems are based on the standard LCLS control architecture, significant additional development was undertaken to address the unique operational and integration challenges presented by the DREAM endstation.

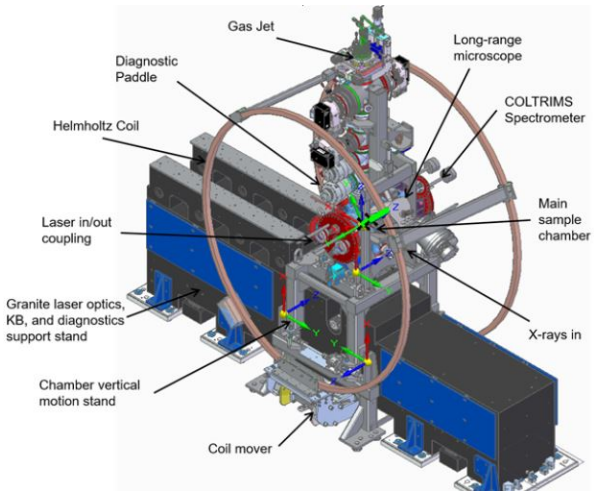



Figure 1: Overview of the DREAM endstation.

## VACUUM CONTROL PLC SYSTEM

The DREAM endstation vacuum infrastructure consists of three systems: a gas jet differential pumping system starting at $1 \times 10^{-4}$ Torr, a catcher differential pumping system, and a static main chamber maintained at $3 \times 10^{-11}$ Torr during continuous gas injection, see Fig. 2. A critical constraint is that the main chamber must remain at the ultra-high vacuum, requiring all turbopumps at full speed; any pump trip lasting more than 15 minutes must trigger a controlled venting to protect sensitive detectors. A further complication stems from the magnetic field of the Helmholtz coils, which can

* This work was supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-76SF00515.

† jyin@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG015

# TOWARDS THE IMPLEMENTATION OF A NEW EUROPEAN XFEL SCIENTIFIC DATA POLICY – CHALLENGES AND CHANCES

J. Malka*, S. Aplin, I.J. Bermudez Macias, F. Dall'Antonia, L. Gelisio, L. Maia, P. Schmidt, E. Sobolev, J. Szuba, W.T. Tsai, O. Turkot, G. Winterhoff, K. Wrona, European X-Ray Free-Electron Laser Facility, Schenefeld, Germany

*Abstract*

As scientific data volumes at European XFEL continue to grow rapidly, the need for sustainable storage, access, and preservation has become increasingly critical. In response and despite operating within an established data management environment, the facility has introduced a new scientific data policy to address rising demands and align with evolving best practices of scientific data management in the photon science community. The policy emphasises long-term sustainability and adherence to FAIR principles (Findable, Accessible, Interoperable, Reusable), promoting greater transparency, collaborative sharing, and sustainable reuse of research data. To implement it, several key concepts have been integrated. Data Management Plans (DMPs) are now mandatory from the experiment planning stage onward, providing a framework for researchers in defining data workflows and lifecycle strategies. Data reduction techniques are essential for managing storage demands without compromising scientific value. The data retention model has been changed to optimise usage of a tiered storage infrastructure, combining high-performance and mass storage systems for short and mid-term needs with cost-effective long-term archival. Metadata tools have been upgraded to improve discoverability, and access to open data has been refined to ensure secure and collaborative use in research. This contribution reports on these projects and outlines their significance for scalable, efficient data management in accordance with global standards.

## INTRODUCTION

European XFEL [1], a world-leading research facility in X-ray science, generates immense volumes of data from its cutting-edge experiments in physics, materials science, biology, chemistry, and many other disciplines. As experimental techniques grow in complexity and number of applications, the scale of resulting datasets increases, often reaching petabyte levels per experiment.

These data volumes, while offering outstanding opportunities for scientific discovery, pose major challenges for storage, access, processing, and long-term preservation. Thus, effective data management is not just a logistical concern but a scientific imperative. Recognising this, the European XFEL has updated its scientific data policy [2] to ensure that research data is managed sustainably and in accordance with international best practices, particularly the FAIR[1] principles [3].

* janusz.malka@xfel.eu

[1] Findable, Accessible, Interoperable, and Reusable

## THE NEED FOR AN UPDATED SCIENTIFIC DATA POLICY

Since the first user experiment in 2017, the facility has operated within a mature data management framework, guided by the scientific data policy [4] established beforehand. This framework is supported by high-performance computing resources, dedicated user support services, and strong collaborations with European and worldwide initiatives. However, the accelerating pace and scale of data generation required an update to the policy to address future scalability, evolving expectations for data access and usability, compliance with shareholder requirements, and emerging trends in the photon and neutron science community.

The updated scientific data policy places sustainability, transparency, and reusability of scientific data at its core. It aims to:

- Enhance the research data lifecycle - from experiment proposal submission to open data release and long-term preservation - through improved user support and systematic implementation of Data Management Plans (DMPs);
- Ensure that data management practices comply with FAIR principles;
- Deliver open data and encourage open science while protecting sensitive data;
- Optimise use of storage and computational resources to ensure the sustainable operation of the facility.

## KEY DEVELOPMENTS IN DATA MANAGEMENT

### *Data Management Plans*

One of the most significant updates in the data policy is the formal introduction of mandatory Data Management Plans (DMPs) for all user experiments, starting in the experiment planning phase. These DMPs are dynamic and evolving plans containing flexible and actionable documents that define responsibilities and processes for managing data throughout its entire lifecycle – from acquisition [5] and processing [6] to sharing, archival, and reuse. The purpose of requiring early and systematic planning is to promote consistent and sustainable data workflows, improve coordination with support services, and ensure compliance with international data management standards, for example, following the best practices recommended by the European Commission's Horizon Europe Programme Guide [7].

The rollout of DMPs represents the culmination of a multi-phase process initiated several years ago. The initial phase

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG016

# UNIVERSAL LIMS FOR DIAMOND LIGHT SOURCE

I. A. Bush*, B. Adegbaju, R. Adey, M. Comina, J. Irvine,
M. Pritchard, J. M. Wilcoxson, Diamond Light Source, Harwell, UK

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Universal LIMS is a new set of web services being developed at Diamond Light Source, as part of the Diamond-II upgrade. It will provide users and beamline scientists with tools to manage the logistics and scientific metadata for their experiments. For scientific samples it will allow users to ship them to Diamond, track where they are within the experiment hall, and store data about them. Universal LIMS will also allow users to define parameters for non-interactive experiments and will provide an overview of the data they have collected, including metadata from the data acquisition systems and summaries of analysis pipelines that have run. These services will work together to provide a consistent, streamlined workflow to facilitate user experiments at Diamond.

A key part of the approach is to provide flexibility in the data that is stored. Defining a database schema to cover the needs of the eight different science groups in Diamond would be challenging and there is a development cost to updating the database schemas as requirements evolve. Instead, with Universal LIMS we store data as JSON, validated against a JSON schema. This ensures schemas can be updated easily, while the data can still be understood effectively by downstream applications.

## MOTIVATION

Currently at Diamond a Laboratory Information Management Service (LIMS) exists for the Structural Biology beamlines, comprising of the ISPyB database and the SynchWeb web-based front-end. For other beamlines at Diamond no such system exists. A new LIMS system will be delivered as part of the modernised beamline software architecture. This will be called Universal LIMS, to distinguish it from Diamond's existing structural biology focused LIMS[1].

The workflows for Structural Biology are well defined and consistent across the beamlines that use it. Hence the ISPyB database has a schema that fits this specific use case well, but the workflows are tightly aligned with the database schema, and hence it is difficult to adapt to other use cases. The architecture for Universal LIMS is driven by the need to create a flexible system, that can easily be adapted to the needs of different beamlines and technique areas, whilst avoiding the need for the development team to deploy code changes as data structures evolve.

* ian.bush@diamond.ac.uk

[1] eBIC, I19 and some other beamlines also use the structrual biology LIMS system; eBIC and I19 are expected to continue with the current system as it is closely aligned with the Structural Biology workflow, while other beamlines can transition when Universal LIMS has reached a mature state.

## DATA SCHEMAS

To provide the required flexibility a system using schemas to define what metadata is stored is being used. The schemas can be created by a defined set of super-users, and then metadata can be stored that is guaranteed to conform to the schemas. Schemas will be immutable, so all stored metadata will reference a schema that exists in the appropriate schema service. This will mean consumers of the data, for example someone who wants to analyse the data, will understand and be able to make correct use of the data.

Schemas will be versioned, following semantic versioning [1] , as commonly used for APIs. A minor version change will mean that the schema should not break existing consumers that rely on the data in the schema. This means that the only changes allowed as minor changes are to add new fields, and to make optional fields required fields. Adding new required fields would require services creating and editing the data to be updated at the point which they switch to use the new schema, but services downstream that have targeted a specific version would not need to be changed unless they wish to make use of the new field. A major version change is effectively a new schema, so this means there is an expectation that both creators and consumers of the data will need to be updated.

To ensure maximum reusability of schemas we will allow schemas to be composed of other schemas. So, for example, Schema A and Schema B may exist as schemas for two different techniques, and these could be both used within Schema C, alongside other metadata required by Schema C. This approach will also allow for common parts of schemas to be broken out and re-used, for example a list of chemical elements or absorption edges.

### *JSON Schema*

JSON (JavaScript Object Notation) Schema [2] will be used to define the schemas. JSON itself is a standard, widely used format for data interchange. JSON Schema provides a way of creating an additional JSON document that provides a fixed schema for the data in other JSON documents; it enables JSON data consistency, validity and interoperability at scale. JSON and JSON Schema are both standard, widely used approaches with mature tooling in a number of different languages. JSON Schema provides support for basic types such as numbers, strings and arrays. Objects are a collection of these types, and sub-schemas can be used to build-up more complex schema types in a clean way. Additional custom validation can be applied, for example for entities such as email addresses, STFC Federal IDs (unique user accounts) or date-time objects. This kind of custom validation would be created and made available by the development team as needed, not by users directly, and documented for users

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG017

# ADAPTIVE APPROACH TO SPATIAL INTERPOLATION AND VISUALISATION OF SCATTERED MONITORING DATA AT CERN

C. M. Musso*, A. Ledeul, A. Savulescu, G. Segura Millan
CERN, Geneva, Switzerland

*Abstract*

In order to ensure safe operations, CERN leverages an extensive SCADA system to monitor radiation levels and collect environmental measurements across its premises. The Health & Safety and Environmental Protection (HSE) Unit addressed the challenge of visualising radiation fields from non-uniformly distributed sensors across large areas. This paper presents the approach and implementation of a 2D interpolation and visualisation system for such measurements. The system implements two algorithms for two complementary monitoring scenarios. The Inverse Distance Weighting (IDW) interpolation addresses cases where radiation sources are located near sensor locations, assuming maximum values occur at measurement points. The Radial Basis Function (RBF) method handles scenarios with potential measurement peaks between sensors. A region network approach divides large areas into independent regions for optimised performance. Accuracy is evaluated using leave-one-out cross-validation. The architecture relies on TypeScript, React and WebSockets. The system processes measurements across CERN's premises and provides operators with spatial visualisation of radiation levels.

## INTRODUCTION

CERN's HSE Unit operates one of the world's most important radiation monitoring systems; measuring radiological and environmental impact, ensuring workplace safety, and providing data for reporting to public and host state authorities. This monitoring infrastructure has 90 device types, manages more than 125,000 alarms, and archives around 80 billion measurements per year.

While real-time alarm systems effectively detect immediate radiation hazards, post-mortem and long term analysis requires a different approach and remains important to address. When radiation alarms occur, operators need to reconstruct the dynamics and dispersion patterns from the localised measurement values of multiple monitors. The current monitoring system provides detailed temporal data visualisation but lacks a field view for spatial interpretation of scattered sensor data.

The reconstruction of field-maps is beneficial in incident analysis, where understanding the spatial patterns of radiation fields can inform safety and operational decisions.

This paper presents an adaptive approach to spatial interpolation and visualisation of radiation monitoring data at CERN. We address two distinct monitoring scenarios that require two complementary interpolation methods: Inverse Distance Weighting (IDW) for cases where radiation sources are located at or near sensors, and Radial Basis Functions (RBF) when source locations are unknown and may exist between measurement points.

The implementation builds on established academic research: adaptive parameter selection for IDW [2] and space subdivision with adaptive shape parameter for RBF [3]. It also introduces a practical solution to numerical stability issues encountered during real scenarios testing. Similar interpolation approaches have been successfully applied to environmental monitoring in agricultural soil property mapping [4] and groundwater depth distribution analysis [5].

The system is integrated into the REMUS Web application [1], CERN's main lightweight web platform for radiation monitoring, offering a new functionality for spatial visualisation of radiation data.

## SYSTEM REQUIREMENTS AND METHODOLOGY

We establish functional requirements and interpolation strategy selection based on physical monitoring characteristics and operator workflow needs.

### *Functional Requirements*

The solution addresses a specific valuable requirement: helping operators analyse radiation distribution patterns when time series monitoring is insufficient. This becomes particularly useful during incident investigations, safety assessments, or when unusual radiation patterns appear across the monitoring region. Operators require a flexible tool that handles multiple measurement units, applies temporal data aggregation, offers interpolation method selection, and enables adjustment of sensitivity parameters and spatial resolution to meet their monitoring objectives (Fig. 1).

The system must complete processing within reasonable amount of time (seconds) and consistently deliver either a valid visualisation or an explanatory error message.

### *Interpolation Strategy Selection*

The key insight is matching interpolation methods to the physical assumptions of monitoring scenarios. When sensors are positioned on known sources, we use this additional information to apply a specific computational approach.

We distinguish two monitoring contexts. In controlled environments like CERN's ventilation systems, air is released from underground areas through conduits with sensors installed at exits; the highest value is thus assumed to be exactly at sensor location. In contrast, open monitoring areas distribute sensors across regions without predetermined pat-

* carlo.maria.musso@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMG018

# THE ESS SYNCHRONOUS DATA SERVICE (SDS) DEVELOPMENT AND FIRST RESULTS

J. P. S. Martins*, G. W. Kowalski, J. Brudvik, J. F. Esteban Müller, S. Brewer, S. Armanet, T. Korhonen,
European Spallation Source ERIC, Lund, Sweden

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The 5 MW proton linear accelerator of the European Spallation Source ERIC is designed to accelerate the beam at a repetition rate of 14 Hz, which dictates the refresh rate of most of the relevant data produced by acquisition systems. Each cycle of the 14 Hz timing structure receives a unique cycle ID from the ESS Timing System, which can be used as an index when data is collected and stored. The ESS Linac Synchronous Data Service (SDS) facilitates the collection of high-resolution data from various accelerator subsystems. Currently, SDS consists of an EPICS extension to be included in data acquisition IOCs and a client service (SDS Collector) that collects and stores the data produced by these IOCs. The novel features provided by the EPICS PVAccess protocol and libraries play a crucial role in this project by supporting structured data in the EPICS Process Variable data format. This paper outlines how SDS is designed, how it enables data-on-demand and postmortem collection of large array datasets without overloading the network and describes the results of using SDS during the latest ESS beam commissioning campaign in 2025. From a broader perspective, SDS will be part of the ESS Operational Data Framework, which comprises a set of tools to collect, store, catalog, retrieve, and analyze ESS Linac data to support advanced applications such as machine learning algorithms. This framework is briefly described in this paper.

## MOTIVATION

The motivation for the SDS project is a combination of three aspects: indexing data with a numeric cycle identifier, exploring EPICS [1] (Experimental Physics and Industrial Control System) Normative Types [2] custom structures for PVs (Process Variables) and archiving of large array PVs.

ESS Linac operates in pulsed mode, with data from DAQ (data acquisition) systems synchronized to the 14 Hz proton beam repetition rate provided by the ESS Timing System [3]. Each cycle receives a unique integer cycle ID in addition to the usual timestamp. Using cycle ID as an index aims to simplify offline analysis, facilitating correlation of data across different subsystems.

Because ESS is an EPICS facility, this cycle ID must be integrated into the EPICS PV format. This makes it an ideal use case for the structured data support introduced in EPICS 7, providing SDS with a research component: evaluating these new features in production.

A further driver is the management of large array datasets. DAQ systems for the ESS Linac generate raw and processed arrays at sampling rates from 1 MSPS up to 250 MSPS and with 8000 to 300000 points, published as EPICS array PVs. Archiving these at the beam repetition rate quickly overwhelms network, storage and computing resources. SDS targets these array datasets by offering a specialized archiving approach that reduces bandwidth usage and enables efficient postmortem and on-demand data collection. Figure 1 shows a sequence of high-speed analog acquisitions from a BCM (Beam Current Monitor) sensor, an example of SDS collection after a postmortem event.

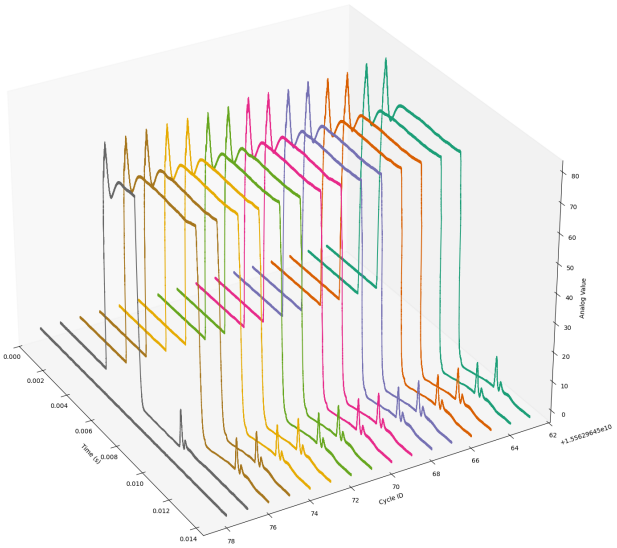



Figure 1: Sequence of BCM acquisitions collected with SDS.

## ARCHITECTURE

The SDS architecture is built around three main components. The first is the SDSIOC module, an EPICS extension that allows IOCs (EPICS Input/Output Controllers) to publish data in the SDS format. The second is the SDS Collector, a client service responsible for monitoring SDS PVs and storing the acquired data for later analysis. Finally, the ESS Timing System, based on the MRF (Micro-Research Finland) [4] platform, distributes the cycle ID to all acquisition nodes, ensuring synchronization across the facility. A simplified schematic of the SDS architecture is shown in Fig. 2.

### *SDSIOC Module*

The SDSIOC [5] EPICS extension module enables a standard DAQ IOC to produce SDS data. Eligible systems must connect to the ESS Timing System, acquiring data triggered by deterministic MRF events. This guarantees that acquired datasets can be matched to the correct cycle ID.

* joaopaulo.martins@ess.eu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR001

# A SHARED VIRTUAL MACHINE FRAMEWORK FOR EPICS HANDS-ON TRAINING

R. Lange*, ITER Organization, St.-Paul-lez-Durance, France
G. Knap, Diamond Light Source, Didcot, UK
K. Shroff, Brookhaven National Laboratory, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Facilities relying on collaboratively developed software projects, like the Experimental Physics and Industrial Control System (EPICS), often face challenges in ensuring consistent skill levels and efficient onboarding of staff.

This paper introduces a new framework for creating reproducible, pre-configured virtual machine (VM) environments, specifically designed for hands-on EPICS training.

A key benefit of this framework is its ability to establish shared, reusable, general training modules. Such shared resources are highly valuable for collaborations, as they provide a standardized platform for skill development, reduce redundant training efforts, cultivate a common understanding of EPICS mechanisms, and ultimately strengthen collective knowledge within and across institutions.

## FRAGMENTED KNOWLEDGE IN COLLABORATIVE STRUCTURES

### *The "Knowledge Silo" Problem*

Named after the term coined by Phil Ensor in 1988 [1], many organizations have consistently reported suffering from "Knowledge Silo" syndrome. Knowledge and insights are not shared or accessible to others, leading to a lack of cross-functional collaboration, insufficient communication, and inhibited learning.

This is especially true for large organizations spread across cultures, continents, and time zones, where communication is already more difficult and less efficient. It also more strongly affects loose collaborative structures with members of different nature and size, where different traditions and organizational procedures create additional hurdles to efficient cooperation.

### *The EPICS Context*

The EPICS (Experimental Physics and Industrial Control System) [2] collaboration is such a large, loosely coupled collaborative structure, between member organizations and laboratories of different kinds and sizes, spread across many continents and time zones. Obviously, the "Knowledge Silo" syndrome is an intrinsic part of this system.

One area where this becomes clear is the scope of EPICS software skill management, that is training courses for existing staff and onboarding of new hires. Many institutions share the goal of ensuring up-to-date, consistent EPICS competencies, but still have each developed their own training materials and environments, adapted to their specific ways and procedures of using EPICS.

This has led to a set of distinct disadvantages:

**Duplication of effort and resources** Creating and maintaining training materials and environments requires a lot of time and effort. Parallel developments that only differ in a few aspects of their content are wasting resources.

**Inconsistent training quality and content** Even in parts covering the same aspects that should be identical, different training materials used across the collaboration show different quality, age, and state of maintenance.

**Difficulty in onboarding high-mobility staff** Moving between institutions or working on multi-institutional projects are common and well-established in the EPICS collaboration. Inconsistent training quality creates difficulties for those practices.

**Stagnation of best practices within isolated groups** Innovative and/or effective solutions to common problems hardly make their way out of an institution and never into another organization's training materials.

## A SHARED HANDS-ON TRAINING ECOSYSTEM

In an effort to start reducing the negative consequences of the siloed approach, this project tries to create a framework for a shared hands-on training infrastructure.

### *Definition*

A shared hands-on training ecosystem provides the technical context for running hands-on training sessions for many trainees, which can be reused by different institutions for different training events.

While the technical feature of sharing a VM (virtual machine) is an important part of this approach, the far more important and hopefully long-lasting aspect is promoting a culture of co-creation and shared learning across the inter-institutional gaps.

### *Core Principles*

The agreed principles of the ecosystem:

**Standardization and Best Practices** The installations on the shared VM shall be configured in a consistent fashion, using the standard, up-to-date options that EPICS provides. When trainees inspect the VM contents, they shall see things being located in obvious places and best practice approaches wherever possible.

* ralph.lange@gmx.de

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR002

# AUTOMATED TEST EXECUTION FRAMEWORK FOR PRE-EXPERIMENT CHECKOUTS

C. Tsoi-A-Sue*, R. Tang-Kong, Z. Lentz, K. Lauer,
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

At the Linac Coherent Light Source (LCLS), maintaining the health and correct configuration of the control system is essential to ensuring successful experimental operations. Automated Test Execution Framework (ATEF) supports this goal by enabling the development and execution of automated procedures to verify control system status before the start of the user's experiment. ATEF focuses on pre-experiment validation of devices and controllers—ensuring they are correctly configured, network-visible, and responsive to control system commands. These verifications range from confirming process variable (PV) status and device communication to setting and reaching predefined setpoints. By providing a framework to establish and routinely verify a configuration baseline, ATEF allows users and system integrators to ensure the control system is ready to commence operation. This paper will present the design of checkouts within the ATEF framework at LCLS, explore real-world examples of its application, and highlight how it supports the broader goals of system integrity and experimental readiness.

## INTRODUCTION

In the past, we evaluated experiment readiness by manually checking that all systems met their specified setpoint limits. This assessment was conducted for nearly every device along our beamlines, including motors, vacuum controllers, and laser systems. The manual checkout of each device was often time-consuming and challenging. To streamline this checkout process, Experiment Control System (ECS) division at Linac Coherent Light Source (LCLS) developed controls architecture of the standardized components known as common components throughout the beamline. At LCLS, a common component is a standard device that is reused multiple times along the beamline. Following a consistent configuration, which incorporates established devices, setpoint states, and naming conventions, has resulted in the development of efficient checkout processes. We wanted to develop a tool that can quickly check the health of our beamline experimental controls system and report back any issues that need to be resolved.

## DESIGNING A CHECKOUT

ATEF classifies checkouts as either passive or active. In passive checkouts, we query the status of process variables (PVs) and validate them against established setpoint values, reporting the findings back to the user. This approach is read-only and provides a safe option for stakeholders to assess the status of a system without altering any setpoints. Conversely, the active checkout feature allows us to command devices to move, change, or enable signals to a specified value. Additionally, we can incorporate passive checks within the active steps to confirm that the devices successfully achieved the desired changes.

To design an effective checkout for a system, it is essential to meet with stakeholders to understand their objectives for the checkout process. When meeting with a stakeholder, it is beneficial to have them compile a list of all the devices within their system, along with the signals they wish to interact with. This information allows us to determine the appropriate type of checkout to construct. If the objective is to validate the state of the system, we instruct them to define the success criteria for each process variable (PV) on their list, and we would then develop a passive checkout. Conversely, if their checkout involves modifying any signals, we will create an active checkout and encourage further discussion about the specific devices included in the process. A well-designed checkout should be applicable to multiple areas of the lab (hutches) and effectively report the status of each hutch's device components. In the ATEF, it is important to note that passive steps can be added to active checkouts, but not the other way around.

To create a passive checkout, it is advisable to have the stakeholder complete a table outlining the desired signals they wish to verify against established setpoints. From this information, we can group the setpoints by device type and, if necessary, create multiple checkouts for various sub-devices to ensure thorough validation.

Figure 1 illustrates the values that a stakeholder would like to valid for an SR-620, a time interval counter [1], a device that laser scientists use to measure laser timing relative to a fixed EVR trigger. Along with the corresponding setpoint values they wish to confirm. Based on this information, we can now proceed to develop a passive checkout to verify that the device is in the correct state. ATEF has a few tools that help run comparisons on sets of PVs, these are, Device Configurations, PV Configurations, and Tool Configurations. Device Configurations allow users to reference our Ophyd [2] device database, and select signals as defined in these python abstractions instead of specifying specific PVs. The Happi [3] search window provides an interface for searching and filtering this database. For PV configurations, PVs can be manually added in plain text, and comparisons can be written based on the provided PVs. Tool Comparisons are different in that you have to first select a tool, currently we have only implemented a ping feature, but you can also add multiple devices to a 'Hosts' list, which will all be pinged

* ctsoi@stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR003

# NO CHILD LEFT BEHIND: MANAGING REQUIREMENTS, INTERFACES, AND COMMUNICATION IN HIGH-IMPACT PROJECTS *

M. Cabral†, K. LaFortune, K. Buehner,
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

Managing requirements and interfaces throughout the project life cycle is a challenging endeavor, regardless of the size or magnitude of the system. At Linac Coherent Light Source (LCLS), the Systems Engineering team has developed practices to ensure that requirements and interfaces are effectively captured, communicated, and validated in an effort to reduce downstream issues during installation or operations. This paper outlines how the Experiment Control Systems (ECS) group at SLAC National Accelerator Laboratory leverages systems engineering methods and tools—particularly Jama—to structure, manage, and communicate requirements and interfaces for high-impact projects.

## SYSTEM ENGINEERING NETWORK

High-impact projects at SLAC such as LCLS-II-HE and MEC-U follow a V-model system engineering framework [1], as illustrated in Fig. 1 below. Project work begins with high-level global requirements that flow down to component-level specifications and ends with integrated system testing. While requirements are essential to guiding designs and defining the system, management of interfaces – roles, responsibilities, and interconnections between teams – is equally critical because they enable the whole system to become greater than the sum of its parts.

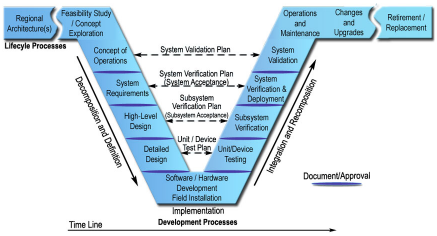



Figure 1: The V-model framework [1]. This diagram illustrates the structured approach taken in high-impact projects like LCLS-II-HE and MEC-U, illustrating the flow from requirements decomposition and design to system implementation and verification.

The ECS group collaborates closely with SLAC's Systems Engineering team to ensure requirements and interfaces are properly documented, reviewed, traced, and validated. This life-cycle is fully supported in Jama [2], SLAC's primary requirements management tool. It's important to note that although LCLS-II-HE follows a similar systems engineering framework, it utilizes a different tool. However, moving forward, all major projects at SLAC will leverage Jama and existing workflows. Before diving into the tool itself, it is useful to clarify SLAC's working definitions of requirements and interfaces in projects.

### *Requirements vs Interfaces*

**Requirements** describe what a system must do. It's the expectations for design results without depicting solutions. They describe the performance, function, or features needed to design a system, answering the fundamental question "what" for a project, and as they are decomposed to lower levels, also address the "how".

Table 1: Requirement Item Type

| Field | Description |
|---|---|
| Project ID | Unique ID within Jama Project |
| Global ID | Unique ID within Jama |
| WBS Level | Project WBS or WBS Level |
| Requirement Type | Functional, Safety, Performance, etc. |
| Name | Displayed Name |
| Description | What must the system do |
| Basis | Why must it do this |
| Verification Method | How will it be verified |
| Owner | Who owns this requirement |
| Status | Proposed, released, etc. |
| Verification Timeline | When requirement needs to be verified |
| Stakeholders | Who cares about this requirement |

Table 2: Interface Responsibility Item Type

| Field | Description |
|---|---|
| Project ID | Unique ID within Project |
| Global ID | Unique ID within Jama |
| Name | Display Name |
| Description | Unique interface description |
| Interface From | Stakeholder 1 |
| Interface To | Stakeholder 2 |
| Interface From Responsibility | Stakeholder 1 responsibility |
| Interface To Responsibility | Stakeholder 2 responsibility |
| Boundary Condition | Assumptions, constraints, etc. |
| References | Reference/supporting doc. |

**Interface Responsibilities** define the boundaries between systems or teams. They are the higher-level answers to "who"

* This work was supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-76SF00515

† mcabral@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR004

# HDB++, A RETROSPECTIVE ON 5+ YEARS USING TIMESCALEDB

D. Lacoste, R. Bourtembourg, Grenoble, France
S. Rubio-Manrique, J. Ramos, ALBA-CELLS, Cerdanyola del Vallès, Spain
L. Pivetta, G. Scalamera, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy
J. Forsberg, D. Egorov, MAX IV Laboratory, Lund, Sweden
T. Juerges, SKA Observatory, Jodrell Bank, United Kingdom
G. Jourjon, CSIRO, Canberra, Australia

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Tango HDB++ project is a high-performance, event-driven archiving system that stores data with microsecond resolution timestamps. HDB++ supports various backend databases to accommodate any infrastructure choice, with TimescaleDB as the default option. TimescaleDB, an extension of PostgreSQL, is selected for its exceptional performance, reliability, and open-source license. After more than five years of using the system in production at major facilities such as the ESRF, MAX IV and SKAO, this paper presents the insights gained from operating HDB++ with TimescaleDB in a large research facility. Results are presented considering various perspectives. From a performance standpoint, the paper examines how the scalability features have maintained low query response times despite the continuous growth in data volume over the years. From the system administration perspective, findings show that standardized and proven technologies have consistently supported high-quality service delivery. Lastly, from the user perspective, we analyze how users can query data stored from the inception of the project up to last week within seconds, either from the Python API or from clients like Grafana. This capability is also enabled by the successful migration and integration of archived data from older or different systems into the database in full compliance with HDB++ standards.

## INTRODUCTION

The Tango Historical DataBase++ (HDB++) archiving system [1] is developed through a collaborative effort among various facilities (Alba, Astron, Byte Physics, DESY, Elettra, ESRF, INAF, MAX IV, Observatory Sciences Limited, S2Innovation, SKAO) utilizing Tango Controls [2]. Since its inception [3], HDB++'s purpose has been to store Tango attribute values, i.e. process variables, in a backend of choice with microsecond resolution timestamps. The HDB++ community meets regularly and maintains a centralized repository with all HDB++ related developments and tools [4].

## HOW DOES IT WORK?

Tango is an object-oriented control system built around the concept of devices. Devices hold attributes and can run commands. These can be executed or queried via a standard request and reply mechanism. On top of that, devices can emit data in the form of events that follow a standard publisher/subscriber scheme. HDB++ is fully event-based, and devices push the data they want to store. A special device, the EventSubscriber, is in charge of subscribing to these events and inserting them into a backend of choice.

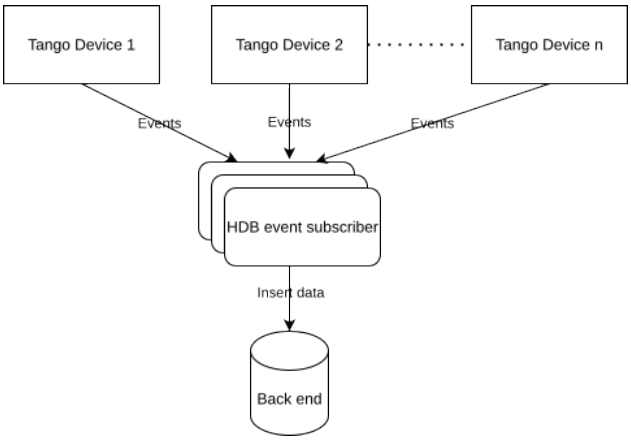



Figure 1: HDB++ Overview.

The design in Fig. 1 is quite efficient as data is pushed asynchronously by the devices, their normal operation is not hindered by the need to read data in order to archive it. It offers a simple way to scale as one can have as many EventSubscribers as needed. If there is a need to absorb more signals, simply instantiate a dedicated EventSubscriber. And finally an abstraction was introduced at the insertion level so that it is possible to insert data into different backends. The same abstraction level is used by the PyHdbppPeriodicArchiver process, to archive periodic data from those systems that for any reason (e.g. firewall rules) are not able to send events.

### *Extraction and Displaying of Data*

HDB++ provides two graphical applications for archiving configuration and data extraction, HDBConfigurator [5] and HDBViewer [6]. In addition, there are several tools for CLI, API, and web extraction. HDB++ can be configured using scripts and yaml files [7], and a unified Python API has been developed for data extraction [8].

The pyhdbpp library provides a common abstract class for developing reader objects, and working implementations for most used backends, MySQL and TimescaleDB. As most of the Tango Controls tools, pyhdbpp is available for download on the Python Package Index [9] and conda-forge

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR005

# EPICS SUMMER SCHOOL: TRAINING FUTURE SCIENTIFIC CONTROL SYSTEM EXPERTS

L. Porzio*[1], S. Vadilonga[1], W. Smith[1], M. Bajdel[1], M. Montis[2], M. Giacchini[2]
[1]Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Berlin, Germany
[2]Istituto Nazionale di Fisica Nucleare - Laboratori Nazionali di Legnaro, Legnaro, Italy

*Abstract*

The EPICS Summer School addresses the persistent demand for expertise in control systems for advanced scientific facilities by providing a crucial training platform for university students and young professionals. The highly successful inaugural event at BESSY II in Berlin offered an effective learning experience through a structured two-week program encompassing one week of foundational lectures and one week of practical application in a hands-on group project with real scientific hardware. This approach demonstrated a significant positive impact, empowering a new cohort from these groups with essential EPICS and distributed control system skills. In an effort to make the event sustainable, a collaborative framework has been established with other institutes, paving the way for a future where hosting the summer school cycles between facilities. This aims to make this training accessible to a wider audience of future scientific leaders, ensuring a continuous supply of skilled engineers capable of supporting and advancing groundbreaking scientific research across multiple facilities.

## INTRODUCTION

Scientific facilities worldwide face an unprecedented challenge in securing qualified control systems engineers. The demand for expertise in frameworks like EPICS (Experimental Physics and Industrial Control System) for distributed scientific control systems far exceeds supply, creating a bottleneck that threatens the operational efficiency of major scientific installations. This shortage stems from three primary factors: insufficient promotion of scientific control systems in universities, facilities preference for experienced engineers due to costly on-the-job training requirements, and documentation that presents steep learning curves for newcomers. The complexity of modern scientific control systems demands specialized knowledge that traditional academic programs struggle to provide.

## THE LEARNING PLATFORM

### *The Idea*

To directly address the identified challenges limiting control systems workforce development, we established a learning platform structured as a two-week intensive summer school format. This event was specifically designed to address the key training challenges hindering the pipeline of qualified engineers. First, to address the lack of promotion and advertisement, we created a highly visible, international program that actively recruits students from diverse academic backgrounds, providing them with exposure to career opportunities in large-scale scientific facilities that they might otherwise never encounter. The program's application process, requiring CVs and motivational letters from a global applicant pool, effectively bridges the gap between traditional industrial control systems education and the specialized requirements of scientific research environments. This approach positions scientific facilities as a viable and compelling professional trajectory distinct from conventional industry roles.

Second, recognizing that facilities often prefer candidates with prior experience due to the high costs of training, our platform front-loads the essential training in an intensive, structured environment. The two-week format provides facilities with candidates who have already overcome the steepest part of the learning curve, making them significantly more attractive hires. This approach effectively socializes the training costs across the community while delivering more job-ready candidates.

### *The Structure*

The summer school is structured into two intensive weeks, each with a distinct educational focus. The first week is dedicated to lectures and training sessions delivered by control system engineers from the hosting institute and participating external laboratories. These sessions present a range of control system implementations in large-scale facilities, demonstrating the broad applicability of these skills to various scientific and technical challenges. A significant portion of the week is devoted to comprehensive EPICS training, where theoretical principles are reinforced through practical exercises, guiding students from working in a simulated environment to hands-on experience with real hardware. Additional hours focus on data acquisition, experiment orchestration, and best practices in scientific data management, including the FAIR principles.

The second week is entirely project-based. Students are divided into small teams and assigned group projects. Each team integrates real scientific hardware into EPICS (Fig. 1), performs calibration, sensor conditioning, and implements graphical user interfaces. They also orchestrate the devices to enable experimental data collection. The summer school concludes with final presentations, where student teams showcase their work to an audience of teachers, tutors, and expert guests from the host and external laboratories (Fig. 2).

* luca.porzio@helmholtz-berlin.de

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR006

# THE ENHANCED LIQUID INTERFACE SPECTROSCOPY AND ANALYSIS (ELISA) BEAMLINE CONTROL SYSTEM PROTOTYPE

M. Brendike*, M. Bajdel, L. Porzio, W. Smith, S. Vadilonga
Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Berlin, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Enhanced Liquid Interface Spectroscopy and Analysis (ELISA) beamline is a new instrument at BESSY II focusing on a novel, integrated approach for the high-fidelity preparation and investigation of liquid interfaces using soft x-ray radiation and infrared radiation light generated simultaneously from the BESSY II light source. As ELISA is part of the BESSY II+ upgrade scheme, it will be the first soft x-ray beamline at BESSY II to use new hardware motion control standards and a novel BESSY II EPICS deployment system.

This contribution focuses on workflows and tools we have developed for beamline control at BESSY II. We demonstrate their application at the ELISA beamline and finish with an outlook on scaling usage across the full experimental hall.

## INTRODUCTION

The Enhanced Liquid Interface Spectroscopy and Analysis (ELISA) beamline is the first large instrument coming into operation in the context of the BESSY II+ project. The wider BESSY II+ project is the latest upgrade to the BESSY II synchrotron light source at Helmholtz-Zentrum Berlin (HZB) and aims to enable BESSY II to meet the growing demands of next-generation research in energy materials [1]. ELISA is currently under construction, and scheduled to see first light in 2026.

In addition to building new instruments, BESSY II+ aims to set new control system standards at the HZB. This includes new motion control hardware as described by Neu et al. [2], together with a novel approach on control system software.

This contribution focuses on a new approach towards control system software and motion control integration. The aim of the presented control system architecture is to automatise and simplify workflows for the end user by using a single tool for development, storage and deployment of all control system related software. ELISA is used as the environment for developing and testing this control system prototype, which will be used for both new and upgraded BESSY II beamlines.

## BEAMLINE OVERVIEW

Figure 1 shows the beamline's optical layout. ELISA is designed as a dual-branch beamline covering an x-ray energy range from 40 to 3500 eV. An infrared (IR) radiation beam, ranging from 0.05 to 0.5 eV, can be used simultaneously with the x-ray beam. Both beams are generated by the BESSY II light source [3].

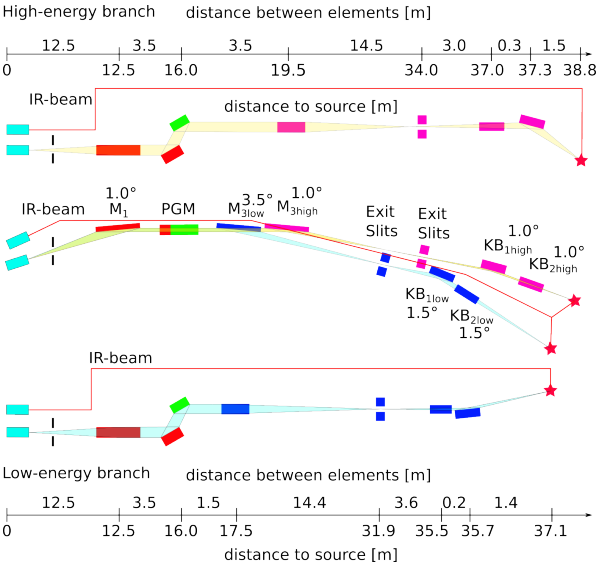



Figure 1: Sketch of the beamline's optical path [4].

The monochromator at ELISA will be fitted with three diffractive gratings. Next to a 400 $l/mm$ laminar grating and a 1200 $l/mm$ blazed grating, a multilayer coated 2400 $l/mm$ blazed grating will be used. This will extend the energy working range up to 3.5 keV. Higher photon energies require the monochromator to work at shallow incidence angles. Keeping a fine energy resolution even at such high energies results in even higher demands in positioning precision for the monochromator. Most of ELISA's mirror units are positioned by hexapod based positioning systems. This combination of high demands in positioning precision for the monochromator and complex hexapod based positioning systems makes ELISA a challenging environment for the new motion control standard.

All beamline components are fully motorized and equipped with absolute encoder feedback. This brings the total to more than 70 motorized axes. To integrate these axes into the control system a highly automated approach is used.

## CONTROL SYSTEM ARCHITECTURE

A sketch of the control system is shown in Fig. 2. To keep the architecture modular, a four-layer-based approach is used. The layers are defined as:

1. mechanics,
2. electronics,
3. communication, and
4. user interface (UI).

* maxim.brendike@helmholtz-berlin.de

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR007

# TOWARDS SUSTAINABLE WORK MANAGEMENT IN SCIENTIFIC FACILITIES: APPLYING KANBAN AT ALBA CONTROLS

Z. Reszela†, F. Becheri, R. Homs Puron, O. Vallcorba, S. Rubio-Manrique
ALBA Synchrotron, Cerdanyola del Valles, Spain

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Controls engineers in scientific facilities manage numerous projects, support multiple Customer Units (CUs), and balance operational demands with new initiatives. This often leads to growing backlogs and staff stress. Meanwhile, managers struggle to allocate resources across CUs, frequently prioritizing short-term goals at the expense of long-term strategy. At ALBA, such pressures prompted a new organizational approach. Building on lessons from successful past transitions - including the 2013 shift to operations and a major staff turnover - the Controls Section adopted the Kanban method to optimize the resource allocation and maximize the throughput. Tasks were categorized by size/complexity and Class of Service (CoS), and a unified board with multiple views was implemented to visualize Work In Progress (WIP) and support follow-up. All work begins in a visible backlog, jointly prioritized with CUs based on CoS. Dedicated engineering teams were formed to improve coordination and knowledge sharing. Policies and metrics are clearly defined and transparent. The implementation was done using Jira and Confluence (Atlassian ecosystem). First results of the new approach are included in the paper. This initiative aligns with broader organizational efforts such as the Activity Plan for ALBA (APA) and project tracking within the Computing division, laying the groundwork toward the ALBA II upgrade to a 4th generation synchrotron light source.

## ALBA CONTROLS AND DAQ SECTION

The Controls and Data Acquisition (DAQ) Section (the Section) is staffed by a team of 19 engineers. They design, develop, operate, and support the control systems of the accelerators, beamlines, and laboratories and the data acquisition systems required for the experiments. In addition, the Section is responsible for personnel and equipment protection systems.

### *Team Composition*

The Section consists of two distinct engineering profiles: controls software (15 engineers) and PLC automation (4 engineers). Four controls software engineers serve as group heads, providing technical leadership to less experienced engineers and coordinating the Section's work. Each year, 3-4 university and vocational students join the Section for 6-12 months internships. The entire Section is managed by the section head.

### *Role and Interactions*

The Section is part of the Computing Division, and collaborates closely with other sections – Electronics, IT Systems, Management of Information Systems (MIS) and Scientific Data Management – to meet the computing needs of the ALBA organization.

For controls and DAQ requests, it provides direct support to multiple Customer Units (CUs) at ALBA, mainly from the Experiments and Accelerators Divisions, but also from Engineering Division and Safety Office.

The Coordination Office is responsible for inter-division project coordination in ALBA. Since the implementation of the Contingency Plan for Controls [1] it has coordinated all CU-requested projects that the Section works on.

The Section also works with subcontractors for some specific projects, such as the CU-requested ones, as well as on certain maintenance or developments.

### *Evolution of Work Organization*

Since the beginning of the ALBA project, the cornerstone of the Section's work organization was assigning each engineer as a contact person for one or multiple CUs. Initially, these contact persons were also the project executors, a model that proved very effective during the ALBA construction and commissioning phase.

However, when ALBA transitioned into operation, the Section faced a major staff turnover, and the reliance on contact persons created single points of failure in the work organization.

To mitigate this, the Section was re-organized into four Scrum teams: Beamlines, Machine, Generic Software and PLCs [2]. Team members began working more closely, designing and developing solutions together, and automatically acting as backups in case of workload peaks or staff absences.

## ANALYSIS AND DESIGN

The second major turnover happened during COVID-19 pandemic. This time, the team-based work organization proved to reduce the single-point-of-failure effect inherent in the previous contact-person model. Thanks to the Contingency Plan [1], the Section recovered very quickly from this challenge. This difficult period was also an opportunity for a thorough retrospective on work organization. During the retrospective meeting we identified planning and lack of clear policies as key areas with significant room for improvement. Additionally, the upcoming ALBA II upgrade will demand unequivocal policies to balance operational support and development projects. These factors convinced

† zreszela@cells.es

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR008

# TRIGGER SYNCHRONIZATION UNIT CONSOLIDATION FOR SPS AND LHC DUMPING KICKERS SYSTEMS

N. Voumard*, M. Blaszkiewicz, L. Felsberger,
M. Niccollini, J. Sanchez, L. Strobino, CERN, Meyrin, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The TSUs (Trigger Synchronization Unit ) goal is to synchronize beam dump requests with the Beam Abort Gap (BAG) upon request from clients, primarily the Beam Interlock System (BIS). Synchronization is crucial to prevent damaging absorbers/collimators and cryogenic ring magnets. The TSU system is part of the Trigger Synchronization and Distribution System (TSDS). The TSU system, per beam, comprises two interconnected, redundant chassis for synchronization. It internally regenerates the Beam Revolution Frequency (BRF) in case of signal loss. A synchronization between both chassis and an internal surveillance ensures high level of redundancy. Dump triggers are issued via three paths: two synchronous and one asynchronous with delay. Software analyses post-mortem all dump triggers and critical signals for interlocking on any unissued or abnormal triggers. The TSU is a safety-critical system for both SPS and LHC Beam Dumping Systems , requiring high reliability. At the LBDS, one asynchronous dump is allowed per year of operation, and one catastrophic failure is acceptable per thousand years. During the design process, a full hardware and principle reliability study was conducted with CERN's Reliability Analysis Working Group (RAWG). In addition, a lab test bench is designed to test all functional requirements using unit tests. This paper presents the TSU system implementation, high-level software, GUI, and reliability study steps and results. The study confirmed the TSU's high reliability, meeting initial specifications, with no single point of failure identified.

## MOTIVATION FOR TSU SYSTEM CONSOLIDATION

The principal motivation to consolidate the TSU System [1] is the redesign coming with the BISv2 [2] upgrade planned for CERN's Long Shutdown 3 (LS3) and the requested change in the interface from the BISv2. Until LS3, the TSU is connected to the BIS optical loops directly and does the decoding of the BIS loop frequencies, as such making the TSU part of the BIS. Following the consolidation of the BIS, the Machine Protection and Electrical Integrity Group (TE/MPE) has requested to implement the frequency decoding themselves to have a unified approach for all BIS actuators; the BIS loops frequencies will be decoded on their dedicated BIS Actuator Board (CIBAB) which asserts BIS loops status to the TSU system. This allows TE/MPE to have a better diagnostic. Moreover, the interface CIBAB-TSU is fail-safe while today's TSU is not able to provide a fail-safe interface. Furthermore, the TSU card components are ageing and some are now obsolete. As well the AMD synthesis and programming tools are deprecated. This leads to a renovation and consolidation of the TSU for LS3; the new boards being expected to last for the whole duration of the HL-LHC, until around 2042. Additionally, the Accelerator Beam Transfer group (SY/ABT) profits from this renovation to enhance and harmonize the TSU systems on both SPS and LHC Beam Dumping Systems (respectively SBDS [3, 4] and LBDS [5]).

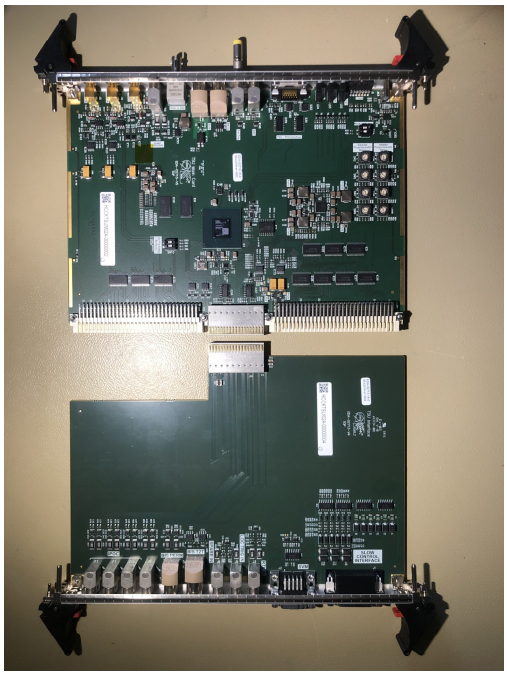


Figure 1: TSU main and RTM boards.

## TSU SYSTEM OVERVIEW

The TSU system, which is part of the TSDS [6, 7] is a critical safety element of both LBDS and SBDS, and concentrates dump requests from all over the LHC/SPS via the BIS, as well as from Beam Loss Monitor Direct Dump (BLMDD) and Beam Dump kickers control state, including its State Control and Surveillance System (SCSS) and Beam Energy Tracking System (BETS). A loss-free extraction requires a particle-free gap in the circulating beam during which the field of the extraction kicker magnets can rise to its nominal value. The beam dump requests need to be synchronized with this particle-free gap using the Beam Revolution Frequency (BRF) generated by the Radio Frequency group (SY/RF) to avoid damaging elements in the dump line.

* nicolas.voumard@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR009

# FAST SCANNING ARCHITECTURE FOR DIAMOND LIGHT SOURCE

G. Nisbet*, T. Cobb, Diamond Light Source, Didcot, Oxfordshire

*Abstract*

As the Diamond II upgrade is delivered, fast scanning will be a key to exploiting the brilliance of the new source. An outline of the fast scanning architecture to be rolled out across the Diamond beamlines is presented along with an explanation of the reasoning behind it.

## INTRODUCTION

There are several fast scanning use cases to be accommodated, for example, SWIFT, which is a beamline for fast operando spectroscopy, requires fast trajectory energy scans with position compare. I13-1 requires snake, spiral, GT scans (Fig. 1) with time based and position compare based triggering. I16 which is the Materials and Magnetism beamline requires both linear and non-linear reciprocal space mapping. I14 and CSXID which are nano-focusing beamlines require snake, spiral, and GT scans with nano-focussing. Important considerations for developing the fast scanning architecture, include the location of coordinate system, detector rates, data rates to the controllers, the types of scan that will be required in terms of shape, speed, dimensionality, complexity and resolution.

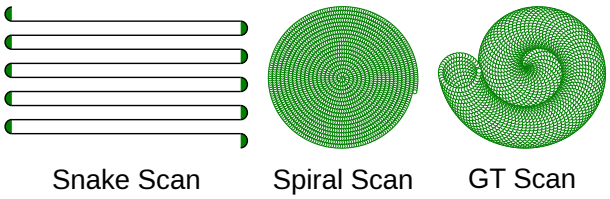



Figure 1: Example trajectories.

## KEY CONSIDERATIONS

- Coordinate Systems
- Ease of rollout
- Fast scanning and high performance control infrastructure (Advanced Models and AI/ML)
- Closing the Loop for engineering solutions

## COORDINATE SYSTEMS

Many of the beamlines at Diamond use hardware with complex coordinate systems such as diffractometers, robot arms and Stewart platforms which can have multiple discrete solutions requiring decision making strategies. When operating close to the edge of these solutions, small perturbations can force a jump from one solution to another. This can lead to discontinuities in the motion and in some circumstances, to hardware damage.

The new paradigm is to define scan paths in the measurement space and to transform to motor or joint space using transformation functions or services accessible by OphydAsync. This will be a bidirectional process and can accommodate analytical and numerical solutions. This approach allows the motion trajectories to be checked, visualised or both, before the motion is triggered. The raw motor PVT points (Position, Velocity, Time) can be calculated and streamed to the controller.

This approach is heavily dependant on the streaming rate that is possible with the controller. The out of the box streaming rate for the PowerPMAC is 400 Hz for two channels. Recent in-house developments suggest this could be pushed to 2 kHz. When very high performance is require, an FPGA based controller will provide the platform for advanced controls techniques employing tools such as Linear Quadratic Regulators (LQR), Predictive Model Control (PMC), Nonlinear Dynamic Inversion (NDI) and ML/AI - Reinforcement Learning.

## EASE OF DEPLOYMENT

The proposed solution eases deployment by simplifying the architecture. This is possible because the solution is cleanly partitioned. Scan paths will be defined in the measurement space, transformation tools make the conversion to motor space PVTs which are then streamed to the controller. The PVT steaming will be common to all beamlines. There will still be the need for instrument specific optimisation.

## ADVANCED CONTROLS INFRASTRUCTURE

Full utility of Diamond II's increased brilliance and coherence will require next generation scanning, capable of operating at high speed in the nano-scale regime. At this scale, motion systems can become highly nonlinear and under actuated creating challenges beyond the capabilities of traditional control methodologies. The new architecture, and more specifically, the high performance controller will push the limits of motion control systems and create a platform for deploying advanced control models. This will begin with system identification and the extraction of data driven models. Control laws will be determined using advanced controls techniques and ML solutions such as agent based reinforcement learning. In addition to improved performance, this process will inform and update the engineering methodologies. The controller could be accommodated in its own unit or on the FPGA in the PandA II. Flow diagrams of the architecture with and without the high performance controller are presented in Fig. 2.

* gareth.nisbet@diamond.ac.uk

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR010

# BEAM-SYNCHRONIZED ACQUISITION AND ENHANCEMENTS TO ASSOCIATED SERVICES*

K. Kim†, M. Weaver, M. P. Donadio, E. L. Williams Jr.
SLAC National Accelerator Laboratory, Menlo Park, CA, United States

*Abstract*

The Linac Coherent Light Source (LCLS) has developed a pulse-by-pulse data acquisition system, Beam-Synchronized Acquisition (BSA). BSA evolved from a 360 Hz software-based system to a 1 MHz firmware-based architecture tightly integrated with the LCLS timing system and beam rate. In parallel, the EPICS control platform transitioned from Channel Access (CA) to PV Access (PVA), enabling BSA to meet modern acquisition requirements—particularly for high-rate, high-volume applications requiring timestamping, pulse ID tagging, and precise cross-system alignment across the facility. BSA includes a fault buffer mechanism for each monitored variable, with four rotating buffers per variable, each capable of storing one million samples. One buffer continuously collects data at beam rate (1 MHz) in a ring configuration, while the others remain on standby. When the Machine Protection System (MPS) detects a fault, the active buffer is instantly frozen, and a standby buffer takes over, preserving a one-second snapshot of data. This snapshot is synchronized across the facility and available for all BSA variables system-wide. This paper presents the architecture, firmware and software components, and supporting services developed to meet the demanding requirements of SC operation, enabling machine learning and real-time feedback capabilities.

## INTRODUCTION

### *Timing System for Normal-Conducting (NC) Linac*

The original NC timing system for the LCLS facility was based on an AC power-line synchronized trigger [1]. This approach was chosen because the legacy high-power RF modulators required pulses to occur on the same AC phase to minimize jitter. The system provided 360 Hz beam-pulse aligned timestamps, with the lower bits of the nanosecond portion of the timestamp encoding the AC time-slot information. An additional 18 bits of the timestamp were reserved for the beam pulse ID, enabling reliable pulse identification.

The timing system delivered these beam-pulse aligned timestamps to EPICS process variables (PVs) through the general-time facility, which was registered as an event provider. As a result, PVs such as beam position, RF phase, RF amplitude, and other diagnostics carried a consistent pulse-aligned timestamp that reflected the beam pulse timing.

### *Introduction to Beam-Synchronized Acquisition (BSA)*

Although EPICS PVs could provide beam-pulse aligned timestamps, aligning data sets across multiple EPICS IOCs and subsystems for each beam pulse remained a significant challenge. To address this, the Beam-Synchronized Acquisition (BSA) system was developed as a common service for all subsystems connected to the timing system.

BSA introduced several key capabilities:

- **Scalable acquisition:** Support for up to 20 simultaneous independent acquisitions across IOCs for the entire NC facility.
- **Flexible operation:** Acquisition could be run on-demand or continuously.
- **Configurable acquisition:** Users can specify start/stop conditions, rates, and number of entries. A key feature is decimation, where multiple beam-synchronized samples are grouped into one acquisition entry. Each entry includes the average value as well as standard deviation, minimum, and maximum, reducing data volume while preserving statistical information about beam-pulse variability.
- **Waveform PV support:** Each acquired data value may be published as a waveform PV with up to 1800 entries. In continuous mode, waveform PVs act as ring buffers—new data overwrites the oldest entries.
- **Fault tolerance with NaN insertion:** If a source fails to provide data for a beam pulse, BSA inserts NaN (Not a Number) markers, ensuring downstream analyses can detect missing data.
- **Software integration:** BSA functions were implemented in software within the timing module and deployed across all EPICS IOCs in the facility.

Through these features, BSA transformed beam-pulse aligned timestamps into a consistent facility-wide acquisition service, laying the foundation for high-precision synchronized analysis of accelerator operations.

## IMPROVEMENT OF BSA SERVICE IN THE SUPERCONDUCTING (SC) SYSTEM

### *SC Timing and Hardware Platform*

As part of the LCLS upgrade, a superconducting linac was introduced with a beam rate of up to 1 MHz. This required a

* Work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Contract NO. DE-AC02-76SF00515.
† khkim@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR011

# BUILDING THE BACKBONE: CABLE PLANT PLANNING, DESIGN, AND PROGRESS FOR LCLS*

M. Cabral†, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

The cable plant is the backbone of the Linac Coherent Light Source (LCLS) Control System as it provides essential power, data, and control connections to enable all experiments and upgrades. Over the past years, LCLS has experienced some challenges with managing cable plant best practices and looks to refine its approach to cable plant design and management. This paper reviews current cable plant status, identifies ongoing challenges and outlines future strategies to improve maintainability and efficiency of the cable plant.

## INTRODUCTION

The nervous system is a combination of tissues and cells, essential for human control and coordination while the backbone provides necessary support and structure. Similarly, the cable plant and tray systems form a facility's nervous system and backbone – delivering and protecting power, data, and control signals to experiment areas and enabling science.

The LCLS cable plant, designed and installed over 10 years ago, proved to be a robust system due to its centralized planning process and a well-maintained database. However, as new projects have been added and ownership has become decentralized, challenges have emerged which include overfilled trays, inconsistent cable routing, and diverging database systems.

This paper reviews the current state of the LCLS cable plant and practices, highlight key challenges, and look forward to future improvements that may support LCLS design, install, and maintain the cable plant for upcoming large-scale installations.

## CURRENT LCLS CABLE PLANT

The LCLS cable plant is comprised of cable trays and racks layouts that were specified for the LCLS-I design. Designed with foresight, it has supported LCLS through multiple installations and demolitions while maintaining sufficient spare capacity for future endeavors. Paired with a well maintained cable database, CAPTAR, and strong installation team, it was a recipe for success and has served as a reliable foundation for a decade.

However, with the backbone complete, the cable plant planning slowly moved away from a centralized system to a distributed responsibility. Stakeholders began routing and installing their own cables while the cable database split into two – CAPTAR and neo-CAPTAR. Although this split and dispersal of responsibility gave groups more flexibility and autonomy, several challenges arose such as tray capacity, installation inconsistencies, thinned resources, and differing QA standards leading to difficulties maintaining tray fill, maintaining cable separation requirements, and unclear authority over existing and future cable routes.

## CURRENT CHALLENGES

Some of the challenges with the present-day state of the cable plant include: maintainability as well as excess labor and materials required for planning, installation, and contingency.

### *Maintainability*

The long-term reliability and efficiency of a cable plant depends heavily on its maintainability. As systems evolve, routine upgrades and reconfigurations become increasingly complex when infrastructure have not been designed or managed with life-cycle considerations in mind. Key challenges currently impacting maintainability include:

- **Tray Overfill:** Many trays are congested and at or beyond recommended tray fill capacities (50 percent of cross-sectional area for typical IC [Instrument and Control] ventilated trays [1]), complicating both new installations and demolition work.
- **Cable Management:** Poor organization, tangled runs, and dead cables increase labor time during installations and reconfigurations.
- **Code Compliance:** Inconsistent routing standards of HV (High Voltage), LV (Low Voltage), DC (Direct Current), and IC cables have led to issues with maintaining proper cable segregation in congested areas. This issue has triggered improvement projects to reorganize cabling and install proper dividers in existing installations and prompted revised planning efforts for future installations.

### *Planning*

Effective planning is essential to ensure that cable plant installations remain scalable, conflict-free, and manageable for installation teams. In practice, however, decentralized workflows have introduced some inefficiencies that hinder coordination and long-term reliability. Current planning challenges include:

- **Decentralized Ownership:** Without a central authority, conflicts between routes, priorities, and QA standards often emerge.

* This work was supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-76SF00515.

† mcabral@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR012

# NOVEL DISTRIBUTED FAST CONTROLS ARCHITECTURE FOR THE CONSOLIDATION OF CERN'S PS KICKERS

L. Strobino*, Ş.F. Er, N. Magnin, P. Van Trappen, N. Voumard, CERN, Geneva, Switzerland

*Abstract*

The control of fast pulsed magnet systems at CERN requires often a common set of fast digital electronics subsystems to perform tight timing control and fast protection of high-voltage pulse generators. Although the generators architecture is mainly modular, these control systems were until now mostly centralized, i.e., several generators per equipment, but one global and equipment-specific control system.

With the upcoming consolidation of CERN's PS kicker magnets controls, a new distributed architecture is proposed. Instead of one global control crate per functionality (timing, fast protection, acquisition, etc.), this new approach incorporates one control crate per generator, merging several functionalities together. That crate becomes more generic, offering higher flexibility in terms of system size (number of generators or magnets). It also allows to reduce the cabling costs but comes with new challenges in terms of complexity, data transmission bandwidth, and software latency.

This paper presents the new Distributed Kicker Fast Controls (DKFC) solution based on CERN ATS Distributed I/O Tier (DI/OT) ecosystem, including new Open Hardware electronic boards (ADCs, DACs, I/Os, dry contacts, etc.) and gateware architecture with high-speed board-to-board data exchange. Advantages and drawbacks of this new architecture and possible future extensions are also discussed.

## INTRODUCTION

The choice for a generic, modular, and distributed architecture for an ABT (Accelerator Beam Transfer, i.e. the author's group) control system has many advantages, such as increased flexibility and reduced cabling costs, while its main disadvantage of increased complexity should be manageable. In addition, the availability of the Distributed I/O Tier (DI/OT) motivated this design choice even more: the project goals match the ABT ones and since it was in an early design phase, lacking features could be identified in time. Due to the nature of an open, common effort, any user contribution helps the project forward, and although initially at a higher personnel and financial cost, it's worth the investment thanks to a decrease in technical debt and assured long-term support at the level of hardware, gateware, and software.

## DISTRIBUTED I/O TIER (DI/OT)

The Distributed I/O Tier project was initiated at the CERN ATS (Accelerator & Technology Sector) with the goal to standardize the distributed I/O layer of the control system (the closest layer to the accelerator equipment), dominated until now by custom-made electronics in many form factors designed by different teams [1]. The aim of this standardization is to increase sharing and reuse of design efforts between equipment groups [2].

The DI/OT ecosystem comprises a hardware kit capable of accommodating both radiation-exposed and radiation-free environments, based on the same custom 3U CPCI-S (CompactPCI Serial) crate and passive backplane. The non-radiation-hardened part of the kit consists of two redundant off-the-shelf CPCI-S 300 W power supplies, a 1U fan tray, and a set of pluggable front and rear boards.




Figure 1: DI/OT crate (middle) with fan tray (bottom) and FIDS patch panel (top).

The crate, shown in Fig. 1, hosts one System Board (SB) in the leftmost front slot (slot #0) and up to 8 Peripheral Boards (PBs) in the adjacent slots (slots #1 to #8). Each peripheral slot possesses a direct connection to a corresponding slot at the back of the crate, which accepts an optional 160 mm deep Rear Transition Module (RTM). The backplane is electrically compliant with the CPCI-S standard, but offers wider (6 HP) and deeper (220 mm) front slots for additional component space and improved cooling. It provides a star topology of 16 LVDS lanes (differential pairs that can also be used single-ended), one LVDS clock, and two high-speed differential pairs for Multi-Gigabit Transceivers (MGT) per peripheral slot, in addition to presence detection signals, a shared $I^2C$ bus, 5 multi-drop lines between the system slot and the peripheral slots, and 8 multi-drop lines between the peripheral slots and the RTM slots.

The base DI/OT kit components were designed by CERN's Controls Electronics and Mechatronics (CEM) group and are centrally provided to CERN users, i.e. the equipment groups. However, to fully replace ABT existing hardware, additional functionalities were needed, such as fast comparators, high-resolution low-bandwidth ADC and DAC channels, interfaces to industrial control systems

* lea.strobino@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR013

# AUTOMATIC BUNCH TARGETING FOR SINGLE-BUNCH BEAM POSITION MONITORS

R. Hong*, W. Cheng, A. Brill, G. Shen, Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Single-bunch beam position monitors (BPM) are used to track the trajectory and measure turn-by-turn positions of a single bunch of electrons while multiple bunches are present in the storage ring. In the Advanced Photon Source Upgrade (APS-U), 20 single-bunch BPMs have been installed in the first three sections right after the injection point. For each BPMs, an RF switch is used to select the signal of the target bunch to the BPM electronics. The timing signals to the BPM electronics and the RF switch are provided by the Fast Event System at APS-U, a global event-based trigger distribution system based on the hardware components developed by Micro-research Finland (MRF). The requirement of the timing signal to the RF switch is stringent to reliably select the target bunch. The GTX output function of the MRF event receiver can fine-tune the delay of the output signal to achieve the desired timing resolution for the RF switches. In this presentation, the hardware configuration of the timing signals and the software developed for automatic bunch targeting are described. The performance of the system and example applications are also discussed.

## INTRODUCTION

The Advanced Photon Source Upgrade (APS-U) [1] project includes the installation of new beam position monitors (BPMs) to improve the measurement of beam parameters and the reconstruction of beam trajectories. Following the upgrade, the APS storage ring has a significantly lower emittance. A swap-out injection scheme is used due to the storage ring's small dynamic aperture. To better understand the trajectory of the injected bunch, 20 single-bunch BPMs [2] have been installed along the beam pipe, starting from the injection point and covering approximately 7.5% of the storage ring. These single-bunch BPMs share the same pick-up buttons with the BPMs used for the stored beam.

To isolate the signal of the injected bunch from other stored bunches, a fast radio-frequency (RF) switch [3] was developed. This switch selects any specific bunch and directs its signal to either the Libera Spark electronics [4] for the single-bunch BPMs or the Libera Brilliance+ (LB+) electronics for the stored-beam BPMs. Normally, the RF switch directs the signal to the LB+ electronics. When provided with a gate signal in NIM logic, the switch directs the beam signal to the Libera Spark electronics. During normal APS operations, the gate width is set to 10 ns, as the minimum time difference between bunches in the APS storage ring is 14.72 ns.

The Libera Spark electronics require both a machine revolution clock and a trigger signal. The machine revolution clock period is the time it takes for a beam bunch to travel through the entire storage ring. This clock is used as a reference to phase-lock the local machine clock within the Spark electronics, and the beam signal is sampled at the local machine clock frequency. The Spark electronics can process a configurable number of turns before and after the trigger and saves the data to a buffer. The trigger signal also serves a synchronization purpose. The first trigger received after the Spark electronics is armed defines the beginning of a turn's digitization window. Since the beam signal may arrive at different BPM electronics at different times, the trigger signals are delayed so that signals from the same bunch appear at the same position within the digitization windows of all BPM electronics. This trigger delay can be set within the Spark electronics or by the timing system that provides the signal.

## FAST EVENT SYSTEM CONFIGURATION

The fast event system at the APS is an event-based timing system that distributes synchronized clock signals and triggers throughout the storage ring using the VME-300 series modules manufactured by Micro-research Finland [5] (MRF). Event codes are generated by event master (EVM) modules VME-EVM-300 in the main control room (MCR) and distributed via optical event links to 22 local timing stations around the storage ring. The events are transmitted at the reference clock frequency. At the APS, the event clock frequency is 117.352 MHz, which is one-third of the storage-ring RF frequency.

Within the local timing stations, an EVM module receives the event link from the MCR and distributes the event links to event receiver (EVR) modules VME-EVR-300 and devices with an embedded EVR, such as the LB+ electronics. The EVR module can generate output triggers in response to incoming events. The logic for generating an output is customizable. Each EVR module can provide up to 24 universal output channels with configurable signal types, along with two CML output channels fixed on the front panel. MRF provides daughterboards for the universal output channels to support different signal types. An EVR can only generate delays and pulse widths in steps of the event clock period. For regular output channels, the generated pulse is output directly. The last two front-panel universal output channels and the CML channels operates in GTX mode [6], where the EVR's output pulse subsequently generates a waveform that can be configured at a clock rate 40 times higher than the event clock frequency. This allows for finer delay adjustments and pulse widths smaller than the event clock period. Some universal output daughterboards can insert an additional delay adjustable in 10-ps steps, but they cannot adjust

* rhong@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR015

# MODERNIZATION OF PLC-BASED CONTROL SYSTEMS AT SNS*

I. Beaushaw†
Oak Ridge National Laboratory, Oak Ridge, TN, USA

*Abstract*

When the SNS site was built around 20 years ago, the Conventional Facilities (CF) control systems were designed using 2 communication protocols to allow programmable logic controllers (PLCs) to interface with motors, variable frequency drives (VFDs), and remote input and output (I/O) devices. The protocol chosen to control motors and VFDs is DeviceNet, a CANbus-based protocol developed by Allen-Bradley, a subsidiary of Rockwell Automation. The protocol chosen to communicate with remote I/O is ControlNet, another protocol developed by Allen-Bradley. Both of these protocols are obsolete and present reliability and maintainability issues, particularly DeviceNet. As the Control Systems Section at SNS is working to modernize control systems throughout the machine, a major goal for PLC-based systems is to remove the obsolete communication protocols in favor of standard, ubiquitous Ethernet. To this end, any new VFDs installed use Ethernet communication. Many VFDs are currently being replaced in the Central Utilities Building (CUB) and the Central Exhaust Facility (CEF) and are being removed from DeviceNet in favor of Ethernet communication. Planning is underway to retrofit Eaton Intelligent Technology motor control centers (MCCs) in the Target Building to remove particularly troublesome DeviceNet adaptors and replace them with Ethernet adaptors for each motor starter. The ControlNet network in the CUB has been demolished, with I/O drops integrated into a local Ethernet network, improving sustainability and maintainability.

## BACKGROUND

The greatest vulnerabilities in Conventional Facilities (CF) control systems at Spallation Neutron Source (SNS) currently lie with the communication protocols in use. Both DeviceNet[1] and ControlNet[2] were part of the original designs and are obsolete. DeviceNet networks were installed to facilitate programmable logic controllers (PLCs) communicating with variable frequency drives (VFDs) and motor starters in CF control systems. DeviceNet has, however, proven itself to be unreliable, exhibiting issues that can often bring down large portions of the network or even the whole network. It is sensitive to network changes and physical connections are often broken over time, possibly due to heating and cooling cycles in motor control centers (MCCs). These issues often take more time to correct than other control system issues because of the caution and coordination needed to troubleshoot within live MCCs. ControlNet was implemented to allow PLCs to communicate with remote input and output (I/O) hardware. It is deterministic, meaning communication with devices happens at established times for established time intervals. This factors into scheduling the ControlNet network, but most of the work is taken care of by the software (RSNetworx for ControlNet). One disadvantage of ControlNet is that it is not ideal for future I/O additions, which are not always predictable. If any new I/O hardware is added to the network, the network must be rescheduled, which requires putting the PLC into program mode, stopping the process. While this may not endanger run time, some processors require great coordination before putting them in program mode, making this unideal. In addition to this, some of the ControlNet hardware used in CF control systems has been discontinued, while other critical hardware will also be discontinued in the near future. Because of the vulnerable state these protocols put control systems in, whether from a reliability or maintainability standpoint, the Controls Systems Section at SNS has been working toward replacing them with Ethernet/IP[3] (EIP). EIP is a superior communication protocol for use with Allen-Bradley PLCs and communicates via the Ethernet medium.

## LOCAL ETHERNET NETWORKS

One step that has been crucial in the change to using EIP is setting up local Ethernet networks for each PLC that is modernized (Fig. 1). The local Ethernet network replaces the functionality of both DeviceNet and ControlNet networks simultaneously. Another benefit of setting up a local Ethernet network is to restrict access to each device or I/O drop to the PLC only. It also cuts down on unnecessary traffic on the Integrated Control System (ICS) network.

## VFD REPLACEMENTS

A major ongoing project in the Central Utilities Building (CUB) at SNS is the replacement of 16 VFDs that support Chilled Water and Tower Water supplies across the site. These VFDs run pumps for these systems as well as fans for cooling towers. An effort to replace these VFDs due to obsolescence has provided the opportunity to remove them from the DeviceNet network in the CUB and add the new VFDs to the local Ethernet network. The existing VFDs were ABB drives, while the replacements are Danfoss (Fig. 2).

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under contract number DE-AC05-00OR22725.
† beaushawiz@ornl.gov
Notice: This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (https://www.energy.gov/doe-public-access-plan).

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR016

# RTST VACUUM CONTROL SYSTEM DESIGN*

J. Yan†, D. Humphreys, A. Chaires, L. Boyd
Oak Ridge National Laboratory, Oak Ridge, TN, USA

*Abstract*

This paper presents the design and implementation of the Integrated Control Systems (ICS) Vacuum Control System for the Second Target Station (STS) within the Ring-to-Second Target Beam Transport Line (RTST). The RTST vacuum system is crucial for maintaining a high-vacuum environment necessary for the operation of a high-intensity proton beamline, extending from the existing Ring to Target Beam Transfer (RTBT) to the new STS. The system is composed of various components, including vacuum assemblies, sensors, pumps, and an architecture based on established control systems utilizing Allen-Bradley Programmable Logic Controllers (PLCs) coupled with EPICS Soft Input/Output Controllers (IOCs). The design emphasizes reliability and safety, incorporating sector gate valves for isolation, remote controls for turbomolecular and ion pumps, and pressure monitoring through advanced gauge systems. This paper details the architectural framework, instrumentation, control layers, and operational interfaces to ensure robust management of the vacuum conditions necessary for the successful operation of the RTST.

## INTRODUCTION

The RTST is a 230-meter-long transportation beamline to transport the proton beam from the existing RTBT beamline of the Spallation Neutron Source (SNS) to the Second Target Station (STS). It extends the RTBT vacuum infrastructure to the STS and requires a robust high-vacuum system to ensure reliable operation under high-intensity beam conditions. The RTST is divided into seven isolated sections, separated by seven sector gate valves. These sections are grouped into three engineering focus areas: Extraction Region, Straights and Bends, and the Final Stretch. The Extraction Region interfaces with the RTBT vacuum system, while the Final Stretch connects the beamline to the Second Target Proton Beam Tube assembly. The required operational vacuum level for the RTST is below 1x10$^{-7}$ Torr, with a leak rate of less than $10^{-9}\frac{Torr*L}{s}$. To achieve this stringent vacuum level, a reliable vacuum system has been designed, incorporating critical components such as vacuum assemblies, sensors, pumps, and associated instrumentation. Additionally, a dedicated vacuum control system has been developed to effectively regulate and monitor the vacuum equipment throughout the RTST.

* This manuscript has been authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy. (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-accessplan).
† yanj1@ornl.gov

As a subsystem of the Accelerator System controls [1], the preliminary design of the RTST Vacuum system has been successfully completed, and efforts are on track for advancing to the final design phase.

## SYSTEM ARCHITECTURE

The vacuum control architecture is modelled on the well-established RTBT vacuum control system [2], which utilizes Allen-Bradley(A-B) Programmable Logic Controllers (PLCs) integrated with EPICS Soft IOCs. The system follows the modern three-layer architecture of the SNS distributed control system [3]: field device layer, control layer, and supervisory layer. Figure 1 illustrates the architecture of the RTST vacuum control system. The field device layer includes vacuum instrumentation, such as vacuum valves, gauges, pumps, residual gas analyzers (RGA), and their respective controllers. The valves, gauges, and pumps are installed on the beam pipe in the RTST tunnel, while the controllers for these devices are located in the RTST Central Utility and Service Building (CSB) Power Supply Room. The control layer is built on A-B ControlLogix PLCs, which monitor and control the vacuum devices. This layer also includes Ethernet-Serial port servers for communication and an EPICS SoftIOC to manage multiple subsystems. The supervision layer provides the machine operator interface in the Central Control Room (CCR) and hosts data servers in the Central Equipment Room (CER).

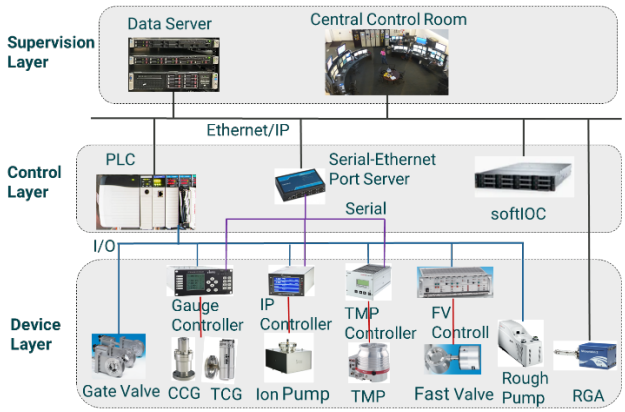



Figure 1: STS RTST Vacuum Control Architecture.

### *Device Controls*

The Vacuum system devices include valves, gauges, ion pumps, turbomolecular pumps, rough pump, and RGAs. These devices communicate with the PLC via hardwired I/O and with the EPICS SoftIOC over serial or Ethernet links.

**Sector Gate Valve Control** VAT gate valves are used to isolate specific sections of the RTST vacuum system, enabling flexibility for maintenance and repairs. Each

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUMR019

# PROOF OF CONCEPT OF A PLC BASED EMITTANCE METER FOR THE NEWGAIN PROJECT*

C.Hocini†, C.Haquin, C.Jamet, P.Salou,
Grand Accélérateur Nat. d'Ions Lourds, Caen, France
T.Adam, E.Bouquerel, P.Graehling, C.Maazouzi,
Institut Pluridisciplinaire Hubert Curien, Strasbourg, France

*Abstract*

The SPIRAL2 accelerator features several diagnostic devices used to characterize, adjust, and monitor the beam. As part of its NEWGAIN (New GANIL Injector) upgrade project, SPIRAL2 will be equipped with a new source and a new injector. Therefore, new diagnostic tools will be developed, including an ALLISON-type emittance meter. To manage the emittance meter, we opted for a modern industrial PLC solution, leveraging its expanding capabilities and established maintenance advantages over a traditional PLC/VME combination. This paper details the architecture and programming concepts of our proof-of-concept system. It further outlines the test campaign conducted to validate the PLC's capabilities in several key areas: controlling the motors for measurement head positioning within the beam, managing high-voltage ramps, acquiring experimental data, and communicating results to the EPICS control system. The paper will also discuss findings related to current measurement accuracy, measurement rate, and synchronization, as well as the repeatability of the overall measurement process.

## PROJECT NEWGAIN

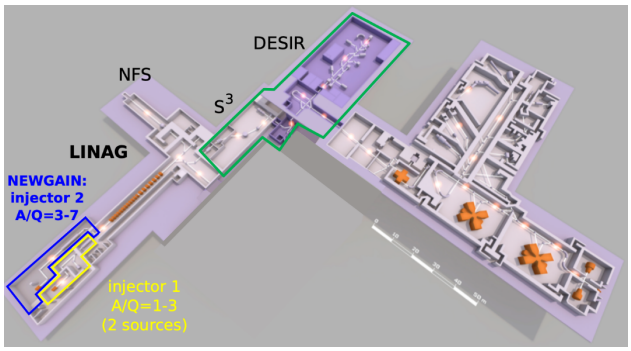



Figure 1: GANIL and SPIRAL2 facility.

The Newgain project involves the installation of a new superconducting source capable of producing high-intensity heavy ion beams, with an A/q ratio ranging from 3 to 7. This project also includes the addition of a new RFQ (Radio Frequency Quadrupole). The entire system will be integrated into the SPIRAL2 facility (Fig. 1), which currently includes two experimental areas: NFS (Neutrons For Science) and $S^3$ (Super Separator Spectrometer), as well as a new facility under construction, DESIR (Decay, Excitation and Storage of Radioactive Ions) [1].

## FUNDAMENTALS OF EMITTANCE MEASUREMENT

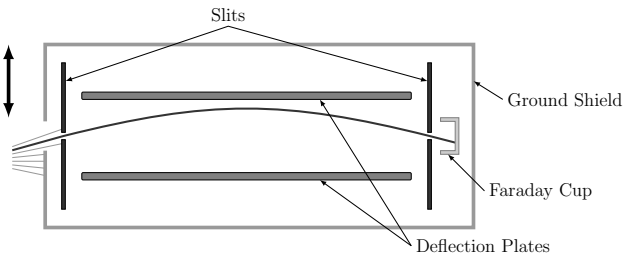



Figure 2: Allison method measurement.

The emittance meter enables the measurement of beam emittance based on the ALLISON principle [2]. Transverse emittance provides a mapping of the transverse angles of beam particles as a function of their positions. It characterizes the divergent or convergent nature of the beam and is thus used to determine the beam focus, study particle propagation, or tune the accelerator system.In this method, the beam is scanned using an analysis head (Fig. 2), which consists of a set of electrostatic deflection plates and a selection slit. By sweeping the voltage applied to the plates, the angular dispersion of the beam can be measured. This technique allows for the emittance profile to be measured in a single transverse phase space plane (either $xx'$ or $yy'$, where $x' = v_x/v_z$ or $y' = v_y/v_z$). For this reason, a complete transverse emittance characterization typically requires the use of two orthogonally positioned analysis heads.

## OVERVIEW OF THE SYSTEM ARCHITECTURE

Typically, a VME-based system is used for acquiring current measurements from the Faraday cup, in combination with a PLC that handles motor control. The aim of this proof of concept is to consolidate all functionalities within a single control unit, the PLC (Fig. 3). This unit will be responsible for driving the servomotors, managing the high-voltage power supply for the beam scanning, acquiring real-time current measurements using Picolin-type electronics [3], and transmitting the data to the EPICS control system.

* NewGain project
† clement.hocini@ganil.fr

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD001

# MODERNIZING THE RF CONTROL SYSTEMS AT THE ADVANCED PHOTON SOURCE

N. DiMonte, Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Advanced Photon Source (APS) recently completed a significant upgrade to its storage ring, replacing all existing components with new ones. However, the RF systems and their control mechanisms were not part of this upgrade. The APS operates four RF systems: the Linac, the Particle Accumulator Ring (PAR), the Booster Ring, and the Storage Ring. These systems have traditionally relied on analog controls utilizing a combination of VME/VXI and Allen-Bradley PLC5 hardware, which are now obsolete. This paper discusses the ongoing transition to digital controls, the integration of new RF hardware, and the progress achieved thus far.

## INTRODUCTION

The APS operates four RF systems: the Linac, the Particle Accumulator Ring (PAR), the Booster Ring, and the Storage Ring. The legacy RF control systems were designed in the early 1990s using a mix of commercial and custom hardware: VME/VXI instrumentation for acquisition and timing, Allen–Bradley (A&B) PLC-5 for interlocks and discrete I/O, and VxWorks-based VME CPUs interfaced to EPICS [1]. After three decades of service, much of this hardware—and several custom modules—are obsolete. In parallel, the APS is transitioning its high-power RF sources from klystrons to SSAs due to the declining availability of 1 MW klystrons.

This paper summarizes the subsystem-by-subsystem modernization: migration to soft IOCs, replacement of analog LLRF with digital LLRF, installation of a commercial timing system, and adoption of ethernet-based device communications (e.g., Modbus/TCP). We also discuss operational implications and the current status.

## LINAC RF CONTROLS

The Linac RF control system was commissioned in 1992 with a design philosophy to use as many commercial off-the-shelf components as possible.

The main elements included:

- Allen–Bradley PLC-5 hardware for modular and reliable discrete I/O and interlocks. An A&B VME module provided a convenient interface to the PLC I/O space.
- VMEbus for general-purpose control and compute. Motorola 68k VME CPUs (e.g., MVME172) running the VxWorks RTOS hosted EPICS IOCs.
- VXIbus for analog data acquisition with specific timing and triggering. This included envelope and phase detector modules (originally developed at LANL) for LLRF monitoring.
- The APS Event Timing system (custom), which distributed timing events used by Linac-specific VME timing modules to generate triggers.



Figure 1: ScandiNova K400.

Although this architecture remains functional, it is obsolete and difficult to maintain. The Linac must support operations between 400 MeV to 500 MeV and is undergoing comprehensive refurbishment.

The key upgrades are:

- Modulators: Replacement of legacy PFN modulators with ScandiNova K400 solid-state pulsed power modulators [2] (Fig. 1). The K400 interfaces via ethernet (Modbus/TCP) and requires only a soft IOC, eliminating VME and PLC hardware previously needed to integrate PFNs into EPICS.
- LLRF: Replacement of the analog LLRF (VXI-based envelope/phase acquisition and A&B control) with a digital LLRF system (Libera LLRF from Instrumentation Technologies [3]). The Libera platform includes an embedded soft IOC and has been customized for APS Linac needs.
- Timing: Replacement of the APS custom event system with the MRF timing system [4], which maintains compatibility with APS timing concepts while using a supported, commercial platform.

Further details on the Linac upgrade are documented in Ref. [5].

## PAR RF CONTROLS

The Particle Accumulator Ring (PAR) employs a control architecture similar to the Linac (VME/VXI, A&B PLC-5, and APS Event Timing). The legacy LLRF is entirely analog, built from APS-developed VXI modules (envelope detectors, phase detectors, phase shifters, modulators, and feedback modules), all now obsolete.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD002

# APS-U STORAGE RING POWER SUPPLY INTERLOCK AND TEMPERATURE MONITORING CONTROL SYSTEM*

S. Xu†, Y. Ruan
Advanced Photon Source, Argonne National Laboratory, Lemont, IL 60439, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The APS Upgrade (APS-U) storage ring power supply interlock and temperature monitoring system plays an important role for personnel and equipment safety. The system utilizes the Experimental Physics and Industrial Control System (EPICS) input/output controllers (IOCs) to interface with remote programmable logic controllers (PLCs) for controlling the storage ring power supply interlock system and monitoring various equipment temperatures. This paper will present how the system is implemented and operating successfully.

## INTRODUCTION

The APS-U storage ring [1] comprises 40 sectors organized into six zones. Each sector includes two longitudinal-bending (L-bend) dipole magnets (M1 and M2) [2], and each double-sector is powered by four raw power supplies. To protect these assets and support safe operations, the APS-U implements a power supply interlock and temperature monitoring control system based on the EPICS framework. PLCs provide deterministic interlock execution and sensor acquisition; EPICS soft IOCs furnish supervisory control, centralized monitoring, configuration management, and operator user interfaces.

To prevent magnet overheating, each dipole features a thermal switch circuit monitored by the PLC. An open-circuit condition triggers the corresponding power supply interlock, removing output to the affected magnet. Additional PLC-monitored conditions include deionized (DI) water flow and pressure, magnet power emergency-off chains, maze door status, and various device temperatures.

Temperature monitoring across the ring uses a combination of K-type thermocouples (TCs) and resistance temperature detectors (RTDs). K-type TCs monitor magnet coils and other devices—including beam position monitors (BPMs), vacuum absorbers, RF-lined gate valves, and warm vacuum chambers in 31 insertion device (ID) sectors—while RTDs capture girder and ambient tunnel temperatures. Additional TCs cover Zone F specialty systems, including transverse and longitudinal feedback (TFB/LFB), the decoherence kicker (DK), DC current transformers (DCCTs), and fast injection/extraction pulsed devices (IK1–IK3).

## SYSTEM ARCHITECTURE OVERVIEW

- Front end: Allen-Bradley ControlLogix PLCs aggregate digital and analog signals and execute interlock logic. External relays provide fail-safe, dry-contact trip outputs to power supplies.
- Supervisory layer: One EPICS IOC per zone communicates with the zone PLC controller via EtherNet/IP, exposes EPICS process variables (PVs), provides health/status, logs operator actions, and retains calibration parameters.
- Networking: Within each zone, PLC interlock relay racks are connected in a device-level ring topology for high resilience. A separate ring connects the six zone networks to facilitate cross-zone data exchange. EPICS communications use a dedicated EtherNet/IP interface that is not part of the interlock ring.
- Operator interface: MEDM screens provide interlock status, temperature trends, and diagnostics, organized by double-sector and zone.

## PLC HARDWARE AND NETWORKING

- Platform: Allen-Bradley ControlLogix 1756 series.
- Topology: Each double-sector hosts a PLC interlock relay rack. Racks within a zone are linked in a ring to improve reliability and resilience. A separate external ring interconnects the six zones for cross-zone interlock dependencies and shared status.
- Controller placement: The zone PLC controller resides at the first interlock relay rack in each zone and accesses remote modules in that zone. Interlock logic is processed in the controller; EPICS supervision uses a separate EtherNet/IP network interface.
- Communication modules: EtherNet/IP modules support both intra-zone ring communications and inter-zone data exchange.

Figure 1 shows the ring network for Zone A and the additional ring network that connects zones together. Other zones have a similar PLC ring network configuration as Zone A.

* Work supported by the U. S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357
† xusf@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD003

# CONTROL SYSTEM UPGRADE OF ARGONNE WAKEFIELD ACCELERATOR FACILITY*

W. Liu†, J. Power, G. Chen, C. Whiteford, A. Ody, E. Wisniewski, S. Doran, P. Piot, J. Hlavenka, R. Margraf-O'Neal, G. Shen, A. Johnson, D. Jarosz, S. G. Wang, L. Xiao

Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

As the Argonne Wakefield Accelerator (AWA) facility expanded, its original in-house control software could not keep up with the increasing complexity and scale of operations. To improve maintainability, reliability, and support future development, the AWA controls group undertook a major upgrade by adopting the Experimental Physics and Industrial Control System (EPICS). With support from the Advanced Photon Source (APS) Controls group, the AWA control system has been successfully upgraded from a centralized, difficult-to-maintain architecture into a flexible, scalable, and maintainable distributed system. This paper presents the status of AWA's new EPICS-based control system and describes the experiences and lessons learned during the upgrade.

## INTRODUCTION

The Argonne Wakefield Accelerator (AWA) is a premier electron accelerator with the world's highest bunch charge to carry out fundamental accelerator research with an emphasis on wakefield acceleration. It started with a simple electron beamline consisted of one L-band RF photocathode electron, one RF linac, some beam diagnostics and one RF station in the 80s. And developed into a small research and user facility with 3 relatively complex and flexible beamlines, 5 RF stations. Along with the evolving of the facility, the control and data acquisition system hardware of AWA also evolved. To accommodate the future expansion of the facility, the old inhouse control and data acquisition software architecture needs to be upgraded into a flexible, scalable and maintainable distributed system.

### *Why EPICS*

EPICS [1] is an open source framework for distributed control system used in most accelerator facilities around the globe. It makes sharing control algorithm and manpower much easier, which is crucial for collaborations.

As the result of the growth of AWA facility, the control system has evolved into a more or less distributed system which fit perfectly well with EPICS.

As a small research and user facility, AWA cannot afford to have its own team to cover every aspect of the control and data acquisition system. Collaboration and resource sharing are crucial for the future of AWA.

With the growing interest in the adoption of AI/ML in accelerator research and facility operations, the open-source ecosystem of EPICS made it the perfect choice for AWA control system upgrade so that AWA can better support its users and contribute to the accelerator community in AI/ML studies.

### *Path of the Upgrade*

AWA is a small but busy facility with about two experiments every month on average. We can not afford to rebuilt the control system with new hardware that has existing EPICS supports. It will not only require substantial amount of funding in purchasing new hardware, but also great amount of time and efforts in installation and commissioning. For the best interest of AWA facility, a phased approach had been chosen to minimize the impact of the control system upgrade on the facility schedules and our collaborators.

For phase 1, we focused on providing EPICS support on top of our existing control system by integrating EPICS IOCs into our control system programs. Doing so will allow AWA facility to continue its operation as usual while the control system upgrade. The goal of phase 1 is to provide EPICS interface of AWA core control system hardware.

For phase 2, we focused on creating EPICS OPI replicas of the old AWA control system GUI to smooth the transitions for the operators and phasing out the old control system GUIs. Then we focused on adding more EPICS OPIs to improve the operator experiences. Creating more EPICS IOCs for hardware which were typically operated with vendor provided programs in the old control system.

With the choice of these two phases, the control and data acquisition system EPICS upgrade of AWA is successfully completed without causing any interruption of the facility schedules.

## BEFORE THE UPGRADE

AWA control system started as a single CAMAC crate with few ADC, DAC and delay generator on the bus at the beginning. As the facility evolved, more control hardware were needed and we rebuilt AWA control and data acquisition system with PXI based hardware at the early 2000s. After the upgrade from CAMAC based system to PXI based system. After that, we just add more and more PXI based hardware to the system when possible to accommodate the expansion of the facility. As the complexity of beamline and facility evolves, we started adding hardware with different interfaces and it was getting harder and harder for our inhouse control system programs to keep it up.

* Work supported by the U.S. Department of Energy, Office of High Energy Physics

† wmliu@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD006

# TECHNICAL DESIGN CONCEPT AND DEVELOPMENT OF THE NEW ACCELERATOR CONTROL AND TIMING SYSTEMS FOR PETRA IV

T. Wilksen†, V. Andrei, J. Behrens, K. Brede, L. Fröhlich, J. Jäger,
H. Kay, H. Lippek, M. Pawelzik, J. Wilgen
Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

*Abstract*

At DESY, extensive technical design and prototyping work is currently underway to upgrade the PETRA III synchrotron light source to PETRA IV, a fourth-generation low-emittance machine. The realization of the new machine necessitates a redesign of the accelerator's control and timing systems. The primary hardware components will be based on the MTCA.4 standard, which has established itself as a reliable platform at DESY. The control system framework will be modernized to accommodate the demands of a fourth-generation light source. This paper presents the key decisions made in this context and provides an overview of the design and development process.

## INTRODUCTION

The PETRA IV project [1] will upgrade the existing 3rd generation synchrotron light source, PETRA III at DESY, into an ultra-low emittance 4th generation synchrotron light facility. This transformation will involve a complete renewal of the 2,304-meter-long accelerator ring and the entire injector chain. To accommodate 18 additional beamlines - bringing the total to 42 - half of the PETRA ring tunnel will be newly constructed. The PETRA IV storage ring will operate at an energy of 6 GeV with a 500 MHz RF frequency. The DESY II booster synchrotron will be upgraded to DESY IV to meet the requirements of a low emittance beam injected into PETRA IV. DESY IV will be complemented by an alternative injector chain with a linac based on the laser-plasma wake field acceleration principle (LPA).

The project is currently in the preparation and development phase, with initial construction scheduled to begin in 2027. PETRA III will be shut down at the end of 2029, followed by a 2.5-year dark period with no operations. Commissioning of the new PETRA IV facility is expected to begin in 2032, with the first users anticipated in 2033.

## CONCEPT

Rebuilding an accelerator almost from scratch, while already operating several large accelerator facilities on campus, naturally suggests the use of synergies across the various control system instances operated on campus at DESY. The aim is to consolidate the existing landscape of accelerator controls at DESY.

The new PETRA IV accelerator control system will be based on the control system for the European XFEL, which is a free-electron X-ray laser running in burst mode at 10 Hz. For PETRA IV, the front-end electronics will be completely redesigned and based on same, well-established MTCA.4 standard for high-speed data acquisition as used at the European XFEL, while industrial technologies such as PLCs will be used for slow controls. Since the European XFEL control system was designed using a similar architecture, existing solutions can and will be re-used for PETRA IV. Many of the low-level applications and implementations will remain suitable as-is, however, the high-level software connecting the middle layer of the control system with operations will need to be adapted to the specific demands of running a synchrotron.

Due to the entirely new front-end electronics based on MTCA.4, the necessity of overhauling the timing and RF synchronization systems is required for the PETRA IV accelerator. That includes not only the electronics but also the controls software and its architecture. Again, here the idea is to borrow as much as possible with respect to concepts and implementations from the existing European XFEL.

The following chapters will outline the current design and architecture of the timing and RF synchronization systems, as well as the overall accelerator control systems envisioned for the PETRA IV project.

## THE TIMING AND RF SYNCHRONIZATION SYSTEMS

### *RF Synchronization System*

The basic requirement of the RF synchronization system is to provide a continuous reference RF signal generated by a unique, stable main oscillator to drive local, low-noise oscillators with high precision. The fundamental RF frequency for the PETRA IV storage ring will be 499.664,3 MHz (hereafter referred to as 500 MHz), while the 3rd harmonic RF system requires an RF frequency of 1.498,992,9 GHz (hereafter referred to as 1.5 GHz), derived from the PETRA IV reference. The DESY IV booster will have the same frequency as PETRA IV, while the LPA injector chain needs a frequency of 2.997,998,580 Hz (hereafter referred to as 3 GHz). Both will be synchronous to the PETRA IV reference frequency. The RF synchronization system will consist of two RF synthesizers for PETRA IV and the DESY IV injector chain and a third one for the LPA injector chain. In normal operation, the synthesizers are coupled and generate a phase-synchronous frequency. The PETRA IV system will supply an additional 1.5 GHz reference for the 3rd harmonic system. The 500 MHz will be adjustable in the range of ± 1.65 kHz, and the 3rd harmonic frequency in the range of ± 4.95 kHz. This feature will be used for non-standard operations like the dispersion measurement.

† tim.wilksen@desy.de

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD007

# DESIGN AND PERFORMANCE OF A NOVEL DATA ACQUISITION AND PROCESSING SYSTEM FOR THE APS UPGRADE FRONT END XBPM*

B.X. Yang, S. Xu, H. Shang†, Argonne National Laboratory, Lemont, IL 60439, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The APS-Upgrade (APS-U) ID beamline front ends are equipped with next-generation X-ray beam position monitors (XBPMs). Each XBPM utilizes 16-element array detectors to simultaneously capture beam distribution information from undulator and bend magnet (BM) sources. A novel IOC has been designed to handle the following tasks in real time: estimating and subtracting BM signals, applying undulator-gap-dependent calibration to calculate undulator beam center positions, collecting beam position statistics, and estimating X-ray RMS motion and XBPM resolution. Initial data from early APS-U user runs demonstrate significant performance improvements. The BM signal contribution is reduced by more than a factor of 10, observed RMS beam motion is within the micrometer range, and XBPM resolution is consistently below 1 micrometer, typically less than 10% of the RMS beam motion.

## INTRODUCTION

The Advanced Photon Source (APS) has completed an upgrade recently [1]. This upgrade implemented two new types of insertion device front ends: a high heat load front end (HHLFE) for a dual inline undulator source, and a canted front end (CFE) for two canted undulators with a cant angle of 1 mrad [2]. In these front ends, two novel types of X-ray Beam Position Monitors (XBPMs) are used to monitor the beam stability, HHLFE-XBPM and CFE-XBPM.

In these XBPMs, the high-power undulator beams are intercepted by two water-cooled Glid-Cop absorbers in grazing incidence. The absorbers are located on the inboard and outboard side of the beam axis for HHLFE-XBPM and above and below the beam axis for CFE-XBPM. The beam footprint on each absorber is imaged on an array detector with 8 (2 × 4) scintillator-photodiode assemblies for reading out the intensity and spatial distribution of the x-ray footprint. The 16 outputs from the array detector elements are grouped into 12 current signals and acquired by three TetrAMM low-noise picoammeters from CAEN Els [3].

Figure 1 shows the effective zones of the array detector elements in HHLFE-XBPM. The undulator beam axis is at the center and two bend magnet (BM) radiation fans from the upstream and downstream sources are shown on the inboard and outboard sides. Note that in each APSU CFE-XBPM, the undulator beam coexists with three BM fans.

During APS user operations, the picoammeter current signals are digitized at 100Ks/sec and integrated over 0.1 second before sending the data to the IOC for real-time data processing. In the following sections, we will describe the data processing algorithm in detail and discuss how the calibration constants were chosen during commissioning of the APSU.

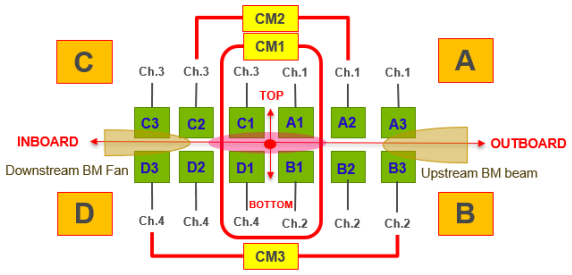



Figure 1: Concept of XBPM with array detectors.

## REALTIME XBPM DATA PROCESSING

The XBPM IOC performs the following tasks:

- Evaluate and subtract the BM background signals when undulator power is low with undulator gap above predefined values.
- Renormalize the undulator signal for correcting detector asymmetry.
- Calculate the horizontal and vertical undulator beam positions.
- Collect statistics of the undulator beam position: calculate filtered beam position data, and estimate the RMS beam motion and XBPM resolution of the unfiltered and filtered XBPM position outputs.

### *BM Background Subtraction*

The 12 detectors can be grouped into four quads: Q1 (A1, B1, C1 and D1) is nearest to the undulator axis and least sensitive to the BM radiation; Q2 (A2/B2/C2/D2) is the next neighbour, and Q3 (A3/B3/C3/D3) is the farthest from the undulator axis and farthest into the BM fans. Hence in the first approximation, we can estimate the BM signals in Q1 in the absence of undulator radiation as linear combinations of the Q3 signals,

$$\begin{pmatrix} A_1^{bm} \\ B_1^{bm} \\ C_1^{bm} \\ D_1^{bm} \end{pmatrix} = \begin{pmatrix} b_{11} & b_{12} & b_{13} & b_{14} \\ b_{21} & b_{22} & b_{23} & b_{24} \\ b_{31} & b_{32} & b_{33} & b_{34} \\ b_{41} & b_{42} & b_{43} & b_{44} \end{pmatrix} \begin{pmatrix} A_3 \\ B_3 \\ C_3 \\ D_3 \end{pmatrix}, \quad (1)$$

* Work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

† shang@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD008

# MOTION CONTROL SYSTEMS FOR INSERTION DEVICES IN DIAMOND-II

R. Mercado*, P. Amos, K. Bolt, A. Brine, L. Hudson, K. Jones A. Lyle,
B. Nutter, Z. Patel, A. Ramezani Moghaddam, G. Sharma, X. Tran, S. Tripathi
Diamond Light Source, Didcot, UK

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Diamond Light Source has been operating since 2007, and currently has installed 26 motion-controlled insertion devices that produce synchrotron light for the majority of the 36 beamlines in operation. The Diamond-II upgrade will reduce the emittance, increase the energy of the electron beam, increase the number of straights available, and includes the delivery of three flagship beamlines. As part of delivering Diamond-II we plan to build and procure 12 new insertion devices of which 10 will be motion-controlled using in-house designed and built control systems. We also plan to upgrade three control systems to manage obsolescence and enable software upgrades. This paper describes the various generations of motion control systems present, and outlines the upgrade plans, controls challenges, and special requirements.

## INTRODUCTION

Diamond Light Source is the UK's national synchrotron science facility, located at the Harwell Science and Innovation Campus in Oxfordshire [1]. The *Diamond-II* upgrade has been in development for several years and reached a milestone in 2023 when its funding was approved by the government. The upgrade project will modify the accelerator lattice to offer brighter beams and increase capacity using a 6-bend-achromat design that provides 24 new "mid-straight" sections [2].

Diamond has 29 insertion devices in operation [3, 4] as shown in Table 1. The column *phase* indicates a chronology of the various ID control system generations.

As an illustration of the capacity upgrade the tables below show 8 systems planned to be removed (entries not in bold in Table 1) and 12 are to be installed (Table 2).

Further discussion in this paper will ignore the I12 and I15 superconducting multi-pole wigglers (SCMPW) and the K11 fixed-gap wiggler that don't require motion control.

Table 1: IDs in Operation as of Aug. 2025

| No | Beamline | ID type | Phase | Notes |
|---|---|---|---|---|
| 1 | I02 | IVU | Phase I | a |
| 2 | I06-A | Apple-II | | |
| 3 | I11 | IVU | | |
| 4 | **I15** | SCMPW | | b |
| 5 | **I16** | IVU | | |
| 6 | **I18** | IVU | | |
| 7 | **I12** | SCMPW | Phase II | b |
| 8 | I06-B | Apple-II | | |
| 9 | I10-B | Apple-II | | |
| 10 | I10-A | Apple-II | | |
| 11 | **I07** | CPMU | | |
| 12 | **I22** | IVU | | |
| 13 | **I09** | IVU | | |
| 14 | **J09** | Apple-II | | |
| 15 | **I13** | IVU | | |
| 16 | **J13** | IVU | | |
| 17 | **I19** | IVU | | |
| 18 | **I20** | Wiggler | | |
| 19 | **K11** | Wiggler | Phase III | b |
| 20 | I05 | Apple-II | | |
| 21 | I08 | Apple-II | | |
| 22 | **I14** | IVU | | |
| 23 | **I21** | Apple-II | | |
| 24 | **I23** | IVU | | |
| 25 | **J02** | CPMU | Upgrade program | |
| 26 | **I03** | CPMU | | |
| 27 | **I04** | CPMU | | |
| 28 | **J04** | HPMU | | |
| 29 | **I24** | CPMU | | |

Notes: (a) upgraded controls, (b) no motion control, **bold** indicates ID staying for Diamond-II

## PHASE I IDS

The motion control system for these insertion devices consists of the following hardware [5]:

- VME based EPICS IOC.
- OMS58 motion controller card.
- Parker Hannifin Compax servo drives, motors, brakes and resolvers.
- TR Electronics linear absolute 0.1 µm resolution encoders and interfaces.
- Omron protection PLC.
- Applied Geomechanics tilt sensors.

The OMS58 motor controller cannot read absolute encoders, therefore the system uses a dedicated *Active PLC* module capable to read the SSI signals together with a mechanism to load the encoder positions before enabling motion.

Figure 1 shows an example Phase-I ID.

* ronaldo.mercado@diamond.ac.uk

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD009

# INTEGRATED DATA ACQUISITION AND PROCESSING PIPELINES FOR USERS AT ELETTRA 2.0: A CASE STUDY AT SYRMEP, THE μCT BEAMLINE

A. Contillo*, M. Belletti, R. Borghes, V. Chenda, A. Hafner, G. Kourousias, E. Longo, A. Olivo, M. Prašek, M. Turcinovich, Elettra-Sincrotrone Trieste S.C.p.A., Trieste, Italy

*Abstract*

Elettra, the Italian Synchrotron in Trieste, is about to undergo a major upgrade of the facility. To effectively exploit such improvement, data acquisition and processing are being integrated into single automated pipelines, subject to a facility-wide standard and yet flexible enough to accommodate the specific usage at each beamline. The entire procedure of data acquisition and processing spans a vast ecosystem of different structures and frameworks, which are being standardized across the whole facility: the GeCo control system, handling the safety and the operation of the beamlines; the TANGO Controls framework, which allows distributed control of the data acquisition process; the architecture of data storage and processing, whose accessibility is mediated by VUO, the Elettra unified portal; a modular adaptive processing infrastructure (MAPI) for analysis workflows; and an overview of the data lake data@Elettra. The intertwining of all these components results in the integrated pipeline experienced by the users. The acquisition/processing sequence presently in place at SYRMEP, the microtomography beamline, is presented as a case study of the standard structure that is being designed. Beamline and acquisition control, on-the-fly, and post-acquisition processing are described in the light of the general landscape proposed for Elettra 2.0, the upgraded facility.

## INTRODUCTION

The third generation Italian Synchrotron facility Elettra has been running since the first light in 1993, until the last shutdown in July 2025. During these decades, it has served the international scientific community by providing X-ray light to almost thirty beamlines, offering a variety of thechniques such as imaging, spectroscopy, diffraction, scattering, etc, for 160 thousand hours of operation. More than 11 thousand beamtimes were conducted during this period, resulting in almost 10 thousand peer-reviewed publications. Despite the constant upgrades and improvements, the facility has finally turned old, requiring a deep renovation in order to maintain its appeal to an international audience. Thus the upgrade to a fourth generation synchrotron light source, with highly enhanced brilliance and coherence. The first light is expected to be produced by the fall of 2026 and the first external users are expected by early 2027.

SYRMEP (SYnchrotron Radiation for MEdical Physics) [1] is the beamline of Elettra dedicated to microtomography, mainly devoted to life sciences research, but also to material science, geology and other neighboring fields. Tomography is one of the setups with the highest throughput, producing data in the order of ten megabytes per projection (the exact figure depending on the surface area of the detector). With the currently available fluxes, SYRMEP was able to acquire up to 100 projections per second, which brought the data acquisition rate to about 800 MB/s. Depending on the specific details of the beamtime (sample size, quickness of sample switching, etc.), this has allowed at times the acquisition of several terabytes per day. Following the upgrade, the photon flux available to the beamlines is expected to increase by three orders of magnitude. Therefore, even in the most conservative estimates (in which the limiting factor, rather than the signal integration time, is the dead time due to motor movements and/or data readout), the rate of data acquisition is expected to increase by at least a factor 10.

SYRMEP-LS (Life Sciences) is the upcoming embodiment of the imaging beamline in the upgraded facility. The sample-to-detector distance will be increased drastically, so as to make the best use of the increased beam coherence in a free space propagation-based phase contrast setup [2, 3], thus reducing the dose of X-rays needed to produce meaningful images. The surface area of the deployed detectors will also be enlarged, possibly even doubled. All things considered, the expected data throughput is expected to increase to 10-20 GB/s. Such rates pose challenges at several different levels: I) the data transmission infrastructure, II) the data storage capacity, III) the data access speed, IV) the data processing power. Each of these aspects must be tackled in such a way to avoid bottlenecks that would slow the whole pipeline of data acquisition and processing.

The Information Technology department of Elettra is enforcing a facility-wide standard for all components of the above-mentioned pipeline. The massive requirements of the beamline SYRMEP in terms of data throughput, storage, and processing make it an exemplary case study to describe these components in detail. In particular, the vacuum and beam transport elements are handled by the GeCo control system [4]. The control of the instrumentation and their coordination into data acquisition sequences are governed through the TANGO Controls [5] infrastructure of the beamline. The processing of raw data is handled by STP3 Web, a distributed software based on a Modular Adaptive Processing Infrastructure (MAPI) [6]. The access to stored data is mediated by the Virtual Unified Office (VUO) of Elettra, a multipurpose portal which (among many other things) handles the data access permissions via user authentication.

* adriano.contillo@elettra.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD010

# UPGRADE OF THE CONTROL SYSTEMS FOR THE CERN LHC GAS SYSTEM DISTRIBUTION MODULES

P. Vanslambrouck†, G. Rigoletti‡, R. Guida, B. Mandelli, M. Quilichini, G. Thomas
CERN, Geneva, Switzerland

*Abstract*

At the experiments of the CERN Large Hadron Collider (LHC), 31 gas systems are used to deliver a precise gas mixture to the corresponding detectors with high reliability and availability. The Distribution Modules of these gas systems control and monitor gas flow across thousands of channels. The Embedded Local Monitor Board (ELMB), an I/O board using the CANopen communication protocol, is used to read mass flow meters in 16 of these gas systems. After years of operation, the PLC hardware and software of these Distribution Modules have become obsolete, raising the need for complex support and impeding improvements. This led to a comprehensive upgrade which was first tested on a retired LHC gas system and subsequently deployed on the production gas system of the Monitored Drift Tube (MDT) detector of the ATLAS experiment. The upgrade allowed the addition of several features, for instance the reporting of detailed sensor error codes and diagnostic information about the ELMBs. System maintainability and upgradability were improved, ensuring uninterrupted operation for the upcoming years of LHC operation.

## LHC GAS SYSTEMS

The gas systems at the LHC experiments are responsible for supplying a stable and controlled gas mixture to gaseous detectors [1, 2]. The gas mixtures must meet stringent requirements since small fluctuations in composition, pressure or the presence of impurities can influence detector performance. However, there are several processes that are constantly degrading the gas quality, such as leaks, diffusion of $H_2O$ and $O_2$ through permeable materials or the breakdown of gas molecules caused by radiation and the electric field within the detector. To maintain a proper gas mixture, a constant influx of fresh gas is therefore needed.

The LHC gas systems are generally distributed across several locations. In the surface gas buildings, the primary gas supply lines are situated and the appropriate gas mixture is formed. The gas mixture is then delivered to the underground service cavern, where it is split into several pre-distribution lines. Next, the gas is supplied to distribution racks located in the experimental cavern where it is further split into channels connected to the detector chambers. To minimize gas consumption, most gas systems recycle the gas that exits the detector. These gas recirculation systems have a pump and purifier module that allow reusing and purifying the gas mixture to minimize the amount of gas that is exhausted into the atmosphere. A diagram of a typical LHC gas recirculation system is shown in Fig. 1.

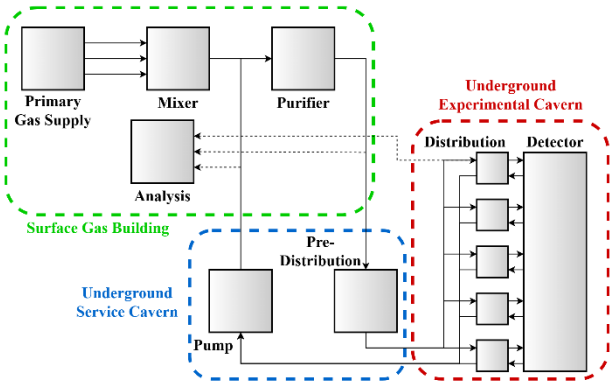



Figure 1: Simplified diagram of an LHC gas system.

The gas distribution racks are the only part of the gas system that is in direct contact with the detector. In addition to splitting the gas in tens to hundreds of channels, these racks are also responsible for monitoring the gas that is sent into the detector. For instance, every channel has a flow cell (a custom thermal mass flow meter) at both the input and the output of the detector. This ensures that the gas flow is properly distributed and serves as the primary tool for detecting leaks at the detector level.

## THE FLOW CELL AS A THERMAL MASS FLOW METER

In the early 2000s, the CERN gas team needed an appropriate gas flow meter for the LHC gas system distribution modules. The device had to be accurate, compact, have a low pressure drop, be resistant to magnetic fields up to 0.1 Tesla and – most importantly – be cost-effective due to the large number needed [3]. At the time, industrial gas flow measurement devices violated at least one of these requirements, which led to the development of a custom thermal mass flow meter suitable for various gases and flow ranges: the flow cell.

### *Operating Principle*

A flow cell consists of an aluminium block with an internal gas channel. In the middle of the channel, a Pt100 resistance temperature detector (RTD) is mounted which is used for both measuring temperature and heating (Fig. 2). The Pt100 satisfies the following relationship [3]:

$$R = R_0 + 0.385\,\frac{\Omega}{\mathrm{K}} \cdot (T - 273.15\ \mathrm{K}) \ \text{ with } \ R_0 = 100\ \Omega.$$

Hence, the temperature $T$ can be determined by calculating the resistance $R$ based on applying a fixed current and measuring the resulting voltage:

† pieter.vanslambrouck@cern.ch
‡ gianluca.rigoletti@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD011

# ACTIVE MAGNETIC BEARINGS ELECTRONIC SYSTEM UPGRADE FOR CERN'S (HL)-LHC CRYOGENIC 1.8 K COLD COMPRESSOR UNITS

M. Pezzetti, CERN, Geneva, Switzerland
A. Kral, SKF Magnetic Mechatronics, Saint-Marcel, France

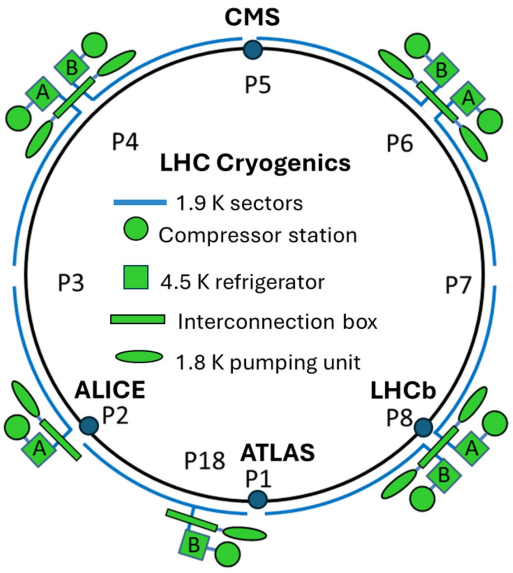
Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Large Hadron Collider (LHC) at CERN relies on superfluid helium, supplied by eight large refrigeration units, each providing 2.4 kW at 1.8 K. These units, developed by two specialized cryogenic industrial suppliers, consist of serial hydrodynamic cold compressors equipped with radial centrifugal impellers, high speed electric motors and Active Magnetic Bearings (AMB), coupled with volumetric warm screw compressors. The AMB electronic units, delivered by suppliers, have been in reliable operation for over 20 years. As part of the CERN consolidation project, the need to meet evolving LHC cryogenic operational requirements along with the obsolescence of electronic components—has driven the upgrade of the entire electronic and electrical control system. A step-by-step analysis, beginning with an operational risk assessment, led to the design and development of a prototype. This effort was undertaken at CERN in collaboration with the French company SKF Magnetic Mechatronics™ as one of the AMB supplier. The prototype underwent extensive testing, first in a dedicated CERN cold compressor test bench and later in real-system operation, where it was successfully validated. This paper presents the complete upgrade process, the positive test results obtained, and the outlook for full deployment during the LHC Long Shutdown 3 (2026–2029).

## INTRODUCTION

The CERN LHC cryogenic system architecture (Fig. 1) ensures the refrigeration of the superconducting magnets and RF cavities, providing stable operation conditions at cryogenic temperatures. It is organized into four cryoplant islands (P18-2, P4, P6, and P8), each responsible for the refrigeration of two machine sectors. In the sector under consideration, the system comprises two plants: an 18 kW refrigerator at 4.5 K, delivering supercritical helium at 4.5 K, and a 2.4 kW refrigerator at 1.8 K. Helium at 4.5 K and 3.5 bar is distributed through a valve box and a 3.3 km cryogenic transfer line, which also interfaces with the 1.8 K pumping plant and the heat exchangers embedded in the cold mass. There are 2 types of suppliers for the cryogenic system but built on the same principle for 1.8 K refrigeration setup, based on a combined cycle with warm screw compressors working at sub-atmospheric pressure, cold centrifugal compressors. The difference lies in design, Air Liquide™uses 3 stages of cold compressors, Linde-IHI™uses 4 stages. In both cases, the cold compressors use Active Magnetic Bearings and Variable-Frequency Drives. Each compressor set is controlled by a dedicated PLC with advanced software that manages them individually while coordinating the overall pumping demand. The 1.8 K refrigeration units are fully integrated into the overall cryogenic system. The CERN LHC cryogenic equipment associated with its control system [1] has been commissioned back in the 2006-2007 and put in regular operation since 2008.



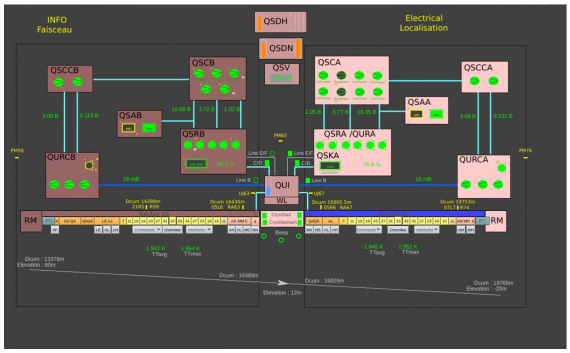




Figure 1: CERN LHC cryogenic system architecture.

## EFFECT ON THE GLOBAL AVAILABILITY OF THE 1.8 K UNITS

At CERN, the installed base of cold compressors consists of 12 cartridges from Air Liquide™and 16 from Linde IHI™. In addition, seven spare cartridges (3 from the first vendor and 4 from the second) are kept in reserve. These spares are also operated on a dedicated test bench, developed specifically to validate Active Magnetic Bearing (AMB) and Variable-Frequency Drive (VFD) upgrades.

During the LHC lifetime, the global availability of the cryogenic system has consistently remained between 95%

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD012

# A NEW PLC BASED CONTROL SYSTEM TO ORCHESTRATE THE PS MAIN POWER CONVERTERS

M. Marin*, O. Minaiev, W. Niinimaeki, J. M. de Paco, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Power for the PS (POPS) system supplies the Proton Synchrotron (PS) magnets by exchanging energy between capacitor banks and magnet loads, minimizing direct power draw from the electrical network. The POPS+ upgrade improves availability through redundancy by integrating additional power converters, introducing a new PLC-based control system to manage increased complexity and retrofit legacy turnkey controls infrastructure. This distributed architecture features a central PLC orchestrating all power converters and their FGC4-based controllers. The PLC enables the custom integration of the FGC4 platform, originally designed for generic converters, into POPS+. A key challenge is the implementation of a real-time Ethernet communication protocol (FGC4PN) on a Multi-Functional Platform (MFP) attached to the PLC. The PLC also manages the state machine, including startup, charging and stopping sequences, as well as interlocks and safe start conditions, leveraging CERN's UNICOS framework for standardized control and SCADA supervision. This paper presents the design and implementation of the control system, detailing the distributed architecture and the communication strategies. It also highlights the use of the Hardware-in-the-Loop (HIL) simulation platform for the system integration, development and virtual commissioning.

## INTRODUCTION

The Power for the Proton Synchrotron (POPS) is the converter system responsible for supplying and controlling the magnetic fields of all main magnets of the Proton Synchrotron (PS) accelerator at CERN. It supplies energy to the main magnets of the PS accelerator by exchanging energy between large capacitor banks and magnet loads. This configuration significantly reduces the direct power draw from the electrical network, limiting peak demand and enhancing grid stability. The existing POPS system comprises two Active Front End (AFE) converters interfacing with CERN's 18 kV electrical network and six DC/DC converters regulating the output voltage and power to the PS magnets [1].

Originally, the POPS control system relied on a proprietary solution provided by CONVERTEAM. This legacy control system featured the P80 VME-based Main Controller (MC) with distributed WAGO IO modules for signal acquisition and control, and an FGC2 module responsible for generating voltage references and precise regulation of the magnetic field [2]. The MC acted as a central node for receiving high-resolution voltage signals and allocating voltage references across converters [3].

* marcos.marin.rodriguez@cern.ch

Due to obsolescence, manufacturer support discontinuation, and maintainability challenges associated with proprietary hardware, a modernization became critical to mitigate operational risks.

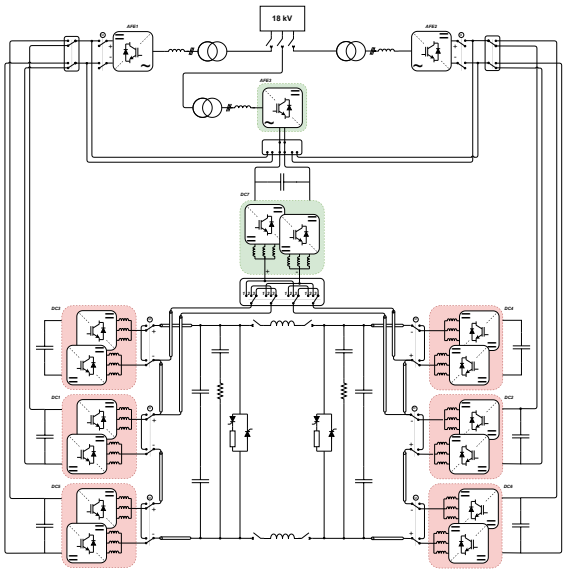



Figure 1: POPS+ power layout. In green, the added converters providing redundancy and availability to the system [4].

The POPS+ upgrade adds one AC/DC and one DC/DC converters to provide redundancy and improved availability, ensuring continuous operation even under degraded (n-1) conditions, as shown in Fig. 1. A centralized Siemens S7-1500 PLC replaces the legacy P80-based Main Controller (MC), together with the new CERN's in-house developed FGC4 controllers which generates the voltage reference and precisely regulates the magnetic field [5]. The PLC's role is to orchestrate the overall system operation, managing state machines of both master and local FGC4 controllers, motorized switches, interlocks, and monitoring all switch positions. This centralized orchestration simplifies system diagnostics, enhances maintainability, and ensures robust operation [6].

The installation and commissioning of POPS+ will occur during CERN's third Long Shutdown (LS3) [4], whereas the PLC deployment will consist of two phases:

- **Phase 1:** During LS3, Siemens IO modules will initially be installed exclusively to interface wired signals with the FGC4 controllers. Concurrently, existing WAGO IO modules will remain operational to maintain backwards compatibility, with the addition of a PROFINET communication header to interface these modules with the Siemens S7-1500 PLC. This tem-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD014

# INSTALLATION AND COMMISSIONING PROGRESS OF THE 2PACL CO2 COOLING CONTROL SYSTEMS FOR PHASE II UPGRADE OF THE ATLAS AND CMS EXPERIMENTS

L. Zwalinski, A. Baran, V. Bhanot, M. Ciupinski, J. Daguin, L. Davoine, M. Doubek, S. Galuszka, Y. Herpin, P. Petagna, K. Sliwa, D. Teixeira, B. Verlaat, M. Zimny, CERN European Organization for Nuclear Research, 1211 Geneva 23, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

In the scope of the High Luminosity Program of the Large Hadron Collider at CERN, the ATLAS and CMS experiments are progressing in the installation and commissioning of their environmentally friendly low temperature detector cooling systems for their new trackers, calorimeters and timing detectors. The selected "on-detector" cooling solution is the $CO_2$ pumped loop concept which is the evolution of the successful 2PACL technique allowing for oil-free, stable, low-temperature control. These systems are of unprecedented scale and largely more complex for both mechanics and controls than installations of today. This paper will present a control system overview, applied PLC architecture and the installation and commissioning progress achieved by the EP-DT group at CERN over the last years. We will describe in detail homogenized solutions which spreads between surface and underground and have been applied for future $CO_2$ cooling systems for silicon detectors at ATLAS and CMS. We will describe in detail applied multi-level redundancy for electricity distribution, mechanics and controls. We will discuss numerous controls-related solutions deployed for electrical design organization, instrumentation selection and PLC programming. We will finally present how we organized early control system commissioning as initial step for LHC Long Shut down 3.

## INTRODUCTION

The detectors of the two largest experiments at the LHC, ATLAS and CMS, are getting upgraded to cope with challenging requirements set by the High-Luminosity LHC. The core part, composed of tracking and timing layers, together with calorimeters for CMS, will be fully silicon-based detectors requiring cold operation at temperatures as low as -40°C. To cope with challenging cooling needs, Detector Technology Group at CERN in the Experimental Physics Department, since 2009 has been developing $CO_2$ cooling systems, of various size and complexity, based on the 2-Phase Accumulator Control Loop (2PACL) [1] concept. Multiple smaller installations with cooling powers up to 15kW, have been in continuous operation at CERN for the LHCb Velo and UT, ATLAS IBL and CMS Pixel Phase-I detectors. The Phase-II upgrade detectors will have much higher cooling power needs and complexity with 310kW for ATLAS and 550kW for CMS and about 1000–1800 evaporators each. These new, more efficient systems will replace the older cooling circuits that used synthetic refrigerants, significantly reducing the environmental impact of detector cooling.

The cooling systems are located in the underground service caverns where the 2PACL cooling plant and most of the control cabinets are located, displaced more than 120m away from the detector proximity distribution manifolds and interconnected via concentric insulated transfer lines, shown in Fig. 1. Due to the larger system volume, and to ensure smooth operation in all temperature conditions that span from +15°C to -40°C, additional Surface Storage tanks will be installed at surface of ATLAS and CMS, interconnecting with underground via two thermally insulated transfer lines. This 'Surface Storge' is 12-cubic-meter storage vessel ensuring $CO_2$ fluid exchanges in a thermosiphon loop between itself and the underground cooling infrastructure following the detector power needs.

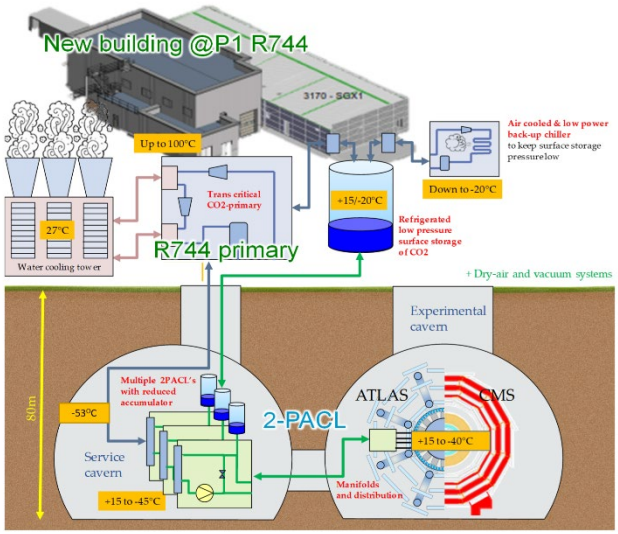



Figure 1: 2PACL system general layout.

The detector's dissipated power removed by the oil-free 2PACL will be rejected to an industrial-like R744 ($CO_2$)-based primary refrigeration system operated by the Cooling and Ventilation Group of the Engineering Department at CERN. The R744 plants are in newly constructed surface buildings and they interconnect with 2PACL via two warm, long transfer lines of about ~90m length. The R744 systems pre-cool the 2PACL to about -53°C (very close to the freezing point of $CO_2$).

In total, the number of cooling plants to be installed is as follows:

ATLAS:

- 6 (+1 backup) 2PACL systems, and,
- 6+1 R744 slices,

CMS:

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD015

# CONSOLIDATION OF THE STATE CONTROL AND SURVEILLANCE SYSTEM OF THE LHC BEAM DUMP SYSTEM

C. Monier*, C. Boucly*, N. Magnin, V. Senaj, L. Strobino, O.Y. Yagci
CERN, Geneva, Switzerland

*Abstract*

The Large Hadron Collider (LHC) Beam Dump System (LBDS) includes 15 extraction kickers (MKD) and 10 dilution kickers (MKB), each powered by a High Voltage Pulse Generator (HVPG), controlled by the State Control and Surveillance System (SCSS) based on industrial PLC technology. After almost 20 years of reliable operation, a software and hardware consolidation is planned during Long Shutdown 3 (2026–2029) to meet the increased demand for diagnostics and functionalities, and to ensure component longevity until the end of LHC operation in 2041. This paper describes the analysis conducted through a detailed review of the existing hardware, software, network layers, and aging field-bus components. It presents the motivations for modernising the SCSS and the new control architecture, including improvements to the safety functionalities implemented. It provides an overview of the new system's interlock state machine with its integration in CERN control middleware.

## MOTIVATION

The LBDS control system, though highly reliable, faces increasing challenges due to hardware obsolescence and limited communication capabilities. These constraints restrict diagnostics and hinder the implementation of advanced functionalities—particularly as data requirements grow to support predictive maintenance and system optimisation. To address these limitations, the Accelerator Beam Transfer (ABT) group has validated a consolidation strategy aimed at modernising the control architecture. This initiative will replace aging field-bus technologies with high-bandwidth communication protocols, enabling richer data exchange and enhanced system supervision. The consolidation also aligns with an ABT strategy to harmonise control across CERN's accelerator infrastructure. By upgrading both hardware and software, the LBDS will be equipped to support future functionalities such as fault analysis tools and automatic conditioning of kicker systems—ensuring operational efficiency and system longevity through 2041.

### *Introduction to the Current LBDS System*

The LBDS controls ensure safe beam extraction from the LHC. Its control architecture is built around four subsystems [1].

- **State Control and Surveillance System (SCSS)**: Based on a fail-safe multi-PLC architecture. Manages global and subsystem states, integrates safety functions and expert test features.
- **Trigger Synchronisation and Distribution System (TSDS)**: Coordinates timing and synchronisation of beam dump triggers.
- **Beam Energy Tracking System (BETS)**: Tracks beam energy and adjusts HVPG output accordingly.
- **Post Operational Checks (IPOC/XPOC)**: Verify that all subsystems reacted correctly during execution of the last dump.

## SCSS OVERVIEW

The SCSS ensures safe and reliable control of the beam extraction and dilution systems using a fail-safe multi-PLC architecture composed of one master PLC and 25 generator PLC. The PLC technologies underpinning the SCSS were initially deployed in 2006, forming the foundation for its safety and control functions throughout the system's operational life-cycle.

### *SCSS Functionalities*

- Person protection with Electrical Discharge Switch (EDS), Manual Earthing Switch (MES), Emergency Stops (AUE).
- Machine protection with a communication with the Trigger Synchronisation Unit (TSU [2]), Beam Interlock System (BIS), TSDS.
- Reception of LHC energy with Beam Energy Tracking System (BETS [3]), computation of HVPG voltage reference through reference tables.
- Voltage control for Power Trigger Module (PTM).
- Acquisition and checking of HVPG voltages.
- Check of trigger signals throught the Power Trigger Controller (PTC [4]).
- Check of HVPG voltages with tracking tables.
- Sequential test with simulation of pulses, cycles, energy scan, IPOC/SAM trigger, retrigger lines.

### *Safety System*

- Designed to meet SIL2 and SIL3 standards (IEC 61508, EN 954-1, IEC 62061, EN ISO 13849-1).
- Implements safety functions for both personnel and machine protection.
- Includes emergency stop monitoring, HV interlocks, and power distribution safety.

### *Current Control Architecture*

- Centralised master controller (Siemens S7-400) with distributed generator controllers.
- Communication via MPI, PROFIBUS-DP, AS-Interface and CERN Ethernet (TN).

* colin.monier@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD016

# ONLINE ANALYSIS FOR KICKER MISSING PULSE DIAGNOSIS

C. Boucly*, M. Algelly, N. Magnin, C. Monier, D. Voscek, O.Y. Yagci
CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The PS beam extraction system includes 12 kicker magnet modules, nine in section 71 and three in section 79, designed to deliver full kick strength for ejecting a 28 GeV/c beam. Since 2020, sporadic missing pulses caused by aging HV generators linked to old electronic control equipment have reduced performance and have been challenging to diagnose. This led to the development of a Missing Pulse Detection Analyser to assist expert diagnostics. Started in 2021, the offline tool correlates kick pulse waveforms with timing data logged in CERN's data logging system (NXCALS), providing an analytical and statistical overview. It has since become an online pulse-to-pulse analyzer that uses data from post-mortem acquisition, the Internal Timing System, and the Generator State Controller, all accessed through Front-End Software Application (FESA) classes. A compact feed-forward neural network, added in 2024, improves early detection of waveform deviations and missing pulse patterns. Developed in Python within CERN's Unified Control Application framework (UCAP), the analyzer interfaces seamlessly with FESA and the Java API for Parameter Control (JAPC), publishing diagnostics through control middleware. This paper details its architecture and initial deployment on the PS Complex (KFA71/79), highlighting operational experience, diagnostic advantages, and plans for integration within the Efficient Particle Accelerator (EPA) framework, including expansion to additional subsystems for the upcoming control consolidation during the 2026 long shutdown.

## INTRODUCTION

The PS extraction is an old system deployed in 1973 and partially consolidated in 2000 [1] to respond to the LHC's requirements with limited diagnostic and remote control. Early in 2021, a new consolidation of the Kicker Generator Controller (KGC) [2] was introduced, enhancing remote diagnostics and improving functionalities for remote control. With this improvement, the identification of sporadic missing pulses that triggered radiation alarms in the PS extraction system appeared easier, but was a challenge in diagnosing these events using existing tools, which initiated the development of a dedicated software-based analyzer. Initially designed in 2021 as an offline analysis tool using Jupyter notebooks and NXCALS [3] data, its primary goal was to help experts correlate generator timing, kick settings, and permit conditions. Through successive iterations, the analyser evolved into an online pulse-to-pulse diagnostic system. It now ingests data from multiple sources, including the Internal Post-Operation Check (IPOC) [4], the Kicker Internal Timing System (KITS) [5], and the Kicker generator Controller (STATE) status. In 2023, the online analysis was deployed into the operational environment, featuring a PyQt6-based graphical interface for live monitoring and automatic publishing of diagnostic results into CERN's control infrastructure. In 2024, a feed-forward neural network was introduced to enable predictive assessment of pulse quality with objectives such as providing early warnings of waveform degradation or anomalous behavior. This paper presents the system architecture, analysis logic, GUI components, and introduces the initial predictive results, along with online analysis, and the roadmap toward future deployment across other kicker systems.

## SYSTEM OVERVIEW

The missing pulse detection system integrates several essential components. A real-time analyser handles signal acquisition and anomaly detection. An online GUI enables visualization to support end-users with diagnostics and decision help. The predictive model enhances this diagnosis by scoring potential pulse degradations based on specific features. Finally, the system interfaces seamlessly with CERN control frameworks, including FESA [6], JAPC, and UCAP [7].

## OFFLINE ANALYSIS DESIGN

The development of the missing pulse analyser originated with an offline prototype created in 2021 using SWAN [8] (Service for Web-based ANalysis) platform. The SWAN service, developed at CERN, provides an interactive cloud platform for scientific data analysis. It integrates Jupyter notebooks and Docker-based environments with existing IT services such as virtualized cloud resources, user authentication, dedicated clusters, and long-term storage, as integration infrastructure with components platform are presented in fig. 1. The offline analyzer objective was to perform post-mortem analysis of kicker behavior during PS extraction cycles. This early tool aimed to correlate beam permit conditions, kick requests, and actual waveform characteristics to identify missing or degraded pulses.

### *Variable Acquisition for Analysis*

The offline analysis used archived control data from the **NXCALS** platform, specifically targeting Expert Acquisition arrays and timing events from the KFA71 controller and kicker generators:

- `arrKickHWInterlockToPlay` – Interlock during kick requests.
- `cyclePlayed` – Identifies the cycle ID triggered during the sequence.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD017

# REAL-TIME CONTROL SYSTEM UPGRADE OF THE CERN LINAC4 PRE-CHOPPER

M. Niccolini*, T. Gharsa, F. Vaga, N. Magnin, C. Boucly, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The CERN Linac4 pre-chopper, installed right after the H- ion source in the Low Energy Beam Transport (LEBT) section, plays a crucial role in providing the 45 keV H- beam to the first accelerating structure, the Radio Frequency Quadrupole (RFQ). By applying a pulsed electric field of -20 kV, the pre-chopper deflects the beam when not required and sharps its head and tail in order to remove the long rise time of the source and avoid transmission losses. The existing pre-chopper controller was implemented in 2015 using National Instruments (NI) LabVIEW and PXIe hardware, relying on their proprietary Real-Time (RT) operating system (Phar Lap) and a secondary Linux Front-End Computer (FEC) for the integration in the CERN control system. Phar Lap is EOL since 2025 and will be discontinued during the upcoming Long Shutdown 3 (LS3). This paper presents an upgrade of the control system, aimed at replacing the LabVIEW-RT control layer with standard CERN solutions, leveraging the new Debian-based Linux RT OS, Front-End Computer Operating System (FECOS), and consolidating all functionalities into a single computer. The goal was achieved using the CERN Front-End Software Architecture (FESA) 3 framework and C++ libraries to interface with the NI hardware via NI Linux drivers deployed on FECOS. A new PyQt-based graphical user interface will be developed to ease system monitoring and operation. Installation of the upgraded system is expected for LS3, using a custom PXIe crate and CPU from CERN instead of NI solutions.

## INTRODUCTION

The pre-chopper is located in the Linac4 [1] LEBT and deflects the beam away from the RFQ input aperture when it is not required. It serves to reduce the beam start-up time from the H- source, to shorten its pulse length, and to stop the beam due to machine interlocks. The beam is chopped by applying -20 kV to a plate in vacuum, which is switched to 0 V when the beam should pass. Figure 1 shows the position of the system in the accelerator. The beam is chopped (-20 kV applied) before the H- ion source is started by receiving a first trigger signal (Pre-chop1, see Fig. 2). To start to deliver the beam to the RFQ, the voltage is switched to 0 V (Pre-chop2). Once the PSB (Proton Synchroton Booster, the next accelerator in the chain after LINAC4) has finished the injection sequence, which is of variable length depending on the type of beam produced, the pre-chopper deflects the beam again (Pre-chop3). Once the ion source stops, the pre-chopper plate returns to 0 V (Pre-chop4). The pre-chopper has also the role of deflecting the beam if it is not required due to an interlock condition provided by the Beam Interlock System (BIS) [2]. Since the application of -20 kV onto a deflection plate is inherently not fail-safe, the pre-chopper must be used in conjunction with additional beam inhibiting devices, managed at the level of BIS.

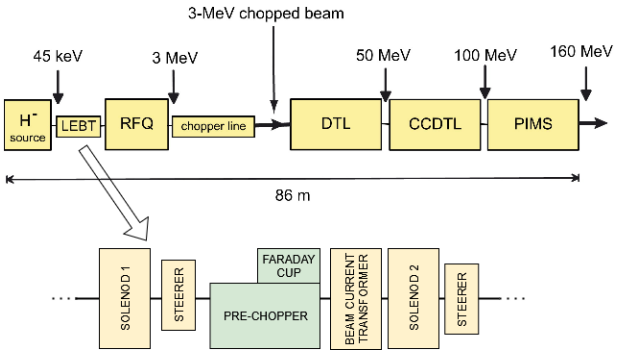



Figure 1: Schematic layout of the Linac4 accelerator. The LEBT section is shown in an exploded view, highlighting the the pre-chopper position.

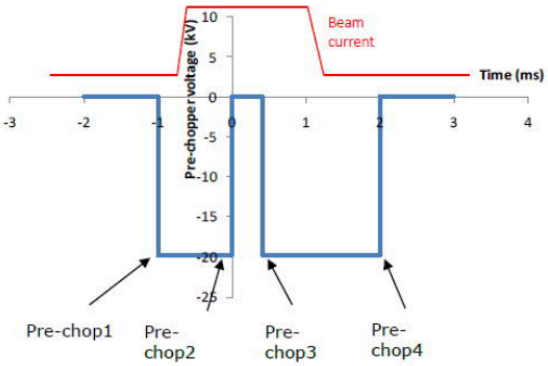



Figure 2: Pre-chopper high voltage pulse timing sequence.

### *Motivation for the Upgrade*

The eradication of the LabVIEW RT is motivated by the obsolescence of Phar Lap, no longer supported by NI since 2025 [3] and of the NI library used with SILECS (Software Infrastructure for Low-level Equipment Controller, formerly known as IEPLC [4]) to interact with the CERN control system stack. A control solution based on FESA 3 [5] and LabVIEW FPGA is proposed as an upgrade that simplifies the integration in the CERN control system architecture.

## CONTROL SYSTEM LAYOUT

### *The Current Control Layout*

The pre-chopper control functionalities, depicted in Fig. 3, are deployed on an NI PXIe controller (NI PXIe-8100) as-

* marco.niccolini@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD019

# CONTROLS OF THE NEW EDDY CURRENT SEPTUM FOR THE CERN PS FAST EXTRACTION

L. Strobino*, S. Klomsten, C. Lolliot, N. Magnin, S. Pavis,
P. Van Trappen, N. Voumard, O.Y. Yagci, CERN, Geneva, Switzerland

*Abstract*

The CERN PS fast extraction septum deflects protons and ions towards the experimental areas of the PS complex and the SPS. With the increased number of extractions per year since it was first put into service in 1994, the magnet lifetime is nowadays estimated at two years, implying bi-annual rebuilds of the septum, significant costs, and non-negligible radiation doses taken by the personnel. Additionally, the present power converter is approaching the end of its life.

In view of its superior robustness, an eddy current septum system was chosen to replace the original direct-drive septum. Due to the different technology of the magnet, the existing power converter is replaced by a new fast pulse generator, which implies a complete new control system.

This paper describes the different units and functionalities of this new control system, covering a wide range of technologies such as high-voltage switch triggering modules, slow interlocks based on PLC, fast interlocks and timing implemented in FPGAs, temperature-compensated acquisition chains, and software-based regulation algorithms. Preliminary results of the system performance are also presented.

## INTRODUCTION

The PS fast extraction septum (PE.SMH16) deflects protons and ions into the extraction transfer line TT2 towards the experimental areas of the PS complex and the SPS. Since its commissioning in 1994, the SMH16 has seen its number of extractions increase significantly, reducing the lifetime of each magnet tank from the original 5-6 years to $\leq 2$ years. This implies regular refurbishment of the septum and a significant dose taken by the personnel.

For its robustness, an eddy current septum system was seen as the most suitable design choice for consolidation. While reusing the existing vacuum vessels and remote displacement systems, the magnets will be switched from a direct-drive septum to an eddy current septum. Due to the different technology of the magnet, the existing power converter needs to be replaced by a fast pulse generator. In addition, this consolidation is an opportunity to improve the system to meet future operational requirements.

## PULSE GENERATOR

The septum magnet requires a current with an amplitude of up to 29 kA and a flat-top length varying from 2 µs (with a tight flatness tolerance) to 11 µs (with slightly relaxed tolerances). This flat top, generated by superimposing a fundamental sine wave of 2.5 kHz and a third harmonic with 11% of the amplitude of the fundamental, must be stable and reproducible from pulse to pulse [1]. The two sine waves are created by discharging two capacitor banks into the magnet through solid-state switches built from a stack of 5 GTOs (Gate Turn-Off thyristors) in series to achieve the required maximum blocking voltage. Each GTO also has an antiparallel diode. The amplitude of the flat-top is directly proportional to the capacitor charging voltage, and the slope can be fine-tuned by changing the delay between the triggering of the two switches. The typical magnet and GTO stacks current waveforms are shown in Fig. 1.

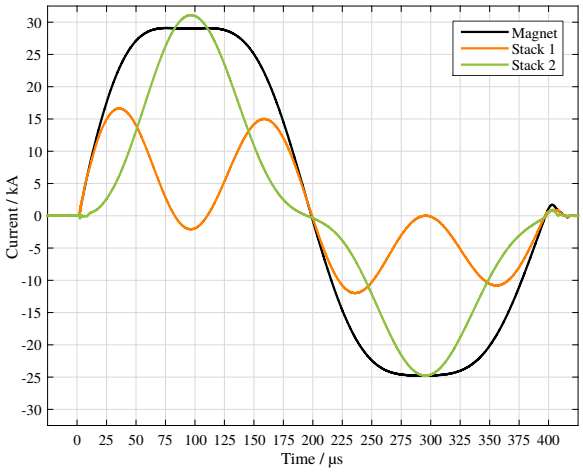



Figure 1: Typical SMH16 generator current waveforms.

## CONTROLS

The consolidation of the septum magnet required a complete new control system for two reasons. First, the current power converter had been designed by another group and is at the end of its life, making it hard to consolidate. Additionally, the change of magnet technology required a different generator topology, leading to a different controls architecture. For this reason, the new generator and its controls are seen more as a new equipment rather than a consolidation.

The control system is based on standard technologies and generic solutions already in use in the group [2], such as PLC systems for personnel safety and slow interlocks and FPGA-based systems for timing and fast interlocks.

### *Power*

Figure 2 shows the racks of the test system installed in building 867. The leftmost blue rack (located adjacent to

* lea.strobino@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD021

# CERN SCADA SYSTEMS 2024 LARGE UPGRADE CAMPAIGN RETROSPECTIVE

L. Goralczyk, A. F. Kostopoulos, M. Bes, O. Kazepis,
CERN, Geneva, Switzerland

*Abstract*

This paper presents a recent upgrade campaign of supervisory control systems within CERN's Accelerators and Technologies Sector. The effort covered over 240 WinCC OA SCADA applications across more than 120 servers, spanning core accelerator systems such as Power Converters, the Quench Protection System, and the Power Interlock Controller, along with essential technical infrastructure including Cryogenics, Vacuum, and Gas Control. These systems are crucial for machine protection, performance, and the reliable operation of the accelerator complex. Building on experience from previous upgrade efforts, this campaign introduced important advances in automation and process optimization. For the first time, a fully unattended upgrade workflow was achieved through the use of Ansible. In addition, the campaign involved a major operating system migration and the upgrade of several supporting satellite systems. This paper details the improvements made in this iteration, discusses the challenges and compares the current campaign with earlier ones. The analysis highlights the evolution of automation strategies and reflects on both successes and difficulties. The work offers valuable insights for future upgrade initiatives and demonstrates how automation tools can significantly enhance the maturity and reliability of large-scale software maintenance in complex operational environments.

## INTRODUCTION

At CERN, WinCC OA-based SCADA (Supervisory Control And Data Acquisition) systems undergo regular upgrades dictated by new versions of the underlying software ecosystem. The upgrades introduce new features, resolve existing bugs and improve security. They also move us away from end-of-life versions, mitigating risks associated with lack of security patches and maintenance.

Uninterrupted operation of the CERN accelerator complex and experiments relies on a diverse range of industrial control systems. These systems encompass critical areas such as electrical power distribution, machine protection, cryogenics, cooling and ventilation and vacuum. These systems are managed through a SCADA infrastructure comprising approximately 240 applications running on 120 physical servers. Importantly, these applications do not operate in isolation; they are interconnected and rely on external services for functionality such as data archiving, alarm distribution and system monitoring.

This paper details a SCADA software upgrade campaign conducted mainly in 2024. Building on experience from previous upgrades, we present an automated upgrade methodology and discuss the technical and logistical challenges encountered during implementation. We provide quantitative results demonstrating the campaign's success and conclude with recommendations for streamlining future SCADA upgrades at CERN.

For details regarding the previous campaigns the reader is referred to [1] and [2].

## PREVIOUS UPGRADE CAMPAIGNS

CERN's SCADA upgrade campaigns have evolved significantly over the past decade and a half, transitioning from manual processes requiring extensive documentation to fully automated workflows. This evolution reflects both technological advances and lessons learned from operational experience.

Table 1: Automation and Interface in Past Upgrades

| Version | Year | Automated? | Interface |
|---|---|---|---|
| 3.6 → 3.8 | 2010 | No | – |
| 3.8 → 3.11 | 2013 | Partially | GUI |
| 3.11 → 3.15 | 2017 | Mostly | GUI |
| 3.15 → 3.16 | 2020 | Mostly | CLI |
| 3.16 → 3.19 | 2024 | Fully | CLI |

Upgrades happen every 3-4 years. The upgrade start date is influenced by many factors but the major ones are the Large Hadron Collider (LHC) schedule, the WinCC OA (SCADA software toolkit used at CERN) new version availability and the major releases of the underlying operating system (OS). The date is carefully chosen and discussed well in advance. The preparations usually start one year before the planned start date, which is enough time to agree on the upgrade schedule and prepare tooling and procedures.

Table 1 reveals a clear trend toward increasing automation over successive campaigns. The shift from GUI-based automation (2013-2017) to command-line interfaces (CLI, 2020-2024) further reduced manual intervention requirements. CLIs enable easier scripting, consume fewer resources and provide more consistent cross-platform behavior than GUIs.

Although successive campaigns share fundamental phases of planning, tooling preparation and testing, each upgrade campaign must be tailored to accommodate evolving operational requirements, changing system architectures, and lessons learned from previous implementations.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD022

# PROGRESS UPDATE ON THE SUPERCONDUCTING UNDULATOR CONTROL SYSTEM FOR THE EUROPEAN XFEL

M. Yakopov†, S. Abeghyan, S. Casalbuoni, S. Karabekyan, European XFEL, Schenefeld, Germany
A. Hobl, A. Sendner, Bilfinger Nuclear & Energy Transition GmbH, Wuerzburg, Germany
M. Gretenkord, D. Pieper, Beckhoff Automation GmbH, Verl, Germany

*Abstract*

The development of a superconducting undulator (SCU) afterburner for the European XFEL represents a major upgrade to the facility's photon generation capabilities. Following the first report on the conceptual design of its control system [1], substantial progress has been achieved in recent years. The factory acceptance test (FAT) of the pre-series prototype (S-PRESSO) is scheduled to take place at Bilfinger Nuclear & Energy Transition GmbH in 2026 with the dedicated control rack (SCR), which is fully produced, and currently, the commissioning steps are being performed at European XFEL. The samples of the main and auxiliary power supplies have been delivered, and the integration tests are progressing. In parallel, the design of the control system architecture has been finalised, ensuring seamless compatibility with the existing permanent-magnet undulator (PMU) control system. This contribution provides an update on the current status of the S-PRESSO and future SCU afterburner control system, highlighting the design and production results, and the roadmap towards integration into the SASE2 undulator line in the future.

## INTRODUCTION AND PROJECT OVERVIEW

The European XFEL has established itself as one of the world's leading X-ray free-electron laser facilities, providing unique photon beams to a diverse international user community. To maintain its competitive edge and further expand its scientific capabilities, a program to develop and implement superconducting undulators (SCUs) in the EuXFEL is being pursued [2]. Compared to conventional permanent-magnet undulators, SCUs offer the potential for higher magnetic fields at shorter periods, which translates into the delivery of an increased photon energy range. As part of this program, a dedicated SCU afterburner called FESTA (Free-Electron laser SuperconducTing Afterburner) is planned for the SASE2 beamline. The project foresees the installation of up to six superconducting undulator cells directly following the last PMU segment. The first step is the fabrication and commissioning of the Superconducting undulator Pre-Series prototype (S-PRESSO). This prototype not only validates the magnet and cryogenic designs but also serves as the critical testbed for the associated control and protection systems. The earlier paper [1] focused largely on the conceptual design of the SCU control system, drawing on the experience gained with the existing XFEL undulators. Since then, the project has transitioned from design and planning to hands- on hardware production, commissioning, and system integration. This paper presents the latest progress achieved in these areas.

† mikhail.yakopov@xfel.eu

## CONTROL SYSTEM ARCHITECTURE

As for the PMU systems, the control system is divided into local and remote subsystems. The local control system provides more detailed functionality and fine-tuning components, while the remote system offers a subset view tailored to operators' usage patterns and requirements. The low-level control will be implemented in Beckhoff TwinCAT3 PLC, with a TwinCAT HMI web-based interface for local operation. Remote control will be integrated into the existing DOOCS server, with JDDD [3] serving as the control system interface. Figure 1 illustrates the main control panel of S-PRESSO for local and remote controls. The display highlights two main superconducting coils called SCUCoil1 and SCUCoil2, each powered by dedicated high-stability main power supplies produced by OCEM Power Electronics. The set value can be defined either in K, the undulator deflection parameter directly related to the photon energy, or in amperes. The currents of each of the SCU coils, powerable up to currents of about 1000 A, are ramped to reach a certain magnetic field. Operator's view provides both the setpoint and the actual current values with ppm precision, enabling continuous verification of power supply stability during operation. Around the main coils, a set of auxiliary correction coils AUX1–4 is shown. These coils are powered by auxiliary power supplies and allow fine-tuning of the magnetic field as well as compensation of asymmetries. The panel also integrates horizontal and vertical Helmholtz coils HH1, HH2, as well as a phase shifter located between the two main coils, ensuring the electron beam and the emitted radiation stay in phase. All correction coils (AUX and HH) as well as the phase shifter are superconducting and are powered by power supplies delivered by Deutsches Elektronen-Synchrotron (DESY), providing up to 20 A. Given a certain K value or corresponding current in one of the SCU coils, those are ramped synchronously with the correction coils to ensure that the first and second magnetic field integrals are compensated and remain below the specified limits. This is needed for the undulator to be transparent with respect to the electron beam. The corresponding values of the current in the correction coils are obtained from lookup tables derived during the magnetic field characterization of the system. Furthermore, the control architecture allows the application of

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD023

# SUPPORTING INJECTOR OPERATION WITH THE FAIR SETTINGS MANAGEMENT SYSTEM

J. Fitzek, H. Hüther, R. Mueller, A. Schaller
GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The FAIR Settings Management System is in productive use at the GSI accelerator facility. The core of the system is developed in collaboration with CERN and is based on CERN's Software Architecture (LSA) framework. It is currently used at GSI for operating the synchrotrons, storage rings, and transfer lines. As part of the Injector Controls Upgrade project, which aims to integrate the UNILAC linear accelerator into the new Control System, concepts for scheduling parallel particle beams are being further realized. In preparation for the integration of the FAIR facility, existing concepts are being reviewed, refined and implemented. A key milestone for the upcoming implementation steps is the operation of the UNILAC with the new control system during the beam time 2026, with preparatory beam tests starting in summer 2025. This paper describes the current implementation status.

## REQUIREMENTS OF PARALLEL BEAM OPERATIONS AT FAIR

The existing GSI accelerators, together with the new FAIR facility—currently under installation and entering its first phase of commissioning—present a set of unique and unprecedented challenges for the Control System. Key challenges include the simultaneous management of up to five parallel beams and the capability to switch between different particle species on a pulse-to-pulse basis. In addition, the Control System must provide a high degree of operational flexibility, enabling modifications of parallel operation schemes on a daily basis in order to adapt to the needs of the experimental programs. To meet these requirements, a range of new concepts and technical solutions are being developed and systematically implemented within the new Control System.

For FAIR, the modeling of the beam within the Control System—describing the entire particle from source to target—becomes a central concept. Such a model is indispensable for ensuring efficient operation, since it provides the foundation for scheduling beams into complex parallel operation patterns (see Fig. 1). In day-to-day use, device settings and timing parameters must remain fully modifiable at any point in time, with both high flexibility and minimal delay, so that operators can respond quickly to changing experimental demands.

Beyond these general challenges, the integration of the UNILAC linear accelerator as an injector into the new Control System introduces additional requirements. These include supporting substantially higher repetition rates, enabling genuine parallel beam operation within a single accelerator, and ensuring the continuous distribution of precise timing signals to sensitive equipment across the facility. Each of these aspects demands careful architectural design as well as robust and reliable implementation.

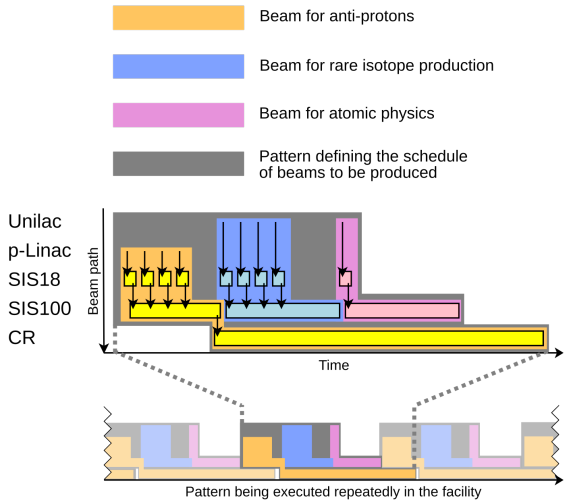



Figure 1: Combination of beams into patterns for execution.

This contribution begins by outlining the milestones and key components of the Control System and the Settings Management System in particular, then introduces the central concept of beam modeling, and concludes with the challenges of integrating the UNILAC injector.

## CONTROL SYSTEM MILESTONES

The new Control System for GSI and FAIR has been operational since 2015 at CRYRING [1], since 2018 at the synchrotron SIS18 [2] and the transfer lines, and since 2019/2020 at the storage ring ESR [3]. The Settings Management System as one core component has undergone major feature enhancements, including extended support for transfer lines with optics, implementation of specific storage-ring functionalities, significant performance improvements through parallelized calculations, and initial realizations of beam-scheduling concepts.

In 2023, the Injector Controls Upgrade project [4] was launched with the goal of integrating the injector into the Control System. The underlying concepts were revised and further developed, leading to successful test runs in 2024 and 2025. A further test run with beam is scheduled for February 2026 to verify major functionality in preparation for the following beamtime, which will be operated exclusively with the new Control System across the entire facility. The legacy system, which had continued to support the UNILAC linear injector, was decommissioned in summer 2025.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD024

# A PYTHON-BASED SERIAL COMMUNICATION FRAMEWORK FOR LEGACY PMAC CONTROLLERS

P. Sreelatha Devi *, A.Balzer †, Helmholtz Zentrum Berlin, Berlin, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Many beamlines at BESSY-II still operate using VME-crate-based PMAC motor controllers that rely on proprietary Windows software for communication and configuration. However, these programs often require old licenses and are not compatible with modern operating systems, making maintenance increasingly difficult. To address this, we have developed an open-source Python tool that communicates with legacy PMAC controllers over serial interfaces. The tool uses a modular manager-client architecture where multiple client programs can send commands concurrently without conflicts, using a managed queue and locking system. Dedicated clients have been created for terminal interaction, batch upload of command files, and a watch window monitoring of PMAC variables, with configurable fetch and display intervals. The programs are lightweight, installed via Python package management, and accessible through simple command-line interfaces. While the serial communication is not real-time, it is sufficient for configuration, monitoring, and motion program uploads. Extensive logging is provided for traceability. A GUI interface and integration into the broader Geobrick framework is implemented and tested. This tool provides a modern, open-source alternative for maintaining legacy motion control systems and ensures continued support without reliance on outdated commercial software.

## INTRODUCTION

A recent cyber attack at the BESSY-II facility has accelerated the obsolescence problems of software dependencies, including security issues of Windows operating systems. There are continuous attempts to modernize the existing infrastructure [1–3]. However, the long operational history of BESSY II has resulted in a diverse and legacy infrastructure, with many beamlines continuing to rely on discontinued VME based motion systems, some based on Delta Tau PMAC VME controllers [4]. These controllers, though functionally reliable, are increasingly difficult to support due to their dependency on proprietary Windows-based software environments, such as Delta Tau's IDE. These environments are no longer actively maintained, often require outdated license mechanisms, and are tightly coupled to specific operating systems and hardware configurations.

From a system maintenance perspective, this creates significant liabilities. The legacy software is incompatible with modern Linux-based beamline workstations and poses additional risks in terms of cyber-security. In particular, updating or re-configuring these systems often involves manual steps, non-portable tools, and physical access to a Windows machine with the appropriate license—a process that contradicts the goals of modern control system standardization and automation.

To mitigate these limitations and extend the lifecycle of critical beamline infrastructure, we present a lightweight, modular, and platform-independent Python software for interacting with PMAC motor controllers over serial interfaces. The design emphasizes clean abstraction, multi-client coordination, and compatibility with both VME-based PMAC and GeoBrick systems. In contrast to closed IDE solutions, our software is fully open-source, maintainable, and suitable for long-term deployment across the BESSY II facility.

## ARCHITECTURE

The developed tool chain follows a modular client-manager architecture, designed for asynchronous serial communication with legacy Delta Tau PMAC motor controllers as illustrated in Fig. 1.

The central component is a Python-based manager process that handles exclusive access to the serial port and mediates all communication between clients and the controller. This approach ensures that concurrent client processes can operate safely without contention or protocol violations.

Communication between clients and the manager is implemented over TCP sockets on the local host. Each client establishes its own socket connection and submits commands prefixed with a functional tag (e.g., TERMINAL, BATCH, WATCHDOG, PLOT). These requests are queued by the manager and serialized using a locking mechanism, ensuring that only one command is processed at a time.

The manager then returns the response back to the requesting client. The backend communication with the PMAC is performed via a lightweight serial protocol using configurable baud rates and port names. A dedicated Python class encapsulates low-level serial I/O, including handling for command-response matching, termination characters, and timeout control. This abstraction allows the system to support multiple PMAC variants (e.g., VME, Geo Brick) without changes to the client programs.

- Each client tool is implemented as a standalone Python script or module:
- The terminal client allows raw command entry and response viewing, with optional decoding of PMAC error/status codes.
- The batch client supports upload of .pmc motion programs, including rudimentary pre-processing of include directives.

* parvathi.devi@helmholtz-berlin.de
† andreas.balzer@helmholtz-berlin.de

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD029

# DEVELOPMENT OF AN EtherCAT-BASED CONTROL SYSTEM FOR AN IN-VACUUM UNDULATOR FOR SPring-8-II

K. Yamakawa[1, †], T. Fukui[2], M. Ishii[1, 2], M. Yamaga[1, 2]
[1]Japan Synchrotron Research Institute, Sayo, Japan
[2]RIKEN SPring-8 Center, Sayo, Japan

*Abstract*

SPring-8, a third-generation light source, has been operating for almost three decades. To transition to a fourth-generation light source, SPring-8 will undergo a major upgrade to its storage ring, which will be renamed SPring-8-II, with the upgrade scheduled for 2027-2028. This upgrade includes the installation of new In-Vacuum Undulators for SPring-8-II (IVU-IIs) and a corresponding modernisation of the accelerator control systems. As part of this, we are replacing the legacy VME-based systems with new EtherCAT-based slow control systems. The installation plan calls for maximum seven IVU-IIs to be deployed annually between 2023 and 2026, each with an EtherCAT control system. In 2024, we successfully implemented the first fully EtherCAT-based control system. This system meets the required 50 um synchronisation accuracy and ensures stable operation of the steering magnets and reliable temperature monitoring, confirming its suitability for the SPring-8-II upgrade program.

## INTRODUCTION

SPring-8 has been operating for approximately thirty years as a third-generation light source. The mainstream of storage ring light sources has shifted towards generating higher brilliance beams with low-emittance, a characteristic of fourth-generation light sources. Several such facilities have recently come into operation. SPring-8 will renew its storage ring, to be called SPring-8-II, to achieve a beam 100 times brighter than the current one by decreasing the beam emittance from 2.4 nmrad to 100 pmrad [1]. The scheduled shutdown will last approximately one year, with SPring-8 shutting down in 2027 and the new SPring-8-II commencing operation in 2028.

The upgrade of Insertion Devices (IDs) is one of the key components of the SPring-8-II upgrade programme. IDs, especially In-Vacuum Undulators (IVUs), are crucial instruments for generating light at SPring-8. In this programme, an upgraded IVU is referred to as the IVU for SPring-8-II (IVU-II) [1, 2]. The IVU-II features the cancellation of the attractive force between magnetic arrays. Installations of these devices began in 2023.

Furthermore, the accelerator control systems are also being upgraded prior to the SPring-8-II project. One of these control system upgrades involves the implementation of MADOCA 4.0 [3]. We have also decided to replace legacy VME modules for slow controls with EtherCAT [4] devices, which employ a master-slave relationship, due to the shrinking market for VME boards [5, 6].

† kosei.yamakawa@spring8.or.jp

## REQUIREMENTS

For an IVU-II, we are required to use EtherCAT to control the following elements:

- Gap control and monitoring
- Power supplies for steering magnets
- Magnet temperature monitoring

The gap is defined as the distance between the upper and lower magnetic arrays. For the IVU-II, we must control these arrays independently. To ensure beam stability and prevent the generation of unwanted spectra, both arrays are required to drive in sync. This synchronisation is critical to prevent the electron beam from being kicked off the center of the gap. The required accuracy is a difference of less than 50 um, a level of synchronisation that is sufficient to neglect any detrimental effects on the beam.

Steering magnets are essential for generating and correcting magnetic fields that cancel out erroneous components from the magnetic arrays. This precise correction is vital for maintaining the high-quality beam needed for advanced experiments.

The magnets' temperatures must be continuously monitored to ensure the operational stability and safety of the IVU-II. This is because temperature fluctuations can directly affect the magnetic field properties, which may lead to deviations in the beam's trajectory and stability. The typical magnet temperature range is 24 to 32 degrees Celsius.

Figure 1 shows the schematic view of the IVU-II magnetic arrays, driving mechanism, and control devices. Two steering magnets are positioned both upstream and downstream of an IVU-II.

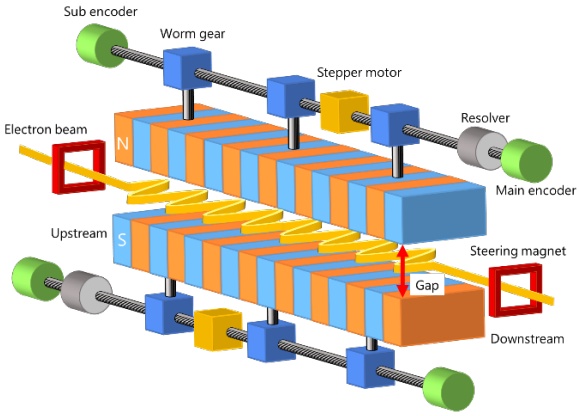



Figure 1: Schematic view of IVU-II magnet arrays, driving mechanism, measuring devices, and steering magnets.

The new system is VME-free and employs EtherCAT, because a lot of VME boards, used for IVUs at SPring-8, were discontinued and EtherCAT is widely used in industry,

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD030

# DESIGN AND COMMISSIONING OF THE UPGRADED ALS SR RF CAVITY WATER CONTROL SYSTEM*

N. Us Saqib†, A. Jurado, C. Toy, D. A. Nawaz, E. Andrade
Lawrence Berkeley National Laboratory, Berkeley, CA, United States

*Abstract*

The Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory, a pioneering third-generation synchrotron light source, has been operational since 1992. The ALS storage ring RF cavity water control system, originally a 30-year-old relay-based chassis system, regulated the temperature of the two storage ring RF cavities while managing interlocks and operational functions. Aging instrumentation and the need for enhanced features, however, led to operational challenges. To address these issues, the system was upgraded to a PLC-based solution. This paper presents the design and commissioning results of the upgraded cavity water control system.

## INTRODUCTION

The Storage Ring (SR) Radio-Frequency (RF) system at the Advanced Light Source (ALS) is a critical subsystem, providing the RF power necessary to accelerate and sustain the 1.9 GeV electron beam. It comprises two single-cell RF cavities operating at 499.642 MHz. For the past 30 years, a relay-based chassis system regulated the RF cavities' temperature while managing operational functions and interlocks. Over time, aging instrumentation and the need for additional features created operational challenges and limited maintainability.

To address these issues, the system underwent an upgrade during the summer shutdown in 2024. This paper details the design and commissioning of this upgrade, which is the first phase of a larger project to replace all of the legacy Horner PLC systems in the SR RF control system [1].

## EXISTING SYSTEM

The existing ALS SR RF cavity water control system was a complex setup, distributed across four racks and encompassing 27 individual chassis and devices. Its relay-based design limited flexibility and made modifications cumbersome. All control logic was hardwired, requiring over 250 cables and cross-connections for installation. Figure 1 illustrates the legacy cavity water control system racks which include portion of the whole system. The main components for each RF cavity are described in the subsections below.

### *Pump Interlock Panel*

Implements the interlock chain required to enable VFD and pump operation.




Figure 1: Existing cavity water control system racks.

### *Cavity Water Control Chassis*

Implements two interlock chains:

- Heater Interlock: Enables the heater when permitted.
- Cavity Cooling Interlock: Aggregates signals to confirm acceptable cavity temperature conditions, forwarding permits to the SR RF master interlock PLC.

### *HOM Probe Chassis*

Implements HOM probes flow and temperature interlocks and generates a permit signal that is included in the cooling interlock chain of the cavity water control chassis.

### *Water Control Amplifier*

Performs PID-based temperature regulation: the operator manually sets a target temperature on the front panel, and the amplifier automatically adjusts the cooling Kates valve and/or heater to maintain the setpoint.

### *Pump Selector Panel*

Implements pump-to-cavity selection: three pumps serve two cavities, with one maintained as a spare. Cavity 1 may be supplied by Pump 1 or Pump 2, and Cavity 2 by Pump 2 or Pump 3.

## UPGRADE SYSTEM DESIGN

The upgraded cavity water control system features a single PLC chassis and an HMI. Figure 2 shows the cavity water control system racks after the upgrade, illustrating the reduced footprint.

* Work supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231

† nusaqib@lbl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD031

# ALS STORAGE RING RF CONTROL SYSTEM UPGRADE PLAN AND STATUS*

N. Us Saqib†, A. Jurado, E. Andrade, Q. Du, J. H. Lee, M. Dach, B. Flugstad
Lawrence Berkeley National Laboratory, Berkeley, CA, United States

*Abstract*

The Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory, a third-generation synchrotron light source operational since 1992, is undergoing a comprehensive upgrade of its storage ring RF control system. The legacy Horner PLC controllers and remote I/O modules, now at end-of-life, are being replaced with an Allen-Bradley PLC platform to improve maintainability, reliability, and long-term support. This paper presents the planning, design, and current status of the upgrade project.

## INTRODUCTION

The Storage Ring (SR) Radio-Frequency (RF) system at the Advanced Light Source (ALS) is a critical subsystem that delivers the RF power required to accelerate and sustain the 1.9 GeV electron beam. It consists of two single-cell RF cavities operating at 499.642 MHz, driven by a pair of klystrons [1]. While the core storage ring RF hardware—including cavities, waveguides, circulators, and high-power switches—remains central to the ongoing ALS-U upgrade project, the associated end-of-life Horner PLC-based control system represents an operational risk, creating substantial challenges for long-term reliability and maintainability of the SR RF control system.

To address this, a phased upgrade plan has been developed to replace the existing PLC controls with a modern, maintainable architecture based on Allen-Bradley PLCs, consistent with ALS-U PLC-based systems [2]. The upgrade is executed in multiple phases, each designed for completion within a two-month shutdown to minimize operational disruption. This paper presents the current SR RF control system, the upgrade plan and status, and the design of the new system.

## CURRENT SR RF CONTROL SYSTEM

The SR RF control system manages the monitoring and control of the storage ring RF system according to its operational configuration. The SR RF system can be configured in six modes which are essential for the control system to manage operation and enforce interlocks effectively:

1. Dual Klystrons: Klystron 1 drives Cavity 1 and Klystron 2 drives Cavity 2.
2. Klystron 1 Drive: Klystron 1 to both cavities, power split via the magic Tee.
3. Klystron 2 Drive: Klystron 2 to both cavities, power split via the magic Tee.
4. Klystron 1 Test: Klystron 1 to dummy loads.
5. Klystron 2 Test: Klystron 2 to dummy loads.
6. Cavity Measurement: Low-power VNA connections to the cavities.

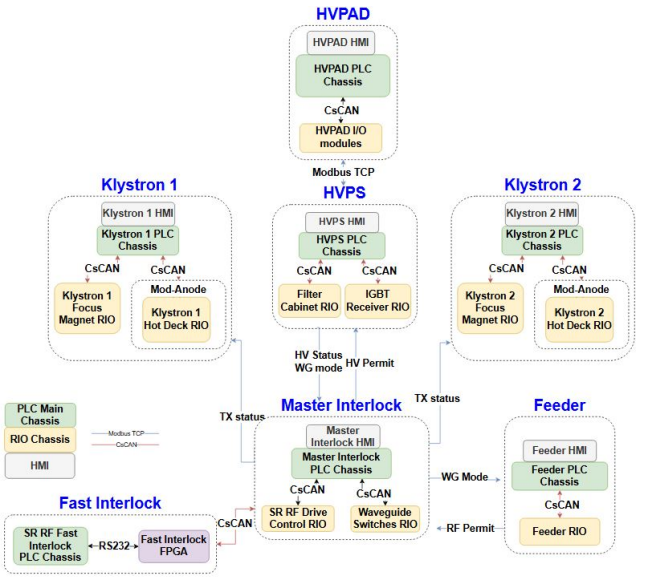



Figure 1: Existing system architecture.

Figure 1 shows the architecture of the existing SR RF PLC-based control system, which comprises the following subsystems:

- High Voltage Power Supply (HVPS) subsystem – Selects configuration mode, controls high-voltage regulation and manages sequencer-driven startup of the klystron high-voltage power supply system.
- High Voltage Pad (HVPad) subsystem – Provides local control of the Variable Voltage Transformer (VVT) position, temperature, AC line, and cabinet monitoring.
- Klystrons subsystem – Manages klystron operating states, monitors inputs, and controls supporting subsystems, including cooling, ion pumps, magnets, and mod-anode circuits.
- Master Interlock subsystem – Aggregates critical signals from multiple subsystems and enforces interlock chains to ensure safe RF operation.
- Feeder subsystem – Monitors water-cooling systems for circulators, circulator loads, and dummy loads, providing operational permits to the Master Interlock PLC based on configuration mode.

* Work supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231

† nusaqib@lbl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD032

# UPGRADING THE ALS BEAMLINE EQUIPMENT PROTECTION SYSTEM: SOFTWARE DEVELOPMENT AND IMPLEMENTATION*

N. Us Saqib†, H. Huang, M. Lopez
Lawrence Berkeley National Laboratory, Berkeley, CA, United States

*Abstract*

The Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory, a third-generation synchrotron light source, has been in continuous operation since 1992. The ALS beamline equipment protection system is being upgraded by replacing legacy end-of-life Modicon PLCs with Allen-Bradley PLCs. This paper presents the software architecture of the upgraded system, highlighting the Allen-Bradley PLC programming approach and the development of the EPICS database that integrates the protection system into the facility's control system infrastructure.

## INTRODUCTION

At Advanced Light Source (ALS), a beamline is defined as the photon delivery system that begins at the storage ring vacuum chamber and extends onto the experiment floor, including the endstations [1]. The three major component groups of a beamline are Front End, Branchlines, and Endstations. Figure 1 illustrates a generalized combination of beamline configurations.

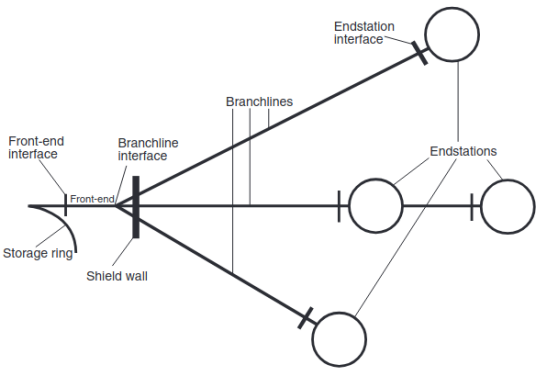



Figure 1: Generalized beamline configurations [1].

The beamline equipment protection system (BL EPS) prevents damage to critical beamline components and safeguard the storage-ring vacuum. The major components in a beamline that interface with the BL EPS are as follows:

- Water-cooled apertures: Transmit full or partial beam width while protecting equipment from thermal load.
- Vacuum valves: Isolate beamline vacuum from the storage ring or endstations for servicing or protection.
- Water-cooled photon shutter: Interrupts synchrotron-radiation beam.
- Personnel safety shutter: Shields beamline or branchline from bremsstrahlung radiation.

Currently, the ALS is undergoing a major facility upgrade known as ALS-U. As part of the ALS-U efforts, a significant portion of the beamline equipment protection system is being relocated, creating an opportunity to modernize the end-of-life equipment. The existing system, based on obsolete Modicon PLC hardware, must be replaced to ensure long-term reliability and maintainability. Allen-Bradley PLCs have been selected for the upgrade to maintain consistency with other ALS-U PLC-based systems. This paper presents the design and implementation of the upgraded beamline equipment protection system, with emphasis on the PLC logic and its integration with EPICS.

## CURRENT SYSTEM ARCHITECTURE

Figure 2 shows the existing BL EPS Modicon PLC sector architecture, implemented as a symmetrical network across the twelve ALS beamline sectors, each following a similar layout. The main BL EPS components in a sector are:

- Front-End PLCs: Two stations, in the accelerator pit and beamline floor.
- Branchline PLCs: Vary by sector configuration.
- End-Station PLCs: Dedicated or combined with branchline PLCs.
- Master PLC: Provides read interfaces for EPICS/LabVIEW, mapping tags from all PLCs.
- MOXA Gateway: Enables write access to PLCs.
- CORBA Server: Links EPICS/LabVIEW with the PLC network

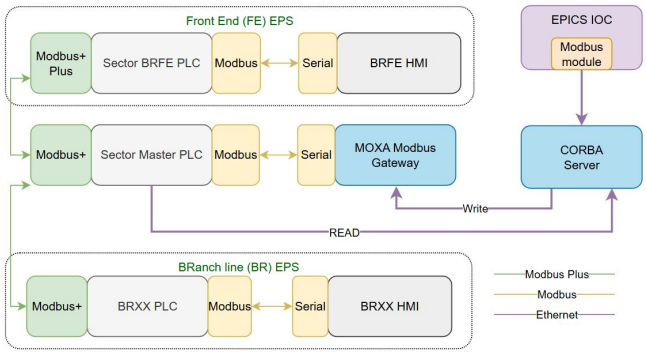



Figure 2: current system architecture.

* Work supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

† nusaqib@lbl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD033

# UPGRADING HARDWARE AND SOFTWARE FOR THE LANSCE PSR ExBPM WITH FPGA BASED HIGH SPEED DIGITIZERS*

T. Ramakrishnan [1†], A. Braido [1], H.Watkins [1], S.A. Baily [1], L. Graham[2]
[1]Los Alamos National Laboratory, Los Alamos, NM, USA
[2]Emerson, Albuquerque, NM, USA

*Abstract*

This paper presents the current approach to upgrade the readout instrumentation for the Proton Storage Ring (PSR) Extraction Line Beam Position Monitor System (ExBPMs) that are located between PSR and Lujan Center target at the Los Alamos Neutron Science Center (LANSCE) accelerator. NI's PXIe platform with high-speed digitizers has been chosen as the hardware which will replace the original CAMAC/NIM/VME system which is obsolete. The beam position algorithm for this project has been demonstrated with the slower real-time Laboratory Virtual Instrument Engineering Workbench (LabVIEW) software using NI-Scope. The software is currently being written and tested using LabVIEW FPGA to make the position processing algorithm match the acquisition speed. The beam position algorithm is designed to handle larger data streams so that it can also be used on the LANSCE Linac BPMs. FPGA design of the algorithm is inherently complex and requires multiple interacting factors such as the LinuxRT OS and supporting packages deployed on the PXIe.

## INTRODUCTION

This work is related to the upgrade of the high-speed digitization of ExBPMs. These Beam Position Monitors (BPMs) receive a large stacked beam pulse from the PSR that is approximately 200 times longer than a typical beam pulse circulated in the LANSCE [1] accelerator. This project gives us an opportunity to develop a state-of-the-art high-speed digitization platform using the NI PXIe platform. This change will give us the opportunity to enhance the capabilities of the ExBPM system and act as the basis of the next generation BPM platform for LANSCE.

### *Current System*

The current ExBPM system relies on obsolete CAMAC/NIM/VME technologies installed at LANSCE between the 1980s and 1990s. Figure 1 shows the external hardware configuration, featuring 50 ohm termination on the left side and the BPM electrode with impedance matching boxes on the right. Figure 2 presents a block diagram of the analog components from the system's original 1991 implementation [2]. Since then, the system underwent partial

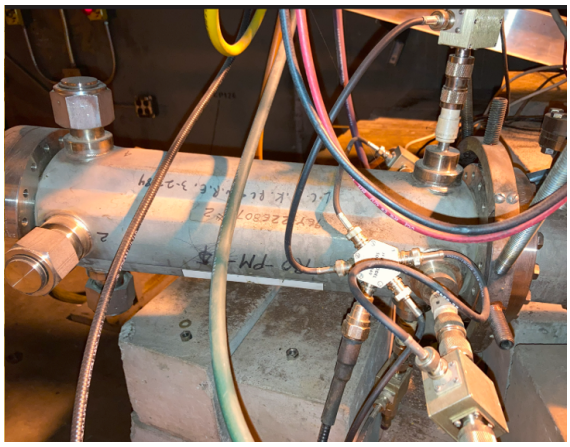


Figure 1: Extraction line BPPM in the field.

modernization, replacing the microVAX and VAX computers with a parallel bus driver and Virtual Machine Environment (VME) Input/Output Controller (IOC). The hardware evolved from terminated stripline to capacitive pickup design, while Nuclear Instrumentation Modules (NIM) continue to perform log-ratio position calculations [3, 4]. Position measurements still rely on a CAMAC delay generator to select precise sample points with nanosecond accuracy.

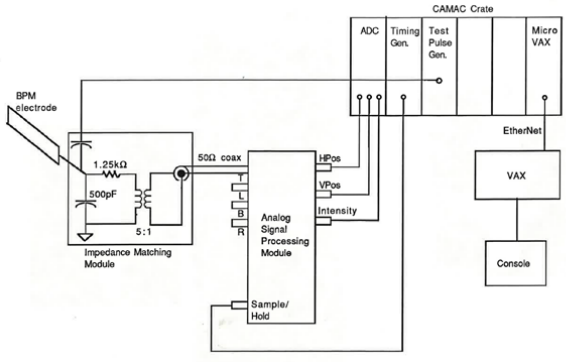



Figure 2: The system that was put in place for the ExBPM in 1991. The current system upgraded the microVAX and VAX to a parallel bus driver with the VME IOC.

### *Proposed System*

Our prototype deployment utilizes the NI PXIe-1088 chassis running LabVIEW Real-Time Linux (LinuxRT) operating system, integrated with a 4-channel PXIe-5764 data acquisition card featuring a Xilinx Kintex KU035 FPGA.

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). LA-UR-25-29139

† tyagir@lanl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD034

# HOMOGENIZING A CONTROL SYSTEM WITH A LONG HISTORY*

S. A. Baily†, C. Hatch, H. Watkins, L. Walker, E. Westbrook
Los Alamos National Laboratory, Los Alamos, NM, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Los Alamos Neutron Science Center (LANSCE) Accelerator is over 50 years old. The original control system was full replaced with an Experimental Physics and Industrial Control System (EPICS) a few years ago. However, the upgrade process was slow, and the "new" control system has processor boards that are 30 years old, as well as new soft-core FPGA processors. We have eliminated VXI and almost eliminated CAMAC crates, but also have VME, VME64, cPCI, and VPX. The control system uses 4 types of VME processor boards and 3 real-time operating systems. The diversity in the control system makes it difficult for new personnel to support everything, as well as making maintenance more difficult. We are in the process of consolidating to fewer hardware platforms and rewriting software that has incurred too much technical debt, while also developing new replacement technologies for upgrade projects. Most of our hardware I/O has been converted to compact RIO based systems. Our timing system cards are either VME or cPCI. VME has had a very long support life, VPX also looks promising for a long support lifetime.

We are currently planning to develop a solution with an IP-based event receiver so that it will be easier to migrate to new commercial off-the shelf FPGA boards. We are also hoping to move away from VxWorks. Newer device support is being converted to operating system independent software using asynPortDriver. Simplifying the hardware and software is the key to sustainable maintenance.

## INTRODUCTION

The LANSCE Accelerator is over 50 years old. The original control system, Remote Instrumentation and Control Equipment or RICE [1], was fully replaced with an EPICS control system a few years ago. However, the upgrade process was slow, and the "new" control system has processor boards that are 30 years old, as well as new soft-core FPGA processors. Working on new upgrades is attractive, and partial deployment of upgrades reduces the maintenance risk for obsolete systems because items removed from service provide a readily available supply of spare parts. While this is very cost efficient in the short term, it poses a long-term risk to maintainability. The mixed system is more complex, needs more documentation, requires more specialized skills, and takes longer to learn. This additional complexity makes it more difficult to maintain and support. Recently, LANSCE has been focusing on simplifying its control system.

## PRESENT STATE

Currently, the LANSCE control system has a large variety of equipment. RICE and associated VAXen have been removed entirely. Our one VXI crate was replaced with a network controlled arbitrary waveform generator. This allowed us to permanently retire the Solaris systems that were used for EPICS 3.13 IOC development. Additional waveform generators were used to replace a CAMAC crate with similar functionality and add features. However, we still have control equipment of varying ages. We have 4 different VME processors (MVME 8110, MVME 6100, MVME 2700, and MVME 167), with the oldest VME processor (with VxWorks 5.5) the most common. We have some cRIOs running VxWorks 6.3 and some running LinuxRT. NI hides many of the differences (although the latest versions of LabVIEW cannot be used), but troubleshooting is still quite a bit different. Nearly all the CAMAC crates have been upgraded. Mostly they were controlling wirescanners, and those have been upgraded to cRIO systems, with individual motor controls and waveform digitizers rather than muxes and triggered I/O. We are working to replace older VME crates with cRIO systems and upgrade VME64 systems to use 8110 boards running VxWorks 6.9. We have many cPCI crates with CR12 processors, but we are investigating running LinuxRT or RTEMS 5, as well as looking for a newer processor that is still available. Realtime performance is particularly important for the low-level RF systems as their adaptive feedforward capability requires obtaining an interrupt from the MRF Event Receiver, reading the event log, and writing flavour information to a register on the bus within 200 µs. Our newest systems run RTEMS 4-11 on a NIOS-II soft-core processor. This means LANSCE runs 3 different real-time operating systems, not counting different revisions of VxWorks. However, as part of the upgrade process, we have managed to reduce the number of different EPICS releases in use to two, R3.15.9 and a 3.15-based custom Data Access EPICS.

## RECENT UPGRADES

Replacing the original RICE system was a significant project that took place over 20 years [2-3]. This system with centralized computing resources and scheduled, triggered readout was replaced with a distributed system of over one hundred cRIO-based IOCS and ~50 FPGA-based IOCS with waveform digitizers and the ability to filter reported data [4]. In addition, many FLECS Fortran programs were converted to Java programs. During this time the accelerator continued to operate, and many other EPICS systems were upgraded or improved as needed, but the

* Work supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). LA-UR-25-28886
† sbaily@lanl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD037

# MODERNIZING CONTROL SYSTEM FOR KLYSTRON TEST STAND AT LANSCE *

J. Quemuel †, A. Waghmare, D. Rai, T. Ramakrishnan
Los Alamos National Laboratory, Los Alamos, NM, USA

*Abstract*

Modernizing scientific test stands is essential for improving data acquisition, control precision, and integration with contemporary research workflows. This paper presents our approach to upgrading the legacy klystron test stand at Los Alamos Neutron Science Center (LANSCE) by implementing EPICS (Experimental Physics and Industrial Control System) for real-time control and monitoring, as well as an overhaul of the diagnostic hardware systems. The transition to EPICS enables scalable, network-distributed control, standardizes communication protocols, and enhances compatibility with the rest of LANSCE's control systems. The improved control system provides intuitive, customizable interfaces for experiment configuration, live visualization, and automated data logging. This upgrade significantly increases maintainability, user accessibility, and automation capabilities, while reducing system downtime and improving experimental reproducibility. The work demonstrates a practical, extensible model for upgrading test infrastructure in research environments where flexibility, openness, and precision are essential.

## INTRODUCTION

Klystrons are a critical component of the LANSCE accelerator facility [1]. They are responsible for providing high power RF to the accelerator, however each run cycle is planned with a certain acceptable klystron failure rate in mind.

When a klystron fails, it is taken to the Electronic Test Lab (ETL) where it is then cleaned and repaired. Once proper repairs have been complete, the klystron is installed at the klystron test stand and run under similar conditions as if it were deployed in the LINAC.

Repairing and redeploying a klystron generally takes about one to two months to complete. This process has been mostly manual in the sense that there had been no overarching control system framework and had procedures dictating manual log entries twice a day [2]. Additionally, the lack of network controls often required personnel to arrive on-site to make adjustments to the system.

In an effort to reduce the time klystrons spend in testing, Accelerator Operations and Technology Instrumentation and Controls (AOT-IC) team and Radio Frequency Engineering (RFE) team have collaborated on modernizing the klystron test stand. This paper details the improvements we have implemented.

* This work was supported by the U.S. Department of Energy and Laboratory Research and Development(LDRD), LA-UR-25-29138

† jkq@lanl.gov

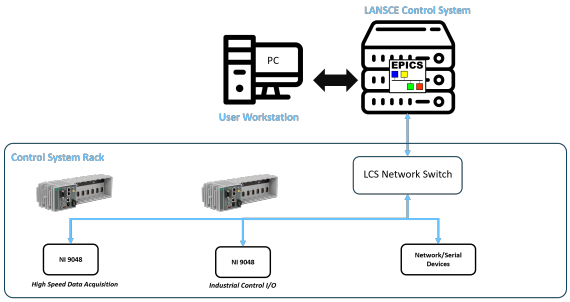



Figure 1: Hardware design overview diagram.

## HARDWARE/DIAGNOSTIC UPGRADES

The hardware implementation adopts a three-platform approach consisting of:

1. Data Acquisition (DAQ)
2. Industrial Input/Output (IIO)
3. Network/Serial Device Integration (NET/IO)

This architecture provides a modular, maintainable, and fully integrated solution for diagnostic waveform capture, process variable monitoring, and network-addressable device control at LANSCE. Figure 1 illustrates the system design.

### *Legacy Hardware Limitation*

Previously, waveform capture was performed using standalone oscilloscopes. While sufficient for single-shot troubleshooting, these instruments were limited by processing cycles, providing only partial or intermittent waveform records. Industrial I/O functionality was handled by a legacy cRIO chassis that was nearing the end of its lifecycle, lacked sufficient channel density, and was not integrated with the LANSCE control system. Additionally, several key diagnostic instruments such as vacuum pump controllers, RF generators, and power meters were operated as standalone units with no centralized monitoring. Together, these limitations reduced diagnostic fidelity, increased operator workload, and compromised long-term maintainability.

### *cRIO for Data Acquisition (DAQ)*

The DAQ chassis is based on the NI 9048 controller, which integrates a 1.30 GHz dual-core processor and 2 GB RAM. For waveform capture, two NI 9223 modules are installed; each provides four simultaneously sampled analog input channels with 16-bit resolution, ±10 V input range, and 1 MS/s per channel throughput. This configuration allows capture of all diagnostic waveform channels with 1 µs time

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD038

# EPICS ARCHIVER APPLIANCE AT LOS ALAMOS NEUTRON SCIENCE CENTER*

J. Quemuel†, S.A. Baily, A. Georgiou, E. Huang, E. Westbrook, N. Yampolski
Los Alamos National Laboratory, Los Alamos, USA

*Abstract*

The Accelerator Operations and Technology Instrumentation and Controls (AOT-IC) group has made the decision to adopt the EPICS Archiver Appliance (EAA) as its primary archiver system. The ability to archive millions of process variables (PVs) has been a topic of interest at Los Alamos Neutron Science Center (LANSCE) as the demand for large data sets and machine learning (ML) workflows increases. The open-source nature of the archiver appliance, the user-friendly web interface, and the proven ability to archive hundreds of thousands of process variables were decisive factors in this choice. This work presents experiences and lessons learned when implementing the EPICS Archiver Appliance at LANSCE.

## INTRODUCTION

At LANSCE, archiving of control system data is essential for efficient troubleshooting, and long-term documentation of the accelerator's operational parameters. As the facility continues to expand its capabilities, our struggle has been the maintenance of a scalable, reliable PV archiver software. EAA [1] is an open-source software tool that meets this demand, offering a flexible, modular architecture tailored towards EPICS-based control environments.

The Archiver Appliance consists of several coordinated components, including modules for real-time data ingestion, data retrieval, data transformation, and system management. It supports tiered storage, configurable retention policies, and a web-based interface for users and administrators. Additionally, it is actively maintained by the EPICS community and allows us the opportunity to collaborate on and develop new features.

We previously relied on an internally developed archiver software to store and manage PV data [2]. While this system served our needs for many years, it grew increasingly difficult to maintain as our requirements expanded. The system's custom architecture, limited documentation, and lack of long-term support hindered troubleshooting and further development.

Beyond addressing maintainability issues, our facility has growing requirements for large data analysis and the application of AI/ML techniques to optimize system performance and reliability. If we were to adopt a new archiver software, it would have to be able to archive and retain the necessary data volume and fidelity to support future AI/ML workflows, including anomaly detection [3] and operational optimization.

This paper presents our experience implementing the EPICS Archiver Appliance at LANSCE. We describe our approach to deployment and configuration, and the system's integration with our operational workflows. We also share performance observations, encountered challenges, and lessons learned during the process. Our goal is to provide practical insights for other facilities considering a similar migration or planning to enable data-driven capabilities through modern archiving infrastructure.

## HARDWARE ARCHITECTURE

Testing and deployment began on a single Dell R640 server. As our user base grew, we commissioned two addition Dell R760 servers to our cluster and have three more which will be added in the future. The following are the specifications for these machines:

- **1x Dell R640 Series**
  - **CPU:** 1x Intel Xeon Gold 6134 @ 3.20GHz (8 cores)
  - **Hard Drive:** 4x 10TB HDD SATA
  - **RAM:** 64GB
  - **Network:** 4x 1GbE, 1x 10GbE
- **2x Dell R760 Series**
  - **CPU:** 2x Intel Xeon Silver 4510 @ 2.40 GHz (12 cores)
  - **Hard Drive:** 24x 16TB HDD SAS
  - **RAM:** 256GB
  - **Network:** 4x 1GbE, 4x 10Gb SFP slot
- **3x Dell R760 Series**
  - **CPU:** 2x Intel Xeon Silver 4510 @ 2.40 GHz (12 cores)
  - **Hard Drive:** 24x 16TB HDD SAS
  - **RAM:** 256GB
  - **Network:** 4x 1GbE, 4x 10Gb SFP slot

### *Network*

To ensure accessibility and usability beyond the confines of the control system's internal network, an NGINX reverse proxy is employed. This service runs on a machine which lies on the controls network, but is also exposed to LANL's Open Collaboration Enclave (OCE) network. User requests

* This work was supported by the U.S. Department of Energy and Laboratory Research and Development (LDRD).

† jkq@lanl.gov, LA-UR-25-29102

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD039

# ASSESSING WinCC OA PROJECT LIMITS TO GUIDE DCS ARCHITECTURE FOR THE Phase-II ATLAS UPGRADE

N. Kanellos*, T. Alexopoulos, National Technical University of Athens, Athens, Greece
P. Moschovakos, S. Schlenker, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

In preparation for the Phase-II upgrade of the ATLAS experiment, detector subsystems that will be upgraded to cope with the new operational conditions imposed by the High-Luminosity Large Hadron Collider (HL-LHC), are required to develop a Detector Control System (DCS) tailored to their specific needs. A key consideration for this upgrade is the size of WinCC OA projects. Understanding how large a WinCC OA project can be, without compromising performance, is critical for ensuring DCS stability and efficiency. This paper presents a series of studies conducted on WinCC OA (v. 3.19) test projects in order to assess its limits, based on the servers that are used within the ATLAS experiment. The findings provide practical insights into the factors that influence system scalability, such as the number of datapoint elements (DPEs), the number of control managers etc. These results aim to support detector DCS groups in planning and optimizing their DCS architectures, helping them decide on the appropriate number and size of WinCC OA projects based on their future operational requirements.

## INTRODUCTION

### *ATLAS DCS Architecture*

The DCS ensures the reliable and safe operation of ATLAS by providing a unified interface for all sub-detectors and the experiment's technical infrastructure [1]. Its architecture is divided into two main categories: the Front-End (FE), which consists of commercial or custom-made hardware devices, and the Back-End (BE), where an industrial Supervisory Control and Data Acquisition (SCADA) software package is used to implement the control system. The communication is realised through various technologies and protocols, with the most commonly used in ATLAS being OPC Unified Architecture (OPC UA), a cross-platform, open-source standard for data exchange in a server-client model.

### *WinCC OA*

WinCC OA is a software package designed for use in automation technology [2]. It provides a comprehensive set of tools and features for real-time monitoring, control, and visualisation of industrial processes. Its operation is based on gathering information from FE equipment and offering supervisory control functions such as data processing, execution of control procedures, alert handling, trending, archiving, and a web interface.

* nikolaos.kanellos@cern.ch

WinCC OA architecture is divided into layers, with each layer consisting of specific units that utilize the necessary functionalities depending on their assigned tasks. These units, that are called *managers*, communicate using a special WinCC OA protocol over TCP/IP.

The central processing unit of a WinCC OA project is called *Event Manager* (EV) and is responsible for maintaining an up-to-date image of all process variables in memory. All other managers, query data from EV's process image rather than communicating directly with each individual unit, responsible for the change, making the EV the central data distributor of a project [3].

This communication is being performed in the form of messages. To ensure smooth project operation, two message queue containers are defined with specific limits on the number of messages that can be sent to and received from the EV. If the number of messages that are sent to EV exceed the limit, the EV becomes saturated. In such case, EV will automatically disconnect from the manager responsible for exceeding the container limit, and that manager will be stopped. The upper limit of message queue container, in the receiving direction of the EV, is defined by the config file entry `maxInputMsgCount` with a default value of 100 000 messages. The upper limit of Buffer Control Manager (BCM) output queue in the sending direction (to all other managers) is defined by the config file entry, `maxBcmBufferSize` with a default value of 10 000 kB [4, 5].

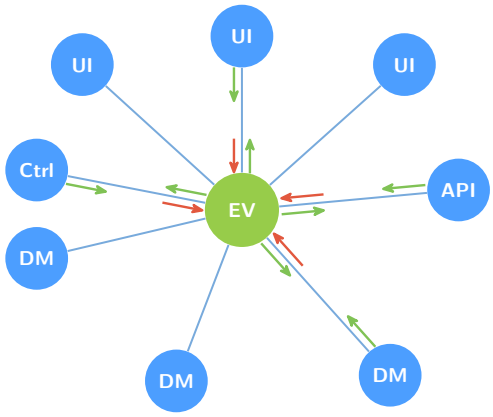



Figure 1: The flow of messages to and from the EV.

### *DCS Usage*

In the simplest approach, a typical project in the DCS context is responsible for monitoring and controlling hardware (such as a high-voltage mainframe, a gas system etc.).

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD049

# SUPERCONDUCTING WIGGLER CONTROL AND QUENCH MONITORING SYSTEM

T. Meng, Y. Ding, S. Wang, Y. Liu†
Shanghai Advanced Research Institute, Chinese Academy of Sciences, Shanghai, China

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The SSRF (Shanghai Synchrotron Radiation Facility) superconducting wiggler consists of four parts: a superconducting multipole magnet, a cryostat system, a magnet power supply and a control system. The superconducting multipole magnet generates a cosinoidal magnetic field profile with a peak value of 4.2 T. The voltage monitoring system can monitor the voltage of each coil section of the superconducting multipole magnet through the voltage taps. Therefore, the voltage trend of each coil section can be recorded during a quench event. The voltage monitoring system collects the voltage data of each coil section through PLC (programable logic controller) analogue input modules. To realize the quench detection, the system records the voltages cycle by cycle and judge by a delay threshold. Based on the PLC system both the equipment control function and the voltage monitoring function are achieved. The quench voltage of each coil section is displayed and discussed.

## INTRODUCTION

The Shanghai Synchrotron Radiation Facility (SSRF) is a third-generation light source. It is based on a 3.5 GeV electron storage ring with a circumference of 396 m [1]. The superconducting wiggler (SCW) is an insertion device used for generating hard X-Rays with superconducting multipole magnets [2, 3]. The magnets of SCW1 in SSRF can generate a cosinoidal magnetic field profile with a peak value of 4.2 T, and the generated magnetic field alternates positively and negatively along the direction of electron motion in the storage ring. The SCW1 superconducting wiggler is the first SCW which is running under long-term normal operating conditions to produce hard X-rays for the experiment stations in SSRF. The device is installed in the BL12 unit of the SSRF storage ring (Fig. 1) [4]. Table 1 shows main parameters of the SCW1.

This paper presents the control system designed for the superconducting wiggler and the quench monitoring results obtained by the system.

The control system has six main functions listed below.

1. Monitor cryostat temperature, helium pressure & level.
2. Control the compressors of helium cryocooler.
3. Monitor the voltage of superconducting coil.
4. Interlock with other systems of the storage ring.
5. Record the high-speed sampling voltage of each coil section when a quench occurs.
6. Record the low-speed sampling voltage of each coil section continuously when the machine is running.

† email address: liuyy@sari.ac.cn

Table 1: The Main Parameters of the SCW1

| Parameters | Value |
|---|---|
| Maximum magnetic field on beam axis | 4.2 T |
| Period length | 48 mm |
| Number of main dipoles | 45 |
| Magnetic length | 1080 mm |
| K-value | 18.8 |
| Critical energy | 34.2 keV |
| Vertical aperture | 9 mm |
| Horizontal aperture | 50 to 70 mm |
| Height | 2621 mm |
| Length | 2496 mm |



Figure 1: SCW1 in SSRF storage ring.

## SYSTEM HARDWARE DESIGN

### *Hardware Architecture*

The superconducting wiggler control system architecture diagram is shown in Fig. 2. The signal cables for cryogenic temperature sensors, magnets voltage sense wires, helium level sensors pass through the vessel of the cryostat via feedthroughs due to the vacuum environment inside the cryostat. The cables are connected to the terminals in the terminal box near the device in the tunnel. Then other cables connecting to the Instruments and Control Cabinet (ICC) are connected to the other side of the terminals. The helium pressure sensor and the vacuum gage are connected

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD050

# STATUS OF THE CONTROL SYSTEM FOR THE DELTA SYNCHROTRON LIGHT SOURCE

D. Schirmer*, A. Althaus, A. Erpelding
Center for Synchrotron Radiation (DELTA), TU Dortmund University, Dortmund, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The 1.5 GeV electron storage ring DELTA, operated by the University of Dortmund in Germany, celebrated its 30th anniversary last fall. Over the past three decades, the control system has undergone many different IT infrastructure development cycles. It was commissioned between 1994 and 1998, utilizing a series of command-line-based in-house applications that operated on individual, low-performance networked HP workstations and VME-based real-time CPUs, initially without the support of graphical user interfaces (GUIs). These GUIs were gradually implemented later with the introduction of the EPICS software package (1999-2001). Based on a combination of EPICS and a newly installed Linux PC-based client/server architecture, a variety of software tools and hardware extensions were introduced in the following years. Today, the DELTA control system utilizes an open-source virtual environment with a server management platform that integrates kernel-based virtual machines (KVM), software-defined storage and network functions on a single platform. In addition, web-based user interfaces simplify the configuration of the integrated disaster recovery tool and enhance the management of high availability and redundancy within the server cluster. Furthermore, machine learning algorithms have been incorporated into the control and optimization of the storage ring. This report gives a historical review, summarizes the most important developments in recent years and provides an outlook on future projects.

## INTRODUCTION

The DELTA synchrotron light source, which is operated by the TU Dortmund University in Germany, consists of a linear pre-accelerator (70 MeV), a full energy booster synchrotron (70 MeV to 1.5 GeV) and a storage ring with 115.2 m circumference (see Fig. 1). The maximum stored beam current amounts to 140 mA (nominal multibunch operation) with a lifetime of about 60 hours (at 100 mA). The facility is operated 24 hours a day, 5 days a week, where 2000 hours/year were used for synchrotron radiation (SR) studies and 1000 hours/year for maintenance work and accelerator physics research. To save electricity costs, accelerator operation has been cut by 50% since 2023 [1–3].

The control system of the storage ring facility is based on the open-source framework EPICS (Experimental Physics and Industrial Control System [4]) and comprises a typical three level client/server architecture. A field bus level, a real-time control level, and the process control level. Over the past 30 years, the systems have been continuously adapted to the latest hardware and software developments.

* detlev.schirmer@tu-dortmund.de

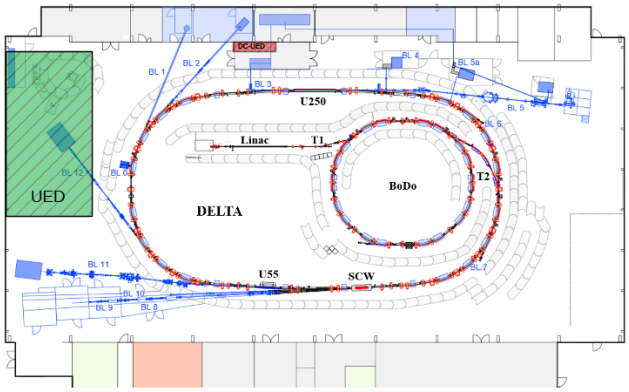



Figure 1: DELTA floor plan: The green area shows where the upcoming UED project will be located. The red area depicts the site of the currently installed 100-keV DC-UED demonstration project. The SR beamline infrastructure is shown in blue [1, 3].

## HISTORICAL REVIEW

From 1999 to 2001, the control system underwent a step-by-step transformation, changing from a non-standard, in-house developed system to one based on the EPICS framework [5]. At that time, it consisted mainly of a few HP-UX-based workstations and about two dozen VMEbus single-board computers (Force/MicroSys CPU types) running the real-time operating system (OS) VxWorks [6]. They were primarily connected to accelerator input/output (I/O) devices via control area network (CAN) nodes (ESD digital/analog I/O modules [7]) or a general purpose interface bus (GPIB). Initially, all computers used a shared 10 Mbit/s 'Thick Ethernet' network consisting of a 10 mm thick, rigid, yellow coaxial cable (10Base5).

The HP-UX workstations were replaced successively by standard Linux-based PCs running the EPICS client software, along with a migration to VME-PowerPC CPUs operating as EPICS servers. As a result, DELTA became one of the first storage rings to operate a synchrotron radiation facility based on Linux PCs in conjunction with the EPICS framework.

Since October 2001, all ramped magnet power supplies of the booster synchrotron (dipole, 6 quadrupoles, 2 sextupoles, steerers), RF power, and all booster diagnostics (beam position measurement, tune control, beam loss measurement, power supply currents, etc.) were handled by VME-based digital signal processing (DSP) boards developed at DELTA ('DeltaDSP'). These DSP cards were designed with generic EPICS driver/device support [8]. This setup has not been modified and operates flawlessly up to now.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD052

# THE WHITE RABBIT COLLABORATION: AN INNOVATIVE MODEL OF PUBLIC-PRIVATE PARTNERSHIP

J. Serrano*, I. Caño López[1], A. Díez Fernández, B. Frisch, E. Gousiou,
M. Lipiński, D. Tacchini, A. Wujek, CERN, Geneva, Switzerland
[1]now an independent researcher

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

White Rabbit (WR) is an open-source synchronisation technology developed at CERN in collaboration with other institutes and industry. It is commercially available from multiple vendors, and its adoption in industry and academia grew significantly after its standardisation as part of the Precision Time Protocol (IEEE 1588-2019). This resulted in a substantial increase in the number of feature and support demands.

To support this broader adoption, ensure the continued development of the open-source core, and establish effective governance over the evolving technology, the White Rabbit Collaboration (WRC) was formed. WRC members, through an annual fee, contribute to funding the CERN-based Bureau, responsible for maintaining WR's core components, providing support to members, and collaboratively shaping the technology's future.

The WRC represents an innovative model of public-private partnership and knowledge transfer of an open-source technology and can serve as a template for similar initiatives. This paper will describe WRC's establishing and first year of operation.

## INTRODUCTION

White Rabbit (WR) is a technology that allows different pieces of equipment to synchronise very precisely with one another. By synchronisation, we mean that electronics in different physical locations, potentially many kilometres apart, share a common notion of time within some uncertainty, which in the case of a WR system is guaranteed to be below 1 ns.

WR-compliant equipment is connected using optical fibres, forming a WR network. WR switches (henceforth referred to as "switches") ensure that users can build networks that can scale easily. A port in a switch can connect to another switch or to a WR node (referred to as "node" from now on). By combining switches and nodes, one can build large control and data acquisition networks in which all nodes are precisely synchronised.

This paper does not aim at describing the mechanisms used by WR to achieve high-accuracy and high-precision synchronisation. The interested reader will find many references on the WR website [1]. Instead, we focus on how the WR project and community have evolved to maximise the impact of this technology outside of physics laboratories while ensuring that all hardware, gateware[1], firmware, software and documentation are published under appropriate open-source licences.

## MAXIMISING THE DISSEMINATION OF KNOWLEDGE AND ITS IMPACT

CERN, the European Organization for Nuclear Research, is the world's largest laboratory for particle physics. It pursues a four-fold mission: a) to perform world-class research in fundamental physics, b) to provide a unique range of particle accelerator facilities that enable research at the forefront of human knowledge, c) to unite people from all over the world to push the frontiers of science and technology, and d) to train new generations of physicists, engineers and technicians, and engage all citizens in research and in the values of science [2].

Similar to many other public institutes and laboratories, CERN has a strong dissemination mandate. It publishes its scientific results, along with methods and tools developed in the course of its research, following Open Science best practices. This is formalised in its Open Science Policy [3]. It simultaneously seeks to maximise the technological and knowledge return to the Member States and promote CERN's image as a centre of excellence for technology, prioritising societal impact over revenue creation [4]. This mandate is managed at CERN by the Knowledge Transfer (KT) group.

A policy that prioritises societal impact requires the application of a wide array of dissemination methods through a variety of commercial and non-commercial strategies, adapting the management of CERN's intellectual property (IP) to the chosen strategy. CERN licenses proprietary IP, makes many of its technologies available through open-source licences, enters collaborations, service and consultancy partnerships, supports spin-off and start-up companies, and enters international collaborations such as through EU projects [4].

Selecting the appropriate strategy can be challenging. Proprietary IP allows to manage the access to the technology and can create incentives for investments that facilitate its industrial uptake and its dissemination into products, yet it prevents widespread access to the technology itself. Conversely, open-sourcing a technology provides widespread access, yet it might reduce the competitive advantage of an industrial partner if competitors can easily access the same technology. It can also create a support and maintenance overhead in the CERN team responsible for the open-source

* Javier.Serrano@cern.ch

[1] Field-Programmable Gate Array (FPGA) designs using Hardware Description Languages (HDL)

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD053

# DEVELOPMENT OF MTCA.4 AT ESS

F. Chicken, J.J. Jamroz, G.S. Fedel, J.P.S. Martins, European Spallation Source ERIC, Lund, Sweden

*Abstract*

As commissioning activities for the Linear accelerator continues at the European Spallation Source (ESS), it is highlighting important issues and development features for the MTCA.4 systems used in several key subsystems. With over 300 MTCA.4 systems deployed, it has proved a valuable platform for fast controls. The use of MTCA.4 has also allowed for novel solutions to provide flexible but reliable controls for the facility, including expanding existing controls design and developing orphan scope. However this is not without its issues, as MTCA.4 for the scientific community is still very much on the smaller scale and reliant on only a few suppliers, with risk of obsolescence of components and changes to the standards. This makes it challenging but fun to continue developing the hardware to ensure the desired reliability for the Linac controls. This paper will describe the procedures and processes that the hardware team has taken to mitigate these issues, the current status of our MTCA infrastructure and the long term plans for its maintenance.

## MTCA.4 STANDARD OVERVIEW

### *Foundation: MicroTCA*

MTCA [1] is based on the AdvancedTCA technology and uses AMC (Advanced Mezzanine Card) as its fundamental modules. The main components are the chassis (crate), which hosts AMC cards, and the MCH (MicroTCA Carrier Hub), which acts as a switch and manages communication. The standard suggests Ethernet and PCIe as the main communication backplanes, complemented by M-LVDS trigger lines and LVDS clock distribution lines.

### *MTCA.4 Extension – Support for Scientific Applications*

MTCA.4 was developed mainly by DESY [2] (Hamburg) for use in control and data acquisition systems in accelerators (e.g., XFEL). This standard, together with the AMC.0 standard from PICMG, recommends the usage of PCIe and Ethernet as the main technologies for communication over the backplane, which is the most common scenario for MTCA products nowadays. It also recommends the topology of the trigger distribution over the backplane. MTCA.4 is a standard designed specifically for science and high–technology industries, extending the baseline MTCA with the possibility of using Rear Transition Modules (RTM). It enables high–quality analog and digital signal handling. An example of an AMC and RTM is shown in Fig. 1.

### *MTCA.4 Desy AMC/RTM Recommendations*

- **Class A** – pin assignment for analogue applications.

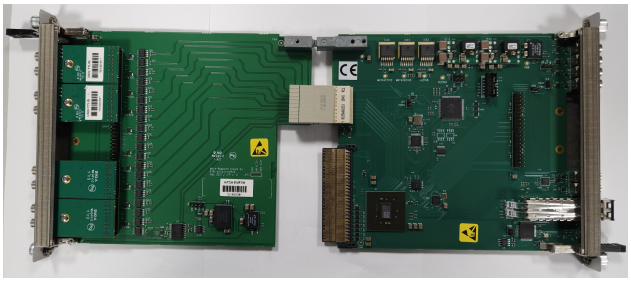


Figure 1: Showing how the AMC and RTM are connected.

- **Class RF** – pin assignment for high performance analogue applications via two 6-pin coaxial connectors.
- **Class D** – pin assignment for digital applications.

## MTCA AT ESS

ESS is progressing with commissioning activities as part of the planned testing schedule ahead of expected first science in 2027. Installation of all 300 needed MTCA systems for beam on target has been almost completed, and the control systems have been used throughout the testing phases. The current setup of MTCA at ESS is a base configuration comprising of either a 3U or 9U chassis, a power supply, a MCH - NAT-MCH-PHYS, single board computer from Concurrent Technologies AM-G64 (CPU), and a MicroResearch Finland Event Receiver (EVR). The AMC, FPGA Mezzanine Card (FMC) and RTMs' are selected based on the application the MTCA system is being deployed for. ESS is currently using various AMCs (as demonstrated in Fig. 2), FMCs and RTMs from a small number of manufacturers: CAENels, Concurrent Technologies, IOxOS Technologies S,A and Struck Innovative Systems.

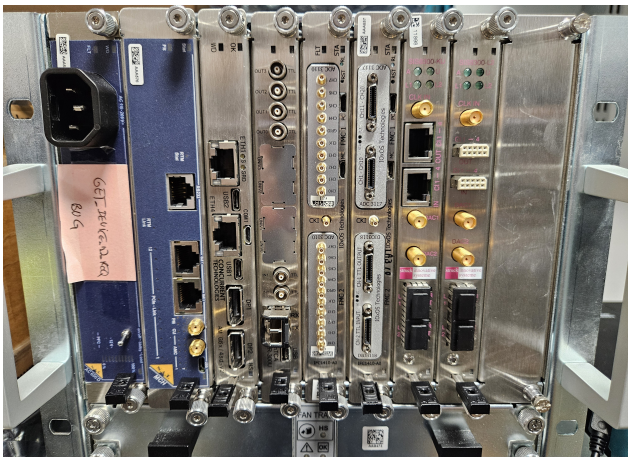


Figure 2: Examples of AMCs in use at ESS.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD055

# REFURBISHMENT STRATEGY OF THE GANIL CYCLOTRONS CONTROL SYSTEM*

C. Haquin†, O. Delahaye, C. Treangle, GANIL, Caen, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Grand Accelerateur National d'Ions Lourds (GANIL) began operating in 1983, since the scientific demand is still there, a wide renovation project called Cyclotrons Renovation (CYREN) is ongoing, aiming at maintaining the accelerator up and running for at least two more decades. One of the work package of the CYREN project is related to the GANIL Cyclotrons control system (GCS). Current control system date back to the early 90's, it relies on Versa Module Eurocard controllers, software were developed with ADA language and MOTIF/XRT services for graphical user interface (GUI). The MOTIF/XRT technology is more and more difficult to operate with latest Linux distribution and hardware generation, user interfaces are difficult and time consuming to maintain compared to state of the art GUI development environment, input output VME cards are less available. The project has decided to take measures against these risks. This paper will present the strategy defined and the status the refurbishment.

## GANIL ACCELERATOR FACILITY

GANIL research center actually hosts two accelerators, the first one, called original GANIL, is the accelerator from the 80's on which the purpose of this paper, the second one is the SPIRAL2 accelerator with its superconducting LINAC and EPICS control system.

Original GANIL accelerator relies on two ions sources, two compact cyclotrons (C01, C02) for injection of the beam, two separated sectors cyclotrons (CSS1, CSS2) for the main acceleration (both cyclotrons are identical with a stripper in between to enable acceleration with the second cyclotron), and a last cyclotron (CIME) for post-acceleration when exotic nuclei are produced by sending the primary beam in the SPIRAL1's target/source system, see Fig. 1. About 2000 devices must be controlled to operate the GANIL accelerator.

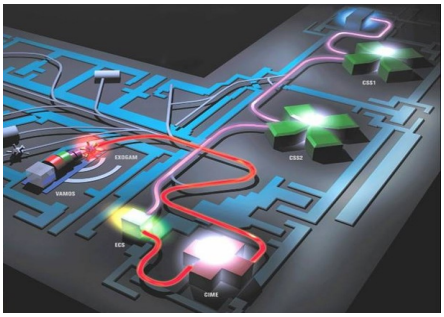



Figure 1: Beam path though GANIL accelerator.

## CONTEXT & MOTIVATIONS

In 2023, GANIL celebrated the 40th anniversary of its first beam. With strong demand from physicists, operations must continue for at least two more decades. To ensure long-term sustainability, an ambitious renovation project is underway to mitigate obsolescence and modernize the control system.

The current system, though still operational, relies on aging VME hardware and obsolete graphical libraries, while Linux platforms evolve rapidly. This combination increases technological debt and undermines stability. In addition, graphical interface development is slow and inefficient due to the absence of modern frameworks and tools.

The renovation project therefore aims not only to remove obsolescence risks but also to introduce a modern development and testing environment, reducing maintenance needs and ensuring reliable operations for the future.

## LEGACY CONTROL SYSTEM

The original GANIL control system is a combination of hardware and software components that allows operators to manage, from the control room, the equipment of the ion sources , the 5 cyclotrons and all beamlines up to the entrance of the experimental halls. It was developed in the early 90's and relies on:

- 31 Versa Module Eurocard (VME) real time controllers (RTC) running with VxWorks 6.9.4.12 operating system.
- 10 real-time applications (RT-app): each RTC executing one instance of these RT-app.
- "equipment" Data Base (EDB) to store equipment description and settings. A dedicated software extracts data from EDB to produce a configuration file for each RTC.
- 230 VME cards: 150 are off-the-shelf (OTS) cards and are of 8 kinds. remaining 80 cards are of 5 kinds but are home made cards embedding specific data acquisition and equipment control.
- 54 GUI using MOTIF/XRT services to build graphics.
- "parameters" Data Base (PDB) to store accelerator tuning values. GUI can read from or write into PDB.
- a home made TCP request standard between GUI and RTC
- 6 software hub between GUI and RTC.
- 1 alarm server with its dedicated database.
- 1 archiving server.

All the software was developed with ADA language (ADA 2012), see architecture diagram Fig. 2.

* Workpackage of the CYREN project.

† haquin@ganil.fr

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD058

# KEK ELECTRON/POSITRON INJECTOR LINAC CONTROL SYSTEM

M. Satoh†,[1,2], D. Wang[1,2], I. Satake[1], M. Sakamoto[1], F. Miyahara[1,2], T. Kudou[3], S. Kusano[3], H. Saotome[4], K. Hisazumi[3], Y. Mizukawa[3], T. Ohusa[4], T. Suwada[1], and K. Furukawa[1]
[1]High Energy Accelerator Research Organization (KEK), Accelerator Laboratory, Tsukuba, Japan
[2]The Graduate University for Advanced Studies (SOKENDAI), Accelerator Science Program, Tsukuba, Japan
[3]Mitsubishi Electric System & Service Co., Ltd (MSC), Tsukuba, Japan
[4]Kantou Information Service (KIS), Tsuchiura, Japan

*Abstract*

Since May 2019, the KEK electron/positron injector Linac has successfully carried out simultaneous top-up injection into four independent storage rings as well as a positron damping ring. Ensuring long-term stable beam operation under such a complex operational scheme requires maintaining a high availability of the entire control system. The original Linac control system, developed in the 1990s, was built on a proprietary in-house software framework utilizing Unix servers, VME bus systems, and PLCs. The operator interface was implemented with the Tcl/Tk scripting language and supported approximately 3,500 control points.

To improve development efficiency and ensure greater compatibility with other accelerator control systems, the Linac control system has been gradually migrated from its legacy platform to an EPICS-based architecture. This migration has expanded its capability to approximately 200,000 control points. Furthermore, Linux-based hyperconverged servers, a PXI bus system, and custom embedded front ends have also been deployed. This paper presents a comprehensive overview of the current status of the KEK injector Linac control system and discusses prospects for its future development.

## INTRODUCTION

The KEK electron/positron injector Linac began operation in 1982 as a dedicated injector for the PF ring. Since then, it has provided electron and positron beams to various rings at different energies, including TRISTAN, PF-AR, and KEKB. Initially, beam injection into each ring was performed in a time-sharing mode based on scheduled injection. However, the demand for top-up injection increased among PF users since top-up injection was already common operational mode at third-generation light sources. There was also a strong request from the KEKB rings and Belle detector group to implement top-up injection for increasing the integrated luminosity. In response to this situation, a project team was launched in 2004 with the goal of achieving simultaneous top-up injection into all rings. In April 2009, simultaneous top-up injection into three rings (the KEKB electron/positron rings, and the PF ring) was successfully realized. Subsequently, with the construction of a dedicated injection beam transport line for PF-AR in 2012, simultaneous top-up injection into four rings (the SuperKEKB electron/positron rings, the PF ring, and PF-AR) was achieved in May 2019 [1].

Currently, electron and positron beams generated at a maximum repetition rate of 50 Hz can be injected into any of the rings according to arbitrary preprogrammed injection patterns. To realize simultaneous top-up injection, a large number of developments and upgrades have been conducted, including the event based timing system [2], the low emittance photocathode rf electron gun with the complex laser system for SuperKEKB injection [3], a flux concentrator for positron capture system [4], high precision beam position monitors [5], and pulsed magnet systems [6, 7]. Currently, various modifications are being undertaken to further improve injection efficiency into both the light source and the SuperKEKB rings [8-17]. With the increasing complexity of beam operation at the injector Linac, the beam control system has also been steadily improved.

## LINAC CONTROL SYSTEM

### *Overview*

The injector Linac control system adopts a three-tier architecture that conforms to the standard client/server model. Often referred to as the standard model of control systems, it consists of the client layer corresponding to the Operator Interface (OPI), the server layer serving as the middleware, and the local control layer that directly controls the equipment. By employing this architecture, modifications to equipment at the local control layer can be absorbed by the middleware, thereby avoiding changes to the client layer. This design ensures flexible expandability of the entire system.

In the late 1990s, at the beginning of injection into the KEKB rings, the injector control system was developed based on a set of in-house libraries utilizing Remote Procedure Call (RPC). At the OPI layer, a Command Line Interface (CLI) using shell scripts and a Graphical User Interface (GUI) using the Tcl/Tk scripting language were employed.

Around the mid-operation period of KEKB, these legacy systems were gradually migrated to a control system based on the Experimental Physics and Industrial Control System (EPICS) [18]. As the KEKB rings had employed an EPICS-based control system from the outset of their operation, this migration improved compatibility between the injector Linac and ring control systems. Consequently,

†masanori.satoh@kek.jp

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD061

# ENSURING RELIABILITY AND FEASIBILITY OF SOFTWARE ALARMS WITH PHOEBUS IN EPICS NETWORK FOR ALS-U*

S. Ryu†, J. H. Lee, Lawrence Berkeley National Laboratory, Berkeley, CA, United States
S. I. Lee, Osprey Distributed Control Systems, Ocean City, MD, United States

*Abstract*

The Advanced Light Source Upgrade (ALS-U) project requires a reliable and efficient alarm system. This presentation examines the reliability and feasibility of a software alarm system implemented using PHOEBUS within the EPICS network. We will discuss its architecture, configuration strategies, and management techniques. Testing results highlight the system's robustness. Furthermore, we introduce a compact software demo environment for public use, offering key insights for comparable high-reliability environments.

## INTRODUCTION

The ALS-U project requires a robust and efficient alarm management system due to its complex operations and the critical nature of maintaining optimal performance and safety. The alarm management system needs to serve multiple purposes: ensuring early detection of potential issues, facilitating rapid responses to prevent operational disruptions, and safeguarding sensitive equipment. With an anticipated control system encompassing over 200,000 process variables (PVs), of which more than 50,000 require alarm configurations, the scale presents challenges in operation and management. To meet these requirements, we introduce the Phoebus Alarm System (PAS). Developed collaboratively over years by the EPICS community, PAS encompasses accumulated experiences in alarm handling from facilities worldwide [1]. Despite its maturity and widespread development, PAS has no official record of comprehensive performance verification or full-scale site deployment. Prior to actual deployment at ALS-U, it is necessary to evaluate the PAS in aspects of user operation, service management, and critically, service stability and performance under severe network conditions. This paper presents a comprehensive approach to alarm management for ALS-U, addressing three key areas: (1) an operational-level alarm philosophy that extends beyond device-level configuration to reduce operator cognitive load, (2) a validated PAS demo package that demonstrates system reliability under extreme load conditions, and (3) a hierarchical filtering methodology implemented through EPICS record logic.

* Work supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

† soo.ryu@lbl.gov

## ALS-U CONTROLS ALARM MANAGEMENT CONCEPTS AND PAS SOLUTION

The alarm management concepts for ALS-U described in Ref. [2] establish both operational and technical requirements that extend beyond simple threshold monitoring. Table 1 demonstrates how PAS addresses these requirements through specific implementation features. The operational requirements define how alarms should behave from the operator's perspective. The technical requirements ensure the system can reliably deliver these operational capabilities. The mapping shows that PAS provides not just the basic alarm functionality, but a complete ecosystem supporting the full lifecycle of alarm management from configuration through operation to historical analysis.

## PHOEBUS ALARM SYSTEM ARCHITECTURE, DEMO PACKAGE, AND PERFORMANCE VALIDATION

### *System Architecture Overview*

The Phoebus Alarm System represents a mature, production-ready solution for large-scale accelerator alarm management. As depicted in Fig. 1, The PAS consists of Java middleware components customized using a suite of open-source technologies, including EPICS, CSS/PHOEBUS, Kafka, Elasticsearch, and Git systems. This architecture employs the properties of EPICS process variables (PVs) to initiate software alarm events [3], while organizing and managing alarm events and their histories through dedicated Java applications integrated with Kafka and Elasticsearch services.

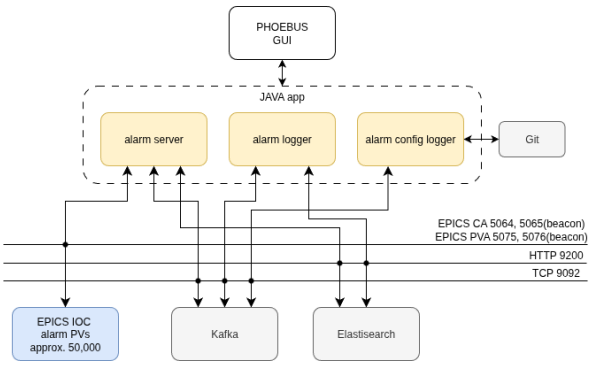



Figure 1: Phoebus Alarm System overview: network connectivity between the components and services

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD067

# A REMOTE-CONTROLLED HIGH VOLTAGE POWER SUPPLY USING OPEN-SOURCE FIRMWARE MODULES

R. E. Spirin*, M. M. Donatti, P. H. Nallin, G. R. Cruz and F. H. Cardoso,
Brazilian Synchrotron Light Laboratory, CNPEM LNLS, Campinas, Brazil

*Abstract*

The High Voltage Power Supply (HVPS) is a programmable, high voltage precision power supply developed at the Brazilian Synchrotron Light Laboratory (LNLS) for use in beamline instrumentation and experimental setups. Specifically designed for integration with ionization chambers, the HVPS provides stable unipolar output voltage up to 5 kV optimized for load regulation and low ripple, ensuring high reliability and measurement precision. The system supports multiple control interfaces, allowing flexible operation in local or remote environments, where real-time voltage and current monitoring is available via a built-in LCD display, analog outputs, and remote monitoring with internal analog-to-digital converters. A dedicated embedded firmware, based on an RTOS, has been developed. It ensures remote configuration, monitoring capabilities, and firmware update over the TCP/IP stack, leading to a full integration with any control system standard and architecture. This paper describes HVPS software, hardware and firmware topology, as well as first use cases at Sirius, the 4th generation light source.

## INTRODUCTION

The HVPS instrument is designed for use in laboratory bench setups, test applications and experimental stations. It provides unipolar high-voltage DC output with excellent load regulation (up to 5 kV with 1 V resolution and up to 1.2 mA current output) and low ripple (823.65 µV). In synchrotron facilities beamlines, high-voltage power supplies play a critical role in the operation of ionization chambers [1]. These detectors are used to measure photon beam intensity and are based on the principle of gas ionization induced by high rate photons in X-ray range [2, 3]. To generate the electric field necessary for collecting ions produced during ionization process, a stable high voltage—typically ranging from few hundreds to thousands volts-is applied to the chamber electrodes. In the context of synchrotron sources, voltage stability and low ripple are particularly important, as they directly affect the accuracy of radiation intensity measurements [4]. Moreover, the ability to finely adjust the voltage allows optimization of the chamber response depending on the gas type, detector geometry and the energy of the incident radiation [5–7]. In addition, the HVPS is commonly integrated into the experiment control system, enabling automation, measurement reproducibility and data archiving. The HVPS is already integrated into QUATI (Quick X-Ray Absorption Spectroscopy for Time and space resolved experiments) and EMA (Extreme condition Methods of Analysis) beamlines at Sirius facility [8,9] with planned installation on the SAPE (Angle-resolved PhotoEmission) and IPE (Inelastic scattering and PhotoElectron spectroscopy) [10] beamlines. In this case, an ethernet interface is considered for the beamlines's EPICS control system [11], providing remote control and diagnostics. The present manuscript describes the main characteristics of the device, outlines its operation modes, and provides general information about instrument usage.

* E-mail: roman.spirin@lnls.br

## DEVICE DESCRIPTION

The general block diagram is shown in Fig. 1. The device consists of several main parts: internal power supplies, the HV module, controller, interlocks, user interface elements, and communication interfaces. There are three operating modes: local, remote, and external. Any mode can be selected using a slide switch on the front panel. In local mode, the user can set the output voltage and current, as well as operate the device, using the navigation buttons on the front panel. In remote mode, the navigation buttons are disabled, and all operations are performed via Ethernet (TCP/IP), RS-485, or USB interfaces. In external mode, the HV output is enabled manually using the front panel button, while the voltage is set via an analog voltage input. An LCD display on the front panel shows the current status of the device and displays the measured output current and voltage values. Voltage and current analog monitor signals are also available in any operating mode via SMA connectors. On the rear panel, there are two interlock input channels designed as dry contacts and output relay with both NC and NO contacts in order to signaling enabled output to external systems.

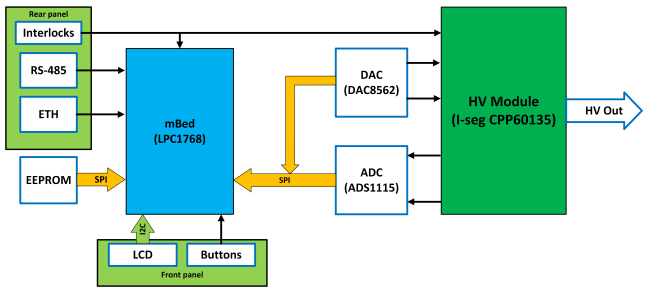



Figure 1: General block diagram.

## MECHANICAL DESIGN

The HVPS is designed to be accommodated in an extruded aluminium enclosure Hammond 1457U2201 (220 mm x 191 mm x 68 mm) [12]. The single 2-layer printed circuit

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD068

# FIRST PICOSECOND DIAGNOSTIC OF THE LASER MEGAJOULE TIMING SYSTEM

T. Somerlinck, CEA CESTA, Le Barp, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The "Laser MégaJoule" (LMJ) timing system has been fully provisioned at the end of 2022. With 19 bundles, 120-beam laser, at the end of 2025, the operational capabilities of the LMJ facility are increasing gradually until its full completion by 2027. At present, there is no global control and measurement of the synchronization system. It's up to each individual subassembly to check that the trigger or fiducial signals are correct. To ensure that the LMJ synchronization system works perfectly, a new picosecond diagnostic was developed and tested on several 1 $\omega$ and 3 $\omega$ laser bundles in comparison with central chamber synchronization and Master clock. This article presents the LMJ timing system with a description of the measurement diagnostic prototype, the main values measured and an overview of the future deployment of this diagnostic on all 22 laser bundles.

## INTRODUCTION

The LMJ facility is a high-power laser designed to deliver about 1.4 MJ of laser energy to target for high energy density physics experiments, including fusion experiments [1]. This energy is produced by 176 laser beams gathered in quadruplets of 4 beams. Each quadruplet is equipped with an Arbitrary Waveform Generator (AWG) that generates the desired temporal pulse shape (lasting typically 3 ns). Synchronization of LMJ's 176 laser beams is crucial to compress symmetrically the millimeter-size target in order to ignite the deuterium and tritium filled capsule. The most demanding experiences need to synchronize the quadruplets to better than 40 ps rms despite the fact that the quadruplet laser sources are separated within the building by several hundred meters. In addition to laser beams synchronization, the LMJ timing system is in charge to deliver, with the same or lower accuracy, two kinds of signals: fiducials for both temporally marked signals and plasma diagnostics, and triggers signals for various devices (sources, amplifiers, Pockels cells, diagnostics...).

The synchronization is therefore one of the most important components for shot experiment, from the laser sources to the target inside the chamber, as shown in Fig. 1.

The calculation of the error budget to reach the 40 ps rms specification has showed that the requirements for the timing systems [2,3] were, for the most precise delay generator, less than 5 ps rms jitter between 2 outputs and 10 ps p-p drift / 1 month.

However, there is no diagnostic available to monitor and control the synchronization signals generate by the timing system using delay generators. Nevertheless, when a synchronization issue is observed, it is very important to quickly identify the problem and repair it.

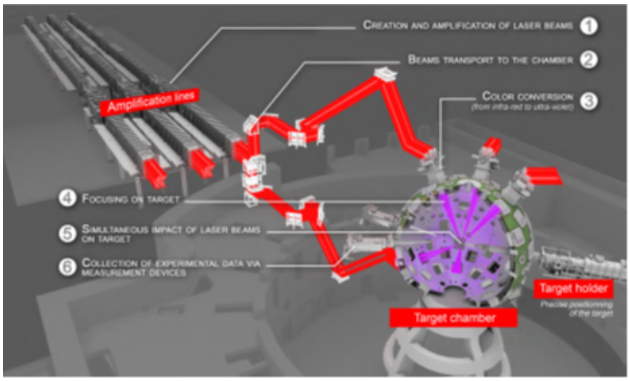



Figure 1: From laser amplification to experiment in the chamber.

Until now, the only solution was to check the signal from the timing system on each delay generator, which was possible using the T0 reference output. However, the main drawback was that it was very difficult to compare this signal with another synchronization signal from another delay generator (which could be located far away in the facility). In the same way, the overall evolution of the LMJ timing system is not really monitored over time.

So a synchronization diagnostic system was therefore studied with the following specifications:

- Monitor the absolute position of each 1 $\omega$ and 3 $\omega$ laser synchronisation signal
- Monitor the absolute LMJ reference Master signal
- Minimize cost
- Minimize the amount of equipment (fiber optic, electrical cable, digitizer...)

The principle is based on the use of reference signals from each precise delay generator (T0 output), which are then multiplexed before being grouped together on a single oscilloscope (~70 signals on 6 channels). Once the experimental phase has been completed and validated, in particular to determine feasibility and the expected level of accuracy, this diagnostic system must be deployed on all 22 laser bundles.

## LMJ TIMING SYSTEM ARCHITECTURE WITH PICOSECOND SYNCHRONIZATION DIAGNOSTIC

The architecture of the synchronization system (Fig. 2) is based on:

- The delay generators and their reference output (T0) which is the "Time 0" for each output (not programmable),
- The natural delay of the optical fiber between all generators due to their position in the LMJ facility,

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD070

# Evolution of White Rabbit Network for large-scale deployment at CERN

M. Suminski†, I. Kozsar, M. M. Lipinski, J. Palluel, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The White Rabbit (WR) network at CERN is undergoing a significant expansion with the introduction of a new branch dedicated to General Machine Timing (GMT), Beam Synchronous Timing (BST), Safe Machine Parameters (SMP) and Radio Frequency (RF) distribution.

Prior to 2024, the WR network consisted of approximately 35 switches and 70 nodes. The new deployment aims to add nearly 200 switches and 1800 nodes to the network, representing a considerable increase in operational complexity. Such a rapid expansion requires careful network design and configuration, as well as the development of more advanced tools for continuous monitoring and efficient diagnostics. This paper presents the design and configuration of the new network, as well as diagnostic methodologies and monitoring strategies implemented to support and supervise the evolving network infrastructure.

## INTRODUCTION

White Rabbit (WR) [1] timing has been in use at CERN for many years, handling several different applications, each constituting a small, often accelerator-specific, network of switches and nodes. With the timing system renovation project in progress, the new branch of the White Rabbit network will reach nearly all control front-end computers in the entire accelerator complex. Such a solution offers great possibilities but also requires smart maintenance to ensure high availability.

In the past, there has been a proposal [2] regarding the diagnostics and monitoring approaches. Most of the methods described in [2] (e.g. configuration management, diagnostics over SNMP (Simple Network Management Protocol), Grafana dashboards) have been implemented and new ways to further enhance these tools are being explored.

## NETWORK ARCHITECTURE

The new branch dedicated to the timing system is essentially independent from the other branches currently in operation at CERN. The sole exception is the grandmaster switch, which serves as the authoritative time source for the entire network.

The architecture adopts a multilayer structure, with layers defined by the number of hops from the grandmaster (with the grandmaster itself defined as layer 0). A switch at layer 1, directly connected to the grandmaster, acts as the primary injection point for the distributed data required by the General Machine Timing (GMT), Safe Machine Parameters (SMP) and Radio Frequency (RF) systems. Beam Synchronous Timing (BST) does not require a dedicated transmitter, as it is derived from GMT and RF within the receiving nodes. At layer 2, the network diverges into two strictly separated sub-branches: the laboratory network and the operational network. For development purposes, the data can also be injected into a layer 2 switch within the laboratory branch. This arrangement ensures that potential misconfigurations or faults remain confined to the laboratory environment, leaving the operational branch unaffected. Furthermore, altering the injection point requires only a simple fibre reconnection between switches.

Beyond layer 2, the architecture will be extended with three additional layers of switches to span the entire CERN infrastructure and to provide the required number of ports for connecting all nodes. This deployment is an ongoing process, gradually extending coverage across all relevant areas.

The structure of the new branch is presented in Fig. 1.

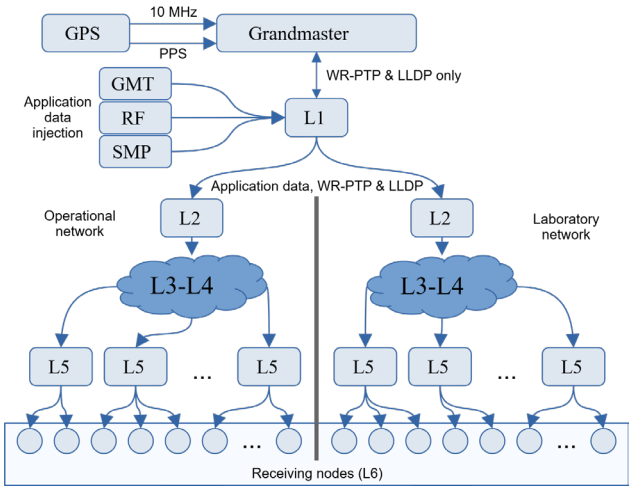



Figure 1: Structure of the new WR network branch

## CONFIGURATION MANAGEMENT

As proposed in [2], for every switch, a subset of the most useful settings is stored in the Controls Configuration Database (CCDB) and managed through its web interface, the Controls Configuration Database Editor (CCDE). The settings are translated into a switch configuration file (also known as *dot-config*) using a generator script and then downloaded by the switch from a boot server at startup.

Thanks to this approach, even if certain options change between firmware versions, the generator script can handle the differences and produce valid configuration files for the requested firmware version or refuse to create a configuration file, if an option is not supported in a particular version. Each generated configuration file is also committed to a git repository, facilitating change tracking and rollback.

† maciej.suminski@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD072

# THE HARDWARE AND SOFTWARE ARCHITECTURE OF THE MTCA.4 BPM ELECTRONICS FOR THE DESY'S PETRA IV ACCELERATOR

M. A. Škoberne*[1], A. Bardorfer[1], M. Cargnelutti[1], P. Leban[1], M. Oblak[1], P. Paglovec[1], B. Repič[1], R. Černe[1], G. Kube[2], J. Lamaack[2], H.-T. Duhme[2], F. Goeppert[2], and T.-J. Lensch[2]

[1]Instrumentation Technologies, Solkan, Slovenia
[2]Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

*Abstract*

The PETRA IV accelerator at DESY represents the forefront of synchrotron radiation science, demanding advanced instrumentation to deliver the requested beam quality. This paper describes the hardware and software architecture of the MTCA.4-based BPM electronics, prototyped collaboratively by Instrumentation Technologies (I-Tech) and DESY. The project was based on the legacy of the Libera Brilliance+ (LB+) BPM system and the Libera Base software architecture. While LB+ is based on the MTCA.0 architecture, the project required full compatibility with MTCA.4. By leveraging the existing front-end electronics and the mature software framework, the project focused on the platform compatibility challenges while keeping risks under control and expediting prototyping. The new key hardware components include a BPM-optimized RTM module developed by I-Tech and the DAMC-UNIZUP AMC processing board from DESY. After two years of tests at PETRA III, the validated prototype is ready to be industrialized and deployed for PETRA IV. The BPM application became hardware-agnostic, supporting multiple BPM boards while maintaining platform independence. This flexibility enhances system versatility and establishes its role in PETRA IV's FOFB architecture. By harmonizing innovative hardware design with robust software solutions, this work underscores a seamless transition from legacy technology to next-generation systems, offering valuable insights for future accelerator advancements.

## INTRODUCTION

The PETRA IV initiative at DESY in Hamburg, Germany, is focused on developing a next-generation, diffraction-limited synchrotron light source with ultra-low emittance, operating at an energy of 6 GeV. The new PETRA IV storage ring will be housed within the existing PETRA III tunnel, preserving the original eight-fold symmetry of the former PETRA collider. Its accelerator design is based on a modified hybrid six-bend achromat (H6BA) cell, and with a total circumference of 2.3 km, it is engineered to deliver electron beams with an emittance of just 20 pm rad. The PETRA IV facility layout is shown in Fig. 1.

This compact beam emittance results in extremely small beam sizes (approximately 7 µm horizontally and 3 µm vertically) at the insertion device source points, demanding exceptional machine stability. To achieve precise beam position monitoring and maintain orbit stability, PETRA IV will deploy a high-resolution BPM system comprising of 788 monitors. These will be supported by MTCA.4-based readout electronics [1].

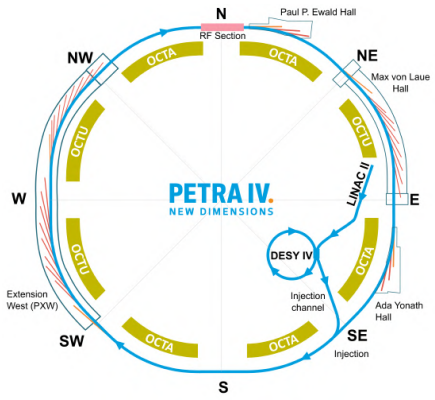



Figure 1: Layout of the PETRA IV facility. Courtesy: DESY.

During the commissioning phase, BPM accuracy is critical. Before beam accumulation can begin, the system must meet the stringent requirements of beam-based alignment, which involves identifying offset errors between the magnetic centers of nearby quadrupoles and the electrical centers of the BPMs. These offsets arise from both mechanical alignment tolerances and intrinsic electronic deviations [2].

To meet PETRA IV's long-term drift stability requirements (less than 1 µm over 6 days) the facility's unique layout, with infrastructure spread across former experimental halls, must be considered. This results in extended cable runs between BPMs and their electronics, which are vulnerable to environmental fluctuations. Stabilizing these cable paths is essential. Among the available drift compensation strategies, the crossbar switching method has proven effective. Additionally, the pilot tone (PT) technique, which injects a reference signal into the RF chain, has been considered. However, tests at PETRA III revealed limitations of the PT method under certain beam conditions.

Consequently, PETRA IV adopted an enhanced approach: external crossbar switching, in which the analog switching component is relocated closer to the BPM pickup to further stabilize the cable paths. In a collaboration between DESY and Instrumentation Technologies, a prototype was developed consisting of a Libera Brilliance+ unit and an external

* martin.skoberne@i-tech.si

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD074

# GPS IRIG-B OVER FIBER LINKS: AN ALTERNATIVE TO NTP

A. J. Hernandez, M. Bobrek, Oak Ridge National Laboratory, Oak Ridge, TN, United States

*Abstract*

The timestamping of neutron events and environment parameters for Spallation Neutron Source (SNS) instruments is based on Network Time Protocol (NTP) to distribute Global Positioning System (GPS) time to the instruments front ends. We are presenting an alternative way of distributing GPS time using Fiber-optical channels. Instead of millisecond accuracy for NTP, we distribute GPS IRIG-B signalling to remote IRIG-B receivers with an accuracy in nanoseconds. The new approach is implemented in VHDL and demonstrated on an FPGA development board using Vivado® design tools.

## INTRODUCTION

The current system in place at the SNS follows the NTP structure. NTP is a time protocol which is utilized by different servers throughout the facility to send and receive timing information to users and components. The network utilizes different servers which are rated by distance to know the accuracy of time being distributed by each server. These ranks are known as stratum values and the higher the value is, the more distance the server has from its master server, meaning there is an increase in delay to be expected at the receiving end. The source server (stratum 1) receives its timing from a GPS receiver which is then distributed to lower ranked servers throughout the facility. Time-of-flight calculations are made between servers to ensure timing information being sent has accounted for any delay the servers have between each other.

A custom NTP-based timing distribution system is used at the SNS to distribute timestamps of sensors and triggers from the accelerator. Components like the accelerator ring require high resolution timing as this becomes critical when proton acceleration occurs. Even more-so once event signals have been sent and their timestamps are required to ensure the system is functioning correctly. The accuracy for the time at which events are sent is crucial. Thus, requiring a need for as little delay between servers as possible. The current delays expected between servers are around 10ms. The proposed alternative in this paper offers delays within the nanoseconds while transmitting non-stop timing and accelerator data.

## GOALS, BENEFITS & CHALLENGES

The end goal of this system is to explore an alternative to the current protocol in place at the SNS. This is to be achieved through demonstration of generation, distribution, and reception of the accelerator's timing using FPGA Multi-Gigabit Transceiver (MGT) Fiber links. Depicted in Fig. 1, the end-goal of the system is to develop a timing master which takes GPS timing information and an RF clock source to timestamp and distribute accelerator events to timing distributors and receivers.

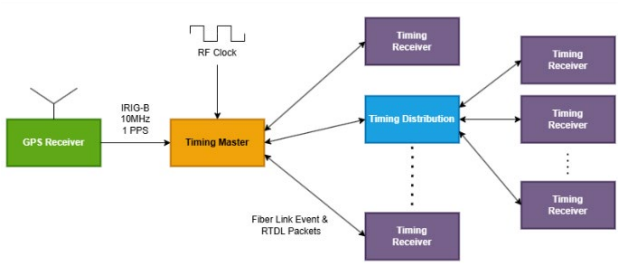



Figure 1: A diagram of the distribution of timing packets from a timing master to timing receivers after receiving source clock and GPS settings.

However, the aim of this paper is to go over the current achievements and steps taken to achieve this end goal. We focused on the development of an IRIG-B transmission environment using a Zynq-7000 SoC FPGA. Beginning development in June of 2024, the objective of this system was to transmit and handle IRIG-B time codes over a fiber optic and copper wire link. The system developed can be seen in Fig. 2. Similarly to the end goal of the timing master, our IRIG-B transmitters were able to send timing information through fiber optic channels to IRIG-B receivers and distributors if needed.

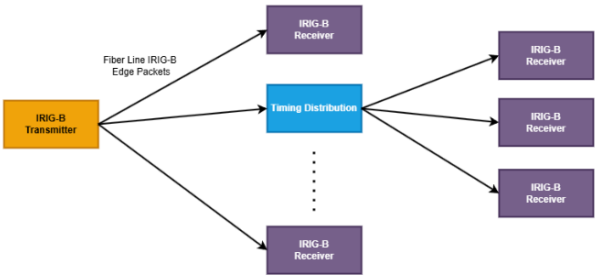



Figure 2: A diagram of the transmission of IRIG-B signals from transmitter to receivers.

The benefits of achieving an IRIG-B based time distribution network that utilizes MGT fiber links are:

- Increased Data Bandwidth.
- Time-of-Arrival Control.
- Commercial-Off-The-Shelf (COTS) vs. Custom Hardware.

The increase in data bandwidth is granted from the upstream and downstream capability of Fiber links for continuous data flow. The current protocol runs at an output of 100 Mbps, implementing the proposed system could increase this by 10 to 100 times. Current implementations of the system run at 2.5 Gbps to 2.8 Gbps but can be increased if needed.

The bidirectional output of the Fiber links allows for the measurement of the time taken for data to be sent between transmitter and receiver. This allows for adjustments to be made pre-transmission to account for delays made once

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD075

# Design of the Korea-4GSR Timing system*

Sohee Park[†,1], Jinsung Yu[1], Yunho Kim[1], Sunwoo Kang[2]
[1]Pohang Accelerator Laboratory (PAL), Pohang 37673, South Korea
[2]Korea Basic Science Institute (KBSI), Daejeon 34133, Korea

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

KOREA-4GSR is a fourth-generation synchrotron radiation accelerator located in Pohang, South Korea. Its event timing system is a system that coordinates and synchronizes events in a precise sequence over time and provides precise timing to local devices. The system consists of an event master (EVM), an event fanout (EVF), and an event receiver (EVR), and each local device receives event, trigger, and timestamp information through EVR. Korea-4GSR transfers 200 MeV electrons to the booster through a linear accelerator, and provides trigger signals to synchronize the injection order into the storage ring after ramping acceleration from the booster to 4 GeV. The system will be configured using MRF's MTCA products, and the EVG and EVR configuration and layout are being designed by investigating the specifications of local devices that receive trigger signals. In addition, the system will be tested by transmitting event codes and trigger signals to embedded EVR equipment such as BPM and FOFB.

## INTRODUCTION

KOREA-4GSR is a large accelerator consisting of a 772.9-meter-circumference booster and a 799.3-meter-circumference storage ring. Compared to existing third-generation accelerators, the fourth-generation accelerator achieves ultra-low emittance, delivering significantly brighter and more consistent X-rays. The electron beam generated by the gun passes through four accelerator tubes, reaching a beam energy of 200 MeV. It is then transferred to the booster via a kicker. From the booster, it enters the storage ring, where it is deposited into a bucket of choice. Within the storage ring, the beam is accumulated according to an operator-defined fill pattern.

The event timing system synchronizes all devices within the accelerator, providing trigger signals in a precise sequence that matches the time of the electron beam's passage. Furthermore, all devices are synchronized on a few picoseconds timescale, ensuring stable beam operation, making the timing system crucial.

## SYSTEM DESIGN

The timing system consists of an EVM, EVF, and EVR, and each local device receives event, trigger, and timestamp information from the EVR. Event timing output signals are provided as TTL or optical signals, and when receiving optical signals, both event information and timestamps can be received. The EVR must be embedded in the local device, and a protocol developed by MRF is used for this purpose. The timing system is designed based on the mTCA.4 form factor and requires components such as a CPU, MCH, chassis, and power supply. The installation of optical and coaxial cables is also required for the overall system configuration.

### *Architecture*

Figure 1 represents the basic architecture of the event timing system. The event timing system operates with clock synchronization to synchronize events occurring between multiple EVRs and precisely coordinate data communication and control operations, thereby ensuring the reliability and accuracy of the entire system. The clock synchronization signal is located close to the Master Oscillator and receives the signal, and the EVM receives the RF clock and AC 1-phase signal from the Master Oscillator and operates in synchronization. The synchronized event is transmitted from the EVM to the EVR through the Fanout. Devices with built-in EVRs receive events and implement their own trigger function. Devices without built-in EVRs are provided with a hardware trigger signal output from mTCA.4.[1]

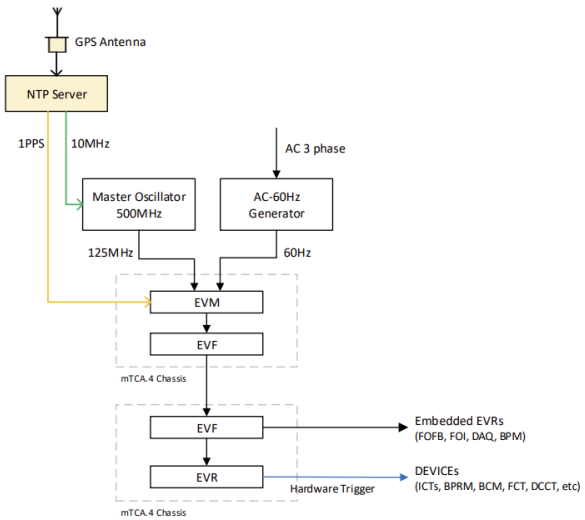



Figure 1: Event Timing System Architechure.

### *Software &Hardware*

EPICS middleware, the core of the system that provides an integrated environment for device control, runs on embedded Linux and uses EPICS 7 or later. Real-time patches are applied in the Linux environment to optimize the software environment by implementing real-time applications. Furthermore, the drivers provided by MRF play a crucial role in maximizing the performance of the timing system and ensuring compatibility with various hardware. The control environment was implemented by linking MRF equipment with EPICS applications, and system perfor-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD076

# FIBER SIGNAL ATTENUATION DUE TO TEMPERATURE

D. Sanchez, C. Bianchini Mattison, R. McClanahan
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The SLAC National Accelerator Laboratory distributes its timing signal throughout the beam via a mix of fiber cables. It is known that radio frequency (RF) over fiber signal is affected by temperature fluctuations causing a phase shift in the signal. That is why special consideration was taken in the design for the type of fiber and type of connector to be used for the LINAC Locking system. As the communication of the system would be a RF over fiber signal. However, there have been some instances over the years where higher temperatures have caused issues with the continuous signals going through the fibers. The cause of the degradation of the signal would usually be a damaged fiber or degraded transceiver but increased temperatures would lead to a high enough signal loss before complete failure of the devices happened. This paper describes how SLAC takes into consideration temperature environment for troubleshooting efforts and possible design considerations to make for future developments.

## INTRODUCTION

For delivery of the timing signal within SLAC National Accelerator Laboratory fiber is used to distribute the signal from the timing pattern generator (TPG) to the various timing pattern receivers (TPR). However, due to the layout of the accelerator a majority of the fibers ran to deliver the signal are ran through non-temperature-controlled environments. As a result of this the fibers get a change of characteristics depending on time of day and season of the year. Thus, some compensation is built into the triggers set to activate the various devices that get connected to the Timing System to offset any temperature related signal attenuation.

### *Timing System*

There are 3 timing systems used at SLAC National Accelerator Laboratory, Linac Coherent Light Source (LCLS), Linac Coherent Light Source II (LCLS-II), and FACET. The timing system for each of these systems used to send specific triggers at specific rates for synching or giving information to the appropriate subsystems within the accelerator. Such as the Machine Protection System (MPS), Beam Position Monitors (BPM), Camera Controls, and RF controls.

### *Fibers*

The fibers used to deliver the signal within the timing signal fall into 2 categories, single mode or multi-mode. The single mode fiber OS2 used is the industry standard fiber of a fiber strand that has a core of 9 micrometers and outer cladding of 125 micrometers. That allows for signals to be transmitted at farther distances and faster data rates. Multi-mode fibers in general transmit signals in shorter distances and slower data rates compared to single mode fibers. SLAC uses multiple standards of multi-mode fibers, OM1 is the first standard of multimode fiber used with a core size 62.5 micrometers and has the shortest transmit length and slowest data rate. The up-to-date multi-mode standards now being used at SLAC are OM3 through OM5 these standards share the same core size of 50 micrometers and have the same data transmit lengths and data rates. The key difference between these fiber types is the ability to send multiband signals withing the core. For the timing system this distinction between OM3 and higher is currently ignored as the same type of signal is transmitted to each sub system that requires a trigger. The only reason different type of fibers are used aside for legacy reasons is due to different systems only being able to use transceivers that restrict the type of core size being used either the single mode 9 micrometers or multimode 50 micrometers. Knowing the type of fiber being used for each subsystem allows for easier understanding to know how to compensate for any signal attenuation caused to temperature changes withing the environment. As, the core of the fiber is typically built with glass as temperatures in the environment changes so does the properties of the fiber. This causes the signal to change usually a phase delay and the equation below can be referenced to see how the change in phase is correlated to a change of temperature in the environment [1].

$$\frac{\Delta\phi}{\phi\Delta T} = \frac{1}{n}\left(\frac{\partial n}{\partial T}\right)_\rho + \frac{1}{\Delta T}\{\epsilon_z - \frac{n^2}{2}[(P_{11}+P_{12})\,\epsilon_r + P_{12}\,\epsilon_z]\}. \quad (1)$$

It is also important to note that aside from temperature changes caused by the environment there can also be changes cause by bends within the fiber. So consideration need to be taken to any loops and the power of the transceivers being used as that can lead to fibers that get warmer than the environment they are routed. The following equation can be used to calculate maximum temperature once equilibrium is reached [2].

$$\Delta T = \frac{P\max(1-\exp(-\beta\Delta P)}{\pi dlh} \quad (2)$$

### *Troubleshooting*

For the timing system and each other system receiving a timing signal for triggers. The signal is constantly being monitored and archived. If the signal would be lost a flag within the software us usually set from 1 to 0 causing an engineer to check the signal. First is to check the software and set resets to see if that causes a return of signal. When software adjustments do not bring back the signal then an engineer is sent out to the field to check on the hardware. The hardware to check are the Small Form Pluggable (SFP) transceivers, fiber connectors, and fiber strands. To troubleshoot the SFPs it is best to just swap out the transceiver

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD079

# EPICS SUPPORT FOR UNIVERSAL ROBOTS*

N. Marks†
Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The integration of modern robots into large-scale control systems, such as synchrotron beamlines, has traditionally required custom, installation-specific solutions, hindering collaboration and reusability. This work presents an EPICS support library for controlling Universal Robots (UR) e-Series robotic arms. Built on the widely used EPICS asyn module, the software provides a standardized interface for controlling UR robots from an EPICS soft IOC. This greatly simplifies deployment to new installations and promotes collaboration amongst robot users. In addition to exposing low-level robot functions, such as joint control and status monitoring, the library offers higher-level features, including waypoint and path creation, to facilitate common tasks like pick-and-place operations. Finally, packaged separately is a 3D visualization program to view the live position of the robot remotely.

## MOTIVATION

Universal Robots (UR) robotic arms, a leading platform in collaborative robotics ("cobots"), are used across several beamlines at the Advanced Photon Source (APS). Unlike most beamline devices such as motors and detectors, which are controlled through EPICS, no standardized EPICS support has been available for UR robots. This gap has led each user group to develop its own ad hoc control solutions, which are typically unique and not easily transferable across beamlines or installations. Yet many groups share similar goals with these robots—often centered on simple, repetitive tasks such as pick-and-place—so the absence of a common framework has hindered the development of reusable, higher-level task implementations. In addition, there has been no convenient way to visualize robot motion other than pointing a web-based camera at the robot, a workaround that is not always practical.

## METHODS AND FEATURES

### *EPICS Support Module*

The robot control software is implemented as a standard EPICS support library, intended to be included within IOC applications [1]. Control is handled using the ur_rtde C++ library, which communicates with the robot through Universal Robots' Real-Time Data Exchange (RTDE) interface. Each RTDE interface is encapsulated in a dedicated asynPortDriver class, allowing EPICS IOCs to connect selectively depending on the intended use case—for example, for monitoring robot state only or for issuing external control commands. To make the system accessible without requiring user-side programming, graphical user interfaces (GUIs) were developed in both MEDM and caQtDM, providing full interaction with the robot through familiar EPICS client tools.

The UR robot support module provides EPICS process variables (PVs) for all fundamental robot features, including the UR dashboard, digital I/O, status monitoring, control interfaces, and the Robotiq Hand-E gripper. The system is designed for full operation through EPICS GUI clients, eliminating the need for user-side coding (Fig. 1).

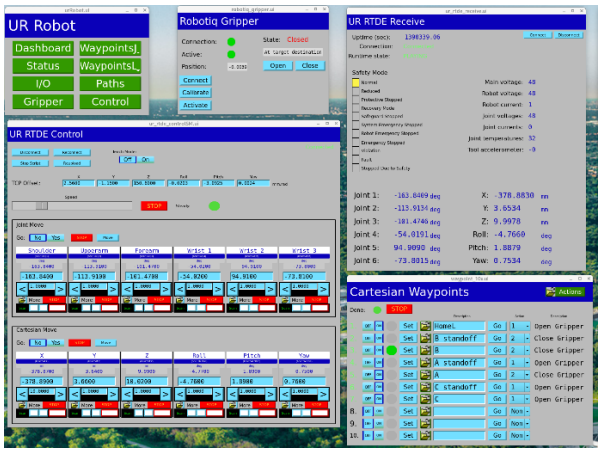



Figure 1: EPICS caQtDM GUIs for controlling the robot.

Users can jog individual joints or the tool position using EPICS soft motors, and waypoints can be saved directly as EPICS PVs in either joint or tool space, with optional associated actions. These actions are implemented using a combination of EPICS forward links and *calcout* records, which both trigger the desired operation and signal when it has completed, ensuring that motion to the next waypoint does not begin until the current action is finished. More complex motion can be achieved by chaining waypoints together to build paths. Paths are defined by specifying a series of waypoints to visit in sequence, with the option to perform an action after reaching each waypoint, temporarily overriding the action already defined for that waypoint if desired. Since most robot tasks can be expressed as moving the tool to a series of positions and performing an operation at each, path definitions provide a flexible and powerful mechanism for building reusable robot procedures.

### *URViewer Visualization Tool*

The URViewer [2] 3D visualization tool (Fig. 2) was developed alongside the EPICS support module as a minimal application that renders a 3D model of a UR robotic arm based on measured joint angles. While many similar visualization tools exist, none met the combined requirements of being easy to compile from source, which enables straightforward installation on any computer or system

* Work supported in part by the U.S. Department of Energy under contract DE-AC02-06CH11357
† nmarks@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD080

# DCCT CONTROLS AT THE ADVANCED PHOTON SOURCE

S. G. Wang*, T. Fors, R. Hong, R. Keane, W. Cheng, G. Shen
Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Advanced Photon Source Upgrade (APS-U) implements a novel swap-out injection scheme. The Direct Current Current Transformer (DCCT) in the Advanced Photon Source (APS) Storage Ring is used to measure the storage ring current. This measurement is digitized using an HP3458A Multimeter and synchronized with beam injection through the Micro-Research Finland (MRF) event system. This paper describes the setup and timing of the control system for the DCCT, highlighting its integration and functionality within the APS infrastructure.

## INTRODUCTION

APS-U [1] is an ultra-low emittance storage ring that has been constructed and commissioned at Argonne National Laboratory. To accommodate the small dynamic aperture of the storage ring, a swap-out injection [2, 3] scheme is used. Accurate measurement of beam current synchronized with injection is essential for the operation, protection, and optimization of particle accelerators. DCCTs provide non-invasive and high-precision measurements critical for storage rings and synchrotrons. This paper describes the implementation of a DCCT system that has been successfully deployed in a production environment.

The system architecture ensures high availability through redundancy and seamless integration with accelerator timing systems. It supports both operational monitoring and diagnostic studies, offering remote control and monitoring capabilities. EPICS Access Security is implemented for critical functions to maintain security and reliability.

## SYSTEM ARCHITECTURE

### *Hardware Configuration*

The DCCT system implements a dual-redundant architecture. The Fig. 1 shows a simplified version of only one half of the system's redundancy.

The DCCT system is equipped with two Current Monitors, CM1 and CM2, each featuring NPCT (New Parametric Current Transformer) [4] signal conditioning to ensure precise measurements. It utilizes two HP3458A [5] digital multimeters (DMM), DMM1 and DMM2, for high-precision measurement tasks. The control system is housed within a VME chassis and operates on an MVME2500 processor board running EPICS [6] IOC. Instrument communication is facilitated by two NI1014 GPIB controllers, designated as L0 and L1. For interrupt handling, the system employs a TVME200 carrier with IP408 digital I/O. Additionally, the MRF [7] timing system components, including the EVG and multiple EVRs, provide synchronization capabilities.

The active DCCT is selected via EPICS process variables, allowing operators to switch between units without physical intervention. This redundancy ensures continuous operation during maintenance or in the event of a fault.

### *Signal Flow and Interfaces*

The DCCT toroid output connects to the NPCT module, which provides signal conditioning and calibration capabilities. The NPCT output voltage, proportional to the beam current, is measured by the HP3458A DMM.

The VME-based IOC communicates with the DMM over GPIB via NI-1014 interface modules. Triggering and synchronization are provided by TTL signals from the MRF event system, and the interface box's calibration is controlled through an IP408 digital I/O module.

## OPERATIONAL MODES

### *Injection-Synchronized Measurements*

The system utilizes an MRF event as an external trigger; this 1 Hz trigger synchronizes measurements with beam injection, enabling precise current monitoring immediately after injection. The newly measured DCCT value is used for injection efficiency calculation and the bunch current monitor.

The measurement sequence involves:

1. MRF TTL triggers DMM measurement
2. IOC software trigger delayed by 200 ms triggers the GPIB read of measurement data
3. Data processing and distribution via EPICS
4. Due to the limited CA connections of the VME IOC, there is a soft IOC to broadcast the DCCT current

The 200 ms delay ensures the DMM has completed its measurement and the data is ready for retrieval. Disabling AutoRange prevents measurement delays when beam current changes significantly during fill-from-zero or slow beam kickout. The Autozero is moved after the read is completed, so it won't delay the measurement.

### *Waveform Mode for Diagnostics*

As we advance development of the DCCT system for the Storage Ring (SR), we are also planning to upgrade the Booster DCCT. Because diagnostic waveform needs differ by application, the Booster DCCT will be sampled at least at 20 kHz over the beam-on portion of the booster cycle (approximately 200 ms), yielding about 4,000 samples per cycle. For the SR, a longer waveform will improve power spectral density (PSD) analysis and help identify spurious

* grass@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD082

# LABVIEW-BASED MODULAR CONTROL SYSTEM FOR THE NOVEL ELECTRON WINDOW TEST STAND*

E. Galetti, L. M. Golino, P. Geißler, C. Charrondiere, A. Rossi, G. Khatri, O. Andreassen,
CERN, Genève, Switzerland
M. Sameed, Universal Quantum, Hamburg, Germany

*Abstract*

Electron Beam Flue Gas Treatment (EBFGT) offers a solution to reduce SOx and NOx emissions coming from the maritime industry. The accelerator setup must be compact and energy-efficient enough to fit onboard a ship. To study EBFGT efficiency, a membrane window has to be placed between the electron gun and the beam-gas interaction area. A Novel Electron Window Test Stand (NEWTS) has been designed to test and validate these ultra-thin membranes by studying the electron beam attenuation and deflection. A dedicated control system for NEWTS has been designed and deployed, which ensures reliable real-time operation in both continuous and pulsed (from milliseconds to nanoseconds) modes. In this paper, we will explain the implementation and architecture of the control system and its suitability to operate in a harsh environment. By combining the capabilities of LabVIEW, RADE, GitlabCI, Python, NI cRIO and FPGA, we will show how the control system for NEWTS was able to be put in place within a few months to control a variety of "off-the-shelf equipment" and integrate them within the CERN infrastructure. We will also demonstrate the control system's versatility, such that it can be reused to control other electron gun test benches at CERN.

## INTRODUCTION

The Novel Electron Windows Test Stand (NEWTS) is a test facility built for the UTMOST CLEEN project [1], which was funded by the Knowledge Transfer (KT) Group at CERN. At NEWTS, an electron beam produced by a commercial electron gun, ES-100-S, manufactured by STAIB Instruments GmbH [2], is used to characterise the transmission of different membrane windows. The same set-up is used to test the Beam Gas Ionisation (BGI) monitor [3].

This section will outline the principles of the Electron Beam Flue Gas Treatment (EBFGT), the project's challenges, and the NEWTS facility. The "Control system overview" section introduces the control system from hardware and software perspectives. The "Control system validation" section provides details of the tests and validation of the control system. The "Other E-Beam Test Benches" section shows an example of how the control system could be reused in other electron beam test facilities. Finally, the "Conclusions" section summarises the work and outlines future developments. This work allows for remote operations and data logging for post-processing and is well integrated into the CERN infrastructure, allowing users to benefit from all the available platforms at CERN.

### *EBFGT Principles*

Electron Beam Flue Gas Treatment (EBFGT) is an advanced air pollution control technology designed to tackle the emission of various noxious compounds from industrial processes [4–6]. This technology utilises a high-power electron beam to interact with the exhaust gases, triggering chemical reactions that yield radicals and other reactive particles within the flue gas stream. These reactive species drive the conversion of pollutants present in the exhaust gases into by-products that can be more effectively filtered before the treated exhaust gas is released into the environment.

Adapting the EBFGT technology to the maritime industry is a promising solution to address the pressing environmental concerns associated with ship emissions [7]. This technology holds the potential to significantly reduce harmful pollutants (e.g. SOx, NOx). EBFGT's versatility in operating with various marine fuels, such as heavy fuel oil and liquefied natural gas, enhances its adaptability to diverse vessel types and fuel preferences. Moreover, EBFGT has demonstrated effectiveness in controlling particulate matter emissions, crucial for improving air quality, especially in port areas significantly impacted by ship emissions.

### *Challenges*

The implementation of EBFGT in the maritime sector presents several challenges. Mechanically, the system must be robust to withstand the rigours of the maritime environment. The EBFGT system should be compact for space-optimisation on ships, reliable, and cost-competitive with existing emission control solutions. As a consequence, the control system needs to be compact, resistant to vibration, and capable of managing beam pulses shorter than 1 µs.

### *The NEWTS Facility and its System Specifications*

To address some of the challenges of introducing the EBFGT in the maritime industry, feasibility studies are being conducted at CERN, where a test facility has been built to simulate the operational set-up, as shown in Fig. 1. The test stand will be used to develop and validate the control system for the compact electron gun[1]. The NEWTS source is developed for high current, low energy applications (~ 0.1 mA electron beam at up to 100 keV), and is equipped with an electron optical system to focus the beam. Electrons are produced by a cathode (a heated filament) and accelerated and

* Project funded by Knowledge Transfer at CERN for the UTMOST CLEEN project, in the frame of the CIPEA (CERN Innovation Programme on Environmental Applications).

[1] ES-100-S electron source from STAIB Instrument GmbH

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD083

# BEAM POSITION MONITOR CONTROL SYSTEM AT THE EUROPEAN SPALLATION SOURCE

J. F. J. Murari*, R. A. Baron, J. P. S. Martins, H. Kocevar
European Spallation Source, Lund, Sweden

*Abstract*

The European Spallation Source facility is divided in three main parts: linear accelerator (LINAC), target and neutron instruments. The Beam Position Monitor (BPM) system is installed along the LINAC and enables accelerator teams to characterize the proton beam properties and optimize the phase tuning of the RF cavities, among other diagnostics. A total of 98 BPM sensors are distributed along the machine, with data acquisition and processing handled by 49 AMC digitizer cards housed in 18 MicroTCA (MTCA) crates. This paper describes the control system architecture surrounding the BPM system, including the EPICS integration along with the graphical user interface developed in Control System Studio. In addition, it presents high-level applications built on top of this framework, such as the BPM Manager, Gain Controls, Archiver, and Synchronous Data Service. The paper also outlines the software strategy adopted for system deployment, maintenance, and updates. Auxiliary systems supporting the BPM infrastructure are briefly discussed, including the BPM components controls, MTCA IPMI management and the MRF (Micro Research Finland) Timing System integration.

## SYSTEM ARCHITECTURE

The BPM system [1] is responsible for measuring the position and phase of the proton beam along the 600 meters ESS linear accelerator [2], enabling important beam characterization and other beam diagnostics activities. It was developed and designed by the *Proton Beam Instrumentation* (PBI) group and integrated into the main EPICS control system in collaboration with the *Integrated Control System* (ICS) group at ESS. Different types of BPM sensors were designed to accommodate the different beam characteristics and to satisfy the beam physics specifications that have different requirements in terms of accuracy and resolution depending on the section of the machine.

The BPM controls hardware is implemented accordingly to ESS ICS strategy for control system hardware, which defines MTCA as the standard for systems with high-speed (above 100 kHz) acquisition requirements. Figure 1 shows the diagram of a generic implementation of the system with the sensors, RF front-end electronics and the MTCA platform.

The MTCA systems at ESS [3] do not share functions among different sub-systems of the LINAC, therefore the BPM has its own dedicated crates installed across the facility. A BPM MTCA crate contains different types of cards installed to satisfy all the demands of the BPM measurements

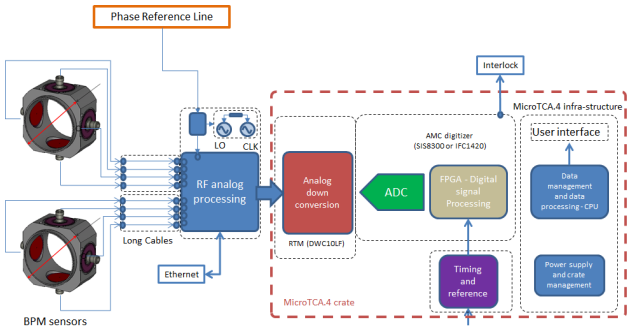



Figure 1: BPM system block diagram.

and the integration with other systems of the facility, as the Timing System and Machine Protection System. The main components are the following:

- Wiener 1000W MTCA.4 Power Supply.
- NAT PHYS-80 MicroTCA Carrier Hub (MCH).
- Concurrent Technologies AM G64 AMC CPU.
- Micro Research Finland MTCA-EVR-300.
- Struck SIS8300KU 250 MSPS ADC AMC.
- Struck DWC8300 RTM.

For the BPM systems, most of of the MTCA crates are the three rack unit size (3U) from *nVent/Schroff*, with six (6) slots for AMC cards. There is also one 9U MTCA crate, but in both cases a maximum of four (4) *SIS8300KU* digitizers is used. Since this digitizer has 10 analog inputs, each unit can measure two (2) BPM sensors, thus, a single BPM MTCA system can control and monitor up to eight (8) BPM sensors. Figure 2 shows a typical BPM MTCA system at ESS.



Figure 2: BPM system MTCA 3U crate at ESS.

The software infrastructure for the EPICS integration of the BPM systems at ESS benefits from a modular framework built on top of the *asyn* and *areaDetector* EPICS device support libraries. Since most of Beam Instrumentation MTCA

* juliano.murari@ess.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD086

# DEVELOPMENT OF A BEAM GATE CONTROL SYSTEM FOR PROTON BEAM IRRADIATION AT KOMAC LINAC*

Y.-G. Song†, J.-H. Kim, S.-Y. Cho, H.-S. Jung, H.-J. Kwon
KOMAC, KAERI, Gyeongju, South Korea

*Abstract*

A total of ten beamlines have been designed for the KOMAC 100-MeV proton linear accelerator (LINAC), among which five are currently operational and delivering beams to users. The LINAC system comprises an ion source, radio frequency (RF) system, high voltage converter modulator (HVCM), and a beam diagnostic system, all synchronized through a timing system to enable precise beam acceleration. A dedicated monitoring and control environment has been established for experiments conducted in each target room. Prior to beam irradiation, beam uniformity and dose per pulse are measured, and users input the desired total fluence into the control panel to initiate beam service. The beam gate control system is designed to automatically stop the beam trigger once the user-defined fluence is achieved, ensuring accurate and safe beam delivery. The beam gate control system was implemented with a redundant architecture to safely control the trigger signals output from the timing system. This paper presents the design and implementation of the beam gate control system developed for beam irradiation applications at KOMAC.

## INTRODUCTION

The Korea Multi-purpose Accelerator Complex (KOMAC) operates a high-power proton accelerator that provides beam pulses to users in pulsed mode [1]. In this operation, the beam pulse repetition rate is inherently constrained by the modulator and high-frequency (RF) repetition conditions. User experiments, however, demand flexible control of the beam fluence, requiring a specified number of beam pulses at a given repetition rate. Since the allowed beam repetition rates are often restricted to sub-harmonics of the RF frequency, experimental efficiency is limited. To overcome this limitation, we developed a beam pulse control system based on the KOMAC timing system. The system provides users with the flexibility to freely select beam repetition rates within the constraints of the RF operation, while also ensuring accurate delivery of the requested number of pulses.

Two complementary methods have been developed for precise beam pulse control as shown in Fig. 1. First, an EPICS SNL sequence monitors the event codes corresponding to beam pulses and automatically disables the event Pulser and EVR output ports once the requested number of pulses has been delivered [2]. Second, a dedicated hardware Beam Counter counts output gates from EVR. The counted gates are supplied to a high-voltage and a bias power supply of an ion source. The Beam Counter also accepts external interlock signals to ensure immediate beam termination. Through these methods, the KOMAC timing system enables flexible adjustment of beam repetition rates within RF operating conditions, accurate delivery of user-specified beam pulses, and real-time monitoring of missing pulses, thereby improving the efficiency and reliability of beam services.

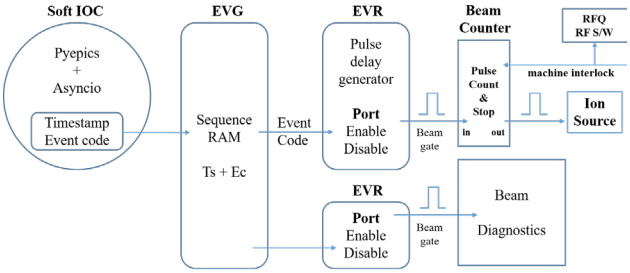



Figure 1: Beam pulse generation using event-code-based control and beam counter system.

## KOMAC TIMING SYSTEM

As shown in Fig. 2, The KOMAC timing system is constructed using Event Generator (EVG-230) and Event Receiver (EVR-230) modules [3][4]. The EVG generates the event clock and transmits event streams to the EVRs. Each EVR decodes the received events and provides synchronized trigger outputs to subsystems such as the RF systems, modulators, ion source, and diagnostic devices.

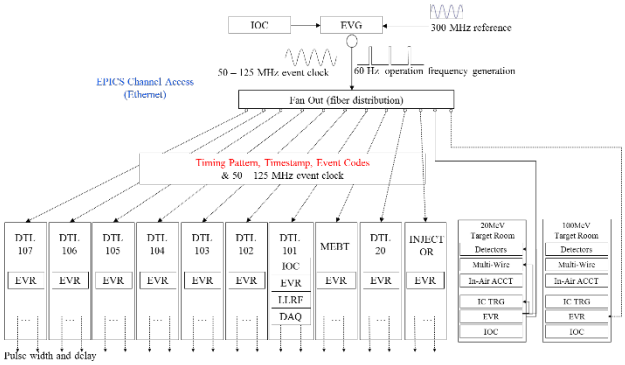



Figure 2: Event timing system in the klystron galley and beam experiment hall.

The timing system plays a central role in synchronizing beam extraction and diagnostics. When the requested number of beam pulses is entered, the timing system generates beam pulses and provides synchronized beam pulses to all diagnostic devices as well as to the RF feed-forward system. All EVRs involved in beam delivery are configured consistently to use the same event codes and pulse structures. This uniform configuration ensures that beam-related signals can be managed coherently across the facility,

* Work supported by KOMAC operation fund of KAERI by MSIT
† ygsong@kaeri.re.kr

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD087

# DRIVING FUSION SUCCESS: TARGET DIAGNOSTICS SOFTWARE CONTROL AT NATIONAL IGNITION FACILITY (NIF)

S. Heerey, J. Kehl, K. Burns
Lawrence Livermore National Laboratory, Livermore, CA USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Achieving controlled nuclear fusion at the National Ignition Facility (NIF) depends not only on precision lasers and target engineering, but also on a robust suite of Target Diagnostics systems. These diagnostics capture critical physics data—including neutron yield, burn width, symmetry, and x-ray emissions—essential for guiding the path to ignition. Central to this capability is the Target Diagnostics software control system synchronized to within tens of picoseconds of each laser shot. It coordinates nearly 100 diagnostics and over 600 software-controlled instruments, including high-speed sensors and radiation-hardened devices operating in extreme environments. A major challenge is supporting diverse, evolving hardware from multiple vendors using protocols like serial, Ethernet, GPIB, and proprietary APIs. The software architecture addresses this through centralized orchestration, abstraction layers, and hardware adapters. The system's flexibility is demonstrated in the successful deployment of the VISAR (Velocity Interferometer System for Any Reflector) diagnostic, which integrates varied vendor instruments — cameras, lasers, delay generators — with custom interfaces and precise synchronization. Long-term use has revealed challenges with aging, vendor-specific systems and managing radiation-hardened devices. Future improvements will enhance resilience, sustainability, and integration of next-generation diagnostics.

## INTRODUCTION

Achieving controlled nuclear fusion at the National Ignition Facility (NIF) requires more than powerful lasers and precision target engineering. A critical component of this effort is the suite of Target Diagnostics systems, which provide the data necessary to evaluate each experiment and guide the path toward ignition [1]. These diagnostics capture a wide range of physics parameters including neutron yield, burn width, hot spot symmetry and x-ray emissions (Fig. 1). Without accurate, reliable diagnostic data, it would be impossible to validate experimental outcomes and refine future designs.

Coordinating nearly 100 diagnostic systems, supported by more than 600 software controlled instruments, presents significant technical challenges. Devices originate from multiple vendors, employ diverse communication protocols, and operate in harsh experimental environments that demand radiation-hardened designs. The Target Diagnostics software control system was developed to meet these challenges by providing precise timing synchronization, flexible hardware integration, and robust orchestration of all diagnostic activities.

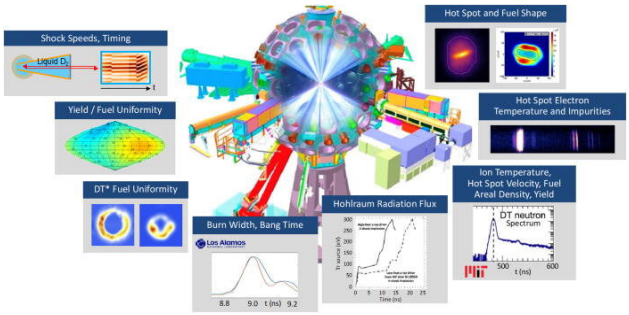



Figure 1: Target Diagnostics measure a variety of parameters to help achieve fusion.

## TARGET DIAGNOSTICS SOFTWARE CONTROL SYSTEM

The Integrated Computer Control System (ICCS) is the comprehensive, distributed software and hardware platform that automates and synchronizes all major operations at the NIF including laser power, alignment, timing, diagnostics, and data archiving [2] (Fig. 2). It is responsible for orchestrating the complex sequence of events required to prepare, execute, and analyse high energy laser experiments.

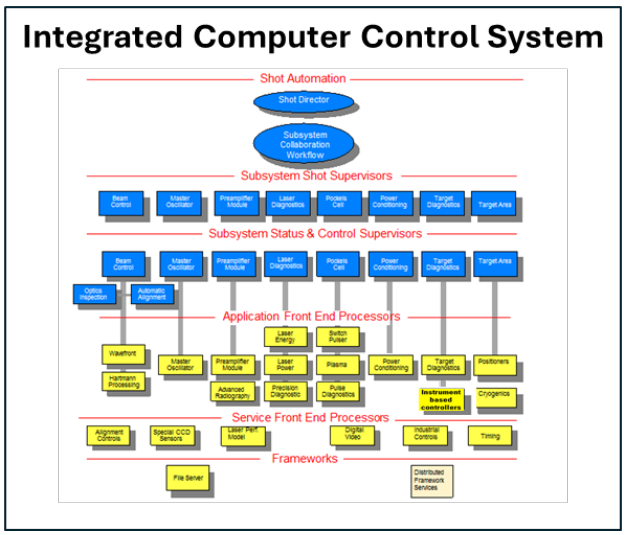



Figure 2: Various subsystems which make up the Integrated Computer Control System (ICCS).

The Target Diagnostics control system is composed of Instrument Based Controls (IBCs). IBCs serve as intermediaries between the high-level ICCS framework and the instruments attached to their host machines, utilizing the standardized Diagnostic Control Protocol (DCP) for communication with ICCS as well as low-level device communication protocols for direct control of hardware [3].

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD090

# POWER PMAC DRIVE CURRENT CASE STUDY*

O. Ivashkevych†, J. Sinsheimer
Brookhaven National Laboratory, Upton, NY, USA

*Abstract*

We analyse current regulation in the Power PMAC motor controller, a new generation of the Turbo PMAC line. Using oscilloscope and probe measurements, we compare commanded and delivered currents, quantify deviations, and performance against the predecessor Geo Brick LV. The results demonstrate improved stability, accuracy, and suitability including for low-current stepper motors.

This is particularly important in high-performance scientific instruments that demand elevated drive currents, such as when mechanical upgrades increase load requirements or as motors experience wear over time

## INTRODUCTION

NSLS2 is gradually transitioning from its predecessor, the Geo Brick LV 5A motor controller, to the new Power Brick LV [1] motor controller, part of the Power PMAC family of controllers. The Power PMAC was originally developed by Delta Tau Data Systems, which is now part of Omron Automation. Power PMAC offers great flexibility, though its wide range of configuration options adds complexity. The new controller provides an object-oriented, human-readable interface for specifying configuration parameters. Among the most important parameters for stepper motors are the commanded direct, or magnetization current, Motor[x].IdCmd, and the integrated current limit, Motor[x].I2TSet [2]. Both can be explicitly defined within the configuration and readily adjusted as needed. The question we asked is how the configured value of Motor[x].IdCmd corresponds to the actual current delivered to the motor. This issue is particularly important in high-performance scientific instruments that demand elevated drive currents, such as when mechanical upgrades increase load requirements or as motors experience wear over time, and for in-vacuum motors where heat load must be minimized.

## EXPERIMENTAL SETUP AND DATA

Table 1 lists the 2 phase stepper motors we tested.

Table 1: Tested Motors

| Maker | Motor | Current |
|---|---|---|
| Oriental Motors | PK264-03B-P1 | 2.1 A |
| Faulhaber | AM2224-V-12-75-10 | 125 mA |
| Faulhaber | AM1020-V-12-250 | 45 mA |

*This research used resources of the National Synchrotron Light Source II, a U.S. Department of Energy (DOE) Office of Science User Facility operated for the DOE Office of Science by Brookhaven National Laboratory under Contract No. DE-SC0012704.

The current was measured using a Tektronix MD3014 oscilloscope with a TCP0030A current probe, and RMS values were calculated using the built in software. This was compared with the expected current derived from the configuration parameters on the motor controller. Motor[x].IdCmd is a measure of the peak current in units of bits. For our model of controller, the conversion of bits to mA should be

$$I_{peak} = \frac{33{,}850}{2^{15}} \times IdCmd(bits), \quad I_{RMS} = \frac{I_{peak}}{\sqrt{2}}.$$

We refer to this as RMS-Theory. Note that here we only test at very small velocities, so the quadrature current contribution to the total current output is negligible. The total current output is the vector sum of the direct ($I_d$) and quadrature ($I_q$) currents.

$$I = \sqrt{I_d + I_q}$$

At the very small velocities used in this study, motor back-EMF and torque demand are negligible, so the quadrature current contribution vanishes, and the delivered current is dominated by the direct ($I_d$) component.

The probe-measured RMS current was marginally higher than expected, with smaller deviations seen at smaller currents for a small current motor as shown in Table 2.

The direct way to boost the current delivered to motor is by raising Motor[x].IdCmd. Figure 1 confirms that we can safely increase IdCmd by a factor 1.3 above motor rated RMS value if needed.

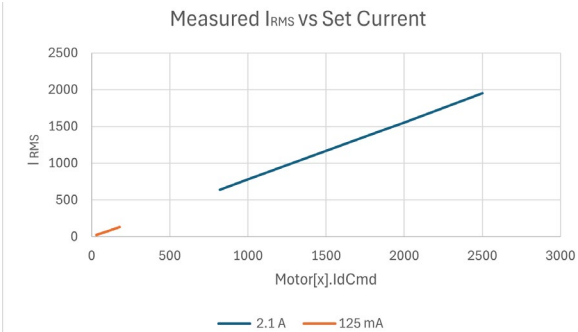



Figure 1: Measured RMS current vs applied via IdCmd.

The Power PMAC IDE configuration software includes a plotting (Gather) tool capable of sampling a wide range of parameters at rates up to 5 kHz, prior to their application to the motor. Using this tool, we collected Motor[x].IqCmd, the commanded quadrature, or torque producing current, demonstrating the negligible output at these small velocities, as shown in Fig. 2.

The current reported by the PMAC IDE gather tool, adequately reflects the real current applied to the motor. Figure 3 presents the corresponding current measured directly with the probe.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD094

# NANOSCALE PRECISION MULTI-AXIS MOTION CONTROL FOR THE CBXFEL PROJECT*

M.A. Montironi†,1, W. Lewis[2], S.J. Stein[2], N. Balakrishnan[1], T. Thayer[1], X. Permanyer[1], M. Balcazar[1], D. Zhu[1], A. Halavanau[1], A. Lutman[1], H.-D. Nuhn[1], Z. Huang[1], Y. Shvyd'ko[3], R. Lindberg[3], D. Shu[3], M. White[3], G. Lanza[1]
[1]SLAC National Accelerator Laboratory, Menlo Park, CA, USA
[2]Osprey DCS, MD, USA
[3]Argonne National Laboratory, Lemont, IL, USA

*Abstract*

The Cavity-Based Free Electron Laser (CBXFEL) project is proposing to produce a recirculating X-ray cavity and deploy it to the SLAC LCLS (Linac Coherent Light Source) Hard X-ray (HXR) undulator line. The shape of the cavity is defined by four diamond crystals which must be positioned with nanometer-level accuracy in four Degrees-Of-Freedom (DOF). Additionally, several electron and X-ray beam diagnostic components need to be precisely positioned to achieve, monitor, and maintain the cavity alignment. These functions are accomplished by a total of sixty-nine motion axes, eight of which are actuated by lead-screw stages operated by stepper motors, thirty-seven by Ultra High Vacuum (UHV) SmarAct piezo stages, and twenty-four by custom designed flexure stages actuated by UHV piezo linear actuators. The flexure stages are driven by UHV piezo actuators and real-time position feedback is provided by capacitive sensors and optical interferometers. A motion control system based on the CK3M PMAC architecture was developed to drive the different motion stages. This paper describes the main requirements to be met, how the technologies were integrated into the accelerator control system, and the main lessons learned.

## INTRODUCTION

The Cavity-Based Free Electron Laser (CBXFEL) project is proposing to produce a recirculating X-ray cavity and deploy it to the SLAC LCLS (Linac Coherent Light Source) Hard X-ray (HXR) undulator line [1, 2]. Figure 1 shows the experimental layout where the optical cavity, defined by four Bragg-reflecting crystal mirrors (C1, C2, C3, C4), encloses the first seven HXR undulator segments. Dipole magnets in a chicane configuration divert the incoming electron beam around C4 and re-insert it in the first undulator segment to generate x-rays. Another chicane diverts the electron beam exiting the seventh undulator segment, around C1, and re-inserts it in the eighth undulator segment just downstream of C1.

Beyond the C1-C4 crystals, the CBXFEL system is composed of several in-vacuum and in-air beam diagnostics and movers organized in six stations: A, B, C, D, E, and F as shown in Fig. 1. Design for an additional station, G, planned for the experimental hutches, will be finalized once Station B installation is completed.

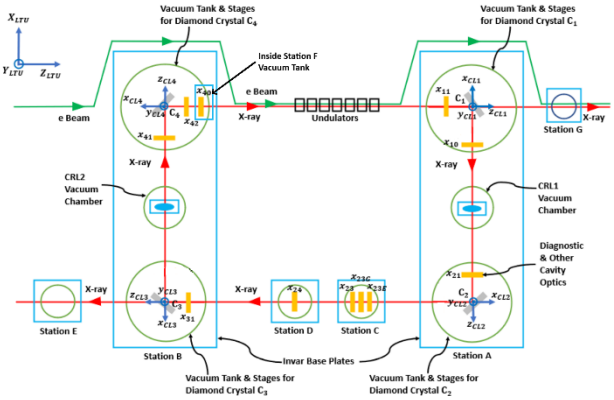



Figure 1: CBXFEL experimental layout.

### *Station A*

Station A, located in HXR cell 21, hosts the C1 and C2 diamond crystals, three X-ray Beam Position/Profile Monitors (X11, X10, and X21), and one Compound Reflective Lens (CRL1)

### *Station B*

Station B, located in HXR cell 13, hosts the C3 and C4 diamond crystals, three X-ray Beam Position/Profile Monitors (X31, X41, and X42), and one Compound Reflective Lens (CRL2)

### *Station C*

Station C, located along the return line, hosts three X-ray Beam Position and Profile Monitors (X23E, X23G, and X23) as well as a beam sampling diffraction grating.

### *Station D*

Station D, located along the return line, hosts one X-ray Beam Profile Monitor (X24).

### *Station E*

Station E, located along the return line, hosts a Beam Position/Profile monitor and two photodiodes.

### *Station F*

Station F, located downstream of Station B in HXR cell 13, hosts a Beam Overlap Diagnostic (BOD).

* Work supported by US Department of Energy Office of Science Under Contracts DE-AC01-06CH11357 and DE-AC02-76FO0515
† alexmon@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD095

# CONTROL SYSTEM IMPLEMENTATION FOR TEMPERATURE MONITORING AND CONTROL FOR CBXFEL PROJECT*

N. Balakrishnan[†,1], M. A. Montironi[1], S.J. Stein[2], W. Lewis[2], P. Liu[3], S. Karimian[1], I. Ali[1], Y. Shvydko[3], X. Permanyer[1], Z. Huang[1], M. Balcazar[1], D. Zhu[1], A. Halavanau[1], D. Shu[3], R. Lindberg[3], Z. Huang[1], M. White[3], G. Lanza[1]
[1]SLAC National Accelerator Laboratory, Menlo Park, CA, USA
[2]Osprey DCS, MD, USA
[3]Argonne National Laboratory, Lemont, IL, USA

*Abstract*

The Cavity-Based Free Electron Laser (CBXFEL) project, a collaboration between Argonne National Laboratory (ANL) and SLAC National Accelerator Laboratory, aims to produce a recirculating X-ray cavity for SLAC Linac Coherent Light Source (LCLS) Hard X-ray (HXR) undulator line. The cavity formed by crystal mirrors will facilitate FEL amplification and Bragg monochromatization [1]. To minimize losses, all cavity crystals must be maintained at the same temperature within +/- 0.4 °C. To support this requirement, a calibrated heater and temperature sensor will be installed at each crystal to ensure this stability. This paper focuses on the architecture and design implemented for the project including sensor calibration and control loop tuning.

## INTRODUCTION

The Cavity-Based Free Electron Laser (CBXFEL) project, a collaboration between Argonne National Laboratory (ANL) and SLAC National Accelerator Laboratory, aimed at advancing the capabilities of the SLAC Linac Coherent Light Source (LCLS) through the installation of a recirculating X-ray cavity in the SLAC Undulator Hall along the Hard X-ray (HXR) undulator line [1].

As shown in Fig. 1, the experimental layout features an optical cavity that encloses seven HXR undulator segments. The cavity consists of two crystal mirrors (C1 and C2; Station A) mounted on the eighth girder (HXR Cell 21) and two additional crystal mirrors (C3 and C4; Station B) installed on the first girder (HXR Cell 13). The electron beam is initially diverted around the crystal mirror C4 through a chicane before passing through the undulator segments. A second chicane then bends the electron beam around C1 [2]. The X-rays produced are reflected within the cavity by crystal mirrors mounted on nano-positioning motion stages [3].

Due to the thermal expansion of the Bragg-reflecting crystals, it is imperative to maintain consistent temperature across the crystals to achieve a stable output [4, 5]. The temperatures of crystals C1, C2, C3, and C4 need to be maintained within +/- 0.4 °C temperature of one another [6]. To achieve this, a closed-loop heating system is implemented, with a temperature sensor and a heater at each crystal holder. This system utilizes an integrated temperature monitor controller, equipped with independent Proportional-Integral-Derivative (PID) controlled heater outputs. Furthermore, the controller interfaces with the EPICS control system [7], allowing for coordinated temperature settings across the four diamond crystals. Cooling is managed by thermal straps via dedicated feedthroughs along the vacuum chamber [3].

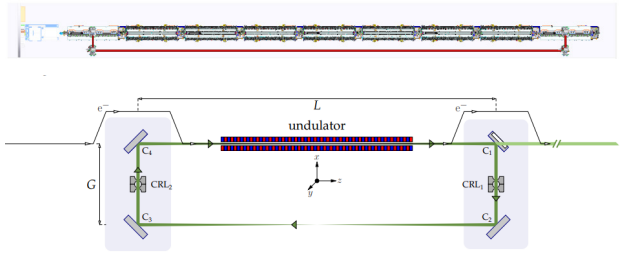



Figure 1: CBXFEL experimental layout [4].

## REQUIREMENTS

To maintain the optimal performance of the crystal mirrors, it is essential to regulate the temperature precisely. A temperature change of the crystal leads to a shift in the center of the angular Bragg reflection curve [4]. Therefore, the temperature of all cavity crystals must be kept within +/- 0.40°C of the setpoint, with an accuracy of +/- 0.10°C and a precision of <= 0.01°C [8]. Furthermore, the temperature of all the four crystal holders must be maintained 3˚C above the nominal temperature setpoint of the SLAC undulator hall, which is 20˚C [3].

The general thermal management requirements are summarized in Table 1.

Table 1: General Requirements

| Description | Value with units |
|---|---|
| Accuracy | 0.10 °C |
| Precision | 0.01 °C |
| PID Response Time | 1 minute |
| Control deadband | +/-0.40 °C |

It is critical that the selected temperature controller is capable of detecting sensor faults and inhibiting heater operation during such events. The temperature controller should include a configurable maximum output for the

* Work supported by U.S. Department of Energy under contract number DE-AC02-06CH1 1357 (ANL) and DE- AC02-76SF00515 (SLAC)
† namrata@SLAC.Stanford.EDU

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD097

# VACUUM MONITORING AND CONTROLS ACCEPTANCE TESTING FOR LCLS-II HE CRYOMODULES AT SLAC*

R.S.P. Chillara†, Z. Huang, A. Wilson, D. White, C.D. Lai, S. Saraf
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

## *Abstract*

The LCLS-II High Energy (LCLS-II HE) upgrade at SLAC National Accelerator Laboratory involves the delivery and long-term storage of superconducting cryomodules, manufactured at partner laboratories and shipped to SLAC. To maintain their performance, these cryomodules must preserve ultra-high vacuum (UHV) conditions prior to installation. Upon arrival, each unit undergoes a comprehensive controls acceptance test to validate instrumentation functionality, including sensor accuracy, signal integrity, and system connectivity.

To ensure vacuum integrity during extended storage, a vacuum monitoring system was developed to continuously track Cold Cathode Gauge pressure, integrated into the EPICS control system for real-time data, alarms, and archiving. Despite using a single gauge per cryomodule, the system enables early detection of vacuum degradation with minimal hardware. This paper outlines the acceptance testing, vacuum monitoring system design, and operational experiences, along with data trends, alarm scenarios, and lessons learned to improve future cryomodule storage practices.

## CONTROLS ACCEPTANCE TESTING METHODOLOGY

Upon arrival at SLAC, each cryomodule is subjected to a rigorous controls acceptance test. This procedure is not a functional performance test, but rather a methodical verification of the integrity and connectivity of the crymodule's embedded instrumentation to make sure they survived the transport intact. This process is critical: early detection of mechanical or electrical issues spares costly delays during installation or commissioning [1, 2].

### *Visual and Mechanical Inspection*

The inspection begins with a meticulous visual inspection of all electrical feedthroughs and connectors, commonly referred to as flanges (Fig. 1). We pay close attention to these interfaces for any signs of mechanical stress, such as hairline cracks, bent pins, or thread damage, which can compromise both electrical and vacuum integrity. Documenting these initial observations is a crucial first and is escalated to respective departments in case we find any issues [2].

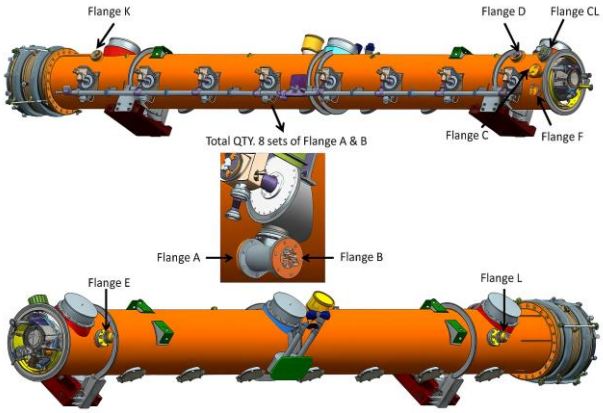



Figure 1: Instrumentation feedthrough flanges mounted on the vacuum vessel.

### *Systematic Electrical Validation*

The core of the acceptance test is systematic electrical validation of the various sensors and actuators. This is performed using a mobile test enclosure equipped with calibrated instruments such as a Lakeshore 240 temperature monitor, a Bartington Mag-01H fluxgate magnetometer, and a Fluke 87V multimeter. We use a portable test kit (Fig. 2) for convenient testing at multiple storage locations.



Figure 2: Mobile controls test enclosure.

- Temperature Sensors: All cryogenic sensors (Cernox RTDs, silicon diodes, Pt100 RTDs) are evaluated with the Lakeshore 240 (via MeasureLINK™). Resistance, voltage, and temperature readings are compared against vendor calibration curves and acceptable operational windows. This evaluation also includes a

* Work Supported by U.S. Department of Energy under contract number DE-AC02-76SF00515
† priyach@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD100

# DEVELOPMENT OF EPICS DEVICE SUPPORT ON THE MELSEC iQ-R EMBEDDED LINUX CONTROLLER

A. Uchiyama†, M. Komiyama, K. Yamada, N. Fukunishi
RIKEN Nishina Center, Wako, Japan

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Commercial programmable logic controllers are used for various accelerator subsystems at the RIKEN Radioactive Isotope Beam Factory (RIBF). The Yokogawa FA-M3 series with a Linux CPU module running an Experimental Physics and Industrial Control System (EPICS) input/output controller has been successfully adopted. In response to user requests, we later implemented an embedded Linux controller (RD55UP12-V), which has a similar architecture and is part of the MELSEC iQ-R series from Mitsubishi Electric. The conventional MELSEC integration with EPICS typically relies on asynchronous device support over TCP/IP. However, this configuration can suffer from reliability issues such as socket disconnections and failed reconnections after network disruptions. To address these problems, we implemented EPICS device support that runs on the MELSEC iQ-R embedded Linux controller; therefore, no TCP/IP socket connection to an external host is required. A Python-based prototype was created using PyDevice for rapid development and verification by incorporating a thread-safe worker architecture. Based on its success, we re-implemented the support in C++ using the asynPortDriver class of the EPICS asynDriver module. The C++ implementation adopts a standard asynDriver pattern, improves the readability and maintainability, and adds support for record types that are not available with the PyDevice prototype. This paper presents the development process, differences between the Python and C++ implementations, and the results of system integration within the RIBF control environment.

## INTRODUCTION

The control system used at the RIKEN Radioactive Isotope Beam Factory (RIBF) is based on the Experimental Physics and Industrial Control System (EPICS). Many front-end devices, particularly commercial programmable logic controllers (PLCs), are attached to control networks over Ethernet [1]. These devices are integrated through EPICS extensions such as NetDev [2], asynDriver [3], and StreamDevice [4], as reported in [5]. Although Ethernet is widely used and easy to handle in the physical and data-link layers, communication between EPICS input/output controllers (IOCs) and TCP/IP devices uses asynchronous socket connections. Unexpected power losses or switch reboots can cause disconnections, and automatic recovery is not always guaranteed. This behavior has raised operational concerns.

To address this issue, the RIBF has adopted Yokogawa FA-M3 PLCs with Linux CPU modules (F3RP61/F3RP71), on which an EPICS IOC runs directly [6–10]. Placing the IOC on the same backplane as the PLC removes the external TCP/IP path, which avoids the reconnection failures described above and increases the reliability of communication between the IOC and devices. RIBF also introduced Mitsubishi Electric MELSEC PLCs, which have a larger market share in the general industry. Historically, EPICS has accessed these PLCs from a separate Linux host via asynchronous device support over TCP/IP, typically using the MELSEC communication protocol [11-13]. In this study, we sought to improve the communication reliability of such systems by executing the EPICS IOC directly on the MELSEC iQ-R C-language intelligent function module (RD55UP12-V), which provides an embedded Linux environment. We focused on the device-support design and initial implementation of the RD55UP12-V.

## OVERVIEW OF THE RD55UP12-V

The RD55UP12-V is a C-language intelligent function module for the MELSEC iQ-R PLC series (Fig. 1). It is designed to be used together with a sequencer CPU running ladder logic, and exchanges data with the CPU via buffer memory. The module can be installed in any slot on the standard base unit; no dedicated slot is required and it does not interfere with other modules. A Debian-based Linux distribution can be installed, enabling application development in C and Python using standard tool chains and libraries, which lowers the barriers to introducing EPICS and other control frameworks. For PLC communication, the vendor-supplied C API supports the MELSEC Datalink Library for iQ-R (MDR) and direct buffer-memory access, which are suitable for industrial control with real-time requirements. The hardware used in this study features an ARM® Cortex®-A9 dual-core CPU and 1 GB RAM, running Debian 10 (Buster) [14, 15].




Figure 1: MELSEC iQ-R test setup used for evaluation. The R08 sequencer CPU module is installed at the left, and the RD55UP12-V C-language intelligent function module (embedded Linux controller) is installed in the rightmost slot of the base unit.

†a-uchi@riken.jp

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD103

# DCS DATA TOOLS: A TIME-SERIES DATABASE SOLUTION FOR ATLAS DETECTOR CONTROL SYSTEM*

P. Moschovakos†, D. Matakias‡, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.


## *Abstract*

The ATLAS Detector Control System (DCS) has accumulated over 68 TB of time-series data since 2007, archived in the Oracle production database. While Oracle ensures long-term reliability, the existing workflow consisting of Oracle, the DCS Data Viewer and custom scripts has become a bottleneck for large-scale analysis, visualization, and accessibility. To address these limitations, we present DCS Data Tools (DDT), a modern time-series database solution based on TimescaleDB that complements the current infrastructure while maintaining compatibility. The architecture includes an Oracle-to-TimescaleDB bridge, deployed independently to avoid risk to production. TimescaleDB compression reduces storage requirements and optimized ingestion makes full migration practical. The system provides APIs in Python and C++, a backend server that abstracts database complexity, and Grafana integration for visualization. DDT has already been adopted by Monitored Drift Tube (MDT) and Liquid Argon (LAr) experts for new monitoring and predictive-maintenance applications. Fully aligned with the NextGen Archiver, DDT ensures continuity with existing infrastructure while positioning ATLAS for scalable, low-latency, and analysis-driven data access.


## INTRODUCTION

### *Background and Limitations*

The ATLAS experiment at CERN has been continuously operational since 2007, generating extensive time-series data from its Detector Control System (DCS). These data, essential for detector operation, monitoring, and analysis, have been archived in the ATLAS Oracle production database (ATONR), which now exceeds 68 TB. While Oracle has ensured reliable long-term storage and data integrity, limitations have emerged from the overall data-access chain. The original DCS data architecture was designed primarily for operational monitoring and basic historical queries. Over the years, incremental extensions were introduced, including the DCS Data Viewer (DDV) [1] and custom extraction scripts, to meet emerging analysis needs. While these additions provided short-term solutions, they introduced complexity and scalability challenges. Users frequently encounter slow query performance, especially for multi-year ranges, and must rely on specialized knowledge of Oracle SQL and its schema. As a result, extracting meaningful insights often requires substantial effort and technical expertise.

### *Community Requirements and Solution*

These limitations are increasingly at odds with the scientific community's growing demand for flexible access to DCS data. To systematically address these challenges, ATLAS DCS conducted a comprehensive requirements survey across the experiment community. Key findings revealed that 86% of users preferred Grafana for visualization, 92% required flexible search functionality, and over 90% demanded continuous service availability with full historical data access. Modern use cases include multi-detector correlation studies, advanced monitoring applications, and integration with external analysis frameworks.

DCS Data Tools (DDT) was developed as a comprehensive solution based on these systematically collected requirements. Built on TimescaleDB, a PostgreSQL-based time-series database, it introduces modern database technology. The system design follows three objectives: (1) improve query performance, (2) lower barriers for end users through simple APIs, and (3) establish a scalable and future-proof architecture. Importantly, DDT is aligned with the NextGen Archiver (NGA) developed by CERN and Siemens [2], ensuring continuity with existing infrastructure while preparing for a seamless transition to next-generation data flows. By combining Oracle reliability with TimescaleDB performance and modern interfaces, DCS Data Tools provides the ATLAS experiment with a robust foundation for both current and future data-analysis needs.

## SYSTEM ARCHITECTURE & DATA INGESTION

### *Architecture Overview*

The DCS Data Tools architecture provides a comprehensive solution for time-series data retrieval, designed specifically to address the limitations of the traditional data-access chain combining the Oracle database, the DCS Data Viewer (DDV), and custom extraction scripts. The system consists of several interconnected components that work together to provide high-performance data ingestion, storage, and retrieval.

The primary architecture follows a pipeline approach, with data flowing from WinCC OA through the NGA to Oracle, then to TimescaleDB, through the backend server, and finally to client APIs (Fig. 1). This design maintains backward compatibility with existing Oracle and WinCC OA infrastructure while introducing modern time-series capabilities through TimescaleDB. This compatibility encompasses three key aspects: (1) the future capability to configure WinCC OA to archive directly to TimescaleDB, (2) the current database schema and software chain maintain full compatibility with existing DCS data structures without limitations, and (3) the

* Work supported by CERN
† paris.moschovakos@cern.ch
‡ dimitrios.matakias@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD104

# SEVEN YEARS OF PROGRESS: OPTIMISING CONTROLS HARDWARE DATA MANAGEMENT AT CERN FOR ENHANCED EFFICIENCY AND RELIABILITY

E. Fortescue†, I. Kozsar, J. Palluel, M. Suminski, E. Van der Bij, CERN, Geneva, Switzerland

*Abstract*

The CERN Controls Hardware Data Management (CHDM) initiative was launched in 2018 with the goal of improving the efficiency of the Hardware Installation Team and equipment groups by allowing them to easily manage and exploit their data when installing, maintaining and dismantling large-scale electronics installations.

After seven years of refining tools, optimising data flows, both internally and between systems, as well as cleaning and defining remaining data within the centralised systems, we have successfully achieved our primary objectives: obtaining valuable reliability statistics and providing auto-generated schematics and data-driven diagnostics tools.

This paper outlines the approach taken throughout this process and examines the costs and benefits associated with this investment.

## INTRODUCTION

The Hardware Installation Team (HIT) within CERN's Controls Electronics and Mechatronics (CEM) group is responsible for the installation and maintenance of a vast array of hardware components for the control of the CERN accelerator complex. This is comprised of 300 distinct parts and 37,000 individual assets, for the Front-End Computers (FECs) and their associated networks.

In 2018, the data that documented the controls hardware installations was distributed across poorly integrated database systems and files, and was often duplicated, inconsistent, incomplete, and difficult to maintain. Visualising complex installations, particularly electronics and cabling layouts, posed significant challenges, due to the reliance on manually updated schematics, resulting in inefficiencies and an increased risk of errors. Furthermore, there was a growing demand for data-driven diagnostic tools to support maintenance and troubleshooting. However, it was recognised that such tools and visualisations could only be effective if built upon a foundation of reliable data.

To address these issues, the Controls Hardware Data Management (CHDM) initiative was launched, with the mandate to establish a coherent data management policy. The scope of each system was clearly defined, the tools were improved and the data flow between the systems was streamlined. In parallel, a significant effort was invested into improving data quality, striving to eliminate inconsistencies and errors by cleaning, consolidating and completing existing datasets, whilst enforcing clearer processes and standards for new data.

By correctly structuring the installation data in CERN's centralised database systems, the long-term vision was to obtain exploitable information and actionable insights, via reports, visualisations, statistics, auto-generated schematics and diagnostic tools to support decision-making and facilitate maintenance operations. Seven years later, these sustained efforts have culminated in the successful achievement of CHDM's primary objectives.

## DATABASES

Figure 1 illustrates that the foundation of CHDM is Quality Assurance, with a central component being the dictionary of CERN-wide naming conventions. This dictionary is maintained in the Accelerator Naming Database and provides a unifying data thread between the various database systems.

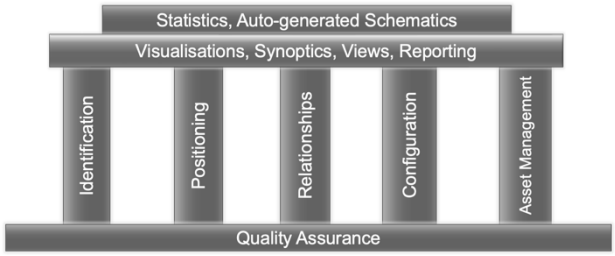



Figure 1: The 5 pillars of CHDM.

Data related to the identification and positioning of hardware, such as racks, crates and modules, as well as the physical connections between them, including cabling, is recorded in the Layout database. Hardware elements are represented as Functional Positions, which act as placeholders defining the type, name, size and location of accelerator components. These Functional Positions are organised in hierarchies, known as assembly breakdown structures (ABS). Each ABS is valid for a specified time period, thus facilitating accurate modelling of infrastructure changes over time.

Through seamless integration between the systems, the physical data model of the installations in the Layout database is exposed to the Controls Configuration Database (CCDB), enabling it to be supplemented with logical configuration data. Together with a declaration of the computers in the Network Database, this allows the control system to operate correctly.

Finally, the Functional Positions representing the theoretical hardware installation are published to CERN's Enterprise Asset Management (EAM) system, so that assets can be assigned. Assets are registered in the EAM system for any operational equipment that must be individually tracked throughout its lifecycle, such as electronic boards, chassis, power supplies, fans, etc. Upon arrival at CERN, each asset is labelled with a corresponding barcode or QR

† Eve.Fortescue@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD109

# FROM TERABYTES TO PETABYTE: SCALING OF THE ARCHIVING SYSTEM FOR FAIR

V. Rapp†,1, K. Klimczyk[2], J. Kerssemakers[1], P. Salabura[3]
[1]GSI, Darmstadt, Germany
[2]S2Innovation, Krakow, Poland
[3]UJ, Krakow, Poland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

With the recent rise of AI and various machine learning models, the importance of storing and managing data generated by control systems is greater than ever before. In 2016, GSI began developing an archiving system to collect, store, and retrieve data from the diverse accelerator devices managed by the GSI control infrastructure. The system was successfully deployed in production in 2021. To evaluate its capabilities and suitability for operational needs, the system was initially launched with a limited storage capacity of 50 TB and reduced computing power. With a current data volume of over 100 GB per day, the archiving system quickly exceeded its initial limits. However, the experience gained in day-to-day operations thus far has allowed us to better understand our use-cases and identify areas for further improvement. In preparation for the anticipated start of FAIR operations in 2027, the system will require significant scaling to meet future demands. Therefore, this is an opportune moment to review and refine the system's architecture based on the experience gained so far. This paper outlines the challenges encountered with the current implementation and presents the solutions that will be incorporated into the system for FAIR operations.

## INTRODUCTION

As part of the controls renovation for the FAIR project, GSI initiated the development of an archiving system in 2016 [1]. After several years of development, the first production-ready version was released in 2021. The primary objective of the system is the collection, storage, and presentation of data generated by control devices and other relevant infrastructure. This system was the first of its kind and scale at GSI and posed several challenges in terms of requirement analysis and assumptions regarding scalability. Consequently, the architecture was designed with a high degree of flexibility with respect to incoming data structures and the ability to accommodate high data rates. Elasticsearch was selected as the storage backbone, leveraging the department's prior experience with several self-hosted Elasticsearch installations at GSI, which provided the necessary operational expertise.

In order to build operational knowledge, evaluate system behavior under realistic conditions, and prepare for FAIR-scale operation, the system was initially deployed with limited resources. The storage backbone comprised three data nodes with a total capacity of 30 TB, while the application components were deployed on two dedicated nodes. As storage rapidly emerged as a bottleneck, capacity was extended to 50 TB. After resolution of initial operational issues, the system demonstrated stable performance and is currently employed in several important use cases [2]. Presently, the system ingests over 100 GB of data per day. To mitigate storage constraints, data is decimated after each beam time, which presently spans approximately six months. In 2024, following modifications to the storage format and migration from Elasticsearch to OpenSearch, the system was integrated with Grafana, thereby providing users with an efficient graphical interface for data access and visualization (Fig. 1).



Figure 1: One of the Grafana dashboards.

While the system has proven adequate for numerous use cases, operational experience has also revealed several architectural limitations. The design philosophy of maximum flexibility has resulted in comparatively high maintenance demands. This flexibility requires careful configuration and user discipline to ensure efficient operation. Furthermore, although Elasticsearch exhibits good performance with smaller datasets, query latency increases substantially with growing index sizes. In addition, the analytical capabilities of the system remain limited and do not fully meet contemporary expectations.

With the commencement of FAIR operation anticipated in 2027 and an extended beam intervention scheduled between 2025 and 2026, GSI has initiated a comprehensive refactoring of the system to meet future requirements. The primary objectives of this effort are to develop a simplified, maintainable architecture and to enable scalability to the petabyte regime.

## SYSTEM ARCHITECTURE OVERVIEW

Similar to the current implementation, the refactored system will retain a layered architecture with a relatively straightforward data flow (Fig. 2). Data is collected by

† V.Rapp@gsi.de

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD110

# Bluesky NeXus: A SOLUTION FOR NEXUS-COMPLIANT DATA ACQUISITION IN BLUESKY*

D. Tomecki[1,†], M. Bajdel[1], D. Burke[2], A. Dillmann[1], E. Lott[1], S. R. Patel[1], L. Porzio[1], W. Smith[1], S. Vadilonga[1]
[1]Helmholtz-Zentrum Berlin für Materialien und Energie GmbH, Berlin, Germany
[2]Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

*Abstract*

Modern scientific experiments require rich, standardized metadata to ensure that data is Findable, Accessible, Interoperable, and Reusable (FAIR [1]). The *NeXus* format [2], a hierarchical data standard widely used in neutron, x-ray, and muon science, provides a robust framework for organizing such metadata. However, automating its integration into data acquisition workflows remains a challenge.

We present Bluesky NeXus [3], a Python package that enables automated, standards-compliant *NeXus* file generation within *Bluesky* [4]—a modular data acquisition framework deployed at synchrotron and neutron facilities. Users of the package define the desired *NeXus* device structure—including groups, datasets, and attributes—using human-readable YAML configuration files. These are validated with *Pydantic* [5], a Python library for schema validation, ensuring consistency with *NeXus* definitions.

Bluesky NeXus captures both static metadata (e.g., instrument configuration) and dynamic measurements, consolidating them into a complete *NeXus* file automatically archived with each experiment. It integrates with the containerized *Bluesky* environment used at BESSY II [6], supporting a wide range of experimental setups.

Developed as part of the ROCK-IT [7] project, Bluesky NeXus simplifies the creation of FAIR-compliant metadata and enhances the reproducibility and interoperability of scientific data.

## INTRODUCTION

Efficient data collection and storage are essential for ensuring reproducibility, accessibility, and interoperability in modern scientific research. Large-scale facilities such as synchrotrons and neutron sources rely on complex instrumentation that produces vast amounts of data. The *Bluesky* framework is widely used to manage experimental workflows, enabling researchers to design and execute scans while capturing structured metadata. However, storing this information in a consistent, standards-based format remains challenging, particularly when multiple instruments contribute to a single experiment.

The *NeXus* data format, built on the HDF5 (Hierarchical Data Format) standard, provides a hierarchical, self-describing structure for organizing, sharing, and preserving scientific data. Integrating *NeXus* with *Bluesky* runs ensures that both metadata and measurements are stored in a standards-compliant manner, enabling seamless retrieval and promoting reproducibility across research environments.

Capturing the configuration of all participating instruments—such as detectors, motors, and environmental sensors—is essential. A structured *NeXus* file preserves this information alongside experimental results, supporting downstream analysis and aligning with FAIR data principles.

This paper introduces Bluesky NeXus, a software package that automates *NeXus* file generation within the *Bluesky* framework. It employs a *Pydantic*-based mechanism to define and validate *NeXus* structures, ensuring compliance with international conventions while supporting user-defined metadata extensions. By standardizing storage and automating metadata capture, Bluesky NeXus strengthens data integrity and facilitates collaboration across experimental facilities.

## OBJECTIVE

The primary objective of this work is to design and implement a robust, extensible mechanism for generating a *NeXus*-compliant file from a *Bluesky* run. The resulting file serves as a structured, self-describing representation of experimental data, in alignment with community standards in neutron, X-ray, and muon science. Our goal is to serialize the data while ensuring it adheres to the hierarchical structure and naming conventions defined by the *NeXus* format.

The file structure is defined to include two principal groups: `NXinstrument` and `NXcollection`, both of which are *NeXus* base classes [8]. This hierarchical layout is illustrated in Fig. 1, which shows an exemplary NeXus file structure including four instruments.

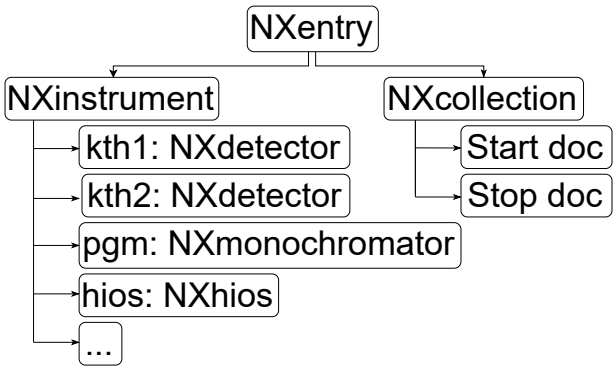



Figure 1: Illustration of the *NeXus* file structure.

The `NXinstrument` group is central to our implementation. It encapsulates all relevant information about the

* Work funded by ROCK-IT (Remote, Operando Controlled, Knowledge-driven, and IT-based) project.
† Corresponding author: daniel.tomecki@helmholtz-berlin.de

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-TUPD113

# UNSUPERVISED ANOMALY DETECTION IN ALS EPICS EVENT LOGS

A. Sulc*, T. Hellert, S. Hunt, LBNL, Berkeley, CA, USA

*Abstract*

This paper introduces an automated fault analysis framework for the Advanced Light Source (ALS) that processes real-time event logs from its EPICS control system. By treating log entries as natural language, we transform them into contextual vector representations using semantic embedding techniques. A sequence-aware neural network, trained on normal operational data, assigns a real-time anomaly score to each event. This method flags deviations from baseline behavior, enabling operators to rapidly identify the critical event sequences that precede complex system failures.

## INTRODUCTION

The Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory is a third-generation synchrotron light source facility serving a diverse international community of scientific users [1]. Its operation relies on a complex, distributed control system based on the Experimental Physics and Industrial Control System (EPICS) toolkit [2]. This system comprises hundreds of Input/Output Controllers (IOCs) that manage over 200,000 process variables (PVs), each corresponding to a specific signal for monitoring or control—from magnet power supplies and vacuum gauges to insertion devices and beamline optical components.

Maintaining high availability and operational efficiency in such a complex environment necessitates rapid fault diagnosis. When the electron beam is lost or a critical subsystem fails, operators face the challenge of identifying the root cause amidst a flood of alarms and status changes that often precede or accompany the actual fault. A single fault can trigger a cascade of secondary events, generating thousands of log entries within seconds and obscuring the initial cause. The ALS employs an "Event Logger" system, which chronologically records state changes for thousands of critical PVs. While this logger provides an invaluable, high-fidelity record of system activity, its sheer volume and unstructured, text-based nature make manual post-mortem analysis time-consuming and heavily reliant on the accumulated expertise of senior staff.

Conventional monitoring often relies on predefined limits and alarms set on individual PVs. Although essential, this approach may not capture subtle precursor events or faults arising from complex, emergent interactions between subsystems. This paper proposes a machine learning framework designed to overcome these limitations. We reframe the problem by treating the Event Logger data as a sequential language and apply modern natural language processing (NLP) and deep learning techniques to learn the "grammar" of normal machine operation. Our system automatically detects and scores anomalous event sequences, providing a quantitative measure of deviation from normal operation for control room operators.

This work aims to shift the diagnostic paradigm from reactive, alarm-based monitoring to a proactive, data-driven approach, with two main contributions: (i) the application of NLP embeddings to create meaningful numerical representations of EPICS events for visualizing and interpreting the machine state; and (ii) a sequence-aware, one-class neural network model to detect anomalies in these event streams.

## METHODOLOGY

Our framework consists of two main stages: (1) data preprocessing and representation, where raw event logs are converted into numerical vectors, and (2) anomaly detection, where a deep learning model scores events based on learned normal patterns.

### *Channel Name Interpretation*

To better understand the inherent structure of the PV naming convention at the ALS, we developed a method to represent the relationships among channel name components as a graph. Although the EPICS naming scheme is mostly hierarchical, this structure can be obscured when viewing thousands of individual PVs. Our visualization technique helps reveal the underlying "grammar" of the control system's nomenclature, confirming that channel names are both parseable and scalable.

**ALS EPICS Channel Name Convention** The established naming convention follows the structure `SysSubSys:DeviceID:Signal`, with an optional `SubDevice` component included as `SysSubSys:DeviceId-SubDevice:Signal`. This format is designed for programmatic parsing, with the colon (`:`) as a reserved delimiter. The segment between the first and second colons constitutes the `DeviceID`, and everything following the second colon is the `Signal` name. To ensure consistency, `Sys` (System) and `Device` names are restricted to an agreed-upon list and cannot contain colons. Furthermore, device names must not end in a number to avoid ambiguity with the `DI` (Device Instance), which must be a number without leading zeros.

Several rules govern channel name components. Names are case-sensitive, and capitalization is encouraged for readability. The use of underscores should be minimized, though a signal name ending with an underscore (`_`) specifically denotes a private signal within an IOC. While the `SubSys` and `DI` fields are not always required, using a `DI` of 1 is recommended if more devices of the same type might be added later. The maximum length of a channel name is still under discussion, but developers should avoid characters such as `{}, [], ., -, +, =, ;, *, |, <, >`.

* asulc@lbl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEAG001

# THE TANGO CONTROLS COLLABORATION STATUS IN 2025

T. Juerges, SKA Observatory, Jodrell Bank, United Kingdom
R. Bourtembourg, A. Götz, D. Lacoste, ESRF, Grenoble, France
S. Rubio-Manrique, ALBA-CELLS Synchrotron, Cerdanyola del Vallès, Spain
B. Bertrand, A.F. Joubert, MAX IV Laboratory, Lund, Sweden
Y. Matveev, DESY, Hamburg, Germany
L. Pivetta, G. Scalamera, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy
G. Kowalski, T. Madej, M. Celary, L. Zytniak, S2Innovation, Kraków, Poland
G. Abeillé, G.Pichon, P.Madela SOLEIL, Gif-sur-Yvette, France
T. Braun, byte physics e. K., Annaburg, Germany
R. Auger-Williams, T. Ives, Observatory Sciences Ltd, St Ives, United Kingdom

*Abstract*

Since the last status update in 2023, the Tango Controls collaboration has undertaken a major effort to add new features to cppTango, the core of Tango Controls, and two other official language bindings, JTango and PyTango. Significant development efforts have been dedicated to the implementation and prototyping of community-requested features. Observability is a trending topic in software development, and we have listened to our community adding OpenTelemetry support. Continuing with cppTango refactoring, we switched to C++17 and catch2 as the new testing framework to improve code quality and test coverage. PyTango has undergone a major overhaul by switching from boost-python to pybind11, which has been a welcome modernization of the code base and has allowed us to remove obsolete APIs. Special Interest Group (SIG) meetings continued to be a great success. Several have been held, among them one that addressed the request of our users for a much improved documentation. Encryption has also been a SIG topic, and a prototype for complete end-to-end encryption of all communication in Tango Controls has been developed. CI/CD has again received major updates and gained more computing power to run more tests in less time, thanks to the Gitlab runner contributions of the collaboration members. Thanks to the continuous community effort on keeping a modern and well maintained core, the future road map of Tango Controls looks promising and achievable.

## INTRODUCTION

In 1998, the European Synchrotron Radiation Facility (ESRF) [1] submitted a paper on Tango, a control system framework based on the paradigm of distributed objects, to ICALEPCS 1999 [2]. Today, Tango is much more than a piece of software. First, it is a well organized, friendly, and amazingly fruitful collaboration of ten major facilities [3] and a growing community of users and developers. It is also the place of constant development, enhancement, refactoring, and innovation [4].

The success of the Tango collaboration relies on its organization, the members and developers, and also on the underlying contract that governs it. The financial contribution by the core members allows the collaboration, represented by its Steering Committee, to oversee its own budget, finance events and technical subcontracting. The latter definitely boosts the development of Tango. It is now much more than just a framework. It is a rich and very active software ecosystem that is constantly evolving. That means Tango improves and tries to benefit from the latest technical developments while maintaining strong backward compatibility for its users.

The present paper provides the reader with the latest news from the Tango Controls collaboration and its ecosystem. We aim to present a broad overview of both the organizational and technical activity around Tango.

## CPPTANGO

cppTango is the C++ implementation of the Tango kernel and is considered the reference implementation of Tango. Due to the performance and efficiency of C++, cppTango is often used for device servers (DS) that interface with hardware that require high speed communication or have to handle high volume data throughput. Examples are detectors or motors in synchrotron beamlines.

The latest major release (10.0.0) added observability features by integrating with the OpenTelemetry [5] framework for unified monitoring across many device servers. The goal is to make it easier to trace and diagnose problems in complex installations by correlating logs and metrics from multiple sources. Another major feature that has been implemented is the alarm event, which provides an automatic signaling by the device server when the device state changes to or from Alarm. New is also that device server developers have the possibility to add extended version information to device servers. Examples could be the source code version, versions of used libraries or firmware revisions of connected hardware. The latest cppTango patch release (10.1.0) has added performance monitoring of the ZeroMQ event subsystem, new fully asynchronous event subscriptions, and has fixes for all known memory leaks. Despite all the bug fixes, improvements, changes of internal and external APIs and ABIs, cppTango has maintained backward compatibility with existing Tango device servers which were developed

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEAG002

# REDIS-BASED STREAMING ARCHITECTURE FOR ACCELERATOR BEAM INSTRUMENTATION DAQ SYSTEMS*

S. Joshi†, D. Steinkamp
Fermilab, Accelerator Division, Beam Instrumentation Department, Batavia, IL, USA

*Abstract*

The Fermilab Acceleraor Division, Beam Instrumentation Department, is always adopting modern and current software methodologies for complex DAQ architectures. This paper highlights the Redis Adapter (RA) as the key software component enabling high-performance, modular communication between digitizers and distributed control systems by leveraging Redis and containerization. The RA provides a unified, efficient interface between Redis-based data streams and consumer systems. In the legacy architecture, digitized data flowed through the custom, UDP-based Distributed Data Communication Protocol in the middle layer. In the current system, DDCP remains the ingestion path, while the RA serves as the decoupling layer. The proposed system replaces old VME digitizers with a SOM-based digitizer that communicates with Redis using the RA. The RA acts as both a performance-critical bridge and a protocol-agnostic adapter, ensuring compatibility with legacy control frameworks while enabling future scalability and modularity. This restructuring of the middle layer also helps the system achieve high throughput, reduce latency, and simplify the data path. Finally, we will demonstrate how RA is utilized in our two core products to deliver both legacy compatibility and future flexibility.

## INTRODUCTION

The Fermilab Beam Instrumentation (BI) Department monitors lifecycle diagnostics from end to end of beam systems across four major categories: Beam Current Monitors (BCM), Beam Loss Monitors (BLM), Beam Position Monitors (BPM), and Beam Profile Monitors (BProM). These systems provide critical data for accelerator stability, experimental optimization, and fault detection using instrumentation software that has evolved over decades to interface with VME-based digitizers and custom UDP protocols dating to the 1980s. While this traditional software lineage provides proven operational reliability, it increasingly limits the flexibility necessary for contemporary accelerator programs due to its tightly coupled nature. With the VME system's monolithic architecture and hardware constraints, there are development bottlenecks that make it no longer possible to meet the modern operational needs. The Proton Improvement Plan II (PIP-II) serves an urgency to modernization. The PIP-II will deliver world-record beam intensity for the DUNE neutrino experiment as the first U.S. particle accelerator to be built with substantial contributions from around the world and will facilitate diverse on-site physics programs over fifty years of operation. The extremely high-performance requirements of a superconducting linear accelerator, including never-before-seen intensities for beams, operational flexibility, and international collaboration workflows, reveal the fundamental limitations of legacy DAQ architectures. This paper discusses a strategic software solution for PIP-II that preserves existing control system expertise like ACNET (Accelerator Control Network) while enabling more options like EPICS (Experimental Physics and Industrial Control System) and architecture flexibility for the DAQ system. Through General Redis ACNET Front End (GRAFE) and General Redis EPICS Frontend (GREFE) implementations, we illustrate how RA bridges legacy systems with modern streaming architectures, providing evolutionary migration that simultaneously fulfils PIP-II's bold imperatives.

### *Legacy Daq Architecture & Limitations*

Traditional VME-based systems feature tightly coupled components, including controller cards, digitizers, and PMC-UCD modules. Data from digitizers flows to control systems through specific VME drivers, kernels, and DDCP (Distributed Data Communication Protocol) implementations. Each consumer embeds device-specific protocol logic, duplicating effort and complicating deployments.

This architecture presents several major limitations:

- Scalability Constraints: Addition of new consumers or increasing sample counts required intrusive changes across multiple system layers and usually required coordinated downtime.
- Protocol Complexity: UDP-based implementation of DDCP was performant but required maintenance overhead and limited interoperability with modern containerized environments.
- Metadata Handling: Waveform metadata such as timestamps and sampling increments used ad hoc conventions, resulting in inconsistent data integration.
- Resource Allocation: Aging VME platforms required substantial sustaining engineering resources and diverted resources from new instrument development.
- Monolithic Coupling: Direct hardware-to-consumer connections led to brittle dependencies, making it harder to test, debug, or evolve.

Strict coupling between the data acquisition hardware and control system consumers in the legacy architecture imposed a maintenance burden that ultimately limited operational flexibility and development agility.

### *Modern System Overview & Container Architecture*

The modernized architecture introduces clear separation of concerns through a Redis-based streaming system combined with comprehensive containerization and automated deployment pipelines. This integrated approach addresses both the technical limitations of legacy systems and the

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEAG003

# INTRODUCING WEB BASED TECHNOLOGIES AT GANIL SPIRAL2 CONTROL SYSTEM

O. Delahaye*, J. Girault, C. Vivier, C. Tréangle, GANIL, Caen, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The SPIRAL2 accelerator began operating in 2021. One of the key application of the control system is the management of all devices parameters (magnets, RF ...), roughly 80000 EPICS variables. That application is fundamental for optimizing the setup time of the accelerator and for easily reproducing the configuration of a given beam from year to year. Because web-based technologies are believed to offer many advantages—such as portability, easier maintenance, optimized use of hardware resources, and centralized security—we decided to evaluate this technology in order to form our opinion in the perspective of a wider renovation project. This paper will explain how the software architecture is designed, both on client and server side, what technologies we used (web framework, REST API, web server, database and ORM). It will also describe the outcomes we achieved in terms of features of the application such as beam characteristics management, reload of a given beam configuration and application to the devices, storage of the accelerator setup, calculation of parameters based on the concept of optic configurations. After 4 months of operation in 2024 with that new application, we will also discuss around the question: are web based technologies a good choice for SPIRAL2 control system user interface?

## GANIL ACCELERATOR FACILITY

The GANIL research center currently hosts two accelerators. The first, known as the original GANIL, dates back to the 1980s. The second is the SPIRAL2 accelerator, which features a superconducting LINAC and an EPICS control system — the focus of this paper.

The SPIRAL2 accelerator includes a deuteron source and an ion source with a charge-to-mass ratio (Q/A) of 1/3, along with a LINAC composed of 26 superconducting cavities. Currently, only the NFS (Neutrons For Science) experimental area is operational. The second area, S3 (Super Separator Spectrometer), is in the commissioning phase. See Fig. 1.

Operating the SPIRAL2 accelerator requires control of approximately 3,000 devices, representing around 80,000 physical parameters.

## CONTEXT AND MOTIVATIONS

In 2021, the SPIRAL2 facility entered operational mode following the successful commissioning of a deuteron beam in the NFS experimental area for neutron production. In 2022, the LINAC delivered its first oxygen beam. Over the years, the increasing diversity of available beams (argon, proton, oxygen, etc.) has made the use of a dedicated tuning application essential. The operations team requires a powerful tool capable of efficiently tuning the accelerator's 3,000 devices within a short time frame. This application must be able to:

- Facilitate complex parameter calculations using the concept of optical configurations and transmit the results in a single operation to the targeted devices
- Store all parameter values associated with a given beam in a database and easily reload them for future experiments
- Retrieve live values from a set of devices in one operation and compare them to those stored in the database
- Manage changes in energy, optics, and beam characteristics

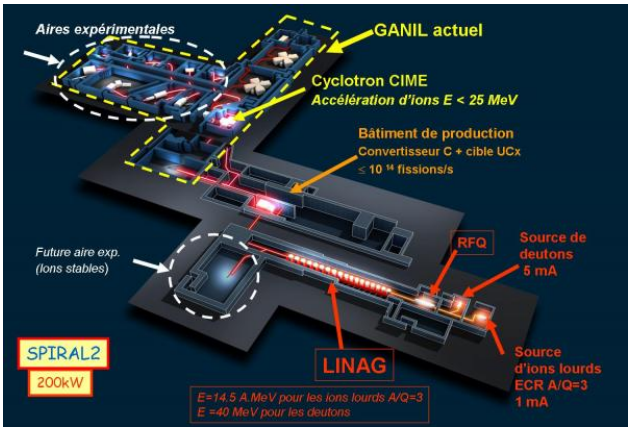



Figure 1: SPIRAL2 beam path.

## TECHNOLOGIES CONSIDERED FOR THE APPLICATION

The SPIRAL2 control system is based on the EPICS framework. Several architectural options were considered for developing this application:

- **Thick client**, which is the current solution at SPIRAL2, using either Phoebus or a Java Swing framework interfaced with the Java Channel Access (JCA) library to access Process Variable (PV)
- **Thin client**, such as Python QT (PyQt) application, interfaced with web services to access both the PV and databases
- **Web client**, developed with frameworks like Google Flutter or React, communicating via web services

For the thick client approach, using EPICS CSS-Studio Phoebus was not feasible due to the complexity of the application, which involves database integration, advanced algorithms, and the use of sophisticated GUI components. Although an alternative could have been our existing Java

* olivier.delahaye@ganil.fr

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEAG004

# ROBUST AND MAINTAINABLE OPERATION USING BEHAVIOR TREE SEMANTICS

W. Van Herck†, R. Gunion, R. Lange, G. Pospelov
ITER Organization, Saint-Paul-lez-Durance, France

*Abstract*

The ITER machine's inherent complexity and the diverse operational phases, such as commissioning and engineering operations, present significant challenges in balancing operability, integration, and automation.

The open-source software framework *oac-tree* was designed to create, maintain, and execute operational procedures in a robust and repeatable manner. Its semantics are based on a behavior tree model, which inherently supports reactive behavior, making it well-suited for goal-oriented tasks.

To accommodate diverse use cases — from small-scale tests to fully integrated operations — the framework's architecture was built with composability and extensibility in mind. System-specific interfaces and user interactions are fully decoupled from the core library, ensuring flexibility and adaptability.

The oac-tree library is deployed in production at ITER, offering a command-line interface for executing procedures as system daemons or for interactive use. Its maintainability is ensured by a minimal, low-complexity codebase with well-encapsulated third-party dependencies.

The oac-tree ecosystem also includes a Qt-based graphical user interface and a server enabling multiple clients to interact with running procedures. A plugin mechanism supports framework extensions, with existing plugins available for EPICS-based control systems, mathematical expression evaluation, common control system logic, and ITER-specific network protocols.

## INTRODUCTION

During the different phases of the ITER machine, many operational procedures need to be defined, verified and executed in a traceable and maintainable way.

For high-level operations, these procedures roughly fall into one of the following categories:

- Monitoring: low level signals are aggregated into higher level statuses that can then be consumed by central control applications.
- Control: high-level goals are expressed as specific combinations of low-level states or goals. This allows applications to reuse those control procedures, avoid duplication of detailed implementations and facilitate their maintenance.
- Automation: this typically combines both monitoring and control to be able to react to changing machine state in the most appropriate way.

Regardless of their category, these procedures will need to evolve during the lifetime of ITER for several reasons:

- Change of environment: as more plant systems are being added, their presence may require adapting monitoring and control logic.
- Increased automation: manual procedures can be replaced by automated ones as they mature and their interactions with the ITER plant are better understood.
- Improved performance: better understanding will also lead to opportunities to fine-tune procedures to maximize operational efficiency.

Furthermore, during integration and commissioning of plant systems, procedures will need to be executed to carry out specific tasks in a repeatable and traceable way. Those procedures are likely to be adapted even more often.

Oac-tree [1] is a software tool meant to facilitate the creation, adaptation, approval and execution of operational procedures. It also defines a text-based format for these procedures, allowing version control and easy traceability. It provides a means to replace procedural documents with unambiguous instructions that can be executed by the tool.

Using extensions that implement the ITER protocols for plant system configuration and monitoring, oac-tree can be effectively used as the main tool for central control and automation of the ITER plant. With proper extensions, it can just as easily be deployed in other environments, such as during commissioning of plant systems as mentioned above.

## GOALS AND REQUIREMENTS

One of the main goals for the ITER control system is to maximize operational efficiency, i.e. to reduce the time spent on tasks that do not directly contribute to the scientific mission of the machine.

Automation of manual procedures is a critical step towards this goal, but care has to be taken that the benefits are not offset by increased costs, due to maintenance and manual interventions.

While high-level programming languages, like C++, Java or Python, provide almost unlimited flexibility, they come at a high cost in terms of maintenance, since small code changes can introduce unexpected instabilities that may be hard to detect with unit testing. Hence, the use of these languages at the level of operational procedures would require extensive test plans, often involving complex test environments and considerable human resources.

Software frameworks that expose some higher-level semantics, like finite state machines, solve this problem by having a more constrained interface that can be tested more easily during automated tests. This trade-off between flexibility and robustness is quite common and finding the right balance between these can be challenging and requires clear understanding about the functional requirements of the framework within its operational context.

† walter.vanherck@iter.org

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEAG005

# FACILITATING USER CODE INCLUSION IN A SUPPORTED ENVIRONMENT TO ENHANCE THE FLEXIBILITY OF BEAMLINE CONTROL FRAMEWORKS

D. Perl*, K. Wakonig, J. Wyzula, C. Appel, S. Augustin, I. Usov
Paul Scherrer Institute, Villigen, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Beamline control system software is complex: it provides a common interface to a diverse range of hardware components, data services, and output file formats, any of which may not have been designed for ease of software integration. It should provide both simple and expert operation levels. End users are not expected to know the intricacies of the software, and facility staff will engage with it at a variety of levels. To enable users of the Beamline Experiment Control (BEC) software for SLS 2.0 at PSI to define custom and even one-off routines, while sustainably integrating them with the framework, we introduced two parallel systems.

First, a plugin system, mainly intended for beamline staff. BEC plugins can be new devices, scan routines, data output formats, and GUI widgets. They are managed through a bespoke plugin-creation application which provides the user with the required boilerplate. It is based on highly configurable code templating, allowing us to update the template without exposing users to the peculiarities of the underlying libraries or the difficulties of resolving merge conflicts.

Second, a system to run arbitrary user-supplied scripts, designed to limit the security issues this would normally entail. Scripts are executed in a container permissioned only to access the central message broker. The execution environment closely mimics the beamline iPython terminal environment, so that a procedure can be interactively developed, then submitted to the server.

## INTRODUCTION

The Swiss Light Source (SLS) at the Paul Scherrer Institute (PSI) is currently coming online following a major upgrade. Alongside the new accelerator and beamline hardware, a new centrally developed software framework has been introduced: Beamline Experiment Control (BEC). BEC [1] provides a common infrastructure for experiment control across the facility. This represents a shift from the historical model of beamlines developing their own control systems towards a shared, general-purpose framework inspired by software developed at other synchrotrons.

While BEC provides the central infrastructure, beamlines still require the flexibility to implement their own experiment-specific routines. There are two main reasons for this. First, the core BEC team comprises only four developers, responsible for 16 beamlines. Second, the beamline scientists are the true experts in their respective experimental techniques and must be able to define new procedures as needed.

This paper describes our strategy for maximizing the impact of the central team while providing beamline staff with the tools they need. Our approach balances the need for integration into the central framework with the flexibility required for cutting-edge experiments. Two particular examples are highlighted to showcase our attempt to provide a set of development and build tools which make it easier for scientists to develop their plugin repositories in a sustainable manner.

**BEC Architecture** BEC follows a modular service-oriented architecture [2] with communication mediated by a central Redis message broker. Key services include a device server which mediates instructions to hardware (using the ophyd [3] hardware abstraction layer), a scan server which manages a queue of measurements, and a file writing service, among others. Each of these services communicates with the others only through the central message broker.

There are two user interfaces:

- A widget-based graphical user interface (GUI) built with PySide6 (Qt), BEC Widgets [4].
- An iPython-based [5] console for terminal operation.

The device server and scan server publish their available components to the message broker, such that the interactive iPython session can be initialised with access to the various hardware devices of the beamline, and available scan routines. The components are then discoverable through "tab-completion" and the user can interactively explore them and inspect their methods and documentation.

Beamline-specific components are implemented on top of the central BEC infrastructure. These include custom macros, scan definitions, device integrations, and GUI widgets.

## STRATEGY FOR EXTENSION

Given limited central resources, our strategy is to empower beamline staff to extend BEC themselves in a sustainable and consistent manner. The guiding principles are:

- Beamline scientists should have the flexibility to write the experiments they require.
- Newly implemented scans should integrate seamlessly with the existing pipelines for data collection, file writing, and analysis.
- No domain-specific language (DSL) is introduced; extensions are written in Python.

* david.perl@psi.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEAG006

# BETTER SOFTWARE OBSERVABILITY USING TANGO CONTROLS WITH OPENTELEMETRY - EXPERIENCE AT MAX IV

A. F. Joubert*, L. Zhu, B. Bertrand, MAX IV, Lund, Sweden

*Abstract*

Distributed software systems are complex and the interactions across multiple machines can be difficult to debug and monitor. Log messages are not enough for observability. We need more information about the communication between applications, how each one is executing, and its internal state. In practice, applications can be made more observable using software frameworks such as OpenTelemetry. The Tango Controls framework has built-in support for OpenTelemetry in C++ and Python since version 10.0.0. We are using it operationally at the MAX IV synchrotron. We provide examples of the traces, trends, and other data available when running at scale on a beamline with hundreds of devices. We report on the compute and performance impact for client and server software applications, as well as practical issues. For the backend servers that ingest and query the telemetry data (running Grafana Tempo for traces and Grafana Loki for logs) we report on the compute resources required..

## INTRODUCTION

The distributed control systems used at large experimental physics facilities are complex with a large number of applications interacting between many different machines. The MAX IV synchrotron in Lund, Sweden, is no exception in this regard. Our Tango Controls [1] system includes nearly 400 machines hosting 26k Tango devices. This can be challenging to debug using just log messages.

We need to understand the communication between applications, how each one is executing, its internal state, and details about any errors, in order to get a picture of how the complete system is operating. The software term for this is *observability* [2]. The three types of data, or signals, that applications can emit to achieve this are *traces*, *metrics* and *logs* [3]. Recently, a fourth signal, *profiling*, has started to become available [4].

We will focus on traces. These allow us to see the causal relationship between remote calls across various applications in the control system, in other words, distributed traces. Traces show what is happening within each application as it processes the request, and how different applications interact.

Figure 1 shows an example of a distributed trace from a Tango Control system. Each coloured rectangle is a *span*, and these show which method in the application was executing and how long it took. Different colours indicate different actors in the call. The client application, the PyTango library [5], the `Leader` device and the `Follower` device. Each span includes many details such as filepath and line number, host name, thread ID, version information, etc.

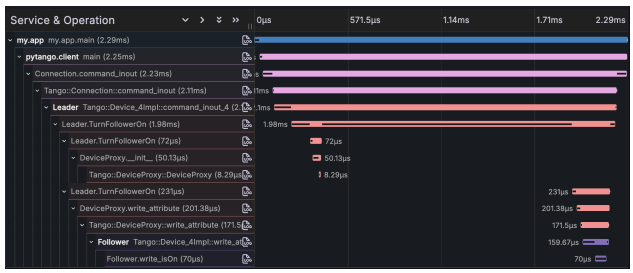



Figure 1: Example of a distributed trace across three applications (client script and two servers). The client script commands the `Leader` server to turn on a `Follower` server. The vertical axis shows a simplified call stack, and the horizontal axis shows the elapsed time since the first call.

### *OpenTelemetry*

OpenTelemetry [6] is an open-source and vendor-agnostic software framework that provides a practical way to realise observability. The project maintains specifications and conventions, and then implements this in many different programming languages. Their software libraries allow applications and 3rd-party libraries to be instrumented.

In order to instrument a software library, one has to add OpenTelemetry function calls at key places of interest. For example, starting a new span, with the necessary attributes, when an external call is received. Once a library is instrumented, applications that use the library can typically do so without any changes to their own code. When the application runs, the telemetry signals are automatically emitted.

### *Telemetry backends*

Given an instrumented application, the telemetry signals have to be sent somewhere. The immediate destination is a *collector* service, which then passes it on to a telemetry backend. The backend typically maintains a database for all the signals (traces, metrics, etc.), and provides a web interface to view the information. Some backends include their own collectors. The basic architecture is shown in Fig. 2.

There are many companies providing such backends [7–11]. We use Grafana [12], specifically their *Tempo* product for traces and metrics, and *Loki* for logs).

## IMPLEMENTATION IN TANGO CONTROLS

Tango Controls added support for OpenTelemetry starting in version 10.0.0. Traces and logs can be sent to a telemetry backend. Metrics and profiling have not been added yet. A high-level specification of features in Tango Controls is in the Request For Comment 19 document [13].

* anton.joubert@maxiv.lu.se

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBG001

# 20 YEARS (TRYING TO) SECURING CONTROLS

S. Lüders*
European Organization for Nuclear Research (CERN), Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Computer security is a marathon ran by some of us since decades. Trying to keep the malicious evil out while not inhibiting or strongly impacting accelerator operations and data taking. This talk shall review what worked and what less, and takes a view onto the upcoming challenges in maintaining a fair balance between operations and "security".

## THE ANCIENT PAST

In the ancient past, control systems were used to be stand-alone systems, usually decoupled from the larger enterprise networks, using their own network cabling infrastructures (based on a variety of proprietary buses and protocols), and using programming languages dedicated to control systems. But no more.

With the rise of "standard" IT technologies, interconnectivity and IT services made it also down to the plant floor. Early tests at CERN [1] and elsewhere, however, have shown that with this rapprochement control systems also inherit the genuine and inherent flaws, weaknesses and vulnerabilities of those IT technologies. Contrary, however, to IT, control systems cannot be fixed and patched as quickly as IT systems used to. Maintenance cycles need to be respected, certifications need to be reissued, controls loops need to be revalidated.

Still, more and more control system manufacturers integrated "IT" means into their products despite the fact that their expertise lies in "controls" and not "IT". Their brilliance centres in sophisticated (and technologically "beautiful") data acquisition hardware, extremely accurate measurement devices like oscilloscopes, precise timing systems, etc. "IT" is usually not their turf. Similarly, and without surprise, control system experts and operators centre their knowledge, skills and expertise on control systems. They are not necessarily network experts, web server managers, software programmers, or data base administrators. Still, both vendors/manufacturers and experts/operators embrace the benefits of IT without considering risks and mitigative measures.

The series of, so far, nine Control System Cyber-Security workshops [2] attached and prior to the ICALEPCS conference is supposed to raise awareness among the control system community, and bring them together with cyber-security experts, in order to holistically and in-detail improve the cyber-security of the control systems of accelerators and (large) experiments. While many accelerator labs and experiments started to adopt standard cyber-security protections, OT (operational technology) environments and their complex implementation and deployment requirements and timelines make it significantly more difficult to thoroughly adapt OT/control systems towards a more secure architecture and operation.

## NOTHING CHANGED?

20 years fast forward, the control system world has not much changed in that aspect of complexity, timelines and security. Actually, control systems have become even more complex and inter-interconnected, creating even more cross-dependencies and making the deployment of security measures even harder. In parallel, hardware manufacturers and vendors have now entirely embraced "IT".

Recent penetration tests by the CERN Computer Security of a plethora of control systems, embedded devices, and the "Internet of Things" (IoT) have repeatedly surfaced basic areas of security sub-optimalisms:

- Commercial CCTV cameras stored their default passwords, secrets and private certificates in an unprotected fashion within their hardware while secure-chip technologies, i.e. so-called Secure Access Modules or SAM, already exist. In other embedded IoT devices, changing default passwords even turned out to be impossible, or access control was entirely missing such that anyone with network access was able to take over full control;
- Intelligent, cloud-controlled lockers allow clients to open or close their lockers by scanning a QR code instead of using an RFID wristband, or coin & key. Unfortunately, that corresponding cloud service came with a vulnerable set-up and allowed anyone on the Internet to manipulate the open/close functionality of any subscribed locker world-wide. While the risk of theft is eventually controllable in this case, a similar vulnerability on defibrillators can be life-threatening. …and the review of cloud-connected washing machines serving the CERN hostels is in the pipeline;
- Custom-designed PCB (Printed Circuit Board) modules for fast feedback control connected to the CERN Intranet turned out to lack any level of robustness when receiving malformed or inappropriate network packages and, instead of just silently dropping those packages, completely broke down and stopped working. In some cases, restarting required the re-flashing of the local firmware;
- The authentication mechanism of a web-based SCADA (Supervisory Control and Data Acquisition) system has been found flawed in such a way that it was allowing to bypass the authentication process as a whole;
- Furthermore, shared passwords are used across vast access control system infrastructures and transmitted unprotected and in clear-text instead of using encryption as well as per-device credentials or certificates;
- In other cases, expert and operator passwords were accidentally exposed on public websites. Unfortunately, systematically changing those passwords turned out to

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBG003

# CENTRALIZED EPICS CHANNEL ACCESS FOR VDI USERS AT NSLS-II VIA CA GATEWAY ARCHITECTURE

A. Derbenev†, P. Shafer, S. Wilkins
Brookhaven National Laboratory, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

At National Synchrotron Light Source II (NSLS-II), Experimental Physics and Industrial Control System (EPICS) servers for the accelerator and beamlines reside on dedicated Virtual Local Area Networks (VLANs) isolated for security and network bandwidth. Since clients must run applications within respective networks, this poses a challenge for enabling centralized observability and control for staff with various roles. We have created a portal to access EPICS process variables (PVs) across the facility, using Virtual Desktop Infrastructure (VDI) and a dual Channel Access Gateway (CAGW) architecture on a dedicated "EPICS VDI" network. For each beamline, and for the accelerator, we deploy a pair of Channel Access Gateways: a VDI-side gateway that listens for client connections on the EPICS VDI network, and a Controls-side gateway that communicates with the Input–Output Controllers (IOCs) on the Controls VLAN. The controls-side gateway bridges the isolated "Controls" network and the routable "Services" network.

CAGW security enforces PVs as read-only by default, with Active Directory group membership granting beamline-specific write access. Any EPICS Channel Access (CA)-based client can run in the VDI environment, including Control System Studio Phoebus – the primary tool enabling staff to interact with PVs across the facility from a single session. PV access via VDI removes the need to run client software in the Controls environment, reducing system exposure and improving architectural separation. CAGW deployment is automated by Ansible using templated generation of network settings, PV lists, and access rules. This approach builds on a proven accelerator-beamline communication model and has shown stable performance.

## INTRODUCTION

Segmented network topologies have become a standard practice at NSLS-II facility over the last few years. They help protect IOCs and hardware from unrelated traffic and help constrain broadcast-oriented protocols such as EPICS Channel Access (CA) [1]. CA's discovery and connection model has proven effective when properly configured, but it is not well-suited for cross-VLAN reachability and provides only weak host/user access control, which historically pushed users to run clients "inside" trusted Controls networks. The EPICS Channel Access Gateway (CAGW) can address these issues: it aggregates client connections to the same PV, can bridge subnets, and adds a second layer of access security in front of IOCs while preserving the standard CA protocol operation supported by existing EPICS client and server applications [2].

At NSLS-II, several factors motivated the deployment of a CAGW-based solution of centralized PV access. A legacy remote-desktop tool was deprecated; there was a need for a supported, centrally managed user environment accessible to beamlines; and engineers and technicians increasingly needed a single session that could span accelerator and multiple beamlines to monitor machine parameters. We responded by pairing an on-premises Horizon VDI solution with a dual-gateway design so that any CA-capable tool (such as widely used CS-Studio Phoebus) can run in a VDI session and reach the PVs across the facility without the need for users to directly access the Controls VLANs. Phoebus' trajectory as a modern EPICS client framework in the facility reinforced this approach.

## RELATED WORK AND CONTEXT

The CAGW is both a CA client and a CA server; EPICS clients discover PVs from and connect to the gateway as if it were an IOC, and the gateway multiplexes those requests to the actual IOCs, reducing connection load while also providing independent configuration capabilities for PVs being served, and access modes for these PVs. At NSLS-II, CAGWs have been routinely used to enable EPICS communication between beamlines and the accelerator, for monitoring and limited control across corresponding control systems.

The usage of CA in segmented networks requires suitable scoping of PV discovery (by appropriately setting EPICS_CA_ADDR_LIST), configuring EPICS servers to bind to a proper network interface, and the need to allow UDP/TCP communication on CA ports through firewalls. NSLS-II EPICS servers, clients, and hosts are configured to meet these requirements.

While the present work focuses on enabling CA protocol communication, EPICS 7 introduced pvAccess (PVA) for structured data and high-throughput streams [3]. PVA is increasingly used for images and arrays at NSLS-II, and existing community implementations (e.g., PVA Gateway) offer capabilities analogous to CAGW. Adding PVA communication is therefore a natural extension of the centralized access setup.

## FACILITY DRIVERS AND REQUIREMENTS

The immediate operational catalyst was the deprecation of a legacy remote-access solution. VDI provides a resource-managed, patchable, and uniform desktop image that can be administered centrally – an advantage over scattered "access machines." Equally important, we wanted to

† aderbenev@bnl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBG004

# TIPPSS FOR NAVIGATING A CHANGING CYBERSECURITY LANDSCAPE AT THE ELECTRON-ION COLLIDER AND OTHER SCIENTIFIC RESEARCH FACILITIES*

L. K. Nguyen†,1, F. D. Hudson[2], J. Jamilkowski[1], K. Kulmatycski[1]
[1]Brookhaven National Laboratory, Upton, NY, USA
[2]Columbia University, New York, NY, USA

*Abstract*

The Electron-Ion Collider (EIC) aims to unlock the secrets of the strong nuclear force and revolutionize our understanding of the fundamental structure of visible matter. It is being built at Brookhaven National Laboratory (BNL) and could possibly be the only large collider built in the world in the next 20-30 years, during the "Age of AI". This creates the very unique opportunity for a complete AI/ML lifecycle of a large-scale state-of-the-art scientific research facility, but also many challenges, as this lifecycle overlaps with a rapidly changing cybersecurity landscape. Standards, regulations, and guidance are likely to be released (and then possibly revised) at the same time that design, construction, and then finally operations of the EIC must proceed. We present the use of the new Trust, Identity, Privacy, Protection, Safety, and Security (TIPPSS) framework from the IEEE/UL 2933 TIPPSS standard as a framework for scientific research facilities. This will enable us to design and build a safe and secure infrastructure, and robust trust and identity architecture, to protect the scientific instrument ecosystem as we enable "AI readiness" and AI/ML deployment (especially at scale) in the face of increasing cybersecurity challenges, using the EIC as a case study.

## INTRODUCTION

The next decade will see the materialization of two very exciting scientific and technological endeavors. The first is the facility known as the Electron-Ion Collider (EIC), being built at Brookhaven National Laboratory (BNL) in New York State, in partnership with the Thomas Jefferson National Accelerator Facility (JLab) and other domestic and international partners. Construction of the EIC is scheduled to begin in 2026, following the end of the Relativistic Heavy Ion Collider (RHIC) program, and will conclude around 2031. See Fig. 1 for a schematic layout of the future facility, which will reuse the hadron injector chain (e.g., Booster and AGS) and 3.8-kilometer RHIC tunnel but also include many new buildings and machines, such as the Rapid Cycling Synchrotron (RCS). The second is the widespread use of machine learning (ML) and artificial intelligence (AI) in research and discovery. In particular, these approaches will increasingly move out of purely data- and simulation-based contributions and into real-time, bidirectional interactions with systems as part of the overall experiment. As AI/ML becomes more sophisticated, large-scale "human-in-the-loop"—if not "human-ON-the-loop"—scientific research facilities are becoming more realizable.

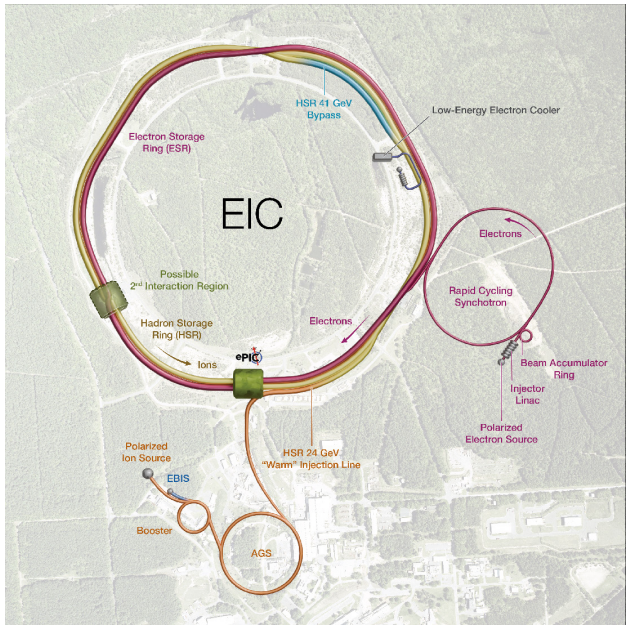



Figure 1: The future Electron-Ion Collider (EIC) at Brookhaven National Laboratory (BNL).

In its most recently released Long Range Plan—used to advise both the Department of Energy (DOE) and National Science Foundation (NSF) on research priorities in nuclear science over the next decade—the U.S. Nuclear Science Advisory Committee (NSAC) explicitly identified "the expeditious completion of the EIC as the highest priority for facility construction" and, in a separate identification of strategic opportunities, called for "a coordinated effort to integrate AI/ML technologies into the nuclear physics research programs [to] accelerate discoveries" [1]. The convergence of these two priorities in the form of an "AI-ready" EIC is thus a very natural step—and as possibly the only large collider built in the world in the next 20-30 years, the EIC could uniquely explore and exploit the possibilities of AI/ML for science with AI readiness deliberately incorporated into its design. Moreover, given that it is still in its design phase, the EIC uniquely presents an entire lifecycle for AI/ML at a large-scale scientific research facility.

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy. TIPPSS efforts have been supported by the IEEE Standards Association, U.S. National Science Foundation (Awards #1623931, #1550284, #1748395, #1916585), Columbia University, and FDHint, LLC.
† lnguyen@bnl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBG005

# CYBER SECURE EXPERIMENTAL PHYSICS AND INDUSTRIAL CONTROL SYSTEM

G. S. McIntyre[1,2,3], G. White[2], K. Kasemir[4], M. Davidsaver[3], B. Dalesio[3], Dr. I. Finch[5], E. Williams[2], J. Einstein-Curtis[6]
[1]Level N Ltd., Lydney, UK
[2]SLAC National Accelerator Laboratory, Menlo Park, CA, USA
[3]Osprey DCS LLC, Ocean City, MD, USA
[4]Oak Ridge National Laboratory, Oak Ridge, TN, USA
[5]Science and Technology Facilities Council, Swindon, UK
[6]RadiaSoft, Boulder, CO, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

**Experimental Physics and Industrial Control System (EPICS)** version 7 introduced **PVAccess (PVA)** as a network protocol supporting request/response and publish/subscribe for process variables. **Secure PVAccess (SPVA)** is an enhanced PVA protocol that adds **Transport Layer Security (TLS)** encryption and X.509 certificate authentication while preserving PVA's core operations (GET, PUT, MONITOR, RPC).

SPVA uses **TLS 1.3** to secure communications end-to-end, ensuring confidentiality and integrity of control system data. Peers (clients and servers) authenticate each other via X.509 digital certificates, enabling **mutual authentication** or at least server-side authentication in every connection. A new **Certificate Management Service (PVACMS)** handles certificate provisioning, status monitoring (encapsulating the Online Certificate Status Protocol, **OCSP** payloads within SPVA), and lifecycle operations (issuance, renewal, revocation).

The SPVA implementation remains **backward-compatible** with legacy (unencrypted) PVAccess so that secure and non-secure clients/servers interoperate seamlessly. It also integrates with site authentication systems (e.g. Kerberos and LDAP) to leverage existing user credentials.

This report presents the SPVA design and implementation in EPICS Base, **PVXS** (the C++ PVA implementation), and the Java PVA client (as used in CS-Studio/Phoebus), with detailed technical verification of its TLS handshake roles, certificate status mechanism, certificate management (including PKCS#12 keystore usage), and integration with EPICS Access Security and external authentication infrastructure.

## SECURE PVACCESS

### *Transport Layer Security*

**Secure PVAccess (SPVA)** modifies the PVA network transport to run over **TLS 1.3** [1] connections instead of plain TCP sockets. All PVA protocol messages are encapsulated inside **TLS** records for encryption and integrity protection. In effect, SPVA adds a security layer beneath the PVA message layer, without altering the high-level PVA semantics. Figure 1 illustrates how PVAccess messages are wrapped in TLS (which itself runs over TCP/IP) to provide confidentiality and authentication:

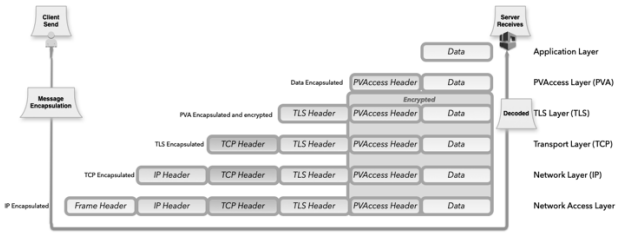



Figure 1: PVAccess message encapsulated within a TLS message (ensuring encryption and integrity) before transmission over TCP/IP.

In an SPVA session, the EPICS **PVA client** acts as the TLS **client** and the EPICS **PVA server** acts as the TLS **server** during the handshake. The TLS 1.3 handshake is always initiated by the client, which sends a *ClientHello* proposing cryptographic parameters. The server responds with a *ServerHello* (selecting the cipher suite) and presents its **X.509 certificate** [2] to the client for authentication. The server uses its private key (corresponding to the certificate) in the handshake to help derive a shared symmetric session key. In SPVA, the server's certificate is typically an RSA or ECDSA certificate issued by the facility's Certificate Authority (the **Trust Anchor**), and the client verifies this certificate against its copy of the Trust Anchor's certificate. If the server's certificate is valid and chains to the trusted root, the client will trust the server's identity.

Optionally, **mutual TLS** can be enabled: the server can request a client certificate, and the EPICS client will present its own X.509 certificate for the server to verify. This allows **peer authentication** in both directions (both sides providing proof of identity via certificates), or **server-only authentication** if the client has no certificate. TLS 1.3 provides strong forward security via ephemeral key exchange and has removed legacy features like handshake renegotiation (which is a security risk). SPVA connections therefore do **not support renegotiation** – if a certificate is revoked or expires, the secure connection is simply terminated and must be re-established with a new certificate.

A critical requirement for trust in SPVA is that both client and server possess the same **Trust Anchor** certificate. The Trust Anchor is the root certificate authority that signed all legitimate client and server certificates. As long

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBR001

# FAST EVENT SYSTEM FOR THE ADVANCED PHOTON SOURCE

R. Hong*, A. Pietryla, G. Shen, A. Johnson, S. Wang, M. Smith, R. Diviero, S. Farrell, R. Koldenhoven, Argonne National Laboratory, Lemont, IL, United States

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Fast Event System, a global time base and event-based trigger distribution system, has been developed and commissioned for the Advanced Photon Source Upgrade (APS-U) and the linear accelerator (LINAC) refurbishment projects. The hardware components developed by Miro-research Finland (MRF), including event masters (EVMs), event receivers (EVRs), and event fan-out modules, are installed in VME Input/Output Controllers (IOCs) to generate and distribute timing signals to client devices. The firmware of the MRF EVM and EVR has been updated by the manufacturer to meet the needs of the Fast Event System at APS. The software for these IOCs has been developed using the EPICS framework. In this presentation, the overall structure, key functions, and recently developed features of the Fast Event System are introduced. Particularly, the performances of the delay compensation and event synchronization are discussed. The experiences and lessons learned during the commission and operation will be discussed as well. .

## INTRODUCTION

The Advanced Photon Source (APS) at Argonne National Laboratory has become a fourth-generation synchrotron light source following its major upgrade, the Advanced Photon Source Upgrade [1] (APS-U). A new electron storage ring was constructed to dramatically increase X-ray brightness and coherence. To meet the demanding requirements of the APS-U project, a new event-based timing system, the Fast Event System, has been designed and commissioned. The system generates and distributes synchronous clock and trigger signals to devices such as beam position monitors (BPMs), the data acquisition (DAQ) system, and the machine protection system (MPS). The fast event system also provides a time base so that data from different systems can be time-correlated in online or offline studies.

Besides the APS-U, there are other refurbishment efforts at the APS to modernize the injector complex. A similar fast event system was deployed for the Linear Accelerator [2] (Linac) before the APS-U storage-ring construction. It replaced the 30-year-old Linac timing system, which provided triggers for the radio-frequency (RF) sources, beam kickers, and beam diagnostic devices.

In this paper, the basic architecture of an event system, the hardware components used at the APS, and the software development are described. New features of the event systems that have not been widely implemented at other light sources are highlighted. The design and commissioning experiences of both event systems at the APS are presented in detail.

* rhong@anl.gov

## EVENT SYSTEM ARCHITECTURE AND HARDWARE COMPONENTS

A typical event system consists of event generators (EVGs), event receivers (EVRs), and fan-out modules. An EVG creates a bit stream of event codes at a programmed frequency or upon an external trigger. This bit stream, typically sent over an optical link (the so-called event link), can be fanned out to multiple destinations by a fan-out module, eventually reaching the EVRs. Upon receiving an event code, an EVR executes a programmed function, such as generating an output pulse, digitizing a waveform, or creating an interrupt in the host system. An event system usually has a tree-structured topology, with a top-level EVG as the root and end-point EVRs as the leaves. The generation and reception of an event are synchronized with a common reference clock. This clock is typically derived from the accelerator's RF source and fed to the top-level EVG, while EVRs can recover the reference clock from the input event link.

The state-of-the-art 300-series event-system components from Micro-Research Finland [3] (MRF) are used in many accelerator facilities. The European Spallation Source [4] (ESS) used the 300-series μTCA event masters (EVMs) and EVRs. The fast event systems at APS are constructed using 300-series VME event master VME-EVM-300 (Fig. 1), event receiver VME-EVR-300 (Fig. 2), together with other supporting modules like the 1-to-11 fan-out module VME-FOUT-12 and transition module VME-UNIV-TB.

The core function [5] of the EVM module is an EVG, which can generate events by a clock divider, an external trigger, or a sequencer with programmable time-stamped event codes. The EVM can take external LVTTL triggers from its 3 input ports in the front panel, or additional 16 input ports from a transition board populated with LVTTL input daughter boards. In each clock cycle, a 16-bit word is transmitted, with the lower 8 bits used for the event code and the upper 8 bits for a distributed bus (DBUS) or generic data transmission as shown in Fig. 3. The DBUS and transmitted data alternate in each clock cycle. The DBUS bits can be set to follow a clock divider or the logic state of an input channel. The generic data transmission function can be used to distribute customized data, such as beam status. The EVM also uses internal data transmission to pass the measured delay value to a downstream module, a crucial step in the delay compensation mechanism discussed later in the paper. An EVM fans out the event stream to eight downstream event links. It also has an upstream event link, allowing multiple EVMs to be daisy-chained to generate more events, and the event clock from the upstream link can be recovered and used as the reference clock. The MRF EVR module receives data from its upstream event link and recovers the event clock.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBR002

# INTER-GENERATIONAL COMPATIBILITY STUDY OF MRF EVENT TIMING MODULES

D. Wang*, M. Satoh, K. Furukawa
KEK, Ibaraki, Japan

*Abstract*

Event timing systems are critical for the synchronization of beam diagnostics and accelerator control at KEK LINAC. Such systems have historically relied on VME-based modules since 2008, such as the MRF 230 series event generator and event receiver. However, with some VME modules approaching their market end-of-life, transitioning to modern platforms like MicroTCA is becoming imperative. This work addresses the challenges of such a migration by focusing on a phased upgrade approach, where new MicroTCA-based MRF 300-series timing modules will coexist with and eventually replace the legacy VME timing modules. A primary concern during this transition is ensuring the functional compatibility between the old and new generations. This paper presents a comprehensive evaluation of critical timing function across coexisting VME and MicroTCA systems. Core compatibility aspects evaluated include event code transmission and reception accuracy, timing precision and jitter, trigger output characteristics, event rate handling, sequencer operation, and timestamp consistency. The findings aim to provide a quantitative assessment of compatibility, identify potential limitations, and offer practical insights for other facilities planning a similar upgrade of their MRF event timing systems, thereby minimizing risk during the accelerator operation.

## INTRODUCTION

The KEK LINAC relies on a distributed event-based timing system, which has historically been implemented using VME form factor modules from Micro-Research Finland (MRF) which has demonstrated high stability and reliability in routine operations. However, the VME standard is now approaching obsolescence. In particular, essential components such as the MVME5500 and MVME6100 CPU modules have reached end-of-life and are no longer supported by their manufacturers. At the same time, the MRF 230-series timing modules that form the backbone of the system are being phased out, replaced by the new MRF 300 series available in multiple form factors, including MicroTCA which demands an VME to MicroTCA upgrade.

During the upgrade, legacy and next-generation modules must coexist within the same network, and full functional compatibility must be guaranteed to safeguard accelerator operations. At the same time, the new MicroTCA modules introduce advanced features such as delay compensation, which require careful evaluation in the context of existing operational patterns.

* di.wang@kek.jp

This paper presents a phased migration strategy for replacing the legacy VME modules with MicroTCA-based MRF 300-series hardware. Particular emphasis is placed on verifying cross-generation compatibility between EVG and EVR modules, assessing performance metrics such as jitter and event transmission accuracy, and developing automated test utilities to streamline large-scale validation. The findings provide both a roadmap for transition and practical insights for other facilities undertaking similar upgrades of event timing systems.

## VME TO MicroTCA TRANSITION

### *Overview of the Current Timing System*

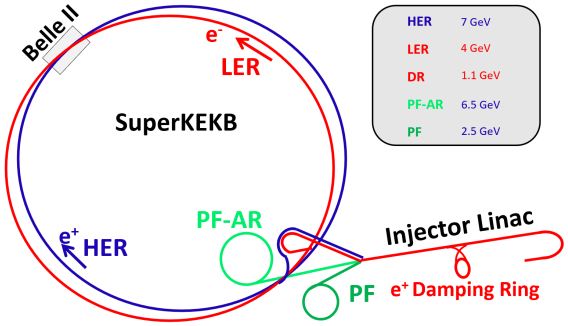



Figure 1: Overview of the KEK LINAC.

As Fig. 1 shows, the LINAC simultaneously provides beam to four storage rings at a repetition rate of 50 Hz. Besides two light sources, PF and PF-AR, the LINAC accelerates electrons to 7 GeV for injection into the SuperKEKB high energy ring (HER). For positrons, the LINAC first injects them into a damping ring (DR) located centrally within the LINAC. After a damping time ranging from 40 ms to 200 ms, the positrons are extracted and continue to be accelerated within the LINAC to 4 GeV before injection into the SuperKEKB low energy ring (LER) [1]. This operational requirement mandates that the timing system delivers stable and reliable synchronization, along with flexible Pulse-to-Pulse Modulation (PPM) operation, to switch the outputs of various accelerator devices every 20 ms [2].

Two operational requirements further constrain the timing system design. First, the trigger sequence must remain phase-locked to the 50 Hz AC mains to stabilize the output from high power kicker magnet power supply used for PF injection. Second, precise RF bucket selection is required across the injector-to-ring chain. While the bucket selection cycle (BSC) for SuperKEKB HER is short (493 µs) and easily accommodated at 50 Hz, inclusion of the positron damping ring enlarges the combined DR–LER BSC to 11.34 ms, complicating AC-phase synchronization [3]. The timing system

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBR003

# WHITE RABBIT TIMING: THE NEW CERN ACCELERATOR TIMING SYSTEM

G. Moscardi*, M. Cejp, A. Dujović, T. Gingold, E. Gousiou,
F. W. Hoguin, I. Kozsar, G. Kruk, A. Zeising,
European Organization for Nuclear Research (CERN), Geneva, Switzerland

*Abstract*

After more than 30 years of service, CERN's accelerator timing system is being renovated, moving from the existing distribution infrastructure based on the RS-485 technology and legacy hardware modules, to a new one based on White Rabbit.

Developed at CERN, White Rabbit Timing (WRT) is a generic toolkit composed of the White Rabbit Event Node (WREN) - a System-on-Chip based hardware module, and the corresponding software stack. WRT allows transmission and reception of messages, along with an arbitrary payload (key-value pairs). The received messages enable the generation of triggers in the form of software interrupts and electrical pulses, with sophisticated and highly configurable triggering patterns. WRT seamlessly integrates time derived from the radio frequency used for particle acceleration, with WRENs capable of locally generating beam orbit and bunch clocks, as well as broadcasting Beam Synchronous Timing (BST) streams over dedicated optical links.

We present the key concepts of WRT, its architecture, multi-layered distribution network layout, functionalities and usage at CERN. We also draw a potential path towards a turn-key timing system based on WRT that could be deployed anywhere for scientific or commercial applications.

## THE CURRENT CERN TIMING SYSTEM

A particle accelerator is composed of many pieces of equipment which must work together to produce, accelerate, maintain and collide particle beams. It is necessary to synchronize the operation of devices such as magnets, beam instrumentation or radio-frequency (RF) cavities.

At CERN, synchronization is achieved by sending messages on dedicated timing networks. These messages are used to generate electrical pulses triggering accelerator equipment as well as to instruct the corresponding software to perform appropriate actions depending on the beam parameters. To cover the wide range of required timing resolution (from milliseconds down to about 25 ns), multiple timing systems are currently in operation (Fig. 1).

### *General Machine Timing*

General Machine Timing (GMT) distributes so-called *Central Timing Events* (CTIMs), such as *Cycle Start*, *Beam Injection*, *Cycle End*, etc. These events are received by Front-End Computers (FECs) equipped with a Central Timing Receiver (CTR) module [1].

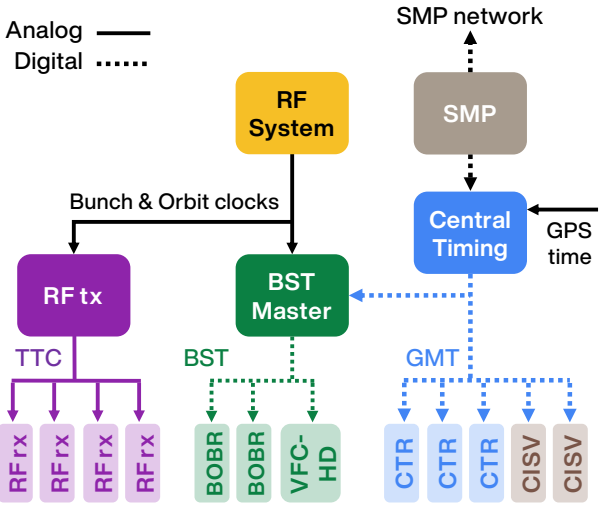



Figure 1: Overview of the current timing systems at CERN.

CTRs are used to decode GMT frames and to generate *Local Timing Triggers* (LTIMs) relative to the Central Timing Events in the form of pulses or software interrupts, that can be used to precisely control software running on a FEC or hardware devices connected to it.

At the physical level, GMT is based on a RS-485 serial link running at 500 kbit/s on a copper cable [1], with one transmitter, hundreds of receivers, as well as signal amplifiers and converters.

### *Radio-Frequency Clocks*

Many equipment types require synchronization with Radio-Frequency (RF) signals so as to be synchronous with the circulating beam. These signals consist of the *Bunch Clock* and the *Turn/Orbit Clock* (also known as $F_{rev}$) and, in the case of the LHC, they are distributed as analog signals over a dedicated optical network.

### *Beam-Synchronous Timing*

The Beam-Synchronous Timing (BST) system is CERN-specific and solely dedicated to the SPS and LHC accelerators. Its purpose is to broadcast timing signals that are required to synchronize the different beam instrumentation systems with the circulating beam. For this purpose, messages are encoded using the Bunch Clock provided by the RF system. Additionally, the BST system rebroadcasts selected GMT events and data, for the benefit of those receivers that are not simultaneously connected to the GMT network.

* giorgio.moscardi@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBR004

# TCLK MUST STAY! CAMAC MUST GO! HOW DOES FERMILAB MOVE FORWARD?*

M. R. Austin†, L. Carmichael, D. McArthur, E. Milton, A. Quilty
Fermi National Accelerator Laboratory, Batavia, IL, USA

*Abstract*

The current Timing System at Fermilab has been around for 40 years and currently relies on 7 CAMAC crates and over 100 CAMAC cards to produce the Tevatron Clock (TCLK). Thanks to the ingenuity of those before us, this has allowed Fermilab the flexibility to change the timing and Events for its accelerator as beamlines and projects have changed over the years. With the advent of the Proton Improvement Plan-II (PIP-II), the Timing System at Fermilab is being reimagined into a single chassis with even greater flexibility and functionality for decades to come while tackling the ever-challenging task of maintaining backwards compatibility.

## INTRODUCTION

TCLK at Fermilab is used to transmit important accelerator timing information to all major systems throughout the accelerator complex with up to 256 unique triggered 8-bit "Events" that are used as timing markers. TCLK is a 10 MHz serial signal which uses an encoding scheme called bi-phase or modified Manchester coding to transmit 10 bits (1 start bit, 8 data bits, and 1 parity bit). There are then two additional 1's transmitted before the next Event is allowed to be transmitted [1]. TCLK begins with a 10 MHz source and the Master Clock Oscillator (MCO). The MCO causes triggers synchronous to the 60 Hz line frequency including the trigger for the Timeline Generator (TLG) to play out certain Events that are generally related to Machine Resets. Some of these Machine Resets then trigger other hardware that triggers the playing out of other Events. The timing of all these Events is based on programmable delays, external Event requests from remote hardware, as well as delays that synchronize Event requests to the 60 Hz line frequency. To accomplish all of this, there are 7 Computer Automated Measurement and Control (CAMAC) crates and over 100 cards with settable devices for about half the cards and hard coded settings for the other half, a few custom chassis, and a VME Front End for the TLG interface. The hardware involved with creating TCLK is shown in the collage of CAMAC Crates in Fig. 1.

## MOTIVATION

PIP-II necessitates the first major modification to Fermilab's timing system since the Tevatron in the 1980s. The new system must provide inter-bunch precision and provide full backwards compatibility with the rest of the complex. This means that the backwards compatible TCLK output from the Accelerator Clock (ACLK) Generator that will interface with older hardware, needs to look and act exactly like the existing TCLK signal. The new system must meet numerous specifications for user interface, Permit inputs, Event triggers, and enabling, disabling, and Timestamping of Events. This results in two systems being necessary: The ACLK Generator, to provide general machine timing and the Linac Clock (LCLK) Generator to provide beam synchronous timing to PIP-II.

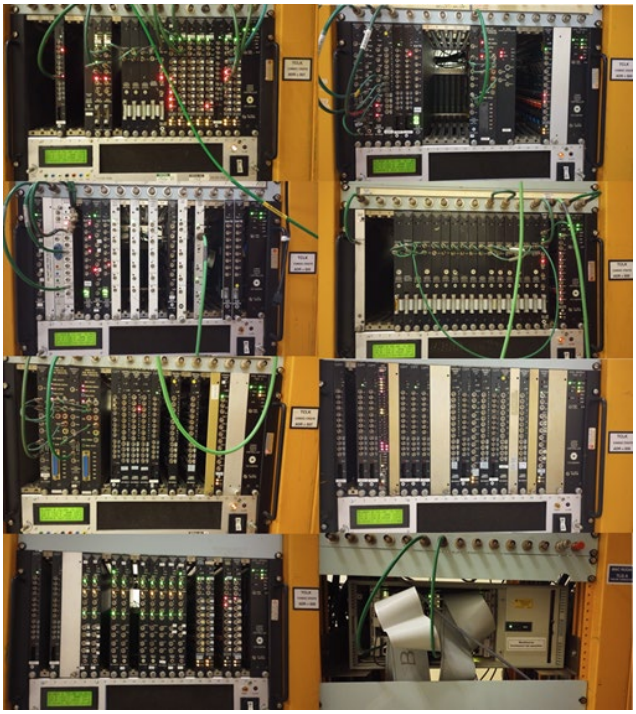


Figure 1: CAMAC hardware for creating TCLK.

## SPECIFICATIONS

Given the similarity between the needs of ACLK and LCLK, it was determined that a system could be designed to share common physical hardware and logic implementation. This system is colloquially known as a Clocking System. This common hardware is shown in Fig. 2. A significant portion of the logic on the inside is similar as well. The Clocking System consists of:

- A System-on-Module (SOM) with an Arria 10 Field Programmable Gate Array (FPGA) containing a Hard Processor System (HPS) to handle all network communication for the user interface. This includes setting devices for various Event triggers and delays, loading new Timelines, and Machine Data (MDAT) settings [2].

* This work was produced by FermiForward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.
† maustin@fnal.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEBR006

# DISTRIBUTED I/O TIER: FROM CONCEPT TO OPERATIONAL READINESS - A MODULAR PLATFORM FOR CUSTOM RADIATION-TOLERANT AND RADIATION-FREE ELECTRONICS

G. Daniluk*, V. Amoiridis, A. Arias Vazquez, P. Peronnard, A. Pinho, B. Guncic, M. Lettrich
CERN, Geneva, Switzerland

*Abstract*

The Distributed I/O Tier (DI/OT) project was launched to develop a common, modular hardware platform for custom electronics at the lowest layer of the CERN control system. Traditionally, this layer—closest to the accelerator and often exposed to radiation—relied on highly specialized, custom-designed devices with little reusability across subsystems. DI/OT addresses this limitation with a 3U crate based on the CompactPCI Serial industrial standard, along with high-performance (non-radiation-tolerant) and radiation-tolerant (lower-performance) modules. Key components include two System Boards (featuring Igloo2 FPGA and AMD Zynq UltraScale+ MPSoC), a radiation-tolerant switched-mode power supply, a fan tray, an FMC WorldFIP module, and a radiation-tolerant RISC-V soft-core. DI/OT users can tailor the platform to their needs by designing application-specific Peripheral Boards, FPGA configurations, and low-level software for the System Board. Standardizing on a common platform enables different equipment groups to benefit from centrally supported hardware while facilitating the sharing of application-specific peripherals. Over the past few years, DI/OT has evolved from a prototype to a production-ready platform. This paper presents DI/OT's hardware modules, radiation qualification results, initial pilot deployments at CERN, and its adoption by the quantum computing community.

## INTRODUCTION

The CERN control system comprises multiple layers of hardware and software, ranging from devices installed close to the accelerator to supervisory applications running on operator workstations for control and monitoring. The lowest hardware tier is typically built from highly specialized, custom-designed electronics. Conventional industrial platforms—such as Industrial PCs, Programmable Logic Controllers (PLCs), VME, or MTCA crates—are generally not used at this level because of radiation exposure or the need for hard real-time processing of large numbers of I/O channels. While the upper layers of the control system benefit from industrial solutions and standardization, including modular ecosystems for custom hardware [1], the tier closest to the accelerator has traditionally relied on fully custom electronic modules tailored to specific instrumentation tasks (e.g., interlocks, position measurement, and power converter control). As a consequence, reusability across subsystems has been limited.

* grzegorz.daniluk@cern.ch

The Distributed I/O Tier (DI/OT) project was launched within the High Luminosity Upgrade project for the Large Hadron Collider (HL-LHC) Work Package 18 [2] to conceive a generic, modular, and reusable platform for custom electronics, offered in two variants: radiation-tolerant (lower performance) and high-performance (non-radiation-tolerant). Its main objectives were to: (1) build on industrial standards; (2) provide robust inter-module communication; (3) offer common monitoring and configuration/discovery services; (4) leave broad freedom for customization by end users; and (5) develop the platform as open hardware and open software in collaboration with industry.

The project has advanced substantially in recent years through hardware design, prototype validation, FPGA development, low-level software, and radiation qualification. Prototype platforms have been delivered to users to enable customization and adaptation to final applications. DI/OT has also found use cases beyond CERN and is now entering series production and first pilot deployments.

In this paper, we first describe the components of the DI/OT hardware kit and the results of radiation qualification. We then describe the reuse model for end users and conclude with the initial pilot deployments and ongoing application developments at CERN and outside.

## DI/OT HARDWARE KIT

The goal of the DI/OT project is to extract the common needs of various users who traditionally designed custom electronics for their specific applications, and to address them by providing a set of centrally supported hardware modules (Fig. 1). The common DI/OT hardware kit is based on a 3U Eurocard chassis compatible with the industrial CompactPCI-Serial (CPCI-S) standard [3]. The standard defines a 9-slot backplane in which the left-most System Slot hosts the crate controller (System Board), while the remaining eight Peripheral Slots host application-specific Peripheral Boards. Unlike other popular crate standards (VME, PXIe, MTCA), CPCI-S specifies a fully passive backplane with a compact, multi-pin Amphenol AirMax connector and a star topology of LVDS differential lanes from the System Slot to each Peripheral Slot. These features make it well suited to DI/OT use cases (radiation tolerance and freedom to choose inter-board communication).

The modular DI/OT hardware kit is provided in two variants: radiation-tolerant and non-radiation-tolerant. The former offers lower compute performance (features a Flash-based FPGA) but uses active components screened for radiation resilience and incorporates mitigation techniques in the

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WECG001

# BEHAVIOR TREE SEQUENCING AND AUTOMATION FRAMEWORK AT LCLS*

R. Tang-Kong†, J. Cohen, Z. Lentz
SLAC National Accelerator Laboratory, Menlo Park CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Linac Coherent Light Source (LCLS) has a distributed array of control systems operating in concert to deliver bright, coherent x-ray lasers to a variety of experiment endstations. The automation systems at these endstations have different goals and maturity levels, which has resulted in a diverse set of sequencing and automation needs. To optimize uptime and flexibility, we are leveraging behavior trees as an automation framework. Proven in industry, behavior trees provide a comprehensive, shared, no-code language that invested parties can communicate and iterate upon. This system formalizes operator experience in a formalized, version-controlled document and gives system owners a structured way to induce state transformation or react to off nominal events.

## INTRODUCTION

The Linac Coherent Light Source (LCLS) is home to nine specialized instrument hutches, each performing unique experiments utilizing a world class X-ray free-electron laser. In such a cutting-edge facility, enabling users and scientists to work efficiently is of paramount importance. Automation is a pivotal aspect of not only enabling operations but also ensuring safety and consistency in the controls system. Automating beamline operations at LCLS has historically been a difficult task due to the wide variety of sequencing and automation needs across the instrument hutches. Historically, automation solutions have been developed organically as processes were iteratively defined. While expedient, these solutions have been as varied as the hutches they originated from, and maintaining these solutions has proven challenging. Smaller scale automation tools have been successfully developed and deployed, but in order to support automation at the subsystem level a more general framework is necessary.

Reactivity is critical for automation at a dynamic facility like LCLS, with changing instrument configurations and beam conditions. This is compounded by the largely colinear configuration of the instrument hutches, where downstream experiments can be impacted by unexpected upstream activities. While the configuration of each hutch is different, they share common components (motors, cameras) that can and should be automated in similar ways. Developing automation routines in a modular fashion is necessary to ensure consistency and leverage work across the lab. Approachability is also critical to enable the users and scientists to interact and understand these automation tools, regardless of their expertise. Thus, for any automation tool to be successful at LCLS, it must be responsive, modular, and approachable.

In order to support these automation needs, LCLS has been developing a high-level sequencing framework based on Behavior Trees (BTs). BTs are a robust sequencing framework that also serves as a shared no-code language that users can communicate with. In this paper we will briefly present Behavior Trees as a sequencing framework and describe how it fulfils our main automation needs (Responsive, Modular, Approachable). A description of the software architecture and design will be provided, along with a description of how Behavior Trees fit in the wider controls system technology stack at LCLS. Finally, the plans for future development will be discussed.

## BEHAVIOR TREES AS AN AUTOMATION FRAMEWORK

Originally developed in the computer game industry as a tool for programming Non-Player Characters, Behavior Trees are a hierarchical model for plan execution. This framework excels at enabling agents to react to dynamic environments [1]. Since then, BTs have received significant attention in academia and robotics for their modularity and reactivity. BTs have been used in applications like robotic arms [2], UAVs [3], brain surgery robots [4], and more [5].

A BT is built from a small set of simple components, allowing practitioners to easily understand and encode complex structures. A BT is a directed rooted tree, where internal nodes are “control flow nodes” and leaf nodes are “execution/action nodes”. Each node can only have one parent, except for the root node which has no parent. Execution of a BT starts from the root node, which sends a “tick” command to its children. Nodes immediately return one of *Running*, *Success*, or *Failure* statuses to their parent, and are executed if and only if it receives a tick command. Control flow nodes route tick commands to their child nodes, and report status based on the statuses returned by its children. For example, a Sequence node routes ticks to its children from left to right until it finds a child that reports either *Failure* or *Running*, and itself only returns success if all its children report *Success*. A tree typically ticks repeatedly at some fixed frequency. See Figure 1 for a description of the standard node types and Fig. 2 for an example BT.

* Work supported by U.S. D.O.E. contract DE-AC02-76SF00515
† roberttk@slac.stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WECG003

# LESSONS LEARNED REFACTORING THE EuXFEL'S CENTRAL DATA ACQUISITION SYSTEM

S. Hauf*, N. Al-Qudami, D. Boukhelef, S. Esenov, G. Flucke, D. Goeries, A. Klimovskaia, P. Schmidt, A. Silenzi, J. L. Vazquez-Garcia, K. Wrona, European XFEL, Schenefeld, Germany

*Abstract*

The European XFEL, a world-leading X-ray light source, delivers high brilliance and fast burst repetition rates supported by bespoke MHz-imaging detectors, resulting in data rates close to 20 GB/s. Correlating auxiliary data from various sources with these images requires a centralized Data Acquisition (DAQ) system that ingests data from the control system, Karabo, and outputs aggregated data into a common HDF5 format. Facility-side calibration and processing rely on the stability of this data format. We discuss refactoring this major, mission-critical software to scale for future data reduction needs. The system transitioned from a blackboard to a pipelined design modeling unidirectional dataflows. Refactoring was prepared with extensive tests for verification. Unit tests, often following test-driven development, accompanied refactoring and new features. A high continuous integration coverage is enhanced by system-level tests. The refactored system was introduced incrementally over 3 months, following a "test early, fail early" philosophy. The new system successfully implements critical petabyte-scale data reduction, and enhanced status monitoring capabilities.

## INTRODUCTION

At the European XFEL [1], scientific data can be considered the primary "product" of the facility, akin to most light sources. Even the most intricate infrastructure is typically designed to support the execution of scientific research, and the critical input to this research is the data measured, encoded into bits and bytes. To interpret these measurements, various auxiliary data sets are required, which are also essential for tuning and maintaining the facility. Data are collected from a wide range of sensors and diagnostics. Notably, the EuXFEL's bespoke MHz-rate imaging detectors output data rates of up to 15 GB/s [2]. Additionally, many of the approximately 4.5 million control variables at the facility may be necessary to provide context to such data.

Data acquired at the European XFEL are represented in one of three time bases: pulse-resolved, characterized by MHz rates; pulse train-resolved, at a rate of 10 Hz; or event-driven, updating at slower rates. In practice, pulse data are aggregated into train-rate data chunks. Consequently, the time structure of the data closely mirrors the unique time structure of the facility, which comprises sub-millisecond bursts of MHz-rate photon pulses organized into a 10 Hz train structure.

The Karabo SCADA system [3] facilitates a standardized approach to ingesting this data by making it accessible through self-describing, pluggable components, known as devices, which expose control variables as properties with a well-defined set of data types. The facility's Data Acquisition system (DAQ) is also implemented in Karabo, enabling it to leverage this standardization to output scientific data into a well-defined HDF5-based data format [4, 5].

## MOTIVATION

As outlined in the Introduction, a functional and scalable Data Acquisition (DAQ) system is crucial for the successful execution of experiments at research facilities. In the case of the European XFEL, the DAQ was designed and commissioned concurrently with the majority of the facility, including the Karabo SCADA system in which it is integrated. Consequently, the system around 2023 could be considered a system with numerous "second-system effect" features [6]. The initial "first" system was designed exclusively to handle data from the MHz-detectors. It was replaced by a Karabo-integrated system responsible for data acquisition in general, including all types of "slow" data, i.e., auxiliary data accessible within the SCADA system. Integrating these aspects, alongside online stream outputs and initial data reduction capabilities, diluted the original design and destabilized the system. Stability was regained in the 2021-2023 time frame. However, it became evident that the scalability and maintainability of the system had reached its limits. Critically, this limited the options for online data-reduction implementation, a topic of high strategic relevance to a facility at which a single experiment can produce multiple Petabytes of data within a few days [7]. In mid-2023, a critical design review was undertaken, including the option for a major redesign or refactoring. Notably, an operational DAQ system would need to be provided to the facility throughout the entire upgrade process.

## PROCESS CONSIDERATIONS

The decision to refactor the extensive DAQ codebase came after the responsibility for the system internally transitioned from the IT&Data Management group to the Control software group in July 2023. A transition task force consisting of 5 individuals was formed. Initially, a clean-up and survey of the code base as of July 2023, in the following identified as the *legacy* code, was to be performed, which was to inform the further approach. This survey identified the following main areas of work:

- Remove unreachable code, that was conditionally compiled depending on preprocessor statements. This was to be carried out first, as to reduce the overall code size

* steffen.hauf@xfel.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WECR002

# THE MAX IV KITOS EXPERIENCE: ORGANISING RAPID RESPONSE OPERATION SUPPORT IN A USER FACILITY

M. Lindberg*, M. J. P. Alcocer, J. Brudvik, D. Ermakov, J. Forsberg, Á. Freitas,
M. Leorato, J. Lidón-Simón, S. Soltani, D. Spruce, M. Yazdi-Rizi
MAX IV Laboratory, Lund, Sweden

*Abstract*

To enhance operational stability and user satisfaction, the ICT groups at MAX IV Laboratory (formerly KITS) have implemented a structured support model inspired by the European XFEL's Data Department Operations Centre (DOC). The KITOS (KITS Operation Support) initiative ensures expert coverage during working hours and a best-effort support model during evenings and weekends, staffed by a rotating team of two on-duty specialists. The team leverages proactive system monitoring and dedicated on-duty time for cross-training and knowledge sharing.

KITOS supports all critical phases of user operation—including accelerator operations, beamline commissioning, experiment setup, and active beamtime—through a combination of real-time response and structured on-call coverage. Central to its success is detailed issue logging, which enables the identification and resolution of recurring problems, whether technical, procedural, or related to documentation and training gaps.

Since its inception in 2022, the KITOS model has demonstrably improved control system reliability. This presentation will share our implementation approach, key lessons learned, and impact metrics with the broader scientific facility community.

## THE STORY OF KITOS

The MAX IV Laboratory is located in Lund, Sweden and started user operations in 2016 as the world's first 4th generation synchrotron light source. It consists of a linear accelerator, a short-pulse facility and two storage rings with 16 beamlines in operation and a 17th beamline under construction. MAX IV receives 2000 users each year from countries all over the world, performing experiments involving spectroscopic, imaging, and scattering/diffraction techniques. Evolving endstations and complex sample environments pose a continuous challenge to delivering robust and reliable beamtime to our users.

### *ICT Group Organisation*

Setting the scene, the ICT (IT and Communications Technologies) department at MAX IV [1] encompasses 4 groups, IT Infrastructure, Electronics, Software and Scientific Data and Detectors, in total 55 staff. These four groups share responsibility for developing and maintaining the MAX IV control system and IT infrastructure. Together the ICT groups maintain and staff the support organisation known as KITOS[1]. This serves the accelerator and beamline staff and addresses control system issues that impact user operation.

### *Inspiration Strikes*

In 2021 Darren Spruce, then manager of the ICT groups, spent time at the European XFEL. There he worked with Steve Aplin, Head of the data department to set up a similar organisation aimed at having a more live and pro-active support for running experiments. After working with the data department management team on a draft description which defined the mission, he assigned it to Steffen Hauf to run a task force to set up the Data Department Operations Centre (DOC). This agile run project created a high motivation among the different teams and was a huge success from its inception. This inspired him to bring the concept to MAX IV. Key persons from each ICT group were asked to form a taskforce responsible for establishing a support structure, that could deliver an immediate response to operational issue, and which could run in parallel to all the other work activities required of the various groups. Putting this focus on the operational aspect gave a better insight into operational issues and also to the implementation of tools for continuous improvement of the control system reliability.

### *Struggling to Support*

MAX IV had gone into operation in 2016 with seven beamline projects being developed, three of which were taking users in 2018 [2]. The second wave of beamline projects was coming online for a total of 16 beamlines in operation by 2024. Focus was split between rapid development and supporting reliable operations and the organisation was struggling to keep up. Most members of the ICT groups were working in the beamline installation projects while concurrently supporting commissioning activities and user operation. The size of the organisation could not support a split into a project group and an operations group.

The control system of a large scale research facility such as MAX IV is both highly complicated and highly complex. Thousands of diverse components need to be integrated, across multiple sub-systems. Configuring and calibrating each one, while at the same time managing interdependencies, dynamic behaviour and unpredictable inputs, poses real challenges to ensuring that the control system is reliably operational. The reliability of the control system directly impacts the main objective of a user facility — user beamtime availability.

* mirjam.lindberg@maxiv.lu.se

[1] A play on the original department name KITS (Kontroll och IT System) and the Finnish word for thank you, kiitos.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WECR003

# ACHIEVING SUSTAINMENT: FROM DAILY OPERATIONS TO LONG-TERM STRATEGIC PLANNING*

M. Pieck†, G. E. Dale, M. S. Gulley, E. L. Kerstiens,
C. Lindenbaum, M. J. McMillen, W. J. Peck, H. A. Watkins,
Los Alamos National Laboratory, Los Alamos, NM, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Particle accelerator facilities support a wide array of applications and vary greatly in type, purpose, size, construction cost, and both the sources and levels of operational funding. While the lifespan of each accelerator is shaped by a range of technical and organizational factors, many continue to operate well beyond their originally projected lifetimes. Sustaining long-term operations requires organizations to bridge the gap between the immediate demands of daily activities and broader strategic goals. This includes aligning short-term actions with long-term objectives, fostering a culture of continuous improvement, and implementing a clear, well-communicated strategic plan. Although each facility faces its own unique challenges, this article offers an initial framework for supporting operational sustainability. It also seeks to inspire facilities to define and pursue their own paths toward lasting success in beam delivery—while encouraging collaboration and knowledge exchange across the accelerator community. The framework presented here draws on ongoing efforts within the Accelerator Operations & Technology (AOT) Division, which oversees the operation and maintenance of the Los Alamos Neutron Science Center (LANSCE) accelerator and is further illustrated through examples from Control and Instrumentation systems.

## INTRODUCTION

Particle accelerators such as LANSCE support a broad spectrum of critical applications, including high-energy physics research, cancer treatment, industrial materials testing, security screening, materials science, and energy research. While many of these accelerators and their supporting facilities were built decades ago, they continue to play essential roles in national and scientific missions—LANSCE, for instance, has been in operation for 53 years and is expected to continue operating at least until 2050 [1, 2]. However, the organizations responsible for maintaining these aging infrastructures now face a variety of significant challenges, including those as outlined below:

- **Aging Infrastructure:** Older accelerators require substantial investments to maintain safe and reliable operations, particularly if maintenance and upgrades have been historically deferred. Aging systems increase the risk of equipment failures, technological obsolescence, compatibility issues, and elevated safety and cybersecurity vulnerabilities.
- **Financial Constraints:** Limited budgets frequently force the deferral of critical maintenance and upgrades, resulting in reduced operational efficiency, higher long-term costs, increased downtime and disruptions, and harm to the organization's reputation.
- **Prioritization and Resource Allocation Issues:** Without a clear understanding of current needs and a defined vision for the future, resources—labor and materials—are allocated inefficiently. This reduces maintenance effectiveness, contributing to decreased beam availability, lower overall system reliability, and diminished returns on investment.
- **Difficulty Demonstrating Funding Needs:** Many organizations struggle to effectively document and communicate the importance of maintenance and upgrade funding, reducing possible leadership support. This further exacerbates resource shortages, resulting in greater deferred maintenance, increased downtime, and compromised performance.

To address these challenges organizations must take strategic actions:

- **Technology Utilization:** Implement preventive and predictive maintenance practices, supported by facilities-management software, to track and extend equipment life, minimize downtime, streamline maintenance processes, enhance asset tracking, and ensure sustained operational reliability and efficiency.
- **Integrated and Data-driven Planning:** Adopting a comprehensive planning approach—spanning short to long-term horizons—that leverages data to optimize objectives, manage resources efficiently, effectively coordinate modernization projects, meet stakeholder expectations, and ensure regulatory compliance.
- **Alignment of Short-term Actions with Long-term Objectives:** Foster a culture of continuous improvement through the implementation of a clear, well-communicated strategic plan that effectively balances immediate operational demands with future goals. Regularly review and update the plan to maintain relevance in response to evolving requirements.
- **Securing adequate funding:** Accurately determine funding needs through credible assessments and clear estimates, effectively communicate the importance of facility sustainment and potential consequences of insufficient funding to leadership and funding agencies while proactively explore alternative funding sources.

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001).

† pieck@lanl.gov, LA-UR-25-29145

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG001

# THE LASER MEGAJOULE FACILITY STATUS REPORT

S. Palmier[1,†], J.P. Airiau[1], G. Boutoux[1,2], C. Lacombe[1], J. Neauport[1], P. Prunet[1,2]
[1]CEA, CESTA, Le Barp, France
[2]CEA, DAM DIF, Arpajon, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Laser Megajoule (LMJ), a 176-beam laser facility developed by CEA, is located at the CESTA site near Bordeaux. The LMJ facility is part of the French Simulation Program, which combines improvement of theoretical models and data used in various domains of physics, high performance numerical simulations and experimental validation. It is designed to deliver about 1.3 MJ of energy on targets, for high energy density physics experiments, including fusion experiments.

In this paper, a review of the LMJ facility project is given with details on the status report and an update of the activities on the laser system, diagnostics and control system. Afterwards, some recent experimental results are presented.

## INTRODUCTION

Since it definitively abandoned nuclear testing, France relies on the Simulation Program to guarantee the operational performance and safety of its nuclear deterrent weapons throughout their lifetime.

Successful simulation requires both:

- Qualified computer codes that integrate laboratory-validated physics models to simulate weapon functioning;
- Teams of qualified physicists to use these codes.

In this respect, the Megajoule Laser (LMJ) [1] plays a vital role, as it is used to validate the numerical codes and certify the skills of French physicists.

The LMJ is designed to perform indirect drive inertial confinement fusion (ICF) experiments on deuterium tritium targets. Aside this major objective, a broad spectrum of experiments can also be addressed from high energy density physics to basic science [2]. This keystone of the Simulation Program combines physics models improvement, high performance numerical simulation, and experimental validation, in order to guarantee the safety and the reliability of French deterrent weapons. When completed, the LMJ will deliver a total energy of 1.3 MJ of 0.35 µm (3ω) light and a maximum power of 400 TW.

This facility is dimensioned to accommodate 176 beams grouped into 22 bundles of 8 beams. These beams are located in the four laser bays arranged on both sides of the central target bay of 60 meters diameter and 40 meters height. The target chamber and the associated equipment are located in the center of the target bay.

The first bundle of 8 beams has been commissioned at the end of 2014. The second bundle at the end of 2016 following the same commissioning process. Nineteen bundles are now operational by the end of 2025.

The laser facility offers also an ultra-high intensity capacity with the PETAL (Petawatt Aquitaine Laser) beamline thanks to the financial support of Nouvelle Aquitaine Region and European Union (FEDER). PETAL provides a single short pulse peak power laser beamline that broadens the facility experimental capabilities [3]. The LMJ and PETAL laser facilities are both hosted in a dedicated building featuring dimensions of 300 m in length, 150 m in width, and 40 m in height.

The first LMJ bundle was operated in 2014 for the first time. Over the past decade the facility has been progressively upgraded with additional bundles. It has also been used at the same time as an experimental facility. To limit the maintenance of final optics of operating bundles while mounting and commissioning novel bundles, the LMJ beamline energy and power operation has been limited by half, i.e., 3.75 kJ/1.25 TW per beamline at the wavelength of 351 nm [4]. Some laser campaigns are performed to gradually demonstrate laser performances at elevated energy and power up to the laser design limit in order to prepare full energy and power operation.

This paper summarizes the report status of the LMJ/PETAL facility. The first section gives an overview of the main laser system. The second section is dedicated on diagnostics status. The third section discusses the control system with emphasing on the last updates. The latest experimental capabilities implemented on the LMJ/PETAL platform are then presented in the last section.

## LMJ LASER SYSTEM

### *Laser Schema*

The 176 beams (35 x 35 cm$^2$ each) are grouped into 22 bundles of 8 beams. In the switchyards, each individual bundle is divided into two quads of 4 beams, the basic independent unit for experiments, which are directed to the upper and lower hemispheres of the target chamber.
Schematically the laser system is divided into four main sub-assemblies represented on Fig. 1.

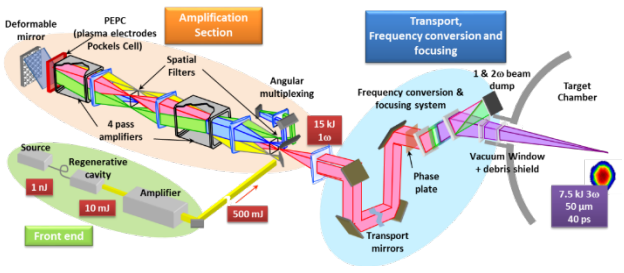



Figure 1: LMJ laser beamline.

(i) a front-end that creates and pre-amplifies the laser pulse to the joule. It provides the desired temporal pulse

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG003

# FIRST PHASE OF CONTROL SYSTEM FOR COMPACT MUON LINAC AT J-PARC

M. Yang†, N. Kamikubota, S. Yamada, M. Kimura, M. Otani
High Energy Accelerator Research Organization (KEK), Ibaraki, Japan

*Abstract*

A muon linear accelerator (linac) for the muon *g-2*/EDM experiment is currently under construction at Japan Proton Accelerator Research Complex (J-PARC). The objective of this project is to accelerate thermal muons (25 meV at 300 K) to 212 MeV, marking the world's first implementation of muon acceleration.

Development of the control system for the Muon Linac began in 2024. Core functionalities of the Ultra-Slow Muon section have been tested during beam time in May 2025.

The system adopts the standard EPICS framework and features a compact architecture consisting of (a) a QNAP NAS for disk storage and LDAP-based user authentication, (b) two operator terminals, and (c) two commercial micro servers serving as the EPICS IOC and the archiver server, respectively.

This paper reports on the status of the control system development for the Ultra-Slow Muon section, as part of the first phase of the Muon Linac project. Toward the full commissioning of the entire Muon Linac in 2028, the prospects for extending the present control system to the main linac components are discussed.

## INTRODUCTION

The E34 experiment at the Japan Proton Accelerator Research Complex (J-PARC) [1] is designed to explore fundamental muon properties, including the anomalous magnetic moment (*g-2*) and the electric dipole moment (EDM), with unprecedented precision [2, 3]. To achieve this goal, a new linear accelerator dedicated to ultra-slow muons, the Muon Linac, is being developed in the Materials and Life Science Experimental Facility (MLF). It will accelerate muons up to the energy of 212 MeV [4-6].

The staged acceleration scheme is summarized in Fig. 1. Initially, muons from the Ultra-Slow Muon (USM) section are captured and accelerated to 0.34 MeV using a radio-frequency quadrupole (RFQ) [7]. The beam is then transported through successive accelerating structures: an interdigital H-mode drift tube linac (IH-DTL), a disk-and-washer (DAW) coupled cavity linac, and finally a series of disk-loaded traveling wave structures (DLS). This configuration was chosen to balance compactness with beam quality, enabling the production of clean, high-energy muons suitable for precision experiments.

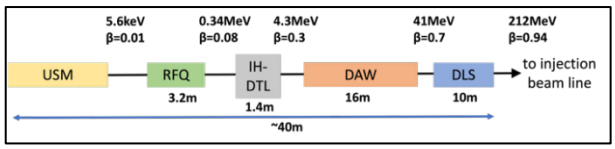



Figure 1: Structure of Muon Linac.

Although the Muon Linac itself is compact compared with other J-PARC machines, its operation still requires integration of various subsystems, including RF, magnets, vacuum, diagnostics, and laser equipment. The control system for it is therefore designed within the EPICS (Experimental Physics and Industrial Control System) [8] framework, consistent with the standards adopted across the J-PARC accelerator complex [9]. In this project, the large-facility framework serves as the basis for a compact-scale implementation, ensuring consistency with existing infrastructures while allowing adaptation to the specific requirements of the Muon Linac.

In April 2025, the control system for the USM section was nearly completed as the first phase. The system was used successfully during the muon beam commissioning in May 2025. This paper reports on the Muon Linac control system, focusing on the practical aspects of its initial construction. It describes the implementation of key control system components, such as the naming rules for devices, directory and file structures, LDAP-based user authentication, and so on. These efforts lay the foundation for future expansion of the system while ensuring compatibility with J-PARC-wide operational practices.

## OVERALL CONTROL SYSTEM

The Muon Linac control system is designed based on the standard architecture of the J-PARC accelerator control system. The system has been independently built to suit the scale and constraints of the Muon Linac project.

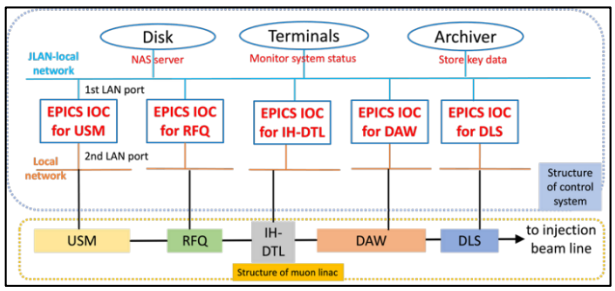



Figure 2: The architecture of overall control system.

† yangmin@post.kek.jp

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG004

# THE EUAPS BETATRON RADIATION SOURCE CONTROL SYSTEM*

V. Dompè†, M. Ferrario, G. Latini, S. Pioli
INFN - Laboratori Nazionali di Frascati, Frascati, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

EuAPS (EuPRAXIA Advanced Photon Source) is a project carried out in the EuPRAXIA context and financed by the Italian Ministry in the Recovery Europe plan framework. A new advanced betatron radiation source, obtained by exploiting plasma LWFA, is currently being realized at the Laboratori Nazionali di Frascati of INFN in Italy and will be operated as user facility. Several elements of EuAPS are remotely controlled, such as laser diagnostic devices, motors and vacuum system components. In order to efficiently run the facility, the realization of a robust and performing control system is crucial. The EuAPS control system is based on EPICS (Experimental Physics and Industrial Control System) open-source software framework. Functional safety systems such as Machine Protection System (MPS) and Personnel Safety System (PSS) in accordance with IEC-61508 standards are also integrated for interlocks control, anomalies monitoring and for protecting the personnel from hazardous areas. In this contribution, details on how the EuAPS control system has been projected and realized will be provided.

## INTRODUCTION

The EuAPS (EuPRAXIA Advanced Photon Source) project, carried out in the framework of EuPRAXIA [1], aims at the realization of innovative facilities in Italy, paving the way to advanced applications in the field of photon science [2, 3]. Among those, a new betatron radiation source that will be operated as a user facility is currently being constructed at Laboratori Nazionali di Frascati of INFN, involving the whole SPARC_LAB facility [4].

The EuAPS betatron radiation source will exploit the plasma Laser Wakefield Acceleration (LWFA) technique to generate the photon radiation. To this purpose, a high power laser, such as in the FLAME facility [5], will be transported and focused on a gas target positioned into a vacuum chamber to produce plasma by gas ionization, accelerated electrons by the self-injection mechanism, and the betatron radiation from their oscillatory motion.

The realization of such a facility requires, therefore, the presence of advanced and high precision laser and vacuum systems. Alongside these, a reliable high-performance and usable control system is crucial for the continuous and efficient operation of the facility.

The EuAPS betatron radiation source control system is based on EPICS, as this represents a common standard for many other scientific facilities around the world and has extensively demonstrated its robustness and scalability, the latter being particularly important also with a view to EuPRAXIA. These characteristics ensure that requirements such as longevity, maintainability and usability are met as well. In this contribution, we describe the architecture of this control system, we report on the elements integration method and current status, and we lastly discuss developments foreseen in the near future.

## ARCHITECTURE

The EuAPS control system architecture is structured in three tiers:

- User interface (Fig. 1 - upper right) - referring to graphical and non-graphical user interfaces;
- Central services (Fig. 1 - upper left)- representing the computer services that need to continuously run regardless of user activities, e.g. for process variables' values archiving, alarm states monitoring, etc.
- Equipment interfaces (Fig. 1 - bottom) - relating to the system responsible for the interaction with equipment and devices, the elementary entity being the EPICS Input Output Controller (IOC), which provides an abstract representation of the devices to higher layers, through which they are monitored and controlled.

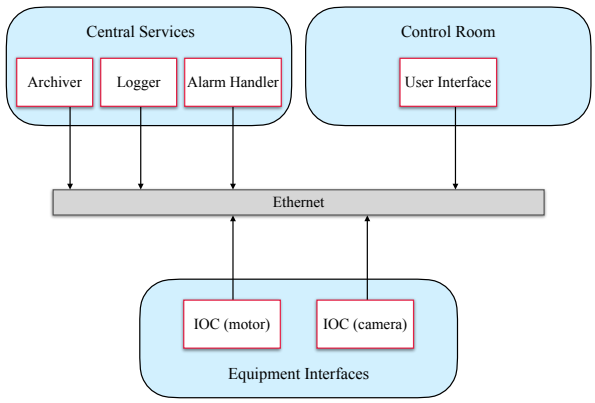



Figure 1: Scheme of the three tiers architecure of the EuAPS control system.

The architecture comprises 140 local control systems belonging to two main subsystems, namely the laser and vacuum ones, as well as to diagnostics and to functional safety systems (machine protection and personal safety systems). The latter constantly communicate with the control system alarm

* This work was funded by the European Union - Next Generation EU, Mission 4 Component 1, CUP I93C21000160006
† valentina.dompe@lnf.infn.it

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG005

# LCLS-II CAVITY HEATER CONTROLS: DESIGN, OPERATION, AND ACCURACY

A. Wilson†, R.S.P. Chillara, M. Keenan, F. Moguel, E. Nkwocha, S. Shrishrimal
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

The SLAC National Accelerator Laboratory's upgrade to the LCLS-II, featuring a 4 GeV superconducting linear accelerator with 37 cryomodules and two helium refrigeration systems supporting 4 kW at 2.0 K, represents a significant advancement in accelerator technology. Central to this upgrade is a 2K system with five stages of centrifugal cold compressors, operating across a pressure range from 26 mbar suction to 1.2 bara discharge. These dynamic centrifugal compressors have a limited operational envelope hence maintaining stable pressure and flow is critical for its operation. This paper describes how SLAC achieved stable LINAC pressures in each of the 37 Cryomodule using electrical heaters compensating actively to the changes in RF power to maintain constant flow through the system. Additionally, this paper details the power accuracy of these heaters, which can be useful not only for control, but also when measuring cavity efficiency.

## INTRODUCTION

The LCLS-II project at SLAC will produce a nearly continuous X-ray laser beam, when it achieves a repetition rate of one million pulses per second at 4 GeV. This performance is made possible by 37 cryomodules each containing eight nine-cell TESLA superconducting radio frequency (SRF) cavities cooled to 2.0 K by a powerful helium refrigeration system. Central to the cryogenic system is a 2K cold box (2KCB) with five stages of centrifugal cold compressors (CC) operating from 26 mbar suction pressure to 1.2 bara discharge [1]. These CC have a limited operating range making it critical to control suction pressure, suction temperature, and flow [1]. Suction pressure varies with the heat load in the system. The primary heat load in the system is the SRF cavities, which produce a variable heat load. The heat load and thereby suction pressure is controlled by adjusting the power of electrical heaters installed on the helium vessel of each cavity for each cryomodule in the LINAC (cavity heaters).

## CAVITY HEATER CONTROL

To achieve the desired suction pressure, a PID controller adjusts the total cavity heat power for the LINAC to bring the measured suction pressure to the chosen setpoint. This total power is evenly distributed to each of the 37 cryomodules. For each cryomodule the control system deducts the heat power generated by the 8 SRF cavities as calculated by the low level radio frequency (LLRF) control system and distributes the remaining power to each cavity heater in operation (Fig. 1). A voltage command is calculated for each heater based on this requested power.

A 4-wire variable power supply delivers the commanded voltage to the heater regardless of voltage drop across the cable. The current through the heater is monitored by the control system, but no adjustments are made to the requested voltage based on feedback. This ensures that rapid adjustments due to a changing RF heat load are applied immediately and consistently. The PID controller has an adjustable ramp rate, but all other changes to the cavity heater power are applied immediately.

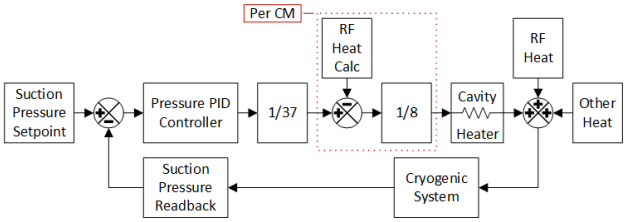



Figure 1: Pressure Control Diagram.

The measured current and requested voltage are used to calculate the delivered power. If the delivered power and requested power diverge, an alarm is activated, but no automatic action is taken.

### *Interlocks*

Cavity heater power supplies are deenergized when interlocked. In an interlock scenario, the PLC deenergizes removes input power from the heater power supplies.

Two interlock scenarios exist specific to the cavity heaters:

- For each cryomodule, liquid helium does not fill the helium vessel for each cavity
- For the entire LINAC, system pressure is above the maximum operating pressure

The heaters are protected from burn out by the cooling provided by the liquid helium, and the cryogenic system avoids pressure relief scenarios caused by heater power. In all other cases during normal operation, the heaters remain energized and responsive to requested power.

### *Hardware*

Each cryomodule contains one cavity heater per cavity, 8 in total. The heaters are polyimide strip elements affixed to the bottom of the helium vessel. The nominal resistance is 46.3 Ω.

Each heater is powered by its own 0-50 VDC variable power supply over a 4-wire connection (Fig. 2). Acopian VY050MX120M linear, fanless, regulated power supplies are used (Fig. 3). The power supply adjusts its output so that the commanded voltage is delivered at the heater terminals, compensating for voltage drop across the cable.

* Work supported by Department of Energy contract DE-AC02-76SF00515
† andreww@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG006

# AN FPGA-BASED AUTOENCODER MODEL FOR REAL-TIME RF SIGNAL DENOISING FOR INDUSTRIAL ACCELERATORS*

V. Rajesh†, Liberty High School, Frisco, TX, USA
J. Edelen, J. Einstein-Curtis, RadiaSoft LLC, Boulder, CO, USA

*Abstract*

A challenge that industrial particle accelerators face is the high amounts of noise in sensor readings. This noise obscures essential beam diagnostic and operation data, limiting the amount of information that is relayed to machine operators and beam instrumentation engineers. Machine learning-based techniques have shown great promise in isolating noise patterns while preserving high-fidelity signals, enabling more accurate diagnostics and performance tuning. Our work focuses on investigating the challenges associated with the implementation of a noise reduction autoencoder that operates in real time on a Field Programmable Gate Array, which we do by creating firmware to run on a Xilinx ZCU104 evaluation kit with the intention of being deployed on industrial particle accelerators in the near future.

## INTRODUCTION

Particle accelerators have wide-ranging applications in industrial processes such as medical device sterilization, food irradiation, and polymer cross-linking. Industrial applications tend to have more relaxed requirements on stability and beam quality, while having high up-time requirements. As a result, they tend to be more robust systems that lack the high degree of sophistication associated with their laboratory counterparts. Recent advances in accelerator technology have brought to light new commercial applications for accelerators in imaging and defense that have increasingly stringent tolerances to control. In addition, applications such as FLASH radiotherapy require precise beam control to achieve the best outcome for patient care.

A problem that arises within industrial particle accelerators is that, as opposed to their research counterparts, their operating conditions are far more electromagnetically "dirty," leading to their sensor readings to accrue a lot of noise, which creates a set of data that does not clearly reflect the actual conditions that are occurring within the particle accelerator. The high amounts of noise in the readings from recording points, such as the inside of the RF cavity, greatly reduce the ability for machine operators to accurately control the accelerator. Our work is focused on the application of machine learning technology to improve signal processing for RF systems to achieve a higher degree of stability.

Our previous work demonstrated various autoencoder-based approaches are well suited to remove signal noise from RF data while accurately reconstructing the baseline signal [1]. In this work, we expand on our previous effort by addressing the challenges associated with the deployment of an ML algorithm on a real-time system operation on a field-programmable gate array (FPGA).

We employ machine learning model optimization methods such as pruning and quantization to improve model performance and reduce its size to fit within the hardware constraints of the FPGA acceleration target. Finally, a software tool called hls4ml [2, 3] is utilized to generate a firmware implementation of the autoencoder.

## DATA GENERATION

The dataset used to train the autoencoder was generated using RadiaSoft's RF signal simulator [4], a tool designed for the research and development of RF control algorithms. The simulator utilizes a scattering matrix approach in order to describe RF signal propagation and an RLC circuit to model cavity dynamics, accounting for parameters such as coupling/quality factors, frequency, drive amplitude/phase, pulse duration, and cavity detuning. The foundational dynamics are derived from prior work [5–7], and additive white Gaussian noise applied to measurements to replicate the noise accrued in an industrial setting. Parameters that mimic the conditions that an industrial particle accelerator operates under were chosen. These include a cavity resonance frequency of 2856 MHz, pulse lengths that vary between 3 µs and 7.5 µs, and quality ($Q$) and coupling ($\beta$) factors of $1.00 \times 10^4$–$2.25 \times 10^5$ and 1–3, respectively. The detuning values varied in a half bandwidth above or below the cavity resonance frequency. Each signal recording lasts a total of 200 samples.

## PERFORMANCE EVALUATION METRICS

We utilized two metrics to evaluate the performance of a given value that we created and trained. One was a mean squared error (MSE) term, which was calculated using the residual term between the model's reconstructions $Y$ and the noiseless ground-truth data $\hat{Y}$ given by

$$\frac{1}{n}\sum_{i=1}^{n}(\hat{Y} - Y)^2$$

where $n$ is the number of data points that were generated. Although our model training is unsupervised, meaning that clean signals were not used as recreation targets during training, we evaluated performance using MSE against the clean signals made available by the data simulation. MSE was chosen due to its high sensitivity to large differences between the reconstructed signals and the ground-truth. If the

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Accelerator R&D and Production Award Number DE-SC0023641.

† viksharrajesh@gmail.com

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG008

# OBJECT-ORIENTED INDUSTRIAL I/O FOR EPICS ON NI cRIO: REUSABLE LabVIEW-FPGA BITFILES VIA THE NI C API

H. R. Martin*, S. A. Baily†, C. D. Hatch‡
Los Alamos National Laboratory, Los Alamos, NM, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Los Alamos Neutron Science Center (LANSCE) has completed a significant modernization effort, migrating from the legacy RICE (Remote Instrumentation and Control Equipment) control system to an infrastructure based entirely on the Experimental Physics and Industrial Control System (EPICS). A key enabler of this transition has been the development and deployment of a modular, object-oriented Industrial I/O (IIO) architecture on NI CompactRIO (cRIO) platforms. The IIO framework provides a reusable and scalable system for controlling and monitoring sensors and instruments. It is built around precompiled field-programmable gate array (FPGA) bitfiles accessed through the NI C application programming interface (API). Where necessary, LabVIEW Real-Time code integrates seamlessly with EPICS input/output controllers (IOCs). This architecture cleanly separates control logic from hardware interfaces, supports long-term maintenance with minimal overhead, and accommodates both modern NI Linux RT cRIO and legacy VxWorks systems. The result is a flexible and resilient method for managing and improving complex control architectures across LANSCE.

This contribution outlines how IIO enables hardware reuse by treating NI cards as modular components with shared logic, abstracting low-level FPGA interaction, and standardizing configurations through parameterized bitfiles and EPICS startup files. The paper focuses on how this approach supports object-like behavior to improve maintainability, scalability, and cross-platform deployment of EPICS-compatible systems.

## INTRODUCTION

LANSCE slow controls and monitoring are implemented with an IIO stack that exposes field signals to EPICS [1] and, where useful, NI's EPICS integration options [2]. At the edge, NI cRIO chassis host precompiled FPGA bitfiles. EPICS device support calls the NI C API to read and write FPGA registers deterministically, while optional LabVIEW Real-Time (LV-RT) glue code is used where it adds value (e.g., device-specific sequencing). This separation keeps control logic independent of hardware specifics and enables repeatable deployments.

* Email: rociomartin@lanl.gov
† Email: sbaily@lanl.gov
‡ Email: chatch@lanl.gov
LA-UR-25-29042.

## RELATED WORK

EPICS+LabVIEW integration has followed several patterns. One line embeds the IOC alongside LabVIEW using `lvPortDriver` (an `asynPortDriver` subclass) [3–7]. A second uses Channel Access directly from LabVIEW (DSC Module / pure-LV CA) [8–10]. Our approach emphasizes precompiled LabVIEW-FPGA bitfiles addressed via the NI C API from EPICS device support.

Additional community material documents early practices and options for interfacing commercial off-the-shelf (COTS) hardware to EPICS via LabVIEW, as well as practical connection notes [11–14].

A complementary pattern for certain diagnostics embeds an IOC alongside LabVIEW RT using `lvPortDriver` so the IOC exchanges data with the LabVIEW process rather than the hardware directly [3,4,7]. That path is out of scope here; this paper focuses on the IIO (slow controls) path in which EPICS talks directly to the FPGA via the NI C API.

Finally, as noted in community reports, the most challenging RICE feature to replicate was custom high-rate acquisition; those use cases were moved to dedicated waveform-digitizer IOCs while preserving the same naming and archiver conventions [4].

### *Clarifications and Additions*

While pure-LabVIEW Channel Access implementations exist [8], the LabVIEW DSC module's EPICS I/O servers use the EPICS *portable Channel Access server* (PCAS) (the same C code used by the EPICS Gateway) and therefore provide CA endpoints without EPICS records or database processing [9,10]. A fourth method, *CALab* [15], offers a CA client interface for LabVIEW; in that pattern a soft IOC runs on the NI device and LabVIEW interacts with it solely via CA client calls. The Industrial I/O (IIO) library underlying our approach traces back to LANSCE work documented in [14]; see also [16] for a broader description of IIO software functionality.

## WHAT'S NEW IN THIS WORK

- Object-oriented (OO) refactor of LabVIEW–FPGA: earlier procedural code was upgraded to an OO design to reduce coupling and simplify module onboarding and maintenance.

- Reusable bitfiles: a slot-map–parameterized LabVIEW-FPGA design producing a single `.lvbitx` per cRIO model/slot-map; reuse is guarded by a build/signature check.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG009

# MODERNIZING FPGA DEVELOPMENT USING THE DESY FPGA FIRMWARE FRAMEWORK *

D. M. Vassallo†, M. Costanzo, J. Jamilkowski, K. Kulmatycski, K. A. Fahey, K. Mernick, A. Singh
Brookhaven, National Laboratory, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Brookhaven National Laboratory (BNL) is currently developing new hardware description language (HDL) code and embedded software for the Electron-Ion Collider (EIC) control system. Part of this effort is modernizing the development process itself, leveraging methodologies and tools that were initially targeted at the software world. These methods include effective source control and project management, modularization and rapid deployment of updated code, automated testing, and in many cases automated code generation. HDL designers additionally face unique challenges compared to software designers, particularly with vendor locking and dependency on particular tools and IP. The FPGA Firmware Framework (FWK), developed by DESY, is a set of tools that helps to both apply these modern methods and to overcome some of those unique challenges. This paper will cover the workflow, successes, and challenges faced when using the FWK. In particular, we will focus on the experience using this workflow to develop a customizable delay generator IP targeting a Zynq FPGA.

## INTRODUCTION

The Electron-Ion Collider (EIC) project has chosen to use the DESY FWK Firmware Framework (FWK) [1] to support the development of Hardware Description Language (HDL) code and associated software and firmware for the Common Hardware Platform (CHP). The FWK framework was originally developed by DESY and is part of the larger trend of digital logic design becoming more like software design. FWK was designed to give gateware developers easier access to recent advances in the software development workflow. Examples of this include, but are not limited to, modern source control and repository management, modularization, rapid deployment, automated testing, and automated code generation.

Brookhaven National Laboratory (BNL) is currently using FWK to design gateware components for CHP. One such component is a CHP equivalent for the Relativistic Heavy Ion Collider (RHIC)-era V202 delay generator board. It is currently being developed under the working name C202, and it will be deployed on the latest generation of Xilinx FPGAs. This development is happening in parallel with the new CHP Timing Data Link (TDL) [2] and CHP carrier and daughter boards. Development of the C202 also serves as a "live test" of the new gateware development workflow using FWK, emphasizing shared gateware modules across the EIC development team.

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
† dvassallo@bnl.gov.

Development and testing of the underlying gateware code has proceeded on physical hardware while key CHP components (new carrier board revisions) are ironed out. Specifically, the C202 is currently being developed on the ZCU-104 Development board, which utilizes the same FPGA family as the final carrier board, the Zynq Ultrascale. When the carrier and/or TDL are in a more mature state, development will be switched to these platforms, facilitated by the modular nature of FWK. Furthermore, the C202 itself will be a FWK compliant git submodule, allowing for easy integration into future designs.

## TRENDS BACKGROUND

Over the past decade the software development process has changed dramatically due to the introduction of new tools and new methodologies. Code development has become more collaborative, more modular, less tool-dependent, and testing has become more automated. Gateware development historically lags behind in the adoption of these new tools and workflows. BNL is seeking to modernize its Gateware development process, and has adopted FWK to help facilitate this.

The widescale adoption of Git is one of the key drivers of this trend. Forges such as Github and Gitlab have become testing grounds for new development methodologies and technologies. FWK is built around a git-based workflow. Gateware components are explicitly git repos and designed to be included as git submodules. While git-based workflows have existed in gateware development before, FWK provides a flexible repository organization scheme and isolates the source files from the tool specific build artifacts.

### *CI/CD*

Continuous Integration/Deployment (CI/CD) is one of the latest trends in software development. The full concept is to speed up development and deployment of new software by automating much of the testing, building and deployment process. Numerous tools exist to enable this, primarily Jenkins, Github actions, and Gitlab CI. They all rely on automated build, test, and deployment tools that directly interact with git repos, and in many cases are built into the source forges themselves.

Gateware does pose some unique challenges in a CI/CD pipeline compared to software. First and foremost, gateware is inherently tied to an FPGA (or an ASIC, but that is outside the scope of this paper), so full automated deployment requires some kind of hardware-in-loop, which is often impractical. Additionally, build times are both much longer and the build process is often much more complicated. One other challenge for CI/CD in gateware is that gateware tends to be more "tool dependant" than software.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG010

# PROTON PULSE CHARGE CALCULATION ALGORITHM IN BEAM POWER LIMITING SYSTEM AT SPALLATION NEUTRON SOURCE

Miljko Bobrek[1], Trent Allison[2], Craig Deibele[1]
[1]Oak Ridge National Laboratory, Oak Ridge, TN, USA
[2]Osprey Distributed Control Systems, Ocean City, MD, USA

*Abstract*

A proton pulse charge calculation algorithm in Beam Power Limiting System at the Spallation Neutron Source was developed and implemented in FPGA. The algorithm calculates one-minute running average of the pulse charges and issues a fault to the Personal Protection System (PPS) and the Machine Protection System (MPS) when a limit is reached. A bit-accurate model of the algorithm was first developed and tested in Matlab® and then implemented and simulated in VHDL using Vivado® design environment. Finally, the algorithm was verified on a µTCA-based hardware platform.

## INTRODUCTION

The proton power upgrade at the Spallation Neutron Source doubled the average proton pulse power from 1.4 MW to 2.8 MW. Considering a 2.145 MW power safety envelope for the First Target Station, the Beam Power Limiting System (BPLS) was developed and implemented to prevent equipment damage [1], [2]. The system uses fast current transformers placed at Ring to Target Beam Transport (RTBT) to measure the proton pulse current [3]. The current pulses are conditioned and digitized at 100 MHz with 11.5 Effective Number of Bits (ENOB). The pulses are approximately Gaussian shape, less than 5µs wide at the base, and around 60 Hz repetition rate.

To estimate the 60-second sliding proton pulse average power based on acquired current pulses, the pulse charge is calculated first. To calculate the pulse energy and subsequently the proton power, the pulse charge is multiplied by beam energy measured by a separate PLC-based subsystem. The pulse charge calculation algorithm design followed best design practices for safety critical systems such as BPLS.

## ALGORITHM MODEL

Matlab® bit-accurate model of the algorithm is developed first to test and characterize the algorithm performance. As the input, the model uses simulated current pulses shown in Fig. 1. Below are listed critical properties of the current signal that guided the algorithm development:

- Pulse period is around 60 Hz and is adjusted every time the beam energy setting changes
- Pulse width is less than 5 µs
- Duty cycle is very small, less than 1/3300 as the accelerator often runs at sub-60 Hz rates
- There is a slowly changing baseline in the signal that needs to be removed

Figure 2 shows a three-step algorithm model including baseline removal, pulse charge selection, and pulse charge running average steps.

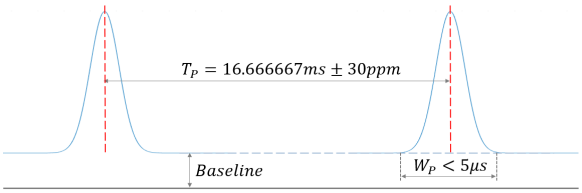



Figure 1: Simulated Current Pulses.

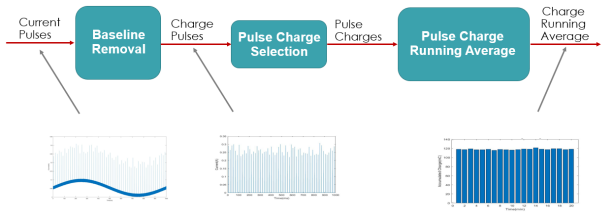



Figure 2: Three-Step Algorithm Model.

At the first step, the algorithm removes the baseline and integrates the current pulses to produce the charge pulses. At the second step, the maximum sample of the sampled charge pulses is selected for each pulse. The third step includes a sliding integrator to calculate 60-second (3600 samples) running average. The final step not shown in Fig. 2 compares the 60-second running average to a programmed threshold and notifies accelerator protection systems if the threshold is crossed.

Figure 3 shows the block diagram and the processing flow of the first step. Two sliding windows 5 µs (500 samples) wide and 5 µs apart integrate the input current samples and subtracts their outputs to remove the baseline. The signal processing path on the right side does not strictly follow the block diagram. It first subtracts the input samples 10 µs apart and then applies a single integrator to get the charge pulses without the baseline. The data path size is scaled accordingly to prevent overflow/underflow.

The main window (solid lines) in Fig. 4 is a fixed, 8.333333 ms wide acquisition window used to find the maximum in the incoming charge pulse samples. The window size is chosen to be half the distance between the pulses to make sure there is no more than one valid pulse in the window. There are two side windows 10 µs wide (dashed lines) used to resolve the boundary cases. The input charge pulses in Fig. 4 are 10 µs wide because of the convolution with a 5 µs wide rectangular window as shown in Fig. 3.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG011

# A HIGH-PRECISION MOTION SCOPE DATA STREAMING PIPELINE FOR LCLS-II FAST WIRE SCANNER*

Z. Huang†, N. Balakrishnan, J. Bong, J. Bryce, T. Thayer
SLAC National Accelerator Laboratory, Menlo Park, CA, United States

*Abstract*

The Linac Coherent Light Source II (LCLS-II) is the first X-ray Free Electron Laser (XFEL) to utilize continuous-wave superconducting accelerator technology (CW-SCRF), capable of delivering X-ray pulses at repetition rates up to 1MHz. The LCLS-II fast wire scanner motion control system, based on the Aerotech Ensemble controller, is designed to measure the beam profile across both high and low repetition rates. To effectively and timely analyze the motion trajectory of the fast wire scanner, we have developed a data streaming pipeline that transmits high-precision scope data from the Ensemble controller to the LCLS-II server. This system integrates the motion scope profile into the EPICS control system, displaying the scope profile in real time via a PyDM GUI. This paper outlines the design of the data transmission pipeline and the software development process.

## INTRODUCTION

The Linac Coherent Light Source II (LCLS-II) at SLAC is the first X-ray Free-Electron Laser (XFEL) facility based on continuous-wave superconducting (SC) accelerator technology (CW-SCRF). This enables the generation of X-ray laser pulses up to 10,000 times brighter than those produced by a normal-conducting (NC) linac [1, 2].

The LCLS-II fast wire scanner is one of the primary diagnostic tool for measuring the transverse profile of the electron beam. Each scanner uses three metallic wires mounted on a card attached to a linear motor stage. Since 2019, twenty fast wire scanners have been installed along the beamline from the RF gun to the Beam Transfer Hall (BTH). These scanners, controlled by an Aerotech Ensemble motion controller, can rapidly intercept the beam to deliver high-resolution beam profiles [3–6]. Fig. 1 is the photo of LCLS-II fast wire scanner in BTH.

In 2023, the motion control software was upgraded to support a broad range of scan velocities, enabling beam profile measurements at electron bunch repetition rates from 1 Hz to 1 MHz. The upgraded system is compatible with both NC and SC beamlines, and control is now accessible through a PyDM-based graphical user interface (GUI) [7].

Despite these enhancements, motion scope data, including wire position, scan velocity, and motor drive current, remain accessible only via archived EPICS process variables (PVs) in the Archiver Appliance (AA) at a low sampling rate. This low-resolution, non-real-time data limits its usefulness for motion diagnostics and troubleshooting. While the Aerotech controller offers high-resolution digital scope tool, accessing it requires a direct network connection to the controller, making it impractical for routine operations.

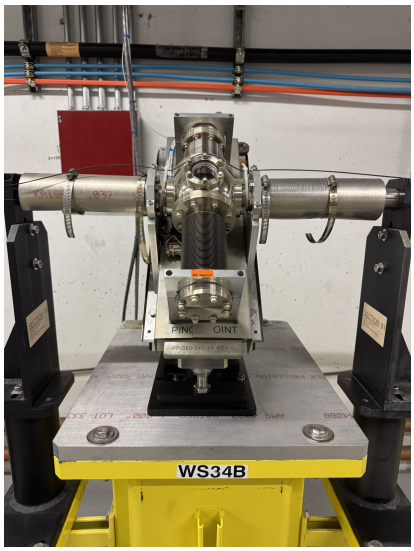



Figure 1: Photo of LCLS-II fast wire scanner.

To address this limitation, we developed a high-precision motion scope data streaming pipeline for the LCLS-II fast wire scanner control system. Implemented in Python, the pipeline acquires high-resolution scope data directly from the Aerotech motion controller and streams it to a newly designed GUI for real-time visualization. This paper describes the software architecture of the data streaming pipeline, details the GUI design, and presents performance test results and commissioning experience. Ongoing work is also discussed.

## DESIGN REQUIREMENT

For the LCLS-II fast wire scanners, the designed travel range between the low and high limits is 50 mm. Following the software motion control system upgrade in 2023, all fast wire scanners can measure the profile of the electron beam with a required maximum wire velocity of 540 mm/s. The wire scanner can be used in two scan modes based on scan velocities. Scan velocities less than 10 mm/sec are called step scan and velocities over 10 mm/sec fall under continuous scan category. This dual-mode configuration ensures compatibility with varying beam rates.

Although the existing LCLS-II fast wire scanner motion control system integrates basic motor parameters into the EPICS environment via PVs, such as motor position readback and motor current, the PV update rate is limited due

* Work supported by U.S. Department of Energy under contract number DE-AC02-76SF00515.
† zyhuang@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG012

# USE CASES FOR CONSISTENT ROBUST PROCESSING OF DATA MODELS

P. Schnizer*, G, Rehm, W. Sulaiman Khail
Helmholtz-Zentrum Berlin, BESSY, Berlin, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Many control algorithms or optimisation procedures profit from a consistent set of data which is available with a high frequency: e.g. machine learning or automated commissioning. Modern distributed control systems allow combining and presenting data based on data models, which are then transported consistently over the network: e.g. EPICS7 introduced these data models as normative types or their combination.

The authors present use cases that can profit from a consistent robust combination of data sub-models of many devices to a higher order model. Finally common patterns are presented which could be reasonable to implement independently.

## INTRODUCTION

Distributed control systems are used for operating complex machines e.g. accelerator complex or large scale experiments. These machines have hundreds to thousands of actuators (e.g. the number of magnets of a large 4th generation light source) or hundreds of dedicated diagnostic instruments (e.g. beam position monitors for modern light sources [1–5]).

All these devices work together for a common target: to provide e.g. stable beam with low emittance for a fourth generation light source. Thus one not only needs to control the individual device or collect data from an individual instrument but rather one wants to operate them as a whole or at least a subset of it: e.g. to get an update of the current orbit or to apply a new consistent settings to the steerers to drive the orbit closer back to its desired positions. Here one can see that these instruments work as an ensemble or need to be orchestrated.

These instruments exhibit their state over typically an extensive set of process-variables where only a certain subset of them is typically needed for studying and describing the machine: these are typically provided in a flat name space, which is sufficient when addressing a single device or instrument.

For real machines, however, hundreds of devices are combined e.g. beam position monitors to an orbit object. One could see it as a line camera producing data over time. Combining this data is rather trivial as long as communication issues can be neglected. Typical instruments of today provide their measurement together with a time stamp which is synchronised over all of them. Furthermore, modern devices deliver data in a data model.

* pierre.schnizer@helmholtz-berlin.de

Data from many devices in a family can be combined, especially if all data packages that are part of a measurement contain the same tag. Handling it in a simple and efficient manner is not so trivial any more, in particular in case soft real time latency is required. Furthermore, such combiners should be able to deal with non-responsive devices or devices sending obviously faulty data or with delay.

Modern distributed control systems support communication using data published in models. EPICS 7 calls these "normative types" [6–9]. Thus an orbit model can be built as an aggregate of the beam position monitor sub-models. Using these data-models changes of the sub-model could then propagate through the whole on-line or even off-line analysis chain.

Similarly, data with time tags could be combined, filtered, analysed and delivered to the data-sink. These data analysis pipelines are typically based on streams. Although these could handle delays and jitters, their end users can depend on a swift delivery: Control loops require a timely flow of data to achieve their optimisation tasks within a required frequency band. Training reinforcement learning agents typically needs many steps (1 000 – 10 000 steps), thus fast swift data combination can make tasks feasible, which can be out of reach if data combination takes too long.

This article is based on the authors' experience combining data from their BPM's into an orbit object using an in-house developed application. We have the impression that it could be worthwhile to consider that certain aspects of this task are common tasks which would be best provided within the distributed control system environment.

## HANDLING DATA MODEL DATA: AN OVERVIEW

Distributed control systems of today allow publishing data in a structured manner: e.g. as normative types for EPICS. Then these structured data or data models are published as a consistent set by a single node.

We think that it would then be helpful if we could work on this data in a straightforward manner to

- **A**ggregate single data models to a consistent larger one
- **T**ransform contents of the data models
- **M**odify the data model e.g. enriching the data model with metadata or remove sub parts of the model that could be confusing further down the processing line.

## USE CASE: ORBIT OBJECT

We see creating an orbit object as an example of aggregation. We build a combined data model from single submodels ensuring that these are consistent.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG014

# PERFORMANCE AND RELIABILITY IMPROVEMENT OF A PYTHON-BASED EPICS IOC BY SWITCHING TO pyDevSup

G. S. Reis*, A. C. S. Oliveira, É. N. Rolim †
Brazilian Synchrotron Light Laboratory (LNLS/CNPEM), Campinas, Brazil

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The power supplies used for Fast Orbit Feedback (FOFB) correctors in SIRIUS expose only electrical current values, making it necessary to perform conversions to and from beam kick values. To take advantage of the canonical Python implementation of this conversion, a separate Input/Output Controller (IOC) was developed using pyEPICS and PCASPy. This technology stack imposed some limitations, making it necessary to limit the update rate, and, even then, requiring one independent instance of the IOC per ring sector (20 in total) to avoid process variable (PV) timeouts and disconnects; disconnection events when one of the power supplies was down also had cascading issues with reconnection and memory corruption. This motivated us to pursue more modern alternatives for integrating Python code into an IOC, specifically one that could take advantage of the Channel Access (CA) integration already present in Experimental Physics and Industrial Control Systems (EPICS) databases, avoiding any of the bridges between CA and Python. We evaluated the pythonSoftIOC project and the pyDevice and pyDevSup support modules, which we present in this work. We settled on pyDevSup due to the development experience it provided. This work also presents benchmarks comparing performance gains with the new IOC and aims to explore the architecture differences that enabled them.

## INTRODUCTION

SIRIUS is a 3.0 GeV synchrotron light source built by the Brazilian Synchrotron Light Laboratory (LNLS) [1]. As part of the effort to reach SIRIUS's orbit stability specification, a Fast Orbit Feedback (FOFB) system [2] was developed in-house.

Even though FOFB modeling and control relies on beam kick values, the power supplies used by the FOFB correctors in SIRIUS expose only electrical current values through Experimental Physics and Industrial Control Systems (EPICS) process variables (PVs), making it necessary to perform conversions to and from beam kick values. This conversion was implemented and added to the siriuspy [3] Python library, and an Input/Output Controller (IOC) using pyEPICS [4] and PCASPy [5] obtains the current measurements through Channel Access (CA) [6] and exposes the converted kick values through PVs of its own. Each ring sector has 8 FOFB correctors (horizontal and vertical planes for each of the 4 subsections). For each corrector, we have one setpoint PV and 4 readback current PVs (`Current-RB`, `CurrentRef-Mon`, `FOFBAcc-Mon`, `Current-Mon`). Each one of these electrical current PVs has a kick counterpart. Due to performance limitations related to the use of pyEPICS for PV connection, it was necessary to impose a 10 Hz update limit in the implementation in order to avoid CPU overload or timeouts; even then, the CPU usage is still considerable, and one independent instance of the IOC was deployed per ring sector to avoid PV timeouts and disconnects. Disconnection events when one of the power supplies was down also had cascading issues with reconnection taking a long time and occasional memory corruption. The PCASPy dependency is another consideration, since it uses the Portable Channel Access Server (PCAS) module that implements its own CA server and is no longer actively maintained.

These issues motivated us to look for a new way to implement this conversion IOC, bridging the existing Python conversion library with an EPICS server. Our experience building traditional IOCs that connect to external PVs over CA, which use the database implementation from EPICS Base, has been very positive, with great performance, short reconnection times and stable connections. That motivated us to try and take advantage of EPICS database links for this IOC as well. Our research led us to three candidates to fulfill these requirements: pythonSoftIOC [7], pyDevice [8], and pyDevSup [9].

## EVALUATING THE CANDIDATES

EPICS support modules provide a mechanism to extend IOC functionality by connecting database records to external code. For Python integration, these modules typically work by bridging the `INP` and `OUT` fields of EPICS records to Python functions or methods. The `INP` field specifies the input source for a record (where it gets its data), while the `OUT` field specifies where processed data should be sent. Python support modules intercept these fields with special device types that route data between the EPICS database and Python code.

Each of the three modules we evaluated implements this bridging differently, with distinct architectural approaches:

### *pyDevice*

pyDevice is an EPICS device support module that allows Python code execution within the IOC shell. It enables variables and functions defined in Python to be accessed directly from EPICS database records by adding Python snippets to their `INP/OUT` fields.

Based on our experience, pyDevice was more suitable for those who want to use Python code in the startup script (`st.cmd` file) or call simple Python functions in records' `INP/OUT` fields (e.g., to send PV data through a socket),

* gustavo.reis@lnls.br
† erico.rolim@lnls.br

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG015

# UPGRADING THE ATLAS TUNE ARCHIVING SYSTEM*

K. Bunnell†, M. Torres, A. Ramaswamy, B. Blomberg, C. Dickerson,
G. Dunn, D. Novak, D. Stanton
Argonne National Laboratory, Lemont, IL, USA

*Abstract*

The Argonne Tandem Linear Accelerating System (ATLAS) is a U.S. DOE national user facility that delivers stable and radioactive ion beams from hydrogen to uranium for low-energy nuclear physics research [1]. Operators routinely expedite setup by restoring previously optimized machine parameter sets ("tunes"). The legacy tune archiving system, implemented in Corel Paradox (1999), has become a maintenance and operational bottleneck due to recurrent table corruption, single-user access, limited integration, and proprietary language.

We present the ATLAS Time Machine (ATM), a modern replacement comprised of PySide6 for the UI, FastAPI for backend services, MariaDB for experiment metadata, and InfluxDB v2 for time-series device data. ATM supports multi-user access, direct integration with the ATLAS control system, and automated beamline-aware data collection based on a dynamically generated beam path view. Initial results from beta operations indicate improved reliability, streamline operator workflows, and more convenient operator access. We conclude with lessons learned and a roadmap toward full production deployment.

## INTRODUCTION

Experiments at ATLAS typically start every 3-5 days with beam being created at one of three ion sources and being delivered to one of nine target areas. Beams are on target 24 hours a day, seven days a week excluding tuning and down time. At the start of an experiment, the accelerator requires operators to tune the accelerator for the specific ion species used. During this process operators must find the optimized value for all beam influencing devices to deliver beam to the target efficiently. The faster the beam can be tuned onto the target, the more hours that are able to go towards the experiment.

One way to improve the speed of the tuning process is to load values from a previous experiment that is similar to the current one and attempt to improve the tune from there. The tunes still require minor tweaks to improve transmission which takes 24 hours on average.

The current solution to this is the Tune Archiving System (TAS) [2]. TAS utilizes legacy software, is difficult to maintain, and has limitations due to the software and programming practices used. Due to these limitations, a replacement for the old TAS is being developed that uses Python to make the solution modern, simpler to maintain, and lower the learning curve.

* This work was supported by the U.S. Department of Energy, Office of Nuclear Physics, under Contract No. DE-AC02-06CH11357. This research used resources of ANL's ATLAS facility, which is a DOE Office of Science User Facility.
† kbunnell@anl.gov

## ATLAS CONTROL SYSTEM

Real-time control of the accelerator is actively being upgraded from a legacy system to an upgraded system both using Vista Controls VSystem software as shown in Fig. 1. The legacy system consists of a Digital Equipment Corporation AlphaServer that relies on a CAMAC (Computer Automated Measurement and Control) serial highway to control and monitor devices. This system supports an Oracle Relational Database (RDB) that allows systems outside of VSystem to read control system data including TAS and websites [3].

The upgraded modern hardware consists of 14 Linux servers. Figure 1 shows that each computer includes VSystem and the drivers required to monitor and control devices using Ethernet [4], USB, or VME. ATM makes use of MariaDB to hold many VSystem channels commonly used amongst the programs at ATLAS. Device data is stored in InfluxDB so that historical trends can be displayed in Grafana.

This upgrade approach allows for a slow migration of beamline devices and software systems to move from the old hardware to more modern hardware.

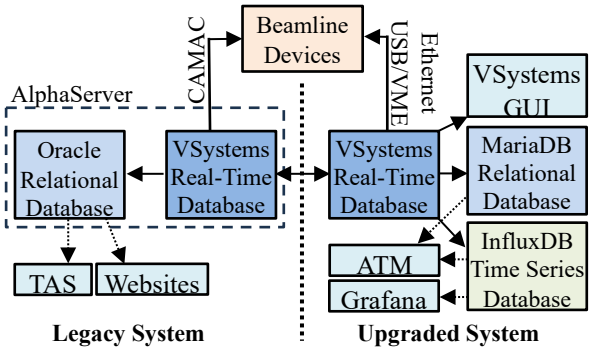



Figure 1: ATLAS Control System configuration. Upgraded system summarizes network of 14 Linux-based computers.

The MariaDB contains several tables that include data about each device on a small section of beamline. At the start of each experiment, the tables determined to be on the beampath are automatically combined into an SQL view that describes the entire beamline from source to target. This view simplifies the code that determines the devices on the beamline and their order. The beampath table encodes a globally unique, order-preserving device id and several other tags to enable efficient sorting and filtering.

## MOTIVATION

TAS was developed in 1999 and uses Corel's Paradox, a relational database management software which includes a UI designer and integrated programming language.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG016

# STATUS OF DEVELOPMENT AND APPLICATION OF THE PYAPAS AT HEPS

X. Lu, Y. Zhao, Y. Zhang, H. Ji, Y. Jiao, X. Huang
N. Li, C. Meng, Y. Peng, H. Xu, D. Ji, Y. Y. Wei
Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China

*Abstract*

To meet the requirements of beam commissioning at the High Energy Photon Source (HEPS) [1], China's first fourth-generation high-energy synchrotron light source, a new high-level application (HLA) framework named Pyapas was developed in pure Python. Designed for flexibility and maintainability, Pyapas serves as the foundation for all HLAs at HEPS, supporting tasks such as orbit correction, optics measurement, and machine modeling. Since early 2023, Pyapas-based HLAs have been successfully applied during the commissioning of the Linac, booster, and storage ring, contributing to key milestones including first light in October 2024. Recently, the storage ring has achieved a beam current of 100 mA with an emittance of 100 pm, marking another significant step in the commissioning progress.

## INTRODUCTION

The High Energy Photon Source (HEPS) is a 6 GeV, 1.3 km, fourth-generation light source [1]. It began construction in Beijing, China in mid-2019 and is expected to become one of the world's smallest emittance light sources after completion. The Linac of HEPS began beam commissioning in March 9 of this year, and subsequent beam commissioning work on the booster and storage ring will be carried out in succession. In order to meet the beam commissioning needs of HEPS, we began developing high-level applications (HLAs) in early 2021 and a new framework was designed for the development of HLAs.

As a fourth-generation light source, HEPS adopts a compact multi-bend achromat lattice design for the storage ring and uses a large number of combined magnets [2, 3]. The cross-talk effect between magnetic fields is significant, and the number of magnets has increased by an order of magnitude compared to the third-generation light sources, which poses great challenges and new requirements for controlling beam parameters. Existing HLA development frameworks are mostly monolithic and not easily scalable [4, 5], which cannot meet the needs of HEPS' complex parameter tuning. To address this issue, we have designed a new HLA framework named *Pyapas* [6]. It adopts a modular design philosophy and increases overall scalability. A dual-layer physical model module has been designed to meet the replaceability of online calculation models. In addition, the communication module, database module, and server module have all been specially designed to meet the needs of adjusting a large amount of parameters.

* Work supported by NSFC12505227
† luxh@ihep.ac.cn

Based on *Pyapas*, we have completed the development of HLAs for HEPS [7-9] and successfully applied them to beam commissioning, verifying the practicality and reliability of the applicatoins. This paper will briefly introduce the design of *Pyapas* and the progress of the development of HLAs for HEPS.

## FRAMEWORK DESIGN

At the core of Pyapas lies a dual-layer physical model module (see Fig. 1), which consists of a device mapping layer and a model calculation layer. The two layers are decoupled and linked through connectors, which only require the specification of details such as the calculation model class name, unit conversion rules, and parameter names to be invoked. The device mapping layer corresponds to the actual machine and is responsible for communication with it, while the model calculation layer performs online computations based on the data obtained from the device mapping layer [10]. This modular design allows different calculation models to be switched seamlessly, enabling flexible adaptation to various online computing scenarios.

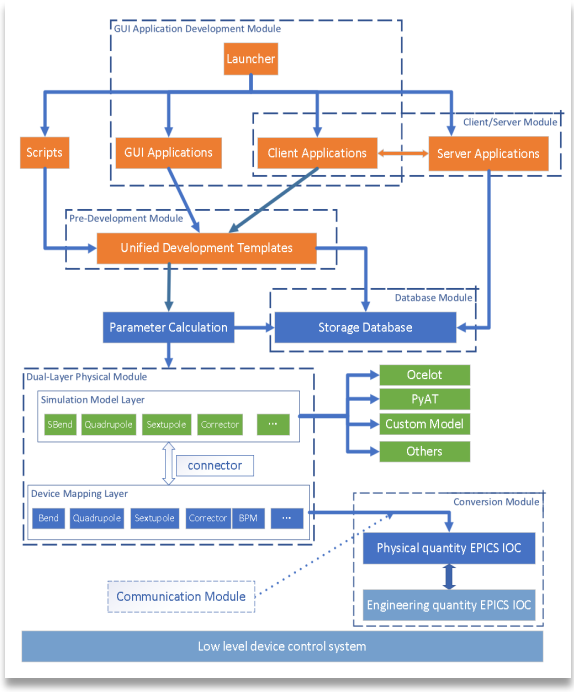



Figure 1: The framework of Pyapas.

In addition, to improve the scalability of the framework and reduce the development difficulty of developers, we adopt a modular design. Specifically, we analyze the functions involved in HLAs, and divide them into different modules for separate development. The modules are

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG017

# MODULAR SCIENTIFIC SCADA SUITE WITH SARDANA AND TAURUS – LATEST DEVELOPMENTS

O. Vallcorba*, J. Aguilar, R. Homs-Puron, E. Morales, M. Navarro (on-leave),
J. Ramos, J. Moldes, F. Becheri, S. Rubio, Z. Reszela, ALBA-CELLS, Barcelona, Spain
M. T. Núñez Pardo de Vera, DESY, Hamburg, Germany
B. Bertrand, J. Forsberg, V. Da Silva, M. Lindberg, MAX IV Laboratory, Lund, Sweden
M. Piekarski, SOLARIS, Krakow, Poland
W. Kitka, S2 Innovation, Krakow, Poland
A. Hoffstadt, ESO, Garching, Germany
T. Braun, byte physics, Berlin, Germany
D. Schick, M. Schneider, MBI-Berlin, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Sardana and Taurus are community-driven, open-source SCADA solutions that have been used for over a decade in scientific facilities, including synchrotrons (ALBA, DESY, MAX IV, SOLARIS) and laser laboratories (MBI-Berlin).

Taurus is a Python framework for building both graphical and command-line user interfaces that support multiple control systems or data sources. Sardana is an experiment orchestration tool that provides a high-level hardware abstraction and a sequence engine. It follows a client-server architecture built on top of the TANGO control system.

In the last two years, significant developments have been made in both projects. Sardana focused on enhancing continuous scans, introducing multiple synchronization descriptions to support passive elements (e.g. shutters) and detectors reporting at different rates. The configuration tool has also been extended, following the roadmap defined by the community. Taurus has seen substantial performance gains, particularly in GUI startup times, as part of an optimization effort that started nearly three years ago. Latest improvements take profit of new TANGO event subscription asynchronous modes. Continuous codebase modernization is underway, and support for Qt6 is planned for the July 2025 release.

This paper reviews recent advances in Sardana and Taurus and outlines the current development roadmap.

## SYSTEM OVERVIEW

Supervisory Control and Data Acquisition (SCADA) systems are essential in scientific facilities, providing the abstraction layers needed to integrate heterogeneous hardware, orchestrate experiments and offer usable interfaces.

Sardana [1, 2] and Taurus [3, 4] are Python-based, community-driven and open-source solutions currently used in production at several facilities. The main adopters of the Sardana–Taurus stack include ALBA, DESY, MAX IV, and SOLARIS synchrotrons and the MBI-Berlin and HI-JENA laser laboratories. Taurus is also used at the ESO observatory and, to a lesser extent, at ESRF, SOLEIL and ELETTRA synchrotrons. Together, Sardana and Taurus provide a modular and extensible SCADA suite covering the range from user interaction to low-level hardware orchestration:

- User interfaces: Taurus provides tools to create graphical interfaces (GUIs) that expose experimental setups, live data and control elements across multiple control systems through a unified abstraction model. Sardana adds extensions to Taurus to enhance the interaction with its specific components and in addition, it offers *Spock*, an IPython-based command-line interface (CLI).
- Experiment orchestration: Sardana includes a sequencing engine and high-level abstractions to coordinate devices. The *MacroServer* allows complex experimental procedures to be expressed as *macros*, while the *Device Pool* defines generic interfaces for motors, detectors and other equipment. This layer synchronizes elements and manages scans and data acquisition.
- Control system: Sardana is built on top of TANGO [5], a distributed client–server middleware widely used in large scientific facilities. TANGO provides standardized access to hardware and software devices and acts as Sardana's communication backbone.
- Hardware: Physical devices such as motors, detectors, shutters and sample environments exposed through the control system (e.g. TANGO devices).

This paper updates the ICALEPCS 2023 status report [6] and is organised from higher-level user interfaces (Taurus) down to orchestration layers (Sardana), followed by the roadmap agreed by the community.

## USER INTERFACES

Taurus uses Qt [7] to build GUIs, supporting both PySide [8] and PyQt [9] bindings via the `taurus.external.qt` module. For plotting live and archive data, the `taurus_pyqtgraph` extension [10] provides 1D widgets, while 2D support is offered through `guiqwt`-based widgets [11].

* ovallcorba@cells.es

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG018

# HL-LHC INNER TRIPLET STRING CONTROLS AND SOFTWARE ARCHITECTURE

S. Blanchard†, A. Antoine, M. Bajko, D. Bozzini, M.-A. Galilée, J.-C. Garnier, R. Gorbonosov, L. Herrero Alvarez, A. Kostopoulos, A. Mnich, T. Podzorny, B. Schofield, D. Wollmann, S. Yammine, M. Zerlauth
CERN, Geneva, Switzerland

*Abstract*

The High Luminosity-Large Hadron Collider (HL-LHC) project at CERN aims to increase the integrated luminosity of the Large Hadron Collider (LHC) by an order of magnitude compared to its initial design. As an important intermediate milestone of the HL-LHC project, the scope of the Inner Triplet (IT) String test facility is to study and validate the collective behaviour of the complete final focusing region of the machine upgrade, and to represent the various operation modes and the controls environment. As for the HL-LHC, the IT String operation requires a wide-ranging set of control systems and software for magnet powering, magnet protection, cryogenics, insulation vacuum, and the full remote alignment system.

An overview of the control systems and their interfaces is presented with a particular focus on the software layers essential for the powering and magnet protection tests during the IT String validation program. Ensuring integration of the new HL-LHC device types and their operational readiness requires close collaboration between development teams, equipment owners and the IT String operation team which is validated by dedicated Dry Run tests. These tests aim at validating the functionalities of new device types within the control and software applications, with the goal of achieving a smooth transition to the magnet powering phase. The IT String facility presents a unique opportunity to validate all control and software layers ahead of the HL-LHC hardware commissioning (HWC) within the LHC complex and their operation in the High Luminosity era.

## INTRODUCTION

The HL-LHC IT String facility [1] aims to validate the superconducting magnet circuits under conditions as close as possible to their operation in the LHC High Luminosity era. The IT String aims to validate the collective behaviour of key components developed for the HL-LHC, such as superconducting magnets based on $Nb_3Sn$ and NbTi technologies, a superconducting link based on $MgB_2$ superconducting cables, novel quench detection and circuit protection systems, high current powering equipment and a fully remote alignment system.

The IT String facility features a total of 17 superconducting magnet circuits. The magnets string is composed of the focusing quadrupoles circuit, orbit correctors circuits, high-order correctors circuits, and the separation and recombination dipoles circuit. A circuit of a superconducting magnets includes all the powering and protection equipment electrically connected to the coils.

Different systems are required to operate the IT String magnets:

- Cryogenics system provides liquid helium and pumping lines for the cooling of the magnets at 1.9 K.
- Vacuum system for the thermal insulation of the superconducting equipment.
- Powering equipment which includes 20 power converters (PC) rated from 120 A to 18 kA DC.
- The Quench Protection System (QPS) ensures the integrity of the superconducting magnet circuits following the loss of superconducting conditions in any part of the circuit. It includes new detection electronics, so-called universal Quench Detection System (QDS) [2], energy extractions systems (EE), and other protection mechanisms.
- Powering Interlock Controller (PIC) gives or removes the permit to start powering of the circuits and triggers power aborts according to a set of interlock inputs received.
- Full Remote Alignment System (FRAS) [3] includes alignment sensors to continuously determine the positions of the magnets during the different operational phases. The Cryo-Assemblies are supported by motorised jacks to perform their remote adjustment.

The controls and software layers are not specific to the IT String but they are based on frameworks already in use at the LHC or developed for the High-Luminosity LHC era.

## ARCHITECTURE

As shown in Fig. 1, the architecture of IT String controls and software can be represented with three layers:

- The first is the Front-End layer, composed of either Programmable Logical Controllers (PLC) or Front-End Computers (FEC).
- The second layer is the Systems application layer either based on a Supervisory Control and Data Acquisition (SCADA) application or Web and Desktop applications dedicated to the Experts and used by the operation team for monitoring purposes.
- The third layer is the Operation Layer composed of applications and services to operate the superconducting magnet circuits, manage the power permit, start & control the powering tests, identify the systems with issue, proceed with first diagnostics and event analysis.

† sebastien.blanchard@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMG019

# IMPLEMENTATION AND SCALABILITY ANALYSIS OF TSPP FOR VACUUM FRAMEWORK

R. Ferreira, A. Rocha, A. Giannoulas, L. Cantu , CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

SCADA (Supervisory Control and Data Acquisition) systems traditionally acquire data from PLCs through polling. The Time Stamped Push Protocol (TSPP), on the other hand, enables a PLC to timestamp and push data to the SCADA at its own discretion. The Vacuum Control Systems for CERN accelerators are primarily built on a dedicated Vacuum Framework, which relies on polling and is therefore subject to its limitations. Implementing TSPP would thus be an important improvement.

TSPP needs software on the PLC – a Data Manager - to determine what data to push, when to push it, and how to package it into the correct format. Due to its particular data model, implementing TSPP for the Vacuum Framework required the development of a dedicated Data Manager. Additionally, while most current systems with TSPP have a single PLC per SCADA instance, Vacuum Framework applications often involve hundreds. Given that no data was available on the impact that large numbers of PLCs pushing data to a SCADA system might have, extensive testing was required. In particular, the relationship between server load and the effective rate of received values was studied to assess performance at scale.

This paper details the implementation of TSPP for the Vacuum Framework, its Data Manager design, and the testing carried out to validate the protocol and assess its performance limits in order to ensure a smooth deployment.

## INTRODUCTION

SCADA systems traditionally acquire data from PLCs by polling them at regular intervals. This approach, while simple and widespread, has important drawbacks:

- The acquisition period is limited to that of the SCADA polling time, which is typically much higher than the PLC cycle time.
- Quick events – i.e. pressure spikes or fast discrete flips – might be missed entirely.
- It can be impossible to establish cause-and-effect relations when analyzing fast and complex chains of events.

The Time Stamped Push Protocol (TSPP) addresses all these limitations by allowing a PLC to timestamp data at the exact moment of acquisition and push it to the SCADA at its own discretion. This means that with TSPP data acquisition frequency is limited only by the PLC cycle time, and not by the SCADA polling time. It also means, however, that additional software is required on the PLC – a Data Manager – to decide which data to push, when to push it, how to filter it, and how to pack it in the standard format defined by the TSPP specification [1].

### *UNICOS and the Vacuum Framework*

TSPP was jointly developed by CERN and Siemens (the manufacturer of the WinCC-OA SCADA) for UNICOS, CERN's general purpose Industrial Control Framework [2, 3]**.** The aforementioned Data Manager is thus built into the UNICOS PLC baseline [4], whilst the SCADA driver required to process TSPP frames is part of WinCC-OA.

Most Vacuum Control Systems, however, resort to the Vacuum Framework, which was purpose built for Vacuum applications and is used in all major accelerators at CERN. It essentially consists of a set of databases, web applications, and PLC/SCADA code libraries [5-7]. The Vacuum and UNICOS frameworks differ in several aspects, two of which are particularly relevant here. Firstly, the data model of the Vacuum Framework is very different from that of UNICOS, requiring the development of a dedicated Data Manager. Secondly, while most UNICOS applications have a single PLC per SCADA application, Vacuum applications often have several hundreds. As a result, extensive testing was required to assess how TSPP scales under these conditions and to evaluate the potential strain that this might put on the SCADA servers. This paper addresses both subjects.

## VACUUM FRAMEWORK DATA MODEL

Devices (e.g. valves, gauges) in the Vacuum Framework are represented by Control Types. A Control Type includes a PLC Function Block, a Datapoint Type of a standard structure on the SCADA, the C++ libraries and CTRL scripts that implement the device's behaviour within the SCADA, and a set of widgets and panels with which to interact with the equipment. Instances of Control Types are created on Machine Databases (e.g. the *LHC Machine DB*) and a dedicated application, the *Vacuum Controls Configurator Exporter,* uses this data to generate source files which can be imported into the PLCs and SCADA. On the PLC, each device is imported as an Instance DB of its respective Function Block. On the SCADA, the equipment is represented as a Datapoint of the appropriate type.

Each PLC contains one or more Read Registers (RD_DB), large data blocks consisting of byte arrays. Every Control Type has an associated Read Register Definition, specifying how its data is mapped in memory.

As an example, Fig. 1 illustrates how the data for an instance of VG_A_RO (a read-only analog gauge) is stored in RD_DB, in the memory area between *n* and *n+8*. The first 4 bytes (RR1) form a DWORD containing several status values (On, Off, Under Range, Over Range), and other discrete data (Error and Warning Codes). At offset *n+4* is *ObjStat*, a BYTE that contains a Global Status value. The gauge pressure is stored in *PR*, a REAL value that occupies the last 4 bytes.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR001

# EXPLORING AI-BASED MODELS IN ACCELERATORS: A CASE STUDY OF THE SOLARIS SYNCHROTRON

M. Piekarski *, J. Biernat, M. Floras, M. Mleczko, M. Wrobel,
National Synchrotron Radiation Centre, SOLARIS, Jagiellonian University, Kraków, Poland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The National Synchrotron Radiation Center SOLARIS, third generation light source, is the only synchrotron located in Central-Eastern Europe, in Poland. The SOLARIS Center, with seven fully operational beamlines, serves as a hub for research across a diverse range of disciplines. The most important aspect of such research infrastructure is to provide stable working conditions for the users, operators and the conducted projects. Due to its unique properties, problem complexities, and challenges that require advanced approaches, the problem of anomaly detection and automatic analysis of signals for the beam stability assessment is still a huge challenge that has not been fully developed. To address this problem, different AI-based projects are under discussion and development, i.e. automatic analysis of diagnostic signals on the example of transverse beam profiles or beam position FFT windows classification. The best proposed solution, based on the InceptionV3 architecture, can assess beam quality automatically, based solely on the image itself with

## INTRODUCTION

### *SOLARIS Synchrotron*

SOLARIS is a third generation light source (shown in Fig. 1) operating at the Jagiellonian University in Krakow, Poland. This advanced and complex scientific infrastructure offers new highly innovative research opportunities for areas including physics, medicine and nanotechnology. Currently at SOLARIS seven experimental beamlines offering various techniques, e.g.: photoemission electron microscopy, X-ray absorption spectroscopy, ultra angle-resolved photoemission spectroscopy or multi-scale X-ray and multimodal imaging, are available to the scientific community whereas another three are already at advanced level of construction or commissioning. Moreover, SOLARIS is also a National Cryo-EM Centre, with two latest generation cryo-electron microscopes enabling life science researchers to unravel life at the molecular level [1].

### *AI in Accelerators*

Artificial intelligence (AI) is a field that is developing very dynamically. It is an important part of research in industry, where it is used from production supervision in factories through monitoring the condition of infrastructure, the analysis of network traffic, up to the national research infrastructures and the science itself. The progress of knowledge and techniques in this area has significantly accelerated the development of AI, which is currently experiencing a real golden age. The proposed solutions in this field include novel machine learning (ML) algorithms, simple neural networks (NN), and sophisticated deep learning (DL) architectures. Therefore integration of automation or semi-automation AI-based tools for maintenance, early fault detection and daily tasks is a natural process that can significantly enhance accelerator operation.



Figure 1: SOLARIS facility in Krakow (photo: Aleksander Koczur).

Artificial intelligence based methods, in particular machine and deep learning, are commonly used in accelerators. Examples include the UNet architecture for beam loss (fault) origin estimation [2], the autoencoder architecture for accelerator electronic circuit failure prediction [3], as well as the incorporation of AI into daily control system operation [4,5] (mainly data mining and automatic control of simple optic elements). Mercato et al. [6, 7] compared the LTSM-based model and various machine learning models to analyze and predict runtime faults in the radio frequency (RF) subsystem that controls superconductive RF cavities.

## SOLARIS AI DEVELOPMENTS

### *Beam Profile for Quality Assessment*

Stable operation is a key factor for all large research infrastructures. Therefore, a thorough analysis of diagnostic signals and their assessment is crucial. In this project, we reviewed accelerator subsystems for data utilization in deep learning and beam quality assessment. Later, we proposed the benchmark for Pinhole [8] image classification including data preprocessing, model implementation, training process, hyperparameter selection as well as testing phase. A database of more than 1 million transverse beam profiles was prepared from scratch and each image was binary labeled in a semi-automatic way on the basis of the emittance value, one of the key parameters of the electron beam. Profile within

* michal.piekarski@uj.edu.pl

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR002

# AI AND ML INTEGRATION FOR BEAMLINE OPTIMIZATION AND VIRTUAL ASSISTANCE AT THE SOLARIS SYNCHROTRON

M. Szczepanik†, M. Piekarski, M. Fałowski
National Synchrotron Radiation Centre SOLARIS, Kraków, Poland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

This paper presents two complementary initiatives currently under consideration at the SOLARIS synchrotron facility, aimed at enhancing beamline operations and user support through artificial intelligence (AI) and machine learning (ML). The first initiative explores the potential of ML-based mirror alignment and beam stabilization, while the second investigates the development of a virtual assistant powered by large language models (LLMs). Both approaches are in the conceptual phase and leverage existing control infrastructure and modern AI frameworks to improve experimental efficiency, reproducibility, and accessibility.

## INTRODUCTION

The SOLARIS synchrotron, located in Kraków, Poland, is a national research infrastructure dedicated to advanced photon science. As synchrotron facilities evolve toward higher brilliance and more complex experimental configurations, the need for intelligent, adaptive control systems becomes increasingly critical. This paper presents two technical solutions for integrating AI and ML into SOLARIS operations: automated beamline optimization and a virtual assistant for scientific and operational support.

## MACHINE LEARNING-BASED BEAM OPTIMIZATION

The beamlines at SOLARIS are equipped with multiple mirrors of varying geometries and monochromators, forming a complex optical path that requires precise alignment. The control system is based on the Tango Controls framework, with Sardana used for high-level experiment orchestration. All control software runs on AlmaLinux 9, ensuring compatibility with modern scientific computing libraries and ML frameworks. The following ML-based solution is currently being evaluated and has not yet been implemented at SOLARIS.

The ML-based solution addresses two distinct but interrelated optimization goals. The first is the enhancement of beam quality, defined by maximum intensity, tight spatial focus, and optimal energy resolution for a given photon energy. This is achieved by analyzing signals from a soft X-ray photodiode, which serves as a proxy for beam intensity and spectral purity. The second goal is the stabilization of the beam position on the sample, ensuring that the beam consistently hits the same location regardless of energy changes. This is monitored using the induced current on the exit slit blades, which reflects the beam's spatial alignment.

Due to the physical constraints of the beamline, simultaneous acquisition of photodiode and slit current signals is not feasible. Therefore, the system will employ two separate ML models. The first model will be trained on data collected during photodiode-based measurements and will predict mirror and monochromator adjustments to maximize beam quality. The second model will use slit current data to maintain beam position, compensating for energy-dependent shifts in the optical path. Both models will be updated online using data collected during routine measurements, allowing them to adapt to changing beamline conditions without manual intervention.

Recent research supports the feasibility and effectiveness of such approaches. Morris et al. [1] demonstrated a general Bayesian optimization framework for autonomous beamline alignment, capable of handling high-dimensional, nonlinear parameter spaces with minimal prior information. Similarly, Morris et al. [2] applied online learning techniques to optimize beamline parameters at NSLS-II, showing that real-time ML can outperform traditional alignment methods in both speed and precision. Rebuffi et al. [3] developed an AI-driven system for autonomous alignment of nanofocusing mirror systems, highlighting the role of digital twins and wavefront sensing in achieving sub-micron precision. Xie et al. [4] proposed a deep learning method for estimating alignment errors in Kirkpatrick–Baez mirror systems, demonstrating that convolutional neural networks can extract meaningful features from speckle-modulated X-ray images. Finally, Yazak [5] reported on the integration of AI agents at NSLS-II for real-time anomaly detection and experiment monitoring, underscoring the broader trend toward autonomous synchrotron operations.

The implementation of machine learning for beam optimization at the SOLARIS synchrotron is supposed to integrate seamlessly with the existing control infrastructure, which is based on Tango Controls and Sardana, running on AlmaLinux 9. The system will be developed in Python 3.9, ensuring compatibility with both the control framework and modern machine learning libraries such as scikit-learn, PyTorch, and NumPy.

The proposed architecture consists of two independent machine learning models, each addressing a distinct aspect of beam optimization. The first model is responsible for maximizing beam quality, defined by intensity, focus, and energy resolution. It operates on data acquired from a soft X-ray photodiode and is trained to predict optimal mirror and monochromator settings for a given photon energy. The second model ensures spatial stability of the beam on the sample, compensating for energy-dependent shifts in the optical path. It uses data from the induced current on the exit slit blades to maintain consistent beam positioning.

Due to the physical constraints of the beamline, simultaneous acquisition of photodiode and slit current signals is

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR003

# PLANS AND STRATEGY FOR EDGE AI/ML AT THE ELECTRON-ION COLLIDER AT BROOKHAVEN NATIONAL LABORATORY*

L. K. Nguyen†, P. Bachek, J. Jamilkowski, K. Kulmatycski
Brookhaven National Laboratory, Upton, NY, USA

*Abstract*

Scheduled to begin Operations in 2035, the Electron-Ion Collider (EIC) is being built at Brookhaven National Laboratory (BNL) and will be the only operating particle collider in the United States. It may also be the only large collider built in the world in the next 20-30 years, during the "Age of Artificial Intelligence (AI)". Recognizing the potential for AI and machine learning (ML) to enhance operations and create more research opportunities, the EIC is being envisioned and designed as a large-scale AI-ready state-of-the-art facility. Specifically, it will support three core areas of AI/ML capabilities, referred to as Edge, End-to-End, and Bottom-Up. Edge capabilities are intended to address what are expected to be some of the most demanding AI/ML applications in the world in terms of timescales by anticipating the infrastructure, hardware, and local compute resources needed for success. At the same time, considerable care must be taken to ensure that these capabilities are manifested in an efficient, safe, and secure Controls ecosystem and Operations environment. We report on our plans and strategy for high-performance edge AI/ML at the EIC.

## INTRODUCTION

The Electron-Ion Collider (EIC) is being built at Brookhaven National Laboratory (BNL) in New York State in partnership with the Thomas Jefferson National Accelerator Facility (JLab) and other domestic and international partners. Construction is scheduled to begin in 2026, following the end of the Relativistic Heavy Ion Collider (RHIC) program, and EIC Operations is scheduled to begin in 2035. The EIC will be the only operating particle collider in the United States and will collide high-energy polarized electron beams with polarized proton or ion beams (hadron beams) to study the strong nuclear force and hence the fundamental structure of visible matter with unprecedented capability. A comprehensive case for the science at the EIC can be found in Ref. [1], while the current status of EIC design and construction can be found in Ref. [2].

In its most recently released Long Range Plan—used to advise both the Department of Energy (DOE) and National Science Foundation (NSF) on research priorities in nuclear science over the next decade—the U.S. Nuclear Science Advisory Committee (NSAC) explicitly identified "the expeditious completion of the EIC as the highest priority for facility construction" and, in a separate identification of strategic opportunities, called for "a coordinated effort to integrate AI/ML technologies into the nuclear physics research programs [to] accelerate discoveries" [3]. These two priorities are thus converging at the EIC with the EIC being envisioned as a large-scale AI-ready state-of-the-art facility with "AI readiness" deliberately incorporated into its design. Much like how the EIC's high-level performance requirements flow down and dictate details about the machine, a unified vision of an "AI-ready" EIC with specific core AI/ML capabilities was created to guide decisions regarding infrastructure and architectures. Figure 1 shows a graphical representation of this unified vision next to a schematic layout of the EIC. The three core AI/ML

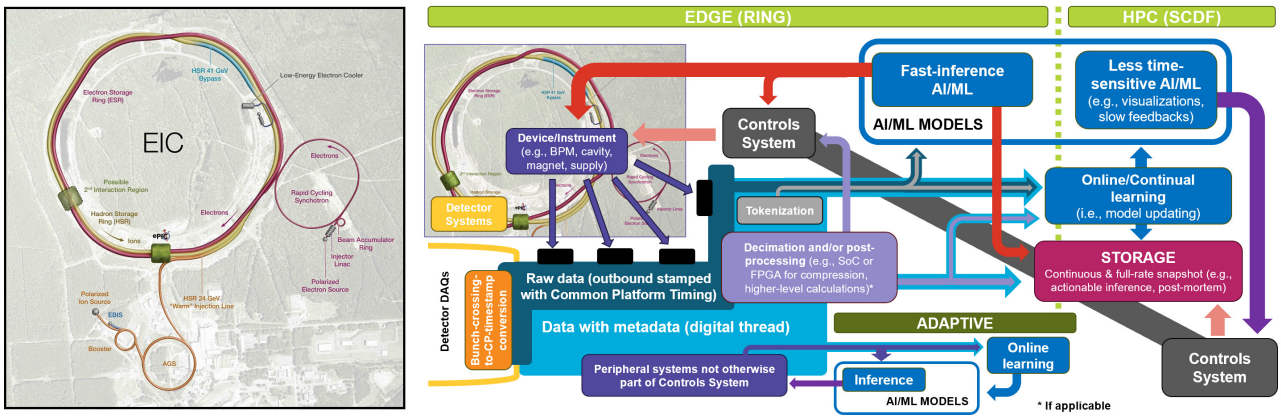



Figure 1: The future Electron-Ion Collider (EIC) at Brookhaven National Laboratory (BNL), and overview of the EIC's planned AI/ML ecosystem and infrastructure. The EIC will reuse the RHIC hadron injector chain (e.g., Booster and AGS), but because the latter will remain on a legacy controls system until upgrades later in the EIC program, the planned AI/ML infrastructure and capabilities herein only have the EIC Project for scope. SCDF (Scientific Computing and Data Facilities, formerly the Scientific Data and Computing Center) is BNL's high-performance computing (HPC) center.

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
† lnguyen@bnl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR007

# DESIGN OF AN INTELLIGENT INSPECTION SYSTEM FOR PARTICLE ACCELERATOR FACILITIES*

Y. C. He†,1,2, Y. L. Zhang [1,2], K. J. Xue [1,2], X. Wu [1,2],
X. H. Lu [1,2,3], L. Wang [1,2,3], M. T. Li [1,2,3], P. Zhu [1,2], W. L. Huang[1,2,3]
[1]Institute of High Energy Physics, Chinese Academy of Sciences (CAS), Beijing, China
[2]Spallation Neutron Source Science Center, Dongguan, China
[3]University of Chinese Academy of Sciences, Beijing, China

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

## *Abstract*

Aiming at the limitations of the traditional manual inspection method for particle accelerator facilities, an intelligent inspection system for particle accelerator facilities based on multi-modal sensors and artificial intelligence technology is provided. Multi-modal sensors can collect various types of data from the accelerator facilities, such as temperature, audio, images, and water leakage information, which provides comprehensive information for a thorough understanding of the equipment status. Artificial intelligence technology is capable of conducting in-depth analysis of massive amounts of data, uncovering fault patterns and rules. Through technologies such as fault modeling and data analysis, it can achieve early warning and diagnosis of faults. The intelligent inspection system for particle accelerator facilities effectively addresses the limitations of traditional inspection methods. It enables real-time monitoring of the accelerator facilities and rapid alarming of abnormalities, significantly improving work efficiency and the operational safety of the equipment. This system provides strong support for the stable operation of particle accelerator facilities.

## MANUSCRIPTS

This paper presents the design and implementation of an intelligent inspection system specifically tailored for particle accelerator facilities. Developed to address the inherent limitations of traditional manual inspection methods, the system provides a comprehensive solution for the real-time, accurate, and automated monitoring of accelerator facilities.

## *Background and Motivation*

Particle accelerators, such as the China Spallation Neutron Source (CSNS) [1, 2], are vital large-scale scientific facilities that underpin interdisciplinary research in materials science, life science, and nuclear physics. Traditional manual inspection of these accelerator facilities suffers from high labor costs (due to the need for extensive human resources at large, distributed sites), delayed fault detection (long 6-hour inspection cycles fail to capture real-time anomalies, leading to unplanned downtime), and low accuracy with inconsistent data (fault judgment depends on personnel experience, and manual data recording is error-prone and difficult to standardize). To address these limitations, this study develops an Artificial Intelligence (AI)-enabled, multi-modal intelligent inspection system. Integrating Internet of Things (IoT) sensing, edge computing, and cloud analysis technologies, the system achieves real-time monitoring of accelerator facilities, accurate fault warning, and automated data management, thereby overcoming the shortcomings of manual inspection and improving the efficiency and reliability of accelerator facility operation and maintenance.

## *System Architecture Overview*

The intelligent inspection system for particle accelerator facilities adopts a four-tier distributed structure with clear division of labor, ensuring efficient data flow and functional collaboration. The schematic diagram of the system architecture is shown in Fig. 1.

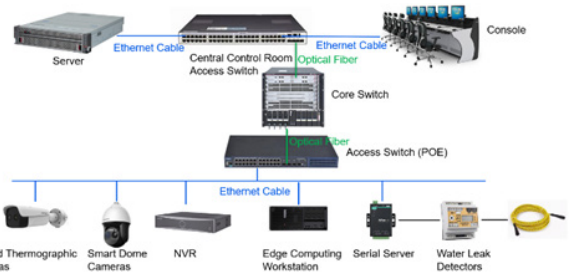



Figure 1: Schematic diagram of the system architecture.

- Frontend Devices Tier. Serves as the data acquisition layer, responsible for collecting real-time data related to accelerator facilities, including temperature, audio, images, and water leak information. This tier lays the foundation for subsequent data analysis by capturing comprehensive environmental and equipment state data.
- Edge Computing Tier. Undertakes local data processing tasks, including data preprocessing (to improve data quality), preliminary anomaly detection (to identify potential issues in real time), and Experimental Physics and Industrial Control System (EPICS) Process Variable (PV) conversion (to standardize data for transmission and analysis) [3, 4].
- Server Tier. Acts as the core computing and management layer, conducting in-depth AI analysis of collected data, building fault models, pushing real-time alarms, and archiving historical data. Its microservice architecture ensures high availability and scalability to adapt to the changing needs of accelerator facilities.

* Work supported by: The National Natural Science Foundation of China (No. 12275294)
† heyongcheng@ihep.ac.cn

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR009

# IMAGE PROCESSING WITH ML FOR AUTOMATED TUNING OF THE NASA SPACE RADIATION LABORATORY BEAM LINE*

L. Hajdu†, K. A. Brown, M. McCarthy, T. Olsen, W. Lin, Y. Gao
Collider-Accelerator Department, Brookhaven National Laboratory, Upton, NY, USA
E. Hamwi, G. Hoffstaetter de Torquat, Cornell University, Ithaca, NY, USA
R. Roussel, Z. Zhang, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

Research conducted at the NASA Space Radiation Laboratory (NSRL) seeks to increase the safety of space exploration. The NSRL uses beams of heavy ions extracted from the Booster synchrotron at Brookhaven National Laboratory (BNL) to simulate the high-energy cosmic rays found in space. To accomplish this, the source machines provide many potential beam species, ranging in atomic number (Z) from 1, hydrogen/protons, to 83, bismuth and previously have used up to thorium. To test large-area samples, beams can be shaped to the user's specifications from a small-format 1 cm radius circular beam up to 20 cm by 20 cm uniform-area rectangular beams. This requires a complex transfer line of 24 magnets, including 9 quadrupoles and 2 octupoles. Given the wide range of beam rigidity and size possibilities, operators tune the optics by hand while observing the beam profile on a phosphor screen imager. In this paper/project, we demonstrate successful tuning using a machine learning (ML) workflow. We capture the live beam image, then process and parameterize the beam to assess centroid, shape, tilt, edge thickness, and uniform area size. These parameters are fed to the Badger software stack to benefit from the accessible UIs, followed by an Xopt-based Bayesian optimization algorithm for real-time iterative tuning. The requirement to start from a real beam image (which can be very noisy) and quantify it, makes the workflow more complex than the standard ML cookie-cutter approach of reading from traditional beam instrumentation fed to an algorithm.

## OBJECTIVE

Our aim is to develop an operator interface based on machine learning which tunes the NSRL [1] beam line for their users' required experiments. A uniformly distributed beam with the characteristic shape should be built for any given ion species and beam size. The process would ideally converge well within 100 samples/iterations.

## BEAM-LINE DESCRIPTION

The transfer line optics [2] between the Booster synchrotron and the target room, illustrated in Fig. 1, expand and shape the beam into a uniform rectangle. The beam is slow-extracted from the Booster at the extraction septum and a small downward slope during the "flat-top" portion of the main magnet's current waveform.

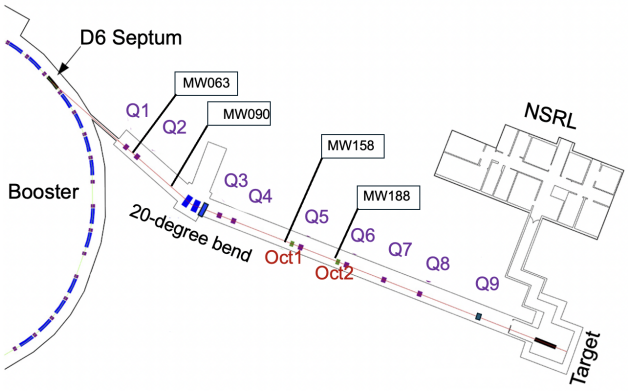



Figure 1: NSRL beam line.

Details of the transfer line optics are shown in Fig. 2 (the bend is omitted). Quadrupoles Q1 and Q2 bring dispersion to zero and make the rest of the line achromatic for beam shaping.

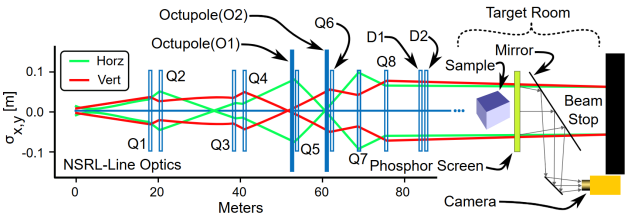



Figure 2: NSRL beam line optics.

Quad 4 widens the beam horizontally and flattens it vertically so that, when it passes through octupole O1, the horizontal edges are folded into the beam core to improve uniformity. The octupole strength increases with the fifth power of distance from the center. Q5 then widens the beam vertically and narrows it horizontally so that, when it passes through O2, the vertical Gaussian beam edge is redistributed to normalize the particle distribution. This results in a beam with a square profile.

Quads 7 and 8 adjust the size of the square beam, and dipoles D1 and D2 adjust the X and Y offset at the target. An example of transverse beam profiles at the target with and without octupoles on are shown in Fig. 3.

After some of the beam is absorbed by the sample, the remainder passes through a phosphor screen, where some

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy and by NASA (Contract No. T570X).
† LBHajdu@BNL.GOV

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR011

# LEVERAGING LOCAL LLM FOR LOG ANALYSIS: INTEGRATING OLLAMA INTO ELECTRONIC LOGGING SYSTEMS*

W. Fu†, S. Nemesure, J. Morris
Brookhaven National Laboratory, Upton, USAAbstract

Modern systems for electronic log keeping and trouble log management generate vast data sets of issues, events, solutions, and discussions. However, extracting actionable insights from this information remains a challenge without advanced analysis tools. This paper introduces an enhancement to two log programs—Electronic Log (elog) and Trouble Log (TroubleLog) - used in the RHIC control system at Brookhaven National Laboratory. The enhancement integrates Ollama, an open-source, locally deployed large language model (LLM) AI tool, to facilitate intelligent log data analysis. We present a framework that combines database queries with long chain split techniques, enabling users to generate period-based summaries (e.g., daily, weekly) and retrieve topic-specific information- such as issues and solutions - through natural language queries. By pipe lining data from databases and interfacing with Ollama's API, the system provides accurate, conversational responses while ensuring data privacy by avoiding external data sharing. The paper details implementation aspects, including SQL query optimization and prompt engineering, and evaluates performance using real-world log data sets. Results demonstrate improved usability and significant reductions in manual analysis time. This work investigates the potential of running local LLMs for log management tailored to accelerator control systems. We examined different methods for streaming data from our control system into Ollama AI engines, considering various source data types and focusing on a scalable and privacy-preserving solution.

## INTRODUCTION

The rapid progress of artificial intelligence (AI) has spurred the development of numerous open-source tools that democratize access to powerful language models, enabling researchers, developers, and organizations to deploy AI solutions locally. Open-source AI tools such as Ollama, Llama.cpp, LM Studio, GPT4All, and LocalAI etc. provide flexible, privacy-focused alternatives to proprietary cloud-based models like ChatGPT. Among these, Ollama [1] stands out as a robust and user-friendly platform for running large language models (LLMs) locally, offering seamless integration, extensive model support, and enhanced data privacy. By operating entirely on local hardware, Ollama eliminates reliance on external servers, ensuring that sensitive data remains secure—a critical advantage for many research and industrial companies. Its compatibility with a wide range of AI models, including Gemma3, Llama 4, and Mistral etc., combined with its support for GPU acceleration and customizable model configurations, makes it an ideal choice for both research and production environments. Furthermore, Ollama's command-line interface and HTTP API facilitate integration into diverse workflows, enhancing its appeal for developers seeking efficient, scalable AI solutions without incurring cloud-based inference costs. Research leveraging Ollama has demonstrated its versatility across various applications [2], particularly in privacy-sensitive and resource-constrained settings. For instance, studies have explored Ollama's ability to support local Retrieval-Augmented Generation (RAG) pipelines, enabling context-aware responses by integrating external data sources without cloud dependencies [3]. Additionally, research has highlighted Ollama's efficacy in tasks such as text summarization, chatbot development, and code generation, with models like Llama3 and Gemma3 showing strong performance in specialized domains [4, 5]. These studies underscore Ollama's capacity to deliver high-performance AI solutions while maintaining data sovereignty and cost efficiency.

To effectively utilize Ollama, several methods exist for pipelining data into its AI models, ensuring seamless integration into existing systems. The most common approach involves using Ollama's HTTP API, which allows developers to send input data and receive model outputs programmatically. This method supports batch processing and real-time interactions, making it suitable for applications like chatbots and automated text analysis. Another method is the command-line interface, which enables direct interaction with models using commands like ollama run model name, ideal for quick prototyping

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.
† Email: fu@bnl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR012

# ACCELERATOR DIGITAL TWIN DEVELOPMENT THROUGH SIMULATION MODELING AND MLOPS USING THE LUME ECOSYSTEM AT SLAC

G. Bhardwaj*, S. Miskovich, E. Williams, L. Lewis, C. Bisegni, J. Bellister,
C. Mayes, R. Roussel, A. Edelen, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Accelerator digital twins in conjunction with machine learning (ML) models enable real-time optimization and predictive control. This is accomplished by leveraging prior experiments, known parameters and real-time measurements. For successful calibration and control, robust infrastructure is crucial along with open-source tooling. This ensures accurate beam modeling and system integration. SLAC has developed the Lightsource Unified Modeling Environment (LUME) for large-scale modeling of X-ray free electron laser performance. The team is developing an ecosystem of tools towards start-to-end simulations for a much broader set of accelerators beyond light sources. This infrastructure is currently being developed and deployed at SLAC, alongside collaborators from other laboratories (especially LBNL and ISIS Muon and Neutrino Source). Here we provide an overview of the LUME ecosystem and example use cases. In these use cases, we have deployed ML models that ingest live accelerator data from the control system, run inference on the local HPC system at SLAC (the SLAC Shared Scientific Data Facility - S3DF), and write predictions back to the control system.

## INTRODUCTION

Particle accelerators are complex machines that require precise calibration and fine-tuning. An accelerator digital twin is a high-fidelity, continuously updated virtual representation of the physical machine while a surrogate model is a simplified, computationally inexpensive approximation of digital twin. For example, the LCLS Copper Injector Surrogate model [1] was trained on simulation data, but deployed to take real inputs from the machine and predict real beam properties including emittance. There is need for standards, wrappers and physics simulation models to develop high-fidelity digital twins.

SLAC is developing the Light Source Unified Modeling Environment (LUME) for efficient modeling of X-ray free electron laser (XFEL) performance [2]. LUME is a modular, scalable, toolkit with Python wrappers for physics simulations and live snapshots from the accelerator through Kafka. Alongside that, LUME has LUME-model for unified data structures and standardised model interfaces that support metadata-driven real-to-sim transformations and LUME-EPICS which is an API for serving LUME-model variables via EPICS.

* gopikab@slac.stanford.edu

While LUME provides an extensible modeling toolkit with standards, we need to streamline machine learning development lifecycle with Machine Learning Operations (MLOps). MLOps are a set of practices for automating ML development and retraining workflows. MLOps provides support for changing data dependencies and varying operations needs by making the entire ML lifecycle reproducible and accessible to all. To demonstrate a proof-of-concept of simulation modeling and MLOps, we deploy the LCLS Cu Injector Surrogate model which is a neural network model trained on simulation data with 16 input parameters. We will discuss this project in greater detail in the next section. Similar work is being done at the ISIS Muon and Neutrino Source in the UK where they have developed Polylithic [3] to automate the deployment of any model.

Another use-case for LUME is the Coincident Anomaly Detection work [4] at SLAC where the system continuously monitors RF station process variables (PVs), generates anomaly candidates using streaming metrics through the Kafka-to-EPICS (K2EG) gateway [5] and confirms them with a trained ML model. Here, LUME-model is used to update/generate the model files. We will present this work in a later section.

## DEPLOYING A SURROGATE MODEL OF THE LCLS INJECTOR

Linac Coherent Light Source (LCLS) is a linear accelerator-based X-ray free-electron laser at SLAC. In this work, the LCLS injector surrogate model is used as a proof of concept. This is a neural network (NN) based surrogate model that provides fast, non-invasive predictions of electron-beam properties. The model was trained on simulation data using 16 input parameters that define the shape and the properties of the beam at an observation screen downstream.

Since the screen is a destructive measurement, having a live surrogate model for accelerator operators to refer to while tuning machine parameters enhances the facility's ability to determine and improve the machine's status while minimizing beam interruptions to scientific users.

As shown in Fig. 1, The model was trained in [6] to predict beam properties at the profile screen using injector PVs for LCLS. As the model was trained with normalized data, input and output transformations have to be applied to use it on simulation data. Another layer of transformations is required for using it with EPICS data.

Figure 2 shows a simplified version of the machine section that is modeled by the surrogate model. Six of the most

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR013

# INTEGRATING CODAC IN ITER PLANT SIMULATOR

R. Balp†, R. Lange
ITER Organization, St. Paul Lez Durance, Cedex, France

*Abstract*

The use of Digital Control Systems (DCS) with process simulators for engineering purpose, control system validation, virtual commissioning, or operator training is increasingly demanded in large and complex industrial and scientific projects.

Coupling a DCS with a process simulator requires to support specific functionalities: ability to operate on a simulated time basis and to save/restore states to load different scenarios starting point, or to jump back in time, which is traditionally achieved by emulating or simulating the DCS.

The ITER Control, Data Access and Communication (CODAC) system uses EPICS at its core, which is not designed to operate with such constraints. In the frame of the ITER Plant Simulator project (IPSi), we leveraged advanced Linux features (libfaketime, namespaces, and CRIU) combined with a custom interface between CODAC and the simulator to meet these requirements. This approach allows integration of a wide range of CODAC tools (HMI, Archive, Alarms, Logbook, Operations Sequencer), synchronized with the simulator, with a lightweight and efficient solution.

## INTRODUCTION

The aim of ITER Plant Simulator (IPSi) is to provide a flexible platform and models to support various needs for the International Thermonuclear Experimental Reactor (ITER), including:

- Engineering simulation of single and integrated plant systems,
- Prototyping of ITER Control, Data Access and Communication (CODAC) [1] supervisor (SUP) [2], which will be used to coordinate and operate the plant systems for various Operational Tasks (Vacuum Vessel Baking and Pump-down, Cryo-plants cooldown, …)
- Operator training.

IPSi focus is not plasma pulses simulation, for which other tools such as Plasma Design Simulator (PDS) [3] are being developed, but the plant supporting the Tokamak.

IPSi simulator runs on Windows and has been based on TeamSuite simulation software by Westinghouse [4] to perform the work presented in this paper. It relies on the Functional Mock-Up Interface 2.0 standard (FMI) [5] to co-simulate Functional Mock-up Units (FMUs) of various plant systems models physics and control logics (Cooling Water systems, Cryoplant, Vacuum…) developed by the Iter Organization (IO).

The ability to integrate CODAC, which runs on Linux, on top of these models in the simulator is key for SUP prototyping, integrated operation preparation and operators training.

In many industries, it is common practice to have simulators using emulation or high-fidelity simulation of Digital Control System (DCS) coupled with process models and providing features such as save/restore or backtracking. One will often pick between using a commercial emulator of the complete DCS (logics and HMI) when available, and between simulating the DCS in a simulation framework, which can represent a significant effort for a realistic replica.

CODAC was not designed to be used in such a simulator context. No emulator or simulator compatible version of it exists as far as we know. We present in this paper an approach which leverages Experimental Physics and Industrial Control System (EPICS) [6] features combined with Linux tools libfaketime [7] and Checkpoint/Restore In Userspace (CRIU) [8] to encapsulate CODAC and make it fulfil simulator requirements.

## SIMULATOR REQUIREMENTS

IPSi project was initially inspired by Nuclear Power Plant (NPP) Full Scope Simulators (FSS) which follow the ANSI 3.5 standard [9]. Many ANSI 3.5 requirements are relevant for interactive operator training on ITER, like the possibility to save/restore simulator state at any time, and the possibility to backtrack if the operator made an error.

But operation scenarios on ITER can be significantly longer than on an NPP, and operating strategies and procedures design is a work in progress which is likely to see lots of changes throughout ITER lifecycle. We are thus looking for a simulator with ANSI 3.5 like features, but that can also operate much faster than Real Time, even when CODAC is integrated.

### *Functional Scope*

CODAC has many components and provides lots of tools, operator training and SUP prototyping lead us to target the following elements, which frontend can be seen in see Fig. 1, in priority:

**CODAC Core Software (CCS)**

- Control System Studio (CSS) for operator HMI
- Archive Engine (BEAUTY)
- Alarm server (BEAST)

**CODAC Operation Application (COA)**

- Supervisor SUP sequencer (oac-tree), a behaviour tree engine to automate operation [2].

**Data Visualization Tool**

- MINT, an interactive data visualization tool [10, 11]

† romain.balp@iter.org

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEMR015

# SCHEDULER FOR COOLING AND VENTILATION PLANTS: FEEDBACK ON EASY AND LOW COST METHOD FOR ENERGY SAVINGS*

N. Bunijevac†, D. Monteiro, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.


## *Abstract*

In industrial engineering, scheduling is a well-established strategy for optimizing resource use and minimizing operational costs. At CERN's Engineering department, the Cooling and Ventilation (CV) group has implemented an automatic scheduling solution to reduce electricity consumption by selectively shutting down HVAC plants during nights and weekends, when their operation is not required. Given that CV systems account for a significant share of CERN's total electricity use, even simple scheduling strategies can yield substantial energy savings - up to 75% in some cases. This paper presents the motivation, methodology, and preliminary results of scheduler deployments across multiple CV plants between 2023 and 2025, including recent pilot projects at Point 5 of the Large Hadron Collider (LHC). The categorization of scheduler strategies has been defined depending on the risks: *no savings*, *savings*, *savings+* and *on demand*. For each, control modifications have been standardized. Additionally, operational safety is carefully assessed; a $CO_2$ measurement campaign is carried out at the access station Point 5 of CERN's largest accelerator, to confirm compliance with environmental and safety requirements. The preliminary results from several sites show a significant reduction in consumption while the controlled process variables, such as zone temperature, remained within required limits. This low-cost approach demonstrates how simple digital solutions can lead to significant energy savings in large-scale technical infrastructures.


## INTRODUCTION

In response to increasing awareness regarding growing environmental consciousness, together with energy rising costs, many governments, companies and organizations have been taking actions to reduce their energy footprint, including Organization for Nuclear Research (CERN). In a report published in 2021, CERN management set a goal "to establish itself as the model for transparent and environmentally responsible research organizations," which includes implementing "actions and technologies aiming at energy saving and reuse" [1]. Among many initiatives, Cooling and Ventilation group (CV), a part of engineering department (EN), focuses on improving efficiency of cooling and ventilation systems. Cooling and ventilation systems maintain air quality for accelerator machinery and experiments equipment, but are, like other systems of this type, very energy intensive. During the experiments run phase, the electricity consumption dedicated to cooling and ventilation makes 12% (75GWh) [2] of the total electricity consumed by the Large Hadron Collider (LHC), CERN's flagship accelerator. The initiative in CV group presented in this paper is scheduler, a control system modification method for energy savings in HVAC plants. Historically, CV plants have been operating continuously even though ventilation services are not required continuously, unnecessarily wasting energy. This is where the energy savings potential lies. Scheduler introduces the possibility of dynamically switching on/off the plant according to users' needs. This strategy has been deployed in several installations at CERN and is proven to be a simple and effective method for energy savings. In fact, its implementation is mandatory for the host states of CERN: France [3] and the Swiss cantons Geneva [4] and Vaud [5]. This paper first introduces the logic behind the scheduling method, then presents the study to evaluate scheduler performance in a pilot project from perspectives on energy savings and possible impact on controlled ambient temperatures and air quality. The conclusion summarizes the findings and presents plans for future work.

## METHODOLOGY

HVAC plants are the focus plants to implement scheduler, unlike water-based cooling plants that are typically more critical, and hence less suitable to be occasionally stopped. For each of the HVAC plants, depending on different critical and noncritical aspects, one of the four different scheduling strategies can be deployed. The strategy is defined by different conditions that set the control signal to allow the plant to stop and save energy. First, four scheduling strategies are introduced and explained, and then the implementation steps are discussed.

### *Scheduling Strategy*

There are four groups of scheduling strategies: *no savings*, *savings*, *savings+* and *on demand*. Each plant is assigned to one of these depending on certain risk factors, e.g., presence of people, dangerous gases, need for pressurization. In the following subsections for each group, there are graphs that present binary conditions:

- parameter *working hours* (by default with value 1 from 7 am to 7 pm, and 0 elsewhere),
- signal *T outside limits* (triggered when controlled temperature - usually the average zone temperature - exceeds desire limits),
- *demand signal* (typically activated and deactivated by a button).

The signal *Run Allowed* is generated by the logical conjunction of these conditions. It is integrated to the existing plant's control strategy as an additional control step.

* Work supported by CERN

† nikolina.bunijevac@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD002

# ATLAS DEMO INHERITANCE COMMISSIONING AND PERFORMACE TESTING USING SCADA AND PLC BASED AUTOMATIZED PROCEDURES

L. Zwalinski†, A. Baran, V. Bhanot, M. Ciupinski, J. Daguin, L. Davoine, M. Doubek
S. Galuszka, Y. Herpin, Y. Penders, P. Petagna, K. Sliwa, D. Teixeira, B. Verlaat, M. Zimny
CERN European Organization for Nuclear Research, Geneva, Switzerland

*Abstract*

By the end of 2024, to cope with needs of future ATLAS tracking detector (ITk), surface integration and testing facility at CERN was fully equipped with a new large power and low temperature cooling system nick-named "DEMO Inheritance" based on the 2PACL technique. This paper will discuss implemented control system solutions for both the cooling plant and the distribution system installed in the detector proximity at the assembly clean room. The PLC and SCADA software has been fully deployed, following CERN UNICOS framework and allowed for successful and rapid initial performance test with 50kW dummy load. In this paper we will describe in details SCADA and PLC based procedures used for automatic system performance tests with no operator attendance as the first step in future commissioning of the final detector cooling systems of ATLAS and CMS detectors at the High Luminosity era of the Large Hadron Collider at CERN. These procedures running unattended allow for significant commissioning time reduction which is a key requirement for very challenging schedule.

## INTRODUCTION

The ATLAS DEMO Inheritance is an evolution of the first preproduction large scale low temperature $CO_2$ detector cooling system nicknamed DEMO [1], constructed by the Experimental Physics Department, Detector Technology Group (EP-DT) at CERN to prove correctness of technical choices and thermal performance before launching serial production of the underground final systems for Phase II Upgrade of the LHC Detectors.

DEMO Inheritance system built in collaboration of ATLAS cooling team and EP-DT is based on the 2-Phase Accumulator Control Loop (2PACL) [2] concept selected for ATLAS and CMS tracking, timing and calorimeter detectors for High Luminosity LHC era. By the end of 2024 the ITk surface integration and testing building (SR1) was fully equipped with its final 50kW cooling plant and accumulator unit, shown on Fig. 1, together with the first two out of six distribution racks, planned for complete installation later in 2025. The 2PACL cooling plant, accumulator, control cabinets and filling accumulator are located outside of the clean room. The cooling system interconnects with the distribution racks shown on Fig. 2, inside the clean room via two nonconcentric thermally insulated transfer lines allowing to cool down detector as low as -40°C. DEMO inheritance is using a two head cooling plant with a mass flow of about 1kg/s which sets it as the middle size system in Phase II cooling systems family.



Figure 1: ATLAS DEMO inheritance.



Figure 2: SR1 distribution rack inside the clean room.

The extra non usual feature applied in SR1 and never used for underground applications is the filling accumulator. This small accumulator allows for smoother and faster detector elements reconnection to the cooling system. It is required as only one 2PACL large accumulator is present in the system, hence only one common detector evaporation temperature can be set for all six-distribution rack at the same time.

The primary refrigeration system nicknamed "System D" based on R744 [3] (not oil free $CO_2$ industrial refrigeration technology), delivered by Cooling and Ventilation Group of the Engineering Department at CERN, has been installed in suitable outside container.

The DEMO inheritance will be exploited at least until ITk detector installation in the experimental cavern foreseen in 2029.

### *Control System*

The control software conforms to CERN's Unified Industrial Control System (UNICOS) [4] framework and uses Siemens WinCC OA as Supervisory Control And Data

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD003

# EXTENSION OF THE SPIRAL2 PLC-BASED CONTROL SYSTEM FOR THE INTEGRATION OF DESIR AND NEWGAIN

Q. Tura, T. Beauchef, B. Luce, C. Hocini, M. Lochon, A. Trudel
Grand Accélérateur Nat. D'Ions Lourds, Caen, France
P. Alfaurt, L. Daudin, Laboratoire de Physique des deux Infinis Bordeaux, Gradignan, France

*Abstract*

The SPIRAL2 heavy ion accelerator is currently undergoing several extension projects under development. In Phase 1+, the DESIR experimental hall will receive very low energy radioactive ion beams from S3 or SPIRAL1. In Phase 1++, a new cryogenic ion source and a new injector will be implemented to broaden the range of heavy ions currently accelerated. This contribution presents the integration of these new facilities into the SPIRAL2 PLC-based control system. The integration builds on recent technological upgrades while maintaining system consistency and ensuring compliance with the specific constraints related to radiation protection and operational safety. Some existing PLCs have been extended, and new ones have been added to automate processes related to beam operation (machine protection system, equipment insertion, vacuum, RF, cryogenics), building infrastructure (nuclear ventilation, refrigeration, alarm management), and nuclear safety (access control units, radiation monitoring). This integrated approach ensures a coherent, maintainable control system that supports safe and efficient operations for the DESIR and Newgain projects.

## INTRODUCTION

Since 2019, the SPIRAL2 Linac (linear accelerator) for heavy ions has been in operation at GANIL, operating in parallel with the laboratory's historical cyclotrons. The SPIRAL2 project includes several development phases aimed at further enhancing GANIL's competitiveness on the international stage of nuclear physics research. This paper outlines the work carried out by the automation engineers of the laboratory's control system group to design and develop PLC-based control systems required for the automation of processes in these new facilities.

### *Phase 1+*

The DESIR [1] project (Decay, Excitation, and Storage of Radioactive Ions) aims to deploy new beamlines, online beam purification devices and a new experimental hall dedicated to the exploitation of very low-energy radioactive ion beams. These beams will originate from the SPIRAL2 facility's S3 (Super Separator Spectrometer) installation, as well as from GANIL's SPIRAL1 facility. The construction of the associated buildings is nearing completion, with the first stable beams expected by 2027.

### *Phase 1++*

The NEWGAIN project (New Ganil Injector) [2] involves the commissioning of a new injector and a new ion source (A/Q=7) ASTERICS (Advanced Spiral Two Electron cyclotron Resonance Ion source at Caen with Superconducting magnets) upstream of the SPIRAL2 linear accelerator. This upgrade will significantly increase the intensity of the heaviest ion beams currently accelerated at GANIL. The installation of the new beamlines is expected to take place progressively over the coming years, within the infrastructure initially planned during the construction of SPIRAL2. First beams are anticipated by 2028.

## AUTOMATION AT GANIL

GANIL has historically chosen to work with Siemens PLCs, for which the laboratory's automation engineers have developed strong expertise in both development and maintenance. These PLCs are used in various accelerator subsystems: machine protection systems (MPS) [3], vacuum systems, HF operating condition control, cryogenic regulation, thermal regulation of the RFQ, access control systems for personnel safety with respect to high voltage or radiation protection, enhanced alarm managers, nuclear ventilation control, AutoFill, and emergency shutdown systems.

All generations of Siemens PLCs are represented. SPIRAL2 [4] was mainly developed using S7-300 and S7-400 PLCs, together with HMI panels running WinCC Comfort and WinCC Pro. In the original facilities, some S7-135U and S7-95U models are still in use, though they are gradually being replaced with latest-generation S7-1500 PLCs, particularly as part of the CYREN renovation project (CYclotrons RENovation). All PLCs on-site are networked and accessible through a VLAN dedicated to automation. Some PLCs also use Profibus fieldbuses, either for distributed I/O architectures or for direct equipment control—for example, vacuum pumps. Similarly, certain components belonging to the same type of function can communicate via the same specific VLAN.

However, due to the high beam power involved, safety-related information exchanges—whether for personnel or machine protection—are traditionally handled via hardwired digital (on/off) signals. These are sometimes duplicated when the associated function is considered particularly safety-critical. Following the same principle, functions identified as important for nuclear safety are not handled by PLCs, but by fully redundant systems as illustrated in Fig. 1, hardwired with non-programmable logic elements such as relays or optocouplers [5, 6].

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD005

# DEVELOPMENT OF AN EPICS-BASED CONTROL PLATFORM FOR ELECTRON BEAM COMMISSIONING AT ELI-NP*

G. Chen†, R. Patrascoiu, A. Cana, A. Iancu, A. Ionescu, P. Moralis, and C. Matei
Extreme Light Infrastructure – Nuclear Physics, Horia Hulubei National Institute for R and D in Physics and Nuclear Engineering, Magurele-Bucharest, Romania

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The ELI-NP (Extreme Light Infrastructure - Nuclear Physics) gamma beam system employs a 234-742 MeV linear accelerator to generate gamma rays via laser-electron interactions. Accelerator operation requires coordinated control of multiple subsystems (RF, vacuum, magnets, diagnostics). During pre-commissioning, we developed an EPICS-based platform to enable equipment debugging and beam tuning simulations before control system completion. The platform integrates both soft IOCs (for device emulation and algorithm validation) and device IOCs (for hardware control). This dual-IOC architecture allows: 1) High-level control software testing without physical hardware, and 2) Physics application interfaces for advanced functionality. The system successfully simulated control operations and supported tuning algorithm development. Testing demonstrated the platform's effectiveness for subsystem debugging and realistic beam tuning simulations. By bridging the gap between hardware availability and control system readiness, the platform is expected to accelerate the commissioning process and enhance the operational efficiency of the ELI-NP gamma beam system in the future.

## INTRODUCTION

The Gamma-ray Beam Source, under construction at Extreme Light Infrastructure – Nuclear Physics (ELI-NP), is a dedicated system for delivering gamma-ray beams with a variable energy range from 1 MeV to 19.5 MeV to the users [1–3]. Taking advantage of the concept of Inverse Compton Scattering, the GBS consists of a linear electron beam system coupled at the interaction point to a high-finesse optical cavity. The electron beam system comprises a Radiofrequency (RF) electron source [4], a linear accelerator system, and a transport line. The aim of the gamma source requires producing gamma rays from 1 MeV to 19.5 MeV. To obtain this range of photon energies, the range of electron beam energy has been limited from 234 to 742 MeV, and two laser wavelengths of $\lambda_1 \simeq 1.03\ \mu$m and $\lambda_2 \simeq 0.515\ \mu$m have been used.

The linac of the gamma beam source is a high-gradient linear accelerator system. The schematic layout of the LINAC is shown in Fig. 1 and Fig. 2. The electron bunches are generated within the RF electron source and then accelerated in the linear accelerator LINAC. There are five acceleration modules in total, each capable of adding up to 160 MeV of energy to a single bunch of variable charge. Figure 1 includes the electron gun, the first and second modules, as well as the first diagnostics station. Figure 2 shows the third to fifth modules, where the blue block on the far right represents the dipole magnet of the second diagnostics station.

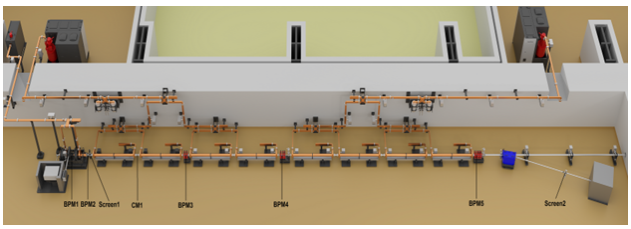



Figure 1: The first half of the ELI-NP linac.

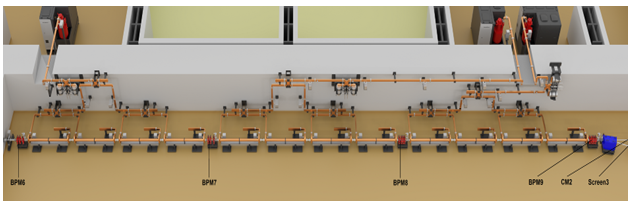



Figure 2: The second half of the ELI-NP linac.

Each acceleration module consists of the following main components: a 50 MW, 4 µs klystron (Canon E3730A or equivalent); a pulse modulator (Scandinova model K-300 or equivalent); an SLED system for RF pulse compression; four accelerator structures, each 2 m in length; and a waveguide distribution system including hybrids and bends. In addition, several RF waveguide loads are employed, along with a group of three or four quadrupole magnets for beam focusing between modules. Beam position monitors are installed within each magnet group, and RF power and phase diagnostics are integrated into the waveguide system to support tune-up and monitoring.

The linac is currently under installation: RF waveguides are still being installed, and diagnostics including BPMs with their electronics, as well as the profile and charge monitors, are still in the procurement process. Since the gamma beam source at ELI-NP is expected to start beam commissioning by 2026, an EPICS-based control platform has been established. While developing the control systems for the existing devices, including IOCs and OPIs, multiple simulated IOCs have also been created to serve as a foundation for the high-level physics commissioning software. As the linac will require

* This work is carried out under the contract PN 23 21 01 06 sponsored by the Romanian Ministry of Research, Innovation and Digitalization.

† guangling.chen@eli-np.ro

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD007

# COMMISSIONING AND OPERATION OF VACUUM CONTROL SYSTEM FOR SPES PROJECT

L. Antoniazzi*, D. Bortolato, A. Conte, F. Gelain, C. Roncolato, INFN/LNL, Legnaro, Italy
G. Savarese INFN/BA, Bari, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The SPES (Selective Production of Exotic Species) project aims to create a facility based on particle accelerators to produce radioactive ion beams. The second phase of the project foreseen the transport of the non-reaccelerated radioactive ion beam from the TIS (Target Ion Source) to the low energy experimental hall. This part of SPES beam lines involves the major requirements in terms of complexity for the vacuum control system, as interface with GRS (Gas Recovery System), GSS (Global Safety System) and MPS (Machine Protection System), management of centralize pumping system for the exhaust, and different configurations of each section.

The VCS (Vacuum Control System) of the TIS and of the following beam line sections are based on modular control units which are highly configurable to be used in different installation and with different equipment. Each unit is constituted by a SIEMENS S7-1500 PLC and 15" touch panel for the local/remote configuration and operation by expert operator, while high level control system is done in EPICS (Experimental Physics and Industrial Control System) and CSS (Control System Studio).

This paper describes the commissioning phase of the VCS for the SPES facility, and the operation of the about 20 systems running at LNL (Laboratori Nazionali di Legnaro)

## VACUUM CONTROL FOR PHASE TWO OF SPES PROJECT

The SPES facility [1], [2], under construction at INFN LNL (Istituto Nazionale di Fisica Nucleare Laboratori Nazionali di Legnaro), aims at the production of intense RIB (Radioactive Ion Beams). It employs the Isotope Separation On-Line (ISOL) technique, for interdisciplinary research scopes. The radioactive isotopes of interest are produced by the interaction of a multi-foil uranium carbide target with a 40 MeV 200 $\mu$A proton beam generated by a cyclotron proton driver.The fission products, are ionized, extracted and accelerated by an electrostatic potential of 40 kV. A series of mass separation, tuning, focusing and charge breeding stages are implemented prior to the injection in the ALPI (Acceleratore Lineare Per Ioni) LINAC (LINear particle Accelerator) [3] for the post acceleration.

The second phase of SPES project foreseen the transport of the non re-accelerated radioactive ion beam from the TIS (Target Ion Source) to the low energy experimental hall. There the diagnostic tape station is designed both for the TIS development/optimization and for the beam characterization before the delivery to the final users (see Fig. 1). The goal of the test executed in the last part of 2024, was to validate the functioning of all the involved system and devices with low intensity beam to minimize the residual activation.

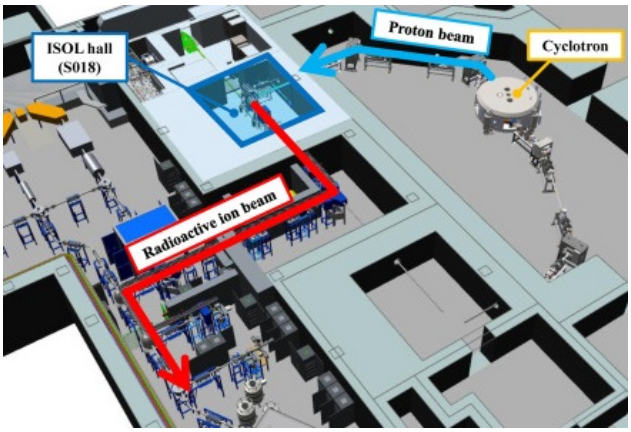



Figure 1: Layout of SPES phase 2.

The TIS is the core of the SPES project, here the radioactive nuclei are extracted, ionized and accelerated. Due to efficiency reasons, the TIS unit needs to be replaced periodically. In this highly radioactive environment all the related activity need to be automated. This means that the involved systems as vacuum, cooling, high voltage and high current power supplies, handling and coupling systems has to be properly coordinated [4]. Safety during beam operation and the coordination of the subsystems are entrusted to the Machine Protection System (MPS) which monitor the status, receives the requests to operate, distributes consents and interlocks to the various devices. The MPS-VCS interface is implemented via Modbus TCP, the exchanged variable together with the bypass possibility are presented in Fig. 2.

Accelerated radioactive beams and their contaminants carry radioactivity along the subsequent transport and selection beam lines. Event though, the level of contamination attenuates gradually as one moves away from the source, the output of the backing pumps can contain radioactive isotopes. Consequently, the delivered gases must be checked and, in case, stored up until their radioactivity has dropped to acceptable levels before being released. At this scope all the gasses from TIS and Low Resolution Mass Scanner (LRMS) beam lines are collected by the High Activity Gas Recovery System (HAGRS). In a similar way, all the VCSs of the low energy experimental area are connected to a second set of centralized pumping lines for fore-vacuum and backing pumps. Centralized pumping system are controlled by one VCS which is interfaced with the HAGRS (in the case of TIS

* loris.antoniazzi@lnl.infn.it

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD012

# WAVEFORM MONITORING SYSTEM

D. Paskvan†, G. Wang, T. Fors
Argonne National Laboratory, Lemont, IL, USA

*Abstract*

The Advanced Photon Source (APS) recently completed a significant upgrade to its storage ring, replacing all existing components with new ones. One of the new systems is the waveform monitoring system for the high voltage pulsers. This system is a 2U rackmount chassis with 8 ADC inputs, MRF event link and 16 channels of TTL I/O. This is an FPGA based 8 channel 4 GS/S digitizer that monitors the decoherence and injection high voltage pulser waveforms. The main function of the system is to qualify each pulse and decide whether to proceed with injection or not. This paper will discuss the details and development.

## INTRODUCTION

The recent upgrade to the Advanced Photon Source (APS) at Argonne National Laboratory, significantly enhances beam brightness and stability to support next-generation scientific research. This upgrade introduced new challenges in beam injection and stability, requiring precise monitoring and of high-speed, high-voltage signals.

To address these challenges, a new waveform monitoring system was developed. This system is designed to capture and analyze critical pulsed waveforms in real time, enabling the accelerator control system to make informed decisions about whether to proceed with beam injection. The system integrates high-speed digitization, FPGA-based processing, and event-synchronized triggering to ensure accurate and deterministic operation within the APS timing infrastructure.

High-voltage (HV) pulse generation plays a central role in accelerators, beam instrumentation, and pulsed power experiments. Precise evaluation of these pulses—particularly their timing, shape, and stability—is essential to ensure reliable operation and to protect experimental hardware. Traditional evaluation methods rely on offline data acquisition followed by CPU-based post-processing. However, as repetition rates increase, real-time methods become necessary.

Using the Xilinx ZCU111 evaluation board, which integrates RF-class data converters and a large FPGA fabric, we developed a streaming cross-correlation engine in Vitis HLS to process digitized pulses at high sample rates. The system enables fast, in situ pulse evaluation with reduced latency. Results demonstrate that the method provides robust timing alignment, noise rejection, and shape evaluation, making it suitable for online monitoring of pulsed accelerator systems.

* Work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No, DE-AC0206CH11357
† dpaskvan@anl.gov

This paper presents the design, implementation, and performance of the waveform monitoring system, highlighting its role in improving injection reliability and providing enhanced diagnostic capabilities for APS operations.

### *System Architecture*

The ZCU111 was the chosen system architecture based on the specifications listed below:

- 8-channel 4.096 GSPS ADCs and 8-channel 6.554 GSPS DACs.
- XCZU28DR RFSoC FPGA fabric
- Embedded ARM Cortex-A53 processors for control and monitoring.
- Custom 8 channel analog front-end
- Custom FMC 8 channels of TTL inputs and 8 channels of TTL outputs.
- DAQ interface for archiving data and streaming data.
- MRF event system for timestamping DAQ data with specific events.

In the images below Fig. 1 shows the internal chassis layout and wiring of the waveform monitoring system. Figure 2 is front panel of the waveform monitoring system.



Figure 1: Image of Waveform Monitoring chassis.




Figure 2: Image of Waveform Monitoring Front Panel.

### *Implementation*

On each injection into the Storage Ring data is digitized and captured in the RAM of the FPGA on the ZCU111. There is a custom IP using Vitis HLS that retrieves that data and performs a normalized cross correlation on the data using the equation below. Below in Fig. 3 is a simple block diagram of the FPGA.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD014

# DESIGN OF AN UPGRADED ANALOG SIGNAL DIGITIZER TO REPLACE THE MADC SYSTEM AT RHIC*

P. Bachek†, M. Costanzo, J.Morris
C-AD, Brookhaven National Laboratory, Upton, NY, USA

*Abstract*

A new general-purpose analog signal digitizer has been designed and prototyped to serve as an upgrade to the legacy Multiplexed Analog to Digital Converter (MADC) system currently in use around the RHIC accelerator and injector complex at BNL. The new system is a standalone rackmount chassis with an embedded System on a Chip (SoC). This is a departure from the traditional VME form factor used by most legacy controls equipment within the Collider Accelerator Department. New features include completely independent channels, real time digital signal processing, large sample buffers, built-in timing links, and high bandwidth network connectivity. Support is included for the legacy timing links as well as future compatibility with the EIC Timing Data Link. The core features, system architecture, and scheme for integration with the controls system network is presented.

## INTRODUCTION

The MADC system at RHIC was originally designed with 64 analog input channels multiplexed to a single ADC for sampling [1]. It is being replaced by a new general-purpose digitizer called MADC2. MADC2 has 32 independent analog input channels each with its own ADC. It is a standalone 1U rackmount chassis instead of a VME card. The system is based on an AMD SoC allowing for flexible digital signal processing and high-speed data transfer via Ethernet to the controls system infrastructure.

## HARDWARE DESIGN

The ADC front end has a high impedance analog signal input to a fully differential ADC. The ADC is then connected to an FPGA for direct control of sampling. The FPGA has access to DDR memory for buffering ADC samples before transmitting the data to the controls system network.

### *ADC Front End*

Each of 32 analog input channel accepts a differential input signal within the range of ±10.24 V on a differential LEMO connector at the device's front panel. The positive and negative signals are each buffered by a noninverting high input impedance OPA328 amplifier with a bootstrap connected power supply to increase the input range and slew rate. This provides an overall differential input impedance of 1 TΩ || 4 pF.

The buffered signal is then driven to an ADA4945-1 differential ADC driver amplifier through an RC filter network. The differential amplifier is used to implement a 2nd order Multiple Feedback Topology antialiasing filter with a 3 dB cutoff at 1 MHz. The overall amplifier gain is set to 0.4 V/V to scale the input signal for input to the ADC. The differential amplifier signal common mode voltage is set to 2.048V which sets the output signal common mode level.

The conditioned signal from is then fed into a AD4003 ADC converter through a charge kickback filter. The AD4003 is a fully differential 18-bit 2 MSPS converter. The reference voltage used is 4.096 V to use the full range of the converter with an input signal range of ±10.24 V.

The ADC has a SPI interface which is connected to the FPGA fabric with an IO voltage of 1.8 V. To achieve the maximum sampling rate of 2 MSPS while observing the ADC quiet periods and meeting the setup and hold requirements of the SPI interface, the SPI clock must run at 75 MHz. This leaves very little timing budget at the FPGA pins when accounting for PCB trace delays of several inches depending upon the physical location of the channels.

### *Processor*

The MADC2 uses an AMD Zynq Ultrascale+ XCZU4CG SoC. This includes a dual-core Arm Cortex-A53 MPCore™ up to 1.3 GHz and a dual-core Arm Cortex-R5F MPCore up to 533 MHz processor. It also has FPGA fabric which will interface directly with each ADC. The processor is connected to 4 GB of external DDR4 memory with a 64-bit bus at 1600 MHz, 128 MB of NOR flash, and 256 Bytes of EEPROM with a 48-bit preprogrammed unique ID.

The FPGA is configured with 32 independent signal chains each with its own SPI interface, sample rate control, digital filtering, and triggered buffer capture. It also implements DMA transfer of ADC samples directly to DDR memory.

### *Connectivity*

The MADC2 has four SFP+ ports connected to the GTH transceivers on the FPGA capable of running up to 16.3 Gbps. The SFP+ connections can communicate with the 8 Gbps EIC Timing Data Link for future compatibility. It also has Gigabit Ethernet directly connected to the processor. There are two Twinax connectors available for Event Link and RTDL connection with existing RHIC infrastructure.

## SIGNAL PROCESSING

The FPGA is used to implement a custom signal processing chain suitable for various applications at the accelerator complex. The biggest challenge with signal filtering

* Work supported by Contract Number DE-AC02-98CH10886 with the auspices of the US Department of Energy
† pbachek@bnl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD017

# WREN: A VERSATILE WHITE RABBIT EVENT NODE FOR CERN'S TIMING SYSTEM RENOVATION

T. Gingold*, M. Cejp, A. Dujović, E. Gousiou, F. W. Hoguin, F. Irannejad, I. Kozsar, G. Kruk, T. Levens, G. Moscardi, M. Murillo Moya, P. Peronnard, J. Serrano, T. Wlostowski, A. Zeising,
European Organization for Nuclear Research (CERN), Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

WREN is a versatile White Rabbit (WR) node developed for CERN's event-based timing system renovation. Thousands of WRENs are expected to be deployed across the whole CERN accelerator complex from 2027 onwards. Equipped with dedicated hardware and gateware, WREN integrates synchronisation in both TAI (International Atomic Time) and RF (accelerator Radio Frequency) timing. It can function as a TAI event transmitter and receiver, a Beam Synchronous (RF) transmitter and receiver, and is also capable of FPGA-based time-to-digital conversion and fine-delay generation. WREN is highly adaptable for various timing and trigger distribution systems. It is available in multiple form factors, including PCIe, VME, PXIe, and μTCA. All boards are based on the Zynq UltraScale+ System-on-Chip (SoC), designed using the open-source KiCad tool, and licensed under the CERN Open Hardware License (OHL). The gateware and software are also open source. This paper presents the WREN hardware modules, the gateware architecture, and potential customisations for applications beyond CERN. It also shares insights from the initial pilot deployments at CERN.

## INTRODUCTION

### *International Atomic Time and Accelerator Radio Frequency Timing*

CERN's accelerator complex delivers particle beams to a wide range of experiments, each with highly specific requirements in terms of beam properties such as intensity, energy, and size. To meet these demands, beams pass through several accelerators and undergo complex manipulations in the CERN accelerator complex, including particle production, bunching, cooling, steering, acceleration, transfer, and eventual delivery to the experimental points. These processes are orchestrated through precisely timed sequences within each accelerator's operational cycle.

Two timescales are fundamental for these operations and are often used together by different types of equipment: TAI[1] (International Atomic Time, our ordinary everyday time) and RF (accelerator Radio Frequency) timing. TAI provides nanosecond precision when defining Central Timing Events (CTIMs), such as Cycle Start, Beam Injection, Beam Extraction and Cycle End. The RF timing, in contrast, is specific to a given accelerator and derived directly from the radio frequency that acts on the particle beams. It is defined by (a) the Turn clock, or revolution frequency, which is the rate at which particles complete one turn in the accelerator, and (b) the Bunch clock, a unit fraction of the base acceleration frequency. In the Large Hadron Collider (LHC), the Bunch clock (40 MHz) is one tenth of the base acceleration frequency (400 MHz) and represents the maximum frequency at which particle bunches can cross at the experimental points, making it the key reference for synchronising all electronic systems involved in bunch measurements and event detection.

Building on these two timescales, CERN currently operates three distinct timing systems [1] across its accelerator complex:

- **General Machine Timing (GMT):** Based on the TAI timescale, GMT distributes both data and a common notion of time throughout the entire accelerator chain. All accelerator equipment makes use of this system. At the physical layer, GMT relies on RS-485 serial links operating at 500 kbps over copper, with a single transmitter feeding hundreds of receivers.
- **Radio Frequency (RF):** Based on the RF timescale, this system distributes the Turn clock and Bunch clock through optical links with very low jitter, and without any data payload. There are three instances in operation: one for the Super Proton Synchrotron (SPS) and one for each LHC beam. It is used primarily by equipment that depend on knowledge of the arrival time of individual bunches, such as the experiments and injection / extraction kickers, while slower systems (e.g., power converters for dipole magnets) typically do not require it.
- **Beam Synchronous Timing (BST):** Also based on the RF, BST combines a less precise RF reference with GMT data. The GMT data are encoded using the RF clocks into a 160 Mbaud optical stream, allowing both RF clock and data distribution through relatively low-cost electronics. BST is used mainly by beam observation equipment to align their measurements with individual circulating bunches.

### *White Rabbit Timing (WRT)*

The WRT project [1] aims at consolidating these systems (GMT, RF and BST) by using the same architecture and the same infrastructure: a physical White Rabbit network. The creation of such a global, CERN-wide network not only eliminates the need of maintaining the multitude of existing

* tristan.gingold@cern.ch

[1] TAI is the basis for UTC, which is used for civil timekeeping all over the Earth's surface and which has leap seconds.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD018

# OPEN SOURCE EVENT TIMING SYSTEM AT ALS-U *

L. M. Russo†, A. Amodio, J. M. Weber, LBNL, Berkeley, CA, USA
W. E. Norum, Osprey DCS LLC, Ocean City, MD, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Advanced Light Source Upgrade (ALS-U) is a major upgrade project for the existing light source at Lawrence Berkeley National Laboratory. One of the key challenges of the upgrade is that the new accelerator sections cannot operate at the existing RF frequency. The injector will operate at the current RF frequency ($f_1$ = 499.64 MHz) and the new Accumulator Ring and Storage Ring will operate at a rationally related frequency ($f_2 = \frac{p}{q} * f_1$ = 500.39 MHz). As a result, the Timing Event Generator (EVG) must be synchronous to both RF frequencies and generate separate event streams for each frequency domain. To support beam transfer from the injector ($f_1$) to the Accumulator Ring ($f_2$), the EVG must detect the coincidence of the two frequencies and synchronize counters to the coincidence rate. Due to extensive need for timing channels, Event Fanout (EVF) and Event Receiver (EVR) chassis were also designed as part of this effort. This paper describes the main concepts of the Dual EVG, EVF and EVR projects, open source gateware, embedded software and EPICS IOC implementations, as well as the tests and deployment at the current ALS.

## INTRODUCTION

The Advanced Light Source (ALS) currently operates with a commercial off-the-shelf Timing system from Micro-Research Finland Oy (MRF) [1]. It uses a combination of VME MRF Event Generators (EVG), VME MRF Event Fanouts (1 input, 11 outputs) and VME MRF Event Receivers (EVR), in-house developed receivers (i.e., User Timing Chassis) and embedded event receivers. The latter was initiated by NSLS-II and has since been modified to provide deterministic latency (i.e., fixed phase relationship between the EVG and EVR) [2].

With CD-1 approval of the ALS Upgrade project in 2018 [3], there was recognition that the commercial solution used did not support (or had any plans to do so in the future) the requirement of synchronization of 2 rationally related RF frequencies. Moreover, there was a desire to completely own the Timing System implementation from the hardware, gateware, and embedded software perspectives, which up until that point had only been partially achieved with the embedded EVRs. The main reasons for doing so were (i) the ability to change the system as the requirements evolved and (ii) the cost of the hardware.

In owning the implementation, it was also possible to start thinking about benefiting the accelerator community by open sourcing the design. This is still a work in progress, but much has been achieved by releasing the entirety of the gateware and embedded software for all three system components - EVG [4], Fanout and EVR [5] - as well as the main FMC Carrier board, Marble [6]. The authors are also in the process of releasing two in-house designed FMC modules as open source projects. The first one is called Event I/O (EVIO), and its purpose is to fan-out timing links, specifically 1 to 35 using three Samtec Firefly modules [7]. The second one is called Event Receiver I/O (EVRIO), and it's designed to generate triggers, patterns and clocks synchronized with the master oscillator [8]. The EPICS IOCs for all systems are also going to be released soon.

* Work supported by the Director, Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

† lucasrusso@lbl.gov

## THEORY OF INJECTION

ALS-U will operate with two rationally related RF frequencies: the Injector (Linac and Booster) RF, $f_1$ = 499.642 MHz, and the new accelerator sections (Accumulator and Storage Rings) RF $f_2 = \frac{p}{q} * f_1$ = 500.392 MHz. Because $f_1$ is the frequency that ALS currently operates on this is also known as $f_{1,nominal}$, $f_{ALS}$ or $f_{BR}$. Similarly, $f_2$ is the projected ALS-U Accumulator (AR) and Storage Ring (SR) frequency and is also known as $f_{2,baseline}$, $f_{AR}$ or $f_{ALSU,SR}$. The dual-frequency design is necessary because the Injector is not part of the upgrade and will therefore continue to operate at its current RF frequency $f_1$. Conversely, the Storage Ring physical path length was modified to accommodate the new lattice design, requiring the RF frequency to be increased by 0.15%.

Generically, the $RF$ frequency $f_{RF}$ can be calculated using Eq. (1):

$$f_{RF} = \frac{c * \beta}{C} * h \tag{1}$$

where $c$ is the speed of light, $\beta$ is the fraction of ligthtspeed at which the beam moves, $h$ is the harmonic number and $C$ is the vacuum chamber circumference. For our purposes, the product $c * \beta$ = 299 792 448.2 m/s

For ALS-U SR, $f_{ALSU,SR} = f_{AR}$ can be calculated by using the nominal $C_{ALSU,SR}$ as shown by Eq. (2):

$$\begin{aligned} C_{ALSU,SR} &= 196.509\,678\,9\ \text{m} \\ f_{ALSU,SR} = f_{2,baseline} &= 500.392\,263\ \text{MHz} \end{aligned} \tag{2}$$

For ALS, $f_{ALS,SR}$ can be calculated by using the nominal $C_{ALS,SR}$ as shown by Eq. (3):

$$\begin{aligned} C_{ALS,SR} &= 196.804\,758\,2\ \text{m} \\ f_{ALS} = f_{1,nominal} &= 499.641\,999\ \text{MHz} \end{aligned} \tag{3}$$

In order to enable deterministic injection between the two frequency domains, specifically between the Booster ($f_1$)

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD020

# MODERNIZING SOFTWARE AND HARDWARE FOR LANSCE EVR WITH FPGA AND REAL TIME LINUX*

Z. M. Sauer†, J. Duran, Los Alamos National Laboratory, Los Alamos, NM, United States

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

This paper presents the approach to modernizing the Los Alamos Neutron Science Center (LANSCE) distributed timing system using open source Micro Research Finland (MRF) Event Receiver (EVR) [MRF Open EVR, https://github.com/jpietari/mrf-openevr]. The MRF EVR will be replaced by a Xilinx UltraScale+ Multi-Processor System on a Chip (MPSoC). The VME-EVR-300 system has been discontinued and replacing it is a critical part of the future of LANSCE.

This project will utilize Vivado for FPGA design, compilation and test benching. Real Time (RT) Linux development will primarily use the Yocto build project [Yocto Build Project, https://www.yoctoproject.org/about/project-overview/] along with Xilinx supporting files. Yocto is an open source toolset for RTLinux development.

## INTRODUCTION

### *Current System*

The EVR is responsible for:

- Syncing with the distributed timing system of LANSCE
- Converting event codes from Event Generators (EVG) to hardware
- Storing and maintaining the Per Cycle Data Buffer (PCDB)
- Receiving and relaying fast protect information
- Forwarding timing information to downstream systems

Table 1: Example Event Codes

| Event code | Offset | Internal functions | Pulse Triggers | 'Set' Pulse | 'Reset' Pulse |
|---|---|---|---|---|---|
| 0x00 | 0x0000 | 4B/32b | 4B/32b | 4B/32b | 4B/32b |
| 0x01 | 0x0010 | 4B/32b | 4B/32b | 4B/32b | 4B/32b |
| 0x02 | 0x0020 | 4B/32b | 4B/32b | 4B/32b | 4B/32b |
| ⋮ | ⋮ | ⋮ | ⋮ | ⋮ | ⋮ |
| 0xFF | 0x0FF0 | 4B/32b | 4B/32b | 4B/32b | 4B/32b |

EVRs receive event code information (Table. 1) from EVGs and then control hardware gates to produce the specified beam mode. EVRs are also responsible for storing and maintaining the PCDB (Table. 2), which updates at 120.0 Hz. The EVR will store a copy of the PCDB in onboard Random Access Memory (RAM) for reference when performing event code actions. The EVR also acts as a "bent pipe" forwarding cycle, gate, time, and PCDB information to a downstream board over compact Peripheral Component Interconnect (cPCI) backplane. These transactions are facilitated through a CR-12 processor [1].

### *Proposed System*

The proposed upgrade system will adhere to all timing requirements presented by the previous design. This includes 100 ns gate resolution and 10 ns gate precision [2]. The upgraded device will continue support for all LANSCE timing modes, see Table 3 for more information on requirements.

The current development plan is create a Minimum Viable Product (MVP) as a proof of concept before adding extra features. To create a MVP the upgraded EVR will not need to provide backplane communication support, this can be added later on. The upgraded system will include the following improved capabilities:

- Improved update system allowing for remote updates of RTLinux, FPGA image, and user space applications like EPICS
- Improved diagnostics by utilizing custom RTLinux FPGA driver
- Improved clocking frequencies for FPGA allowing for even higher time constraints in the future
- Additional four CPU cores allow for increased computational power
- Increased on board memory allowing for long term diagnostic data storage

The additional computational power could allow for hardware simplification by replacing the previous and obsolete CR-12 processors duties. The previous CR-12 processor was responsible for hosting the EPICS application and facilitating cPCI backplane communications [1]. These tasks could easily be moved to RTLinux.

The current design keeps all timing dependent tasks within the FPGA. This choice was made due RTLinux being unable to 100 ns gate resolution and 10 ns gate precision [3].

Using Direct Memory Access (DMA) to pass information from the FPGA to RTLinux will allow for additional statistics and diagnostics, decreasing debug time. The upgraded EVR will forward flags and counters from the FPGA, providing useful insight without recompilation. For more information on the planned design see Fig. 1 that contains conceptual block design of the project.

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). LA-UR-25-29063

† zmsauer@lanl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD023

# DEVELOPMENT STATUS OF FPGA-BASED FOFB SYSTEM FOR PLS-II

Wooseong Cho†, Jaeyu Lee, Ilmoon Hwang, Changbum Kim
Pohang Accelerator Laboratory, POSTECH, Pohang, Gyeongbuk, Republic of Korea

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The third-generation synchrotron accelerator Pohang Light Source-II (PLS-II) at Pohang Accelerator Laboratory uses a Fast Orbit Feedback (FOFB) system to maintain beam orbit stability in the storage ring. The FOFB system operates in real time to suppress orbit perturbations in both horizontal and vertical directions. Currently, the system uses VME-based Single Board Computers (SBCs) and Reflective Memory (RFM) technology, achieving a feedback repetition rate of about 1 kHz. However, the aging hardware is causing difficulties in maintenance and performance upgrades. To solve this issue, a new FOFB system based on Zynq UltraScale+ FPGA high-speed digital processing technology is under development, aiming to increase the repetition rate to 10 kHz. The new design distributes twelve independent FOFB controllers throughout the storage ring to minimize latency from Beam Position Monitor (BPM) Fast Acquisition (FA) data reception to the output of control signals to the Fast Power Supply (PS). The system is being developed to work stably with the existing Fast PSs at 1 kHz and to support future high-performance supplies capable of operating at higher rates. The FPGA-based FOFB system is currently under development, with a goal of achieving a suppression bandwidth greater than 100 Hz and significantly improving maintainability and scalability. This paper introduces the design concept and the current development status of the new system.

## INTRODUCTION

In the PLS-II storage ring, a FOFB system to suppress disturbances occurring in the frequency band below 100 Hz was introduced in 2015 to ensure beam orbit stability [1]. The system has a feedback repetition rate of 833 Hz, and each cell is equipped with 8 BPMs and 4 Fast Correctors. The FOFB system is based on the Versa Module Europa bus (VME) platform and consists of 1 master and 12 nodes. The master handles timing signal processing and manages and monitors control parameters via the EPICS, while each node, composed of RFM and a SBC, shares the X and Y positions of the BPM FA data and performs orbit correction calculations. However, the discontinuation of the SBCs and the increasing maintenance cost of the VME platform have created difficulties in scalability and maintenance. To address this, a next-generation FOFB controller based on a Zynq UltraScale+ FPGA is currently under development using the ZCU102 evaluation board [2]. Using a ZCU102 evaluation board, this controller targets a 10 kHz repetition rate. The internal processing latency of a single controller was measured by a loopback test to be about 8 µs, and the bidirectional data transfer time for 6 hops across all 12 cells, based on 6.25 Gbps communication, is estimated to be about 4 µs, leading to an expected total system latency of about 12 µs a dramatic improvement compared to the existing VME-based system (about 1200 µs). This is expected to make it possible, in conjunction with a Fast PS, to achieve a suppression bandwidth above 100 Hz.

† wscho@postech.ac.kr

## FOFB SYSTEM ARCHITECTURE

The overall system architecture of the FPGA-based FOFB network at PLS-II is shown in Fig. 1.

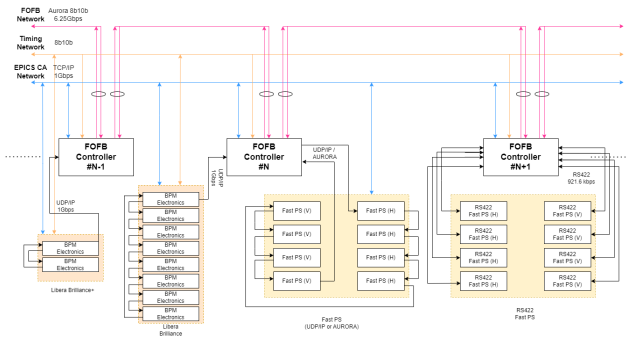



Figure 1: Architecture of the FPGA-based FOFB network at PLS-II.

The FOFB network is configured in a ring topology, and each controller performs bidirectional ring communication with its neighboring controllers. For synchronization, an event code decoding logic was implemented using the SINAPv2 event timing network, and the controller is designed to operate upon the input of specific event codes. At present in the PLS-II storage ring, Libera Brilliance and Libera Brilliance+ are both in use as BPM electronics, and RS-422 Fast PS devices are in actual operation. For the Fast PS required for future high-performance FOFB operation, the figure shows a UDP/IP or Aurora protocol based Fast PS.

To validate the FPGA-based FOFB system, a testbench was configured to connect BPM electronics and the Fast PS. BPM electronics used a Libera Brilliance+, and the Fast PS interface was designed to support switching between a UDP-based CAENels Fast PS and an RS-422 Fast PS through EPICS control (see Fig. 2).

In addition, to expand the FPGA communication interfaces, an FPGA Mezzanine Card (FMC) module compliant with the ANSI/VITA 57.1 standard was used. To provide additional SFP+ ports for the network configuration, a 10-port SFP+ module from HiTech Global [3] was added, and to control the RS-422 Fast PS, a TFMC684-20R module from TEWS [4] was used (see Fig. 3).

These extensions are intended to provide enough ports needed so that a single controller can process 2 cells in the future.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD027

# EPICS IOC EXTENSION POINTS: OLD, RECENT, AND PROPOSED*

A. N. Johnson†
Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The EPICS Input/Output Controller (IOC) has always been extensible, enabling applications to add functionality without modifying the core software. Since the early EPICS 3.14 releases in 2002 the Core Developers have introduced ten new extension mechanisms that IOC application developers can use in individual IOCs or shared support modules. This paper reviews the plugin interfaces available in EPICS 7.0.9 and suggests a couple of areas where new extension points could be added in the future.

## INTRODUCTION

EPICS (the Experimental Physics and Industrial Control System [1]) is a set of open-source software tools, libraries and applications used worldwide for building distributed soft real-time control systems for particle accelerators, telescopes and other large scientific instruments. From its beginnings in the early 1990s, the IOC software was designed for extensibility, so developers could add the site-specific functionality needed at their installation without changing the core [2].

### *Extensibility*

To support extensibility, the common IOC components use standard interfaces so new parts can connect to existing ones without needing special glue code. Network links between IOCs use the same protocols and libraries that Graphical User Interfaces (GUIs) and other client programs use to access IOC data.

Over time the IOC's interfaces have been extended, and new ones added to support user needs and provide extra functionality. This paper highlights the plugin interfaces that exist, showing where the IOC can be extended without modifying the core software at all, and suggesting where some additional interfaces could be added.

## HISTORY OF IOC EXTENSIONS

### *Early Releases*

Versions of EPICS released before 3.14 in 2002 were configurable and extensible in several ways. The most frequently used was the ability to add new device support to the existing record types, allowing IOCs to be able to read and write to many kinds of hardware. Some sites created new record types for locally tailored functionality, but most found it quicker to combine sets of standard records (using existing hardware device support) than to develop new record types in C.

The following extension points existed prior to the development of the EPICS 3.14 release series:

- Database Definition (DBD) file: The IOC discovers component types from this configuration file.
- Database scanning: IOC-specific periodic scan rates, software events, and I/O event scanning.
- Database structures: Record, device and driver support from Base, support modules or IOC applications.
- User-defined breakpoint unit conversion tables.
- User C code called from subroutine records.
- Hook routines called during iocInit.
- Registerable error status codes with descriptions.
- Register a listener for error log messages.
- Locally added IOC and application templates.

### *EPICS 3.14: IOC Portability*

The primary goal of the early 3.14 releases was to port the IOC code from VxWorks to run on other Operating Systems (OSs) [3-5], initially RTEMS, Solaris, Windows NT and Red Hat Linux. At the time most IOCs were based on VME and similar bus standards that those OSs didn't support, so no attempts were made to port device drivers.

Many changes replicated features provided by VxWorks, but with a generic interface. For example, the IOC Shell [6] replaced Wind River's Target Shell, adding an extension point for plugins to register commands and support for access to variables used for debugging and configuration.

### *IOC Construction and Configuration*

To link all the necessary components into an IOC executable, the 3.14 build system generated C++ source code from information in the IOC's DBD file. That file specifies the components to be included, giving symbol names of the functions or data structures needed to access each component.

Table 1 shows examples of every DBD statement type that provides symbolic information, and the EPICS version that introduced it.

Table 1: DBD File Statements With Symbols

| DBD Statement | Introduced |
|---|---|
| `recordtype(mbbi) { … }` | < 3.13 |
| `driver(drvAsyn)` | < 3.13 |
| `device(ao, INST_IO, devAoStats, "IOC stats")` | < 3.13 |
| `registrar(rsrvRegistrar)` | 3.14.2 |
| `variable(dbRecordsOnceOnly, int)` | 3.14.3 |
| `function(generalTimeCheck)` | 3.14.6 |
| `link(const, lnkConstIf)` | 3.16.1 |

* Work supported by the U.S. Department of Energy Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.

† Andrew Johnson <anj@anl.gov>

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD029

# AUTOMATED SAMPLE IDENTIFICATION AND REGISTRATION SYSTEM FOR THE MOGNO BEAMLINE AT SIRIUS

L. E. P. Vecina*, L. B. C. Campoi, G. S. R. Costa, G. R. de Lima, G. R. dos Santos, N. L. Archilha
Brazilian Center for Research in Energy and Materials (CNPEM), Campinas, Sao Paulo, Brazil

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

MOGNO is a tomography beamline at the Sirius Light Source with two experimental stations, both equipped with robotic arms and automated sample magazines. Managing hundreds of samples under fast-measurement and high-throughput conditions requires reliable traceability and minimal user intervention. To address this need, an automated sample identification and registration system was introduced, enabling consistent association of samples with their metadata throughout preparation, acquisition, and reconstruction. Identification is based on Data Matrix codes read by a fixed scanner and robotic arm, while registration and cataloging are handled through a FastAPI service backed by PostgreSQL and integrated with the beamline control environment. By automating these processes, the system reduces operator errors, simplifies experiment organization, and supports large-scale tomography measurements with improved efficiency and reproducibility.

## INTRODUCTION

MOGNO is a micro- and nano-tomography beamline at the Brazilian Synchrotron Light Source, Sirius. It performs fast tomographies with tender (22 keV, 39 keV) and hard (67 keV) X-rays supports spatial resolutions down to 120 nm. A broad set of acquisition modes is supported, including classical tomography, 4D (time-resolved) acquisitions, continuous zoom (variable magnification) and high-throughput operation [1]. Two experimental stations are currently available: a nano-station, open to external users, and a micro-station that is in the scientific commissioning phase. Both stations are equipped with automated sample exchange systems based on robotic arm manipulators and magazines capable of holding 21 and 88 samples, respectively.

Handling and tracking large numbers of samples is operationally challenging, particularly because a typical acquisition at MOGNO may be completed in a matter of minutes and therefore produces a high throughput of raw data and reconstructed images. Historically, users were instructed to record manually which sample was placed in a given holder and in which magazine slot it was stored. This manual approach proved to be error-prone and heavily dependent on users' attention and organizational discipline. Furthermore, beamline staff lacked efficient tooling to manage defective sample holders (for example, holders exhibiting mechanical misalignment, damaged or unreadable Data Matrix codes, or other physical defects).

To address these operational shortcomings, an automated sample identification and registration system was developed. The system was designed to improve sample traceability by mapping each sample from registration through acquisition and post-processing, associating reconstructed datasets with the corresponding sample names and metadata. From the staff perspective, the system provides mechanisms to track holders currently in use and to manage holder availability and health, thereby reducing the risk of failed acquisitions and limiting experimental downtime.

## SERVICE ORIENTED ARCHITECTURE

During system design two primary requirements were identified. First, the solution needed to integrate seamlessly with the existing experiment orchestration and post-experiment processing pipeline. Second, the system had to be optional from the user's perspective: users must retain the choice of whether to register a sample, and the experiment orchestration must tolerate both registered and unregistered samples. This second requirement follows from typical user workflows — for example, when only a small number of samples will be measured or when the user's objective is to collect statistical data across many similar samples without associating each measurement with a distinct sample name.

The MOGNO experiment orchestration ecosystem exposes user functionality via PyQt (Python binding for the Qt cross-platform C++ framework) graphical clients. These clients are backed by Python control packages that implement the device and experiment logic; EPICS (Experimental Physics and Industrial Control System) soft IOCs (Input/Output Controller) are used to expose temporary beamline state; and dedicated APIs (Application Programming Interfaces) run on a separate server to provide operation support [2]. Given this ecosystem, a service-oriented architecture (SOA) based on a Relational Database Management System (RDBMS) and a CRUD API (service with fundamental functionalities to manage data: create, read, update and delete) was selected because it enabled straightforward integration with existing control software while satisfying the optional usage constraint. The chosen architecture therefore exposes a registration/cataloging interface for pre-experiment use and provides API endpoints consumable by the beamline experiment control software at runtime.

### *Database*

Schema design began by identifying the minimal set of fields required to track a sample across the experimental workflow. The initial minimum viable product (MVP) used a single table in a local SQLite database containing only the sample name and the sample-holder Data Matrix label. As requirements matured — notably the need for more complex

* lucas.vecina@lnls.br

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD031

# EPICS-IN-DOCKER: A SMALL FRAMEWORK FOR BUILDING SLIM IOC AND EPICS TOOLING CONTAINER IMAGES

É. N. Rolim*, G. R. Lima, G. S. Reis, H. F. Simões
Brazilian Synchrotron Light Laboratory (LNLS), Campinas, Brazil

*Abstract*

The SIRIUS accelerators have used containers for Input/Output Controllers (IOCs) for years, but build definitions and launch scripts were often duplicated, and image sizes could be over 3 GB. On the other hand, the SIRIUS beamlines, until recently, used IOCs installed in a shared Network File System (NFS), which complicated application management, especially across different Operating System (OS) versions.

To address these issues, we have developed a framework for building slim IOC container images (e.g. ADAravis takes 300 MB) using a curated set of dependencies (and their versions) and simple and short build definitions. We avoid duplicating shared information by using git submodules, which aids in versioning the base images used. The resulting container images include a standard set of installed packages and scripts, making them ready for deployment in a wide range of container orchestration setups. The shared interface provided by the EPICS build system allows us to also create images with EPICS tools, including Channel Access (CA) and PVAccess (PVA) gateways and epics-base utilities.

For beamlines, it was necessary to adapt the IOC orchestration to also support containerized applications, keeping the same user interface for managing IOCs for beamline and support staff.

This article aims to explain the epics-in-docker architecture, the user experience, and how SIRIUS manages containers. It also aims to explore the different tradeoffs made in epics-in-docker and other frameworks, such as our choice to not support different versions of dependencies.

## INTRODUCTION

SIRIUS is the Brazilian 4th generation synchrotron light source, which has been in operation for external users since March 2023. Active operation imposes a priority ranking, where availability ranks highest; other activities, such as control system and infrastructure maintenance and updates, must adapt to the operation schedule. This is further complicated by the heterogeneous environment, where the SIRIUS accelerators and 14 beamlines (10 fully available to users) [1] can have different requirements, priorities and maintenance schedules.

The control system used both by the SIRIUS accelerators and beamlines is based on EPICS, but their infrastructure differs significantly. The SIRIUS accelerators use a small amount of servers which centralize all possible services for all accelerators (some devices can't be directly controlled over the network and require a field server), while SIRIUS beamlines have mostly independent infrastructure per beamline, with a Virtual Machine (VM) for most services, and field servers wherever local device access is required. This leads to a great amount of servers to be managed.

Given this situation, being able to decouple application updates from infrastructure updates can greatly simplify these activities, allowing them to occur in tighter time windows, and with a smaller impact to operations. This been done in the SIRIUS accelerators for years, by using *Docker Swarm* [2] to launch and monitor EPICS Input/Output Controllers (IOCs) inside containers across multiple hosts. This brings many benefits: isolated environments simplifying dependency management and decoupling applications from the host OS; a broad ecosystem of management and monitoring tooling; the possibility of setting up new hosts by simply installing a container orchestration tool.

Historically, however, the accelerator and beamline infrastructure and the applications running on it have been managed and developed by many disjointed teams, with little communication between them. SIRIUS beamlines traditionally used IOCs installed into a shared Network File System (NFS), which had some disadvantages, since application updates were disordered global events, and servers running different OS versions using a same IOC could face dependency issues.

Since November 2022, these responsibilities have been centralized into one team (two teams since November 2023), enabling developers to apply similar standards to accelerators and beamlines and develop common solutions. One of these solutions was the epics-in-docker [3] project, which generated base images and defined simple build specifications to easily containerize any IOC, simplifying container deployment to the SIRIUS accelerators and introducing containers to the SIRIUS beamlines.

## PRIOR DEVELOPMENTS

The container images previously used at SIRIUS accelerators [4] were based on a multi-layer architecture, with separate layers for dependencies, EPICS Base [5], synApps [6] modules, and site-specific configuration. This container build scheme featured multiple branches, and aimed to support multiple base Linux distributions, EPICS Base versions, and synApps versions, leading to a combinatorial explosion of supported base images. This meant IOC integrators had to choose the right set for their IOC, and updates depended on a concerted effort to define an upgrade path, since there were so many possibilities. Even though the usage of layers should amortize the storage requirements for different images that shared a common base, these savings weren't

* erico.rolim@lnls.br

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD033

# THE TANGO AlarmHandler: ADVANCEMENTS IN CORE FUNCTIONALITY AND TOOLS

G. Scalamera*, L. Pivetta, G. Strangolino, Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy

*Abstract*

The AlarmHandler system is a key component for ensuring operational safety and efficiency in complex control systems. Key updates include improved support for array data types within the alarm evaluation logic, enabling more sophisticated and flexible condition definitions directly involving array or matrix data. Furthermore new tools, designed to extend the AlarmHandler's reach and usability, are now available. A dedicated notification service allows for configurable, multi-channel alerting, such as email and messaging platforms, facilitating timely operator awareness and response. Complementing this, new management utilities have been created to streamline the configuration, deployment, and maintenance of alarm definitions across distributed systems, significantly simplifying administrative tasks. This contribution details the architectural changes, implementation specifics, and the benefits these advancements bring in terms of system robustness, operator efficiency, and overall monitoring capability.

## INTRODUCTION

Modern large-scale scientific facilities, from particle accelerators to astronomical observatories, depend on sophisticated control systems that monitor and manage thousands of devices, sensors, and subsystems operating continuously. The complexity and scale of these installations create critical challenges in maintaining operational efficiency, preventing equipment damage, and ensuring optimal performance. Within this context, alarm management systems are a fundamental tool, providing real-time monitoring capabilities that detect anomalous conditions and alert operators to potential issues before they escalate into serious problems. The TANGO Controls framework has established itself as a widely adopted solution for distributed control systems across many scientific facilities worldwide. As part of this ecosystem, the AlarmHandler system [1] plays a pivotal role in maintaining operational integrity by continuously evaluating alarm conditions based on device attributes and system states.

More and more often control systems rely on array-based and matrix data from advanced sensors, high performance detector systems, and multi-channel devices. Multi-dimensional data led to the development of array support in alarm formula evaluation. The enhanced system introduces native support for array data types within alarm evaluation logic, enabling direct processing of multi-dimensional data and sophisticated condition definitions that were previously impossible or required complex workarounds.

Furthermore, the operational demands of 24/7 scientific facilities require robust and flexible notification systems that extend beyond simple console alerts. Operators rely on multi-channel communication capabilities that can deliver contextual information through various platforms including email, messaging services, and mobile applications. At the same time flood protection mechanism is a key feature to prevent important notifications from being lost when too many are received.

Managing alarm configurations across distributed systems is a significant operational challenge. As structures become more complex, maintaining consistency, implementing updates, and ensuring adequate alarm coverage across hundreds or thousands of devices becomes increasingly difficult without dedicated management tools and automated implementation mechanisms. Significant advancements have been introduced to the TANGO AlarmHandler system that address this operational scenario: a comprehensive notification service architecture provides configurable, multi-channel alerting capabilities that improve operator awareness and response times.

Finally, graphical interfaces have been renewed, with updated widgets to improve overall usability, and new management utilities have been developed to simplify the entire alarm life-cycle, from initial configuration to deployment and ongoing maintenance.

## THE AlarmHandler: SYSTEM OVERVIEW

The TANGO AlarmHandler represents a sophisticated, event-driven alarm management system that has evolved significantly since its initial implementation. Based on the EN IEC 62682 [2] standard, it provides a robust foundation for alarm management in distributed control systems. The AlarmHandler is implemented as a modular TANGO device server, written in C++, designed around an event-driven architecture that exploits TANGO publish/subscribe capabilities. The system operates asynchronously, triggering alarm evaluations only when relevant attribute values change, ensuring optimal performance and resource utilization.

A critical component is the formula parsing engine, implemented using the Boost Spirit Parser Framework. This choice provides an object-oriented recursive descent parser that builds an Abstract Syntax Tree (AST) for each alarm rule. Each formula is parsed once at device initialization, with subsequent evaluations performed using the optimized AST representation, resulting in superior performance compared to traditional parsing approaches.

* graziano.scalamera@elettra.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD035

# AN OVERVIEW OF THE FGC4 – CERN'S NEW POWER CONVERTER CONTROLLER

D. Zielinski*, R. Murillo-Garcia, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The CERN's Electrical Power Converters group manages over 5000 power converters, 4000 of which are controlled, monitored, and diagnosed by a few generations of the Function Generator/Controller (FGC) devices. However, the current generation (FGC3) is now facing performance limitations and component obsolescence. To address this and accommodate future installations at CERN and other laboratories, a fourth generation of FGC is under development. Built with cutting-edge technology and modern standards, the FGC4 features a Linux-based System-on-Chip (SoC), delivering an order-of-magnitude improvement in regulation rate, extensive configuration options, and significantly enhanced diagnostics. While designed to fit CERN's accelerator control system, reusability beyond CERN has been a core design principle from the outset, enabling compatibility with EPICS and TANGO frameworks. This paper provides an overview of the FGC4 project, with a primary focus on its software architecture and highly modular design, which facilitates extensibility and ensures a future-proof solution. Additionally, it discusses the hardware architecture, including a CERN-developed System-on-Module hosting a Xilinx SoC.

## SYSTEM OVERVIEW

The FGC4 is a modular control system that combines a Linux-based System-on-Module (SoM) with application-specific I/O to regulate, monitor, and diagnose power converters. An FGC4 "cassette" contains three boards: the SoM, a digital board that is also a SoM carrier, and an analog board that implements conditioned analog I/O. Cassette's connectors (both from the digital and analog boards) plug into a passive backplane; from there, a backpanel (or a combined backplane/backpanel) routes signals to the power converter via Burndy connectors. The crate also hosts a dedicated FGC4 PSU, optional x2V measurement cards (I2V current-to-voltage and V2V voltage-to-voltage scaling), and when needed, I/O adapter boards that translate FGC4 I/O to converter-specific formats (e.g., optical).

Networking is split by role. One RJ45 port on the front panel connects the cassette to the accelerator technical network for controls and management. In addition, four SFP cages provide deterministic, low-latency links for the FGC private network and optional White Rabbit [1]. A USB-C debug port and a special 15-pin Sub-D connector for debugging and analog channels calibration are also available on the cassette. At system scale, the backplane remains passive (except for identification); the PSU provides power and health monitoring (PMBus), and cards are geo-addressed for unambiguous system identification. The mechanical concept follows IEEE-1101 / CompactPCI-Serial practices (3U cards/cassettes with front panels/handles). On the software side, FGC4 runs a Linux front-end directly on the device (two Cortex-A53 cores). Two additional A53 cores execute bare-metal real-time control; one dedicated to voltage control and the other to current or field control. Deterministic exchange between the FPGA programmable logic (PL) (ADCs, PWM, digital banks) and the A53 cores, known as the programmable system (PS) is achieved with a custom Data-DMA path into PS level 2 (L2) cache.

Boot and serviceability are aligned with CERN's platform: in normal operation the kernel, rootfs, and the main application are network-booted. In case of failure at this stage, a reliable, local "golden" image is used. Standard front-end computer monitoring systems and tools are used to track system-level health and parameters.

## HARDWARE ARCHITECTURE

### *Backplane, Backpanel, and Rack Integration*

FGC4 systems are built around a passive 3U backplane that distributes power, management buses, and user signals between the cassette(s), the x2V measurement cards, and any I/O adapters. A separate backpanel (or a single PCB that combines both roles) brings those signals to converter-side Burndy connectors and field cabling. The PSU is installed to the left-most slot to simplify routing of power entry. The architecture scales from compact one-cassette controls to multi-slot systems (FGC4 + PSU + x2V + up to eight I/O adapter cards), and supports special layouts (e.g., tall/narrow sub-racks) when converters require it.

The point of specification for user I/O is at the cassette's rear connectors to the backplane; from there the backpanel reaches the converter or, for small integrated systems, the backplane itself may double as the backpanel.

Management and identification are handled on the backplane via a PMBus segment and on-board ID devices; geographic addressing rules and reserved addresses ensure predictable discovery and monitoring. Differential interfaces have length and impedance matching.

### *The FGC4 Cassette: SoM and Carrier*

The SoM hosts the heterogeneous AMD/Xilinx Zynq UltraScale+ MPSoC [2] that combines a quad-core Arm Cortex-A53 processing system (PS) and large FPGA programmable logic (PL). At CERN, a cross-departmental project called *ATS SoM* is being developed to provide a well-defined, reusable SoM. In this particular module the specified device is XCZU17EG-1FFVC1760E (A53 at 1.2 GHz), mounted on a module no larger than 94 × 94 mm,

* dariusz.jakub.zielinski@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD036

# VALIDATING VSLIB, THE VOLTAGE SOURCE CONTROL LIBRARY USED IN THE FOURTH GENERATION OF FUNCTION/GENERATOR CONTROLLERS AT CERN

D. Arominski*, R. Murillo, D. Zielinski, CERN, Geneva, Switzerland

*Abstract*

The fourth generation of power converter control at CERN, known as Function Generator/Controller 4 (FGC4), is currently under development. The chosen hardware is based on a quad-core A53 ARM-architecture CPU within an AMD Zynq UltraScale+ MPSoC System-on-Chip (SoC), where one core is dedicated to voltage source control. This architecture necessitates a new approach to voltage control software, previously implemented on a Digital Signal Processor (DSP) card and an FPGA, where high reliability and performance are crucial.

To achieve these goals in a highly integrated environment, a new library called VSlib (Voltage Source library) has been developed. VSlib serves as a toolkit, providing all necessary building blocks for user-defined voltage regulation algorithms, such as filters, controllers, and lookup tables. Additionally, it supports digital and analogue signal logging, inter-core data exchange, and communication with other FGC4s using an in-house Ethernet-based protocol.

The development process was test-driven, focusing on performance, determinism, and reliability. The library adheres to best industrial practices, including version control, static analysis, and automated testing. Tests were conducted against power convert models running on a Speedgoat Hardware-in-the-Loop system. The application of the library in a real-time hardware-in-the-loop representation of the POPS+ power converter system at CERN is discussed along with the commissioning process.

## INTRODUCTION

The power converter controls software (CCS) section at CERN manages control software for over 5'000 power converters spread out over the CERN's accelerator complex. The main products are Function Generator/Controllers (FGC), embedded computers allowing for real-time control of power converters with the first generation having its roots in the 1990s, and the most recently used, FGC3, can no longer be procured due to component obsolescence. With the new challenges in terms of the control frequency, of up to 100 kHz, the fourth generation of FGC, henceforth referred to as FGC4, is currently in final steps of its development. The CCS team uses this opportunity to upgrade the software architecture for the new platform [1]. The FGC4 hardware is based on AMD Zynq Ultrascale+ MPSoC, featuring quad-core Cortex A53 CPU that the software development is targeting. Two cores are dedicated to running Linux OS and serve, among other tasks, to communicate with CERN IT infrastructure, and two cores run without an operating system ("bare metal"). One will run current or magnetic field regulation, while the other voltage source control.

The Voltage Source control library (VSlib) is a modular toolkit library for power converter experts to develop their own regulation and monitoring workflows in hard-real time. It is developed in modern C++20, in a test-driven way and according to the needs and requirements of the power converter community at CERN. To achieve the required determinism, the library is intended to run bare-metal on core #3 of the FGC4, a decision that posed a number of challenges on the design and development of the library. Users are intended to use the library modules to build their own bespoke workflows, called *vloops*.

The modules responsible for the communication between the *vloop* and the two cores dedicated to running Linux, e.g., to allow the user to configure their binary at run-time, has already been described in Ref. [2].

### *Implemented Components*

The library implements a number of classes intended for users to be used in their control workflows. The interdependency of runtime-settable Parameters of Components, the Components themselves, including examples of a class derived from it (*FIRFIlter* and *RootComponent*) are shown in Fig. 1.

The Components implemented in VSlib fall to one of the following categories:

- filters
- limits
- controllers
- electrical transforms
- look-up tables
- hardware abstraction layer interfaces
- application specific

All of these classes are derived from the VSlib's Component class and share its interface, with optional runtime-settable Parameters. Each class implements a single interaction method to streamline user experience and provide an intuitive interface.

The FGC4 platform already contains first-order finite-impulse (FIR) and infinite-impulse (IIR) filters on analogue inputs implemented on the FPGA side. The filters implemented in VSlib are slower than those in the FPGA but offer more flexible interactions, namely the order of the filters as well as the filter coefficients can be specified by the user. These filters include three Components: *BoxFilter*, *FIRFilter*, and *IIRFilter*. Each of them implements a single interaction method: *filter*, taking one argument.

* dominik.arominski@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD038

# ENABLING HIGH-PERFORMANCE PLC COMMUNICATION THROUGH OPEN STANDARDS: OPC UA PUBSUB

L. Gutierrez Prendes*, L. Milicevic, A. Schifferer, B. Schofield, B. Farnham
CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The growing diversity of PLC models and brands in industrial controls systems is increasing the complexity of the communications at the supervisory (SCADA) and control layers. At CERN, the in-house framework UNICOS-CPC manages this diversity through the integration of proprietary and open communication protocols, as well as bespoke implementations at the SCADA level to manage the incoming process data received by different communication drivers, including Modbus, S7, and S7Plus. To reduce this complexity, this paper proposes unifying PLC-SCADA communications across all platforms using OPC UA PubSub, a lightweight and highly performant publisher-subscriber protocol specified in the IEC 62541 industrial standard. This approach simplifies integration with new vendors and technologies, while enabling direct communication between PLCs. It positions OPC UA once more as a homogenizing middleware layer on top of heterogeneous hardware, which has proven to be a reliable and scalable solution for other use cases at CERN, such as power supplies, powered crates and custom electronics. The paper outlines the design, prototyping and testing phases involved in integrating the OPC UA PubSub protocol into industrial applications. It also presents the challenges encountered in the integration process, and concludes with the promising results achieved in both PLC-PLC and PLC-SCADA communication setups.

## INTRODUCTION

At CERN, PLC-based industrial control systems are usually designed and implemented using the UNICOS-CPC framework [1], which supports different PLC models, including Siemens S7-300s, S7-400s, S7-1500s and most Schneider PLCs. This leads to a high diversity of communication protocols to transfer process data from and to the Supervisory Control and Data Acquisition (SCADA) systems, that have accumulated in the framework throughout the years, according to the needs at the time of implementation. With Schneider PLCs using Modbus, and Siemens PLCs using the proprietary protocols S7 and S7Plus, the configuration of the PLCs and the SCADA systems is strongly tied to the hardware that is used. The ultimate goal of the framework is to unify the platforms and abstract the user from the specific device in use. However, this diversity of underlying protocols increases the complexity of the framework and hinders its maintainability.

Moreover, the use of proprietary protocols can complicate maintenance due to limited access to software and dependency on vendor development cycles. In the case of Siemens PLCs in UNICOS-CPC, the S7 and S7Plus protocols are used with the CERN-made Time Stamp Push-Protocol (TSPP) [2], enhancing the overall performance of the communications. However, TSPP needs custom implementations in the Siemens SCADA software used at CERN (WinCC OA). This imposes a strong dependency on external tools and can lead to costly integration with alternative SCADA layer technologies, involving a complete refactoring of the PLC communication routine.

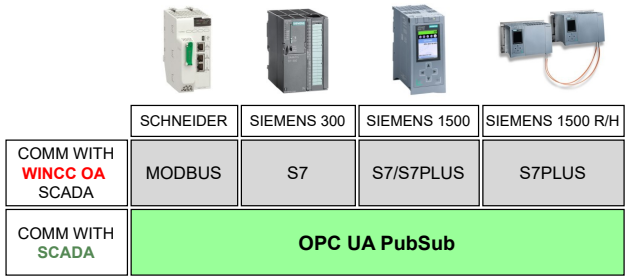



Figure 1: Unification of PLC communication protocols in UNICOS-CPC.

The aim of this paper is to present a prototype that uses an open communication protocol, with the goal of ultimately unifying the communication mechanisms for all PLC models supported by the framework, as depicted in Fig. 1. This change would also open the path to new features, such as data-sharing between PLCs, and further decouple the framework from external dependencies.

## LIMITATIONS OF THE CURRENT SOLUTIONS: WHY OPC UA PUBSUB?

Traditional PLC communication mechanisms primarily rely on client-server architectures, where a supervisory system actively polls PLCs for data using protocols such as Modbus TCP or OPC UA Client/Server. Other legacy protocols used at CERN solve some of the problems found in these solutions, which are discussed later in this section, but are proprietary and therefore limited to a specific manufacturer. While adequate on a small-scale, these methods become restrictive in modern industrial environments that evolve towards real-time, many-to-many communication, hence the growing need for more flexible, scalable, and decoupled communication paradigms.

To date, efforts at CERN to implement open standards in PLC-SCADA communication have focused on OPC UA Client-Server, with the OPC UA server running on the PLC and the SCADA client subscribing to the data. This solution

* loreto.gutierrez.prendes@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD039

# CANModule: A LIGHTWEIGHT, VENDOR-NEUTRAL CAN BUS ABSTRACTION LIBRARY FOR SIMPLIFIED INTEGRATION AND DIAGNOSTICS

L. Miguens Fernandez, B. Farnham, G. Thomas, M. Ludwig, P. Moschovakos
CERN, Geneva, Switzerland

*Abstract*

This paper presents CANModule, an open-source cross-platform library designed to provide a unified abstraction layer for vendor-specific Controller Area Network (CAN) bus implementations. It supports Ethernet-CAN gateways from Analytica (GmbH) and Linux's SocketCAN out of the box, and an open architecture for extending support to additional CAN gateways. CANModule offers a lightweight, framework-independent solution, requiring only standard C++17, distinguishing it from alternative approaches such as Qt CAN Support, which introduces additional dependencies. CANModule standardises CAN communication through a generic API available in both C++ and Python, thereby reducing the complexity associated with integrating vendor-specific APIs. Additionally, built-in diagnostic tools offer functionality, similar to SocketCAN's *canutils*, with the key advantage of supporting multiple vendors seamlessly across multiple operating systems. This approach enhances interoperability and simplifies software development for heterogeneous CAN environments. CANModule is successfully integrated into CERN's Quasar Framework, where it has facilitated the development of numerous OPC UA servers used in control systems for large-scale scientific experiments and infrastructure. Deployed in applications such as detector control within the ATLAS experiment and power supply control, CANModule demonstrates reliability and versatility in demanding operational environments. Beyond CERN, the library presents potential applications in industrial contexts (e.g. automotive, robotics etc.), offering a scalable and extensible solution for CAN-based systems. By abstracting vendor-specific complexities, CANModule streamlines CAN bus integration, providing a flexible, dependable and efficient foundation for both research and industrial use cases.

## INTRODUCTION

The Controller Area Network (CAN) bus is a field-proven standard for reliable real-time communication, originally developed for automotive systems but now common in many industrial and research applications. As CAN adoption grew, a variety of hardware interfaces and vendor-specific APIs emerged, making cross-platform development challenging. Engineers often face incompatible software libraries for different CAN devices – for example, Linux's SocketCAN interface versus proprietary APIs for Ethernet-to-CAN gateways like AnaGate. Bridging these disparities in a consistent, high-level API can greatly simplify development and integration of CAN-based systems.

CANModule is a lightweight, cross-platform library created to address this need [1, 2]. It abstracts vendor-specific CAN implementations (such as AnaGate's TCP/IP CAN interface and Linux SocketCAN) behind a unified interface, accessible in both C++ and Python. The library requires only modern C++17 features and minimal dependencies, making it lean and easy to integrate into existing projects. CANModule originated as an optional module for CERN's quasar framework [3], providing plug-in CAN integration for OPC-UA server applications, however CANModule can be used as a stand-alone component without quasar contexts. This paper provides an in-depth look at CANModule's design and features, highlighting how it simplifies multi-vendor CAN development. Real-world use cases from CERN's accelerator and experiment control systems are discussed, alongside potential applications in industry. The paper is organised into sections covering the CANModule architecture, built-in diagnostic tools, integration experiences, and example deployments, followed by considerations for wider industrial usage.

## NEED FOR A UNIFIED CAN ABSTRACTION

Multiple vendors provide CAN bus hardware – including, for example, PCI/USB adapters and IP networked gateways – each with their own API or driver. For instance, Linux offers the SocketCAN subsystem as a common interface for CAN devices, however Windows provides no such equivalent, furthermore vendors like Analytica GmbH (AnaGate) supply proprietary APIs for their CAN-over-Ethernet gateways. Without a unifying layer, developers are forced to write code specific to each CAN interface, increasing complexity and maintenance burden. Traditionally, one solution is to use Qt's SerialBus module (QCanBus) which supports various CAN plugins [4]. However, adopting Qt for CAN integration only may be overkill in projects aiming to remain lightweight or framework-agnostic. An abstraction, ideally independent of large frameworks, is therefore highly valuable.

CERN's control systems exemplify this challenge. Different experiments and facilities utilise diverse CAN hardware – from Linux servers using SocketCAN to Windows machines with Ethernet-CAN bridges – all needing to run uniform control software. Previous solutions required extensive manual testing for each change and had complex configuration scripts, making it difficult to automate and maintain. This environment drove the development of CANModule, a fresh abstraction layer focusing on multi-platform, multi-vendor support with minimal dependencies and straightforward configurability. By standardising how CAN devices are accessed in code, CANModule frees

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD041

# CONTROL SYSTEM DESIGN FOR THE NEW SMH16 PULSED CURRENT GENERATOR IN THE CERN PS EXTRACTION SYSTEM

O.Y. Yagci, L. Strobino, N. Magnin, G. Gräwer, M.C. Bastos, N. Beev, C. Boucly
CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The existing direct-drive septum magnet (PE.SMH16), in operation since 1994, is reaching the end of its lifetime under increased extraction frequency and will be replaced by a new eddy-current septum magnet, requiring a redesigned 30 kA pulsed generator. To meet the demanding flat-top stability requirement of ±0.05% over an 11 µs window and ±0.05% pulse-to-pulse repeatability over a year, a dedicated regulation control system was developed and validated. The objective of this work is to demonstrate that the control system achieves the required closed-loop performance, delivering repeatable magnet current waveforms under representative test bench conditions. The control system employs real-time regulation of flat-top amplitude and flatness, supported by a thermally stabilized enclosure to mitigate acquisition drift. Measurements confirm that the closed-loop system consistently maintained flatness requirement over an 8 µs window, with long-term repeatability and all observed deviations, including spurious glitches up to 400 ppm, remaining within specification. The restriction of 8 µs stems from test-bench limitations rather than control capability, which is discussed in detail. These results show that the regulation system fulfils its role up to specification; final confirmation of absolute flat-top accuracy will come from further qualification of acquisition chain elements in a dedicated test campaign, followed by beam-based validation in operation.

## INTRODUCTION

The direct-drive septum magnet (PE.SMH16) is being replaced due to the reduced lifetime under increased extraction frequency. Its successor, an eddy-current septum [1], requires a high-current pulsed generator capable of delivering up to 30 kA with a flat-top stability of ±0.05% during a pulse and with pulse-to-pulse repeatability of ±0.05% throughout the year [2]. Meeting these stringent requirements, together with the technical challenges described in this paper, necessitates a dedicated regulation control system.

Beam extraction in the septum magnet relies on the magnetic field generated by eddy currents in the magnet, induced by a high-current pulse from the pulsed generator. The generator consists of two capacitor banks, a power inductor, two switch assemblies, and snubber circuits [2]. The high-voltage DC power supply (DCPS) charges the two capacitor banks, which in combination with the magnet's impedance generate the principal 2.5 kHz sine wave and its 3$^{rd}$ harmonic. These banks are discharged through the switches, allowing to adjust the trigger time of both stacks. The sum of these two signals form the magnet current with an elongated flat-top. Parameters such as the amplitude reference and switch delays are controlled through real-time software, using the FESA framework [3], in which instances of real-time software are conventionally referred to as "devices" in the accelerator controls domain.

The magnet current is measured with a Commercial-Off-The-Shelf (COTS) current transformer (CT) - see further - and acquired by a COTS digitiser card, the Spectrum M2p.5933-x4 [4]. To meet the ±0.05% stability requirement throughout the year, mitigating acquisition errors, such as the drift induced by ambient changes, is imperative. Therefore, a thermally-regulated enclosure was designed, employing a COTS thermo-electric controller (TEC), the Meerstetter TEC-1162 [5], which regulates a Peltier element to maintain the temperature within ±0.005% of the reference value on average.

The digitised signals are analysed in real-time using the IPOC framework [6]. The analysis results are propagated to a Programmable Logic Controller (PLC) based State Control and Surveillance System (SCSS) [7], the CERN Controls Middleware (CMW) [8], and to the NXCALS [9] database for long-term storage. The flat-top regulation algorithm is deployed on a UCAP Node [10], a centrally-managed computing platform that allows device access and control.

The control system comprises a waveform analyser, which dynamically calculates the flat-top and computes waveform analytics as well as error metrics based on live device data acquired through CMW. The corrections are then propagated to the real-time devices that manage the hardware settings, e.g. switch delays, as described previously.

## KEY TECHNICAL CHALLENGES

**1. High flat-top stability:** To meet beam quality specification, the flat-top current must remain within ±0.05% of the nominal amplitude over an 11 µs window. Meeting this specification requires active correction of small deviations arising from circuit losses, component tolerances, and measurement errors.

**2. Acquisition precision:** High-precision acquisitions are required to meet the stringent stability specification. Achieving such precision requires high-resolution sampling combined with good repeatability of acquisition chain elements such as the current transformer (CT) and the digitiser.

**3. Pulse-to-pulse repeatability:** With up to two pulses per PS cycle and a nominal repetition rate of 1.2 s, flat-top stability must be maintained during continuous operation throughout the year. Generator induced perturbations are generally limited to a single pulse and remain within beam tolerances, making them acceptable for operation. The control system, therefore, must regulate the long-term, pulse-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD042

# PERFORMANCE CHARACTERISATION OF REAL-TIME SOFTWARE IN C++. A REAL-LIFE EXAMPLE OF DIGITAL CAMERA-BASED ACQUISITION SYSTEMS AT CERN

A. Topaloudis*, E. Balci, S. Jackson, CERN, Geneva, Switzerland

*Abstract*

The performance of real-time software is critical in accelerator control systems, where precision and reliability are essential. This paper presents a method for performance characterisation of real-time software developed in C++, using a digital camera-based acquisition system at CERN as a case study. Key performance metrics, including execution time, latency, memory footprint and network use are analysed to evaluate the system's ability to meet strict real-time and scalability constraints. In addition, the profiling of task execution serves as a means of ensuring that the software continues to adhere to agreed specified behaviour, particularly as the system evolves and undergoes changes. The impact of software architecture, multithreading strategies, and hardware optimization on overall system performance is also discussed. Finally, the tools used to extract key performance metrics are presented, emphasising their generic nature and potential use for other real-time software developed in CERN's accelerator sector.

## INTRODUCTION

Many accelerator complexes around the world, including CERN, divide their control system into a 3-tier architecture [1]. The top tier is usually dedicated to operational and high-level expert client applications, while the middle tier incorporates the business and supervision logic. The lower tier is devoted to real-time C++ software controlling and monitoring the accelerator equipment [2] that span heterogeneous hardware, numerous devices, which often ingest high-rate and high volume data.

Software controlling the accelerator equipment must respond to inputs and events in a predictable and deterministic manner, as a late delivery of data may lead to beam loss, equipment stress or costly downtime. The combination of tight latency constraints, reliability requirements, and scalability demands makes careful performance characterisation of real-time software indispensable.

## DIGITAL CAMERA-BASED ACQUISITION SYSTEM OVERVIEW

In order to demonstrate the impact of software benchmarking and profiling on the long-term stability of the operation of the accelerator complex, the performance analysis of the software controlling digital camera-based acquisition systems was chosen as a case-study.

Since cameras are fundamental in beam imaging systems, they are installed across the whole accelerator complex in various configurations based on the operational needs and specifications of each instrumentation [3]. It is therefore crucial to know the limits of each hardware-software combination to choose the solution that guarantees the desired performance for each new or existing installation.

The acquisition system consists of GbE digital cameras from Basler [4]. They are most often Powered over Ethernet (PoE) via a 10 GbE RUCKUS [5] 12 or 24 port switch. Their trigger is synchronous with the beam passage at a rate that can vary from 0.01 to 35 fps. Their images are sent within a local network to a computer used for data processing that can be one of the following:

- a rack-mounted industrial PC [6]
- a MEN A25 CPU [7] installed in an VME Front-End Computer (FEC)
- a high-performance SuperMicro server [8]

Each platform's specifications are summarised in Table 1.

Table 1: Platform Specifications

| Resource | Industrial PC | MEN A25 | SuperMicro server |
|---|---|---|---|
| CPU | Intel(R) Core(TM) i5-8500 | Intel(R) Pentium(R) | 2xAMD EPYC 7302 |
| CPU clock | 3 GHz | 1.5 GHz | 3 GHz |
| CPU cores | 6 | 4 | 32 |
| CPU threads | 6 | 8 | 64 |
| RAM | 8 GB | 8 GB | 196 GB |
| NIC | 1x1 GbE, 1x10 GbE | 1x1 GbE, 1x10 GbE | 2x10 GbE |

The Operating System (OS) for all three options until now has been CentOS Linux 7 [9], a downstream derivative of Red Hat Enterprise Linux (RHEL) 7) [10]. However, there is an up-coming campaign to migrate the OS for embedded systems at CERN to Debian OS [11] due to OS end-of-life, the associated lack of security updates, and incompatibilities between new hardware (e.g. embedded servers, FECs, etc) and the old OS [12]. The OS migration affects especially the SuperMicro servers which for organisational reasons will not be available with the new OS and will have to be replaced during the third Long Shutdown (LS3 - 2026 - 2029), implying that an alternative hardware must be identified.

The real-time C++ server controlling such systems is designed with the Front-End Software Architecture (FESA) framework [13] and uses the *pylon* Software Development Kit (SDK) [14] from Basler for low-level camera access.

* athanasios.topaloudis@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD043

# PLC INTEGRATOR: A MODERN TOOL FOR PLC-EPICS INTEGRATION AT ESS

A. Fontoura†, S. Jäger, A. Rizzo, H. Ligander, J. Choi, E. Foy, ESS ERIC, Lund, Sweden


## *Abstract*

The *PLC Integrator* is a new tool developed at the European Spallation Source (ESS) to modernize the integration of PLC-based control systems with the EPICS framework. It replaces the legacy *PLC Factory*, removing outdated dependencies such as the CCDB database, while introducing support for modern technologies like OPC-UA and Beckhoff ADS. The software retains key features familiar to integrators - such as automatic PLC-IOC communication code generation, alarm handling, and archiving - but is built on a sustainable, extensible, and easier-to-maintain foundation. By adopting JSON files instead of the older *.def* format, it improves standardization, enables automatic validation, and increases memory efficiency on PLCs. This article presents the context and motivations behind its development, describes the architecture and main features of the PLC Integrator, highlights the challenges faced during implementation, and discusses its impact at ESS as well as the path forward.


## INTRODUCTION

The European Spallation Source (ESS), under construction in Lund, Sweden, is one of the world's largest scientific infrastructure projects, aiming to become a global leader in neutron scattering research. To achieve this, its operation relies on highly robust, scalable, and interoperable automation systems capable of handling the complexity of large-scale scientific instruments. Within ESS, the ICSHWI group (Instrumentation & Control Systems Hardware and Integration), and specifically its Automation Section, plays a central role by bridging Programmable Logic Controllers (PLCs) with the EPICS (Experimental Physics and Industrial Control System) framework [1]. EPICS underpins communication and supervision across thousands of devices. For years, this integration relied on an in-house tool called PLC Factory [2], which simplified tasks such as generating PLC-IOC communication [3] code for Siemens controllers, managing alarms, and supporting device hierarchies. While effective early on, the tool accumulated significant technical debt over time. Its main limitations were:

- Reliance on an outdated version of Python [4].
- Dependence on the CCDB database [5], which was scheduled for decommissioning.

Faced with these limitations, ESS needed a new solution that would ensure continuity, remove obsolete dependencies, and prepare the infrastructure for future challenges. This need gave rise to the PLC Integrator.

## DEVELOPMENT AND IMPLEMENTATION

The *PLC Integrator* was designed to be backward-compatible with existing workflows, while offering completely modernized architecture (Fig. 1).

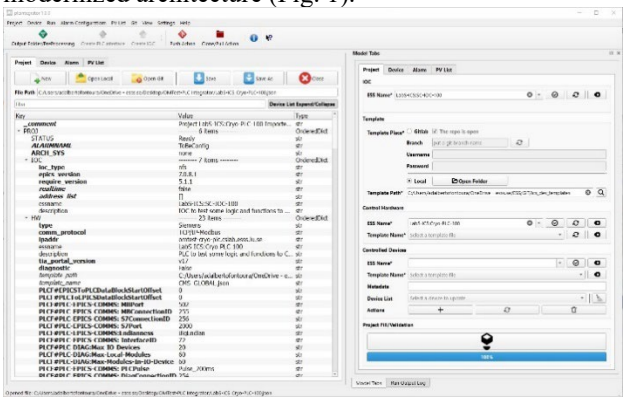


Figure 1: PLC Integrator UI.

### *Interface and Platform*

Its graphical interface was built from the ground up using PySide6 [6], resulting in a modern, responsive, and intuitive application. The goal was to shorten the learning curve for engineers familiar with PLC Factory, while also adding new features and improved usability.

### *Modular Architecture*

All functions were restructured into independent modules (Fig. 2), each responsible for a specific task such as template management, preprocessing, or code generation.

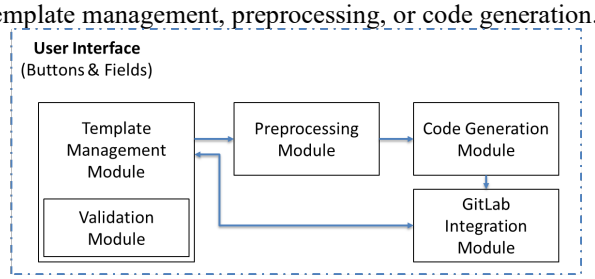



Figure 2: Modular Architecture

- **Template Management Module:** Manage Device Templates and Project Templates.
- **Validation:** Performs real-time validation during configuration, checking user inputs, JSON syntax, and PV consistency, running continuously in the background rather than only at preprocessing.
- **Preprocessing Module:** Validate data before code generation. Detect errors in PV names (missing, duplicated, or invalid).

† adalberto.fontoura@ess.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD045

# PHYSICS APPLICATION SOFTWARE FOR FRIB: FROM COMMISSIONING TO OPERATIONAL EXCELLENCE*

T. Zhang†, K. Fukushima, K. Hwang, T. Maruta, A. Plastun, Q. Zhao, P. Ostroumov
Facility for Rare Isotope Beams, Michigan State University, East Lansing, MI, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The physics application software is a critical part of the FRIB accelerator's control and beam tuning infrastructure. Development of high-level applications and online modeling tools began well before initial beam commissioning to support early machine setup, diagnostics, and operational readiness. As the accelerator transitioned to routine beam delivery, a broader suite of applications was developed and iteratively refined through close collaboration between accelerator physicists, control system and operation engineers, and software developers. These applications leverage model-based control techniques, online beam dynamics simulations, and automated optimization algorithms to enhance tuning efficiency and improve beam delivery reliability across a wide range of operating conditions. Robust data management has also become essential, enabling the capture, organization, and rapid access to operational and diagnostic data critical for prompt decision-making and post-run analysis. This paper presents an overview of the current software ecosystem, highlights key applications for lattice characterization, beam tuning, data management, and outlines future directions, including expanded automation, tighter integration with machine learning frameworks, and improvements in scalability and maintainability.

## INTRODUCTION

The Facility for Rare Isotope Beams (FRIB) at Michigan State University is a unique modern user facility, featuring a large-scale superconducting LINAC capable of accelerating heavy ion beams from oxygen to uranium. It supports the generation of a wide range of rare isotope beams for nuclear science research [1]. A sophisticated high-level physics control software system has been developed prior to early commissioning, continually refined during staged commissioning phases, and remains under active evolution during user beam operations.

Built on top of the EPICS distributed control system [2], and leveraging efficient physics modeling codes, a systematic software framework named *PHANTASY* has been designed to address both the immediate and long-term requirements of accelerator commissioning and operation [3]. A model-based virtual accelerator engine was developed to simulate the EPICS environment, enabling early application development independent of hardware availability. The information flow within the accelerator, particularly for essential physics configurations, is carefully abstracted into high-level programming interfaces, allowing applications to operate seamlessly against either virtual or real machines. The Python scripting framework is implemented with consistent unit management to ensure accurate interpretation between machine controls and physics modeling environment.

To better support commissioning, user-friendly graphical interface applications have been developed by encapsulating the validated scripts. Modularized Qt widgets have been created to handle general use cases ranging from data visualization to the integration of physics models, streamlining the development of Python-based graphical user interface (GUI) applications.

To manage various types of control system data, multiple libraries and tools have been developed to integrate live and archived data into the Python ecosystem, facilitating the unified interface for advanced data analysis application development.

The rapid advancement of modern computing technologies enables the adoption of powerful new methods into the development workflow. For example, data-driven modeling techniques using machine learning are being actively explored, along with best practices for software engineering to ensure long-term code quality and maintainability.

These combined efforts contribute to the continuously improving efficiency of beam development and delivery at FRIB.

## EARLY DEVELOPMENT AND COMMISSIONING SUPPORT

The FRIB accelerator began commissioning in 2017, while development of high-level physics applications started earlier, around 2015. Two critical components emerged during this phase: the physics modeling code and the overall software architecture for physics applications.

### *Physics Modeling*

A key requirement at the time was a simulated EPICS environment with configurable virtual devices to support engineers and physicists in developing and testing applications. To meet this need, and to provide simulation responses more realistic than random noise, *FLAME*, a fast accelerator simulation engine, was proposed and developed.

*FLAME* (Fast Linear Accelerator Modeling Engine) is a novel beam envelope tracking code capable of efficiently modeling non-axisymmetric superconducting RF cavities and multi-charge state acceleration [4]. It features matrix-based models for major lattice components such as solenoids,

* Work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, and used resources of the FRIB Operations, which is a DOE Office of Science User Facility under Award Number DE-SC0023633 and DE-SC0024707.

† zhangt@frib.msu.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD051

# LIBERA INSTRUMENTS INTEGRATION WITH CONTROL SYSTEMS

A. Kete*, P. Leban, M. Babič, and R. Černe, Instrumentation Technologies, Solkan, Slovenia

*Abstract*

Libera instruments have been used with various control systems for several years. In line with the latest security and functionality upgrades, Libera control system interfaces have also been upgraded. EPICS interface in Libera instruments has been upgraded to support the latest EPICS BASE version 7.0.9, which enables users to use the PVA protocol and retrieve more signal data in a single call. Group PVs, allowing atomic access to all signal components, were also added. Furthermore, similar parameters, such as sensors, can be grouped on the Libera side already and provided to PVA clients in a single PV. The TANGO interface has been upgraded to version 9.5. It supports flexible configuration where device classes can be configured for each board type individually. The interface has also been extended with TANGO alarm and logging functionalities. Both interfaces, EPICS and TANGO, can run on the Libera instrument or they can now be compiled and run from an external server station. This network architecture enables easier maintenance and upgrades. This paper details all recent updates and improvements to the Libera control system interfaces and presents possible use cases.

## INTRODUCTION

Instrumentation Technologies has developed Libera family of instruments, which offer remote access and data acquisition through dedicated tools. Various adapters were introduced to simplify integration of Libera instrument into control systems. In particular, *libera-ioc* and *libera-ds* were created for the two primary frameworks, EPICS [1] and TANGO [2], respectively. Both adapters provided stable means for integration of Libera instruments into control systems, they received only several updates until lately.

The adapters can now run either directly on the instrument or remotely on an external system. This is especially beneficial for less performant instruments with lower CPU processing power and lower RAM capacity (see Fig. 1).

## LIBERA IOC (EPICS ADAPTER)

The libera-ioc was updated to support the latest EPICS BASE release (7.0.9) In addition to the existing Channel Access interfaces, PV Access interfaces were also enabled to improve performance, scalability, and support for modern EPICS features. This enhancement allows for more efficient data communication and better integration with current control system infrastructures.

The two key challenges were successfully addressed:

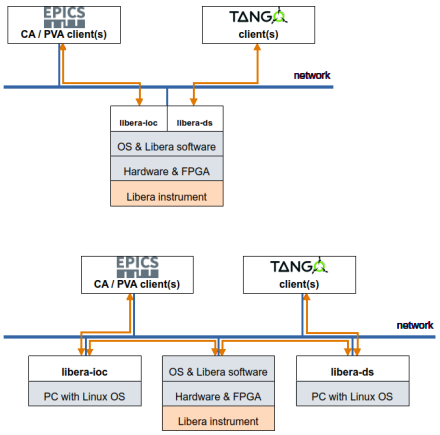



Figure 1: Control system adapters can be deployed locally or remotely.

- atomic acquisition of multiple signal components,
- enabling archive (log) monitors.

The atomic acquisition of multiple signal components ensures consistent and synchronized data capture that improves performance and enhances system reliability and analysis. Enabling archive monitors greatly simplifies the integration of Libera instruments with EPICS archiver services, such as the EPICS Archiver Appliance.

Atomic acquisition is enabled via new group PVs which are automatically instantiated for each signal when IOC is built with PV Access support - no changes to database files are required (Fig. 2). Group PV uses standard `epics:nt/NTTable:1.0` layout [3].

Further customizations can be done via usual group entries in customized database files [4].

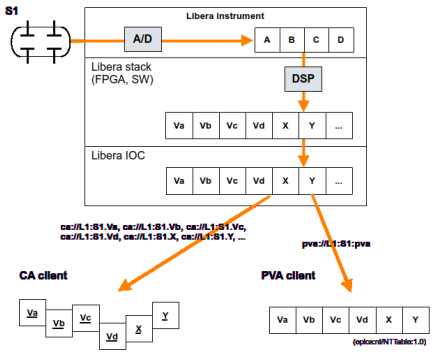



Figure 2: With PV Access group PV, atomic acquisition of signal components is much simpler.

To simplify integration into custom IOCS, Libera device and record types support were separated. This allows IOC maintainers to extend their own IOCs with support for Libera

* ales.kete@i-tech.si

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD052

# DESIGN STUDY FOR INTEGRATING EPICS-BASED CONTROL SYSTEMS WITH MEDICAL TREATMENT APPARATUS

M. Giacchini†, C. Baltador, F. Grespan, M. Montis, A. Pisent
Istituto Nazionale di Fisica Nucleare - Laboratori Nazionali di Legnaro, Legnaro, Italy

*Abstract*

This contribution proposes a design study focused on the future integration between an Experimental Physics and Industrial Control System (EPICS) based accelerator control system and a medical treatment environment within the ANTHEM (AdvaNced Technology for Human centrEd Medicine) project. As the accelerator subsystems (source, RFQ, MEBT, etc.) evolve toward operational readiness, a conceptual architecture is needed to bridge high-level beamline control with treatment room systems. The goal is to anticipate integration challenges by identifying critical requirements, data flows, and safety constraints typical of medical environments. This study explores design approaches to minimize bottlenecks at the final integration phase, particularly where patient positioning, safety interlocks, and real-time synchronization with beam delivery intersect. The study will not implement a prototype, but rather provide a high-level blueprint and system requirements, laying the groundwork for future harmonization between scientific instrumentation and clinical operation contexts.

## INTRODUCTION

The ANTHEM project is an Italian initiative aimed at improving care for patients with chronic diseases and fragile conditions. It adopts a patient-centred approach in which treatment pathways and technologies are customised to patients' specific health needs, living environments, and desired outcomes. To achieve these goals, the initiative brings together twenty-three partners from universities, research centres, hospitals, and private companies across Italy. Activities are organised into four "Spokes" covering smart monitoring, prevention and diagnosis, personalised medical treatments and theragnostic. One of the main applications of ANTHEM is Boron Neutron Capture Therapy (BNCT), an experimental radiotherapy that uses a high-flux neutron beam to destroy boron-loaded tumour cells [1]. Spoke 4 of ANTHEM aims to develop a hospital-based BNCT facility.

The facility will produce thermal neutrons by bombarding a beryllium target with a 5 MeV, 30 mA proton beam, which generates nuclear reactions that produce neutrons. In this way that beryllium target can be seen as a proton to neutron converter. The main accelerator parameters of the ANTHEM BNCT facility are summarized in Table 1.

The high-power proton beam is initially generated by an ion source and transported through a Low-Energy Beam Transport (LEBT) line. It is then accelerated by a Radio-Frequency Quadrupole (RFQ) and further conditioned in the Medium-Energy Beam Transport (MEBT) section before being delivered to the target. Following acceleration, the beam passes through a moderator designed to shape, convert, and moderate the particle energy, ensuring it is suitable for safe and effective clinical application. The sequence of these steps, shown Fig. 1, is critical to transform a high-energy proton beam into a neutron source suitable for BNCT treatment.

Table 1: ANTHEM BNCT Main Parameters

| | | |
|---|---|---|
| **Ion source** | MDIS | Proton energy 80 keV |
| **Accelerator type** | RFQ | 352.2 MHz, 7.13 m long |
| **Final Energy** | 5 Mev | |
| **Beam current** | 30 mA | Accelerated |
| **Duty cycle** | CW | 50% duty cycle possible |
| **p-n Converter** | Beryllium | 150kW, 1kW/cm2 |
| **RF power** | 1.0 MW | |
| **Power consumption** | 3.5 MVA | |

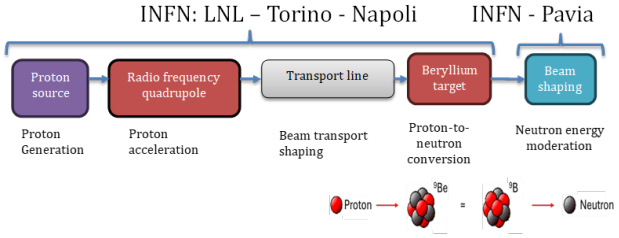



Figure 1: ANTHEM accelerator layout.

Building such a facility requires not only state-of-the-art accelerator engineering but also integration with clinical workflows. The control system for ANTHEM's BNCT accelerator will be built on the EPICS framework [2, 3]. EPICS is widely used at research laboratories to coordinate complex instruments; it offers a toolkit of software components for building distributed controls. Using EPICS will help ANTHEM manage development schedules and reduce maintenance costs by providing a common platform for equipment control and data management.

However, most EPICS installations are used to control scientific apparatus rather than clinical equipment, and experienced developers emphasize that EPICS is never used for safety-critical functions. These responsibilities are instead handled by dedicated safety systems, such as certified

† mauro.giacchini@lnl.infn.it

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD053

# SWARM AND BAYESIAN OPTIMIZATION STRATEGIES FOR THE PIAVE-ALPI ACCELERATORS AT LNL

M. Montis*, L. Bellan, D. Bortolato, M. Giacchini, Y. K. Ong[1]
INFN - Laboratori Nazionali di Legnaro, Legnaro, Italy
[1]also at University of Sapienza, Rome, Italy

*Abstract*

The ALPI linear accelerator at the Legnaro National Laboratories serves as the final superconducting stage in a complex chain designed to accelerate heavy ions—from carbon to uranium—for nuclear and applied physics experiments. It also plays a key role in the SPES project, aimed at re-accelerating exotic radioactive ion beams. Within the TANDEM-PIAVE-ALPI (TAP) complex, the PIAVE injector provides superconductive acceleration of very low velocity ions before they enter ALPI. Managing the interface between these two systems poses significant operational challenges: manual tuning is often required, resulting in lengthy setup procedures and reduced transmission efficiency. Beam instabilities further complicate operations, requiring frequent manual re-adjustments. To address these limitations, advanced optimization strategies based on swarm intelligence and Bayesian algorithms have been applied. These methods enable coordinated control of multiple subsystems, including beam optics, RF settings, and ion source parameters, offering a more autonomous and adaptive tuning process. Experimental results demonstrating the effectiveness of this approach will be presented.

## INTRODUCTION

The heavy ion accelerator complex at LNL is composed of three main accelerators: Tandem, ALPI, and PIAVE. The Tandem XTU is a Van de Graaff accelerator that normally operates at terminal voltages of up to about 14 MV. It can function either as a stand-alone machine or as an injector for the linear booster ALPI (Acceleratore Lineare Per Ioni). ALPI is a linear accelerator built with superconducting resonant cavities. It is organized into three sections: a low-beta branch, which plays a crucial role in accelerating heavy-mass ions, a medium-beta branch, and a high-beta branch. ALPI can also be coupled with the PIAVE injector, which includes a superconducting RFQ (Radio Frequency Quadrupole) and an ECR (Electron Cyclotron Resonance) source.

Due to the wide variety of ion beams required, machine tuning is a critical activity. As shown by the TAP performance data of the last three years (Table 1) [1], beam preparation has a significant impact on time and costs, making its optimization a key asset for the facility.

## THE LNL CASE

As previously mentioned, the Legnaro CW heavy ion accelerator facility is divided into two main sections: the injectors and the superconducting independent cavity Linac, ALPI [2]. Stable ion beams achieve a final output energy of about 10 MeV/u (for stable ion beams), accompanied by an average output current of roughly 100 nA. The provided ion species provided can span in a range from protons to 208Pb ions.

Table 1: TAP Performances in 2022–2024

| Activity | Year | Time [hours] |
|---|---|---|
| Beam on Target | 2024 | 3277 |
| | 2023 | 3475 |
| | 2022 | 2560 |
| Beam Preparation | 2024 | 850 |
| | 2023 | 1137 |
| | 2022 | 541 |
| Accelerator Division Test | 2024 | 480 |
| | 2023 | 214 |
| | 2022 | 240 |
| Unscheduled Maintenance | 2024 | 481 |
| | 2023 | 717 |
| | 2022 | 461 |
| Working Time | 2024 | 4607 |
| | 2023 | 4826 |
| | 2022 | 3341 |

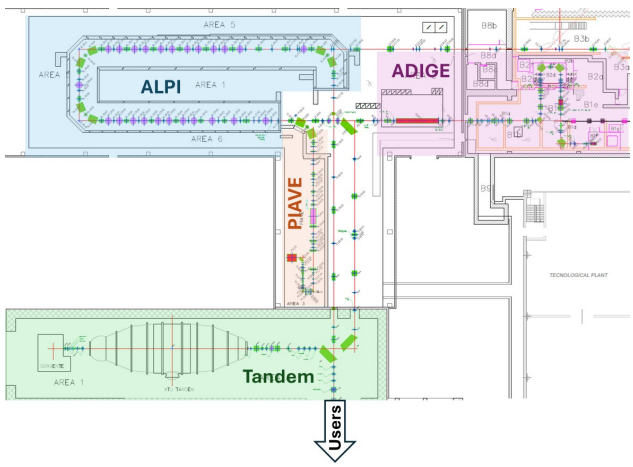



Figure 1: The TAP plant schema identifies Tandem XTU (green), ALPI linear accelerator (blue), PIAVE injector and line (red) and the newer ADIGE line (purple).

ALPI utilizes two injectors for stable ions: the TANDEM-type electrostatic accelerator responsible for light ion acceleration, and the PIAVE superconductive RFQ [3]. A third

* maurizio.montis@lnl.infn.it

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD055

# MODERNIZING LEGACY PYTHON APPLICATIONS AT LCLS*

Z. Lentz†, R. Tang-Kong, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

Between the start of LCLS in 2009 and Python 2's end of life in 2020, many control system tools and user interfaces were created using in-house Python 2 environments. With the end of official support for Python 2 comes a host of maintenance issues for the legacy applications that survive to this day. This paper will contain techniques for modernizing these applications, common pain points in Python 2 to 3 conversions, and advice for testing and redeployment with recent examples from LCLS.

## INTRODUCTION

The Linac Coherent Light Source (LCLS) has been operating since 2009. Shortly after first light, Python 2.7 and 3.1 were released in quick succession [1]. In the wider Python community, most developers stuck with the stable 2.7 series for many years before switching to Python 3. In the intervening years at LCLS, many software tools were written using Python 2 and many useful applications have persisted all the way up through 2025.

Python 2 reached end of life in 2020, which leaves security and maintainability issues going forward. Python 2 is no longer a standard install in modern Linux systems and Python 2-compatible modules are no longer written or maintained by the community at large. To fully leverage the weight of the open-source community and to keep up to date with security patches, LCLS had a lot of software to update.

This paper will go through various strategies and lessons learned from the LCLS Experiment Control Systems (ECS) team's efforts to migrate mission-critical applications from Python 2 to Python 3 and modernize their architectures.

## DECIDING TO START

It's not an efficient use of time to modernize every application that has ever been written or used by a team. Most software is meant to be used within a particular context and then retired. For example, some hyper-specific code might only be run during a single experiment, and some driver software might only be run while a certain device is still in use.

Migration efforts should be reserved for applications that are used on a regular basis, that need to be kept compatible with modern hardware, and have relevant feature requests or bug reports. In general: supported software should use supported languages and tooling.

For applications that are used rarely, only need to run on older hardware, or are largely in their final forms, this effort is misplaced. One question to consider is "how much longer is the total lifetime of this application?" De-scoping a software tool will free up bandwidth to work on other projects, replacing the old support scope.

In addition, before beginning any migration effort, it's important to look for open-source and commercial software alternatives. For example, switching to a well-known open-source framework that accomplishes similar goals to a legacy application will save time and yield better results with less time investment.

As always, performing a migration like this will be contingent on competing project priorities. Migration work typically has only deferred benefits, and no short-term benefits. Following the above thought exercises is necessary to justify migration in the first place.

## PREPARING FOR MIGRATION

Before starting, there must be a plan. The plan must at minimum include two elements: the scope of migration, and how the software will be tested.

This shouldn't be revolutionary or controversial, but it's especially easy to launch head-first into a large migration and lose sight of the goals, or to believe the migration is complete but have no understanding of how to test that the project is complete.

As much as possible, the scope must be kept small. It can be useful to target one "killer feature" that will help justify the work, but staying focused on simply "upgrading to Python 3" will improve the chances of success. No team wants to "refactor code" for a month with no other discernible output, which sets up the project for failure.

The biggest hurdle to a large refactor or code migration is proving that the result works as good as or better than the original. Ideally, verification is rerunning the same automated tests before and after the migration. More likely, these applications that still run Python 2 won't have followed modern best practices such as continuous integration (CI) or automated testing. Therefore, migrations should always include augmenting the test infrastructure.

Some applications are not feasible to test automatically, and this needs to be planned for appropriately. When preparing for a migration, consider asking yourself the following or similar questions if they are applicable:

- What will you click to confirm that it works?
- Is there a suitable development system to test on?
- What do you expect to happen?
- What do you expect not to happen?
- What is the exhaustive list of features that must work?

## HOW TO CONVERT AN APPLICATION TO PYTHON 3

Python 3 itself (until 3.13) has tools for this. There is a module called "2to3" [2] that automates the basic changes for the migration. We recommend starting by using this tool,

* Work supported by U.S. D.O.E Contract DE-AC02-76SF00515.

† zlentz@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD056

# PVPLOT: A LIVE SOFTWARE OSCILLOSCOPE LIBRARY FOR ACCELERATOR CONTROL SYSTEMS*

E. Westbrook†, T. Ramakrishnan, A. Braido, C. Hatch, H. Watkins
Los Alamos National Laboratory, Los Alamos, NM, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

When operating a complex accelerator facility such as the Los Alamos Neutron Science Center, it is often necessary to observe live signal waveforms for various purposes. Traditionally, this has been done using dedicated physical oscilloscopes, whether permanently installed alongside equipment or temporarily deployed on rolling carts. At times, the screens of these oscilloscopes were even transmitted by video link over coaxial cable to secondary television monitors, which was a remarkable convenience at the time, but is considered cumbersome and limited today. With modern control system software and network infrastructure, the inconvenience of physical co-location and dedicated long-distance cabling with dedicated secondary equipment can be eliminated in favor of a flexible and dynamic distributed software approach which reduces complexity while adding significant capability. Here we present a solution using the Experimental Physics and Industrial Control System (EPICS) and Python as part of a comprehensive control system UI library that allows connection to arbitrary signal sources, simultaneous viewing from multiple remote networked locations, and instant reconfiguration or selection of alternate signal sources. Library architecture and various other available UI tools are also discussed.

## SOFTWARE ARCHITECTURE

### *Overview*

EPICS [1] is an open-source architecture and implementation of a highly distributed instrumentation and control systems architecture used by accelerator and industrial systems worldwide, including Los Alamos Neutron Science Center (LANSCE) [2]. EPICS provides software, protocols, and libraries for hardware interfaces, distributed network communication, and software automation and user interfaces.

This project relies on the *PyEpics* [3] library for Python [4] for access to continuously updated waveform data associated with distributed EPICS Procss Values (PVs) via the Channel Access (CA) network protocol. PVs configured to serve data in a waveform format can be selected through the library, which subscribes to updates from the relevant IOC, makes the latest data available to the GUI, and triggers live GUI updates through *PyQt* [5] as needed. The *PyQtGraph* [6] library is used to generate the GUI visualization of the data in a direct and extremely efficient manner, thanks to the common use of *NumPy* [7] primitives by both *PyEpics* and *PyQtGraph*. Figure 1 shows a block diagram of this architecture.

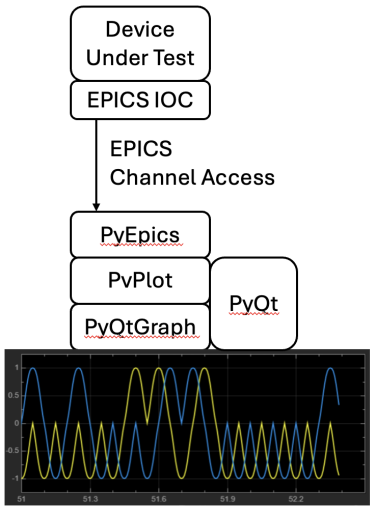



Figure 1: Architecture block diagram.

### *Object-Oriented Design*

In addition to the common Python PEP 8 [8] conventions, this project follows the object-oriented style, conventions, and design philosophies of the *PyQt* library, which is also used as the primary supported GUI framework. Great care is taken to keep classes minimal, complete, and well-organized. Inheritance is heavily and carefully leveraged to maximize reusabiliity, readability, and maintainability, as well as to minimize complexity and redundancy.

## ESSENTIAL LIBRARY COMPONENTS

Figure 2 shows the inheritance and containment relationships of the essential components. Understanding these relationships is important for optimal use of the libray.

### *Pv*

The *Pv* class provides the essential interface to the *PyEpics* library and therefore to EPICS itself via the Channel Access (CA) protocol. Passing an EPICS PV name and an optional *PyQt* signal to the *Pv* constructor results in an object that asynchronously receives and maintains updated waveform data from the associated signal under test. Each

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001), LA-UR-25-29161. No AI or ML was used in any way in the creation of this document nor the work product described.

† westbrook@lanl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD057

# OVERVIEW AND CURRENT STATUS OF THE SKA-LOW MONITORING, CONTROL AND CALIBRATION SUBSYSTEM (MCCS)

E. L. Arandjelovic*, A. Child†, G. Chira‡, A. Clemens§, T. Moynihan¶
Observatory Sciences Ltd., UK
D. Devereux[1‖], M. Waterson**, SKA Observatory, UK
M. Davies††, J. Harvey‡‡, S. Melhuish§§, University of Manchester, UK
[1]also at CSIRO, Australia

*Abstract*

SKA-Low is the low frequency radio telescope currently under construction in Western Australia. At its final extent, it will consist of 512 stations up to 74 km apart, each containing 256 antennas which can be used in different combinations to digitally "point" the telescope. The Monitoring, Control and Calibration Subsystem (MCCS) is responsible for performing calibration and providing local monitoring and control of all of the LFAA (Low Frequency Aperture Array) hardware components. This includes managing the allocation of resources for an observation, and the aggregation of health status. SKAO has adopted the Tango control system framework, and the MCCS software comprises upwards of 18 different Tango devices, some of which are replicated dozens of times for a single station. This complexity poses a significant challenge to computing resources and reliability when considering how to scale up the system, first to the 16 stations to be constructed and integrated by January 2026, then 68 stations at the end of 2026, and targeting 307 stations by mid-2028. This paper will describe the MCCS architecture, report on our latest performance profiling, and discuss how we are preparing for the AA 2 construction milestone which will need to support 68 stations.

## INTRODUCTION

The stated mission of the Square Kilometre Array Observatory (SKAO) is to provide two world-leading radio telescopes: SKA-Low in Western Australia's Shire of Murchison on the traditional lands of the Wajarri Yamaji targeting the 50-350 MHz observing range, and SKA-Mid located in the Karoo Region, South Africa, targeting 350 MHz - 15 GHz. This paper focuses on the Monitoring, Control, and Calibration Subsystem (MCCS) being developed to support SKA-Low, a "mathematical" telescope which when complete will comprise 131,072 log-periodic antennas to receive the radio signals [1], grouped into 512 stations and distributed in three spiral arms radiating from a central core. The signal is first boosted by a low noise amplifier and then converted to an analogue optical signal for transmission to a processing facility. Here the RF signals are digitized and transmitted 800 km off-site to the correlator / beamformer (CBF) where signals from different stations are combined, allowing high-resolution sky images to be formed from the interferometric visibilities or high-sensitivity narrow beams to be directed for deep pulsar searches [2].

This paper provides a brief overview of the MCCS software and its role in the operation of the SKA-Low telescope. It begins with a high-level description of the purpose of MCCS and how it fits into the overall telescope control system. The next section outlines the software architecture, and then provides an overview of some critical functionality including calibration and health monitoring. The paper is concluded with a discussion about performance profiling and how we are tackling the challenges involved in scaling up to 68 stations by the end of 2026.

## A BIRD'S EYE VIEW OF MCCS

Functionally, MCCS is part of the Low Frequency Aperture Array (LFAA) telescope subsystem, and sits underneath the Telescope Manager (TM) which integrates all the engineering entities in the system to form the telescope device [3]. MCCS interfaces to the Telescope Monitoring and Control (TMC) software used by the telescope operators to control the telescope. The purpose of MCCS is to:

- Control and manage the LFAA receiver electronics
- Acquire, process and apply signal chain calibration data
- Acquire, process and provide diagnostic data products
- Monitor, aggregate and provide resource health to TM
- Provide an interface to TMC to acquire and transmit science data to the correlator and beamformer in the Central Signal Processor (CSP)

This context is illustrated in Fig. 1. SKAO has adopted the Tango [4] control system framework, and at the time of writing MCCS consists of 18 different Tango devices which are deployed using Kubernetes [5]. These Tango devices implement the SKA Control Model, a standard set of state transitions and mode indicators developed to unify the design of control system components across the telescope [6]. The device hierarchy that has been implemented reflects the layout of the physical components of the SKA-Low telescope for which MCCS is responsible, namely:

* ela@observatorysciences.co.uk
† abc@observatorysciences.co.uk
‡ gsc@observatorysciences.co.uk
§ ajc@observatorysciences.co.uk
¶ tdm@observatorysciences.co.uk
‖ drew.devereux@skao.int
** mark.waterson@skao.int
†† mark.davies-6@manchester.ac.uk
‡‡ joseph.harvey@manchester.ac.uk
§§ simon.melhuish@manchester.ac.uk

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD060

# AI-DRIVEN DEVICE DRIVER GENERATOR*

K. Fugiel, K. Klimczyk, Sz. Kupiecki, T. Madej, L. Zytniak†
S2Innovation, Krakow, Poland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Developing device drivers for control systems can be a time-consuming task, often requiring specialists to write boilerplate code and implement device-specific logic by hand. In frameworks like TANGO Controls, tools such as POGO exist to auto-generate skeleton code for new devices, but these produce only generic templates that still demand significant manual coding to implement the actual device logic. The rise of large language models (LLMs) presents an opportunity to automate and accelerate this process. Recent work has shown that LLMs like GPT-4 can produce working code for instrument control by interpreting natural language descriptions or documentation. In the embedded systems domain, for example, new AI-driven tools can convert hardware datasheets into digital models and automatically generate low-level driver code, greatly reducing the need for hand-coding and speeding up development. Building on these advances, we introduce an AI-driven device driver generator that integrates LLMs into the development of TANGO device servers. This approach aims to go beyond template generation by using device documentation to implement device-specific functionality, thereby producing more complete drivers with minimal human intervention.

## RELATED WORK AND EXISTING SOLUTIONS

Traditional code generation tools for control system drivers have largely been template-based. In the TANGO Controls ecosystem, the POGO code generator provides a graphical interface for defining a device class model and outputs a device server skeleton in languages like C++, Java or Python [1]. POGO simplifies initial class creation but does not fill in the device-specific behaviour – developers must manually code the logic for each attribute and command based on the device's manual. Similarly, in other environments (e.g. EPICS or vendor-specific frameworks), one can find wizards or templates to scaffold driver code, yet the heavy lifting of implementing protocol specifics remains manual.

More recently, attention has turned to using AI to assist or automate driver development. Developers have begun experimenting with general-purpose AI coding assistants (e.g. OpenAI's Codex/ChatGPT or Anthropic Claude) to generate driver code from datasheets or problem descriptions. In one reported case, an engineer used an AI assistant to help write a QNX driver for a thermal sensor, finding that the AI could suggest register operations and fix errors through iterative prompts [2, 3]. While such ad-hoc use of chatbots shows promise, it still requires the developer to guide the AI and verify the code through multiple rounds of Q&A.

Dedicated solutions are emerging to make this process more systematic. Embedd.it, for example, has demonstrated a tool that parses component datasheets into a structured digital twin model and then generates production-ready driver libraries "at the click of a button," using deterministic code generation to avoid AI hallucinations [4]. This approach has been shown to quickly produce drivers for sensors and peripherals and even supports frameworks like Zephyr RTOS, substantially accelerating R&D workflows [5]. Likewise, researchers in materials science reported using ChatGPT-4 to create a complete control software module (including a GUI) for a Keithley 2400 SourceMeter instrument solely from the natural-language instrument manual, with only minimal human corrections needed [6]. These examples underline a trend: by leveraging AI with domain knowledge (e.g. device documentation), one can automate much of the coding for device control.

However, a gap remains in applying such AI-driven code generation specifically to TANGO Controls device servers, which have a particular structure (attributes, commands, properties) and rely on high-level device classes rather than low-level register tweaking. The tool presented in this paper addresses this gap by marrying the template-based convenience of existing tools with the intelligence of LLMs to implement device-specific logic.

## BENEFITS OF AN AI-ASSISTED DRIVER GENERATOR

Adopting an AI-driven device driver generator yields several benefits over traditional development or template-only approaches:

- Accelerated Development: By automatically writing the device-specific code, the tool can shorten the driver development cycle from weeks to hours. Early tests showed that for devices like power supplies and teslameters, the generated TANGO server code required only minor edits before it was functional, indicating a dramatic speed-up in getting a working driver. This allows faster integration of new instruments into the control system.
- Reduced Manual Coding and Errors: Automating boilerplate and parsing device documentation spares developers from writing repetitive code and reduces the chance of human error. The generator reads the official device manual (if provided) and uses it to produce correct command sequences and data conversions, which means fewer omissions or mistakes compared to manually coding from a spec sheet. It effectively acts as an expert assistant that doesn't "forget"

* Work supported by S2Innovation
† lukasz.zytniak@s2innovation.com

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD061

# SIMPLIFYING CRYOGENIC PROCESS CONTROL AT ESS LINAC THROUGH AUTOMATION: DEVELOPMENT AND INTEGRATION OF AN AUTOMATIC CONTROL SEQUENCE

W. Bińczyk†,1,2, D. Domagała[1], N. Elias[2], M. Skiba[2], E. Asensi[2], A. Fontoura[2], H. Cardona[2,4], N. Baskar[2], A. Rizzo[2], P. Arnold[2], J. Zhang[2], P. Nilsson[2], P. Hałczyński[3], W. Gaj[3]
[1]S2Innovation, Krakow, Poland
[2]European Spallation Source, Lund, Sweden
[3]IFJ-PAN, Krakow, Poland
[4]AFRY, Stockholm, Sweden

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

This paper presents the Automatic Control Sequence (ACS) developed and implemented to simplify the control of cryogenic processes in the linear accelerator (LINAC) at the European Spallation Source (ESS), which includes 27 cryomodules and 43 valveboxes. The main objectives of the project were to reduce the risk of human error and minimize manual operations — while maintaining full decision-making authority in the hands of the operator. The sequence is designed through interdisciplinary collaboration, with Excel serving as a database for information exchange in human readable form. A custom Python script is then used to generate PLC code in SCL programming language based on the defined logic. The final sequence is deployed on a master PLC and 43 dedicated PLCs, fully integrated with the EPICS control system and interconnected via Profinet for optimized system synchronization. A user-friendly operational interface was developed using CS-Studio, serving as both a monitoring and control layer. It provides visibility across all levels of the control system — from individual devices, through local PLC sequences, up to inter-system synchronization. This paper provides an overview of the development of the Automatic Control Sequence and discusses key lessons learned in this process as well as possible future improvements in cryogenic control at ESS.

## INTRODUCTION

Cryogenic systems are indispensable for superconducting accelerators. At the European Spallation Source (ESS), stable cryogenic conditions ensure that superconducting radio-frequency (SRF) cavities remain below their critical temperature, enabling efficient acceleration of high-intensity proton beams. The accelerator complex ultimately aims for 5 MW of beam power, supported by 27 cryomodules and 43 valveboxes integrated into the linear accelerator (LINAC). Each cryomodule must be carefully cooled to 2 K, maintained within tight stability margins and synchronized with the ACCP [1].

† wojciech.binczyk@s2innovation.com

### *Manual Control Challenges*

Prior to the introduction of ACS automation, cryogenic operations were almost entirely manual. Operators carried out lengthy procedures consisting of valve actuation, heater adjustments, and pressure monitoring. These operations were often repeated, resulting in a substantial burden during testing in Test Stand 2 (TS2), which is a CRMs commissioning installation [2, 3]. As reported during commissioning campaigns and activities performed in cryomodules, manual procedures carried several risks:

- High workload for control room operators, who had to track dozens of concurrent actions.
- Increased risk of human error, such as mis-timed manipulations.
- Inconsistency between different operators, which sometimes led to variations in cooldown and warm-up results.
- Increased risk of failure during night shifts.

This operational reality motivated the design of an Automatic Control Sequence (ACS). The ACS sought to simplify the control of cryogenic processes, reduce operator workload, and improve reliability and reproducibility. Importantly, its design philosophy emphasized that automation should support, not replace, operator authority.

## SYSTEM OVERVIEW

As shown in Fig. 1, the cryogenic infrastructure of the ESS LINAC is composed of three interlinked subsystems:

- Accelerator Cryogenic Plant (ACCP). Provides helium at the necessary temperatures and pressures.
- Cryogenic Distribution System (CDS). Transports helium from the ACCP to the accelerator tunnel, including transfer lines and 43 valveboxes, one for each cryomodule.
- Cryomodules (CRM). Distributing helium flows and regulating cavity conditions.

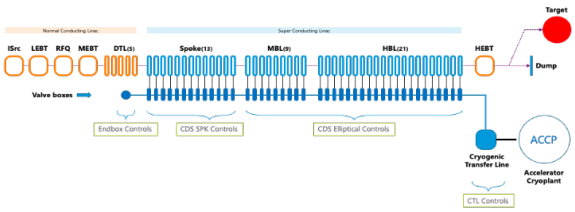



Figure 1: Overview of the system parts.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD062

# FIELD DEPLOYMENT AND ITERATIVE ENHANCEMENT OF THE DISH STRUCTURE QUALIFICATION (DiSQ) SOFTWARE FOR SKA-Mid

U. K. Pedersen*[1], Observatory Sciences Ltd, St. Ives, United Kingdom
J. Engelbrecht[1], Vivo Technical, Cape Town, South Africa
T. Juerges, SKAO, Jodrell Bank, Lower Withington, Macclesfield, Cheshire, UK
O. Skivington[1], Critical Software, Chilworth, Southampton, United Kingdom
[1]also at SKAO, Jodrell Bank, Lower Withington, Macclesfield, Cheshire, UK

*Abstract*

As part of the construction of the SKA-Mid telescope in South Africa's Karoo desert, each of the 133 new mid-frequency radio dish structures, supporting a 15 m diameter dish, must undergo a thorough qualification process before they are integrated into the array. To support this work, the Square Kilometre Array Observatory (SKAO) Wombat team has developed the Dish Structure Qualification (DiSQ) software: a tailored suite of tools designed to interact with the dish structure's PLC-based control system via an OPC-UA interface. DiSQ comprises a user-focused engineering GUI, a synchronous Python API for automated testing with bespoke scripts, and a high-performance data logger that captures engineering parameters in HDF5 format. Since 2024, DiSQ has been successfully deployed during testing and commissioning activities by Dish Structure engineers, operating in the field on the first delivered dish structures. Its modular design enabled rapid adaptation to differences between simulation environments and real hardware, with updates informed by continuous feedback from SKAO and SARAO Dish Structure engineers. This paper presents the current status of DiSQ, highlights lessons learned from deployment, and details enhancements made to improve usability, resilience, and compatibility with evolving control interfaces. DiSQ's evolution exemplifies the value of iterative development and close collaboration between software development teams and end-users in delivering robust tools for complex scientific engineering tasks.

## INTRODUCTION

The SKAO is a global-scale project which aims to build the world's largest radio telescope. The SKAO will be constructed as a single observatory with two radio telescopes: SKA-Mid and SKA-Low. SKA-Low is located in Australia and is designed to monitor low-frequency radio waves from the early universe using arrays of static antennas. SKA-Mid on the other hand is designed to observe mid-frequency radio waves, and the first few of an array of 133 new radio dish structures are currently under construction in South Africa's Karoo desert. The scientific goals of the SKAO [1] are pushing the boundaries of science and aim to answer significant questions about fundamental physics and the origin of the universe. These aims translate to high demands on the performance and reliability of the telescope, and for the first few delivered dishes the design must be carefully qualified to demonstrate compliance with strict mechanical requirements after installation. The Dish Structure Qualification (DiSQ) software [2] was developed by SKAO's Team Wombat in close collaboration with Dish Structure Engineers from the South African Radio Astronomy Observatory (SARAO) [3] to support this process. This paper presents an overview of the design and implementation of the software, the iterative development process in collaboration with Dish Structure Engineers, and reports on its current status and lessons learned from use in the field.



Figure 1: Dish Structure assembly in South Africa's Karoo Desert. Photo credit: Keith Browne, SKAO.

## DISQ SOFTWARE OVERVIEW

The DiSQ software was developed to address the specific needs of the Dish Structure Engineers at SKAO and SARAO. A particular emphasis was made on the requirement that the software could run locally on a laptop in the field, without reliance on network connectivity other than direct connection to the Dish Structure Controller (DSC). The lack of

* ukp@observatorysciences.co.uk

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD064

# LOGGING INFRASTRUCTURE FOR EPICS-BASED CONTROL SYSTEMS USING LOKI AND PROMTAIL

J. Bellister*, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.


## *Abstract*

To support multiple scientific facilities at SLAC, a modern logging system capable of high message throughput and easy record filtering is required. A new logging system deployed at SLAC meets these requirements by blending an existing EPICS logging server with open-source technologies for log storage and retrieval. To save log data, Loki was chosen as the data storage technology. Promtail was selected to push existing log files generated by various EPICS clients into Loki. To encourage increased accelerator operator engagement, multiple interfaces are provided for interacting with logs. Grafana dashboards offer a user-friendly way to build displays with minimal code, while a custom PyQt-based user interface displays the results of direct queries to Loki in a table-based display. A wrapper script around LogCLI called qlog provides a command-line interface for interacting with the logs. The rationale behind these decisions and their integration into our controls infrastructure will be considered.


## INTRODUCTION

Efficient logging systems are important for modern accelerator facilities, enabling operators, engineers, and scientists to monitor system behavior, diagnose issues, and optimize performance. At SLAC, we recently undertook an upgrade to our existing logging infrastructure in order to both modernize the technology being used and increase user engagement with the logging system.

The existing system was based on the old EPICS iocLogServer that would write to both a flat file and an Oracle database. This system had difficulties with performance when it came to messages older than thirty days. And there were not a lot of convenient off-the-shelf options for interacting with the logs.

Thus, a search for a system that was performant for both high ingest throughput for handling bursts of data, as well as for retrieving records from many months or years ago was conducted. The system should be capable of both historical queries and live data streaming, with easy ways of filtering for both cases.

## SYSTEM ARCHITECTURE

Since there are many modern free and open-source logging solutions available today, the search started with a comparison of these options. Reducing the amount of custom solutions used minimized both the upfront development cost and ongoing maintenance costs that are instead absorbed by the team maintaining the chosen software. Figure 1 illustrates the end-to-end flow of the chosen technologies.

* jesseb@slac.stanford.edu

### *Loki for Log Storage*

Loki [1] is a log storage solution created and maintained by Grafana Labs. It is a free, open-source technology that provides a number of nice features. These include:

- No formatting requirements for log ingestion
- An easy to use API that includes tailing logs in real time
- Native integration with both Grafana and Kubernetes

The first point is particularly important, as it means access to modern visualization and data-processing tools without actually needing to change the log generation process. So although modern tools for log aggregation such as Logstash are available, they wouldn't be necessary if keeping the older EPICS log server code is considered sufficient, which for our cases it was.

The Loki API exposes a number of HTTP endpoints that support posting, retrieving, and filtering log entries. As mentioned above, perhaps the most useful is an endpoint for tailing logs in real time. There is no need to string together multiple queries or API calls to make this work, just one simple call with optional filters will allow for retrieving data as it arrives.

Elasticsearch was also considered, as the Elasticsearch-Logstash-Kibana (ELK) stack is a widely-used suite of logging and data-processing tools as well. But the tighter focus of Loki on managing logs only coupled with the existing deployment of it for monitoring our Kubernetes-based infrastructure made it a more natural fit.

### *Promtail for Log Ingest*

Promtail [2] is the technology used to ingest logs into the Loki data storage layer. Because it is highly flexible to configure, we were able to use it to ingest logs as written by the EPICS iocLogServer without needing to make any changes to the format it output logs in. The only modification was to disable the writes to the Oracle database, leaving only the writes to the flat file.

We then pointed a separate Promtail instance to each facility's file. The configuration file for Promtail specifies how to read the file, with a series of regular expressions parsing out each needed piece of information.

## DEPLOYMENT

The Loki instance is deployed using Kubernetes into the SLAC Shared Science Data Facility (S3DF). Because of the flexibility in labeling and data storage as mentioned above, the same Loki instance can hold not just the logs of our accelerator facilities, but also logs of S3DF itself.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD066

# ADVANCED P4P USAGE AT THE ISIS NEUTRON AND MUON SOURCE

I. D. Finch*, A. Al Shafei, K. R. L. Baker, A. Kurup[1]
ISIS, STFC, Rutherford Appleton Laboratory, Chilton, Oxon, UK
[1]Also at Imperial College London, London, UK

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The p4p library is a Python wrapper for the C++ pvxs library allowing Python developers to access client functionality to put, get, and monitor pvAccess PVs. Server functionality allows the creation of PVs and implements the structure of the most commonly used Normative Types (e.g. NTScalar) and their fields (e.g. alarm, control, etc.). To facilitate the transition to EPICS underway at the ISIS Neutron and Muon Source accelerators, an implementation of the logic of the Normative Type fields and a subset of other IOC functionality such as CALC records has been developed. We present our uses of this work and highlight parts which may be applicable to other facilities interested in using Python.

## INTRODUCTION

At the ISIS Neutron and Muon Source the control system for the accelerators is migrating [1,2] from Vsystem [3] to the open-source EPICS [4] control system. (ISIS Experiment Operations have completed their migration to EPICS [5].) A key technology in this transition is pvecho [6, 7], software developed to allow Vsystem and EPICS to be used simultaneously. The first fully EPICS system deployed for the transition from Vsystem was an EPICS interface to Omron PLCs using the CIP protocol. Both pvecho and the CIP interface were written in Python using the PvaPy library [8].

A key task within this migration to EPICS is to create EPICS interfaces to the CPS system [9] developed in-house at ISIS. These systems constitute the majority of the PVs remaining in Vsystem. The EPICS interface to the CPS systems is planned for implementation in Python [10].

Using Python to replicate functionality of EPICS IOCs has continued with other interfaces. An alternative MQTT interface to Omron PLCs [11] and an interface to Moxa terminal servers developed in Python have been deployed. Currently in planning is an interface to OPC-UA similar to that implemented for the CIP and MQTT protocols. These interfaces read configuration information from the PLCs and use it to create or update PVs.

The machine learning team embedded in our Controls team also make heavy use of Python. It is used for purposes such as model training and for interfacing with the control system for optimisation studies. Testing of machine learning based Bayesian optimisation of the Low-Energy Beam Transport (LEBT) of the ISIS linear accelerator [12] is completing and deployment is expected by the end of the year. This system uses the p4p library [13] to interface with the control system.

* ivan.finch@stfc.ac.uk

### *p4pillon*

Development of these Python interfaces by different developers has meant duplication of effort and a failure to transfer experience. To address this, a new Python library called p4pillon has been developed. The p4pillon library adds to the server functionality of p4p.

Previously p4p for ISIS has been described [10]; p4pillon is an open-source release of code from that library for potential use at other sites. Development of p4p for ISIS will continue in parallel for tools and configuration details which are ISIS site specific. We note where p4p for ISIS and p4pillon differ.

## P4P

Although the first projects (pvecho, CIP interface) in the transition from Vsystem to EPICS used PvaPy [8], subsequent development uses p4p [13]. PvaPy is no longer under active development but continues to receive security patches and bug fixes.

The p4p Python library is recommended for EPICS interfaces in Python. It implements pvAccess client and server functionality. (p4p is an initialism of "pvAccess for Python".) It is a wrapper around the PVXS C++ library [14] (implementing the same functionality) but for Python developers on Windows, Linux and macOS operating systems it appears as a native library.

More fully, p4p implements the pvAccess protocol [15] and EPICS Normative Types (NTs) [16]. It provides this implementation for several Python concurrency models - asyncio, cothread [17], and threads. Roughly speaking, the pvAccess protocol specifies how data is transferred between EPICS servers and clients using UDP and TCP, while the Normative Types specify the format of the data payload. We will not discuss the pvAccess protocol further but give a brief introduction to Normative Types.

### *Normative Types*

There are a number of Normative Types (e. g., NTScalarArray, NTEnum, NTTable) in the specification document, but NTScalar is illustrative. It specifies the data format for the transfer of a single scalar value with optional metadata. In its simplest form a single scalar value (e. g., a boolean, integer, or floating point number or a string) is the entire payload. However, best practice is to include the optional timestamp field.

The other optional fields in an NTScalar are alarm, display, control, and valueAlarm. Strictly, valueAlarm is not listed as an optional field of NTScalars in the current revision of the specification. The specification says that alarmLimit

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD073

# USING COMPUTER VISION FOR ONLINE CALIBRATION OF BEAM INSTRUMENTS AT CERN

J. Martínez Samblas*, S. Burger, M. Gonzalez-Berges
European Organization for Nuclear Research (CERN), Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Accurate calibration of beam instrumentation is critical for the optimal operation of particle accelerators. This work presents a case study of a beam imaging system at CERN's Antiproton Decelerator (AD) target, composed of a light-emitting screen interacting with the beam and an observation camera. During operational use, the system required frequent online recalibrations to address temperature-induced image drifts. To resolve this issue, a fully automated procedure was developed that periodically acquired images and applied multiple computer vision techniques. These techniques included custom curve-fitting methods applied to pre-processed regions of interest and SIFT-based (Scale-Invariant Feature Transform) feature detection to track and correct positional shifts. By automatically performing recalibrations at regular intervals, the approach has significantly enhanced consistency and reliability, enabling continuous and precise beam monitoring in varying environmental conditions. This stabilization technique has subsequently contributed to the optimization of antiproton production at the AD facility. The paper first introduces the challenges associated with calibrating the beam imaging instrumentation of the AD target. It then presents the chosen image analysis techniques, followed by a discussion of the results and measurement errors of the tested methods. Finally, an outlook on potential future improvements is provided.

## INTRODUCTION

Beam instrumentation provides the real-time data needed to monitor, steer, and protect the particle beam. To remain trustworthy, the associated measurements must remain calibrated. Unfortunately, effects such as timing drifts, mechanical shifts, radiation damage, and gradual changes in temperature or humidity can alter an instrument's response within hours or even minutes.

To address this challenge, we present an automated method that continuously calibrates the camera of the Antiproton Decelerator (AD) [1] Target beam imaging system at CERN. The solution relies on computer vision techniques, including curve fitting, Scale-Invariant Feature Transform (SIFT) [2], and Random Sample Consensus (RANSAC) [3], to track geometric offsets with sub-pixel precision, all without interrupting normal operation.

## AD TARGET BEAM IMAGING SYSTEM

The AD produces and decelerates antiprotons for a broad physics program. In the target area, a high-energy proton beam strikes a dense target, and a beam imaging system observes an optical-transition-radiation (OTR) screen to monitor the proton beam profile and position; information that is essential for optimizing antiproton production [4].

### *Why Is Calibration So Important at the AD Target?*

The installed imaging system employs a 20 m-long optical line: a sequence of mirrors and lens assemblies routes the light vertically through shielding blocks to an upstairs service area where the camera sits outside the irradiated zone (Fig. 1). While this protects the sensor and eases maintenance, it also amplifies mechanical sensitivities; a shift of less than 1 mm in any mirror can move the image by several pixels. Precise and frequent calibration is therefore indispensable.

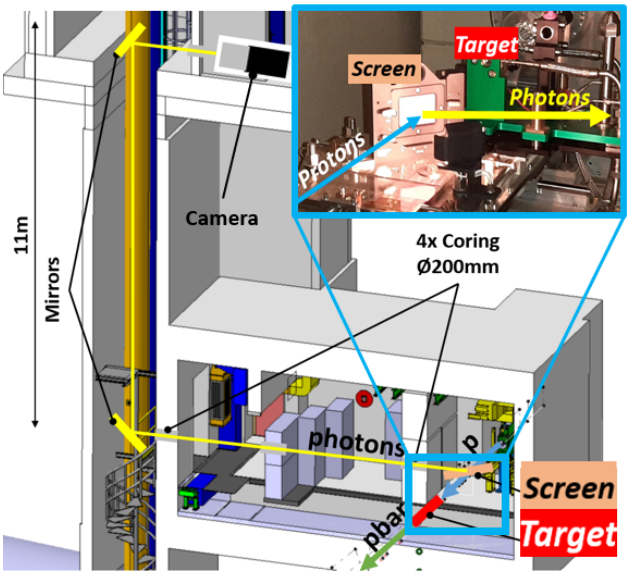



Figure 1: Optical layout of the AD target beam imaging system, highlighting the target/screen station. The long optical path increases sensitivity to mechanical and thermal shifts [4].

### *Calibration Issues*

In the spring of 2024, the operations crew observed large, sporadic shifts in the measured beam position. Evidence pointed to drifts potentially introduced by temperature changes along the optical line, although this had not been conclusively verified. To mitigate such effects, a fully automated calibration scheme capable of real-time compensation was deemed necessary.

* javier.martinez.samblas@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD075

# BAYESIAN ACTIVE LEARNING FOR CONVERGING POSTERIORS IN LATENT VARIABLE INFERENCE FOR CONTROL SYSTEMS*

K. Hwang†, K. Fukushima, T. Maruta, P. Ostroumov, A. Plastun, Q. Zhao, T. Zhang
Facility for Rare Isotope Beams, East Lansing, MI, USA

*Abstract*

Inferring latent variables, such as Courant-Snyder parameters in particle accelerators, is challenging due to noisy and indirect observations that often produce multi-modal posterior distributions, despite the true latent variable being unique. We present a Bayesian Active Learning (BAL) framework to enhance latent variable inference in simulation-equipped control systems. BAL actively selects control settings (e.g., quadrupole magnet configurations) to maximize information gain, efficiently refining multi-modal posteriors into unimodal ones for improved inference accuracy. Using an ensemble of physics-informed beam envelope simulations in PyTorch, our approach approximates posterior sampling and mutual information to guide data acquisition. This interpretable framework holds broad potential for improving latent variable inference in control systems.

## INTRODUCTION

Latent-variable inference is fundamental to many control and diagnostic tasks in physics and engineering, where the underlying system state $z$ is not directly observable but governs system dynamics. Instead, inference relies on indirect, noisy measurements $y$ generated through a forward system

$$y = F(x, z, c) = f(x, z) + \delta(x, z) + \xi, \tag{1}$$

where $f(x, z)$ is a simulation model that approximates the true system $F(x, z, c)$. Here $x$ denotes controllable inputs, $c$ represents unmodeled nuisance factors, $\delta(x, z)$ denotes the discrepancy between simulation and real system, and $\xi \sim \mathcal{N}(0, \sigma_\xi^2)$ is measurement noise. If we model the discrepancy as a random function

$$\delta \sim \mathcal{N}\big(\mu_\delta(x, z),\, \sigma_\delta^2\big), \tag{2}$$

where $\mu_\delta(x, z)$ is an unknown mean discrepancy, and $\sigma_\delta^2$ is the domain expert knowledge on expected size of homoskedastic model error, the likelihood becomes

$$\begin{aligned} p(y \mid x, z) &= \int p(y \mid x, z, \delta)\, p(\delta \mid x, z)\, d\delta \\ &= \mathcal{N}\big(f(x, z) + \mu_\delta(x, z),\, \sigma_\xi^2 + \sigma_\delta^2\big), \end{aligned} \tag{3}$$

Because $\mu_\delta$ is generally unknown, we will treat it as an additional latent quantity to infer jointly with $z$

A canonical example arises in particle accelerators, where reconstructing Courant–Snyder (CS) parameters from profile monitors (PM) or virtual diagnostics (VD) is essential for beam tuning and stability. However, a single diagnostic measurement of $y$ (e.g. one PM scan) under a fixed control setting $x$ typically produces a multi-modal posterior, implying several candidate solutions even though the latent parameter $z$ (e.g. CS) is unique. To resolve this ambiguity, accelerator operations employ control variations—most commonly quadrupole magnet scans (q-scans) in conjunction with PM or VD measurements—which yield sufficient constraints to recover a unique solution for the latent parameters [1]. This illustrates a broader principle: strategically chosen controls can collapse wide, multi-modal posteriors into narrow, unimodal ones, thereby enabling more accurate and reliable inference (see Fig. 1).

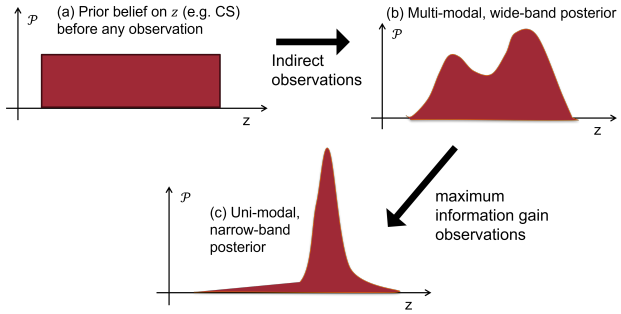



Figure 1: Illustration of posterior refinement through informative control. Horizontal axis: latent variable $z$. Vertical axis: posterior density $\mathcal{P}(z \mid \mathcal{D})$ given dataset $\mathcal{D}$. (a) A weak prior provides only broad bounds on $z$. (b) Limited or uninformative measurements yield a multi-modal posterior, reflecting multiple candidate solutions. (c) Informative control actions (e.g., quadrupole scans) generate observations that disambiguate these candidates, collapsing the posterior into a narrow, unimodal distribution centered near the true latent variable.

However, selecting the most informative controls is nontrivial: poorly chosen settings may waste valuable experimental time, while optimal ones can rapidly collapse the posterior onto the true latent variable. Bayesian active learning (BAL) offers a principled framework for this task. By maximizing the expected information gain of each new observation, BAL guides control selection to most efficiently refine the posterior. Yet, exact Bayesian experimental design is mathematically intractable and computationally prohibitive in control systems, especially when physics or system model $f$ is complex and costly.

* Work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, under Award Number DE-SC0024707 and used resources of the FRIB Operations, which is a DOE Office of Science User Facility under Award Number DE-SC0023633.

† hwang@frib.msu.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD077

# LONG SHORT-TERM MEMORY OF RECURRENT NEURAL NETWORK TO THE ANALYSIS OF TEMPERATURE RISE ON THE PRODUCTION TARGET IN HADRON EXPERIMENTAL FACILITY OF J-PARC

K. Agari†, Y. Sato, A. Toyoda, Y. Morino, H. Akiyama
KEK, Ibaraki, Japan

*Abstract*

Hadron Experimental Facility (HEF) is designed to handle intense slow-extraction proton beam from 30-GeV Main Ring (MR) of Japan Proton Accelerator Research Complex (J-PARC). The production target in the HEF was operated under severe conditions in which its temperatures periodically varied from 30 to over 300 °C by heat deposit of irradiated beams. In a long term, a careful evaluation of damage to the target and its lifetime is quite important. If the target temperatures are accurately predicted from the existing data, including beam intensity, duration of beam extraction and beam positions, it could be possible to verify the cumulative damage by comparing the predicted temperature rise with the measured one. The predicted temperature rise has been calculated from the existing data using linear regression with a machine learning library, scikit-learn. However, the predictions of the temperature rise by linear regression are not fully satisfactory due to changes of beam optics and accelerator conditions. Therefore, an enhanced prediction method for the temperature rise on the proton target has been developed using Long Short-Term Memory (LSTM), a type of Recurrent Neural Network (RNN) architecture. A systematic study was also conducted to investigate the effects of hyperparameters, including a sequence length and hidden layer. The present paper reports the status of the prediction system of the temperature rise on the production target using the LSTM analysis in detail.

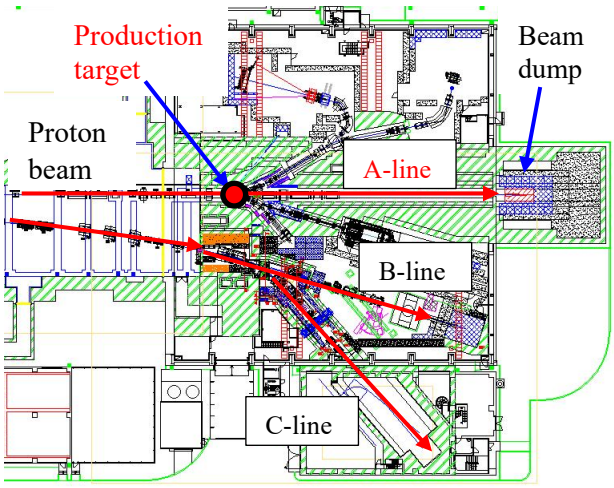



Figure 1: Drawing of J-PARC Hadron Experimental Facility.

† agari@post.kek.jp

## INTRODUCTION

Hadron Experimental Facility [1] (HEF) at Japan Particle Accelerator Research Complex (J-PARC), as shown in Fig. 1, is designed to handle intense slow-extraction proton beam from 30-GeV Main Ring (MR). The period of beam extraction from MR to HEF is 2 seconds and the operation cycle was 5.2 seconds.

### *Production target*

The production target [2], operated in the HEF from 2015 to 2019, was made of a gold and a copper block with coolant stainless pipes, as shown in Fig. 2. The production target was designed to be capable for up to 50-kW proton beams. The dimension of the gold target was 15W ×6H ×66L mm. The gold target was divided into 6 blocks to reduce thermal stress by beam irradiation. A double-headed design enabled to switch the one target to the other quickly and remotely on demand. The water coolant pipes, embedded in the copper block, were made of stainless steel to avoid erosion and corrosion. The production target was enclosed with an airtight chamber filled with circulating helium gas. In order to monitor the temperature rise of the target, six thermocouples, named as TC1 to TC6, were mounted on top of each gold block along the beam direction. Currently, a new production target [3] designed to be capable for up to 95-kW proton beams has been operated since 2020.

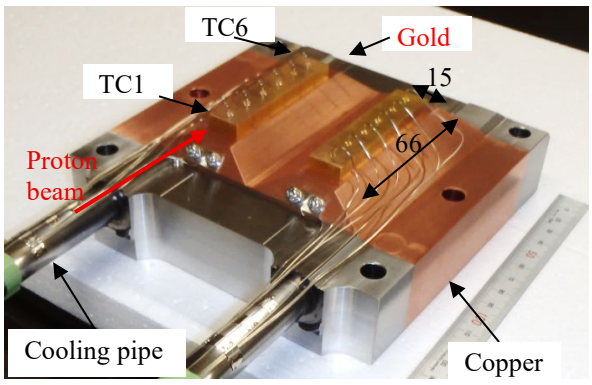



Figure 2: A photograph of production target.

### *Monitoring Device*

The monitoring device is based on Programmable Logic Controller (PLC) with a Linux-CPU working as an embedded EPICS IOC on Yokogawa's FA-M3 PLC platform [4]. Temperatures of the gold target, the copper block, and cooling water are monitored with K-type thermocouples

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD078

# MACHINE LEARNING FOR ISIS CONTROLS

M. Leputa*, K. R. L. Baker, M. Romanovschi, R. Banerjee
ISIS Neutron and Muon Source, Didcot, United Kingdom

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

ISIS Controls Group is continually developing and improving its use of Machine Learning (ML) across the accelerator. In this paper, we discuss its applications as an optimisation tool, anomaly detection, and surrogate modelling tool. In addition, we present a machine learning operations (MLOps) framework for rapid prototyping and reliable and repeatable deployment. We also outline the current progress of work towards a larger-scale model of the accelerator. Finally, we discuss the current and future impact of ML on the control system.

## INTRODUCTION

The operation of particle accelerators presents significant challenges, ranging from stable and efficient operation to managing, fault-finding, and diagnosing complex high-dimensional systems. At ISIS, these challenges are amplified by the increasing volume of data generated as the control system migrates from VSystem [1] to EPICS [2, 3].

The scale and complexity of the data far exceed the capability of humans to monitor and analyse. This drives a need for automation and ML-driven analysis, as well as anomaly detection, particularly unsupervised methods, which offer an opportunity for a more hands-off approach that reveals relationships between subsystems that may not be otherwise apparent.

The ISIS Controls Group has developed several virtual diagnostics for the LEBT and MEBT (Low and Medium Energy Beam Transports, respectively) [4, 5]. The LEBT and MEBT are responsible for monitoring and matching the beam characteristics to the requirements following sections of the linear accelerator [6]. Their layout is shown in Fig. 1 [7]. The virtual diagnostics provide operators with a risk-free and rapid way to test machine settings, eliminating the need for hardware trials or physics simulations, which take longer to execute than model inference. The surrogate models are deployed into the control system as EPICS devices using Poly-Lithic [8] and integrated into Phoebus as GUIs (Fig. 2) [9]. Beyond their role as operator tools, these models also serve as a basis for evaluation of optimisation methods such as Bayesian optimisation and its variations using the Xopt optimisation framework [10].

The remainder of this paper will discuss optimisation, anomaly detection, surrogate modelling, the MLOps infrastructure and practices, as well as discuss the future direction of the development.

* mateusz.leputa@stfc.ac.uk

## SURROGATE MODELLING

There are two virtual diagnostics currently in use at ISIS with the linear accelerator controls: ASTRA **Medium Energy Beam Transport** Surrogate model (MEBT) and **Low Energy Beam Transport** Optimisation model (LEBT).

The **ASTRA MEBT Surrogate** estimates beam descriptors for the ISIS Medium Energy Beam Transport, trained on outputs from the ASTRA simulation code [11]. Using a combination of incoming beam properties and machine settings, the model predicts the beam envelope, energy spread, Twiss parameters, and beam loss along the MEBT. Inputs include the machine settings and the position along the transport line, with predictions available at discrete steps [4, 12]. This surrogate provides operators with rapid diagnostics, enabling a wide range of settings to be evaluated quickly. Because it is trained on simulation data rather than in-place measurements, it can generate beam metrics that would otherwise be unavailable. Recent iterations incorporate uncertainty quantification, allowing model confidence to be assessed and guiding targeted re-training [13].

The **LEBT Optimisation** model, in contrast, is trained on real machine data and is used primarily for offline optimisation studies as ISIS moves towards wider automation in control systems. The model accepts ten adjustable machine settings as inputs and predicts transmission efficiencies through subsequent components of the linear accelerator. Deployed as an EPICS device, it interacts with the control environment through the same middleware as the real hardware. It provides a safe sandbox for testing optimisation routines quickly and without risk [5, 14].

The models are moderately sized by ML standards (less than 100,000 parameters) and are implemented in PyTorch [15]. Early iterations used TensorFlow [16], but we have found that PyTorch provided better community support and flexibility. The models are version-controlled using MLflow [17] and deployed into the control system using Poly-Lithic [8]. The Machine Learning Operations section provides more detail on the deployment infrastructure.

The current surrogate models deliver predictions in milliseconds, making them practical for control-room operations and fast optimisation. Both models are in stable long-term deployment and have been integrated into operations, where they demonstrate reliable performance. The near-term work will focus on developing additional surrogates and exploring methods to enable them to simulate larger sections of the linear accelerator. Beyond their immediate diagnostic and optimisation roles, surrogates also provide the foundation for the optimisation and anomaly detection workflows discussed in the following sections.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD082

# TOWARD PARTICLE ACCELERATOR MACHINE STATE EMBEDDINGS AS A MODALITY FOR LARGE LANGUAGE MODELS

F. Mayet*, Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany
T. Hellert, A. Sulc, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
C. Tennant, Thomas Jefferson National Accelerator Facility, Newport News, VA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Understanding and diagnosing the state of a particle accelerator requires navigating high-dimensional control system data, often involving hundreds of interdependent parameters. We propose a novel multimodal embedding framework that jointly learns representations of machine states from both numerical control system readouts and natural language descriptions. This enables the translation of complex machine conditions into human-readable summaries while maintaining fidelity to the underlying physical system. The obtained embeddings are subsequently adapted to an open-weights large language model via cross-attention conditioning. We demonstrate a first implementation trained on European XFEL machine state data. This work covers the embedding model architecture, training methodology, and presents initial examples demonstrating the model's capabilities in action. Due to the general concept of machine state, the model can be easily adapted to other facilities and control system environments.

## INTRODUCTION

Particle accelerators are complex machines with numerous interconnected subsystems. Their state at any moment is defined by parameters like magnet currents, beam positions, and RF voltages, all monitored by control systems. We introduce the concept of *Machine State* as a novel modality for large language models (LLMs). Unlike traditional multimodal models that use e.g. images, we feed the machine state as tuples of control system addresses and values. This enables LLMs to interpret accelerator states and provide informed responses.

In this study, we use machine states from the European XFEL [1], derived from *Machine Catalog Files* containing thousands of tuples. These states are enriched with metadata like operation modes. We also present initial examples demonstrating the capabilities of our current model, which enhances the open-weights model `Mistral-7B` [2]. These examples showcase how the integrated system can provide state-informed responses to user queries.

Integrating machine state data with LLMs can revolutionize accelerator operations, enabling real-time troubleshooting and optimization. The following sections detail the embedding model for converting machine states into LLM-compatible formats and the cross-attention based adapter facilitating this integration.

* frank.mayet@desy.de

## RELATED WORK

In this work, we adapt the intermediate layers of a LLM by adding a vector generated by a multimodal adapter. This process of input adaptation is functionally a form of output control, which directly links our approach to the field of activation steering, also known as representation engineering.

Activation steering modifies an LLM's behavior at inference time by adding steering vectors to its internal activation states to guide generation towards desired concepts. For a survey of the field, see [3,4]. This technique has been successfully applied to control a range of abstract behaviors, with some examples like mitigation of toxicity [5] and improvements to model safety [6].

Our work demonstrates that a cross-attention adapter layer makes this steering connection explicit and dynamic by using machine states, passed as the adapter's input, to steer the model toward specific outputs that characterize the machine's state.

## MACHINE STATE

The state of the particle accelerator at a given point in time is approximated by all available readback and setpoint values the control system provides. These values include mostly machine-related properties such as magnet currents, but also beam properties and various other performance measures. A snapshot of all of these values constitutes what we consider, in the context of this work, a *Machine State*. As stated in the introduction, the goal is to introduce machine state as a second modality to the prompt of an LLM, in order to enable the LLM to interpret the state and generate state-informed answers to a user query. Existing multimodal models, such as GPT-5 [7], accept image data as a second modality. Here, the machine state is fed to the modified LLM as a list of tuples of the form `(address [str], value [float, int, bool, str])`. Each tuple contains a control system address, as well as the corresponding value, which can be numerical, boolean, or a text string.

For this particular study machine states of the European XFEL are used. The states are constructed from so-called *Machine Catalog Files*, which are machine state files used to document the machine state, as well as to restore machine setups. These files initially contain approximately 15000 `(address, value)` tuples, which are reduced to $\sim$ 5000 through redundancy removal and filtering. Additionally, we enrich each machine state with operational metadata including the requested operation mode and other user requirements.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD083

# OVERVIEW AND STATUS OF THE MACHINE LEARNING DATA PLATFORM PROJECT*

C. K. Allen†, C. McChesney, M. Davidsaver, M. Frauenheim, and L. B. Dalesio
Osprey Distributed Control Systems, LLC, Ocean City, MD, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

## Abstract

Osprey DCS is developing the *Machine Learning Data Platform* (MLDP), a public-domain, open-source product to support data science, machine learning and artificial intelligence applications for large particle accelerator facilities. Funding for MLDP development was obtained from the United States Department of Energy (DOE) through a Small Business Innovative Research (SBIR) grant. The project Phase II activity was completed in July of 2025; we present an overview then summarize project activity, achievements, and status.

## INTRODUCTION

### Background

The project was formally started in July 2022 with an SBIR Phase I award from the United States DOE, Office of High Energy Physics [1]. After the 6-month Phase I effort Osprey DCS was awarded a Phase IIA grant for continued MLDP development. The Phase IIA duration was from July 2023 to July 2025. Reported here are a review of the research and development efforts and accomplishments for the full 2.5-year project duration (Phase I and II).

### Product Description

The *Machine Learning Data Platform* (MLDP) is a product providing full-stack support for data science, Machine Learning, and Artificial Intelligence (ML/AI) applications at particle accelerator and large experimental physics facilities. It supports ML/AI applications from front-end, high-speed acquisition of heterogeneous, time-series data, through data archiving and management, to back-end analysis. The MLDP embodies a "machine-learning ready" platform for data analysis and ML/AI applications in diagnosis, modelling, control, and optimization of these facilities. It provides data scientists and applications a consistent, data-centric interface to archive data standardizing implementation and deployment of ML/AI algorithms to different operations configurations within the same facility, or between facilities. Being an open-source, public-domain project, the MLDP is intended for broadest possible impact by increasing accessibility and minimizing the required expertise for installation and operation.

The MLDP can also be deployed at user facilities for experimental data collection, archiving, and analysis. It is capable of acquisition and archiving of heterogeneous data from experimental equipment (e.g., images, arrays, structures, etc.) along with system hardware configurations (e.g., scalars, tables), control system process variables, and any metadata required for provenance. Thus, the MLDP can manage experimental data through its entire lifecycle, from acquisition and archiving, through analysis and investigation, to release and final publication.

Below is a succinct outline of the MLDP design, more complete descriptions are available in previous articles (e.g., see [2-3]). The *Data Platform* (see below) main repository site also contains substantial online documentation, including papers, reports, and presentations [4].

Referring to Fig. 1, the MLDP is composed of 3 components: 1) the *Aggregator* performing front-end data acquisition, correlation, and staging, 2) the *Data Platform* responsible for data archiving, management, and analysis, and 3) the *Web Application* providing remote archive interaction using a web browser. The Aggregator [5] is deployed within Experimental Physics and Industrial Control System (EPICS) [6]. The Data Platform is a standalone system featuring an online installation utility and simple deployment [7]; it provides the data science functionality post-acquisition. The Web Application loads into web browsers via a separate server, providing remote access and interaction with the MLDP archive without any programming or scripting requirements [8].

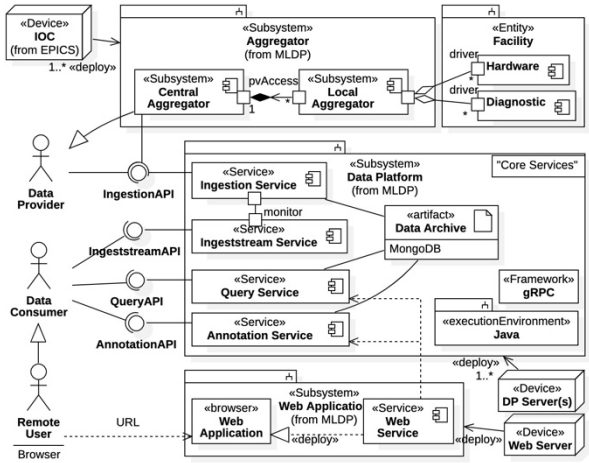



Figure 1: MLDP composite diagram.

Figure 1 provides a view the MLDP basic architecture, along with deployment and client relationships. The Aggregator system is distributed within the EPICS environment containing a *Central Aggregator* for coalescing and staging of acquired data. The Data Platform consists of multiple *Core Services* [9] interacting through an independent gRPC communication framework [10]. Each service has a gRPC Application Programming Interface (API) for interaction with clients, the archive, or each other. Data Platform *clients* are divided into *Data Providers* and *Data*

* Work supported by the U.S. Department of Energy, Office of High Energy Physics, SBIR Grant DE-SC0022583.
† email address allenck@ospreydcs.com.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD095

# PARTICLE ACCELERATOR SIMULATION USING GPUS IN ACCELERATOR TOOLBOX

J. L. Pons†, S. White, T. Rozier, European Synchrotron Radiation Facility, Grenoble, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

In storage ring light sources, two important performance parameters are injection efficiency and lifetime. Accelerator particle tracking in digital twins and simulators is an approach to efficiently derive these quantities and guide the optimization of the physical accelerator machine. These calculations are resource-intensive processes that often require high-performance computing. GPU usage can be an effective strategy to achieve good performance at a reasonable cost. This paper presents the current state of GPU particle tracking code implemented in Accelerator Toolbox (AT) using OpenCL or CUDA. The GPU kernel code is dynamically generated according to the lattice definition, allowing for optimal performance and selection of the symplectic integrator. We present accuracy and performance comparisons obtained with GPUs versus those obtained with AT CPU code. Trade-offs necessary to achieve optimal results are also discussed. Additionally, we discuss GPU kernel synchronization techniques needed by collective effect elements, as well as the object-oriented code structure to facilitate the integration of other GPU APIs.

## INTRODUCTION

Accelerator Toolbox (AT) is a collection of tools to model storage rings and beam transport lines [1]. 6-D modular tracking engine has been written in C/C++ for efficiency. However, some algorithms, such as acceptance calculations, require tracking a large number of particles over a large number of turns. Graphics Processing Units (GPUs) are especially efficient to solve this kind of problems where lots of identical tasks, using different input data, need to be performed. GPUs use the single-instruction multiple-thread (SIMT) execution model where multiple independent threads execute concurrently using a single instruction. GPU APIs (OpenCL and CUDA) also offer a way to dynamically compile a program and then to execute it on the hardware. Just-In-Time compilation (JIT) allows to generate an optimized code for a given lattice in order to get optimal performances as done in UFO [2]. Such strategy has also been chosen for AT. AT GPU uses a C++ GPU abstraction layer allowing the same code to be run on both OpenCL and CUDA platforms. Due to lack of double-precision floating point support, Metal API was not considered.

## ARCHITECTURE

### *Dynamic Code Generation (DCG)*

Unlike static code, which is written, compiled, and executed in a fixed manner, dynamic code can be generated, compiled, and executed on the fly according to the needs of the moment. For a given lattice only needed code is generated while minimising the number of execution paths by removing unwanted conditional branches or useless code. However, in order to maintain compatibly with CPU code and to avoid a strong compilation overhead (see Table 1), the AT *keep_lattice* flag can be used as shown below. At each *ring.track()* call, the 3 tasks (Table 1) are performed. When *keep_lattice=True* is used, only the third one is executed. It is especially useful for turn by turn tracking. When *keep_lattice=True* is used, only *T1, T2 ,R1 ,R2 ,[E|R]Aperture, Polynom[A|B]* and *KickAngle* coefficients of each element of the ring can vary between two *ring.track()* calls:

```
rout,  *_  ring.track(rin  ,nturns=1  ,use_mp=
at.MPMode.GPU, refpts=refs, keep_lattice=True)
```

Table 1: DCG Overhead (EBS, 4434 Elements)

| Task | OpenCL [msec] | CUDA [msec] |
|---|---|---|
| Code generation | 87.6 | 87.3 |
| Code compilation | 850.3 | 443.7 |
| Lattice loading on GPU | 3.1 | 0.6 |

OpenCL performs caching to avoid dynamic recompilation when the GPU program does not change. It explains, at least partially, why the compilation phase is slower than CUDA for a new lattice.

### *Object Oriented Implementation*

AT GPU is implemented in C++. An incomplete but representative element class diagram can be shown in Fig. 1. These classes are in charge of generating GPU code for the corresponding PassMethods. An AT PassMethod is basically a random function that transforms particle coordinates according to an element (drift, magnet, RF cavity, etc…). AbstractElement handle low level GPU memory transfers and access to python or Matlab variables. IdentityPass handle translation (T1, T2), rotation (R1, R2) and aperture checks ([E|R]Aperture). StrMPoleSymplectic4Pass handle multipoles (Polynom[A|B] and KickAngle). Other classes handle their own specificities.

### *Symplectic Integrators*

AT CPU uses the well know Forest/Ruth 4$^{th}$ order symplectic integrator [3]. AT uses (semi-)explicit implementation of this integrator. An advantage of DCG is the possibility to generate very easily an optimized code for a random integrator. AT GPU proposes several integrators including optimal 4th order (asymmetric) integrator [4]

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD096

# APOLLO: A FACILITY-SCALE DIFFERENTIABLE VIRTUAL ACCELERATOR AT FERMILAB FAST/IOTA*

N. Kuklev†, M. Wallbank, N. Banerjee, J. Jarvis, A. Romanov
Fermi National National Laboratory, Batavia, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

As the design complexity of modern accelerators grows, there is more interest in using advanced simulations that have fast execution time or yield additional insights like gradients. The FAST/IOTA facility has been working on implementing and experimentally validating an end-to-end digital twin that is both fast and gradient-aware, allowing for rapid prototyping of new software and experiments with minimal beam time costs. Our framework integrates physics and ML codes for linac and ring simulation through a set of generic interfaces between surrogate and physics-based sections. To reproduce device inputs and outputs, system state is exposed as a deterministic event loop in a specialized discrete event simulator architecture. Because Fermilab is undergoing control system transition, several APIs were implemented as final user interfaces - a fully asynchronous EPICS soft IOC, a gRPC-based Data Pool Manager (DPM), and legacy ACNET protocols. We discuss implementation details as well as challenges handling live data assimilation and future plans to extend modelling to main complex proton accelerators like PIPII and Booster.

## INTRODUCTION

Particle accelerator projects face increasing performance demands, resulting in tighter tolerances on lattice accuracy and stability. Achieving these tolerances requires both more advanced diagnostics and more commissioning time, which contributes to rising costs and lower beam availability.

A common approach to reduce commissioning risks is to implement realistic versions of key accelerator systems as simulations and use them to validate commissioning procedures. A recent successful example of this approach is the Advanced Photon Source Upgrade Project (APS-U), which had a commissioning plan that included advanced ELEGANT simulations, detailed task scheduling, and many validation studies, including within a dedicated EPICS control network [1, 2].

So far most commissioning simulation projects have focused on the specific needs and quirks of their respective machines, avoiding complexity through implementing only the physics models and control system behaviors relevant to a particular project. This approach is not scalable, especially for a research facility like FAST/IOTA, since our machines are in constant flux and undergoing upgrades as the various experiments and elements are installed. Moreover, Fermilab main complex is currently undergoing several major upgrades, PIP-II (Proton Improvement Plan II) and ACORN (Accelerator Controls Operations Research Network). These will result in a mix of ACNET and EPICS controls systems, with brand new beamlines and devices having to interface with those over 40 years old.

Given the above complexity, it is clear that significant validation and practice will be needed to ensure smooth PIP-II commissioning and operation. We require a tool capable of testing classic methods (such as optics measurements via LOCO [3, 4], orbit feedback, multi-wire measurements, and many others), but also amenable to training and validating a new generation of machine learning (ML) tools. Notable examples of the latter include Bayesian optimization (BO) [5–8], reinforcement learning (RL) [9], and phase space reconstruction [10]. ML methods are especially promising to address issues that have previously relied on operational experience or incomplete, qualitative simulations. This paper presents our recent progress in development of a virtual accelerator (digital twin) toolkit as a potential unified facility-scale solution for Fermilab and potentially other accelerators. We will first describe the key requirements and resulting architectural details. Then, we will show use cases and our first upcoming project, the IOTA proton injector commissioning.

## VIRTUAL ACCELERATOR

Throughout this paper, we will use the term '(surrogate) model' to describe the specific (approximate) implementation of an accelerator or system. For example, an ImpactX simulation and a neural network (respectively) that can track beam with space charge. We will use the term 'virtual accelerator' to describe a framework that exposes and coordinates both standard and surrogate models to form a complete simulated control and physics environment. In the industry, the term 'digital twin' is also popular. One can think of virtual accelerators as backend engines for a virtual test stand, whereby various programs and users can interact with the models in a configurable environment with an appropriate mix of control systems, simulation fidelities, and computational requirements.

Our plan for virtual accelerator development has three main goals aligned with both the Fermilab project priorities and the more global challenges described in the recent DOE GARD accelerator and beam physics roadmap [11, 12]:

- To provide a robust virtual test stand for testing tools and procedures, including commissioning and operational applications as well as research codes

* This manuscript has been authored by FermiForward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.
† nkuklev@fnal.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD104

# MeerKAT ANTENNA POSITIONER EMULATOR TEST BENCH PROJECT*

B. Ngcebetsha†, South African Radio Astronomy Observatory, Cape Town, South Africa

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Antenna Positioner Emulator (APE) project is being built to imitate the functionality of an antenna positioner (AP) servo system. The AP is the mechanical component mainly responsible for the mounting and pointing of the antenna structure on the MeerKAT radio telescope. The project aim is fault finding, obsolescence mitigation, software debugging, and technical training. APE will use existing AP building blocks with a new mechanical design. The project will be designed for remote access, involving network configuration. Internal interfaces include monitoring with electrical and internal infrastructure. The external interfaces include connections to Control And Monitoring software (CAM), time and frequency reference (TFR), development servers, and on-site hardware. The project is also focused on electronics design review, cable manufacturing, power supply considerations, mechanical aspects such as load tests and weight distribution. Project management activities include discussions on asset transfers, procurement, foundation design, and remote access considerations.

## INTRODUCTION

The Antenna Positioner Emulator (APE) is also known as the MeerKAT AP Testbench, and its purpose is to emulate AP (Antenna Positioner) servo system functionality. The aim is to do fault finding and root cause investigations in the case of issues on the live MeerKAT system [1]. Plans to mitigate obsolescence are going to be tested on APE as well as verification procedures for upgrades of the AP make the APE useful. The APE will also be the apparatus in routine start-up of AP spares such as the Antenna Control Unit (ACU), servo amplifiers, motors. The aim is to minimize time-loss in the case of component failure. The debugging of AP system software failures and software verification will take place as an additional layer of control software life-cycle on the APE. The APE will also be used in the training of maintenance personnel. The project follows the systems engineering approach, systems engineering involves documentations, discussions around user interface, and coordination with various stakeholders.

### *Interfaces*

The APE will interface with the physical infrastructure on-site, requiring a space in the area near to the location of the MeerKAT antennas in the Karoo. A complex with a three-phase power supply has been dedicated for this. It is also a requirement that the APE is sheltered from the Karoo climate environment (temperature, humidity, dust) as well as Radio Frequency Interference (RFI) in a similar manner to the MeerKAT AP. The APE will connect to the SARAO network for time service, remote control, and testing. When connected to the Control And Monitoring (CAM), it will run as a receptor within an observation subarray to test its behavior as is the case for any other AP. In Fig. 1 is an illustration of external interfaces of the APE.

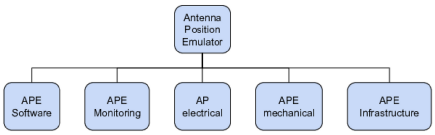



Figure 1: APE interface overview.

In Fig. 2 is included the external and internal interfaces of the APE. The summary in Table 1.

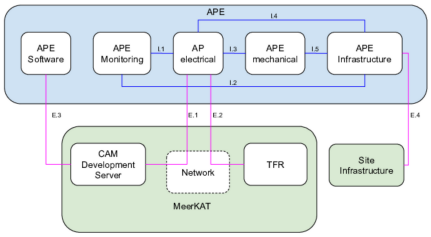



Figure 2: External and internal interfaces overview.

Table 1: Summary of Internal and External Interfaces in APE, see Fig. 1 and Fig. 2

| Identifier | Source | Destination |
|---|---|---|
| I.1 | APE monitoring | AP electrical |
| I.2 | APE monitoring | APE Infrastructure |
| I.3 | AP Electrical | APE mechanical |
| I.4 | AP Electrical | APE Infrastructure |
| I.5 | APE Mechanical | Infrastructure |
| E.1 | AP Electrical | CAM |
| E.2 | AP Electrical | TFR (or PTP source) |
| E.3 | APE Software | CAM development servers |
| E.4 | APE Infrastructure | Site Infrastructure |

The external interfaces are indicated by the teal boxes that sit outside the internal blue enclosing box.

## ARCHITECTURE

The APE architecture consists of several components with internal and external interfaces [2]. The APE is primarily controlled through the ACU and sniffers (embedded microcontroller devices on their own Virtual Local Area Network (VLAN) to minimize impact on networking, and a switch that handles the networking. There is also a local server in the same container as the APE where technicians can plug in laptops.

* Work supported by South African Radio Astronomy Observatory.

† bngcebetsha@sarao.ac.za

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD106

# ENHANCING SIRIUS FAST ORBIT FEEDBACK ACTUATORS USING IIR FILTERS

G. R. Cruz*, E. N. Rolim, A. F. Giachero, A. C. S. Oliveira, L. A. Pelike, G. S. Ramirez, D. O. Tavares, Brazilian Synchrotron Light Laboratory (LNLS/CNPEM), Campinas, Brazil

*Abstract*

Studies with the Sirius Fast Orbit Feedback (FOFB) system revealed that its power supplies were operating near saturation, which limited the achievable FOFB controller gain. Further analysis identified avoidable noise sources in Beam Position Monitors (BPMs) and adjacent magnet power supplies as the main causes of this behavior. This paper presents a model-based approach employing Infinite Impulse Response (IIR) filters to attenuate the effects of such noise sources on the control effort as well as equalization of the actuators' dynamic responses without any hardware modifications. The proposed method effectively reduces actuator effort and provides response equalization adding negligible phase shift in the band of interest, resulting in higher FOFB loop gain and improved rejection of orbit disturbances up to 1 kHz.

## INTRODUCTION

Detailed analyses of the Sirius Fast Orbit Feedback (FOFB) [1] performance have identified two main causes of operating with its fast corrector power supplies near saturation: (i) noise originating from the Beam Position Monitors (BPMs) switching scheme and (ii) noise in fast corrector power supplies coupled from other nearby magnet power supplies — both of which undesirably caused an increase in the control effort.

Although hardware modifications - such as redesigning the amplifier electronics - could mitigate the second source of noise, they are expensive, time-consuming, and potentially disruptive to user operation. As an alternative, this work presents a model-based digital compensation technique employing Infinite Impulse Response (IIR) filters to address both issues entirely within the existing infrastructure.

## CONTROL EFFORT REDUCTION

Operational data revealed that several fast corrector power supplies [2] operated close to their voltage saturation limits. Figure 1 (blue trace) shows the maximum absolute control effort ($u_{PS}$ in Fig. 2) across all units prior to any noise mitigation.

Spectral analysis showed that a substantial portion of the control effort was consumed by high-frequency noise within the internal proportional-integral (PI) regulation loops of the power supplies (inner loop in Fig. 2). Figure 3 (blue traces) shows the spectrum of a representative unit operating near saturation. Although the high-frequency content is far from the bandwidth of interest and naturally attenuated by the magnet inductance ($A$ in Fig. 2), not affecting the beam directly, it unnecessarily increases the voltage demands on the actuators and restricts the achievable FOFB performance due to saturation.

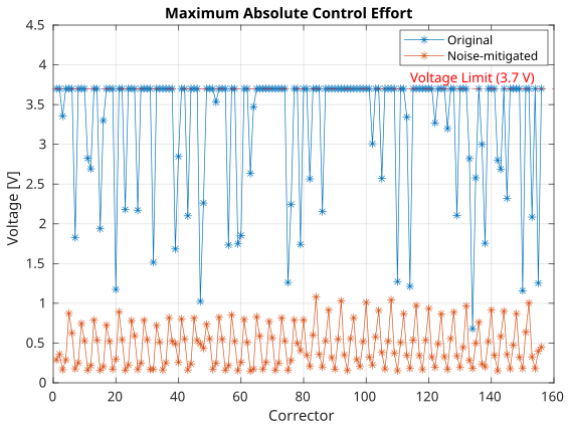



Figure 1: Maximum absolute control effort (voltage) across all fast corrector power supplies before and after noise mitigation. FOFB gain is set to $0.052/T_s$.

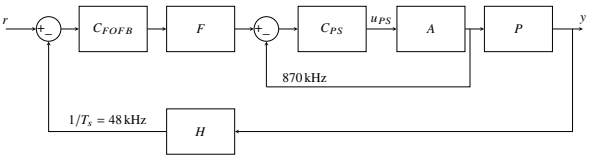



Figure 2: Block diagram of the FOFB loop. $C_{FOFB}$: FOFB controller. $F$: configurable IIR filters. $C_{PS}$: power supply controller. $A$: fast corrector magnet and power supply responses. $P$: vacumm chamber and beam orbit responses to varying magnetic field in the fast corrector magnets. $H$: BPMs response and network data distribution delay. $T_s$: FOFB loop sample time.

The nature of the observed noise and the strategies using IIR filters ($F$ in Fig. 2) adopted to mitigate its effects are described in the following subsections.

### *Notch Filters*

BPM signals are conditioned by analog front-end electronics [3] which perform RF channel switching between opposite BPM antennas at one-quarter of the FOFB loop update rate. This introduces deterministic noise at the switching frequency and its harmonics, as seen in Fig. 3 (blue traces) at frequencies multiple of 12 kHz.

To suppress these components, a narrow-band IIR notch filter was designed to attenuate both the fundamental switch-

* guilherme.ricioli@lnls.br

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD109

# STUDY DESIGN OF A MODEL-BASED CONTROLLER FOR TIME VARYING DELAY COMPENSATION IN A CRYOGENICS PROCESS

M. Pezzetti, T. Barbe, K. Buch, A. Onufrena, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

For the CERN High Luminosity HL-LHC project at CERN, a dedicated cryogenic test facility IT String is being commissioned. This cryogenic system includes a helium refrigerator, a 100 m-long cryogenic distribution line, and a low-pressure pumping system with a cold compressor Unit (CCU). A key challenge is the management of time-varying delays in the system's thermal and pressure responses, especially in the presence of dynamic components, which requires control due to its axial-centrifugal configuration and thermal constraints. This paper presents a study on the design of a model-based controller employing a Smith Predictor architecture to compensate for time-varying delays in the cryogenic process. The controller is specifically tailored to the cold compressor's operating sequence, which includes conditioning, regulation, and controlled stop phases—all of which impose strict operational and safety requirements due to the thermal and mechanical sensitivity of the system. Simulation results using simulation software are presented alongside experimental data from the commissioning of the IT String. The proposed control strategy enables stable operation of the pumping line during transient phases, such as magnet preloading and cooldown, and ensures the preservation of system integrity by limiting thermal gradients and compressor acceleration. The study gives some perspective in advance control model based delay compensation in complex cryogenic infrastructures like those at CERN.

## INTRODUCTION

For the CERN High Luminosity HL-LHC project at CERN, to validate the design and to assess the collective performance of the several new systems of the HL-LHC IT, a string of magnets (the HL-LHC IT string [1]) is installed in CERN main cryogenic test facility SM18. The objective of the HL-LHC IT String is to reproduce the operating conditions of the magnets in the LHC tunnel, to validate the technological choices and to determine the behaviour of the various subsystems, before the actual installation in the LHC tunnel. In addition to the technical validation [2], it is expected that the anticipated experience gained on the String will allow a shorter commissioning time when restarting the LHC. The IT STRING cryogenic process schematic is shown in Fig. 1.

The cryogenic system is required to provide a mass flow of around 8 g/s at a pumping pressure of 10 mbar. This allows to ensure a pressure of 16 mbar in the HL-LHC IT magnets, which corresponds to the saturation temperature of superfluid helium at 1.9 K necessary to ensure the superconductivity of the magnets [3]. During a magnet ramp-up, the helium flow from the pumping line increases from 8 g/s

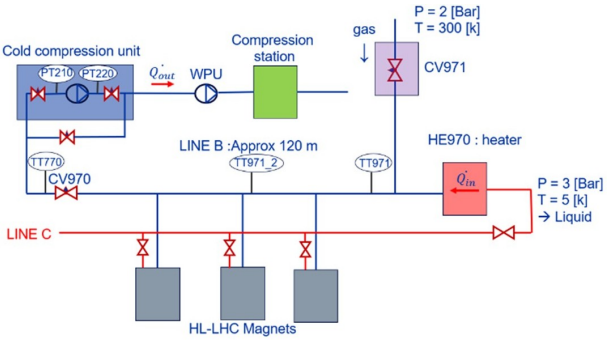



Figure 1: CERN IT STRING cryogenic process.

to 18 g/s; to compensate and stabilize the pressure, the Cold Compression Unit (CCU) is activated. A secondary control objective is to regulate the temperature at the inlet of the CCU in order to comply with the mechanical constraints of the system. This temperature regulation is achieved via the CV971 valve, which injects helium at 300 K to mix with the cold helium returning from the magnets.

## HELIUM COLD COMPRESSOR UNIT

In the CCU, pressure must track flow, defining three regions: surge, choke, and the desired operating region (Fig. 2). Surge, caused by excessive pressure for a given flow, can induce recirculation and damage the CCU. Choke, resulting from insufficient pressure, can overstress and break the blades. Both surge and choke are forbidden operating regions. Maintaining operation within the desired region requires controlling the inlet temperature.

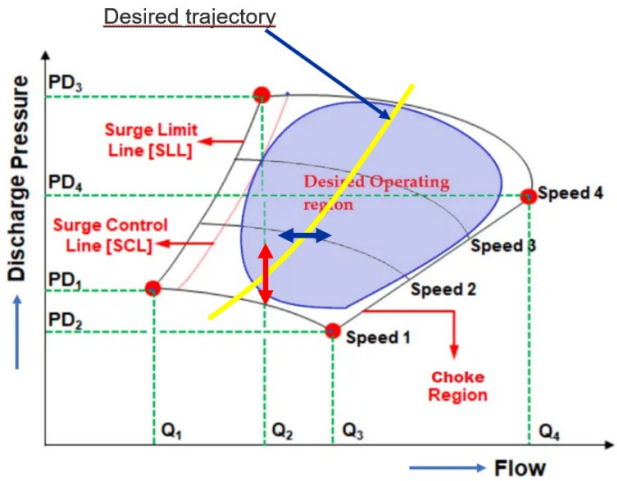



Figure 2: Desired operation regions.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WESV001

# XOPT AND BADGER: A MACHINE LEARNING ECOSYSTEM FOR REAL-TIME ACCELERATOR CONTROL AND OPTIMIZATION

R. Roussel* D. Kennedy, A. Edelen, S. Miskovich, Z. Zhang,
SLAC National Accelerator Laboratory, Menlo Park, CA, USA
N. Kuklev, Fermi National Accelerator Laboratory, Batavia, IL, USA

*Abstract*

Machine learning (ML)-based black-box optimization algorithms have demonstrated significant improvements in accelerator optimization speed, often by orders of magnitude. However, deploying these algorithms in real-time facility control remains challenging due to the specialized expertise and infrastructure required. To bridge this gap, we introduce the Xopt ecosystem, a versatile suite of tools designed to make advanced ML-based optimization accessible to the broader accelerator community. This ecosystem includes Xopt, a modular Python framework that facilitates the integration of ML-based optimization algorithms with arbitrary control problems, and Badger, a graphical user interface built on top of Xopt, which enables seamless deployment of ML algorithms in real-time control systems. The Xopt ecosystem has been successfully applied towards solving challenging real-time control problems at leading international accelerator facilities, including SLAC, LBNL, Argonne, Fermilab, BNL, DESY, and ESRF, demonstrating its effectiveness in real-world optimization tasks. In this presentation, we provide an overview of Xopt's capabilities and illustrate its impact through case studies from SLAC accelerator facilities including LCLS, LCLS-II, and FACET-II.

## INTRODUCTION

The need to solve complex optimization problems is widely prevalent in the field of accelerator physics. For example, accelerator control parameters must be tuned during accelerator operations to improve performance (so-called "online" tuning) or accelerator design parameters must be optimized in simulation to implement novel operational modes ("offline tuning"). These problems can be solved using a wide variety of conventional optimization algorithms, including Nelder-Mead Simplex [1] and Robust Conjugate Direction Search [2]. In addition, evolutionary algorithms, such as NSGA-II [3], can leverage the power of parallelized computation to solve multi-objective optimization problems. More recently, advanced machine learning based algorithms have been used to solve both online and offline accelerator optimization problems, most notably Bayesian Optimization (BO) [4].

Despite their wide applicability, connecting these algorithms to optimization problems can be challenging due to the diverse algorithmic and measurement interfaces. Additionally, as novel algorithms are developed, a standard set of community-implemented algorithms should be used as a benchmark to which algorithmic performance should be compared to. Finally, as algorithms become more complex, it becomes necessary to make these algorithms easily accessible to non-experts such that they can be used in a wide variety of applications. To facilitate the robust development, implementation, and delivery of complex optimization algorithms to accelerator operations teams, the machine learning community has contributed to the Xopt Ecosystem, a collection of modular python packages that provide a framework for developing and deploying advanced optimization algorithms towards solving arbitrary problems.

These packages (shown in Fig. 1) include Xopt [5] which provides a framework for connecting optimization algorithms to arbitrary optimization problems. Xopt also provides a collection of community developed algorithms that can be used off-the-shelf by non-experts or adapted to specific use cases by experts. These algorithms include those historically used by accelerator physicists today for online optimization (Nelder-Mead Simplex, RCDS, Extremum Seeking), algorithms used in massively parallelized contexts such as simulations running on high performance computing clusters (evolutionary algorithms including NSGA-II, CNSGA), and a wide variety of Bayesian optimization based algorithms. The other major package in the Xopt ecosystem is Badger [6], which provides a convenient graphical user interface for Xopt algorithms that can be used in facility control rooms. By separating the implementation of optimization algorithms and the user interface into separate packages, the same optimization algorithms can be used in both online and offline contexts. This modular design, coupled with a flexible framework for interacting with arbitrary problems has contributed to widespread adoption of Xopt across domestic (SLAC, LBNL, LANL, Fermilab, Argonne, BNL, Cornell) and international (STFC Rutherford Appleton Laboratory, ERSF, DESY, Pohang) accelerator facilities.

In this work, we provide an overview of the Xopt ecosystem with an emphasis on usage inside accelerator control rooms for solving online optimization challenges. We then highlight recent use cases at SLAC and other facilities. Finally, we describe future development directions, namely the integration of community developed algorithm specification standards and the addition of advanced algorithms from control theory and machine learning such as reinforcement learning.

* rroussel@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAG002

# MULTIMODAL DATA ACQUISITION SYSTEM AT MAX IV

V. Da Silva*, J. Just, M. Alcocer, K. Klementiev, A. Freitas, M. Lindberg
MAX IV Laboratory, Lund, Sweden

*Abstract*

The Balder beamline at MAX IV Laboratory, a state-of-the-art 4th generation synchrotron, is designed for X-ray absorption and emission spectroscopy. Delivering a high photon flux ($10^{13}$ ph/s), it supports *in situ* experiments, which require fast, high-quality data acquisition and support for sequential multi-technique measurements. This work presents a data acquisition (DAQ) system that combines X-ray absorption spectroscopy (XAS) and X-ray diffraction (XRD) within a single, synchronized experiment. At the core of the system is a Double Crystal Monochromator, operated by ACS SPiiPlusEC motion controller. This controller enables stable and rapid energy scanning via programmable motion trajectories, allowing sequential acquisition of energy spectra and diffraction patterns. Experiment synchronization is achieved via FPGA-based PandABox, which generates TTL signals based on the real-time motor position, enabling technique-specific pulse trains to be sent to the respective XAS and XRD detectors, precisely gated to the energy scan. The entire experiment workflow is orchestrated using Sardana through dedicated macros and controllers. User interaction is streamlined through a Taurus GUI, providing an intuitive drag-and-drop interface for sequencing and configuring each technique. This contribution outlines the system architecture, integration challenges, and benchmarking results, highlighting the enhanced experimental capabilities made possible by this advanced DAQ system at MAX IV.

## INTRODUCTION

The Balder beamline [1] is located at the 3 GeV storage ring of the Swedish 4th generation synchrotron, MAX IV, and is dedicated to X-ray absorption and emission spectroscopy in the energy range of 2.4–40 keV. With a photon flux of up to $1 \times 10^{13}$ ph/s, Balder is optimized for experiments under *in situ* and *operando* conditions. These experiments increasingly demand fast energy scanning with reliable data quality, while also benefitting from the integration of complementary techniques for improved analysis of sample dynamics.

Following the MAX IV standard, the beamline control system is based on TANGO [2], with Sardana [3] used for scan configuration and execution. To address the demand for multimodal and time-efficient data acquisition (DAQ), a system was developed that integrates X-ray absorption spectroscopy (XAS) and X-ray diffraction (XRD) into a single synchronized workflow.

A major challenge for such experiments is synchronization across multiple detectors and techniques, especially when scans must run quickly while minimizing acquisition overhead. Conventional approaches can lead to significant dead time, limiting experiment throughput and data quality.

By addressing synchronization challenges and reducing acquisition overhead, this multimodal DAQ system enables faster and more efficient XAS experiments at Balder, while establishing a framework for broader multimodal applications.

## SYSTEM OVERVIEW

The combined XAS–XRD experiment at Balder [1] is executed as a sequential, repeated measurement program: one or more XAS energy scans are interleaved with one or more XRD acquisitions taken at fixed energies. The general setup for the multimodal XAS-XRD system is shown in Fig. 1.

The central element of the Balder beamline energy selection is a fixed-exit Double Crystal Monochromator (DCM) supplied by FMB Oxford. The monochromator hosts two crystal pairs, Si(111) and Si(311), covering the energy range from 2.4 to 40 keV. The crystals are indirectly cooled through liquid-nitrogen-cooled copper plates to maintain thermal stability under the high heat load from the wiggler source. Energy selection is achieved via the Bragg motor rotation axis combined with the perpendicular translation of the second crystal (vertical beam offset), ensuring a constant exit beam height. Both axes are controlled by an ACS Motion Control SPiiPlusEC [4](EtherCAT-based) motion controller, with high-resolution incremental encoders providing real-time positional feedback for precise trajectory execution.

Synchronization between the beamline energy position (Bragg motor motion) and detector acquisition is performed by PandABox [5], an FPGA-based timing and control platform. At Balder, PandABox is configured to generate gating signals derived from the Bragg axis encoder. These gates define acquisition windows during which detectors integrate data. This hardware-based gating eliminates software latency, minimizes dead time, and ensures deterministic alignment of detector readout with the monochromator position.

For intensity monitoring of ionization chambers, ALBA Electrometer [6] is used as a current amplifier and to record the incident ($I_0$) and transmitted ($I_T$) beam intensity, enabling acquisition of XAS data. For diffraction measurements, Balder employs Eiger 1M (2 kHz readout), a two-dimensional hybrid pixel X-ray detector, which combines high frame rates with large dynamic range, making them suitable for *operando* XRD acquisition and as complementary technique during multimodal experiments.

## HARDWARE INTEGRATION

The beamline control system is based on TANGO Controls [2] framework, which provides a distributed architec-

* vanessa.silva@maxiv.lu.se

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAG003

# COMPREHENSIVE CONTROL SOFTWARE FOR TIMEPIX3 DETECTOR SYSTEMS IN EPICS ENVIRONMENT*

Kazimierz J. Gofron[1†]
[1] Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA

*Abstract*

This paper presents the development of a comprehensive control software suite for Timepix3 pixel array detectors within the EPICS environment. The system integrates ASI vendor software with four custom-developed EPICS components: (1) ADTimePix3, an areaDetector driver for direct detector control; (2) TPX3 Emulator IOC, for software-based testing; (3) TPX3 Serval IOC, for server management; and (4) TPX3 Histogram IOC, for real-time histogram processing. The system supports multi-chip configurations (1-4 chips), real-time data acquisition with TCP streaming, pixel masking, and detector health monitoring. Real-time event streaming is compatible with Spallation Neutron Source (SNS) ADARA infrastructure. The emulator supports RASNIK pattern simulation and .tpx3 data replay for training and testing. The system has been deployed across multiple Linux platforms, supporting Serval versions 3.0.0-4.1.1, providing a production-ready solution for Timepix3 detector control.

## INTRODUCTION

The Timepix family of pixel detector readout chips, developed by CERN, represents a significant advancement in hybrid pixel detector technology. Timepix3 detectors [1] are the latest generation in this family, offering high-resolution imaging capabilities with advanced event-mode readout and precise timing measurements. These detectors are increasingly being deployed in scientific facilities for applications ranging from X-ray imaging and neutron detection to particle physics experiments.

Timepix3 detectors support two distinct readout modes: event-mode and frame-based mode. Event-mode readout involves counting individual events, such as single photons, single X-rays, or single neutrons, with precise timing information. This mode enables raw data acquisition in .tpx3 file format and real-time TCP socket streaming, making it ideal for applications requiring detailed event analysis and real-time processing. In contrast, preview and image modes operate in frame-based mode, providing integrated images over specified time periods.

The Timepix3 detector provides comprehensive event characterization through two key timing measurements. Time-of-Arrival (ToA) information records the precise arrival time of particles, enabling temporal analysis and correlation studies. Time-over-Threshold (ToT) measurements record the energy deposited by particles in the detector. ToT records the duration, in clock cycles, that the signal generated by the deposited charge in a pixel remains above a user-defined threshold. This duration is proportional to the deposited charge, providing valuable energy information for each detected event.

The detector's advanced timing capabilities include event timestamping down to 1.5 ns resolution with a 260 ps Time-to-Digital Converter (TDC), enabling precise temporal measurements. The system employs one lower threshold for event detection, allowing users to define the minimum energy threshold for particle detection. The Timepix3 detector system typically consists of several key components: the detector hardware itself, a data acquisition server (Serval), emulation software for testing and development, and specialized processing algorithms for data analysis.

Integration into EPICS [2] control systems presents significant challenges due to the complexity of detector control, data acquisition, and real-time processing requirements. Each component requires sophisticated control interfaces and real-time monitoring capabilities to ensure reliable operation in scientific environments.

This paper presents the development of a comprehensive control software suite that addresses these challenges by providing a complete EPICS-based solution for Timepix3 detector control. The system consists of four integrated components that work together to provide seamless detector operation, data acquisition, real-time processing, and monitoring and testing capabilities. Previous work on the ADTimePix3 driver has been reported at ICALEPCS 2023 [3] and LINAC 2024 [4].

## SYSTEM ARCHITECTURE

The complete Timepix3 control system integrates ASI vendor software with four custom-developed EPICS IOCs (Input/Output Controllers) that work together to provide comprehensive detector control and data processing capabilities.

### *ASI Vendor Software Components*

The system utilizes two key ASI vendor software components:

**Serval Data Acquisition Server:** ASI's Serval server [5] (versions 3.0.0-4.1.1) provides the core data acquisition functionality, serving as the communication bridge between the detector hardware and control software. Serval handles raw data collection, detector configuration, and provides a REST API for external control. The server supports multiple

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under contract number DE-AC05-00OR22725. This research used resources at the High Flux Isotope Reactor and Spallation Neutron Source, a DOE Office of Science User Facility operated by the Oak Ridge National Laboratory.
† gofronkj@ornl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAG004

# HASMI: A CONFIGURABLE PYTHON FRAMEWORK FOR AUTOMATED HARMONIC ANALYSIS AND SCAN ORCHESTRATION AT EMIL

A. Balzer, A. Efimenko, P. Sreelatha Devi, K. Holldack
Helmholtz-Zentrum Berlin (HZB), Berlin, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

We present HASMI (Harmonic Analyzer State Machine Interface), a highly configurable Python framework designed to automate data acquisition and analytical tasks for optimizing coordinated undulator and monochromator movements at synchrotron beamlines. Built upon EPICS, HASMI features a comprehensive scan library and a command-line interface for managing multi-actuator scans. The framework offers extensive customization through a user-configurable database for defining scan types, actual scan parameters, and user preferences. Integrated analytical tools enable precise identification of undulator harmonic positions using advanced techniques such as convolution and customized cross-correlation. Beamline operations are automated and parameterized through a generalized state machine, which provides both an intuitive command-line interface and a dedicated EPICS interface. The state machine operates efficiently as a continuous background service alongside the monochromator IOC. By integrating multiple alignment scans, visualization, and automated data analysis into a streamlined 'one-button' procedure, HASMI significantly enhances the speed, performance, and reliability of beamline operations for users of the four hard X-ray CPMU17-DCM branches at EMIL.

## INTRODUCTION

At BESSY II, efforts are underway to establish a more holistic and unified software environment [1] based on Bluesky [2]. However, the current controls landscape is still heterogeneous: software such as SPEC, SpecsLab Prodigy and Igor Pro® as well as many other frameworks are deployed in day-to-day operations. With few exceptions, the Experimental Physics and Industrial Controls System (EPICS) [3] and its Channel Access (CA) protocol [4] provide the common denominator.

In this context, it is natural to seek reusable solutions for core tasks that operate independently of a particular experimental framework. An interoperable and configurable approach offers the greatest benefit, as it can be adapted to diverse requirements across beamlines. HASMI addresses this need by exposing its functionality through multiple access layers: an EPICS CA interface, a command-line interface (CLI), and a Python API. This flexibility allows operators to use HASMI directly through standard operator interfaces (OPIs), while integrators can embed it seamlessly into existing automation frameworks.

### *Undulator Lookup Table (ULT) Correction*

A central task in undulator–monochromator synchronization is the maintenance of a reliable lookup table (ULT) that maps undulator gap values to monochromator energies. In practice, the energy scale of the ULT often requires fine-tuning, which can be achieved by introducing a linear correction with slope and offset. The gap is then related to the setpoint energy $E$ by

$$\text{gap} = f(\text{slope} \cdot E + \text{offset}), \tag{1}$$

and the corrected photon energy follows from the inverse mapping as

$$E = \frac{f^{-1}(gap) - \text{offset}}{\text{slope}}. \tag{2}$$

Two reference scans are acquired at the beginning and end of a selected energy range. The harmonic analyzer evaluates both scans, identifies the respective harmonic peak positions with sub-sample precision, and derives an updated calibration in terms of slope and offset. These values are a proposed *correction*, which is stored in a manifest file together with metadata and the previously active calibration values.

The correction can then be applied to the corresponding EPICS process variables. This mechanism ensures that the undulator–monochromator relation remains consistent over time, compensating for drifts or mechanical changes in the beamline environment.

## TECHNICAL BACKGROUND AND DESIGN

### *Software Foundations*

HASMI builds on two modular Python frameworks under active development at HZB: a multi-actuator *scan* framework and an *asyncio*-based device orchestration framework, the *scheduler*. Both are lightweight, designed for beamline production tests and automation.

**Scan Library** The scan library provides flexible multi-actuator scanning with the capablilty of

- parallel scans over independent axes,
- concurrent execution of move sequences or other devices, and
- continuous-mode scans where one or more axes run at constant velocity while others are synchronized (continuous mode performance depends on the capabilities of the motion controller and IOC.)

A plugin architecture exposes hooks before, during, and after scan steps, enabling integration of custom controllers,

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAG006

# ADVANCING POSITION-BASED CONTINUOUS ENERGY SCANS AT MAX IV: EXPANDED BEAMLINE COVERAGE AND ENHANCED CONTROL INTEGRATION

L. Zhu*, Á. Freitas, B. Bertrand, J. Lidón-Simón, N. Al-Habib, M. Lindberg
MAX IV Laboratory, Lund University, Lund, Sweden

*Abstract*

The position-based continuous energy scanning system at MAX IV continues to evolve, delivering significantly faster and more consistent data acquisition for X-ray beamlines. Since its initial implementation on BioMAX, FlexPES, and FinEstBeAMS, the system has now been successfully deployed on three additional beamlines: SPECIES, HIPPIE, and NanoMAX. Non-linear trajectories for undulator have been implemented for FinEstBeAMS and FlexPES beamlines. Integration efforts are also underway at SoftiMAX and MAXPEEM beamlines. A major recent advancement includes the implementation of non-linear trajectories gap motion for both APPLE-II and in-vacuum undulators (IVU), supporting coupled-axis scanning in both soft and hard X-ray regimes. In addition to its native compatibility with Sardana/TANGO-based orchestration, the system now offers broader integration capabilities, allowing for seamless operation within alternative experimental control environments such as *contrast*. These developments further underline the system's modularity, flexibility, and potential to become a facility-wide standard for efficient, low-dose, and high-resolution scanning.

## INTRODUCTION

The 4th generation synchrotrons, compared to conventional X-ray sources, deliver tunable and extremely bright X-ray that enables non-destructive, high-resolution imaging and spectroscopy across disciplines such as materials science, biology, and nanotechnology. From an experimental perspective, one of the most common measurements at synchrotrons facilities is performing scans to measure some quantity as a function of X-ray energy. The efficiency and speed of scanning and data acquisition are of critical importance.

Previous works at MAX IV [1, 2] introduced the continuous energy scan, in which data acquisition is performed while continuously varying the energy at a constant rate (in eV/s), with coupled motions of the undulator and monochromator. The continuous energy scan significantly improves scan efficiency by dramatically reducing the dead time for starting and stopping of the scan axis (typically several motors), compared to the step scans. This allows for more efficient use of X-ray beam time in storage-ring facilities, enabling high-quality measurements even with millisecond exposure times. Shorter acquisition times also reduce radiation sample damage, a critical factor for X-ray sensitive materials such as biological materials and 2D materials. On the other hand, the high brilliance of 4th generation synchrotrons and the availability of fast detectors further enhance the achievable measurement speed by minimizing exposure and readout times. In the first implementation at MAX IV, only monochromator was operated using parametric trajectories [3], while the undulator gap was varied linearly during scanning. The non-linear parametric trajectories for the undulator was under consideration.

However, the assumption of a linear relationship between the undulator gap and its harmonic peak energy is no longer valid, particularly at low energy ranges or scanning a wide energy range. To address this, non-linear parametric trajectories for undulators have been developed as a Sardana [4] MotorController. This is deployed on the accelerator network, separate from the beamline network used for monochromator trajectories. With the ultimate goal of synchronizing both undulator and monochromator in parametric mode to maximize photon flux by ensuring the undulator harmonic peak tracks the monochromator pass energy. The trajectory motions of undulator and monochromator are started by software. A proof-of-concept prototype has also been demonstrated, where PandABox [5] triggers are used to start both trajectories motions simultaneously, thereby eliminating network latencies, ultimately pushing the boundaries of efficiency and speed. In parallel, efforts are underway to decouple components of the software from Sardana into standalone TANGO [6] devices, enabling integration with alternative experimental control environments.

## SYSTEM IMPLEMENTATION

### *Non-Linear Trajectories in Gap Motion*

The photon flux of different energies produced by an undulator is concentrated in a series of narrow harmonic peaks. The photons reaching endstation are filtered by the monochromator, which selects the desired energy. To maximize photon flux at the sample, the target energy must be aligned with the corresponding harmonic peak before passing through the monochromator. At high photon energies (hard X-rays), the relationship between the undulator gap and photon energy is approximately linear. Consequently, linear gap motion performs well in this regime, particularly when experiments require scanning only within a narrow energy window [7]. Since beamlines aim to maintain maximum flux, undulator tapering is not considered in this implementation.

* lin.zhu@maxiv.lu.se

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAR001

# THE INTEGRATION OF CUSTOM RADIATION TOLERANT ELECTRONICS IN INDUSTRIAL CONTROL SYSTEMS

M. Kalinowski†, A. Arias, M. Bes, E. Blanco, G. Daniluk,
J. Dumerger, A. Dutruel, B. Guncic, P. Klasa, A. Kostopoulos,
R. Mompo, J. Palluel, S. Pennock, I. Romera, J. Uythoven
CERN, European Organization for Nuclear Research, Geneva, Switzerland

*Abstract*

In the CERN accelerator complex, the conventional magnets are protected against overheating and power converter failures by a PLC-based system, the so-called Warm magnet Interlock Controller (WIC). In 2026, the systems installed in the SPS-LHC transfer lines will reach end-of-life after 20 years of successful operation. Furthermore, Siemens' announcement regarding the phase-out of the S7-300 series necessitates development of a second-generation magnet protection system. Initially, a solution based on the existing purely industrial configuration was considered, involving a simple upgrade of the PLC modules to the new Siemens S7-1500 series. However, due to the susceptibility of this new series to radiation-induced electronic effects, this approach was deemed unfeasible. As a result, a new control system architecture was explored, integrating both an industrial control processor and custom in-house designed radiation-tolerant electronics. To maintain overall system integration and ensure seamless interfacing with the existing SCADA layer, it was decided to retain the industrial CPU – Siemens S7-1516 PLC – as the process control master. For the deported units, operating in radiation-prone environments, a CERN-developed platform, Distributed I/O Tier, was chosen. The new system was recently installed in one of the beam transfer lines. The technical challenges and chosen solutions are described in this paper.

## INTRODUCTION

The Warm magnet Interlock Controller (WIC) is a PLC-based system. It protects the normal conducting magnets from overheating by switching off the corresponding power converter when a fault occurs (i.e. magnet overheating or low water flow) (Fig. 1). It can also be interfaced with the Beam Interlock System (BIS) to inhibit a beam injection or request a beam dump in case of a magnet or power converter fault in order to prevent beam-induced damage to accelerator equipment. The WIC is presently installed throughout the CERN accelerator complex, protecting in total 3100 magnets with 43 system instances. Only for a small fraction of the magnets, the WIC electronics are exposed to radiation.

It should be noted that no analogue values are transmitted or processed by the WIC. Only digital information is considered (either OK or NOT OK). Signals are implemented in fail safe logic, meaning that a logic OK is always represented by an active input/output.

† michal.kalinowski@cern.ch

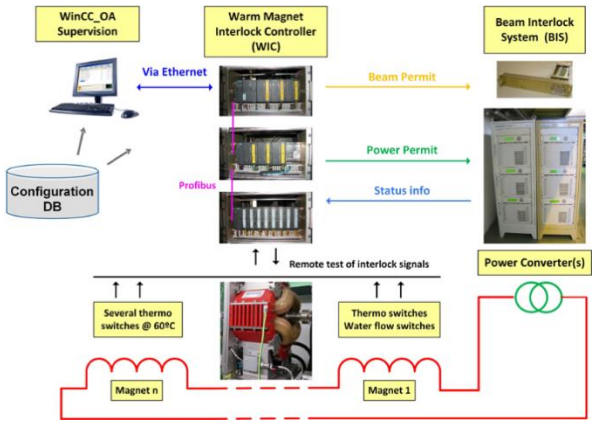



Figure 1: WIC system layout.

For magnet protection, two different types of digital protection signals are mainly used: thermo-switches (several in series per magnet) and/or pressure switches / flow switches. In both cases, they are connected via a magnet interlock box installed on top of the magnet, providing standard connectivity as well as remote testing functionality of the input signals.

The communication between the master crate and its deported units is ensured via a Profibus DP fieldbus. Depending on the configuration of the installation, the crates can be geographically distributed over several buildings or between surface and underground areas [1].

All protection functions are executed within the PLCs at the control layer. Additionally, the WIC system is provided with a SCADA (Supervisory Control and Data Acquisition) to monitor the status of an installation and logging of the data (i.e. faults are archived with accurate timestamps to retrace sequence of events).

During LS3 (Long Shutdown 3, July 2026 – June 2030) [2], the WIC's S7-300 PLC modules will reach the end of their operational life, after 20 years of successful operation. Furthermore, Siemens announced the phase-out of the S7-300 series in October 2023, with discontinuation effective 1 October 2025. After this date, it will no longer be possible to purchase new components from the S7-300 family; they will only be available as spare parts (e.g., to replace broken modules). This situation prompted the development of a second-generation WIC system CERN wide – further referred to as WIC2. The first systems to be addressed for obsolescence were those located in the SPS-LHC transfer lines, which protect over 700 resistive magnets. These were the earliest WIC systems deployed in 2004 (TT60, TI2, TT40, TI8) and 2005 (TT41) (Fig. 2). These systems pose an additional challenge for the design, as the deported

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAR002

# ENABLING HIGH REPETITION RATE SCIENCE: UPGRADING HARD X-RAY EXPERIMENTAL INSTRUMENTS CONTROLS FOR LCLS-II-HE*

M. Ghaly†, V. Esposito
SLAC National Accelerator Laboratory, Menlo Park CA, U.S.A

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory has been undertaking a major project that builds on the foundation of the LCLS-II project. The LCLS-II High Energy (LCLS-II-HE) upgrade is designed to push the capabilities of LCLS-II even further by increasing the energy of its superconducting accelerator to 8 GeV (up from 4 GeV in LCLS-II), enabling the production of even shorter and more intense X-ray pulses for cutting-edge scientific experiments. One of the project's Key Performance Parameters (KPPs) is the delivery of a high-repetition-rate capable Hard X-ray (HXR) experimental instrument that can perform experiments with the LCLS-II-HE beam. Meeting this requirement has driven the need to upgrade the existing HXR experimental control system. This paper will focus on the scope and progress of that upgrade effort, which builds on the LCLS-II controls architecture integrating new hardware and software components. A major part of the upgrade includes the implementation of the Preemptive Machine Protection System (PMPS), which is integrated across all key control subsystems including motion, optics and vacuum to ensure safe beam delivery. The upgraded system is also designed to support dual-mode operation and beamline multiplexing to meet evolving experimental demands. In addition, as the project approaches installation and transitions into its final phase, the challenges encountered and mitigation implemented to ensure successful delivery are presented.

## INTRODUCTION

The Linac Coherent Light Source at SLAC National Accelerator Laboratory has recently completed the LCLS-II upgrade, delivering a superconducting linac capable of operating at repetition rates up to 1 MHz. The ongoing LCLS-II High Energy (LCLS-II-HE) project will extend this capability by increasing the accelerator energy to 8 GeV, enabling even shorter and brighter X-ray pulses for scientific experiments.

Specifically, the X-ray Pump Probe (XPP) instrument in the Near Experimental Hall represents the project's Threshold KPP — the minimum scope required to achieve Critical Decision 4 (CD-4), the formal Critical Decision for project completion as defined in DOE O 413.3B [**doe-4133b**]. Meeting this threshold means delivering a fully upgraded, high-energy-capable instrument, including all main controls systems and a functional endstation ready for user operations.

In contrast, the upgrades for the Dynamic X-ray Scattering (DXS), Macromolecular Femtosecond Crystallography (MFX), and Coherent X-ray Imaging (CXI) instruments in the Far Experimental Hall are Objective KPPs — they represent additional capability goals that, while highly desirable, are not required to meet the minimum project deliverable. As a result, these instruments are progressing through design, procurement, and production of control system components, but their installation may be phased or deferred to align with available resources and overall project schedule.

This distinction allows the project to focus resources on completing XPP to full operational readiness while still advancing the broader upgrade path for the remaining HXR instruments.

To fully leverage the capabilities of LCLS-II-HE, the upgraded controls must provide safe and reliable high-repetition-rate operation with machine protection and efficient beam delivery. This includes support for dual-mode operation and beamline multiplexing, both of which impose new demands on the control system architecture. This paper presents the scope, design, and installation progress of the XPP instrument control system upgrades. The upgrades build directly on the LCLS-II Experiment Control System (ECS) architecture, extending its hardware and software components to meet the increased demands of high-repetition-rate science.

## UPGRADE SCOPE

### *Background*

The LCLS Experiment Control System (ECS) team has developed the LCLS-II controls systems to provide a robust and standardized control architecture for experimental instruments in preparation for the higher power and complexity introduced by LCLS-II beam[**lcls-ii-controls**]. ECS integrates industrial control platforms, Experimental Physics and Industrial Control System (EPICS) SCADA, and Python-based user interfaces to deliver safe and flexible operation across vacuum, motion, optics, and laser subsystems.

Core features of LCLS-II Experiment Controls include:

- Beckhoff PLC-based control with ad-hoc TwinCAT libraries for vacuum, motion, and preemtpive machine protection.
- EPICS for process variable management and SCADA functionalities
- Python toolchain and GUIs (PyDM, Typhos, Lightpath and Hutch Python) for experiment automation and diagnostics.

* THIS WORK WAS SUPPORTED BY THE U.S. DEPARTMENT OF ENERGY, OFFICE OF SCIENCE, UNDER CONTRACT NO. DE-AC02-76SF00515.

† mghaly@slac.stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAR003

# UPGRADE OF THE LHC VACUUM CONTROL SYSTEM TOWARDS THE HIGH LUMINOSITY LHC ERA

G. Pigny†, A. Gutierrez, A. Giannoulas, A. Rocha, B. Bayliss, I. Lobato, J. Rebelo, J. De la Gama, L. Zygaropoulos, L. Cantu, N. Chatzigeorgiou, R. Ferreira, S. Soares
CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The High Luminosity LHC (HL-LHC) project has initiated a comprehensive upgrade of the LHC vacuum control system. Much of the vacuum control hardware, installed at the beginning of the LHC in 2008 or even dating back to the LEP era in the 1990s, was becoming obsolete and required modernization.

Additionally, the new HL-LHC operating conditions will induce higher radiation levels; therefore, new radiation-tolerant electronics are required in the arcs and dispersion suppressor areas, as well as radiation-hard equipment in the matching sections.

Moreover, during the third long shutdown (2026-2030), the matching sections around the ATLAS and CMS experiments will need to be extensively modified with the insertion of new systems. The vacuum control system will be relocated to new underground galleries, with a significant increase in vacuum control equipment.

This upgrade, progressively implemented during each year-end technical stop and long shutdown, spans from 2019 to 2030. Hundreds of controllers are involved, along with significant control software design and refactoring, while keeping the vacuum control system fully operational.

This paper provides an overview of the new designs and technological solutions chosen, the radiation hardness assurance applied, the evolution of the vacuum control system architecture, and the main control software upgrades to ensure decades of reliable operation. Furthermore, it presents the current progress, challenges, and outlines future activities planned until 2030.

## UPGRADES OF VACUUM CONTROL EQUIPMENT

### *Fixed Pumping Group Controllers*

228 fixed pumping groups are installed in the LHC tunnel. They are composed of turbo molecular pumps, roughing pumps, intermediate valves and vacuum gauges. They are almost all used for insulation vacuum (IV) pumping for the cryogenics distribution line (QRL), the cryogenic magnets or the cryogenic RF cavities. Only a few are used for the beam vacuum (BV) pumping in the dump lines for the dilution kicker tanks (MKB), or other smaller installations such as the Beam-Gas Curtain, Gas Injection or Beam-Gas Ionization Systems. Table 1 provides the number of fixed pumping groups per area and per system in the LHC.

The initial architecture implemented at the beginning of the LHC was based on Siemens® S7-300 PLCs. Each controller crate, consisting of one CPU and I/O modules, controlled two fixed pumping groups with their associated gauges. In the arcs, these controllers also read the insulation vacuum gauges of the QRL and magnets, as well as the beam vacuum gauges. These controllers were also responsible for generating vacuum alarm signals for the cryogenics system in the event of deteriorated insulation vacuum levels. Several of these controllers were then connected as slave PLCs on an MPI network to a master Siemens® S7-400 PLC before connecting to the Ethernet Technical Network (TN) and SCADA [1].

Table 1: Number of Fixed Pumping Groups per Area and per System

| Area | QRL (IV) | Magnets (IV) | Cavities (IV) | MKB (BV) | Others (BV) |
|---|---|---|---|---|---|
| LSS | 17 | 54 | 4 | 14 | 11 |
| DS | 16 | 16 | - | - | - |
| Arc | 48 | 48 | - | - | - |

Because the controllers were located in radiation-safe areas far from the fixed pumping groups, the long cable lengths led to significant voltage drops. To address this issue, a local crate was installed in the tunnel for each pumping group. This crate relayed the control commands from the controller and provided local 24VDC in both the Long Straight Section (LSS) and arcs. In the arcs, the local crate also housed the radiation-tolerant turbomolecular pump controller.

Due to increased radiation levels, the obsolescence of PLC modules, and an architecture that combined the control of pumping groups for the QRL and magnets, gauges for both beam and insulation vacuum, as well as vacuum alarms for the cryogenics system, an upgrade of the fixed pumping group controllers started during the second long shutdown (LS2). This upgrade was carried out in anticipation of the third long shutdown (LS3) activities and continued during the subsequent year-end technical stops YETS22-23 and YETS23-24, finishing in YETS24-25.

The new controller crate is based on cost-effective Siemens® S7-1200 PLCs, with each unit dedicated to controlling a single pumping group with its associated gauges, if present. A full decoupling of the QRL and magnet pumping groups was achieved, albeit at the cost of a higher number of fixed pumping group controllers. Each controller is now directly connected to the TN for communication with the SCADA, replacing the previous master-slave architecture (see Fig. 1 – LSS + ARC VPG PLC). The local crate was also redesigned to be radiation-tolerant by removing active

† gregory.pigny@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAR004

# INTEGRATION OF NEW CRYOGENIC PLANTS INTO AN EXISTING CONTROL SYSTEM: A SCALABLE AND STANDARDIZED APPROACH

J. Fernandez-Cortes*, E. Blanco Viñuela, O. Kozinski, Ph. Durand, J.M. Beckers,
CERN, Geneva, Switzerland

*Abstract*

The High-Luminosity upgrade of the Large Hadron Collider (HL-LHC) requires the integration of new cryogenic plants without interrupting the operation of the existing infrastructure. This paper presents a scalable and standardized approach to modernizing the control system to meet this challenge, ensuring flexibility, interoperability, and long-term maintainability. The solution leverages commercial off-the-shelf (COTS) components, including PLCs programmed using IEC 61131-compliant languages, to enable modular deployment and reduce development complexity. The adoption of CERN's UNICOS framework streamlines implementation, enforces consistency, and facilitates seamless interaction between new and legacy subsystems. We describe the control system architecture and the design principles that support scalability and adaptability. Operational constraints, reliability requirements, and automation consistency are addressed through formalized engineering practices, including database-driven specification, virtual commissioning, and continuous integration workflows. This work provides a reference model for large-scale cryogenic system upgrades, demonstrating how industry standards and automation frameworks can support integration while ensuring long-term operational sustainability.

## INTRODUCTION

The Large Hadron Collider (LHC) at CERN is the world's largest particle accelerator, a 27-kilometer underground ring where protons collide at unprecedented energies to explore the fundamental laws of nature. Since the first beams circulated in 2008, the LHC has delivered groundbreaking discoveries, including the observation of the Higgs boson. To extend its physics reach, CERN and its international partners are now pursuing the High-Luminosity LHC (HL-LHC) project [1], an ambitious upgrade that will increase the collider's luminosity by a factor of five to seven. By producing far more collisions, the HL-LHC will allow researchers to collect larger datasets, probe rare processes with greater precision, and push the boundaries of our understanding of the universe.

Meeting the induced thermal loads in high-radiation areas requires two new cryogenic plants. Their integration must occur without disrupting ongoing LHC operation and must preserve long-term maintainability and full compatibility with the existing cryogenic system, including the control system.

The cryogenics team has operated the LHC cryogenic plants for nearly two decades, accumulating extensive operational and maintenance experience. This expertise has shaped the methodology applied to the new HL-LHC installations. To ensure consistent operator experience and streamlined life-cycle support, the new systems reproduce the proven functional layout: a warm helium compressor station, a helium dryer, a surface cold box, and a tunnel cold compressor. Beyond equipment compatibility, controls and operational conventions are also preserved to guarantee seamless interaction with the legacy systems.

This paper focuses on the approach taken to engineer the new cryogenics control systems for the LHC cryogenics production plants using commercial off-the-shelf (COTS) components and CERN's Unified Industrial Control System (UNICOS) framework to ensure standardization, scalability and maintainability. Starting from process functional specifications provided to the supplier, the cryogenics team transformed the requirements into structured, traceable documents aligned with UNICOS conventions. A repeatable integration method based on template-driven code and supervision configuration, model-based virtual commissioning, and GitLab-based quality control ensured compatibility and reduced integration risks. The following sections describe the project from the point of view of controls emphasizing the role of the UNICOS framework, integration methodology, and engineering practices. For clarity, this project employed the UNICOS-CPC package for Continuous Process Control, although for simplicity it will be referred to as UNICOS throughout the paper.

## LHC CRYOGENICS AND THE HL-LHC PROJECT

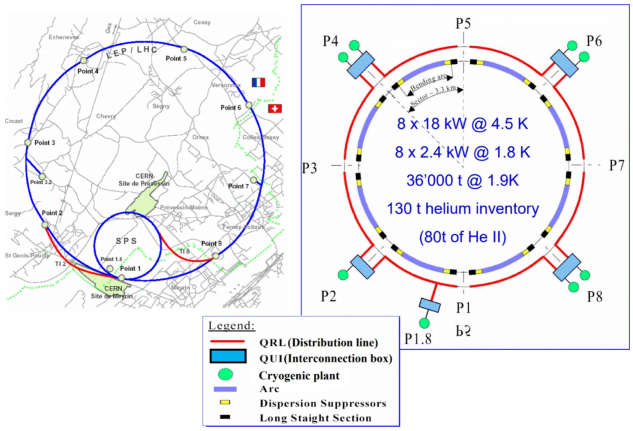



Figure 1: LHC cryogenics architecture

The cryogenic system of the Large Hadron Collider is one of the most extensive ever built, providing the ultra-low

* jesus.cortes@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THAR006

# BEAMLINE CONTROLS EXPERIENCES UNDER THE APS UPGRADE PROJECT

K. Lang, Argonne National Laboratory, Argonne, IL, United States

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The upgrade of the Advanced Photon Source (APS) included the design and construction of eight brand new beamlines and significant reconstruction of an additional fifteen existing beamlines. With a monumental amount of new support being required, the APS Beamline Controls group took the opportunity to evaluate the hardware we were using and the ways in which we develop and deploy beamline support.

This is an overview of the outcomes of those discussions, alongside lessons learned during the actual implementation of such planning. We discuss the hardware chosen, changes in the configuration and management of Input-Output Controllers (IOCs), the challenges in development of support tools, and the setup of data acquisition to fully take advantage of our new beam.

## INTRODUCTION

The Advanced Photon Source (APS) instrument at Argonne National Laboratories (ANL) recently finished a complete replacement of the electron storage ring in order to increase brilliance and coherence of the beam. As part of properly taking advantage of this new electron beam, a number of experimental beamlines were dismantled and rebuilt from the ground up, alongside other beamlines which were provided funding for new instrumentation.

This represented the single largest on-boarding of new controls equipment since the construction of the APS and required a unified Beamline Controls (BC) group strategy to be able to reasonably deploy support for all these items, while still enabling techniques that could leverage the capabilities of the new accelerator. Further, beamlines would be receiving xray detectors larger and faster than any had been previously, leading to an expected order of magnitude increase in the rate of data generation by the APS beamlines (see Fig. 1). To make sure the APS was equipped to handle this data load, members of the BC group were brought in alongside members of several other groups to create the Beamline Data Pipeline (BDP) group, which developed additional strategy in regards to data acquisition, data management, and data storage.

## BEAMLINE CONTROLS

The primary concern underlying all decisions on what hardware would be useful to the future beamlines was what would be the controls platform that was integrating those devices. For the entirety of the APS's operation, the standard system for operating beamline controls was an embedded controller using the Versa Module Eurocard (VME) computer bus and hosting the vxWorks real time operating system. On this, the Experimental Physics and Industrial Control System (EPICS) is then run, altogether known as an Input/Output Controller (IOC), being able to communicate and control attached devices.

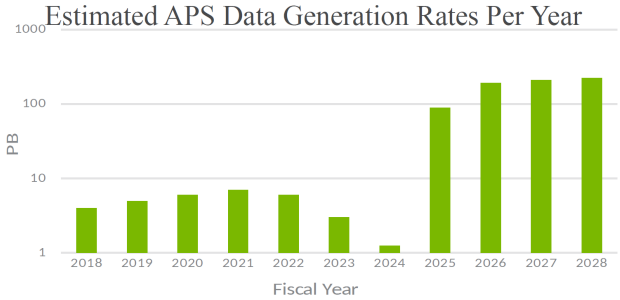



Figure 1: Estimated data rates.

While vxWorks IOCs had operated extremely reliably during the APS's operation and had significant existing support for hardware, a number of potential challenges were identified for continuing to use them going forward.

As a legacy platform, the supply chain developing new controls boards for VME had become constrained. It could become difficult to source the hardware necessary for any new deployments or to quickly find replacement controls when needed. Additionally, the licensing cost structure for being able to compile IOCs for the vxWorks operating system was looking to change, increasing the cost of active development for those systems. And finally, the group in charge of controls for the APS electron storage ring had already themselves decided on moving away from VMEs, reducing the capability to share development across the institution.

Because of this the BC group ultimately decided against continuing to use the VME environment as a primary development platform. With the EPICS software having expanded out onto many different platforms and the vast majority of our applications not requiring a real time operating system, Linux was determined to be a suitable host for IOC production.

### *Linux IOCs*

In transitioning to Linux-based IOCs, it was necessary to not only determine hardware that would be compatible, but also how to manage large-scale deployment of these IOCs. VME boards exist to connect or communicate with numerous different devices, but the APS does not have a critical mass of similar connections for Linux computers.

In surveying the equipment that scientists required in order to do their experiments, it was found that the vast majority could be accomplished through either Ethernet or serial communication standards. Through the use of Moxa networked serial servers, serial communications could be routed

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBG001

# ESRF'S SOFTWARE STRATEGY FOR HIGH-THROUGHPUT DETECTORS: SCALABLE DATA ACQUISITION WITH REAL-TIME PROCESSING

M. Dubouchet, S. Debionne*, A. Homs Puron, L. Claustre, V. Rey Bakaikoa[1]
ESRF, Grenoble, France

*Abstract*

The ESRF EBS upgrade dramatically increased the X-ray photon flux for the beamlines, thereby opening up new experimental opportunities but also requiring faster and more advanced 2D detectors. At the ESRF, 2D detectors have been integrated through LIMA [1], a framework that provides a unified control and acquisition Tango APIs for our control software, BLISS. To meet the challenges of new 2D detectors, a comprehensive data strategy has been developed. This includes the redesigned data acquisition framework, LIMA2 [2], which aims to support an increasing data throughput with a scalable design, facilitates Online Data Reduction (ODR) tasks and provides data access for Online Data Analysis (ODA) software. In addition, new detectors introduce new data models, such as multi-band, sparse, sub-byte or event-based data which need to be handled appropriately. In the framework of ESRF's data strategy, we give an update on the recent developments of the LIMA2 project components: the core library architecture and the hardware plug-ins for Dectris, Rigaku XSPA-1M detectors, in addition to the PSI Jungfrau and the ESRF Smartpix. We also showcase the processing plugin model illustrated by the Fast Azimuthal Integration (FAI) and X-ray Photon Correlation Spectroscopy (XPCS) techniques. We demonstrate LIMA2's new capabilities with real-world applications on ID10-COH and ID15A and present the results. We conclude with our vision on future data models and their integration with the control, visualization and processing software.

## INTRODUCTION

ESRF's data strategy covers the need for IT infrastructure and software throughout the life cycle of the data produced on the beamlines from the acquisition, processing, storing, cataloguing to publishing on the user data portal. In this paper, we focus on two-dimensional (2D) data from high-performance detectors for photon science applications, which pose a major challenge given the sustained high data rates (of a few gigabytes per second, GB/s) they produce.

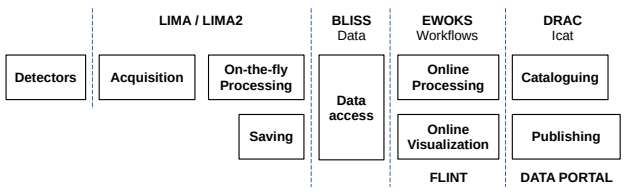



Figure 1: Data flow and software responsibilities.

In this context, Fig. 1 shows that LIMA2 is tasked with acquiring two-dimensional (2D) data from high-performance detectors. While retaining the primary objective of providing a unified control interface for 2D detectors, LIMA2 is a significant redesign of LIMA that addresses the need for scalability in response to the continuous increase in the data rate of high-performance 2D detectors. Furthermore, LIMA2 is designed to cope with new on-the-fly data processing needs, either for Online Data Reduction (ODR) or Online Data Analysis (ODA) to provide a fast, active feedback to the user and the experiment control system.

Blissdata [3] is the common interface that gives access to the scientific data during and after experiment. The data are accessed either from the LIMA/LIMA2 buffers or from file.

Further down the line, EWOKS [4] is used to automate additional data processing. LIMA2 and EWOKS workflows complement each other: while LIMA2 provides computation and data storage close to the data sources with dedicated infrastructure to keep latency low, EWOKS is more flexible and leverages the processing power of our computing cluster.

This paper presents the recent evolution of the design, the additions to the core, detector support and pipeline implementations as well as two applications on ESRF beamlines.

## SOFTWARE DESIGN

The LIMA2 design has been driven by two fundamental principles:

- the ability to scale the number of back-end computers depending on acquisition and processing needs, and
- the implementation of dedicated, optimised processing pipelines to provide reduced, low-latency data.

The framework is meant to be extensible either to support new detectors, new pipelines and/or new topology configurations, and provides the corresponding hardware and pipeline interfaces.

### *Distributed Acquisition*

Depending on the capabilities of the detector, different system topologies and frame dispatching are foreseen in LIMA2 [2]. The simplest topology includes a single back-end computer receiving all data from the detector.

Some modular detectors enable sending data from subsets of the detector pixels to different computers, what we call partial frame dispatch. Other detector systems can send complete frames alternating among multiple computers, referred to as full frame dispatch. Figure 3 represents the data processed by each receiver. Finally, we foresee the case of legacy systems that do not allow multiple back-end

* samuel.debionne@esrf.fr

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBG002

# FELIX, THE ATLAS READOUT SYSTEM: FROM LHC RUN 3 TO RUN 4

R. Luz * on behalf of the ATLAS TDAQ Collaboration
Argonne National Laboratory, Lemont, IL, USA

*Abstract*

After being successfully deployed to read out a subset of the ATLAS detectors during LHC Run 3 (2022-2026), FELIX will serve all ATLAS detectors in LHC Run 4 (2030-2033). FELIX is a router between custom serial links from front-end ASICs and FPGAs to data collection and processing components via a commodity switched network. FELIX is also capable of fixed-latency forwarding of the LHC clock, trigger accepts, and resets received from the TTC (Timing, Trigger and Control) system to front-end electronics. FELIX uses FPGA-based PCIe I/O cards installed in commodity servers. To cope with the increased data rate expected after the major upgrade to the LHC and the detector in between runs, the FLX-712 Run 3 PCIe Gen3×16 card, based on an AMD Kintex Ultrascale XCKU115 FPGA, will be replaced with the FLX-155, a bifurcated 2× PCIe Gen5×8 card equipped with an AMD Versal Premium VP1552 FPGA/SoC. Firmware installed on the FPGA and software running on the FELIX server are also being upgraded to handle the increased data rate.

## INTRODUCTION

The ATLAS experiment [1] is located at the Large Hadron Collider (LHC) at CERN. Particle bunches are accelerated and collided at the LHC. The time between collisions of the bunches can be as low as 25 ns so that the 40.079 LHC clock is synchronous to the bunch collisions. The ATLAS experiment uses the clock to sample the detector signals, and every bunch crossing can be identified uniquely with an identifier (BCID) and Orbit ID. The ATLAS experiment detects and analyses the collision results, generating vast amounts of data that need to be filtered out and processed. ATLAS has been operational since 2009, alternating between running periods and long shutdowns. The latest long shutdown (LS2) took place between 2019 and 2022, after which the current run, Run 3 [2], followed. Run 3 will end in June 2026, at which point Long Shutdown 3 (LS3) will start in preparation for Run 4, also known as, High Luminosity LHC (HL-LHC) [3]. The upgrades will include, among others, a new, all-silicon, Inner Tracker (ITk) detector [4, 5]. The expected start date of Run 4 is 2030.

During Run 3, the 40 MHz collision event - data from the ATLAS detector that corresponds to a single collision of bunches - rate is reduced first to 100 kHz of potentially interesting events and then down again to 3 kHz, at which point the data is permanently stored. This is possible due to a two-level trigger system used to select the events: the first, called Level-1 (L1) and based on custom-built electronics, uses information from the calorimeter and the muon detectors; the second, a software-based High-Level Trigger (HLT), uses higher granularity calorimeter and muon detector information, as well as, inner detector tracking information. In Run 4, the event rate of the first-level trigger will increase to 1 MHz, while for the second-level, a new commodity hardware-based high-level trigger will be implemented to achieve a maximum output rate of 10 kHz [6].

FELIX, the Front-End Link eXchange readout system, has played a significant role during Run 3, reading out the new detectors and trigger systems installed during LS2, that includes the new muon detector systems New Small Wheel (NSW) [7] and Barrel Inner Small (BIS7/8) [8], the new trigger system for Liquid Argon (LAr) calorimeter digital readout [9], and the calorimeter hardware trigger (L1Calo) [10]. Other detector readouts continued to use the legacy readout [11], already part of Run 2. In Run 4, FELIX will become the sole readout system interfacing with all ATLAS detector and trigger systems. The installation of FELIX as the single ATLAS readout system will culminate the collaboration-wide efforts aimed at harmonizing the detector readout.

FELIX consists of commodity servers equipped with custom PCIe FELIX I/O cards. The main components of the FELIX cards are: the fiber optic transceivers, the PCI Express bus, and an FPGA. Custom firmware handles data transmissions between the detector front-ends and the PCIe Interface. The optical links connected to the detector front-ends are bi-directional, carrying readout data upstream (detector to FELIX), and configuration, trigger decisions and other LHC control signals, such as bunch counter resets (BCR), downstream (FELIX to detector). FELIX also interfaces with the Timing, Trigger and Control (TTC) system that distributes trigger, clock and LHC information to FELIX with fixed deterministic latency. Both the TTC and the downstream and upstream detector links (downlink and uplink) use custom fully synchronous data transmission protocols. Once in the FELIX servers, the data is routed to network switches interfacing with another set of commercial servers, where a software Readout Driver (SWROD) processes data and performs event building and aggregation. The routing of the large amounts of data between the FELIX card and the network switches is facilitated by the FELIX software.

FELIX takes advantage of the significant developments in switched networks, server CPUs and FPGAs over the last decade. Tasks that previously required custom hardware can now be moved into FPGAs and CPUs that can handle more flexible configurations, as well as higher data throughput. Furthermore, the move to commercial components not only simplifies the overall system design, particularly through the use of network switches, but also enables easier maintenance and upgrades.

* ricardo.luz@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBG004

# THE KARABO MIDDLELAYER API AND MOTION CONTROL SYSTEMS AT EUROPEAN XFEL

D. Hickin*, N. Anakkappalla, V. Bondar, R. Costa, W. Ehsan, S. Esenov,
R. Fabbri, G. Flucke, A. García-Tabarés Valdivieso, G. Giovanetti
D. Goeries, S. Hauf, I. Karpics, A. Klimovskaia, A. Nasri, A. Parenti,
A. Samadli, H. Santos, M. Sharon, F. Sohn, M. Teichmann, J. L. Vazquez-Garcia,
European XFEL, Schenefeld, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Karabo is a device-based distributed control system toolkit used to implement the control and data acquisition systems of European XFEL. A feature of Karabo is its Middlelayer (MDL) API - a powerful, flexible and easy-to-use asynchronous Python API which can be used to implement Karabo devices. Such devices may interface hardware directly, or they may communicate with other Karabo devices to provide coordination between them or derive functionality from them. Lightweight middlelayer devices, called macros, allow routines for coordinating or monitoring Karabo devices to be quickly created. A command-line interface using the MDL API is provided by the iKarabo utility. In this contribution, an overview of the Karabo Middlelayer API is given, presenting key parts of the API, and macros and iKarabo are introduced. Examples of MDL devices are presented, with an emphasis on motion systems. A framework for multi-axis motion, Virtual Motor Base, and the scan tool, Karabacon, are described. The Karabo MDL API is compared to EPICS solutions, including pythonSoftIOC and Asyn, and to Bluesky.

## INTRODUCTION

Karabo [1] is an open source SCADA (supervisory control and data acquisition) framework which can be used in building distributed control systems, particularly those targeted at the operation of large science experiments.

It was developed at the European X-ray Free Electron Laser facility (European XFEL), a free-electron laser (FEL) user facility providing soft and hard X-ray FEL radiation, currently to seven scientific instruments [2].

At European XFEL, Karabo is the main SCADA framework used in the implementation of the control system for the three photon tunnels (SASE1, SASE2, SASE3) and the seven instruments (see Fig. 1). For the accelerator control system, including control of the undulators, the main choice of framework is DOOCS [3], with some use of TINE [5] and EPICS [4]. (DOOCS and EPICS are also used in some parts of the photon tunnels and instruments.)

Karabo is device-based - monitoring and control of a SCADA system implemented using Karabo is achieved using a system of named distributed objects, called devices, which can communicate with each other.

A device may provide direct hardware control - for example it may represent a single piece of equipment, such as a motor, pump, valve, camera or digitizer or part of a piece of equipment, such as a digitizer channel, or be used to control a set of equipment. Other devices may provide software routines such as data analysis. Some devices may derive functionality from other devices or provide coordination between devices ("middlelayer"). Or, a device may provide an interface to a service. Some devices provide system services, such as devices for data logging, GUI server devices, or devices to interface project databases used for device configuration.

Communication between devices is routed through a message broker, and devices are grouped into broker topics. At European XFEL, separate topics exist for each tunnel and for each instrument, as well as three separate topics for lasers, one for the instruments of each of the three tunnels.

Each device has an identifier, unique within its topic, its *device ID*. Devices run within a software process called a *device server*. Multiple devices can run in a single device server. Each device is an instance of a *device class*. For devices providing equipment control, a device class may correspond to a particular hardware type, for example, a hardware model, or related hardware, such as a set of cameras with a common interface specification or controlled via the same vendor library. Device classes are provided via plugins called device packages. Typically, a device package provides one device class or a collection of related device classes.

### *Device Schemas*

Devices have properties and slots. *Properties* are essentially key-value pairs that represent named values within a device. Examples include a position of a motor, the current of a power supply or the pressure of a pressure gauge. Each property has a key name and a value. The value of a property may be of any of the Karabo data types. These include scalar types (Boolean, signed and unsigned integer types, floating point types, strings), vector types such as vectors of scalar types, N-dimensional arrays and tables thereof. Properties may receive further specification using attributes (for example, maximum and minimum values for a numerical property or the properties units).

Properties may be configurable after device instantiation, read-only or configurable only at instantiation. Properties may have an access level defined. Users without sufficient

* david.hickin@xfel.eu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBR001

# THE ELT PRIMARY MIRROR FAULT DETECTION, ISOLATION, AND RECOVERY SOFTWARE

L. Andolfato†, J. Argomedo, M. Dimmler, T. Grudzien, M. Pasquato,
E. Pozna, A. Reinacher, B. Sedghi, P. Sivera
European Southern Observatory, Garching bei Muenchen, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The 39 m diameter primary mirror of the ESO Extremely Large Telescope, currently under construction at Cerro Armazones in Chile, is composed of 798 hexagonal segments. Each segment is equipped with several sensors and actuators to measure and adjust its position. The control algorithm uses 24000 I/O points, distributed over 1195 square meters, to dynamically maintain the alignment and the shape of the mirror. The reliability of this large number of devices is improved by the introduction of a failure detection, isolation, and recovery strategy. The main goal of the FDIR strategy is to enable the control loop to continue working even in case of failures by identifying and masking faulty devices in real-time and avoid failure propagation. This paper provides an overview of the system, summarizes the main availability requirements, and illustrates how the FDIR software has been designed, implemented, and tested. The detection and isolation of failing edge sensor devices, responsible for measuring the difference in piston, shear, and gap between adjacent segments, is taken as running example. Advantages and limitations of the presented design are summarized in the conclusions.

## INTRODUCTION

The Extremely Large Telescope (ELT) is one of the largest optical/near-infrared telescopes being built by the European Southern Observatory on top of Cerro Armazones in the Atacama Desert of northern Chile. The light is collected by the primary mirror (M1) which has a diameter of 39 m, and it is made of 798 quasi-hexagonal mirror segments of approximately 1.45 m in size (Fig. 1).

Each segment is equipped with six Edge Sensors (ES), three Position Actuators (PACT) and nine Warping Harness (WH). The ES can measure with nm accuracy the out-of-plane (piston) and the in-plane translation displacements (gap and shear) between two adjacent segments. The PACT, with a stroke of 10 mm and a nm accuracy, are used to dynamically position the segment in piston, tip and tilt to keep the segments aligned under variable load conditions and disturbances. The WH can change the shape of the segment by applying different axial support forces to the mirror segment (Fig. 2).

The M1 mirror is organized in six sectors, and each sector contains 22 flowers made of up to seven segments each. All segments within a flower are connected to a Segment Concentrator (SegC) cabinet that hosts one control module per segment per type of device for a total of 21 control modules. The control modules provide the electronics to acquire the measurements from the sensors and to apply the target positions to the actuators. In addition, the SegC cabinets provide services such as power distribution, temperature monitoring, liquid cooling, and connectivity to the communication infrastructure via a network switch.

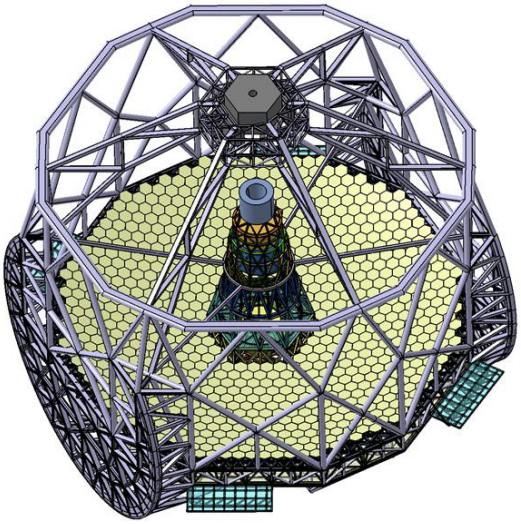


Figure 1: The 798 segments of the M1 primary mirror.

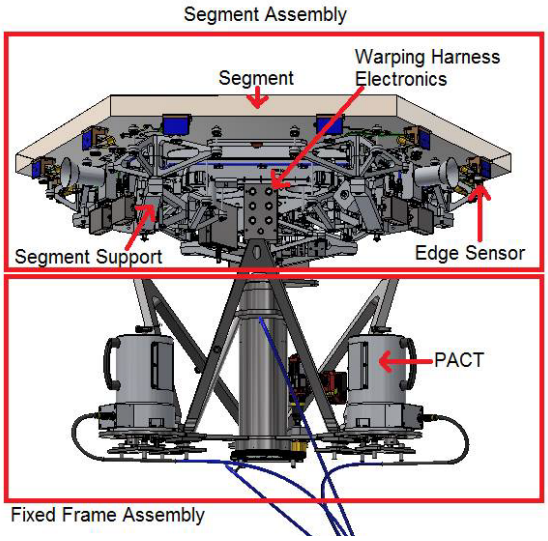



Figure 2: M1 Segment Assembly with ES, PACT, WH.

The 132 SegC cabinets are controlled by the M1 control software (M1 CS) that consists of two layers: the Local Control System (LCS) and the Local Supervisor (LSV). The main task of M1 CS is to obtain and maintain, for the duration of the observation, a given mirror optical quality.

† landolfa@eso.org

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBR002

# EXPERIMENTAL SUPERVISION TOOLS SOFTWARE UPGRADE ON THE LASER MEGAJOULE FACILITY

N. Roux*, J. Airiau, M. Aubin, S. Palmier, F. Tournemenne
Centre d'Études Scientifiques et Techniques d'Aquitaine, Le Barp, France

*Abstract*

The Laser MegaJoule (LMJ), a 176-beam laser facility developed by CEA, is located near Bordeaux. It is part of the french *Simulation* Program, which combines improvement of theoretical models used in various domains of physics and high performance numerical simulation. It is designed to deliver about 1.4 MJ of energy on targets, for high-energy density physics experiments, including fusion experiments.

The LMJ technological choices were validated on the Ligne d'Intégration Laser (LIL), a scale-1 prototype composed of 1 bundle of 4-beams. The first bundle of 8 beams was commissioned in October 2014 with the first experiment at the LMJ facility. Operational capabilities increase gradually every year until the full completion by 2026. By the end of 2025, 22 bundles of 8-beams will be assembled, and 19 bundles are expected to be fully operational.

We present a successful modernization strategy of operational tools for the LMJ, enhancing usability and maintainability. Our efforts focused on replacing legacy systems by migrating from a monolithic client application to a Django and React web app, providing a more accessible and maintainable interface.

Additionally, we streamlined the XML (Extensible Markup Language) tree-based configuration of calculation scenarios by migrating toward a robust PostgreSQL database.

Furthermore, we overhauled our deployment process using Ansible, increasing operational efficiency and consistency.

## INTRODUCTION

The Laser MegaJoule (LMJ) is one of the largest scientific facilities in Europe, dedicated to high-energy density physics experiments and inertial confinement fusion. The scale and complexity of the facility require sophisticated supervision software for configuration, monitoring, and operatsion of its laser chains.

The LMJ is primarily designed to perform indirect-drive inertial confinement fusion (ICF) experiments on deuterium–tritium targets. Beyond this major objective, the facility also enables a wide spectrum of experiments, ranging from high energy density physics to fundamental science [1]. As a cornerstone of the French *Simulation Program*, it brings together three complementary components: the continuous improvement of physics models, large-scale high performance numerical simulations, and experimental validation. Together, these elements ensure the reliability and safety of the French deterrence strategy.

* nicolas.roux@cea.fr

The facility is dimensioned to accommodate 176 laser beams, organized into 22 bundles of 8 beams each. These bundles are distributed across four laser bays, arranged symmetrically around a central target bay measuring approximately 60 meters in diameter and 40 meters in height. The target chamber and its associated experimental equipment are installed at the center of this bay (see Fig. 1).

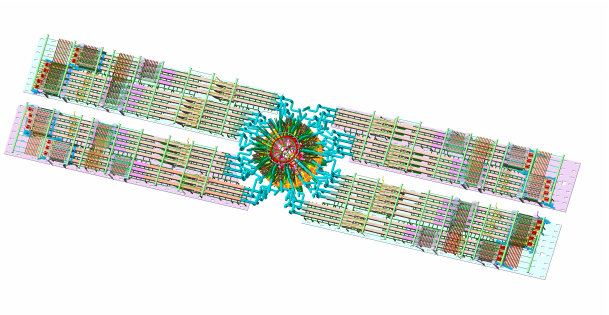


Figure 1: Laser bays and target chamber.

Each laser bundle is composed of four main subassemblies (see Fig. 2):

- the front-end, responsible for generating the initial laser pulse,
- the amplification section, located in the laser bay, which boosts the energy of the beam,
- the transport system, which guides the beam toward the target chamber,
- the frequency conversion and focusing system, which adapts the laser wavelength and concentrates the energy on the target.

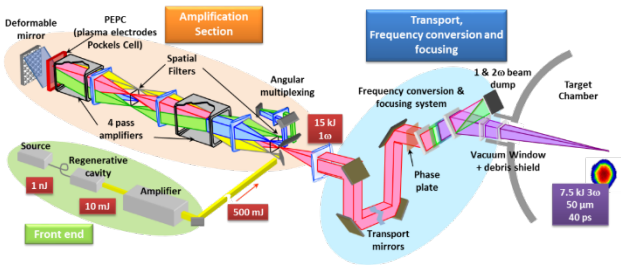



Figure 2: LMJ laser bundle.

Operating such a facility requires precise coordination of thousands of parameters that define the laser chain configuration, the timing of operations, and the diagnostic setup

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBR003

# EVALUATING FUNCTION-AS-A-SERVICE (FAAS) FRAMEWORKS FOR THE ACCELERATOR CONTROL SYSTEM

A. Jaikar*, J. Diamond, A. Tiradani, B. Harrison
Fermi National Accelerator Laboratory, Batavia, USA

*Abstract*

The Fermilab Accelerator Complex is the largest national user facility in the Office of High Energy Physics (DOE/HEP) program which uses a linear accelerator (Linac), Booster rapid cycling synchrotron, Recycler storage ring, and Main Injector to produce proton beams. This particle accelerator produces high-volume, high-velocity telemetry and controls data that must be ingested, processed, and retained for operations and analysis. As particle accelerator control systems evolve in complexity and scale, the need for a responsive, scalable, and cost-effective computational infrastructure becomes increasingly critical. Function-as-a-Service (FaaS) offers an alternative to traditional monolithic architecture by enabling event-driven execution, automatic scaling, and fine-grained resource utilization. This paper explores the applicability of FaaS frameworks in the context of a modern particle accelerator control system, with the objective of evaluating their suitability for short-lived and triggered workloads.

## INTRODUCTION

The particle accelerator complex consists of thousands of devices like magnets, RF cavity, vacuum, power supply, and many other devices, which are coordinately controlled by the control system. This accelerator control system has been designed, developed, and maintained for around 3 decades. When the system was designed, there were many limitations, such as network latency, RAM capacity, CPU frequency, and network protocol as well. The system was designed with the ACNET (Accelerator Network) protocol [1], which is based on the User Datagram Protocol (UDP). Although it has performed very well over the years, modern technology landscapes demand higher efficiency, flexibility, and scalability, which the existing framework struggles to provide. This framework has been designed without the consideration of large storage, processing, and machine learning-based operations. Therefore, modernization is needed to support future operations of the complex.

In the accelerator complex, data is captured either at regular intervals or in response to specific events. Periodic data can be processed at regular intervals and event-based data can be processed at regular intervals or at the same time. To support both periodic and event-driven processing, the Function-as-a-Service (FaaS) has the ability to process the data with a scalable computing infrastructure. Kubernetes hosted FaaS platform helps the developer concentrate on developing a function which will process the data without managing the required infrastructure.

* amol@fnal.gov

## BACKGROUND

The particle accelerator produces heterogeneous high-volume data streams and operates under strict latency, reliability, and safety constraints. Therefore, data generated by these properties have the following characteristics:

- **High cardinality**: Slightly more than half million properties
- **Mixed frequency**: Different frequencies for different devices and their properties
- **Low latency**: Some application requires low latency data
- **Long retention**: months to years for trend analysis, reliability studies on historical data

Once the raw data has been ingested, it undergoes processing steps such as filtering, tagging, unit conversion, and other transformations. The FaaS platform offers event-based data processing that can be used for accelerator data processing. It also depends on Kubernetes for container orchestration, automatically handling scaling, networking, and resource allocation.

In this paper, we evaluate three Kubernetes based FaaS platforms; fission [2], Knative Serving/Eventing [3], and nuclio [4] with respect to their architecture, characteristics, and fit for the accelerator operations.

## ARCHITECTURAL OVERVIEWS

This section covers the architecture for all three FaaS frameworks.

### *fission*

Fission is an open-source Function-as-a-Service (FaaS) framework built on Kubernetes, designed to provide a lightweight, event-driven execution environment. Instead of traditional server-based deployments, fission functions are short-lived, stateless, and triggered by events, making the framework highly suitable for cloud native microservices, event processing, and real-time data pipelines. Fission consists of multiple components that build up the architecture.

**Controller** The controller is the component which accepts requests from clients. All Custom Resource Definitions (CRDs) of fission, which are defined in Kubernetes, are used for Create, Read, Update, and Delete (CRUD) operations. A CRD is a powerful feature in Kubernetes that allows you to extend the Kubernetes API and define your own custom resource types. The controller uses CRUD APIs for functions, triggers, and environments. A cluster-wide admin permission is required for the controller to access all fission CRDs from all namespaces.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBR004

# EBPFCAT — AN OPEN SOURCE ETHERCAT IMPLEMENTATION FOR EXPERIMENTAL CONTROL

Martin Teichmann*, European XFEL, Schenefeld, Germany

*Abstract*

EtherCAT has become a popular communication protocol for controlling equipment in large facilities. Although it is an open standard that can be used on common hardware, commercial solutions on specialized hardware are usually employed. Here we introduce EBPFCat, a fully open-source implementation of the EtherCAT protocol that can be used on any modern Linux system without the need to write kernel modules. In EBPFCat, non-time-critical parts are implemented in Python, while real-time performance and reliability are ensured through the use of EBPF, a virtual machine embedded within the Linux kernel. EBPFCat contains a code generator that allows EBPF programs to be generated directly from Python without the need of separate compilation tools. These programs are verified for correctness by the Linux kernel and executed therein, allowing for feedback bandwidths exceeding 10 kHz. The EBPF code generator is a generic solution that can also be used for unrelated EBPF programs. EBPFCat is a standalone project that operates independently of any specific high-level control system. Nevertheless, it has already been integrated into Karabo, the control system used at the European XFEL. We show applications for a large variety of control problems, from simple vacuum systems to intricate feedback loops for motion control.

## INTRODUCTION

The computer control of hardware traditionally followed many standards, as the constraints for different use cases were very different, whether it be high speed, low latency, high reliability or low cost. These considerations become less important with the advent of modern electronics, as there are bus systems which fulfill all constraints at the same time. However, a major disadvantage of those systems is often a loss of flexibility, since these bus systems are usually bound to vendor specific hard- and software. They usually also require special operating systems to fulfill hard real time conditions, which makes its usage more cumbersome, especially if those real time conditions are not actually needed.

EtherCAT [1] is such a bus system. As it is publicly documented, there are many software drivers available, some of which are even open source and free of charge [2]. They are, however, still cumbersome to program, as they run on real-time operating systems, which usually means that code has to run in the operating system's kernel, requiring to write kernel extension modules, which is usually very error-prone and risks crashing the entire system.

This is where EBPF [3] comes into play: this Linux kernel technology, originally called the Extended Berkeley Packet Filter, has nowadays grown into a fully-fledged virtual machine inside the Linux kernel, that allows user space programs to load code into kernel space, which runs, thanks to just-in-time compilation, at full processor speed.

EBPFCat [4] binds all these technologies together. Fully written in Python, it allows for computer control applications where real time conditions are not needed, to be written in Python. For real-time applications, an EBPF generation library is included, that allows for the programming of in-kernel code within the same Python code, with a user experience as close to normal Python programming as EBPF allows.

EBPFCat can already be used from within the Karabo SCADA system, but at it is a standalone library, other SCADA systems may be used as well.

## ETHERCAT

EtherCAT is an application of the EtherNET protocol. A regular computer with an EtherNET card, called the *MainDevice* in EtherCAT jargon, controls other devices which are daisy-chained on the EtherNET port. These other devices are usually called *SubDevices*, but as they mostly come in shape of input/output terminals, we just call them *Terminal*s in EBPFCat and throughout this paper. The Main Device sends specially formatted EtherNET packets to the terminals, which hand them from one to the next, while finally sending it back to the origin. The size of said packets is not changed during this game of telephone, yet the content may be modified by a terminal for it to send data to the Main Device.

When EBPFCat is started, it will first send inquiry messages via the EtherNET port, which allows it to gather a good picture of the connected terminals, as most terminals contain full self descriptions.

The whole point of a bus system is to control several devices at the same time. However, these devices are often not just the terminals on the EtherCAT bus. Especially if the device to be controlled does not (yet) have a direct EtherCAT control, one usually uses generic EtherCAT terminals like digital or analog in- and outputs, which are connected to the target device.

## BOOTSTRAPPING THE BUS

Upon starting EBPFCat, the bus is scanned for all connected terminals, and the self-description relevant for running the bus is read out. This includes the serial number of the terminals, which is stored in each terminal's EEPROM, which is used to identify the individual terminal. This procedure is different from other implementations of the EtherCAT bus, which often assume the correctness of a user-

* martin.teichmann@xfel.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBR005

# UPDATES ON THE KARABO CONTROL SYSTEM

S. Hauf*, N. Anakkapalla, V. Bondar, R. Costa, W. Ehsan, S. Esenov, R. Fabbri, G. Flucke, A. García-Tabarés Valdivieso, G. Giovanetti, D. Goeries, D. Hickin, K. Jahangir, I. Karpičs, A. Klimovskaia, A. Muennich, A. Nasri, A. Parenti, A. Samadli, H. Santos, D. Sequeira, M. Sharon, A. Silenzi, F. Sohn, J. L. Vazquez-Garcia, European XFEL, Schenefeld, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Karabo is a SCADA system developed at the European XFEL to facilitate operations in a complex and flexible control environment, and to ingest and process data from diverse sources, including MHz-rate detectors, frequently with tight time correlation requirements. In this contribution an overview of the current state of the Karabo ecosystem is given and recent developments are presented. These include the introduction of authentication and authorization, the transition from *Boost.Python* to *pybind11* and from Boost to the C++ standard library, and additional modernization of the codebase. These significant changes are reflected in a new major version, Karabo 3, which replaces the 10-year-old Karabo 2 branch.

## INTRODUCTION

Karabo [1] is a SCADA system designed to meet the specific requirements of the European X-ray Free Electron Laser (EuXFEL) [2]. These requirements include the need for high-data-rate acquisition and online processing [3], capable of handling up to 15 GB/s, and in-built concepts to accommodate flexible installations, which may undergo significant changes on a weekly basis due to the diverse experimental techniques and user communities supported by the facility.

Karabo was conceived in 2011 following a survey that revealed that the then dominant scientific control systems, EPICS [4], and TANGO [5], could only cope with these requirements with substantial custom development. Karabo 1.X [6] was capable of supporting the component-level commissioning of the evolving facility. In 2016, Karabo 2 [7] was released, improving on core concepts such as a substantial simplification of the C++ event loop, the introduction of a project database, and a complete refactoring of the main user interface, the Karabo GUI, towards type safety using the Python *traits* package [8]. This system was designed to support production-scale usage of an event-driven, asynchronous distributed SCADA system at a large scientific facility. It carried EuXFEL within its DNA, exemplified by timing identifiers that mirror EuXFEL jargon, and was commissioned alongside the facility's transition into user operation. The facility's requirements are exemplified through the in-built point-to-point interfaces for large data rates and a fully self-describing system topology, which can accommodate large-scale reconfiguration of the system at runtime, without the need for a centralized database or coordinating entity.

* steffen.hauf@xfel.eu

By 2020, the facility and the SCADA system had matured to the point where user operation could be sustained without constant attention to system stability. During the next five years, Karabo 2 became the primary workhorse for operating the scientific instruments and photon systems of the facility.

A fundamental principle within the EuXFEL Control software development team is to prevent the accumulation of substantial technical debt. Consequently, even during periods of stable operation, the software has been consistently revitalized, for instance, by migrating its core communication backend from a JMS broker [9] to AMQP/RabbitMQ [10,11] and by incorporating novel features, such as the logging of all broker-exchanged data updates in an InfluxDB timeseries database [12, 13].

To date, Karabo, an open-source software platform, has been exclusively developed at the EuXFEL facility. However, recent initiatives have been undertaken to streamline the installation process beyond the facility's boundaries. These efforts are further elaborated upon towards the conclusion of this article.

### *Programming APIs*

Karabo supports three programming APIs: C++ and two Python APIs. The Python Bound API is implemented by binding the C++ code and is intended for high-performance computing tasks, requiring high-speed, low latency I/O as part of near-realtime data processing pipelines. EuXFEL's online calibration system [14], which can process images at a throughput of 1 Megapixel at >5000 Hz or 5 Gpixel/s, is implemented in this API. Conversely, the EuXFEL Data Acquisition System (DAQ) is implemented in the C++ API.

The second Python API, Karathon, is a middle-layer API that does not share significant code with the other APIs. It is implemented fully in Python, does not aim to be feature-complete with respect to these APIs (and vice versa), and strongly emphasizes asynchronous and pythonic concepts. It also avoids domain-specific language extensions and aims to be as concise as possible. These aspects frequently make it the API with the shortest time to production, and especially when raw computational performance is not required, the preferred choice for many tasks (see Fig. 1).

In the past two years, Karathon has also demonstrated that it is particularly well-suited for AI-assisted coding, since in essence it is *just another Python package*. While hardware is frequently directly integrated through Karathon, the API was initially conceived and still excels in multi-device coordination within the distributed system, and procedural tasks such as scanning. It thus covers significant aspects of

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THBR006

# TOWARDS ASYNCHRONOUS CONTROL SYSTEMS, AN ASYNCIO IMPLEMENTATION OF OPC UA USING TANGO GREEN MODES

E.J. Morales*, X. Mercadal, J.A. Ramos, Z. Reszela, A. Rubio, S. Rubio-Manrique†
N. Serra, J. Villanueva, ALBA Synchrotron, Cerdanyola del Vallès, Spain
G. Cuni, Paul Scherrer Institute, Villigen, Switzerland

*Abstract*

The ALBA Synchrotron (Barcelona, Spain) has been operating as a 3 GeV facility for over 10 years and is now preparing its transition to ALBA II, a fourth-generation light source. As part of this planned upgrade, we are evaluating state-of-the-art technologies that could shape the future of our Tango Control System. In particular, we investigate how asynchronous programming can enhance system responsiveness while reducing latency and resource usage. This study focuses on applying asynchronous communication paradigms at all levels between our Taurus SCADA UIs, Tango Control System and PLC-based systems — used for Equipment (EPS) and Personnel (PSS) Protection as well as automation. In this context, we explore the adoption of OPC Unified Architecture (OPC UA), a self-descriptive industrial standard for secure, platform-independent communication, alongside asyncio, the Python standard library for coroutine-based asynchronous programming, as supported by the FreeOpcUa library and "green" modes of PyTango, the Python binding for Tango Controls. Our goal is to demonstrate a modern, flexible, vendor-independent and high-performance control strategy for ALBA II Control System. We provide a comprehensive comparison and benchmark between the proposed solution and existing PyPLC Tango Device Servers.

## INTRODUCTION

The ALBA synchrotron (Barcelona, Spain) [1], a third-generation electron accelerator, is preparing for its transition to a fourth-generation light source. A scheduled shutdown in 2030 will enable the upgrade and restart of the facility. Over the past decade, ALBA has relied on its original design technologies, demonstrating robustness through more than ten years of continuous service. The forthcoming upgrade requires modernizing supporting technologies and adapting the infrastructure to new operational demands.

### *ALBA Equipment Protection System*

The Equipment Protection System (EPS) [2] safeguards the accelerator and beamline subsystems through a network of PLCs [3] interconnected via `Ethernet PowerLink` [4]. It monitors PLC variables and can interlock machine operation in under 50 ms. PLCs are integrated into the ALBA control system using Tango [5], a distributed control framework developed by the Tango Controls Collaboration [6, 7]. Each EPS PLC manages over 400 nodes and 1500 parameters, grouped into analog, floating-point, digital, and actuator types. The `EPS_LIBRARY` [5] standardizes access using 16-bit `StatusRead` and `StatusWrite` words for diagnostics and control, while additional words report CPU and network health.

Communication currently relies on `Modbus TCP` [8, 9], a simple and robust protocol widely used in industry. While effective for small to medium systems, its scalability in large or dynamic environments such as ALBA's EPS is limited, motivating the evaluation of newer approaches. All EPS configuration is centralized in a CSV [10] file mapping nodes to `Modbus` addresses and parameters.

Integration with Tango is handled by the `AlbaPLC` class, which parses the CSV at startup to generate Tango attributes [11]. The lower communication layer `PyPLC` manages `Modbus` transactions, while `AlbaPLC` exposes EPS states, parameters, and events. This modular design separates hardware-level management (`EPS_LIBRARY`) from control-system integration (`AlbaPLC`), enabling reliable and scalable operation.

### *OPC Unified Architecture*

OPC UA (IEC 62541) [12], developed by the OPC Foundation [13], offers a modern, cross-platform, secure, transport agnostic and event-driven communication framework. It supports client–server and publisher-subscriber models, advanced information modeling, and hierarchical data representation. Being self-descriptive, OPC UA can export PLC internal structures directly, eliminating intermediate files and avoiding transcription errors. In this work, we rely on `asyncua` [14], an open-source Python library providing asynchronous OPC UA clients and servers.

### *Asynchronous Programming and Tango Green Modes*

Asynchronous programming puts its focus on allowing multiple operations to progress concurrently without blocking, a paradigm suited for large event-driven systems. `asyncio` [15] is Python's standard library for asynchronous programming. Based on coroutines and event loops, it allows to manage tasks and asynchronous events concurrently while keeping resource usage low. The Python binding of Tango Control System, `PyTango` [16], and its `asyncio` Green Mode have been selected for scalable data handling at ALBA. This approach is focused on improving responsiveness, reducing CPU load, and replacing polling-based communications with data-change event subscriptions.

This article presents the evaluation and benchmarking of a modernization path for EPS communication, comparing the

* emorales@cells.es
† srubio@cells.es

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG001

# FAST ORBIT FEEDBACK USING THE GSVD FOR SYSTEMS WITH MULTIPLE SLOW CORRECTOR ARRAYS *

H. Kim, H. Studholme, I. Kempf†, University of Oxford, Oxford, UK
W. Heath, Bangor University, Bangor, Wales
S. Banerjee, L. Bobb, Diamond Light Source, Oxfordshire, UK

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Advances in detector speed and resolution at 4th generation light sources make electron beam stability a critical requirement. At Diamond-II, the fast orbit feedback (FOFB) will stabilise the beam using 252 beam position monitors and two actuator arrays of 252 slow and 138 fast correctors. We previously proposed an approach based on the generalised singular value decomposition (GSVD), a two-matrix factorisation, to decouple the system into two-input modes controlled by both slow and fast correctors, and single-input modes controlled by slow correctors alone. This approach assumed identical dynamics for all correctors within each array. However, recent developments have shown that variations in vessel geometries and cooling channels introduce significant differences in corrector bandwidths, particularly for the slow correctors in the horizontal plane. Specifically, Diamond-II will have at least three distinct types of slow correctors in the horizontal plane. To address this, we extend the GSVD-based approach with balancing input filters to incorporate multiple slow corrector arrays with different dynamics. We also introduce a new regularisation matrix that preserves controller properties between the original and mode spaces for any choice of regularisation parameter. We analyse the resulting control system and present simulation results from preliminary Diamond-II data, demonstrating that the control specifications are met.

## INTRODUCTION

At Diamond Light Source (Diamond), the fast orbit feedback (FOFB) attenuates disturbances of the electron beam in the storage ring. These disturbances include ground and girder vibrations, power supply noise, and RF noise. At Diamond, the FOFB uses 173 electron beam position monitors (BPMs) measuring horizontal and vertical beam position, and 172 vertical and 172 horizontal corrector magnets (CMs) to correct the electron beam position at 10 kHz. The FOFB suppresses root-mean square (RMS) beam deviations to 10 % of the beam size for frequencies up to a closed-loop bandwidth of 140 Hz, corresponding to the frequency at which disturbances are attenuated by −3 dB [1]. For Diamond-II, advances in detector speed and resolution will require increasing the closed-loop bandwidth to ≥ 1 kHz and reducing the RMS beam deviation to 3 % [2].

* This work was supported in part by funding from the Engineering and Physical Sciences Research Council Impact Acceleration Account (Grant number EP/X525777/1) and in part by the Engineering Undergraduate Research Opportunities Programme (EUROP) scheme of the Department of Engineering Science, University of Oxford, Oxford, UK.

† Corresponding author: idris.kempf@eng.ox.ac.uk

Table 1: Slow and Fast Corrector Magnet (SCM and FCM) Types at Diamond-II.

| Type | Number | BW (H) | BW (V) |
|---|---|---|---|
| SCM1 | 12 | 52 Hz | 177 Hz |
| SCM2 | 54 | 93 Hz | 167 Hz |
| SCM3 | 186 | 118 Hz | 187 Hz |
| FCM1 | 132 | 9331 Hz | 9331 Hz |
| FCM2 | 6 | 1895 Hz | 9331 Hz |

H: horizontal, V: vertical, BW: bandwidth

To meet these requirements at Diamond-II, the sampling frequency will be increased to 100 kHz and the time delay will be reduced to less than 100 µs. The FOFB will use $n_y = 252$ BPMs and two types of CMs: an array of $n_s = 252$ slow CMs (SCMs) and an array of $n_f = 138$ fast CMs (FCMs). The SCMs will attenuate DC and lower-frequency disturbances up to 100 Hz, such as ground vibrations and drifts, whereas the FCMs will attenuate disturbances from 100 Hz to 1000 Hz, which will include disturbances from 300 Hz to 500 Hz due to RF noise and water cooling from 200 Hz to 400 Hz. While in principle the two-array system allows higher-frequency disturbances to be attenuated, it prohibits the current FOFB control algorithm from being reused. This control algorithm decouples the electron beam dynamics using a singular value decomposition (SVD) of the orbit response matrix [1], an approach that fails with more than one array of CMs. However, if the two-array system is not decoupled, the SCMs and FCMs can interfere and lead to disturbance amplification instead of attenuation [3], even when the targeted bandwidths do not significantly overlap [4].

To accommodate the two CM types, we previously extended the SVD-based approach to the generalised SVD (GSVD), a two-matrix factorisation technique [5]. The GSVD-based approach decouples the system into $n_f$ modes, controlled by both SCMs and FCMs with a closed-loop bandwidth up to 1 kHz, and $n_s - n_f$ modes, controlled only by SCMs with a closed-loop bandwidth up to 100 Hz [4]. For the $n_f$ fast modes, the control effort is distributed between SCMs and FCMs using mid-ranging control [7]. This approach is successfully being tested on the existing Diamond storage ring [8], but it assumes identical bandwidths within all SCMs and within all FCMs. However, preliminary simulations for Diamond-II showed that vessel geometries and cooling channels will introduce significant differences in both SCM and FCM bandwidths [6], with one possible scenario summarised in Table 1. In the horizontal plane,

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG003

# DESIGN OF ROBUST CONTROLLER BASED ON LOOP-SHAPING FOR QUICK-SCANNING MONOCHROMATOR

Zongyang Yue*, XiaoBao Deng, Lirong Zheng,
Hao Liang, Yuanshu Lu, Ge Lei†, Gang Li‡
Institute of High Energy Physics, Chinese Academy of Sciences, Beijing, China

*Abstract*

Quick-scanning Double Crystal Monochromator (QDCM) is core equipment of the Quick-scanning X-ray Absorption Fine Structure (QXAFS) experimental method. To realize the sub-second temporal resolution of QXAFS method, QDCM is responsible for achieving a reproducible oscillatory movement of monochromator crystals with a given frequency and amplitude around Bragg angle. A new robust controller design method is proposed for QDCM system. This controller, based on the loop-shaping principle in frequency-domain, considers the effect of nonlinearities and uncertainties in low frequency region and the suppression of vibration mode in high frequency region. The results are shown to demonstrate the feasibility of this method for the QDCM fast-scan seeking.

## INTRODUCTION

High Energy Photon Source(HEPS) [1] X-ray absorption spectroscopy beamline(ID46) is a high-performance hard X-ray beam line optimized for X-ray absorption spectroscopy and derivative experimental methods. The basic aim of the this beamline is to build a high universality, high-brightness XAFS experimental system with high versatility and wide energy range to provide the users with a reliable, stable and easy-operated platform for X-ray absorption spectroscopy experiment. The layout of the ID46 beamline is shown schematically in Fig. 1.

Quick-scanning extended X-ray absorption fine structure (QEXAFS) is one of the most significant applications of HEPS, as a powerful structural analysis tool for time-resolved in situ investigations of materials on the atomic scale during physical and chemical reaction processes. The QEXAFS technique uses a quick-sacnning X-ray monochromator to enable rapid and repetitive energy scans by oscillating the monochromator crystals, which requires the Bragg axis of the fast-scanning monochromator to be capable of tracking periodic signals with fixed angles at frequencies exceeding 25 Hz. Although the Bragg axis driven by traditional stepper motors can satisfy stability under static conditions, it cannot achieve high-dynamic scanning functions. In contrast, the driving method based on permanent magnet synchronous motors(PMSM) can achieve both steady-state accuracy and high-dynamic tracking performance under closed-loop control. For Bragg axis close-loop control, the traditional PID control cannot suppress high-frequency resonant modes and realize the error-free tracking control for periodic signals. Therefore, these demanding requirements pose a major challenge to optimize of the Bragg axis closed-loop structure.

Accordingly, advanced control algorithms for DCM have been proposed to improve the system control performance. Reference [2] propose the FxLMS-PID hybrid algorithms to supress the vibration and improve the stability for double crystal monochromator. In [3], predictive modeling through dynamic error budgeting was applied in the HD-DCM which improve the closed-loop performance. The multiple loop control algorithm was applied in [4], which realize the fly-scan real-time synchronization bewteen insertion device and monochromator. Despite these control algorithms mentioned above, most of existing methods didn't fully address the robustness issue for the Bragg axis of DCM in the presence of crystal loads and external disturbances, and cannot satisfy the tracking control requirement of DCM for high-frequency signals. In this paper, the dynamic characteristics of the bragg axis of QDCM were studied, and system identification was conducted using the DFT algorithm to establish the dynamic model of the system. A 2-DOF control algorithm structure was be adopted, including acceleration feedforward and mixed sensitivity feedback controller [7] based on H-infinity. In the feedback controller design, this paper based on the system's dynamic equations obtained from experiment results, design weighting functions to weight the system's sensitivity function and complementary sensitivity function. The proposed control algorithm can effectively trade off the close-loop bandwidth and system uncertainty.

This paper is organized as follows: In section II, the dynamic model of Bragg axis of QDCM is discussed. In Section III, the design method of feedback controller is proposed. The results of fast-scan seeking of QDCM are shown in Section IV. Finally, the conclusion is given in Section V.

## THE DYNAMIC MODEL OF BRAGG AXIS OF QDCM

The mechanical model of Bragg axis [5] is depicted in Fig. 2. The rotor of the Bragg axis is held by two ball bearings. The crystal assembly is attached at the head of the spindle. An encoder(angular sensor) is placed at the end of spindle for the measurement of the rotation angle, and can be used for positioning control.

In order to ensure the required performance, it is crucial to establish the linear dynamic model of the system. Due to the rigid connection between crystal assembly and the spindle of Bragg axis, we can consider Bragg axis motor

* yuezy@ihep.ac.cn
† leige@ihep.ac.cn
‡ li75gang@ihep.ac.cn

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG004

# ENERGY-OPTIMAL PREDICTIVE CONTROLLER FOR CHILLER-BASED COOLING PLANTS OF ACCELERATOR FACILITIES

D. Monteiro*, N. Bunijevac, CERN, Geneva, Switzerland
A. La Bella, L. B. de Giuli, L. Nigro, Politecnico di Milano, Milan, Italy

*Abstract*

Cooling plants are energy-intensive systems which enable ideal thermodynamical conditions for the smooth operation of accelerator facilities. Among this class of plants, chilled water production systems, used to bring cooling water from high to low temperatures, are commonly found. The automatic control system of such plant is usually complex due to the large number of correlated control inputs available, rendering it particularly challenging to minimize its energy consumption when solely relying on conventional control methods (e.g., PID and/or if-based logic). In 2024, CERN's Engineering department partnered with the university Politecnico di Milano to develop an energy optimal model-predictive controller (MPC) for one of the critical chilled water production plants used for the cooling of CERN's flagship accelerator, the Large Hadron Collider. This 18-month collaboration is well underway, and this paper explores the motivational and organizational aspects of the project, as well as highlights the technical solution proposed, the challenges faced to date, and how these were overcome. Simulation-based results are presented for a detailed performance comparison between MPC and the currently used rule-based logic controller. Finally, the architecture of controls and operator interface for the MPC deployment in the real plant are discussed, in view of extending this optimal control solution to numerous similar systems at CERN.

## INTRODUCTION

In the follow-up of previous successful initiatives for energy savings by controls optimization [1], the Engineering department of the European Organization for Nuclear Research (CERN) has decided to pursue further this endeavor. This time, the team decided to focus on chiller-based cooling plants, which are widely present in the laboratory, and are accountable for a significant part of the electricity consumption of the cooling and ventilation (CV) equipment of the organization. As a reference, CV plants account for up to 12% of the consumption of the flagship accelerator of the organization, the Large Hadron Collider (LHC).

Chiller-based cooling plants are particularly energy intensive plants as they incorporate high-power equipment, such as chillers and pumps, and are usually subject to high cooling loads. As an example, the overall installed cooling capacity of chiller-based cooling plants at CERN is estimated at 85 MW[1]. These plants, which are widely used in accelerator and industry facilities, are the default mechanism to produce chilled water to cool down equipment and enable air-conditioning in technical facilities. Not less important, from the industrial control's perspective, they compose a very interesting problem as the available control variables are highly correlated, and hence challenging to manage in parallel. As detailed later in the paper, these plants have historically been controlled by a rule-based strategy combined with Proportional-Integrative-Derivative (PID) controllers. This strategy does not incorporate any kind of formal optimization, being for energy savings or any other type of performance-based metric.

This paper focuses on the work developed to date on the development of a new energy-optimal control solution for chiller-based plants at CERN. In the remaining of the paper, we delve into the foundation and organization of this research collaboration, describe the basic understanding of a generic chiller-based cooling plant, and provide a high-level description of both the legacy control strategy and the new energy-optimal control. In addition, the methodology and key design considerations for the development of the energy-optimal controller are highlighted. Finally, simulation-based comparative results on energy performance are presented, as well as some concluding remarks on the definitive deployment of the new solution in the case-study plant and plants alike.

## A SUCCESSFUL RESEARCH COLLABORATION

At the core of the project discussed in this paper, there is a research collaboration gathering several scientific partners: CERN's Engineergin deparment, CERN's Knowledge Transfer group, and the Systems & Control group of Politecnico di Milano (PoliMi). As introduced, the objective set for the project was the development of an energetically optimized control system using a chilled water plant at CERN as case-study. The challenge posed by the topic was the key motivation to bring together and leverage the expertise of each participating partner: CERN contributing primarily with its experience on innovation and research project management, controls optimization for energy savings, and year's of field-experience operating chiller-based cooling plants; PoliMi with a key contribution to the solution development given its long track record of research in advanced control methods for energy systems, including techniques as Model Predictive Control (MPC), and its application to real-world problems.

The project, which is planned to last 18 months (July 2024 to December 2025), was organized by a well-defined set of sequential tasks and deliverables, each with a pre-

* diogo.monteiro@cern.ch

[1] To put in perspective, the usual capacity of a nuclear reactor is 1000 MW [2].

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCG005

# PETRA-IV FOFB SYSTEM INTEGRATION TEST SETUP

B. Dursun*, S. H. Mirza, S. Pfeiffer, A. Eichler, H. Schlarb, H. T. Duhme, G. Kube
Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

*Abstract*

The PETRA-IV Fast Orbit Feedback (FOFB) system will be a large-scale Multi-Input Multi-Output (MIMO) control system, utilizing 790 Beam Position Monitors (BPM) and 560 Fast Corrector Magnets (FCM) to maintain the desired orbit trajectory. Data acquisition and distribution will be managed across 16 supply areas and connected via an extended star network topology. This contribution focuses on system integration test setups while describing the Model-Based Design (MBD) methodologies that are being used for developing, verifying, and commissioning the complete system step by step. Integration of the components of such a large system must be systematic so that potential issues can be isolated and fixed efficiently. The setups will also be used for the characterization of sensors, actuators, and transmission lines. Subsystem identification is of utmost importance for a comprehensive understanding of the system dynamics, which will guide the design of appropriate filters and control strategies to ensure optimal orbit stabilization performance. The analysis will also precisely assess the overall system latency, which is critical for feedback bandwidth and stability.

## INTRODUCTION

The PETRA-IV project at DESY aims to upgrade the present PETRA-III synchrotron into a fourth-generation, ultra-low emittance source operating at 6 GeV. The target emittance is 20 pm·rad horizontally and 4 pm·rad vertically, which translates into significantly smaller beam sizes, approximately 7 µm horizontally and 3 µm vertically at the insertion device (ID) source points. The FOFB system is a crucial component for rejecting disturbances up to 1 kHz, thereby ensuring the required stability of 10% of the beam size [1]. This represents a significant improvement over third-generation light sources, where disturbance rejection bandwidths were typically limited to a few hundred Hertz.

This work focuses on the methodology being used in the development process of the FOFB system. It describes the model-based workflow, which utilizes the Model-In-the-Loop (MIL), Software-In-the-Loop (SIL), and Hardware-In-the-Loop (HIL) simulations throughout the design, implementation, and verification stages. The paper also details the setups being designed for integration tests and system identification. In addition to scope-specific real-time emulators, such as the beam emulation used for BPM system identification and BPM-FOFB integration, the complete system will feature an operational mode to run a real-time plant emulator for final verification before commissioning tests.

* burak.dursun@desy.de

## PETRA-IV FOFB SYSTEM OVERVIEW

The FOFB system is designed in an extended star topology, featuring one global node (GOC - Global Orbit Controller) and 15 local nodes (LOC - Local Orbit Controller) distributed across supply areas (Fig. 1). Each LOC is the center of regional sensor and actuator stars. Feedback from BPM processors ($BPM_P$) will be aggregated through sensor hubs ($BPM_H$), and the control action will be distributed to the FCM drivers ($FCM_D$) through actuator hubs ($FCM_H$).

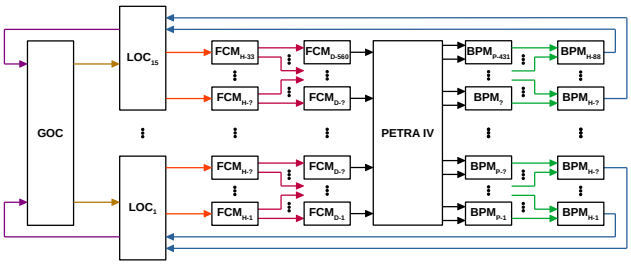



Figure 1: PETRA-IV FOFB block diagram.

The complete signal chain is planned to be realized in the MicroTCA.4 form factor excluding $FCM_D$. Fig. 2 illustrates the hardware layout at each supply space. The GOC will be located at Supply Space 1, with direct optical links to all LOCs. This standardized approach enhances modularity and maintainability while supporting the high-speed data processing required for the FOFB system.

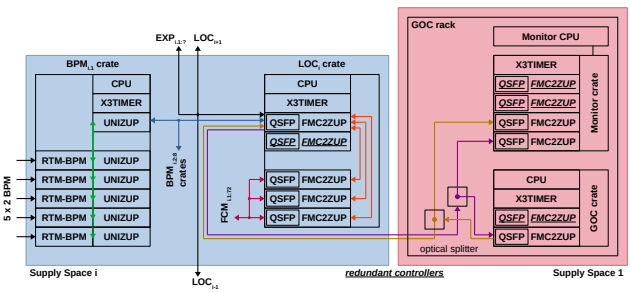



Figure 2: Hardware layout in the MicroTCA.4 form factor.

The rate of the closed-loop control operation will be 130 kHz, equal to the revolution frequency. Data propagation is optimized to minimize the overall latency. The total loop latency is estimated to be a maximum of 80 µs, which includes delays from sensors, actuators, communication links, and the controller. This latency budget is critical for achieving the desired feedback bandwidth and stability.

Data propagation in the processing chain is visualized in Fig. 3. The communication links from $BPM_P$ to $BPM_H$, and from LOC to $FCM_D$ are latency-optimized, while other links

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCR001

# DATA MANAGEMENT INFRASTRUCTURE FOR EUROPEAN XFEL

J. Malka*, S. Aplin, D. Boukhelef, K. Filippakopoulos, L. Maia,
T. Piszczek, G.Previtali, J. Szuba, K. Wrona,
European XFEL, Schenefeld, Germany
S. Dietrich, M. Gasthuber, J. Hannappel, K. Hoyos, M. Karimi, Y. Kemp, R. Lueken,
T. Mkrtchyan, K. Ohrenberg, M. Sahakyan, F. Schluenzen, K. Schwarz, S. Sternberger,
P. Suchowski, Ch. Voss, Deutsches Elektronen-Synchrotron (DESY), Hamburg, Germany

*Abstract*

Effective data management is critical to ensuring that research data remains readily accessible, efficiently processed, and optimally usable. In this article, we will describe the design and implementation of the data management infrastructure at the European XFEL, which supports high-level data management services. The system architecture is organised into four layers of storage, each addressing specific challenges associated with data handling.

The first layer, known as the Online storage, acts as a high-speed cache to accommodate the extremely high data rates generated during experiments, up to 15 GB/s from a single scientific instrument. The second layer, High-Performance Storage, supports both real-time data processing during experiments and offline post-experiment analysis. These two layers are interconnected through a single InfiniBand fabric, which is linked by a 4.4 km long, 1 Tb/s connection, facilitating rapid data transfer from the European XFEL experiment hall to the DESY computing centre.

The third layer, Mass Storage, significantly expands the system's capacity, enabling mid-term data access for more detailed analyses. Finally, the Tape Archive layer ensures secure, long-term preservation of data for a minimum retention period of 10 years. The high-performance and mass storage systems are connected to computing cluster, allowing users to conduct both near-online and offline data analysis, or export data from the European XFEL facility as needed.

The European XFEL's data management infrastructure is capable of handling and processing up to 2 PB of data per day, showcasing the exceptional performance and reliability of the system and its associated sub-services.

## INTRODUCTION

The European XFEL Facility is among the world's most advanced sources of pulsed, extremely intense, and coherent radiation in both the hard and soft X-ray regimes. These unique characteristics make it highly attractive to a wide range of scientific communities, enabling diverse experimental research. The facility spans a total length of 3.4 km, stretching from the DESY campus in Hamburg to the town of Schenefeld in Schleswig-Holstein, where the experimental stations are located. At present, experiments can be conducted using one of seven[1] available instruments, each powered by one of three self-amplified spontaneous emission (SASE) X-ray sources. The configuration of the scientific instruments and SASE sources is shown in Fig. 1. It is important to note that, at any given moment, only a single instrument is in operation on a SASE. Depending on the specific SASE, experiments are conducted either in a continuous 24-hour mode or in a 16-hour experimental mode followed by 8 hours dedicated to commissioning or in-house research on other instruments. This means that all resources are dedicated exclusively to the current experiment, and the facility operates 24 hours a day.

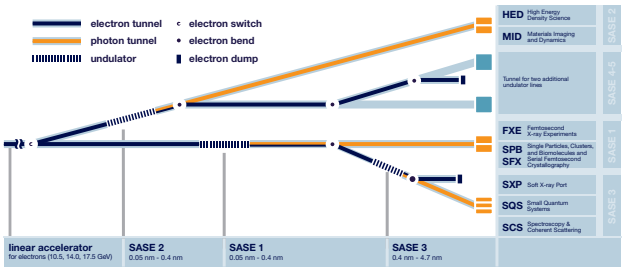



Figure 1: Beamlines and instruments of the European XFEL.

## EXPERIMENT DATA FLOW

At the European XFEL, a wide variety of data sources generate data. These sources include sensors, cameras, motors, and both 1D and 2D detectors. Detector data is transferred over a dedicated network to data acquisition (DAQ) servers, where it is aggregated by separate DAQ devices called data aggregators. Typically, multiple data aggregators run on a single server, with each aggregator producing one HDF5 [1] file that is written to disk. In parallel with writing to storage, the data is also sent over the network to the online computing cluster for real-time processing. The period during which data is collected is called a run. Each experiment consists of multiple runs, and the duration of a run depends on the experimental technique being used. Typically, a run lasts a few minutes, but runs longer than 30 minutes also occur. Breaks between runs can be used to change a sample or adjust other critical parameters. Once the setup is complete and stable, the pause between runs can be reduced to just a few seconds. As a result, the high load generated by the DAQ system on the file system becomes nearly continuous during active data collection. For example, when the AGIPD [2] detector is used in a run with other smaller detectors and data sources, a throughput of about 15 GB/s is generated from all servers

* janusz.malka@xfel.eu

[1] The eighth instrument is under construction.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCR002

# DATA AUTOMATION STRATEGY AT ESRF

V. Rey-Bakaikoa†, P.O. Autran, G. Berruyer, A. Beteva, A. Bocciarelli, M. Bodin, L. Broche, L. Claustre, A. De Maria Antolinos, W. De Nolf, S. Debionne, S Delcamp, M. Dubouchet, L. Felix, S. Fisher, J.B. Florial, G. Gaisné, A. Goetz, P. Guillou, C. Guilloud, A. Homs, L. Huder, M.C. Lagier, A. Mauro, J. Meyer, C. Muzelle, D. Naudet, A.M. Noordheen, M. Oskarsson, E. Papillon, H. Payno, S. Petitdemange, M. Ruyer, M. Spitoni, O. Svensson, V. Valls, T. Vincent, Y. Walesch
European Synchrotron Radiation Facility, Grenoble, France

*Abstract*

The ESRF has initiated an ambitious program to automate data processing and management across all 47 of its beamlines. Since the successful installation and commissioning of the ESRF-EBS new storage ring in 2020, there has been a significant increase in X-ray flux. A comprehensive strategy has been developed to optimize the use of the increasing quantities of data generated by the new advanced detectors. This strategy consists of several components, including the library for image acquisition in its second version (LIMA2) [1], which provides scalable detector control and early data processing and reduction. The new beamline control support software (BLISS) [2], coordinates synchronization and data acquisition. Additionally, the Data Automation software, the Extensible WorKflow System (EWOKS) [3], processes and exploits the produced data by running data processing workflows either triggered by BLISS online or run offline. The Data Repository for Advancing open sCience (DRAC) [4] complements the strategy by implementing a database and associated tools for managing all data and metadata generated during experiments and the associated automated pipelines. We will examine how DRAC, EWOKS, and BLISS work together and the orchestration function of Blissdata.

## BACKGROUND

In recent years, and particularly since the operation of the EBS began, the ESRF has experienced a significant increase in data production volume and the number of stored datasets (Fig. 1). However, this increase has not resulted in a corresponding rise in published scientific papers. Several factors contribute to this situation. For example, as data collection capacity has grown, so has the amount of raw data supporting each scientific publication. While a direct proportional relationship between data volume and science output is not expected, it is evident that users face growing challenges in managing and analysing the data recorded during a session.

A strategy incorporating data processing automation, early data reduction, on-site data analysis, high-performance computing (HPC) facilities, and standardisation of data management, data formats, and metadata annotation aims to streamline the transition from raw data to scientific results. The approach described in this paper covers all stages, including detector control, data acquisition, automated data processing, data management, curation, and the publication of both raw and processed data. This strategy is being implemented across all beamlines at the ESRF.

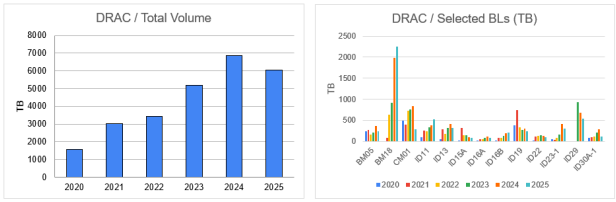



Figure 1: Total data stored in DRAC catalogue (left), with detail on largest producers (right) – covers data up to September 2025.

We are currently assessing the potential of novel techniques, including the use of AI tools where applicable and the adoption of new data transfer models. Since the strategy has been implemented across most beamlines, any advancements will be accessible to all scientific techniques as they are introduced.

The data strategy covers all phases of the data lifecycle, by integrating software components as illustrated in Fig. 2.

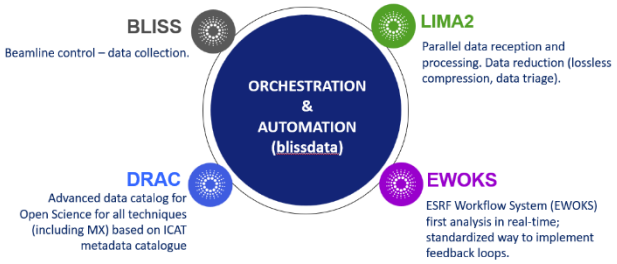



Figure 2: Key components of the ESRF data strategy.

## DATA ACQUISITION

### *BLISS*

The BLISS [2] software suite for data acquisition was developed at the ESRF as the main platform for configuring and conducting experiments. Its primary interface is the Python programming language command line, which provides users with the benefits of a high-level programming environment, robust data structures, and extensive module and extension support. BLISS enables comprehensive control of all beamline devices through a device abstraction layer, complemented by a rich set of features. The software facilitates the execution and coordination of data

† rey@esrf.fr

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCR003

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

# A LONG TERM STORAGE SOLUTION FOR TANGO ATTRIBUTE DATA AT SKAO

M. Zambrano*, T. Juerges, SKA Observatory, Jodrell Bank, United Kingdom
J. Venter, SARAO, Cape Town, South Africa

*Abstract*

At the Square Kilometre Array Observatory (SKAO), monitoring data is ingested from distributed subsystems via the Tango Controls archiver, with attribute data stored in the Engineering Data Archive (EDA). The EDA uses a PostgreSQL database with the TimescaleDB extension, offering a performant solution for time-series storage. However, as SKAO infrastructure scales, PostgreSQL becomes impractical for long-term retention due to cost and operational complexity. This paper outlines a long-term storage strategy based on S3-compatible object storage. The solution decouples operational and archival storage by exporting and serializing Tango attribute data into efficient formats like Apache Parquet for storage in S3. Metadata indexing ensures the data remains discoverable and retrievable over time. The approach draws from the MeerKAT telescope's experience, a precursor to SKAO operated by SARAO. MeerKAT faced similar challenges archiving large volumes of telemetry data and adopted a database and long term storage model. We also describe supporting tools and processes for managing data lifecycle transitions. The paper concludes with open challenges and future directions for integrating this approach into observatory-wide data access frameworks, ensuring engineering telemetry remains accessible throughout the SKAO system lifecycle.

## INTRODUCTION

Modern telescopes are complex distributed systems. They are made of thousands of networked independent devices created by different developer teams. Each single component in the telescopes can generate a myriad of continuous steams of telemetry data that can be used for:

- Operational Reliability & Troubleshooting: to check if a device is functioning according the the specification and take a corrective approach if needed. Historical data could help to detect failures due to environmental variables and schedule maintenance.

- Commissioning & Calibration: during construction each system has to be characterized and tuned to match its specification. Long term archiving could help in detecting deviation from the expected performance.

- Performance Optimisation: by having all the telemetry data, the telescope team could find operational improvements.

- Root cause and forensic analysis: after an issue is reported, the story on how it was triggered can be reconstructed from the archive and lessons can be learn from it.

- Data-driven maintenance: by using predictive algorithms in large time series, components failures can be scheduled.

The Square Kilometre Array Observatory (SKAO) [1] is currently constructing two large radio telescopes, SKA-Mid [2] and SKA-Low [3]. Its global headquarter is in Jodrell Bank, UK. The SKA-Mid telescope, which is currently constructed in the Karoo desert in South Africa, will be composed of 197 15-meters dishes. The SKA-Low telescope in Western Australia will have 512 stations that consist of 131,072 log-periodic antennas. During its construction and production, the telescope will produce an enormous amount of engineering data that will support scientists and engineers building the Observatory. All of these data need to be preserved for fifty years, which is the planned lifetime of the observatory, but there is no specification about the final format of the data. The SKAO Software component in charge of doing this task is the SKAO Engineering Data Archive (EDA) [4].

## THE MEERKAT ENGINEERING DATA ARCHIVE

MeerKAT [5], an SKAO precursor, consists of 64 radio telescopes that collectively generate approximately 6,000 engineering data sensor samples every second. The following architecture is used to support the long-term storage of these samples.

Each dish includes several KATCP device components, each generating samples at various rates and writing them to a shared messaging queue (NATS [6]) as shown on figure 1.

A separate server runs multiple processes that subscribe to the message queue and write the incoming data to a buffer table in a database. Every few hours, all samples older than one day are compressed and archived into block storage (CEPH [7]) from the database. The key used in CEPH for samples follows the following format *<days since epoch>:<sensor name>*, and the data is compressed when writing.

Clients access a single REST query interface that transparently searches and merges results from both the database and the object store, returning the data in JSON format. The REST interface supports custom filters.

* mauricio.zambrano@skao.int

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCR004

# STATUS OF THE ITER DATA HANDLING ECOSYSTEM

L. Abadie[†,1], R. Castro[2], Y. Makushok[3], P. Abreu[1], S. Simrock[1], D. Dequidt[1], G. Ferro[1], P. Martin Villares[3], O. Hoenen[1], J. Hollocombe[4], P. P. Sawantdesai[1], J. S. Cabanilla[3]
[1]ITER Organization, Saint Paul Lez Durance, France
[2]CIEMAT, Madrid, Spain
[3]INDRA, Madrid, Spain
[4]UKAEA, Culham, UK

*Abstract*

Data is one of the key deliverables of the ITER machine. Since 2019, the ITER data handling system has gradually been extended to cope with the commissioning of new plant systems and new needs. This contribution gives an overview of the different sub-systems which compose the data handling ecosystem from data archivers to visualization tools. We will summarize the short- and long-term storage for data archiving and its challenges. Different data archivers based on the Hierarchical Data Format (HDF5) [1] have been developed to collect the data coming from the fast and slow systems. Techniques to speed-up data retrieval has been implemented via data processors. We have extended the Unified Data Access library (UDA) [2], developed by UKAEA for MAST, to support our data models by adding a new plugin which uses the client-server model that UDA is based on. The same server with a different plugin is also used offline to retrieve data for processing, which is based on the (Integrated Modelling & Analysis Suite (IMAS) [3] framework. Our ecosystem includes a web application, which allows plotting both archived data and live data. For interoperability purposes with other systems, a plugin for Grafana [4] has been developed to create dashboards which mix archived data and monitoring metrics. In collaboration with our Science and Diagnostic teams, we are developing visualization tools to allow plotting offline data and experimental data. Finally, we will conclude with key takeaways and an outlook on the impact of AI on the data handling system.

## INTRODUCTION

The Data Handling system has been built by pieces. In 2010, the COntrol, Data Access and Communication (CODAC) team decides to develop and release on a yearly basis a "CODAC Core System", (CCS) [5] which allows the Instrumention and Controls (I&C) developers to interface with the Control System. One of the main objectives of the CCS is to ease the integration with the CODAC system once the system is delivered on site, and one of its the main requirements is to ensure backward compatibility of the libraries. The data acquisition system via the DAN (Data Archiving Network) library and its archiver has been integrated in CCS 5.0, released in 2015.

## DAN

The main objective of DAN is to allow archiving of data at high resolution (at least kHz rate). DAN does not permit control or monitoring of the data. Other networks fulfil this role.

### *Data Format Selection*

The development of the DAN acquisition system and archiving started in 2013. The first stage was the selection of 7the data format. In 2011, two parallel surveys by two independent entities were conducted to select the most suitable data format for ITER. The main candidates at that time were MDSPlus [6], ROOT [7] and HDF5. Developing an ITER dedicated format has been ruled out for budget reasons. Both reports concluded to opt for HDF5. The main strengths of HDF5 are self-described format, support by many scientific communities and the release of SWMR (Single Writer Multiple Reader), allowing to read the data while the file is open.

Also, unlike the other formats, HDF5 is a data format only. Archivers and data access must be developed from scratch.

### *Architecture*

The DAN consists of two parts: the client that publishes data and the archivers which receive and write the data to HDF5 files. The DAN protocol has been developed in-house based on TCP/IP. Referring to Fig. 1, the I&C developer uses the DAN library in the acquisition program to publish the data in a shared memory that it is read by the DAN streamer. The DAN streamer connects to the DAN archiver. Splitting the DAN into 2 sub-modules on the client part makes the interface with the I&C developer more straight forward as only data need to be put on the shared area. Communication with the archivers is handled by the streamers.

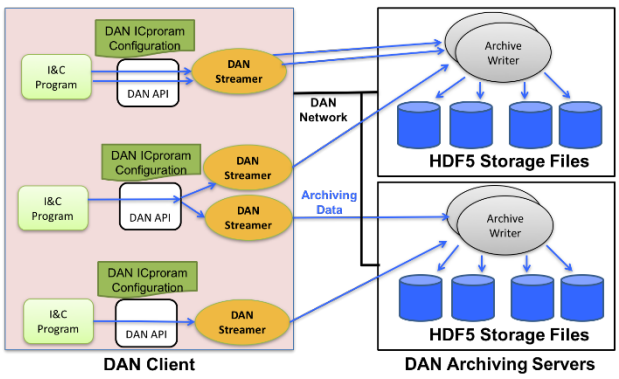



Figure 1: Outline of the DAN architecture.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THCR005

# myMdC: A CENTRALISED METADATA CATALOGUE FOR SUSTAINABLE DATA MANAGEMENT AT EUROPEAN XFEL

L. Maia*, S. Aplin, D. Boukhelef, F. Dall'Antonia, I. Derevianko, L. Gelisio, D. Goeries, S. Hauf, N. Liss, J. Malka, M. Manetti, G. Previtali, P. Schmidt, V. Singh, J. Szuba, W.-T. Tsai, B. Vanganuru, K. Wrona, European X-Ray Free-Electron Laser†, Schenefeld, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Large-scale research facilities such as the European XFEL rely on multiple distributed services to manage complex experimental workflows. Consistent, reliable and timely metadata exchange between these services is essential to ensure automation and data integrity.

This paper presents **myMdC**, the centralised metadata backbone that supports scientific data services in European XFEL. *myMdC* unifies and orchestrates metadata flows across core facility services, enabling interoperability and ensuring alignment with the facility's long-term data management strategy.

In operation since the first day of user experiments in 2017, myMdC ensures real-time injection, documentation, and cataloguing of scientific datasets, forming the basis of data and metadata operations at European XFEL. Its integration with facility services enables immediate dataset tracking from acquisition through publication. It supports automation across the data inventory, control access during embargo periods, metadata preservation, Digital Object Identifier (DOI minting and integration of Findable, Accessible, Interoperable, and Reusable (FAIR) data workflows.

By consolidating metadata management into a unified platform, myMdC facilitates scalable, transparent, and future-proof data practices, ensuring that scientific output remains robust and interoperable within the global research community.

## INTRODUCTION

Modern accelerator-based facilities such as the European XFEL operate a diverse ecosystem of services responsible for experiment scheduling, configuration, control, as well as data acquisition and processing. Ensuring that these services have consistent and validated metadata at the right time is essential for automation, interoperability, reproducibility, and overall efficient operation of the facility.

Historically, metadata management has often relied on domain-specific tools or custom scripts, which can lead to duplication, inconsistencies, and operational complexity [1]. At the scale of the European XFEL, such limitations become critical: experiments involve many interdependent services, and produce large volumes of data and metadata such as configuration parameters and provenance information.

To overcome these challenges, *myMdC* [2] was developed at European XFEL as a centralised metadata backbone, designed to deliver a unified and reliable infrastructure for consistent metadata preservation, seamless integration, and orchestration across facility services. *myMdC* serves as the authoritative source of metadata, ensuring consistent and reliable information for all dependent systems. Since its deployment in 2017, it has become a cornerstone of the European XFEL infrastructure, providing high availability across all experiment stations and demonstrating flexibility through seamless integration with newly introduced services.

This paper describes the most important design and operational features of myMdC. We first present its layered architecture, followed by details of the services integration approach and operational lessons from seven years of production use. Finally, we outline future developments.

### *European XFEL Overview*

The European X-ray Free-Electron Laser (European XFEL) is a facility featuring a 3.4 km long linear accelerator producing pulsed, ultra-intense, and coherent X-rays in hard and soft regimes. Spanning from DESY[1] (Hamburg) to Schenefeld (Schleswig-Holstein), it hosts seven scientific instruments powered by three self-amplified spontaneous emission (SASE) sources [3,4]. At any time, one instrument operates per SASE, in either 24 h continuous or 16 h user beamtime / 8 h commissioning cycles. This scheduling maximises beamline efficiency, ensuring full-capacity delivery for each experiment.

In this context, myMdC is responsible for gathering and documenting experiment metadata and for ensuring that proposal-related resources and team assignments are prepared to support a successful experimental campaign.

### *Data Management*

European XFEL generates several petabytes of experiment data per week, amounting to tens of petabytes annually, driven by MHz-repetition-rate X-ray large-area detectors [5]. Data systems perform real-time capture, annotation, and quality control, feeding high-throughput pipelines for calibration, reduction, and feedback during experiments [6–8]. Raw and processed data are stored on European XFEL data systems integrated into the DESY data centre and are managed jointly with DESY [9]. The tiered storage architecture balances rapid access for active experiments and cost-efficient long-term archival, but it requires coordinated data migration between layers. Common services (such as data analysis clusters) are available for users who can access them on-site and remotely.

* Corresponding author: luis.maia@xfel.eu

† https://www.xfel.eu/

[1] Deutsches Elektronen-Synchrotron

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG001

# ACCESS CONTROL INTERLOCK SYSTEM UPGRADE AT ANL/APS*

J. Stevens†, G. Shen, K. Belcher, J. Lenner, V. Nguyen, S. Thakkar
Argonne National Laboratory, Lemont, IL, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

During the 2023 APS Upgrade (APS-U) 'dark-year' shutdown, the Advanced Photon Source Access Control and Interlock System (ACIS) was overhauled. The modernization replaced programmable logic controllers (PLC) and redesigned the read-only status interface to the Experimental Physics and Industrial Control System (EPICS). This paper summarizes ACIS architecture and details the upgraded EPICS interface.

## INTRODUCTION

The Access Control Interlock System (ACIS) protects personnel from prompt radiation at the Advanced Photon Source (APS). ACIS enforces controlled access to accelerator enclosures and terminates the particle beam when required. An ACIS trip disables radio-frequency (RF) amplifiers and modulators, steering and focusing power supplies, and the electron gun, inhibiting delivery of user beam [1].

Commissioned in 1996, the original implementation used redundant programmable logic controllers (PLCs) tied to independent, hardwired safety chains and delivered reliable performance for decades [2].

By the 2020s, Allen-Bradley 1771 Series PLC-5 obsolescence pushed the architecture to end-of-life. The EPICS interface used text-based database files with legacy record names. The 2023 APS-U dark-year shutdown enabled migration to Allen-Bradley GuardLogix 5580 safety-certified PLCs and a redesigned, read-only EPICS status interface.

## ACIS OVERVIEW

ACIS protects personnel from prompt radiation in five accelerator sections:

- Linac Extension Area (LEA)
- Linac and Positron Accumulator Ring (LP)
- Booster (BOO)
- Radio-Frequency Wing, Zone F (RF)
- Storage Ring, Zones A–E (SR)

Each section has a dedicated ACIS PLC that monitors local field devices and exchanges interlock status with associated ACIS PLCs. System integrity is enforced by two independent PLC chains (A and B) that supervise enclosure security.

In addition, more than 50 experimental beamlines use PLCs for their Personnel Safety Systems (PSS). The beamline PSS employs a dual-chain, PLC-based architecture to monitor safe entry and exit of experimental stations. The PSS PLCs interface with the Zone F and storage ring ACIS PLCs to control downstream safety shutters and beam stops.

ACIS field hardware includes emergency-shutdown buttons, controlled-access key switches, radiation and beam-current monitors, relays, limit switches, watchdog timers, message displays, magnetic locks, and control panels wired into remote I/O chassis.

## LEGACY EPICS HARDWARE

The original Access Control and Interlock System (ACIS) used Allen-Bradley PLC-5 processors located in the Main Control Room (MCR), with PLC fiber-optic links fanning out to remote 1771 I/O chassis (Fig 1). The EPICS read-only status-interface hardware, consisting of dedicated ACIS PLC "Read Only" modules, also resided in the MCR.



Figure 1: Legacy ACIS racks in the Main Control Room.

The ACIS status interface used an EPICS Input/Output Controller (IOC)—a VME single-board computer running the VxWorks real-time operating system—hosting Allen-Bradley 6008-SV2R VME Remote I/O scanner cards. The scanner Remote I/O ports were wired via "blue-hose" aka "Data Highway" twisted-pair cable to nearby MCR PLC-5 controllers (Fig. 2).

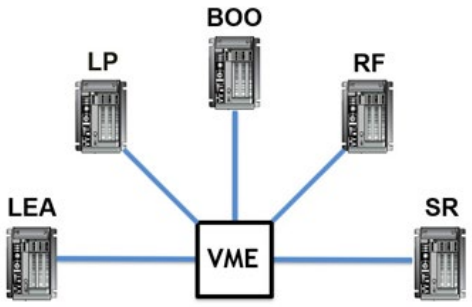



Figure 2: Allen-Bradley VME Scanner to PLC-5.

* Work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No, DE-AC0206CH11357
† jbs@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG002

# THE NOVEL AND ROBUST DESIGN OF FAST PROTECTION SYSTEM FOR CSNS-II

Peng Zhu[1, 2, *], Lei Zeng[1, 2], Jianmin Tian[1, 2], Yuliang Zhang[1, 2], Renjun Yang[1, 2, 3], Xiaohan Lu[1, 2, 3], YongCheng He[1, 2], Kang Jia Xue[1, 2], Xuan Wu[1, 2], Lin Wang[1, 2, 3], Mingtao Li[1, 2, 3], SiNong Cheng[1, 2]

[1] Institute of High Energy Physics, Chinese Academy of Sciences (CAS), Beijing, China
[2] Spallation Neutron Source Science Center, Dongguan, China
[3] University of the Chinese Academy of Sciences, Beijing, China

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

To ensure the long-term stability and reliability of the hardware platform under continuous operation, a key innovation lies in the custom design and implementation of a high-performance architecture, specifically, the "FPGA + Rocket IO + ATCA" solution, tailored for the Fast Protection System (FPS) of CSNS-II. This self-developed cooperative structure, incorporating both backplane and front panel functionalities, achieves functional decoupling and modularity while offering superior scalability and maintainability. The design effectively minimizes transmission delays and signal degradation associated with intermediate logic components and shared bus architectures, thereby enhancing timing determinism and signal integrity in inter-board communication. This paper will detailed description that the system is well-suited for fast protection and real-time control systems that require nanosecond-level synchronization precision, simplifying hardware expansion and augmenting maintainability and adaptability.

## INTRODUCTION

The China Spallation Neutron Source (CSNS) is the fourth large-scale pulsed high-intensity proton accelerator globally, following the Spallation Neutron Source (SNS) in the USA, the Japan Proton Accelerator Research Complex (J-PARC), and the ISIS facility in the UK. This milestone signifies a significant leap for China in the realm of major scientific facilities, transitioning from follower to peer status. CSNS successfully passed national acceptance tests in 2018 and has since been operational, providing a powerful neutron probe platform for cutting-edge research in materials science, life sciences, physics, chemistry, and other multidisciplinary fields [1].

Since its inception, CSNS has continuously enhanced beam power and operational stability, achieving stable beam delivery exceeding 100 kW and steadily advancing toward higher power levels of 170 kW. To further increase neutron yields and expand experimental capabilities, the National Development and Reform Commission officially approved the launch of CSNS Phase II (CSNS-II) in 2024, marking a new phase of transformative development for the facility.

One of the most critical upgrades in CSNS-II involves a substantial overhaul of the existing LINAC. Specifically, after the initial 80 MeV radio frequency quadrupole (RFQ) and drift tube Linac (DTL) sections, a new superconducting(SC) segment has been introduced. This upgrade utilizes high-efficiency Spoke cavities and elliptical cavities, boosting the output beam energy from 80 MeV to 300 MeV, as shown in Fig. 1. This enhancement significantly reduces the acceleration burden on the rapid cycling synchrotron (RCS), improves injection efficiency, and enhances overall transmission efficiency within the accelerator complex. It lays the groundwork for achieving beam powers of 500 kW.

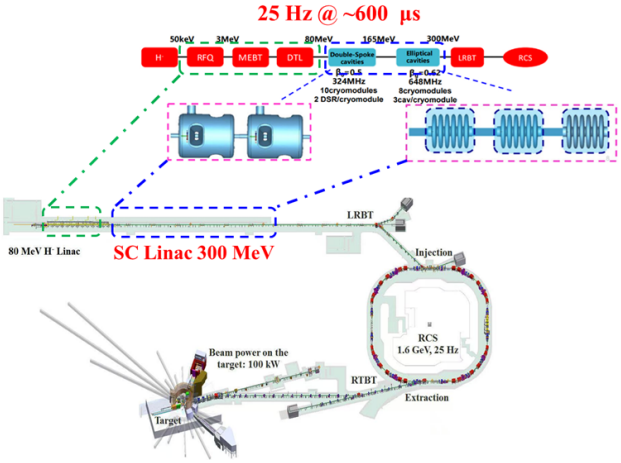



Figure 1: Overall layout of the CSNS-II LINAC.

## REQUIREMENTS OF THE FPS

### *The Response Time*

With the significant increase in beam power, the safety risks associated with facility operation have correspondingly risen, imposing unprecedentedly stringent requirements on the fast protection system(FPS) capabilities of critical systems. In particular, the FPS must accurately detect abnormal operational conditions and promptly terminate the beam within microsecond-level response times to prevent equipment damage and other safety hazards caused by beam loss. Therefore, the development of a FPS featuring high reliability, rapid response, and intelligent diagnostics has become a key technological challenge for ensuring the long-term stable, safe, and reliable operation of CSNS-II, as shown in Fig. 2a. Comprehensive upgrades and innovative break-

* zhup@ihep.ac.cn

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG004

# ESS ACCELERATOR PERSONNEL SAFETY SYSTEM JOURNEY TOWARDS STEADY STATE OPERATION

V. Harahap, A. Farshidfar, A. Abujame, A. Nordt, A. Andersson, A. Petrushenko, C. Webber, D. Plotnikov, D. Daryadel, J. Gustafsson, J. Grönvall, J. Lastow, L. Nunes, M. Zmuda, M. Carroll, M. Eriksson, M. Mansouri, N. Naicker, P. Holgersson, R. Foroozan, Y. Takzare
European Spallation Source ERIC, Lund, Sweden

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

## *Abstract*

The European Spallation Source (ESS) Personnel Safety System (PSS) is a critical, integrated safety system designed to protect personnel from hazards associated with the operation of the ESS facility. One of the core components of the ESS is the linear accelerator that accelerates protons to high energies. The Accelerator PSS (ACC PSS) serves the overall safety interlock and access control functionalities for the linear accelerator.

The safety interlock mitigates various hazards (ionising radiation, high voltage) associated with the proton beam operation and Radio Frequency (RF) equipment. The access control is achieved by implementing an integrated access management system, where a Personal Access Station (PAS) is merged with a Physical Access Control System (PACS), allowing fully automated passage. This paper provides lessons learned throughout the ESS ACC PSS commissioning journey

## ACCELERATOR PERSONNEL SAFETY SYSTEM

The purpose of the Accelerator Personnel Safety System (ACC PSS) is to prevent personnel from being harmed by exposure to prompt ionising radiation as generated by the proton beam or by sending Radio Frequency (RF) power to the cavities in the accelerator tunnel. This is achieved by restricting access to the 600-metre accelerator tunnel.

The ACC PSS provides two main functionalities: the Safety Interlock System and the Access Control System [1].

### *Safety Interlock System*

The safety interlock system is a series of safety interlocks installed to ensure the safety of personnel working in the linear accelerator. The safety interlock system consists of sensors, a logic solver, and actuators that de-energise the proton beam and RF systems under conditions that lead to ionising radiation hazards.

The ACC PSS establishes interfaces with the RF system to ensure an effective interlock. The primary interfaces are the interface with the Moulded Case Circuit Breaker (MCCB), which interlocks the main incoming power to the modulator, and the interface with the Radio Frequency Distribution System (RFDS), which disconnects the RF systems from the accelerator tunnel.

### *Access Control System*

Access to the accelerator tunnel is regulated by the ESS Physical Access Control System (PACS). The PACS is integrated with the ACC PSS through the use of the Personnel Access Station (PAS) to ensure access management. A PAS is located in each entry point, through which personnel enter and exit the accelerator tunnel.

The PAS is a double-gated access system in place to ensure single-person entry to and exit from the accelerator tunnel. In order to ensure a single passage, a PAS is equipped with multiple means, such as infrared sensors on the sides and at the top of the PAS, and a pressure floor mat with a marked area where a person needs to stand.

## RADIO FREQUENCY DISTRIBUTION SYSTEM BYPASS

### *Radio Frequency Isolation*

The disconnection of an RF system from the accelerator tunnel is realised through the ACC PSS interfaces with its RFDS. A typical RF isolation is removing the RF waveguides and installing the short circuit plates, see Fig. 1

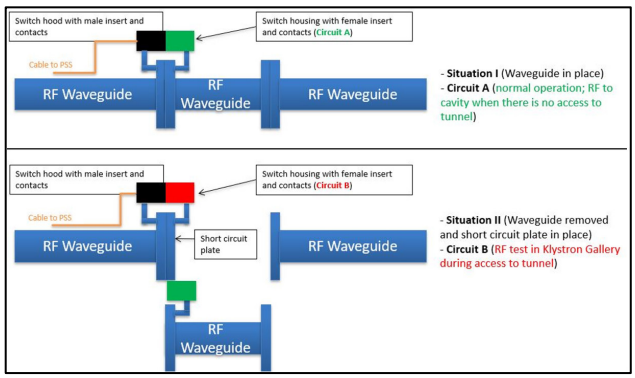



Figure 1: Typical RF isolation.

ACC PSS monitors the position of the waveguides and the short circuit plates on the RFDS and energises or de-energises the respective RF system accordingly. Only if a dedicated section of the waveguides is removed and the short circuit plates are in place, the power to the respective RF systems be enabled by the ACC PSS after an acknowledgement from an operator is performed on the Human Machine Interface (HMI). This state is well-known as RF isolation. At this state, there is no possibility for RF to reach the cavity; therefore, ACC PSS grants its permit for RF generation while the accelerator tunnel is open for access.

Based on the experience gained in the Normal Conducting Linac (NCL) trial operation, the ACC PSS monitoring function of the RFDS might need to be bypassed to enable certain test activities of the RFDS. These test activities

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG006

# DESIGN AND PERFORMANCE OF THE SNS CREDITED BEAM POWER LIMIT SYSTEM*

K.L. Mahoney[1†], A, Aher[1], T. Allison[2], M. Bobrek[1], P.A. Bong[1],
C.E. Deibele[1], N. Evans[1], K. Kasemir[1], C, Michaelides[1]
[1]Oak Ridge National Laboratory: Spallation Neutron Source, Oak Ridge, TN, USA
[2]Osprey Distributed Control Systems, Ocean City, TN, USA

*Abstract*

The US Spallation Neutron Source (SNS) is the world's most powerful pulsed spallation neutron source. The recently completed Proton Power Upgrade (PPU) project doubles the available average beam power from 1.4 to 2.8 MW. However, 0.8 MW of that is intended for a future second target station (STS), which is in the preliminary design phase. The mercury based first target station (FTS) has a safety design basis limit of 2.0 MW, thus a precision credited safety system is needed to ensure the proton beam power to the FTS cannot exceed this limit. The Beam Power Limit System (BPLS) is a novel FPGA/PLC based safety credited instrument that measures beam energy and charge to calculate beam power and will shut off the proton beam if the power exceeds 2 MW + 10% for a specified period of time. This paper will discuss the design challenges and operational performance of the BPLS as a credited system.

## MOTIVATION

The Spallation Neutron Source (SNS), Fig. 1, uses short pulses of protons (~0.75μs) to generate intense pulsed neutron beams for applied and basic research. In 2019 the facility started the proton power upgrade (PPU) project to double the average proton beam power from 1.4 to 2.8 MW. This is accomplished by increasing the beam energy by 30% to 1.3 GeV and the average beam current by 70% to 2.3mA [1]. One key deliverable of the PPU project is to enable routine operation of the First Target Station (FTS) at its original design limit of 2.0 MW. The SNS accelerator safety envelope (ASE) defines absolute limits and equipment and processes that must be in place to ensure the facility is operated safely. These limits, equipment and processes are "credited" as necessary and accountable for safe operations. Exceedance of ASE parameters requires immediate shut down of facility operations and approval of the US Department of Energy for restart.

The ASE limit for beam power to the FTS is defined as no more than 2.0 MW + 10% for one minute [2]. Because the post PPU beam power could significantly exceed the 2.2 MW for one minute safety limit, a credited system is required to ensure this limit is not exceeded. A Beam Power Limit System (BPLS) was included in the PPU project scope to perform this function.

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under contract number DE-AC05-00OR22725
† mahoneykl@ornl.gov

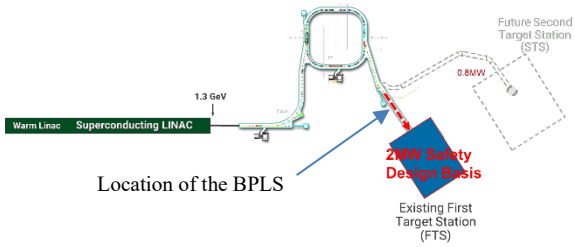



Figure 1: SNS Accelerator showing the BPLS location in relation to the First Target Station (FTS) and future Second Target Station (STS).

## REQUIREMENTS

The BPLS must independently measure beam power in real time and shut off the proton beam if it exceeds the ASE requirement. Average beam power is the product of repetition rate, beam kinetic energy and integrated pulsed beam charge [3] as given in Eq. (1).

$$P_{avg} = R_{rate} E_b \int_{t_0}^{t_0+w} dt\, I(t) \ \text{W} \qquad (1)$$

In the SNS, beam energy is normally fixed for a given run but may vary from run to run. Beam energy may also be reduced when accelerating elements are out of service. The BPLS must be able to perform its safety function with energies between 0.5 and 1.5 GeV.

Beam charge is varied as needed to reach a given beam power at the operating energy. Therefore, the BPLS must be able to work with peak currents as high as 150 Amps.

Complicating this is the need for routine operations very close to the limit while maintaining sufficient margin to avoid false trips. Because the ASE limit is an absolute number, any BPLS measurement and trip threshold must allow for uncertainties in the power measurement. Uncertainties from all sources not only include equipment and signal processing errors sources but also errors attributable to factors such as test equipment calibration. This has the effect of reducing the margin between the operating point and the trip setpoint.

The worst-case time for the BPLS to shut off the proton beam was calculated to be 4 seconds [4], therefore the specification for the BPLS integration time was reduced to 56 seconds to remain within the one minute limit.

As a credited system, the BPLS design and implementation must follow best practices for a safety instrumented system. This means use of fault tolerance, diagnostics, and rigorous Verification and Validation (V&V) for credited hardware and software. Credited functions should be

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG007

# LOGGING A NEW ERA AT THE APS USING BELY*

D.P. Jarosz†, L. Xiao, E. Chandler, G. Shen, S. Veseli, N.D. Arnold
Argonne National Laboratory, Argonne, IL, USA

*Abstract*

As the "Dark Year" of Advanced Photon Source Upgrade (APS-U) concluded, a new logbook was essential to document the process of bringing the facility back online. The Best Electronic Logbook Yet (BELY) has been developed and deployed as a solution to fulfill this requirement. This paper dives into the development process and technologies used to create BELY. Additionally, it will explore the features BELY provides to address all of its operational requirements. One of the significant strengths of BELY is its broad adoption across the APS, driven by its well-organized structure. The widespread use at the APS significantly enhances communication between teams responsible for maintaining the machine, ensuring that information is easily accessible, and collaboration is seamless. Furthermore, the paper discusses various uses of BELY. Finally, it presents ideas for the future development and enhancement of BELY.

## BACKGROUND

As the facility was being brought back online, it was a great opportunity to unify the communication amongst various technical groups within the APS. Historically, many groups have deployed their own logbooks to fulfill their logging requirements. This is not ideal since it is hard to keep track of all the logbooks, especially for intergroup communication. Logbooks from other facilities were considered, such as the Olog [1] and Fermilab's Elog. The functionality of Elog seemed quite good, but the project wasn't open source. This made it difficult to fully evaluate this project. The Olog was also a promising option but lacked the hierarchical structure and customization required by the groups. During the evaluations the findings were that neither project fulfilled all the requirements and would require lots of development work. After the considerations were made, a new option became more obvious, to investigate the possibility to implementing a whole new logbook application.

### *Journey to BELY*

Implementing a new logbook is not the obvious choice because it is very time-consuming. Taking a step back to evaluate an existing project that can be modified into this logbook was the answer. This existing project is the Component Database (CDB) [2–5]. Starting with a component centric application doesn't seem like an obvious choice for a logbook application, however; it accelerated the development to create a working logbook proof of concept. More importantly it exposed some strengths of a component centric system such as its organization, permission model, and generic properties. This project has been implemented in a very generic way to allow adding functionality that can represent many concepts such as component, inventory, machine design, cables, apps, etc. With the expertise and experience adding a new representing concept to CDB is not as time-consuming as creating a whole new application. This new concept is to represent log documents. It was possible to create a proof of concept CDB based logbook within about a month. After implementation of the critical functionality such as log centric document views, search, templates, shift creation utility and cleanup of component centric domains the first production release of BELY [6] was ready.

## TECHNOLOGY STACK

BELY is a java application that runs within Payara server. It uses MariaDB as its database backend. This stack is very similar to the application its based on, the Component Database (CDB). Below is a high-level overview of some additional technical changes such as selecting the log entry text format and changes to the database schema. These are changes that needed to be made to support the requirements of BELY.

### *Log Entry Format*

Since this is a logbook application, some way of storing formatted text is needed. Markdown [7] was the format of choice. This format was chosen due to its widespread adoption. It is currently the standard format that is used on many code repositories such as Gitlab, Github and others. It is also easy to read in its unformatted state. This makes it a good choice to feed to Large Language Models (LLM). Markdown seems to also be the output of choice for most LLM. Unfortunately selecting a popular format does not mean that everyone knows how to use it. To aid in ease of use some features have been added such as ability to paste links directly, which get converted to Markdown link. Another feature that was added to ease Markdown usage was the ability to paste images or files. These get uploaded and automatically attached either as a link or an image based on its format. Uploading files can also be done using a file browser widget with similar user experience however the file reference get attached to the end of the log entry. Finally, a simple syntax dialog is also provided for guidance as shown in Fig. 1.

### *Database Schema*

Though much of the database schema is based on CDB, a few modifications were made to support the logbook appli-

* The work is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357

† djarosz@anl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG009

# STANDARDIZING UI/UX ACROSS ACCELERATOR LABS

T. Tran†, T. Summers, SLAC National Accelerator Laboratory, Menlo Park, CA, USA
S. De Leon, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
M. Polzin, Fermi National Accelerator Laboratory, Batavia, IL, USA
R. Hill. Z. Spielman, C. Kovesdi, Idaho National Laboratory, Idaho Falls, ID, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

During February 26–28, 2025, the first-ever particle accelerator user interface/user experience (UI/UX) workshop was held at SLAC. Attendees had backgrounds ranging from software development to control systems management and human factors (HF) science. The workshop began with participants discussing the current state of UI/UX procedures and practices at their respective laboratories to share experiences and learn from one another. Additional discussions focused on how to effectively integrate UI/UX best practices into actionable goals for developers, managers, and operators when working on new or existing interfaces. The goal of the working group is to create a website that will guide developers, managers, scientists, and end users at accelerator laboratories in incorporating UI/UX best practices into software development. The working group continues to meet virtually toward this goal, and is planning a second workshop for next year.

## INTRODUCTION

### *Importance of Shared Tools & Their Usability*

In accelerator labs, success is heavily dependent on deep collaboration between many groups including software engineers, operators, control systems specialists, and physicists—all with very different expertise, specialized skills, and unique perspectives. They align with one another through their tools and workflows. Their environment demands consistency, efficiency, and above all, safety. Therefore, the tools they use should aid in meeting these demands.

### *Usability, not Aesthetics*

When we talk about standardizing UI/UX across accelerator labs, we don't mean forcing everything to look or work the same. Sure, design may be subjective—but usability isn't. Aesthetic design is a matter of taste, whereas usable software is a matter of function. Things like readable font sizes, consistent behaviour, and error prevention are not aesthetic preferences but they're functional necessities. Figure 1 illustrates this point: the left and middle displays show typical HMIs that can cause confusion and limit effectiveness, whereas the rightmost example applies high-performance principles to create a display that is consistent and usable.

Too often, scientific control systems are designed around the system's architecture, while the user experience is overlooked. Our goal is to shift that mindset. As many labs modernize their systems, now is the time to bring UI/UX considerations to the forefront.

This paper shares not just why this matters, but also how we're starting to address it through intentional design using UI/UX guidelines tailored for accelerator labs. We're consolidating proven practices, encouraging knowledge sharing, and building resources that make it easier for teams to create intuitive, usable software.

## MOTIVATION

### *Not Standard UI/UX, but HPHMI*

Designing for control systems is very different from designing for consumer applications. In our case, it takes the form of human-machine interfaces (HMIs). We're not designing for paying customers but we're designing for our colleagues: the operators, engineers, physicists, and experts who need to do their jobs safely and efficiently using safety-critical systems.

### *Impact*

Accelerator labs, like refineries and other industrial companies, depend on well-designed HMIs to ensure the right decisions are made for dangerous equipment. Other industries have learned the hard way that poor design has consequences. In refineries, poorly designed control displays contributed to tragic incidents including 15 deaths and 180 injuries at BP's Texas City Refinery [2],

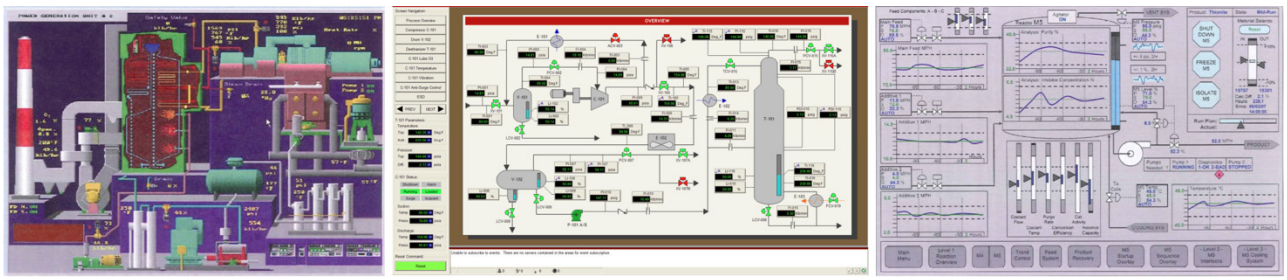


Figure 1: Poor HMI (left), more recent but still ineffective HMI (middle), effective high-performance HMI (right) [1].

† tiffany@slac.stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG011

# ENHANCING USER EXPERIENCE THROUGH GUIS AND HELPER ROUTINES AT MOGNO, THE MICRO AND NANO TOMOGRAPHY BEAMLINE AT SIRIUS

L. B. C. Campoi*, G. S. R. Costa, L. E. P. Vecina, B. V. P. Carlos, N. L. Archilha
Brazilian Synchrotron Light Laboratory (LNLS), Campinas, Sao Paulo, Brazil

*Abstract*

Since its inception at the first Brazilian synchrotron (UVX), the tomography beamline group has been concerned with the beamline's usability, recognizing that operation can be a significant challenge for users. This has been carried forward to MOGNO (micro and nano tomography beamline at SIRIUS), being one of the cornerstones for user side software development, aimed at a diverse user pool that encompasses researchers from various fields and employing a range of tomography techniques (e.g., classic, phase contrast, zoom, and in-situ 4D). To effectively support this heterogeneity of users and techniques, the development of MOGNO's operation software stack focused on three key objectives: providing an easy-to-use set of tools, enabling user experiments; ensuring a smooth learning curve for new and experienced users; and enable the beamline's staff to efficiently prepare it for experiments. The resulting software stack is composed of a set of libraries and helper scripts built with Python (e.g., making sample alignment easier), that are run from GUIs (Graphical User Interfaces), developed with PyQt and PyDM. More recently, new possibilities for improving user experience are being explored, using web applications, which can lead to more accessible beamline experiment information.

## INTRODUCTION

MOGNO beamline (micro and nanotomography) at Sirius is a fast and versatile experimental station, capable of housing experiments from many backgrounds and several different tomography types, with zoom-in (variable pixel size and FOV, afforded by the cone beam geometry) and 4D (time-resolved) experiments being flagship use cases, using two different endstations [1]. Furthermore, it also works as an entry point for synchrotron analysis at Sirius, providing users with a more straightforward experiment and easier sample preparation, coupled with a relatively quick experiment execution time (in the order of minutes), cementing a secondary responsibility for the beamline: working as a nucleus for a heterogeneous community to flourish around.

To this extent, ease of use and efficiency have been cornerstones in the development, which have been approached at different stages, from beamline preparation done by the staff, support by local contacts (staff researchers that support external users during beamtime) or operation by users. Historically, these issues have been tackled via the creation of GUIs and automated routines to be executed during the experiment. More recently, with the need to serve new users almost every day, standardized routines were also implemented during beamline preparation and user onboarding, and new frontiers are being explored with the planning of the integration of new tools.

## STANDARDIZATION OF ROUTINES

During the first regular user calls, difficulties while explaining the beamline's functions and limitations became a point of attention, as many users didn't have previous tomography or synchrotron experience. On the other hand, the rapid development of new features by support groups and the beamline's engineering team during this period, have made the need to define a secondary point of contact as a support for the local contact, a so called technical contact from the beamline engineering team.

Alongside the addition of the technical contact, this need to improve the knowledge transfer during the first moments at the beamline prompted the definition of a standardized onboarding strategy. This procedure is illustrated in the diagram in Fig. 1a, and can be executed by either the local contacts or the technical contact allocated for each proposal. As shown in the left side of 1a, this process begins with tasks outside of beamline experiment control, validating if the users need and have access to the facilities for sample preparation and IT resources, and ensuring each user in the proposal is able to perform the search safety procedure. The right side of Fig. 1a illustrates the procedures where the users are oriented on how to operate the beamline during the experiment, and is directly connected to the configuration and execution of the automated routines available through control and reconstruction GUIs.

In parallel, the technical contacts also execute a pre-experiment check, described in the flow diagram shown in Fig. 1b. This procedure was inspired by the pre-flight checks performed before an aircraft can take off and before surgeons start a medical procedure, ensuring safety in each field [2]. In MOGNO's pre-experiment case, the checks performed were defined to ensure equipment safety and the quality of experimental results. The procedure is performed using a mix between the "do-verify" and "do-list" methods described in [2], where each step is done by a single person following a to-do list of checks available in an internal group resource wiki. The check starts by ensuring that anything related to the previous experiment or user (sample or equipment) is removed from the beamline, which is shown in the left side of the diagram in Fig. 1b. The right side of the diagram shows the checks related to insuring that the

* lucca.campoi@lnls.br

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG012

# A MULTI-LEVEL MONITORING INTERFACE FOR THE SKA CENTRAL SIGNAL PROCESSOR USING THE TARANTA SYNOPTIC VIEW

G. Marotta*, C. Baffa, E. Giani, INAF-OAA, Firenze, Italy
V. Alberti, INAF-OATs, Trieste, Italy
M. Canzari, INAF-OAAb, Teramo, Italy
G. Brajnik, University of Udine and IDS, Udine, Italy

*Abstract*

Graphical user interfaces (GUIs) play a critical role in the operation and maintenance of large-scale distributed control systems. In this work, we present a synoptic-based visualization for the Central Signal Processor (CSP) of the SKA (Square Kilometre Array) telescope, developed using Taranta, a web-based visualization tool for TANGO systems. The synoptic view provides an intuitive, multi-level representation of CSP, from the upper-level control managed by CSP.LMC down to individual data processing subsystems, including CBF, PSS, and PST devices. The synoptic diagrams are created using Inkscape and exported as SVG files, enabling flexible and straightforward integration with Taranta. A key focus of this work is the collaboration between different subsystem teams, which was essential to accurately model and visualize the full CSP hierarchy. We also define a strategy to validate the effectiveness and the usability of the GUI, ensuring it delivers tangible value in improving the monitoring and troubleshooting capabilities of complex control systems such as SKA.

## INTRODUCTION

The Square Kilometre Array (SKA) is currently under construction and represents the most ambitious radio telescope ever built, both in scale and in scientific potential. When finalised, it will consists of two large interferometers located in South Africa and Australia, that will search the cosmos at Mid (350 MHz – 15.3 GHz) and Low (50 MHz – 350 MHz) frequencies ranges respectively [1].

With the first milestones already achieved and the system progressing through its commissioning phase, the project is now transitioning from the earliest integration stage, known as Array Assembly 0.5 (AA0.5), toward a gradual scale-up of its full capabilities. In this context, Graphical User Interfaces (GUIs) play a crucial role as tools to support the testing and commissioning of software components prior to their deployment in production. They provide developers with control panels and monitoring screens that deliver timely feedback during exploratory testing sessions, helping to identify unexpected code behavior. At the same time, they serve Assembly, Integration, and Verification (AIV) teams as collaborative platforms for synchronizing and validating system functions. Furthermore, the feedback from this early user experience will support the development of interfaces tailored for operators and maintainers, facilitating their work with the telescope during astronomical observations.

UI development has also proven to be an effective tool for communication between teams working on different software components, as it encourages the exchange of knowledge, discussions about interfaces, and clarifications of code behavior. In previous phases, collaborative GUI development was successfully adopted with limited telescope configurations, demonstrating its value for both technical validation and team coordination. [2, 3].

## THE CENTRAL SIGNAL PROCESSOR

At the core of the SKA lies a highly complex computing system designed to process and reduce, in real time, the immense flow of data generated by receptors. The Central Signal Processor (CSP) constitutes a sophisticated computing infrastructure within the project, integrating multiple software and hardware components to enable the first stage of raw antenna data reduction [4].

The main design of the CSP is shared between the two interferometers (Mid and Low) and comprises the following:

- CSP.LMC (Local Monitoring and Control): the CSP's control entry point, which interfaces with the Telescope Manager and orchestrates control over the entire system.
- CBF (Correlator and Beam Former): processes raw antenna signals into scientifically usable data streams.
- PSS (Pulsar Search): an engine dedicated to identifying new pulsars and transient signals.
- PST (Pulsar Timing): optimized for precise timing of known pulsars.

In this context, CBF, PSS, and PST represent the subsystems dedicated to data reduction and analysis, while CSP.LMC coordinates requests received from the higher-level Telescope Monitoring and Control (TMC) system and provides responses on behalf of the entire CSP.

A fundamental feature of SKA operations is the ability to partition telescope resources, such as antennas and computing hardware, into subarrays, allowing up to 16 simultaneous and independent observations or technical/engineering activities. To achieve this, a large portion of CSP resources must operate flawlessly in parallel, and any error or disruption must be detected and resolved as quickly as possible.

For these reasons, the design and implementation of GUIs must be regarded as a comprehensive and integral part of the scaling process toward the full implementation of CSP

* gianluca.marotta@inaf.it

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG013

# myLog: A MODERN, INTEGRATED LOGBOOK FOR SCIENTIFIC COLLABORATION AT EUROPEAN XFEL

L. Maia*, A. Anakkappalla, S. Aplin, D. Boukhelef, F. Dall'Antonia, I. Derevianko, L. Gelisio, D. Goeries, S. Hauf, J. Malka, M. Manetti, T. Michelat, G. Previtali, P. Schmidt, V. Singh, J. Szuba, W.-T. Tsai, B. Vanganuru, K. Wrona, †European X-Ray Free-Electron Laser, Schenefeld, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Scientific experiments at large-scale research facilities require flexible and collaborative tools to document, discuss, and track the progress of the experiment. We present *myLog*[1], a modern logbook solution developed at European XFEL following four years of iterative prototyping with users and staff of the facility.

Designed to address the evolving needs of scientists and support teams, *myLog* combines a user-friendly interface with a robust, modular architecture. It uses *Zulip* for real-time discussion-thread-based communication and leverages the facility's metadata catalogue, *myMdC*, to orchestrate workflows and deliver enriched interfaces. This tight integration enables seamless linkage between discussions, control system events, experiment metadata, and data analysis artifacts.

*myLog* emphasises user-driven organization of communication: experiment groups can configure information structure and notifications, while Principal Investigators maintain full access control. Real-time integration with the facility's metadata catalogue *myMdC*, control system *Karabo* and Data Analysis tool *DAMNIT*[2] enables the automatic logging of key events, complemented by manual entries as needed.

Data governance is enforced through embargo policies, ensuring secure handling of sensitive information. By merging communication, metadata, and automation in a single platform, *myLog* provides a scalable, future-proof approach to scientific logging and collaboration. This paper presents an overview of the solution architecture and its integration into the facility infrastructure.

## INTRODUCTION

Large-scale scientific facilities such as the European XFEL conduct highly complex experimental campaigns that involve diverse interdisciplinary teams. Their experiments often generate not only petabytes of raw data but also a wealth of contextual information like sample details, instrument settings, operator notes, and real-time observations and communication, which are essential for interpreting the results. Traditional logbooks capture some of this information, but not all, and often fail to provide the structure and traceability needed for complex, distributed workflows [1].

Replacing such a long-standing tool required more than a technical upgrade: it involved introducing new concepts and workflows while preserving continuity for users. The development of *myLog* [2] was therefore driven by the need for a unified solution that integrates logging, context-aware metadata, and automation within the facility's data management ecosystem.

In this paper, we present the design principles, integration strategies, and deployment of *myLog*, a solution built on Zulip [3] and extended by *myMdC* [4] to meet the specific requirements of experiment and operational logbooks at European XFEL.

## MOTIVATION AND REQUIREMENTS

The development of *myLog* was driven by several key factors:

- **Complexity of Experiments:** Experiments combine diverse diagnostics, detectors, and data sources with complex control infrastructures, all producing high-throughput data and metadata streams. Beyond raw data, configuration changes such as detector selection, parameter tuning, and adjustments to monitoring and control data sources occur frequently and must be logged with precision to maintain experiment integrity and reproducibility.
- **Collaboration Across Teams Before Beamtime:** Experiment preparation often begins months in advance and involves coordination between multiple groups, including instrument scientists, sample specialists, data and software experts, facility support teams, and user groups. Effective collaboration requires structured communication channels to capture discussions, decisions, and evolving configurations.
- **Collaboration During Beamtime:** Experiment shifts typically involve multiple groups working in parallel, 24/7, requiring collaborative continuous communication to maintain alignment of information.
- **Collaboration After Beamtime:** Post-experiment communication remains essential, especially for complex data analysis workflows that may take weeks or months.
- **Data Governance and Compliance:** Sensitive information as embargoed results and proprietary procedures, must be securely managed according to facility policies.
- **Integration with Existing Infrastructure:** The solution must integrate seamlessly with available data

* luis.maia@xfel.eu
† https://www.xfel.eu/

[1] available at https://mylog.connect.xfel.eu/

[2] Data And Metadata iNspection Interactive Thing

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMG014

# AN AGILE/XP SOFTWARE DEVELOPMENT PROCESS FOR MODERNIZING THE ACCELERATOR CONTROL SYSTEM AT FERMILAB

J. S. Diamond[†,1]
[1]Fermi National Accelerator Lab, Batavia, United States of America

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Fermilab is undergoing the most ambitious upgrade to its accelerator control system of the 21st century. As part of the ACORN project, hundreds of legacy control system applications written in C/C++ will be re-imagined and developed from the ground up. In addition, applications to support Fermilab's new super-conducting linear accelerator are already under construction. To manage the development of modern controls applications, the Controls department has adopted an Agile software development process based on eXtreme Programming. In this paper we will describe our process and detail our experience applying it to the development of two case studies.

## INTRODUCTION

Fermilab is the premiere particle accelerator facility in the United States hosting five particle accelerators spread out over 6800 acres. A sophisticated, distributed control system known as ACNET has been at the core of operations for this facility since the era of PDP-11s. Over the decades the control system computing hardware has migrated from mini-computers to modern day multi-core Intel-based servers, but an alarming amount of the original software remained unchanged.

To address this build-up of legacy software and ensure a sustainable controls system for the next generation of scientific discover at Fermilab, the ACORN project was initiated. ACORN is a DOE O413b project with an estimated cost of $160MM to modernize the accelerator control system. Roughly $30MM will be dedicated to modernizing the software components of control system – end user applications, central services, and data acquisition pipelines.

PIP-II is another DOE O413b project to build a new super-conducting linear accelerator to replace Fermilab's existing linear accelerator and boost proton intensities that will support next-generation neutrino experiments. PIP-II requires new, modern control system applications to support this advance

## THE CAVEMAN SOFTWARE DEVELOPMENT WORKFLOW

The existing software development workflow in the Controls department at Fermilab is like the Caveman workflow. If the reader is not aware, the Caveman workflow goes like this: the Software Developer receives a set of vague requirements from their customer, usually in oral form and then retreats to their office (the "cave") where they work tirelessly on the project for several weeks or months and emerge once the software is "complete" and ready for testing. At this point the real requirements can be gathered from the customer and (if the customer is lucky) the software developer begins the second phase of software development – twisting the code into a usable product. This phase continues until the customer is satisfied or the software developer is assigned a different task.

This type of software development workflow cannot scale to meet the volume of software that ACORN aspires to deliver. A more modern, scalable approach to developing software is necessary. Most software development teams use some form of Agile framework to build enterprise software.

However, one does not have to search far or long to discover that Agile is controversial in the software industry. Just go to TikTok or YouTube and search for Agile or Scrum if you are unaware of this controversy. Most of the negativity related to Agile software development is related to the number of meetings. Software developers dislike meetings, or rather the meetings that are more about collecting status and less about productivity. There is nothing in the Agile Manifesto about preferring meetings over writing code, so there is a disconnect between how some teams practice Agile and what being Agile really means.

## IS AGILE AN IMPROVEMENT

First recognize that Agile is not a software development process. The original signatories of the Agile Manifesto (cite) did not publish a book or even a "how to Agile" blog post. As most are aware by now, the Agile Manifesto is more of a philosophy for software development. Several software development frameworks that claim to be Agile do exist and it's these frameworks that are used to guide a team of software developers through completion of a software product. Common Agile frameworks in use today include Scrum, Scaled Agile Framework (SAFe) and eXtreme Programming.

One of the co-authors of the Agile Manifesto, Robert C. Martin, published a book on what he believes went wrong with the Agile movement in the 2010s. According to Martin, the overemphasis on the management practices (burndown charts and daily meetings) and the de-emphasis of the technical practices (test-driven development, pair programming and refactoring) results in the development of brittle code that's difficult to extend and maintain after reaching a critical mass of about thirty thousand lines of code.

## OUR AGILE PROCESS

Our Agile process leans heavily into the technical practices of eXtreme Programming – continuous integration and continuous delivery, refactoring, simple design, test-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR001

# SECURE INTEGRATION OF SLAC S3DF WITH EPICS CONTROL SYSTEMS FOR MACHINE LEARNING-DRIVEN ACCELERATOR OPTIMIZATION*

E. Williams†, J. Zhou
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

We present a Secure EPICS PVAccess (SPVA) deployment framework developed at SLAC to enable authenticated, encrypted, and authorized access to control systems from external scientific networks. In Phase 1, SPVA has been deployed to connect HPC clients and services on SLAC's Scientific External Network to internal PVAccess gateways supporting production accelerators.

SPVA enforces strong mutual authentication using Kerberos service principals, which establish the runtime identity of services and clients. These identities are used to request short-lived X.509 certificates from the SLAC-managed PVAccess Certificate Management Service (PVACMS). Certificates enable TLS-secured PVAccess communication, ensuring cryptographic trust between peers. Authorization is enforced through Access Security Files (ACFs) that define PVAccess security groups (ASGs) mapped to User Access Groups (UAGs) and Host Access Groups (HAGs) managed centrally in LDAPS, allowing fine-grained control.

This framework provides secure, traceable access to EPICS PVs across administrative domains while maintaining compatibility with PVXS-based IOCs and tools. This paper outlines the architectural design and operational lessons from the Phase 1 rollout, providing a model for deploying secure control system access in federated scientific computing environments.

## INTRODUCTION

SLAC's accelerator control system, built on EPICS, supports mission-critical operations for LCLS, LCLS-II, and XFEL. These facilities require precise, real-time control — often at rates up to 120 Hz and 1 MHz for beam-synchronous data — to achieve world-class scientific performance. Meeting these requirements demands a secure, scalable, and maintainable integration with SLAC's Shared Science Data Facility (S3DF), which provides GPU-accelerated HPC resources, containerized Kubernetes clusters, and high-performance storage through WEKA-NFS

## SCIENTIFIC AND TECHNICAL MOTIVATION

Modern accelerator science increasingly relies on data-driven methods such as machine learning (ML) to optimize beam performance, automate tuning, and support real-time fault detection. Traditional physics simulations, while accurate, are often computationally expensive and cannot always run in real time. Surrogate models — lightweight ML approximations — enable predictions in milliseconds, making them suitable for feedback loops during LCLS-II operations [1].

Virtual accelerators provide risk-free environments for testing algorithms by simulating EPICS PV behaviour. Digital twins extend this concept further by continuously synchronizing with live machine data (e.g., BPM trajectories, MPS rate/destination signals, and cavity parameters) to provide real-time predictive insight into machine performance [2]. The integration with S3DF enables training these models at scale using MLflow and MinIO for versioning and artifact management, while SLAC's high-performance Kubernetes infrastructure ensures flexible, containerized execution of ML workloads.

## SYSTEM ARCHITECTURE OVERVIEW

The integration design bridges S3DF and SLAC's EPICS-based control system through a demilitarized zone (DMZ) as shown in Fig.1. Key components include:

- EPICS PVA Gateways deployed in the DMZ for secure, TLS-enabled read/write access.
- CA Gateways restricted to read-only operation for legacy support.
- Kubernetes-native ML services in S3DF, orchestrated for scalability and resilience.
- Data pipelines linking control-system PV data with S3DF ML workflows and the ML Data Platform (MLDP).

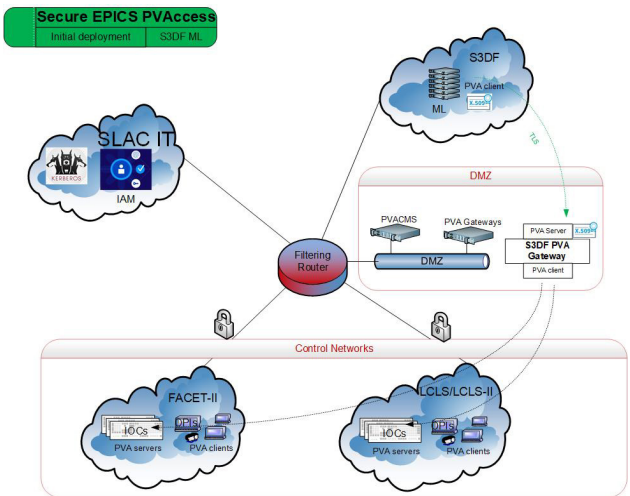



Figure 1: Secure Integration of S3DF ML with control systems via SPVA.

* Work supported by the U.S. Department of Energy, Office of Science
† ernesto@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR002

# hkl COMPUTATIONS AS AN EPICS IOC*

Alexander Baekey[1,2], Kazimierz J. Gofron†[1]
[1]Oak Ridge National Laboratory, Oak Ridge, TN, USA
[2]University of Central Florida, Orlando, FL, USA

*Abstract*

We present an Experimental Physics and Industrial Control System (EPICS) Input/Output Controller (IOC) that integrates an open-source `hkl` crystallographic library with EPICS through a sophisticated `PyDevice`-based architecture. This IOC enables real-time calculation of elastic scattering peak positions from instrument motor rotations, UB matrix transformations, and visualization within the control system environment. By embedding crystallographic knowledge directly into the IOC using GObject introspection, users can perform experiment-critical computations without external scripting or offline tools. The implementation leverages a layered architecture that combines EPICS process variables, Python bindings, and C library performance. We describe the design and use of the IOC, which supports multiple diffractometer geometries (2-circle, Eulerian and kappa 4-circle, 6-circle) along with addition of new instrument geometries and provides a powerful reference for single-crystal scattering experiments with enhanced real-time capabilities.

## INTRODUCTION

Modern X-ray and neutron scattering experiments increasingly require tight integration between crystallographic calculations and beamline control systems. Traditional workflows separate calculation from motion control and data acquisition, creating inefficiencies and potential synchronization issues.

Reference spectra are essential for matching peaks with diffraction planes, as the d-spacing collected from experiment is not unique to a crystallographic plane. A reference spectrum for a given sample can be expensive or difficult to obtain, and it may be non-trivial to convert motor rotations from one instrument to another. Quickly identifying diffraction planes from reflections enables users to conduct scattering experiments efficiently.

The `hkl` library [1], developed by Synchrotron SOLEIL, provides a modular C framework for mapping real to reciprocal space for various diffractometer geometries. This Input/Output Controller (IOC) [2] builds upon our initial work [3] and uses Python bindings within the `PyDevice` [4] framework to call functions from the `hkl` package directly into the Experimental Physics and Industrial Control System (EPICS) [5] control layer. The integration leverages GObject introspection to provide seamless C-to-Python API exposure, enabling high-performance crystallographic calculations within the EPICS framework.

The IOC maps real and reciprocal space, determines crystal alignment, and visualizes peak positions for multiple diffractometer geometries, with the ability to configure the IOC for new geometries and communicate with real-time experiment data. The IOC takes advantage of EPICS process variables (PVs) for communication to complement scattering experiments, providing a unified interface for both control and calculation operations.

## IOC DESIGN

The IOC is organized into five main functional tabs in CSS Phoebus [6], corresponding to the natural workflow of setting up and running a diffraction experiment. The architecture implements a layered design pattern, separating the EPICS control layer, the `PyDevice` integration bridge, and the crystallographic computation layer.

### *Initials*

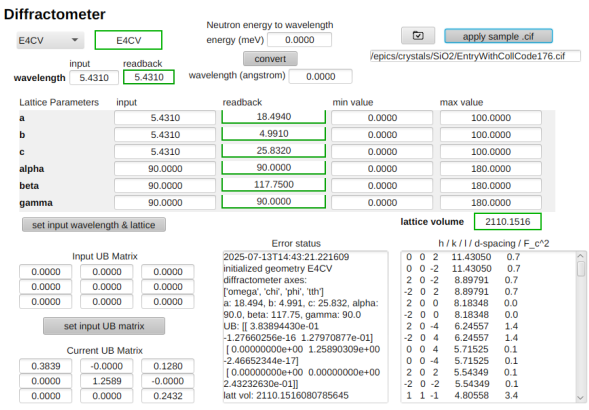



Figure 1: Initials tab.

This tab (Fig. 1) provides the entry point for the experiment configuration. Users set the diffractometer geometry, beam wavelength, sample lattice parameters, and sample orientation if known. The interface communicates with the

* This manuscript was authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).
This research was supported in part by an appointment to the Oak Ridge National Laboratory Research Student Internships Program or Technical and Professional Internships Program, sponsored by the U.S. Department of Energy and administered by the Oak Ridge Institute for Science and Education. (RSI/TPI).
† gofronkj@ornl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR003

# THE STATUS OF HEPS BEAMLINE CONTROL SYSTEM*

G. Li†, AI.Y. Zhou, Y. Liu, Z.Y. Yue, CH.X. Yin, X.B. Deng, Q. Zhang, D.SH. Zhang, ZH.H. Gao, ZH. Zhao, N. Xie, G. Lei†, Institute of High Energy Physics, CAS, Beijing, China

*Abstract*

HEPS (High Energy Photon Source) is the first fourth-generation high-energy (6 GeV) synchrotron radiation facility in China, which is mainly composed of accelerator, beamlines and end-stations. In the phase I, "14+1" beamlines and corresponding experimental stations have been constructed. The beamline control system is responsible for the high-precision automated control, status monitoring, and data acquisition of all optical components, diagnostic devices, and sample stages on the beamline, ensuring the delivery of stable, reliable, and high-performance beam to the end-stations. The EPICS-based control system has been designed and built, with commissioning now underway. Here, the status of beamline control system is presented.

## INTRODUCTION

The High Energy Photon Source (HEPS) [1] is the first high energy diffraction-limited storage ring (DLSR) light source in China (Fig. 1). HEPS will serve as a cutting-edge platform for multidisciplinary research in fields such as materials science, condensed matter physics, chemistry, life sciences, and engineering.

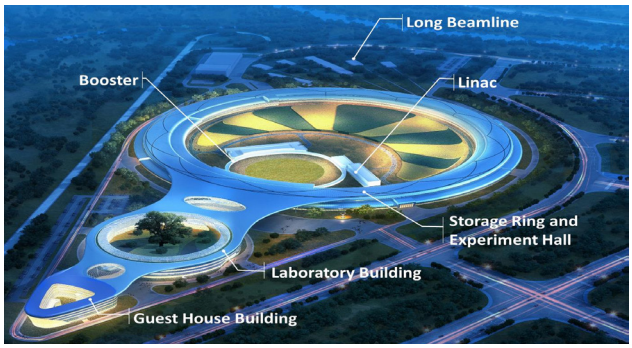



Figure 1: A bird's eye view of HEPS.

To satisfy some major demands of the nation and to further cutting-edge scientific development, the first phase will deliver 15 beamlines, including 14 for user experiments and one for testing purposes [2] (Fig. 2).

In order to simultaneously complete the construction of 15 beamlines control system on schedule, a set of standards is developed to design and build control systems for beamlines, minimizing heterogeneousness, to ensure that the control system has good stability, reliability, flexibility, availability, reliability etc.

* Work supported by HEPS project

†li75gang@ihep.ac.cn; leige@ihep.ac.cn

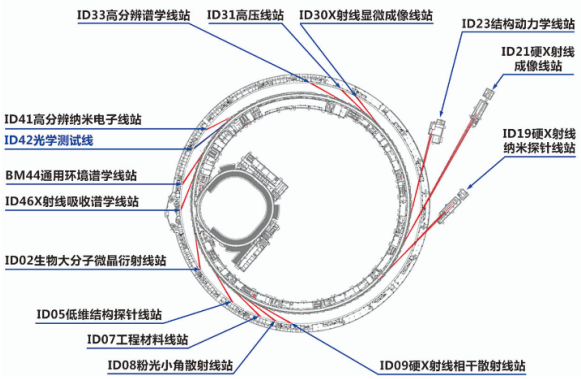



Figure 2: Beamlines layout in HEPS Phase I.

## BEAMLINE CONTROL SYSTEMS

EPICS is an open-source software, which is widely used worldwide to create distributed soft real-time control systems. It has also a large collaborative and active community. EPICS and its application components are designed and developed modularly.

The beamline control system is critical for ensuring stable, reliable, and precise beam delivery to experimental end-stations. It integrates heterogeneous hardware and software components into a unified architecture, based on the EPICS [3] framework (Fig. 3).

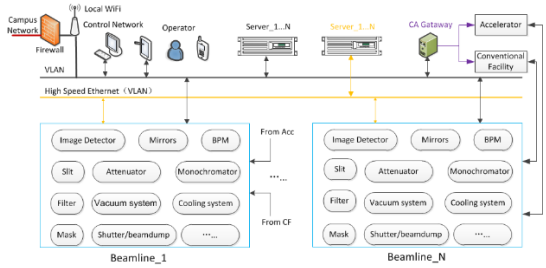



Figure 3: Overall structure of beamline control system.

The HEPS beamline control system adopts a typical three-tier architecture:

Operator Layer: Located in the experimental station or control room, it provides users with a graphical user interface (GUI).

Control Layer: Composed of high-performance servers and IOC (Input/Output Controller) servers, it forms the core of the EPICS architecture, running databases and control logic.

Device Layer: Includes all hardware devices and their controllers (motion controllers, PLCs, Intelligent I/O Module, digital multimeters, various point/line/area detector controllers, etc.)

This paper presents the current status of HEPS beamline control system.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR004

# LIQUID SAMPLE DELIVERY CONTROLS AT LCLS*

J. Zamudio Estrada†, J. Joshi, J. Cohen, C. Hampton, S. Dehe, R. Sierra,
SLAC National Accelerator Laboratory, Menlo Park, CA, United States

*Abstract*

The Linac Coherent Light Source (LCLS) at SLAC National Accelerator Laboratory is a pioneering X-ray free-electron laser that provides researchers with the ability to investigate matter at atomic and molecular scales with unprecedented temporal resolution. Its applications span a wide range of scientific disciplines, including materials science, chemistry, biology, and physics. A vital aspect of conducting successful experiments is the precise delivery of samples into the X-ray beam. Depending on the nature of the sample—whether liquid, gas, or solid—various delivery systems are employed to ensure accurate positioning, high repetition rates, and minimal sample waste. This paper will present an overview of the control system developed to support liquid sample delivery with focus on two advanced systems that have significantly enhanced experimental capabilities. The first is a recirculating liquid sheet jet system that enables the generation of tunable liquid sheets with minimal sample volume, making it ideal for experiments with limited material availability. The second is a Droplet-on-Demand (DoD) robot designed for high-throughput pump–probe experiments. This system allows precise sample placement, low sample consumption, and efficient mixing, which are essential for time-resolved measurements.

## INTRODUCTION

In LCLS II, most if not all control systems use Beckhoff Programmable Logic Controllers (PLC) for interfacing with 'dumb' devices and motors without controllers. The Beckhoff PLCs are programmed using TwinCAT 3 software [1]. The control systems are designed to be modular and scalable, allowing for easy integration of new components and functionalities as needed. Using Pytmc [2], a Python package developed at SLAC National Laboratory, we can generate Experimental Physics and Industrial Control System (EPICS) process databases records from TwinCAT .tmc files streamlining the twincat to EPICS workflow.

To allow users to interact with the control systems, Experiment Control Systems (ECS) group, developed custom graphical user interfaces (GUIs) using PyQt. ECS has created a collection of widgets and ophyd [3] device classes written in Python. These packages are called *pcdswidgets* [4] and *pcdsdevices* [5] respectively. The ophyd device classes group PVs into logical units that represent physical devices. Pcdswidgets assigns a user interface (UI) elements to a process variable (PV). These libraries provide a consistent interface as well as a consistent look and feel across GUIs in LCLS hutches. For template and reusable GUI components, ECS has developed *Typhos* [6], a tool for generating control screens from ophyd device classes.

Not all hardware devices use Beckhoff PLCs. For devices that have common interfaces like TCP/IP or serial, we can develop or configure an EPICS Input/Output Controller application. Then create custom screens using the PyQt.

## CHEMRIXS SAMPLE DELIVERY SYSTEM

In the Chemical Resonant Inelastic X-ray Scattering (chemRIX) instrument, we have developed a sample delivery system that allows for the generation of tunable liquid sheets with minimal sample volume requirements. This system is particularly useful for experiments where the sample is scarce or expensive, as it minimizes waste while still providing a stable and reliable delivery method.

The chemRIXS sample delivery system consists of two main parts. A sample injector, and a sample recirculation system. The sample injection system utilizes a load-lock chamber to allow for rapid exchange of microfluidic injector tips. The recirculation system utilizes a "catcher" that collects the sample into a storage system beneath the main chamber for re-circulation and re-use in the experiment, see Fig. 1 for a overview of the sample delivery system. Both the sample injector and the recirculation system are motorized, allowing for spatial overlap between the sample jet and the XFEL, as well as the catcher and the jet.

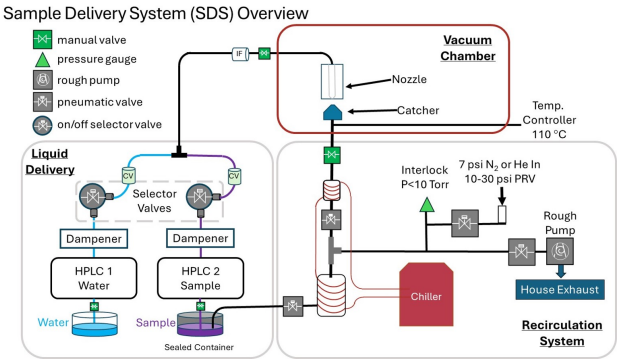



Figure 1: Schematic of the chemRIXS sample delivery system.

### *Vacuum Chamber*

In the vacuum chamber, we have a motorized sample injector that allows for precise positioning and tuning of the liquid sheet jet. There are many opportunities for the injector to bump or crash into other components in the sample chamber, so we have implemented an equipment protection system (EPS) to limit the motion of the injector and prevent damage to the equipment. To see into the chamber a

* Work supported by U.S. D.O.E. Contract DE-AC02-76SF00515.
† jozamudi@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR005

# THE LCLS-II MODULAR OPTICAL DELIVERY SYSTEM: LESSONS LEARNED*

A. J. Berges†, T. Johnson, C. Pino, M. Britton, J. Joshi, B. Fishler,
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

The LCLS-II optical delivery system supports multiple interaction points across multiple experiment hutches using only a handful of laser sources. This reduces financial burden and space usage at the cost of increased complexity for the optical laser systems. To ameliorate this complexity, each interaction point is supplied with a Modular Optical Delivery System (MODS) to inject, shape, and compress the beam before it is further conditioned for the experimental use. To meet operational demands, these MODS must be highly configurable, flexible, and robust while supporting 140+ control points in a dense enclosure. With control points spanning piezoelectric motors, optical imaging, digitizers, and more, the EPICS control system framework simplifies driver maintenance and allows growth of community-driven solutions. Each control point is accessible remotely via PyDM GUIs which enables the operator to control these various alignment and diagnostic tools. Managing the deployment and operational stability of these modular systems is nontrivial and has presented several challenges in recent LCLS experiments that inspired significant design changes for the future of the MODS.

## INTRODUCTION

The Linac Coherent Light Source superconducting accelerator (LCLS-II) takes the next step in ultrabright, ultrafast x-ray generation, bringing with it demanding requirements for optical photon delivery. To meet the accelerator's maximum repetition rate of ≈ 1 MHz, a hutch's optical laser source needs to support dynamic changes to power, pulse width, wavelength, and more to meet experimental requirements. LCLS-II simultaneously introduces more interaction points per hutch, increasing the density and complexity of the experimental design while enabling cutting-edge science. Centralizing laser sources provides an attractive solution to this problem, LCLS-II's Near Experimental Hall (NEH), shown in Fig. 1, supports a parallel delivery system, providing uncompressed, high-power laser pulses to up to 4 hutches simultaneously. This leaves the hutch responsible for conditioning the laser source to meet the experiment's demand, precipitating a need for a compact, remotely controllable optical delivery system.

## THE CURRENT GENERATION OF MODS

LCLS-II fulfills this need with the Modular Optical Delivery System (MODS), where the incoming uncompressed optical pulse is tailored to the experiment's science requirements at the hutch's interaction point. Each core step of this tailoring is divided into discrete Table-top Integrated Laser Enclosures (TILEs), where the pulse is injected (INJ), compressed (COM), tracked with an Arrival Time Monitor (ATM), and ejected (EJX) to the interaction point. The pulse may be optionally modulated with non-linear TILEs, utilizing harmonic generation (HRM) and optical parametric amplification (OPA), before it is ejected. In recent runs, the MODS enjoyed success in facilitating laser delivery and remote alignment while also experiencing pitfalls borne from design constraints. This paper explores the shortcomings of the current MODS design, incorporating feedback from operations and proposing an opportunity for improving the system.

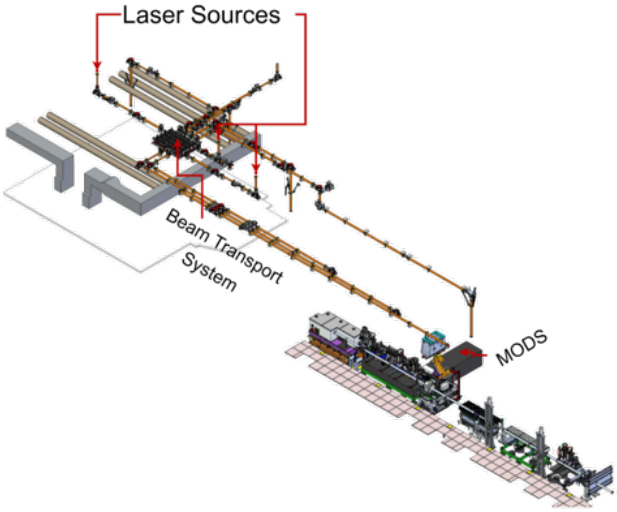



Figure 1: Overview of the LCLS-II Near Experimental Hall laser delivery system. Laser sources on their respective tables are routed to the Beam Transport System (BTS), where the beam is redirected to the corresponding hutch's MODS.

While science initiatives spark the inclusion of additional interaction points in LCLS hutches, control and optical density increase dramatically to meet evolving demands. These spatial constraints necessitate a compact design for optical alignment and delivery. A typical deployment of the MODS can be seen in Fig. 2 below, where the uncompressed pulse enters the INJ TILE from above from the Beam Transport System (BTS). Coordinating beam pointing across the BTS and through the MODS requires a myriad of diagnostic and motion systems to successfully deliver the laser pulse. In just this interaction point alone, there are nearly 150 different control points across the general categories of motion, char-

* This work is supported by Department of Energy contract DE-AC02-76SF00515

† aberges@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR008

# EPICS IN PRACTICE AT LCLS*

K. Malapati†, Z. Lentz, SLAC National Accelerator Laboratory, Menlo Park, CA, United States

*Abstract*

At LCLS (Linac Coherent Light Source), EPICS (Experimental Physics and Industrial Control System) plays a central role in our controls architecture. Input-Output Controllers (IOCs) are used to interface directly or indirectly with almost all experimental hardware supporting our heterogeneous requirements. EPICS network protocols are used to make devices available over the network, data acquisition, and device security and safety. These ultimately enable a rich environment of control tools built around standardized communication protocols such as Channel Access (CA) and pvAccess, such as alarm systems, software interlocks, and data analysis. This paper will detail how EPICS is used in LCLS and the tools built around or on top of it with specific examples and applications used on a day-to-day basis to accomplish basic and advanced needs of a complex controls system. This includes tools developed at LCLS or SLAC specifically to fulfill general needs that other users may be able to take advantage of. This paper will also describe what EPICS currently does well along with its limitations.

## INTRODUCTION

In LCLS, there is an essential need for a distributed control system capable of supporting a wide variety of hardware and operations. To do this, we use IOCs and EPICS network protocols [1]. IOCs are programs that maintain databases of process variables (PVs) and are able to interact with hardware, either locally or via network protocols like TCP/IP. PVs hosted by IOCs can be accessed with Channel and PV Access, EPICS-specific protocols with shell tools for manual use, and bindings in popular programming languages such as C++ and Python for programmatic or graphical usage [2]. With just these two EPICS-based capabilities, LCLS is able to implement the core of our control system.

IOCs distributed among several servers on different subnetworks and for different instruments perform local I/O with a wide variety of hardware - motion systems, detectors, environmental and vacuum systems, and more. IOCs are extensible and can be augmented by EPICS modules to support new hardware or add new capabilities to existing support [3]. They also provide a layer of abstraction to simplify management of hardware; cameras from different vendors may communicate with IOCs in different ways but can be controlled similarly from a user's point of view despite design differences. No matter how IOCs are implemented or what they control, they almost invariably expose access to hardware indirectly via PVs.

PVs range in their design and complexity, from simple records that are only used to read and write scalar values to records capable of managing more than a hundred different motion-related fields for a positioner [4]. The most direct way to interact with an IOC is via the EPICS IOC shell, which is a simple CLI with a variety of commands to manipulate PVs [5,6]. This may be sufficient for minor uses such as light debugging, but in practice it is unnecessary and unscalable - this is because of EPICS-provided network protocols like Channel Access (CA).

The implementation details of CA (or its successor, pvAccess) are of secondary importance to its simplicity - all PVs from all IOCs, and by extension all the hardware it is linked to, can be controlled in the same way using this protocol. Any tool built with or around CA can immediately augment the capabilities of all devices with IOCs. Every new IOC made will immediately benefit from all CA-based software that already exists for free, excluding any costs associated with developing the IOC which EPICS and its modules generally facilitate. The rest of this paper will detail the specific ways EPICS is used in LCLS.

## IOC INFRASTRUCTURE

Once built, an individual IOC is a program that can be easily run in any terminal; however, there is a desire to formalize the management of IOCs, such as deployment and maintenance. This is in due in part to the scale of devices to be controlled at LCLS - there are well over a thousand IOCs, many of which may control multiple pieces of hardware (multiple gauges and pumps, a controller with several motors, etc.), but there are other considerations such as maximizing uptime without manual intervention and preventing changes to existing IOCs from unauthorized users.

Our first layer of convenience is procServ, a wrapper program used to run IOCs in the background while providing a telnet port to interact with the IOC and automatic logging [7]. Having IOCs run in procServ sessions makes it trivially easy to stop, start, and access the shell of an IOC with just the procServ port and host the IOC is running on; it also allows for CLI tools that streamline interaction with IOCs such as updating versions by interaction with the relevant procServ sessions.

That being said, procServ sessions do not start themselves - for that we rely on a chain of services started at boot as illustrated in Fig. 1. Instruments at LCLS own configuration files with sufficient information about each IOC to start procServ sessions for each one at boot time or interact with them using a human-readable name mapped to a procServ ID. These configurations are typically only updated using a user-friendly GUI called IOC Manager [8]. This GUI uses PyDM, an extension of QT developed by another group at SLAC that integrates PVs into GUI elements using CA [9]. IOC Manager provides capabilities such as a display of a heartbeat PV for each IOC, a basic indicator of liveness, or

* This work was supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-76SF00515.

† kaushikm@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR010

# SIGNAL RESPONSE AND ANALYSIS OF LARGE MICRO CHANNEL PLATE DRIVEN DELAY LINE DETECTORS

P. Sjöblom*, V. Ekholm, L. Kjellsson, O. Grimm
MAX IV Laboratory, Lund University, Lund, Sweden

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

For soft X-ray spectroscopy beamlines, delay line detectors are often the main system for detecting the photons from the sample and hence also a component determining the overall beamline performance as it might be a limiting factor of both measurement speed, noise, artifacts, and resolution. As such, and even more with larger micro channel plate driven delay line detectors, the signal readout must be fast and robust to minimize noise and artifacts while still accommodating even the flux from 4th generation synchrotrons. This paper studies the signal response of a delay line detector and how the ns current signal pulses can be filtered, amplified, and converted to voltage before the digitization. The digitizer is a 12 bit 2.5 GSPS 6 channel system, which is set up in a manner to minimize noise and enable post signal analysis integrated into the Sardana control system and live view. The early results indicate that many of the currently present image artifacts are, to a very high degree, suppressed due to analog signal treatment and proper triggering. The digitized signals are fitted using the python tool LMFIT to different signal models, such as the exponentially modified Gaussian, to extract the peak of the main signal after identifying the common background response in all channels with the aim to even further improve the resolution of the detector. To optimize sampling, the system is also stress tested with regards to e.g. sampling length and out of range measurement.

## INTRODUCTION

A common type of soft X-ray detector is the delay line detector (DLD) with a micro channel plate (MCP) placed in front of the DLD, see Fig. 1. Photons leaving the sample under study are allowed to, in a focused and dispersive manner, impinge on the MCP, where free, but few for every photon, electrons are generated. Through a few hundred volt difference between the top and the bottom of the MCP, the free electrons accelerating through the holes create an avalanche of electrons. In the case of Rowland spectrometers where the incoming light on the detector has the same gracing incidence angle as the gracing outgoing angle from the grating, a repeller plate can and is often placed closely above the MCP with additional voltage to push (repel) electrons down and back on the MCP and into the holes, drastically increasing the event rate. [1] As the avalanche of electrons leaves the MCP channels, it continues to accelerate in a trajectory towards the meandered lines of the DLD. Each cloud of electrons induces a pulse in the DLD that starts to travel to both ends of the line. Pulses arriving simultaneously at both ends indicate a pulse with equal distance to travel, while the time difference is proportional to the distance from the center. As the line is meandered and the time traveling orthogonal to the physical coordinate being measured, time is binned into channels. With two delay lines rotated 90°, an (X,Y) coordinate is possible to reconstruct corresponding to the position of the incoming photon. As the photons are dispersed by a grating, position in the y-direction corresponds to the energy of the light.

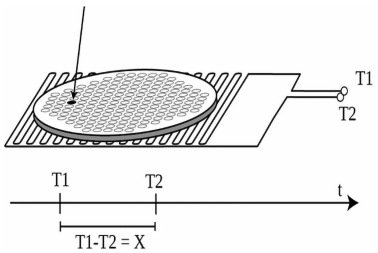
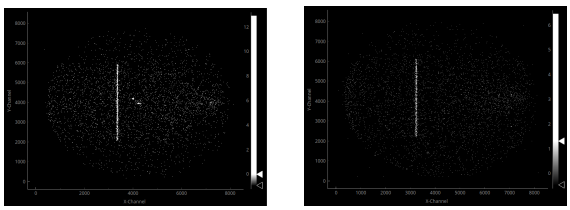



Figure 1: The electron cloud from the MCP passing through the DLD meanders will induce a two pulses, one in each direction. The path length differences between the two pulses to the end will translate into a time difference that corresponds to the x-coordinate of the event.



(a) (b)

Figure 2: The energy from the spectrometer is imaged as a line from the DLD detector and the narrower the line the better energy resolution the system has. Also of importance is that artifacts, such as visible in a) are suppressed during readout, as in b).

The characteristics of the detector must match the performance of the rest of the system, which is, in this paper, the performance of the Veritas beamline at MAX IV described in [2–5].

The standard for many years has been based on a four channel time to digital converter (TDC) that creates sets of (X,Y) coordinates for every photon hitting the MCP. The TDC receives two signals from the line determining the X-position, X1 and X2, and two from the Y-line, Y1 and Y2,

* peter.sjoblom@maxiv.lu.se

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THMR013

# INTEGRATED CONTROL SYSTEMS FOR TIME-RESOLVED RIXS AT LCLS-II: DESIGN AND OPERATIONAL CHALLENGES*

J. Joshi†, P. Oppermann, G. Dakovski, K. Kunnus, Q. Yougoubare
SLAC National Accelerator Laboratory, Menlo Park, USA

*Abstract*

The newly enhanced LCLS-II X-ray laser at SLAC National Accelerator Laboratory represents a major advancement in X-ray science, providing unprecedented capabilities for probing ultra-fast dynamics in chemistry, materials science, biology, and beyond. Among the new beamlines, the Resonant Inelastic X-ray Scattering (RIXS) beamline leverages the high repetition rate of LCLS-II to investigate the energy distribution and evolution of occupied and unoccupied molecular orbitals in complex and catalytic systems, particularly in liquid environments. This beamline features two dedicated endstations, qRIXS (upstream) and chemRIXS (downstream), each optimized for distinct scientific goals. This paper will detail the design and implementation of the experimental controls and data systems that unify beamline hardware and instrument automation. In addition, this paper will discuss the challenges of synchronizing operations across two endstations on a single beamline for time-resolved spectroscopy under demanding experimental conditions.

## INTRODUCTION

The LCLS-II (Linac Coherent Light Source II) X-ray laser facility is built on the legacy of LCLS, and is capable of delivering X-ray pulses at repetition rates up to a million per second to enable experiments that were previously limited by photon flux, data throughput, or temporal resolution (see Fig. 1). This capability opens unique opportunities to probe structural and electronic dynamics in real time, advancing our understanding of chemical reactions, energy materials, and biological processes at the most fundamental levels. Among the flagship scientific instruments of LCLS-II, the Resonant Inelastic X-ray Scattering (RIXS) beamline is designed to leverage high-repetition-rate operation for the study of electronic excitations and orbital dynamics in complex systems. To maximize its scientific reach, the RIXS beamline incorporates two specialized endstations, qRIXS [1] for momentum-resolved spectroscopy and chemRIXS [2] for solution chemistry, both sharing a common photon delivery system. The complexity of this dual endstation design presents new challenges for experimental controls, automation, and data handling, which must ensure seamless operation under demanding experimental conditions.

* Work supported by U.S. D.O.E. Contract DE-AC02-76SF00515.
† jjoshi@slac.stanford.edu

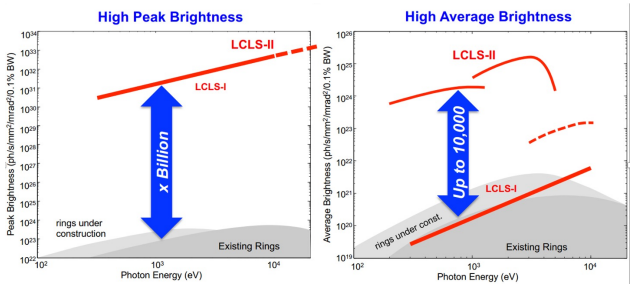



Figure 1: Plots showing Peak Brightness and Average Brightness comparison between LCLS and LCLS-II.

## SCIENTIFIC AND TECHNICAL BACKGROUND

Time-resolved Resonant Inelastic X-ray Scattering (tr-RIXS) [3], which is very photon-hungry, is the technique to observe the electronic excitations of the material, like charge, spin, or lattice-dynamics, change and evolve in real time after being excited by an ultrafast X-ray laser pulse. By using ultrashort X-ray laser pulses that act as a probe at precisely controlled delays after the initial optical pump, tr-RIXS allows one to capture snapshots of the transient states, revealing the dynamics of a system on extremely short timescales, such as femtoseconds.

The high repetition rate of LCLS-II, reaching up to a megahertz, unlocks new opportunities for RIXS experiments that were previously impractical at X-ray free-electron lasers. To serve this broad scientific program, the LCLS-II RIXS beamline was designed with two dedicated endstations, supported by a common photon delivery system. The upstream qRIXS endstation is optimized for momentum-resolved measurements, and the downstream chemRIXS endstation is tailored for liquid jets, enabling time-resolved studies of solution chemistry and heterogeneous catalysis.

The beamline layout, as shown in Fig. 2, reflects this dual-purpose mission. Monochromators and focusing optics deliver tunable, high-flux X-rays to the first interaction point, with beam conditioning and diagnostics ensuring stability before the photons reach the downstream endstation. Sharing a single beamline across two sophisticated instruments introduces unique technical challenges, including precise photon energy control and beam steering. In addition, experimental controls and data systems must support seamless switching between endstations, unify hardware interfaces, and handle high-rate data streams without compromising reliability.

Together, the scientific drivers and technical design considerations outlined here form the foundation of the RIXS

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD001

# IMPROVING CYBERSECURITY ON THE LASER MEGAJOULE FACILITY

J.P. Airiau, J. Hannier, B. Nuhrich, S. Palmier, V. Raffin, N. Roux
CEA CESTA, Le Barp, France

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Laser MegaJoule (LMJ), a 176-beam laser facility developed by CEA, is located at the CESTA site near Bordeaux [1]. The LMJ facility is part of the French Simulation Program, which combines improvement of theoretical models and data used in various domains of physics, high performance numerical simulations and experimental validation. It is designed to deliver about 1.4 MJ of energy on targets, for high energy density physics experiments, including fusion experiments.

The LMJ technological choices were validated on the LIL, a scale-1 prototype composed of 1 bundle of 4-beams. The first bundle of 8-beams was commissioned in October 2014 with the implementation of the first experiment on the LMJ facility. The operational capabilities are increasing gradually every year until the full completion by 2026. By the end of 2025, 22 bundles of 8-beams will be assembled (full scale) and 19 bundles are expected to be fully operational.

The computer systems are facing many threats with increasingly sophisticated attacks to break through the defenses in place. In addition to existing measures on the LMJ facility, an action plan has been implemented to improve our protection against cyber threats. This article provides an overview of the different measures considered and the impact of their implementation on existing systems.

## INTRODUCTION

The LMJ is designed to perform indirect drive inertial confinement fusion (ICF) experiments on deuterium tritium targets. Aside this major objective, a broad spectrum of experiments can also be addressed from high energy density physics to basic science [2]. This keystone of the Simulation Program combines physics models improvement, high performance numerical simulation, and experimental validation, in order to guarantee the safety and the reliability of French deterrent weapons.

This facility is dimensioned to accommodate 176 beams grouped into 22 bundles of 8 beams. These beams are located in the four laser bays arranged on both sides of the central target bay of 60 meters diameter and 40 meters height. The target chamber and the associated equipment are located in the center of the target bay (see Fig. 1).

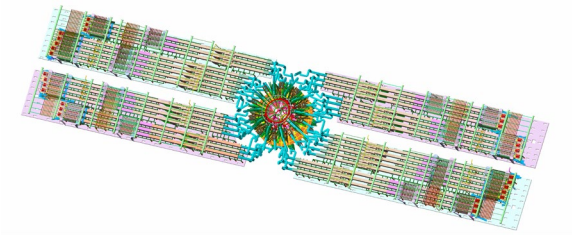


Figure 1: Laser bays and target chamber.

The 176 beams are grouped into 22 bundles of 8 beams. The bundles are distributed evenly across the laser bays (5 to 6 per bay). A laser bundle is divided into four main sub-assemblies (see Fig. 2), the front end and the amplification section located in laser bay and the transport and frequency conversion & focusing system in the target chamber [3].

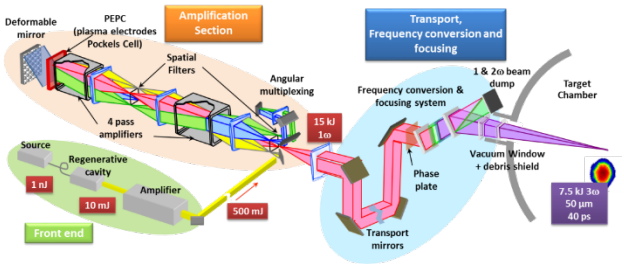



Figure 2: LMJ laser bundle.

Target Diagnostics (TD) are a key tool for numerous physical data acquisition during LMJ shots.

There are actually 22 operational TD [3]. They can be mounted around the target chamber depending on the experiment requirements. Some of them are positioned permanently around the target chamber, while others are mounted as needed (see Fig. 3). This kind of TD are inserted in the target chamber using a System for Diagnostics Insertion (SID).

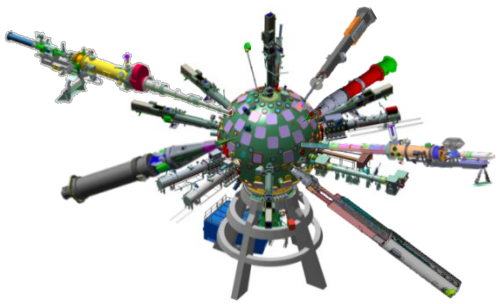


Figure 3: Target diagnostics on the target chamber.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD002

# A DUAL-NETWORK CENTRALIZED CONTROL SYSTEM FOR ELETTRA 2.0

L. Pivetta*, A.I. Bogani, G. Gaio, C. Scafuri, G. Scalamera, I. Trovarelli
Elettra-Sincrotrone Trieste S.C.p.A., Basovizza, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Elettra 2.0 storage ring, scheduled for commissioning in 2026, introduces a novel control system architecture, departing from distributed front-end computers. High processing power and intelligent devices, such as magnet power converters, beam position monitors, beam loss monitors, low-level RF systems, and fast interlocks, support centralization. These devices connect through a dual-network design: a standard Ethernet link, managed by virtual servers running Tango, handles supervision, while a high-speed fibre optic link with custom protocols enables real-time control. Fast interfaces feed a centralized, multi-core real-time (RT) server, where hardware partitioning and the Data Plane Development Kit (DPDK) allow low-latency processing up to 1.1 MHz. Conceptually, each core mimics a front-end computer, while the server RAM acts as a high-speed communication bus minimizing delay and latency. This architecture simplifies the infrastructure, improves scalability and enables advanced machine-wide control, ensuring reliability for next generation accelerators.

## REQUIREMENTS

A conventional control system, based on the TANGO Controls framework, is in charge of interfacing and integrating all the devices and equipment needed for normal operation in an homogeneous way. Sampling periods in the order of 1 second, down to 100 ms are foreseen, depending on the equipment. Similarly, typical timing for command execution or actuator setting is in the order of 100 ms, down to 1 ms. Hard real-time behavior is not required. These performance numbers can be achieved with a conventional network and standard servers deployed as virtual machines or containers.

Novel generation beam position monitors (BPM), low level radio-frequency systems (LLRF), beam loss monitors (BLM), equipped with 1 Gb or 10 Gb ethernet interfaces, are capable of very high data rates. The hard real-time control system has been designed to be able to acquire all the sensors, streaming continuously without trigger, perform the necessary calculations, and drive the actuators, e.g. the power converters of the magnets. Additionally, the global orbit correction system will run at 50 kHz repetition rate, or higher. For specific diagnostic purposes, the system will acquire some BPMs at the turn-by-turn revolution frequency of the stored electron beam which, for Elettra 2.0, will be 1.15 MHz and mark each sample with a timestamp/turn number, provided by the timing system, allowing on the fly correlation of data from different sensors.

## REAL-TIME PROCESSING

The pursuit of ultra-low latency in control systems has driven a shift toward centralized architectures, where large computational resources and high-bandwidth networks converge to handle time-critical tasks with unprecedented performance. Modern server motherboards and multi core processors offer options and features designed to minimize latency an maximize determinism, in addition to high bandwidth. These platforms, used in conjunction with hardware partitioning and operating system kernel bypass techniques provided by DPDK [1], can deliver deterministic performance, rivaling dedicated hardware like FPGAs, while offering greater flexibility for algorithm development and system scalability. Special care has been taken analyzing and defining hardware partitioning and cpu isolation; this guarantees that network adapters are installed in specific PCIe slots, corresponding to the desired NUMA node. SIMD microprocessor instructions, such as AVX2 and AVX512, provide advantage in computing the control and feedback algorithms.

To make all the integration work easily available, Voltumna Linux [2], a custom distribution based on Yocto/OpenEmbedded, that provides pre-built operating system images higly optimized for each specific hardware platform, have been synthesized at Elettra.

As anticipated, this in-memory control paradigm is the core of the hard real-time control system for Elettra 2.0. All critical data, from diagnostics to actuator setpoints, reside in RAM, effectively transforming the server memory into a virtual high speed bus. This approach minimizes communication overheads, enabling seamless correlations across subsystems and supporting advanced features like real-time digital twins and predictive analytics. Sophisticated hashing-based APIs for direct read/write access to the shared memory allow for very low latency for intra-processor transfers, fostering an environment where complex control loops can operate at high frequencies without the rigidity of hardware-centric designs [3]. Central to this approach is DPDK, integrated to facilitate poll-mode driver operation and user-space networking that eliminates kernel induced jitter [4]. A test facility is in operation to validate the architecture and measure the actual performance of this solution [5].

Scaling to the full Elettra 2.0 deployment will be based on a multi-socket server, featuring a Supermicro H14DSH motherboard [6] with two AMD EPYC 9755

* lorenzo.pivetta@elettra.eu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD003

# MANAGEMENT OF SERVER AND NETWORK INFRASTRUCTURE AT SuperKEKB

H. Sugimura*, F. Ito, A. Morita, T. T. Nakamura, S. Sasaki
High Energy Accelerator Research Organization (KEK), Tsukuba, Japan

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The SuperKEKB accelerator employs an EPICS-based control system for its operation. This paper presents the design and implementation of the core server infrastructure and network environment that support the accelerator control system. The servers are provisioned and managed using Ansible, enabling infrastructure-as-code practices that ensure idempotency, reproducibility, and maintainability across the system.

We describe the configuration of physical and virtual servers, the monitoring infrastructure based on Zabbix, iDRAC, and the Elastic Stack, as well as our operational experiences, including the smooth OS migration from CentOS 7 to AlmaLinux 9. The lessons learned and the benefits of code-based infrastructure management are also discussed.

## INTRODUCTION

At the SuperKEKB accelerator facility, documentation related to computing and networking infrastructure has traditionally been maintained in the form of internal technical notes. These records, inherited from the preceding KEKB accelerator project, have supported the continuity of technical knowledge despite personnel changes over time. In particular, the dedicated accelerator control network required systematic management of configuration files for machines running core services such as DNS, DHCP, and NTP. These configuration files have been carefully preserved across operating system upgrades, allowing seamless continuity in system administration.

In addition to these core service servers, various user-facing machines have been provided to support the development of accelerator-related software tools. These systems required the installation and maintenance of application environments, which have been updated as needed in accordance with changes in the underlying operating systems.

In recent years, however, modern technologies have emerged that dramatically improve the maintainability and reproducibility of infrastructure. One such advancement is Infrastructure as Code (IaC), which enables system configurations to be described programmatically. By adopting this approach, server setup and configuration tasks can be performed consistently by any administrator, thereby reducing reliance on specific individuals. Ansible, a representative tool for IaC, has been introduced in our environment to automate and accelerate server provisioning processes [1]. Its adoption has also proven beneficial in disaster recovery scenarios, allowing rapid reinstallation and configuration of servers from scratch following hardware failures. At SuperKEKB, Ansible-based automation has been employed to deploy and manage a variety of servers essential to the control system.

In the past, server monitoring was often performed by manually saving and searching command-line output logs to determine who had executed which operations on which machines. This method has since been replaced by a more modern and robust monitoring framework using Zabbix, which stores monitoring data in a structured database and provides intuitive visualization [2]. Zabbix is used primarily for OS-level monitoring, while hardware-level monitoring is achieved using tools provided by the server vendor. In our system, all servers are standardized on Dell hardware, and monitoring is performed using iDRAC (Integrated Dell Remote Access Controller), which enables centralized management and hardware status tracking.

For network monitoring, we have adopted the Elastic Stack to collect, store, and analyze communication logs [3]. These data are visualized using Grafana and Kibana, enabling comprehensive visibility into network activity and facilitating efficient troubleshooting [4, 5].

## SYSTEM OVERVIEW

The server infrastructure of the SuperKEKB accelerator can be broadly categorized into three types: servers for core infrastructure services, user-facing servers, and administrator-exclusive servers. Each of these three server types is deployed across two network layers—the accelerator control network and the higher-level institutional network. In particular, the core service servers are configured with redundancy by deploying two servers per network layer, resulting in a total of four core servers in operation. These servers host essential services directly (i.e., without containerization or virtualization layers). The operating system selected for these servers is AlmaLinux 9 [6]. Previously, CentOS 7 had been used [7]; however, with its end of life (EOL), AlmaLinux was chosen as a binary-compatible alternative to RHEL.

User-facing servers and administrator-exclusive servers are operated as virtual machines (VMs). The physical host machines also run AlmaLinux 9 and utilize KVM (Kernel-based Virtual Machine) for virtualization, hosting multiple VMs on each server. This virtualization approach was adopted in response to recent advances in server hardware capabilities, which now allow a single physical server to reliably support multiple VMs. Currently, over twenty VMs are in operation across the system. VM management was performed using both the virsh command-line interface and

* hitoshi.sugimura@kek.jp

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD004

# TEN YEARS EXPERIENCE OF OPERATING TINY FANLESS SEVERS AS EPICS I/O CONTROLLERS FOR J-PARC MAIN RING

Shuei Yamada*, High Energy Accelerator Research Organization / J-PARC Center, Tsukuba, Japan

*Abstract*

J-PARC (Japan Proton Accelerator Research Complex) Main Ring (MR) deployed tiny fanless servers as EPICS I/O controllers (IOCs). During the construction phase of J-PARC MR in 2007, VME single board computers (VME SBCs) were introduced as IOCs. Subsequently, it was found that the majority of the control targets were network devices, meaning that VME bus was not mandatory. Pilot installation of tiny fanless servers as IOCs was carried out in 2014. Full-scale deployment began in 2015 and the transition from VME SBCs to those tiny fanless servers was completed by 2023. This paper reviews operational experience with tiny fanless servers and discusses future prospects.

## INTRODUCTION

The control system for the J-PARC MR accelerator was constructed in 2007 using EPICS [1] as the framework and Linux as the operating system for the computers [2]. It has been in operation since beam operation began in 2008 [3]. The front-end computers in EPICS-based control system are reffered to as I/O Controllers (IOCs).

The MR accelerator control group requires that the IOC has at least two network interfaces (NICs). The majority of controlled devices are network-connected equipment such as PLCs and oscilloscopes. Their network processing capabilities are not always as high as typical PCs. They may hang if a broadcast storm occurs due to causes such as incorrect network cable connections or equipment failure. By separating the higher-level control network from the controlled device network, network failures are localised, making the control system more robust.

## PINON SABATARO® TYPE-P

PiNON Sabataro® Type-P is a tiny fanless server equipped with two NICs (Fig. 1, [4]). At the MR accelerator, Sabataro was introduced as an IOC on a trial basis in 2014. From 2015, a full-scale deployment of Sabataros began [5], initiating the migration from VME SBCs [6] introduced in 2007 at the time of construction. By 2023, migration to Sabataro was completed with the exception of 37 IOCs using the VME bus (Table 1). As of 2024, 86 Sabataro units are in operation as IOCs for the MR. The specifications of the VME SBC and Sabataro are shown in Table 2.

Sabataro is compact (W80.6 mm × D110.6 mm × H34.4 mm) and costs approximately one-fifth that of a VME SBC, making it easier to deploy compared to VME SBCs. The number of VME SBCs that can be installed is limited by the number of available slots in the VME crate. Therefore, it was previously necessary to operate multiple IOCs on a single piece of hardware. The introduction of Sabataro enabled the separation of hardware per IOC. Furthermore, as multiple VME SBCs installed in the same VME crate shared a power supply, switching off the entire crate was necessary to power down a single IOC. Migration to Sabataro allows the power to be switched off only for the intended IOC. To reduce the cost of VME SBC deployment as much as possible, four hardware variants with different CPU and memory specifications were used according to the number of signal points and data size of the controlled device. In contrast, only one Sabataro variant is being deployed. Consequently, the introduction of Sabataro has enhanced maintainability and reduced management costs.




Figure 1: Appearance of PiNON Sabataro®.

In case of IOC hardware failure, a spare unit is ready to be replaced so that the accelerator control system will be recovered promptly. Unexpectedly, Sabataro had few faults. Of approximately 100 Sabataro units procured so far, four have failed, all exhibiting symptoms of being unable to power on. Compared to the 19 out of 32 control terminals introduced between 2013 and 2014 that failed, reliability of Sabataro is remarkably high.

## OPERATIONAL STATUS OF IOCS

Usage profiles of Sabataros were summarised in histograms. Each IOC used in the MR accelerator operation exposes data on its own usage status, such as CPU usage, amount of memory used, and network traffic, as EPICS PVs. These are monitored and archived by the control system. Of the 86 Sabataros operated as IOCs at J-PARC MR, 84 units were used. The remaining two units were omitted here as they were mistakenly not archived.

* shuei@post.kek.jp

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD006

# REAL-TIME EPICS-INTEGRATED RAON TIMING SYSTEM ENHANCEMENTS*

E. Kwon†, H. M. Jang,
Institute for Rare Isotope Science (IRIS), Institute for Basic Science (IBS), Daejeon, Korea
H. Yoo, Department of Electronics Engineering, Chungnam National University, Daejeon, Korea

*Abstract*

The RAON timing system enables precise synchronization for accelerator operations. During actual use, limitations such as restricted sub-event generation, lack of emergency timing signals, and inflexible frequency/width adjustments were identified. To address these, an analog feedback system using a Digilent ADP3450 and Raspberry Pi 5 is under development to enhance real-time monitoring and control. This system achieves real-time monitoring of key timing signals over the EPICS network and implements pulse-train, ramp-up, emergency-halt, and event-switching functions, directly addressing requirements identified during 2024 operations. These enhancements pave the way for safer, more flexible, and highly efficient high-energy beam delivery at RAON.

## INTRODUCTION

The RAON timing system ensures precise synchronization across all accelerator devices by distributing clock and event signals, embedding GPS time information directly into protocols [1]. RAON's timing system is designed to minimize timing errors and ensure reliable synchronization over the large 500-meter accelerator facility.

The timing system consists of an Event Generator (EVG), Event Fan-out (EVF), and Event Receiver (EVR). Modules are interconnected with optical fibers to minimize timing errors over the 500-meter accelerator. The EVG, which generates multiple clocks and manages event sequences, is implemented using a Zynq UltraScale+ MPSoC. The EVF and EVR modules, responsible for event distribution and reception, use Zynq-7000 series devices. This modular hardware design enhances precision, reliability, and scalability.

## OPERATIONAL CHALLENGES

RAON successfully completed internal system tests and initial beam delivery for users in the first half of 2024, but additional timing system functionalities were found necessary to deliver higher energy beams. The following operational challenges were identified:

- Frequency ramp-up for smooth power increase
- Pulse-train for segmented event control
- Emergency-halt and dynamic event switching

These operations currently require significant manual intervention and time. Research and development efforts are ongoing to address and improve these operational challenges.

## ENHANCED FEEDBACK SYSTEM

We implemented the feedback system by pairing a Zynq-7000–based ADP3450 [2] with a Raspberry Pi 5 (Fig. 1). Unable to buffer 10 MSPS high-speed samples in the ADP3450's internal memory, we stream data in real time to the Pi 5 for storage and processing, then forward it over the EPICS [3] network. As a result, beam frequency is now monitored live via the CSS Phoebus GUI in the control room.

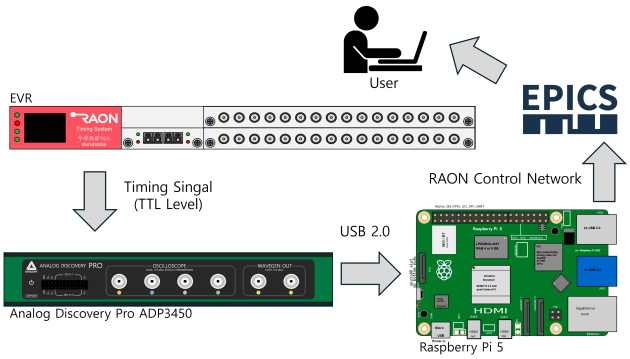



Figure 1: System infrastructure for real-time streaming: ADP3450 captures high-frequency beam data and transmits it to Raspberry Pi 5 for processing and EPICS network integration.

Additionally, we have developed and integrated the "pulse train" and "ramp-up" features. The pulse-train feature generates multiple sub-events within a single event trigger, while the ramp-up feature smoothly varies the output frequency over a specified duration between user-defined start and end frequencies (Fig. 2).

## TIMING SYSTEM CONFIGURATION

The current RAON timing system configuration utilizes an optical fiber network connecting EVG, EVF, and EVR modules (Fig. 3). This configuration ensures reliable synchronization across the entire RAON facility.

## IMPLEMENTATION RESULTS

The enhanced timing system now achieves real-time monitoring of key timing signals over the EPICS network. We have successfully implemented and validated the following functions:

* Work supported by the Institute for Basic Science (IBS-I001-D1).
† eunsang.kwon@ibs.re.kr

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD008

# OPERATING SYSTEM SECURITY UPDATES AND NETWORK BOOT SUPPORT FOR LIBERA INSTRUMENTS

M. Bajič, A. Bardorfer, R. Černe, A. Kete, M. Škabar
Instrumentation Technologies, Solkan, Slovenia

*Abstract*

Libera instruments are widely used in particle accelerators for applications like beam position monitoring, beam loss monitoring and control of radio frequency fields. The instruments rely on embedded Linux operating systems to ensure stable and precise operation. Maintaining security and operational reliability across a fleet of such instruments presents a significant challenge, especially in facilities with limited physical access. This paper presents a robust and automated solution for managing operating system security updates and network-based booting for Libera instruments. We detail an approach that integrates secure, version-controlled Linux OS updates with a centralized network boot infrastructure, enabling consistent and traceable deployments across all instruments. The network boot process further simplifies device provisioning and recovery, reducing downtime and minimizing maintenance overhead. This combined strategy improves system security posture, ensures reproducibility and supports scaling. Implementation experiences from production accelerator environments are also discussed.

## INTRODUCTION

In today's increasingly connected world, cybersecurity is a fundamental concern for all technology users. Protecting instrumentation systems from evolving threats is essential to ensure data integrity and system reliability. To address this, Instrumentation Technologies provides regular security updates for the Linux operating system (OS) that runs on Libera instruments. These updates help users keep their systems up to date with the latest security patches. Security updates are software updates designed to address vulnerabilities in the OS or applications.

Booting a Linux OS from the network, known as PXE (Preboot Execution Environment) booting, allows a computer to load its operating system from a remote server instead of a local hard drive [1]. This is especially useful for deploying Linux OS across multiple machines efficiently, which is well-suited to Libera users, who typically manage multiple instruments. Key advantages include centralized management, faster deployment, reduced hardware requirements (no need for local storage), and easier system maintenance or updates since everything is managed from a central server.

## LINUX OS SECURITY UPDATES

At Instrumentation Technologies, we have established a Security Update Process (SUP) to maintain and enhance the security of Libera products [2]. This process ensures that critical security updates are applied regularly to protect our devices from emerging cybersecurity threats.

In Fig. 1, we show a simplified SUP, where the Linux OS receives security updates without the need to recompile the Libera application software. The process includes automated development testing and regression testing by the quality assurance department. It can be executed periodically or on demand in case of high-severity security updates.

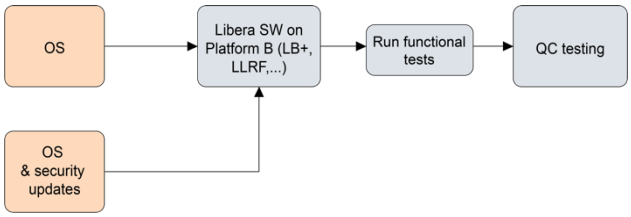



Figure 1: Diagram of a simplified security update process [3].

In Fig. 2, we present a detailed SUP that includes support for network boot. It outlines two possible scenarios, depending on the type of security updates to be applied. The process flow proceeds from left to right. In the first scenario, the security updates do not include Linux Kernel updates. As a consequence, there is no need to re-compile the Libera driver. In the second scenario, the security updates include Linux Kernel updates, and the Libera driver needs to be recompiled. The delivery of updated Linux OS to the customer also includes a script for generating the network boot folders. This simplifies the deployment of the updated Linux OS for users relying on network-based booting.

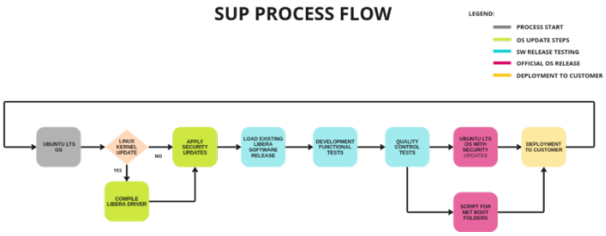



Figure 2: SUP flow [2].

Changes to the Linux OS or application libraries may require the Libera application software to be recompiled for compatibility. This would require updating the build environment (not shown in Fig. 2).

## NETWORK BASED BOOT OF LINUX OS

The network boot infrastructure contains three different servers: the DHCP server, TFTP server, and NFS server. Figure 3 shows an example DHCP server configuration that defines the mapping between the MAC address and the IP address of the Libera instrument is defined [3]. The configuration also includes information about the TFTP server

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD009

# PROMETHEUS SYSTEM MONITORING STACK

T. Gatsi, N. Makhoba, X. Baloyi, S. Hoosen
South African Radio Astronomy Observatory, Cape Town, South Africa

*Abstract*

The MeerKAT radio telescope, a 64-dish instrument located in South Africa, represents a significant leap in Southern Hemisphere radio astronomy, providing unprecedented sensitivity prior to its integration with the Square Kilometer Array (SKA). The operational efficiency of complex projects like MeerKAT relies heavily on robust Control and Monitoring (CAM) systems, which is underpinned by over fifty Linux server infrastructures. Ensuring the stability and performance of these servers is paramount to maximizing scientific output. This paper details the implementation and benefits of a Prometheus-based monitoring stack, designed to provide comprehensive surveillance of the MeerKAT CAM system's hardware, operating systems and CAM services. The system proactively detects issues, triggers alerts, speeds remediation, and improves troubleshooting, ultimately minimizing downtime and protecting critical data.

The monitoring stack for CAM is deployed with Ansible. The Node Exporter application extracts the operating system metrics such as `node_memory_MemFree_bytes`, `node_filesystem_free_bytes` as well as our custom CAM services metrics such as `karoo_camlog_status`, `karoo_vault_status` and feeds them into Prometheus. Prometheus processes these metrics and generates alerts using Python scripts that the Alertmanager application then uses to generate notifications and alarms via the Mattermost messaging application. Grafana transforms stored metrics in Prometheus into interactive dashboards.

## INTRODUCTION

The efficient operation of the MeerKAT radio telescope, a cornerstone of Southern Hemisphere radio astronomy, hinges on the continuous reliability of its underlying Control and Monitoring (CAM) infrastructure, composed of over fifty Linux servers. This paper introduces the implementation of a comprehensive, Prometheus-based monitoring stack, deployed with Ansible for scalability and resilience. Fig. 1 illustrates the Prometheus stack data flow, that will be further discussed in following sections: Node Exporter and Custom Metric Collection, Prometheus for Time-Series Data, Grafana for Data Visualization, Grafana for Data Visualization and Ansible for Automated Deployment.

## NODE EXPORTER AND CUSTOM METRIC COLLECTION

The Node Exporter is a crucial software component in the CAM Prometheus Monitoring stack, designed to collect and expose a wide range of operating system metrics. It is essentially an agent that runs on a machine and provides a wealth of information about its health

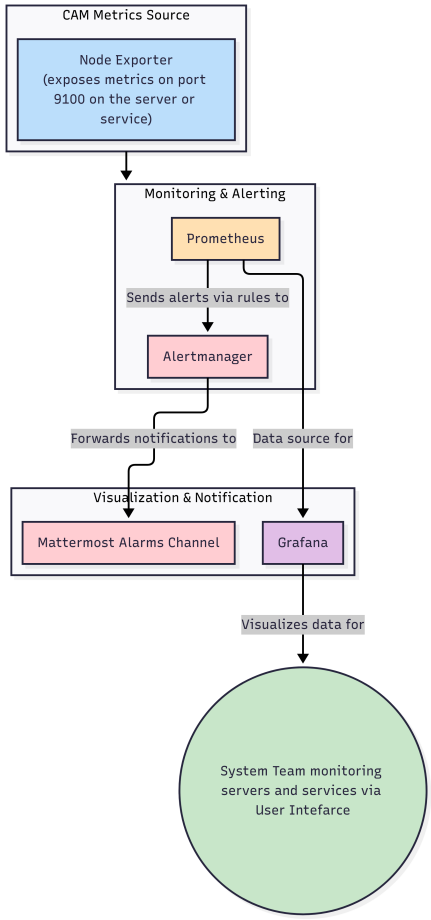



Figure 1: Overview of Prometheus Stack components in CAM.

and performance. The Node Exporter application extracts metrics such as `node_memory_MemFree_bytes` and CAM-specific service metrics like `karoo_camlog_status`, `karoo_vault_status` from monitored nodes and feeds them into Prometheus. CAM services are monitored by means of python scripts that poll an endpoint of each CAM service to determine the status of the service.

At its core, the Node Exporter is a Prometheus-style exporter for machine metrics. It's an executable program that, when run, creates an HTTP endpoint, typically on port 9100. An example URL `https://<node_exporter_url>:9100/metrics`.
When Prometheus scrapes this endpoint, the Node Exporter dynamically generates a list of metrics in a specific format that Prometheus can understand, digest and store in its internal time-series database (TSDB) .The metrics are presented as a simple text file, where each line represents a single metric, its value, and any associated labels.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD012

# SMART WATCH AND CONTROL IN ACCELERATOR CONTROL SYSTEM*

W. Fu†, S. Clark, S. Nemesure, J. Morris, Brookhaven National Laboratory, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The controls system at Brookhaven National Laboratory's RHIC complex contains millions of control points. Many of these produce alarms when a fault condition is present. The severity of alarms often depends on a combination of factors within the Controls system. To provide Operations with condition-specific alarms, it is sometimes necessary to monitor and evaluate multiple control points simultaneously. A server/client-based software architecture was developed which monitors multiple control points, and dynamically generates appropriate alarms based on user-defined programmatic relationships. The user interface tools are provided which allow Operations staff to create, update, enable, and disable conditions for any specific case on the fly. The flexibility of the system can help Operations to simplify diagnostics during some complex failure situations.

## INTRODUCTION

The RHIC accelerator control system has more than 2 million control parameters. Many of them are critical for the normal operation of the accelerator system. The ability to monitor the values of these critical control points and send alarm or take action accordingly is an important control system feature. This capability can greatly improve the overall performance of an accelerator control system.

Monitoring a single control parameter with target range is easy. Monitoring multiple control parameters simultaneously with desired programmatic logic in a way which can be applied to many control situations is a challenge. Much research has been devoted to this topic and different approaches have been explored. Some are hardware focused, and some others are software based or combined. For example, A "Smart Alarm" system based on trained neural network was tested for the Continuous Electron Beam Accelerator Facility at Jefferson Lab [1]; A machine learning based smart alarm system for CEBAF injector control was developed and tested [2]; Liang *et al.* developed and tested an intelligent alarm system which integrates control alarm signals with LabView, mobile device messages, and emails etc. [3] and improved the stability of their control alarm system. We take a different approach. We take advantage of our existing control system features and extend their capability to achieve the smart watch and control of the systems.

In a generic sense, a smart watch and control system should do more than just issuing alarms. It would be very helpful to include the following parts:

1. Smart Monitoring: Monitors a list of control parameters, integrates the monitored results from all parameters and implements desired programmable monitoring logic (mathematical manipulation of results from each parameter) and gives more meaningful results.
2. Smart control actions: Specify desired control conditions, and use these conditions as trigger points to do many things such as: send alarms; send emails or text messages; change the values of other control parameters accordingly; and/or trigger the execution of a program.
3. Log the details of events from monitored devices with time, IP, user, action taken, etc. so all things can be track-able for post mortem data analysis.
4. User can easily setup new watching case or update existing watching cases with graphic user interface (GUI) Tools.

In the RHIC control system, we have developed a General Purpose Monitor (GPM) system many years ago [4]. This is a very powerful system which allows the monitoring of any list of control parameters of accelerator control devices with concise GPM programming language codes. Complicated monitoring logic can be implemented with the GPM programming language. This makes real-time device control data manipulation easy. To combine the GPM with the powerful python mathematical packages, a new PYGPM engine was developed at Brookhaven National Lab [5]. With this PYGPM Engine, we can:

1. Monitor any control parameters (for settings or measurements) with asynchronous IO connections and get real-time live data;
2. Take the advantages of more than 350 mathematical functions. These include all GPM system functions plus the python numpy [6] functions.
3. A simple syntax GPM programming language [7] can be used to setup a GPM file which user can use to implement any monitoring control logic and mathematical manipulation.

Our watch and control programs were created based on this PYGPM engine, which mainly does the smart monitoring work. The programs also implement the features of all smart control actions.

This paper presents the software system we designed and used in the RHIC control system with these features and capability, describes how the programs (servers and GUI tools) monitor multiple control parameters with a rich programmatic relationship, send alarms and/or take actions according to the watching data configuration.

The Software system consists of servers and GUI tools for front end users. The systems are capable of monitoring multiple cases. In each individual case, the system can monitor multiple control parameters (setting points or measurement points) with a user defined control logic. This makes the system very flexible to configure and extend.

* Work supported by Brookhaven Science Associates, LLC under Contract No. DE-SC0012704 with the U.S. Department of Energy.

† fu@bnl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD013

# ENSURING THE FAST RESPONSE IN THE POWERING INTERLOCK CONTROLLER (PIC) UPGRADE: A UNICOS-BASED APPROACH

J. Fernandez-Cortes*, E. Blanco, B. Schofield, A. Antoine, M. Kalinowski, I. Romera
CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Powering Interlock Controller (PIC) plays a critical role in supervising the powering conditions of the superconducting magnet circuits of the Large Hadron Collider (LHC), reducing the risk of severe equipment damage. It ensures that conditions are met before granting the powering permit and reacts within milliseconds to remove it if conditions become unsafe for operation. To enhance maintainability, modularity, and long-term sustainability, the PIC interlock system has been redesigned within CERN's UNICOS framework. This modernisation posed a major challenge: preserving the system's stringent real-time constraints while adopting high-level programming in Structured Control Language for improved maintainability. This paper presents the strategies adopted to ensure the required response time, including execution path optimisations and efficient handling of interlocks within the UNICOS framework. We discuss the trade-off between standardisation and performance, the impact on system diagnostics and operation, and the validation process ensuring compliance with LHC safety requirements. This work serves as a reference for applying UNICOS in protection-critical systems while preserving real-time capabilities.

## INTRODUCTION

The LHC relies on a complex machine protection system to safeguard its superconducting magnets, power converters, and auxiliary equipment from damage caused by failures or magnet quenches (when the superconducting magnets transition to a normal conducting state). With over 10,000 superconducting magnets powered through more than 1,600 electrical circuits, the stored magnetic energy per sector can exceed 1 GJ, making fast and reliable protection systems essential.

Two complementary subsystems form the backbone of this protection: the Beam Interlock System (BIS), responsible for beam-related risks, and the Powering Interlock Controller (PIC), which supervises the powering conditions of superconducting magnet circuits. The PIC authorizes or inhibits powering based on hardwired interlock loops with the Quench Protection System (QPS), power converters, and auxiliary systems such as UPS and cryogenics. In case of unsafe conditions, it either triggers a controlled ramp-down of the current in the magnets or issues a fast power abort, and can request a beam dump via the BIS.

When the LHC was commissioned, 36 PICs were deployed across the eight access points of the accelerator, with between two and six PLC-based systems per Point, depending on the number of powering subsectors to protect. This distributed architecture met the stringent safety requirements, but also introduced complexity in maintenance and inter-PIC communications, particularly for beam dump requests [1].

To date, the PIC has been implemented with the CERN's Unified Control System (UNICOS) dedicated package designed to meet the specific constraints of LHC magnet protection. Although this customized solution has operated reliably, it has accumulated considerable technological debt over time. The reliance on custom PLC and SCADA (Supervisory Control and Data Acquisition) code has led to high maintenance costs, as it requires in-depth design knowledge across multiple teams. In addition, support of specialized components, such as the time-stamp push communication protocol, has further limited flexibility for future upgrades. To overcome these challenges, the new version of the PIC has been re-engineered as a standard UNICOS application using the Continuous Process Control (CPC) package.

This paper describes the upgrade of the PIC within the UNICOS framework, emphasizing the strategies adopted to balance maintainability and standardization with demanding performance requirements.

## POWERING INTERLOCK CONTROLLER FUNCTIONALITY

The Powering Interlock Controller is a security-critical component of the LHC machine protection system, ensuring that superconducting magnet circuits can only be energized under safe conditions (Fig. 1). Its main role is to authorize or inhibit the operation of power converters, based on the status of multiple subsystems, and to trigger protective actions in case of detected failures.

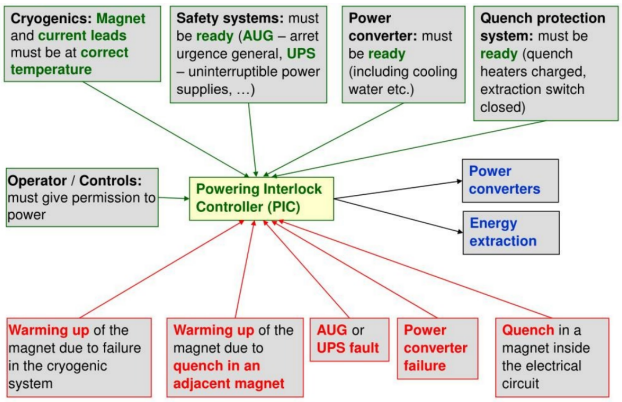



Figure 1: Conditions for powering.

* jesus.cortes@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD014

# BUILDING ON EXPERIENCE: RENOVATION OF THE SNIFFER GAS AND FIRE PROTECTION SYSTEM FOR THE LHC EXPERIMENTS

F. Havart, A. Andruszkiewicz, N. Broca, F. Deperraz, S. Grau,
T. Hakulinen, P. Ninin, R. Nunes, N. Venzin
CERN, Geneva, Switzerland

*Abstract*

The LHC SNIFFER system was commissioned in 2006 to protect the personnel working in the large LHC experiment caverns from hazards caused by fire, $CO_2$/flammable gas leaks, and oxygen deficiency. Currently, SNIFFER operates within the ALICE, ATLAS, and LHCb experiments. The system is composed of custom-built modules, each housing an aspiration pump and a combination of sensors tailored to the specific needs of each experiment. Sensor control units are connected to a central control PLC via a PROFIBUS link. The SNIFFER system enables air sampling and analysis from the most remote locations within the LHC experiments, including inside the detectors, using a dedicated air piping network. After nearly 20 years of continuous service, a complete renovation has become necessary. The aim of the modernization project is to address the identified shortcomings, resolve the obsolescence of key components, and significantly reduce the overall cost of system ownership and maintenance. The new system, known as SNIFFER NG (New Generation), is designed to guarantee the safety of personnel in the LHC experiments up to the announced end of operation of the HL-LHC era. This paper presents the new design philosophy, the rationale behind the selected design choices based on the operational experience gathered over the past two decades, and the key challenges associated with the development and deployment of the new system.

## INTRODUCTION

The CERN Large Hadron Collider (LHC) [1] accelerator has been designed to provide particle beams to four experiments, each hosted in independent underground caverns. These vast and complex structures pose several potential safety concerns which must be addressed. Some issues of note which are discussed within this report are: limited space within several locations of concern; quantity of electrical circuits and connections and the presence of gases, both flammable and oxygen deprivation capable. These hazards mandate a careful and extensive monitoring. To date, there is no commercial off-the-shelf (COTS) solution which incorporates the necessary sensors. Moreover, most COTS sensors are not designed to tolerate the pronounced magnetic influence, or spatial constraints present within the experimental magnets. Consequently, it was necessary to design and manufacture a custom solution, which includes specific modules, capable of providing remote analysing chamber to shelter the sensors.

Since their construction, three of the four LHC experiments (ALICE, ATLAS and LHCb) have seen a custom-made detection and interlock combined system installed and operated to cover smoke detection, oxygen deficiency hazard (ODH) and gas leaks ($CO_2$, $CH_4$, $nC_5H_{12}$).

Based on a modular design, SNIFFER consists of a control application, which uses a local SCADA (Supervisory Control and Data Acquisition) and PLC (Programmable Logic Controller), and a variable set of detection modules. The latter are built to a standardised size, designed with a mechanical enclosure to slide in a custom rack. Each SNIFFER module contains a vacuum pump, a selection of sensors, and a microcontroller-based electronic card. This card runs an application-specific firmware and provides a PROFIBUS interface for communication with the control PLC.

Dedicated aspiration networks are deployed within hard-to-reach areas of the experimental magnets, allowing continuous air sampling from these zones. The collected samples are analysed in real time by the sensors within the SNIFFER modules. The latter block diagram is presented in Fig. 1. Pre-alarms and alarms are generated, by comparing the current measurements with the embedded predefined values. When those predefined thresholds are exceeded, the control PLC reports them to the CERN Fire Brigade, who may initiate personnel evacuation if necessary. Additional protective measures, such as power cuts and other machine-protection actions can also be triggered.

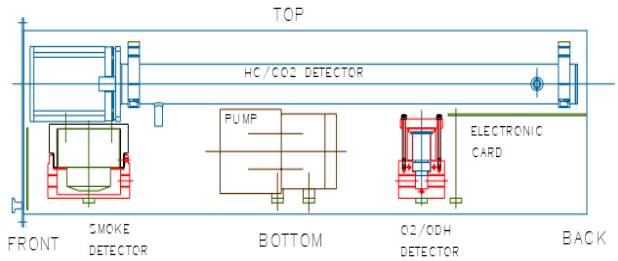



Figure 1: Module block schematic, side view.

Each SNIFFER module is subject to mandatory annual maintenance. During this process, pump valves, silent blocs, and air intake filters are replaced, sensors are calibrated, and the module airtightness is assessed. Consequential of a proposed pool of spare modules of each configuration, maintenance can be conducted with minimal impact on the safety of the chambers by module substitution: those requiring maintenance are exchanged for pre-

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD015

# CERN SAFETY SYSTEM MONITORING – ENTERING A NEW ERA

T. Riesco, T. Hakulinen, F. Havart, P. Ninin, R. Nunes, CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

For the past 20 years, the CERN Safety System Monitoring (SSM) framework has safeguarded the operational health of CERN's access and personnel safety systems. Built on the Zabbix monitoring platform, Grafana, and in-house developments, SSM provides real-time diagnostics, alerts, issue escalation, and predictive analytics for a wide range of critical infrastructure, operating systems, network devices, storage, and specialized equipment like video cameras and UPS units. The objective of SSM is to enhance maintenance and operational efficiency by delivering timely and reliable system feedback, enabling rapid identification of both immediate failures and gradual degradation before they can impact operations. SSM also supports long-term data analysis for post-incident investigations, statistical evaluations, and trend forecasting, thereby contributing to the optimization of safety system designs. It also provides operational statistics and graphs to CERN management and site services offering valuable graphical insights. Ongoing developments aim to expand SSM's capabilities through integration of AI modules for predictive maintenance, enabling pre-emptive interventions and reducing downtime. These enhancements also include automatic generation of reports and notifications to operations teams. By continuously assessing the safety status of operational systems, SSM plays a crucial role in mitigating risks and ensuring long-term reliability of CERN's technical infrastructure.

## INTRODUCTION

CERN's mission to explore fundamental physics depends on a vast and sophisticated network of technical and scientific systems, ranging from particle accelerators and detectors to complex computing and safety infrastructure. To ensure these systems operate reliably and safely, real-time data collection and long-term analysis are essential. The Safety System Monitoring (SSM) [1] framework, developed and operated by CERN, provides continuous supervision of the access and personnel safety infrastructure. It serves both technical experts and operational staff by enabling fast incident response, performance monitoring, and risk mitigation like the service down, safety and physical operational risks, etc.

Over the past two decades, SSM has grown from a basic alarm collection system into a fully integrated monitoring platform based on open-source software and cloud infrastructure. In this paper, we describe the implementation, components, and strategic direction of SSM, with particular attention to virtualization, automation, resilience, and data management.

## SYSTEM ARCHITECTURE AND DEPLOYMENT

SSM (Fig. 1) is deployed primarily on CERN IT's OpenStack-based infrastructure [2], leveraging virtual machines (VMs) to ensure scalability, flexibility, and integration with CERN's broader computing ecosystem. The main virtual machine hosting the core services runs AlmaLinux 9 [3], a stable and enterprise-grade Linux distribution. This OS supports a diverse software environment where most automation and data manipulation are handled using Python and BASH scripting.

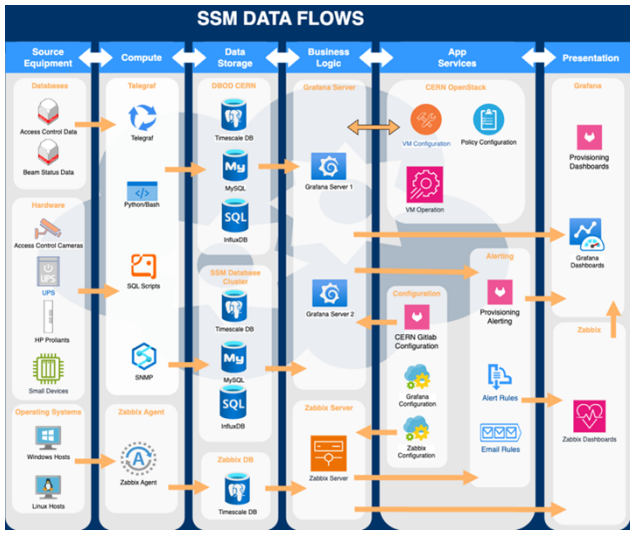



Figure 1: SSM architecture.

Key applications — including Zabbix [4] for monitoring and Grafana [5] for visualization — are deployed using Docker containers orchestrated by Kubernetes. This containerized architecture simplifies maintenance, improves fault isolation, and enables rapid updates without interrupting the entire system. Updates to Zabbix, Grafana, or internal scripts can be rolled out and rolled back with minimal impact, ensuring continuous monitoring availability.

To monitor the health of SSM itself, we use Healthchecks.io [6] probes configured internally to validate key components such as data pipelines, database status, and alerting mechanisms. If anomalies are detected—such as stalled data ingestion or unresponsive dashboards—alerts are triggered automatically and sent to the operations team.

## RESILIENCE AND REDUNDANCY

Given the criticality of the safety systems monitored, SSM is designed with redundancy in both data paths and computing infrastructure.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD017

# THE ESS FAST BEAM INTERLOCK SYSTEM – DESIGN, DEPLOYMENT AND COMMISSIONING TILL THE BEAM ON DUMP

S. Gabourin*, A. Abujame, A. Nordt, E. D'Costa, M. Carroll, R. Xavier da Silveira, S. Pavinato
European Spallation Source, Lund, Sweden

*Abstract*

The European Spallation Source (ESS), is a linear accelerator located in Lund, Sweden. It is currently under completion and will be the most powerful neutron source. A key system ensuring the safe operation of the machine is the Fast Beam Interlock System (FBIS), which is the brain of the Machine Protection System (MPS) at the ESS. FBIS is both modular and distributed, designed to react to approximately 250 input signals from critical accelerator and target subsystems at the time of this commissioning. The commissioning phase is until beam on dump in 2025. Its role is to assess the beam clearance conditions in real time, ensuring a fast beam stop when necessary to prevent unsafe operation. To meet the requirements of the protection integrity level, FBIS operates at a high data throughput and ultra-low latency. This paper provides an overview of the FBIS control system and the most significant challenges faced during the last commissioning phase. The focus was on integrating several new systems and automating integration tests across the site. The strategies used to validate and deploy over 20 newly installed crates, as well as the important role automation plays in ensuring reliable and efficient commissioning under increasingly complex system conditions.

## INTRODUCTION

At ESS [1] the Machine Protection (MP) system has been designed in order to have an optimum balance between appropriate protection of equipment from damage and high beam availability [2]. Since the proton beam power of 125 MW per pulse (5 MW average) will be unprecedented and its uncontrolled release can cause serious damage to equipment within a few microseconds an FBIS is required to provide a minimum latency. The FBIS is responsible of collecting several types of information from different kinds of so-called Sensor Systems, which generally fall under two categories: slow systems (Vacuum, Magnets and position of Insertable devices) or fast systems (Radio Frequency Control Systems and Beam Instrumentation). Sensor Systems provide signals feeding FBIS signal processing units that make the decision on maintaining or not the beam production.

## ARCHITECTURE OF FBIS

The FBIS architecture is fully redundant to ensure it can reach the Protection Integrity Level requested by the Protection Functions. Three hundreds of Sensor System connections are installed along the 600 m Linac, leading to a specific architecture with two component types: the "Signal Conversion Unit" (SCU) and the "Decision Logic Node" (DLN).

The SCU, Fig. 1, is a concentrator based on cPCI form factor with a set of mezzanine cards (MCs), it serializes all its connected Sensor Systems data on optical link using an AMD Aurora 64/66b protocol [3]. Several kinds of MCs were designed to interface with various types of Sensor Systems. One of them manages PLC-based Sensor Systems, and two others fast electronic Sensor Systems, mainly FPGA-based. For simplification purpose, an effort was made to use only one of the two MCs to connect all FPGA-systems during the Beam on Dump (BoD) commissioning phase.

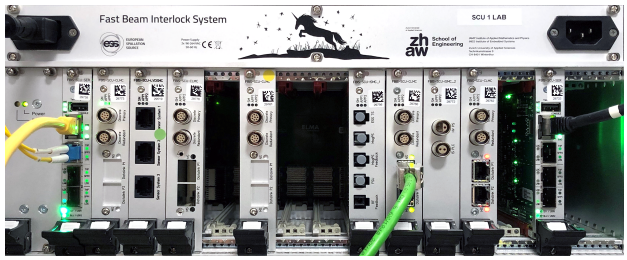



Figure 1: Signal Conversion Unit - SCU.

The DLN, Fig. 2, is a 3U chassis MTCA standard [4] that gathers data from SCUs and performs the protection logic of the FBIS to maintain or switch off the beam by using the Actuators. It connects to the high level control system EPICS and timing system using a Timing Event Receiver (EVR) [5].



Figure 2: Decision Logic Node - DLN.

More detailed information can be found in Ref. [6]. The full deployment will be of 7 DLNs connecting the 25 SCUs as shown on the schematic in Fig. 3, where the equipment used for the previous commissioning phase is highlighted.

## FULL LINAC DEPLOYMENT

In 2022, the Beam commissioning concerned the first 50 m of the Linac, driving the Beam up to the DTL4 faraday cup (FC) in the Normal Conducting Linac (NCL). The FBIS equipment to be deployed was the two first DLNs with their

* stephane.gabourin@ess.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD018

# THE ESS FAST BEAM INTERLOCK SYSTEM HISTORY BUFFER FOR POST-MORTEM ANALYSIS AND ACCELERATOR STATISTICS

S. Gabourin*, A. Nordt, E. D'Costa, M. Carroll, S. Pavinato
European Spallation Source, Lund, Sweden

*Abstract*

The European Spallation Source (ESS), is a linear accelerator located in Lund, Sweden. It is currently under completion and will be the most powerful neutron source. A key system ensuring the safe operation of the machine is the Fast Beam Interlock System (FBIS), which is the brain of the Machine Protection System (MPS) at the ESS, gathering all information to make decision on keeping or stopping beam production. Being the only system stopping the Beam, it is the central point where transit all Beam stop requests. The latter, with a multitude of events, are recorded to be used for the so-called post-mortem analysis. This consists of keeping an history of events preceding a Beam Stop, identifying what was its root cause, then ensuring the machine can be reliably restarted. To achieve such an analysis with events from fast systems, a resolution of few nanoseconds is necessary and some firmware component have to be used for the data collection. This component is called history buffer and is present in each of the 56 FPGAs of the FBIS actually in operation. This paper explains what are the history buffers, their implementation, and how to exploit their information for post-mortem analysis and accelerator statistics.

## INTRODUCTION SCENARIO

Starting from an operational scenario where a Regular Beam Interlock (RBI) stopped the beam, and Operator Public Interfaces (OPIs) show that a superconducting cavity has tripped and some beam loss tripped Beam Loss Monitors (BLMs). Operators want to resume operation, as it looks like the cavity trip induced a loss of velocity of the beam that was partially lost and detected by BLMs. These events are fast, occurring within a few hundreds of nanoseconds, and are difficult to distinguish. The reason could be beam instability making some bunches touching an insertable device, like a collimator, spread part of the beam in the superconducting cavity and inducing the cavity trip. In this case the beam loss may be detected by the BLMs a few microseconds before the cavity trip. Operators should then focus on tuning the beam or moving the collimator out to avoid more trips, instead of re-tuning the cavity.

How to understand what happened from the beginning? With "something" recording events occurring with a very low time resolution and in a deterministic way, like a history buffer (HB) implemented in firmware, where the event "BLM loss" appears slightly before the event "cavity trip". And the firmware needs to have access to these information, which is the case of the one of the Fast Beam Interlock System (FBIS) at ESS [1].

* stephane.gabourin@ess.eu

## ESS AND FBIS

The FBIS collects data from Sensor Systems, mainly binary information like Beam Permits (BP) and Ready (Rdy) signals. When set to "Not OK" (NOK), these signals can induce a Beam Switch Off (BSO) request and stop the beam production.

Moreover, the ESS accelerator is a machine that is almost 600 m long, with standard sections of a linear accelerator. The beam can be stopped at nine places called Proton Beam Destination (PBD) to facilitate intermediate commissionings. One of them is the Ion Source itself, four are Faraday Cups (FC) in the Normal Conducting Linac (NCL), two are Intermediate Beam Stoppers (IBS) in the Super Conducting Linac, one is a Tuning Beam Dump (TBD) at the end of the straight accelerator, and the last one is the Target which produces neutrons. Except TBD and Target, all other PBDs can be inserted or removed from the beam path.

In order to tune the beam sent to any destination, a second parameter is defined: the Proton Beam Mode (PBM). Twelve of them are defined for normal beam operation, from "No Beam", "Probe Beam" to "Production". Except the "No Beam" for which all is zero, these modes are used to fix limits for 3 beam parameters of each beam pulse: the length of the pulse, from 5 us to 2.86 ms, the current from 6 mA to 62.5 mA, and the pulse repetition rate, from 33 mHz to 14 Hz.

Many systems use these 2 parameters PBD and PBM that are delivered by the Timing Distribution System (TDS), like the FBIS and some of its Sensor Systems. In the case of Machine Protection (MP), they are used to tune protection thresholds, mask interlocks, etc. Therefore the FBIS needs to read them from all the Sensor Systems using them, and ensure they are properly configured.

BP, Rdy signals, PBD and PBM are 95 % of the signals read by the FBIS, the rest being flags at specific places from Sensor Systems or Actuators.

The FBIS is not a simple piece of electronic, but a full set of electronic crates of 2 kinds - Signal Conversion Unit (SCU) and Decision Logic Node (DLN) [2] connecting with an optical link called Slink, which uses the Xilinx Aurora core as the communication protocol. Seven DLNs and twenty five SCUs are to be installed for the final configuration (Fig. 1), each of them being fully redundant. Therefore it consists of 64 FPGAs, each one containing an HB. Almost all the FBIS electronics were in operation during the commissioning to the Beam on Dump that took place on the first half of this year [3].

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD019

# EVENT-DRIVEN ALARM MANAGEMENT IN THE KARABO SCADA SYSTEM

F. Sohn*, G. Flucke, A. García-Tabarés Valdivieso, D. Goeries, S. Hauf, I. Karpičs,
European XFEL, Schenefeld, Germany

*Abstract*

In this contribution a flexible software for situational awareness and alarm management within the Karabo Supervisory Control and Data Acquisition (SCADA) system is presented. The supervision of hardware and software components is an essential function of a SCADA system and includes alarm management as a key aspect. This means that a SCADA system should detect components that are running in abnormal conditions and trigger alarms to alert operators. The presented software allows operators to define alarm conditions and, in turn, tracks dependent device properties. The evaluation of alarm conditions is event-driven: an alarm condition is evaluated once an updated value of a dependent property is received. If an alarm condition is met, alarms can be issued in various formats such as text messages, sounds, or visual indicators. The software is widely used at the European X-ray Free Electron Laser facility and decreases down-times of instruments and software modules by allowing staff to address problems immediately or even proactively.

## INTRODUCTION

The successful operation of large-scale research facilities like the European X-ray Free Electron Laser (XFEL) necessitates reliable and precise operation of a vast number of scientific instrumentation as well as technical infrastructure. Failure of instrument components can result in a considerable loss of hardware, beam time, and facility reputation if experiments become unfeasible to execute within the allotted beamtime. The scientific hardware components are typically integrated into a Supervisory Control and Data Acquisition (SCADA) system, providing access to operational parameters. Addressing the "supervisory" aspect of "SCADA", an alarm and notification system is an important part of a SCADA system to alert operators if components within the control system operate under unusual conditions. Real-time monitoring with immediate notification on detected problems allows facility staff to act quickly to prevent or shorten downtimes, to prevent damage to scientific instrumentation, and to do preventive maintenance. Common ways to issue alerts are text notifications, such as emails, as well as sounds or visual indicators.

In this contribution, a new software for the SCADA system at the European XFEL is presented, which is used to monitor components of the control system and alert the responsible personnel to detected problems.

* florian.sohn@xfel.eu

## ALARM MANAGEMENT IN THE EVENT-DRIVEN SCADA SYSTEM KARABO

This section first gives an overview of the Karabo SCADA system [1–3] and its new alarm and notification system, the so-called *Event Management System*. Afterwards, details on the implementation of the Event Management System are presented.

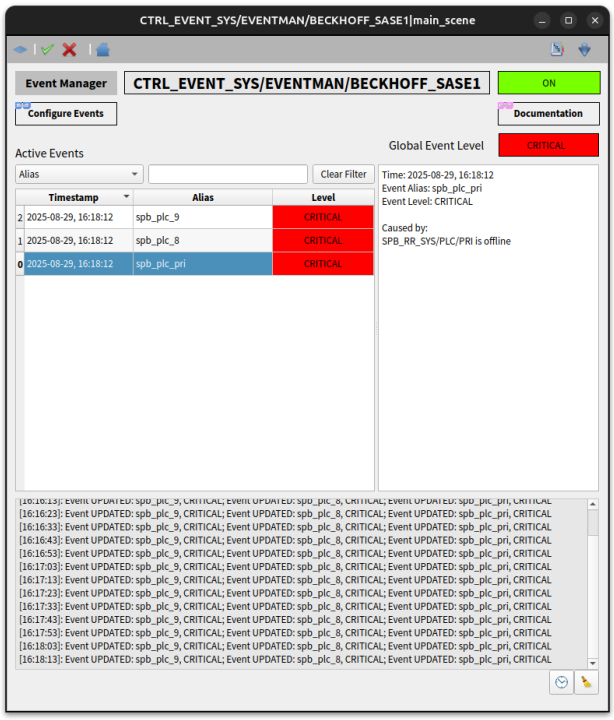



Figure 1: The main scene of the event manager device. The scene contains a list of currently active events and, for the selected event, shows the conditions that triggered the event. At the bottom of the scene, a log widget provides a history of event updates. A 'Global Event Level' is set according to the highest active event level.

### *The Karabo SCADA system at the European XFEL*

The SCADA system at the European XFEL research facility is used to operate thousands of hardware components within the photon beamlines and experiment end stations. It is built on top of the Karabo SCADA framework and integrates a wide variety of so-called Karabo devices, i.e. software modules that are tailored for specific tasks, such as controlling a specific piece of hardware. Communication between control system components relies heavily on

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD020

# THE CONTROL AND SAFETY SYSTEM FOR THE JULIC NEUTRON PLATFORM TARGET STATION

H. Kleines†, F. Suxdorf, A. Möller, A. Steffens, M. Glum, P. Kämmerling, J. Voigt, J. Baggemann, U. Rücker, P. Zakalek, T. Gutberlet, JCNS-2, Forschungszentrum Jülich, Jülich, GermanyY. Beßler, R. Hanslik, D. Marshall, F. Palm, ITE, Forschungszentrum Jülich, Jülich, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

For the future high current accelerator-driven neutron source HBS (High Brilliance Neutron Source) at Forschungszentrum Jülich, a prototype for a target station has been developed, which was operated successfully at the JULIC neutron platform and will be relocated to the ARGITU accelerator at ESS Bilbao, in future. A major safety-related feature of the target station is the automatic motorized opening of the shielding gate, which potentially could expose humans to radiation or crushing hazards. Based on a risk assessment according to EN ISO 12100, a safety system hast been designed that fulfills the requirements of the European machinery directive, which is necessary for CE marking. The safety system relies on safety edges, enabling switches, emergency stop switches and an interface to the personal safety system of the accelerator. The functional safety achieving PL d according to EN ISO 13849-1 has been implemented with a Siemens S7-1500 safety PLC, which is integrated into an overall control system based on TANGO and NICOS.

## TARGET STATION PROTOTYPE AT THE JULIC NEUTRON PLATFORM

The Jülich HBS project [1] aims at the development of a novel High Current Accelerator-driven Neutron Source that fills the gap between neutron lab sources and high flux spallation sources. In the HBS, protons from a LINAC (70 MeV, 100mA peak current) will be distributed to 3 target stations where they hit tantalum targets, where they produce neutrons by a nuclear reaction. In more than twenty neutron instruments mounted at the target stations beam ports these neutrons will be used for the study of samples in various research applications.

A target station prototype for the HBS was built at the COSY facility at Forschungszentrum Jülich to demonstrate the feasibility of the concepts and to test essential components like target, reflector, thermal moderator and the cold moderators [2]. JULIC, the injector cyclotron for cooler synchrotron COSY, was also used as a proton source for the target station demonstrator providing 45 MeV proton pulses at 25 Hz repetition rate and 8 µA peak current. The target station setup included two neutron instruments, the reflectometer HERMES provided by LLB and the diffractometer TOAD. A liquid hydrogen cryogenic moderator has been built that is used by HERMES and a solid methane cryogenic moderator has been built that is used by TOAD. Several successful beam times have been provided at this so-called JULIC Neutron Platform from December 2022 to December 2023 during which additional neutron extraction channels have also been used by external groups. Fig. 1 shows the prototype target station at the JULIC neutron platform.

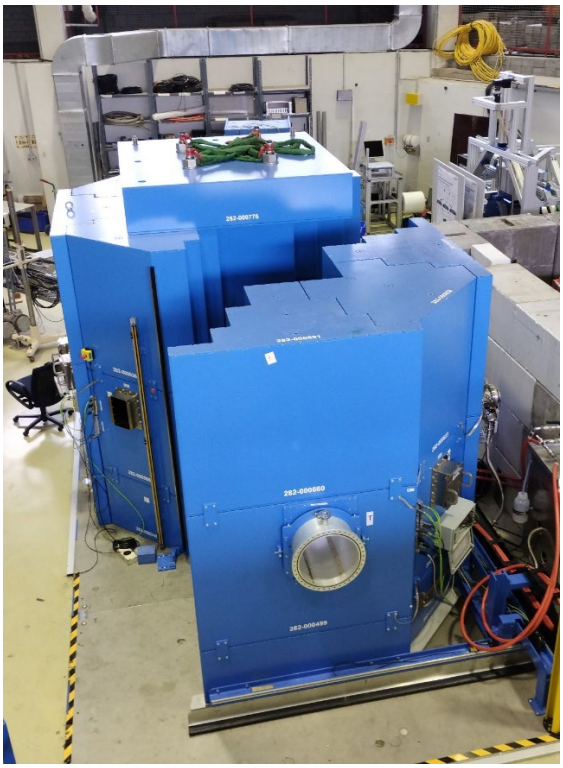


Figure 1: JULIC neutron platform target station.

In the target station, the target is mounted on a removable target plug sitting in an L-shaped tube that is connected to the proton beam line via a vacuum window. A cooling system for the target and a vacuum system for the L-shaped tube as well as some diagnostic devices have been implemented. For easy modifications of the moderator-reflector assembly the target station shielding is equipped with a shielding gate that can be opened and closed automatically with a Parker ETH100M electric cylinder equipped with a Parker SMH115 servo motor.

## TARGET STATION CONTROL SYSTEM

A control system has been implemented for the target station that consists of several subsystems, e.g. for target cooling or target vacuum. Since the target station has an automatically moveable shielding gate, personnel protection is a major issue of the control system. Possible hazards

† email address h.kleines@fz-juelich.de

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD021

# DESIGN OF A LASER SAFETY SYSTEM FOR S3 PROJECT*

N. Simon-Bauduin†, B. Osmond, L. Rousseau
Grand Accélérateur Nat. d'Ions Lourds, Caen, France

*Abstract*

The SPIRAL2 accelerator began operating in 2021 deliver ion beams to the experiment Room S3 (Super Separator Spectrometer). This hall hosts the S3 spectrometer, currently under commissioning, as well as the Low Energy Beam Line (S3LEB). The S3 project also includes the installation of a laser system in the building housing the S3 separator. The laser beams required for laser spectroscopy experiments are produced either in Laser Room by a TiSa (Titanium/sapphire) laser system and/or in Laser Room (by a dye laser system). If necessary, both laser systems can produce beams simultaneously for the same experiment.

The purpose of the new laser setup is to enable resonant laser ionization of nuclei of interest produced by the S3 spectrometer, or to perform laser spectroscopy experiments.

## LASER SYSTEM

### *TISA System*

The Ti:Sa laser system (Class 4) consists of several laser cavities and laser transport lines. The laser cavities are located on two optical tables (See Fig. 1) surrounded by an aluminum structure that prevents the laser beams from leaving the optical table up to a height of 20 cm. The laser systems are connected to power supplies and water-cooling units.

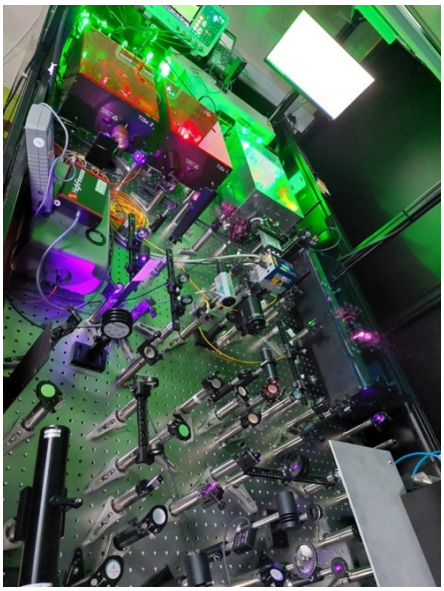


Figure 1: Optical table.

### *Laser Installation*

The laser beams produced in Room 41 are transferred to Room 52 before being sent to their final destination (Fig. 2). The different laser beams originating from Room 52, or from Room 41 (via Room 52), can then be directed either to the gas cell of the Low Energy Beam Line, to S3, or offline to a multiphysics Room.

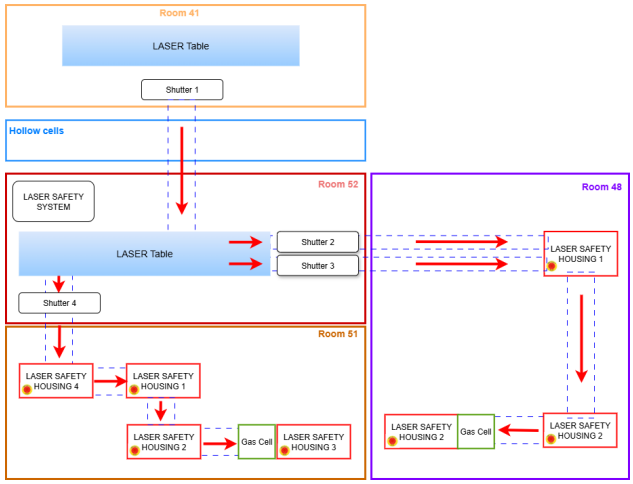



Figure 2: Laser installation.

## LASER SAFETY SYSTEM

As indicated above, the installed lasers are classified as Class 4. Class 4 laser devices are high-power lasers that are extremely hazardous to eyes and skin. Several safety measures are therefore implemented to protect personnel and equipment. The lines can be opened within the "laser zones" for beam handling by authorized personnel.

In public areas, the beam lines are enclosed, safety Laser housing allow authorized laser personnel to perform optical adjustments. In addition to ensure personnel safety, a safety system is required. It manages signaling, beam-blocking elements, and the laser power supplies. It ensures that safety conditions are met before allowing the production of laser beams and their transmission within the relevant areas.

### *Requirements*

The main safety features are associated with the laser protective housings, which are equipped with door sensors to monitor their closure and with an emergency stop to shut off the laser beam.

These safety systems are used when the LASER system is in operation and ensure the protection of personnel present in the S3 or multiphysics rooms. The laser shutdown triggered by the opening of a protective housing must be bypassed to

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

* S3LEB Project.
† nicolas.simon-bauduin@ganil.fr

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD022

# RADIATION MONITORING SYSTEM USING A PIN PHOTODIODE AT THE KEK ELECTRON/POSITRON INJECTOR LINAC

I. Satake[†,1], H. Iwase[1], T. Suwada[1], S. Kusano[2], D. Asano[3], M. Tokiyoshi[3], Y. Matsumoto[4], Y. Ishigaki[5], M. Sato[6],
[1]High Energy Accelerator Research Organization (KEK), Tsukuba, Japan
[2]Mitsubishi Electric System & Service Co., Ltd, Tokyo, Japan
[3]TAISEI CORPORATION, Tokyo, Japan
[4]Keio University, Yokohama, Japan
[5]The University of Electro-Communications, Chofu, Japan
[6]Yaguchi Electric Corporation, Ishinomaki, Japan

*Abstract*

We have developed a real-time radiation monitoring system based on a PIN photodiode gamma ray detector. This compact and portable system consists of a small sensor unit and a Raspberry Pi, with control software developed using an EPICS-based framework. The software is designed to support a wide range of radiation measurements. To evaluate the system's measurement accuracy and reproducibility of the system, the sensor was calibrated with Co-60 sources at the KEK radiation irradiation facility. In addition, radiation measurements were conducted during both operation and shutdown of the KEK electron/positron injector LINAC. This paper presents a detailed description of the system and its control software, along with the results of calibration and measurement tests carried out at the facility.

## INTRODUCTION

We used a PIN photodiode gamma ray detector (developed by Taisei Corporation, manufactured by Yaguchi Electric Co., Ltd., originally designed by Radiation-Watch Co., Ltd., PocketGeiger (TM) [1]; hereafter referred to as the PIN photodiode sensor) to conduct calibration tests using radiation sources as well as radiation measurements at the KEK electron/positron injector linear accelerator [2] (hereafter referred to as the KEK Linac). The PIN photodiode sensor can be connected to a microcomputer (Raspberry Pi Foundation, trademark Raspberry Pi [3]), enabling real-time radiation monitoring. Furthermore, the system is low-cost, compact, and portable.

The PIN photodiode sensor used in this study is an improved version designed to operate under higher dose-rate environments. The sensor size was reduced to 7.5 × 7.5 mm², less than one-thirteenth of the original, thereby lowering its sensitivity. In calibration tests performed in 2024, the upper measurement limit was approximately 83 μSv/h [4]. In the present study, to further explore the capabilities of the PIN photodiode sensor, we attempted to extend the measurable range by changing the resistor used in the system. Calibration tests were carried out using a Co-60 radiation source, and radiation dose measurements were conducted during both operation and shutdown of the KEK injector Linac. This paper reports the details of these measurement tests in 2025.

† itsuka.satake@kek.jp

## SYSTEM CONFIGURATION

The system uses a PIN photodiode sensor as the detector unit, with a Raspberry Pi 4 Model B handling control and data processing. The signal input and output are provided via D-sub connectors, allowing flexible connection between the Raspberry Pi and the sensor. The sensor signal lines from the Raspberry Pi's 40-pin general-purpose input/output (GPIO) header are converted to D-sub connectors. The distance between the Raspberry Pi and the sensor can be adjusted by changing the cable length. Figure 1 shows a photograph of the assembled system with all connections in place.

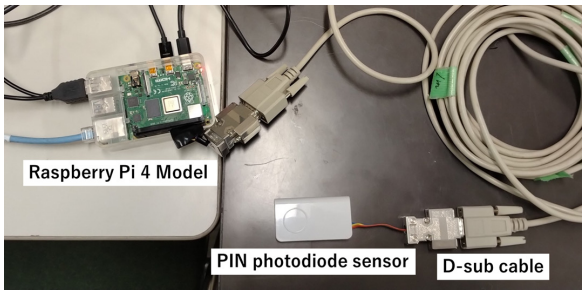



Figure 1: Photograph of the assembled system, showing the PIN photodiode sensor (detector unit) connected to a Raspberry Pi 4 Model B via D-sub cable.

The active area of the PIN photodiode sensor is 7.5 mm². Although the sensor is sensitive to both beta and X/γ rays, a brass plate is placed in front of it to shield beta rays. As a result, the system effectively functions as an X/γ-ray detector in environmental radiation measurements.

## SIGNAL ACQUISITION

One PIN photodiode sensor is connected to a single Raspberry Pi, on which an Experimental Physics and Industrial Control System (EPICS) IOC [5] is running. When radiation is detected, the radiation detection signal output is pulled Low through a 10 kΩ pull-up resistor. This signal is acquired via a GPIO pin and the pulses are counted. Within the EPICS IOC, the number of pulses accumulated over one second is calculated and converted into CPS (Counts Per Second) and CPM (Counts Per Minute). The measured CPM is then multiplied by a conversion factor to obtain the dose rate (μSv/h). The information collected by

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD023

# KEK ELECTRON/POSITRON INJECTOR LINAC SAFETY SYSTEM

M. Satoh†,1,2, I. Satake[1], A. Shirakawa[1], K. Hisazumi[3], T. Kudou[3], S. Kusano[3]
[1]High Energy Accelerator Research Organization (KEK), Accelerator Laboratory, Tsukuba, Japan
[2]The Graduate University for Advanced Studies (SOKENDAI),
Accelerator Science Program, Tsukuba, Japan
[3]Mitsubishi Electric System & Service Co., Ltd (MSC), Tsukuba, Japan

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

At the KEK electron/positron injector Linac, simultaneous top-up injection into four independent storage rings and a positron damping ring has been successfully carried out since May 2019. To ensure long-term stable beam operation under such a complex operational scheme, high availability of the control system is critically essential. In addition, a reliable safety system is also required to prevent personnel radiation accidents and to protect machine components. At the KEK injector Linac, the control system has been developed based on the EPICS framework, which is widely adopted in the accelerator community. However, the safety system was originally developed independently of EPICS. To enhance compatibility between the control and safety systems, an EPICS IOC for the safety system was implemented using the netDev device support. Furthermore, a newly developed beam-operation logic status GUI enables rapid identification of invalid linac operating conditions. This paper presents an outline of the safety system and its recent improvement at the KEK injector Linac.

## INTRODUCTION

The KEK electron/positron injector Linac began operation in 1982 as a dedicated injector for the PF ring, and subsequently has supplied electron and positron beams to rings of different energies, including TRISTAN, PF-AR, and KEKB. Initially, beam injection into each ring was carried out using a time-sharing scheme based on fixed-time injections. However, the demand for top-up injection, which was becoming common in third-generation synchrotron light sources, increased among PF users, and the KEKB ring also requested top-up injection to improve integrated luminosity. In response to this situation, a project team aiming for simultaneous top-up injection into all rings was established in 2004, and in April 2009, simultaneous top-up injection into three rings (the KEKB electron/positron rings and PF ring) was realized. After the construction of a dedicated injection beam transport line for PF-AR in 2012, simultaneous top-up injection into four rings (the SuperKEKB electron/positron rings, PF ring, and PF-AR) and a positron damping ring was achieved in May 2019 [1]. Currently, the injector Linac can deliver electron and positron beams generated at up to 50 Hz to any ring according to prescheduled injection patterns.

To realize simultaneous top-up injection, an event timing system [2], a laser-driven low-emittance RF electron gun [3], a flux concentrator for positron capture [4], high-precision beam position monitors [5], and a pulsed magnet system [6, 7] have been developed and implemented. Various modifications are currently underway to improve injection efficiency into both synchrotron light sources and the SuperKEKB rings [8]. As beam operations have become more complex, the beam control system has also become more sophisticated. Under such a complex operational scheme, ensuring the soundness of both personnel and accelerator equipment in the injector Linac and downstream rings increasingly requires a robust safety system. Although the safety system is constructed independently of the control system, as beam operation logic becomes more complex, it is essential to promptly identify and resolve any failures of beam operation conditions to maintain the operational availability of downstream rings. Therefore, to ensure seamless operation, the compatibility between the safety system and the control system should be enhanced, enabling safety system data to be integrated within the control system framework.

## INJECTOR LINAC SAFETY SYSTEM

### *Overview*

In general, accelerator safety systems are broadly classified into the Personnel Protection System (PPS), which ensures human safety, and the Machine Protection System (MPS), which protects accelerator equipment. During injector Linac operation, which accelerates electron and positron beams, radiation exceeding the permissible dose for humans can be generated inside the tunnel where the accelerator is installed. Therefore, a necessary condition for injector Linac operation is that the PPS requirements are satisfied, thereby ensuring human safety. When human safety is guaranteed and injector operation is permitted, the system is said to be in the "LINAC READY" state. This indicates that the PPS requirements of not only the injector Linac but also all downstream rings permit injector Linac operation.

Figure 1 shows the interlock signal logic for injector Linac operation. The following conditions must be satisfied within the injector Linac PPS:

- The accelerator tunnel is unoccupied.
- All doors involved in interlocks are closed.
- Emergency stop switches are not pressed.
- The beam key at the safety control console is in the "READY" state.
- Emergency stop switches are not pressed.
- The integrated beam charge does not exceed the limit.

The injector Linac can operate only when all of the above conditions are satisfied, 'READY' signals are

†masanori.satoh@kek.jp

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD024

# UPDATING AND OPERATING THE MACHINE PROTECTION SYSTEM AT J-PARC MR

T. Kimura†
J-PARC Center/KEK, Ibaraki, Japan

*Abstract*

The Machine Protection System is a group of devices used to ensure the safety of the accelerator and experimental facility by automatically stopping beam operation and aborting the beam in the event of equipment failure. The Machine Protection System in the J-PARC Main Ring (MR) is called MR-MPS and has been in operation since 2008, when the MR started operation. Development of a new MR-MPS to replace the ageing MR-MPS started around 2020 and was completed in April 2022. New MR-MPS introduced to protect those equipment when the main magnet power supply is renewed and new RF power supply is introduced from autumn 2022. The good results obtained in its operation have led to the completion of mass production of the new MR-MPS, which is now progressively being replaced by the new MR-MPS. This presentation will report on the configuration of the MR-MPS, its EPICS-based operation and the timeline for its complete replacement to be completed in the future.

## INTRODUCTION

The Japan proton accelerator research complex (J-PARC) consists of three accelerators: a 400-Mev LINAC, a 3-GeV rapid cycling synchrotron (RCS), and the 30-GeV Main Ring (MR) synchrotron. The 3-GeV beam from the RCS is delivered to the Materials and Life Science Experimental Facility (MLF), while the 30-GeV beam from the MR is delivered to the Neutrino Experimental Facility (NU) or the Hadron Experimental Facility (HD) [1].

The Machine Protection System is a group of devices used to ensure the safety of the accelerator and experimental facility by automatically stopping beam operation in the event of equipment failure. Also, the MPS unit that integrates the MPS signals from all J-PARC facilities is called J-PARC MPS, whereas the MPS unit that integrates the MPS signals from the four facilities—350BT (the 3-GeV beam from the RCS is delivered to the MR through the 3-50 beam transport line), MR, NU, and HD—is called MR-MPS. The MPS units for J-PARC MPS, LINAC, RCS, and MLF were developed by Japan Atomic Energy Agency (JAEA), while the MR-MPS units were developed by High Energy Accelerator Research Organization (KEK). This paper explains the functions and configuration of MR-MPS and introduces the new MR-MPS developed to update the current MR-MPS.

## MR-MPS

### *Configuration*

The role of MPS is to automatically halt beam operation when equipment failure. Additionally, MR-MPS possesses the capability to request beam dumping to the Beam Abort Dump. The MR-MPS system consists of multiple MR-MPS units, as shown in Fig. 1, connected in a tree-like configuration via an optical network. The MR-MPS unit can accept interlock signals from various equipment via the contact input module, supporting up to 28 channels. Contact inputs are normally closed when the equipment is operating correctly and open when an interlock occurs. The interlock signal input is processed by the FPGA integrated into the CPU module, which then outputs MPS signals for beam stop and beam abort requests via the optical signal input/output module. The MR-MPS unit integrates MPS signals from four facilities—350BT, MR, NU, and HD—by optically transmitting these MPS signals to higher MR-MPS units. The equipment for these four facilities is located in separate areas as shown in Fig. 2. Specifically, the equipment for the 350BT and MR is located in six power supply buildings across three areas on the ground level of the MR tunnel's inner perimeter. Since each building is separated by up to several hundred meters, the MPS signal is transmitted via optical fiber. As shown in Fig. 3, interlock signals are consolidated by installing one to three MR-MPS units in each building.

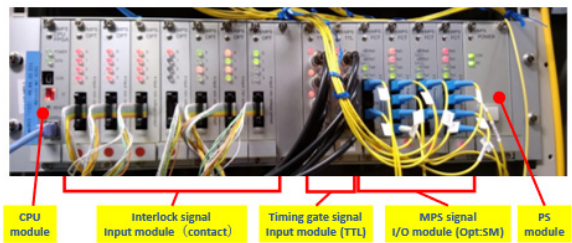



Figure 1: Current MPS unit configuration.

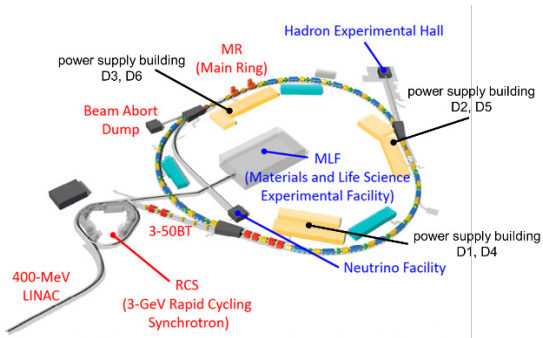



Figure 2: Layout of the J-PARC accelerators and experimental facilities.

† email address kimurata@post.j-parc.jp

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD025

# IFMIF-DONES ADVANCED MACHINE PROTECTION SYSTEM

R. Lorenzo[1,†], C. Carvajal[1], I. Podadera[1], A. Jokinen[2], M. Cappelli[3], E. Botta[4]
[1]IFMIF-DONES Spain Consortium, Granada, Spain
[2]Fusion For Energy (F4E), Garching, Germany
[3]ENEA Frascati Research Center, Rome, Italy
[4]Ansaldo Nucleare, Genova, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The International Fusion Materials Irradiation Facility-DEMO-Oriented Neutron Source (IFMIF-DONES), a cutting-edge accelerator-based neutron source for fusion materials research. IFMIF-DONES Facility Project involves a particle accelerator that produces a deuterium beam of 40 MeV and 125 mA, impacting on a flowing liquid lithium target, generating the neutron source by called nuclear stripping reaction, therefore a robust central control system for safe and efficient operation is required. IFMIF-DONES Central Instrumentation and Control System (CICS) focuses on three groups of systems: Control, Data Acquisition, and Communication (CODAC), responsible for overall coordination and data management; Machine Protection System (MPS), ensuring machine protection; and Safety Control System (SCS), implementing safety functions for personnel and the environment. This work presents the current design focused only on robust Central MPS interlock signals for safe, reliable and efficient protection of the machine. The MPS design considers advanced technology and fast interlock propagation, based on hierarchical, modular, and scalable architecture with built-in redundancy. The MPS system shall be able to respond in a maximum time of 30 µs, acting over the accelerator systems and the lithium systems.

## INTRODUCTION

The IFMIF-DONES facility is currently under construction. The plant will produce a 125 mA deuteron beam, that accelerated up to 40 MeV and achieving a beam shaped to of cross section in the range from 100 x 50 mm to 200 x 50 mm. The beam will then strike on a liquid lithium curtain 25 mm thick, flowing at about 15 m/s. The stripping reactions generate a large amount of neutrons that interact with the materials samples located immediately behind the Lithium Target, on the Test Modules.

The systems devoted to produce and deliver the high-power beam are grouped under (Fig. 1) the *Accelerator Systems (AS)*. The systems related to the Lithium Target management constitute the *Lithium Systems (LS)*. The systems in charge of the irradiation test module(s), the Test Cell and their support systems which constitute the *Test Systems (TS)*. The systems in charge of performing the global control of the Plant are grouped under the *Central Instrumentation and Control Systems (CICS)*. Finally, the *Site, Building and Plant Systems (SBP)* include the buildings and the systems providing power, cooling, ventilation, remote handling of components and services to the other systems.

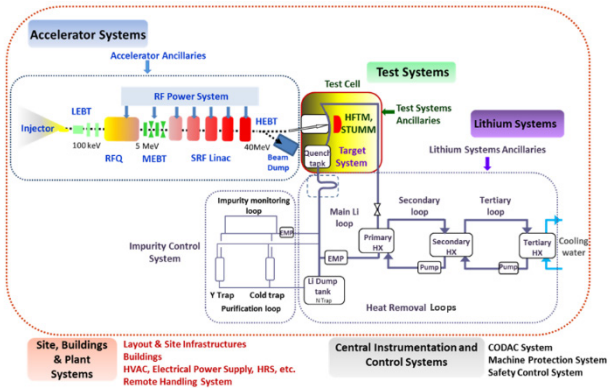



Figure 1: IFMIF-DONES schematic plant configuration.

### *Accelerator Systems*

The main functions of Accelerator Systems (Fig. 2) are to generate and accelerate to an energy of 40 MeV an average current of 125 mA $D^+$ ions in continuous wave (CW) mode and direct them onto the lithium target to produce neutrons by nuclear stripping reactions.

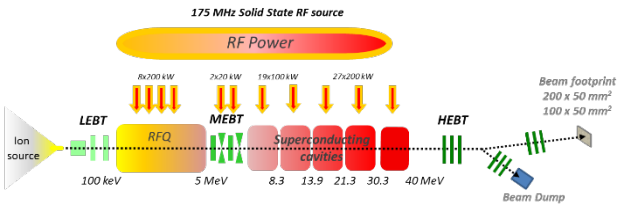



Figure 2: IFMIF-DONES accelerator systems schematic.

### *Lithium Systems*

The main functions of the Lithium Systems are to provide the liquid lithium jet whose interaction with the $D^+$ beam generates the neutron flux for the irradiation of the test modules; and to evacuate the heat power deposited in the lithium by the deuterons. The 125 mA, 40 MeV $D^+$ beam is injected into the lithium target, with a footprint of 200 mm nominal width and 50 mm nominal height.

### *Test Systems*

The main function of Test Systems is to house the Target Assembly, and to allow controlled irradiation of the small steel specimens in the high neutron flux region, just behind the Target Assembly by means of the High Flux Test Module (HFTM) system. Also, to provide sufficient shielding to the surrounding rooms and areas against the intensive

† ruben.lorenzo@ifmif-dones.es

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD028

# DESIGN AND IMPLEMENTATION OF A SCALABLE AND SAFETY-ORIENTED MACHINE PROTECTION SYSTEM FOR THE EuAPS PARTICLE ACCELERATOR

G. Latini, V. Dompè, S. Pioli, Frascati National Laboratories of INFN, Frascati (RM), Italy

*Abstract*

The evolution of research infrastructures in the field of particle accelerators requires increasingly intelligent, integrated and safety-compliant protection systems. In this context, the development of a Machine Protection System (MPS) is underway for the EuAPS project, currently being implemented at the SPARC_LAB facility of INFN National Laboratories in Frascati. The system is designed as a distributed network based on real-time embedded industrial controllers, with FPGAs on cRIO modules and dedicated devices, and is intended to monitor the operational conditions of critical subsystems such as diagnostics, vacuum and safety in real time, responding promptly in the event of interlocks. The prototype is being developed in accordance with the IEC-61508 standard, including the identification of Safety Instrumented Functions (SIFs) and a structured strategy encompassing hardware/software design, implementation, verification and validation. The goal is to provide a prototypal, standards-compliant platform that is scalable and suitable for future machine expansions, such as the EuPRAXIA project. In addition to protection, the MPS aims to enhance operational management and prevent critical failures, contributing to a safer, more autonomous, modular and efficient accelerator infrastructure.

## INTRODUCTION

The EuAPS (EuPRAXIA Advanced Photon Sources) [1] project, under development at the INFN National Laboratories of Frascati (LNF) in collaboration with CNR and the University of Rome Tor Vergata, aims at realizing a compact user oriented X-ray source based on laser driven betatron radiation [2, 3]. Hosted at the SPARC_LAB facility and funded by the Italian PNRR program with over 22 million euros, the infrastructure will provide ultrafast X-ray pulses (femtosecond duration, 1–10 keV photon energy) for applications ranging from material science to biology, industrial diagnostics and cultural heritage studies.

The development of next-generation infrastructures such as EuAPS, also conceived as a step toward the broader EuPRAXIA project [4], requires reliable machine protection systems compliant with functional safety standards, able to monitor critical subsystems in real time and to promptly react in case of abnormal conditions.

In this context, a dedicated Machine Protection System (MPS) is being developed for EuAPS. The system is based on real-time industrial controllers, FPGA modules and dedicated devices. It is designed in compliance with the IEC-61508 functional safety standard.

This contribution presents the design approach, prototypal implementation, and future perspectives of the EuAPS MPS, highlighting its scalability and suitability as a reference platform for next-generation accelerator infrastructures.

## MACHINE PROTECTION SYSTEM

The Machine Protection System (MPS) [5, 6] for EuAPS is conceived as a central element to ensure the safe, reliable, and continuous operation of the accelerator.

Its main objective is to protect critical subsystems, such as vacuum and safety-related devices, from abnormal operating conditions or unexpected failures that could cause severe damage to components or prolonged interruptions of experimental activity.

The system is designed to continuously monitor the machine's vacuum status and interlocks from the controllers; if predefined safety thresholds are exceeded, it can promptly close the vacuum zones by actuating the gate valves.

In this way, not only is hardware damage prevented, but downtime is also reduced, ensuring greater facility availability for users.

Another goal is to provide a scalable and modular platform capable of adapting to future machine developments and expansions, particularly in view of the EuPRAXIA project.

Following the guidelines of the IEC-61508 standard and through the definition of Safety Instrumented Functions (SIFs), the MPS also aims to serve as a reference prototype for the development of standard-compliant protection systems in next-generation accelerator infrastructures.

In addition to protection, the MPS contributes to improving the overall operational management of the accelerator, fostering a safer, more autonomous and efficient infrastructure, where failures are not only mitigated but also anticipated and prevented through structured and continuous monitoring.

## COMMUNICATION ARCHITECTURE

To provide advanced automation and safety capabilities, the core of the MPS is an NI cRIO-9057 controller [7], equipped with a 1.33 GHz Intel dual-core processor, 4 GB SSD, a Xilinx Artix-7 40 MHz FPGA, and configurable I/O interfaces. Programming is performed using the G language (LabVIEW), which translates block diagrams into VHDL code.

The central controller is integrated into a distributed architecture based on a master-slave EtherCAT communication network, ensuring reliability and synchronization in signal transfer. The overall layout, shown in Fig. 1.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD029

# UPDATE OF THE LINAC AND RCS MPS FOR BLM

H. Takahashi[†,1], T. Miyao[2], T. Suzuki[3], Y. Suzuki[4], M. Yoshimoto[1]
[1]Japan Atomic Energy Agency, J-PARC Center, Tokai-mura, Japan
[2]High Energy Accelerator Research Organization, J-PARC Center, Tsukuba, Japan
[3]Mitsubishi Electric System & Service Co., Ltd, Tsukuba, Japan
[4]Total Support Systems Co., Ltd, Tokai-mura, Japan

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Machine Protection System (MPS) for the J-PARC Linac and 3 GeV Rapid Cycling Synchrotron (RCS) consists of the standard MPS modules and the BLM (Beam Loss Monitor) MPS modules. As these modules have been in use since the beginning of J-PARC operation, it is necessary to address their aging. We therefore began designing and developing new MPS modules. Currently, we are constructing parts of the MPS for the Linac and the RCS using the new standard MPS modules.

On the other hand, development of the BLM MPS module has not progressed because the current module includes a comparator function, and the required specifications for the Linac and the RCS differ. Therefore, we decided to separate the comparator function from the BLM MPS module and implement it with different hardware. The MPS for BLM will be updated by combining this hardware with the new standard MPS module.

In this paper, we report on the current status of updating the MPS for BLM at the Linac and RCS, and outline the future plan.

## INTRODUCTION

The Machine Protection System (MPS) has been configured to immediately inhibit the beam in the event of irregularities, in order to minimize damage caused by beam loss, as the J-PARC Linac and 3 GeV Rapid Cycling Synchrotron (RCS) are high-intensity accelerators [1-3]. The MPS for the J-PARC Linac and RCS consists primarily of the standard MPS modules and the BLM MPS modules, which monitor analog signals from the BLM (Beam Loss Monitor) and include a comparator function to detect excessive beam loss. However, the existing modules have been in use since the beginning of J-PARC operation, increasing the risk of malfunction due to aging. Therefore, we have developed new MPS modules and are replacing parts of the Linac and RCS MPS with the new standard MPS modules [4, 5].

On the other hand, development of the BLM MPS module has not progressed because the current module includes a comparator function, and the beam loss detection methods for the Linac and the RCS are different. Therefore, we decided to separate the comparator function from the BLM MPS module and implement it with separate hardware. A beam loss detection method suitable for the Linac will be implemented in the BLM comparator module (NIM), while a different method will be implemented in the BLM digitizer for the RCS. It was therefore decided to proceed with updating the Linac and RCS MPS for BLM by combining these hardware components with the new standard MPS module.

This paper reports on the development status of the comparator modules for the Linac and the BLM digitizer for the RCS and outlines the update plan for the MPS for BLM.

## DESIGN POLICY OF MPS FOR BLM

Beam measurements using various monitors are performed when a timing signal is received. Therefore, if an incorrect timing signal is input, accurate beam measurements become impossible. On the other hand, since the MPS for BLM is a safety system, the situation in which beam loss measurement is impossible must be avoided. Therefore, either continuous beam loss measurement without a timing signal or accurate beam loss measurement through timing signal monitoring is necessary.

In the Linac, it is possible to detect irregular beam loss by monitoring the instantaneous value of the BLM signal for beam pulses with a maximum width of 500 μs, since a beam pulse is accelerated along a single path. Therefore, we decided to monitor the BLM signal continuously without the timing signal and output the MPS signal when the BLM signal exceeds the threshold.

On the other hand, because the RCS accelerates the 400 MeV beam pulse injected from the Linac to 3 GeV in a circular orbit for 20 ms, the BLM signal level differs between beam injection and beam extraction, even during normal beam acceleration. Therefore, the method of monitoring the instantaneous values of the BLM signal adopted in the Linac is not suitable for beam loss detection in the RCS. Following our investigation, we decided to measure the integrated beam loss value from beam injection to extraction, using this measurement to monitor beam loss.

From the following section, details of the updating Linac and RCS MPS for BLM will be provided.

## UPDATE OF THE LINAC MPS FOR BLM

In accordance with the design policy, the current Linac MPS for BLM is configured to input signals from the BLMs into the current BLM MPS module, which has the comparator function.

While the current system benefits from a simplified configuration, the comparator function being included into the MPS module makes it difficult to support new requirements. Furthermore, since the BLM MPS module receives analog signals such as the BLM signal, it is difficult to standardize the interface with standard modules, which has become an issue during the update.

† takahashi.hiroki@jaea.go.jp

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD033

# LIFECYCLE MANAGEMENT TOOLS IN THE DEVELOPMENT OF SNS CREDITED SAFETY SYSTEMS *

K.L. Mahoney†, A. Aher, P.A. Bong
Oak Ridge National Laboratory: Spallation Neutron Source, Oak Ridge, TN, USA

*Abstract*

A new FPGA/PLC based credited Beam Power Limit System (BPLS) ensures beam power to the SNS First Target Station (FTS) does not exceed a safety envelope of 2 MW + 10%. A new SNS beamline, VENUS, uses a single safety PLC to perform safety interlock logic. In order to link and track the extensive documentation required to manage a safety credited system, the commercial tool Reqtracer™ was used to develop and link dozens of documents from the system requirements, specifications, code, down to specific test results. This level of traceability helps to ensure the pedigree of the SNS safety systems remains intact over the life of the system.

## MOTIVATION

The Spallation Neutron Source (SNS) recently completed the Proton Power Upgrade (PPU) project which doubled the available accelerator average power from 1.4 to 2.8 MW [1]. Only 2 MW of this power can be used by the existing first target station (FTS) with the additional 0.8 MW intended for the future second target station (STS). Part of the PPU project was the design of a credited beam power limit system (BPLS) to ensure power to the first target station (FTS) does not exceed its safety basis limit of 2 MW+10% over one minute [2].

In order to calculate beam power, the BPLS needs to independently measure both beam energy and current. Energy is relatively stable and can be calculated in a PLC using dipole magnet measurements. Average current must be calculated from integration of individual fast pulses which is outside the capability of a PLC. The selected architecture is shown in Fig. 1.

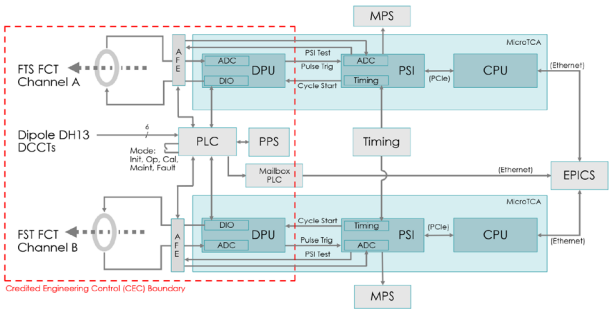



Figure 1: BPLS architecture.

A commercial certified safety PLC is used to calculate beam energy and then set the maximum current threshold based on the maximum power trip setpoint. A certified safety commercial-off-the-shelf fast beam current measurement device is not available. Therefore, best practices for custom designed safety hardware must be used to implement this part of the system. This includes complete specifications, fault tolerance (redundancy) and diagnostics, Verification and Validation (V&V), and strict configuration management. An FPGA based signal processing architecture was chosen for this part of the system.

The BPLS would be the first implementation of an FPGA based credited safety device at SNS. Therefore, it was paramount that design and V&V documentation was complete and auditable by independent reviewers. Requirements needed to be tracked to specifications down through implementation and testing. Test records would become auditable evidence that the system was implemented faithfully. All of this would form a "pedigree" for the BPLS system.

## IMPLEMENTATION

Multiple document types would need to be linked during the project and later during operations. The basic structure and connection between documents are shown in Fig. 2. To facilitate traceability, the Siemens Reqtracer™ product lifecycle management (PLM) tool was chosen early in the project [2].

SNS held in-house Reqtracer™ training for all project engineers to learn the basics of document development. Use of the tool also required training an in-house expert on the system to centrally manage the documents and address issues. This effort added approximately 0.3 full time equivalent (FTE) labor to the project cost. Templates of each major type of document were developed and distributed to engineering staff to facilitate document development.

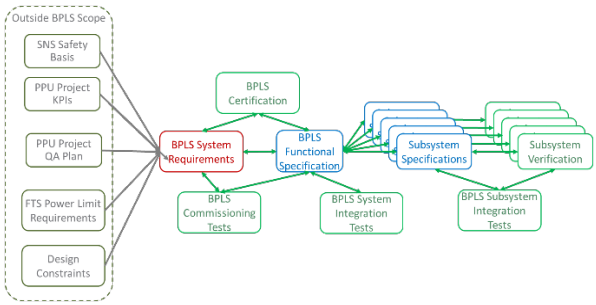



Figure 2: Document types needing to be tracked for the BPLS.

Reqtracer™ uses keywords and phrases to link documents. A requirement or specification in a higher-level document is "covered" in one or more lower-level documents. Figure 3 shows an example where a BPLS accuracy requirement is linked to a functional specification down through a subsystem specification and finally test results.

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under contract number DE-AC05-00OR22725
† mahoneykl@ornl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD034

# DEVELOPMENT OF A PROTOTYPE FAST ORBIT INTERLOCK SYSTEM FOR THE 4GSR

J.-S. Yu[†,1], W.-S. Cho[1], Y. Kim[1], S. Park[1], S. Kang[2]
[1]Pohang Accelerator Laboratory, Pohang, Republic of Korea
[2]Korea Basic Science Institute, Daejeon, Republic of Korea

*Abstract*

This study aims to develop the Fast Orbit Interlock system (FOI) that can quickly determine situations caused by beam deviation during accelerator operation at the 4th generation synchrotron radiation facility. The FOI system is designed to detect minute beam orbital deviations in real time and thereby determine beam deviation during accelerator operation. The hardware and software of the system are currently being designed to implement high-speed data processing, network communication, and real-time control algorithms. Designed to provide higher accuracy and faster response time, performance evaluation will be conducted through prototypes. This study is expected to provide an important technical foundation for quickly stopping by judging the critical situation caused by the failure of the periphery due to beam deviation and radiation emissions above the allowable amount at the accelerator facility.

## INTRODUCTION

Korea-4GSR (4th Generation Synchrotron Radiation source) provides a photon beam that has beam energy of 4 GeV and a stored beam current of 400 mA. Since the beam current is continuously moving in the storage ring, variations in the beam trajectory can damage the banding magnet (BM) and the insertion device (ID). The purpose of the FOI is to protect the device by preventing such a damaging situation and generating a signal to stop the beam in a short time. To find out the trajectory state of the beam, turn by turn (TBT) data about the position and angle of the beam is provided per 375 kHz from the beam position monitor (BPM) located before/after the BM and before/after the ID. The FOI (Fast Orbit Interlock system) collects and analyzes these beam trajectory states, determines the beam deviation situation, quickly outputs the beam dump signal, and transmits the interlock signal to peripheral devices. The state of all data and devices is expressed on the monitor in the operator's room so that the beam's deviation can be notified while driving. In addition, when an interlock occurs due to a beam deviation during the accelerator operation, TBT data before and after the occurrence of beam deviation is collected and provided to researchers to analyze the cause of beam deviation.

## ARCHITECTURE

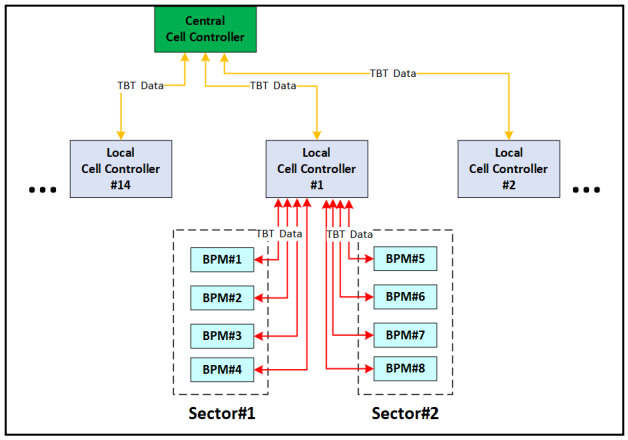



Figure 1: Connection with central CC, local CC and BPM.

Figure 1 represents the basic connection of FOI, and the controller in charge of the cell of 28EA of 4GSR is classified into Local CC (Cell Controller) and Central CC to show the connection with the Beam Position Monitor (BPM) for each sector [1]. The FOI of 4GSR is installed to cover the 4GSR storage ring to monitor the beam trajectory of the entire storage ring, and is located for each 2EA sector in the 28EA sector of the storage ring in consideration of connection and peripheral devices. Each local CC has the function of transmitting/receiving data to/from the Central CC and transmitting an interlock signal to the peripheral device in the event of an Orbit Interlock.

Figure 2 represents interconnection with Timing system, BPM, MIS (Machine Interlock System), EPICS (Experimental Physics and Industrial Control System) and others. FOI's operation is synchronized with 375 kHz timing event from 4GSR Timing System. It is designed based on reliable and proven technology for this connection [2]. All FOI controllers have EPICS IOC to support EPICS protocol that is used on the control network for operation. All signals generated during operation are connected to the EPICS control network for expression in the driver's compartment through EPICS, allowing the user to monitor and command the status of the FOI. Additionally, Phoebus OPI will show data for BPM during 1,000 ms, after/before the orbit interlock.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD035

# DESIGN OF THE MACHINE INTERLOCK SYSTEM FOR KOREA-4GSR

Y. Kim†, J. Yu, S. Park
Pohang Accelerator Laboratory, Pohang University of Science and Technology
Pohang, Kyungbuk, Republic of Korea

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

MIS (Machine Interlock System) for Korea-4GSR is a system designed to prevent operations that may cause damage to related devices when faults occur in any of the devices configured throughout the accelerator during its operation. It is an essential system of accelerator operation. In the MPS (Machine Protection System), MIS mainly covers signals with relatively slow response speeds among various classifications of signals. For signals generated by an interlock situation, which are recognized by the central device from local devices and transmitted to the final destination, differences in speed exist depending on the route. However, MIS is designed with minimized signal routes so that the difference in speed by signal will be minimal. The controller of MIS is a PLC (Programmable Logic Controller), which is widely used in various fields of industry and experimental physics. Since various PLC guidelines are provided for securing system reliability and responding to failures, configuring a more robust system is more accessible with PLC than with other types of controllers. This poster presents the design to minimize beam downtime due to device faults and maximize device protection.

## INTRODUCTION

MIS is the system that restricts operations so that operation can be performed according to established procedures and minimizes damage(mitigation) to the device in the event of a device failure. Therefore, the goal of MIS is to maximize reliability and availability to ensure normal operation under any situations. To achieve this, MIS adopted PLC CPU redundancy and a redundant ring topology between the PLC CPU and remote IOs. 4GSR will use EPICS as middleware and MIS will also be integrated with EPICS to enable monitoring of MIS status, data storage and control of MIS through EPICS in the main operation centre. In the fault of the field device during beam operation, only the sector related to the sector where the field device is located will be given the trip signal, maximizing the availability of the storage ring. The signal wiring is being designed to be able to respond to failures such as a short circuit or signal transmission failure by wiring the signal to the Source Output and Sink input (PNP type, normal high).

## HW ARCHITECTURE

Figure 1 shows the hardware architecture of MIS. MIS is designed to protect 4GSR by consisting of three layers: central layer, local layer, and Device & System layer.

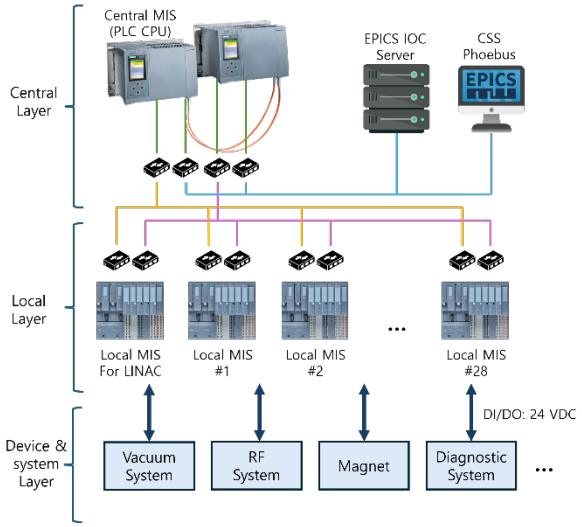



Figure 1: Hardware architecture.

### *Central Layer*

This layer monitors and analyses collected data from Local Layer by purpose (sector, function) protect devices. It is also connected to the EPICS network, allowing the main operation centre to monitor the status of the MIS and perform controls such as device status reset.

### *Local Layer*

This layer collects status signals provided by field devices by purpose (area, function) and transmits them to the Central MIS. It also receives the result signals analysed by the Central MIS and sends them to the Device & Equipment Layer.

### *Device & System Layer*

This layer protects itself and transmits the status to other related systems when it detects a problem by watching itself. It also receives interlock signals from related systems for control.

## SW ARCHITECTURE

Figure 2 expresses the SW architecture of MIS, showing the relationship between the software used in each layer.

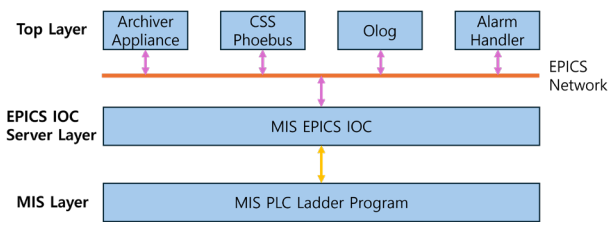



Figure 2: Software architecture.

† kim.yh@postech.ac.kr

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD036

# FACILITY MONITORING SYSTEM

N. R. Waters†, J. Nguyen
SLAC National Accelerator Laboratory, Menlo Park, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

As the SLAC FEL Beam enters its new era of LCLS-II SC beam and plans ahead for LCLS-II-HE, interest has grown in the area of facility monitoring. A system is needed for the early detection and response to adverse conditions affecting critical infrastructure. An engineered solution aims to minimize the impact on facility operations due to overheated or over humidified enclosures, low cooling water flow rates, high cooling water temperatures, and cooling water leaks. The newly deployed monitoring system uses two main hardware systems: Raritan and Beckhoff. The former consists of a suite of sensors distributed throughout an area, up to 30 meters apart, all connected to a single controller, which itself occupies only 1 rack unit. The latter utilizes a Beckhoff EK9000 Bus Coupler which allows Ethernet networks to connect to EtherCAT terminals. The EtherCAT terminals utilize 4-20mA signals which can easily reach across large areas. These signals are then communicated through software to the user in three ways. The first relies on EPICS Process Variable Alarm Status and the Next Alarm System (NALMS). The second uses the EPICS Archiver, Grafana, and Slack Chat. The last alerting mechanism, a multi-functional LED alarm annunciator, has been installed to notify staff with audio and visual signals.

## INTRODUCTION

An engineered system was created at Linac Coherent Light Source (LCLS) with the goal of detecting abnormal conditions in the facility and initiating a response to those conditions by LCLS staff. This system does not include any automatic hardware response, as a result, a response from staff is crucial. Achieving this response is an important consideration in our design principles which were:

- Highly Maintainable and Extensible System
- Alerting Mechanisms
- Identifiable Incident Locations

With these design principles in mind, a set of hardware and graphical interfaces were chosen. These elements when combined, form a system which affords a user the ability to detect and respond to any malfunction at LCLS quickly. Next, we explore the hardware chosen to support these design choices.

## SENSOR HARDWARE

A highly maintainable system is important to reduce the likelihood of missed malfunctions and ignored alerts. This kind of system is characterized by a high uptime and low occurrence of false positives. The latter is of particular importance to avoid degrading the users' trust in the system.

Extensibility was also considered to be important in the hardware selection of the system. Staff members with the greatest knowledge of critical failure points may have little to no understanding of controls. To remedy this issue, the hardware selection was made to enable staff members to add additional monitoring devices on their own.

With these principles in mind, the first and main piece of hardware chosen was the Raritan Smart Rack Controller [1], as shown in Fig. 1.



Figure 1: Raritan Smart Rack Controller (SRC).

This controller was chosen for its capability of supporting multiple types of sensors over a simple Ethernet interface. Each sensor is daisy-chained with its origin at one of 8 different ports. Each port is capable of supporting up to 12 sensors. The types of sensors chosen are illustrated in Figs. 2-4.

The DX2-T1H1, in Fig. 2, is utilized to measure temperature and humidity conditions in our instrument server racks. These server racks house an extensive array of hardware such as PLCs, Timing Servers, EPICS hosts, and camera hosts. The older server racks do not have temperature control and are acutely affected by the closing or opening of rack doors. Our newer server racks receive cooling water and incorporate active cooling. This sensor allows LCLS to detect an overheat condition because of a closed door or the failure of a server rack cooling system.



Figure 2: DX2-T1H1 temperature and humidity [2].

The DX2-T1, in Fig. 3, is utilized to measure ambient temperatures in the general instrument space and the differential temperature between the internal and external environments of the laser enclosures. Together, sensor measurements enable laser staff to correlate environmental data with laser pointing drifts.

* Work supported by U.S. D.O.E. Contract DE-AC02-76SF00515
† nrw@slac.stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD037

# THE OPERATION MODE DESIGN FOR SHINE MACHINE PROTECTION SYSTEM (MPS)

C. Yu, Y. Yan, G. Chen, J. Ding, J. Chen, Y. Gao, X. Liu
Shanghai Advanced Research Institute, Chinese Academy of Sciences, Shanghai, P.R. China

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Shanghai HIgh repetitioN rate XFEL and Extreme light facility (SHINE) is a large-scale scientific facility under construction in China. Due to 1 MHz repetition rate and 3 thousand kilometers length of SHINE, the accelerator interlock system should be a large system with response speed in microsecond magnitude and capable of processing tens of thousands of signals simultaneously. In order to satisfy the requirements of different segments commissioning and further reduce the probability of misoperations, the design of operation mode is introduced in MPS. The operation mode can achieve mode and logic switching in the interlock system through "one click". Different logic can be automatically matched in different modes. The mode switch and management is controlled by the main station comprehensively, and it can be distributed to various sub-stations of the system. According to the structure and debugging requirements of SHINE accelerator, the accelerator has been designed in segments in the operation mode design. The injector section is divided into two segmented modes based on the electron gun and the injector section. The linac section is divided into five segments according to the position of the beam bump, including BC1 (bunch compressor 1), BC2 (bunch compressor 2), break, the branch line, and the total linear accelerator.

## BACKGROUND

Shanghai High Repetition rate XFEL and Extreme light facility (SHINE) is a large-scale scientific facility designed as a high-repetition-rate hard X-ray free-electron laser. It is capable of generating X-ray pulses with the maximum repetition rate of 1 MHz and electron energy of 8 GeV [1]. With broad applications across numerous fields, SHINE serves as a critical tool for research in physics, such as chemistry, materials science, life sciences, and beyond. SHINE is currently designed with two undulator beamlines, with a total length of approximately 2 km from the injector to the end of the undulator section. A general schematic of the SHINE facility is shown in Fig. 1. The accelerator interlock system provides timely and effective protection for critical equipment in the event of a failure, preventing machine damage. Additionally, the interlock system is responsible for transmitting alarm information to the upper-level control system to alert operators for appropriate response and handling. It also offers functions such as machine status indication and historical record query [2].

Due to its high repetition rate and high stability, SHINE imposes stringent requirements on the interlock system: the single-node response time must be less than 10 μs, and the inter-node response time should be under 100 μs. Traditional PLC-based interlock systems only achieve response times on the order of milliseconds. Hence, a fast interlock system employing FPGA-based technology is necessary. To distinguish it from the fast interlock system, the traditional PLC I/O module-based interlock system is referred to as the slow interlock system.

Furthermore, with a total length of 3,100 meters for the SHINE facility, the number of interlock nodes, the quantity and variety of signals, and the geographical distribution have all significantly expanded. The number of interlock stations is expected to reach approximately 100. This scale introduces challenges in the transmission and sharing of specific signals within a large-scale system architecture, which must be addressed in the interlock system design.

To meet all these requirements, SHINE needs to implement an interlock system based on a combined FPGA and PLC architecture. This includes the design and development of signal interactions between stations, incorporating operational mode management of the machine [3].

## SYSTEM DESIGN

The SHINE interlock system adopts a system solution based on “mature hardware boards coupled with fully independent development of interlock functions and software.” The development scheme for the fast interlock system has been determined utilizing Yokogawa's FPIO (Field Programable High-speed I/O) modules. The FPIO module is a high-function I/O module installed and used within Yokogawa’s FA-M3/e-RT3 platform, offering I/O response times on the order of microseconds.

Equipped with a built-in FPGA, the module allows users to develop and compile FPGA functions internally, providing open access for implementing custom logic circuits. Each module supports 24 input channels and 24 output points (5–24 V).

The advantage of mounting the fast interlock modules on the PLC baseboard is that the data communication between the fast and slow interlock systems via the IO bus can be implemented. By developing code for both the CPU and the FPGA sides, bidirectional data reading and writing between the fast and slow interlock systems can be achieved. This significantly facilitates information exchange and sharing between the two systems, effectively resolving the challenge of data integration between fast and slow interlocks.

The FPGA module supports register-based data transfer for communication with the PLC CPU. Through the development of interface programs, it is possible to read from the CPU side all input and output points, the status

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD038

# DEVELOPMENT AND OPERATION OF THE NEW BEAM INTERLOCK SYSTEM OF RIKEN RI BEAM FACTORY

M. Komiyama[1,†], M. Fujimaki[1], A. Kamoshida[1,3], M. Hamanaka[2], M. Nishimura[2], R. Koyama[2], K. Kaneko[2], H. Yamauchi[1], A. Uchiyama[1], N. Fukunishi[1]
[1]RIKEN Nishina Center for Accelerator-Based Science, Wako, Japan
[2]SHI Accelerator Service Ltd., Tokyo, Japan
[3]National Instruments Japan Corporation, Tokyo, Japan

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

A successor system to the existing machine protection beam interlock system for the RIKEN Radioactive Isotope Beam Factory (RIBF) using CompactRIO, a platform developed by National Instruments, has been under development since 2021. The current beam interlock system for RIBF, RIBF-BIS, consists of 10 I/O stations distributed throughout the RIBF facility. It monitors signals from such sources as the magnet power supplies and vacuum equipment, as well as the beam loss detected by the baffle slits in the cyclotrons, and it stops the beam supply just downstream of the ion source approximately 15 ms after detecting a signal anomaly. The successor system, RIBF-BIS2, inherits the logic of RIBF-BIS, while improving the response time by a factor of approximately 100. During the development phase of RIBF-BIS2, its I/O stations are being installed in parallel with those of RIBF-BIS to ensure operational safety. Test operations are initiated sequentially for each station once its installation is completed. The installation of 80 % of the I/O stations of RIBF-BIS2 has been completed. It now can stop the beam service approximately 130 μs after detecting a digital input signal anomaly, and approximately 470 μs after detecting an analog input signal anomaly. This response speed is sufficient to ensure the safety of the RIBF equipment against the beam intensity currently supplied by the RIBF. In 2025, the remaining 20 % of the I/O stations will be installed, enabling RIBF-BIS to be updated to RIBF-BIS2. During the test operation period, several improvements to the system, including enhancing the interlock logic and upgrading the station setup to make maintenance easier, were made. In addition, we have improved the RIBF-BIS2 front-end circuit installed to shorten the time from the detection of an abnormality in the digital input signal to stopping the beam supply was improved. Moreover, the power consumption was reduced to one-fifth while maintaining equivalent performance.

## INTRODUCTION

The RIKEN Radioactive Isotope Beam Factory (RIBF) at the RIKEN Nishina Center for Accelerator-Based Science is a complex accelerator facility, comprising two linear accelerator systems (one of which includes a superconducting linac), four normally conducting cyclotrons, and the largest superconducting cyclotron in the world. At RIBF, various radioactive isotopes (RIs) can be delivered as a beam with optimized energy for each experiment by adjusting combinations of multiple ion sources and accelerators. RIBF provides the most intense RI beams in the world across the entire atomic mass range through fragmentation or fission of high-energy heavy ions [1]. Almost all the components of the RIBF accelerator complex are controlled by the Experimental Physics and Industrial Control System (EPICS) [2], and all essential operation datasets of EPICS and other control systems have been integrated into the EPICS-based control system [3]. The RIBF, safety system is divided into two independent subsystems: one for protecting personnel [4] and another for protecting the accelerator and beamline components from hardware instabilities, troubles, etc. (the machine protection system). The machine protection system operates in three configurations, depending on the type and location of the experimental facility [5]. The beam interlock system for RIBF (RIBF-BIS), established in 2006, is the oldest and most extensive of the three machine protection systems. It covers one of the two linear accelerator systems, four of the five cyclotrons (including the superconducting cyclotron), and their associated beamlines. The other two systems, operational since 2020, manage the remaining cyclotron and linear accelerator, including the superconducting linac, along with their respective experimental facilities.

## DEVELOPMENT AND IMPLEMENTATION OF THE NEW BEAM INTERLOCK SYSTEM

### *Logic of RIBF-BIS*

Figure 1 shows a schematic diagram of the interlock logic of RIBF-BIS. This logic is common to all machine protection systems in RIBF. RIBF-BIS monitors digital signals that indicate the status of various equipment associated with both the accelerator and the beamlines. In addition, it continuously monitors the beam loss measured by baffle slits installed at the entrance and exit of the magnetic channels of the cyclotrons. When a beam loss exceeds the limit value specified for each baffle slit, RIBF-BIS issues an abnormal signal. When an abnormality is detected in the signals, RIBF-BIS responds by sending a signal to a beam chopper, which immediately deflects the beam just downstream of the ion source being operated. This response time is referred to as “the system response time”. Simultaneously, the system inserts a beam stopper (Faraday cup) upstream of the faulty component. This dual action ensures

† misaki@riken.jp

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD039

# BL31-FaXToR, Hard X-ray Micro-Tomography and Radiography at ALBA: Current status and ongoing improvements

S. Wohl, J. Aguilar, F. Becheri, R. Homs, X. Mercadal, A. Ollé, Z. Reszela,
N. Serra, O. Vallcorba, F. Cova, C. Organista, A. Patera
ALBA Synchrotron, Cerdanyola del Vallès, Spain

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

BL31-FaXToR is the hard-X-ray micro-tomography and radiography beamline at the third-generation ALBA synchrotron. It enables 3D imaging with sub-second temporal resolution under either monochromatic or white-beam conditions. The control system is based on Tango, Sardana and Taurus, providing an efficient, distributed Python environment with full user access to hardware via both graphical and command-line interfaces. This paper outlines the BL31-FaXToR control system architecture, presents implementation examples, and discusses upcoming planned features.

## BL31-FAXTOR

BL31-FaXToR is fed by an in-vacuum multipole wiggler installed at port 31 of the ALBA storage ring. The main optical element is a Double Multilayer Crystal Monochromator (DMM), chosen to provide tunable energy and quantitative 3D analysis of tomografic datasets [1, 2, 3]. Other main components in the optics hutch are: set of filters in the multifilter unit acting as beam attenuators, two blade vertical slits, diagnostic elements and bremsstrahlung attenuators. The experimental hutch includes a beam conditioning table with four blade sample slits, a second CVD diamond window, and an in-house–developed fast shutter that ensures precise timing during dynamic acquisitions and prevent radiation damage to sensitive samples. To minimize air absorption at lower X-ray energies, movable pipes guide the beam safely from the conditioning table to the tomography stage, located 36.5 m from the source.

### *Endstation*

The endstation hosts two microscopes equipped with four cameras, which can be positioned either close to the sample for absorption measurements or at longer distances to exploit phase propagation in soft materials. Configurations can be switched easily depending on user requirements, offering a flexible design that supports multiscale imaging (Fig. 1).

This setup enables simultaneous access to different spatial and temporal resolutions, facilitating experiments that demand both high throughput and dynamical studies with temporal resolution better than one second. The high-speed detector has been designed to provide live reconstructions during scans with minimal latency [4], offering immediate feedback to users and paving the way for time-resolved in-situ experiments.

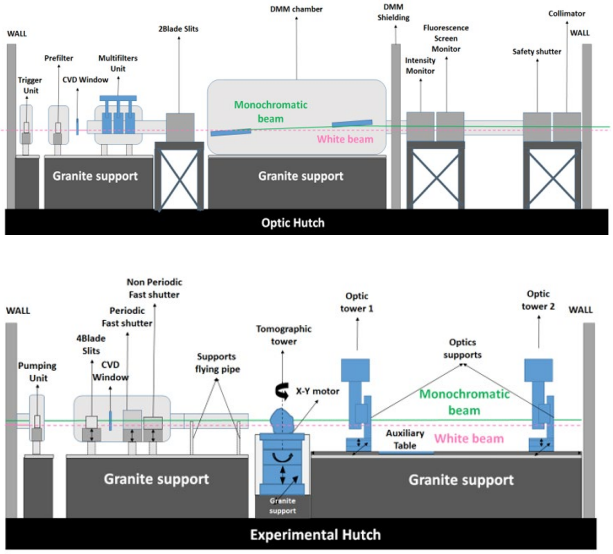



Figure 1: BL31-FaXToR layout of the optical path on the optic (top) and experimental (bottom) hutches. X-ray beam enters from the left side. Both MB (green) and WB (pink) are allowed [3].

## CONTROLS SYSTEM ARCHITECTURE

The BL31-FaXToR control system follows the common architecture adopted for all ALBA beamlines, designed to maximize similarity and reuse. Whenever possible, the same hardware and controller types are employed across subsystems such as vacuum, motion, and diagnostics, which in turn allows for shared software interfaces. As part of the Tango community [5], we benefit from existing device server developments while contributing new ones when required.

Sardana serves as the common SCADA layer across ALBA [6, 7], making control systems comparable and simplifying the implementation of complex scans. Although specific hardware controllers may differ between beamlines, their software access remains standardized. The Taurus [8] framework further supports this approach by providing consistent graphical user interfaces, enabling development to be shared across projects. For exclusive hardware, such as the FaXToR goniometer, dedicated Tango device servers have been designed and integrated so that higher-level control layers can access them through standard Taurus widgets or Sardana macros (Fig. 2).

High-speed data acquisition is a key requirement at FaXToR. The detector subsystem uses multiple frame grabbers in combination with an IBM Storage Scale clustered file

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD040

# MICRO-TCA-BASED DATA ACQUISITION FOR APS RADIO FREQUENCY SYSTEMS

T. Madden*, T. Berenc, E. Chandler, S. Veseli, Y. Yang
Argonne National Laboratory, Argonne, IL, USA
K. Vodopivek, Oak Ridge National Laboratory, Oak Ridge, TN, USA
J. Breeding, Cutting Edge Communications, LCC, Huntington Beach, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

A data acquisition (DAQ) system has been developed for Radio Frequency (RF) systems at the Advanced Photon Source (APS). The hardware is based on the micro-TCA platform and includes FPGA carrier boards equipped with FMC cards featuring RF ADC and DAC capabilities. These components interface with a Linux blade via PCIe over a backplane. The software, built on the EPICS framework, facilitates control and real-time DAQ of RF systems. The EPICS DAQ software streams data to distributed services on the APS network. The dataflow is presented, starting at the RF hardware, through the FPGA to the Linux blade, through the EPICS software, and finally to the network services. This project was a collaborative effort involving industry and multiple institutions.

## INTRODUCTION

The Advanced Photon Source (APS) at Argonne National Laboratory is an X-ray light source with four accelerator machines: a Linear Accelerator, Particle Accumulator Ring (PAR), Booster Synchrotron, and Storage Ring [1]. The APS was upgraded to increase its brightness, and a new Data Acquisition System (DAQ) was developed to collect data from all accelerator machines. The DAQ uses numerous EPICS IOCs that interface with accelerator hardware to collect data and transmit it over the computer network via the PVAccess protocol [2].

DAQ hardware has been installed for RF systems in the Particle Accumulator Ring (PAR), Booster, and Storage Ring. The PAR has two RF DAQ systems for its two cavities and RF systems. The Booster has a single DAQ system for collecting data. Finally, the Storage Ring will have 12 Solid-State Amplifier (SSA) RF systems in the future, each with a DAQ. It currently has DAQs in each Klystron RF system and its Bunch Lengthening Cavity [3, 4].

## MICRO-TCA PLATFORM

The RF DAQ systems are based on a Micro-TCA platform manufactured by Vadatech [5]. This platform consists of a Micro-TCA crate that houses an FPGA carrier card. This card carries a daughter board housing RF electronics for acquiring data from the RF systems. The crate also houses a Linux blade that communicates with the FPGA carrier over a PCIe bus on the crate's backplane. RF data is streamed from the FPGA carrier card to the Linux blade over this PCIe backplane.

Refer to Fig. 1 for a diagram of the entire DAQ system. The three accelerator machines at APS that have DAQ systems installed are the PAR (a), Booster (b), and the Storage Ring (c). In Fig. 1, items "d," "e," "f," and "g" are Micro-TCA systems, each complete with a Linux blade, FPGA, and RF electronics. These systems stream DAQ data to the computer network, marked as "i" in Fig. 1.

### *Software Running on the Linux Blade*

The Linux blades run both control IOCs for controlling the RF hardware and DAQ IOCs for gathering data from the RF systems. The reason for separating control and data acquisition into separate IOCs is to increase reliability, ensuring that if the DAQ IOC crashes, the control IOC is still operational. The IOCs interface with the FPGA hardware through a custom Linux driver that uses the PCIe bus. The FPGA hardware appears in the /dev directory as a device file, and the EPICS IOCs simply open this file to perform their DAQ and control functions.

## NETWORK INFRASTRUCTURE

The Linux blades in the Micro-TCA systems have two network interfaces to separate EPICS control and DAQ data into two distinct networks. The reason for this separation is to prevent high-bandwidth DAQ data from corrupting the EPICS control communication. The two networks are shown as "h" and "i" in Fig. 1. The EPICS control network, item "h" in Fig. 1, carries data for controlling and monitoring accelerator systems. Control screens, based on MEDM (item "o" in Fig. 1), connect to this network to provide a user interface for system control [2].

### *RF DAQ Server*

The RF DAQ server, shown as item "j" in Fig. 1, has two primary functions. The first is to store DAQ data to disk from all the RF DAQ systems. This server is on a virtual LAN dedicated exclusively to RF DAQ data. A file service application captures the DAQ data, which is in the PVAccess protocol, and streams it to a local disk. A set of constantly running Python scripts then monitors the local file system for new DAQ data and transfers it from the RF DAQ server to a main cluster called "Curzon," shown as item "m" in Fig. 1. The network connecting the RF DAQ server to the Curzon cluster is a high-speed "backbone" network, shown as item

* tmadden@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD042

# DEVELOPMENT OF A HIGH-PRECISION ELECTROMETER AT NSLS-II*

K. Ha†[1], Y. Tian[1], J. Mead[1], S. Kongtawong[1], H. Yan[1], D. Gavrilov[1],
J. Keister[1], J. Sinsheimer[1], D. Padrazo[1], S. Wilkins[1], R. Lee[2,3], C. Nam[2,4]
[1]National Synchrotron Light Source II, Brookhaven National Laboratory (BNL), Upton, NY, USA
[2]BNL, High School Research Program, Upton, NY, USA
[3]Stony Brook School, Stony Brook, NY, USA
[4]Earl L. Vandermeulen High School, NY, USA

*Abstract*

A new high-precision electrometer has been developed at NSLS-II to support X-ray Beam Position Monitoring (XBPM) and current measurements across six ranges (100 nA–10 mA) with 200 fA resolution. The third-generation design features a low-noise analog front end, 20-bit ADCs, and FPGA-based signal processing optimized for beamline current readout. The control system, integrated with EPICS via a hybrid IOC using quadEM and pscDriver, enables monitoring and feedback. We describe the NSLS-II electrometer's hardware architecture, EPICS integration, and performance in beamline environments.

## INTRODUCTION

The National Synchrotron Light Source II (NSLS-II) is a state-of-the-art 3rd-generation electron storage ring that began beam operation in 2015. Currently, the machine delivers a 500 mA top-off stored beam current with 29 operational beamlines, and is actively expanding with the construction of NEXT-II and NEXT-III beamlines [1]. Modern X-ray beamlines demand exceptional stability and precision in beam delivery to support advanced experiments such as coherent diffraction imaging, in situ spectroscopy, and nanofocused X-ray scattering. Achieving and maintaining sub-micron beam stability requires accurate, high-resolution, and low-noise current measurements from beamline detectors and beam position monitors (XBPMs). Even small drifts or instabilities in the beam can degrade data quality, reduce measurement repeatability, or affect sample exposure conditions. To meet these performance requirements, we developed a high-precision, high-stability electrometer system known as NSLS2EM, which provides fast and accurate readout of photocurrents from XBPMs or other beamline diagnostics. The system features a 4-channel current input with a 20-bit ADC, programmable gain stages, and a wide input range from ±100 nA to ±10 mA. It includes a real-time data averaging engine that supports fast acquisition (FA) rates of up to 10 kHz, enabling both real-time monitoring and closed-loop feedback control. The EPICS IOC simultaneously uses processed data for user readout and by embedded PI controllers for beam stabilization. The integrated FPGA firmware supports flexible X-ray beam position calculation (square or diamond geometry), timestamped event synchronization via the facility timing system, and high-speed DAC outputs for feedback. The EPICS IOC is based on a hybrid of the pscDriver [2] and quadEM frameworks [3], enabling real-time control, data streaming, and HDF5-based logging. We present a technical overview of the NSLS2EM hardware and firmware, including beamline integration examples such as the Hard X-ray Nanoprobe (HXN) beamlines, to demonstrate the local feedback and global feedback system.

* Work supported by DOE contract No: DE-SC0012704
† kha@bnl.gov
Not Export Controlled.

## SYSTEM DESCRIPTION

Figure 1 shows NSLS2EM, a high-precision FPGA-based electrometer built for advanced beamline diagnostics and fast feedback at NSLS-II. The instrument integrates a custom low-noise analog front end, high-speed digital processing in an FPGA, and a modular EPICS IOC interface for reliable, scalable operation across diverse experimental beamlines. At its core is an Avnet PicoZed 30 System-on-Module combining a Xilinx Zynq-7030 (dual-core ARM Cortex-A9) with programmable logic, providing a flexible platform for real-time acquisition, processing, and feedback control. The analog front end comprises four independent current channels; each uses an ADA4350 transimpedance amplifier followed by a 20-bit LT2378 ADC. Six programmable gain settings cover a dynamic range of ±100 nA to ±10 mA, with resolution down to 200 fA. A daughterboard interface adds four voltage-input channels for monitoring auxiliary signals such as detector bias or reference voltages. For actuator control and real-time feedback, the system provides four 16-bit bipolar DAC outputs (±10.8 V, AD5754) and a 16-bit AD5060 DAC for detector biasing (±12 V), all exposed as EPICS PVs. Additional GPIO lines supply programmable clock, interlock, and trigger signals with configurable frequency and duty cycle, enabling adaptation to beamline-specific requirements.

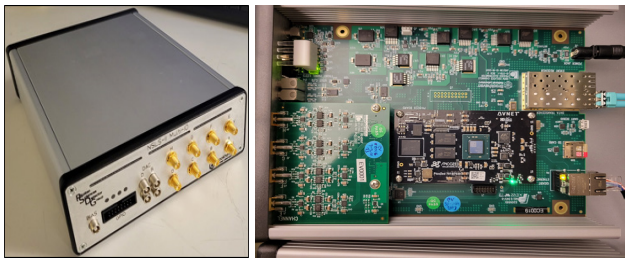


Figure 1: NSLS-II electrometer, left: hardware top view, right: inside chassis and board top view.

The NSLS2EM hardware includes:

- PicoZed 30 Field Programmable Gate Array (FPGA) and SoM
  - 20-bit 4-channel ADC for current signals

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD043

# TOWARDS NEXT-GENERATION INSTRUMENT CONTROL FRAMEWORK FOR NEUTRON EXPERIMENTS AT MLF, J-PARC

K. Moriyama[1], H. Hasemi[2], Y. Inamura[2]
[1]Comprehensive Research Organization for Science and Society (CROSS), Tokai, Japan
[2]Japan Atomic Energy Agency (JAEA), Tokai, Japan

*Abstract*

At the Materials and Life Science Experimental Facility (MLF) of J-PARC, many neutron instruments currently operate using IROHA2, the standard instrument control framework. After more than a decade of operation, architectural obsolescence and diversifying user requirements have revealed its limitations. To address these issues, we initiated the development of a next-generation instrument control framework as a post-IROHA2 system in 2023, which is currently in the prototype phase.

The prototype adopts a data-centric distributed architecture with Redis as its backbone, integrating heterogeneous back-end systems through a unified Device Manager. Asynchronous communication via Redis Pub/Sub ensures loose coupling among components, and a hybrid cloud configuration enables secure remote experiments while providing a foundation for future integration with cloud-based AI services.

This paper outlines the conceptual framework, the current status of the prototype system, and the remaining technical challenges toward full-scale deployment.

## INTRODUCTION

At J-PARC MLF, many neutron instruments have adopted IROHA2 [1] as the standard framework for instrument control. IROHA2 is a distributed control system composed of multiple server components that together provide device-level control, measurement management, automated sequencing of experiments, and unified monitoring of instrument states. IROHA2 also offers a web-based user interface through which users can access and operate instruments.

Although continuous modifications and feature enhancements have been made, more than ten years have passed since its initial development. Architectural obsolescence and limitations in scalability and flexibility have become increasingly evident, making it difficult to address emerging requirements with incremental upgrades.

In particular, IROHA2 suffers from several shortcomings:

- Monolithic server components that hinder maintainability and extensibility.
- Synchronous request/response communication that propagates failures and disrupts ongoing system operations.
- Limited flexibility in integrating data acquisition and device control within a unified framework.
- Proprietary device modules with low standardization and poor interoperability.
- Restricted remote access due to legacy architecture incompatible with current facility security policies.

To overcome these issues, we have initiated the development of a next-generation instrument control framework (post-IROHA2) with the following key requirements:

- **Data-centric architecture:** Unified management of experimental parameters and metadata through a central database.
- **Microservice architecture:** Independent, loosely coupled services for greater extensibility.
- **Standardized device control:** Adoption of EPICS [2], the de facto standard for device control, replacing proprietary modules to ensure interoperability.
- **Remote access support:** A cloud-based front-end to securely and efficiently operate and monitor experiments.
- **Cloud integration:** Native support for cloud environments and advanced services such as AI.

## ARCHITECTURE

### *System Overview*

The system adopts a data-centric distributed architecture with Redis [3], referred to as the Experimental Information Store (EIS), as its central hub. Redis provides both a key–value store and a messaging backbone, managing parameters, states, events, and Run Info in an integrated manner, and forms the foundation for the system's distributed design. The overall structure of the system is illustrated in Figure 1 (see Figure 1 on the following page).

This architecture is organized into three parts:

- **Interface and Monitoring:** User Interface (Web GUI/CLI) and Log Server (OpenSearch).
- **Measurement Control:** Measurement Manager as the core of experiment orchestration, and Authentication Server for user and group authorization.
- **Device Control:** Device Manager and diverse back-end systems including EPICS IOCs, legacy IROHA2 modules for existing MLF devices, and DAQ systems and related components.

The legacy IROHA2 modules have been developed over the past decade to support a wide range of devices used in neutron instruments at MLF. By introducing EPICS alongside legacy components, the Device Manager can reuse EPICS device drivers while maintaining compatibility with existing implementations. This promotes the standardization of device control and enhances interoperability with other experimental systems.

By placing these components on top of the loosely coupled messaging backbone provided by Redis, the system

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD045

# INTEGRATION OF MICROCHANNEL PLATE DELAY LINE DETECTOR WITHIN THE KARABO CONTROL SYSTEM AT THE SXP INSTRUMENT OF THE EUROPEAN XFEL

I. Karpics[1], D. Doblas-Jiménez[1], A. Münnich[1], P. Grychtol[1], M. Heber[1,2], E. Tikhodeeva[1], V. Vardanyan[1], M. Turcato[1], S. Hauf[1], and M. Izquierdo[1]
[1]European XFEL, Schenefeld, Germany
[2]Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany

*Abstract*

The Soft X-ray Port (SXP) instrument at the European X-ray Free Electron Laser (EuXFEL) is designed to enable versatile, time-resolved X-ray Photoelectron Spectroscopy (tr-XPES) experiments on solids, surfaces and interfaces. At the core of the tr-XPES experimental station is a custom-built Time-of-Flight (ToF) momentum microscope photoelectron spectrometer, equipped with a commercial MicroChannel Plate Delay Line Detector (MCP-DLD). This contribution outlines the key steps and technical challenges involved in integrating the MCP-DLD into the EuXFEL's Karabo control system, highlighting both hardware and software considerations.

## INTRODUCTION

The SXP, located downstream of the third Self-Amplified Spontaneous Emission (SASE3) undulator at the European XFEL [1–4], is the seventh scientific instrument commissioned at the facility. It is designed to serve as a flexible FEL beamline where user-provided experimental setups can be integrated, thereby expanding the range of techniques available beyond those offered by the existing soft X-ray baseline instruments—Spectroscopy and Coherent Scattering (SCS) [5] and Small Quantum Systems (SQS) [6]. The first integrated user station is aimed at probing solid-state materials, surfaces and interfaces [4] leveraging the broad photon energy range of SASE3—from 0.25 keV to 3.5 keV—and enabling "complete" tr-XPES experiments. This includes simultaneous characterization of the electronic, chemical, and structural properties of materials with femtosecond time resolution.

Information about the material properties is obtained by accurately analyzing the photoelectrons emitted from a solid sample. To this end, we use a Momentum Microscope (MM) photoelectron spectrometer optimized for soft X-rays [7]. This system combines a cathode lens (an objective lens with a strong immersion field in front of the sample), a ToF tube, and a detector. The cathode lens is optimized for collecting electrons over a large solid angle, enabling imaging of a wide in-plane/parallel 2D momentum space (momentum mode) with high momentum (k) resolution. Additionally, the lens system can image the 2D real space in Photo-Emission Electron Microscopy (PEEM) mode. Energy discrimination is performed in the ToF tube, with the ultimate energy resolution determined by the drift tube length and the time resolution of the detector system.

The detector system must allow for the simultaneous measurement of energy and 2D momentum of photoelectrons (momentum mode), or energy and 2D real-space coordinates (PEEM mode) [8]. To reconstruct the full three-dimensional map of the distribution of photoelectrons, the detector must accurately discriminate the ToF of the photoelectrons and determine their point of impact with a spatial resolution of at least 70 µm. As illustrated in Fig. 1 [9], the MCP-DLD assembly consists of a silicon-based microchannel plate and two orthogonally oriented thin wires meandering across the detector area. Incoming photoelectrons are guided by electrostatic lenses and deflectors through the column of the MM into the MCP, where they are amplified, producing electron avalanches. These avalanches induce charge pulses on the delay lines—one corresponding to the x-axis and the other to the y-axis—which propagate to both ends of each line. By measuring the time difference between signals arriving at either end of a given delay line, the spatial coordinate of the electron impact is determined. This process, known as a "DLD event", yields the (x, y) position of the impact. Simultaneously, the mean arrival time of the hits provides the absolute time of detection. All timing and spatial data are processed by a Time-to-Digital Converter (TDC), which, together with the MCP-DLD and its associated readout electronics, constitutes the complete detection unit. This integrated setup enables synchronous data acquisition, streamlining operations by avoiding disparate software stacks and ensuring that the detector functions as a fully integrated component of the broader experimental infrastructure.

At the soft X-ray photoemission experimental station, a commercial MCP-DLD system provided by Surface Concept GmbH was deployed [10]. In addition to the hardware, the vendor supplies software for real-time visualization and data acquisition, storing the output in HDF5 format. However, to enable precise correlation of detector signals with other experimental subsystems on sub-millisecond timescales, it was necessary to integrate the detector directly into the Karabo framework developed and used at EuXFEL.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD046

# INTEGRATION OF A TAPE DRIVE SAMPLE DELIVERY SYSTEM AT THE EUROPEAN XFEL

A. García-Tabarés, E. Brambrink, G. Giovanetti,
C. Strohm, M. Tang, European XFEL, Schenefeld, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

High-throughput and reliable sample delivery plays a key role in high-energy laser experiments using low cross-section techniques such as X-ray Thomson scattering (XRTS), which require, due to the destructive laser-sample interaction, a large number of reproducible samples to ensure statistical accuracy. The tape drive—a conveyor-based system capable of continuously delivering samples at high speed—has been applied in structural biology and high-intensity laser experiments. The STFC-developed tape drive has been fully integrated into the facility's control infrastructure and is now used to perform highly accurate XRTS measurements. In this contribution, we describe the integration of the tape drive into the European XFEL control system, Karabo, enabling a level of sample delivery throughput that would not be achievable with alternative methods. We discuss how this integration improves efficiency by enabling precise synchronization and minimizing downtime.

## INTRODUCTION

High-energy laser experiments provide unique opportunities to probe matter under extreme conditions of temperature and pressure, enabling insights into states of matter relevant to planetary interiors, inertial confinement fusion, and astrophysical phenomena. Among the diagnostic techniques available, XRTS [1] has emerged as a particularly powerful method due to its sensitivity to electron density, temperature, and ionization state in dense plasmas. However, the scattering cross section in XRTS is inherently low, resulting in weak signals from individual shots [2]. To achieve statistically meaningful results, it is therefore necessary to accumulate data from a large number of repeated measurements. The High Energy Density (HED) instrument at the European XFEL provides unique opportunities to investigate matter under extreme conditions, offering a combination of high-power optical lasers and the high brightness, ultrafast, and high-repetition-rate X-ray pulses delivered by the XFEL [3].

Each shot, however, involves the interaction of a high-power laser and an intense X-ray probe with the target, which irreversibly damages the sample at the interaction point, precluding its reuse in subsequent measurements. As a result, high-throughput experiments require a system capable of delivering a fresh, reproducible sample for each shot, with high reliability and at a rate that matches the repetition frequency of the facility. This requirement is particularly acute at modern X-ray free-electron lasers (XFELs), such as the European XFEL, where the high repetition rate and demand for precise synchronization impose stringent constraints on sample delivery methods.



Figure 1: View of the tape drive rolls on position 0.

Conveyor-based tape drive systems, such as the one shown in Fig. 1, have been successfully employed in fields including structural biology and high-power laser experiments to meet similar challenges of high-speed, automated, and reproducible sample delivery. Building on this concept, the STFC-developed tape drive has been adapted for the European XFEL to meet the specific demands of high-accuracy XRTS measurements.

Here, we describe its integration into the XFEL control infrastructure, highlighting how precise synchronization with the facility's timing system enhances operational efficiency and minimizes experimental downtime. This implementation represents an important step toward fully exploiting the high-throughput capabilities of XFELs for studying matter under extreme conditions.

The Methodology section summarizes the technical implementation of the tape drive, while the Results section presents the results obtained in two experiments performed in 2025.

## METHODOLOGY

### *Tape Drive Sample Delivery System*

A conveyor-based tape drive system, developed by STFC, was employed to meet the requirement for delivering new and reproducible samples at the full repetition rate of the XFEL. The system consists of a motorized reel-to-reel mechanism advancing a continuous polymer tape at a controlled speed, positioning a fresh section of tape into the interaction region before each XFEL shot. The tape material and

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD047

# IMAGE LOGGING AND ANNOTATION IN KARABO FOR IMPROVED EXPERIMENTAL MONITORING

A. García-Tabarés Valdivieso*, R. Costa, D. Goeries, S. Hauf, I. Karpics, F. Sohn
European XFEL, Schenefeld, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

In the Karabo framework, efficient visualization and data logging are essential for monitoring and optimizing experiments. The Image Logger device captures and compresses images from cameras or other imaging sources, storing them for future reference. Image logging can be performed opportunistically—without a dedicated DAQ setup—by treating image data as slow data, meaning it is transmitted asynchronously through broker-based communication rather than via the real-time pipeline used for full-size image streams. This approach offers flexibility and adaptability to a variety of experimental conditions. Users can retrieve past images and view them alongside live ones within the same scene, enabling straightforward visual comparison and reducing reliance on manual logbook entries. Complementing this, the Image Annotation device tracks regions of interest (ROI) such as crosshairs and rectangles, preserving alignment and diagnostic information over time. Additionally, Grafana Panels have been configured to support historical inspection and troubleshooting, providing a structured way to review and interpret logged image data across different timeframes.

## INTRODUCTION

Effective image annotation is essential in experimental workflows, both to highlight and analyze regions of interest (ROIs) in imaging data and to keep track of specific features of interest over time. Various methods have been developed to support annotation, each offering different levels of complexity and capability. For image annotation, there were a few existing approaches, from very simple to more advanced:

- At the most basic level, users added marks directly to printed or displayed images. These marks were temporary and couldn't be saved or reviewed later.
- A more advanced, option allowed users to save markers within a project, keeping annotations linked to the image data for future use.
- An additional, yet invasive, approach embedded the annotations directly into the image files, permanently altering them in the process.

However, keeping track of annotations alone was sometimes not sufficient; furthermore, image logging relied on manually taking screenshots and attaching them to logbook entries, a method that lacked automation and limited data accessibility.

In some cases, users would also record annotations or note interesting features manually in the logbook to keep track of observations over time.

The initial image annotation workflow enabled users to define, record, and review regions of interest (ROIs) on experimental images using the Karabo GUI [1], which allowed marking areas with shapes and saving annotations alongside metadata such as labels, positions, and timestamps. Based on feedback from this early workflow phase [2], an integrated image logging capability was developed to preserve the full visual context by saving images together with their annotations in a structured way. Additional features, including an external trigger for automated overnight monitoring, were subsequently added to further enhance the system's functionality.

### *Slow and Fast Data at European XFEL*

At the European XFEL, the Karabo framework distinguishes clearly between fast and slow data: fast data for high-rate scientific measurements, and slow data for low-rate operational and diagnostic signals. Cameras are typically treated as fast data sources, producing continuous high-rate streams handled outside the standard slow-data infrastructure. In this work, however, compressed and selectively stored camera images are managed as slow data due to their reduced rate and format. This allows integration with monitoring and visualization tools like Grafana [3], which rely on time-series databases such as InfluxDB [4]. Storing images and metadata in the slow-data framework allows users to analyze visual data alongside operational parameters in real time using Grafana's querying and alerting tools. To support this capability, an Image Logger device has been developed which captures and stores full image frames in compressed, encoded form. The Image Logger supports both dedicated single-camera operation and multi-camera setups.

The Methodology section summarizes the technical implementation of these devices, while the Results section presents their operational performance from early 2025 through the beginning of the installation and maintenance period.

## METHODOLOGY

The approach and methodology for both annotating and logging images are outlined below.

### *Image Annotation*

As shown in Fig. 1, the Karabo GUI allows users to define regions of interest (ROIs) interactively by selecting from two shape tools: rectangle or cross. These can be drawn directly

* ana.garcia-tabares@xfel.eu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD048

# REAL-TIME EXPERIMENT STEERING AT EUROPEAN XFEL WITHIN KARABO

C. Carinan*, D. Hammer, M. B. Jakobsen, P. Schmidt, G. Flucke, D. Göries, I. Bermudez, T. Michelat, E. Sobolev, J. Wrigley, L. Le Guyader, J. Möller, A. Scherz, A. Madsen, S. Hauf, L. Gelisio
European XFEL GmbH, Schenefeld, Germany

*Abstract*

European XFEL is an X-ray free-electron laser user facility that generates high-throughput data streams. To ensure optimal steering of experiments, the real-time provision of compact metrics and live visuals for immediate feedback is critical. In this contribution, we describe how the facility's supervisory control and data acquisition system Karabo is exploited to this end. Two use cases, time-resolved X-ray absorption spectroscopy at the Spectroscopy and Coherent Scattering (SCS) instrument and X-ray photon correlation spectroscopy at the Materials Imaging and Dynamics (MID) instrument, exemplify typical workflows.

## INTRODUCTION

X-ray free-electron lasers (XFELs) produce ultrashort pulses of high-peak brilliance, which enable time-resolved measurements of samples of interest down to the atomic level. At European XFEL, X-ray imagers capture megapixel frames at megahertz rates, producing hundreds of gigabits of data streams which must be processed and analyzed in real-time—within seconds—to optimally control experiments, so as to maximize the scientific output of experiments performed at the facility. This contribution highlights how Karabo—the European XFEL's distributed supervisory control and data acquisition (SCADA) system—enables real-time monitoring and rapid decision-making by integrating data processing elements, and visualizing relevant analysis metrics as part of the environment used to control every aspect of ongoing experiments.

## EUROPEAN XFEL: OVERVIEW

European XFEL consists of a superconducting linear accelerator that can deliver electron energies up to approximately 17.5 GeV, followed by undulators which generate X-ray pulses spanning the soft and hard X-ray regimes [1, 2]. Three beamlines operate in parallel, each with multiple end-stations that support diverse experimental techniques, including imaging, scattering, and spectroscopy.

### *Pulses, trains, and data organization*

The facility operates in a burst mode that comprises up to 2700 X-ray pulses organized into a so-called train (Fig. 1). While the temporal spacing of trains is 100 ms (that is, their repetition rate is 10 Hz), the spacing of X-ray pulses is as short as 222 nanoseconds (4.5 MHz). Individual pulses are generally much shorter than 100 femtoseconds [1, 2].

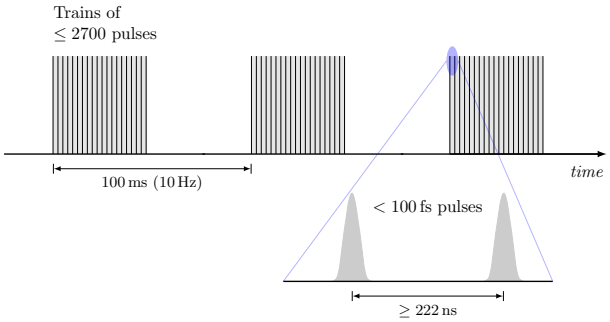



Figure 1: Time structure of European XFEL. A train contains up to 2700 X-ray pulses, generated at a maximum repetition rate of 4.5 MHz. The train repetition rate is 10 Hz.

### *Detectors for Burst-Mode Operation*

European XFEL mostly utilizes custom-made two-dimensional pixel detectors, including AGIPD [3] and DSSC [4], that record multiple frames per train, matching the facility's time structure. These devices come in a range of configurations, including variations in pixel size, sensor technology, readout concept, and dynamic-range strategy (e.g., adaptive gain or signal compression). These factors enable instruments to select the optimal approach for different experimental techniques, such as imaging, diffraction, or spectroscopy. Two-dimensional pixel detectors consist of different modules, each generating data streams of approximately 9 Gbit/s. When aggregated into a single component, often consisting of 16 modules, the resulting stream can reach about 140 Gbit/s. These data volumes require dedicated high-performance, low-latency data processing.

## KARABO: THE SCADA SYSTEM OF EUROPEAN XFEL

Karabo is a SCADA system developed at European XFEL [5, 6]. It provides a unified environment for hardware control, time-synchronous monitoring, and high-throughput data acquisition. The framework is device-oriented and distributed so that systems running across many hosts appear as a single control space. It differentiates between a low-rate control plane and a high-throughput data plane to ensure control responsiveness while large data flows move peer-to-peer.

### *Architecture and Communication*

Karabo consists of a distributed set of devices that encapsulate hardware (e.g., motors and detectors) as well as

* cammille.carinan@xfel.eu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD049

# THE STATUS OF HEPS BEAMLINE CONTROL SYSTEM*

G. Li†, AI.Y. Zhou, Y. Liu, Z.Y. Yue, CH.X. Yin, X.B. Deng, Q. Zhang, D.SH. Zhang, ZH.H. Gao, ZH. Zhao, N. Xie, G. Lei‡, Institute of High Energy Physics, CAS, Beijing, China

*Abstract*

HEPS (High Energy Photon Source) is the first fourth-generation high-energy (6 GeV) synchrotron radiation facility in China, which is mainly composed of accelerator, beamlines and end-stations. In the phase I, "14+1" beamlines and corresponding experimental stations have been constructed. The beamline control system is responsible for the high-precision automated control, status monitoring, and data acquisition of all optical components, diagnostic devices, and sample stages on the beamline, ensuring the delivery of stable, reliable, and high-performance beam to the end-stations. The EPICS-based control system has been designed and built, with commissioning now underway. Here, the status of beamline control system is presented.

## INTRODUCTION

The High Energy Photon Source (HEPS) [1] is the first high energy diffraction-limited storage ring (DLSR) light source in China (Fig. 1). HEPS will serve as a cutting-edge platform for multidisciplinary research in fields such as materials science, condensed matter physics, chemistry, life sciences, and engineering.

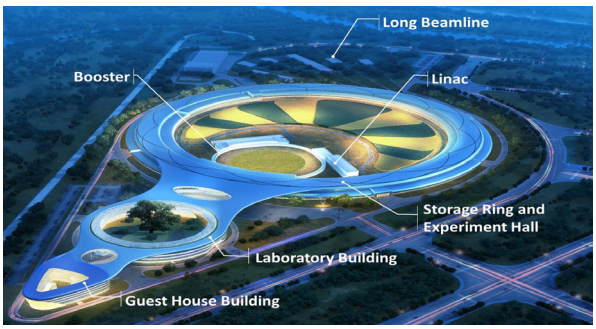



Figure 1: A bird's eye view of HEPS.

To satisfy some major demands of the nation and to further cutting-edge scientific development, the first phase will deliver 15 beamlines, including 14 for user experiments and one for testing purposes [2] (Fig. 2).

In order to simultaneously complete the construction of 15 beamlines control system on schedule, a set of standards is developed to design and build control systems for beamlines, minimizing heterogeneousness, to ensure that the control system has good stability, reliability, flexibility, availability, reliability etc.

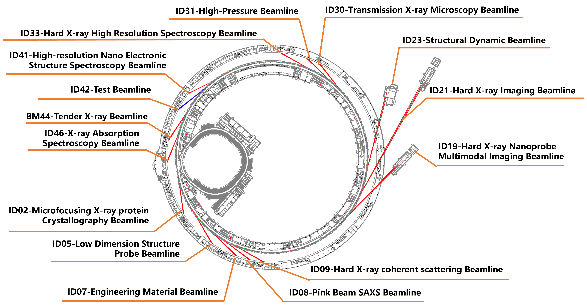



Figure 2: Beamlines layout in HEPS Phase I.

## BEAMLINE CONTROL SYSTEMS

The beamline control system is critical for ensuring stable, reliable, and precise beam delivery to experimental end-stations. It integrates heterogeneous hardware and software components into a unified architecture, based on the EPICS [3] framework (Fig. 3).

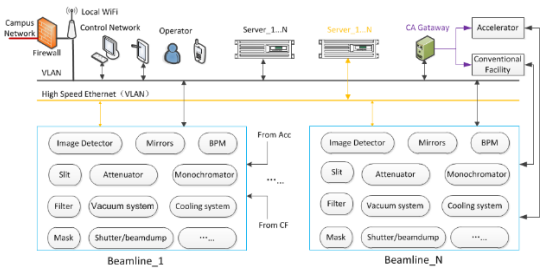



Figure 3: Overall structure of beamline control system.

The HEPS beamline control system adopts a typical three-tier architecture:

Operator Layer: Located in the experimental station or control room, it provides users with a graphical user interface (GUI).

Control Layer: Composed of high-performance servers and IOC (Input/Output Controller) servers, it forms the core of the EPICS architecture, running databases and control logic.

Device Layer: Includes all hardware devices and their controllers (motion controllers, PLCs, Intelligent I/O Module, digital multimeters, various point/line/area detector controllers, etc.)

This paper presents the current status of HEPS beamline control system.

### *Automated Deployment of the Software*

Each beamline has many front-end controllers that need to deploy operating system (OS) and EPICS applications, so it is necessary to design and build a software deployment system to complete the above tasks.

* Work supported by HEPS project
†li75gang@ihep.ac.cn
‡leige@ihep.ac.cn

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD053

# BLUESKY QUEUE SERVER FOR BEAMLINE CONTROL AT NSLS-II

D. Gavrilov, T. Caswell, A. Sligar, M. Rakitin
NSLS-II, Brookhaven National Lab, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Bluesky is a Python-based framework for experiment orchestration that is widely used at synchrotron facilities around the world. Queue Server (QS) is an essential component of Bluesky software stack that supports high-level functionality, such as control over the environment for executing Bluesky plans, enqueueing plans, executing and managing the plan queue, monitoring and controlling running plans, etc. The functionality is exposed via comprehensive set of APIs, which are designed to support wide range of workflows. QS is successfully used in applications, such as GUI-based and remote control, AI-driven and multimodal experiments.

## INTRODUCTION

Bluesky [1] is a Python-based open-source software developed at NSLS-II as part of the effort to implement standard approach to experiment control at the beamlines. Source code for all components of Bluesky software stack is available from Bluesky Collaboration GitHub organization [2].

The components of Bluesky stack are shown in Fig.1. The purpose of the Ophyd layer at bottom of the stack is to represent controlled devices as Ophyd objects, i.e. instances of Python classes with standard interface. The interface hides the underlying device implementation and allows standard access to the device. Ophyd devices may communicate with hardware directly or using higher level protocols, such as EPICS, REST API etc. The Bluesky layer supports execution of step-by-step logic of the experiments and communicates with the hardware using Ophyd objects. The experiment logic is expressed in the form of Bluesky Plans executed by Run Engine, which is part of Bluesky package.

Ophyd and Bluesky packages are Python libraries that provide components for developing Ophyd objects and experimental plans and Bluesky Run Engine for execution of the plans. In the traditional Bluesky workflow, plans were executed by manually running commands in IPython environment. While being the most flexible and convenient for development purposes, this workflow was found not intuitive by many beamline users. Several custom GUI applications were also developed to support workflows for selected beamlines. Those applications are running the control and GUI code in the same process and suffered from reliability issues.

Queue Server (QS) was recently added to the stack as part of transitioning to service-based architecture for experiment control. The purpose of QS is to implement server-client framework, in which execution of Bluesky plans on the server is fully controlled by client software using rich set of APIs. The client software may include GUI applications, user-developed scripts, autonomous agents, etc. Strong separation between the clients and the server performing low-level control operations increases reliability and robustness of the system. The architecture implements security features, including authentication and access control.

## QUEUE SERVER OVERVIEW

### *Components of Queue Server*

QS consists of the following components: Run Engine (RE) Manager [3], Bluesky QS API Python package [4], Bluesky HTTP Server [5]. Also, Bluesky-Widgets package [6] includes a number of PyQt widgets for controlling QS.

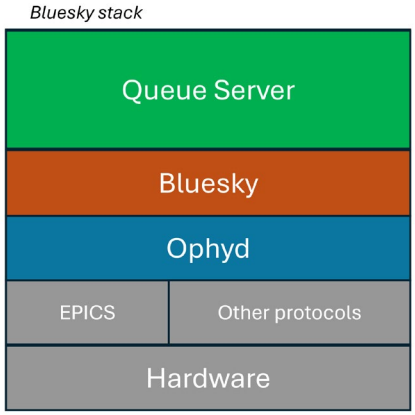



Figure 1: Components of the Bluesky Stack.

RE Manager is an execution engine for Bluesky plans. It is the core component of QS and can be started as a service or as a standalone application. It also supports complementary functions, including simple queueing of plans, execution of Python scripts and functions, etc. RE Manager exposes a rich set of APIs to client applications. Bluesky HTTP Server implements REST API for securely accessing RE Manager and exposing the QS on a wider network. The HTTP Server is easily extensible with custom routers to implement workflow-specific logic. Bluesky QS API is a Python package for communicating with QS using REST API via the HTTP server or directly with the RE manager. The package implements features for efficient communication with QS, such as local data caching, status monitoring, etc. Both thread-based and *async* versions of the API classes are available. Bluesky-Widgets package provides developers with general purpose PyQt widgets for controlling QS. The package also contains a generic *queue-monitor* application, which is useful for monitoring of QS in simple workflows.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD054

# INTEGRATION OF SAMPLE TEMPERATURE CONTROL AND X-RAY DATA ACQUISITION AT THE TPS 13A BIOLOGICAL SMALL- AND WIDE-ANGLE X-RAY SCATTERING BEAMLINE

C.-Y. Lin[1], Y.-Q. Yeh[1], K.-F. Liao[1], C.-C. Liang[1], Y.-H. Jiang[1], U–S. Jeng[1,2]
[1]Synchrotron Radiation Research Center, Hsinchu, Taiwan
[2]Department of Chemical Engineering, National Tsing Hua University, Hsinchu, Taiwan

*Abstract*

Integration of sample temperature control with X-ray scattering data acquisition is essential for the biological small- and wide-angle X-ray scattering (BioSWAXS) beamline 13A at the 3 GeV Taiwan Photon Source (TPS), National Synchrotron Radiation Research Center (NSRRC) [1–3]. Based on the Experimental Physics and Industrial Control System (EPICS) framework [4, 5], we implemented two independent temperature control systems—chiller and heater—into the X-ray data acquisition workflow through the Control System Studio (CSS) and the beamline data collection graphical user interface (DC-GUI). The chiller system employs a JULABO FCW-2500T unit for a temperature range of –20 °C to 80 °C, while the heater system (ambient to 300 °C) uses heating rods regulated by an SRS PTC10 controller, complemented by a gas-flow cooling module. Through the DC-GUI, the X-ray data acquisition system coordinates seamlessly with the sample environment control, including sample positioning, cooling, and both chiller- and heater-based GUIs. This integration enables programmable X-ray data acquisition under synchronized, programmable temperature control. A temperature-dependent SAXS measurement is demonstrated, showing that the stability and interoperability of the combined systems significantly enhance beamline performance and experimental reliability.

## INTRODUCTION

Small- and wide-angle X-ray scattering (SWAXS) is a versatile technique for probing structures of soft matter and biological macromolecules, including proteins, nucleic acids, and complex assemblies. Reliable temperature control is critical for studies involving temperature-sensitive biomolecules, phase transitions of polymer composites, or ligand interactions. The TPS 13A biological SWAXS beamline (BioSWAXS) is one of world-advanced SAXS-WAXS beamlines [1-3], powered by the high-brilliance 3 GeV Taiwan Photon Source (TPS) at the National Synchrotron Radiation Research Center (NSRRC). The beam-line operation system is based on the Experimental Physics and Industrial Control System (EPICS) [4, 5] framework, which enables seamless integration of diverse devices such as motors, vacuum systems, beamline optics, and sample environment controls, as detailed in the previous reports [1-3].

In this work, we report the integration of two sample temperature control systems to the X-ray data acquisition system of the beamline (the latter was reported previously [3]). The temperature control systems include a JULABO FCW-2500T chiller [6] and a heater system with two heating rods regulated by an SRS PTC10 controller [7], further complemented by a gas-flow cooling module. The temperature monitoring system employs on a Delta DTB4824 [8] temperature readout module (Table 1). Two independent graphic user interfaces are developed using the Control System Studio (CSS) for graphic operations of these two temperature control systems [9]. Furthermore, integration of the chiller and heater systems to the X-ray data acquisition systems are enabled through the coordination of the beamline's data collection graphical user interface (DC-GUI), the SPEC [10] system, the EPICS, and the Chiller and Heater GUIs, as detailed below. This integration provides programmable and synchronized temperature regulation and automated X-ray data collection, enhancing the reliability and efficiency of the beamline SWAXS data acquisition.

## SYSTEM ARCHITECTURE HARDWARE

Table 1: Devices for the Temperature Control Systems

| Device | Chiller (JULABO FCW-2500T) | Heater (SRS PTC10 with heating rods) | Temp. readout (Delta DTB4824) |
|---|---|---|---|
| **Function** | Coolant circulation cooling & heating | Precision heating control/ Gas-flow cooling | Temp. readout |
| **Range** | -20 –75°C | RT–300°C | 10–300°C, |
| **Practical Stability** | ±0.1 °C | ±0.5°C | ±0.5°C |
| **Sensor Type** | K-type TC | K-type TC | K-type TC |
| **Interface** | RS-232> RJ-45 | RJ-45 | RS-485> RJ-45 |
| **Moxa** | NPort 5150A | N/A | NPort 5150A |

The temperature control system of the TPS 13A BioSWAXS beamline is designed for dual-range operation, integrating a chiller of programmable temperature control in a temperature range of –20°C (coolant-dependent) to 80°C and a heater system (ambient to 300°C) using heating rods regulated by an SRS PTC10 controller. These two systems are integrated under the Experimental Physics and Industrial Control System (EPICS) frame-work through dedicated IOCs, enabling synchronized and programmable operation with the beamline X-ray data acquisition system through data collection graphical user interface (DC-GUI). All temperature sensors used are K-type thermocouples

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD055

# EFFICIENT MULTI-REGION FLY-SCAN ARCHITECTURE USING PLC INTEGRATION FOR SYNCHROTRON XAS

C.-C. Liang†,[1], Y.-J. Lu[1], J.-L. Tu[1], M.-H. Lee[1], Y.-C. Huang[2],
C. Chan[1], H.-W. Chen[1], K.-L. Yu[1], P.-Y. Peng[1], C.-L. Chen[1]
[1]National Synchrotron Radiation Research Center (NSRRC), Hsinchu, Taiwan
[2]Curtin Centre for Advanced Energy Materials and Technologies,
Curtin University, Perth, Western Australia

*Abstract*

A high-speed fly-scan system has been developed and implemented at the Taiwan Photon Source (TPS) 32A Tender X-ray Absorption Spectroscopy (TXAS) beamline to enhance data acquisition efficiency for X-ray Absorption Spectroscopy (XAS). Unlike conventional fly-scan systems that rely on Field Programmable Gate Array (FPGA) or CompactRIO (cRIO) platforms, this system adopts a programmable logic controller (PLC)-based architecture, offering high modularity, low development complexity, and seamless integration with existing beamline infrastructure. The system enables continuous scanning of the double-crystal monochromator (DCM) while simultaneously acquiring synchronized data from multiple detectors. Analog-to-digital converters (ADCs) digitize current signals for transmission and total electron yield (TEY) modes, while high-speed digital outputs from the PLC provide external triggers for silicon drift detectors (SDDs) used in partial fluorescence yield (PFY) measurements. This approach simplifies the coordination of asynchronous devices and eliminates the need for complex timing hardware. To further optimize performance, a multi-region scan strategy was implemented, allowing dynamic adjustment of scan speed across different energy segments. This significantly reduces total scan time by up to a factor of six compared to traditional step-scan methods, while maintaining comparable spectral quality. The PLC-based fly-scan system also supports real-time data logging, automated sample switching, and energy-axis registration, making it a robust and scalable solution for time-resolved and in-situ XAS experiments.

## INTRODUCTION

The TPS 32A is a new beamline recently established at the TPS of the National Synchrotron Radiation Research Center (NSRRC) [1]. Utilizing a bending magnet source, it covers a photon energy range of 1.7-11 keV and delivers a highly purified X-ray beam with a small focal spot and minimal higher-order harmonic contamination. Fig. 1 illustrates the beamline schematic.

A significant challenge for this beamline is its inability to employ the common quick-scanning monochromators used elsewhere for rapid data acquisition [2, 3], due to its specific energy range. Consequently, reliance on the traditional step-scan method for collecting an Extended X-ray Absorption Fine Structure (EXAFS) spectrum is highly inefficient, with each scan requiring 25 to 50 minutes to complete. This limitation severely restricts experimental throughput and the optimal use of the synchrotron source.

To overcome this bottleneck, we have developed and commissioned a new "fly-scan" data acquisition module for the TPS 32A beamline. This system successfully reduces the time for a full EXAFS scan to under one minute while maintaining exceptionally high data quality. The module supports multiple detection modes, including transmission, TEY, and PFY. Its development hinged on the precise synchronization of the monochromator's Bragg angle ($\theta_B$) movement with various detectors, while also overcoming challenges related to electronic noise and signal triggering. The resulting fly-scan system now provides data quality comparable to the conventional step-scan method, thereby significantly enhancing the research capabilities and efficiency of the beamline.

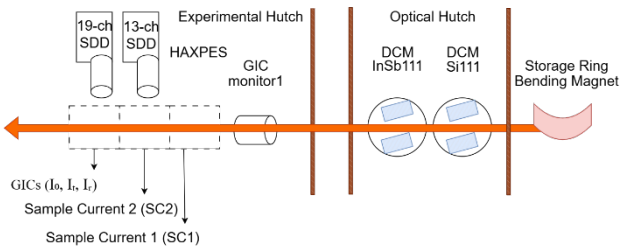



Figure 1: The schematic of the TPS 32A beamline.

During the initial phase of beamline commissioning, a conventional step-scan system was systematically established, along with several optimization tools, to verify the beamline design and support alignment and calibration procedures [4]. In traditional transmission XAS, the incident photon energy is varied by adjusting the pitch of the crystals ($\theta_B$) in the DCM (Kohzu TWC-2A/B). At each energy point, the system pauses to record the current signals from gridded ionization chambers (GICs), which correspond to $I_0$, $I_t$, and $I_r$. The scan then proceeds to the next energy position, repeating the measurement cycle. After numerous steps, a complete absorption spectrum is obtained. For PFY measurements, the system must dwell at each designated energy point to allow the SDD to integrate the signal over a defined exposure time, typically ranging from 0.4 to 4 seconds. Only after multiple exposures can a full fluorescence spectrum be acquired. The current signals from the GICs are amplified using NF-CA5351 amplifiers, while the sample current is measured using a Stanford Re-

†changeliang@nsrrc.org.tw

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD059

# GAS SUPPLY CONTROL SYSTEM FOR STRAW TUBE DETECTORS IN THE PANDA EXPERIMENT

T. Madej†,[1], J. Smyrski[2], L. Zytniak[1]
[1]S2Innovation Sp. z o.o., Krakow, Poland
[2]Jagiellonian University, Krakow, Poland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

In the PANDA experiment designed for studies with antiproton beams at the FAIR (Facility for Antiproton and Ion Research) accelerator complex, gaseous detectors of straw tube type are used for precise reconstruction of charged particle tracks. Developed multi-channel gas mixture supply system for the PANDA straw detectors meets high requirements regarding, among others, precision of mixing the component gases, stabilization of gas pressure in the detectors and protection of the detectors in the event of a system failure.

The hardware architecture of the gas system integrates gas supply components, atmospheric condition sensors, and automation based on Siemens S7-1200 PLC controllers to maintain optimal working conditions. The software structure, based on EPICS, enables modular control, real-time monitoring, and efficient data collection. The user interface, developed using Phoebus, provides an intuitive graphical environment for system operation and diagnostics.

This work presents the design, implementation, and functionalities of the gas supply control system, highlighting its role in ensuring precise and reliable operation of the PANDA straw detectors. Potential improvements and future extensions are discussed, focusing on optimizing system performance for upcoming experimental phases.

## INTRODUCTION

The PANDA [1] experiment (antiProton ANnihilation at DArmstadt) is one of the projects at the FAIR accelerator complex in Darmstadt, Germany. By exploiting high-quality antiproton beams with momenta between 1.5 and 15 GeV/c, PANDA will enable precision studies of the strong interaction. Its broad physics program includes the spectroscopy of charmonium states, the search for exotic hadrons such as glueballs and hybrids, investigations of hadrons embedded in nuclear matter, and detailed studies of hadron structure. The detector system is divided into two major parts: the Target Spectrometer, which surrounds the interaction region, and the Forward Spectrometer, optimized for forward-going particles.

Within this complex setup, charged-particle tracking plays a central role, and PANDA employs straw tube trackers as its main tracking devices in both the central and forward regions [2]. These detectors are based on thin Mylar tubes, operated at overpressure, which makes them self-supporting and eliminates the need for heavy mechanical frames. Their low material budget is particularly important to minimize multiple scattering and thereby preserve tracking precision. In addition to reconstructing trajectories with high spatial resolution, the straw detectors also provide particle identification at momenta ($p < 1$ GeV/c). The combination of accuracy, low mass, and dual functionality explains the choice of straw tubes as the baseline technology for PANDA.

The reliable operation of these detectors critically depends on a dedicated gas system that supplies a well-defined mixture of argon, carbon dioxide, and oxygen [3]. This system must ensure the stability of gas composition, maintain constant pressure to support the mechanical integrity of the straws, and include safety mechanisms to prevent damage in case of failure. High precision, stability, and robustness are therefore essential requirements. Meeting these criteria guarantees that the straw tube detectors will achieve their expected performance and contribute to the overall success of the PANDA physics program.

## HARDWARE ARCHITECTURE OF THE CONTROL SYSTEM

The hardware architecture consists of three subsystems: gas mixing, gas distribution, and environmental monitoring. Together they ensure accurate gas preparation, safe delivery to the detectors, and continuous supervision of ambient conditions.

### *Gas Mixing Subsystem*

The gas mixing subsystem (see Fig 1.) is the first stage of the control system and is responsible for preparing a defined gas mixture delivered to the gas distribution system. It consists of three independent supply lines dedicated to argon (Ar), oxygen ($O_2$), and carbon dioxide ($CO_2$).

At the entrance of each line, a high-pressure gas cylinder (left side of Fig. 1) serves as the primary source of the component gases, which is controlled with manual valve, providing operator-level control for enabling or isolating the gas flow during start-up, shutdown, or emergency procedures.

After the valve, the gas passes through a pressure gauge (PG), which provides continuous feedback on the actual supply conditions. This ensures that abnormal pressure levels can be identified early, and corrective measures taken.

The next stage in each line is a filter unit, which removes potential contaminants.

Each supply line is then equipped with a mass flow controller (FC1 – FC3). In the current setup, Bronkhorst mass flow controllers are used, connected directly to the main control computer via an RS232 interface. These controllers allow for accurate regulation of the individual gas flows.

† tomasz.madej@s2innovation.com

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD061

# INTERFACING CUSTOM (EPICS) SOFTWARE WITH IN-HOUSE DEVELOPED FRAMEGRABBERS & COMMERCIAL CAMERAS*

M. Skoufis†, C. Curtis, B. Hill, R. Herbst, L. Ruckman, P. C. O'Grady, J. Thayer, M. Weaver
SLAC National Accelerator Laboratory, Menlo Park, CA, United States

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Integrating machine-vision system cameras in an experimental setup can be a tedious and perilous process. First, there is the challenge of adapting a commercial solution to a slightly different environment or mission than the one for which it was originally designed. Second, the design may involve both commercial and internally built hardware and software, with all pieces required to be compatible and interoperable. Furthermore, since separate engineering teams usually develop and deploy the different layers of the final architecture, this could easily lead to miscommunication in defining the various interfaces, configuration and even operational settings and parameters. As a result, failures may occur on various layers in the resulting architecture. At the SLAC National Accelerator Laboratory, such cameras are used, among other things, to monitor the profile of the Linac Coherent Light Source (LCLS) X-ray laser beam, but also for modeling purposes. A custom data-acquisition solution that includes an in-house designed framegrabber board has been developed to interface with those cameras. The primary motivation of a custom solution is the integration of the timing system into the camera readout. In essence, the LCLS timing system is used by firmware to trigger the camera, but also to construct events consisting of camera image data, along with the associated timing information. In this paper, we present a summary of this work from a primarily software perspective.

## MOTIVATION

The main motivation of this work comes from the pressing need to deploy a machine-vision solution tailored to SLAC's needs. In particular, part of such a solution is required to encompass the internal Linac Coherent Light Source (LCLS) timing and ensure that the obtained image data gets timestamped. This will allow us to support ongoing efforts on beam diagnostics, beam modeling, beam optimization, machine-learning algorithms for various purposes, data visualization and physics-based analysis.

Also, having a common solution that can be applied to a broad category of cameras, ensures the ability to select from various commercial cameras those that comply with the mission requirements the best. Moreover, being able to add functionality to a machine-vision system whenever a need arises – and with a short turnaround time – is extremely valuable. Otherwise, the overall mission capabilities would most likely be limited by what commercial products have to offer at the time.

## TIMING OVERVIEW

The LCLS timing systems aim to provide the means for the synchronization of critical beam events, but also for the triggering of devices for data acquisition purposes. The two main LCLS experiments, LCLS-I and LCLS-II, use different timing systems, as the beam rate requirements are significantly different between these two [1–4].

In particular, the original LCLS-I timing system supports rates up to 120 Hz, while the LCLS-II timing system substantially raises that to approximately 1 MHz. Also, the former relies on timing events for the propagation of relevant information to the various devices, whereas the latter uses timing patterns for the same purpose [5, 6].

In terms of the targeted capture rates for beam monitoring, the limitations are primarily due to the maximum attainable frame rates supported by each camera and its interface for the desired pixel resolution and frame size. Secondary latencies can also be introduced elsewhere, depending on the communication links used and the processing time complexity in firmware and software.

## FIRMWARE OVERVIEW

Although the developed firmware is not the primary focus of the present manuscript, it is useful to present a brief summary of its overall capabilities.

The designed and developed PCIe board and associated firmware form a generic framegrabber solution used to interface with various commercial cameras. Besides the in-house designed board, commercial boards have also been used successfully for the same purpose.

The only requirement for the commercial cameras is that they are Camera Link compatible cameras, meaning that they utilize the Camera Link protocol to establish external communication [7]. Also, it is anticipated that all of these cameras will provide a common set of capabilities, independent of manufacturer.

The firmware incorporates LCLS-I or LCLS-II timing in the camera readout, and all acquired frames are timestamped with this internal timing. Following acquisition, the obtained data can be processed and used in meaningful ways, such as for standard beam modeling & simulation, in machine-learning algorithms to train prediction models, etc.

As part of the deployed solution, a conversion box is incorporated between the framegrabber PCIe board and the camera in order to convert the Camera Link protocol to the proprietary Pretty Good Protocol (PGP) communication

* Use of the Linac Coherent Light Source (LCLS), SLAC National Accelerator Laboratory, is supported by the US DOE, Office of Science, Office of Basic Energy Sciences (contract No. DE-AC02-76SF00515).

† skoufis@slac.stanford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD065

# UPGRADES TO THE FACET-II DATA ACQUISITION SYSTEM*

S. Perez†, S. Gessner[1], I. Rajkovic, SLAC National Accelerator Laboratory, Menlo Park, CA, USA
[1]also at Stanford University, Stanford, CA, USA

*Abstract*

The Data Acquisition System (DAQ) at FACET-II collects and saves synchronized, time-stamped data from various diagnostics on the linac and experimental area, including digital cameras and devices on the EPICS control system. During data acquisition, the camera input-output controllers (IOC) save image data to network-attached storage (NAS). While large files are able to save to the NAS efficiently, writing many small image files to the NAS is 10x slower. Additionally, the DAQ cannot acquire data from several important devices, such as BPMs, klystrons, and magnets, because they are only accessible via the legacy control system, the SLC Control Program (SCP). To address these issues, the FACET-II DAQ has been upgraded to save data faster by packaging multiple small image files as larger HDF5 files. The improved DAQ is also able to acquire BPM data from SCP by parallelizing the data acquisition process from both control systems, then comparing timestamps across cameras, EPICS devices, and SCP BPMs to ensure data is fully synchronized.

## INTRODUCTION

At FACET-II [1], a 1-kilometer-long test beam facility at SLAC, digital cameras are the most important diagnostics. During experiments, users request time-synchronized images from many cameras at beam rate (10 Hz). FACET-II uses EPICS [2] to control and retrieve data from cameras, and communication in EPICS is managed by Channel Access (CA). However, data rates are limited by the network. The Data Acquisition System (DAQ) [3] was developed to solve the problem of acquiring and saving time-synchronized data from multiple sources, without overloading the network and disrupting accelerator applications. It does this by coordinating camera input-output controllers (IOCs) to save camera data to a network-attached storage (NAS) drive. The DAQ also saves scalar data from various devices along the linac, such as beam position monitors (BPMs) that update at beam rate. These devices are included in the beam-synchronous acquisition (BSA) [4] infrastructure at SLAC. The DAQ finds matches between timestamps from the different sources in order to return time-synchronized data.

Some desired improvements to the original DAQ included increasing the speed and efficiency of the software, as well as adding additional functionality to support data acquisition from devices on the legacy control system. The DAQ has been upgraded to address these issues.

* Work supported by the U.S. Department of Energy under DOE Contract No. DE-AC02-76SF00515.

† shperez@slac.stanford.edu

## OVERVIEW OF DAQ FUNCTIONALITY

The DAQ runs on a dedicated server, called facet-srv20. The control server and the NAS are on the same network switch as the camera servers. The DAQ has a graphical user interface (GUI) component, shown in Fig. 1, where the user selects desired parameters. For example, the user specifies the number of shots (data points) per step and cameras to be included in the data acquisition. The user does not have to select any scalars, but at least one camera must be selected. Once the setup is done, the user initializes the DAQ program by pressing the "Run" button. First, the DAQ completes basic checks, such as ensuring that the DAQ is not already running. Then, the DAQ creates a data path on the NAS drive and sets up the Event class that manages BSA data acquisition. Next, the DAQ starts acquiring images from the requested cameras, followed by scalar BSA acquisition. During the acquisition, the DAQ can also collect data from non-BSA variables at a rate close to the beam rate. The DAQ repeats this process for each step. While the DAQ is running, a log updates with timestamped messages.

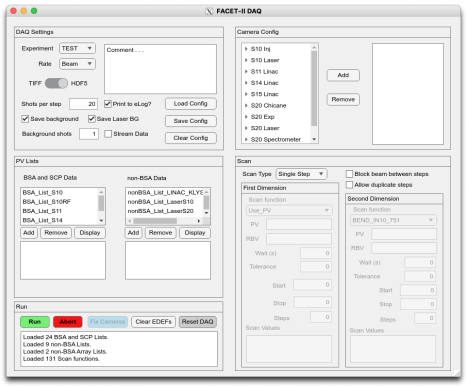



Figure 1: Graphical User Interface for the FACET-II DAQ.

## EFFICIENCY IMPROVEMENTS

The principal motivation for upgrading the DAQ was to increase the efficiency of the software. In particular, there was a delay in the DAQ when writing camera image files to the NAS during data acquisition. Originally, each image was saved as a TIFF file directly to the camera disk, and files were subsequently written to the NAS one by one. We found that this step took longer when there were many small files than for one singular, larger file.

For example, in one particular test, we found that saving 450 images (2.4 MB each) took longer to write to the NAS than one 1 GB file. The effective speed was about 30 MB/s when writing many small files, compared to 300 MB/s

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD066

# AUTOMATING X-RAY BEAM ALIGNMENT PROCESSES WITH THE SPLIT AND DELAY GRAPHICAL USER INTERFACE AT THE LINAC COHERENT LIGHT SOURCE*

C. Gee†, C. Melendrez, S. Rafiqui, S. Song
SLAC National Accelerator Laboratory, Menlo Park, CA, United States

*Abstract*

At SLAC National Accelerator Laboratory's Linac Coherent Light Source (LCLS), a series of optics and diodes in the X-ray Correlation Spectroscopy (XCS) Split and Delay Chamber divide the beam into two equal intensity pulses. One pulse is intentionally delayed, facilitating X-ray Photon Correlation Spectroscopy techniques that operate at nanosecond time scales for experiments in biology, chemistry, and materials science. However, achieving precise beam alignment with this setup poses significant challenges because meticulous adjustments to each of the beam line optics are required in succession each time the Split and Delay chamber is used, and this process is very time and effort intensive.

To address these challenges, the implementation of the Split and Delay graphical user interface (GUI) streamlines x-ray beam alignment by enabling remote control of the optical components and displaying user-friendly, live-time monitoring of the beam diagnostics. Work is ongoing to allow automatic optimization of beam alignment by stepping the motorized stages holding optical components through various positions, measuring the beam intensity, and moving optics on motorized stages to the preferred positions. This advancement will further streamline alignment of the beam through the Split and Delay chamber and thus increase the available time for data collection at XCS.

## XCS SPLIT AND DELAY CHAMBER

The XCS Split and Delay Chamber is installed along the LCLS [1] hard x-ray beamline in the x-ray transport tunnel (XRT) alcove, which is upstream of the XCS hutch. It is an enclosed chamber containing Split and Delay optics as shown in Fig. 1 and it operates under helium gas.

The crystal optics are used to split the x-ray beam into a pair of x-ray pulses that can be continuously adjusted to have a time separation as short as femtoseconds. The crystal beam splitter splits the x-ray beam into two paths shown in Fig. 2. The green line in Fig. 2 shows the Delay Diode (DD) path, which is adjusted by changing the relative positions of the two stages. The red line in Fig. 2 shows the fixed Channel Cut (CC) beam path. The beams are then recombined with a crystal beam combiner at the sample interaction point to achieve two closely-timed coherent scattering patterns that can be measured by a pixel detector in the XCS hutch.

The Split and Delay system is composed of 24 motors with crystals mounted on them. Six of the motors used frequently for beam alignment are included in the Split and Delay GUI (t1.x, t2.x, t3.x, t4.x, t1.th2, and t4.th2), as well as eight diagnostic diodes (d08-d15) that are between the crystals as shown in Fig. 3.

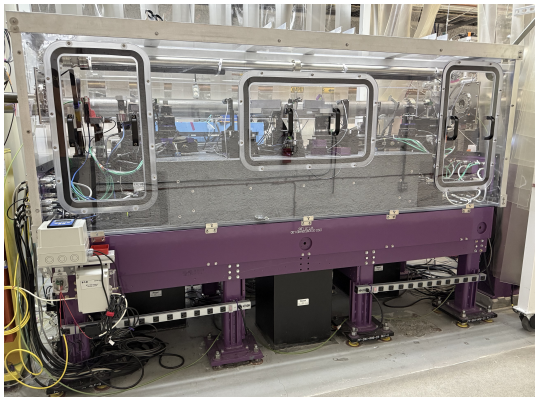


Figure 1: The XCS Split and Delay Chamber on the LCLS hard x-ray beamline.

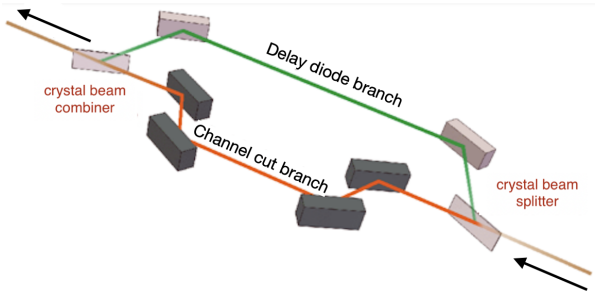



Figure 2: Diagram of the crystal beam splitter separating the channel cut (red) and delay diode (green) paths and then recombining them with the crystal beam combiner through the Split and Delay chamber. The beam travels from right to left in this diagram [2].

## MONITORING

Monitoring the eight diagnostic diodes in the Split and Delay chamber is essential for efficient beam alignment through it. Diodes 8, 9, 11, 13, and 15 are used for x-ray beam intensity diagnostics, while diodes 10, 12, and 14 are used for beam intensity and profile diagnostics. The monitoring screens in the Split and Delay GUI make it easy to view the relevant graphs of diode intensity over time and diode intensity ratios in a single diagnostic monitoring window. Examples of some plots in the monitoring screen are shown in Fig. 4. Each diode measurement is time stamped and

* This work was supported by the U.S. Department of Energy, Office of Science, under Contract No. DE-AC02-76SF00515
† cagee@slac.stanford.edu

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD067

# EPICS IOC DEVELOPMENT FOR HIGH-SPEED AND HIGH-RESOLUTION IMAGING WITH THE AXIS CAMERA

K. Kim†, B. Hill, E. Williams, K. Kunnus, K. Kim, P. C. O'Grady
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

At SLAC National Accelerator Laboratory, we recently integrated a high-performance AXIS-SXRF-60 camera into the EPICS control system for advanced imaging applications. The AXIS camera features a 6144 × 6144 pixel sensor with 16-bit depth per pixel, generating approximately 75 MB per image. With a maximum frame rate of 26 frames per second, the system must handle a sustained data throughput of nearly 1.9 GB/s. This paper presents the development and deployment of EPICS IOC designed to interface with the AXIS camera, focusing on software optimization strategies to manage the substantial data volume efficiently. Leveraging the areaDetector framework, the ADGenICam and ADEuresys modules, we expedited the integration process while ensuring compatibility with the Euresys frame grabber hardware. Key topics include multithreaded data pipelines, network optimization, and techniques to maintain stable, real-time image acquisition. The work demonstrates how high-speed imaging systems can be effectively supported within an EPICS environment, providing a template for similar large-scale instrumentation efforts.

## INTRODUCTION

High-performance cameras are increasingly critical in modern accelerator facilities for both machine protection and scientific productivity. High-brightness, short-pulse beams demand single-shot, low-latency diagnostics to characterize beam size, shape, position, and pointing on every pulse. These measurements underpin fast feedback loops for orbit stabilization, emittance tuning, and jitter reduction, enabling reliable delivery of tightly specified beam parameters to users. In addition, advanced experiments—from soft X-ray imaging and spectroscopy to pump–probe studies—benefit from large field-of-view, high-dynamic-range sensors capable of capturing weak features alongside bright structures in a single exposure. The combination of high resolution, high frame rate, and low read noise offered by modern sCMOS/CMOS sensors makes them well-suited to these tasks [1, 2].

Across the community, facilities have adopted high-speed area cameras for beam profile monitors, optical transition radiation screens, scintillator-based diagnostics, and X-ray imaging. EPICS areaDetector is widely used to integrate such devices with control system infrastructure, data acquisition, and data reduction pipelines [3-6]. CoaXPress frame grabbers, GenICam-compliant camera interfaces, and robust EPICS plugins have enabled reliable deployments for high-data-rate applications at multiple labs [3, 6-9].

At SLAC, the AXIS-SXRF-60 camera was chosen for this project, which provides the largest sensor currently available at our lab, a 36 Mpix, 6144 × 6144 sCMOS sensor with 16-bit depth [10, 11]. Its combination of very large field-of-view, soft X-ray optimization, and CoaXPress connectivity addresses our requirements for (1) full-aperture imaging in a vacuum environment, (2) high sustained bandwidth for continuous acquisition, and (3) compatibility with existing SLAC timing, controls, and data systems [10-12]. The camera's integration with a Euresys CoaXPress frame grabber and its GenICam compliance made it a strong fit for rapid integration with EPICS via areaDetector, ADGenICam, and ADEuresys [3, 7, 8].

## SYSTEM ARCHITECTURE

### *Hardware*

Figure 1 shows the wiring diagram of the complete system (modified from [10]). The main elements are:

- A: Duplex optical cable from the camera head
- B1: Optical converter
- B2: Four coaxial cables (CoaXPress links)
- B3: Euresys CoaXPress frame grabber (Coaxlink Quad CXP-12) installed in the host PC
- C: Host PC, which runs EPICS IOC and areaDetector on RHEL7
- D: DB15 to BNC cable to receive trigger signals
- E: 12 VDC universal power supply
- F: Water cooler for camera thermal management
- G: F-mount adaptor for test and development

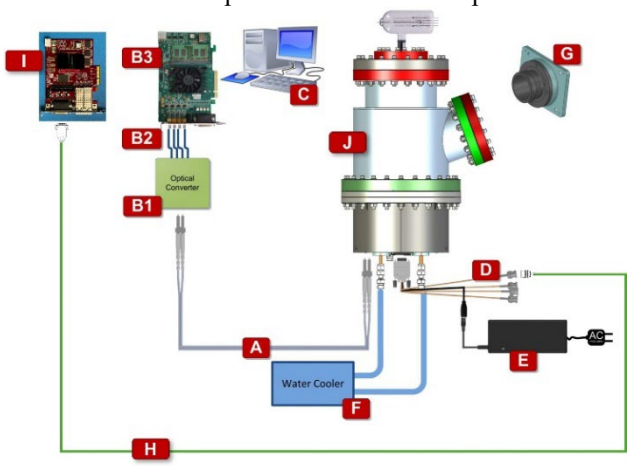



Figure 1: Wiring diagram of the system.

† ktkim@slac.staford.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD072

# INTEGRATED DISPLAY MANAGEMENT FOR CONTROL SCREENS AT APS ACCELERATOR COMPLEX*

L. Xiao†, A. Johnson, N. Arnold, E. Chandler, D. Jarosz, G. Shen
Advanced Photon Source, Argonne National Laboratory, Lemont, IL, United States

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Advanced Photon Source (APS) accelerator employs both MEDM and CSS Phoebus as display managers to support legacy control screens and the new displays introduced during the APS Upgrade project. This dual-system setup presents challenges in the organization, deployment, and maintenance of accelerator control displays. To address these issues, a unified, centralized framework has been developed by the APS accelerator for managing the control screens, regardless of format, purpose, or update frequency. This paper describes the strategies used by APS accelerator controls to organize, deploy, and maintain control displays. It also presents the tools and practices used to manage display managers across the APS accelerator control system.

## MOTIVATION

The APS Upgrade Project has introduced new control system displays to support modernized infrastructure and enhanced functionality, creating a more complex environment that includes both legacy and modern technologies. MEDM [1], the legacy EPICS display manager, remains in use for its simplicity and reliability, while CSS Phoebus [2,3], a modern Java-based display manager, provides advanced features and integration with tools like ChannelFinder and Archiver Appliance.

Managing these two display systems poses challenges in organization, deployment, and long-term maintenance. To address this, the APS Accelerator Controls group has implemented a unified strategy for managing legacy and modern displays. This framework ensures seamless interoperability, robust version control, efficient development and deployment, and streamlined operator access, while maintaining operational continuity across the APS accelerator complex.

## CURRENT DISPLAYS

APS employs a diverse set of display applications to monitor, control, and diagnose its accelerator subsystems. These include legacy MEDM (ADL) displays, high-performance EDM (EDL) panels, and modern BOB files implemented in CSS Phoebus, as summarized in Table 1. The following sections describe the primary display formats, their version control and maintenance practices, and their usage within the APS accelerator infrastructure.

Table 1: Applications and Associated Display Formats

| Application | Display Format |
|---|---|
| MEDM (Motif Editor and Display Manager) | ADL (.adl) |
| EDM (Extensible Display Manager) | EDL (.edl) |
| CS-Studio Phoebus | BOB (.bob) |

### *Legacy Displays*

The APS maintains a substantial collection of legacy MEDM displays, which primarily support the injector and other critical systems. These display files are located in directories defined by the `EPICS_DISPLAY_PATH` environment variable—a colon-delimited list that MEDM uses to resolve relative paths [4]. To ensure long-term stability and preserve institutional knowledge, these displays remain under CVS version control for consistent display management [5].

A dedicated set of EDM (Extensible Display Manager) [6] displays is utilized for high-performance applications. These EDL files are provided by Dimtel and supplied with the devices, serving as operator panels for precise control and monitoring of the iGp12 bunch-by-bunch feedback processors in the lepton storage rings. All EDL displays have been converted to BOB format to streamline management and integrate with modern tools.

### *New Displays*

As part of the APS-U project, the APS accelerator has adopted CSS Phoebus for modern control system displays. It was introduced to support devices that provide only BOB files (e.g., Libera Brilliance+ and Libera Spark) and systems requiring advanced features, such as image plotting for cameras and DAQ [7].

New ADL and BOB files are developed for each IOC top application and distributed alongside the corresponding IOC software. These display files are stored in the same Git repository as the IOC support code, ensuring that any changes are automatically included in the relevant software release.

In addition to IOC-specific displays, new ADL and BOB screens are also developed for general or user-centric purposes. These may support a particular subsystem, such as Timing or Libera LLRF, but are not tied to a specific IOC. All modern display files are maintained in GitLab, enabling collaborative development and robust change tracking [5].

* Work supported by U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, under Contract No. DE-AC02-06CH11357.
† xiaol@anl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD074

# ELETTRA COMMANDER: A FLEXIBLE WEB INTERFACE ON LARGE TOUCHSCREENS

L. Zambon*, G. Gaio, Elettra Sincrotrone Trieste, Trieste, Italy

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Elettra Commander is an innovative web interface designed to centralize particle accelerator operations. It integrates web-based and native user interfaces into a unified platform, simplifying control tasks. A key feature is the dynamic integration of machine synoptics, which display accurately scaled curves and data with rapidly updating graphs and machine control applications, enabling real-time interaction. Optimized for large touch-screens exceeding 50 inches, Elettra Commander offers intuitive navigation and responsive performance. This presentation will explore its architecture, highlight key features, and outline future developments, demonstrating its potential to advance accelerator control with a flexible, centralized, and effective solution.

## INTRODUCTION

Elettra Commander draws inspiration from the console of a vessel commander, serving as a proof-of-concept development aimed at demonstrating an innovative approach to organizing the graphical user interfaces (GUIs) of a control system. The foundation of Elettra Commander is a full-screen browser page, which can also operate in kiosk mode (Fig. 1).

This application effectively displays and summarizes all essential parameters while providing the capability to launch any GUI, whether web-based or native. By integrating the primary synoptic of a control system with an application launcher, Elettra Commander merges critical functionalities into a single, cohesive application.

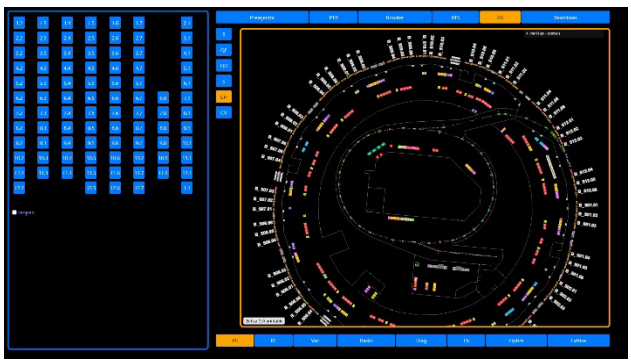


Figure 1: Elettra commander.

## SYNOPTIC

Elettra Commander utilizes a large touch screen divided into two or three thematic areas. By default, the main area displays the primary synoptic, but it can seamlessly switch to a series of specialized synoptics as needed. Each button in the first level modifies the content of either the main or secondary area, enabling control over specific geographical regions of the accelerator and/or logical functions.

The secondary area is dedicated to launching specialized applications, such as atomic panels for device servers. Each section, as well as the entire page, is designed to be responsive, functioning effectively on smaller screens and tablets, although Elettra Commander is primarily optimized for large touch screens. The implementation of tables is achieved through flexbox [1], ensuring a flexible layout. A significant feature of Elettra Commander is the use of iframes, which allow for the integration of different independent web pages into a single screen.

Iframes can be viewed as multiple browser tabs condensed into one, and they are managed carefully to mitigate security vulnerabilities, including clickjacking and cross-site scripting (XSS) attacks.

The multi-tab functionality, including tab grouping, remains a valuable feature. For instance, a tab that is not in focus can draw the user's attention by displaying a blinking LED-like sphere if any parameter requires urgent attention. This design enhances user interaction and ensures critical information is readily accessible.

## LAUNCHER

Elettra Commander facilitates the launching of native applications through Firefox using a custom protocol. The chosen keyword for this protocol is "puma," which serves as a valid identifier. When Firefox encounters an unknown protocol, it searches its configurations for an associated application to execute. In this case, we have configured the application to be a Python script that manages the launching of local applications.

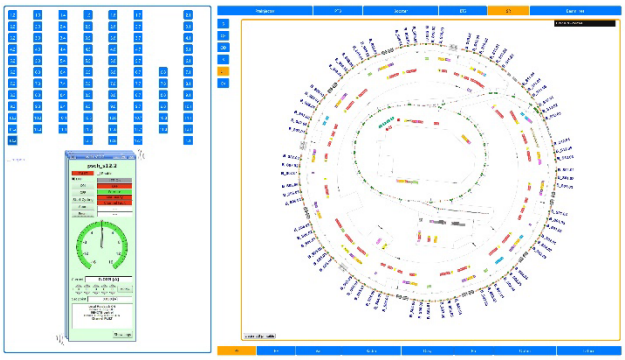


Figure 2: an already launched application is raised and vibrated for 1.5 seconds.

If the specified application is already running, the script brings it to the foreground and activates a vibration feature for 1.5 seconds to highlight its presence (Fig. 2). This

* lucio.zambon@elettra.eu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD076

# PYCUMBIA: BRIDGING HIGH-PERFORMANCE CONTROL LOGIC AND PYTHON SIMPLICITY FOR MODERN UIS

G. Strangolino, L. Zambon, Elettra Sincrotrone Trieste, Italy

*Abstract*

Pycumbia is a Python binding to the high-performance C++ cumbia framework, designed to simplify the development of control system applications without sacrificing responsiveness or scalability. It offers a user-friendly interface while maintaining the speed, concurrency, and low memory footprint of its C++ backend. By releasing Python's GIL, pycumbia ensures that GUI applications and data workflows remain smooth and responsive even under heavy load, a key requirement in control system environments.

From a user experience perspective, pycumbia significantly reduces the complexity typically associated with integrating control systems into custom applications. Developers can build advanced data visualization tools and synoptic panels with just a few lines of Python code, without dealing with polling logic, event dispatching, or thread management. Pycumbia also ships with PYI stubs for full IDE code completion.

A key architectural advantage is its flexibility in deployment: pycumbia can run inside an isolated miniconda environment, allowing developers to use up-to-date Qt and Python packages independently of the operating system, or system-wide, provided the base OS supports a recent enough software stack. This enables modern Qt-based graphical applications to be developed and run consistently across platforms, bypassing limitations of outdated system packages. Pycumbia in real-world control system applications improves both the developer experience and application performance.

## INTRODUCTION

In modern accelerator and laboratory environments, control system applications need to balance high performance with rapid development. While C++ frameworks such as cumbia provide robust concurrency, low latency, and minimal memory overhead, they can be less accessible to developers who prefer dynamic languages or rapid prototyping. Pycumbia bridges this gap by exposing the cumbia [1] C++ libraries directly to Python, combining the execution speed of C++ with the simplicity and flexibility of Python. This makes it possible to create responsive, real-time control system user interfaces (UIs) with minimal code, without compromising scalability.

## BACKGROUND TECHNOLOGIES

### *Qt Framework*

Qt is a cross-platform application development framework widely used in scientific software for building rich, interactive graphical interfaces [2]. The pycumbia project targets Qt 6, which brings C++17 support, improved rendering, better high-DPI scaling, and enhanced performance for both widgets and QML-based UIs.

### *Python for Control Systems*

Python [3] enables fast iteration, scriptable automation, and the integration of data processing, visualization, and UI logic in a single language. Pycumbia releases the GIL during C++ operations to maintain smooth GUI updates even under heavy loads.

### *Shiboken Binding Generator*

Shiboken generates Python bindings for C++ libraries with minimal overhead. Pycumbia uses it to map cumbia's C++ API to Python efficiently.

### *CMake Build System*

CMake manages cross-platform compilation of cumbia and pycumbia, simplifying integration with other dependencies.

## PYCUMBIA ARCHITECTURE AND ADVANTAGES

Pycumbia wraps the C++ cumbia libraries into Python, preserving their multithreaded architecture. Advantages include concurrency without complexity, near-native performance, UI responsiveness, reduced boilerplate, and IDE integration via `.pyi` stubs.

## AVAILABLE MODULES

### *PyCumbia (Base Library)*

The base module implements cumbia's multithreaded activity management in Python. Developers define activities (init, execute, exit), assign them to threads, and exchange data via CuData key–value bundles. This enables safe communication between background tasks and the main UI thread without manual locking.

### *PyCumbiaTango*

Integrates cumbia with the Tango control system [4], providing asynchronous read/write access to device attributes and commands. It supports both event-driven and polled updates. Connections, database access, and event subscriptions are managed transparently.

### *PyCumbiaQtControls*

Provides Qt widgets for control and display.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD077

# WRAP: INTEGRATING AN EVENT PROCESSING FRAMEWORK

E. Galatas*, S. Deghaye†, K. Tsiaras ‡, C. Roderick §, CERN, Geneva, Switzerland

*Abstract*

The Web Rapid Application Platform (WRAP) provides a centralized, low-code environment for building Graphical User Interfaces (GUIs) through an intuitive drag-and-drop workflow. These GUIs act as high-level user applications to the complex network of devices within CERN's accelerator control system. Some are relatively simple, displaying device data or setting control parameters as entered by the user. However, more advanced scenarios require correlation and processing of multiple asynchronous events triggered by independent devices. In such cases, where data from different sources must be synchronized and transformed before presentation, a more sophisticated processing layer is essential. To this end, an event processing framework has been implemented in TypeScript and integrated into the WRAP front-end. This framework allows users to express data correlations via event building logic and to implement lightweight data processing scripts. Given that the primary WRAP users are not software developers, a simple Domain Specific Language (DSL) was designed to express this logic in an accessible and declarative manner. This paper presents the motivation for the event processing framework, the design philosophy and architectural choices made, the technical implementation, and integration into WRAP. Challenges encountered are also described and future directions are outlined, including how the framework positions WRAP as a successor for GUIs developed on older platforms.

## INTRODUCTION

The development of CERN's Web Rapid Application Platform (WRAP) aims to replace the costly and time-consuming process of specialized application development, as well as reduce the personnel effort required for maintenance [1, 2]. A key missing feature limiting adoption is the ability for users to carry out computations and analyses on data streaming into their applications from control system devices [3]. These operations range from simple aggregations, such as maximum or average, to more advanced scenarios computing overall system statuses or applying multi-dimensional image filters.

### *Data Model*

WRAP interfaces with data through its Data Source model. This model defines a protocol and the interactions to which data providers must conform, ensuring a consistent interface for downstream consumers. A data source must provide a metadata description and support one or more of the following interactions:

- **Subscription to live data**: Streaming updates as device state changes occur.
- **Historical data query**: Retrieval of archived data points within a defined time range.
- **Control commands**: Ability to set values back to the device or source.

Currently, WRAP supports the following data source types:

- **Device data source**: Interfaces with physical or virtual control system devices via their property model [3]. If the target is a control field, its state can also be set. For devices registered in the Accelerator Logging Service (NXCALS) [4], historical data points can be fetched and used to backfill live updates.
- **NXCALS variable**: Provides access to archived data over a fixed time range, even if the device is inaccessible or no longer exists.
- **WRAP variable**: A user-defined value, specific to an application instance, that can be used to control or parametrise other data sources or widgets. Optionally, it can take its value from the URL, enabling generic applications operating over sets of devices with similar property structures, where the device name can be dynamically assigned.

## SCRIPTING CONSIDERATIONS

At its core, a script is a function over the values of one or multiple underlying data sources. However, several technical challenges must be addressed to meet control system requirements.

Combining values from multiple sources is non-trivial, as the values may be emitted independently with different timestamps and rates. No fixed strategy suits all scenarios; ideally, users or system experts tailor the synchronization approach per use case.

Once the data is correctly synchronized and batched, the specification of the processing logic remains open. Possible strategies vary widely: from simplified user interfaces with a limited hard-coded function set, to visual graph-based frameworks like Node-Red [5] and LabView [6], to full programming languages.

## EXISTING SOLUTIONS

We first considered existing tools and frameworks, both from industry and within CERN's control system, to provide such data interaction capabilities.

* epameinondas.galatas@cern.ch
† stephane.deghaye@cern.ch
‡ konstantinos.tsiaras@cern.ch
§ chris.roderick@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD078

# PREPARING THE MIGRATION OF PYTHON GUIS FROM QT5 TO QT6 AT CERN

I. Sinkarenko*, S. Deghaye†, C. Roderick‡, P. Mantion§, CERN, Geneva, Switzerland

*Abstract*

The use of Python in CERN's particle accelerator domain has grown steadily over the last decade. For non-software developers, it integrates well into the reality of a physics laboratory, where people who develop hardware or design operational workflows often author associated software, including Graphical User Interfaces (GUIs). While the general strategy is to standardise GUIs using no-/low-code solutions, complex use cases require more customised GUIs. For this, Python and Qt are combined through PyQt bindings.

Qt5 has been used for several years but with its end-of-life already reached, migration to Qt6 is a priority. At the same time, GUI applications in the CERN Controls ecosystem require long-term stability that exceeds the software industry average. Therefore, the migration requires a lot of care, taking into account both maintenance costs and long-term risks associated with reliance on a 3rd-party technology. To plan the upgrades, alternatives to the current PyQt approach have been reviewed. In the past, PyQt5 was an obvious choice, but investment into the PySide library by the Qt Company now presents more options.

We outline the criteria considered in our comparison of PyQt, PySide, and QML, and present the resulting decisions and rationale behind them. Finally, we discuss long-term risks associated with Qt bindings for Python and describe how we will manage the migration to future versions of the Qt framework, such as Qt7.

## INTRODUCTION

Since Python became an officially supported language of the accelerator controls group at CERN in 2019, it has catalysed the development of a substantial body of new software [1]. This includes Graphical User Interfaces (GUIs) of varying complexity, powered by PyQt, the Python bindings for the Qt framework. PyQt5, in particular, has been a mainstay for several years, facilitating user applications in control rooms, hardware development in laboratories, and personal research initiatives. Despite the fact that Qt5 (the underlying C++ library upon which PyQt is based) has reached its official end-of-life, it remains the predominant version of the toolkit within CERN's Accelerators and Technology Sector (ATS). The prerequisites for migrating to the newer Qt version have only been fulfilled recently, with the move away from CentOS7 [2] which was too old for Qt6.

At CERN, PyQt initially emerged as an alternative to the long-established GUI technology, Java Swing, which served as the primary GUI toolkit for over two decades. However, Oracle's lack of investment in Swing, combined with the steep decline in its wide-spread usage, raised significant concerns and prompted the search for viable alternatives. Simultaneously, the growing volume of in-house control room software and the associated challenges of maintaining an expanding portfolio encouraged a reassessment of the GUI strategy at a higher level. After several iterations, the new strategy took shape, advocating for no-code or low-code solutions for standardised or simpler scenarios, typically represented in a dashboard-like format. This approach left more complex use cases to be addressed through traditional imperative programming, either in Python (PyQt) or Java (Swing). A new web-based solution, WRAP [3], emerged to address the no-code requirements. The CERN community of equipment experts and operations teams soon began to favour Python over Java for new GUI projects. This was driven by the increasing global popularity of Python, which was particularly apparent with new recruits, and also coincides with the rise of AI.

In light of these developments, it was anticipated that Qt6 would become the sole recommended framework for complex desktop GUI applications. For most CERN teams, this implies migrating their existing Swing-based software. The planning for such substantial changes is typically aligned with CERN's accelerator schedule, which allocates 2-3 year periods for Long Shutdowns (LS) dedicated to upgrades of the Large Hadron Collider (LHC) and its injectors, followed by several years of physics operation, known as a *Run*. The upcoming third Long Shutdown, LS3 (starting mid-2026), represents the deadline by which the migration strategy, framework, and supporting tools must be fully developed. However, certain critical applications will undergo migration prior to this, subject to their operational criticality, driven by the fact that Qt5 has already reached its end-of-life. Concurrently, an evaluation of the long-term stability and associated risks of the chosen framework is crucial, as the software must remain stable and operational throughout the duration of the next Run, with the likelihood of continued use well beyond that time frame.

## PYTHON QT LANDSCAPE

### *PyQt and PySide*

In the context of Qt for Python, two primary options are available: PyQt, a well-established third-party library, and PySide, which is currently developed by The Qt Company, the organisation behind the Qt framework itself. At first glance, one might expect PySide, being developed by the same entity responsible for Qt, to offer superior integration. However, PySide was effectively abandoned for a number

* ivan.sinkarenko@cern.ch
† Stephane.Deghaye@cern.ch
‡ Chris.Roderick@cern.ch
§ pierre.mantion@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD079

# ENHANCING CERN EXPERIMENTAL AREAS BEAMLINE CONFIGURATION MANAGEMENT WITH DATA-DRIVEN SCHEMATICS

A. Moshenska*, G. Romagnoli , P. Le Roux, L. Burdzanowski,
CERN, Geneva, Switzerland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The beamlines in CERN's experimental areas are continuously evolving to meet the demands of a rich and diverse physics program. Various data management systems are used to efficiently plan, document, and track this evolution, describing past, present, and future beamline layouts, including equipment positions, dimensions, connections, assemblies, and their relationships with underlying physical assets.

To support this work, a data-driven web application was developed. The application generates schematics for experimental area beamlines and particle accelerators. The paper highlights the added value of this solution, its technical design and key implementation aspects, emphasizing the role of effective data-driven visualization.

A case study demonstrates the practical application of the software. Additionally, the paper discusses future improvements, including the ongoing development towards data-driven synoptics with associated intuitive beamline controls.

## INTRODUCTION

*Schematics* is a web-based, data-driven application that automatically generates accurate, up-to-date beamline diagrams showing each device's position, dimensions, classification and related details.

Developed as part of the North Area Consolidation Project [1], Schematics facilitates access to relevant configuration data, enhancing quality and configuration management for both data managers and end-users. It supports planning and daily operation of CERN accelerators and Experimental Area beamlines [2] by providing a clear, interactive view of accelerator configurations.

Integrated as a module of the Layout web application [3], Schematics retrieves all related data directly from the Layout Database [4]. The data-driven nature of the tool is designed to reduce the time required for reviewing and correcting underlying configuration data. Additionally, the tool aims to improve accuracy and increase efficiency by generating printable drawings that have to date been manually prepared.

## PROJECT ORGANIZATION

The development of Schematics was organised as a cross-team project, with well-defined roles engaging project members from different organisational units to fulfil the typical software engineering roles linked to:

1. Business Analysis: Addressing end-user engineering requirements, functional specifications, and validation, acting on behalf of end-users and providing the necessary business know-how to the development team.
2. System Architecture and Data Engineering: Ensuring technical consistency with existing solutions such as the Layout Service, and anticipating future development of related projects such as data-driven Synoptics.
3. Software Engineering: Developing, testing, and supporting the evolution of the tool.

Traditional agile methodologies were followed, incrementally delivering new features and improvements in iterations of 2–3 weeks long. Milestone-oriented planning, with a clearly defined scope of the Minimum Viable Product (MVP) at each stage, was essential to deliver a fully working solution within a span of several months. The milestones were decomposed into specific tasks, varying from several hours to 2–3 days of estimated development time. Project tracking was facilitated by Atlassian JIRA. Each iteration was accompanied by continuous integration and testing. Regular sprint reviews and retrospectives facilitated early feedback, ensuring that technical and business objectives were aligned.

## INTERFACE DESIGN AND PROJECT IMPLEMENTATION

### *Technology Stack and Architecture*

Schematics is built following a standard three-tier architecture, which organizes an application into three logical computing tiers: presentation, application, and data [5].

The presentation tier is the user interface of the application and is developed using HyperText Markup Language (HTML), Cascading Style Sheets (CSS) and the Angular framework (a standard across all CERN Accelerator Controls web applications). The application tier processes the information collected in the presentation tier. For Schematics the application tier is developed using the Spring Boot framework. The data tier is based on the Oracle Relational Database Management System (RDBMS), with the application tier using Structured Query Language (SQL) to manage the data.

### *User Controls*

The main goal of Schematics is intuitive beamline visualization and hence there is a strong focus on the presentation tier. This tier renders beamlines, and offers users a number of functionalities to control rendering and to maximize the customization of the beamline visualization for user convenience. One of the first mockups of basic functionality of the Graphical User Interface (GUI) is shown in Fig. 1.

* anastasiia.moshenska@cern.ch

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD082

# STREAMLINING PHOEBUS OPI SCREEN CREATION AND MAINTENANCE WITH OPI-GENERATOR *

T. Zhang†
Facility for Rare Isotope Beams, Michigan State University, East Lansing, MI, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

With the transition from Eclipse-based Control System Studio (CS-Studio) to the Phoebus platform now complete, the focus has shifted to developing and maintaining Operator Interface (OPI) screens for Phoebus. The *opi-generator* package offers a streamlined solution for generating properly formatted OPI XML files directly from Python scripts. Designed to simplify the creation and ongoing maintenance of OPI screens, *opi-generator* provides unified customizable screen styling and fine-grained control over widget layout at the pixel level. By automating OPI screen generation, this tool minimizes manual editing, enhances consistency, and accelerates development workflows for Phoebus-based CS-Studio environments.

## BACKGROUND AND MOTIVATION

CS-Studio [1] has been a cornerstone in the development and operation of accelerator control systems, which provides a set of tools for monitoring and controlling accelerator components. However, with the transition to the Phoebus [2] platform, which is developed on JavaFX as the new graphics solution, and for new features, there is a need to adapt and enhance the existing tools and workflows, for instance, to maintain the legacy OPI pages originally created for the CS-Studio Eclipse platform [3].

FRIB utilizes a modern accelerator facility based on superconducting radio-frequency (SRF) linear accelerator (linac) technology, to accelerate a wide range of ion species to produce rare isotope beams [4]. EPICS [5] is used as the control system solution to manage the complex interactions and operations of the accelerator components. Debian Linux is the operating system of choice for the FRIB control system infrastructure. Initially, CS-Studio Eclipse was adopted as part of the tools for high-level physics control application development on Debian 9 (Stretch). As the OS evolved and Phoebus matured as a replacement, the transition to Phoebus began with Debian 11 (Bullseye), and since 2023, only Phoebus is supported on Debian 12 (Bookworm).

*opi-generator* [6, 7] was developed at a time when both CS-Studio Eclipse and Phoebus coexisted in the controls network, to minimize the maintenance effort required to support both platforms.

It is worth mentioning that Phoebus provides tools to convert legacy OPI screens from the CS-Studio Eclipse format to the new Phoebus format. However, these tools do not always work perfectly, and the converted screens often require manual adjustments. *opi-generator* was developed to address these challenges by providing a more efficient workflow for OPI screen creation and maintenance.

On the other hand, both Phoebus and CS-Studio Eclipse provide GUI tools for OPI building, which are user-friendly and allow for intuitive screen design. However, they are not suitable for all use cases, particularly for large-scale or highly customized OPI screens involving scripts and rules. Maintaining these UI-generated OPIs is also a challenge, as they may not easily accommodate changes in requirements or design.

Instead, *opi-generator* offers a way to define the OPI layout and internal logic with Python scripts, providing a more flexible and powerful approach to OPI development. Most importantly, everything is maintainable through code.

As *opi-generator* is actively developed, support for other display managers [8] is under consideration. In this context, the developed scripts could produce proper OPI files for different display managers.

## SYSTEM ARCHITECTURE

The *opi-generator* project was forked from *cssgen* [9] in the summer of 2022, extending support to both CS-Studio Eclipse and Phoebus. It also adds advanced features and customization options, such as script and rule support and new data visualization widgets. As such, *opi-generator* is a new project with its own roadmap and design. It is also worth noting that, at the time, a project called *phoebus-gen* [10] existed for generating OPI files exclusively for Phoebus.

### *Widget Attribute Mapping*

An OPI file is a structured XML document that defines the layout and behavior of the user interface elements (widgets) within the screen. Each widget in the OPI file can have various attributes that control its appearance and functionality.

For instance, the widget for monitoring text updates is called `TextUpdate` in both Phoebus and CS-Studio Eclipse, but uses different XML attributes, reflecting differences in their underlying architectures. The XML widget element is named `org.csstudio.opibuilder.widgets.TextUpdate` in CS-Studio Eclipse (as `typeId`), while in Phoebus it is `textupdate` (as `type`). However, the base attributes like `x`, `y`, `width`, and `height` are nearly the same, so

```
text = TextUpdate(x, y, width, height,
pv_name)
```

would represent the same widget displaying the value of a PV named `pv_name` in both platforms.

* Work supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics, and used resources of the FRIB Operations, which is a DOE Office of Science User Facility under Award Number DE-SC0023633 and DE-SC0024707.

† zhangt@frib.msu.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD083

# Redesign of the Timeline Generator at Fermilab using a web-based Flutter application, GraphQL API and an IOC*

L. Carmichael[†,1], M. Austin[,1], S. Johnson[1], S. Wang[2]
[1]Fermi National Accelerator Laboratory, Batavia, U.S.A.
[2]Argonne National Laboratory, Lemont, U.S.A.

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The control system at Fermilab is undergoing an evolution with a shift towards web-based applications with connections to the EPICS infrastructure. The Timeline Generator (TLG) is an application that serves to coordinate events across the lab using different timing links. These links include the Tevatron clock (TCLK), a 10 MHz serial link with events encoded at 20Hz and Machine Data (MDAT), a communication link with states encoded at 720Hz. This paper covers the redesign of the major components of the TLG. This includes a web-based Flutter application for building timelines. A placement service is in use that has a GraphQL interface and uses a timeline input to compute a schedule of events and states. The Flutter application sends this computed schedule to the TLG IOC

via a GraphQL interface to the Data Pool Manager (DPM). The TLG IOC runs on an Arria FPGA, the Accelerator Clock Generator (ACLK-GEN), which is responsible for writing the events and states on to the different timing links.

## INTRODUCTION

The Fermilab control system is evolving to a multi-tier control system [1] whose user interface consists primarily of web displays. Many of these web displays will be implemented as Flutter applications. There is a middle tier which utilizes the Data Pool Manager (DPM) middleware to provide access to both ACNET and EPICS frontends. Figure 1. provides a simplified view of the control system where the Role-Based Access and Control (RBAC) module and DPM, an in-house development, are key components of this infrastructure.

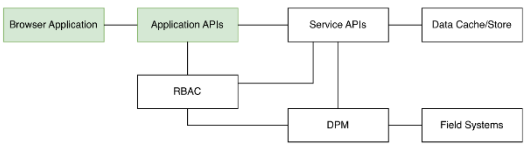



Figure 1: User Interface Components.

RBAC controls access to system resources by using individual user roles to authenticate users.

Fermilab utilizes multiple timing links to coordinate events across the lab. These links include the Tevatron clock (TCLK), a 10 MHz serial link with events encoded at 20Hz and Machine Data (MDAT), a communication link operating at 720Hz with states encoded at 20Hz. Additionally, the Accelerator Clock (ACLK) [2] is a new timing link that is currently under development. ACLK will include events encoded at 20 Hz with a data rate of 1.2Gbps..

The current TLG has a console application as its user interface and communicates with the TLG VME board that plays events on the TCLK and MDAT timing links.

The new TLG will use a thin client Flutter application as its web interface. Its business logic will be contained in a separate placement service, and it will use GraphQL, an open-source architecture, to communicate with this service. The application will use the DPM middleware to access the TLG IOC. This IOC serves to front end the ACLK Generator (ACLKGEN), a FPGA that writes events and states out to the ACLK, TCLK and MDAT timing links.

All the Flutter applications and services will be run in Kubernetes containers allowing for high availability and robustness. GitHub has been chosen to manage software versioning.

This paper will explore the different layers of the TLG architecture and provide status on the ongoing development.

## SYSTEM OVERVIEW

The TLG infrastructure includes a high-level Flutter application to build timelines, a placement service to compute event placement and an IOC to handle the playing of these events. The TLG architecture consists of 3 tiers and is illustrated in Fig. 2.

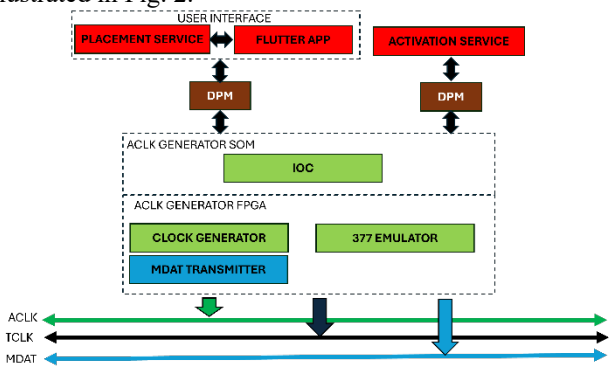



Figure 2: TLG System Overview.

The top-level Flutter application serves as the user interface. Operators use this app to build a schedule of events and states to be sent to the TLG IOC to be played out. It uses the DPM middleware to communicate with the TLG IOC via readings and settings. Additionally, its business

* This work was produced by FermiForward Discovery Group, LLC under Contract No. 89243024CSC000002 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics.
† carmichael@fnal.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD084

# AN AGILE/XP SOFTWARE DEVELOPMENT PROCESS FOR MODERNIZING THE ACCELERATOR CONTROL SYSTEM AT FERMILAB

J. S. Diamond, Fermi National Accelerator Lab, Batavia, IL, United States of America

*Abstract*

Fermilab is undergoing the most ambitious upgrade to its accelerator control system of the 21st century. As part of the ACORN project, hundreds of legacy control system applications written in C/C++ will be re-imagined and developed from the ground up. In addition, applications to support Fermilab's new super-conducting linear accelerator are already under construction. To manage the development of modern controls applications, the Controls department has adopted an Agile software development process based on eXtreme Programming. In this paper we will describe our process and detail our experience applying it to the development of two case studies.

## INTRODUCTION

Fermilab is the premiere particle accelerator facility in the United States hosting five particle accelerators spread out over 6800 acres. A sophisticated, distributed control system known as ACNET has been at the core of operations for this facility since the era of PDP-11s. Over the decades the control system computing hardware has migrated from mini-computers to modern day multi-core Intel-based servers, but an alarming amount of the original software remained unchanged.

To address this build-up of legacy software and ensure a sustainable controls system for the next generation of scientific discover at Fermilab, the ACORN project was initiated. ACORN is a DOE O413b project with an estimated cost of $160MM to modernize the accelerator control system. Roughly $30MM will be dedicated to modernizing the software components of control system – end user applications, central services, and data acquisition pipelines.

PIP-II is another DOE O413b project to build a new super-conducting linear accelerator to replace Fermilab's existing linear accelerator and boost proton intensities that will support next-generation neutrino experiments. PIP-II requires new, modern control system applications to support this advance

## THE CAVEMAN SOFTWARE DEVELOPMENT WORKFLOW

The existing software development workflow in the Controls department at Fermilab is like the Caveman workflow. If the reader is not aware, the Caveman workflow goes like this: the Software Developer receives a set of vague requirements from their customer, usually in oral form and then retreats to their office (the "cave") where they work tirelessly on the project for several weeks or months and emerge once the software is "complete" and ready for testing. At this point the real requirements can be gathered from the customer and (if the customer is lucky) the software developer begins the second phase of software development – twisting the code into a usable product. This phase continues until the customer is satisfied or the software developer is assigned a different task.

This type of software development workflow cannot scale to meet the volume of software that ACORN aspires to deliver. A more modern, scalable approach to developing software is necessary. Most software development teams use some form of Agile framework to build enterprise software.

However, one does not have to search far or long to discover that Agile is controversial in the software industry. Just go to TikTok or YouTube and search for Agile or Scrum if you are unaware of this controversy. Most of the negativity related to Agile software development is related to the number of meetings. Software developers dislike meetings, or rather the meetings that are more about collecting status and less about productivity. There is nothing in the Agile Manifesto about preferring meetings over writing code, so there is a disconnect between how some teams practice Agile and what being Agile really means.

## IS AGILE AN IMPROVEMENT

First recognize that Agile is not a software development process. The original signatories of the Agile Manifesto (cite) did not publish a book or even a "how to Agile" blog post. As most are aware by now, the Agile Manifesto is more of a philosophy for software development. Several software development frameworks that claim to be Agile do exist and it's these frameworks that are used to guide a team of software developers through completion of a software product. Common Agile frameworks in use today include Scrum, Scaled Agile Framework (SAFe) and eXtreme Programming.

One of the co-authors of the Agile Manifesto, Robert C. Martin, published a book on what he believes went wrong with the Agile movement in the 2010s. According to Martin, the overemphasis on the management practices (burndown charts and daily meetings) and the de-emphasis of the technical practices (test-driven development, pair programming and refactoring) results in the development of brittle code that's difficult to extend and maintain after reaching a critical mass of about thirty thousand lines of code.

## OUR AGILE PROCESS

Our Agile process leans heavily into the technical practices of eXtreme Programming – continuous integration and continuous delivery, refactoring, simple design, test-driven development and pair programming. We seek to build software that can be easily changed and maintained

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD087

# CSS TO PHOEBUS TRANSITION AT KARLSRUHE INSTITUTE OF TECHNOLOGY

E. Blomley*, J. Gethmann, J. L. Steinmann, M. Schuh, A.-S. Müller
Karlsruhe Institute of Technology, Karlsruhe, Germany
S. Marsching, aquenos GmbH, Baden-Baden, Germany

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The two accelerator test facilities KARA and FLUTE at the Karlsruhe Institute of Technology have been using Control System Studio as the main GUI for over ten years. The migration to Phoebus allows us to renew the full middle-layer stack of services, including alarm server, logbook, channel finder, and save-and-restore system with an emphasis on a modular and containerized setup. The panels are served via a central, highly available web server stack, omitting the need to maintain local versions of our panels. This allows us to use continuous integration (CI) and continuous deployment (CD) workflows to gather and publish panels from various different sources, such as EPICS modules, without having to copy and maintain such panels into our panel tree. This paper describes the technical setup of our full Phoebus stack and reports on the migration process of our existing panels.

## MOTIVATION

### *Software Environment*

The two accelerators KARA (Karlsruhe Research Accelerator) [1] and FLUTE (Ferninfrarot Linac und Test Experiment) [2] each operate in separate, self-sufficient network environments, including separate services such as DNS, DHCP and NTP. But over the last few years, we have started migrating as many services as possible to central and automated deployments, decreasing overhead and maintenance costs, while improving efficiency and deployment efforts. By now, most systems are virtual machines, which are managed via a Proxmox VE cluster [3]. Dedicated computer hardware is only used if required, e.g. for the operator terminals. Most computers run Ubuntu long-term support releases as the operating system (OS). Windows is only used if absolutely necessary, for example, to run certain third-party software.

### *Accelerator Operations*

An on-site operator is only required during beam injection, which takes about one hour a day. The beam is then available for all beam lines for 23 hours until the next injection cycle. The operator on duty is only present in the control room as long as it takes to complete the injection process. Although the operator continues to be on call throughout the day, the control room is typically not staffed during the 23 hours of beam operations.

FLUTE is a smaller linear accelerator without a fixed beam schedule and dedicated operators. Operations at FLUTE focus on individual experiments. The experimenter is responsible for both operating the overall accelerator and the experiment itself.

### *Requirements*

This handling of operations has important implications and defines some of the requirements for our Phoebus distribution. It is crucial to run our control system panels from the office environment, to be able to monitor the accelerators, and provide remote support by domain experts. The office and accelerator environments should run on the same distribution, without the need to manage, maintain, and deploy different code projects. The transition between accelerators and offices should be as seamless as possible, always providing access to the latest version of the panels. The installation of Phoebus and the deployment of the panels need to be as simple as possible for the users, but it cannot rely on the software orchestration tools available inside the accelerator network. And it needs to run on Mac, Windows and Linux operating systems without depending on any pre-installed third-party software.

## OVERVIEW

The basic operational model of the accelerators outlined above has not changed dramatically over the last 10+ years. Some of the requirements were already addressed by the previous Control System Studio (CSS) setup. But the experience gathered during that period helped to identify key areas where improvements were possible.

### *Previous CSS Setup*

The existing CSS setup focused on two separately managed distributions using multi-level Subversion (SVN) repositories making use of the *svn:externals* feature. It used Python scripts to enable users (and developers) to start CSS in an agnostic way, independent of the OS, also taking differences in the network environment into account. These scripts took care of retrieving the correct binaries of CSS and the Java Development Kit (JDK) for the respective OS and network environment.

But some of the key issues and limitations were the following:

**Code Base:** Customizations of the core code base were never integrated upstream, leading to complexities in building newer versions and ultimately resulting in using an outdated version of CSS.

**Panel Development** In order to be able to share panels between both accelerators (e.g. panels for the same type

* edmund.blomely@kit.edu

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD088

# MAINTAINABLE WEB INTERFACES AT BEAMLINES*

S. De Leon[1†], F. Rodolakis[2], E. Gann[3], D. Allen[3], A. Kaufman[1], R. Hamlyn[4], E. Crumlin[1,4], N. Manha[1], K. Iliev[1], J. Mahl[1], P. Glans[1], A. Islegen-Wojdyla[1], A. Hexemer[1], D. McReynolds[1]
[1]Advanced Light Source, Lawrence Berkeley National Laboratory, Berkeley, CA, USA
[2]Advanced Photon Source, Argonne National Laboratory, Lemont, IL, USA
[3]National Synchrotron Light Source II, Brookhaven National Laboratory, Upton, NY, USA
[4]Chemical Sciences Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

At the Advanced Light Source (ALS) new strategies are being employed to create maintainable web-based controls and data analysis applications for beamline users. Our approach centers on creating modular library components available in different software languages that address requirements shared across multiple beamlines. A combination of modern tooling and packaging systems has enabled a development procedure that encourages reuse, consistency, and long-term maintainability.

## BACKGROUND AND CONSIDERATIONS

A large driver behind creating a new web based graphical user interface (GUI) for beamlines at the ALS is an increasing usage of the Bluesky/EPICS controls stack for beamlines at the facility [1, 2]. Bluesky (a Python library for data acquisition [3]) does not come with an overarching GUI that encompasses all its features, instead there are several GUIs split up by different services. This web interface will serve as a means to utilize multiple major features of Bluesky in a single cohesive interface, and is being actively developed as a contribution to the Bluesky project.

This web GUI must provide Bluesky compatible beamline controls, but also clean integration with analysis and AI powered web applications. The initial development is therefore sensitive to multiple considerations (Fig. 1) discussed further here.

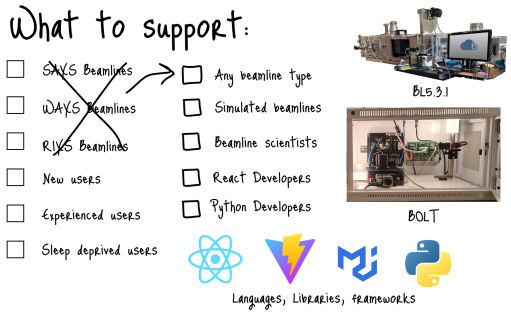



Figure 1: A requirements list for a controls web interface.

### *Consideration 1: Beamline/Facility Variety*

Beamlines require vastly different interfaces to suit not only the type of x-ray techniques that a beamline specializes in, but the different physical devices used in the experiment process which could range from simple 1D stages to 6-axis motors. There is even significant variation in the few beamlines that are the first recipients of the web GUI software.

Early adopters of the web GUI software include the fully operational Small-Angle X-ray Scattering Beamline 5.3.1 [4], and the Resonant Inelastic X-ray Scattering Beamline 6.0.1 under commissioning. Another physically smaller recipient is the Beamline Optics Learning Testbed (BOLT) [5], a small testbed running EPICS with motors and cameras running photogrammetry.

Current and planned beamlines at the ALS have different user workflows and interface needs, so the corresponding web GUI framework must gracefully support this customization. Additionally, different facilities besides the ALS should be able to benefit from this development without needing to rework their software implementation.

### *Consideration 2: End User Variety*

Beamline users represent a large variety of people in the scientific community, including students, postdocs, visiting scholars, developers, and more. Users typically have a limited window of time to access and run the beamline, making every hour of productivity valuable.

Due to limited time and beamline variation, it is unrealistic to assume that the end user is an expert in that beamline's interface or the underlying controls system. Therefore a successful interface must be able to abstract away complexity and lead unfamiliar users through successful experiments.

### *Consideration 3: Developer Variety*

Software developers themselves add significant variety to the ecosystem. A single beamline can utilize different applications created by different developers. One app could handle running the beamline itself, and another could be used to process and group detector images with machine learning.

If the main goal is to meet the needs of beamline users, then allowing the best language or framework for the job promotes a better solution. The challenge is then enabling consistent look and feel between different applications, as well as making code reusable across languages/frameworks.

* Work supported by the Advanced Light Source, a U.S. DOE Office of Science User Facility under contract no. DE-AC02-05CH11231
† seijdeleon@lbl.gov

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD089

# AUGMENTED REALITY FOR ACCELERATOR OPERATIONS: A VIRTUAL CONTROL ROOM PROTOTYPE*

T. Hellert†, M. Dach, A. Hessler, E. Lam, F. Sannibale
Lawrence Berkeley National Laboratory, Berkeley, CA, United States

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

Particle accelerator control rooms have traditionally relied on fixed, on-site workstations, which limits operational flexibility and remote collaboration. While a previous proof-of-concept demonstrated the feasibility of a Virtual Control Room (VCR) using a web-based client and an intermediary laptop, this new work presents a fully implemented, native visionOS application that eliminates the need for an external host. The new system leverages native Apple frameworks such as SwiftUI and RealityKit to provide a standalone, high-performance solution with integrated secure remote access and real-time holographic collaboration. This paper outlines the native system's architecture, details its implementation, and presents key performance metrics, demonstrating a significant improvement in usability, latency, and security. We conclude that this native application establishes Augmented Reality HMDs as viable and effective platforms for next-generation interfaces for accelerator operations.

## INTRODUCTION

Particle accelerator control rooms are critical hubs where operators monitor and adjust accelerator parameters to ensure optimal performance. At facilities such as the Advanced Light Source (ALS) at Lawrence Berkeley National Laboratory [1], operators must continuously coordinate diagnostics, feedback loops, and machine settings to deliver stable, high-quality beam for a large user community. The demand for remote accessibility has grown, but traditional remote setups often lack sufficient screen space, leading to poor situational awareness and reduced efficiency.

To address these limitations, our previous work [2] presented a proof-of-concept for a Virtual Control Room (VCR) using an Apple Vision Pro headset [3], an intermediary laptop, and a web-based VNC client. This prototype successfully demonstrated the viability of immersive interfaces for accelerator operations, but its architecture introduced dependencies and potential latency issues.

Building on that foundation, this paper presents a fully-fledged VCR implemented as a native visionOS application. By eliminating the need for an external laptop, the new system allows the headset to directly handle remote connections and display management. Implemented entirely with Apple's native frameworks, the application provides a more robust, secure, and performant solution. We detail the system's architecture, describe its core features—including real-time holographic collaboration—and discuss results from its initial deployment at the ALS.

* This work is supported by the U.S. Department of Energy (DOE) under Contract No. DE-AC02-05CH11231.

† thellert@lbl.gov

## SYSTEM IMPLEMENTATION

The native visionOS application represents a major architectural shift from the previous proof-of-concept. Instead of relying on an external laptop, the Apple Vision Pro now acts as a standalone client that directly establishes secure connections and renders virtual displays, see Fig. 1. The system consists of the following core elements.

**Native App Frameworks:** The application is developed entirely in Swift, using SwiftUI for declarative interface design and RealityKit for spatial computing. SwiftUI simplifies the creation of complex layouts, while RealityKit manages the positioning and rendering of multiple virtual screens with precise control over scale, placement, and orientation.

**VNC Client:** Replacing the earlier web-based noVNC client, the application integrates a native implementation based on the libvnc library [4]. Processing the VNC stream directly within the app minimizes latency and enables fluid real-time visualization of accelerator data.

**Secure Remote Access:** The system incorporates a built-in SSH client that establishes encrypted tunnels from the headset to the control-room servers, removing the need for an external bastion host. This reduces the attack surface and improves performance by eliminating intermediate hops.

**Input and Interaction:** Bluetooth keyboards and mice are natively supported, with input automatically directed to the active virtual display. The Vision Pro's hand-tracking provides an alternative input mode, ensuring basic navigation even without peripherals.

**User Interface (UI):** The application supports customizable, multi-station layouts that can be saved and recalled, enabling consistent operator setups. Multiple stations can be accessed with a single password, simplifying session management. In the current prototype, screens are arranged in a fixed grid, which limits 3D flexibility but allows the entire workspace to be repositioned or resized as one unit.

**Holographic Collaboration:** While being a key motivation of the VCR concept, holographic collaboration is not yet implemented natively in the application. Instead, the prototype leverages visionOS's built-in FaceTime shared-session capabilities: on a FaceTime video call, one participant shares the VCR app, which then becomes visible across the connected sessions. Remote collaborators appear as avatars with realistic facial expressions and hand gestures, enabling spatially accurate pointing and a strong sense of co-presence. Future development will focus on integrating these capabili-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD090

# STREAMLINING AND UPDATING USER INTERFACES FOR LANSCE OPERATIONS USING PYTHON*

G. DeLaTorre†, Los Alamos National Laboratory, Los Alamos, NM, USA

*Abstract*

This paper describes the work and design process for updating user interfaces used for operations at Los Alamos Neutron Science Center (LANSCE) from Tcl (Tool Command Language) to the Python language. Python has become a golden standard in the software industry and has a wide variety of libraries and plugins that can be leveraged for a variety of projects. Utilizing the pyQt5 library, alongside internal libraries designed to interface with the Experimental Physics and Industrial Control System (EPICS) that primarily drives the accelerator's technologies, we can greatly improve the usability and visuals for its legacy user interfaces. These updates are a multi-year effort, having already replaced several legacy Tcl interfaces in production.

## INTRODUCTION

Many of the graphical user interfaces that operators use every day at LANSCE were developed using Tcl along with its graphical user interface toolkit (Tk) in the mid-to-late 1990s. While at the time the language may have provided a quick solution for the creation of these screens, accelerator operations have outgrown them. The development environment of today uses new and more modern tools to implement user interfaces, tools that are more flexible and easier to improve. While Tcl is by no means a dying language, its inherent characteristics and the use case it was originally intended for severely limit it. We are now in the process of migrating these existing applications to the Python language, which is not only improving the existing functionality of the software, but also future proofing our environment for decades to come.

## TCL/TK

Tcl was created in 1988 [1] as a scripting language intended to be embedded within other existing applications. The language itself gained widespread praise for its ability to be used dynamically and be easily deployed and began to be used for many use cases beyond the original intent. Tk is a toolkit that expands the capabilities of Tcl to create native applications. While many things can be created and achieved using Tcl/Tk, many of its characteristics and issues have become a hindrance in our development environment.

### *Characteristics and Issues*

Despite its widespread use, Tcl has a much smaller community around it. This leads to less support for the language and a smaller number of direct or third-party libraries compared to its peers in the developer language space. The user interfaces produced using Tcl/Tk also tend to lack polish and feel outdated. Figure 1 shows an existing Tcl/Tk application currently in use at LANSCE.

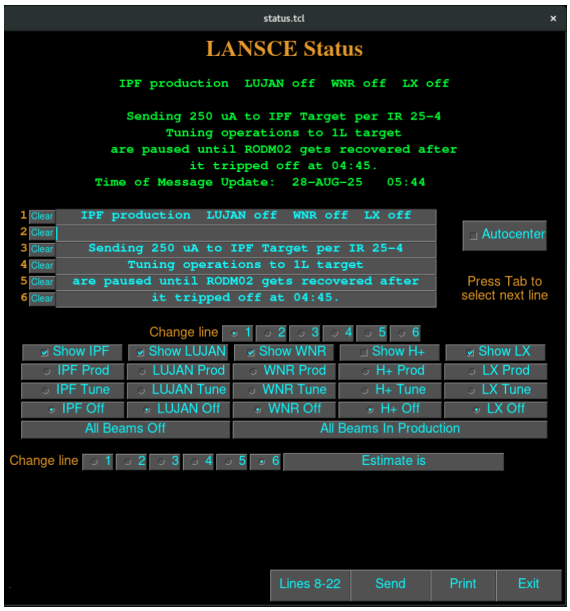



Figure 1: Tcl/Tk implementation of LANSCE Status screen.

Since it is a scripting language, its performance is much slower and unpredictable, and its handling of data is suboptimal. When writing code in Tcl, all data is considered a String, and it is the responsibility of the user to ensure the translation of the data is handled correctly when interacting with other applications or systems.

The language itself lacks the readability of more modern, object-oriented languages. Most applications consist of a single, large file that any developer without exposure to the language beforehand can struggle to navigate.

Alongside these shortcomings, perhaps the most pressing matter for operations here at LANSCE is the loss of support for some currently used Tk functionality in the release of the Rocky 9 Linux operating System, specifically the BLT toolkit. This obstacle alone is a major factor in our migration of these applications to Python. While there are several benefits for us to move to Python, future proofing by avoiding any further loss of support for Tcl tools is substantial.

* This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). LA-UR-25-29103

† gerg@lanl.gov

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

# PYDM: UPDATES AND ROADMAP

N. Stelter, J. Bellister, Y. Yazar, SLAC National Accelerator Laboratory, Menlo Park, CA, USA

*Abstract*

PyDM (Python Display Manager) is an open-source software platform designed to simplify the development of user interfaces for control systems, leveraging Python's power, flexibility and large 3rd party library support. PyDM provides a no-code, drag-and-drop system to make simple screens, as well as a straightforward Python framework to build complex applications. Over the past two years, significant enhancements have been made to PyDM, including adding support for PySide6 and steady improvements to the existing framework. We will highlight key advancements achieved over the past two years and provide an overview of planned developments aimed at maintaining PyDM going forward.

## UPGRADES

### *PySide6 Support*

Significant changes have been made to the PyDM codebase to support PySide6, the latest version of the PySide Python wrapper for Qt6. This has been implemented alongside the existing support for PyQt5. Swapping which wrapper PyDM is actively using can be done by setting the environment variable 'QT_API' equal to either 'pyside6' or 'pyqt5' (assuming PyDM is running in an environment with both wrappers installed). The 'About PyDM' window displays which wrapper is actively in use and it's version number, along with the version numbers of some other significant Python libraries actively in use by PyDM (see Fig. 1).




Figure 1: The "About PyDM" screen lists which Qt Python wrapper is currently active, along with library version info.

PyDM previously only supported PyQt5, a wrapping around the older Qt5. A primary reason for the upgrade to PySide6 was so PyDMcould be ran with Qt6, the newest version of Qt. Qt6 offers some new features and improvements that can potentially be utilized in PyDM. Some of these new features include improved font rendering, automatic scaling to high DPI monitors, and improved 3D rendering and concurrency, among others. Qt6 is also likely to provide improved performance due to its upgrade to C++17, but PyDM still needs to be profiled while running with PySide6 to confirm this.

When upgrading PyDM to support Qt6, it was necessary to choose an updated Python Qt wrapper, the most popular of which being PyQt6 and PySide6. PyQt6 and PySide6 are quite similar overall, but ultimately PySide6 was chosen for two reasons. First, PySide6 is offered under the LGPL license which avoids the need to purchase a commercial license or apply the GPL license to the codebase. Second, PySide6 is developed and supported directly by the Qt Company as their official Python Qt wrapper. This hopefully means PySide6 will stay more inline with Qt's development and be get more frequent updates than PyQt6 in the future. Qt5 has also reached the end of its support lifetime as of May 2025, and will no longer receive new maintenance patches [1]. This means Qt5 could have un-patched bugs and security issues in the future.

Previously, PyDM was updated to utilize the 'QtPy' abstraction layer [2] in an attempt to ease in upgrading to new Qt versions. QtPy provides an abstraction over the different Qt Python wrappers, so that PyDM code could ideally run unchanged across the different Qt wrappers and Qt versions. This would have allowed a straightforward upgrade to PySide6, but in practice significant changes were ultimately needed. One reason for this was that QtPy had not yet implemented some of Qt's features into it's abstraction layer, forcing PyDM developers to implement wrapper-specific sections in the code-base. Also to blame was existing PyDM code that either bypassed the abstraction layer entirely or had errors only evident when running on PySide6, requiring additional workarounds. Lastly, the upgrade was complicated by our requirement to keep support for PyQt5 alongside PySide6.

## NEW FEATURES

### *LED Multi-State Indicator Widget*

The PyDMMultiStateIndicator is a new widget that displays a color changing circle or rectangle, and sets its color using data from a binary PV. The widget changes color to represent one of 16 possible states, set by a connected channel's value. The design goal of this widget was to provide a concise indicator widget that displays the state of a binary PV, acting similar to a simple LED that changes color to display status on a hardware device. And PyDMMultiStateIndicator is already implemented into some PyDM screens, for example as a visual indicator of sensor status (see Fig. 2).

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD096

# MANAGING SHIFTS IN THE NA61/SHINE EXPERIMENT: DEVELOPMENT OF THE SHIFT SCHEDULER SYSTEM*

A. Dyśko, M. Słodkowski† for the NA61/SHINE Collaboration
Faculty of Physics, Warsaw University of Technology, Warsaw, Poland

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The NA61/SHINE experiment (SPS Heavy Ion and Neutrino Experiment) is a high-energy particle spectrometer at CERN, using the Super Proton Synchrotron (SPS). With the SPS providing protons and various ions, NA61/SHINE can study a wide range of nuclear systems. Its physics program focuses on strong interactions and supports research on cosmic rays and neutrinos. The experiment requires 24/7 shift coverage by collaboration members. To ensure continuity, quality, and efficient staffing, the NA61/SHINE shift scheduler system was developed. This web-based system, integrated with the collaboration's infrastructure, provides access to member roles, contact details, and assigned activities. Coordinators can manage runs and ensure optimal use of available staff. Members can independently book shifts based on their permissions. The system includes user and institute profiles with shift statistics and charts, visible within the collaboration for transparency. A key feature is the interactive shift calendar, which supports data-taking planning and duty assignments. The system can also generate calendar printouts with assigned shifts, helping to coordinate participation and maintain balanced staffing throughout data-taking periods.

## INTRODUCTION

NA61/SHINE [1, 2] is a multipurpose, fixed-target experiment at the CERN Super Proton Synchrotron (SPS), approved in 2007 following requests from the strong interactions, neutrino, and cosmic-ray communities. The detector is a large-acceptance hadron spectrometer, and its programme, ongoing in 2025, covers several domains. In the field of strong interactions, NA61/SHINE studies the onset of deconfinement and searches for the critical point of strongly interacting matter, with recent results including open-charm production and meson asymmetries in nucleus–nucleus collisions. For neutrino physics, the experiment provides precise hadron production data to reduce flux uncertainties, in particular for the T2K experiment. In cosmic-ray physics, NA61/SHINE measures hadron production and nuclear fragmentation to improve the understanding of cosmic-ray propagation, with reference data supporting the Pierre Auger Observatory and KASCADE. Furthermore, measurements of $\pi$–C and p–C collisions also serve neutrino experiments at Fermilab [3].

The tracking detectors of the NA61/SHINE experiment, in particular the large Time Projection Chambers (TPCs), measure the electric signals of charged particles produced in nucleus–nucleus collisions. The collected data are processed by a reconstruction procedure, which allows the identification of particles and the reconstruction of their tracks. The detector is about 13 meters long and 4 meters high, which is rather small compared to the experiments at the Large Hadron Collider (LHC). The NA61/SHINE facility is installed in one of the large experimental halls of the CERN North Area.

The NA61/SHINE experiment is not only a set of scientific instruments but also a network of researchers forming a Collaboration that must be effectively managed. The Collaboration currently consists of about 140 physicists from 28 institutions in 14 countries. Within this Collaboration, it is essential to organize the workflow efficiently across all stages of research: detector construction and upgrade, measurement periods, software development, and data analysis. Each member of the Collaboration is responsible for both scientific research and service work, and each institution is required to actively participate in data-taking.

Data is collected during specific periods of the year, called runs, when the CERN accelerator complex provides the beam to the experiment. Typically, there are two such periods each year: a shorter run for tests and hardware/software calibration, followed by a longer run dedicated to recording physics data. During these runs, the measurements must be supervised continuously without interruption. Institutions belonging to the Collaboration are obliged to provide personnel to staff the experiment. The crew is responsible for monitoring and maintaining the detector to ensure that it remains fully operational. A shift is defined as an 8-hour working period assigned to crew members (shifters), and corresponds to a full working day in the experiment's rotation scheme. Shifters are in charge of supervising the experiment.

## SOFTWARE REQUIREMENTS

The Shift Scheduler application supports the organization of shifts[1] during upcoming runs[2] of the NA61/SHINE experiment. The most natural solution for providing easy and widespread access is to implement it as a web application. The system must be scalable to ensure consistent performance over time, regardless of the amount of incoming data. The software architecture is designed to be modular and adaptable, with containerization ensuring reproducible deployments, simplified maintenance, and seamless integra-

* Work supported by IDUB-POB-HEP project granted by Warsaw University of Technology.

† marcin.slodkowski@cern.ch

[1] Working periods assigned to crew members.

[2] Periods of the year when the CERN accelerator complex provides the beam to the experiment.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD097

# A SERVER FOR THE ATLAS TIME MACHINE*

M. Torres, K. Bunnell, A. Ramaswamy, B. Blomberg, C. Dickerson,
G. Dunn, D. Novak, D. Stanton
Argonne National Laboratory, Lemont, IL, USA

*Abstract*

The Argonne Tandem Linear Accelerating System (ATLAS) facility at Argonne National Laboratory is a National User Facility capable of delivering ion beams from hydrogen to uranium. The existing Tune Archiving System (TAS), which utilizes Corel's Paradox relational database management software, is responsible for retrieving and restoring machine parameters from previously optimized configurations. However, the Paradox platform suffers from outdated support, a proprietary programming language, and limited functionality, prompting the need for a modern replacement.

The new system is composed of a modular architecture featuring a separate user interface, a time-series storage database, and a backend that connects the two. The new backend employs the FastAPI framework with WebSockets for asynchronous communication, and integrates Pydantic and SQLAlchemy ORM to enable a type-safe, object-oriented interface with SQL databases. This upgraded system significantly improves upon the legacy Paradox-based solution by offering a more robust, open-source architecture with enhanced reliability, maintainability, and ease of use.

## INTRODUCTION

The operators of ATLAS use historical beamline device information, as well as data concerning their present experiments, to quickly tune the beam to run a new experiment with high accuracy. However, TAS suffers from limited functionality, a proprietary programming language, outdated support, and software that is regularly prone to error. The user interface also merits substantial improvement.

To address the many issues inherent in the implementation of Paradox for the ATLAS control system, we have developed a system involving several databases, and a python-based server and user interface. This new system is called the "ATLAS Time Machine," (ATM) because it allows operators to easily replicate beamline configurations from past experiments on the current beamline, and change these configurations as needed- as if they were sending the beamline back in time.

The focus of this report is the ATM's server, which takes requests from client endpoints and writes or reads to the ATM's databases accordingly.

* This work was supported by the U.S. Department of Energy, Office of Nuclear Physics, under Contract No. DE-AC02-06CH1357. This research used resources of ANL's ATLAS facility, which is a DOE Office of Science User facility.

## THE ATM: A BIRD'S EYE VIEW

In this section, we will briefly describe a high-level overview of the ATM. For a more detailed-yet-holistic overview, see "Upgrading ATLAS's tune archiving system" [1].

The ATM is made up of three main parts: the databases, the server, and the clients.

### *Databases*

ATM uses MariaDB v10.3 to store experiment parameters and comments. The experiment parameters are added after operators enter the information for a new experiment in the control system. Comments are stored by operators while using ATM to describe the status of the tune and operations

InfluxDB v2.7 is used to store all historical device data. 25 years' worth of historical data from TAS was imported into InfluxDB. New device data is saved every 5 minutes, or by operator demand.

### *The Server*

The server handles the majority of the logic of ATM, by performing read and write requests to InfluxDB and MariaDB, and organizes the data from these databases to send to the client, using FastAPI v0.116.1 as a framework.

### *The Client*

The client is the interface for the users of the ATM. It features a GUI built with PySide6 v.12 that allows the users of the ATM to access all of its features, such as finding stored experiments that meet certain criteria, viewing the data in those experiments, and loading data to the beamline. The client code also handles the user-heavy logic of building configurations for loading to the beamline.

## COMPONENTS OF THE ATM SERVER

The ATM server can be broken up into the following main components: A FastAPI framework which utilizes SQLAlchemy, WebSockets, Open ID Connect (OIDC), and InfluxDB. In this section, we will examine the properties of each of these components and how they contribute to the server.

### *FastAPI*

FastAPI is a platform based on the OpenAPI standard that uses Uvicorn to implement a server based on the Asynchronous Server Gateway Interface (ASGI) standard [2], with Starlette as the basis for frameworking. FastAPI also uses Swagger UI to generate documentation, and Pydantic for fast type validation and JSON-data formatting [3].

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD098

# EPICS CONTROL SYSTEM FOR BEAM OPERATION AT SARAF LINAC ACCELERATOR

A. Gaget†, A.Chance, A. Choquet, G. Desmarchelier, J. Dumas, G. Ferrand, F. Gougnaud, T.Joannem, V. Nadot, S. Monnereau, CEA Saclay IRFU, Gif sur Yvette, France
E. Reinfeld, I. Shmuely, L. Golan, H. Isakov, A. Perry, Y. Solomon, N. Tamim, L. Weissman SNRC, Yavne, Israel

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The CEA Saclay Irfu is responsible for most of the EPICS control system for the Saraf Linac accelerator located at Soreq in Israel. This scope includes the control and the tuning of the beam. The accelerator is commissioned by disciplines (vacuum, RF, cryogenics…) and sections. For this purpose, we have developed a high-level application that can activate a set of subsystems to prepare the accelerator for the desired configuration. This application, called Beam Operation Mode (BOM), transmits and controls the operator requests of beam mode and destination to a set of subsystems. The Beam Destination Master (BDM), part of the Machine Protection System (MPS), is responsible for ensuring, among other things, that interceptive elements are not in the beam path, that power converters are in the expected state and that the vacuum conditions are correct. The Section Beam Current Transmission (SBCT), also an element of the MPS, monitors the current transmission along the accelerator using ACCTs and Faraday cups must be configured for the desired destinations. The Cavity Phasing system verifies that the cavities are properly conditioned. The Timing System (TMG), another element of the MPS, relies on fibre optics and its network must be monitored and operational. The Max Pulse Width (MPW) application limits the beam power by restricting the pulse width, according to the inserted devices. This paper presents the different components and their functionalities.

## INTRODUCTION

SNRC and CEA collaborate to the upgrade of the SARAF accelerator to 5 mA CW 40 MeV deuteron and 35MeV proton beams (Phase 2) at 176MHz [1]. The CEA control team has been actively contributing to the EPICS-based development of the accelerator, with major applications such as the timing system [2] and the Machine Protection System (MPS) [3]. This paper focuses on the set of applications that enable beam operation.

The Beam Operation Master (BOM) coordinates all other applications required for machine commissioning and protection. The Beam Destination Master (BDM) verifies that all the conditions necessary to deliver the beam to a given destination are correctly fulfilled. The Section Beam Current Transmission (SBCT) monitors the beam current along the accelerator and detects any losses. The Timing System (TMG) shapes the beam by defining the pulse widths of the chopper, magnetron and RFQ.

The Cavity Control (CAV) application ensures that the RF cavities (the RFQ, 3 rebunchers and 4 cryomodules accelerating cavities) are properly configured and operational according to the desired destination, providing a simplified setup process for operators.

Finally, the Max Pulse Width (MPW) module limits the maximum allowed beam power depending on the destination and the state of the inserted diagnostics, by acting on the beam pulse width.

## BEAM OPERATION MODE (BOM)

### *Overview*

The main function of the BOM is to assist the operator in applying the pilot beam and configuring the beam settings. The pilot beam is a low-intensity beam used during beam start-up to validate various parameters without risking damage to the accelerator. During beam start-up, the pilot beam operates at 50 µs, 1 Hz. The BOM executes a sequence to ensure every aspect of the accelerator is ready for the pilot beam. In rare cases the pilot beam may be lower but such exceptions are managed by the MPW, which will be presented later.

When the pilot beam is applied, the BOM allows for an increase of the duty cycle.

The BOM is implemented as a full EPICS IOC, and it uses the SNL sequencer to manage sequence given in Fig 1.

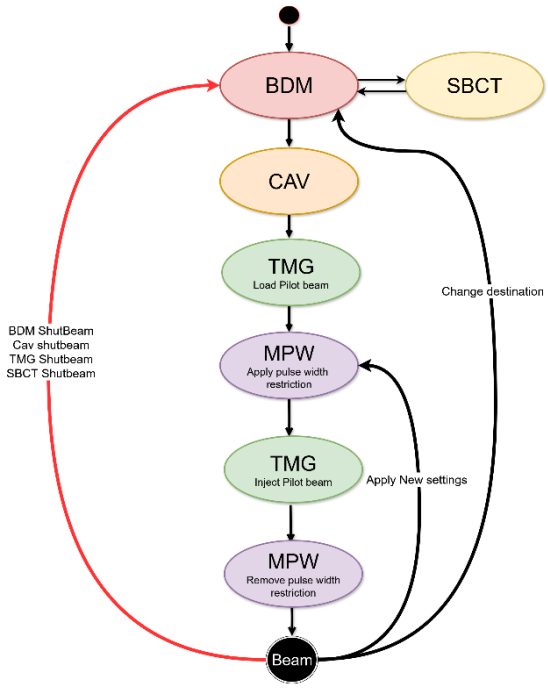



Figure 1: BOM pilot beam sequence.

† alexis.gaget@cea.fr

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD100

# SMART GUIDED COMMISSIONING OF INDUSTRIAL CONTROL SYSTEMS

W. Niinimäki*, Z. A. Morrissey-Ralevic, M. Marin, H. B. Dennis, J. M. de Paco
CERN, Geneva, Switzerland

*Abstract*

This paper proposes a web application to assist in the commissioning of industrial control systems with guided functional tests. The tests verify that trigger conditions produce expected responses within the control system by automatically reading process states from the controllers (e.g. PLCs). The commissioning engineer will receive step-by-step guidance according to the test scripts, real-time feedback, and a final report at the end of the test. The reproducible tests aim to reduce human error and speed up commissioning by automating the tedious task of process state verification.

## COMMISSIONING TESTING OF INDUSTRIAL CONTROL SYSTEMS

Commissioning testing refers to the functional testing of a control system in production before the machine is handed over to the operators. During commissioning testing the commissioning engineers verify the full functionality of the machine.

### *Commissioning Testing Procedure in the Converter Controls Software Section at CERN*

The commissioning testing procedure is dependent on the project at hand, but it at least covers the following points:

- Hardware IO tests.
- Triggering every alarm and confirming their effects and interlocks.
- Traversing the state machine(s) of the industrial control system.
- Confirming the most important conditions for actuations or transitions in the state machine(s).
- Operational mode tests (if applicable).

## AUTOMATED FUNCTIONAL TESTING

The main idea of automated testing is to catch any issues and bugs in software before they make it to the production environment, and it's been the industry standard practice for a long time in IT. Automated tests are scripts used to verify that a software component behaves as expected. These tests vary in their scopes all the way from testing an isolated function to full integration tests between interfaced systems.

In contrast to IT, Operational Technology (OT) has traditionally relied on manual testing methods, such as Factory Acceptance Testing (FAT) and Site Acceptance Testing (SAT). However, in the past few years, automated software testing has started to emerge in OT development cycles. For example, Siemens provides a way to unit test TIA Portal applications with Test Suite add-on, and more recently integrated function tests in the Simatic AX development environment. Other vendors also provide similar tools for their ecosystems. This paper explores the use of functional testing of a full system, including the control system, actuators and sensor readings, most closely aligned with end-to-end testing.

## APPLYING AUTOMATED TESTING TO IMPROVE MANUAL COMMISSIONING

Automated software testing can support and enhance manual commissioning through a method we call *smart guided commissioning*. The goal is not to replace manual commissioning entirely, but to speed it up, reduce errors, and ensure consistency by automating repetitive tasks and enforcing correct test procedures.

In smart guided commissioning, test scripts guide the commissioning engineer step-by-step. They indicate the required actions, automatically verify system responses by monitoring live process values, log state changes with timestamps, and generate test reports automatically. This approach offers several benefits:

- Faster commissioning in production environments.
- Traceable and repeatable test results.
- Reduce the chance for human error by automating the verification of live process values.
- Suite of reproducible tests that can be used for example for regression testing during maintenance of the system to confirm that everything is functioning as expected.

However, preparing high-quality test scripts takes time. Mistakes in these scripts can disrupt commissioning or result in false positives.

## SMARTCO

SmartCo is a web application that was developed to prove the concept of smart guided commissioning. It allows users to create and run tests, as well as store and view test reports. A test in SmartCo connects to all relevant controllers, provides real-time feedback and instructions, verifies process states and generates test reports with time stamps. SmartCo also manages and stores all the test files and reports centrally.

* walter.niinimaki@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD101

# A FLEXIBLE VALIDATION FRAMEWORK FOR STREAMLINING HARDWARE TESTING IN ACCELERATOR CONTROL SYSTEMS

A. Martínez Landete, O. Ø. Andreassen, B. Banas, L. Cordier–Temple, R. Ferraro, I. Kozsar, A. Masi, J. Palluel, J. Tagg, CERN, Geneva, Switzerland

*Abstract*

A new high-performance Sensors, Acquisition and Motion Control system (SAMbuCa) is under development at CERN to address the challenging requirements of motion control for beam intercepting devices such as the collimators of the Large Hadron Collider. These requirements include high precision in extremely radioactive environments, millisecond-level synchronization, and long-term operational reliability. To meet these stringent demands, rigorous testing and validation is essential throughout the development process. The abstraction models used by existing frameworks for executing and defining tests often lack the flexibility required as they are software oriented. To address this, a generic framework for Production Test Suites (PTS) has been developed, helping to validate all the SAMbuCa components before mass production. It is an open-source hardware validation framework written in Python, with the Accelerator Controls community in mind. The framework targets different types of tests, such as end-of-line, reliability and calibration. The PTS framework addresses common challenges faced by accelerator and industrial teams. It provides an adaptable and scalable testing solution relevant across multiple control-system environments. In this paper, the architecture and rationale behind PTS is explained and its functionality compared with existing solutions.

## INTRODUCTION

Commercial off-the-shelf components cannot always meet CERN requirements. The LHC tunnel is a harsh environment, and poses many significant challenges for the systems that are installed in there. Finding products that can receive support for decades, that are radiation tolerant, and that feature appropriate in-field repair possibilities are almost impossible. Moreover, these mission-critical systems need to meet high reliability requirements, as failure often means downtime for the LHC. As a result, CERN often designs custom systems, and every component targeted to be installed in the underground facilities must be tested and validated at different stages of its development and production. Testing is a very broad concept, but there are some relevant references to take into account to better understand the scope of a generic framework for Production Test Suites. The ISO9000:2015 standard defines validation as "confirmation, through the provision of objective evidence, that the requirements for a specific intended use or application have been fulfilled" [1]. In summary, testing allows engineers to ensure compliance to standards and requirements, and can also help determine the necessary time between scheduled interventions for preventive and corrective maintenance. This paper presents the study of existing testing platforms, compares them, and provides the motivation behind developing a new software solution to cover all the requirements for testing hardware in the Accelerator Controls community.

## COST AND RISK OF INADEQUATE TESTING

Every test and validation strategy has a cost, although usually lower than the long-term cost of not having one. Design defect remediation costs increase dramatically throughout the development lifecycle, with fixing a defect at the design stage offering saving factors up to 20,000 compared to fixing it once deployed [2]. In the LHC context, one hour of downtime is estimated to cost approximately 150 kCHF, considering the total operational expenses including maintenance, personnel, and infrastructure costs [3]. Proper testing and validation is therefore a fundamental part of the development for every custom critical device designed at CERN. Although the 'Production Test Suite' name may suggest that it is only targeted for production validation, its flexibility allows using it for validation at every stage of a project.

### *Testing Stages*

These are some of the testing stages through which the accelerator controls systems go, and where the framework can help with automation [4].

Table 1: Summary of Testing Stages

| Testing Stage | Key Activities / Focus |
|---|---|
| Design validation | Validate design functionality against requirements; gather tested samples for statistical analysis. May require sophisticated instrumentation. |
| Prototype validation | Assess correct assembly; benchmark performance of prototypes. |
| Reliability testing | Estimate MTBF and survival rates; eliminate early-failure units using HALT and HASS techniques [2]. |
| Characterization and calibration | Estimate tolerances for production; determine calibration and maintenance intervals. |
| End-of-line testing | Validate each unit after production; may require manual intervention. |

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD102

# CONTINUOUS INTEGRATION MEETS INDUSTRIAL AUTOMATION: TESTING INDUSTRIAL CONTROLS FRAMEWORKS

L. Gutierrez Prendes*, J. Castro Pinto, B. Schofield, J. Ortola Vidal, CERN, Geneva, Switzerland

*Abstract*

Continuous Integration (CI) pipelines are a cornerstone of modern software development, enabling early bug detection and robust validation across all system layers. Extending this concept to industrial automation introduces unique challenges due to the involvement of real hardware and vendor-specific proprietary tools. This paper presents how CERN has built a fully automated testing pipeline for one of its industrial controls frameworks, covering every stage from framework releases (traditionally manual developer tasks), to end-to-end PLC-SCADA validation, simulating user workflows. The pipeline automates the generation of platform-specific PLC projects, the configuration of SCADA systems, and the establishment of the PLC-SCADA communication, achieved through a combination of in-house and third-party tools to programmatically execute these steps. Implementing such CI pipelines requires establishing a permanent testing infrastructure, including PLCs, development machines and SCADA servers. While the setup demands significant initial effort, the maintenance cost remains low, yielding high returns in the form of time saved for developers and testing engineers during the pre-release validation phases. This approach bridges the gap between modern software testing and industrial automation, enhancing reliability in complex, multi-platform environments.

## INTRODUCTION

UNICOS-CPC is a CERN framework to develop control applications for industrial process control. It provides developers with the means to design and develop full control applications, and operators with ways to interact with all items of the process from the most simple to the higher level objects. In addition, UNICOS offers tools to diagnose the process and the control system [1]. Throughout the years, several tools have been produced to automate the instantiation of the objects in the supervision and process control layers and to generate the final Programmable Logic Controller (PLC) programs. The UNICOS-CPC package can be deployed to different platforms. For the control layer, Siemens and Schneider PLCs are supported, while WinCC OA is used as SCADA at the supervision layer.

Due to the nature of the UNICOS-CPC framework, the testing of the software has traditionally been performed manually through the generation of reference PLC-SCADA applications. While the relevant unit tests and integration tests are added with each new feature, the amount of code that can be generated for different device types and configurations makes full manual coverage nearly impossible. Moreover, the codebase has been continuously expanding since the framework was created in 1998, making it increasingly difficult to ensure that new developments do not unintentionally affect other older features, given the complexity of the system.

Manually testing the framework involves generating the PLC program, loading it to the PLC, creating a SCADA application, importing the SCADA configurations, establishing the PLC to SCADA connection and, finally, performing the tests. Considering the effort involved in all these steps, this is usually only done once a year, right before the release. This highly time-consuming process reinforces the need for an automated testing pipeline that automatically executes all the aforementioned steps after every small change in the framework or any of its dependent packages, rather than testing one year's developments at once, as has been done up to now.

## CHALLENGES TO ACHIEVE FULLY AUTOMATED TESTING

Achieving fully automated testing of PLC systems presents different challenges, as it involves real hardware such that applying state-of-the-art software workflows is not trivial. In addition, the software tools for PLC systems are often very specific and not always open source, as they are usually provided by the PLC vendors, who sell these tools as an accessory to their hardware products.

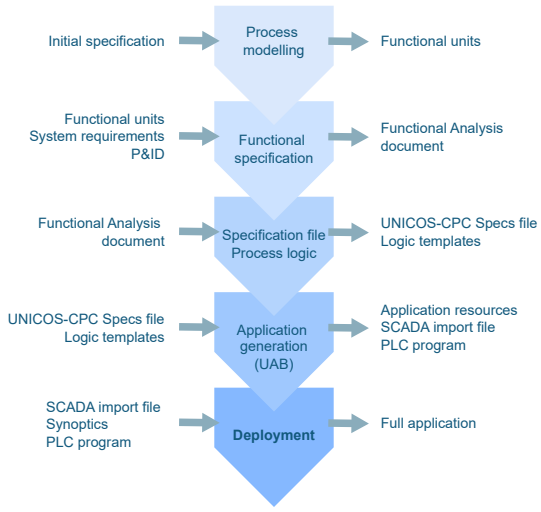



Figure 1: UNICOS CPC application workflow from modelling to deployment.

To successfully automate the testing of these systems, tests must ideally cover all the steps that a user of the framework would take in the creation or upgrade of a UNICOS-CPC application, shown in Fig. 1.

* loreto.gutierrez.prendes@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD103

# FROM TESTING TO VERIFICATION: A COMPREHENSIVE METHODOLOGY FOR INDUSTRIAL CONTROLS FRAMEWORKS

J. Castro Pinto*, L. Gutierrez Prendes, B. Schofield, B. Fernandez Adiego, X. Fink[1], J. Ortola Vidal,
CERN, Geneva, Switzerland,
[1]also at RWTH Aachen University, Aachen, Germany

*Abstract*

Industrial control frameworks frequently rely on object representations to manage complex devices found in the field. However, traditional testing methods often struggle to comprehensively validate internal state transitions of these objects, particularly as their state spaces expand due to increasing complexity and configurability of the objects. This paper introduces a novel testing suite for one of CERN's Industrial Controls Frameworks (UNICOS-CPC), capable of systematically verifying internal state transitions of objects in industrial applications by modeling their state space as a graph. The framework enables developers to define test cases by specifying start states, end states, and commands, while autonomously navigating between states to place objects in the correct initial states of each test case. The test suite also stands out for its extensibility: it decouples test syntax from PLC platform-specific implementations through object-oriented design and OPC UA communication, making it adaptable to diverse industrial control systems and frameworks outside CERN that follow a similar device representation approach. By focusing on individual state transitions rather than complex command sequences, this work simplifies testing of UNICOS-CPC objects and enhances their robustness, allowing for the validation of complex object configurations and behaviours in a transparent manner. Furthermore, this novel testing framework enables an automated workflow to validate scenarios determined by formal verification methods where UNICOS-CPC objects arrive at invalid and undesirable states.

## CONTEXT & PURPOSE

At CERN, industrial controls are widely used for the operation of the particle accelerator complex. The experiments that take place at CERN, and their accelerator infrastructure, are operated by a large set of industrial control applications. In this context, the UNified Industrial Controls System (UNICOS) framework was developed by CERN over the past decades in order to accelerate the production of industrial control systems [1–3].

UNICOS-CPC (Continuous Process Control) is a package of UNICOS [2] which offers a set of generic devices (also referred to as objects) that users can instantiate to build their own industrial control systems. These devices range from low to high level abstractions; the former represent physical I/O signals, and the latter can represent an equipment, for example a diesel generator, or even a whole industrial plant. The devices are hierarchically organized: an object with a higher abstraction is composed of a set of lower-abstraction devices. Furthermore, the framework is capable of generating the code and configurations necessary to prepare the control and supervision layers of a classical control system [4] for operation. Figure 1 illustrates the code generation process.

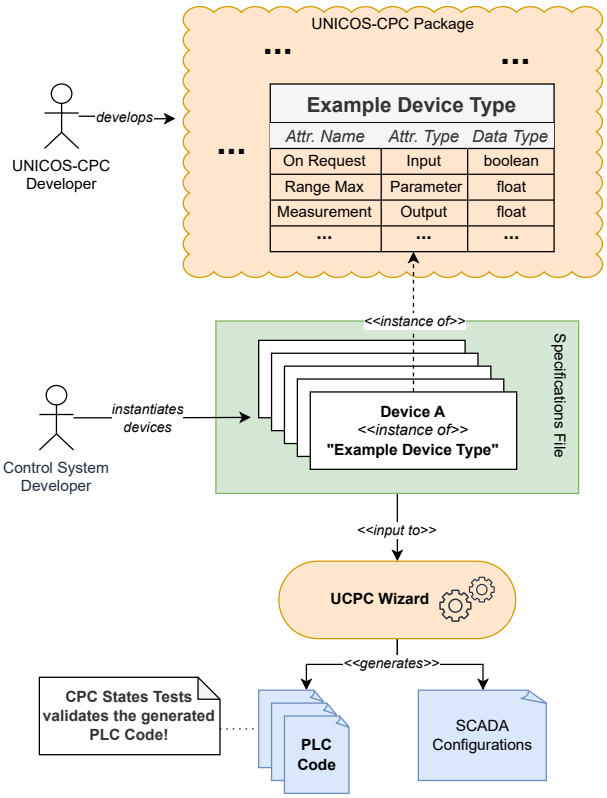



Figure 1: Representation of device instantiation in a UNICOS-CPC environment.

Since UNICOS-CPC is widely used across CERN's particle accelerator complex, it is critical to guarantee the safety and reliability of the internal logic of the framework's devices.

In this context, we propose the CPC States Test suite, a novel manner of testing the internal logic of UNICOS devices. This solution promotes test granularity, maintainability and understandability, proposing a way to theoretically test every transition between two logical states of a device. Furthermore, this concept is extensible to non-UNICOS environments.

* joao.castro.pinto@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD104

# SCALABLE AND STANDARDIZED PLC DEVELOPMENT: AN AUTOMATED FRAMEWORK FOR CERN'S COOLING AND VENTILATION SYSTEMS

R. Figueiras dos Santos*, E. Blanco, J. Cortes, CERN, Geneva, Switzerland

*Abstract*

This paper presents an automated development and testing framework for PLC-based control systems operating CERN's Cooling and Ventilation plants, addressing the challenge of engineering numerous new systems annually while maintaining over 500 existing ones with constrained resources. The framework enhances scalability through standardized, reusable control system components integrated with installation-specific logic to form complete PLC code. These components, stored in a centralized repository, undergo rigorous unit testing and comprehensive validation via GitLab pipelines, utilizing PLC manufacturer-specific Command-Line Interfaces and CERN's proprietary tools to ensure reliability, quality, and regression-free updates. This approach significantly reduces time and resource demands by optimizing development and maintenance processes, providing a scalable model for domains requiring frequent, standardized control system development under high system-to-resource ratios. This work provides insights for engineers seeking efficient, automated solutions for complex control system environments.

## INTRODUCTION

Industrial Control Systems play a crucial role at CERN, forming the backbone of operations that support cutting-edge particle physics research. These systems ensure the safe and efficient operation of complex infrastructure across the entire facility like Cryogenics, Gas distribution or Cooling and Ventilation Systems.

To monitor, control, and automate these processes, CERN relies on thousands of Programmable Logic Controllers (PLCs) distributed across the accelerator complex and experimental facilities. To manage such a vast number of systems and facilitate both development and maintenance, a central Industrial Control Systems team provides the UNified Industrial Control System (UNICOS) framework.

The Cooling and Ventilation (CV) domain is currently the largest user of UNICOS at CERN in terms of the number of PLC-controlled systems. In 2010, CV initiated the migration of its systems into the framework [1]. Since then, new systems have been developed directly in UNICOS at CERN rather than outsourced to external companies. This migration has steadily grown, reaching about 500 PLCs and 60,000 input/output signals by 2025 (see Fig. 1).

Over the past six years, an average of 40 new PLCs have been deployed annually. With a development team of only four members, most of them on short-term contracts, automating the development workflow has become essential. This automation is the only way to sustain such a high development rate while ensuring reliable and maintainable applications.

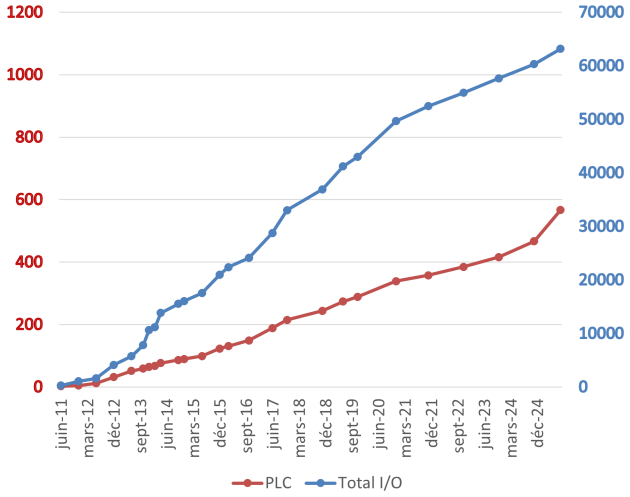



Figure 1: Deployed CV control systems using the UNICOS framework.

The CV development workflow presented in this paper illustrates a scalable approach that can be applied to other domains requiring standardized and efficient engineering practices. It is also directly relevant for future large-scale projects, such as the control systems for CERN's Future Circular Collider (FCC). Furthermore, tools initially developed for CV have already demonstrated their versatility in other projects, including Primary Gas control systems and the AWAKE experiment's vapour source temperature control [2].

## CERN INDUSTRIAL CONTROLS FRAMEWORK

The standard framework adopted at CERN for industrial control systems is UNICOS [3]. Within this framework, the Continuous Process Control (CPC) package provides a set of predefined device types, an engineering methodology, and supporting tools that simplify the development of control applications [4].

A CPC application is developed by defining a configuration file, known as the Specs file, and by implementing parametrizable PLC code templates. These files form the input to the UNICOS Application Builder (UAB), which automatically generates the PLC source code as shown in Fig. 2.

* rafael.j.f.santos@cern.ch

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD105

# AUTOMATED TESTING OF THE INTEGRATED COMPUTER CONTROL SYSTEM AT THE NATIONAL IGNITION FACILITY

D. Gibbard, M. Paul, S. Sauter, J. Castro-Morales, M. Fedorov,
Lawrence Livermore National Laboratory, Livermore, CA, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

This paper describes the Automated Shot Tester (AST), a test automation framework designed to comprehensively test experiments performed using the Nation Ignition Facility's (NIF) Integrated Computer Control System (ICCS). The AST enables automatic testing of diverse experiment configurations on an emulated test system instead of real hardware and eliminates the need for human intervention. While the actual control system is operated by a team of 12, AST acting on their behalf represents significant effort savings while assuring testing fidelity. The AST considerably enhances testing efficiency and expands the range of test configurations compared to the manual method.

The AST is a complete end-to-end framework that manages and monitors the state and condition of ICCS software throughout an experiment. This approach is made possible by leveraging ICCS's distributed architecture and middleware, which enables the AST to receive state updates via the ICCS pub-sub system and trigger commands based on a user-specified configuration file. This file creates modularity and expandability, allowing the AST to exercise a library of test case scenarios, and facilitates the creation of new experiments to be added to integration tests. This testing, along with unit, component, and manual tests, ensure software quality at the NIF. This paper will focus on the design of the AST, the benefits gained from automation and conclude with proposed future enhancements.

## INTRODUCTION

The NIF is the world's largest and most energetic laser experimental system, providing an international center to study inertial confinement fusion and physics of matter at extreme densities and pressures. The NIF is operated by the Integrated Computer Control System, which is a layered architecture of 66,000 devices, over 2,300 front end processors, and embedded controllers. These are all coordinated by higher-level supervisory subsystems in the main control room. In order to run a system the size of NIF, a highly data-driven, modular, automated, and distributed controls architecture was needed. This control system is operated with oversight by 12 operators [1].

A shot automation framework has been developed and deployed to automate shots performed at the NIF. The shot automation framework is designed to automate a 4 to 8 hour sequence of shot activities including automatic creation of shot settings from goals defined in the experiment, setup of the laser and diagnostics per those goals, automatic alignment of laser beams, and a countdown to charge and fire the lasers. On December 5th, 2022 such an experiment achieved fusion ignition at NIF for the first time in human history and has been repeated many times as of time of writing [2].

This software has been successfully deployed to the NIF facility and is constantly maintained and upgraded in order to improve shot rate and system usability. In order to ensure consistent software quality, this shot automation framework is tested daily by AST on multiple experiment types, which will automatically run multiple experiments from end-to-end as if controlled by a team of operators in our emulated offline environments. AST can set up a suitable test environment, load an experiment, progress through the shot cycle, clear and override errors, and log any observed items of note for immediate review. By testing in this manner, NIF is capable driving new software through realistic scenarios under realistic concurrent load [3].

## ICCS STRUCTURE

The shot cycle at ICCS consists of 10 shot lifecycle states, which each represent a specific concept for the overarching experimental goal. For example, Begin Shot does basic configuration checks while Implement Plan prepares for the shot by aligning beams. While some transitions between shot lifecycle states are automatic, some state transitions are manual and must be selected by an operator in order to continue.

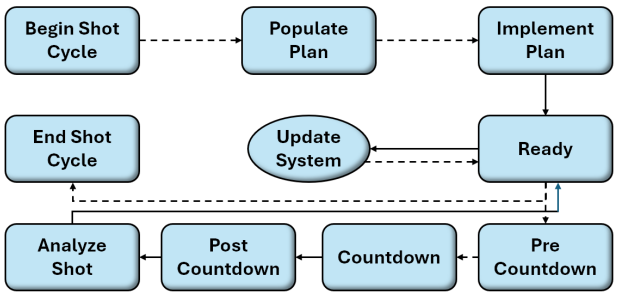



Figure 1: ICCS Shot Lifecycle State Diagram. Dashed arrows represent manual transitions, solid arrows represent automatic transitions.

In order to transition through these shot lifecycle states, ICCS consists of one/two Shot Directors (SD) each communicating with 27 Collaboration Supervisors(CS) which each communicate with 7–10 Shot Supervisors(SS). The shot director advances the shot by commanding its collaboration supervisors to execute a shot lifecycle state, and will only register a state as completed when all collaboration supervisors individually report state completion. These Collaboration Supervisors command their subscribed shot supervisors to each run small, automated checklists known as macrosteps.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD107

# AN OBSERVATION PLANNING TOOL FOR THE MEERKAT RADIO TELESCOPE

Z. Kukuma*, R.L. Schwartz, S.S. Sithole
South African Radio Astronomy Observatory, Cape Town, South Africa

*Abstract*

The South African Radio Astronomy Observatory (SARAO) allocates time on the MeerKAT Radio Telescope to the scientific community to maximize its impact on radio astronomy while fostering South African scientific leadership and human capital development. To streamline and optimize this process, SARAO has developed an Observation Planning Tool (OPT), which allows astronomers to define and plan observations on the telescope. The submitted observations are then processed by an Astronomer on Duty (AOD), before being scheduled.

In this paper, we detail how the OPT supports SARAO's broader mission of effectively operating its radio telescope to produce usable scientific data by enhancing the efficiency, transparency, and scientific utility of the scheduling process. We describe the tool's functionality, design rationale, and ongoing improvements. Key features include a calibrators catalogue; the ability to simulate or dry-run observations; a scheduling assistant to aid the scheduler's optimizations efforts and future schedule plan of observations.

## INTRODUCTION

The MeerKAT Radio Telescope is a project of the South African Radio Astronomy Observatory (SARAO), operating under the National Research Foundation (NRF). Officially inaugurated in 2017, MeerKAT consists of 64 interlinked receptors located in the Karoo region of the Northern Cape, South Africa. Originally conceived as the Karoo Array Telescope (KAT) with only 7 receptors, the project was expanded when the South African government increased funding. This expansion led to the new name "MeerKAT" meaning "more of KAT".

MeerKAT was developed as a precursor to the intergovernmental Square Kilometre Array (SKA), a next-generation facility that will combine thousands of receptors and associated infrastructure across South Africa and Australia. Once complete, the SKA will be the largest and most powerful radio telescope in the Southern Hemisphere, with an effective collecting area of one million square meters.

Since its commissioning, MeerKAT has delivered groundbreaking science and continues to operate at the forefront of radio astronomy. Its achievements include producing the most detailed image yet of the Galactic Centre region. Ultimately, MeerKAT will be integrated into the SKA, extending its legacy as a critical stepping stone toward one of the most ambitious scientific instruments ever built.

Like other astronomy observatories, SARAO makes its telescopes available to interested parties, typically professionals in the field of radio astronomy, through Open Time Calls (OTCs). An OTC is the process SARAO uses to invite and collect radio observation proposals from astronomers worldwide. These proposals undergo a rigorous review and selection process, resulting in a set of approved observations that the MeerKAT telescope will conduct over a designated period.

The Open Time Call process is managed via SARAO's Proposal Workflow System (PAWS) [1], an information system designed to handle proposal submissions and reviews. Once proposals are approved, the planning and scheduling of these observations are carried out using the Observation Planning Tool (OPT), a tool developed by SARAO. The OPT allows astronomers to define, plan, and refine their observation strategies, ensuring that the MeerKAT telescope is used efficiently to meet the scientific objectives of each approved project.

This paper outlines the functionality, design rationale, and continuous improvements of the OPT. Key features include an integrated calibrators catalogue, the ability to simulate or "dry-run" observations, and a scheduling assistant that supports the Scheduler in optimizing observation sequences. Together, these capabilities enable SARAO to better manage valuable telescope time, improve operational planning, and ensure the delivery of high-quality, scientifically valuable data to the global astronomy community.

## BACKGROUND

Before the development of the OPT the process of preparing and executing observations on the MeerKAT Radio Telescope was largely manual and fragmented. Astronomers and operators relied on submitting individual schedule blocks by hand, with supporting tasks spread across multiple disconnected systems. For example, elevation plots had to be generated through custom Jupyter notebooks (see Fig. 1), while scheduling was managed through static documents containing observation times and dates.

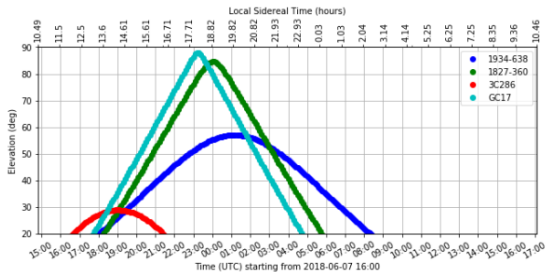



Figure 1: Elevation plot from jupyter notebook.

* zkukuma@sarao.ac.za

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-THPD108

# PROCESS ORCHESTRATION AND SYSTEM CONFIGURATION IN THE MEERKAT RADIO TELESCOPE

R. Ebrahim, L. Williams, South African Radio Astronomy Observatory, Cape Town, South Africa

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The Control and Monitoring (CAM) system of the MeerKAT telescope is highly distributed, necessitating a reliable and automated framework for configuring, deploying, and managing the lifecycle of its many software processes. System configuration follows a hybrid approach using static and dynamic configuration that define the telescope's operational parameters and hardware setups. General functionality can be extended or customized through configuration adjustments, minimizing the need for code modifications. The **katlauncher** application serves as the entry point for initialising the CAM software infrastructure processes for logging and serving configuration information prior to the rest of the CAM system. The **katsyscontroller** application acts as the overall system coordinator, responsible for the sequenced startup and controlled shutdown of the entire system, as well as managing interventions and operator commands. The processes on each node are monitored and managed by the **katnodemanager** application, which exposes sensors and requests to control processes via a Karoo Array Telescope Communication Protocol (KATCP) interface. These components work together to ensure that the MeerKAT telescope operates efficiently and reliably, with clear separation of responsibilities for configuration, process management, and system-wide orchestration.

## INTRODUCTION

The MeerKAT radio telescope is a 64-receptor instrument, with multiple components and subsystems within each receptor, as well as high-level telescope array systems. Effective coordination of these systems is achieved through the MeerKAT CAM system, enabling the telescope to operate as a coherent instrument [1].

The MeerKAT CAM is a distributed control and monitoring system comprising numerous nodes, software components, and interdependent processes, which adds complexity to the challenge of coordination and managing process lifecycles [2]. To address these challenges, the CAM system employs a novel approach to application configuration and process orchestration, utilizing purpose-built software libraries and applications developed specifically for these functions.

## MEERKAT CAM SOFTWARE OVERVIEW

The following section provides an overview of the MeerKAT CAM architecture, which has been presented in detail at ICALEPCS2013 [3] and ICALEPCS2015 [1]. The architecture of the MeerKAT CAM can be divided into three hierarchical layers: CAM Control Components at the top, followed by Proxies, and then Devices at the lowest level [1]. Communication between components in these layers is via KATCP client-server connections. While most communication is done via KATCP, a separate category of components interfaces with CAM's control components, enabling communication using standard web-based technologies such as HTTP Representational State Transfer (REST) APIs and Websocket connections that are consumed by the web frontend **KATGUI**.

### *KATCP*

KATCP is a lightweight, text-based messaging protocol developed by SARAO for communication between hardware devices and their controlling software systems within MeerKAT. Structured around a client-server model over Transmission Control Protocol / Internet Protocol (TCP/IP), KATCP supports three primary message types: requests, replies, and informs. KATCP provides support for both single-client and multi-client configurations, fine-grained sensor sampling strategies, and structured logging, enhancing its adaptability to a range of device interfaces in the telescope system.

### *CAM Software Components*

CAM Software Components consist of various Python libraries and application processes, with the exception being KATGUI which was developed in AngularJS. The foundation of CAM is the **katcp-python** library, which is the Python implementation of the KATCP protocol. The Tornado framework is an asynchronous networking library and web framework that is used as the core engine for asynchronous operations in **katcp-python** [4]. Tornado is also used in the backend web servers for the MeerKAT GUI.

**Devices** Devices are KATCP servers that represent hardware or subsystems, exposing a KATCP interface to the CAM system [1]. In subcontracted systems and third-party hardware without KATCP interfaces, Device Translators are used to translate communication protocols to KATCP so that they can be interfaced with CAM [3].

**Proxies** Communication between higher-level CAM components and other telescope systems or devices is mediated through proxies. Proxies provide a single interface to logically related devices, allowing monitoring, control and logging to be aggregated on a single KATCP interface. Proxies further allow multiple lower level device commands for specific functionality to be abstracted as a higher-level Proxy command. There are multiple instances of proxies for different systems in a functional telescope array.

**CAM Control Components** CAM control components provide telescope system level control and monitoring in-

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-FRAG001

# ML-ASSISTED BEAMLINE OPTIMIZATION AT LCLS*

Z. Lentz†, R. Tang-Kong, S. Miskovich, SLAC National Accelerator Laboratory, Menlo Park, CA, USA
T. Morris, M. Rakitin, Brookhaven National Laboratory, Upton, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

LCLS is currently developing and deploying beamline optimization techniques at our x-ray end-stations. This is an increasingly important topic at LCLS as it fully leverages its new high rep-rate superconducting beam. The increased throughput of LCLS-II shifts the performance bottleneck to on-shift setup time. As part of the Illumine collaboration, LCLS is leveraging Bayesian optimization techniques with on-the-fly machine learning in conjunction with more conventional iterative alignment and digital twin techniques to automatically optimize the beam quality and streamline common elements of experiment setup. This paper will go over how it works, what worked well, challenges faced, and more from a controls perspective.

## INTRODUCTION

In recent years, tools like Xopt [1] and Blop [2] that leverage Bayesian optimization (BO) [3–5] through botorch [6] have become indispensable for optimization of accelerator outputs at LCLS. The relative usability of these tools enables non-experts to implement BO routines for real beamlines without a background in machine learning (ML). This has been applied to the Macromolecular Femtosecond Crystallography (MFX) instrument at the Linac Coherent Light Source (LCLS) in an effort to automate beam alignment.

This work is part of the Illumine [7] project, a cross-lab artificial intelligence (AI) and ML initiative that encompasses a wide range of collaborative projects with the goal of expanding and applying these technologies to improve areas of common interest.

This paper will focus on the application of these tools from a controls and software development perspective, rather than from a scientific perspective. It will step through the learning process of applying Xopt to real-world applications to optimize the alignment of an x-ray beam, including lessons learned.

* This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award Number FWP-101101 for the project "ILLUMINE - Intelligent Learning for Light Source and Neutron Source User Measurements Including Navigation and Experiment Steering".
Use of the Linac Coherent Light Source (LCLS), SLAC National Accelerator Laboratory, is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Contract No. DE-AC02-76SF00515.
This research used resources of the NSLS-II, a US Department of Energy Office of Science User Facility operated for the DOE Office of Science by Brookhaven National Laboratory, under contract no. DE-SC0012704. The involvement of the ARD ML team is supported by the Integrated Scientific and Data Intensive Computing Initiative funding from SLAC.
† zlentz@slac.stanford.edu

## HIGH-LEVEL EXPLANATION OF BO

Even without a rigorous mathematical description, we can understand the practical applications of BO models by examining their capabilities.

These models are designed to make predictions about an N-dimensional system with M optimizable objectives. The M objectives are independently measurable, and the N inputs are both independently controllable and measurable. We can take three critical actions with the model:

1. Submit data about measured inputs and objectives to update (train) the model.
2. Receive suggested inputs from the model that are likely to optimize the objectives.
3. Receive suggested inputs from the model that are likely to improve the model.

Given these actions, we can implement a series of inputs and measurements that efficiently explore a multidimensional parameter space and optimize the objectives. In practice, this might look like moving five motors to various unrelated points and maximizing the scalar output on a diode.

Exploration of such a parameter space can be expensive and this expense scales with the number of points explored. This expense has a time component that arises from a combination of factors such as physical movements, detector readout time, and analysis involved in the optimization. The expense may also have a financial or practical component, such as labor or materials needed for sample preparation. Therefore, exhaustive N-dimensional mesh scans have rapidly scaling costs as we increase the dimensionality of the space. In contrast, a sparse scan that leverages BO will have fewer points and save resources that can be spent elsewhere.

Note that this is not a technique that exists in a vacuum. For example, one can implement hybrid optimizations that use both a BO ML algorithm and a physics-based backpropagation model at the same time.

## POPULAR TOOLS IN THE ACCELERATOR SPACE

Xopt is a Python library developed by the Accelerator Research Division's (ARD) machine learning group at LCLS. Blop is a Python library developed by the Data Acquisition and Detectors (DAD) group of the Data Science and Systems Integration (DSSI) Division at the National Synchrotron Light Source II (NSLS-II). The teams behind these tools collaborate together via the Illumine project and externally in other collaborative efforts. In the long-term, the intent is to make the two libraries interoperable. Both of these tools interface with BO models using botorch in remarkably similar ways. Both are easy to use, both are technically

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-FRAG002

# DEVELOPMENT AND EXPERIMENTAL VALIDATION OF A MACHINE LEARNING-BASED METHODOLOGY FOR CYCLOTRON BEAM CONTROL: RESULTS FROM THE PSI HIPA FACILITY*

M. Haj Tahar†, M. Bocchio, M. Busch, W. Joho, S. Marquie, E. Solodko,
Transmutex SA, Vernier, Switzerland
A. Barchetti, C. Baumgarten, J. Grillenberger, M. Sapinski, M. Schneider, J. Snuverink,
Paul Scherrer Institute, Villigen, Switzerland

*Abstract*

Transmutex SA is developing an accelerator-driven system (ADS) designed to generate clean energy while reducing the lifetime of radioactive waste. Such a subcritical reactor concept requires high reliability and a high degree of accelerator automation to ensure operational effectiveness.

To address these demands, a machine learning (ML) methodology was developed and experimentally validated for automatic beam control in cyclotrons. This work reports the first practical demonstration of machine-learning-based beam control in a high power cyclotron, representing a significant step for this class of accelerators.

The validation experiments were performed on the injector ring of the High Intensity Proton Accelerator (HIPA) at the Paul Scherrer Institute (PSI), whose design closely matches the injector concept developed by Transmutex. Key challenges were addressed, including the identification of suitable observables and actuators, adapting the ML model to the accelerator response dynamics, and integrating ML-based control with existing feedback loops. The approach reliably aligned the beam with the reference trajectory, improving extraction efficiency while minimizing losses.

Over an extensive 12-day operational test campaign, remarkably long in the context of real-time ML experiments, the model demonstrated robust performance across a range of operational scenarios, including varying beam currents and different turn numbers.

These results show that machine learning can enhance operational efficiency, reduce operator workload, and increase automation in cyclotron-driven systems.

## INTRODUCTION AND MOTIVATION

The operation of Accelerator-Driven Systems (ADS) requires accelerators to be tuned with high stability, reproducibility, and efficiency. Traditional manual approaches, although effective in research environments, cannot meet these industrial demands.

Machine learning (ML) offers a promising path forward. By enabling automated, data-driven tuning, ML has the potential to simplify operational control, reduce reliance on expert operators, increase reliability, and improve efficiency of accelerator operation.

At Transmutex, these ideas are being developed within the framework of the START (Subcritical Transmutation by Accelerated Reliable Technology) system and particularly its cyclotron design. Given the architectural similarity between the START cyclotron design and the HIPA accelerator complex at the Paul Scherrer Institute (PSI), the Injector 2 at PSI serves as an ideal experimental platform. PSI granted twelve days of dedicated beam time, offering a unique opportunity to validate ML-based beam control under realistic operating conditions.

## EXPERIMENTAL SETUP

### *HIPA Injector 2 Overview*

Injector 2 is a four-sector cyclotron designed to accelerate protons from an injection energy of 870 keV to a final energy of 72 MeV. Its main components include four large sector magnets (SM1–SM4), two double-gap RF resonators (CI1 and CI3), and two recently installed single-gap resonators (CI2 and CI4). Extraction is performed by the septum magnet (AXA), which deflects the last orbit toward the extraction channel. The beam is then bent by the next dipole (AXB) into the transfer line where final beam current is measured by current transformer (MXC1). Finally the beam is stopped by a dedicated beam dump (BX2). A schematic layout is shown in Fig. 1.

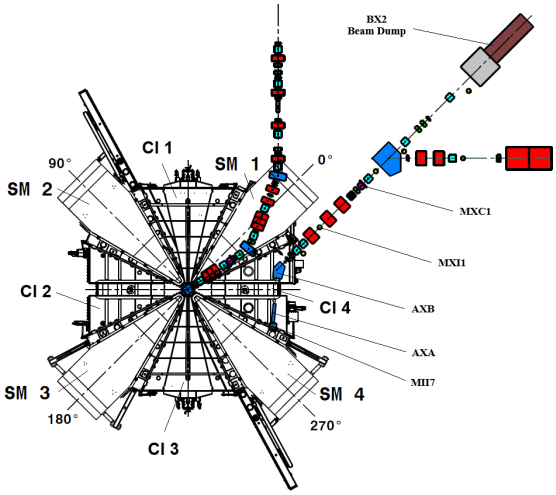



Figure 1: Schematic layout of Injector 2 and its main components.

* Work supported by Paul Scherrer Institute.
† malek@transmutex.com

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-FRAG004

# ADAPTIVE MODEL-BASED REINFORCEMENT LEARNING FOR ORBIT FEEDBACK CONTROL IN NSLS-II STORAGE RING

Z. Dong*, Stony Brook University, Stony Brook, NY, USA
Y. Tian†, Brookhaven National Laboratory, Upton, NY, USA
Y. Sun‡, Sunrise Technology Inc., Stony Brook, NY, USA

Content from this work may be used under the terms of the CC BY 4.0 licence (© 2025). Any distribution of this work must maintain attribution to the author(s), title of the work, publisher, and DOI.

*Abstract*

The National Synchrotron Light Source II (NSLS-II) uses highly stable electron beam to produce high-quality X-ray beams with high brightness and low-emittance synchrotron radiation. The traditional algorithm to stabilize the beam applies singular value decomposition (SVD) on the orbit response matrix to remove noise and extract actions. Supervised learning has been studied on NSLS-II storage ring stabilization and other accelerator facilities recently. Several problems, for example, machine status drifting, environment noise, and non-linear accelerator dynamics, remain unresolved in the SVD-based and supervised learning algorithms. To address these problems, we propose an adaptive training framework based on model-based reinforcement learning. This framework consists of two types of optimizations: trajectory optimization attempts to minimize the expected total reward in a differentiable environment, and online model optimization learns non-linear machine dynamics through the agent-environment interaction. Through online training, this framework tracks the internal status drifting in the electron beam ring. Simulation and real in-facility experiments on NSLS-II reveal that our method stabilizes the beam position and minimizes the alignment error, defined as the root mean square (RMS) error between adjusted beam positions and the reference position, down to ~1 μm.

## INTRODUCTION

NSLS-II is a third-generation storage ring producing synchrotron radiation through laser-electron interactions. Electrons are accelerated through a synchrotron and injected into the storage ring. Low emittance in a light source facility requires stable electron beam orbit [1]. Figure 1 shows a simplified electron orbit that remains in the beam position and emits X-ray radiation at multiple X-ray experiment locations.

The beam operators rely on beam monitoring and control systems to interact with orbits. In each unit of the storage ring, beam position monitors (BPMs) measure the relative position of the beam. Each storage ring also includes corrector control to adjust beam dynamically. Ideally, the corrector currents in orbit controls are initialized at the beginning of experiments, and their initialization depends only on the design of the light source facility. In reality, noise and environmental change cause the beam to gradually drift away from the reference position. Orbit feedback control systems [2] apply corrections to corrector controls and regains the beam reference position.

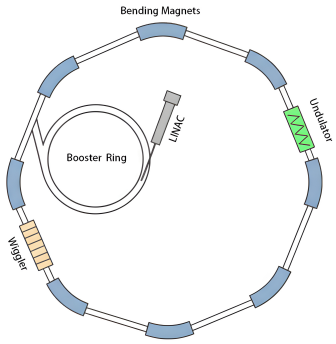



Figure 1: An illustration of a storage ring.

The orbit response matrix (ORM) reflects how BPMs respond to the correction. The matrix is static and belongs to the original machine design. In practice, beam operators periodically measure the ORM by tuning the beam close to the golden beam position and changing the current setting on one corrector one by one. However, we cannot obtain precise measurement on the ORM because of the following reasons:

1. the machine dynamics drifts slightly over time, due to external environment influence, for example, hysteresis of the correctors [3], room temperature;

2. Nonlinear beam; the linear approximation of the ORM measured at the previous time incurs a large bias in modeling the beam system that is non-linear and evolves continuously during facility operation;

3. Environment factors, including vibration and electronic noise, introduce measurement error (noises) in ORM.

Traditionally, SVD-based algorithms produce feedback signals by filtering out high-frequency components in ORMs [4]. Modern NSLS-II design for fast orbit feedback control combines SVD-based feedback control with other controllers, such as PID (proportional–integral–derivative), to ensure robust and stable beam orbits [5]. However, this method has several limitations in practice. Empirically, a large $\lambda$ value leads to robust control to the ORM errors while generating biased RMS values. On the other hand, because of the machine's internal drifting, the controller still has a degenerated performance, even with sufficiently large $\lambda$.

* Zeyu.Dong@stonybrook.edu
† ytian@bnl.gov
‡ yu.sun@sunriseaitech.com