ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TU3I01

COMMISSIONING OF THE LCLS-II MACHINE PROTECTION SYSTEM FOR MHz CW BEAMS*

J. A. Mock†, A. S. Fisher, R. T. Herbst, P. Krejcik, L. Sapozhnikov
SLAC National Accelerator Laboratory, Menlo Park, United States

Abstract

Beam power at the LCLS-II linac and FEL can be as high as several hundered kW with CW beam rates up to 1 MHz. The new MPS has a latency of less than 100 µs to prevent damage when a fault or beam loss is detected. The MPS architecture encompasses the multiple FEL beamlines served by the SC linac and can mitigate a fault in one beamline without impacting the beam rate in a neighboring beamline. The MPS receives inputs from various devices including loss monitors and charge monitors as well as magnet power supplies and BPMs to pre-emptively turn of the beam if a fault condition is detected. Link nodes distributed around the facility gather the input data and stream it back to a central processor that signals other link nodes connected to beam rate control devices. The system design and some experience during initial commissioning are discussed.

INTRODUCTION

The Linear Accelerator Facility at the SLAC National Accelerator Laboratory (SLAC) has seen a significant upgrade to the Linac Coherent Light Source (LCLS) facility with the installation of LCLS-II. This new accelerator and undulator complex was designed to increase capacity at SLAC for photon science using the free electron lasers. A high level diagram of the new facility is shown in Figure 1.

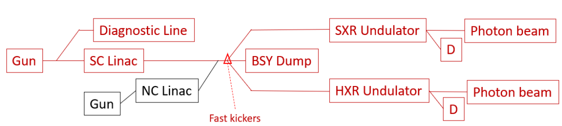

Figure 1: A high level map of the new LCLS-II facility (red), alongside the old NC linac. The new LCLS-II MPS has mitigation devices to inhibit beam or reduce beam rate at either the gun or the fast kickers, to each of the destinations shown.

Because it uses superconducting RF cavities, the LCLS-II accelerator is capable of delivering electrons with a minimum bunch spacing of 1 µs and maximum electron energy of 4 GeV, leading to a maximum overall beam power of 120 kW. An upgrade is also in progress to raise the energy to 8 GeV which will further double this power. The high rate and high power potential of these beams requires safety systems capable of reacting to beam events such that operations remain safe. The personnel protection and beam containment systems (PPS and BCS) are safety systems designed to protect people from radiation-based hazards, and the Machine Protection System (MPS) is designed to protect the accelerator from damaging itself. The MPS is not a credited safety system, so it is designed to be more agile to allow a variety of operating conditions while still protecting the accelerator.

SYSTEM REQUIREMENTS

The scope of the MPS is confined exclusively to shutting off the electron beam when a fault condition occurs that can potentially damage beam line hardware. Other systems that protect high power devices such as power supplies, RF power sources, vacuum systems, or cryogenic systems are handled separately as equipment protection.

Response Time

A key driving parameter of the MPS is the maximum allowable time interval in which the beam must be shut off before damage can occur. The MPS requirement for the original LCLS dictated that the electron beam be shut off within one beam pulse at the full repetition rate of 120 Hz. This is not possible in LCLS-II where the minimum bunch spacing is only 1 µs and propagation delay for a signal in a cable from one end of the accelerator to the other can be as long as 20 µs, not including additional processing delays incurred from electronics. The MPS baseline beam shutoff time, defined as the time between detection of fault and suppression of the electron beam, is required to not exceed 100 µs to avoid catastrophic damage to the beam line, though in principle the MPS physics requirement is as low as reasonable achievable. Not every fault condition requires the fast shutoff time of 100 µs. For example, a slow change in some state, such as a temperature rising, allows ample time for the control system to warn of the impending change. Therefore, MPS responses to prompt events such as beam loss mitigate within the fast response window, and other, slower events and more complicated logic process within a 360 Hz window, equal to the processing time of the legacy LCLS-I MPS.

As shown in Figure 1, the LCLS-II accelerator can deliver to some combination of an injector diagnostic line, one of two undulator beamlines, or through the linac to a high power dump in the SLAC Beam Switchyard (BSY). The default destination for the electron beam is this high power dump. Pulsed kicker magnets are used to kick the beam into one of the other beamlines, which means if the pulser does not kick, the beam will need to travel to the BSY beam dump. Therefore, the MPS uses this high power dump as its primary mitigation device for downstream faults. It grants

* SLAC is supported by the U.S. Department of Energy, Office of Science, under contract DE-AC02-76SF00515
† jmock@slac.stanford.edu

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TU3I01

and revokes permits to the various kicker magnets in order to control beam power or completely inhibit beam to the other destinations. Additionally, the MPS can inhibit the beam within the entire facility by mitigating at the electron gun (through an interaction with the drive laser of the phot-injector). The mitigation device for each destination can be treated independently and all can operate at the response times listed above.

The LCLS-II accelerator can operate side-by-side with the old normal conducting (NC) linac which has a separate MPS operating independently, so that faults in one system do not interact with the other beam.

During the LCLS-II commissioning phase, overall beam power will be limited. The maximum repetition rate through the superconducting linac is 93 kHz, and the maximum repetition rate through the undulator complex is 1 kHz. This limitation allows the MPS to grow into its role as the system comes online and is better understood on the scale of the facility. During this phase, the MPS must shut off the beam in the linac within the 100 μs window at the gun, but in the undulator complex the system can respond as slowly as 1 ms and still inhibit the beam before the next bunch would enter those beamlines.

Inputs

The MPS takes inputs from many types of beamline devices which it uses to determine the state of the accelerator. Types of inputs are:

- Obstructions. Objects like vacuum valves or diagnostic screens can impede the path of the electrons. If these objects are struck, they can be damaged. Depending on the type, an obstruction will cause the beam to be rate limited or shut off. Obstructions all have at least one physical limit switch associated with the OUT position that is an input to the MPS. If the obstruction is not in the OUT position, the beam is shut off to preemptively avoid damage to the component. Some obstructions, such as profile monitor screens, allow beam when they are IN, and those devices have IN switches in MPS as well. When the device gets to its IN position, the MPS fault clears and beam can be allowed again, typically at a lower repetition rate with lower power.
- Beam Loss Monitors [1–3]. MPS brings in beam loss measurements from three types of beam loss monitors. First, a series of optical fibers run the length of the accelerator, divided into approximately 200 m segments. These fibers interact with beam loss to produce light which is detected and integrated into a measure of the beam loss within a zone. Second, point beam loss monitors are placed at strategic locations such as collimators to measure beam loss from that point. These two types of beam loss monitor are shared between the BCS and MPS systems. Finally, MPS has a fast beam loss monitor installed at every undulator. These are detectors that detect Cherenkov light produced by the interaction of beam loss with the material. Each of these beam loss inputs has an adjustable threshold, beyond which the beam will be tripped.
- Destination Dipole and Chicane Bend Magnets. The MPS ensures the beam makes it to one of the beam dumps. Independent MPS magnet current monitors are used to ensure certain bend magnets are at appropriate levels so the beam can safely traverse the beam pipe. The MPS magnet current tolerance bands are set automatically and derived from the expected beam energy.
- Bunch Charge Monitors (BCMs). Toroid detectors are placed at strategic locations in the accelerator complex to measure beam charge to the various destinations. The MPS uses these data to determine that the beam is being delivered to the proper destination. Additionally, the MPS verifies that the maximum beam power is not being violated nor is charge being lost along the linac.
- Beam Position Monitors (BPMs). The measured X and Y values from each BPM are used by MPS to verify the electron orbit is within an acceptable envelope. This is used to preemptively prevent beam loss in the event of large orbit excursions. Additionally, the measured charge at each BPM can be used as a secondary measure of the beam charge at any location in the same way as a beam charge monitor. The timing pattern for triggering the BPM acquisition carries meta data with the expected bunch charge on each pulse allowing a fast comparator to trip the beam if the measured charge is not within a specified tolerance of the expected charge [4].
- General Digital Inputs. The MPS is designed to read in any signal that can be turned into a present / absent electrical signal. This allows maximum flexibility for what can be interlocked with the system. As an aside, faults from the photon systems (MPS for the photon beamlines) are handled as generic digital inputs. Due to the changing nature of photon systems needed to accommodate many different users, the photon groups maintain a more flexible MPS system and uses digital inputs to send data to the electron systems to mitigate rate or inhibit the photon beam. The photon beam transport lines do not have any fat shut-off means, so it is necessary to suppress the electron beam with the MPS to protect sensitive photon systems needing a fast reponse.

SYSTEM DESIGN

The MPS is comprised of a collection of distributed devices known colloquially as Link Nodes that collect and process data and send them off to a central processor referred to as the central node. The central node collects these data and compares them against a pre-programmed logic table to determine the overall state of the accelerator. The output of this calculation is then distributed to other link nodes that grant or revoke permits to mitigation devices to allow beam to certain destinations. The link nodes and cen-

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TU3I01

tral nodes process their data and send out new messages with each 1 µs clock cycle from the LCLS-II timing system. The link nodes connect to the central node via a dedicated, low-latency communication network with a bandwidth of 5 Gbps. The central node is capable of receiving and processing data from 12 of these links. The MPS network is unidirectional - the link nodes send messages when the timing system trigger is present. To conserve bandwidth and avoid many km-long fiber runs, the link nodes are arranged into a daisy chain whereby one link node feeds its information to the next until no more than half of the available bandwidth is consumed. A diagram of the MPS network is shown in Figure 2.

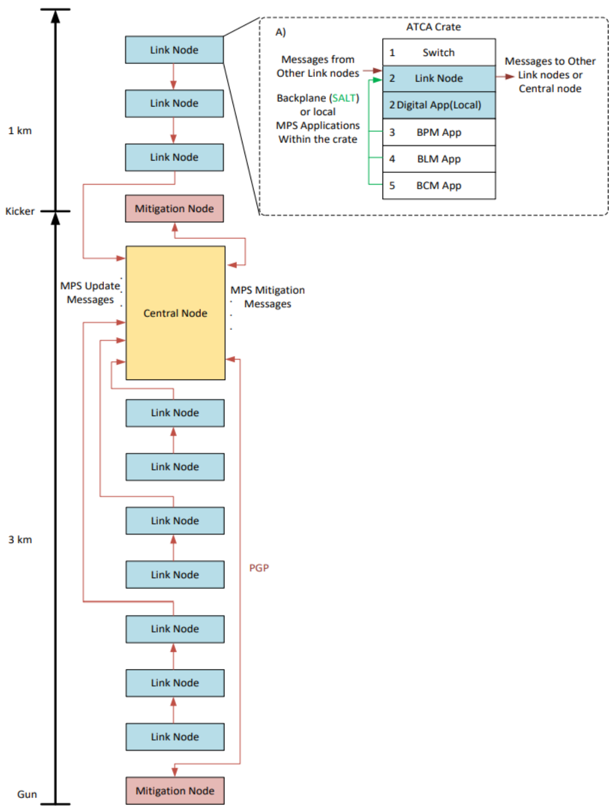

Figure 2: The low-latency MPS communication network.

Because the MPS operates at the full 1 MHz repetition rate of LCLS-II, the core functionality of the link nodes and central node is achieved within an FPGA. The MPS is built upon the SLAC Common Platform, a generalized facility for beam instrumentation based on the ATCA shelf format [5]. The various beamline diagnostics such as BPMs, BCMs, and BLMs also use this platform. The insert in Figure 2 shows a typical ATCA crate installation at SLAC. Each crate includes a network switch in slot 1, a link node in slot 2, and application payloads in the other slots. A computer with local 10 Gbps network is connected to the crate switch to provide a mechanism for asynchronous data transfer from the FPGAs to the control system [6]. The application payloads all contain an MPS firmware block which handles applying thresholds to analog data, converting the analog signals into digital bits to be sent to the central node. The full complement of logic is distributed across the entire system to preserve bandwidth to ensure future expand-ability. The LCLS-II project saw the installation of approximately 100 link nodes (ATCA crates) and 500 discrete applications (slots within the crate).

Central Node

The central node is built on the same SLAC common platform ATCA-based architecture as the link nodes and other accelerator diagnostic systems. The high-level central node architecture is shown in Figure 3. The central node

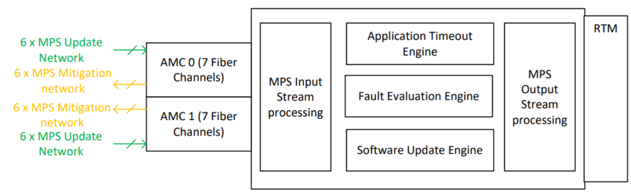

Figure 3: The high-level central node architecture.

is responsible for receiving data from 12 input streams and process them to determine the overall permit state of the accelerator. After the input stream processing, three different fault processing engines are run.

- Application Timeout Engine. This engine is responsible for ensuring that every application in the system sends a unique message on each clock cycle. The allowed latency of this engine is 7 µs, and if it takes longer than that to receive a unique message, the application timeout engine will issue a timeout fault and revoke the permits from all destinations since it does not know which destinations are possibly affected by the missing data.
- Fault Evaluation Engine. This engine is responsible for evaluating the fast faults against the pre-programmed firmware rules tables. Beam loss monitors, BPMs, BCMs, and in principle any fault with a single input (OK / Not OK) is evaluated as a fast fault. The input data are evaluated against the fast fault rules table by this engine every accelerator clock cycle.
- Software Update Engine. Faults that have inputs that change slowly (for example temperature) or more complicated logic tables (for example, a screen with an IN and OUT position) are evaluated as slow faults by a software algorithm. The central node streams synchronous input data via the software update engine from the FPGA to a process running on the attached computer at the slow 360 Hz update rate. The software algorithm receives these input data and evaluates the more complicated logic tables to compute a software permit set. This permit set is sent back to the firmware to be included in the final permit calculation.
- Timing System Verification. The MPS uses the timing system to impose rate limits, and it must verify the timing system is doing as asked. The timing pattern is read in by the central node so the beam rate and charge

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TU3I01

expected at each destination can be computed to verify all is as expected. If information in the timing stream is found to be out of compliance with what the MPS has dictated, the MPS will revoke all permits from all destinations.

The permits computed by these engines are all combined to form the final permit state for each destination. The final permit per destination is defined as the lower allowed state from each of these. These permits are broadcast back out to certain link nodes which are also connected to mitigation devices as well as the timing system. The shutoff mechanisms are discussed in the next section.

Shutoff Mechanism

The MPS has two mitigation paths. It can inhibit the electron beam completely with its own, independent mechanisms or it can impose a power (rate) limit by requesting the timing system reduce the beam rate to a particular destination. The MPS treats the timing system as an untrusted partner - the MPS requests a rate from the timing system and then verifies the timing system has delivered. This rate reduction mechanism allows the MPS to be proactive in some situations. For example, all Beam Position monitors (BPMs) are input devices to the MPS. The BPMs measure the position and bunch charge of the beam within the beam pipe and feed this data to the MPS at the full 1 MHz rate. If the MPS detects that the beam orbit (trajectory) or bunch charge is deviating outside its pre-determined tolerance band, the MPS can first impose a rate reduction so the accelerator operators can correct the deviation and resume high power operation. If the orbit deviation is not corrected or gets worse, the MPS will completely inhibit the beam before damage were to occur. The same mechanism can be applied for beam loss - power can be reduced when beam loss is detected, allowing operators to address the issue before the complete shutoff of the beam.

Each beamline destination shown in Figure 1 has a mitigation device that receives permits from the MPS via one of the link nodes. Pulsed kicker magnets direct beam from the main linac line into the hard X-ray and soft X-ray beamlines as well as the injector diagnostic beamline. The pulsed power supplies are triggered by the LCLS-II timing system, so rate MPS rate reduction is achieved by adjusting the trigger patterns driving these magnets. In the event that rate reduction is not good enough, an operational permit from MPS to the pulsed kicker power supply is revoked, effectively shutting off the magnet. The main linac mitigation is governed by the laser at the photo-injector. Rate reduction is achieved by adjusting the laser triggering timing pattern, and full shut off of the beam is achieved when the MPS revokes a permit and shuts off the laser completely. In this way, the MPS uses the timing system to mitigate power but does not rely on the timing system if the beam must be completely shut off. Each mitigation device has a direct permit from the MPS that can be revoked if needed.

Central Node Concentrator

In principle, the MPS can function as a distributed set of link nodes sending data to a single central node located near the center of the linear accelerator facility which can then issue permits based on one overarching calculation. The full extent of the SLAC linear accelerator facility is longer that 4 km, and the fast beam shutoff time would be dominated by signal propagation if the MPS were constructed with this scheme. To reduce the propagation latency, the central node duties were split from one node into three. This design is shown in Figure 4. Each central node is independent and

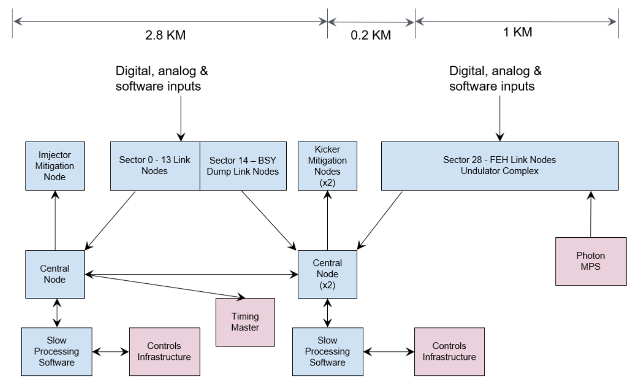

Figure 4: The overall design of the LCLS-II MPS system.

contains its own set of inputs and rules tables. After computing the final permits from the fast and slow fault engines, the final permit states are exchanged between all central nodes, keeping them synchronized with each other. Two central nodes - one near the gun and one near the undulator kicker magnets - handle the superconducting linac and the beam transport line through the high-power dump in the beam switchyard. The third central node, also located near the undulator kicker magnets, handles inputs pertaining to the undulator complex including the undulator transport line between the kicker magnets and undulators.

Software

The software supporting the MPS central node and link node is built upon the EPICS toolkit. An EPICS IOC connects to and provides synchronous and asynchronous control and readout of the firmware and software algorithms contained within each node. EPICS channel access and PV access are used to exchange status data from each node to the control system for operations. The high level software architecture is shown in Figure 5. The central node publishes its fault history to a MongoDB database which can be queried by control room tools so operations can reconstruct faults to determine why beam may have tripped off. The timing system operates Beam Synchronous Acquisition (BSA) data buffers for every fast diagnostic on the machine. When an MPS fault occurs the circular BSA buffers are frozen and contain the last 60 seconds of data from the accelerator before the trip. This allows operators to delve into the root

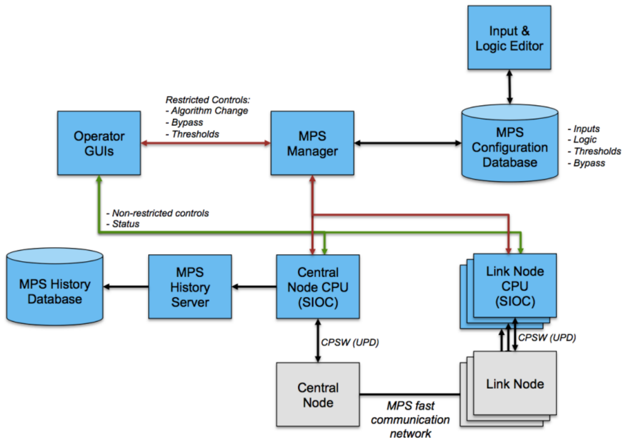

Figure 5: The MPS software architecture.

cause of difficult-to-diagnose machine trips by correlating orbits, energy, losses etc. leading up to the trip.

The overall MPS configuration is stored in a sqlite database. Python scripts are used to insert and extract data into this database. The extraction process produces all necessary EPICS database support as well as the logic tables that are eventually published into the central node firmware rules. Additionally, input and logic reports are generated as output from this configuration database. In this way, the configuration database is a single source of truth for the MPS configuration.

The control room tools provided include software diagnostic screens that show status information about each link node and the central node. Additional screens show status of inter-link node communication status, network bandwidth, and message latency. Information about input status are provided at the link node and central node levels. All of these data are provided as individual EPICS variables and stored in the EPICS archive appliance. The output of the logic tables and permits are shown in a GUI with searchable interface and ability to query the history as well as live faults. To minimize interruptions to the software stack in the central node that computes the slow logic, fewer EPICS variable pertaining to the output of the logic are published in lieu of the logic GUI querying the configuration database with a single PV per fault to reconstruct the status.

OPERATIONAL EXPERIENCE

The MPS was installed and checked out over a multi-year period and was completed before the LCLS-II commissioning activities began. Thus, it was a tool in use for commissioning activities. Before major activities began to each beamline, the fast MPS shutoff times were measured. These are shown in Table 1. The times shown represent a measurement of the time from detection of fault to the loss of the laser or kicker pulse. The kicker shutoff time is slower than the laser shutoff time, and has not yet reached the 100 µs fast shutoff requirement. The measured maximum time of 110 µs is still faster than the beam minimum beam spacing of 1 ms in use for commissioning. Further investigation will be undertaken to understand and correct this slower-than-desirable signal before high power operation commences.

Table 1: Measured Shutoff Times

Destination (device)	Commissioning required shutoff	Measured shutoff
Linac (laser)	100 µs	20 µs
Diagnostic (kicker)	10 ms	1 ms
Hard X-ray (kicker)	1 ms	70 µs
Soft X-ray (kicker)	1 ms	110 µs

The superconducting linac commissioning demonstrated 93 kHz beam rate operation. During this demonstration, the MPS relied on BPM orbit interlocks and BLM interlocks to ensure all the beam reached the high power dump in the beam switch yard. The MPS did its job to stay out of the way of operations unless intervention was necessary at which point the control room tools provide information about what the problems are.

The undulator complex commissioning utilized the undulator beam loss monitors to verify clean trajectory through the undulators. Electron bunches were sent one at a time through the undulator, with analysis of each bunch trajectory happening before the next bunch was sent. On the soft x ray undulator beamline, the MPS beam loss monitors showed that the first shot attempted was lost at the first girder. After trajectory tuning, the undulator beam loss monitors showed the second electron bunch traversed cleanly through the undulators with minimal beam loss measured.

CONCLUSIONS

The LCLS-II MPS has been designed to be as logically simple as possible to maintain robust and stable operation while still providing a flexible platform that can conform to the needs of the facility. The system has been fully installed and commissioned. It is currently in use to support commissioning activities and provides protection of beamline components from damage by the accelerator itself. As the facility advances and accelerator power levels increase, the system will evolve to meet these needs as necessary.

REFERENCES

[1] A. S. Fisher, R. C. Field, and L. Y. Nicolas, “Evaluating Beam-Loss Detectors for LCLS-2”, in *Proc. IBIC’16*, Barcelona, Spain, Sep. 2016, pp. 678–681. doi:10.18429/JACoW-IBIC2016-WEPG23

[2] A. S. Fisher *et al.*, “Commissioning Beam-Loss Monitors for the Superconducting Upgrade to LCLS”, in *Proc. IBIC’22*,

ISBN: 978-3-95450-236-3 ISSN: 2673-5350 doi:10.18429/JACoW-IBIC2023-TU3I01

Kraków, Poland, Sep. 2022, pp. 207–210. doi:10.18429/JACoW-IBIC2022-TU2C3

[3] A.S. Fisher *et al.*, “Commissioning the Beam-Loss Monitoring System of the LCLS Superconducting Linac”, in *Proc. IBIC’23*, Saskatoon, Canada, Sep. 2023, this conference.

[4] C. Bianchini Mattison *et al.*, "LCLS-II Timing System and synchronous bunch data acquisition”, in *Proc. IBIC’23*, Saskatoon, Canada, Sep. 2023, this conference.

[5] J.C. Frisch *et al.*, “A FPGA Based Common Platform for LCLS2 Beam Diagnostics and Controls”, in *Proc. IBIC’16*, Barcelona, Spain, Sep. 2016, pp. 651–654. doi:10.18429/JACoW-IBIC2016-WEPG15

[6] J.A. Vásquez, J.M. D’Ewart, K.H. Kim, T. Straumann, and E. Williams, “YCPSWASYN: EPICS Driver for FPGA Register Access and Asynchronous Messaging”, in *Proc. ICALEPCS’17*, Barcelona, Spain, Oct. 2017, pp. 1707–1709. doi:10.18429/JACoW-ICALEPCS2017-THPHA138