DEVELOPMENT OF THE MACHINE PROTECTION SYSTEM FOR LCLS-I*

J. Dusatko#, M. Boyes, P. Krejcik, S. Norum[1], J. Olsen
SLAC National Accelerator Laboratory, Menlo Park, CA 94025, U.S.A.

Abstract

Machine Protection System (MPS) requirements for the currently operating Linac Coherent Light Source (LCLS-I) demand that fault detection and mitigation occur within one machine pulse (1/120th of a second at full beam rate). The MPS must handle inputs from a variety of sources including loss monitors as well as standard state-type inputs. These sensors exist at various places across the full 2.5Km length of the accelerator and beam lines. A new MPS has been developed based on a distributed star network where custom-designed local hardware nodes handle sensor inputs and mitigation outputs for localized regions of the LCLS accelerator complex. These Link-Nodes report status information and receive action commands from a centralized processor running the MPS algorithm over a private network. The individual Link-Node is a 3u chassis with configurable hardware components that can be setup with digital and analog inputs and outputs, depending upon the sensor and actuator requirements. Features include a custom MPS digital input/output subsystem, a private Ethernet interface, an embedded processor, a custom MPS engine implemented in an FPGA and an Industry Pack (IP) bus interface, allowing COTS and custom analog/digital I/O modules to be utilized for MPS functions. These features, while capable of handling standard MPS state-type inputs and outputs, allow other systems like beam loss monitors to be completely integrated within it. To date, four different types of Link-Nodes are in use in LCLS-I. This paper describes the design, construction and implementation of the LCLS MPS with a focus on the Link-Node, which has proven to be a very useful and flexible component for the MPS.

HISTORICAL PERSPECTIVE

The SLAC National Accelerator Laboratory accelerator complex has hosted several different Machine Protection Systems throughout its varied history. The original 1960s era SLAC Linac contained two protection networks [1]. The primary one, capable of responding within 1ms, consisted of a transmitter broadcasting two tones through a network of series connected Tone Interrupt Units (TIUs) to a receiver unit, which gated beam production on/off at the injector. The TIUs used mechanical relays actuated by various system sensors, so that if a sensor tripped, the tone signal path would get interrupted, inhibiting triggers and indicating a fault alarm.

The SLAC Linear Collider (SLC) era in the 1980s ushered in a more sophisticated MPS [2] where programmable logic algorithms and more sophisticated sensors allowed automatic limiting and control of beam repetition rates without switching off the entire accelerator. This system was based on both CAMAC and VME bus platforms and used the MIL-STD-1553 serial bus standard as a fast, dedicated link between the sensor units and fault processor units. This became known as the "1553 MPS". It interfaced with the accelerator timing system, dictating beam rates to it. The 1553 MPS was largely independent and parallel to the TIU MPS.

The development of LCLS-I presented a unique opportunity and challenge to replace both of these legacy systems by combining and integrating their functionality. The new LCLS-I MPS had to connect to legacy accelerator systems as well as to completely new installations built solely for LCLS-I. The development effort also provided the opportunity to use modern technology.

THE LCLS-I MPS

SLAC is building a second facility, LCLS-II, so we now refer to the present machine as LCLS-I. The LCLS-I Free Electron Laser is a pulsed machine with a maximum repetition rate of 120Hz. LCLS-I uses the last 1/3rd of the SLAC LINAC, the existing switchyard beamline and then extends this to a newly constructed undulator hall e-beam dump, photon transfer lines and experimenter areas. An overview of the LCLS-I MPS is given in [3]. The system requirements for the LCLS MPS [4], state that the MPS will only turn off the electron beam and not any other devices. In addition to turning off the beam, this system also needed to be capable of rate limiting beam production upon certain conditions. System inputs include standard (mechanical and electronic switch state) sensors (e.g. valve position) as well as direct input from beam loss detectors. Presently the input device count is approximately 2100.

The LCLS MPS has outputs to control beam mitigation and shutoff actuators. These devices are mechanical shutters at the photoinjector laser and laser heater, as well as the gun timing trigger permit. Another mitigation device is a fast pulsed kicker magnet located just prior to the undulator magnet string.

The minimum required response time (sensor in to mitigation out) is 8.3mS. As implemented, the system actually samples, processes and acts on sensors at 2.78ms.

The architecture is that of a star network with a central MPS Link Processor connecting to individual Link-Node units. The interconnect channel is a private Gigabit Ethernet network which uses commodity Ethernet switches (Cisco Catalyst 3750). The Link Processor (LP)

*Work supported by U.S. Department of Energy under Contract Nos. DE-AC02-06CH11357 and DE-AC02-76SF00515.
[1]Now at Varian Medical Systems, Palo Alto, CA
#jedu@slac.stanford.edu

Copyright © 2011 by the respective authors — cc Creative Commons Attribution 3.0 (CC BY 3.0)

implements the MPS algorithm and is responsible for all decision making based on the states of device inputs it receives from each of the Link-Nodes (LN). The LNs provide the I/O interface between the MPS and the outside world; with a LN connecting to sensors (input devices) and mitigation units (output devices). The LP (a Motorola MVME-6100 CPU) also communicates rate limiting commands to the LCLS Timing System [5] via a point-to-point Ethernet link. It receives timing pattern and fiducial information from the timing system by way of a connected timing Event Receiver (EVR) module. The EVR (MicroResearch Finland PMC-EVR-200) generates an interrupt to the LP at the timing system 360Hz fiducial rate. The LP also connects to the regular control system Ethernet network (via a second Ethernet port) for status and control communication.

Communication between the LP and LNs over the private Gigabit Ethernet network uses the UDP protocol. The upper layers of the Ethernet protocol were not used in order to keep the overhead low. The communications sequence is as follows: 1) The LP is woken up by the 360Hz interrupt generated by the timing fiducial received by the EVR. 2) The LP sends a sync broadcast message to all LNs, asking them to return their status. 3) The LNs then send their device state status information to the LP. 4) The LP sends a permit command message telling specific mitigation devices to actuate based on the previous (one 2.78ms tick) sensor status. 5) The LP processes the new status and sends out a status clear message back to the LNs, telling them to clear any latched faults. 6) The LP, based on its algorithm, processes the current sensor states and decides if mitigation devices need to fire on the next fiducial. 7) The LP then waits for the next interrupt. Software and FPGA watchdogs ensure a fail-safe communications link. A maximum of 254 LNs can be connected to the LP. The LN IP address is set with DIP switches. Packet size is modest (< 1000 bytes) and the data payload is small enough to handle communications at 360Hz without dropped or split packets. Maximum link communication times were calculated and later empirically confirmed (< 1ms). The LNs also connect (via its embedded CPU) directly to the control system network for control, status and housekeeping.

THE LINK-NODE

The Link-Node is the central functional element of the LCLS MPS and its unique design merits further elucidation. In order to accommodate a variety of input signals distributed across the full accelerator complex, a local signal collection and processing point was required. It was desirable to develop a solution that could accommodate a variety of inputs, which led to the idea of a modular platform. Several crate-based platforms were considered including VMEbus and micro-TCA. The chassis design ultimately won out based on cost and the need for customizability. This gave a common platform that could be configured in multiple ways with low overhead. A block diagram of the Link-Node chassis is given in Fig. 1. The chassis consists of several circuit boards, and depending upon the flavour of LN (Mitigation, BLM, PIC, ByKIK) the configuration and types of these boards varies. The general functions of each of these boards will now be described

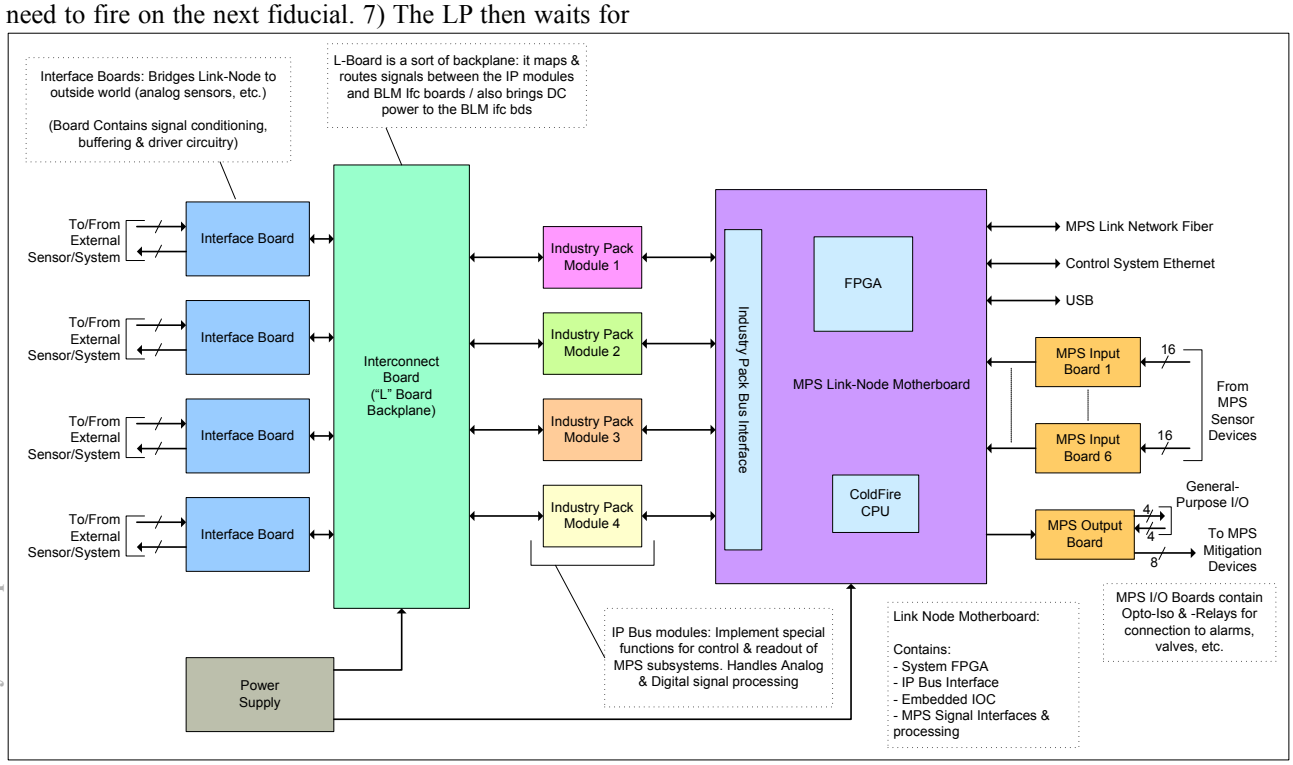

Figure 1: MPS Link-Node Block Diagram.

Copyright © 2011 by the respective authors — cc Creative Commons Attribution 3.0 (CC BY 3.0)

Motherboard

The Motherboard (MB) contains the main functions of the LN chassis. It contains a Xilinx XC4VFX20 Virtex-4 FPGA, a ColdFire embedded CPU (Arcturus Networks uC5282), the MPS digital I/O interface, a USB 1.0 interface and the Industry Pack bus interface. The USB interface permits local control and status readout for development and maintenance purposes. The FPGA implements the MPS engine, the private Ethernet interface (using its embedded MAC and Rocket I/O cores), the USB interface and the IP bus interface. The MPS sensors and mitigation devices connect to the motherboard via the MPS Digital Input and MPS Digital Output boards, respectively.

Input and Output Boards

The MPS Digital Input Board accommodates 16 optically-coupled inputs. It is capable of supplying up to 2A of fuse-protected standard +24VDC for interface logic. A total of six Input Boards can be installed into the MB, giving the L-N chassis up to 96 digital inputs.

The MPS Digital Output Board contains eight optical relay outputs for driving mitigation devices. It is also provides up to 1A of fused +24VDC for interface logic. The MB can accommodate only one Output Board. This board also contains four each, general-purpose digital inputs and outputs; available on Lemo connectors. The inputs are 50Ω terminated TTL compatible schmitt-trigger types, while the outputs are high current TTL drivers, capable of driving up to 67mA into 50Ω. These I/Os are configurable for different functions inside the MB FPGA and are used for dedicated I/O in the BLM, PIC and ByKIK L-N models.

Industry Pack Interface and Modules

A standard Industry Pack (IP) Bus interface is available on the MB that can accommodate up to four IP modules. This interface allows the use of custom and COTS IP modules to enable a signal processing, status and control path between sensor elements and the MPS. With this capability we can completely integrate a sensor readout subsystem within a LN chassis. This enables the MPS to directly connect to the beam loss and pulsed kicker magnet systems. The IP bus side of the IP modules connect to LN MB, while the I/O side of the IP modules connect to the L-board. The IP modules receive their power from the MB. The type of IP board used depends on the flavour of LN, and can be analog and/or digital I/O.

L-Board and Interface Board

The L-Board, so named for its physical shape, is essentially a backplane that provides interconnect and power between the IP modules and the interface boards. There is one L-Board per chassis. It contains mostly connectors for power, the IP and Interface boards and I/O of key IP module signals. To date, two L-boards have been designed: the BLM and PIC L-boards.

The Interface Board contains signal conditioning, receiver, driver and protection circuitry that bridges the LN to external devices. It receives its power from the L-board. The Interface Board is application-specific and to date, two Interface Boards have been designed: the BLM and PIC Interface Boards.

Motherboard FPGA

The MB FPGA implements the core functions of the LN. A block diagram of the FPGA is given in Fig. 2.

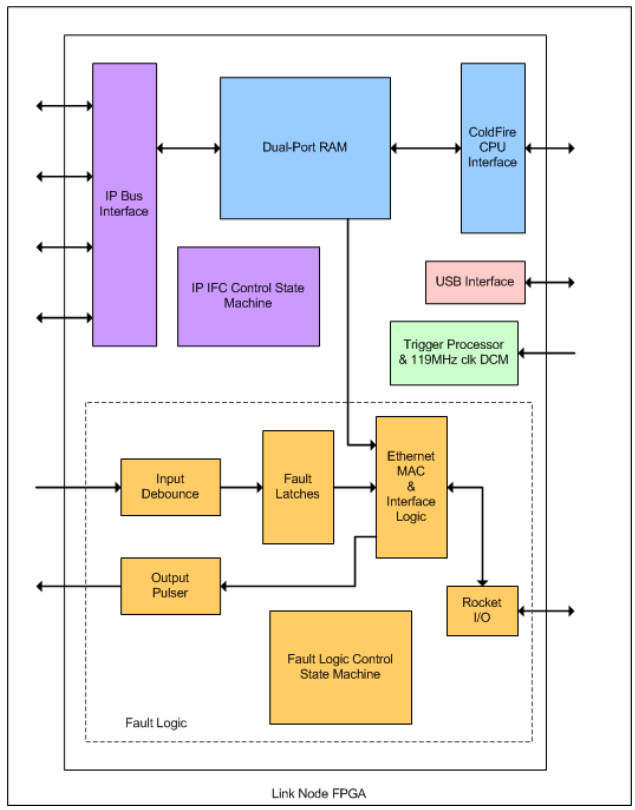

Figure 2: LN Motherboard FPGA Block Diagram.

Conceptually, the FPGA is divided into sections. The Fault Logic, inside the dotted box, implements the MPS engine. This logic contains all of the fault processing and communications functions including the Ethernet interface. The sensor inputs are passed into a debouncer block which has eight programmable debounce times ranging from 10µs to 5s; this block helps filter out noise (e.g. bubbles on waterflow switches), preventing false trips. The debouncer feeds the fault latch logic, whose output is sent to the LP. The output logic consists of a pulse generator which generates a 9ms pulse when a mitigation command is received. This output must be refreshed at each communications sequence, otherwise it will time out. This is done to avoid a latched failure mode in the case of a link failure. The fault logic is controlled by a state machine which sequences operations and communication. The fault logic also receives inputs from the IP interface dual-port RAM block. For certain applications, the IP modules can provide fault information. This fault data is passed to the Ethernet

Copyright © 2011 by the respective authors — cc Creative Commons Attribution 3.0 (CC BY 3.0)

interface so that the LP can receive it. The ColdFire interface is used to control the IP modules and other LN chassis housekeeping, control and status functions. In addition, the ColdFire has access to the fault latch data in a read-only manner.

LINK-NODE FLAVORS

This section describes each of the types of Link-Node currently being used in the system. It should be noted that in all configurations, the MPS Digital I/O boards can be installed.

Mitigation Link-Node

The Mitigation LN is the simplest configuration. It consists of just the motherboard and MPS Digital Input/Output boards. The configuration of I/O boards is done on a per-instance basis depending upon device requirements.

BLM Link-Node

The Beam Loss Monitor LN is the most complicated. It is part of the BLM readout system [6]. This LN can connect to up to eight undulator BLM detector head and front end electronics box pairs. It is a triggered device, receiving triggers from the timing system [5]. A custom eight-channel IP module, the IP-QINT-ADC (QADC), was developed at SLAC to digitize and process the detector signals. The BLM LN sets and reads back the PMT high voltage power supplies (located in the front end electronics) using COTS DAC and ADC IP modules. The LN also generates a test pulse to exercise the BLM PMT & readout chain. The BLM's high voltage and fault thresholds are setup and controlled from EPICS via the ColdFire CPU. Detector analog readout data is also sent over EPICS. Thus, the BLM LN can also be used as a radiation monitoring instrument.

PIC Link-Node

The Protection Ion Chamber LN is similar to the BLM LN. It interfaces the PIC loss monitors to the MPS. The output of the PIC chamber is connected directly to the LN. The chamber signals are digitized and processed using the same custom QADC module, operating in PIC mode. The LN can accommodate up to four IP-QINT-ADC modules, allowing for a total of 32 PIC chambers to be connected to a single PIC LN. This LN is a triggered device and also controls the PIC high voltage power supply chassis. Fault thresholds and analog data are setup and read out respectively over EPICS via the ColdFire.

ByKIK Link-Node

A single bunch beam dumper magnet, called ByKIK, is located upstream of the Undulator. This magnet interfaces to the MPS through the ByKIK LN. This LN receives Standby and Abort triggers from the timing system and provides a conditioned trigger to the kicker pulse generator, qualified with an MPS permit signal. The kicker outputs an analog waveform that indicates the quality of the kick pulse. This waveform, is digitized (using a COTS IP ADC) by the ByKIK LN and the sampled waveform is compared against thresholds. If the comparison exceeds these thresholds, an MPS fault is generated and the beam is disabled. Specially configured PIC L- and Interface Boards connect the kicker signals to the ADC.

LINK-NODE SOFTWARE

The LN code is EPICS running on the RTEMS RTOS. EPICS applications code is written in C. Information from the LN is available as EPICS Process Variables (PVs) displayed on specialized EDM panels. The PVs are read out of the LN at a 0.5Hz polling rate. Critical PVs are save/restored, archived and connected to the alarm handler.

OPERATIONAL EXPERIENCE

The LCLS MPS was commissioned in 2009. Currently, there are a total of 32 Link-Node chassis in the system. Prior to commissioning, an interim MPS was setup using legacy hardware. By summer of 2010, all of the legacy TIU and 1553-MPS inputs were moved over to LCLS MPS Link-Nodes; fully transitioning MPS functionality for LCLS to the LCLS MPS. The system has been running satisfactorily for almost two years now.

FUTURE DIRECTIONS

This same system will be used for machine protection in the upcoming LCLS-II machine. Plans exist to upgrade the LN ColdFire CPU with a more powerful COM Express platform unit. Based on recent needs, the design of two more flavours of LN are being considered: thermocouple input and general-purpose analog input units. The flexibility of the platform makes these latter items relatively easy to implement.

REFERENCES

[1] R. B. Neal, et al., SLAC Blue Book "The Stanford Two Mile Accelerator", chapter 21, section 21-4 "Equipment Protection Systems"; 1968

[2] S. Clark, et al., "Smart Machine Protection System" ICLAPECS 1991 Conference, SLAC-PUB-5688.

[3] S. Norum et.al., "The Machine Protection System for the Linac Coherent Light Source", PAC09, Vancouver B.C. Canada, FR5REP039. May, 2009.

[4] P. Krejcik, "LCLS Machine Protection System Requirements" SLAC/LCLS Internal Document, ESD 1.3-312-r1, March 2, 2006

[5] J. Dusatko et.al., "The LCLS Timing Event System", 2010 Beam Instrumentation Workshop, TUPSM083; May, 2010.

[6] J. Dusatko et.al., "The LCLS Beam Loss Monitor Readout System", 2010 Beam Instrumentation Workshop, TUPSM084; May, 2010.

Copyright © 2011 by the respective authors — cc Creative Commons Attribution 3.0 (CC BY 3.0)