THE LCLS TIMNG EVENT SYSTEM[*]

J. Dusatko[#], S. Allison, M. Browne, P. Krejcik,
SLAC National Accelerator Laboratory, Menlo Park, CA 94025 U.S.A.

Abstract

The Linac Coherent Light Source requires precision timing trigger signals for various accelerator diagnostics and controls at SLAC-NAL. A new timing system has been developed that meets these requirements. This system is based on COTS hardware with a mixture of custom-designed units. An added challenge has been the requirement that the LCLS Timing System must co-exist and "know" about the existing SLC Timing System. This paper describes the architecture, construction and performance of the LCLS timing event system.

INTRODUCTON

The Linac Coherent Light Source (LCLS) at the SLAC National Accelerator Laboratory is a fourth generation X-ray light source. It uses the last $1/3^{rd}$ of the SLAC Linac to create 4.6 to 13.6 GeV, electron bunches at a nominal charge of 0.05 to 1.0 nC which are sent into a single-pass permanent magnet undulator FEL. Through the Self-Amplified Spontaneous Emission (SASE) process, the undulator produces tuneable 1.5 to 15 Angstrom hard X-rays. A precision timing system is required to coordinate the acceleration process, as well as the triggering of diagnostics to measure beam performance. In addition, experimental data acquisition systems require precision timing reference and triggers to properly coordinate their measurements in coincidence with the LCLS beam.

The comprehensive LCLS Timing System consists of three levels of timing and synchronization, correlating to three different, but interlinked timing subsystems [1]. The first is a precision RF phase reference generation and distribution system needed to drive the injector gun and acceleration klystrons as well as the gun laser system and keep the phase within few fractions of a degree of S-band. The second is an ultra-stable timing phase reference distribution system that is used to synchronize the LCLS electron beam with experiment pump-probe laser systems to a level of 10fs at the photon end of the machine [2]. The third is the timing trigger generation and distribution system which produces precision timed triggers for use by acceleration, diagnostic and experimental DAQ subsystems. This has been designated the Timing Event System to distinguish it from the other two timing subsystems, and is the focus of this paper.

SYSTEM OVERVIEW AND DESIGN

The Timing Event System was developed to meet the requirements of the LCLS [3]; table 1 lists these requirements. It should be noted that the original trigger jitter requirement was specified as 2ps rms. This requirement was driven by the bunch length measurement. The requirement was later relaxed to the

10-ps value after it was determined that level of trigger jitter was not needed and would require a larger effort in the system design to achieve.

Table 1: LCLS Timing Trigger System Requirements

Requirement	Value
Maximum trigger rate	360 Hz
Clock frequency	119 MHz
Clock Precision	20 ps
Coarse step size	8.4 ns ± 20 ps
Delay range	>1 sec
Fine step size	20 ps
Maximum timing jitter (w.r.t. clock)	10 ps rms
Differential error, location to location	8 ns
Long term stability	20 ps

The previous generation SLAC Linac timing trigger system is CAMAC based and developed in the 1980s for the Stanford Linear Collider (SLC) [4]. This system had served the machine well for many years, but was not extensible for LCLS from a requirements, HW obsolescence and SW development point of view. It was not practical to replace the entire SLC timing system (where existing system requirements were sufficient) for LCLS from a cost and retrofitting effort standpoint. Thus, the LCLS Timing Event System had to be layered on top of and function alongside the SLC Timing System; at least for the initial running of LCLS. Because of this, certain architectural tradeoffs were made and the LCLS event timing system was slaved to the SLC system, receiving reference and synchronization signals from it, which are ultimately used to form the LCLS timing triggers.

During the design phase of the LCLS Timing System, several timing trigger solutions were considered including COTS (Commercial Off The Shelf) units. COTS was emphasized as a implementation solution for LCLS controls as a way of cost savings by leveraging industrial solutions that could meet system requirements. An accelerator timing component supplier, MicroResearch Finland (MRF) [5], was identified and selected. MRF had developed the Event System, originally for the Swiss Light Source and later commercialized it. The MRF Event

---

System design was conceptually based on the ANL APS Event System, but upgraded and improved.

The EVG and EVR

The MRF hardware is central to LCLS Event Timing, and is now described in detail. The Event System consists of one or more EVent Generator (EVG) timing masters, capable of being daisy-chained to send data in a priority-encoded manner. The EVG broadcasts serialized timing data, along with an embedded clock to multiple EVent Receiver (EVR) timing clients. The system uses the 8b10b linecode standard to broadcast serial timing data over fiber optic links. The timing data consists of a 16-bit word, which is further divided into two bytes. The upper byte contains timing trigger bytes called event codes, giving 256 distinct combinations. The event codes are used to create triggers in the EVR. The value of the event codes broadcast depends on certain pre-defined conditions setup by external information and triggers into the EVG. The event codes are received by the EVR and, depending on the setup, certain selected codes are used to generate triggers. The other half of the 16-bit timing word sent across the serial link contains two multiplexed data streams. The first contains the distributed databus data, and the second contains the EVG data buffer data. The distributed databus provides an 8-bit HW signalling interface between the EVG and EVRs; it is not used for LCLS. The EVG contains a 2K data buffer, whose contents can be sent to a receiving data buffer in the EVRs. The transmission of the buffer is triggered by software. The EVG is a 6U VME module, which is setup and controlled via a VME CPU.
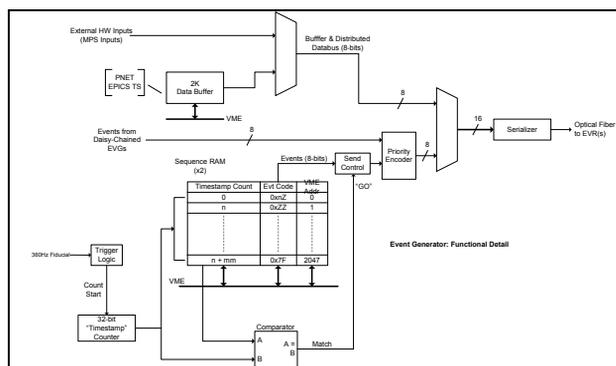

Figure 1: EVG block diagram.

A detailed block diagram of the EVG is shown in Fig 1. At the core of the EVG is the Sequence RAM. This RAM contains the event codes which get broadcast to the EVRs. Each event code in the sequence RAM is associated with a timestamp value. This timestamp is not to be confused with the EPICS timestamp. One sequence RAM timestamp is a value associated with one event code. The timestamp value is 32-bits wide. When the EVG receives a trigger telling it to begin sending out event codes, said trigger actually starts the timestamp counter running. The counter value is presented to the

sequence RAM and when a match between the counter value and the event code's associated timestamp value is made, that event code is sent out. This mechanism allows the user to have temporal control over when the event code is sent out following a trigger. There are two sequence RAMs, permitting one to used for actively broadcasting while the other is being setup for the next trigger. The Sequence RAM is 2K deep and can be setup to either broadcast the full contents per trigger, or continuously loop, sending out the same event codes after one trigger has been received. The EVR requires a system clock for timing and sequencing its operation, this is typically the accelerator system reference clock. This clock is used for all internal logic and is multiplied up by serialization hardware for the serial bit stream. The clock is embedded into the serial data stream sent to the EVRs from the EVG. The EVG logic and serializer are implemented in a Xilinx FPGA.
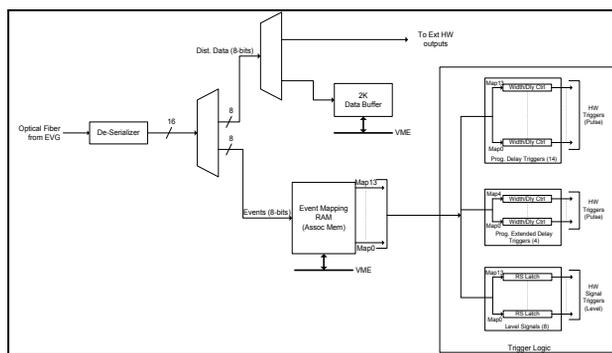

Figure 2: EVR block diagram.

A block diagram of the EVR is shown in Figure 2. The EVR's logic is also implemented in a Xilinx FPGA. The EVR contains an associative memory, called the Mapping RAM. The mapping RAM is programmed with a desired event code associated with one of fourteen mapping bits. The mapping bits are used to form individual hardware triggers. When the event code broadcast by the EVG is received and de-serialized, it is presented to the mapping RAM. If the event code matches a location in the mapping RAM, a hit is generated on the associated map bit. This map bit generates a trigger. The triggers are fully programmable in delay (from the assertion of the map bit), width and polarity. The EVG clock embedded into the serial data stream is recovered and used to clock all of the EVR trigger logic. Thus, the triggers are synchronous with the system reference clock. The EVR also contains a 2K data buffer which receives the EVG's buffer. The buffer data, once received, can signal this condition by sending an interrupt to the EVR's controlling CPU. EVRs are available in VME and PMC formfactors. Each type can supply 14 output triggers with several configurable signalling levels. LCLS uses TTL level triggers for all client subsystems. MRF also supplies a special version EVR (VME-EVR-RF230) which will supply the recovered system clock to external applications.

System Architecture

A block diagram of the LCLS Timing Event System is shown in Fig 3. The MRF EVG and EVRs are arranged in a star topology with one EVG broadcasting timing data to multiple EVR subsystem clients across the accelerator complex. When triggered, the EVG broadcasts these event codes serially at 2.38Gb/s to the EVRs over a fiber optic distribution system. The items in Fig 3 connected to the EVG support the forming of the event codes and generation of the clock and trigger needed to broadcast them. The EVG receives special timing data in the form of a 128-bit timing pattern from the SLC Master Pattern Generator (MPG). The pattern data is broadcast over a dedicated network at SLAC called PNET. The PNET data contains information about the machine rate, beam states, and so on. When a PNET is pattern is received, the PNET module generates an interrupt to the VME CPU located in the EVG crate. This CPU loads the PNET data along with data from the EPICS control system and forms the event codes which get loaded into the EVG's sequence RAM. PNET data is broadcast in a pipelined fashion, with data for three up coming machine cycles being sent. Once the data is loaded into the sequence RAM, it can be sent out when the 360Hz trigger is received. The other hardware provides a clock and trigger for the EVG. LCLS runs at a maximum of 120Hz, with sub-multiple rates being available. The Linac timing system broadcasts a *Fiducial* pulse at 360Hz. The Fiducial is the fundamental metronome of the entire timing system complex. The 360Hz fiducial is derived from the AC mains, and is used to sequence the firing of various accelerator systems in a consistent manner so as not to perturb the AC line. This pulse is amplitude-modulated onto the accelerator 476Mhz phase reference carrier that is driven onto the Main Drive Line that runs the entire 3000 meters of the Linac. The MDL is tapped off at the LCLS Linac Sector 20 RF hut and the Fiducial is demodulated by an SLC Timing FIDO (FIducial DetectOr) chassis. The raw 360Hz fiducial is passed into the Sync/Div chassis which was designed for LCLS. This chassis re-synchronizes and conditions the fiducial and outputs it to the EVG. The fiducial triggers the EVG to broadcast the event codes in its sequence RAM to the EVRs. At SLAC, the 360Hz fiducial is divided down into six 60Hz "timelslots", which correspond to different beam types. LCLS, running at 120Hz occupies timeslots 1 & 4 (the MPG keeps track of the timeslots and sends out special codes corresponding to each one). Other timeslots in the system were reserved for legacy systems such as PEP-II, SLC and others. The timeslots provide a way of intermixing different beam types in the machine.

Figure 3: LCLS Timing Event System diagram.

The EVG also receives a 119Mhz system clock from the Sync/Chassis. This clock is a divided-by-4 version of the 476Mhz accelerator clock. This clock comes from the LCLS master oscillator which performs two important operations: 1) It provides a very low phase noise clock phase-locked to the MDL clock for the LCLS RF system.

And 2) It provided a stable clock source that masked the +/-720 deg phase shift present on the MDL for PEP-II storage ring injection. Item 2 is no longer relevant since PEP-II has been shut down, but was important when LCLS first began operation. Although the EVG could receive the 476Mhz clock and divide it down directly, there is no way to synchronize the clock division process so that is consistently occurs on a fixed phase of the 119MHz. This is necessary to ensure a fixed phase relationship between the Fiducial (which is synchronous to the 476MHz), the 476MHz and output triggers of the EVRs. Without this fixed phase relationship, the phase of the EVR triggers could vary by several 8.4ns steps depending upon which phase the system was initialized at and could differ from initialization to initialization (e.g. power cycles of the system components). The division process is synchronized by a 120Hz trigger, running on timeslots 1 & 4 from the PDU (SLC Programmable Delay Unit) going into the Sync/Div chassis. The EVG uses the 119Mhz clock to generate the 2.38Gb/s serial data link clock, which is embedded into the data stream.

TRD Subsystem

Another component of the timing system, is the Timing Reference Distribution (TRD) system. This is located in the upper right hand corner of Fig 3. This system distributes the 119Mhz clock with a phase-modulated 360Hz fiducial encoded onto it. This timing signal is provided for systems requiring a basic timing reference clock and trigger without the overhead of the EVR. The LCLS Undulator Cavity BPM, Toroid Charge Monitor, Beam Containment and Machine Protection systems use this signal for clocking and/or trigger sequencing. The TRD timing signal is generated by the Sync/Div chassis. The phase modulation consists of one "missing" 119MHz clock cycle and is easily demodulated using purely digital techniques. The TRD signal is distributed over fiber optic links, using spare dark fibers from the timing cable plant. It should be noted that this signal is not phase-stabilized. The goal of this system was simplicity and it takes advantage of commodity Small Formfactor Pluggable (SFP) electrical to optical transceiver modules for signal transmission and reception. The outputs from the TRD chassis are TTL-level for all client subsystems.

System Software

The LCLS Timing Event system is controlled by 360hz high priority software tasks running on the RTEMS real-time operating system. The user interface and diagnostics for these tasks are implemented using EPICS. On the EVG side, the system software coordinates receiving the 360hz PNET pattern as well as other pulse-unique information, including the EPICS timestamp. It determines which event codes need to go out for the pulse and loads them into the sequence RAM. In addition, it determines if the pulse applies to any beam-synchronous-acquisition (BSA) user request and sets BSA flags appropriately. It then immediately sends 360hz data, consisting of the PNET pattern, EPICS timestamp, BSA

flags, and other pulse-unique information, using the EVG data buffer. At reception of a fiducial trigger (about 2 ms later), the sequence RAM is sent over the link to awaiting EVRs.

On the EVR side, the 360hz software task, which is activated on the fiducial trigger (about 1 msec before beam), processes the buffer data received from the EVG and makes it available in memory for other data acquisition tasks on the CPU, i.e. to timestamp the data and determine if there is beam on this pulse. The BSA flags are also stored and used by data acquisition tasks, later on the same pulse, to update BSA arrays. At a lower priority and on user request, the EVR EPICS software controls the mapping of event codes to HW triggers and the set up of the trigger parameters.

The BSA software allows a user to gather beam-dependent scalars across multiple CPUs over multiple pulses up to 120Hz. The data is gathered in BSA arrays (maximum of 2800 values) on each CPU, where each array is assigned to a user. Each value in the array can be an average of up to 1000 values. The user can specify the machine conditions of interest and only pulses that satisfy those conditions will be updated in that user's BSA arrays. The user interface for reserving, releasing, setting up, and activating BSAs is done by EPICS on the EVG CPU. The EVG 360hz task determines if the pulse matches each user request and sends flags to the EVR 360hz tasks for data gathering. The EVR CPU data acquisition tasks look at the BSA flags and update the appropriate data arrays, with optional averaging. Once the BSA is finished, the EPICS software on the EVG CPU notifies the user. The user then reads the assigned data arrays on all CPUs of interest and then releases the BSA slot for another user.

Trigger Generation Process

The following sequence of operations occurs in the formation of a trigger:
1) The PNET message is broadcast by the master pattern generator (MPG) following a fiducial. The PNET message is for 3 cycles ahead which requires all software to pipeline the data.
2) The VME PNET receiver in LCLS master timing crate receives the PNET broadcast and interrupts the VME CPU.
3) The master timing VME CPU takes the PNET data and uses it to assign proper event codes in the sequence RAM which is then ready for the next fiducial.
4) The external fiducial signal (at 360hz) is received by the VME EVG.
5) The EVG begins to send out the event codes in its Sequence RAM.
6) The event codes get sent across the serial fiber links to the EVRs.

7) When the event code matches the same value in the mapping RAM, a hit is generated and after a programmed delay, a HW trigger is output from the EVR to the device.

SYSTEM IMPLEMENTATION

The system was designed for integration within the LCLS control system framework; which is based on the VME-64x bus platform, running EPICS. A key design decision was that each accelerator control and diagnostic subsystem would have its own CPU and VME crate. Thus, each subsystem has its own EVR. The EVG, along with the CPU, PNET receiver and fiber Fanout modules reside in the master timing crate. This, along with the other supporting hardware, occupy one 19inch rack in the Linac sector 20 RF Hut, adjacent to the LCLS injector. All subsystems use PMC-EVRs, with the exception of Low-Level RF and Beam Position Monitors, which use VME EVRs (which have lower jitter). Physical connection of the triggers to subsystem units are typically made using RG-174 lemo cables going from either the EVR front panel or the rear trigger transition module. A dedicated fiber distribution system was designed and implemented to supply timing over the full 2KM of the LCLS machine. Cost was a significant factor in designing the cable plant. One key design requirement was that no more than four layers of timing fanout could be used, in order to keep within the system jitter budget. For the Linac, timing had to be distributed to ten Linac sectors (Sectors 20…30). Beyond the end of the Linac (sector 30) timing was required in the Beam Switchyard, the Linac to Undulator line, the Undulator, Main Dump, the X-ray Line and finally, the Near and Far Experimental Halls. Due to the distances involved, multi-mode Fiber could not be used (at 2.38Gb/s the maximum distance is ~300meters) and required the use of single-mode fiber and transceivers. The system uses Single-Mode fibers for all long-haul distribution trunks and Multi-Mode for all local distribution (e.g. connections to EVRs within one service building). The decision to use mixed fiber types was based on cost, with single-mode being higher. The MRF components use industry-standard Small Formfactor Pluggable (SFP) fiber optical transceivers. A single-mode SFP transceiver was identified (Avago AFCT-57R5APZ), evaluated and used. One item to note is that the fiber distribution system is currently uncompensated for fiber phase drifts due to temperature. While this was not a fundamental requirement, it may become important for future applications.

PERFORMANCE

The trigger jitter of the system was measured using the EVG fiducial input as a reference and measuring multiple time intervals between it and an EVR (MRF VME-EVR-200) output trigger (setup to trigger at the same rate at the fiducial). This measurement was performed using an Agilent model 54845A Infinium digital oscilloscope in time interval histogram mode. The resulting measurement showed the trigger jitter to be better than 10ps rms, over thirty minutes. Factoring out the instrument's internal jitter, the measured jitter is closer to 7ps rms.

FUTURE DIRECTIONS

The LCLS Event Timing System has performed satisfactorily to date over the run of LCLS. Some future upgrades of the system are being planned. The largest of which is the eventual separation of the LCLS Timing Event system from the legacy SLC Timing System. This mainly involves no longer slaving LCLS timing to the SLC MPG. A project is underway to make the EVG operate stand-alone from the MPG. This work involves implementing software and hardware to determine various timing conditions based on machine states and timeslots for formation of the EVG Event Codes. Additional upgrades are planned including the eventual replacement of SLC/CAMAC controls and timing across the accelerator complex. In addition, some work and planning is underway to implement a phase stabilization mechanism to deliver stabilized triggers for experiments.

ACKNOWLEDGEMENTS

REFERENCES

[1] P. Krejcik et.al., "Timing and Synchronization at the LCLS" DIPAC 2007, Venice, Italy May 2007, SLAC-PUB-12593

[2] J. Byrd. "Development of a Sub-Picosecond Timing and Synchronization System for LCLS", Project Overview Slides, LCLS Laser Synchronization Preliminary Design Review, Feb 13, 2008.

[3] P. Krejcik, "LCLS Timing System Requirements", LCLS Physics Requirement Document #1.1-305-r0

[4] L. Paffrath, et.al., "A New Timing System for the Stanford Linear Collider", IEEE-NSS, Orlando, FL October 31, 1984; SLAC-PUB-3508.

[5] MicroResearch Finland Oy, http://www.mrf.fi