

Document Approval:

Date Approved

Originator: Matthew Weaver, Timing System Lead		9/9/16
Approver: Bob Dalesio, Controls Systems Manager	Email Approval	10/4/16
Approver: Matt Boyes Controls Account Manager		7 Sep 2016
Approver: Darren Marsh, Quality Assurance		9/8/16
Approver: David Schultz, Project Technical Director		10-3-16

Revision History

Revision	Date Released	Description of Change
R0	October 4, 2016	Original Release.

Table of Contents

1	Purpose	2
2	Scope	2
3	Definitions	2
4	References	2
5	Responsibilities	2
6	Overview	2
	6.1 External Agents	2
	6.2 Behavior	3

1 Purpose

This document describes the required behavior of the LCLS-II timing system.

2 Scope

The requirements are derived from the timing system physics requirements document and operational experience from LCLS-I. The requirements reflect the current understanding of division of responsibilities as laid out in the interface control document. The requirements reflect some estimate of their impact on design without detailed assumptions of the design.

3 Definitions

Term	Definition
CAN	Channel Access Network
MPS	Machine Protection System
BCS	Beam Containment System
BPM	Beam Position Monitors
BSA	Beam Synchronous Acquisition
XTES	X-ray Transport and Experimental Systems

4 References

Document Number	Document Title
LCLSII-2.7-PR-0345	Timing System
LCLSII-2.7-IC-0314	LCLS-II Phase reference and Timing Systems to Accelerator Systems, Cryogenic Systems, Photon Systems, Infrastructure Systems and LCLS-I

5 Responsibilities

Person(s) or Areas Responsible	Define Responsibility
--------------------------------	-----------------------

6 Overview

6.1 External Agents

The behavior of the timing system is described below in terms of its interaction with the “external agents”. The external agents for the timing system are:

1. The AC Power Line
2. The Channel Access Network (CAN)
3. The Machine Protection System (MPS)
4. The Beam Containment System (BCS)
5. The Beam Position Monitors (BPM)
6. The Low Level Radio Frequency (LLRF) System
7. The Beam Synchronous Acquisition System (BSA)

8. X-ray Transport and Experimental System (XTES)
9. LCLS-I timing system

The collection of devices which acquire sensor data with input from the timing system are referred to as “timing clients”. These include components from the MPS, BCS, BPM, LLRF, BSA, XTES, and others (including the common HPS architecture). Timing clients are expected to have a local host CPU interface.

6.2 Behavior

The master source of the timing system shall generate the set of phase locked reference timing signals listed in Table 1.

Table 1. Master source reference signals

Frequency	Purpose
1300MHz	LLRF phase reference
3900MHz	3 rd harmonic LLRF phase reference
185.7MHz (1300MHz/7)	Timing generator reference
476 MHz	LCLS-I phase reference and timing generator input
71.43kHz (1MHz/14)	Common subharmonic of LCLS-I and LCLS-II. Resync all phases of derived timing.

- The phase reference line shall distribute the 1300MHz LLRF reference and the 3900MHz LLRF reference. The stability and jitter requirements on these signals are listed in Table 2.

Table 2. Phase reference timing requirements.

Parameter	Value
Jitter (integrated 50Hz to 5kHz)	10fs
Stability (1 sec)	1fs
Stability (1 day)	1ps

- The master source will eventually provide the 476MHz phase reference for the LCLS-I timing system in order to achieve synchronized operation between the two accelerators. The jitter requirement on the 476MHz timing reference is 30fs integrated from 10Hz to 10MHz.
- The 71.43kHz synchronization reference will be distributed among the master source, the LCLS-I timing master, and the LCLS-II timing master. It shall be used to synchronize all clock dividers to fix the phases of all derived timing. This includes the derivation of the 185.7MHz reference from 1300MHz, 119MHz reference from 476MHz, and 360Hz reference from powerline sampling. The combined jitter and stability of this signal shall be less than 150ps.
- The timing system shall distribute the 185.71 MHz clock (1300MHz/7) from the master source to all timing clients. This clock is hereafter labeled the “distributed 185.71 MHz clock”. Its jitter requirement is 30ps, and its stability requirement is 500ps.
- The timing system shall derive a “base rate trigger”, the maximum beam rate, and distribute it to all timing clients synchronous with the distributed 185.71 MHz clock. The base rate trigger frequency is determined by input from the CAN as a particular sub-harmonic of the distributed 185.71 MHz clock. The default base rate is 1300MHz/1400 which is 185.71MHz/200.

- The base rate trigger is accompanied by a monotonic 64-bit identifier, hereafter referred to as the pulseID, so that timing clients may tag data acquired on that base rate trigger. The pulseID may be read and set (restored) via the CAN, but it must remain monotonic over the lifetime of LCLS-II.
- A global EPICS timestamp shall be distributed to all timing clients for tagging EPICS records from scalar dump requests, for example, described later.
- The timing system shall derive from the AC power line input a 360 Hz timeslot, values ranging from 1 to 6, indicating phase with respect to the 60 Hz AC power line. Timeslot 1 shall be defined as one of the three phases of the AC mains consistent with existing systems. It is believed this is currently phase C of the AC mains.
- The timeslot value shall be distributed with the base rate trigger to all timing clients with an accuracy of 100 μ sec.
- The timing system shall generate a trigger to the LCLS-I Main Drive Line Fiducial Generator for each transition of the 360 Hz timeslot and another trigger denoting timeslot 1, replacing that function in the LCLS-I timing system when the two systems operate in synchronized mode.
- For each base rate trigger, the timing system shall distribute a timing pattern with the information listed in Table 3. The timing pattern may be compared or latched by timing client interactions described later. Its contents are enumerated here for illustration only.

Table 3. Timing pattern fields

Field	Description
PulseID	64-bit monotonic base rate trigger counter
Beam present	true/false
Timeslot	AC power line phase w.r.t. 60 Hz {1-6}
Beam destination	HXR dump, LXR dump, D10 dump, diagnostic line, injector spectrometer
Bunch information (DES)	Charge, Energy at { spectrometer bend, BC1, BC2, end of LINAC }, photon wavelength
Timing generator status	OK/fail
Array reset request	Clear the snapshot arrays
Array snapshot requests	Capture data for a set of the snapshot arrays
Array dump request	Freeze the snapshot arrays
MPS rate, destination, fault status	Limitations imposed by the MPS for the injector gun, HXR fast kicker, and LXR fast kicker.
BCS fault status	OK/fail by destination
Experiment control states	Sequenced states for XTES
User-defined states	States forwarded from a feedback (midstream)
Constant rate markers	Fixed rate selections from Table 4
Power line synchronized rate markers	Rate selections fixed to power line frequency from Table 5
EPICS timestamp fragment	Decomposed EPICS timestamp for periodic update

- The timing pattern distribution shall be reliable and uninterrupted by other communication.
- The fixed rate markers in Table 4 shall classify base rate triggers as belonging to a recurring subset with fixed period. Note that the patterns marked as 1 Hz are a subset of those marked as 10 Hz, which are subset of those marked as 100 Hz, and so on.

Table 4. Fixed rate markers

Fixed Rate	1300 MHz Prescale
1 Hz (actually 0.92857Hz)	1,400,000,000
10 Hz	140,000,000
100 Hz	14,000,000
1 kHz	1,400,000
10 kHz	140,000
71.43 kHz	18,200
half-rate	2,800 (nominal)
full-rate	1,400 (nominal)

- The power line synchronized markers in Table 5 shall classify base rate triggers synchronized to a particular timeslot and also occurring at an approximate constant rate. For example, the patterns with this 1 Hz marker and timeslot 1 are a subset of those marked with this 10 Hz marker and timeslot 1, which are a subset of those marked with this 30 Hz marker and timeslot 1, and so on. Similarly, for timeslot 2 and so on. As a consequence,
 - the 1 Hz marker shall appear in approximately 6 patterns each second, once with each timeslot value;
 - the 10 Hz marker shall appear in ~60 patterns each second, ten times with each timeslot value;
 - the 30 Hz marker in ~180 patterns each second, thirty times with each timeslot value;
 - the 60, 120, and 360 Hz markers shall appear in ~360 patterns each second, 60 times with each timeslot value. They are redundant.

Table 5. Power line synchronized markers

Rate	360 Hz Prescale
1 Hz	360
10 Hz	36
30 Hz	12
60 Hz	6
120 Hz	3
360 Hz	1



- The received timing pattern shall be qualified by transmission error and version checking. Base rate triggers from errant timing patterns shall be ignored. A “keep-alive” condition, described later, may be configured to engage after a predefined number of consecutive errant patterns.
- The beam present flag in the distributed timing pattern shall result in the injector gun generating an electron bunch for an RF bucket associated with the base rate trigger. Similarly, the bunch information meta-data shall control beamline devices to generate those bunch parameters. The bunch information meta-data may also be used by beamline monitoring devices in the MPS, for instance, to validate the beam quality against the desired settings.
- For each destination, the CAN may request the sequence of desired beam conditions to distribute. The sequence shall consist of up to 1k steps. The start of each sequence may be selected to synchronize with a fixed rate marker, an AC rate marker, including timeslot choice, or the end of the current sequence for that destination. A sequence may be requested to repeat a fixed number of times, up to 100k, or until requested to stop. Each step in the sequence designates a starting delay {N occurrences of a fixed rate marker from Table 4 or a power-line synchronized rate marker from Table 5, including timeslot choice}, the bunch information parameters from Table 3, and the number of times to repeat the step including delay. The CAN shall also set the order of priority among sequences to resolve instances where beam is requested to more than one destination; i.e. destination arbitration. Prepared pattern sequences may also be requested shorthand by the CAN automating common processes. Up to 2 future sequence requests may be queued for macro-sequences that require seamless sequence changes.


An example sequence is outlined in Table 6.

Table 6. Example sequence

	Condition	Action	Repeat
Sequence Start	Wait for 60Hz TS1	Go to step 1	
Step 1	Wait for 1x100kHz	Generate ½ bunch charge	10 times
Step 2	Wait for 1x100kHz	Generate full bunch charge	20 times
Step 3	Wait for 1x100kHz	Generate ½ bunch charge	10 times
Sequence Stop	After 5 iterations	Go to next sequence	


The required use cases include:

- Selecting the rate of beam to each possible destination. The selected rates can either be fixed rates from Table 4 or power line synchronized rates from Table 5;
- Selecting a train of bunches by the spacing between bunches in the train, the number of bunches in the train, and the spacing between trains;
- Loading the superconducting cavities after a fault with bunch trains of increasing duty cycle;
- Scanning the AC power line with bunch trains of fixed length;
- The CAN shall request experiment control sequencing for the XTES. A separate programmable sequence of control patterns, up to 20 bits in length, for each experimental endstation shall be executed within the distributed timing pattern. Each sequence shall consist of up to 1k steps. The start of each sequence may be selected to synchronize with a fixed rate marker, an AC rate marker, including timeslot choice, or the end of the current sequence for that endstation. A sequence may be requested to repeat a fixed number of times, up to 100k, or until requested to stop. Each step in the sequence

	Functional Requirements Specification Document	
	Document Title: LCLS-II Timing System Functional Requirements	
	Document Number: LCLSII-2.7-FR-0527-R0	Page 7 of 9


designates a starting delay {N occurrences of a fixed rate marker from Table 4 or a power-line synchronized rate marker from Table 5, including timeslot choice}, a 20-bit field, and the number of times to repeat the step including delay. Beam may be synchronized to the execution of the experiment control sequence by requesting a beam condition sequence, described above, whose start is linked to the start of the experiment control sequence.

- The MPS shall provide an input link to the timing system master indicating fault status and limited beam rates to particular destinations. Upon receipt of an MPS fault, the fault shall be latched into the timing pattern and output to the CAN. An array dump request, described later, shall be generated for diagnosis of the conditions leading up to the fault. The limited beam rate and destination shall be distributed to all timing clients to reflect the condition that beam is being diverted. Because the MPS executes the beam disruption with less latency than the distribution of that information in the timing pattern, these fields in the timing pattern necessarily represent information that is stale by several microseconds. The limited rate imposed by the MPS should correspond to one of the markers from Table 4 or Table 5 (including timeslot choice). The fault latch must be reset by the CAN. An invalid input link from the MPS shall also be reflected in the distributed timing pattern for diagnostics.
- The BCS shall provide an input link that signals a BCS fault. Upon receipt of a BCS fault, the fault shall be latched into the timing pattern and output to the CAN, and an array dump request shall be generated. The fault latch is reset by the CAN.
- A manual fault may be received over the CAN. The result shall be a prompt switch over of the indicated beam destination sequences to their default sequence (no beam). A manual array dump request may also be generated upon demand.
- An internal fault may be generated by the timing system generator itself, resetting all beam destination sequences to their default, until the fault is reset by the CAN.
- The timing clients participating in BSA shall reserve buffer sets each capable of caching a history of data points. The timing system shall distribute commands in the timing pattern for clearing those buffers (an array reset request), capturing sensor data to those buffers (an array snapshot request), and latching those buffers (an array dump request) awaiting retrieval of the data by the local host CPU. The latched shall be cleared by an array reset request. Each buffer set shall be individually capable of clearing, capturing, and latching, with the exception that an MPS or BCS fault shall latch all array buffers. The capture shall coincide with a trigger configured by the local host CPU as described below. In addition, one timing client shall capture the full timing pattern for each data point going into the snapshot buffers. The latched data shall be retrievable by a higher level application via the CAN. The number of buffer sets shall be 50.
- For each experimental endstation in the XTES, at least ten additional inputs (bits) shall be provided for distribution with the timing pattern without delay. In order to minimize delay that is critical to the acquisition rate, those inputs should be inserted into the closest common distribution point upstream of the endstation timing clients. The endstation timing clients shall use the additional timing pattern fields in their trigger selection criteria.
- Each timing client shall configure the timing system through its local host CPU interface to generate a set of trigger pulses from base rate triggers matching a selection criteria described below. Either the distributed or high-precision 185.71 MHz clock may be chosen for the base rate trigger synchronization and delays described below. Synchronization of the generated trigger with respect to the beam must satisfy the accuracy requirements in Table 2. Each pulse set shall have the following programmable parameters:
 - One or two configurable delays, i.e. one or two pulses per selected base rate trigger, of from 0 to 1 second, in steps of the chosen 185.71 MHz clock;
 - A pulse width of from 0 to 1 second, programmable in steps of the chosen 185.71 MHz clock;

	Functional Requirements Specification Document	
	Document Title: LCLS-II Timing System Functional Requirements	
	Document Number: LCLSII-2.7-FR-0527-R0	Page 8 of 9

- Positive or negative polarity
- If a second base rate trigger match is found for a given trigger before the delay and pulse width execution is complete for that trigger, the behavior should be configurable as either:
 - The trigger shall not be “re-armed”, and the second base rate trigger ignored; or
 - The trigger shall be “re-armed” and the current pulse reset.
- Trigger pulses may either be external TTL signal outputs or signals on a local carrier by some other standard. The pulseID for each trigger pulse must either be latched for retrieval by the local host CPU or encoded with the trigger signal on the local carrier.
- Trigger delay adjustments common to a timing client, due to unique timing distribution delay, will be compensated by a single delay counter accessible by the CAN. A common application to monitor and set those delays is required.
- The selection criteria for timing client requests shall apply to a base rate trigger at a chosen integer from 1 to 100 base rate triggers in the future, and shall be one of:
 - Some logical combination of matches in the timing pattern fields from Table 3. Commonly used combinations will be given a shorthand representation, “event codes”, for a simpler IOC interface.
 - A “keep-alive” condition, in case of timing distribution failure or commissioning. The base rate trigger for this condition will occur at a fixed rate unsynchronized to the timing system.
- Sensor data acquired from a timing client with a timing system trigger shall be tagged with the 64-bit pulseID either by software using the pulseID retrieved by the local host CPU or by hardware using the pulseID encoded on the local trigger signal carrier.
- A timing client’s local host CPU must be able to receive the timing pattern, the global EPICS timestamp, the 185MHz clocks since timeslot transition, and bunch number within a train for each base rate trigger which generates a trigger pulse. The local host CPU must also be able to request receipt of the same set of information for selection criteria matches without configuring generation of a trigger pulse.
- A scalar dump request shall be generated at specific intervals determined by input from the CAN. Participating timing clients shall synchronously latch sensor data, possibly with a configured timing system trigger, and cache the sensor data conditionally into a set of rate-limited EPICS records { 1Hz, 10Hz, and 100Hz }. For example, the 1 Hz records will only be updated by a selected sample of scalar dump requests to limit changes to once per second.
- The timing system shall update an EPICS record for diagnostics with the pulseID, phase reference input status, and internal status at a rate of < 10 Hz.
- Each timing client shall be able to cache up to 100 timing patterns, synchronously latch that cache based upon a distributed user request, and dump that cache to the local host CPU for diagnostics.
- The timing system in the hard x-ray line must support both LCLS-I and LCLS-II operation. One of two scenarios is possible:
 - The LCLS-I timing distribution system in the hard x-ray line remains intact. The timing receivers are replaced with units with inputs from both the LCLS-I and LCLS-II timing systems to decode either timing; or
 - The LCLS-II timing distribution system also carries LCLS-I timing and the timing receivers are capable of decoding timing from the one input for both accelerators.

The LCLS-I event receiver IOC interface should remain intact.

	Functional Requirements Specification Document	
	Document Title: LCLS-II Timing System Functional Requirements	
	Document Number: LCLSII-2.7-FR-0527-R0	Page 9 of 9

- The timing system must provide a graphical interface for editing collections of beam generating sequences. Each collection shall contain one sequence for each beam destination. The collections may be saved and loaded for editing or copying. A collection shall be retrievable by a meaningful name attached by the creator of the sequence collection. The interface shall provide a visualization of the collection of sequences for the longest period of repetition, maximum of 1 second. The collection shall also contain set points for the injected charge and electron energy at the spectrometer bend, after BC1, after BC2, and at the end of the LINAC. The interface should display the resulting powers at each beam destination for comparison against allowed machine limits.
- The timing system shall provide a graphical interface for retrieving saved beam generating collections (or “beam delivery mode”) for loading into the generator during operation. Transition from one beam delivery mode to another should be possible at well-defined times within the previous and next mode; i.e. no unintended gaps in delivery. Loading of selected favorite delivery modes may be tied to operator console buttons.
- The timing system shall provide a tool for capturing a sequence of generated timing patterns during operation for occasional verification against the intended timing pattern sequence.
- The timing system shall provide an agent for reserving control of the BSA arrays. This agent shall have an associated graphical interface showing the status of control for each BSA array; i.e. its current owner, running state, configuration, and some periodic update of the acquisition’s progress. The interface may also provide privileged control to remove a reservation for an acquiring BSA array.
- Engineering displays shall also be provided for the purpose of diagnosing operation. They shall periodically update with fault status and detailed statistics of operation.