LCLS Experience and Comments on the First Workshop on IT Infrastructure and Control System for SwissFEL

Patrick Krejcik
pkr@slac.stanford.edu

Outline

- Reviewers comments and opinions from the first workshop:
  - *Patrick Krejcik* — *SLAC*
  - *Hans-Joachim Kreidel,* — *Uni. Mainz*
  - *Mark Plesko, Rok Sabjan,* — *COSY LAB*
  - *Nick Rees,* — *Diamond*
  - *Javier Serrano, Jean-Claude Bau* — *CERN*
- *Include some lessons learned from LCLS commissioning and operation that might also be relevant here*

The First Topical Workshop on IT Infrastructure and Control Systems for SwissFEL

- Project planning should account for manpower in both the construction phase of the project and the long-term maintenance of the control system.
  - Users have high expectations on uptime
- Don’t under estimate the scope of the Controls system effort
  - LCLS spent $40M or 10-12% on controls
  - LCLS Controls staff number >80 people
  - “Be prepared to spend 80% of commissioning time on controls”.
    - P. Emma

Patrick Krejcik
pkr@slac.stanford.edu

A high level view on what to change for the SwissFEL CS

- SLS whole system availability is 99%
  - exceptional for an accelerator
  - performance may be related to high price in manpower and hardware
  - => management decision of whether to apply the same approach also to SwissFEL or not
- REC: set a lower limit on the needed accelerator and CS availability and quality of service/support response
  - Then decide acceptable cost for such performance

Mark Plesko

On Manpower

- PSI CS group has more responsibilities than other groups at similar projects
  - 38 people for 3 machines and 20 beamlines, vs. DLS with 37.5 FTE in 2008 for Diamond and 22 beamlines
  - internal economies are always possible, but definitely not enough to build another machine like the SwissFEL.
- REC: don’t focus just on the price of the equipment but also on manpower:
  - if use existing personnel only, must balance between delay of the project and reduction in existing services
  - Reduction of price by changing the existing HW platform, results in additional strain to personnel, amplifying the dilemma above
- additional hiring/outsourcing will have to be accounted for somewhere.

Mark Plesko

Collaboration with LCLS and also ITER

- main challenges:
  - Pulsed operation,
  - Synced data handling included to experiments,
  - several feedbacks needed,
  - interlocking.
- LCLS has solved all these before
- Look also to fusion community
  - As we work for ITER, we have learned that the fusion community always had much more emphasis on those issues - and also has solutions for them.
  - As ITER will be EPICS based, a direct collaboration would be possible.

Mark Plesko

Planning

- Design for flexibility and expandability
  - evolving needs of accelerator physicists and photon beamline users.
- LCLS already has
  - 326 IOCs in total
    - 124 iocs
    - 155 embedded iocs
    - 47 soft iocs
  - 366,902 PVs in total

- Designing for just one purpose might seem like a good idea at the time

Patrick Krejcik
pkr@slac.stanford.edu

Architecture

- PSI well placed to take advantage of EPICS infrastructure
  - make rapid advances by copying some systems from other FEL laboratories (LCLS)
- High level applications important for commissioning
  - Matlab for rapid prototyping
    - changes can be made on the fly in the control room.
    - Widespread at other facilities (LCLS) makes collaboration during commissioning easier.
    - Once requirements for operation become stabilized one should consider using free application programming languages such as Python, Java, C++, QT and CAFÉ.

Patrick Krejcik
pkr@slac.stanford.edu

On Matlab:

- Not needed if issue is operation screens
  - Very few screens need sophisticated mathematical/physics calculations
  - will be most likely run by physicists anyway.
    - They can set up a calculation server that returns the data via TCP/IP (or CA) to a thin client operator screen.
- Is CAFE an alternative to Matlab?
  - as much proprietary, because it is being developed by a single person
  - probably little documentation available
- Open source alternatives see e.g. http://www.math.tu-berlin.de/~ehrhardt/matlab_alternatives.html
- Matlab has large community
  - any new student can be expected to have at least a basic knowledge of Matlab.
  - Matlab/MCA is the XFEL community standard - so one can hire people from those other labs.

Mark Plesko

LCLS Example of high level application

- Matlab application menu driven to select wire scanners along the machine and collect beam size measurements
- Online model data accessed to compute emittances from beam sizes
- Requires middleware support from the Controls Group

Emittance Application

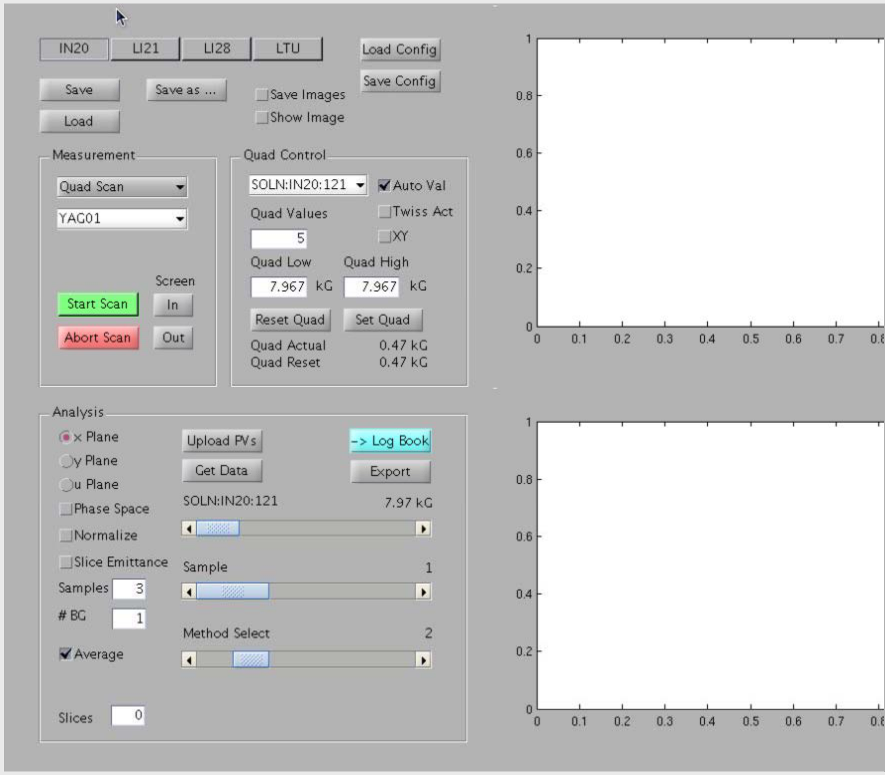

Wire scanner application

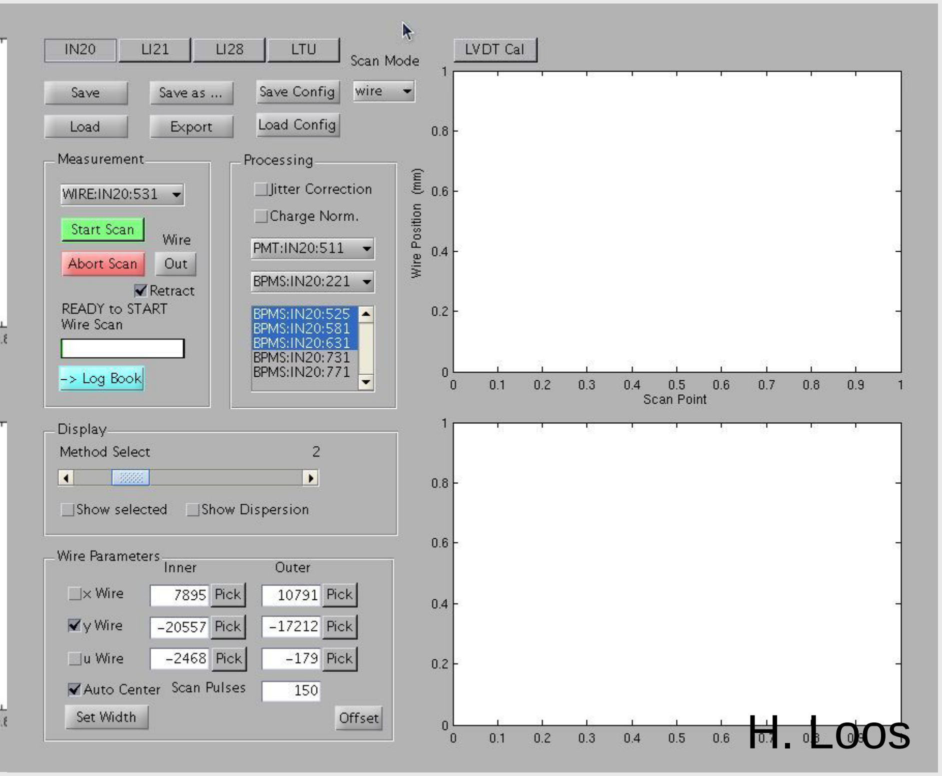

H. Loos

Patrick Krejcik
pkr@slac.stanford.edu

eLogs

- The DOOCS electronic logbook software is a valuable documentation tool during commissioning
- needs some minimal level of system support (unix operating system, disc space allocation, backups etc.).

Repetition rate

- 100 Hz initially and 400 Hz are both multiples of the 50 Hz AC power.
- Consider 300 Hz instead of 400 Hz
  - remains synchronized to the 6 harmonics of the 3-phase AC power distribution.
  - These 6 harmonics inject 6 separate sources of noise (offsets) which are noticeable in a single-pass pulsed machine.
  - difficult to measure until the facility is actually operating at this full rate.
  - at LCLS “time slot” noise offsets are compensated since beam is synchronized to the 6 zero phase crossings of the AC power distribution.

Patrick Krejcik
pkr@slac.stanford.edu

Technology Choices

- Invest in GB Ethernet infrastructure for high data bandwidth requirements at 300 Hz beam repetition rates

Patrick Krejcik
pkr@slac.stanford.edu

Feedback Systems with GB Ethernet?

- Can GB Ethernet be more than a physical layer?
- With the right SW (RT-Linux and RTnet) eth and UDP may become a feasible technology
  - It is long lasting, given its ubiquity.
  - This is another topic for a potential collaboration with ITER, which made some investigations into this solution.

Mark Plesko

Platforms

- Evaluate the advantages of a state-of-the-art Linux-Intel IOC versus real-time operating systems
  - Real-time, fully deterministic requirements driven by beam synchronous acquisition and feedback requirements
  - A Linux IOC platform, possibly with Intel rather than PowerPC, is a favored choice representing state-of-the-art technology amongst many of the reviewers.
    - When weighing up these choices the Controls Group should keep in mind the requirements imposed by a high repetition rate, pulsed FEL machine for a real-time, deterministic operating system.
    - VXWorks is a widespread commercially available real-time operating system, and LCLS uses the free, open-source RTEMS operating system.

Patrick Krejcik
pkr@slac.stanford.edu

SwissFEL – Control System Workshop

Contribution by H.-J. Kreidel (MAMI)

Concerning our (the reviewers) discussion about IOCs (though it is not my opinion they should change the another OS), two slides about the hardware platform:

If one really thinks about changing to another operating system (instead of VxWorks), it may be interesting how we control our VME-crates at MAMI. The main point is, that one needs no VME-driver for the operating system.

VME-CPUs used at MAMI

We mounted an ETX-board (Kontron, Intel-CPU) on a single-width VME-card and programmed a CPLD (Lattice) to convert the PCI-bus signals directly into VME-bus standard.
There is no bridge involved !

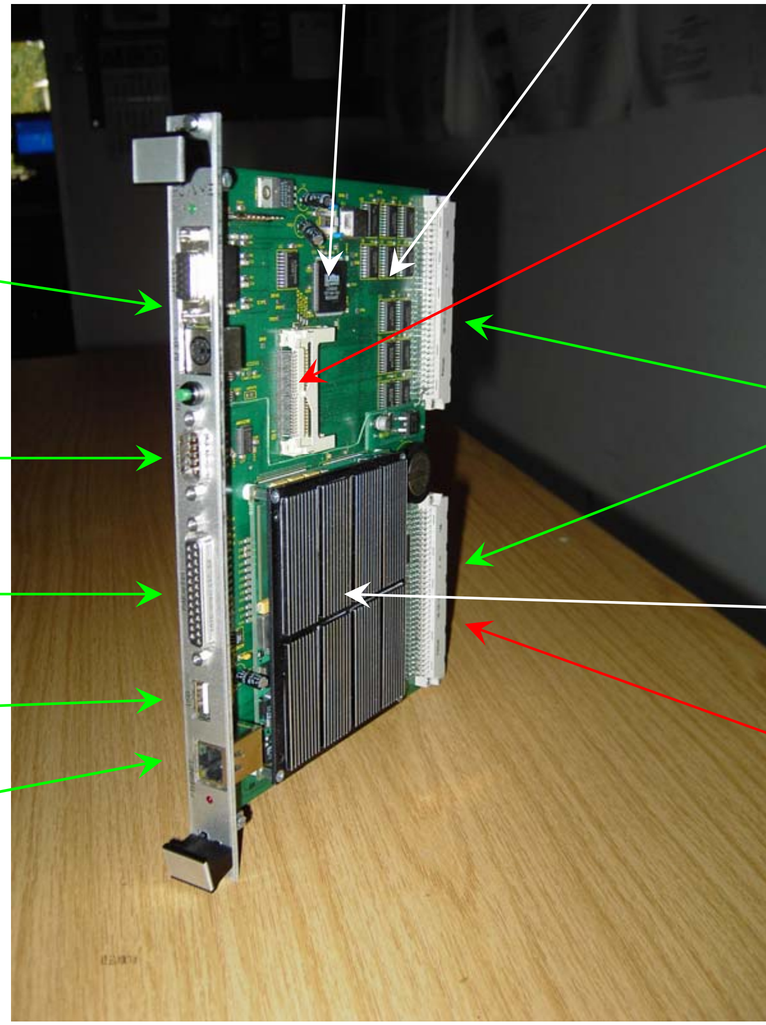

Some minor constraints resulted:

- No bus-arbitration; means: Only 1 CPU per VME-crate
- The address space mapped at once is limited to 28 Bit

Advantages:

- Low hardware costs (<1000€/CPU)
- No operating-system driver for the VME-bus access is required. Thus we are practically free in the choice of the OS.
  Presently we use FreeBSD (also Linux is possible) accessing VME via the mapped /dev/mem device, but one may use even Windows together with any free multi-purpose driver.
- VME-access is very fast (like ext. memory access, OS is not involved)
- OS may be on-board (in a Flash-memory or on a hard-disk connected via the VME-P2-connector)

Disadvantage:

- Without a VME-driver one cannot process interrupts.
  But using a multi-core CPU polling may be acceptable and will enable faster response than a standard OS-interrupt.

Intel platform solution

- Board Support Package (BSP)
  - BIOS is well defined (standardized) → Suppress vendor dependency
  - For PowerPC platforms : No standardization → You have to pay for BSP, Operating System upgrades could be postpone to wait BSP upgrades
- Active Management Technology (Intel®AMT)
  - Allows remote, power and security management features
  - OS independent (firmware in the board)
  - Defined by Intel → Only on Intel platforms
- Intel to PowerPC migration
  - Problems of endianess must be fixed
- Reference platform for Linux

Intel vs. PowerPC ?
The question should be : Why a PowerPC ?

Linux for Real Time application

- Hard RT with RT_PREEMPT patch
- Is Open→ you know how things tick inside and you *can* change them
- Is free → Choice between a free or a commercial distribution
- Is widespread and supported by a large and enthusiastic community
- TCP/IP stack is rock solid → Problems have been discovered in LynxOS causing Real Time errors.
- Service Management Interrupts (SMI) can introduce unexpected latencies
  - Interrupts are not seen by the OS→ Problem in all OS
  - Can be blocked for non critical services. Must be seen with the board vendor.
  - Some tests conducted by A. Steinbacher for EBG MedAustron highlight this problem. Irrespective of the OS, interrupt latencies beyond the limits have been detected (1 time for 10 millions of interrupts)

Diagnostics

- Traditionally, the Diagnostics group has developed diagnostics FPGAs, but if EPICS is embedded in them then the controls group needs to be involved.

Patrick Krejcik
pkr@slac.stanford.edu

LCLS Beam Diagnostics

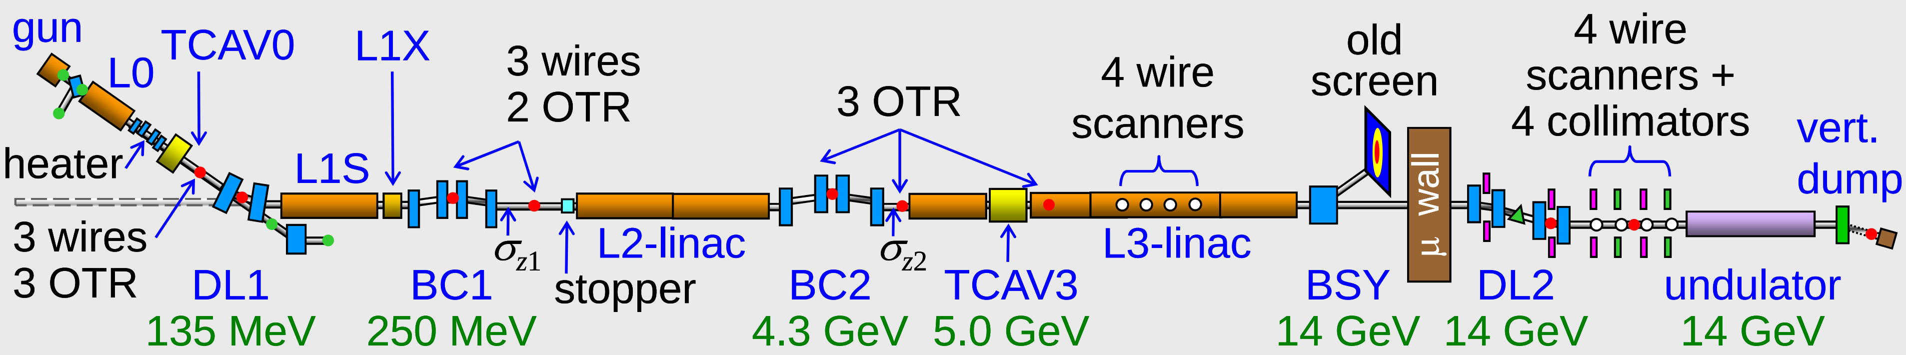

- 2 Transverse RF cavities (135 MeV & 5 GeV)
- 179 BPMs
- 13 Toroids
- 7 YAG screens (at $E \leq$135 MeV)
- 12 OTR screens at $E \geq$ 135 MeV
- 15 wire scanners (each with x & y wires)
- CSR/CER pyroelectric bunch length monitors at BC1 & BC2
- 4 beam phase monitors (2856 – 51 MHz)
- 3 Energy spectrometers: Gun, injector, dump

YAG screens
OTR screens
Wire scanners

Patrick Krejcik
pkr@slac.stanford.edu

Diagnostics

- BPMS and other diagnostic devices will require a synchronous beam data acquisition system.
- At LCLS the requirement was stated that every beamline device should be able to be read out on the same beam pulse and that data should be available at the full beam rate.
- This is necessary to measure and pinpoint sources of jitter in the beam.

Patrick Krejcik
pkr@slac.stanford.edu

BSA at LCLS

- Beam Synchronous Acquisition user setup
  - Reserve one of the 20 buffers available on all IOCs
  - Define which Event Codes are to send data to the buffer
    - Includes beam rate information (1 Hz, 10 Hz full rate etc.)
    - Includes conditional timing expressions
- Initiate data acquisition
- When finished, read buffer contents via CA
  - Gives 2400 samples from EVERY diagnostic device

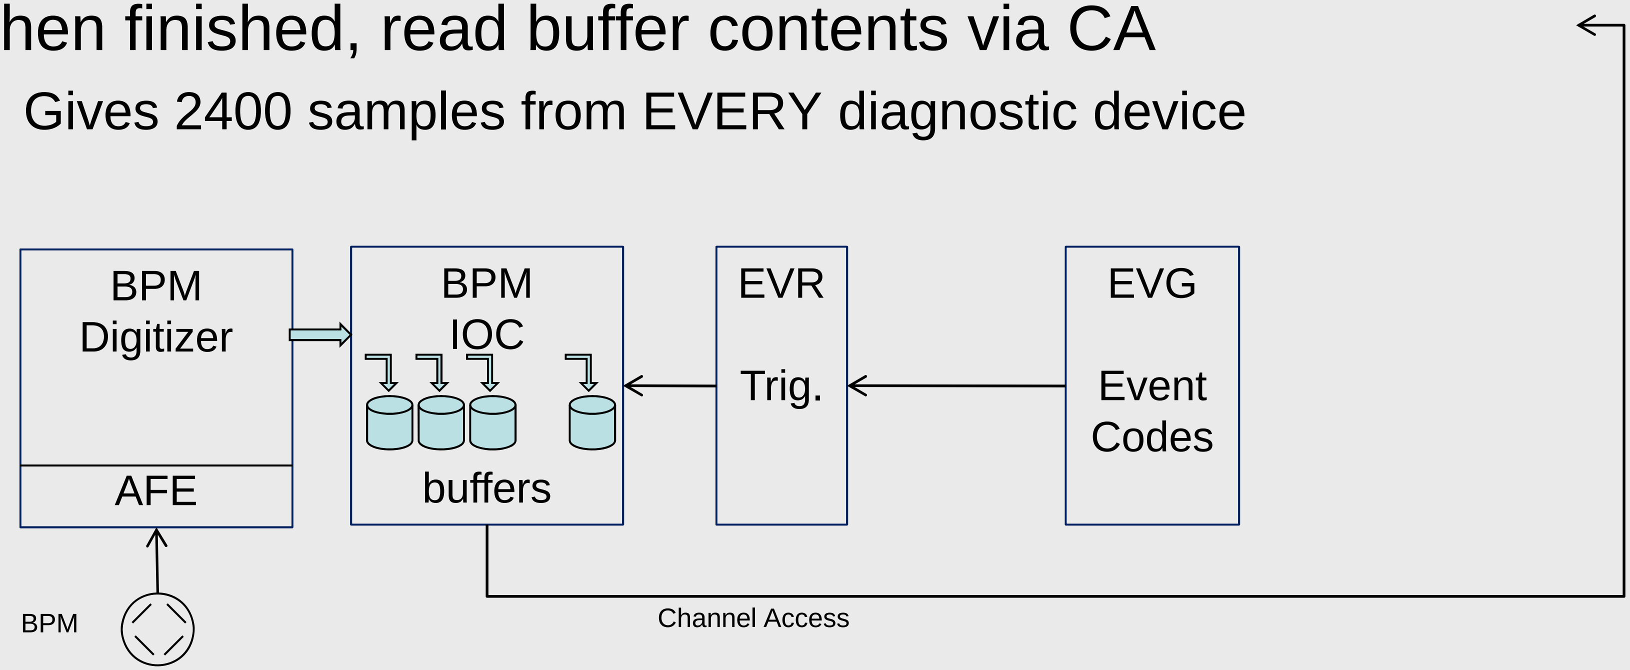

Patrick Krejcik
pkr@slac.stanford.edu

Timing

- The SLS already has an advanced timing system that can be adapted to the FEL. Since LCLS is using the same EVG and EVR system it may be possible to adapt the synchronous data acquisition software mentioned above.

Patrick Krejcik
pkr@slac.stanford.edu

LCLS Timing/Event System Architecture

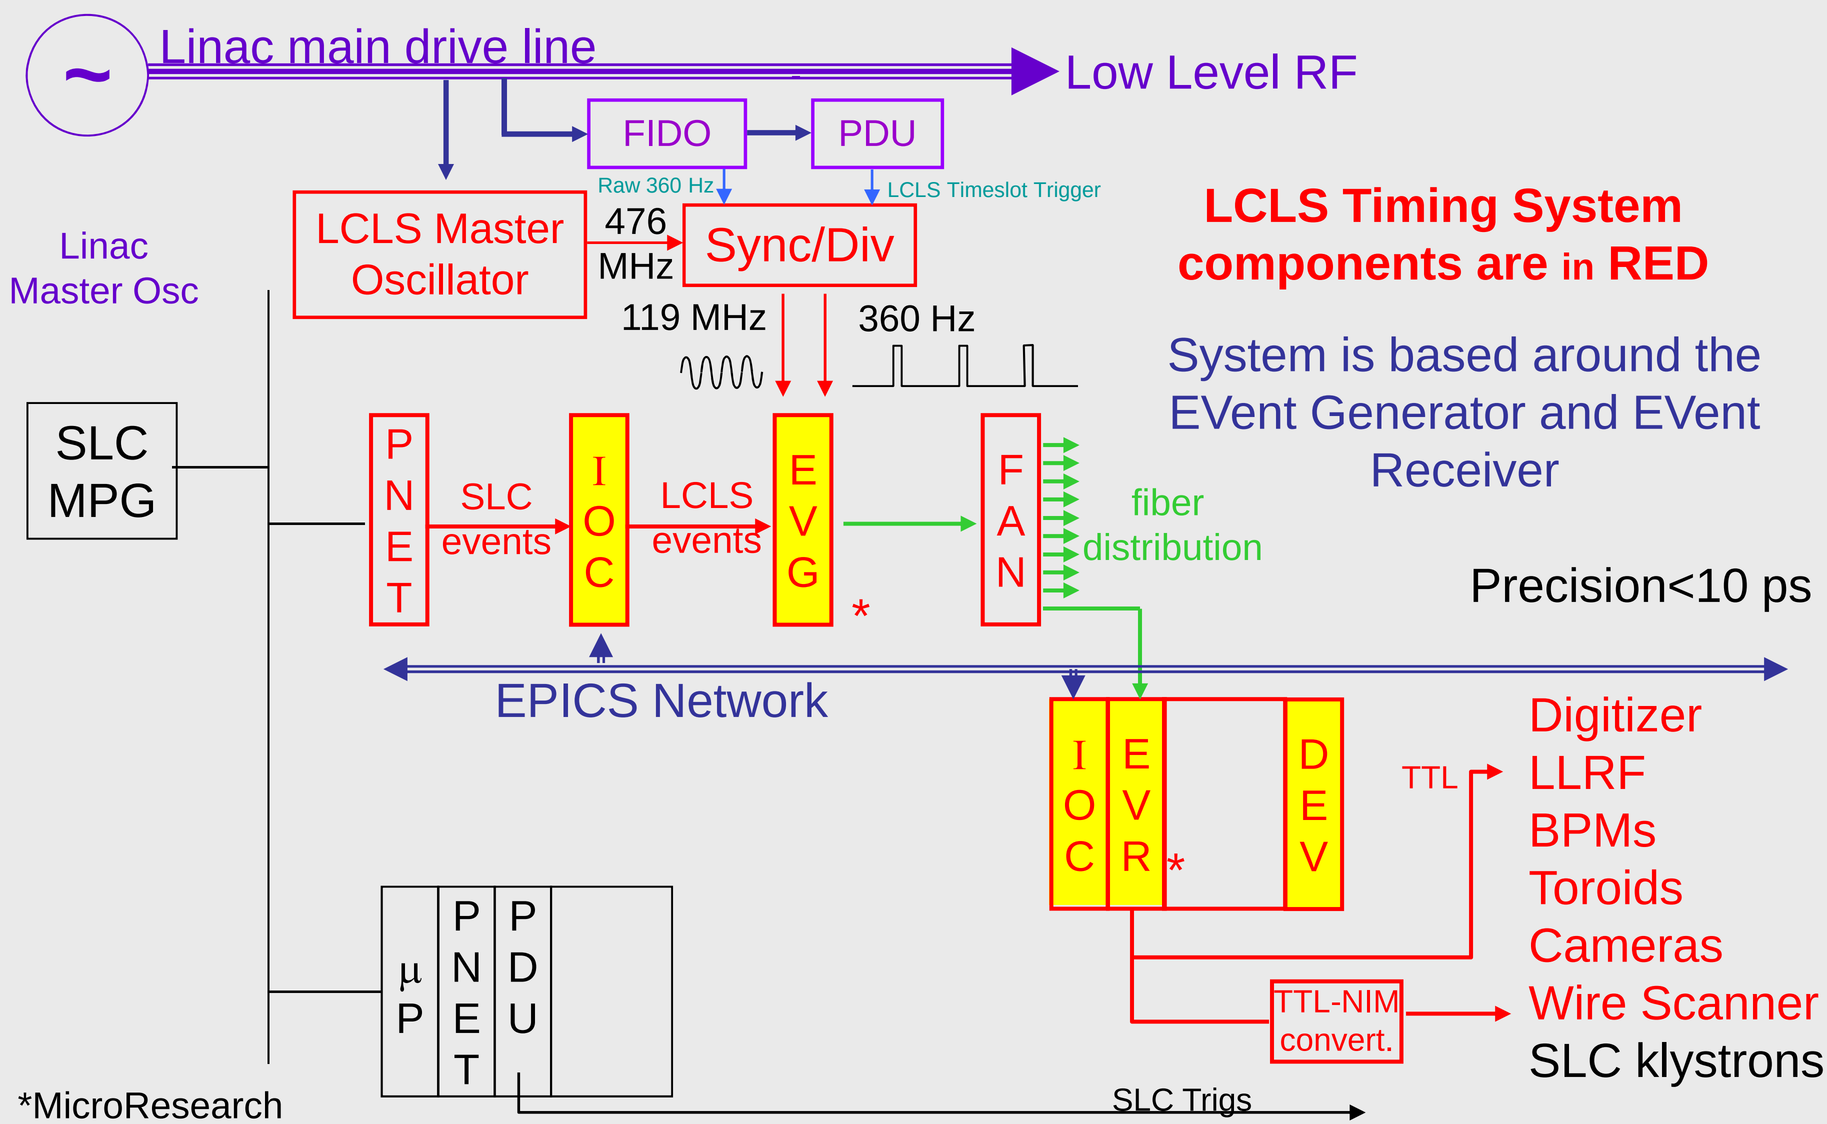

Patrick Krejcik
pkr@slac.stanford.edu

Feedback

- Many aspects of the SLS orbit feedback system will be adaptable to the requirements for a pulse-by-pulse “feedforward” system for the FEL.
- Beam-based measurement of energy and bunch length is the final arbiter in correcting for drifts in the LLRF reference system.

Patrick Krejcik
pkr@slac.stanford.edu

Laser & Electron-Based Feedback Systems

Transverse Loops:

- Laser spot on cathode
- Gun launch angle
- Injector trajectory
- X-band cavity position
- Linac trajectory
- Undulator launch

Longitudinal Loops:

- DL1 energy
- BC1 energy
- BC1 bunch length
- BC2 energy
- BC2 bunch length
- Final energy

- Bunch Charge

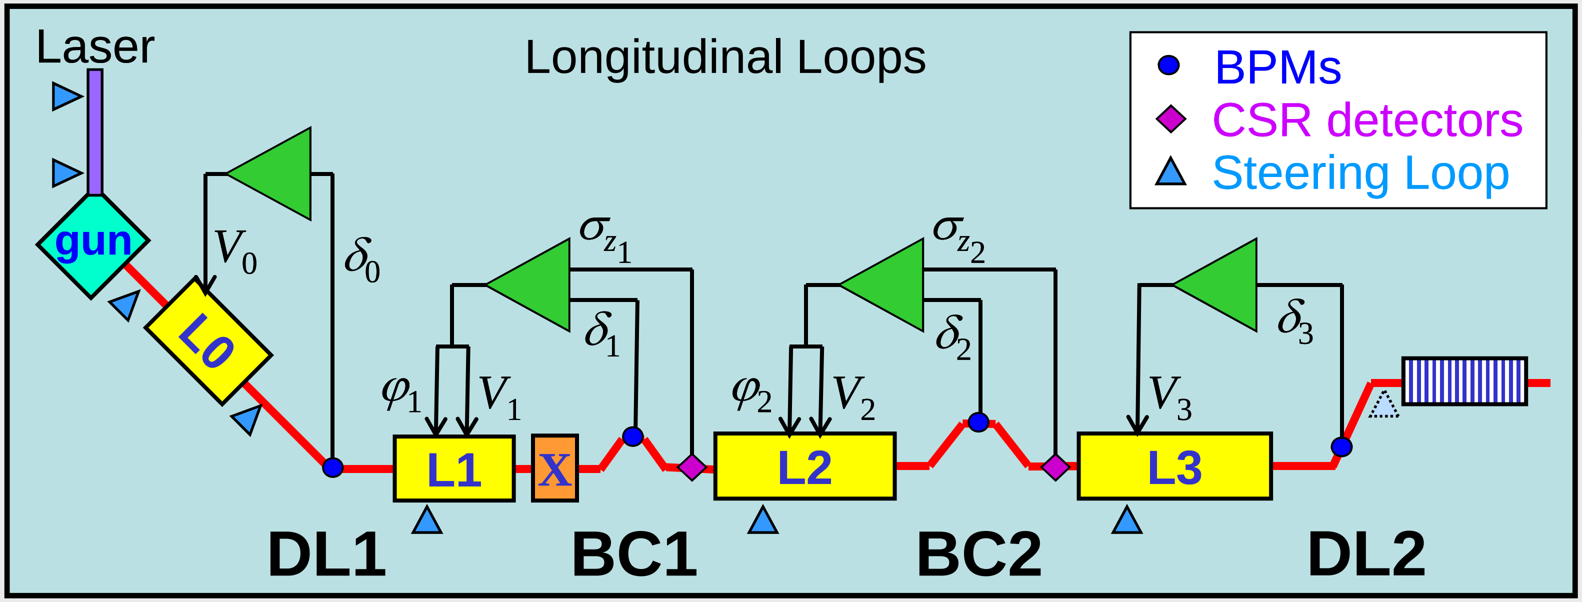

Patrick Krejcik
pkr@slac.stanford.edu

Fast Feedback Conceptual Block Diagram

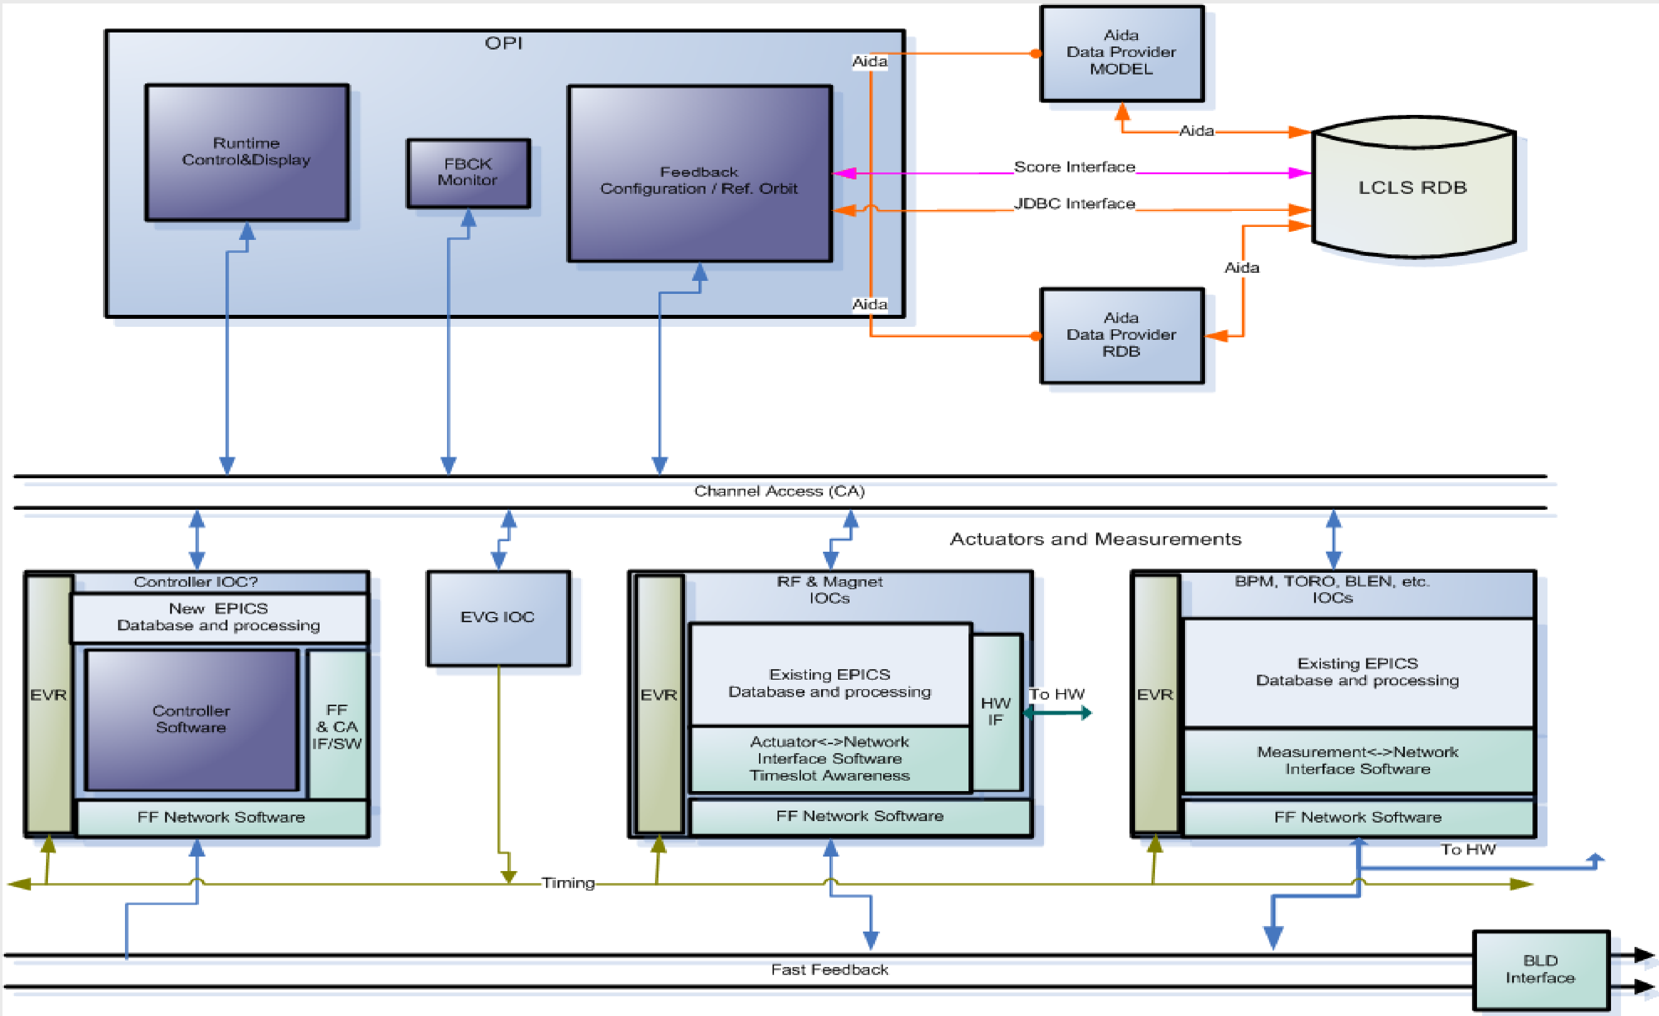

Patrick Krejcik
pkr@slac.stanford.edu

LLRF and synchronization

- The tightest phase and amplitude stability requirement comes from the x-band system
  - the stability of the LLRF distribution system will be dictated by this.
- A reasonable analysis of data rates was given for the number of LLR waveforms to be digitized for every pulse.
  - Some thought should still be given as to how this data may be archived, either through decimation or averaging.

Undulator control

- lack of support buildings directly above the undulators means long cable runs will have to be supported.
- some electronics being located in the tunnel is inevitable (BPMS, motion controllers).
- Variations in heat load need to be calculated for temperature stability of the undulator.
  - Temperature stability tolerances should be clearly defined.
- Experience with SLS insertion device controls will be of great benefit.

Patrick Krejcik
pkr@slac.stanford.edu

LCLS Undulator Temperature History

(stable to **0.05 degC** over 24 hrs)

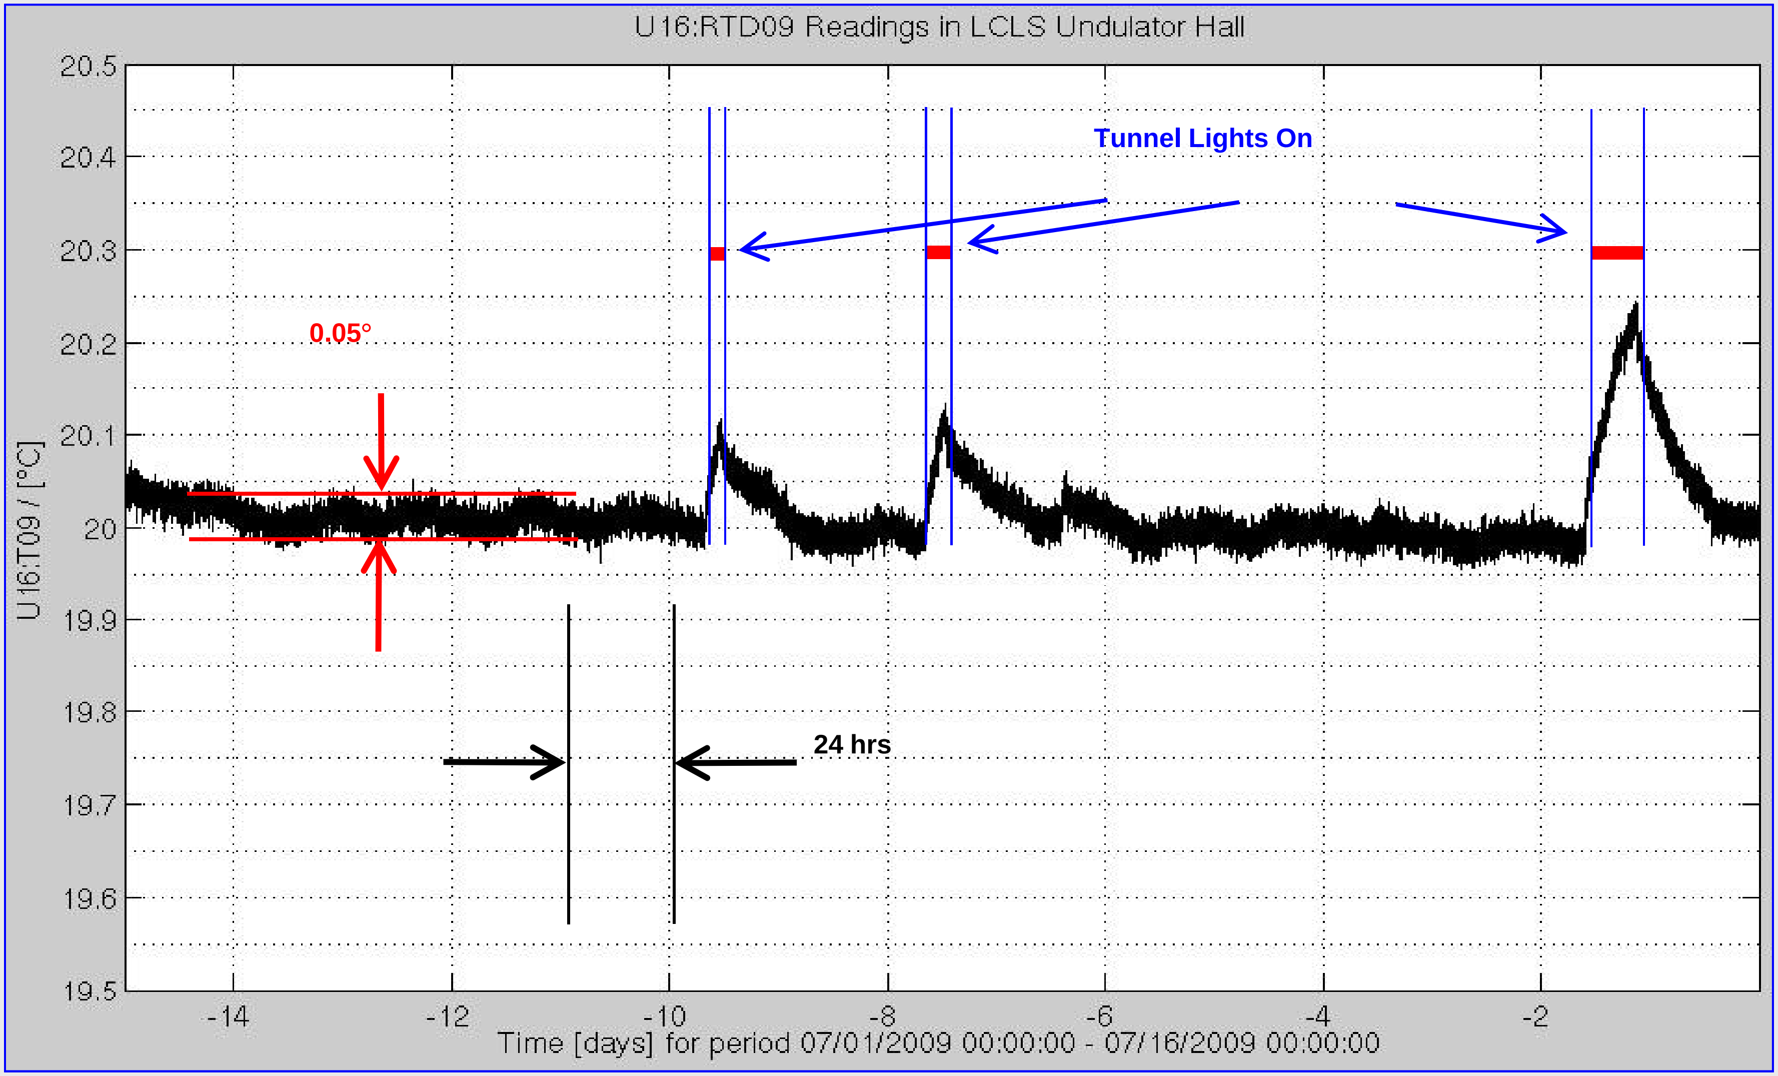

Patrick Krejcik
pkr@slac.stanford.edu

Undulator Motion Control

- Users want sophisticated GUIs for undulator motion control
- At LCLS control physics parameters like taper and pointing

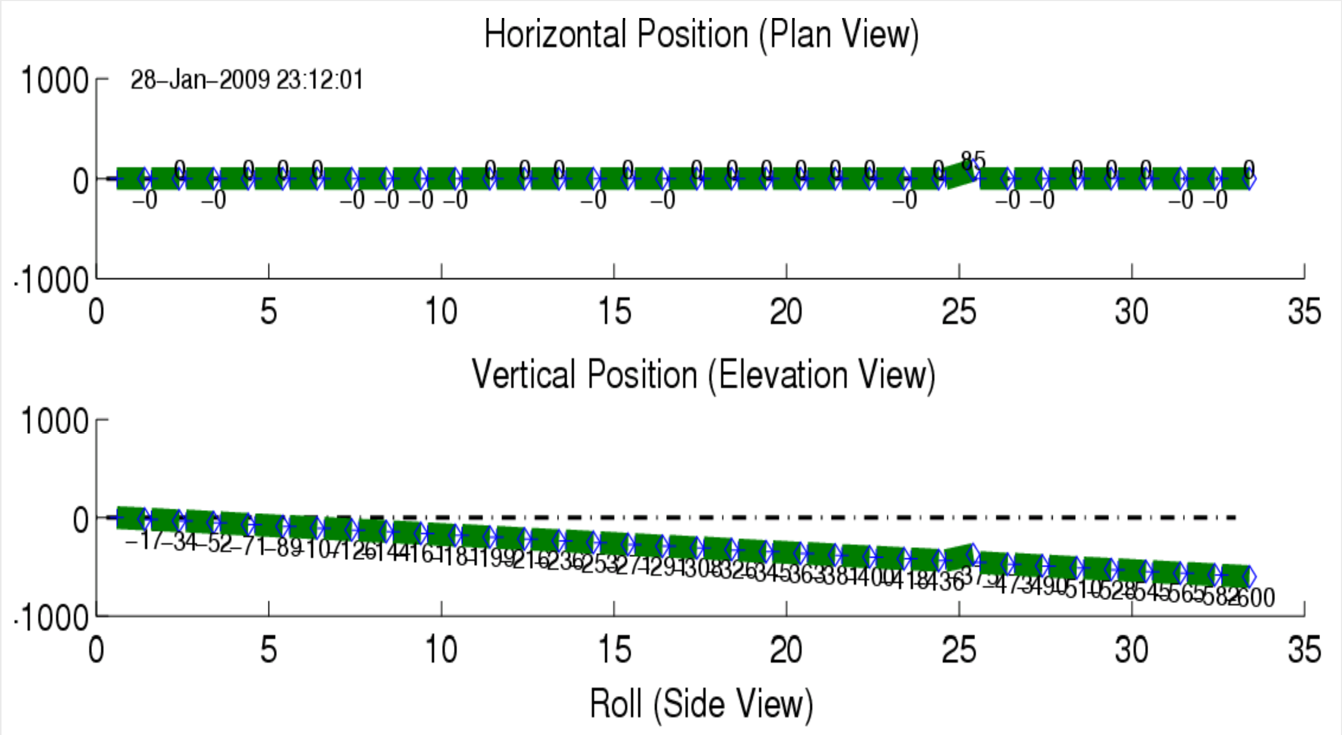

Patrick Krejcik
pkr@slac.stanford.edu

on undulator control:

- gap/shift will be calculated by a semi analytical model
  - planned Siemens PLCs is not for calculations
- look into other, more powerful motor controllers, such the Delta Tau as used by Diamond.
  - provide inverse kinematics, RT synched motion
  - might help to reduce the CS from two full cabinets to one.

Mark Plesko

DAQ

- image acquisition requirements for x-ray pixel detectors thoroughly researched.
  - how will detectors be integrated with the accelerator control system,
  - electron and photon data must be synchronously acquired and analyzed.
- implementation of a high speed or synchronized link between the accelerator control system and the user stations.
- At LCLS "Beam Line Data" is computed on a pulse-by-pulse basis by the steering, energy and bunch length feedbacks and is streamed to the photon users so that they may tag each shot.

Patrick Krejcik
pkr@slac.stanford.edu

To be considered still …

- Machine Protection Systems
- management of interlocks by users.
  - who has the rights and responsibility to operate a certain device, who has to give permission, who can hold the interlock, etc.
- data analysis infrastructure
  - since users don’t have infrastructure at their home institutes.
  - Needs for long term, large volume data storage should also be addressed.

Patrick Krejcik
pkr@slac.stanford.edu

LCLS Machine Protection System

- protecting undulator and beamlines
- Fast beam shutoff at photoinjector & kicker

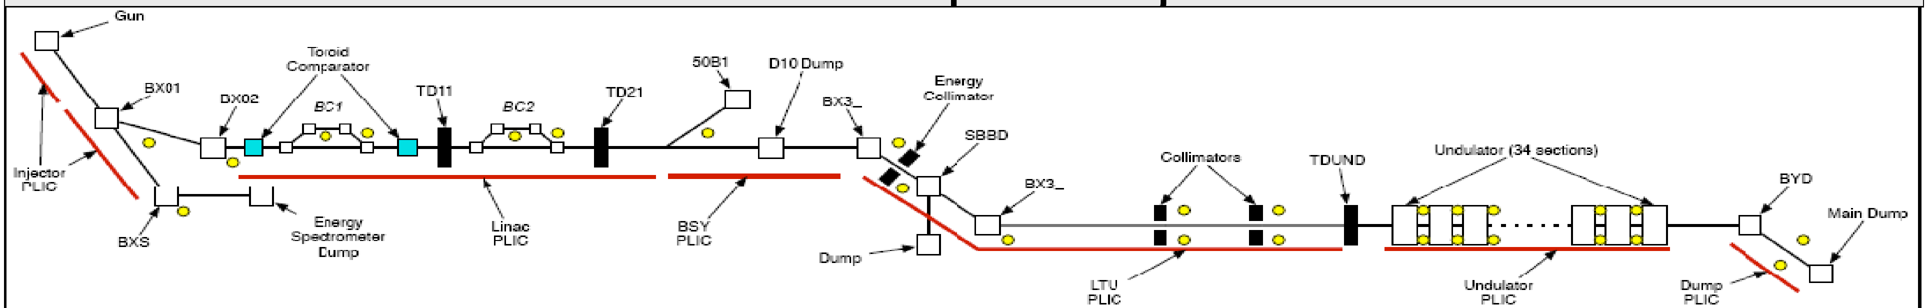

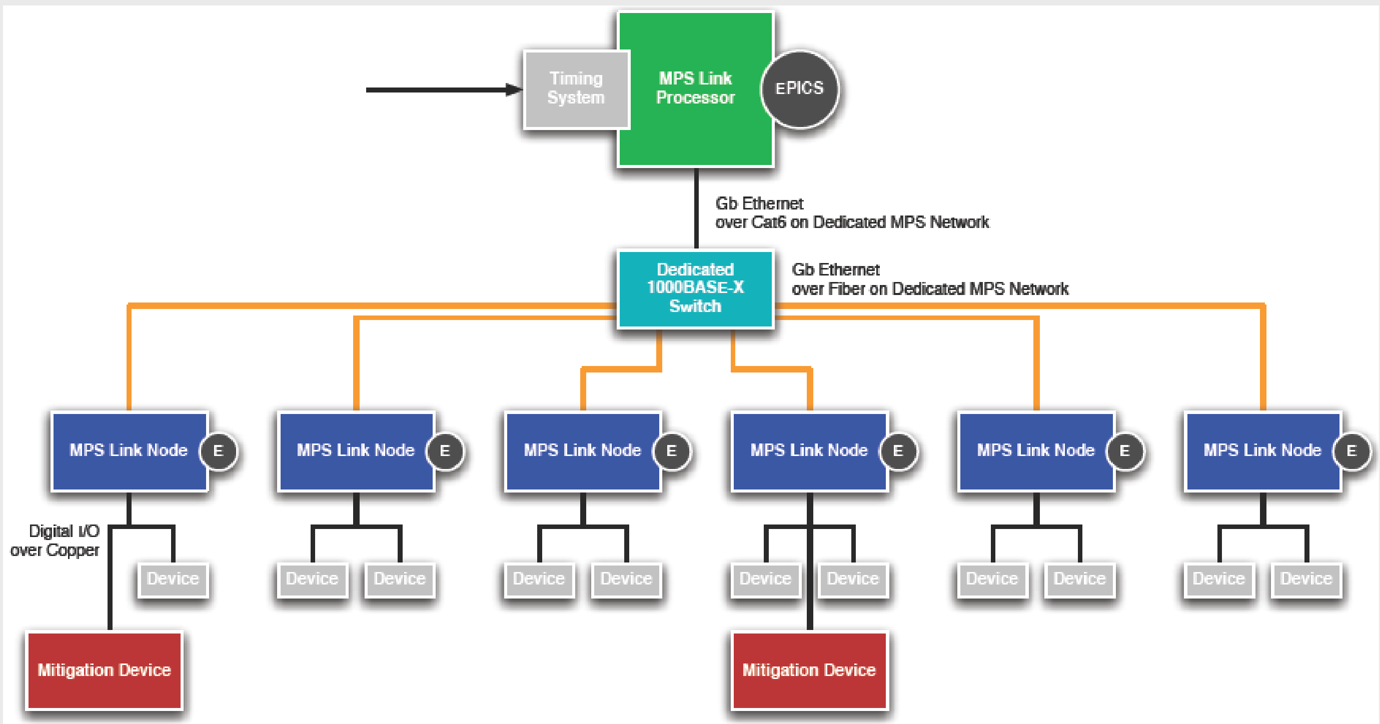

Patrick Krejcik
pkr@slac.stanford.edu

MPS User Interface

- Multiple views
  - Summary
  - Inputs
  - Logic
  - History

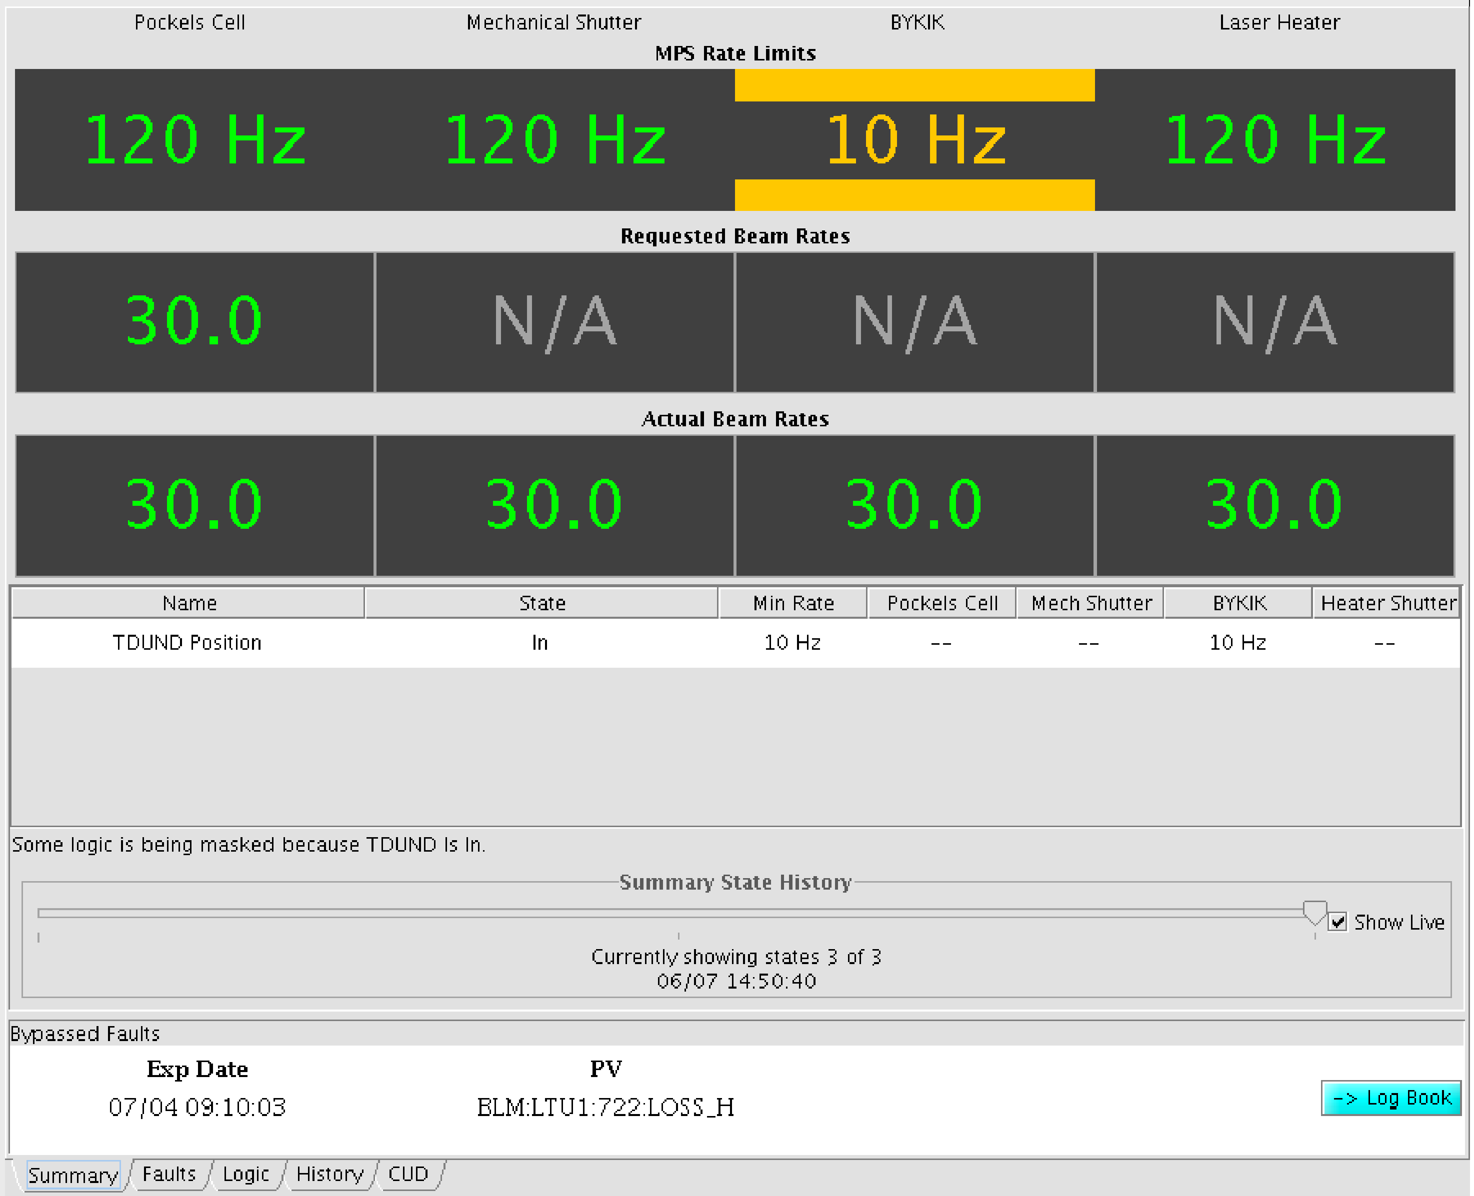

Patrick Krejcik
pkr@slac.stanford.edu

More discussions to come in the working groups

THANKS FOR INCLUDING US!

Patrick Krejcik
pkr@slac.stanford.edu