~~Progress~~ Status of Super*B* Electronics, Trigger, DAQ and Online (ETD/Online)

Steffen Luitz, SLAC 12/3/2010

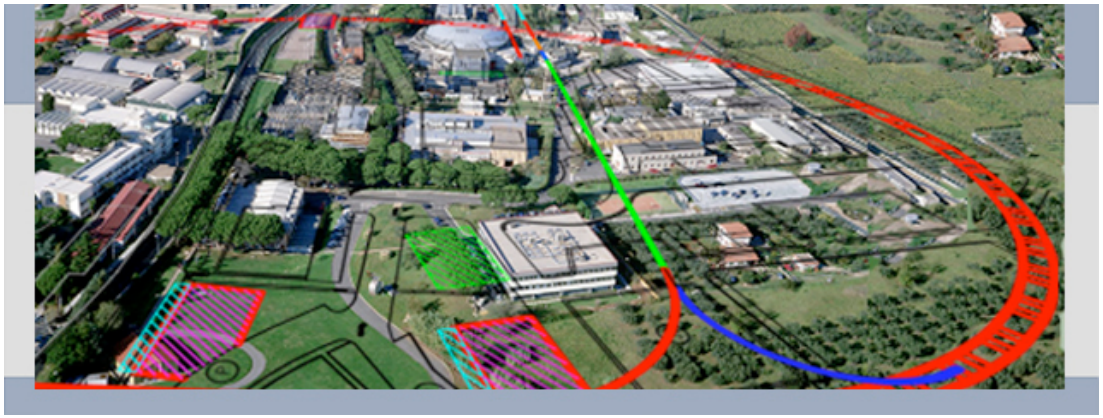

Content

- Introduction to Super*B*
- The Detector
- Rates and Event Sizes
- Design Principles
- The ETD/Online System
  - Trigger
  - Fast Control and Timing System
  - Common Front-End Electronics
  - Experiment Control System
  - ROMs
  - Event Builder
  - HLT & Logging
  - Other Systems
- Open questions
- Conclusions & Outlook

- What is Super*B*?
  - http://web.infn.it/superb
    - "International enterprise aiming at the construction of a very high luminosity asymmetric e^+e^- flavour factory"
    - "… promises to be the flagship experiment in flavor physics …"
    - Links to "progress reports"
  - To be built at or near LNF (Frascati)
    - … or elsewhere … ?
  - (Still) Waiting for approval from Italian government
  - Goal: collect at least 50-75ab^{-1} (50-100x existing B-Factory data samples)
  - To do that on a reasonable timescale (~5 years)
    - L >= $10^{36}cm^{-2}s^{-1}$
  - State-of-the art accelerator design
  - Detector based on the BaBar design
    - Reusing components from BaBar where possible

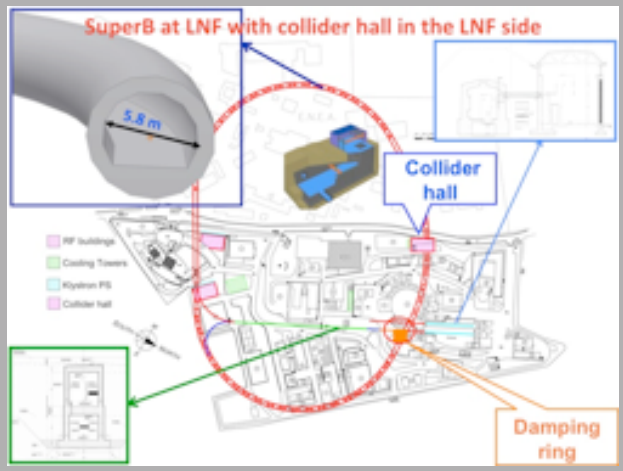

Introduction to Super*B*

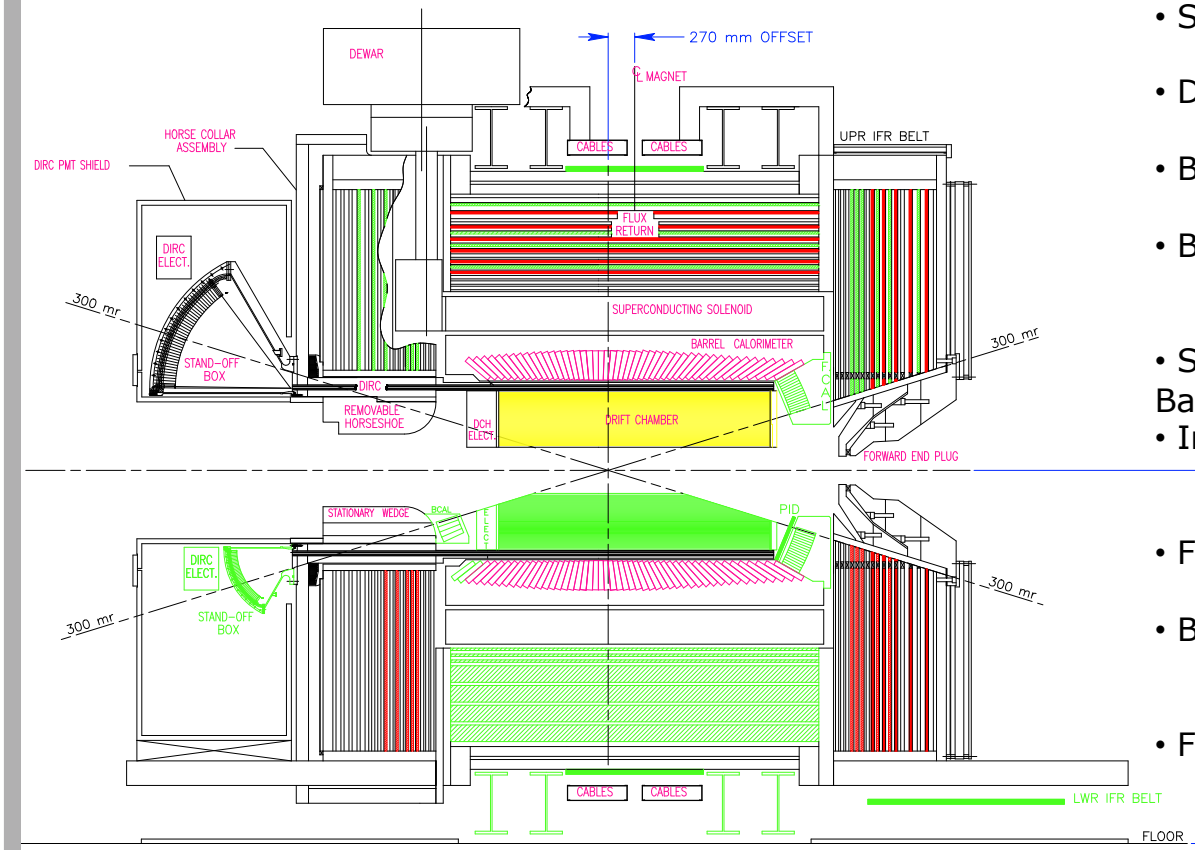

- Silicon Vertex Detector
  - Striplets? MAPS?
- Drift Chamber
  - 40 cell layers
- Barrel PID
  - FDIRC (MaPMTs)
- Barrel EMC
  - CsI (from BaBar) + PIN
  - Faster shaping
- Superconducting Magnet (from BaBar)(
- Instrumented Flux Return
  - Plastic scintillators + WLS fibres + APD pixels
- Forward EMC
  - L(Y)SO + PIN or APD
- Backward EMC
  - Lead+Scintillator+ WLS fibres + Multi-pixel devices
- Forward PID?
  - TOF? RICH?

The SuperB Detector

Expected Trigger Rates and Event Sizes

- Estimates extrapolated from BaBar for a detector with BaBar-like acceptance
- Bunch crossing instantaneous rate: 476MHz
  - At 10^{36} the average rate about half that (only half the RF buckets are filled)
- Level-1 trigger rates (scaled from BaBar)
  - At 10^{36}: 50kHz Bhabhas, 25kHz beam backgrounds, 25kHz "irreducible" (physics+backgrounds)
  - → 75kHz with a Bhabha veto at L1 rejecting 50%
  - → 100kHz without Bhabha veto
  - 50% headroom desirable (from BaBar experience)
- → baseline: 150kHz rate capability
- Event size: 75kByte (extrapolated from BaBar)
  - Raw (pre-FEX) sizes understood
  - Still some uncertainties for post-FEX/ROM sizes

Logging Rates

- From BaBar:
  - Expect do be able to achieve 25nb logging cross section with a safe real-time HLT
    - Could be improved by maybe 5-10nb with a more aggressive filter (storage & processing cost vs. risk)
  - → Have to log 25kHz of 75kByte events
  - → almost 2 Gbyte/s

ETD/Online System Design Principles

- Apply lessons learnt from BaBar & LHC experiments
- Keep it simple
  - Synchronous design, fixed-latency trigger
  - No "untriggered"/push readouts
    - Except for trigger data stream and subdetector internals
    - Minimizes the # of links
- Use off-the-shelf components where applicable
  - Links, networks, computers, other components
  - Software (look into reusing Online frameworks from other experiments)
- Modularize the design – share across the system
  - Common building blocks and modules for common functions
  - Implement subdetector-specific functions on specific modules
  - Carriers, daughter boards / mezzanines
- Design with radiation-hardness in mind
  - Where necessary
- Design for high-efficiency and high-reliability factory mode
  - Where affordable ☺ - BaBar experience will help with the tradeoffs
  - Design for minimal intrinsic dead time (goal ~1%)
  - Automation and minimal manual intervention, physical access to hardware

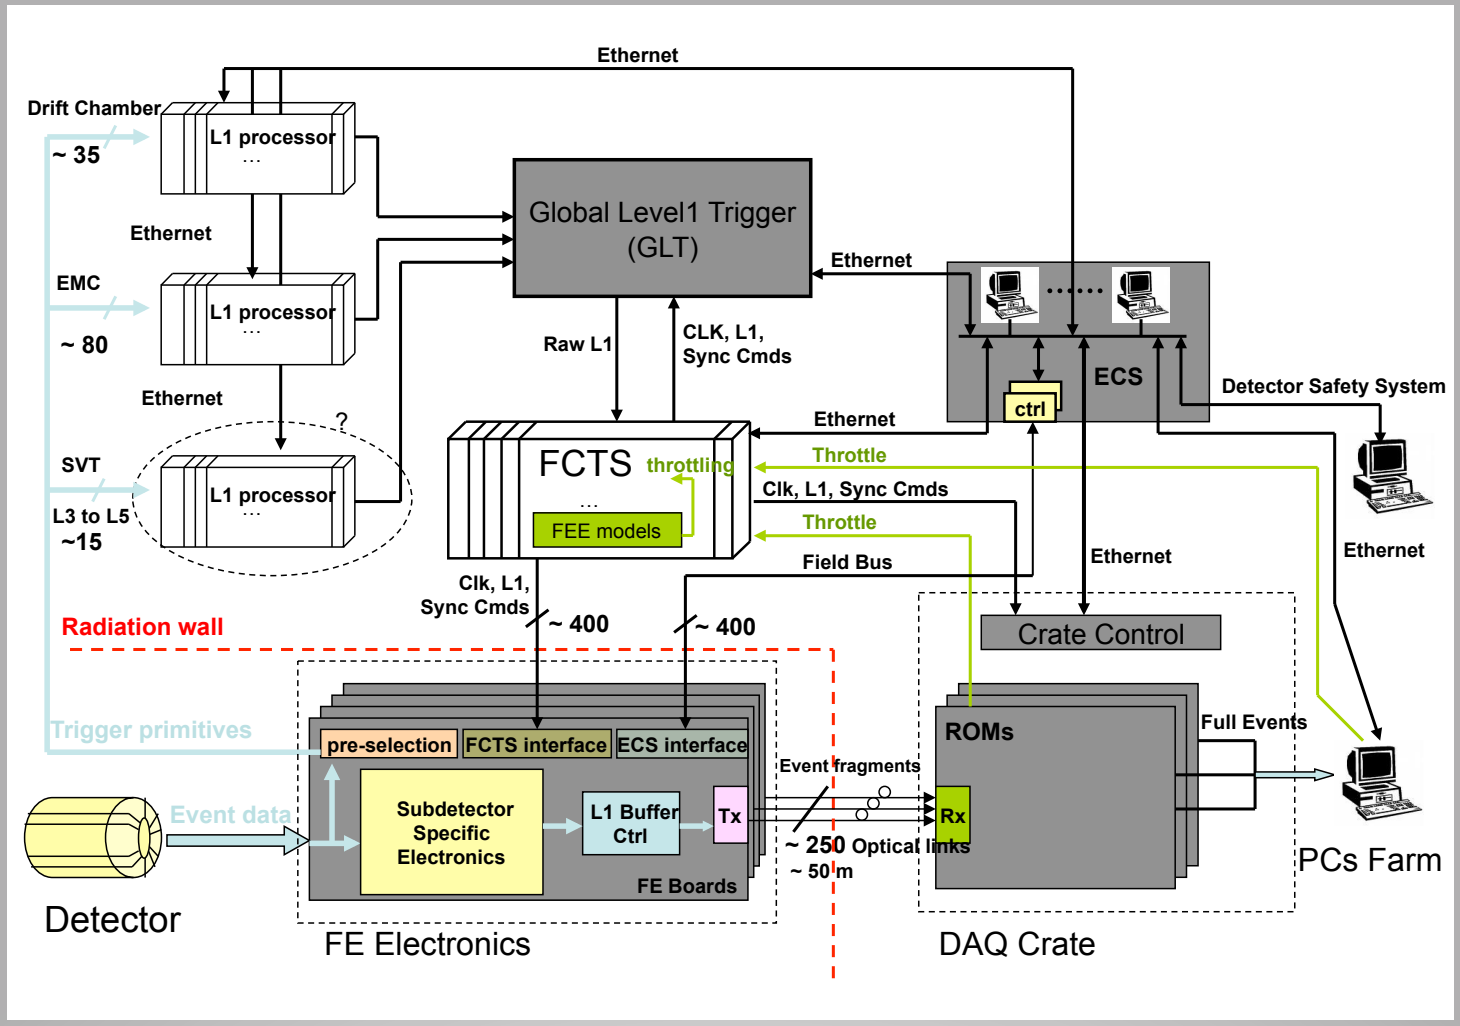

SuperB ETD System Overview

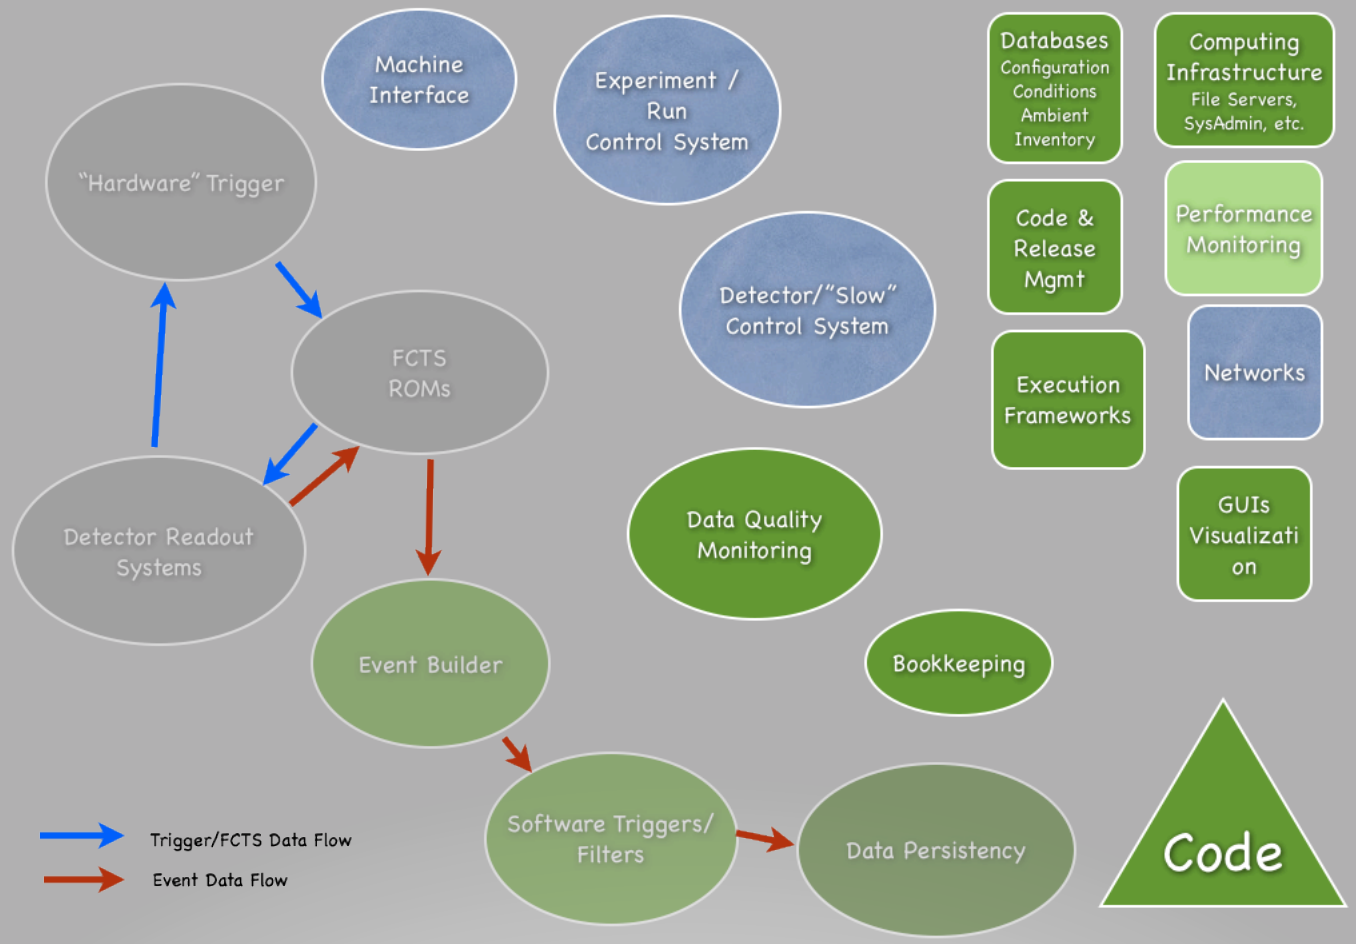

ETD/Online Components and Data Flow

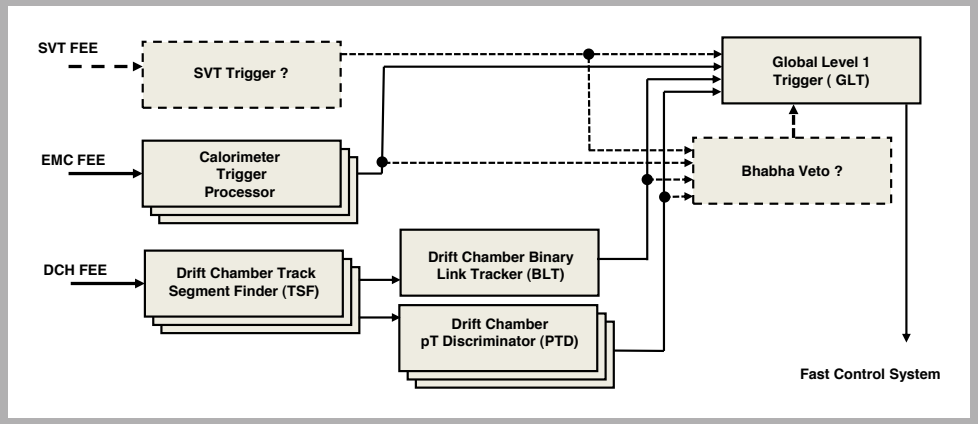

- Fully pipelined
- input running at 7MHz (maybe 14MHz?)
  - continuous reduced-data streams from sub-detectors over fixed-latency links
    - EMC crystal sums (in the FEE)
    - DCH hit patterns (in the FEE)
- output maybe 14 MHz (fine time fit)
- Total latency goal: ~5us
  - Includes trigger readout, FCTS, propagation
  - leaves about 1-2 us for the trigger itself

"BaBar-like L1 Trigger"

- Calorimeter Trigger
  - cluster counts and energy thresholds
  - but 2-d map (what granularity?)
- Drift chamber Trigger
  - Track counts, p_T, z-origin of tracks
- Highly efficient, orthogonal

To be studied:

- SVT trigger
  - # tracks, # tracks not from IP, # back-to-back tracks in phi
- Bhabha veto
  - → HLT?

Level-1 Trigger

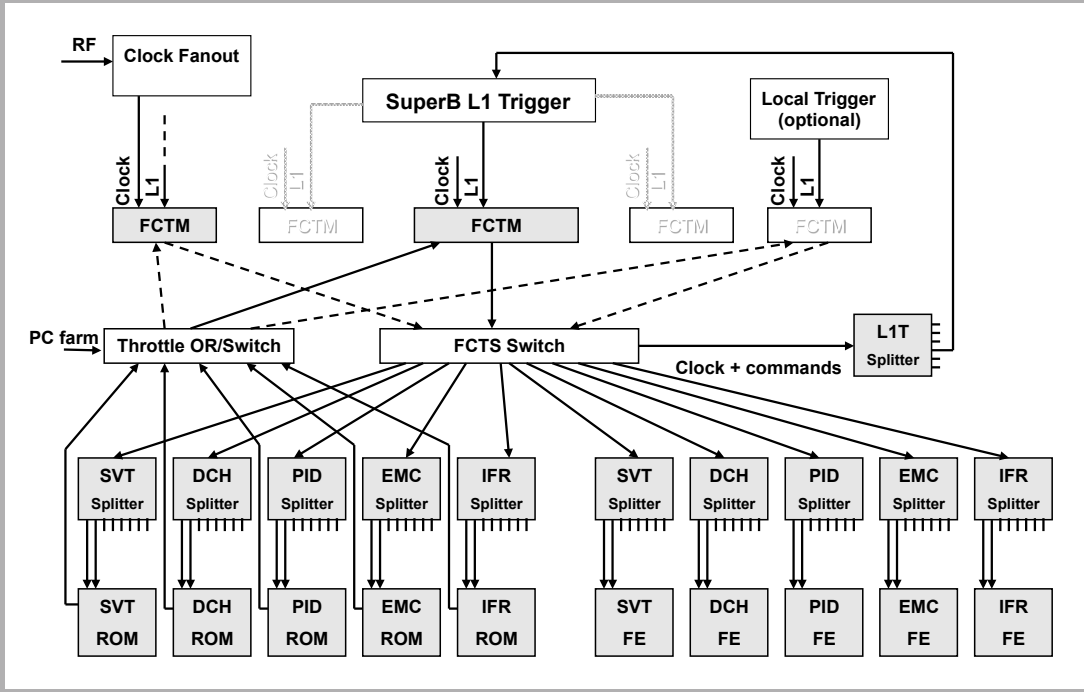

- Clock distribution
- System synchronization
- Command distribution
  - L1-Accept
- Receive trigger decisions from L1
  Participate in pile-up and overlapping event handling
- Dead time management
  - Fast throttle (emulates front-ends in FCTM)
  - Slow throttle via feedback
- System partitioning
  - 1 partition / subdet.
- Event management
  - determine event destination in farm

Links carrying clock, commands and trigger data need to be synchronous and fixed-latency - ~1 GBit/s

Readout data links can be variable latency (even packetized) ~2 GBit/s

Fast Control and Timing System

Dead Time

- Target: ~1% event loss
- Assume exponential pdf of event inter-arrival time.
- Assume continuous beams (2.1ns between bunch crossings)
- No simulation of derandomizer buffers yet
- 1% event loss due to dead corresponds to 1/150kHz -- ca. 70ns maximum per-event dead time.
- Places hard constraints on trigger, trigger output and FCTS command length!
- Challenge: Detector time resolutions achievable in a hardware trigger and per-channel dead times are also in the same order of magnitude!

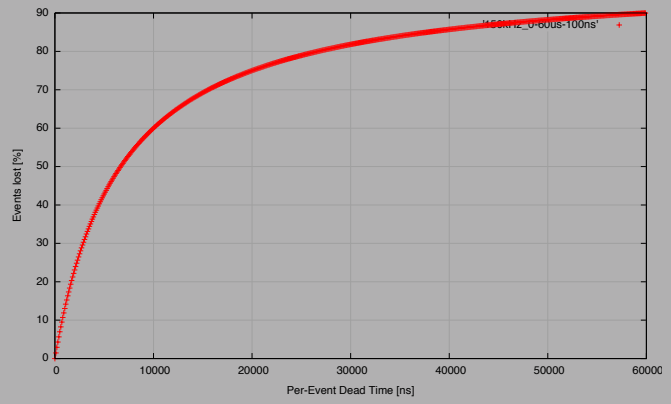

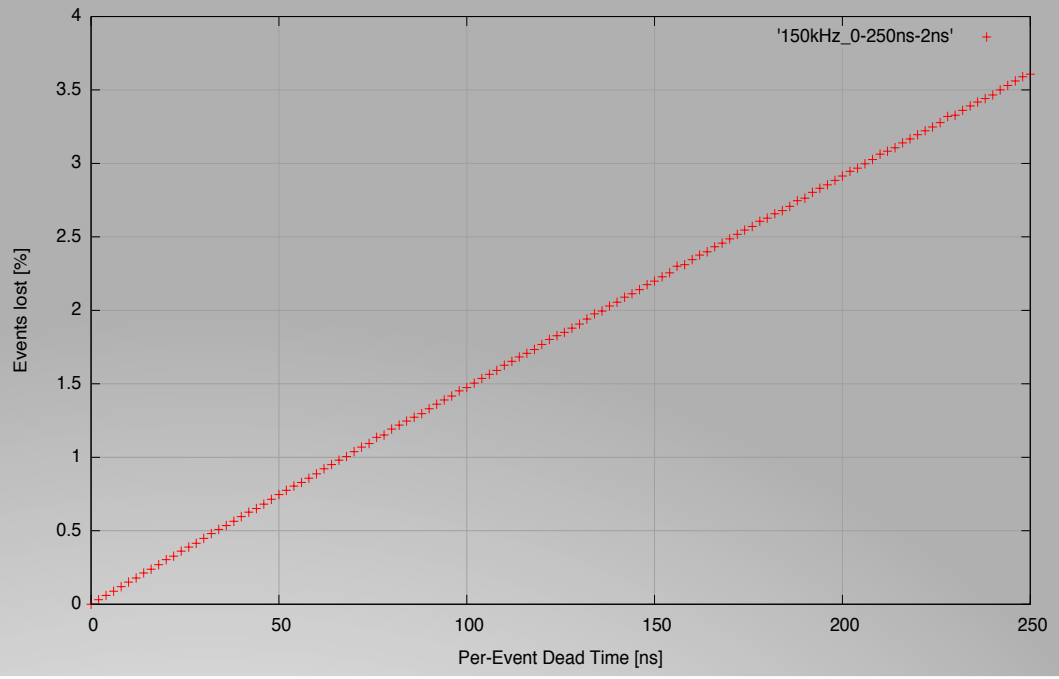

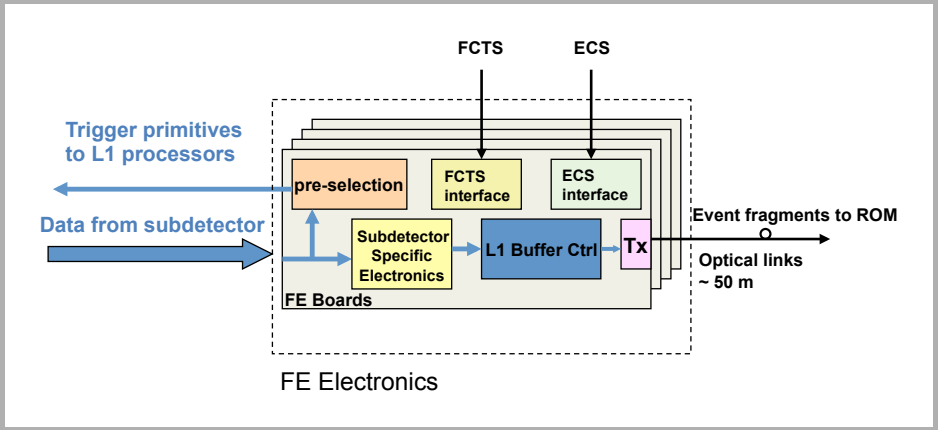

Provide standardized building blocks to all sub-detectors, such as:

- Schematics and FPGA "IP"
- Daughter boards
- Interface & protocol descriptions
- Recommendations
- Performance specifications
- Software

- Digitize
- Maintain "circular" latency buffer
- Maintain de-randomizer buffers & output mux
- Generate reduced-data streams for L1 trigger
- Interface to FCTS
  - receive clock
  - receive commands
- Interface to ECS
  - configure
  - calibration
  - spy
  - test
  - etc.

Common Front-End Electronics (FEE)

Experiment Control System

- Configure System
  - Upload configuration into FEE
  - Should be fast!
- Monitor System
  - Spy on event data
  - Monitor power supply, temperatures, etc.
- Testing the System
  - Using software specifically written for the FEE
  - We do not foresee ECS-less self-test capabilities for the FEE
- Proposed implementation
  - SPECS (Serial Protocol for Experiment Control System)
  - Bidirectional 10MBit/s bus designed for LHCb

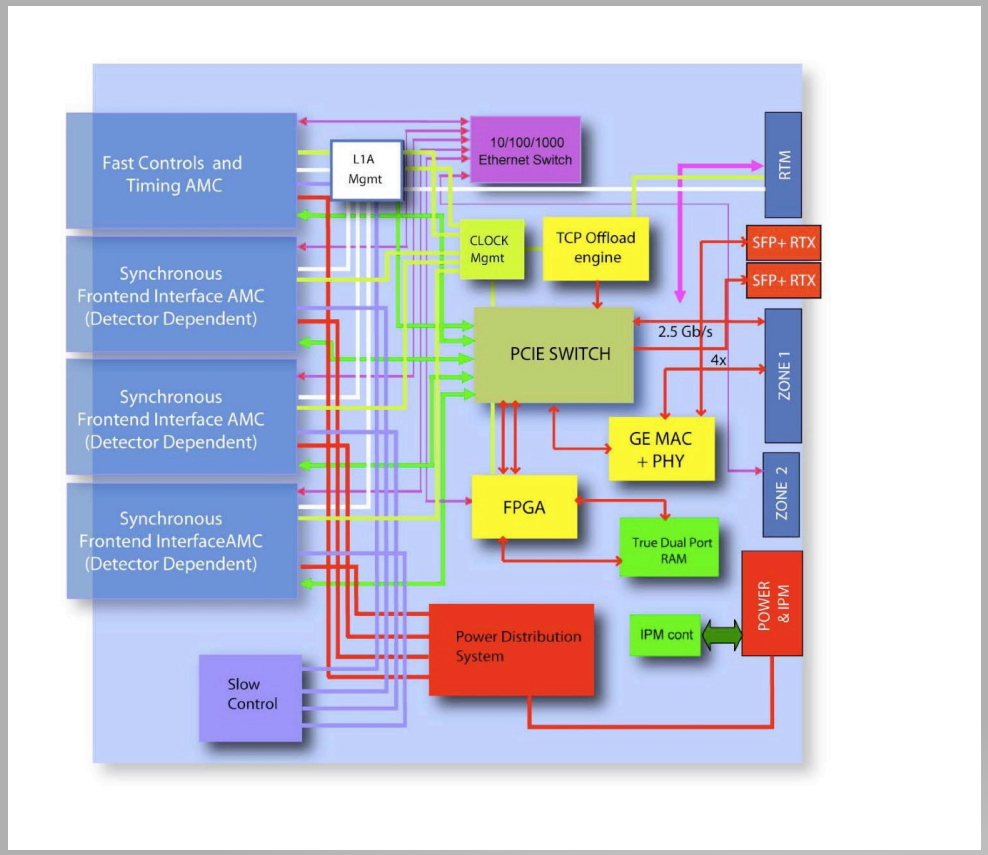

- Process and forward FCTS information to FEE, implement FEE-specific requirements
- Receive data from the sub-detectors over optical links
  - 8 links per ROM?
- Reconstitute linked/pointer events
- Process data (FEX, data reduction)
- Send event fragments into HLT farm (network)

- Would like to use off-the shelf hardware as much as possible (i.e. off-the shelf computers with OL PCIe cards?) → R&D
- Will need to determine processing requirements from sub-detectors.

Readout Modules (ROMs)

Event Builder and Network

- Combines event fragments from ROMs into complete events in the HLT farm
  - In principle a solved problem ☺
  - Prefer the fragment routing to be determined by FCTS
  - FCTS decides to which HLT node all fragments of a given events are sent (enforces global synchronization), distribute as node number via FCTS
  - Choice of network technology
    - Combination of10Gbit/s and 1GBit/s Ethernet prime candidate
    - UDP vs. TCP … a long contentious issue?
      - Pros and cons to both
    - Can we use DCB/DCE for layer-2 end-to-end flow control / reliable layer-2 networking in the EB network?
  - Design choices for protocol and network / node congestion control
    - My personal preferences:
      - Connection-less UDP-based event builder (events are small)
      - Event-to-event destination decisions taken by FCTS firmware (using a table of node #s)
      - Node availability slow-feedback protocol (over network in software)
      - Use DCB/DCE layer-2 reliability to avoid packet loss in network
  - Can we re-use some other experiment's event builder?
    - Interaction with protocol choices

HLT Farm and Logging

- Standard off-the shelf rack-mount servers
- Receivers in the network event builder
  - Receive event fragments from ROMs, build complete events
- HLT trigger (L3)
  - 10ms/event (baseline assumption, almost 10x BaBar)
  - → 1500 cores needed (~150 servers 16 cores, 10 usable for HLT purposes)
    - We now already have 24-core servers! → farm will shrink ☺
- Data logging & buffering
  - Few TByte/node
  - Local disk (e.g. RAID1 as in BaBar)?
  - Storage servers accessed via a back-end network?
  - Probably 2 day's worth of local storage (2TByte/node)? Depends on SLD/SLA for data archive facility.
  - No file aggregation into "runs" → bookkeeping
  - Back-end network to archive facility

Data Quality Monitoring

- Envision same concept as in BaBar:
  - Collect histograms from HLT
  - Collect data from ETD monitoring
  - Run fast and/or full reconstruction on sub-sample of events, collect histograms.
    - May include specialized reconstruction for e.g. beam spot position monitoring
  - Could run on same machines as HLT processes (in VMs?) or a separate small farm ("event server clients")
  - Present to operator via GUI
  - Automated histogram comparison with reference histograms and alerting

Control Systems

- Run Control
  - Coherent management of the ETD and Online systems
    - User interface, managing system-wide configuration, reporting, error handling, start and stop data taking
- Detector Control / Slow Control
  - Monitor and control detector and detector environment
- No effort has gone into even a high-level design of these systems, but we assume that we can use/re-use LHC experiment and commercial technology

Auxiliary Systems

- Electronic Logbook
  - Web based – integrated with bookkeeping
- Databases
  - Configuration, Conditions, Ambient
- Configuration Management
  - Authoritative source of configuration
  - Log trail of configuration
  - "Provenance light"
- Software Release Management

- "ETBD" (eventually to be designed)

Open Questions / R&D Topics

- Upgrade paths to a luminosity of $4x10^{36}$
  - What do design upfront, what do upgrade later, what is the cost?
  - Accelerator plans already existing
- Data links
  - Jitter, clock recovery, coding patterns for error detection, radiation qualification, performance of embedded SERDES
- ROM
  - 10GBit/s networking technology, I/O sub-system, using a COTS motherboard as carrier with links on PCIe cards, FEX & processing in software or FPGA
- Trigger
  - Latency, physics performance, details of event handling, time resolution and intrinsic dead time, L1 Bhabha veto, use of SVT in trigger, HLT trigger, additional L4 filter, safety vs. logging rate
- ETD performance and dead time
  - Trigger distribution through FCTS (length of commands), intrinsic dead time, pile-up handling/overlapping events, depth of de-randomizer buffers
- Event builder
  - Anything re-usable out there?, network and network protocols, UDP vs. TCP, applicability of emerging standards and protocols (e.g. DCB, Cisco DCE), HLT framework vs. Offline framework (any common grounds?)
- Software Infrastructure
  - Sharing with Offline, reliability engineering and tradeoffs, configuration management ("provenance light"), efficient use of multi-core CPUs

Outlook & Conclusion

- There has been significant progress on the SuperB ETD/Online system design in the last two years (unfortunately less in the last few months ☹)
- We are at the end of what we can do without positive news towards approval
- There is a lot to be done – help is very welcome!
- Hoping for a positive decision within the next two weeks
  - If so, I'd be happy to come back here and talk about the larger scale computing plans