Thoughts on Si-W ECAL DAQ

Paul Dauncey

For the Imperial/UCL groups

D. Bowerman, J. Butterworth, P. Dauncey,
M. Postranecky, M.Warren, M. Wing, O. Zorba

How to get from here to there…

- Where "there" is ~ 2009
  - Detector TDR should be written
  - Prototyping of final ECAL system 2010-2011
  - Production of final ECAL system 2012-2015
- Five years to define what the ECAL looks like
  - Which detectors to use
  - Readout system architecture
- UK groups beginning to look at DAQ for ECAL
  - Within context of CALICE, so concentrate on Si-W ECAL (at least for now)
  - Try to identify if any R&D "bottlenecks" which need work
  - Order of magnitude values only here

Generic Si-W ECAL

- TESLA ECAL design
  - Si-W sampling calorimeter
  - 8-fold symmetry
  - 40 silicon layers
- Any LC ECAL will probably be
  - Radius ~2m
  - Length ~5m
  - Thickness ~20cm
- Sensitive detector of silicon wafers
  - $1 \times 1 cm^2$ p-n diode pads
  - Circumference ~12m ~1000 pads
  - Length ~5m ~500 pads
  - Number of layers ~30 (?)
  - Barrel total ~1000×500×30 ~ 15M pads
  - Including endcaps, total ~ 20M pads

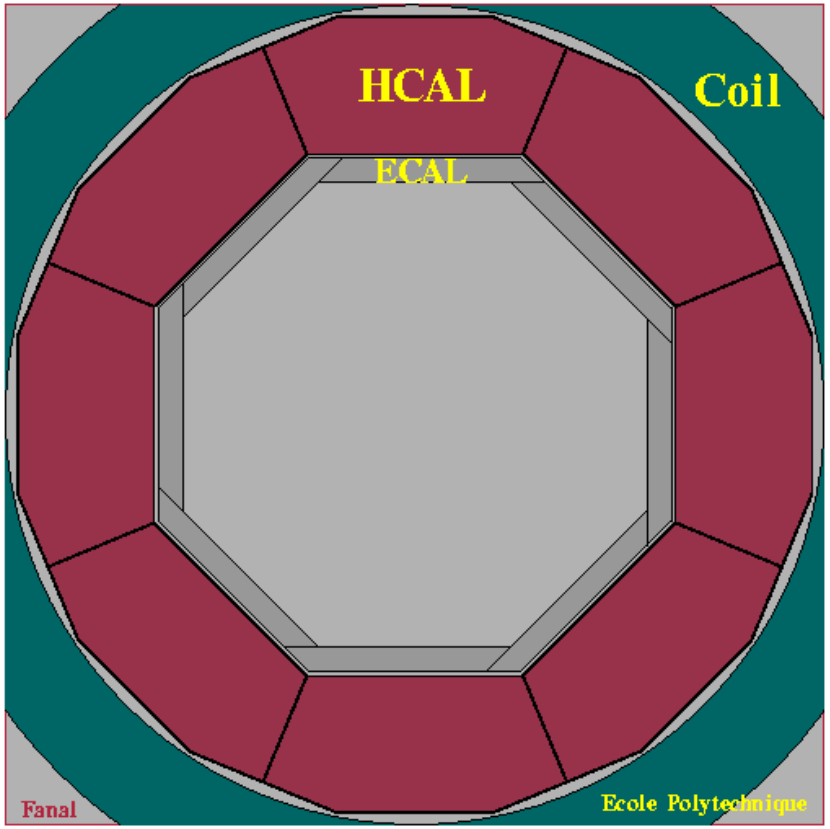

Generic Si-W ECAL (cont)

- Design likely to be modular
  - Mechanically and electrically
  - Assume this will be true here also
- TESLA ECAL made from slabs
  - Inserted into alveoli of carbon fibre/tungsten mechanical structure
  - Each slab was two-sided, 16 wafers long, each wafer having 10×10 pads
  - Total number of pads per slab was 3200
  - Readout all along one (short) edge
- Generic ECAL; assume modularity will be similar
  - Could be tower of several layers, etc.
  - Each slab ~ 4k pads
  - Total detector ~ 5k slabs

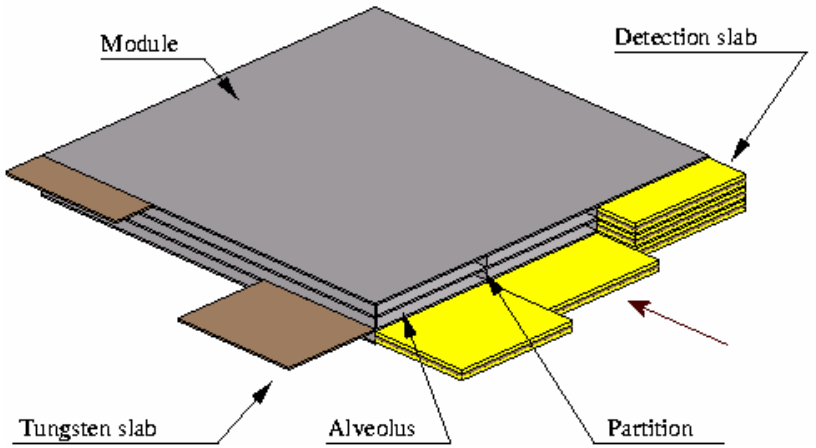

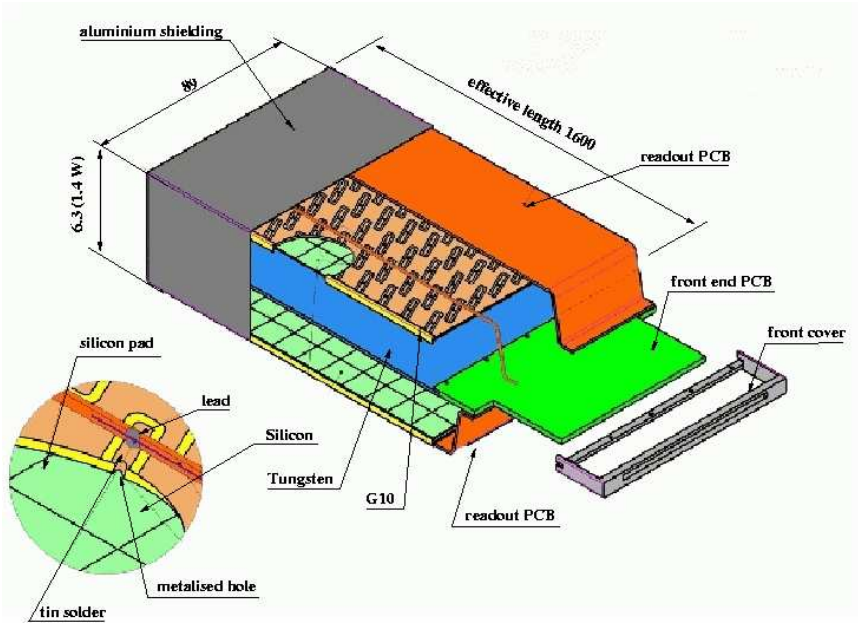

Slab electrical structure

- Now likely there will be ASICs mounted on/near wafers
  - Preamplification and shaping
  - Possibly digitisation and threshold suppression
  - ASIC will handle something like 128 channels ~ 100
  - Need ~ 40 ASICs per slab, ~ 200k total for ECAL
- Very restricted space between layers
  - Need small gap to take advantage of tungsten Moliere radius
  - Cooling difficult to get in; must minimise ASIC power consumption
  - Power down electronics between bunch trains?
- All communication to outside via edge electronics
  - Integration still difficult but assume sufficient cooling possible

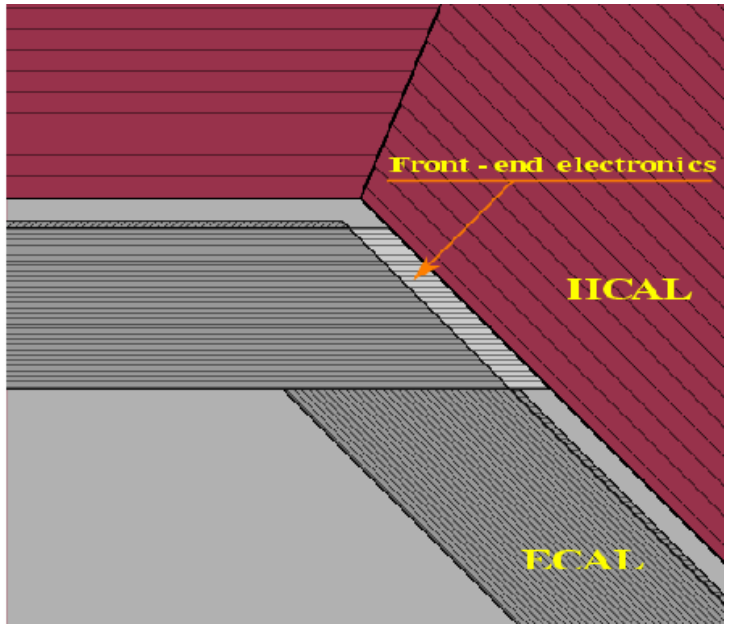

Assume something like 800 GeV TESLA

- Parameters were
  - Bunch crossing period within bunch train = 176ns ~ 200ns
  - Number of crossings per bunch train = 4886 ~ 5000
  - Bunch train length = 860μs ~ 1ms
  - Bunch train period =250ms ~ 200ms
  - Results in front end electronics duty factor ~ 0.5%
- Also assume ECAL needs to
  - Digitise signal every bunch crossing
  - Readout completely before next bunch train
- However, a need for cosmics may change this
  - See later…

Fast control and timing

- Beam crossing time is slow ~5MHz
  - Will need some finer adjustment within this period for sample-and-hold, etc.
  - Will need faster clock for FPGAs, ASICs, etc.
- Assume order of magnitude sufficient, i.e. ×8 of clock ~ 40MHz
  - Straightforward to distribute by fibre to slab
  - Count multiples-of-2 on this fast clock to recover beam crossing times
  - Need synchronisation signal to start multiple-of-2 counters
  - Need programmable offset to adjust phase of multiple-of-2 counters
- Configuration would be done via fibre also
  - Load constants (pedestals, thresholds, etc)
  - Load firmware? This would need to be absolutely failsafe in terms of recovery from errors; the slabs will be inaccessible for years
- Can receivers on slab be made generic for whole detector?
  - Counters plus configuration data protocol needed by all systems

Data volumes and rates

- Each silicon pad needs ADC conversion
  - Basic measurement is number of particles which went through pad
  - Usual measure is in units of Minimum Ionising Particle (MIP) deposit
  - Highest density expected in EM shower; up to 100 particles/mm^2
  - In 1×1cm^2 pad, expect up to 10k MIPs; sets dynamic range needed = 14 bits
- Assume 2 bytes per pad per sample
  - Raw data per bunch train ~ 20M × 5000 × 2 ~ 200GBytes ECAL total, per slab ~ 40MBytes, per ASIC ~ 1MByte
  - This appears within ~1 ms; per ASIC ~ 1GByte/s
- Clearly do not want to ship out 200GBytes per bunch train
  - Threshold around 2σ gives ~ 1% output due to noise
  - Results in ~ 1B samples above threshold per bunch train
  - Completely dominates physics rate (TESLA TDR ~ 90MBytes per train)
  - Each sample above threshold needs channel and timing label ~ 4 bytes
  - Assume ~ 5GBytes ECAL total, per slab ~ 1MBytes

Should suppression be done in ASIC?

- If no ADC or suppression in ASIC
  - Need to transmit 14-bit precision signal several meters
  - Need to do ~ 4000×5000 ~ 20M conversions at slab end in ~ 1ms; e.g. with 10 ADCs, need 2GHz, 14-bit ADCs…
  - …or have analogue buffer to store all 20M analogue values and use slower ADCs
- Analogue suppression in ASIC without ADC?
  - Severe problems with setting and monitoring threshold level, as well as its stability
  - Assume ASIC has ADC; power is "only" downside
- If pads suppressed within the ASIC
  - Need buffering to smooth out gaps or require ~ full-speed bandwidth
  - Need configuration of threshold values, etc; must be non-volatile to not lose values during ASIC power cycling between bunch trains
  - ASIC significantly more complex (and higher power?)

Unsuppressed ASIC data

- If not suppressed within ASIC
  - ASIC simplified; sends every sample immediately after digitisation
  - All decisions on threshold, etc, made at slab end
  - Allow histogramming, pedestal/noise monitoring, deconvolution, etc, in FPGAs at slab end
  - But need ~ 1GByte/s rate; i.e. fibre
- Use of fibres for chip-to-chip data transfer is not (yet) standard
  - How to connect fibre to ASIC? Part of ASIC itself?
  - How to drive fibre? Too much power within slab?
  - Mechanical size of laser driver? Too large to fit within layer gap?
  - Use modulation with laser mounted at end of slab?

Data transfer from slab

- If no cosmics required, rate from slab is slow
  - ~ 1MBytes per bunch train read out within ~ 200ms means ~ 5MBytes/s
- Would probably use fibre anyway
  - Smaller, electrically isolated
  - Although wireless has been mentioned
- Issue is what is at the other end of the fibre?
  - Custom receiver or generic network switch?
- Need to get data for a bunch train from all slabs to same location
  - For trigger and reconstruction
  - Generically called Event Builders here
  - Assume ~ 1000 PCs (from TESLA TDR)

Network switch

- Run TCP/IP directly in FPGA on slab end
  - Allows data out to go directly to network switch
  - Hence goes to event builders without any further layer
- Issue is network/switch performance
  - Unlikely to have major buffering capability
  - Would need to get all ~ 5GBytes of data to event builders before next bunch
  - Network has to handle ~ 25GBytes/s continuously from ECAL alone
  - Single event builder input ~ 25GBytes/s for 200ms from ECAL alone
  - But then no data for ~ 200 seconds
- Alternative is to buffer data from multiple bunch trains at slab end
  - Drip feed data from ~ 1000 trains simultaneously into the network
  - Memory needed on slabs ~ 1000×1MByte ~ 1GByte on each slab
- Need to do some work on how this compares to LHC
  - We (UK) have little expertise in this area

Network switch (cont)

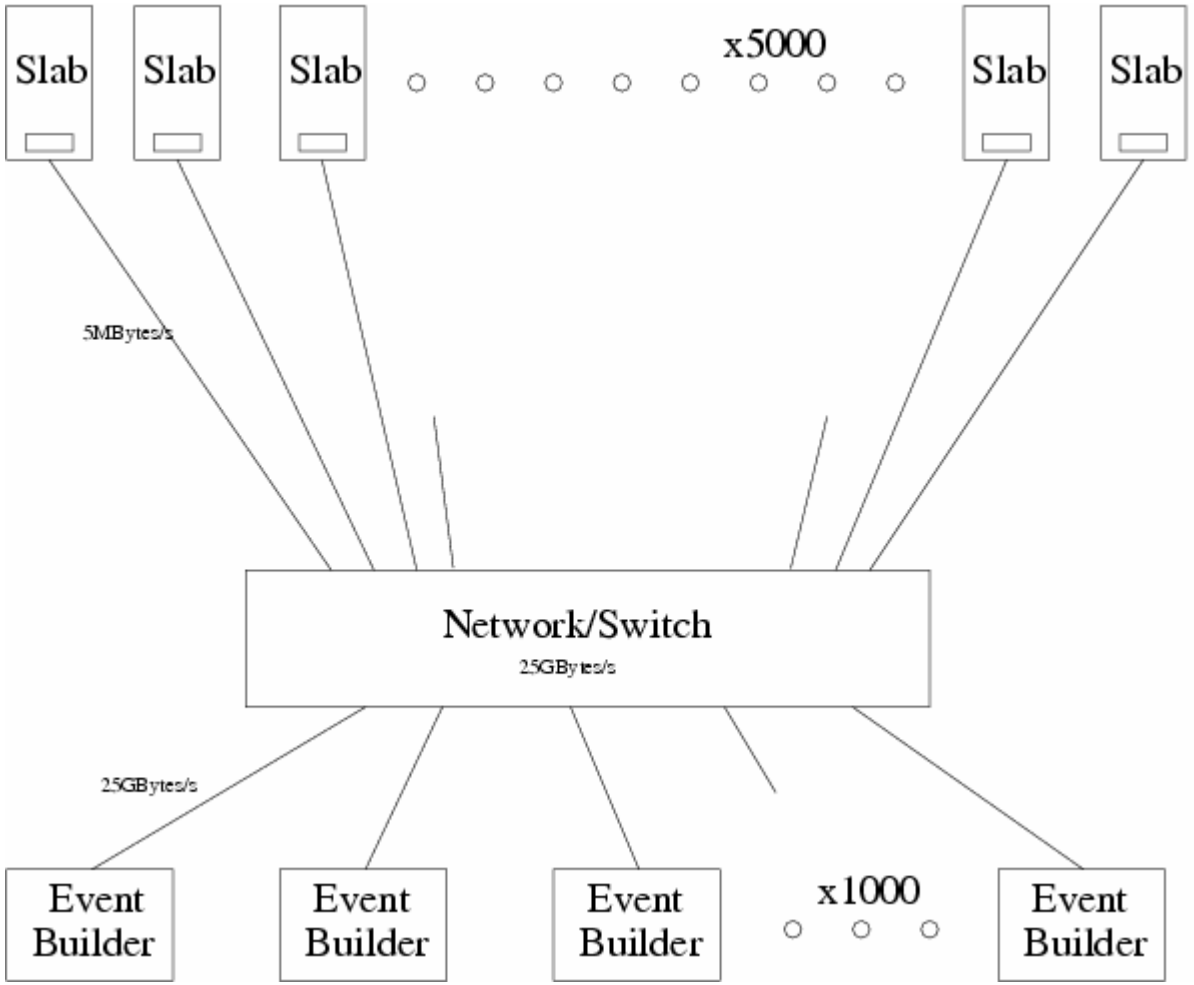

Custom receiver

- Go directly to a PCI card sitting in a off-the-shelf PC
  - PCI card would need to be custom
  - Around 4 slab output fibres per PCI card; total of 1200 PCI cards
  - Around 4 PCI cards per PC
  - Total of ~ 300 PCs needed; each receives ~ 20MBytes per bunch train
- PCs act as (cheaper) buffering to network
  - Much easier to upgrade and maintain; also gives CPU power
- Data rates are reasonable
  - Rate into PCI card is ~ 4×5MBytes/s ~ 20MBytes/s
  - Rate onto PCI bus is ~ 4×20MBytes/s ~ 80MBytes/s
  - Should be straightforward with new PCI standards (PCI Express…)
- Need to buffer train data during transfer to event builders
  - Store ~ 1000 trains, need ~ 1000×20MBytes ~ 20GBytes

Custom receiver (cont)

- PCI card would also act as fast control distribution card
  - Driver for fast control fibre mounted on it also
- Timing, clock, control and configuration comes from PCI card
  - PC determines configuration
  - Fast control system determines timing and control

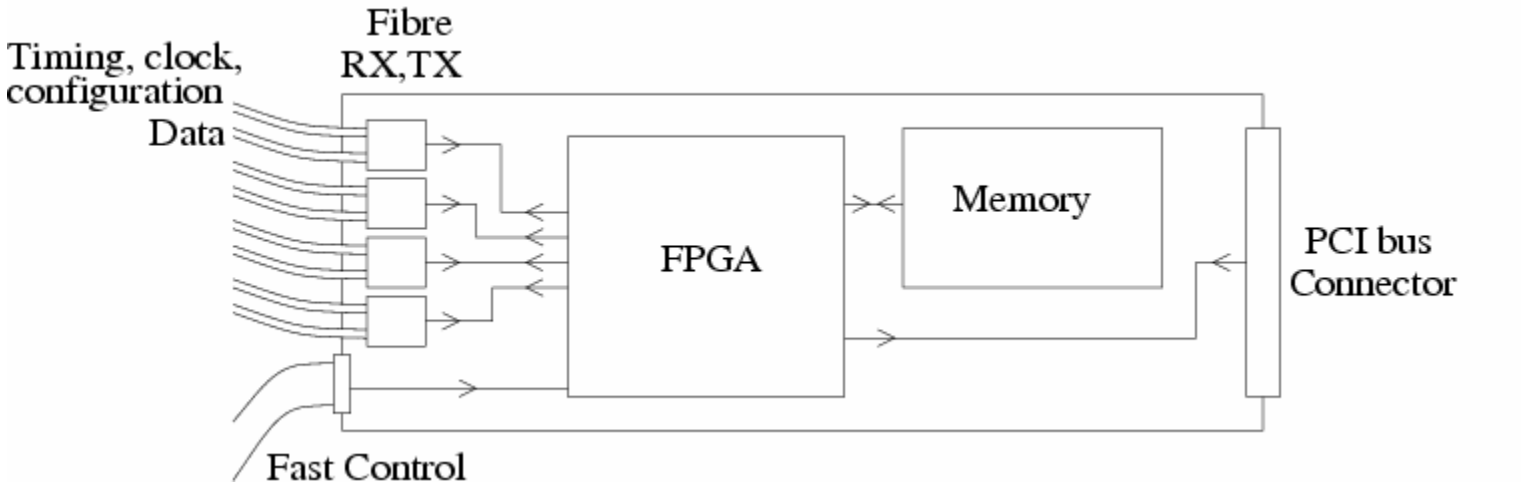

Custom receiver (cont)

- Data rate out of PC is uncertain
  - PC could do a lot of preprocessing of data…
  - …but each only sees ~0.3% of the ECAL so significant pattern recognition not very easy
  - Will certainly do histogramming, monitoring, etc.
- Difficult to estimate how much of the ~ 5GBytes per bunch train sent out over network to event builders
  - If all, then still need network capable of handling an average ~ 25GBytes/s
  - Each individual connection can be significantly slower
- Assuming effectively no suppression
  - Rate from each PC = rate on PCI bus ~ 80MBytes/s

Custom receiver (cont)

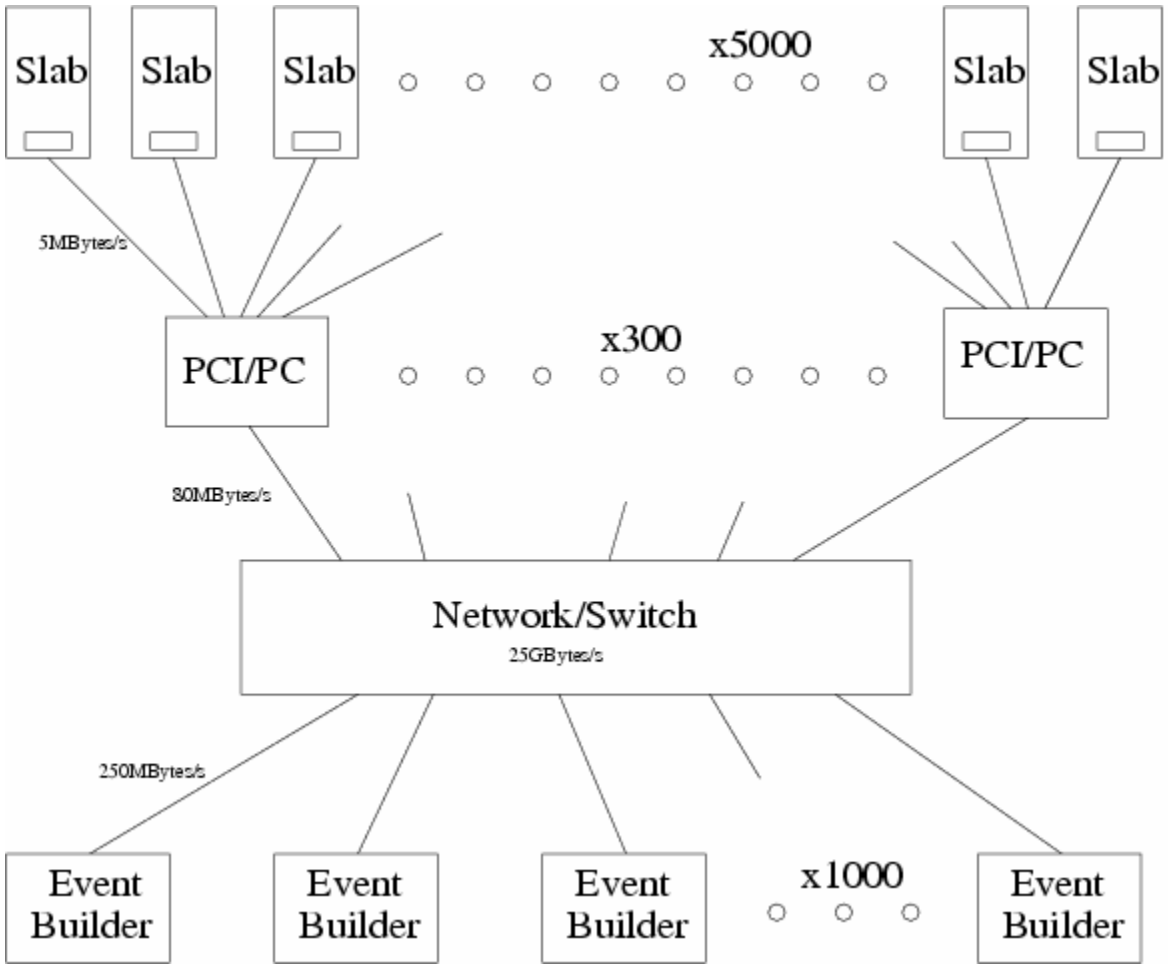

Cosmics?

- Is there a requirement for cosmic data?
  - Monitor pedestal and noise changes? Alignment?
  - For ~ 20 usable pad hits per cosmic track, need ~ 1M cosmics to average one hit per pad
  - With a cosmic rate of ~ 100Hz, this takes 10ksecs ~ 3 hours
- "Simplest" method is to run fake bunch train data-taking between real bunch trains
  - Need longer ASIC shaping time to give multiple samples per waveform; allow reconstruction of time and peak for arbitrary time of cosmics
  - At least 3 samples/waveform; shaping time ~ 3× bunch crossing ~ 500ns
  - Slower shaping also helps reduce intrinsic signal/noise; do anyway?
- Changes requirements for getting data off slab
  - Need to not rely on power-down between bunch trains to reduce cooling requirements
  - Need to speed up readout of bunch train data

Cosmics (cont)

- E.g. for 10% cosmic duty factor
  - This only gives ~ 1 hit/pad per day
- Front end (wafers and ASICs) have 10% power duty factor
  - Up from ~ 0.5%, makes cooling more difficult
- Need to read slab within a time equal to 10× the bunch train length
  - Transfer time 10ms not 200ms
  - Data rate from slab now ~ 100MBytes/s, not ~ 5MBytes/s
  - Still easily within rate possible for fibre
- Data rates through PCI cards and into PCs also up by 20×
  - Rate on PCI bus ~ 2GBytes/s; unclear if feasible
  - Need to suppress most data in PCs; otherwise 500GBytes/s onto network
  - Probably kills network idea without significant slab buffering (20GBytes)
- Important to determine whether cosmics are needed
  - Significant change to DAQ system requirements

Future plans

- Identified some areas which need work
  - Some of it is bringing us up to speed on what exists
  - Some areas we already know definitely need R&D
- UK groups plan to write proposal for ECAL DAQ
  - Sub-bid within next round of CALICE funding
  - Cover various of the R&D items discussed here
  - Submit early next year
- Need to firm up ideas between now and then
  - Input, collaboration, etc, very welcome