

LC Note 05-XXX

January 2005

# **A proposed DAQ system for a calorimeter at the International Linear Collider**

CALICE-UK

contributors from Cambridge(?), Edinburgh, Imperial, Manchester(?), UCL

9 November 2004

## **Abstract**

This note currently describes all thoughts and ideas on the subject of R&D for a data acquisition system for a calorimeter at the future linear collider. It is not yet a technical note nor a proposal draft to the funding agency; both are contained. It is partly a copy of Paul's ECFA talk, the minutes from the brainstorming meeting and subsequent brainstorming at UCL.

# 1 Introduction

With the recent decision on the accelerator technology to be used for a future International Linear Collider (ILC), detector R&D can become more focused. The time-line for the an R&D programme is also clearer with, assuming a technical design report to be written by 2009, five years to define the make-up of a given sub-detector. Within the CALICE collaboration, which is designing a calorimeter for the ILC, a collection of UK groups (CALICE-UK) are part of the initial effort to prototype a calorimeter composed of silicon and tungsten [1]. The electromagnetic section of the calorimeter (ECAL) is expected to take test-beam data using single electrons at DESY later this year. The UK has designed and built electronics to readout the ECAL [2] and will play a lead rôle in the coming data-taking period. Building on this expertise, CALICE-UK is defining its R&D programme with a significant part of it being the design of the data acquisition (DAQ) system for a future calorimeter.

This document details such a programme for three years of research effort which is split into several parts. The main aim is to start an R&D programme which will work towards designing the *actual* DAQ system of the future calorimeter. The lead technology for this is the combination of silicon and tungsten as proposed by the CALICE collaboration. Hence the generic DAQ system proposed will aim to read out data from such a calorimeter. The parameters of the superconducting accelerator design and calorimeter structure and properties which impinge upon considerations of the DAQ system for a calorimeter are discussed in Section 2. The generic DAQ system is described in Section 3 outlining the concept, the ideas for future R&D and possible technology solutions.

Two further aspects to this three year proposal are detailed in Sections 4 and 5. The very front end (VFE) chip is being designed in France and we propose to collaborate with the French groups and provide the read-out for the VFE chip for test-beams sometime in 2006. The other aspect of the DAQ programme is related to another work-package in this proposal which will investigate a new design of calorimeter using monolithic active pixel sensors (MAPS) instead of silicon diodes. This is a challenging new concept and is described in detail in that work-package. Such a design will also require a different DAQ concept which we, in the UK, will provide with the goal of again taking test-beam data to see if this design would be possible for a future calorimeter.

## 2 General detector and accelerator parameters

The current favoured design for the sampling calorimeter for the ILC is composed of 40 layers of silicon interleaved with tungsten. The calorimeter, shown in Fig. 1, has eight-fold symmetry and dimensions: a radius of about 2 m, a length of about 5 m and a thickness of about 20 cm. The silicon is organised in  $1 \times 1 \text{ cm}^2$  p-n diode pads which given the dimensions of the calorimeter, leads to a total of about 24 million pads. Mechanically, the calorimeter will be composed of about 6000 slabs each of which contains about 4000 silicon diode pads.

It is likely that ASICs will be mounted on or near wafers with about 200 000 required in total. The ASICs will perform pre-amplification and shaping and possibly digitisation and threshold suppression. The ASIC power consumption has to be minimised as they are difficult to cool due to the small gaps between layers which are required to take advantage of tungsten's Molière radius. Communication to outside the detector is via electronics at the end of modules as shown in Fig. 1.

The machine parameters which will affect the calorimeter design are the following. Machine running at a centre-of-mass energy of 800 GeV is assumed as any lower energy running should be less




Figure 1: View of the barrel calorimeter modules and detail of the overlap region between two modules, with space for the front-end electronics.

demanding for reading out the calorimeter. The bunch crossing period within a bunch train is 176 ns with 4886 crossings per bunch train. The bunch train length is 860  $\mu$ s with a period of 250 ns. The ECAL is expected to digitise the signal every bunch crossing and readout completely before the next bunch train.

### 3 Generic data-acquisition design

Any data acquisition system for the calorimeter is going to have to handle extremely large data volumes due to the large number of channels. In a shower, up to 100 particles/mm<sup>2</sup> can be expected, which in a 1×1 cm<sup>2</sup> pad equates to 10000 minimum ionising particle deposits. The ADC therefore needs a dynamic range of 14 bits. Assuming that 2 bytes per pad per sample, then the raw data per bunch train is  $24 \cdot 10^6 \times 4886 \times 2 = 250$  GB which equates to 1 MByte for each ASIC. Transferring 200 GB of data per bunch train is non-ideal and reducing due to application of noise thresholds could reduce this to a few GB.

#### 3.1 Transmitting digitised data from the VFE chip to FE

This is very heavily influenced by what can be done within the slab given the low heat load requirements due to the difficulties of cooling. It is not yet known how much will be put in the VFE ASIC, so various possibilities were considered.

In general, somewhere in the readout system for the pad options, there will have to be an ADC and a threshold discriminator. These could in principle be done in either order and could be done in the VFE or in the FE. There are then four possibilities:

1. Neither is done in the VFE
2. Only the ADC is done in the VFE
3. Only the threshold is done in the VFE
4. Both are done in the VFE

It is considered that any threshold discrimination is best done after the ADC step rather than before. This allows much easier control of monitoring the pedestals and noise, etc., by allowing some readout at a low rate even when below the threshold. In addition, setting a stable analogue threshold is not easy; any drifts will change the level. The uniformity over all channels might not be good enough which would then require a large number of trim DACs.

1) If neither an ADC or threshold discriminator is built into the VFE ASIC (due to them taking too much power), then the raw analogue signals will be sent out to the FE. This is 2k analogue channels which require around 14 bits precision, which is not at all trivial to achieve. Even if this can be done, digitising the data at the FE would be hard. The space is limited and so it is likely only a restricted number of ADCs could be mounted in this area. Assuming 20 channels of ADCs would be possible, then each would have to handle 100 pads, with these being multiplexed in turn into the ADC. To keep up with the sampling rate needed, i.e. 176 ns for each channel would therefore require the ADCs to sample at 1.76 ns. Finding a 14-bit FADC which can do this would not be easy. The alternative would be to use an analogue pipeline; assuming one for each of the 20 ADC channels would result in each pipeline having to store about 500k analogue samples which is difficult. Putting an analogue threshold in front of the ADCs would clearly cut the rate down but would need a major ASIC development to handle this; a variable length analogue pipeline with timestamps would be needed. This is in addition to the pedestal monitoring problems mentioned above.

2) This seems a much more reasonable option. The 14-bit requirement is much easier to achieve with a short signal path before the ADC. The digitised data can be transmitted from the VFE to the FE more easily than analogue data. The rates are not trivial however; these would be around 20 GB/s per slab, or 1 GB/s from each wafer/ASIC. This is at the level where a fibre would be needed; commercial fibres now carry 5 GB/s. Fibres are also less noisy than copper. This use of fibres within the slab would raise many other issues such as the power needed to transmit the light out (or could it be supplied by an external laser and then only modulated on the ASIC?), how to reliably attach the fibres at each end (a total of 200k fibres would be needed), how large the fibre connectors would be (the total thickness within the slabs is limited to some mm only), etc.. Although this is an active area of commercial development, it is not clear if opto-electronics inter-PCB communications will become standard enough on the timescale needed [3].

It is clear some development would be needed for this to be an option; the equivalent system in Atlas has three fibres transmitting a total of 10 MB/s with a 2 mm high connector needed. Self-aligning silicon-fibre interfaces are possibilities; while we could not do significant R&D compared with the commercial sector, we could test industrial prototypes and it might be possible to apply for a PIPPS award along these lines.

Once the data are on a fibre direct from the ASIC, the idea of whether any FE electronics would be needed at all was raised, as the fibre could go 100s of metres, bypassing the FE completely. However, shipping out all the raw data to the offline seemed a very expensive overkill.

3) This suffers from the same problems as mentioned above; there is a difficulty of monitoring the

pedestals as well as the complexity of the ASIC needed to handle the channels. Here, the ASIC would be incorporated into the VFE ASIC, but the effort needed would still be substantially bigger.

4) This option places the easiest requirements on the FE, with a corresponding increase in difficulty for the VFE (assuming the threshold is applied after the ADC). It would still need some communication of the threshold configuration from the FE to the VFE. The data rate out is clearly reduced; it would be around 400 MB/s for the slab, or 20 MB/s for each wafer/ASIC. Although easiest for transferring data from the VFE to FE, due to the low rates, it is not clear if the threshold can be reliably applied in the VFE. This scenario also looks like the situation if the MAPS were used rather than silicon diodes.

We, therefore, propose R&D for both scenario 2) and 4) because they both provide realistic solutions and have complementary applications. In favour of 2), any threshold suppression can be performed more accurately in the FPGA at the FE rather than in the VFE. For 4), the data transfer rate from the VFE and FE is significantly smaller. Scenario 4) will also provide the read-out for the MAPS technology as discussed in Section 5. A schematic of scenario 2) is shown in Fig. 2.

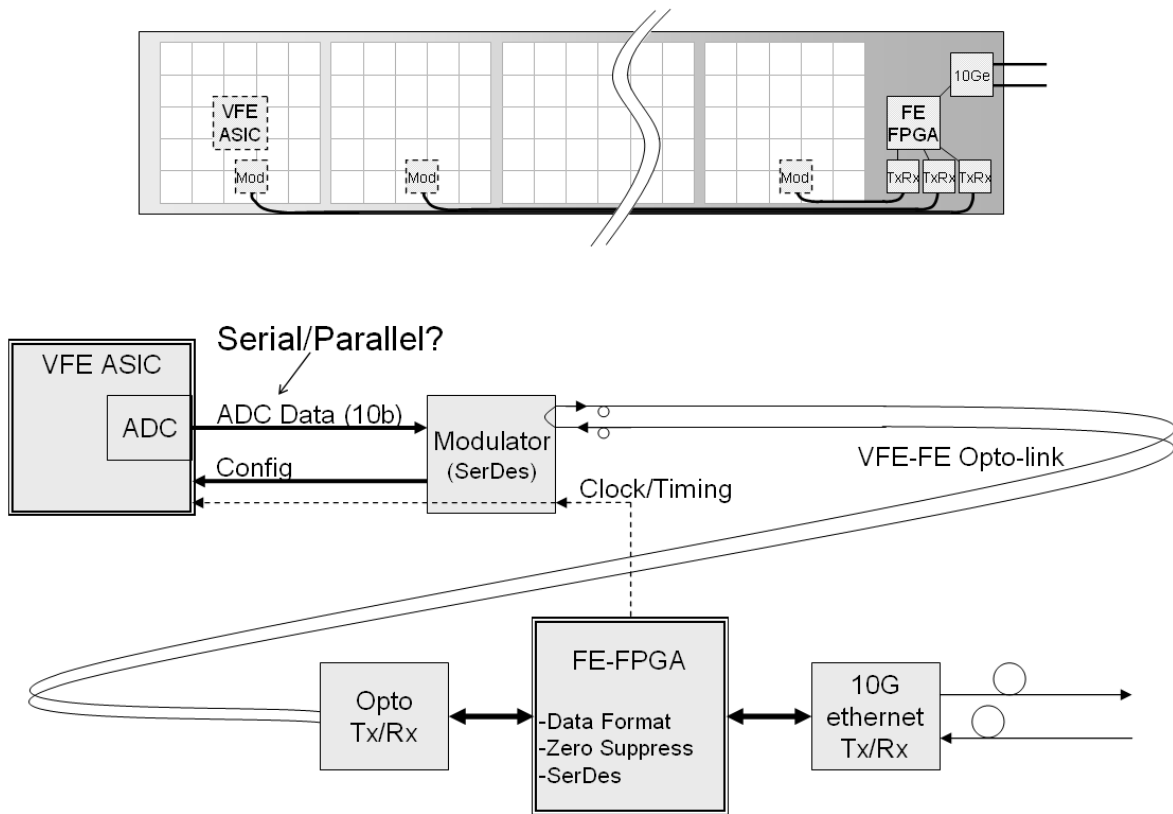


Figure 2: Design of VFE to FE link.

In both scenarios, we intend to set-up a mock data transfer system which requires having a test board with FPGAs linked by fibres. This will simulate a link between the VFE ASICs and the FE FPGAs. Any developed system e.g. a new VFE chip design or the MAPS set-up could also be tested in our prototype system. We will also demonstrate that the system would work for the hadronic calorimeter as well as the ECAL. This would require modifying the system to have a more

links but a lower rate. The prototype will incorporate, wherever possible, commercially available components such as Virtex-4 FPGAs [4] which has multi-gigabit serial transceivers and is compatible with 10/100/1000 Mb/s ethernet and PCI express x16 and higher.

### 3.2 Data from the front-end to off-detector

This is assumed to be fibre, but could be either dedicated point-to-point to a PCI card or a direct connection to a network, using something like TCP/IP protocol running in an FPGA in the FE. The former would be straightforward and would be guaranteed to send the data volumes needed in a fixed time. The rates themselves are in all cases less than 100 MB/s so there is no problem there. There would be 10k fibres coming off the detector, which is not a large number compared to LHC experiments. Typical fibre sizes are 150  $\mu\text{m}$  plus cladding, or 250  $\mu\text{m}$  diameter total. Hence, 10k fibres would take up 2.5 m of the 12 m circumference, leaving ample room for power and cables for other detectors. For the direct connection to a network (e.g. using Gigabit ethernet), then there is a lot more flexibility about where the data are sent, allowing for hardware failures, etc., but there is no guarantee that the data will be sent before the next bunch train. No currently available network switch could in fact sustain the throughput required.

As the network solution currently sounds unfeasible, we will concentrate on the point-to-point connection from the FPGA to a PCI card. The set-up would be as follows. A fibre would be used to transport data to the PCI card, but this would first go through a passive router. The passive router would have a fast control system which knows which PCs are alive and therefore is able to send the data to a sensible place. Geographically local information would be sent by the passive router to the same PC to allow clustering of calorimeter data in the PCI card. It could also have copies of data which is sent to more than one PC allowing, for example, overlaps in clustering regions. The passive router would also allow a PCI card to send a busy signal such that data is rerouted elsewhere. The busy signal could come from some particularly large sample of hits in the calorimeter which could be of interest and need a long processing time.

Again this scenario will be prototyped and a test system set-up.

### 3.3 Off-detector receiver

Having a standard PC with a PCI Express bus [5], would be a candidate for a point-to-point connection (see Fig. 3). PCI express cards have up to 32 lanes each of which provides transfers of 2.5 Gbit/s in each direction. Therefore a 32 lane link will have a total bandwidth of 20 GB/s. A possible scenario would be to have 4 cards for each PC with 16 lanes, i.e. fibres, per card giving a total of about 300 cards. Each PC would therefore process 64 slabs requiring around 100 PCs to handle the total of 5000 slabs which need to be read out. This complete system would not be too expensive depending on the price of such PCI express cards.

Buffering and/or data reduction would also be done on the PCI cards. The data reduction would be achieved, using the FPGA, by doing local clustering and removing isolated hits. This would be a challenge in which the technical aspects would go hand-in-hand with physics simulation efforts. Any algorithm for clustering at such a stage would need to be simulated and its effect on jet finding and general reconstruction of physics events understood and quantified. This would be the task of a new RA and would be coupled with the simulation workpackage.

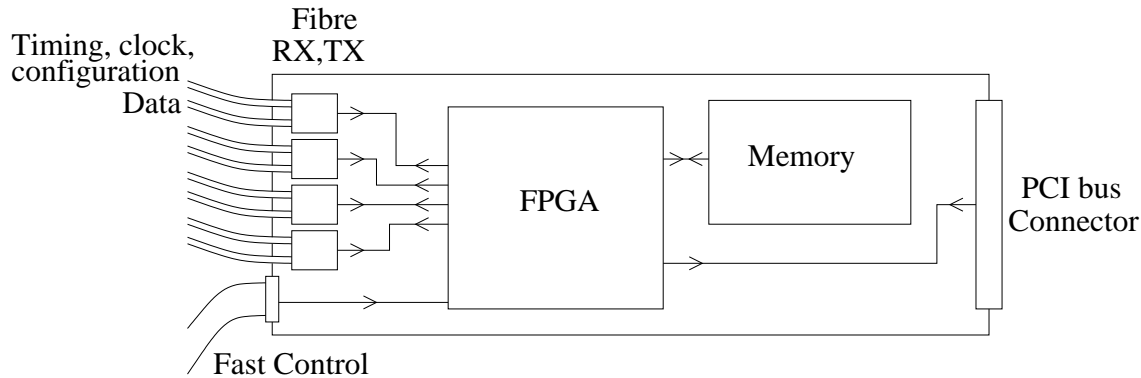


Figure 3: Design of a customised PCI card.

The timing, clock, control and configuration may also come from the PCI card. The PC would then determine the configuration and the fast control system would determine the timing and control. Another scenario is where the clock is controlled along with the switching of data between the FE and PCI cards. Knowing the bunch crossing time of 176 ns the box would feed out the clocking with the switch to the FE and PCI cards. We have many different components which have different clocking so getting them all in phase is a challenge which will be investigated. This is exacerbated by using commercial components, all of which may have different clocks. However, the ease of use, availability and cost mean that the advantages of commercial outway those of proprietary components even though the problem of clock control needs to be understood.

The reliability of large PC farms is an issue for reading out the data; if one PC goes down, all of the data in that region of the calorimeter is lost. This could be overcome by having a network switch as shown in Fig. 4 which decides which PC gets the data depending on whether it is alive or not. Current PC farms show a rate of 1 PC failure per day in farm of 200. This is not large but is also not small and would require some surplus of PCs (say 10%) above the number required based just on the number of detector channels. For a final working calorimeter readout system these PCs would need to be repaired and put back into the system on a regular basis. Therefore our program would be to add in a “off-detector” receiver onto our test system detailed in the previous two subsection. This would require modified PCI cards in a small farm of say 5 PCs. These three sub-components; VFE to FE link, FE link to off-detector and off-detector layout, would form a data acquisition system for a calorimeter at the ILC. The design will be able to test the realistic rates and dataflow expected and be scalable to a full detector system. New calorimeter designs and prototypes can be tested using our DAQ prototype system.

## 4 Reading out current design of the VFE chip

French collaborators within CALICE are designing the very-front-end (VFE) chip which is expected to be that used in the actual calorimeter. A prototype version of this is expected at the beginning of 2006. The chip will require test beam to study its performance in a prototype calorimeter. The data from the VFE needs to be read out and our expertise will allow us to design and implement this. As well as providing necessary R&D, this will also maintain the close collaboration between French and UK groups.

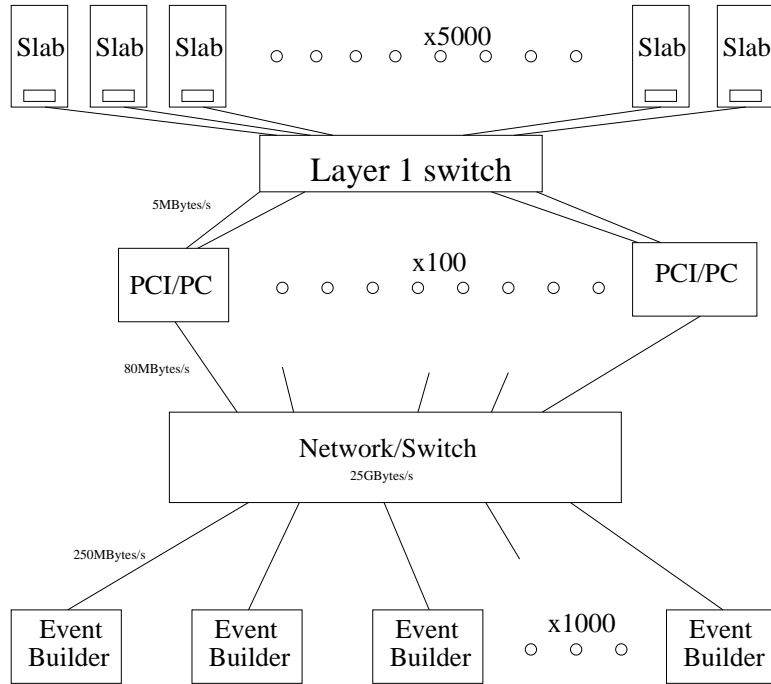


Figure 4: Design layout for transferring data from slabs to an event builder.

Current designs of the VFE chip [6] in Orsay, France have the following specifications. The final chip should have  $\sim 36$  channels to match a wafer and would be embedded inside the detector. The ADC(s) should be included in the chip in order to output digital data serially at high rate (typically 1-2Gb/s). The DAQ would thus look more like "an event builder" than a traditional DAQ. It would perform the data reformatting (from "floating" gain + 10bit to 16 bit), calibration, possibly linearisation and some digital filtering. It is possible that at this level, some event processing be performed. The other task of the DAQ is to load all the parameters needed by the front-end, control the power cycling and run the calibration. These specifications fit in well with our current generic system.

## 5 The DAQ system for a calorimeter based on MAPS technology

Instead of using silicon diodes, the feasibility of using the MAPS technology is to be investigated. The use of this technology would also have an impact on the design of the DAQ system. Here, there would be no ADC and a threshold has to be applied on the wafer, by definition. The data rate would be 3 GB/s per slab, or 150 MB/s per wafer, which is low enough for non-fibre communication.

Questions:

- How much to go in here? More input from MAPS work-package needed.



## References

- [1] TESLA: The superconducting Electron-Positron Linear Collider with an integrated X-Ray Laser Laboratory. Technical Design Report. DESY 2001-011. March 2001.
- [2] CALICE-UK: Proposal 317 - The CALICE collaboration: calorimeter studies for a future linear collider.
- [3] <http://www.intel.com/update/contents/it04041.htm>
- [4] [http://www.xilinx.com/xlnx/xil\\_prodcats/landingpage.jsp?title=Virtex-4](http://www.xilinx.com/xlnx/xil_prodcats/landingpage.jsp?title=Virtex-4)
- [5] [http://www.intel.com/update/contents/st11041.htm?iid=labs\\_homepage+update\\_st11041&http://www.intel.com/technology/pciexpress/downloads/pci\\_ei\\_pcb\\_guidelines.pdf](http://www.intel.com/update/contents/st11041.htm?iid=labs_homepage+update_st11041&http://www.intel.com/technology/pciexpress/downloads/pci_ei_pcb_guidelines.pdf)
- [6] C. De La Taille, VFE electronics design group, private communication.

People effort

Name	Position	Institute	Funding	2005/6	2006/7	2007/8
M. Goodrick(?)	Engineer	Cambridge	PPARC (RG)			
D. Ward(?)	Academic	Cambridge	HEFCE			
?	?	Edinburgh				
P. Dauncey	Academic	Imperial	HEFCE	?	?	?
O. Zorba	Engineer	Imperial	PPARC (RG)	?	?	?
New RA	RA	Imperial	PPARC (new)	?	?	?
R. Barlow(?)	Academic	Manchester	HEFCE			
R. Hughes-Jones(?)	Engineer	Manchester	PPARC (RG)			
S. Kolya(?)	Engineer	Manchester	PPARC (RG)			
M. Lancaster	Academic	UCL	HEFCE	0.2	0.2	0.2
M. Postranecky	Engineer	UCL	PPARC (RG)	0.3	0.3	0.5
M. Warren	Engineer	UCL	PPARC (RG)	0.3	0.5	0.5
M. Wing	Academic	UCL	HEFCE	0.3	0.3	0.3
New RA	RA	UCL	PPARC (new)	0.6	0.6	0.6

Table 1: Table of FTE effort from all Universities

Costs and miscellaneous

- Expect travel to be low as collaboration is mainly within the UK. However, collaboration with French groups requires some trips and beam tests are expected for all aspects of the programme. Assume £8k per year and £6k for travel to test beams, giving a total of £30k.
- Hardware costs and manufacture:
  - Fabrication of prototype boards can be done at UCL by the Mullard Space Science Laboratory (MSSL).
  - stuff

Suggested workpackages

- DAQ 1: VFE to FE  
VFE interface, FE board
- DAQ 2: FE to off-detector  
Optical network, switching, control and clock distribution
- DAQ 3: Off-detector receiver  
PCI card
- DAQ 4: Off-detector farm  
PC issues, software and optimisation for physics