

LC Note 05-XXX

September 2005

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

CALICE-UK

contributors from Cambridge, Imperial, Manchester, RHUL, UCL

25 August 2005

Abstract

This note describes R&D to be carried out on the data acquisition system for a calorimeter at the future International Linear Collider. A generic calorimeter and data acquisition system is described. Within this framework modified designs and potential bottlenecks within the current system are described. Solutions leading up to the Technical Design Report will to be carried out within CALICE-UK groups.

1 Introduction

With the 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 an R&D programme is also clearer with, assuming a technical design report to be written by 2009, four 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) will take test-beam data using single electrons at DESY in 2006. The UK has designed and built electronics to readout the ECAL [2] - these will also now be used by the analogue hadronic calorimeter - and will take part in the coming data-taking period. Building on this expertise, CALICE-UK is defining its R&D programme. A significant part of this programme is the design of the data acquisition (DAQ) system for a future calorimeter.

In this proposal, DAQ equipment will be developed which attacks likely bottlenecks in the future system and is also sufficiently generic to provide the readout for new prototype calorimeters. The main aim is to start an R&D programme which will work towards designing the *actual* DAQ system of the future calorimeter. Alternative designs of a DAQ system which could affect the layout of the final detector or functionality of components are also considered. For example we consider a setup where more information is processed on the ASIC mounted on the silicon pads inside the detector or off the detector and hence there is no need for FE electronics. Such a design could save space, and hence cost, and reduce the amount of heat dissipated inside the detector. A strong under-pinning thread here is to attempt to make use of commercial components and identify any problems with this approach.

This note is organised as follows. 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 main body of the note in Section 3 discusses the DAQ design and proposes areas of R&D within it. The proposal is split into five tasks which investigate the three principal stages of the DAQ system: the connection from the VFE to the FE; the connection from the on- to off-detector; and the off-detector receiver. The note ends with a brief summary in Section 4.

The programme detailed below will allow CALICE-UK groups to continue to assist in development of new technologies for the DAQ system. We would expect to write a chapter in the future technical design report on the DAQ system for the calorimeter. For the final calorimeter, the DAQ should ideally be the same for the ECAL and HCAL. Although CALICE-UK has so far concentrated on the ECAL, our proposals for R&D contained in this document are sufficiently generic that both calorimeter sections should be able to converge to use the DAQ system we design. This will place us in a position to build the DAQ system for future large-scale prototype calorimeters and the final system. Indeed the principle of a generic design using commercial components should be applicable to many detector sub-systems. Therefore, the R&D to be performed here may have consequences or applications to the global DAQ system for a future detector.

2 General detector and accelerator parameters

The design [1] for a calorimeter for the ILC poses challenges to the DAQ system mainly due to the large number of channels to be read out. The current favoured design for a sampling electromagnetic calorimeter 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. Mechanically, the calorimeter will consist of 6000 slabs, of length 1.5 m, each containing about 4000 silicon p-n diode pads of $1 \times 1 \text{ cm}^2$, giving a total of 24 million pads.

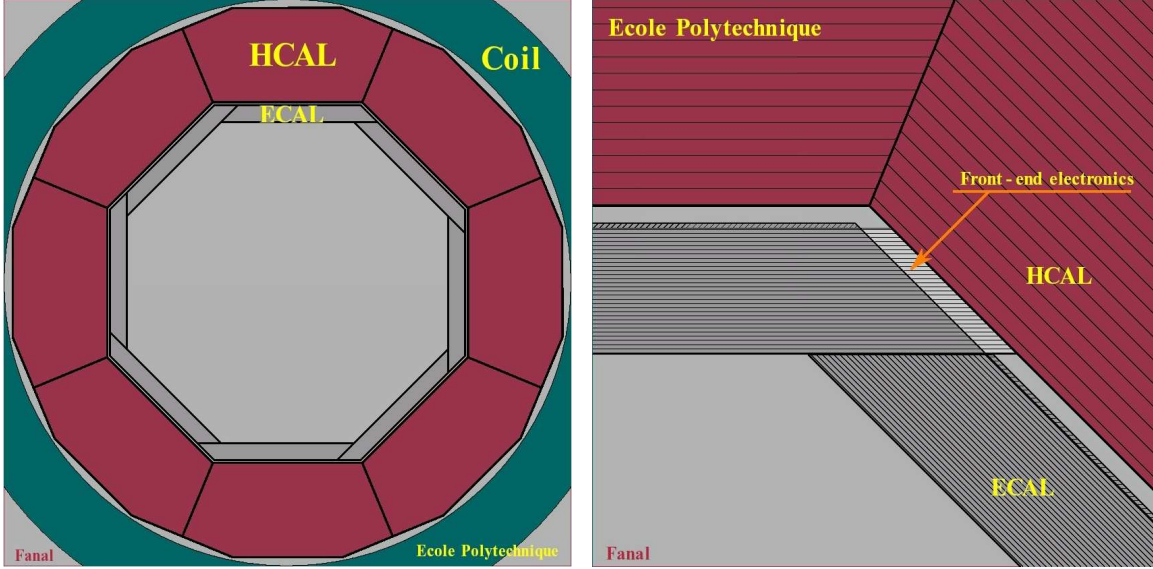


Figure 1: View of the barrel calorimeter modules and detail of the overlap region between two modules, with space for the front-end electronics.

A generic scenario for the DAQ system is as follows. At the very front end (VFE), ASIC chips will be mounted on the silicon and will process a given number of pads. The ASICs will perform pre-amplification and shaping and should also digitise the data and *may* even apply a threshold suppression. The current design [3] of such chips has each containing 16 channels, although this may be higher in the final calorimeter. 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. The data will be digitised in the ASIC and transferred to the front-end (FE) electronics which are placed in the detector at the end of the slab as shown in Fig. 1. It is expected that zero suppression will be done in the FE (using FPGAs) to significantly reduce the rate. The data will then be transferred off the detector, probably via a network switch, to a receiver of many PCI cards in a PC farm.

If we assume the TESLA design for data taking at 800 GeV, the following parameters have to be considered. There will be a 4886 bunch crossings every 176 ns in a bunch train, giving a bunch train length of about $860 \mu\text{s}$. The bunch train period is 250 ms, giving a duty factor between trains of about 0.35%. The ECAL is expected to digitise

the signal every bunch crossing and readout completely before the next bunch train. In a shower, up to 100 particles/mm² can be expected, which in a 1×1 cm² pad equates to 10000 minimum ionising particle deposits. The ADC therefore needs a dynamic range of 14 bits. Assuming no threshold suppression and that 2 bytes are used per pad per sample, then the raw data per bunch train is $24 \cdot 10^6 \times 4886 \times 2 = 250$ GBytes which equates to 0.3-2.5 MBytes for each ASIC depending on whether they process between 32 and 256 channels. The data appears within a bunch train length of 860 μ s giving a rate out of the ASIC of 0.4-3 GBytes/s, which we take to be 1 GBytes/s from now on. Threshold suppression and/or buffering (to allow readout between bunch trains) within the ASIC could reduce this rate. However, suppression in the ASIC may not be flexible enough compared with doing this in the FE and buffering requires some ASIC power to remain on between bunch trains, potentially generating too much heat. Hence the rates after the VFE depend on the assumptions made and system layout and will be discussed for each individual case where necessary.

3 Design of a DAQ system

3.1 Transmitting digitised data from the VFE chip

The transmission of digitised data from the ASIC 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 what the capabilities of the VFE ASIC will be, so various possibilities were considered.

In general, somewhere in the readout system, 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 ADC nor thresholding is done in the VFE
2. Only the ADC is done in the VFE
3. Only the thresholding is done in the VFE
4. Both are done in the VFE

It is considered that threshold discrimination is best done after the ADC step rather than before. This allows much easier 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 storing 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 time-stamps would be needed. This is in addition to the pedestal monitoring problems mentioned above.

2. Only doing the ADC on the VFE 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 50 GByte/s per slab, or 1 GByte/s from each wafer/ASIC. This is at the level where a fibre would be needed; commercial fibres now carry 5 GBytes/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 300 000 fibres would be needed for 6000 slabs each with 50 ASICs), 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-electronic intra-PCB communications will become standard enough on the time-scale needed [4].

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 MBytes/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 do R&D in conjunction with industry.

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 would go 10s of metres, bypassing the FE completely. However, shipping out all the raw data to the offline seems an expensive overkill, but is considered as this may change with commercial development.

3. Only doing the threshold in the VFE 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.

4. Doing both ADC and threshold in the VFE places the easiest requirements on the FE, with a corresponding increase in difficulty for the VFE. Assuming the threshold is applied after the ADC, some communication of the threshold and other configuration data from the FE to the VFE will still be needed. The data rate out is clearly reduced; it would be around 400 MBytes/s for the slab, or 20 MBytes/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 technology were used rather than silicon diodes. For the diode option in this scenario, it is also questionable as to whether FE electronics are needed. As the ASIC performs both ADC and threshold suppression, the data could be transferred directly off the detector.

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. A schematic of scenario 2) is shown in Fig. 2.

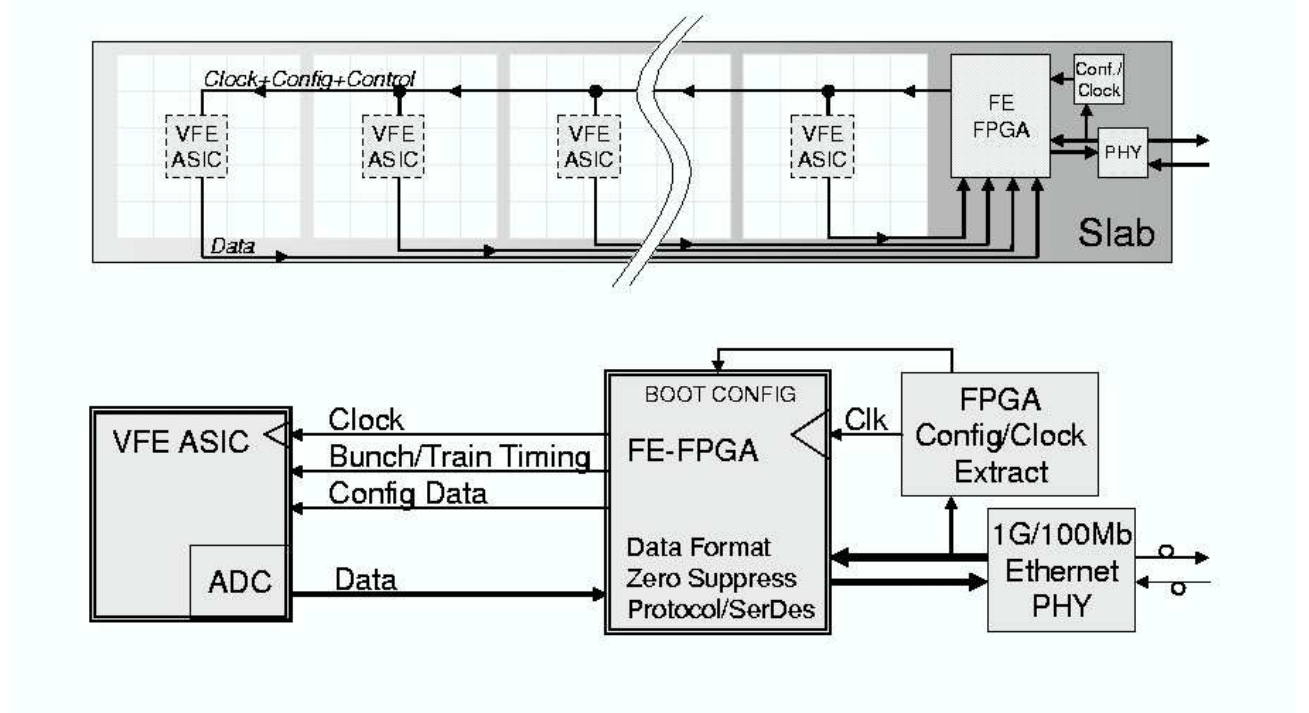


Figure 2: Design of VFE to FE link.

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 [5] which has multi-gigabit serial transceivers and is compatible with 10/100/1000 Mbit/s ethernet and PCI express x16 and higher.

Task 1: Reading out prototype VFE ASICs

The final chip should have at least 16 channels 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 1-2Gbit/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 will 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.

The current version of the VFE ASIC chip [8] is being used to read out the existing CALICE ECAL. This chip does not meet the requirements for the ILC ECAL and the development of the design is an ongoing project in LAL/Orsay [3]. In the next 1–2 years, they expect to have a version of such a chip with low enough power and noise that would serve as a realistic prototype. This ASIC is expected to have 16 channels, an internal ADC per channel, multiple gain ranges, and optional threshold suppression and digital buffering to reduce the required output rate. They expect to submit the ASIC for a fabrication round roughly once per year in early summer. We propose to build simple readout systems to study these ASICs.

The steps in this task are as follows:

- The LAL/Orsay group produce single-ASIC PCBs to study the chip performance. They will supply us with some of these PCBs with chips from each production round. This will allow us to understand the chip clocking and control.
- We propose to design and build a PCB to hold at least 32 chips, thereby reading out about 500 channels. This would need to use the threshold or buffering option to keep the rate low enough for straightforward readout.
- We will then mount and test multiple VFE ASICs on the PCB to give us practical experience in the associated issues of clock, control and configuration of these prototypes in bulk, as well as allowing studies of pedestal and noise stability. How these chips behave will significantly influence the design (and cost) of the final on-detector DAQ readout system and so early data here will be invaluable.

This UK activity is obviously paced by the work of the LAL/Orsay design team so is anticipated that activity will be mainly concentrated in 2008. A similar PCB readout system as used to readout the MAPS PCBs would be used, thereby consolidating firmware effort. Instead of using silicon diodes, the feasibility of using the MAPS technology is to be investigated [7]. 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 for a final detector would be 3 GBytes/s per slab, or 150 MBytes/s per wafer, which is low enough for non-fibre communication.

Task 2: Study of data paths over 1.5 m slab

If threshold suppression or buffering could be done in the VFE ASIC, the rate to the FE would be reduced by two orders of magnitude. Current designs cannot do this and it may not even be desirable or practical, so we have to allow for data rates of order GByte/s needing to be transferred out of each VFE ASIC during the bunch train. Whether an electrical or optical connection would be needed has to be investigated. Although chip-to-chip fibres are not yet standard technology, this is an active area of industrial research [4].

The construction of a 1.5 metre long fully operational calorimeter prototype is unlikely on timescales of the next few years. However, issues of how the data would be transported from the VFE to FE have to be considered and can be done already without a real prototype. Transporting of order GByte/s of data over 1.5 m in a very limited space is a challenge. The conventional approach would be to use copper but here the effects of noise and interference will have to be considered. There is also the possibility of using optical fibres although here there are also design considerations: the size of connectors would have to be investigated as the vertical clearance at the VFE is of the order of mm and the power needed to transmit light out would also need to be investigated. This work ties in closely with the mechanical and thermal aspects of the design.

In preparation for a real prototype, a test system will be built with a 1.5 metre PCB containing FPGAs linked optically or electronically. The data transfer would then be considered as a function of the number of VFEs, whether zero suppression is done in the VFEs and whether data is buffered during the bunch train. The bandwidth and cross-talk of the data transfer will be simulated using CAD tools. The clock and control distribution from the front-end to the VFE chips will be investigated as to whether one transmission line per chip is needed or multi-drop is possible. The resources required would initially be two multilayer 500 mm PCBs for initial tests and investigations before moving onto three 1600 mm PCBs for final tests.

The steps in this task are as follows:

- In preparation for a real prototype, we propose to build a test system with a 1.5m PCB containing FPGAs to emulate the VFE and FE, linked electrically. This would begin with two 500 mm PCBs for initial tests and investigations before moving onto three 1600 mm PCBs for final tests.
- The bandwidth and cross-talk of the data transfer will be simulated using CAD tools and compared with measurements from the PCBs.
- The clock and control distribution from the front-end to the VFE chips will be investigated to determine whether one transmission line per chip is needed or multi-drop is possible.

3.2 Connection from on- to off-detector

In this section, we consider two widely differing scenarios, (shown in Figure 3):

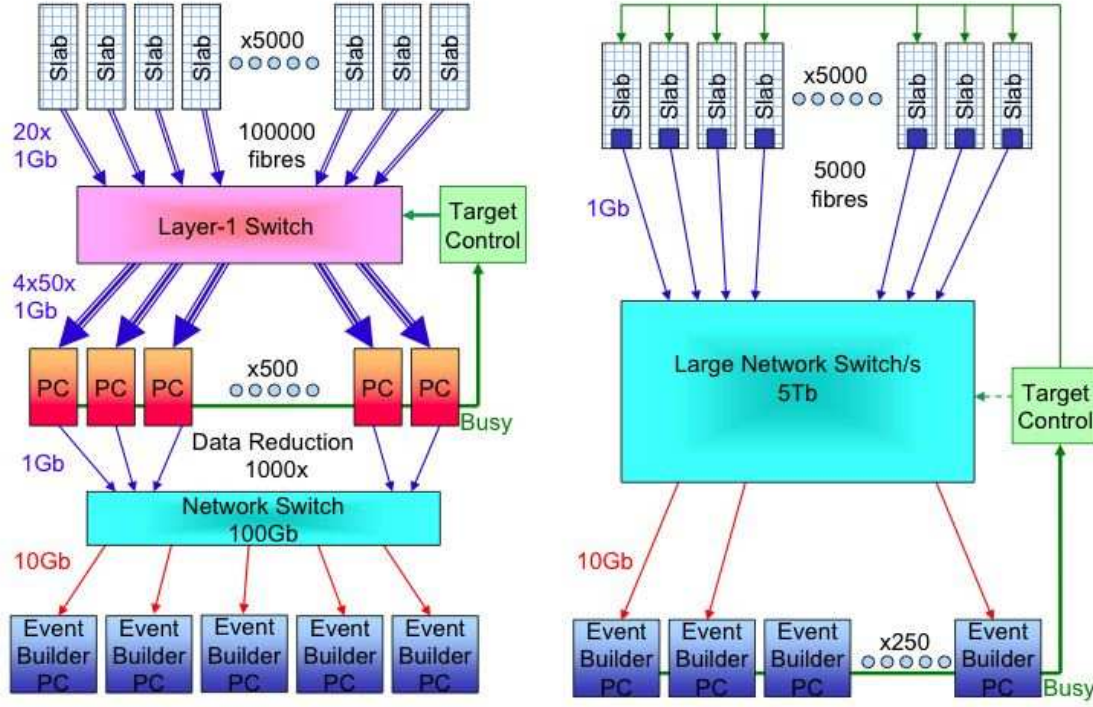


Figure 3: A comparison of the two scenarios for the on- to off-detector DAQ system.

Standard Configuration

In our assumed standard detector configuration, communication from the VFE will pass via the electronics at the FE to an off-detector receiver. We assume that threshold suppression will be done at the FE, and hence the rate would be significantly reduced from that at the VFE. Assuming that the rate is reduced to 1% of the original data volume of 250 GBytes per bunch train and each sample above threshold needs a channel and timing label, the total data volume to be read out from the calorimeter is about 5 GBytes or about 1 MByte per slab. These data have to be read out within a bunch train period of 250 ms, giving a rate of 5 MBytes/s.

Alternative configuration

Here we imagine that the FE is removed and the communication is directly from the VFE to the off-detector. We assume here that the ASIC only digitises the data and 250 GBytes has to be transported off the detector per bunch train. This will require a high-speed optical network. It should be noted, however, that the need for FE electronics also becomes questionable if more processing is done on the ASIC chip, such as threshold suppression. In such a scenario, transporting the data directly from the ASIC off the

detector could be done and so the FE would become redundant. The number of fibres required to read-out the 24 million channels would vary between 750 000 to about 90 000 depending on whether the ASIC handles 32 or 256 channels. If we assume that the diameter of a fibre is $150\text{ }\mu\text{m}$, or with cladding $250\text{ }\mu\text{m}$ then if half of the circumference of 12 m had fibres running along it, the bundle would be up to 1 cm in depth, but could be as little as 1 mm, depending on the number of ASICs and hence fibres. This would leave ample room for other cables and power supplies. This concept would revolutionise the whole calorimeter design and so needs to be considered now when changes in its general structure could be considered. Our research in this area will provide important feedback to groups designing the ASIC chips.

The off-detector receiver, as described later, is assumed to consist of PCI cards housed in PCs. 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. Current PC farms show a rate of 1 PC failure per day in a 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. The standard scenario would require less high-speed equipment off the detector, whereas the alternative would require many optical fibres with dedicated optical switching. The alternative scenario would, however, remove material from inside the detector which would ease construction and have a potentially advantageous impact on event reconstruction and, hence, physics measurements. It would also reduce the number of processing components within the detector which could be attractive since they would be inaccessible for long periods of time.

Using the experience gained from the two scenarios above, a hybrid path can be explored. By optimising the functionality in the VFE and matching it to an FE, overall instantaneous data rates can be reduced as well as lowering fibre count. For example, if data is buffered in the VFE, some form of passive multiplexer can be envisaged that will combine data from all the VFEs. This could take the form of spliced optic-fibres, or an OR gate. This would not necessarily remove the need for an FE, but the reduced size and complexity will have many benefits (thermal, configuration time, SEUs, cost, etc.).

Task 2.3: Connection from on-detector to the off-detector receiver

We propose to set up a test system to validate the effectiveness of high-rate switching. This will investigate the issues in directing multiple 10 Gbit data flows with the linear collider time structure to high-performance processors.

The steps in this task are as follows:

- The test system, for the standard configuration, will consist of four 10 Gbit ethernet interfaces each housed in a high-performance PC. These will be connected to six 1 Gbit ethernet PCs capable of driving data through the system. The 1 Gbit network interface cards will be a mixture of different cards and characterising them

will be part of the goals of the project.

- After building the system, a series of tests will be done. Firstly the network interface cards, motherboards and network switch will be characterised considering jitter, latency, throughput, etc.. Then simple data flows will be examined for different scenarios. Finally the data flow will be examined for rates expected from the accelerator timing structure.
- To investigate the alternative, we propose to set up a test system to validate the use of fibres to transport the necessary data rates from the VFE to the off-detector receiver and to ascertain whether an optical (“layer-1”) switch can re-route data with high efficiency. The test setup would contain a bundle of fibres coming into a switch from an effective VFE. A bundle would then come out of the switch to the off-detector receiver. The off-detector receiver, to be produced as part of Task 2.5, would be a set of PCI cards housed in PCs. The same PCI card would be used in our test system to emulate the VFE chip. As the PCI card can be customised, we are also able to explore new and bespoke protocols.
- The optical switch efficiency would then be tested when PCs fail and the data have to be swapped to another PC. Also if a PC contains a large amount of interesting data and takes up a large amount of processing time, a busy signal would be sent to the optical switch manager which would again re-route new incoming data. As it is assumed that some clustering will be done on the PCI cards, overlaps in geometrical areas of the calorimeter will be necessary. Therefore, the optical switch would also have to be able to send the same data to different PCs. The test system would consist of an optical switch connected via optical fibres to the PCI cards housed in PCs. We assume scalability of the optical switch, so we would use a relatively small 16×16 port switch.

Task 2.4: Transport of configuration, clock and control data

To transmit data onto the detector, we will attempt to use the same commercial hardware used for on- to off-detector communication. The requirements are different, though, as the detector front-end requires clock and synchronisation signals as well as low-level configuration and data signals. Commercial network hardware is not ideally suited to synchronous tasks but the significant cost and reliability benefits make it worthy of investigation. This is split into two areas: Failsafe configuration prior to normal running, and clock and control signal distribution.

Failsafe Configuration

In a scenario requiring an FPGA on-detector, it is imperative that the device can be re-configured remotely. This is necessary not only due to the number of FPGAs but also because of the uncertainty in the detector performance, in terms of data suppression, pedestal drifts, bad channel identification etc., all of which have to be implemented in the FE FPGA. The optimal algorithms for these tasks will only be determined after some experience of operating the calorimeter. As the FPGAs are relatively inaccessible, it is important that a failsafe method of resetting and reconfiguring the devices exists.

To boot an FPGA remotely requires either non-volatile memory with a base configuration, or additional hardware on the board, enabling external communication. Using devices with internal non-volatile memories may have advantages here, but increased cost and limited re-programming cycles make these less desirable. Questions surrounding the reliability of these memories in the detector environment over the lifetime of the experiment also need to be addressed. In contrast, datasheets of modern SRAM-based FPGAs boast that they can re-configured an infinite number of times, and could therefore be ‘refreshed’ at regular intervals to ensure a corrupt configuration is cleared. The option of using tailored firmware for different run types is also attractive.

Reducing the number of components on the front-end-module is advantageous from a cost, reliability and dead material point of view, so we will focus on methods of customising the datastream to the detector such that it requires the minimum of components to extract signals from the physical network interface module.

Clock and Control

To ensure the on-detector electronics captures data from actual collisions, all components need to be synchronised to a bunch-crossing clock. Bunch-train start, stop and ID signals will also be required. Traditionally this is done with a bespoke, low-latency, highly synchronous system, but here we will look at using a commercial network switch for signal distribution.

Switches are not designed for synchronous signal fanout - in fact the now obsolete forerunner to the switch, the hub, is much better suited to the task. Modes for signal broadcasting do exist, but these need to be closely examined. Studies on latency and a port-to-port skew will be undertaken.

To obtain maximal control over timing, the lowest level protocols will need to be used. Directly accessing the network at the physical layer (i.e. as a simple serial link) would facilitate complete control over data-packet composition, and hopefully timing. For this a specialist and customisable network interface card is required, such as the off-detector receiver detailed in Task 2.3 above.

To regenerate the clock at the detector end, the time structure of the data transmitted needs to be such that a tuned circuit can extract the clock. Board-level signals also need to be decoded from the data stream (as with the failsafe system above) and finally payload data needs to be made available to higher level logic.

In many ways this task can be thought of as an investigation to verify that a clock, control and configuration system can be constructed using commercial network hardware for distribution.

The steps in this task are as follows:

- We will consider a literature survey of the response of various FPGAs to radiation and rates of single event upsets and how to circumvent them.
- A small test system for the failsafe studies will be set up, requiring a PC with a specialist network card and an FPGA development board with cables and connectors.

- The clock and control test system will build on the previous FPGA failsafe connectivity work, the same equipment, but additional network cards and a switch will also be required.

3.3 Off-detector receiver

An ideal system would have all data from the detector for each bunch train sent directly to a single PC where full event reconstruction would be performed. However, considering that the electromagnetic calorimeter alone will contain 24 million channels, even after data suppression this seems unfeasible.

It would, however, be desirable to geographically group as many channels as possible at a processor. This would permit broader local clustering which could be used as an input to full event reconstruction later. The fundamental questions to be answered are how much of the calorimeter can be received into one PC and how much needs to be received to make local clustering effective. The impact this has on full event reconstruction and on simulated physics needs to be evaluated. This task therefore involves a combination of hardware and simulation work.

Task 2.5: Prototype off-detector receiver

We propose a system comprising PCI cards (see Fig. 4) mounted in PCs. These cards will provide timing, clock, control and configuration data, be capable of performing local clustering, as well as being able to be reconfigured to act as data transmitters to exercise the receivers. We will use the 'new' PCI Express bus standard [6] as this offers increased bandwidth and flexibility by using many high bandwidth serial links (called lanes) instead of a single parallel bus. It is expected to be scalable due to future increases in the number of lanes. Aside from providing an interface to the processor, lanes can be setup between cards, allowing advanced clustering prior to processing.

Although it is hard to predict what technology will be available for the final system, gaining experience with cutting-edge technologies now, will prove valuable for estimates of how much data can be realistically processed by a computer and help us understand the benefits and bottlenecks, as well as the scalability.

Much of this DAQ workpackage relies on novel methods in using network infrastructure, so a custom card providing both direct access to the network hardware and room to build any other functionality is useful. This applies to the PCI Express interface too, where the card will be able to test-drive the bus and provide debug feedback and monitoring.

The steps in this task are as follows:

- Design and build a PCI Express based interface card suitable for all the tests we would like to perform.
- Build a test system consisting of four PCs with one PCI card in each for communications testing. Four further PCI cards, housed in two PCs, will be required

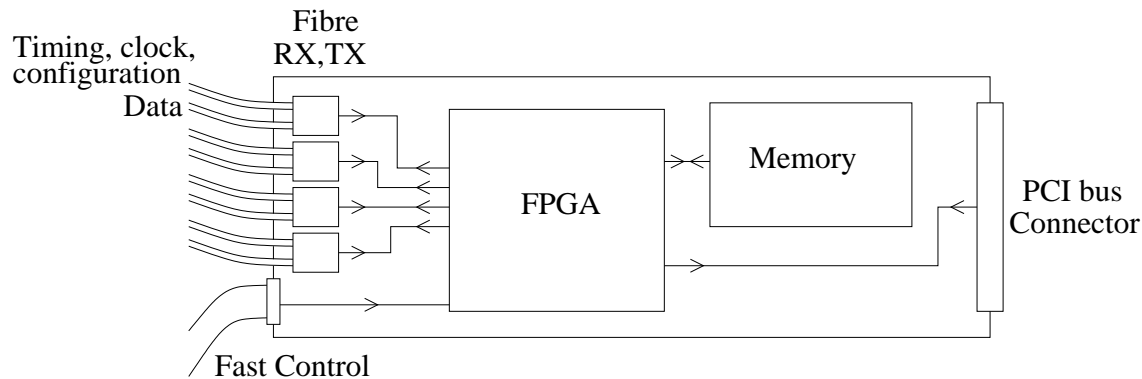


Figure 4: Design of a customised PCI card.

to act as VFE- or FE-like data sources. Tests designed to maximise throughput, such as four cards in a single PC, driven by four PCs with a single card in each will be performed. The cards will be produced as two pairs of prototypes, with a production run of eight for the full test system, so testing will be structured accordingly.

- In conjunction with the network workpackages discussed previously, tests will be made to see how efficiently data destinations can be changed if a PC becomes unreachable or busy.
- With sufficient data transported to one computer, local clustering and removal of isolated hits can be performed before full event reconstruction is done online. This would involve firmware programming of the FPGA on the PCI card, its simulation and subsequent effect on full reconstruction and performance in a physics analysis.

4 Summary

The conceptual design of a data acquisition system for the ILC calorimeter has been detailed. The concept relies heavily on commercial equipment and is generic such that it could be applied to other detector systems. Potential bottlenecks have been identified and form a programme of research and development of the next three to four years for the CALICE-UK groups. Bench tests and real-time tests with prototype calorimeters will be undertaken. After this period, a technical design of the data acquisition system will be possible.

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] C. De La Taille, VFE electronics design group, private communication.
- [4] <http://www.intel.com/update/contents/it04041.htm>
- [5] http://www.xilinx.com/xlnx/xil_prodcataloglandingpage.jsp?title=Virtex-4
- [6] [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)
- [7] P.D. Dauncey et al., “MAPS conceptual design”, LC Note in preparation.
- [8] J.Fleury, et al., “Front End Electronics for Si-W Calorimeter,” proceedings of the LCWS04 Conference, Paris, April 2004.