CALICE - Calorimetry for the International Linear Collider

C.G.Ainsley[2], R.J.Barlow[4], G.Boorman[5], D.Bowerman[3], P.D.Dauncey[3]*, C.Fry[3], B.J.Green[5], M.G.Green[5], C.M.Hawkes[1], M.Kelly[4], M.Lancaster[7], N.M.Malden[4], G.Mavromanolakis[2], M.Postranecky[7], P.-F.Salvatore[5], R.J.Thompson[4], M.A.Thomson[2], R.Turchetta[6], D.R.Ward[2], M.Warren[7], N.K.Watson[1], M.Wing[7], O.Zorba[3],

[1]University of Birmingham, [2]University of Cambridge, [3]Imperial College London, [4]Manchester University, [5]Royal Holloway, University of London, [6]Rutherford Appleton Laboratory, [7]University College London

December 13, 2004

1 Introduction

The CALICE collaboration is undertaking a major programme of R&D into calorimetry for the International Linear Collider (ILC). It is the only collaboration within the ILC community studying both electromagnetic (ECAL) and hadronic (HCAL) calorimeters in an integrated way. We believe this overall calorimetry approach is the only way to obtain a calorimeter system which will be capable of meeting the demanding physics requirements of a high energy ILC detector.

The CALICE program covers beam tests of several technologies and simulation studies based on the results of these tests to design an ILC calorimeter optimised for both performance and cost. The breadth of the project and the size of the collaboration ensure that CALICE cannot be ignored in a future ILC detector.

It is expected that there will be a call for ILC detector Technical Design Reports (TDR) around 2009. The programme proposed here should allow the UK to define its contribution to the TDR by that time. It is intended that the UK should have a clear lead in the areas where it is contributing, so as to give us influence for the future ILC detector programme.

1.1 General status of CALICE

The CALICE collaboration [1] now has XXX members from YYY institutes worldwide. It is by far the largest group studying calorimetry for the ILC. The collaboration intends to test pre-prototypes of an electromagnetic calorimeter (ECAL) along with at least two types of hadronic calorimeter (HCAL) in electron and hadron beams over the next two years. In addition, the collaboration serves as an umbrella organisation for longer-term ILC calorimeter projects where developments can be tested together.

The CALICE ECAL is a silicon-tungsten sampling calorimeter and consists of 30 layers of silicon wafers interspersed between tungsten sheets. Each wafer layer contains a 3×3 array of silicon wafers, each containing 36 $1 \times 1\,\mathrm{cm}^2$ diode pads. There are around 10,000 channels in total in a volume of approximately a cube with 18 cm sides. The ECAL is close to completion;

*Spokesperson

more than half of the silicon wafers are now manufactured. Around one third of the layers will be ready by the end of the year and the rest are expected to be complete by April 2005.

The analogue HCAL (AHCAL) is a tile scintillator sampling calorimeter with 40 layers of steel absorber sheets. The total volume is approximately a cube with 1 m sides. The tiles are of varying sizes, with the high granularity central region being $3 \times 3\,\mathrm{cm}^2$. As the name implies, the readout will be analogue, with the off-detector electronics being common to the ECAL. The AHCAL has around 8,000 channels and is scheduled to be completed by September 2005

The digital HCAL is a binary readout sampling calorimeter. The sensitive layers will be mainly resistive plate chambers (RPC) although for some of the tests, one or more layers may be replaced with gas electron multiplier (GEM) foils. The RPC (or GEM) pads will be $1 \times 1\,\mathrm{cm}^2$, giving 400,000 channels, each reading one bit. As one of the main aims of the beam tests is to compare the performance of these HCALs, the same absorber structure as for the AHCAL will be used, so as to eliminate any spurious differences which might otherwise arise. Hence, the DHCAL is also around 1 cubic metre in volume. The DHCAL is being developed by US groups, who are applying for funds for the readout electronics early in 2005. Assuming funding is secured, the DHCAL should then be completed by summer 2006.

The first ten layers of the ECAL will be put in the low energy electron test beam at DESY in January 2005. This will be a technical commissioning run to debug the system and provide a first look at the ECAL performance. This test is expected to continue intermittently over the following six months as the rest of the ECAL is assembled and installed. The completed ECAL will then be run continuously for June and July.

In September 2005, the ECAL will move to FNAL to take hadron beam data. Even without any HCAL behind it, it has been shown [2] that significant differences between hadronic shower models can be seen in the ECAL alone. These tests will continue until around November, when the AHCAL will be ready. The two detectors are then expected to take data for six months. Come summer 2006, the DHCAL should then arrive at FNAL and again, around six months is expected for data-taking with the ECAL and DHCAL together.

1.2 Background to the proposal

The CALICE-UK groups got approval from the PPRP in December 2002 for an initial programme with the CALICE collaboration. This was to provide the readout electronics for the ECAL and perform simulation studies. The approval covered the first two years of the CALICE programme.

It was expected that the groups would return to bid for continuation of the funding to complete the current CALICE programme as well as defining a longer-term programme to lead more directly to the ILC detector TDRs. This is the purpose of the current proposal.

The proposal is organised as follows. Section 2 describes the status of the work funded in the original proposal. Sections 3 to 7 outline the five proposed workpackages for the current bid. Section 8 gives some details of the management of the UK groups

2 Status of the current UK effort

2.1 Status of the UK electronics

The UK commitment was to design, build and commission the VME readout system for the ECAL. This is essentially complete.

Upstream of the UK electronics, the silicon wafer pads are amplified and read out with the very front end (VFE) ASIC chip, designed by the LAL-Orsay group. The wafers and VFE chips are mounted on the VFE PCB which is connected to the UK electronics via a mini-SCSI cable. The UK is responsible for all the readout electronics downstream of the VFE PCB.

2.1.1 Readout electronics

The readout system consists of six 9U VME boards which provide the control and the digitisation of the analogue signals from the VFE PCB. They also provide local buffering of data for around 2000 events, which should be more than sufficient for each spill. One of the boards also provides trigger logic and control, with the trigger being fanned out from this central board to the other boards in the system (including itself) using point-to-point connections on a custom-built PCB which attaches to the VME J0 connectors across the crate backplane. The maximum allowable trigger latency is set by the peaking time of the shaped VFE preamplifer signal and is around 180 ns.

The boards were designed as a modification of the CMS silicon tracker Front End Driver (FED) design [?]. The CALICE design required a complete replacement of the FED front end (FE), which receives the input data, with a new design. The back end (BE) and VME interface were less substantially modified. Because of this commonality, there has also been some firmware shared between the two boards, reducing the engineering effort required. A photograph of one of the boards is shown in figure **??**.

Two prototype boards were fabricated by November 2003 and were thoroughly tested over the following year. These tests included a suite of stand-alone tests and calibrations using the on-board DAC. In addition, tests with the VFE PCB preproduction boards were done during early summer 2004. These involved using the DAC to calibrate the VFE PCBs, as well as using radiative sources and cosmic rays. These tests showed the boards were working very well [3] and only minor improvements were implemented in the design for the production versions.

Nine boards were fabricated for the production run in October 2004, of which only two were populated immediately. These two were again extensively tested before the remaining seven were released for assembly. They are expected to be completed by February 2005. These two boards were taken to Paris for cosmic ray tests in December 2004. They will be transfered for the initial beam test at DESY which will start in January 2005 and will continue for several months.

The development of these boards has been seen as a major success within the collaboration. The AHCAL has decided to use the same boards for its readout and is currently placing an order with RAL to produce another seven of them.

2.1.2 Data acquisition software

The UK groups also provided the software data acquisition (DAQ) system for the whole of the CALICE readout.

The aim of the CALICE beam tests is to get around 10^6 events for each configuration of beam energy, particle, type, angle of incidence, etc. This implies a complete sample of order 10^8 events. To get this in a reasonable time, the DAQ system should achieve an average of 100 Hz readout rate. Given the spill structure at FNAL, this means an instantaneous rate during a spill of at least 1 KHz. There will be no threshold suppression applied online as studies of pedestal stability and noise are fundamental to this work. Each event is therefore expected to be around 50 kBytes in size, giving a total raw data volume of around 5 TBytes.

The DAQ was designed from scratch to allow a very lightweight, fast system to be developed which was tailored to the needs of CALICE and so achieve the high rates needed. It was implemented in C++ on Linux and uses the Hardware Access Library software from CERN [4] and ROOT [5].

The DAQ system has been used for all the UK board tests so far and is currently running in the cosmics tests in Paris. The trigger rate here is limited by the physical rate of cosmics to around 1 Hz. However, in calibration runs, readout of the two production boards at rates of 150 Hz have been achieved. This implies a sustained rate with six boards of around 50 Hz would be possible now. With further firmware and software developments this is expected to be improved. The biggest limiting factor currently, however, is the CPU power of the PC being used; it is an old 400 MHz machine which has been borrowed from the Imperial group. The purchase of a fast, multi-processor DAQ PC was included in the original bid but has been postponed until now in order to get the best possible system for the available money. This new PC will be bought before the end of the FY and will give large increase in CPU speed. Extrapolating from the rates observed now, the 100 Hz requirement should be straightforward to achieve.

The software is still under development. The ECAL tests performed so far only required a single VME crate system, which allowed synchronisation relatively simply. However, when the AHCAL tests begin later this year, a multi-crate system will be needed, which will most likely be implemented using multiple readout PCs. The DAQ has been written to allow a straightforward extension to multi-PC operation, where the trigger itself is used as the synchronisation marker and unsynchronised inter-PC communication and control uses socket calls. This additional functionality exists but has not been extensively used and so will need to be thoroughly tested over the next six months.

2.2 Status of the UK simulation effort

The PPRP, when they first considered the UK groups' application to join CALICE in 2002, strongly encouraged a vigorous participation in various areas of simulation and other software, to complement our electronics contribution to the prototypes. We report here on some of the activities in this area. The PPRP recognised the importance of this field by awarding an RA post to Cambridge (filled in August 2003 by G. Mavromanolakis, who has contributed strongly to the simulation work and preparations for the test beam). We were also fortunate to attract the services of a PPARC Fellow (C.G. Ainsley) who has made a great contribution to the shower reconstruction problem.

2.2.1 Hadronic shower modelling

Apart from testing the hardware concepts, one of the most important goals of the CALICE test beam campaign is to gather data on the properties of showers, especially hadronic showers, in highly granular calorimeters. We therefore embarked on a program of systematic comparison between different shower packages.

The standard Monte Carlo program for simulation of the Calice calorimeters (both for the test beam, and for the full ILC detector) is `Mokka`, which is based on `Geant4`. `Geant4` provides a toolkit approach, whereby a variety of different interaction models can be combined in different energy ranges. The `Geant4` authors provide a number of standard packages combining these models which are tailored to suit different "use cases". We have examined all those which are considered appropriate for high energy calorimetry. In addition, `Mokka` provides a facility for writing out a FORTRAN description of the detector geometry suitable for inclusion in a `Geant3` program. In this way we gain access to the physics models available in `Geant3`, namely `Gheisha`, `Fluka`, `Gcalor` and `Micap`. However, the old version of `Fluka` interfaced to `Geant3` is now deprecated by the authors, and the current version is not yet interfaced to `Geant4`. In order to gain access to this interesting model, we have employed a package called `Flugg`, which provides an interface between the `Geant4` geometry and materials, and the physics and transportation code of `Fluka`.

In total, seventeen different models or combinations of models have been studied, at different energies, for different particle species, and for two different HCAL detector technologies: scintillating tiles operated in analogue mode and RPCs operating in digital mode. Fig 2.2.1 shows some typical comparisons between these models for the case of 10 GeV π^- at normal incidence to the prototype calorimeter. In all cases the results have been normalised to the "LHEP" model which is the default in `Mokka`. A few of the significant findings are these:

- The response of the ECAL varies by up to $\sim 5\%$ between different hadronic models in `Geant4`. The response in `Geant3` is systematically higher, but this is likely to be correlated with similar differences in the electromagnetic response seen for electron showers in the ECAL (but not for the HCAL). This is under study, but is thought largely to arise from the different algorithms used to treat multiple Coulomb scattering in each model.
- The response of the HCAL varies considerably more from model to model.
- The scintillator and RPC HCAL technologies show different levels of sensitivity to the hadronic model, with the scintillator being more sensitive. This seems largely to result from the interaction of low energy neutrons in the hydrogen-rich scintillator.
- The `Flugg` model shows a significantly greater ECAL response from the `Geant` models, both for hadrons and electrons; this is still under investigation, but is suspected to be associated with issues such as tracking cutoffs and precision.
- The shower width is an important determinant of the spatial resolution between showers, and we see sizeable differences between models of up to 50%. The models which use the Bertini cascade model (which is favoured by many experts) predict particularly broad showers. The scintillator HCAL is particularly sensitive to the modelling of low energy neutron transport.

All in all, these studies reinforce the need for appropriate data against which to validate these models. It may well turn out that none of the models accords with data, but this will either allow us to tune models, or to assign realistic systematic uncertainties to their predictions when optimising the final detector design.

2.2.2 Pattern Recognition

The primary motivation for a highly granular calorimeter design for an ILC detector is to separate the energy depositions of different particles within hadronic jets. This permits the use of a "particle flow" algorithm, in which the charged particles will be measured using the tracking chambers, and the neutrals by appropriate combinations of ECAL and HCAL information. Experience from LEP, and early studies for the Tesla TDR, showed that this was the most promising way to achieve the $\sim 30\%/\sqrt{E}$ jet energy resolution which is demanded by linear collider physics.

Such energy flow algorithms have already been written for early LC studies. However, these codes tend to be tied to specific simulation packages and detector geometries. In view of the need to optimise global detector designs for cost and performance, it is clearly desirable to have a flexible algorithm which can readily be adapted to new setups.

A necessary first stage is to develop a versatile and robust calorimeter reconstruction algorithm. Inspection of simple event displays in the CALICE calorimeters shows that, with the 1 cm^2 cells and fine longitudinal sampling envisaged, substructures like tracks or small clusters are commonly observed within showers. This suggests that the optimal clustering algorithm for such a calorimeter may not be a simple "merge contiguous hits" algorithm. Within the UK we have pursued two complementary approaches to these problems. These have been presented at

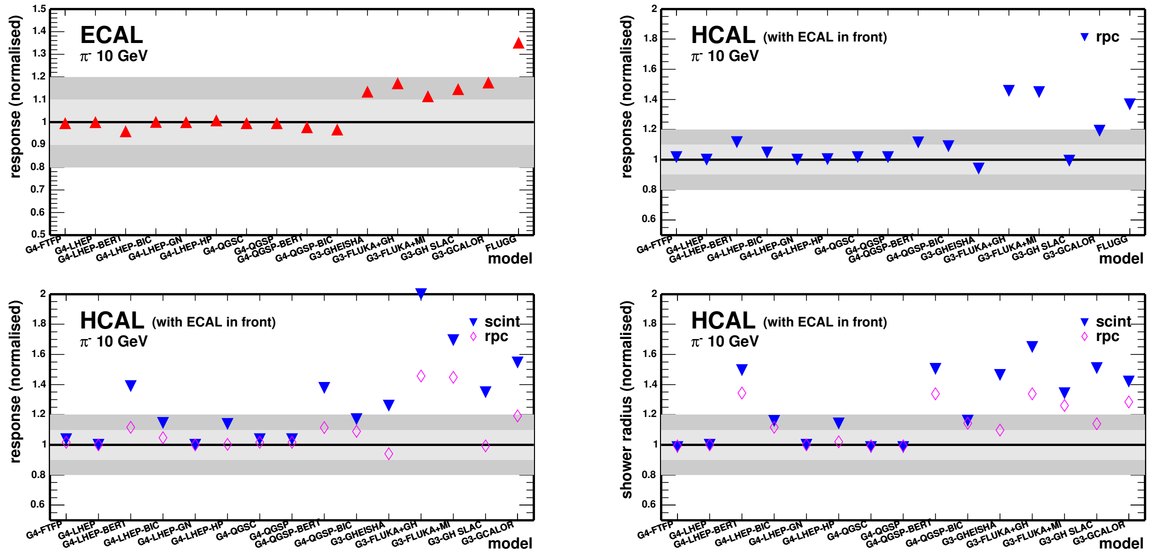

Figure 1: *Comparisons between different hadronic models for 10 GeV π^-. In each case, the results are normalised to the predictions of the LHEP model in* `Geant4`*. The upper left figure shows the response (i.e. mean energy deposited) in the Ecal; upper right the response (number of cells hit) in the RPC DHcal; lower left compares the tile AHcal response with the DHcal; lower right compares the shower radius in the HCal. The shaded bands denote 10% and 20% differences.*

several CALICE meetings and Linear Collider workshops, and been well received. Both algorithms are interfaced to the LCIO data format, which is becoming the agreed standard for ILC studies. In this way, it should be straightforward for others to use our code, and to integrate it with tools (such as tracking) developed elsewhere.

The first algorithm takes a tracking approach. The calorimeter is treated in concentric layers. Seed clusters are formed in the first layer of the calorimeter, and then hits in each layer travelling outwards are considered in turn. Each hit is compared with extrapolations of hits in previous layers, taking account of knowledge of the direction of the cluster from previous layers, and the best match found. If no acceptable match is found, a new cluster is seeded. In coding the algorithm, care is taken to encapsulate geometrical information and cuts in a single place. The algorithm was developed and tested on a variety of different samples – single particles, pairs of particles, and physics events such as hadronic Z^0 and W^+W^- events – and geometries. The performance seems to be quite robust. An example of the performance is given in Fig. 2.2.2. We see that photons can already be well separated from hadrons for separations above $\sim 3-5$ cm, while the separation power is a little worse for neutral hadrons. The separation power has been found to be significantly degraded if the Bertini hadronic model, which yielded wider showers (see Fig. 2.2.1), is used.

The second approach to clustering, based on Minimal Spanning Trees (MST), is quite different. Each calorimeter hit is regarded as a node in a tree, and the MST represents the way of connecting all the nodes, with no loops, which minimises the sum of the "lengths" between nodes. Geometrical information only occurs in the metric which defines the "length", which does not have to be the geometrical distance. There are standard algorithms which efficiently compute the MST. This effectively forms the whole event into one cluster; an algorithm is then applied to form clusters by cutting the tree.

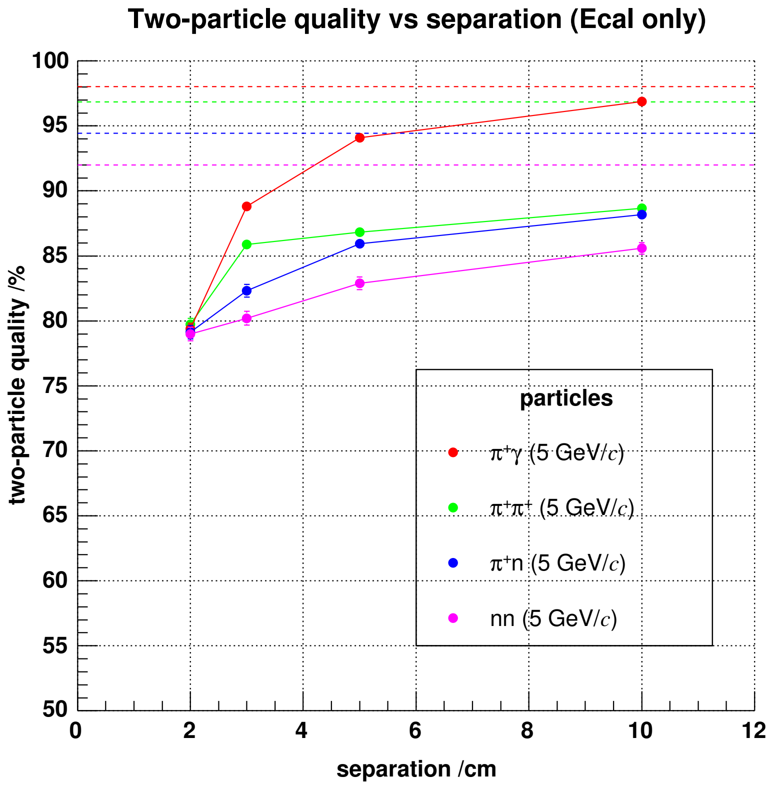

Figure 2: *Reconstruction Quality (defined as the fraction of energy which has a one-to-one relation between truth and reconstruction) for various pairs of 5 GeV particles, as a function of separation at the front of the calorimeter. The dashed lines indicate the quality for single particles.*

2.2.3 Preparation for test beam

The simulation studies, comparing different models, have had a direct bearing on the test beam proposal for CALICE. For example, we have examined the statistical precision required to distinguish between models, the particle species and energy regimes required. We have also performed studies of the possible usefulness of exposing the Ecal alone to a hadron beam.

In addition we are contributing strongly to the software infrastructure for analysis of the test beam data. Obviously, because of our responsibilities for the DAQ, we are in a prime position to understand the data. We are also preparing the code for converting the raw data from the DAQ into LCIO format for subsequent analysis and comparison with simulation.

2.3 General contributions

D.R. Ward is the now the Software Coordinator for CALICE. In addition to their work directed towards the CALICE test beam, some UK members of CALICE are contributing to generic international software activities for the ILC. For example, D.R. Ward is a convener for the Simulation group (covering simulation and reconstruction tools) for the ECFA LC study, and also for the worldwide LCWS workshop in Paris (April 2004), while M.A. Thomson is a convener of the Global Detector Performance group for the ECFA study.

3 Workpackage 1: Completion of CALICE test beam program

In our original 2002 application to the PPRP, we requested three years funding to support the construction of the CALICE ECAL prototype, and its exposure to test beams with a variety of HCAL prototypes behind. In this work package, we are seeking the third (AND FOURTH?) year's funding, which essentially consists of travel funds to participate in test beam running and related CALICE meetings, and also the third year of the Cambridge RA post so that the UK simulation expertise built over the past two years can be fully exploited in the analysis of the data.

The simulation work performed in the UK has contributed significantly to the planning for CALICE test beam running. Some of the key points are as follows:

- Our broad objective is to characterise the properties of hadronic showers to 1% precision, which therefore demands samples of at least 10^4 events. Monte Carlo studies have shown that significant differences between models are seen with such a level of statistics.
- Differences between the predictions of models have been seen over a wide energy of energies, from 1 to 50 GeV. This is therefore the range of hadron energies we would wish to study, which covers the range of typical energies of hadrons in jets at the ILC.
- The differences between models are different for baryons and mesons. It is therefore desirable to obtain data with both proton and pion beams. We have checked that the effects seen in Monte Carlo for neutral hadrons are similar to those for charged hadrons.
- We shall wish to perform position scans across the calorimeters (probably by exploiting the natural width of the beam) and to record data at several angles of incidence in the range 0° to 45°.

This work package divides naturally into two phases – tests of the ECAL alone in an electron beam, and combined tests with the CALICE Hcal prototypes in hadron beams.

3.1 Task 1.1: Analysis of DESY test beam data

The CALICE ECAL prototype is being moved to an electron test beam in DESY in January 2005. We expect the complete prototype to be in place and fully equipped by Easter 2005. During this period, we shall record electron data with energies in the range 1–6 GeV. This phase of operation will be of great importance in commissioning the hardware, and in starting to understand the problems associated with real data.

The actual data taking at DESY should be largely completed within the currently funded CALICE program, though it may be that some running will continue after April 2005. However, the data analysis will clearly continue through much of 2005. The electron beam data, together with cosmic muons, will be of great importance in understanding the detector response, and in tuning the Monte Carlo to describe it. The current Monte Carlo describes the raw $\mathrm{d}E/\mathrm{d}x$ energy deposits in the silicon pads, but clearly the UK groups will have a key rôle to play in adding realistic simulation of electronics effects such as noise and crosstalk. First versions of this exist, but will require tuning and adjustment in the light of experience.

The electron data will also be of importance in tuning parameters of the Monte Carlo simulations. Our early studies have shown that the detector response is somewhat sensitive to details of the particle transport (such as tracking cutoffs or the precision with which multiple scattering is followed); these details can easily change the response by a few percent. Before assessing the hadronic response, it will be important to tune the simulations to model the electromagnetic response in data.

Work Plan

The work required may be summarised in the following tasks:

1.1.0 We assume that the generation of preliminary Monte Carlo samples, the establishment of an analysis framework, and first comparisons between data and Monte Carlo will already have been done by the end of FY'04.

1.1.1 Work will be needed to understand the properties of the DESY beam, to implement a realistic simulation of the beam, and of the beam defining detectors (MWPCs).

1.1.2 Work will be needed to understand and tune the detector response in the light of real data. In particular the UK groups will be needed to simulate the behaviour of the electronics.

1.1.3 Generation of sizeable Monte Carlo samples will then be able to begin. This will use RAL CSF as well as cpu resources in the Universities. Work is also starting at Imperial to port the simulation code to the Grid.

1.1.4 Data and Monte Carlo comparisons will be carried on throughout. The aim is to ensure that both `Geant3` and `Geant4` and probably `Fluka` are all tuned so as to be capable of describing the electromagnetic response (electrons and muons) of the CALICE ECAL.

1.1.5 The results of this work should be published, probably along with a technical description of the detector.

Gantt chart

DESY test beam	**FY'05**				**FY'06**				**FY'07**			
Quarter	1	2	3	4	1	2	3	4	1	2	3	4
1.1.1 Beam simulation	==											
1.1.2 Tune detector simulation	==											
1.1.3 Large MC samples	==	==										
1.1.4 Data/MC comparisons	==	==	==									
1.1.5 Publish results			==	==								

3.2 Task 1.2: Analysis of hadron test beam data

CALICE has applied to Fermilab for time on the MTBF (Meson Beam Test Facility) test beam. This is a general purpose facility which can provide a proton beam of 120 GeV and secondary hadron beams in the energy range $\sim 5 - 80$ GeV which matches our needs well. The beam will contain both pions and protons, and is equipped with Čerenkov counters for particle identification, and MWPCs for beam definition. Electrons should also be available.

Our present expectation is that we shall expose the ECAL prototype in the Fermilab beam in summer/autumn 2005. For this run, it will be placed in front of the analogue Hcal prototype (iron, instrumented with scintillating tiles, 3×3 cm^2 in the shower centre) and a tail catcher (iron-scintillator), so as to provide complete instrumentation of hadronic showers. For this configuration, we shall wish to record samples of data, typically of order 10^6 events, over a wide range of energies, for example 10-15 energy points ranging from 5 to 50 GeV. In addition, data at non-normal incidence should be recorded at a few energies.

As part of the preparations for the test beam, we have also performed studies which demonstrate that, in the unlikely event that there are problems or delays with the Hcal prototypes, an exposure of the ECAL prototype to a hadron beam (particularly if additional absorbers were placed in front) would yield valuable information, especially for low energy hadrons. It is also important, of course, to record data on the performance of the Si-W ECAL in higher energy electron beams than available at DESY.

A key part of the CALICE program is to compare different detector technologies for the Hcal. Specifically, RPCs and GEMs with digital readout in cells of 1×1 cm^2 are being designed. The main limitation on this is funding for the RPC electronics (which would also be used for thr GEMs), and for this reason we envisage that the digital Hcal tests will not occur until 2006. We shall have a clear responsibility to keep our DAQ and electronics working throughout the test beam program, and an obvious intellectual interest in contributing to the interpretation of the results, using our expertise in the simulation area. We can therefore envisage some limited activity continuing well into 2007.

Work Plan
The work required may be summarised in the following tasks:

1.2.1 We shall need to maintain our capability of using the `Geant3`, `Geant4` and `Fluka` hadronic codes, and also of using relevant new models which become available in `Geant4`.

1.2.2 Work will be needed to understand the properties of the hadron beam, and how to use the beam detectors for particle identification and definition of the beam position. We shall need to implement a realistic simulation of the beam, and of the beam defining detectors (MWPCs and Čerenkovs).

1.2.3 Generation of sizeable Monte Carlo samples will then be able to begin. This will use RAL CSF as well as cpu resources in the Universities and on the Grid.

1.2.4 Comparison of data and Monte Carlo, for many different hadronic models. Tuning of beam simulation and detector response as required, and iteration. This process will clearly start as soon as data are recorded, with the aim of identifying those models which provide an adequate description of data.

1.2.5 Repeat for digital Hcal.

1.2.6 Publish findings on the properties of hadronic showers, and on comparisons with models.

Gantt chart

Hadron test beam	**FY'05**				**FY'06**				**FY'07**			
Quarter	1	2	3	4	1	2	3	4	1	2	3	4
1.2.1 Maintain hadron codes	==	==	==	==	==	==	==	==	==	==		
1.2.2 Understand beam line		==	==									
1.2.3 Generate MC samples		==	==	==								
1.2.4 Compare data/MC		==	==	==	==	==						
1.2.5 Repeat for AHcal						==	==	==	==			
1.2.6 Publish results									==	==		

3.3 Task 1.3: DAQ preparations for HCAL beam tests

Software effort for DAQ, firmware effort for CRCs; from Imperial.
NEEDS TO BE ADDED.

3.4 Summary

Deliverables

1. Published analysis of electron beam data

2. Published analyses of hadron beam data

Milestones

Effort
Need input from each group on names, fractions of time, and what they expect to do.

Name	Institution	Type	FTE '05	FTE '06	FTE '07	Funding
D.R.Ward	Cambridge	Academic	0.5	0.4	0.0	HEFCE
M.A.Thomson	Cambridge	Academic	0.x	0.x	0.0	HEFCE
C.G.Ainsley	Cambridge	Fellow	0.x	0.0	0.0	PPARC
G.Mavromanolakis	Cambridge	RA	1.0	1.0	0.0	PPARC
N.K.Watson	Birmingham	Academic	0.x	0.x	0.0	HEFCE
C.M.Hawkes	Birmingham	Academic	0.x	0.x	0.0	HEFCE
P.Dauncey	Imperial	Academic	0.1	0.1	0.0	HEFCE
D.Bowerman	Imperial	Fellow	0.2	0.0	0.0	PPARC
New Imperial RA	Imperial	RA	0.8	0.2	0.0	PPARC
O.Zorba	Imperial	Engineer	0.2	0.0	0.0	PPARC
M.Wing	UCL	Academic	0.x	0.x	0.0	HEFCE
R.J.Barlow	Manchester	Academic	0.x	0.x	0.0	HEFCE
etc etc			0.x	0.x	0.0	

Do we need to summarise the rôles of each of these people?
Total costs to PPARC

Item	FY05/6	FY06/7	FY07/8	Total
Equipment	0	0	0	0
Effort	50	50	0	100
Travel	70	70	0	140
Total	120	120	0	240

Risk assessment
Requested new Costs
We request the continuation of the RA post in Cambridge for FY'05. We seek travel costs for FY'05 and FY'06. *Any request for M&O or consumables?*

4 Workpackage 2: Data acquisition

The design [6] for a calorimeter for future ILC poses challenges to the data acquisition system (DAQ) mainly due to the large number of channels to be read out. A generic scenario for the DAQ system is as follows. Mechanically, the calorimeter will consist of 6000 slabs, of length 1.5 metres, each containing about 4000 silicon diode pads of $1{\times}1$ cm^2, giving a total of 24 million pads. At the very front end (VFE), ASIC chips will be mounted underneath 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 [7] of such chips will each contain 32 channels, although this may be higher in the final calorimeter. The data will be digitised in the ASIC and the data transferred to the front-end (FE) electronics which are placed in the detector at the end of the slab. It is expected that a zero suppression will be done at the FE in the FPGAs to significantly reduce the rate. The data will then be transferred off the detector, probably via a switch, to a receiver which would be PCI cards in a PC farm.

This proposal highlights areas of R&D where bottlenecks may arise in such a DAQ system. Alternative designs of a DAQ system which could affect the layout of the final detector or functionality of components are also considered. The proposal is split into three sub-workpackages, each of which has a few distinct tasks. These workpackages investigate the different 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. An under-pinning thread here is to strongly attempt to make use of commercial components and identify any problems with this approach.

This programme will allow us to continue aiding development of new technologies by providing the data acquisition system. We will also write a chapter in the future technical design report on the data acquisition system for a calorimeter. This will place us in lead position to build the DAQ system for future large-scale prototype calorimeters and the final system.

If we assume the parameters of 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 μs. The bunch train period is 250 ms, giving a duty factor between trains of about 0.35%. In a shower, up to 100 particles/mm^2 can be expected, which in a 1×1 cm^2 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 $24\cdot10^6 \times 4886 \times 2 = 250$ GB 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 μs giving a rate out of the ASIC of 0.4-3 GBytes/s, which we take to be 1 GB/s from now on. Details of rates after the VFE depend on the assumptions and system layout and will be discussed for each individual case where necessary.

4.1 DAQ1: Connection from the VFE to the FE

If threshold suppression can be done in an ASIC, the rates to the FE would be reduced by two orders of magnitude. Current designs cannot do this and may not even be desirable or possible, so we should assume the threshold suppression will be done elsewhere and 1GB/s has to be transferred. 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 [8].

4.1.1 DAQ1.1: Readout for calorimeter prototypes

A new VFE ASIC chip is currently being designed in Orsay, France. In the next 1-2 years, they expect to have a version 0 of such a chip with 32 channels and an ADC incorporated. We propose to build a circuit board to hold 16 chips, thereby reading out about 500 channels. The PCB will be designed and layed out at Imperial College where they will also write the required firmware with the appointment of a new Research Associate. Imperial College, in collaboration with the French groups, will then perform tests of reading out all chips. The PCB will be connected to a development board in a PC. This UK activity is obviously closely coupled to the work of the French design team so is anticipated to be concentrated towards the end of the second financial year. The same set-up will be used to read-out the MAPS system as detailed in Section 5, thereby consolidating costs and effort.

For the final calorimeter, the DAQ should ideally be the same for the ECAL and HCAL, but CALICE-UK has so far concentrated on only the ECAL. Our proposals for R&D contained in this document are generic such that both calorimeter sections should be able to converge to use the DAQ system we design. Although it is unrealistic for us to provide the hardware electronics for the HCAL beam tests in 2005-7, we will set up discussions to advise on our design and methods. Therefore, we will collaborate with groups who are building the HCAL prototypes such that we can direct the DAQ system to be a common system in the future. This will be an effective paper study between UCL and the HCAL groups with the only costs being occasional travel to collaborating institutes and test-beams.

4.1.2 DAQ1.2: Study of data paths over 1.5 m slab

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 about 1GB/s of data over 1.5 metres is a challenge. One could think of doing it conventionally using copper but here the effects of noise and interference would have to be considered. There is also the possibility of doing it using optical fibres although here there are also design considerations. The size of connectors would have to be investigated as the space at the VFE is of the order of mm. This would also involve understanding the power needed to transmit light out. This work also ties in closely with the mechanical and thermal aspects of the design as discussed in Section 6.

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. Hardware for bit error rate testing could be re-used from the ATLAS design. CAD tools would also be required particularly in the initial phase. Cambridge and Manchester Universities will work together on this project, sharing costs and effort. This would be a project which would run over the three years of the proposal.

4.2 DAQ2: Connection from on to off detector

In the standard detector configuration, communication from the VFE will be via the electronics at the FE to an off-detector receiver. As a threshold suppression will be done at the FE, the rates 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 and each sample above threshold needs channel and timing label, the total data volume to be read out from the calorimeter is about 5 GBytes or about 1 MBytes per slab. This data has to be read out within a bunch train period of 250 ms, giving a rate of 5 MBytes/s. This is significantly slower than the 1 GBytes/s from the VFE.

In the following, we consider two scenarios: the first is the standard detector configuration and the other is where the FE is removed and the communication is directly from the VFE to the off-detector. The off-detector receiver, as described later is assumed to be PCI cards housed in PCs. As PCs are not 100% reliable, a switch will be necessary to send data efficiently to working PCs. The first scenario would require less high-speed equipment off the detector, whereas the latter would require many optical fibres with dedicated optical switching. The second 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. This would also reduce the number of processing components within the detector which would be inaccessible for long periods of time.

The need for FE electronics is also questionable if more is done on the ASIC chip, such as zero suppression. In such a scenario, transporting the data directly from the ASIC off the detector could be done as the FE would become redundant. This again highlights the need to consider having no FE electronics and thereby radically changing the design of the detector. Our research in this area will then provide feedback to groups designing the ASIC chips.

4.2.1 DAQ2.1: Connection from FE to off-detector receiver via network switch

We propose to setup a test system to validate the effectiveness of high-rate switching. This will investigate the issues in directing multiple 10 Gigabit data flows with the linear collider time structure to high-performance processors. The test system 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 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 and motherboards will be characterised considering jitter, latency, throughput, etc.. The network switch will also be characterised. Then simple data flows will be examined for different scenarios. Finally the data flow will be examined for rates expected from the linear collider's accelerator timing structure.

The cost of the system would include a 10 Gbit ethernet card with four ports, four optical data transceivers and the six 1 Gbit ethernet cards, costing in total £ 23000. This sub-workpackage is to be led by Manchester with input from UCL who will lead the alternative system. UCL will provide the FPGA to fibre communication interface, related to the clock and control package discussed later.

4.2.2 DAQ2.2: Connection from VFE to off-detector receiver via optical switch

We propose to setup a test system to validate that fibres can transport the necessary data rates from the VFE to the off-detector receiver and an optical ("layer-1") switch can re-route data with high efficiency. This would contain a bundle of fibres coming into a switch from an effective VFE. A bundle would the come out of the switch to the off-detector receiver. The off-detector receiver, described in more detail in the next section, 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. The optical switch efficiency would then be tested when PCs fail and the data has to be swapped to another PC. Also if a PC contains a large amount of data which could be of interest and takes up a large amount of processing time, a busy signal would be sent to the optical switch 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 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 μm, or with cladding 250 μm 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 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 so needs to be considered now when changes in its general structure could be considered.

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$\times$16 lane switch. Such a switch currently costs about £ 12500. Other hardware costs would be covered in building the PCI card in another sub-workpackage. As part of the bid, Nicola Pezzi (UCL) would be funded at a 10-15% level for three years. Pezzi brings the necessary networking expertise to the project. This sub-workpackage is to be led by UCL with input from Manchester who will lead the alternative system. To consolidate costs, equipment would be shared between Manchester and UCL for this and the previous sub-workpackage. This then sets constraints on the time-line of the two projects which will run serially.

4.2.3 DAQ2.3: Failsafe FPGA

In a scenario requiring an FPGA on-detector at the front end, it is imperative that the device can be re-configured remotely. This stems both from the number FPGAs, but also the relative inaccessibility of the devices. To facilitate this some sort of non-volatile memory that boots the device into a communicative state, or an external control system that provides an interface between the unconfigured FPGA and the outside world is required.

Using devices with non-volatile memories built-in may have advantages here, but increased cost and limited re-programming cycles make these less attractive. Questions surrounding the reliability in the detector environment of these memories over the lifetime of an experiment also need to be addressed. In contrast modern SRAM based FPGAs boast that they can re-configured infinitely and can therefore be 'refreshed' at regular intervals to ensure their configuration is not corrupted. The option of running with different firmware for different types of run is also attractive.

Reducing the number of components on the front-end-module would be advantageous from a cost, reliability and dead material point of view. To this end we propose a task of researching a means of connecting a commercial network receiver package to an FPGA with the absolute minimum of components, so as to reliably reset and configure the device remotely. We will also consider a literature survey of the response of various FPGAs to radiation and rates of single event upsets and how to circumvent them.

The test system will be simple and require a PC with a network card and a development board with cables and connectors. This work is closely connected to the clock and control in the next section and will be carried out at UCL.

4.2.4 DAQ2.4: Transport of clock and control and configuration data

To ensure front-end detector electronics captures data from actual collisions all components in a detector system need to be synchronised to a bunch-crossing clock. Although a 'trigger-less' system is envisaged it is still useful to consider a signal that moderates read-out, such as 'bunch-train-start' signal.

One of the systems in this proposal make use of standard network protocols to tag data at the front-end such that a commercial network can deliver it to a waiting event-builder PC. To co-ordinate this requires that busy information from the event-builder system is used to derive available data destinations that needs to be sent to the front-end. This can be called a trigger ID at some level. Although the need to build a large scale system is still some way off, the means of interfacing these time-dependent signals into the more asynchronous commercial network arena is not understood. The aim of this task is to investigate means of distributing clock and control signals using commercial networking equipment and components, as opposed to a more traditional bespoke system.

This couples well with the previous task of investigating the FPGA failsafe in that it attempts to minimise components and connections on the front-end-module. In many ways this will be an investigation to verify that a single clock, control and configuration system can be constructed using a commercial hardware for both receiver and transmitter hardware. Although it is envisaged that fibre optic will be the dominant communications technology in a final design, lower speed copper based systems are sufficiently similar to allow the first stage of development to use this technology.

The test system will build on the previous FPGA failsafe task and use much of the same equipment. Additional network cards, a fanout board and switch will be required. As stated before, this work will be undertaken at UCL, with a new Research Associate contributing along with Academic and Engineering staff.

4.3 DAQ3: Off-detector receiver

An ideal system would have all data from the detector sent directly to one PC where full event reconstruction is performed. Given the over 20 million channels in the calorimeter, this seems unfeasible. However, it is desirable to get as much of the calorimeter into one processor as possible. This would allow for local clustering which could be used as a seed to full event reconstruction offline. We propose a system which consists of PCI cards mounted in PCs which receive the data and perform local clustering. The fundamental questions to be answered is 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.

4.3.1 DAQ3.1: Prototype off-detector receiver

The off-detector receiver will consist of a set of PCI cards in PCs. The PCI cards will be part of both of the network systems, being configurable to act as a receiver for either. In a test system, they would also act as the VFE or FE sending data off detector. Recent developments in PCI cards have led to high bandwidth which is scalable due to future increases in the number of lanes using PCI Express technology. We will use cards based on this technology using a PCI Express bus. We will design and build our own PCI card to maximise flexibility for the tests we want to perform. In conjunction with the network workpackages discussed previously, tests will be made to see how efficiently data is transferred if a PC is unreachable or busy. Although it is hard to predict what technology will be around for the final system, gaining experience of cutting-edge technologies now, such as PCI Express, will prove valuable now for estimates of how much data can be received by one computer. So, whilst not trying to absolutely maximise the amount of data which can be transferred into a PC in the test system, the model should be scalable.

The final system would consist of 4 PCs each with a PCI card in them and the possibility to move PCI cards between PCs so as to maximise the data rate to one computer. To act as the VFE or FE sending the data, 4 further PCI cards would be required housed in two PCs. Assuming 4 prototypes at £ 2500 per board, £ 1500 per board for the final eight and six PCs each costing £ 1200, the total price comes to about £ 29200. The design and build of the PCI card would be undertaken by Cambridge, UCL (with the new RA contributing) and RHUL (with new RA contributing) in conjunction with Manchester and UCL's work on networking. The PCI cards and PCs would be shared by Manchester and UCL for the network system tests.

4.3.2 DAQ3.2: Local clustering on PCI cards and simulation

With sufficient data transported to one computer local clustering and removal of isolated hits can be performed before full event reconstruction is done offline. This would involved firmware programming of the FPGA on the PCI card, its simulation and subsequent effect on full reconstruction and performance in a physics analysis. This would be undertaken by Cambridge and UCL (lead by the new RA) and would have a large overlap and close collaboration with their interests in the simulation workpackage (see Section 7).

4.4 Summary

Milestones

DAQ1.1: Readout for current calorimeter designs

1. Design of PCB and development board for reading out 16 chips.
2. Layout of boards.
3. Write firmware.
4. Readout prototype chip.
5. Readout 16 chips.
6. Collaboration with HCAL groups to form common system for DAQ.

DAQ1.2: Study of data paths over 1.5 m slab

1. Choice of electrical standard.
2. Clock and control distribution from FE to VFEs.
3. Build test boards to check validity for 1.5 m paths.
4. Data transmission as a function of number of VFEs.
5. Data transmission as a function of whether zero suppression is done on VFE.
6. Data transmission as a function of whether buffering done during bunch train.
7. Simulation of bandwidth and cross-talk.
8. Estimate of electrical pick up at pads.

DAQ2.1: Connection from FE to off-detector receiver via network switch

1. Investigation of technologies and test criteria
2. Set up 10 Gigabit test-bench.
3. Characterise the NICs and motherboards - throughput, latency, jitter, packet loss, cpu overhead, PCI and memory bus usage.
4. Characterise the switch.
5. Examine the behaviour of simple traffic flows, i.e one-to-one and multiple one-to-one, many-to-one (classic bottleneck) and mixtures of 1 Gbit and 10 Gbit flows.
6. Generate CALICE traffic pattern (860 μs burst every 250 ms), examine performance of directed flows and request-response data transfer methods.

DAQ2.2: Connection from VFE to off-detector receiver via optical switch

1. Investigation of technologies and test criteria.
2. Link Test. Setup point to point link between 2 PCs and investigate throughput, bit error rates, etc..
3. Understand switch configuration. Connect the 2 PCs above via the switch and verify link; next crudely control switching between 2 different receivers.

4. Busy/failover system. Setup means of generating a busy and automatically switching without data loss.
5. Explore synchronous switch-over. Can it be done in between bunch-trains and how reliably.
6. Explore broadcast and one-to-many options. Try configurations that send the same data to more than one receiver.

DAQ2.3: Failsafe FPGA

1. Become familiar with a modern FPGA reset/configuration implementation (e.g. Xilinx Virtex-4).
2. Investigation of FPGA radiation tolerance and rates of single event upsets.
3. Build a test system using signals generated from a digital signal source (e.g. another FPGA) to test reset and configuration protocols.
4. Explore capabilites of commercial networking devices.
5. Build a test system to interface a commercial networking devices to the signal required in 2) above.
6. Upgrade the system to use modern very-high speed network links (10 Gbit Ethernet).
7. Examine integration of such a system into a commercial network based DAQ system.

DAQ2.4: Transport of clock and control and configuration data

1. Evaluate the clocking requirements for a front-end-module in the front-end.
2. Investigate means of regenerating a clock from commercial network links. Build test module to do this with standard (100 Mbit) components.
3. Investigate means of extracting accurate trigger/spill timing signals.
4. Explore using commercial network equipment to distribute triggers. Build a small network, driven from a PC for testing.
5. Examine problems associated with scaling to a 10000 link system.
6. Look at cost effective ways of distributing time sensitive signals using a more proprietary system, still making use of commercial transmitters and receivers. Would involve building a small test board.
7. Make use of a proprietary (the CALICE PCI card) to generate these signals in a point-to-point or low-level fanout configuration.

DAQ3.1: Prototype off-detector receiver

1. Investigation of latest PCI technology.
2. Design and build PCI card probably based on PCI Express.
3. Test PCI card.
4. Fabricate batch of PCI cards and setup test system for prototype off-detector receiver.
5. Efficiency of data transfer when PCs unreachable or busy
6. Amount of data, i.e. how much of calorimeter, can be transported to a PC.

DAQ3.2: Local clustering on PCI cards and simulation

1. Program FPGAs to perform local clustering and do hit removal.
2. Simulate in Monte Carlo.
3. Effect of local clustering on final event reconstruction.
4. Effect of local clustering on physics process.
5. How much of calorimeter needs to be contained in a PC for clustering to be effective.

GANNT chart

DAQ Workpackage	FY 05/06				FY 06/07				FY 07/08			
	1	2	3	4	1	2	3	4	1	2	3	4
DAQ1.1: Readout for VFE prototype chip												
1. Design of PCB and Dev. board							=	=				
2. Layout of boards								=	=			
3. Write Firmware								=	=			
4. Readout prototype chip									=			
5. Readout 16 chips									=	=		
6. Collaboration with HCAL groups	=	=	=	=	=	=	=	=	=	=	=	=
DAQ1.2: Study data paths of 1.5 m slab												
1. Choice of electrical standard	=	=	=	=								
2. Clock and control from FE to VFEs			=	=	=	=						
3. Build test boards for 1.5 m paths			=	=	=	=			=	=	=	=
4. Data transmission as a function of number of VFEs					=	=			=	=		
5. Data transmission with zero suppression on VFE						=	=			=	=	
6. Data transmission with buffering during bunch train							=	=			=	=
7. Simulation of bandwidth and cross-talk					=	=	=	=	=	=	=	=
8. Estimate of electrical pick up at pads									=	=	=	=
DAQ2.1: Connection from FE to off detector												
1. Investigation of technologies and test criteria	=	=	=	=								
2. Set up 10 Gigabit test-bench				=	=	=						
3. Characterise the NICs and motherboards					=	=						
4. Characterise the switch						=	=					
5. Examine the behaviour of simple traffic flows							=	=				
6. Generate LC timing structure and examine performance							=	=	=			
DAQ2.2: Connection from VFE to off detector												
1. Investigation of technologies and test criteria			=	=	=	=						
2. Link Test							=	=				
3. Understand switch configuration								=	=			
4. Busy/failover system									=	=	=	
5. Explore synchronous switch-over										=	=	=
6. Explore broadcast and one-to-many options										=	=	=
DAQ2.3: Failsafe FPGA												
1. Familiarisation with FPGA configuration	=	=										
2. FPGA radiation tolerance and SEU rates	=	=	=	=								
3. Test reset and configuration protocols		=	=									
4. Explore commercial networking			=	=								
5. Test system linking signals and network				=	=							
6. Upgrade to high-speed networks					=	=						
7. Integration into DAQ system						=	=					

DAQ Workpackage	FY 05/06				FY 06/07				FY 07/08			
	1	2	3	4	1	2	3	4	1	2	3	4
DAQ2.4: Clock and control and configuration data												
1. Evaluate the clocking requirements					=	=						
2. Regeneration of clock from commercial network links						=	=					
3. Extracting trigger /spill timing signals							=	=				
4. Build small network for testing								=	=			
5. Problems scaling to a 10000 link system									=	=		
6. Proprietary systems on a small test board										=	=	
7. Test of system											=	=
DAQ3.1: Prototype off-detector receiver												
1. Investigation of latest PCI technology	=	=										
2. Design and build PCI card		=	=	=								
3. Test PCI card				=	=							
4. Fabricate batch of PCI cards and setup test system					=	=	=					
5. Efficiency of data transfer when PCs unreachable or busy							=	=	=	=		
6. How much of calorimeter can be transported to a PC										=	=	=
DAQ3.2: Local clustering on PCI cards												
1. Program FPGAs	=	=	=	=								
2. Simulate in Monte Carlo				=	=	=	=	=				
3. Effect on final event reconstruction							=	=	=	=		
4. Effect on physics process									=	=	=	=
5. How much of calorimeter needed											=	=

Deliverables

- Demonstration of functionality of new ASIC chips.
- Properties of data flow along 1.5 metre board.
- Most reliable data transportation and routing method off the detector.
- Need (or not) for front-end electronics.
- Method for rebooting and reconfiguring inaccessible FPGAs on the front-end electronics.
- Method of synchronisation of commercial standard components to an accelerator clock.
- Scalable model of off-detector receiver to answer how much data can be sent to one PC.
- Amount of data needed in one PC for local clustering to be effective.
- Write chapter on data acquisition for calorimeter for technical design report.

Effort

Name	Position	Institute	Funding	2005/6	2006/7	2007/8
M. Goodrick	Engineer	Cambridge	PPARC (RG)	0.2	0.5	0.6
ELECTRONICS	Technician	Cambridge	PPARC (RG)	0.2	0.5	1.0
MECHANICAL	Technician	Cambridge	PPARC (RG)	0.0	0.15	0.15
P. Dauncey	Academic	Imperial	HEFCE	0.1	0.1	0.1
O. Zorba	Engineer	Imperial	PPARC (RG)	0.0	0.1	0.1
D. Price	Engineer	Imperial	PPARC (RG)	0.0	0.2	0.2
Electronics	Technical	Imperial	HEFCE	0.1	0.5	0.6
New RA-1	RA	Imperial	PPARC (new)	0.0	0.2	0.2
R. Barlow	Academic	Manchester	HEFCE	?	?	?
R. Hughes-Jones	Engineer	Manchester	PPARC (RG)	?	?	?
S. Kolya	Engineer	Manchester	PPARC (RG)	?	?	?
G. Boorman	Engineer	RHUL	PPARC (RG)	0.0	0.1	0.2
B. Green	Engineer	RHUL	PPARC (RG)	0.2	0.2	0.2
M. Green	Academic	RHUL	HEFCE	0.2	0.2	0.2
New RA-2	RA	RHUL	PPARC (new)	0.5	0.5	0.5
M. Lancaster	Academic	UCL	HEFCE	0.2	0.2	0.2
N. Pezzi	Engineer	UCL	PPARC (new)	0.1	0.15	0.15
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-3	RA	UCL	PPARC (new)	0.6	0.6	0.6
Total FTE				3.1	5.5	5.1
Total new				1.2	1.75	1.25

Table 1: Table of FTE effort from all Universities

Costs per task

	2005/6	2006/7	2007/8
DAQ1.1: Hardware	£0 k	£9 k	£0 k
DAQ1.2: Hardware+CAD	£10 k	£12.5 k	£12.5 k
DAQ2.1: Hardware	£0 k	£23 k	£0 k
DAQ2.2: Hardware	£0 k	£12.5 k	£0 k
DAQ2.3: Hardware	£0 k	£5 k	£0 k
DAQ2.4: Hardware	£0 k	£0 k	£6 k
DAQ3.1: Hardware	£5 k	£17 k	£7.2 k
DAQ3.2: Hardware	£0 k	£0 k	£0 k
DAQ: Total equipment	£15 k	£78.5 k	£25.7 k

Table 2: DAQ: Total costs

Total costs to PPARC

Item	FY05/6	FY06/7	FY07/8	Total
Equipment	15	79	26	120
Effort	250	450	400	1100
Travel	10	14	16	40
Total	275	543	442	1260

5 Workpackage 3: Monolithic active pixel sensors

5.1 Introduction

All studies for a Si-W ECAL for the ILC currently consider diode pads as the sensitive elements in the silicon wafer. However, mainly outside the ILC community, work on monolithic active pixel sensors (MAPS) have developed this technology to a level of maturity so that they can now be considered as a possible alternative. The main goal of this workpackage is to study the use of MAPS for an ILC ECAL and, over the three year period of the grant, to either validate this concept or show that it is not a feasible option. If MAPS are a viable alternative to diode pads, then working to this timescale would allow the UK to contribute to the detector TDR in this area and would place MAPS as a technology to be considered where the UK would have clear leadership.

MAPS have been developed over the last decade for imaging purposes but their application to PPARC science really took off with the collaboration lead by R.Turchetta. This group has been supported by PPRP seedcorn funding for the last two years [9] and this sucessful project has attracted interest for tracking applications not only from the vertex detector groups involved in the ILC [?] but also for SLHC applications [?]. However, as far as we are aware, there is no other proposal to use MAPS in a calorimeter.

MAPS implement both the sensitive silicon detector and the readout electronics onto the same wafer, using standard commercial CMOS processes for fabrication. The actual charged particle detection is done using electron-hole creation in the epitaxial layer. This is of order $10\,\mu$m thick and is located immediately below the surface layer where standard CMOS circuitry can be implemented for readout. The charge liberated diffuses within the layer and is collected by contacts passing from the surface electronics to the epitaxial layer.

The basic concept which this workpackage will develop is to use MAPS with binary readout on the sensor pixels. The channel count is extremely high and so analogue readout would produce prohibitive amounts of data. By making the pixels small enough, the probability of more than one minimum ionising particle (MIP) passing through each pixel can be kept small. Hence, a simple bit to indicate the presence of a MIP is sufficient information. The data would then be stored during the beam train in memory implemented within the pixel and read out before the next train. This idea is described in more detail in the following subsection; a more complete description of the conceptual design of a MAPS ECAL can be found in [10].

5.2 MAPS concept

To be able to use binary readout for MAPS pixels, the probability of multiple charged particles per pixel must be small. The density of charged particles in the core of a 500 GeV electromagetic shower is around $100/\mathrm{mm}^2$, which implies pixel sizes of at most $100 \times 100\,\mu\mathrm{m}^2$, although an actual size of $50 \times 50\,\mu\mathrm{m}^2$ is more realistic. This results in a huge pixel count; for an ECAL the size of the one in the TELSA TDR [**?**], this would be around 10^{12} pixels. The only possible way to handle such a large number of channels is to do a huge amount of data processing directly at source and so build data suppression into the MAPS sensor itself. This can then reduce the data volume which needs to be read off the sensor to a managable level.

Since the MAPS sensors have similar mechanical and thermal properties to diode pad wafers, then the structure of the ECAL is somewhat independent of which would be used. However, since the MAPS actual measure different physical quantities and have no external readout electronics, there are several advantages to using them over diode pad wafers.

- Energy resolution. The proposed MAPS would basically measure the number of particles passing through the silicon, rather than the energy deposited as is measured by the diode pads. While these two would be expected to be related, the energy deposited is smeared out due to Landau fluctuations, non-normal incidence and the velocity dependence of the ionisation rate. Preliminary simulation studies show the higher quality MAPS information improves the energy resolution for electromagnetic showers from single electrons by around 20% [**?**].

- Pattern recognition. The granularity of the MAPS ECAL is much finer than for the diode pad option, by two orders of magnitude. The natural spread of the showers prevents this from having as dramatic an impact on the cluster reconstruction, but improvements on quantities such as two cluster separation are seen in simulation [**?**].

- Size. The absence of a VFE readout chip on the detector results in a significant reduction of the VFE PCB module thickness, potentially by around factor of two. This impacts in several areas. The gaps between the tungsten layers can be correspondingly reduced, giving a smaller shower spread between layers and hence decreasing the effective Molière radius. This contributes directly to improved pattern recognition between layers. In addition, the radial thickness of the ECAL can be reduced, by around 10-20% with a subsequent decrease is cost of not only the ECAL itself, but also the HCAL and solenoid magnet which are further out. This could be a significant saving; the magnet cost alone is estimated to change by around \$2M per cm radius [**?**].

- Thermal issues. The main heat source in the ECAL is in the readout electronics. For the diode pad option, that is concentrated in the small, packaged VFE chip. For the MAPS option, the heat production is spread evenly over the whole surface of the bare sensor, which has a larger area by something like a factor of 100. This makes a good thermal contact much easier for the MAPS case, reducing the problem of heat removal.

- Cost. The cost of a Si-W ECAL is large and is dominated by the surface area of silicon needed. Even with estimates of around $\$2/\mathrm{cm}^2$ for processed high resistivity silicon by 2010 (a factor of five below current prices), the silicon alone in the TESLA TDR ECL would cost \$60M. While the price of such wafers in several years' time is hard to predict, the relative cost of diode pad-type wafers compared with standard CMOS processes, as used for MAPS, is potentially less uncertain. Currently, ignoring NRE, high resistivity diode wafers cost around a factor of two more than CMOS. Thus, there is potentially a large reduction in the dominant cost of such an ECAL through using commercially standard, widespread fabrication methods.

However, despite these potential advantages, there are several issues which need to be studied before MAPS can be considered to be a viable alternative to diode pad wafers. The main ones are

- Signal/noise. To operate with binary readout, a good signal/noise ratio much be achieved. To achieve a high efficiency, as required for good pattern recognition, then the threshold must be set well below the most probably MIP level. Specifically, the actual deposited energy distribution can be as low as 70% of the most probably value. To be efficient, it is necessary to set the threshold at least two sigma below this level. However, to not get too many noise hits, then the threshold must be at least four or five sigma. This leads to a requirement of a signal/noise ratio of at least 10/1, which means a a noise hit probability of around 10^{-6}; a non-trivial requirement. The noise in MAPS tends to dominated by the reset; schemes allowing "soft reset" to minimise this are being studied. Given the assumed relatively long time of several hundred ns between bunch crossings and the low physics rate at the ILC, then an alternative approach would be to never reset but let the charge leak off with a characteristic time of the same order.

- Crosstalk. As stated above, the binary nature of the proposed MAPS readout gives a good measure of the number of particles passing through the silicon and improves the energy resolution. However, if each charged particle often resulted in more than one binary hit being recorded per layer, this potential improvement could be lost. Such crosstalk effects would result from inefficient charge collection in the epitaxial layer resulting in longer diffusion distances and hence charge being collected by neighbouring pixels. In addition, a thick epitaxial layer compared with the pixel size dimension would result in more charge sharing between pixels for tracks at non-normal incidence, as they would be more likely to cross the epitaxial layer beneath two pads. The design of the collection pads needs to be optimised to reduce the former effect. Minimising the epitaxial layer thickness will reduce the latter effect, but at a cost of directly reducing the signal size; the balance of these two will be a major part of this study.

- Uniformity and stability. The binary readout needs a threshold to determine if the bit should be set or not. There are far too many pixels to allow each threshold to be set separately for each pixel, so the MAPS must be made uniform enough at least so the pixels within each sensor can operate with a common threshold. In addition, there are likely to be thermal fluctuations as the electronics is powered on and off with the bunch trains. If the threshold is not stable with temperature then this would result in excess noise or loss of efficiency.

- Power. One of the major problems with the mechanical structure of a Si-W ECAL is how much cooling will be required and how it would be implemented; indeed the mechanical workpackage described in section 6 proposes that the UK becomes involved in these crucial studies. It is clear the MAPS option would need to be at least similar to the diode pad option in terms of total power dissipation to be taken seriously. Techniques to reduce the power consumption, such as powering on the comparators for the shortest possible time, will form part of the work programme.

The aim of this workpackage is to investigate all of these and evaluate whether MAPS sensors are feasible for this application.

5.3 Outline of proposed work

It is foreseen that two rounds of sensor fabrication will be required to get a design close to the specifications needed. The first design will test several different ideas within the same sensor, while the second fabrication round will produce uniform sensors iterating on the best design found in the first round. Each fabrication round will produce the minimum run quantity of six eight-inch wafers using $0.35\,\mu$m AMS technology and will cost around £60k per round.

The first round will produce small sensors, around $1 \times 1\,\mathrm{cm}^2$, which means around 40k pixels total per sensor. Six wafers would result in over a thousand such sensors if the yield was good. The sensor will contain pixels with around nine different designs, allowing comparisons of various ideas without needing a production run for each. It is estimated that the design of this first sensor will take around 13 months, with the fabrication of the sensors being completed 17 months into the grant period. The sensors will be tested electrically for correct functionality and then studied to measure their properties using sources, cosmics, a 5 T magnetic field, etc. Control, clocking and readout would be supplied by using a commerical FPGA development board connected directly to a PC. These tests will take the next six months to complete.

The iteration of the design for the second fabrication round will take around six months, much of which will happen in parallel with the tests of the first batch of sensors. The fabrication run itself should be completed by month 30. The sensors will be the maximum size allowed within standard CMOS processes without stitching, which is $2 \times 2\,\mathrm{cm}^2$. The sensors will have a uniform design for all pixels and the six wafers could produce up to 300 sensors, again if the yield is good. These sensors would again be subjected to similar tests to the first fabrication round.

However, in addition, it is planned to put several of these sensors into a test beam. If they are mounted on a PCB with the same dimensions as the CALICE ECAL wafer PCBs, then one or two layers of diode pad wafers could be exchanged for MAPS sensors within the tungsten structure used for the CALICE beam tests. This would allow a direct comparison of MAPS against diode pads. It would also allow some study of possible single-event unset effects resulting from dense electromagnetic showers passing through the MAPS sensor circuit. To cover the same surface area as a full VFE PCB, i.e. $12 \times 18\,\mathrm{cm}^2$, would require 54 working sensors per layer, so realistically at most two layers would be equipped in this way. The PCBs to hold and read out such a number of MAPS sensors would require significant design effort. The control and readout would also be more complicated than for the single sensor tests. However, the existing CALICE readout boards have the flexibility to do completely digital I/O, by setting jumpers to bypass the ADC inputs and connect directly to the FE FPGAs. Hence, with some firmware development, the same boards can be used for both the diode pads and MAPS readout.

5.4 Responsibilities and effort

The sensor designs will be done by a full-time engineer at RAL, supervised by R.Turchetta. RAL PPD personnel will perform device simulations to guide the design (TBC).

Initial tests will also be done at RAL but the main bulk of testing, including setting up the readout systems, will be done at Imperial and Birmingham. The bulk of this test work will be undertaken by a new RA at Imperial which is being requested as part of this proposal.

The test PCBs to mount individual MAPS from the two fabrication rounds will be designed, manufactured and assembled through the Imperial group. The design and layout for these PCBs will be done by Imperial personnel, which fabrication and assembly will be done commercially. Firmware and software for the readout of the test MAPS will be supplied by Imperial. A mechanical holder for the high magnetic field tests will be manufactured at Imperial.

The large PCB for the beam test will be designed by Birmingham, being laid out and manufactured commercially.

5.5 Milestones

- First fabrication round complete by month 18.
- First fabrication testing complete by month 30.
- Second fabrication round complete by month 30.
- Second fabrication testing complete by month 36.

5.6 Effort

Name	Position	Institute	Funding	FY05/6	FY06/7	FY07/8
N.K.Watson	Academic	Birmingham	HEFCE	0.1	0.2	0.2
J.Wilson	Academic	Birmingham	HEFCE	0.1	0.2	0.2
R.Staley	Engineer	Birmingham	PPARC/RG	0.0	0.2	0.2
Electronics support	Technical	Birmingham	PPARC/RG	0.0	0.1	0.2
P.D.Dauncey	Academic	Imperial	HEFCE	0.2	0.2	0.3
New Imperial RA	PGRA	Imperial	PPARC	0.2	0.5	0.5
O.Zorba	Engineer	Imperial	PPARC/RG	0.0	0.2	0.2
Electronics support	Technical	Imperial	HEFCE	0.0	0.1	0.1
Mechanics support	Technical	Imperial	PPARC/RG	0.0	0.1	0.1
R.Turchetta	Engineer	RAL ID	PPARC	0.1	0.1	0.1
Sensor designer	Engineer	RAL ID	PPARC	1.0	1.0	1.0
???	???	RAL PPD	PPARC	0.?	0.?	0.?

5.7 Costs

Equipment costs

Item	FY05/6	FY06/7	FY07/8
First MAPS fabrication round	0	60	0
Test PCB for first round	0	3	0
PC, DAQ electronics and miscellaneous test equipment	0	6	0
Second MAPS fabrication round	0	0	60
Test PCB for second round	0	0	3
Beam test PCB	0	0	15
Total	0	69	78

Effort costs to PPARC

Institute	New/RG	FY05/6	FY06/7	FY07/8
Birmingham	RG	0	20	25
Imperial	New	10	25	25
Imperial	RG	50	50	50
RAL ID	New	100	100	100
Total		160	195	200

Total costs to PPARC

Item	FY05/6	FY06/7	FY07/8	Total
Equipment	0	69	78	147
Effort	160	195	200	555
Travel	5	5	15	25
Total	165	269	293	727

5.8 Deliverables

The main deliverable is an evaluation of whether MAPS are a vaiable option for silicon-tungsten ILC ECAL.

If the concept is proved valid, then a further deliverable is a working MAPS pre-prototype for this application.

5.9 Risks

Fail of design within a fabrication round. This would incur a one year delay, resulting in non-verification of MAPS concept by the time of the TDR. This may mean a possible loss of consideration of MAPS as an option for an ILC ECAL. Mitigation: the designs will be reviewed extensively before submission. If TDR itself is delayed, the this project could be extended into another funding round.

6 Workpackage 4: Mechanical and thermal studies

The activities here are centred on the Manchester group, but activities by others can be incorporated. *Must talk with UCL people about how they want to fit in.*

6.1 TM1: Thermal Studies

Heat is generated in the front end electronics, and flows through the readout boards, the tungsten and the carbon fibre. This process is critical to operation, as the electronics requires stable operating temperatures, and critical to the design, as putting the electronics processing as near the front end as possible minimises the background pickup problems associated with small analog signals going down long cables, and the extent to which this can be done depends on heat and temperature considerations.

These considerations are interlinked with each other, and also with the electronics design, making the deliverables interdependent.

6.1.1 TM1.1: Power Dissipation

We need to check the power output actually produced by the existing readout chips, making actual measurements and cross-referencing them with manufacturers' data sheets.

This includes the actual Silicon detector pads themselves and the VFE ASIC which does preamplification and shaping, and possibly digisation and zero suppression. The later FE FPGA circuitry will also require cooling, though this is less critical as there are fewer circuit elements and they are more accessible.

6.1.2 TM1.2: Thermal design

Investigate the chip and circuit design to see what can be done to minimise heat production. In particular, look at the possibility of cycling the chips to they're only on during the bunch train passage

This activity must continue as an input to the design process for the VFE and FE. In particular, if we move to a MAPS readout the issues of thermal stability and control will need to be re-examined.

6.1.3 TM1.3: Thermal modelling

The heatflow from the source to the outside can be modelled using software already available and licensed, and the process is in principle simple, although the non-isotropic thermal conductivity properties of the Carbon Fibre will require care. The key difficulty of this activity will be to ensure that the detector being thermally modelled really does contain all the elements and properties of the real design, requiring those doing the modelling to be well integrated with the rest of the design team.

6.1.4 TM1.4: Thermal Measurements

The predictions of the thermal model must be validated against measurements to give them credibility. This will involve assemblies that mimic the actual CALICE silicon/Tungsten and carbon fibre as realistically as possible, with thermal generators and sensors.

6.1.5 TM1.5 Cooling Technology

There is a provision for a cooling pipe within the current design. This needs to be refined (for example, by the use of binary freon rather than water, so that the whole pipe could be at the same temperature) and the practical details of temperatures and flow rates worked out, backed by measurements as appropriate.

6.2 TM2: Glue

The 1 cm square silicon detector pads are connected to the board by conducting glue. This connection has to persist mechanically, thermally and electrically for the duration of the experiment for each of the 24 mllion pads. Ageing techniques can be studied by thermal cycling in a programmable oven.

6.2.1 TM2.1: Data collection

Adhesives are widely used in mechanical design, and considerable knowledge must exist as to their properties, both in general behaviour and specific examples. A first task must be a literature search to ensure that we benefit from existing research on the performance issues facing us, both for guidance in experimental technique and to ensure that do not repeat measurements that have already been performed.

6.2.2 TM2.2: Glue characterisation

Studies of existing adhesive through thermal cycling in a programmable oven, investigating the mean time between failures and the nature of such failures.

We also need to check whether the performance of the Silicon as a particle detector is adversely affected by glue molecules diffusing into the semiconductor.

6.3 TM3: Assembly

The assembly of 6000 boards, each involving the gluing of 4000 pads, is a major task. The detectors have to be positioned accurately and glued firmly, connections have to be made for the readout electronics, and the integrity of the whole board has got to be verified. Information on the whole process has got to be stored in a database for monitoring the progress of the build, and for later use.

This is in several ways similar to the building process for the ATLAS SCT currently being undertaken at Manchester, and we can use this expertise for the CALICE detector assembly planning.

6.3.1 TM3.1: Robot design

Design an automatic, or semi-automatic, process that will locate the detectors on the PCB with the desired position accuracy, applying adhesive and curing. It should also insert whatever reaodout chips are required and make connections.

6.3.2 TM3.2: Robot prototyping and testing

Based on the design, construct a proof-of-principle device that can demonstrably perform the assembly to the required precision and at the required rate.

6.4 Endcap

The design of the endcap for the ECAL would be a major contribution for the collaboration. It would position the UK to play a visible and leading rôle in the experiment. But it is a major task and we do not currently have the academic resources to pursue it. It is possible that a new academic appointment at Manchester could work on this, perhaps working with the engineering office at the Daresbury Laboratory. So we are not putting it in the request at this time, but may come back with a request to get involved if circumstances develop favourably.

6.5 Tables

6.5.1 Effort

Please note that these numbers have yet to be negotiated within Manchester and the ATLAS ring fence has to be respected.

Name	Position	Institute	Funding	2005/6	2006/7	2007/8
R Barlow	Academic	Manchester	HEFCE	0.1	0.1	0.1
R Thompson	Engineer	Manchester	PPARC(RG)	0.1	0.2	0.3
S Snow	Engineer	Manchester	PPARC(RG)	0.1	0.1	0.1
J Freestone	Technician	Manchester	PPARC(RG)	0.2	0.2	0.2
A Elvin	Technician	Manchester	PPARC(RG)	0.2	0.2	0.2
M Perry	Technician	Manchester	PPARC(RG)	0.1	0.1	0.1

ID	Start	Finish	Predecessors
1	**4/1/05**	**4/12/06**	
2	4/1/05	5/12/05	
3	5/13/05	6/28/05	2
4	4/1/05	5/12/05	
5	5/13/05	8/4/05	4
6	8/4/05	4/12/06	
7	**12/20/06**	**4/25/07**	
8	12/20/06	1/31/07	
9	1/31/07	4/25/07	8
10	**4/1/05**	**12/20/06**	
11	4/1/05	8/3/05	
12	8/5/05	12/20/06	11
13	4/1/05	4/1/05	

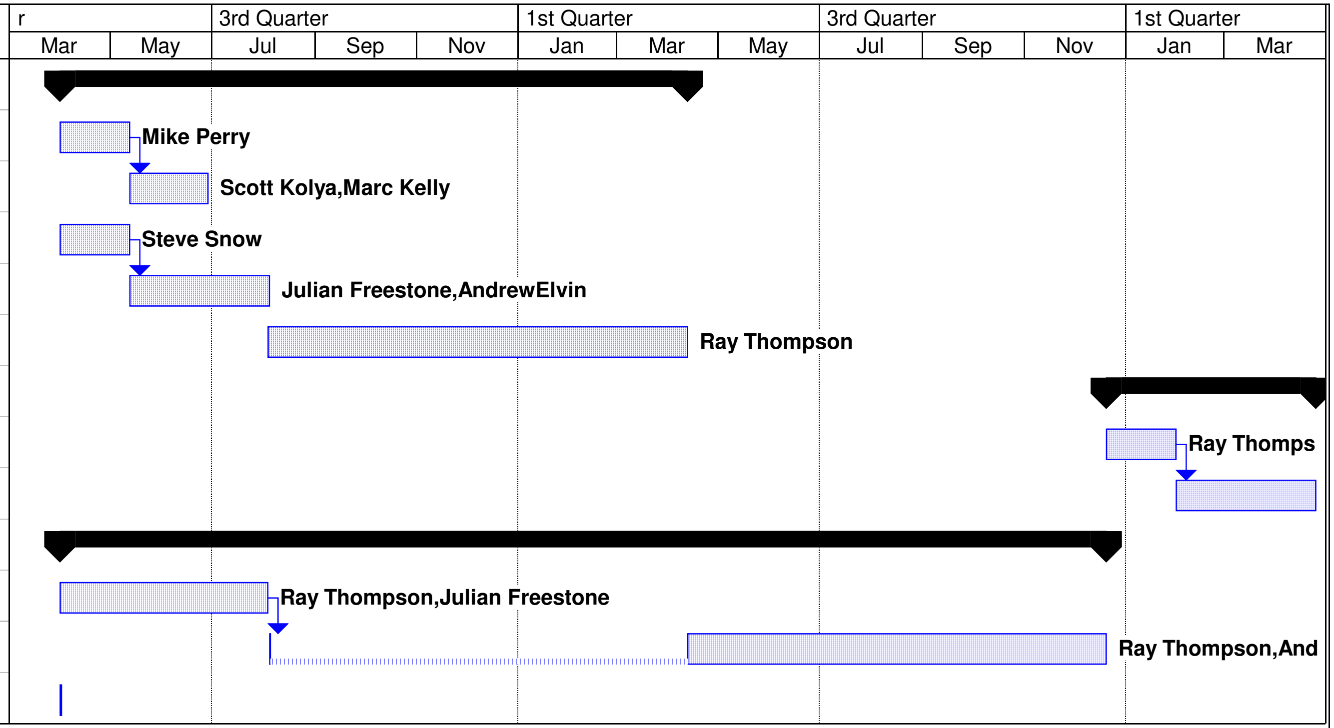

Project: Project
Date: 12/13/04

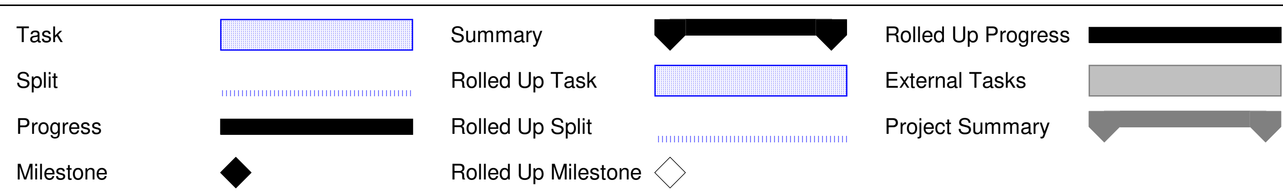

ERROR: stackunderflow
OFFENDING COMMAND: restore

STACK: