INCLUDING SOME RESPONSES TO
RMR FEEDBACK

October 30, 2009

ATLAS UK Upgrade Proposal

ATLAS UK
Birmingham, Cambridge, Edinburgh, Glasgow, Lancaster, Liverpool, Oxford,
Manchester, QMUL, RAL, RHUL, Sheffield, Sussex, UCL
Spokesperson/Contact J.Butterworth@ucl.ac.uk

Contents

Chapter 1

Introduction

The ATLAS upgrade is a complex particle physics project which encompasses a number of detector, trigger, software and computing developments required to continue the exploitation of ATLAS at the Large Hadron Collider (LHC) well into and beyond the next decade. The current international ATLAS collaboration is by necessity already engaged in planning and R&D directed at ensuring the success of the project, and ATLAS UK is an integral part of this activity, with substantial leadership, as discusion in Section 1.1. The UK context for this proposal is discussed in Section 1.2.

The LHC and upgrades are likely to be staged and gradual to some extent as experience is gained with operating the accelerator complex. Current planning divides detector upgrades into two phases. Phase-I encompasses upgrades to luminosities of around 3×10^{34} cm^{-2}s^{-1} around 2015. This would be followed by a long shut down (in which a new tracking detector can be installed) around 2019 and start of Phase-II (10×10^{34} cm^{-2}s^{-1}) around 2020. Upgrades required for this operation are described as "Phase-II". Despite the long time-scales, in the period of this three year bid much work is required both on Phase-I and on R&D for Phase-II, due to the long lead-times. In fact R&D, on the Phase-II tracker upgrade in particular, is already well underway both internationally and in the UK.

The principal areas of UK involvement centre on providing a successor to the current Semiconductor Tracker (SCT), which has a radiation-dictated lifetime finishing towards the end of the coming decade, Trigger, and Computing Infrastructure, all of which must be able to cope with an order of magnitude more luminosity than the current design. These three major areas are discussed in the three main chapters of this bid, and are broken down into Phase-I- and Phase-II-related sub-areas and workpackages as appropriate.

One area which is somewhat different is the ATLAS Forward Physics project, which is on the same timescale as Phase-I luminosity upgrades but which in fact is a pure instrumentation upgrade, largely independent of the rest of the programme. As requested by the PPAN, this project is presented such that the case and resources may be treated independently of the luminosity upgrade activities. Nevertheless the synergies in the physics, technology and personnel, as well as the requirement to manage a coherent UK programme, lead us to present in the same document, with its own physics case (Section 1.3.1) but with the major hardware deliverables inluded under Phase-I tracking upgrades (Section 2.2).

1.1 International Context

The LHC machine and its associated experiments are the end result of a two decade long programme of research, development, construction, installation and commissioning. As the LHC takes over as the world's highest energy facility, detailed planning has been going on for several years on the full exploitation of this remarkable new facility, including operation well beyond the 10 years foreseen for achieving the LHC's original physics goals.

Experience with all previous accelerator complexes and collider experiment programmes teaches that planning to fully exploit the capabilities through major upgrades allows a much greater return on investment and, as experience at the Tevatron in particular has illustrated, the benefits outlined above really do lead to new physics reach. A difference from previous upgrades is that since the LHC represents such a hostile environment in terms of radiation, event rate and hit occupancy, it already required detector technologies which were state-of-the-art for their time to guarantee a 10 year lifetime. The experiments therefore need considerable investment in newer technologies to meet the challenges of sLHC operation.

The priority of the LHC luminosity upgrade has been emphasised by many international reviews, but the planning is still limited by a number of factors, not least the timing of the LHC ramp-up itself and the dates when physics priorities and real running experience (including accurate dose predictions) are likely to be available.

The current thinking in ATLAS is to aim for a Letter of Intent (LoI) in 2010, followed by a Technical Proposal (TP) with Memorandum of Understanding (MoU) in 2012 for the full Upgrade programme. In parallel, Technical Design Reports (TDRs) for sub-systems will be developed with some, such as that for the tracker upgrade, needed soon after the TP to allow adequate time for procurement construction and commissioning by 2018, assuming an 18 month shutdown for installation and resumption of operation around 2019/2020. For the triggers, the Level-1 Calorimeter Trigger and Level-1 Track Trigger upgrade work proposed here is all contained in a TDAQ system LoI, which was approved by the TDAQ Institute Board (70 institutes) and submitted in February 2009. The next step is the submission of an ATLAS internal TDAQ R&D proposal, which is expected to start roughly at the same time as the work proposed here. We have throughout the various internal ATLAS discussions taken steps to to bring new participating institutes into the activities.

The UK continues to play a prominent role in many aspects of installation, commissioning and testing at CERN, with many personnel now based there associated with the SCT, Trigger and Computing. From July 2005, members of the UK community have started participating in discussions on a possible luminosity upgrade with the hit density and radiation levels implying the need for a complete replacement of all tracking layers and much of the triggering architecture and computing infrastructure (as well as major changes to other detector sub-systems in which the UK has not historically been involved). Many major international ATLAS upgrade workshops have been held at CERN and elsewhere: Genova (2005), Liverpool (2006) Valencia (2007), and NIKHEF (2008). Four out of the 14 highest-level positions in the Upgrade Steering Group (those for the Radiation Simulation, Tracker, Trigger/DAQ and Computing) are held by UK personnel, as are many of those for sub-task convenors (six out of the 25 defined roles in the Tracker Upgrade for example). The contributions of the UK to all aspects of the sub-systems where it has involvement are already highly visible, influencing the appointments to the leadership roles discussed above.

1.2 UK Context

ATLAS is the largest UK particle physics collaboration. Fourteen institutes from the UK participate. We are currently around 10% of ATLAS, and aim to maintain this level of involvement into the upgrade and for the duration of the world-leading physics programme at the LHC.

The ATLAS tracker upgrade received PRD funding in 2007, and the current bid covers the next three years of the whole upgrade project, from April 2010 to March 2013. In the tracking area (Chapter 2), this bid is a request for resources to contine the programme. In the case of Trigger and Computing areas (Chapters 3 and 4), this bid is to initiate funding for UK activity.

The ATLAS project involves approximately 90 UK academics (for whom the majority of funding comes from non-STFC sources) and we would anticipate the eventual involvement, including exploitation of the upgraded detectors, to be comparable to the current ATLAS complement of around 70 PDRA, 85 technical/engineering/applied physicists and 150 students.

1.3 Overview of Physics Motivation

1.3.1 Forward Physics

Although forward proton detectors have been used to study Standard Model (SM) physics for a couple of decades, the benefits of using proton detectors to search for New Physics at the LHC have only been fully appreciated within the last few years [1–6]. By detecting both outgoing protons that have lost less than 2% of their longitudinal momentum [7], in conjunction with a measurement of the associated centrally produced system using the central components of the ATLAS detector, a rich programme of studies in QCD, electroweak, Higgs and Beyond the Standard Model (BSM) physics becomes accessible, with the potential to make unique measurements at the LHC. A prime process of interest is Central Exclusive Production (CEP), $pp \rightarrow p + \phi + p$, in which the outgoing protons remain intact and the central system ϕ may be a single particle such as a Higgs boson. In order to detect both outgoing protons in the range of momentum loss appropriate for central systems in the 100 to 200 GeV mass range during nominal high-luminosity running, high-precision proton tagging detectors must be installed close to the outgoing beams in the high-dispersion regions at 220 m and 420 m from the interaction points on each side of ATLAS.

The physics case for such detectors was discussed in detail in the AFP Letter of Intent to ATLAS, which is largely based on the FP420 R&D report [8]. The initial FP420 R&D programme was endorsed by the LHCC in 2005: "The LHCC acknowledges the scientific merit of the FP420 physics program and the interest in its exploring its feasibility." The addition of proton detectors in the 220 m region as well as at 420 m further strengthens the physics case by increasing the acceptance over a wide mass range of potential resonances. The addition of AFP to ATLAS enables the LHC to operate as a gluon-gluon or photon-photon collider with known centre-of-mass energy in the range of 70 GeV to 1.4 TeV. The mass acceptance for 420-420, 220-420 and 220-220 configurations is shown in Fig.1.2.

AFP provides a rich QCD and electroweak physics programme extending the current ATLAS physics capabilities. Figure 1.1 shows Feynman diagrams that can be explored using AFP - CEP and photon-photon scattering.

The physics of photon-photon scattering is described in Refs. [9, 10]. Diagram 1.1a shows

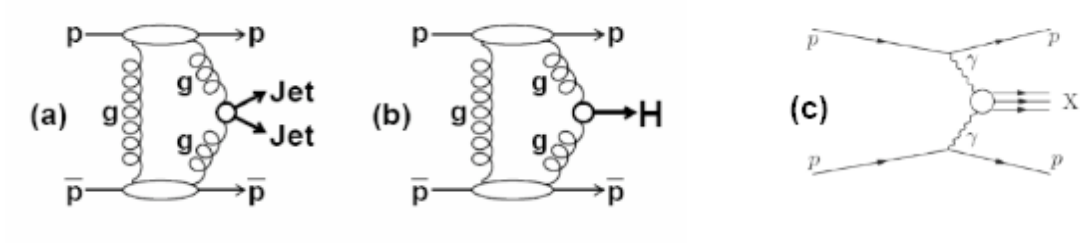

Figure 1.1: Key Feynman diagrams that can be studied by AFP. a) Di-jet production, b) Higgs production by CEP, and c) photon-photon physics.

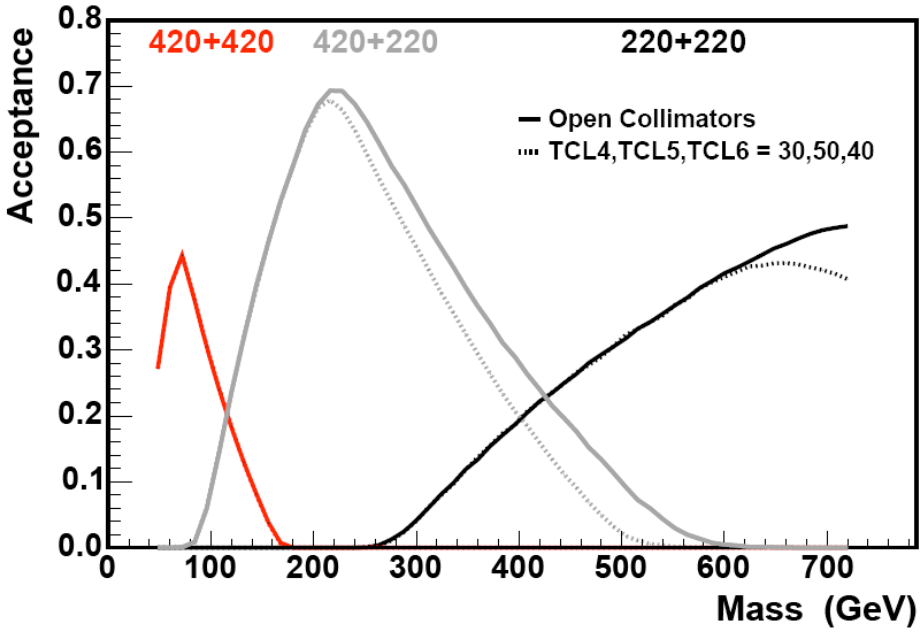

Figure 1.2: Mass acceptance plot for protons detected in stations at 420 + 420, 420 + 220 and 220 + 220. Various LHC collimator settings are indicated.

di-jet production. This process has been found by CDF [11], and cross-sections are compatible with theoretical calculations. CDF have also found evidence for $\gamma\gamma \to \mu^+\mu^-$ and exclusive charmonium production, also with rates that are compatible with calculations [12, 13]. Photon-photon physics complements the ATLAS strategy for precision limits on (or measurements of) anomalous couplings and the search for supersymmetric (and other BSM) charged particles. Gluon-gluon physics allows the study of Higgs bosons in the CEP channel. In addition to a mass measurement of about 3 GeV on an event-by-event basis, the observation of a Higgs boson in this channel is equivalent to a quantum number measurement because resonance production is heavily suppressed for particles that do not have $J^{PC} = 0^{++}$. If a Higgs signal is observed in the central detector, it will be important to confirm its properties, especially the quantum numbers. There are two approaches to determine the quantum numbers - either by production or decay. The methods are summarised in Table 1.1.

Channel	**Quantum number determination due to production mechanism**
CEP (AFP only)	If the Higgs is created via CEP, then selection rules require it to be $J^{PC} = 0^{++}$.
Vector Boson Fusion	Some channels have been demonstrated [13].
	Quantum number determination from analysis of Higgs decay
Central Inclusive Production	$H \to ZZ$ angular correlations for quantum numbers possible for a heavy Higgs. Observation of $H \to \gamma\gamma$ requires that C must be positive and $J = 0$.

Table 1.1: Quantum Number Determination

AFP offers a complementary approach to the few existing standard approaches in central inclusive production. The CEP technique is likely to work with standard model or BSM production of Higgs bosons. In addition, CEP measures the Higgs vector boson couplings and the Higgs quantum numbers as two independent measurements, whilst in other approaches the two determinations and their interpretations are interconnected. The capability of precision forward proton detectors to measure Higgs boson properties has been extensively studied in the literature over the last four years [1].

- The Standard Model $H \to WW^*$ channel has been studied, with the conclusion that the semi-leptonic and fully leptonic decay channels can be measured for 140 GeV $< M_H <$ 200 GeV. For the small α_{eff} scenario in the MSSM, it is perhaps possible to study the Higgs boson down to 120 GeV.

- The $h, H \to \tau^+\tau^-$ and $b\bar{b}$ channels have been studied for the MSSM, tri-mixing, triplet, and 4^{th} generation Higgs sectors, with the conclusion that AFP could observe these channels for a large area of parameter space in each model. There are also other models (NMSSM, CPX) in which the $bb/\tau\tau$ channels could be studied using AFP. Figure 1.3 shows mass plots for a MSSM Higgs with $M_A = 120$ GeV and $\tan\beta = 40$. More integrated cross-section is required at high luminosity because of overlap background (see

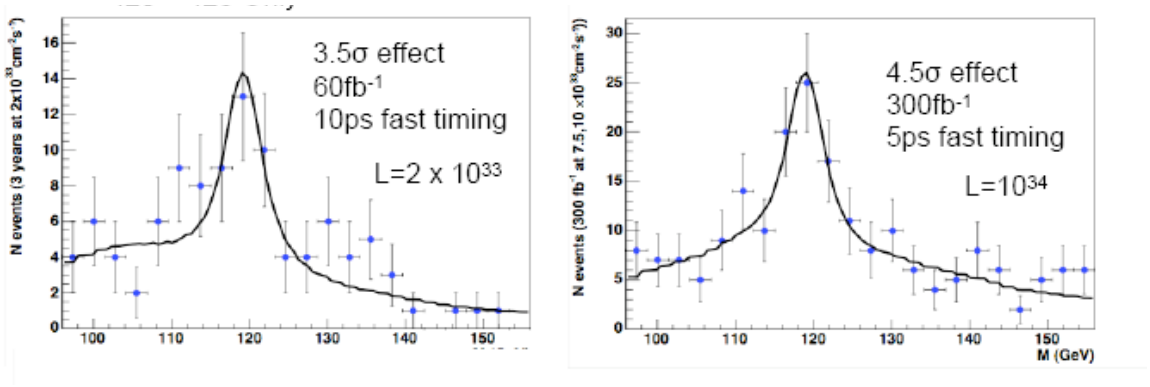

Figure 1.3: Higgs signal for MSSM $H \to b\bar{b}$ with $M_A = 120\,\mathrm{GeV/c^2}$ and $\tan\beta = 40$. See Ref. [1] for further details.

Section 2.2.2). A strategy which alters the use of trigger bandwidth dependent on the actual luminosity which drops considerably during an LHC fill will improve the sensitivity at high luminosity by a factor two.

- The $h \to aa \to 4\tau$ decay channel has been studied, and demonstrated to be viable, for the NMSSM. It should be noted that similar topologies occur in other models such as CPX.

In recent months, much effort has been put into developing a trigger for the produced in CEP. For a low mass standard model Higgs this is a very challenging channel for the GPDs at the LHC. This work has been written up in Ref. [14]. At Level-1, by using a proton tag at 220 m and topological cuts based on information from the electromagnetic calorimeter, it may be possible to detect a low mass standard model Higgs, and a BSM Higgs is observable. Further comments are added later. This will require an upgrade to the Level-1 trigger which is described in Section 3.2.

We have augmented previous studies on the Higgs physics case using ATLFAST with the ExHuME Monte Carlo and various Monte Carlo samples with either full simulation if available, or ATLFAST if the full simulation samples did not exist. A detailed note is in preparation, and we provide a summary of the important points here. We have studied two decay modes of the Higgs boson, $H \to b\bar{b}$ for $M_H = 120$ GeV and $H \to WW^*$ in the semi-leptonic channels with $WW^* \to l\nu jj$ and $M_H = 160$ GeV (with $l = e$ or μ). Our current understanding after detailed studies is consistent with the previous generator-level studies, namely that under realistic assumptions, the enhanced cross sections predicted by BSM Higgs models are required to observe the Higgs Boson in the $H \to b\bar{b}$ channel, while, if the Higgs is heavier, the SM Higgs is observable in the WW^* channel. The simple observation of a Higgs signal establishes its quantum numbers, and only a few events are needed for a precise mass measurement. In addition to a sufficiently large data sample, the observation of a 120 GeV Standard Model Higgs would require further improvements in the trigger capabilities combined with improved background rejection, both of which may be possible given experience operating the detectors. We also note that the WW^* cross section is only a slow function of Higgs mass - the cross section at 140 GeV is only reduced by 20% relative to 160 GeV - and that tau decay modes and dilepton channels will give further enhancements to the significance. Combining channels and applying realistic running scenarios has the potential to extend our Standard Model Higgs reach to 130 GeV and

possibly below.

Finally, in addition to the Higgs physics and photon-photon physics, there is much standard model QCD physics and various exotic processes to study which are described in the FP420 Design Report [8].

1.3.2 Phase I Trigger

The obvious consequences of raising the luminosity of LHC are higher detector occupancy, increased trigger rates at fixed transverse-momentum thresholds (or higher thresholds for fixed rates), and higher levels of radiation that could damage or perturb the detectors and the on-detector electronics. Increased occupancy has two important consequences for the trigger and data acquisition (TDAQ) system: degraded performance of trigger algorithms due to the increase in pileup, and a larger event size to be read out. Examples of the degradation of the trigger performance include reduced rejection at fixed efficiency from isolation requirements on electron/photon candidates, and increased muon-trigger background rates arising from accidental coincidences between radiation-induced 'noise' hits in the muon detectors.

Most ATLAS detectors plan to keep the same electronics and readout for Phase-I operation. Therefore, the designed maximum Level-1 Trigger rate of 100 kHz can certainly not be exceeded. Indeed, for those detectors whose event size increases with occupancy, the maximum allowed Level-1 Trigger rate for fixed readout bandwidth has to be correspondingly reduced. The trigger rate can potentially be controlled by raising the transverse-momentum thresholds on candidate electrons, photons, muons, etc., with the increase in thresholds compensating for the higher interaction rate, and also for the degradation in algorithm performance due to the higher occupancy (less rejection for fixed efficiency). However, studies show that the muon trigger rate declines only slowly with increasing threshold, while low thresholds are essential for parts of the physics programme. The preferred technique is therefore to bring forward selections previously used only in later stages of the trigger, notably by using topological information and increasingly exclusive trigger signatures at Level-1. The triggers required at higher luminosity cannot of course be known accurately until significant running has been completed at LHC design luminosity. In addition, it is not yet possible to predict what new or unexpected physics might emerge from initial running, nor whether there will be unforeseen backgrounds that require changes in trigger strategy. However, three types of triggers are likely to be required:

- Triggers to complete the LHC physics programme, e.g. precise measurements of the Higgs sector. These require thresholds on leptons, photons, and jets to be as low as those used at design luminosity. As some of the final states will be known, more exclusive menus could be used (e.g. one lepton plus two b-jets plus missing energy) targeted to the final states to be studied.

- Triggers for very high-p_{T} discovery physics. These should not cause big rate problems since thresholds can be as high as several hundreds of GeV.

- Control and calibration triggers with low thresholds, selecting, e.g., W, Z and top events. Here again, exclusive menus can be used, with pre-scaling to limit the accepted rate.

Monte Carlo simulation of trigger and detector performance with pileup is key to understanding which selections are effective. All detectors are influenced by the instantaneous rates, but for many the performance also depends on the interaction history extending tens of bunch-crossings before the interaction of interest. Initial data samples have recently been produced and validation is underway. This work is covered in detail in Section 4.

Just as the event rate at Level-1 is limited by the readout bandwidth, the rate after the High-Level Trigger (HLT) is fixed by the bandwidth that can be accepted for offline processing. Because of the increased event size it is estimated that the HLT will have to retain essentially the same rejection factor as for design luminosity. To achieve this, substantial changes are needed in the High-Level Trigger to adapt and optimise the algorithms for the high pileup environment, to upgrade software for the detector and Level-1 changes, and to implement and tune additional, more exclusive, HLT selections. The tracking algorithms and selection software will need to be retuned for the higher occupancies and greater complexity of events. There will be some new aspects, such as the inclusion of a new pixel layer (the Insertable B-layer) and adapting the Level-2 tracking to profit from the track parameters determined by the Fast Track Processor (FTK) , if installed. But much is an extension of ongoing work to prepare for the evolution of luminosity up to the design value and will profit from the existing UK experience and expertise.

Predicted trigger rates will be determined as a function of luminosity, including output rate and the rates at each stage of selection. Optimisations will be made to selection cuts and ordering of steps within the trigger chains to achieve the earliest possible rejection, and so minimise CPU resources. The development of menus for Phase-I is a continuation of ongoing work to adapt menus as luminosity increases. In order to control the trigger rate, in addition to increasing pre-scales and raising thresholds, new selections will be developed including increased use of multi-object and topological triggers (combining two or more features identified at Level-1).

1.3.3 Phase-II Upgrades

The LHC will be the first accelerator to operate well above the energy regime of electroweak symmetry breaking. Therefore, data from the ATLAS detector will revolutionise our understanding of high energy physics by either finding the Standard Model (SM) Higgs particle - in which case many theoretical problems relating to the Higgs mass will still need to be addressed by detailed measurements - or probing whatever other mechanism nature has chosen to generate the masses of the fundamental particles. Many other discoveries are also possible - amongst them supersymmetry, extra spatial dimensions or TeV-scale gravity. Many of these could explain the huge difference between the electroweak symmetry breaking scale and the GUT/Planck energy scale.

The exact physics goals of the Phase-II upgraded ATLAS detector will clearly be much better known after first data are analysed. In general terms, there are two benefits:

- Increased data volume means that the statistical precision with which rare processes may be measured improves. In practice, many systematic uncertainties are also determined by statistically limited control samples, and so in fact the overall precision can in improve very substantially. This provides a step up in the ability of the data to challenge new and existing fundamental theories.

- The energy scale probed is dicated by the energy of the incoming quarks and gluons

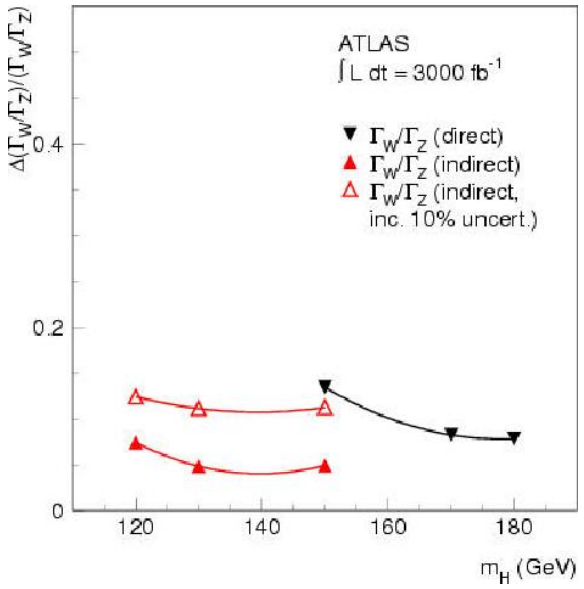

Figure 1.4: Fractional uncertainty on the relative partial Higgs widths to WW and ZZ as a function of the Higgs mass for 3000 fb^{-1}.

which make up the proton. These have a distribution which falls as the fraction of the proton's energy they carry increases. Thus an increase in luminosity means an increase in statistics for the very highest energy quarks and gluon collisions, significantly extending the energy reach of the collider, and hence its ability to probe new physics at the high energy frontier.

Even now however we can identify the main areas in which a ten-fold increase in statistics will have the most profound impact.

Precision measurements of electroweak symmetry breaking. Should a Higgs boson be discovered at the LHC, detailed measurement of its properties will be a major part of the ATLAS Phase-II physics program. Increased data sets in different final states will enable more accurate measurements of the Higgs couplings to SM fermions and gauge bosons (e.g. for WW and ZZ final states see Fig. 1.4 [15]), which provide a sensitive test of the Standard Model . For instance the larger data sets will enable for the first time precise determination of the relative branching ratios of the $\tau^+\tau^-$ and $\mu^+\mu^-$ final states. This provides a vital test of the nature of the Higgs boson as the predicted relative decay rates depend only on the muon and tau masses. μ and τ lepton identification and measurement with the required precision will be crucially dependent on accurate reconstruction of charged particle tracks in the ATLAS inner tracker.

Should a Higgs boson not have been observed at design-luminosity, the main focus of the physics programme will turn to further investigation of the scattering of gauge bosons at large invariant mass (Fig. 1.5 [15]), which will provide information on the dynamics of electroweak symmetry breaking. This process violates unitarity at the TeV-scale without a light Higgs boson, and so the Standard Model must break down at this scale. The processes of interest are $qq \rightarrow qqVV$ where V can be either W or Z, which probe a possible strongly coupled VV final state. The VV final states will be detected where at least one of the gauge bosons undergoes a leptonic

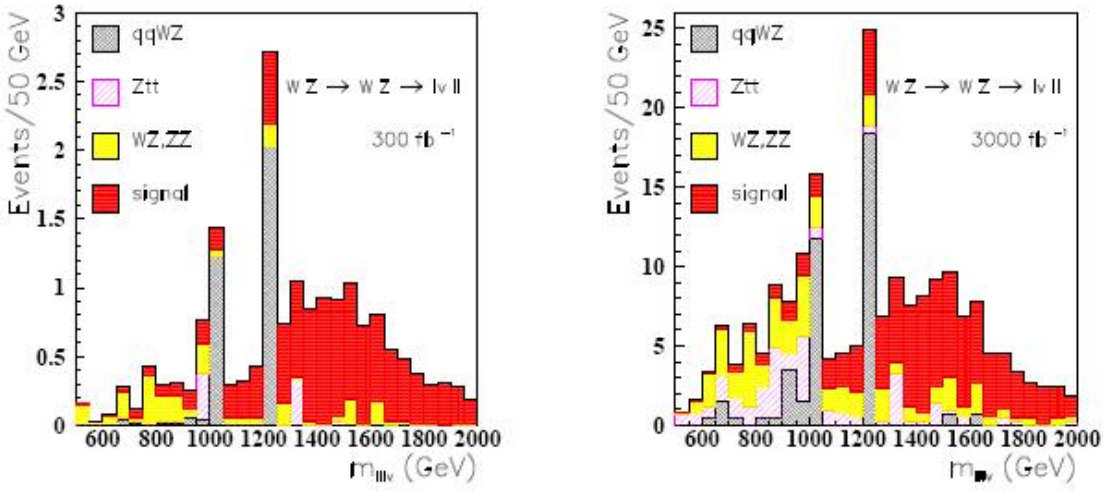

Figure 1.5: Expected signal and background for a 1.5 TeV WZ resonance in the leptonic decay channel at 300 fb^{-1}(left) and 3000 fb^{-1}(right). (Note the scaling of the vertical axis)

decay. Efficient lepton triggers will therefore be crucial, as will accurate lepton measurements with the new inner tracker.

Further investigation of new physics signatures. ATLAS will give unprecedented sensitivity to physics beyond the Standard Model. Many models of such physics predict the existence of a wide spectrum of new particles. However, some of these may not be accessible to detailed study with either the design-luminosity LHC or a TeV-scale linear collider. For these states, a ten-fold increase in LHC statistics has a major impact, enabling the discoveries and precision measurements which would be crucial for developing a complete understanding of the new physics.

One of the main classes of new physics model which ATLAS hopes to discover is supersymmetry, which seeks to solve the SM gauge hierarchy problem and at the same time provides a natural candidate for the dark matter. If SUSY signals are seen at ATLAS the role of the Phase-II upgraded ATLAS detector will be to find evidence for decays of supersymmetic particles into final states with small branching ratios (e.g. containing Higgs bosons see Fig. 1.6 [16]) and to observe the direct production of those sparticles without strong couplings (such as sleptons and gauginos) whose production rates may be too small to observe at design-luminosity. Observation of such processes will strongly constrain the SUSY model by enabling cross-section and mass measurements, the latter via end-points in lepton and jet invariant mass distributions. Another exciting possibility enabled by enhanced LHC luminosity and pioneered by ATLAS UK is the measurement of characteristic SUSY particle spin-statistics through subtle differences in the shapes of these invariant mass distributions. Such measurements will be vital for distinguishing between SUSY and alternative interpretations such as Universal Extra Dimensions (UED) (see Fig. 1.7 [17]).

Alternative new physics models sought by ATLAS include those generating new gauge bosons such as a Z' or a W'. These particles can occur in models with extra gauge symmetries or extra spatial dimensions. If such particles are found then it is vital that their masses

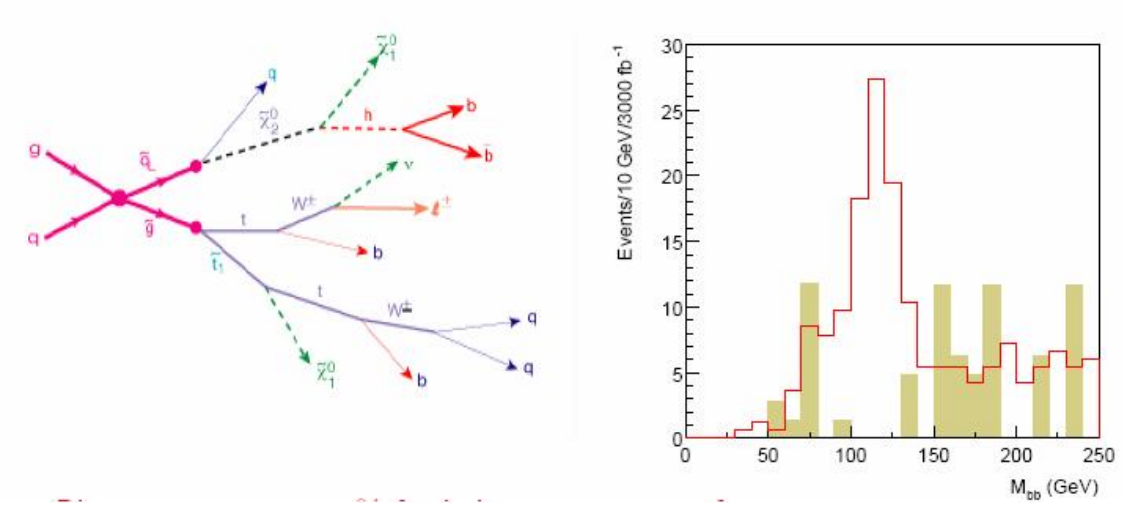

Figure 1.6: Sensitivity to SUSY particle couplings to the Higgs boson. Higgs bosons produced in cascade decays (left) can be observed with 3000 fb^{-1}(right - signal(background) is open(filled) histogram).

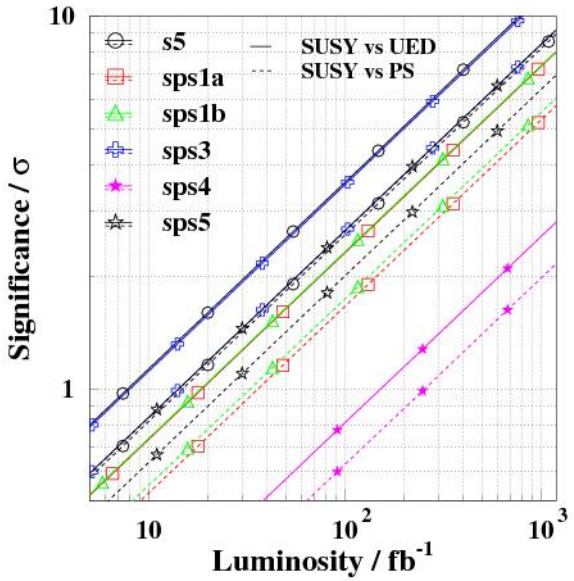

Figure 1.7: Sensitivity to SUSY particle spin-statistics in the di-slepton channel. Full (dashed) lines indicate the significance of a deviation of a SUSY signal from phase-space (UED) models as a function of integrated luminosity.

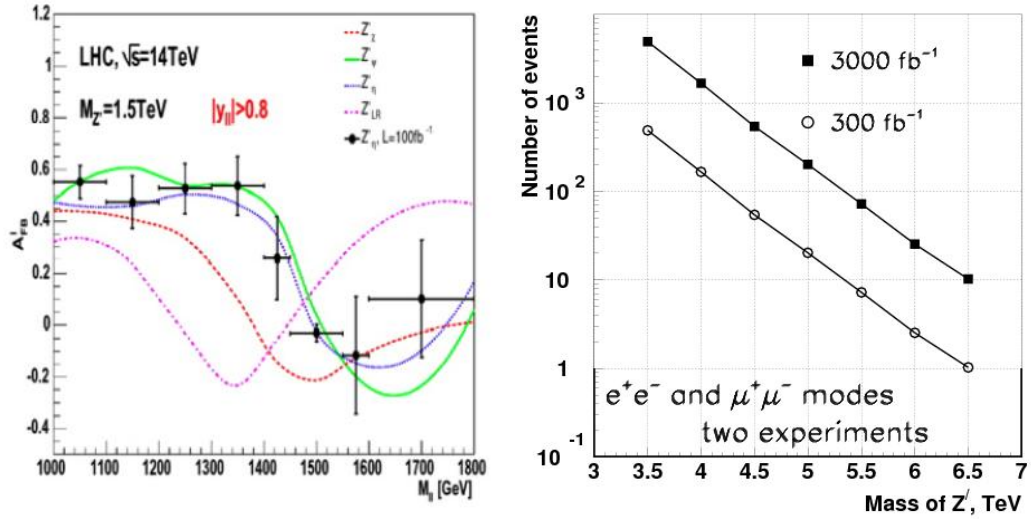

Figure 1.8: Z' measurements and searches at LHC and the Phase-II upgrade. The left-hand figure shows that 300 fb^{-1}will be insufficient to discriminate between some classes of Z' model. The right-hand figure shows the number of Z' events expected as a function of Z' mass for one Z' model in the leptonic decay channel, given 300 fb^{-1}or 3000 fb^{-1}.

and couplings are measured, much as was done with the Z boson at LEP. These analyses will require highly precise lepton and b-jet measurements enabled by the new tracker. An example of a Z' leptonic forward-backward asymmetry measurement requiring high luminosity is shown in Fig. 1.8 (left) [18].

Extended mass reach for new particles. The impact of increased luminosity exceeds that of simply improving measurements of observed states; it opens up new energy regimes by raising the statistical significance of low-probability, high-energy constituent collisions. This correlation between energy-reach and luminosity is characteristic of hadron colliders.

In a wide range of channels the discovery reach of the Phase-II LHC upgrade is increased by $\sim$0.5 - 1 TeV over that of the design-luminosity LHC. Figure 1.8 (right) [19] shows the expected improvement in the number of Z' events in the leptonic decay channel. Similar plots can be shown for high mass SUSY states or for the characteristic energy reach defining the distance scale in compactified extra dimensions models. Many SUSY models also predict the existence of additional heavy Higgs bosons too massive to be produced at a TeV-scale Linear Collider. Searches for these states at the LHC are limited by statistics, but are crucial to a complete understanding of SUSY. The discovery reach in terms of the Higgs mass is extended by up to 75% by a ten-fold increase in LHC luminosity.

Precision Standard Model measurements. It should not be forgotten that the Phase-II upgraded LHC will act as a copious source of Standard Model particles. The resulting very large samples will enable ATLAS to place strong constraints on the anomalous couplings and rare decay modes of these particles. Of particular interest are trilinear and quartic couplings of

Process	WWW	WWZ	ZZW	ZZZ	$WWWW$	$WWWZ$
N(m_H=120 GeV)	2600	1100	36	7	5	0.8
N(m_H=200 GeV)	7100	2000	130	33	20	1.6

Figure 1.9: Quartic gauge coupling event statistics for 6000 fb^{-1}.

the SM gauge bosons (Fig. 1.9) which probe directly the non-Abelian nature of the SM gauge group. Flavour changing neutral current decays of the top quark represent another area where several design-luminosity analyses will be statistics rather than systematics limited. High statistics measurements of rare decays of tau leptons or hadrons containing bottom quarks also offer potential sensitivity to departures from the Standard Model.

If the increased data sets provided by the Phase-II upgraded LHC are to be fully exploited, the performance of ATLAS must not be allowed to degrade. Specifically, any changes to the trigger efficiency and trigger thresholds must not lower the trigger acceptance, and the inner tracker resolution and background rejection must not result in a reduced signal to background ratio for the processes of interest. In particular accurate measurements of systems in the few hundred GeV range could be significantly affected by the large minimum-bias event pile-up, and reduced efficiencies or increased backgrounds could spoil the advantage of the higher luminosity.

The UK's expertise and leadership in the tracker, trigger and computing areas is essential for successful realisation of all of these physics goals and indeed non-participation in even one would threaten the viability of the entire project. The UK needs to maintain its leadership in these key sub-systems in order to capitalise fully on the original LHC investment and to guarantee continued access to world-leading, energy-frontier science over the next twenty years.

Chapter 2

Upgrade Tracking

2.1 Overview

2.1.1 Heritage

The groups bidding for funding to do the work described in the tracking sections of this document have a long history of success in building silicon tracking detectors. In particular there is a very large overlap with the current team and that which successfully built the existing silicon strip tracking detector in ATLAS. The ATLAS SCT is a large international project, which was, and is, led by UK team members. During the construction phase, the project leader and several of the sub-project leaders were from the current team, and financial oversight was supplied by the UK. All four cylinders and half of the forward discs for the detector were assembled in the UK. Final assembly and installation at CERN was undertaken predominantly by UK personnel. Leadership and much of the detailed work for the DAQ for the detector has come from the UK.

Amongst other things, the UK team designed modules and support structures; completed numerous FEA and fluid dynamical calculations; developed cooling systems, connectors, data transmission and clock and control distribution systems. The breadth of experience vested in the team is world leading and the proposed programme will exploit that for the next generation experiment and also serve to preserve the UK as an international leader in the field.

2.1.2 Outline of Proposal

As described in Section 1, upgrade tracker construction can be divided into two phases. In the first phase ATLAS-UK is enthusiastic about making a significant contribution to the ATLAS Forward Physics programme, AFP. UK leadership in this programme is well established and the physics case is strong. The UK has leadership in 3D technology and the leader of the 3D sensor development within ATLAS is from the UK. The requested new resource is modest and the programme is excellent value. As requested by STFC, this programme has been treated slightly differently to the other tracking requests and the sections below on AFP represent a self-contained description of the proposed programme, including the physics case as this differs from that for Phase-II. It has been included as part of the proposed tracking upgrade to emphasise that the effort is integrated with the other aspects of the proposal. Further discussion of AFP is left to the detailed sections.

In the second phase of the upgrade programme, ATLAS-UK intends to bid to play a major role in the construction of the barrel inner detector strip sLHC upgrade tracker, specifically the short strip cylinders, though the design of long and short strip barrels will be similar. The work proposed here builds on the R&D programme already funded by STFC. The UK has been responsible for the design and prototyping of the first ATLAS short strip Tracker Upgrade barrel modules. Bringing UK expertise to building the upgraded strip tracker will be vital to ensuring that the detector is built in an optimised way to a tight schedule. Beyond the role of defining the detector modules and support structures, it is crucial that the extensive experience with systems be applied to the upgrade tracker. Lessons from both design and installation of the existing ATLAS silicon strip tracker must be applied at the detailed level to ensure that the sLHC tracker is not only as successful as the existing ATLAS silicon tracker, but also more robust and therefore capable of much faster installation and commissioning. Amongst other things, this results in a need for more parallelism in construction, a different strategy for alignment and manufacture tolerances, and much more attention to services and clearances at the design stage.

The global plan for production of the short strip tracker in the UK involves two clusters of institutes, each of which is to be capable of producing identical detector elements, termed 'staves'. These are described in Section 2.3.2 and are sent to CERN for final assembly into the macro support. Evolving the optimised organisational structure to accomplish this, including evaluation of the potential of the new STFC Detector Systems Centre, is one goal of the programme proposed here.

Whilst the short-strip tracker will be the largest part of the UK Phase-II tracker upgrade programme, it is important that it is not the only UK role. Strategically it is vital that UK particle physics becomes as strong in hybrid pixel detectors as it currently is in strip technology since future projects are likely to rely more heavily on pixels than in the past. In addition, pixel detectors are proving to be of considerable importance to other disciplines, making it particularly timely to increase emphasis on this technology. There is UK leadership in sensor technologies aimed both at the short strip region and for pixels where there is leadership in development of 3D technologies and in proving that planar p-type detectors can work to sLHC doses. The commissioning of pixel assembly (flip-chip) systems at STFC laboratories is therefore a natural development and an excellent way to play to UK strengths.

ATLAS-UK intends to bid to be a major force in the development of forward Pixel Discs for the Phase-II tracker upgrade. This can offer significant synergies with the UK proposed upgrade of the LHCb VeLo and, in terms of support services and other ancillary systems, also with the short-strip work. ATLAS-UK participated in two tracker projects for the current ATLAS inner detector and believes that playing a major role in both the short strip barrel and in the forward pixel discs represents a viable and exciting programme which will preserve the UKs leading role in silicon tracking, and in the ATLAS inner detector upgrade in general. Currently the UK holds a disproportionately large number of international leadership roles in the ATLAS Project Office and Upgrade Steering Group. These include cooling, powering, opto-electronics and DAQ. As these are common projects across the tracker, it is easier and natural for the UK to bid to be involved in two technology areas.

2.2 ATLAS Forward Physics (Phase-I)

2.2.1 Overview

The current ATLAS Forward Physics proposal grew out of the FP420 R&D collaboration, which was formed in 2005 and included members from ATLAS, CMS, TOTEM and the accelerator physics community, with support from theorists. This aimed at assessing the feasibility of installing proton-tagging detectors 420 m from the interaction regions of ATLAS and CMS. This UK-led collaboration was funded by PPARC/STFC and reported on its work in May 2008 [8]. The ATLAS FP420 groups joined with the ATLAS RP220 groups in 2008 and submitted a LOI to ATLAS to add forward detectors at 220 m and 420 m to perform the additional physics programme in late 2008. This new sub-system grouping is known as ATLAS Forward Physics (AFP). Internal ATLAS review meetings were held in February and June 2009 with the submission of relevant paperwork. A decision is expected from ATLAS later this year to proceed to a TDR to the LHCC. Note that to fully cover the physics programme, detectors are required at both 220 m and 420 m. The physics case for the programme has been given in Section1.3.1.

2.2.2 Detector Layout

The LHC is used as a magnetic spectrometer to analyse the forward scattered protons which have a momentum that is just a few percent lower than the beam particles. These emerge at 220 m and 420 m from the interaction point. Tracking detectors are required to measure the position and angle of the protons at these two positions. The detector stations have to be moved to within a few millimetres of the LHC beams. The solution to this is to use a “Hamburg pipe” . This is a beam pipe with a slot into which the detectors are placed. The whole pipe is then moved towards the beam. The 220/420 regions are drift spaces in the LHC, but the 420 m position is currently cold. To insert detectors, this region must be warm. A conceptual design for a New Connection Cryostat (NCC) was a key output of the PPARC funded FP420 project and enables detectors to be placed at 420 m. Figure 2.1(left) shows the NCC design at 420 m. Two detector stations are required, placed about 8 m apart, to measure the forward protons. Figure 2.1(right) shows the detector station design in more detail. It consists of the Hamburg pipe with motors to move the system to the beam, a silicon tracking station inside a secondary vacuum, and a timing detector (GASTOF). Precise timing detectors (10 - 20 ps resolution) are required to reduce the overlap background. Overlap backgrounds fake the exclusive signal through the combination in one bunch crossing of three separate interactions that contain two opposite arm protons and a central hard scatter. This background type scales with the square of the number of interactions. Two types of timing detector have been developed by AFP collaborators - a quartz bar Cherenkov with spatial segmentation (QUARTIC) and a gas Cherenkov with no segmentation (GASTOF). The GASTOF and QUARTIC will be used in the front and back tracking stations respectively.

Figure 2.2(left) shows the silicon tracker in more detail. It consists of several tracking layers, called ‘superlayers’. A superlayer consists of an AlN substrate onto which ATLAS pixel sensors are mounted. These are bump-bonded to front-end readout chips. Signals are taken via flex connections to a board containing an ATLAS MCC chip. Signals from this are transmitted out of the secondary vacuum box to local electronics and then via optical links to the data acquisition board (ROD) 220/420 metres away. Heat is conducted via the substrate to a copper block and from this to an external heatsink. An active silicon area of around 24×5.5 mm^2

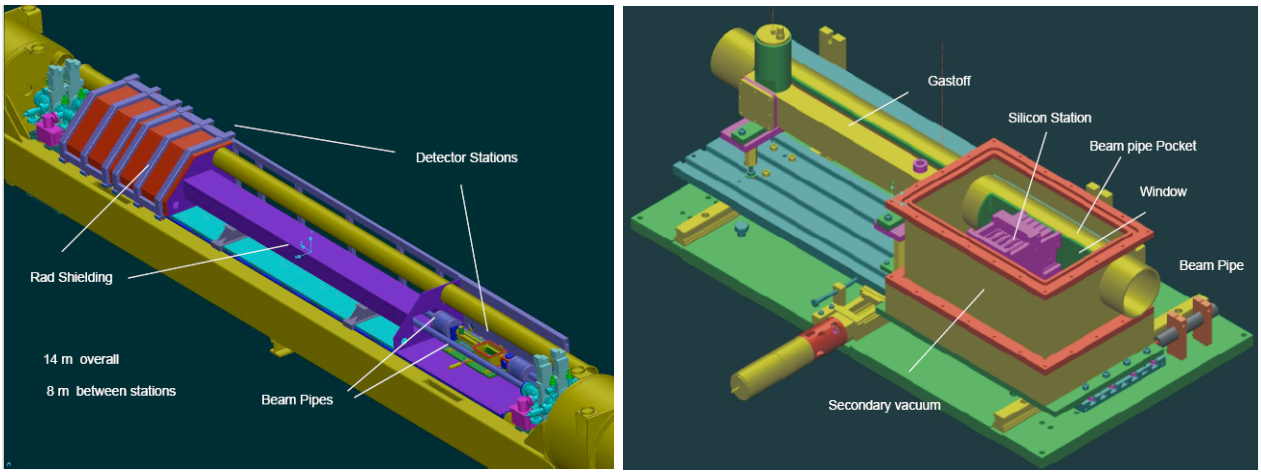

Figure 2.1: Left: The New Connection Cryostat at 420 m indicating the position of the detector stations. These are placed about eight metres apart. Right: A detector station showing the Hamburg Pipe, secondary vacuum for the silicon tracker, and a timing detector (GASTOF).

and 24×12 mm^2 has to be covered at 420 and 220m respectively. The full design is described in detail in the FP420 Design Report. This includes details of the assembly techniques, jig designs, and thermal modelling. The silicon tracker must provide an angular resolution of 1 μrad - roughly 10μm precision in two stations placed 10m apart. The collaboration has already performed much thermal modelling both by computer and on physical models. Pre-production jigs have been manufactured and a test assembly using thermal-mechanical mimic chips rather than expensive sensor/FE-I3 detectors will commence shortly. The thermal-mechanical chips were manufactured at Stanford. A "blade", which is half a superlayer, was tested on a CERN test beam on 2007. Figure 2.2(right) shows thermal/mechanical modelling performed at Glasgow. The sensor has to be held at -20°C and the whole arrangement must be mechanically stable. The figure shows the flex circuits that take the signals to the MCC boards.

2.2.3 Sensors and the Superlayer Design

3D sensors are the chosen technology for AFP because of their radiation hardness and the fact that the dead silicon at the edge can be made as small as 5 microns or 200 microns wide for full-3D or 3D-DDTC technology respectively (see Section 2.4.1.3). For acceptance reasons, the sensors must operate close to the beam and every mm gained is vital for the acceptance.

Using the current ATLAS pixel detector design read out by an FE-I3 chip requires three and six sensors per layer at 420 and 220 m respectively. The current FE-I3 420 m superlayer design actually uses four sensor/FEchips read out by one MCC ship. This is shown in Fig. 2.3.

The sensors match the front end chip whose dimensions are shown in Fig. 2.4. An FE-I4 superlayer design would use two chips and would be suitable for both 220 and 420 tracking stations. The radiation levels in AFP are most intense closest to the beam - around 10^{15} charged hadrons cm^{-2} per year for a luminosity of 10^{34} cm^{-2}s^{-1}. The innermost sensors are most affected. The 3D sensor is radiation hard to 10^{16} charged hadrons cm^{-2}, but the FE-I3 is only tolerant to 10^{15} cm^{-2}. Calculations show that by moving the superlayers vertically, one can spread the dose and gain a factor three extra lifetime for the FE-I3 chips - or three years at

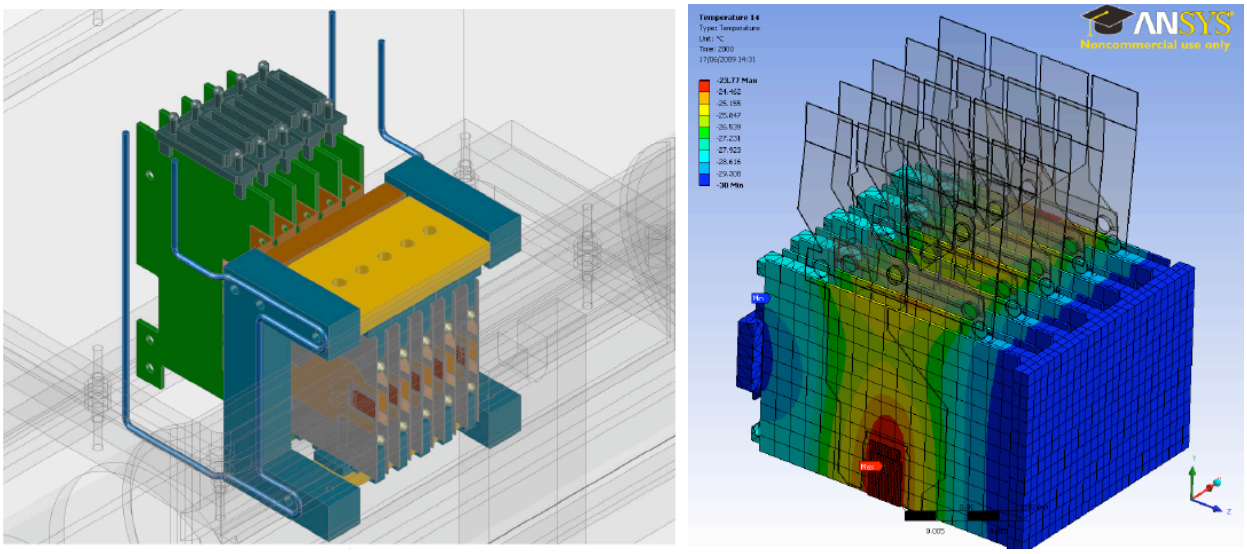

Figure 2.2: Left: Silicon tracker design. Each silicon sensor has dimensions of approximately 7×8 mm^2. The figure shows the MCC boards, cooling and superlayers. Right: Computer Finite Element modelling of the thermal distribution in the tracker. This figure shows the flex lines that take signals from the silicon detectors to the MCC board.

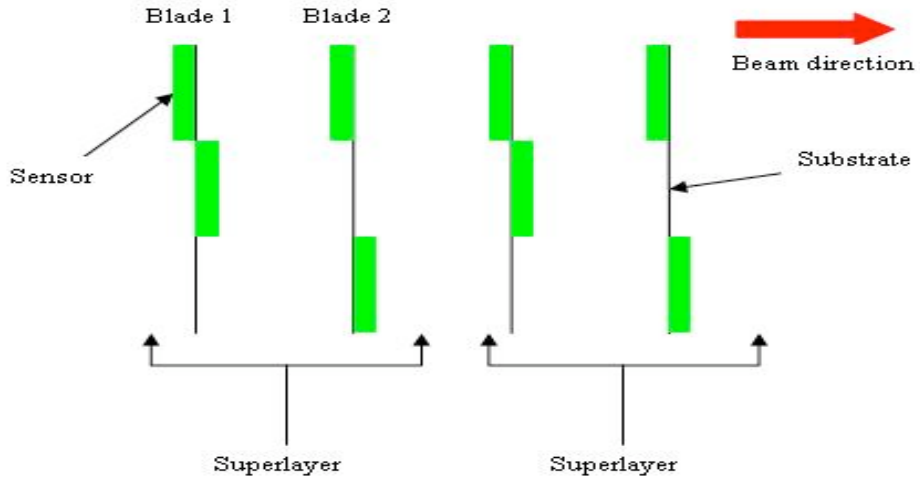

Figure 2.3: Dimensions of the FE-I3 and FE-I4 front end pixel readout chips.

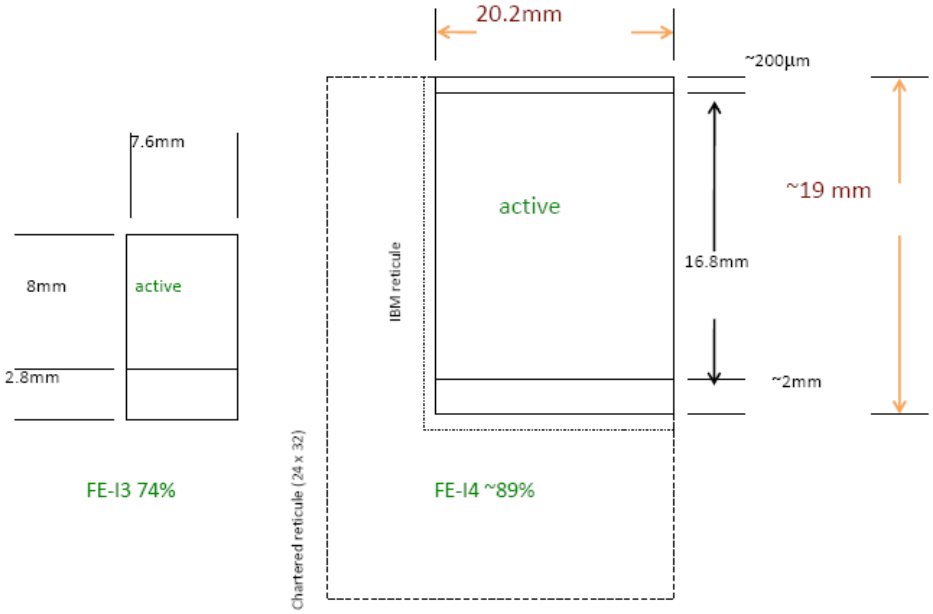

Figure 2.4: FE-I3 superlayer design. It consists of two blades each with two sensors which are readout by a single MCC board. Note that more sensors are used closer to the beam.

$L = 10^{34}$ cm^{-2}s^{-1}. FE-I4 chips are a factor four more radiation tolerant due to the use of 0.13 micron technology. Together with vertical movement of the superlayers, innermost detectors using the FE-I4 would survive for ten years at $L = 10^{34}$ cm^{-2}s^{-1}.

Two R&D matters require completion before the AFP production starts. Firstly, to evaluate FE-I4 sensor/FE combinations. This will require a new superlayer design and development of software and hardware to readout the FE-I4 chip. If this is not complete by early 2011, then the FE-I3 design will be used. Secondly, to finalise the process for bump-bonding chips to 3D sensors with etched active edges. Devices closest to the beam will require this process. Both these R&D issues are part of the 3D workpackage described earlier.

2.2.4 Tasks

To build the AFP subsystem various parts need to be constructed. These are detailed in Section 5.3,WP1.

The UK will take responsibility for building the silicon tracking stations at 420 m including the electronics and data acquisition, the wire alignment system, and monitoring equipment for the Hamburg pipe. With CERN machine physicists and the University of Louvain, we will finalise the design of the Hamburg pipe and construct a pre-production pipe. This will also be used to construct and test a full instrumented AFP pre-production station. The UK will take the lead in assembling the final stations, including all the parts made by other countries, and their test and commissioning on a CERN test beam and in the LHC.

2.3 Strip Tracker (Phase-II)

2.3.1 Overview

The existing STFC supported R & D project into the tracking upgrade at ATLAS has resulted in major progress in understanding the requirements and technologies needed in the challenging sLHC environment. One focus of the programme has been development of the short-strip region of the barrel inner detector tracking upgrade. Many issues have been addressed, with the UK having major input into the international decision to adopt a 'stave' approach to the short-strip region. The anticipated final deliverables of the existing programme are a fully tested realistic thermo-mechanical prototype assembled in the UK and a first fully instrumented electrical prototype assembled as part of an international programme (STAVE09). The UK has leadership in several areas of the tracking upgrade and the work is described in detail in the reports to the oversight committee [20].

The next phase of the short-strip programme, which covers the period 1/4/2010 to 31/3/2013, is a major component of this request for resources. After the final confirmation of the physics case for an upgraded machine and detectors, which will result from the early running of the LHC, the ATLAS UK collaboration intends to bid for resources to manufacture a large fraction of the short-strip tracker for sLHC. The short strip tracker is based on a large area (0.12 m^2) composite double-sided detector element termed a 'stave', which is described below. The largest fraction of the requested resources, and one of the major purposes of the programme proposed here, is to fully define the final 'stave' in the international arena, and to manufacture the first production staves using mass-production tooling at the sites which will do the final assembly work. Developing an efficient and cost effective strategy for division of work amongst available sites will be one of the significant outcomes of the programme. At the end of the planned work ATLAS-UK should be in a position to start mass-manufacture of the final short-strip design, contingent only on the international situation (See Section 1.1). It is the goal that two clusters of institutes be formed, a northern cluster with Liverpool, Lancaster, Sheffield, Glasgow, Edinburgh, ATC and a southern cluster with Birmingham, Cambridge, Oxford, QMUL, RAL, UCL, as was done for the barrel and forward systems for the current ATLAS SCT. Each cluster will be capable of the full production process, which will ensure that the UK will be robust in meeting its production rate obligations and provide cross-checks on quality control standards. Production consists of taking sensors, FE ASICs, hybrids, supports, and services and producing modules and then mounting these modules onto staves. Fully tested staves will then be sent to CERN for assembly into the barrel support structure.

To reach these goals it will be necessary to take the existing R&D results and, in close collaboration with international colleagues, refine the prototype staves into an optimised, cost-effective, detailed design which is compatible with the other components of the upgraded inner detector and with realistic assembly scenarios. All the issues associated with alignment, radiation damage, power distribution and heat removal need to be fully resolved. These latter demand continued work on the services internal to the ID, but the full range of required activity is described in the following sections.

At the end of the programme, five full scale pre-production short strip staves will have been manufactured, using all the final tooling and QA procedures (one confirming the final production process, then two at each cluster for qualification).

To aid clarity of purpose in this document, the work associated with the forward pixel pro-

gramme is described in a dedicated section. However because of the overlap between elements of the strip and pixel work, work of the strip sensors will be undertaken in close association with that for planar pixels (Section 2.4) in order to make most efficient use of resources.

2.3.2 Short Strip Detector Design

The following sections relate to work on a 'stave', which has been the subject of extensive R&D over the past $\sim 2\frac{1}{2}$ years. To help clarify the relevance of the work described in those sections, a brief description of the 'stave' concept, as currently pursued in the UK, is given here.

Figure 2.5, below, shows a schematic cut-away diagram of a stave. It consists of a stave core which is composed of two three-ply 270 μm thick carbon fibre skins measuring 120 cm by 12 cm sandwiching a carbon fibre honeycomb of 5 mm thickness to form a rigid 'plank'. This core has embedded in it cooling pipes surrounded by a thermally conductive carbon foam (Pocofoam). The cooling for the detectors and electronics is via thermal conduction to these pipes, which in the baseline design are stainless steel of 3.2 mm internal diameter and 200 μm wall thickness. The pipe is hard-bonded using epoxy into 5 mm $\times$ 10 mm pocofoam blocks. The sides of the core are closed using carbon fibre c-channels.

During their manufacture the carbon fibre skins are co-cured with kapton tapes which carry the electrical signals and power lines needed to service all the detectors which will be mounted on the stave core. At one end of the stave the core and kapton tape width is increased by 40 mm, producing a 'shelf' onto which the opto-electronic and power distribution interfaces are mounted.

Directly glued to the stave core are the silicon 'modules', which are also shown in Fig.2.5. The sensor is a 100 mm $\times$ 100 mm silicon wafer patterned with strip detectors. Mounted on the wafer are electrical hybrids, which carry the ASICs which read-out the strips. Electrical signals from the ASICs run along the kapton tape to the MCC interface mounted on the 'shelf' where the information is multiplexed and encoded into optical data and transmitted off the detector.

The stave is attached to a carbon support cylinder through a mounting and locking mechanism which is an integral part of the side of the stave opposite the service shelf. A crucial part of ensuring a viable design for the stave is to understand how the services can be connected, installed and maintained. Figure 2.6 shows a design of the current end-of barrel concept to illustrate the complexity inherent in the region between barrel and endcap strip detectors. It is planned for the final integration into ATLAS that there be no cooling connectors in this region, but rather that pipes are welded in-situ using orbital TIG welding. This is a technique which the UK is pioneering in this context; excellent results have been obtained to date, but final qualification will be an extended process due to the proximity of sensitive front-end electronics to the weld locations.

2.3.3 On-Detector Systems

2.3.3.1 Strip sensors

During the R&D phase from Spring 2007 to date, a large number of miniature (roughly 1×1 cm^2) p-type bulk, n-strip readout microstrip sensors have been manufactured, along with 27 large sensors with the current baseline final geometry.

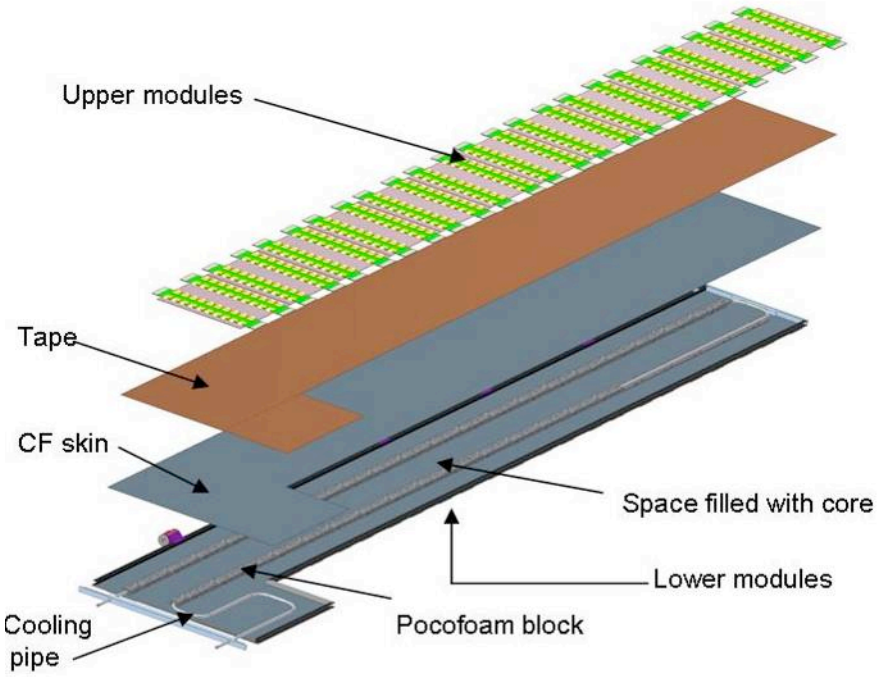

Figure 2.5: Schematic of a Short Strip Stave

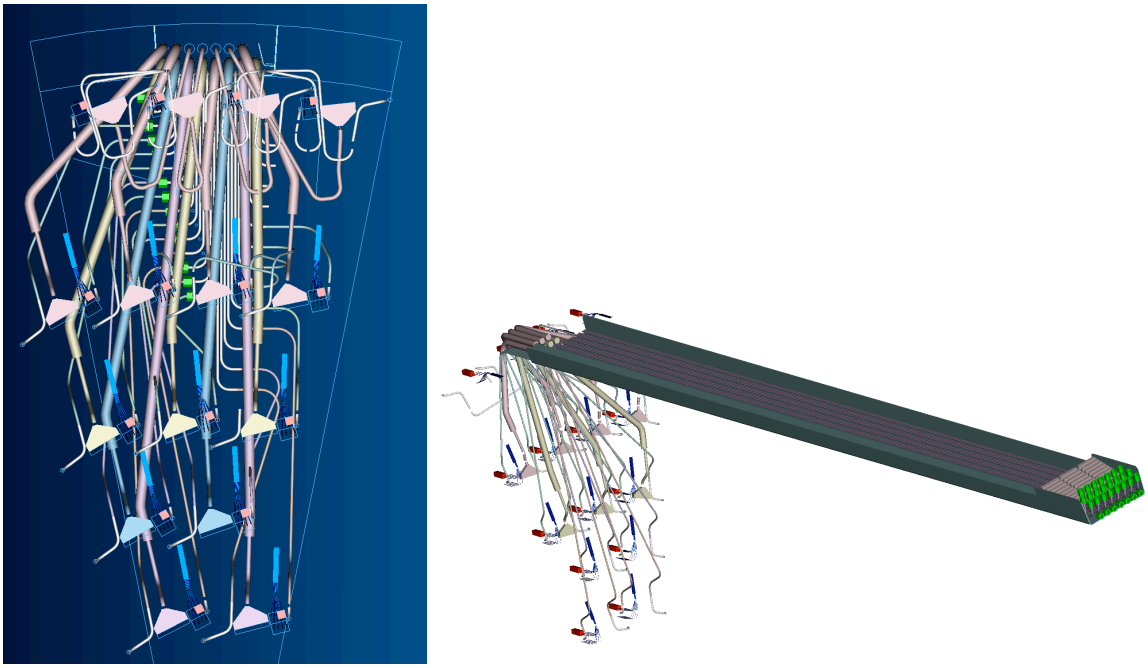

Figure 2.6: Short Strip End Region - showing complexity of services for one octant of barrel (left) and the service module proposed to simplify assembly (right).

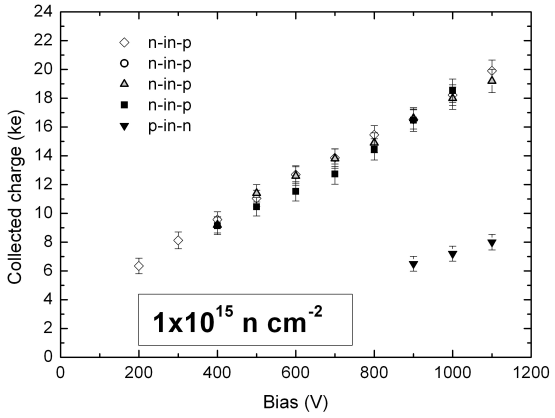

Figure 2.7: CC(V) of n-in-p and p-in-n sensors after $1 \times 10^{15} n_{eq}$ cm^{-2} [23].

The miniature detectors have been irradiated to twice the expected dose of the innermost short-strip layer (at 38 cm radius in the present tentative layout of the experiment) i.e. to $1 \times 10^{15} n_{eq}$ cm^{-2} or 1 MGy for 3000 fb^{-1}of physics luminosity. The miniatures have the same readout type and geometry (strip pitch and width) as the final large area sensors but with a variety of interstrip isolation structures (p-spray only, p-spray + p-stop and p-stop only with different doses of implantation, individual or common p-stop geometries etc.); this has allowed investigation of the implant strip isolation quality after irradiation. The results of the tests (CV, IV, Interstrip capacitance and resistance, Charge Collection Efficiency and Noise measurements with LHC speed electronics) performed after proton and neutron irradiations, have led to the conclusion that this silicon substrate choice is adequate for the strip sensors in terms of radiation hardness. The use of the p-type single crystal silicon with implanted n-type readout strip (n-in-p geometry) was first proposed and pioneered by UK institutes [21] and it has shown the expected advantages in terms of radiation tolerance and lower production cost with respect to the possible alternative technologies [22]. (The p-in-n geometry of the present ATLAS-SCT detectors does not provide the required radiation hardness, and the n-in-n geometry of the current pixel sensors requires a much more expensive processing). Figure 2.7 shows the charge collected as a function of the bias voltage (CC(V)) with n-in-p and p-in-n sensors after $10^{15} n_{eq}$ cm^{-2}, which is the qualification dose for the inner short strip layers in the ATLAS upgraded barrel tracker. A signal to noise in excess of 15 for a 500 V bias is achieved after this dose, after allowing for the performance of the present ABCN25 ASIC.

The comparison of p-type sensors to other substrates and geometries has been extended to very high doses of hadron irradiation. Figure 2.8 shows the signal degradation measured at 500V and 900V after various proton and neutron irradiation, and demonstrates remarkable signal collection by n-in-p devices after the highest doses [24].

The remaining issue to be further investigated as part of the current tracker upgrade R&D programme is the interstrip isolation structure, which must be optimised for best performance before and after irradiation. Further irradiation campaigns with miniature detectors are planned for this purpose with protons, neutrons and mixtures of these to reproduce as closely as possible the real experimental conditions. A small number of the Hamamatsu large area sensors will also

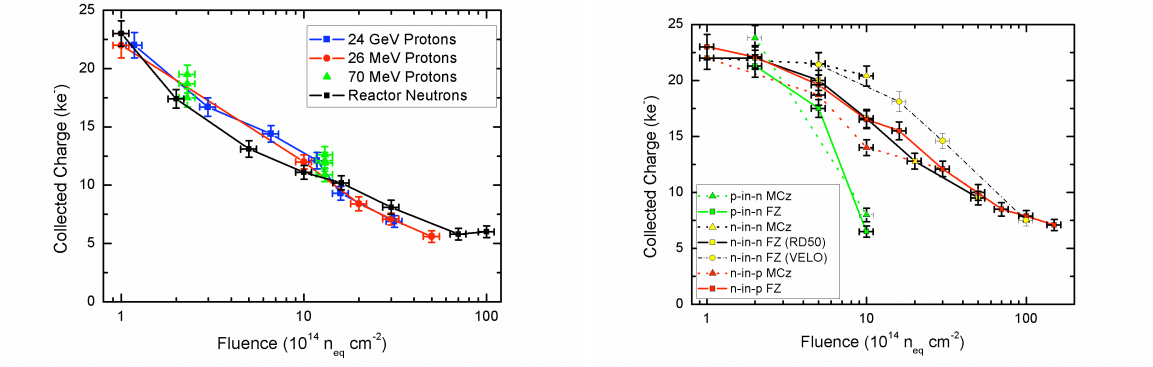

Figure 2.8: Degradation of the signal as a function of the proton and neutron fluence (left, measured at 500V) and of the neutron fluence (right, measured at 900V) for various type of microstrip sensors up to $2 \times 10^{16} n_{eq}$ cm^{-2}.

be irradiated with protons to qualify realistically sized detectors. The operating voltages for the sLHC will be selected.

The work reported above, will merge into the next phase of the programme, for which funding is requested here. It will be focused on preparation for production of the final staves for the experiment. The large number of sensors and high channel density implies the set up of well defined procedures for acceptance, quality control and handling of the detectors. These will be required from delivery of sensors from the foundry, through module production to mounting on the staves. Of particular importance will be the definition of the quality tests to be performed on the devices and the preparations for efficiently performing these tests. The quality assurance (QA) tests have the function of confirming that delivered devices conform to the specifications issued by the ATLAS upgrade community. Such a QA programme was successfully undertaken in the UK for the existing ATLAS SCT, however the increase in the number of read-out channels per device and in the number of devices requires an enhanced rate of testing. The peak rate of detector testing in the UK is anticipated to be 25-30 per week ($>$ 150000 individual readout channels).

To enable the QA to be undertaken, extensive tooling must be assembled and commissioned. This includes mechanical jigs and fast automatic probe stations equipped with semiconductor characterisation instrumentation. The details of these procedures will be developed as part of this programme. They will be based on experience accumulated during the characterisation and radiation tests of the large area devices acquired as part of this programme. The number of large area devices to be acquired will allow optimisation of the speed and efficiency of the procedure for rejecting unsuitable devices. It can be anticipated that the tests will check for; strip quality, defined in terms of the number and defect type of non-functional strips; compliance with the specification on leakage current at full voltage (absence of microdischarge currents); and capacitance characteristics.

The programme foresees production of five complete barrel staves, which is 1% of the total, requiring 120 full size silicon sensors, of which half will have straight and half stereo angled strips. The plan is to acquire 150 silicon sensors (75 of each type) to accommodate for

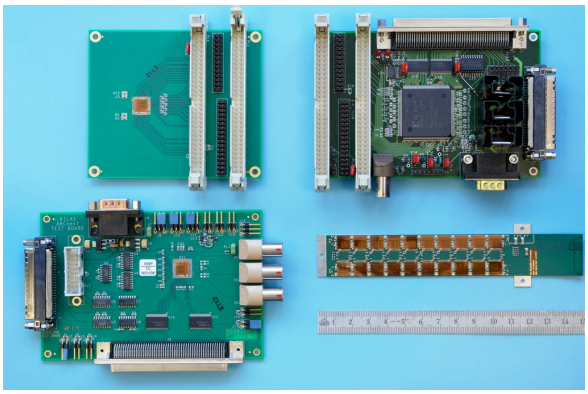

Figure 2.9: Printed circuits for ABCN25 ASIC testing. Top left and right: probe card compatible single chip board and driver PCB. Bottom left: single chip test card. Bottom right: 20 chip hybrid.

spares and for sensors to be dedicated to initial module studies and further irradiation testing. These are essential further studies as the new ASICs, in contrast to the existing versions, will be fabricated in 0.13 μm CMOS technology. This number of devices will also provide sufficient statistics for the optimisation of the QA procedures.

2.3.3.2 ASICs

The readout architecture of the short-strip tracker is a development of the binary readout architecture used in the present SCT, but with increased readout bandwidth for use at higher luminosity. The chipset comprises three chips on the hybrid - the ABCN readout chip, the MCC module controller chip and a power protection chip (yet to be named, here referred to as the SPi chip). An additional stave controller chip is envisioned for use at the end of each stave. Each of these components must have sufficient radiation tolerance to withstand the lifetime dose expected during sLHC operation.

Last year, the community successfully delivered the ABCN25 readout chip in 0.25 μm CMOS technology for use during the R&D phase of the short strip tracker. The UK made significant contributions to this programme, having provided specifications for some of the chip's circuit blocks, delivered the test PCB used for initial evaluation of the chip and been the first to demonstrate the chip's functionality, Fig. 2.10.

This year, the community has delivered a prototype version of the Module Controller Chip known as the Basic Controller Chip (BCC), also fabricated in 0.25 μm CMOS technology, which has just sufficient functionality to enable full stave read-out. The UK specified the chip and delivered a hardware description model of its digital core which was successfully synthesised and included in the design. As part of the programme proposed here, a Field Programmable Controller Chip (FPCC) will be developed in the UK. Based on a commercial FPGA device, this will allow a thorough evaluation of options for the MCC to take place before the final, radiation-hard ASIC design is submitted.

The target technology for the sLHC tracker upgrade is 0.13 μm CMOS, a decision driven by the power savings associated with smaller feature size processes coupled to the reduced

availability of older processes beyond 2010. Test structures are presently being designed for delivery and evaluation in 2010. These ASICs are crucial to the UK programme, and their progress will be monitored closely.

Specification of the first ABCN13 and MCC circuit blocks will take place through year 10/11, one element of the programme proposed here is to develop the concepts required with regard to powering, data formats and possibly the prioritised readout of track trigger events. It is expected that first versions of the complete 0.13μm ABCN13 and MCC chips will be delivered for evaluation in year 11/12, when the UK will again lead the testing of these chips. PCBs will be manufactured to be used during initial evaluation. Later in the programme screening of the pre-production ABCN13 wafers will be performed.

Similarly the UK's involvement with the SPi chip will lead naturally through to the specification, design, delivery and testing of the proposed protection chip, which is described in more detail in Section 2.3.4.2. Approximately 6000 ABCN13 and 300 MCC ASICs are required for use in the proposed pre- production phase described in this document.

2.3.3.3 Hybrids

The hybrids provide the needed electrical, mechanical and thermal interfaces between the silicon sensors, the readout ASICs, and on-detector electronics. The UK is leading the international hybrid programme within the ATLAS strip upgrade. The hybrids for the current international R&D phase have been designed, fabricated, assembled and tested in the UK. These have been the primary test bed for understanding the effects of different powering schemes and have improved the understanding of the 0.25 μm ABCN25 readout ASIC.

The hybrids have been designed and fabricated in consultation with UK industrial partners (Stevenage Circuits, Ltd and Hawk Electronics, Ltd.) in order to ensure high reliability, high yield and low costs. The current design is a substrate-less, multi-layer kapton hybrid, and is based on previous work done by the ATLAS barrel pixel group. The technique has been modified to enable production of industrial quantities. In this model 6-8 hybrids are produced in sheets approximately the size of A4 paper. The hybrids then have the passive components populated using a reflow method. Finally, the hybrids have die-attached, are wire bonded and tested in the sheets in house, at which point they are ready to be cut from the sheet and used for module assembly.

Until the new 0.13 μm chipset is available (year 2011/2012), the focus of hybrid development will be on improving the quality and speed of large-scale production and testing the processes using the 0.25 μm ABCN25. This will include migrating more of the powering and control circuitry to the hybrids. Several rounds of panel submission, population and testing are required. Schemes for automated testing of entire panels in a serial power chain will be developed. This will allow full utilisation of the advanced features of the HSIO DAQ (see Section 2.3.4.1), increasing testing speed and decreasing manpower needs for final production using the 0.13 μm chip set. Automated testing of entire panels will be the first necessary step to migrating die-attachment, wire bonding and testing of hybrids to UK industry. Tooling for these processes will also be required at each of the two proposed production clusters.

It will be essential to work closely with the designers to ensure that the specifications and layout of the 0.13 μm ASICs are conducive to large-scale production and that any added track-trigger functionality is compatible with hybrid development

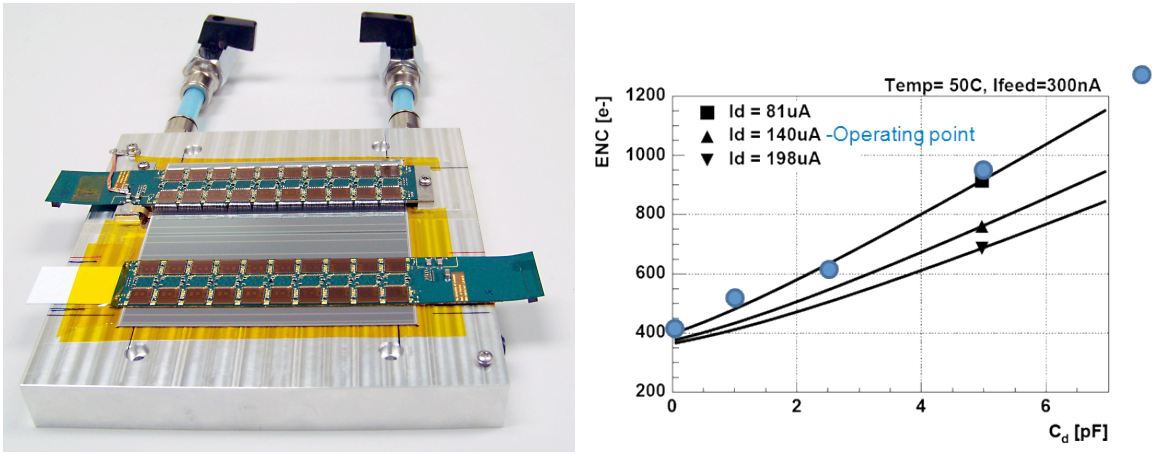

Figure 2.10: Left: Prototype module using ABCN25. Right: Noise vs. capacitance from prototype module.

The second major phase of the hybrid programme starts when the 0.13 μm chipset is available. All the lessons learned from the 0.25 μm developments will be incorporated into a 0.13 μm design and one or two small-scale 'proving' submissions of the hybrids made. It is expected that these will incorporate the 0.13 μm ABCN13, MCC and power protection ASICs and include a prioritised readout chain for a track trigger. Confirmation of the success of the switch to 0.13 μm and of incorporation of any track trigger driven design changes is a crucial step and the hybrids will again be one of the primary test vehicles for the chip sets. After full testing, $\sim$ 300 hybrids will be manufactured for the UK pre-production stave programme. It is anticipated that the UK will also facilitate the supply of hybrids to the international upgrade programme.

As part of the pre-production exercise, it will be necessary to demonstrate that the required throughput of one finished hybrid panel per day can be accomplished. This demands that fixtures and testing systems be built.

2.3.3.4 Modules

The modules are the first assemblies which can be tested as detector elements. They consist of 2 hybrids directly glued to $\sum 10 \times 10$ cm^2 silicon sensors. The UK is leading the international barrel module upgrade project and has built the first, and currently only module for any sLHC silicon upgrade project (Fig. 2.10). Design and prototyping of all the module assembly, wire bonding, and testing fixtures for the international ATLAS stave upgrade project is being done in the UK as part of the existing funded R&D programme.

The first phase of the programme proposed here is to produce modules using the 0.25 μm chip set and work towards devolving module production for the STAVE09 programme at 4-6 international sites, including in the UK. Modules produced in the UK will be evaluated with those from other international sites in order to enhance yield, reliability and throughput for future module designs and productions.

The second phase of the programme proposed here will be the manufacture and testing of a small number of modules using the 0.13 μm hybrids. These modules will be used to verify the performance of the final system and to determine if modifications are required to fixtures or test

equipment.

The final phase of the programme will be focused on developing the tools needed to assemble, wire bond and test production quantities of modules at each of the two proposed clusters and to demonstrate that production rates of 4-6 modules per day per cluster can be achieved. To enable such production rates, it is planned to enhance the single module fixtures to be capable of producing a day's throughput per fixture. When trials have demonstrated the viability of the evolved jigging, several sets will be fabricated and ~ 120 modules will be fabricated at the two clusters. Electrical tests of mass produced modules will be necessary. Experience with schedules has shown that it is also necessary to develop methods to safely store large numbers of tested modules.

For mass production, both clusters will need 2-3 modern wire bonders with a rate capability of at least 2-3 wire bonds per second, state-of-the-art pattern recognition and remote touchdown capabilities.

2.3.3.5 Tapes

Low voltage, high voltage and temperature monitoring signals will be distributed on the staves by low mass Cu/Kapton flat circuits (tapes). The tapes will also have transmission lines for the readout of data from the hybrids as well as for sending the Timing, Trigger and Control (TTC) to the hybrids.

Realistic prototype bus tapes for the short strip staves have been designed and fabricated (see Fig. 2.12). The tapes have been used to study data transmission properties of this construction method. Tests have shown very clean data transmission is possible at the required rate of 160 Mbits/s. Successful trials of multi-drop LVDS have demonstrated that the TTC data can be distributed to hybrids on a bus, which reduces the 36 pairs that would be required for point to point links to 3 pairs.

As part of the current R&D programme, bus tapes will be fabricated for STAVE09. They will be based on the 0.25 μm ASICs and will be integrated with an aluminium screen layer and the carbon fibre stave-core to become a fully realistic prototype.

In order to be able to launch full scale production at the end of the programme proposed here, it will be necessary to investigate yield, reliability and cost issues. There were many problems found with the current ATLAS SCT electrical services which were only discovered during production because they were rare problems. An insufficient number of prototypes had been built to fully explore overall yield issues. Because of the high value of an assembled stave, which is inherent in the stave design concept, it is therefore essential to produce a large number of prototype stave-cores and tapes in order to understand the yield and reliability. This work will form the early part of the tape development in the proposed programme and will require manufacture of tapes in both institute based facilities and in industry. Developing a mass production process suited to the co-curing of tapes to skins, and potentially cover-layer to tape, will be actively pursued and if proven effective the methodology transferred to industry.

Critical to the success of mass manufacture will be development of QA procedures, which will be required, both upon receipt of the tapes from the vendor and also at several stages of stave manufacture. Due to the complexity, high number of connections, and the small size of some of the contact points on the tapes only automated test facilities with high positional accuracy are suitable for this task. Similar devices in industry are usually not large enough

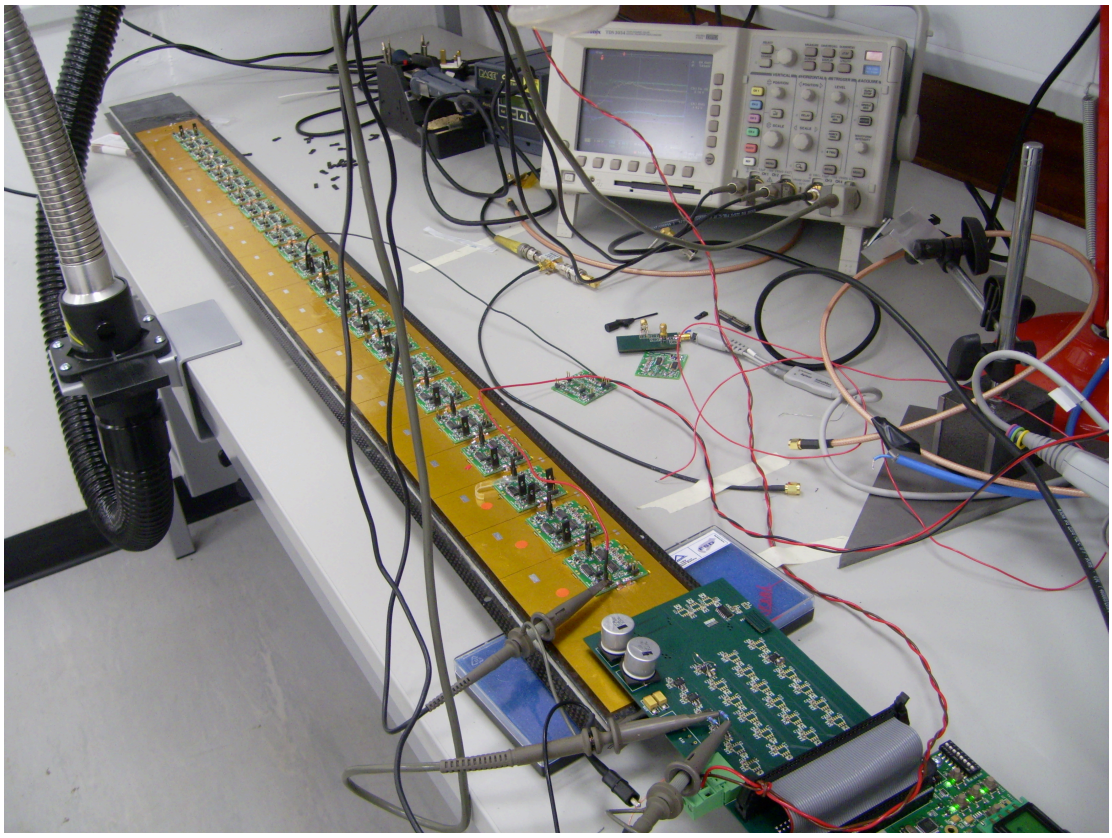

Figure 2.11: Photograph of the prototype stave showing the interface PCB on the right and 24 dummy hybrids.

to cover the full area of a tape as currently envisaged for the future ATLAS strip tracker, so developing such tape testers is part of the programme proposed here. One of these tape test facilities is required per assembly cluster.

All the tapes in the existing R&D plan are based on the 0.25 μm chip set and the detailed layout will need to be redone to suit the 0.13 μm ASICs when they are available. In addition the End of Stave (EoS) interface is not yet fully defined and it will be necessary during the latter half of the programme to adapt the tape geometry to suit the final chips and optical interfaces.

2.3.3.6 On-stave Optical interface

UK personnel are pivotal members of both the ATLAS sLHC strip and pixel architecture working groups, ensuring that the UK is placed to use its expertise with the Versatile Link (VL), to drive the design of the EoS interface. This will include the GBT (Giga Bit Transceiver) which performs the high speed multiplexing, drives the laser and the VL to transmit strip data off-detector. The VL also contains a high speed photo-diode which together with a trans-impedance and a limiting amplifier enables the VL to operate as a bi-directional link; the plan is that it will be used for the transmission of the Timing, Trigger and Control (TTC) data from the control room to the staves.

As part of the programme proposed here collaboration on the VL project will continue,

which will enable sourcing of prototype VL and GBTs for the EoS. These will be crucial for the final pre-production staves that the two UK clusters will assemble. It will also ensure a source of the appropriate versions of the VL for the counting room end of the system. UK knowledge of GBT and VL will be used to perform the thermo-mechanical studies of the EoS, starting with thermal FEA calculations to optimise the cooling for the GBT and the VL. The cooling must remove the heat from the GBT and the VL but avoid cooling the lasers below -20°C if practical. In case the lasers will have to operate below -20°C, it will be necessary to work on qualifying laser wafers for this temperature range.

To validate the FEA calculations for the final GBT and VL designs, it will be necessary to manufacture thermo-mechanical prototypes. This thermo-mechanical design and test work is a pre-requisite for completing the layout for the stave Cu/kapton bus-tape, including the EoS region.

For the final prototype stave with optical readout, test systems to perform the QA for the optics for the 2 UK clusters will be required. These will purchased from our VL collaborators and qualified as part of the programme proposed here. The systems will be based on FPGAs and perform BER tests. These test systems will then be operational for the full stave production.

2.3.4 Off Detector

2.3.4.1 Data Acquisition (DAQ)

Within the upgrade programme, the tasks to be fulfilled under the DAQ heading are threefold: Firstly, in the course of the R&D towards the upgrade tracker, hardware, firmware and software should be provided to read out newly developed components based on the ABCN front-end chip: hybrids, sensors and various testboards; Secondly, the DAQ should support Q&A measurements during module and stave production, using application specific firmware and software builds; Thirdly, R&D towards a final DAQ for the upgrade tracker requires performance evaluation of candidate components, leading to a scalable, full-speed demonstrator DAQ. This would include studies to establish the requirements a SCT track trigger would pose on the DAQ system. Compared with the DAQ for the SCT, that for the upgrade will have to cope with at least an order of magnitude higher data rates, driving the need for a re-design of the system.

Given the prominent role that UK institutions have assumed in the ATLAS SCT DAQ, it follows naturally that the UK continues to lead the DAQ developments for the upgrade tracker. The knowledge and experience gained during the development of the SCT DAQ system are essential inputs to the upgrade.

During the current R&D programme, the UK institutions have successfully provided a DAQ for the readout of hybrids and test-boards hosting the ABCN25 front-end chip. This was accomplished efficiently by upgrading the SCTDAQ software and using the existing MuSTARD and SLOG VME based DAQ hardware. For testing of individual ABCN chips, a DAQ system based upon commercial components (National Instruments PXI series) was deployed. SCTDAQ software, specifically adapted by UK institutes, is used for the gathering and analysis of the data.

Future developments towards the upgrade tracker include the next revision of the front-end ASIC, and the inclusion of the MCC (see Section 2.3.3.2 for ASIC definitions) into the readout chain. Last year, a BCC and FPCC were developed to serve as stepping stones towards a MCC ASIC solution. These developments allow for high-speed data transmission, and hence far

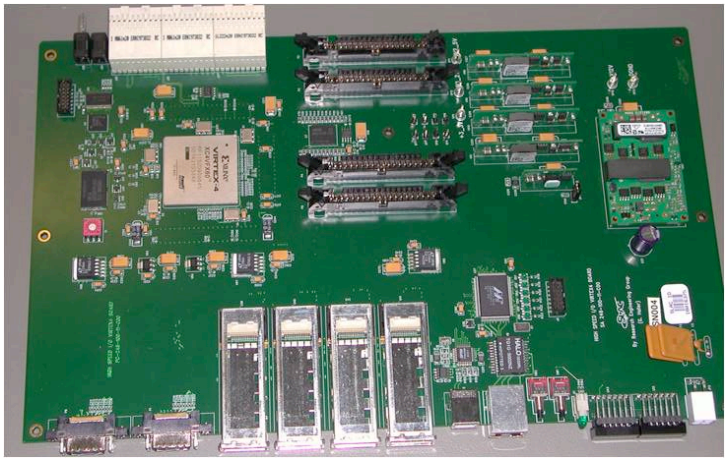

Figure 2.12: HISO board.

exceed the capabilities of the existing MuSTARD/SLOG hardware, necessitating an upgraded hardware platform.

Through a collaboration set up with SLAC, Berkeley and UK institutes, DAQ hardware with sufficient capabilities to cope with future developments has been specified and developed. The main component is the Hi-Speed Input/Output (HSIO) board, hosting a powerful FPGA, and several high-speed interfaces, along with a large connector with many user-defined pins (see Fig. 2.12. For ease of use, low-cost commodity interfaces (copper Ethernet, USB) are used for connecting the HSIO to a PC.

To provide specific interfaces, optimal matching and damage protection between the development hardware and the HSIO, the HSIO is extended with an intermediate board. When necessary, a re-issue of the intermediate board provides a relatively low-cost and rapid way of customising the DAQ to new hardware developments.

As part of the programme proposed here the HSIO will be developed, in continued collaboration with our US colleagues, as the cornerstone component for the DAQ requirements for the inner detector upgrade. As well as supporting module R&D, the HSIO is capable of performing all tests necessary for quality assurance and partial system tests during module and stave production. Therefore the intention is to deploy HSIO based DAQ systems at all institutions participating in module production, and to continue the development of firmware and software for support of both quality assurance tasks during module production as well as the readout of newly developed components.

Due to its high-performance interfaces, the HSIO provides an excellent testbed for DAQ performance studies. These studies will be particularly valuable should it be decided that a track trigger is a requirement for the inner detector upgrade.

A major milestone will be to use multiple, synchronised HSIOs for the readout of the fully instrumented stave prototype, including the radiation-hardened optical link hardware being developed within the VL project (see Section 2.3.3.6). Here, the HSIO will provide the ideal testbed, as the off-detector receiver is designed such that it can be a firmware block residing in a user's FPGA. Also, the possibilities for implementing a trigger interface can be explored using the HSIO as a development platform.

The build-up of knowledge and experience with a well-established high-performance FPGA-based platform will prove indispensable when defining requirements for the DAQ system of the upgraded tracker. By the use of standard interfaces and VHDL firmware code, firmware and software are very likely to be portable to a future DAQ platform tailored to the upgrade tracker.

2.3.4.2 Power

Distributing power to the electronics of the sLHC silicon trackers has been recognised to be a major challenge. The UK has been leading the international power distribution R&D. Because it has been recognised as critical, power distribution has grown into a field with about 20 active groups world-wide, the majority in ATLAS. UK R&D in the current upgrade programme has been focussed on serial powering, with limited contributions to DC-DC conversion. All critical elements of a serially powered silicon supermodule have been developed and prototyped, including: the constant-current power supply; shunt regulators; power management blocks in the readout integrated circuits; and protection circuits. Custom electronics for the AC coupling of multi-drop and multi-point signals have also been developed, as has a serial powering system architecture and a grounding and shielding concept. Prototypes of circuit blocks have been characterised.

The conclusion of our current R&D is that serial powering is a very strong candidate technique for powering the sLHC trackers efficiently and reliably. While serial powering has not yet been selected to be the ATLAS sLHC baseline, it is the leading contender. In the work proposed here the aim is to develop the final system components to be ready for production (see Fig. 2.3.4.2).

This requires: transformation of the powering blocks from the ABCN25 to the ABCN13 chip (see Section 2.3.3.2); integration of the serial power control (SPi) chip's 0.25 μm shunt regulator's op-amp, specified in UK, into the 0.13 μm MCC; implementation of a hybrid voltage and shunt control voltage monitoring system into the MCC; design of a radiation-hard protection chip to replace the prototype discrete protection circuits; design of a voltage limited programmable 'current' source with capabilities to track varying module currents for maximum efficiency and performance; and characterization of pre-production staves and development and integration of power control systems.

The UK groups will specify the power related circuitry blocks in close collaboration with the ASIC designers, and will perform functionality tests and full characterisation studies of components. Characterisation will be performed on each single die, hybrid and system.

The UK contribution to the SPi ASIC will include specification, functionality tests, characterisation and circuit design, in collaboration with one of the ATLAS IC design groups. The design of the 'current' source will proceed as a continued close collaboration with ASCR, where the UK work will focus on layout, construction and development of firmware for maximum efficiency, reliability and minimum noise performance. The programme proposed calls for manufacture of pre-production staves, an integral part of which will be characterisation of their performance, including test and evaluation of the powering scheme.

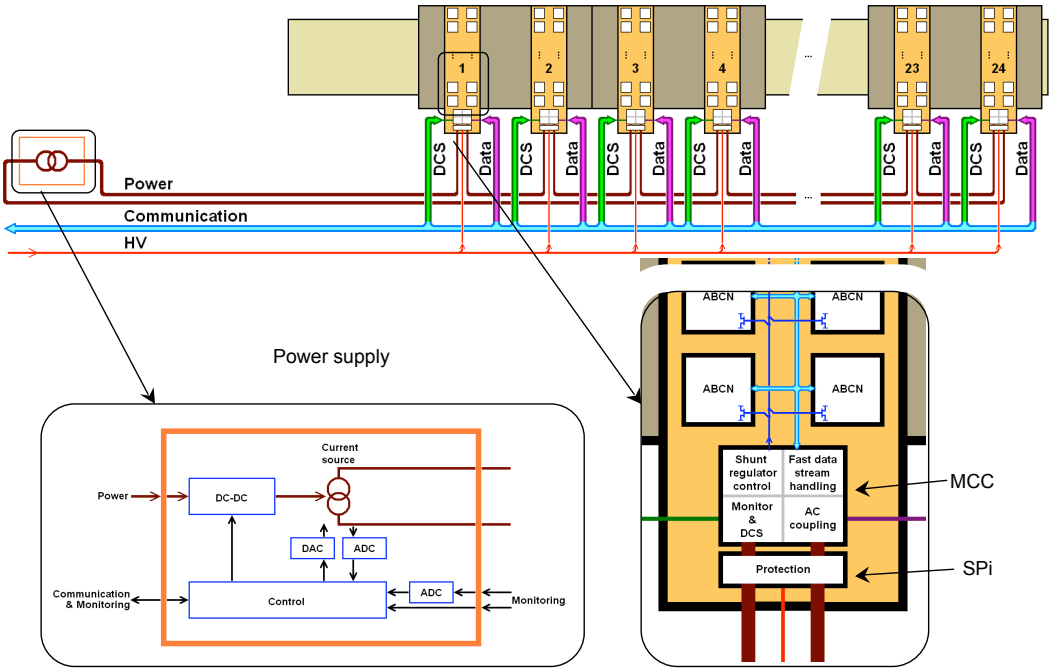

Figure 2.13: Sketch of serial powering elements for sLHC production stave.

2.3.4.3 Passive Optics

The UK is active in the Versatile Link project which aims to produce suitable radiation tolerant optical links for the ATLAS and CMS detectors at the sLHC [25]. Within this project the UK has responsibility for the passive optical components, which include fibres, fibre jackets and cables, optical connectors, optical couplers and optical multiplexers, MUX.

The first radiation tests for fibres to sLHC doses were very successful and have identified commercial fibres as good candidates [26] for the upgraded tracker (see Fig. 2.14). The tests showed some evidence for the temperature sensitivity of the radiation tolerance and further tests and irradiations at cold temperatures are being performed. Because the environment in the future ID will be cooled to temperatures near -25°C, further cold testing will be required to qualify these fibres for use inside the tracker volume. Tests in high-rate gamma ray sources with the Draka RHP-1 prototype multimode fibre gave particularly good results. However this particular prototype requires modification for high-bandwidth operation; therefore these tests will need to be repeated when the higher bandwidth commercial version is released.

Optical fibre is pulled from a silica melt called a ‘pre-form’. Past experience has shown that the quality of optical fibre within one pre-form is highly uniform, but there can be significant differences between pre-forms. For QA in the production phase, it will be necessary to perform radiation tests on samples of fibre from each pre-form. The tests undertaken have shown that annealing effects are very important, so it is essential to monitor the fibre attenuation during the irradiation. The test equipment that has already developed for the R&D stage will need to be expanded to allow for many fibre samples to be irradiated and tested in parallel.

The mechanical performance and reliability of the fibres will need to be evaluated after radiation. Experience of the methodology required has been gained using a pull test machine

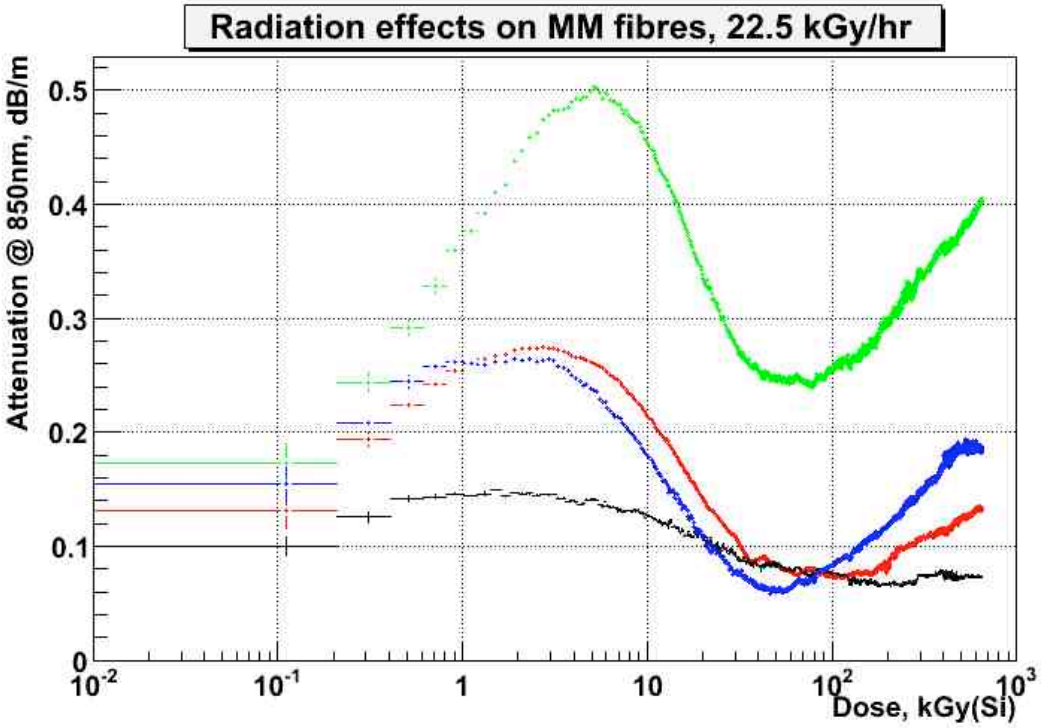

Figure 2.14: Attenuation (dB) per meter of fibre length vs. radiation dose. The Black curve is the Draka RHP-1 An unofficial market survey of optical fibres has been carried out and similar surveys will be done for other passive components. This first R&D phase has resulted in close contacts with the fibre companies Ericsson, Fibresson and Draka.

at CERN. However in order to get sufficient statistics for a suitable reliability study, it will be necessary to purchase a pull test machine in the UK. This will allow study of the reliability before and after irradiation and therefore to validate the mechanical suitability of the fibres in a radiation environment. These tests will need to be repeated on fibre samples from each production pre-form. Similar testing will also be required for the fibre jacketing and cables. Our industrial partners (Draka and Fibersson) are interested in the possibility of working closely with us on a full reliability analysis of the final production fibres as well as commercial fibres.

The first passive and active tests of a small form-factor connector (LC) gave encouraging results. More tests are required to validate the use of passive radiation tests for connectors. Passive tests only require measuring the changes in performance before and after radiation and are thus less expensive. In production, batch testing of connectors will be required; therefore it is essential to develop appropriate jigs for a semi-automatic measurement system. For fibre ribbons, the preferred connector is the MT-12 and equivalent jigs and test procedures will need to be developed for these connectors.

Although neutrons are only expected to degrade the optical properties of the fibres by secondary ionisation, they will degrade the mechanical strength of the fibres, jackets, cables and connectors. Therefore measurements will be required to compare the performance of these components before and after exposure to neutron irradiation.

Passive Optical Networks (PONs) are being considered as a way to reduce the fibre count and therefore the UK will need to qualify suitable optical couplers (also called optical splitters). Such couplers come in two main types; fused taper and Planar Lightwave Chip (PLC) couplers. The current upgrade R&D programme has demonstrated that the radiation tolerance of fused taper couplers is simply related to the choice of fibre used to make them. However PLC couplers are more compact and can achieve larger split ratios. Demonstrating the radiation tolerance of PLC couplers will be part of the programme proposed here. Systematic tests in a well controlled temperature environment will be required. Part of the capability to do this has already been developed in the current R&D phase but modifications will be needed for these tests.

2.3.5 Mechanics

In the R&D programme towards the detector mechanics of the future ATLAS barrel strip tracker a local support or ‘stave’ has been developed (Fig. 2.5). This supports 12 modules on each surface at an azimuthal position over half the barrel length. The UK activities have been fully integrated into the framework of our international collaboration in which the UK is established as one of the key contributors.

In the early stages of the past R&D on this project the UK groups have developed and promoted a super-module design based on a stiff carbon fibre structure with dismountable Silicon modules and a ‘wiggly’ metal cooling pipe to control the stresses introduced by differential thermal expansion. Key concepts from this work have been incorporated into the baseline stave design. Significant effort has been invested in understanding the mechanical challenges inherent in the stave concept. Potential difficulties included: the differential thermal expansion of carbon-fibre structure and metal cooling pipe; potential consequences of bonding hybrids directly to the active sensor side; the location of stave end electronics; the mechanism to lock stave to support cylinder; and the complexity of a tightly bonded sandwich of directional composite materials. Considerable understanding of these issues has been acquired, with the design

of STAVE09 and thermo-mechanical staves as the result. However, it will be necessary to devote some attention to these issues in the future.

There are three main goals to the mechanical programme proposed here. The first is to develop and finalise the fabrication methods for the staves including integration with the requisite ancillary systems. This includes developing the mass manufacture methodology. The second major goal is development of the process for attaching strip modules and optical interface components to the support structure which, because of the high value of assembled staves, must be accomplished with extreme reliability. The third goal is to develop the UK infrastructure required and the quality control methodology for the mass manufacture stage of the strip tracker build. This third goal includes establishing the organisational structure and manufacturing base for the mass assembly phase.

To accomplish these three major goals a number of preliminary steps need to be taken. The prototype staves which will be extant at the end of the existing R&D programme will demonstrate the viability of the concept, but will need extensive refinement in a number of areas as the final, internationally agreed, design is evolved. To achieve the requisite tolerances, low mass, thermal performance and mechanical stability in a cost effective way, the methods required to manufacture and assemble such structures need significant further R&D. In addition there is scope for enhancing the performance and reducing the cost of the final stave by further investigation of materials and details of design. Because of the high stresses that are inherent in a hard-bonded stave design, it is believed that careful finite element analysis (FEA) study and prototyping of the minutiae of the engineering detail will result in the highest probability of success and the lowest overall cost to the programme.

2.3.5.1 Finite Element Analysis

FEA calculations have been undertaken as part of the existing R&D tracker upgrade programme. These are essential to guide the design process and radically increase the probability of a prototype being successful. FEA is used for thermal and stress analysis for modules, hybrids, stave-cores and ensembles. Potential areas of failure are identified as predicted high stresses in cooled structures and these are studied in detail both computationally and in prototypes built specifically for the purpose. The FEA studies have been successful and crucial in designing the stave. For example Fig. 2.15 shows simulations of the thermal performance of the modules near the side-mounted SMC.

In the future programme, where fine-tuning of designs is anticipated, it will be crucial to maintain FEA support. Experience shows that small changes in design can have a big impact on performance. For example the failure of glue bonds on carbon fibre cylinders in the existing SCT was ascribed to glue layers being too thin, something which is clearly understandable using FEA. The final, internationally agreed, design will evolve in numerous small ways over the three years of the programme and it is essential that personnel with requisite FEA skills be available to compute the consequences of changes.

2.3.5.2 Materials Selection

There are two main areas to be addressed in materials selection which are: the thermal performance of interface materials and materials in the thermal path; and the dimensional stability and strength of the structural items. These characteristics must be measured both before and

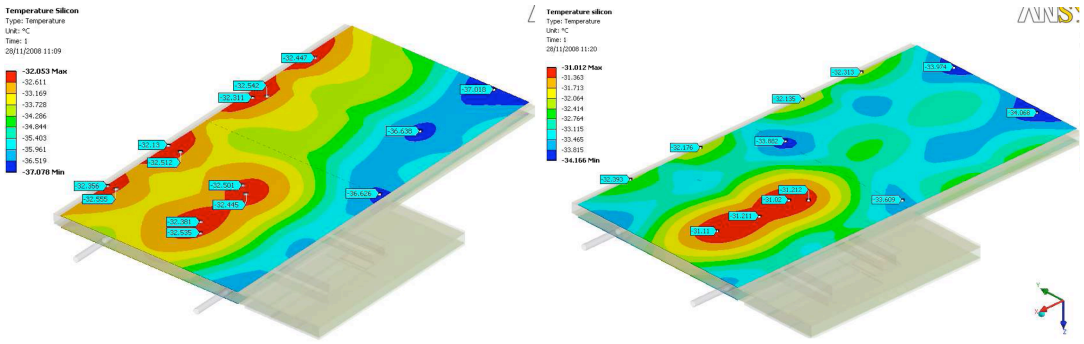

Figure 2.15: Showing thermal impact of side-mounted SMC on last module with coolant temperature dropping round cooling circuit (left) and with constant temperature coolant (right).

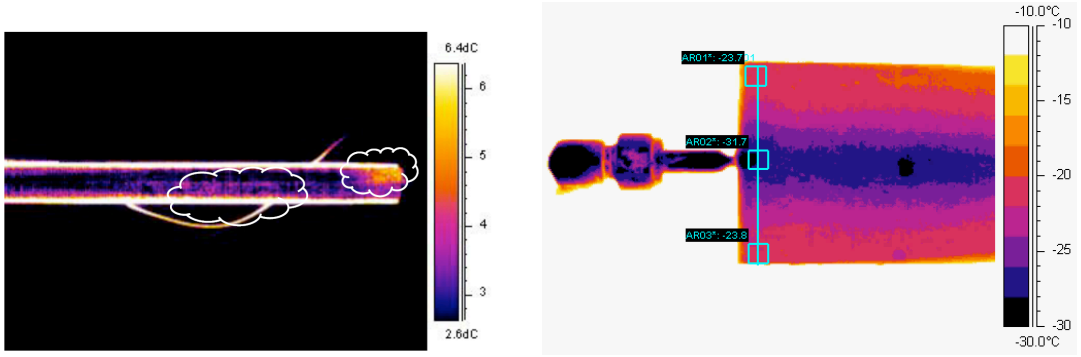

Figure 2.16: Left: Evidence of poor thermal connection between PocoFoam and the stave's skin when bonded with CGL (side view). Right: Thermal profile across stave section with coolant at $-32°$C (frame of video clip).

after the expected radiation dose. The performance of the materials in the thermal path from the module to the cooling pipe is crucial to enable the lowest possible thermal resistance for the smallest amount of material. The dimensional stability and strength of the materials that form the stave are required to be high for as small a number of radiation lengths as possible so that the structure is stable whilst introducing a minimum of dead material into the tracker.

The present ATLAS upgrade project focused on the development of the apparatus and techniques to measure the thermal and mechanical properties of materials that were candidates for the stave build. Several of the proposed materials have been measured (see for example Fig. 2.16(left) which shows a measurement of the performance of a candidate 'glue', CGL; and Fig. 2.16(right) which shows a frame from a thermal image video study of the cool-down performance of a stave prototype).

The results from measurements performed have been included in the finite element analysis models of the stave for thermal and mechanical simulations of the device. During the course of the work proposed here developments of these methods will enable the full exploitation of material technologies, and measurements will be made during development both before and after irradiation and as part of the QA processes. The effect of moisture in the environment will also be investigated in a systematic way. A crucial element of the current proposal is a study

of nominally similar materials from a range of suppliers and measurement of a statistically significant number of different batches to understand likely QA issues. Some new materials will also be investigated.

Whilst there is an existing design for the stave, there remains the strong possibility of improvement, and certainly for optimisation, which will improve the physics capability of the tracker by reducing its total radiation lengths. A portion of the programme proposed here is orientated towards that goal.

The present measurement apparatus, developed at three sites in the UK, has reached a degree of maturity which will allow the detailed study of adhesives to progress with reasonable speed during the proposed grant period. Systematic measurement of the thermal performance of the many glue types used in the present STAVE09 design will be concluded and the proposed co-cured interfaces have to be fully understood, especially after the expected radiation doses. The scope for further improvements in the performance of interface materials will be investigated; for example the use of carbon nano-tubes as an additive to Araldite. The adhesives used in the carbon fibre skins will also be investigated via the study of samples from the pre-preg manufacturers, as process which will also be developed for use in mass-production QA. The structural, as well as thermal, performance of these items will be examined.

The measurement of Young's modulus at strains as low as 0.003, (which for typical samples corresponds to the measurement of extensions of less than 5 μm), has been demonstrated on aluminium standards and used on the carbon fibre samples for STAVE09. The value of Young's Modulus at these low strains is important as this is the range of strain of interest to the tracker construction. If Young's Modulus is not the same as measured at higher (more commonly studied) strains then the stave/tracker system may be over designed and have more mass than required. The developments to measure Young's Modulus at these low strains will be fully exploited during the period of the new grant with a full program of measurements of Young's Modulus on the stave's component materials and possible alternative low mass foams.

The measurement of the CTE with the Perkin Elmer dynamic mechanical analyser at RAL has only just started and will be continued throughout the present grant period. These measurements of dimensional stability will be extended with the use of thermogravimetric analysis to study moisture absorption of the materials of interest. During the programme proposed here, changes in the dimensional stability and strength (CTE and E) as a function of moisture uptake and loss will be studied.

All the measurements will be performed before and after the expected radiation dose of the ATLAS tracker upgrade.

2.3.5.3 Cooling System

The ATLAS Tracker Upgrade stave construction concept relies on embedded cooling pipes constrained within a composite structure for removal of heat load from the Silicon modules (Fig. 2.5). Following a study of pressure handling characteristics, corrosion resistance, mass, mechanical strength and stiffness; 3.175 mm $\times$ 0.22 mm wall 316 L Stainless Steel seamless tube has been adopted as the suitable baseline material for the inner detector cooling tubes. Successful management of the cooling services requires the tubes to be bent, joined and capable of accommodating a range of working pressures. Each of these areas has been studied as part of the existing ATLAS tracker upgrade R&D programme.

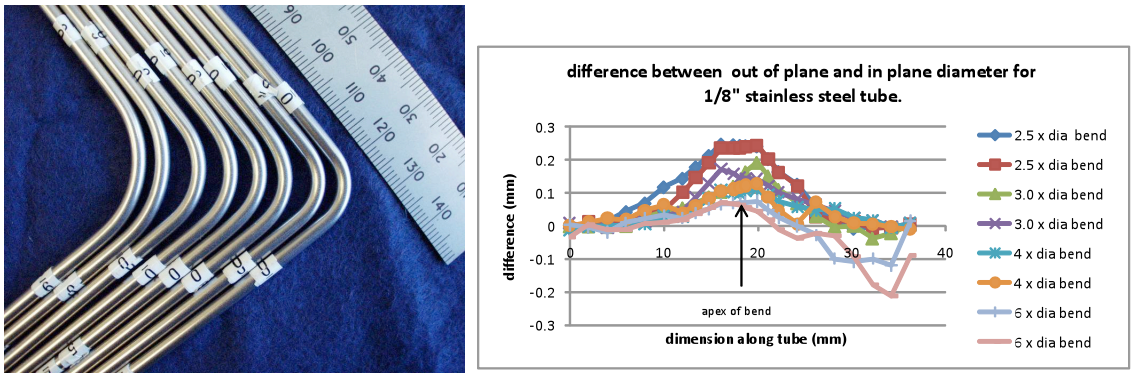

Figure 2.17: Photograph showing bends of varying radii and (right) results Connectors have been a perennial problem in the existing ATLAS SCT, so it is planned to make joints internal to the tracker envelope using in-situ orbital welding.

One of the main tasks within the proposed programme will be to supply the cooling pipes needed for stave pre-series production. These pipes will need to be bent to the specified geometry and supplied with fittings to connect the pipe to an external cooling plant. A stock of pipes with realistic dimensions was procured during the current ATLAS tracker upgrade R&D programme, however it is foreseen that another stage of pipe procurement will be required once the final stave design has been agreed internationally.

A pressure test facility has been built, capable of testing pipes and welds to 150 bar, which is sufficient for the baseline CO_2 cooling system. The pipe diameters can be measured before, after and during pressurisation. The baseline pipe, butt welded joints and bent corners have all been shown to withstand the relevant pressures. Pipes have been bent at a range of radii, down to $2\times$ pipe diameter, with excellent roundness. As an illustration Fig. 2.17 shows some results of measuring pipe bends of different radii.

Despite the considerable success with the baseline pipe design, there are many issues which remain to be pursued during the programme proposed here. The cooling system in the existing ATLAS tracker caused many problems and there is considerable determination that this situation not be repeated. In addition, the possibility of using titanium pipes, which convey the particular advantage of reduced CTE and concomitant lowering of stress in the stave core, has not been fully explored. A programme of titanium pipe studies is planned.

Use of welding with thin wall pipes in this kind of application to replace connectors is a novel approach. Several aspects of this require long-term study. It is planned to investigate welding dissimilar materials; pipes of different diameter; pipes of different wall thicknesses; and welding to commercial fittings. It is also possible to consider the use of pre-placed filler rings for a non-autogenous process, where the rings are composed from alloy powders which allow the fusion of two normally dissimilar materials. Each of these offers the possibility of simplification or improvement in the final assembly process and needs to be carefully understood. A long term programme of testing of the results of welding and bending is needed to confirm that the procedures are reliable. Figure 2.18 shows a remarkable weld between two 220 μm wall pipe, welded using parameters developed as part of the current tracker upgrade R&D programme.

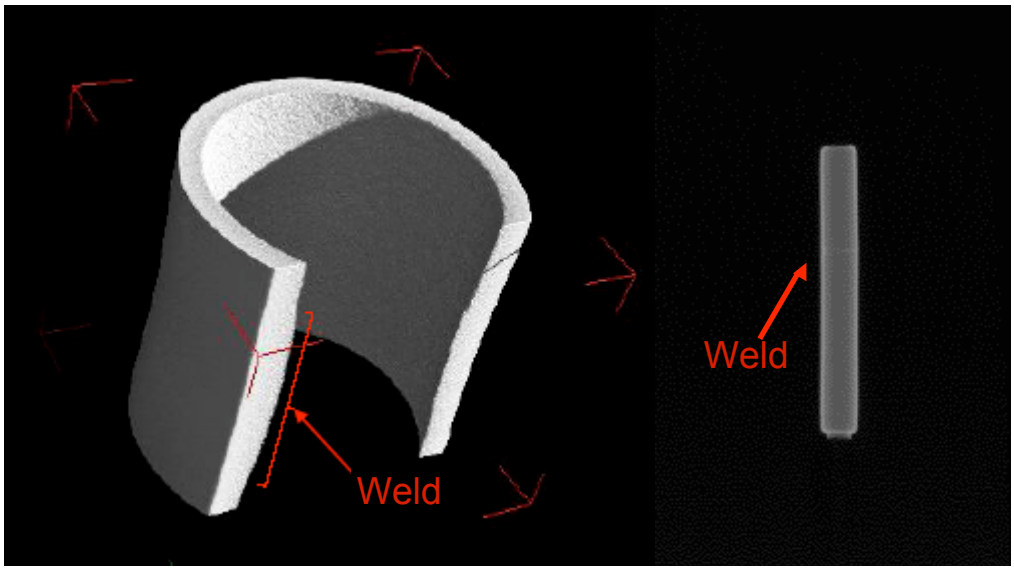

Figure 2.18: Micro-focus C.T x-ray image (left) & conventional x-ray image (right) of autogenous butt weld on 316 L Stainless Steel, illustrating the excellence of the weld.

In addition it is critical to understand the impact of welding in the vicinity of sensitive electronics; a test programme has been instigated to investigate this, but the programme will extend past the end of the current upgrade programme.

Beyond investigation of paths to improve the baseline design, it is essential that QA procedures be developed, and mass production tooling fabricated. Batches of pipe will need qualification and aggressive destructive procedures may be required on samples from each batch. Pipe bending tooling will need to be adapted to mass manufacture and several sets produced. In particular the end-of-stave bends are quite complex and extended studies of reliability are essential.

Prototype work on modules and staves will require a few low-capacity (< 1 kW) evaporative cooling plants. In case of CO_2 cooling these are likely to be simple blow-off systems. It is intended to build these plants based on concepts developed by our international collaborators and from past projects (LHCb).

Critical parameters for the ultimate thermal performance of a stave design are the heat transfer coefficient from the coolant pipe wall to the coolant and the pressure drop over the length of the cooling pipe. We are involved in a programme to study these for the two candidate coolants, CO_2 and C_3F_8. These studies will need to be continued on final geometry cooling pipes.

2.3.5.4 Stave Assembly

By the start of the R&D phase covered by this document the UK will have built and characterized at least one full-scale thermo-mechanical stave prototype, which will also confirm the validity of our mechanical models of these designs.

While the basic design of a stave is established, further aspects of the stave design remain to be optimised. The design of the stave core and associated components will be reviewed in the light of the data from the thermo-mechanical prototype and the stave design modified as

needed to optimise performance whilst minimising radiation length. It is expected that further optimisation is possible of the proposed mounting of the SMC on a sideways extension of the stave plank ('side-mounted SMC'). This saves space in the critical service gap between the active areas of the barrel and the endcap strip detector. The UK locking strategy, where the stave will be firmly connected to the overall support structure along one edge and the stave ends, needs further development and test. Special emphasis has been put in this design on a good definition of the relation between the local (stave) and global coordinate frame. Experience with the existing SCT strongly motivates good service management, which also has important implications on the layout of the on-stave services and the EoS region, and work will continue on this as the international project evolves.

The assembly of a stave consists of the construction of the 'stave plank', and the subsequent mounting of the modules. The stave plank is a carbon fibre sandwich structure with embedded cooling pipes, locking mechanism and electrical tapes. Vacuum jigs will be used to locate the skins whilst building up the core. Use of co-curing techniques, laminating the electrical tape together with the CF pre-preg layers look promising and more extensive co-curing may be possible. Mass production techniques must be developed to bond the Pocofoam blocks needed for good thermal coupling around the cooling pipe. This bond is critical for both the thermal performance and for the mechanical strength of the stave design. For the passive core material different materials are still being considered: CF honeycomb; graphite foam; or custom carbon-fibre reinforcement. The impact of this choice on mass production needs to be understood.

After construction of the core, the elements needed for the connection to the overall support structure need to be attached. The required positional precision for this, as well as for all the other stave assembly steps, still needs to be fully understood. This is a complicated issue bound up in assumptions about software track reconstruction techniques.

Prior to attaching modules, the construction of the plank will be concluded by mechanical, thermal and electrical QA procedures, all of which still need to be developed. Significant infrastructure suited to aspects of the QA procedure have been built, including cooling systems and deformation measuring methods, but adaptation to mass QA is needed.

As described earlier, at the end of the programme, the UK expects to have two stave production clusters operational. In consequence, the main emphasis in the coming three years in the work related to strip staves will be the development of the production and QA methods (procedures, tools, training) for this assembly. This will require standardisation of assembly methods within the UK and internationally. It is expected that the UK will lead these developments therefore and demonstrate their applicability within the international effort. The required tooling will be fabricated.

2.3.5.5 Module Mounting

The details of the process of mounting silicon modules onto carbon-fibre structures depend on the module's size and handling restrictions, the accuracies required during mounting, and the adhesive used for final fixation. Much experience has been gained in the UK in all of these aspects, most recently during the production of the ATLAS SCT, however there are many differences of detail in the upgrade tracker assembly.

During the current R&D programme, it is planned to develop a mounting process using a program of trial assembly using silicon sized glass sheet. These dummy modules will be

mounted onto a carbon fibre stave. The intent is to simulate the careful handling and manipulation process required with delicate silicon components. Extreme care is required from the first operation of extracting the module from its delivery box up to the final gluing onto the carbon fibre structure. Using experience gained in the past and reusing existing silicon cleanroom facilities has allowed this development work to progress steadily and the first milestone of providing full manipulation for dummy silicon module mounting is about to be reached.

The size of the test stave is limited to suit the measuring equipment available in the assembly facility and one element of the new R&D programme will be to progress to a full size model for assembly trials. The frame used for handling the stave during assembly will also need to be adapted, both to the final designs where the details of component location are crucial, and also to mass production. This frame is likely to be part of the suit of parts described in Section 2.3.5.7. For the existing test program the dummy stave frame is mounted onto a stable base plate which is fixed to the optical inspection equipment. For mass production this arrangement will need to be reviewed.

The dummy silicon modules are picked up in a vacuum carrier and transferred to and located on the base plate. The vacuum fixture will need evolution during the programme proposed, particularly when the hybrids switch to the 0.13 μm ASIC suite. The vacuum carrier is mounted on a four-axis manipulation stage. With the silicon in place over the stave optical inspection equipment is used to locate the module to its specified location on the stave. This process depends on the fiducials used to define the location of the module and the optimal way to achieve the necessary precision will be developed in collaboration with module designers.

Towards the end of the programme period, when the assembly process is fully understood, two sites will be equipped to assemble modules to staves. Assembly of dummy staves followed by full assembly of pre-production staves will be undertaken, with the goal of achieving assembly of two staves per assembly site with at least some of the process done at the required mass-production pace.

2.3.5.6 Quality Control

The bulk of the content of this section is covered elsewhere, but to emphasise the importance attached to quality control, the diverse parts are drawn together here. Whilst the technical details of the QA are very much part of the individual programs it is vital that a consistent approach to quality is maintained.

At the end of the programme described, it is expected that ATLAS will have an upgrade TDR, which will enable production to commence. The schedule shows that the production phase of the project is quite tight. To achieving such a schedule, high yields, and a smooth system of testing and rework will be required. To be ready for this in three years from now, preparations must start very soon.

All aspects of the production program will involve industrial partners. They will deliver components which past experience has shown will need reception QA tests. These tests will be dependent on the outcome of the individual work programmes described in this proposal. However, it is important that at all stages the global context of each putative reception QA test is considered. This is necessary to ensure that tests needed are undertaken without duplication at various stages of assembly.

The module program will need a great deal of qualification and multiple stages of testing. For this some engineering effort will be needed throughout the pre-series stages. With appro-

priate care and attention directed at these processes it will be possible to avoid the need to iterate for mass production. This work will be essential for the pre-series work, however this will not naturally provide data collection systems, numbering and cataloguing control, or any interface to the centralised data bases. This will need developed explicitly towards the end of the programme proposed.

The mechanical aspects of the upgrade programme will generally lag somewhat behind the module program. Nevertheless full QA/QC provision for the staves must be developed in the next three years. Cooling provision, metrology and part qualification must all be established. A comprehensive qualification program for all of the completed staves must be developed. This will require internationally agreed methodologies.

All of the services integral with a stave will require testing and again the techniques and equipment require development. Whilst this work is largely geometry independent, it is unlikely that the engineering work will start before the final international design and development programs are complete. However, development and refinement of these techniques must commence much earlier, in parallel with the final stave development. Funding is requested for the development and prototyping of these techniques within this program.

Over the last three years a facility has been developed to qualify pipes and pipe fittings, once the R&D phase of this project is complete this facility will be adapted to testing and approval of staves. The tape and opto programs are starting from a lower base in terms of QA/QC, but in both cases some infrastructure already exists.

2.3.5.7 Test/Shipping Container

During construction of the existing SCT the UK supplied a variety of precision boxes to ATLAS, ranging from sophisticated boxes used in beam tests to large numbers of simpler containers for global transport of modules and sub-assemblies. For the upgrade, the final deliverable is a relatively large stave, of high value and for which a system must be developed for shipping and storage as well as support during a variety of tests required for QA and potentially some assembly steps. Far from mundane, the containers required will be multi-purpose and become quite sophisticated. Experience has shown that these containers must be designed with care from the outset and UK experience will be essential in ensuring that their design accomplishes all the required functions.

The containers will be designed to act as handling tools during the module attachment process, and probably during parts of the stave-core manufacture. Adding and removing additional components will accommodate various stages of manufacture and QA: stave-core assembly; module attachment; wire-bonding; testing both warm and cold, with and without optical interfaces; storage; and safe shipping.

Finally, the container will be used during the insertion into the barrel support structure during the final barrel integration at CERN. In this way staves never need to be handled on their own, with the containers providing support and protection continuously until this role is taken over by the tracker support structure.

The anticipated design will have detachable elements, some of which will be precision parts and expensive; these will be made in small quantities and reused. However the basic handling and storage functions will be accomplished with relatively cheap structures and will be manufactured in large numbers as experience has shown that it is likely that one storage case per discrete element is likely to be required.

Development of the containers will proceed in concert with the development of the mass production methods and QA procedures they serve. At the end of the R&D programme proposed here a fully specified and tested design for the containers will exist and several prototypes have been manufactured.

In addition to development of the production systems, various prototype containers performing similar functions to the final versions will be required for handling and shipping during the development of the mass production process.

2.3.5.8 Integration of Staves into Inner Detector

In integrating the existing ATLAS SCT, significant problems were encountered supporting and positioning the various detectors and their services. The upgrade layout has been developed with extensive consideration given to making the integration simpler. However significant work remains to be undertaken. Given the UKs heavy involvement with the services and integration in ATLAS it is crucial that this experience be applied to the new detector. Over the last three years the UK has engaged in this program peripherally, but has developed a proposal for service integration inside the ID, based on the service module concept (Fig. 2.6). To avoid repetition of past problems, it is important that the UK service and integration proposal for the region from the end of the staves to the end of the inner detector be actively pursued as part of the programme proposed here. In collaboration with CERN and Annecy this concept will be developed, prototyped and a program of robustness testing undertaken.

The final design and manufacture of services will be outside the time scale of this program; however the next three years are the time when the future practicality and reliability of the system will be determined and it is extremely important that the UK experience be brought to bear on the design phase.

A small UK engineering effort in this direction will secure a large influence on the whole of the inner detector integration including the support structures, ID support and endcap insertion. This will enable the UK to make optimal decisions about the end of stave region, whilst ensuring there is sufficient engineering input into the layouts. The goal is to have a complete mechanical prototype of the service region for one quadrant of one end, by April 2013, as part of an international collaborative effort.

It is not planned to work on services outside of the ID, however it is important that the UK engineers work closely with the other areas of the inner detect community to ensure that all the services work efficiently together and the space is appropriately utilised.

2.3.6 Pre-production System Tests

The UK played a significant and leading role in the prototyping, installation and commissioning of the current SCT, and UK members continue to provide leadership and coordination roles in its present operation. In particular, the UK developed internationally recognised expertise in a broad range of SCT subsystems and work packages from sensor and module prototyping/production to opto-electronics, DAQ and system tests.

The system tests program for the upgrade will draw heavily on these UK strengths as it will provide a test-bed for the stave pre-production prototypes and for the integration of work packages including mechanics, powering, modules, DAQ and detector control systems.

It is envisaged that 2010/2011 will cover the development of the system test infrastructure for the build-up and test of the STAVE09 short strip prototype stave using ABCN25 chips. That infrastructure will include evaporative cooling, mechanical jigs and hardware for DAQ and control systems, and will be established at CERN under UK leadership as a central facility for the international upgrade community. 2011/2012 will be used to consolidate the system test setup with thorough tests of the stave and the various associated support systems, including shielding and grounding. 2012/2013 will build on this experience to test the pre-production staves using the new ABCN13 chips, and is likely to involve an expansion of the test facility to provide for a multi-stave test setup.

The systems test facility will evolve as needed according to progress and developments within the various work packages. A strong UK involvement at this central test facility will ensure the UK remains at the heart of developments and decision making within the international upgrade program.

2.4 Silicon Pixels

Pixels are an essential ingredient for the upgraded tracking detector. They provide both pattern recognition and also the precision points close to the interaction that determine the resolution of primary and secondary vertices.

Milestones, along with cost and effort tables for the Pixel programme, are given in Section 5.3 (WP5). Here we describe the work programme in detail.

The pixel system has many requirements that make it challenging to develop and construct but the UK has expertise in many of the areas and would make a significant impact in the pixel system. These include world leading expertise in ultra-radiation hard technologies for sensors, including planar and 3D technologies; readout electronics and modules, extensive experience designing and building minimum mass modules and support structures; interconnect expertise based on MCMd work, and synergy with UK built LHCb Vertex Locator (VeLo) developments. The UK also has long standing experience with detector layout, construction and commissioning, both from the barrel and forward SCT, and many years of experience of module construction for many projects.

The forward pixel system has been identified by ATLAS as a critical area requiring experienced groups to take leadership. This provides the UK with an opportunity to extend its experience of constructing silicon strip based trackers into the area of pixels which, as described in the introduction, is considered a strategic necessity for the UK particle physics community.

The aims of the work package described in this section are: to undertake the R&D necessary to develop and qualify the technologies required for the forward pixel disks; to develop the global and module design for the forward pixel detector system; and to place the UK in a position to take a leading role in the construction, commissioning and running of the upgrade forward pixel system.

In the short term the sensor technology will be used in the Insertable B-Layer project (IBL), which provides an essential test-bed for the technology as well as improving the physics performance of the ATLAS pixel detector. The 3D R&D is also critical to the AFP WorkPackage described in Section 2.2 and Section 5.3, WP1.

2.4.1 Sensor Development

2.4.1.1 Sensor requirements

The sensors for the pixel system will have to operate under extreme irradiation. The expected dose varies with radius up to $\sim 2 \times 10^{16} n_{eq}\,\mathrm{cm}^{-2}$ over five years for the innermost layer at 3.7 cm. This is over a factor of ten higher than that anticipated for short strips (Section 2.3), which is a significant challenge demanding further dedicated R&D. The other critical requirement for a pixel system is that, because it is the detector which measures the event vertices, it is particularly important that it be low mass. This requires the sensors to have the largest possible active area compatible with the module geometry in order to minimise the material added by sensor overlap.

Two technologies are considered to be the main contenders by ATLAS, 3D and planar, and the UK has international leadership in both areas. The work proposed here builds on generic work funded by PPRP PRD awards to institutes in the current team. The technology must be developed to the specific requirements of ATLAS. Both options are extensively studied within the RD50 collaboration and by R&D consortia within ATLAS of 14 and 17 institutes for 3D and planar respectively. UK sensor results are dominating discussions in both planar and 3D work. It is the intention within ATLAS that the R&D culminate in sensors of both technologies, tailored to the ATLAS geometry and with pixel readout, being studied in a CERN testbeam, ideally after irradiation to the required doses. This will allow a realistic comparison to be made and allow ATLAS to determine where these technologies are required in the upgrade. As leaders in the development of both technologies, UK expertise will be key to such comparative studies

2.4.1.2 Planar sensor technology

Planar detectors are fabricated in either n-type or p-type doped bulk silicon, with n-type structures in both cases, ‘n-in-n’ and ‘n-in-p’. These sensors are based on established planar technology and suppliers (Hamamatsu, Micron) who have provided sensors for existing large area silicon detectors such as ATLAS SCT and LHCb VeLo. This is therefore a ‘conservative’ technology choice, offering low cost per unit area and good security of supply. However, the key to using this technology in a high radiation environment is, however, the ability to operate with significantly higher voltages (1000 V) than in previous detectors, which is necessary to achieve adequate signal to noise ratio, S/N. A few pixel detectors compatible with the ATLAS FE-I3 ASIC have been included in a $6''$ mask set processed as part of a PRD funded programme on p-type wafers at Liverpool with Micron Semiconductor (UK) Ltd. The first measurements of the reverse current characteristics of the pixel devices from Micron exhibit very encouraging results, sustaining bias voltages up to 1100 V (Fig. 2.19), which hopefully points the way to establishing Micron as a contender for future pixel orders for particle physics and other disciplines.

Figure 2.20 shows the CC(V) properties of 140, 200 and 300 μm thick n-side readout sensors after doses of 1.5 and $2 \times 10^{16} n_{eq}\,\mathrm{cm}^{-2}$ (the highest anticipated dose for the innermost layer) [27]. These results have been obtained in earlier research led by the UK, with 1 cm-long strip detectors made by Micron Semiconductor (UK) Ltd. The strips are AC coupled to 40 MHz readout analogue electronics. These results demonstrate that planar silicon detectors can provide adequate signals to be a possible option for the future pixel system at the sLHC.

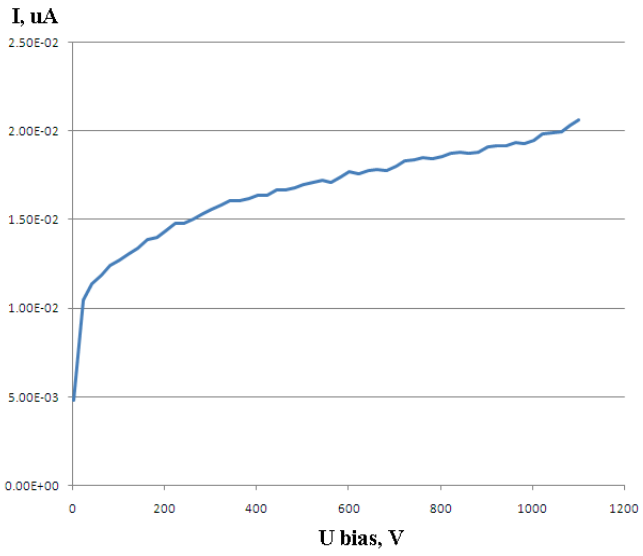

Figure 2.19: Reverse current as a function of the bias voltage for a planar p-type pixel detector produced with Micron Semiconductor (UK) Ltd to be compatible with the FE-I3 readout ASIC.

The measured collected charge is significantly higher than predicted by current modelling of irradiation effects in silicon detectors. This effect is now well established through several measurements performed with different readout systems and with detectors made by at least two manufacturers. One interpretation is that it could be due to a charge multiplication component of the signal and this is now a key topic of research internationally. One further topic of research being proposed here is to investigate the possibility of achieving the signal enhancement at lower bias voltages. Additionally, the implementation of thin edge technology, to achieve the required voltage with minimal dead region at the device boundary, is under investigation.

Due to the much smaller area of the individual sensitive diodes, the capacitance-related electronics noise per channel is lower and pixel detectors are expected to exhibit a lower noise, enhancing the signal over noise ratio. Therefore enhanced performance compared to strips is expected after these extreme doses. An added advantage is that the signal is further enhanced when using thinner detectors for a given voltage, reducing the overall required material budget.

2.4.1.3 3D technology

3D technology is a relatively recent development in sensor technology based on traditional planar technology coupled with micro-machining. The micro-machining is employed to create electrodes inside the bulk of the Si sensor (Fig. 2.21, [28]). This results in charge collection distances as short as $\sim 50\ \mu$m while the charge generated by the traversing particle is determined by the substrate thickness which can be up to 300 μm. This decoupling of the charge collection and generation processes provides the improved radiation tolerance. The short carrier drift distance, and higher average electric field for a given voltage, can improve the response time by a factor of ten. Moreover the arrival time spread for all the charge can be lower since the particle path from a minimum ionizing particle is parallel to the electrodes. 3D geometry has

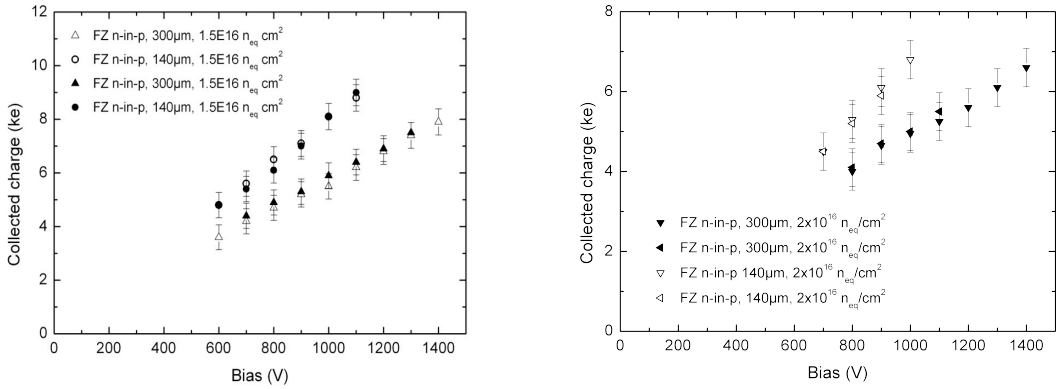

Figure 2.20: Signals with minimum ionising particle for n-in-n and n-in-p sensors with various thicknesses (from 140 μm to 300 μm) after 1.5 and $2 \times 10^{16} n_{eq}$ cm^{-2} as a function of voltage.

shown great flexibility in the use of readout electronics, and since both electrodes are accessible from the front and back side of the wafer, it is possible to process the readout electrodes to be compatible with both pixel and micro-strip readout chips for both input polarities and on either side of the wafer.

Low depletion voltages and short drift distances result in the 3D structures being able to achieve good S/N to noise at relatively low voltages at sLHC fluences, although the noise is greater relative to planar devices due the increased capacitance of the sensing structures. The lower power impacts on the system aspects. Devices have been characterised after irradiation to sLHC fluences (Fig. 2.22, [29]) and laser diode and mip results show that large signals can be achieved. Several different layouts have been developed, led by the UK in collaboration with small specialist foundries (Sintef, CNM and FBK), and their response has been studied in various testbeams. The columnar structure of the 3D detectors results in dead regions where the implants are located. Measurements have shown that this can result in hit efficiencies $\sim 95\%$ across a detector for perpendicular tracks. This is mitigated by tilting the detector by 15°, giving hit efficiencies $\sim 99\%$. Optimisation of the final detector design will allow for this constraint. Active edges are relatively easily implemented in 3D sensors. Dead edge regions as small as 4 and 100 micron are possible for full-3D and 3D-DDTC technologies respectively. This helps to improve the tiling leading to reduced detector mass.

To apply 3D technology to the ATLAS pixel system requires the fabrication of large sensors, with high yield, by a supplier capable of producing sufficient sensors for the detector in an appropriate time scale. To date, sensors for the current ATLAS pixel chip have been manufactured which are around 0.5 cm^2. Existing devices have been made by SINTEF(Oslo), FBK (Trento) and CNM (Barcelona) and an investigation of the capability of manufacturers to produce large sensors with sufficient yield will form part of this project.

Development of 3D for ATLAS is essential as the applicability of planar technology to the innermost layers of the ATLAS pixel upgrade, where very high voltage operation will be required, is not sufficiently established.

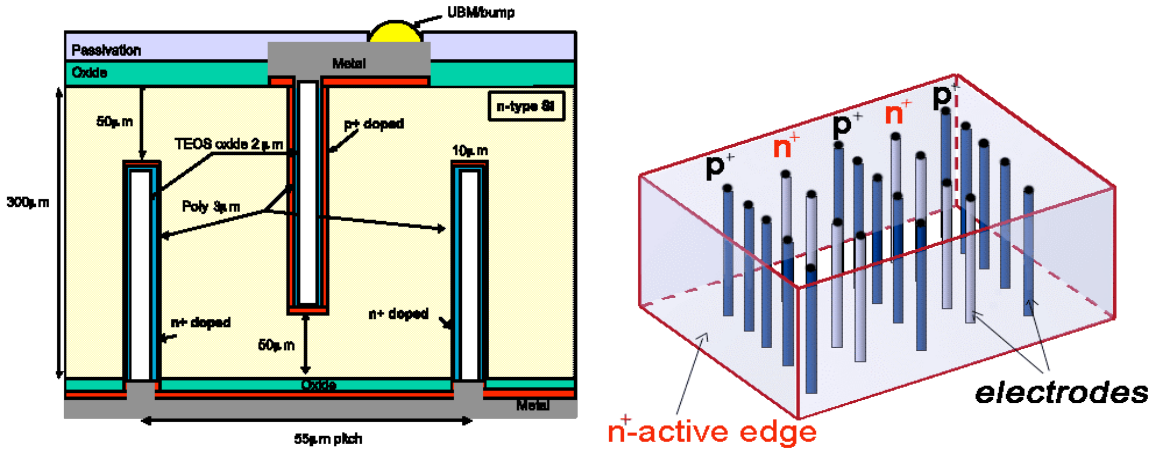

Figure 2.21: Double side process (left) and full with active edge 3DC (right) sensors. These layouts will be both studied within this R&D.

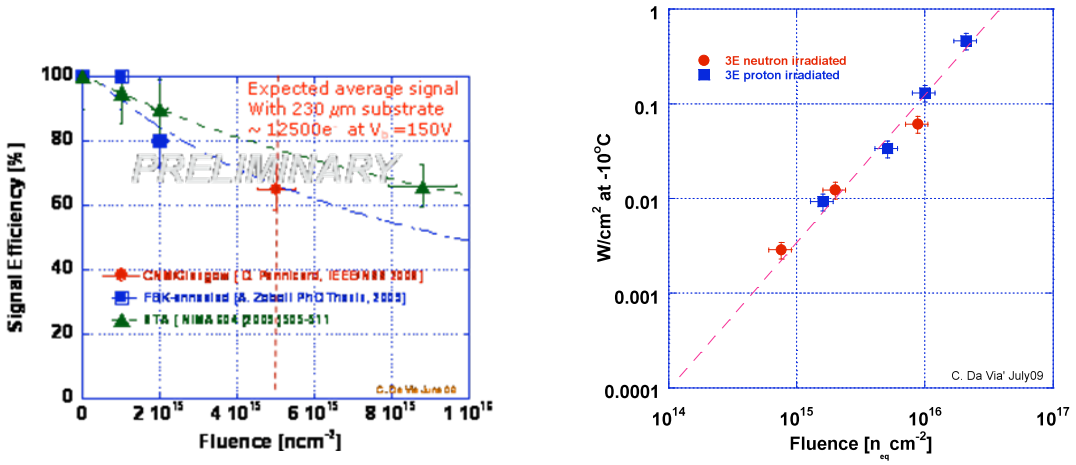

Figure 2.22: Left: Signal efficiency vs Fluence for 3D-DDTC and full-3D technologies showing good compatibility between the two 3D designs and measurements using different experimental methods. The green triangles are the efficiencies of full-3D sensors and were measured using an infra-red laser diode. The red dot and blue squares are for measurements with MIPs for CNM/Glasgow and FBK-Trento 3D devices respectively. Right: Power dissipation measured at $-10°$C of neutron and proton irradiated full-3D sensors up to $2 \times 10^{16} n$ cm^{-2}.

2.4.1.4 ATLAS Sensor Development

For both planar and 3D sensors in the context of ATLAS it will be necessary to design sensors compatible with the FE-I4 readout chip, which promises sufficient radiation hardness to be used in technology comparisons in test beams. For planar this requires focus on HV design and active edges. For 3D this means development of large sensors and an understanding of yields and costs from commercial suppliers.

Sensors must be procured and bump-bonded to FE-14 read-out chips and tested both before and after irradiation. A significant number of sensors must be irradiated.

The CERN testbeam programme anticipates un-irradiated and irradiated modules to be studied as a function of angle (and ideally magnetic field) with evaluation of S/N, geometrical acceptance, efficiency and spatial resolution.

For the UK programme it is also important that experience is built up at several sites in the UK in handling and quality control of pixel sensors, modules and systems. This will require a significant number of sensors be procured, tested, and used.

The work will be carried out within the global ATLAS upgrade project and form input to the TDR that will define the applicability of planar and 3D technologies for the pixel upgrade programme.

The sensors developed as part of this programme will also be used to develop the novel interconnects described in Section 2.4.2.

2.4.1.5 The Insertable B-layer

A new pixel layer, the IBL will added to the ATLAS detector, possibly coincident with the Phase-I upgrade, as the existing detector will eventually become inoperable due to radiation damage. It will also improve the rz impact resolution by a factor two, improving the b-tagging. Both planar and 3D technologies are being considered by ATLAS for IBL. This provides a low cost opportunity to demonstrate sensors from the UK development programme in a real detector environment and establish the UK within the ATLAS pixel community, which will be critical in playing a leading role within the Phase-II upgrade.

As described in Section 2.4.1.1 it is important that several UK groups gain experience in the handling, testing and QA of pixel sensors. Ideally the sensors being developed under UK leadership should be used as part of a real detector, providing final confirmation that the procedures employed are valid.

The IBL has stringent radiation tolerance requirements ($2 \times 10^{15} n_{eq}$ cm-2), although not at the level of the final phase of the sLHC. Here it is proposed to contribute to the design and the acquisition of $\sim 30\%$ of the sensors for the IBL module production. This is a sufficient number to enable UK institutes to build the requisite experience in pixel sensor technology. Both the possible planar n-in-n and n-in-p, or the 3D solution could be chosen by the IBL ATLAS community. The cost will be modest and will provide an essential real-life test-bed for the technologies foreseen for sLHC.

2.4.2 Connectivity

The UK plans to develop modules and support mechanics for the forward region of the pixel system. Here there has been much less effort than devoted to the sLHC pixel barrel design,

leading to a concept based on the square barrel modules of $250 \times 50\mu$m pixels arranged in an approximately pointing geometry with very large overlaps leading to significant material and a rather poor match to the pattern recognition requirements of the region. In the UK it is proposed to develop a more appropriate solution to the region. This effort will be aided by exploiting the similarities with the LHCb VeLo pixel upgrade. There is a high degree of overlap in the R&D required on interconnectivity, radiation tolerance, and geometric coverage. Given the lead role of the UK in the LHCb VeLo and the likely earlier timescale of its smaller scale upgrade, this provides a good opportunity to develop a coherent effort, with advantages for the UK in avoiding duplication of investments in key aspects such as read-out ASICs and interconnects better suited to symmetric geometry pixels and assembly into low-mass disc structures.

One of the crucial aspects of the R&D towards improved pixel sensors is the reduction of the dead area for a system of sensors tiled on a given plane. The need to access the wire-bonding pads on the ASICs demands that the contacts be outside the area of the sensor, which prevents efficient tessellation on the contact side. Moreover, the active part of the pixel sensors and the overlapping electronics must have the same dimensions. This requirement makes it particularly difficult to efficiently cover circular surfaces with minimum overlap of square or rectangular sensors (to match the shape of the readout ASIC chips).

A different approach is to produce pixel sensors with a more appropriate geometry to cover wedge surfaces (and therefore wheels). To achieve this, the shape of the sensor needs to be decoupled from that of the electronics by mean of a dedicated routing network connecting pixels not positioned directly underneath the readout channel. This result can be achieved with multi-metal layers on sensors or with interposed routing layers such as MCMd routing. UK institutes have already studied both options. Limited application of such routing is already implemented in current ATLAS pixels to gang together pixels between ASIC boundaries. These techniques will be developed for the ATLAS upgrade geometry as part of this programme.

Irrespective of connectivity method, a key requirement for flip-chip bonding is to be able to wire-bond to the ASICs from the opposite side to the bump bond contacts. This important feature can be achieved with through silicon vias (TSV), or low resistivity column wires etched to bring the contact to a metal pad deposited on the back-side of the ASIC. Thinned electronics is needed for producing TSV with a relatively small column diameter (10 μm). This requirement naturally reduces the ASIC mass, which is desirable for pixel sensor systems in all vertex detector applications. These techniques will be developed to be suitable for ATLAS, through collaboration with CERN and participants in the international Medipix/VIPS/AIDA consortia.

2.4.3 Forward Pixel Layout & Mechanics

Development of pixel sensors will be performed in parallel with that of the corresponding forward module and layout concepts. There is an intimate connection between the two activities; this is due to the fact that the sensor geometry dictates the amount of overlap required for full coverage which in turn impacts on the engineering design of the mechanical supports, services interconnect and cooling.

Layout studies will be undertaken to optimise the geometry of the pixels, (e.g. square versus rectangular), and of the layout of the pixel modules on the forward disks to achieve the best possible resolution on the primary vertices, with the minimum possible mass. These studies will define the design of and the requirements for the single-chip pixel prototype that will be

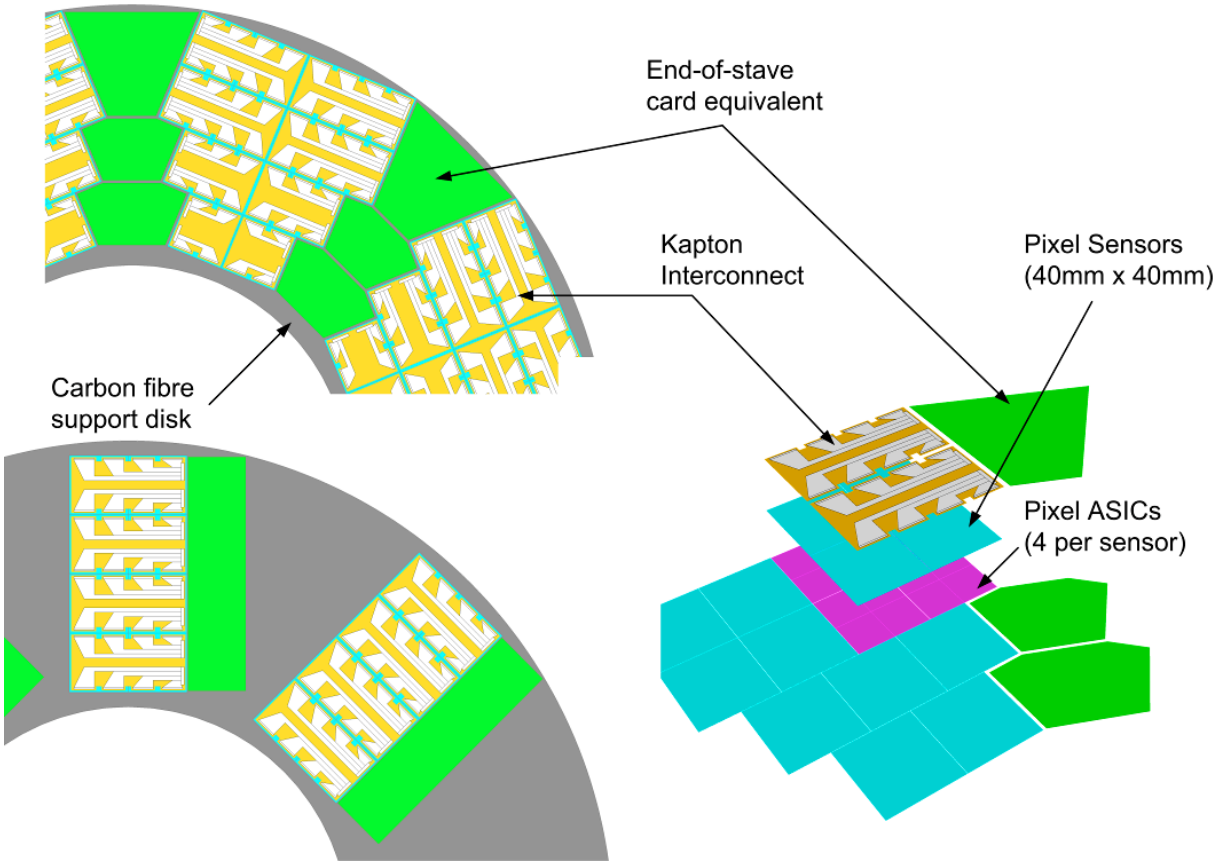

Figure 2.23: Generic Forward Pixel Disc Design

built using advanced connectivity techniques (see Section 2.4.4).

It can be anticipated that the minimum overlap of adjacent modules can be obtained with square pixel electronics or with Pixel with Displaced Readout (PDR).Square (or hexagonal) pixels could have an individual channel size of order $100 \times 100 \mu m^2$, compatible with the area required for the presently proposed FE-I4 rectangular pixels ($50 \times 250 \mu m^2$). The PDR solution could also lead to pixels with different size at different radii, to keep the occupancy $\ell^+\ell^- 1\%$ at all radial locations.

In either scenario one could envisage the possibility of having arrays of non-overlapping sensors with minimal dead spaces and an initial conceptual forward pixel generic mechanical design, based on this topology, is shown in Fig. 2.23. Sensors are mounted in groups on both sides of a disk with the signals from each group being bussed on copper/Kapton circuits to peripheral boards whose functionality replicates that of the End of Stave Card in the barrel pixel detector.

To implement either of these designs, technologies emerging from the connectivity R&D will be crucial (Section 2.4.2). It is anticipated that both approaches will share common solutions for low mass supports, flex hybrids and efficient cooling. The latter is a particularly crucial feature for successful operation of severely irradiated sensors.

Proof-of-principle designs, backed up with realistic prototypes, will be essential to demonstrate the significant improvements possible to the existing basic pixel concept for the forward

region and the performance improvements quantified as part of the simulation programme.

There will be significant overlap in the development of prototypes with the work being undertaken in the context of the short strip programme such as use of Ti pipes and carbon fibres support structures. Initially investigations will be made into applying the technologies and techniques developed for the barrel to the forward geometry, for example bending pipes into arcs or wiggles. Based on these studies and the layout studies, a thermo-mechanical model of part of a disk will be manufactured and populated with services and dummy modules that reproduce the thermal load of the pixel modules. The thermo-mechanical performance of the prototype disk will then be evaluated. As this proposal is directed towards preparing for production, a bare disk will be manufactured in collaboration with industry based on the tests on the thermo-mechanical prototype to establish the manufacturability of the disks.

2.4.4 Single-chip pixel prototype

The layout studies in Section 2.4.3 will define the optimum pixel geometry and module geometry. This will assume the use of the advanced connectivity technologies studied in Section 2.4.2 to achieve a low mass pixel system. The logical conclusion of these studies is production of a single-chip pixel prototype that will demonstrate that such a module for ATLAS forward pixels can be manufactured using the advanced connectivity techniques developed in collaboration with CERN and participants in the international Medipix/VIPS/AIDA consortial.

2.4.5 Status of Irradiation Facilities

The irradiation campaigns for the ATLAS tracker upgrade are performed at various irradiation facilities. The small sensors (1×1 cm^2) have been irradiated with neutrons at the neutron TRIGA reactor in Ljubljana by the ATLAS collaborator of the J. Stephan Institute, and in protons at the KEK-PS and CERN-PS. We have also benefited from parasitic access to the Karlsruhe 25 MeV cyclotron, through the RD50 collaboration. The dose rate in the proton facilities does not provide the high doses required for the irradiation of the pixel sensors in a short time (from 30 to 80 irradiation days in the CERN PS, which is often available for periods not longer than 15 days).

An alternative irradiation source has been identified in the Cyclotron at the University of Birmingham. It can provide an intense proton beam with monochromatic particle energies up to 38 MeV. It is planned to tune the beam energy to 26 MeV because at this energy the damage caused has been demonstrated to be equivalent to that caused by high energy protons. This facility will make it possible to achieve the maximum anticipated qualification dose for pixel sensors ($2.5 \times 10^{16} n_{eq}$ cm^{-2}) over a 4 cm^2 area within a day. This facility is being equipped with scanning table and irradiation cool-box for use in the near future, well in time for the pixel sensor studies.

The UK ATLAS Tracker Upgrade group is developing a low temperature scanning box for the CERN-PS T7 irradiation area that makes possible tests of the large area (10×10 cm^2) sensors and modules of the tracker. This facility is important for final qualification and quality monitoring of the assembled detector-electronics and will be used for testing the new sensors and ASICs of the pre-production phase of the Tracker Upgrade. It allows the upgraded SCT qualification dose to be reached uniformly over the 10×10 cm^2 surface, within two weeks

of exposure to the beam, provided that the CERN-PS accelerator delivers its maximum intensity. The use of high energy protons for final testing is required because of lower penetration with lower energy sources, which is an issue for modules in particular. However, access to the CERN-PS is at a premium and the dose rate low compared with cyclotrons (or reactors) making the latter ideal for radiation studies once calibration against CERN-PS protons has been confirmed.

Chapter 3

Trigger Upgrades

3.1 Trigger and Data Acquisition Overview

The ATLAS trigger system has three distinct levels: Level-1, Level-2, and the Event Filter. Each trigger level refines the decisions made at the previous level and, where necessary, applies additional selection criteria. The first level uses a limited amount of the total detector information to make a decision in about 2.1 μs, with an accept rate of up to about 75 kHz. The two higher levels access more detector information to select events at a final rate of up to 200 Hz with an event size of approximately 1.3 Mbyte. The Level-2 Trigger and Event Filter together form the High-Level Trigger (HLT).

Data for events passing the Level-1 Trigger selection are transferred off the detector and subsequently to the data acquisition system via point-to-point links. Here the event data are received and buffered for subsequent Level-2 and Event Filter processing prior to being recorded or discarded.

An overview of the ATLAS Trigger and Data Acquisition system is shown in Figure 3.1.

3.1.1 Level-1 Trigger

The Level-1 Trigger searches for high transverse-momentum muons, electrons, photons, jets, and τ-leptons decaying into hadrons, as well as large missing and total transverse energy. Its selection is based on information from a subset of detectors. High transverse-momentum muons are identified using trigger chambers in the barrel and end-cap regions of the spectrometer. Calorimeter selections are based on reduced-granularity information from all the calorimeters. Results from the Level-1 Muon and Calorimeter Triggers are processed by the Central Trigger Processor, which implements a trigger 'menu' made up of combinations of trigger selections. Pre-scaling of trigger menu items allows optimal use of the bandwidth as luminosity and background conditions change. In each event, the Level-1 Trigger also defines one or more Regions-of-Interest (RoIs), i.e. the geographical coordinates in η and ϕ of those regions within the detector where it has identified interesting features. The RoI data include information on the location, type of feature identified, and the criteria passed, e.g. a threshold. This information is subsequently used by the High-Level Trigger.

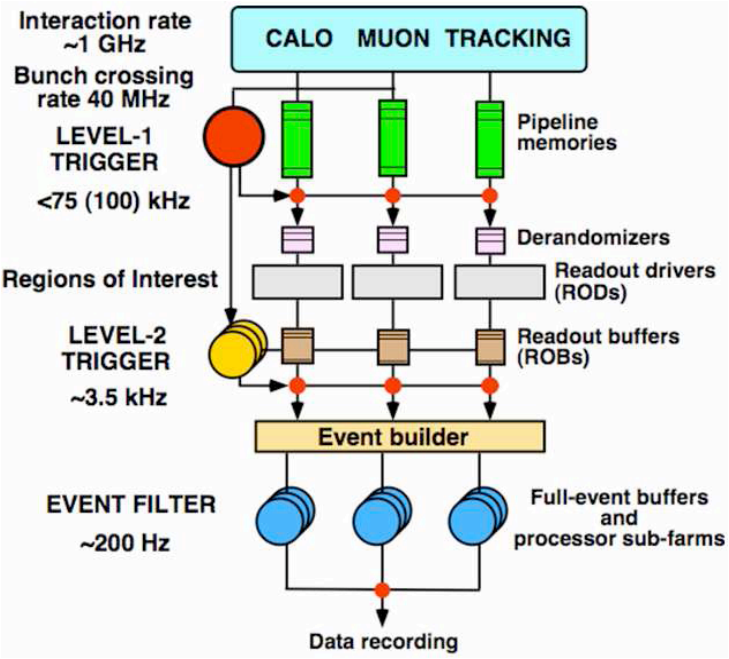

Figure 3.1: Overview of the ATLAS Trigger and Data Acquisition system.

3.1.2 High-Level Trigger

The Level-2 event selection is seeded by the RoI information provided by the Level-1 Trigger. Level-2 selections use all the available detector data within the RoI geographical areas. The Level-2 menus are designed to reduce the trigger rate to approximately 3.5 kHz, with an average event processing time of about 40 ms. The final stage of the event selection is carried out by the Event Filter, which reduces the event rate to roughly 200 Hz. Its selections are implemented using offline analysis procedures within an average event processing time of the order of four seconds.

The HLT algorithms use the full granularity and precision of calorimeter and muon chamber data, as well as the data from the Inner Detector, to refine the trigger selections. Better information on energy deposition improves the threshold cuts, while track reconstruction in the Inner Detector significantly enhances the particle identification, for example by distinguishing between electrons and photons. The event selection at both Level-1 and Level-2 primarily uses inclusive criteria, for example high-E_{T} objects above defined thresholds. One exception is the Level-2 selection of events containing the decay of a B-hadron, which requires the reconstruction of exclusive decays into particles with low momentum.

3.1.3 Trigger and Data Acquisition Issues at Phase-I

Few detector changes are planned for Phase-I running, with the exception of the Inner Detector Insertable B-Layer (IBL). The LHC machine and ATLAS front-end electronics timing will be unchanged, and the detector interface to TDAQ will be largely remain as at present. Trigger hardware R&D will focus on addition of topology to Level-1 processing, with the trigger latency

increasing but remaining within the original 2.5 μs design envelope. For planning purposes, it is assumed that the HLT will continue to be RoI-guided. Monte Carlo simulation will be used to develop the improved Level-1 and HLT algorithms. From the data acquisition perspective, some detector readout links are expected to run at close to 100% capacity, and will probably limit the L1A rate to 60 kHz or less. Adjustments will be needed inside the TDAQ system to transport the larger events.

Early trigger upgrades are likely to be needed to support the forward physics programme. Although not directly related to the luminosity upgrade programme, they require topology algorithms at Level-1 similar to those planned for the Phase-I upgrade.

3.1.4 Trigger and Data Acquisition Issues at Phase-II

Major ATLAS detector changes are planned for Phase-II. The LHC machine timing is likely to be different, the Inner Detector and forward calorimetry will be replaced, and modified or new front-end electronics will be needed for most other detector subsystems. There will also be major changes to the trigger and timing distribution system. The Level-1 Trigger electronics will be completely replaced, and additional muon trigger chambers and a Level-1 Track Trigger are under consideration. The High-Level Trigger will require further changes to algorithms, with the RoI-guided strategy possibly no longer viable, and the HLT software infrastructure will also require updating. The high data volumes will require new detector readout, and probably significant changes to TDAQ internal dataflow, with further improvements to the internal networking infrastructure.

The long lead-time for custom electronics on the detector and in the trigger mean that key parameters (including maximum Level-1 Trigger rates and latency) must be fixed during the next one to two years. This requires development and testing of an integrated trigger strategy for Phase-II covering all trigger levels as part of the R&D work in this bid.

3.2 ATLAS Forward Physics Triggering

Although not directly connected with running above LHC design luminosity, the forward physics programme requires some adaptation in the Level-1 Calorimeter Trigger during the period covered by this proposal. It also requires significant use of event topology in the Level-1 trigger.

The main goal of the AFP programme is to observe Higgs production in the diffraction process pp $\rightarrow$ ppH. Assuming a low-mass Higgs (115–140 GeV), the signal in the ATLAS detector is two jets from the decay of the Higgs accompanied by the scattered protons observed in the proposed forward detectors at 220 m and 420 m from the interaction region. The observation of the scattered protons, essential to tag the Higgs decay, restricts the kinematics of the detectable Higgs boson; the low boost of the Higgs boson means that the two jets from its decay will be relatively low in energy (50–100 GeV) and will be produced with low acoplanarity (almost back-to-back in ϕ). To trigger on such diffractive events requires incorporating signals from the 220 m detector into the ATLAS trigger. (Because of the signal propagation time the 420 m detector cannot be used in the Level-1 Trigger.) Topological information is likely to be needed to reduce the trigger rate for two jets plus a forward proton to an acceptable level. For example, requiring the two jets to be central and back-to-back requires the mean η and $\Delta\eta$ of the two

jets to be used in the trigger. Furthermore, it may also be possible to utilise the ϕ correlation between the tagged proton and that of the di-jet system. As such processing would be implemented in the Level-1 Topological Trigger Processor for Phase-I running, the requirements of the AFP programme would impact on the detailed design of the L1Calo Topological Processor. However, the likely timescales for the proposed AFP programme mean that some L1Calo adaptation is needed before construction of the main topological processor is complete. This may be possible with firmware changes using the flexibility built into current modules. From the perspective of the L1Calo R&D, it is essential to understand the requirements of the AFP programme so that they can be incorporated into the design of the L1Calo upgrade. Simulation studies are required to understand the AFP trigger requirements.

3.3 Level-1 Calorimeter Trigger Upgrade

The current ATLAS Level-1 Trigger system has three main components: the Level-1 Calorimeter Trigger (L1Calo) for triggering on electrons, photons, tau leptons, jets and missing energy; the Level-1 Muon Trigger (L1Muon) for triggering on muons; and the Central Trigger Processor (CTP) in which the final trigger decision is made. The system was designed to operate up to the nominal LHC luminosity of 10^{34} cm^{-2}s^{-1}, with a Level-1 Trigger (L1) rate limited to a maximum of 100 kHz, matching the maximum readout rate of the front-end electronics in the ATLAS detector systems. Due to the finite on-detector buffer sizes, the Level-1 Trigger decision must be made within a maximum latency of 2.5 μs.

Triggering at the LHC is complicated by the event pileup, which for nominal luminosity corresponds to approximately 25 minimum-bias events per 25 ns bunch-crossing. At the higher luminosities foreseen for the Phase-I and Phase-II LHC upgrades, the effect of increased pileup will degrade the trigger performance. Increasing the trigger thresholds, which would significantly compromise the physics performance of ATLAS, is unlikely to be sufficient to control trigger rates. For the Phase-I upgrade, the addition of a new L1Calo Topological Trigger Processor is therefore envisaged. The design, development, and production of the Topological Processor form the main part of the L1Calo upgrade proposal. For the Phase-II upgrade, the L1Calo system will need to be replaced, both to cope with the increased pileup and to operate with the new calorimeter front-end electronics. Whilst the timescales are longer, the potential scope of the Phase-II upgrade requires that the initial design study starts in the period covered by this proposal.

3.3.1 Overview of the Current L1Calo System

The current L1Calo system is shown schematically in Fig. 3.2. The inputs to L1Calo are analogue signals from the calorimeters, summed into trigger towers of 0.1×0.1 in (η, ϕ). The summed analogue signals are processed in the Preprocessor Modules (PPMs), which sample the calorimeter signals at 40 MHz (reflecting the 25 ns bunch structure of the LHC). The PPMs produce processed digital values representing the transverse energy in each trigger tower, aligned in time with a particular bunch-crossing. The digital signals from the PPMs are sent to two processors working in parallel, the Cluster Processor (CP) and Jet/Energy-sum Processor (JEP). The individual Cluster Processor Modules (CPMs) and Jet/Energy Modules (JEMs) process signals from a particular ϕ quadrant and η of the detector.

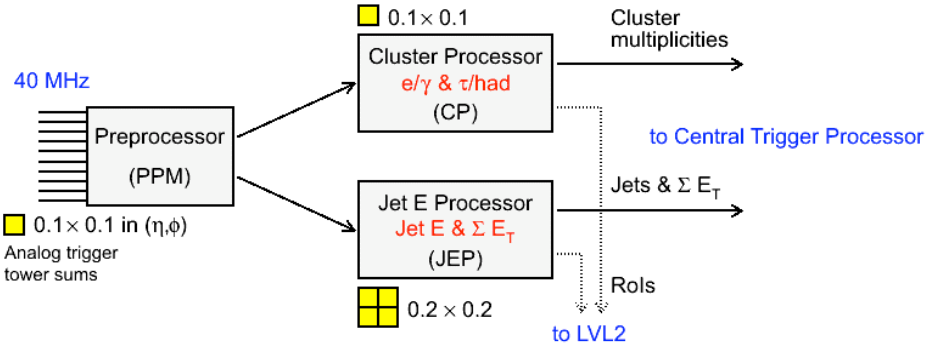

Figure 3.2: Schematic of the current L1Calo architecture. There are 124 PPMs. The Cluster Processor consists of 56 CPMs and 8 CMMs. The Jet/Energy-sum Processor consists of 32 JEMs and 4 CMMs. The Readout Drivers (RODs) and timing infrastructure are not shown.

L1Calo is based on trigger features which are identified in real time from the pattern of energy deposits in the calorimeters. For example, electrons or photons are identified as isolated clusters predominantly in one or two trigger towers. The CPMs identify high transverse-energy (E_{T}) clusters (electrons/photons and taus), and the JEMs identify high-E_{T} jets and perform transverse-energy sums. A set of eight Common Merger Modules (CMMs) are used to merge the signals from the different CPMs, and four perform merging of signals from JEMs. The output of the Calorimeter Trigger is twofold: counts of identified physics objects passing a pre-defined set of energy and isolation requirements are sent to the Central Trigger Processor, where the final Level-1 Trigger decision is made and a Level-1 Accept (L1A) is generated; and data describing regions of interest (RoIs) are sent from the CPMs, JEMs and some CMMs to the Level-2 Trigger. The RoI data include information on the position (η, ϕ) and E_{T} threshold passed by the candidate objects. Once detector readout has been completed following an L1A, the Level-2 Trigger uses the RoIs to access detailed data from the relevant parts of the detector.

Because of the high data rate (approximately 2000 Gbit/s), the Level-1 Trigger is a complex, highly parallel system which uses high-speed links and custom electronics modules throughout. In addition, the limited latency budget implies severe constraints on the data volume per event which can be passed between the various elements of the trigger. For example, the real-time output information from the CPMs describing electron/photon candidates is restricted to a 25-bit data word per LHC bunch-crossing. The Cluster Processor was designed and built entirely by the UK groups, as were all Common Merger Modules, the Readout Driver Modules and a significant part of the common L1Calo infrastructure.

3.3.2 Phase-I L1Calo Upgrades

For Phase-I of the LHC luminosity upgrade, the calorimeter readout will remain unchanged. In addition, given the timescale, it is neither possible nor desirable to change the main architectural layout of the L1Calo system. Hence, to cope with the increased luminosity, L1Calo needs to make better use of existing information. The present data sent to the CTP are essentially counts of physics objects, e.g. electrons, jets, etc., for a number of configurable thresholds. In addition, the Cluster Processor and Jet/Energy-sum Processor operate independently. Hence it is possible

for the same energy deposit to fire independently the electron/photon, tau and jet triggers. This limits the ability to trigger cleanly on combinations such as a high-E_{T} electron and a high-E_{T} jet. Given the constraints, the only realistic Phase-I upgrade path is to perform further processing of the trigger objects already identified by the CPMs and JEMs.

For this phase of the LHC upgrade, the ATLAS L1Calo collaboration proposes to construct a new Topological Processor, which would take more detailed information from the CPMs and JEMs and, in addition to the current trigger signals, would implement topological triggers and provide additional inputs to an essentially unchanged CTP. Level-1 Trigger decisions could then include requirements such as a pair of acoplanar high-E_{T} jets, angular correlations between electrons and jets, and potentially di-jet and di-electron invariant-mass triggers. The topological processor would also allow the possibility of including information from other triggers, e.g. L1Muon or a forward proton trigger. The overall L1Calo architecture including the Topological Processor is indicated in Fig. 3.3.

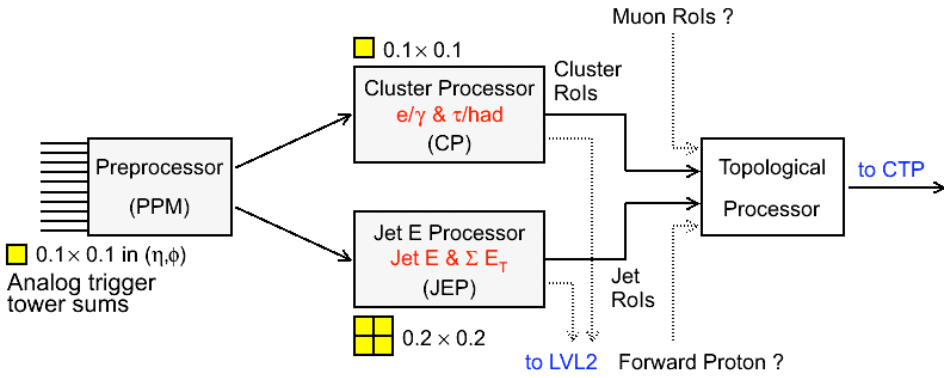

Figure 3.3: Schematic of the proposed L1Calo architecture for the Phase-I LHC luminosity upgrade.

3.3.2.1 The Topological Processor

The new hardware associated with the Topological Processor is relatively modest. The current system takes a little under 2.2 μs to form the trigger decision, leaving a latency reserve of approximately 0.3 μs which can be used for additional processing. To remain within this low latency, it is likely that the new processor would consist of a single crate. The Topological Processor would work on the basis of spatial information rather than simply counts of trigger objects. This requires more information to be transferred per clock cycle across the backplanes of the CP and the JEP. It has recently been demonstrated that by over-clocking the backplanes, a factor four in bandwidth can obtained (160 Mb/s). The best use of this additional bandwidth, and the exact functionality and architecture of the Topological Processor cannot be specified at this time; detailed Monte Carlo studies of triggers at Phase-I upgrade luminosities, including the effects of pileup, are needed to determine which topological features are most essential for the overall Level-1 Trigger.

A possible schematic layout of the Topological Processor is shown in Fig. 3.4. Here the counts and spatial information from the four azimuthal quadrants of the ATLAS detector

are brought together in four quadrant-merger modules. This enables overlaps between electron/photon, tau, jet, and muon triggers to be identified. The spatial information from the four quadrants is then brought together such that η and ϕ maps for the whole detector can be constructed (it may also be possible to form a two-dimensional $\eta - \phi$ map). This information would allow the implementation of the proposed topological triggers. Output data from the new topological trigger (e.g. indicating a separate high-E_{T} jet and isolated electron) would be sent to the Central Trigger Processor, along with the current trigger counts. In addition to the improvements to L1Calo, the Topological Processor would greatly enhance the flexibility of the L1 system.

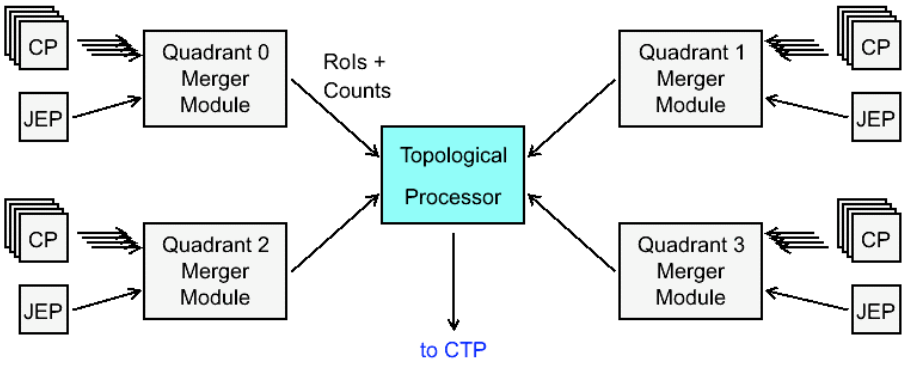

Figure 3.4: The currently favoured layout for combining the cluster and jet information for use in the Topological Processor. The Topological Processor consists of four input boards combining information from the four quadrants of the ATLAS detector, and FPGA-based processors for the implementation of the topological algorithms.

The impact on the existing hardware of the Topological Processor is relatively small. The majority of the L1Calo system (Preprocessor, CPMs and JEMs) remains unchanged. Of the existing hardware, only the 12 Common Merger Modules need to be redesigned. The Topological Processor would be housed in a single crate. A likely layout of the new crate would include several types of FPGA-based modules: four input boards (one per quadrant) processing the signals from the new CMMs, two core algorithmic processing boards to implement the new trigger algorithms, and one or more readout boards for monitoring and RoI output to the Level-2 Trigger. The crate will also need to be equipped with modules for control and timing. In addition to the hardware, new firmware is required to: i) extract and format the relevant information in the CPMs and JEMs, ii) repackage and transmit data to the Topological Processor, iii) merge the cluster and jet (and muon) information in the new quadrant merger modules, iv) implement the Topological Processor itself, and v) adapt the RODs to read out the new data formats.

3.3.3 Phase-II L1Calo Upgrade

The current L1Calo system will not cope with the high rates at the Phase-II LHC upgrade luminosity of 10^{35} cm^{-2}s^{-1} without seriously compromising the trigger efficiency. In addition, for Phase-II luminosities there will be a major upgrade of the calorimeter readout, in which digitisation will be performed directly on the detector. This means that the L1Calo front-end

needs to be replaced to handle digital rather than analogue signals. The high level of pileup, up to 400 minimum bias events per bunch-crossing, implies that L1Calo requires significantly better background rejection. To operate at Phase-II luminosities, it is likely that L1Calo will need to form trigger objects using calorimeter data at a finer level of granularity than the current trigger towers. It should also have greatly enhanced exclusive triggering capabilities based on topological information. This cannot be implemented in the current L1Calo system. It is also likely that the timing of the LHC machine will change, and that ATLAS will replace its internal timing and data acquisition systems, which by then will be at least ten years old. All of these changes imply that for the LHC Phase-II upgrade the entire Level-1 Trigger system will need to be replaced.

At this time it is impossible to specify the upgraded L1Calo system fully. Nevertheless a number of requirements can be identified. The L1Calo system must:

- Process the digital signals for the calorimeters. It is envisaged that the Preprocessor Modules (which currently perform the analogue-to-digital conversion and timing of signals) are replaced by Digital Preprocessor (DPP) modules which would perform bunch-crossing identification (BCID), baseline shift correction, E_{T} calibration, and any mini-tower summation.

- Operate with finer granularity features due to the high level of pileup, better utilising the spatial and depth granularity of the calorimeters. For example, the L1Calo trigger could use information from the full granularity of the ATLAS Liquid Argon ECAL: 0.003×0.1 (η, ϕ) in the first sampling layer, 0.025×0.025 in the middle sampling layer, and 0.025×0.1 in the rear sampling. The current 0.1×0.1 trigger towers would be replaced by smaller towers, e.g. 0.05×0.05. In addition, information from the finer granularity samplings, such as the location of the maximum of the shower and consistency with the expected EM profile, could be extracted and propagated to the rest of the system. This use of finer granularity information would greatly enhance the ability to identify narrow electromagnetic showers from electrons and photons in the high occupancy environment of the sLHC.

- Possess significant topological processing capability. To keep the Level-1 trigger rate below approximately 100 kHz, some of the current functionality of the Level-2 trigger would be moved to Level-1. This is only possible if the new processors provide detailed information on the identified features (electrons/photons, τ-jets, jets) rather than simple counts of features.

- Interface with the Level-1 Track Trigger. If there is a track trigger, and depending on the chosen architecture, L1Calo (and L1Muon) features will need to be passed to the track trigger to initiate local high-speed readout and processing.

Given the timescales for the Phase-II LHC upgrade, it might be tempting to defer the Phase-II L1Calo upgrade R&D for a number of years. However, the present trigger has taken more than ten years to produce, and due to the scope of the Phase-II upgrade project, i.e. the design, prototyping and construction of a completely new system, it is necessary that the initial design work proceeds in parallel with the Phase-I upgrade work. Specifically, design studies using simulated data are required to understand fully the requirements of L1Calo at the sLHC. In addition, R&D in the underlying technology used to implement the system needs to be completed.

One such possibility is to move away from a VME-based system to one based on emerging telecommunications standards such as ATCA (envisaged for the Topological Processor). This conceptual design work needs to commence as soon as possible.

3.3.4 UK Contribution to the L1Calo ATLAS Upgrade

The UK comprises approximately half (by groups and by FTEs) of the ATLAS L1Calo collaboration, and requests funding to contribute to the ATLAS L1Calo upgrade at a proportionate level. The L1Calo upgrade work naturally falls into four topics: design studies, hardware, firmware and online software, and engineering project management. The UK has the essential expertise and experience to make leading contributions in all four areas.

3.3.4.1 Design Studies

Detailed design studies studies using simulated data are needed to understand fully the necessary functionality of the L1Calo system at high luminosities; this understanding will form the basis of the detailed design of the upgrade. These simulations are technically challenging; it is not sufficient to simulate a single bunch-crossing. For example, the pulses from the Liquid-Argon (LAr) calorimeter are approximately $0.75\,\mu$s (~ 30 bunch-crossings) long, and pileup from previous bunch-crossings affects the digitisation and bunch-crossing identification. The MC studies for nominal LHC luminosity (10^{34} cm^{-2}s^{-1}) used for the Level-1 Trigger TDR did not fully account for the effect of pileup in previous bunch-crossings as there was no pulse history in the Geant3/Fortran-based simulation at that time. In addition, a number of other effects were not included in the MC trigger simulation used for the L1Calo studies in the TDR, and the modelling of the underlying physics has changed significantly. Consequently, the limitations of the current trigger close to and above the nominal LHC luminosity are unclear. Nevertheless, it is almost certain that for high luminosities, the current L1Calo could only achieve the required rate reduction with high thresholds which would significantly impact ATLAS physics. Understanding the performance of the current L1Calo trigger for Phase-I luminosities and beyond is an essential part of this proposal. Soon it will be possible to run the full ATLAS Monte Carlo (MC) simulation for a luminosity of 3×10^{34} cm^{-2}s^{-1} including the full effects of pileup. Nevertheless, with the current simulation software this will be extremely time-consuming. Hence, the L1Calo trigger design studies will benefit greatly from the work to optimize and rework the simulation process which forms part of WP9.

Whilst understanding the limitations of the current L1Calo trigger is a necessary first step in the upgrade work, the main goal of the Phase-I design studies are to study and define the topological signatures to be implemented in the Topological Processor. There are many possibilities, including invariant mass combinations, gaps between jets, angular correlations, and missing-E_{T} from objects. There is also the possibility of combining electron and jet information with spatial information from the Level-1 Muon Trigger. In considering any new topological trigger, the implementation in firmware needs to be taken into account. These studies are non-trivial; understanding and mitigating the effects of pileup will be both technically and intellectually challenging.

In parallel with the Phase-I design studies, we propose to start to study the necessary functionality of the L1Calo trigger at the sLHC luminosity of 10^{35} cm^{-2}s^{-1}.

For the three-year period of this proposal, the proposed design work is:

- To gain a detailed understanding of the luminosity limitations of the current Level-1 Calorimeter Trigger, based on full simulation of the system including pileup (design task 1).
- To identify the most important topological algorithms and to study their impact on ATLAS physics (design task 2).
- To start to study the necessary functionality of the L1Calo trigger at Phase-II luminosities (design task 3).

3.3.4.2 Hardware

For the Phase-I upgrade, the majority of the L1Calo hardware (Preprocessor, Cluster Processor Modules and Jet/Energy Modules) is unchanged, although CPM and JEM firmware modifications will be required. The 12 CMMs will be redesigned to use the increased bandwidth gained by over-clocking the backplane. This will enable spatial information to be sent to the Topological Processor, which will be housed in an ATCA (or similar) crate. The detailed design of the Topological Processor will depend on the requirements identified by the MC simulation studies. As currently envisaged, the Topological Processor will consist of approximately ten boards. Four input modules will receive the electron/photon and tau cluster information and the jet information from the four quadrants of the detector, along with Muon Trigger information and potentially information from the forward-proton trigger. The input modules will form the quadrant hit sums and will resolve overlapping trigger features. The global topological algorithms will be implemented in two (or possibly more) dedicated FPGA-based algorithmic processing boards. The Topological Processor crate will also require a timing module. Because of the number of connections needed, two output boards will be required to feed the existing trigger sums and the new topological triggers to the Central Trigger Processor. Hence, in this model the ATCA crate contains:

- Four input modules;
- Two (or more) algorithmic processing modules;
- Two output modules; and
- One timing module.

It is worth noting that the Topological Processor is a non-trivial system, with a potential input data rate of approximately 700 Gbit/s. The proposed UK hardware contribution is the development and production of approximately half of the Topological Processor. In the above model, the estimated total cost of the hardware for the Topological Processor, including prototypes and spares, is £410k. On the assumption that the L1Calo Phase-I upgrade will be installed at CERN in 2014 for operation in 2015, construction of the Topological Processor will need to be completed in 2013. Consequently, the main spend on equipment will take place in the period covered by this grant application, and for this proposal we are thus requesting £200k for the UK hardware for the Topological Processor. Details are given in Section 5.3. Given that the final architecture of the Topological Processor will depend on the detailed simulation studies, a working allowance of 30 % on the equipment spend is requested.

For the Phase-II upgrade, it is expected that the current Level-1 Trigger will be entirely replaced. The upgraded L1Calo will receive fine-grain calorimeter information at roughly ten times the data rate of the present trigger, and will be expected to provide a very high level of selectivity. Experience constructing the present L1Calo shows that the usable link speeds are a key factor in overall system design, so it is important to start R&D on very high speed fibre and backplane links as soon as possible. The goal would be to produce demonstrator modules with multiple low-latency backplane and fibre links, each running at 510 Gbit/s, ten times the speed of current low-latency links in L1Calo. The hardware infrastructure would be provided by the platform (e.g. ATCA) chosen for the Phase-I upgrade.

For the three-year period of this proposal, the proposed hardware tasks are:

- To produce a conceptual design of the Phase-I Topological Processor based on the results of the MC simulation studies (hardware task 1).
- To design and construct approximately half the Topological Processor (hardware task 2). This is likely to correspond to the production of five FPGA-based boards for the ATCA crate.
- To commission the Topological Processor ready for possible installation at CERN in 2013 (hardware task 3).
- To start to investigate the potential of an ATCA-based system to replace the current VME-based L1Calo system for the Phase-II upgrade, and to develop test modules to evaluate high-speed backplane and fibre link technologies in the 510 Gbit/s range (hardware task 4).

3.3.4.3 Firmware

The Phase-I upgrade will require significant firmware development. Based on experience with the current L1Calo trigger, this aspect should not be under-estimated. There are a number of essential tasks, most of which build on existing UK firmware expertise and experience:

- Develop firmware to exploit fully the flexibility built into the existing CMMs. This work is needed independent of the final Topological Processor architecture. For LHC design luminosity, it may also provide some of the functionality required by the forward physics programme until the topological processor is complete.
- Modify the firmware in the existing CPMs and JEMs to extract more information, and to send this data over the current backplane at higher speeds to the new CMMs.
- Modify the readout firmware for the current Readout Drivers to capture the expanded CPM and JEM data before it goes to the backplane and the new CMMs.
- Develop the firmware for the input modules of the Topological Processor, which combine the cluster and jet information from the detector quadrants.
- Implement the new algorithms in the Topological Processor, including those relevant for the ATLAS forward physics programme.

In addition, the implementation of the new Topological Processor will require development of online software. Although this may sound like a relatively small element of this proposal, it is, nevertheless, essential for the project. These tasks will be shared between the L1Calo institutions. For the UK part of the Phase-I L1Calo upgrade we propose to develop the firmware for the systems where the UK has existing responsibilities, and to develop firmware for the Topological Processor. This latter work is essential to maintain UK expertise in cutting-edge real-time processing.

For the three-year period of this proposal, the UK firmware tasks in this proposal are:

- To develop the firmware to enable existing modules (CPMs and CMMs) to run with and utilise the higher bandwidth allowed by over-clocking the backplane (firmware task 1).
- To develop readout firmware for the current Readout Drivers to handle new data formats for the existing L1Calo modules (firmware task 2).
- To develop the firmware for the input modules of the Topological Processor, which combine the cluster and jet information from the detector quadrants (firmware task 3).
- To implement the new algorithms in the Topological Processor, including those relevant for the ATLAS forward physics programme (firmware task 4).
- To develop the online software associated with the Topological Processor (firmware task 5). This is an essential part of the project.

The first two of these tasks, both of which use existing hardware, can and need to commence in the near future. For the period of this proposal, we do not foresee any significant firmware development for the Phase-II upgrade.

3.3.4.4 Engineering and Scientific Project Management

The UK expects to contribute to the international scientific and engineering project leadership for the L1Calo and overall TDAQ upgrades. From the perspective of project management, coordination between the tasks in WP2 and WP3 described in this proposal, and between the work in the UK and the rest of the L1Calo collaboration, is essential. We estimate that 0.3 FTE of senior engineer effort is sufficient for the engineering aspects of the management of this project. It is essential that effort is earmarked for this function as the proposed work will occur in parallel with the operation of the existing system.

3.3.4.5 Track Record

The UK part of the L1Calo collaboration has the necessary expertise and experience to complete the above programme of work successfully. ATLAS-UK has led the L1Calo collaboration, including the initial R&D and the writing of the Level-1 Trigger Technical Design Report. In the design, construction and commissioning of L1Calo, the three original UK groups (Birmingham, QMUL and RAL) were responsible for roughly half the project's hardware, firmware and software. These groups were central to delivering the complete system in full, on time and on budget. The Cambridge group, who have recently joined the L1Calo collaboration, bring further expertise in simulation and in design and construction of high speed electronics. Within the

L1Calo collaboration, the UK has an exceptionally strong track record. All the L1Calo project leaders have come from the UK groups. Most of the other major positions of responsibility in the management have been filled by UK people, including key coordination posts. Below we list some of the roles and responsibilities fulfilled by ATLAS-UK.

Hardware: The UK is entirely responsible for the Cluster Processor which implements the electron/photon and tau triggers, and is responsible for the Readout Drivers used to format and send the trigger data to the Data Acquisition system and the RoI information to L2. The UK also supplied the Common Merger Modules (CMMs) for the Jet/Energy-sum Processor and the Timing Control Modules for the entire trigger system.

Firmware: The UK wrote all the firmware for the Cluster Processor and most of the firmware for the CMMs. For the Readout Drivers, the UK developed all the firmware, including that needed to handle the various types of data used for readout and the RoIs.

Online software: The UK has led and made major contributions to the online software for controlling, testing, running and monitoring the L1Calo system, including the complex interactions with the ATLAS Run Control system, databases, and the Detector Control System.

Offline software and physics simulation: The UK has led the offline software development for the trigger: the raw data decoders, the event data model, and the offline simulation tools. Much of the software for understanding the trigger operation and performance was designed and written by UK physicists and programmers. We have led the work in physics simulation and performance studies of the trigger, and carried out many of the studies for the design and operation of L1Calo as well as optimisation of trigger menus. A UK physicist is the main contact for the ATLAS trigger menu group and physics community.

3.3.4.6 Relationship with Non-UK Institutes in L1Calo

For the upgrade work, the original L1Calo collaboration (Birmingham, Heidelberg, Mainz, QMUL, RAL, and Stockholm) has been strengthened by the addition of new collaborators from Argonne National Laboratory (ANL), Cambridge University, and Michigan State University (MSU). The nine groups have agreed that all will contribute to operation of the installed trigger system. This will release modest effort from the original six groups so that all nine groups in the enlarged L1Calo collaboration can contribute to the Phase-I and Phase-II L1Calo upgrades.

The proposed UK work on the L1Calo upgrade project is consistent with the plans of the non-UK institutes. In the three years covered by this proposal, the Heidelberg group plan to work on Preprocessor Module upgrades to improve the accuracy of the digitisation in high pileup conditions. ANL and MSU are both contributing to Monte Carlo studies, and together with Mainz and Stockholm will contribute to trigger hardware and firmware development. Given the need to develop the Phase-I upgrade in a relatively short period, while at the same time performing initial R&D for Phase-II, the overall international effort is a reasonable match to the project requirements.

Milestones and other details are given Section 5.3,WP6.

3.4 Level-1 Track Trigger

Given the harsh environment at sLHC, especially in Phase-II, the proposed upgrades of L1Calo and L1Muon may not by themselves be sufficient to extract all the desired physics from ATLAS.

The large number of minimum-bias interactions per bunch-crossing will inevitably degrade the performance of isolation-related variables in L1Calo, and the increased cavern radiation will lead to high occupancies in the muon chambers and a consequent increase in the rate of muon triggers due to fakes. Therefore, without additional information at the Level-1 Trigger, there is the risk that the p_{T} thresholds, particularly for single or combined lepton triggers, will have to be raised to values that will compromise the physics reach of ATLAS.

The only other source of information that has the potential to enhance the purity of the events selected at Level-1 is the Inner Detector. Given that a new tracker will be built, there is the opportunity to instrument it with Level-1 triggering capabilities. Indeed, in the current ATLAS, where tracking is not used at Level-1, tracking information is an essential ingredient in the event selection in the High-Level Trigger. Matching tracks found in the Inner Detector with objects in the calorimeters or the muon detector is a key handle for achieving the required background rejection at the current Level-2 Trigger. In addition to this matching, tracking information would have two more important uses for the Level-1 Trigger at sLHC. First, it would allow verification of whether multiple Level-1 trigger objects came from the same proton–proton interaction. This would significantly reduce the rate of double lepton triggers, which is expected to be dominated by events where two objects (most often a true lepton and a fake Level-1 lepton) come from different z positions along the beam line. Second, it may provide the capability to perform track-based isolation for Level-1 electron, muon or tau candidates, again based on the z-positions of tracks along the beam line. Finally, a track trigger would provide much needed flexibility and redundancy, and therefore robustness, to the Level-1 Trigger system. This is vital, especially given that the detailed physics goals of the sLHC will only be known after some years of running at the LHC. This flexibility should ensure that the full physics potential of sLHC can be exploited, irrespective of what the LHC discovers.

ATLAS-UK initiated the Level-1 Track Trigger project (L1Track) within ATLAS, and is currently providing leadership and most of the intellectual and technical input. Given that many aspects of the design of the tracker upgrade depend critically on the design of L1Track, as will be described in the following, and that many architectural decisions for the tracker have to be taken in the coming years, it is imperative that the R&D described below is performed during the period of this proposal. The aim of the three-year programme in this proposal is to lead the work on the design and specification of the full L1Track system, building on the established ATLAS-UK leadership and expertise in the tracker and trigger upgrade projects. This work will also bring closer the ATLAS-UK tracker and trigger upgrade projects, and will strengthen the UK role within ATLAS.

3.4.1 L1Track design options

Given the colossal number of channels in the Inner Detector, it is not feasible to read out the entire detector at 40 MHz to produce L1Track information. Two different approaches to reducing the volume of tracker data to be read out for L1Track are currently under consideration in ATLAS. The first is a regional readout of the Inner Detector, based on Region-of-Interest (RoI) information from L1Calo and L1Muon; the second is a standalone approach, using dedicated tracker layers to select hits from high-p_{T} tracks and read out only those for further processing.

The regional readout approach assumes that L1Calo and L1Muon would identify potentially interesting features at a few hundred kHz. They would issue fast readout requests to specific

regions in the tracker at this rate, providing the (η, ϕ) position of the objects identified as interesting. In this way, only a small fraction of the tracker would be read out, and only at a reduced rate, so the required additional bandwidth for the Inner Detector would be very modest. Several variations are possible in this approach, depending on how fast the regional Inner Detector data can be read out and processed, and on the overall Level-1 Trigger latency envelope. Ideally, tracking information should be used directly within the Level-1 Trigger. However, ATLAS has also discussed an option for a two-stage Level-1 trigger, for use if the Inner Detector readout is too slow. This would require additional buffers for all ATLAS detectors, in which data would be held until the slower, definitive Level-1 Trigger decision was available.

The alternative, standalone track triggering approach relies on the use of doublets of closely spaced layers of silicon wafers. These are electrically connected and instrumented with coincidence logic, so that relatively straight (high-p_{T}) tracks give a coincidence while those bending more (low-p_{T}, the vast majority of tracks in minimum-bias events) do not. This would give a significant on-detector data reduction, since only the coincidence data would be sent off-detector to the L1Track processors. Nonetheless, being independent of L1Calo and L1Muon, it would have to be running at the bunch-crossing frequency, hence the overall bandwidth is still very large. This idea was first proposed in [30] and is currently being pursued as the baseline solution by CMS.

While a standalone track trigger is attractive, CMS studies have shown that for it to be robust and achieve the required performance, such a system would be extremely complex and require significantly more material in the tracker than otherwise. There is, therefore, a concern that the physics benefits from the track trigger would be at a cost of degradation in the calorimeter and offline tracking performance. As there are important differences between ATLAS and CMS, different track-trigger designs may be optimal for each experiment.

For ATLAS, preliminary feasibility studies indicate that the UK-proposed, RoI-based approach can lead to a light, flexible, non-invasive track trigger design that will satisfy the experiment's needs at sLHC. Therefore, the ATLAS-UK L1Track R&D programme will focus in the next three years on developing current ideas into a concrete full design of the track trigger system, based on the RoI-driven approach. Meanwhile, the alternative standalone approach will continue to be studied by our international collaborators within ATLAS, whose progress will be followed closely so that we can converge on a common design that is optimal for ATLAS.

3.4.2 L1Track simulation studies

Simulation studies will have a vital role in comparing the various options and converging on the best design for L1Track. These studies will range from a discrete event simulation to single particle and full high-luminosity event Monte Carlo simulations, addressing the major physics and technical questions of the design.

3.4.2.1 L1Track-specific software tools for high pileup studies

As many of the studies will have to be performed using high-pileup Monte Carlo simulations, it is important to develop flexible tools and techniques that will allow a wide range of ideas to be tested rapidly. Some tools have already been developed in the UK, for example a fast pileup framework whereby simulated events in single p-p collisions are merged after digitisation, and the effects from out-of-time bunch-crossings are ignored. This technique has been shown to

be a good approximation, within 15%, of the full simulation in the tracker in terms of overall occupancies.

In the period of the proposed bid, more software tools and techniques will be developed that are essential for L1Track design studies:

- Infrastructure to produce reduced-size datasets by stripping off unnecessary event information from fully simulated, high pileup events, hence making the data more portable and faster to process (DPDs in the ATLAS jargon);
- a lightweight framework operating within the ATLAS offline software for L1Track pattern recognition studies to facilitate evaluation of different pattern recognition strategies;
- if producing and analysing fully simulated, high-pileup events continues to be very resource demanding, a range of fast simulation alternatives will need to be investigated, including the fast tracking parameterisation package FATRAS and the fast calorimeter shower parameterisation package FastCaloSim, and will be adapted to the needs of the L1 trigger simulation. This work will be performed in close consultation with the ATLAS-UK fast simulation experts.

3.4.2.2 Physics-based design studies

The parameters that will largely determine the design of the RoI-based L1Track include the event rate at which regional readout requests are sent, the number of RoIs per event, the size and shape of the different types of RoIs (electron, muon, tau, etc.), and the desired track parameter resolutions. All these questions will be studied in close collaboration with the L1Calo and L1Muon experts on common issues.

The size and shape of the different RoI types is particularly important, as it has a direct impact on the data bandwidth and readout efficiency. An example lepton RoI in the tracker is shown in Fig. 3.5, with dimensions $\Delta\phi = 0.2$, $\Delta\eta = 0.2$ at the outer tracker radius, increasing to $z = 40$ cm near the beam line to account for the spread of the luminous region in z. This spread in z means that modules in layers closer to the beam line, and especially the innermost Pixel barrel layer, are more frequently inside an RoI than other layers, potentially causing bandwidth or deadtime issues. It is therefore important to understand how many layers of the Inner Detector (starting from the outer layers) are necessary to achieve the required track parameter resolutions in L1Track for the different types of RoIs. The size of the RoIs will also depend on whether track-based isolation proves useful for Level-1 lepton triggers, and studies will be performed to evaluate this possibility.

3.4.2.3 Pattern recognition studies

Pattern recognition studies will be performed, bearing in mind the need to develop an algorithm which can be ported to hardware, and must therefore be simple, robust and adaptable for low-latency execution. The work will draw on existing expertise in pattern recognition for the current ATLAS Level-2 tracking, as well as experience with Level-1 tracking systems in previous experiments. As in the current Level-2 tracking, one of the key parameters is the minimum p_{T} threshold for track reconstruction. In addition, the requirements on the level of fake tracks vary for the different RoI types and will depend strongly on whether track-based isolation is

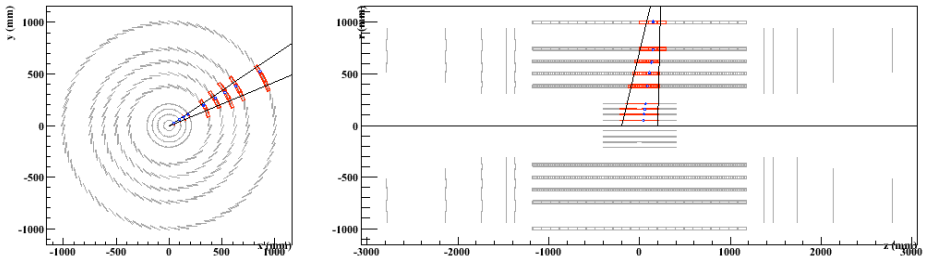

Figure 3.5: Example lepton RoI in the Inner Detector.

required, as well as on the number of layers used. Studies will also investigate whether there is adequate information in the r-ϕ layers alone, so as to reduce the amount of data read out per RoI, and whether in the very busy environment of sLHC an outside-in pattern recognition strategy is more appropriate.

All these pattern recognition studies will be performed in conjunction with investigating the design and various hardware alternatives for the track-trigger processor boards, as detailed in section 3.4.3.4 on the hardware for the track trigger processors.

3.4.2.4 Discrete Event Simulation

A Discrete Event Simulation (DES) of the dataflow in the L1Track system will be essential for specifying the bandwidth requirements and latencies in the various parts of the system. This will enable us to giagnose possible throughput bottlenecks and deadtime due to busy components, and to identify those parts of the system which need to be faster so that such bottlenecks can be avoided. An existing DES framework such as SimPy [31] will be used, in which all the L1Track-specific components to be optimised will be embedded. Many of the current ideas for L1Track involve sharing common lines, fibres, etc., with the normal readout system of the tracker, so this work will be done closely with the ATLAS-UK tracker upgrade readout experts to ensure that the DES is developed in a coherent fashion.

3.4.2.5 Level-1 trigger rate studies and interface to the HLT

The performance of L1Track will ultimately have to be assessed as part of the overall ATLAS trigger. Work will proceed in close consultation with the (ATLAS-UK) L1Calo and (ATLAS) L1Muon experts to take ideas that are used in the current ATLAS Level-2 Trigger and adapt them for the Level-1 Trigger in order to estimate trigger rates and optimise the performance of the overall system. These studies will have to be performed with various levels of detail throughout the period of this R&D. Initially, they will be based more on parametrising/extrapolating the performance of the current (Level-1 and levelii) trigger algorithms. As the understanding of the potential improvements to the various sub-systems and to the capabilities of pattern recognition get refined, the trigger rate studies will give more guidance in how to optimize the overall performance. Finally, the use of L1Track information to seed the HLT tracking will be studied in collaboration with the ATLAS-UK HLT experts, as this may improve the performance and speed of HLT tracking pattern recognition.

3.4.3 Off-detector hardware and data paths

3.4.3.1 Overall Design

The track trigger builds on the existing Level-1 Trigger architecture, in which a potentially interesting event is identified, and a signal synchronous with that event is sent to the detector front-end (FE) modules. The FE ASICs transfer the event data from their pipelines to a readout buffer where they are queued and sent off-detector.

For regional data readout, the process has two important differences:

- The trigger in this case is a regional-readout-request (R3) only, and is not broadcast to all FE modules. Instead it is sent only to the Inner Detector modules inside the RoI.
- Readout is minimally buffered — when an FE module receives an R3, it must return the data as fast as possible (with known latency) employing prioritised multiplexing or a separate data path.

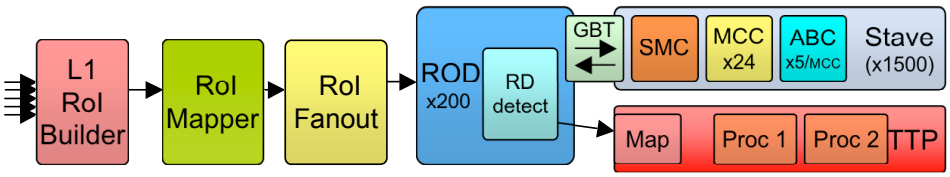

Figure 3.6: Overall system layout.

The Track Trigger process starts with the receipt of one or several RoIs from the L1Calo or L1Muon system by the Track Trigger RoIMapper (RoIM). This information is decoded and synchronised, generating readout requests to be sent to the modules within the RoI. At this stage the physical geometry of the detector is used to send R3 signal to the specific hardware chain connected to the desired FE module. The RoIM is described in detail below.

Using the topology of the current ATLAS SCT and Pixels, all FE modules are connected to off-detector Readout Drivers (RODs), so the RoIM will send groups of R3 signals identifying individual ROD channels. On each ROD, the R3 is integrated into the control link and sent to the stave. The SuperModule Controller (SMC) decodes the signal onto dedicated R3 signal lines connected to each module, identifying which should be read out.

The FE modules comprise a Module Control Chip (MCC) and many ABC FE ASICs. Upon receiving the R3, the MCC prepares for readout of track trigger data while forwarding the R3 signal to the ABCs. The ASICs process the event data and send it off-detector as fast as possible. In the case of dedicated track trigger links, these data would go directly to the Track Trigger Processor (TTP), otherwise they are identified on the ROD and forwarded on to the TTP.

3.4.3.2 RoIMapper design

The RoIMapper provides the hardware interface from the L1Calo and L1Muon Trigger systems to detector geometry-specific signals. It is tasked with three functions: interface to the L1Calo/Muon RoI system, map RoIs to physical module connections, and interface to the control system that drives the modules.

- Level-1 RoI Interface: The trigger system will probably send each RoI as a number giving the $\eta - \phi$ location and the type of RoI (which essentially defines the shape). These data will arrive at the RoIM as parallel data on a differential copper or a GBT (or similar) optical link. It is unclear whether each trigger subsystem will be connected separately or share a fibre/bus, but it will be necessary to time-align all RoIs for a given bunch-crossing. As the RoIM is at the top of the Track Trigger tree, it should also be able to function as a standalone R3 signal generator, using random or preloaded RoI information.

- RoI $\rightarrow$ Module map: The RoIM uses the RoI type and location to generate a list of target modules. Connection to these modules is via a channel on a ROD, so the module-stave-ROD mapping needs to be used. This needs to happen as fast as possible (< 25 ns), so the RoI-to-module mapping will be pre-loaded in RAMs, CAMs or associative memories. These could reside on an FPGA or in dedicated hardware — evaluation is needed. This system needs to be programmable to account for changes in layout or module-to-ROD mapping.

- ROD Interface: The RoIM sends signals to all the connected RODs in the system. As the final system could have ~ 3000 readout links (strips and pixels) connected to ~ 200 RODs, a fan-out system will be required. Additionally, the system must be synchronous, being able to fit all R3 information inside a 25 ns packet. Fan-out hardware will be needed and will be located in the ROD crate or next to the RoIM board. The distribution of lookup logic will to be optimised between RoIM, Master and Crate fanouts and RODs.

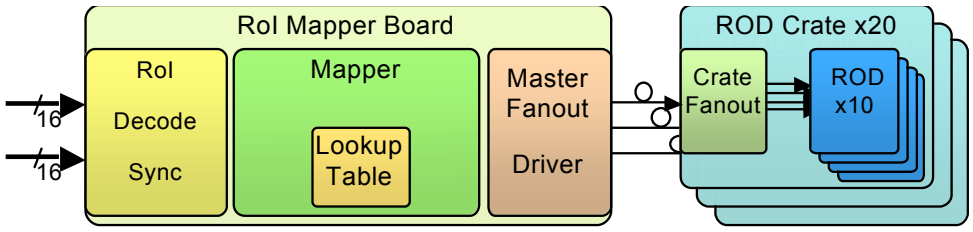

Figure 3.7: RoIMapper functional schematic.

3.4.3.3 RoIMapper Hardware Program

While requirements for a complete system required by the upgrade will be evaluated, it will only be possible to study a fraction of the system in hardware. Development will be consistent with larger system requirements, but scaled for bench testing. The use of RAM, ROM, CAM, or AM architectures will be investigated in FPGAs, with a view to understanding if ASICs are needed.

Construction of a demonstration system is planned. This will be integrated with the Tracker upgrade DAQ/Control prototype that ATLAS-UK plans to develop, as described in the tracker upgrade sections of this proposal. Depending on their availability, the early SMC prototype and modules will also be integrated.

Three units are required:

1. An RoIM development platform, using a high-power FPGA board with optical links. A good candidate is the ATLAS DAQ WP development board (HSIO), which is well known to the project.

2. An RoI generator (Level-1 RoI emulator), to generate signals simulating the Level-1 systems, to drive the RoIM. This could use a cheap development board or a subset of signals in the available on the HSIO.

3. An R3 Fan-out, to look at options for sending many R3s to multiple RODs. This could be a plug-in for HSIO, or use existing HSIO optical interfaces.

3.4.3.4 Track trigger processor

Although it is not proposed to build the Track Trigger processor (TTP) boards at this stage, it is important to understand the requirements for these boards and include them in the overall system specification. This work will be carried out in close connection with the pattern recognition studies, since the pattern recognition strategy influences and depends on the hardware technology on which it would be implemented. The aim of this task is to evaluate the alternative hardware choices for pattern recognition, including FPGAs, network search engines, and associative memory-based chips, and to specify the I/O requirements and interfaces of the TTPs to the rest of the system.

3.4.4 On-detector electronics and readout

The Level-1 Track Trigger relies on most other systems of the Tracker upgrade project integrating the required functionality. In this proposal, the tracker upgrade DAQ, Stave Interface (SMC), Tapes, Modules and Hybrid will all need to include aspects of L1Track. Additionally, it will be essential to work closely with the non-UK projects, particularly the MCC and ABC ASIC designs, as well as the CERN GBT/Versatile Link projects, to ensure that the required functionality for a track trigger is incorporated in these designs.

3.4.4.1 Stave Tapes Signal Distribution

Options will be investigated for distribution of regional-readout-request (R3) signals. To target a subset of modules/MCCs on a stave, ideally a separate line to each MCC is required. As stave resources are limited, investigations into trade-offs between increased numbers of control-links and higher transfer rates are needed. Regional data (RD) is more demanding in latency and synchronisation of stave readout resources than normal event data. It requires fast links which are always available. Merging this data into the existing readout architecture may be possible, and would probably be more efficient in terms of material and power. On the other hand, dedicated fast readout links would improve performance, which in turn could reduce the front-end pipelines and buffer resources and add efficiency there instead. These two alternatives will be evaluated in order to converge to the optimal solution.

3.4.4.2 End of Stave

Implementation of regional readout needs dedicated signals decoded from the control (optical) link. It is unclear whether this functionality will reside in the GBT hardware, SMC hardware or need additional components, but at the very least drivers will be needed for the regional-readout-request lines. In the opposite direction (data path) the option of dedicated regional-readout lines would require separate hardware either to multiplex into the existing optical link or to drive a second link. These issues need to be developed to understand better the additional needs.

3.4.4.3 Hybrid/Module

The hybrid has very tight constraints mainly due to its size. Any changes need to be carefully investigated. As the regional readout design advances, especially with respect to ABC ASIC, its integration with the hybrid must be considered. The implications of additional lines between the ASICs and the MCC as well as between the MCC and the tape need to be investigated. The feasibility and efficiency of the various design options will be studied. It is anticipated that a new hybrid prototype will need to be developed, and this has been included in the capital cost of the Track Trigger project.

3.4.4.4 DAQ (Off-Detector)

The prototype DAQ is tasked with reading out a large variety of hardware, from single-chip test boards to entire staves. As regional readout will be incorporated into some of these designs, the DAQ will need to be able to deal with both testing and co-existing with the system. This will include acting as a stand-in for SMC/GBT, and handling R3 command decoding and de-multiplexing.

The possibility of having dedicated regional-readout links from the MCCs means that the number of LVDS links being received by the DAQ could double. This needs to be factored into the DAQ design (probably requiring two HSIOs for stave testing).

Once off the detector, regional data needs to be split from event data and sent to a Track Trigger Processor (TTP). Although it is not planned to build a dedicated TTP in this bid, the use of the readout hardware as a test-bed for algorithms and for architecture evaluation would be beneficial.

As the DAQ will control or interface to trigger sources/generators, a regional-readout interface will also be needed.

3.4.4.5 ASICs

The Track Trigger requires support in the Module Controller Chip and ABC-X FE ASICs, although these developments are not UK projects. The UK work will include simulation studies into data compression, data formats and shared readout, and it is proposed to develop these into firmware as a possible contribution to the ASIC projects, and as a demonstrator.

- ABC: Regional readout introduces a new signal to the ASIC design. When an R3 is received, data must be copied from the pipeline and prepared for transmission. Dedicated signalling or “queue jump” logic are required to ensure the data leave the ASIC quickly.

To lower latency and bandwidth requirements, methods for reducing regional data volume on the chip will be investigated.

- MCC: The MCC passes data off detector from the ABCs via the stave tape and SMC. In the case of shared readout lines, it plays the important role of arbiter. It will also be responsible for adding headers, etc. to label the regional data. When an R3 is received by an MCC, a "hold" signal is sent to the FE ASICs, temporarily freezing normal readout. The regional data packet is then transferred from the ASICs with the requisite header and trailer, and then normal readout resumed.

3.5 High-Level Triggers

3.5.1 Background

The HLT comprises the hardware and software of the HLT farms (Level-2 and Event Filter) and the dataflow system. The farms consist of a total of some 2300 PCs running the trigger selection software. The dataflow system includes the Readout System (ROS), Event Builder, and network infrastructure. The ROS stores event data in Readout Buffers (ROB) pending a Level-2 decision, and serves the data to the Level-2 processors and Event Builder. The Event Builder assembles the data fragments, for events accepted by Level-2, prior to processing by the Event Filter and, for events passing the trigger, subsequent output to storage.

The ATLAS-UK HLT groups have played a leading role in designing, building and commissioning the HLT hardware and developing the event-selection software. All the construction-phase deliverables of the ATLAS-UK HLT project were completed on time and were in full use at the LHC start-up, as well as in the combined cosmic runs of the last quarter of 2008 and summer 2009. ATLAS-UK led the design, production, installation and commissioning of the ROB-in boards that physically implement the ROBs. Half of the 700 boards were manufactured and tested in the UK. The UK also played a significant part in developing the ROS software, which handles data requests and data movement within the ROS PCs. The UK contributed to the purchase of the PCs for the HLT processor farms and the TDAQ network equipment, having significant input to the specification and procurement of all these commercial hardware items, and currently provides the TDAQ resource coordinator responsible for these procurements.

The UK has made a major contribution to the Event Selection software, providing leadership in the areas of HLT track reconstruction, core HLT software, and electron and B-physics trigger selections, as well as providing overall coordination of the work on trigger selection algorithms. The UK provides the current Trigger coordinator. For Level-2 tracking, the UK has developed all the components of the Level-2 tracking chain, from data preparation to pattern recognition and track fitting, as well as fast secondary vertex reconstruction. For Event Filter tracking, the UK was largely in charge of developing the wrapper tools for embedding the offline tracking code in the HLT framework. In addition, the UK has led the work for developing the online monitoring tools and the validation tools for both the Level-2 and Event Filter tracking. The UK played a leading role in the design of the core HLT software, including the Trigger steering software. This component controls the running of the trigger algorithms in step-wise processing chains, which are configured so as to implement the trigger menu, and forms the trigger decision. Other key components of the core HLT software, developed in the

UK, include: the TriggerTool, which is the interface to the Trigger Database and is used to construct/retrieve/modify the trigger menus and configure the trigger for online and offline operations; the RegionSelector, which uses lookup tables to provide a list of detector modules contained inside a Region of Interest; and the Serialiser, which converts the trigger objects and trigger decision into bytestream format when storing events. Finally, the UK has been playing a leading role in developing and optimising the algorithms which implement the electron and B-physics trigger selections as well as contributing to the development of the overall trigger menus for the different periods of running.

The UK HLT M&O commitment covers major responsibilities for core trigger software, validation, trigger configuration, inner detector software, and the selection software for electron, muon, B-physics, tau and missing-energy triggers. In addition, the UK M&O commitment includes support of the ROS and HLT farm infrastructure, including effort for a rolling replacement programme. The programme of work for the HLT upgrade builds on these core areas of expertise. The Phase-I developments are closely integrated with the ongoing work for the evolution and tuning of the trigger as luminosity increases up to design values. For Phase-II, the overall software and hardware architecture will be re-evaluated and substantially upgraded. During the period covered by this bid, the various architectural options will be studied and different trigger strategies reviewed. This will allow key decisions to be made that will enable the detailed design and implementation work to start soon after the time-frame of this bid. Having played a leading role in designing, building and commissioning the ATLAS High-Level Trigger hardware and event-selection software, the UK groups have experience and expertise vital to this work.

3.5.2 Phase-I

3.5.3 Track-based Selection Software

The high level trigger code using Inner Detector tracking information must be upgraded to cope with the increased occupancy at luminosities above design values. This means tuning and optimising the Level-2 pattern recognition to ensure good quality track reconstruction in this more complex environment, possibly benefiting from track parameter information calculated prior to the HLT if FTK[1)] is adopted. In particular, care must be taken to minimise the number of fake tracks produced by the increased background, while preserving high efficiency. The tracking at the Event Filter is based on offline reconstruction tools. The performance of these tools must be evaluated for the trigger and, where needed, specific configurations developed which are optimized for the Event Filter.

The scaling of execution time with occupancy will be measured, for both Level-2 and the Event Filter, in order to predict the CPU requirement as a function of luminosity. Various optimisations will be studied, such as raising the minimum p_{T} for tracking, minimising RoI sizes, and tightening cuts for reconstructed features. Tracking will be studied, including the effects of pileup, radiation damage and the addition of a new pixel layer (the Insertable B-

[1)] The FTK is a pre-processor running before Level-2, which would process SCT and Pixel data to produce track parameters approaching Level-2 quality in less than about 1 ms. This project is currently preparing a technical design report prior to requesting ATLAS approval to move to a technical proposal (no UK resources are being requested for the FTK project itself.)

Layer) between the beam-pipe and the current inner layer of the Inner Detector.

These performance studies and optimisations will cover the full range of RoI types (electron/photon, muon, tau, B-jet, B-physics), which all have different characteristics and make different requirements on the tracking. The studies will use high-luminosity data simulated by Monte Carlo, and also emulated by overlaying real collision events. As part of the overall speed optimisation of the code, the possibility of performing some data preparation (clustering and association of stereo information) at the level of the SCT Readout Drivers will be investigated in conjunction with the UK SCT experts. The exploitation of Graphical Processing Units (GPUs) will be studied for the Trigger, following the developments for the offline code described in Section 4.

3.5.3.1 Selection Software and Menus

The developments for Phase-I are an extension of ongoing work to adapt menus and selections as luminosity increases. The work will focus on the electron and muon triggers, where the UK has key expertise. Predicted trigger rates will be determined as a function of luminosity, including output rate and the rates at each stage of selection. Optimisations will be made to the selection cuts and to the ordering of steps within the trigger chains to achieve the earliest possible rejection, and so minimise CPU resources. This includes the development and improvement of optimisation tools, and studies of new techniques such as the use of neural networks.

The inclusion of the FTK processor would allow greater flexibility in the ordering of the steps in the HLT selection chains, giving the possibility to perform track-based selections (including B-jet tagging, primary and secondary vertex reconstruction, and invariant mass reconstruction) from the start of the HLT selection. Where necessary to achieve the required rate reduction, new selections will be developed including increased use of multi-object and topological triggers (for example geometrical cuts, or selections based on the reconstructed invariant mass of intermediate-state particles).

Changes are needed to the HLT steering software to support the option of full event reconstruction, to enable a comparison of this with RoI-based processing at high luminosity. Extensions to the steering software will also be required for input from the FTK, if approved, and to allow for the possibility of running both Level-2 and Event Filter algorithms in the same processor. Ways to adapt the code to make optimum use of future multi-core processors are being investigated as part of WP9, and these will be evaluated for the HLT.

3.5.3.2 Dataflow and farms

There will be an evolution in the hardware (via the rolling replacement programme) and software used in the dataflow and farms, but no major architectural changes are planned. The rolling replacement programme of hardware is foreseen for the long-term viability of the system and will lead naturally to higher performance components. However, some effort will be required to provide the optimisations, including some reconfiguration of the hardware, to reach the system performance required for Phase-I. For example, higher data rates will be required from the readout System (ROS); we expect this to be provided by higher performance in the host PCs and possibly the use of the direct network output port on the ROB-in cards. (This latter option, the so-called switch-based ROS, is currently being studied for use with the initial full luminosity.) Some effort will be needed to tune the system to the higher performance, to

balance the resources to match the needs of the different detectors and to provide the additional readout channels required for Phase-I (e.g. for the Insertable B-Layer). We also foresee that the rolling replacement programme will provide for the higher performance network and farm components. Here again some tuning and reconfiguration of the hardware will probably be required, to match the very different requirements across the system.

Similarly, it is foreseen that there will be a steady evolution in the dataflow software; indeed studies have already started into a variant of the current software with the possibility to combine the Level-2 and Event Filter processing into the same farms. Whether or not this is adopted, effort will be required to continue the evolution beyond the needs of the full initial luminosity.

3.5.4 Phase-II

A significant upgrade of the HLT system, both hardware and software, will be required for Phase-II. Various options must be studied in detail and compared in order to enable key architectural choices to be made. This will enable the work on the detailed design and implementation of the system to start soon after the time-frame of this bid.

The specification of the trigger hardware and dataflow architecture is intimately linked to the chosen selection strategy, and the design of the selection software depends on these choices. ATLAS-UK HLT has greatly profited in the past from the spread of expertise across the areas of hardware and event selection software, and this was instrumental in the successful completion of the HLT construction project. The UK played a key role in the design, implementation and commissioning of the selection software, elements of the Readout System hardware and software, and specification and procurement of the HLT processor farms and TDAQ network infrastructure. It is important to maintain this range of expertise for the upgrade work.

3.5.4.1 Selection Software and Menus

Substantial changes will be needed both to the HLT framework software and to the selection software itself. In the upgraded trigger, some of the steps currently performed in the HLT may be performed at Level-1, such as the use of finer granularity calorimeter information, the inclusion of Inner Detector information, and the use of topological information, i.e. associating information from different detectors and different points in space. It will probably be possible to have a higher output rate from the HLT to storage (and possibly a higher input rate from Level-1) and it is likely that there will no longer be a sharp divide between Level-2 and Event Filter in terms of data access. As a result, the HLT selection strategy will need to be redesigned to give the optimum performance in the upgraded system, including re-assessing the relative performance of RoI-based and full-event processing. In its current implementation, HLT reconstruction is performed in Regions of Interest. This represents a saving of execution time when average RoI multiplicities are low, but may not be the optimum strategy at very high luminosities.

Performance studies will be made using simulated data. The various options will be assessed using benchmark physics channels for efficiency measurements, and background samples to determine CPU times and measure trigger rates. An optimisation of the use of tracking information in the trigger, including a comparison of the different options for use of pre-HLT track information, will be a particular focus for the UK groups. Another important area of work,

in conjunction with the UK Level-1 effort, is an overall optimisation of electron-based triggers, including topological triggers and the information available from the upgraded Level-1.

3.5.4.2 Dataflow and Farms

A variety of trigger architectures will be studied, ranging from an upgrade of the existing architecture to a more substantial redesign. If the architecture of the TDAQ-detector interface is retained, with point-to-point data links between detector front-ends and the TDAQ Readout Buffers, it will be necessary to start development of a new high-bandwidth data transfer link exceeding the current bandwidth limitation of 160 Mbytes/s. A detector modelling exercise will also be needed to evaluate the characteristics of detector event data at higher luminosity, and decide how to adapt the granularity and arrangement of the detector to TDAQ connections. The present segmentation was chosen to equalise the bandwidth per readout link, and the same principle should guide the definition of a new layout.

At the same time, alternative architectural approaches will be developed and assessed in discussion with the detector communities. The current Readout Driver (ROD) implementation is output limited, as it acts as concentrator (many input links and one output link); there are suggestions that merging functions of the ROD with those of the Readout Buffer (ROB) and ROS may help deal with the increase in detector event bandwidth, either by providing more latency for improved Level-1 data selection (if possible), or by allowing an RoI-like selective readout approach directly to the detector front-end boards. A study in this direction should also be driven by the same principle of a uniform, detector-independent, design of the TDAQ system and of each of its elements.

The other main line of investigation has to deal with the probable increase in data volume transferred from the ROB to the Level-2 Trigger. Here the main studies are to follow technology developments in the areas of PC and network hardware to implement a ROS system with a significantly larger bandwidth capability towards the Level-2 Processing Units, and to model possible changes in the ROS-to-HLT connection scheme, both in terms of hardware topology and data transfer protocols.

Key input comes from software-based modelling of the dataflow within the various architecture options. This requires input from measurements made with the selection software, both running online on the existing system and running on high-luminosity data simulated for the upgraded detector. These measurements are strongly dependent on the selection strategy and menus, and this work requires close interaction with the groups working in those areas

Chapter 4

Software, Simulation and Computing

An essential component of the upgrade project is the provision of simulation and reconstruction tools to allow the study of the detector and physics performance. It is important that this is closely integrated with the hardware R&D and build projects. At the lowest level, this requires software descriptions of the proposed detectors/triggers. Work is proposed here on the barrel and forward wheel tracking simulation, and the adaptation of the tracking code where required. The UK has a central role in this work from the current ATLAS UK upgrade tracking project.

Another important component to the upgrade studies is the correct description of the radiation environment. The understanding will progress hugely with real collisions, and will also depend on the other detectors, such as the position of the forward calorimeter. Sheffield intends to continue its lead in this area, established in the existing ATLAS UK tracking upgrade project.

The upgrades present new and distinct challenges for ATLAS software and computing. Concurrent with the operation of the current detector, preparatory work is required to investigate high luminosity running, requiring the development of simulation and reconstruction codes. By a luminosity of 3×10^{34} cm^{-2}s^{-1} this requires special techniques and arrangements, as the memory profile for that degree of pile-up of events far exceeds the memory per core of available resources (Fig. 4.1). At present, and only after great effort already made to minimise the memory profile of the events, the only approach is to run in a highly inefficient mode, effectively wasting most of the available CPU. Much of the early work of the project under Phase-I represents attempts to adapt the ATLAS software to use the likely existing resources more effectively to make practical the high luminosity studies; there is much established expertise in Edinburgh and Oxford.

By Phase-II, another and more profound issue must be addressed. Experimental HEP data are organised in 'Events' each corresponding to one particle collision. Simulation, Reconstruction and Analysis programs all process one event at a time. Events are independent of each other and therefore trivial to process in parallel. In this respect Experimental HEP has cost-effectively exploited high data-throughput computing on $\sim$100k cores. However, the development in CPU technology is towards very many cores per processor. While the work in Phase-I will give some mitigation of the effect, simply farming events to each core will result in a huge memory profile in the processor, even for no-pile-up and reconstruction. In addition, there will be a problem with the input/output bandwidth into the processor. This problem may actually be even worse for the usually simpler case of analysis, as this is a particularly I/O bound activity. Two possible broad strategies present themselves: to trust that adapting the code and the kernel to allow multi-

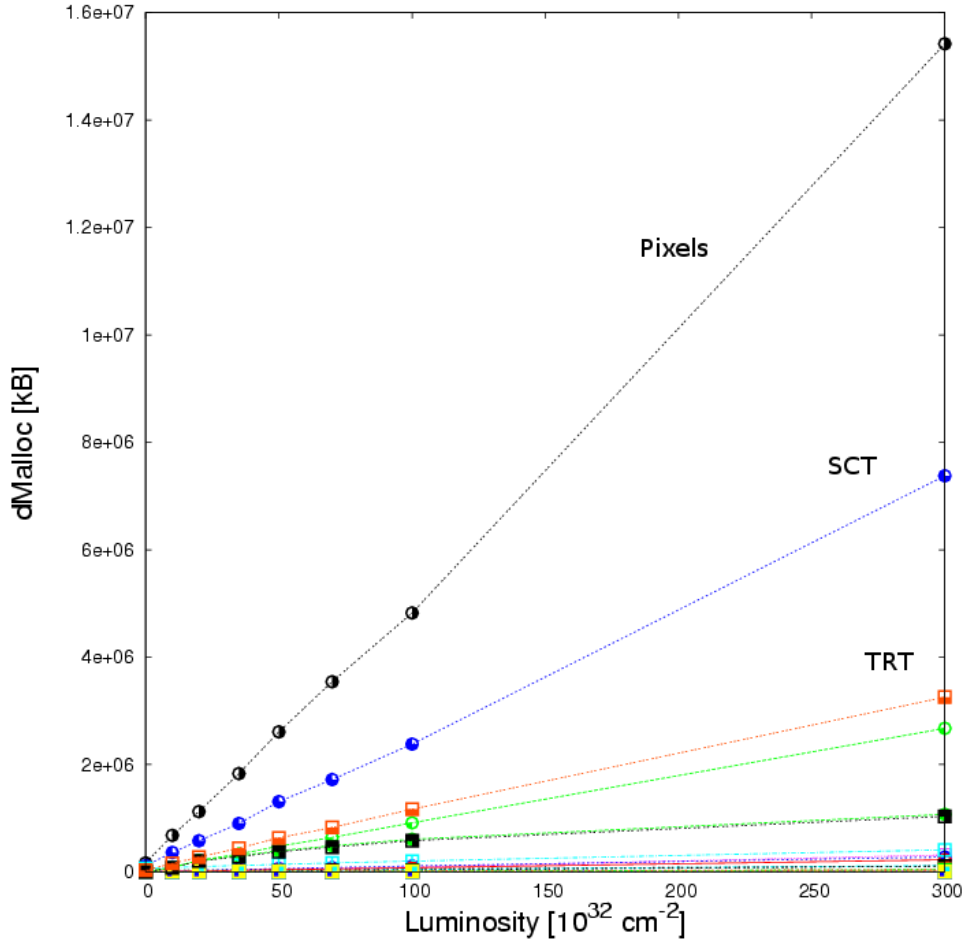

Figure 4.1: Digitisation event loop memory allocation. Pixels and SCT are already consuming 15 GB and 7 GB of memory respectively at 3×10^{34}.

core chips to be used more effectively will deliver the required performance; or to turn instead to parallel code, especially Graphical Processing Units (GPUs), which are already designed for high throughput.

Both strategies need investigation. It should be stressed that there are indeed several use cases, e.g., simulation, reconstruction, analysis (event selection and parameter fitting), which are quite distinct. There is obviously overlap with problems faced for the high level trigger with those in the reconstruction, but the problems in simulation and analysis are quite distinct. There will therefore be a Phase-I process of R&D and technology evaluation, followed by a Phase-II programme of implementation of the selected technology for the appropriate use case. For both tracks, close contacts will be made with technology providers. We have already made contact with NVIDIA developers, who are keen to help with the work on GPUs and may provide sponsorship. The proposed work will benefit greatly from such contacts, and those with EPCC in Edinburgh and Computing in Lancaster; it will depend on the expertise in the Edinburgh (parallel computing, code optimisation, data formats and tracking simulation) and Lancaster groups (tracking and reconstruction, ATLAS Grid environment, virtualisation, project management). The work will inform and require close collaboration with the High Level Trigger activity, which represents a distinct and important use case for these strategies.

4.1 Phase-I Programmes

4.1.1 Initial tools for Large Pile-up Studies

By the conclusion of the current UK ATLAS tracking upgrade project, the Inner Detector simulation will have reached a state where it can implement layouts with simplified support structures and service volumes, simulate physics events and overlay them with pileup derived from minimum bias events, and identify and reconstruct charged tracks using state-of-the-art ATLAS tracking algorithms. Several layout studies will have been performed in order to investigate the effect on performance of certain classes of layout modifications, such as adding or removing layers of particular granularity. The level of pileup will have been investigated beyond 400, in accordance with leading machine scenarios. Layout options will have been narrowed down to several general categories and a baseline layout design for the LoI will have been studied extensively.

It is vital that layout investigations continue, narrowing the options down to one design, especially as new information becomes available from latest detailed engineering, LHC machine scenarios, and initial LHC data. The process will proceed in stages, with a few major options in early 2010, and one soon afterwards. It may be expected, based upon the experience of the construction of the current ID, that the endcap design and the barrel-endcap transition, along with services and support structures, will require sustained design analysis. There will be many concurrent sources of information, including improvements of the existing detector simulation in the light of real data experience and so the tools developed for the next few years will have to retain great flexibility.

Given the UK's strong interest in Pixel technology it will be important to focus additional effort towards simulations which investigate and compare the various emerging hardware design and layouts. In particular, the initial design layout of the forward pixel disks is looking very complex with many different device radii. This area is also particularly challenging due to the

competing requirements for many services to come out in this direction. Careful design studies, using the recently implemented volumes, are key to the success of this project. It is critical to make the correct decisions, and to estimate correctly the material budget, which argues for the close integration of the design, construction and simulation teams.

A particular problem to be addressed by the simulation is to investigate the robustness of proposed layouts to the reality of detector construction and operation. It has been shown, for instance, that a tracker with 4 pixel layers and 5 doubled strip layers can maintain tracking efficiency with an acceptable fake rate up to 400 pileup by requiring at least 11 hits out of a nominal potential 14. A random hit inefficiency of up to 5% has been introduced in these simulations in order to simulate the effect of ageing. It is less clear whether the layout can perform effective pattern recognition if a small region, or even a whole stave is lost, due to cooling or powering issues.

Another area that remains to be investigated more thoroughly is the behaviour of the z vertex resolution along the beam line. With 400 pile-up events spread over a luminous region of less than ≈ 30 cm, the separation between vertices is not large and could lead to significant overlap between pileup and hard processes. Such overlap could pollute important quantities such as missing ET and thus diminish the detector's sensitivity to physics beyond the Standard Model. At present, the smallest pixels to have been considered have a z length of $250\,\mu m$. To achieve maximum resolution, Pixel tiling, turning some sensors 90 degrees, or 3D pixel sensors in the innermost layer, may need to be brought into play in order to increase the ATLAS physics capability.

It is also essential for the ATLAS upgrade that combined performance studies be performed with the Inner Detector along with the calorimeters and muon systems. Even though those systems do not envisage wholesale replacement, as in the ID case, the ID material will affect their ability to identify and analyse signatures at both trigger and offline levels. The ID simulation developers will collaborate actively with experts in these other areas in order to produce a flexible and adaptable simulation, ready to evaluate the performance of the whole ATLAS detector in sLHC conditions.

If it is found that ground-breaking physics is impossible to separate from background in the presence of 400 pileup, other LHC machine scenarios such as luminosity levelling must be considered. Such solutions, as currently understood, could reduce the pileup to a relatively constant 75 events, as opposed to an exponentially falling distribution from the peak at 400. These studies thus form a critical input to LHC as well as ATLAS upgrade decisions.

Even when efficiencies and resolutions for single particles (muon, electrons etc.) are known from full combined performance simulation, it is not practical to fully simulate all physics channels. Parametrisation of the performance of ATLAS is a key method which is already implemented in the fast simulation (Atlfast). Atlfast is a UK led project, and UK developers will need to drive the adaptation of Atlfast to the requirements of upgrade studies.

Tracking for isolated tracks and tracks within jets have been, and will continue to be, investigated. Over time, these studies will be refined with real data. Early data will be able to probe actual detector response, the character of LHC backgrounds, and, to some extent, minimum bias and Standard Model backgrounds. Clearly, the earlier we can achieve 14 TeV collisions, the more reliable the extrapolation that can be performed to sLHC conditions.

4.1.2 Radiation simulation verification

The unprecedented levels of radiation background expected at the LHC have had a major impact on the design of the machine and experiments. The backgrounds will degrade detector and trigger performance, damage detectors and electronics, induce digital upsets and data corruption, and give rise to radio-activation which impacts experiment access and maintenance scenarios. The situation at the sLHC will be even more challenging. Given the significant resource implications of upgrading ATLAS, it is crucial that radiation background issues are assessed well in advance.

The Sheffield ATLAS group has played a leading role in understanding and solving radiation problems for the ATLAS experiment at the LHC [32–34], and is the only ATLAS group with the relevant expertise and experience in understanding the complex radiation backgrounds expected in the upgraded ATLAS detector. This requires an understanding of particles interacting with matter from TeV energies down to thermal energies for neutrons, and experience of running Monte Carlo particle transport codes such as FLUKA [35]. Until 2005, Dr. Ian Dawson was a principal member of the Radiation Task Force, set up by ATLAS Management to investigate all aspects of the expected radiation environment in and around the experiment. A summary of much of this work can be found in the Radiation Task Force Summary Document [34]. Dawson is currently a member of the ATLAS High Luminosity Upgrade Steering Group (HLUSG), and editor for the radiation background chapter of the ATLAS Upgrade Letter of Intent.

ATLAS management has stressed the importance of cross-checking the predictions at the LHC with real data measurements, so that the situation at the sLHC can be better assessed. Unfortunately expertise to perform these comparisons is severely limited. The Sheffield ATLAS group took on the responsibility to analyse and compare measurements of 1 MeV silicon damage fluence, ionising-dose and thermal-neutron fluence made with radiation monitors in the inner detector volume. This work is being performed by Dr. Ludovic Nicolas, who is supported until the end of 2010 by the UK ATLAS Tracker Upgrade R&D project, with some travel support from a successful EU-sLHC preparatory phase bid [36]. The monitors to be used for this work are sufficiently sensitive to make useful measurements before nominal luminosity is achieved. It has now been established that first collision data will occur at the lower centre of mass of 7 TeV, so new FLUKA predictions are being obtained to allow comparison at these energies. A recent summary of this work can be found in Ref [37].

Beyond the ATLAS inner detector it is not clear how the comparison of measurements with predictions will be performed, due principally to the lack of expertise outside the UK. In some cases this is critical. For example, the radiation background uncertainties in the ATLAS cavern are so large that the level of muon system upgrade cannot be determined until radiation background measurements have been made. Therefore requests have been made to ATLAS UK to extend their responsibility to beyond the inner detector regions.

We are proposing to maintain the UK's expertise and leadership in the very important area of comparison of measurements with simulated predictions. Nicolas supervised by Dawson will accomplish the benchmarking of ATLAS radiation simulations by analysing data from radiation monitors installed throughout the inner detector and other ATLAS sub-systems. This work will directly improve radiation models for the proposed new inner tracker and enable detailed optimisation of tracker design (see Section 4.2.3). Nicolas already leads in this area.

4.1.3 Visualisation

The event display (ATLANTIS) was a critical tool for commissioning the day-one ATLAS detector and continues to be so for the current detector. It is anticipated that visualisation will also be vital for upgrade studies and the understanding and evaluation of different detectors layouts and beam conditions. ATLANTIS is a UK led project and M&O responsibility. A modest amount of effort from existing ATLANTIS developers will be required to ensure it can cope with the requirements of upgrade studies.

4.1.4 Near-term computing upgrades for sLHC Phase-I

Given the large cross-section for inelastic, non single-diffractive, pp interactions we expect for the upgrades to have up to 400 simultaneous overlapping events piled up per beam crossing (varying even higher due to Poisson statistics). ATLAS computing is already at the performance limit (using the current design) of the available computing architectures. Therefore, planning for the upgraded LHC environment requires a number of major design changes and is a serious computing challenge. We present here some ideas, and the R&D, which are necessary for the upgrade transition to happen successfully.

The ATLAS computing currently taxes all aspects of the so-called "performance triangle" (Fig. 4.2). For full ATLAS detector simulation CPU time is the major bottleneck, for digitisation (under pileup) both CPU and memory are major issues. For reconstruction memory is the current limiting factor and for Athena analysis jobs I/O is the problem. To meet the ATLAS upgrade computing requirements, all three crucial facets must be tackled: CPU, Memory and I/O, which we now consider in turn.

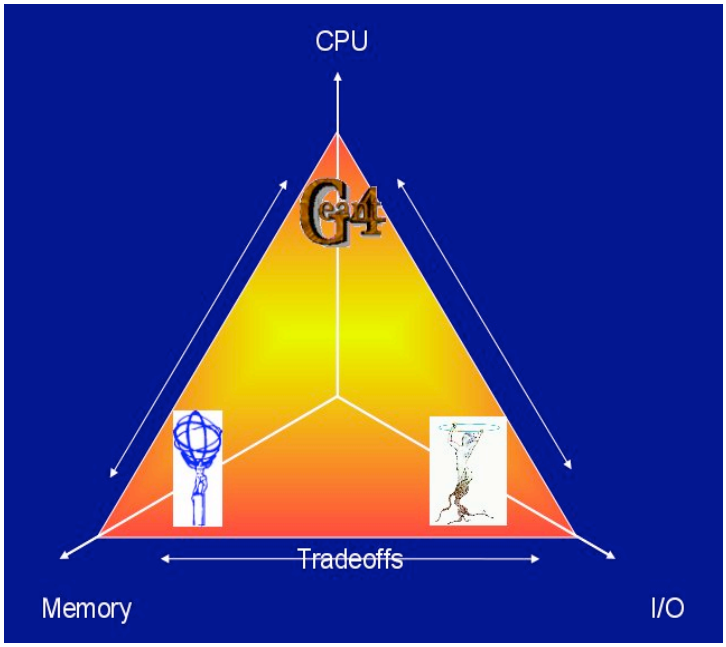

Figure 4.2: The ATLAS Computing Performance Triangle.

4.1.4.1 The CPU evolution from multi-core to many-core processors

The only way to maximise our CPU effectiveness in ATLAS is to better harness the CPU power available on the x86_64 architectures that we will likely be running on. Multi-core processors, which contain two or more independent cores/CPUs on a single integrated circuit (or multiple dies in a single chip) are already prevalent and we are now moving to the many-core paradigm, where traditional techniques (e.g. serial event processing in many jobs), which have worked for us so far, will be less effective. If we look at the recent processor evolution the problem becomes clear. If we define the number of dispatch job-slots on a typical desktop (or Grid worker node) to be the no. of sockets $\times$ no. of cores $\times$ no. of simultaneous multi-threads (e.g. Intel hyper-threads) then we can capture the recent processor evolution:

Processor	Dispatch slots		
Dual-socket Intel quad core Harpertown:	2*4*1	=	8
Dual-socket Intel quad core Nehalem:	2*4*2	=	16
8-socket AMD 6-core Istanbul Opteron:	8*6*1	=	48
Quad-socket Intel Nehalem 8-core:	4*8*2	=	64
8-socket Intel Nehalem 8-core:	8*8*2	=	128
Quad-socket SUN Niagara (T2+):	4*8*8	=	256
AMD MagnyCours Opteron (12-core):	x*12*y	=	?

Table 4.1: Recent processor evolution towards many-core paradigm.

Therefore, we are soon to have hundreds of dispatch slots (perhaps thousands by 2018) on each machine. This increase has already started to drive the number of dispatch slots well above the level for which the current ATLAS software has been designed, i.e., serial processing of events within independent jobs.

It will affect all areas of ATLAS upgrade computing and it is important that the UK programme is on top of the new processor developments. We need to plan ahead to avoid hurried migrations and to mitigate the semi-imminent performance bottleneck. The many-core paradigm must be utilised effectively, otherwise we will be using the wrong software design for the hardware environment we are running.

One of the major steps will be to develop the Athena framework to be multi-threaded so we can run events and parts of events in parallel. Of course gains will be limited by the software fraction that can be run in parallel. We are fortunate here because events are inherently parallel, and subsystems could possibly be processed in separate threads, or another approach is to segregate different types of pattern recognition (tracking, clustering etc.), or even individual tracking algorithms.

No matter what the approach, the limit of performance gain is given by Amdahl's Law:

$$S = \frac{1}{(1-p) + p/n}$$

where S is the speed-up factor of the code, p is the fraction of code that can be made parallel and n is the number of parallel processes. It is therefore important to keep the fraction of non-

parallel code to a minimum, perhaps even making parallel the start up and initialisation of the event loop for much faster turn-around of jobs.

As general software moves evermore towards highly multi-threaded systems, e.g., the OS itself, virtual machine hypervisors, and Java virtual machine implementations. Then it is natural that the processor board and memory architectures start to cope with this. However, if the ATLAS computing is still running serial monolithic jobs, then we will not be able to exploit the future computing environment effectively.

4.1.4.2 The memory wall

With the advent of multi-core machines the amount of memory per dispatch slot has remained relatively static at 2GB per core. However, pressures such as those to increase the processing time for ATLAS upgrade simulation, reconstruction and analysis drive the Athena software to become more memory intensive. Even without this steep increase in software memory profile, the underlying event complexity at the LHC will increase linearly with luminosity, regardless of the software we use. This, coupled with the already-large memory profile poses another separate, but equally serious, challenge. While the speed of memory access has increased exponentially its exponent is quite small and has not matched the much more rapid improvement of the CPU. There is a performance divide, which is partially mitigated by various levels of on-die caching. As the number of cores and thus simultaneous jobs/threads increases this memory management becomes very complex and traditional SMP methods break down. Typically, SMP architecture is used to handle memory usage in multi-core systems by allowing for two or more CPUs to access to the same shared memory bus. However, this does not scale well for newer systems with tens (and shortly hundreds) of cores.

Memory bottlenecks at this level may alleviated by taking advantage of NUMA (Non Uniform Memory Access) architecture. Here the memory is organised in number of distinct "nodes" to which CPUs are attached to depending on their locality (e.g. if the CPU is on the same bus as the memory). Tasks are then given a node affinity to allow processes to run on the same CPU as the local memory, giving the benefit of much faster access to the cache. In addition, the full memory space is still cache coherent and accessible to all processes albeit with a higher latency depending on the speed of interconnection between the memory regions.

Performance improvements depend primarily on the access pattern of the application. It is therefore vital that ATLAS task groups and applications are profiled and tuned correctly otherwise relative performance can be reduced, due to higher latency data sharing between remote nodes. If a system is NUMA capable then both a command line interface (numactl) and shared library (libnuma) are available to configure the expected memory usage. The former enables tasks to be attached to a specific CPU either as a preference or as part of a stronger binding policy. The latter allows access to an API that can be linked in with applications to enable fine-grained policies to be defined by the developer.

If an application or process group can be optimised for this architecture then this could yield a significant performance improvement over traditional SMP use; potentially equivalent to the number of cores (or memory regions) in the system. An investigation into the possible benefit of optimising the current software for multi-core NUMA systems is therefore very important.

Currently ATLAS is exploring the forking of events as a way forward (Athena MP), where jobs are forked simultaneously using Linux `fork()`. They then share common data, e.g., the

detector description saving around 600MB of memory this is done by then letting the OS employ "Copy-on-Write" memory algorithms. This has been shown to reduce the memory consumption, but there are some deployment issues, e.g., merging the events afterwards and the scheduling model.

Another planned improvement is the use of a new RedHat kernel module (KSM) which can do dynamic sharing of kernel memory pages. If many jobs are submitted in batch then a number of identical pages are naturally created. KSM then sorts pages by raw page content and if the memory is identical then it is shared. It was developed to aid the KVM hypervisor and initial tests have shown large amounts of real memory gain. However, given the steep profile of memory usage within ATLAS and rising event complexity, this can only go so far. Whilst it is worthwhile deploying it will not address the problem of running 100 jobs on a single machine. However, it is possible with the rise in virtualisation that several such technologies may emerge; their degree of success with ATLAS applications remains to be seen, but it is likely that it will still require application-specific work to achieve good performance.

4.1.4.3 The input/output challenge for ATLAS

ATLAS upgrade analyses running over much larger datasets on many-core processors will increase the I/O overheads already described and limit the possibility of reaching optimum computation efficiencies given by Amdahl's law. There are many potential bottlenecks, and the challenge is to identify and address them, and to then ensure that I/O performance keeps pace with networking and processing improvement. Solutions will, of course, depend on the particular disk access patterns, which will in turn depend on the computing model and data format employed. A recurrent source of pressure arises from the typical end-user pattern of (effectively) random access of large data files as well as of particular 'hot' files that are accessed by multiple competing processes. Approaches to help deal with such challenges include improvements to the hardware, methods of file access, and the underlying file system. These are described in some more detail below. Given the extent of the challenges, a successful solution will have to incorporate elements of each of these.

Hardware improvements offering better access times than current SATA systems include Solid State Devices (SSDs) and alternative hard drive solutions such as Fibre Channel (FC) or Serial Attached SCSI (SAS). The current cost of these solutions is decreasing, but nonetheless the best performing technologies would need to be reserved for frequently accessed hot files or portions of the data file. Ways this may be achieved include changes in methods of file access, as discussed below, or using tiered storage solutions that automatically assign or move data from drives of different formats according to their usage. Prototype high-performance storage solutions will be needed for optimisation studies.

In terms of file access, the variety of protocols currently in use can be configured to buffer a portion of the file either in the RAM or local disk of the processing machine. This latter option can help reduce pressure on ATLAS storage file systems or network bandwidth, but it can cause significant issues for local disk access on multi-core nodes as multiple processes competing for these resources. Thus, one should also explore alternative file systems for the data servers themselves or Direct Memory Access (DMA) methods could also be utilised to bypass CPU and disk utilisation when copying data between nodes.

As I/O limitations arise primarily from access by multiple processes to a disk, this limitation can be mitigated by spreading the load through implementation of a cluster file system such as

GPFS, Lustre, or Hadoop (HDFS), which are commonly used in the HPC and Cloud computing environments. Fully testing these in a HEP environment, as well as evaluating the other configurations discussed above, will be an important part of dealing with the computing requirements of the LHC upgrade.

4.2 Phase-II Programmes

There are several important tasks that must be performed in the next years to plan the upgrade activity in Phase-II. The most obvious of these are physics performance studies to guide the upgrade detector design. These in turn require the simulation and reconstruction code to be able to handle the very high pile-up rates expected in Phase-II. Similarly, the radiation environment must be predicted, based on the studies from the existing data and from tuned Monte Carlos. The planning must be in place for the very large computing requirements to handle the growing data sets. The projections for the required resources indicate that by the middle of the decade, important choices will have to be made about the technological path for offline computing for each of several use cases (simulation, reconstruction, analysis; and in WP6, in the trigger). The work must be done in the next few years to evaluate the best approach for the various use cases in order to begin the significant work of implementing the computing for Phase-II. These activities are described in more detail in the following sections.

4.2.1 Physics Performance Studies

The UK represents only 10% of the ATLAS membership, but already has a standing in physics that exceeds this level. As well as a recent physics co-ordinator (now a deputy spokesperson), the UK currently provides convenors for four of the eight physics groups. This leadership will continue into the upgrade physics studies, notably in the area of the influence of trigger decisions on physics outcomes and in studies of rare and exotic decays. The studies will require large samples of simulated data to guide both the physics case and the upgrade detector decisions.

The UK also plays a major role in the existing ATLAS Grid computing & productions, and will need to contribute similarly to the production of samples for the upgrade studies. The requirements for upgrade productions will be discussed below; the samples to be produced will include simulations of selected physics channels for example physics studies that will test the potential of the renovated detector.

As the actual physics content of 14 TeV collisions remains as yet unknown, the simulation program will use several stand-ins as performance benchmarks. The primary set includes isolated electrons and muons, Standard Model $t\bar{t}$ production, $Z' \rightarrow b\bar{b}$ events and $B_s \rightarrow \mu\mu$ and $\tau \rightarrow \mu\mu\mu$ rare decays, as well as minimum bias events. As an example, in addition to a suite of high-pT physics studies, the $B_s \rightarrow \mu\mu$ decay, which is rare in the Standard Model, is sensitive to many new physics models. Given suitable triggers, ATLAS is very well suited to this study. It therefore provides an interesting test in Monte Carlo studies for multi-muon triggers. A related study is generally regarded as an exotic, the lepton flavour violating decay $\tau \rightarrow \mu\mu\mu$; the production of τ leptons will be plentiful, and the multi-muon trigger will be needed; the boost of the τ will mean the muons would tend to be closely clustered, and hence the triggering opportunities and hence physics potential is dependent on the capabilities of the the Level-1 muon trigger system in the upgrade. Naturally, when it comes to real data, standard candles

such as the J/ψ will form the basis of performance studies. These latter studies come from the B-physics group, which is led by the UK.

4.2.2 Simulation of Large Pile-up

As noted previously, consolidation of a final design will require a number of detailed benchmark and physics-based studies grounded in a simulation of increasing realism and taking into account what has been learned from early ATLAS data from pp and heavy ion collisions.

Thus, Phase-II in the detector & trigger simulation program forms a continuation of the Phase-I activity. At the same time, the simulation must evolve from a prototyping tool, with simplified geometries and models, to one with increasing detail as better understanding is gained of engineering and machine scenarios. Indeed, at some point the simulation must become part of the standard toolkit for ATLAS analysis in the sLHC-era computing model. At the very least, this transition will entail several major infrastructure changes presently specific to upgrade simulation itself to be propagated throughout the ATLAS software suite. An example already underway is the requirement to use 64-bit rather than 32-bit channel identifiers.

Until that time, the upgrade simulation must be kept up to date with special effort outside the normal maintenance of current ATLAS software infrastructure. Coding standards have simplified some aspects of adapting the software infrastructure for upgrade use, but enforcement of such standards such as not hard-coding the number of layers or channels expected in software requires the development and maintenance of special upgrade-related run-time tests, as well as constant monitoring for any number of unphysical results arising from these analyses.

It is expected that as the standard simulation become more detailed, it will consume more CPU and memory resources; indeed, it is already clear that upgrade simulations push against the limits of computing systems which are currently widely available. These limitations, and potential solutions, will be discussed in subsequent sections. For the full, Geant4-based simulation, there is considerable scope for optimising CPU and memory usage on individual computers, and the simulation can be arranged to use multi-core and more novel architectures more effectively. In addition, more effective use of distributed computing technologies for simulation should be investigated, keeping in mind CPU, memory, and network bandwidth limitations of the systems involved.

In the meantime, it will be necessary to generate large samples for physics studies with the upgraded detector. For this purpose, a simplified simulation will be needed, a development of the Phase-I Atlfast. The standard tools at present incorporate hit smearing as well as models of electromagnetic and hadronic secondaries. However, while the simplified hit-level simulation increasingly matches the occupancies of the Geant4-based simulation, the computing burden of pattern recognition remains largely the same; at present, simulation and reconstruction take similar times. There is substantial room to improve the performance of the standard pattern recognition tools in the expected sLHC environment, much as it has been tuned for pp as well as heavy ion collisions. Even so, for studies geared towards exploiting 3000 fb^{-1}integrated luminosity, there will be a need for Atlfast's parametric simulations, building on the work already pioneered by UK researchers.

4.2.3 Radiation Environment Projections

The investigation of radiation background issues for ATLAS at the sLHC has been dominated by the efforts of the Sheffield ATLAS group, both in terms of leadership [38] as well as publications [39]. Simulations to date have focused on inner tracker issues, as defined by work package WP2 of the UK ATLAS Tracker Upgrade R&D project, and studies have included:

- A new design of the poly-moderator shielding used for reducing silicon damage fluences. This was partly motivated by the loss of the TRT in an upgraded inner tracker, which contains a high percentage of hydrogenous material in its construction which acts as an effective moderator for the silicon detectors. We investigated the impact of different moderator thicknesses in different regions, and liaised with the upgrade Project Office to ensure engineering constraints were respected. A full description of these studies can be found in Ref [40].

- An investigation into the impact of positioning machine magnet elements close to the interaction point. In order to achieve an order of magnitude increase in the LHC luminosity, machine design scenarios have been proposed which require machine magnets positioned inside the ATLAS experiment. It is important to assess the implications of such machine upgrade scenarios on ATLAS radiation backgrounds. A full description of these studies can be found in [41].

- Simulations for the Phase I insertable B-Layer upgrade [42], which looked at fluences and doses close to the interaction point in much greater detail and precision than previously performed.

There are two components for achieving final predictions for ATLAS radiation backgrounds at the sLHC. First measurements made at the LHC will be validated against the predictions, as described in Sec. 4.1.2. This is necessary to determine the level of confidence in the simulation codes. Second, simulations need to be performed iteratively with the subsystems to achieve a final ATLAS upgrade layout. It is important that the current expertise is maintained to allow successful completion of these tasks.

ATLAS management have also requested that the Sheffield ATLAS group extend their unique radiation background expertise to go beyond the inner tracker. Examples of issues not currently being addressed by ATLAS, but considered critical for its successful operation include:

- Dose rates from activation. This is already a major issue at the LHC, and at the sLHC it will have serious consequences on any plans for intervention. This is particularly important for the inner tracker, which sits close to the beam-line where the dose rates are highest. Strategies will have to be defined to allow intervention which will involve studies to investigate feasibility of various access scenarios, and whether or not remote-handling will be necessary.

- Fluences and doses in the cavern have very large uncertainties which impact critically on the design of the upgraded muon spectrometer.

- Radiation backgrounds in the adjacent cavern USA15, which hosts most of the off-detector electronics, are likely to be too high for standard access to be granted. This has serious implications on the day to day running of the experiment and will need to be addressed. One possibility is to improve the shielding capability of the shielding wall that divides UX15 and USA15, which would require a dedicated simulation study.

We propose to maintain our vital UK expertise and leadership in dealing with the serious and unprecedented radiation background issues expected for ATLAS. In addition to the benchmarking responsibilities outlined in Section 4.1.2, we propose that Nicolas, supervised by Dawson will:

- Improve the new inner tracker radiation model using benchmarked simulations and use them to further optimise tracker design. The new inner tracker radiation field is already significantly better understood than that in other ATLAS sub-systems through the work of Nicolas. This understanding will need to be improved further using the benchmarking results outlined in Section 4.1.2. This will lead to further optimisation of the tracker layout, services and neutron moderator design. It will also be necessary to assess the impact on the new tracker of major changes to the design of upstream detector and machine elements, including the FCAL and final-focus magnets.

- Maintain responsibility for development of radiological access scenarios. The possibility of large dose rates from material activation is already a major issue at the LHC and at the sLHC will impact strongly on intervention plans. This is particularly important for the new inner tracker, which sits close to the beam-line where the likely dose rates are highest. Intervention strategies will be defined using the benchmarked radiation models from the benchmarking tasks discussed in Section 4.1.2, considering questions of maximum allowed access time and need for remote-handling in particularly active regions. It should be noted that Sheffield is the only ATLAS group with any expertise in this area.

It should be stressed that the LHC/sLHC experiments are entering radiation background regimes unprecedented in high energy collider physics, in terms of both the detector impact and radiological consequences to personnel. At present, only UK ATLAS has the skills and expertise to deal with both these issues.

4.2.4 sLHC Computing Requirements

The existing resource projections from ATLAS include only the current ATLAS experiment with a 200 Hz trigger output and a maximum luminosity of 10^{34}. Additional computing resources are required for the simulation and reconstruction of higher luminosity events with upgraded detector designs. When the design phase is over, simulated events will be needed to develop the software and computing for the upgraded detector. Furthermore, there is a strong case for an increased output rate to allow proper study and precision measurements of semi-rare channels.

The requirements will effectively represent a new experiment in terms of requests to GridPP, although the requests will be presented coherently with those from the current ATLAS. One possible resource scenario is presented here. This is quite aggressive, but can still be met with a

	12	13	14	15	16	17	18	19	20
Specific luminosity (10^{34} cm^{-2}s^{-1})	1	1.5	2	2.5	3	0	5	8	10
Integrated luminosity (fb^{-1})	108	198	318	468	648	648	948	1428	2028

Table 4.2: The assumed specific and integrated luminosity profile from 2012-2020. The sum of the required resources at CERN, the Tier 1s and the Tier 2s for ATLAS is given in Fig. 4.3. An attempt has been made to convert the processing requirement into a number of cores.

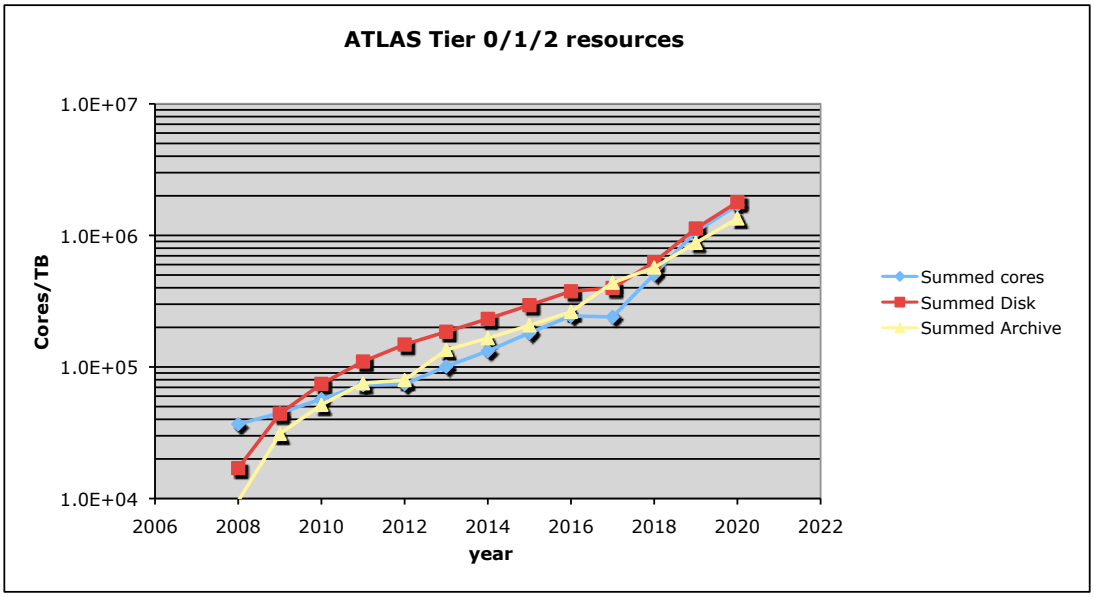

Figure 4.3: The projected resource needs for ATLAS at CERN, in the Tier 1s and the Tier 2s. The capacity in local Tier 3 facilities is not included.

flat budget and a Moore's Law doubling over 18-24 months. Should a less aggressive scenario come to pass, the storage and CPU requirements will still rise more than linearly, with increasing requirements per event for each step in luminosity. The issues require new architectures to arise in Phase-I and Phase-II in all scenarios.

The assumed luminosities are given in Table 1; they assume that there will be a LHC shut-down for detector and accelerator upgrades during 2017, after which a new phase of data taking would begin and reprocessing of earlier data would largely cease. To broadly characterise the results, there is a good degree of agreement between the outcomes for the two experiments, and at CERN, in the major national centres ('Tier 1') and shared university facilities ('Tier 2') there is a growth of between two and three orders of magnitude in the required processing power, active and archival storage. While no explicit calculation was made for the non-shared ('Tier 3') resources required, a similar scaling to that seen in the Tier 2s is to be expected. It is worth noting that the assumption that the output rate from the detectors scales with the specific luminosity after 2013 affects the projections over a decade by only a factor of around three, although would become more significant thereafter.

To illustrate the extreme scale of the requirements, the Tier 1 capacity by 2020 of each of

the experiments would be about 1 million cores and one thousand Petabytes of active storage. The tier 2 capacity numbers are similar. While these projections would be met by a Moores Law growth with a doubling every 1.5-2 years, they present many technical challenges to allow the resources to be exploited, and will entail significant support from the host institutions.

Consideration has also been give to the networking and bandwidth requirements implied by the growth in resource requirements described above. Some of the issues will be discussed in later sections, but one obvious challenge that presents itself is the need to reprocess the raw data stored on archival media at the Tier 1s. This must be done to benefit from detector calibration and alignment improvements and new reconstruction algorithms, which will all have advanced during study of the initial processing of the data. At a typical Tier 1, this will start by requiring O(100) MB/s average recall rate from archive, but after a decade this will be more like O(10) GB/s. This would seem to pose a challenge for any sequential access medium like tape, and indeed may be hard with any mechanical storage medium.

A similar consideration applies to data-access for on demand user analysis at the Tier 2 facilities. In the present architecture, the data is typically accessed locally from a storage element by processes on the cores. Each process reads at between 2 and 30 MB/s. As the number of cores grows, the aggregate average rate will grow from of O(1) Gb/s per experiment today to O(20) Gb/s in a decade. This can face many potential bottlenecks, some of which will be addressed below.

The ATLAS near-term computing resource projections and requests remain a UK responsibility; it is intended that the longer-term projections for the sLHC will form an continuing extension to this work.

4.2.5 Computing upgrade path approaches

The following computing techniques and hardware opportunities will all help the considerable computing challenges previously outlined. The true performance, effort required to exploit the approach must be evaluated for the major computing use cases in order to plan the computing to meet the upgrade requirements. The UK has been at the forefront of identifying the challenges to be faced and the strategies that offer solutions.

4.2.5.1 Virtualisation

An area of increasing development is virtualisation. At present, this is of particular interest for server consolidation and in service provision within cloud computing environments. These techniques can also extend the sites available to run ATLAS tasks on the Grid, which will partially mitigate the very large increase in the required resources.

Given the rapid increase in the number of CPU cores and large amount of direct access memory available within a single system it is desirable to pin and furthermore dynamically adjust these resources for dedicated services and applications. This is possible through virtualisation methods which allow for CPU and memory utilisation tuning without the need for resource downtime. This is particularly important for upgrade simulations where the amount of memory available for a virtual instance can be adjusted to the exact amount needed for various pile-up scenarios. Further, there are prospects that kernel-level virtualisation options that will partially aid the use of large numbers of cores per processor. Given the large issues to be addressed, all such options must be investigated, and combined with optimisations in the ATLAS code base.

There are many further important opportunities for ATLAS upgrade analysis. Firstly, persistent storage can be managed very efficiently by having a scalable amount of virtual instances branch from a base virtual instance using file system device mapping techniques such as CoW (copy-on-write) so only differences between the base OS - and underlying data - and the child OS require disk space. This is similar to the use of CoW for effective memory management discussed earlier (which can again be used here), but in this context can be applied to ATLAS software libraries. It is also useful when applications require host partitioning, but need to only to modify small elements of large data sets.

A further advantage to virtualisation is the possibility of live host migration. For example, running instances of the ATLAS software can be moved to intransient data sets or even follow data in a computing cloud which will incur no downtime and only a small performance impact on the running process.

4.2.5.2 Computing on graphics processing units (GPU)

The potential performance of a Graphics Processing Unit (GPU) for running code is staggering compared with its low cost and power consumption. Mainly due to the inbuilt nature of a modern graphics chip, where massively parallel and extremely fast processing is completely normal.

As an example to elucidate this, if we consider one NVIDIA Tesla S1070 chip (GPU targeting compute applications), it can provide up to 4.147 TFLOPS in single precision. This is roughly equal to 345 3.0 GHz Xeon processor cores. The next GPU model to be released is expected to have a similar number of double precision units and error correction checking. With this, and future developments, it is expected that the performance improvement will far surpass CPUs [43]. This can be seen in Fig. 4.2.5.2, where in a similar fashion to CPUs, the performance is increasing exponentially (currently 4 TFlop/s), but the doubling time is roughly half that of CPUs.

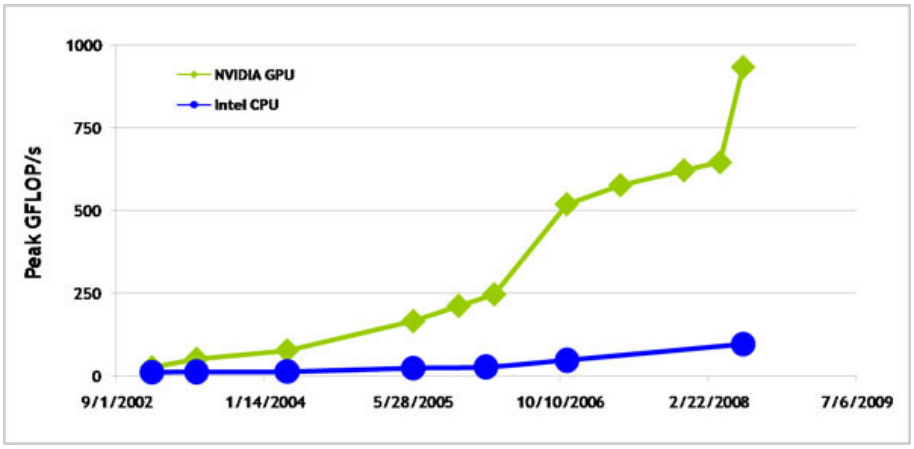

Figure 4.4: NVIDIA GPU versus Intel CPU performance in GFlops/s.

The main strengths of GPU programming are:

- Many more floating point units than a CPU, if these units can be kept busy then code can go much faster

- It has many more threads in flight than a CPU. One thing this can do is to effectively remove the latency of memory access. While one thread is fetching from memory, other threads are executing on the floating point units.
- The memory interface is much faster, utilising a wider interface than that of a CPU. In addition, the GPU uses much quicker DRAM (GDDR3). The memory available is large, easily accommodating several thousand ATLAS events in one go. The total memory for the Tesla S1070 is 4x4 GBs with a bandwidth of 408 GB/s.

Together, the GPU is ideal for data processing which requires a lot of floating point calculation, such as analysis, simulation and some parts of reconstruction & triggering.

Nvidia is the only major independent GPU manufacturer left. Therefore not surprisingly they have the most advanced architecture for GPU programming called CUDA, see Fig. 4.5.

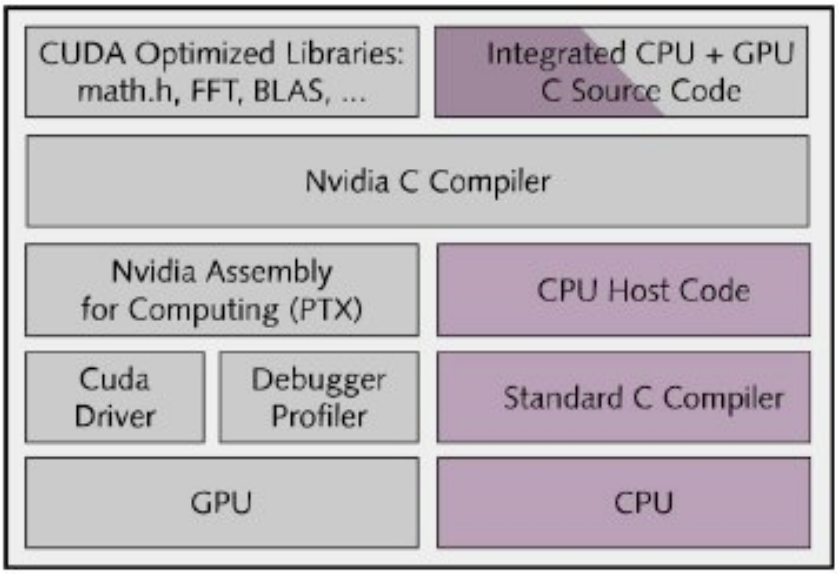

Figure 4.5: Nidia's CUDA platform with a C/C++ source-code compiler for CPUs.

There is also the OpenCL framework (Open Computing Language) which is the most likely framework we may adopt. AMD, Intel and NVIDIA have all committed to this and is the likely emerging standard. OpenCL consists of functions called kernels and it provides APIs to define and then control the CPU and GPU platforms. It provides both task-based and data-based parallelism.

Lastly, on ECDF (ScotGrid-Edinburgh) a local queue "gpgpu" has been set up to run GPU programs using a NVIDIA Tesla S1070. The queue comprises of two worker nodes (IBM x3550). It can do two times 4 CPU threads and 480 GPU threads on each of the two nodes.

4.3 Software, Simulation and Computing Summary

We have outlined above a large number of different areas that must be addressed by ATLAS. Given the limited manpower resources that we will have available, we have targeted carefully

areas that match the strength and expertise present within the UK, and will seek industrial support and funding. Collaboration with the particular ATLAS software experts and the CERN working groups on virtualisation and parallelisation will be important. We have already identified the appropriate contacts necessary within core Geant4 simulation, ATLAS tracking and some initial supportive NVIDIA industrial contacts. It is extremely difficult in computing to predict exactly what is the right future technology so we focus on what is achievable in a 3 year time horizon with the best possible ideas for the Phase-I and -II upgrades.

Chapter 5

Resources and Organisation

5.1 ATLAS UK Upgrade Organisation

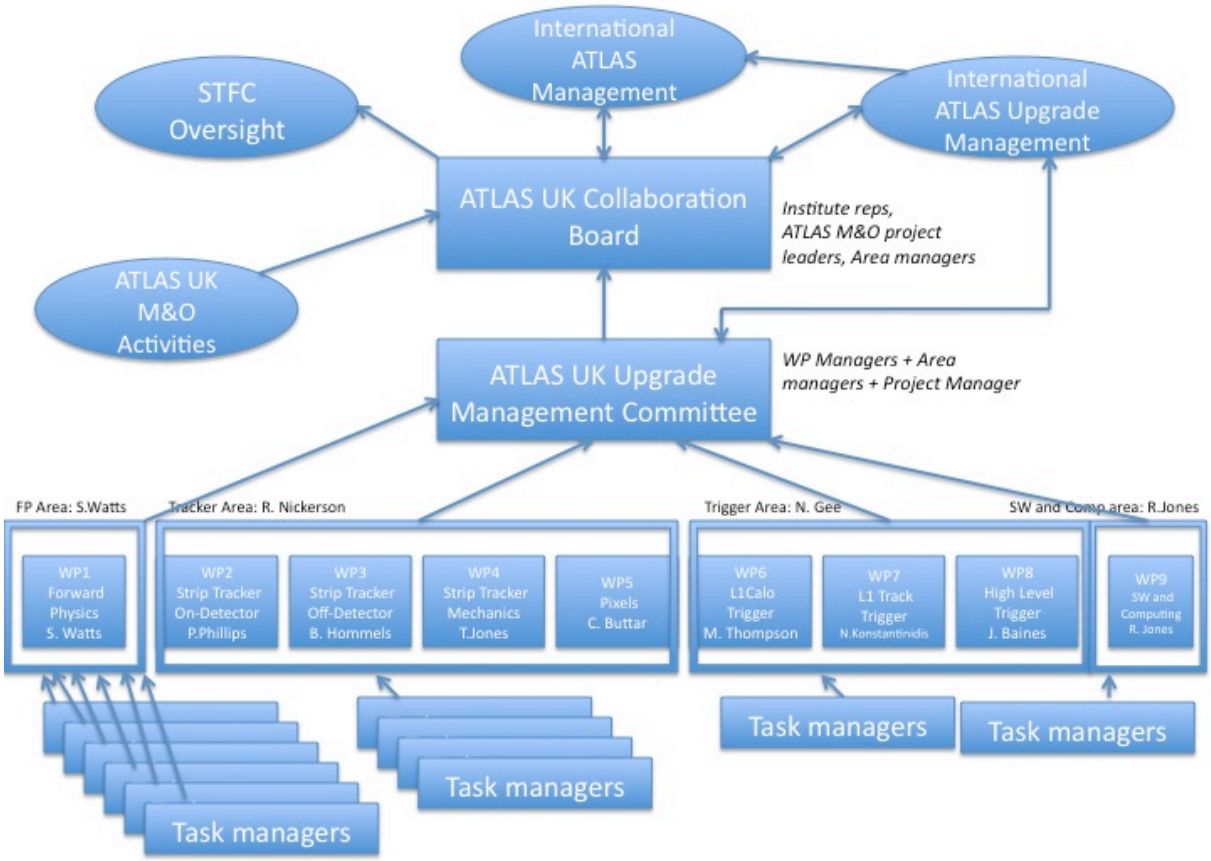

The above diagram summarises the ATLAS UK upgrade organisation. Arrows indicate the direction of reporting. The roles of the different bodies are as follows:

ATLAS UK Collaboration Board This existing body is the overall decision-making body for all ATLAS UK activities. It is made up of one voting representative from each ATLAS UK institute, the project leaders from the ongoing maintenance and operation projects (SCT, L1Calo, HLT, Software/Computing), and the upgrade area representatives, with the upgrade project manager invited as required. The chair is elected for a four year term (two years as deputy, two years as chair), and serves also as National Contact Physicist for the international ATLAS collaboration, and as Principle Investigator for the purposes of STFC. The UK CB is formally responsible for reporting to the STFC oversight committee and (via the National Contact Physicist) for financial and managerial relations with ATLAS. It is at the UK CB where all decisions involving transfer of resources between groups or between projects must be ratified before approval by the oversight committee. The UK CB will regularly review the project progress and the development of the risk register, in synchronisation with the oversight committee meetings. It meets roughly twice a year for face-to-face meetings, with frequent teleconferences and email communication as required.

ATLAS Upgrade Management This is the existing international steering group coordinating ATLAS Upgrade activities. The UK is strongly represented on this body. Financial relationships (MoUs etc) are dealt with by this body and the UK CB (via the ATLAS CB). Project management issues are dealt with by this body and the ATLAS UK Upgrade Management Committee.

ATLAS UK Upgrade Management Committee This is the working project management comittee of the UK upgrade project, consisting of the project manager, the area managers and the WP managers. It will be chaired by the UK CB deputy chair. It receives reports on all workpackages and reports to the UK CB on project development, finances and risks. It makes recommendations to the UK CB on decisions involving significant transfer of resources or changes in scope. It will meet approximately once per month.

Area and Workpackage Organisation Some of the larger workpackages define further internal organisational structure, which is given in the appropriate part of Section 5.3. Where there are common issues between workpackages, the Area Managers will convene meetings within their areas are needed.

5.2 Costs Summary

The process for deriving the project costs was as follows: The scope of the project and the project organisation was defined. The project is divided into Workpackages, each of which has a manager. Several of the workpackages are subdivided into tasks which also have defined managers. Where they exist, the task managers were asked to evaluate the costs of the tasks for which they had responsibility. These estimates were reviewed by the project engineer and the workpackage manager, with a view to confirming the estimates and to ensure there was no overlap between task areas. The estimates were adjusted, in consultation with workpackage managers, where necessary. The appropriate working margin was evaluated taking into account the type of expenditure, again by the task managers, but reviewed by the project engineer. Inflation was not included, the costs are in £2009. The cost estimates are therefore derived bottom-up, but reviewed for consistency and coherence top-down. Cost and effort estimates are based on the activities recently completed as part of the main ATLAS construction project.

ATLAS Upgrade All WPs			Year 1	Year 2	Year 3	Total
Staff						
HEI / Universities	Birmingham	Rolling Grant	48.9	101.4	126.6	276.8
		New Money	88.8	100.8	115.2	304.9
	Cambridge	Rolling Grant	180.8	219.5	242.8	643.2
		New Money	47.7	48.5	49.0	145.2
	Edinburgh	Rolling Grant	3.2	7.9	8.0	19.0
		New Money	77.2	142.1	160.0	379.3
	Glasgow	Rolling Grant	118.6	133.7	142.1	394.4
		New Money	54.5	55.6	60.3	170.3
	Lancaster	Rolling Grant	130.9	147.4	157.3	435.6
		New Money	47.5	48.6	49.8	145.9
	Liverpool	Rolling Grant	266.6	326.6	358.3	951.5
		New Money	7.9	8.1	8.2	24.1
	Manchester	Rolling Grant	403.1	387.7	385.5	1,176.2
		New Money	235.4	235.4	234.0	704.8
	Oxford	Rolling Grant	508.6	495.5	504.4	1,508.5
		New Money	22.0	88.3	89.1	199.4
	QMUL	Rolling Grant	194.6	214.5	235.9	644.9
		New Money	81.2	81.2	81.2	243.6
	RHUL	Rolling Grant	5.7	14.9	24.2	44.8
		New Money	141.3	142.9	145.0	429.2
	Sheffield	Rolling Grant	98.4	126.6	162.9	388.0
		New Money	25.6	95.7	104.1	225.4
	Sussex	Rolling Grant	3.5	9.7	7.0	20.2
		New Money	35.4	35.8	36.3	107.5
	UCL	Rolling Grant	135.5	146.5	146.5	428.4
		New Money	133.0	141.5	141.5	416.1
Total HEI / University Staff						
Total - Rolling Grant			2,098.3	2,331.7	2,501.6	6,931.5
Total - New Money			997.6	1,224.5	1,273.6	3,495.7
STFC	RAL PPD		954.3	1,088.5	1,249.7	3,292.5
	RAL TD		869.5	1,173.4	1,368.3	3,411.1
Total STFC Staff			1,823.8	2,261.8	2,617.9	6,703.6
Recurrent						
Equipment			1,352.3	2,245.9	1,476.7	5,075.0
Consumables			85.0	78.5	81.5	245.0
Travel			285.4	312.3	321.2	918.8
Other						-
Total Recurrent			1,722.7	2,636.7	1,879.4	6,238.8
Project total			6,642.3	8,454.8	8,272.6	23,369.7
Working Margin			141.6	181.8	163.1	486.4
Contingency						
Project Total (including margin/contingency)			6,783.9	8,636.5	8,435.7	23,856.1
Rolling grant			2,098.3	2,331.7	2,501.6	6,931.5
Project Total (less rolling grant)			4,685.6	6,304.9	5,934.1	16,924.6

The totals compare to the estimate for these three years in the SoI of £M 6.7, 8.0, and $\sim$ 12.0, totalling £26.7M.

5.3 Workpackages

WP1: Forward Physics

Overview

See Section 2.2.

The goal of this workpackage is to deliver the forward physics upgrade detectors. The UK activities are as follows:

Task 1: Machine interface, Hamburg pipe, and associated equipment

AFP is integrated with the LHC. Efficient liaison with the LHC machine physicists is vital. During the FP420 project regular Technical Integration Meetings (TIM) were held to coordinate the work that led to the NCC and detector layout that is described in detail in the FP420 Design Report. The meeting was coordinated by CERN machine physicists and Keith Potter from Manchester/Cockcroft. During the construction phase, the following work will be required.

- To finalise the design of the Hamburg Pipe and build a pre-production version. This will include the moving mechanisms and beam-pipe monitors (BPMs). Much R&D work on the Hamburg pipe was performed by the Louvain group for FP420. This work needs to be completed and the engineering design verified by the LHC machine division. This will be coordinated by the Technical Integration Meeting. The pre-production unit will be thoroughly tested for reliability.

- Liaison between collaboration and machine, including layout, up-to-date optics, interaction with collimation. Updating of the detector acceptance and chromacity diagrams.

- Background environment at 220m and 420m. Calculation of background at 220m and 420m with e.g. FLUKA, G4. Interface to detector models. Measurement of background with beam (e.g. 20 MeV CH, Muons, 1 MeV Equivalent neutrons). Perform a radiation review of all mechanical and electrical systems.

- Beam transport and Optics. Validation of proton tracking tools and any input needed for acceptance studies. Development of collimator strategies. Special topics like local bumps or optimised optics to increase acceptance.

- Machine protection studies. Circulating beam failure scenario (beam dynamics in quench, PC failure, collimator shadow etc). Protection against local bumps. It has been agreed that there will be a joint safety committee for CMS and ATLAS Forward Physics detectors. A joint risk review will be conducted in 2010 along similar lines to the LHCb VELO Risk Review.

- Beam Loss Monitor (BLM) studies. Placement and signal in BLMs. Setting of thresholds for machine protection including link to BIS. Damage threshold of detector and housing.

- Final study of RF impact of Hamburg pipe.

Task 2: Silicon Tracker

This task is in three parts.

- Complete the R&D and finalise the design

  As noted in the sensors section, there are two options for the sensor; FE-I3 or FE-I4 compatible design. This will be developed within the 3D work-package and a decision on which option to take will be taken in late 2010/early 2011. The FE-I4 is the baseline, however, the FE-I3 will be used if there are significant delays in the FE-I4 chip or FE-I4 sensors.

  To finalise the cooling design. The heat from the front-end chips has been transferred to a copper block. This now has to be transferred out of the station. Various designs of heatsink source suitable for use in the LHC tunnel (e.g. pulse-tube, dry-air cooling etc.) need to be evaluated by a design review and one option chosen and tested. This will be performed in collaboration with the Prof. Vacek at the Czech Technical University in Prague, who has worked on similar problems for the TOTEM experiment.

  To test pre-production superlayers, and verify assembly and burn-in procedures.

  To select one of the three options given in the FP420 design report for the LV/HV system and test a pre-production system. We will collaborate with AFP colleagues in Cracow on this system.

- Production

  Production of the sensors will commence at the start of 2011. Once processed and tested, they will have to be bump-bonded and then re-tested prior to use in the construction of superlayers. The superlayers then need to be assembled with the mechanics into the final tracking station.

  We will start with a pre-production tracker which will test all the assembly procedures, burn-in, testing and commissioning.

  After the pre-production system has been successfully tested, then the four production tracking stations plus a spare, required for the 420 AFP position, will be assembled.

- Commissioning and installation.

  The tracker has to be integrated with the Hamburg pipe, alignment system, and timing detectors, and the whole system tested and installed. This is a major task and is described below.

Task 3: DAQ, Monitoring and Slow Control

There are two aspects to this task.

- Signals from the detector station are sent via optical links to a readout board in the control room (ROD). This has to be programmed for the AFP data and then integrated with the ATLAS DAQ. DAQ for both the tracker and timing detectors will be developed. Trigger and DAQ issues will also be resolved in this task.

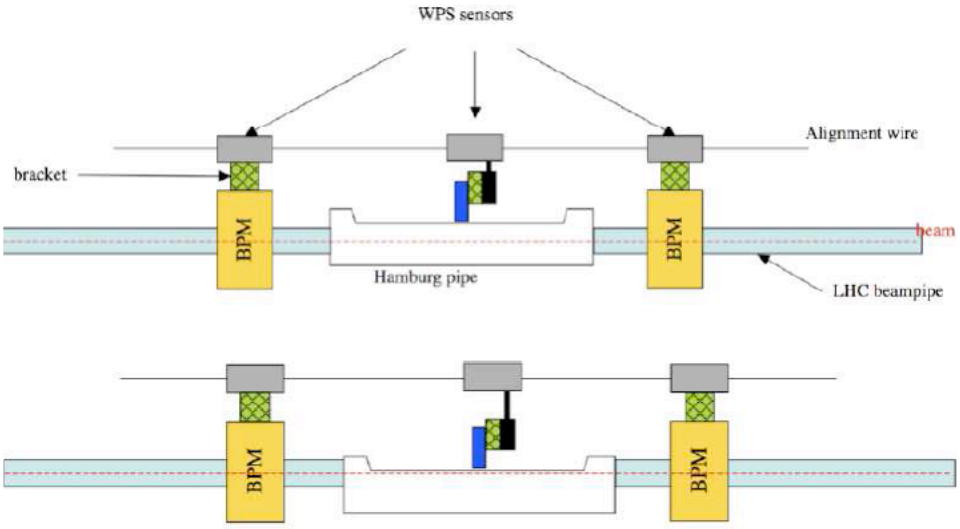

Figure 5.1: Schematic of the wire alignment system and its relationship to the Hamburg Pipe, Beam Position Monitors (BPMs) and the silicon tracker.

- There are many parameters of the system (e.g. temperature, motor position, alignment wire, voltage settings etc.) that have to be set and monitored. Software for this monitoring system has to be written. This system is important to the safety of both the AFP detectors and the LHC. Note that the Hamburg Pipe will be treated like a collimator by the LHC operators and will be under their control.

Task 4: Alignment and Calibration

Alignment of the detector is in three parts:-

- Careful assembly and survey of the tracking detector. This is part of Task 2. The final detector station will be mounted onto the Hamburg Pipe and commissioned and aligned with tracks on a CERN test beam (Task 6).

- Measurement of the beam position and relationship to the tracking system.

  There are two elements to this part beam position monitors (BPMs) and a wire alignment system. These were studied for the FP420 project and conceptual designs are in the Design Report. The BPMs are modified LHC monitors with improved readout electronics to provide few micron precision. The BPMs are related to the silicon tracker by a wire alignment system which is shown in Figure 5.1. The UK will complete the design of this system and build them for the AFP stations. The work is split between Manchester (wire-alignment) and UCL (BPM readout).

- To calibrate the system to obtain the absolute momentum from the tracking information. This is performed when installed in the LHC.

To obtain the planned performance for the 420 m silicon system, considerable attention must be paid to its calibration. The calibration involves two aspects: the accurate alignment

of the positions of the detectors, and the use of the measurements, once the alignment has been achieved, to obtain accurate values of the momenta of the protons recorded in the system. Accurate beam position monitors will be installed by which we can know the position of the detectors with respect to the beam, the position of the beam with respect to the magnets and the position of the counters with respect to the magnets, which is the desired goal. There will be considerable hardware redundancy here. By recording diffractively scattered protons, the known correlation between the distance of deflection of the protons and the angle of their trajectory can be used to check the information from the hardware devices. This requires care and logical thinking, and the necessary monitoring must take place continually during data taking, since the beam can move and the silicon detectors certainly will move during runs. The calculation of the momentum of a given proton is done in terms of its lateral deflection and the angle of its trajectory, and at 420 m has greater sensitivity to the angle. At present, tracking calculations are used to derive a formula by which the momentum can be obtained, using numerical fits to obtain the momentum as a function of the position and angle of the tracked protons. It will be necessary to optimise this formula and extend it to cover variations in the beam conditions and the position of the interaction point. Movements of the magnets are hoped to be small, but the possibility of this must be investigated. Finally, we anticipate that the true momentum of the protons will differ from the calculated by small amounts whose origins must be investigated. For this, and in any case, an empirical calibration will be essential. We can trigger the central detector on muon pairs produced by photon-photon interactions from which the momentum of the protons, after radiation, can be calculated. Several hundred of these events should be available per day of running, and will enable the momentum of the protons as measured at 420m to be compared with the accurately known momentum values derived from the muon pair measurements. In this way, a calibration map as a function of momentum can be established, something that will need to be repeated frequently while we gain understanding of the performance of the system. Our collaborators at the 220 m system will be studying methods of calibrating the measurements there, and there will be a need to work closely with them since it is likely that data can be taken in which protons traverse both systems, enabling a cross calibration.

Task 5: Offline Software and Physics Studies

There will be a continuing need to perform physics studies to optimise the detector design and trigger strategy.

- Proton tracking and acceptance studies have been performed to date by Peter Bussey at Glasgow. This work will continue and is very relevant to Task 4. This work also feeds into physics studies.
- Generator level physics studies. Much of these have been performed at Manchester in recent years. These studies will continue as they are necessary to check new ideas on physics and trigger options.
- In recent months, more detailed Monte Carlo simulations of the physics channels have been performed. More work is required, especially on pile-up issues at high luminosity.
- GEANT simulations of the detector have been performed. These need to be continued with additional UK support.

- Offline analysis software for AFP needs to be developed.

Task 6: Assembly and Commissioning

The pre-production tracker will be mounted onto the pre-production Hamburg pipe and integrated with other systems and then tested and aligned on a test beam at CERN. Only the front station will be fully constructed. A dummy back station will be made as part of a prototype eight metre mechanical system. This will test the mechanical design and alignment between the front and back stations. This system will be commissioned and tested on a beam at CERN.

All the production trackers will be mounted on a Hamburg Pipe, integrated with the rest of the AFP instrumentation and then tested, commissioned and aligned on a test beam at CERN. The whole detector system will then be transferred to the LHC and installed. The UK will support this process for both the 220 and 420 stations. This stage will be mainly outside the three year period of this proposal period because it will not start until late 2012.

As noted above, a pre-production station will test all the assembly and commissioning procedures in the period of this proposal.

Task Summary

WP1	Workpackage Manager: Project Engineer:	Stephen Watts Andy Nichols
	Task	Task Organiser:
0	Organisation	
1	Machine Interface, Hamburg Pipe	R. Appleby K.Potter
2	Si Tracker R&D Production	R. Thompson, S. Kolya
3	DAQ and DCS	M. Campanelli, B. Gallop
4	Alignment and Calibration	J. Pater, P. Bussey
5	Offline and Physics Studies	M. Campanelli, A. Pilkington
6	Assembly and Commissioning	A. Nichols

Inputs

Results of R&D from the FP420 Collaboration Refer to FP420 Design report, reference [8]

Major outputs and deliverables

Task(s)	output/deliverable
1	Hamburg Pipe Design and pre-production unit manufactured and tested
2	R&D completed, pre-production tracker manufactured and tested, and production started
3	Operational data acquisition and data monitoring system for AFP
4	Production alignment system
5	Proton tracking software, Monte Carlo studies, GEANT4 simulation and offline analysis software
6	Services installed, pre-production station commissioned, production trackers started

Major Milestones

Milestone	proposed date
Hamburg Pipe Design and pre-production complete	Jan 2011
Tracker thermal design complete	Apr 2011
Pre-production tracker complete	Sep 2011
DAQ and DCS ready for pre-production station	Sep 2011
Pre-production alignment system ready	Nov 2011
Services installed during LHC 2011 shutdown	Dec 2011
Pre-production station and eight metre prototype ready	Feb 2012
Commissioning and test of pre-production system complete	Oct 2012
Commissioning tests of pre-production system in LHC at 220 m	April 2013
Two production trackers ready	Apr 2013

Note: Milestones associated with the FE-I4 sensor development are in WP4

Long Term Time Profiles

2010	Design and build of pre-production Hamburg Pipe. Complete R&D required to finalise the Tracker Design. Safety Review and Radiation Review. AFP Technical Design Report written and submitted to the LHCC by end of 2010. At end of 2010 place order for sensors.
2011	Sensors processed. Bump bond sensors for pre-production tracker. Install cables during 2011 LHC shutdown. Install Hamburg Pipe at 220 m with background monitoring detectors. Build and test the pre-production tracker and assemble with a Hamburg pipe and 8 metre prototype.
2012/3	Beam test and commissioning of pre-production system and build production trackers. Assemble, test and align all stations. Two trackers built by April 2013. Remainder completed by end of 2013.
2014	Finish commissioning and installation and operation in LHC of all detector stations.

The LHC schedule is not well known beyond 2011. Stations will be constructed, tested, aligned and stored. They will be installed in the LHC during maintenance periods as these occur. It looks likely that these machine maintenance periods will be for several months and occur after a running phase of 9-12 months.

Staff

	Type/Funding	**Institute**	**2010-11**	**2011-12**	**2012-13**	**Activities**
Bates	AP/RG	Glasgow	0.2	0.2	0.15	2,6
Bussey	Ac/RG	Glasgow	0.2	0.3	0.3	4
Buttar	Ac/RG	Glasgow	0.05	0.05	0.05	0,4
D'Auria	Ph/RG	Glasgow	0	0.1	0.2	4,6
Di Mattia	T/RG	Glasgow	0.1	0.1	0.1	2,4,6
McEwan	T/RG	Glasgow	0.1	0.2	0.2	2,4,6
Melone	T/RG	Glasgow	0.3	0.3	0.3	2,4,6
O'Shea	Ac/RG	Glasgow	0.05	0.05	0	0,2
AN Other	Ph/Project	Glasgow	1	1	1	4,6
Appleby	Ac/RG	Manchester	0.1	0.1	0.1	1,6
Cox	Ac/RG	Manchester	0.35	0.35	0.35	1,5,6
Da Via	Ac/RG	Manchester	0.1	0.1	0.1	2,6
Loebinger	Ac/RG	Manchester	0.1	0.1	0.1	5,6
Watts	Ac/RG	Manchester	0.5	0.5	0.5	0, All
Elvin	T/Project	Manchester	0.4	0.4	0.4	2,4,6
Freestone	E/RG	Manchester	0.4	0.4	0.4	2,4,6
Hasi	AP/RG	Manchester	0.2	0.2	0.2	2,6
Kelly	AP/RG	Manchester	0.3	0.25	0.2	2,3,6
Kolya	AP/RG	Manchester	0.3	0.15	0.1	2,3,6
Pater	Ph/RG	Manchester	0.8	0.8	0.8	2,4,6
Perry	T/Project	Manchester	0.35	0.35	0.3	2,4,6
Snow	Ph/RG	Manchester	0.3	0.3	0.3	2,6
Thompson	E/RG	Manchester	0.4	0.4	0.4	1,2,4,6
New RA	Ph/Project	Manchester	1	1	1	1,6
Watts - buy out	Ac/Project	Manchester	0.3	0.3	0.3	0, All
Pilkington	Ph/Other	Manchester	0.2	0.2	0.2	5
PPD assembly tech	T/Project	RAL (PPD)	0.1	0.1	0.3	2
M. Gibson	E/Project	RAL (PPD)	0.1	0	0	2
PPD analog engineer	E/project	RAL (PPD)	0.2	0.2	0.2	2
B. Gallop	Ph/Project	RAL (PPD)	0.1	0.2	0.1	3
J. Matheson	Ph/Project	RAL (PPD)	0.1	0.1	0.1	3,6
P. Philipps	Ph/Project	RAL (PPD)	0.1	0.1	0.1	3,6
M. Weber	Ph/Project	RAL (PPD)	0.1	0.1	0.1	2,6
PPD Physicist	Ph/Project	RAL (PPD)	0.15	0.15	0.15	2,6
Project engineer	E/Project	RAL (TD)	0.5	0.5	0.5	1,2,4,6
Eng support	E&T/Project	RAL (TD)	1.5	1.5	2	1,2,6
M. Campanelli	Ac/RG	UCL	0.3	0.3	0.3	3,5,6
P. Sherwood	Ph/RG	UCL	0.2	0.2	0.2	5,6
M. Warren	E/RG	UCL	0.15	0.15	0.15	3,6
G. Crone	E/RG	UCL	0.2	0.2	0.2	3,6
A. Lyapin	E/RG	UCL	0.1	0.2	0.2	3,4,6

Effort by task

The required FTE of staff effort across main areas of activity are:

- 36 FTE-yrs in total. This is broken down into 16.9 FTE-yrs of engineering effort plus 12.8 FTE-yrs of physicist/applied physicist effort plus 6.3 FTE-yrs of academic effort.
- 1.6 academic FTE-yrs and 1.5 engineering FTE-yrs are required to provide scientific/engineering high-level coordination within the UK and to the AFP project globally.

The estimated WP1 staff effort is broken down into tasks below. Engineering support from RAL TD has been put in the Technician and Engineering Support (T and ES) column. This amounts to 5 FTE-yrs and is a mix of engineer and technician support.

Task	Description	FTE-yrs				
		Ac	Ph/AP	E	T and ES	Total
0	Organisation	1.6	–	1.5	–	**3.1**
1	Machine Interface/HP	0.25	2.7	0.2	–	**3.15**
2	Si Tracker	1.0	5.95	3.1	7.5	**17.55**
3	DAQ and DCS	0.9	1.0	1.6	–	**3.5**
4	Alignment	1.0	1.0	0.6	1.0	**3.6**
5	Physics Analysis	1.1	1.15	–	–	**2.25**
6	Assembly/commissioning	0.45	1.0	0.4	1.0	**2.85**
Total WP1 Effort		6.3	12.8	7.4	9.5	**36.0**

Costs

ATLAS Upgrade WP1			Year 1	Year 2	Year 3	Total
Staff						
HEI / Universities	Glasgow	Rolling Grant	28.9	39.2	42.0	110.0
		New Money	47.5	48.4	49.4	145.3
	Manchester	Rolling Grant	259.6	241.9	233.4	734.8
		New Money	170.3	170.3	168.9	509.5
	UCL	Rolling Grant	49.9	58.7	58.7	167.2
		New Money				-
Total HEI / University Staff						
Total - Rolling Grant			338.4	339.7	334.0	1,012.1
Total - New Money			217.8	218.8	218.3	654.9
STFC		RAL PPD	65.0	65.4	75.4	205.8
		TD	158.1	167.3	211.8	537.2
Total STFC Staff			223.1	232.7	287.2	743.0
Recurrent						
Equipment			287.0	661.0	422.0	1,370.0
Consumables			10.0	5.0	5.0	20.0
Travel			61.0	80.0	85.0	226.0
Other						-
Total Recurrent			358.0	746.0	512.0	1,616.0
Project total			1,137.3	1,537.2	1,351.5	4,026.0
Working Margin			28.7	66.1	42.2	137.0
Contingency						
Project Total (including margin/contingency)			1,166.0	1,603.3	1,393.7	4,163.0
Rolling grant			338.4	339.7	334.0	1,012.1
Project Total (less rolling grant)			827.6	1,263.6	1,059.7	3,150.9

WP2: Strip Tracker On-Detector

Overview

See Section 2.3.3.

The overarching goals of the on-detector systems workpackage are to take the short strip stave prototype (STAVE09) to a finalised design, using the correct ASIC technology and then develop mass production module building techniques. By the time of the TDR the goal is to be ready to launch full scale module, and stave sub-system, production. This is expected to closely coincide with the April 2013 end of the proposal.

- Fully qualify sensor performance in the probable radiation environment.
- Prototype, iterate and finalise hybrid and module design.
- Set up two UK clusters to build modules during mass production and establish the build techniques.
- Develop the other on-stave systems (stave end card and tapes) to their final designs.
- Establish the processes, process control and QA that will be required during production.

Task Summary

WP2		Workpackage Manager: Project Engineer:	Peter Phillips Ian Willmut
#	Section	Task	Task Manager:
0		Management	
1	2.3.3.1	Strip Sensor Development*	G. Casse
2	2.3.3.2	Front end ASICs	P. Phillips
3	2.3.3.3	Hybrids	A. Greenhall
4	2.3.3.4	Strip Modules	T. Affolder
5	2.3.3.5	Service Tapes	A. Weidberg
6	2.3.3.6	Optical Interface	T. Huffman

Inputs

Baseline design for large area strip sensors and radiation study results
Tested DAQ and ASICs
UK experience with design and fabrication of tapes
UK collaboration in Versatile Link Project

Milestones

Milestone	Date
First Gbit optical links available.	Q1/2011
100 Full size final sensors available.	Q4/2011
Tested 130nm die available for use.	Q3/2012
First module produced with multi-module fixtures.	Q4/2012

Outputs

Tested Strip sensors and 130nm hybrids
Tested Modules and tapes to Stave programme
Mass manufacture tooling for WP2 areas
EoS card suited to stave programme

Staff Summary by Task

Task		FTEyears					
		Ac	Ph	AP	E	T	Total
1	Strip Sensor Development	0.6	1.25	0.8	0.4	2.2	5.25
2	Front end ASICs	0.7	0.4	1.1	0.8	2.1	5.1
3	Hybrids	0.95	0.5	0.15	1.8	1.4	4.8
4	Modules	1.75	3.4	0.8	2.1	4.5	12.55
5	Service Tapes	0.3	0.85	0.0	2.5	0.1	3.75
6	Optical Interface	0.3	0.25	0.0	2.1	0.1	2.75

Costs Summary

WP2 (On detector) Item / Cost	FY 2010/11 k£	FY 2011/12 k£	FY 2012/13 k£	Total k£
2.3.3.1 strip sensors	107.0	281.0	6.0	394.0
2.3.3.2 ASIC	36.0	44.0	106.0	186.0
2.3.3.3 Hybrids	20.0	42.0	67.0	129.0
2.3.3.4 Modules	25.0	355.0	90.0	470.0
2.3.3.5 Tapes	122.0	49.5	132.0	303.5
2.3.3.6 On stave opto	-	12.6	22.6	35.2
Use of University Facilities	37.4	48.1	55.7	141.1
Travel	41.5	41.5	41.5	124.4
WP2 (On detector) Total costs	388.8	873.6	520.7	1,783.1

Long Term Time Profile

2010-2012	Final R&D, Develop Production Methods
2012-2013	TDR, Pre-production Stave Manufacture
2014-2016	Procurement, Stave Production
2016-2018	Surface Assembly of ID
2018-2019	Installation of ID into ATLAS, ATLAS closed

Staff

	Type/Funding	**Institute**	**2010-11**	**2011-12**	**2012-13**	**Tasks**	
C.M Hawkes	Ac/RG	Birmingham	0	0.1	0.1	3	
J.A.Wilson	Ac/Project	Birmingham	0.05	0.1	0.1	2,3	
R.J Staley	E/RG	Birmingham	0.1	0.1	0.1	2,3	
S. Pyatt	T/RG	Birmingham	0.25	0.25	0.25	2,3	
X.Serghi	T/Project	Birmingham	0.25	0.25	0.25	2,3	
LBA Hommels	Ac/RG	Cambridge	0.1	0.1	0.1	2,3	
MA Parker	Ac/RG	Cambridge	0.2	0.2	0.2	0,1,4	
MJ Goodrick	E/RG	Cambridge	0.2	0.3	0.2	1,4	
D Robinson	Ph/RG	Cambridge	0.3	0.3	0.3	4	
RJ Shaw	T/RG	Cambridge	0.4	0.4	0.5	1,4	
O'Shea	Ac/RG	Glasgow	0.1	0.1	0.15	0,2,4	
Bates	AP/RG	Glasgow	0.05	0.05	0.1	1,4	
Eklund	AP/RG	Glasgow	0.1	0.2	0.25	2	
Doherty	T/RG	Glasgow	0.3	0.4	0.4	1,2	
Di Mattia	T/RG	Glasgow	0	0.1	0.2	1,4	
Melone	T/RG	Glasgow	0.2	0.2	0.15	4	
McEwan	T/RG	Glasgow	0	0.1	0.1	2,4	
H. Fox	Ac/RG	Lancaster	0.05	0.05	0.05	0,1,4	
A. Chilingarov	Ph/RG	Lancaster	0.7	0.7	0.7	1,4	
I. Mercer	T/RG	Lancaster	0	0	0.2	1,4	
J. Statter	T/Project	Lancaster	0.5	0.25	0.25	1,4	
Allport P	Ac/RG	Liverpool	0.3	0.3	0.3	0,1,2,3,4	
Affloder A	AP/RG	Liverpool	0.3	0.3	0.3	1,2,3,4	
Casse G-L	AP/RG	Liverpool	0.25	0.25	0.25	1,2	
Dervan P	AP/RG	Liverpool	0.1	0.15	0.2	4	
Carroll JL	E/RG	Liverpool	0.05	0.05	0.05	4	
Greenall A	E/RG	Liverpool	0.6	0.6	0.6	2,3,4,5	
Sutcliffe P	E/RG	Liverpool	0.05	0.05	0.05	4	
Tsurin I	E/RG	Liverpool	0.1	0	0	1,4	
Whitley M	T/RG	Liverpool	0.05	0.1	0.1	4	
Wormald MP	T/RG	Liverpool	0.4	0.4	0.4	4	
Workshop	T/RG	Liverpool	0.2	0.3	0.5	1,2,3,4	
LSDC	T/RG	Liverpool	0.05	0.08	0.13	1,2,3,4	
Workshop	T/Project	Liverpool	0.1	0.2	0.2	1,2,3,4	
LSDC	T/Project	Liverpool	0.05	0.1	0.1	1,2,3,4	
A.R. Weidberg	Ac/RG	Oxford	0.2	0.2	0.2	5,6	
R. Wastie	E/RG	Oxford	1	1	1	5,6	
P. Lau	E/RG	Oxford	0.2	0.2	0.2	5,6	
SRF Mech	T/RG	Oxford	0.34	0.34	0.34	5,6	
SRF Elec	T/RG	Oxford	0.34	0.34	0.34	5,6	
SRF Mech	T/RG	Oxford	0.17	0.17	0.17	5,6	
SRF Elec	T/RG	Oxford	0.17	0.17	0.17	5,6	
Continued...							

	Type	Institute	2010-11	2011-12	2012-13	Tasks
M.Bona	Ac/RG	QMUL	0.2	0.2	0.3	4
A.Bevan	Ac/RG	QMUL	0.15	0.15	0.15	4
A.Martin	Ac/RG	QMUL	0.1	0.1	0.1	0,4
F.Gannaway	E/RG	QMUL	0.1	0.1	0.1	4,5,6
J.Morris	E/Project	QMUL	0.2	0.2	0.2	4,5,6
G.Beck	Ph/RG	QMUL	0.1	0.1	0.1	4,5,6
J.Mistry	T/E	QMUL	0.1	0.1	0.1	4,5
R. Holt	E/PPD	RAL (PPD)	0.2	0.2	0.2	2,3
J. Matheson	AP/PPD	RAL (PPD)	0.1	0.1	0.1	3,4,5
P. Phillips	AP/PPD	RAL (PPD)	0.2	0.2	0.2	0,2,3,4,5,6
M. Weber	Ac/PPD	RAL (PPD)	0.1	0.1	0.1	2,3,4,5
M. Tyndel	Ac/PPD	RAL (PPD)	0.3	0.2	0.15	1,2,3
eng support	E/TD	RAL (TD)	0.6	0.6	0.8	2,3,4
E. Paganis	Ac/RG	Sheffield	0.1	0.2	0.2	3,4
R.S. French	E/RG	Sheffield	0.2	0.3	0.3	3,4,5
I. Dawson	Ph/RG	Sheffield	0	0.1	0.3	4
P. Hodgson	Ph/RG	Sheffield	0.3	0.3	0.3	4,5
P. Johansson	Ph/RG	Sheffield	0	0	0.2	4

Costs

ATLAS Upgrade WP2			Year 1	Year 2	Year 3	Total
Staff						
HEI / Universities	Oxford	Rolling Grant	134.7	134.7	134.7	404.1
		New Money				-
	Sheffield	Rolling Grant	41.8	61.7	95.8	199.3
		New Money				-
	Glasgow	Rolling Grant	28.4	41.5	47.0	117.0
		New Money	-	3.6	7.3	10.8
	Cambridge	Rolling Grant	65.5	77.9	68.6	212.0
		New Money				-
	Birmingham	Rolling Grant	14.3	16.6	16.9	47.7
		New Money	7.4	8.6	8.9	24.9
	Liverpool	Rolling Grant	95.8	99.8	107.3	302.9
		New Money				-
	Lancaster	Rolling Grant	43.1	45.9	58.2	147.3
		New Money	16.6	8.8	9.4	34.8
	QMUL	Rolling Grant	24.2	24.2	25.7	74.0
		New Money	9.3	9.3	9.3	28.0
Total HEI / University Staff						
Total - Rolling Grant			447.8	502.3	554.2	1,504.4
Total - New Money			33.4	30.3	34.9	98.6
STFC		RAL PPD	111.5	101.0	97.1	309.7
		TD	59.6	61.7	163.6	285.0
Total STFC Staff			171.1	162.8	260.8	594.7
Recurrent						
Equipment			317.3	793.1	424.8	1,535.2
Consumables			21.0	21.0	31.0	73.0
Travel			41.5	41.5	41.5	124.4
Other						-
Total Recurrent			379.7	855.6	497.2	1,732.5
Project total			1,032.1	1,551.0	1,347.0	3,930.1
Working Margin			15.2	7.2	28.5	50.9
Contingency						
Project Total (including margin/contingency)			1,047.3	1,558.2	1,375.5	3,981.0
Rolling grant			447.8	502.3	554.2	1,504.4
Project Total (less rolling grant)			599.5	1,055.8	821.3	2,476.6

WP3: Strip Tracker Off-Detector

Overview

See Section 2.3.4, 2.3.6.

The work serves two major purposes; firstly, to develop systems so that modules and other prototype and production parts can be tested during final development of the mass production capability. In addition the work involves development of the production versions of off-detector systems, including the passive optics, power and DAQ. There are relatively few milestones but the work is very closely linked with the on-detector systems (Section 2.3.3).

- Develop a fully functional DAQ system that can read out a full stave at the data rates the sLHC requires.
- Take the present powering work forward so that there is a tested robust powering scheme ready to be implemented in the final stave designs.
- Provide the powering and DAQ support to the two UK building clusters.
- Drive forward the off detector support work, providing the information to make informed decisions about powering, readout and DAQ.
- Develop Final QA procedures for passive optics.

Task Summary

WP3		Workpackage Manager: Project Engineer:	Bart Hommels Ian Willmut
#	Section	Task	Task Manager
0		Management	
1	2.3.4.1	Data Acquisition	B. Hommels
2	2.3.4.2	Power	G. Villani
3	2.3.4.3	Passive Optics	T. Huffman
4	2.3.6	Pre-Production System Tests	D. Robinson

Inputs

Baseline HSIO hardware, with performance figures, built according to UK specifications.
Specification of stave09 and stave11 interfaces
Power source prototype hardware, specifications and power protection chip prototypes
Optical fibres with reliability analysis andsetup for evaluation in controlled environment
Stave09 prototype hardware, software and controls

Milestones

Milestone	Date
Stave11 built and tested	Q3/2011
Prototype programmable supply evaluated	Q4/2010
Evaluation reports completed on passive optical components.	Q4/2011
ID TDR Complete	Q4/2012

Outputs

Interface boards specifications for stave09/11 prototypes
Hardware between HSIO and End-Of-Stave controller, including Versatile Link
Software and hardware integration of HSIO with module production facilities
Test procedure and setup for production fibres & parts, pre and post irradiation
Specifications and evaluation test set up for the programmable power source
Evaluation of the protection chip prototypes.
Provision of Stave09/11 services systems: cooling, power, readout. Installed at CERN.

Staff Summary by Task

Task		FTEyears Ac	Ph	AP	E	T	Total
1	Data Acquisition	3	3.25	0.0	2.2	0.9	9.35
2	Power	0.6	0.9	0.0	2.7	0.7	4.9
3	Passive Optics	1.3	0.0	3	2.1		6.4
4	Pre-production System tests	1.45	2.1	0.0		0.35	3.9

Costs Summary

WP3 (Off detector) Item / Cost	FY 2010/11 k£	FY 2011/12 k£	FY 2012/13 k£	Total K£
2.3.4.1 DAQ	13.0	18.0	38.0	69.0
2.3.4.2 Powering	30.0	30.0	39.0	99.0
2.3.4.3 Passive optics	28.5	16.5	16.8	61.8
Use of University Facilities	10.8	10.8	10.8	32.3
Travel	26.2	26.2	26.4	78.8
WP3 (Off detector) Total costs	108.5	101.5	131.0	341.0

Long Term Time Profiles

2010-2012	Final R&D, Develop Production Methods
2012-2013	TDR, Pre-production Stave Manufacture
2014-2016	Procurement, Stave Production
2016-2018	Surface Assembly of ID
2018-2019	Installation of ID into ATLAS, ATLAS closed

Staff

	Type/Funding	Institute	2010-11	2011-12	2012-13	Tasks
LBA Hommels	Ac/RG	Cambridge	0.4	0.5	0.5	1,4
CG Lester	Ac/RG	Cambridge	0.1	0.15	0.2	1,4
JC Hill	Ph/RG	Cambridge	0.25	0.3	0.4	2,4
D Robinson	Ph/RG	Cambridge	0.2	0.2	0.2	4
MJ Goodrick	E/RG	Cambridge	0.2	0.3	0.2	1,2
RJ Shaw	E/RG	Cambridge	0.1	0.1	0.2	1,4
Greenshaw T	Ac/RG	Liverpool	0.1	0.1	0.1	4
Jackson JN	Ac/RG	Liverpool	0.05	0.1	0.1	0,4
Vossebeld J	Ac/RG	Liverpool	0.2	0.2	0.2	1,4
Affloder A	AP/RG	Liverpool	0.1	0.1	0.1	2,4
Dervan P	AP/RG	Liverpool	0.2	0.25	0.3	1,4
Greenall A	E/RG	Liverpool	0.1	0.1	0.1	1,2,4
King BT	Ph/Project	Liverpool	0.05	0.05	0.05	4
LSDC	T/RG	Liverpool	0.02	0.02	0.02	1,2,4
LSDC	T/Project	Liverpool	0.01	0.01	0.01	1,2,4
A. Barr	Ac/RG	Oxford	0.04	0.15	0.2	1
B.T. Huffman	Ac/RG	Oxford	0.2	0.3	0.3	3
C. Issever	Ac/RG	Oxford	0	0	0.2	3
A.R. Weidberg	Ac/RG	Oxford	0.1	0.1	0.1	3
M. Jones	E/RG	Oxford	0.5	0.5	0.5	3
S. Yang	E/RG	Oxford	0.2	0.2	0.2	3
Project Student	St/Project	Oxford	1	1	1	3
SRF Mech	T/RG	Oxford	0.34	0.34	0.34	1,3
SRF Elec	T/RG	Oxford	0.34	0.34	0.34	1,3
SRF Mech	T/Project	Oxford	0.17	0.17	0.17	1,3
SRF Elec	T/Project	Oxford	0.17	0.17	0.17	1,3
A.Martin	Ac/RG	QMUL	0.05	0.1	0.15	0,4
G.Beck	Ph/RG	QMUL	0.1	0.1	0.1	4
G. Villani	E/PPD	RAL (PPD)	0.5	0.5	0.5	2
M. Weber	Ac/PPD	RAL (PPD)	0.3	0.4	0.4	2,4
B. Gallop	AP/PPD	RAL (PPD)	0.4	0.3	0.4	1,2,4
R. Holt	E/PPD	RAL (PPD)	0.7	0.7	0.7	1,2
J. Matheson	AP/PPD	RAL (PPD)	0.3	0.3	0.3	2
P. Philipps	AP/PPD	RAL (PPD)	0.2	0.2	0.2	0,1,2,4
Tech support	E/TD	RAL (TD)	0.8	0.2	0.3	1,2
B Green	E/Project	RHUL	0.2	0.1	0.1	4
M. Wing	Ac/RG	UCL	0.2	0.2	0.2	1,4
M. Warren	E/RG	UCL	0.3	0.3	0.3	1,4

Costs

ATLAS Upgrade WP3			**Year 1**	**Year 2**	**Year 3**	**Total**
Staff						
HEI / Universities	Oxford	Rolling Grant	109.6	114.1	119.1	342.8
		New Money	0.0	0.0	0.0	0.0
	Cambridge	Rolling Grant	72.7	89.5	93.7	255.9
		New Money				-
	Liverpool	Rolling Grant	25.6	29.5	32.5	87.7
		New Money	4.9	5.0	5.0	14.9
	UCL	Rolling Grant	17.5	17.5	17.5	52.6
		New Money				-
	RHUL	Rolling Grant	-	-	-	-
		New Money	9.5	4.9	5.0	19.4
	QMUL	Rolling Grant	9.6	10.5	11.3	31.4
		New Money				-
Total HEI / University Staff						
Total - Rolling Grant			235.0	261.1	274.2	770.4
Total - New Money			14.4	9.9	10.0	34.3
STFC		RAL PPD	214.9	226.1	241.8	682.8
		TD	71.2	18.8	29.2	119.2
Total STFC Staff			286.0	244.9	271.0	802.0
Recurrent						
Equipment			75.7	68.7	98.0	242.5
Consumables			4.0	4.0	4.0	12.0
Travel			26.2	26.2	26.4	78.8
Other						-
Total Recurrent			106.0	99.0	128.4	333.4
Project total			641.4	614.9	683.7	1,940.0
Working Margin			7.3	6.8	8.8	23.0
Contingency						
Project Total (including margin/contingency)			648.7	621.7	692.5	1,963.0
Rolling grant			235.0	261.1	274.2	770.4
Project Total (less rolling grant)			413.7	360.6	418.3	1,192.6

WP4: Strip Tracker Mechanics

Overview

See Section 2.3.5.

The mechanics workpackage is designed to draw together, and refine, the design taking into account mass production requirements. The results will be essential for preparation of the upgrade TDR which will be produced close to the end of this program. The approval of the TDR will initiate the mass production of components. including support structures.

This workpackage also needs to support the on-detector systems program, ensuring that appropriate hardware and tooling are provided for the manufacture of module-carrying support structures. Many aspects of the mechanics programme interface extensively with the international programme, and in consequence there are many milestones that are external to the UK.

- Iterate and optimise the existing mechanical staves, reducing material and improving performance.
- Fully develop the tooling and processes needed to begin production at the end of the programme.
- Understand how the inner detector will be integrated together and be pro-actively in this development work with international collaborators.
- Have a fully developed quality policy for the build program that will ensure high quality production.

Task Summary

WP4		Workpackage Manager: Project Engineer:	Tim Jones Ian Willmut
#	Section	Task	Task Manager
0		Organisation	
1	2.3.5.2	Materials Selection	R. Bates
2	2.3.5.3	Cooling System	R. French
3	2.3.5.4	Stave Assembly	G. Viehhauser
4	2.3.5.5	Module Mounting	I. Wilmut
5	2.3.5.7	Test/Shipping System	G. Beck
6	2.3.5.8	Integration	I. Wilmut

Inputs

Candidate materials and options offering improved performance
Successful demonstration of bending and butt-welded stainless tubes
Experience and results from preceding stave programme
Requirements for shipping and testing.

Milestones

Evaluation of titanium tube completed	Q4/2010
Prototype test/shipping containers evaluated	Q3/2011
Functional prototype of full-scale module mounting system	Q4/2011
Final stave core selected and prototyped	Q2/2012

Outputs

Final set of material choices with fully characterised properties
Evaluation of titanium tube, spec and plans for mass manufacture
Optimised stave geometry, design, manufacturing process, production facilities
Two sites qualified to mount modules to staves
Designed, prototyped test/shipping system with tooling to fabricate

Staff Summary by Task

Task		FTEyears					
		Ac	Ph	AP	E	T	Total
0	Organisation	0.9	0.1	0.1	0.3	0.0	1.4
1	Materials Selection	0.3	0.2	0.4	1.0	0.8	2.7
2	Cooling System	0.5	0.0	0.0	2.5	2.0	5.0
3	Stave Assembly	1.0	0.8	0.6	8.6	2.3	13.3
4	Module Mounting	2.0	1.9	0.4	8.6	2.3	15.2
5	Test/Shipping System	0.7	2.1	0.4	4.9	1.1	9.2
6	Integration	0.8	0.6	0.0	4.7	0.3	6.3

Costs Summary

WP4 (Mechanics) Item / Cost	FY 2010/11 k£	FY 2011/12 k£	FY 2012/13 k£	Total k£
2.3.5.2 Materials Selection	22.5	12.0	-	34.5
2.3.5.3 Cooling	128.0	60.0	57.0	245.0
2.3.5.4 Stave assembly	155.0	115.0	44.0	314.0
2.3.5.5 Module mounting	30.0	88.0	37.0	155.0
2.3.5.7 Test Shipping Box	8.0	15.0	15.0	38.0
Travel	54.0	61.9	65.7	181.6
Use of University Facilities	17.9	28.4	42.5	88.8
WP4 (Mechanics) Total costs	415.4	380.4	261.2	1,056.9

Long Term Time Profiles

2010-2012	Final R&D, Develop Production Methods
2012-2013	TDR, Pre-production Stave Manufacture
2014-2016	Procurement, Stave Production
2016-2018	Surface Assembly of ID
2018-2019	Installation of ID into ATLAS, ATLAS closed

Staff

	Type/Funding	Institute	2010-11	2011-12	2012-13	Activities
ATC TD effort	E/TD	ATC	2	2	2	3,4
P. Clark	Ac/Project	Edinburgh	0.1	0.1	0.1	3,4
A. Main	T/Project	Edinburgh	0.05	0.1	0.2	3,4
Bates	AP/RG	Glasgow	0.15	0.1	0.05	1
Di Mattia	T/RG	Glasgow	0.1	0	0	1
McEwan	T/RG	Glasgow	0.3	0.1	0.1	1
H. Fox	Ac/RG	lancaster	0.05	0.05	0.05	2,3
I. Mercer	T/RG	lancaster	0.9	0.9	0.7	2,3,5
J. Statter	T/Project	lancaster	0	0.25	0.25	2,3,5
Allport P	Ac/RG	Liverpool	0.05	0.05	0.05	0,3,6
Greenshaw T	Ac/RG	Liverpool	0.05	0.05	0.05	1,4
Jackson JN	Ac/RG	Liverpool	0.15	0.2	0.3	0,3,4
Carroll JL	E/RG	Liverpool	0.05	0.05	0.05	3,4,5
Jones TJ	AP/RG	Liverpool	0.4	0.6	0.6	0,1,2,3,4,5
Muskett DA	T/Project	Liverpool	0.1	0.1	0.1	2,3,4
Sutcliffe P	E/RG	Liverpool	0.05	0.05	0.05	3,4,6
Whitley M	T/RG	Liverpool	0.05	0.1	0.1	3,4,5
Workshop	T/RG	Liverpool	0.2	0.2	0.2	2,3,4
Workshop	T/Project	Liverpool	0.1	0.1	0.2	2,3,4
R.B.Nickerson	Ac/RG	Oxford	0.5	0.5	0.5	0,3,6
G. Viehhauser	Ac/RG	Oxford	0.6	0.6	0.6	2,3,5,6
M. Dawson	E/RG	Oxford	0.38	0.13	0.1	2,3
P. Lau	E/RG	Oxford	0.4	0.4	0.4	2,4
W. Lau	E/RG	Oxford	0.43	0.4	0.4	2,6
Senanayake	E/RG	Oxford	0.43	0.55	0.5	2,5
S.Yang	E/RG	Oxford	0.2	0.2	0.2	2,6
SRF Mech	T/RG	Oxford	0.44	0.84	1.01	2,3,5
SRF Elec	T/RG	Oxford	0.34	0.84	1.34	2,3,5
SRF Mech	T/Project	Oxford	0.22	0.42	0.50	2,3,5
SRF Elec	T/Project	Oxford	0.17	0.42	0.67	2,3,5
A.Bevan	Ac/RG	QMUL	0.15	0.15	0.15	1,4,5
A.Martin	Ac/RG	QMUL	0.2	0.2	0.2	0,3,4,5
G.Beck	Ph/RG	QMUL	0.7	0.7	0.7	1,3,4,5
F.Gannaway	E/RG	QMUL	0.7	0.7	0.7	3,5
J.Morris	E/Project	QMUL	0.8	0.8	0.8	3,5
J.Mistry	T/RG	QMUL	0.7	0.7	0.7	3,5
Continued...						

	Type	Institute	2010-11	2011-12	2012-13	Activities
Ass Tech	T/PPD	RAL (PPD)	0.3	0.3	0.7	1,2,4
J. Matheson	Ph/PPD	RAL (PPD)	0.25	0.25	0.25	3,4,5
M. Gibson	E/PPD	RAL (PPD)	0.31	0.3	0.3	4
M. Tyndel	Ac/PPD	RAL (PPD)	0.2	0.3	0.15	3,4,5,6
M. Weber	Ac/PPD	RAL (PPD)	0.23	0.2	0.2	4,5
R. Preece	E/PPD	RAL (PPD)	0	0	0.5	4
S. Haywood	Ac/PPD	RAL (PPD)	0.1	0.2	0.2	0,4,5
Project engineer	E/TD	RAL (TD)	1	1	1	0,4,6
Eng support	E/TD	RAL (TD)	2.7	2.85	3.15	2,2,3,4,6
D. R. Tovey	Ac/RG	Sheffield	0	0.1	0.1	0,2,6
R.S. French	E/RG	Sheffield	0.5	0.5	0.5	2,5

Costs

ATLAS Upgrade WP4			Year 1	Year 2	Year 3	Total
Staff						
HEI / Universities	Oxford	Rolling Grant	213.9	220.7	229.9	664.5
		New Money				-
	Sheffield	Rolling Grant	41.4	44.3	45.3	130.9
		New Money				-
	Glasgow	Rolling Grant	17.9	9.2	6.2	33.3
		New Money	3.5	-	-	3.5
	Liverpool	Rolling Grant	43.9	64.3	76.6	184.7
		New Money	3.0	3.1	3.2	9.2
	lancaster	Rolling Grant	38.0	40.5	33.8	112.3
		New Money	-	8.8	9.4	18.2
	Edinburgh	Rolling Grant	-	-	-	-
		New Money	3.3	4.9	8.2	16.3
	QMUL	Rolling Grant	124.2	124.2	124.2	372.6
		New Money	37.3	37.3	37.3	112.0
Total HEI / University Staff						
Total - Rolling Grant			479.2	503.1	516.0	1,498.4
Total - New Money			47.0	54.1	58.1	159.2
STFC		RAL PPD	138.8	168.1	231.0	537.9
		TD	489.5	526.3	569.0	1,584.7
Total STFC Staff			628.4	694.4	800.0	2,122.7
Recurrent						
Equipment			334.0	295.1	173.4	802.5
Consumables			24.0	23.0	21.0	68.0
Travel			54.0	61.9	65.7	181.6
Other						-
Total Recurrent			412.0	380.0	260.1	1,052.1
Project total			1,566.6	1,631.5	1,634.1	4,832.3
Working Margin			43.1	43.2	36.3	122.6
Contingency						
Project Total (including margin/contingency)			1,609.7	1,674.8	1,670.4	4,954.9
Rolling grant			479.2	503.1	516.0	1,498.4
Project Total (less rolling grant)			1,130.5	1,171.7	1,154.4	3,456.5

WP5: Pixels

Overview

See Section 2.4.

One major goal of this workpackage is development of ultra-rad hard pixel sensors for the forward pixel upgrade for ATLAS, exploiting UK-leadership in sensor technology and utilising the IBL as an intermediate test of the technologies. The second major area of work is development of layout, mechanics and single-chip pixel module prototypes, using advanced connectivity techniques, for the forward pixel system upgrade; The goal is to place the UK in a position to take the leading role in this area. There are strong synergies with several of the other tracker work packages (WP2, 3, 4) and use will also be made of the simulation work in WP9.

- Development and demonstration of pixel sensors capable of operating in the ultra-high radiation levels anticipated in ATLAS during sLHC running,
- Delivery of tested pixel sensors for the construction of the intermediate B-layer,
- Delivery of a layout of the forward pixel system that minimises mass and optimises the reconstruction of primary vertices of forward going particles,
- Production of prototype forward pixel discs, based on the optimal layout, to demonstrate thermo-mechanical performance,
- Delivery of a single-chip pixel prototype module using advanced connectivity techniques.

Task Summary

WP5		Workpackage Manager: Project Engineer:	Craig Buttar Ian Willmut
#	Section	Task	Task Manager
0		Organisation	
1	2.4.1.2	Planar Sensor Technology*	G. Casse
2	2.4.1.3	3D Sensor Technology	C. daVia
3	2.4.2	Connectivity	M. Tyndel
4	2.4.3	Forward Pixel layout, Mechanics	T. Greenshaw
5	2.4.4	Single chip Module Prototype	R. Bates

* WP2 2.3.3.1 and WP5 2.4.1.2 have many synergies and a common task manager. There is no duplication of effort.

Inputs

Sensor design requirements for IBL UK-ATLAS/ATLAS
Simulation software and results for tracker studies
Results of tracker and pixel Readout chip studies

Milestones

Milestone	Date
Design for disk prototype ready (thermo-mechanical)	Q1/2011
Results from thermo-mechanical disk prototype	Q3/2012
Design of single-chip pixel prototype from layout studies	Q2/2012
Results from prototype disk	Q1/2013
Results from single-chip pixel prototype	Q1/2013

Outputs

Tested sensors for IBL 3D and planar single chip
Prototype modules for sLHC pixel upgrade
Performance results from lab, irradiation and testbeam studies
Performance results from thermo-mechanical disk prototype
SLHC forward pixel module and disk design

Staff Summary by Task

Task		FTEyears					
		Ac	Ph	AP	E	T	Total
0	Organisation	0.35	0	0.65	0.5	0.6	2.3
1,2	Sensor technology	0.55	0	3.55	0.4	0.75	8
3	Connectivity	3.3	0	2	0.1	0.4	2.8
4	Forward pixel layout, mechanics	0.2	0.1	0.8	0.8	0.15	2.4
5	Single pixel module	0.45	0.2	1.1	0.75	0.9	3.3

Costs Summary

WP5 (Silicon Pixels) Item / Cost	FY 2010/11 k£	FY 2011/12 k£	FY 2012/13 k£	Total K£
2.4.1.2 Forward Pixel	93.0	117.0	151.0	361.0
2.4.5 Irradiation Studies	9.0	11.0	8.0	28.0
2.4.2 Conectivity	71.0	116.0	55.0	242.0
2.4.1.3 3D Pixels	97.0	97.0	27.0	221.0
Use of UniversityFacilities	6.0	21.2	21.2	48.4
Travel	29.0	29.0	29.0	87.0
WP5 (Silicon Pixels) Total costs	305.0	391.2	291.2	987.4

Long Term Time Profile

2010-2013	R&D, Production method development
2012-2014	TDR, pre-production manufacture
2013-2015	Procurement, Start Manufacture
2015-2017	Manufacture complete, detector assembled
2017-2018	Integrate into ID on surface
2018-2019	Installation of ID into ATLAS, ATLAS closed

Staff

	Type/Funding	Institute	2010-11	2011-12	2012-13	Activities
Bates	AP/RG	Glasgow	0.2	0.25	0.3	0,5
Buttar	Ac/RG	Glasgow	0.1	0.1	0.1	2,5
Doherty	T/RG	Glasgow	0.3	0.2	0.2	0,4,5
Di Mattia	T/RG	Glasgow	0.1	0.1	0.1	1,2,5
Eklund	AP/RG	Glasgow	0.15	0.15	0.1	4,5
McEwan	T/RG	Glasgow	0.2	0.2	0.2	3,5
Melone	T/RG	Glasgow	0.1	0.1	0.15	3,5
O'Shea	Ac/RG	Glasgow	0.05	0.1	0.1	2,5
J.A.Wilson	Ac/RG	Birmingham	0.05	0.1	0.1	1,2,3,5
Affloder A	AP/RG	Liverpool	0.05	0.1	0.1	1,3
Allport P	Ac/RG	Liverpool	0.05	0.05	0.05	0,1,3
Burdin S	Ac/RG	Liverpool	0.2	0.2	0.2	1,4,5
Carroll JL	E/RG	Liverpool	0	0.2	0.15	4,5
Casse G-L	AP/RG	Liverpool	0.25	0.25	0.25	0,1,3
Greenshaw T	Ac/RG	Liverpool	0.2	0.2	0.2	0,3,4,5
Jones TJ	AP/RG	Liverpool	0.05	0.2	0.2	4,5
Sutcliffe P	E/RG	Liverpool	0	0.05	0.05	4,5
Tsurin I	E/RG	Liverpool	0.5	0.5	0.5	1,3,5
Whitley M	T/RG	Liverpool	0	0	0.1	1,4,5
Wormald MP	T/RG	Liverpool	0.1	0.1	0.1	1,5
Workshop	T/RG	Liverpool	0.1	0.2	0.2	2,4
LSDC	T/Project	Liverpool	0	0.1	0.1	1,3
Workshop	T/Project	Liverpool	0.1	0.1	0.1	2,4
Da Via	Ac/RG	Manchester	0.6	0.6	0.6	0,2
Watts	Ac/RG	Manchester	0.1	0.1	0.1	0,2
Elvin	T/RG	Manchester	0.1	0.1	0.1	2,5
Freestone	E/RG	Manchester	0.1	0.1	0.1	2,5
Hasi	AP/RG	Manchester	0.8	0.8	0.8	2
Kelly	AP/RG	Manchester	0.1	0.1	0.1	2,5
Kolya	AP/RG	Manchester	0.1	0.1	0.1	2,3,5
Pater	Ph/RG	Manchester	0.2	0.2	0.2	2
Perry	T/RG	Manchester	0.05	0.05	0.05	2,5
Snow	Ph/RG	Manchester	0.1	0.1	0.1	2,4,5
Thompson	E/RG	Manchester	0.1	0.1	0.1	2,3,5
Da Via - buy out	Ac/Project	Manchester	0.3	0.3	0.3	0,2
A. Nomerotski	Ac/RG	Oxford	0.2	0.2	0.2	1,2
Eng Support	E/TD	RAL (TD)	0	0.3	0.5	3,4,5
M. Tyndel	Ac/PPD	RAL (PPD)	0.2	0.2	0.2	0,3
M. Gibson	E/PPD	RAL (PPD)	0	0	0.2	0,3
J. Matheson	Ph/PPD	RAL (PPD)	0.25	0.25	0.25	3

Costs

ATLAS Upgrade WP5			Year 1	Year 2	Year 3	Total
Staff						
HEI / Universities	Glasgow	Rolling Grant	43.4	43.8	46.9	134.1
		New Money	3.5	3.6	3.6	10.7
	Birmingham	Rolling Grant	-	-	-	-
		New Money	1.0	2.0	2.0	5.0
	Oxford	Rolling Grant	4.4	4.4	4.4	13.1
		New Money				-
	Manchester	Rolling Grant	141.5	141.5	141.5	424.5
		New Money	65.1	65.1	65.1	195.3
	Liverpool	Rolling Grant	75.0	102.0	110.4	287.4
		New Money				-
Total HEI / University Staff						
Total - Rolling Grant			264.3	291.7	303.1	859.1
Total - New Money			69.6	70.7	70.8	211.0
STFC		RAL PPD	50.3	52.1	73.3	175.7
		TD	-	22.4	38.6	60.9
Total STFC Staff			50.3	74.5	111.9	236.7
Recurrent						
Equipment			254.3	328.0	235.5	817.8
Consumables			23.0	22.5	17.5	63.0
Travel			29.01	29.01	29.01	87.0
Other						
Total Recurrent			306.3	379.5	282.0	967.8
Project total			690.5	816.3	767.8	2,274.6
Working Margin			25.4	32.8	23.5	81.8
Contingency						
Project Total (including margin/contingency)			715.9	849.1	791.3	2,356.4
Rolling grant			264.3	291.7	303.1	859.1
Project Total (less rolling grant)			451.7	557.5	488.2	1,497.3

WP6: Level-1 Calorimeter Trigger

Overview

See Section 3.3.

Task Summary

WP6	Work-package Manager	Mark Thomson
0	organisation	Engineering and overall project management
1	design activity 1	Understand impact of pileup
2	design activity 2	Study physics impact of topological algorithms
3	design activity 3	Understand Phase-II requirements
4	hardware activity 1	Develop conceptual design of Topological Processor
5	hardware activity 2	Design and fabricate Topological Processor boards
6	hardware activity 3	Commission and test Topological Processor
7	hardware activity 4	Study technology options for Phase-II
8	firmware activity 1	Upgrade CPM and CMM firmware
9	firmware activity 2	Implement new data formats in RODs
10	firmware activity 3	Develop firmware for TP merger boards
11	firmware activity 4	Develop firmware for Topological Processor algorithms
12	firmware activity 5	Develop online software

Inputs

Previous ATLAS TDR Monte Carlo studies
Improved understanding of minimum-bias events based on first ATLAS data

Major outputs and deliverables

Task(s)	Output/Deliverable
1	Understanding of L1Calo trigger rates at $> 10^{34}\text{cm}^{-2}\text{s}^{-1}$
2	Definition of design of useful topological triggers
4	Design of topological processor architecture
5,6	Commissioned topological processor in L1 trigger
8	Increased flexibility in current L1Calo system
9,10,11	Implementation of topological algorithms into L1Calo
12	Operational online software/monitoring for topological triggers
3	Understanding of requirements of L1Calo trigger at $10^{35}\text{cm}^{-2}\text{s}^{-1}$
7	Understanding possible technology choices for Phase-II upgrade

Major Milestones

Milestone	proposed date
Identification of topological trigger algorithms	Jun 2011
Design of Topological Processor	Mar 2011
Prototype Topological Processor	Jun 2012
Delivery of system to CERN	Jun 2013

Long Term Time Profile

2010–2012	MC studies of L1Calo at Phase-I luminosity, and identification of most effective topological triggers
2010–2011	Firmware upgrades to allow current system to operate at higher backplane speeds
2011–2012	Design and construction of Topological Processor prototype
2012–2013	Manufacture and testing of Topological Processor
2013 (end)	Delivery of Topological Processor components to CERN
2014	Installation and commissioning of Phase-I upgrade
2012–2013	MC studies of requirements of L1Calo at Phase-II luminosities
2014	Decision on design of Phase-II upgrade
2014–2017	Prototyping of Phase-II upgrade
2017–2019	Construction of full Phase-II L1Calo
2019	Installation of Phase-II L1Calo upgrade
2020	Commissioning of Phase-II upgrade

Staff

Name	Type/Funding	Institute	10/11	11/12	12/13	Activities
D. Charlton	Ac/RG	Birm	0	0.1	0.2	3,4
A. Watson	Ac/RG	Birm	0.1	0.1	0.2	1,2,3,4
J. Bracinik	AP/RG	Birm	0	0.2	0.3	6,8,10,12
S. Hillier	AP/RG	Birm	0.1	0.3	0.3	4,6
M. Krivda	E/Project	Birm	0	0.2	0.6	5,6,7
R. Staley	E/RG	Birm	0.3	0.4	0.5	4,5,6,7,8
S. Pyatt	T/RG	Birm	0.1	0.15	0.15	5,6,7
X. Serghi	T/Project	Birm	0.1	0.15	0.15	5,6
M. Thomson	Ac/RG	Camb	0.3	0.3	0.3	0,1,2,3
B. Hommels	Ac/RG	Camb	0.1	0.1	0.1	10,11
M. Goodrick	E/RG	Camb	0.1	0.1	0.3	4
S. Sigurdsson	T/Project	Camb	0.4	0.4	0.4	5,6
J. Chapman	Ph/RG	Camb	0.3	0.3	0.3	1,2,3
N. Gee	Ac/Project	RAL (PPD)	0.4	0.4	0.5	0,4,6,9,12
R. Middleton	Ac/Project	RAL (PPD)	0.65	0.65	0.65	1,2,3
B. Barnet	AP/Project	RAL (PPD)	0.35	0.5	0.6	4,6,7,12
D. Sankey	PP/Project	RAL (PPD)	1	1	1	4,6,9,10,11
W. Qian	E/Project	RAL (PPD)	0.7	0.8	0.85	5,6,7,10,11
Prieur/cont.	E/Project	RAL (PPD)	0.2	0.2	0.2	1,2,4,6
I Brawn	E/Project	RAL (ID)	0.8	0.8	0.8	0,4,5,8
A Davis	E/Project	RAL (ID)	0	1	1	5,6,10,11
Design Engineer 1	E/Project	RAL (ID)	0	1	1	10,11
Design Engineer 2	E/Project	RAL (ID)	0	0.3	0.2	10,11
Draw. Off. Eng.	E/Project	RAL (ID)	0.15	0.6	0.4	5
ESS Test Engineer	E/Project	RAL (ID)	0	0.2	0.2	6
ESS Test support	E/Project	RAL (ID)	0	0.2	0.1	6
L. Cerrito	Ac/RG	QMUL	0.12	0.12	0.12	1,2,3,7
E. Rizvi	Ac/RG	QMUL	0	0.1	0.2	2,3
M. Landon	AP/RG	QMUL	0.4	0.6	0.8	4,6,7,12
New Post 1	AP/Project	Birm	1	1	1	8,9,10,11
New Post 2	Ph/Project	Camb/QMUL	1	1	1	1,2,3
Total			**8.9**	**13.5**	**14.6**	

Effort by task

The required FTE of staff effort across main areas of activity are:

- 7.9 FTE-yrs (3.7 Ac, 3.9 Ph, 0.3 PP) for the L1Calo upgrade design studies.
- 14.6 FTE-yrs (0.8 Ac, 1.5 PP, 2.8 AP, 7.5 E, 2 T) for the conceptual design, construction and testing of the topological processor.
- 12.9 FTE-yrs (0.8 Ac, 1.8 PP, 4.6 AP, 5.6 E) for firmware/online software development for the topological processor and the associated modifications to the existing L1Calo system.

- 1.0 FTE-yrs (0.5 Ac, 0.5 E) of effort to provide ATLAS-wide TDAQ upgrade leadership and the scientific/engineering coordination of the UK project.

The estimated WP6 staff effort is broken down into tasks below. The staff effort for firmware and hardware activities is well matched to the likely needs based on the UK L1Calo experience of designing boards of a similar complexity to those proposed for the Topological Processor.

Task	Description	FTE-yrs						
		Ac	Ph	PP	AP	E	T	Total
0	Organisation	0.5	–	–	–	0.5	–	**1.0**
1	Understand pileup	0.9	1.3	0.2	–	–	–	**2.4**
2	Physics algorithms	1.4	1.8	0.1	–	–	–	**3.3**
3	Phase-II requirements	1.4	0.9	–	–	–	–	**2.3**
4	TP conceptual design	0.3	–	0.5	0.5	0.9	–	**2.1**
5	TP boards	–	–	–	–	3.6	1.4	**5.0**
6	TP commissioning	0.4	–	1.0	1.5	2.1	0.4	**5.4**
7	Phase-II technology	0.1	–	–	0.9	1.0	0.2	**2.1**
8	CPM/CMM upgrades	–	–	–	1.2	1.2	–	**2.4**
9	ROD data formats	0.3	–	0.6	0.5	–	–	**1.3**
10	TP merger boards	0.2	–	0.6	0.9	2.2	–	**3.9**
11	TP algorithms	0.2	–	0.6	0.9	2.2	–	**3.9**
12	Online software	0.8	–	–	1.3	–	–	**1.5**
Total L1Calo Effort		5.8	3.9	3.6	7.5	13.6	2.0	**36.4**

Costs

Costs

From the design and construction of the current L1Calo system it is possible to provide a reasonable estimate of the likely resources (both equipment and engineering effort) required for the proposed L1Calo Phase-I upgrade project. Based on the current assumptions for the Topological Processor, the total cost of the hardware, including prototypes and test rigs at four international sites plus the production system and spares, is estimated to be approximately £410k. The assumptions used to make this estimate are shown below. The UK contribution is taken to be half this total cost.

Item	Cost	Prototypes & Test Rigs		Production System		Spares	
	£k	Qty	Total £k	Qty	Total £k	Qty	Total £k
New CMM	4	4	16	12	48	4	16
ATCA Crate	4	4	16	1	4	1	4
ATCA B/P	6	4	24	1	6	2	12
Input Boards	6	4	24	4	24	2	12
Processing Boards	6	4	24	2	12	2	12
Timing Board 6	4	24	1	6	2	12	
Readout Board 6	4	24	1	6	2	12	
CPU, Cables Misc	10	4	40	1	10	2	20
Subtotal			192		116		100
Grand Total in £k							**408**

On this basis the requested UK contribution to the Phase-I upgrade hardware is £200k. Due to the uncertainty in the final design of the system we request a working allowance of 30 %, i.e. £60k. We request an additional £25k for R&D into high-speed backplanes and links for the initial Phase-I and -II upgrades. Finally, we request £50k for travel during the period of the proposal. This is in addition to travel for running the existing L1Calo, and is needed for international meetings, and for work on testing, installation and commissioning of the upgraded system. It includes, per annum, approximately 30 international trips (£500 each) and ten visits (£200 each) for travel inside the UK to cover work on tests and effective collaboration between the UK institutes in the project.

Item	Request £k
Phase-I upgrade hardware	200
Working allowance	60
High-speed backplane R&D	25
Travel	50
Total	**335**

ATLAS Upgrade WP6			Year 1	Year 2	Year 3	Total
Staff						
HEI / Universities	Birmingham	Rolling Grant	34.6	84.8	109.7	229.1
		New Money	80.4	90.2	104.3	274.9
	QMUL	Rolling Grant	36.6	55.6	74.7	166.9
		New Money	34.5	34.5	34.5	103.6
	Cambridge	Rolling Grant	40.6	40.9	60.0	141.5
		New Money	47.7	48.5	49.0	145.2
Total HEI / University Staff						
Total - Rolling Grant			111.8	181.3	244.3	537.5
Total - New Money			162.6	173.3	187.8	523.7
STFC		RAL PPD	319.6	368.2	397.5	1,085.4
		TD	91.0	376.9	356.2	824.1
Total STFC Staff			410.7	745.1	753.7	1,909.4
Recurrent						
Equipment			67.0	67.0	91.0	225.0
Consumables			-	-	-	-
Travel			16.7	16.7	16.7	50.0
Other						-
Total Recurrent			83.7	83.7	107.7	275.0
Project total			768.8	1,183.4	1,293.5	3,245.6
Working Margin			20.1	20.1	19.8	60.0
Contingency						
Project Total (including margin/contingency)			788.9	1,203.5	1,313.3	3,305.6
Rolling grant			111.8	181.3	244.3	537.5
Project Total (less rolling grant)			677.1	1,022.1	1,069.0	2,768.1

WP7: Level-1 Track Trigger

Overview

See Section 3.4.

Task Summary

WP7 Manager:	N. Konstantinidis.
0	Management - Coordination
1	Off-detector hardware - overall design
2	RoIMapper - Design and develop a demonstrator module
3	Evaluate hardware technologies for trigger processor boards
4	Provide L1Track requirements/constraints and give input to on-detector electronics and readout
5	Develop simulation tools for high pile-up L1Track studies
6	Perform simulation studies for overall design
7	Perform pattern recognition studies
8	Develop and use discrete event simulation in the L1Track design
9	Perform physics/trigger rates studies to determine the benefits/optimal use of L1Track

Inputs

Milestones

Milestone	Date
DES framework operational	Q1/2011
Rates for lepton triggers with/without L1Track	Q2/2011
First implementation of pattern recognition adapted for hardware	Q4/2011
Complete design of RoIMapper	Q2/2012

Major outputs and deliverables

L1 trigger rates at $10^{35}\text{cm}^{-2}\text{s}^{-1}$ to build the physics case for L1Track
Overall design of the L1Track system based on the regional readout concept
A hardware demonstrator system for the RoIMapper

Long Term Time Profile

2010-11:	Demonstration of the benefits of L1Track and proof of principle for regional readout
2011-13:	Conceptual design, evaluation of technologies, development of hardware demonstrators
2013-14:	Technical Design Report
2014-16:	Detailed design and prototyping
2016-19:	Production, testing, system assembly
2019-20:	Installation at CERN and system commissioning

Staff

	Role/Funding	**Institute**	**2010-11**	**2011-12**	**2012-13**	**Tasks**
G Crone	Pr/RG	UCL	0.30	0.30	0.30	3,6,8
M Kauer	AP/Project	UCL		0.20	0.20	1,3
N Konstantinidis	Ac/RG	UCL	0.30	0.30	0.30	0,1,6,7
M Lancaster	Ac/RG	UCL	0.20	0.30	0.30	0,3
M Warren(*)	E/RG	UCL	0.30	0.30	0.30	1,2,4
Elec. Eng.	E/Project	UCL	0.25	0.25	0.25	1,2,4
V Boisvert	Ac/RG	RHUL	0.20	0.20	0.20	0,6,9
B Green(*)	E/Project	RHUL	0.30	0.40	0.40	1,2,4
A Misiejuk	AP/RG	RHUL	0.00	0.10	0.10	2
P Teixeira-Dias	Ac/RG	RHUL	0.10	0.10	0.10	0,9
V Bartch	PP/Other	Sussex	0.30	0.30	0.30	5,6
A De Santo	Ac/RG	Sussex	0.10	0.20	0.20	0,6,9
F Salvatore	Ac/RG	Sussex	0.10	0.10	0.20	0,6,9
T Affolder(*)	AP/RG	Liverpool	0.10	0.10	0.10	4
P Allport(*)	Ac/RG	Liverpool	0.05	0.05	0.05	0,4
A Greenall(*)	E/RG	Liverpool	0.10	0.10	0.10	4
B King(*)	Ph/Project	Liverpool	0.15	0.15	0.15	5,8
A Mehta	Ac/RG	Liverpool	0.00	0.10	0.10	6,8
P Sutcliffe(*)	E/RG	Liverpool	0.00	0.05	0.05	4
D Costanzo	Ac/RG	Sheffield	0.05	0.05	0.10	0,5
M Sutton	Ph/Project	Sheffield	0.10	0.10	0.20	3,7
New PDRA1	Ph/Project	UCL	1.00	1.00	1.00	1,3,7
New PDRA2	Ph/Project	RHUL/Sussex	1.00	1.00	1.00	5,6,9
Total			**5.00**	**5.75**	**6.00**	

(*) Staff marked with asterisk are also involved in the Tracker Upgrade WPs of this proposal.

The detailed distribution of FTEs to the tasks is summarised in the following table:

Task	Ac	E	Ph	PP	PR	AP	Total
0	1.2						1.2
1	0.2	0.9	1.0			0.2	2.3
2		0.9				0.2	1.1
3	0.5		1.2		0.3	0.2	2.3
4	0.1	1.4				0.3	1.8
5	0.2		1.2	0.5			1.9
6	0.8		1.0	0.4	0.3		2.5
7	0.2		1.2				1.4
8	0.1		0.2		0.3		0.6
9	0.8		1.0				1.8
Total	4.1	3.2	6.8	0.9	0.9	0.9	16.8

In summary, the required FTE for the different broad areas of activity, summed over the 3-year period of the proposal, are:

- 4.1 FTE of electronics engineer and applied physicist expertise for the overal system design, the design of the RoIMapper and the development of a hardware demonstrator, and to interact with the Tracker upgdrade WPs, providing inputs about the L1Track requirements and contributing to the work for implementing the additional required functionality;
- 3.8 FTE of physicist and programmer expertise to develop the tools for the L1Track simulation studies and the Discrete Event Simulation framework;
- 7.7 FTE of physicist and academic effort to perform physics simulation studies and develop pattern recognition suitable for hardware implementation; these activities will provide the physics justification for L1Track, as well as all the necessary inputs to the full system design; and
- 1.2 FTE of academic effort to provide ATLAS-wide leadership for L1Track and coordinate the UK project.

Costs

Item	Cost £k
2 HSIO development boards at £3k each	6
4 Interface Boards at £1k	4
Hybrid prototype construction	15
Miscellaneous	1
Travel £15k/year	45
Total	**71**

ATLAS Upgrade WP7			**Year 1**	**Year 2**	**Year 3**	**Total**
Staff						
HEI / Universities	Liverpool	Rolling Grant	26.2	31.0	31.5	88.7
		New Money				-
	Sheffield	Rolling Grant	4.6	1.0	1.9	7.5
		New Money	3.6	7.4	15.0	26.0
	Sussex	Rolling Grant	3.5	9.7	7.0	20.2
		New Money	35.4	35.8	36.3	107.5
	UCL	Rolling Grant	37.2	39.3	39.3	115.8
		New Money	92.6	101.1	101.1	294.8
	RHUL	Rolling Grant	5.7	14.9	15.2	35.8
		New Money	53.4	59.0	60.1	172.5
Total HEI / University Staff						
Total - Rolling Grant			77.1	95.9	95.0	268.0
Total - New Money			185.0	203.3	212.4	600.7
STFC						
Total STFC Staff			-	-	-	-
Recurrent						
Equipment			5.0	6.0	15.0	26.0
Consumables			-	-	-	-
Travel			15.0	15.0	15.0	45.0
Other			-	-	-	-
Total Recurrent			20.0	21.0	30.0	71.0
Project total			282.1	320.2	337.4	939.7
Working Margin			0.5	0.6	1.5	2.6
Contingency						
Project Total (including margin/contingency)			282.6	320.8	338.9	942.3
Rolling grant			77.1	95.9	95.0	268.0
Project Total (less rolling grant)			205.5	224.9	243.9	674.3

WP8: High-Level Trigger

Overview

See Section 3.5.

The WP8 work falls into the three areas described in Section 3.5: Tracking Software, Selection Software and Menus, and Dataflow and Farms. The main outputs from this work are upgrades to key components of the trigger software for Phase-I and performance measurements that will provide vital input to inform crucial decisions on the Trigger Selection Strategy and Trigger Architecture for Phase-II.

Task Summary

WP8 Manager:	John Baines.
0	Management of the HLT Upgrade project
1	Optimise HLT ID Tracking software for Phase-I
2	Optimise electron, muon and tau trigger selections for Phase-I
3	Upgrade Trigger steering software
4	Reconfigure and tune trigger farm hardware for Phase-I
5	Perform simulation studies to optimise the trigger selection strategy for Phase-II
6	Develop and optimise the trigger architecture for Phase-II

Staff Summary by Task

Task		FTEyears			
		Ac	Ph	PP	Total
0	Organisation	0.3			0.3
1	HLT Tracking Software	0.75	1.0	1.7	3.55
2,3,5	HLT Selection Software	1.45	1.45	2.7	5.7
4,6	Dataflow and Farms	0.85		0.4	1.25

Inputs

Production of simulated datasets for Phase-I and Phase-II.
Improvements and optimisations of the offline tracking software used at the Event Filter.
Decision on whether to proceed with FTK.

Notes:

1. The production of the simulated datasets relies on work by the ATLAS core software and production teams.

2. The work to upgrade the HLT software does not depend on the FTK decision, but the final optimisation of the system is influenced by this choice. A decision on whether to proceed to a Technical Design Report is expect in 2010.

Major outputs and deliverables

Task(s)	output/deliverable
1	HLT Tracking software upgraded and optimised for Phase-I
2	Electron, muon and tau trigger selections optimised for Phase-I
3	Trigger steering software modified for the LVL1 upgrade, for FTK and to provide the option of full event processing.
4	HLT farm hardware and software tuned for Phase-I.
5	Performance measurements with the upgraded trigger selection software forming the input to critical choices for the Trigger Strategy for Phase-I and Phase-II.
6	Results from dataflow modeling as input to key decisions on the Trigger Architecture for Phase-II.

Major Milestones

Task	Milestone	date
1	HLT Tracking Code updated for IBL and FTK	Mar 2011
2	Trigger Selections and Menus Defined for MC Productions	Dec 2011
3	Steering Code updated for upgraded LVL1, FTK and full-event reconstruction	Mar 2012
1	Measurements made of speed-up from using GPU for L2 tracking	Dec 2012
2	HLT Selections updated for upgraded LVL1	Dec 2012
1	HLT Tracking code optimised for Phase-I	Mar 2013
5	Performance measurements available for Phase-II selections	Dec 2013
6	Results available from Dataflow modeling	Dec 2013
2	HLT Selections Optimised for Phase-I	Jun 2014

Long Term Time Profile

2010–2014	Upgrade ID trigger software, upgrade selection software, study architectural choices, study different trigger strategies
2014(end)	Deploy phase-I upgrade
2014	Major choices made for Trigger architecture and Trigger strategy
2014–2015	Trigger Software Design
2016–2018	Trigger Software implementation
2019–2019	Testing & commissioning
2019–2020(beg.)	Software deployed ready for running mid/late 2020

Staff

	Type	**Institute**	**2010/11**	**2011/12**	**2012/13**	**Activities**
A Oh	Ac/Other	Man	0.1	0.1	0.1	2
P Bell	Ph/RG	Man		0.1	0.1	2
J Masik	Ph/RG	Man	0.1	0.1	0.2	1
M Owen	Ph/RG	Man		0.1	0.1	2
T Wengler	Ac/RG	Man	0.1	0.1	0.1	2
U Yang	Ac/RG	Man	0.1	0.1	0.1	2
S. Farrington	Ac/Other	Oxford	0.1	0.1	0.05	2
F Wickens	Ac/Project	RAL (PPD)	0.2	0.3	0.5	0,4,6
J Baines	Ac/Project	RAL (PPD)		0.2	0.2	0,1,5
M Wielers	Ac/Project	RAL (PPD)		0.1	0.1	2
Kirk	Ph/Project	RAL (PPD)		0.1	0.15	2
D Emeliyanov	PP/Project	RAL (PPD)		0.1	0.1	1
S Burke	PP/Project	RAL (PPD)	0.5	0.5	0.5	1
S George	PP/RG	RHUL			0.1	3
R Goncalo	Ph/RG	RHUL			0.1	5
E Nurse	Ac/Other	UCL	0.2	0.2	0.2	1,5
New Post 1	Ph/Project	UCL	0.5	0.5	0.5	1,5
New Post 2	PP/Project	RHUL	1	1	1	3,5,6
Total			**2.9**	**3.7**	**4.2**	

Costs

Item	Cost £k
Contribution to CERN development/evaluation system	10
UK development/evaluation systems	10
Travel £7k/year	21
Total	**41**

ATLAS Upgrade WP8			**Year 1**	**Year 2**	**Year 3**	**Total**
Staff						
HEI / Universities	UCL	Rolling Grant	-	-	-	-
		New Money	40.4	40.4	40.4	121.3
	RHUL	Rolling Grant	0.0	-	9.0	9.0
		New Money	78.4	79.0	79.9	237.3
	Manchester	Rolling Grant	11.2	30.3	37.9	79.4
		New Money				
	Oxford	Rolling Grant	8.6	8.8	4.4	21.8
		New Money				-
Total HEI / University Staff						
Total - Rolling Grant			19.8	39.1	51.3	110.2
Total - New Money			118.8	119.4	120.4	358.6
STFC		RAL PPD	66.8	123.4	154.8	344.9
Total STFC Staff			66.8	123.4	154.8	344.9
Recurrent						
Equipment			-	15.0	5.0	20.0
Consumables			-	-	-	-
Travel			7.0	7.0	7.0	21.0
Other			-	-	-	-
Total Recurrent			7.0	22.0	12.0	41.0
Project total			212.5	303.9	338.4	854.8
Working Margin			-	3.8	1.3	5.0
Contingency						
Project Total (including margin/contingency)			212.5	307.6	339.7	859.8
Rolling grant			19.8	39.1	51.3	110.2
Project Total (less rolling grant)			192.7	268.6	288.4	749.6

WP9: Software, Computing and Physics Studies

Overview

See Section 4.

This work package will deliver several distinct, but interconnected, simulation outcomes that are closely related to the tracking detector and trigger developments. Initial tools will be provided, making short term adjustments to fit the simulation and reconstruction into currently available processor memories. Detailed simulations will be provided for both the short strips and the pixel detectors, integrating with the inner B-layer. These advanced codes must take the simulation of high pile-up into account in their design. Equally, the radiation environment must be accurately modelled, using real data and GEANT simulations. The radiation environment will also depend on the position and geometry of the upgraded ATLAS sub-detectors. At each stage, the full simulation will be captured in geometry and parameterised response in the ATLFAST fast simulation. The geometries will also be incorporated into the ATLANTIS event display to allow detailed studies of individual events, tracks and clusters.

In order to effect the simulation and reconstruction, and to prepare the computing for the post-TDR detector, framework and infrastructure developments must occur. Various concurrent programming techniques will be applied to speedup the code and minimise the memory profile, allowing up to about 16 cores per processor to be used more efficiently. In preparation for the next phase, more detailed parallelisation of the code and virtualisation will be applied to example parts of the ATLAS code base and the performance and bottlenecks evaluated. At the same time, the potential of GPUs for other tasks (including analysis, where throughput looks particularly challenging) will be investigated using test cases. This will allow the appropriate technology choices to be made for the TDR and later development. As part of these technological evaluations, the performance on the Grid is important, as is the interaction with local storage; a test bed (with appropriate monitoring for performance evaluation) will be constructed and CE developments for GPU usage made.

Task Summary

WP9 Manager:	R Jones
1	Initial tools for large pile-up studies
2	Short strip simulation
3	Pixel simulation
4	Fast simulation
5	Radiation simulation verification
6	Visualisation
7	Near term computing and optimisation
8	GPU and parallelisation (simulation and tracking)
9	Analysis parallelisation
10	Virtualisation
11	GPU/CE development

Inputs

Previous ATLAS upgrade Monte Carlo studies
Existing leading role in radiation environment models
Leadership in Atlfast
Leadership in Atlantis visualisation
Tracking software development
Grid environment developments

Major outputs and deliverables

Task(s)	output/deliverable
1,2,3,4	Effective detector simulation at $> 10^{34}\text{cm}^{-2}\text{s}^{-1}$
5	Radiation environment model
6	Visualisation for upgrade detector
7	Near term computing optimisation for upgrade studies
8, 11	GPU options demonstrated for future ATLAS computing
8, 9, 10	Parallelisation options demonstrated for future ATLAS computing

Major Milestones

Milestone	proposed date
Initial tools for large pile-up	Sep 2011
Upgrade short simulation default version	Apr 2013
Pixel simulation release for TDR	Jan 2013
Atlfast release for TDR	Apr 2013
Final Fluka release	Dec 2012
Radiological assessment	Apr 2013
Atlantis release for TDR	Apr 2013
GPU application optimised release	Jan 2013
Parallelised analysis study	Apr 2012
Virtualisation report	Apr 2013
Performant GPU on grid	Feb 2013

Longer Term Time Profile

2010–2011	Initial tools preparations; first radiation environment validation; Grid CE and GPU testbed established
2012-2013	Deployment of upgrade simulation version; reconstruction and analysis parallelised example; Atlfast and Atlantis for TDR
2014–2014	Evaluation of technologies and architectures for computing use cases
2014–2015	Upgrade Reconstruction Software & Data Design
2016–2018	Upgrade Computing implementation
2019–2019	Testing & commissioning
2019–2020(beg.)	Running mid/late 2020

Staff Summary by Task

Task		FTE years				
		Ac	Ph	PP	Pr	Total
1	Initial tools for large pile-up studies	0.24	0.73	0.0	0.0	0.97
2	Short strip simulation	0.72	2.37	0.3	0.0	3.39
3	Pixel simulation	0.3	0.9	0.45	0.0	1.2
4	Fast simulation	0.15	0.3	0.0	0.0	0.9
5	Radiation simulation verification	0.3	2.75	0.45	0.0	3.05
6	Visualisation	0.15	0.3	0.22	0.0	0.9
7	Near term computing and optimisation	0.12	1.56	0.0	0.0	1.84
8	GPU and parallelisation (simulation and tracking)	0.15	1.55	0.0	0.0	1.73
9	Analysis parallelisation	0.12	0.948	0.16	0.0	1.236
10	Virtualisation	0.105	1.72	0.0	0.2	2.025
11	GPU/CE development	0.045	1.522	0.02	0.2	1.787

Staff Summary by Theme

Task		FTE years				
		Ac	Ph	PP	Pr	Total
1	Initial tools for large pile-up studies	0.2	0.7	0.0	0.0	1.0
2-4	Simulation	1.2	3.6	0.8	0.0	5.5
5	Radiation Environment	0.3	2.8	0.0	0.0	3.1
7,9	Optimisation and parallelisation	0.2	2.5	0.4	0.0	3.1
8,10,11	GPU and virtualisation	0.3	4.8	0.02	0.4	5.5

Staff

Name	Type/Funding	Institute	2010/11	2011/12	2012/13	Tasks
R. Batley	Ac/RG	Cambridge	0.10	0.10	0.10	2
P. Ward	PP/RG	Cambridge		0.10	0.20	2
P. Clark	Ac/RG	Edinburgh	0.10	0.10	0.10	1, 7-11
V. Martin	Ac/RG	Edinburgh		0.15	0.15	3
A. Buckley	Ph/Other	Edinburgh	0.20	0.20	0.20	1, 3, 9
A.N. Other	Ph/Project	Edinburgh	1.00	1.00	1.00	3,7,8,9
A. Washbrook	Ph/Other	Edinburgh	0.20	0.17	0.00	7
A. Washbrook	Ph/Project	Edinburgh		0.83	1.00	7-11
R. Jones	Ac/RG	Lancaster	0.10	0.10	0.10	8-11
P. Love	Pr/RG	Lancaster	0.15	0.15	0.15	10, 11
R. Henderson	PP/Project	Lancaster	0.15	0.15	0.15	7, 9, 11
A.N. Other	Ph/Project	Lancaster	1.00	1.00	1.00	10, 11
J. Tseng	Ac/RG	Oxford	0.07	0.07	0.07	1, 2
vice A. Abdesselam	Ph/Project	Oxford	0.70	0.70	0.70	1, 2
J. Ferrando	Ph/RG	Oxford		0.50	0.50	1, 2
J. Butterworth	Ac/RG	UCL	0.10	0.10	0.10	4, 6
P. Sherwood	PP/RG	UCL	0.20	0.20	0.20	4, 6
B. Waugh	PP/RG	UCL	0.10	0.10	0.10	4, 6
A.N. Other	Ph/Other	UCL	0.20	0.20	0.20	4, 6
D. Tovey	Ac/RG	Sheffield	0.10	0.10	0.10	5
I. Dawson	Ph/RG	Sheffield	0.10	0.20	0.20	5
L. Nicolas	Ph/RG	Sheffield	0.25	1.00	1.00	5

Costs

Item	Cost £k
UK development/evaluation systems	36
Travel £35k/year	105
Total	**141**

ATLAS Upgrade WP9			**Year 1**	**Year 2**	**Year 3**	**Total**
Staff						
HEI / Universities	Oxford	Rolling Grant	37.4	12.8	11.9	62.1
		New Money	22.0	88.3	89.1	199.4
	UCL	Rolling Grant	31.0	31.0	31.0	92.9
		New Money				-
	Cambridge	Rolling Grant	2.0	11.3	20.6	33.8
		New Money				-
	Sheffield	Rolling Grant	10.7	19.6	19.9	50.3
		New Money	22.0	88.3	89.1	199.4
	Edinburgh	Rolling Grant	3.2	7.9	8.0	19.0
		New Money	74.0	137.3	151.8	363.0
	Lancaster	Rolling Grant	49.8	61.0	65.3	176.1
		New Money	31.0	31.0	31.0	92.9
Total HEI / University Staff						
Total - Rolling Grant			133.9	143.5	156.7	434.1
Total - New Money			148.9	344.8	361.0	854.7
STFC						
Total STFC Staff			-	-	-	-
Recurrent						
Equipment			12.0	12.0	12.0	36.0
Consumables			3.0	3.0	3.0	9.0
Travel			35.0	35.0	35.0	105.0
Other						-
Total Recurrent			50.0	50.0	50.0	150.0
Project total			332.8	538.3	567.7	1,438.8
Working Margin			1.2	1.2	1.2	3.6
Contingency						
Project Total (including margin/contingency)			334.0	539.5	568.9	1,442.4
Rolling grant			133.9	143.5	156.7	434.1
Project Total (less rolling grant)			200.1	396.0	412.2	1,008.3

5.4 Cases for New Posts

To follow.

5.5 Risk Register

5.6 Gantt Charts

AFP Gantt Chart (WP1)

ID	Task Name	Duration	Start
1	FE-I4 Prototypes	103 days?	Thu 10/09/09
2	FE-I4 Module Prototypes	216 days?	Tue 02/02/10
3	FE-I4 Engineering Run	62 days?	Wed 01/12/10
4	FE-I4 Test Eng Run Samples	68 days?	Fri 25/02/11
5	FE-I4 Production Run	61 days?	Wed 01/06/11
6	FE-I4 QA of production	92 days?	Thu 10/09/09
7			
8	FE-I4 3D Sensor R&D	342 days?	Thu 10/09/09
9	FE-I4 3D Sensor AFP IBL prod	173 days?	Mon 03/01/11
10	FE-I4 Sensor Testing	88 days?	Thu 01/09/11
11	FE-I4 Sensor Bump bond	86 days?	Tue 03/01/12
12			
13	WP1 R&D	342 days?	Thu 10/09/09
14	WP1 TDR Write and Approval	342 days?	Thu 10/09/09
15	WP1 Pre-production	260 days?	Mon 03/01/11
16	WP1 Production	522 days?	Mon 02/01/12
17	WP1 Commissioning and Installation	435 days?	Wed 01/08/12
18			
19	WP1.1 Hamburg Pipe Design	189 days?	Thu 10/09/09
20	WP1.1 Build Pre-Production	130 days?	Wed 02/06/10
21	MS#1 HP design and pre-production complete	0 days	Fri 31/12/10
22	WP1.1 Safety Review	333 days?	Thu 10/09/09
23	WP1.1 Radiiation Review	333 days?	Thu 10/09/09
24	WP1.1 Liasion with LHC, background and em modelling	1189 days?	Thu 10/09/09
25	MS#2 Commission tests in LHC with LHC Machine Physicists	0 days	Mon 01/04/13
26			
27	WP1.2 FE-I4 Superlayer Design, build and test	87 days?	Thu 01/04/10
28	WP1.2 FE-I4 Mechanical Design	197 days?	Thu 01/04/10
29	WP1.2 Thermal Design, build and test	197 days?	Thu 01/04/10
30	MS#3 Thermal Design finished and prototype tested	0 days	Fri 01/04/11
31	WP1.2 Tracker local electronics	450 days?	Thu 10/09/09
32	WP1.2 LV/HV Design and build	452 days?	Thu 10/09/09
33	WP1.2 Construction of pre-production tracker station	151 days	Mon 03/01/11
34	MS#4 Pre-production station complete	0 days	Thu 01/09/11
35	WP1.2 Construction of tracker station Left-Front	109 days?	Wed 02/05/12
36	WP1.2 Construction of tracker Left-Back	109 days	Tue 02/10/12
37	WP1.2 Construction of tracker Right-Front	109 days	Mon 04/03/13
38	WP1.2 Construction of tracker Right-Back	109 days	Fri 02/08/13
39			
40	WP1.3 DAQ development and implementation	493 days?	Thu 10/09/09
41	WP1.3 Optical link,cables, development and build	320 days?	Thu 10/09/09
42	WP1.3 Monitor and slow control software	493 days?	Thu 10/09/09
43	MS#6 DAQ/DCS ready for pre-production tracker	0 days	Thu 01/09/11
44			
45	WP5.4 Build pre-production wire/bpm alignment	348 days?	Thu 01/04/10
46	MS#7 Pre-production alignment system ready	0 days	Thu 01/09/11
47	WP5.4 Build production wire/bpm alignent	218 days?	Thu 01/12/11
48	WP5.4 Calibration software development and test	1189 days	Thu 10/09/09
49			
50			
51	WP5.6 Assembly of preproduction detectior station	88 days?	Tue 02/08/11
52	MS#8 Pre-production station and 8m prototype ready	0 days	Tue 31/01/12
53	WP5.6 Installation of cables and optical fibres	85 days?	Mon 03/01/11
54	MS#5 Services installed in LHC tunnel	0 days	Thu 01/09/11
55	WP5.6 Assemble, commission, test pre-production station	174 days?	Fri 02/12/11
56	WP5.6 Assemble, commission and test production stations	522 days?	Tue 02/10/12
57	MS#9 1st production tracker ready for commission	0 days	Wed 31/10/12
58	MS#10 2nd production tracker ready for commission	0 days	Mon 01/04/13

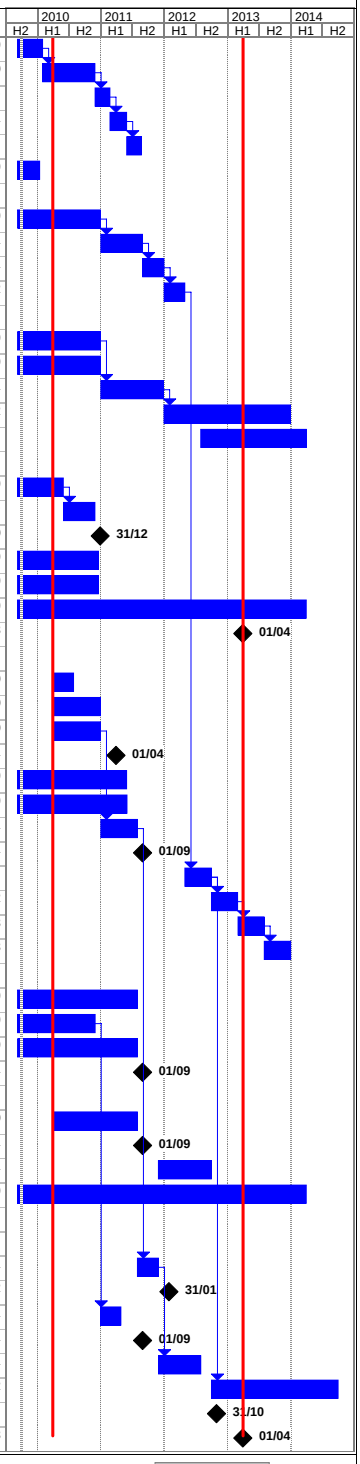

Project: ATLAS Forward Physics
Date: Wed 30/09/09

Task — Milestone — External Tasks
Split — Summary — External Milestone
Progress — Project Summary — Deadline

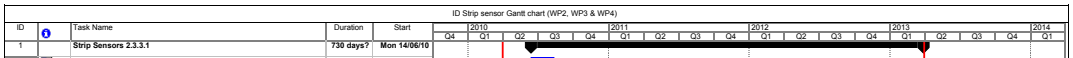

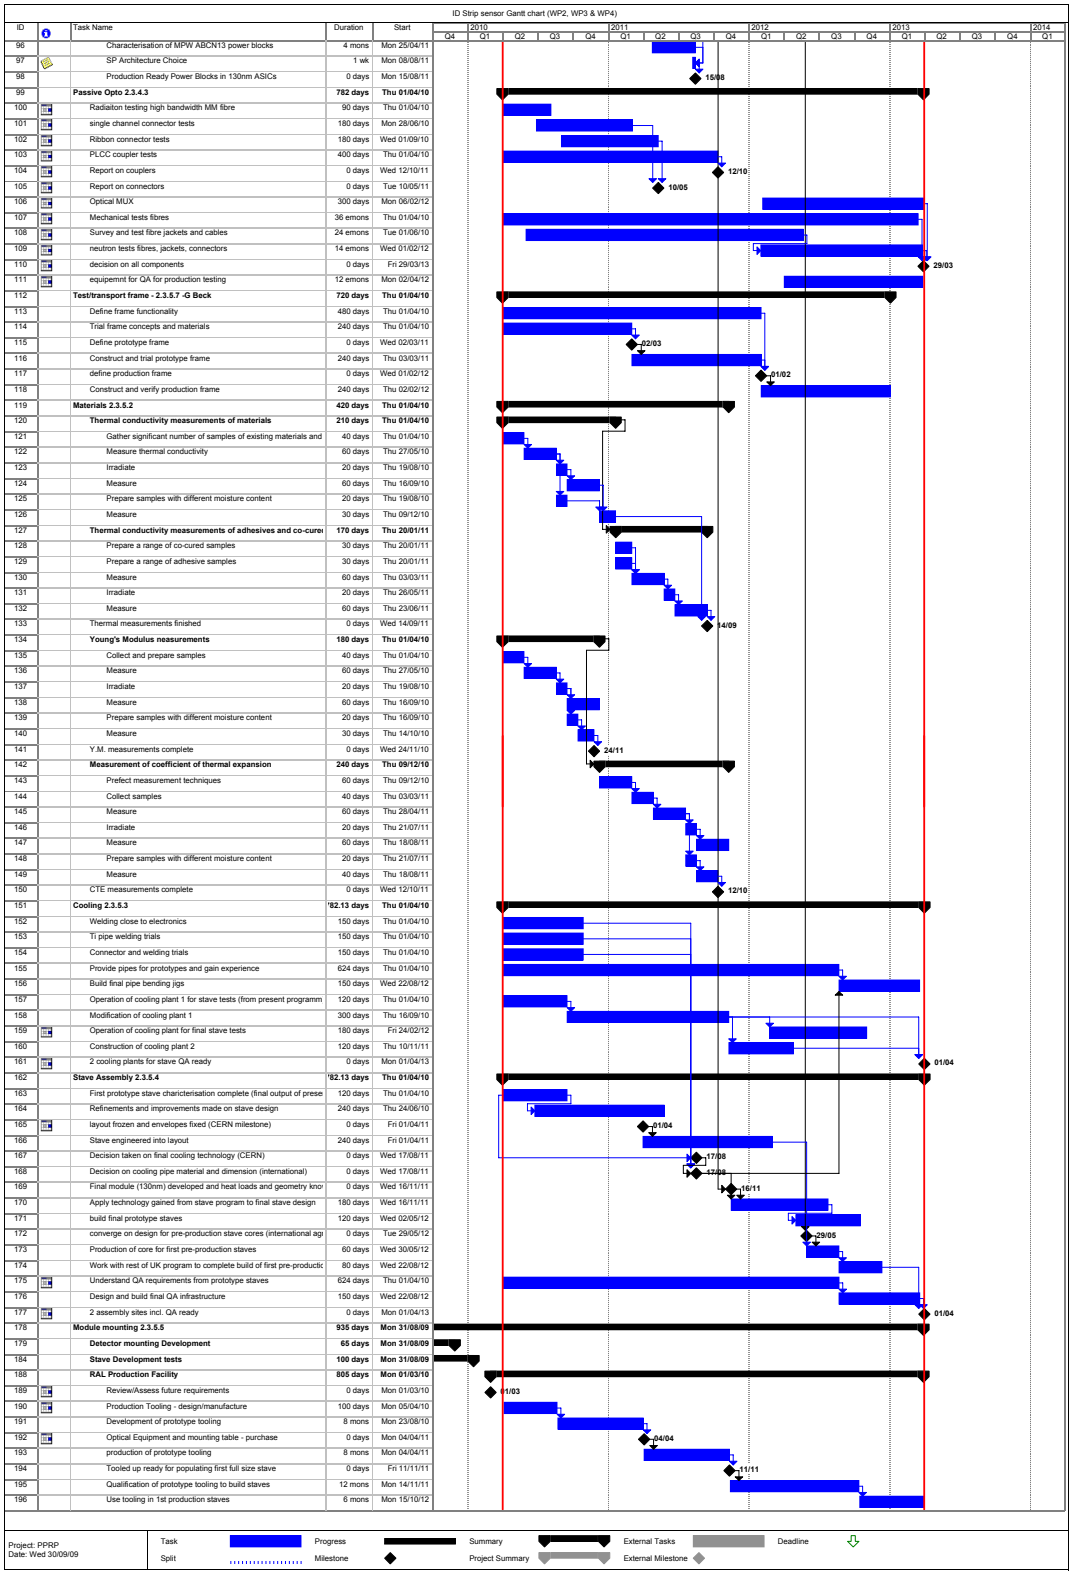

ID	Task Name	Duration	Start
96	Characterisation of MPW ABCN13 power blocks	4 mons	Mon 25/04/11
97	SP Architecture Choice	1 wk	Mon 08/08/11
98	Production Ready Power Blocks in 130nm ASICs	0 days	Mon 15/08/11
99	**Passive Opto 2.3.4.3**	**782 days**	**Thu 01/04/10**
100	Radiaiton testing high bandwidth MM fibre	90 days	Thu 01/04/10
101	single channel connector tests	180 days	Mon 28/06/10
102	Ribbon connector tests	180 days	Wed 01/09/10
103	PLCC coupler tests	400 days	Thu 01/04/10
104	Report on couplers	0 days	Wed 12/10/11
105	Report on connectors	0 days	Tue 10/05/11
106	Optical MUX	300 days	Mon 06/02/12
107	Mechanical tests fibres	36 emons	Thu 01/04/10
108	Survey and test fibre jackets and cables	24 emons	Tue 01/06/10
109	neutron tests fibres, jackets, connectors	14 emons	Wed 01/02/12
110	decision on all components	0 days	Fri 29/03/13
111	equipemnt for QA for production testing	12 emons	Mon 02/04/12
112	**Test/transport frame - 2.3.5.7 -G Beck**	**720 days**	**Thu 01/04/10**
113	Define frame functionality	480 days	Thu 01/04/10
114	Trial frame concepts and materials	240 days	Thu 01/04/10
115	Define prototype frame	0 days	Wed 02/03/11
116	Construct and trial prototype frame	240 days	Thu 03/03/11
117	define production frame	0 days	Wed 01/02/12
118	Construct and verify production frame	240 days	Thu 02/02/12
119	**Materials 2.3.5.2**	**420 days**	**Thu 01/04/10**
120	**Thermal conductivity measurements of materials**	**210 days**	**Thu 01/04/10**
121	Gather significant number of samples of existing materials and	40 days	Thu 01/04/10
122	Measure thermal conductivity	60 days	Thu 27/05/10
123	Irradiate	20 days	Thu 19/08/10
124	Measure	60 days	Thu 16/09/10
125	Prepare samples with different moisture content	20 days	Thu 19/08/10
126	Measure	30 days	Thu 09/12/10
127	**Thermal conductivity measurements of adhesives and co-cure**	**170 days**	**Thu 20/01/11**
128	Prepare a range of co-cured samples	30 days	Thu 20/01/11
129	Prepare a range of adhesive samples	30 days	Thu 20/01/11
130	Measure	60 days	Thu 03/03/11
131	Irradiate	20 days	Thu 26/05/11
132	Measure	60 days	Thu 23/06/11
133	Thermal measurements finished	0 days	Wed 14/09/11
134	**Young's Modulus neasurements**	**180 days**	**Thu 01/04/10**
135	Collect and prepare samples	40 days	Thu 01/04/10
136	Measure	60 days	Thu 27/05/10
137	Irradiate	20 days	Thu 19/08/10
138	Measure	60 days	Thu 16/09/10
139	Prepare samples with different moisture content	20 days	Thu 16/09/10
140	Measure	30 days	Thu 14/10/10
141	Y.M. measurements complete	0 days	Wed 24/11/10
142	**Measurement of coefficient of thermal expansion**	**240 days**	**Thu 09/12/10**
143	Prefect measurement techniques	60 days	Thu 09/12/10
144	Collect samples	40 days	Thu 03/03/11
145	Measure	60 days	Thu 28/04/11
146	Irradiate	20 days	Thu 21/07/11
147	Measure	60 days	Thu 18/08/11
148	Prepare samples with different moisture content	20 days	Thu 21/07/11
149	Measure	40 days	Thu 18/08/11
150	CTE measurements complete	0 days	Wed 12/10/11
151	**Cooling 2.3.5.3**	**'82.13 days**	**Thu 01/04/10**
152	Welding close to electronics	150 days	Thu 01/04/10
153	Ti pipe welding trials	150 days	Thu 01/04/10
154	Connector and welding trials	150 days	Thu 01/04/10
155	Provide pipes for prototypes and gain experience	624 days	Thu 01/04/10
156	Build final pipe bending jigs	150 days	Wed 22/08/12
157	Operation of cooling plant 1 for stave tests (from present programm	120 days	Thu 01/04/10
158	Modification of cooling plant 1	300 days	Thu 16/09/10
159	Operation of cooling plant for final stave tests	180 days	Fri 24/02/12
160	Construction of cooling plant 2	120 days	Thu 10/11/11
161	2 cooling plants for stave QA ready	0 days	Mon 01/04/13
162	**Stave Assembly 2.3.5.4**	**'82.13 days**	**Thu 01/04/10**
163	First prototype stave charicterisation complete (final output of prese	120 days	Thu 01/04/10
164	Refinements and improvements made on stave design	240 days	Thu 24/06/10
165	layout frozen and envelopes fixed (CERN milestone)	0 days	Fri 01/04/11
166	Stave engineered into layout	240 days	Fri 01/04/11
167	Decision taken on final cooling technology (CERN)	0 days	Wed 17/08/11
168	Decision on cooling pipe material and dimension (international)	0 days	Wed 17/08/11
169	Final module (130nm) developed and heat loads and geometry kno	0 days	Wed 16/11/11
170	Apply technology gained from stave program to final stave design	180 days	Wed 16/11/11
171	build final prototype staves	120 days	Wed 02/05/12
172	converge on design for pre-production stave cores (international ag	0 days	Tue 29/05/12
173	Production of core for first pre-production staves	60 days	Wed 30/05/12
174	Work with rest of UK program to complete build of first pre-productic	80 days	Wed 22/08/12
175	Understand QA requirements from prototype staves	624 days	Thu 01/04/10
176	Design and build final QA infrastructure	150 days	Wed 22/08/12
177	2 assembly sites incl. QA ready	0 days	Mon 01/04/13
178	**Module mounting 2.3.5.5**	**935 days**	**Mon 31/08/09**
179	**Detector mounting Development**	**65 days**	**Mon 31/08/09**
184	**Stave Development tests**	**100 days**	**Mon 31/08/09**
188	**RAL Production Facility**	**805 days**	**Mon 01/03/10**
189	Review/Assess future requirements	0 days	Mon 01/03/10
190	Production Tooling - design/manufacture	100 days	Mon 05/04/10
191	Development of prototype tooling	8 mons	Mon 23/08/10
192	Optical Equipment and mounting table - purchase	0 days	Mon 04/04/11
193	production of prototype tooling	8 mons	Mon 04/04/11
194	Tooled up ready for populating first full size stave	0 days	Fri 11/11/11
195	Qualification of prototype tooling to build staves	12 mons	Mon 14/11/11
196	Use tooling in 1st production staves	6 mons	Mon 15/10/12

Pixel Development Gantt Chart (WP5)

ID	Task Name	Duration	Start
1	**IBL**	**458 days?**	**Thu 01/04/10**
2	Planar sensor prototype production	65 days?	Thu 01/04/10
3	3d Sensor design prototype production	65 days?	Thu 01/04/10
4	Module preparation	44 days?	Thu 01/07/10
5	Testbeam and analysis	43 days	Fri 01/10/10
6	**Technology selection**	0 days	Tue 30/11/10
7	Sensor production	173 days?	Mon 03/01/11
8	**Sensor production complete**	0 days	Wed 31/08/11
9	Sensor testing	95 days	Tue 23/08/11
10	**Tested sensors delivered**	0 days	Mon 02/01/12
11	**Forward Pixel sensor R&D Planar and 3D**	**730 days?**	**Mon 14/06/10**
12	**First R&D run**	**655 days?**	**Mon 14/06/10**
13	Mask design for R&D	65 days?	Mon 14/06/10
14	**Design submitted for fabrication**	0 days	Fri 10/09/10
15	**First wafer delivery**	0 days	Mon 17/01/11
16	Delivery and evaluation of first R&D wafers	75 days	Mon 17/01/11
17	**Wafers ready for bump bonding**	0 days	Fri 29/04/11
18	Bump bonding and evaluation with readout	151 days?	Fri 20/05/11
19	**Results from lab studies of first R&D wafers**	0 days	Fri 16/12/11
20	Irradiation and lab and testbeam studies	337 days?	Thu 01/09/11
21	**Lab and testbeam results from irradiated sensors**	0 days	Fri 14/12/12
22	**Second R&D run**	**355 days?**	**Mon 21/11/11**
23	Mask design for second run	65 days	Mon 21/11/11
24	**Design submitted for fabrication**	0 days	Fri 17/02/12
25	**First wafer delivery**	0 days	Fri 25/05/12
26	Sensor testing implementing 'production' procedures	55 days	Mon 28/05/12
27	Bump bonding and evaluation with readout	165 days?	Mon 13/08/12
28	**Results from studies of 2nd R&D wafers**	0 days	Fri 29/03/13
29	**Pixel modules and support structures**	**782 days?**	**Thu 01/04/10**
30	Pixel layout studies	782 days	Thu 01/04/10
31	Preliminary R&D (overlap with short-strip mechanics)	257 days?	Thu 01/04/10
32	**Design for disk prototype ready (thermo-mechanical)**	0 days	Fri 25/03/11
33	Fabrication and evaluation of disk prototype (thermo-mechanical)	255 days	Mon 11/07/11
34	**Results from thermo-mechanical disk prototype**	0 days	Fri 29/06/12
35	**Design of single-chip pixel prototype from layout studies**	0 days	Fri 23/03/12
36	Fabrication of bare disk	195 days	Mon 02/07/12
37	**Results from prototype disk**	0 days	Fri 29/03/13
38	**Connectivity**	**664 days?**	**Mon 13/09/10**
39	Flex circuit design	330 days?	Mon 13/09/10
40	Through Silicon Vias R&D	315 days?	Mon 04/10/10
41	MCMD circuit R&D	315 days?	Mon 04/10/10
42	Design of single-chip pixel prototype	120 days	Mon 09/01/12
43	Fabrication and evaluation of single-chip pixel prototype	185 days	Mon 02/07/12
44	**Results from single-chip pixel prototype**	0 days	Fri 29/03/13
45	**Irradiation**	**785 days?**	**Mon 12/04/10**
46	Irradiation studies of sensors and hybrids	698 days?	Mon 12/04/10
47	Irradiation with new designed strip sensors	393 days?	Thu 12/05/11
48	Irradiation for pixel R&D	421 days?	Fri 02/09/11

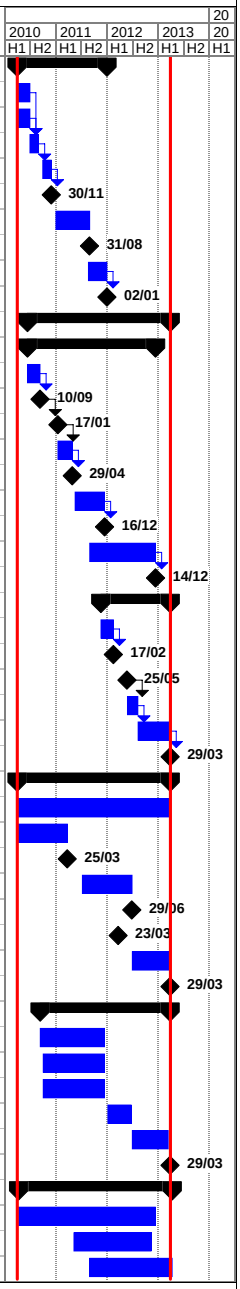

Project: Pixel WP.mpp
Date: Wed 30/09/09

Task
Progress
Milestone
Summary
Rolled Up Task
Rolled Up Milestone
Rolled Up Progress
Split
External Tasks
Project Summary
Group By Summary
Deadline

Trigger Activities Gantt chart (WP6, WP7 & WP8)

ID	Task Name	Duration	Start
1	**WP6 - L1Calo design studies**	**766.13 days**	**Fri 23/04/10**
2	Initial design studies	9 mons	Fri 23/04/10
3	Detailed design of topological algorithms	9 mons	Fri 31/12/10
4	Phase II design studies	19 mons	Mon 17/10/11
5	**WP6 - L1Calo Hardware/firmware**	**830 days**	**Thu 01/04/10**
6	Technology evaluations	9.5 mons	Thu 01/04/10
7	Firmware development in preparation for Topological Processor	9.5 mons	Thu 01/04/10
8	Conceptual hardware/firmware design of Topological processor	3 mons	Thu 23/12/10
9	Conceptual hardware/firmware design of Topological processor complete	0 days	Thu 17/03/11
10	Construction of prototype	16 mons	Thu 17/03/11
11	Construction of prototype complete and tested	0 days	Thu 07/06/12
12	Construction of system/spares	11 mons	Thu 07/06/12
13	Delivery of system to CERN	0 days	Thu 06/06/13
14	Phase II technology evaluations	29 mons	Mon 03/01/11
15	**WP7 - L1Track Simulation Studies**	**782 days**	**Thu 01/04/10**
16	Development of the Descrete Event Simulation framework	218 days	Thu 01/04/10
17	DES framework operational	0 days	Mon 31/01/11
18	Use of DES in system design studies	608 days	Wed 01/12/10
19	L1Track-specific software tools for high pile-up studies	326 days	Thu 01/04/10
20	Physics and trigger rate studies with/without L1Track	717 days	Thu 01/07/10
21	Rates for lepton triggers with and without L1Track	0 days	Thu 31/03/11
22	Pattern recognition studies	673 days	Wed 01/09/10
23	First implementation of pattern recognition adapted for hardware	0 days	Wed 30/11/11
24	Use of L1Track info at L1 trigger and propagation to HLT	260 days	Mon 02/04/12
25	**WP7 - Off-detector hardware and data paths**	**782 days**	**Thu 01/04/10**
26	RoIMapper design work	522 days	Thu 01/04/10
27	RoIMapper design complete	0 days	Fri 30/03/12
28	Construction of RoIMapper demonstrator testbench	260 days	Mon 02/04/12
29	Track Trigger Processor - evaluation of technologies	348 days	Wed 30/11/11
30	**WP7 - On-detector electronics and readout**	782 days	Thu 01/04/10
31			
32	**WP8.1 Optimise HLT ID Tracking sw for Phase-I**	**782 days**	**Thu 01/04/10**
33	Update L2 tracking for IBL & FTK	240.13 days	Thu 01/04/10
34	Tracking Code adapted for IBL & FTK	0 days	Thu 03/03/11
35	Optimise L2 Pat. Rec. for Phase-I lumi	180.13 days	Thu 03/03/11
36	Optimise L2 Pat. Rec. for different RoI types	180 days	Thu 02/02/12
37	Optimize L2 code to exploit GPU	240.13 days	Thu 10/11/11
38	Develop Data Prep Code for SCT RODs	240.13 days	Mon 02/08/10
39	Adapt HLT Tracking for RoD data prep.	60 days	Mon 04/07/11
40	Evaluate Perf. of Data Prep at SCT ROD	60.13 days	Mon 19/12/11
41	Optimise config. of EF tools	180 days	Mon 06/02/12
42	Performance mesurements & optimisation	60 days	Mon 07/01/13
43	HLT Tracking code optimized for Phase-I	0 days	Fri 29/03/13
44	**WP8.2 Optimise HLT Selection sw for Phase-I**	**780.25 days**	**Fri 01/07/11**
45	Define Trigger Menus for MC productions	120.13 days	Fri 01/07/11
46	Adapt Selections for upgraded LVL1	160.13 days	Fri 09/03/12
47	Optimise Selections & Menu	140 days	Fri 11/01/13
48	Performance measurements & optimisation	180 days	Fri 18/10/13
49	Trigger Selections optimised for Phase-I	0 days	Fri 27/06/14
50	**WP8.3 core HLT sw updates**	**861 days**	**Fri 30/04/10**
51	Update Core sw for Upgraded LVL1 + FTK	240 days	Fri 30/04/10
52	Update Core sw for Full ev. proc.	261 days	Fri 01/04/11
53	Steering code upgraded for fulle event reconstruction	0 days	Fri 30/03/12
54	Update Core sw to run L2 & EF on same proc.	120 days	Mon 02/04/12
55	Optimise core HLT sw	240 days	Mon 17/09/12
56	**WP8.4 Reconfigure and tune farm hw for Phase-I**	**640.25 days**	**Fri 01/07/11**
57	Evaluation of replacement h/w	400.13 days	Fri 01/07/11
58	Optomise and tune hardware	240.13 days	Fri 11/01/13
59	Evolution of dataflow sw for Phase-I	400.13 days	Fri 01/06/12
60	**WP8.5 Define Phase-II Trigger Selection Strategy**	**600 days**	**Wed 01/06/11**
61	Define Trigger Menus for Phase-II sim.	120 days	Wed 01/06/11
62	Optimise Selections	180 days	Wed 08/02/12
63	Performance Measurements	240 days	Wed 17/10/12
64	**WP8.6 Define Phase-II Architecture**	**240 days**	**Wed 09/01/13**
65	Dataflow Modelling	240 days	Wed 09/01/13
66	Results available from simulation & dataflow modeling	0 days	Wed 11/12/13

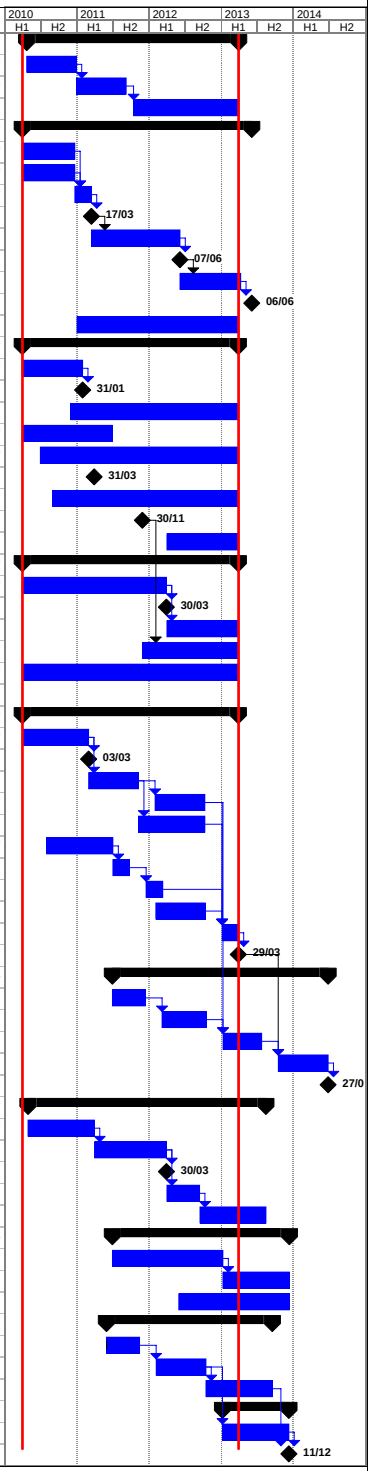

Project: Untitled
Date: Wed 30/09/09

Task
Split
Progress
Milestone
Summary
Project Summary
External Tasks
External Milestone
Deadline

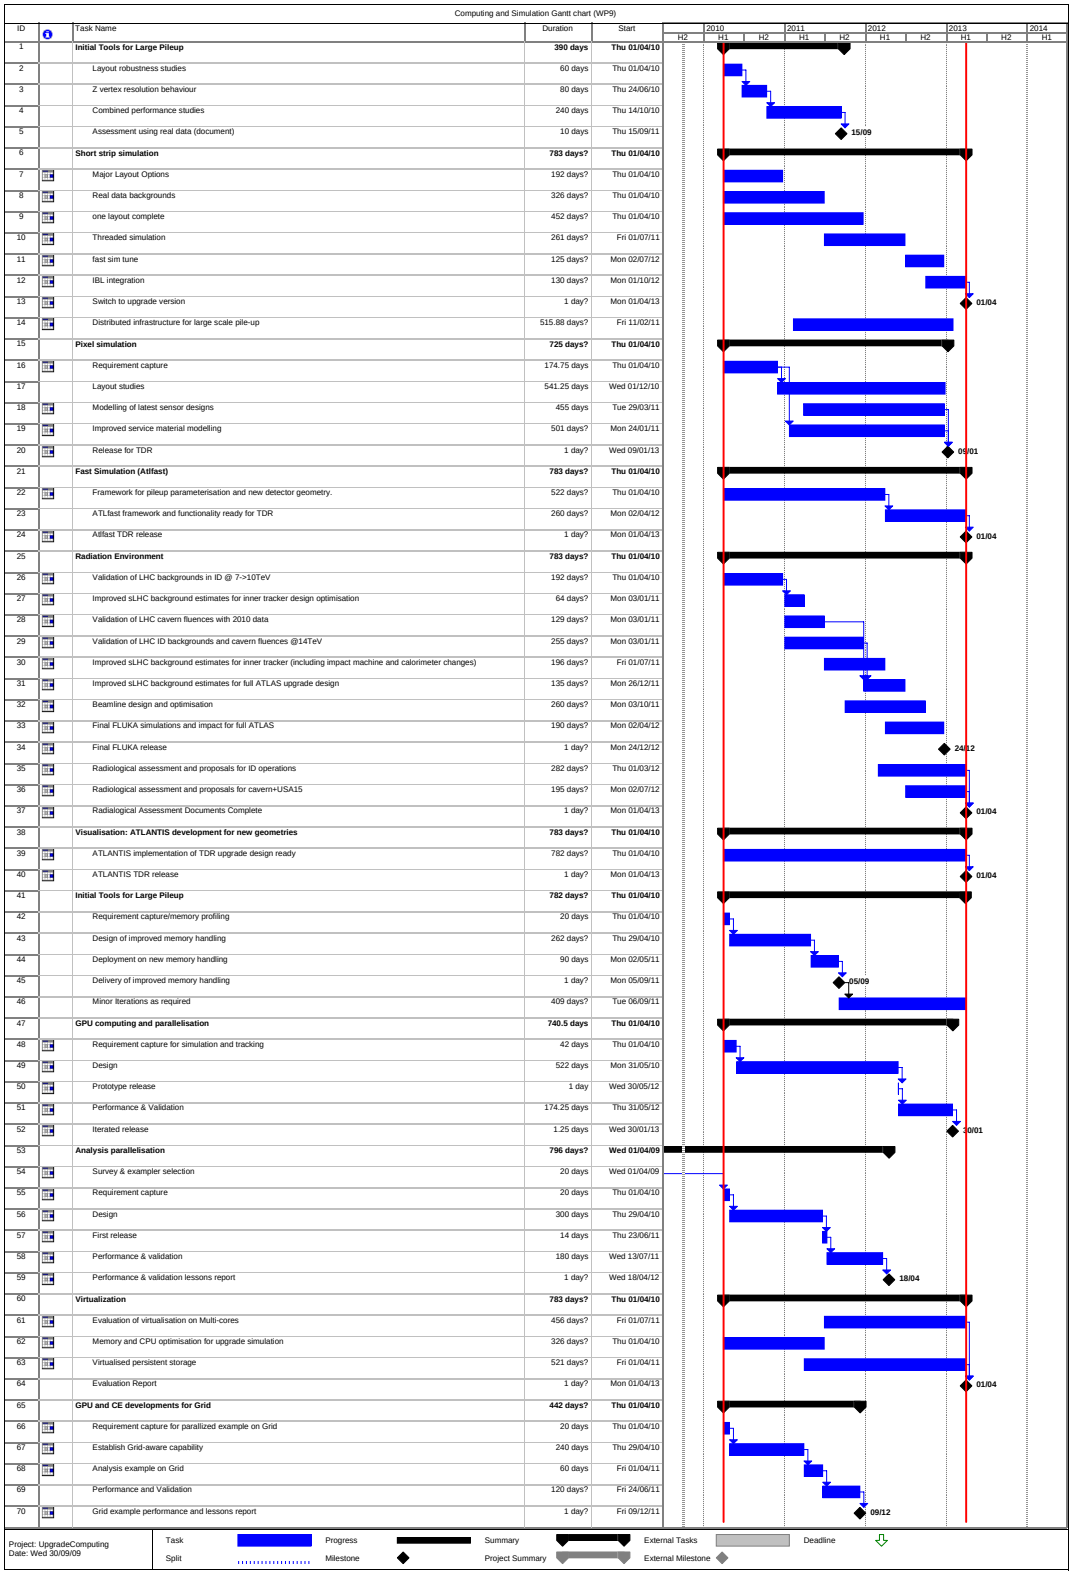
Computing and Simulation Gantt chart (WP9)
ID
Task Name
Duration
Start
2010
2011
2012
2013
2014
H2
H1
1 Initial Tools for Large Pileup 390 days Thu 01/04/10
2 Layout robustness studies 60 days Thu 01/04/10
3 Z vertex resolution behaviour 80 days Thu 24/06/10
4 Combined performance studies 240 days Thu 14/10/10
5 Assessment using real data (document) 10 days Thu 15/09/11
15/09
6 Short strip simulation 783 days? Thu 01/04/10
7 Major Layout Options 192 days? Thu 01/04/10
8 Real data backgrounds 326 days? Thu 01/04/10
9 one layout complete 452 days? Thu 01/04/10
10 Threaded simulation 261 days? Fri 01/07/11
11 fast sim tune 125 days? Mon 02/07/12
12 IBL integration 130 days? Mon 01/10/12
13 Switch to upgrade version 1 day? Mon 01/04/13
01/04
14 Distributed infrastructure for large scale pile-up 515.88 days? Fri 11/02/11
15 Pixel simulation 725 days? Thu 01/04/10
16 Requirement capture 174.75 days Thu 01/04/10
17 Layout studies 541.25 days Wed 01/12/10
18 Modelling of latest sensor designs 455 days Tue 29/03/11
19 Improved service material modelling 501 days? Mon 24/01/11
20 Release for TDR 1 day? Wed 09/01/13
09/01
21 Fast Simulation (Atlfast) 783 days? Thu 01/04/10
22 Framework for pileup parameterisation and new detector geometry. 522 days? Thu 01/04/10
23 ATLfast framework and functionality ready for TDR 260 days? Mon 02/04/12
24 Atlfast TDR release 1 day? Mon 01/04/13
01/04
25 Radiation Environment 783 days? Thu 01/04/10
26 Validation of LHC backgrounds in ID @ 7->10TeV 192 days? Thu 01/04/10
27 Improved sLHC background estimates for inner tracker design optimisation 64 days? Mon 03/01/11
28 Validation of LHC cavern fluences with 2010 data 129 days? Mon 03/01/11
29 Validation of LHC ID backgrounds and cavern fluences @14TeV 255 days? Mon 03/01/11
30 Improved sLHC background estimates for inner tracker (including impact machine and calorimeter changes) 196 days? Fri 01/07/11
31 Improved sLHC background estimates for full ATLAS upgrade design 135 days? Mon 26/12/11
32 Beamline design and optimisation 260 days? Mon 03/10/11
33 Final FLUKA simulations and impact for full ATLAS 190 days? Mon 02/04/12
34 Final FLUKA release 1 day? Mon 24/12/12
24/12
35 Radiological assessment and proposals for ID operations 282 days? Thu 01/03/12
36 Radiological assessment and proposals for cavern+USA15 195 days? Mon 02/07/12
37 Radialogical Assessment Documents Complete 1 day? Mon 01/04/13
01/04
38 Visualisation: ATLANTIS development for new geometries 783 days? Thu 01/04/10
39 ATLANTIS implementation of TDR upgrade design ready 782 days? Thu 01/04/10
40 ATLANTIS TDR release 1 day? Mon 01/04/13
01/04
41 Initial Tools for Large Pileup 782 days? Thu 01/04/10
42 Requirement capture/memory profiling 20 days Thu 01/04/10
43 Design of improved memory handling 262 days? Thu 29/04/10
44 Deployment on new memory handling 90 days Mon 02/05/11
45 Delivery of improved memory handling 1 day? Mon 05/09/11
05/09
46 Minor iterations as required 409 days? Tue 06/09/11
47 GPU computing and parallelisation 740.5 days Thu 01/04/10
48 Requirement capture for simulation and tracking 42 days Thu 01/04/10
49 Design 522 days Mon 31/05/10
50 Prototype release 1 day Wed 30/05/12
51 Performance & Validation 174.25 days Thu 31/05/12
52 Iterated release 1.25 days Wed 30/01/13
30/01
53 Analysis parallelisation 796 days? Wed 01/04/09
54 Survey & exampler selection 20 days Wed 01/04/09
55 Requirement capture 20 days Thu 01/04/10
56 Design 300 days Thu 29/04/10
57 First release 14 days Thu 23/06/11
58 Performance & validation 180 days Wed 13/07/11
59 Performance & validation lessons report 1 day? Wed 18/04/12
18/04
60 Virtualization 783 days? Thu 01/04/10
61 Evaluation of virtualisation on Multi-cores 456 days? Fri 01/07/11
62 Memory and CPU optimisation for upgrade simulation 326 days? Thu 01/04/10
63 Virtualised persistent storage 521 days? Fri 01/04/11
64 Evaluation Report 1 day? Mon 01/04/13
01/04
65 GPU and CE developments for Grid 442 days? Thu 01/04/10
66 Requirement capture for parallized example on Grid 20 days Thu 01/04/10
67 Establish Grid-aware capability 240 days Thu 29/04/10
68 Analysis example on Grid 60 days Fri 01/04/11
69 Performance and Validation 120 days? Fri 24/06/11
70 Grid example performance and lessons report 1 day? Fri 09/12/11
09/12
Project: UpgradeComputing
Date: Wed 30/09/09
Task
Split
Progress
Milestone
Summary
Project Summary
External Tasks
External Milestone
Deadline

Ref	Risk description	Potential impact on project	Inherent risk score			Existing controls	Mitigating factors	Residual risk score			Comment	Proposed action
			L'hood	Impact	Total			L'hood	Impact	Total		
WP1	AFP approval at CERN	Project delayed	2	5	10	In review since early 2009	Continued work on physics and detector R&D	1	5	5		R&D plan in the proposal
1.1.1	Serious compatability issue between LHC and AFP	Prevents data-taking	2	5	10	Technical Integration Meetings between AFP and LHC for last four years	Machine physicist on AFP staff	1	5	5	See proposal	Post request (Cockrofit/Manchester)
1.2.1	FE-I4 chip is delayed	Delay to R&D and production by 6-9 months	2	4	8	FE-I4 is critical to ATLAS IBL project so has significant external resource	Schedule would be modified in short term. If long term issue, then AFP would use FE-I3	1	4	4	See proposal	
1.2.2	FE-I4 sensors have yield problem	Delay to R&D and production	2	4	8	R&D plan with devices being manufactured in three different companies	Use FE-I3 sensors or FE-I4 planar sensors in region far from the beam.	1	4	4		
1.2.3	Detector is unreliable	Access in LHC tunnel is very limited. Data would be lost.	2	5	10	R&D and design is geared towards a system that is reliable and fast to replace.	Burn-in of detectors	1	5	5	See proposal	Part of Post Request (Glasgow)
1.2.4	Radiation tolerance of local electronics	Lifetime of detector	2	4	8	Detailed work from FP420 design report	Radiation review in 2010. R&D will choose from various options.	1	4	4	See proposal	
1.2.5	Local cooling compatible with LHC tunnel	Delay to installation.	2	4	8	Conceptual design exists.	Part of the R&D to complete design and make a prototype	1	4	4	See proposal	
1.3.1	DAQ/SC software is delayed	Delay operation of pre-production system	1	4	4	DAQ is based on ATLAS standards	In the planning and has clear delivery date	1	3	3	See proposal	
1.4.1	Momentum precision poor due to alignment	Physics potential	2	4	8	Significant work performed in recent years	Continuing effort by dedicated team	1	4	4	See proposal	Part of post request (Glasgow)
1.6.1	Delays in assembly and commissioning	Delay operation of detector	2	5	10	Built a good team and have an outline plan.	System engineering support from RAL and consultant	1	5	5	See proposal	

Ref	Risk description	Potential impact on project	Inherent risk score			Existing controls	Mitigating factors	Residual risk score			Comment	Proposed action
			L'hood	Impact	Total			L'hood	Impact	Total		
Strip, Pixel Trackers	UK funds insufficient	Completion of project delayed	2	3	6	Review by UK project management	Contingency and International Collaborators	1	3	3	Significant cost uncertainties	Call on contingency, look to international collaborators
Strip, Pixel Trackers	Project definition changes dramatically	TDR delayed. UK leadership and impact diminished	3	3	9	Review by international Project Office	Major UK roles in decision	1	3	3	Many UK roles	Adapt scope of programme and, if needed, call on contingency
Strip, Pixel Trackers	Key UK staff leave project	UK leadership and impact diminished	3	3	9	Review by UK project management	Pool of expertise	2	2	4	Delays and potential risk to deliverables	Examine responsibility and scope in light of funding
Strip, Pixel Trackers	Shortage of appropriate staff	Risk to quality and schedule	2	3	6	Monitoring by UK project management and regular reviews	Pool of expertise	1	3	3	Technical staff do not FEC and are less expensive	Hire additional technical staff
Strip, Pixel Trackers	Experience from ATLAS changes project specification	Higher hit rate than simulation	2	3	6	Review by international Project Office	Safety factors in design	1	2	2	Higher radius = larger	Revisit layout and/or technologies
Strip, Pixel Trackers	Experience from ATLAS changes project specification	Higher radiation levels than simulation	2	3	6	Review by international Project Office	Safety factors in design	1	3	3	area = cost	Revisit layout and/or technologies
Strip, Pixel Trackers	Experience from ATLAS changes project specification	Different physics priorities	2	2	4	Review by international Project Office	High luminosity bound to be desirable	1	1	1	Technologies developed likely to be needed in any scenario	Revisit layout and/or technologies
Strip, Pixel Trackers	Currency fluctuations result in a dramatic reduction in our buying power	Major purchases will consume far more of a WP budget than expected	3	4	12	Reduce scope of UK orders to match new price regime	Contingency possible through greater sharing of projects with collaborators	3	2	6	Sensors and CFRP are most likely to be a problem	Reduce scope if appropriate, call upon contingency
WP2	ASIC development falls behind project office schedule	Overall schedule delayed	3	3	9	Review by ATLAS Upgrade PO, SG	Many developments still able to proceed	3	2	6	Costs of extra time	Adapt programme and schedule to accommodate any revised ASIC dates
WP2	STAVE09 geometry (side mounted SMC) not final stave adopted	Schedule impacted. Rework of tooling needed.	1	2	2	Review by ATLAS Upgrade Layout, SG and PO	Alternatives already studied	1	2	2	Potential costs of extra time	Adapt hardware programmes to accepted final geometry
WP3	Serial Powering not proven robust in high statistics studies.	Requirement to switch to DC to DC	2	3	6	Review by ATLAS Upgrade PO, SG	DC to DC already being studied	2	2	4	rework of tapes, hybrids	Adapt to DC to DC as required
WP4	Welding pipes near ASICs rejected	Redesign of EoS region needed	2	2	4	Review by UK project management, PO, SG	none	2	2	4	none	Redesign EoS region
WP4	Mass production techniques prove problematic	Overall delay to programme	1	2	2	Review by UK project management	Integration studies raise interface problems with other subsystems	1	2	2	none	Redesign of stave to improve mass manufacturability
WP4	Integration studies raise interface problems with other subsystems	Programme delays	1	3	3	Review by UK project management, PO, SG	none	1	3	3	Important to be active to keep this risk minimised	Work with other systems to avoid such issues.
WP5	Sensors for highest dose regions prohibitively expensive for full detector.	Final detector cost untenable	1	3	3	Review by UK project management, PO, SG	Geometry Dependent	1	2	2	Compressed R&D timescales	Redesign Layout to move sensors to lower does regions
WP5	TSV, or other methods prove infeasible	Detector Layout requires revision	1	3	3	Review by UK project management	Other layouts possible	1	2	2	more massive detector	Revert to undesirable geometry.
L1Calo	Simulation fails to identify viable topological triggers when detailed pileup is taken into account	Topological trigger technique not viable. ATLAS performance compromised at Phase I. Need to seek alternatives.	1-2	5	5-10	Detailed simulation studies underway	Early studies show likely benefit from topology	1	5	1-5	Outcome depends on results from simulation with pileup. This is a developing study within ATLAS.	Continue simulation studies
L1Calo	Trigger processor too complex or processing outside latency	Topological technique not viable. Need to seek alternative technique	1-2	5	5-10	Engineering studies will start when funding is available	Firmware testbenches will establish viability and latency early in project	1	5	1-5	Initial studies already show some topological options with existing hardware	Initiate firmware studies when funding is available
L1Track	sLHC high pile-up simulation too slow for L1 trigger rate studies	Not possible to evaluate benefits of L1Track at sLHC.	2	5	10	Major effort to improve the performance of ATLAS full simulation	Investigate the use of fast simulation tools for Level-1 trigger rate studies	2	3	6		Situation will improve with additional effort
L1Track	L1Track trigger leads to significant increase in the tracker material	ATLAS does not approve L1Track, due to impact on offline tracking and calorimeter performance.	1-2	5	5-10	Simulations and evaluation of alternative designs	Investigate the use of fast simulation tools for Level-1 trigger rate studies	1	5	5	Initial estimates of the RoI-based L1Track approach encouraging	Careful evaluation of alternative designs
HLT	Phase-I simulated pile-up data samples not available in time.	Not possible to optimize HLT tracking and selection algorithms with appropriate data samples	2	4	8	Review by ATLAS and UK project management	Alternative solutions based on overlay of real events and extrapolation from lower luminosity data.	1	2	2	Potential cost of extra effort	Pursue multiple options for generating datasets
HLT	HLT tracking software too slow or fails to achieve required performance	Potential delay will alternative techniques are developed	2	4	8	Review by UK project management	Potential speed-up from GPU, RoD-level pre-processing and pre-HLT tracking info.	1	3	3	Potential cost of extra effort	Pursue hardware and software options for speed-up
Computing	Lost of key skilled personnel	Difficulty in delivering more technical aspects of package	3	4	12	Spread of institutes and skills	Training and cross-coverage of tasks	3	2	6		Avoid reduction in cross-coverage
Computing tasks 7-10	Dependence on operating system features	Refactoring required to retain performance gains	3	4	12	Use of widely-accepted OS	Use of open standard languages	1	3	3		Enforce open software etc
Computing	Prohibitive cost of new technologies	Inability to provide capacity or develop on new technology	3	5	15	More than one technology path under investigation	Strong links with providers	1	3	3		Avoid premature technology choice
Computing - radiation environment	Delay in the LHC running and energy ramp	Knock on delays in deliverables	3	5	15	Multi-strand activity	Focus on later simulation and radiological milestones and return to benchmarking when data available	3	3	9		Constant review of schedule

Chapter 6

Appendices

Risk Register

Glossary

ABC	Chip used to read-out existing SCT strip detector
ABCN	one of either the ABCN25 or ABCN13 chips next generation
ABCN25	ABCN chip in .25mm CMOS
ABCN13	ABCN chip in 13 mm CMOS
AFP	ATLAS Forward Physics
ASIC	Application specific integrated circuit
ATCA	Advanced Telecommunications Architecture (a crate and bus standard)
ATLAS	One of two general purpose LHC experiments
BCC	Basic Controller Chip
BER	Bit Error Rate
BSM	Beyond the Standard Model
CEP	Central Exclusive Production
CMOS	Complementary Metal Oxide Silicon
CMS	Compact Muon Spectrometer
CMM	Common Merger Module (part of L1Calo)
CPM	Cluster Processing Module (part of L1Calo)
CV	Capacitance Voltage
CTE	Coefficient of Thermal Expansion
DDTC	Double sided Double type columns
Ethernet	Standard communication bus
EoS	End of Stave reference to the regions close to service gap
FPCC	Field Programmable Controller Chip
FEA	Finite Element Analysis
FPGA	Field Programmable Gate Array
GBT	Giga Bit Transceiver
GPU	Graphics Processing Unit
GPGPU	General Purpose Graphics Processing Unit

HLT	High-Level Trigger
HSIO	Hi-Speed Input/Output
IV	Current Voltage
JEM	Jet-Energy module (part of L1Calo)
L1Calo	Level-1 Calorimeter trigger
LLVM	Low level virtual machine
LVDS	Low-Voltage Differential Signals
MCC	Module Control Chip
MSSM	Minimal Supersymmetric Standard Model
MuSTARD	SCT test read-out board
NUMA	Non Uniform Memory Access
OpenCL	Open Computing Language
PON	Passive Optical Networks
PLC	Planar Lightwave Chip
PXI	PCI Extensions for instrumentation - modular instrumentation bus
QA	Quality Assurance
ROD	Readout Driver (connects detector to ATLAS readout system
SATA	Serial ATA (Advanced Technology Attachment)
SCTDAQ	Silicon Tracker Data Aquisition
sLHC	Super Large Hadron Collider (upgraded LHC)
SLOG	SCT test read-out board
SM	Standard Model
SPi	Serial Power monitor chip
STAVE09	International prototype short strip tracker stave
TOTEM	LHC experiment measuring Total Cross Sections, Elastic Scattering and Diffraction
TTC	Trigger and Timing Control
USB	User Serial Bus
VL	Versatile Link (optical link)
VME	intrnational standard crate backplane control bus
YM	Young's Modulus (of elasticity)

Bibliography

[1] M. G. Albrow and A. Rostovtsev, hep-ph/0009336.

[2] V. A. Khoze, A. D. Martin and M. G. Ryskin, Eur. Phys. J. **C23**, 311 (2002), [hep-ph/0111078].

[3] B. E. Cox, hep-ph/0501064, AIP Conf Proc. 753 (2005) 103.

[4] V. A. Khoze, A. B. Kaidalov, A. D. Martin, M. G. Ryskin and W. J. Stirling, hep-ph/0507040.

[5] J. R. Forshaw, hep-ph/0508274.

[6] K. Piotrzkowski, Phys. Rev. **D63**, 071502 (2001), [hep-ex/0009065].

[7] M. G. Albrow *et al.*, CERN-LHCC-2005-025.

[8] FP420 R and D, M. G. Albrow *et al.*, 0806.0302, Accepted for publication in Journal of Instrumentation. This is referred to as the FP420 Design Report in the text.

[9] V. A. Khoze, A. D. Martin, M. G. Ryskin and W. J. Stirling, Eur. Phys. J. **C38**, 475 (2005), [hep-ph/0409037].

[10] K. Piotrzkowski, Nucl. Phys. Proc. Suppl. **179-180**, 2 (2008), [0807.1624].

[11] CDF, T. Aaltonen *et al.*, Phys. Rev. **D77**, 052004 (2008), [0712.0604].

[12] CDF, T. Aaltonen *et al.*, Phys. Rev. Lett. **102**, 242001 (2009), [0902.1271].

[13] T. Plehn, D. L. Rainwater and D. Zeppenfeld, Phys. Rev. Lett. **88**, 051801 (2002), [hep-ph/0105325].

J. R. Ellis, J. S. Lee and A. Pilaftsis, Phys. Rev. **D71**, 075007 (2005), [hep-ph/0502251].

B. E. Cox, F. K. Loebinger and A. D. Pilkington, JHEP **10**, 090 (2007), [0709.3035].

J. R. Forshaw, J. F. Gunion, L. Hodgkinson, A. Papaefstathiou and A. D. Pilkington, JHEP **04**, 090 (2008), [0712.3510].

M. Chaichian, P. Hoyer, K. Huitu, V. A. Khoze and A. D. Pilkington, 0901.3746.

B. E. Cox *et al.*, Eur. Phys. J. **C45**, 401 (2006), [hep-ph/0505240].

[14] G. J. A. Brown *et al.*, ATL-DAQ-PUB-2009-006.

[15] G. Azuelos *et al.*, J. Phys. **G28**, 2453 (2002), [hep-ex/0203019].

[16] I. Hinchliffe and F. E. Paige, Prepared for APS / DPF / DPB Summer Study on the Future of Particle Physics (Snowmass 2001), Snowmass, Colorado, 30 Jun - 21 Jul 2001.

[17] A. J. Barr, JHEP **02**, 042 (2006), [hep-ph/0511115].

[18] M. Mangano, Prepared for CARE-HHH-ADP Workshop on Scenarios for the LHC Luminosity Upgrade, Arcidosso, Italy, 31 Aug - 3 Sep 2005.

[19] F. Gianotti *et al.*, Eur. Phys. J. **C39**, 293 (2005), [hep-ph/0204087].

[20] ATLAS-UK, Reports to STFC Oversight Committee.

[21] M. Hanlon, The Development of p-type Silicon Detectors for the High Ra- diation Region of the LHC. PhD thesis, University of Liverpool (1999).

[22] G. Casse *et al.*, Nucl. Instrum. Meth. **A487/3**, 465 (2005).

P. Allport *et al.*, IEEE Trans. Nucl. Sci. **48**, 1007 (2001).

[23] A. Affolder *et al.*, Nucl. Instrum. Meth. **A604**, 250 (2009).

[24] A. Affolder *et al.*, IEEE Trans. Nucl. Sci. **56**, 765 (2009).

G. Casse *et al.*, Nucl. Instrum. Meth. **A598**, 671 (2009).

[25] USCMS, L. Amaral *et al.*, FERMILAB-PUB-09-335-CD.

[26] B. Arvidsson *et al.*, Journal of Instrumentation **4**, P07010 (2009).

[27] G. Casse *et al.*, IEEE Trans. Nucl. Sci. **55**, 1695 (2008).

G. Casse *et al.*, Can we claim multiplication effects in irradiated silicon? 4th RD50 - Workshop on Radiation hard semiconductor devices for very high luminosity colliders, Freiburg, Germany, 3-5 June 2009.

G. Casse *et al.*, Enhanced efficiency of segmented silicon detectors of various thicknesses after hadron irradiations up to 21016 neq cm2, 11th European Symposium on Semiconductor Detectors, June 7 - 11, 2009 Wildbad Kreuthmy (GE).

[28] G. Pellegrini *et al.*, Nucl. Instrum. Meth. **A592**, 38 (2008).

D. Pennicard *et al.*, Nucl. Instrum. Meth. **A592**, 16 (2008).

D. Pennicard *et al.*, IEEE Trans. Nucl. Sci. **54**, 1435 (2007).

[29] C. Da Via *et al.*, IEEE Trans. Nucl. Sci. **56**, 505 (2009).

M. Mathes *et al.*, IEEE Trans. Nucl. Sci. **55**, 3731 (2008).

D. Pennicard *et al.*, Nucl. Instrum. Meth. **A598**, 67 (2009).

D. Pennicard *et al.*, Nucl. Instrum. Meth. **A604**, 412 (2009).

C. Fleta *et al.*, Nucl. Instrum. Meth. **A607**, 89 (2009).

C. Da Via *et al.*, Nucl. Instrum. Meth. **A587**, 243 (2008).

C. Da Via *et al.*, Nucl. Instrum. Meth. **A604**, 505 (2009).

C. Da Via *et al.*, Nucl. Instrum. Meth. **A603**, 318 (2009).

[30] J. Jones *et al.*, (2005), Proc. LECC 2005 Workshop. CERN Report CERN-2005-011.

[31] http://simpy.sourceforge.net/.

[32] I. Dawson *et al.*, Nucl. Instrum. Methods Phys. Res. **A510**, 219 (2003).

[33] I. Dawson *et al.*, Nucl. Instrum. Methods Phys. Res. **A515**, 422 (2003).

[34] B. et al., ATLAS Radiation Background Task Force Summary Document, ATL-GEN-2005-001, January 2005.

[35] The FLUKA particle physics MonteCarlo simulation package: http://www.fluka.org.

[36] EU SLHC Preparatory Phase project: http://slhcpp.web.cern.ch/SLHCPP/index.html.

[37] L. Nicolas, Presentation to ATLAS Cavern Background meeting: http://indico.cern.ch/conferenceDisplay.py?confId=66817.

[38] Dawson member of High Luminosity Upgrade Steering Group; Member of EU-SLHC-PP governing board; LoI chapter editor.

[39] See for example EDMS ATLAS Upgrade notes: https://edms.cern.ch/nav/ATLAS/ATU-0000000031.

[40] Nicolas and Dawson, Neutron Moderator Studies for the ATLAS Inner Tracker Upgrade, EDMS Id 982957.

[41] Dawson and Shupe, Radiation background studies for the LHC upgrade: Impact of machine magnets inside the ATLAS experiment. EDMS Id 932316.

[42] I. Dawson, Presentation to IBL general meeting by I.Dawson: http://indico.cern.ch/conferenceDisplay.py?confId=54685.

[43] J. Huang, (2009), Symposium on High Performance Chips, http://www.hotchips.org/hc21.