


University of Bristol
Brunel University
Imperial College London
Rutherford Appleton Laboratory



R&D in preparation for an upgrade of CMS for the Super-LHC

1	Executive Summary	2
2.	Project history and recent developments	2
2.1	LHC upgrade schedule and planning	2
2.2	CMS planning	3
2.3	UK adaptation to CMS planning	4
3.	Work Package 1: Tracker and Level-1 Trigger Simulations and Software	5
3.1	Objectives	5
3.2	Progress to date	5
3.3	Deliverables	6
3.4	Staff on project	7
3.5	Expenditure	7
3.6	Comparisons with CMS activities elsewhere	7
4.	Work Package 2: Outer Tracker Readout	7
4.1	Objectives	7
4.2	Progress to date	7
4.3	Deliverables	9
4.4	Staff on project	10
4.5	Expenditure	10
4.6	Comparisons with CMS activities elsewhere	10
5.	Work Package 3: Design of Level-1 Calorimeter Triggers at the SLHC	10
5.1	Objectives	10
5.2	Progress to date	11
5.3	Deliverables	12
5.4	Staff on project	12
5.5	Expenditure	12
5.6	Comparisons with CMS activities elsewhere	12
6	Risk register	13
6.1	Risk Management	13
7	Finances	14

1 Executive Summary

The proposal was approved in Spring 2008, and funding announced in January 2009. The project start date has been taken to be April 2009 for reporting purposes. Much of the work depends on the recruitment of new staff which took place in the months following. However some preparatory work was possible using Rolling Grant effort which was foreseen to be allocated to upgrade activities.

There has been significant progress in the first year of the project and several milestones have been met or brought forward, which has reduced some of the potential risk in challenging technical areas.

Full simulation code for the Phase-1 upgrade of the CMS pixel detector has been completed ahead of schedule, and is being used in the preparation of the upgrade Technical Proposal. Physics studies of the upgraded CMS detector are under way. Online software has been delivered in support of the Mini-T card, and a new architecture for CMS online software developed with UK leadership.

A 128 channel CBC (CMS Binary Chip) has been submitted for fabrication in July 2010 after a single design iteration. This was originally foreseen to take place by the end of the second year of the project.

The Mini-T card was successfully manufactured and demonstrated to provide the design functionality for trigger requirements, with early firmware development objectives met. These deliverables were also achieved some months earlier than foreseen. One important outcome was the demonstration that a custom backplane is not required, which was a high risk item and depended on a potentially costly US source.

2. Project history and recent developments

The original proposal was submitted to STFC in October 2007. A review in January 2008 proposed modest changes to reduce costs and recommended approval. This followed, with financial allocations in January 2009. Project funded activities began in the following months, with staff recruitment. In December 2009, we were informed of cuts to already approved projects but it was later confirmed that the CMS upgrade project should be reprofiled within the approved funding envelope.

Meanwhile both the LHC planning and overall CMS objectives have been adapting to the realities of machine operation and performance as well as future physics needs, as summarised below.

2.1 LHC upgrade schedule and planning

At the time of the original proposal, in October 2007, CERN plans for the LHC machine upgrade were not specified in detail, except that improvements needed to the accelerator complex were identified and staged implementation envisaged. The goal was to achieve a luminosity of $10^{35} \text{ cm}^{-2}\text{s}^{-1}$, an order of magnitude higher than the nominal LHC maximum.

In January 2008, the CERN DG foresaw that the machine upgrade would proceed in two stages: Phase I, which would increase the luminosity to $2\text{-}3 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, and Phase II, which would deliver the full upgrade. Approximate dates for both phases were indicated. Since that time there have been significant developments affecting the upgrade schedule.

In September 2008, the LHC was operated for the first time, followed by the well known machine incident instigated by a quenching superconducting splice connector which led to significant damage. Since then CERN began to reconsider both how it would operate the LHC and implement luminosity upgrades. The new plan announced in June of 2010 explains how the LHC will operate, what shutdowns are needed to accomplish key steps to reach design energy and luminosity and then to raise the luminosity, and what work must be done on the machine in each shutdown. Long data runs are followed by long machine shutdowns in which major changes are made. This takes into account the difficulty, time, risk and expense of warming up a large cryogenic system in order to be able to work on it. The new plan proceeds in the following steps

- The LHC will run at 3.5 TeV per beam until the end of 2011 with the goal of providing an integrated luminosity of 1 fb^{-1} to the CMS and ATLAS experiments.

- The LHC will then shut down for 1-1.5 years to make revisions necessary to run at 14 TeV.
- When operation resumes in 2013, there will be a long run of two or three years at a peak luminosity ultimately reaching the nominal value of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$.
- In 2015 or 2016, there will be another long shutdown to install the collimation required to operate at and above the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$.
- There would then be another long run with the luminosity achieving the LHC design value of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and then rising gradually to $\sim 2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The goal for this period, summarized by S Myers at ICHEP 2010, is by 2019 to integrate $\sim 340\text{fb}^{-1}$ of luminosity.

The whole period between now and the end of the long run described above is now referred to as “Phase 1 of LHC Operation”. Following it, there will be a long shutdown for installation of components to enable the LHC to produce up to 300 fb^{-1} per year of integrated luminosity. There would follow up to ten years of operation at this new higher luminosity. This is referred to as “Phase 2 of LHC Operation” and the goal would be for the experiments to integrate up to 3000 fb^{-1} of luminosity.

This is the baseline plan to which CMS and the UK Upgrade R&D activities should therefore adapt.

2.2 CMS planning

CMS is presently preparing a Technical Proposal for the Phase I upgrade and a long term vision for the Phase II period. The Technical Proposal will explain the physics motivations and describe plans for carrying out improvements, repairs and upgrades and installing them in the shutdowns foreseen in 2012-2013 and 2015-2016. It is due for submission to the LHCC in September 2010.

The Phase I upgrade will correct any deficiencies discovered in initial CMS operation. It will implement efficiency improvements that will at least partially compensate for lower luminosity in the early years of operation. It will ensure that CMS can take full advantage of luminosity increases that may take place at the end of Phase I. More specifically, operation at the highest Phase I luminosity will require upgrades to detector elements that will deteriorate because of radiation damage or will be affected by higher occupancies that complicate pattern recognition or lead to increased dead time. Trigger performance will also degrade as increased detector occupancy reduces the effectiveness of isolation and correlation algorithms. These considerations lead to the plan for CMS detector upgrades for the first phase of LHC operation. The plan must be flexible because the order of construction and installation may depend on details of how the CERN programme develops in the next few years.

The projects that CMS proposes to carry out between now and 2016 are:

Pixel System: Replacement of the current FPIX system with 3 forward disks on each side of the IR, redesigned to match a new 4 layer barrel system; to implement more radiation hard sensors; to reduce dead time in the readout; and to lower the material budget. This upgrade is driven by radiation damage (integrated luminosity) and peak instantaneous luminosity (data loss at full trigger rate). The reduced amount of material in the pixel detector results in improved vertex reconstruction and tracking efficiencies and better resolution at all luminosities.

EMU: Addition of chambers in the fourth Endcap Muon layer (ME4/2) to add redundancy to reduce the accidental trigger rate and to preserve a low p_T threshold for the Level 1 Muon Trigger; upgrade of the layer 1 (ME1/1) electronics with a new CSC Front End Board (CFEB) to include it in the trigger for added redundancy. Deployment of new muon trigger primitive electronics to extract additional muon segments for the upgraded CSC Trigger Track-Finder. The EMU upgrade is driven by considerations of peak instantaneous luminosity on the muon trigger.

HCAL: Implementation of depth segmentation to cope with the higher luminosities; use of a new sensor, the Silicon Photomultiplier (SiPM) that provides the high gain needed for segmentation; and new electronics, required to provide enhanced information to the upgraded Regional Calorimeter Trigger (RCT). This upgrade is directed at handling instantaneous luminosity, integrated luminosity, overall robustness and efficiency and providing opportunities to make improvements to the trigger at all luminosities.

Trigger: Rebuilding the Regional Calorimeter Trigger (RCT) using new technologies, such as μ TCA, which will permit fully flexible clustering and implementation of isolation algorithms; rebuilding the CSC Trigger Track-Finder to use more input segments and combine a greater variety of tracks to enhance performance amidst greater occupancy and backgrounds; and use of the new flexibility to permit eventual use of both upgraded Calorimeter and CSC triggers in combination with L1 tracking trigger information to produce triggers for very high luminosity operation. This upgrade is driven by peak instantaneous luminosity and overall efficiency.

DAQ: Increasing the bandwidth of the DAQ by a factor of 2–5 necessary to cope with increased peak luminosity.

The two phases of the upgrade, taken together, are expected to be similar in scale and in cost (in CHF of the year) to the original detector construction project. Perhaps 25–35% of this will be devoted to the Phase 1 Upgrade. Because SLHC upgrades are complex and technically challenging, significant R&D efforts are required. This will necessitate new concepts and developments that will keep CMS scientists on the cutting edge of technology. Because of the technical complexity of the project, the long lead time needed for construction, and the desire to avoid lost beam time while commissioning new devices, aggressive R&D and careful planning are mandatory.

A master schedule for the CMS upgrades does not yet exist, except as outlined above.

Up to now, the UK has focused its attention primarily on Tracker and Trigger activities, although there is obvious overlap with central DAQ and calorimeter changes.

2.3 UK adaptation to CMS planning

The UK must develop a plan which is matched to the needs of the overall CMS upgrade project. The accompanying Gantt chart summarises the CMS Tracker and Trigger upgrade planning which has been evolving rapidly of late, as well as the major steps in the LHC machine plans. It includes the main elements of WP1, WP2 and WP3.

Trigger and Tracker Technical Design Reports are expected in the next 1-2 years, following the evaluation of the CMS Technical Proposal, which will start in September 2010. There are a number of issues which must be resolved before a UK upgrade proposal is submitted and commitments taken.

The CMS central DAQ is considering upgrades to the S-links used between FEDs and the central switch (via FRLs). This has implications for these components, which are installed on the 500 Tracker FED transition cards, and the associated S-links. However, information is sparse and it is not yet possible to judge what this might require in effort or cost, or how it might be shared.

For the Phase I Tracker, the UK expects to contribute. There are two principal areas of interest, which could overlap: DAQ and online software and module assembly and qualification. A contribution to the DAQ would build on existing expertise in the Tracker DAQ, including firmware and online software. The new pixel system, with its higher speed digital links, could also require an upgraded FED for which the Mini-T card (WP3) might be a very good candidate or starting point. Module assembly could profit from existing infrastructure in the UK and strengthen certain core skills in RAL Technology Department. The cost and detailed implications of both options will be assessed in the coming year so a decision can be reached.

The prototyping of new modules for the Phase II Tracker and the expected lengthy evaluation and qualification period means that the CBC development remains very timely and there is little slack in the schedule, despite the expected installation date of about 2020 for the replacement Tracker. The modules will use CO₂ cooling, DC-DC conversion and very high speed, and expensive optical links, so must be well optimized. All these new features are very novel so a significant development period is expected. The same is true of the track-trigger requirements, which are still at a very early stage despite being thought very important for the functioning of the future CMS trigger.

The Phase I Trigger is expected to be based on μ TCA hardware, where the UK developments are more advanced than in the few other places where there is similar expertise. There are both new ideas for implementation of the trigger and much more advanced hardware which must be mastered, and software and firmware developed to meet the requirements. UK progress has been good and estimates

and outline plans must be developed so that the initial developments can be turned into a credible and affordable plan for the upgrade.

3. Work Package 1: Tracker and Level-1 Trigger Simulations and Software

3.1 Objectives

The objective of WP1 is to support with software tools the design, optimisation and prototyping of replacement detector and electronic systems for the CMS tracker and L1 trigger. Whereas a key focus for WP1 activities in the original proposal was the development of a L1 track trigger, this can take place only during Phase-II of the CMS upgrade. The project goals have therefore been refocused on support of Phase-I upgrades, along with software-focussed R&D in support of the overall CMS upgrade programme. Work on the long-term Phase-II track trigger development continues at a lower priority, but is not a key task for the current phase of the project.

The updated top-level work package goals are as follows:

- Development of new tools for simulation and optimisation of upgraded tracker and trigger systems in a very high luminosity environment
- Investigation and optimisation of tracking detector layout and inclusion of tracking data into the overall CMS trigger decision
- Provision of online and offline software tools to support the design and operation of upgraded L1 trigger systems for Phase-I, including the construction and operation of hardware prototypes produced within WP3. [Updated since last report, see below]
- Assessment of the performance of the upgrade CMS detector against key physics requirements and with realistic background conditions.

The overall strategy for WP1 has not changed; however, deliverables within the trigger software component of the work package have been adjusted to reflect the focus on Phase-I trigger upgrades. This has allowed an extremely effective synergy between the WP1 and WP3 development within the project, with the UK groups working to produce both generic modular trigger hardware and the software tools to allow this to be tested and further exploited both within and outside the UK project.

3.2 Progress to date

Simulation tools: A key CMSUK contribution has been the extension and reorganisation of the existing tracking software packages to allow the simulation and optimisation of a range of new candidate geometries for both inner (pixel) and outer tracking detectors. The original software was highly optimised for simulation of the existing CMS tracker. Over the last six months, the updated software has been used extensively in the simulation of the upgraded Phase-I pixel detector. A recent achievement has been the re-integration of upgrade simulation software with mainstream CMS releases, leveraging recent developments in the handling of pileup and offering more flexibility in the configuration of the event processing chain. This paves the way to the production of very large simulation samples which will be needed for detailed tuning of detector design and post-Technical Proposal physics studies.

The ultra-fast simulation tools for study of Phase-II stacked tracking layout and performance have been ported to the latest CMS software versions, which will allow their future integration with layout parameterisation tools produced elsewhere by collaborators in CMS.

Tracking simulation studies: In addition to development and validation of the overall tracking simulation, the UK has taken the lead in study of electron tracking performance in the upgraded Phase-I detector, providing estimates of tracking efficiencies, fake rates and impact parameter resolution. We have also contributed to the development of realistic material budgets and their inclusion in the simulation. Work has now started on the optimisation of track reconstruction algorithms, comparing new seeding techniques (taking advantage of the proposed addition pixel layers) with an optimised version of the existing proven algorithms. We have also begun to look at the b-tagging performance improvements offered by the upgraded detector, including the provision of fast

simulation code which can offer access to the parameterised b-tagging efficiency for TP physics studies.

Trigger software: Over the last six months, the effort on track trigger simulation has been reduced to a minimum, and results obtained from both full simulation and ultra-fast simulation studies are being written up for publication. The focus has turned to support of the demonstration of the WP3 modular hardware as a candidate platform for the L1 trigger upgrade. To this end, WP1 staff took responsibility for rapid production of basic online software tools to configure, control and monitor the hardware, and these tools are now being used successfully in the Mini-T testing. The same software is being used in development of a simplified test platform for the WP2 CBC ASIC.

This effort is first part of a more generic R&D programme into new online software tools for control of the new μ TCA hardware across CMS. In collaboration with other groups in CMS (Minnesota, Boston, CERN, Wisconsin) a new online software architecture has been proposed, which leverages ‘off-the-shelf’ hardware and software components to provide control and local DAQ functions via commodity gigabit Ethernet networks. The new architecture is designed from the outset to be highly scalable, using lessons learnt from the development of very high performance storage and network systems at the CMS computing centres; it uses software techniques and components from the telecommunications industry to provide the required combination of reliability and performance. Layered on top of this software stack is a generic object-oriented toolkit for control of generic modular trigger and DAQ systems, which will allow the WP3 hardware to be exploited for a range of applications across the upgrade project. We propose to pursue R&D in this direction over the coming months, culminating in the construction of a prototype system using many instances of both the next generation modular hardware and a large-scale simulation online environment; this will inform the online software technology choices that will need to be made within the next year by CMS.

Detector physics performance: This area of the project is becoming increasingly topical, as the decisions on Phase-I detector configuration have become concrete during the TP preparation. D. Newbold is co-author of the TP chapter on physics justification, and there is full input from UK physicists in the decisions on ‘benchmark channels’ for the upgraded detector. The UK has responsibility for overall coordination of upgrade simulation activities and is working with US colleagues to provide the large full and fast simulation samples required for the initial physics studies for the TP; this is technically challenging owing to the large computing resources required for simulation of high luminosity events. The physics simulation programme will continue beyond the TP publication later this year, with UK involvement mainly from RG-funded academics and PhD students.

3.3 Deliverables

Of the proposed milestones, the following have been met or effectively retired, as summarised above.

Year 1:

The contribution to the stacked-tracking simulation programme has been successfully completed, and further work assigned a low priority. The results are currently being written up.

The studies of overall tracker / trigger performance, and the identification of benchmark SLHC physics channels, have been subsumed into the overall CMS activity towards preparation of the Phase-I Upgrade Technical Proposal, with substantial UK leadership. This milestone will therefore be met with the publication of the TP in the autumn of 2010.

Year 2:

The development of full simulation code for the upgraded tracker has proceeded ahead of schedule, and the corresponding milestone is expected to be met in Autumn 2010.

The technical design study for the overall stacked tracking detector layout is now not required on this short a timescale by CMS. We propose to replace this milestone, and the corresponding milestone in Year 3, with a Year 3 target more relevant to the R&D programme towards the design of detectors and electronics for a L1 tracking trigger module.

First steps towards a common trigger emulation package have taken place under the aegis of the TP simulation programme. We are on schedule for a first software release by the end of Year 2.

3.4 Staff on project

Reported in tables. Note that the level of new RAL effort allocated to the project is likely to consume most of the accumulated WP1 staff underspend by the end of this project phase in December 2011.

In the last WP1 report, it was noted that it had not been possible to fill certain of the funded posts at RAL due to STFC recruitment restrictions. This situation persisted for most of the last six months, but has recently been successfully addressed via the transfer of existing staff from other RAL projects. These new staff are now being integrated into the project, working within the tracker detector simulation and online software areas.

3.5 Expenditure

The expenditure to date is dominated by staff costs, with some travel. It is reported in accompanying tables. Expenditure required in the integration of WP1 online software with prototype trigger and readout hardware in WP3 and 2 is covered in the budget of those WPs.

3.6 Comparisons with CMS activities elsewhere

The majority of WP1 deliverables are being addressed in close cooperation with CMS collaborators. In many cases, the UK is leading the relevant parts of the programme; in others, it is playing a vital technical role.

Simulation tools: The UK has driven the technical adaptation of the CMS tracker geometry description and simulation code to the new upgrade geometries. It has provided the first tools allowing studies of track trigger performance in the endcaps, and of overall data rates within the track trigger system.

Tracker simulation studies and Detector physics performance: The UK has joint responsibility for key aspects of tracker performance studies (electron tracking, b-tagging), in collaboration with US and INFN groups. CMS UK collaborators are active in the preparation of the corresponding chapters of the Technical Proposal, with D. Newbold a co-author of the Physics Justification section.

Trigger software: The UK is the project leader for trigger online software, and has been responsible for the design and delivery of all system components so far; there is interest from other groups (CERN, US) within the CMS trigger upgrade project in joining this project in the coming year.

4. Work Package 2: Outer Tracker Readout

4.1 Objectives

The objectives of WP2 are to develop a readout chip suitable for the outer tracker, to study options for providing Level 1 trigger data from a new Tracker, and to contribute to development of a complete readout system including off-detector components.

The immediate goals are to build a characterised pre-production prototype front-end ASIC, within the duration of the project, define and prototype an off-detector readout board (in CMS terminology, a Front End Driver) and prepare for testing the key parts of the readout in both the laboratory and test beams.

4.2 Progress to date

Front end chip

The readout chip development is for short silicon microstrips (2-5 cm) to be used in the outer tracker region ($r > \sim 50$ cm). These short strip outer layers would probably not contribute to the Level 1 trigger, although there have been recent suggestions to incorporate logic in future to allow to build “two in one modules”, where two closely spaced radially separated sensors are read by one ASIC and provide some of the functionality which would be achieved by pT modules.

Design of the CBC (CMS Binary Chip) started in March 2009. Specifications were developed in collaboration with the CMS Tracker in regular meetings. The architecture uses unsparisified binary readout for reasons of chip and system simplicity, and should offer the lowest possible power consumption. The front end amplifier has been designed to match both possible sensor polarities, and can tolerate DC leakage currents up to 1 μ A.

It has been designed in 130 nm CMOS in collaboration between Imperial College and RAL TD. Up to now only one RAL engineer was available to the project, less than foreseen. A full chip prototype (128 channels) has recently been submitted (5 July) after a single design iteration, a few months delayed compared to the plan summarised at the OSC of March 2010, but well in advance of the original project schedule. At the time of writing, wafers are expected in early October. The submission is part of a MPW run, dominated by a RAL XFEL chip, but including several other smaller circuits.

The CBC includes

- a fast front end amplifier, with 20ns peaking time,
- comparator with programmable threshold,
- 256 deep pipeline,
- 32 deep buffer for triggered events,
- output MUX and driver, using a low power signalling standard (SLVS),
- fast (SLVS) and slow (I2C) control interfaces.

It also incorporates a power block and a low drop-out regulator from CERN, which can both be disabled and by-passed, to explore the impact of local DC-DC switched voltage regulation.

Preparations are now under way for components needed to study the CBC when it returns from manufacture.

The first priority for the coming year will be to evaluate the CBC and decide what aspects require more attention or improvement. There are several options under consideration with CMS collaborators for further development to match it to the long term requirements of future Tracker modules. These include coarse pitch bump bonding for ease of module assembly, consideration of a 256 channel layout and possible trigger functionality, in addition to further development of interfaces to the readout and control, which is at an early stage compared to the final requirements. A critical issue will be evaluation of the DC-DC conversion and power regulation, which are expected to need work to optimise them for a final chip. In short, although the CBC will not be installed into CMS until about 2020, the development work which must precede this is considerable, especially given the lengthy module studies expected.

Power and module developments

DC-DC conversion was adopted by CMS as the baseline powering scheme for the future. This has already proved beneficial, since some of the developments can be exploited in the Phase I pixel system where it would have been difficult to supply sufficient power at 2.5V for a 4-layer barrel detector.

Bristol have investigated alternative magnetic components for DC-DC converters that might have intrinsically less stray field and EMI fabricated in a printed circuit board. This work has now ended as the potential show-stoppers foreseen in the buck conversion scheme seem to have diminished in importance and progress with the alternatives did not look promising in comparison.

Bristol procured a TTP board designed to test VFAT2, which is a CERN-designed TOTEM front-end chip with similar functionality to the proposed CBC. We obtained the circuit schematics, source-code for the firmware and readout software. It has become clear that this is a complex system which may be difficult to maintain. An alternative based on an FPGA evaluation board is being studied, and may be a better and cheaper alternative. It also has potential synergies with the online software effort within WP1 and WP3.

SFED developments

During year 1 of the project a shortage of design effort at RAL meant there was little development, or resource usage, for the off-detector readout. However, in view of the evolution of the LHC planning

and developments in this R&D project, this has not been disadvantageous. In particular, the rapid progress with the WP3 Mini-T development provides a potentially better starting point for FED development than acquiring commercial development boards for a prototype FED which in any case will not be needed for some years, although certain important functions should be studied in the interim.

During the period since our proposal, the CERN Gigabit Transceiver (GBT) system has evolved considerably further. The GBT provides the ASIC family and system design required for readout, control, timing and trigger transmission for future standard data and control links and will be implemented along with components developed by the CERN-led collaborative VOL (Versatile Optical Link) project. We expect the GBT and VOL to provide the basic infrastructure for many of CMS links, especially for the Tracker.

CERN have recently developed GBT ASIC logic in an FPGA; the source code is now at RAL and under study by a TD engineer. The objective would be to run this on a commercial development board to study a working GBT link on the bench in RAL. This is essential to master the non-trivial protocols and data transmission arrangements for the CBC and other parts of the Tracker.

L1 trigger system

The importance of providing information to the level 1 trigger in CMS remains a priority and the components, such as tracker modules, to achieve this are a formidable challenge. In the last year, CMS simulation studies, including those of WP1, have continued to improve our understanding of the requirements, for which the PT module concept remains the most potential means to deploy.

Working groups have met regularly to develop plausible design concepts and the key elements have been identified. These include commercial assembly using through-via technologies and the design of a suitable pixellated ROC and sensor for coarse pitch bump bonding and evaluation of an overall architecture, where power dissipation and material are of great concern. Work on this topic was scheduled for the second year of the programme and given the progress with the CBC, which was the top priority, some attention can now shift to trigger module developments, as TD microelectronic staff should become more readily available after September, as other projects are completed.

Preparations for test-beam activities

Evaluation of modules in test beams will be needed, especially when trigger modules are developed. This task can be shared with other CMS collaborators, for example the study of a detector module read out by CBC chips. However, it is a lengthy task to prepare the ancillary equipment needed to instrument a beam test, in particular the data acquisition system and software, and it is profitable to begin now since many suitable items of equipment are, or can be made, available now but will not be so accessible in future being based on actual CMS hardware.

We are assembling a beam telescope of up to 5 stations (pairs) of x-y planes using CMS components (namely APV25 hybrids, optical links and opto-hybrids and Front End Driver) with custom control electronics. Two stations would be mounted upstream of a module under study, and two or three downstream. Sensors acquired from the D0 upgrade project from FNAL, which was cancelled, which have finer pitch than those used by CMS, have been used. The DAQ is based on the CMS Tracker system. Two stations were tested successfully in a CERN North Area 400 GeV/c proton beam in June. We plan to use them in a test beam in September 2010 in collaboration with UA9 for measurements of crystal channelling for LHC collimation development.

In a later phase, we will develop the corresponding DAQ to read CBC modules, read with GBT and VOL high speed digital links.

4.3 Deliverables

Of the proposed milestones, the following have been met or effectively retired, as summarised above.

Year 1:

Documented system conceptual design and performance specifications; although the outer tracker readout system is not specified in final detail, it has been agreed to use the GBT link system which is

documented in detail by CERN. The specifications for the CBC have been agreed with collaborators and are available by URL.

The requirement for front end and other test structure circuits has been removed.

Documented results of preliminary investigations of powering schemes have been reported.

Year 2:

The deliverable of a full chip prototype designed and submitted for fabrication has been brought forward considerably and thus already met. This minimises some of the longer term risk by offering more potential for a subsequent submission for tuning or modifying the design. Test setups are in preparation and will be available when devices are delivered.

4.4 Staff on project

Listed in the tables. One significant change is that C. Hill will be leaving Bristol at the end of the present academic year. The contributions of the Bristol group to WP2 and the project as a whole will be reviewed as the staff situation becomes clearer.

4.5 Expenditure

Listed in the tables, up to April 2010. The most notable recent commitment has been the MPW run contribution, 57kCHF.

4.6 Comparisons with CMS activities elsewhere

The ASIC work in CMS is limited to a few institutions. At present, the main ongoing effort is at PSI in the pixel system for Phase I so with no overlap with the UK effort. CERN has a significant capability in microelectronic design but does not yet have any major commitment to CMS, although there is interest in the PT module development and very much we hope CERN will become an active collaborator. The US has an engineering team in FNAL but they are currently not active in the Tracker, except to a limited extent in the pixel system. There is interest in a few other institutes, such as Bergamo, Lyon, Cornell, in ASIC development for the long term but there is no serious work presently underway. Thus the CBC is a unique activity in CMS, which would be hard to replace and is also significantly advanced compared to any other design effort.

Regarding off-detector electronics, Vienna was responsible for the present pixel FEDs and will continue to support them, though resources are very limited. CERN developed the Tracker FECs, used also by other CMS sub-detectors. US groups have expertise in board design and firmware, though not currently applied to the Tracker. In the Tracker collaboration, there are no other institutes yet active, although there is interest from Strasbourg in contributing some engineering effort for PT module readout. The comparison between the Mini-T module work and work elsewhere in CMS is explained in section 5.6, but the UK effort is also in the vanguard of such developments and seems likely to remain so.

5. Work Package 3: Design of Level-1 Calorimeter Triggers at the SLHC

5.1 Objectives

Work Package 3 is developing a demonstrator for upgraded Level-1 Trigger electronics and studying new trigger algorithms using this system. The main goal is to provide a standard device for all calorimeter Level-1 triggers which can be adopted during the Phase I upgrade in the first instance.

The platform now widely adopted in CMS without controversy is the μ TCA standard which is used by the telecom industry and whose data transfer capabilities are in the multi-Gbps range. Our system consists of generic modular devices based on FPGAs and optical links and is the first prototype of off-detector trigger electronics for SLHC.

5.2 Progress to date

Since the last status report there has been good progress with the system design and underlying hardware. In addition a good working collaboration has been established with other institutes to advance the core infrastructure.

The prototype main processing card (Mini-T5) was delivered in Feb 2010 and has been successfully validated in a μ TCA crate with 32 optical links operating at the rated specification of 5Gb/s. As discussed in the March 2010 report, the design of a custom backplane was postponed until test results from Mini-T5 were available. The successful tests, plus the extra I/O capability of this card, allow us to use a commercial backplane reducing risk and dependence on specialist manufacture. They also enable more cards to be built so that firmware and software engineers have a solid platform to develop against while the final card is in design, manufacture and test. This is essential if the firmware and software are to keep pace with the hardware.



Fig. 5.1. The Mini-T5 card, which was designed and laid out at Imperial College and manufactured by eXception PCB Ltd (Tewkesbury UK).

The system design has benefitted enormously from implementing trigger algorithms, and consequent improved knowledge of the FPGA resources required. In addition recent announcement from Xilinx of the Virtex 7 FPGA series allows simplification of the system design by not requiring modifications to the standard μ TCA backplane. This has other benefits, of which the most notable is that increase the bandwidth and processing capability of the trigger can be achieved simply by adding more hardware, albeit with a small increase in latency each time.

There is coalescence within CMS around a common vision of how to use the latest high speed serial backplanes, such as μ TCA, in which the UK has taken a leading role. One result will be a CMS Note to be published in the coming months.

The CMS implementation of μ TCA is based on COTS (Commercial Off The Shelf) parts wherever possible. The backplane is a standard redundant telecom μ TCA crate and MCH (μ TCA Controller & Hub). However, the crate is operated in a non-redundant mode with the additional MCH slot used to provide the additional connectivity necessary for CMS. E Hazen (Boston University) and R Hirosky (University of Virginia) have built, or are in the process of building, prototype cards to provide the service functionality required. The key aspects of the hardware therefore seem well supported.

A collaborative effort has also sprung up to support a simple common communication protocol running over GbE (Gbit Ethernet) that was inspired by firmware from J Mans (Minnesota University).

The UK has again taken a leading role by writing software to ensure that key components have the robustness and high performance demanded of them.

For the rest of the year our focus will be on building a μ TCA based demonstrator system and inserting the algorithms into it. The card already has GbE communication within a μ TCA crate and a basic software interface provided by WP1. A scalable software structure is under development via a collaboration of WP1 and WP3 staff. Algorithms for the electron and photon identification are written, but further work has been postponed until the software structure is in place.

5.3 Deliverables

Of the proposed milestones, the following have been met or effectively retired, as summarised above.

Year 1:

Design of the main processing card ready for production.

Backplane design ready for layout. (no longer relevant)

Preliminary version of the firmware ready.

Year 2:

Working prototypes of processing card and backplane ready

Further deliverables have been added to the overall plan in view of the overall progress to manage intermediate objectives.

5.4 Staff on project

Reported in accompanying tables. The WP manager C. Foudas will leave Imperial in August 2010. Meanwhile the PI has taken responsibility for overall supervision of WP3 activities.

5.5 Expenditure

The major items of expenditure are the optical connectors and further Mini-T card orders, which are not yet invoiced, but the overall spend is still well within the foreseen envelope.

5.6 Comparisons with CMS activities elsewhere

At present the Mini-T card is a technically more advanced card than anything else in CMS (and perhaps beyond, in LHC) but there other (US) groups who are aiming to increase their technical capabilities in the digital hardware and firmware areas.

The Wisconsin group have manufactured a μ TCA card, currently with only DAQ functionality in mind. It does not use optical links and is limited to 16 transceivers at 3Gb/s so significantly less powerful than the Mini-T. Wisconsin have played a big role in the present trigger and are responsible for the Regional Calorimeter Trigger (RCT), based on GaAs ASICs and electrical ECL links, so rather dated and with limited spares as components become obsolescent. They are gaining firmware expertise, although it has not so far been used in CMS.

The Minnesota group have developed a μ TCA FED for the HCAL upgrade. The Minnesota card (MINI-CTR2) is also based on a Virtex5 with GTX transceivers (up to 6Gb/s). However the MINI-CTR2 has 16 links whereas the Mini-T has 40; the Mini-T also a large parallel connection to provide a low latency interface that makes it more challenging to design and manufacture (but so far very successfully).

Regarding the trigger philosophy, we have advocated a new approach of time multiplexing the data which has significant benefits while others favour a more conservative approach that would retain the current structure. Time multiplexing concentrates the entire calorimeter trigger data for a given bunch crossing into a single FPGA and dispense with boundary conditions between processing nodes and thus the custom backplane. (It is analogous to event building by the CMS DAQ which transfers data from many inputs to a single location for processing.) It makes the system extremely flexible (e.g. there are no limits on algorithms due to limited input data as in a conventional trigger). The entire trigger would remain a very powerful image processor that could be reconfigured by simply changing

the data selected within the multiplexing units. It also allows special triggers to be placed in parallel with the main triggers for custom searches (e.g. with dedicated hardware).

6 Risk register

The current version of the project risk register is v2.1, revised from the version presented at the last meeting.

We note the guidelines, in particular: “Risk analysis at the start of the project helps managers to consider what might have an impact on their plan and stop it being completed and so allows them to put mechanisms into place that stop this happening. However as the project progresses new risks will occur and some may recede and so the register needs to be reviewed regularly.”

It is notable that some conspicuous risks to this project concern issues over which the PI and managers have little influence, namely CERN plans for the machine evolution and STFC funding. The corresponding actions have been to make the planning responsive to changes and to deploy a working margin which can be called upon.

In addition, given the technical nature of the project, it relies on specialist expertise and key individuals. Although this risk cannot be eliminated, new staff are further trained during the project, e.g. in firmware skills or online software, and as the ASIC designs evolve, the expertise is distributed.

However, it should be noted that this is an R&D project, designed to permit the UK to contribute significantly to a future upgrade of the CMS detector in areas where the UK and its teams already have significant expertise. As such, the definition of risk should be considered differently than, for example, a construction project for a number of reasons:

commitments have not yet been made, either to CMS construction responsibilities or for deliverables to be used on a large scale,

the purpose of the R&D is to develop objects such as trigger hardware or readout ASICs, with ancillary supporting services, such as firmware and software, on a best-efforts basis, which necessarily includes a risk of a non-successful outcome, because of novelty and technical challenge,

the risk of a non-successful outcome is mitigated during the project by identifying the most promising routes forward and adapting the component development and the requirements of CMS, in the most challenging cases. Progress so far demonstrates that, to date, this has been effective.

6.1 Risk Management

The guidelines recommend the following actions:

- Modify the project plan
 - elimination – risks are removed so they are no longer a threat to the project
 - reduction – action taken immediately to minimise the risk
- by putting in place an action to be followed if the risk materialises
- by transferring risk to someone else
- by accepting the risk – where the costs of taking action outweigh the benefits

The first action is built into the planning from the outset while considering the final action, since otherwise some challenging parts of the project such as the CBC and trigger hardware (endorsed strongly by referees) would not have been started.

The transfer of risk to others is in some cases feasible but not at this stage desirable in view of the long term strategy for UK future participation in the upgraded detector. For example, there are a few CMS collaborators in the trigger hardware who are interested in developing significant skills in the same area and have begun complementary developments. For the ASIC work, CMS has few collaborators presently resourced to take on the CBC development but it is not inconceivable that this might change in a future difficult situation.

Thus, the most reasonable action at this stage in most cases is the second proposal, to put in place a suitable mitigating action. The “existing controls” itemises the proposed actions being followed. Technical progress has been good on the two major costly items (CBC and trigger hardware)

indicating a sensible appreciation of the risks. There is a relatively generous spending reserve in the Working Allowance to recover from unsuccessful ASIC runs or trigger hardware failures. Each line of the risk register has been reviewed recently and comments have been added where appropriate.

7 Finances

The financial report is summarised in the attached tables. WP2 and WP3 spending is well within the foreseen envelope for requisitions and staff spending is roughly in accordance with expectations given the late start. However, staff spending at RAL is slower than anticipated because of availability of engineering staff.

For WP1, staff spending within University groups is approximately on target. A substantial staff underspend was accumulated at RAL PPD due to recruitment issues; however, new staff are now in place, and this underspend is due to be absorbed over the remaining duration of the project. The WP1 equipment funding has not yet been utilised, though we anticipate that this will be necessary during the coming year in order to accommodate computing requirements.

We note that, at the time of writing, the new STFC finance reporting system provided via the Shared Services Centre is not yet operating so reports of spending in RAL are presently not accessible.

A Working Allowance of £230k has been set aside using categories where some of the spending is anticipated to come late in the project. These funds will be released as further progress is made and can be redeployed if needed as a protection against problems with ASIC submissions or trigger hardware manufacture.

The future spending profile has been estimated making assumptions about staff in post and likely timing of large purchase commitments. Although this should be adequate for future estimates, it is kept under review. In particular, the principal limitation to extending the profile comes from staff availability for those paid from the project.