


University of Bristol  
Brunel University  
Imperial College London  
Rutherford Appleton Laboratory



## **R&D in preparation for upgrades of the CMS detector for High Luminosity LHC**

1	Executive Summary .....	2
2.	Project history and recent developments .....	2
2.1	LHC upgrade schedule and planning .....	3
2.2	CMS planning .....	4
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 .....	6
3.3	Deliverables .....	8
3.4	Staff on project .....	8
3.5	Expenditure .....	9
4.	Work Package 2: Outer Tracker Readout .....	9
4.1	Objectives .....	9
4.2	Progress to date .....	9
4.3	Deliverables .....	11
4.4	Staff on project .....	12
4.5	Expenditure .....	12
4.6	CMS activities elsewhere .....	12
5.	Work Package 3: Design of Level-1 Calorimeter Triggers at the SLHC .....	12
5.1	Objectives .....	12
5.2	Progress to date .....	12
5.3	Deliverables .....	13
5.4	Staff on project .....	13
5.5	Expenditure .....	13
5.6	CMS activities elsewhere .....	14
6	Risk register .....	14
7	Finances .....	14

## 1 Executive Summary

The proposal was approved in Spring 2008, and effectively began in April 2009. Notable progress since the last report includes:

A CMS Upgrade Technical Proposal was prepared and submitted to the CERN LHCC; it is now under review.

There has been further evolution of the LHC schedule for the next decade, whose implications are not yet fully clear, after a very successful first full year of LHC operation. Valuable insights into issues likely to be of importance for the future have been gained as a result of experimental operation.

There has been good progress in developing software and simulation tools for this R&D project, as well as CMS studies in general. A strategy for coherent development of online software for the trigger upgrade has been proposed, and a first prototype demonstrated. The close coupling between software and firmware seems to be a good and effective model to follow. Demonstrating the physics benefits of proposed upgrades to the experiment requires further work to which UK groups are contributing.

The 128 channel CBC (CMS Binary Chip) was submitted for fabrication in July 2010 after a single design iteration. However delivery from the foundry was delayed; wafers arrived in CERN just before the end of 2010 and die arrived back from cutting this week, so are now awaited for testing. A beam telescope for module tests was built and operated in CERN beam tests with excellent results. Work has begun on firmware to utilise CERN-developed high speed links expected to be widely used in CMS in future.

The Mini-T5  $\mu$ TCA card has been in use in test setups, expediting software and firmware development. A more advanced prototype is in design for manufacture in mid-2011. A proposal for a new time-multiplexed trigger architecture which offers many benefits over a conventional architecture has been developed and is under discussion in CMS. A US-UK collaboration is making more detailed plans for the Phase I trigger upgrade.

## 2. Project history and recent developments

The original proposal was submitted to STFC in October 2007 with an upgrade of the LHC peak luminosity from  $10^{34} \text{ cm}^{-2}\text{s}^{-1}$  to  $10^{35} \text{ cm}^{-2}\text{s}^{-1}$  about ten years after start-up in mind. The accelerator upgrade has since been proposed to take place in two main stages, with an increase in luminosity to  $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$  being Phase I, occurring in 2016 in the most recent CERN ten-year plan.

Both the LHC planning and CMS continue to adapt to machine operation and performance as well as physics needs. The LHC had a very successful year of operation in 2010, reaching  $47\text{pb}^{-1}$  of integrated luminosity delivered to CMS, with  $43\text{pb}^{-1}$  recorded. The instantaneous peak luminosity rose by several orders of magnitude during the year, culminating at about  $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ . CMS systems ran very efficiently and the operations team was able to concentrate much of their effort on identifying ways to avoid losses of data taking time and ensuring high quality performance.

There was also about a month of heavy ion running with Pb ions and  $9.2\text{mb}^{-1}$  was delivered by the LHC and  $\sim 8.3\text{mb}^{-1}$  of data collected by CMS. Overall data taking efficiency was over 90%.

In both heavy ion and proton running various difficulties were encountered, which were successfully overcome. However, we can expect that there will be additional challenges ahead, such as handling the electron cloud effect or collimation issues, as the luminosity grows. Some of them may require modifications to CMS operations to adapt so vigilance will continue to be essential.

The hardware for which the UK is responsible continues to be generally very reliable. It needs strict, though not excessive, attention. It is interesting to note that in large systems such as the tracker off-detector readout FEDs that individual boards and firmware and software can be thoroughly tested and debugged, yet small issues at sub-1% level can appear which need lengthy attention to isolate and correct. Most of them to date have been related to subtle timing issues, for example the phase of the FED clock relative to that used to drive the S-links from FEDs to the DAQ FRLs (Front End Links). Occasional resynchronisations are required in CMS (an example is discussed below) which brought to

light inconsistencies in the prescribed reactions by most sub-systems and necessitated firmware modifications. It is not easy to prove definitively but it seems likely that in such large systems small manufacturing variations at board level, or within components such as FPGAs, give rise to very subtle internal timing differences which only become evident during full system operation and hence are even more difficult to trace. Nevertheless this has been achieved to a remarkable degree.

Issues of a similar nature appear in the trigger system. Some optical components have shown less reliability than expected but with too small a sample size to indicate whether this could become a concern. Known problems with quality of connectors on heavy SCSI cables, transferring data to the GCT, are one source of occasional difficulties, as are links from the GCT to the Global Trigger, which run very close to the manufacturer's specifications. The GCT-GT links will be replaced by optical links developed by the GCT team, but these are still being commissioned in CMS.

The importance is in the real-life operational experience which emerges. However well a system is designed, unexpected features arise from its complexity or scale which are impossible to plan for and it is easy to ignore such points in planning for upgrades, especially when extrapolating to performance or scale well beyond the present.

A classic example of an unexpected problem in a very well designed and constructed system was the appearance of "PKAM" events in the pixels, which also affected ATLAS. In early LHC running, large events were noted where shallow particle trajectories close to the beam line lit up a whole row or column of pixels, mainly in the barrel region. These were soon attributed to beam-gas events consistent with the high quality vacuum present in the LHC, so were not anomalous in rate. However, they caused significant problems because the pixel readout did not cope well with very high occupancy events, which required very long readout times, and the DAQ effectively gave up waiting for data when a module did not appear to be responding.

The solution adopted was to pause data taking and resynchronise the experiment. This was initially a time-consuming process and firmware modifications in the pixel FEDs were introduced to speed it up. However, for a time it was found impossible to make all pixel FEDs behave identically, despite all working individually. Eventually internal CMS discussions led to the suggestion to slow down the response of the pixel FED FPGAs, as done on the Tracker FEDs. This immediately solved the problem (and has been passed onto ATLAS).

The point of this preamble is to emphasise the importance of engaging operations experts in upgrade developments, as our team is doing, and to ensure that data quality monitoring and feedback is followed closely. The efforts in firmware and online software have important consequences for new developments, and it will be essential to maintain the full breadth of expertise accumulated in the past years. New problems can be expected in the future as the LHC pushes to higher luminosity, with greater pileup in data taking. Modifications for heavy ion operation may be needed. Experience at the system level is vital to integrate into developments for the upgraded systems.

## 2.1 LHC upgrade schedule and planning

Since the last report, there have been further changes in the overall planning. The new machine operation schedule released in July 2010 has been revised, although all details are not yet clear. It is confirmed that the LHC will run in 2012 so the shutdown foreseen during that year, to replace current-carrying splices and other maintenance tasks, will be postponed for one year. This will prepare the machine for operation at full energy of 7 TeV per beam and was foreseen to require about 15 months; it now seems likely to exceed that by a few months because of the volume of work needed.

The second major shutdown, foreseen for the end of 2015, will slip accordingly. This would prepare the LHC for higher luminosity operation beyond the nominal  $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ .

A new ten-year plan will be prepared by CERN but will take a couple of months. One of the issues to be better understood is how to carry out the maintenance work to be done on or around the cold machine elements. It is undesirable to warm up the whole machine, which would be subject to mechanical expansions and changes in stresses.

Both of these changes have an impact on CMS planning, since the Technical Proposal prepared in late 2010 made plans based on the July 2010 schedule. Several of the early CMS activities, such as completion of the muon stations and modification of the HCAL photosensors might be brought

forward and be fully complete in time for resumption of running in 2014 or 2015. However, the speed with which several of these modifications can be implemented depends also on availability of funding, in some cases mainly in the US, and this is also uncertain. However, these uncertainties do not yet have a major impact on the UK activities and planning

## 2.2 CMS planning

The CMS Technical Proposal for the Phase I upgrade, including a summary of activities underway for the Phase II period, was given to the LHCC in final draft form during the autumn and presented to them in November. Written questions were received and answers prepared which have recently been sent to the committee.

Naturally, the LHCC requested more information on details of the upgrades and planning, responses to which have been given but some of which will be provided in full at the TDR stage. However, there were also questions about the physics motivation, where the justification was considered to be incomplete. Since there are high expectations of discovering new physics in the coming few years, this is one reason. Another is the difficulty of completing sufficiently detailed simulation studies in a new detector configuration. For example, it is clear – and demonstrated in studies in the Technical Proposal - that a significantly reduced mass 4-layer pixel system will be more effective than the present one. However, given the challenges involved in building it, it is hard to demonstrate a simple cost-benefit ratio. Similar arguments can be applied to the trigger, or other sub-systems.

More information will be provided on the physics impact of the proposed upgrades using a couple of physics channels for the March LHCC meeting, but this effort must continue, including gaining help from some of those contributing to topical analyses in CMS.

The channels chosen should contain final state b-jets and tau-leptons. The channels should be relevant for  $>100 \text{ fb}^{-1}$  operation at full centre-of-mass energy. Higgs Boson decays to b-jets and taus have been chosen with the following two physics topics:

- MSSM heavy Higgs ( $M_A=500 \text{ GeV}$ ,  $\tan \beta = 30$ ) produced in association with two b-jets and decays to a tau pair. TDR studies showed this will need substantial luminosity to close the high  $M_A$  phase space, while CMS has already started to cover the low mass space now.
- A standard model low mass Higgs produced in association with vector bosons in order to measure the Higgs Yukawa coupling to b-jets.

The choice of these two channels was made because there is expertise within the upgrade community in analyses being prepared with 2010 data, and for 2011. It is important to profit from existing analyses because it takes several months to develop and validate a new analysis. However, challenges remain as availability of physics analysts to work on simulated data is in competition with opportunities to study real LHC data, and preparing software with tuned settings for upgrade configurations is a major task.

## 2.3 UK adaptation to CMS planning

A master schedule for the CMS upgrades does not yet exist and, given the LHC planning, is not expected to be clarified for some months. However, broad-brush plans for most sub-detectors were summarized in the Technical Proposal. Trigger and Tracker Technical Design Reports are expected in the next 1-2 years but no definite dates have yet been proposed.

Up to now, the UK has focused its attention primarily on Tracker and Trigger activities, and this is expected to continue. Clearly the WP1 work is highly relevant to providing a strong physics case and we would like to strengthen this activity if possible. This is limited by staff available. The main areas for our future contributions are:

The CMS central DAQ will upgrade 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. This will be a small hardware project but with implications for the FED firmware and maintenance, and for our long term commitment to the Tracker DAQ which is crucial to CMS.

For the Phase I Tracker, the UK has some interest to contribute. An effort on the pixel DAQ would mesh naturally with our expertise and existing commitments to the Tracker DAQ and is one item under discussion. Pixel DAQ efforts could also be expanded elsewhere within the UK, especially RAL, and could evolve from the activity described later in the WP1 report. A more ambitious project would be to engage in pixel module assembly, and profit from existing infrastructure in the UK and strengthen certain core skills in RAL Technology Department. The cost and detailed implications of this will be assessed in the coming half year so a decision can be reached. If it is realistic, a new Work Package should be defined.

Although Phase II, or the evolution of the LHC to high luminosity running, is still remote, the need to replace the entire Tracker means it is of the utmost importance to maintain the R&D and prototyping activities of WP2.

The Phase I Trigger upgrade should be based on  $\mu$ TCA hardware, where UK developments are very advanced and there is leading-edge 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. We are collaborating closely with Wisconsin to define the future calorimeter trigger and expect to be joined by Vienna to harmonise efforts for the full implementation of the global trigger system.

The date for a submission of a UK construction project grant is under discussion, with the aim to be coherent with the rest of CMS. As usual, most national circumstances are slightly, or even considerably, different. The main concern at present is the duration of WP1, especially given the importance of trigger software and physics simulations. It has become very evident that the software development and simulation studies are essential to the overall project development yet WP1 has a shorter duration than the other work packages. It would be helpful to rationalise this anomaly. We seek comment from the Oversight Committee on the possibility of submitting a bridging project proposal in which all UK groups would participate, or other means by which interim funding might be provided, to sustain the software and simulation activities until the longer term plans for construction can be better defined. This would reduce the pressure on some critical staff posts.

In view of the uncertainties, the Gantt chart has not been revised. The WP activities are coupled but not heavily interdependent. The critical items to clarify are the new LHC schedule and the areas where the UK wishes to concentrate. Much of this, especially the Phase I trigger upgrade, is clear.

### **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. The project is now mainly focussed on the Phase-I upgrade; however, the tools and ideas for Phase-II which were developed in the early part of the project continue to be exploited as a background activity.

The top-level work package goals are as follows:

- Development of 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 CMS trigger decision
- Provision of online and offline software tools and firmware to support the design and operation of upgraded electronics systems for Phase-I, including the construction and operation of hardware prototypes.
- Assessment of the performance of the upgraded CMS detector against key physics requirements and with realistic background conditions.

Since the last report, the key objectives of WP1 have not changed. However, the project has developed and expanded in two areas, which link closely to the goals above:

- The strategy for online software for the upgraded electronic systems in Phase-I has been developed, and a first prototype demonstrated in conjunction with WP3. This approach mandates a very close coupling between software and firmware components, which has resulted in an increased firmware activity within WP1.
- Interest in support and development of readout components for the upgraded pixel detector has become more evident, with the UK taking some initial actions in this area. Details are given below. We report this activity within WP1, though it is clearly cross-project in nature.

### 3.2 Progress to date

*Simulation tools:* A significant contribution has been made by CMSUK to the refactoring of the CMS tracking software, for both simulation and reconstruction. The original software was highly optimised for the existing CMS tracking system, and was insufficiently flexible to allow the study of a range of new upgrade geometries. This contribution is now reaching completion, with the new tracking tools in constant use, mainly for the tuning of upgraded Phase-I pixel geometry.

The CMS offline strategy calls for a convergence during the first part of 2011 between the ‘mainstream’ CMS software and the new code for upgrade simulation and reconstruction, and this work is under way. This will improve performance and flexibility of the existing software, and ensure that upgrade tools are always synchronised with the latest versions. However, this does put pressure on providers of upgrade software, and the UK in particular, to support and maintain the new code within the regular CMS release cycle. Dedicated effort will be required for this task in the future.

*Tracking simulation studies:* The UK has taken a lead in the study of electron tracking performance in the upgraded Phase-I pixel detector, providing estimates of tracking efficiencies, fake rates, and impact parameter resolution. It is clear that the ‘busier’ environment at the upgraded LHC will require new and more efficient reconstruction algorithms, especially if information from a fourth pixel layer is to be used optimally. The UK work now encompasses improvements to code efficiency (both for reconstruction and detailed validation of tracking performance), which will be essential for exploitation of the new pixel detector both offline and within the High-Level Trigger system.

The UK has also taken responsibility for preparation and testing of the simulation and reconstruction steps for the next round of Technical Proposal physics studies (see below).

*Trigger software:* The effort started in 2010 in online software R&D has begun to produce concrete results. The goal of this project is to simplify the structure of the lowest-level online software (‘hardware drivers’) by implementing a common framework usable by all groups adopting the new CMS-wide  $\mu$ TCA electronics standard. Since almost all new developments rely upon large FPGAs and fast serial links, we have proposed a very close coupling between software and firmware components, with both sitting within a standardised framework. The overall goal is to strongly reduce the duplication of effort in software and firmware, and the assist in the long-term maintenance of those components. This approach is likely to be adopted by CMS as part of the overall upgrade electronics strategy.

A first demonstrator system for this approach has been constructed, and demonstrated in the context of the Mini-T hardware from WP3. This work was documented in a TWEPP conference publication in late 2010. Work is now focussing on the packaging and first formal release of the higher level driver software framework (‘Redwood’) so that other groups may begin to test it. The lower-level component of the framework (the ‘control hub’), which allows the system to scale to very large numbers of crates and boards via a switched gigabit Ethernet network, is undergoing testing at realistic scale using off-the-shelf FPGA boards. Test setups have been constructed at Bristol, Imperial and RAL. In addition, there is a new effort to investigate the optimal transport protocol layered above Ethernet. Novel ideas such as the use of ATAEOE are under study.

*Detector physics performance:* The upgrade Technical Proposal was submitted in late 2010, focussing mainly on the mechanical and electronic requirements of the new pixel, HCAL and muon detectors for the Phase-I upgrade. Subsequent discussions with LHCC have indicated that a wider effort to study the physics performance of the detector defined in the TP is now appropriate. Such a

study naturally makes significant demands on the availability and quality of simulation and reconstruction software for each subsystem; in some cases, this is still work-in-progress.

Three benchmark physics channels have been defined for study at  $2.10^{34}$  luminosity. The focus is on Higgs observation in channels involving  $b$ ,  $\tau$  and photons. Whilst a wide range of physics scenarios may be of interest at high luminosity, these have been chosen in order to most clearly demonstrate and tune the performance of the Phase-I upgrade (including that of the L1 trigger hardware). The UK is leading the preparation of large simulation samples in support of the physics studies, with a first report due to be submitted to the LHCC in March 2011.

*Support of pixel module readout:* As an additional development in the Phase-I pixel upgrade project, RAL has agreed to begin some initial activity in support and development of lab-based readout systems for single-module testing. Such work is important for practical support of pixel detector construction, as module-level testing was already proven to be a vital component of quality control. The project has potential links with online software R&D within WP1 and, possibly, hardware developments within WP3. It could also help to build expertise for any future involvement in Phase-I or Phase-II pixel readout or detector involvement.




Figure 1: Pixel module readout system under test at RAL

This work is taking place in the context of the new ‘pixel lab’ established at RAL PPD, which supports work across a number of projects inside and outside HEP. In conjunction with CMS collaborators, a hardware / software setup for single-module readout has been assembled. This system is based upon a board previously produced at PSI, and used during construction of the existing pixel detector. The adaption of this system for the upgraded detector modules is currently under study, and is likely to involve a combination of hardware and software development.

The pixel readout project is clearly in its infancy, and is explicitly intended as a possible route to participation in the future construction or integration effort for the upgraded pixel detector. Since it draws upon WP1 staff resources, which are limited, it represents a significant broadening of scope for what was previously a software-centric work package. It is therefore intended that, after the current initial tests are completed, a detailed proposal justifying the use of project resources in this area be written. This proposal will cover the developments anticipated in 2011, and provide a forward look to plans beyond the scope of current WP1 funding. It should provide a better understanding of the viability of future activities on the pixel project and allow other UK groups to decide if they wish to contribute.

*Tracker layout tools:* In view of the long term needs to develop a complete concept for a new Tracker which is likely to include triggering capability, there have been a number of ideas presented in CMS, several of which have had substantial UK involvement. One of the most productive means to evaluate objectively the different proposals has been a ‘Layout Tool’, which is software to compute the detailed parameters of various layouts and make estimates of power, cost, channel count, material budget, cable types and numbers, optical fibre, transmitter and receiver distribution, and many other

important features. The merit of this tool is the objective assessment it provides and, as it was developed by members of the CERN team led by Duccio Abbaneo, who played large roles in the original tracker design and construction, mechanical and electrical integration, it includes realistic assessment of material and components. Examples are reasonably detailed models of module designs, and constraints in assembling a real detector, such as feasible routes and material implications in choice of cables and their parts, and practical physical layout, where space is very limited. It has been supplemented to include up to date information on power-conversion and new cooling systems.

The tool evidently cannot displace the need for full simulations to estimate physics performance. However, full detector simulations have proven, unsurprisingly, to be very time-consuming and subject to concerns about realistic modelling of the proposed layouts. In addition, they do not really allow to compare alternative layouts, for example for cost, power or channel count. In view of this, the layout tool has an important role and we have discovered that it is possible to extract important features of performance from it.

Work done by G. Hall demonstrated that it is possible to reproduce quite accurately major details of performance, closely matching full simulations, by means of simple numerical calculations. Building on previous semi-analytical methods traceable to modelling of large tracking detectors from the past and used in the early days of CMS, it was found possible to extend these calculations to correctly include multiple scattering. This allows accurate estimates of momentum and impact parameter resolution over a wide range of momenta and angular coverage. The methods have now been incorporated into the software tool to be applied to future layouts.

### **3.3 Deliverables**

The progress against the project milestones is summarised below. Depending on progress with the pixel module readout project, it will be appropriate to introduce a new milestone. We note here that milestones indicated as ‘Year 3’ are at present due for delivery before the end of WP1 funding at 1 January 2012.

#### *Year 1:*

The contribution to the stacked-tracking simulation programme has been successfully completed, and further work assigned a low priority.

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 was with the publication of the TP in late 2010. Detailed physics performance measurements are ongoing, with a next report to LHCC due in March 2011.

#### *Year 2:*

A CMS-wide release of the full simulation code for the upgraded tracker, and the emulation package for the upgraded L1 calorimeter trigger has been made, with significant UK contribution in each case.

A new milestone has been added (1.3.5) for a feasibility study for a UK contribution to the Phase I pixel upgrade, which should allow an estimate of resources required.

#### *Year 3:*

The full-simulation performance study of the upgraded detector against physics benchmarks is ongoing, and the results will be documented in a series of increasingly detailed reports leading up an upgrade Physics TDR.

The proof-of-principle upgrade trigger algorithms are being defined in conjunction with WP3, and will be demonstrated during 2011.

### **3.4 Staff on project**

Reported in tables. We note that all funded staff effort will currently expire at or before the end of 2011.

The situation with RAL posts remaining unfilled (due to recruitment freezes), detailed in the last report, has now been partially rectified. RAL staff within WP1 have been engaged in several areas of

the project, including simulation, online software / firmware developments and pixel readout. Roles undertaken by RAL staff in the reporting period include:

- K. Harder has contributed to Phase-I pixel simulation code, providing tools for the parameterisation of b-tagging performance for physics studies. He has also driven the assembly of the module test stand in the PPD pixel lab.
- D. Sankey has undertaken the setup of a online software test environment at RAL, and has carried out initial investigations on the use of non-IP Ethernet protocols as the transport layer for online control of hardware.
- S. Worm has carried out studies of the material distribution in the upgraded pixel detector, and tuned the simulation.

### **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.

## **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.

### **4.2 Progress to date**

#### ***Front end chip***

The CBC readout chip (CMS Binary Chip) development is for short silicon microstrips (2-5 cm) to be used in the outer tracker region ( $r > \sim 50$  cm). It was originally thought that these short strip outer layers would probably not contribute to the Level 1 trigger, but there is growing interest in incorporating 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. The current chip does not include this functionality, but conceptual designs of the required logic now exist, including mechanisms for transmitting the tracking information off-detector, so a subsequent version of the CBC could be produced.

The 130nm CMOS CBC chip was designed in a collaboration between RAL TD and Imperial College. At the time of the last OSC report (Sept 2010) the CBC had recently been submitted (in July) and a description of the functionality was included there. The main functionalities of the 128 channel prototype are:

- a fast front end amplifier, with 20ns peaking time,
- ability to match both sensor polarities
- ability to tolerate leakage currents up to 1 uA
- comparator with programmable threshold,
- 256 deep pipeline,
- 32 deep buffer for triggered events,
- unsparisified binary readout for chip and system simplicity
- output MUX and driver, using a low power signalling standard (SLVS),
- fast (SLVS) and slow (I2C) control interfaces.
- on-chip DC-DC switched capacitor power block
- low dropout regulator to supply analogue stages

A setup has been prepared for detailed evaluation of the CBC prototypes in the lab. This completes milestone M2.2.2.

At the time of the previous OSC report we were expecting chips back in October 2010. We have incurred an unexpected delay of approximately three months, mainly due to the IBM foundry being exceptionally busy in the run-up to Christmas. Eventually wafers left the foundry just before Christmas, but the holiday season meant that wafer cutting (organized by CERN) could not begin until January. This is now complete and we expect chips in hand imminently. Nevertheless there will be some implications of this delay for milestone M2.2.5 (review options for further CBC development) due March 2011. The next milestone after that for the CBC development will be M2.3.2 (documented preliminary chip results) due March 2012. If the chip works well it may be possible to meet that milestone at its current date, but this will be kept under review and an update can be provided at the time of the next OSC report.

The first priority for the coming year will be to evaluate the CBC and decide what aspects require more attention or improvement. Assuming it works adequately this will include test beam activity including evaluation of the SEU sensitivity of the internal circuitry, which is important to understand to inform further development of the chip and therefore meet future milestones.

With regard to milestone M2.2.5 (due March 2011), we have discussed several options for further development with CMS collaborators 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. While we are close to taking decisions on the above items, one outstanding issue will be evaluation of the DC-DC conversion and power regulation, which requires input from the CBC prototype testing which has been delayed.

### ***SFED developments***

An engineer from the Technology system design group began full time on the project in September 2010. He will concentrate initially on implementing basic SFED FPGA firmware modules and testing them in simulation and on commercial development boards. He has started by working with the CERN Giga-Bit Transceiver (GBT) kit which provides an FPGA implementation of the GBT ASIC logic. The aim of this work was to gain understanding and practical experience of the GBT system in general and the working of the GBT protocols in particular.

The GBT kit is targeted to work on a particular Xilinx Virtex5 series FPGA commercial development board. The design was tested at RAL on this platform and then ported to the mainstream Virtex5 development board. The port required changing the I/O from electrical links to optical transceivers. This exercise enabled a good understanding of the working of the GBT design to be obtained. The design was then ported to a commercial board using the latest generation of FPGA Virtex6. As the Gigabit transceivers are different in Virtex6 the work required modification of the VHDL source code. We have verified that the ported designs work as the original. A report has been written describing the work carried out which includes instructions for other users wishing to run the GBT kit on the Virtex5 or Virtex6 boards. The design files have been submitted to the CERN GBT group for inclusion in their central repository.

Work has now begun on developing FPGA VHDL designs for fast data buffering to external memory on a Virtex6 development board with a view to interfacing with multiple GBT data streams. CERN intends to use the Virtex6 on their new GBT test platform GLIB board. With a view to using the WP3 Mini-T board in future SFED tests we have provided advice for adding fast external memory to the next iteration of the Mini-T board design. We are also familiarising ourselves with the CMS IPbus standard which provides a common interface to FPGA firmware from PC networks and which is used for control of the Mini-T board.

### ***L1 trigger system***

CMS remains committed to the idea of using Tracker data in the level 1 trigger at high LHC luminosity. Ideas for this are still under discussion and simulation. One proposal which matches naturally to the CBC activity, is to construct “two-in-one” strip modules to be deployed in the outer tracker. Strips of ~5 cm length installed in modules assembled as two layers could provide a means of transverse momentum selection, similar to long-pixel modules, but in a technically safer, lower cost and easier to implement manner. Modules of that kind have been prototyped using APV25s in Pisa

and appear to be viable. We are discussing the options for a variant of the CBC, followed by suitable modules, with our collaborators. More TD microelectronic staff are now available to develop these ideas.

### ***Preparations for test-beam activities***

In view of evaluation of modules in test beams, we prepared equipment for this purpose, with emphasis on much higher speed readout than conventional in beam tests but with no sacrifice of spatial resolution.

A charged particle telescope was constructed and used in September 2010 for data taking at high rates in the CERN North area H8 400 GeV/c proton beam line. It utilises ten planes of silicon microstrip sensors, arranged as five pairs each measuring two orthogonal coordinates, with an active area of  $3.8 \times 3.8 \text{ cm}^2$ . The objective was to provide excellent angular and spatial resolution for measuring the trajectories of incident and outgoing particles. The apparatus has a long baseline, of approximately 10m in each arm, and achieves an angular resolution of  $5.2 \text{ } \mu\text{rad}$  with performance limited by multiple scattering in the sensor layers. The sensors are instrumented by a system based on the CMS Tracker electronic readout chain, including analogue signal readout for optimal spatial resolution.

The system profits from modified CMS software and hardware to provide almost deadtime-free data acquisition capable of sustained trigger rates of up to 9 kHz. This can be further improved in future and was important for applying the system to crystal channeling studies in collaboration with the UA9 project. It also provided a test bed to demonstrate high speed data acquisition, which is important for evaluating future components and systems under realistic operating conditions.

One notable effect illustrating the precision obtained was the discovery of a hitherto unnoticed movement of the beam during a spill. The incident angle of the beam, relative to a nominal initial value, varies with a period equal to the beam spill (flat-top) period. The effect is observed consistently in every sensor, indicating that the origin is related to the beam optics upstream of the telescope. The telescope is able to resolve that the beam direction shifts by approximately  $6.5 \text{ } \mu\text{rad}$  horizontally and  $2.5 \text{ } \mu\text{rad}$  vertically over the course of a spill. This would be undetectable in a typical telescope, much shorter in length, but is important for the precision of the channeling studies.

The system is described in a paper shortly to be submitted for journal publication. It has been agreed with the CMS Tracker upgrade team to develop the system further in collaboration with others in CMS to instrument beam tests for a wide range of tracker upgrade activities including sensor evaluations.

We are now beginning to develop ideas for CBC modules and the required DAQ for readout, including suitable high speed digital links which could be based on GBT and VOL prototypes.

### **4.3 Deliverables**

The milestone status is

*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 system discussions continue regularly. This is a long term activity.

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

Preliminary investigations of powering schemes have been reported but a final written report is not needed, in view of the decision by CMS to adopt DC-DC conversion.

*Year 2:*

The deliverable of a full chip prototype submitted for fabrication was brought forward considerably and already met. However, as mentioned above, delivery has been delayed. Meanwhile test setups have been prepared and are ready for use.

The milestone (2.2.4) referring to the development of FPGA logic to be used for the CBC readout has been deferred to match the expected CBC evolution. Meanwhile, prototyping of link interfaces continues which is one of the key areas to be explored.

#### **4.4 Staff on project**

Listed in the tables but no significant changes since the last report. At Imperial College, Osman Zorba has resigned for personal reasons, which affects WP2 and WP3. His post has been advertised, with a replacement appointment expected in the near future.

#### **4.5 Expenditure**

Spending was modest in view of the CBC late delivery. Test equipment has been the main contribution.

#### **4.6 CMS activities elsewhere**

The CBC is a unique activity in CMS, which would be hard to replace and is also very advanced compared to any other design effort. Lyon have recently dedicated some design effort to front-end ASIC development and have showed some ideas for implementing cluster-finding logic and a basic amplifier, which they have submitted as a prototype on the same MPW run as the CBC. We continue to work closely with CERN collaborators on the future system design, especially links, and have discussed options for trigger chips. We expect to maintain this collaboration.

### **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.

#### **5.2 Progress to date**

In the last status report it was recognised that a robust, high performance and scalable Ethernet Hardware Access Layer (HAL) was essential to the project. A collaboration with WP1 staff led to the development of the IPbusClient software and refinements to the original IPbus firmware supplied by the Minnesota group. This looks likely to be adopted throughout CMS as it already has significant support (CERN & US).

In parallel to the software development, the core firmware infrastructure was completed, the electron and tau algorithm implemented and results presented at the TWEPP conference in Aachen in September 2010. System tests with multiple cards have started, but will be paused to allow the migration of existing software to the new IPbusClient. The UK also coordinated the purchase of a common MicroTCA crate for CMS that will be adapted to CMS requirements (e.g. vertical airflow).

Lastly, a proposal for an alternative trigger architecture (Time-Multiplexed Trigger - TMT) has been distributed. It potentially offers many benefits over a conventional trigger architecture and is under discussion in CMS.

The prototype main processing card (Mini-T5-R0) was delivered in Feb 2010 and a revised version with alternative optics in August 2010 (Mini-T5-R1). It has been successfully validated in a  $\mu$ TCA crate with 32 optical links operating at the rated specification of 5Gb/s.

The system design has benefitted enormously from implementing trigger algorithms, and consequent improved knowledge of the FPGA resources required. This has led to the addition of extra RAM in the latest iteration of the card (Mini-T5-R2) that will go out to manufacture in April 2011.

WP3 progress slowed in the Autumn of 2010 because it was necessary to adapt the GCT to CMS operations in a new, and unexpected, manner during heavy ion collisions and to exploit the flexibility of the GCT to implement an improvement to the transverse energy sum algorithm. The changes required are close to completion, but will still have an impact on WP3 and to a lesser extent WP1 for a

couple of months. Integration tests of the Optical Global Trigger Interface card were completed in January. It potentially offers a way into the Global Trigger for an upgraded Calorimeter Trigger.

We have been fortunate to have received a generous donation by Xilinx that will allow us to build 12 Mini-T5-R2 cards. This will allow system and algorithm development to continue over the next year while a final hardware design is designed, manufactured and tested.




Figure 2: The Mini-T5 card, undergoing system tests in a  $\mu$ TCA crate at CERN.

### 5.3 Deliverables

The milestone table has been updated. The main relevant points are:

#### *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.

Core firmware infrastructure implemented (Ethernet, link control, pattern RAMs and DAQ) and tested with electron and tau trigger time multiplexed trigger algorithm.

Control software infrastructure developed.

Alternative trigger architecture note distributed; for Time-Multiplexed Trigger.

### 5.4 Staff on project

Reported in accompanying tables. The WP manager C. Foudas left Imperial in August 2010. The PI has taken responsibility for overall supervision of WP3 activities. Alex Tapper will gradually play a larger role in the project and assume some responsibilities for the longer term.

### 5.5 Expenditure

The major items of expenditure have been optical connectors and further Mini-T card orders.

## 5.6 CMS activities elsewhere

The new approach of time multiplexing the data has significant benefits. It concentrates the entire calorimeter trigger data for a given bunch crossing into a single FPGA and dispenses with boundary conditions between processing nodes and thus the custom backplane.

The Mini-T card remains a technically more advanced card than anything else in CMS and is providing a significant platform for firmware and software developments. The Wisconsin group have favoured a more traditional approach to the trigger upgrade, using the existing architecture but with FPGA hardware yet to be fully defined. This is the main subject of our ongoing discussions.

## 6 Risk register

The current version of the project risk register is v3.2, revised from the version presented at the last meeting. Each line of the risk register has been reviewed recently and comments have been added where appropriate.

Overall progress continues to be good, although we have noted the delay in delivery of the CBC chip, which clearly has knock-on effects, even though progress in the early stages of the project was much faster than originally expected. Until the chip has been tested some uncertainty is added to the remaining schedule.

The major risks remain unchanged, namely possible shortfalls in funding and loss of key staff.

The overall timescale of the CERN accelerator project (no longer usually referred to as Super-LHC but High Luminosity LHC) has changed, as noted, with the new schedule not yet announced. This seems to be a likely occurrence in future in view of the novelty of the LHC and potential challenges to be overcome in operating such a ground-breaking accelerator. An explicit risk (no.38) has been added to the register. The implications for this R&D project may not be great, but there are greater ones in the long term given the inevitable uncertainties in defining the actual construction dates. However, some of the extra costs will simply form part of future operating costs for which future STFC planning must take account.

## 7 Finances

The financial report is summarised in the attached tables. There are a number of noteworthy issues.

The new STFC finance reporting system provided via the Research Councils Shared Services Centre has presented well known difficulties in accessing budgetary information and ensuring it is accurate and up to date. This affects mainly the reporting of RAL PPD and Technology spending. Travel spending is so far rather limited; this is mainly due to the fact that many trips can and have been combined with other CMS business. While the SSC remains in its present state there has been no value in trying to attribute travel more accurately to different budgets. Hopefully this will be possible in the next financial year, when a US CMS upgrade workshop is again foreseen and conference presentations are very likely.

Rolling grant spend is likely to be lower in total than originally foreseen in the proposal and approval, especially affecting Bristol and Brunel. This is a straightforward consequence of the last grant settlement and subsequent prioritisation and cuts exercise. We also recall the long delay between approval and allocation of funds.

Equipment spending, which covers mainly ASIC manufacture costs and trigger board development, lags behind early forecasts. Now WP3 demonstrators have been produced and debugged, with small faults corrected, it is expected the production of further trigger boards will increase the rate of expenditure in the coming year. Once the CBC has been tested, the next investments in ASIC manufacture will generate much greater expenditure. In the first year of the project, we were also fortunate that the CBC production was possible on a Multi-Project Wafer (MPW) run, whose cost was dominated by a large prototype chip designed by RAL for XFEL. In

future runs, we expect to pay a much larger share of the cost of each production run, possibly even the entire cost, depending on other projects and the feasibility of sharing the wafer surface.

WP2 and WP3 spending remains well within the envelope foreseen for requisitions and staff spending. However, engineering staff spending at RAL has continued to be lower than anticipated because of availability of engineers for design work. This has been changing over the last six months and spending will accelerate in 2011 with progress of the CBC.

For WP1, staff spending within university groups is approximately as expected. The Brunel allocation is close to exhaustion. In Bristol, some details of the original project funding allocation categories are being confirmed with STFC which explains the apparent low ratio of estates and indirect costs to project effort in the current financial year. In the last financial year, RAL PPD was unable to spend at the predicted rate due to issues related to implementation of cuts throughout STFC. More staff are now allocated as explained in the WP1 report, and the earlier underspend is expected to be absorbed over the remaining duration of the project.

If a Phase I pixel upgrade project looks viable, then it should be the subject of a new Work Package which would require a bid for additional funding. At present the financing is at the expense of other, approved, WP1 activities.

The £230k Working Allowance remains untouched.