22 October 2013

University of Bristol
Brunel University
Imperial College London
Rutherford Appleton Laboratory

Upgrades of the Tracker and Trigger of the CMS experiment at the CERN LHC

1 Executive Summary

The new project started in April 2013; it continues work on the tracker and trigger begun in the previous R&D project. For the Phase I upgrades of the L1 calorimeter trigger, this now becomes a construction project.

In the last six months there has been further excellent progress with the UK activities:

- Further tests on the CBC2 chip have been carried out demonstrating the chip to be working very well.
- The first prototype of a small 2S-module has been constructed and laboratory tests have demonstrated track-stubs using a beta source and cosmic data. A beam test is scheduled for late November.
- The MP7 has been successfully used in two trigger integration tests in CERN.
- A demonstration of a slice of the Time Multiplexed Trigger was made in late September, meeting all the objectives defined for the test.

The project is on the envisaged schedule but the long term plan is under review to ensure it is best matched to overall CMS objectives.

2. Project history and recent developments

As reported at the final OSC meeting of the previous R&D project in March 2013, a new proposal for R&D for Tracker upgrades for Phase II of the HL-LHC and for construction of the Phase I calorimeter upgrade had been reviewed by the PPRP and recommendations for approval made. The funding allocations were confirmed by STFC in March in time for the expected project start on 1 April 2013. There are two technical work packages: WP2 for the outer tracker readout R&D; WP3 for the calorimeter trigger.

The LHC upgrade is proposed to take place in two main stages, with an increase in luminosity to ~2.5 10^{34} $cm^{-2}s^{-1}$ from 2015 over a couple of years then, after a two to three year shutdown around 2022 in the most recent CERN ten-year plan, to ~5 x 10^{34} $cm^{-2}s^{-1}$ levelled luminosity, denoted as Phase II. A total of 3000 fb^{-1} in integrated luminosity over a decade is the goal.

CMS and the LHC are currently in a Long Shutdown (LS1) with the expectation that the LHC will restart in early 2015 with first collisions at ~13 TeV expected around March. The exact operating conditions are still under discussion; the machine will not aim initially for the full 14 TeV which should eventually be achieved and the collision frequency should be 40 MHz for the first time. Both targets, especially the bunch crossing rate, are likely to be kept under review during early operation and it is not excluded that the machine will revert to 50 ns bunch spacing if difficulties emerge. However, both ATLAS and CMS are continuing to press strongly for 25 ns operation if at all possible in view of the implications for pileup as the luminosity rises and hopefully exceeds the design value in the coming few years.

Recent presentations on progress with the machine have been very positive with the work being very close to the expected schedule and no significant problems identified. An update is due in the next RRB meeting on 28 October, before the OSC meeting.

CMS shutdown activities are also proceeding broadly according to schedule. One significant occurrence was a break in the new beam pipe – twice – at a weld between straight and conical beryllium sections. This is attributed to handling errors at the supplier's premises due to non-optimal support jigs and will delay delivery of the pipe to CERN by about two months, but is not expected to have a significant overall schedule impact. Among other major progress, the tracker has been effectively sealed for future operation at the minimum possible temperature which is important to restrict future radiation damage; until now the operating temperature was limited by humidity in the tracker volume. Recommissioning is beginning and the first cosmic ray data taking runs will take place in the next few weeks.

2.1 LHC upgrade schedule and planning

There have been no real changes in the plans presented previously. Of course, until there is experience of operation at the new energy, better predictions of progress cannot easily be made. The machine objective is to reach the design luminosity of 10^{34} $cm^{-2}s^{-1}$, then increase the maximum instantaneous luminosity to about twice that value over the subsequent two years.

The current LHC planning foresees three long shutdowns: LS1, LS2, and LS3. Following LS1, the centre of mass energy will increase to ~14 TeV, and the bunch spacing should be reduced to 25 ns. In LS2 in 2018 the injector chain will be improved to deliver high intensity and low emittance bunches. In LS3 (2022-) the LHC will be upgraded with new low-β triplets and crab-cavities to optimize the bunch overlap at the interaction region.

The original goal for the LHC, to operate at an instantaneous luminosity of 10^{34} $cm^{-2}s^{-1}$ with 25 ns bunch spacing, is likely to be achieved soon after LS1. Under these conditions, CMS will experience an average pileup of about 25 inelastic interactions per bunch crossing. This is the operating scenario for which CMS was designed. Based on the excellent LHC performance to date and the upgrade plans for the accelerators it is now anticipated that the peak luminosity could reach 1.6 10^{34} $cm^{-2}s^{-1}$ before LS2, and above 2 10^{34} $cm^{-2}s^{-1}$ after LS2, providing an integrated luminosity of over 300 fb^{-1} by LS3.

Following LS3, the high luminosity programme with the HL-LHC begins. The operating scenario is to level the instantaneous luminosity at 5 10^{34} $cm^{-2}s^{-1}$ from a potential peak value of 10^{35} $cm^{-2}s^{-1}$ at the beginning of fills, and provide 250 fb^{-1} per year for a further 10 years of operation. Under these conditions event pileup will be a major challenge for the experiments, and performance degradation due to integrated radiation dose must be addressed.

The exact timing and duration of the next long shutdown (LS2) are still under discussion; there is no complete consensus from the experiments on this but it will depend both on machine progress and the ability of the experiments, including LHCb and ALICE, to meet their objectives for changes to their detector systems, some of which are major and not yet fully funded. CMS is aiming for an extended technical stop at the end of 2016 to install the new pixel detector. This looks feasible but also depends on completion of the new detector and its successful commissioning before installation inside CMS as well as agreement with the other experiments and the LHC on the duration and timing of the necessary shutdown.

2.2 CMS planning

In 2007 CMS submitted an "Expression of Interest in the SLHC", where SHLC referred to an LHC delivering a peak luminosity of 10^{35} $cm^{-2}s^{-1}$. The plans for upgrading the accelerator complex have significantly advanced since and the luminosity is now expected to steadily increase throughout the next decade. To maintain current performance CMS has developed an upgrade programme in two phases.

Phase I comprises completion of the original scope for 10^{34} $cm^{-2}s^{-1}$ operation, and a targeted set of upgrades to allow operation at the luminosities and radiation levels anticipated up to LS3. The full programme for Phase I was described in a Technical Proposal in 2011 that also included a discussion of work towards Phase II for HL-LHC in an appendix. Following this, Technical Design Reports were prepared for three major Phase I upgrades: replacement of the pixel detector, upgrades to the Hadron Calorimeter electronics, and upgrade of the Level-1 trigger system.

In the last six months the remaining CMS Technical Design Report for the Phase I trigger upgrade was reviewed and approved by the CERN LHCC and this forms the basis for WP3 planning.

A CMS Phase II Technical Proposal is expected in 2014 and a Tracker Technical Design Report in 2016.

What has changed over the last year or so is the view of improvements needed to other parts of the detector than the tracker. Studies and projections of performance degradation due to radiation damage are ongoing, but it is clear that the barrel calorimeters will perform well at HL-LHC. The Forward HCAL (HF) will degrade at the highest η and further studies are needed to determine if an upgrade will be required. If so, it may be possible to stage this upgrade later than the main Phase II upgrades in LS3. It is now clear, however, that the endcap calorimeters EE and HE must be upgraded in LS3.

Longevity of the Electromagnetic Calorimeter is driven by radiation damage to constituent crystals and photodetectors. The behaviour of both has been studied, analysing dedicated beam and laboratory tests and CMS data. Transparency loss in the crystals is the dominant aging factor in the ECAL endcap, where the collected signal drops as a function of η and integrated luminosity. By 3000 fb^{-1}, almost half (in terms of η coverage) of the EE will have light collection reduced by more than a factor of 10, leading to degradation in the measured energy resolution and progressive loss of the trigger coverage. Maintaining the signal collection at a level higher than 10% (compared to undamaged crystals), over the full detector coverage, will only be possible up to 300-500 fb^{-1}. The feasibility of partial replacement, in which only crystals with the highest radiation damage are replaced, was investigated in detail. The mechanical mounting of the present detector, the layout and connectivity structure of the electronics, and the exposure of personnel to significant radiation levels make this scenario untenable. It therefore seems necessary to replace the full EE during LS3.

In addition physics arguments have evolved following the Higgs discovery. The distribution of jets in the Vector Boson Fusion (VBF) process peaks at $|\eta| \sim 3$. To improve VBF jet-tagging, extending the EE and HE coverage from $|\eta| < 3$ to $|\eta| < 4$ is under consideration. This would avoid the transition to HF at $|\eta| = 3$ and increase coverage for photon and electron measurements.

It is anticipated that extension of the endcap coverage would make a significant improvement to measurement of VBF processes, an important element of the HL-LHC physics programme. It may however affect radiation backgrounds in the tracker and muon chambers. Studies are ongoing to demonstrate the performance improvement that can be attained, and to determine whether the radiation levels are acceptable for the calorimeter and the backgrounds are tolerable for neighbouring detectors.

This would be complemented by extension of tracking to higher η. In the last few months a baseline tracker geometry was adopted, based on extensive modelling and studies of the effects on material, power, cost and performance. It is illustrated in figure 2.1. There are two types of pT-module for triggering using the silicon tracker: 2S-modules based on short microstrip sensors read out by the CBC ASIC, and PS-modules which are conceptually similar but combine long pixels with short microstrips. The pixel layout is based upon the Phase I detector with 4 barrel layers, but with a system of 10 disks per end rather than the 3 disks of the Phase I design. The outer three disks on each end consist of the outer ring only, in order to avoid interference with the conical section of the beam pipe.

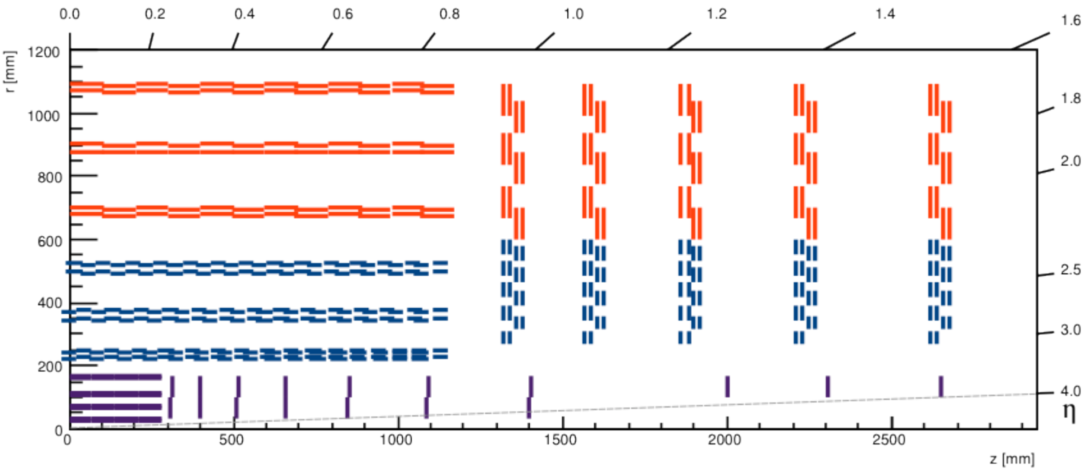

Figure 2.1: A quadrant of the Tracker layout. Outer Tracker: blue lines represent PS modules, red lines 2S modules. The Inner Pixel detector, with forward extension, is shown in purple.

A document summarising CMS objectives has been submitted to the next RRB to ensure the Funding Agencies, as well as CERN, are aware of these ideas. The overall cost is estimated at 270 MCHF, plus R&D. Meanwhile ATLAS had submitted a Letter of Intent detailing their plans earlier in the year so a reasonably clear picture of the overall experiment ambitions is available; when this is combined by CERN with requirements from other experiments, the LHC machine and the overall

CERN programme presumably there is a basis for forward planning. An ECFA-sponsored meeting was held in early October in Aix-les-Bains which exposed all these subjects to a wide community and it was expected that the next RRB in late October would provide an opportunity for deeper discussion of the future with representatives of the funding agencies to define the way forward.

2.3 UK adaptation to CMS planning

Excellent progress has continued with UK R&D activities, with more details given later in the WP reports. At the same time, uncertainties in the overall planning persist and this may require adjustment of UK planning.

The main uncertainty concerns the Phase II upgrades. An ECFA workshop was recently held where the experiments outlined their proposed upgrades and there were detailed presentations from the entire LHC community on physics objectives, machine progress and plans and issues concerning the upgrade implementation, such as electronics and detector R&D, potential radiation exposure and infrastructure and experimental areas. However questions of funding and constraints which might arise for economic reasons were not discussed and the overall strategy for obtaining sufficient funding to undertake all the required work, including the machine changes, remains unclear. This may be further addressed in the forthcoming RRB meeting. It is important to do so since it is difficult to prioritise activities without, e.g., knowing the complete scope of the CMS upgrade in view of the overall economic situation. It will certainly be challenging to complete all the required R&D in time and some of the upgrade programme, for example in the forward region, is interlinked. The forward tracker is most useful if new calorimeter systems in the same region are provided.

The overall design of the Tracker is, however, reasonably clear and the forward pixel detectors are not expected to be different than other endcap detectors so do not involve major changes to the R&D already planned. The basic design of the tracker, outlined above, was agreed prior to a CMS workshop in DESY in June and is based on a Barrel-Endcap geometry which is clearly significantly better than alternatives which have been under discussion for some years. The design makes use of “stacked” tracker modules throughout, which originate in ideas from the UK, of which the 2S-modules now in prototype form in WP2 would form a major part.

However, to provide an effective trigger it is certain that 2S-modules alone would be insufficient because of their limited spatial resolution in the beam direction. PS-modules have always been foreseen as essential to the overall tracker design and here progress has not been so fast. In contrast, two successful iterations of a full-size CBC have been carried out, despite significant differences between them. The CBC development is not complete but the features which require further development are modest in comparison with the overall design which has been done so far. Therefore we have begun to question if the work on further CBC iterations foreseen to complete the development should be given a lower priority compared to contributions to design of the PS-module, to which it was always foreseen that we would contribute, provided progress with the CBC was sufficient. We note that the PT-module concept originated in work done at Imperial College and most of the early studies of the idea were carried by PhD students there, two of whom (M Pesaresi, A. Rose) remain on the upgrade project.

An important point to keep in mind is that R&D expenditure is not formally recognised as a cost contribution to the delivery of the experiment, although pre-prototype manufacture is. However, at what point R&D becomes pre-prototype work is not well defined; certainly everyone would be happy to have a complete and final chip in hand at the time of the TDR, especially if it had been fully paid for. The UK has contributed substantially to the R&D to date and it is not clear that every agency has been able or willing to do likewise.

In addition the specifications for the future tracker have changed relatively recently in an important way. Simulation studies are not far enough advanced to be certain that the L1 rate can be reduced to 100 kHz within 6.4 μs latency and a possible objective of L1 readout with up to 1MHz and a latency of 10-20 μs is under study. New ASIC developments are requested to design for up to 1 MHz and 10 μs. It appears this need not have major implications for the CBC, especially if the rate is constrained to about 500 kHz, because it is probable that the data compression required to transmit trigger data off-detector can be implemented in a concentrator ASIC which has not yet been designed and which is currently being modelled in an FPGA before a digital chip is produced. However, the high L1 trigger rate and latency do have cost implications for the rest of CMS, especially for the ECAL and end-cap muon electronics.

This has few implications for the actual 2S-module development since we have plenty of CBC2 chips available and they provide all of the functionality for module and stub-finding evaluation studies, test beam work and test system development. Other components, such as the concentrator ASIC, low power GBT, realistic DC-DC converter and optical transceiver are not foreseen to be available in more realistic form for some time anyway.

Overall, it seems unwise to commit to further iterations of the CBC at present until the situation becomes clearer, and we can also explore if some of the submission costs might be shared with others. We propose to review the situation over the next six months and adjust milestones accordingly.

There are also motives to review the trigger schedule, but for different reasons.

The TDR plan to implement the new calorimeter trigger foresees parallel operation of the new and existing systems during 2015, with the new trigger foreseen to take over from early 2016. In order to provide calorimeter data to the new trigger system, existing copper serial links, which transfer energy information from ECAL and HCAL to the calorimeter trigger, will be replaced with fibre-optic links. This modification to the CMS systems will take considerable time, and must take place during this year and next.

Meanwhile, as reported later, a very successful demonstration of the Time Multiplexed Trigger was carried out in September, fully meeting all the specified objectives. To make this possible the first layer of the trigger (Pre-processors transmitting data to the Main processors) was implemented using the UK-designed MP7s. In the final system, the pre-processor boards are a US deliverable but these boards (CTP7) are only now at the layout stage. From our experience with the MP7, this is not a small job and there are serious questions about whether this can be achieved to meet the trigger upgrade schedule, especially in view of resource allocations in the US, where funding constraints have been a concern for some time.

In addition, there has been pressure within CMS for improvements to the trigger as early as 2015 in view of the discovery potential at 13 TeV and the chance that luminosity will increase quickly during the next running period. A proposal has been made, which is not described in the TDR, to incrementally upgrade the Level-1 system during the LS1 shutdown. The present GCT system will be replaced by a small number of MP7s. The current Regional Calorimeter Trigger (RCT) will be augmented by US collaborators with a new final output board (ORSC), to allow optical data transmission to the MP7 modules at sufficiently high rates. These changes (“Stage 1”) are intended to provide a significant improvement in the performance of the CMS calorimeter trigger whilst requiring relatively modest resources, and minimal disruption to the existing CMS trigger system.

A test was carried out at CERN in late July which was supposed to fully demonstrate the integration of the ORSC with the MP7. Basic performance parameters were measured, especially eye diagrams and Bit Error Rate measurements, but the ORSC was not available for sufficient time to complete all the tests envisaged and it was returned to the US after the test. An internal review in August expressed concern about progress. The UK developments are fully on schedule but how best to adapt to the overall trigger situation is still unclear.

A CMS Engineering Design Review (EDR) in November will consider the overall state of the L1 trigger developments and recommend actions and priorities for the future. The CMS management is paying close attention to this issue.

3. Work Package 1: Management

The project is divided into two main work packages concerned with implementation. Each WP has two managers, listed in the table below. One main reason for this is the need for hardware design and layout to be closely supervised by an Imperial College team member, but it also has the merit of distributing responsibilities between institutes and providing continuity and sharing of information and duties. There is a significant level of interaction regularly required with CMS, and it is essential that the WP managers do not become overloaded with reporting and budget management duties, and are able to actively supervise others working on the project in their direct vicinity.

WP	Manager	Institute	Role
1	G Hall, PI	Imperial	Overall management, budgetary responsibility and supervising procurements, interface to CMS, as UK CMS PI and CMS Management Board and Tracker Management Board member.
2	M Raymond	Imperial	Overall responsible for CBC specifications, interface to module design team, chip testing and module evaluation and CMS planning
	M Prydderch	RAL TD	Manager of ASIC design team in RAL
3	A Tapper	Imperial	CMS Upgrade Trigger Project Manager, currently based in CERN with supervisory responsibilities for G. Iles, Imperial College engineer, also based in CERN.
	D Newbold	Bristol	UK firmware and software coordinator. Trigger Institution Board chair.

The costs of WP1 are the fraction of PI time funded through the project and the support staff assisting. However, in the effort table, the PI time is all reported in WP2 although his activities have covered all three WPs.

4. Work Package 2: Outer Tracker Readout

4.1 Objectives

The objectives of WP2 as stated in the proposal are:

- To complete development of a readout and triggering chip suitable for the 2S-PT module, bringing the chip to a final state ready for mass production.
- To develop the hardware and software required for the large-scale production testing procedures, and to deliver tested wafers to the CMS experiment.
- To play a major role in construction, definition and evaluation of prototype modules.
- To contribute to development of ancillary chips required for the 2S-PT module, and to participate in the PS-PT module development.
- To contribute to the future large-scale module production programme, and to participate in integration and commissioning activities.

The 2S-PT module concept meets the HL-LHC challenge of providing tracking information to the level 1 trigger decision in CMS by providing coordinates of high pT stubs formed by correlating signals occurring in closely spaced sensor layers. The final deliverable of the previous programme, the CBC2, is a 130 nm prototype CMOS chip containing all the functionality required to allow prototype 2S-PT modules to be constructed.

4.2 Progress to date

CBC2 was submitted for fabrication in July 2012. The main features are:

- 254 readout channels to allow correlation between two sets of 127 sensor channels.
- Cluster width discrimination logic: wide clusters cannot be consistent with high pT tracks.
- Correlation logic: A prompt trigger pulse is produced if a cluster in one layer correlates with one found within a window in the other layer.
- Bump-bond layout
- Test pulse system. Signals can be injected into all front end channels.

Eight CBC2 wafers were delivered in January 2013. Early test results were reported in the previous OSC report (February 2013) where basic functionality had already been demonstrated. Since then we have continued with the test programme, and are now at a point where, although there is still much to do, we can confidently state that the CBC2 chip is working well.

The CBC2 has been designed for C4 bump-bonding, early test results being obtained using the wire-bondable pads which were included to allow wafer-probing. To properly characterise performance it is necessary to mount chips using the technique for which they were designed, so suitable carriers are needed to which the chips can be bump-bonded. Since we would like to have a clear distinction between failures that might arise from faulty connections and those due to naturally occurring chip defects (yield) it is necessary to wafer-probe to provide known good chips.

Figure 4.1 shows one of the C4 wafers and the insets show some magnified images of the CBC2 chips and ball bonds on the wafer and the wafer probe needles in contact with the pads. The yield of working chips is high, and for the first wafer the only chips rejected were damaged accidentally during the probing procedure. A subsequent wafer has now been screened with a more comprehensive set of tests, and the yield remains high at > 95%.

Figure 4.2 shows one of the first mini-pT-module constructed using CBC2 chips bump-bonded on semi-rigid hybrids manufactured by Endicott. CERN has primary responsibility for the design of bump-bondable hybrids and the example shown in the figure uses 2 CBC2s to read out 254 channels from each sensor layer.

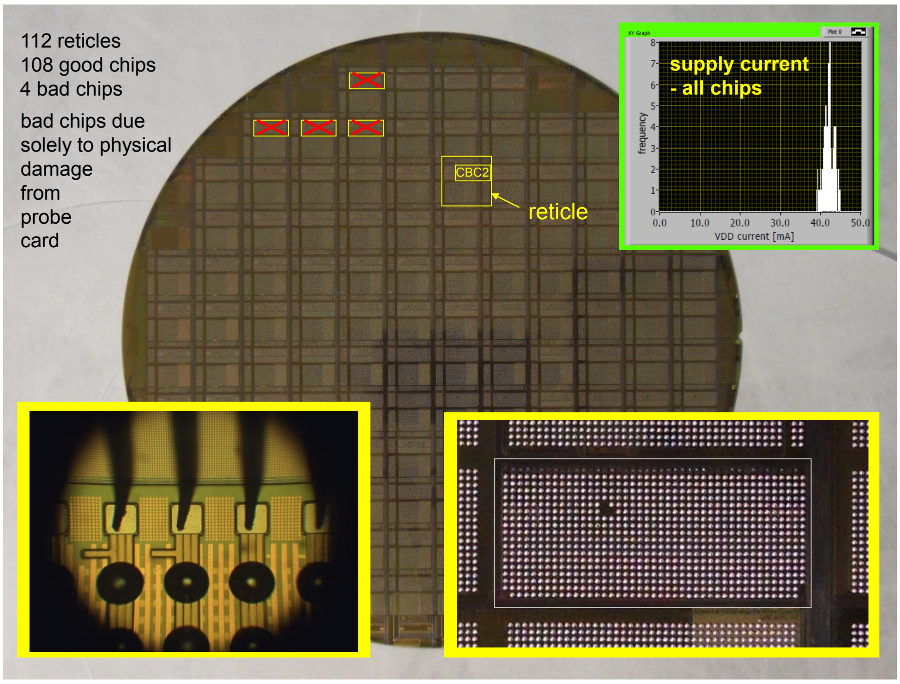

Figure 4.1. C4 CBC2 wafer. Insets show the C4 solder balls visible on the surface of the CBC2, some of the wafer-probe card needles in contact with the probe-able pads, and a histogram of the supply currents of all chips

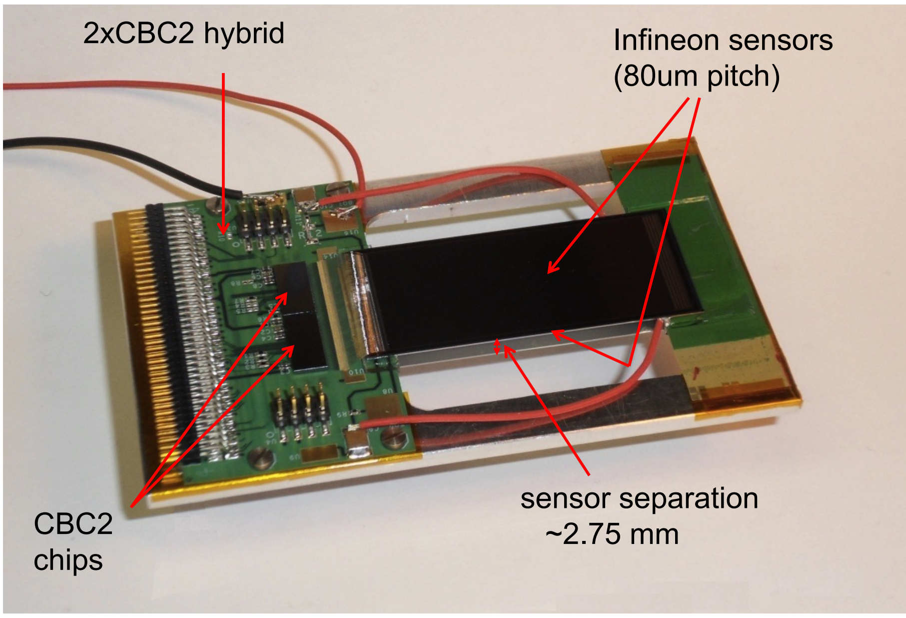

Figure 4.2. Prototype mini-pT module based on the 2xCBC2 hybrid

Figure 4.3 shows the correlation logic in operation on the mini-pT module. In this example the CBC2 chips have been programmed to generate a trigger if there is a coincidence between a cluster in one layer and a cluster occurring within a programmable window in the second layer (a cluster is a group of hits in up to 3 neighbouring channels). The clusters resulting from the passage of a cosmic muon through both layers has resulted in the chip generating a prompt trigger pulse. The full-chip readout has subsequently been triggered and the output data frame shows the hit channels that generated the trigger.

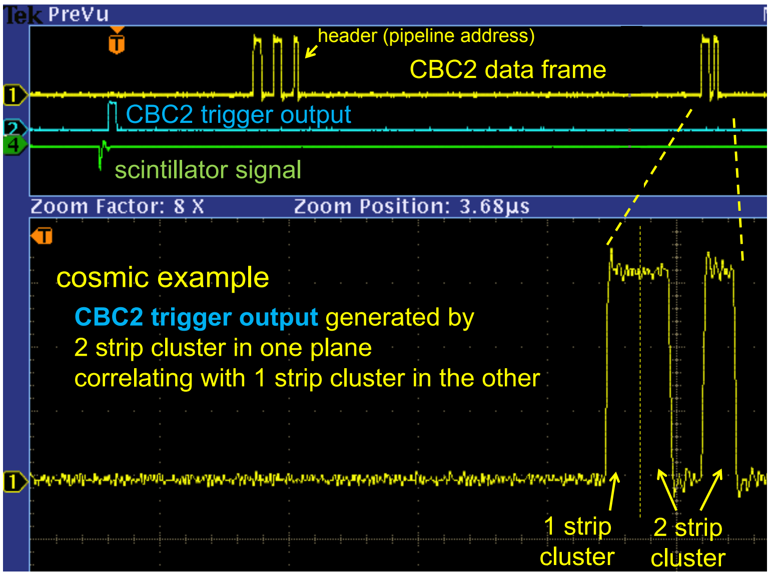

Figure 4.3. An example of oscilloscope observation of a triggered stub from a cosmic ray traversal of the prototype 2S-module.

The 2xCBC2 hybrids have also been used for electrical characterisation of chip performance. All core features have been verified functional, and front end performance is consistent with that measured for the first prototype, the front end circuitry in both chips being essentially the same apart from modifications to fix a few minor problems.

The next version of the chip, CBC3, is intended to be a final prototype containing all features required for the outer tracker system, including fast off-chip transmission of up to three stub addresses and associated information. Design of the CBC3 was intended to be underway by now, with a design ready for submission at PM12. There have been some recent discussions in the CMS Upgrade community concerning modifications to other sub-detectors to allow increased Level 1 trigger rate and latency. These discussions are on-going, but likely to result in an increase in both of these numbers and some adjustments to the CBC design will be necessary to cope with these changes. Fortunately we can accommodate some delay in the schedule because of the success of the CBC2, which will meet all our current needs for construction and evaluation of prototype modules.

The hybrid and mini-module production process has exposed some issues which are under investigation. Faulty channels appear after module construction, presumably caused during module assembly and bonding of the sensors. Better yield of working channels has been obtained with hybrids where chips have been under-filled. The flex-rigid hybrid solution appears more fragile than we would like, and alternative hybrid technologies are under investigation.

Constructional problems aside, the mini-modules are certainly working well enough to allow us to evaluate details of performance and the first test-beam study of these modules will take place at DESY in November. This is a collaborative effort, with the UK providing key components of the hardware, firmware, software and overall coordination. For this test the DAQ will be based around the CERN GLIB card, which has been produced in sufficient numbers, but it is expected that this will migrate soon to the UK developed FC7 card.

The FC7 (FMC carrier - Xilinx Series 7) is based on the MP7 and GLIB developments and can host two FMC modules and support signalling rates up to 10 Gbps. It is a flexible, general purpose card allowing it to be deployed across many applications in CMS (e.g. TTC upgrade, Phase I pixel FED upgrade) and its development has been shared with CERN-PH/ESE engineers who will eventually take responsibility for distribution and support. The board will form the basis of a scalable DAQ system both for testing of the CBC2 and for readout of 2S-PT module prototypes.

Work on the FC7 began in November 2012 with first prototypes delivered August 2013. Tests have shown that all IO lines are operational at full speed, all board peripherals are functioning, and the boards operated correctly in a µTCA crate. Minor issues have been corrected in a subsequent design revision to be submitted shortly. This submission will be treated as a pre-production run but also aims to survey the manufacturers before launching production next year. Once fabricated and tested, pre-production boards will be made available to start development of the DAQ system for 2S-PT modules.

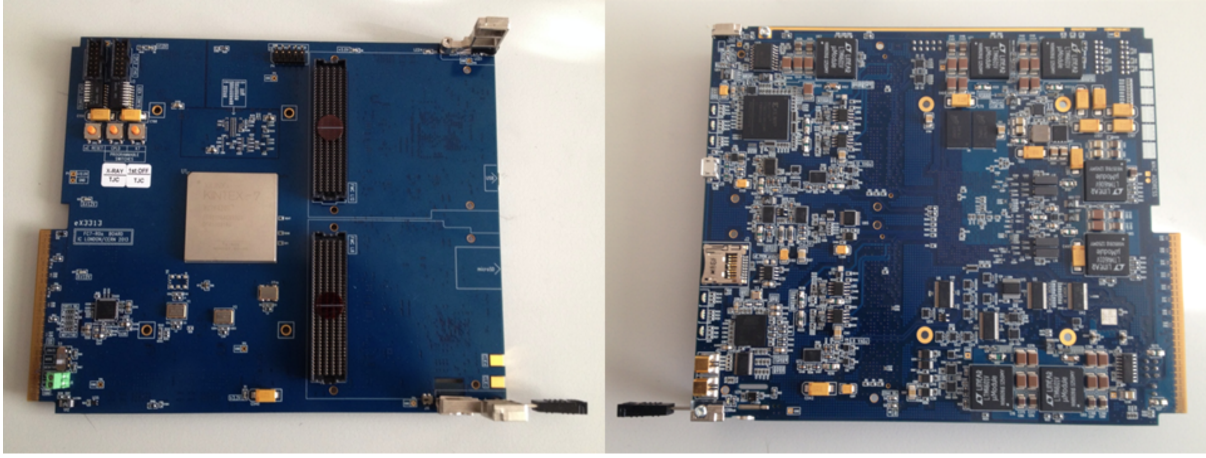

Figure 4.4. FC7 card prototype. Top face (left) shows the Kintex XC7K420 FPGA and the two connectors to take the FMCs. Bottom side (right) accommodates power regulators, clocking logic and services for FMCs and front panel interfaces (USB, µSD)

Meanwhile Bristol has taken on responsibility for developing test stands for CBC hybrids. The first of these stands will be installed at a handful of institutes to facilitate prototyping of the hybrids and modules. Ultimately, the stands will be used in production centres to certify the final components as they are built. The design of the stands utilises as much common framework as possible: the hardware is based on the CERN GLIB, and the software is built on the generic CMS DAQ chain (adapted for the CBC by Strasbourg). Bristol has produced a mezzanine card to connect the GLIB to the hybrid interface board, and already has written software to carry out basic tests on the first version of the CBC. Updated firmware and software from Strasbourg is expected by the end of October, which will allow the full testing of the CBC version 2.

4.3 Deliverables

The Work Breakdown Structure prepared at the time of the proposal has not changed and is listed below. There are three milestones due at the end of the first year (PM12). We continue to make progress in the system specification, having regular meetings with CMS colleagues responsible for other components on the 2S-pT module which have been very fruitful in developing the system ideas and sharing work. CERN and Strasbourg are the most prominent collaborators from outside the UK, but several other groups contribute and more are expected to become involved in module evaluation as the construction period approaches. There is also an overlap with the long term developments for track-trigger implementation and development of the off-detector processing.

CBC2 test results are regularly presented at collaboration meetings[1] and have also been presented at the recent 9th International "Hiroshima" Symposium on the Development and Application of Semiconductor Tracking Detectors HSTD9, and Topical Workshop on Electronics for Particle Physics, TWEPP13, with associated papers in preparation.

As explained above, there is some uncertainty about the Phase II upgrade while this is still formally an unapproved project. This is an evolving situation and we propose to review the milestones and deliverables during the next six months with a view to reprioritising the effort to match the overall CMS objectives. For example, at present the CBC and 2S-module are well ahead of the development of other modules, such as the PS but also the future pixel system, and it may be desirable to accelerate work in those areas to ensure CMS will be construction-ready at the time of the TDR.

[1] http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation

WBS	WBS L2	Start	Finish	Months	Task Description
2	Phase II tracker Readout	04/13	09/21	102	
2.1	**system**	04/13	03/14	12	**definition of the CBC-based 2S-Pt module readout**
	2.1.1 specification definition	04/13	03/14	12	regular meetings with CMS collaborators to define overall system specification and interfaces
2.2	**CBC2 test**	04/13	03/15	24	**CBC2 is final deliverable of the UK upgrade R&D**
	2.2.1 CBC2 ongoing testing	04/13	03/14	12	complete the detailed studies of the CBC2 chip, including irradiation and SEU tests
	2.2.2 CBC2 SS-Pt module prototype studies	04/13	03/15	24	a programme of 2S-Pt module studies, in collaboration with CMS, including test beam
2.3	**CBC3**	04/13	09/14	18	**CBC3 is specified for the final system**
	2.3.1 CBC3 design	04/13	03/14	12	design period
	2.3.2 CBC3 production	03/14	09/14	6	production period
	2.3.3 test setup preparation	03/14	09/14	6	wafer and chip test setup preparation
2.4	**CBC3 test**	09/14	03/17	30	**CBC3 chip and module testing**
	2.4.1 early tests	09/14	03/15	6	chip verification tests to prior to module tests
	2.4.2 ongoing testing	03/15	09/15	6	complete characterization, including irradiation and SEU tests
	2.4.3 CBC3 SS-Pt module studies	03/15	03/17	24	CBC3 based module studies in collaboration with CMS in lab and test beam
2.5	**CBC4 design and test**	03/15	09/16	18	**CBC4 is the final version of the chip, fixing any remaining bugs found in the CBC3**
	2.5.1 CBC4 design	03/15	09/15	6	design period
	2.5.2 CBC4 production	09/15	03/16	6	production period
	2.5.3 testing	03/16	09/16	6	tests to verify full and final functionality
2.6	**CBC4 mass production preparations**	09/16	12/17	15	**a full wafer engineering run is required for CBC4 in preparation for mass production**
	2.6.1 CBC4 final masks	09/16	03/17	6	mask preparation for full wafer engineering run
	2.6.2 CBC4 engineering run	03/17	09/17	6	production period
	2.6.3 CBC4 final production readiness verification tests	09/17	12/17	3	final functionality check
	2.6.4 procurement planning	01/17	12/17	12	detailed financial plans for mass production
2.7	**Production phase activities**	12/17	09/21	45	**wafer production, testing, modules assembly, integration and commissioning**
	2.7.1 mass production	12/17	06/20	30	wafer production
	2.7.2 production test	01/18	06/20	30	wafer testing
	2.7.3 modules assembly	01/18	06/20	30	module assembly
	2.7.4 integration activities	04/18	09/21	42	integration

4.4 Staff on project

The key people contributing to this WP are M. Raymond and M. Pesaresi at Imperial, supported by technician effort. A new RA, K. Uchida, took up her post in October and will contribute to this WP, at about 50% level, with the remainder being spent on track-trigger simulation activities in conjunction with WP3. In Bristol, J. Goldstein and D. Cussans are contributing to the module evaluation, along with effort on software from M. Grimes, also in WP3.

4.5 Expenditure

Recent expenditure has been concerned mainly with FC7 manufacture, most of this via CERN, and minor purchases in support of beam tests and evaluation studies.

5. Work Package 3: Level-1 Trigger

5.1 Objectives

- Improvement of the current CMS calorimeter trigger in preparation for above-design-luminosity conditions.
- Provision of infrastructure to allow testing of an entirely new calorimeter trigger in parallel with the existing system.
- Design, construction and testing of a time-multiplexed hardware trigger for CMS, capable of implementing new and more selective algorithms.
- Design of a track trigger architecture for HL-LHC running, and construction of a technology demonstrator.

5.2 Progress to date

The end of September saw a slice test of the time-multiplexed calorimeter trigger; a significant milestone in the calorimeter trigger upgrade schedule. The test was conducted in the dedicated CMS electronics integration area in Prevessin, France. It used two MP7 cards to simulate the ECAL and HCAL detectors and time-multiplex the data; a single MP7 as the main processor node; and another MP7 card for de-multiplexing the data (see photo).

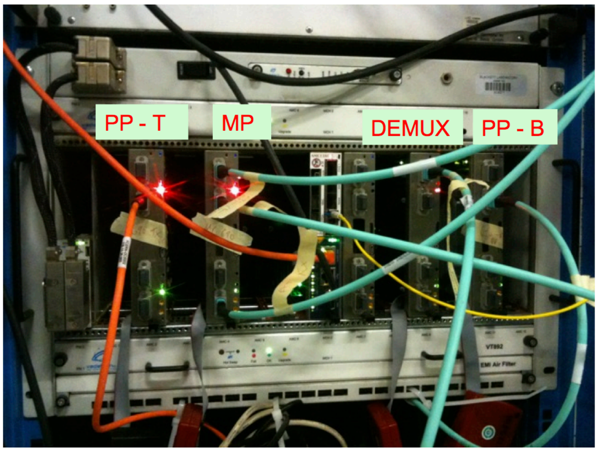

The test verified the stability of the 10 Gbps high-speed serial links; the latency of the system; and the ability to implement the large, complex algorithms given the high throughput of 720 Gbps. Results will be presented at the forthcoming Engineering Design Review in November (see below). This was a crucial test, scaling up the previous demonstrator system at Imperial to use close to the final hardware and realistic algorithms. All aims of the test were met and this success should have a very significant positive impact on the project.

An important development has been greater understanding of how to write the firmware for such large FPGAs so that internal timing is easily met. The complexity of the routing is illustrated in the figure, showing a screenshot of the routing tools. This will probably result in the algorithm code being slightly modified so that it can be more easily mapped onto the physical structure of the FPGA.

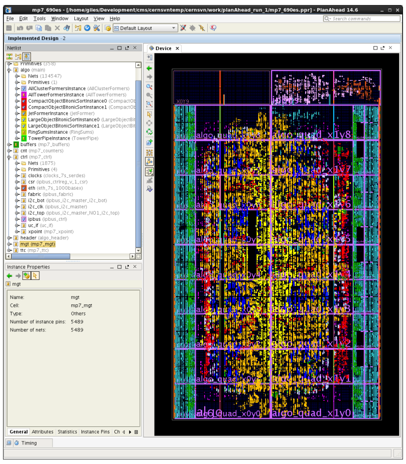

While risks remain for the TMT approach, they are substantially less than when the project was first conceived. The coming months will focus on the final algorithm development and the resource usage required – the latter determining whether the main processor node should be split over two MP7 cards. The current results looks close to that expected with algorithm resource usage running at ~30% with the bulk of the algorithms implemented.

In parallel with the milestone has been the delivery of six MP7-R1 cards. These form a pre-production batch. The order was placed in May, but several manufacturing failures resulted in the cards only arriving in September, despite several successful previous orders from the same company. The batch was assembled in two parts to minimise risk and while the first two cards were perfect three out of the last four had some minor defects. Tests are currently underway with simulated manufacture of seven cards at the assembly company, which will subject the cards to thermal loads. It is likely that the quality assurance process for the PCBs will need to be strengthened, with additional thermal stress tests.

Version 2.0 of the UK developed IPbus control system for µTCA has been released. This is the CMS-wide standard for upgraded systems, and has also been adopted by other collaborations and experiments. Work continues at a modest level on optimisation of the system to achieve higher bandwidths and bug fixing.

The Imperial group has developed a jet-finding algorithm which is able to estimate the energy contributed by pile-up interactions, on an event-by-event basis, and subtract this from the trigger objects. This algorithm gives very significant improvements in performance for multi-jet and H_T triggers and may also be used to improve the performance of isolated lepton triggers.

The RAL group are developing improved trigger algorithms for electrons and photons in collaboration with LLR. The focus of the RAL work is on isolation criteria, requiring similar pile-up subtraction to allow robust identification of electrons and photons in high luminosity running conditions. Initial studies indicate that using such an algorithm high efficiency for electron and photon triggers can be maintained despite the high levels of pile-up expected in future running.

The Bristol and Imperial groups are studying the expected rates and thresholds for different trigger objects under different LHC operating conditions. This work was started to study important physics channels, which are documented in the L1 Technical Design Report (see below) and is now continued to study scenarios after the Phase 2 upgrades.

Activity has begun on planning the next phase of the trigger upgrade, including tracking information. The UK groups are working closely with other CMS groups, for example, FNAL and CERN, to plan detailed studies of dataflow and first discussions on hardware demonstrators have started.

5.3 Overview of CMS plans

The Technical Design Report for the first phase of the L1 trigger upgrade was approved by the LHCC at their June meeting (see https://cds.cern.ch/record/1556311) and the draft budget was presented to the RRB.

The trigger upgrade project will have an internal CMS Engineering Design Review on 12-13 November at CERN. This partly overlaps with the OSC, which will necessitate limited representation at the OSC as it is essential that the trigger WP managers participate in the EDR in full.

Work to duplicate the input to the L1 trigger from the ECAL and HCAL, which can only realistically be done during the current shutdown period, has been delayed due to technical problems. The problems have now been solved, but much of the contingency in the schedule for this work is exhausted.

Preparations are already underway for CMS running despite the end of the LS1 upgrade still over a year away. The GCT, the current UK commitment to the CMS Calorimeter Trigger is currently being made ready for the first cosmic-muon runs scheduled in November. The computers used for control will be replaced and several key software packages will change.

5.4 Deliverables

The UK project is on track to meet upcoming milestones, which were revised in February 2013 prior to the final funding allocation. As explained earlier, the overall planning will be reviewed following the November EDR.

L1	L2		Start	Finish	PM	Task description
3.1		**Stage-1 calorimeter trigger upgrade**				
	3.1.1	Hardware development		07/13	6	Finalisation of production hardware module (48-link version)
	3.1.2	Procurement and testing	07/13	10/13	3	Procurement, production and acceptance tests of hardware
	3.1.3	uTCA infrastructure		07/13	6	Completion of baseline IPbus / uHAL
	3.1.4	Online software development	04/13	10/13	6	Development of system-specific and trigger-wide online software (control, monitoring, DAQ)
	3.1.5	Algorithms and offline software	04/13	04/14	12	Development of stage-1 algorithms and corresponding emulator and DQM software
	3.1.6	Integration	07/13	01/14	6	Integration tests with other trigger components, DAQ, TTC
	3.1.7	Commissioning	09/14	03/15	6	Commissioning with cosmics and beam
	3.1.8	Support	03/15	01/16	9	Ongoing expert support and optimisation of Stage-1 system
3.2		**Stage-2 calorimeter trigger (TMT) upgrade**				
	3.2.1	Hardware development	10/13	04/14	6	Development and finalisation of production hardware module (72-link version)
	3.2.2	Procurement and testing	04/14	10/14	6	Procurement, production and acceptance tests of hardware
	3.2.3	Online software development	10/13	04/14	6	Development of system-specific and trigger-wide online software (control, monitoring, DAQ)
	3.2.4	Algorithms and offline software	04/14	04/15	12	Development of stage-2 algorithms and corresponding emulator and DQM software
	3.2.5	Integration	04/14	10/14	6	Integration tests with other trigger components, DAQ, TTC
	3.2.6	Commissioning	04/15	04/16	12	Commissioning with cosmics and beam
	3.2.7	Support	04/16	04/19	36	Ongoing expert support and optimisation of stage-2 system
3.3		**Post-LS3 trigger R&D**				
	3.3.1	Design studies	04/13	10/14	18	Simulation studies of track trigger performance, and decision on final concept
	3.3.2	Dataflow design	10/14	10/15	12	Detailed simulation, architecture design and technology choices for track trigger
	3.3.3	Hardware development	04/16	10/17	18	Development of next-generation hardware modules for integrated L1 trigger
	3.3.4	Algorithms and offline software	10/15	04/17	18	Development of algorithms and firmware for integrated L1 trigger
	3.3.5	Integration and demonstration	10/17	10/18	12	Hardware slice test of integrated L1 trigger
	3.3.6	Final system design	10/18	04/19	6	Production planning for final version of integrated L1 trigger

Progress with the calorimeter trigger was presented recently at the Topical Workshop on Electronics for Particle Physics, TWEPP13, and the associated paper is in preparation.

5.4 Staff on project

Reported in accompanying tables. The project is now fully staffed. One of the new RAs (S. Paramesvaran, Bristol CG post since July) has been appointed CMS Calorimeter Trigger Commissioning Coordinator.

5.5 Expenditure

Most of the expenditure to date has been committed to MP7 manufacture and component purchases. Some finds have been committed by individual groups for µTCA infrastructure. Overall spending is well within the budget foreseen at this stage.

5.6 Comparisons with CMS activities elsewhere

The MP7 remains well ahead of other comparable developments in both CMS and ATLAS. There is no other board which is operating at 10 Gbps and certainly not at the same level of development, with a significant number of boards produced and in operation in labs and a successful system test achieved. We plan to provide MP7 cards to other groups within the L1 trigger project with whom we wish to collaborate, for example the LLR group, who collaborate with RAL on the development of electron/photon algorithms and the Vienna group who plan to base their design for the CMS Global Trigger on the MP7 card.

The first layer of the calorimeter trigger is planned to be built by the University of Wisconsin, based on cards with the same FPGA as the MP7 but fewer optical links. It is currently in design. A CMS progress review in August identified this card as a “significant schedule risk” for the calorimeter trigger upgrade, while concluding that the MP7 was working to specification. It is recognised by CMS that the MP7 card is capable of fulfilling the role of the CTP7 card should this be necessary.

7 Risk register

A new risk register was prepared for the PPRP proposal in 2012 using the criteria and numerical evaluation recommended by STFC. It has recently been reviewed to remove risks associated with work not now included in the project and to add some extra risks associated with board manufacture. Before undertaking further revision, we seek feedback on the preferred way of reporting and handling risks especially since several of them are not directly under UK control. One of them concerned schedule changes and for the Phase II upgrades this is still a matter under active discussion.

8 Finances

The financial report is summarised in the attached tables, approximately following the format used in the past. The expenditure to date includes expenditure in universities and also commitments made at CERN which have not all yet been paid in the UK but are imminent.

A few questions were raised in preparing the summary:

Grant announcements are sent to individual institutes but are not provided to the PI. Because of adjustments made for indexation and Wakeham efficiency factors, the allocations are not quite the same as the final proposal submissions made to STFC. For the present, only the Imperial entries match the grant award.

Reporting Rolling or Consolidated Grant spending, or existing staff costs in the case of RAL PPD, has been a confusing issue in the past, and it is difficult to track precisely with overhead costs. Some guidance would be helpful.

For RAL expenditure it is most convenient to take figures from SBS but actual PPD overhead charges are reported to be different than what SBS records.

For monitoring actual spending, it might be more useful to distinguish expenditure by individual institute. The table has been slightly reorganised to allow this to be done better, but the OSC and STFC views on how best to track expenditure would be helpful.

It may be helpful to organise the table differently but a number of options are possible – e.g. by institute, allowing to compare the expenditure against the individual grants,. If alternatives are preferred, it would still be helpful to be able to compare total individual group expenditure with allocations.

In the first six months equipment spending has been primarily MP7 and FC7 components, board fabrication and assembly and institute spending on μTCA crates and ancillary equipment.

Engineering staff spending at RAL has deliberately been slightly lower than the original plan while the options for the future of WP2 are considered.

There have been appointments to all the posts opened which in most cases required some months to complete the recruitment process. M. Grimes (Bristol) was available from the outset, since he was already employed by Bristol but others (S. Dasgupta [Bristol, July], A. Thea [RAL PPD, July], A. Bundock and K. Uchida [Imperial, October]) started later. In view of the excellent progress elsewhere this has not had a negative impact and in any case, is to be expected.

The Working Allowance is held at Imperial as part of the equipment budget. However, since most new appointments took several months to arrange, there is a natural margin in staff budgets.

We hope to use some of the first meeting to discuss the best format for financial reporting.

9 Gantt charts

The Gantt charts prepared for the proposal, and updated in February 2013 following the provisional award to reflect the allocations, are included.

10 Milestones

The deliverables from each work package are listed below. As the first of them is not yet due, we simply record them here.

Deliverable	Date	Description
M2.1	PM12	System specification document produced
M2.2.1	PM12	Documented CBC2 detailed test results
M2.2.2	PM24	Documented 2S-Pt module results
M2.3.1	PM12	CBC3 ready for production
M2.3.2	PM18	CBC3 produced & test setups ready
M2.4.1	PM24	Documented early CBC3 test results
M2.4.2	PM30	Documented CBC3 detailed test results
M2.4.3	PM60	Documented CBC3 2S-Pt module results
M2.5.1	PM42	CBC4 ready for production
M2.5.2	PM48	CBC4 produced
M2.5.3	PM54	Documented CBC4 test results
M2.6.1	PM60	Final production masks prepared
M2.6.3	PM69	CBC4 ready for mass production
M2.7.3	PM72	First production modules available
M3.1	PM9	Stage-1 calorimeter trigger hardware tested and installed
M3.2	PM18	Stage-2 calorimeter trigger hardware tested and installed
M3.3	PM23	Stage-1 calorimeter trigger commissioned & system ready for physics
M3.4	PM30	Post-LS3 trigger dataflow design completed
M3.5	PM35	Stage-2 calorimeter trigger commissioned & system ready for physics
M3.6	PM54	Post-LS3 trigger prototype trigger modules produced and tested
M3.7	PM66	Demonstration of post-LS3 trigger slice
M3.8	PM72	Post-LS3 trigger construction plan delivered