


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	5
3.3	Deliverables and Milestones	6
3.4	Staff on project	6
3.5	Expenditure	6
4.	Work Package 2: Outer Tracker Readout	8
4.1	Objectives	8
4.2	Progress to date	8
4.3	Deliverables	10
4.4	Staff on project	10
4.5	Expenditure	10
5.	Work Package 3: Design of Level-1 Calorimeter Triggers at the SLHC	11
5.1	Objectives	11
5.2	Progress to date	11
5.3	Overview of CMS plans	12
5.4	Deliverables	13
5.4	Staff on project	13
5.5	Expenditure	13
5.6	Comparisons with CMS activities elsewhere	13
6.	Work Package 4: Pixel tracker development	15
6.1	Objectives	15
6.2	Progress to date	15
6.3	Deliverables	17
6.4	Staff on project	18
6.5	Expenditure	18
6.6	Comparisons with CMS activities elsewhere	18
7	Risk register	20
8	Finances	20
9	Gantt chart	20
10	Milestones	20
	Appendix: PPRP proposal	20

1 Executive Summary

The R&D project is now mid-way through the final year. Notable progress since the last report includes:

The LHC continues to run extremely well, to date (late August) accumulating 12.8 fb^{-1} of integrated p-p luminosity at 8 TeV centre of mass energy in 2012, with 11.9 fb^{-1} recorded by CMS, so high operational efficiency. The principal Higgs search results were widely publicised in early July when a 125 GeV mass [± 0.4 (stat.) ± 0.5 (syst.)] boson discovery was announced by CMS, with complementary results from ATLAS. The object has many similarities to a Standard Model Higgs but it will be some time and (much) more data before its full properties are elucidated. However, the motivation for experiment upgrades has strengthened as a result. Not only does the discovery warrant deeper investigation with a more advanced detector, but the challenges of detector operation under present LHC conditions are becoming more apparent.

The first LHC Technical Stop, LS1, has been delayed to allow extra data taking with protons until the end of this year, with heavy ion running in early 2013. The stop should now commence in March 2013 so LHC operation should begin again around March 2015.

There has been steady progress in developing and improving software for upgrade simulations, especially for Phase I tracking studies. This has been exploited for TDR studies for the pixel upgrade and to study the trigger requirements.

Earlier this year, CMS adopted IPbus, the software and firmware framework largely developed in this R&D project, as a common standard control system for all upgraded μ TCA crate-based electronics and the latest full version of the software and firmware has recently been released.

The first official CMS-wide release of MicroHAL - the robust, high performance and scalable Ethernet Hardware Access Layer (HAL) has been made after a period of extensive beta testing with groups in the UK, CERN and the US.

There has been significant progress with calorimeter trigger simulations, complemented by contributions from UK non-upgrade project funded staff.

The MP7 trigger board was received back from assembly in late July. There were some problems with the manufacture, in particular with power supply components, but they have not prevented initial evaluation and some of the important features have been shown to work successfully.

The CBC2 was submitted for manufacture in July 2012 and is expected back in October. It is a coarse-pitch bump-bonded chip suitable for track-trigger modules. There have been further developments with CERN on modules matching the CBC2 layout.

The first beam test to measure inefficiencies of the upgraded digital pixel ROC at high rate has taken place using UK hardware and firmware, to provide a beam telescope with 8 pixel modules.

A PPRP proposal for construction and continued R&D funding for six years to March 2019 was submitted at the end of June for presentation at the 20 September meeting. This was a substantial effort, which should lead to a balanced programme of construction and further R&D exploiting the developments of the last three years.

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 up to $2.5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ being Phase I, occurring in 2018 in the most recent CERN ten-year plan, and a further major increase in the next decade, probably to $\sim 5 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ levelled luminosity.

The recently announced delay in the first Long Shutdown (LS1) may also delay LS2 but the long term CERN operations plan has not recently been updated and probably awaits operating experience at full energy and with 25 ns bunch spacing.

The experiment and accelerator continue to operate extremely well, continuing regularly to break luminosity records. This has had a big and well publicised impact on the Higgs search and CMS is

intensifying its efforts to exploit this for the Higgs studies and expand the range of searches for other new physics, particularly SUSY where much of the phase space of many of the simpler, constrained, models has been excluded. However, it is well recognised that the SUSY parameter space is huge and the likely existence of the Higgs around 125 GeV appears to strengthen some of the arguments for certain kinds of SUSY, or mass hierarchies, so search strategies are under review to profit from the flood of data and the most up to date theoretical responses to present results. UK groups have been very active in Higgs, SUSY and heavy vector boson searches, among other topics, as explained in the CMS grant request of February this year.

2.1 LHC upgrade schedule and planning

The main change to the LHC operations schedule since the last OSC meeting has been to delay the start of LS1 from September 2012 to March 2013. This was largely motivated by the wish to acquire the maximum amount of p-p collision data to ensure the boson discovery and its properties can be made as definitive and comprehensive as possible using all potential decay channels, although it also seems likely that the accelerator can profit from a few months delay to be better ready for their modifications during LS1.

The most significant signal excesses are seen by both CMS and ATLAS in the $\gamma\gamma$ and $ZZ \rightarrow 4l$ channels. CMS was again successful in analysing all major decay modes ($\gamma\gamma$, $\tau\tau$, $b\bar{b}$, $WW \rightarrow 2l2\nu$, $ZZ \rightarrow 4l$, $2l2\nu$, $2l2q$, $2l2\tau$), in July using then available 2012, plus 2011, data. There are intriguing hints of possible differences from the simplest Standard Model Higgs predictions but only additional data can allow experimental studies to pass more definitive judgements. The extra data, including some which is being parked (i.e. stored but not analysed for the present, as it exceeds our real time processing capacity) can be further studied during LS1 and should enhance other searches for new physics, particularly SUSY. There have not been recent projections of the total integrated luminosity expected in 2012, but it seems likely that 30 fb^{-1} might be reached, which is a significant gain compared to the 5 fb^{-1} of 7 TeV data in 2011, and also the 5 fb^{-1} of 8 TeV data available in July.

The LHC continues to operate with 50 ns bunch spacing and, although it is expected that 25 ns operation will be the norm from 2015, it is unclear how challenging this will be for the accelerator staff to achieve. To maximise data collected in 2012, 25 ns operation was restricted to machine development sessions; it is not clear how many studies have recently been carried out. The electron cloud effect remains the major concern and how much this will limit the achievable luminosity.

The penalty of 50 ns operation for the experiments is increased pile-up and loss of efficiency in data analysis. However, this was estimated to be a modest effect (e.g. $\sim 15\%$ impact on the $\gamma\gamma$ channel with 25 events pileup) compared to the loss of integrated luminosity expected by operating using 25 ns spacing. In 2012, pileup, at least at the start of a run, has exceeded 35 events per crossing and there are examples of much higher levels. Typically the decay time constant is several hours but runs are terminated when the luminosity is around 50% of the initial value, or earlier for other reasons. The design assumption, using 25 ns and $10^{34} \text{ cm}^{-2}\text{s}^{-1}$, was for 20-25 events average pileup so the experiments are already operating beyond their design parameters.

So far, the High Level Trigger has been able to cope, as extensive efforts have been made to use all available means to limit non-linear effects increasing processing times. It is not clear how far this can continue in the short and intermediate-term, although the future upgrade of the pixel detector will contribute significantly to track seeding improvements. The latest instantaneous luminosity record is $7.6 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, i.e. corresponding in pileup terms to over $1.5 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ if the machine were running at 25 ns. This is also at only half the design energy and machine experts predict future operation at 13-14 TeV with up to $2.5 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ before LS2 in 2018.

Even more of a challenge is the Level 1 trigger. Some short extra-high pileup runs have been carried out to evaluate the effect on trigger performance. The results are still under study, including possible optimisations of the hardware or firmware processing, but it is clear that CMS faces a major challenge in the next period of operation. Some recent UK studies demonstrated that jet rates could be reduced significantly by requiring a jet seed energy deposit of a few GeV and firmware modifications were incorporated in the GCT to implement this. While beneficial, it still only scratches the surface of the overall rate issue and each increase in luminosity is applying new trigger menus designed to

optimise data selection and thresholds without compromising the acceptance of interesting events too much. However, it reinforces the need for the calorimeter trigger upgrade to be carried out expeditiously with the best performance.

2.2 CMS planning

The CMS Technical Design Reports for sub-detector upgrades are progressing steadily. The pixel TDR is in its final draft and will be submitted imminently for the LHCC September meeting. The HCAL TDR, which is mostly not directly relevant to UK activities, follows closely behind the pixel TDR. The trigger TDR is expected to reach the LHCC about six months later. Longer term Phase II Tracker planning has not changed recently.

There is an ongoing discussion of the actual implementation of the CMS trigger upgrade, partly influenced by the current US funding situation which foresees their expenditure coming relatively late. US representatives at recent RRBs have endorsed the model of CMS sharing, according to PhD authors, but have been unable to fully commit funds in the present economic climate. The US approach has been to optimise their use of existing operations funding (which has a wider scope than UK M&O, for example) but this is clearly insufficient to commit to all their chosen projects so forces prioritisation. Presently the pixel upgrade is well supported and expected to remain so, along with significant commitments to HCAL and muon system upgrades, but the trigger activities seem to have been assigned lower priority.

CMS is attempting to confront this situation but progress has been slow. However, in view of the TDR, decisions are expected soon, such as on the calorimeter architecture and any intermediate steps towards the trigger upgrade. Several internal reviews of the upgrades have already taken place and others are imminent. It is hoped this will help resolve the remaining uncertainties. The delay in LS1 does help in moving the window for procurement and installation of key components so this works in our favour.

2.3 UK adaptation to CMS planning

In parallel with the R&D activities, the last six months have seen a substantial effort devoted to preparation of the funding bid for the next, mainly construction, stage of the upgrade. Rather than attempt to summarise this here, it seemed better to provide the document to the Oversight Committee. We apologise for the possible information overload, but the Gantt charts, milestones and cost and risk assessments give a better idea of our expectations, as well as the possible concerns, for the next phase of the project than those associated only with the present R&D.

In the last meeting we raised the question of continuity of funding and a possible request to STFC to put in place bridging funds for a six-month period from the end of the present grant. It is not yet clear what the final approval timescale will be so we reiterate this request.

We have received confidential referee reports, which seem very positive, and we would be happy to share these with the OSC if desired. We would also welcome comments from the OSC to the PPRP, perhaps via the STFC office, to enable the PPRP to have an independent view of our progress in the R&D phase in addition to our own explanations.

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. For the last year, the project has focussed entirely on the Phase-I upgrade. 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.

3.2 Progress to date

WP1 is now in the ‘extension period’ to the original work plan, supported by new funding released from June 2012. Work during this period has focussed on the support of deliverables from the earlier phases of the project, and on preparations of Technical Design Reports for the pixel and L1 trigger upgrades, thus fulfilling the main objective of WP1: to facilitate and conduct design studies of the upgraded detector systems.

At the conclusion of the current project, in March 2013, the activities of WP1 will be reorganised as parts of the overall pixel and trigger activities of a successor project, with certain elements (e.g. online software) bridging the two. For the current round of trigger performance studies in preparation for the L1 TDR, WP1 and WP3 are already acting as a single entity overseen by both WP managers; we report progress on this area mainly in the WP3 section.

The TDR for the pixel system is currently at the editing stage, with the UK having made a substantial contribution to both simulation studies and detector prototype tests. In the latter case, the project status is reported in the WP4 section, but has depended substantially on the online firmware / software framework produced in WP1. The TDR is due to be published in September 2012. For the trigger system, the TDR is currently being written, with the key ideas for the L1 upgrade to be reviewed in a Conceptual Design Review in late 2012.

As in previous reports, we note that all aspects of our work have regularly been presented at international workshops and conferences during the reporting period.

Simulation tools: A key step in the CMS event reconstruction chain, identified as a bottleneck in performance studies, is pixel track seeding in the High Level Trigger (HLT). The rapid identification of track seeds is essential for adequate trigger performance to be maintained at high pileup, with the complexity of this operation increasing after the pixel detector upgrade due the availability of a fourth tracking layer. I. Reid has investigated the use of advanced data structures (k -dimensional trees) to improve the performance of seed-finding code without jeopardising efficiency. Results so far show a radical improvement in code performance, which will be beneficial for HLT performance even before the replacement of the pixel detector.

The upgrade L1 menu studies under way in WP1 and WP3 have been approached so far via a combination of simulation studies and use of real data from special runs of the LHC, giving extremely high pileup levels in single bunches, but with very low overall luminosity. Simulation of extremely high pileup scenarios via simulation is demanding, and the data-driven technique has limited statistical reach. We have therefore turned to a ‘data-mixing’ approach, whereby recorded events are overlaid at the energy-deposit level and then re-digitised and reconstructed as single events. Integration of all three techniques into a common trigger menu framework, including changes to the CMS low-level event processing framework code in some places, has been undertaken by J. Brooke and C. Lucas and US collaborators. Results from this new package are now being evaluated, though more work remains

to validate the treatment of out-of-time pileup signals from bunch crossings preceding the triggered one.

Physics performance studies: M. Grimes has continued his work on the demonstration of the upgrade pixel detector performance via the exemplar Z+H production channel, with the Z decaying to muons and the Higgs to b-quarks. Results from this study appear in the pixel upgrade TDR. Since this work has been in progress for some time, it has been optimised as a Higgs boson search rather than a measurement. However, results from the study are still a useful guide for optimisation and performance estimation for the new detector. The conclusions of this first study are that the upgraded detector allows around a 65% improvement in signal efficiency at high luminosity compared to a three-layer detector of identical design. Improvements in the efficiency of double-muon-triggered event selection is dramatic, with a 175% increase of signal event count. Since the electron reconstruction in the new detector has not yet been fully tuned, a further efficiency improvement is expected. Work in this area will now turn to development of realistic data-handling and reconstruction software to allow the fully exploitation of the new detector.

Online software: The IPbus system developed by CMSUK has now been adopted by CMS as the common standard control system for all upgraded electronics. A major milestone in this subproject has been passed with the release of the next full version of the software and firmware components. This release contains many features identified in an extended review of the system by future CMS users, and incorporates a simplified and rationalised API at the top level. Work in this area is now focussing on improvements to the performance and robustness of the system, and on support for the first tranche of users for the system. In order to achieve the first goal, some simplifications to the protocol have been proposed, which will now be implemented and released as the next major version of the system. In addition, work has begun on the integration of IPbus as part of the overall CMS hardware control and monitoring system, including interfaces to the Detector Control System and existing higher-level control software. The IPbus development has attracted new collaborators from within CMS, and is also now in use or under evaluation by subprojects within the ATLAS, LHCb and COMPASS experiments.

3.3 Deliverables and Milestones

The project is now reporting against the additional milestones proposed to cover the period up to March 2013.

- M1.3.9 has been achieved, with the release of the next major version of the IPbus system, though large-scale testing at the CMS Electronics Integration Facility originally foreseen has been postponed until 2013.
- M1.3.8 is due for completion in the coming weeks, and is on schedule. The work in this area is proceeding as part of the overall L1 upgrade TDR programme, and the results will be published as part of that document.
- M1.3.10 is essentially complete, though results are currently being reviewed by the collaboration and further work may be necessary in response to collaboration feedback. This work will be documented as part of the pixels upgrade TDR.
- M1.4.1 and M1.4.2 are on track, though the deliverables and scope of the corresponding subproject within CMS are currently under discussion as part of a new upgrade online systems project.
- M1.4.3 and M1.4.4 concern the publication of pixels and L1 TDRs, and are on track.

3.4 Staff on project

Reported in tables. We note that the tables have been updated to reflect the extension to WP1, with the additional personnel costs reported as part of the overall planning.

3.5 Expenditure

31 August 2012

The expenditure to date is reported in the financial tables; WP1 is dominated by staff costs, with some travel. Expenditure required in the integration of WP1 online software with prototype trigger and readout hardware in WP3 and WP2 is generally covered in the budget of those WPs.

4. Work Package 2: Outer Tracker Readout

4.1 Objectives

The objectives of WP2 were to develop a readout chip suitable for the Phase II outer tracker, to study options for providing Level 1 trigger, and to contribute to development of a complete readout system, including off-detector components for the front-end DAQ.

As explained in the previous report, the DAQ component of this work package has become integrated into the pixel FED activity, where effort is required sooner than Phase II. Also, the outcome of studies into the tracker level 1 trigger contribution, where the SS-pT module concept has become established as a viable option to provide high-pT stub information, has had significant implications for the readout chip architecture definition which has now evolved to include prototype triggering functionality, building on the successful first iteration of the triggered-readout-only CBC.

4.2 Progress to date

Front end chip

The 130 nm CMOS CBC is a 128 channel wire-bondable chip for readout of short (~ 5 cm) microstrips. Good performance has been successfully demonstrated in the lab and test beam, and results have been presented at conferences¹², most recently at the WIT2012 Workshop on Intelligent trackers in Pisa in May³, published⁴, and reported extensively at CMS meetings⁵. A further test beam study using the CBC prototype is planned this October using a 5 cm long, 80 μm pitch, microstrip sensor that more closely matches the geometry proposed for the Phase II outer tracker SS-pT modules (the previous test beam sensors were remnants from CMS production and had much larger pitch strips). Test beam results with this prototype CBC/sensor system can then be compared directly with results from the next CBC iteration, CBC2, which will also be used to read out the same sensors.

CBC2 (fig. 4.1) has been designed to match the requirements of the SS-pT modules proposed for the Phase II outer tracker region (fig. 4.2). SS-pT modules will contribute 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 (see previous report for a more complete description of the SS-pT module concept). The CBC2 concept was also presented at the WIT conference⁶. Much of the CBC prototype core circuitry remains with only minor modifications, including front end amplifier architecture, analogue bias generator, slow control interface, pipeline and triggered data readout circuitry, powering circuitry (DC-DC switched capacitor circuit and low-dropout linear regulator). Major modifications to implement functionality required for SS-Pt module readout include:

- Expansion to 254 channels to allow correlation between two sets of 127 sensor channels.
- Cluster width discrimination logic: a programmable circuit will exclude events where signals over threshold are detected in more than three neighbouring channels (in each layer), as these cannot be consistent with high pT tracks.
- Correlation logic: If a cluster is detected in the inner layer (closest to the beam interaction point) and a corresponding cluster is found in the upper layer, within a window of programmable width and offset, then a prompt trigger output is transmitted off-chip to signal the presence of a high pT stub. The window width defines the pT threshold in conjunction with the layer separation.
- Bump-bond layout. The chip has been laid out to use C4 bump bonding which will allow robust multi-chip substrates to be commercially produced during the large scale production

¹ http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/Lawrence_FEE2011_Bergamo.pdf

² http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/TWEPP2011.pdf

³ http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/Pesaresi_CBC1.pdf

⁴ <http://iopscience.iop.org/1748-0221/7/01/C01033/>

⁵ http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/

⁶ http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/Braga_CBC2.pdf

phase, with fine pitch routing of tracks to wire-bond pads on both sides of the substrate for sensor mounting.

- Test pulse system. A circuit has been included to allow programmable timing and amplitude signals to be injected into all front end channels.

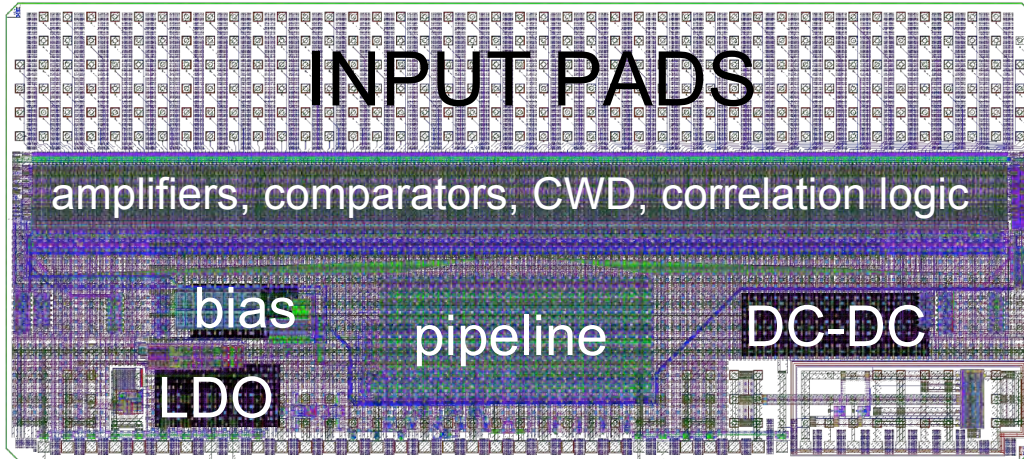


Figure 4.1. CBC2 layout [overall dimensions 11 mm x 5 mm]




Figure 4.2. SS-p_T module concept

Design of the CBC2 has dominated WP2 activities for the last six months. The chip was submitted in July and wafers are expected to be back in October. A number of test setups are now in preparation.

To produce the 8-chip substrates required for SS-p_T modules with acceptable yield it is necessary to probe wafers to identify good chips. It is not possible to probe C4 pads, so the powering, control and data readout pads have been duplicated as wire-bond pads (lowest row of pads in fig. 4.1) which can be probed. But this means that the chip can also be wire-bonded, so a chip carrier board has been designed. Since the layout has been optimised for C4 bonding, best performance is not expected that will be achieved, but basic functionality can be confirmed and this test system can also be used to develop wafer probe software.

The UK was responsible for wafer probing all the APV25 chips used for the existing CMS micro-strip tracker, so equipment and expertise exists to perform this task for CBC2 wafers. The hardware

interface has to be adapted to the CBC2, a probe card has to be designed and procured, and test software must be developed. The target completion date for this is the end of 2012, so that tested wafers for the manufacture of 8-chip substrates can be provided early in 2013.

Detailed characterization of CBC2 performance can only be achieved once chips have been bump-bonded. CERN has taken primary responsibility for the design of suitable chip carrier substrates and we have worked closely with them, cross-checking to ensure compatibility with the CBC2 design and that the layout will meet our needs. In the first instance we will use a substrate (2CBC2) that takes 2 chips, to verify the inter-chip communication interface where neighbouring sensor strips are bonded to different CBC2 chips and signals have to be exchanged to avoid discontinuities in the correlation operation at chip boundaries. The 2CBC2 substrates are already designed and will shortly be submitted for manufacture. The plan is to populate them with untested die, to avoid delays while wafer probing is developed. Provided individual chip yield is high, the yield of fully working 2-chip substrates should be sufficient. Only when populating 8-chip substrates do known good die become essential.

The 2CBC2 substrate assemblies will be the main vehicle for evaluating the CBC2 performance in the medium term, and have been designed to allow electrical characterisation as well as with sensors, where a mini 2S-pT module will be assembled using the same sensors to be used for the CBC prototype in the October test beam. We are responsible for providing interface electronics to adapt low voltage CBC2 signals to higher voltage standards compatible with existing test systems interfaces.

Documented results of testing the CBC2 will be the final deliverable of this phase of the CMS UK Upgrade programme. A working chip will allow construction of mini (2CBC2) and full-size SS-pT module prototypes, the evaluation of which will begin in 2013 and extend throughout and well into 2014. Definition of the final specifications required for these modules continues within CMS and we continue to be active in that process. In the next version of the chip, CBC3, it will be necessary to prototype all features required to fully implement the system, the most important being a high speed interface between the CBC and concentrator & controller chips on the modules to transfer stub addresses. The design phase of CBC3 should begin in Spring 2013.

4.3 Deliverables

The only remaining milestone for WP2 is **M2.4.1** (09/12) *Documented results of final prototype*. As explained in the last report, this milestone will be delayed as we do not expect the CBC2 wafers to arrive before mid-October. We expect to verify functionality soon after delivery, on chips cut from the wafers and wire-bonded to chip carriers. A deeper understanding of the performance can only follow once chips are bump-bonded on the 2CBC2 substrates, and this is currently foreseen to take place around the end of the year. We therefore expect to meet this milestone partially before the end of this year, and fully early in 2013.

4.4 Staff on project

Listed in the tables. No significant changes.

4.5 Expenditure

The ASIC submission was shared with two other designs, the CBC2 share being \$182k.

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

5.1 Objectives

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

Our system consists of generic modular devices based on FPGAs and optical links and is the first prototype of off-detector trigger electronics for HL-LHC, as well as the Phase I trigger.

5.2 Progress to date

The UK was requested by the Calorimeter Trigger Review (Oct 2011, Report Dec 2011) to continue “*development of a common L1 trigger hardware platform*” given that the UK had the “*only available prototype with performance approaching the requirements*”.

This has led to development of the final trigger processing card – the MP7. It has been specified with the help of the CMS trigger community and recently returned from manufacture.

It uses optical interconnects rather than the backplane for trigger data paths, thus allowing it to implement any type of trigger architecture, in contrast to the alternative under development. This approach also keeps the standard backplane ports free, thus allowing users to make full use of standard commercial μ TCA capabilities (e.g. 10GbE, SRIO & PCIe).

The MP7 will feature an optical interface of up to 1.44 Tb/s and a 64 Gb/s electrical interface on the front panel with GbE, DAQ, TTC, etc supported on the backplane. The design is scalable so that, depending on the requirements, the most appropriate FPGA and optics may be selected and thus cost minimised.




Fig. 5.1. The MP7 Trigger Processing Card (before mounting of optical components).

Prior to MP7 manufacture, a careful study was made using some simple prototype boards of 10 Gb/s signal propagation across the board. It was demonstrated the direction of propagation relative to the weave of the board substrate has an important effect on delays, which was far from a well established fact. This has been incorporated into the final layout to ensure differential signals retain the correct relative phase.

Initial results on the MP7 which was delivered are promising, although an assembly error (two power modules swapped) has not allowed the full gamut of tests to be completed yet. For flexibility during testing the card was designed so that power units could be isolated. We have taken advantage of this capability in order to get testing immediately underway and to verify that core card functionality is present; the CPLD and FPGA can all be programmed; board communication via GbE is operating and a subset of the 48 optical links have been operated successfully at 10Gb/s. This has allowed us to place an order for assembly of a second card.

A small delamination in the PCB appeared after the second reflow during assembly (yellow arrows on the picture). The most likely cause is moisture absorption, but the only way to confirm this is to cut up the PCB (other causes are resin starvation or contamination). This will eventually be done, but at present the card is far more valuable to us in one piece, particularly as the delamination is very small and does not seem to be causing any adverse effects. The recommendation from the PCB manufacturer has been to modify the PCB pre-assembly bake time and to take additional safeguards during PCB manufacture due to the hygroscopic nature of the material. A simulated assembly of 2 further PCBs (i.e. subjecting the card to the thermal stresses of reflow) and which used the extended bake time showed no delamination.

The review committee also recommended “*the definition of a program of additional [algorithm development] studies using simulated data and collision data (high pile-up runs) to assess the performance of different algorithm options*”. In response the UK groups have recruited effort from existing grant posts and students to pursue a programme of algorithm simulation studies, in the context of the planned Technical Design Report (see below).

Imperial has focussed on developing a jet-finding algorithm which is able to estimate the energy contributed by pile-up interactions, and subtracted this from the trigger objects. Simulation studies show that the resulting trigger performance for jets is independent of the number of pile-up interactions in a bunch crossing. Such an algorithm will be essential in high luminosity running and is expected to give very significant improvements in performance for jet triggers and can also be used to improve the performance of isolated lepton triggers.

The Bristol group is studying the expected rates and thresholds for different trigger objects under different LHC operating conditions. High pile-up data from LHC test runs have been used to explore the performance limits of the current trigger in high luminosity running. These studies provide an important benchmark for evaluating the performance improvements an upgrade will bring as well as being a strong and quantitative motivation for the upgrade.

These two studies will come together, also including work on electron/gamma and muon algorithms by other groups, to benchmark the performance improvements provided by the upgrade based on high pile-up data. Data from high pile-up test runs are also being used to validate the CMS trigger simulation at high pile-up and the mechanism for mixing event data together. These tools will be necessary for future studies of trigger performance.

The first official CMS-wide release of MicroHAL - the robust, high performance and scalable Ethernet Hardware Access Layer (HAL) has been made after a period of extensive beta testing with groups in the UK, CERN and the US.

5.3 Overview of CMS plans

CMS aims to deliver a Technical Design Report (TDR) for the first phase of the L1 trigger upgrade by the end of 2012, allowing a review by the LHCC to begin early in 2013. Preparations for this started with a series of extended meetings in February 2012 to understand the status and requirements of the muon, calorimeter and global trigger sub-components and the central software. Since then there has been a period of discussion, during which plans, cost estimates, etc have evolved. These will be presented to a Trigger Workshop on 12-13 September in preparation for the Conceptual Design Review, which will be composed of two parts (i) a review of strategy and timeline in late September or early October and (ii) a review of the project plan at the end of November.

Defining physics requirements for the upgrade is one of the most urgent topics, which until now has had limited support because physics teams have been focused on analysis of current CMS data. This is changing as the upgrades gain prominence and the challenges of operating in a high pile-up

environment become apparent. A detailed programme of simulation studies is planned for the TDR in cooperation with the physics analysis groups, in particular Higgs and Supersymmetry.

Commissioning the upgraded trigger must not put efficient data-taking at any risk. To ensure this, it is intended to duplicate incoming data upstream of the trigger electronics to allow the new system to be commissioned in parallel with the existing trigger. Installing the splitting components to achieve this is a key task which must be completed during the LS1 shutdown and is now a critical issue. In particular in the calorimeter trigger system, signals from the ECAL must be converted from electrical to optical and duplicated. To maintain the function of the existing calorimeter trigger system requires a large number of new mezzanine boards to transmit and receive these optical signals. As noted in the Calorimeter Trigger Review report, there are considerable uncertainties remaining in the design, costing and schedule for the production of these items, which need to be addressed.

5.4 Deliverables

The delivery of a prototype final card (M3.3.7) was scheduled for June 2012. This has been delivered, one month late due to misunderstanding between the PCB manufacturers and Imperial over nominal and final board thickness. This required a change in design and the purchase of long lead time material.

Testing of the MP7 (M3.3.8) is scheduled for Sep 2012 and it seems likely that we will meet this objective despite the delay and the assembly error (assuming no more surprises).

For the remaining year, the objective is a pre-production batch of ~12 MP7 processing cards in November. There are currently parts for 13 more cards in hand, excluding RAM, FPGAs, optics & fibre. These are expensive parts and purchase has been delayed (i) until the prototype has been validated (ii) to benefit from price reductions. However, it does lead to schedule risk unless managed very carefully.

5.4 Staff on project

Reported in accompanying tables. Since the beginning of 2012 A. Tapper has taken over as CMS Trigger Upgrade Project Manager, and overall responsibility for WP3. The UK groups have recruited additional effort for algorithm studies through existing grant posts. M. Pioppi (Imperial) and J. Brooke (Bristol) are working with students from Imperial and Bristol on these studies.

5.5 Expenditure

The major items of expenditure are the optical connectors, FPGAs and MP7 card orders, which are not yet invoiced, but the overall spend is still well within the foreseen envelope. The planning for the final year of the project is being studied to ensure sufficient MP7 cards are produced in a timely fashion within the budget.

5.6 Comparisons with CMS activities elsewhere

The MP7 is significantly in advance of other comparable developments. It has a factor 4 greater optical bandwidth (1.44 Tb/s v 370Gb/s) and a factor 2 greater logic capacity than any other previous AMC card in CMS. Furthermore, the optical interface keeps the backplane ports free so that the system can take advantage of all that μ TCA has to offer.

Up till now Wisconsin have done more simulation work than the UK; however this is quickly changing. In the September workshop results will be presented on improvements to the angular and energy resolution of the jet path at L1 and the impact this has on the trigger rates and thresholds. This will be analogous to the electron/gamma trigger work completed by Wisconsin last year.

Studies of the expected performance of the current trigger, in high luminosity running are under way at both Bristol and Ohio (a new member of the TriDAS group). The UK is therefore matching the US in terms of simulation work, but in terms of software (μ HAL), firmware (GCT & TMT demonstrator system) and hardware (MP7) and the ability to deliver a new trigger it is significantly ahead.

The Time Multiplexed Trigger concept remains attractive and, apart from weak latency arguments against it, the main objection from some is that it is a new idea, so might include hidden pitfalls. The

31 August 2012

lab demonstrator should allow us to dispense with such objections as well as implement more algorithms in firmware. So we remain optimistic that the TMT architecture will eventually be the CMS preference, although the UK hardware is agnostic on the architecture choice.

6. Work Package 4: Pixel tracker development

6.1 Objectives

The main objective of this work package, pending approval of follow-on UK funding for the CMS upgrade, is development and commissioning of a DAQ system for the new pixel detector. Along the way to this goal, nearer term objectives have been defined to prepare for prototyping and production of new FEDs, and to establish UK groups within the CMS pixel community. These objectives are

- development of a DAQ test-board for upgrade pixel ROCs (2012)
- DAQ hardware and firmware for high rate beam tests of pixel ROCs (2012)
- provision of a DAQ system for the pixel pilot blade system, based on prototype FEDs in μ TCA form factor, with associated firmware and software (2014)

The latter item is beyond the scope of the present grant, but current activities will ensure the UK is well on track for delivery of the longer term objectives.

6.2 Progress to date

DAQ test-board

Development of a new CMS pixel DAQ test-board is divided into two stages. Stage one uses a new custom board to extend functionality of the existing PSI DAQ test-board to cover upgrade readout chips (ROCs) and token bit manager ASICs (TBMs). Configuration of ROCs, triggers and data handling is controlled by the PSI DAQ test-board. A first prototype of a new test-board is used to receive data from the upgrade ROC and convert it to a format suitable for the PSI DAQ board. All hardware, firmware and software is available and has been tested in a test beam.

Stage two of the test-board development transfers control and readout functionality from the PSI board to the new UK board, which will then be capable of running standalone and assume the role of a firmware development test bed for the Phase 1 FED. This requires additional hardware (mostly programmable power supplies to operate ROCs), which has been available since July. The majority of the work required for standalone operation of this board is the firmware needed to replace relevant functionality of the PSI test-board. The required firmware modules and their status are listed in Table 1. The majority of the work is done, and we expect fully functional standalone UK hardware, firmware and software in September 2012.

Decoding of ROC data	Done
Memory management (data buffering)	Done
Generic IPbus interface	Done
Readout with IPbus protocol over Gbit Ethernet	Done
ROC configuration with fast I2C	Partly implemented
ROC/TBM control (reset, trigger, calibration)	Done
Trigger processing (adjustable latency, prescales, CMS trigger rules)	Done
Control and decoding of ROC power supplies with I2C	Not yet implemented
Decoding of dual channel TBM 4/5 bit NRZI encoded data	Not yet implemented

Table 6.1: Status of firmware blocks required for standalone operation of the UK pixel DAQ test-board. Testing of all blocks is ongoing.

Basic software, consisting of a C++ application for data acquisition, and a graphical user interface for control of the test-board, is available. It was put together with minimal effort, thanks to software libraries and APIs provided as WP1 deliverables. Two main software components are missing. One is low level functions to configure ROCs and the test-board itself, which are trivially implemented based on the IPbus APIs. The other missing component is ROC calibration algorithms, so far only available for the PSI test-board. This will be provided CERN and

National Taiwan University. All missing software components are expected to be available for the October beam test.

High rate beam test

A first beam test to measure inefficiencies of the new ROC took place 30 July - 3 August. UK groups were responsible for hardware and associated firmware, comprising the prototype UK test-board (operated with the PSI test-board for ROC control) and a simple beam telescope to carry eight ROCs. Other participating groups (CERN, National Taiwan University, Karlsruhe Institute of Technology) were responsible for mechanical support and alignment, beam area infrastructure, trigger scintillators, trigger processing hardware and firmware, data quality monitoring, reconstruction and simulation software.

Fig. 1 shows the beam telescope, designed by UK groups specifically for this beam test, which will continue to serve an important role for FED development, being electrically compatible with the new pixel detector module, the first actual prototypes of which are not expected in significant quantities until 2013. Thanks to its modular concept, components (ROCs, TBM) can be tested individually, and their configuration varied within a wide range of parameters. The rest of the DAQ system provided by the UK is not specific to beam tests: the test-board will continue to serve as a firmware development test bed. UK-contributed modifications to the PSI DAQ test-board, for example a USB2 interface improving its readout speed by almost a factor of 10, will generally be useful for the remaining lifetime of this board.

The system was prepared under time pressure, with a full set of digital ROCs for testing the telescope only available 1.5 months before the test. Several unexpected problems further delayed completion. For example, the CERN group failed to provide working firmware to adjust trigger pre-scales, latency and to enforce trigger rules – including an important workaround for a design error in the ROC, requiring a trigger veto during readout of a previous event. This block was instead provided by University of Bristol. Furthermore, this was the first attempt at reading out more than one digital ROC simultaneously, and existing firmware and software from PSI and ETHZ failed to deal with this correctly. Lastly, the PSI and the UK test-board needed synchronized clocks. Initially, the UK board was programmed to run on the PSI-provided clock. However, it turned out to be of insufficient quality, which forced us to revise the clock distribution to use a signal from the UK board at the last minute.




Figure 6.1: The RAL beam telescope used in the 2012 high rate beam test. Bump-bonded assemblies of ROCs with sensors were produced by PSI and wire-bonded to carrier PCBs (yellow/golden). One PCB per ROC (green PCBs with square outline) connects control signals and readout to a bus.

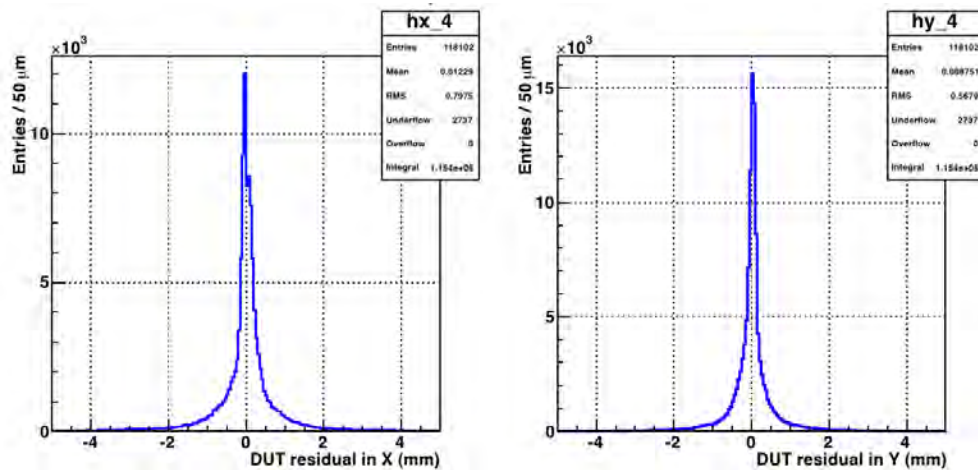


Figure 6.2: Tracks were reconstructed from telescope data using seven out of eight ROCs. The residuals of hits on the eighth ROC with respect to the interpolated trajectory are shown on the left (horizontal) and right (vertical) plot. Results obtained by R-S Lu from National Taiwan University.

Despite the problems, we did provide a working DAQ system. Analysis of data is on-going; fig. 2 shows preliminary results. Some headers in the raw data have specific errors, also seen by PSI using different hardware, and might indicate a ROC problem. Also, some 10% of pixel addresses reported by ROCs are invalid in ways that do not seem easily explained by bit errors. The rate of pixel decoding errors is similar to that in X-ray tests at PSI and ETHZ, again making ROC issues a more likely explanation than DAQ problems. The main objective of the beam test, measurement of pixel inefficiencies at high trigger rates, requires further analysis and should be completed in September. However, final results are expected only after additional measurements from a second beam test scheduled for late October 2012.

The telescope hardware will be reused without modification in October and will be run by a standalone UK test-board. This removes the complexity and problems associated with the PSI board, and it pushes the UK groups further along the road towards a FED prototype.

FED development for pilot blade test and CMS pixel detector upgrade

This part of the project is mostly foreseen for the next upgrade grant period. Nevertheless, the firmware and software development associated with the test-board and beam test activities is directly applicable to the initial phase of the proposed project. In addition to this development effort, progress was made on planning and definition of the μ TCA FED project.

The UK proposal has been explained in pixel meetings, allowing collaborating institutes to adjust their plans to the proposal. The need for full scale system testing at the CERN Tracker Integration Facility (TIF), emphasized by the UK, is now widely recognized. Discussions with potential impact on FED design routinely involve us. For example, there is an ongoing discussion about the line encoding scheme used to transmit optically pixel data from TBMs to FEDs.

Work has begun at Imperial to evaluate how to modify the basic MP7 layout to accomplish the pixel FED functionality. It is not yet clear if the layout is best done starting from scratch or by modifying the MP7. In either case, it will profit from much experience already gained. It is likely to require a few months of effort but better estimates require decisions and a little more time.

6.3 Deliverables

The development of a new test-board for digital ROCs and TBMs is about to be completed. Test beam DAQ electronics has been made available as planned. Formal deliverables for the μ TCA FED development are part of the next funding period, but preliminary work on some aspects relevant to the first year of the FED project is starting. This covers many firmware modules and some software needed to operate the first prototype FEDs.

The original proposal for pixel DAQ work during the current grant period, dated 20 July 2011 and submitted to the Oversight Committee, defined a work breakdown structure for three separate parts of the project. The status of those tasks is:

6.3.1 Development of DAQ test-boards for PSI46dig

Essentially complete. The last item (schematics, layout, assembly instructions, firmware, software and documentation for production of new test-boards) is on hold pending further discussion about whether the board will be used for FED development only, or if it will be needed for all module production and testing centres.

6.3.2 Test beam DAQ

All tasks have been completed, with the exception of building and testing the adapter board with the digital TBM and evaluating performance of PSI46dig using test beam results. Both the ROC and TBM were delivered late. Participation in test beam data taking and data analysis was undertaken in July and there will be a second beam test in October.

6.3.3 DAQ Integration for pixel pilot modules

This was replaced by the UK proposal to build new FEDs for the pixel upgrade, which contains the original pilot DAQ integration proposal as part of its activities.

6.4 Staff on project

Reported in accompanying tables. Staff levels are unchanged with respect to the previous report.

6.5 Expenditure

The main expenditure is RAL PPD staff costs. PCB production and parts needed for the beam test were paid for by CERN, University of Kansas, and some contributions from general RAL PPD operating budgets.

6.6 Comparisons with CMS activities elsewhere

Two near-term test-board options are available for readout of digital ROCs. Temporarily, the existing PSI board for analogue ROCs can be used, with only minor modifications to hardware and firmware. The crucial modification is deserialization of high speed digital data. PSI use a small add-on board to feed digital data into general purpose I/O ports of the FPGA on the main board; the deserializer is then implemented in firmware. This is the simplest solution, but is known to have signal quality limitations. The second option uses the same PSI board coupled to the UK prototype, and deserializes in an external FPGA; this is a little more complex but allows standalone use of the external deserializer as a separate test-board.

In early 2013 the first dual-core TBM chip is expected, which more than doubles the input data rate. The PSI modification will not cope with this but the second option will. However, the hybrid UK-PSI solution is unlikely to be needed, because the UK board will be capable of standalone operation, making the system much simpler.

For the longer term, for example for test setups during module production, two options are available. The UK board could take this role, as explained above. The PSI group developed a digital successor to their previous test-board. Both boards are justified: the UK board serves as a FED development platform until μ TCA prototype boards are available, and possibly beyond due to lower cost. Emphasis of the UK design is thus on very high data throughput, among other priorities. The new PSI board is mostly for special ROC tests, including wafer testing. Both boards could become the standard test-board. There has not been a decision but the barrel upgrade project manager has favoured the UK solution.

Until recently, the situation was similar for FEDs. The UK μ TCA FED was always seen as the solution for the upgraded detector while modification of the existing VME FED, was seen as a fall-

31 August 2012

back. Detailed planning of the μ TCA FED has convinced CMS that it will be available in time, with fully functional prototypes ready even for the pilot blade system.

7 Risk register

The current version of the project risk register is v6.0, revised from the last version. Although it is late in the project, risks 19, 25 and 37 have been updated with comments. These concern the CBC2 submission and MP7 progress. Sufficient funds remain to alleviate or solve any significant problem. Risk 33 was retired and others might possibly be too, in view of the time remaining.

8 Finances

The financial report is summarised in the attached tables.

The SSC Oracle data base no longer reports PPD staff costs. These have been provided to us separately but we were advised not to include overheads. The costs since April 2012 have therefore been inserted in an additional line. The total PPD staff costs have exceeded the originally allocated budget.

The additional allocations to Bristol and Brunel are included as extra lines as contingency.

Upgrade travel spend continues to increase, with beam tests and more frequent meetings in CERN.

Equipment spending has recently been substantial, although later than expected for reasons explained. The cost sharing of the CBC-2 submission was more favourable than expected mainly because of a RAL XFEL development which shared the wafer area.

Engineering staff spending at RAL has also continued to be high in view of the activities on the CBC-2 design and in support of WP4, which is using effort originally planned for WP2 FED development.

The Working Allowance currently stands at £120k. It has been reduced to accommodate the expected use of TD staff in the remainder of the project. It is expected the remainder will be called on for CBC2 and MP7 production, especially in case of the MP7 production planned. We will not know the success of the CBC2 until October or later.

There is an additional £20-30k of expenditure incurred in CERN, temporarily charged to the M&O CMS account, which has yet to be repaid to the SSC system from Imperial College.

9 Gantt chart

The Gantt chart has been updated only for the overall LHC schedule with new LS1 dates, but the LS2 date is left unchanged in the absence of hard information. There are no significant other changes, except small delays in a few milestones for WP2 and WP3 which were mentioned above. We refer to the PPRP Gantt charts for our best estimates for the future.

10 Milestones

The milestone table has been updated as agreed at the last OSC, i.e. highlighting recently completed and imminent milestones in the main table. However, it was not agreed to identify the fractional completion of milestones and, in fact, this is impossible. It is not feasible to accurately estimate the fractional completion of software, firmware or board design and, from experience, such estimates, e.g. provided by a responsible engineer, are often notoriously inaccurate. For example, firmware or software completion cannot be measured by the number of lines written but by whether or how well it works. Similarly for layout of an electronic board or ASIC layout. If further task breakdown is required, then suitable intermediate milestones are needed, but we believe those within the project have been sufficient, as demonstrated by the fact that most were completed in a timely manner.

Appendix: PPRP proposal

Included as separate document