University of Bristol
Brunel University
Imperial College London
Rutherford Appleton Laboratory

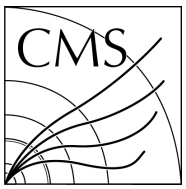

R&D in preparation for upgrades of the CMS detector for High Luminosity LHC

1 Executive Summary

The R&D project is now in its final few months. There has been further good progress since the last report. As an R&D project, evolving over time to adapt to a better understanding of the CMS requirements as these also evolved, a comparison between the proposal and the present status is not completely straightforward but the project has met, and exceeded, some of the major objectives, such the Tracker ASIC and the trigger demonstrator. The expected contributions from simulation studies to CMS TDRs have been made, while the TDRs are not identical to those envisaged in 2007 when only one phase of a machine upgrade to high luminosity was foreseen. Goals have been achieved which were not formally anticipated at the outset, such as the implementation by us and adoption by CMS of IPbus, the innovation of the Time Multiplexed Trigger concept and the inclusion of trigger data in the second successful version of the Tracker CBC ASIC (when only one full size chip iteration was originally envisaged) which has been designed for a very new assembly technique.

The LHC ran very successfully with protons in 2012, finally accumulating 23.3 fb^{-1} of integrated p-p luminosity at 8 TeV centre of mass energy, with 21.8 fb^{-1} recorded by CMS, and further improvements in data taking efficiency. The first LHC Technical Stop, LS1, has just begun following heavy ion running until mid-February. LHC operation should begin again around March 2015. CMS continues to produce physics results and has been refining the results on the newly discovered boson with greater statistics, as well as steadily publishing results on a wide variety of other topics.

Scrubbing runs and LHC machine developments with 25 ns bunch spacing in December gave very encouraging results and it seems most likely that this will be the baseline operating mode, as hoped, from 2015.

There were valuable physics simulation contributions to the Pixel TDR and a UK RA was recognised by a CMS annual achievement award. There has also been further progress with calorimeter trigger simulations focussing on L1 performance and trigger menus.

IPbus is now well established in CMS but, as a result, there is greater interest in influencing the software and details of functionality, which had an impact on one milestone.

The CBC2 took much longer in manufacture than anticipated but wafers arrived in January and testing began using cut die in early February. The early results are extremely positive. The module developments have continued with CERN and it should be possible to assemble commercially the first, simple, dual-CBC hybrids beginning in March, once wafer probing, presently underway, has selected good die.

The Time Multiplexed Architecture was adopted in December as the baseline for the main upgrade of the calorimeter trigger. It will be implemented starting in 2015 with commissioning and operation of the new trigger in parallel with the existing trigger. There will be an intermediate upgrade of the present trigger to ensure sufficient performance margins in 2015, in case of very high pileup conditions.

Extensive tests of the MP7 trigger board have been very successful; it is significantly in advance of other comparable developments anywhere and is operating in demonstrator systems, with additional boards being assembled.

The second 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.

The PPRP proposal for construction and continued R&D funding for six years to March 2019 has been reviewed and recommended for approval. Following the final recommendation from Science Board in early January, a revised work plan and financial plan have been submitted and should lead to the next grant commencing on 1 April.

A no-cost extension to the existing grants was requested from STFC, and has been agreed. This was largely motivated by the slower spending towards the project end, in view of the late CBC2 delivery and uncertainties about final FPGA availability for the MP7. It also helps with a smooth

transition into the new upgrade project, especially with existing staff, and avoids the need for bridging funds.

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} $cm^{-2}s^{-1}$ to 10^{35} $cm^{-2}s^{-1}$ about ten years after start-up. The accelerator upgrade has is now proposed to take place in two main stages, with an increase in luminosity to ~2.5 10^{34} $cm^{-2}s^{-1}$ being Phase I, up to 2022 in the most recent CERN ten-year plan, and a further major increase in the next decade, probably to ~5 x 10^{34} $cm^{-2}s^{-1}$ levelled luminosity, Phase II.

The experiment and accelerator continued to operate extremely well during 2012. The instantaneous luminosity record is 7.7 10^{34} $cm^{-2}s^{-1}$ at 50 ns bunch spacing, corresponding in pileup terms to over 1.5 10^{34} $cm^{-2}s^{-1}$ if the machine were running at 25 ns. This is at half the design energy and machine experts expect future operation at 13-14 TeV with up to ~2 10^{34} $cm^{-2}s^{-1}$ before LS2.

From the physics perspective, CMS continues to produce publications at a high rate and has been consolidating results on the Higgs-like boson. No further surprises are yet evident but there is a lot of unanalysed data "parked" for processing during LS1 and the refinement of results on current analyses will continue for many months exploiting the full data sets from 2012.

2.1 LHC upgrade schedule and planning

There were no further changes in the machine operations plan in 2012 and the LHC continued to run with proton-proton collisions to the end of the year. Machine tests with 25 ns spacing were postponed to December but appear to have been quite successful. A report was given to the Evian LHC meeting which was summarised for CMS:

- A scrubbing run lowered the Secondary Electron Yield in the arcs to less than 1.45, resulting in reduced heat load as well as improved beam quality (lifetime, emittances),
- In spite of the high heat loads, close to the cryogenics limit, further scrubbing and ramps to 4 TeV did not exhibit any clear improvement in conditioning of the arcs,
- Emittance blow-up ascribable to electron clouds is still observed at injection energy with long and closely spaced trains,
- At 4TeV, there was no indication of further emittance degradation caused by electron clouds.

This led to the conclusion that a concentrated scrubbing run will be probably be insufficient to fully suppress the electron clouds from the arcs for 25 ns beams in future operation. However, what was more notable was the much increased confidence in the ability to operate readily with 25 ns and achieve luminosity objectives. In three physics fills with 25 ns bunch trains, very low emittances were achieved, from 1.85μm to 2.03μm (half of the nominal emittance for 25ns beams) and an observed rate of UFOs (Unidentified Falling Objects, constituted of tiny material fragments, ~μm in size) which, while 5-10 times higher than at 50 ns, did not generate beam dumps by causing losses sufficient to trigger beam loss monitors. A peak pileup of 12 events/crossing at the beginning of the fill, with β^*=1m, was reached.

The first Long Shutdown (LS1) may delay LS2 but the long term CERN operations plan has not recently been updated; informal feedback suggests LS2 coming one year later, in 2019. LS3 will not be before 2022, but also seems likely to be later. Its duration is not known but it seems likely to be at least two years in view of provisional experiment work plans and possibly longer.

2.2 CMS planning

The CMS Technical Design Report process for sub-detector upgrades is almost complete. The trigger TDR is in its final draft and will be submitted for the March LHCC meeting. The HCAL and pixel TDRs were approved by the LHCC in September. Longer term Phase II Tracker planning has not changed recently. Much of the recent focus has been on financial matters with the cost-books for each upgrade activity being thoroughly internally scrutinised for accuracy, consistency, adherence to agreed standards and adequacy, for example spares and other provision. The objective has been to be able to present final expected costs to funding agencies at the RRB with a high degree of confidence that the

budgets are realistic and will not be exceeded. The process is complete for muons, HCAL and pixels, and underway for the Trigger.

The cost sharing at the sub-detector level is straightforward, since agencies commit to their contributions following approval in their own countries, and this is part of the TDR process. CMS agreed a common fund, representing about 10% of the total Phase I upgrade cost; 74% of this is either paid or committed to be paid, representing 79% of the value of the common fund. Remaining common items are partly funded, in some cases using in-kind or special contributions; discussions are ongoing to enable eventual full funding.

Attention has been shifting to Phase II upgrades, both in developing the physics motivation in more detail and a financial plan, including the scope of the upgrades. There are good physics arguments for the increase in statistics using the High Luminosity LHC but these need to be made more quantitative through more simulation studies and this activity will ramp up considerably during 2013. The main physics goals are clear:

- extend tests of the Standard Model, in particular measuring properties of the newly discovered particle and longitudinal components of the massive vector bosons with the highest possible precision, aiming to establish whether there are deviations from Standard Model predictions;
- confirm whether the Higgs is accompanied by other new particles at the TeV scale; not only additional resonances that might be evidence for an extended Higgs sector, but also other particles that may play a role in the global picture of electroweak symmetry breaking or in the solution of the dark matter puzzle.

Further discoveries after 2015 will certainly enhance the case but the likely precision on Higgs couplings which can eventually be expected with 3000 fb^{-1} looks encouraging, potentially 1-4%, and there is a wide range of further studies, including SUSY, heavy vector boson resonances, top studies and much more.

Technically, the challenges which are emerging in addition to the well known problems of a more granular tracker operating in an even harsher radiation environment, include longevity, in view of the expectation of operating the LHC beyond 2030, and maintaining trigger rates with low enough thresholds. A track-trigger is part of CMS strategy, and our progress with outer tracker module developments supports the realism of this approach. However, it is not yet clear that this can achieve the reduction factors necessary to preserve the 100 kHz L1 rate or if the advanced pixel or strixel detector modules needed can be delivered in time.

The alternative is to consider altering the trigger specifications and recent proposals for CMS to have the capability to operate at up to 1 MHz rate and ~20μs latency, instead of 100 kHz and 6.4μs, have been made. The major constraint on these changes is the implication for the ECAL; changes to the front end electronics (developed by the UK) would be needed to extend the latency. This would involve accessing ECAL supermodules which are buried under layers of tracker services on the CMS cryostat. The work could be completed in a shutdown lasting more than about 26 months and the electronic changes are technically feasible, although no work has been undertaken on circuit development.

More problematic is the rate increase. This also seems to be principally limited by the ECAL, by the volume of data required to be transferred to the trigger for the L1 decision. At present 10 samples per trigger channel are sent using 800 Mbps fibre optic data transmission, using variants of tracker optical links. Without replacing optical transmitters, this sets a hard limit of about 150 kHz unless the number of samples is reduced; a factor of two would increase the achievable rate to ~300 kHz. The viability of this, and the cost and resource implications, are under discussion.

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

2.3 UK adaptation to CMS planning

The next phase of the upgrade activities is now clearer following review of the PPRP proposal and the recommendation following the Science Board meeting in December. A revised work plan and financial plan has been submitted based on the recommendations and provisional budget allocations.

This foresees work continuing on the two major activities in WP2 and WP3 as requested, but concluded that WP4 (pixel DAQ) was a lower priority and that we should withdraw from this activity. The OSC undoubtedly understands the financial background and the influence this, and future STFC planning, has had on the decision.

This news arrived at the start of 2013 and we have been winding down the activities on WP4 since then, which mainly affects RAL PPD; most of the effort will transfer to WP3, which strengthens the activities, both short and long term, on the delivery of the Phase I calorimeter trigger. The FC7 board which was foreseen for the pixel FED is continuing as it will be required for the CBC module readout. We are sharing some responsibilities with CERN who foresee its use in the future TTC system and for other users, and who will support that.

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.

Pixel system: The CMS Pixel Upgrade TDR was presented to the September meeting of the LHCC, and contained numerous results either produced by UK collaborators, or based upon UK-developed software. In particular, the principal analysis effort used to illustrate the physics impact of the proposed upgrade (Z+H production channel, with the Z decaying to muons and the Higgs to b-quarks) was led by M. Grimes (Bristol). Substantial improvements in physics event yield were demonstrated. The work in this area was recognised by CMS through the award of the annual achievement award for tracking upgrades to Mark.

In order to achieve the substantial improvements in simulation performance required to support full simulation and optimisation of the pixel detector, substantial changes have been made to the 'core' of the simulation package, including the mechanism for the overlay of individual minimum-bias events to represent the overall incoherent background underlying all signal events. This in turn necessitated changes to the pixel tracking software, which were made principally by the UK expert team: I. Reid (Brunel) and M. Grimes. As a result of this and other changes, very large performance increases have been achieved for the new 4-layer detector in both simulation and track reconstruction – an important area of work, since pixel reconstruction is used at an early stage of the CMS High Level Trigger System.

The prototype pixel readout ASIC for the upgrade was taken to test beam during the reporting period, with the UK taking lead responsibility for both the readout hardware and the associated firmware and readout software. This activity is reported under WP4, but was essentially a joint WP4-WP1 effort, building upon the IPbus control system developed earlier in the project.

We recently learnt that the pixel upgrade activity will be discontinued from April 2013. Efforts therefore now focus on finishing with current commitments, and on the handover of responsibility to other non-UK collaborators. The current software effort will be retained in the next project phase, and deployed in support of tracker and track-trigger software tools.

L1 trigger system: Detailed simulation work towards the L1 Trigger Upgrade TDR was a major focus of WP1 during the reporting period. These studies were require to show both the physics performance improvements achievable with the new L1 hardware (in addition to the technical advantages of the upgrade), and to fully document the likely impact of *not* carrying out the planned upgrades. In the UK, these studies were organised in conjunction with WP3, with J. Brooke and D. Newbold (Bristol) heavily involved in both planning and execution. The simulation programme

comprised three stages: (a) simulation of the baseline algorithms proposed for the L1 upgrade (which comprises multiple stages of hardware replacement), resulting in trigger object-level performance estimates; (b) development of exemplar trigger menus, demonstrating a range of strategies for balancing the allocation of a fixed total trigger bandwidth between physics channels, and (c) physics studies showing the impact of alterations in trigger efficiencies and rates on event yields. The UK took responsibility for the central part of this programme, the menu development, in addition to contributing substantially to object-level studies.

In November 2013, CMS decided to institute a L1 Performance Task Force, with responsibility for coordinating all studies towards the L1 TDR. J. Brooke was asked to serve as co-leader of this effort, and has subsequently taken responsibility for delivery of all results for the TDR. An example trigger menu, showing the performance in terms of object thresholds of running the existing trigger algorithms at 14TeV and <n>=50 pileup levels, is shown in the figure below.

Algorithm	8 TeV 7E33 ~25 PU		14 TeV 2E34 50 PU	
	Thresh (GeV)	Rate (kHz)	Thresh (GeV)	Rate (kHz)
Single EG	22	10	46	10
Single IsoEG	18	9	31	9
DoubleEG	13, 7	9	22, 12	9
Single Muon	16	9	50	9
Dble Muon	10, open	5	35, open	5
EG+Mu	12, 3.5	3	21, 6	3
Mu+EG	12, 7	2	25, 15	2
SingleJet	128	2	188	2
DoubleJet	56	10	132	10
QuadJet	36	2	96	10
Double Tau	44	2	56	2
MET	36	7	84	7
HTT	150	2	511	2

Fig. 1.1 An example trigger menu, showing the performance in terms of object thresholds of running the existing trigger algorithms at 14TeV and <n>=50 pileup levels.

Online software: Work towards a new release of the IPbus support software was largely concluded during the reporting period, though the milestone of full integration with the CMS online software framework was not met. This was due to a decision by CMS to carry out a review of the software system, which is now intended to support not just the trigger, but also the upgraded HCAL and pixel readout hardware. This review resulted in some changes in functionality, but has also led to the interest and involvement of a wider community within CMS, to the benefit of the project. The associated milestones have now been replanned, with the aim of having a complete large-scale demonstrator in place in the CMS electronics integration centre in Q2/13, operating the general CMS µTCA test stand. The technical aspects of the software are essentially complete, with the remaining functionality (reliable operation over unreliable networks, higher performance) put in place as the last

contribution of R. Frazier (Bristol) who has now left the project, in conjunction with A. Rose (Imperial) and D. Sankey (RAL).

3.3 Deliverables and Milestones

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

- M1.3.8 has been achieved, as an integral part of the preparations for the L1 TDR. Detailed work on menu development and physics impact are still under way, but the comparison with algorithms is complete.
- M1.3.10 has been achieved, as part of the overall Pixels TDR effort.
- M1.4.1 and M1.4.2 experienced a delay, as described above, due to a change in CMS planning, and also delays in the setup of the μTCA test stand in the electronics integration centre. The UK deliverables towards this effort are essentially complete.
- M1.4.3 and M1.4.4 concerned the publication of Pixels and L1 TDRs. The former has been achieved with the presentation of the Pixel TDR to the September LHCC. The L1 TDR will be previewed in draft form to the March 2013 LHCC, and considered in full at the September meeting.

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

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.

The DAQ component of this was work package was integrated into the pixel FED activity. The 2S-pT module concept has been developed to provide high-pT stub information to a level 1 trigger decision. The final WP2 deliverable of this phase of the UK CMS Upgrade programme, the CBC2 chip, prototypes this stub-finding functionality in a bump-bondable chip.

4.2 Progress to date

The 130 nm CMOS CBC prototype 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 conferences, published, and reported extensively at CMS meetings[1]. During the last six months a further test beam study using the CBC prototype was carried out successfully using a 5 cm long, 80 µm pitch, microstrip sensor that more closely matches the geometry proposed for the Phase II outer tracker 2S-pT modules. Test beam results from this prototype CBC/sensor system will be directly comparable directly with results from the next CBC iteration, CBC2, which will also be used to read out the same sensors.

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.

CBC2 wafers were expected back in October, but production delays resulted in wafers leaving the foundry only just before Christmas, being finally received in the UK mid-January. The submission was shared with other projects which did not require C4 bump-bond processing and these wafers were the first to arrive (eight C4 wafers arrived a few days later). One of the wire-bond wafers was diced and the CBC2 chips separated out. It is possible to verify functionality of these chips without bumping, because a row of wire-bondable pads on the readout edge of the chip is provided to allow testing of chips on the wafer with a probe card. Figure 4.1 shows one of these chips wire-bonded to a dedicated test setup.

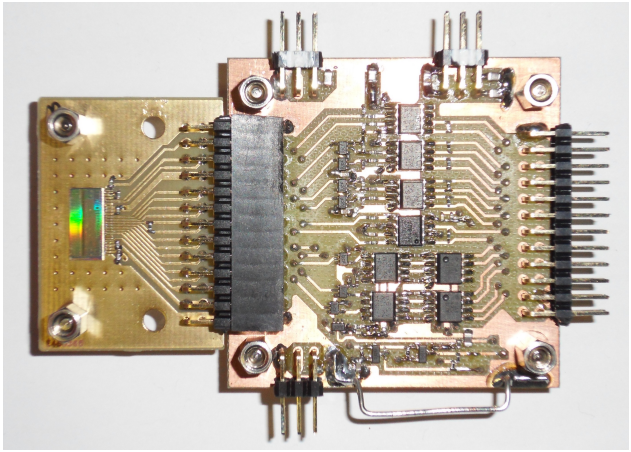

Figure 4.1. Wire-bonded CBC2 chip in dedicated test setup

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

Figure 4.2 shows one of the first results from the wire-bond CBC2 test setup. The output data frame shows that the new internal test pulse circuitry is functioning correctly, and the trigger output shows that the stub-finding correlation circuitry is also operational. Detailed studies of performance will take time, and true performance can only be measured after bump-bonding to a suitable carrier, but it is already fairly certain that the chip is working sufficiently well to allow us to progress.

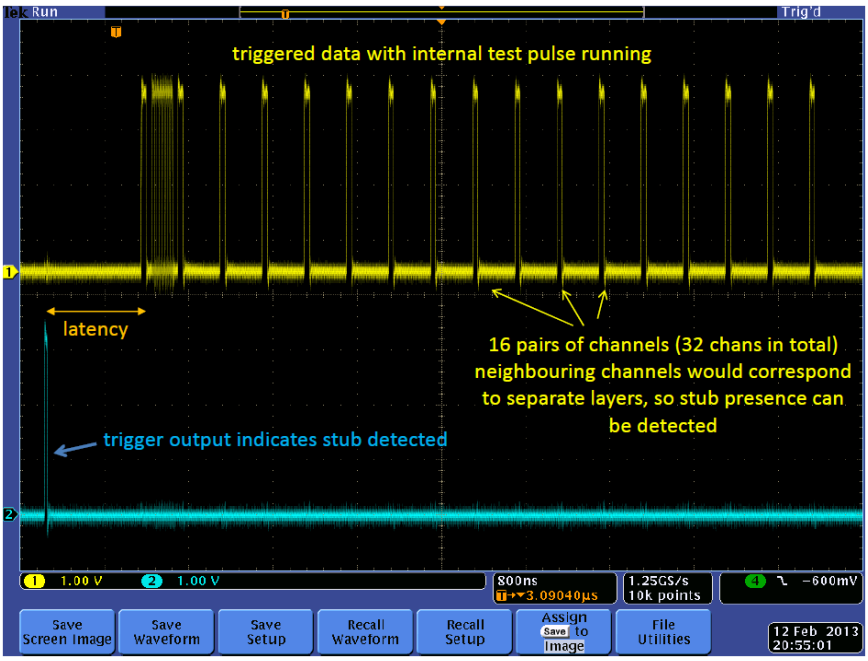

Figure 4.2. CBC2 first result

A number of test set-ups have been prepared. The wire-bond set-up in figure 4.1 not only allows functionality to be verified, but also provides an ideal environment for the development of the wafer test procedures required to verify chip operation before wafers are diced, since in the probing environment the wire-bond connections are replaced by probes, and the same interface hardware and software can be utilised. A probe card has been procured, and in figure 4.3 it can be seen equipped with the same interface card visible in figure 4.1.

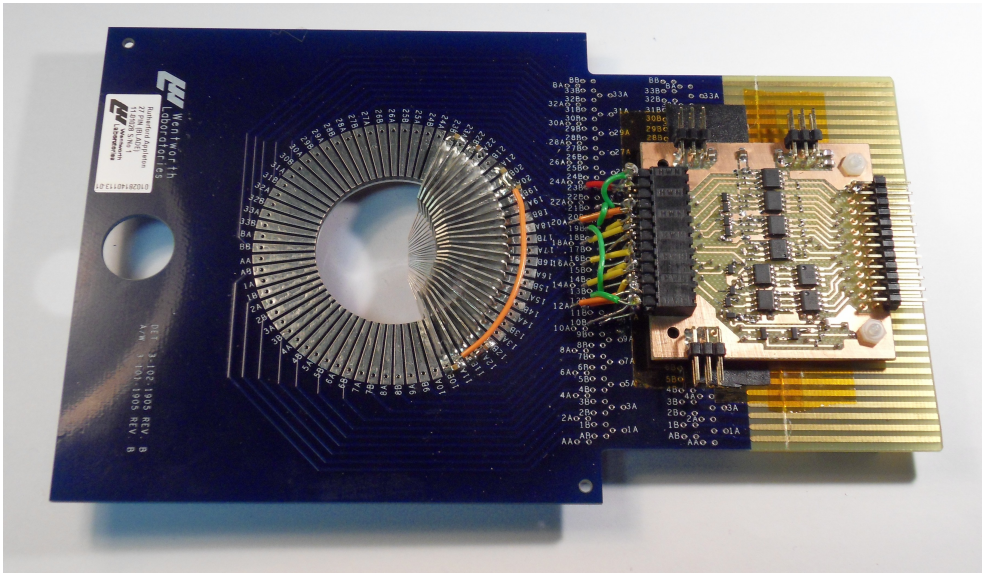

Figure 4.3. CBC2 probe-card

CERN has taken primary responsibility for the design of bump-bondable chip carrier substrates. In the first instance we will use a substrate (dual-CBC2) that takes two chips, allowing to verify the inter-chip communication interface where neighbouring sensor strips are bonded to different CBC2 chips. The dual-CBC2 substrate assemblies will be the main vehicle for evaluating CBC2 performance in the medium term, and will be used for electrical characterisation as well as with sensors.

The dual-CBC2 substrates have been produced. The original plan was to populate them with untested die in the first instance, but because of the novelty of using this type of substrate (in the HEP community) this has been revised 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). We therefore plan to quickly probe one C4 wafer to provide known good chips for the dual-CBC2 substrate assembly. This work is underway and we expect to deliver chips for dual-CBC2 substrate assembly in March.

In the UK we have developed an interface board for the dual-CBC2 substrates, which translates CMOS level signals to standard differential levels which can be transmitted over longer distances, and also an FMC (FPGA mezzanine card) which can interface the same signals to an FMC carrier platform (e.g. the CERN GLIB card).

A working CBC2 allows us to construct mini 2S-pT module prototypes using the dual-CBC2 substrate and the same sensors used with the CBC prototype in the recent test beam. This will become available around the middle of the year. A full-size module using 16 CBC2 chips will follow. The next version of the chip (CBC3) must contain all features required to implement the final outer tracker system, particularly a high speed interface to transfer stub information. Definition of the next version of the chip continues in regular systems meetings held at CERN approximately every 6 weeks. The design phase of CBC3 will begin later this year in the next phase of the UK CMS Upgrade programme.

4.3 Deliverables

The only remaining milestone for WP2 is **M2.4.1** (09/12) *Documented results of final prototype.* This milestone has been affected by the delays in CBC2 wafer production. CBC2 functionality has now been verified, but deeper understanding of performance will still take a few more months to establish and document.

4.4 Staff on project

Listed in the tables. No significant changes.

4.5 Expenditure

Updating of VME bench-top and probe-station DAQ systems (new PCs). Procurement of PCBs and associated components for CBC2 test systems. CBC2 wafer probe cards.

4.6 Comparisons with CMS activities elsewhere

The CBC is still the only Phase II ASIC which has actually been built. Other ASICs remain at the conceptual level, with VHDL simulations of some components.

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 Time Multiplexed Trigger TMT architecture proposed by the UK has been adopted by CMS. This is a significant achievement and was only possible because of the support from STFC over the last three years. This allowed prototype systems to be built to validate the new triggering concept and enabled a significant portion of the core firmware to be written.

The final trigger processing card - the MP7 - returned from manufacture in September. Since then the card has been extensively tested; in particular, all 48 links have been operated simultaneously at 10 Gb/s during which 10^{17} bits were transmitted without error.

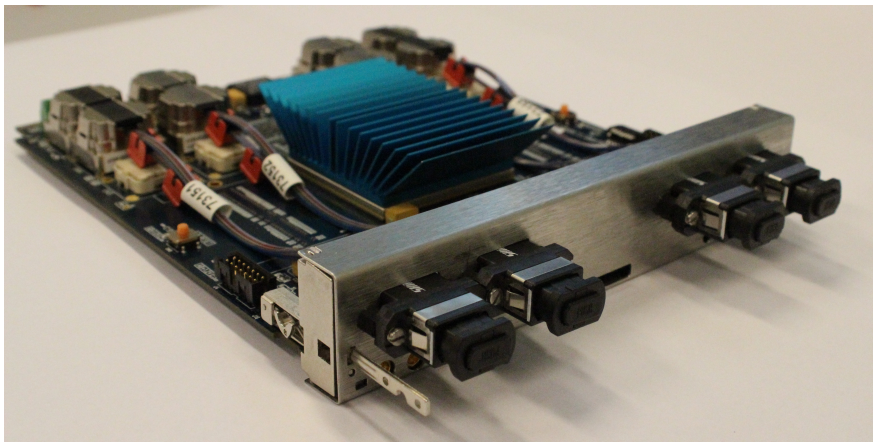

Fig. 5.1. The MP7 Trigger Processing Card with 48 link FPGA and optics fitted and the original heat sink.

The successes have not been without problems. Most of these are now well understood and either fixed in firmware or with modifications implemented for the pre-production card. There are two critical aspects that are still outstanding and there is continued effort to solve both these issues.

The first concerns the optical fibres external to the card. Normally these fibres would contribute a negligible amount of jitter to the high speed serial signal; however, when we moved to a supplier capable of providing the quantities of fibre required for the final system there was a degradation in performance. While the links still operate at 10 Gb/s there is a marked reduction in the operating margin. We are working closely with USCONEC, the developers of the MTP connector to understand the cause of the problem. The most recent analysis by USCONEC indicates that the MTP connector end face was somehow damaged, although the cause has yet to be established. A report is due in the next few days. USCONEC will re-terminate the fibres and we will test again at CERN to verify the performance of the MPT48 to LC fibre harness . When the report conclusions are known we will purchase additional fibre from the new supplier and also seek out a third company. Two potential companies have been identified and quotes are being obtained. We have also asked optical experts for advice, but at present they are as surprised as we are; however, few are working at 10 Gb/s and none with such high density systems.

The second issue concerns one of the transceiver power supplies. While the power supply is adequate for the 48 link FPGA it may no longer be so for the 72 link FPGA. This is due to a recent revision of the Xilinx Power Estimator (XPE 14.2), which increased the MGTAVcc supply from 7.6A

to 10.5A. The power supplies are currently rated to 12A. The MGTAVtt power supply has also been increased which increases the total power dissipation by 4W. This has been mitigated with a new heat sink design which has an absolute thermal resistance 40% less than the original design.

There are currently 4 prototype cards available for test and development. In the next week we expect back the first cards fitted with the larger 72 link FPGAs. These will allow us to perform detailed power consumption tests to determine whether the power supplies need to be upgraded. Alternative power supplies have been built, which can be tested, if necessary, by bypassing the on-board power supplies. At present none of the alternative supplies is ideal (i.e. high efficiency and low noise in a simple and compact design) and thus we are continuing to explore the different options.

The pre-production cards have been delayed while we await power supply test results from the 72 link part. The objective is to submit the pre-production design for manufacture at the end of March. The modifications, all minor, to the prototype have already been completed. Parts are currently available or ordered for up to 6 pre-production cards.

Firmware and software

As the hardware verification tests reach a conclusion the focus has shifted to providing the base firmware and software for the algorithms. Firmware developed for the MINI-T5, with a Xilinx Virtex 5 FPGA, is being ported to MP7 Virtex 7 FPGA. In many instances porting the firmware is trivial, but in two cases significant effort has been required:

- A fully working simulation of the low latency link automatic alignment and checking firmware now exists. The firmware includes asynchronous link capability, CRC integrity checks and spy buffers. It has been possible to bypass parts of the transceiver to improve latency. Hardware tests will commence next week.
- The MP7 firmware images will be stored on a μSD card to reduce the time taken to write new images from many minutes to just a few seconds and also to allow many images to be stored on the card rather than just one. This allows CMS far more flexibility because the system can quickly boot from a number of different preloaded images at runtime. The most challenging tasks have already been tested and the performance met (i.e. loading of the FPGA within a few seconds). The full firmware and software package is now being finished.

The base firmware and software are therefore progressing well. It will also shortly incorporate the final version of IPbus (v2.0) that is expected to be released imminently. It improves error handling and allows faster readout.

As the base firmware becomes more mature and extra cards become available a test system of MP7 cards will be created for algorithm testing and development. The extra cards will also allow interface tests to other systems to commence."

Trigger Simulations

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 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 and 14 TeV centre-of-mass energy Monte Carlo events have been used to explore the performance limits of the current trigger in high luminosity running and compare to the expected performance of the upgraded trigger. 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 come together, also including work on electron/gamma and muon algorithms by other groups, to benchmark the expected performance improvements provided by all aspects of the upgraded trigger based on high pile-up data and Monte Carlo events, and form the input to studies of important physics channels, which are documented in the L1 Technical Design Report (see below).

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. In the course of these studies puzzling

results were found and a CMS-wide task force has been assigned to understand the performance of the CMS detector simulation at high pile-up. These tools will be necessary for future studies of trigger performance.

5.3 Overview of CMS plans

A Technical Design Report (TDR) for the first phase of the L1 trigger upgrade was approved by the TriDAS Institution Board on 31 January and is currently undergoing internal CMS review. It will shortly be submitted for review by the LHCC in readiness for the LHCC meeting in mid-March. A review of the Cost Book by the Finance Board will take place this month.

Defining the physics requirements for the upgrade has been one of the most urgent topics. A task force led by Jim Brooke (Bristol) and Brian Lee Winer (Ohio State) has provided detailed, quantitative information on the performance of the upgraded detector (i.e. the thresholds for different trigger objects) and has also coordinated with the physics groups to study the impact of the upgrade on the different physics channels. This has been achieved, in part, by broadening the number of groups in the trigger, which bodes well for the longer term.

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 much 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 schedule for the production of these items despite prototypes now existing.

5.4 Deliverables

The remaining two milestones are:

- M3.3.8, 9/12, Prototype cards tested
- M3.3.9, 3/13, Final pre-production card manufactured and tested

The testing undertaken in M3.3.8 was significantly more than expected due to the issues listed above and also board assembly errors, board delamination, excessive power consumption connected to the MGT firmware (different from revisions to the Xilinx Power Estimator) and the quality of the AMC edge connector. Most of these have now been resolved, but the pre-production batch has been delayed until we have results from the larger 72-link part. This in turn has been delayed due to the long lead-time in procuring the PCB base material (3-4 weeks) coupled with the complete failure of a PCB batch at final inspection that also required the order of extra PCB base material. Managing stock of the PCB base material is complicated by the material shelf life and manufacturing yield.

As a consequence of the challenges faced in milestone M3.3.8 we now expect pre-production manufacture in April with testing in May. Milestone M3.3.9 is therefore expected to be completed in 5/13.

5.4 Staff on project

Reported in accompanying tables. Since the beginning of 2012 A. Tapper has been CMS Trigger Upgrade Project Manager, and has 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) have been 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 remains well ahead of other comparable developments in both CMS and ATLAS. 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. 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. We only await the availability of the final FPGA with all 72 links to demonstrate what should be the final module for the trigger upgrade.

6. Work Package 4: Pixel tracker development

6.1 Objectives

The main objective of this work package, aiming at approval of follow-on funding for the UK CMS upgrade activities, was development and commissioning of a DAQ system for the replacement CMS pixel detector. A series of objectives had been defined to prepare for prototyping and production of new FEDs:

- 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 a µTCA form factor, with associated firmware and software (2014)

In view of the UK decision not to fund the proposed pixel detector DAQ, the long term objectives are no longer relevant. Contributions to the pixel pilot blade system also fall victim to this change of plan. The DAQ test-board and high rate beam test analysis are about to be completed.

6.2 Progress to date

DAQ testboard

The transition from a digital adapter of a PSI analogue pixel test-board to a standalone test-board was completed. A revised version of the UK test-board with DC/DC converters to supply power to pixel readout chips had been available since the previous report. These boards are now fully tested, and all required firmware and software is available.

The assembly of the majority of second revision UK test-boards was not successful at the first attempt. There was a high rate of bad bump bond connections of DC/DC converter modules. This was traced to a problem with the CERN footprint libraries, where pad sizes for the bump bonds were too small. Most boards were recovered by reworking. The board design was improved accordingly, leading to a third and final version about to go to production. This version also implements a few minor improvements resulting from experience using the second revision board in a beam test environment. The changes include providing a dedicated LEMO connector for external trigger input, making the board mechanically more robust, and making the board smaller by removing a few components only required for coupling the board to a PSI Altera test-board, which will not be done again.

The end of UK involvement in the Phase 1 pixel detector led to modifications of our work plan:

- We will not offer the UK test-board as a general test-board to CMS. This would require long-term support of hardware and firmware. CMS has a backup - a new board developed by PSI; the change can be accommodated easily. The only exception is for groups involved in high rate beam tests (CERN, National Taiwan University, Karlsruhe Institute of Technology). Because the next beam test is scheduled to take place in a few months, switching to different hardware and software is not feasible. We will therefore produce about five boards, sufficient for the beam test, but there will not be a full production run.
- With the scope limited to beam test activities, there is no need to provide hardware and firmware to support Token Bit Manager (TBM) chips. This simplifies the firmware significantly and reduces the remaining firmware work almost to zero.
- We will hand over hardware design files, firmware and documentation to the CERN and Taiwan groups for possible future use of the board without UK support. The associated software (calibration code, DAQ, DQM etc) was mostly written by these groups with limited UK contributions, and thus there will be no need for UK software support.

The UK board was fully operational and adequate for the requirements. It was successfully operated in beam tests with particle rates sufficiently high to push the pixel ROC to its limits. Thus, though there is still a small amount of close-out activity ongoing, the project was a success.

High rate ROC beam test

The UK pixel beam telescope with test-board was used in a second CERN high rate beam test in October 2012. The first beam test in July/August 2012 used the UK board as adapter to a PSI analogue test-board. The second beam test was successfully performed with standalone UK hardware. The beam telescope, illustrated in the last report, was re-used without modifications from the first beam test.

While the first beam test DAQ used the USB connection of the PSI analogue test-board for transferring event data to a PC, the standalone UK board used Gbit optical Ethernet with the IPbus protocol. The much higher transfer speed and much improved reliability of the DAQ link allowed us to increase the amount of data recorded during a similar number of days by several orders of magnitude compared to the first beam test. This beam test was the first application of IPbus under fully realistic conditions and made use of the full available IPbus bandwidth. No IPbus problems were experienced, although the memory management required to buffer pixel data before sending them over IPbus turned out to be difficult to implement due to bugs in the Spartan-6 FPGA.

Analysis of the test beam data is performed by CERN/Taiwan/Karlsruhe. The main result is the efficiency to find a pixel hit on one ROC associated with a reconstructed particle track that was found using all other telescope planes, as a function of beam intensity. This efficiency measurement is complicated by a strong dependence of the beam intensity on time within a single spill. Normalization of the beam intensity is done using trigger scintillator rates. Final numbers are not available yet, but it appears clear that the digital pixel ROC performs as expected. A publication is in preparation, with all UK contributors to the project on the author list.

It is therefore fair to conclude that the UK part of the beam test project was completed successfully. All design files and documentation of the beam telescope will be made available to interested groups, especially those still involved in pixel upgrade beam tests.

FED development for pilot blade test and CMS pixel detector upgrade

STFC decided to withdraw from this project. The pilot blade test will be performed with a modified VME pixel FED; the strategy and contributors to the full DAQ upgrade remains under discussion.

6.3 Deliverables

The original work description submitted to the Oversight Committee in July 2011 defined a work breakdown structure for three parts of the project. The status of each of these is given in the following:

Development of DAQ test-boards for PSI46dig

WBS 1.1–3.2 are all completed. The last item, 4.1, “provide schematics, layout, stuffing instructions, firmware, software and documentation for production of new test-boards” has been on hold pending further discussion about whether the UK test-board will be used for FED development only, or whether it will be made available for all CMS module production and testing centres. Following the PPRP proposal decision this item is now obsolete, although we do plan to make documentation available to interested groups such as CERN, Taiwan and Karlsruhe.

Test beam DAQ

All tasks have been completed, with the exception of 1.2 (“build and test adapter board with digital TBM”) which is now obsolete. For the test beam itself, using a TBM chip is optional. The additional risk associated with the TBM carrier PCB development, and the extra effort required to implement communication with a TBM in the test-board firmware is not required without the need to develop a general purpose pixel DAQ test-board and FED.

Item 4.1 (“evaluate performance of PSI46dig in light of test beam results”) was outsourced to other groups, and is not on our list of deliverables, although we will benefit from publication of results.

Item 3.1 (“participate in test beam data taking and data analysis”) has been completed. Most of the FPGA firmware used in the beam test was delivered by Bristol. A beam test is planned at Fermilab in summer 2013, but our involvement will be limited to willingness to answer questions and remote help with unexpected problems on a best effort basis.

DAQ Integration for pixel pilot modules

This part of the project had already been replaced by initial work towards a pixel FED. Other progress on this is described below.

FED-related work

In addition to Bristol FPGA firmware developments for the beam test, work was carried out by RAL TD to provide an event buffer for the IPbus-based readout in view of the longer term. A memory buffer design developed for a Xilinx Virtex 6 platform was adapted to a Spartan 6 development board chosen by the pixel group. This mainly involved changing the external memory interface and integrating it with the serial receiver and IPbus modules developed by PPD and Bristol.

Between October and December firmware was developed for a demonstrator 10 Gb Ethernet readout; an optical link system developed for the XFEL project was ported to a Xilinx Virtex 6 development board and reliable data transfer to a PC with a standard 10G network card was demonstrated. The XFEL system uses an FMC mezzanine optical module developed by DESY for the XFEL project; a commercial replacement FMC module which has equivalent functionality was used.

In parallel to the efforts by Bristol and RAL PPD on the test-board and ROC beam test, work began at Imperial on the design of a board which would serve as a µTCA pixel FED for the Phase I upgrade and as the means to read out the CBC2 in future module tests, and potentially a prototype of the future Phase II Tracker FED; it draws heavily on the MP7 developments for WP3. It is based on the FPGA Mezzanine Card (FMC) standard which was designed specifically to increase I/O flexibility for FPGA based boards.

An FMC has the benefit of lowering the costs of generating groups of boards with similar functionality but differing connectors as well as promoting a modular and standardised design applicable to multiple complex systems. For this reason, although it added some delay in developing the board, we discussed its design with the CERN CMS electronics group with a view to devising a standard board which could be used in several CMS applications and potentially outside CMS as well. This is progressing well and a simple MoU was written to clarify contributions; it is hoped CERN would eventually take responsibility for distribution and support of the boards.

The FC7 (FMC Carrier - Xilinx Series 7) is a flexible, general purpose µTCA card. It supports signaling rates up to 10Gbps and can host up to two FMC-compatible mezzanine modules, allowing it to be deployed for many applications. As an example, the board will form the basis of a scalable DAQ system for testing the CBC2 and for readout of multiple 2S-PT module prototypes during module testing and beam tests. It is also planned to use the FC7 as part of the TTC upgrade to a µTCA based system, for distribution of timing and status signals. Although UK effort on the pixel FED has stopped, the development may still provide hardware for the Phase I pixel DAQ.

Work on the board began in November 2012. Schematic work was primarily a CERN task, with help from Imperial College. The first set of schematics was completed in January, after which layout work could begin full-time. The FC7 layout is an Imperial responsibility and now complete; the board is currently 50% routed. The design is undergoing final checking and review in parallel with the routing and it is planned to submit PCBs for manufacture and assembly at the end of March. We expect the first assembled boards at the end of April.

Fig. 6.1 is an illustration of the current layout and annotated description of the top and bottom faces of the FC7. The top face is relatively free of components, mainly because it is designed to accommodate a wide range of FMC mezzanines (including non-standard ones) but also to allow space for adequate cooling, which can become a substantial problem with Xilinx 7 FPGAs. Additionally, this eases routing as the majority of the dense FMC signaling and high speed traces can run on the top. The bottom half of the board accommodates the power regulators, the majority of the clocking logic and all of the services, including those for the FMCs and for front panel interfaces (e.g. USB and µSD). The power architecture, services and peripherals have all been adapted from the MP7 except for the 4Gb DDR3 SRAM, which is a new feature, and the clocking circuitry, which has been adapted from the CERN GLIB.

The FC7 work will continue as part of WP2.

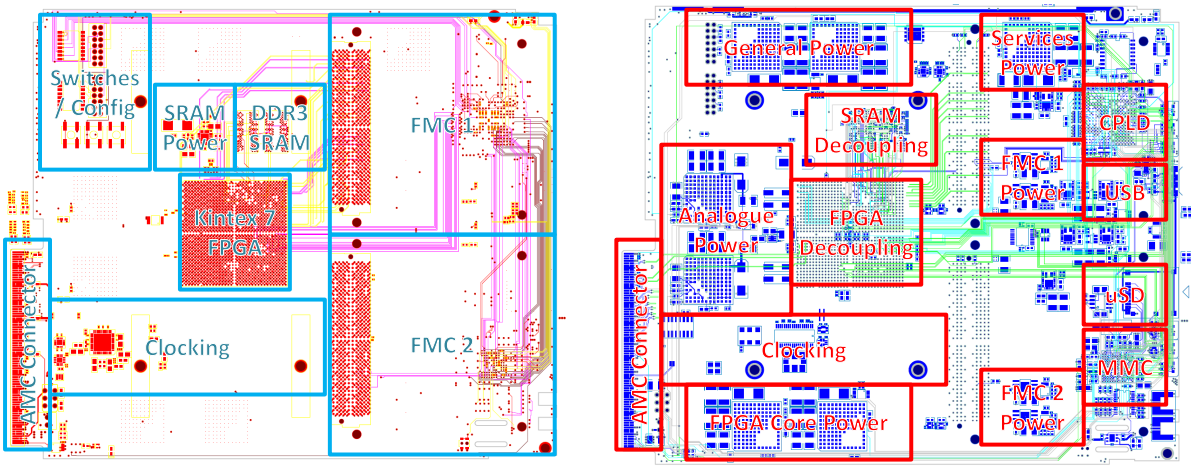

Fig. 6.1. Layout of the FC7, showing top (left) and bottom (right) faces of the board.

6.4 Staff on project

Reported in accompanying tables. RAL PPD staff levels are unchanged with respect to the previous report; G Zhang will no longer contribute to the work in future. The Bristol effort used for the beam test has not been separated from that reported under WP1. Most of the PPD effort is now being transferred to WP3 activities.

6.5 Expenditure

The main expenditure has been RAL PPD staff costs. Participation in the second high rate beam test involved travel expenses for transporting material and personnel to CERN, plus some equipment (such as optical fibres, hard disk for data storage). These expenses were covered by funds diverted from WP2 and WP3.

6.6 Comparisons with CMS activities elsewhere

The high rate beam tests are unique to the CMS pixel community. Other beam tests and X-ray test setups study various aspects of the pixel ROCs and sensors such as radiation hardness, amplifier gain, effect of parameter settings etc., but there is no other programme capable of verifying functionality of the chip with tracking under conditions close to those the phase 1 upgrade pixel detector will be exposed to at maximum instantaneous luminosity.

Two DAQ test-board solutions were developed as candidates for a general purpose pixel upgrade test-board. The UK option had good chances of being adopted as the standard solution because of its earlier availability and use in beam test projects demonstrating full functionality and readiness. However, due to withdrawal of the UK groups the PSI alternative will now be adopted as the standard board.

7 Risk register

The current version of the project risk register is v7.0, revised from the last version. Many of the risks have been retired. The major objectives have been met with the success of the MP7, adoption of the TMT architecture, delivery of IPbus, success of CBC1 and recent delivery of a promising CBC2. Some of the risks remain live, mainly in view of the longer term and next phase.

8 Finances

The financial report is summarised in the attached tables.

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

Upgrade travel spend continues at a steady level, with frequent visits to CERN, including beam tests. We are on target to use the entire budget by the end of March.

Equipment spending primarily on MP7 prototypes and related WP3 hardware, and the CBC2 submission and minor materials, has recently been substantial, although later than expected for reasons explained.

Engineering staff spending at RAL has also continued to be as high as expected 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 estimate for the remaining period, having checked the latest figures to mid-February, is that the final outturn for the final year will closely match the allocation.

The Working Allowance now stands at zero. The remaining budget will be used for MP7 and FC7 prototypes and staff costs.

Over the last year or so a significant expenditure has been made via CERN. For the CBC submission this was handled by a CERN invoice of about £118k to Imperial College in May. For the trigger developments, the expenditure has been by multiple orders which vary in size up to about £20k. The easiest way to handle this has been to charge the costs to the UK CMS Tracker team account, using a specific task code, then request invoices from SSC to Imperial College at regular intervals to repay the amounts. So far invoices for £13k (October) and £26k (December) have been paid and are included in the financial table. A further invoice of £34k was requested in February and there is an additional £52k of expenditure incurred in CERN yet to be invoiced to SSC by CERN.

STFC approval was sought for one month of Imperial RA costs (M Pesaresi) to be covered by some of the remaining funds, since the Imperial College grant would end at the end of February. It is estimated that ~£35k of the materials budget will then remain in the next financial year. A no-cost extension to the existing grants was requested from STFC, and has been agreed.

9 Gantt chart

The Gantt chart has been updated.

10 Milestones

The milestone table has been updated.