30 October 2017

University of Bristol
Brunel University London
Imperial College London
Rutherford Appleton Laboratory

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

1 Executive Summary

The new project started in April 2013; it continues work on the tracker and trigger begun in the previous R&D project. L1 track finding and endcap calorimeter work packages were added in 2016, but they rely on resources already within the R&D project, or on new external resources, from ERC and PRD grants. The Phase I upgrade of the L1 calorimeter trigger was completed in 2016.

In the last six months there has been significant further progress:

- Laboratory testing of the CBC3 ASIC has continued successfully. Some delays in manufacturing and assembling hybrids have been experienced, which limits the possibility for beam tests in 2017.
- The design work to revise the CBC3 is almost complete and we expect to submit the next version (CBC3.1) in early 2018.
- Further x-ray irradiations of the CBC3 have been carried out at CERN at several dose rates and temperatures, and a good model exists to predict the chip behaviour under realistic operating conditions.
- A CBC3 SEU test was carried out showing improvements in performance largely as expected, and well within a tolerable range.
- CMS concluded that an all-FPGA implementation of the L1 track finder, as proposed by the UK, should be the baseline for the project and has implemented a new managerial structure, with significant UK involvement.
- A paper reporting the design and achievements of the UK L1 track finder demonstrator, based on the TMTT concept, has been submitted for publication to JINST.
- The upgraded Level-1 trigger has continued to run operate well in 2017, with gradual adaptation to operational conditions; this is considered a CMS M&O project.
- Technology demonstrators of the next generation trigger hardware have been validated. The next series of prototypes are in design, with the objective of developing a family of boards able to meet the requirements of all those projects to which the UK is committed.
- There has been further progress in the HGCAL project with major contributions to the TDR due for submission in November.
- Good progress has been made on the design of the HGCAL back-end hardware, which is to be based on hardware developments from this R&D project.
- CMS UK has recently submitted its proposal for a five year construction project, from April 2019, to the PPRP, which builds on the foundations established during the present R&D project.

2 Project history and recent developments

The LHC upgrade is proposed to take place in two main stages, with an increase in luminosity reaching ~2 x 10^{34} $cm^{-2}s^{-1}$ in LHC Run II (achieved), then after a two to three year shutdown from 2023 in LS3, to 5 x 10^{34} $cm^{-2}s^{-1}$ levelled luminosity, denoted as Phase II at the High Luminosity (HL)-LHC. A total of 3000 fb^{-1} in integrated luminosity over about a decade is the goal. This should lead to a typical pileup of 140 events/BX but in view of the possibility to increase the luminosity even higher, or accommodate fluctuations without much degradation in performance, CMS aims to be operable at up to 200 events/BX corresponding to 7.5 x 10^{34} $cm^{-2}s^{-1}$ levelled luminosity.

The current phase of the project began on 1 April 2013. The technical work packages are WP2 for Phase II outer tracker readout R&D, WP3 for Phase I calorimeter trigger construction, now complete, and R&D aimed at Phase II, now with a strong common hardware element, WP4 on the high granularity forward calorimeter project and WP5 on L1 track-finding.

2.1 LHC operations and progress

CMS is currently taking high-quality data at very high efficiency. High energy p-p operation will end on 11 December, preceded by a period of machine development and 5 TeV running from 20 November.

- CMS recorded 35.4 fb^{-1} of good quality p-p data at 13 TeV centre of mass energy from the 38.4 fb^{-1} delivered by the machine as of 18 October (Figure 2.1);
- CMS switched to xenon ions for a very short period on 12 October and recorded ~3 μb^{-1} of data;
- Since the last Technical Stop in September CMS has achieved data taking efficiencies of 94-95%, even at the recent very high pileup (up to 60) conditions.

Overall efficiency in 2017 is just below 90% which reflects the impact of commissioning the new pixel detector. Some unexpected issues, in particular with the first pixel layer, were encountered and solutions found. However this still resulted in smaller data losses than the tolerable maximum (10% of 2017 data) agreed before installation.

The performance of the pixel detector, e.g. in terms of spatial hit resolution or hit efficiency, is already at design levels in barrel layers 2-4 and the forward disks, while optimization is still ongoing for the first layer.

The new two-phase CO_2 cooling system has performed very well. However, in the last couple of weeks a new issue has arisen unexpectedly, leading to the effective loss of a number of DC-DC converters, or related components, probably connected to power resets, which is not yet understood.

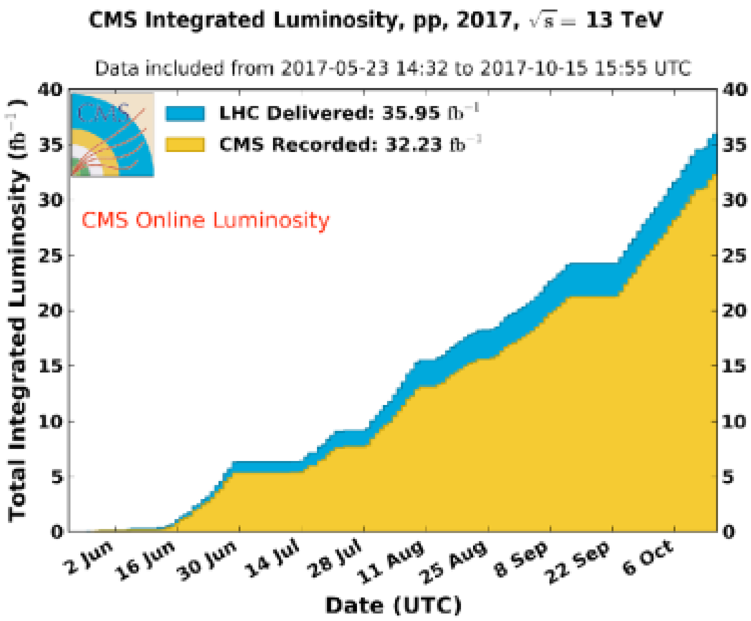

Figure 2.1. LHC and CMS integrated luminosity status in 2017.

After a technical stop in June, some machine instabilities arose which were not understood. An investigation was carried out in August, during which there was an accidental entry of air into a pumping port (possibly as a consequence of a power cut during the intervention). This led to condensation of nitrogen and oxygen in the beam pipe and on the thermal screen and a reduction in the heat load sustainable in this sector. Consequently the machine has been operated with an unusual bunch pattern which has led to higher than anticipated pileup to achieve the target luminosity. LHC operation is limited to around 1900 bunches instead of the nominal 2500. Thus, for the same luminosity, pileup is higher, with averages over 60 and for some bunches over 70. The machine has also changed the bunch structure to ameliorate the problem, which causes higher rates because it is less uniform (8 filled bunches followed by 4 empty bunches).

This potentially has consequences for CMS and ATLAS operations, especially for the trigger. CMS and ATLAS level the luminosity at 1.5 10^{34} $cm^{-2}s^{-1}$ by beam separation. This keeps pileup below 60 so CMS can run efficiently with less than 5% dead time at the start of fills and so far the experiment has coped very well with the new conditions.

Some details of the issue were given by F. Bordry in the RRB presentation in October. It is expected to be fixed during the end of year technical stop.

2.2 CMS planning

TDRs for the Tracking System, Barrel Calorimeters, and Muon Systems have been submitted to the LHCC and Upgrade Cost Group, and the TDR for the Endcap Calorimeter will be submitted by the end of November 2017. Interim documents have been submitted for the L1 Trigger and the DAQ/HLT systems, with their full TDRs planned for 2019-2021. In addition, a reference document on the infrastructure upgrades and the logistics of work during LS3 will be submitted by the end of 2017. CMS is also studying the use of precision timing measurements to optimize performance at high luminosity. Along with simulation studies of the performance benefits and an initial cost estimate, a conceptual design will be presented to the LHCC by the end of 2017 with a proposal to add the MTD (MIP Timing Detector) project to the Phase II Upgrade scope. If agreed, a TDR will be prepared for submission in late 2018.

The Tracker Phase II Upgrade project has achieved several important milestones in recent months. The Technical Design Report was completed and approved by the CMS collaboration and has been reviewed by the LHCC; preliminary feedback is very positive. The cost book, schedule, risk registry, and manpower estimates have been submitted to the Upgrade Cost Group and the review has just started. The first submission of full-size ASICs for the pixel-strip modules of the outer tracker, MPA and SSA, and of the first fully-functional prototype of the ASIC for the pixel detector, RD53A, was done at the end of August. An architecture fully based on FPGAs, as proposed by the UK, has been chosen for the novel Level 1-trigger track finding system of the future outer tracker.

2.3 UK adaptation to CMS planning

Following our submission of an outline construction plan for the period 2019-2024 in September 2016 to STFC, which was reviewed by Science Board, a complete proposal was prepared and submitted to STFC in October for review by the PPRP. It mainly builds on the activities within this R&D project, including the work packages added in 2016 on L1 track finding and the endcap calorimetry. Once this has been approved, the financing of the construction project will be finalised, enabling CMS and UK planning to proceed on a firm footing. This is expected to be achieved by mid-2018.

3 Work Package 1: Management

A reminder of the project management is included below. G. Hall remains PI, since January 2017 working 50%. D. Newbold has been UK CMS PI since October 2015, and will be succeeded by G. Davies in January 2018. Newbold will be PI for the CMS upgrade activity for the next period, with funding requested to start from April 2019.

The spreadsheet explaining assignments of individuals to work packages prepared for the last meeting has been updated, representing a snapshot at the end of September, and accompanies the documents for this meeting. Other staff changes are explained in the document. The principal change affecting management is that J. Borg has taken over the duties of M. Raymond, now retired.

WP	Manager	Institute	Role
1	G Hall, PI	Imperial	Overall management, budgetary responsibility and supervising procurements, interface to UK CMS PI. Tracker Management Board member.
2	J Borg	Imperial	Having taken over from M. Raymond, overall responsible for CBC specifications, interface to module design team, chip testing and module evaluation.
	M Prydderch	RAL TD	Manager of ASIC design team in RAL
3	G Iles	Imperial	Based in CERN with responsibility for L1 calorimeter trigger operation, as well as future hardware development.
	J Brooke	Bristol	Supervision of UK trigger upgrade activities.
4	P Dauncey	Imperial	Jointly coordinating UK HGCAL developments, with G Davies. L1 trigger coordinator for the HGCAL.
	G Davies	Imperial	Also, UK representative on the HGCAL IB and FB.
5	M Pesaresi	Imperial	Coordinating demonstrator integration activities, hardware/firmware specifications, and general project planning.
	I Tomalin	RAL PPD	Maintaining and running the Monte Carlo analysis software, tracking algorithms and oversight of RAL L1 track finder activities.

4 Work Package 2: Outer Tracker Readout

4.1 Objectives

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

4.2 Progress

The CBC3 is the latest prototype of the 130 nm CMOS bump-bondable front end readout chip for 2S-modules in the outer silicon tracker. The previous version of the chip (CBC2) was very successful, allowing R&D on (the rather complex) hybrid technologies suitable for the construction of 2S and PS modules to progress, and proof-of-principle demonstration of the pT stub triggering approach in beam tests. As discussed last time, the CBC3, with its optimized front-end and enhanced stub logic has also been a success. A few minor issues were identified during early tests, including results from total dose irradiations and SEU studies, which have been investigated further and discussed below. Work is progressing on schedule at RAL to modify the CBC3 design, with submission of the CBC3.1 for manufacture planned for January 2018.

Chips from the wafers submitted to a supplier for bump placement have been delivered (Figure 4.1), but the process does not appear to have been as successful as hoped, and include mistakes in the optical inspection step, as obvious errors exist in the lots of “good” and “bad” chips. (The metal bumps on the CBC2 were done by a sub-contractor to the foundry and were very successful. This time the process also used lead-free solder.) This is currently being investigated by our CERN collaborators. In addition, the first two-chip CBC3 hybrid module prototype has been substantially delayed due to procurement and manufacturing issues, and was not delivered from the manufacturer to CERN until October. Thus, they were not available in time for the anticipated October test beam. Meanwhile a simpler prototype was developed at Imperial to be studied in a test beam at CERN.

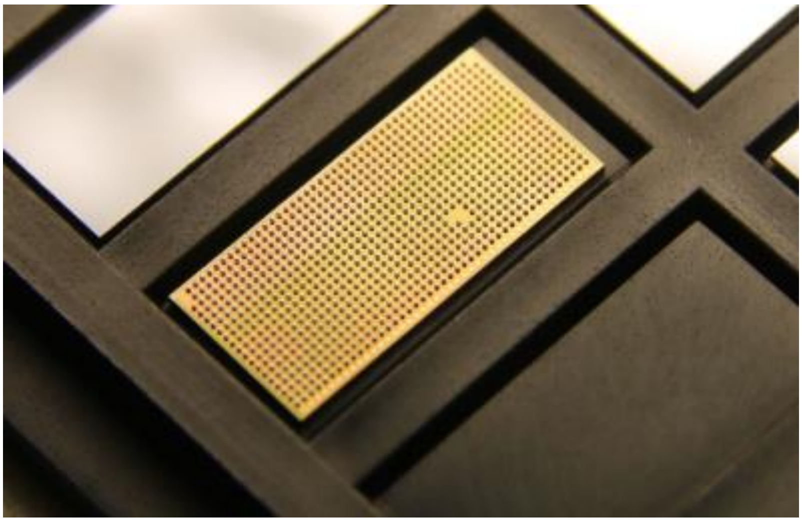

Figure 4.1. Bumped CBC3 chip. The apparent missing bumps are a pattern to permit orientation. Along the bottom right edge of the chip the row of un-bumped pads for wafer probing are visible.

The CBC3 design and test results from single chip testing were presented at TWEPP 2017 and publication in the conference proceedings is forthcoming.

4.3 CBC3 test results

Total ionizing dose (TID) tests performed at CERN showed significant increases in supply current when the chip was irradiated with high (relative to LHC) dose rates, despite changes to the pipeline cells which had been redesigned with enclosed-gate NMOS transistors to eliminate radiation-induced leakage effects. This was not totally unexpected, since there had been a substantial increase in the digital logic related to the stub finding. Irradiations where different parts of the chip were masked confirmed that the leakage originates from the stub logic built from standard cells rather than the pipeline memory.

Since the effect is also enhanced at lower operating temperature, the TID tests were extended to cover a range of temperatures and dose rates to allow the effect to be modelled and the increase over the lifetime of the detector under HL-LHC conditions to be calculated. The final conclusion is that an increase of ~0.4% can be expected during operation at HL-LHC, well within acceptable levels; it is intended that this study be published in the near future.

Single event upset (SEU) tests with 62 MeV protons (representative for the average cross section in the tracker at the HL-LHC) performed at the cyclotron facility at UCL (Louvain) were also performed. These show that the design changes from CBC2 to CBC3 yielded an improvement in SEU sensitivity of about a factor 7 compared to the CBC2. In the whole outer tracker operating in the HL-LHC, about 3% of the front-end chips would experience a bit flip per day that affects readout data and 7% would experience bit flips that affect stub (L1 trigger) data. This can be managed without design changes, e.g. by reloading all registers of each module once per hour (requiring about 160 ms). In addition, potentially sensitive circuit nodes in the latch circuit have been identified and minor design changes that could improve SEU tolerance are also being investigated, and may be implemented if the risk inherent in this design change is considered to be small.

When the extent of the delays in the manufacturing of the two sensor, two-chip CBC3 modules became apparent some additional effort was applied at Imperial to build reduced functionality hybrids using conventional PCBs and in-house assembly. This has been very successful, and single and dual CBC3 modules for a single sensor have been assembled and tested (Figure 4.2). One of the single-CBC3 modules is under test in a beam at CERN at the time of writing with analysis of the data still in progress, although preliminary results are clearly positive.

Figure 4.2. A single sensor 2CBC3 module built from a PCB designed at Imperial, with a small sensor attached and wire-bonded. This version includes some of the inter-chip communication signals and will allow this feature to be tested for the first time.

4.4 CBC3.1 design changes

Under test, the CBC3 was found to output repeated stub addresses and bend codes for a few locations. Additionally, invalid stubs formed by clusters outside the correlation window were allowed to be output with an invalid bend code identifier. The relevant circuit has now been modified to remove the bad addresses and bend codes. The simulated results have been plotted so as to replicate the test results, to be confident that the issue has been corrected. The circuit has also been modified to reject stubs that have an associated bend code that is invalid. It has been simulated and presented at a design review including the Imperial team.

Tests also showed an issue for the Parallel-In Serial-Out (PISO) circuit that outputs the L1 triggered data from the pipeline. This issue only arose for a few phase settings of the on-chip recovered 40MHz clock, and only when triggers were spaced less than 1µs apart, indicating a timing issue. The PISO circuit was modified to improve timing robustness. Additional control functions that allow active clock edge selection for certain settings of the 40MHz phase have also been added. This circuit has been thoroughly simulated and independently checked, and also presented at the design review.

We have investigated the origin of the SEUs in the I2C registers experienced during the Louvain study. Simulations confirm the test results, and show that the cause is related to a single reset buffer in the register design. The decision to correct the circuit or whether to implement a software fix with routine I2C register reloads is currently under consideration.

Two I2C-controlled test circuits have been added to the chip in order to allow testing inter-chip digital inputs and outputs during wafer probing. The e-fuse circuitry has been modified to include a pull-up resistor on the power supply so that the 3V power line is never allowed to float to ground, which has an effect on reliability of the fuse detection circuit, especially under irradiation. The Fast Reset control signal has also been extended in time to ensure consistent reset of the L1 counter circuit.

The new circuits have been incorporated into the top level layout and schematic and we are now conducting high level circuit verification to check that signal timing between circuit blocks is good. Once these simulations have been completed, we will hold a final design review.

4.5 Staff on project

M. Raymond retired in July 2017 and was replaced by J. Borg, who has proven his expertise since joining Imperial in an ERC-funded HGCAL post.

4.6 Expenditure

As usual the main recent expenditure has continued to be on RAL TD staff. The CBC3.1 design work is almost complete and within the expected budget. Whether significantly more effort will be required next year will depend on the success of the CBC3.1 submission, which will not be clear until mid-2018. The working allowance would still permit additional TD staff effort if required.

4.7 Deliverables

The WBS for WP2 is included below. Items in red under 2.3 and 2.4 have been discussed previously, but are left identified since they have had an unavoidable knock-on effect on the schedule. The overall goal remains, which is to be ready for full-scale production by the end of the project.

The delays have mainly been caused by late deliveries of the hybrids and modules and our inability therefore to sign off a fully tested module. There are actually other ASICs in the module schedule which are on the critical path for CMS construction, including especially a concentrator chip (CIC), which will be used by both 2S and PS modules. It is the responsibility of Lyon, and is in a 65 nm technology. This has been a concern to us for some time, but we have emphasised the importance of factorising the CBC and CIC development. Since the operation and data flow from the CBC in the 2S-module is much simpler than that from the PS-module, this should be possible.

WBS item 2.4.1 (early CBC3 tests) has been concluded and a large fraction of the activities of 2.4.2 (on-going tests), such as SEU and TID tests have also been completed. CBC3 2S-PT module studies (2.4.3) have been delayed and will begin in Q4 2017. The main part of the circuit that remains to be tested is the inter-chip communication logic that allows stubs to be created from hits in strips connected to adjacent front-end chips. This has only recently become possible with the completion of the first 2-CBC3 modules.

While CBC3 contains all the final system functionality, WBS item 2.5 is the CBC4 contingency iteration included to fix any remaining bugs. WBS items 2.6 covers the final preparations for mass production of the CBC4 (At present, we are referring to the next version of the chip as the CBC3.1, rather than the CBC4, as the modifications are so minor.) It is probable that the CBC3.1 iteration will be a full wafer engineering run, in which case it would be equivalent to item 2.6 and the separate item 2.5 would not be required.

Assuming WBS items 2.5 and 2.6 remain separate, the WBS has already been modified to fit both in the remaining project period. The WBS includes the original dates together with revised dates shown in bold red. Work has progressed substantially on the circuit redesign necessary to fix the issues so far discovered. CBC3 testing is ongoing, however, and CBC3 based module operation should be confirmed before taking a decision about when to submit a CBC4 follow-up iteration. A likely timescale for the final decision is the end of this year.

The revised WBS allows a longer period (until the end of this year) to finalise any design changes required for the CBC3.1. The period allowed for production is left unchanged, but the period allowed to verify the changes has been halved to three months, since as changes are to the digital logic only it should be possible to confirm correct functionality quickly.

The CBC4 mass production preparations, item 2.6, is harder to revise at present. If, as we hope, the CBC3.1 is the final version and the mask set used for its submission contains only the CBC3.1, then all that remains is to validate the CBC3.1 modules. If some issue is identified, it is likely to be very minor, in which case mask modifications might be trivial. However, funding exists for a complete new submission and we should be able to undertake this by the end of the project.

WBS	**WBS L2**	**Start**	**Finish**	**Months**	**Task Description**
2	Phase II tracker Readout	04/13	03/19	102	
2.1	**system**	04/13	03/14	12	**definition of the CBC-based SS-Pt module readout**
	2.1.1 specification definition	04/13	03/14	12	regular meetings with CMS collaborators to define overall system specification and interfaces
2.2	**CBC2 test**	04/13	03/15	24	**CBC2 is final deliverable of the UK upgrade R&D**
	2.2.1 CBC2 ongoing testing	04/13	03/14	12	complete the detailed studies of the CBC2 chip, including irradiation and SEU tests
	2.2.2 CBC2 2S-PT module prototype studies	04/13	03/15	24	a programme of 2S-PT module studies, in collaboration with CMS, including test beam
2.3	**CBC3**	04/14	03/16	24	**CBC3 is specified for the final system**
	2.3.1 CBC3 design	04/14	09/15	18	design period
	2.3.2 CBC3 production	09/15	03/16	6	production period
	2.3.3 test setup preparation	09/15	03/16	6	wafer and chip test setup preparation
2.4	**CBC3 test**	03/16	03/18	24	**CBC3 chip and module testing**
	2.4.1 early tests	03/16	09/16	6	chip verification tests to prior to module tests
	2.4.2 ongoing testing	09/16	03/17	6	complete characterization, including irradiation and SEU tests
	2.4.3 CBC3 2S-PT module studies	09/16	03/18	24	CBC3 based module studies in collaboration with CMS in lab and test beam
2.5	**CBC4 design and test**	09/16	12/17	15	**CBC4 is the final version of the chip, fixing any**
		07/17	**07/18**	12	**remaining bugs found in the CBC3**
	2.5.1 CBC4 design	09/16	12/16	3	design period
		07/17	**12/17**	**7**	**longer design period while waiting for CBC3 results**
	2.5.2 CBC4 production	01/17	06/17	6	production period
		01/18	**06/18**	**6**	**production period**
	2.5.3 testing	07/17	12/17	6	tests to verify full and final functionality
		06/18	**08/18**	**3**	**expedite functionality tests**
2.6	**CBC4 mass production**	01/18	12/18	15	**a full wafer engineering run is required for CBC4**
	preparations	**06/18**	**03/19**	**9**	**in preparation for mass production**
	2.6.1 CBC4 final masks	12/18	03/19	3	mask preparation for full wafer engineering run
		06/18	**08/18**	3	**overlap mask preparation with CBC4 test**
	2.6.2 CBC4 engineering run	03/18	09/18	6	production period
		09/18	**03/19**	**6**	**production period**
	2.6.3 CBC4 final production	09/18	12/18	3	final functionality check
	readiness verification tests	**03/19**	**03/19**	**0**	**push final verification into post project period**
	2.6.4 procurement planning	01/18	12/18	12	detailed financial plans for mass production

5 Work Package 3: Level-1 Trigger

5.1 Objectives

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

5.2 Progress to date

5.2.1 Phase I trigger operation

The Phase I upgrade to the Level-1 trigger has been running successfully since May 2016. Since our last report, the system has been operating smoothly and reliably with negligible contributions to downtime or data marked as bad. Due to beam-loss problems attributed to frozen O_2 and N_2 in a particular sector of the LHC, the machine has been operating with a different configuration than planned, resulting in higher pileup conditions than anticipated, with a mean of up to 60. This has tested the pileup mitigation scheme implemented in the calorimeter trigger during the 2016-17 technical stop, but trigger rates have remained acceptable. Further optimisation is foreseen during the 2017-18 stop.

The final components of the data quality monitoring system for the Phase I calorimeter trigger have been commissioned, including performance monitoring via comparisons of offline reconstructed objects with L1 trigger objects, as well as validation of the trigger data processing by comparison with offline emulation. In addition, a software framework for prompt analysis of the trigger data in case of unexpected problems has been developed, completing the suite of tools required to operate the trigger system for the rest of its lifetime.

5.2.2 Next generation hardware development

Looking towards the future, work has continued on a successor to the MP7 for use in the Phase II CMS upgrade. The preliminary work for this is now complete (Figure 5.1: left) with the two technology demonstrators validated (a) the **MP-Ultra** - a small, low cost, form factor (PCIe) card to evaluate an alternative PCB stack-up to accommodate the increased density of newer components, and to provide experience with the latest generation of optical engines and Ultrascale FPGAs (b) the **Service card** - a full size ATCA card to test and validate the core services (e.g. embedded CPU, Ethernet, PCIe, power, signal integrity, thermal aspects, etc.).

These designs have been merged into an initial ATCA prototype for Phase II that will be submitted before the year end (Figure 5.1 Right). A recent design review recommended only minor modifications. A key aspect of the design is the use of an interposer array connector (Figure 5.2) that allows the FPGA(s) to be mounted on a daughter card, with the objective to both mitigate risk and increase flexibility. While mounting FPGAs on a daughter card is not novel, the use of an interposer array connector is. The spring contacts reside on the interposer, so a faulty contact requires that only the interposer be replaced.

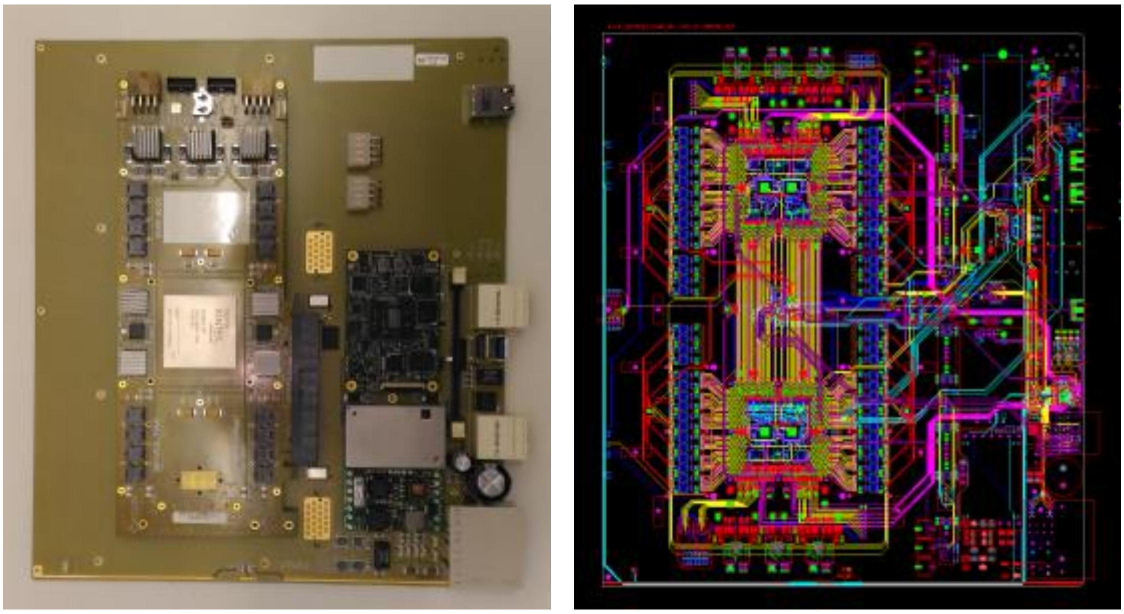

Figure 5.1. Left: The MP-Ultra and embedded CPU module mounted on the Service Card. Right: The MP-Ultra and Service Card merged into a single design. The two large mounting sites for FPGA daughter cards are clearly visible, as are the optical engine mount points on either side (blue columns).

The cooling aspects of the ATCA form factor remain a concern, but at present it is thought manageable, albeit with a large amount of power devoted to fans (up to 20%) and significant noise (in excess of 85 dBA). At the last review there was some concern about operating the FPGAs close to their design limit (e.g. 110°C) for extended periods, but after further investigation this is considered minimal risk. However, the optical engines must be kept at less than 50°C and ideally less than 40°C to have a reasonable lifetime. This is likely to be challenging and thus we have opted to ensure that the optical engines have access to cool air, with most of the heat concentrated in a vertical column over the FPGAs. The fallback position at present is to deepen the crate to make maximum use of the cooling available within the rack. This would be relatively inexpensive, but will need agreement from collaborators. We shall continue to keep cooling options under review because we need to make sure that the full system, which will involve many crates, is a viable option for the underground service cavern.

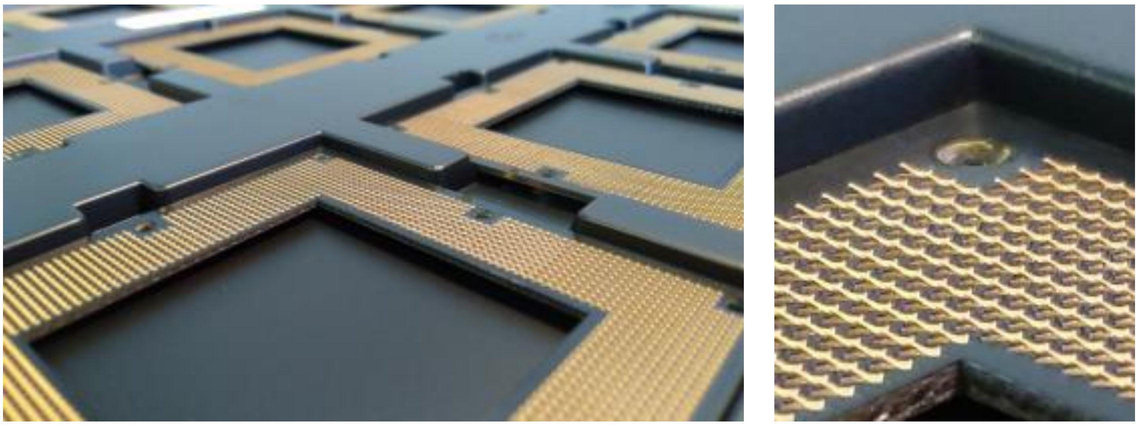

Figure 5.2. Left: The interposer that connects the FPGA daughter card. Right: A close-up of the individual connectors.

5.3 Overview of CMS plans

The main focus of the Phase II trigger project, led by J. Brooke (Bristol) and R. Cavanaugh (UIC/FNAL), has been preparation of the L1 trigger Interim Technical Design Report, which was submitted to the LHCC on 12 September. The main focus of this document is baseline specification of the interface between sub-detectors and the central L1 trigger. Options for the architecture of the future L1 trigger are also presented, including the use of time-multiplexing based on the success of the Phase I calorimeter trigger. The document also describes progress on trigger algorithms, including the use of “particle-flow” techniques, whereby individual particles are reconstructed using the full detector information.

The next major milestone in the CMS Phase II trigger project is publication of a full TDR at the end of 2019. The key decisions required ahead of this milestone include finalisation of the processor specification, system architecture, and baseline algorithms. These will be informed by ongoing algorithm studies, using both physics simulation and firmware emulation, as well as hardware R&D. Good progress is being made on all these fronts, and we do not anticipate problems meeting the milestone. In particular, good progress has been made in showing that particle-flow techniques are feasible and profitable in the L1 trigger, within the constraints of bandwidth, firmware resources, and latency.

Plans for UK contributions to the Phase II trigger have been consolidated, focusing on the Correlator system, where central tracks, calorimeter clusters, and muon tracks are matched to form the final physics objects used in the L1 trigger decision. We will collaborate with US institutes (FNAL, UIC, MIT) and CERN, on the correlator hardware and particle-flow algorithms, and are in the process of building this collaboration. We will contribute to the core particle-flow algorithms via vertex identification, as well as identification of jets, energy sums and electrons. The UK is already contributing to Correlator design studies, with a focus on vertex identification; milestones are described later.

The optimal hardware configuration will ultimately depend on the choice of algorithms and how they are implemented. The current generation of cards, built for the L1 trigger Phase-I upgrade (e.g. MP7), allows the topology to be easily changed because all the L1 trigger interconnects are based on standard multimode optical links. This flexibility, which was not present in past L1 trigger designs, has enabled rapid prototyping of new trigger systems (e.g. UK+KIT L1 track finder).

For Phase-II we will endeavour to add to this flexibility by decoupling the FPGA processing unit from the baseboard. The intention is to make the baseboard a low cost, relatively simple object into which optical modules, embedded processors and the FPGA are plugged. The FPGA will be mounted using a new array-based, high density, interposer connector that should operate at 28Gb/s. This should not only allow FPGA flexibility, but also mitigate risk because should there be a fault with the baseboard all pluggable components (i.e. FPGA etc.) can be remounted on a new baseboard.

As described earlier, the baseboard is almost ready for submission. Our objective is to have a working initial prototype, with a preliminary firmware/software stack in Q2 2018, with a revised final prototype version by Q2 2019. Detailed planning is now underway to decide the programme of tests that will be necessary when the card returns from manufacture (e.g. signal integrity, thermal studies) and the optimum choice of FPGA daughter card for testing and algorithm development. It is likely that we will exploit the baseboard flexibility by developing several FPGA daughter cards to support different UK Phase-II activities. Given the short timescale it will be essential for the basic features of the firmware/software stack are already well developed when the card returns from manufacture in Q1 2018. At present work on this objective has been paused because of the hardware design submission, but it should resume within the next few weeks.

While it cannot be guaranteed that a single baseboard design will satisfy all requirements of future UK Phase-II activities it is expected that a small family of similar cards should meet the requirements of most projects. This approach should ensure that cards have a common firmware and software stack, thus saving significant human resources.

The baseboard features are required by many projects for Phase-II and thus we are in collaboration with others (e.g. CERN DAQ group, KIT) to reuse common solutions where possible.

5.4 Staff on project

There are no staff changes.

5.5 Expenditure

Overall spending is within the budget foreseen.

5.6 Deliverables

The deliverable list is appended below. Blue font means complete. Red font is delayed. PM represents duration of the task.

The original milestones for the long term trigger R&D were defined at the outset of the project well before the present picture evolved and, in particular, the possibilities for common hardware developments which would benefit other sub-projects, some of which were not initially foresee, such as the L1 track finder and the HGCAL. New milestones which are listed below are proposed for the remainder of the project. The longer term objective is to be compatible with the proposed schedule of the UK CMS construction project and the overall CMS sub-detector milestones.

For the hardware development:

- Q1 2018: preliminary Firmware Software Stack for Xilinx UltraScale Devices (build tool, slow control and documentation).
- Q2 2018: prototype v1 hardware available with basic infrastructure (i.e. slow control and payload only).
- Q2 2019: prototype v2 hardware available with full infrastructure set for prototype card (i.e. slow control, payload, link and diagnostics).

For the development of the system design: the CMS milestones for this activity during the remainder of the grant period are to benchmark the simulated performance of selected algorithms by Q2 2018, followed by the full suite of baseline algorithms by Q1 2019, in preparation for a hardware demonstrator programme. The UK will focus on vertex and jet performance for the Q2 2018 milestone, then expand to include energy sum and electron identification for Q1 2019.

- Q2 2018: benchmark performance of selected algorithms.
- Q1 2019: benchmark performance for all core algorithms.

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

6 Work Package 4: High Granularity Calorimeter

6.1 Objectives

- To complete the testing of the UK parts of the design for the front-end electronics HGCROC ASIC.
- To play a leading role in the HGCAL TDR, in terms of both overall project management and technical aspects, specifically in the two areas listed below.
- To develop the trigger primitive generator (TPG), including contributing to the algorithms, firmware and HGCAL-specific hardware.
- To study the physics performance of the HGCAL and optimise the design parameters and, in addition, develop the reconstruction techniques to provide the best overall performance.

6.2 Progress to date

The UK HGCAL effort has been mainly concentrating on consolidating results for two major documents, namely the CMS UK Phase II upgrade proposal mentioned earlier and the HGCAL TDR which will be submitted later this month.

The UK work on the front-end electronics has continued since the last OSC meeting. A UK engineer (J. Borg) designed a low-power (<1 mW) 50 ps resolution TDC circuit for digitizing the time-over-threshold (TOT) signal based on gated ring-oscillators. This circuit was included on the HGCROC V1 ASIC submitted in July 2017. This engineer will test this part of the circuit when the ASIC is returned from fabrication in early November. These tests are expected to take up to two months. In the CMS-UK Phase II upgrade proposal, because of funding constraints, we have chosen to restrict future UK HGCAL effort to the two remaining topics described below. Some limited activity on the front-end electronics will continue for the duration of T. Virdee's ERC Advanced grant, subject to personnel changes detailed later.

Major progress has been made on the design of the HGCAL back-end hardware, which consists of both the DAQ and the trigger primitive generator (TPG) boards. These areas are both led by the UK and the latter is a proposed UK deliverable in the UK upgrade bid. The system-level design has now been firmed up significantly, with a costed baseline for all parts of the system documented as part of the upcoming TDR. An overview of the TPG structure in terms of numbers of boards and data rates is shown in Figure 6.1.

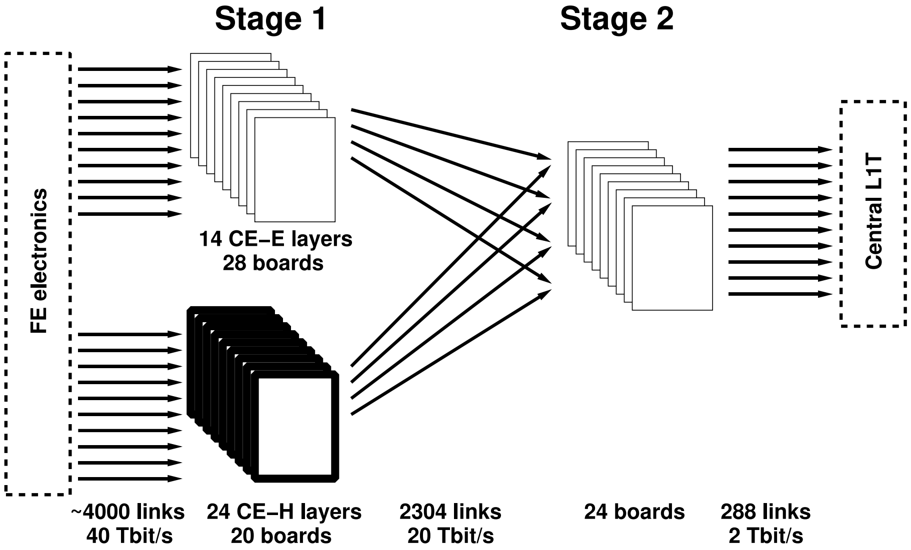

Figure 6.1. Overview of the TPG hardware for one of the two identical endcap systems. Stage 1 forms 2D clusters within a layer, while Stage 2 combines these in depth to give 3D clusters which form the primitives.

The formation and growth in activity of the generic board design team (see WP3) allows us to exploit commonality across CMS subprojects and develop widely usable firmware and software. This has significantly accelerated the firming up of the HGCAL board specifications. The boards are now well defined and the schedule adjusted forward compared to previous expectations to reflect this. There will ideally be one (although in reality probably more) common generic base board holding daughter boards specific to each particular application. A conceptual idea of the generic board and one of the HGCAL-specific TPG boards is shown in Figure 6.2 (a) and (b).

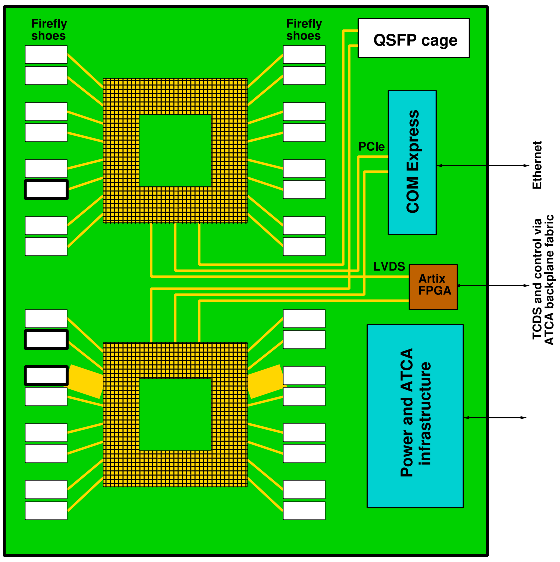

Figure 6.2(a). Generic base board showing main interfaces, the two sites for mounting daughter boards, and the bases for inserting fibre optic transceivers

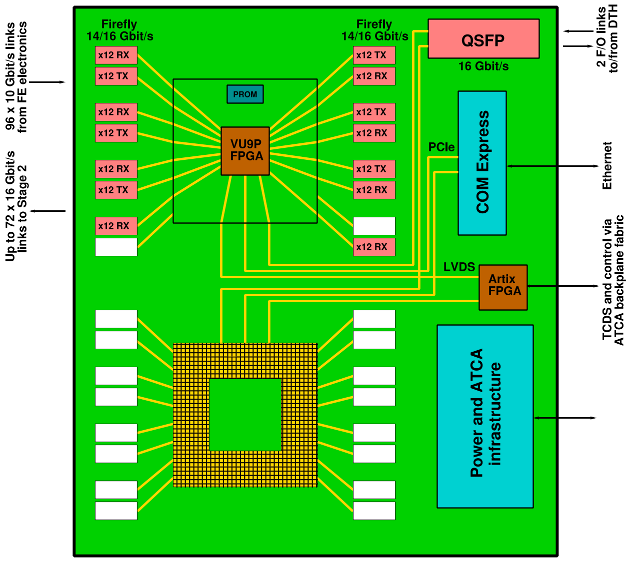

Figure 6.2(b). The generic board with one daughter board and 172 fibre optic transceivers mounted, which would be the required configuration for a TPG Stage 1 board. The Stage 2 boards are identical except they require a smaller number of fibre optic components.

As discussed in WP3, the first prototype of the generic base board will be available early in 2018, so the design and manufacture of the simplest HGCAL-specific daughterboard (used for the DAQ application) will be completed a few months later. The more complex (and more expensive) TPG daughter board should be manufactured by the end of 2018. This is earlier than originally anticipated for the HGCAL and has allowed us to add schedule contingency into the project. A significant effort to produce the firmware and software for thoroughly testing these HGCAL daughter boards will be required over the next 18 months. This will be written in collaboration with groups from France and Croatia, and will be coordinated by the UK.

The primary aim of recent work on simulation and reconstruction has been producing results for the TDR. The UK has contributed major parts of these, in particular producing estimates of the photon resolutions under the most extreme conditions expected, namely an pileup of 200 interactions per bunch crossing, see Figure 6.3, rather than the 140 interaction pileup expected as the average.

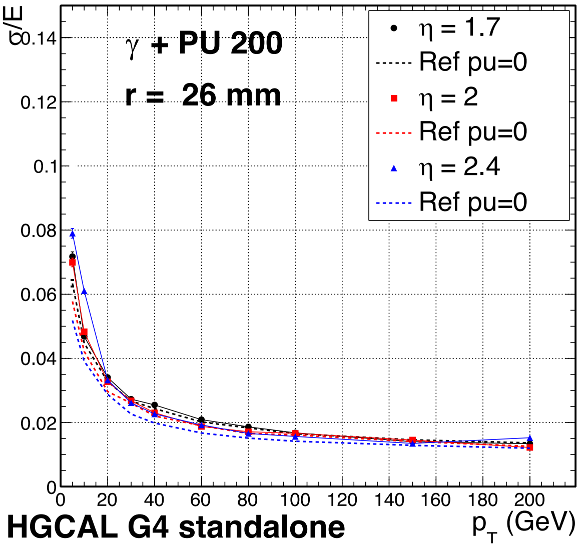

Figure 6.3. Fractional energy resolution, σ/E, as a function of p_T for unconverted photons at η = 1.7 (300 μm silicon thickness), η = 2.0 (200 μm), and η = 2.4 (100 μm), using a region of radius 2.6 cm in each layer to sum the energy. The energy resolution is not degraded very much by the addition of an average of 200 pileup interactions per bunch crossing.

UK personnel have also produced realistic techniques for calibrating the individual calorimeter cells using MIP deposits. This has involved effectively using neighbouring layers to form short “tracks” and this, together with isolation criteria, gives a clean signal even in the highest pileup occupancy regions. A typical MIP peak fit is shown in Figure 6.4.

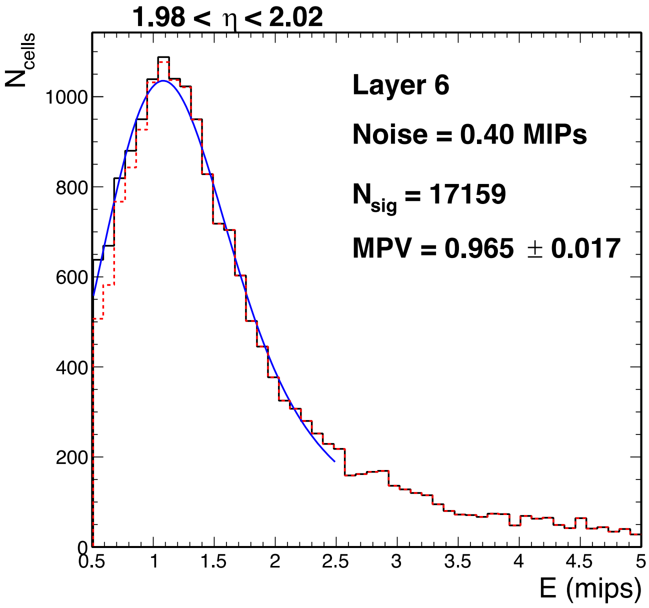

Figure 6.4. Example of a simulated fit to determine the MIP peak position for cells around η = 2. The cut at 0.5 MIP corresponds to the threshold in the HGCROC ASIC for the channel to be read out. The fit includes a tail from pedestal events, for which the distribution is centred at zero and hence not visible in this figure.

As an example of UK results feeding into detector optimization, this calibration work has raised the potential implementation of a “threshold bypass” circuit in the HGCROC ASIC, allowing a low rate of channels to be read out (contributing less than 1% of the total rate) even if below threshold. This would allow an accurate determination of the pedestal shape below the MIP peak, in order to constrain the MIP peak better in the fit.

6.3 Overview of CMS plans

The most important element of the overall CMS planning for the HGCAL is the scheduled Technical Design Review, which is due for submission at the end of November. It is currently under collaboration-wide review. The UK has contributed many of the results for this document and UK personnel have played a major role in writing it; a total of around 60 of the 250 pages were a UK responsibility. Most significantly, the overall editor (C. Seez) is from the UK.

The most notable CMS-wide future HGCAL milestone is the EDR in July 2020. The UK will be heavily involved in documenting the back-end electronics designs for this review, as well as producing new and updated simulation results. Prototypes of boards for both the DAQ and TPG applications will need to be manufactured and tested thoroughly (including the interfaces to the other systems such as central DAQ and L1T) in time for this review.

6.4 Staff on project

Since the last meeting, Imperial have hired a new RA (T. Strebler) funded from the ERC Advanced grant to replace the RA funded from the HGCAL PRD grant (L. Mastrolorenzo), who is leaving at the end of the month when the PRD grant finishes. The engineer hired on the ERC grant (J. Borg) has moved post to take a CG engineering core position (previously held by M. Raymond, who retired in July). As he has picked up the CBC ASIC testing responsibilities in WP2 most of his effort will in future be in that area. The appointment of a replacement engineer on the ERC grant has been put on hold, and hiring may now be done through CERN where the effort is most needed.

6.5 Expenditure

Expenditure so far has been almost entirely on staff, and on travel to a lesser extent, mainly using non-project (ERC and PRD) funds. However, this is likely to change over the next six to twelve months. The prototypes of the TPG daughter-boards will be relatively expensive as they will each hold a large FPGA (Xilinx Ultrascale+ VU9P). We foresee around £50k from the ERC grant being spent on these components plus some infrastructure required for testing.

6.6 Deliverables

These are as follows for the three areas in which the UK is involved:

- The main remaining UK deliverable for the front-end electronics work is the evaluation of the Imperial-designed TDC implemented in the HGCROC V1 ASIC, which is expected to be complete in February 2018.
- Over the next year, the main deliverable for the TPG is the production of a generic base board and daughterboard to produce a functional prototype of the TPG board.
- The near-term deliverables for simulation are: an updated and detailed GEANT geometry description that corresponds to that described in the TDR; clustering algorithms capable of efficient three dimensional pattern recognition of hadron showers in the presence of high pileup; and a fully realistic simulation of the precision timing in the HGCAL electronics, allowing exploitation of this information.

A set of pre-TDR HGCAL milestones were defined for October 2015 to November 2017. Recent milestones relevant to the UK effort are:

- Submission of HGCROC V1 ASIC submission in March 2017: The V1 ASIC was submitted in July 2017. The delay was not associated with UK activities.
- Baseline definition of the TPG architecture in September 2017, Baseline definitions of the event and TPG raw data formats, both in October 2017: All three milestones were required for the TDR and were achieved on schedule.

- Performance results of the TPG system in October 2017: This milestone has not yet been fully achieved at the time of writing. The results are still in preparation and the TDR will be updated after the internal review with improved figures, before submission at the end of November. This is a collaborative effort involving the UK and other institutes and the delay has been mainly due to a shortage of available people.
- Large scale production of fully simulated and reconstructed events for physics and trigger studies in March 2017: A large set of samples is needed for the final results produced for the TDR. Smaller (but still significant) samples were delivered earlier in the year and have been used to prepare the TDR studies. The large samples were simulated at the GEANT4 level in September and reconstruction will be completed in October.

As part of the TDR preparations, a review of the HGCAL schedule has resulting in an update of the future milestones for the detector construction covering January 2017 to December 2025. There are no explicit milestones yet for the future simulation work. The main near-term hardware milestone relevant to the UK effort is:

- Prototype board hardware, firmware and software basic validation in April 2019: This is the next significant milestone for the TPG project. To meet this, we aim to have a functional prototype assembled by the end of 2018. Note, this milestone is just after the end of the current STFC upgrade grant.

7. Work Package 5: L1 track finder

7.1 Objectives

- To design the architecture and technological implementation of a first-level track finder for the CMS Phase II upgrade.
- To demonstrate and document a prototype track-finding system, as required for CMS review purposes, design reports, and integration exercises.
- To generate a construction plan for the CMS track finder and readout system, including any R&D required for final implementation decisions.

7.2 Progress to date

As recorded in the previous report, by December 2016 the UK had successfully demonstrated an FPGA-based system for reconstructing tracks from the future Phase II tracker within a latency requirement of 4 μs, such that they could be used as input to the L1 trigger decision. The demonstrator system used corresponds to a single Track Finding Processor (TFP), consisting of several MP7 boards connected by optical links, where each MP7 is equipped with a Virtex7 FPGA. The firmware algorithms employed in the TFP are capable of reconstructing tracks with $p_T > 3$ GeV from 1/8 of the total angular acceptance of the tracker (known as an octant) for one LHC event in 36. Tracker hit data (known as stubs) from simulated HL-LHC collisions are injected into the TFP, and the tracks reconstructed by it are in almost perfect agreement with those expected from emulation. When processing $t\bar{t}$ events with 200 superimposed pileup collisions, tracking efficiencies in excess of 94% are achieved, and a latency (time of last track out relative to first data in) of 3.7μs. A paper describing this work has just been submitted to JINST.

The track-finding algorithm consists of three blocks: a Geometric Processor (GP) assigns the stubs in an octant into 2×18 regional segments in $\varphi\times\eta$, inside each of which tracking is done independently; a Hough Transform (HT) then does coarse track-finding in the r-φ plane; and finally, Kalman Filter (KF) and Duplicate Removal (DR) algorithms fit the tracks in 3D, cleaning them to eliminate duplicate tracks.

If the full tracker L1 tracking system were built from this technology, then 88 x 36 = 288 TFPs would be needed. However, use of higher speed optical links (16 Gb/s) will allow tracker data to be transferred to the TFPs at twice the rate, meaning that each TFP would receive one event in 18, reducing the number of TFPs required to 144, and the expected system cost to 4.3 MCHF. This would also cut the latency to about 2.5 μs. The TFP could process the increased data rate if its FPGA operates at twice the frequency, (although other solutions exist). Good progress has been made towards demonstrating that this is possible in next-generation (Ultrascale) FPGAs. The mathematical calculations in the GP algorithm run at 500 MHz; and the HT runs at 416 MHz, with improvements currently being implemented to increase this frequency further. Besides this, comparisons are being made of the GP and KF performance when their firmware is written in VHDL, as opposed to a High Level Synthesis (HLS) language. Furthermore, a possible alternative to the KF, known as the Linear Regression (LR) track fitter now gives identical tracking efficiency to the KF, whilst using slightly less latency and FPGA resources.

Two significant changes to the tracker design have been introduced relative to that assumed in the December review: some modules in the tracker barrel are tilted such that they point towards the interaction region; and the tracker readout is now divided into nonants (1/9 of φ) instead of octants. Simulation studies indicate that the tilted tracker geometry reduces the data rate through the L1 track-finding chain (e.g. by ~30% after the HT block) which increases the robustness of the system, whilst neither tracker design change causes any loss of tracking efficiency. Studies performed in the UK of data rates in the nonant design helped define the cabling scheme used to read out the tracker.

7.3 Overview of CMS plans

Following the review of competing L1 track-finding proposals undertaken by CMS in December 2016, and after taking advice from a Task Force of technical experts, Tracker management adopted the FPGA-based L1 track-finding solution as its reference design, relegating the alternative solution based on Associative Memory (AM) chips to a back-up for the unlikely scenario that a fundamental problem would be found with the FPGA-based solution.

Two alternative FPGA-based L1 tracking solutions were presented at the December review: the UK one (“TMTT”) and a US one (“Tracklet”). As a result of discussions within the Task Force, the Tracklet group agreed to adopt a time-multiplexed system architecture based on the UK proposal. This facilitates collaborative work between the two groups, in particular allowing them to work together on a common hardware and system design.

A new tracker working group structure has been defined following this agreement, which is shown in Figure 7.1. The “BE System Development WG”, co-convened by M. Pesaresi, is responsible for the system architecture and hardware design of the TFP and also of the DTC boards, which receive stub data from the Tracker FE and, after pre-processing, forward them to the TFP. It is expected that the DTC and TFP will be variants of a common board. This WG is also in charge of firmware and algorithm development. A Data Processing WG is responsible for liaison with all CMS groups with an interest in the L1 tracking, and in the short term, will also devise tests to confirm the robustness of the FPGA-based solution.

This all reflects a very successful outcome for the UK, following the excellent L1 tracking results it presented in the December 2016 review. The original objectives of this WP have been met.

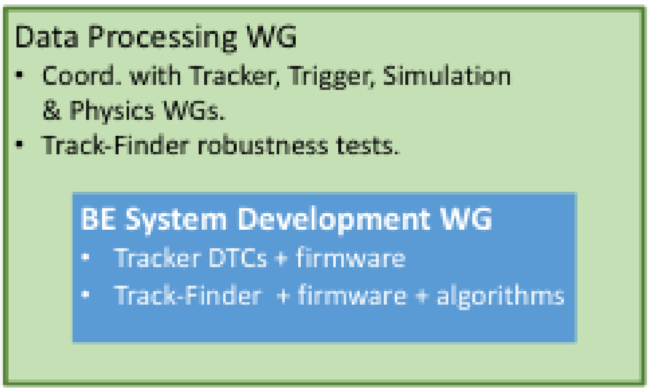

Figure 7.1: The new tracker back-end (BE) electronics working group (WG) structure.

7.4 Staff on project

As mentioned in the previous report, some effort has been transferred from WP5 to WP3 this year, profiting from the significant overlap in hardware R&D between these two WPs and their common interest in L1 tracks. EU-funded effort on WP5 is coming to an end, with L. Calligaris having already left, and D. Cieri completing his PhD soon. However, C. Manolopoulos, M. Pesaresi, I. Tomalin continue to devote the majority of their time to WP5, and several other people remain active. Furthermore, we continue to benefit from a strong collaboration with KIT, particularly in the fields of firmware and hardware development.

7.5 Expenditure

Currently only modest expenditure has been required since the prototyping uses spare MP7s from the L1 trigger project and WP3, and related equipment. Some travel expenditure has been incurred on meetings in CERN and the UK.

7.6 Deliverables

The deliverables list presented in the previous reports, and listed below, is now completed. Blue font means complete.

Milestone Date	Description
16Q4	Presentation of demonstrator at CMS internal review
17Q3	Definition of track finder design
19Q2	Operation of track finder systems

Following the adoption by Tracker Management of the FPGA-based L1 tracking solution, we can formulate a new list. Work will continue to refine the Tracker BE architecture and hardware design; to optimise the L1 tracking algorithms whilst seeking convergence with the Tracklet group; and to improve the firmware so that it can be used on Ultrascale FPGAs operating at higher frequencies. Since the FPGA-based boards for WP3 and WP5 will be variants of each other, R&D on these will proceed in common, and will be presented within the WP3 report. A demonstrator system will be built using the, the first prototype of the generic ATCA base board, equipped with Ultrascale FPGAs, that as explained in WP3, will be available early in 2018. This initially be used to demonstrate the TFP, and later the DTC+TFP chain.

Milestone Date	Description
18Q4	Demonstrator of TFP using Ultrascale FPGAs
19Q2	Demonstrator of DTC+TFP chain using Ultrascale FPGAs

8 Risk register

The risk register has again been reviewed and some risks added and retired, as well as revised.

Risk 8.3 has been added, which refers to the possibility of a laboratory move in the Imperial Physics department interfering with the work of WP2.

Risk 9.1 for WP4 has been retired since the HGCAL ASIC development no longer is a UK responsibility, but Risk 9.3 has been added to cover possible delays in developments of new hardware impacting progress.

Risk 10.2 refers to the L1 track finder demonstrator, WP5. As described, this reached a satisfactory conclusion so was retired.

9 Finances

Expenditure is reported in the usual financial table.

As reported before, RAL PPD staff costs are somewhat higher than originally expected. Bristol staff costs were not available for the financial table at the time of writing to due to absences for overseas travel, so provisional figures are included, which will be updated.

As discussed in the last OSC meeting, £100k of the Working Allowance will be used to cover extra TD staff costs for the remainder of the project. After discussions with STFC office, it appears the best way to manage this, and the request to increase the travel budget from the WA, is to overspend compared to the allocation, but to underspend by a similar amount the equipment budget at Imperial College. STFC will reconcile these at the end of the project.

Despite exchange rate changes the overall situation remains favourable for the remaining expenditure foreseen, and has actually improved somewhat of late. However, this could easily fluctuate in the other direction, but the expenditure margins could accommodate this.

We have made estimates of the major items of expenditure remaining, using current exchange rates:

Item	**£k**	**Comment**
CERN expenditure in pipeline	43.1	To be invoiced shortly
CBC3.1 submission	296.5	NRE and 12 wafers
CBC_final	296.5	NRE and at least 12 wafers: CORE
Ultrascale development	185.0	
RAL TD effort	100.0	Additional to that remaining in financial table
Additional travel	40.0	Compared to proposal
Total required	961.1	£1,137k available

In the above, we still assume that two full wafer submissions are needed to revise and finalise the CBC. The 24 CBC2 wafers mentioned last time will be paid by the tracker project, since they will mainly be used for prototype module construction by others in the collaboration.

10 Gantt charts

Recently most of the work underway aims to conform to CMS planning. The tracker TDR contains a very detailed summary in Gantt form which is included in the documentation, and the UK CMS project plan shows the plan for the period from April 2019, obviously subject to PPRP review; this is also included in the documentation. A revised HGCAL Gantt chart is in preparation to be submitted with that TDR at the end of November so we await that.

11 Milestones

The deliverables from each work package are listed below. The milestones which were due have been highlighted in red font, or those met in blue.

For reference, the reporting date of September 2017 corresponds to PM54.

Deliverable	Date	Description	Rev.Date
M2.1	PM12	System specification document produced	PM12
M2.2.1	PM12	Documented CBC2 detailed test results	PM12
M2.2.2	PM24	Documented 2S-PT module results	PM24
M2.3.1	PM12	CBC3 ready for production	PM39
M2.3.2	PM18	CBC3 produced & test setups ready	PM42
M2.4.1	PM24	Documented early CBC3 test results	PM45
M2.4.2	PM30	Documented CBC3 detailed test results	PM48
M2.4.3	PM60	Documented CBC3 2S-PT module results	PM60
M2.5.1	PM42	CBC4 ready for production	PM54
M2.5.2	PM48	CBC4 produced	PM57
M2.5.3	PM54	Documented CBC4 test results	PM60
M2.6.1	PM60	Final production masks prepared	PM64
M2.6.3	PM69	CBC4 ready for mass production	PM69
M2.7.3	PM72	First production modules available	PM72
M3.1	PM9	Stage-1 calorimeter trigger hardware tested and installed	PM21
M3.2	PM18	Stage-2 calorimeter trigger hardware tested and installed	PM28
M3.3	PM23	Stage-1 calorimeter trigger commissioned & system ready for physics	PM27
M3.4	PM30	Post-LS3 trigger dataflow design completed	PM30
M3.5	PM35	Stage-2 calorimeter trigger commissioned & system ready for physics	PM35
M3.6	PM54	Post-LS3 trigger prototype trigger modules produced and tested	PM54
M3.7	PM66	Demonstration of post-LS3 trigger slice	PM66
M3.8	PM72	Post-LS3 trigger construction plan delivered	PM72

12 Glossary

Following the request at a previous meeting, we compiled a list of acronyms in common use in the report, or during the oral session, or by CMS which we may have referred to.

AM	Associative Memory.
AMC13	A μTCA data concentration and clock distribution card specific to CMS.
AMC	Advanced Mezzanine Card (from the ATCA specification).
APD	Avalanche Photodiode.
ASIC	Application Specific Integrated Circuit.
ATCA	Advanced Telecommunications Architecture.
BER	Bit Error Rate.
BX	Bunch crossing.
CBC(x)	CMS Binary Chip, version x, for the front-end ASIC for the outer tracker
cDAQ	Central Data Acquisition.
CMSSW	Compact Muon Solenoid Software, is the CMS experiment software package.
CPM	Central Partition Manager.
CPU	Central Processing Unit.
CRC	Cyclical-redundancy check, a family of algorithms for identifying data corruption.
CTP7	Calorimeter Trigger Processor 7 card, featuring the Xilinx Virtex-7 FPGA.
DAQ	Data Acquisition.
DAQ2	Upgrade to DAQ system during LS1.
DSP	Digital Signal Processor.
DTC	Data, Trigger and Control board
DPG	Detector Performance Group.
FB	Finance Board.
FC7	FMC Carrier Xilinx Kintex 7, a processor board hosting multiple FMCs.
FED	Front End Driver, a CMS data acquisition board.
FMC	FPGA Mezzanine Card, ANSI/VITA standard for cards which interface to FPGAs.
FPGA	Field-Programmable Gate Array.
FSM	Finite State Machine.
GBT	Gigabit Transceiver Project at CERN.
GBTX	Gigabit Transceiver ASIC developed at CERN.
GCT	Global Calorimeter Trigger.
GLIB	General purpose μTCA card developed by the CERN microelectronics group.
GMT	Global Muon Trigger.
GP	Geometric Processor.
GT	Level 1 Global Trigger.
GTX	A version of the Xilinx high speed serial transceiver, found on the Virtex 7 FPGA.
HDL	Hardware Description Language.
HE	Endcap Hadron Calorimeter.
HF	Forward Hadron Calorimeter.
HGCal	High Granularity Calorimeter, the proposed new endcap CMS calorimeter.
HI	Heavy Ions, at the LHC refers to collisions between lead ions.
HL-LHC	High Luminosity LHC, the planned upgrade of the LHC machine around 2023.
HLT	High Level Trigger, a collection of software trigger algorithms.
HT	Hough Transform Processor.
I2C	Inter-Integrated Circuit chip-to-chip communications protocol.
IB	Institution Board.
IPbus	A protocol to control and communicate with Ethernet-attached xTCA hardware.
IPMI	Intelligent Platform Management Interface, a standardised computer system interface.
JTAG	Joint Test Action Group; test and diagnostic bus standard by IEEE1149.1.
L1A	Level-1 Accept.
LP-GBT	Low power GBT.
LS1	Long Shutdown 1, first LHC long shutdown from beginning 2013 to end of 2014.

LS2	Long Shutdown 2, second LHC long shutdown scheduled for around 2018.
LS3	Long Shutdown 3, third LHC long shutdown scheduled for around 2022.
MGPA	Multi-Gain Preamplifier ASIC, used to readout ECAL photosensors.
MIP	Minimum Ionising Particle
MMC	Mezzanine Management Controller, part of the μTCA specification.
MP7	Master Processor 7 card, featuring the Xilinx FPGA Virtex-7 chip.
MTF7	Muon Track Finder 7 card, featuring the Xilinx FPGA Virtex-7 chip.
MPW	Multi Project Wafer manufacturing submission, for CMOS ASIC production.
μGT	Micro Global Trigger.
μHAL	Micro Hardware Abstraction Layer.
μHTR	Micro HCAL Trigger and Readout Card.
μTCA	Micro Telecommunications Computing Architecture.
O2O	Software to simplify the propagation of configuration online.
oRM	Optical Receiver Mezzanines.
oRSC	Optical Regional Summary Card
oSLB	Optical Synchronization and Link Boards.
PCIe	Peripheral Component Interconnect Express, a high-speed serial computer bus.
SBS	Shared Business Services.
SerDes	Serialiser/Deserialiser chip.
SFP	Small Form-factor Pluggable standard for optical and other transceivers.
SFP+	Extension of the SFP standard to support up to 10 Gbps data rates.
SLINK	CERN specification for an easy-to-use FIFO-like data-link.
TCC	Trigger Concentrator Card.
TCDS	Trigger Control and Distribution System.
TFP	Track Finder Processor.
TMT	Time-Multiplexed Trigger, that processes events in parallel rather than sequentially.
TMTT	Time-Multiplexed Track Trigger
TPG	Trigger Primitive Generator.
TriDAS	Trigger and DAQ.
TTC	Trigger Timing and Control, a system for distribution of clocking and control.
UCG	Upgrade Cost Group.
uHTR	μTCA HCAL Trigger and Readout card.
VHDL	VHSIC Hardware Description Language
VTRX	Versatile Link Transmitter/Receiver, optical transceiver developed by CERN.
VTTx	Versatile Link Dual Transmitter, optical transmitted developed by CERN.
XDAQ	Cross DAQ, a data acquisition software framework.
YETS	Year-End Technical Stop, a brief stop of the LHC during the winter holidays.