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 3
2 Project history and recent developments 3
2.1 LHC operations and progress 3
2.2 CMS planning 4
2.3 UK adaptation to CMS planning 5
3 Work Package 1: Management 6
4 Work Package 2: Outer Tracker Readout 7
4.1 Objectives 7
4.2 Progress 7
4.3 CBC3 test results 8
4.3.1 Further test plans 10
4.4 Staff on project 11
4.5 Expenditure 11
4.6 Deliverables 11
5 Work Package 3: Level-1 Trigger 13
5.1 Objectives 13
5.2 Progress to date 13
5.2.2 Next generation hardware development 13
5.3 Overview of CMS plans 14
5.4 Staff on project 15
5.5 Expenditure 15
5.6 Deliverables 15
6 Work Package 4: High Granularity Calorimeter 17
6.1 Objectives 17
6.2 Progress to date 17
6.3 Overview of CMS plans 19
6.4 Staff on project 19
6.5 Expenditure 19
6.6 Deliverables 20
7 Work Package 5: L1 track finder 21
7.1 Objectives 21
7.2 Progress to date 21
7.3 Overview of CMS plans 23
7.4 Staff on project 23
7.5 Expenditure 24
7.6 Deliverables 24
8 Risk register 25
9 Finances 25
10 Gantt charts 26
11 Milestones 27
12 Glossary 28

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. The Phase I upgrade of the L1 calorimeter trigger was completed in 2016.

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

- The latest version of the CBC ASIC (CBC3) was returned and lab testing began in November. A couple of minor features requiring revision have been identified but the chip is performing very well.
- Total dose irradiations of the CBC3 have been carried out at CERN at several dose rates and temperatures, to allow inference of performance under future LHC operating conditions. Expected small effects were observed, but there are no concerns about radiation tolerance.
- Preparations are under way for a CBC3 SEU test in May.
- Most of the work on the hardware demonstrator of the Time Multiplexed Track Trigger (TMTT) for the L1 track-finder was concluded very successfully in December. The objectives of measuring the latency and demonstrating high track finding efficiency were achieved. A paper has been completed reporting the system and results and is presently undergoing CMS internal review.
- The three demonstrators were reviewed by a CMS internal expert committee in December; the report is included in the documents for this meeting. The most important conclusions are that the TMTT demonstrator was the only one which completely achieved the objectives set, and that a low-risk implementation of the L1 track finder can be built using only FPGAs as the UK proposed.
- The review of the alternatives for the L1 track-finder has not yet concluded and has occupied much of our time since December.
- The upgraded Level-1 trigger ran smoothly during 2016, with excellent results; this is now considered to be largely a CMS M&O project.
- Progress with next generation trigger hardware has been slower than hoped but no serious problems have been encountered.
- There has been further good progress across the HGCal project, in the areas with UK involvement of front-end electronics, trigger design and simulation and performance studies.

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, WP4 on the high granularity forward calorimeter project and WP5 on L1 track-finding.

2.1 LHC operations and progress

Since the last report data taking during 2016 was concluded. Both the LHC and CMS operated extremely well.

- CMS recorded 38.3 fb^{-1} of good quality p-p data at 13 TeV centre of mass energy from the 41.1 fb^{-1} delivered by the machine (Figure 2.1);

- CMS switched to heavy ion running in November and recorded 186 nb^{-1} of p-Pb data at 8.16 TeV/nucleon;
- As reported last time, no further problems with the CMS magnet were experienced;
- An Extended Year-End Technical Stop (EYETS) started on 5 December for 21 weeks to allow the installation of the upgraded CMS Pixel Tracker, which is now complete.
- Most other upgrades foreseen for the EYETS, such as photodetectors and readout for HF, improved luminosity monitors, and installation of demonstrator GEM muon detectors, are either complete or well on schedule. The only exception is the HE upgrade, where after an extensive readiness review, CMS decided on the reduced goal of installing 1/36 of the upgrade, which is now planned for the YETS in 2017/18.
- CMS is on schedule to be ready for beam, as planned, by 1 May.

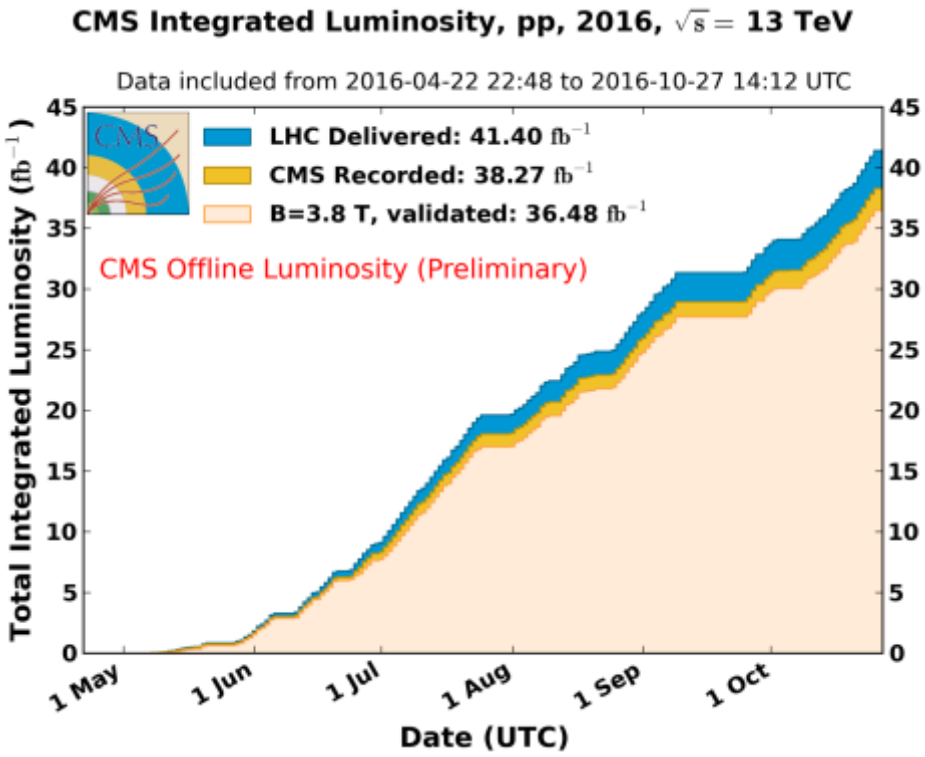

Figure 2.1. LHC and CMS integrated luminosity status in 2016.

2.2 CMS planning

The Phase I upgrade to CMS, which included the UK calorimeter trigger upgrade, is now essentially complete, on time and within budget, except for the HE installation already mentioned. A summary of the situation is to be presented in the April 2017 RRB. The document [CERN-RRB-2017-033] also includes a summary of activities underway for the Phase II upgrade, some of which is reported below.

The Phase II upgrade will start at the beginning of 2019 and is expected to be largely completed in 2026. Proposals for sub-detector upgrades have been reviewed by the LHCC, based on the Technical Proposal [CERN-LHCC-2015-010].

The financing of the Phase II upgrade project is agreed to be a common responsibility of the CMS collaboration. The overall sharing of costs is based on the principle of equity as defined by the proportion of PhD physicist authors supported by each funding agency, similarly to defining the sharing of M&O A costs.

The first of the Phase II projects is the Gas Electron Multiplier (GEM) GE 1/1 project described in a TDR [CERN-LHCC- 2015-012]. The GEM GE 1/1 installation is foreseen ahead of other Phase II projects and should take place in Long Shutdown 2.

CMS is preparing Technical Design Reports (TDRs) for the remaining projects, to provide details on final technical choices and timelines for execution. The Tracker TDR is planned for mid-2017, with the others following by the end of 2017. Interim documents for L1 trigger, DAQ/HLT, Infrastructure and Common Items will be provided to the LHCC and Upgrade Cost Group in the same period, as TDRs for the L1 trigger and DAQ/HLT are foreseen only between 2019 and 2022.

Regular contacts of CMS Management with funding agencies are taking place via link-persons to the Finance Board. The information obtained so far is considered to be reassuring and, although related discussions at national level are at different stages, overall funding is thought to be assured. A growing number of agencies have earmarked funds to cover their participation. Some internal discussions are taking place to try to redirect funds from projects which might be overfunded to others that might experience problems.

The infrastructure and common systems upgrade will be secured through a Common Fund; a list of items with estimated costs is in the Technical Proposal. It includes magnet power and cryogenics, infrastructure, test facilities, surface facilities, safety systems, electronics integration, engineering integration and technical support. The estimated Common Fund value is 25 MCHF. It will be managed by the CMS Resources Manager and CMS Technical Coordinator, with advice from CMS Management; all expenditure will be reported to the RRB.

2.3 UK adaptation to CMS planning

As reported last November, STFC requested from CMS UK an outline construction plan for the period 2019-2024, which was submitted in September 2016. We received positive feedback from Science Board (the letter is included in the documentation for this meeting).

Science Board noted that the Phase II Upgrade proposal had well defined goals and the case was scientifically excellent, building coherently on previous R&D in strategically important areas with considerable potential for UK leadership and impact.

Science Board noted that CMS UK responded positively to the funding challenges and proposed a significantly focused programme. It was noted that difficult decisions had been taken to rebalance the programme within the 2014 planning guideline whilst maintaining a technically challenging, interesting and relevant project. They noted that further reductions to the programme would result in significant loss of leadership in strategically important areas and recommended funding as requested.

Science Board agreed that the CMS-UK Phase II Upgrade proposal could proceed as submitted to the 2016 review, but it also stressed that this does not guarantee funding for the project, and any award will be dependent on peer review of a full proposal and the future financial scenario. A Statement of Interest was submitted to Science Board on that basis, and further input has recently been requested.

As also discussed in the last report, a significant uncertainty for our planning is a decision on the L1 track-finding processing architecture, where three alternatives, one being a UK time-multiplexed design, were under consideration, with demonstrators in construction. This is reported later.

The second review took place in early December, and the report is included with the documents. It was clear both from oral feedback and the report itself that the UK developments were considered by far the most successful and that an all-FPGA solution was considered to be the safest and least costly route to follow. A couple of tracker management board meetings have since taken place and the situation remains controversial. A task force was set up to discuss the technical implications and a recommendation has been made to adopt an all-FPGA approach, and to build the collaboration to develop that. This will be further discussed in a tracker management meeting on 26 April, although it is expected that any decision will require at least one more meeting.

3 Work Package 1: Management

A reminder of the project management is included below. G. Hall remains PI for this project; since January 2017 he has been working 50%, with almost all his time devoted to the upgrade project. D. Newbold has been UK CMS PI since October 2015. A. Tapper has been CMS L1 trigger project manager since September 2016.

At the last meeting, we were asked about assignments of individuals to work packages, in view of the new WPs. A separate spreadsheet has been produced accompanying the documents for this meeting which attempts to provide a snapshot of present and recent contributors to each activity; it therefore does not list all those who contributed to the now complete L1 trigger upgrade.

At Imperial, some staff changes are soon to take place to optimise the use of funds and ensure future continuity when the present grant ends. M. Pesaresi will move to the CG, as will J. Borg, who will replace M. Raymond on his retirement.

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	M Raymond	Imperial	Overall responsible for CBC specifications, interface to module design team, chip testing and module evaluation and CMS planning
	M Prydderch	RAL TD	Manager of ASIC design team in RAL
3	G Iles	Imperial	Based in CERN with responsibilities for operation of L1 calorimeter trigger.
	J Brooke	Bristol	Supervision of UK trigger upgrade activities.
4	P Dauncey	Imperial	Jointly coordinating the UK HGCal developments, with G Davies. L1 trigger coordinator for the HGCal.
	G Davies	Imperial	Also, UK representative on the provisional HGCal IB and FB.
5	M Pesaresi	Imperial	Coordinating demonstrator integration activities and hardware/firmware specifications, and general project planning.
	I Tomalin	RAL PPD	Maintaining and running the Monte Carlo analysis software, improving tracking algorithms based on offline experience and day-to-day oversight of RAL track trigger 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 final prototype of the 130 nm CMOS bump-bondable front end readout chip for the 2S modules in the outer silicon tracker. The previous CBC2 version of the chip was very successful, allowing R&D on hybrid technologies suitable for the construction of 2S and PS modules to progress, and proof-of-principle demonstration of the p_T stub triggering approach in beam tests. CBC2 chips continue to be used for hybrid construction and performance studies.

As discussed last time, the CBC3 contains an optimized front end and new logic to produce stub addresses with half-strip resolution by assigning 2 and 4 strip seed clusters to the mid-position between two strips. Additional bend information on the stub direction is produced using the position of a correlating cluster in the window layer. Stub address and bend information for up to three stubs can be passed off-chip within a single 25 ns bunch crossing period on five outputs running at 320 Mbps. The CBC3 pipeline is 512 cells deep allowing a L1 trigger latency of up to 12.8 µs. Following a L1 trigger the unsparsified pipeline data are transmitted off-chip on a dedicated 320 Mbps output, in a 950 ns long data frame allowing an average trigger rate of up to 1 MHz.

The CBC3 was submitted for manufacture in a joint run with another chip, to share costs, in July 2016. At the time of the last report, delivery of the first wafers was imminent. Six wafers with wire-bondable pad finish were delivered and one was sent for dicing, so we have had chips under test since November. A study of performance under ionizing radiation has begun, and a setup for single event upset performance is ready. Five wafers have been probe-tested and sent to a company to have solder bumps deposited and subsequently diced.

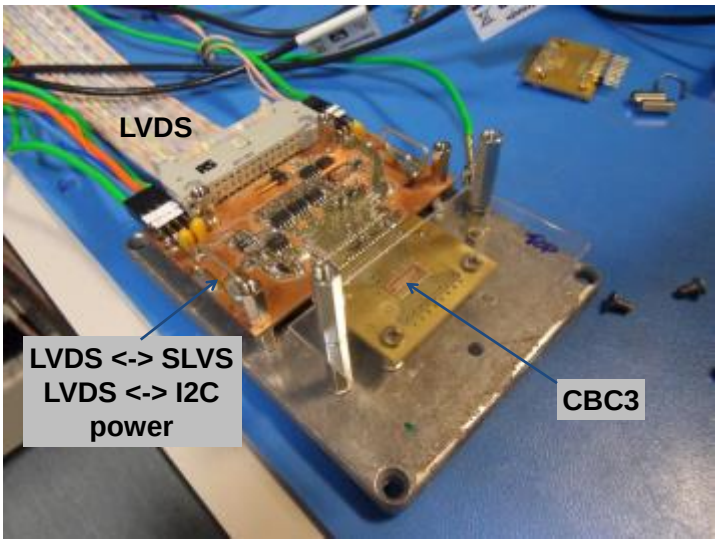

Figure 4.1. CBC3 chip in single chip test setup

4.3 CBC3 test results

Figure 4.1 shows the single CBC3 chip test setup. The chip can be powered, controlled and read out via a column of wire-bondable pads at the back edge. The chip is mounted and bonded face-up on a small carrier board which is plugged into a support card which interfaces power and provides buffering of input and output signals, converting SLVS (low voltage differential standard) signals to standard LVDS levels on a twisted pair cable connected to the DAQ and control system.

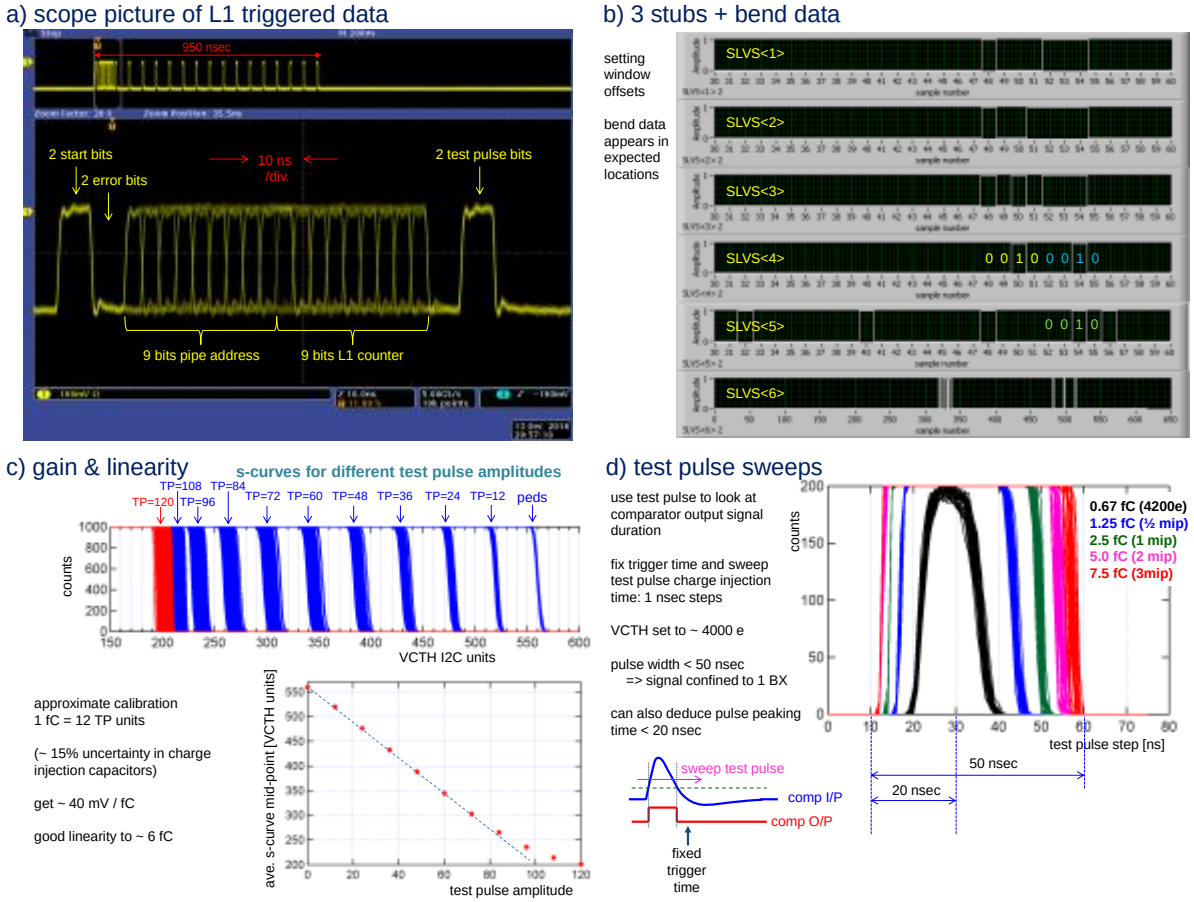

Figure 4.2. CBC3 single chip test results. a) Scope picture of L1 triggered data shows a complete data frame with test pulse active, and zoomed in picture of the digital header. b) CBC3 output data acquired by DAQ system shows stub and bend data samples in top 5 panels, together with subsequent L1 triggered data in bottom panel. c) S-curve data, acquired by sweeping global comparator threshold setting VCTH, for different test pulse amplitudes, together with average gain vs. test pulse amplitude plot. d) Front end pulse duration measurements obtained by sweeping the test pulse charge injection time for various signal amplitudes.

Figure 4.2 shows a selection of results from the single chip test setup, illustrating some of the CBC3 functionalities, taken from talks given at CMS outer tracker electronics systems[1] and tracker week[2] meetings. Figure 4.2(a) shows an oscilloscope picture of the L1 triggered output of the CBC3. The on-chip test pulse can be programmed to fire all 254 channels in 7 groups of 32 and 1 of 30, and the top panel shows the resulting data frame with a digital header followed by the channel data where the 32 channels (16 pairs) fired by one of the test pulse groups can be seen active. The lower

[1] http://www.hep.ph.ic.ac.uk/~dmray/systems_talks/2016/CBC3_first_results_Dec_2016.pdf
[2] http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/CBC3_status_March_2017.pdf

panel shows a zoom of the digital header portion of the frame, consisting of two start bits (always high), followed by two error bits (activated by the chip itself if specific error conditions are detected), followed by two 9-bit pipe address and L1 counter fields. The chip is pseudo-randomly triggered so these fields show all bits active with the scope operated in persistence mode. The pipe address field contains the pipeline column address in which the L1 data were stored. The L1 counter is incremented by the L1 trigger and its value, at the time of the trigger which gave rise to the data in the output frame, is returned in the counter field. All CBC chips in the experiment will be operating synchronously, so all should return the same values of pipe address and L1 counter in response to a given trigger. This is a strong check that every chip in the system is operating correctly.

Figure 4.2(b) shows all 6 data outputs of the CBC3 after acquisition by the DAQ system. The panel labelled SLVS<6> shows the L1 triggered data. Panels labelled SLVS<1> to SLVS<5> show the stub data. A test pulse group has been used to generate only three pairs of hit channels by masking all other channels (on-chip mask registers) and only allowing outputs from those channels through to the stub finding logic. The test pulse wiring is such that consecutive pairs of channels in the output data correspond to different layers (seed or window) in the stub finding logic, so that the test pulse can exercise that logic. The panels corresponding to stub data outputs SLVS<1>, <2> and <3> show the 8 bit stub addresses. SLVS<4> and <5> contain the bend information (b0010 for all three stubs). SLVS<5> also produces a synchronization pulse in every eighth bit position to indicate the start of a new 8-bit data field for the stub data lines.

Figure 4.2(c) shows front end gain and linearity measurements acquired by sweeping the global comparator threshold VCTH for different values of test pulse amplitude to generate families of s-curves for all channels on a chip. The picture shows well-behaved s-curves and good linearity, even though linearity is not a strong requirement for a binary chip. Note that the increasing spread in s-curve response with test pulse amplitude reflects the unavoidable and expected spread in the very small test charge injection capacitors.

Figure 4.2(d) shows a measurement of the front end performance obtained by sweeping the time of test charge injection while keeping the readout trigger time fixed, which effectively gives the duration of the comparator output pulse. For a realistic operating threshold of 4000 electrons the comparator output pulse duration is below 50 ns for normal amplitude signals, so single bunch crossing resolution is achievable, which was one of the requirements for the CBC3 front end optimization.

It is not expected that best performance can be achieved for a chip that has been designed for bump-bonding, but the selection of results in figure 4 already indicate satisfactory operation in the single chip test setup, allowing more detailed tests of digital functionality to be performed.

Some faults in the digital logic of the CBC3 have been discovered. Some mistakes in the Verilog description of the stub finding logic result in a few incorrect stub addresses and wrong bend information being output. This can be easily fixed in the final production chip and will not be difficult to correct for in tests using the CBC3. Another digital malfunction has been discovered which is associated with the programmable DLL required to synchronize the chip to the particle arrival time in the final experiment. For some DLL settings the L1 triggered readout data gets corrupted if a trigger arrives during the readout period of a previous trigger. Most DLL settings are unaffected, and these “good” values are consistent from chip to chip, so this problem should not adversely affect the use of the CBC3. The problem has been traced to the control logic associated with retrieving data from the buffer memory used to store data awaiting readout to the output serializer. Work has begun on circuit redesign to rectify the problem for the final production version of the chip.

Six wafers arrived in the first delivery, of which one was diced. Once satisfactory functionality was confirmed we purchased an additional three wafers remaining at the foundry. Five wafers have been probe-tested and sent to a company to have bumps deposited and subsequently diced.

The wafer probe tests were developed using the single chip test setup of figure 4.1 since the wirebond pads are those which are probed during wafer test. It has been found feasible to run the chip at full speed on the wafer. A test protocol has been developed which includes extensive tests of digital functionality including stuck bits in registers or pipeline or buffer memory, and exhaustive

tests of the stub-finding logic. Analogue tests included spread of offsets after tuning, gain uniformity, and bias current and voltage register sweeps. Results for the five wafers sent for bumping can be seen in figure 4.3. There is a small but noticeable patch of failing chips in the centres of the wafers, but the average yield of chips passing all tests is 85%, which is high.

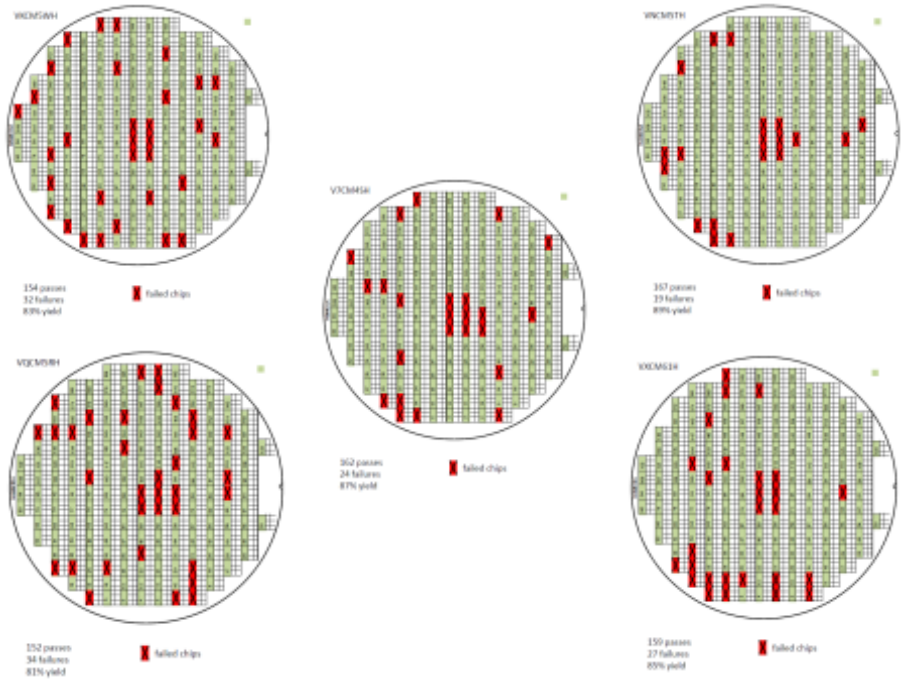

Figure 4.3. Wafer maps for first five CBC3 wafers probed.

A first ionizing radiation test of the CBC3 using X-rays was carried out in March. For the CBC2 an increase in the digital power consumption was observed, which peaked in the 1-2 Mrad region, decaying away (annealing) when irradiation was interrupted, and falling back to pre-irradiation values at higher doses. It was traced to the pipeline region of the chip, where non-enclosed NMOS transistors were used. While it was doubtful whether the extra power consumption would be visible at the low dose-rates in the experiment it was decided to adopt the use of enclosed NMOS devices in the CBC3 pipeline to remove the effect altogether. Nevertheless the CBC3 irradiation has shown some similar behaviour, which we now think is likely be occurring in the extensive additional synthesized digital logic in the CBC3 where non-enclosed NMOS devices are used in the library components. While preliminary analysis indicates that the digital current increase will be negligible in the experiment, we are following up with more tests to confirm.

The CBC3 SEU immunity has been designed to improve on that exhibited by the CBC2, where some vulnerability was found in the I2C registers. A test setup has been developed for this, and we are ready to conduct the test using the Cyclotron facility at UCL (Louvain). This has been delayed due to a problem with the accelerator and we are currently awaiting re-scheduling.

While the prototype and wafer testing has been carried out using a legacy VME-based DAQ system, the ionizing and SEU test setups have been read out using FC7 hardware with appropriate firmware and software. It is currently planned to replace the wafer probe readout with an FC7 based system where significant improvement in speed, necessary for the final production, can be expected.

4.3.1 Further test plans

In the short term testing plans presented last time we assumed that CBC3 wafers would be sent immediately for bump-bond processing (bumping) once single chip studies had confirmed functionality, with wafer-testing carried out afterwards, before dicing. The sequence of operations has been modified, with wafer testing performed first, because the bumping and dicing will be carried out

by the same company, and so the need for the wafers to travel back and forth between these operations is avoided. CERN collaborators have responsibility for the selection of a suitable company to carry out the bump-bond processing, since this is a common technology which is required for other chips in the 2S and PS systems, and PacTech have been chosen for the CBC3 wafers. They received the wafers at the beginning of April, and we are expecting chips in June.

As explained last time, the plan is to populate a small 2-chip hybrid in the first instance, allowing a small-scale but fully functional module to be constructed, which can be lab-tested and should also be ready for a beam test in October. The hybrid and interface electronics is also the responsibility of the CERN team and is currently under design.

4.4 Staff on project

Mark Raymond will retire in July 2017. Johan Borg, who has proven his expertise since joining Imperial in an ERC-funded HGCal post, will replace him.

4.5 Expenditure

As usual the main recent expenditure has continued to be on RAL TD staff, and the cost in the last financial year was larger than expected to complete all the design changes, most of which were driven by system choices. It became clear in early 2016 that the modifications had generated more work than anticipated. Additional TD staff effort is required to avoid risk to the remaining work; this is discussed further in the Finance section.

4.6 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.

WBS item 2.4.1 (early CBC3 tests) can now be considered to be complete. We are in the ongoing testing phase 2.4.2 where the ionizing irradiation tests are already underway and the SEU test will be done as soon as the facility becomes available. CBC3 2S-PT module studies (2.4.3) will begin in the summer.

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. We now know that the CBC3 is fully functional and performing well enough to proceed, but we also know that another iteration is required to fix the few issues so far identified. At the moment these changes are confined to the digital logic only. It is possible that the CBC4 iteration could 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 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 already begun on the circuit re-design 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 decision would be at the end of this year.

The new WBS allows a longer period (until the end of this year) to finalise any design changes required for the CBC4. The period allowed for production is left unchanged, but the period allowed to verify that the changes have worked has been halved to three months, since if the changes are to the digital logic only it should be possible to confirm correct functionality quickly.

The CBC4 mass production preparations item 2.6 has also been modified. The period allowed for any final mask preparations is the same, but has been overlapped with the CBC4 prototype test period (in principle, if the design changes have worked then the final masks should not need to be changed at all). The production period remains the same but now extends right up to the end of the project (previously it ended in December 2018). Final verification tests ought to be quick, but will now take place in the next phase of the CMS UK Upgrade programme.

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 SS-Pt module prototype studies	04/13	03/15	24	a programme of SS-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 SS-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	03/18	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

The Phase I upgrade to the Level-1 trigger has been running successfully since May 2016. There was steady improvement in algorithms and configuration to optimise data quality through the first months of the 2016 LHC run. During 2016 the Level-1 trigger contributed to less than 10% of the small amount of collision data that CMS failed to record. In addition, only around 3% of the small amount of data recorded by CMS but later marked as bad, due to detector problems, was as a result of the Level-1 trigger. This level of reliability and high quality data from a completely new system shows the wisdom of the decision to commission it in parallel during the 2015 run in parallel with the old system.

During the extended end of year shutdown, the opportunity was taken to improve robustness of the trigger algorithms in the presence of pileup. In particular, pileup mitigation was included in the missing energy calculation, which had seen significant rate increases as pileup increased during 2016. Results obtained with existing data indicate that the new missing energy trigger will be highly robust against pileup in 2017. In addition, substantial effort has been devoted to offline data quality monitoring. For the coming run, comparisons of online and offline reconstruction, including resolution and efficiency will be produced automatically, as well as detailed comparisons of the hardware output with bit-level emulation, providing real time hardware diagnostics.

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. Development is taking place along two paths: the **MP-Ultra** - a small, low cost, form factor (PCIe) card that will allow us to validate new technologies (e.g. PCB build-up, optical engines, FPGA, etc.) and the **Service card** - a full size ATCA card that will allow us to test and validate the core services (e.g. embedded CPU, power, signal integrity, thermal aspects, etc.).

Since the last review the MP-Ultra card infrastructure has been validated. The full assembly of three further cards, including Xilinx Ultrascale FPGA, was requested in November. Cards were shipped in February, significantly later than the preliminary planned December shipment date. From visual inspection, the quality of the assembly appears to be very good. Testing is now progressing in earnest, but was delayed by approximately six weeks due to illness. Programming and debug access to the main FPGA is now available, PCIe communication is being established, with RAM and optical interface testing to follow.

The Service card was submitted for PCB manufacture in February, whereas we had originally targeted a submission before the year end. The additional time was required to correct several errors detected during the pre-submission review. PCB production proceeded as planned and delivery included a full set of quality assurance documentation. Assembled cards were to be delivered mid-April, but has been delayed by non-availability of 25 Gbps backplane connectors. These are now forecast for delivery in May, with assembly and delivery a couple of weeks later. The requirement for these specialised connectors is driven by a wish in the collaboration to run DAQ over the backplane, but at 25Gbps this may be misguided. If there are further delays, we will proceed without the

connectors, and attempt to have them fitted later. This will enable 90% of the testing to be completed. The MP-Ultra can be mounted on the Service card.

Recently, as Xilinx characterise their new Ultrascale+ parts, it has become possible to start making reasonably accurate power predictions. These suggest that FPGA power may increase from 40W for the MP7 to up to 130W for the largest parts, running at the highest clock frequencies. In light of this we have been reviewing power and cooling delivery within ATCA crates, the proposed standard for CMS. In essence, ATCA provides extra power compared to MicroTCA (400W vs 80-100W), but additional cooling is limited to the use of more powerful fans. The height available for heatsinks remains unchanged and while the card depth increases from 18 cm to 28 cm some of this real estate is used to bring in power and reduce it from 48V to 12V. The net effect is that we may need to operate crates with high airspeeds, which may generate significant noise (in excess of 85 dBA) and require additional power for fans. Alternatively, we operate the components at temperatures close to their design limit (e.g. 100°C for FPGAs), but in this case, we must understand the potential impact on reliability.

The full ramifications are still being assessed, but we need to make sure that the full system, which will involve many crates, is a viable option for the underground service cavern.

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 the interface between the sub-detectors and the central trigger system. These interfaces are currently being finalised, and will be documented in an Interim Document, to be submitted to the LHCC in Q3 2017. Some concerns arise due to inhomogeneity of this interface between sub-detectors, which will need to be resolved over the next two months. Although the design of the Phase II trigger system will not be finalised until the TDR (due in Q4 2019), two possible architectures are under study and will be described in the Interim Document. One of these uses the time-multiplexing technique developed for the Phase I Calorimeter upgrade, which has also been used for both FPGA-based Track Trigger systems, and in the AM proposal. Finally, the document will describe progress that has been made recently in object identification techniques at Level-1, in particular the use of “particle-flow” inspired reconstruction.

Completely new back-end electronics for high luminosity is essential. It is necessary to handle higher bandwidth interfaces to the front-end electronics within CMS, and to process the resulting data for delivery to the DAQ system and to extract a trigger signal. This problem materialises in several places in the upgraded CMS, including tracker, HGCal and L1 trigger. Therefore it is natural to ask if common solutions are possible. We and CERN have particularly drawn attention to this question, and it is included in our long term plan for the UK upgrade construction contributions.

We intend to build on our successful R&D accomplishments, of general purpose state-of-the-art digital electronic boards accompanied by comprehensive firmware and software, much of which has been adopted throughout CMS. We are among the most experienced developers of advanced programmable logic and high speed optical technology.

Therefore we have begun to view the Ultrascale development as a step towards common hardware, rather than targeted solely at the trigger, which was the picture at the time of the proposal in 2012. We propose its evolution towards a generic, programmable processing card with high speed optical I/O, which will then be used for applications proposed in all UK work packages. We have been in discussions with CERN and a few CMS technical experts about how this might happen, but it is still at an early stage. It has not yet evolved into a project, but should do so in 2017, once clarification of the future of the L1 track finder activity has taken place.

It is not expected that a single board will satisfy all objectives; however, adaptations into a small family of very similar boards should meet the requirements of most projects. In particular, the power system, embedded computing, crate control, optical technology and cooling system will remain unchanged or with slight modifications. They require significant design and testing effort, and the aim will be to share the load with partners at CERN and elsewhere. These parts should be interchangeable so that the base card becomes a low risk object. The infrastructure firmware and core software

libraries should also be applicable to other projects and only require additional modules for covering project-specific functionality.

A possible roadmap for the future is shown below, Figure 5.1.

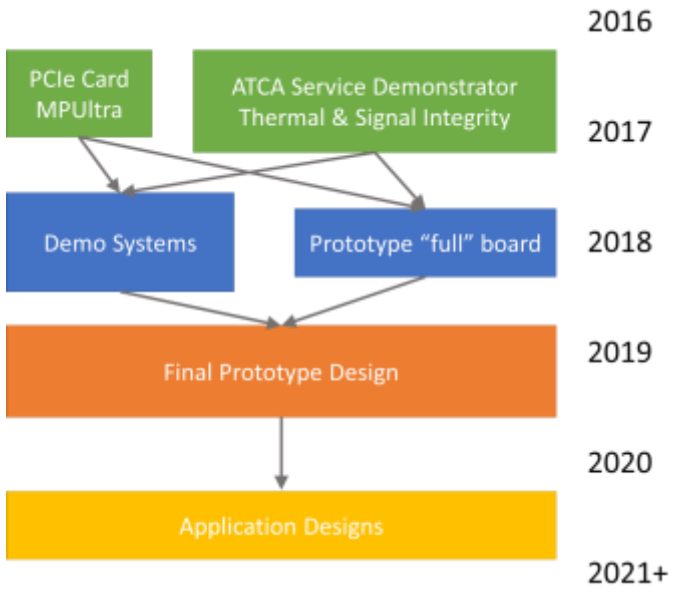

Figure 5.1: Possible roadmap for common hardware developments resulting from the present R&D.

5.4 Staff on project

Sarah Greenwood, the Imperial College electronic technician who became expert at multi-layer digital board layout, retired in March 2017. Duncan Parker joined the group to replace her in January 2017 and has already very successfully begun to contribute.

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.

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 play a leading role in producing the HGCal TDR at the end of 2017, in terms of both overall project management and technical aspects, specifically in the key areas listed below.
- To contribute to the design and testing of the front-end electronics ASIC for the HGCal in collaboration with the Omega/Ecole Polytechnique group in Paris.
- To develop the trigger primitive generator (TPG), including contributing to the algorithms, firmware and common 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

There has been good progress across the HGCal project, including all three areas with UK involvement, namely the front-end electronics, the trigger and simulation and performance studies.

The work on the analogue front-end electronics has proceeded along two paths. During 2016 the main effort was on developing a test system for, and thoroughly evaluating the performance of, the SKIROC2-CMS front-end ASIC designed by the Omega group. The tests show noise and time-of-arrival (TOA) jitter comparable to simulations, as well as channel gain matching and stability comparable to expectations, but a smaller linear range than expected. The lower end of the measurement range of the time-over-threshold (TOT) method for digitizing large signals proved to be higher than expected, leading to a decision to evaluate the TOT circuit relative to the low-gain ADC path (rather than the high-gain path as originally intended). The chip was also found to exhibit substantial time-walk, shown in Figure 6.1(a)., and has a significantly non-linear TOA transfer function. However, it was found that if the temperature of the front-end chip remains within a few degrees the TOA measurement, accuracy can be restored to a level close to the measurement jitter limit, see Figure 6.1(b). These measurements were completed in December 2016, and provided as input into the HGCAL test beam project which will be operating a SKIROC2-CMS-based calorimeter during the summer of 2017.

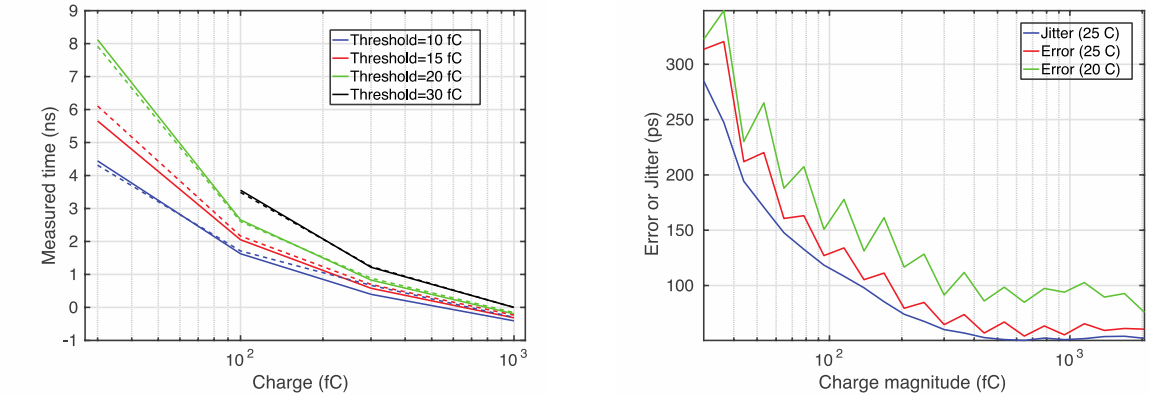

Figure 6.1 (a) Time-walk measured for different TOA threshold settings. (b) Corrected time walk compared to the measured jitter of the TOA circuit. The jagged shape is due to using a coarser resolution when performing the characterization measurement than was used for validation, and applying linear interpolation between the measured characterization levels. The 20˚C graph corresponds to a measurement where the circuit was characterized at 25˚C but the time-walk correction was applied to measurements performed at 20˚C. A TOA threshold of 15 fC was used.

The second aspect of this work is analogue electronics design work aimed at optimizing the electrical performance of the front-end electronics. This work commenced while waiting for the SKIROC2-CMS test boards, but was on hold from mid-July to December 2016. Results include a framework for optimizing charge-preamplifier front-end circuits and hybrid analogue/digital filtering

pulse shapers. The current work is focused on developing a low-power (<1 mW) 50 ps resolution TDC circuit for digitizing the TOT signal based on gated ring-oscillators. This circuit will be included on the V1 HGCal ASIC, which is expected to be submitted in June 2017.

Much of the recent work within the UK on the HGCal trigger has been focused on the creation of the raw data for the trigger primitive generator, which are produced on-detector in the front-end electronics. This is a priority as the design of these electronics will need to be frozen much earlier than the rest of the trigger generator. The front-end electronics architecture was redefined in December 2016, and this significantly affects the transmission of both the event and TPG raw data. The front-end is now based on "panels" which can hold up to 6 large or 12 small wafers and from which all data are transmitted optically directly to the off-detector electronics. The size of the motherboards is limited by the PCB fabrication process but is large enough that the optical components are in a sufficiently low radiation environment that they will survive throughout the HL-LHC running. While this new architecture brings a lot of electrical advantages, it means there is no longer any merging of different areas of the HGCal into one link, so the link occupancies vary by an order of magnitude. The occupancies and resulting bandwidths as a function of the position within the HGCal have therefore been re-evaluated and a new baseline architecture for the off-detector electronics for both the event readout and TPG has been defined, where an example for the event readout is shown in Figure 6.2.

Work on the firmware for the off-detector trigger primitive generator has also made progress. The two-dimensional (2D) clustering per layer is the first of the two main steps of the trigger processing. Preliminary firmware to implement an algorithm for this step has been written and is found to fit within both the FPGA resource and latency requirements. The second step, where the 2D clusters in each layer are combined to give a 3D overall shower, is now being developed in collaboration with the Saclay group.

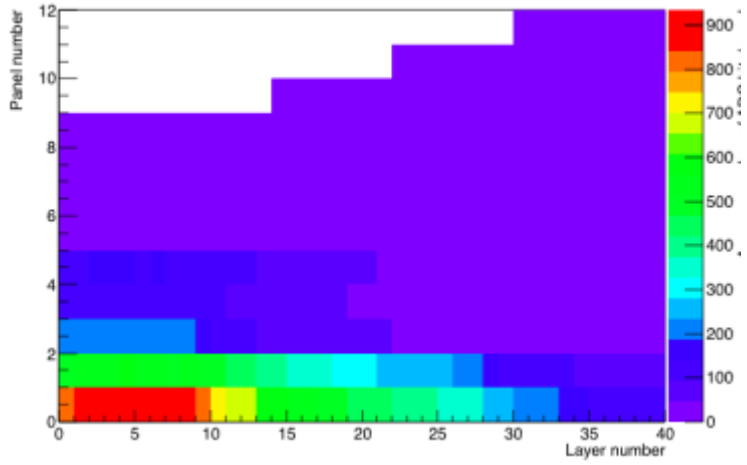

Figure 6.2 Number of channels above threshold per event, resulting from events with 200 minimum-bias interactions, as a function of layer and panel number. The panel numbering is approximately related to distance from the beam line within each layer.

The primary aim of recent work on simulation and reconstruction is to complete a proof-of-principle particle-flow reconstruction for the calorimeter and integrate it into the standard CMS software, obtaining plots and results for the forthcoming TDR; an example is shown in Figure 6.3. Implementation of a reasonable approximation of the technical proposal geometry was completed at the end of 2016.

The two- and three-dimensional clustering has been subjected to a first round of optimization, where the most serious flaws in the original implementation have been corrected. Work is proceeding on the "superclustering", meaning the gathering together of 3D clusters that originate from the same particle and are separated because the particle started showering in the tracker material (in the case of photons and electrons), or because of the irregularity and inhomogeneity of hadron showers (in the case of hadrons). The final step will be to associate these "superclusters" with the charged particle tracks reconstructed in the inner tracker, and feed the resulting entities into the current CMS particle

flow reconstruction of jets and missing transverse energy. At the same time, an alternative reconstruction, where all the above steps of pattern recognition use knowledge of Monte Carlo truth information to gather the simulated energy deposited in the calorimeter, has been put in place to feed identified entities into the current CMS particle-flow reconstruction, and thus provide useful and informative results for comparison. Meanwhile a steady stream of results has been produced using standalone GEANT simulation models of the calorimeter, developed within the UK, to provide guidance for many of the design choices that have been made.

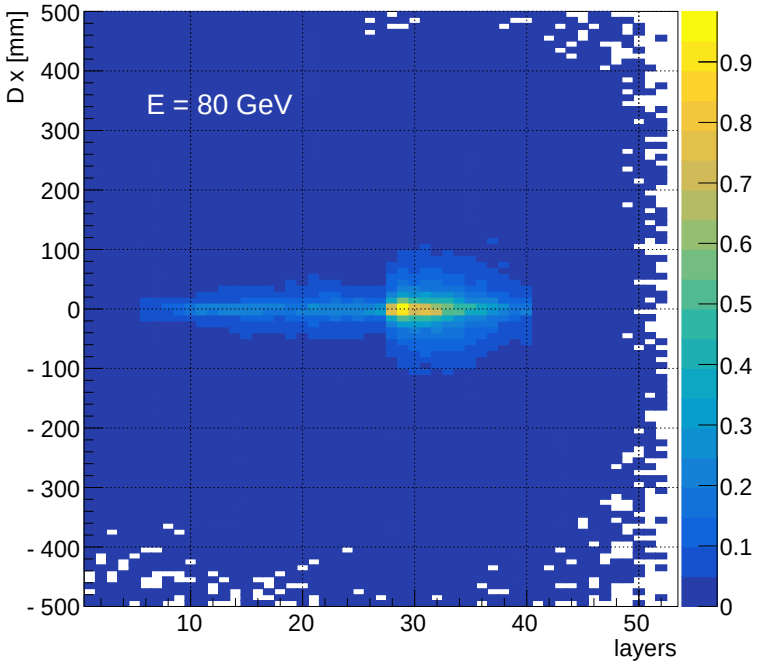

Figure 6.3 Energy deposited by a sample of 80 GeV charged pions in the 52 layers of the HGCal. The distance of the deposited energy from the intersection of the pion direction with the detection layer, in one of the lateral dimensions, is plotted versus the layer number.

6.3 Overview of CMS plans

The most important element of the overall CMS planning for the HGCal is the scheduled Technical Design Report, which is due for submission in November 2017. All the UK work will contribute to that document and UK personnel are expected to play a major role in writing it, including the overall editor (C. Seez).

In addition, the central L1 trigger project will produce an Interim Design Report in June, presenting the concept of the global L1 trigger design. This will ensure that the HGCal (and other upgrade subdetector) TDRs are compatible with the overall CMS trigger design. The UK HGCal trigger effort will interact closely with, and contribute to, the production of this IDR.

6.4 Staff on project

Imperial are in the process of hiring a new RA and firmware physicist-engineer to join the HGCal effort, funded from T. Virdee's ERC grant.

6.5 Expenditure

Expenditure so far has been almost entirely on staff, and on travel to a lesser extent. We do not foresee significant equipment expenditure for some time in the projects with UK involvement. The ASIC fabrication will be funded by the French groups and, for the trigger, the first prototype based on a common hardware board is not expected to be produced until at least 2020.

6.6 Deliverables

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

- The main deliverable for the front-end electronics work is the complete design, testing and characterization of the front-end ASIC for the HGCal. The shorter-term deliverable relevant on the timescale of the current grant is the completion of testing of the SKIROC_CMS2 chip (concluded in December 2016), with work currently under way to prepare a TDC circuit for the HGCAL V1 prototype to be submitted in June 2017. Effort on developing a test board for the HGCAL V1 prototype and evaluating the Imperial-designed TDC implemented in this chip is expected during the autumn of 2017.
- Similarly, the long-term deliverable for the trigger work is to produce a trigger system which will provide a usable trigger from the HGCal. Shorter-term, the aim is to define this design sufficiently for the TDR next year.
- The long term deliverable is particle flow reconstruction that optimally exploits the unprecedented granularity of the HGCal. More immediate deliverables are adequate demonstration of the promise of such reconstruction for the TDR, together with results of studies to assist detailed design choices for the TDR.

HGCal milestones were defined from October 2015 onwards. Upcoming milestones relevant to the UK effort since then and over this year are:

- Submission of the front-end ASIC test structures in March and September 2016 and the V1 ASIC submission in March 2017: The test structures were submitted in April and December 2016, and the V1 ASIC is currently planned for submission in June 2017.
- Baseline definition of the TPG architecture in September 2017: This will be based on common hardware boards which are being developed within WP3. Baseline definitions of the event and TPG raw data formats, and also performance results of the TPG system, both in October 2017: The software and firmware developments are being kept synchronized so that the firmware algorithms are the ones being used for the performance studies. All these milestones are required for the TDR.
- Large scale production of fully simulated and reconstructed events for physics and trigger studies in April 2016 and March 2017. The former has now been achieved and these were used for the trigger and performance studies reported above. A bigger set of samples will be needed for the results to be produced for the TDR and these are now expected in May 2017.

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-2 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

The most significant development since the last report has been the completion of the demonstrator slice, followed by the presentation of the results extracted from hardware at a CMS internal track-finding review.

The demonstrator slice, based at the Tracker Integration Facility at CERN, is designed to process stub data from a region of a future Phase II tracker, reconstructing tracks with p_T > 3GeV/c with maximum efficiency and resolution for use in a Level-1 triggering system, within an overall latency of 4μs. As presented in previous reports, the slice, also known as a Track Finding Processor (TFP), reconstructs tracks within an angular coverage of 1/8 in φ, and the entire tracker acceptance in pseudorapidity. Since the slice is part of a time-multiplexed system, a TFP is also defined to receive data for one event in 36. Therefore, to reconstruct tracks in all events, a fully scaled system would require 288 (8 x 36) processors.

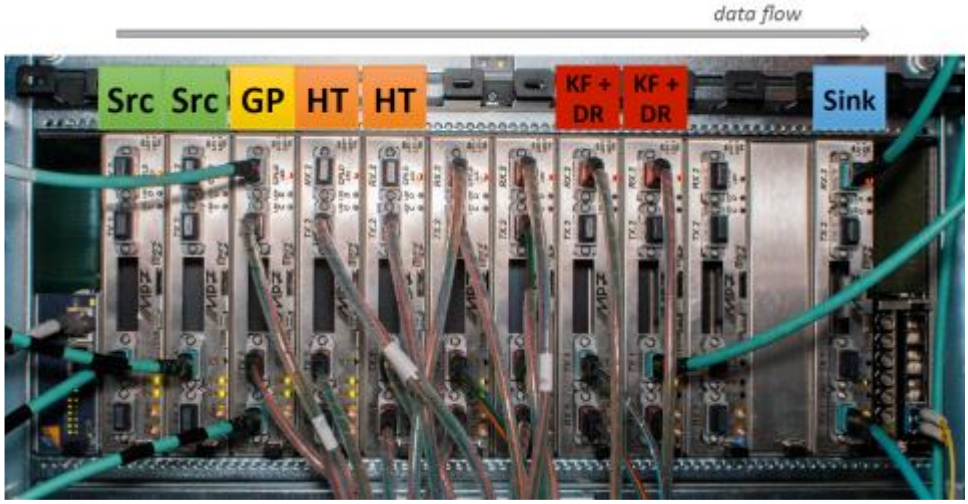

Figure 7.1: The demonstrator slice at CERN as used for the review, consisting of several MP7s in a MicroTCA crate, each corresponding to an element of the Track-Finding Processor.

The demonstrator is shown in Figure 7.1. The TFP is currently implemented on five MP7 boards, which should be replaced by a single more powerful board in a final system by profiting from advances in FPGA technology. Additionally, two *Source* cards are required to take simulated LHC collision tracker data, corresponding to all stubs generated within a tracker octant, and inject them over optical links into the Track Finding Processor. The reconstructed tracks from the TFP are captured in a *Sink* card at the end of the chain. Both Source and Sink cards are implemented on MP7 boards, where data can be injected or read out via IPbus from/to a local PC.

Within the Track Finding Processor, MP7s fulfil the following roles:

- Geometric Processor (GP) – assigns the stubs within an octant to 2×18 regional segments in φ×η, and transmits them on specific links.

- Hough Transform (HT) – coarsely reconstructs tracks in the r-φ plane using a Hough transform technique, grouping stubs into track candidates.
- Kalman Filter & Duplicate Removal (KF+DR) – a candidate cleaning and precision fitting algorithm that removes fake tracks and improves track parameter resolutions, followed by a unique but simple algorithm to eliminate duplicate tracks.

Over the course of the demonstrator development, alternative algorithms have also been pursued and tested on hardware, highlighting the flexibility of our demonstrated system. This includes other methods of filtering and fitting the candidates generated by the Hough Transform. For the review however, only the baseline design above described was presented.

The performance of the demonstrator has been evaluated by injecting simulated LHC events (typically $t\bar{t}$) at up to 200 pile-up into the Source, and comparing the tracks arriving at the Sink against those predicted by our C++ emulator software.

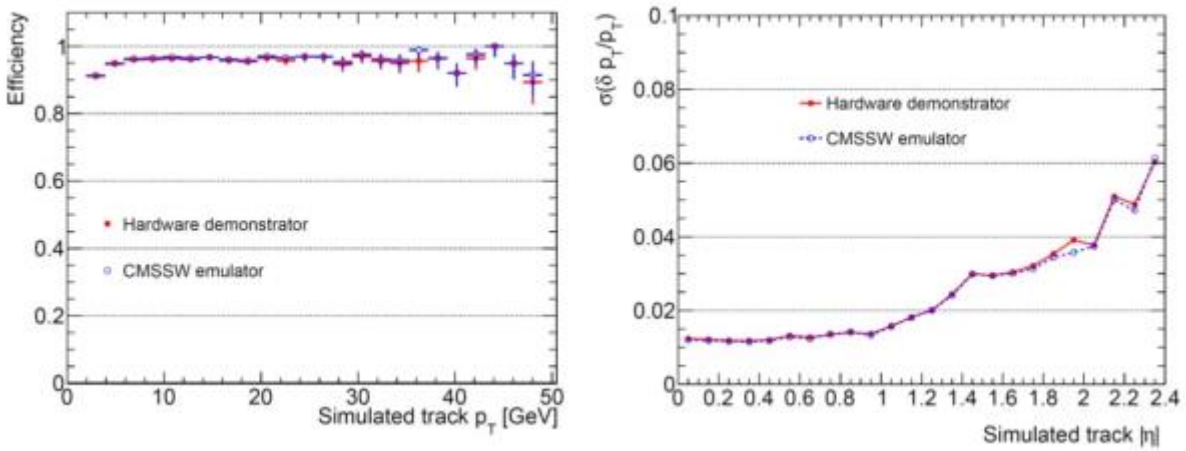

Figure 7.2: Track reconstruction efficiency (left) and relative pT resolution (right), measured in both hardware and emulation, for tracks originating from the primary interaction in $t\bar{t}$ events with 200 pile-up. Only the UK demonstrator was able to show efficiencies as measured in the hardware, in the review.

As shown in Figure 7.2, the overall tracking efficiency for $t\bar{t}$ events with 200 pile-up, as measured in hardware, is good (>94% over the entire acceptance) and agrees well with the emulator (>99.5% overall). The efficiency to reconstruct muons from the $t\bar{t}$ system was measured to be greater than 97% over the entire acceptance.

Figure 7.2 also shows the resolution of one of the four track parameters for reconstructed primary tracks in both hardware and emulation for the same data sample. The level of agreement between hardware and emulation is good, with remaining differences due to the use of floating-point arithmetic in parts of the emulator code. The degradation in resolution with increasing pseudorapidity is expected, due to a combination of the shorter lever arm available, the reduced effective precision for hits in the endcap and the impact of increasing material traversed by particles. It was noted at a late stage that the z0 resolution could be improved by a factor of two through a simple change to the stub data format, though this has not been implemented in the demonstrator yet.

The demonstrator latency was measured to be 3.5μs from the time the first stub leaves the Source to the time the first track arrives at the Sink, fixed for every event. The last track arrives an additional 0.2μs later. When scaling the demonstrated design to the full system, assuming no improvements to the algorithms, the latency would be similar. The Track Finder system would consist of 288 boards at an extrapolated cost of 6.9MCHF.

However, a reasonable estimate of the final system can be applied instead, where: the time-multiplexed period is reduced to 18 (instead of 36), algorithms can be improved and optimised to reduce overall logic requirements, and advances in FPGA technology and faster optical transmission

speeds will allow us to run our designs faster, reducing overall latency. In this scenario, we extrapolate the final system to consist of 144 boards, at an overall cost of 4.3MCHF, and a latency of less than 2.5μs.

While further results have been extracted from the demonstrator, and additional studies performed, only a few key results can be highlighted in this report. To this end, we have documented our Track Finding concept, algorithms and main findings in a paper, which is intended for publication. The paper is currently under review in CMS. In addition, we have also summarised our work and extracted results, for the Tracker Phase II Technical Design Report, due for submission to the LHCC in July. At a lower level, we have also started work on optimising the designs implemented in the demonstrator and understanding their performance when applied to the latest available FPGAs from Xilinx, particularly at higher clock speeds.

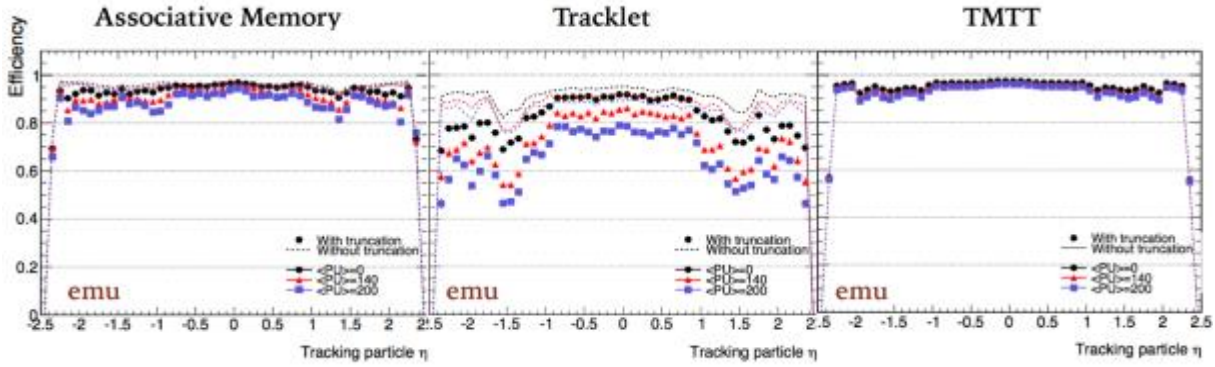

Figure 7.3: Comparison of standardised efficiency results, required for the review, across the three approaches for $t\bar{t}$ events with 0, 140 and 200 pile-up where the emulator is required to either simulate or ignore truncation effects (where the hardware does not have time to process the event before new data arrives). The UK demonstrator is the one labelled TMTT.

7.3 Overview of CMS plans

Despite previous indications, a decision on the preferred L1 track-finding architecture was not taken at the end of 2016 with the publication of the report after the review. The report itself acknowledges the UK demonstrator to be the most complete of the three presented. Following the presentations, in a closed session with the review committee we were commended on the 'fantastic results and performance' we were able to demonstrate for the review (Figure 7.3). Nevertheless, the report concluded that any of the three approaches could meet the requirements of a track-finder for CMS, although highlighted several reservations about the development and use of AM ASICs in the final system.

In January, Tracker Management agreed to formulate a Task Force comprising technical experts from each of the three track-finding approaches, in order to attempt to factorise the project and constrain the differences to a small part of the system. The Task Force concluded in April having partially succeeded in constraining the differences, but not enough to allow a workable project to be formulated. As a result, the Tracker MB has recently proposed that the baseline for the track-finder system should be based on FPGAs only, constructed as a common project, with an invitation for the AM groups to join. Follow-up discussions on this proposal, and hopefully ratification, will take place in the next month after which discussions on sharing of costs and responsibilities can begin.

7.4 Staff on project

The names of contributing staff and students were given in the previous two reports. Since January 2017, two people have left the project: F. Ball (Bristol, final year student), T. Schuh (KIT, final year student returned to Karlsruhe). Like Schuh, L. Ardile-Perez was funded by an EU ITN grant to be a

RAL employee for a several month secondment; he is about to return to KIT and resume his studentship there.

Others have temporarily reduced their involvement to concentrate on WP3 in particular (D. Cieri, K. Harder, G. Iles, A. Rose) where hardware R&D for WP3 overlaps significantly with WP5, or the demonstrated track finder results can be used as input to Level-1 simulation studies at Phase II. As mentioned, the project has profited from additional effort from our collaborators in Karlsruhe, working on developing and further optimising the track finding algorithms.

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 report is now completed, including the final item, the demonstrator note (5.4.8) which was delivered in 04/17. At this stage, it would not be sensible to formulate a new list without clearer direction from CMS on the next steps following the review, and a better understanding of the sharing of responsibilities between interested groups.

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

8 Risk register

The risk register has again been reviewed and some risks added and retired, as well as revised. There was a significant discussion of risks in the last report which we do not repeat.

No new risks have been identified for WP2 or WP4.

Risk 5.4, referring to MP7 manufacture, has been retired since further production seems unlikely. Risks 6.1 and 6.2, referring to trigger firmware and software, have also been retired, since the trigger has now been operating smoothly for over a year. Risks concerning trigger operation now belong to UK CMS M&O.

Risks 10.1 and 10.2 refer to the L1 track finder demonstrator. Although 10.1 was only added in the last iteration, the outcome of the work was very successful and it can be retired. However, the risks for the future concern the actual implementation, which now depend much more on the organisation of the future project. As observed last time, the problem is also political and strategic and is subject to some strong opinions, so the outcome is uncertain. For this reason, risk 10.2 has been added, which really concerns the implementation of results from the R&D work.

9 Finances

Expenditure is reported in the usual financial table. For those who consult it directly, the SBS Oracle database does not include income under the CMS codes, which affects those items for which repayments have been made to RAL from the Imperial grant, but not the summary financial table. The problems experienced with the Oracle system reported at the last OSC meeting were solved by RAL shortly after the meeting; they were attributed to assignment of management responsibilities by the system.

As discussed last time, RAL PPD and TD staff expenditure has been somewhat higher than expected for two main reasons:

- A small overspend compared to forecast affects PPD staff, thought to be partly due to STFC overheads but also due to a promotion of one staff member during the year.
- TD staff usage has, as explained before, been higher than anticipated due to requirements of the CBC design. The expenditure in 2016-17 was close to the revised forecast.

As anticipated in the last OSC meeting, it is now desirable to use some of the Working Allowance to cover extra TD staff costs for the remainder of the project. This is discussed further below.

Travel expenditure exceeded the original forecast for the year by about £10k. This was mainly due to the track finder demonstrator activity and associated meetings in CERN, and some conference attendance. It has been offset by refunding the materials expenditure (£9.3k) in RAL from the Imperial grant. However, given the activities under way and foreseen, such as SEU and beam tests of the CBC and modules, and development of the track-finder system, and conference reporting of results, it now seems that future travel will probably remain at the present level, for which the original budget would be insufficient.

Overall equipment expenditure to date is still less than originally predicted for, essentially positive, reasons explained in the last report; cautious development of trigger hardware and savings in CBC fabrication runs because of shared wafer submissions. The outturn at Imperial for the present year would have been close to the £400k forecast at the last meeting, but an invoice from CERN for about £91k (112.5kCHF) arrived later than planned so will be paid in the next period.

Despite exchange rate changes the overall situation remains favourable for the remaining expenditure foreseen. The Working Allowance of £386k has not been used to date, so it remains as part of the equipment allocation held in the grant at Imperial College.

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

Item	£k	Comment
CERN expenditure in pipeline	91.5	Invoiced April 2017
CBC3 wafers (24)	25.8	24 wafers, assume pay half
CBC4 submission	273.9	NRE only
CBC_final	308.4	NRE + 12 wafers: CORE
Ultrascale development	185.0	As in the proposal
RAL TD effort	100.0	Additional to that remaining in financial table
Additional travel	40.0	Compared to proposal
Total required	1,024.6	£1,227k available

In the above, we have made the very conservative assumption that two full wafer submissions are still needed to revise and finalise the CBC. The 24 CBC2 wafers are required for probe testing development and for module construction to test assembly procedures, and for evaluation until the final CBC is available.

Thus, we would like to transfer £140k to RAL. We have a proposal to arrange this, but would like now to seek OSC endorsement and plan it with the help of STFC.

10 Gantt charts

The Gantt charts for the projects remain unchanged. Some updates are expected once the TDRs are submitted.

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 March 2017 corresponds to PM48.

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.
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.