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

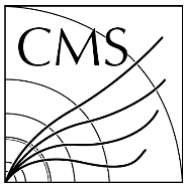

Upgrades of the Tracker and Trigger of the CMS experiment at the CERN LHC

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. The project is now in the final half-year.

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

- Delays were experienced with several aspects of the recent orders of CBC ASICs, all due to issues in the MOSIS interface to Global Foundries. However, as reported below, wafers have now been delivered, and the resolution of these administrative issues means that hopefully they will be avoided in future, during larger scale production.
- New hardware, firmware and software for CBC wafer probing was successfully applied to chip qualification, with excellent results and fast throughput.
- The additional lot of 24 CBC3 wafers ordered at the end of 2017 for module development, mainly funded by CMS, was delivered in September and successfully tested on the Imperial College wafer prober.
- The CBC3.1, which should be the final CBC version, ordered in February, was also delivered in September. Testing is still ongoing, but the indications are that all changes have been successful.
- CBC3.1 chips have been mounted on Imperial-designed hybrids, to be used for total-dose validation and detailed studies of performance.
- The first set of prototype cards for the Serenity ATCA platform delivered in May have been successfully tested, with all 120 bi-directional SerDes links operating at 16 Gb/s. Results were reported at TWEPP 2018, in September.
- Serenity 1.1, with minor bug fixes, was submitted for manufacture in July. Manufacturing time has increased a lot, as companies seem to be extremely busy and some component lead times are substantial.
- Several technical concerns for future ATCA hardware were previously highlighted and are under study. Encouraging results have been obtained from thermal modelling, with good support from collaborators.
- The CMS common hardware development initiated and managed by the UK seems to be evolving well and now involves several groups collaborating in complementary activities, covering the full range of hardware, firmware and software.
- Serenity cards will soon be shared among the collaboration, including processing sub-units and optical links, allowing evaluation of the Serenity capabilities as well as algorithm development. A Tracker Integration Facility will have multiple Serenity cards, for track finder and DAQ studies.
- The final milestone for the Serenity project to provide the v1.1 prototype with firmware and software infrastructure is on track but, as expected, overlaps with the construction project.
- Alex Tapper has been reappointed to a second two-year term as L1 trigger project manager, with effect from September 2018; the task includes overall management of the Phase II upgrade as well as operations and maintenance needed during LS2.
- A 3D clustering algorithm for the HGCAL trigger primitive generation has been successfully realized in firmware, overcoming challenges to implement the same logic as the TDR software studies.
- Backup studies to confirm the performance of the TPG firmware algorithm show that 3D clustering can be done on raw trigger data, rather than on 2D clusters. This has advantages for FPGA resources, but also gains in physics performance.

- HGCAL work is on track to utilise the Serenity hardware when new boards arrive to begin to validate HGCAL-specific daughterboards and interfaces which must be developed, in preparation for 2019 milestones.
- HGCAL simulation studies continue, focusing recently on optimisations of space and material necessary for a practical detector design.
- All elements of the track reconstruction chain of the L1 track finder have been adapted to the Ultrascale FPGA family and studies continue to profit from the technology enhancements over previous FPGA generations.
- Clock-speed improvement from 240MHz (Virtex 7) to 320MHz (Kintex Ultrascale) has been demonstrated. Further gains should be possible with the Ultrascale+ family.
- Xilinx High Level Synthesis firmware in the Kalman Filter Fitter has been shown to reduce logic use, improve memory usage and latency significantly.
- The CMS UK proposal for a five year construction project from April 2019 was approved by Science Board and now awaits final endorsement at government level. Efforts are beginning to make a smooth transition to the construction project.
- Although the UK R&D work is proceeding to a timely conclusion, mostly well in advance of other CMS R&D needed to be completed before construction starts, the delivery of some major items, such as the final CBC and the Serenity family, is somewhat later than originally foreseen. Procurements have been carried out with caution, in case of unexpected problems, and some significant items of expenditure are still to be completed. Therefore invoices are expected to arrive late. It would be desirable to extend the duration of the grant, at no extra cost, by a further six months to complete these activities.

2 Project history and recent developments

The LHC upgrade is proposed to take place in two main stages, with an increase in luminosity reaching 2.1 x 10^{34} $cm^{-2}s^{-1}$ in LHC Run 2 now completed, then after a shutdown from the end of 2023 to mid-2026 in LS3, to 5 x 10^{34} $cm^{-2}s^{-1}$ levelled luminosity, denoted as Phase II at the High Luminosity LHC (HL-LHC). During LS2, which starts at the end of 2018, the LHC energy will also increase to 14 TeV. In runs following LS3, 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.

It now also seems likely that the machine will reach almost 400 fb^{-1} of integrated luminosity in Run 3, compared to about 300 fb^{-1} previously projected. This may have some implications for the experiment, although CMS was designed for 500 fb^{-1}.

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

At the time of the last report, the LHC had only been operating for a short time this year; now the p-p running has just been completed (24 October) having delivered 163 fb^{-1} of integrated luminosity in Run 2, and exceeding the goal of 150 fb^{-1}. During 2018, CMS operated with an overall efficiency of 94% (Figure 2.1) to reach 150.5 fb^{-1} from the three years of Run 2, so an average efficiency of 92.5%. In 2018 levelling was not required to limit the instantaneous luminosity.

Following a week of machine development and a short technical stop, the year will conclude with a month of Pb-Pb operation with a target of 1.2-1.8 nb^{-1}. Several changes have been made in the

trigger and DAQ, including firmware modifications to the tracker FEDs and changes to trigger algorithms and running conditions, both involving the UK, to allow data taking rates of up to 30 kHz, which is about a factor 3 higher than previously reached.

Following the heavy ion run, the LHC will switch off for a lengthy shutdown (LS2) until 2021.

The problem with pixel DC-DC converters mentioned in the last report was resolved by further studies of the FEAST ASIC, which identified an increased radiation-induced leakage current magnified by a current-mirror as the origin of the problem. By charging a capacitive node, this caused a part of the circuit to exceed the safe operation voltage under some conditions, and therefore fail. Once understood, operational conditions could be adjusted so that this situation did not occur and the converters will be replaced with a revised ASIC during LS2. As well as important for operation, this is reassuring for the future, since DC-DC converters of this type will be widely deployed in the Phase II upgrade (including outside CMS). Revision of parts of the pixel system, with new DC-DC converters and a revised TBM ASIC, and a replacement inner layer (which was anyway anticipated because of radiation damage) with the revised PROC600 chip, will take place during LS2.

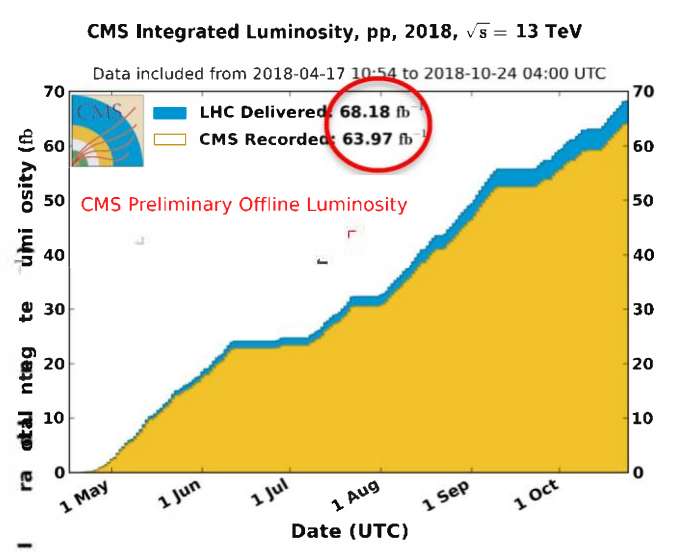

Figure 2.1. LHC and CMS integrated luminosity status in 2018.

The Level-1 Trigger system has run very reliably in 2018, with operational inefficiency halved compared to 2017. The calorimeter trigger, built and run by the UK, has performed very reliably, with very low downtime, and provided high-quality data for CMS. Thresholds have generally been lower than in 2017, giving higher acceptance for physics. Preparation for the Pb-Pb run involved extensive optimisation of algorithms and substantial firmware changes which have been successfully completed.

A new CMS spokesperson, Roberto Carlin, and team took office in September, for a two year term, so will oversee the long shutdown soon to begin.

The CMS Phase I upgrade is close to completion; the final components of the upgrade, the front-end electronics and photosensors of the hadron barrel calorimeter, are being prepared for installation during LS2. Other Phase I systems have been fully commissioned and are in operation.

Interestingly, given the UK trigger work and efforts on the Phase II track finder, a new track-finding algorithm for barrel muons based on a Kalman Filter has been studied with the aim of improving the muon trigger in the Phase II Upgrade of CMS. Studies in 2018 have shown that the algorithm can be run in the existing CMS Level-1 trigger hardware (MP7s) and it may become the baseline in Run 3, also providing triggers on displaced vertices.

The LS2 planning is considered to be very well studied, with a critical path passing through pixel and beam-pipe removal; muon system upgrade and maintenance; installation of the new Phase II

beampipe plus the revised Phase I Pixel tracker, and, finally, after closing the yoke, commissioning of the magnet with the Phase II Upgrade powering system.

The HCAL Phase I barrel upgrade (the last step of the Phase I program), the Phase I pixel revision, and several other tasks will take place in the shadow of this critical path. However, contingency is limited and additional work may be needed. CERN technical department resources are also very stretched and some infrastructure upgrades may not prove possible within the current LS2 plan so rescheduling may be required. This may have an impact on Run 3 year-end shutdowns.

2.2 CMS planning

Several infrastructure and facility milestones in relation to the Phase II upgrade of the P5 site have been successfully achieved. A new PM54 elevator became operational in mid-July. The CERN Finance Committee approved the SXA5 facility, which will provide dedicated workshop and storage areas and will initially host construction of the electromagnetic part of the two new endcap High Granularity Calorimeters. Work has also begun on the foundation of a new SXT1 storage building. Design studies have started for a new CMS control room which will address drawbacks of the existing location: powering, climate control, and space to work. Delivery by end 2021 is targeted.

The initial plan for funding of the CMS Phase II Upgrade systems was presented at the April 2018 RRB. The Phase II Upgrade projects have proceeded with Technical Design Reports (TDRs) and a detailed package describing cost estimates, schedule, resource information and management. This documentation was reviewed by the LHCC and the Upgrade Cost Group (UCG) in April 2018 for all projects, except for the Trigger and DAQ systems whose final TDRs are planned in 2020-2021 in order to profit from technology advances. For these, interim TDRs were included for the April 2018 round of reviews to provide complete updated information for the full upgrade scope.

For each project with an approved TDR, MoUs are being prepared to describe agreements on funding and on construction responsibilities. These MoUs will be submitted to the Funding Agencies for signatures at the end of 2018 and beginning of 2019.

2.3 UK adaptation to CMS planning

Our construction plan for the period 2019-2024 submitted to STFC in October 2016 has been through the review process and approved in full. It simply now awaits final approval at government level, and formal announcement of the awards, which we understand is expected soon. This will allow us to move to the next stage of planning, such as appointments and spending.

In the last report, there were some comments on the how the scope of the upgrades has evolved in the past years. At this point, it is probably sufficient to reiterate that the UK R&D phase is on track to conclude successfully, with the Phase I upgrade of the L1 trigger completed and operational since 2016, deploying a novel time-multiplexed architecture, which already has had a wide impact. This has led to the development of new Serenity hardware evolved from the MP7s of the Phase I trigger, and a model for a common CMS hardware development initiated and led by the UK, which has already been beneficial in promoting collaboration and minimising direct competition, and allows a focus on firmware and software, rather than on alternative hardware variants. In the longer term, maintenance effort and operational cost savings should result from such shared work.

In the tracker, the CBC development began about ten years ago as an evolution of the APV25 readout chip used in the current CMS tracker. Within a few years, this evolved, again initiated as a result of UK work, into one of the ASIC workhorses for the Phase II tracker, in which it was necessary to provide suitable data for the L1 trigger, to constrain trigger rates to tolerable levels. Although these have increased from the 100 kHz of the present CMS, which was initially expected also to be the limit in Phase II, the increased scope of the upgrade revision of the ECAL in particular has allowed this limit to rise to 750 kHz. Meanwhile, it gradually became clear that, while the on-detector p_T discrimination we proposed as the basis for the new tracker was viable, it was necessary to deploy it in all layers of the outer microstrip tracker (i.e. not pixels) and aim for complete reconstruction of tracks

exceeding about 2 GeV, rather than provide the limited number of points which was originally foreseen.

The consequence of this was twofold: the CBC design had to evolve to provide L1 trigger data, as well as full hit data readout, and a demonstration was required to prove that the very challenging, and never before achieved, target of online reconstruction within a few μsec, of huge volumes of tracker hit data, at the LHC rates, was actually possible. This was achieved by us at the end of 2016, with a time-multiplexed demonstrator based on the MP7s developed for the L1 trigger, and showed that a hardware solution based on programmable, and therefore highly flexible, FPGAs would be possible for Phase II.

The CBC development can now be considered almost complete, as some final verification of its performance remains to be fully proven and ionising dose tests should be repeated on the CBC3.1 even though they have been very thoroughly studied on the CBC3.0. Similarly we will repeat SEU tests and studies at the module level to ensure nothing has been overlooked. At present, the CBC is the only complete tracker readout ASIC in CMS (or ATLAS either, we believe) and has been vital for the 2S-modules, which are the only pT-modules to be constructed so far, and provided evidence that on-detector track-stub measurements could be made, justifying the new tracker design.

The CBC development has been hampered by the slower than expected development of tracker hybrids, which have proven to be more difficult to manufacture than hoped. This is mainly because of aggressive material budget targets, combined with high-density layout to handle the high data rates and large data volumes, and the use of C4, coarse-pitch, bump-bonded assembly. This method should in the longer term result in benefits from the commercial assembly to be deployed (an advanced version of what was done for the present tracker) but, perhaps unsurprisingly, issues at the development stage have required a good deal of attention and some delays. (These also affect the PS-modules, which are assembled in a similar way, but are some way behind compared to the 2S-modules, and are more complex). It is still believed the construction phase has sufficient contingency to absorb the module prototyping delays which have been experienced.

As an aside, one tracker ASIC has given some cause for concern, which is the digital CIC concentrator chip, which takes data from CBCs and the SSA/MPA front-end chips on the 2S- and PS-modules and formats the data and transmits it to the Gigabit (GBT) data links. It is therefore essential for full demonstration of tracker modules. The CIC is the responsibility of the Lyon group, but its development has taken longer than expected, and their efforts are now being supplemented by additional effort from experienced CERN staff to recover delays. The first prototype, CIC1, is due to be delivered in November, and it is hoped that CIC2 can be submitted in mid-2019.

The CIC is essentially the interface between tracker modules and the FPGA-based off-detector track-finder system and Data Trigger and Control (DTC) boards which are the first stage of the tracker DAQ. All of these are foreseen to be based on the UK Serenity hardware family.

Elsewhere in CMS, the developments from this R&D project will have a major impact, particularly in another UK-inspired and led project, the (forward) High Granularity Calorimeter (HGCAL) system which will deploy Serenity hardware to process the high-rate, high-volume data emerging from the silicon calorimeter and provide clusters and triggerable data to the L1 system.

Finally, the Serenity family will underpin the HL-LHC version of the L1 trigger by providing suitable processors to be applied to trigger systems and algorithms, probably especially for the Correlator, which should be the key element of the system where data from calorimeters, muon systems and, for the first time, the tracker will be processed and decisions made.

As this project reaches its end, we can look back with satisfaction on some major achievements, which we believe will project CMS in the direction it requires to produce and handle the Phase II physics primitives. Without exaggeration, the UK achievements can hardly be matched anywhere, and certainly not in the trigger, tracker, and data processing areas where we have focused our attention.

The financial situation will be discussed later in more detail. Because of the delays which have been described elsewhere, we approach the end of the project, especially for the CBC, with sufficient funds, but having necessarily been cautious in our spending until we could be sure that the design could be considered complete and no further iterations would be necessary. In the Serenity area too, caution may be less evident, but, as for the MP7 development, all the prototyping has been done with

great care to avoid unnecessary risk and excess costs. Many of the delays experienced (CBC manufacture, hybrid production, MP7 and Serenity PCB manufacture and assembly, etc) have effectively been outside our control, but nevertheless have been contained, by planning ahead, rigorous design methods, focused partial prototyping, working closely with external vendors where possible, and, perhaps most important, accumulated experience. Thus, the expenditure profile has been different than expected at the outset, but largely completed within the grant period, though now with somewhat more funds in hand than originally anticipated. Some expenditures have exceeded the original expectation, mainly RAL TD effort for the CBC design work, and travel, as a consequence of more comprehensive beam and system tests, and collaborative meetings, which have been vital to establish and maintain our leadership positions.

3 Work Package 1: Management

A reminder of the project management is included below. G. Hall remains PI, since January 2017 working 50%. G. Davies has been the UK CMS PI since January 2018. D. Newbold will be the PI for the CMS upgrade construction project, which begins in April 2019.

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	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. DAQ and 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 CBC3.0, a bump-bondable 130 nm CMOS front end readout chip for 2S-modules in the outer silicon tracker, was extensively tested in various ways in 2017 and, based on the findings, some minor issues were addressed by RAL and a new prototype, the CBC3.1, was submitted for manufacturing in early 2018. After some delays the finished wafers were finally delivered – for bump metallisation – in early August. Whereas the first order of CBC3.0 wafers was bumped for soldering by PacTech, with substandard results, the bumping of the CBC3.1 (and a new order of CBC3.0 wafers) was instead processed by Amkor in Singapore (subcontracted by Global Foundries, i.e. the same process offered by IBM in the past). As expected, the new wafers were delivered to us in September with the same type of uniform and high quality bumps (Figure 4.1) we received on the CBC2s. (NB: CBC2 wafers were metallised with leaded solder, with a relatively high melting point around 300°C, while CBC3.0 and CBC3.1 use a lead-free SnAg solder with a melting point around 220°C.)

After testing on a wafer probe station, the first two wafers were sent for dicing. The tests of chips on wafers and individual chips soldered to purpose-made test modules (Figure 4.2) both show encouraging results: the modifications implemented in the CBC3.1 described in the last report appear to have been successful, and do not appear to have introduced any new faults into the design. Nevertheless, the testing continues, and major remaining tasks include tests to verify that the radiation tolerance of the CBC3.0 remains for the CBC3.1 (scheduled for mid-November 2018), and to what degree the single-event upset immunity has improved due to design changes (planned for early 2019).

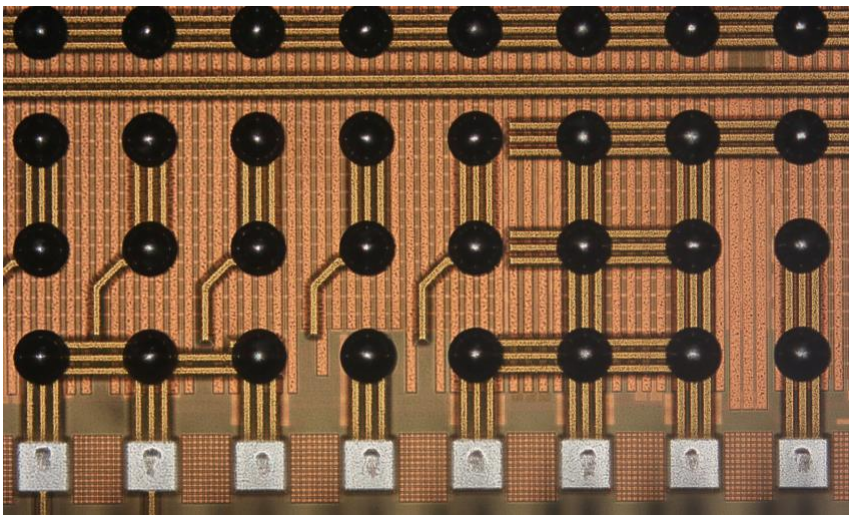

Figure 4.1. Smooth solder balls of uniform size manufactured by Amkor for GF (here on a CBC3.0 from the second order), but slightly smaller than we received from PacTech.

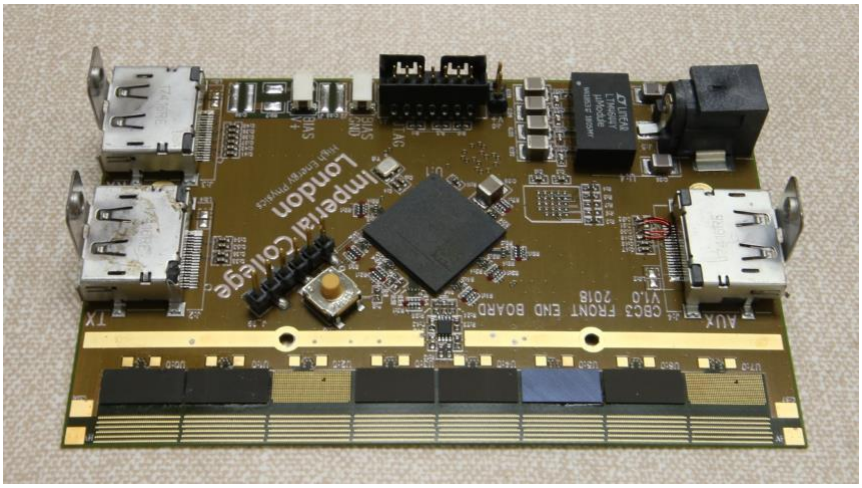

Figure 4.2. A 8CBC3 test board used for testing the inter-chip communication logic. The two chips on the left are CBC3.0, whereas those on the right are CBC3.1. The component at the centre of the board is an FPGA which transfers data from the CBCs, via the links on the left, to the DAQ.

Note that the boards used for these prototypes were designed at Imperial College and manufactured by Somacis (Italy), who we have used regularly for MP7, FC7 and Serenity boards. They do not replicate all the features of the CERN-designed CBC hybrids for 2S-modules, but allow us to mount bump-bonded chips in-house, using a BGA solder-reflow station, and avoid commercial assembly and handling, including adding a (single) sensor if we wish. We have been doing this since October 2017 as it only became possible with the lower melting point lead-free solder. It obviously has a big impact on the time needed for manufacture, allowing us to evaluate the ASICs and modules much more quickly, and in greater numbers if needed, and has so far proven to be very successful.

In parallel with the CBC3.1 evaluation, 24 wafers from a lot of CBC3.0 chips (shared with a CERN SCA chip, which is discarded because of the bump metal processing for the CBC) were qualified, so the chips can be used in module prototyping. Despite submission dates differing by several months both sets of wafers arrived in September only a week apart. Significant development of the hardware, firmware and software supporting the wafer prober had been undertaken since the previous deliveries, and so this was used in a production mode for the first time. The time required to fully test each chip had been reduced from a few minutes to 30s, so a throughput of several wafers per day was possible (these wafers contain 186 CBC3.0 sites, compared to 478 for the CBC3.1). The observed yield was encouraging, but still had a spread from 58-90%, which is not an insignificant variation.

4.3 Overview of CMS plans

The main issues which directly affect the UK relate to the construction schedule and particularly the 2S-module construction, since the track-finder and DAQ activities are not on the critical path and should not become so. While large scale module assembly does not get underway until 2021, a significant amount of prototyping remains to be carried out to achieve that and some of the recent developments have had an impact on progress. The early start on the CBC development (partly motivated by a much earlier Phase II LHC schedule than is now the case) was vital to prove the concept of the novel pT-modules, which necessitated working ASICs, and the overall electronic system design has evolved considerably, such as the increase in the overall trigger rate from 100 kHz to 750 kHz and the consequent increased data bandwidth at the module level.

These changes have had an impact on the CBC design, but have also made the hybrids more challenging, with higher density traces and clock speeds and the progress in prototyping modules has been slower than anticipated. Some of the difficulties over the last year or so have been:

- There have been technical challenges in manufacturing the hybrids required for both 2S- and PS-modules, including delamination issues, for which reason production was suspended for a lengthy period until solutions were found.
- Bump-bond metallisation of the ASICs was prototyped using a vendor (PacTech) whose process did not use photolithographic definition of the electroplated bumps, but so-called electroless metal deposition. This essentially involves depositing bumps individually, and is suitable for prototyping but costly. However, the work at PacTech was not very successful, of relatively poor quality, and resulted in low yield, generating significant delays.
- The CIC concentrator ASIC already mentioned has taken longer to design than expected. It was not available to be included in the front-end hybrids so far.
- The service hybrid, which contains DC-DC converters, GBTs and optical links was also delayed because of various technical challenges.
- The sensors required for the tracker were being prototyped by two large manufacturers (HPK, Japan and Infineon, Austria), with mostly excellent results from both. However, there has been a limited supply of prototype sensors. More importantly, in the last year, Infineon decided to withdraw from future sensor business, which means that several major HL-LHC projects must rely on a single vendor. While HPK has an outstanding record, a single source is not ideal, but is being managed attentively at high level in CERN.

Each of the issues has had an impact on the CBC development since we were unable to assemble bump-bonded hybrids and modules to fully validate the CBC3.0 design until November 2017, by which time we had produced alternative hybrids of our own design and assembled working modules using a local ball-grid array reflow station, primarily intended for assembling or repairing FPGA-type boards.

Despite these issues, the contingency in the tracker schedule is believed to be sufficient to accommodate the delays, and recently good progress has been made in overcoming the problems, although full proof awaits assembled modules. From the UK perspective, the main concern has been to stress that issues elsewhere should not have an impact on the CBC design, i.e. by proposing unexpected changes at a late stage. This has been accepted by the tracker, and understood to be very unlikely, especially given the large efforts which have been made to test all features of the CBC. In this respect, the CBC is far ahead of other developments and our tests have been comprehensive. In addition, the CBC and 2S-modules are much simpler than the MPA/SSA ASICs and PS-modules, so, for example, the interface to the CIC is considerably less demanding.

Unfortunately, the delays elsewhere have held up the CBC development and added extra costs (mainly in RAL TD design effort as previously explained). In contrast, during the CBC development we have been more successful than originally expected in sharing ASIC engineering runs (primarily with a RAL XFEL ASIC and a CERN GBT-related chip), which allowed us to contain that part of the expenditure and apply it to engineering effort.

Thus, in conclusion, the CBC development is ending later than originally planned, but in a timely way and within the constraints of both the UK R&D project, and the overall tracker R&D and construction schedule.

4.4 Staff on project

No recent changes.

4.5 Expenditure

The expenditure on RAL TD staff has been low for the past months as only some work on updating the documentation of the CBC3.1 remained after the design had been submitted for manufacturing. As this new prototype appears to have been successful, we expect a continued modest

expenditure as the role of RAL has shifted to support and documentation of the CBC3.1. However, full verification of the CBC3.1 will not be possible until Q4 2018.

Other expenses involved are the manufacture of the hybrid PCBs, and components, and in upgrading the probe station and surrounding facilities. The Imperial basement laboratory is scheduled to be handed over to another research group in the coming months and for the laboratory to be re-established on the Blackett Laboratory 5th floor. No major problems are foreseen, but this has taken longer to carry out than expected. It will include an improved clean environment, supplemented by an enhanced local clean air facility surrounding the prober.

In addition to documentation and support for CBC3.1 testing from RAL TD engineers, there is one further modest task we would like to complete, which is development of a model of CBC functionality, based on the actual ASIC design, which can be used in simulations. This refers to electronic simulations, using System Verilog, although the work could easily be the basis for later implementation in CMS software. A simple behavioural model of the chip has been used by CERN engineers to help verify the CIC functionality, and it is desirable to put this on a firmer footing. About 2-3 FTE months of engineer effort are estimated to be needed.

4.6 Deliverables

The WBS for WP2 is included below. The work we described earlier is expected to confirm that the CBC3.1 should be the final version, so item 2.5 can be considered complete. The CBC3.1 masks will be those needed for production, hence part of our CORE construction contribution will have been delivered.

We therefore believe that item 2.7 (a further CBC version) will not be required and item 2.6 (2S-module based on CBC3.1) is on track to be completed on schedule if the new hybrids are successful, although the service hybrid containing DC-DC converters, GBT ASIC and optical links will not be final.

WBS item 2.8 is ongoing. The delivery and wafer probing of the CBC ASICs belongs to the construction project and the Global Foundry orders needed will proceed after the end of this grant. However, we plan to place orders for several 24 wafer lots of CBC3.1s using funds from the present grant. There are a number of reasons for this:

- we can only properly evaluate the yield with more statistics and several production lots are necessary to permit this,
- although the recent 24 wafer CBC3.0 delivery shows a reasonable yield, the wafers were shared with another design, so may not be representative,
- yield issues have been experienced by a Medipix CERN design. This chip probably has a denser design than the CBC and has more metal in upper layers. Studies at Global Foundries have led to an optimisation of the process, whose impact on the CBC manufacture is unknown,
- manufacture of a couple of 24 CBC3.1 wafer lots will allow us to complete the commitment of UK funds for part of our CORE contribution to the tracker (500kCHF), and give a fairer representation of UK funding of the tracker project than only counting funding from the new construction project. (The UK has been unique in completely funding such a major development as the CBC ASIC over a ten year period. This has amounted to approximately 1MCHF in ASIC manufacturing costs and 2MCHF in RAL TD engineering design effort, plus substantial additional activities in Imperial College and Bristol.)

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 2S-PT module readout**
	2.1.1 specification definition	04/13	03/14	12	regular meetings with CMS collaborators to define overall system specification and interfaces
2.2	**CBC2 test**	04/13	03/15	24	**CBC2 is final deliverable of the UK upgrade R&D**
	2.2.1 CBC2 ongoing testing	04/13	03/14	12	complete the detailed studies of the CBC2 chip, including irradiation and SEU tests
	2.2.2 CBC2 2S-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 2S-PT module studies	09/16	03/18	24	CBC3 based module studies in collaboration with CMS in lab and test beam
2.5	**CBC3.1 design and test**	09/16	12/17	15	**CBC3.1 is the final version of the chip, fixing any**
		07/17	**11/18**	**12**	**remaining bugs found in the CBC3**
	2.5.1 CBC3.1 design	09/16	12/16	3	design period
		07/17	**01/18**	**7**	**longer design period while awaiting CBC3 results**
	2.5.2 CBC3.1 production	01/17	06/17	6	production period
		02/18	**09/18**	**6**	**production period**
	2.5.3 CBC3.1 testing	07/17	12/17	6	tests to verify full and final functionality
		09/18	**12/18**	**3**	**expedite functionality tests**
2.6	**Final validation**	12/18	04/19	4	
	CBC3.1 2S-module	12/18	04/19	4	
2.7	~~**CBC3.2 iteration**~~	~~11/18~~	~~04/19~~	~~5~~	~~**Only if correction is required**~~
	2.7.1 CBC3.2 final masks	12/18	01/19	2	mask preparation for full wafer engineering run
	2.7.2 CBC3.2 engineering run	02/19	08/19	6	production period
	2.7.3 CBC3.2 final production	09/18	12/18	3	**final verification in post project period**
2.8	**Production planning**				
	2.8.1 procurement planning	01/18	04/19	15	detailed 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

As reported earlier, the Level-1 Trigger system has run very reliably in 2018, with excellent efficiency. The Pb-Pb run will begin shortly, and required extensive optimisation of algorithms and substantial firmware changes by UK staff.

5.2.2 Phase II hardware development

At the last meeting in May, the first set of prototype cards (v1) for the Serenity ATCA platform had been delivered and the initial power tests had just started. Testing continued throughout June, culminating with all 120 bi-directional SerDes links operating at 16 Gb/s – a total of 3.84 Tb/s, and with a presentation at the Topical Workshop on Electronics for Particle Physics, in September.

In early July a version of the card with minor bug fixes (v1.1) was submitted for manufacture, with the objective of having a relatively fast new iteration; however, PCB manufacture has taken 10 weeks, 2 weeks longer than for v1.0. Furthermore, PCB assembly is still underway, but is estimated to be 8 weeks, 4 weeks longer than v1.0 and currently scheduled for delivery Nov 9th. Total manufacture time will be 18 weeks, which is far beyond the expected 6-8 weeks. Some of the additional time is due to the summer shutdown at the PCB manufacturer, but the companies also seem to be extremely busy, with component lead times now reaching 6 months for several parts.

The biggest impact of these component lead times has been on the supply of the FPGA-based daughter cards, which mount on the base board and provide the processing capability and high speed SerDes I/O. These were ordered at the same time as Serenity and shared the same PCB run to reduce costs; however, the two types of FPGA (KU15P, KU115) are not due for delivery till end November and the middle of January respectively. We have ordered 8 of each type of FPGA (one for logic intensive applications and a lower cost part for interfacing to the detector). We will assemble 3, followed by 5 more. The optical modules are also on long lead times of ~6 months, with parts not expected to arrive till January or February. An open question remains whether we should order additional optical modules now to mitigate supply issues, albeit with a risk that the parts are somehow not suitable.

Some applications, such as the HGCAL trigger or Tracker track finder may require one of the largest FPGAs available from Xilinx (VU7P). The daughter cards for these are currently under design at IC and TIFR, with the aim to submit the design in November with parts available from Xilinx in February.

To speed up the assembly process we are considering a new approach that would purchase parts in advance and require the assembly company to have insurance against manufacturing defects, although the robustness of this purchasing process in case of difficulties remains unclear. Consideration should also be given to providing additional support to the assembly company to flag up potential issues (e.g. they have no automated detection mechanism for part changes, which have therefore gone undetected

in the transition from v1.0 to v1.1, causing significant delays to us and presumably cost to the company involved).

5.2.3 Phase II hardware evaluation and analysis

In the last report we highlighted a series of technical concerns for the hardware, which we felt needed to be addressed. These are expected to be consolidated into the following sections:

- Thermal Analysis
- Optical Module Evaluation
- Mechanical Analysis
- Slow Control Evaluation
- Procurement and Quality Control

While the procurement process is not a direct technical aspect of the board it is closely coupled when considerations of part availability, performance, cost and manufacturer start to be considered. It is therefore critical that it be integrated with the hardware development at an early stage to incorporate purchasing requirements (e.g. market survey) and to avoid design changes late in the design cycle.

The thermal analysis studies, which we were planning in May, are now largely complete. They benefited from substantial support from INFN, Pisa. Agreement at the level of ~5°C was achieved between measurements and simulation for several different heatsink configurations without any significant tuning, which has given us confidence in our ability to model the system accurately. This has enabled us to design heatsinks that should enable us to operate the LHC crates in a manner that reduces fan power from ~60% of the electronic card power consumption to ~15% and reduces fan noise from levels that exceed the health and safety limit of 85dBA to more manageable values. It has yet to be seen whether when scaled up to a system of many crates the noise levels remain tolerable or whether special safety precautions will be necessary. The only caveat is whether it is economical to manufacture the optimised heatsink design for relatively few parts (~1000).

Importantly, these results also suggest that we should be able to maintain the temperature of the optical modules at 50°C or less, which is essential for their longevity. Only when we model systems with many 25Gb/s links, which dissipate more power, do we start to have temperatures that exceed 50°C, but only by a few degrees, which is acceptable.

Of the remaining sections, the optical module and slow control evaluations are both starting to get underway, but it will be several months before results are presented. The mechanical analysis and procurement sections have yet to start.

The evaluation of high bandwidth, high density mid-board optical modules (i.e. placed close to the FPGA for optimal signal integrity) is necessary because the parts are a niche product. At present we are reliant on a single supplier and we would like to lessen that risk by looking at alternative companies. This would potentially require that another version of the Serenity card should be manufactured, although the evaluation of alternative parts can be completed with existing development cards and small test cards.

An alternative slow control approach is being evaluated by our collaborators at KIT, which may offer a more integrated and flexible solution to the existing off-the-shelf components, but it has yet to be proven. They have been supplied with a Serenity v1.0 card to undertake these tests.

At the end of 2019, when the systems and card designs are finalised, we aim to have all sections documented so that there is a clear technical decision process, involving all the interested groups. As hoped, the development has become international with KIT, Germany; University of Belgrade, Serbia; TIFR, India and CEA Saclay, France all collaborating with the UK on hardware.

5.2.4 Phase II firmware and software

Many of the groups supporting the hardware development are also contributing to the firmware and software, but there is additional support from Chulalongkorn University, Thailand and Ioannina, Greece. These groups are extending the existing basic firmware and software infrastructure (i.e. slow control and payload only) to include a back-end link protocol at 16 and 25 Gb/s, while others will

integrate the existing front-end link design. In parallel to this the UK is both simulating and developing the required DAQ functionality.

.

5.2.5 Plan for 2019

Our plans for 2019 have not changed significantly since May. We aim to equip ourselves and our collaborators with the Serenity platform, including processing sub-units and optical links. There will be cards for ourselves at Imperial, RAL and CERN, plus a common integration/development area in the existing Tracker Integration Facility at CERN. Cards will also be distributed to our key collaborators so that they can expand and test the capabilities of Serenity (e.g. testing of an additional processing sub-unit at KIT and timing analysis at Saclay), as well as develop algorithms for the projects we are involved in (e.g. HGCAL firmware by TIFR). Some locations, such as the Tracker Integration Facility, will have multiple Serenity cards, but populated with different optics and processing units, depending on the target application.

A first batch of 5 cards (v1.1) is due in November with an order in place for up to 15 additional cards. Some of our own facilities would need to be equipped with ATCA crates, but most would use our benchtop adapter unit, which avoids the cost of additional crates and provides better access for debugging.

Our final milestone is to provide a v2 prototype (renamed v1.1) with full firmware and software infrastructure set (i.e. slow control, payload, link and diagnostics) at the end of Q2 2019. At present this seems achievable, except for the Phase II TCDS interface, which is expected to operate somewhere between 5 and 10Gb/s (particularly, at Gen 2 and Gen 3 PCIe line rates), whereas Serenity was designed to operate up to 6.25Gb/s. The TCDS does have a legacy interface that should be sufficient.

A summary of the plan is given below in tabular form, along with the expected costs of the major components; these sums have been committed but mostly not yet invoiced.

We intend to share the boards where appropriate with collaborators, both to expedite progress but also to cement this model of development and ensure that everyone is committed to it, as there is an inevitable risk that other groups may prefer to compete, rather than collaborate. At present, partly because of excellent UK progress, this risk seems slight and there are no visible alternatives at present. Where boards are delivered outside the UK groups, we expect to be recompensed, or balance our contribution against reciprocal expenses.

Date	**Deliverable**
May 2018 - Complete	All electrical and optical testing complete. Revised design – v1.1 ready for manufacture.
June 2018 - Complete	Basic firmware and software infrastructure (i.e. slow control and payload only)
July 2018 - Complete	PCB manufacture complete Preliminary results from thermal mock-up and simulation
Sep 2018	Assembly of 3-5 base and daughter cards complete Preliminary results from mechanical mock-up
Nov 2018	Validate v1.1 Cards
Jan 2019	Assembly of remaining base cards and daughter cards complete
Mar 2019	Full firmware and software infrastructure (i.e. slow control, payload, link and diagnostics)

	Cost (£k)
Card manufacture, including components and assembly	83
FPGAs	66
Optical Components	33

There are some issues about how this activity should evolve and be managed, which mainly concern the construction project from April 2019, rather than the remaining R&D activity. While the common hardware development has generated enthusiasm in CMS, as an activity spanning multiple sub-detectors, it does not have a very specific mandate and may prove challenging to coordinate in the long term. At present, the two main issues are:

- the technical choices facing the UK for Phase II back-end systems and the timely delivery of the technology for our own interests - i.e. tracker, HGCAL and L1 trigger.
- the wider international context, in which we form a collaboration to provide additional support to the Serenity platform and encourage others to develop common solutions to the hardware, firmware and software tasks facing us.

A bi-weekly meeting has been held since November 2017 to coordinate the activities between the international groups contributing and has been working well. It has led to a common code repository and associated documentation. It has ensured that the Q1 2018 milestone for a preliminary firmware and software stack was met on schedule. The group is now active in extending the firmware/software stack for the Q2 2019 milestone, developing additional daughter cards (TIFR and KIT) and investigating the thermal aspects (INFN).

5.3 Overview of CMS plans

A. Zabi (LLR) and J. Berryhill (FNAL) were appointed to Phase II L1 Trigger Upgrade coordinator positions till Q2 2020. They will be responsible for the publication of a full TDR at the end of 2019.

For the upcoming TDR milestone, the key decisions required 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 for the Correlator will ultimately depend on the choice of algorithms and how they are implemented. A series of workshops will be held to determine the optimum Correlator configuration and to discuss the simulations and implementations of the Phase-II Level-1 algorithms.

5.4 Staff on project

There are no staff changes.

5.5 Expenditure

Overall spending is within the budget foreseen. Delays in manufacture, such as those mentioned, have an inevitable knock-on impact on invoicing, which therefore needs care as we approach the formal end of the grant period. We foresee the need to commit soon approximately another £62k for optical components, FPGAs and PCBs for evaluation studies, also because of long lead times on parts which will be needed in 2019. Please see additional remarks in the financial section.

5.6 Deliverables

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, some of which were not foreseen, such as the L1 track finder and the HGCAL. The milestones which are listed below were proposed for the remainder of the project. The overall 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 (now called v1.1) 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.

6. Work Package 4: High Granularity Calorimeter

6.1 Objectives

To play a long-term leading role in the HGCAL, in terms of both overall project management and technical aspects, specifically for the latter in the two areas listed below:

- To develop the trigger primitive generator (TPG), including contributing to the algorithms, firmware and HGCAL-specific hardware design;
- 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

Much of the HGCAL TPG work since the last OSC has focused on a more complete firmware implementation and improving the simulation performance. At the time of the TDR in Nov 2017, there was an almost complete firmware implementation for 2D clustering of trigger hits within each layer, but no firmware for the subsequent 3D clustering to form complete particle-like objects. Hence, the TDR had latency and FPGA resources estimates for the former but not the latter. A significant effort has been put into this 3D clustering implementation in the last six months. During this work, one issue discovered was that the assumed "cone" 3D clustering algorithm used for the TDR software results could not be realistically implemented in firmware so a different algorithm had to be developed. This is based on histogramming the hits over all layers to find energy peaks and then using these as 3D cluster seeds. An almost complete implementation of the new "histo" algorithm now exists in firmware and a flow chart of this algorithm is shown in Figure 6.1. The design fits comfortably into the resources of the assumed FPGA, using around half of each resources type, and takes around one-third of the total latency budget of 1 μs, so this approach looks very feasible.

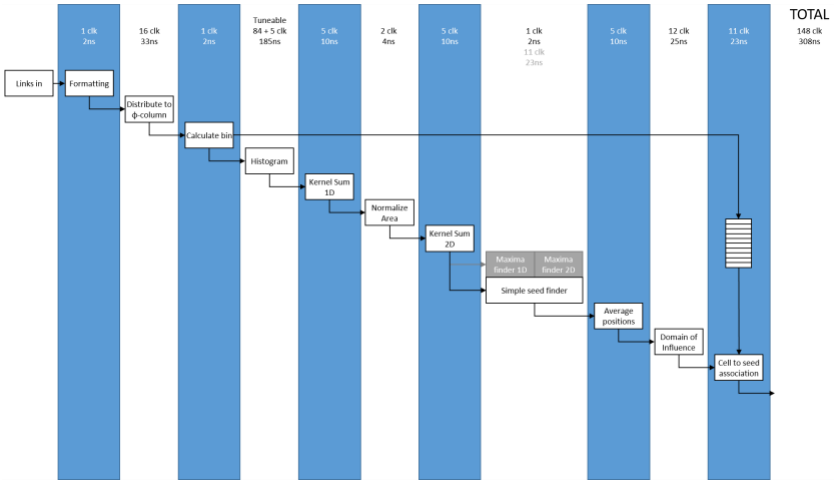

Figure 6.1: Flow chart of the "histo" 3D clustering algorithm. The estimated total time required for the processing of all data in each event is 308 ns, which is well within the required limit.

This "histo" algorithm then had to be implemented in software and checked in simulation to ensure its performance would not be worse than the "cone" algorithm, such that the TDR results would still hold. A comparison of the two methods for jet energy resolutions is shown in Figure 6.2 (left) for a variety of "histo" parameter values; with reasonable parameters the two have effectively the same resolution. Furthermore, the processing time and resources of the "histo" algorithm were found to depend only weakly on the volume of data entering the clustering algorithm. This allows the 3D clustering to be done on the trigger raw data, rather than on 2D clusters; it had previously been assumed that 2D clustering would be needed to cut down the data volume into the 3D clustering processor. Using the raw data has the advantage that no energy is lost in the 2D clustering step from

unclustered hits or low energy 2D clusters below the seeding threshold. A comparison of the jet energy resolution when using 2D clusters and raw hits is shown in Figure 6.2 (right) for both the "cone" and "histo" algorithms. The raw hits give a better resolution by 10-20% in all cases. An alternative cost-neutral architecture for the TPG which would allow the raw hits to be transmitted directly to the 3D clustering boards has been developed in terms of Serenity boards and looks feasible, so this is an interesting and realistic option which will be investigated further.

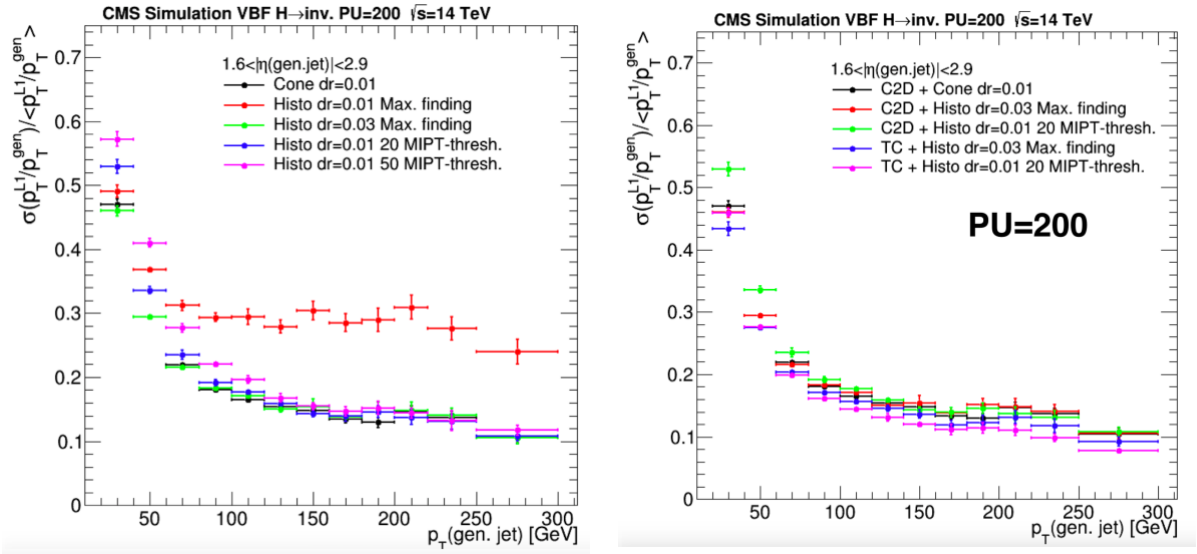

Figure 6.2 Left: Comparison of the performance of the TDR "cone" and new "histo" algorithms for several parameter sets in terms of jet energy resolution. Right: Comparison of the performance of the "cone" and "histo" algorithms using 2D clusters ("C2D") or raw trigger cell ("TC") hits as inputs, again in terms of jet energy resolution. The best performance is given by the "histo" algorithm using raw hits.

A number of detailed decisions must be made before the final engineering design of the HGCAL. It has become clear that the space allocated to the electronics between the calorimeter layers needs to be increased compared to that specified in the TDR. A range of options was suggested and simulations were made to compare the impact of the different options on the energy resolution and shower size in the electromagnetic section. It was shown, for example, that an idea to maintain the calorimeter density by including some of the absorber of each layer in the space between the two PCB boards, in the form of a copper plate with several recesses in it to accommodate space-hungry components, resulted in an unacceptable resolution loss due to the variation of absorber thickness resulting from the recesses.

The study then went on to examine the variation of shower size result from the different strategies. Figure 6.3 compares the radius of 68% containment of photon showers in layers of the calorimeter versus distance from the front face, for different versions of the calorimeter. To good approximation the shower size versus distance from the front face follows the same curve for the different versions, allowing easy estimation of the effect of any further proposed small changes on the shower size in any given layer. Dedicated standalone simulation studies are under way on a number of other issues where decisions must be made in the near future.

Progress towards a future full offline reconstruction in the CMS software continues with the development of an iterative cluster framework, TICL, by a small CERN group. One of the ingredients of this will be a "tracking" approach to the association of the 2D clusters in each layer, which has been explored by an Imperial physicist. Figure 6.4 shows an event display of individual silicon cell energy deposits in a charged hadron shower in the HGCAL, collected in the presence of 140 pileup events, using a rudimentary tracking algorithm. The trajectory of the p_T = 20 GeV charged pion is represented by the yellow line.

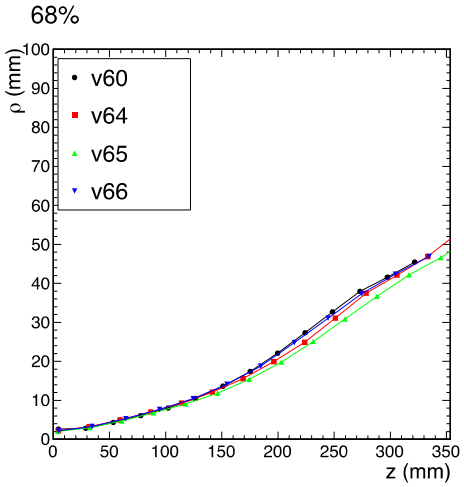

Figure 6.3: Radius of 68% containment of electromagnetic showers in layers of the HGCAL versus distance from the calorimeter front face for different versions of the calorimeter.

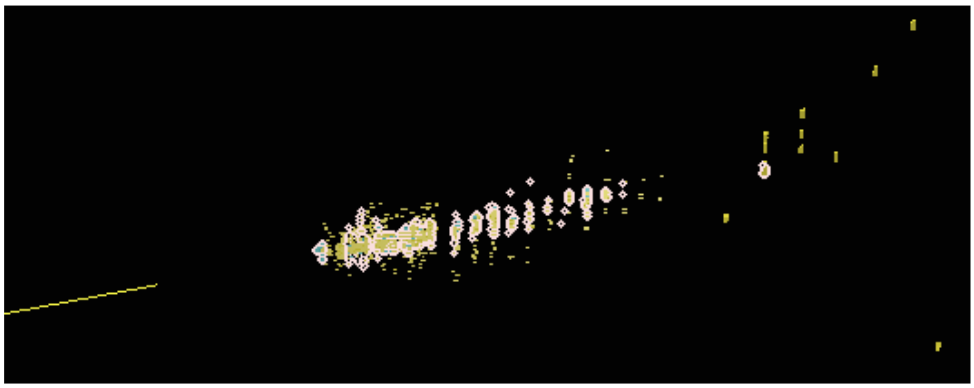

Figure 6.4: Event display of individual silicon cell energy deposits in a charged hadron shower in the HGCAL, collected in the presence of 140 pileup events using a rudimentary tracking algorithm. The trajectory of the $p_T = 20$ GeV charged pion is represented by the yellow line.

6.3 Overview of CMS plans

We are now looking ahead to two major milestones, namely the completion of the prototype back-end board tests in October 2019 and the Engineering Design Review (EDR) in July 2020. The former is of course a critical step; while much of the testing will be done within the context of the common hardware design group, there are HGCAL-specific daughterboards and interfaces which must also be verified. This work will increase steeply when the next version of the Serenity boards arrives later this year. Following these tests, we will be heavily involved in documenting any changes to the back-end electronics designs for the EDR, as well as producing new and updated simulation results.

6.4 Staff on project

Since the last OSC, the Imperial RA hired on T. Virdee's ERC grant, Thomas Strebler, has moved to a permanent position in France. A new RA, Samuel Webb, has been hired as a replacement and he started work on 17 September.

6.5 Expenditure

Expenditure so far has been almost entirely on staff and travel, mainly using non-project ERC funds. However, this is beginning to change as the HGCAL funds have contributed to the prototypes

of the common hardware baseboards and HGCAL-specific daughterboards. There will be a new order of the latter within the next few months and this will be relatively expensive, as these will each hold a reasonably large FPGA (Xilinx Ultrascale+ VU7P). We foresee around £30k from the ERC grant being spent over the next six months on these components plus some infrastructure required for testing.

6.6 Deliverables

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

- Over the next year, the main deliverable for the TPG is the manufacture and testing of a generic baseboard and several HGCAL-specific daughterboards to produce a functional prototype of the TPG board;
- The main near-term deliverables for the UK within the simulation are the detector optimisation study of the mechanical layer thicknesses and single-channel calibration methods.

The HGCAL milestones cover the detector construction from January 2017 to December 2025. There are no explicit milestones yet for the future simulation work. The full list of remaining TPG high-level milestones is shown in Table 6.1. The main near-term hardware milestones relevant to the UK effort are:

- Medium level: 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 start testing a functional prototype within the next month. Note, this milestone is just after the end of the current STFC upgrade R&D grant.
- High level: Prototype board functions and interfaces validated in October 2019. To meet this, all I/O from the HGCAL-specific daughterboard, including the interfaces to the DAQ and L1T, will need to be checked. Fulfilling this milestone will allow us to document the design for the EDR review in 2020.

TPG and back-end electronics (BE)		
Prototype functions and interfaces validated	CE.BE.4	07-Oct-19
Specification of BE system documented in EDR	CE.BE.5	20-Apr-20
Preproduction functions and interfaces validated	CE.BE.6	01-Nov-21
BE hardware production readiness review	CE.BE.7	11-Jul-22
Production functions and interfaces validated	CE.BE.8	25-Dec-23
BE electronics production complete	CE.BE.9	23-Dec-24
HGCAL integration with central DAQ and L1T complete	CE.BE.10	30-Dec-25

Table 6.1. Remaining high level milestones of the HGCAL back-end electronics project.

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

While the original objectives of this WP have been met, the UK continues to build upon the concepts and results presented to date in preparation for the anticipated construction project. Mark Pesaresi is System Development WG co-leader within the CMS Tracker project, and now UK representative in the Tracker Upgrade Steering Group.

As the next-generation hardware development in WP3 progresses, we continue to port and improve upon the track reconstruction algorithms demonstrated to date. This process is aided significantly through the use of common tools to debug the algorithm firmware as it is developed. These tools are provided by a combination of the infrastructure firmware framework discussed in section 5.2.4, which provides integrated simulation test-bench functionality as if the design were running in real hardware, and a software demonstrator framework that handles the interface between the firmware and simulation software, and automatically compares the firmware result with its emulation.

All elements of the track reconstruction chain have been now been adapted to the Ultrascale family of FPGAs (currently the KU115 only) and we continue to investigate how to take advantage of the technology enhancements over previous FPGA generations. A key parameter is the operating frequency of the firmware and we have demonstrated that an improvement from 240MHz (Virtex 7) to 320MHz (Kintex Ultrascale) is feasible for almost all elements so far (Figure 7.1). Further gains should be possible by targeting the Ultrascale+ family, which includes the VU7P, where a frequency of 360MHz would be a realistic target.

The use of Xilinx High Level Synthesis (HLS) language for firmware continues to be an avenue of interest for WP5 and potentially other work packages too. Firmware written in HLS should enable faster development turn around, modularity and potentially faster and smaller logical implementations. This has now been demonstrated to impressive effect in the state update calculation of the Kalman Filter Fitter, where logic has been cut by 30-50% (and memory usage by a factor of 4) and latency by 20%. This comes with the added benefit of easier validation against simulation so that improvements or fixes can be rapidly applied. We will continue to explore other areas of the design which might benefit from adapting to HLS, particularly if it aids future portability and maintainability of code.

These improvements, in combination with significant optimisation of the algorithms taking into account changes of the tracker geometry, the latest system design and latest estimated data rates, currently yield a provisional factor of 2 reduction in FPGA resources compared to the estimates provided in 2017.

While these firmware investigations continue, we are also exploring if adapting parts of the track reconstruction chain to use alternative algorithms can offer improved performance, resource usage or latency. At present these hybrid algorithm studies are focused on understanding performance enhancements, but firmware designs will be considered in the near future (Figure 7.2).

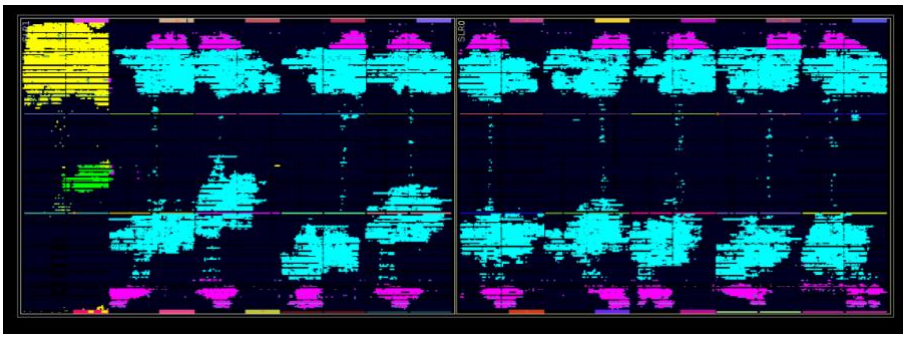

Figure 7.1. Implementation of the full Kalman Filter Fitter algorithm (cyan) in a KU115 design at 320MHz, with 64b-capable infrastructure and service blocks derived from WP3 (pink, yellow, green).

7.3 Overview of CMS plans

It was reported in our last update that CMS had tasked a committee with ensuring that the FPGA-based approach could be considered sufficiently robust under non-standard operating conditions. A number of scenarios (extreme pileup up to 300PU, increased tracker material by 50%, tracking in high p_T jets, multiple dead layers and modules, stuck bits/noisy strips) were proposed where studies and cross-checks were carried out by a team of L1 track finder experts, including some from the UK. In October the committee concluded that it "found the system demonstrated the desired flexibility in handling the various stress tests" and a closing report is in preparation. As a result the "top-level" Data Processing WG, responsible for interacting with the committee, will now disband. The System Development WG will continue in its present role of steering the continuing R&D for the Tracker off-detector readout and processing system, which includes the L1 track-finder.

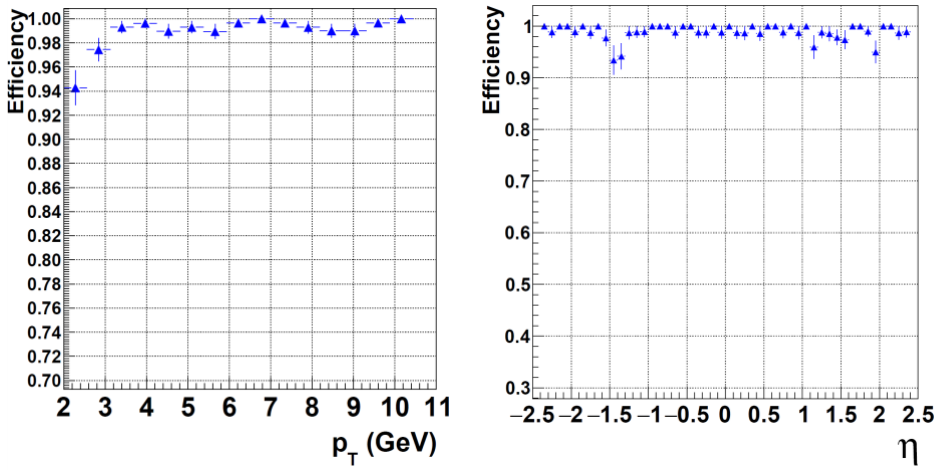

Figure 7.2. Performance estimates in software simulation for a hybrid algorithm implementation, where tracking down to a transverse momentum (p_T) of 2 GeV/c looks feasible.

In 2019 the System Development WG will turn its attention to the first integration tests of hardware at the Tracker Integration Facility (TIF) where the present tracker and readout electronics was assembled and commissioned. The initial test stand foreseen will be modest, with a handful of ATCA prototype cards installed in a crate, including 2-3 Serenity boards with different daughtercards and optics. The area will be used as a test bed for focusing efforts on different aspects of infrastructure, from evaluating the ATCA environment, to assembling the necessary low-level firmware and software to support operations, and gaining experience with the hardware prototypes including one of the first DAQ Timing Hub (DTH) prototypes. Many of these test stand goals, deliverables and expertise overlap significantly with those of WP3. For WP5, it is expected that a first

integration test between boards implementing some of the functionalities of the final track-finder system will be successfully completed in order to help meet the System Development group's first major milestone of defining the system specification towards the end of 2019. In addition, the test stand will provide the group with the opportunity to demonstrate and refine track reconstruction algorithms, where the results will also be used to feed into the system specification.

The most relevant upcoming CMS-wide milestones for the project are the Outer Tracker EDR and the L1 TDR submission, both in March 2020. The UK will be involved in documenting the back-end electronics designs in preparation for these milestones, as well as producing updated simulation results for the L1 track-finding system.

7.4 Staff on project

As mentioned in previous reports, most effort has been transferred from WP5 to WP3, profiting from the significant overlap in hardware R&D between these two WPs. The only staff change is the addition of T. Schuh at RAL on a joint post with KIT, who is contributing in the area of firmware development.

7.5 Expenditure

There has been no significant expenditure.

7.6 Deliverables

The new deliverables list, as set out in the report after Tracker Management approved the FPGA-based L1 track finding solution, is listed below. Blue font means complete. These milestones are consistent with CMS project planning as defined by the System Development WG.

Milestone Date	Description
16Q4	Presentation of demonstrator at CMS internal review
18Q4	Demonstration of track-finder elements using Ultrascale FPGAs
19Q2	Demonstration of the track-finder chain using Ultrascale FPGAs

8 Risk register

The risk register has again been reviewed and revised. Some minor modifications have been made but at this point, there seems little justification for any significant changes, and the overall goals of the project have essentially almost been met. In principle, there are limited risks associated with the final assessment of the CBC3.1 and the Serenity v1.1 versions and the main issue is that it will be challenging to complete all the work considered as part of the R&D project within the remaining time to March 2019, but it should be possible at no extra cost to the project within a further six months.

9 Finances

Expenditure is reported in the usual financial table.

As reported before, RAL PPD staff costs are somewhat higher than originally expected and some errors occurred in the most recent SBS reports. The errors have been identified by RAL and will be corrected but had not been at the time of writing. Nevertheless, it appears that final PPD expenditure will exceed the grant allocation. However, part of this is due to the overheads rate being lower at the time of the proposal compared to during the grant (62% cf. 73%). Inflation was not included in the proposal bid, but was to be applied by STFC. We believe this was done in the university grants but not in the PPD (and TD?) figures. Hence, we believe that the apparent PPD overspend is mostly not real.

At Bristol, the staff effort line also seems to predict an overspend. It is thought that this must be an error, and is under investigation. As is the case elsewhere, the accounts produced by university administrations do not match the format of the financial table, so someone is required to compute the figures for the table, which is sometimes prone to error.

We have made estimates of the major items of expenditure remaining, using current exchange rates, and this is listed in the following table. As discussed in the previous OSC meetings, £100k of the Working Allowance (which is effectively in the Imperial College grant) was used to cover extra TD staff costs on CBC design. Travel expenditure was also higher than originally foreseen, partly because of work carried out on L1 track finder activities, and £40k was committed to this. Note that in the financial table, overspend on TD effort and travel is intended to be balanced by an underspend in the Imperial College grant on capital. STFC agreed to manage this balancing at the end of the grant.

At Imperial, expenditure is classified as capital or consumables within the ordering system, so the consumables "overspend" is believed to be due to misclassification of some items.

The additional items of CBC3.1 orders and Serenity developments were described earlier.

As can be seen from the table, we propose small additional expenditure on TD effort and travel before the end of the grant.

Item	**£k**	
CERN commitments in pipeline	118.8	Not yet invoiced
Commitments at Imperial	85.7	Awaiting deliveries and invoices
3x24 CBC3.1	148.1	estimate @ $2675 per wafer
Serenity further evaluation studies	62.0	optical components, FPGAs, PCBs
"Overspend" on consumables	15.0	Some capital items mis-classified
subtotal	**429.6**	
Overspend on RAL TD effort	100.0	Agreed at past OSC
Overspend on travel	40.0	Agreed at past OSC
Additional RAL TD effort	20.0	
Additional travel FY 2018	44.0	
Total	**633.6**	
Remaining Working allowance	**54.6**	
Original Working allowance	386.6	For comparison

Note that any additional expenditure mentioned for the HGCAL from the ERC grant is not included in these estimates, as the ERC funds are expected to be applied to subsequent HGCAL-specific prototyping after this grant, such as for dedicated high-range FPGAs or board manufacture.

10 Gantt charts

There have been no major revisions of CMS sub-detector schedules compared to those produced in the TDRs. As discussed for the tracker, there have been some updates to reflect certain delays, and a Merlin version of the (large and complex) plan is posted with the upgrade documents. If required, we will endeavour to extract any parts of it in a suitable format for the OSC on request. In addition, a very recent plan for procurement and assembly of the next series of hybrids with 8 CBC3s mounted on them is posted.

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 2018 corresponds to PM66.

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	PM66
M2.5.1	PM42	CBC3.1 ready for production	PM58
M2.5.2	PM48	CBC3.1 produced	PM66
M2.5.3	PM54	Documented CBC3.1 test results	PM72
M2.6.1	PM60	Final production masks prepared	PM66
M2.6.3	PM69	CBC3.1 ready for mass production	PM72
M2.7.3	PM72	First production modules available	***
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.
CIC	Concentrator Integrated Circuit, which follows the CBC.
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 CMS Phase II endcap 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 Gigabit Transceiver.

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.