4 November 2015

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

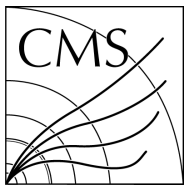

Upgrades of the Tracker and Trigger of the CMS experiment at the CERN LHC

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. For the Phase I upgrades of the L1 calorimeter trigger, this is a construction project.

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

- Work continues steadily on the final CBC design. Manufacturing submission is still foreseen in early 2016.
- Following the understanding of problems with the FC7 reported last time, a new design has been produced and more boards of the existing design (to be corrected by a simple wire) have been manufactured to ensure the CMS TCDS system is not at risk. The manufacture of the new version of the design was delayed until it was clear the assembly problems had been solved, using the existing design but with a new source of the PCBs. Deliveries have recently been received and the quality looks promising.
- Manufacturing of the MP7 is still under way. Since the end of 2015, a series of problems have been experienced, mostly related to the PCB quality. Steps have been take to address this with two new companies, who have successfully provided similar boards to other projects, including the FC7. The first set of boards is emerging from production and the quality looks promising.
- The development of algorithms and firmware for the Level-1 trigger is almost complete and some remarkable work has been done to optimise the firmware designs to minimise the FPGA resource usage.
- Both phases of the calorimeter trigger upgrade project (so-called Stage-1 and the TDR trigger) have been brought into operation in CMS, with Stage-1 now providing the actual L1 trigger and the TDR trigger recently operating in parallel, as intended. While the time remaining for evaluation is short, since the switch to heavy ion running (using Stage-1) is due from 4 November, it now looks probable that the TDR trigger will be satisfactorily demonstrated during 2015 running.
- There has been further progress in developing the Time Multiplexed Trigger architecture for the L1 track-finder needed for the future track-trigger, and steady progress is being made in developing a demonstrator system. It now seems timely to plan a new work package to cover this activity, with defined milestones, schedule and a suitable budgetary estimate, and we aim to do so before the next OSC in 2016.
- A new WP is also planned to report UK activities for the proposed High Granularity Calorimeter, which is largely funded by a substantial recently started ERC grant, with additional support from an STFC PRD post for two years. This WP will also be defined before the next OSC.
- There have been CMS requests for the UK to undertake the development of a new ASIC to replace the UK-designed MGPA for the barrel ECAL readout following the Phase II upgrade. This is under cautious consideration, in view of the likely resources which may be available within the UK.

Outside the project, the evaluation of the university Consolidated Grants was recently carried out by the PPGP. Since most of the resources for this R&D project were already allocated, the impact of the new CG award is limited. However, the long term implications, in view of the reductions to the support for the LHC programme, and CMS in particular, remain to be digested. There was also an impact on a few posts, where individuals were expected to work on CG and upgrade activities and so were supported on both grants. We hope to absorb the effect of these reductions and are in the process of doing so in collaboration with UK group leaders. RAL PPD is presently under review as well.

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} $cm^{-2}s^{-1}$ a couple of years after 2015 in LHC Run 2, 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. There are currently two technical work packages: WP2 for the Phase II outer tracker readout R&D for HL-LHC; WP3 for the calorimeter trigger construction for Phase I, starting in 2015, and further R&D aimed at Phase II.

There have been no further changes reported to the CERN plans for the longer term.

2.1 LHC upgrade schedule and planning

LHC commissioning and operation has gone a little more slowly this year than hoped and it is clear that the 10 fb^{-1} target will not be reached for 2015 p-p running (as of 21 October, LHC had delivered 2.6 fb^{-1} and the integrated luminosity may reach 4 fb^{-1} by the end of the run). The plan shown in fig 2.1 has been maintained, but the luminosity delivered has not been as high as hoped. However, it has been emphasised in the past that 2015 should be seen as a commissioning year for the machine and that operating under the new conditions of higher energy and with 25 ns bunches would be more challenging than in the past. CMS has been commissioned and is operating now with good efficiency during data taking periods but has not been free of problems. Naturally, after two years without operation and new hardware and systems to be commissioned, there were teething issues, most of which have been overcome. While it might be invidious to identify culprits for data losses which have occurred, at least we can report that none of the UK systems, including the trigger, and despite the new components, have contributed much to them.

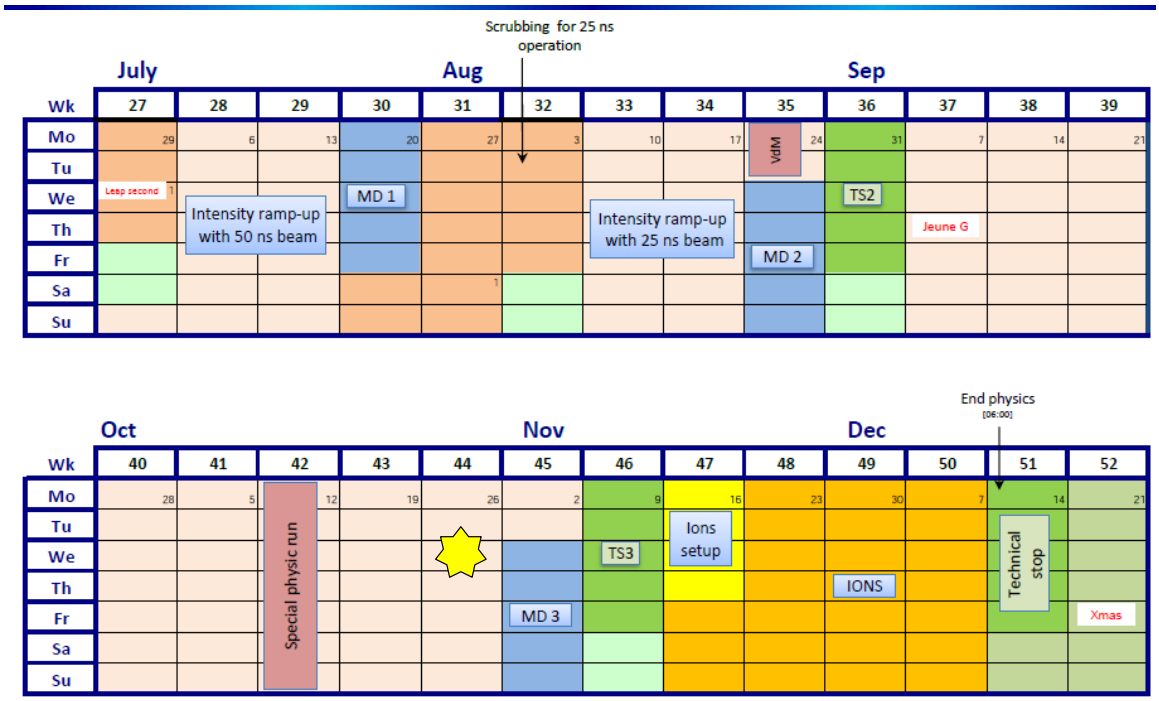

Figure 2.1: The LHC schedule for 2015.

The major concern for several months in CMS has been due to the magnet, where there have been repeated problems since March with the cryogenic system leading to a loss of cooling power and the need to ramp down the field. Overall, this has not led to major losses of data because the LHC operating periods have been adapted where possible to match the need to reset the magnet. However, a fraction of the data has been taken at B = 0 T instead of 3.8 T. The origin of the problem is still not completely understood but believed to be related originally to the introduction of contaminants into

the system because of incorrectly installed filters. This led to clogging of filters in the heat exchanger system which could be corrected only by warming up parts of the cryogenic system and replacing them.

It has been a complicated and lengthy story but mitigating actions were taken in early September which appeared to be very effective. The magnet had then been stable at full field for about a month until the weekend of 24/25 October, where an unexpected problem developed in the early hours of Sunday morning. At the time of writing, this is still under investigation and previous remedies had not been as effective as in the past.

The main objective now is to get the magnet back into operation for the final days of p-p operation, and then the heavy ion run to follow. Deep cleaning of the system is anticipated but can only be done in a lengthy stop such as the 2015/2016 YETS, after which it is hoped that normal operation will be possible. It may be worth noting that such contamination of other superconducting systems has been experienced and corrected so, although there are a number of poorly understood issues, both in CMS and elsewhere, there is good reason to expect normal operation again next year.

As regards this project, the magnet problems have not significantly hampered trigger commissioning.

2.2 CMS planning

CMS delivered the Upgrade Technical Proposal[1] to the LHCC in June, in its final form; it had been held back to allow the decision to be made on the endcap calorimeter options in May, as reported last time. Shortly afterwards another document was prepared summarising the R&D required to deliver the detector, with an estimate of the personnel resources needed, which agencies plan to contribute to them, and major milestones. This was requested by the RRB in April and was provided at the October RRB meeting, as was the likely schedule for the construction stages and dates of TDRs (Fig. 2.1)

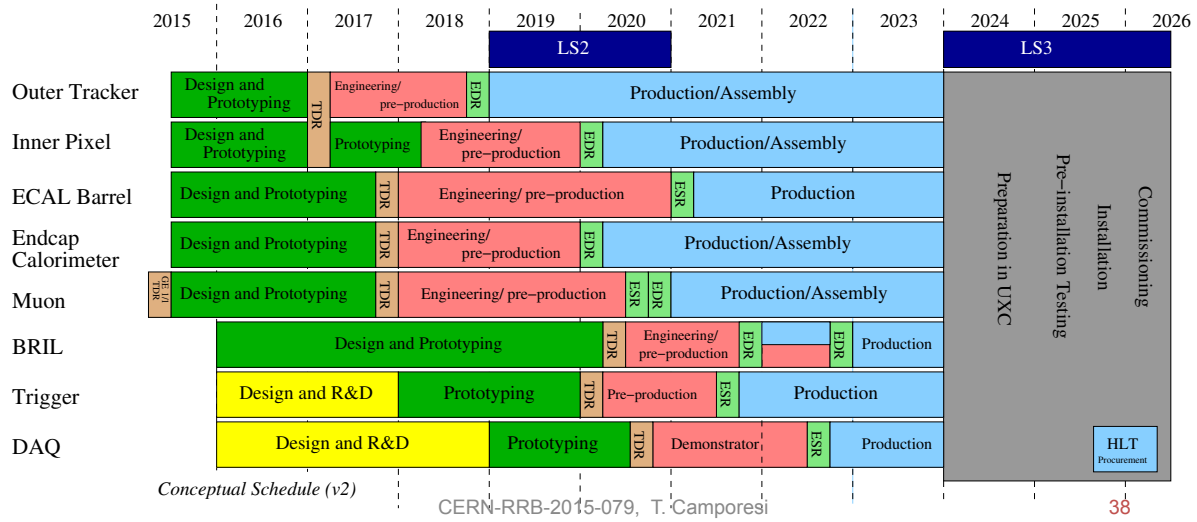

Fig. 2.1 The overall schedule foreseen in CMS Phase II upgrade planning, including TDR dates.

In September another document was prepared for the RRB, again reporting the cost estimates for the upgrade (265 MCHF), the so-called reference design, along with two other scenarios which would reduce the cost, one to approximately 90% and the second to 75% of the full cost, with a discussion of the implications. They lead to cost estimates of 242 MCHF and 208 MCHF respectively.

The document summarises the cost optimization in the reference design and aimed to complement the Technical Proposal by assessing the capability of the baseline design to profit from ultimate

[1] CERN-LHCC-2015-010 https://cds.cern.ch/record/2020886

luminosity (200 PU), i.e. to ensure that CMS would not fail under such conditions and assess its performance. This required full simulation studies of physics object reconstruction for selected representative crucial channels, such as VBF $H \to \tau\tau$, $H \to \mu\mu$ and $H \to 4e$.

The studies also compared the performance of the reference design and the two lower cost configurations at 140 PU. The reduced configurations were distributed across sub-systems, using the reference detector with specific elements reduced or removed, with a relatively more substantial downgrade of the High Granularity Calorimeter, for which the design is at an earlier stage of optimization. Representative configurations were defined in an attempt to minimize degradation of the standalone and combined detector performance to preserve the physics reach, but without full optimisation of the designs, which would be a lengthy process. The reduced configuration also reduces the computing power available to online event selection, limiting the trigger rate capability.

As expected, since the reference design is for 140 PU, the upgrade provides sufficient, rather than optimal, performance at 200 PU, with slightly higher degradation observed for some of the calorimeter measurements, which it is envisaged may be mitigated by new pointing and possible timing capabilities of the High Granularity Calorimeter which have not yet been studied in simulations.

While the lower cost designs might be further optimised, this method is expected to give a realistic picture of the likely performance with a reduced detector. Some of this might be recovered later, if future funding were to permit, but several of the items would be non-recoverable, e.g. if layers were removed from the tracker. It is inevitable that, in the harshest cost reduction scenario, significant penalties on the design must be incurred and that an impact on performance should be expected.

The studies at the baseline luminosity show that compensating the performance degradation of the 90% configuration would require about 20% more LHC operation time, for important sections of the physics programme. Partial recovery of the losses may be possible with further design and descope optimizations, and the use of additional computing power as the luminosity increases. Nevertheless, it is very likely that, for some crucial signals that are limited by statistics, it will not be possible to benefit significantly, or even marginally, from the higher luminosity potential of the accelerator.

The configuration costing about 208 MCHF was also evaluated. It implemented increments in scope reduction of the upgrade elements. Simulations of this detector configuration demonstrate that already at the baseline luminosity the physics program will be adversely affected in all thematic areas, with some major physics goals becoming unachievable. It was also demonstrated that, in addition to the expected loss of performance, substantial descoping or downgrading of the upgrade elements will also create important operational risks arising from the loss of redundancy in the detector measurements, which perhaps should be of even greater concern.

2.3 UK adaptation to CMS planning

The major question which has emerged recently has been the planning of commitments and responsibilities for the Phase II upgrade, including especially an assessment of the resources likely to be available to sustain it. This is being addressed by seeking input from the sub-detector project managers from their discussions with likely participants and with regional representatives to confirm potential availability of funding and the status of, and procedures for, allocation of resources in each country. Obviously this is delicate process but it is important to distinguish between possibly optimistic estimates and eventual reality, while it is well understood that funding agencies are probably unable to make commitments at present until the process for approving the experiments, defining the schedules, allocating resources in CERN both to the accelerator and the experiment upgrades are on a completely sound footing. The present economic climate, of course, does not make this easier.

The UK has had a quite well defined process for allocating R&D funds and is probably well ahead of most other agencies in that respect. In fact, there seem to be few other agencies, unless CERN is included in this category, who have had such a well-planned process to anticipate the necessary R&D and ensure it was reviewed and resourced. The situation for the future, which in our case means from 2019 onwards when the present R&D funds end, is less clear but that has certainly been partly due to the lack of clarity in CERN planning. The discussions over the last couple of years, and especially in the last year, in the RRB seem gradually to be bearing fruit by clarifying the CERN picture and

encouraging agencies to try to provide information which would allow a better assessment of the overall resources likely to be available to CMS.

CMS has compiled a table of expected CORE cost contributions, where each entry by nation has been checked with the agency representative before or during the last RRB. This has been kept confidential so that agencies only know their own estimates but the CMS Spokesperson and a few others have seen the entire table. This has recently been shared in confidence with the RRB Chair (the CERN Director of Research), the LHCC Chair, and the small committee charged with assessing the upgrade cost estimates. This is expected to lead to a decision by CERN and the funding agencies of the target budget which should be assigned to each experiment.

Meanwhile in the UK, we have held two meetings of the collaboration since July to discuss our potential commitments, which we are basing on the indications of possible funding from 2019 onwards based on the outcome of Science Board deliberations in 2014, as explained to us by STFC. We are well aware of the significant uncertainty which may affect these estimates in view of possible changes to government support for science. At present, we have an idea of which projects the UK hopes to participate in and unsurprisingly the areas on which we have been concentrating our R&D are the principal ones.

However, over the last year or so another candidate for significant resources has emerged, in addition to a coherent project including tracker, track-trigger and possible L1 trigger contributions, which is the replacement of the endcap calorimeter, the High Granularity Calorimeter described last time. As a reminder, fig 2.2 is repeated.

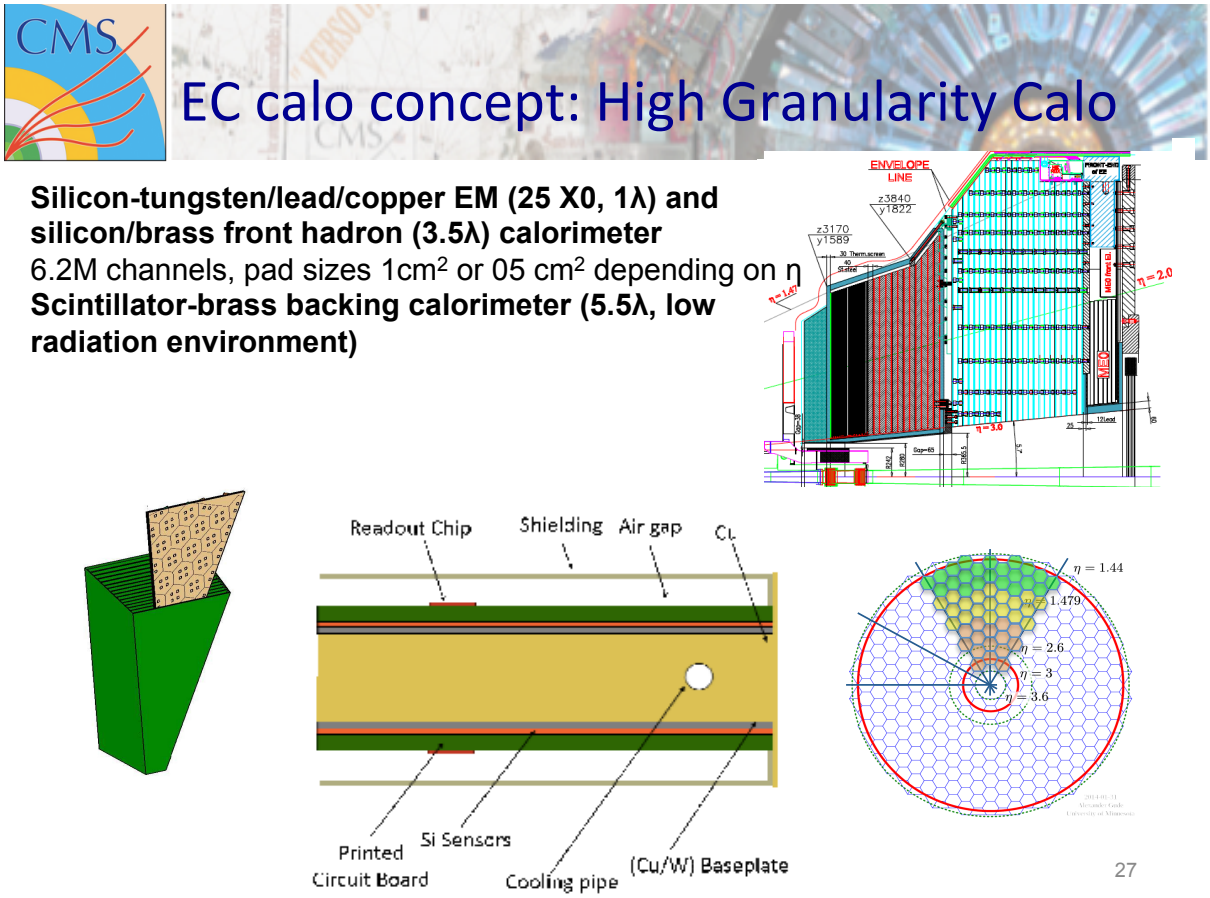

Fig 2.2 High Granularity Calorimeter concept

As mentioned in our submission to STFC in March 2014, CMS realised relatively recently that it needed to change the endcap calorimeters for the HL-LHC phase. Effective endcap calorimetry is essential to the CMS physics programme, and has a particular impact on Higgs physics. Given our responsibilities for the Endcap ECAL, and involvement in Higgs physics, we were clearly interested in

this subject. After a two-stage down-selection process, the HGC was chosen in spring of this year as the technology of choice and included in the Technical Proposal. The UK has made major contributions to this project, from the instigation of the initial concept (T. Virdee) and significant contributions to initial physics studies, up to now mainly from Imperial College. Virdee is the interim project leader. The full organisation chart is given below, fig. 2.3.

Key issues that need to be addressed for the TDR, planned for late 2017, include detector optimisation, the design and testing of the front-end electronics and the trigger / back-end electronics and the mechanical / cooling aspects. We have been successful in attracting new funding to allow us to address several of these issues, exploiting the existing UK leadership and expertise. A two year STFC PRD will address detector optimisation and development of trigger algorithms. An Advanced European Research Council award for ~3M€ over five years has funding for several positions, including a postdoctoral position to continue the PRD work longer term, as well as electronics engineering resources to tackle aspects of both the back-end / trigger and front-end electronics. Funding for prototyping was also requested, and we expect to be able to contribute ~0.5 MCHF to the construction cost book.

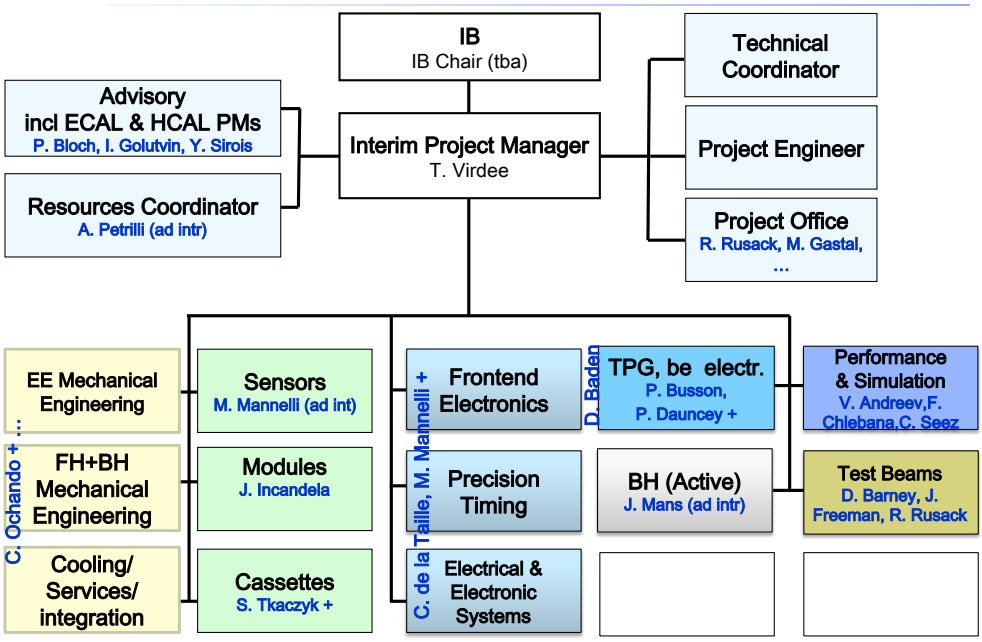

Fig. 2.3 HGC Organisation Chart.

The overall HGC project is being broken down into work packages, with responsibility being assigned to interested institutes. A proto-money-matrix is also being crafted. We are in the process defining the UK roles, and hence developing an appropriate UK work-package. There is a natural synergy, both conceptually and practically, between the HGC trigger work, and our existing upgrade activities and we would expect this to be a central aspect of our HGC construction commitment. One of the advantages of the HGC is the greater information available, and hence involvement in the development of the reconstruction software is planned. We currently convene both relevant groups (P. Dauncey, C. Seez); see the organisation chart (fig. 2.3). While we will contribute to the development of the front-end electronics or its associated testing over the 5-year period of the ERC grant, an actual construction period commitment is less clear and will depend on the interplay with other CMS upgrade commitments. We will present a detailed UK work package to the Oversight Committee at the next meeting in 2016.

In addition there has been a request from the ECAL project to consider the task of designing and testing a new amplifier for the APD photosensors of the barrel ECAL, in view of the expertise which exists in the UK in ASIC design in RAL Technology Department and that the existing amplifier, used by both endcaps and barrel, the MGPA, was designed by the UK, at Imperial College, with the circuit layout done by RAL TD. Redesign of the front-end ASIC, to incorporate increased functionality such

as anomalous heavily ionising APD "spike" rejection, is a key element of the ECAL barrel upgrade, which is described in the Technical Proposal. We are currently investigating the implications of such work, including whether it might be compatible with front-end designs for the HGC or if it would need to be a design specifically for the barrel ECAL, and especially the resources required. Provisional answers should be available before the next OSC meeting.

2.4 Progress with track-triggering for Phase II

Following the publication of the Phase II Technical Proposal in June, new conveners have been appointed to coordinate the L1 Track Finder Working Group within CMS. As part of this reorganisation, a set of specifications for the track finder demonstrator systems has become better defined with a strong emphasis on common software tools for event generation and comparison of results. While coordination of the group remains US dominated, we continue to provide technical input towards the common tools, and feedback on demonstrator specifications, definitions and milestones.

The UK has presented a concept for a L1 track finder for Phase II, based on a time-multiplexed architecture. The demonstrator system we propose is based on a small number of MP7 boards and is required to transfer Monte Carlo generated trigger primitives (high transverse momentum hit candidates, or *stubs*, eventually provided by the Phase II tracker) through a set of processor cards to generate fitted L1 tracks for a time slice of the full system. The demonstrator slice will also be regional, but we plan to ensure the demonstrator can be easily reconfigured to process data from any region of the tracker.

As described in our previous report, the first part of our demonstrator system is focused on the Hough Transform track finder processor which implements an algorithm to collect stubs compatible with real high transverse momentum tracks into candidates before further processing. To this end the first demonstrator system is a single board that tests a fraction of the track finder. Data are loaded into the link buffers over IPBus and played through the processor before the results are stored in the output buffers, ready for readout. At this stage a small but scalable system such as this is important to operate, and requires us to ensure that the crucial infrastructure to drive the system and verify its results is in place. However we also expect that such a system will already help us start to understand the impact of the track finder on latency, and specify the hardware that might be required to implement a full scale design.

Recent progress includes:

- The Monte Carlo simulation software we have developed has been used to test the phase space of the Hough Transform algorithm we want to operate, identifying a minimum set of specifications for the track finder in hardware.
- Firmware has also been written implementing a 2D (r, ϕ) Hough Transform for binning stubs into candidates and is ready to be deployed on the demonstrator MP7s. We expect further firmware optimisation in due course to fully exploit the resources on the MP7, including variations on the implementation of the array, so as to minimise both FPGA resource usage and processing latency.
- Online software is ready for both loading stub data and configuration parameters into the MP7s, and for reading out binned candidates. Emulator software has also been developed to ensure that we have an accurate representation of the performance of the track finder hardware to feed back into our simulations. In addition, a generic software framework that could be used either to easily construct emulators in future, or quickly test alternative firmware concepts for dataflow and latency limitations, has been in development and will be put to use on the track finder in the next few months.
- Two MP7s and associated infrastructure have been assembled for a first demonstrator system at RAL, while a parallel system will be set up at CERN in future.

- Increased emphasis and effort has been placed towards simulation of the algorithms required following the Hough Transform candidate finder, particularly on the track fitting, which will eventually allow us to compare the track finder performance against official CMS benchmarks.

In the next few months the majority of our work will include: operation of the demonstrator systems bringing together the strands of work described above; exploring more optimised firmware designs making use of the demonstrator and common software/infrastructure we have developed for testing the concepts; and further investigation of the algorithms that the track finder would require in order to demonstrate reasonable performance.

Work on the track finder simulations and firmware development has been undertaken in collaboration with individuals from CERN, Vienna and Karlsruhe.

Over the last six months and more this activity has attracted an increasing number of UK participants, and external collaborators, and we have seen excellent progress in simulations, defining specifications and firmware development. Because of the pressure from the trigger upgrade, some individuals who will make valuable contributions in future, for example on the complex firmware development, have not so far been able to devote much time to this activity. However, as the commissioning of the calorimeter trigger converges, as it now appears to be doing, we expect more effort to be available for the track-finder work. In view of this, and other factors, such as the potential evolution of hardware for the future described in the WP3 report, it is timely to structure the track-finder work more clearly to define objectives and responsibilities, and maintain them. For this reason, we propose to define a work package specifically for this activity and will present this in more detail at the next Oversight Committee meeting in 2016, with a suitable list of milestones, and identify the individuals who will be responsible for managing it.

3. Work Package 1: Management

A reminder of the project management is included below. G. Hall stepped down as UK CMS PI at the end of September 2015, and succeeded by D. Newbold, but he will remain responsible for this project. Hall has been asked to act as an advisor to the CMS spokesperson, so will continue to attend CMS MBs for another year.

J. Brooke from Bristol has taken over the role of D. Newbold as co-manager of WP3. He has previously been heavily involved in the trigger offline software, including emulation activity, as well as commissioning. The WP3 managers are not intended to have different roles, although they bring complementary expertise to the project, but to share the duties of steering the activities. This is expected to evolve further as the Phase I trigger now enters the operational phase.

Two new work packages are planned, to be described in detail in the next OSC report, with milestones and schedules. One, WP4, should report activities in the calorimeter area, which is expected to be dominated by the High Granularity (Endcap) Calorimeter activity. A summary of the status is given below, but schedules, milestones and resources will be presented next time.

WP5 will cover the ongoing work on the L1 track-finder, which will be part of the CMS Tracker project, though in our case will integrate activities from the UK tracker and trigger teams who have already been collaborating for some time. Now that the delivery of the Phase I trigger upgrade is almost complete, we can profit from the availability of several individuals who until now have been dedicated primarily to that task, and define a schedule and suitable milestones for a reasonable period ahead.

WP	Manager	Institute	Role
1	G Hall, PI	Imperial	Overall management, budgetary responsibility and supervising procurements, interface to CMS, as UK CMS PI and CMS Management Board and 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	A Tapper	Imperial	Based in CERN with supervisory responsibilities for G. Iles, Imperial College engineer, also based in CERN. Currently project manager for the calorimeter part of the trigger upgrade.
	J Brooke	Bristol	Supervision of UK Phase I trigger upgrade activities, with A Tapper. Previously DPG convenor and software coordinator in the trigger project.
4	tbd	Imperial	To cover calorimeter activities, primarily the HGC developments
5	Two co-managers tbd.		To oversee and plan the 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 design of the final prototype of the readout chip for the 2S-PT modules, the CBC3, is nearing completion. All the design modifications and additions are mostly complete, with mainly layout and post-layout simulation tasks remaining.

The CBC3 front end amplifier design has been optimised for final requirements, one of these being a shorter overall pulse shape. While the confinement of a signal to a single bunch crossing can be achieved by a front end amplifier output pulse with a sufficiently fast rise time, there is a dead time penalty if the return to baseline takes too long. Figure 4.1 shows that a pulse shape is achieved which rises to a maximum well within 25 ns and returns to the baseline (with minimal overshoot) within approximately 50 ns, for a wide range of simulation conditions. The simulated noise performance is also shown, where the required specification of less than 1000 electrons is met for external capacitances up to 10 pF, compatible with 2S module sensor strip lengths of up to 5 cm.

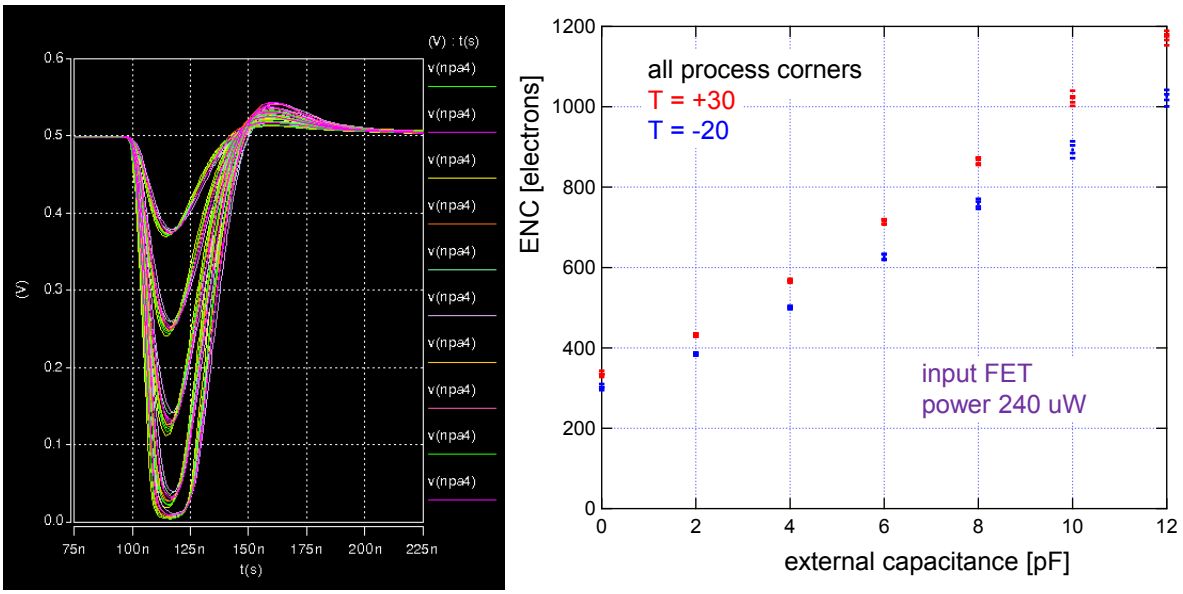

Figure 4.1 On the left is the simulated CBC3 front end amplifier pulse shapes for signals in the range 1 to 5 MIPs (2.5 to 12.5 fC). Pulse shapes are shown for all process corners, for two temperature values (-20 and +30 deg. C). The plot on the right shows the noise dependence on external capacitance for the same simulation conditions.

There is a possibility that longer sensor strip lengths of up to 8 cm may be employed in the 2S module region of the tracker, if sensors of satisfactory quality can be produced on 8 inch wafers. A

higher current in the input transistor will be required to keep the noise low, and the CBC3 bias circuitry has been adjusted to allow for this.

The analogue front end circuits have been modified to run at a supply voltage as low as 1V, which is supplied by an on-chip low dropout regulator (LDO). The input to the LDO powers the digital circuitry as well, and can be as low as 1.1V, which allows good tolerance to module power supply variations. To achieve the lower voltage operation the option of compatibility with both sensor polarities has now been removed and the chip is now compatible with n-on-p technology only (the CMS HL-LHC tracker choice).

The front end optimisation has mainly involved relatively minor changes to the existing design, while there are completely new digital functionalities required in the CBC3. Most of these blocks are now complete, through to post layout simulation, including the stub gathering logic, the data packet assembly and transmission logic, the serial output shift register, and the fast control interface.

We have revisited the design of the on-chip reference voltage with a view to improving the stability to radiation, particularly displacement damage, and we are planning to use a PMOS based voltage reference designed at CERN, which has been shown to be very robust to both ionizing and displacement type radiation effects. While stability is good, the absolute value of the reference voltage is strongly process dependent, so we are planning to trim the value during wafer test using e-fuses (electrically programmable) which can be one-time-programmed. At the same time (during wafer test) we propose to include an e-fuse programmable chip identity on the chip so that every chip in the system will have traceability, with associated test records in a database.

Since the last report we have participated in a beam test of an irradiated mini-PT-module at CERN in June, where the module had been subjected to a maximum HL-LHC dose and had therefore to be kept at a low temperature (-20 deg. C) during operation. Full size 2S module sensors have recently become available and a full-size module using the 8CBC2flex hybrids has just been constructed (figure 4.2) and we are currently preparing for its evaluation in a test beam in November.

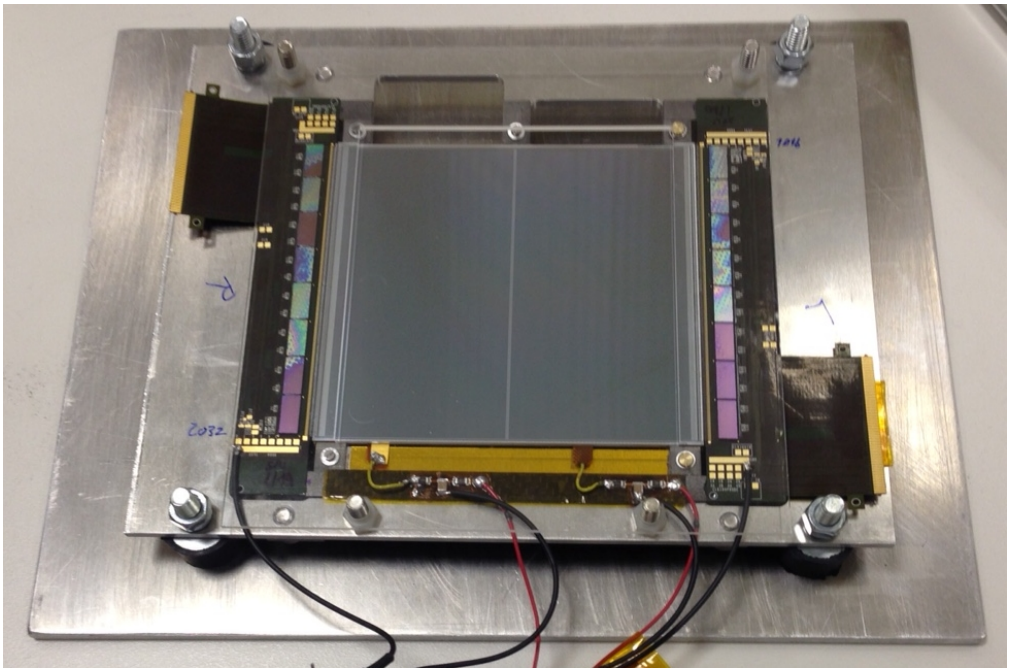

Figure 4.2 Full-size 2S-PT module prototype. Two layers of silicon sensors (only top layer visible), ~10 x 10 cm^2, each segmented into 2032 5cm microstrips, read out at each side by two 8CBC2flex hybrids.

FC7 developments

The FC7 (FMC carrier - Xilinx Series 7) AMC is based on the MP7 and CERN GLIB boards with the ability to host up to two FMC (FPGA Mezzanine Card) modules and support signalling rates up to 10Gbps. It is a flexible, general purpose card allowing it to be deployed across many applications in CMS and its development has been shared with engineers from CERN.

In our last report we described the failures associated with an overvoltage condition on one pin of the FC7 FPGA and steps taken to prevent further FPGA failures, and correct the design for the future.

- In TCDS, where an initial 12 failures were observed out of 50 R0[2] prototypes at P5, all remaining functioning boards were corrected for the design error. Since then no further failures have occurred except when boards were power cycled once, suggesting they have been significantly weakened by the operating condition and replacement will be necessary during the YETS as planned. However there are sufficient spares available for urgent use if necessary.
- No failures have been observed on the R1 version of the FC7 (total 25), specifically because the boards were corrected before any significant long term use and therefore neither Phase I Pixel DAQ nor Phase II Tracker DAQ have been affected by this – although a small number of R1s have been lent to TCDS to act as spares.

We also reported on ongoing manufacturing issues with our supplier (Hapro, Norway). From detailed investigations of the PCBs it was clear that the boards were not being manufactured to the electrical class specified and we concluded that the fault partially lay with Hapro for relying on an intermediary company to manage supply and control of PCBs. (There is a similarity here with the MP7 contract.) As a result:

- Given the urgency for more boards at the time, a further order of 30 R1 FC7s was placed without pressing this issue further. However, no more boards have been successfully manufactured to date due to delamination during assembly, leading to an additional 3 month delay. Further PCB issues have recently been reported with the subsequent batch.
- At the time, to reduce the risk to TCDS in the event of a delay such as this, but also to trial new manufacturers, an order for an additional 35 R1 boards was placed with BBElectronics (Denmark), who were highly successful with the manufacturing of the GLIB (200+ pieces over 3 years). We also requested that the PCBs be manufactured by SOMACIS (Italy), who are now also manufacturing the MP7 XE.
- The R1 boards from BBE were recently delivered and not only is the assembly of very high quality, the PCBs demonstrate excellent characteristics (thickness, rigidity, edge connector tolerances) while meeting the electrical specification. The yield is very good (33/35) from initial testing, and the more comprehensive feedback provided by both SOMACIS and BBE has already proved useful for future orders.
- CERN has now placed a long term purchase order with BBE for manufacture of 250 boards with additional options. It was decided to delay the placement of the order until the yield from the first batch was known.
- This order will be for R2 boards, which includes the overvoltage correction to the R0/R1, as well as other simplifications to the design in order to guarantee compliance with Xilinx recommended guidelines for the FPGA. The first batch of 65 is currently in production with an expected delivery end-2015.

[2] R0 is the FC7 prototype, although since the TCDS project had a very tight schedule it was decided to manufacture R0s before they were fully qualified; these R0s were subject to the overvoltage fault over many months because of the pin misconnection.
R1 was a minor update of the R0, implementing small feature changes to ease the acceptance testing process – however since the order was submitted a few months before the R0 design error was spotted it is also subject to the same fault. Fortunately we implemented the workaround on the R1s soon after they were manufactured so that the risk of FPGA failure due to overvoltage is negligible.
R2 is an updated version of the R1, both fixing the design error but also implementing conservative changes to ensure that there are no deviations from Xilinx recommended practice, along with other small layout and feature updates.

4.3 Deliverables

The WBS for WP2 is included below. As explained in the previous report, we were planning to submit the CBC3 for manufacture through MOSIS in February 2016, but we are now expecting to join a CERN organized submission with some other groups developing chips in the same process, which will provide us with many more chips than we would get through MOSIS and may also lead to cost savings. We are still working towards being ready for submission in February, which is later than planned in the WBS, but we are well aligned with our collaborators developing other module components (concentrator ASIC, low-power GBT and optical transceiver, DC-DC conversion), where prototypes are expected in the second half of 2016.

WBS	WBS L2	Start	Finish	Months	Task Description
2	Phase II tracker Readout	04/13	03/19	72	
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 2S-PT module studies, in collaboration with CMS, including test beam
2.3	**CBC3**	04/14	03/16	24	**CBC3 is specified for the final system**
	2.3.1 CBC3 design	04/14	09/15	18	design period
	2.3.2 CBC3 production	09/15	03/16	6	production period
	2.3.3 test setup preparation	09/15	03/16	6	wafer and chip test setup preparation
2.4	**CBC3 test**	03/16	03/18	24	**CBC3 chip and module testing**
	2.4.1 early tests	03/16	09/16	6	chip verification tests to prior to module tests
	2.4.2 ongoing testing	09/16	03/17	6	complete characterization, including irradiation and SEU tests
	2.4.3 CBC3 2S-PT module studies	09/16	03/18	18	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 remaining bugs found in the CBC3**
	2.5.1 CBC4 design	09/16	12/16	3	design period
	2.5.2 CBC4 production	01/17	06/17	6	production period
	2.5.3 testing	07/17	12/17	6	tests to verify full and final functionality
2.6	**CBC4 mass production preparations**	01/18	12/18	12	**a full wafer engineering run is required for CBC4 in preparation for mass production**
	2.6.1 CBC4 final masks	12/18	03/18	3	mask preparation for full wafer engineering run
	2.6.2 CBC4 engineering run	03/18	09/18	6	production period
	2.6.3 CBC4 final production readiness verification tests	09/18	12/18	3	final functionality check
	2.6.4 procurement planning	01/18	12/18	12	detailed financial plans for mass production

4.4 Staff on project

Michelle Key-Charriere and Mark Prydderch continue to work together on the project and in late autumn should be joined by Przemek Mroszczyk and Lawrence Jones, to ensure the completion of the CBC3 design in time for submission in early 2016.

One post affected by the CG award was J Fulcher, from Imperial College, who is the online DAQ expert dedicated to the maintenance and operation of the Tracker FEDs and online software, and who was 50% funded by this project and the CG. His CG support was removed, which is a major concern, although STFC have helped us re-establish the 50% upgrade funding. His contributions have been principally in support of beam tests, with the expectation of significant contributions to online software for the upgrade in future. We are still in the process of discussing how to handle this.

4.5 Expenditure

The main expenditure continues to be on RAL TD staff. The submission is foreseen in early 2016, as mentioned above, so the so the manufacturing cost is expected to be invoiced in FY2016-17.

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 Stage 1 calorimeter trigger for 2015 was commissioned in parallel with legacy trigger operation during the 50ns run, subsequently taking over from the legacy calorimeter trigger for 25 ns running. The UK made very significant contributions to this project in the areas of offline software, jet finder firmware and commissioning, in addition to providing the hardware and corresponding core software and firmware. In particular the commissioning of the DAQ link required extended work at P5 from several UK personnel. The Stage 1 trigger is now running reliably and selecting physics events for CMS, requiring very little support from UK personnel. A second MP7 board was supplied for the development of heavy ion algorithms by the MIT group. This development is now complete and being commissioned ready for the LHC heavy ion run later this year.

Since the last report, commissioning of the TDR trigger system has progressed significantly. After a lengthy programme of testing using Monte Carlo generated patterns, injected into the inputs of the trigger system, the algorithms were shown to be working satisfactorily. Testing with the CMS DAQ system, benefiting from the experience of Stage 1, was completed in early October and since then the system has entered parallel global running in with CMS on several occasions, often running for several hours without problems. The highlight so far has been collecting more than 1 billion events in global running (see Figure 5.1 for an example). These data are being used to validate simulations of the trigger and measure efficiencies and rates. This will allow the tuning of trigger conditions for the start-up of the LHC next year. In addition the latency of our part of the trigger system has been measured and meets specification.

Work on testing and refining algorithms for the calorimeter trigger in conjunction with the LLR group has been ongoing. As a reminder, the responsibility for jet and energy sum algorithms belongs to the UK, while the electron/photon and tau algorithms are the responsibility of LLR. The integration of the firmware from the two groups proved an extremely involved task, requiring intense effort from the UK side. The algorithms have now been integrated and optimised such that the latency and resource usage in the system is acceptable. Expected performance results, for these algorithms, and commissioning progress were presented at the TWEPP 2015 conference in September.

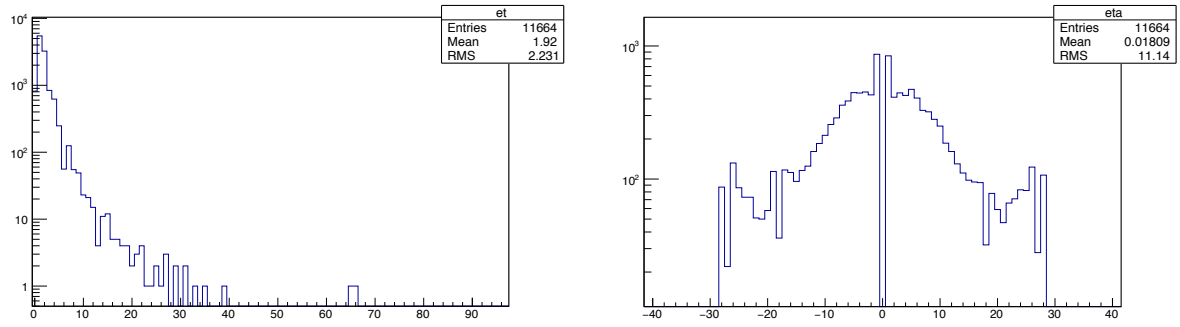

Figure 5.1: Jet E_T in GeV (left) and η (right, scaled units) from data taken in global CMS running.

Two hardware aspects remained outstanding at the time of the last report: the redundant power system which did not operate quite as expected, and the last part of the fibre installation between the processing nodes and the demultiplexing card. The former was solved with the help of the manufacturers and the latter installation is postponed until after the current pp collision run.

All the boards necessary for the UK contribution to the CMS trigger system, including spares, have been installed and successfully commissioned. However a number of extra MP7 boards were ordered for the muon and global trigger projects. Again a significant challenge over the last six months has been that the MP7 procurement has not proceeded as expected.

At the time of the last report a formal warning had been sent to Exception PCB, which led to improved communication and eventual completion of the contract.

Procurement of the additional MP7s for the global trigger and muon trigger projects has also encountered manufacturing issues, despite the same company, Hapro (Norway), successfully building earlier batches of MP7s.

In both cases the primary issue has been the availability of high quality PCBs rather than the assembly. A series of meetings with PCB experts at CERN, in addition to visits to PCB manufacturers has resulted in two alternative manufacturers being identified (SOMACIS (Italy) & TTM (US)). An additional thermal shock test has been added to those normally undertaken at the PCB company to make sure that the PCBs are robust enough for the assembly process. Orders have been placed with both companies to mitigate the risk of one company failing to deliver. PCBs are expected back in the first two weeks of November. Assembly will commence after the PCBs have undergone an extended inspection process at Hapro so that first batch of two cards will be delivered towards the end of November. The second and third batches will follow shortly afterwards.

Developments for the future

Because of the very high data volumes, the trigger and readout electronics used in particle physics detectors has traditionally been application specific, but over the last decade advances in FPGA and optical technology mean that a common processing board could now be used in the vast majority of cases. This idea is now (slowly) gaining traction within the wider collaboration. Possible approaches are a common CMS-wide project, with the workload spread among participating institutions, or a path of commercialization where hardware is bought in, with a price to cover the cost of on-going support.

As previously reported, we are currently in the process of designing next generation hardware. Given the upgrade timescales, there is no need for “final” hardware immediately, and the emphasis must, instead, be on evolving a common-hardware philosophy: developing something sufficiently similar to final hardware to allow construction of demonstrator systems, keeping pace with technology evolution but developing software and firmware infrastructure which abstract away the evolution of the hardware.

Hardware design

We have now embarked on a successor to the MP7 board whose key limitations were

- the 100W limit on the power supply of the µTCA form-factor
- repeated difficulties related to manufacturing to the µTCA specification
- PCB yield issues related to the chosen stack-up
- relatively high cost
- less importantly, some concerns about the amount of RAM available on the board.

The successors to the Xilinx 7-series are the Xilinx Ultrascale and then the Ultrascale+, and so the successor to the MP7 is (currently) named the MP-Ultra. Our roadmap for this board spans these next two generations of FPGA, utilizing the cross-generational pin-compatibility which Xilinx has planned.

The first generation of card (fig 5.2) is based on the Kintex-Ultrascale FPGA, providing 48+48 optical links at 16Gbps. For typical particle physics trigger requirements, this FPGA is the most desirable of the Ultrascale series, having the highest logic-bandwidth ratio and the lowest cost-

bandwidth ratio. The low cost of the FPGA makes this the most cost-effective for the development cycle of the card.

The second generation of the card will be based on a much larger Virtex-Ultrascale FPGA, providing 96+96 optical links, again at 16Gbps. This board provides far more bandwidth than the first generation, but a lower logic-bandwidth ratio, approximately the same as that on the MP7. Crucially, features such as the additional optics, the power requirements, etc. have been designed into the first generation card, and should not need modification for the second or third generation

The third generation of card is identical to the second, but with a Virtex-Ultrascale+ FPGA when these devices are released, providing around four times the logic resources of the second generation.

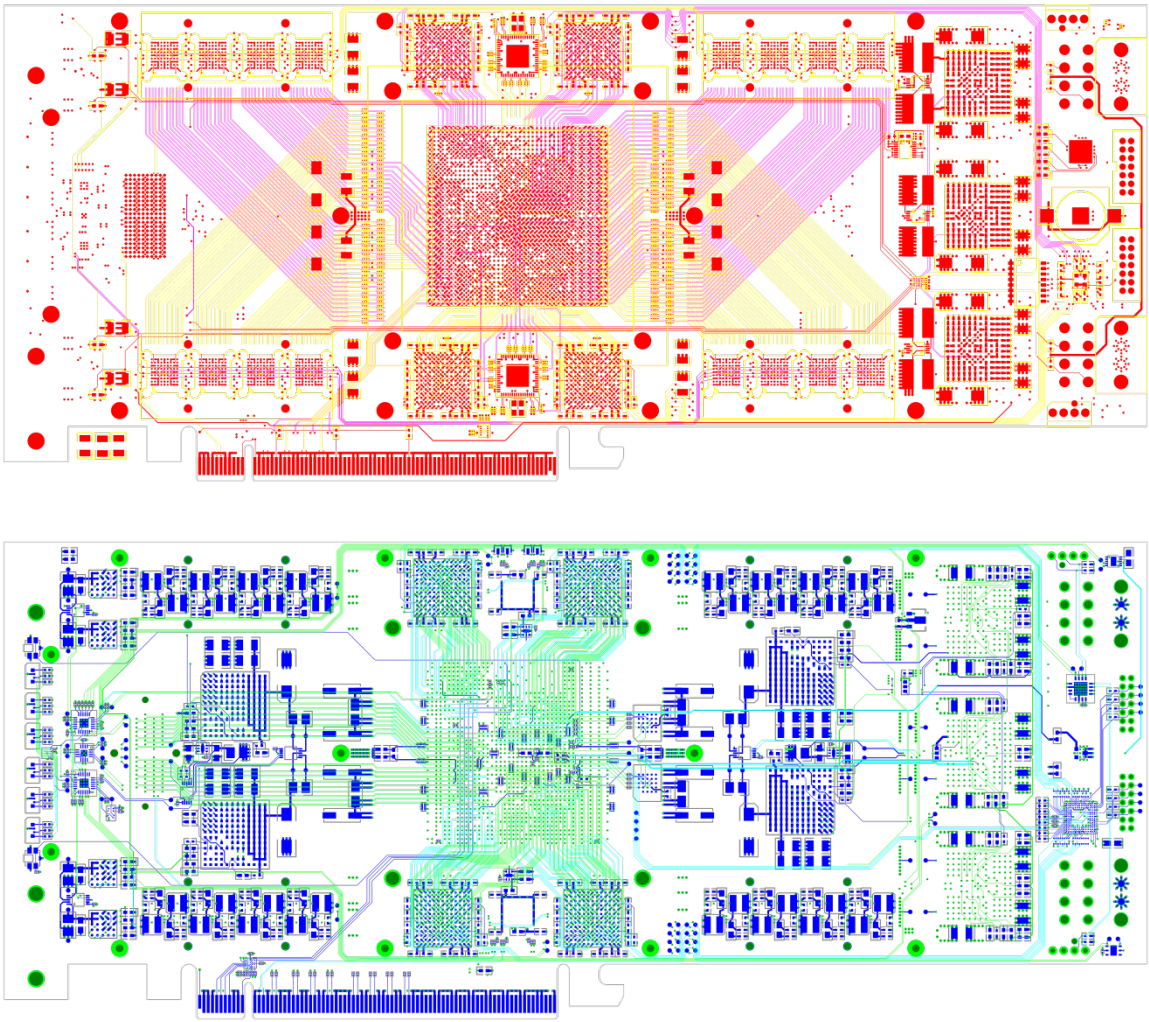

Figure 5.2: Current state of the top and bottom stack layouts of the MP-Ultra. The key features on the top surface are the Ultrascale/Ultrascale+ generation FPGA, four 1Gb RLDRAM3 chips and two ultra-high performance clock synthesizers surrounding it, and four groups of four optical transceivers, each providing 12 optical links, giving 96+96 TX+RX in total. Exposed pads provide 32 differential or 64 single-ended general-purpose I/O through a daughter-card by means of a one-piece connector. On the bottom surface there are an additional four 1Gb RLDRAM3 chips and, at the front of the board, the peripheral control circuitry: The rest of the bottom surface and the back (right) of the top surface are mainly used for powering the various components.

Rather than target µTCA, the board has been designed in a PCIe form-factor, alleviating constraints on power and mechanical tolerances, making manufacturing simpler. This has advantages that developers do not need additional, expensive infrastructure, as the card can fit any standard

desktop PC, and we can utilize off-the-shelf components designed for the mass-computing market. For installation in the experiment, it may be desirable to have use a crate-oriented form-factor, such as ATCA. The conversion of schematics from PCIe to an ATCA form-factor has been tested and demonstrated to take a few hours so, for the present, the smaller, more ubiquitous PCIe format is preferred.

The board has been designed to provide up to 8Gb of RLDRAM-3 (compared to the 144Mb of QDR-II+ RAM on the MP7) arranged in 4 independent banks, which is far cheaper and has higher capacity than the competing QDR-technology used on the MP7, at the cost of slightly lower bandwidth.

The board stack-up has been changed from that of the MP7 to one independently identified by three PCB manufacturers as being the easiest to manufacture.

Progress

Initial layout is almost complete (fig. 5.2) with only minor changes, tweaks and corrections expected to layout and schematics over the next few weeks. It is currently proposed to manufacture the PCBs through the preferred MP7 manufacturer (SOMACIS, Italy) and assemble the cards in the UK, so we would expect to have first prototypes in hand in January 2016.

5.3 Overview of CMS plans

At the time of the last report the Layer-1 electronics (the responsibility of University of Wisconsin) were not yet complete as originally envisaged. Eighteen of the thirty six boards required were installed. In discussion with the CMS management it was decided that this was sufficient for 2016 running. The remaining boards will be available at CERN to potentially expand the system in the future, offering increased flexibility.

The delivery of the hadronic calorimeter (HCAL) back-end electronics, on which the trigger relies, was completed in July.

A review of the project was undertaken by CMS Technical Coordination in March. The report was finally released in September. There are no significant recommendations for the UK area of the project. As expected it highlights a lack of manpower in online software generally. The delayed release of the report has meant that several of the recommendations have been rendered irrelevant and although the UK has acted to strengthen the online software area where we were already contributing significantly, other groups have not, resulting in a delay to the completion of work in online software area.

Overall CMS planning foresaw the upgrade trigger system to be operated in parallel from September, in order to be fully commissioned and ready for physics data taking from the start of the 2016 LHC run. Delays to the online software have affected the planned parallel running, restricting this to a number of test runs rather than a more extensive period of continuous use. It is expected that these runs will be sufficient to demonstrate that the trigger system is functioning correctly.

5.4 Staff on project

The vacant RA position at Bristol will be filled by Luke Krezcko, who brings significant software development expertise to the online software project.

5.5 Expenditure

The expenditure to date continues to be committed to MP7 manufacture and component purchases. Overall spending is well 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	uTCA 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

The status of the milestones which have changed or are late is given below in a little more detail:

- 3.1.7 Stage-1 commissioning was completed in June with the 50ns data and the Stage 1 trigger system has been used for the 25ns data.
- 3.1.8 is ongoing but the major work on our part is complete, and Stage-1 will be used for the heavy ion data taking in November.
- 3.2.3 As reported last time, online software development has proved to be a substantial task requiring more effort than so far available.
- 3.2.4 Algorithms and offline software is complete.
- 3.2.5 The integration task is complete.
- 3.2.6 TDR trigger commissioning is under way, though delayed compared to our objective.

- 3.3.1 Design studies for the future are overdue but well under way. However, they are simply adapted to the global CMS schedule, so not a cause for concern.
- 3.3.2 Simulation and design of the architecture is well underway and ahead of other studies within CMS. It is described as dataflow design but includes technology choices for the track-trigger which are effectively delayed, compared to what we expected, by CMS. These choices will be influenced by the behaviour of demonstrators currently under construction.
- 3.3.3 Hardware design, is probably ahead of our original plan but should certainly continue to evolve for some time.

6 Risk register

The risk register has again been reviewed. Some risks have been re-evaluated (4.1, 4.2, 5.4, 6.2). No new risks have been identified but one (4.2, referring to the FC7 manufacturing difficulties) has been reinstated; in view of the design flaw (4.1, FC7 not meeting specifications) and the difficulties in ensuring manufacturing quality and timeliness this was retired prematurely. However the impact is really on users of the board, not much on our R&D activities, except for effort usage.

Other risks whose likelihood has been raised are 5.4 (MP7 manufacturing difficulties), which is partly correlated with 4.2, and 6.2 (software development problems). Risk 5.4 has been addressed by identifying new PCB manufacturers, one of whom is producing FC7 boards, although we will only be sure of success when the next series of MP7 boards is delivered in a few weeks. Risk 6.2 arises from the major effort to develop new online software for the Phase I trigger upgrade which was significantly underestimated and under-resourced by the trigger project as a whole. The UK has contributed very effectively to addressing this and raising the issue so that CMS sought extra resources elsewhere; it remains a challenging job.

It is worth noting one risk which has not been increased, which is 6.1 (FPGA firmware complexity). The trigger algorithm development took longer than hoped because of poor specifications from some of our collaborators. It is easier to invent algorithms which can be implemented in software than those which are well adapted to firmware, especially subject to timing and resource constraints. However, our firmware experts (particularly A. Rose, G. Iles, D. Newbold, with support from several others in evaluation, emulation and infrastructure support) have done magnificent work to meet the objectives.

7 Finances

Expenditure is reported in the usual financial table.

The travel expenditure reported is, as always, a little behind actual expenditure due to delays in university charges appearing in SBS. This is effectively corrected at RAL for the end of the financial year by estimating the final quarter expenditure, with small corrections then applied in the following quarter. Therefore travel spend seems to be on target to match the estimate for this year with additional LTAs in place.

As reported last time, the Imperial equipment expenditure increased considerably, as expected, but still lags behind the original plan because of the slow completion of MP7 and FC7 production. This now seems to be converging. However, outstanding commitments in CERN are not large since the bulk of the remaining MP7s are due to be delivered to Austrian, Greek and Italian collaborators, who have already committed the necessary funds in CERN.

The CBC3 manufacture is still foreseen for February 2016, and is likely to be part of a shared MPW run. This would increase the number of available chips and probably reduce the cost below the MOSIS submission which was estimated to be ~$250k.

The system of purchasing via CERN and invoicing Imperial College, with the aid of the CMS resource management team and a dedicated account in CERN, continues to work well and several invoices have again been paid by Imperial in the last six months.

Some mistakes in the previous table from last time affecting Imperial (totals to date were taken to be the expenditure for the last financial year) have been corrected.

It may be noted that the Imperial consumable spend (DI) was negative in the last half year. This was due to payment by the COMET project for two FC7 boards, which was attributed by Imperial accounts to consumables, rather than equipment.

Use of engineering design effort in RAL TD is running at about the same rate as last year but extra design effort is now available and is needed to complete the CBC3 submission by early next year, so the total spend for the year will be considerably more than double the current amount. However, this is well within the allocated budget although actually being spent later than our early estimates.

RAL equipment expenditure to date reflects some procurements made via SBS, which will soon be refunded from the Imperial grant.

8 Gantt charts

The Gantt charts for our activities have not changed since the last report.

9 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. The column of revised milestone dates has not required changes since the last meeting, although M3.4 will certainly evolve further.

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	PM30
M2.3.2	PM18	CBC3 produced & test setups ready	PM36
M2.4.1	PM24	Documented early CBC3 test results	PM42
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	PM45
M2.5.2	PM48	CBC4 produced	PM51
M2.5.3	PM54	Documented CBC4 test results	PM57
M2.6.1	PM60	Final production masks prepared	PM60
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

10 Glossary

Following the request at the last meeting, we have 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.

AMC13	A μTCA data concentration and clock distribution card specifically designed for the CMS experiment.
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 widely-used family of algorithms for identifying data corruption.
CTP7	Calorimeter Trigger Processor 7 card, featuring the Xilinx FPGA Virtex-7 chip.
DAQ	Data AcQuisition.
DAQ2	Upgrade to DAQ system during LS1.
DSP	Digital Signal Processor.
DPG	Detector Performance Group.
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 mezzanine cards which interface with FPGAs, as supported by the FC7 and GLIB.
FPGA	Field-Programmable Gate Array.
FSM	Finite State Machine.
GBT	Gigabit Transceiver Project developed at CERN, source of the GBTX and associated chips.
GBTX	Gigabit Transceiver ASIC developed at CERN.
GCT	Global Calorimeter Trigger.
GLIB	μTCA card developed by the CERN microelectronics group for testing and operating with the GBT link.
GMT	Global Muon Trigger.
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.
HGC	High Granularity Calorimeter, the proposed new endcap calorimeter for HL-LHC providing both electromagnetic and hadronic calorimetry.
HI	Heavy Ions, at the LHC refers to collisions between lead ions.
HL-LHC	High Luminosity Large Hadron Collider, the planned upgrade of the current LHC machine around 2023.
HLT	High Level Trigger, a collection of software trigger algorithms.
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 used by system administrators.
JTAG	Joint Test Action Group; test and diagnostic bus standard by IEEE1149.1.
L1A	Level-1 Accept.
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 expansion bus standard.
SerDes	Serialiser/Deserialiser chip.
SFP	Small Form-factor Pluggable standard for optical and other transceivers which connect to a standard-defined socket.
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.
TMT	Time-Multiplexed Trigger, trigger design 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.
uHTR	µTCA HCAL Trigger and Readout card.
VTRX	Versatile Link Transmitter/Receiver, optical transceiver developed by the CERN Versatile Link project.
VTTx	Versatile Link Dual Transmitter, optical transmitter developed by the CERN Versatile Link project which combines two transmission channels in a single SFP module.
XDAQ	Cross DAQ, is a data acquisition software framework
YETS	Year-End Technical Stop, is a relatively brief stop of the LHC, typically less than three months in length, during the winter holidays.