13 May 2015

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

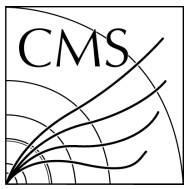

Upgrades of the Tracker and Trigger of the CMS experiment at the CERN LHC

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:

- There has been further evolution of the final CBC design. Manufacturing submission is foreseen in early 2016, and results from total dose radiation and SEU tests have been reported.
- The FC7 general purpose data acquisition board was installed in significant numbers in the upgraded CMS TTC (TCDS) system, where after some months boards began to fail unexpectedly. This was the subject of an extensive evaluation between CERN and Imperial, with an important contribution from Xilinx and the design fault has been traced, and corrected.
- Manufacturing of the MP7 is still under way with two companies. The complement of boards needed for the UK part of the calorimeter trigger is available and installed but there have been delays in the manufacture of further boards.
- The development of algorithms and firmware for the Level-1 trigger is not yet complete but in a very good state.
- The overall project plan for the Level-1 trigger has not been further revised but deliveries of hardware on the US side were delayed and there is little contingency in preparing the new trigger for parallel operation in CMS later this year, and then operation in 2016.
- A Time Multiplexed Trigger architecture for the track-trigger needed for the Phase II upgraded CMS has been further studied, and we are making good progress in developing the ideas for implementation of a demonstrator system over the next year, including algorithms in firmware.

The UK project is on a schedule consistent with CMS overall planning. Progress with the Phase I trigger upgrade has continued to be contentious and remains the subject of greatly increased scrutiny by CMS management.

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, 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. Recently we have been warned that the machine is talking seriously about a higher target of 7.5 x 10^{34} $cm^{-2}s^{-1}$ levelled luminosity, which of course has implications for pileup and radiation levels.

The current phase of the project began on 1 April 2013. There are 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.

As reported by the CERN management at the April RRB, Long Shutdown 1 (LS1) is now considered completed, and a great success, in terms of safety, quality and schedule. All systems operational and LHC commissioning with the beams has started, with experiments ready for beam.

A fairly detailed summary of CMS LS1 activities was given in the May 2015 CERN Courier[1] including explanations of some of the details mentioned in previous reports about tracker cooling and beam pipe installation, as well as other problems which were encountered and overcome.

At the same RRB meeting, a proposal was also presented by the CERN management for a multi-step process of proceeding to approval of the experiments, including assessing the available funding.

[1] http://cerncourier.com/cws/article/cern/60877

- In the first step, the overall scope and cost for the entire upgrade programme for each experiment will be defined, with the possibility to maintain different options depending on technical issues and on funding availability.
- In the second step, detailed technical design reports (TDR) for the subsystems will be reviewed. They are expected to arrive at different times depending on the maturity of the projects, and will be reviewed individually, with the requirement that each fits in the overall approved plan.
- In the third step, the final design and construction readiness of the major detector components will be reviewed, with the requirement that they are compatible with the overall construction and installation plan.
- In the fourth step, as sub-systems are completed, an operations readiness review should be held to evaluate the capability to provide the expected performance and mark the end of the Phase-II upgrade construction project.

This is mostly fairly conventional, although only the first two steps are probably well planned, but the present more delicate issue concerns available funding and how the agencies can endorse the process and allow construction to commence.

Options for detector scoping with a total cost for three benchmark figures in the range 200-275 MCHF with a discussion of the impact on the physics performance have already been requested, and CMS described the status of this activity at the RRB. Then the CERN Management foresaw to propose to the RRB a reference number for the total CORE cost of the Phase-II upgrade of each experiment. Clearly the agencies were not ready to accept this immediately and would like explanations of details; this may take place in the next few months before the next RRB meeting.

CMS also presented a summary of the R&D costs expected to be needed, following detailed appraisals of the detector costs and the R&D. Overall it is estimated that ~22.6 MCHF is needed, representing 8-10% of the construction cost, and from lengthy informal iterations within the experiment it is judged that this funding is either in place or achievable.

2.1 LHC upgrade schedule and planning

As widely publicised, the LHC is now in the start-up phase but the schedule (fig. 2.1) shifted slightly since last reported due to a small number of issues. Beams have been routinely ramped to 6.5 GeV with beam splash events observed, and injection of nominal intensity bunches. The first collisions at injection energy of 450 GeV + 450 GeV have recently taken place. However, minor problems are still being investigated, of which the latest is an obstruction traced to a so-called ULO (Unidentified Lying Object) in one sector, which moves with the beam and may be a fragment of insulator which becomes charged and is subject to electrostatic forces. The machine staff are still exploring if an intervention is required, which may add delays for the sector to be warmed up, or if the beams can be steered around the item. If this is possible there is a few weeks delay with respect to the previous schedule.

At the recent April RRB the expectation from CERN had not changed; 2015 will be a recommissioning year to prepare physics production in 2016, with an integrated luminosity target of 10 fb^{-1}.

From the CMS perspective, there are no major obstacles to physics data taking in 2015, but the expected challenges being overcome to recommission the experiment.

There have been no further changes reported to the CERN plans for the longer term but clearly the financial planning is evolving, and was discussed at the recent April RRB. No doubt the new DG will also influence the plans when she takes up her post.

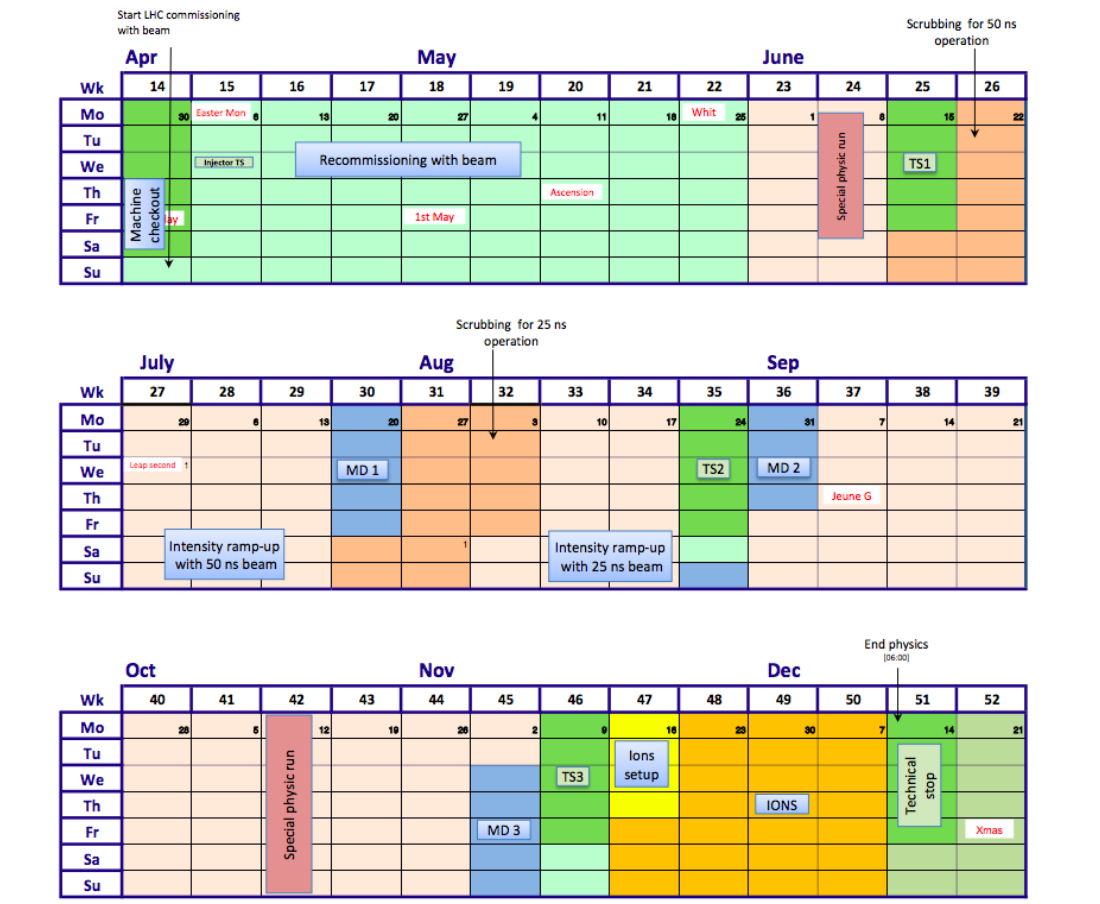

Figure 2.1: The most recent LHC schedule for 2015 as of 4 May.

2.2 CMS planning

The main issue for the UK in the last six months has continued to be progress with the L1 trigger, which is explained later.

The Phase I trigger has continued to be subject to delays on the US side, with hardware for Layer- of the TDR trigger well behind the original plan. The commissioning period has continued to shrink although it should still be feasible to be ready for parallel running of legacy and new triggers in the autumn.

The review of the 2016 calorimeter trigger project foreseen for late January 2015 did not take place as planned, largely because of other CMS activities at Point 5, which limited access to infrastructure in January and therefore made evaluation of the situation difficult. The review was postponed until April, by which time the US group declared delivery of their hardware close to the final schedule. However, the report has not yet been circulated. It will certainly comment on shortages of effort for online software but it is less clear how it will now react to the late deliveries of US hardware, since the emphasis now is on implementation of the system, not on discussing past history.

A CMS Phase II Technical Proposal was prepared for submission to the LHCC last year but full submission was delayed following discussions with the LHCC. This was to allow for difficulties in assessing better with simulations the physics performance of the upgraded detector, and to align it to the proposed decision on the forward calorimeter replacement. It is due for internal CMS approval of the final document in late May and presentation at the June LHCC meeting.

The decision on the calorimeter technology for the endcap region was delayed slightly compared to the schedule reported in the last meeting, to allow more details to be addressed and simulations to be completed, but has now been concluded with endorsement of the management recommendation by

the Collaboration Board in its 8 May meeting. The choice is for the High Granularity Calorimeter (HGC), which is a silicon pad-tungsten calorimeter supplemented with a scintillator-brass backing system (fig 2.2). The performance of the two options under consideration, HGC and Shashlik, seems to be comparable, although possibly with greater potential for the HGC which can exploit the depth sampling and granularity; this needs more simulations to confirm. However, the technical risks of the Shashlik were judged to be more significant, with potential showstoppers, if for example scintillator radiation hardness under LHC conditions were found not to be sufficient.

While the overall cost of the Phase II upgrade does not increase, this option may have implications for other parts of the upgrade because of the resources needed, and because of technology similarities elsewhere, even if it is the lowest priority after Tracker and pixel reconstruction, and barrel ECAL electronics upgrades. The amount of silicon to be procured is significant and the electronics design is challenging; cooling with CO_2 is also a major new project. However at least one manufacturer (Hamamatsu) is confident they can deliver sufficient detectors for the LHC programme in a timely way, and good progress has been made by the tracker in qualifying another large vendor (Infineon) in Europe. New collaborators have been identified to strengthen substantially the electronics development and CERN has a strong interest in global application of the CO_2 cooling technology, which is in any case required to be successful for delivery of the new Tracker.

The UK has an interest in the HGC since one of the main proponents has been T. Virdee from Imperial College and significant extra funding has recently been secured to support the R&D phase: 2 FTE years for studies via an STFC PRD and a much larger grant from the ERC which should provide about 12 FTE years as well as a large materials budget. CMS UK plans a detailed discussion of the future upgrade options and our commitments to it in early July to try to judge how to best deploy the available effort and configure its activities for the long term future.

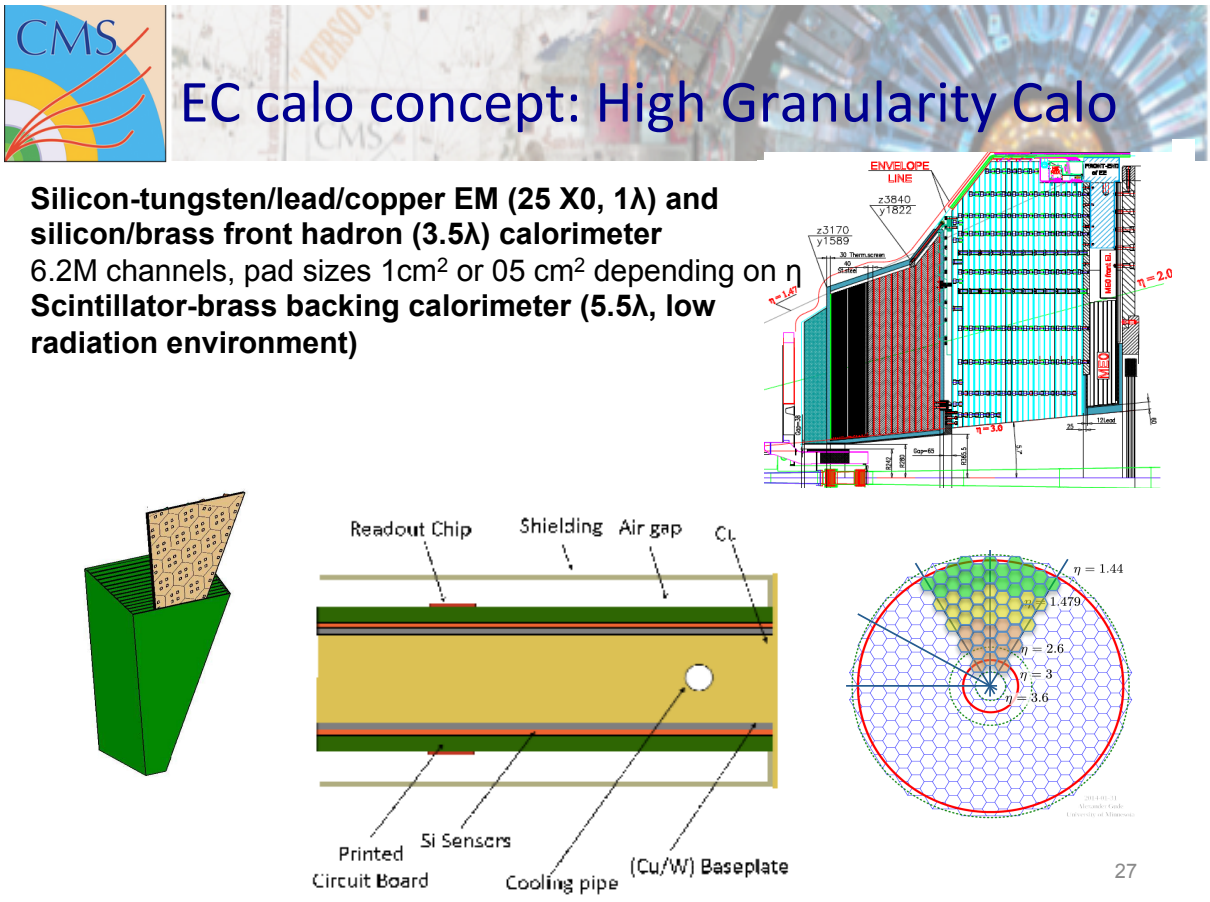

Fig 2.2 High Granularity Calorimeter concept

2.3 UK adaptation to CMS planning

Again, the main issues have been uncertainties in the Phase I trigger project which are discussed in several other places in this report.

2.4 Progress with track-triggering for Phase II

In our previous report, we summarised the possible alternatives for the eventual track-trigger which is needed in the Tracker upgrade in LS3, which will take data from the 2S- and PS-modules and transfer them off-detector for Level-1 track finding and incorporation into the rest of the trigger, which is not yet well defined.

The UK has presented a concept for a track finder for a Phase II trigger, based on a time-multiplexed architecture. A demonstrator system using MP7s is planned for mid-2015, where a slice of the track trigger will be constructed to help gauge the performance, and estimate the requirements for a full system. (fig. 2.3)

A time-multiplexed track trigger based on MP7s would be able to transfer trigger primitives (high p_T hit candidates or stubs) from the full Phase II tracker to a set of processor cards, using a similar number of boards as used for readout of the current silicon strip tracker. These primitives could be processed, ordered and time-multiplexed on a first layer of cards before transmission to a second layer where tracks are identified. Simulations of track finding algorithms across the UK institutes have been focusing on a promising approach which groups tracker primitives, consistent with a real track, together so that track parameters can be easily deduced.

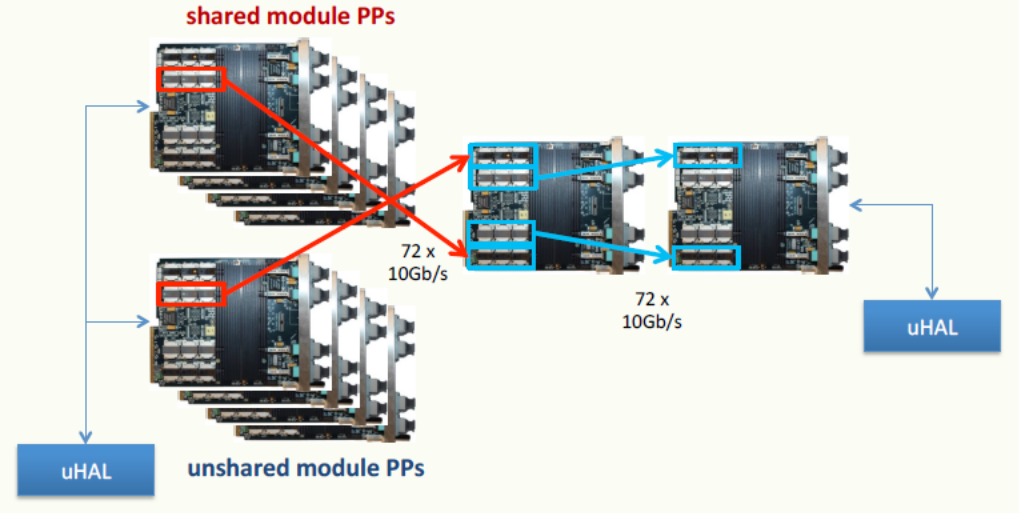

Fig. 2.3. Proposal for a slice of a track trigger demonstrator using MP7s.

The concept, known as the Hough Transform, allows the processor to bin groups of primitives which follow straight lines in a 2D projection of the tracker geometry and is typically used in image manipulation and enhancement (fig 2.4). Simulations show that by pipelining tracker primitive data into the processor, one can fill a Hough Transform array and identify high transverse momentum tracks with >95% efficiency (for muons), even at high pileup.

Effort has been devoted to try to understand the operating margin of this concept, including Monte-Carlo simulations to test the robustness to varying pileup conditions and missing layers, estimating the number of output candidates, fakes and efficiency for differing event topologies, and small scale firmware designs to measure the impact of the algorithm on FPGA resources. Additionally the results of these studies are allowing us to refine the approach in preparation for the demonstrator system. The firmware to encode the Hough Transform array in an FPGA and the logic required to pass primitives through the array is well under way and is currently being integrated with the infrastructure firmware that is provided for use with MP7. Software previously used to test the L1 calorimeter trigger demonstrator is being reused to download and play events through the track finder

MP7, and monitor the results at the output. This demonstrator will be expanded over the next six months to attempt to emulate a more realistic slice of a future track trigger system.

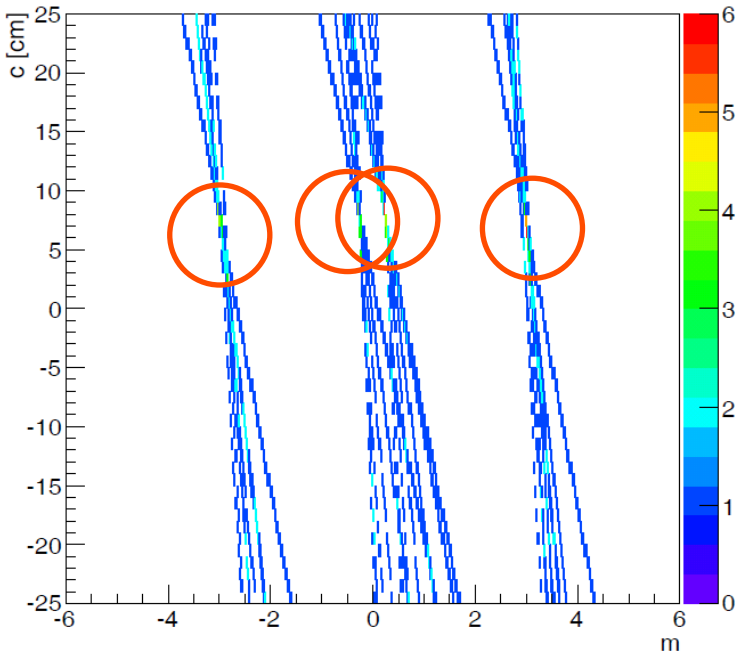

Fig 2.4. High transverse momentum track trigger primitives from a 4 muon particle gun event represented in Hough space. The red circles indicate where the primitives line up on a straight line projection in normal space and a track is found.

These initial studies on the firmware implementation have been undertaken in collaboration with the CERN group. We have also received enquiries about collaboration from the Vienna and Karlsruhe CMS groups and are discussing future plans.

3. Work Package 1: Management

A reminder of the project management is included below. G. Hall is stepping down as UK CMS PI at the end of September 2015, but will remain responsible for this project. Another Imperial College member of the administrative support team has recently left; Carol Barlow, at Easter; there are departmental plans for a replacement in the coming months. Paula Consiglio was appointed to a support position and has taken on the duties previously carried out by Piera Brambilla.

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.
	D Newbold	Bristol	UK firmware and software coordinator. Trigger Institution Board chair.

4. Work Package 2: Outer Tracker Readout

4.1 Objectives

As usual we repeat the objectives of Work Package 2 as stated in the proposal:

- 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.

The 2S-PT stacked tracking module concept is the established solution for providing tracking information to the CMS Level-1 trigger decision in the outer tracker regions at the HL-LHC. Developed in the UK within the framework of this upgrade project, the CBC2 130 nm CMOS bump-bondable chip is our existing prototype which has demonstrated the stacked tracking principle. The CBC2 development has proven to be a vital step in understanding and advancing the 2S module design. The system requirements have evolved and are now mature, allowing us to develop the final version of the chip

4.2 Progress

The design of the CBC3 is now fully under way. The main architectural changes from the design of the CBC2 are:

- capability to transmit up to 3 stub addresses off-chip, per 25 ns bunch crossing period, with ½ strip resolution of the seed cluster (stub defined by a cluster in the lower seed layer correlating with a cluster occurring within a window in the upper layer)
- each stub also has associated bend information indicating the location of the cluster within the window in the upper layer
- 512 cell pipeline depth, allowing a Level-1 latency of up to 12.8 μsec
- Level 1 trigger rate capability of 1 MHz

A functional block diagram of the CBC3 is shown in figure 4.1. To achieve the data transmission capability of stubs information and triggered data the chip will drive 6 differential output pairs at 320 Mbps (48 bits per 25 ns bunch crossing). In addition there will be a single 320 Mbps differential input pair carrying the trigger, test pulse trigger, fast synchronous reset and Level-1 counter reset signals (the Level-1 counter is also a new feature which simply keeps track of the number of triggers since the last reset, the value being included in the output data packet). The chip will receive the LHC 40 MHz clock and a second synchronous 320 MHz clock for the data reception and transmission blocks. The sampling time of the discriminated input signals will be tuned using an on-chip DLL locked to the 40 MHz clock. Data must be passed between the 40 MHz and 320 MHz clock domains.

A new stub gathering logic block is required to generate the stub addresses and bend information and pass the data (8 bits address + 5 bits bend) to the 320 MHz region of the chip. Up to 3 stubs worth of data can be passed per 25 ns bunch crossing (average stub occupancies per chip are simulated to be $<< 1$).

The cluster width discrimination block will now pass clusters up to (but not exceeding) 4 strips wide (it was 3 for the CBC2) and will produce an additional output for 2 and 4 strip clusters to achieve the ½ strip resolution requirement.

The front end amplifier and bias generator circuitry is being reviewed and adjusted to fix minor deficiencies found in the CBC2, to adjust for a slightly shorter pulse shape, and to be able to run at an internal power rail (regulated on chip) of 1.0 V. The biasing circuitry is being adjusted to be able to run more current in the input transistor if necessary, because there is a possibility that the 2S module

strip lengths may increase if sensor production could move to 8-inch wafers. The front end will be optimised for the final sensor polarity choice of n-on-p technology.

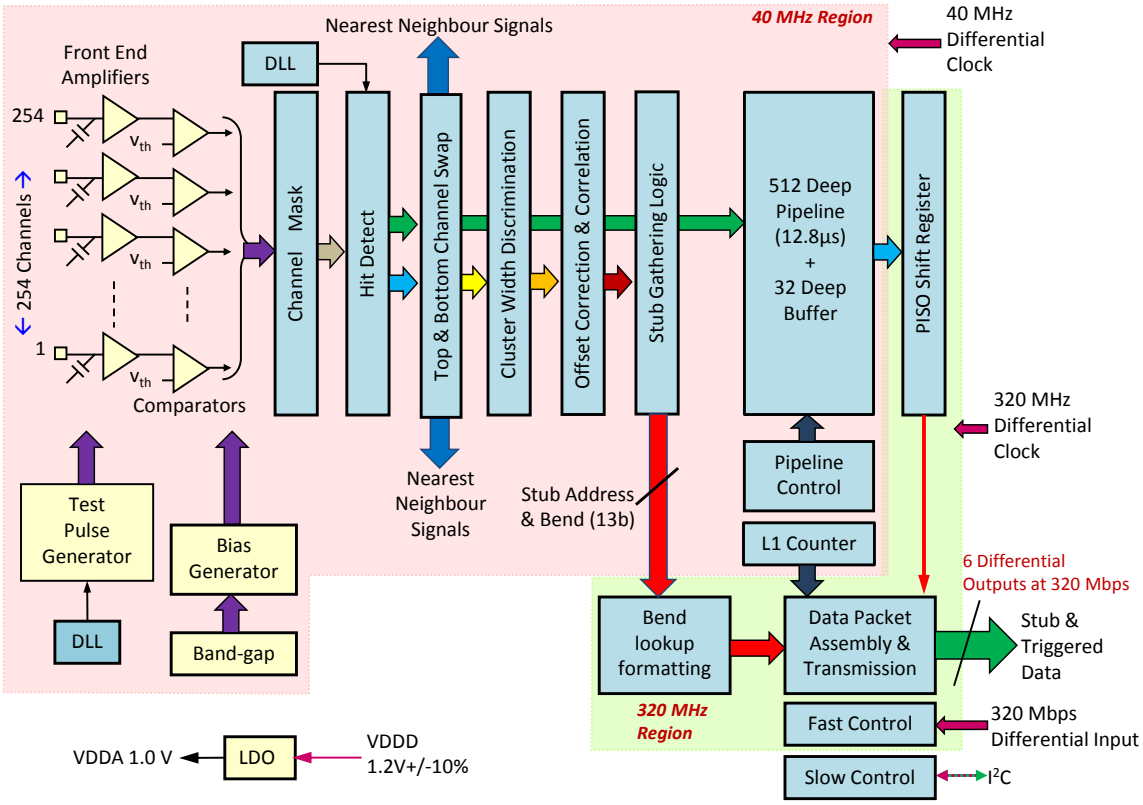

Figure 4.1. CBC3 functional block diagram

As reported last time, the characterization of the CBC2 is complete, including ionizing radiation and SEU testing, and minor issues identified in those tests are being addressed in the CBC3 design. A further ionizing irradiation of the CBC2 up to 180 Mrad has been performed by CMS collaborators in Karlsruhe. No significant effects are observed up to 120 Mrad, with full functionality, but some degradation in noise performance observed at 180 Mrad. The expected dose in the outer tracker region at HL-LHC is 10 Mrad.

A full specification document for the CBC3 has now been produced[2], which results from regular discussions and meetings with our CMS electronics system team collaborators, and circulated for comments in mid-April. The design work is progressing, with modifications to the front end amplifier, hit detect logic, stub-gathering logic, and I2C registers all well-advanced. Review meetings are scheduled at key stages to monitor progress and to check that the project is on track, for example an Intermediate Design Review was held for the CBC3 design in March to review the circuits designed so far, including the Stub Gathering Logic, the Hit Detect Circuit and the new I2C registers.

Tests continue with modules based on CBC2 readout. A beam test with mini-modules built with irradiated sensors is planned for June this year, and a beam test with full-size modules based on the 8CBC2flex hybrid is planned for November. We will be participating in these tests.

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. We envisaged it originally as a

[2] http://www.hep.ph.ic.ac.uk/~dmray/CBC_documentation/CBC3_Technical_Spec_V1p03.doc

dual purpose board for the Phase I upgraded pixel DAQ (which was not supported by our PPRP award) and as a prototype for the Phase II DAQ, both to take data from 2S-modules (and potentially PS-modules) in beam and laboratory tests and position us for a possible role in the Phase II tracker construction project, based on our established expertise and prototype developments.

The fact that the FC7 could be envisaged also as a more general purpose component, and still be exploited by the CMS pixel project with very little support from the UK, was attractive to others; hence our collaboration with CERN to maximise the general purpose, as opposed to project-specific, features of the board. However, some unexpected issues materialised in the last six months, which we believe to now be successfully resolved but which have important implications for many projects in the future.

Over 80 prototype R0 (revision 0, i.e. the very first version) boards have been produced for use in CMS for use in the Trigger, Control and Distribution System (TCDS), which is now fully installed at P5. In December 2014, after several months of operation, two boards started to exhibit odd behaviour whereby the FPGA failed to drive a small number of I/O pins. While a workaround was implemented, further boards appeared to demonstrate logical and I/O failures in the FPGAs in 2015. An increase of the FPGA core current was found to accompany the logical failures. In February, a Failure Analysis procedure was initiated with Xilinx to help determine the cause of the issues, which eventually reported back that one pin responsible for powering the on-chip ADC had been biased above absolute specification, leading to electrical overstress of the silicon. The design error was spotted and corrected on all working R0 prototypes in existence by early April through the use of a single wire connection.

It is worth noting that this is a typical error in the design of a board, that a small number of connections may be logically incorrect. Such occurrences are usually identified at the first design iteration since they give rise to some odd feature, such as an inverted signal or missing functionality. In this case, the misconnection led to overvoltage of a single pin (3.3V cf. 1.8V specified limit), but with no visible repercussions and hence did not show up during testing as explained below.

The Xilinx 7-series FPGA is clearly surprisingly resilient to a significant overvoltage condition, as demonstrated by the relatively low failure rate of less than 1 board per week following 5 months of continuous operation at P5 with no apparent issue. In the past, using earlier generation Xilinx components, a design error such as this would have led to immediate and fatal FPGA failure, and this would have been identified during prototype testing. This implies that a lengthy period of qualification should be considered for FPGA based board designs in future, including provision for accelerated burn in tests, and tests under realistic running conditions with a large enough quantity of boards.

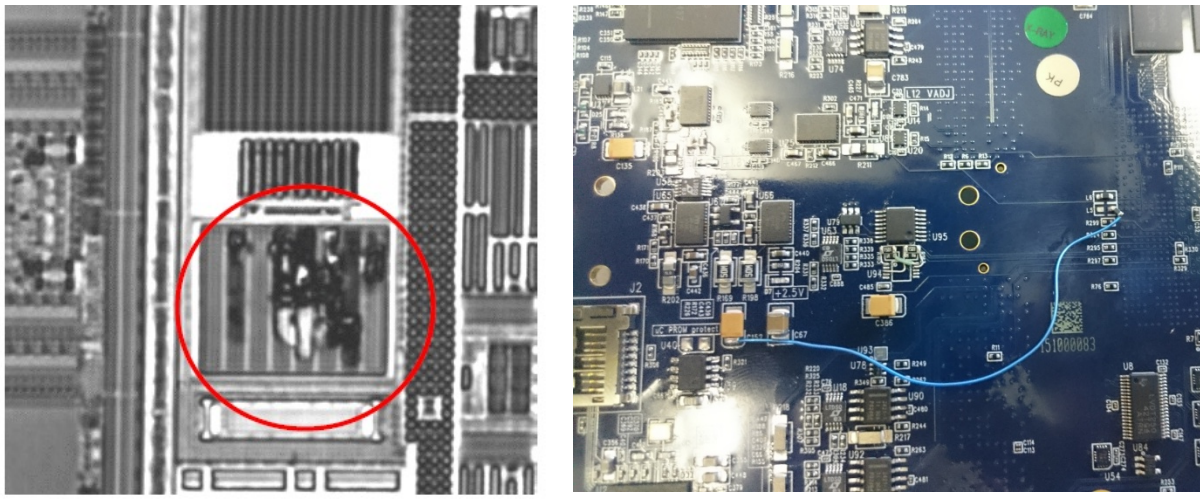

Fig 4.2 Left: Die backside laser imaging revealed diffusion damage in the configuration area, indicating an electrical overstress event. Right: Wire fix to the FC7-R0/R1, correcting the ADC powering design error.

The remaining fixed boards at P5 (approximately 50) are currently working without issue, although it is unclear if the silicon has been significantly weakened by the overvoltage condition or if a definite reduction in the lifetime of these boards should be expected. A set of 10 (revision 1) FC7-R1s (out of a batch of 30) developed for prototyping the CMS Phase I Pixel FED (and also carrying the same design error) is available for substitution into the TCDS if further failures are observed. All R1 boards have now also had the wire fix applied to them and, significantly, the majority of these boards were not powered up beforehand, thereby limiting possible damage to the FPGA.

The R1 differs from the R0 mainly in the use of a new PCB material (Panasonic Megtron6), which has been recommended as a replacement for the Nelco N4000-SI used on the FC7-R0 and MP7 which can suffer from delamination during assembly heating under non-ideal handling conditions. During testing in 2015, the R1 has been verified to have a high speed performance that matches that of the R0 although further measurements are planned.

However, another possible issue which has been identified is from sectioning of test coupons produced alongside the PCBs indicate that the PCB manufacturer (contracted by the assembly company Hapro, Norway) has not manufactured the boards to the electrical class specified, reflecting our experience with the MP7-XE. The impact of the failure to meet the PCB specifications is yet to be understood and investigations are ongoing.

A new version of the FC7 (R2 - revision 2) has been prepared for manufacture and 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. A large scale production order of the R2 is being prepared by CERN and first boards are expected mid-2015.

A few important issues raised by this experience are:

- Because of the urgency of the TCDS project (proposed and managed by CERN) a board on its first iteration was manufactured in quite large numbers and installed into CMS with practically no time for evaluation or even burn-in,
- The project itself had not planned any serious qualification, and efforts from the design team (though extremely effective once triggered) in collaboration with the users, with notable help from Xilinx once it was convinced to do so, were essential to resolve the fault, which took some time,
- Even the most careful scrutiny of the design by experienced engineers did not spot the misconnection and it is questionable whether this would have been found by any review. Nevertheless more emphasis on design reviews, with outsiders, is already being instigated for our projects.
- Provision of common hardware to a wide variety of projects is probably highly desirable. However, the support implications are rarely thought through, especially on the user side, where it is generally assumed that the cost of such boards is the bare manufacturing cost, and yet the boards come fully qualified, possibly even with significant support for implementation in the form of firmware, software and advice with practical problems in use.

Part of the relevance of this experience is that we have already seen significant numbers of requests for provision of FC7s or MP7s often even with little or no contact with the designers. Until both boards are in full use we expect to continue to be cautious in commitments of this kind.

4.3 Deliverables

The WBS for WP2 is included below. The system specification definition task 2.1.1 is now complete, as far as factors affecting the CBC3 design are concerned, and the CBC3 specification document is now available. We are in the middle of the CBC3 design task 2.3.1. The intention is to submit the chip for manufacture through MOSIS, which offers fewer chips, but with costs approximately half those of a full wafer run. The most convenient and achievable submission date for a run with MOSIS is February 2016, which is later than planned in the WBS but which should still lead to chips in hand by the middle 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 SS-PT module readout**
	2.1.1 specification definition	04/13	03/14	12	regular meetings with CMS collaborators to define overall system specification and interfaces
2.2	**CBC2 test**	04/13	03/15	24	**CBC2 is final deliverable of the UK upgrade R&D**
	2.2.1 CBC2 ongoing testing	04/13	03/14	12	complete the detailed studies of the CBC2 chip, including irradiation and SEU tests
	2.2.2 CBC2 SS-PT module prototype studies	04/13	03/15	24	a programme of SS-PT module studies, in collaboration with CMS, including test beam
2.3	**CBC3**	04/14	03/16	24	**CBC3 is specified for the final system**
	2.3.1 CBC3 design	04/14	09/15	18	design period
	2.3.2 CBC3 production	09/15	03/16	6	production period
	2.3.3 test setup preparation	09/15	03/16	6	wafer and chip test setup preparation
2.4	**CBC3 test**	03/16	03/18	24	**CBC3 chip and module testing**
	2.4.1 early tests	03/16	09/16	6	chip verification tests to prior to module tests
	2.4.2 ongoing testing	09/16	03/17	6	complete characterization, including irradiation and SEU tests
	2.4.3 CBC3 SS-PT module studies	09/16	03/18	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

Davide Braga, the principal RAL engineer who was also working on a part-time PhD at Imperial College, actually left RAL at the start of this year to take up a post at Fermilab, having completed the design and layout of the Stub Finding Logic.

Michele Key-Charriere started work to replace him in September working on the design of most of the new digital sections. She encountered an unexpected issue with the Design Kit, which caused some delay, but with that issue resolved, she has been able to make up the lost time. Mark Prydderch has been developing the Hit Detect circuit, the new I2C registers and writing the CBC3 specification document.

Michelle and Mark Prydderch will continue to work together on the project for the foreseeable future and in the autumn should be joined by Przemek Mroszczyk, who joined the RAL group in November last year. The balance of effort should shift to Przemek in the later stages of the design.

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 next major manufacturing cost is likely 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

Commissioning of the system is underway in the Underground Service Cavern next to CMS. The system is essentially complete, with a few small exceptions. Since the last review we have switched to a final crate design, installed 12 MP7s (see Fig. 5.1), added remote debug capability via the addition of a JTAG Switch Module (see Fig. 5.2) and switched to a final network configuration.

The two aspects that remain outstanding are the redundant power system (see Fig 5.2), that does not operate quite as expected, and the last part of the fibre installation between the processing nodes and the de-multiplexing card. We are in discussion with the manufacturers about the former and the latter is waiting a technical decision from CMS.

The Layer 1 – Layer 2 full mesh optical patch panel was validated on schedule in December 2014 and since then pattern tests have been completed to validate new algorithms. The algorithms have been driven from both internal buffers on the MP7 and from 1/9 of the Layer 1 system (the full Layer 1 system is not yet complete). The test was extended to include the final stage of the chain, the upgraded Global Trigger, so that data was passed through the full system.

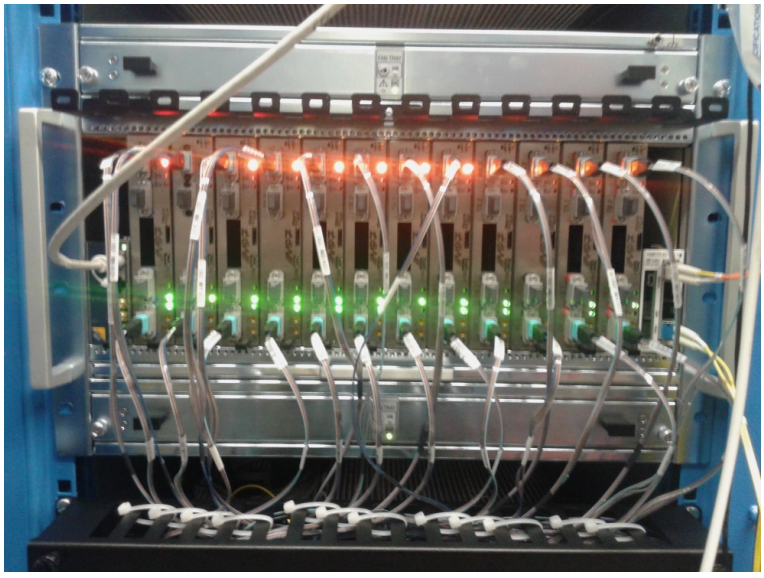

Fig. 5.1 The main processing nodes connected to the L1-L2 optical patch panel. Communication is via a standard Ethernet card on the left. LHC clock distribution, fast control & feedback and DAQ is via the module on the right.

A significant challenge over the last six months has been that the MP7 procurement has not always proceeded as planned, although all cards are in place for the TDR trigger, including those needed for Layer-1 should they be required. The initial MP7 production of 32 cards was split between two companies to mitigate risk and because the two leading bids were within 3% of each other. The

production at Hapro, Norway proceeded smoothly and is finished. All cards have passed through a detailed set of power, interconnect and functional tests.

In both cases, the manufacture was staged into deliveries of 2 (A), 4 (B) and 10 (C) boards with approval from our side at each stage, following acceptance tests, to ensure quality and minimise risk.

At the last review production at the second supplier, Exception PCB (UK), was not proceeding as hoped despite a review with the Managing Director on 1 August 2014. This led to a substantial improvement in communication with several manufacturing changes; however, problems persisted. A formal warning that CERN would start breach of contract proceedings unless certain deliveries were met was issued on 17 December 2014. The company met the deadline for delivery of 6 February 2015 for 4 assembled cards (part B) and 23 January 2015 for the outstanding bare PCBs. The latter was achieved with a significant increase in yield from 30% to 70% due to manufacturing improvements. The subcontractor, Jaltek, has struggled to assemble the final 10 cards. At present 14 of the 16 cards have been delivered with last 2 expected shortly.

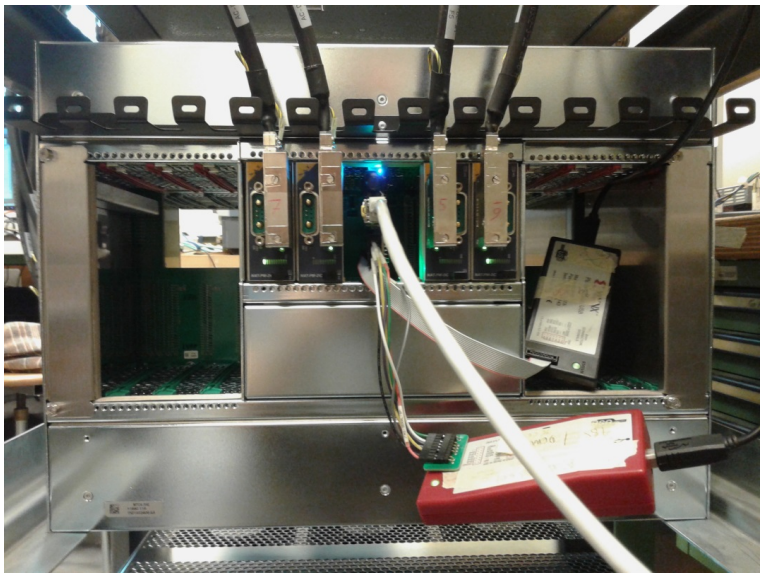

Fig. 5.2. The rear of the crate with space for 4 power modules in a 2+2 configuration (i.e. two primary supplies and two backup supplies). The low level debug/upgrade communication interface is shown in the middle.

A supplementary production of 8 cards were ordered from Hapro using spare PCBs from the first production batch. A failure during production delayed delivery; however, all 8 cards have now been delivered and successfully tested. The final production of 32 cards to meet the needs of our collaborators is underway, with the first 6 pre-series cards at CERN and under test.

Work on testing and refining algorithms for the calorimeter trigger, in conjunction with the LLR and CERN groups is ongoing.

The Global Calorimeter Trigger (GCT), the UK responsibility in the legacy trigger, has been re-commissioned during cosmic ray muon runs and the first beam splashes from the LHC. Several aspects of the infrastructure at CMS had been changed (computers, experiment-wide timing, clock and control system etc.) and also the installation of the Stage-1 trigger upgrade for 2015 necessitated further work to change the source of the inputs to the GCT. This has all been completed successfully and people on the project based at CERN have been trained so they are able to support the operation of the GCT in 2015.

Developments for the future

Ideas have also been developing for a successor to the MP7. A brief summary of work under way is given here, and more details will be reported in future as the design evolves towards prototype manufacture.

The three key limitations of the design of the MP7 are the 100W limit on the power supply of the µTCA form-factor, difficulties related to manufacturing to the µTCA specification, and the high cost, dominated by the FPGA. A fourth comment about the MP7 by some potential users is the "lack" of RAM available on the board.

Instead of µTCA, a new board has been designed in a PCIe form-factor, alleviating constraints on power and mechanical tolerances, making manufacturing simpler. This has the additional advantage that developers do not need additional expensive infrastructure, as the card fits in a standard desktop PC or server, and we can utilise off-the-shelf components designed for the mass-computing market.

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 cheaper and has higher capacity than the QDR-technology used on the MP7, at the cost of slightly lower bandwidth.

The first generation of card is based on the Kintex Ultrascale FPGA, providing 48+48 optical links at 16Gbps. For the requirements of particle physics, 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 at cost-effective for the development cycle of the card.

A 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

For installation in an experiment, it may be desirable to have the board in a crate-oriented form-factor, such as ATCA. The conversion of the schematics from the PCIe form-factor to an ATCA form-factor has been tested and demonstrated to take a matter of hours. For ease of manufacturing and development, however, the smaller, more ubiquitous PCIe form-factor is preferred.

5.3 Overview of CMS plans

The Layer 1 electronics (the responsibility of University of Wisconsin) are not yet complete as originally envisaged. Eighteen of the thirty six boards required are installed. A discussion is under way about whether all 36 are needed to implement the trigger for 2016.

The delivery of the hadronic calorimeter (HCAL) back-end electronics is proceeding well. Two batches of cards have been received and installed at CERN and delivery of the full number of cards is expected by the end of May. Another 2-3 weeks is likely to be necessary for firmware work before testing with the trigger system can begin. This is a U.S. responsibility (primarily University of Minnesota).

Overall CMS planning requires the trigger system to be operating in parallel from September, in order to be fully commissioned and ready for physics data taking from the start of the 2016 LHC run.

A review of the project was undertaken by CMS Technical Coordination in March. The report is not yet available. We do not anticipate significant recommendations for the UK area of the project but it is expected to highlight a lack of manpower in online software generally.

Despite the delays to hardware from outside the UK the schedule has not been significantly changed by the Project Management since reported last time.

5.4 Staff on project

The remaining project studentship at Imperial has been filled, with the student beginning in October.

5.5 Expenditure

Most of the expenditure to date continues to be committed to MP7 manufacture and component purchases. Overall spending is well within the budget foreseen at this stage.

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 is given below in a little more detail:

- 3.1.6 Stage-1 integration can be said to be complete. There are some small outstanding items for DAQ but they are minor details.
- 3.1.7 Stage-1 commissioning is well under way but will not finish until the 50ns data is delivered, which is expected in June.
- 3.2.2 Procurement and testing of the UK Stage-2 boards needed in the trigger is complete.
- 3.2.3 As reported last time, online software development has proved to be a substantial task requiring more effort than so far available and is expected to be the subject of comments in feedback from the review when it is published.

- 3.2.4 Algorithms and offline software is complete, except for DQM, which is a small amount of effort.
- 3.2.5 The UK part of the integration task is complete; what remains to be completed is the full scope of Layer 1.
- 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.

6 Risk register

The risk register has again been reviewed. No new risks have been identified or any more retired, but a couple of risks which had been identified have actually arisen and have been, or are still being, addressed. However, neither of them has a direct impact on UK deliverables, although they could have an impact on other projects relying on our developments.

The two issues concern manufacture and performance of the digital processing boards, the FC7 and MP7. They were described earlier and do not affect our deliverables, either for the Phase I trigger completion or for our future R&D. However, as remarked in the report on the FC7, there are useful lessons to be drawn, both for us and for CMS.

The latest plan still foresees completion of the TDR trigger by August 2015, so it can be commissioned for operation in 2016 and the CMS management has made clear the importance of this objective.

7 Finances

Expenditure is reported in the usual financial table.

The travel expenditure is a little lower than planned but LTAs are now in place from Imperial since February, and at RAL since last year, which means the rate of spend will increase next year as foreseen.

The Imperial equipment expenditure has increased since the last report because of the CERN invoicing, as expected, but still lags behind the plan because of the slow completion of MP7 and FC7 production, which delays the invoices. However, this will gradually be corrected in future, since funds are now transferred to place orders from CERN, rather than be underwritten at the time of order from our team account. Even more significant compared to the plan is the much later submission of the next CBC version, which is now foreseen in February 2016, and using MOSIS. Although this leads to a smaller number of chips being available, it reduces the risk – and a number of significant changes are bing made for the CBC3 – and costs are lower. The estimated cost is ~$250k, which is significant for only a modest number of chips, but is acceptable for evaluation.

Use of engineering design effort in RAL TD is running slower than our original plan but is acceptable from the design perspective, and preserves some margin for the future which, given the uncertainty about the construction phase, both timing and cost sharing, is desirable.

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 is working well and several invoices have been paid by Imperial in the last six months.

RAL equipment expenditure this year was slightly higher than expected because some procurements were made via SBS. The cost table does not reflect this (it shows £10k expenditure) since it will repaid soon from Imperial; the amount outstanding incurred in the last financial year prior to the new arrangements to be invoiced by RAL is £20,379.

8 Gantt charts

The most recent Gantt chart for the trigger project is again included (v11a), in several versions. It is a detailed update to the last plan (v11) included at the last OSC, retaining the same project milestones.

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 added to the table for revised milestone dates has not been changed since the last meeting.

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

The milestone dates for WP2 have not changed since the last meeting and in any case are adapted to the overall planning for the future and tracker (and CMS) HL-LHC upgrade uncertainties described earlier. Therefore adjustments are not the result of delays in the UK activities or technical difficulties but reflect additional work or evolution of the CMS specifications following from physics and detector studies. M2.2.2 has been flagged as complete on the basis of modules constructed using small sensors but will continue as new 8-CBC flex-based hybrids and full size sensors become available.

Steady progress has continued with WP3, with no further delays in the overall milestones but where the US deliverables have consistently been late. The commissioning and debugging period needed for the TDR trigger for 2016 has certainly been compressed.

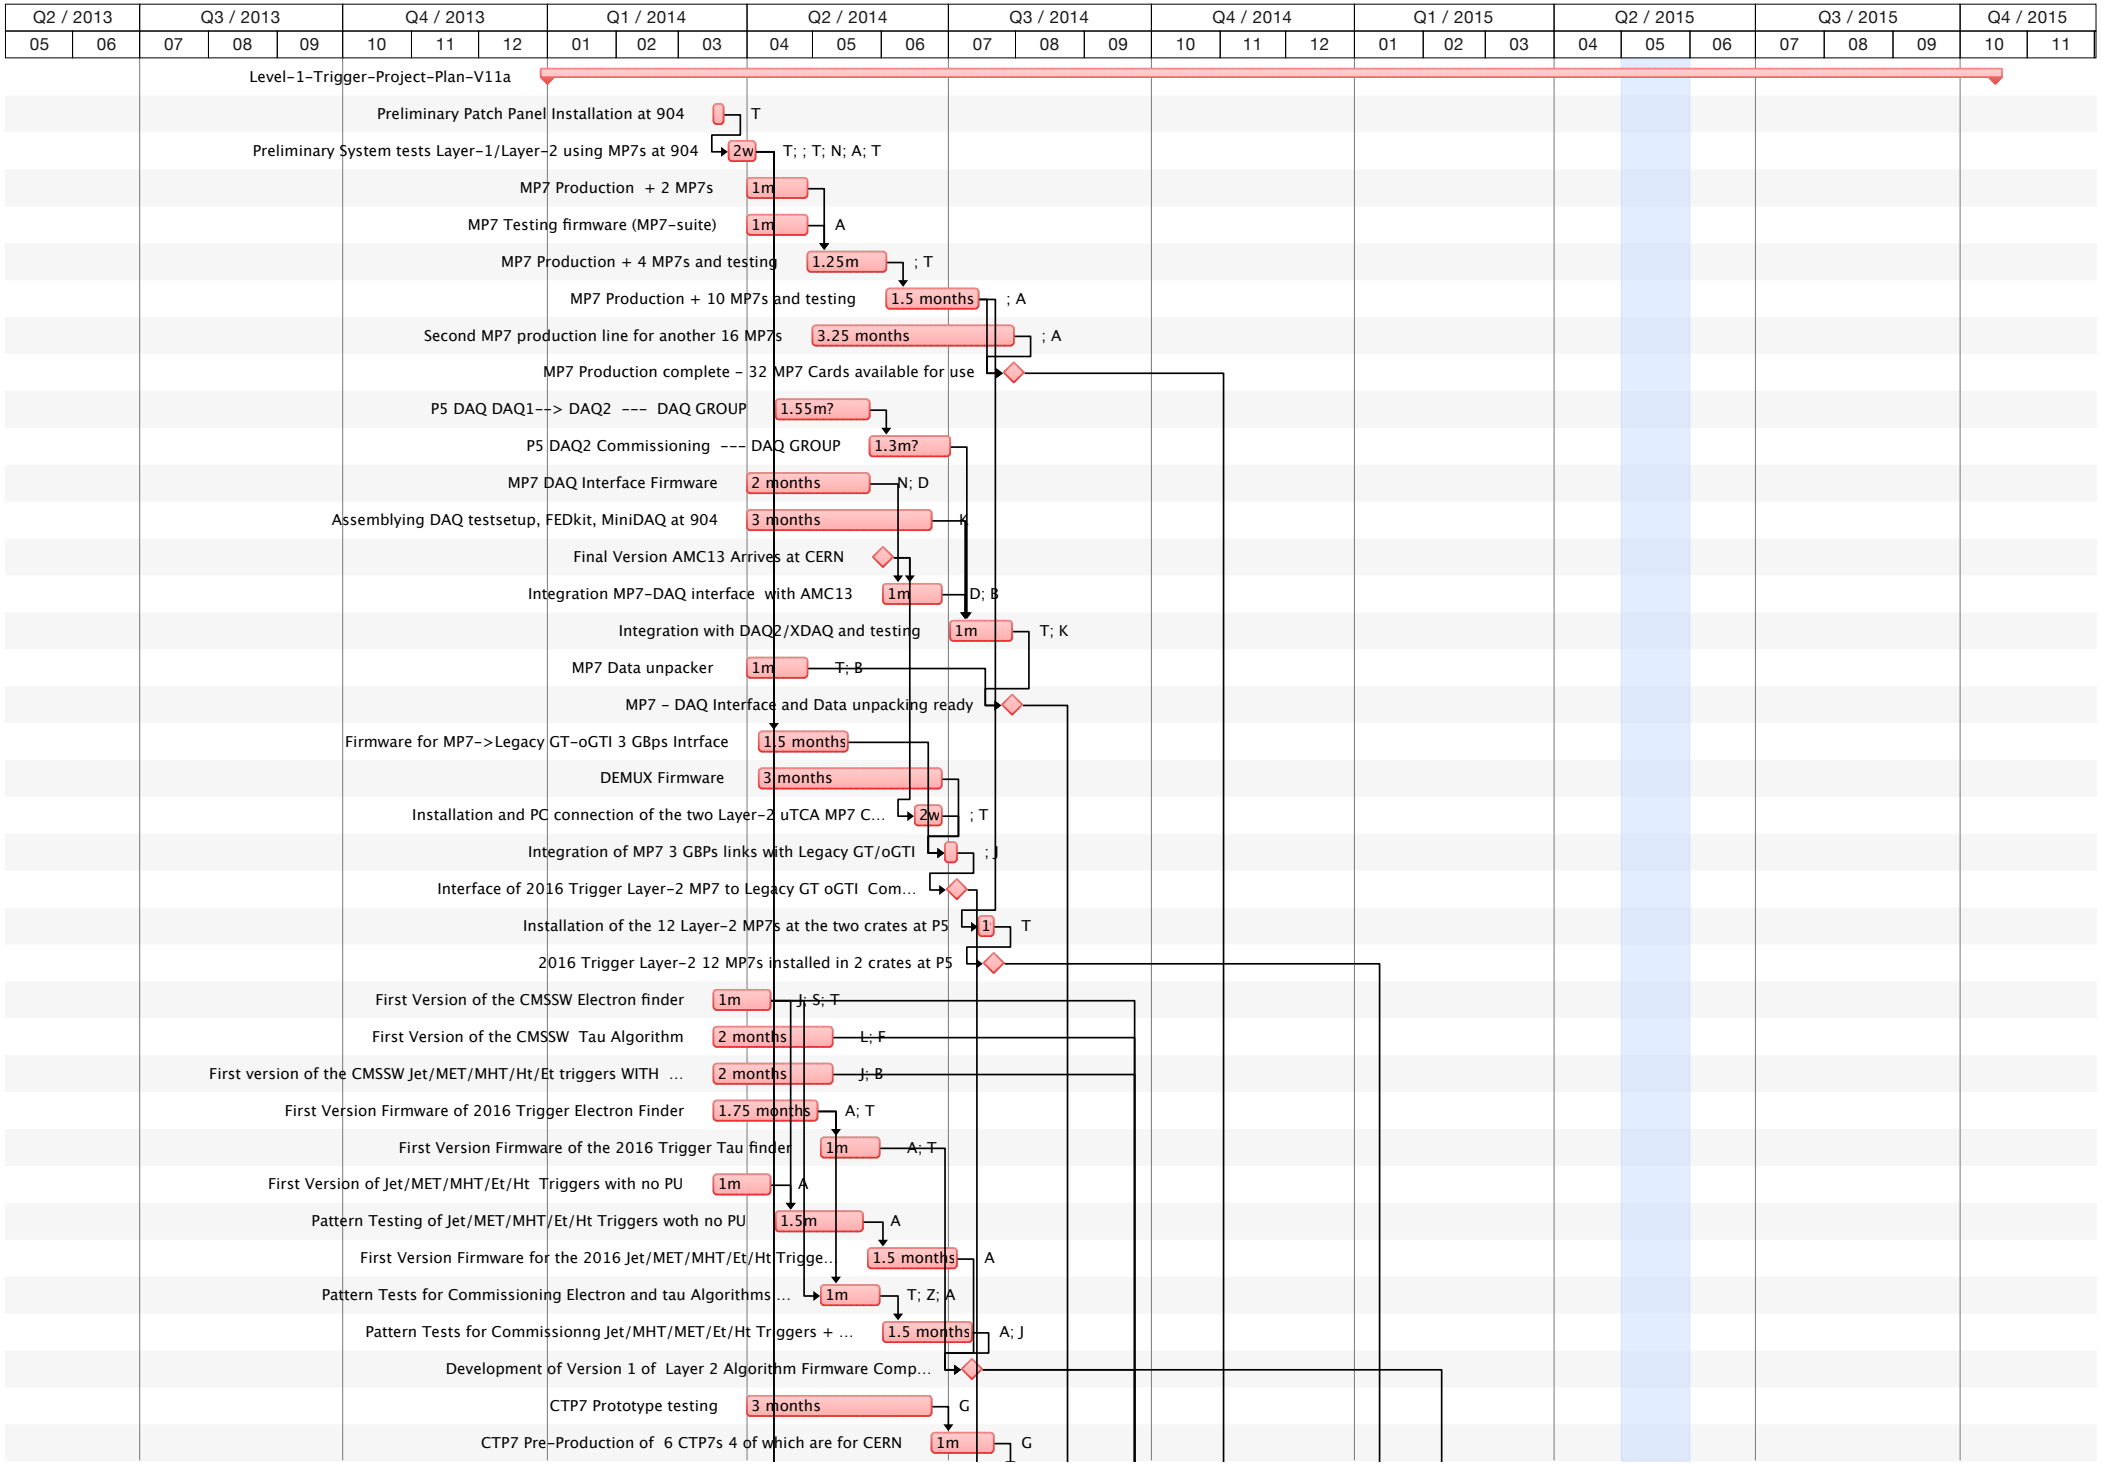
Q2 / 2013
Q3 / 2013
Q4 / 2013
Q1 / 2014
Q2 / 2014
Q3 / 2014
Q4 / 2014
Q1 / 2015
Q2 / 2015
Q3 / 2015
Q4 / 2015
05 06 07 08 09 10 11 12 01 02 03 04 05 06 07 08 09 10 11 12 01 02 03 04 05 06 07 08 09 10 11
Level-1-Trigger-Project-Plan-V11a
Preliminary Patch Panel Installation at 904
T
Preliminary System tests Layer-1/Layer-2 using MP7s at 904
2w
T; ; T; N; A; T
MP7 Production + 2 MP7s
1m
MP7 Testing firmware (MP7-suite)
1m
A
MP7 Production + 4 MP7s and testing
1.25m
; T
MP7 Production + 10 MP7s and testing
1.5 months
; A
Second MP7 production line for another 16 MP7s
3.25 months
; A
MP7 Production complete - 32 MP7 Cards available for use
P5 DAQ DAQ1--> DAQ2 --- DAQ GROUP
1.55m?
P5 DAQ2 Commissioning --- DAQ GROUP
1.3m?
MP7 DAQ Interface Firmware
2 months
N; D
Assemblying DAQ testsetup, FEDkit, MiniDAQ at 904
3 months
K
Final Version AMC13 Arrives at CERN
Integration MP7-DAQ interface with AMC13
1m
D; B
Integration with DAQ2/XDAQ and testing
1m
T; K
MP7 Data unpacker
1m
T; B
MP7 - DAQ Interface and Data unpacking ready
Firmware for MP7->Legacy GT-oGTI 3 GBps Intrface
1.5 months
DEMUX Firmware
3 months
Installation and PC connection of the two Layer-2 uTCA MP7 C...
2w
; T
Integration of MP7 3 GBPs links with Legacy GT/oGTI
; J
Interface of 2016 Trigger Layer-2 MP7 to Legacy GT oGTI Com...
Installation of the 12 Layer-2 MP7s at the two crates at P5
1
T
2016 Trigger Layer-2 12 MP7s installed in 2 crates at P5
First Version of the CMSSW Electron finder
1m
J; S; T
First Version of the CMSSW Tau Algorithm
2 months
L; F
First version of the CMSSW Jet/MET/MHT/Ht/Et triggers WITH ...
2 months
J; B
First Version Firmware of 2016 Trigger Electron Finder
1.75 months
A; T
First Version Firmware of the 2016 Trigger Tau finder
1m
A; T
First Version of Jet/MET/MHT/Et/Ht Triggers with no PU
1m
A
Pattern Testing of Jet/MET/MHT/Et/Ht Triggers woth no PU
1.5m
A
First Version Firmware for the 2016 Jet/MET/MHT/Et/Ht Trigge..
1.5 months
A
Pattern Tests for Commissioning Electron and tau Algorithms ...
1m
T; Z; A
Pattern Tests for Commissionng Jet/MHT/MET/Et/Ht Triggers + ...
1.5 months
A; J
Development of Version 1 of Layer 2 Algorithm Firmware Comp...
CTP7 Prototype testing
3 months
G
CTP7 Pre-Production of 6 CTP7s 4 of which are for CERN
1m
G

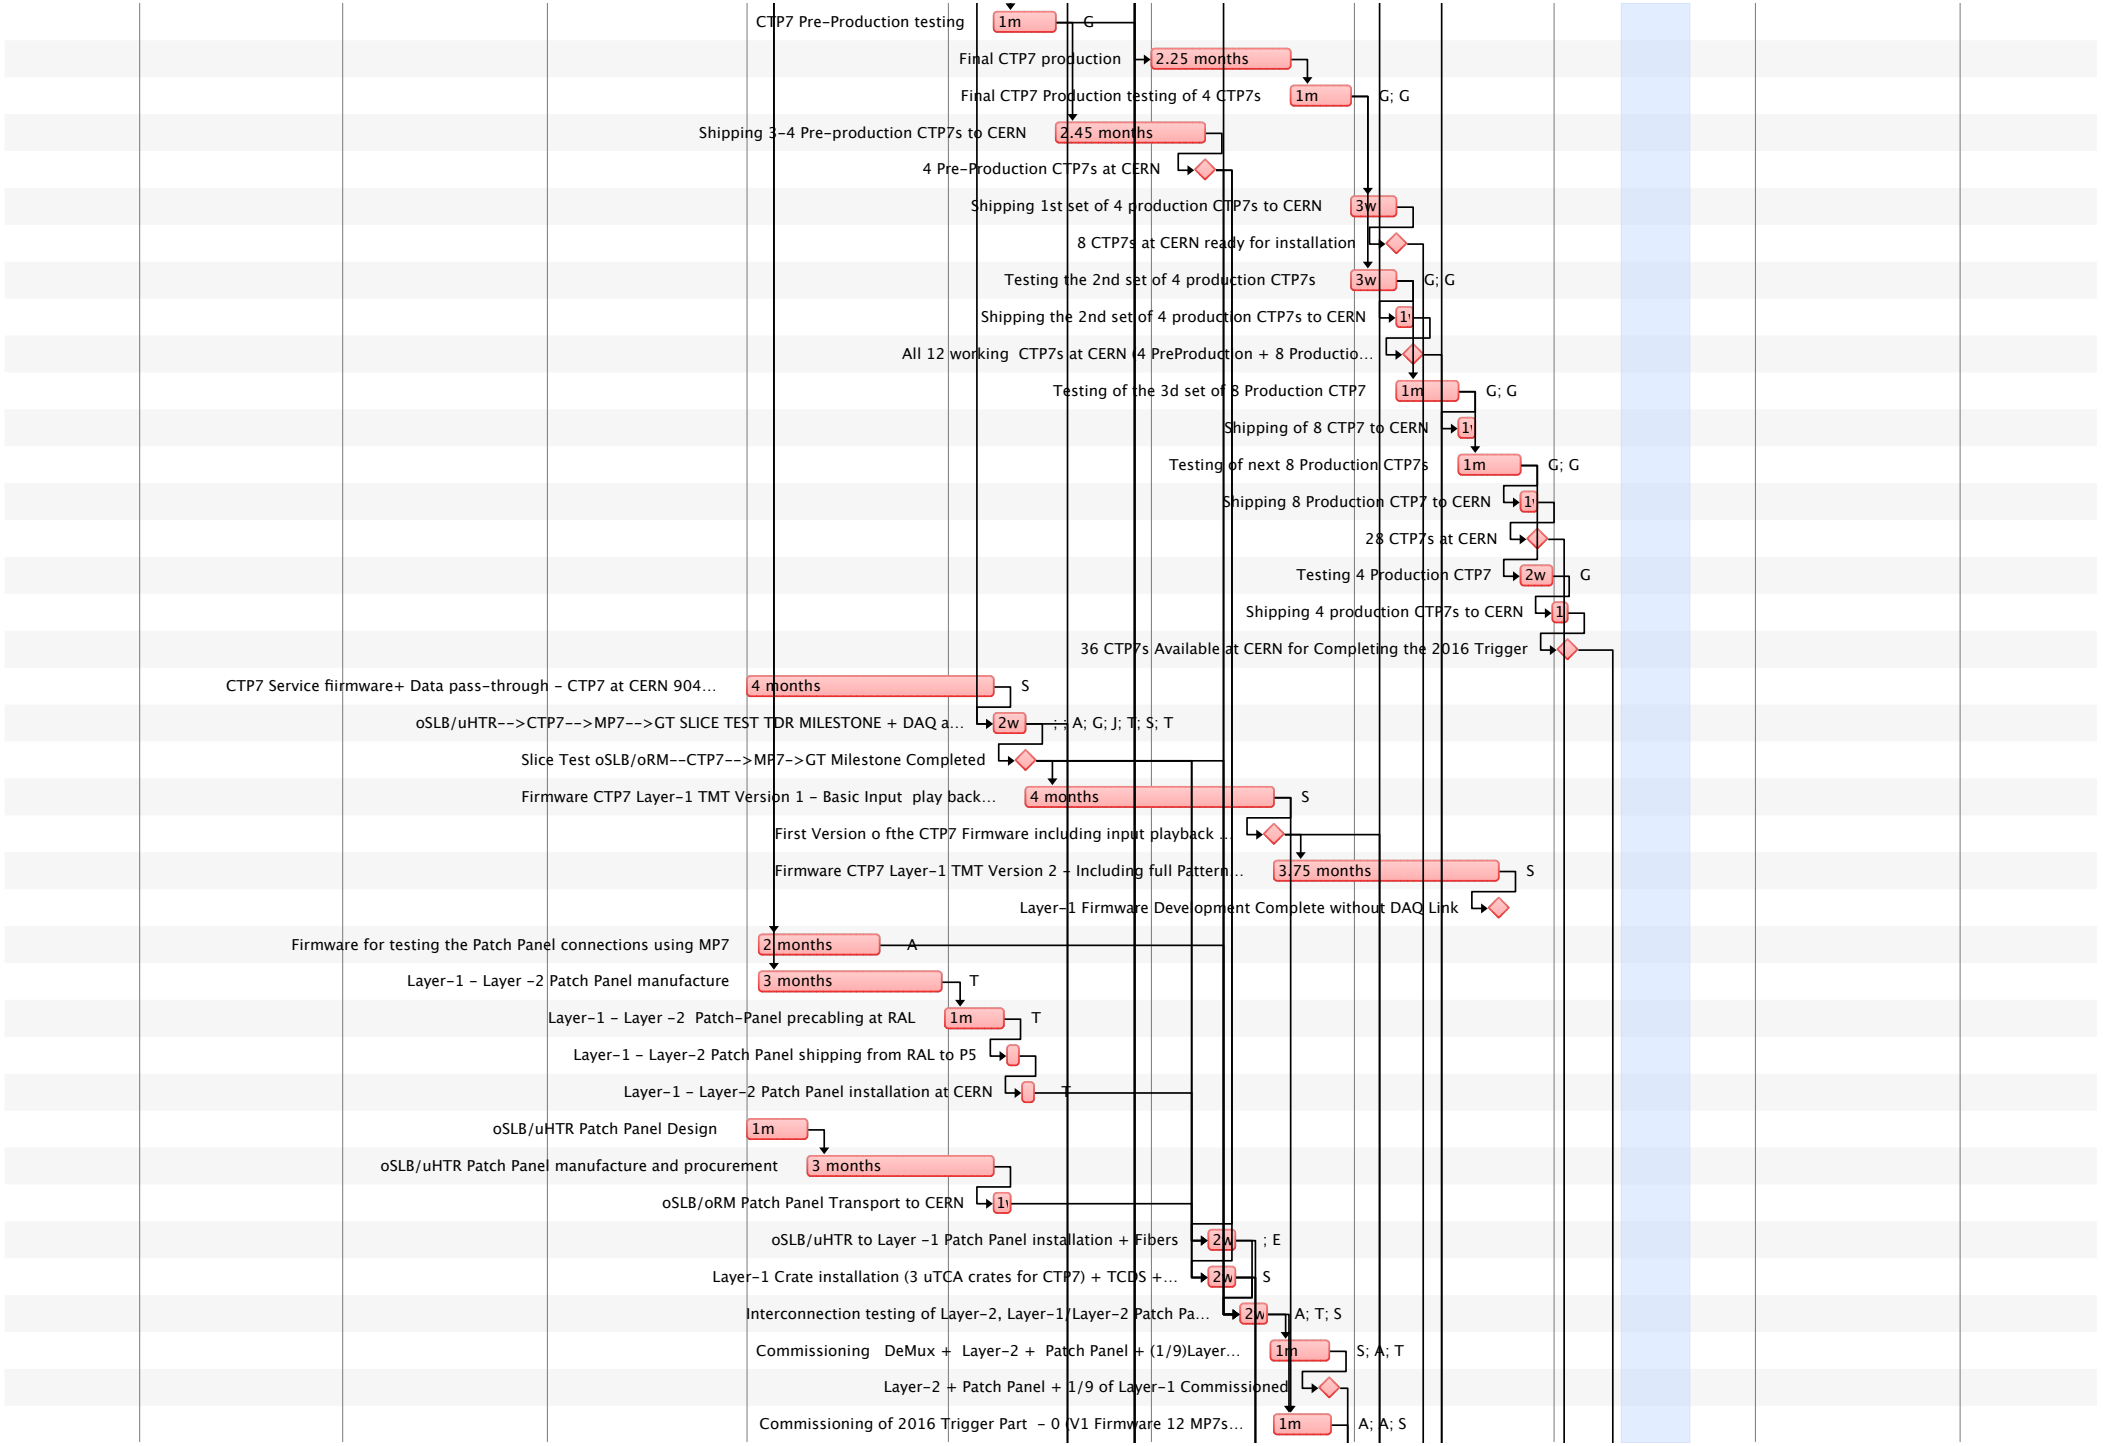
CTP7 Pre-Production testing
1m
G
Final CTP7 production
2.25 months
Final CTP7 Production testing of 4 CTP7s
1m
G; G
Shipping 3-4 Pre-production CTP7s to CERN
2.45 months
4 Pre-Production CTP7s at CERN
Shipping 1st set of 4 production CTP7s to CERN
3w
8 CTP7s at CERN ready for installation
Testing the 2nd set of 4 production CTP7s
3w
G; G
Shipping the 2nd set of 4 production CTP7s to CERN
All 12 working CTP7s at CERN (4 PreProduction + 8 Productio...
Testing of the 3d set of 8 Production CTP7
1m
G; G
Shipping of 8 CTP7 to CERN
Testing of next 8 Production CTP7s
1m
G; G
Shipping 8 Production CTP7 to CERN
28 CTP7s at CERN
Testing 4 Production CTP7
2w
G
Shipping 4 production CTP7s to CERN
36 CTP7s Available at CERN for Completing the 2016 Trigger
CTP7 Service fiirmware+ Data pass-through - CTP7 at CERN 904...
4 months
S
oSLB/uHTR-->CTP7-->MP7-->GT SLICE TEST TDR MILESTONE + DAQ a...
2w
; A; G; J; T; S; T
Slice Test oSLB/oRM--CTP7-->MP7->GT Milestone Completed
Firmware CTP7 Layer-1 TMT Version 1 - Basic Input play back...
4 months
S
First Version o fthe CTP7 Firmware including input playback ...
Firmware CTP7 Layer-1 TMT Version 2 + Including full Pattern ...
3.75 months
S
Layer-1 Firmware Development Complete without DAQ Link
Firmware for testing the Patch Panel connections using MP7
2 months
A
Layer-1 - Layer -2 Patch Panel manufacture
3 months
T
Layer-1 - Layer -2 Patch-Panel precabling at RAL
1m
T
Layer-1 - Layer-2 Patch Panel shipping from RAL to P5
Layer-1 - Layer-2 Patch Panel installation at CERN
T
oSLB/uHTR Patch Panel Design
1m
oSLB/uHTR Patch Panel manufacture and procurement
3 months
oSLB/oRM Patch Panel Transport to CERN
oSLB/uHTR to Layer -1 Patch Panel installation + Fibers
2w
; E
Layer-1 Crate installation (3 uTCA crates for CTP7) + TCDS +...
2w
S
Interconnection testing of Layer-2, Layer-1/Layer-2 Patch Pa...
2w
A; T; S
Commissioning DeMux + Layer-2 + Patch Panel + (1/9)Layer...
1m
S; A; T
Layer-2 + Patch Panel + 1/9 of Layer-1 Commissioned
Commissioning of 2016 Trigger Part - 0 (V1 Firmware 12 MP7s...
1m
A; A; S


Integration of the 2016 Trigger DemMux at P5 with uGT
1m
D; T; B; K; M
CERN CLOSED FOR CHRISTMAS AND NEW YEAR
2.2w
Connecting Layer-1/PP/Layer-2/Demux to uGT
2w
Layer-2 + Patch Panel + 1/9 of Layer-1 Connected to uGT
Commissioning of the oSLB links to the 4 CTP7s at P5
2.5 months
; M
Commissioning the 36 HF uHTR fibers to 4 Layer-1 CTP7s
2.5 months
M; E
HF 36 uHTR and oSLB System Commissioned
Commisioning - 12 MP7 + 8 CTP7 2016 trigger Patterns
1m
A; A; S
System Test test oSLB+HF uHTR-->8xCTP7-->12xMP7-->Demux-->uG...
1m
Layer-1 CTP7 DAQ Link Development work - Common with Stage-...
1m
S; K; ; D
Layer-1 + Layer-2 DAQ Interface test with DAQ
1m
; D; K; S
2016 Trigger with HF+oSLB Ready
Commissioning HCAL B/E uHTR system with trigger 12/36 Layer-...
1.25m
Commissioning of the HCAL B/E uHTR system with trigger 28/36...
1m
Commissioning of the HCAL BE/uHTR systerm with trigger 36/36...
2 months
Barrel and EndCap uHTR systems connected to the Layer-1 CTP7...
Online Monitoring Software for the 2016 Trigger
10 months
2016 Trigger final Emulators
5 months
Final design of Algorithms and Data Format defined
Final Firmware for the 2016 Trigger + Validation + Upgrades
7 months
2016 Trigger Ready for Parallel Data Taking
ORSc Firmware - V1 for the APRIL2014 RCT->ORSc->MP7 test
4.15 months
S; M
ORSc Testing software for the April-2014 RCT->ORSc->MP7 test
4.15 months
K; C
RCT-ORSc-MP7 -- Integration Tests at 904 - 10 GBps Links
3w
S; C; K;
ORSCs Integrated with RCT and MP7 and ready for production
ORSc Firmware for Integration with Legacy GCT
1m
; M
OSCs integration tests with Legacy GCT - 2GBps links
1m
; M
ORSCs Integrated with Legacy GCT
1st Global Run April 7.4 - 11.4
ORSc Production
1m
G
ORSc Testing at Madison
1m
G
Easter Week 18.4 - 21.4
Power off at P5
3.2w?
Commissioning Software testing interconnectivity at P5
1m
G
Installation of uTCA Crate at P5 with MP7
2w
ORSCs Testing complete
ORSc transport US-->CERN
ORSC Installation at P5 including all fibers to GCT+MP7 + P...
1m
C; K
Comissioning ECAL/HCAL/ORSC/MP7 system interconnectivity at ...
1m
C


ORSCs Installation and testing with MP7 at P5 Complete
ORSc firmware for running Stage-1
2 months
G
ORSC Firmware for running the Legacy trigger
2 months
M
2nd Global Run 3.7 - 10.7.2014
1.
Comm the Legacy Running via ORSC 2GBps - Move 108 ECL - No ...
1m
; T; M
Legacy System Can Trigger with ORSc input
Comissioning of the Stage-1 running firmware and integration...
1m
G; C; K;
Stage-1 Comm. In Parallel with Legacy
Stage-1 MP7 DAQ link devel. and integration - Common with St...
1m
; D
Commissioning of the Stage-1 MP7 DAQ LINK
Stage-1 MP7 DAQ Link Commissioned
ORSc firmware for testing+Validation+Monitoring the Legacy s...
2 months
M
ORSc Software for testing+Validation+Monitoring the Legacy S...
2 months
T; M
Legacy System Fully Integrated with ORSc
Electron and Tau Algorithm Firmware development for MP7
6 months
F
Electron and Tau Algorithm Pre comissioning at 904
2 months
F
Electron and Tau Algorithm Commisioning at P5
1m
Electron and Tau Trigger Validation vs Emulator
3 months
F; C
Monitoring Software for Stage -1 Electron and Tau Triggers
4 months
G
Electron/Tau Triggers Ready and Validated
Jet+MET+MHT Algorithm Firmware Development
6 months
F
JET+MET+MHT Algorithm precommissining
3 months
F
JET+MET+MHT Firmware Commissioning at P5
1m
F
JET+MET+MHT Trigger Validation vs Emulator
3 months
F; M
Monitoring Software for Stage -1 JET/MET/MHT (Hardonic Trigg...
4 months
STAGE-1 MET/MHT/JETS AND ALL OTHER TRIGGERS DONE
Heavy Ion Emulator Development
2.8 months
C; A; M
First Version of Heavy Ion Emulators finished
Preparing e/gamma VHDL and gaining exeprince with the HI MP7
2 months
C;
Validate with patterns and MC the E/gammar Algorithm Firmwa...
1m
C
Electron/Photon Trigger Validated on the bench
Instalation of HI Patch Panel, Links and MP7 at P5
2w
M
Test the HI DAQ and TCDS connections
1m
M
Running the HI Algorithms parasiticaly and validating with p...
4 months
Jet Algorithm + PU VHDL Firmware Development for HI
2 months
M
Validate with patterns the Jets and PU algorithm
2 months
M
Jet and PU alogorithm Validated on the Bench
Centrality Algorithm VHDL Development
2 months
C
Validation of the Centrality Algorithm Firmware
2 months
C

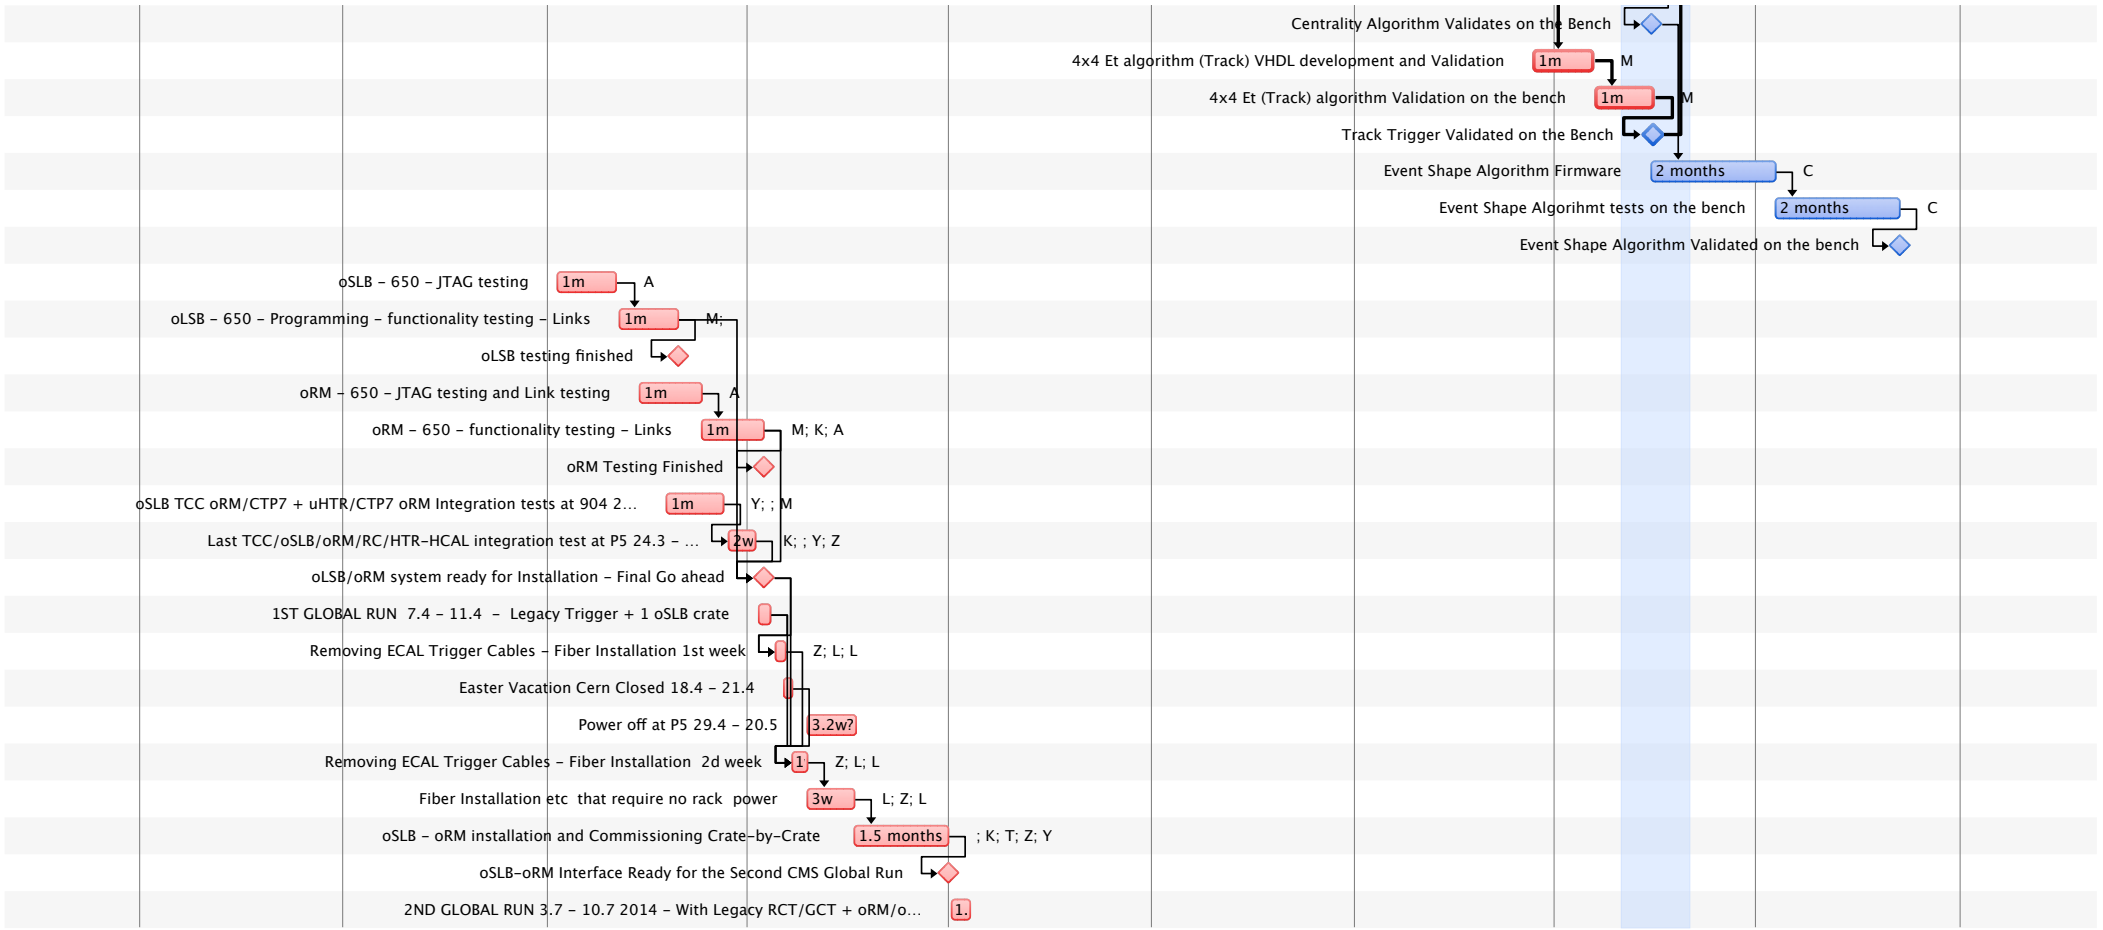
Centrality Algorithm Validates on the Bench
4x4 Et algorithm (Track) VHDL development and Validation
1m
M
4x4 Et (Track) algorithm Validation on the bench
1m
M
Track Trigger Validated on the Bench
Event Shape Algorithm Firmware
2 months
C
Event Shape Algorihmt tests on the bench
2 months
C
Event Shape Algorithm Validated on the bench
oSLB – 650 – JTAG testing
1m
A
oLSB – 650 – Programming – functionality testing – Links
1m
M;
oLSB testing finished
oRM – 650 – JTAG testing and Link testing
1m
A
oRM – 650 – functionality testing – Links
1m
M; K; A
oRM Testing Finished
oSLB TCC oRM/CTP7 + uHTR/CTP7 oRM Integration tests at 904 2…
1m
Y; ; M
Last TCC/oSLB/oRM/RC/HTR-HCAL integration test at P5 24.3 – …
2w
K; ; Y; Z
oLSB/oRM system ready for Installation – Final Go ahead
1ST GLOBAL RUN 7.4 – 11.4 – Legacy Trigger + 1 oSLB crate
Removing ECAL Trigger Cables – Fiber Installation 1st week
Z; L; L
Easter Vacation Cern Closed 18.4 – 21.4
Power off at P5 29.4 – 20.5
3.2w?
Removing ECAL Trigger Cables – Fiber Installation 2d week
1
Z; L; L
Fiber Installation etc that require no rack power
3w
L; Z; L
oSLB – oRM installation and Commissioning Crate-by-Crate
1.5 months
; K; T; Z; Y
oSLB-oRM Interface Ready for the Second CMS Global Run
2ND GLOBAL RUN 3.7 – 10.7 2014 – With Legacy RCT/GCT + oRM/o…
1.