The High-Luminosity Upgrade of the CMS Detector

University of Bristol: P. Baesso, J. Brooke, E. Clement, D. Cussans, H. Flächer, J. Goldstein, H. Heath, L. Kreczko, D. Newbold, S. Paramesvaran, B. Penning, T. Sakuma, S. Seif El Nasr

Brunel University: J. Cole, P. Hobson, A. Khan, I. Reid

Imperial College London: G. Auzinger, J. Borg, R. Bainbridge, O. Büchmuller, A. Bundock, D. Colling, P. Dauncey, G. Davies, S. Greenwood, G. Hall, G. Iles, A-M. Magnan, S. Malik, L. Mastrolorenzo, J. Nash, V. Palladino, M. Pesaresi, M. Raymond, A. Rose, C. Seez, A. Tapper, K. Uchida, T. Virdee, N. Wardle, S. Zenz

Rutherford Appleton Laboratory: D. Cockerill, T. Durkin, K. Harder, S. Harper, K. Manolopoulos, E. Olaiya, D. Petyt, D. Sankey, C. Shepherd-Themistocleous, A. Thea, I. Tomalin, T. Williams

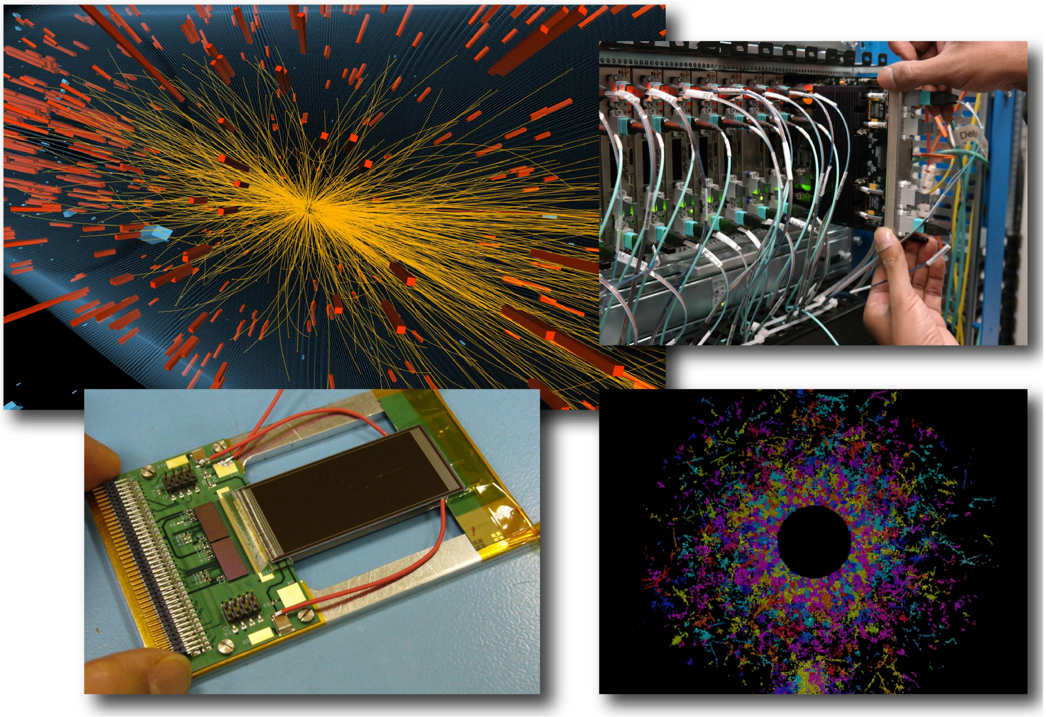

Summary

The CMS detector will be upgraded for high-luminosity operation of the CERN LHC from 2026. We propose a technically ambitious five-year project from 2019 that will deliver state-of-the-art tracking, calorimetry and trigger capabilities for CMS. These improvements will permit an order of magnitude increase in the recorded data set, allowing measurement of Higgs couplings and self-couplings, and greatly enhanced sensitivity to physics beyond the Standard Model. The capital and new staff cost to STFC is estimated at £9.6M, plus a working allowance of £300k.

1 Introduction

The Large Hadron Collider (LHC) at CERN will be the world's highest energy particle accelerator for the foreseeable future, and the only facility capable of investigating some of the highest priority topics in fundamental physics. The UK has made a substantial long-term commitment to the LHC, and to the design, construction and operation of the CMS detector. The discovery of the Higgs boson in 2012 [1] was a triumphant start to the scientific programme, and showed that the design choices and scientific strategy adopted by CMS were valid. The CMS collaboration in the UK (CMSUK) has led the exploitation of LHC data in key areas: in searches for supersymmetry, dark matter and exotic phenomena, as well as in Higgs physics. LHC operation is now expected to continue until at least 2035, with ever-increasing performance required from both the machine and detectors. The high luminosity upgrades to CMS described in this proposal will allow scientific return from the LHC to continue in the long term. The 2013 European Strategy for Particle Physics [2] stated that "Europe's top priority should be the exploitation of the full potential of the LHC, including the high-luminosity upgrade of the machine and detectors".

From 2026, the High Luminosity LHC (HL-LHC) will provide a levelled instantaneous luminosity of $5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, resulting in data sets of up to $3\,\mathrm{ab}^{-1}$, an order of magnitude larger than the CMS detector was designed for. The vastly increased sensitivity to rare processes this offers, underlies each of our long-term scientific objectives. However, it also requires major upgrades to detectors. Key components of CMS will require renewal due to the radiation damage sustained during LHC operation and the need to remain efficient under much harsher conditions, and electronic systems that will be around a quarter of a century old will be replaced to make use of up-to-date technology. As foreseen in our 2012 proposal, the CMS silicon tracker will be replaced by a more granular and radiation-hard system, also capable of providing data to the trigger system. The CMS trigger will require a major upgrade to cope with increased luminosity, and to take advantage of tracking data. Since 2012, it has also become clear that a major upgrade to CMS calorimetry is required; we will replace the current endcap calorimeters with a much higher granularity detector, and replace the electronics of the barrel ECAL to allow full use of its precision in the trigger system. The CMS upgrade plans have been documented in a Technical Proposal [3, 4], approved by the CERN LHCC in late 2015.

The UK groups have an exceptional track record in the construction and exploitation of CMS, and have provided sustained innovation and leadership in the upgrade R&D programme, and in its justification in terms of science goals. In particular, the UK: proposed the concept of stacked sensors, around which the new silicon tracker design is based; led R&D on the endcap calorimetry which led to the adoption of the high-granularity concept by CMS in 2014; led the definition of the ECAL upgrade; and has already successfully delivered a complete Phase-1 replacement of the Level-1 trigger based on a novel time-multiplexed architecture that will form the basis of future developments. In all cases, these achievements build on the UK's acknowledged expertise in advanced electronics, and on the close and coherent connection in CMSUK between detector work and physics exploitation. We are excellently placed to make key contributions to the upgrade construction programme.

We propose here a five-year construction programme that will deliver upgraded tracking, calorimetry and trigger systems for CMS by 2024. This project will begin at the conclusion of

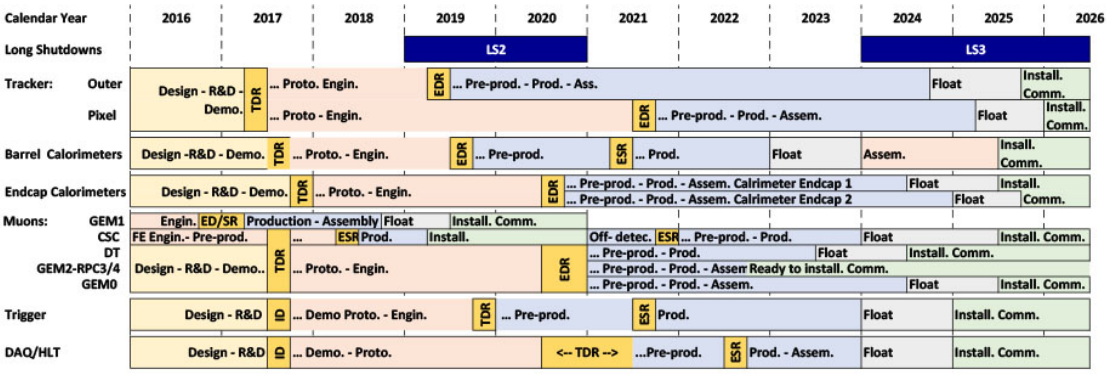

Figure 1: CMS top-level planning for upgrade construction

the currently-funded R&D project, in April 2019, and will be followed by a two-year installation and commissioning period running through the LHC Long Shutdown 3 (LS3), culminating in the start of HL-LHC exploitation in April 2026. We request STFC funding for the capital cost of constructing detectors, the cost of staff to undertake the project, and the costs of related travel and consumables. In all areas, the capital expenditure (i.e. CMS CORE contribution) is strongly linked to UK deliverables, and underpins UK leadership. It is our aim to spend as large a fraction of capital as possible within the UK. In this following sections, we explain the motivation for detector upgrades, and give an updated overview of the technical planning in each area. Formal Technical Design Reports will be produced over the next two years, followed by comprehensive Engineering Design Reviews (Figure 1). In all cases, capital spend will follow the approval of the detailed CMS sub-detector plans by LHCC and STFC.

2 Scientific Motivation

The discovery in early LHC data of a 125 GeV resonance, with the properties and couplings expected of the Higgs boson, marks the final experimental verification of the Standard Model. However, the existence of a light scalar immediately implies physics beyond our current understanding, at a mass scale of less than a few TeV. The observed balance between matter, antimatter and dark matter in the Universe also requires new physics beyond the Standard Model (BSM). The future CMS science programme therefore focuses on the detection and exploration of such phenomena; the discovery of the Higgs boson represents the beginning, rather than the end, of our physics programme. Although several hundred different measurements are planned for the coming years, our future work has three key themes, each with substantial UK leadership:

- Precise measurement of the couplings of the 125 GeV boson will indicate whether it is in fact the Standard Model Higgs boson, and are a powerful tool in understanding physics at the TeV scale. Of particular importance are couplings to charged leptons and the Higgs self-coupling. In the both cases, precise measurements will only be possible with very the high statistics available at HL-LHC (see Figure 2).

- Supersymmetry is a strongly theoretically motivated mechanism for cancelling radiative corrections to the Higgs mass, which can also accommodate a dark matter candidate. CMS has already set mass limits of around a TeV for superpartners of Standard Model particles. These limits can be extended to over 2 TeV at HL-LHC, and hard-to-observe scenarios such as compressed SUSY can be more fully explored.

- Direct searches for production of new particles such as new gauge bosons, dark matter (detected via collisions with 'missing energy'), leptoquarks and other exotic objects will extend to much higher mass scales at HL-LHC, challenge models with ever-lower predicted cross sections, and provide limits that in combination with other approaches will provide decisive conclusions on BSM scenarios.

Each of these physics strands is already being pursued in LHC Run-2, and in the event of BSM physics existing with highly favourable parameters, rapid progress may be made before the start of the high-luminosity period. Nonetheless, in any conceivable scenario, HL-LHC will allow the fullest exploration of the new physics landscape, or the pursuit of new physics across the broadest range of models.

The scientific goals place extreme demands on the reconstruction efficiency for leptons and photons at HL-LHC, and require trigger thresholds that do not compromise the ability to select W, Z and H decays. In addition, excellent vertex reconstruction is required for b-tagging, and jet and energy flow resolution must be maintained in the face of extreme backgrounds. This latter requirement is particularly important in the high-η region, since the use of forward jets to detect vector boson fusion and vector boson scattering processes will be increasingly important.

It is important to note that the identification of all physics objects relies on 'particle flow reconstruction', either for object identification or isolation. This approach makes use of all available information from tracking, calorimetry and muon systems, combining them statistically to achieve optimum efficiency and resolution. The CMS upgrade strategy is therefore to improve the performance of each sub-detector in a comprehensive but balanced way, preserving our ability to investigate any unexpected signals of new physics in a wide range of final states. Moreover, the incorporation of tracking information into the Level-1 hardware trigger, allowing a similarly powerful approach for online selection, is a key objective of the upgrade programme.

3 Strategy and Objectives

Achievement of the required detector performance for HL-LHC will require a number of major technical developments. The most important and challenging of these are:

- A new silicon tracker, offering enhanced reconstruction performance, increased acceptance, and on-detector pattern recognition

- Off-detector electronics capable of real-time track reconstruction at high luminosity

- A highly granular endcap calorimeter (HGCal), providing high reconstruction performance in a high rate environment

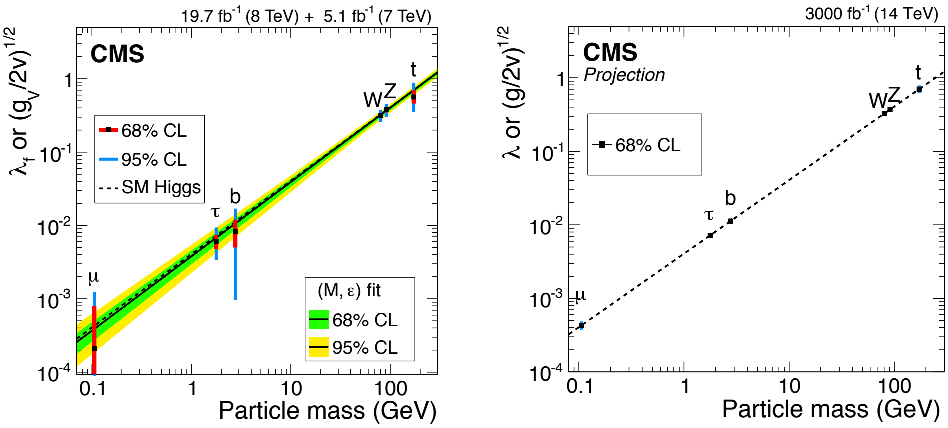

Figure 2: Higgs boson couplings as a function of boson or fermion masses, as measured by CMS (l) and projected measurements for HL-LHC (r)

- A new Level-1 trigger, using precise inputs from tracking, calorimetry and muon detectors.

As documented in the following sections, each of these new systems represents a major improvement on the state-of-the-art, is crucial for future CMS physics goals, and has required an intense development programme over several years. CMSUK proposed the key concepts behind all of these systems, has driven R&D activities, and has leadership in the detailed technical planning of the final system. We currently provide the CMS Project Managers for the Level-1 trigger, ECAL and HGCal, as well as the CMS Run Coordinator.

The single most demanding issue in the high-luminosity upgrade is the handling of extreme data rates resulting from increased detector granularity and much higher background levels. The original CMS electronics systems were start-of-the-art at the time of construction, and we will now be required to push the technological envelope in a similar way for the upgraded detector. In particular, efficient triggering will once again be the central problem. The performance of any hadron collider experiment is ultimately dependent upon the performance of the first-level trigger. At high luminosity, background rates increase enormously, whereas trigger thresholds (driven by the electroweak energy scale) cannot. Radical new approaches are required.

We therefore propose to focus our construction-phase deliverables on closely-interlinked elements of readout and trigger electronics for CMS (Figure 3). We will provide key elements of the readout systems of the tracker (WP3), and calorimeters (WP4), and construct the central part of the Level-1 trigger system (WP5). This will leverage the world-leading CMSUK capabilities in electronics, firmware and software, built up during the original CMS construction and the Phase-1 upgrade. We have identified common solutions to the technical requirements of these multiple systems, maintaining our historically broad contribution to CMS within a constrained funding envelope. Given the cross-project nature of these developments, we propose a common technology work package (WP2) underpinning all other CMSUK deliverables.

The UK plans form a key part of the overall CMS upgrade strategy. In all areas, we will work closely with (usually, long-standing) collaborating groups within CMS, combining capital

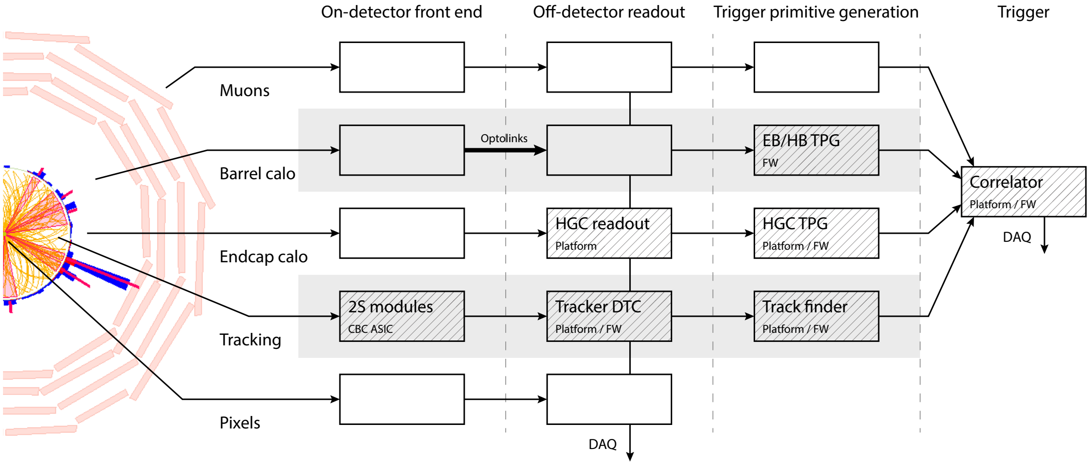

Figure 3: CMS high-luminosity readout electronics, with UK contributions indicated

and staff resources towards a common goal. In some areas (e.g. in the architecture of the track-finding electronics) multiple technical solutions exist, pending a collaboration decision before a TDR. In these cases, final decisions on responsibilities and roles will be made after decisions are known, but our expertise and strong track record assures a leading contribution in each of the areas defined above. The top level objectives of this proposal are:

1. Development and prototyping of a common readout and trigger electronic module for use across the CMS subsystems
2. Production and delivery of the outer tracker readout ASIC, and integration of tracker on-detector readout components
3. Development, production and integration of the tracker off-detector readout electronics and track-finding system.
4. Development, production and integration of off-detector readout and trigger electronics for the CMS calorimeters
5. Development, production and integration of the Level-1 trigger correlator unit.

4 Project Description

4.1 WP1: Project Management

Work package manager: D. Newbold

The project structure is illustrated in Figure 4. The Upgrade PI is a separate role from that of CMSUK PI. He or she works closely with deputies and WP managers to manage resources and personnel, to track progress against milestones, and to liaise with the Oversight Committee.

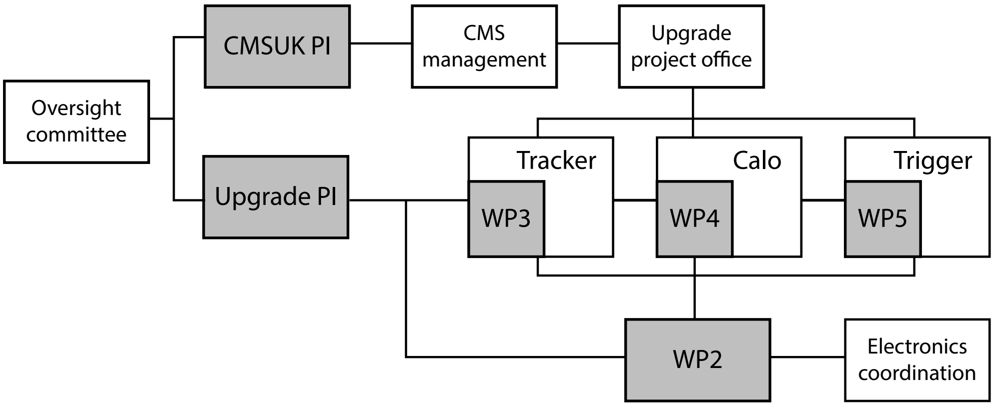

Figure 4: Project management structure (CMSUK roles shaded)

The project is divided into four technical work packages, with a summary of deliverables given in Table 1, and a top-level timeline in Figure 5. All work packages are integrated into the CMS upgrade project for each sub-detector, or in the case of WP2, the electronics integration coordination project. Detailed technical planning is now being put in place in preparation for sub-detector TDRs. A robust project management regime, jointly overseen by the CMS Upgrade Project Office and Technical Coordination teams, will be established using similar methods to those deployed during the original construction programme. UK deliverables are an integral part of the overall CMS plan, and will be tracked within it.

We will continue the successful management approach used during the R&D phase. Each work package is overseen by a pair of managers, typically with complementary skills and experience. In the case of the substantial WP3, we will explicitly appoint deputy managers with responsibility for the off-detector electronics project. Coordination will take place via monthly UK management meetings, with formal reporting of progress against deliverables, evolution of the risk budget, and the financial status, to the Oversight Committee on a six-monthly basis.

Our goal of purchasing several million pounds of electronics boards means that management of tendering, procurement and contracts is a key activity within the project. Project participants (G. Hall, G. Iles, T. Durkin) have significant experience in this area, and we will work closely with CERN and RAL procurement experts to ensure best practice and minimise risk. We will track closely the evolution of procurement rules for UK capital funds in light of our withdrawal from the EU.

There will be several senior personnel changes during the course of the project. In particular, several academic and engineering staff with highly relevant experience will be retiring, in most cases after a 20 year commitment to CMS. We recognise the importance of a succession plan in each case, with sufficient overlap period between new and existing post holders. Our current Upgrade R&D project has provided the opportunity for a new generation to gain experience of organisation and management of complex detector development projects, and several of these individuals will play more senior roles from 2019. CMSUK has ‘strength in depth’ in this respect.

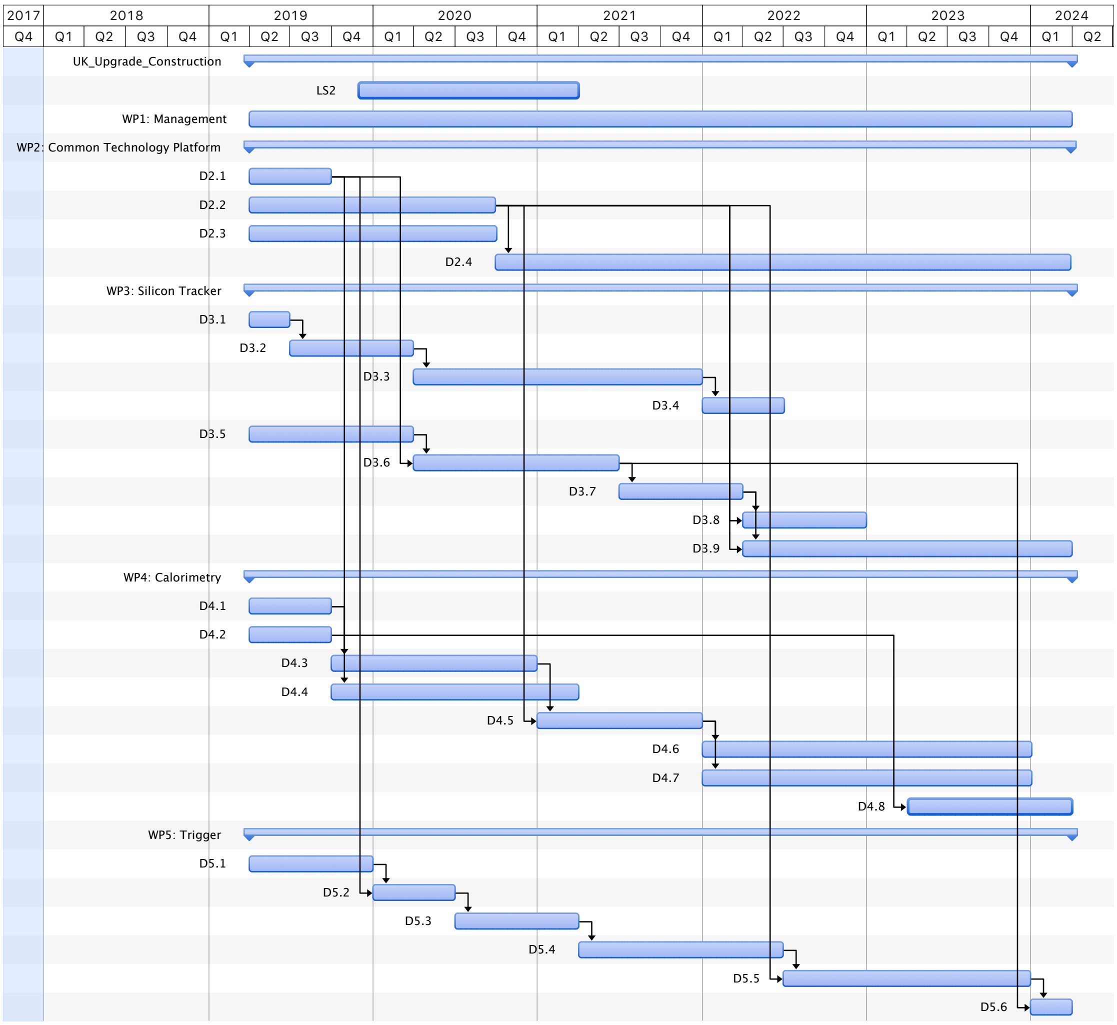

Figure 5: Top level project Gantt chart

Deliverable	Due date	Description
D2.1	Q3 2019	Common platform reference hardware
D2.2	Q3 2020	Common platform system firmware
D2.3	Q3 2020	Common platform system software
D2.4	Q1 2024	Common platform support and maintenance
D3.1	Q2 2019	Tracker CBC production order
D3.2	Q1 2020	Tracker CBC wafer production stand
D3.3	Q4 2021	Tracker CBC wafers
D3.4	Q2 2022	Tracker front end tests completed
D3.5	Q1 2020	Tracker off-detector architecture
D3.6	Q2 2021	Tracker prototype TF
D3.7	Q1 2022	Tracker TF readiness review
D3.8	Q4 2022	Tracker integration tests
D3.9	Q1 2024	Tracker TF final hardware
D4.1	Q3 2019	Calo baseline trigger algorithms
D4.2	Q3 2019	Calo ECAL optolinks design
D4.3	Q4 2020	Calo HGCal slice test
D4.4	Q1 2021	Calo ECAL baseline firmware
D4.5	Q4 2021	Calo HGCal pre-production boards
D4.6	Q4 2023	Calo HGCal hardware production
D4.7	Q4 2023	Calo final firmware
D4.8	Q1 2024	Calo optolinks
D5.1	Q4 2019	L1 trigger TDR
D5.2	Q2 2020	L1 trigger slice test
D5.3	Q1 2021	L1 trigger Run 3 system
D5.4	Q2 2022	L1 trigger final board design
D5.5	Q4 2023	L1 trigger hardware
D5.6	Q1 2024	L1 trigger integration

Table 1: Project deliverables

4.2 WP2: Common Technology Platform

Work package managers: K. Harder, G. Iles

4.2.1 Introduction

Completely new back-end electronics is essential for high-luminosity operation. It is necessary to handle higher bandwidth interfaces to the front-end electronics within CMS, and to process the resulting data for delivery to the DAQ system and to extract a trigger signal. To place the scale of the upgrade in perspective it is useful to consider the tracker readout, which will require an increase in bandwidth from $15\,\mathrm{Tbit\,s^{-1}}$ to $184\,\mathrm{Tbit\,s^{-1}}$.

We plan to build on the very successful R&D accomplishments of the last few years, in which the UK has developed general purpose state-of-the-art digital electronic boards for: the CMS Phase-1 trigger upgrade in 2015–16; the upgrades to the CMS timing and control system; the upgraded pixel and HCAL readout systems to be installed in 2016–7; and our own tracker R&D DAQ requirements. These boards, the MP7 and FC7 [5], were the highest-performance digital electronics in use in particle physics at the time. They have been accompanied by a comprehensive firmware and software package that has been adopted throughout and beyond CMS. This has only been possible due to the availability of increasingly powerful programmable logic (FPGAs) over the last two decades, coupled with a de-facto standard interface between systems in CMS (high speed multi-mode optical links) that has allowed a significant de-coupling between the hardware and the application. We are among the most experienced developers of advanced programmable logic and high speed optical technology, with an internationally recognised stature.

The de-coupling between hardware and application has propagated through the entire hardware, firmware and software stack so that systems with completely different algorithms can use the same, or very similar, online software. This has been proven to bring significant resource savings, in operations and maintenance as well as construction costs. This has enabled application projects to focus on the firmware and software necessary for their specific task, whilst the infrastructure load has been shared amongst collaborators. Both the flexible common hardware approach pioneered by CMSUK, and the IPbus distributed control technology developed to support it [6], are now in use within a wide range of experiments, including ATLAS, LHCb and ALICE, g-2, COMET, COMPASS, SOLID, and DUNE.

In this work package, we propose to design a generic, programmable processing card with high speed optical I/O, which will then be used either directly or with small modifications for the applications proposed in work packages 3, 4 and 5. Infrastructure firmware implementing all board-specific functions except the actual processing algorithms will be developed and distributed alongside the hardware. A software library providing an interface between the hardware/firmware and online software will be part of the package. The proposed WP2 deliverables are necessary prerequisites for WP3–5, and the majority of WP2 work will have to be completed before the final electronics construction period. Subject to approval, this work package will therefore start in 2017, funded by the R&D grant, and we describe here the entire scope of work.

It is not expected that a single board from this WP will satisfy all objectives; however, we believe that adaptations into a small family of very similar boards can meet the requirements of the UK-led projects. In particular, the power system, embedded computing, crate control, optical technology and cooling system will be the same for all variants of the hardware. These parts require significant design and testing effort, and the aim will be to share the activity with partners at CERN and elsewhere. There is a strong desire to make all these parts as interchangeable as possible so that the base card becomes a low risk object to manufacture. The infrastructure firmware and core software libraries will also be fully applicable to other projects and only require additional modules for covering project-specific functionality.

4.2.2 Work plan

The foundations for this common hardware are already being laid by means of prototype development using the latest programmable logic and optical links, building on the design and

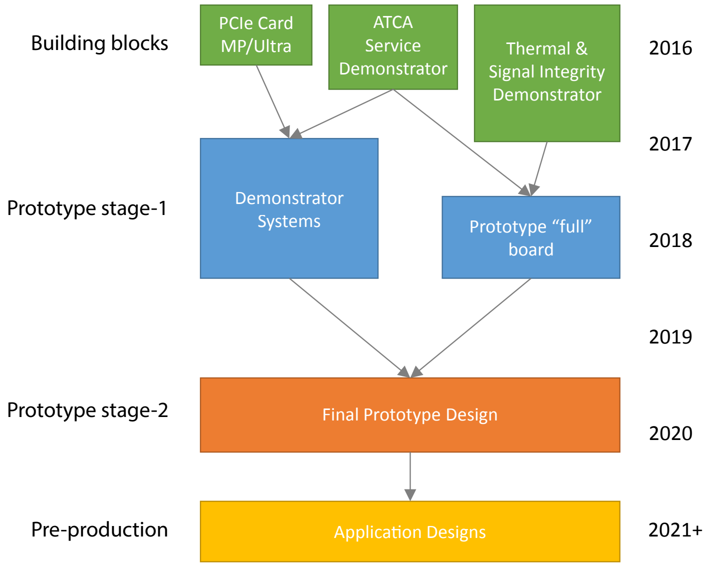

Figure 6: Roadmap for WP2

manufacturing experience of the last several years.

The first pre-production cards for the other work packages will be produced in 2021 with full production commencing around a year later. Before this stage, prototypes must be manufactured and evaluated for signal integrity, cooling capability, etc, and to develop the core infrastructure services. The prototypes will also play a crucial part in the development of the CMS electronics infrastructure and for detector prototype integration studies and test beams.

The roadmap for WP2 (Figure 6) foresees a preliminary stage that builds on a small form factor (PCIe) development card that is currently under test, the MP-Ultra. This card, which is a successor to the MP7 using the latest generation of FPGA, will enable us to validate many of the technologies we wish to use, but at a lower cost and in a simpler form because it requires less infrastructure overhead than a card designed for a managed crate environment. The custom MP-Ultra was developed to prototype technology because off-the shelf cards would be prohibitively expensive and occupy too much real estate in the space-limited underground counting room. The optical link, ram and processing density of the MP-Ultra is significantly beyond commercially available FPGA boards.

The MP-Ultra will be complemented by two other cards, both in the larger ATCA form factor. This form factor, or possibly a variant of it, has been selected by the CMS experiment for high-luminosity upgrade applications, where strict power and PCB floor-space limits will make MicroTCA unviable. The similarity of the MicroTCA and ATCA standards in terms of control and monitoring functionality allows us to build upon the substantial UK hardware and software experience with the MicroTCA architecture, built up during during the Phase-1 upgrade.

One of the two prototype cards proposed in addition to the MP-Ultra, the *service card*, will allow validation of all the ATCA services, such as embedded CPU, power, PCIe sub-system, IPMC, etc. This card will feature all components that a final ATCA blade requires, but without the high performance programmable logic and optical interfaces, and so will be much simpler and faster to design. The *thermal/signal integrity card* will enable validation of some of the key mechanical aspects of the design. Cooling is a particular concern because while simple resistor load cards have shown that up to 400 W per blade can be dissipated, it remains to be validated that this can be applied to a more realistic scenario where heat is concentrated in high power programmable logic and optics. Developments in other experiments have recently encountered significant issues in this area.

These strands of development will merge by 2018, with the service card hosting a mounted MP-Ultra for demonstrator systems and test beam operations, while the service and thermal/signal integrity card designs will be merged into a first prototype full board. After gaining experience from both these development paths, a final prototype will be produced in 2019. Final customisation, if required, for WP3–5 will take place in 2021.

4.2.3 Context and collaboration

Although the UK intends to take a leadership position in this development, engagement with collaborators across CMS is essential, both to share technical expertise and to ensure full take-up of the deliverables. Several institutes, including CERN and Fermilab, are also interested in the development of a hardware platform, including the associated infrastructure firmware & software. The previous informal collaboration in this area has now been formalised within the CMS electronics coordination group. As in Phase-1, the WP2 work will be an ‘open development’, with use beyond the proposed UK projects encouraged and anticipated – though it will be the responsibility of other upgrade sub-projects to determine which parts of the common development they can most effectively use.

CMSUK has a long history of collaboration with UK industry, and a large fraction of the back-end electronic cards currently installed in the counting room were manufactured in the UK. This extended to many production and all prototype cards developed by the UK for the Phase-I upgrade. It is our intention to continue this approach. However, procurements of large amounts of common hardware will involve CERN and non-European CMS users, and thus the tender process becomes fully international in scope. Furthermore, procurements using CORE funds will be made by other work packages in collaboration with their sub-system management and with reference to CERN tendering rules.

As we have done in the past, we will work closely with UK companies throughout the development phase of the project (i.e. prior to any tenders) so they are well placed to make successful bids. This requires a correct balancing of project risk with the desire to supporting possibly new UK partners in understanding HEP requirements. We will build here upon both our own experience and recent industry engagement work by other STFC projects, e.g. SKA.

4.2.4 Deliverables

D2.1: Common platform reference hardware (Q3 2019) Production of final prototype board for adoption by detector WPs.

D2.2: Common platform system firmware (Q3 2020) Final system firmware image, support both the prototype board and any further customisations required by sub-detectors (e.g. a move to 32 $\mathrm{Gbit\,s^{-1}}$ serial links). Note that working firmware is clearly required at all interim points, the deliverable marks an effective feature-freeze point.

D2.3: Common platform system software (Q3 2020) Final system software stack, supporting the prototype board and firmware, and any further features required by sub-detectors or higher level online software systems.

D2.4: Common platform support and maintenance (Q1 2024) Core firmware and software will need to be maintained for the lifetime of the hardware, including the remainder of construction project. Support will be offered for demonstrator systems and test beams, along with a low level of new feature additions all the way to the installation / commissioning period.

4.2.5 Resources

The existing grant will provide for the initial period of prototyping, but the final prototype, including the associated firmware & software package is scheduled for 2019/20 and thus will fall in the first year of the construction project. We require 7 FTE in 2019/20 for testing, firmware and software development, integration and quality assurance. For the subsequent four years, we estimate a requirement of around 12 FTE-year in total, with the earlier years having a larger fraction of this than the later years. This effort will be focused on providing the core features in the firmware and software package that are required by the other work packages; updating the package to remain in-sync with CERN online software and programmable logic vendor tools, and offering support to users. However, the bulk of effort (including expert staff from WP2) will migrate to WP3–5.

There is no CORE capital request for WP2; however, there is a non-CORE cost of £0.1M for prototypes and test equipment, plus some consumables costs at institutes and at CERN. This budget will be complemented by spending from the current R&D grant in early years.

4.3 WP3: Silicon Tracker

Work package managers: J. Goldstein, G. Hall

4.3.1 Introduction

The tracker plays a pivotal role in almost every CMS physics result, including the Higgs discovery and subsequent studies, measurements of rare decays such as $B_s \to \mu\mu$, and searches for supersymmetry and exotic physics, such as searches for long lived hadrons. The UK took a leading role in the design, construction, commissioning and operation of the current silicon strip tracker, and has made crucial contributions to the HL-LHC upgrade R&D. UK expertise will be essential in delivering the highly ambitious tracker upgrade described below, and we propose to focus on the critical areas where our capabilities can be used to greatest impact.

For the HL-LHC era CMS requires a completely new tracking system with reconstruction performance at least as good as at present, but able to cope with higher particle fluxes, detector occupancies, trigger rates, and radiation damage [3]. This necessitates a design with much higher granularity, presenting formidable challenges in the rate of data that must be read out from the detector and processed. However, the major challenge will be the inclusion of tracker data in the Level-1 trigger. The tracker must be able to transmit synchronous, low-latency track data to the trigger system, as well as asynchronous hit data for all selected crossings.

The upgraded detector will contain a pixel detector at small radii and an outer tracker instrumented as two regions. The outermost region (60–120 cm radially) will be populated by so-called 2S p_T-modules with two closely spaced microstrip sensor layers, with strips of 5 cm length. Level-1 triggering information is generated by correlating hits in the two layers *on-detector* to form 'stubs', in the process rejecting hits from low transverse momentum tracks. A similar approach is adopted for the region below 60 cm radius, but with increased segmentation in z, using one coarsely pixelated sensor layer.

Each of the approximately 15000 modules has a dedicated fibre-optic link to a Data Trigger and Control board (DTC). The DTCs receive stubs from the modules at the 40 MHz crossing rate, and send these data to a dedicated Track Finder (TF). Tracks must then be reconstructed sufficiently precisely within a few microseconds for the Level-1 trigger to associate them with information from the calorimeter and muon detectors. For a triggered crossing, the DTC will request and receive the full tracker data from the modules, sending it to the CMS DAQ and High Level Trigger. The processed data rate through the DTCs and TF approaches 200 $\mathrm{Tbit\,s^{-1}}$, making this by far the highest performance electronics system ever built for particle physics. Moreover, the algorithms for robust and efficient track-finding in hardware are complex and entirely novel, and it is necessary to compare the several different possible approaches though a realistic demonstration in hardware. This process is under way, with a collaboration-level decision due to be made early in 2017, before the tracker upgrade TDR.

4.3.2 Work plan

The UK, funded primarily by STFC, has been at the forefront of the CMS Phase II tracker R&D programme. We initiated and developed key concepts that have been accepted by the wider collaboration. In particular the UK proposed the p_T-module concept [7, 8], and has been developing the CBC (CMS Binary Chip) front-end ASIC for the 2S modules [9]. We have been collaborating closely with partners, particularly CERN, on the development and successful testing of the 2S modules themselves. This includes beam, laboratory and total dose radiation and single event effect tests. The UK has also developed the most advanced digital electronic hardware in use in CMS (the MP7 and FC7 microTCA modules) for the CMS Phase-1 upgrades [10, 11]. These developments, and our wide-ranging expertise, underpin our contributions to the new tracker.

The CBC is a 254-channel low-power device developed in 130 nm CMOS by Imperial College and RAL TD, and builds on their successful delivery of the APV25 front-end ASIC for the current tracker [12]. Successive versions of the chip [13] have incorporated all necessary features, and allowed us to gain confidence in the chip concept and implementation. The current version (CBC3) was submitted for manufacture in July 2016. Approximately 150,000 CBC chips are required for the tracker, and the development of the final version ready for production will be complete during the present R&D project.

During the construction project we must test the CBC chips on each of 350 wafers using an automated process. This task is foreseen to require about two years, and most of the preparation is expected to be carried out before March 2019 in the present R&D project. The plan is to carry out the testing at Imperial College, using facilities used for APV25 acceptance, and suitably updated for the CBC. RAL TD engineers will maintain design files, documentation and in-depth understanding of the chip for the long term.

Subsequent to chip production and testing, deep knowledge of the CBC is required to support production testing of complete modules, as well as their evaluation in beam and laboratory tests in collaboration with partners. These crucial quality control tasks are required before signing off the modules for full scale production. Other monitoring must continue beyond the CBC delivery to ensure that no aspects of the performance which might materialise at the module or system level have been overlooked. We will carry out irradiation and SEU studies on final modules, but tests emulating the high-rate environment will also be essential, along with investigation of any anomalies which might arise.

The UK has proposed a highly flexible architecture for the TF based on the time-multiplexed concept successfully developed for the Phase-1 trigger upgrade [14]. We have already assembled a prototype demonstrator system [15] using MP7 modules interconnected by $10\,\mathrm{Gbit\,s^{-1}}$ optical links, which is performing well. This system has clearly demonstrated that a TF based on affordable FPGAs, as opposed to ASICs, is a practical possibility. The use of commercial FPGAs is a key step in ensuring the future flexibility of the system, and reducing technical risk. However, it requires adaptation of algorithms to the particular combinations of memory, logic and IO found on FPGAs; this in turn requires deep expertise in FPGA architectures and capabilities, beyond simply implementing logic in RTL-level firmware. The UK design is one of three possible architectures being considered by CMS [16], and is acknowledged to be the most advanced of the concepts at this stage.

The architecture for the DTC, which interfaces between the modules, TF and DAQ, is also likely to be an FPGA board with a large number of high-speed optical links. Given the UK's expertise in this area we could play a major role in the design and delivery of this crucial component. The baseline plan is for both the TF and DTC to be variants of the board developed in WP2. The design and prototyping of the off-detector system will be accompanied by a comprehensive evaluation and testing programme, and we expect to play a leading role in this effort. The tests will naturally progress from board validation, through system testing with modules, to large-scale integration tests at CERN. The UK's strong participation in this programme will position UK collaboration members as key experts in the future detector commissioning and operations.

4.3.3 Context and Collaboration

The CBC is established as a key UK deliverable within CMS, and we have sole responsibility for its design, production and delivery to module production centres. Given our role and expertise, we also will continue to make a significant contribution to module and system-level testing in collaboration with significant partners such as CERN.

It is not yet possible to fully define the detailed activities to be undertaken by the UK in the off-detector electronics project. When the internal CMS review panel makes its recommendations in early 2017, based on assessment of the three proponent TF systems, a period of negotiation will be necessary to share the project in the most effective way. However, we expect that the system architecture to be adopted will closely resemble the one we have proposed, or be sufficiently close to it that we can adapt our contributions to it without difficulty. It is recognised that whatever TF choice CMS makes, the final readout and trigger primitives system will be designed and assembled by a collaboration formed from the different proponents to exploit similarities between the approaches and complementary resource and skill sets. Even if the time-multiplexed proposal is not chosen, the UK's contributions so far, key expertise, and substantial CORE contribution, ensure that it will play a leading role. In a scenario where we only contribute to the TF, as opposed to the DTC, we would expect to deliver $\sim$50% of the hardware required.

Other CMS teams who have announced already their intention to contribute in this area are the US (where the DoE and NSF expect to contribute about 50% of the track trigger system), Italy, Germany and France. Sizeable contributions are expected from the European partners, although at a lower level than the UK or US. These will all be required since the complete system comprising the DTCs and TF is sufficiently large and complex to mandate collaboration and cost sharing between several agencies.

4.3.4 Deliverables

Production of CBC chips for Outer Tracker:

D3.1: CBC production order (Q2 2019) Final design of CBC submitted for full production run.

D3.2: CBC wafer test stand (Q1 2020) Wafer test stand commissioned and ready to test production wafers.

D3.3: CBC wafers (Q4 2021) Final delivery of fully-tested CBC wafers to module production sites.

D3.4: Front end tests completed (Q2 2022) Full qualification of prototype outer tracker modules complete.

Off-detector electronic system:

We assume a scenario in which we focus on the development of the TF and contribute $\sim$50% of the TF boards.

D3.5: Off-detector architecture (Q1 2020) Definition of full of-detector electronics architecture agreed within CMS.

D3.6: Prototype TF (Q2 2021) Final prototype TF demonstrated.

D3.7: TF readiness review (Q1 2022) TF passes readiness review, enabling hardware production to start.

D3.8: Integration tests (Q4 2022) System integration tests with pre-production components begins at CERN.

D3.9: TF final hardware (Q1 2024) Delivery of final production TF hardware for qualification and installation.

Resources

The effort required for CBC production and support activities is not large provided that we are well prepared and retain expertise of staff trained during the R&D phase. We estimate that we will require about a total of 3 FTE in year 1, falling to 0.7 FTE in the final year of the project. The CBC wafer testing takes place in the first two years, while support for other users and module development will be at its maximum in year 1 and 2, but still be required in subsequent years; this is expert effort from designers. Similarly implementation of module testing, mainly in the form of DAQ and beam tests, will also peak at the outset, but decline to zero in the final year of the project.

In the first two years the track finder will require an average of 6 FTE per year. This will cover definition of the final architecture and hardware design (1 FTE) testing and operation of the demonstrator system using WP2 prototype hardware, including demonstrator firmware development and both online and emulator software (3 FTE), and simulation and performance studies of track finder algorithms based on results with the demonstrator system and compatibility with CMS requirements (2 FTE).

The subsequent three years require increased manpower on hardware construction, including preproduction design and manufacture, functional testing, production manufacture and quality assurance (2.5 FTE). An average of 2 FTE will be required to develop algorithm firmware for the final hardware, while reduced effort on simulation and performance studies will release effort for development of online software for the full system (3 FTE). As production hardware is assembled, 2.5 FTE will be needed for integration tests in the Tracker Integration Facility at CERN for testing with the sensors, DAQ and Level-1 trigger systems.

The UK CORE contribution to CBC production has been the subject of recent negotiations with CMS. The UK has paid for all the design effort and manufacture of the CBC to date. We expect to cover additional expenses through to the final preproduction run, which will provide the final CBC and mask set, from the R&D budget. However, beyond that point CMS proposes that CBC production wafers are paid from a CMS tracker common fund, so that UK resources can be used for the off-detector electronics. The UK would maintain its leadership of the CBC project. This allows the UK to contribute to the TF manufacture at a similar level to other major contributors, notably the US, which we believe will help share the costs, responsibilities and contributions in a balanced way.

4.4 WP4: Calorimetry

Work package managers: P. D. Dauncey, D. Petyt

4.4.1 Introduction

Physics at HL-LHC requires precision calorimetry in both the barrel and endcap regions. In particular, we will exploit Vector Boson Fusion (VBF) processes, which have a characteristic forward jet topology and are crucial to separate out different Higgs production modes. This will be achieved through matching extensions to the coverage of tracking (see WP3) and high precision calorimetry [3, 4]. The HGCal, a highly segmented sampling calorimeter, will replace the existing electromagnetic (ECAL) and hadronic (HCAL) calorimeter endcaps, providing an unprecedented level of granularity for reconstruction and triggering. The new detector will have 52 sampling layers, arranged in one ECAL and two HCAL sections. The ECAL and first HCAL sections will use silicon sensors as the active layers, with scintillator being used in the second HCAL section. The absorber materials are copper/tungsten in the ECAL and steel in the HCAL sections. The silicon sensors will have cell sizes of 0.5 and $1.0\,\mathrm{cm}^2$, while the scintillator tiles will be around 4 or $8\,\mathrm{cm}^2$, resulting in approximately 6M channels, compared with the existing endcaps which have around 20k channels in total. This increase of more than two orders of magnitude in granularity will result in much-improved pattern recognition for shower reconstruction. The aim is to maintain, or improve, on the current endcap performance, in the much higher interaction rate environment of the HL-LHC. Figure 7 illustrates the high occupancy and complexity of the data expected.

The optimal use of data from this novel detector will require a major study. Detailed studies of high granularity calorimetry have been undertaken for some years. However, none have ever been used in a colliding beam experiment (though similar detectors are proposed for practically all future collider experiments) and there is enormous need and scope for innovation and new ideas. One of the primary challenges posed by the HL-LHC is how to upgrade the readout and trigger to deal with the larger interaction rates, with corresponding increases in data volume and complexity. It is these issues that drive the UK efforts for both the ECAL upgrade and HGCal. We note that whilst offline reconstruction algorithms will be tested, refined and tuned with early data, the data-handling systems and trigger must function from the start, and their basic architecture cannot be changed.

The front-end and off-detector readout of the ECAL barrel will be upgraded in order to satisfy the increased CMS trigger performance requirements at HL-LHC. At present, ECAL information is available to the trigger at a granularity of 5×5 crystal groupings. The upgraded detector will provide full-resolution information on a single-crystal basis, allowing highly precise trigger primitive formation. This is essential both for improved stand-alone electron/photon identification and for matching with tracking information in hardware. It will also allow signals resulting from direct ionisation in the ECAL photodetectors to be efficiently rejected, which would otherwise present a major issue at HL-LHC luminosities.

In the HL-LHC era, excellent calorimetry will remain at the heart of the CMS detector concept, and is necessary for essentially all physics analyses. Delivery of the HGCal and ECAL upgrade will underwrite CMS future performance for physics analysis and, crucially, allow the

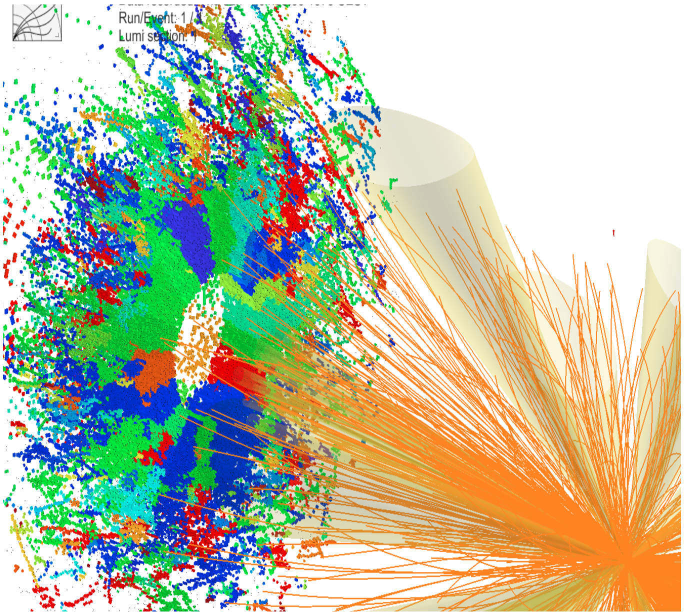

Figure 7: Simulated view of one HGCal endcap, containing particles from the nominal 140 pileup interactions expected at the HL-LHC.

full exploitation of the calorimeter performance in the Level-1 trigger.

4.4.2 UK involvement

The UK has a long history of involvement in the CMS ECAL, motivated by our interest in its exploitation for physics – in particular the Higgs sector. We initiated the original idea of using $PbWO_4$ crystals [17], and the excellent resolution that was subsequently achieved led directly to the Higgs discovery in 2012 [1]. Detailed studies of the Higgs boson are a cornerstone of our HL-LHC physics programme.

The UK led the design and construction of the crystal endcap detectors of the ECAL, and designed the readout ASIC that is used for both barrel and endcaps [18]. We followed up with all necessary steps to ensure the optimum use of the excellent ECAL resolution for physics, involving careful attention to calibration strategy, development of first level trigger algorithms for electron and photon identification, optimisation of advanced clustering algorithms both in the high level trigger and offline, and leadership of crucial physics analyses. The UK currently provides the System Manager for the ECAL, and we seek to maintain our leadership in calorimetry in the high-luminosity era.

We have subsequently made major contributions to the design and optimisation of the HGCal. The UK established the HGCal project. The current Project Manager is from the UK and we lead both the trigger and performance simulation sub-projects. We have so far chosen to focus on three central aspects of the HGCal, namely the front-end electronics, the trigger, and the simulation performance studies, all of which are highly synergetic with the UK's other leading roles in the

CMS upgrade. We have been successful in attracting ERC and STFC PRD funds, allowing a rapid start to the project. However, due to budgetary constraints, we do not propose to take all three areas forward, and only request funds to continue the trigger and simulation studies here.

4.4.3 Context and collaboration

The UK leads both the HGCal performance and trigger efforts. The former will require close collaboration with most of the institutes involved in the HGCal. For the trigger, the other groups involved are mainly from France and the US. It is agreed that the HGCal trigger hardware will be largely designed in the UK as part of WP2.

The ECAL trigger project is mainly led by US groups, who are also responsible for upgrades to the HCAL barrel and the design and production of the off-detector trigger boards for both the ECAL and HCAL barrel. Together, these will allow both parts of the barrel calorimeters to have improved and coherent trigger capability for HL-LHC. Hence, strong collaborative ties will be needed with US institutes.

In both the HGCal and ECAL, we will collaborate strongly with the central Level-1 trigger group, who will implement the electron/photon and tau/jet algorithms that will use the HGCal and barrel trigger primitives. The Level-1 trigger is currently led by the UK, and the HL-LHC upgrade subproject by the UK and US.

4.4.4 Work plan

Learning how to use the huge amount of information available from the HGCal is a major task. The particle flow concept requires individual particles within the calorimeter to be efficiently identified and associated to tracks. Once this pattern recognition problem is solved, putting the pieces together to obtain the excellent energy resolution required by the physics of the HL-LHC will be another significant study, requiring substantial intellectual effort to determine the optimal techniques to achieve these goals. Hence we request funds for effort for HGCal physics performance studies throughout the period of the proposal. As well as continuing to lead the development of the reconstruction code, we will undertake the crucial physics studies required to ensure that the HGC is optimally functional.

We also request funds for trigger work for both the HGCal and the ECAL barrel. This builds upon the significant expertise developed during the successful algorithm and firmware development for the recent Phase 1 CMS Calorimeter trigger upgrade. Each of the sub-detectors must perform substantial data reduction within the off-detector electronics, in order to generate a stream of energy clusters (so-called ‘trigger primitives’) that are sent to the Level-1 trigger for further processing. These data reduction steps will rely on a combination of time deconvolution, clustering, zero suppression, and data compression. The overall goal is to enable particle flow reconstruction in the Level-1 trigger (see WP5), a technique which needs to take place seamlessly across the whole detector coverage to give optimal performance. A consistent treatment and format is therefore needed for barrel and endcap data, and it is therefore important to have a single team defining this.

For the HGCal trigger, we lead the definition of the trigger architecture and algorithms, and subsequently the design, testing and production of the trigger processor board, based on the common platform (see WP2). The complexity of these algorithms and the interplay with detector design optimisation will require close cooperation with the UK physics performance studies.

While the HGCal combines the ECAL and HCAL sections, the barrel ECAL and HCAL are separate sub-detectors with different groups involved. Consistent inputs from the two parts of the barrel are required, and the two sub-detectors will use similar off-detector hardware, provided by US groups. Here, we will focus on applying UK expertise to the development of new clustering algorithms, maximising the use of the full information from the upgraded electronics of the ECAL barrel.

For both the barrel and endcaps, we will develop the firmware implementations of trigger primitive algorithms, firstly on demonstrator systems and prototype boards, and then on the final hardware. This will imply a long-term commitment to installation, commissioning and support up to and beyond the start of HL-LHC.

Finally, we will take responsibility for the optoelectronic links that transmit information from the ECAL on-detector electronics into the off-detector processing boards. This hardware contribution will make full use of the UK technical expertise in fibres, optoelectronics, and mechanical engineering, built up during the ECAL endcap construction and the more recent Level-1 trigger upgrade. It will form a valuable and requested technical contribution, complementing the US provision of the readout electronics.

4.4.5 Deliverables

D4.1: Calo baseline trigger algorithms (Q3 2019) Develop baseline trigger primitive algorithms and implement these in the CMS offline software framework. The conceptual designs will be fully documented in the HGCal EDR (Q2 2019), ECAL barrel EDR (Q3 2019) and Level-1 Trigger TDR (Q3 2019).

D4.2: Calo ECAL optolinks design (Q3 2019) Develop conceptual design and layout of optical fibre links between ECAL on-detector and off-detector electronics.

D4.3: Calo HGCal slice test (Q4 2020) Demonstration of a fully functional prototype HGCal trigger board. This includes a slice test demonstrator system of the complete chain of the trigger, to establish the feasibility of the architecture and allow evaluation of prototype trigger boards in a realistic environment.

D4.4: Calo ECAL baseline firmware (Q1 2021) Firmware implementation of the ECAL barrel algorithms, using project-developed off-detector boards, and full verification of the algorithm and firmware concept prior to the ECAL barrel ESR.

D4.5: Calo HGCal pre-production boards (Q4 2021) Pre-production of HGCal trigger boards. This will be a fabrication of a small sample of the boards for more advanced tests.

D4.6: Calo HGCal hardware production (Q4 2023) Completion of HGCal trigger board production, including production testing of boards when delivered.

D4.7: Calo final firmware (Q4 2023) Develop and test final algorithms and firmware on production off-detector boards from 2021 onwards. Fully functional algorithms and firmware must be available for detector installation and commissioning.

D4.8: Calo optolinks (Q1 2024) Procurement and manufacture of fibres and patch panel for ECAL barrel optical links.

4.4.6 Resources

The HGCal work will require an average of around 6 FTE, consisting of 2 FTE for the simulation and performance work, 2 FTE of 'physicist' effort in trigger algorithm development and software implementation, and the rest in firmware development. There is likely to be a large overlap of the four posts for performance and trigger algorithm development, as the techniques are common and there is clear synergy in these areas. The ECAL barrel project requires an average of around 2 FTE, with approximately 1 FTE devoted to algorithm development, participation in laboratory and beam tests and data analysis, and 1 FTE devoted to firmware development. Engineering effort for the mechanical design of the optical patch panel is required, but is accounted as part of the CORE cost.

The HGCal project is organised with a budget for common items, such as silicon sensors, to which all participants are asked to contribute. We foresee a contribution of 1.0 MCHF, of which 0.5 MCHF would be from non-STFC sources. We will also make a contribution of 2.0 MCHF to the HGCAL trigger hardware, corresponding to around 80% of the system, reflecting our leadership in this area. A contribution of 0.5 MHCF to the CORE cost of the ECAL upgrade is foreseen, to be spent largely on optolink purchase and testing, and associated items.

4.5 WP5: Trigger

Work package managers: J. Brooke, A. Tapper

4.5.1 Introduction

The CMS trigger system is required to select or reject each beam-crossing that occurs within the CMS detector, based on its utility for later analysis. This task must be completed within a few microseconds of the collision, and with a rejection ratio of around 100:1. Decisions made by the trigger are irrevocable, and so this real-time hardware processor effectively represents the first step in data analysis. CMSUK has had a strong role in the trigger project from the start of CMS, and currently holds all the main leadership positions in the project.

The high-luminosity physics programme requires that the efficiency for capturing electroweak scale physics is as good as, or better than that achieved in LHC Run 1. Applying this to the Level-1 trigger, this implies that the thresholds for selection of final-state physics objects at the trigger level must be maintained at the same levels as Run 1, but with many times the background. This in turn leads to the requirement that much of the CMS detector electronics and data acquisition system must be replaced [3]. The trigger is therefore at the centre of the CMS high luminosity upgrade programme, and connected fundamentally to the physics potential of the experiment. The improved technology available by 2024 will allow fundamentally new approaches to be taken to triggering, but will require substantial development before it can be exploited.

4.5.2 UK involvement

The UK has a long and highly successful history of leadership in the CMS Level-1 trigger, and other related projects. The original CMS global calorimeter trigger was a UK deliverable, and despite pushing the technology envelope, provided reliable, high-performance triggering for CMS in LHC Run 1 [19]. The UK led the campaign to replace this with yet higher-performance hardware in the Phase-1 upgrade, which was recently successfully completed [20]. Many novel ideas have originated in the UK and been adopted by CMS, including the 'time multiplexed' trigger architecture [21, 22] used in the Phase-1 calorimeter trigger and the 'stacked tracking' approach to producing track stubs for the Level-1 trigger [23]. The UK also led trigger related activities in the ECAL endcap sub-detector. Our work on the Level-1 trigger links directly with UK physics interests, particularly in the areas of Higgs boson and supersymmetry studies, and many STFC-funded students have earned the right to be CMS authors through their work on the trigger before transitioning to physics analysis.

The STFC-funded R&D programme which led to the delivery of the Phase 1 calorimeter trigger upgrade was hugely successful. Not only was the trigger upgrade delivered successfully and deployed for first running in 2016, but the technology R&D has had far-reaching impact as described in the WP2 section. CMSUK has built up world-leading expertise in state of the art electronics, in particular in FPGA programming and high-speed optical data transmission.

4.5.3 Context and collaboration

CMS will replace the electronics for the barrel sections of the ECAL and HCAL, as well as building an entirely new high-granularity endcap calorimeter for the high luminosity upgrade. In addition, electronics in the muon systems will be replaced, and a new tracker, with Level-1 trigger capabilities will be built. The trigger is a key driver for all these upgrades, which will each deliver a hugely increased amount of data to the trigger system in order to fulfil the physics requirements detailed above. CMSUK proposes to contribute to almost all of these upgrades, and past work and ongoing R&D places the CMS UK group in a unique position at the intersection of these projects. WP5 will focus on the design and implementation of algorithms to use these data to make a trigger decision; this clearly will have a decisive impact on the trigger upgrade strategy, and links closely to our physics interests. The proposed programme will be executed in collaboration with colleagues from a number of US institutes, with a roughly 50/50 split of the project costs and activities foreseen.

4.5.4 Work plan

The CMS high luminosity trigger will, for the first time, receive data from the tracker, offering greater selectivity compared to the current calorimeter and muon based trigger system. Lepton triggers will benefit from track-based identification (in the case of electrons and taus), and improved transverse energy measurement in the case of muons, as well as making use of track-based isolation. Photon, jet, and energy sum triggers will benefit from precise online vertex identification.

The use of tracking information, together with latest generation FPGAs, will allow full event reconstruction to be performed in the Level-1 trigger for the first time. Our choice of algorithms will be guided by the latest offline reconstruction techniques [24]. The extent to which offline algorithms can be executed in hardware within $\sim 5\,\mu\text{s}$ is the subject of intense study, and will depend on the available technology towards the end of the construction period. For this reason, the finalisation of the trigger design will come later than for other sub-detectors (see Figure 1), though sub-detector interfaces must be frozen much earlier in the project. The input to the Level-1 trigger will comprise: fitted tracks from the outer tracker; calorimeter clusters, based on the highest possible granularity information; muon track stubs from the different muon detectors; and possibly some information from the pixel detector.

The central component of the Level-1 trigger will be the correlator unit (see Figure 3), where tracks are matched to calorimeter objects. The design and construction of this system is the focus of this work package. We will develop algorithms for electron, jet, energy sum and vertex identification, inspired by the latest offline reconstruction techniques. We envisage that event reconstruction will be performed in two stages. First, tracks are matched with calorimeter objects to produce candidate photons, electrons, neutral hadrons, and charged hadrons. In parallel, track-based primary vertex identification takes place. The second stage then uses candidates associated with the primary vertex to identify isolated electrons and photons, jets, and to calculate energy sums. These algorithms will be complemented by muon identification, to be developed by other groups.

We will develop these algorithms through simulations in software and firmware emulation, and later implement these algorithms in hardware using a processor card based on the reference designs produced by WP2. Infrastructure to support operation of the processing cards will be developed in WP2, however some application-specific components will be developed in WP5. A detailed bit-level software emulation of the system will be required for simulation and monitoring purposes, and these tools maintained into the commissioning period as the primary means of verifying trigger performance against data.

The lessons learnt in the development of the advanced algorithms for high luminosity may be relevant to trigger performance in LHC Run 3. As part of the ongoing campaign to optimise trigger performance, it may be useful to include some aspects as part of an intermediate upgrade during LS2, using existing or new hardware. In any event, provision will be made in LS2 to install a fraction of the new system, with the objective of commissioning the infrastructure developed in WP2, in situ during Run 3 data taking. Following the Level-1 trigger TDR in 2019 the processor card design will be finalised before the 2021 ESR. Production of final hardware and software will proceed during 2022–23, with system installation in 2024. This schedule is based closely on that of the Phase-1 upgrade, which reinforced the need for an extended period of basic system commissioning for an object as complex as the Level-1 trigger, followed by commissioning with data. These steps will take place in 2024–6, in preparation for first high luminosity data.

4.5.5 Deliverables

D5.1: L1 trigger TDR (Q4 2019) The Technical Design Report will document baseline algorithms and their simulated performance, along with the system architecture and construction planning.

D5.2: L1 trigger slice test (Q2 2020) A full vertical slice test will be carried out using WP2 reference hardware, including interconnection tests to sub-detector prototype hardware.

D5.3: L1 trigger Run 3 system (Q4 2020) Implementation of any possible improvements to Phase 1 calorimeter trigger, installation of Run 3 test system at P5 during LS2.

D5.4: L1 trigger final board design (Q2 2022) Final board designs, ready for production and CMS Electronics System Review.

D5.5: L1 trigger hardware (Q4 2023) All hardware delivered, tested and at CERN.

D5.6: L1 trigger integration (Q1 2024) System operational and ready for integration with sub-detectors and commissioning.

4.5.6 Resources

The UK project requires personnel with strong skills in software and firmware, along with a core familiarity with the CMS physics programme and trigger strategy. Physics-focussed and detector-focussed personnel must be able to work together as a single team. In addition, substantial hardware expertise is required at key times during installation, testing and commissioning. CMSUK has the required skills in all of these areas, and we will take the opportunity to augment the core team from the Phase-1 trigger upgrade with additional input from specialists in the use of calorimeter and tracking information at HLT level. The required size of this team is well understood from the Phase-1 upgrade, though since the project is wider in scope, collaboration with other strong CMS teams is mandatory. We estimate a requirement of around 5 FTE, augmented in the latter two years by student effort for integration at CERN.

We note that the integration, commissioning and tuning of the system in the period 2023–26 is a key opportunity for students and non-specialist physicists to make a deep contribution to the high luminosity performance of CMS during LS3.

The CORE cost estimate of the CMS project is around 1 MCHF. We propose to contribute half of this cost, with the remainder coming from a small number of DOE-funded US institutes. This will ensure a closely-coupled collaboration, which we feel is essential for a small but highly complex development of this type, with strong UK leadership.

5 Key Participants

- Dr Jim Brooke (WP5 manager) is CMS coordinator for the high luminosity trigger upgrade, and co-manager of the current UK trigger upgrade project.

- Dr Oliver Büchmuller (WP5) is convenor of the CMS exotica working group, and was previously coordinator of the Trigger Strategy Working Group for the CMS upgrade.
- Prof. Paul Dauncey (WP4 manager) is head of the Imperial HEP group. He is leading the HGCal trigger project and has previous experience in triggering and high granularity calorimeters through his work in the BaBar and CALICE experiments.
- Prof. Gavin Davies (Deputy PI, WP1) is Imperial CMS PI. He played a significant role in the design of the original ECAL and is now a leading member of the Higgs group. He also has extensive project management and Higgs analysis experience from the Tevatron.
- Prof. Joel Goldstein (WP3 manager) is head of the Bristol HEP group, with experience in tracking detectors at collider experiments. He was responsible for the first generation of CMS upgrade front-end hybrid/module test stands and coordinated their distribution to test centres.
- Prof. Geoff Hall (WP3 manager) is PI on the UK CMS R&D project, and a member of the CMS and CMS tracker management boards. He led the UK contribution to the construction and operation of the current CMS tracker, including the development of the APV25 ASIC and readout boards.
- Dr Kristian Harder (WP2 manager) has over ten years of experience in developing readout electronics, online software, control and monitoring systems. He is one of the main microTCA infrastructure experts in the CMS collaboration. He has considerable experience of Si detectors, and was the D0 Si detector group leader.
- Prof. Peter Hobson (WP3) is CMS group leader at Brunel, and has long experience of calorimetry and tracking developments for particle physics experiments.
- Dr Greg Iles (WP2 manager) is an electronic engineer who was responsible for the Phase-I Calorimeter Trigger Upgrade system design. Working with Dr Rose, he pioneered the use 10Gb/s multi-board optics and MicroTCA systems in CMS. He is a member of the TWEPP conference organising committee.
- Prof. Jordan Nash (WP4) is currently Head of the Physics Department at Imperial College. He has previously been CMS upgrade project manager and ECAL electronics coordinator. He has varied and extensive hardware experience from ALEPH and BaBar.
- Prof. David Newbold (PI, WP1 manager) is CMSUK PI and a CMS MB member (until September 2017), and chair of the Trigger and DAQ Institution Board, with expertise in trigger and readout systems. He was previously co-manager of the UK trigger upgrade project and CMS upgrade simulations coordinator.
- Dr Sudarshan Paramesvaran (WP2,3,5) is CMS Run Coordinator, and previously served as commissioning coordinator for the Phase-1 trigger upgrade and CMS HCAL operations manager.

- Dr Mark Pesaresi (WP3 deputy manager) is joint UK Track Trigger project manager who was responsible for the seminal studies of the CMS Track Trigger. He won an Achievement Award in 2014 for this; his contribution to the CBC ASIC chip validation; and development of the FC7 MicroTCA hardware.
- Dr David Petyt (WP4 manager) is the CMS ECAL System Manager. He previously held the position of ECAL Deputy Project Manager (2013-15) and was ECAL Detector Performance group convenor during the period of the Higgs boson discovery (2011-12).
- Dr Andrew Rose (WP2) is an electronic engineer who implemented the complex algorithms used in the Phase-1 Trigger Upgrade and has contributed extensively to the hardware, firmware and software of the MP7 (MicroTCA card). He won the CMS Young Researcher award in 2015.
- Prof. Claire Shepherd-Themistocleous (Deputy PI, WP1) is CMSUK Deputy PI, and Head of Division in RAL PPD. She has led numerous CMS analyses with particular emphasis on electron reconstruction and triggering. She was Deputy Chair of the ECAL IB.
- Dr Alex Tapper (WP5 manager) is CMS System Manager of the Level-1 trigger, and was previously trigger upgrade coordinator and co-manager of the UK trigger upgrade project.
- Dr Alessandro Thea (WP2) is Online Software group coordinator for the Level-1 trigger. He led the development of the new control and monitoring software framework for the trigger upgrade project. He was previously the ECAL Data Acquisition group coordinator (2009-11).
- Dr Ian Tomalin (WP3 deputy manager) has worked on tracking detectors on CMS and ALEPH and has expertise in track reconstruction. From analyses involving long-lived particles he has a detailed knowledge of how to reconstruct non-standard tracks and the impact of these in triggers. He was editor of the major paper on track and vertex reconstruction in CMS.
- Prof. Jim Virdee (WP4) created the HGCal project and is the overall project manager. He is one of the founding members of CMS and was CMS Spokesperson for several years. He was the initiator of the existing crystal ECAL detector and served as ECAL project manager.

6 Resource Request and Costs

The resource and capital costs of the upgrade construction project are summarised in the tables below. In all cases, the amounts stated are non-inflated costs to STFC (i.e. 80% staff costs). While staff and related costs have essentially a flat profile, capital costs have a back-weighted profile, both in light of the nature of the deliverables (digital electronics is typically purchased as late as possible in order to take advantage of falling prices), and in response to STFC guidance.

We request no new posts in this project. All project-funded posts are either a continuation from the current R&D phase, or will continue from the current ERC-funded HGCal project or the related STFC PRD project. This will ensure a direct continuation of current expertise, with no risk of project delays due to recruitment. Indeed, in order to remain within the available financial envelope, some posts from those projects must be terminated, though plans will be put in place to transfer required expertise. We request a number of project studentships. During the R&D phase, project students have been a key asset; some of the most original technical developments have come from students. Along with ongoing CMS exploitation, the construction project will form an excellent training ground for a new generation of physicists, and the funding of project students is essential to allow us to meet our deliverables. The project-funded staff contingent is complemented by consolidated grant funded staff from Universities and permanent staff at RAL, some of whom will manage and oversee the project.

We note that estimates of required effort for each project deliverable are based on extensive experience of similar projects from the original CMS construction, the R&D project and the CMS Phase-1 upgrade. The most recent construction deliverable undertaken by CMSUK (the calorimeter trigger upgrade) was delivered on time and on budget, but has caused us to re-examine the balance of effort required between hardware, firmware and software deliverables; the lessons learnt are incorporated in our planning here. The total requested project-funded effort is 65.2 FTE-year, of which 18 FTE-year is students, complemented by 77.6 FTE-year of existing staff effort (see Table 2).

Alongside staff costs, we request modest funds for long-term attachment and short-term travel to CERN, and a general laboratory consumables budget for each technical work package, to be used both in institutes and to equip laboratories and test areas at CERN. These costs are based on actual expenditure during the R&D phase, scaled by the number of personnel involved.

Our capital request covers the direct costs of detector construction, i.e. CORE costs in the CMS cost book. A breakdown of the CORE contribution is shown in Table 3. In this table, costs are given in MCHF for consistency with CERN planning tables already disseminated to STFC; in the cost spreadsheets accompanying this document, average exchange rates of 1.4 CHF/£ and 1.2 €/£ have been assumed. We note that some capital costs are covered outside the current proposal, by capital funds already allocated for the R&D project, or from ERC funds. We note that some CORE contributions from the UK do not appear in the total project cost, since it has been agreed that they should be paid outside the period 2019-24, and / or from other sources. In addition to the CORE costs, we request a £100k capital budget for hardware prototyping in WP2 in 2019–21, since only final production prototypes are accounted for within CORE.

Alternative scenarios

This proposal has been constructed under conditions of severe resource constraint. Even in the baseline scenario, the size of our staff will decrease as we go from the R&D phase to construction. The 30% cut already imposed following the 2014 review has necessitated a rebalancing of staff costs against capital, in order to achieve a coherent project which guarantees delivery whilst maintaining UK leadership in key areas. We note that the entire 8 MCHF CORE budget in the baseline scenario (including the non-project contributions) is less than 65% of the UK nominal ‘pro rata’ contribution to the overall CMS upgrade cost, as approved by LHCC. This is the lowest per capita contribution of any major contributing country, and roughly half that of France and Germany. In light of our potential intellectual contribution to the upgrade, CMS has granted the UK considerable flexibility in the profile as well as level of CORE, including payment of the cost of ASIC production in the UK.

A budget increase of $+10\%$ would not be sufficient to prudently increase the scope of the project. Our first priority would therefore be to address the CORE capital budget. In particular, we would seek to increase our contribution to the TF/DTC construction, ensuring UK leadership in the implementation of a system originally proposed and prototyped by us, and excellently matched to our skills. We would also seek to decrease project risk in this and other deliverables by increasing working margin, mainly to allow targeted additional engineering effort where needed.

In a -10% scenario, we would be forced to further reduce our CORE contribution whilst preserving posts, allowing at least some intellectual contribution to systems where we cannot construct hardware. This of course implies a loss of UK leadership, but would allow skills that are unique to the UK to be exploited by CMS. However, a further reduction in CORE may not be acceptable within the CMS collaboration, and so in this scenario a difficult negotiation will be necessary before making a definite plan. A substantial reduction in the scope of the project (e.g. via loss of an entire work package) is not possible without decoupling a substantial part of the CMSUK collaboration from the upgrade programme.

7 Risks

An outline project risk register is shown in Tables 4 and 5. A detailed CMS-level risk register for each subproject will be established before the start of the approved construction phase, and documented in the TDRs.

In general, there are six categories of risks in the technical work packages:

1. Mis-specification or change of requirements or interfaces
2. Mis-costing or cost fluctuation in components
3. Insufficient processing capacity in systems
4. Technical failure during implementation
5. Failure of suppliers or the supply chain
6. Force majeure, acts of God, etc

Of these, we regard 1–3 as risks primarily belonging to the CMS overarching projects, but where the CMSUK has a role in defining specifications, estimates and contingency budgets in concert with CMS management. The use of flexible hardware and programmable logic is a key element in reducing capacity- and specification-related risks. 4–5 are risks primarily belonging to the UK project, and where the mitigation is indicated in the risk register. For this reason, we do not specify working margin or contingency for capital costs. For 6, we assume that STFC holds sufficient contingency to handle large-scale disruptions.

Based on our risk register, we request a working allowance of £300k (around 6% of staff cost). This is sufficient to cover any mitigation of any of the most major events, plus any combination of the less major ones. These mitigations almost uniformly consist of addition of engineering effort to the project, and this we would cover by moving tasks to RAL TD or industry. We exclude from this estimate risk R1.2, concerning overall slippage of the LHC schedule. This is outside UK control, and no practical mitigation can be made. Nonetheless, there is a substantial STFC risk exposure that should be considered in allocating contingency.

References

[1] S. Chatrchyan et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 30 –61. DOI: `http://dx.doi.org/10.1016/j.physletb.2012.08.021`.

[2] “The European Strategy for Particle Physics Update 2013. La stratégie européenne pour la physique des particules Mise à jour 2013. 16th Session of European Strategy Council”. In: (2013). URL: `https://cds.cern.ch/record/1567258`.

[3] D Contardo et al. *Technical Proposal for the Phase-II Upgrade of the CMS Detector*. Tech. rep. CERN-LHCC-2015-010. LHCC-P-008. CMS-TDR-15-02. Geneva, 2015. URL: `https://cds.cern.ch/record/2020886`.

[4] J Butler et al. *CMS Phase II Upgrade Scope Document*. Tech. rep. CERN-LHCC-2015-019. LHCC-G-165. Geneva: CERN, 2015. URL: `https://cds.cern.ch/record/2055167`.

[5] K Compton et al. “The MP7 and CTP-6: multi-hundred Gbps processing boards for calorimeter trigger upgrades at CMS”. In: *Journal of Instrumentation* 7.12 (2012), p. C12024. URL: `http://stacks.iop.org/1748-0221/7/i=12/a=C12024`.

[6] C. Ghabrous Larrea et al. “IPbus: a flexible Ethernet-based control system for xTCA hardware”. In: *Journal of Instrumentation* 10.02 (2015), p. C02019. URL: `http://stacks.iop.org/1748-0221/10/i=02/a=C02019`.

[7] J. Jones et al. “A Pixel Detector for Level-1 Triggering at SLHC”. In: *ArXiv Physics e-prints* (Oct. 2005). eprint: `physics/0510228`. URL: `https://arxiv.org/abs/physics/0510228`.

[8] M Pesaresi and G Hall. “Simulating the performance of a p T tracking trigger for CMS”. In: *Journal of Instrumentation* 5.08 (2010), p. C08003. URL: `http://stacks.iop.org/1748-0221/5/i=08/a=C08003`.

[9] D Braga et al. “CBC2: a microstrip readout ASIC with coincidence logic for trigger primitives at HL-LHC”. In: *Journal of Instrumentation* 7.10 (2012), p. C10003. URL: `http://stacks.iop.org/1748-0221/7/i=10/a=C10003`.

[10] M. Pesaresi et al. “The FC7 AMC for generic DAQ & control applications in CMS”. In: *Journal of Instrumentation* 10.03 (2015), p. C03036. URL: `http://stacks.iop.org/1748-0221/10/i=03/a=C03036`.

[11] G. Auzinger. “Deployment of the CMS Tracker AMC as backend for the CMS pixel detector”. In: *Journal of Instrumentation* 11.01 (2016), p. C01056. URL: `http://stacks.iop.org/1748-0221/11/i=01/a=C01056`.

[12] M Raymond et al. “The CMS tracker APV25 $0.25\mu m$ CMOS readout chip”. In: (2000). URL: `https://cds.cern.ch/record/478270`.

[13] M Raymond et al. “The CMS binary chip for microstrip tracker readout at the SLHC”. In: *Journal of Instrumentation* 7.01 (2012), p. C01033. URL: `http://stacks.iop.org/1748-0221/7/i=01/a=C01033`.

[14] G. Hall et al. “A time-multiplexed track-trigger architecture for CMS”. In: *Journal of Instrumentation* 9.10 (2014), p. C10034. URL: `http://stacks.iop.org/1748-0221/9/i=10/a=C10034`.

[15] C. Amstutz et al. “An FPGA-based track finder for the L1 trigger of the CMS experiment at the high luminosity LHC”. In: *2016 IEEE-NPSS Real Time Conference (RT)*. 2016, pp. 1–9. DOI: `10.1109/RTC.2016.7543102`.

[16] F. Palla, M. Pesaresi, and A. Ryd. “Track Finding in CMS for the Level-1 Trigger at the HL-LHC”. In: *Journal of Instrumentation* 11.03 (2016), p. C03011. URL: `http://stacks.iop.org/1748-0221/11/i=03/a=C03011`.

[17] “The Electromagnetic Calorimeter Project: Technical Design Report”. In: *CERN/LHCC* 97-033 (1997). CMS TDR 4.

[18] M. Raymond et al. “The MGPA Electromagnetic readout chip for CMS”. In: *Proc. of the 9th workshop on electronics for the LHC experiments* CERN-LHCC-2003-005 (2003).

[19] J. Brooke and on behalf of the CMS Collaboration. *Performance of the CMS Level-1 Trigger*. 2013. eprint: `arXiv:1302.2469`.

[20] A Tapper and Darin Acosta. *CMS Technical Design Report for the Level-1 Trigger Upgrade*. Tech. rep. CERN-LHCC-2013-011. CMS-TDR-12. 2013. URL: `https://cds.cern.ch/record/1556311`.

[21] B Kreis et al. “Run 2 Upgrades to the CMS Level-1 Calorimeter Trigger”. In: *J. Instrum.* 11.arXiv:1511.05855. 01 (2015), C01051. 10 p. URL: `https://cds.cern.ch/record/2103853`.

[22] Alexandre Zabi. *Triggering on electrons, jets and tau leptons with the CMS upgraded calorimeter trigger for the LHC RUN II*. Tech. rep. CMS-CR-2015-297. Geneva: CERN, 2015. URL: `https://cds.cern.ch/record/2105511`.

[23] A Tricomi. “Upgrade of the CMS tracker”. In: *Journal of Instrumentation* 9.03 (2014), p. C03041. URL: `http://stacks.iop.org/1748-0221/9/i=03/a=C03041`.

[24] *Particle-Flow Event Reconstruction in CMS and Performance for Jets, Taus, and MET*. Tech. rep. CMS-PAS-PFT-09-001. 2009. Geneva: CERN, 2009. URL: `https://cds.cern.ch/record/1194487`.

WP	Institute	Name	Post	19/20	20/21	21/22	22/23	23/24
WP1	Bristol	Newbold	C	0.2	0.2	0.2	0.2	0.2
	Imperial	Davies	C	0.1	0.1	0.1	0.1	0.1
	RAL	Shepherd-T.	C	0.1	0.1	0.1	0.1	0.1
	Total			**0.4**	**0.4**	**0.4**	**0.4**	**0.4**
WP2	Bristol	Kreczko	P	0.5	0.5	0.5	0.5	0.5
		Newbold	C	0.2	0.2	0.2	0.2	0.2
		Paramesvaran	C	0.5				
		Student	P	0.25	0.5	0.5	0.5	
	Imperial	Bundock	P	0.5	0.2	0.1	0.1	0.1
		Hall	C	0.2	0.2	0.2	0.2	0.2
		Iles	C	0.7	0.3	0.3	0.25	0.2
		Paladino	O/P	0.6	0.3	0.2	0.1	0.1
		Pesaresi	P	0.3	0.3	0.2	0.2	0.2
		Rose	C	0.9	0.3	0.2	0.1	0.1
	RAL	Durkin	S	0.25	0.25			
		Harder	S	0.75	0.25	0.25	0.25	0.25
		Manolopoulos	S	0.25	0.25	0.25	0.25	0.25
		Thea	P	0.5	0.1	0.1	0.1	0.1
		Williams	S	0.5	0.2	0.1	0.1	0.1
	Total			**6.9**	**3.85**	**3.1**	**2.85**	**2.3**

WP	Institute	Name	Post	19/20	20/21	21/22	22/23	23/24
WP3	Bristol	Baesso	C	0.25	0.25	0.25	0.25	0.25
		Goldstein	C	0.3	0.3	0.3	0.3	0.3
		Paramesvaran	C	0.2	0.2	0.2	0.2	0.2
		Seif El Nasr	P	1.0	1.0	1.0	1.0	1.0
		Student	P	0.25	0.5	0.5	0.5	
	Brunel	Cole	O/C	0.08	0.15	0.15	0.15	0.15
		Hobson	O/C	0.05	0.1	0.1	0.1	0.1
		Khan	O/C	0.05	0.1	0.1	0.1	0.1
		Reid	P	0.5	0.5	0.5	0.5	0.5
	Imperial	Auzinger	P	0.5	0.5	0.5	0.5	0.5
		Davies	C	0.1	0.1	0.1	0.1	0.1
		Greenwood	P		0.5	0.25	0.25	0.25
		Hall	C	0.3	0.3	0.3	0.3	0.3
		Iles	C		0.6	0.5	0.4	0.3
		Pesaresi	P	0.7	0.7	0.8	0.8	0.8
		Raymond	C	0.7	0.7	0.7	0.7	0.7
		Rose	C			0.3	0.4	0.4
		Uchida	P	1.0	1.0	1.0	1.0	1.0
		Student	P	0.5	1.0	1.0	1.0	
		Student	P			0.5	1.0	1.0
		Student	P					0.5
	RAL	Durkin	S			0.3	0.3	0.5
		Harder	S	0.25	0.5	0.5	0.5	0.5
		Manolopoulos	S	0.5	0.5	0.5	0.5	0.5
		Sankey	S	0.3	0.3	0.3	0.3	0.3
		Shepherd-T.	S	0.2	0.2	0.2	0.2	0.2
		Thea	P		0.4	0.4	0.5	0.5
		Tomalin	S	0.7	0.7	0.7	0.7	0.7
		Student	P			0.5	1.0	1.0
		TD Engineer	P	0.2	0.2			
	Total			**8.63**	**11.3**	**12.45**	**13.55**	**12.65**

WP	Institute	Name	Post	19/20	20/21	21/22	22/23	23/24
WP4	Bristol	Cussans	C	0.3	0.3	0.3	0.3	0.3
		Heath	C	0.25	0.25	0.25	0.25	0.25
		Kreczko	P	0.5	0.5	0.5	0.5	0.5
		Newbold	C	0.2	0.2	0.2	0.2	0.2
	Imperial	Borg	O	1.0	0.5			
		Dauncey	C	0.3	0.3	0.3	0.3	0.3
		Davies	C	0.2	0.2	0.2	0.2	0.2
		Greenwood	P	0.2	0.4	0.4	0.3	0.2
		Auzinger	P	0.5	0.5	0.5	0.5	0.5
		Magnan	C	0.2	0.2	0.2	0.2	0.2
		Mastrolenzo	O/P	1.0	1.0	1.0	1.0	1.0
		Palladino	O/P	0.4	0.7	0.8	0.9	0.9
		Rose	C	0.1	0.2	0.2	0.2	0.2
		Seez	C	0.6	0.6	0.6	0.6	0.6
		Virdee	C	0.3	0.3	0.3	0.3	0.3
		Zenz	C	0.2	0.2	0.2	0.2	0.2
	RAL	Cockerill	S	1.0	1.0	0.5	0.5	0.5
		Durkin	S			0.3	0.3	0.3
		Petyt	S	0.6	0.6	0.5	0.5	0.5
		Student	P	0.5	1.0	1.0	1.0	
	Total			**8.35**	**8.95**	**8.25**	**8.25**	**7.15**
WP5	Bristol	Brooke	C	0.5	0.5	0.5	0.5	0.5
		Clement	C	0.2	0.2	0.2	0.2	0.2
		Paramesvaran	C	0.2	0.2	0.2	0.2	0.2
		Penning	C	0.2	0.2	0.2	0.2	0.2
		Student	P				0.5	1.0
	Imperial	Bundock	P	0.5	0.8	0.9	0.9	0.9
		Iles	C	0.3	0.1	0.2	0.35	0.5
		Rose	C		0.5	0.3	0.3	0.3
		Tapper	C	0.2	0.2	0.2	0.2	0.2
	RAL	Harder	S		0.25	0.25	0.25	0.25
		Harper	S	0.3	0.3	0.3	0.3	0.3
		Manolopoulos	S	0.25	0.25	0.25	0.25	0.25
		Olaiya	S	0.5	0.5	0.5	0.5	0.5
		Shepherd-T.	S	0.2	0.2	0.2	0.2	0.2
		Thea	P	0.5	0.5	0.5	0.4	0.4
		Williams	S	0.5	0.5	0.4	0.4	0.4
	Total			**4.35**	**5.2**	**5.1**	**5.65**	**6.3**

Table 2: Project staff profile (FTE)

WP	Description	19/20	20/21	21/22	22/23	23/24	24–	Project	Other
WP1	Common Fund						1.0		1.0
WP2									
WP3	Off-detector electronics			0.07	1.18	1.75		3.0	
	HGCal common items	0.33	0.33	0.34				0.5	0.5
WP4	HGCal off-det electronics			0.05	0.95	1.00		2.0	
	EB optoelectronics			0.07	0.18	0.25		0.5	
WP5	Trigger electronics				0.07	0.43		0.5	
Total		0.33	0.33	0.53	2.13	2.68	1.00	6.50	1.50

Table 3: CORE costs (MCHF)

Ref.	Risk Description	Potential impact on project	Owner	Inherent Risk			Existing Controls	Current/Proposed mitigation	Residual Risk			Risk Exposure		Notes
				Likelihood	Impact	Total			Likelihood	Impact	Total	description	Cost (£k)	
R1.1	Loss of key personnel	Delay to deliverables	PI	0.5	50	25	Succession planning; formal and reviewed project documentation	Accept and plan for risk; schedule contingency	0.5	25	12.5	Delay to deliverables		
R1.2	Slippage of the overall LHC / CMS schedule	Additional costs for support of deliverables	PI	0.25	25	6.25		Accept and plan for risk	0.25	25	6.25	Delay before commissioning phase	500	1/2 cost of 1 year extension to all new posts
R1.3	Large fluctations in costs or exchange rates	Additional hardware costs	PI	0.25	25	6.25	Tracking of cost evolution; flexibility in capital spend schedule	Risk transferred to CMS experiment (with larger capital flexibility); eventually, descope	0.1	25	2.5	Overall risk sits with CMS experiment		
R2.1	Interposer technology not viable	More PCB design work to adapt baseline to WP3/WP4/WP5	PI & WP2 mgrs	0.2	40	8		Test board to allow verification of key aspects of the interposer design	0.2	40	8	Additional effort by designers	20-50	
R2.2	Prototype board failure due to PCB design error	Additional iteration needed, associated delays	PI & WP2 mgrs	0.3	60	18	Layout reviewed by several project staff before production	Layout review, rigorous testing of prototypes before next generation is released for production	0.2	60	12	Additional effort by hardware designers, additional cost for prototype production	80-100	
R2.3	Power dissipation of future FPGAs not manageable with standard air cooling	Boards cannot be used unless improved cooling systems applied	PI & WP2 mgrs	0.3	90	27		improved cooling (embedded fans, liquid cooling) with large safety margin being tested; test board should allow validation	0.2	90	18	Additional engineering effort and increased cost	100-200	
R2.4	Failure or delay at manufacturing company	3-9 month delay	PI & WP2 mgrs	0.7	80	56		Extended quality assurance process; appropriate controls at tender time	0.2	80	16	Schedule delay and cost	40-100	
R2.5	Lack of part availability causes design change	3-6 month delay	PI & WP2 mgrs	0.4	80	32		Close contact with suppliers of key components	0.2	80	16	Delay and cost increase	80-100	
R2.6	Design spec change from external source	Additional iteration and delays	PI & WP2 mgrs	0.3	90	27		Build flexible design	0.1	90	9	Additional design effort, prototype cost	80-100	
R3.1	CBC ASIC: Loss of access to preferred foundry	CBC ASIC: Loss of access to preferred foundry	PI & WP3 mgrs	0.1	90	9	New foundry owners intend to preserve access. Remain vigilant to commercial changes. CERN manages access.	Would affect many projects, so negotiate with CERN and CMS	0.1	90	9	Up to 2FTE electronic engineer effort	200	
R3.2	CBC ASIC: Technical problems with 130 nm CMOS process (e.g. failure of run)	Delay to construction project	PI & WP3 mgrs	0.2	50	10	Continue careful and detailed scrutiny of ASIC and module performance.	Design is well advanced with two successful submissions, with no problems to date or experienced elsewhere.	0.2	50	10	Effort required to diagnose origin. 0.5 FTE.	50	
R3.3	CBC ASIC: Access to chosen radiation tolerant technology affected by government decisions	Delay to construction project	PI & WP3 mgrs	0.05	90	4.5	Remain vigilant to any changes in situation. Remain abreast of regulatory requirements. Design has reasonable degree of technology independence.	Negotiate with CERN and CMS.	0.05	90	4.5	Requires switch to alternative foundry and design changes.	200	

Table 4: Project risk register

Ref.	Risk Description	Potential impact on project	Owner	Inherent Risk			Existing Controls	Current/Proposed mitigation	Residual Risk			Risk Exposure		Notes
				Likelihood	Impact	Total			Likelihood	Impact	Total	description	Cost (£k)	
R3.4	CBC ASIC: ASIC design failure (significant flaw in design)	1 year delay of schedule, cost of additional submissions.	PI & WP3 mgrs	0.1	90	9	CBC design well advanced and on final iteration. Design reviews undertaken at all stages in accordance with ISO9001.	Thorough evaluation of ASICs delivered.	0.1	90	9	Up to 2 FTE electronic engineer effort	200	
R3.5	Off-detector system: Difficulties in agreeing on track finder architecture with collaborators	Delay but during a period when time is available.	PI & WP3 mgrs	0.5	20	10	Designs are undertaken with a well defined oversight structure in CMS and consensus should be found.	UK FPGA-based design is flexible and therefore adaptable.	0.5	20	10	Adaptation of existing hardware to meet new requirements.	50	
R3.6	Off-detector system: Difficulties developing common hardware design	Delays in defining system and extra costs incurred.	PI & WP3 mgrs	0.7	30	21	Most likely could arise from differences of opinion on technical details, or definition of interfaces or task sharing.	Mainly requires patience and flexibility, but also good technical expertise.	0.7	30	21	Extra time and effort from our engineers. 1 FTE.	100	
R3.7	Off-detector system: Technical difficulties in manufacturing processing boards.	Delay to the project but mainly in design stage.	PI & WP3 mgrs	0.3	50	15	Patience and forward planning.	Prototype further test boards. Work with manufacturers to solve problems	0.3	50	15	1 FTE	100	
R4.1	Delays in delivery of common technology platform	Corresponding delay in production of HGCal trigger boards	PI & WP4 mgrs	0.4	60	24		Allow sufficient installation and commissioning contingency. Make HGCal board as similar as possible to common board	0.2	60	12	Delay to HGCal-specific board design		
R4.2	Design error in prototype HGCal-specific boards	3-9 month delay	PI & WP4 mgrs	0.4	80	32	Layout reviewed by several project staff before production	Layout review. Aim to minimise design differences from common technology platform	0.2	80	16	Delay to HGCal-specific board delivery	40-50	Purchase of more common technology boards as interim measure
R4.3	Delays in delivery of US barrel calorimeter boards	Corresponding delay in ECAL trigger	PI & WP4 mgrs	0.7	60	42	Monitor progress of US groups closely	Continue firmware and software development on test boards during any delay	0.3	60	18	Delay to ECAL trigger delivery	20-30	Purchase of FPGA development boards
R4.4	Algorithms too complex for available FPGA resources	More algorithm studies and/or reduced trigger performance	PI & WP4 mgrs	0.3	80	24	Thorough studies of performance for many algorithms of varying complexity	Early simulation in VHDL and testing of critical parts of algorithms on commercial test/development FPGA boards	0.2	80	16	Delay to trigger delivery	60-80	Extra RA effort to perform further simulations studies
R5.1	Delay in delivery of common technology platform	Corresponding delay in production of trigger boards	PI & WP5 mgrs	0.3	50	15	WP2 progress monitoring	Schedule contingency	0.3	50	15	Delay to L1 board design		
R5.2	Change in algorithm requirements from external source	Delays due to firmware design iteration	PI & WP5 mgrs	0.6	60	36	Continuous review of algorithm performance	Schedule contingency for iterations on firmware	0.3	60	18	Delay		
R5.3	Algorithms cannot be implemented in FPGA	Reduced performance, delays	PI & WP5 mgrs	0.5	80	40	Design for overcapacity, plan for reduced performance	Use higher performance FPGA at higher cost, or accept reduced performance. Schedule contingency for firmware re-design	0.25	80	20	Delay, extra effort for firmware design 1 FTE	100	
R5.4	Board cost greater than expected	Increased cost	PI & WP5 mgrs	0.5	80	40		Consider lower cost FPGA with potentially reduced performance.	0.25	80	20	Delay, extra effort for firmware design 1 FTE	100	
R5.5	Common technology platform needs significant extension for trigger requirements	Delays due to design iteration	PI & WP5 mgrs	0.3	70	21		Schedule contingency for board re-design	0.3	35	10.5	Extra time for board design, 1 FTE	100	

Table 5: Project risk register (continued)