Status of the LHCb upgrade

CERN-RRB-2015-041

29 April 2015

1 Introduction

The LHCb Upgrade will be installed during the 18 month Long Shutdown 2 of the LHC (LS2). The upgraded detector will be able to read out all sub-detectors at 40 MHz and to select physics events of interest by means of a pure software trigger at the bunch crossing rate of the LHC. This capability will allow the experiment to collect data with high efficiency at a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Flavour-physics measurements will be performed with much higher precision than is possible with the current detector, and across a wider range of observables. The flexibility inherent in the new trigger scheme will also allow the experiment to diversify its physics programme into important areas beyond flavour.

The Upgrade was proposed in the Letter of Intent [1] in 2011, and its main components and cost-envelope were defined in the Framework TDR [2] one year later. Technical Design Reports (TDRs) have been written for all sub-systems [3–6] and approved by the Reseach Board. Addenda to the Memorandum of Understanding (MoU) were presented to the RRB in April and October 2014, covering the division of resources and responsibilities for Common Project items [7] and sub-system items [8], respectively.

In this report we give a brief update on the status of the Upgrade, reiterating the detector choices made in the TDRs and summarising recent progress. Information is also given concerning overall project organisation, infrastructure and funding.

2 Tracking system upgrade

After the approval of the tracking detector TDRs [3,5], the Vertex Locator (VELO), Upstream Tracker (UT) and Scintillating-Fibre (SciFi) Tracker projects are now finalising their R&D activities and preparing for the construction phase. Project organisation structures and detailed schedules with milestones have been devised in order that the construction activities can be closely monitored. The first engineering design reviews (EDRs) are already planned for the middle of this year.

A summary of recent progress and expectations for the next half year is given for each of the three sub-detectors of the tracking system.

2.1 Vertex Locator (VELO)

Work has been progressing on the VELO in all areas of the project, focussing particularly on the tile prototypes, the module design and the mechanical integration.

Hybrid pixel assemblies have been constructed with the Timepix3 ASIC, pre–cursor of the VeloPix, which is the ASIC intended for the final detector. These assemblies form the central component of the high speed Timepix3 telescope, which has already been used to validate this ASIC for particle tracking, and study the performance of other devices. During a testbeam campaign in the second half of last year the telescope was used to evaluate Timepix3 assemblies constructed with 200 μm thick n-in-p prototype sensors. The efficiency and resolution was probed at multiple grazing angles, and the tests repeated with irradiated sensors. The telescope itself was successfully tested at high data-taking rates.

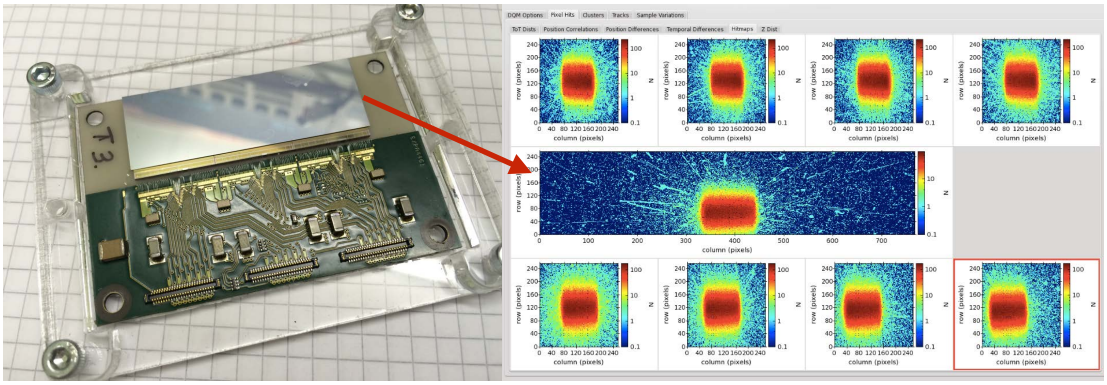

Figure 1: The left photograph shows the prototype Timepix3 quarter module equipped with a triple ASIC sensor. This was successfully operated in the testbeam; the right figure shows the online beamspot monitoring of the 8 telescope planes (squares) together with the test sensor in the telescope centre.

Initial attempts at bonding thinned ASICs led to problems of deformation, which arose from temperature dependent effects associated with the metal stacks in the 130 nm technology. A parallel line of investigation has been launched for thinned ASIC assembly, with the baseline solution being the use of glass-handle wafers during the bonding process, and first results are expected in Q2 2015. The next round of sensor evaluation will include devices from more than one potential vendor, and parylene coating will be added to all n–in–p assemblies to allow tests at up to 1000V in the testbeam environment. The tests will provide input to the sensor review in Q2 2015.

There has been much progress on the design of the VeloPix itself. Following successful irradiation tests at CERN, it has been decided to adopt the 130 nm

TSMC technology, and the high density logic-library has now been fully prepared, with most building blocks now conforming to this choice. Seven designers are working in parallel, aiming for an ASIC EDR in early summer. It is then planned to devote the following two months to an extensive top-level integration and verification period before the final submission. The designers are also working in close collaboration with the hybrid designers in order to optimise the layout of the dense I/O region. In parallel the GWT prototype chip has been submitted separately, and was tested for performance and single-event upset at the end of 2014. The ASIC performance has been re-simulated using as input the most recent Upgrade Monte Carlo data, building on the improved understanding acquired since the TDR. These tests have shown only small differences with respect to previous results, and confirm that the bandwidth is sufficient for comfortable operation at design luminosity, even in the hottest regions. The ASIC test systems have typically a long development track, and extensive preparations are in place for the anticipated VeloPix test programme, followed by incorporation into the official LHCb DAQ slice.

As regards the readout chain, there is steady progress on the continuing integration of VELO firmware code into the LHCb framework. Attention is focussed on evaluating the impact of the new firmware architecture with two data paths, which changes VELO resource usage, and a direct ASIC emulation using the latest simulation data in order to give more realistic GWT data and rate information. The new, thinner (100 μm kapton layer) data tape prototypes have been received and are under evaluation.

A joint workshop with LHC representatives was recently held in order to address mechanical issues, and in particular the design of the secondary vacuum foil. The design of the foil with the rotated modules has been implemented in CAD, including the wake field suppressor region, and finite-element analysis shows excellent mechanical performance. The baseline approach for the manufacturing technique is to sculpt a solid block on the DMF260 5–axis milling machine. The schedule foresees an initial manufacture of a half size box with 500 μm foil thickness, followed by a half size box with 250 μm thickness, followed by a full 500 μm prototype by the end of the year, at which time a foil EDR will be scheduled. The alternative option of chemical etching at the end of the procedure for further thinning is still pursued, with a final decision expected in Q2 2016. In parallel there has been significant progress on estimating the impedance of the new box design, with a full implementation of the structure inside CST Microwave Studio, and quantitative results are expected soon. Vacuum simulations have been performed in Molflow, showing that the effective decreased width of the face section of the rotated foil has decreased the conductance limitation, leading to a potentially improved vacuum performance.

A recent meeting was convened to assess the intermediate stage of module design. Two options are being developed, which differ in the way the microchannel substrates are supported and constrained in the z dimension. Significant progress has been reported on the design and simulation of both prototypes, and the process

steps for the manufacture have been defined, and gluing jigs manufactured. A test setup exists to enable the measurement with LVDT sensors of module displacements after temperature variations in vacuum, and the initial detailed analysis of the first test results is underway. The cooling substrate design is being modified following pressure tests on the connector region, and will be finalised together with the hybrid layout, in order to be able to deliver cooling to both the tile and GBT region in the most efficient way. The final sequence of tests that are needed in order to converge on a design choice has been defined, and a target date for the end of May has been set for this decision.

2.2 Upstream Tracker (UT)

The UT group has recently reviewed the findings of the early R&D effort and external developments associated with ASIC manufacture, to devise a new schedule and set of milestones, updated to that presented in the TDR [5], which is well-tuned to the challenges ahead, and the requirement to be ready for installation in Q1 of 2019.

Test beam studies that took place in the second half of 2014 addressed the key aspects of the silicon sensor technology adopted for the project. Sensors irradiated to different fluences up to about half of the maximum value foreseen were operated successfully in extended test-beam studies. A paper summarizing the results of these tests is currently being finalized. A second stage of R&D, encompassing prototypes featuring all the unique aspects of the UT design, is underway. Prototypes satisfying the design requirements have been produced by the vendor and delivered to the relevant UT institutes. They will be irradiated and then evaluated in test beams and the lab.

The project schedule is dictated by the design and development of the UT front-end ASIC, named the SALT128. This ASIC features 128 channels incorporating a low-noise, fast pulse-shaping analogue processor, followed by a digitisation stage, and a complex digital block performing common-mode suppression, zero suppression, data formatting and serialisation. Early R&D prototyping of the analogue front-end circuit and of the ADC block demonstrated the soundness of the design. A small delay of two months in the original schedule has been introduced by the transition to a different foundry (TSMC as opposed to IBM), but there is confidence that this delay will not impact upon the timely completion of the project. An eight-channel 'SALT8' prototype was submitted in mid February. On receiving this prototype a qualification process will begin that will not only validate the chip, but also the overall structure of the front-end-electronics.

The SALT128 chip will be mounted on hybrid circuits that, together with the silicon sensors, will be assembled on an L-shaped low-mass ceramic substrate. This assembly constitutes the 'UT module'. The design of the UT module is advancing rapidly. Reliability features such as the ability to replace modules on instrumented staves have been implemented. Mechanical assemblies modelling the properties of the silicon-hybrid modules are currently being evaluated in mechanical test-stands.

Plans to extend this work to electrical assemblies are under way. The hybrid circuits are wire-bonded to flex cables that provide connectivity with the near-detector electronics and distribute low-voltage and high-voltage to the modules. The various components of the electronics processing chain were the subject of a rigorous review that occurred in October 2014. The referees noted the impressive progress achieved during the previous year of R&D an gave valuable advice for the final phase of development prior to construction readiness.

The flex cable is a critical component in the overall design. The first iteration of prototyping was only partly successful because of the complexity of the cable and the challenge of finding a company able to provide high quality items with the required features. A new design is now being pursued in collaboration with CERN, which exploits the lessons learned in the first R&D round. Actions are being taken to mitigate the delays accumulated in the initial phase of this progamme.

Considerable progress has been achieved in the mechanical design challenges of the project. Mock-ups of staves with different routing of the Ti cooling tubes have been assembled and characterised. The process of module replacement has been perfected, and the radiation resilience of the different components used in the assembly has been established. Procedures to ensure a reliable and leak-proof connection between the Ti and stainless steel tubes completing the circuit in the CO_2 cooling plant have been established. Progress in the overall CO_2 cooling-plant design has been achieved, where a joint system with the VELO sub-detector is planned.

A set of major EDRs is planned for June 2015. The main goal is to establish the soundness of the 'bare stave' design, thus enabling the start of construction as planned. In addition the silicon sensors, SALT chip, flex cable and overall UT electronics slice design will be discussed.

2.3 Scintillating-Fibre Tracker (SciFi)

The technology and the full detector design of the SciFi system is described in the LHCb Tracker Upgrade TDR [5]. The location of the SciFi is shown in Fig. 2.

The SciFi will consist of 250 μm thick and 2.5 m long scintillating fibres placed next to each other and forming 135 mm wide mats of fibre layers. Eight of these mats will be joined together to form 5 m long and 50 cm wide modules. The fibres will be read out with multi-channel Silicon Photo-multipliers (SiPMs). To ensure that the light yield remains good, even after irradiation, it has recently been decided to build six-layer fibre mats for all detector modules.

In cooperation with the industrial manufacturer, the SciFi group has worked on the production quality of the fibres. Issues are on one hand the light attenuation for which now attenuation lengths above 3 m are constantly achieved, and on the other hand, the uniformity of the fibre diameter. In parallel, the group has performed low-dose irradiation tests for the most recent fibre samples, validating earlier assumptions. The fibre procurement process has been started by preparing a market survey via CERN.

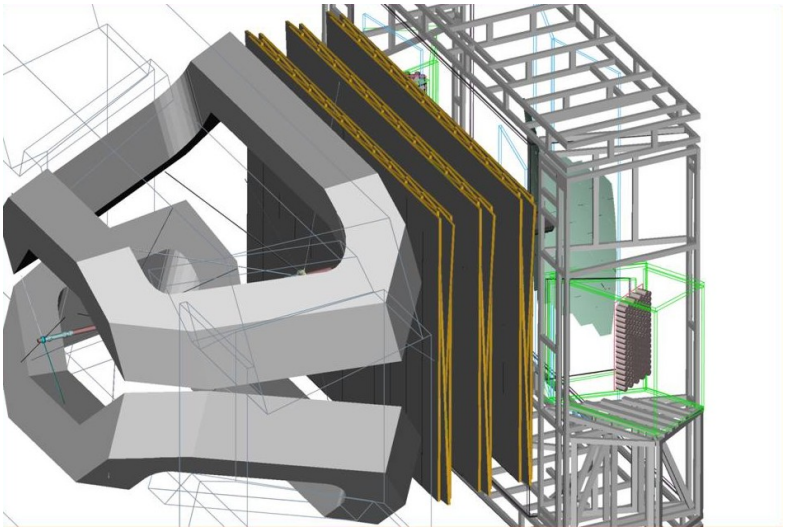

Figure 2: The three stations of the scintillating-fibre tracker shown between the dipole magnet on the left and RICH 2 on the right.

Two different silicon photomultiplier (SiPM) from different industrial suppliers are under investigation. One of these candidates already largely fulfills the SciFi requirements. A final iteration to improve further the electrical properties and the geometrical layout of this SiPM is underway. The performance of the SiPM from the alternative supplier has also improved significantly. The choice between these two devices will be made in early 2016.

Significant progress has been made in the fabrication of the fiber mats and in the module design. The SciFi group is currently building a full-size prototype module in preparation for the EDR of the fibre mats and the modules in July 2015. The construction of this prototype requires tools for fibre winding, for glue casting and for module assembly. The preparation of all tools, in particular the first production-type machine, which is being developed in collaboration with a commercial company, is well on track for the EDR. In parallel to the EDR preparation the envisaged production sites are preparing the infrastructure to start fibre mat and module production which will begin no later than January 2016.

An important component of the future SciFi detector is the Read-Out Box (ROB). The ROB is mounted at each end of the fibre module and houses the SiPMs as well as the necessary readout electronics. It is subdivided into a cold part which allows the cooling of the SiPMs, and into a part containing the front-end electronics. A first version of the complete readout box is currently being built. The challenging ROB design uses modern titanium 3D printing to produce the precise cooling bar.

For the readout of the 600k SiPM channels the SciFi group is developing a fast 64-channel ASIC (PACIFIC) which provides two-bit pulse-height information for every hit by using three adjustable thresholds. A first prototype of this chip, com-

prising the full readout for eight channels, has been successfully tested. Currently the full 64-channel version is being prepared. The design has successfully passed an internal review. The submission of this version of the ASIC is foreseen for late spring 2015.

In autumn 2014, several full size fibre mats, read out by the last generation of SiPM arrays, were tested in the SPS H8 beam. Valuable information on attenuation length, uniformity, hit efficiency and resolution was collected, even if the limited availability of a beam telescope compromised the precision of the results. Two more test beam campaign are scheduled for spring and autumn 2015.

3 Track reconstruction

The proof-of-concept of the pattern recognition algorithms for the LHCb Upgrade was demonstrated in the VELO and Tracker TDRs [3, 5]. In the post-TDR phase the main focus of the tracking software efforts is to implement in detail the geometry, the digitisation and the pattern recognition according to the design choice of the hardware, which is now being finalised.

3.1 Optimisation of SciFi tracker geometry

The focus of the pattern recognition efforts in the last six months has been to implement a more realistic and flexible description of the SciFi geometry and response, and to evaluate potential modifications of the detector design to obtain a more robust tracking performance.

The single hit inefficiencies in the SciFi detector are mainly determined by the photon detection and clustering inefficiency (1%), the size of the gaps between the modules (1% of the active area), and most importantly by the size of the gaps between the SiPM arrays (2% of the active area). These numbers are still subject to evolution, and are dependent on the progress in the detector design and the test of the SiPM arrays. A detailed standalone simulation of light propagation and clustering effects has been developed and will soon be validated with test beam data. The results will then serve as input to the full detector simulation. A hit inefficiency of 1% can be directly translated in a loss of 2-3% in tracking efficiency, depending on the algorithm used. The numbers for single hit inefficiencies that correspond to current best knowledge are now properly taken into account in the tracking performance studies. This result in a 1-2% drop in efficiencies compared to the numbers in the TDR.

The only way to increase the pattern recognition efficiency with a fixed detector design is to loosen the minimum hit requirement in the algorithms, which leads to a significantly increased ghost rate, especially in the inner modules, which suffer from a high occupancy. However, a study implementing additional y-segmentation in the inner region, and adapting the pattern recognition algorithms, showed a significant gain in efficiency (10-15% for low momentum tracks, 3-4% for high

momentum tracks) with similar or even lower ghost rates. The realisation of this segmentation in hardware is quite challenging. Therefore, in parallel, an effort has started to simulate additional detector planes in the inner region, which would be technically easier to be installed. These studies will conclude within the next few months.

4 Particle Identification

The Particle Identification (PID) system of the upgraded LHCb consists of the RICH, Calorimeter and Muon systems. Following the approval of the PID TDR [4] all the sub-systems are progressing rapidly towards the construction phase.

4.1 RICH system

The upgraded RICH system will consist of a re-designed RICH 1 detector, an essentially unchanged RICH 2, and new photodetectors that can be read out at 40 MHz. The key feature of the new RICH 1 is a modification of the optics (and hence also the mechanics of the gas vessel) in order to spread the image compared to the current detector, and thereby reduce occupancy. This re-design can however be performed within the footprint of the current RICH 1, and therefore is compatible with the existing magnetic shielding box, and the evolving plans of the VELO and the UT.

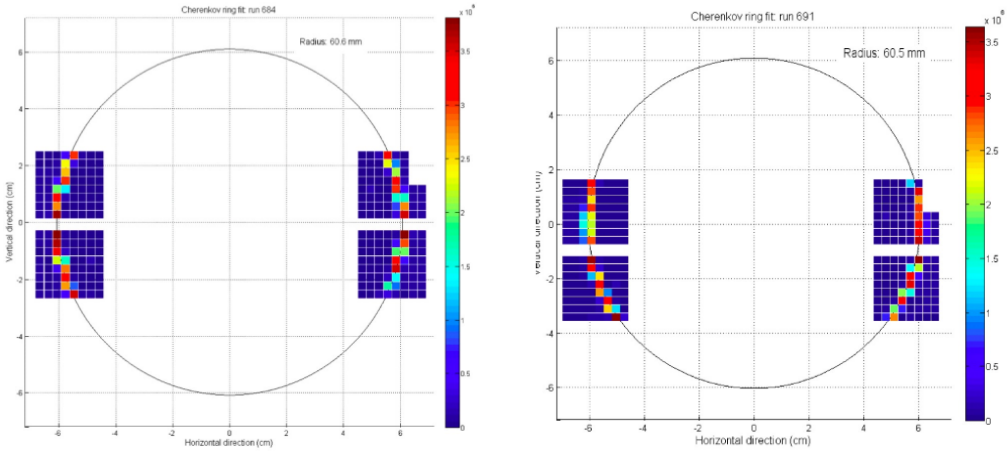

Figure 3: Measured Cherenkov rings with the RICH Upgrade prototype opto-electronic chain. Each blue square represents a MaPMT with its pixels. Colors represent measured Cherenkov photons. The circles are a fit and coincide perfectly with the expected Cherenkov rings.

Simulation indicates that the physics performance of the new RICH system at high luminosity will be similar to that achieved with the existing detector in LHC

Run 1. Studies are focussing on finalising the optical and mechanical concept of RICH 1, which is now entering the phase of detailed engineering designs. The MaPMT is established as the baseline technology for the RICH photon detector. It will be read out by a customized ASIC named the CLARO, which has been produced in its versions 0, 1 and 1.1 and tested in test-beams and radiation areas. The magnetic field robustness of the MaPMTs and radiation tolerance of the CLARO have been fully characterised. All results to date are satisfactory.

The modularity and mechanical design of the opto-electronic chain, arranged in photo-detector planes, its supporting mechanics and cooling system have been mostly defined and are undergoing tests. Successful EDRs for the CLARO, the Front-End Board and the Elementary Cell (the basic module hosting four MaPMTs with their baseboards and the CLARO boards) were held in October 2014. A prototype of the system, made of several MaPMTs together with the accompanying front-end, digital and acquisition electronics and mechanics, was readied for test beam studies, which were carried out late last years. These test beam studies were a great success, paving the way to the 2015 important milestones and full production of the components in 2016. Cherenkov rings produced and detected by two Elementary Cells are shown in Fig. 3. The photodetector order will be placed within the next few months.

4.2 Calorimeter system

The existing electromagnetic and hadronic calorimeters will remain after LS2, but the Scintillating Pad Detector and Preshower will be removed, as they are considered inessential for the most important calorimeter-based physics topics of the Upgrade, such as radiative-penguin studies, and are no longer required for the trigger.

The gain of the calorimeter PMTs will be reduced by a factorfive in order to keep them operational throughout the high-luminosity Upgrade running. The electronics must compensate for this gain reduction and send the data at 40 MHz to the PC farm. Two solutions have been considered for the analogue electronics of the readout: an ASIC, and a circuit based on discrete components. The ASIC has recently been adopted as the chosen solution, after considering factors such as intrinsic performance, ease of implementation and costs. The ASIC EDR is scheduled for May this year. It is anticipated that few if any modifications will be needed after this review, and that the present prototype is very close to our final design. A test bench is planned for June 2015 with the ASIC soldered on the current prototype of the front-end board (FEB). A new version of the FEB, complete with full functionality, will be ready for autumn 2015. It is then planned to have another test beam in November 2015 and/or April 2016. The results of the tests will be used for the EDR of the FEB and the production readiness review (PRR) of the ASIC, in late 2015. The control board design is progressing well, its EDR being scheduled forearly next year.

There has also been significant progress in developing the cluster finding and

reconstruction software for the Upgrade. The energy reconstruction has been modified to form clusters from 2×2 cells, an approach that is less sensitive to pile-up than the current strategy. The reconstruction of the cluster position has also been adapted in order to reduce the effect of the simultaneous collisions at the interaction point.

The ECAL inner modules will suffer from radiation but should remain operational and in specification until LS3. It is planned to replace them during this shutdown by spare modules that are already available.

4.3 Muon system

Over the last six months there has been significant progress in all activities associated with the upgrade of the muon system.

The material required for the construction of spare MWPCs has been received and the production site in Frascati has begun chamber production, as shown in Fig. 4. The PNPI production site is ready and waiting for the delivery of material. The design of the nSYNC ASIC has started in Cagliari using the UMC 130 nm technology: the TDC block was completed in January and was included in a test ASIC that was submitted in February and is expected to be ready for testing in July. In February the LHCb Technical Board decided it is safe to discard the option of a hardware muon low-level trigger (LLT), a backup solution for the fully software LHCb upgrade readout. This decision simplifies the muon readout and allows a more seamless integration of any future higher luminosity design changes of the muon system As a consequence the architecture of the nSYNC and of the nODE boards is now fully defined, and the group is now focussed on optimising the data format required for the upgrade phase.

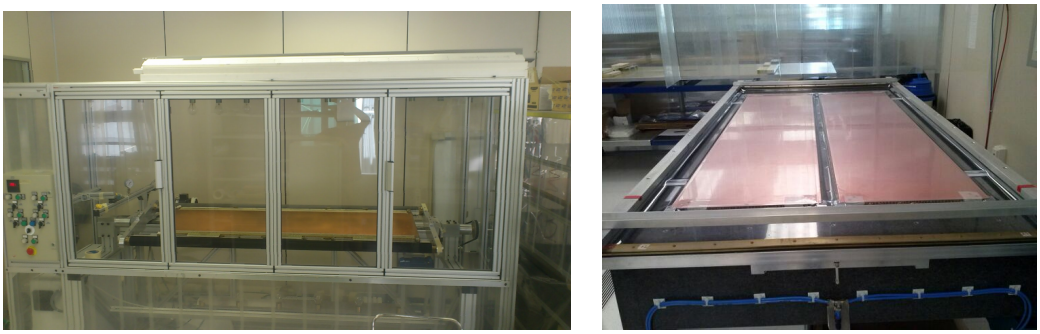

Figure 4: Muon MWPC production at LNF, Frascati. Left: MWPC cathodes ready to be wired. Right: honeycomb cathode panels on the glueing table.

Studies to evaluate and optimise the muon system performance in Upgrade conditions are continuing and promising new muon-identification algorithms are being defined. The redesign of the muon detectors located in the most irradiated regions of M2 and M3 is also under consideration, along with studies devoted to

improving the shielding to reduce the background in these inner regions. Some special technical runs will be performed during the 2015 data-taking period to improve understanding of how to optimise the muon system operation at high luminosity.

5 Data processing

'Data processing' is here defined to encompass the transport of the data from the output of the frond-end (FE) electronics through to the offline reconstruction. It includes data acquisition, triggering and computing.

5.1 Data acquisition

The current challenge in the area of readout electronics is to define and optimise the required number of common electronics components and number of the PCIe40 readout boards. The optimisation criteria for the latter problem are the simplicity of the overall system, the FPGA occupancy and running frequency, as well as the total cost.

The prototype PCBs of the readout board were delivered in April 2015 and are shown in Fig. 5. The delay of few months with respect to the initial planning is due to the limitation on the board width, the high current dissipation close to the FPGA and the optimisation of the layout to facilitate mass production by industry. The market survey and tendering procedure are being prepared with the intention of being able to launch the production of the pre-series by the end of this year.

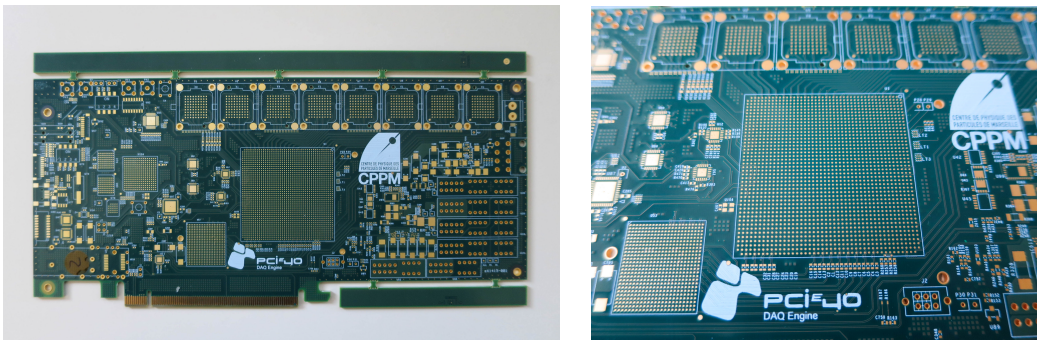

Figure 5: PCBs for prototype PCIe40 readout board.

In the current architecture the front-end electronics and the readout electronics are connected via long distance optical links between the underground area and the surface. A small set of OM3 fibres have already been deployed in the PM 85 shaft. Long-term tests are ongoing to measure the performance of the link in Upgrade conditions.

5.2 Online

The heart of the online system is the event-builder, which assembles the event at a rate of 40 MHz. This system will be based on large bandwidth bi-directional network interconnecting event-builder PC-servers.

Studies are ongoing to select the chassis for the event-builder PC-server. A solution is being sought that provides efficient cooling for the PCIe40 board independently of the CPU load.

Full scale tests of the event-builder are in preparation at INFN-CNAF. About 500 servers will be interconnected by the InfiniBand network running at half of the required bandwidth and the performance evaluated.

5.3 Trigger

The trigger analyses all collisions and selects those to be written for offline storage. The Upgrade trigger consists of a collection of identical software tasks running on the event-filter farm.

The Trigger TDR [6] already laid out an achievable solution for the Upgrade trigger, one that is highly performant and operates within the available timing budget. Further optimisation of this solution is occuring through the commissioning of the high-level software trigger for Run-2 data taking. Here major effort has been invested in optimising the software, which has resulted in a 30% reduction in the number of CPU cycles. The Upgrade trigger will be able to benefit directly from this achievement. It is expected that the experience of Run 2 triggering will be valuable for the Upgrade in many other ways.

5.4 Computing

Twice a year a workshop is organised to discuss all aspects of LHCb computing, including the requirements of the Upgrade. By way of example, out of one of these workshops a small group of experts was convened to evaluate the different techniques for vectorisation, which is an essential skill for Upgrade programming. R&D is also being pursued to determine the benefits of alternative architectures, such as GPUs, on tracking, pattern-recognition and fitting algorithms.

LHCb does not wish to make any offline-computing decisions prematurely, and in any case the Upgrade must learn from the experiences of Run 2. Therefore key dates have been established in this direction and reviewed recently. These include a road-map document to be approved by the end of this year, a computing TDR for the end of 2017, and the finalisation of the computing model during autumn 2018.

6 Preparation for the LHCb upgrade

6.1 Organisation

The Upgrade organisation has been adjusted to meet the requirements of the new, post-TDR phase of activities. The management has created an Upgrade Planning Group (UPG) which meets regularly to review progress. The UPG membership consists of an Upgrade Detector Coordinator, an Upgrade Resources Coordinator, an Upgrade Performance Coordinator and an Upgrade Data Processing Coordinator, as well as the management and a representative of the Physics Coordinator.

Detector upgrade activities are organised within the existing Projects, to ensure efficient sharing of resources between operational needs and Upgrade work. The two exceptions are the UT and SciFi systems, where new Projects have been created.

6.2 Milestones

A list of major milestones for all upgrade projects has been established and is being closely monitored. The first EDRs on this list have already been held, but it will only be possible to assess the true status and progress within the projects after the second quarter of the year, a period during which many key targets should be met. All projects are working to a schedule that assumes the start of the Long Shut Down 2 (LS2) in summer 2018, and with this the beginning of the Upgrade installation. LHCb still considers installing optical fibres for the data transmission in the extended year end technical stops during Run 2.

6.3 LHCb upgrade infrastructure

Infrastructure developments at point 8 for the Upgrade are already underway. The installation of the first optical fibres for test purposes was successfully completed in October 2014. One set of 144 micro cables were blown from the cavern to the surface, close to the future event filter farm. In addition, two sets of trunk cables each with 144 single fibres were installed. First tests of these fibres confirm that the quality is sufficient for the LHCb upgrade, but further evaluation will be required.

A document summarising the requested infrastructure for all sub-systems that will be required for the Upgrade is in preparation. In addition, an Upgrade Infrastructure Workshop was organised in February. Here the sub-systems presented their needs in terms of services, detector integration and assembly. Representatives of the technical departments were invited to discuss constraints and boundary conditions for services and Upgrade installation.

7 Funding

The LHCb collaboration submitted a 'LHCb Upgrade Framework Technical Design Report' (FTDR) [2] in 2012 and then during 2013-14 Technical Design Reports for all the Upgrade subsystems [3–6]. By the end of 2014 all TDRs had been fully approved by the Research Board. In April last year the final version of the Addendum No. 1 to the Memorandum of Understanding (MoU) for Common Projects [7] was submitted for signature to the Funding Agencies, followed in October by an Addendum No. 2 to the MoU for Upgrade of the Sub-Detector Systems [8]. These documents define in detail the technical design and cost of the upgraded detector, as well as the sharing of responsibilities among the institutes and funding agencies in the construction, installation and commissioning of the various sub-systems.

As detailed in the Addenda, the total cost of the LHCb Upgrade of 57.2 MCHF is divided into a Common Project component of 15.7 MCHF and a Sub-Detector System component of 41.5 MCHF. According to the Addendum No. 2 the sharing of responsibilities in the detector upgrade has been fully defined among all participating institutes, and the funding is secured up to a missing 2% of the total cost. This underfunding of 1.2 MCHF, if persisting, would affect the risk reserve and lead to a delay in the purchase of spare parts for the RICH detectors.

As of the beginning of March this year, we had already received positive feedback from a number of Funding Agencies concerning their anticipated commitments to the Upgrade project. These commitments cover at present around 85% of the Common Fund contributions, and around 75% for the core contributions to the various sub-detector components. We are confident that in due course further commitments will follow to ensure a complete and timely installation of the LHCb Upgrade during Long Shutdown 2.

References

[1] LHCb collaboration, *Letter of Intent for the LHCb Upgrade*, CERN-LHCC-2011-001, LHCC-I-018.

[2] LHCb collaboration, *Framework Technical Design Report for the LHCb Upgrade*, CERN-LHCC-2012-007.

[3] LHCb collaboration, *LHCb VELO Upgrade Technical Design Report*, CERN-LHCC-2013-021.

[4] LHCb collaboration, *LHCb Particle Identification Upgrade Technical Design Report*, CERN-LHCC-2013-022.

[5] LHCb collaboration, *LHCb Tracker Upgrade Technical Design Report*, CERN-LHCC-2014-001.

[6] LHCb collaboration, *LHCb Trigger and Online Technical Design Report*, CERN-LHCC-2014-016.

[7] LHCb collaboration, *Addendum No. 01 to the Memorandum of Understanding for Collaboration in the Construction of the LHCb Detector. The Upgrade of the LHCb Detector: Common Project items*, CERN-RRB-2012-119A, revised April 2014.

[8] LHCb collaboration, *Addendum No. 02 to the Memorandum of Understanding for Collaboration in the Construction of the LHCb Detector. The Upgrade of the LHCb Detector: Sub-Detector Systems*, CERN-RRB-2014-105, October 2014.