Status of the LHCb Upgrade

CERN-RRB-2016-111

October 2016

1 Introduction

The LHCb Upgrade will be installed during the two year 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 [8] and sub-system items [9], respectively.

In this report a brief update is given on the status of the Upgrade, reiterating the detector choices made in the TDRs and summarising progress since the previous RRB. No new significant risks have been identified, or concerns arisen, in any of the sub-detector or common projects. Information is also given concerning overall project organisation, infrastructure and funding.

2 Tracking system upgrade

Over the last six months the Vertex Locator (VELO), Upstream Tracker (UT) and Scintillating-Fibre (SciFi) Tracker projects have undergone several important Engineering Design Reviews (EDRs) and Production Readiness Reviews (PRRs). Delivery of key components and the production of detector parts has started. A

summary of recent progress and plans for the next half year are given for each of the three sub-detectors of the tracking system.

2.1 Vertex Locator (VELO)

New VELO sensor prototypes have been received from one of the two vendors currently under consideration and are expected in the coming weeks from the second. The geometry includes the rounded corners and elongated pixels to accommodate the ASIC design and adjustments to the guard ring regions to address the extra charge collection from the edges seen in previous prototypes. The sensors are currently being bump bonded to Timepix3 ASICs for HV and signal tests. A PRR is planned for October 2016.

The VeloPix underwent final simulation verification steps and was submitted at the end of May. In parallel the new chipboard for ASIC testing was produced. The wafers were delivered to CERN at the end of August and were immediately sent for thinning and dicing and then mounted on the test board. The tests so far have indicated excellent ASIC performance in agreement with the simulations. Testing with sensor prototypes and irradiation are planned for the coming months.

The split VeloPix evaluation hybrid ("VeloPix hybrid" and "control" hybrid) system together with the vacuum feedthrough evaluation board, OPB prototype board, and VeloPix specific MiniDAQ system are in the final stages of preparation for a complete readout electronics slice test. This test is important to check the transmission of the complete chain in order to make an early decision on a possible modification to add pre-emphasis to the system.

The data acquisition firmware integration is progressing well, including the VELO specific blocks.

The microchannel testing with the pre-production prototypes is advancing well. A setup with a double cooling plate and complete pipe and valve routing has been developed to test the dynamic performance of the system and to check the pressure flow performance under different powering and temperature scenarios. The connector soldering technique has evolved to a self aligning float solder method, which give reliable results free of voids in vacuum. The microchannel production tender was submitted in May. At the end of the tender no suitable bids had been received and LHCb entered into direct negotiations with potential suppliers. This resulted in the establishment, with one vendor, of a three phase programme for the microchannel order. The initial step will prove the technical feasibility of the final design, most crucially to establish the bonding yield on the 8-inch wafers. A second step will produce an initial batch of cooling plates with the final design, and will be followed by a final phase for the full microchannel cooling-plate order. In order to compensate for the delay in microchannel plate delivery the module assembly and transport procedures are being fully exercised with dummy plates cut from silicon of the correct thickness. In addition, a backup design incorporating steel capillaries embedded in a composite plate is being developed. Full sized prototypes are being produced with the backup design. If the wafer bonding yield

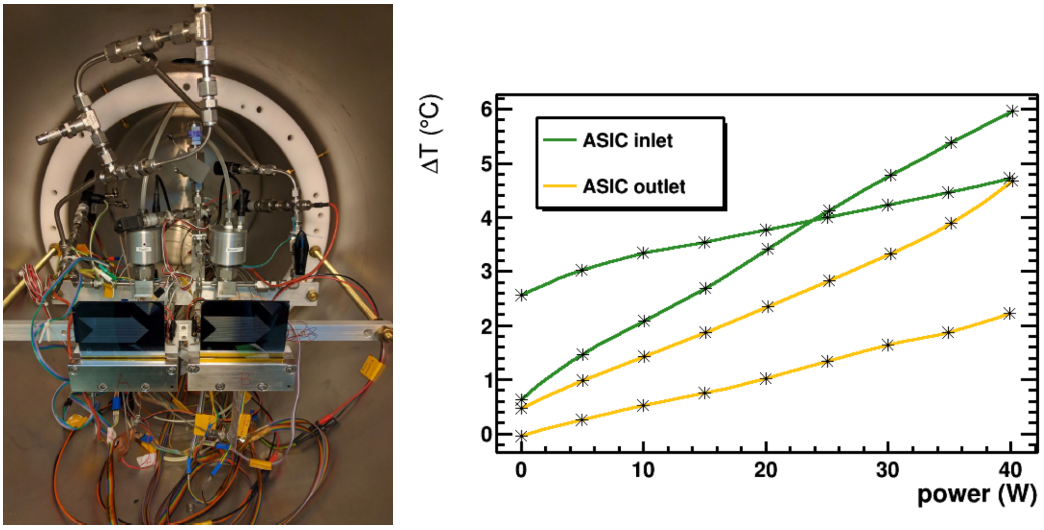

Figure 1: (left) double microchannel prototype cooling setup equipped with heaters to simulate the ASIC heat dissipation; (right) measurement of the difference in temperature between the coolant and the input and output of the prototype ASIC cooling plates for the two snakes. The ΔT is seen to be small for both samples and the cooling performs well beyond the anticipated 30W heat dissipation

is too low for production purposes then the project will switch to adopt the backup solution.

Good progress has been made with the module design. The vacuum tanks and cooling systems have been developed and module deflection measurements in vacuum have been performed. For the mechanics, work has started on the hood design and internal routing of the moveable cooling pipes and electrical services. This has highlighted a number of interface areas where modifications have been necessary to ensure a proper overall optimisation. The major changes include an adjustment to the RF foil to allow greater clearance to the corners of the sensors, the installation of an internal support skeleton for the moving components, and a modification to the vacuum feedthrough to optimise the data signal transmission path. The mechanics EDR is planned for the end of 2016.

A production of the full scale prototype of the RF box is in progress, with the milling of the inside completed. A modified procedure for the wax filling has been adopted, to avoid the effects of the cooling and shrinking wax deforming the thinned foil, as was seen on the previous prototype. New wakefield suppressors have been delivered and are being used to prototype the mushroom attachment on the foil as well as the connection to the exit foil.

2.2 Upstream Tracker (UT)

The UT project is progressing well. The transition toward construction is underway. Several key aspects of the projects have undergone important reviews in June

2016, and we are refining our production plan based on the input received.

The last batch of prototype type-A sensors, which are needed for the outer region of the UT planes, has been studied in the laboratory with lasers and radioactive sources, as well as with extensive test-beam runs. Miniature detectors featuring the two pitch-adapter design were irradiated to different levels of fluence, up to twice that expected at the Upgrade. In addition, full size p-in-n modules were irradiated at the IRRAD facility at CERN in early May and then tested with beams at the end of May. The findings of the extensive test programme were scrutinised in a type-A sensor PRR that took place in June 2016. The UT group is now incorporating the reviewers' suggestions to finalise the pre-production masks.

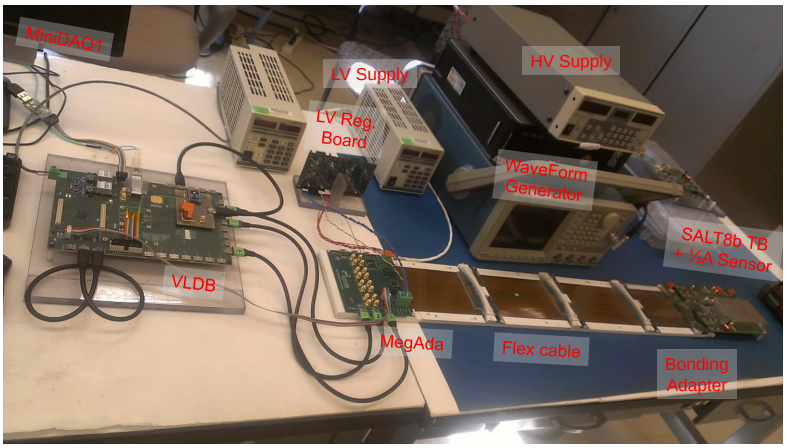

Figure 2: Test stand for the study the performance of the SALT chip.

Considerable progress has been made in the R&D programme of the SALT chip, which is the ASIC that will be used for UT readout. The second eight-channel SALT8b ASIC, which was designed and submitted for fabrication in November 2015, has been thoroughly evaluated. All the ASIC building blocks have been studied in the test stand shown in Fig. 2. In particular, the response to an infrared laser pulse has demonstrated that the behaviour of the analogue block is consistent with expectations from simulation. Figure 3 shows the pulse shapes shown obtained with scans of the digitised information at different sampling times; these curves demonstrate that all the major building blocks of the ASIC work well. In parallel, a first iteration of the full ASIC (SALT128) was submitted in June 2016 and is currently being characterised.

A complete test of an electronics slice is under preparation. This slice will comprise the silicon sensors wire-bonded to SALT8 ASICs, mounted on flex hybrids, which are connected to the recently acquired prototype flex cables that route signal and power to the near detector electronics. The baseline hybrid technology,

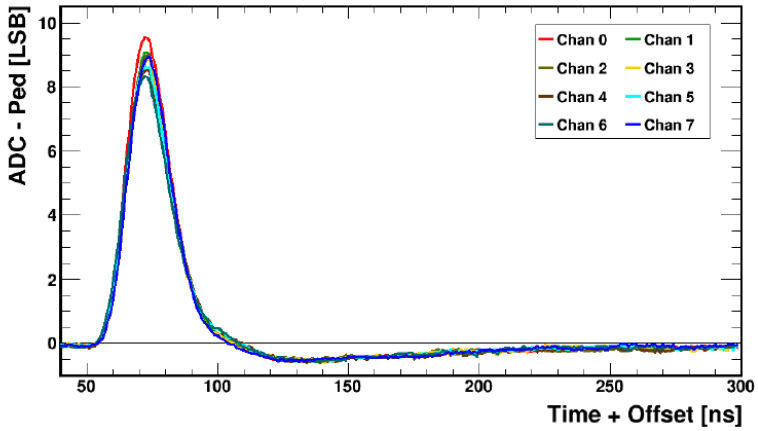

Figure 3: Response to infrared laser pulse from type A sensor wire-bonded to SALT8b prototype. An infrared laser producing light roughly equivalent to 1 MIP is shone on a full size Si microstrip detector in steps corresponding to the location of each connected strip.

initially incorporating SALT8b ASICs, is at an advanced stage of prototyping. In parallel, the design of the SALT128 hybrid is progressing well.

The second round of prototype flex cables is currently being studied, in conjunction with SALT8b test board and the MINI DAQ based data acquisition system used for the slice test. Preliminary tests are very promising, demonstrating robust signal integrity, noise immunity with respect to the I2C slow-control signal, and sound high-voltage and low-voltage distribution design. Based on this experience, the final iteration of flex cable design is in progress.

The peripheral electronics, incorporating GBTx based boards that provide data aggregation and transition to optical transmission and voltage regulator boards, that distribute power to the hybrids, is approaching the production phase. A major electronics review is planned for December, including an EDR of the hybrid circuits and PRRs for flex circuits and low voltage distribution boards.

Significant steps forward have been made in the mechanical design of the detectors. Two bare staves have been constructed and tested. A successful PRR took place in June; the reviewers agreed that the design is sound and provided recommendations on potential improvements of the construction. In addition, developments are ongoing for the CO_2 cooling system. The design of the manifold to distribute the cooling to the various staves is advanced. A recirculating TRACI CO_2 cooling system allows for studies of temperature profiles on mechanical mock-ups.

Finally, there have been achievements in the design of the detector enclosure and mounting frame. The partitioning of the space near the boundary between inner detector volume and near detector electronics, just outside the detector volume, has been optimised. Detailed designs of the mounting frame and detector

enclosure have been examined in an EDR in June and design revisions incorporating the feedback from the referees are being drafted.

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 SciFi will consist of 250 μm thick and 2.5 m long scintillating fibres arranged as hexagonally close-packed six-layer mats of 135 mm width. Eight of these mats will be joined together to form 5 m long and 52 cm wide modules. The fibres will be read out with multi-channel Silicon Photo-multipliers (SiPMs) which have to be cooled down to -40^{o}C to limit the dark count rate after irradiation. The modules will be arranged in 12 stereo layers.

The industrial supplier of the scintillating fibres started the production in May 2016 with a pre-series of 500 km. After verification of the contractual specifications the green light for the production of the remaining 10500 km was given. We have now reached the full delivery rate of 300 km every two weeks. After reception at CERN, these fibres undergo various tests and a geometrical refinement at which large "bumps" (occasional local increases of the fibre diameter) are removed during the fibre testing. The largest fraction of the bumps is treated automatically with a heated shrinkage tool. About 10% of the bumps need to be manually removed.

After successful PRRs, the serial production of the fibre mats has started at two centres (Aachen and EPFL). Figure 4 shows the milling of a fibre-mat endpiece, and an optical image of the fibres after this operation. A third centre (Dortmund) has recently undergone a PRR and will start production in October. The PRR of the last centre (Kurchatov) is foreseen before the end of the year.

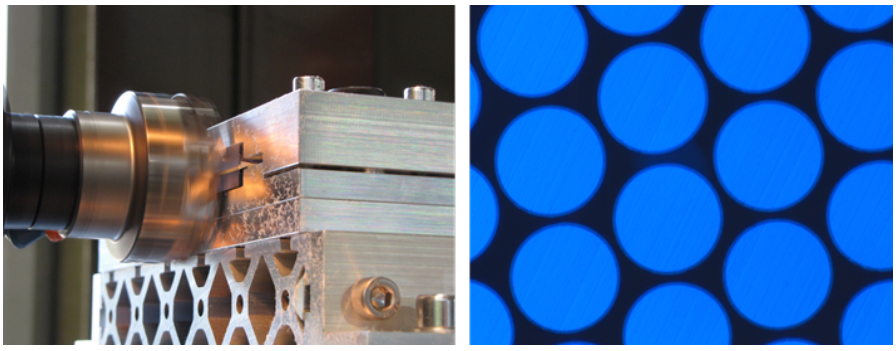

Figure 4: Left: milling of a fibre mat endpice. Right: optical image of the fibre surface after machining.

The PRR of the first module assembly site (Heidelberg) took place in July. Figure 5 shows a module produced at this site. The module production is foreseen to start in October with an initial production speed of 1 module per week which will be increased once all mat winding centres are producing with nominal speed (4

mats per week per centre). A second module assembly site is planned at NIKHEF, for which the PRR is foreseen for the end of the year. According to the production schedule, the module production will continue until spring 2018.

Figure 5: The first full-size scintillating-fibre module being assembled in Heidelberg.

The SciFi group pursued intense R&D on SiPM arrays with two industrial suppliers. The EDR was successfully passed in April this year. The devices from one supplier fully match the requirements of the project. A recent test sample (v2016) of this supplier shows further improved characteristics by combining the benefits of the 2015 model (optimised geometry and a higher photon detection efficiency) with a much lower delayed cross-talk and a lower rate of after-pulses. Recently the call for tender for the procurement of 5500 pieces (plus an option for additional 1000) has been issued.

The 64-channel prototype front-end readout ASIC (PACIFICv3), with full functionality, has been produced and tested. The chip was reviewed in April, and a revised version of the ASIC (PACIFICv4) was submitted to the foundry in June. Unfortunately first tests of the recently received chip revealed a 'trivial' problem in the I2C communication which renders a proper characterisation of the chip impossible. The chip was resubmitted after correcting this bug.

A first version of the subsequent digital readout electronics is available. The readout chain, consisting of a single PACIFIC carrier board, a cluster-board to find the hits, and the master board, comprising the Slow and Fast Control as well as the optical links, has been tested, and the results presented at the electronics EDR. A recommendation of the review committee is to use differential signal transmission between the PACIFIC chip and cluster board. Revised versions of all components taking account of this recommendations are currently under development.

The so-called Read-Out Box (ROB) is mounted at each end of the fibre module and houses the SiPMs in the cold part of the box ("cold-box") and the front-end

electronics. Four prototype cold-boxes, including a vacuum insulated distribution of the cooling liquid (Novec 649), have been built and tested (see Fig. 6). The design successfully passed an EDR in August. The schedule foresees the production of a small prototype series in the course of the next few months.

Figure 6: SciFi engineer (R. Walet from Nikhef) in-front of the cold-box test stand. The distribution lines (not visible) as well as the input and output lines supplying the coolant at -40°C to the four cold-boxes are vacuum insulated.

In November 2015, a full-size fibre mat was irradiated at the CERN PS IRRAD facility over its full length, reproducing the expected dose profile for the lifetime of the detector. The performance of the irradiated mat has been studied in a test-beam using high-energy protons and pions as well as in the laboratory with a beta source. While the test beam measurement performed directly after the irradiation revealed a smaller light yield than expected, the laboratory measurements several weeks later confirmed the photon yield predicted by simulation. The irradiation campaign was repeated with a new fibre mat in July and the light-yield of the module was measured as a function of the time after irradiation. The time-resolved analysis allows two damage components to be clearly identified: a fast-annealing (with a time constant of days) and a permanent component. The latter is as predicted by our simulation. We conclude that the two irradiations (Nov 2015 and July 2016) are consistent and expect the fast-annealing component not to degrade the detector performance when operated at the moderate dose rates expected in LHCb. This assumption will be cross-checked using a fibre mat which has been placed close to the LHC beam pipe and is now being irradiated at the 'natural' dose rate. The light yield of the mat will be measured during the next technical stop.

In summary, the SciFi project proceeds according to schedule: the fibre supply

is at its nominal rate, the mat and module productions are starting, the procurement of the SiPMs has been launched, the development of front-end electronics will soon be concluded and the work on the detector infrastructure is rapidly gaining momentum.

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 & simulation

Significant progress has been made in the detailed design of the SciFi tracker geometry. Constraints from the cryogenic equipment on one side of the detector limit the space for opening the stereo-layers. For maintenance, in particular of the inner modules, it is necessary to fully open the detector to access the modules and the electronics, which would not be possible with all full-width fibre mats assembled. This problem will be overcome by reducing the number of modules in the first two tracking stations by one, yielding an acceptance reduction of 29 mm. The corresponding loss in tracking efficiency is limited (0.3% for standard so-called 'long' signal tracks with higher momentum from B decays, 0.7% for inclusive long tracks). Investigations are ongoing to ascertain if this small acceptance loss can be recovered by alternative design choices.

The shape of the beam-pipe hole of the SciFi tracker has been optimised. To avoid an overly complicated cutting of the fibre mats with round shapes, the same rectangular cut-out for all stations is adopted. The x-width will be 261.6 mm, which is the width of two fibre mats. the y-height will be 230 mm, the diameter of the beam pipe at the outermost tracking station with a margin of 20 mm added. With this design, only two types of beam-pipe modules are needed, one for the x-layers and one for the u and v layers.

Shielding is required to protect the SiPMs that read out the SciFi tracker from the neutron fluence. The proposed location for this shielding is at the z-position of the current M1 station. A scenario with 30 cm of polyethylene in the inner region and 10 cm at the outer region is favoured, and would provide some safety margin for the SiPM operation. Simulation studies are underway to establish the effect of this shielding on the ECAL.

The SciFi track reconstruction performance is found to be stable against the inclusion of noise in the simulation. A workshop has been held to focus efforts on the detector simulation and digitisation. The more realistic simulation of photons in the detector has been successfully implemented in the LHCb software frame-

work. A set of milestones will be defined to monitor the progress in development and implementation of the SciFi detector simulation.

3.2 Timing performance of the tracking algorithms

Work has begun on the optimisation of the timing performance of the track reconstruction algorithms, with the requirement of not losing any track finding efficiency. Already factors of two to three have been achieved in speed with the most complex versions of the algorithms. For optimal use in the trigger sequence, simplified ("fast") tunings of the algorithms are needed. Studies to port the optimisations to the fast algorithm tunings have begun.

A first optimised complete tracking sequence for the first trigger stage has been set up and its performance has been evaluated. This sequence is more realistic than previous versions, *e.g.* it includes noise and the full hit management. Also the primary vertex finding algorithm now contains an internal track fit. A first version of the simplified geometry description for the Kalman-filter based track fit is currently being tested. Most algorithms have so far been tuned with the most complex algorithm versions in mind, so improvements on the execution time are expected. Still, the timing for the reconstruction sequence of the first stage trigger is already comparable to the one outlined in the trigger and online TDR [6].

Further work is ongoing in this direction in view of the studies required for the Software and Computing TDR (Q4, 2017).

4 Particle identification

The Particle Identification (PID) system of the upgraded LHCb detector consists of the RICH, Calorimeter and Muon systems. The design of the main components of the three sub-systems is complete. Several PRRs have taken place since the last report. The EDRs of the remaining components will be completed within the next few months.

4.1 RICH system

The upgraded RICH system will consist of a re-designed RICH1 detector, an essentially unchanged RICH2, and new photo-sensors with new front-end electronics that can be read out at 40 MHz. The key feature of the new RICH1 is a modification of the optics (and hence also some of the mechanics) in order to spread the image compared to the current detector and thereby reduce the occupancy. This re-design can however be performed within the footprint of the current RICH1, and therefore is compatible with the existing magnetic shielding box and the evolving plans of the VELO and the UT. 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 Runs 1 and 2.

The MaPMT is established as the baseline technology for the RICH photon detector, read out by a customized ASIC named CLARO. The order for the MaPMTs was placed in 2015, the pre-series arrived and was accepted in April 2016 after Quality Assurance (QA) tests. The first batch of MaPMTs of the production series arrived in September 2016 and is being qualified, through our two regional centers for MaPMTs QA in Edinburgh and Padova. The MaPMT, CLARO, front-end electronics and system integration have been tested in test-beams and radiation areas. All results to date are satisfactory, PRRs have been carried out successfully during summer 2016 and the tendering process for the electronics components has started.

The photodetector assembly, including the MAPMTs, all on-detector electronics and ancillary systems, is common to RICH1 and RICH2. The modularity and mechanical design of the full opto-electronic chain, its supporting mechanics, cooling and power distribution systems has been defined and final design and tests are underway. EDRs for the RICH1 and RICH2 mechanical systems were held in May 2016.

Prototypes of the system, consisting of several MaPMTs together with the accompanying front-end, digital and acquisition electronics and mechanics (Fig. 7), have been evaluated in a series of test beams in 2014, 2015 and 2016. These test-beam studies were a great success, paving the way towards the important milestones of 2016 and full production of the components.

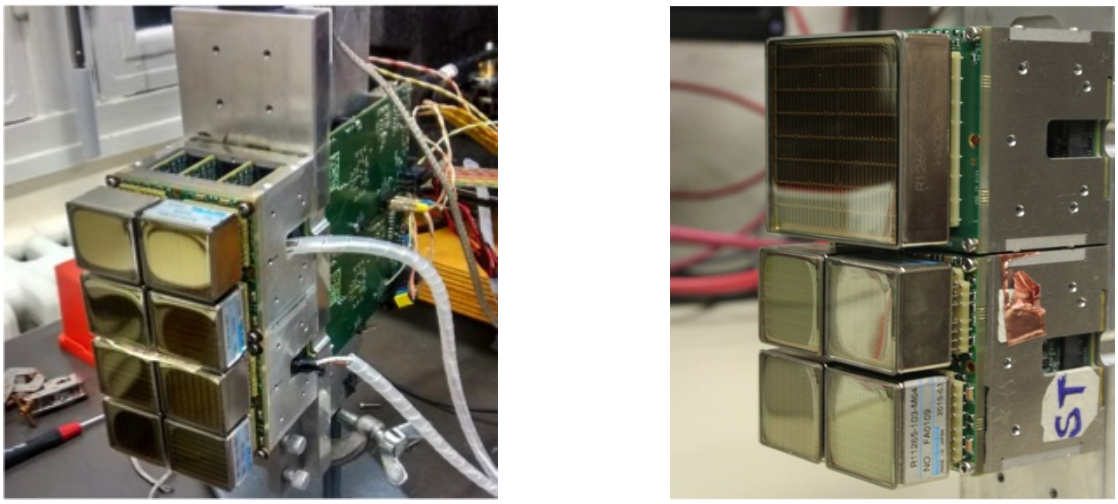

Figure 7: Left: prototype of the photodetector plane of the upgraded RICH used in recent beam tests. Right: prototype where both the 3x3 cm^2 and 6x6 cm^2 devices are mounted. The large MaPMTs are used in the outer region of RICH2.

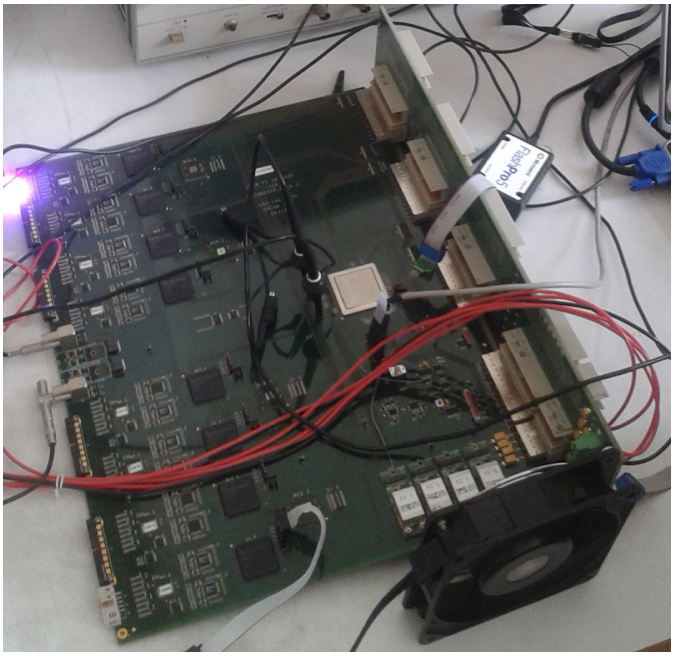

Figure 8: Picture of the latest prototype of the Calorimeter Front-End board under tests.

4.2 Calorimeter system

The existing electromagnetic and hadronic calorimeters will remain in place after the LS2, but the Scintillating Pad Detector (SPD) and Preshower (PS) will be removed, as they are considered inessential for the most important calorimeter-based physics topics of the upgrade and are no longer required for the trigger.

The gain of the calorimeter PMTs will be reduced by a factor five in order to keep them operational throughout the high-luminosity upgrade running and the electronics will compensate for this gain reduction. The upgraded detector will send the full data flow to the counting room at 40 MHz. The present earliest-level trigger calculations performed on the Front-End boards (FEB) will be kept and the result (now an ingredient in the so-called Low Level Trigger) will be sent to the PC farm in order to optimise the software trigger.

The calorimeter upgrade consists in the replacement of the front-end electronics, the upgrade of the high-voltage, monitoring and calibration systems, the dismantling of the SPD/PS, the implementation of the data acquisition firmware in the PCIe40 boards and the improvement of the calorimeter reconstruction in a high-multiplicity environment.

The design of the front–end ASIC, the ICECAL, is now complete. The latest version of ICECAL has been tested, with a FEB prototype, in an electron beam at CERN and it has been irradiated with intense beams of protons and heavy ions at Louvain-La-Neuve. The ASIC PRR will take place mid-October and it will be followed in November by the start of the production.

A new prototype of the Front-End Board (FEB), shown in Fig. 8, based on the final components has been designed and produced and is now under test. The digital components used for the FEB have been irradiated at Louvain-La-Neuve in a proton beam, in order to validate their resistance to the cumulated dose and to the single event effects. A new irradiation test is planned with heavy ions by the end of the year. The firmware for the FPGAs in the FEB is currently being written.

The control board, called the 3CU, is also at the prototyping stage and a new version will be produced by November.

The upgrade of the HV, monitoring and calibration systems is progressing well. Two mezzanine boards equipped with GBT-SCA and a GBT fanout board, which allows an optical link to be connected from the control room to several mezzanines through a GBT-X chip, are under design. Prototypes of the mezzanine boards have already been partially tested. The GBT fanout board should have similar functionalities as those of the VLDB board from CERN, hence its design will be based on the CERN system. The EDR of the FEB, 3CU, HV, monitoring and calibration systems will take place before the end of the year.

The distribution of the optical links on the PCIe40 acquisition boards is now being finalised and the calorimeter specific firmware is in preparation.

First studies for the dismantling of the SPD/PS have started. The tools used for the installation of the detectors have been refurbished and will be reused. It may, however, be necessary to produce duplicates of some of the equipment items in order to fit into the time interval allocated for the task. A draft schedule is being prepared.

4.3 Muon system

Much progress took place in the last six months. The LNF and PNPI construction sites for the spare MWPCs are expected to complete their production and acceptance tests by the end of the year. Some of the new MWPCs for the M5R4 region will be transported to CERN in late November, because of the current unavailability of spares for this region. After the front-end electronics installation and the test with cosmic rays these chambers will be available for installation in the event replacements are already required in this coming extended end-of-year technical stop. The remainder of the spare MWPCs will be transported at CERN in 2017 or 2018.

The nSYNC ASIC, submitted in February 2016, was delivered early September and is currently under test (Fig. 9). First results appear promising. The design of the upgraded Muon readout board that will carry this ASIC, the nODE, has been completed, and the prototype board will be produced after the full validation of the ASIC.

The work on the Muon system specific PCIe40 firmware is continuing. Simulations show that the PCIe40 board should be able to handle 24 optical links in input, but a zero suppression algorithm, currently being implemented, is required

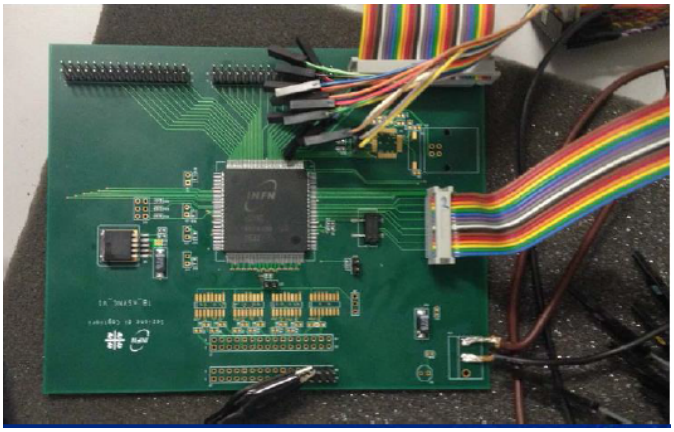

Figure 9: Test board for the muon system nSYNC chip.

because of the limited PCIe bandwidth.

Concerning the new control system of the MWPCs, the first nSB prototype is ready and currently under test. Reliable 80 Mbps communication on the crate backplane point-to-point lines has been demonstrated, as required for the links between the nPDM and the many nSBs populating the Muon system control crates.

The design of a new tungsten M2 shielding-plug was completed before the summer. The design of the other beam plug components is under way, and a global price inquiry will be launched when the final drawings will be available.

Studies to improve the Muon system performance for the upgrade are continuing and the new algorithms are being tested on Run 2 data and Upgrade simulation. They will soon be available for use in Run 2 data analysis

5 Data processing

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

5.1 Data acquisition

In the domain of readout electronics, the current priority is to define and optimise the number of the PCIe40 (PCI Express) readout boards that the experiment requires. The optimisation criteria are the simplicity of the overall system, the FPGA occupancy and running frequency, as well as the total cost. The provisional number of PCIe40 boards for each subdetector was updated in June 2016, with a final decision expected in October 2016.

Prototypes of the PCIe40 readout board were delivered in mid-2015. An extensive campaign of tests were run, which demonstrated the feasibility of a PCI

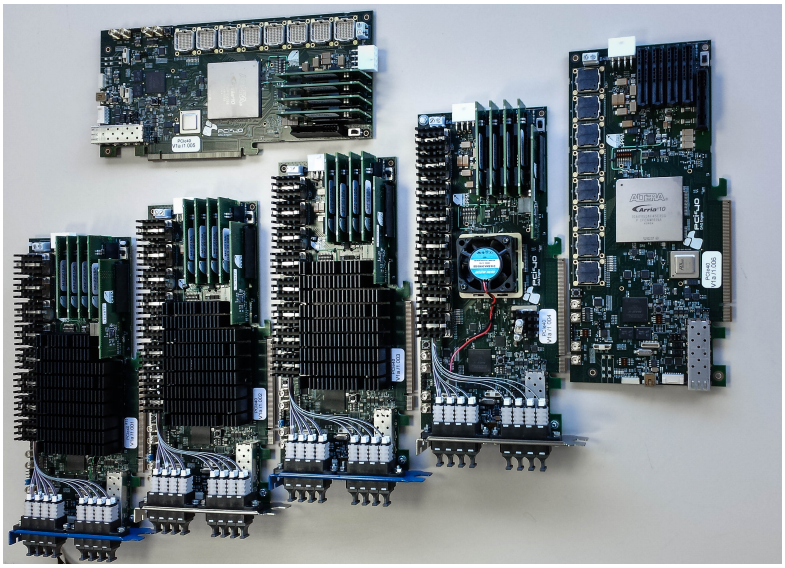

Figure 10: The first six PCIe40 prototypes of 2016, at different stages of their assembly.

Express board handling a total bandwidth of 100 Gbit/s.

Twenty-five copies of the PCIe40 prototype, with the final version of the FPGA Arria 10, have been produced during 2016 (Fig. 10). These boards will be used to equip the MiniDAQ2, which is a PC-server for the testing and validation of the front-end electronics of each sub-detector. The MiniDAQ2 systems will be delivered to the collaboration by November 2016. This date will be eighth months behind the initial schedule, a delay which is caused by the late availability of the Arria 10. However, this new schedule is still compatible with the timelines of the sub-detector projects.

The final version of the PCIe40 will be designed between October 2016 and March 2017. The board will be optimised to ease its manufacturing and to reduce its cost as far as possible. The invitation for tender will be launched by the CERN procurement service when the manufacturing files are ready. The production will start before the end of 2017.

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 a large bandwidth bi-directional network interconnecting event-builder PC-servers.

Software packages have been developed to evaluate the performance of the event building and the transport layers in different configurations, as well as different technologies. Large scale tests have been run at HPC centres, such as CINECA in Italy, in close collaboration with industry. First results are very encouraging. In 2015, an aggregate traffic of 3.2 Tbit/s was measured between 128 servers using

the transport layer Infiniband EDR. In 2016, the bi-directional traffic has been measured at 34 Tbits/s between 512 nodes using OMNI-PATH technology. In the coming years, the remaining factor two that is required will be achieved by improving the software and the transport-layer technologies.

In parallel, the TFC architecture has been revisited, taking into account: the regulation of the data flow using the local throttle mechanism; the required number of PCIe40 boards for each sub-detector; and the evaluation of the PON technology. This architecture will be frozen by the end of 2016.

To house the event-builder and the event-filter farm, a new data centre has to be built at the surface. Several studies have been performed in order to determine the best technology and location. A final decision will be taken by the end of the year.

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. Many novel aspects proposed in this document have already been achieved in the Run-2 high-level trigger. In 2016, the Turbo stream has been improved, becoming more flexible in terms of event content. See Ref. [7] for more details. Therefore, it has been adopted by an increasing number of analyses.

In the remainder of this year we intend to review the strategy for charm physics and evaluate once more the performance of the Upgrade trigger as a whole, taking into account all improvements to the Upgrade detector simulation and reconstruction. The reconstruction sequence was frozen in June 2016 and results will be published by Q4 2016. This will be nine months late with respect to the initial schedule, a delay which is caused by the latest optimisation of the SciFi detector, but one that does not endanger the successful completion of the trigger Upgrade work as a whole.

5.4 Computing

Last year several working groups were established to ascertain the needs and explore new ideas on the framework, the event model, the hardware, the simulation and the analysis model for the Upgrade. These working groups recorded their conclusions in an internal roadmap document which was presented to the collaboration in March 2016.

The roadmap details the work to be done and the decisions to be taken for the computing TDR, including estimates of the required and available effort. The most important goal of the R&D phase are demonstrators that are to be completed by the beginning of 2017. These demonstrators will permit the study of: a new

software framework using a task-based approach; event models in which objects are immutable and composable; and the vectorisation of individual algorithms, taking into account the attributes and capabilities of the hardware platform. Work on others areas, such as the analysis model, simulation and collaborative working, will follow an evolutionary approach in which ideas are already being tested during the current LHC run.

In May, July and September 'hackathons' were organised. During these events developers worked together intensively on a given project for two days. This initiative led to the development of a simplified language to configure application, a tool to manipulate data compliant with a task-based approach, and a new API for users to help in the implementation of the new framework. The first integration of the main components is scheduled to occur during a Computing Workshop that will be held in November in Paris.

The required level of effort for all computing activities is still a concern, but the general situation has improved with the acquisition of new staff, including a software engineer to oversee the overall architecture of the project.

The main upcoming goals concerned with Upgrade computing are the Software and Computing TDR (Q4 2017), and the finalisation of the computing model during autumn 2018.

6 Infrastructure

The outer shell of the new assembly hall at the experimental site has been completed and the cooling and ventilation infrastructure is advancing (Fig. 11). It is expected that all major work will be completed by the end of 2016. This will be followed by the preparation of specific infrastructure required for the assembly of the LHCb Upgrade detectors.

Studies for the new data centre at point 8 are almost concluded and a decision is expected for November this year. In parallel, the CERN management has started to evaluate the possibility to construct a huge data centre at Prevessin to house the CPU farms for LHCb and ALICE, and other facilities. A decision on this option is expected on the same timescale, and if positive will necessitate a change to the LHCb-specific planning.

The position and geometry of the shielding to protect the SiPMs of the SciFi were agreed during the last Technical Board, subject to simulation studies on the impact on the ECAL (see Sect. 3). The production drawings for this shielding are almost ready and candidate companies for the construction will be contacted soon.

The cooling transfer line for the VELO and the UT will be installed in the extended-year-end-technical-stop (EYETS). The types and numbers of copper cables and optical fibres for the read out are being discussed with all sub-detectors and final numbers will be decided by the end of this year.

Preparation of safety documents is required prior to the disassembly of the

Figure 11: The new assembly hall under construction at IP8.

current detector in LS2. The planning for the installation of the new sub-detectors is progressing well.

7 Organisation & oversight of Upgrade activities

7.1 Organisation

The Upgrade Planning Group (UPG) 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.

7.2 Milestones

A series of important milestones has been achieved in the last six months including many EDRs and PRRs of crucial components of all the sub-systems. Overall, the upgrade project is on schedule with a few small delays distributed among the sub-projects. In May 2016 the Upgrade was subjected to an in-depth review from the LHCC, from which it received a positive assessment.

Delays on the production of crucial elements especially in the VELO and UT

projects are becoming critical and are being closely monitored. To this end an internal comprehensive review of the upgrade, with focus on critical aspects, risk-assessment organisation of the construction and preparation for the installation, will be organised in January next year.

PRRs of almost all the components of the sub-detectors will take place within the next six months. The first milestones to monitor major steps in the production are scheduled for the second half of 2017.

8 Funding

The funding requirements of the LHCb upgrade project have been defined in detail in Addendum No. 1 to the Memorandum of Understanding (MoU) for Common Projects [8] and in the Addendum No. 2 to the MoU for upgrade of the Sub-Detector Systems [9], which refer to the LHCb Upgrade Framework Technical Design Report (FTDR) [2] and the Technical Design Reports (TDRs) for all upgrade subdetector-systems [3–6]. These documents define in all details 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 upgraded sub-systems. 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. All TDRs have been fully approved by the Research Board and both Addenda have been submitted for signature to the Funding Agencies.

The LHCb Upgrade project is progressing according to schedule. Price enquiries have been launched, invitations to tender are ongoing and major contracts have been placed. Spending of CORE funds has started for the majority of sub-detector components and a substantial fraction of the anticipated funds are being spent in 2016 and 2017. The Upgrade project is evolving within the agreed cost envelope and there is confidence that full funding will be available in time to ensure a complete and timely installation of the new experiment in LS2.

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, *Status of the LHCb experiment*, CERN-RRB-2016-107.

[8] 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.

[9] 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.