Status of the LHCb upgrade

CERN-RRB-2017-036

26 April 2017

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 [7] and sub-system items [8], 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)

ASIC

The first VeloPix engineering wafers were received at the end of August. After thinning and dicing the first communication with the bare chips was performed in September. This was followed by beam tests in November and irradiation tests in mid-December. The ASIC has shown excellent functionality, low noise and threshold dispersion, and a total power consumption of below 2 W at full rate. Charged particles have been measured in a proof-of-principle testbeam installation and further testbeams are planned for April/May to probe in detail the performance and to evaluate the ASIC at high rates. However, the detailed tests have revealed occasional transmission errors. Even though it is known that these errors affect the data at a very low level it is desirable to correct these in a second submission. In addition, the irradiation tests have shown possible features related to latch up with heavily ionising particles. If these are confirmed, then they will also be corrected in a second submission.

Sensors and cooling

A second round of pixel-sensor prototyping was carried out with Hamamatsu incorporating DRIE-etched rounded corners and uniform 39 μm implants. Sensors from both Micron and Hamamatsu were bump bonded to both Timepix and VeloPix ASICs and irradiated to full and partial fluence. A dedicated test chamber with thin-walled entry and exit windows was designed and installed in the Timepix telescope, allowing the performance with charged particles to be measured in vacuum. The critical performance parameters of the sensors including resolution, charge collection, noise and edge performance and leakage currents were measured. In December the sensor PRR was held, at which the referees recommended that, as the Hamamatsu devices respect the criteria set, the production should be launched with this vendor. It was recommended that further high voltage studies be made in order to fine tune the HV operation and these studies are currently underway.

The first phase of microchannel cooling plate production, which is anticipated to deliver approximately 30 substrates, is proceeding according to schedule and is almost complete. The yield of the wafer bonding step has been assessed with SEM imaging, and the short-loop metallisation has been checked for soldering reliability. Two methods of dicing have been qualified: laser micro-cutting within thin low-pressure water jets, and plasma etching after the application of suitable resist photomasks. A plasma-diced wafer is shown in Fig. 1. Quality issues have cropped up with respect to the setting of the alignment marks on the wafers and edge chipping and these are being addressed with changes to the process steps. Creep and fatigue tests of the substrates have been performed with multiple pressure

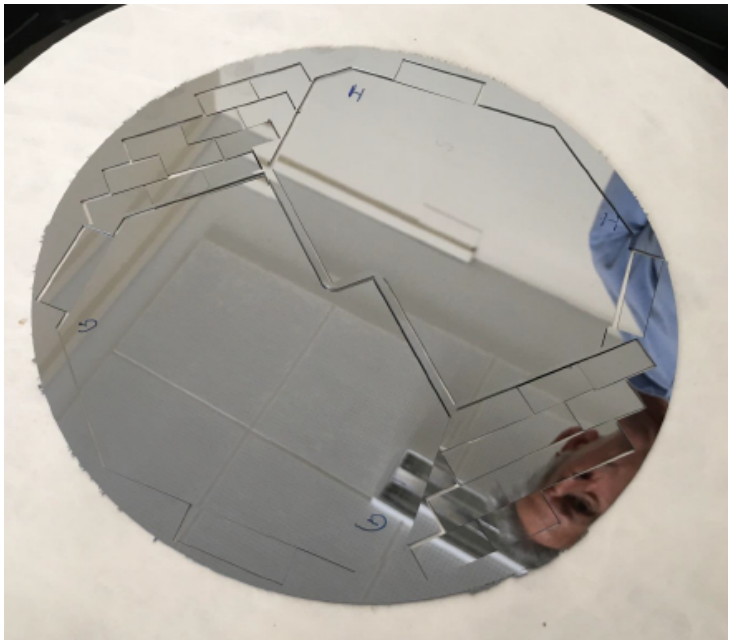

Figure 1: First test of plasma-diced microchannel wafer

and temperature cycles. The soldering of the fluidic connector to the substrate remains challenging and the use of a reducing formic-acid atmosphere is currently under study.

In parallel the backup plan of circulating CO_2 within stainless steel pipes embedded within ceramic substrates has been pursued. This has been shown to deliver a thermal efficiency which is adequate for the experiment, and equivalent pressure-flow characteristics to the microchannels. The design is currently being refined to fit the production site assembly jigs and to optimise the choice of glue and the deposition method in the grooves which host the pipes.

Modules, mechanics and readout

The module design and interfaces has been solidified, with agreement reached on the positioning of the low and high voltage cables and connectors, the cooling loop and VCR placements, and the conceptual placement of the front end and control elements of the copper-kapton hybrid. Much attention has been paid to the module displacement measurements in vacuum with temperature variations, using two parallel setups at the module construction sites, and the results of these tests are being used to modify the module designs. The module construction steps have been defined and have been exercised with dummy components.

The DAQ and readout chain remains a challenging step ahead of the module EDR planned for early summer; even though the VeloPix testing can largely proceed with the SPIDR readout a readout chain with final components will be

essential for full ASIC module readout and characterisation. There has been substantial development of the VELO specific firmware and a close to final readout chain has been set up, using the SPIDR for configuration commands, and the VeloPix hybrid is being debugged. Prototyping work is ongoing for the CTLE emphasis circuits which will be necessary to recover the signal after the copper transmission in vacuum.

The overall mechanical design is progressing and was reviewed at the mechanics EDR at the start of April. Machining of the assembly jig has begun and an assembly mockup has been created of the hood together with a perspex copy of the module support base. This will be complemented with 3D printed parts, including the vacuum flange, and will allow the test installation of parts for fit, access and movement tests before committing to machining a single piece hood. The pipe routing, which must sustain the VELO movement cycles, is being optimised with FEA simulations together with a bending jig that allows many thousands of movements to be excercised with the real pipe designs to check for work hardening.

RF foil

The milling of the first full size RF-foil box was successfully completed in February. The new material used is cast and 3D forged AlMg4.5Mn07, which is expected to give an improved vacuum tightness with respect to previously produced prototypes. The box, shown in Fig. 2, is milled to a thickness of 500 μm and will now go forward to the chemical etching step. Leak tests with the new foil show very good performance, with leaks almost imperceptible above the background rate from the O-ring used in the test setup. The metrology measurements so far show good positional accuracy and a reasonable success in achieving the target thickness. However there are thickness variations around the module slots of around 100 μm, which may eventually limit the total amount of etching which can be applied.

2.2 Upstream Tracker (UT)

Good progress is occuring throughout the UT project in all key areas.

Sensors

The lessons from the May 2016 test beam programme have been incorporated in the development of an improved design of the silicon sensors that is expected to meet all the specifications and overcome some of the problems identified in previous studies. The pre-production samples of the so-called A-type sensors were delivered this month. The sensors will be irradiated and undergo testbeam studies in early June. This schedule will enable detector procurement to begin in the late summer.

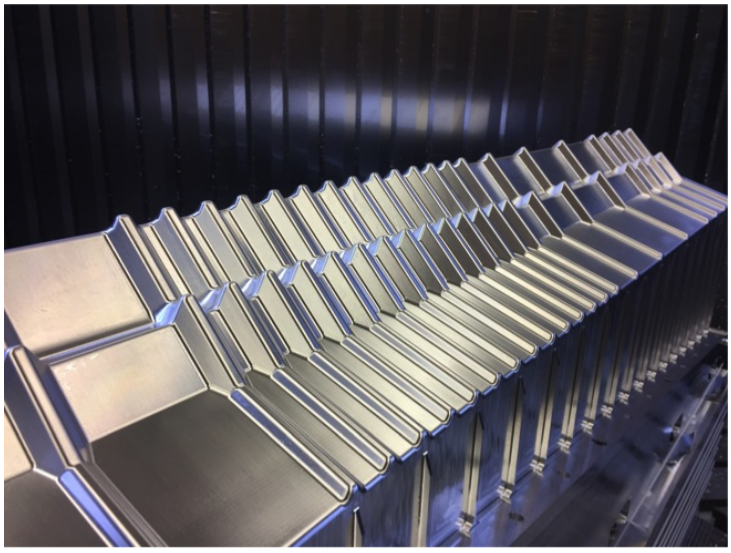

Figure 2: Full sized RF-foil prototype after machining

ASIC and readout electronics and components

A detailed study of the performance of the first full-size SALT chip (SALT-128) has been conducted. Problems have been identified, both in terms of noise immunity, and performance under irradiation, that need to be addressed. A new design, which will correct these deficiencies, will be submitted in late April / early May.

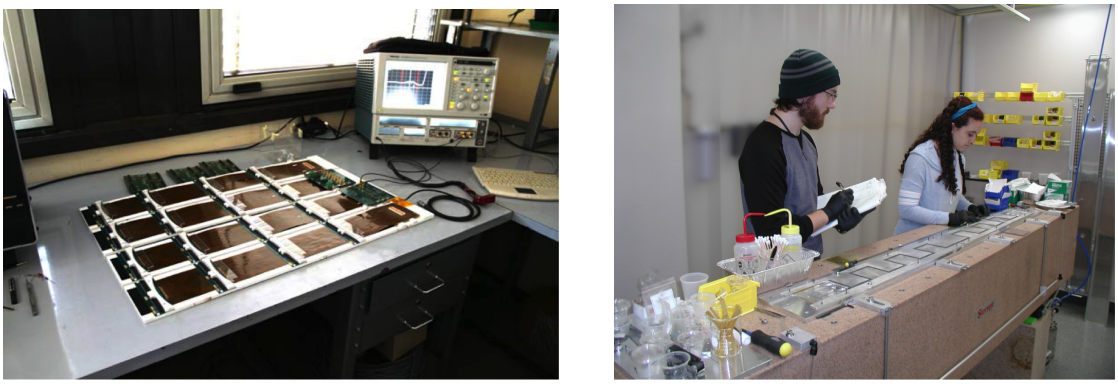

Figure 3: Left: time-domain reflectometry tests of flex cables. Right: assembly of UT staves

A first iteration of the flex hybrid circuit hosting four SALT chips has been fabricated and will be integrated in a slice test of the UT-electronics readout system. Currently SALT-8 chips are being mounted and wire-bonded on these hybrids. They will be mounted on ceramic stiffeners of final design and wire-bonded to preproduction sensors. This work is necessary to validate the electronics

design, but also establishes the methodology for assembling silicon-hybrid modules and instrumented staves. In particular an extensive study of the reliability of the wire-bonding procedure is in progress.

The third generation iteration of the flex circuit connecting the sensor-hybrid modules has been implemented with an industrial partner (ALTAFLEX) and the results have been excellent both in terms of performance (see Fig. 3, left) and delivery times.

Assembly

The pre-production of three bare staves is being completed (see Fig. 3, right). This work is documented on a production manual that describes the procedures to be followed in the construction and testing of the bare staves. This documentation will be reviewed by the referees in the bare stave pre-production readiness review before proceeding to mass-production of the 68 staves needed for the detector. The construction of nine additional staves to be used as spare units is also envisaged.

The planning for the transfer of the staves from the construction site to CERN has started. A detailed assembly and test procedure on the surface, near the experiment site is being planned.

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.

Mat and module production

The industrial supplier of the scintillating fibres started production in May 2016 and is now delivering 300 km of fibres to CERN every two weeks. The fibres undergo various quality checks during which occasionally appearing "bumps" (local increases of the fibre diameter) are automatically removed. Apart from this manageable feature, the received fibres are of excellent quality and exhibit an attenuation length of about 350 cm and a light yield exceeding the values of the pre-series by about 10%. As of mid-March about 50% of the ordered 11,000 km have been received and tested.

The serial production of the fibre mats is running at three winding centres (Aachen, Dortmund and EPFL) and more than 20% of the required mats have already been produced. For the fourth centre (Kurchatov Institute) the PRR wll occur in April. Figure 4 shows the number of the produced mats for the three

active winding centres. The winding centres are trying to increase further the weekly production rate to compensate for delays during the production ramp-up, and to finish by spring 2018.

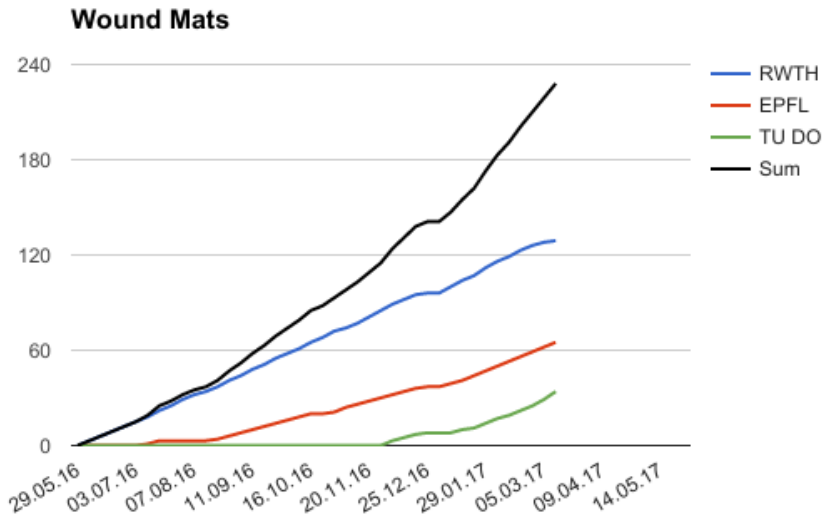

Figure 4: Number of wound mats by the three winding centres Aachen, Dortmund and EPFL.

The first module assembly site (Heidelberg) has demonstrated the envisaged production rate of two modules per week. The site has already produced about 10% of the required modules. The commissioning of the second assembly site (Nikhef) is ongoing (see Fig. 5) and the first module will be produced at the end of March, with a PRR following in April. The delayed start-up of this site has no influence on the overall schedule as the module production is currently limited by the availability of fibre-mats. The module production is scheduled to finish about four weeks after the end of the mat production in spring 2018.

SiPMs

The silicon photo multiplier arrays (5500 pieces) have been ordered at the industrial supplier. The production schedule foresees the delivery of the first production batch in May and the last batch arriving in November 2017. Upon arrival the SiPMs will be mounted on so-called flex-cables and tested. The flex-cables will provide a thermally decoupled, flexible electrical connection between the SiPMs and the front-end electronics. The design of the flex-cables is currently being finalised. Flex-cables to test the first SiPMs will be available.

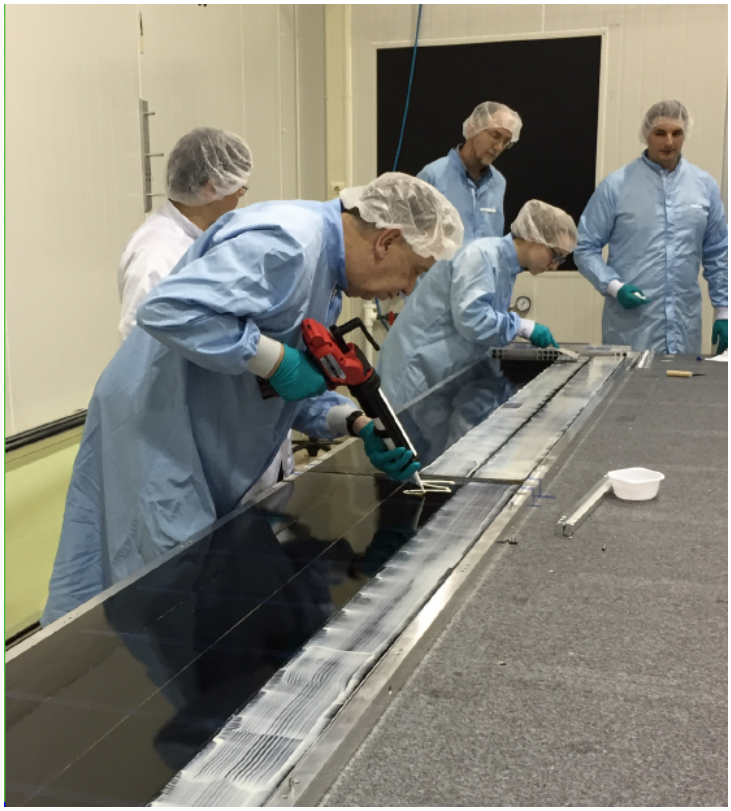

Figure 5: Assembly of a SciFi module.

ASIC and read out

A revised version of the PACIFC readout chip (v4) was received in December 2016. The chip has been intensively tested in the laboratory and, mounted on test-modules, in a testbeam. Both tests show that the chip fulfills the SciFi requirements. This v4 version however features single-ended outputs. On recommendation of the EDR the SciFi group decided to submit in March another development version (v5) to change to differential output signals.

The subsequent readout chain consists of the PACIFIC carrier board, a clusterboard to find the hits, and the master board, comprising the Slow and Fast Control as well as the optical links. The first version of these boards (v1) was tested and fulfills the SciFi requirements. A revised version will account for the differential signal outputs of the final PACIFIC chip and will overcome a problem observed with the commercial FPGA-chips when operated under irradiation. The revised cluster board (v2) has been finished (see Fig. 6), while the revised master board will become available in April. A full test of the revised readout boards together with the new PACIFIC chip (v5) is foreseen for July.

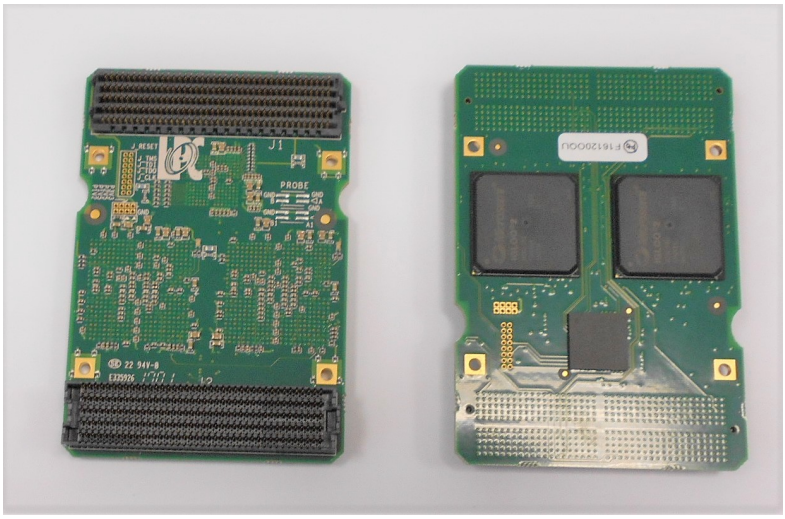

Figure 6: Revised version (v2) of the cluster board.

Read-out box and mechanics

A 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. The first design of the cold-box successfully passed an EDR in August 2016. Modifications suggested by the reviewers have been implemented in the second version. A small pre-series of 20 boxes is currently being produced. Open questions arising from the EDR, such as the humidity management inside the cold-box, the homogeneity of the cold-bar temperature and the thermal insulation, are being studied.

The detector modules and the corresponding read-out boxes of pairs of half layers (one vertical, one stereo) will be mounted on a single C-shaped support frame. In addition to the mechanical support these C-frames will also provide the necessary services to read out and cool the detector elements. The modules of two C-frames closing around the beam-pipe form the detection (stereo) layers. In total 6×2 C-frames will be arranged along the beam-pipe. The first detailed engineering design of the detector C-frame mechanics, including the services, is available. An internal design review is foreseen for spring and a first prototype C-frame should be assembled by the beginning of the autumn.

In summary, the SciFi is project proceeding well: the fibre supply is at its nominal rate; the mat and module productions are ramping up; the SiPMs have been ordered; the development of front-end electronics will soon be concluded; and the work on the detector infrastructure is progressing satisfactorily.

3 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. Mass production of several key detector and readout electronics components is now underway.

3.1 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 RICH1 re-design is performed within the footprint of the current detector, and is therefore compatible with the existing magnetic shielding box and the upgraded VELO and the UT mechanical structures. 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 the 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. So far, seven monthly batches of the MaPMTs production series have arrived, starting from September 2016, and are all being qualified at two regional centres in Edinburgh and Padova (Fig. 7, left). The quality of the production and the collaboration with the manufacturer are excellent. Therefore, a quick feedback can be provided and production adjusted in case of issues.

The MaPMT, CLARO, front-end electronics and system integration have been evaluated in testbeams 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 begun.

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 and PRRs in April 2017.

Prototypes of the system, consisting of several MaPMTs together with the accompanying front-end, digital and acquisition electronics and mechanics (Fig. 7, right), have been evaluated in a series of testbeams in 2014, 2015 and 2016. These testbeam studies were a success, paving the way towards the important milestones of 2016 and full production of the components. A full Photon Detector Module, including its Digital Boards and the LHCb DAQ modules, will be tested and demonstrated during the 2017 beam tests campaign.

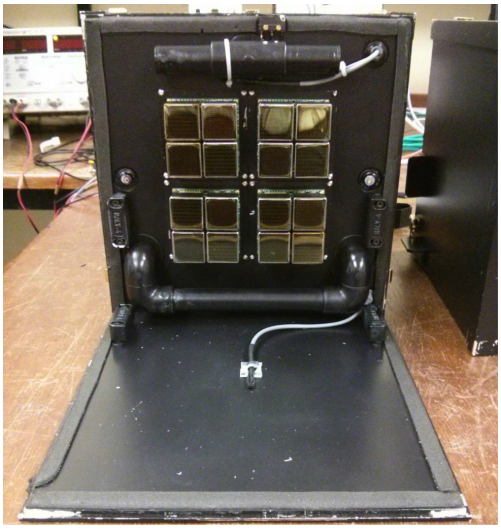

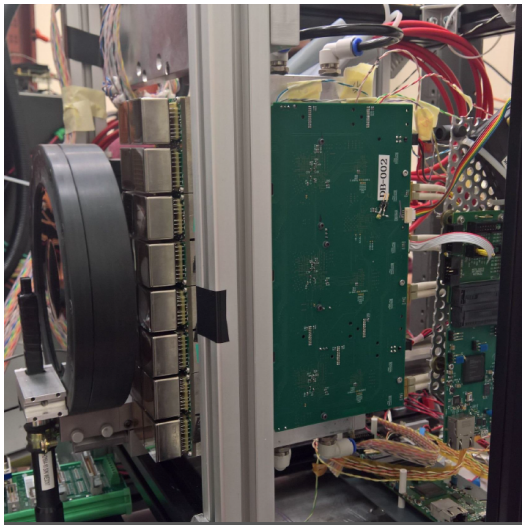

Figure 7: Left: setup for QA tests of MaPMT. Right: RICH readout demonstrator used for beam tests and read out and DAQ development.

3.2 Calorimeter system

The upgrade of the calorimeter system will consist of a replacement of the electromagnetic (ECAL) and hadronic (HCAL) calorimeter readout electronics and the removal of the Scintillating Pad Detector (SPD) and the Preshower (PS). In particular, the gain of the photomultipliers will be reduced by a factor five in order to keep them operational throughout the high-luminosity Upgrade running. The new analogue electronics will compensate for this gain reduction. The upgraded detector will send the full data flow to the counting room at 40 MHz by means of four optical links per Front-End Board (FEB). The present earliest-level trigger calculations performed on the 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. Hence, the current front-end electronics will be fully replaced. The high voltage, monitoring and calibration systems will be adapted to the new slow control based on the GBT driven optical links. A new data-acquisition system relying on the PCIe40 boards will be used, which will require a dedicated firmware adapted to the calorimeter data format.

The analogue electronics is based on an ASIC called the ICECAL. A successful PRR recently took place. The full production of the ASIC should start very soon. The test of the 2000 components that must equip the FEB is being prepared. An automatised bench will be mounted at Barcelona for this purpose.

The current prototype of the FEB, mentioned in the last RRB report, has been intensively tested. The schematics of a new prototype version are almost complete. This new version will integrate all the corrections that have been identified in the previous iteration. The tests of the new prototype will last a few months and the production of the final FEB wil be launched soon afterwards, by the end of 2017.

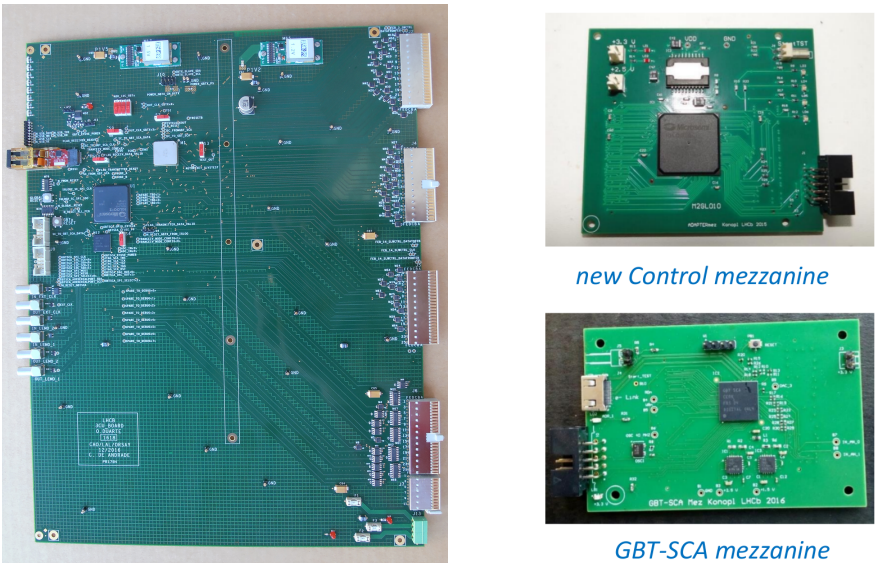

Figure 8: Left: the latest prototype of the calorimeter control board under test. Right: the two prototypes of the mezzanines designed for the high voltage, calibration and monitoring system.

A first prototype of the control board (3CU) has been received (see Fig. 8, left). The tests are well advanced. A consolidation version of this board is planned before the final production, which should take place simultaneously with the FEB production.

The firmware of the FPGA of the FEB and 3CU boards are in their final stage of development. A bench with a MiniDAQ system is being used in the laboratory to test the FEB and 3CU.

The upgrade of the high voltage, monitoring and calibration system of the calorimeter requires the realisation of three boards. Two mezzanines, the control and GBT-SCA ones, will replace mezzanines that are used in the current detector. The former will permit the system to be upgraded to a newer FPGA while the latter is compatible with the GBT optical link, designed at CERN, for the slow control of the system. The third board allows an optical link to be connected from the control room to several mezzanines through copper e-links. This board embeds a GTB component.

Prototypes of the two mezzanines have been produced and tested (see Fig. 8, right). The third one will be very similar to the VLDB board and its design has begun.

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. This will be decided in the coming months.

3.3 Muon system

All aspects of the muon upgrade project have progressed significantly in the last few months.

The production of the spare MWPCs completed at the LNF site and is being finalised at the PNPI site. Sixteen M5R4 chambers, out of a stock of 24 M5R4 produced at LNF, have been transported to CERN late in 2016 for possible installation during the winter shutdown. In the event, this chamber replacement was not necessary and the front-end electronics installation and the final test with cosmic rays have been delayed to summer 2017. The remaining spare MWPCs will be transported to CERN this year and in 2018.

The nSYNC-v1 ASIC was received in September 2016, and tested in the weeks that followed. The bulk functionality (I2C and TFC interfaces, PLL and frame builder) was found to be in agreement with specifications, apart ffrom a missing connection to the LVDS drivers of the GBTX interface. This problem was fixed on two sample chips, and this allowed a successful validation of the chip design. A second submission of the ASIC, incorporating all the required modifications, was made at the end of 2016 and the nSYNC-v2 ASIC was received in February. After tests, this version of the chip will be mounted on the nODE prototype board, which is currently under production. In parallel, irradiation tests are being prepared, which are needed to characterise the nSYNC before submitting the full production. The readout electronics PRR and the nSYNC production submission are scheduled for autumn 2017.

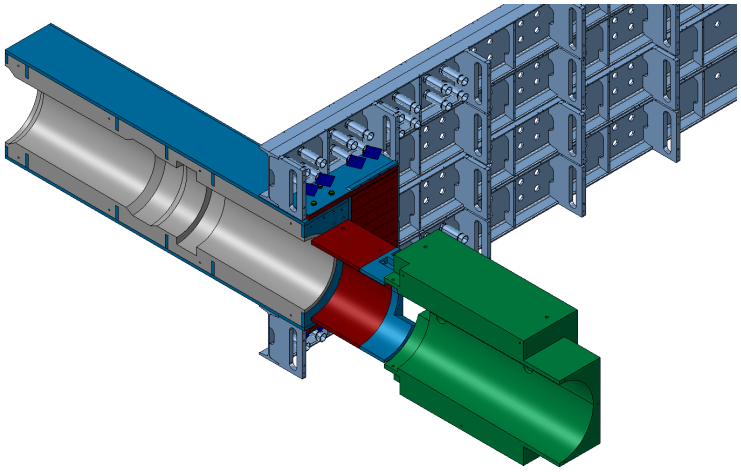

Figure 9: Design of the additional shielding in front of M2, split into three main components: the new HCAL beam plug (grey), the additional shielding behind HCAL (red), and the new M2 beam plug (green).

The work on the muon system specific PCIe40 firmware is continuing. The first test of the firmware using a MiniDAQ setup has been done, and a successful injection of the data frames has been achieved, with proper operation of the zero-

suppression mode. After summer 2017, this setup will be integrated with the nODE and the control board (nSB and nPDM) protoytypes for a full readout chain test.

The design of the new shielding in front of M2, shown in Fig. 9, passed successfully through a full EDR in January 2017. A price inquiry will be launched in the coming months. Survey points were added for installation and alignment of the beam plug during the finalisation of the design. Simulation studies indicate that this new shielding will result in a $\sim 50\%$ background reduction in the inner M2 region is expected. This will be beneficial on the deadtime induced inefficiency, and especially on the input rate to the muon trigger.

4 Online, trigger and reconstruction

4.1 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 and on the generic readout board, PCIe40 (PCI Express), embedded in each PC-Server.

The current priority is to define and optimise the number of PCIe40 boards that the experiment requires. The provisional number of PCIe40 boards for each subdetector was updated in June 2016 and then again in March 2017, with a final decision to be taken in September 2017.

The PCIe40 prototype, with the final version of the FPGA Arria 10, has been replicated, with twenty-four copies being produced during 2016. These boards will be used to equip the MiniDAQ2, which is a PC-server for the testing and validation the front-end electronics of each sub-detector. The MiniDAQ2 systems will be delivered to the collaboration by May 2017. This date will be twelve months behind the initial schedule, a delay which has been caused by the late availability of the Arria 10 and by issues encountered during the manufacturing of the printed circuit board. However, this new schedule is still compatible with the time-lines of the subdetector projects.

The final version of the PCIe40 was designed between January and April 2017. The board is optimised to ease its manufacturing and to reduce its cost as far as possible. A successful review of the board took place at CERN this month. An invitation for tender will be launched by the CERN procurement service in May. The production of the 600 boards will start before the end of 2017.

The TFC distributes timing and fast control to the front-end electronics and readout boards. Its architecture was frozen at the end of 2016. It relies on a dedicated optical network running at 3.2 Gbits/s, as well as the local throttle mechanism implemented in the readout board.

The event builder has to aggregate 40 Tbits/s with a 100 Gbits/s bi-directional network interconnecting 500 nodes. 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. Results are encouraging. In June 2016 a bandwidth of 30 Gbits/s was measured between 512 nodes using the Omni-path technology. In December, 80 Gbits/s was achieved between 64 nodes with the Infiniband technology. In the coming years, the remaining factor of three that is required will be achieved by improving the software and the transport-layer technologies.

To house the event builder and the event-filter farm, a new data centre has to be built a the surface. The tender for a container-based solution located at Point 8 is close to completion. A final decision will be taken by mid-2017, depending on the future of the so-called Green Cube project, which is a CERN-wide data-centre initiative under evaluation by the lab management.

4.2 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 the Upgrade TDR have already been achieved in the Run-2 high-level trigger.

At the end of 2016 the expected performance of the trigger was re-evaluated, taking into account all improvements of the Upgrade detector simulation and reconstruction. The overall performance [9] was found to be unchanged for the first stage (HLT1) with respect to that quoted in the Trigger TDR. However, the evolution in CPU power is lower than had been foreseen at that time. From running the trigger on a server purchased in 2016 it is found that the yearly CPU growth is now 1.1 instead of the factor of 1.37 that had been assumed previously. For this reason the propagation of the current CPU growth rate up to 2020 indicates that a factor of six will be missing in the CPU power of the HLT1 for the Upgrade. The envisaged strategy to mitigate this situation relies on one hand on exploiting the new multi-thread and vectorised computing framework and on the other the re-use of the nodes of the current experiment. The situation will be reassessed in 2018 once the new computing framework is available and a large fraction of the HLT2 algorithms have been vectorised. The re-evaluation will also take account of the requirements of the disk buffer and also the HLT2 throughput.

The strategy to determine the bandwidth division between beauty and charm channels has been revisited [10]. In the proposed procedure, analysts can develop selections in a way that optimises their useful physics output over a broad range of output bandwidth requirements. In the coming months, the bandwidth division will be exercised with three selections from each Physics Working Group in order to cover the full physics programme of the Upgrade.

4.3 Reconstruction studies

The optimisation of the tracking system has been completed and was described in the October 2016 RRB document. Since then, the focus of the reconstruction studies has moved to a detailed analysis of the performance and optimisation of the reconstruction sequence. This analysis is documented in the *biannual performance review document* [9]. Adjustments to the geometry description and a more realistic simulation of the detectors degrade the performance; however, this degradation is mitigated by optimisations of the algorithms. The net performance of the reconstruction sequence for the trigger is hence basically unchanged. Further performance improvements are under active development.

Significant effort will be expended in the coming years to port the current reconstruction sequences to the new software framework that is described in Sec. 5. First encouraging results have already been achieved with the RICH reconstruction (again, see Sec. 5).

Another area where significant development is ongoing is in the Kalman Filter-based track fit. This algorithm consumes a very significant part of the current reconstruction sequence and will in its current form dominate the reconstruction sequence. A simplified version of the LHCb geometry description has been implemented and is currently being tested. Two main approaches are being worked on to speed up the Kalman Filter: a parallelised and a parametrised fit. The fit has been developed in standalone configurations and is currently being ported to the LHCb reconstruction framework.

5 Computing

An internal roadmap document written last year details the work to be done and the decisions to be taken for the Software and Computing TDR, scheduled for the end of this year. The most important goals of the R&D phase are demonstrators that were completed at the beginning of 2017. These demonstrators have permitted 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 in others areas, such as the analysis model, simulation and collaborative working, will follow an evolutionary approach, in which ideas will already be tested during the current LHC Run.

Hackathons are organised every two months. Developers work together intensively on a given project for two days. Such an event made possible the first integration of a reconstruction sequence (MiniBrunel demonstrator) in November 2016. A hackathon in January focused on user-oriented tasks and proved very popular, featuring an advanced C++ course followed by 50 people (see Fig. 10). This approach is therefore found to be a very effective method for spreading the knowledge and tools required for Upgrade computing throughout the collaboration.

Figure 10: Participants at a recent hackathon.

The MiniBrunel demonstrator was successfully running by February of this year. This is a major checkpoint in the progress towards the Software and Computing TDR. It relies on a task-based framework, Gaudi Hive, and runs more than 100 reconstruction algorithms. This achievement demonstrates that the scheduling of a series of tasks for concurrent events using a multi-thread approach works well and that the memory consumption is under control. Precise benchmarking of the performance in terms of timing, cache usage and memory consumption has begun. In parallel, the vectorisation of the RICH ray tracing algorithm shows that factor of up to eight can be gained, in agreement with expectations. How this gain evolves when this algorithm runs in the full reconstruction sequence has yet to be seen.

The main upcoming goals concerned with Upgrade computing are the integration in the new framework of a sequence as close as possible to the upgrade HLT1 by summer 2017, the Software and Computing TDR itself (Q4 2017), and the finalisation of the computing model during the autumn of 2018.

6 Infrastructure

The installation of services for the Upgrade in the cavern continues. During this Extended Year End Technical Stop CO_2 transfer lines for the Vertex Locator and the Upgrade Tracker were designed, fabricated and mounted to the walls between the service cavern and the detector (see Fig. 11 left). The junction box which will distribute the CO_2 cooling from the end of the transfer lines close to the UT and VELO will be installed before the re-start of Run 2.

The new assembly hall is ready and layout of the interior is under discussion (see Fig. 11 right). This area will be of particular importance in the final stages of the SciFi and UT projects.

The technical coordination team is finalising the catalogue of needs concerning the cooling systems, cables and electrical distribution. This requires a careful planning of systems that will have to be removed before the installation of the Upgrade. The design of the neutron shielding, required for the SciFi, has been completed and the tender will soon be launched.

The LHCb integration office continues to update all 3D models as the design and first construction progresses, and supports the Upgrade project groups by designing support structures.

An updated and detailed installation schedule now exists, templates for work packages and safety documents are ready to be distributed, and the first work packages for the cooling systems are in hand.

Figure 11: Left: transfer lines for the VELO and UT cooling. Right: preparations inside the new assembly hall.

6.1 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.

6.2 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. Notable among these were the PRR of the VELO pixel and the UT strip silicon sensors,

the PRR of the UT FLEX cables, the VELO mechanics EDR and the RICH1 mechanics PRR. A snapshot of the global status is given in Fig. 12. This indicates an overall delay of around three months with respect to original planning, but shows that progress continues at a healthy rate. Overall, the Upgrade project is on schedule, despite having some areas of concern that are closely monitored. An internal in-depth review was held in February with focus on the most critical issues and on the organization of the construction and installation phase. A thorough risk analysis was also required for all the sub-projects. Another in-depth review has been planned by the LHCC for May 2017.

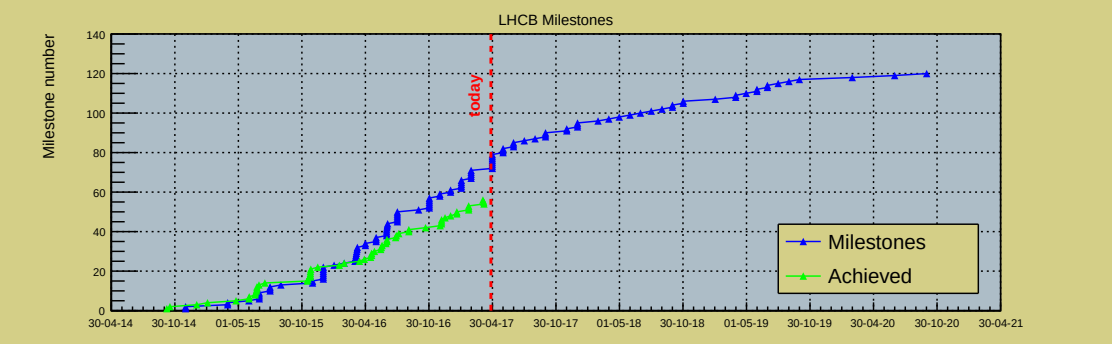

Figure 12: Snapshot of LHCb Upgrade milestones.

Delays on the production of crucial elements, especially in the VELO and UT projects are being closely monitored and decision flow-charts and detailed schedules have been prepared. The first milestones to monitor major steps in the production are scheduled for the second half of 2017.

7 Funding

The funding requirements of the LHCb Upgrade have been defined in detail in Addendum No. 1 to the Memorandum of Understanding (MoU) for Common Projects [7] and in the Addendum No. 2 to the MoU for the Upgrade of the Sub-Detector Systems [8], 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 is ongoing for most of the sub-detector components and a major fraction of the anticipated funds will be spent in 2017 and 2018. The Upgrade project is evolving within the agreed cost envelope and there is confidence that the funding profile will match the spending profile 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, *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.

[9] R. Aaij *et al.*, *Upgrade trigger: biannual performance update*, LHCb-PUB-2017-005

[10] C. Fitzpatrick *et al.*, *Upgrade trigger: Bandwidth strategy proposal*, LHCb-PUB-2017-006