Status of the LHCb upgrade I

CERN-RRB-2018-101

31 October 2018

1 Introduction

The LHCb Upgrade I 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. All sub-detector and common projects are progressing and are on schedule for the installation deadline although some areas of concern are present in some of the sub-projects. Substantial progresses was made in the development of the fully software trigger and an Upgrade Software and a Computing TDR has been released and submitted to LHCC in May 2018 [9].

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

Production Readiness Reviews (PRR). Delivery of components and production of detector assemblies is in full swing, and many quality assurance (QA) stands are in operation. In the UT project, a delay due to the finalization of the SALT ASIC is being mitigated by redefining the installation strategy. 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 VELO

Front end ASICs, Sensors, and Tiles

The sensor order has been completely fulfilled. The ASIC production order was launched and the first 12 wafers showed a 63% yield for class A, yielding sufficient ASICs for VELO installation. Following a successful Price Enquiry procedure for the bump bonding, the sensors have been distributed to the bump bonding company. The pre-production phase for the bump bonding order was successfully fulfilled with a delivery of 20 tiles which underwent detailed tests. The IV (current-voltage) curves are measured at the bump bonding company directly after tile production, and compared with the single pixel IV curve provided by the manufacturer. After reception and testing by LHCb the IV curves are repeated under vacuum conditions. Further detailed checks were performed on the pre-production tiles, with each tile undergoing mechanical survey, equalisation procedures, tests with a radioactive source and search for missing and shorted pixels. The tested tiles were distributed to the module production sites for use in prototyping and were also used in the construction of the first full electrical chain. To date, more than 150 tiles have been produced out of 400 ordered, of which 208 will be required for final installation. A production tile is shown in figure 1 (left). The tile testing procedure and equipment is operating very smoothly; a vacuum jig equipped with a copper stencil with recesses for sensors allows for the accurate positioning and probe card testing of 10 tiles at a time, followed by an easy mechanical transition to the vacuum jar for IV curve measurements. To date, 61 tiles have been fully tested, of which 55 are of installation quality. The losses due to bump bonding failures are very small; there have been a few losses relating to the commissioning of the equipment. All information from the tests is stored in the database, which is fully functional for the tracking of the ASICs, sensors, and bump bonded tiles. The number of available Velopix2 ASICs will be exhausted in a few months; a fresh production order to make up the numbers to cover prototyping and spares has been launched and the ASIC testing, thinning and metallisation timetable has been agreed with the bump bonding company.

Electrical Readout Chain

The externally reviewed module PRR was successfully passed at the end of April. The objectives were to ensure that the complete chain, comprising front end hybrids, flex tapes, vacuum feedthrough board, Opto and Power Board (OPB), to-

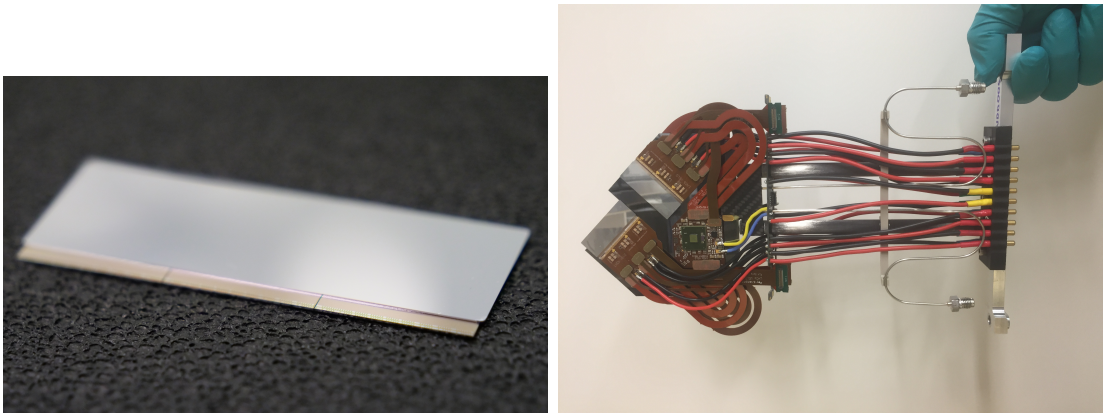

Figure 1: Left: an assembled tile. Right: an assembled prototype VELO module.

gether with the off detector electronics, function in a coherent manner and satisfy the requirements in terms of reliability and radiation hardness. All elements were individually examined, and various revisions were suggested which have been incorporated. Since then a significant milestone has been passed with the complete equipment and readout of a double sided electrical assembly. This included four fully assembled hybrids equipped with Velopix2 tiles, all interconnect cables and full length data tapes, vacuum feedthrough board and an OPB. The setup was configured and controlled from Wincc panels which are being developed within the same structure as will be used for the commissioning and operation of the detector. The control chain is fully working to control, monitor and configure all VeloPix chips and the high speed links have been tested with PRBS31 signals. The setup has been successfully duplicated at the module assembly sites and is in use for module testing.

Microchannel Substrates

The microchannel substrate manufacturer has almost completed two production lots, which will be added to the pre-production prototype supply. Each cooler of the two potential coolers per wafer is classified based on the SAM (scanning acoustic microscope) images from the wafer bonding step as Grade A (no imperfection), Grade B (imperfection in the inlet/outlet area) or Grade C (imperfection in the area of the channels), where only Grade A is considered as production quality. In addition, the standard surface, metallisation quality and mechanical requirements are imposed on the diced coolers. The first production lot had an unexpectedly low yield due to a problem in the final production step which occurred at the subcontracted plasma dicing company. This affected most of the lot, resulting in lower yield. Following this, production was paused while a dicing process flow was defined; the procedure is expected to be relaunched shortly, and it is expected that, due to the overall low yield, a third production run will be necessary; however

for the next four months the supply of coolers is sufficient to guarantee module production. Upon reception of the coolers, the attachment of the fluidic connector is achieved using a soldering technique which involves more than 80 steps and is documented in detail with images, videos and reports. Some highlights of the technique involve the use of formic acid and vacuum environments for different steps in the process, metallisation and sandblasting of different surfaces to promote (or suppress) solder attachment, and detailed quality control including 3D tomography. The production rate is sufficient to match the demands of the module assembly sites. During the prototyping phase a problem was observed at the assembly sites of very small leak rates which would degrade the vacuum quality. Using the solder attachment detailed records it was possible to trace this problem to a heat transfer issue in one of the soldering steps. This has now been remedied and a helium leak test installed to ensure efficient production from now on. The long term tests are ongoing; with the prototype microchannel coolers having now undergone up to 6000 hours of operation up to 70 bar and proving to be robust against thermal cycling and thermal shocks.

Modules

The module PRR was successfully passed in July of this year. In preparation for the review, prototypes were assembled at both module assembly sites. Due to lack of final components a final fully functioning module was not yet available, however no show stoppers were identified and the construction sites are now advancing at full speed into the preproduction phase. A number of detailed modifications and decisions were taken relating to the module construction. The gluing procedure for the tiles - where Stycast is used to optimise the thermal connection - and for the hybrids - where a silicone glue is used in order to provide flexibility against the CTE strain - has been under intense refinement, and has been the subject of many variations in the prototype building. The installation procedure for the modules has been examined and, following this, changes have been made to the routing of the large amount of service cabling and piping which links the module to the base. The most important changes concern the cooling pipe routing and LV supply. The cooling pipe routing has been modified to avoid a the need to thread of the LV supply cables and includes a custom designed clamp which protects the microchannel substrate from forces introduced during the VCR tightening procedure. The LV cable harness has been modified with custom machined connectors to allow an easier insertion of the module into the base and a custom transition PCB which allows connection to finer flexible cables at the level of the module substrate. During the prototyping the survey tools have been finalised and it has been shown that the modules are manufactured with an excellent tile placement precision. The assembly jigs have been finalised and final versions of all the elements of the module are available and documented on EDMS, and production has been launched. A prototype module, with all components, is shown in figure 1 (right). Each assembly site now plans to build 2-3 modules of production quality,

following which there will be site-specific reviews prior to the launch of production. The rate of production will be 1-2 modules per week per assembly site, which allows a comfortable production schedule in parallel with the final assembly.

RF foil and Mechanics

There has been excellent progress on the construction of the major mechanical pieces. The complex isolation vacuum pipe and valve routing manufacture and delivery is complete and the pieces are under test. The first large vacuum hood has been completed, along with the vacuum and transport trollies. The module assembly frame from the VELO I construction is being modified to fit the taller, thinner modules, and the insertion frame is being designed in consultation with the module assembly sites. The construction accuracy of the module base, on which the modules will be fixed prior to insertion of the halves, is particularly important. The base construction proceeded in a phased manner, with multiple surveys in order to refine the mechanical accuracy for the final machining steps. The completed bases were surveyed in a horizontal position supported in the identical manner as in the final experiment, proved to be well within specifications, and have now been delivered to the assembly site. The RF foil construction is proceeding smoothly. One box has been successfully machined to the target thickness of 250 μm and has undergone successful metrology and leak tightness tests and a second one is under preparation. In addition the pilot box is undergoing a number of tests related principally to a possible chemical etching step which may - if proven to work reliably - be performed on the second set of boxes.

2.2 Upstream Tracker (UT)

The UT upstream tracker production plan is proceeding at a good pace: the main challenge remains to absorb the delays related to the SALT ASIC. The SALT ASIC does not operate successfully with the UT silicon sensors. A strong effort is being put in place to refine and optimize the installation and commissioning plans of the UT to mitigate the significant delays in achieving a successful design of the SALT chip.

Silicon sensor-hybrid module

The type A sensor production is continuing according to plans. We have received and qualified the second batch of sensors (80) and we are expecting the third batch (150) in early October 2018. So far, only one sensor out of 105 is rejected by QA. A few sensors do no meet the surface quality requirements. The identification of surface spots (small scratches or residues) is performed with a dedicated pattern recognition which makes the QA step considerably less effort intensive. A preseries of the smaller Si sensors that are needed for the innermost portion of the detector is being ordered.

A second iteration of the hybrid design has taken a considerable amount of time to be implemented successfully. The requirements posed by the SALT designers included a very strong capacitive filtering of the power supplies. Consequently a very dense array of interdigitated capacitors was included in the hybrid design, making the component integration challenging. Recently, a few hybrids have been produced and they are currently under evaluation. These delays have made it difficult to perform a detailed comparison of the two designs. However we are now in a position of being able to perform rigorous studies that should allow the hybrid design finalization shortly after the new SALT design is submitted for production (see below). In parallel, the mechanical aspects of the hybrid production are being studied. The foils will be received and tested, then individual hybrids will be assembled, together with Si sensors on ceramic stiffeners which will be mounted on the bare staves in a double-sided structure called "instrumented staves." The mechanical assembly infrastructure is in an advanced state of prototyping and all the steps have been validated with sensor and hybrid prototypes. We are planning to produce multiple copies of the assembly jigs to expedite as much as possible the instrumented stave assembly phase.

SALT ASIC

The first 128-channel SALT ASIC was submitted in June 2016 and tests started in September 2016. Several problems were uncovered in all the chip's building blocks, affecting the dynamic range, noise performance, and radiation resilience. In addition, focused ion beam editing was needed to resolve some current leak between ground domains that affected the SALT ADC operation. This led to several design changes that have been studied with extensive simulation and were reviewed in a pre-submission readiness review that took place in April 2017. The implementation of the referees' recommendations led us to submit the design for an engineering run in late June 2017. Unfortunately laboratory tests of the received version two ASICs uncovered several new problems that make these ASICs not suitable for the production modules. The solution proposed by the designers was tested on wafers whose metal layer could be processed with a different layout to validate the effectiveness of the changes proposed. These modifications were unsuccessful. However, one change introduced, namely the availability of two test channels that can be connected to a Si sensor and whose analog output can be probed in different processing stages with an oscilloscope, facilitated more understanding of the problems.

The key issue with the SALT ASIC appears to be an instability of the analog section that is sensitive to the capacitive load at the input of the preamplifier and is triggered by the SAR ADC activity. Design changes to improve the stability of the analog section, the power supply rejection ratio, as well as to reduce the dynamic swing of the ADC current are underway. More detailed studies are required to assess the quality of the design to be submitted in October, with the goal of having working front-end ASICs before the end of the year. While the core problem of

the SALT chip prevents us from proceeding to the production of silicon-hybrid modules, we have been able to test multiple chips in a "slice" of the readout system shown in Fig. 2 and thus we have gained valuable experience in operating multiple chips. Example of analog responses of test channels are shown in Fig. 3.

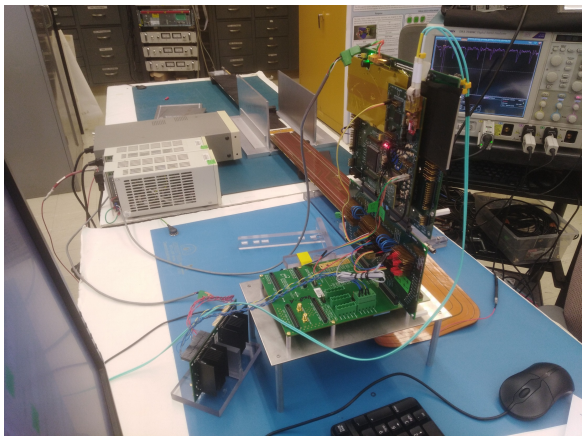

Figure 2: Electronics configuration for the first UT slice test: a hybrid is mounted on the farthest location from the peripheral electronics and communicates with it via a final-design prototype flex cable.

Electronics

The more advanced component of the detector electronics is the flex circuit providing power and bi-directional communication with the data control boards (DCB). There are three different geometries needed. In order to optimize manufacturability, the three cable types are derived from a single design, the one of the long cables used in the innermost staves, which are the ones with the highest number of interconnections. The order for these cables has been placed and the first samples are expected imminently. The QA of the flex cables includes a qualification of the connectivity, DC resistance of the power distribution lines and differential impedance of the signal lines via time-domain reflectometry.

The production of the near detector electronics is expected to start in 2018 and is consistent with the project schedule. The PRR of the PEPI electronics took place in August 2018. The complex backplane, which provides connectivity between the on-detector electronics and the other electronics infrastructure, is turning out to be a very challenging project. As the various electronics components of the project are becoming available, a system test of a first instrumented stave is planned for the fall 2018. This will be a prototype comprising 14 modules with hybrids populated mostly with ASICs not attached to sensors, because of the SALT problems discussed above. The final PEPI electronics will be included, as well as a prototype of the backplane. This test has the purpose of gaining experience with the complex powering scheme envisaged for the UT detector, as

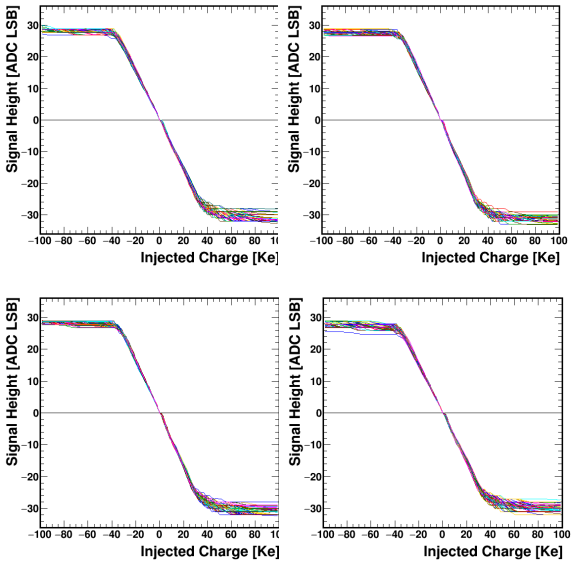

Figure 3: Study of the dynamic range of the four SALT chips integrated in a flex-hybrid prototype included in the slice test shown in Fig. 2.

well as validating the data acquisition model and refining the experiment control panels currently being implemented. In addition, it will allow us to commission the test infrastructure in the construction site at the surface dedicated to the UT detector.

Mechanics

The bare stave production is almost completed. The next stage in the production flow is the gluing of the four flex cables needed to connect the SALT hybrids with the near detector electronics. Two set of fixtures have been prepared to process two stave faces in parallel. Fig. 4 shows a prototype flex cable of the final design glued onto a stave. This flex cable has been used with prototype hybrids to demonstrate that the cable provides all the needed functionality to operate the hybrids correctly.

Integration in the experiment

The design of the frame and box that will support and enclose the UT detector is progressing. A mechanical mock-up of this assembly, including mechanical samples of the various components that occupy the very densely populated region near the ends of the staves, is in an advanced stage of construction. The infrastructure needed to assemble and test the staves in the surface laboratory at LHCb is being finished and commissioned. A comprehensive electronics test of the first fully instrumented stave in this laboratory is planned for the fall. A dedicated CO_2

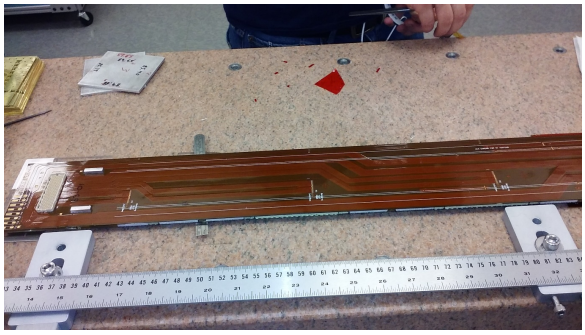

Figure 4: A flex cable glued onto a prototype stave.

system has been commissioned and will be used in the instrumented stave test to operate the electronics at different temperatures. The focused effort towards the instrumented stave test has brought significant impetus for the development of the slow control and experiment control software for the UT, which has made great strides since last spring. A workshop is planned in October to refine the planning for the integration and testing activities planned for the Summer 2019.

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 are joined together to form 5 m long and 52 cm wide modules. The fibres will be read out by 128-channel arrays of Silicon Photo-multipliers (SiPMs), which have to be operated at -40°C to limit the dark count rate after irradiation. The readout electronics is based on a custom-designed ASIC followed by digital boards for further data-processing and the optical data-transmission. The modules including the readout electronics will be mounted on support frames and will be arranged in 12 stereo layers.

Mat and module production

All fibres (more than 11'000 km) have been delivered by the industrial supplier and have been refined and tested at CERN. The serial production of the fibre mats was done at four winding centres (Aachen, Dortmund, EPFL and Kurchatov Institute) and the main production is finished. Production of additional spare mats will continue until the end of 2018.

Module production is progressing well at the two production centres, Heidelberg and Nikhef, and about 90 % of the required modules have been produced. The production of modules including spare modules will be concluded in November 2018. The module production as a function of time is shown in Fig. 5.

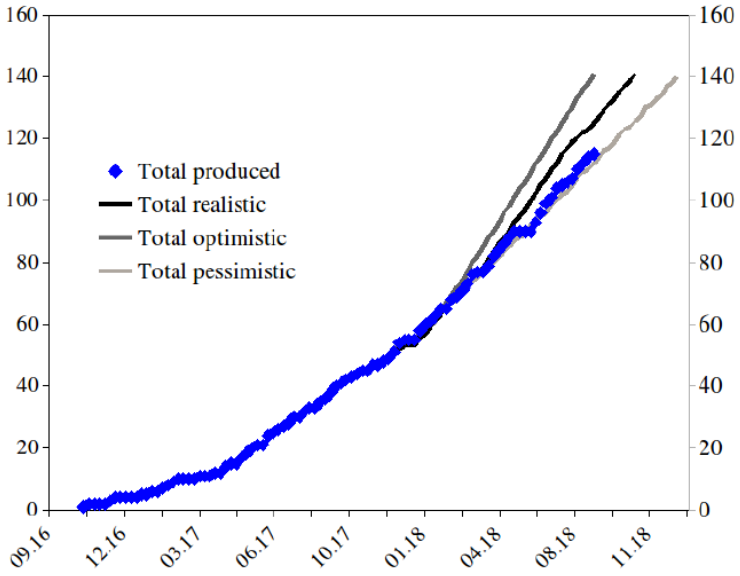

Figure 5: Number of produced modules as function of time.

SiPMs

All silicon photo-multiplier (SiPM) arrays, 5500 pieces, have been received from the industrial supplier and have been inspected at EPFL. The SiPMs are balled and afterwards mounted on flex cables by two different companies. About 2000 assembled detectors have been received and have been tested successfully. The first batch of tested detector assemblies has been shipped to Nikhef where they are mounted on cold-bars. According to the production plan all SiPM assemblies will be produced by January 2019.

ASIC and read-out box

The final ASIC design was submitted for production in January 2018. The first packaged chips from the engineering run arrived in Heidelberg at the beginning of May and were tested before they were sent to China for the production of the PACIFIC Carrier Boards.

The subsequent digital readout chain consists of a Cluster Board to group the hits, and the Master Board, comprising the slow and fast control interfaces as well as the optical links (see Fig.6 for illustration). After a successful PRR of all elements of the readout chain in January 2018, the serial production of all electronics components has been launched in spring 2018. First boards of the pres-series of PACIFIC Carrier Boards, Cluster Boards and Master Boards have arrived and have been tested. The full pre-series, sufficient to equip one C-frame (one twelfth of the final detector), is expected for October. The main board

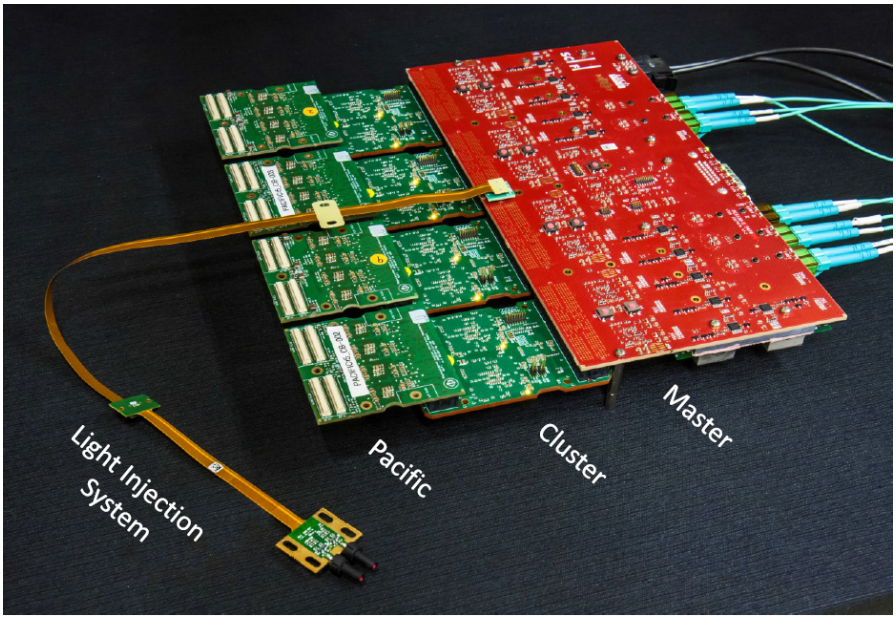

Figure 6: Photograph of the electronics readout chain.

production will start in parallel and will continue in 2019.

The front-end boards will be mounted on cooling frames which are thermally connected to cold-blocks connected to the water circulation system. The cooling concept has been successfully validated under realistic conditions. The first cooling frames should arrive in September, in time to start assembling the front-end boards to so called read-out boxes (ROB) in October. The cold-blocks of the water circulation system are currently being ordered at an industrial supplier.

Cold-box

The SiPMs are not part of the ROB but are mounted in a separate mechanical unit, the so-called cold-box. The SiPMs are carried by a cold-bar which will be cooled down to -40° C using Novec, a modern cooling liquid with small environmental impact. The cold-bar further allows the precise mechanical positioning of the SiPMs on the ends of the fibre modules. Sufficient thermal insulation and gas-tightness to avoid ice building-up is provided by the cold-box. The production of the cold-box outer shells has started and about 120 enclosures have already been produced (about 50 %). The first cold-bars have been equipped with SiPMs and the precision mounting of the SiPMs on the cold-bars have been successfully verified. The first complete cold-boxes have been finished and tested and the production is now ramping up.

The cold-boxes will be mounted on both ends of the fibre modules before installation. The flex cable of the SiPMs will later be connected to the front-end electronics. The module finishing, i.e. mounting of the cold-boxes, is currently

being prepared at CERN.

Mechanical structure, services and detector assembly

Groups of five or six detector modules and their corresponding cold and read-out boxes will be mounted on C-shaped support frames. Each C-frame will carry a vertical and stereo half-layer. 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. In addition to the mechanical support these C-frames will also provide the necessary services to power, read out and cool the detector elements.

The design of the detector C-frame mechanics, including the services, is finished. A first prototype C-frame has been built and is being tested in the new construction hall SXL8 at Point 8 (see Fig. 7). In September the prototype frame will also be equipped with modules and electronics and will serve as a test-bed to develop the necessary debugging and commissioning tools.

The PRR of this rather complex mechanical component will take place at the end of September. In parallel, orders of the first two C-frames have been placed. The current production schedule foresees to start the assembly of the first frame in December 2018. According to this schedule, production, assembly and the test of all 12 detector elements will conclude in March 2020, in time to finish the installation of all components before May 2020.

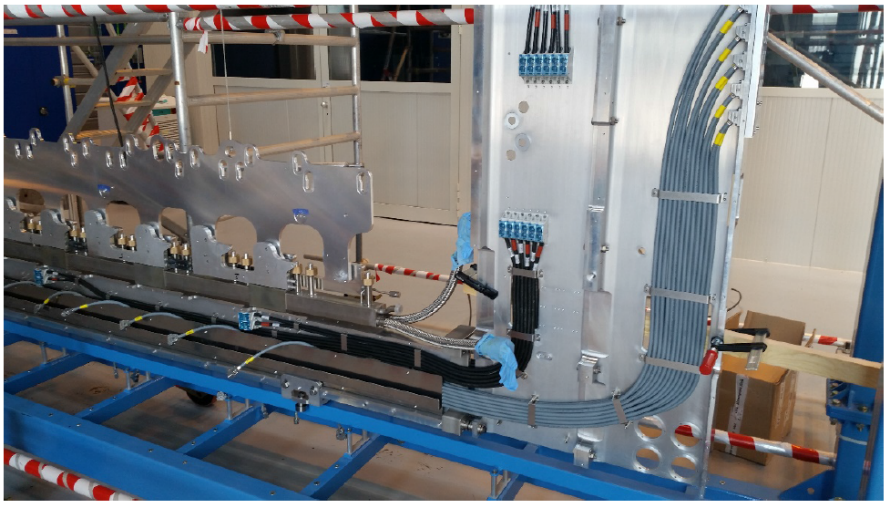

Figure 7: Photograph of C-frame prototype.

Service systems

The operation of the SciFi will require the cooling of the SiPMs down to a temperature of -40°C. A reduced-size demonstrator cooling plant has been built and

is available to perform cold tests during the detector assembly. The distribution of the cooling liquid (Novec) will require vacuum insulated distribution lines. A vacuum station has been built. To prevent icing of the SiPMs the inner cold-box volume will be flushed with dry air. A corresponding dry-air system as been purchased and is available. First cooling tests involving all three systems (Novec, vacuum and dry-air) are being performed with the C-frame prototype.

All other service systems (high-voltage and low-voltage supplies, data-acquisition system) necessary for the full functional tests of the assembled detector frames have also been prepared.

Beam-test

A beam-test of two full-width SciFi modules, with cold-boxes and final read-out boxes, has been performed in July 2018. The full-system comprised 4096 channels and was read-out by 32 optical links. The measured efficiencies (>99%) and the measured spatial resolution (80 μm) agree well with the design specifications and validate the SciFi components at the system level.

3 Particle identification

The Particle Identification (PID) system of the upgraded LHCb detector consists of the Ring-Imaging Cherenkov (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 front-end electronics components is either complete or approaching completion. The projects are focusing on assembling the detector modules and electronics boards.

3.1 RICH

The upgraded RICH system consists of a re-designed RICH1 detector, an essentially unchanged RICH2 structure, and new photo-sensors for both with new front-end electronics that can be read out at 40 MHz. The Cherenkov photons will be detected using Multi-anode Photon Multiplier Tubes (MaPMT) read out by a customized ASIC, the CLARO chip. Simulations indicate 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.

MaPMT, front-end ASIC

The order for the MaPMTs was placed in 2015, the pre-series devices arrived and were accepted in April 2016 after QA tests. The full production has now been delivered and qualified, i.e. 3100 MaPMTs of type R13742 and 450 of type R13743. The quality of the production and the collaboration with the producer was excellent, quick feedback was provided and the production adjusted in case of

issues. Discussions are ongoing with the manufacturers upon replacement of the few remaining non-conforming MaPMTs.

The MaPMT, CLARO, front-end electronics and system integration have been tested in test-beams and irradiation facilities. All results to date are satisfactory, PRRs have been carried out successfully during summer 2016 and the tendering process for the electronics components is finished. Components are being produced. Notably, the full CLARO ASIC production has been received (more than 100'000 chips) and the chips are being tested and qualified.

Digital electronics

Important studies have been carried out to assess the compliancy of all the electronic and mechanical components to the future hostile radiation environment. One important decision was the choice of the FPGA to adopt for the Digital Boards. Following these tests, we have confidence that the Xilink Kintex7 complies with the RICH specifications. and production has been started. The PRR of the Optical link plug-ins was successfully passed in December 2017 and the order has been placed. The last electronics PRR, concerning the motherboard, was passed in May 2018 and production will start soon.

Mechanics, mirrors

The photodetector assembly, including the MaPMTs, all on-detector electronics and ancillary systems, is common to RICH1 and RICH2. The PRRs have been passed and all components are ready to go through the tendering procedures for the RICH1 and RICH2 mechanical systems. The carbon fibre spherical mirrors for RICH1 have been ordered and are now in production and we are ready to order the flat mirrors.

Photon detector modules

The elementary cell (EC) of the RICH hosts either 1 or 4 MaPMTs (depending on type), covering approximately the same detection area. Four ECs are put together to form one Photon Detector Module (PDM). A complete PDM is shown in Fig. 8. Prototypes consisting of several MaPMTs together with the accompanying front-end, digital and acquisition electronics and mechanics, have been evaluated in a series of successful test beams from 2014 onwards, with two sessions in 2017, and the last one will be in October 2018. For the first time in LHCb, we have successfully tested a full PDM in a DAQ system combining MiniDAQ1 and MiniDAQ2 modules. We are now operating the system with two MiniDAQ2 modules to prepare the important activity of assembly and commissioning of the RICH1 and RICH2 columns (each made of 6 PDMs), which will take place at CERN in the RICH commissioning laboratory (ComLab). A complete PDM with its upgrade-compliant DAQ has been installed in RICH2, i.e. in the LHCb exper-

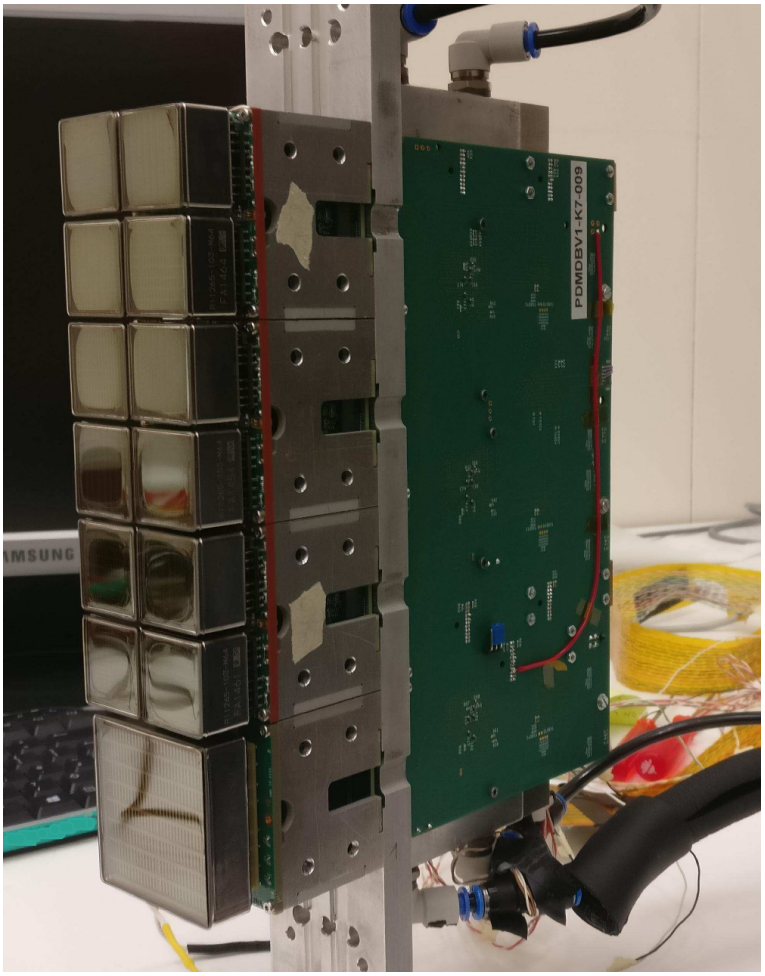

Figure 8: Photograph of a complete Photon Detector Module.

imental hall at Point 8, and is successfully taking parasitic test data during the current LHCb 2018 run.

3.2 Calorimeter system

The upgrade of the calorimeter system will consist in the replacement of the electromagnetic (ECAL) and hadronic (HCAL) calorimeter readout electronics and the removal of the Scintillating Pad Detector (SPD) and of the Preshower (PS). The gain of the photomultipliers will be reduced by a factor up to five in order to keep them operational throughout the high-luminosity upgrade running. The new analogue electronics will compensate for the 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 will be sent to the PC farm in order to optimize the software trigger. Hence, the current front-end

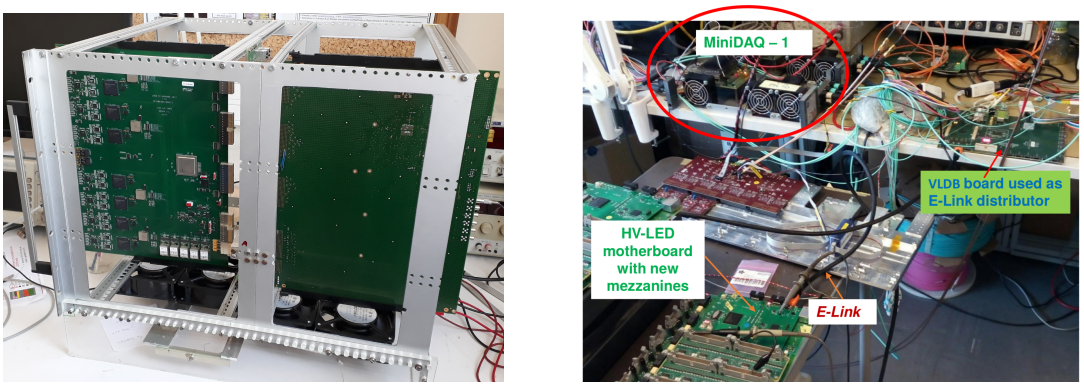

Figure 9: Left, picture of the common tests performed with the FEB and the control board in a specific crate, giving an easy access to the system. Right, picture of the HV/Calibration/Monitoring systems under tests with the MiniDAQ at CERN.

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 ICECAL. The component has been produced in full quantity and tested in Barcelona on a specific automatized test bench with pneumatic suction and pressure sensors. Two versions of the ASIC have been made with two different gains giving us the possibility to choose the transverse energy range of the calorimeter easily just before the production of the Front-End boards (FEB) on which the ICECAL will be soldered. The amount of spares, after tests and validation of the components, for both types of gains is large enough in order to equip all the FEB with their spares.

The FEB has been thoroughly tested. The performance of the prototypes meet the specifications in term of linearity, noise, cross-talk and time-stability of the sampling with the ICECAL chip. The board has been tested in a crate with the final version of the slow control (GBT-X and GBT-SCA system). The integrity of the data sent by the board has also been carefully checked.

The PRR of the FEB was passed in February 2018. A few corrections and improvements have been implemented on a new prototype which was received in the spring. This new board has been tested thoroughly and is fully functional. At present the activity on the FEB is focusing on the firmware. The market survey for the production of the FEB was launched in July and the company in charge of the realization of the board has been chosen. We expect the start of the production to occur in autumn 2018. A first set of two boards will be made and tested in our laboratory. Then, a pre-production batch of 16 boards will be received and tested before the rest is produced.

The control board has also passed its PRR and has been carefully tested during the past months. A few modifications have been implemented on the last prototype

these mainly concerned some extra decoupling capacitors and the grounding of the DC-DC converters of the board. The present prototype is satisfactory and considered to be the last one. The board has been tested altogether with the FEB (see Fig. 9, left). No market survey is necessary for the control board production that should start in autumn 2018.

For the upgrade, we will have to modify the crates and MARATON power supplies that are being used in the cavern at present. Those modifications have been tested on a lab crate and a spare crate at CERN. The procedure to modify all the crates is now well-defined.

A realistic test of our FEB and control board in a final and modified crate is planned for mid-October at CERN. We will use an electron beam at energies ranging from 20 to 120 GeV. Several calorimeter cells equipped with photo-multipliers will be used for this test.

The high-voltage, monitoring and calibration systems have to be upgraded for two reasons. First, some mezzanines are too sensitive to radiation to bear the amount of particles that will be received during the Upgrade data taking. Secondly, the electronics also has to be adapted to the new GBT based slow control system.

A total of 144 boards have to be fabricated, these include 132 mezzanines that will replace old ones on the mother boards. The mother boards are kept from the current system. The upgrade to the new slow control requires that 12 boards are made to convert the optical signal from the counting room into an electric signal that will feed some of the mezzanines.

Prototypes of these boards have been produced and tested (see fig 9, right). The firmware is at an advanced stage. The PRR has been passed altogether with the FEB and the control board and no technical issue has been identified.

The production time for the 144 boards is approximately two months and the manufacturers have already been identified. The production schedule for those boards will follow the schedule of the FEB and control board.

The dismantling of the SPD/PS is being prepared actively. A plan has been defined and the procedure has been documented. The tools used for the installation of the detectors have been refurbished and will be reused. Some new equipment has been designed both for the dismantling and for the storage of some pieces of the SPD/PS.

3.3 Muon system

During the last few months, substantial progress has been achieved in the preparation of the new off-detector electronics for the upgrade. This consists of: a new readout board (nODE), equipped with its custom ASIC (nSYNC), redesigned to be compliant with a 40 MHz readout of the detector; new control boards; the Service Board (nSB); and the Pulse Distribution Module (nPDM), redesigned to be compliant with the new ECS/TFC system.

The PRR for nSYNC and nODE was successfully passed in mid October 2017. All the tests on the final version of the board have been completed successfully,

using the MiniDAQ2 setup, equipped with the PCIe40 readout board. The irradiation tests on a few commercial components have been carried out, showing no failures for the expected lifetime of the boards. The nSYNC/nODE production has been launched, and the first 20 complete boards are expected at CERN by April 2019. The delivery of the full production is expected at CERN between July and September 2019.

The final version of the nSB/nPDM boards has been produced and successfully passed its PRR on May 15, 2018 (EDMS 1974869). The production of the boards is starting now and we expect to have the pre-production boards (20 nSB + 2 nPDM) by the end of 2018. The nSB crate custom back-panel is also being produced and will be at CERN in February 2019. The delivery of the full production of the control boards is expected at CERN in April 2019.

Given the above production schedule, we expect to be ready to start the installation and commissioning of the first quadrant of the Muon Detector in April 2019. The first months of 2019 will be then devoted to the usual maintenance activity on the chambers (we will not change our MPWCs until the end of Run 4), and to the installation of the new optical fibers from the off-detector crates to the patch-panels, where the long cables from the surface are connected.

Once the first readout boards are installed the commissioning will start, which will consist of first checking the connectivity between the chambers and the nODEs by reading the signals with a portable PCIe40 crate down in the cavern. In a second phase, after summer 2019, the signals will be readout directly at the surface using the online farm and the standard readout chain. Furthermore, the design of the additional shielding behind the HCAL, which is made of tungsten, and of the new beam-plugs under HCAL and Muon station M2, made of lead, have been completed and reviewed within the collaboration. The price inquiries have been launched. We plan to launch the order before the end of 2019, to ensure delivery by spring 2019. This improved shielding will reduce the particle flux in the hottest parts of station M2R1 by about 60%.

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. The baseline relies on a large bandwidth bi-directional network interconnecting event-builder PC-servers and on the generic readout module, PCIe40 (PCI Express), embedded in each PC-Server.

The production of the PCIe40 module has started in March 2018. It is organized in four steps: a pre-series of 24 followed by three batches of 50, 280 and 343 modules respectively. The first two modules of the preseries were delivered end of July and the remaining 22 are expected at the beginning of October. The delivery of the first batch of 50 is scheduled at the beginning of 2019. PC-servers to host

the PCIe40s are under evaluation in close collaboration with industry.

The event-builder has to aggregate 40 Tbits/s. The baseline architecture relies on a 100 Gbits/s bi-directional network interconnecting 500 nodes. A recent contract signed by CERN-IT opened the door to a simplified architecture at an affordable cost. It is based on a very large switch routing event fragments directly to each event-filter farm node, in a similar way to the current LHCb and ATLAS event-builders. A hardware loan has been arranged to conduct tests in the coming months. Both architectures can house accelerators, like GPUs. A review of event-builder architectures is foreseen in April 2019, with external referees, and the final decision by the end 2019.

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. In parallel, simulations are used to study the performance of different topologies and to test ideas to improve the scalability. The simulation of the traffic as a function of the algorithm, network bandwidth and topology is now working for 500 nodes and gives confidence in the scalability of the baseline system.

A platform for integration and performance tests is under preparation. It corresponds to a vertical slice of the online system with two 48U racks located at the surface of LHC Point 8. The first rack houses the hardware for the event-builder and the second one the pc-servers for the event filter farm. The vertical slice setup will be ready by the end of October 2019. It is the first step to prepare the commissioning of the online system.

To house the event-builder and the event-filter farm, a new data centre will be be built on the surface at the LHCb experimental site. It is composed of six containers located at LHC Point 8. Civil engineering started in June. The delivery of the first two, for the event filter farm, is scheduled in October 2019. It is followed by two for the event-builder in March 2019 and the last two in September 2019.

The surface and the underground areas are connected via a backbone of long distance optical cables. A contract is now in place with the company producing them as well as optical path cords to interconnect sub-detectors to the backbone. A tender has also been launched to selected the company which will install the long distance optical cables.

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] laid out a solution for the Upgrade trigger, one that is highly performant and can potentially be operated within the timing budget. Many novel aspects proposed in the Upgrade TDR have already been implemented in the Run 2 high-level trigger.

The HLT1 fast reconstruction sequence have been migrated to the new computing framework. It mainly reconstructs long tracks, and primary vertices. The

relevant algorithms have been adapted to the task-based framework, including partial redesign and vectorization, and profiting from an an improved event model. The fast Kalman Fitter is not yet included. Two dedicated servers, bought in 2016, have been implemented to test and benchmark in realistic conditions the performance of the software that is being developed. Modern code profiling and clear benchmarking procedures have been also put in place.

Applying prototype transverse momentum and impact parameter requirements (800 MeV/c and 100 μm respectively) the performance obtained are encouraging and in these conditions a throughput of about 16 MHz is achieved. The benchmarking tools also allowed the identification of major bottlenecks, mainly in the decoding and the clustering of the UT an SciFI data. It is estimated that a proper optimization of these algorithms will lead to speed-up factors in the HLT1 code. Discussions between computing and backend experts started to explore different encodings of the data in the PCIe40 module in order to facilitate the HLT1 processing.

These works will continue in order to prepare the second bi-annual update of the trigger performance. The aim is to review the reconstruction strategy and the bandwidth division, as well as the performance for different architectures at the beginning of 2019.

5 Computing

The new task-based Gaudi framework developed in 2017 has demonstrated that the scheduling of a series of tasks for concurrent events using a multi-threaded approach works well. The memory consumption is also under control and significantly reduced compared to the multi-process case. Demonstrators built on top of the framework have also shown that a gain in performance will come from vectorization techniques, data access optimization and algorithmic improvements. The impact of recent work on these items on the HLT1 fast reconstruction sequence has been discussed in the previous Section.

In the last months, a new Gaudi scheduler was designed and commissioned to support control flow and trigger lines in a multi-threaded context. It is required to design and benchmark HLT2 lines. In addition, the geometry of the detector will be described by using the DD4Hep library instead of XML files. This integration is progressing.

A close collaboration between computing experts and physicists has been established in order to use in an optimal way the task based framework, the C++ language and vectorization libraries. The organisation of hackathons every two months is found to be very effective method for spreading the knowledge and tools required for Upgrade computing throughout the collaboration. Code reviews in gitlab have also become a very effective tool in this respect.

In April 2018, the *Software and Computing* Technical Design Report was submitted to the LHCC [9] and the *Computing Model* Technical Design Report will

be submitted in autumn 2018. The former presents the engineering aspects related to core software and distributed computing, while the latter discusses the offline workflows, resource provisioning and resource requirements. The foreseen trigger output bandwidth and the amount of Monte Carlo samples to be simulated pose stringent requirements on the offline resources. Given the significant increase in data volume with the upgrade experiment, the requests for the Upgrade are significantly larger than those achievable in a constant budget resource evolution scenario that has been employed so far for the current LHCb experiment. Mitigation measures are being considered, though they require the development of tools and associated person-power on a relatively short time scale. Mitigation measures can have an impact on the operations for both the experiment and the data centres or on the physics output of the Collaboration. The approval of the Technical Design Reports and the required computing resources is foreseen by the April 2019 Resources Review Board meeting.

6 Infrastructure

The preparation of the infrastructure that has to be ready for the Long Shutdown 2, to allow the current detector dismantling and the installation of the new subsystems, is well advanced. Most of the tooling and handling equipment has been renovated or re-designed and produced. Only a few storage structures have to be finalized during the next two months.

Assembly areas have been heavily used since summer this year and the SciFi team has started to assemble the prototype support frames. The neutron shielding has been delivered and stored in the assembly hall, and is ready for installation, .

The first module for the Data Centre is scheduled for delivery at the end of October. The civil engineering and infrastructure for the placing of the modules is advancing, although accumulating some delay. The installation of the transformers is planned to start very soon.

The contract for the long distance optical fibres has been signed and the offers for their installation have been recently received.

Projects for the dismantling and installation of detectors and their services have been divided in work packages which are close to being finalized. A workshop on installation was held in May with the participation of all the main actors, both from all participating institutes and from relevant CERN departments. The workshop was very successful and gave the required boost to this work in all projects.

Work on the first piping for the cooling systems have been finished, where the cavern was accessible during technical stops. The tendering of a common chiller for the SciFi, Upstream Tracker and VELO detectors has been delayed due to administrative reasons but is still on track.

7 Organisation and milestones

The upgrade activity organization has been revised in order to better match the needs of the current phase. The former Upgrade Planning Group, now Upgrade Detector Planning Group (UDPG), has been reorganised with the appointment of an Upgrade Coordinator, who is now chairing the UDPG, and the extension of the membership to the LHCb Electronics Coordinator. The UDPG membership consists of an Upgrade Coordinator (chair), an Upgrade Resources Coordinator, an Upgrade Data Processing Coordinator, an Upgrade Electronics Coordinator, as well as the management and a representative of the Physics Coordinator.

All the activities concerning the development of the fully software trigger have been reorganised and are now coordinated by the Upgrade Software Planning Group (USPG).

The UDPG and USPG meet regularly to review progress. Detector and software upgrade activities are organised within the existing Projects and working groups, to ensure efficient sharing of resources between operational needs and Upgrade work. The two exceptions are the UT and SciFi systems, where new Projects were created.

7.1 Milestones

Several important milestones have been achieved in the last six months, most notably the VELO Module PRR, which validated the overall fabrication of the VELO modules based on silicon microchannel substrates, and the start of production of the PCIe40 board, which is LHCb's main custom component in the Online system. The SciFi PACIFIC chip production was completed and delivered. The Upgrade LHCb Software & Computing Technical Design Report was submitted to the LHCC in May 2018 (CERN/LHCC 2018-007). Much attention is being given to the SALT ASIC which is causing concerning delays in the UT detector project. A thorough review with external experts was organized and a submission of the SALT version 3 is currently being prepared. The delays are considered by the project to be recoverable to allow installation in LS2 if this submission is successful.

A snapshot of the global status is given in Fig. 10. The overall delay of around six to twelve months with respect to original planning is mostly absorbed in the foreseen contingency and progress continues at a steady pace. Some areas of concern are being closely monitored and actions have been taken to mitigate the delays. An extensive 1.5-day Installation Workshop took place in May 2018, with broad participation from all the subdetector projects, infrastructure and support teams. It focused on removal, installation and commissioning. An LHCC in-depth review took place end of May 2018. Overall, the Upgrade project is on track for completion and installation in LS2.

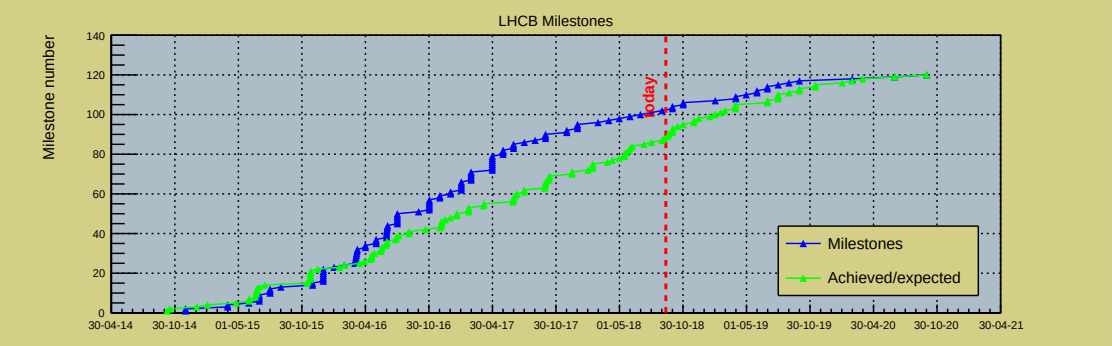

Figure 10: Snapshot of LHCb Upgrade milestones.

8 Future upgrades

As previously reported [10], the LHCb Collaboration submitted an Expression of Interest in future Upgrades of LHCb [11], to the LHCC in February 2017. This addresses consolidation and modest enhancements in LS3 (Upgrade 1b) and a significant upgrade in LS4 (Upgrade II). The collaboration was encouraged by the LHCC to continue studies, particularly providing detail on the physics case and the compatibility of the concept with the high luminosity LHC. Both of these have now been addressed.

A physics case document has been submitted and presented to the recent LHCC [12]. This was prepared by over 100 contributors and signed by the full Collaboration. The LHCb Upgrade II will fully exploit the flavour-physics opportunities of the HL-LHC and study additional physics topics that take advantage of the forward acceptance of the LHCb spectrometer. As an example of the physics reach the constraints on the Unitarity Triangle of *CP* violation in the CKM matrix from LHCb today and projected for after Upgrade II are shown in Fig. 11. Generically, the new physics mass scale probed, for fixed couplings, will almost double compared with the pre-HL-LHC era. The Upgrade II of LHCb will enable a very wide range of flavour observables to be determined with unprecedented precision, which will give the experiment sensitivity to New Physics scales several orders of magnitude above those accessible to direct searches.

A second important document for Upgrade II has also been submitted in this period [13]. This is the document prepared by the HL-LHC team who have studied the options for delivering the required luminosity for the LHCb Upgrade II. This concludes that "a range of potential solutions for operating LHCb Upgrade II...and permitting the collection of 300 fb^{-1} or more" have been identified.

At the time of writing the minutes of the LHCC are not available but the referees expressed their enthusiasm for the programme presented and discussed with us the timescales for proceeding to Technical Design Reports.

The first LHCb Upgrade Resources Board dedicated to Upgrade II was recently held for the national representatives. R&D plans are taking shape in many of

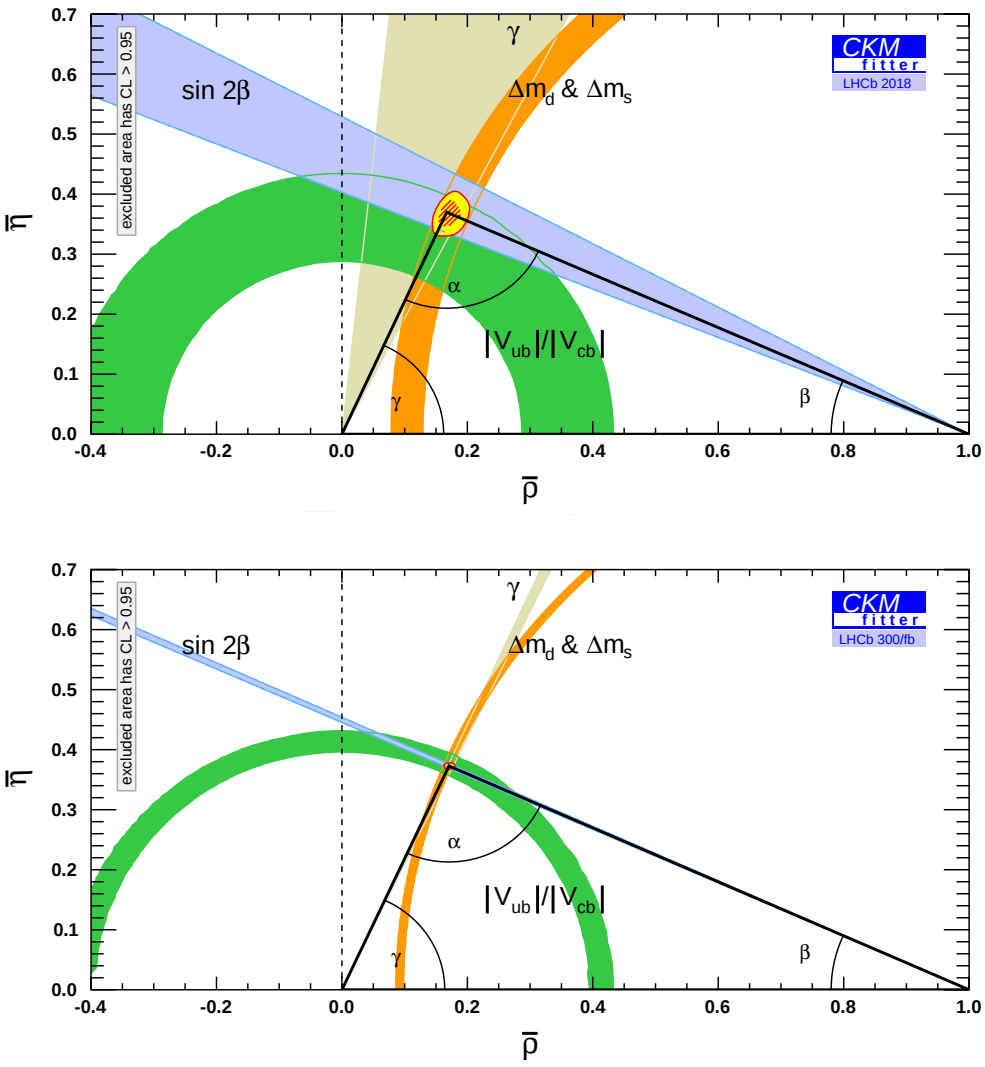

Figure 11: Evolving constraints on the Unitarity Triangle of *CP* Violation in the CKM matrix from LHCb measurements and lattice QCD calculations alone, with current inputs (top) and the anticipated improvements from Upgrade II (bottom).

the contributing nations and a number already have funded programmes with relevance to Upgrade Ib/II. As reported in the collaboration matters, the first technical associate group has recently joined the Collaboration and have done so in order to work on R&D for Upgrade II. The fourth in the sequence of dedicated workshops on Upgrade Ib and Upgrade II is taking place in Amsterdam in spring 2019 and, in light of the positive LHCC discussions, will focus on the detector design and technology.

9 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 detector 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 making major progress. Serial production is ongoing for all sub-detectors and first detector components are being delivered to CERN. Major contracts have been placed and spending of CORE funds is proceeding for all of the sub-detector components. Most of the remaining funds for sub-detector construction will be spent in 2018 and 2019. The Upgrade project continues to evolve 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] *Letter of Intent for the LHCb Upgrade*, CERN-LHCC-2011-001. LHCC-I-018.

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

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

[4] LHCb collaboration, *LHCb PID 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. 2 to the Memorandum of Understanding for Collaboration in the Construction of the LHCb Detector Upgrade of the LHCb detector Sub-Detector Systems*, CERN-RRB-2014-105.

[9] LHCb collaboration, *Upgrade Software and Computing*, CERN-LHCC-2018-007. LHCB-TDR-017.

[10] G. Passaleva, *LHCb Upgrade Status*, Tech. Rep. CERN-RRB-2018-053, CERN, Geneva, Mar, 2018.

[11] LHCb collaboration, *Expression of Interest for a Phase-II LHCb Upgrade: Opportunities in flavour physics, and beyond, in the HL-LHC era*, CERN-LHCC-2017-003.

[12] LHCb collaboration, *Physics case for an LHCb Upgrade II — Opportunities in flavour physics, and beyond, in the HL-LHC era*, `arXiv:1808.08865`.

[13] I. Efthymiopoulos *et al.*, *LHCb Upgrades and operation at* $10^{34} cm^{-2} s^{-1}$ *luminosity A first study*, .