Status of the LHCb upgrade I

CERN-RRB-2019-039

17 April 2019

1 Introduction

The installation of LHCb Upgrade I started immediately after the end of 2018 LHC run, indeed the final beams were dumped of Run 2 at 4:30am on Monday 3rd December and work started on Upgrade I at 8am. The dismantling of the previous detectors is well under way and on schedule. 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 Research Board. At the last Research Board session the Software and Computing and the Computing Model TDRs [7, 8] have been approved as well. 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 [9] and sub-system items [10], respectively. A new sub-project, the Real-Time Analysis Project has been created to organize the complex software developments for the upgrade trigger.

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.

2 Tracking system upgrade

Over the last six months the tracking system has made considerable progress. The Vertex Locator (VELO) is approaching the start of module serial production and the first pair of two RF boxes is available. The Upstream Tracker (UT) is still delayed by the finalization of the SALT ASIC but encouraging results have been obtained with the third version of this chip and other components are being prepared for construction. The large Scintillating-Fibre (SciFi) Tracker is about to be assembled at CERN, in the surface building. A more detailed 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)

Front end ASICs, Sensors and Tiles

The sensor order had already been completely fulfilled last year, and a second production order of the second VeloPix submission was successfully launched and fulfilled in January. Ten out of the twelve wafers have been probed, exhibiting a yield of about 50% slightly lower than the first production round but still well within expectation. The wafers have been sent to the bump bonding company for tile production and the production and reception testing is advancing steadily, with about 70% of the production now fully in hand and with a bump bonding production yield of greater than 95%. The long term storage of tiles in the gel packs will now be avoided and the tiles have been shipped to the assembly sites where they will be transferred to a more gentle long term packaging solution in conditioned atmosphere. The results from the tile testing are automatically uploaded via a tarball file to the database and an automatic grading and reporting procedure has been implemented. This model has been extended to many other production components with the plan to be completed by the start of module production.

Hybrid Circuits

The Front-end hybrids and GBTX hybrids continued to give worries, mainly at the level of the unsatisfactory flatness characteristics and a poor quality of component mounting at the company. The production was relaunched with a new vendor which has yielded very good results, and 80 production quality Front End hybrids (out of a final 204 required for production) are now in hand and are being used for the final module construction verification. The GBTX hybrid underwent some minor design changes to address issues which arose during module construction, in particular the need to simplify the routing of temperature probe information from the module, and will be delivered shortly. The interconnect and High Voltage cables are under order and are expected to be delivered shortly.

The electrical readout chain was been fully assembled for a full module. This includes the OPB (Opto Power Board), vacuum feedthrough, datatapes, low voltage cables, control and VeloPix hybrids and sensor triples. The set up has been used to develop the control panels and testing procedures for the final modules.

Cooling Susbtrates

The microchannel substrate production has some delay. The original production was launched in two Lots; Lot 1 and Lot 2. Lot 1 proceeded with a good yield at the silicon wafer bonding step but then hit a major problem in the final, subcontracted plasma dicing step. This resulted in a very large loss of potential production Grade A substrates due to issues with the surface and edge quality, although the substrates are perfectly useable for connectorisation and prototype module production. There was also a significant loss of time while the problems were investigated together with alternative process flows. Attempts are still underway to recover some pieces from the Lot 1 production (a final yield of 10% can be anticipated) and a Lot 3 production has been launched with high priority, which has just passed the bonding step and is showing promising results. The plasma dicing step has been moved from a subcontracted solution to an in house production at the same company performing the silicon microfabrication of the substrates. On the current planning schedule the microchannel order is expected to be fulfilled in entirety by March 20th.

There were also some challenges in the soldering attachment of the inlet/outlet pipes and connector manifold to the silicon substrate. The voidless fluxless soldering procedure showed an issue at the pretinning step of the connectors which meant that it was too risky to proceed to microchannel soldering. The issue was tracked down to the introduction of a detergent for cleaning the connectors at the company and the problem has now been eliminated. However the supply of connectors from the institute still has only a small number of potential spares, and investigations are ongoing to launch a backup production.

The exercise of soldering good quality connectors to the sub-Grade-A microchannel substrates in hand has gone ahead smoothly, yielding to this date 14 soldered pieces and a well debugged production process.

Readout Chain

The first large production batch of the data cables were delivered from the vendor at the end of January. The quality of the signal transmission lines and the mechanical quality was good, however there were unfortunately issues with the overetching of some traces and connection pads. This has resulted in the rejection of a large fraction of the batch, and the company has relaunched production. The connector mounting and production testing of the remaining tapes has gone ahead successfully. The Vacuum Feedthrough Boards have been fully and successfully produced, together with the component mounting. The installation of the LV ca-

bles is going ahead at an external company and the fully assembled pieces will be tested in the second week of March. The final pieces will be assembled into the vacuum flanges at CERN before shipping to the assembly sites. The six OPB prototype boards produced up to now have been in constant use at the module assembly sites and at the electrical readout chain test sites. Minor bug fixes and improvements will be implemented and a new mechanical air cooling design has been successfully integrated. The tender is currently underway for the final order.

Test Beam

Three pre-production modules (two from Nikhef and one from Manchester) were brought to the H8 beam line in the SPS north hall and installed upstream of the Timepix3 Telescope. The system was readout by the MiniDAQ2 with a special firmware version in order to combine the information from one ASIC in each module. The collected data was analysed using an adaptation of the Kepler/Gaudi software which allowed the verification of tracks and timestamps to those of the Timepix3 telescope. Preliminary results showed that tracks could be reconstructed matching segments from the VELO modules and telescope.

Modules

There have been small changes to the module construction procedures. A new star shaped pattern for the glueing of the tiles to the substrate has been adopted to optimise the coverage and a more elastic glue has been adopted for the hybrids to ensure zero movement due to the CTE mismatch between the kapton circuits and the silicon substrate. The final modules have shown full electrical efficiency and extremely small deviations from the vertical when being cooled. The mechanical precision of the tile placements and glue thicknesses achieved by both module assembly sites has been shown to be well within tolerances. There have been a recent series of sequential problems which continue to delay the launch of the module production. This has included the lack of availability of the final production components, as described above, a problem which is now solved. In addition both assembly sites have suffered from issues in the cooling systems (failure of pumps, need to install pressure sensors for the full module test) and these problems are in the final stages of being addressed. Further final details such as the adaptation of the bond wire programmes to adapt to the final glue layer thicknesses are currently being implemented. Finally an issue has been seen with loss of communication to some modules when fully cooled down. This is being investigated in a parallel full slice test setup which has been established in addition to the module testing facilities at the construction sites to speed up the debugging. While there are no show stoppers foreseen the delay means that the module construction rhythm of up to two modules per week will have to be respected when the construction shortly starts.

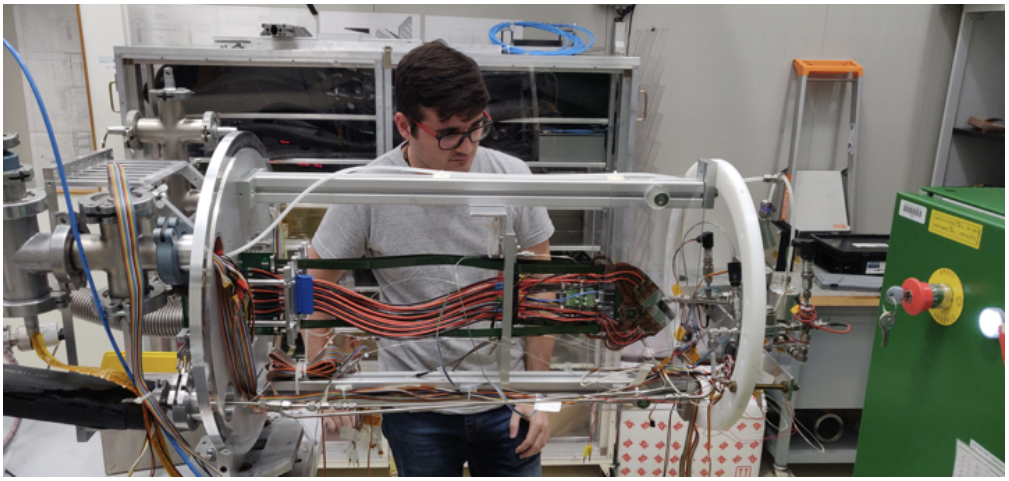

Figure 1: One of three complete module test setups. The setup illustrated here is capable of cooling and operating two modules in parallel and of checking the complete functionality of the system.

RF Foil and Mechanics

The mechanical assembly is proceeding very satisfactorily. The valve assemblies for the isolation vacuum and cooling manifolds have been fully delivered and characterised. The C side hood has been fully machined and the pipe runs are under construction together with the module handling tools and installation frames. The two detector bases have been successfully machined to the required mechanical tolerances and delivered to the assembly site.

The RF foil project is proceeding very well. The first A side box machined to a nominal 250 micron thickness was completed last year and the matching C side box was recently completed. Both boxes respect the mechanical tolerances for both position and thickness, although the C side box shows slightly more thickness variation. The leak tests of both boxes have been successful. Work is continuing on the next box pair, and in parallel studies into the validation of a possible chemical etching to further reduce the thickness to 150 micron are proceeding on the pilot box.

2.2 Upstream Tracker (UT)

The main challenge in the construction flow of the Upstream Tracker (UT) remains to absorb the delays related to the SALT ASIC. A new version of the chip has been produced to address the issues encountered when the ASIC is connected to a full size sensor. In parallel, a strong effort is being put in place to refine and optimize the instrumented stave construction and to optimize the installation and commissioning plans of the UT to mitigate the significant delays in achieving a successful design of the SALT chip.

Figure 2: The first complete pair of RF boxes, ready for coating and installation.

Silicon sensor-hybrid module

The type A sensor production is continuing according to plans. All the type A sensors have been received and two thirds have already been tested. The identification of surface spots (small scratches or residues) is performed with a dedicated pattern recognition which makes the quality assurance step considerably less effort intensive. A pre-series of the smaller Si sensors that are needed for the innermost portion of the detector has been evaluated and the silicon sensor production is expected to be completed in the next few months.

A second iteration of the hybrid design has taken a considerable amount of time to be implemented successfully, due to some design changes requested by the SALT designers that made the hybrid construction and assembly more challenging for the company used. Eventually a simpler solution was identified that is expected to be optimal for the new SALT version and a number of prototypes of this design are currently being manufactured by four different companies.

In parallel, the mechanical aspects of the hybrid production are being studied. Panels containing eight hybrids with associated test circuitry are currently being used to finalize the SALT mounting and testing procedure, as well as the integration in the silicon-hybrid modules.

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 implemented in a submission for an engineering run in late June 2017. Unfortunately laboratory tests of the received version-2 ASICs uncovered several new problems that make these ASICs not suitable for the production modules. Careful studies of the performance in laboratory tests, and in a brief test beam run allowed a diagnosis of the design flaws. The key issue with the SALT ASIC appeared 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 were implemented for a SALT version-3 ASIC. This version-3 ASIC has recently been received, focused ion beam corrections have been performed on some chips to correct an issue that prevented a digital output from the chip. The version-3 ASIC is currently being evaluated. Preliminary results indicate that the main instability in the performance has been susbtantially reduced. Figure 3 shows the ASIC response to a calibration signal equivalent to a minimum ionizing particle most probable energy deposition in 320 μm silicon, and the total and intrinsic chip noise, measured on a dedicated test board hosting a SALTv3 chip wire-bonded to a full size sensor. The pulse shape is reconstructed by scanning the sampling time in a time window containing the injection time.

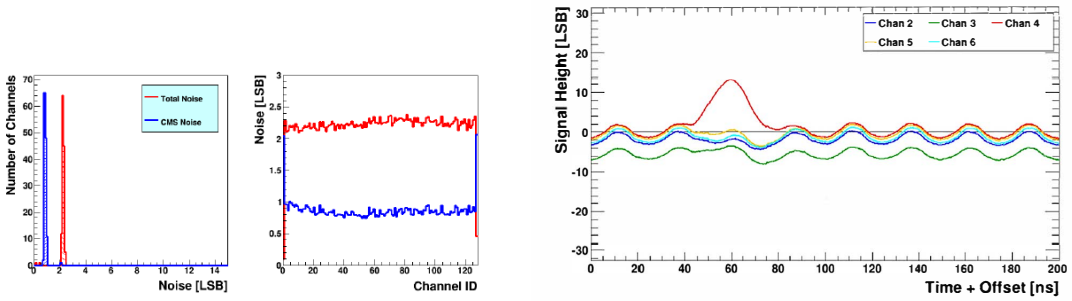

Figure 3: Preliminary performance assessment from SALTv3 wire-bonded to a full-size sensor: (left) total and common-mode subtracted noise as a function of the channel number, (right) time scan of individual channels: a calibration signal equivalent to one minimum ionizing particle in 320 μm is injected in channel 4, the other signals represent the baseline as a function of time for channels without any signal injected.

The first results of these investigations are encouraging and systematic tests of the SALTv3 are ongoing.

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 the manufacturing, 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. A recent set-back was the failure of the selected company to deliver acceptable cables, despite their pre-production samples being of good quality. A new company has been identified and the first samples will be evaluated during week 10 of 2019.

The production of the near detector electronics is starting imminently. The complexity of the backplane, which provides connectivity between the on-detector electronics and the other electronics infrastructure, is the main challenge. As the various electronics components of the project are becoming available, a system test of a first instrumented stave is planned for the spring of 2019 (known as the 'slice test'). This represents a delay of about six months with respect to the original plan, due to the complexity of the infrastructure needed. The final PEPI electronics is included, as well as a prototype of the backplane, thus the first step in the slice test will be a validation of all the components of the near-detector electronics and the overall UT system design.

Mechanics

The bare stave production was completed in November 2018. 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 sets of fixtures have been prepared to process two stave faces in parallel.

In light of the delays in front-end electronics, the instrumented stave assembly has been reconfigured to encompass three operations in parallel: the gluing of flex cables, silicon-hybrid module assembly and module mounting on the bare staves. Fixtures for module assembly and module mounting are currently being designed based on the prototype hybrid panels. Multiple versions of these fixtures will be fabricated to be able to process the components needed for the staves in about one working day.

Integration in the experiment

The design of the frame and box that will support and enclose the UT detector is progressing. In parallel, a test box capable of hosting two staves is being fabricated. Its first use will be during the slice test, which will also allow us to perform detailed tests of the detector operation at different temperatures.

A workshop was held in October 2018 to review the various steps of integration and commissioning and a detailed resource-loaded schedule has been prepared and is being acted upon. The slice test is the first phase of the installation and commissioning procedure: on the basis of this experience the stave mounting and testing procedure will be optimized and start in earnest in the Fall of 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

The serial production of the fibre mats has been finished. Including spares, 1531 fibre mats have been wound by the four production sites (Aachen, Dortmund, EPFL, Kurchatov). Modules have been produced at two production centres, Heidelberg and Nikhef, and Fig. 4 shows the production as a function of time. The target of 140 modules has been reached in early 2019.

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 3500 assembled detectors have been received and have been tested successfully. The tested detector assemblies have been shipped to Nikhef where they are being mounted on cold-bars. The remaining 1200 detector assemblies will be arriving at EPFL in March. After testing and grouping, these detectors will be ready for mounting on the cold-bars in the course of April 2019.

ASIC and read-out box

The readout ASIC (PACIFIC) has been produced and packaged. The complete sample of 25000 chips have already been tested. The best chips have been selected and have been sent to Beijing for the production of the PACIFIC Carrier Boards. The first 750 Carrier Boards have arrived at CERN and are ready for use.

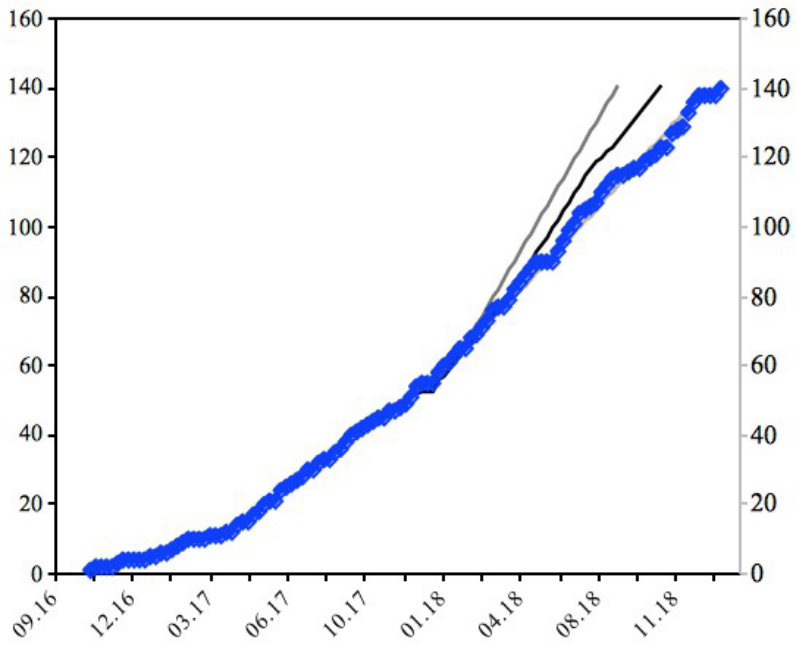

Figure 4: Number of produced modules as function of time.

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. Before launching the full production, the final production version of the boards have been successfully tested by operating a full readout chain, including SiPMs and fibre modules. Meanwhile, the production of the Cluster Board is in full swing and 500 boards have already been produced. For the Master Board the first boards are available and the production of the remaining 500 boards have started.

The front-end boards will be mounted on cooling frames and form the so called readout-box (ROB). First assemblies have been tested and installed on the prototpye C-frame (see below). The mechanical components for all boxes of the full detector have been delivered by the industrial supplier and the assembly of the first serial boxes has started in Clermont-Ferrand.

The ROBs will be mounted on water-cooled aluminium blocks to ensure the cooling of the electronics. The aluminium-blocks and also the water-pipes will be integrated into the C-frame structure. All water-cooling components, blocks and pipes for the full detector, have already been produced.

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 is provided by the cold-box enclosure. The production of the cold-boxes is in full swing and about 30% of the necessary cold-boxes have been finalized.

The cold-boxes will be mounted on both ends of the fibre modules before installation. The flex-cables of the SiPMs will later be connected to the front-end electronics. The module finishing, i.e. the mounting of the cold-boxes onto the modules, has started and the first 12 modules have already been finished and tested. Fig. 5 shows a finished module with mounted cold-box.

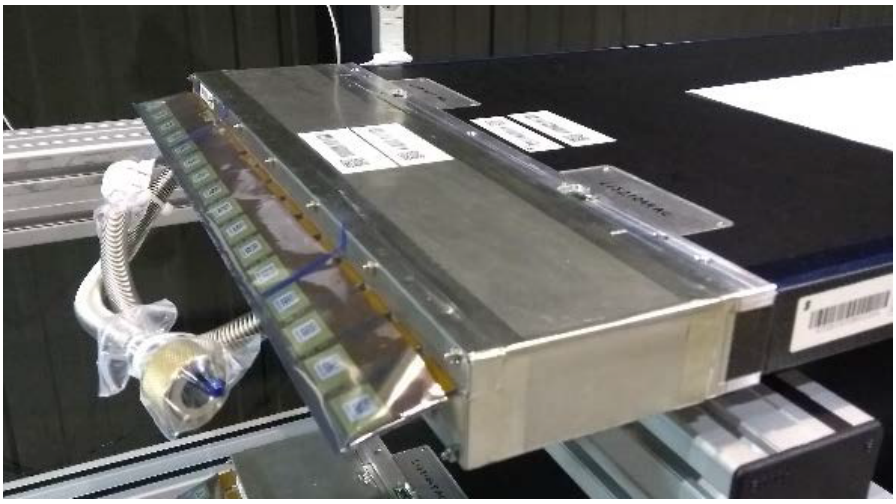

Figure 5: Photograph of a finished module with mounted cold-box. The two vacuum insulated Novec connections and the 16 SiPM flex cables are visible.

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 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.

A first prototype C-frame has been built in the new construction hall SXL8 at Point 8 (see Fig. 6) and has been equipped with all services (water-cooling, vacuum insulated Novec cooling, dry-gas flushing lines, low and high voltages),

with finished modules and with readout electronics (ROBs). The prototype C-frame is serving as a test-bed to develop the necessary readout, commissioning and debugging tools.

The ordering and production of the mechanical components of the final C-frames turned out to be more cumbersome than anticipated and led to a two month delay with respect to the planning. Meanwhile all components for the first serial C-frame have been delivered by the industrial producers and the assembly of the first SciFi detector frame has started. The components for the following 11 C-frames are being ordered. We foresee that the first six fully assembled and tested C-frames are ready for the installation in November 2019. The remaining frames should be finished in March 2020, when their installation is scheduled. Fig. 6 shows the construction hall SXL8 at Point 8 equipped with the mechanical structures to carry up to four C-frames.

Figure 6: Left: Photograph of C-frame prototype showing four mounted modules, with cold-boxes and ROBs. Right: Photograph of the assembly hall equipped with the support structures to carry the C-frames during assembly.

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 in the assembly hall.

The prototype C-frame has been used to successfully test the services at the system level: The vacuum system reached a level of 2×10^{-4} mbar sufficient to guarantee a good insulation of the Novec pipes; By flushing the cold-boxes with dry-gas a dew-point of $-45°$ was reached at a flow rate of 0.5 liter per hour; In a

first Novec cooling-test a temperature of -30° was reached for the SiPM without any build-up of ice.

All other service systems (high-voltage and low-voltage supplies, data-acquisition system) have also been commissioned for the prototype.

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

3.1 RICH

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. 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 MaPMT is established as the baseline technology for the RICH photon detector, read out by a customized ASIC named CLARO. The order for the MaPMTs was placed in 2015, the pre-series arrived and was accepted in April 2016 after Quality Assurance (QA) tests. At this time, the full production has arrived and has been qualified. The collaboration with the producer are excellent and a quick feedback is provided and production adjusted in case of issues.

The MaPMT, CLARO, front-end electronics and system integration have been tested in test-beams and radiation areas. All results to date are satisfactory. Components have been/are being produced. Notably, the full CLARO ASIC production has been received and the chips have been 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 is adapted for the RICH specifications and production has been started.

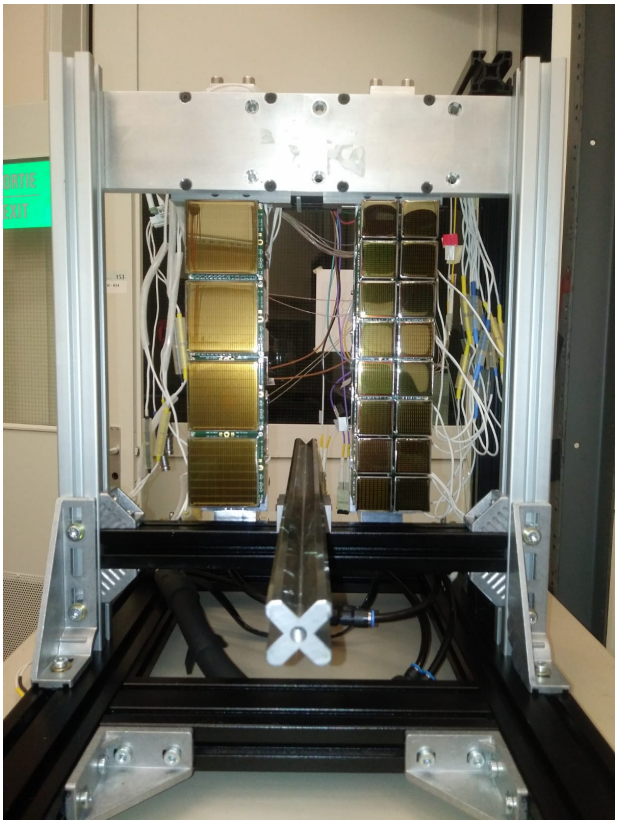

Figure 7: Test setup for the Photon Detector Module. Both MaPMTs types are used in the test

Mechanics, mirrors

The photodetector assembly, including the MAPMTs, all on-detector electronics and ancillary systems, is common to RICH1 and RICH2. All components are going through the production phase 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

Prototypes of the system, 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 to October 2018. For the first time in LHCb, we have successfully tested a full Photon Detector Module in combination with the Mini-daq 1 and 2 DAQ modules (Fig.7). We are now operating the system with two Mini-daq 2 DAQ modules to prepare the important activity of assembly and commissioning of the RICH1 and RICH2 columns, which will take place at CERN in the ComLab.

A complete Photon Detector Module (PDM) with its compliant DAQ was installed in RICH2 to take parasitic test data during the whole 2018 Run II.

3.2 Calorimeter system

The upgrade of the calorimeter system consists 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 partially compensate for the gain reduction by a factor 2.5. The remaining factor 2 will be used to extend the dynamics of the calorimeter system and to extend the physics case to some new topics. 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 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 but the version giving the largest transverse energy range of the calorimeter will be used for the production of the Front-End boards (FEB) on which the ICECAL will be soldered.

The FEB has been thoroughly tested and passed its PRR. 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 digital processing has been tested regarding the time adjustment of the signals, the numerical treatment in the FPGA (pedestal subtraction, low level trigger processing, calibration, etc...), the data exchange between components (proper sampling by the GBTX) and between boards in different slots (through the backplane) or between two crates (data exchange with cables), etc... The schematics of the FEB is frozen and the production should start in April, the market survey being closed since October 2018. The delay was mostly due to a shifting delivery time of the FPGA that equip the board and that we should receive within a few weeks.

At present the activity on the FEB is focusing on the firmware and the control firmware with WinCC. 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. The present prototype is satisfactory and considered to be the

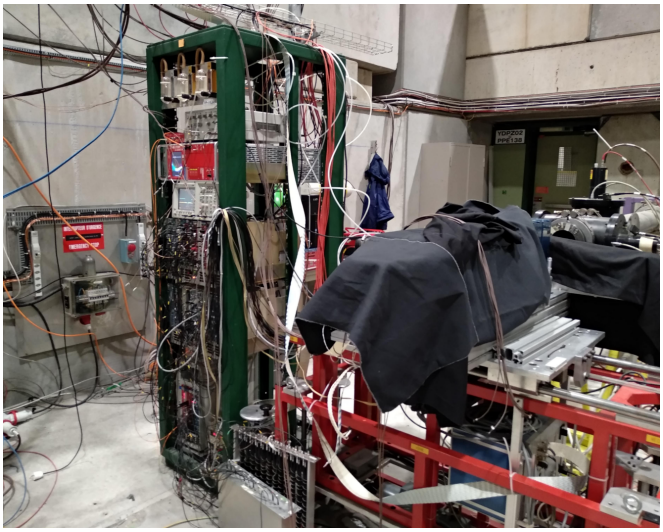

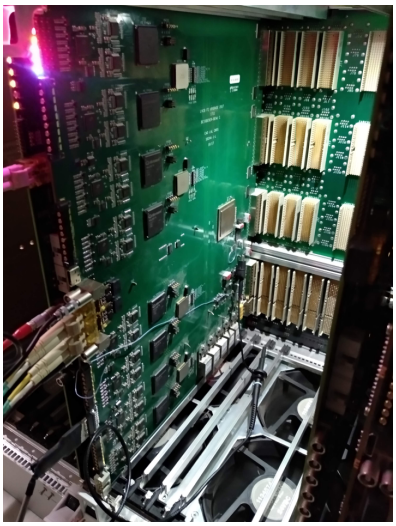

Figure 8: Left: picture of the test beam that took place in October 2018 at CERN on the SPS. A beam of electrons was targeting four standard calo modules (under the dark cover). The rack on the left contains the trigger system (based on two scintillators located between the modules and the end of the beam pipe) and the high-voltage for the Cockcroft-Walton bases used to polarize the photo-multipliers. Right: the front-end board prototype during the test, connected to the modules (inputs) and to the MiniDAQ (output) with four optical links.

last one. The board has been tested altogether with the FEB. No market survey is necessary for the control board production. We still plan to organize a final test in March together with a calibration board as these two systems exchange commands. Then, the production will start. The company has already been chosen.

For the upgrade, the crates and MARATON power supplies must be modified. The power supplies are being modified and the backplanes of the crates will be adapted soon. The procedure has been tested in the past and is now well-defined.

A realistic test of our FEB and control board in a final and modified crate has been done at CERN in October 2018. We used an electron beam at energies ranging from 20 to 120 GeV. Several calorimeter cells equipped with photo-multipliers have been used for this test and the electronics showed a good behaviour (see Figure 8).

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

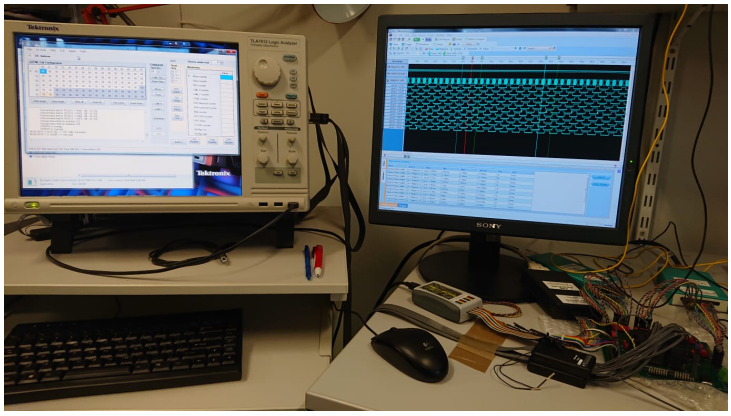

Figure 9: Test bench of the nSYNC ASIC.

that will feed some of the mezzanines.

The PRR, based on final prototypes, has been passed altogether with the FEB and the control board and no technical issue has been identified.

The 132 mezzanines have been produced and cabled. The tests should start in March and last one or two months. A specific test bench has been designed for this purpose and is ready. The firmware is at an advanced stage. The 12 boards have also been produced and the tests will be done in parallel with those of the mezzanines.

The dismantling of the SPD/PS was carefully prepared and started already. A plan has been defined and the procedure has been documented. The tools used for the installation of the detectors have been refurbished and are being reused. Some new equipments have been designed both for the dismantling and for the storage of some pieces of the SPD/PS.

3.3 Muon system

In the last months there has been a substantial progress in the production and test of the new off-detector electronics for the Muon upgrade. This consists of a new readout board (nODE), equipped with four custom ASICs (nSYNC) redesigned to be compliant with a 40 MHz readout of the detector, and of new control boards, the Service Board (nSB) and the Pulse Distribution Module (nPDM), redesigned to be compliant with the new ECS/TFC system.

The nSYNC have been delivered in Cagliari and tests are now ongoing (Fig.9). So far, about 200 of them have been tested and they show the expected performance. We need 80 nSYNC for the nODE pre-production and a total of 800 nSYNCs for the full nODE production. We expect to have all the needed nSYNCs tested in about one month.

The first 4 nODE boards arrived in LNF (Fig.10)and the tests show that the boards work as expected. We are now proceeding with the production of 16 more

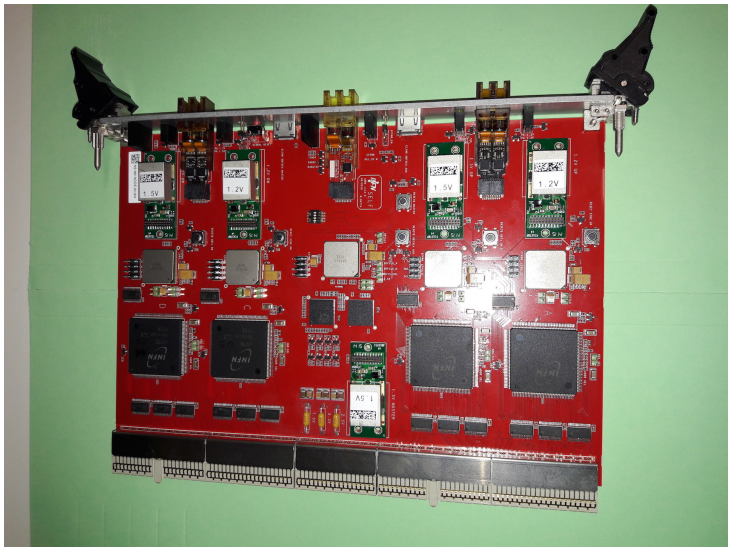

Figure 10: The nODE board.

nODE boards (expected in LNF by the end of March) and to their test. If tests are ok we'll give the green light for the full production (200 nODE) which is expected at CERN after the summer.

Two nSB and two nPDM boards have been delivered in Rome and tested. A few minor issues were found and promptly solved. Eighteen more boards are now under test and, if tests go well, the full production will start (120 more nSB boards and 8 nPDMs). The full nsB+nPDM production is expected to be delivered at the end of April.

Given the above production schedule, we expect to be ready to start the installation and the commissioning of part of the first quadrant (M2-M3) of the Muon Detector in April 2019. The commissioning phase will be split in two parts: in the first part we will check the connectivity between the chambers and the nODEs by pulsing the FEBs and reading back the signals with a portable PCIe40 crate down in the cavern. In a second part, after summer 2019, the signals will be readout directly at the surface using the online farm and the standard readout chain.

4 Online, trigger and reconstruction, computing

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 started in March 2018. It is organised

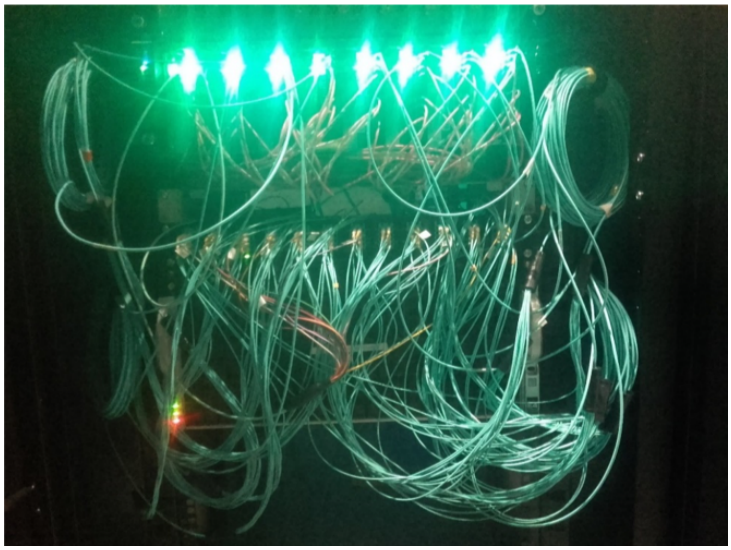

Figure 11: Vertical slice server fully equipped with eight pre-production PCIe40 cards.

in three steps: a pre-series of 24 followed by two batches of 330 and 343 modules respectively. Modules of the pre-series were delivered at the beginning of October 2018. The production was launched at the end of January 2019 and should end by November 2019. The specification of the PC-servers to host the PCIe40s are under preparation in close collaboration with ALICE and industry. The call for tender is scheduled in July for a delivery in November 2019.

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 contract signed by CERN-IT, in 2018, 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 system was loaned to us to conduct tests between the end of 2018 and the beginning of 2019. A review of event-builder architectures is foreseen in April 2019, with external referees, and the final decision by the end 2019.

A platform for integration and performance tests is running since October 2018. It corresponds to a vertical slice of the online system with two 48U racks located at the surface of LHC Point 8 (Fig.11). The first rack houses the hardware for the event-builder and the second one the pc-servers for the event filter farm. It is used for the acceptance test of the PCIe40 modules, for bandwidth measurements of the very large switch, for the study of HLT performance and for the Continuous Integration of software and firmware. 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 is built on the surface at the LHCb experimental site. It is composed of six containers located at LHC Point 8. The first two, were installed in October 2018. It is

followed by two for the event-builder in March 2019 and the last two in June and September 2019. The first two will be used by CERN-IT during the long shutdown 2.

The surface and the underground areas are connected via a backbone of long distance optical cables. All of them were delivered in February 2019 as well as all optical path cords to interconnect sub-detectors to the backbone. The installation of the long distance optical cables will start at the beginning of April 2019 and will last three months.

4.2 Trigger and real-time analysis

The Upgrade trigger [6] consists of a collection of identical software tasks running on the event-filter farm. All collisions are reconstructed then in real time with the best possible quality and then selected to be written for offline storage. This process is done in two steps. In the first one, HLT1, the fast reconstruction sequence is run in order to reduce the rate to about 1 MHz. Data are then stored locally waiting for the calibration and alignment constants. Once ready, the second step, HLT2 , performs the full reconstruction and selects collisions interesting for the physics. In this scheme, the full reconstruction is done once and never redone at a latter stages.

In order to reach this ambitious goal, the collaboration created the *Real Time Analysis* project. It is organized like a detector project and it is supervised by the Technical Board. The project is organized in six working groups covering data structures, reconstruction, calibration and alignments, selections, quality assurance as well as the possible use of accelerators. It will work in close collaboration with the Online and Computing projects. The organization of the project is in place since January 2019 and the people leading work packages identified. Work is in progress to define the main deliverables and milestones, a preliminary version of which were presented to the February LHCC review. The project continues previous developments and is currently reviewing the HLT1 performance, and aims to have a review of a realistic HLT2 by the end 2019 and a rehearsal of the full chain by mid 2020. In addition, "data challenges" using misaligned and miscalibrated detector simulations will already be launched in mid-2019 to make sure that the real-time data processing will be sufficiently robust to cope with a largely new detector from the start of Run 3 data taking.

4.3 Computing

The Gaudi multi-threaded framework constitutes the foundation on which algorithms are developed, integrated and executed in a concurrent way, with much reduced memory consumption with respect to the multi-process case. A new scheduler was recently developed, aimed at providing the foundation of a selection framework that will allow seamless integration of trigger lines. An overhaul of conditions data and detector description is under way in order to make these parts

compliant with thread safety and concurrency. Further work in this area shows that the optimization of data access and the rearrangement of data structures will enable the usage of vectorization techniques, thereby improving the software performance. All these developments are made in strict collaboration with the Real-Time Analysis project mentioned above.

A close collaboration between computing experts and physicists has allowed to use in an optimal way the software framework, the C++ language and vectorization techniques. Software hackathons, organized four times per year, are used for spreading the knowledge and tools required for Upgrade computing throughout the collaboration, and to review software design and quick prototyping. Code reviews in gitlab are effectively used as efforts are moving more and more towards the practical implementation.

In April 2018, the *Software and Computing* Technical Design Report was submitted to the LHCC [7]. The *Computing Model* Technical Design Report [8] was submitted in November 2018. The former discusses 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. Despite an increase of the data volume from the Upgrade experiment of a factor 30, the requested offline resources do not exceed those achievable in a constant budget resource evolution scenario by a large factor. The increase required is only by a few tens of percent in any given year and for any kind of resource. This was possible due to mitigation measures such offline filtering, that will allow to decrease the amount of disk space needed, and the usage of fast and parameterized Monte Carlo simulations. In February 2019, the LHCC referees approved these Technical Design Reports.

5 Infrastructure

Since the beginning of Long Shutdown 2, a number of activities related to detector cooling infrastructure have been carried out in the LHCb cavern, including the dismantling of the existing VELO, IT, and TT cooling plants, the upgrade of the existing RICH and demineralized water plants, and the dismantling and modification of long-distance transfer lines. The installation of the SciFi Novec plant is scheduled for April 2019, followed by the installation of the shared primary chiller in June 2019. The MAUVE CO2 cooling plants for the UT and the upgraded VELO are currently being pre-commissioned at the Meyrin site and are planned to be installed in the LHCb cavern in June.

The transformers for the new data-centre are operational, two modules that will house part of the online PC-farm were delivered last year and passed the acceptance test in February this year (Fig.12. The civil engineering work for the next four data-centre modules is well on schedule. Two modules will arrive in March and the remaining two in summer 2019. Storage areas of almost 500 m^2 are

Figure 12: Installation of the new data center modules.

ready for receiving material. The long distance optical fibres are already available at the experimental area, and their installation will start in April, aiming for a completion of work in August 2019.

6 Organisation and milestones

The upgrade activity is overseen by the Upgrade Detector Planning Group (UDPG). 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 all-software trigger are coordinated by the Upgrade Software Planning Group (USPG). The USPG membership consists of the USPG chair, representatives of the Computing, Online and Real-Time Analysis projects and of the Simulation working group, the Operation Coordinator, as well as the management and the physics coordination team.

The UDPG and USPG meet regularly to review progress of the various projects. 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. Exceptions are the UT and SciFi systems and the all-software trigger, where new Projects were created.

6.1 Milestones

A number of major milestones have been achieved in the last six months, most notably the submission of the Computing Model Technical Design report [8] in November 2018 and the completion of the SciFi module production early in 2019. The UT detector project is still on the critical path due to the difficulties of producing the SALT readout ASIC. The new design (version 3) was submitted in the fall of 2018 and delivered by the manufacturer in early 2019. First results are encouraging and, if the ASIC is considered acceptable, the UT project will move into instrumented stave production around June 2019.

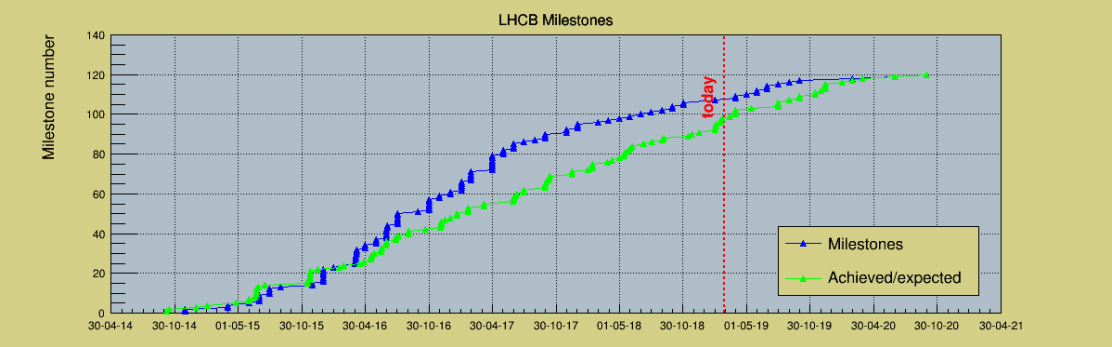

Figure 13: Snapshot of LHCb Upgrade milestones.

A snapshot of the global status is given in Fig. 13. 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. Overall, the Upgrade project is on track for completion and installation in LS2. Most subdetectors are under construction and in line with the installation schedule. Some areas of concern are being closely monitored (in particular UT) and actions have been taken to mitigate the delays.

After installation, the full detector will need to be commissioned and steps have been taken to prepare this phase. A commissioning coordinator has been appointed and regular commissioning meetings are taking place since the fall of 2018 with broad participation from all the subdetector projects, infrastructure and support teams.

7 Upgrade II

As previously reported a physics case document was submitted and presented to the LHCC [11] along with a document prepared by the HL-LHC team who have studied the options for delivering the required luminosity for the LHCb Upgrade II [12]. The LHCC has recommended that we proceed to the preparation of TDRs, and this was further endorsed by the CERN Research Board. The delivery of a

framework TDR on the timescle of around two years was discussed. The primary Upgrade II installation is planned to occur in LS4, with some preparatory upgrades and consolidation of the Upgrade I detector in LS3.

The performance of the LHCb Upgrade II has further been described in documents produced during the HL/HE-LHC Physics workshop process. This process was a year long study from the LHC experiments of the physics potential of the machine. This is summarised in short documents that were submitted to the European Strategy for Particle Physics Update process. Full reports from the working groups of the study are available in Refs. [13–17]. All working groups contained input on LHCb Upgrade II, with particularly dominant contributions in WG4 on flavour physics. We are aware that a number of the documents from contributing nations that have been submitted to the European Strategy for Particle Physics Update are also highly supportive of Upgrade II.

R&D plans on Upgrade II are taking shape in many of the contributing nations and a number already have funded programmes with relevance to Upgrade II. All subdetectors now have regular meetings related to Upgrade II studies. Internal notes are currently in preparation on the sub-detectors which propose changes during the preparatory phase in LS3. These will be reviewed in February and March by the Collaboration's Upgrade II Planning Group. The fourth in the sequence of dedicated workshops on Upgrade II is taking place in Amsterdam in spring 2019 and, in light of the positive LHCC recommendation, will focus on the detector design and technology.

8 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 [9] and in the Addendum No. 2 to the MoU for the Upgrade of the Sub-Detector Systems [10], 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 continues to progress as planned. Long Shutdown 2 (LS2) has started and the removal of old detector components is progressing according to schedule to allow for a timely installation of the new sub-detectors. Serial production is ongoing for all sub-detectors and a number of detector components have been delivered to CERN. All 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 2019 whilst the majority of the Common Project funds (in particular for the acquisition of the Computing Farm) are expected to be spent in 2020. 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 by the end of LS2.

References

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

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

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

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

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

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

[7] LHCb collaboration, *LHCb Upgrade Software and Computing*, CERN-LHCC-2018-007, 2018.

[8] LHCb collaboration, *Computing Model of the Upgrade LHCb experiment*, CERN-LHCC-2018-014, 2018.

[9] 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, 2012.

[10] 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, 2014.

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

[12] I. Efthymiopoulos *et al.*, *LHCb Upgrades and operation at 1034 cm-2 s-1 luminosity A first study*, .

[13] HL-LHC, HE-LHC Working Group, P. Azzi *et al.*, *Standard Model Physics at the HL-LHC and HE-LHC*, arXiv:1902.04070.

[14] Physics of the HL-LHC Working Group, M. Cepeda *et al.*, *Higgs Physics at the HL-LHC and HE-LHC*, arXiv:1902.00134.

[15] X. Cid Vidal *et al.*, *Beyond the Standard Model Physics at the HL-LHC and HE-LHC*, arXiv:1812.07831.

[16] A. Cerri *et al.*, *Opportunities in Flavour Physics at the HL-LHC and HE-LHC*, arXiv:1812.07638.

[17] Z. Citron *et al.*, *Future physics opportunities for high-density QCD at the LHC with heavy-ion and proton beams*, arXiv:1812.06772.