Status of the LHCb Upgrade

CERN-RRB-2016-33

April 2016

1 Introduction

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

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

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

2 Tracking system upgrade

Over the last six months the Vertex Locator (VELO), Upstream Tracker (UT) and Scintillating–Fibre (SciFi) Tracker projects have undergone several Engineering Design Reviews (EDRs) and are now preparing for the construction phase. A summary of recent progress and expectations for the next half year are given for each of the three sub-detectors of the tracking system.

2.1 Vertex Locator (VELO)

VELO sensor prototypes from two candidate suppliers have been evaluated in the testbeam after full, non-uniform irradiation with protons and after full irrradiation with neutrons. All the prototypes satisfy performance requirements in terms of charge-collection efficiency and resolution, and the timewalk behaviour has been characterised. The only cause for concern remains the additional charge collection beyond the pixel matrix in some n-on-p sensors, and possible additional currents drawn on the edges of the thinned n-in-n sensors, both from one supplier. The sensor pre-series orders addressing these issues and adding the rounded corner designs have been placed with both manufacturers. The sensors have been elongated by two pixels to accommodate the larger size of the new TSMC ASICs compared to those from the original foundry. As the HV performance with the new TSMC technology remains unproven, it is considered vital to maintain the n-in-n option, only currently available from one supplier, as a back-up in case sparking is experienced at high voltage.

The submission of the VeloPix, the frontend ASIC, has been delayed with respect to the initial plans and is now scheduled for April. All aspects are being verified with the expectation of a fully functioning ASIC after first submission. This will be critical for the module development work during 2016. All basic building blocks are ready. A detailed planning for the final steps exists and the submission date is realistic. The VeloPix test setups are in production. For the bare ASIC testing a chipboard design has been submitted and the SPIDR readout development is fully on track. The wafer probing setup is under development. For the DAQ slice test, which will evaluate the signal chain quality, all essential components have been designed and sent for manufacture. This includes the split VeloPix evaluation hybrid ("Velopix" hybrid and "control" hybrid), the vacuum feedthrough evaluation board, and the OPB prototype board, which were submitted following their respective EDRs in November. The final stage of the chain will be the VeloPix specific MiniDAQ system, which has already undergone several successful preliminary tests.

The microchannel cooling EDR was held in November and the referee report delivered. The referee report identified no show stoppers and made a number of recommendations to evaluate further and define the performance and technical specifications of the microchannel devices. Endurance and safety tests are underway for critical items in the VELO cooling chain, such as the creep tests recommended in the referee report. The VELO microchannel cooling tender is under preparation and expected to go out in a matter of weeks.

The module design has moved to a more narrow foot which eases the assembly. Simulations show that the new design has the required stiffness and also has the consequence of a more balanced and symmetric electrical cabling and layout on both the hybrid and the OPB. Glueing and handling jigs are under construction for the module assembly and are being excercised with full scale 3D printed module mockups. A detailed study has been made of possible gluing patterns which

provide the necessary strength but avoiding any trapped air pockets.

A new half-size RF-foil box prototype with 500 μm thickness (Fig. 1) has been completed using cast aluminium. Whereas the box was at first leaktight, small leaks developed after a 200 degrees bakeout. By using a dye-penetrant method these leaks were traced to a small number of micro-holes in the flat part of the structure, an issue which is believed to be due to the grain sizes in the material, and which will be eliminated with the use of "forged" aluminium as planned for the first full-size box. The box will be used for etching tests and work is going ahead on the construction of a new 250 μm thick machined box.

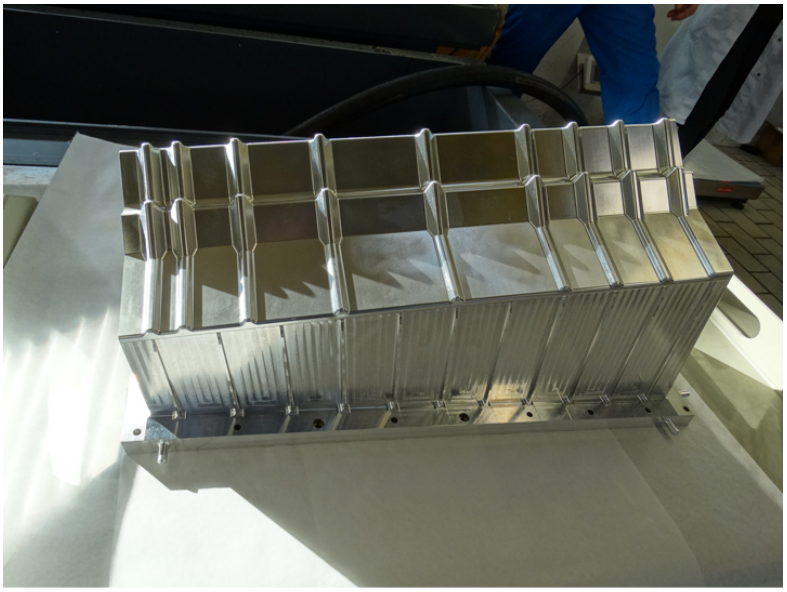

Figure 1: Half-size prototype of the VELO RF foil box.

2.2 Upstream Tracker (UT)

The UT project is progressing well. The transition toward construction, initially planned for the last quarter of 2015, is now underway. The small delay in the start of this phase is due to the time taken to construct two clean rooms, one for stave assembly and metrology, the other for sensor and electronics testing and wire bonding. The latter is now operational and the automatic wire bonder has been used to assemble hybrids with Beetle ASICs, ready for test-beam evaluation with prototype sensors in May 2016. Figure 2 shows a silicon sensor wire-bonded to Beetle devices mounted on hybrid circuits.

Hamamatsu have completed fabrication of the last batch of prototype type-A sensors, which are needed for the outer region of the UT planes. These sensors are implemented in p-in-n technology, which is expected to be adequate for the irradiation levels expected after the nominal Upgrade integrated luminosity is accumulated. A slight modification of the design will allow further studies of the embedded

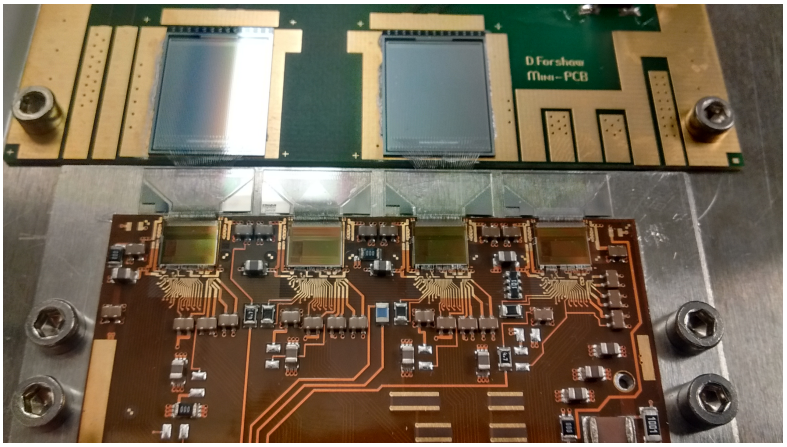

Figure 2: Hamamatsu miniature UT detectors wire-bonded to Beetle ASICs mounted on hybrid circuits to be used in May test beams. The hybridisation work has been done in the bonding laboratory commissioned for instrumented stave assembly.

pitch-adapters planned for these sensors. Analysis of the 2015 test beam data indicates that the"fan-in" pitch-adapter design has charge-collection problems, while the alternative "fan-up" design performs well in this regard. Miniature detectors featuring the two pitch-adapter design have been irradiated to different levels of fluence, up to twice that expected at the Upgrade. Test beam studies of their performance are planned for early May 2016. In addition, full size p-in-n modules will be irradiated at the IRRAD facility at CERN in early May and then tested with beams at the end of May. The findings of the extensive test programme will be scrutinised in a "Type A sensor preproduction readiness review (PRR)" planned for June 2016.

Considerable progress has been made in the R&D programme of the SALT chip, which is the ASIC that will be used for UT readout. The first eight–channel prototype (SALT8), submitted in February 2015, has been fully characterised. Minor problems have been uncovered, but all the main features tested were consistent with expectations from simulations. Electrostatic discharge (ESD) problems at some input pads prevented the designers from testing some of the ancillary structures fabricated in this submission. Therefore a second eight-channel SALT8b ASIC was designed and submitted for fabrication in November 2015. The characterisation of this device is ongoing. In parallel, a first iteration of the full ASIC (SALT128) is currently in an advanced stage of design. This design has been reviewed in the light of the results from the SALT8 tests and will be modified to include the latest ESD protection for the input pads. It will be submitted in May, which represents a slight delay with respect to the previous schedule. This delay will be offset by working in parallel on the design of the printed circuit board that will host SALT128 and provide the necessary infrastructure for the test system.

A complete test of an "electronics slice" is currently being prepared. This slice

will comprise the silicon sensors wire-bonded to SALT8 ASICs, mounted on flex hybrids, which are connected to the recently acquired prototype flex cables that route signal and power to the near detector electronics. The results of this test will be reviewed in an EDR planned for June 2016. In parallel, near-detector electronics boards, which provide data aggregation and the transition to optical signals, are being designed and prototyped. The performance of these boards will also be reviewed in June.

Two technologies to implement the SALT hybrids are under consideration. The baseline approach involves a flex circuit glued onto a ceramic substrate (BN) that provides mechanical support, thermal coupling to the cooling system, and electrical insulation. An alternative solution, combining the electrical and mechanical functions, based on an AlN or BN thick film hybrid, is also being pursued. A prototype hybrid circuit to host the SALT8 ASIC is being designed and will be incorporated in the slice test. The hybrid EDR will take place in autumn of this year.

The second round of prototype flex cables are currently being instrumented with the adapters necessary to interface with the SALT8 test board and the MINI DAQ based data acquisition system used for the slice test. Figure 3 shows one of these samples. Preliminary tests are very promising.

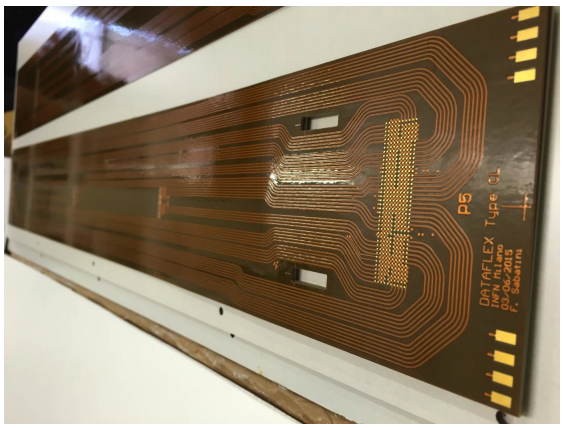

Figure 3: Prototype UT flex cable.

Significant progress has been achieved in the mechanical design of the detectors. Construction of the first two bare staves is imminent. In parallel, the assembly procedure of the instrumented staves is being defined and will be validated with a mechanical model. A PRR is planned for June. In addition, developments are ongoing for the CO_2 cooling system. A successful EDR of the cooling plant has taken place. The design of the manifold to distribute the cooling to the various staves is advanced. A recirculating TRACI CO_2 cooling system allows for studies of temperature profiles on mechanical mock-ups.

Finally, there have been achievements on the design of the detector enclosure and mounting frame. The partitioning of the space near the boundary between

inner detector volume and near detector electronics, just outside the detector volume, is being optimised.

2.3 Scintillating-Fibre Tracker (SciFi)

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

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

After a market survey for the scintillating fibres an invitation to tender has been sent to two suppliers. A contract with the cheapest supplier has been signed for 11,000 km of fibres. According to the contractual schedule, delivery will start in May 2016 and be completed in January 2018.

A persisting issue concerning the fibre quality is the uniformity of the fibre diameter. Occasional "bumps", where the diameter is significantly larger than the nominal value, are problematic for the mat assembly, as they lead to local defects in the winding pattern. A reliable diameter control procedure has been established allowing fast feedback to the supplier and the winding centres. Moreover different methods of removing or reducing the bumps have been developed and will be introduced in the production process.

The serial production of fibre mats and modules is currently being prepared in the four winding centres and the two module assembly centres. All winding centres are equipped with identical winding machines (Fig. 4) and the necessary tools. The first winding-centre (Aachen) will start the fibre mat production in April. A PRR will qualify the production process and quality assurance procedures as well as the readiness of the site. The PRRs of the other three sites (Dortmund, Lausanne and the Kurchatov Institute) will follow soon after. The PRR of the module assembly is foreseen for June.

SiPM arrays from two industrial suppliers are under investigation. The devices from one supplier already fulfill the requirements of the project. However, the most recent version, producted in 2015, which benefits from an optimised geometry and a higher photon detection efficiency, shows an unexpectedly large delayed cross-talk, which is currently being studied. The SiPMs from the other supplier exhibit various production problems. Both options will be pursued up to, or soon after the SiPM EDR, which will take place at the end of April.

For the readout of the 600k SiPM channels the SciFi group is developing a fast 64-channel ASIC, named the PACIFIC, which provides two-bit pulse-height information for every hit by using three adjustable thresholds. A 64-channel prototype chip (PACIFICv3), with full functionality, has been produced and tested; however several improvements are necessary and are currently being implemented.

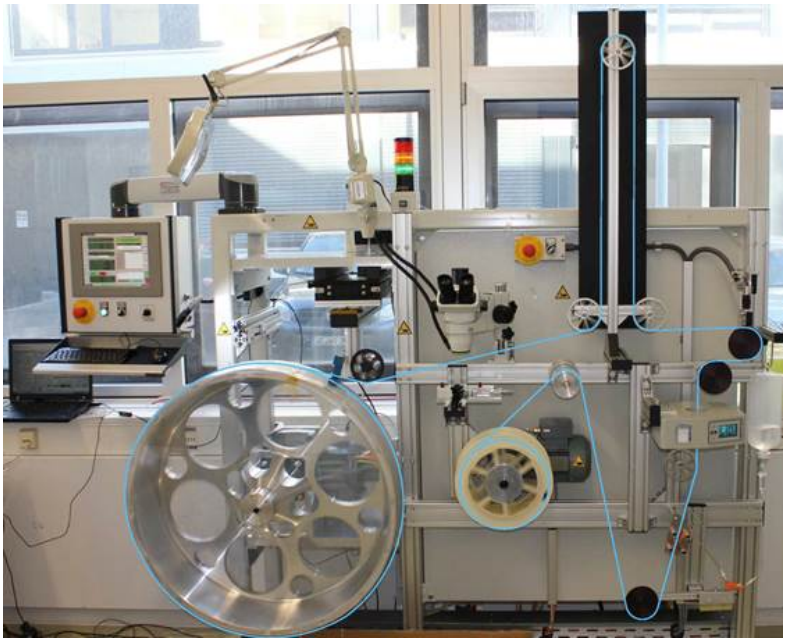

Figure 4: Winding–machine for scintillating fibre mats.

The submission of PACIFICv4 chip is foreseen for June 2016. A first version of the subsequent digital readout boards is also available. The full readout chain, consisting of a PACIFIC carrier board, a cluster-board to find the hits, and the master board, comprising the Slow and Fast Control as well as the optical links, is currently being tested. The EDR of the front-end electronics is also foreseen for the end of April.

The so-called Read-Out Box (ROB) is mounted at each end of the fibre module and houses the SiPMs in the cold part of the box ("cold-box") and the front-end electronics. A prototype cold-box is currently being tested and was successfully operated with the SiPMs maintained at the design temperature of -40°C.

In November 2015, a fibre mat was irradiated at the PS IRRAD facility over its full length, reproducing the expected irradiation profile for the lifetime of the detector. The performance of the irradiated mat has been studied in a test-beam using high energy protons and pions as well as in the laboratory with a beta source. While the test beam measurements reveal a smaller light yield than expected, the laboratory measurements confirmed the photon yield predicted by simulation. The discrepancy between the two measurements is under active investigation.

In summary, the SciFi project proceeds according to schedule: mat and module production is about to begin, the development of SiPM and front-end electronics will soon be concluded and the work on the detector's infrastructure is underway.

3 Track reconstruction

The proof-of-concept of the pattern recognition algorithms for the LHCb Upgrade was demonstrated in the VELO and Tracker TDRs [3, 5]. In the post-TDR phase the main focus of the tracking software efforts is to implement in detail the geometry, the digitisation and the pattern recognition according to the design choice of the hardware, which is now being finalised. In parallel efforts are ongoing to set up a time- and performance-optimised complete tracking and fitting sequence, which will serve as the basis for studies to be performed in the upcoming computing and software TDR.

3.1 Optimisation of SciFi tracker geometry and timing performance of the tracking algorithms

The single-hit inefficiency in the SciFi is the driving parameter of the tracking efficiency and ghost rate of the pattern recognition algorithms. It is mainly determined by geometrical inefficiencies such as gaps between modules ($\sim 1\%$), gaps in the middle and at the boundaries of each SiPM array ($\sim 2\%$) and the inefficiency which is introduced by a combination of photon detection and clustering thresholds ($\sim 1\%$). As a rule of thumb, a single-hit inefficiency of 1% induces a 2-3% tracking inefficiency, depending on the algorithm. In the latest round of production of SiPM arrays the manufacturer optimised the ratio of active to passive area. The corresponding contribution to the single-hit inefficiency went down to below 1%. This improvement helped the performance of all SciFi related algorithms to be raised back to the level quoted in the TDR. Work is ongoing to improve the single-hit and tracking efficiency still further by forming clusters over the boundaries of SiPMs.

A detailed photon transport model has been developed and tuned to match the test beam data. Applying this tuned model to the simulation of pp collisions in the LHCb experiment reduces the average photon yield compared to the version which was previously used for performance studies. This change, however, has hardly any impact on the single-hit efficiency. On the contrary, the lower photon yield reduced the number of hits in the event related to spill-over, which had a very beneficial effect on the ghost rate (-20%) and the timing (-30%).

As well as improving the realism of the SciFi simulation, much effort is also being invested in the study of potential modifications of the SciFi tracker to improve its performance. Among other ideas a y-segmentation of the inner, most populated modules and additional modules in this region have been investigated. The aim of these potential modifications is to increase the signal-over-noise ratio for true track candidates, and thus to allow cuts to be loosened in the pattern recognition, thereby improving tracking efficiency. It transpires that additional modules do not significantly improve the performance. Introducing y-segmentation appears to be a more promising approach; however the engineering challenges make this solution extremely hard to realise. The small gain of about 2% in tracking efficiency is

considered insufficient to justify this effort. Forthcoming studies will include an optimisation of the shape of the beam-pipe hole.

Detailed studies on the CPU consumption of the tracking algorithms have begun. Already factors of two to three have been achieved in speed, without any loss in tracking performance. Further work is ongoing in this direction in view of the studies required for the computing and software TDR.

4 Particle identification

The Particle Identification (PID) system of the upgraded LHCb consists of the RICH, Calorimeter and Muon systems. After a first series of EDRs all the sub–systems are now in the phase of preparation for construction.

4.1 RICH system

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

The MaPMT is established as the baseline technology for the RICH photon detector, read out by a customized ASIC named CLARO. The MaPMT, CLARO, front-end electronics and system integration have been tested in test-beams and radiation areas. All results to date are satisfactory.

The Photo-Detector Assembly, including the photo-sensors, all on-detector electronics and ancillary systems, is common to RICH1 and RICH2. The modularity and mechanical design of the full opto-electronic chain, its supporting mechanics, cooling and power distribution systems have been mostly defined and are undergoing final design and tests.

Current studies are focusing on finalising the details of the mechanical concepts of RICH1 and the common Photo-Detector Assembly, with the project entering the phase of detailed engineering design.

Successful EDRs for the CLARO, the Front-End Board and the Elementary Cell of the Photo-Detector Assembly were held in October 2014.

Prototypes of the system, made of several MaPMTs together with the accompanying front-end, digital and acquisition electronics and mechanics (Fig. 5), was evaluated in test beams during both 2014 and 2015. These test-beam studies were

a great success, paving the way towards the important milestones of 2016 and full production of the components.

The order for the MaPMTs has been placed and the first devices are arriving. The regional centers for the Quality Assurance are starting to become operational.

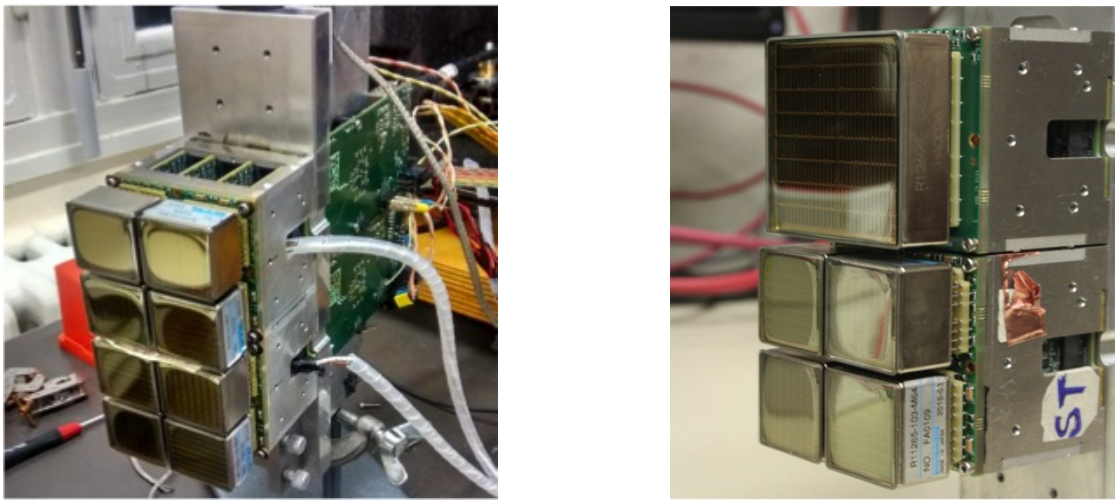

Figure 5: Left: Prototype of the photo-detector plane of the upgraded RICH used in recent beam tests. Right: prototype where both the 3x3 cm^2 and 6x6 cm^2 are mounted. The large MaPMTs are used in the outer region of RICH2.

4.2 Calorimeter system

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

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

The upgraded analogue electronics is based on an ASIC called the ICECAL. The ICECAL underwent a successful EDR during 2015. Several test beams took place in 2015 at CERN with the current digital board prototype and the ICECAL, using electron beams with energies ranging from 20 to 120 GeV. The performance of the electronics was evaluated in terms of linearity, spill-over, stability of the integrated signal sampling and several other parameters. The results are satisfactory

and match the design requirements. The radiation tolerance is now under study. A first measurement on a X-ray machine is planned for Q2 2016 at CERN. Further measurements after exposure to heavy ions will take place soon after, which will involve not only the ASIC but also the ADC candidate for the front-end board (FEB).

In parallel to the work on the analogue electronics, the design of the digital part of the FEB is progressing. Several decisions have been taken recently concerning the data format, which have allowed the number of optical links per front-end board to be reduced from five to four. This modification has led to a reduction in the cost of the FEB and also the number of acquisition boards in the counting room. A new prototype is being designed and is expected to be as close as possible to the final version. It will be equipped with the standard GBT-X, GBT-SCA, VTTx and DC–DC chips from CERN, and the chosen SEL/SEU and dose-safe FPGA. The schematics of the next FEB prototype have been finalised and the board is expected to be ready in May 2016. This new board will be equipped with the ICECAL and tested at CERN in an electron beam during the summer.

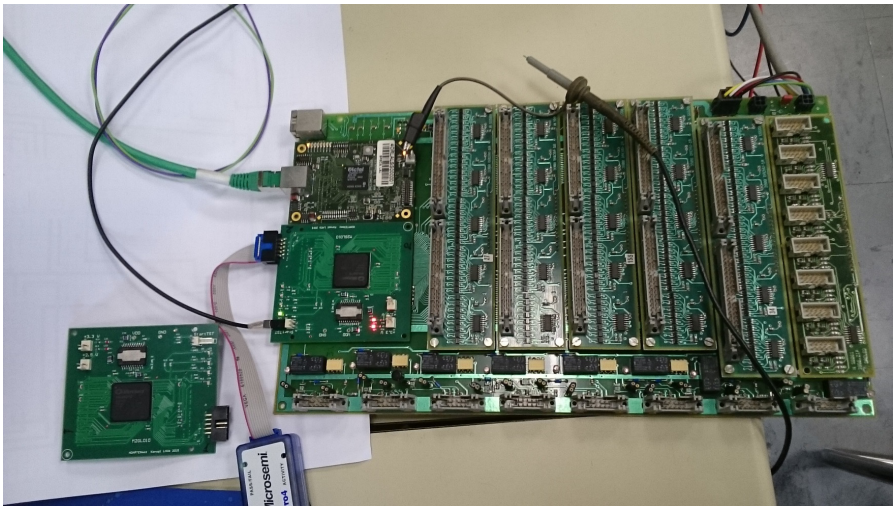

Figure 6: The new calorimeter HV control mezzanine board equipped with its radiation tolerant FPGA, plugged onto the HV motherboard.

The high-voltage, calibration, monitoring and control boards have to be re-designed or modified (for example with the replacement of mezzanines). Prototypes have been designed and produced (Fig. 6), and are now being tested.

A new clusterisation scheme has been developed in the reconstruction to reduce the effect of cluster pile-up at Upgrade luminosity. More recently, new energy- and position-dependent corrections have been determined to improve the calorimeter resolution. The corresponding software has been developed and integrated into the standard LHCb event reconstruction so that they can already be tested with the data taken during Run 2.

The innermost ECAL modules will eventually degrade significantly due to radiation. It is planned to replace them during the LS3 shutdown by spare modules that are already available.

4.3 Muon system

Much progress has taken place in the last six months on the Muon system upgrade project. The LNF construction site is producing replacement and spare MWPCs at the foreseen rate, which is one chamber every two weeks. A system using a radioactive source to scan automatically the new MWPCs has now been put into operation, testing one MWPC per week. The PNPI MWPC construction site is now also operational, and four chambers for the M2R4 region have already been built. In addition a second production line for chambers for the M4R2 region is almost ready to go in production. The MWPC production in these two sites is expected to continue during 2016. In 2017 chambers will be transported to CERN for the installation of the front-end electronics and final tests with cosmic rays.

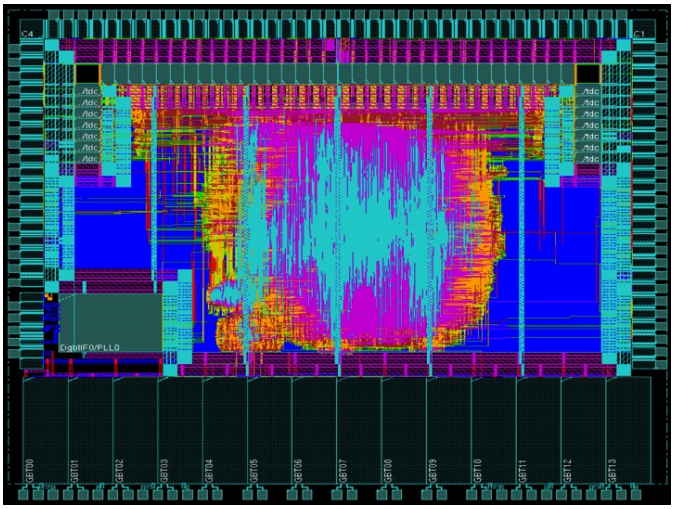

Figure 7: Layout of the muon system nSYNC chip.

An EDR for all the new readout (nSYNC, nODE) and control (nSB, nPDM) electronics took place at the end of November 2015. This in-depth review allowed the functionalities required in the new ASIC and the boards to be finalised. The nSYNC ASIC, designed in UMC 130 nm and using a TDC block extensively tested in autumn 2015, was submitted in February (Fig. 7) and is expected to be ready before the summer break. The work on the construction of the nODE board prototypes, for which the design was finalised after the EDR, will continue in the coming months, with the target of having a first board ready for assembly when the ASIC will be available.

The work on the muon system specific TELL40 firmware is continuing. Current simulation indicates that a TELL40 board based on the ARRIA10 FPGA should

handle 24 optical links at input. A study on possible zero-suppression algorithms is under way to allow the output bandwidth to the CPU to be reduced. It has been demonstrated that the FPGA that will be used (IGLOO2) the new MWPC control (nSB) and pulsing (nPDM) boards perform as required. In addition, laboratory tests showed that a prototype full control chain, composed of a MiniDAQ, a Versatile-link Demo Board, an IGLOO2 driver and muon system front-end boards, can be operated in a reliable manner.

Work is progressing to improve the shielding in the beampipe region and in front of the muon system. The design of a new M2 plug and of an improved tungsten/lead shielding inside the HCAL is under way, and an EDR is foreseen for May. The PRR for the M2 plug will take place at the same time, allowing the order to be placed during the summer and the shielding to be installed during the extended end-of-year shutdown after the 2016 run.

Studies on how to improve the performance of the muon identification algorithms for the upgrade are continuing. New approaches are being tested with full Upgrade simulation and real data collected during Run 1 at high luminosity, and show promising results. These new approaches will already be deployed during Run 2.

5 Data processing

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

5.1 Data acquisition

The current priority in the domain of readout electronics is to define and optimise the required numbers of the PCIe40 (PCI Express) readout boards. The optimisation criteria are the simplicity of the overall system, the FPGA occupancy and running frequency, as well as the total cost. Studies are ongoing, with a final decision expected in June.

The prototypes of the PCIe40 readout board were delivered in mid-2015. An extensive campaign of tests were run, which demonstrated the feasibility of a PCI Express board handling a total bandwidth of 100 Gbit/s.

Twenty-five copies of the PCIe40 prototype will be produced with the final version of the FPGA Arria10. These boards will be used to equip the MiniDAQ2, which is a PC-server for the testing and validation of the front-end electronics of each sub-detector. The MiniDAQ2 systems will be delivered to the collaboration by the summer of 2016. This constitutes a delay of four months with respect to the previous schedule, and is caused by the late availability of the Arria 10. However, a slippage of this duration is still compatible with the timelines of the subdetector projects.

In the current architecture the front-end electronics and the readout electronics are connected via long distance optical links between the underground area and the surface. A small set of OM3 fibres have been deployed in the PM 85 shaft. Long-term tests have shown that a bit error rate of 10^{-18} is achievable with 95% confidence level. These tests use a loop-back between the emitter and the receiver, doubling the optical fibre length. They will be repeated with the nominal fibre length using a front-end board provided by the SciFi subdetector, as well as the GBT transceiver.

5.2 Online

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

Software packages have been developed to evaluate the performance of the event building and the transport layers in different configurations and different technologies. Large scale tests have been run at HPC centres. First results are very encouraging with an aggregate traffic of 3.2 Tbit/s having been measured between 128 servers. Recently, a bandwidth of 86 Gbit/s has been achieved on a 12-node cluster using the transport layer Infiniband EDR.

In parallel, proposals are being evaluated to distribute the timing and fast commands using a PON network, as well to simplify the regulation of the data flow using a local throttle mechanism.

5.3 Trigger

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

The Trigger TDR [6] already laid out an achievable solution for the Upgrade trigger, one that is highly performant and operates within the available timing budget. Many novel aspects proposed for the Upgrade trigger have already been achieved in the Run-2 High-Level Trigger.

In 2016, the strategy for the charm physics will be reviewed and the performance of the Upgrade trigger will be re-valuated, taking into account all improvements to the Upgrade detector simulation and reconstruction. Originally foreseen at the beginning of the year, these studies will be postponed to the second half of 2016 since they require the final optimisation of the SciFi detector layout.

In parallel, the current High-Level Trigger will be improved since the TURBO stream will become more flexible in terms of event content. This improvement can be considered as a prototype for the Upgrade trigger.

5.4 Computing

Twice a year a workshop is organized to discuss all aspects of LHCb computing, including the requirements of the Upgrade. The latest edition took place in November 2015 in Paris and was devoted entirely to the Upgrade. Several working groups were established to explore requirements and new ideas on the software framework, the event model, the computing hardware, the simulation and the analysis model for the Upgrade. These working groups are documenting their conclusions in the so-called computing roadmap document, a first draft of which was presented to the collaboration in March.

The roadmap details the work to be done and the decisions to be taken for the computing TDR, including estimates of the required and available effort. The most important goals of the R&D phase are three demonstrators that are to be completed by the beginning of 2017. The first of these is to study a new software framework using a task-based approach. The second is devoted to event models in which objects are immutable and composable. The third is focused on the vectorisation of individual algorithms, taking into account the attributes and capabilities of the hardware platform. Work on others areas such as the analysis model, simulation, and collaborative working, will follow an evolutionary approach in which ideas will already be tested during Run 2 operation. The required level of effort for all activities is around 19 FTEs, of which only half are currently identified. A campaign of recruitment has begun to meet this shortfall.

The key upcoming milestones in Upgrade computing are the final version of the roadmap document (Q1 2016), the Software and Computing TDR (Q4 2017), and the finalisation of computing model (Q3 2018).

6 Organisation & oversight of Upgrade activities

6.1 Organisation

The Upgrade Planning Group (UPG) meets regularly to review progress. The UPG membership consists of an Upgrade Detector Coordinator, an Upgrade Resources Coordinator, an Upgrade Performance Coordinator and an Upgrade Data Processing Coordinator, as well as the management and a representative of the Physics Coordinator.

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

6.2 Milestones

A series of important milestones has been achieved in the second half of 2015 including many EDRs. A revised schedule with major milestones for the con-

struction phase has also been established: it takes into account the new long-term LHC schedule, a preliminary LHCb upgrade installation programme for LS2, and a better and more refined definition of the construction schedule of the subdetectors.

All the upgrade projects are now moving towards the production phase, with a set of remaining EDRs planned for 2016 and a series of PRRs, some of which already planned for Q2 2016. Some milestones (VeloPix ASIC submission, RICH Elementary Cell and CLARO ASIC PRRs) have been delayed by a few months to account for final design optimization before launching mass production. These delays do not cause concern at the moment and plans to mitigate their effects within the current schedule are in place.

All projects are working to a timescale that now assumes the start of the LS2 at the end of 2018 and, with this, the beginning of the Upgrade installation. The schedule has also been adjusted to reflect the extended duration of LS2 to two years. Consideration is being given to the possible installation of optical fibres for the data transmission during Run 2 in the extended year-end technical stops.

6.3 Infrastructure

The construction of the new assembly hall at the experimental site is progressing (Fig. 8), although a delay of six months has been accumulated. The hall will be completed by early autumn this year. A large semi-clean room will be installed in the existing building at the surface close to the experimental shaft. The floor for this has recently been prepared.

Figure 8: The new assembly hall under construction at IP8 in February 2016.

Construction of the technical gallery is on hold until a decision is made on the final location and housing of the computer farm. The baseline solution is a container-based approach, but other cost-effective solutions are under consideration. A workshop on this subject is scheduled for late April 2016 and the agenda is in preparation.

Studies of the dismantling of PS/SPD detector and the accompanying lead absorber are progressing well. A procedure for the dismantling of the first muon station (M1) has been almost finalised. The wall that currently holds the M1 chambers will be reused to install the neutron shielding that is required for the SiPMs of the SciFi.

The integration of services such as the detector-cooling system has begun. Possible routing of the transfer lines from the cooling plant to the patch panels close to the detector are under study.

7 Funding

The funding requirements of the LHCb upgrade project have been defined in detail in Addendum No. 1 to the Memorandum of Understanding (MoU) for Common Projects [7] and in the Addendum No. 2 to the MoU for 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 Subdetector System component of 41.5 MCHF. All TDRs have been fully approved by the Research Board and both Addenda have been submitted for signature to the Funding Agencies.

The LHCb Upgrade project is now fully underway. Price enquieries have been launched, inviations to tender are ongoing and some major contracts have already been signed. CORE funds are now being committed for a number of sub-detector components and a substantial fraction of the anticipated funds is expected to be spent in 2016 and 2017. There is confidence that full funding of the Upgrade project will be available in due time to ensure a complete and timely installation of the LHCb upgraded detector during Long Shutdown 2.

References

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

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

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

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

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

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

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

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