Phase I Upgrade of the Readout System of the Vertex Detector at the LHCb Experiment

Antonio Fernández Prieto, Pablo Vázquez Regueiro, Karol Hennessy, Jan Buytaert, Martin van Beuzekom, Edgar Lemos Cid, Lars Eklund, Kristof de Bruyn, Sneha Naik, Manuel Schiller, Dónal Murray, Alexander Leflat, Oscar Boente García, Abraham Gallas Torreira, Beatriz García Plana, Themis Bowcock, Francesco Dettori, Karlis Dreimanis, Vinicius Franco Lima, David Hutchcroft, Kurt Rinnert, Tara Shears, Oscar Augusto, Victor Coco, Paula Collins, Tim Evans, Massi Ferro-Luzzi, Wolfgang Funk, Heinrich Schindler, Kazu Akiba, Elena Dall'Occo, Cristina Sanchez Graz, Wouter Hulsbergen, Daniel Hynds, Igor Kostiuk, Marcel Merk, Aleksandra Snoch, Dana Seman Bobulska, Silvia Borghi, Stefano de Capua, Deepanwita Dutta, Marco Gersabeck, Chris Parkes, Peter Svihra, Mark Williams, Galina Bogdanova, Vladimir Volkov, Pawel Kopciewicz, Maciej Majewski, Agnieszka Oblakowska-Mucha, Bartlomej Rachwal, Tomasz Szumlak, Lucas Meyer Garcia, Franciole Marinho, Larissa Helena Mendes, Irina Nasteva, Juan Otalora, Gabriel Rodrigues, Jaap Velthuis, Pawel Jalocha, Malcolm John, Nathan Jurik, Luke Scantlebury-Smead, John Back, Tim Gershon, Tom Latham, and Andrew Morris

Manuscript received October 25, 2019; revised December 18, 2019 and January 16, 2020; accepted January 24, 2020. Date of publication January 30, 2020; date of current version April 16, 2020. This work was supported in part by the Spanish Government, Ministry of Science, Innovation and Universities, in part by the Ministry of Economy and Competitiveness under the Unit of Excellence María de Maeztu, and in part by the Regional Government of Galicia, IGFAE Research Center and University of Santiago de Compostela.

Antonio Fernández Prieto, Pablo Vázquez Regueiro, Edgar Lemos Cid, Oscar Boente García, Abraham Gallas Torreira, and Beatriz García Plana are with the Instituto Galego de Física de Altas Enerxías (IGFAE), and the High Energy Physics Department, Universidade de Santiago de Compostela, 15705 Santiago de Compostela, Spain (e-mail: antonio.fernandez.prieto@cern.ch).

Karol Hennessy, Themis Bowcock, Francesco Dettori, Karlis Dreimanis, Vinicius Franco Lima, David Hutchcroft, Kurt Rinnert, and Tara Shears are with the Department of Physics, Liverpool University, Liverpool L69 3BX, U.K.

Jan Buytaert, Kristof de Bruyn, Oscar Augusto, Victor Coco, Paula Collins, Tim Evans, Massi Ferro-Luzzi, Wolfgang Funk, and Heinrich Schindler are with the European Organization for Nuclear Research (CERN), 1211 Geneva, Switzerland.

Martin van Beuzekom, Kazu Akiba, Elena Dall'Occo, Cristina Sanchez Graz, Wouter Hulsbergen, Daniel Hynds, Igor Kostiuk, Marcel Merk, and Aleksandra Snoch are with the National Institute for Subatomic Physics (Nikhef), 1098 XG Amsterdam, The Netherlands.

Lars Eklund, Sneha Naik, Manuel Schiller, and Dana Seman Bobulskar are with the Department of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, U.K.

Dónal Murray, Silvia Borghi, Stefano de Capua, Deepanwita Dutta, Marco Gersabeck, Chris Parkes, Peter Svihra, and Mark Williams are with the School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, U.K.

Alexander Leflat, Galina Bogdanova, and Vladimir Volkov are with the Skobeltsyn Institute of Nuclear Physics, Moscow State University (MSU), Moscow 119991, Russia.

Pawel Kopciewicz, Maciej Majewski, Agnieszka Oblakowska-Mucha, Bartlomej Rachwal, and Tomasz Szumlak are with the Department of Particle Interactions and Detection Techniques, AGH University of Science and Technology, 30-059 Kraków, Poland.

Lucas Meyer Garcia, Franciole Marinho, Larissa Helena Mendes, Irina Nasteva, Juan Otalora, and Gabriel Rodrigues are with the Department of Physics, Universidade Federal do Rio de Janeiro (UFRJ), Rio de Janeiro 21941-901, Brazil.

Jaap Velthuis is with the School of Physics, University of Bristol, Bristol BS8 1TH, U.K.

Pawel Jalocha, Malcolm John, Nathan Jurik, and Luke Scantlebury-Smead are with the Department of Physics, University of Oxford, Oxford OX1 2JD, U.K.

John Back, Tim Gershon, Tom Latham, and Andrew Morris are with the Department of Physics, University of Warwick, Coventry CV4 7AL, U.K.

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNS.2020.2970534

***Abstract*— This article describes the high-speed system designed to meet the challenging requirements for the readout of the new pixel VErtex LOcator (VELO) of the upgraded LHCb experiment. All elements of the electronics readout chain will be renewed to cope with the requirement of ∼40-MHz full-event readout rate. The pixel sensors will be equipped with VeloPix ASICs and placed at ∼5 mm from the Large Hadron Collider (LHC) beams in a secondary vacuum tank in an extremely high and nonhomogeneous radiation environment. The front-end (FE) ASICs with the highest occupancy will have to cope with pixel-hit rates above ∼900 Mhits/s using up to four 5.13-Gb/s data readout links. Each module comprises six VeloPix ASICs, wire-bonded to two FE hybrid boards, while a third hybrid will employ a GBTx ASIC as the control interface. High-speed data will reach the wall of the vacuum chamber through low-mass flexible copper tapes. A custom board routes the signals outside the vacuum tank. On the air side, an optical and power board converts the electrical high-speed signals into optical signals for transmission from the underground cavern to the off-detector electronics that process data and send them to a farm of computers for further analysis. Several tests allowing the validation of the system are described here with special emphasis on a test with proton beams that confirms the correct operation of the whole readout hardware.**

***Index Terms*— Field-programmable gate array (FPGA), high energy physics (HEP) instrumentation, high-speed readout, real-time data acquisition system, silicon detector, vertex detector.**

I. Introduction

LHCb [1] is a general-purpose physics experiment in the forward region at the Large Hadron Collider (LHC). It is designed to perform high-precision measurements of CP violation and search for new physics using the enormous flux of beauty and charmed hadrons produced at the LHC. The LHCb experiment, in which detector is placed in a cavern 100-m underground, has an excellent tracking resolution $\delta p/p \sim 0.4\%$, particle identification, and a flexible trigger. The trigger system consists of two stages. The first level of a trigger, which is hardware-based, uses information provided by the calorimeter and muon systems to preselect hadron-collision events at a rate of 1 MHz. All of the data that pass through the first trigger are sent to a CPU farm where the

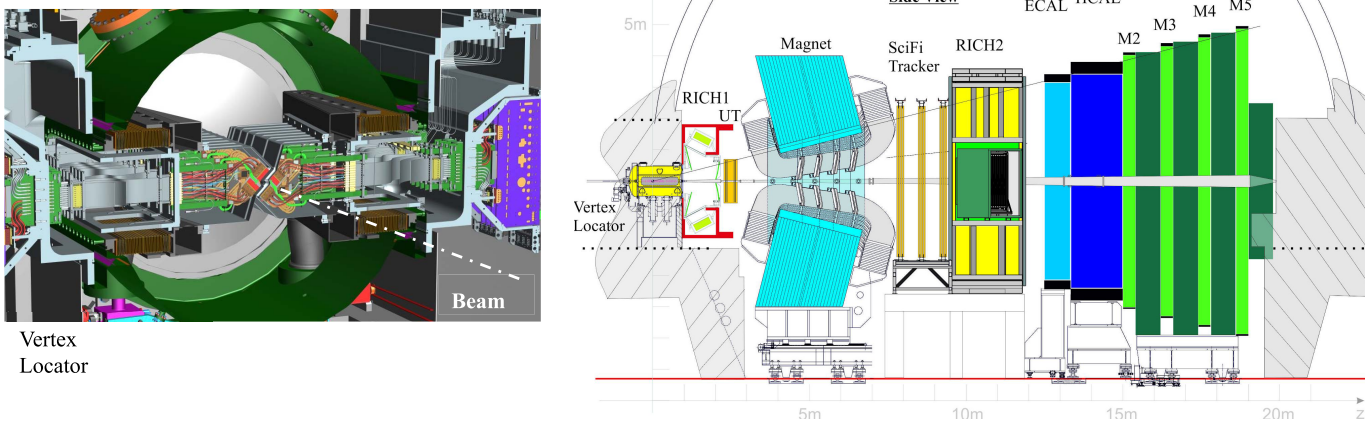

Fig. 1. LHCb upgrade detector and VELO subdetector.

events are reconstructed and selected. Thus, the LHCb output rate is 5 kHz. During Runs 1 and 2 of the LHC (2010–2018), 9.2 fb^{-1} of proton–proton collisions data were accumulated at the center-of-mass energies $\sqrt{s} = 7 - 8$ TeV for Run 1 and $\sqrt{s} = 13 - 14$ TeV for Run 2.

A. LHCb Upgrade

Throughout 2019–2020, the LHCb detector (see Fig. 1) will undergo a major upgrade. The objective is to increase the yields of signal samples during Runs 3 and 4. To achieve this, it is necessary to improve the trigger efficiency on hadronic channels and rare decays. The luminosity will be increased by a factor of 5, up to 2×10^{33} $\text{cm}^{-2}\text{s}^{-1}$, expecting an integrated luminosity of 50 fb^{-1} by the end of Run 4 (2030). Improving trigger rates would quickly saturate the current readout system. Hence, LHCb will remove the actual hardware trigger and replace it with a fully software one [2]. This implies that all of the subdetector front-end (FE) electronics will be replaced in order to handle the LHC peak (40 MHz) and mean (30 MHz) collision rates. The data are driven from the subdetectors FE via a multi-Tb/s readout network into a computing center where events are fully reconstructed prior to the trigger decision. The aim of the full software trigger is to select for storage only the events with beauty and charm particles decaying into a large variety of final states with the highest efficiency and purity, minimizing systematic uncertainties. For the calorimeter and muon detectors, only the FE electronics will be replaced. Meanwhile, the particle identification systems (RICH1 and 2) will be partially replaced, and the tracking system will be completely renewed, as shown in Fig. 1. This includes a new vertex detector [VErtex LOcator (VELO)] and tracking stations before (Upstream Tracker, UT) and after (Scintillating FIber, SciFI) the dipole magnet.

The new VELO will operate closer to the beam replacing the silicon strip technology with hybrid pixels. This will improve the hit resolution and, therefore, the track reconstruction. The tracking station before the magnet (UT) will use silicon strips, as was done in the past, reducing the thickness and granularity, while after the magnet, a new technology, SciFI, was chosen. The hybrid photon detectors (HPDs) of the RICH detectors will be replaced.

The gain of the photomultiplier tubes (PMTs) in the calorimeters will be reduced to keep the same average anode current and reduce the impact of aging. New FE electronics have been designed to cope with this decrease in gain.

This article is an extension of the work presented at the IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC) 2018 [3] with a more detailed overview of the off-detector electronics design and the validation of the readout chain.

II. LHCb VELO Upgrade

The upgraded VELO [4] will be a hybrid pixel detector consisting of 208 silicon sensors, each of those bonded to three FE ASICs using flip-chip technology. The closest sensor will be placed at ~5 mm from the interaction point within a secondary vacuum isolated from the LHC primary vacuum. The sensors are distributed in 52 modules along the beam in two opposite retractable (by 6 cm) halves, as shown in the left part of Fig. 1. Based on the Monte Carlo simulations, the highest occupancy ASICs will have pixel hit rates of about 900 Mhits/s, which is equivalent to 520 Mhits/s in terms of cluster rates (see Section II-B). This translates into a bit rate of ~16 Gb/s adding up to ~3- Tb/s data for the ~41 Mpixels of the whole upgraded VELO. One of the major challenges of the sensors and the readout chips is caused by their orthogonal arrangement to the beam pipe. The irradiation profile along the sensors is inversely proportional to the squared distance to the interaction point with the greatest concentration of hits at the innermost sensor regions.

A. Sensors

The upgraded VELO will be populated with 200-μm-thick Hamamatsu n-on-p sensors optimized for electron collection in order to mitigate the effect of timewalk. This effect is caused by the different spectrum of the sensor collected charge, which

TABLE I
VELOPIX ASIC SPECIFICATIONS

Technology	130 nm CMOS
Pixel matrix	256×256 pixels with pixel size of 55 µm×55 µm
Radiation hardness	4 MGy achieved using triplicated voting registers
Power consumption	< 1.5 W/ASIC
Readout	Unsorted binary data driven at 40 MHz
Data throughput	20.52 Gb/s

generates a nonnegligible spread in the time response at the input of the electronics preamplifier. Hence, hits with small collected charges generate signals with a slower rising edge, generating a late detection of the pixel hits.

The sensor efficiency is higher than 99%, reducing the risk of missing the first measurement points. The pitch of the sensor is compatible with the readout chip and with elongated pixels of 55 μm $\times$ 137 μm in the region between the ASICs. The sensor is built with a 450-μm guard ring and deep reactive-ion etching (DRIE) needed for the foil clearance.

By the end of LHC Run 4 (2030), the hottest sensor will be exposed to a fluence of 8×10^{15} 1 MeV and will withstand bias voltages around 1000 V. Sensors will be operated at -20 °C to protect the silicon from thermal runaway effects after irradiation.

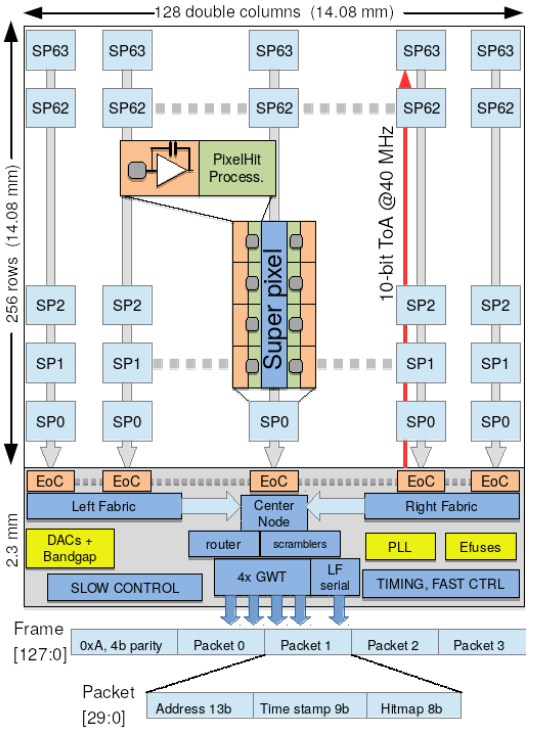

Fig. 2. General architecture of VeloPix ASIC.

B. VeloPix ASIC

To read out the charge collected in the sensors, a new FE ASIC called VeloPix [5], based on the MediPix/TimePix family, has been designed to meet the very demanding readout rate, radiation hardness, and low-power consumption requirements of the experiment (see Table I). A single sensor is bump-bonded to three VeloPix ASICs comprising a "tile." The VeloPix has to timestamp the collected data according to the 3564 bunch crossings of one LHC orbit, for which 12 bits are needed. In order to optimize the data bandwidth, only 9 bits of tagging are transmitted. The full 12 bits are reconstructed in the back-end readout board (see Section III).

The architecture of the VeloPix ASIC is shown in Fig. 2. The VeloPix transmits only binary data generated from Time-over-Threshold (ToT) measurements in each pixel. Pixel hits are clustered in "superpixel packets" of 2 $\times$ 4 pixels and vertically transmitted with a clock rate of 160 MHz over the matrix until they reach the End-of-Column (EoC) block in the periphery of the ASIC. Then, the data proceed to a central node, as shown in Fig. 2, where they are assembled in 128-bit frames and fed to the high-speed transmitters. Superpixel packets are scrambled prior to transmission using the polynomial $X^{30} + X + 1$ to maximize the transitions, allowing for good data transmission.

With the aim of reducing the power consumption and keeping in mind the radiation hardness of the ASIC, a readout link serializer block, named gigabit wireline transmitter (GWT) [6], was designed for a rate of 5.13 Gb/s. This serializer works by feeding a 16-bit input round-robin multiplexer with a byte of data on both edges of a 320-MHz clock (eight times faster than LHC clock). The reading of this multiplexer is triggered by the signal generated on a multiphase delay-locked loop (MDLL), which generates 16 phases evenly spaced from the given 320-MHz clock. Four GWT links provide the VeloPix with a maximum readout throughput of 20.52 Gb/s.

C. On-Detector Electronics

An upgraded VELO slice is shown in Fig. 3. From right to left, what can be seen is a VELO module, the high-speed tapes and low- and high-voltage cables, the vacuum feedthrough board (VFB), and the opto and power board (OPB). The high-speed tapes [7] have a low mass to reduce the radiation length in the detector acceptance, and they are flexible enough to allow the detector opening during LHC beams filling without compromising signal integrity. High-speed transmission lines travel almost 1 m before reaching the optical converters. A feedthrough board is needed to make the electrical connection between the vacuum and air-side electronics. All the electrical readout and control signals are converted into optical in the OPB and driven 300 m to the field-programmable gate array (FPGA)-based back-end board on the surface.

1) VELO Module: A single module (see Fig. 4) is equipped with four sensor "tiles" on a silicon substrate, two per side, with microchannel evaporative CO_2 cooling [8]. Each side of the substrate has a GBTx ASIC [9], which distributes the control and timing signals to the six VeloPix.

The design of the module was made with the aim of mitigating stresses due to the coefficient of thermal expansion (CTE) mismatch and minimizing the material in the acceptance of the subdetector. Consequently, two kinds of PCB interconnected

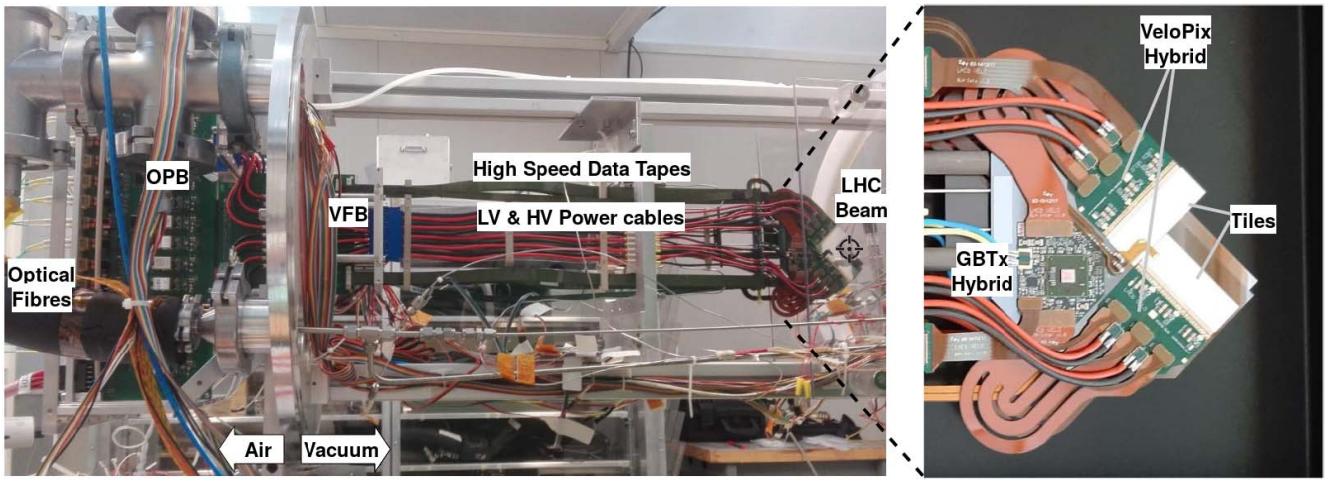

Fig. 3. Left: upgraded VELO slice. Right: close-up of a VELO module.

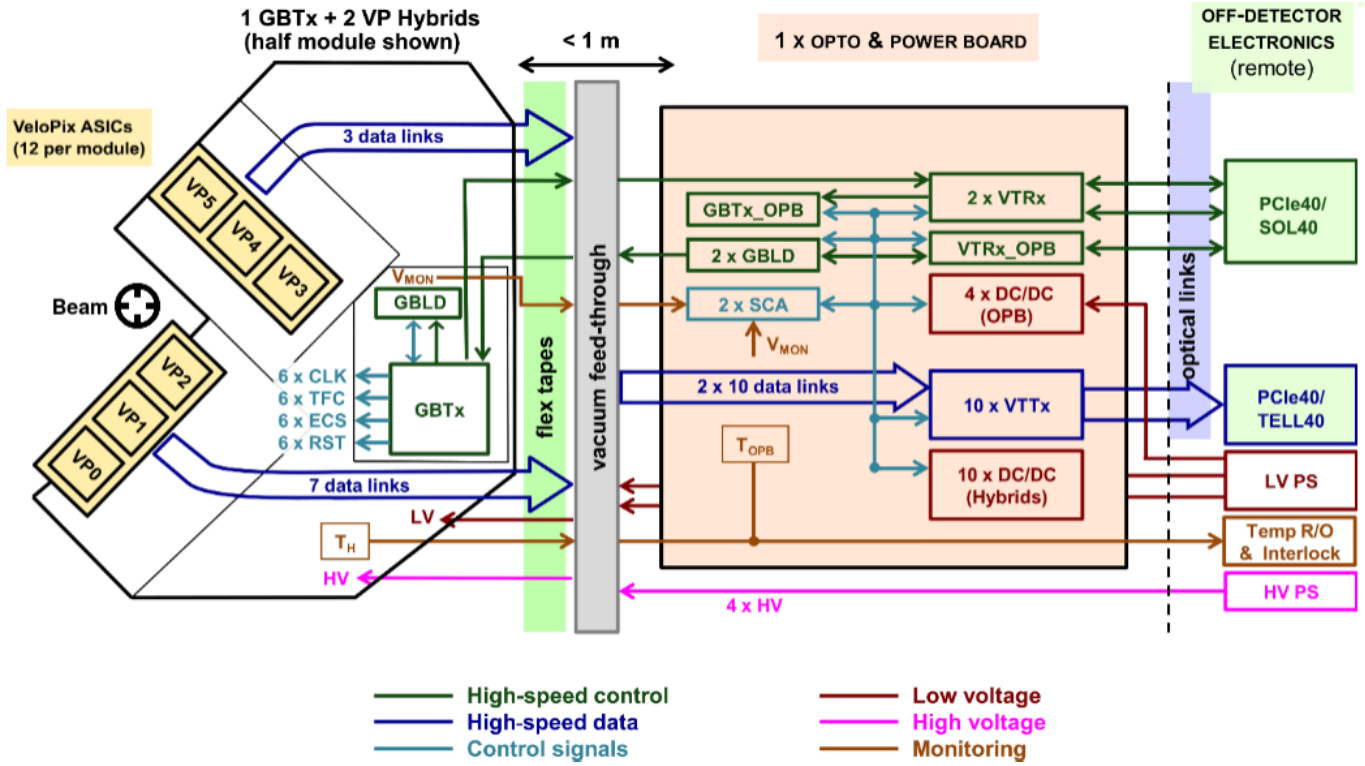

Fig. 4. LHCb VELO electronics overview.

by Kapton tapes were designed: one for hosting the GBTx and one for the VeloPix wire-bonding (see Fig. 3).

The Monte Carlo simulations show that the VeloPix hit distribution for the hottest half module can be grouped in three regions: one with the closest VeloPix to the beam (VP2 in Fig. 4), with occupancies about 16 Gb/s, will require the use of the four GWT readout links. A second region with about half of the maximum occupancy (VP1 in Fig. 4) will need two readout links, and a third region in the periphery (VP0, 3, 4, 5 in Fig. 4), with one-tenth of the maximum occupancy, will only use one readout link per ASIC. As the hit distribution does not change dramatically over the 52 VELO slices, every half module will be read out with ten data links.

2) High-Speed Data Tapes: All module readout and control signals are routed inside the vacuum using low-mass high-speed flex tapes with Kapton microstrip technology in the acceptance region of the detector where low-mass materials are critical to detector performance (∼5 cm). Four flexible PCB tapes built-in stripline technology drive up to seven links. The dielectric used for these flexible tapes has to be radiation tolerant and provide low dissipation loss for high-frequency signals. In addition, high reliability and yield would be required for the impedance control. The best candidate material found, i.e., Pyralux AP Plus [10], was retained. It is also widely used in medical, aerospace, and military applications.

3) VFB and OPB: The VFB is placed at the end of the high-speed tapes. This board has to deal with a pressure of $\sim 10^{-7}$ mbar on its inner side and the atmospheric pressure of the ambient air on the outer side. To help handling the difference in pressure, the PCB was designed with edge metallization.

On the air side of the vacuum, feedthrough is connected to the OPB through a peripheral component interconnect (PCI) express connector. The main functions of the OPB are to provide low voltage to the FE electronics, monitor the FE temperatures, and convert the signals from electrical into optical, and vice versa.

All the active components used on the OPB board are radiation tolerant: the FEASTMP dc/dc converters [11], the VTTx and VTRx [12] electro-optical converters, the GBTx gigabit transceiver, and the GBT-SCA ASIC [13], which is used for control and monitoring.

III. Off-Detector Electronics

Off-detector electronics encompass all the electronic systems placed at the ground level, which are responsible for the control, timing, and readout of the underground detector.

One of the main limitations of the current LHCb, as mentioned in Section I-A, is the hardware trigger, and consequently, it will be removed in the upgrade. In this trigger architecture, all FE electronics record and transmit data continuously at 40 MHz to the readout boards, while the muon and calorimeter subdetectors transmit parallel event information to the software trigger. This stage issues the decision to accept or reject data on a per LHC bunch crossing basis. These decisions are fed back to the FE and back-end boards via the so-called timing system.

The timing system is responsible for handling the readout dataflow from the FE to the event data processor farm. This system is based on two different types of boards with common hardware (see Section III-A): a readout supervisor board and an interface board. The first acts as an interface with the LHC, distributing the timing and synchronization commands to the LHCb detector. The second handles the distribution to all FEs of two kinds of signals: the timing signals from the previously mentioned readout supervisor board and the control commands from the multiple computers used as supervisory control and data acquisition (SCADA).

A. LHCb Back-End Board Hardware

To implement the LHCb upgrade readout scheme, a common board, called PCIe40, was designed to act as a readout supervisor, a control board, and a readout board. A picture of the PCIe40 board based on an Intel Arria 10 FPGA [14] is shown in Fig. 5. The board that has 48 full-duplex optical links is configured and read out from a host computer via a PCIe interface (version 3, 16 lanes) that offers 128 and 112 Gb/s of peak and sustained throughput, respectively. Moreover, two more bidirectional optical links allow the connection of the board to the passive optical network [15], which is responsible for distributing the timing signals to the entire LHCb detector. The role performed by a PCIe40 board (readout supervisor, control interface, or readout node) is entirely determined by the firmware loaded in the on-board FPGA; the hardware remains unchanged.

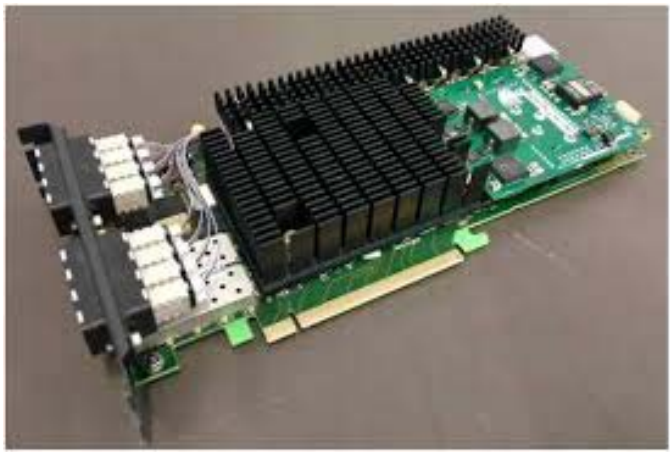

Fig. 5. PCIe40 board.

B. LHCb Back-End Board Firmware

The whole LHCb detector comprises only one readout supervisor board. This board receives the 40-MHz clock of the LHC and distributes it along with the detector but also generates the signals that keep all subdetectors synchronous. Besides that, the supervisor card contains information about the LHC filling scheme that is used to accept events only with collisions in the interaction point.

Readout and control interface board firmwares are developed in an LHCb wide common framework with provision made for subdetector specialization. In order to reduce the firmware complexity in the readout board, two independent datastreams are instantiated, as shown in Fig. 6 for VELO case.

C. VELO Control Firmware

The VELO is the only subdetector within LHCb that needs a specific control interface firmware to act as a bridge between the PCIe base address registers and its FE due to the VeloPix control and timing dedicated protocols.

The upgraded VELO requires four PCIe40 cards configured as control interface boards, each with 39 links connected to their corresponding GBTx. A VELO slice has three GBTx: one for commanding the OPB and monitoring temperatures and analog signals through the GBT-SCA, and two command the 12 VeloPix of a module. Scalable low-voltage signaling (SLVS) links are used between the GBTx and the VeloPix: a full-duplex serial peripheral interface 80-Mb/s link for the control and a simplex 320-Mb/s link for timing. Both links use proprietary protocols.

D. VELO Readout Firmware

1) GWT Interface: The readout board firmware architecture shown in Fig. 6 is a significant modification to the LHCb framework as the VeloPix ASIC is the only FE reading out data over a dedicated protocol (GWT) in reference to the

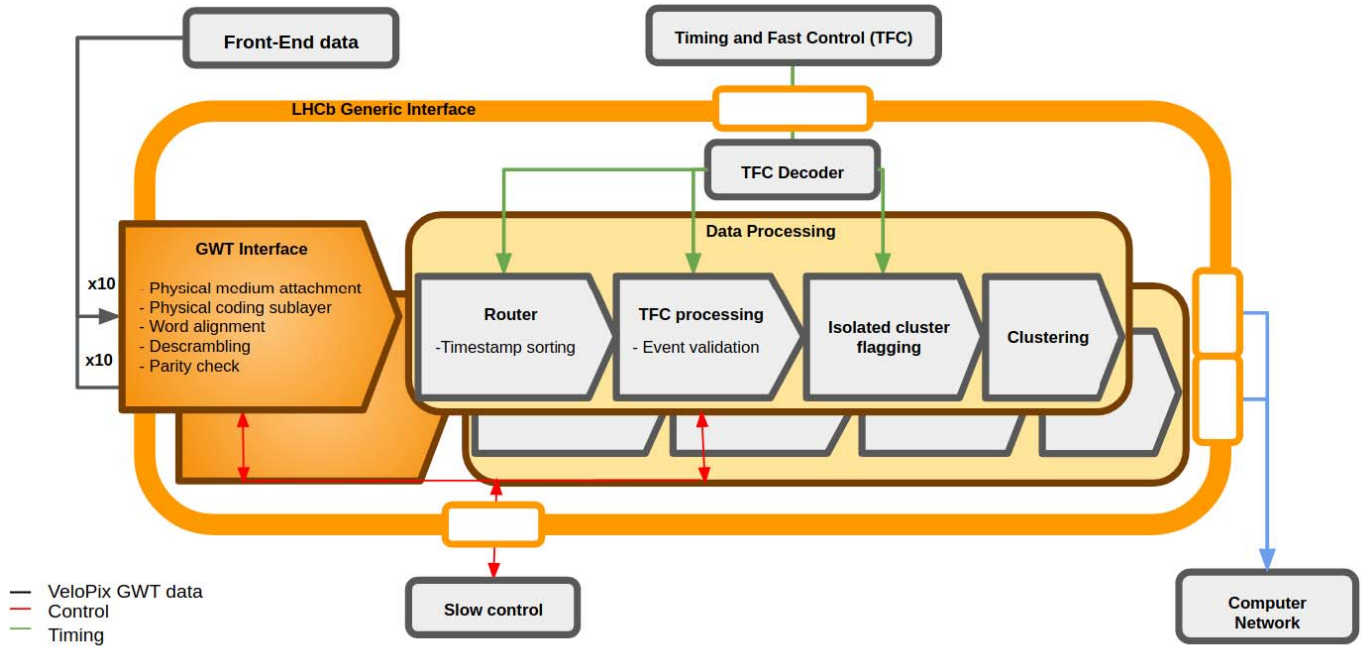

Fig. 6. VELO readout firmware architecture divided in two parallel datastreams.

VeloPix serializer. Hence, an *ad hoc* firmware was designed to drive the serial transceivers embedded in the FPGA. The physical medium attachment of the FPGA transceiver requires an optimization of the dc and ac gains and the continuous-time linear equalization (CTLE). Moreover, in the GWT interface, a custom-made physical coding sublayer was made for the GWT protocol. This requires the design of an alignment block for reconstructing the 128-bit GWT word along with the descrambling and a parity check mechanism to monitor and discard erroneous packets. No forward error correction is available within this protocol.

2) Data Processing: The first task of data processing is to collect all the packets coming from the same event. For the VELO case, each data stream receives on average 22 time-unsorted packets per bunch crossing and half module. The timestamp sorting algorithm, based on a two-stage process, classifies packets by its 9-bit timestamp and stores them in 512 internal RAM segments. The first stage, based on a combination of ten decision trees of 1-bit sorting elements, sorts the four most significant bits. The second stage, based on lookup tables, sorts the five remaining bits. At this point, events declared as valid are propagated to further processing. Memory usage is one of the big issues of the sorting algorithm. Simulations show that the maximum number of packets that can be stored per event fitting in the FPGA is 512, keeping packet loss below 4×10^{-6}.

If more memory resources are needed for extra processing, a still acceptable packet loss of 2×10^{-5} can be reached with the next smallest value, i.e., 256.

The VELO readout board latency is below 26 μs. Faster subdetectors must hold their data fragments in their host computer until the slowest one is ready for the event building. Once the events are created, they pass through the event filter farm that reduces the event rate from 30 MHz to a rate that can be handled for storage between 20 and 100 kHz. This reduction is possible due to a full software trigger based on clustering, full track reconstruction, fit, and particle identification. Part of the clustering, as isolated cluster flagging or the entire clustering, can be implemented in firmware if enough resources are available.

Reading out the whole VELO requires 52 PCIe40 boards configured as readout boards (one per VELO slice) with 20 input links.

IV. Validation of the Electronics

A. VeloPix Tests

A specific carrier board and a readout system based on the speedy pixel detector readout (SPIDR) readout system [16] were designed to validate all the functionalities of the VeloPix ASIC. The SPIDR system is designed for the readout of the Medipix/Timepix family of ASICs.

The VeloPix version of the SPIDR system allows the full control of the signals provided to the ASIC and the reading out of the four GWT links at full speed, driving the data without further processing using 10-Gb Ethernet transmission.

Extensive tests on bare ASICS using the SPIDR system consisted of injecting both analog and digital test pulses at the pixel level. These tests allowed inspecting the chip at most of its design corners with special attention given to the high readout rate. The ASIC performed as expected, reaching readout rates up to 1280 Mhits/s. Tests of ASICs equipped with silicon sensors were also performed with the help of the SPIDR system. As an example, the reading of data coming from radioactive sources or withstanding the maximum hit rate of Fermilab beam facility (about 300 Mtracks/s) was done on a five-plane VeloPix telescope.

In order to validate the radiation hardness of the VeloPix ASIC at the level of dose that is expected in the upgraded LHCb detector, several tests were accomplished with ionizing

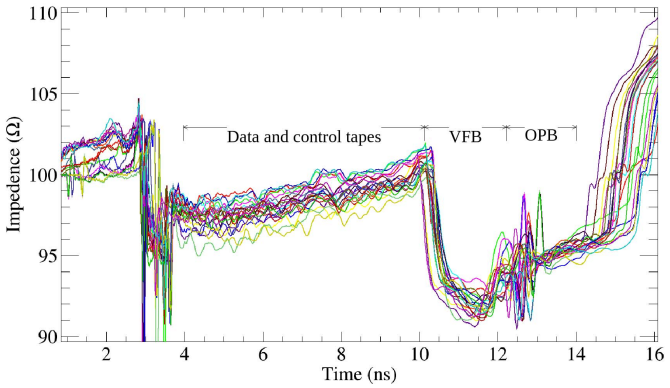

Fig. 7. Characteristic impedance per link for all the FE links.

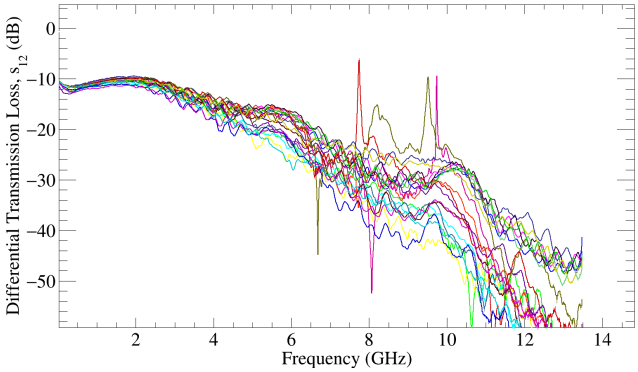

Fig. 8. Data transmission loss for all the FE data links.

particles. Total ionizing dose (TID) tests up to 4 MGy using X-ray tubes took place at the Universities of Santiago de Compostela and Glasgow operating the VeloPix at −20 °C and room temperature, respectively. No change in power consumption and no drift in the analog parameters were found. Three campaigns of heavy ion irradiation, in the Louvain-La-Neuve cyclotron, were needed to validate the chip under single-event effects (SEEs). On the first version of the chip, the first observation of single-event latch-ups (SELs) in a design using 130-nm technology was found. This version also showed single-event upsets (SEUs) in a reset line. The SEE occurrence was later confirmed and debugged in a laser test facility in Montpellier. All these issues were well understood and corrected in a second version of the VeloPix by reducing the large distance between contacts to avoid SEL and applying triple modular redundancy (TMR) in the reset line to avoid SEU.

B. Signal Integrity Measurements

In parallel with the VeloPix tests performed with SPIDR readout, the VELO on-detector electronics were widely tested in many different versions of the hardware. In order to minimize the signal distortion, it is mandatory to guarantee that the differential edge-coupled strip lines have a characteristic impedance close to 100 Ω over the different components of the electrical readout chain (hybrid, tapes, VFB, OPB). Fig. 7 shows the impedance over the line components for the 20 data links of a module, all in the range of 96 ± 6 Ω.

Differential losses in the data transmission lines were measured along the OPB, VFB, and high-speed data tapes using a Vector Network Analyzer, as shown in Fig. 8. These lines are made of copper and are almost 1 m long for all the 24 links (control and data).

Even with a good performance, the transmission lines show a residual signal distortion, addressed by adding hardware CTLE circuits (passive high-pass filters) with a Nyquist frequency of 2.4 GHz for the control links, running at 4.8 Gb/s, and 2.56 GHz for the readout links.

Control link transmission is optimized as well by tuning the preemphasis value of the gigabit laser driver (GBLD) [17] at the beginning of the link, as shown in Fig. 9.

The GWT serializer of the VeloPix is very sensitive to the phase between the 160- and 320-MHz internal clocks.

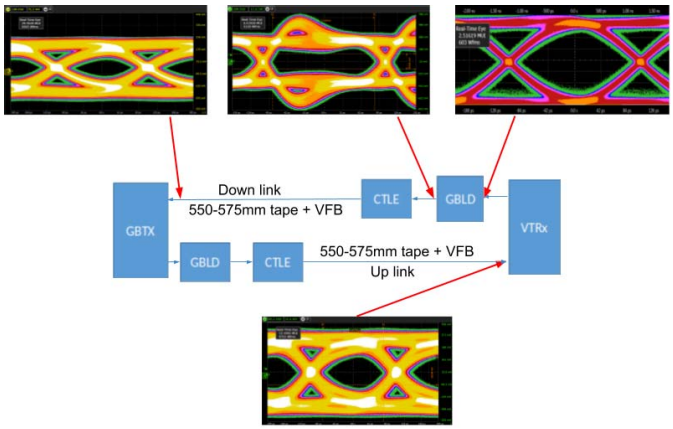

Fig. 9. Optimization of the control link transmission running at 4.8 Gb/s.

This phase can be modified externally; optimal phase values allow a data link transmission bit error rate (BER) after 400 m of optical transmission and below 10^{-12}. This was considered acceptable as it represents less than one error every 200 s.

C. VELO Beam Test

Three slices of the upgraded VELO readout were successfully operated in a beam test at SPS North Area beam facility (CERN) in October 2018. Even though the SPS hit rate is not comparable with the expected rate of the upgraded VELO (-three to four orders of magnitude smaller), the goal of the beam test was to validate the whole readout chain. To do so, 114 GB of data were taken over the GWT links synchronously with a TimePix 3 telescope [18], allowing the measurement of the same tracks in both systems and checking the synchronization between planes with less than 1 ns.

Different tests were performed on the three module slices as follows:

1) sensor bias voltage scan;
2) Velopix analog gain, threshold, and ToT scans;
3) VeloPix GWT clock-phase scan;
4) time-of-flight alignment between slices.

A specific version of the readout board firmware (see Section III) was developed exclusively for this beam test, including system debugging tools as a data bypass (allowing unprocessed data to be saved directly to disk), a 64-bit global timestamp, and a customized telescope synchronicity system.

The synchronization between the VeloPix planes and the TimePix telescope was made by using the VeloPix planes as a master system that triggers the data acquisition of the TimePix telescope with a fixed offset between both systems.

First beam tracks of the VELO upgrade were recorded. In addition, each of the three planes under test was correlated in time and space with the TimePix 3 telescope.

V. Conclusion

The LHCb VELO is being upgraded during the LHC long shutdown 2 (2019–2020) along with the rest of the LHCb tracking system and readout electronics to allow for an increase in the luminosity by a factor of five. Thus, a novel VELO design based on hybrid pixel technology was developed, overcoming a series of challenges in order to read out the huge amount of data produced; FE ASICs are bump-bonded to the sensor and transmit zero suppressed, time-stamped, nonsorted packets with binary hit information over 5.13- Gb/s links with a custom data protocol. Data travel almost 1 m on flexible copper tape and are then converted into optical signals that are sent over 400 m of fibers to the back-end readout board. These data are reordered in time and processed in real time to identify and flag isolated clusters. Processed data are then driven out to a CPU farm at an expected maximum rate of 100 Gb/s per back-end board, corresponding to an aggregate total data rate of 2.85 Tb/s for the ∼41 Mpixels of the whole VELO.

A series of tests have been described focusing on the performance of the VELO control and readout system, assuring the performance of the FE ASIC under the LHC radiation conditions, the quality of the transmission lines, and the integration with the readout board. The validation process culminated in a beam test, certifying the correct operation of the whole readout chain.

Acknowledgment

The authors would like to thank all the people involved in the LHCb VELO Upgrade Project and the Rede de Infraestruturas de Apoio á Investigación e ao Desenvolvemento Tecnolóxico (RIAIDT)-USC analytical facilities of the University of Santiago de Compostela, Spain, where the TID test presented in this article was performed. This work was carried out on behalf of LHCb VELO Upgrade Collaboration.

References

[1] *Letter of Intent for the LHCb Upgrade*, document CERN-LHCC-2011-001, LHCb Collaboration, Mar. 2011.

[2] F. Alessio *et al.*, "The LHCb readout system and real-time event management," *IEEE Trans. Nucl. Sci.*, vol. 57, no. 2, pp. 663–668, Apr. 2010.

[3] A. F. Prieto and P. V. Regueiro, "A terabit readout system for the LHCb VELO upgrade I," in *Proc. IEEE Nucl. Sci. Symp. Med. Imag. Conf. (NSS/MIC)*, Nov. 2018, pp. 1–2.

[4] *LHCb VELO Upgrade Technical Design Report*, document CERN-LHCC-2013-021, LHCb Collaboration, 2013.

[5] T. Poikela *et al.*, "The VeloPix ASIC," *J. Inst.*, vol. 12, no. 01, Jan. 2017, Art. no. C01070.

[6] V. Gromov *et al.*, "Development of a low power 5.12 Gbps data serializer and wireline transmitter circuit for the VeloPix chip," *J. Inst.*, vol. 10, no. 1, Jan. 2015, Art. no. C01054.

[7] E. Lemos *et al.*, "Production, measurement and simulation of a low mass flex cable for multi gigabit/s readout for the LHCb VELO upgrade," in *Proc. TWEPP*, Oxford, U.K., 2012, Art. no. C01018.

[8] O. A. de Aguiar Francisco *et al.*, "Evaporative CO_2 microchannel cooling for the LHCb VELO pixel upgrade," *J. Instrum.*, vol. 10, May 2015, Art. no. C05014, doi: 10.1088/1748-0221/10/05/C05014.

[9] P. Moreira *et al.*, "The GBT project," in *Proc. TWEPP*, Paris, France, 2009, pp. 342–346.

[10] DuPont. (2009). *Pyralux AP-Plus All-Polyimide Thick Copper-Clad Laminates*. [Online]. Available: https://www.cirexx.com/wp-content/uploads/Pyralux_AP-Plus_DataSheet1.pdf

[11] CERN. (2014). *FEAST Radiation tolerant 10 W Synchronous Step-Down Buck DC/DC converter*. [Online]. Available: http://project-dcdc.web.cern.ch/project-DCDC/public/Documents/FEAST%20datasheet.pdf

[12] L. Amaral *et al.*, "The versatile link, a common project for super-LHC," *J. Inst.*, vol. 4, no. 12, Dec. 2009, Art. no. P12003.

[13] A. Caratelli *et al.*, "The GBT-SCA, a radiation tolerant ASIC for detector control and monitoring applications in HEP experiments," *J. Inst.*, vol. 10, no. 3, Mar. 2015, Art. no. C03034.

[14] M. Bellato *et al.*, "A PCIe Gen_3 based readout for the LHCb upgrade," *J. Phys., Conf. Ser.*, vol. 513, no. 1, Jun. 2014, Art. no. 012023.

[15] D. Kolotouros *et al.*, "A TTC upgrade proposal bidirectional 10G-PON FTTH technology," in *Proc. Top. Workshop Electron. Part. Phys.*, Aix en Provence, France, no. 4, Sep. 2014, Art. no. C04001, doi: 10.1088/1748-0221/10/04/C04001.

[16] J. Visser *et al.*, "SPIDR: A read-out system for Medipix3 & Timepix3," *J. Inst.*, vol. 10, no. 12, Dec. 2015, Art. no. C12028.

[17] L. Amaral *et al.*, "A CMOS 0.13 μm, 5 Gb/s laser driver for high energy physics applications," *IEEE Trans. Nucl. Sci.*, vol. 59, no. 6, pp. 3229–3234, Dec. 2012.

[18] D. Hynds, "The Timepix telescope for charged particle tracking," *Nucl. Instrum. Methods Phys. Res. A, Accel. Spectrom. Detect. Assoc. Equip.*, vol. 730, pp. 50–53, Dec. 2013, doi: 10.1016/j.nima.2013.04.060.