PROCEEDINGS OF SCIENCE

The LHCb VELO Upgrade

Daniel HYNDS[*†]
University of Glasgow (GB)
E-mail: daniel.hynds@cern.ch

The LHCb experiment at CERN is dedicated to the study of heavy flavour physics, and more generally to physics in the forward direction. Vital to the analyses which take place is a high precision vertex detector - the VErtex LOcator (VELO). The VELO is also of vital importance in the current online trigger, in the selection of interesting events. With the proposed upgrade of the experiment as a whole in 2018, a replacement for the VELO must be designed and built to cover the even more rigorous demands imposed upon it, particularly in the upgrade of the trigger system. The upgraded detector must be able to sustain full readout at 40MHz and be more radiation hard, necessitating the replacement of both modules and front end electronics. A summary of the current R&D taking place towards this goal is presented.

The 21st International Workshop on Vertex Detectors
16-21 September 2012
Jeju, Korea

[*]Speaker.
[†]On behalf of the LHCb collaboration

1. The LHCb Upgrade

LHCb [1] is a forward detector, built to study heavy quark flavour physics at the LHC. In addition, the complementary coverage of the solid angle compared with the general purpose detectors at the LHC, makes LHCb an ideal place to study general physics in the forward direction. Currently the detector operates at a luminosity of 3.5 - 4 $\times 10^{32} cm^{-1} s^{-1}$, corresponding to 1.4 - 1.7 interactions per colliding bunch crossing (μ). While this is higher than the design luminosity of the detector, with a μ which is a factor of 4 higher due to the current LHC bunch structure, this is still below what the LHC is able to deliver. After a first phase of operation, from 2010 until 2018, LHCb is anticipated to upgrade significant sections of the detector [2] in order to take advantage of the increase in available luminosity, and to further the physics goals of currently unachievable/statistics limited measurements.

This increase in luminosity cannot be exploited with the existing detector. While the detector has shown itself capable of running far beyond the specifications to which it was built, the current trigger scheme is a limiting factor in the usefulness that an increased collision rate would provide [3]. As with other particle physics detectors in high rate environments, LHCb relies on several levels of triggering to reduce the input collision rate down to an event rate compatible with the offline computing and storage resources available - of order 5kHz. The first level of triggering, L0, uses the fast readout of the muon and calorimeter systems to make a decision about each beam crossing at 40MHz. The result of this decision is to trigger all remaining sub detectors to read out the event, and feed this into the computing farm for more sophisticated analysis. The maximum output rate of this level of triggering is 1MHz.

While this L0 trigger is well suited to physics in the current running conditions, for higher luminosities this step is a limitation in the selection of certain event signatures. Shown in figure 1 is the trigger yield for a selection of B decay channels as a function of luminosity - the saturation in yield for those with hadronic final states is clearly visible. The solution envisaged for the upgrade of LHCb is therefore to remove this step from the trigger chain, and to directly read out all sub-detectors at 40MHz. A software trigger will be used to select from all events those which best match the events required in the offline analyses, taking advantage of the full event information to

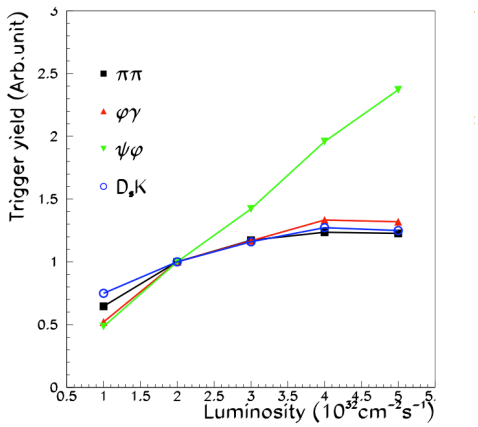

Figure 1: Trigger yield for several decay channels versus luminosity for the current trigger scheme

trigger more efficiently.

2. The Current VELO

The current VELO is a silicon strip detector [4], with modules constructed from back-to-back single sided sensors with the strips oriented to give either a radial or pseudo-azimuthal coordinate. Each module thus provides a single R-Phi space point, used for the reconstruction. The sensors themselves are 300μm thick n^+-on-n float zone silicon, and a station consists of two such modules which form a circular active region perpendicular to the incoming beam, with an aperture radius of 7mm during stable data taking (this is the radius of the module, with the first active strips located at a radius of 8.2mm, and a foil radius, mentioned later, of 5.5mm). During beam injection into the LHC, the two detector halves (21 stations) are retracted by up to 29mm, and are subsequently closed around the beam once stable. Further shielding the detector from beam influence is a 300μm corrugated foil, which prevents the pickup of radio frequency (RF) waves from the beam, allows the passage of mirror currents, and separates the beam and detector vacuua. The foil is shaped around each module such that when the two detector halves are closed there is a small overlap between the active sensor areas, preventing gaps in the acceptance.

The VELO is read out using the Beetle ASIC [5]. The Beetle is an analogue chip, built in 250nm CMOS technology, which samples each input pad synchronously with the LHC 40MHz clock. Local buffering permits the readout of up to 16 consecutive events, with a total output bandwidth limited to 1MHz. Readout is entirely analogue, and pedestal subtraction, common mode corrections and clustering are performed on the off-detector electronics boards, the TELL1 [6]. The channel noise varies with strip pitch, but an overall signal/noise ratio of 18-22 has been maintained across all modules.

3. The VELO Upgrade

As a result of the envisaged trigger scheme for the upgrade, the foremost requirement on the upgraded VELO is full detector readout at 40MHz. The requirements relevant to the analysis of collision data are that the Impact Parameter (IP - the distance of closest approach between a particle's flight path and the interaction point) resolution should improve upon that of the existing detector, and that the track reconstruction has high efficiency and low ghost rates. For the IP resolution both the material budget and geometry will play key roles.

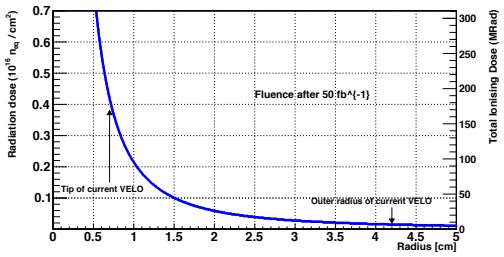

Figure 2: Expected dose for the upgraded VELO after 50 fb^{-1}

Two solutions have been proposed to fulfil these criteria: a strip detector similar to the current VELO, but with increased granularity and faster readout; and a pixel detector building on the TimePix family of ASICs [7]. While both options have their own particular strengths, common issues remain. In particular, the minimum radius of the modules from the LHC beams has to be fixed irrespectively of the module layout, and it is hoped to decrease this from the current 7.0mm to as low as possible (currently estimated of order 5.3mm). The implication of this, aside from increased track density, is radiation tolerance for the active silicon in this region. The expected radiation damage at end of lifetime (corresponding to an integrated luminosity of 50 fb^{-1}) without changing the radius from the existing detector corresponds to a dose of around 5×10^{15} 1 MeV n_{eq} cm^{-2}. This is for the innermost sensitive region of the detector (at a radius of 7.5mm), falling off towards the outer edge of the module as shown in figure 2. With a reduced radius this value sharply increases, leading to a dose of up to 9×10^{15} 1 MeV n_{eq} cm^{-2}.

4. Pixel Detector Options

A pixel detector has been proposed for the upgraded VELO to replace the existing strip detector, with the advantages of a fast and robust pattern recognition for the trigger and radiation hardness. The most challenging aspect of such a detector is the full readout at 40MHz: for the most central chips to the interaction region, hit occupancies of order 6 tracks/chip/bunch crossing are expected. With almost 30MHz of colliding bunches, and the effects of charge sharing on cluster size, this translates to a hit pixel rate in excess of 300MHz. Accounting for a bunch crossing ID of 12 bits, addressing of 16 bits, and ADC values for each pixel of 4 or more bits, the data output from the hottest chips is expected to reach 12 Gbit s^{-1}.

Prototyping of the front end for this chip, VELOPix, has already begun as part of the TimePix3 project [8]. TimePix3 will be the successor chip to TimePix, and contains many features intended for use in LHCb. In particular, the analogue front end for TimePix3 is identical to that required for

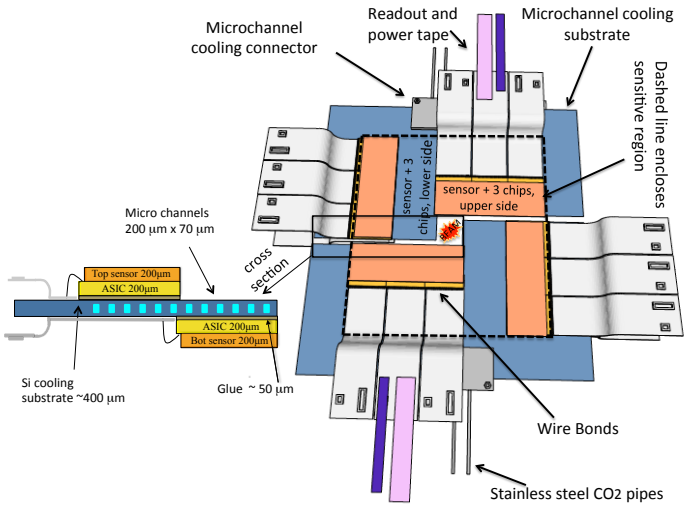

Figure 3: Example layout of a pixel module - shown with microchannel cooling

VELOPix, and will feature data driven readout for a hit pixel rate of up to 40MHz. Submission of the chip is planned for the start of 2013, with chip testing mid-year.

Besides the technically challenging front end, the data driven readout scheme has the drawback that data packets arrive to the periphery of the chip asynchronously. All data from an individual proton-proton interaction must be collected by the off-detector electronics boards, the TELL40 [9], for event processing before data is passed on to the computing farm for reconstruction. This requires a large amount of buffering within the TELL40, in order to account for the maximum latency of the data transmission (somewhere in the region of 10 μs), and the re-ordering into event IDs of the incoming data. Both procedures are costly in terms of the available resources, and second order data pre-processing (such as clustering) may need to be deferred to the computing farm.

The pixel layout is based upon two L-shaped half stations with 12 chips apiece. Each half station contains chips mounted on alternate sides of the mechanical spine, in order to minimise inactive regions at the boundaries between chips and to accommodate the readout lines. Each of the L-shaped modules is additionally separated into two halves, which will be read out separately - further splitting the detector into sets of 3 chips. These 3 ASICs will be mounted on a single silicon sensor, 200μm thick, with elongated pixels at the joint between chips. A schematic of such a station, shown with microchannel cooling option, can be seen in figure 3.

5. Strip Detector Options

An alternative technology choice for the upgrade is a design based on an evolution of the existing strip detector. The main changes are a reduction in the pitch and length of the strips to keep pace with the increased occupancy, and the use of thinner sensors to optimise still further the material budget. This layout has been optimised using the simulated track densities to keep the occupancy per strip roughly even across the full sensor, and the sensor type has been changed from n^+-on-n to n-on-p in order to save in production costs.

One issue yet to be investigated fully in the development of a strip detector is the performance of the pattern recognition, and the subsequent effect that this will have on the online trigger. Initial

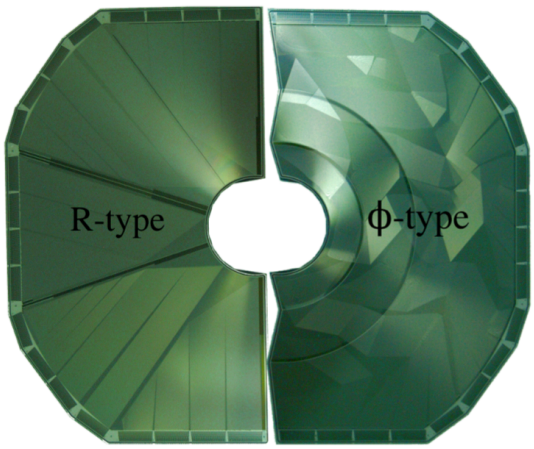

Figure 4: Prototype strip sensors

studies with the present detector indicate that timing issues with the reconstruction are likely be negligible, from both simulations of the current detector with a higher number of interactions per bunch crossing, and from data taken in high luminosity running conditions. The two most important criteria will be the efficiency of track reconstruction, and the number of ghost tracks reconstructed (tracks which are coincidental, and not due to real particles). These will dictate how well the detector will performance in the trigger, and show any effects which may lead to systematics in the offline analyses.

Readout of the strip detector ASIC [11] is expected to be zero suppressed, based on a S/N cut using the locally held threshold values. Each chip will read out 128 channels and, due to the flat channel occupancy, should have a constant output data rate of roughly 2 Gbit s^{-1}. More elaborate on-chip functionality has been considered, such as centre of gravity clustering, though current expectations for the off-detector electronics are such that there will be sufficient resources available to perform these operations without introducing increasingly complex processes on-chip.

Module designs for an upgraded strip detector follow very similarly the design of the current VELO modules: a low-mass paddle with R and Phi sensors mounted on opposing sides of a cooling spine, with the ASICs arranged around the periphery out of the acceptance of the detector. Prototype designs of both sensors, with realistic strip layout, have been fabricated and are in the process of being characterised (received sensors shown in figure 4). It is anticipated to irradiate these devices, as the effects of radiation damage for a strip detector vary slightly compared with the pixel option. This is due to the observed pickup of charge from secondary routing lines, observed in the current VELO [10], which have yet to be quantified at the fluences expected for the upgrade.

6. Cooling

Several cooling options have been proposed for the VELO upgrade, capitalising on the successful performance to date of the bi-phase CO_2 cooling used by the existing detector [12]. The goal for the upgrade will be to re-use the existing infrastructure, with the possibility of upgrading the chiller which controls the cooling fluid temperature. This would place an upper limit of 1kW heat removal from each detector half, with a coolant temperature of between -30°C (current fluid temperature) and -40°C. The requirements on the sensors, for both technology choices, are that the most heavily irradiated parts of the module be kept below -15°C. The evident differences between the two detectors will therefore be the difference in heat dissipation: while the ASICs for a strip detector will sit at the periphery of the module, and far from the inner radius of the sensors, the pixel detector will have heat dissipation until the very edge of the active region. This places a limit of 2 - 3 W power consumption for VELOPix, corresponding to a half detector consumption of 750W.

The most novel solution undergoing testing is the use of microchannel cooling [13]. This technology has been proposed for several upcoming particle physics experiments (NA62 [14], ALICE upgrade [15]), and uses etched channels through a silicon substrate to allow the cooling fluid to pass directly through the mechanical support, with minimal material. A single wafer of silicon is first etched to produce channels of the required size (for LHCb, of order $200\mu\text{m} \times 30\mu\text{m}$), and this wafer is then covalently bonded to a second silicon wafer to seal the channels. Cooling fluid (liquid CO_2) can then be pumped through the channels to allow evaporative cooling directly beneath

the ASICs and other sources of heat dissipation. There are several advantages of this technology over conventional pipe cooling: aside from the fact that the cooling "pipes" are integrated into the substrate, large surfaces for heat transfer are available by using many parallel channels; the use of silicon matches the sensor and ASIC thermal expansion perfectly; the system is low mass, containing no metallic components; there are only small thermal gradients across the module. Despite the many advantageous properties, microchannel cooling has not yet been employed on a large scale experiment, and thorough validation will be required, in particular for use in LHCb where the detector sits within secondary LHC vacuum. Pressure tolerance of > 100 bars must be shown, reliably, in order to validate this and ensure the safety of both the detector and the LHC.

A secondary option for cooling, using more conventional methods, is also under study. Such a detector would consist of a highly conductive substrate, with cooling pipes mounted around the periphery of the module. Several substrate possibilities are available, the most promising of which is high thermal grade diamond. This has been reliably manufactured with thermal conductivities in excess of 1500 $WK^{-1}m^{-1}$, and a substrate thickness of 200μm would be sufficient for mechanical stability and cooling performance.

Both cooling options are, to a great extent, independent of the detector technology. Optimisations for both strips and pixels are under way, with the main consideration being the extent to which the substrate must extend under the module. As the IP resolution is strongly dependent on the material traversed by particles between the first and second measured point [16], retracting the module substrate even just several mm before the second measured point would have an impact on the final detector performance. Thermal simulations to model the temperature gradient and lowest achievable temperature are also under way, to gauge the safety margin below the expected temperature for thermal runaway.

7. RF Foil

One of the most delicate parts of the existing and upgraded detectors is the thin layer of foil which separates the secondary vacuum of the modules from the primary LHC vacuum. This additionally prevents beam-induced effects in the modules, such as pickup of radio frequency (RF)

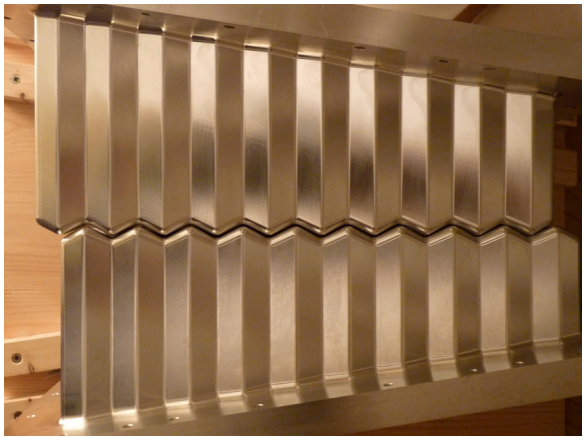

Figure 5: Prototype RF foil, showing both halves (covering 10 stations)

waves, and allows the passage of mirror currents. The current RF foil was produced by mechanically pressing a thin sheet of aluminium into the required shape. With the drive to reduce material within the acceptance for the upgrade, this method was predicted to be unsuitable in the production of even thinner foils. In addition, stresses around the curved inner regions during shaping lead to small variations in thickness uniformity. For the upgrade foil, a new method based on a 5-axis milling technique has been proposed. A single massive block of aluminium is mounted in the milling machine, and a 1.4mm head mills away the material to reveal the inside of the desired RF foil. Once this is complete, the inside of the box is filled with wax, and mounted again in the milling station. The rest of the aluminium is then removed to reveal the outside of the foil, held mechanically stable by the solid wax on the inside. When finished, the whole RF box is heated to remove the wax, with the final foil remaining. This method has been prototyped on several small-scale models, consisting of the full size foil for 4 - 6 modules, and has been very successful. 300μm thick models have been produced, with work being carried out on producing even thinner samples. An example of one such mock-up is shown in figure 5.

8. Testbeams

Testbeams over the last 2 years have focussed on the validation of the Timepix ASIC for charged particle tracking [17], and on investigating performance issues relevant to the design of the ASIC for the proposed pixel detector: VELOPix. To this end, several studies have been undertaken to measure the extent of time walk, charge sharing, and the effects of several front end choices on the spatial resolution. As a vital tool to the many testbeam studies, a dedicated telescope has been built around the Timepix ASIC, and as part of the AIDA project [18] this has been extended as a user facility, with beam tests by a variety of different groups having been already undertaken.

With the aid of the telescope, a range of sensor characterisations have also taken place, investigating sensor technologies relevant to both detector options. Of particular interest has been the investigation of slim edge devices, minimising the pixel-to-edge (PTE) distance in order to reduce the amount of inactive material surrounding the sensors. Several techniques can be used to produce such sensors, including those where the slim edges are created during sensor fabrication [19] and those utilising post-processing of existing sensors [20].

Testbeam plans for the following two years will concentrate on two areas: sensor studies and ASIC validation. In the first instance, LHCb involvement in the RD50 collaboration [21] has allowed the sharing of knowledge on the radiation-induced damage in silicon, and thus sensor irradiations will focus on those issues unique to LHCb, rather than generic R&D. Such considerations will be, as mentioned above, the loss of charge to secondary metal layers, non-uniformity of the received dose, and breakdown characteristics. The more vital line of investigation will be the validation of the different front-ends, particularly at high rate. Of primary interest for the pixel chip will be efficiency measurements made with sustained particle fluxes in excess of 200 MHz. Packet loss rates will need to be verified, in addition to confirming the expected performance in terms of resolution, etc. For the strip chip, pulse shape measurements can be performed outside of testbeams, but the effects of spill-over with high beam intensity and high channel occupancy will need to be performed.

9. Outlook

Realistic module designs for the upgrade of the LHCb VELO have been presented. The baseline technology choice will be defined in the TDR scheduled for Q3 2013, and both options follow roughly the same timescale: prototype ASICs are expected to appear mid-2013, with first full chip iterations expected in 2014. Measurements on radiation-induced effects, in addition to the calculations for necessary data bandwidth and the full detector performance extracted from simulations will be used to gauge which option is most suited to the physics goals of the upgrade. Trigger studies, to determine the online and timing performance, will be vital for this, due to the importance of the trigger in the upgraded detector. Construction of full modules is scheduled to commence 2015/16, with installation during the long LHC shutdown in 2018.

References

[1] The LHCb Collaboration, *The LHCb Detector at the LHC*, *JINST* **3** (2008) S08005

[2] The LHCb Collaboration, *Letter of Intent for the LHCb Upgrade*, CERN-LHCC-2011-001

[3] The LHCb Collaboration, *LHCb trigger system : Technical Design Report*, CERN LHCC 2003-031

[4] The LHCb Collaboration, *LHCb VELO (VErtex LOcator) : Technical Design Report*, CERN-LHCC-2001-0011

[5] D.Baumeister et al., *The Beetle Reference Manual*, LHCb 2001-046

[6] G.Haefeli et al., *The LHCb DAQ interface board TELL1*, *Nucl. Instr. A* **560** (2006) 494-502

[7] X.Llopart et al., *Timepix, a 65k programmable pixel readout chip for arrival time, energy and/or photon counting measurements*, *Nucl. Instr. A* **581** (2007) 485-494

[8] T Poikela et al., *Architectural modeling of pixel readout chips Velopix and Timepix3*, *JINST* **7** (2012) C01093

[9] J. P. Cachemiche et al., *Study for the LHCb upgrade read-out board*, *JINST* **5** (2010) C12036

[10] A.Webber, *Impact of Radiation on the LHCb VELO Sensors*, *Proceedings VERTEX 2012*

[11] C.Parkes et al., *Preliminary Specification of a Silicon Strip Readout Chip for the LHCb Upgrade*, CERN-LHCb-PUB-2012-011

[12] B.Verlaat et al., *CO_2 cooling for the LHCb-VELO experiment at CERN*, *Proceedings of 8th IIF/IIR Gustav Lorentzen Conference on Natural Working Fluids* (2008)

[13] R.Dumps et al., *Micro-channel CO_2 cooling for the LHCb VELO upgrade*, LHCb-TALK-2012-295

[14] M.Fiorini et al., *The Gigatracker: An ultra-fast and low-mass silicon pixel detector for the NA62 experiment*, *Nucl. Instr. A* **628** (2011) 292-295

[15] R. De Oliveira et al., *Development of a polyimide microchannel on-detector cooling system for the ITS upgrade*, CERN-PH-EP-2012

[16] M.Alexander, *Constraints on mixing and CP-violation in the neutral charmed meson system at LHCb*, University of Glasgow Doctoral Theses (2012)

[17] K.Akiba et al., *Charged particle tracking with the Timepix ASIC*, *Nucl. Instr. A* **661** (2011) 31-49

[18] The AIDA Project, *http://aida.web.cern.ch*

[19] J.Kalliopuska et al., *Overview on Measured Properties of VTT's Edgeless Detectors and their use in High Energy Physics*, *Physics Procedia* **37** (2012) 1046-1054

[20] M. Christophersen et al., *Alumina and silicon oxide/nitride sidewall passivation for P- and N-type sensors*, *Nucl. Instr. A* **699** (2013) 14-17

[21] The RD50 Collaboration, *http://rd50.web.cern.ch*