PROCEEDINGS OF SCIENCE

The upgrade of the LHCb Vertex Locator (VELO)

Martin van Beuzekom[*†]
Nikhef, Science Park 105, 1098 XG Amsterdam, the Netherlands
E-mail: martinb@nikhef.nl

The upgrade of the LHCb experiment, planned for 2018, will enable the detector to run at a luminosity of $2{\cdot}10^{33}$ cm^{-2}s^{-1} and explore New Physics effects in the beauty and charm sector with unprecedented precision. To achieve this, the entire readout will be transformed into a triggerless system operating at 40 MHz, where the event selection algorithms will be executed by high-level software in the CPU farm. The upgraded silicon vertex detector (VELO) must be lightweight, radiation hard, vacuum compatible, and has to drive data to the data acquisition system at speeds of up to 3 Tbit/s. This challenge will be met with a new VELO design based on hybrid pixel detectors, positioned to within 5 mm of the LHC colliding beams. The sensors have 55 x 55 μm^2 square pixels and the VeloPix ASIC, which is being developed for the readout, is based on the Timepix/Medipix family of chips. The hottest ASIC will have to cope with integrated hit rates of up to 900 MHz which translates to a bandwidth of more than 15 Gbit/s. Work is in progress to optimise the sensor guard ring design to cope with the irradiation levels, which are highly non-uniform and reach $8{\cdot}10^{15}$ 1 MeV n_{eq} cm^{-2} at the innermost region. The material budget is optimised with the use of evaporative CO_2 coolant circulating in microchannels within a thin silicon substrate. Microchannel cooling brings many advantages: very efficient heat transfer with almost no temperature gradients across the module, no CTE mismatch with silicon components, and a small contribution to the material budget. LHCb is also focussing effort on the construction of a lightweight foil to separate the primary and secondary LHC vacua, the development of high-speed data acquisition boards, and the radiation qualification of the pixel detector modules.

22nd International Workshop on Vertex Detectors,
15-20 September 2013
Lake Starnberg, Germany

[*]Speaker.
[†]on behalf of the LHCb VELO upgrade group

PoS(Vertex 2013)016

1. Introduction

The LHCb detector [1] is a general purpose detector instrumented in the forward region at the LHC. It is specialised in measurements of beauty and charm hadrons and is also well suited for New Physics searches. The key components of the LHCb detector are its tracking system, particle identification system and its flexible trigger. The tracking system consists of a vertex detector around the interaction region, silicon tracking layers in front of the warm dipole magnet, and tracking planes consisting of silicon strips and straw tubes behind the magnet. For particle identification purposes the LHCb detector is equipped with two ring-imaging Cherenkov detectors, a system of calorimeters, and a muon detector. Triggering in LHCb is done in two stages. The first-level hardware trigger (called L0-trigger) uses E_T and p_T information from the calorimeter and muon systems to select 1 MHz of events from the 30 MHz of colliding proton bunches. For every L0-trigger the whole detector is read out and sent to a CPU farm which performs online event reconstruction and selection, and thereby brings down the event rate to 5 kHz which is subsequently written to disk for offline analysis. LHCb has been taking date successfully during LHC run-1 (2010-2013) and has collected a total of 3.5 fb^{-1} of data. The experiment will start its second run (2015-2018) next year in which it is expected to collect another 5-7 fb^{-1}. However, continuing to run at the same luminosity in the years beyond 2018 means that in order to double the statistics we need to run for another five years, which is not very rewarding. Hence it has been decided to upgrade the detector in the long shutdown (LS2) of the LHC in 2018-2019, such that it can cope with a considerably higher instantaneous luminosity.

The current LHCb detector has been designed to run at a luminosity of $2{\cdot}10^{32}$ cm^{-2}s^{-1}, but during the last year of run-1 it has been running steadily at twice this luminosity. Running at an even higher luminosity is not possible due to limitations of the first-level hardware trigger. As shown in Fig. 1, the physics yield in the hadronic decay modes saturates around a luminosity of $4{\cdot}10^{32}$ cm^{-2}s^{-1}.

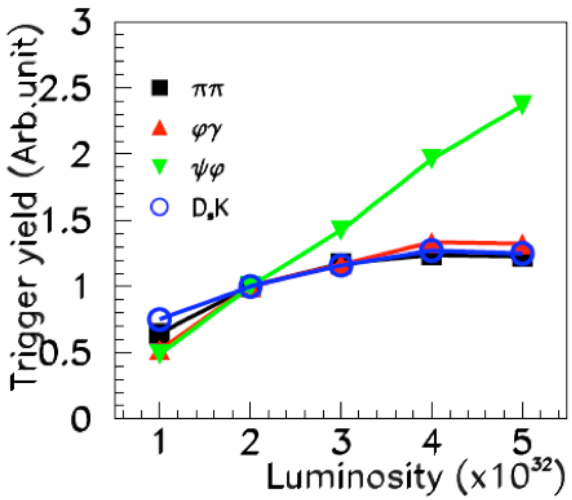

Figure 1: Trigger yield of current the LHCb detector versus luminosity for various decay channels. At higher luminosities the hadronic decay channels saturate which make running at a high luminosity hardly beneficial.

In order to benefit from a higher luminosity the trigger algorithms require, amongst others, information about the decay vertices and hence all the data from the Vertex Locator (VELO). Therefore, the L0 hardware trigger will be removed in the upgraded detector and the data from each sub-detector will be sent to the CPU farm for every bunch crossing. The event selection can then be done by (partial) event reconstruction in the earliest stages of the software trigger. This implies replacing the front-end electronics of virtually every sub-detector. In this paper we will further focus on the upgrade of the Vertex Locator. More information about the upgrade of the complete LHCb detector can be found in [2].

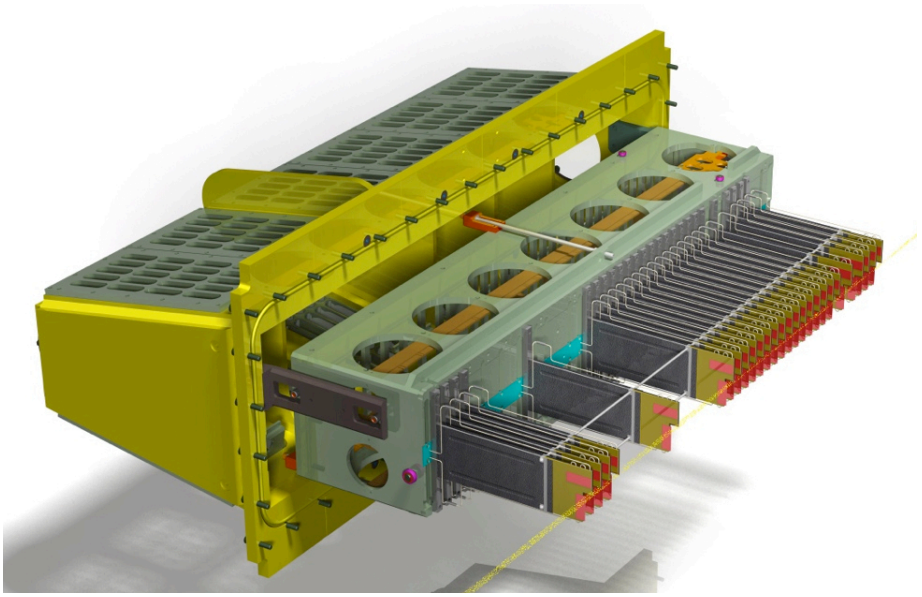

Figure 2: One half of the VELO with 26 silicon pixel modules. Each module has four sensor tiles, two on either side.

2. LHCb and Vertex Locator upgrade

Figure 2 shows the overall layout of the VELO for the upgrade. It consist of 26 stations which are placed perpendicularly to the beam, similar to the current VELO. The 26 modules on either side are placed on a moveable base, which allows us to move the modules away from the beam by 30 mm during injection and energy ramping of the LHC beams. For the upgrade the sensitive elements are silicon pixels with a surface area of 55x55 μm^2 and a thickness of 200 μm. A picture of a single module is shown in Fig. 3. It consists of four hybrid silicon tiles with an active area of 14.1×42.5 mm^2 each. Two tiles are mounted on the top of a mechanical substrate which also provides the microchannel cooling, and the other two on the back. This minimises the gaps in the acceptance which would otherwise be present due to the inactive periphery of the readout chips, and the wirebond pads. Each sensor tile is bumpbonded to three VeloPix ASICs which read out a matrix of 256x256 pixels. The first active pixel is positioned at 5.1 mm from the circulating beams.

The material budget in the active area of a module amounts to about 0.8% X_0 on average. To further reduce the material between the interaction point and the first (and second) measured

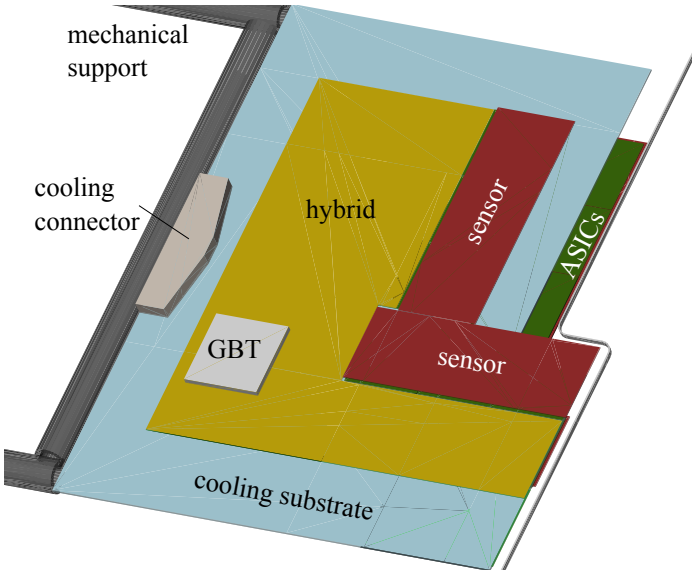

Figure 3: Schematic drawing of a silicon module. A module consist of four silicon sensor (in red), two mounted on either side of the cooling substrate. Each sensor is read out by three VeloPix ASIC (in green).

point the 400 μm thick cooling substrate is retracted by 5 mm. The whole detector is placed in a secondary vacuum which is separated from the beam volume by a thin aluminium box. The total active surface area of the upgraded VELO is 1242 cm^2 and it has about 41 million pixels.

3. Simulated performance

The expected performance of the upgraded VELO has been studied using the default LHCb applications for simulation and reconstruction which are also used in the current experiment. The layout of the upgraded VELO was integrated in the LHCb detector description framework and algorithms for simulating the response of sensor and front-end ASIC, clustering of pixel hits, and pattern recognition were developed. Results are obtained for a minimum bias data set with an average number of interactions per bunch crossing of $\nu = 7.6$ corresponding to a luminosity of $2{\cdot}10^{33}$ cm^{-2}s^{-1}. For comparison, the performance of the current VELO has been simulated in parallel. A detailed description of the simulation and reconstruction software can be found in [3]. The pattern recognition, for which the pixel geometry is highly advantageous over the current R-ϕ geometry, performs with an efficiency of 99%. Two of the key performance figures of a B-physics experiment are the vertex resolution and the impact parameter resolution. An example of the simulated performance is shown in Fig. 4 for both the upgraded and current geometry. The upgraded detector is required to deliver physics performance at least as good as the current detector; as can be seen from these figures it in fact shows an improved performance. The IP resolution of the upgraded VELO shows a significant improvement which is mainly due to the fact that the first measurement is closer to the beam. In addition, the material contribution to the multiple scattering is improved with thinner modules (400 μm compared to 600 μm for the current VELO) and a more

optimal RF foil shape. Note that these plots have been evaluated for running at upgrade conditions, for which the current VELO was not designed.

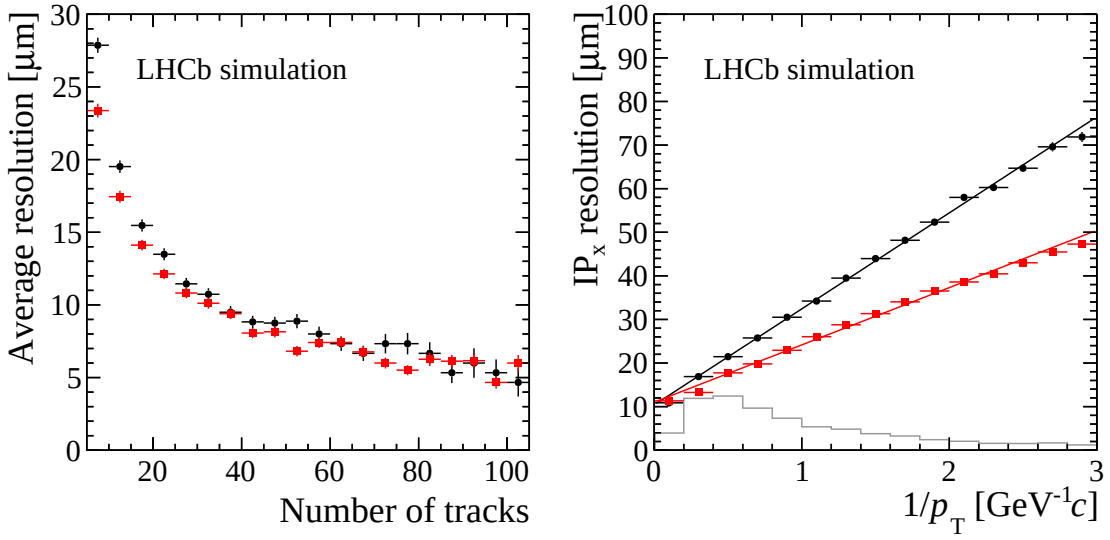

Figure 4: The left figure shows the primary vertex resolution as function of the number of reconstructed tracks. The right figure shows the x resolution of the IP. The current VELO is shown with black circles and the upgrade VELO with red squares, both are evaluated at a luminosity of $2{\cdot}10^{33}$ cm^{-2}s^{-1}, and $\sqrt{s} = 14$ TeV. The resolutions in x and y are similar. The light grey histogram shows the relative population of b-hadron daughter tracks in each p_T bin.

4. Radiation environment and silicon

The silicon modules are positioned perpendicularly to the beam, and the first active pixel is only 5.1 mm away from the beam. This means that even though LHCb runs at a much lower instantaneous luminosity than the general purpose experiments at the LHC, the tip of the sensor will receive a radiation dose which is only a factor 2-3 less than those foreseen for the upgrades of Atlas and CMS. For the upgrade we intend to use 200 μm thick sensors as this thickness is the best compromise between the required maximum depletion voltage at the end of life, and the rigidity of the assemblies which is important to achieve a high yield from the bump bonding. It has not yet been decided whether to use n-on-n or n-on-p type sensors.

A complicating factor for the sensor design is that the fluence across the sensor surface is highly non-homogeneous due to the orientation w.r.t. the beam. At the tip of the sensor, the received fluence at the end of the experiment (after 50 fb^{-1}) is $8{\cdot}10^{15}$ 1 MeV n$_{eq}$ cm^{-2}, while the fluence is a factor 40 lower at the far corner of the same sensor. This means that the full depletion voltage is very different between the most and least irradiated region. When the most irradiated region has to be operated at 1000 V in order to collect a charge of about 8 ke$^-$[4] the depletion voltage in the far corner is only a couple of tens of Volts. The sensor requirement is hence that it has to withstand the maximum operation voltage (1000 V) prior to irradiation. Therefore, we opt for a relatively wide set of guard rings of 450 μm. Currently, the final prototype of the sensor

is being designed which will be manufactured by two different vendors. The performance of the sensors will be verified both prior to, and after irradiation to several fluences.

5. VeloPix ASIC

The readout ASIC for the VELO upgrade, called VeloPix, is a pixel chip based on the Timepix3 ASIC [5]. Both ASICs are designed in the same commercial 130 nm CMOS technology. The VeloPix is a binary readout chip which timestamps each hit when it passes the threshold. Since many tracks give rise to clusters with a size of more than one pixel, it is beneficial in terms of bandwidth to pack the hits into so-called super-pixels. Duplicate information consisting of the 10 bit time-stamp, and most of the address bits then have to be supplied only once per super-pixel packet, thereby reducing the output bandwidth. These super-pixels consisting of 2x4 pixels have fixed boundaries. Although flexible boundaries would further reduce the bandwidth requirements, the technical implementation is too resource intense, and also requires the digital signals to cross the boundaries of the analog front-end, which might introduce additional noise.

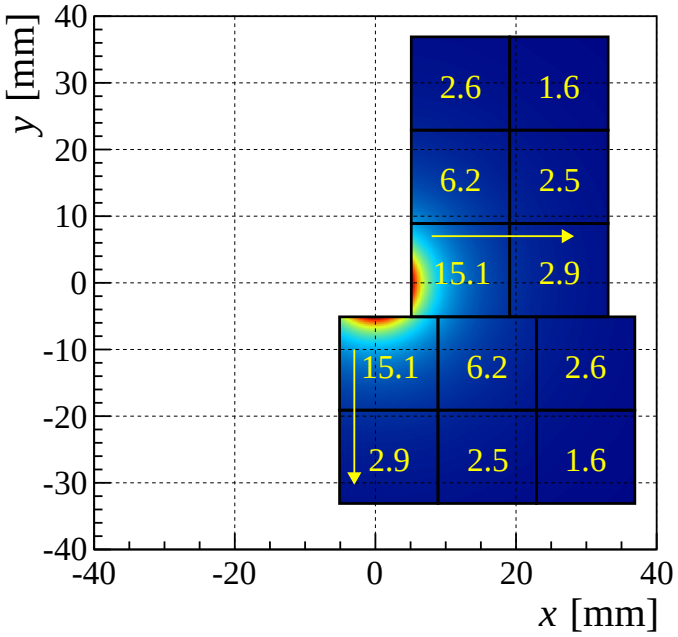

Figure 5: Required bandwidth for the various VeloPix ASICs in a VELO module in Gbit/s. Only the two ASICs closest to the beam have to push out the maximum amount of data and hence require 4 serial links running at 5 Gbit/s. For the other ASICs one or more serial links can be disabled in order to reduce power.

Being close to the interaction region, the hit-rate of the innermost pixels is as high as 40 kHz. For a complete VeloPix ASIC of 256x256 pixels the hit rate amounts to 900 Mhits/s which translates to an output bandwidth of 15 Gbit/s for the ASICs closest to the interaction region. The hit-rate decreases quickly with the distance from the beam ($R^{-1.9}$) and hence only the two ASICs closest to the beam have to deliver the aforementioned performance, see Fig. 5. The readout is done column-wise and hence to reduce the maximum bandwidth in the columns the ASICs are rotated

such that the readout direction always points away from the interaction point, as indicated by the yellow arrows in the figure. Each VeloPix has four high-speed serial links sending data at a speed of 5 Gbit/s over copper. Because of the highly non-uniform track density, only the two (of twelve) ASICs closest to the beam will require four active links to push out the data. For the other ASICs one or more links will be switched off to reduce the power consumption. The available power budget per ASIC is 3 W. At a distance of 5.1 mm from the beam the total ionising dose during the lifetime of the experiment will be about 4 MGy. It has been shown that the chosen 130 nm technology survives this dose [6].

6. Data acquisition

As the silicon modules are close to the beam, and hence face a high radiation dose, the optical to electrical conversion is not done on the silicon modules, but outside the vacuum tank at a distance of about 70 cm. The copper flat cables between module and vacuum flange have been constructed with a special laminate (Pyralux AP-plus) from Dupont that is targeted for these types of high-speed signal transmission applications. The differential pairs are implemented as edge-coupled strip lines, with characteristic impedance close to 100 Ω. Each silicon module with 12 VeloPix ASICs requires 20-22 optical fibres to send its data to the DAQ boards located at a distance of about 200 metres. A DAQ board houses several large FPGAs, one per silicon module. The main task of the FPGA is to decode the incoming data stream (up to 50 Gbit/s), to time-order the pixel packets[1]. Note that the data coming from the VeloPix is already zero-suppressed and hence the data is not reduced but only re-ordered and re-grouped in order to reduce the load on the CPU farm. As final step, data from several LHC bunch crossings are packed into ethernet packets and sent via (up to twelve) 10 Gigabit ethernet links to a large CPU farm. The software trigger will reconstruct the events online and make the selection to bring the event rate down from about 30 MHz to 20 kHz that can be stored on disk.

7. Microchannel cooling

The current VELO pioneered the use of evaporative CO_2 cooling in high energy physics [7]. The sensors are kept at around -7 °C, which is essential to control the leakage current of the silicon sensors after irradiation and prevent degraded performance or even thermal runaway. At the upgrade, there will be a more challenging radiation environment and hence the sensors have to be kept below -20 °C. Moreover there is a marked increase in heat load mainly due to the use of the VeloPix ASICs. The worst case total power of a module is calculated to be 43 W. The cooling remains CO_2 based like in the current VELO detector. For the upgrade, the cooling substrate is integrated directly into the module and forms the mechanical backbone. The front end readout hybrids are glued to either side. In order to efficiently drain the heat the coolant must flow in close proximity to the ASICs. This is achieved by passing cooling liquid through miniature channels etched within a silicon wafer. The trenches are etched to be 120 μm deep in a 260 μm thick wafer

[1] Because of the data-driven readout of the VeloPix, the pixel packets are not sent off chip ordered in time. The order depends on the actual occupancy in the ASIC and the latency can be as large as several hundred bunch crossings (of 25 ns).

and covered with a 140 µm thick lid via hydrophobic bonding. A schematic diagram of the silicon cooling substrate with the routing of the microchannels is shown in Fig. 6.

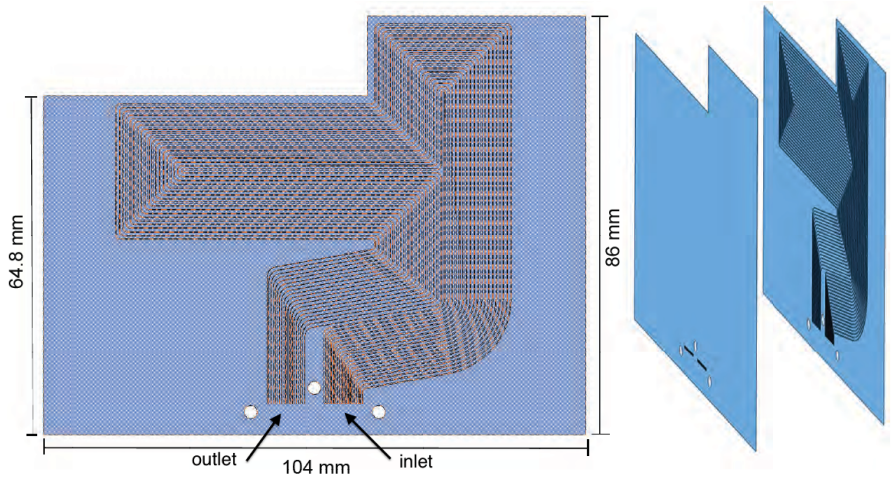

Figure 6: Microchannel layout for one module. The microchannels run underneath the VeloPix ASICs thereby minimising the temperature gradient in the module.

Prototype cooling substrates have been pressure tested and it has been shown that they can withstand a pressure of more than 700 bar, which is a factor ten above the pressure of CO_2 at room temperature. Pressure and thermal cycling tests are ongoing to prove that the substrates do not suffer from fatigue. An additional advantage of a silicon cooling substrate is that there is no difference in the coefficient of thermal expansion because the sensor, readout ASIC and cooling substrate are all silicon. This minimises the thermally induced stress on the module. Tests with heaters glued on a prototype substrate have shown that the heat can be taken out with a temperature difference between heater and coolant of only 7 °C. The nominal operation temperature of the CO_2 is -35 °C and hence the sensor will stay well below the required -20 °C. More information about the microchannel cooling for the LHCb VELO upgrade can be found in [8].

8. RF box

The VELO halves are enclosed in a so-called RF-box. This thin-walled aluminium (AlMg3) box conducts the mirror currents of the beams, shields the silicon modules from the electromagnetic interference of the passing beams, and also acts as a secondary vacuum chamber. The latter is required in order not to pollute the high quality vacuum of the beam volume due to the outgassing of the modules. The side of the box closest to the beam has a complex corrugated shape which minimises the amount of material traversed by the particles. The current box has a 300 µm thick top foil which is pressed into shape. However, this technique can not be applied to create the foil for the pixel detector because of the different geometry. Alternatives have been investigated for many years, and the most promising technique is milling the box out of a solid block of aluminium.

Figure 7 shows a small-scale prototype with a length of about 25% of the final RF-box. The top side of the box has a thickness of 300 μm.

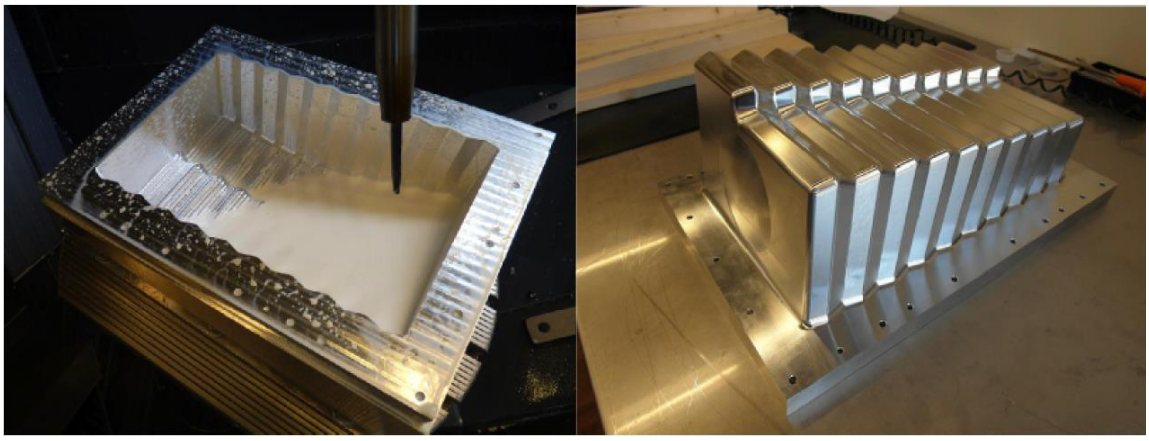

Figure 7: Small-scale prototype RF box milled from a solid block of aluminium

Reducing the thickness to less then 250 μm over the full length of the box by milling is difficult. Therefore, a process is being developed to further reduce the thickness locally by etching it with sodium hydroxide. The first results show that the thickness has been reduced evenly to about 150 μm.

9. Summary

The LHCb VELO has been operating successfully during run-1 of the LHC and will continue to run with its current configuration in run-2 which will end in 2018. In the mean time the experiment is actively working on an upgrade of the detector, to be installed in 2018-2019, which allows us to run at a 5 times higher instantaneous luminosity. The first-level hardware trigger will be replaced by a pure software trigger, and the complete detector will be read out at every bunch crossing. The VELO will be replaced by a pixel detector that is read out by the VeloPix ASICs which have a data output bandwidth in excess of 15 Gbit/s each. The silicon modules will be cooled using microchannels etched in a silicon substrate. Very successful results have been obtained both from the thermal performance tests and the pressure tests. Further challenges lie in the production of the RF box by milling it out of a solid block, followed by local thinning by means of etching. More details about the VELO upgrade can be found in the technical design report that was recently published [9].

References

[1] A. A. Alves Jr et al., *The LHCb detector at the LHC*, JINST **3**, S08005, 2008 [doi:10.1088/1748-0221/3/08/S08005].

[2] R. Aaij et al, *Framework TDR for the LHCb Upgrade*, CERN-LHCC-2012-007. LHCB-TDR-012, 2012.

[3] T. Head et al., *VP Simulation and Track Reconstruction*, LHCb-PUB-2013-018.

[4] G. Casse et al., *Evidence of enhanced signal response at high bias voltages in planar silicon detectors irradiated up to* 2.2×10^{16} $n_{eq}cm^{-2}$, Nucl.Instrum.Meth., **A650**, S56, 2011, [doi:10.1016/j.nima.2010.04.085].

[5] T. Poikela et. al., *Timepix3: a 65K channel hybrid pixel readout chip with simultaneous ToA/ToT and sparse readout*, Proceeding of 15th International Workshop on Radiation Imaging Detectors (IWORID 2013), to be pushlised in JINST.

[6] R. Plackett, X. Llopart, R. Ballabriga, M. Campbell, L. Tlustos, W. Wong, *Measurement of Radiation Damage to 130nm Hybrid Pixel Detector Readout Chips*, Proceedings of Topical Workshop on Electronics for Particle Physics (TWEPP09), CERN-2009-006, [doi:10.5170/CERN-2009-006]

[7] E. Jans, *Operational aspects of the VELO cooling system of LHCb*, Proceedings of this conference

[8] A. Nomerotski et al., *Micro channel evaporative CO_2 cooling for the upgrade of the LHCb vertex detector*, JINST, **8**, P04004, [doi:10.1088/1748-0221/8/04/P04004].

[9] The LHCb collaboration, *LHCb VELO Upgrade Technical Design Report*, CERN-LHCC-2013-021. LHCB-TDR-013, 2013.