



Microchannel cooling for the LHCb VELO Upgrade I

Oscar Augusto de Aguiar Francisco ^{a,b,*}, Wiktor Byczynski ^{a,h,**}, Kazu Akiba ^c, Claudia Bertella ^{a,b}, Alexander Bitadze ^b, Matthew Brock ^d, Bartosz Bulat ^a, Guillaume Button ^a, Jan Buytaert ^a, Stefano De Capua ^b, Riccardo Callegari ^a, Christine Castellana ⁱ, Andrea Catinaccio ^a, Catherine Charrier ⁱ, Colette Charvet ^a, Victor Coco ^a, Paula Collins ^a, Jordan Degrange ^a, Raphael Dumps ^a, Diego Alvarez Feito ^a, Frank Fournel ⁱ, Julian Freestone ^b, Mariusz Jedrychowski ^{a,l}, Vinicius Franco Lima ^e, Abraham Gallas Torreira ^f, Wouter Hulsbergen ^c, Daniel Hynds ^c, Gonzalo Arnau Izquierdo ^a, Pawel Jalocho ^d, Eddy Jans ^c, Malcolm John ^d, Nathan Jurik ^{a,d}, Thomas Latham ⁿ, Alexander Leflat ^g, Edgar Lemos Cid ^f, Rolf Lindner ^a, Alessandro Mapelli ^{a,j}, Abhijit Mathad ^m, Andy Morris ⁿ, Jerome Noel ^a, Andrei Nomerotski ^k, Rui de Oliveira ^a, Martijn van Overbeek ^c, Chris Parkes ^{a,b}, Paolo Petagna ^a, Alexandre Porret ^a, Denis Renaud ⁱ, Erno Roeland ^c, Giulia Romagnoli ^a, Eric Rouchouze ⁱ, Krista de Roo ^c, Freek Sanders ^{a,c}, Thomas Schneider ^a, Heinrich Schindler ^a, Burkhard Schmidt ^a, Andreas Schopper ^a, Luke Scantlebury-Smead ^d, Miranda Van Stenis ^a, Peter Svihra ^{a,b}, Benoit Teissandier ^a, Jean-Francois Teissier ⁱ, Xavier Thery ^a, Eric Thomas ^a, Bart Verlaat ^a, Mark R.J. Williams ^o

^a CERN, Geneva, Switzerland

^b The University of Manchester, Manchester, United Kingdom

^c Nikhef, Amsterdam, Netherlands

^d University of Oxford, Oxford, United Kingdom

^e University of Liverpool, Liverpool, United Kingdom

^f Universidad de Santiago de Compostela, Santiago de Compostela, Spain

^g Institute of Nuclear Physics, Moscow State University, Moscow, Russia

^h Cracow University of Technology, Krakow, Poland

ⁱ CEA, LETI, Grenoble, France

^j EPFL, Lausanne, Switzerland

^k Brookhaven National Laboratory, Upton, NY, United States of America

^l AGH University of Science and Technology, Krakow, Poland

^m Universitaet Zurich, Zurich, Switzerland

ⁿ University of Warwick, Coventry, United Kingdom

^o University of Edinburgh, Edinburgh, United Kingdom

ARTICLE INFO

Keywords:

Microchannel cooling
Bi-phase CO₂
Silicon wafer bonding
Thermal figure of merit
VELO upgrade

ABSTRACT

The LHCb VELO Upgrade I, currently being installed for the 2022 start of LHC Run 3, uses silicon microchannel coolers with internally circulating bi-phase CO₂ for thermal control of hybrid pixel modules operating in vacuum. This is the largest scale application of this technology to date. Production of the microchannel coolers was completed in July 2019 and the assembly into cooling structures was completed in September 2021. This article describes the R&D path supporting the microchannel production and assembly and the motivation for the design choices, together with the achieved fluidic and thermal performance. The Thermal Figure of Merit of the microchannel coolers is measured on the final modules to be between 1.5 and 3.5 K cm² W⁻¹, depending on glue thickness. The microchannel coolers constitute 18% of the total radiation length of the VELO and less than 2% of the material seen before the second measured point on the tracks. Microchannel cooling is well

* Correspondence to: Particle Physics Group, Department of Physics and Astronomy, The University of Manchester, Schuster Laboratory, Brunswick Street, Manchester M13 9PL, United Kingdom.

** Correspondence to: EP Department, CERN, 1 Esplanade des Particules, 1217 Meyrin, Switzerland.

E-mail addresses: oscar.augusto@cern.ch (O.A. de Aguiar Francisco), wiktor.byczynski@cern.ch (W. Byczynski).

<https://doi.org/10.1016/j.nima.2022.166874>

Received 12 January 2022; Received in revised form 4 April 2022; Accepted 13 May 2022

Available online 26 May 2022

0168-9002/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

sued to the VELO implementation due to the uniform mass distribution, close thermal expansion match with the module components and resistance to radiation.

1. Introduction

The LHCb collaboration is replacing its Vertex LOcator (VELO) detector during Long Shutdown 2 of the LHC as part of a major Upgrade programme to allow the experiment to operate at an instantaneous luminosity of $\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, five times higher than that of previous runs [1]. The VELO requires very precise tracking and fast pattern recognition in order to reconstruct collisions and decay vertices in real time as the first step of the LHCb trigger decision. A new hybrid pixel detector replaces the original silicon strip detector [2,3], equipped with a lightweight and highly thermally efficient cooling solution of evaporative CO_2 circulating in microchannels embedded in a silicon cooler.

The sensing elements of each module consist of four planar silicon sensors, each bump-bonded to three VeloPix [4] readout ASICs. A diagrammatic view of the module and the arrangement of the two VELO halves is shown in Fig. 1. The sensors are approximately $43.5 \times 15 \text{ mm}^2$ in size and are arranged with two on each side of the module. The principal sources of heat on the module are the VeloPix ASICs which can dissipate up to 1.9 W ($\sim 1 \text{ W/cm}^2$) each depending on the luminosity and the operation mode, the sensors, which may dissipate about 1 W per tile after irradiation (for an integrated luminosity of 50 fb^{-1}), and the GBTx hybrids [5], one on each side of the module, which dissipate approximately 1.6 W each. The total nominal power budget per module is conservatively 30 W, giving a total of 1.56 kW for the 52 modules of the whole VELO. The heat is dissipated in vacuum, requiring a suitably efficient cooling system. The sensors will experience an integrated irradiation of up to $\phi = 8 \times 10^{15} \text{ 1 MeV n}_{\text{eq}}/\text{cm}^2$, with a highly non-uniform distribution concentrated at the tip of the sensor. The sensor tips need to be kept at $-20 \text{ }^\circ\text{C}$ or below in order to reduce reverse current and reverse annealing, while being directly bump bonded to the heat generating ASICs, placing stringent requirements on the cooling. At the same time the material budget must be kept to a minimum, in order to allow accurate VELO reconstruction of low momentum tracks originating from heavy flavour decays and to be able to identify displaced tracks and vertices. Low material is particularly crucial at the tip of the module, where it affects the extrapolation to the original primary or secondary vertex, and for this reason the cooling substrate is retracted from the tip to minimise the material. The material budget is also constrained over the full module height, due to the fact that the downstream modules are entirely within the LHCb acceptance [1]. The cooling system must be capable of continuous operation such that the sensors can be kept cool to suppress reverse annealing even during long shut down periods of the LHC. Finally, the VELO system is enclosed within the secondary vacuum of the LHC and will be almost completely inaccessible for the decade of scheduled running time, hence the reliability of the system is paramount.

1.1. Cooling options overview

Based on the successful operation of the CO_2 cooling based on a 2-phase accumulator controlled loop [6,7] implemented for LHCb Run 1, this principle is maintained for the cooling plant of the VELO Upgrade. Two identical CO_2 cooling plants, which share a common chiller, have been built, capable of simultaneously removing 1.6 kW of heat from the VELO and 5.4 kW of heat from the silicon upstream tracker at $-35 \text{ }^\circ\text{C}$ [8] for the set temperature at the cooling plant accumulator. As for the previous VELO an evaporator system [9] must be designed to distribute the coolant to the modules, however the previous solution for the thermal interface of cooling pipes embedded in Aluminium “cookie” heat spreaders [10] does not satisfy the pixel upgrade requirements due to excessive additional material in the acceptance. Various

new options were considered and are discussed in the next paragraphs, and the rest of this article is devoted to describing the microchannel option, which was prototyped and eventually adopted.

At the start of the VELO Upgrade R&D programme, a wide range of cooling interfaces were discussed, ranging from the use of LN_2 cooling pipes and cryopanel, which was the solution adopted for BTeV [11], through to the proposal of the use of diamond substrates, which can be effective heat spreaders in a range of applications [12,13]. Other options were prototyped to various levels of production readiness [14, 15].

One option consisted of machining precision grooves in a “Shapal” (aluminium nitride) plate into which stainless steel pipes were embedded, after robot deposition of heat conductive glue. The Shapal plate provided mechanical support and excellent thermal conductivity. Orifices for the expansion step were provided in an inlet block grouping all the channels. A second solution consisted of manufacturing 3D printed titanium substrates. Such an approach is in principle cheap and provides a flexible layout and fast turnaround for prototyping. It benefits from vast experience in industry and the material can be welded or brazed to the cooling pipes. Inlet restrictions can be easily incorporated into the design. At the time of the VELO development feature sizes of $100 \text{ }\mu\text{m}$ were readily available in a suitable material.

Both the Shapal and printed Titanium options reached a mature level, and due to the tremendous challenges posed by the microchannel solution were maintained as backup options until a relatively late stage. Other proposals included integrating the cooling pipe into a “pocofoam” substrate, or force fitting a cooling tube into TPG sheets, however after initial prototyping these ideas were not taken further. After internal review the silicon microchannel solution was chosen for the VELO upgrade, with the other solutions rejected due to the risk of a too long R&D trajectory to reach the required cooling performance.

1.2. Microchannel cooling overview

The idea of integrating coolant in microscopic channels was first proposed [16] as a method of achieving high performance cooling of planar integrated circuits, with water as the coolant. For laminar flow in confined channels, the heat transfer coefficient scales inversely with the channel width, ultimately limited by the coolant viscosity. The microchannel concept with channels integrated directly into a silicon wafer is very attractive for active cooling of planar silicon components for many reasons above the high thermal efficiency provided. It benefits from no coefficient of thermal expansion mismatch between the substrate and the silicon components to be cooled, which would otherwise have to be absorbed by an intermediate glue layer. In this case the silicon components are the hybrid pixel tiles, which are the major source of heat on the module and also the most critical in terms of temperature; the other active components are mounted on a flex circuit with a flexible glue interface. The radiation length is naturally low due to the small size of the channels and the material distribution is very homogeneous compared to a solution incorporating bulky pipes. The silicon itself has a high thermal conductivity; $\sim 150 \text{ W/(m K)}$ at room temperature, improving to $\sim 190 \text{ W/(m K)}$ at a typical operational temperature of $-25 \text{ }^\circ\text{C}$ [17]. By using an appropriate layout of the channels, the cooling can be routed precisely to the areas where it is needed. The natural planar format is an excellent match to the planar components to be cooled, maximising the geometric thermal contact. Using standard silicon processing techniques additional features are available, such as metallisation of alignment marks and solder footprints¹ and the wafers can be diced to produce final coolers with complex shapes.

¹ although not used for LHCb, features such as metallic traces may be used in the future




Fig. 1. Diagrammatic view of the module layout (left), highlighting the hybrid pixel sensors (two on each side) and the optical link chip (GBTx) hybrid (one on each side), which are the main power-dissipating elements. Module layout in LHCb (right). The shaded region indicates the LHCb acceptance, showing that the inactive areas of the downstream modules fall in the acceptance.




Fig. 2. The left side of the figures shows the typical channel shape for the bi-phase CO_2 microchannel cooling implementation. The pressure drop at the point where the channel expands should bring the coolant to the saturation point as it enters the region of the detector to be cooled. The diagram on the right illustrates the principle of the Two-Phase Accumulator Controlled Loop (2PACL) cooling concept used in LHCb [6].

1.3. Microchannel cooling for the LHCb VELO and synergy with parallel R&D efforts

For the LHCb VELO implementation, the choice has been made to use microchannel cooling in combination with bi-phase CO_2 as the coolant. CO_2 is ideal as a coolant due to its high latent heat capacity and low viscosity, and it is radiation hard. Bi-phase cooling, where the heat from the electronics is absorbed by the evaporating liquid CO_2 , gives powerful cooling capability and is stable against changes in power consumption of the module during running. As illustrated in Fig. 2, the CO_2 enters the microchannel structure as sub-cooled liquid through a restriction, before entering the main channel where it starts to boil. A fraction of the liquid is then evaporated while flowing under the heat sources. The VELO modules are designed for a typical flow of 0.4 g/s , which for 30 W heat dissipation would give a final vapour quality of approximately 25%, which gives a good thermal efficiency while still remaining safely away from ‘dry-out’ conditions, where the CO_2 completely evaporates before the end of the channel. The restriction dominates the fluidic resistance and so provides excellent control, ensuring even flow in the main channels and avoiding instabilities. The coolant temperature decreases from the input to the output of the module due to the bi-phase nature of the cooling associated with the variation of pressure to keep the coolant flow. A major challenge of the system relates to the operational pressures; for normal operation at $-30 \text{ }^\circ\text{C}$ the system is at $\sim 14 \text{ bar}$, while at room temperature ($20 \text{ }^\circ\text{C}$) the pressure reaches approximately 57 bar . For safety reasons the entire system must be validated to 186 bar .

A cross section of the microchannel structure as implemented in the VELO modules is shown in Fig. 3. By using the microchannel layout a very homogeneous material distribution can be achieved and the coolant can be delivered directly under the heat sources. There is a distance of just $140/240 \text{ }\mu\text{m}$ between the coolant and the surface of the microchannel cooler. The thermal barrier is dominated by the glue layer, which has a typical thickness of $60 \text{ }\mu\text{m}$, and low thermal gradients can be achieved across the module. The cooling is found to be effective enough that it is possible to withdraw the microchannel 5 mm from the tip of the module, in order to keep the material budget to the minimum for the first measured points on each track.

As this was an emergent technology at the time of the adoption into LHCb, a long R&D path was necessary to overcome the main challenges. At the start a suitable facility had to be found in order to be able to make prototypes and demonstrate the basic principles. Following this a manufacturer had to be identified capable of producing microchannel coolers with such large dimensions with high enough quality to satisfy the pressure requirements. The R&D also had a large focus on the method of attaching the fluidic connector, and establishing suitable quality control processes to ensure long term operational stability.

The work described in this paper was carried out by the LHCb VELO collaborating institutes, and also benefited from a high degree of synergy with other experiments, in particular the NA62 GigaTracker [18], which implements a pioneering microfluidic cooling system with a different coolant and with different requirements, but has many specifications in common. At CERN the R&D was supported by the EP-DT group and the collaboration set up between CERN and EPFL which




Fig. 3. Cross section of the VELO module (not to scale). The sensors overhang the cooler at the tip of the module by 5 mm in order to keep the material budget to the minimum for the first measured points on each track.

allowed quick prototyping at the MicroNano Technology centre. The EP-DT group is continuing to conduct research into questions related to an effective and reliable use of micro-structured thermal management devices for HEP detectors, working in the framework of the CERN strategic R&D programme on technologies for future experiments [19].

2. Microchannel technology prototyping and design optimisation

The principle of microchannel manufacture is to etch trenches into the surface of a silicon wafer then atomic bond a so called ‘cap’ wafer to close the channels. Suitable exit and entry holes are incorporated, typically by etching after bonding, such that liquid can circulate through the capillaries. To fabricate such a system many steps are involved, including photolithography, etching, chemical mechanical polishing, bonding and dicing. The substrates may be silicon, plastic, or glass. For the VELO upgrade application both wafers in the final module are silicon. However glass of various thicknesses has also been used as a cap wafer for several of the R&D steps, since samples can be produced by the more accessible anodic bonding process, avoiding the need to go to industry,² and have the advantage of allowing a clear view of the channels.

2.1. Proof of principle - Snake I

For the reasons described above the first prototypes which were made to test the microchannel geometry and measure the heat and flow characteristics were produced with silicon-pyrex bonding. For these samples the heat load is applied on the silicon side only, due to the much poorer thermal conductivity of the glass. This method has the advantage that the channels can be directly observed through the glass. This prototype was tested with circulating CO₂ supplied by a mobile cooling plant which provided liquid CO₂ at -40 °C with a pressure differential between the inlet and outlet of ΔP~4 bar. It was demonstrated that, with a ΔP~3 bar and a CO₂ mass-flow of 0.15 g/s, 12.9 W of power could be removed from the module before the CO₂ enters the ‘dry-out’ condition. The total heat removal was limited by the pressure difference which could be supplied by the cooling plant. At maximum heat dissipation, the temperature differential between the heater and the output was only 3 °C. From this the thermal figure of merit (TFM) of the system can be determined to quantify the effectiveness of the cooling design. The measured value was

$$\text{TFM} = \frac{\Delta T_{\text{heater-output}}}{\text{Power}/\text{cm}^2} = 1.5 \text{ K cm}^2 \text{W}^{-1} \quad (1)$$

which demonstrates the high cooling efficiency potential of the microchannel design. For comparison, conventional designs with cooling ledges can typically achieve a TFM of 20 K cm² W⁻¹, while highly integrated pipe-structure designs aim to reach values of TFM ~ 12 K cm² W⁻¹ [7]. This work was published in [20].

² The samples used for the VELO R&D steps were produced at the MicroNano Technology centre at EPFL, Lausanne, Switzerland.

2.2. Test structures

After the successful proof of principle a second round of prototyping was undertaken with multiple objectives. A large number of silicon-silicon test samples were produced, to investigate the impact of design choices on the rupture resistance, to qualify different silicon bonding processes, and to test small scale circulating samples. In addition, silicon-glass structures were produced to investigate the resistance of the silicon cover thickness to rupture and the dependence on the channel width.

The LHCb-style silicon wafer layouts produced are illustrated in Fig. 4. For these samples all channels were etched to a depth of 70 μm with a width of 200 μm for the main channels and 30 μm for the inlet restrictions. The total thickness of the samples was 400 μm. The CO₂ is supplied through an inlet hole with diameter 1.6 mm to supply all the channels, and the samples have different sizes for the outlet manifold, with P4 having no outlet manifold at all. In addition a number of smaller samples were produced with a variety of channel widths and silicon thicknesses for more general investigations, referred to below as ‘small samples’. Sample sets from two different bonding processes were received, *hydrophilic* and *hydrophobic* which differ in the bonding strength and the complexity of the surface preparation [21,22] as illustrated in figure Fig. 5:

- The hydrophilic bonding process is an industry-wide technique that typically starts with one surface having a thermal oxide layer and the other surface either a thermal oxide layer or native oxide on top of a silicon substrate. A wet chemical treatment decontaminates the surfaces and promotes the formation of silanol groups (Si-OH), covered with a few monolayers of water. As two wafers are pressed together they spontaneously bind to each other via hydrogen bonds, capillarity and van der Waals forces. The final bonding is achieved by condensation of silanol groups to give strong covalent Si-O-Si siloxane bonds. These bonds initially form mainly at the asperity contact points between the wafers, but their density is drastically increased with annealing steps, which also serve to mobilise the oxide layer to flow and fully close the interface. The siloxane bond formation releases water molecules, some of which diffuse into and oxidise the silicon, forming an amorphous SiO₂ layer between the wafers with H₂ gas produced as a by-product. The presence of the SiO₂ layers helps to prevent bubble formation as it can absorb the H₂. A disadvantage of the hydrophilic bonding is that the presence of the thin SiO₂ layer could lower the sample rupture resistance.
- The hydrophobic bonding process proceeds by removing the silicon surface native oxide with a chemical treatment using hydrofluoric acid and passivating with hydrogen atoms just before the bonding step. The spontaneous bonding adhesion proceeds via van der Waals forces between the Si-H groups, with relatively low attractive forces, followed by annealing steps which allow




Fig. 4. (a) Silicon microchannel six-inch test wafer design produced at Southampton University Nanofabrication Optoelectronics (b) Scanning Acoustic Microscope image of microchannel eight-inch wafer produced at CEA-LETI Grenoble (c) Detail of the P1, P2, P3, P4 test structures on the CEA-LETI wafer.

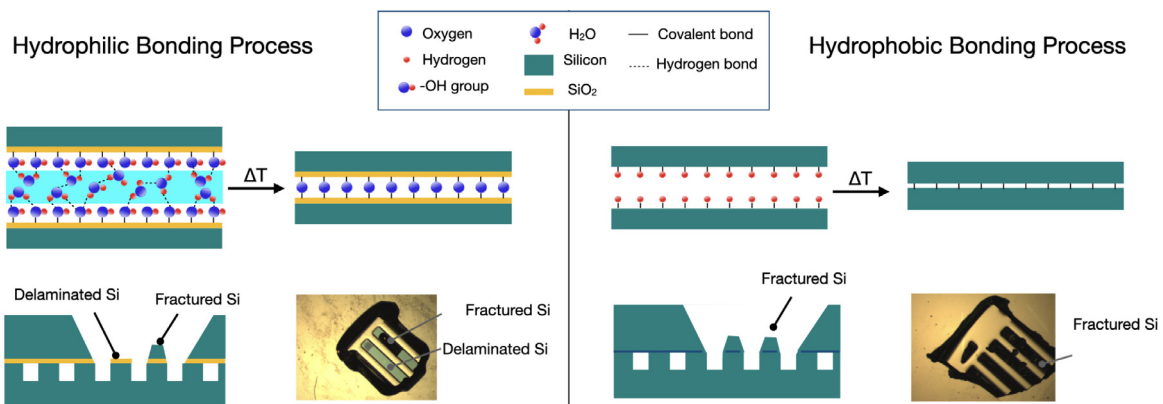


Fig. 5. Illustration of the different characteristics of hydrophilic (left) and hydrophobic (right) bonded samples. The types of bonds connecting the silicon wafers are illustrated before and after annealing. The typical types of breakages observed in the different sample types are illustrated for each bonding type with a diagram and a photo.

the formation of pure covalent Si-Si bonds and promote atomic mobility to close the interface. The disadvantage is that the surface quality must be of a much higher standard and prepared in a higher specification clean room as the adhesion energy is quite low. Moreover specific annealing steps are needed to prevent also hydrogen bubbles formation as there is no oxide layer present to absorb the gas. The advantage is a near perfect homogeneous union of the silicon lattice without a glass layer in between, giving an intrinsically stronger bond.

The small samples were used to provoke ruptures and observe the behaviour of the hydrophilic and hydrophobic samples. The characteristics of the samples and the typical breakages seen are illustrated in Fig. 5. The breakages seen in the hydrophilic bonding samples occurred both in the fusion layer and in the lattice, whereas the breakages from the hydrophobic bonded samples occurred only in the silicon lattice [1].

Pressure tests were performed with water on the P1-P4 test structures. It was possible to provoke breakages, with any breakages below 300 bar being seen at the large output manifold, not within the channels. A similar pattern of rupture characteristics was seen as for the small samples. For the P1-P3 samples a pattern was seen of greater resistance in the hydrophobic samples. For the P4 hydrophilic samples it was not possible to provoke any breakage when testing up to 300 bar, which was the available experimental limit at the time this test was performed. The results are shown in Fig. 6. To permit tests to higher pressure, silicon-silicon samples were clamped to reinforce the fluidic connection such that the integrity of the microchannels could be pressure-tested in isolation. Hydrophilic samples were seen to rupture in the body of the microchannel network at around 400 bar. The hydrophobic samples withstood the maximum pressure of the pump,

which was 700 bar. The results of these tests drove the decision to remove the large output manifold completely (see Section 2.3, and showed that with this design change both hydrophilic and hydrophobic samples were sufficiently strong to be satisfy the LHCb requirements. For production the hydrophilic process was chosen to optimise yields.

Dedicated silicon-glass samples were used to investigate the dependence of the pressure resistance on the thickness of the silicon covering the cooling channel. When measuring the thickness between the coolant and the wafer edge in the etched wafer, this is referred to as $cover_{etch}$, and when measuring the corresponding thickness in the cap wafer, this is referred to as $cover_{cap}$. Rupture tests were performed on a number of samples of different silicon thickness and microchannel width. A clear relationship between rupture pressure and $cover_{etch}$ thickness is for all three channel widths tested. The 500 μm wide channels require a $cover_{etch}$ thickness of approximately 200 μm or more to be able to hold the pressure to the safety limit. However the 200 μm wide microchannels, which is the nominal width for LHCb, tolerate pressures of more than 200 bar even with a 30 μm $cover_{etch}$ thickness. For the nominal LHCb design thickness of 140 μm the rupture pressure obtained by extrapolation is beyond the scale of the plot. A further set of samples were used to investigate the rupture pressure as a function of channel width for $cover_{etch}$ and $cover_{cap}$ thicknesses corresponding to the nominal 140 μm and higher. The rupture pressures rise to high values for lower channel widths, driving the choice of a maximum channel width of 200 μm for the LHCb implementation as being a safe limit. These results are illustrated in Fig. 7.

2.3. Design optimisation - Snake II

Following the first round of prototyping a second large scale prototype, dubbed "Snake II", was produced, with the aim of implementing

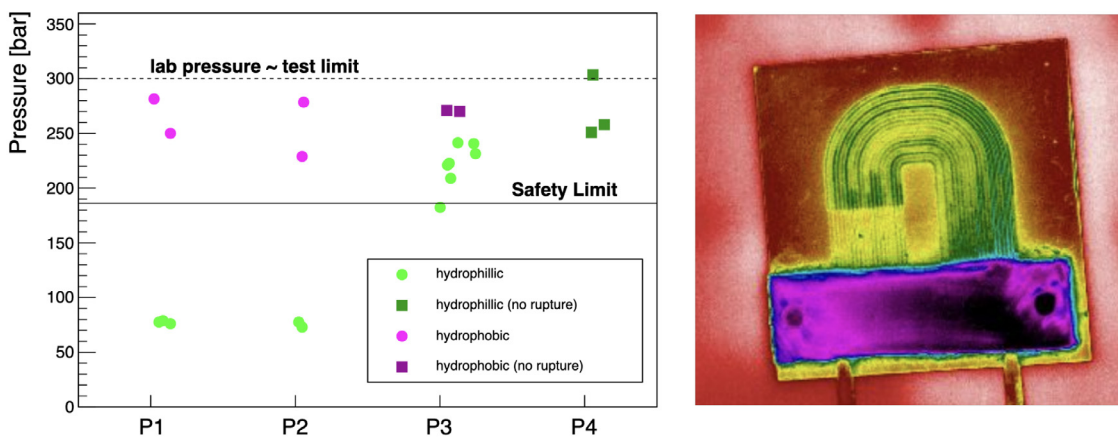


Fig. 6. (left) Rupture tests of samples P1 to P4. The P4 samples surpass the safety test for both hydrophilic and hydrophobic bonding. “No rupture” indicates that the equipment went to the highest pressure possible on that particular test without being able to provoke a rupture. (right) Infra-red camera image of water circulating in silicon-silicon bonded sample.

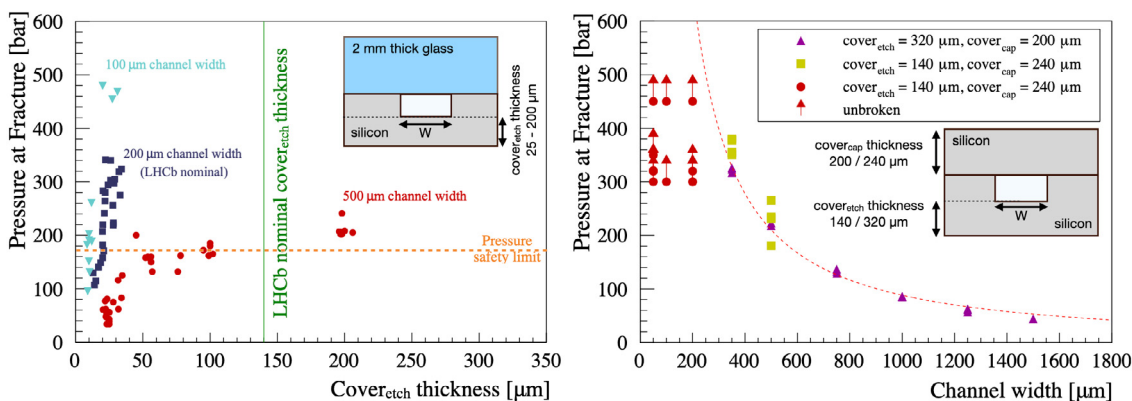


Fig. 7. Results of investigations to determine dependence of rupture pressure on design parameters. Left: Dependence on cover_{etch} thickness. It can be seen that the LHCb cover_{etch} thickness of 140 μm is expected to be very resistant for the nominal channel width of 200 μm. Right: Dependence on channel width, for cover_{etch} and cover_{cap} thicknesses close to the LHCb nominal. The line is an empirical fit to the data. The red points with arrows did not break up to the pressure shown. For this data set, those samples which broke all broke in the cap wafer. Based on these data a maximum channel width of 200 μm is set.




Fig. 8. Design optimisation following prototyping rounds. Left: Optimisation of channel dimensions. Right: Optimisation of manifold and connector region.

and testing the major design changes before the manufacturing step. The first change was to improve by a factor four the fluidic resistance without reducing the rupture resistance. This goal was achieved by moving to a two step etching procedure; in the first step etching the restrictions and main channels, and in a second step, the main channels only. The main channel width was maintained at 200 μm, but the depth was increased to 120 μm, and the restrictions were made square, with dimension 60×60 μm², decreasing the resistance and reducing the chance of clogging. The issue of sensitivity to rupture in the manifold was removed by moving the manifold inside the connector, such that no channel in the silicon has a width greater than 200 μm. The CO₂ is therefore distributed within the connector, and each channel has its own inlet and outlet. These design changes are illustrated in Fig. 8.

The Snake II was successfully produced (see Fig. 9) with silicon-pyrex bonding technology. It featured 19 channels with similar length and fluidic resistance. It was equipped with thermal mockups made of 300 μm thick silicon, metallised to simulate the heat density in the irradiated sensors, and thin film deposit heaters designed to mimic the ASIC heat distribution, all attached with conductive tape.³ The heaters were all mounted on the silicon side of the Snake II. As the test was designed to check the fluidic capabilities of the channels, a simple metallic clamp and O-ring was used for the connector and inlet/outlet cooling tubes. The system was able to provide stable cooling up to 60 W for temperatures between -25 °C and 15 °C set at the cooling

³ Tape 3M 9461P.




Fig. 9. Snake II prototype design and manufacture.




Fig. 10. Layout of the final microchannel design, overlaid with the hybrid pixel tiles and heat producing ASICs (GBTx and the smaller GBLD). The tiles on the front of the cooler are shown in yellow, those on the back in pink, and the position of the GBTx/GBLD chips on both sides of the module in orange. Scanning electron microscope and optical microscope images of details from the production are also shown.

plant accumulator and with 60 W power dissipation giving a factor two safety margin. At the lowest temperature the flow was approximately 0.3 g/s for a ΔP of 4 bar. The flow characteristic was found to be fairly well described, assuming that the pressure drop is dominated by the flow through the narrow section of the microchannels, by $F \sim \Delta P \times C_v$. The flow coefficient C_v for Snake II was found to be approximately 2.1 g/s/ $\sqrt{(\Delta P)}$, three to four times greater than that measured for Snake I. These tests validated the design principles for the LHCb microchannels.

3. LHCb implementation

3.1. LHCb microchannel design

The microchannel cooler design as implemented for LHCb and manufactured by CEA-LETI is illustrated in Fig. 10. Viewed from the connector attachment side, the CO_2 enters as subcooled liquid on the right, passes through a 40 mm long restriction, and ideally starts boiling in the main channels just before flowing under the first ASIC. The channels are routed to pass under the heat dissipating hybrid pixel sensors and the GBTx hybrid. There are a total of 19 channels and every bend is designed with the same radius (0.5 cm) to ensure an equivalent pressure drop due to bending on all the channels. The dimensions of the 19 main channels are $120 \times 200 \mu\text{m}^2$, with a pitch varying between 450 μm (at the outlet), 700 μm (in the main body) and 990 μm (under the GBTx hybrid). The length of the main channels (after the restrictions) varies between 231 and 292 mm. The inlet restrictions have dimensions of $60 \times 60 \mu\text{m}^2$ and a pitch of 450 μm . The total thickness of the wafer is

500 μm , and the channels are positioned asymmetrically in the final wafer, with 140 μm silicon between the channel and the outside of the wafer on the connector side ($\text{cover}_{\text{etch}}$) and 200 μm between the silicon and the outside of the wafer on the non-connector side ($\text{cover}_{\text{cap}}$). The inlet and exit holes are identical for all channels with dimensions of approximately $200 \times 600 \mu\text{m}^2$. The cooler has similar dimensions as Snake II but the positions where the sensors will be glued have been rotated by 45 degrees around the beamspot. This decision, taken to ease the VELO insertion, however poses an additional challenge for the cooler production as the introduction of an internal angle to be cut out prevents saw dicing being used in the cooler production.

The principal process steps are illustrated in Fig. 11. The starting material is double polished eight inch diameter silicon wafers with total thickness variation of $< 1 \mu\text{m}$ and the thickness of the final device is 500 μm . The channels are etched in the first wafer, before the bonding step with the cap wafer followed by an annealing process. The cap wafer, then the channel (etched) wafer are thinned. The metallisation, consisting of a 200 nm titanium layer to provide the interface to the silicon, covered with 350 nm of nickel and a 500 nm protective gold layer, is applied. This provides the solder footprint for the connector and alignment marks which can be used during assembly. On the backside of the wafer, alignment marks are etched for plasma dicing and for additional reference marks during detector assembly. In the final step, singulation (i.e. the extraction of the individual coolers from the wafers) is achieved by plasma dicing, with the final wafer components being delivered on UV tape. Two coolers can be fitted onto one eight inch diameter wafer. The layout of the final wafer is

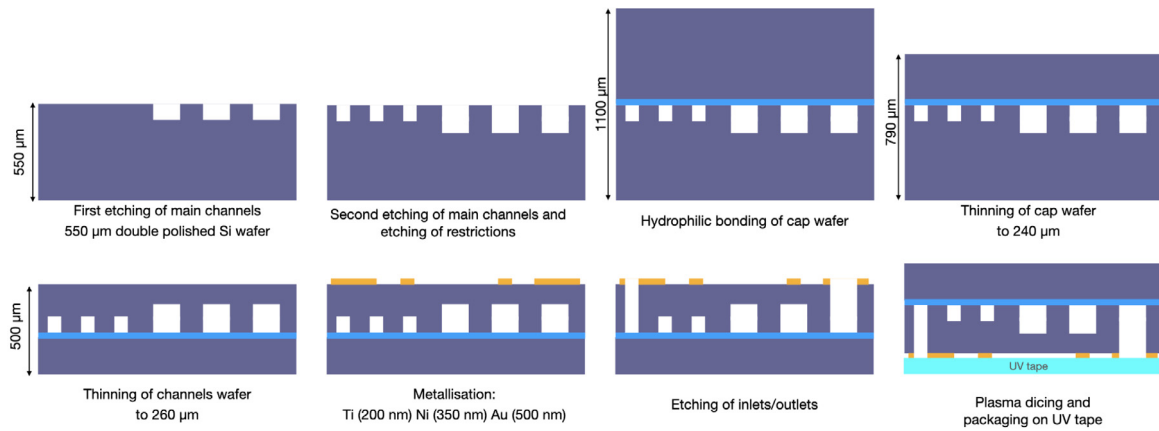


Fig. 11. Schematic of the principal steps of the production process.




Fig. 12. Left: Layout of the final wafer as manufactured. There are two coolers, four additional connector footprints for metallisation and adhesion tests, and 24 pressure control samples. Right: Photograph of the back of a processed and diced wafer.

illustrated in Fig. 12 along with a photo of the back face of a completed wafer. The additional wafer space was filled with dedicated samples to check the metallisation quality for the connector soldering pad and small channels which could be used for pressure testing and bonding quality.

3.2. Quality assurance

The quality assurance at the production stage was based on the following criteria:

- No defects in the high resolution ($< 50 \mu\text{m}$) scanning acoustic microscope (SAM) images taken after post-bonding annealing can be detected in the vicinity (within $500 \mu\text{m}$) of the channels, inlet or outlet areas.
- The edges of the devices do not present any abnormal defects caused by the dicing, such as cracks or chipping.
- The metallisation stack for soldering the fluidic connector has good surface quality and adherence to the silicon surface for subsequent soldering of the fluidic connector. The surface quality, adherence to the silicon surface and wettability can be checked with the additional samples on the wafer.
- The pressure tests of the test structures on the same wafer show consistent pressure resistances with the measurements during the R&D phase.

These requirements were challenging, in particular the first one, which is based on the concept that no channel width should exceed

$200 \mu\text{m}$, as discussed in Section 2 and supported by the measurements in Fig. 7. The identification of defects from the SAM images required painstaking scanning by eye of multiple high-resolution images. A grading procedure was developed, classing perfect coolers with no defects as A grade, coolers with defects in the input/output region as B grade and coolers with defects near channels as C grade. Small dicing imperfections were accepted if the distance perpendicular to the ideal dicing line to the tip of the imperfection is below $300 \mu\text{m}$ and their location did not affect the module assembly procedure afterwards. As the connector region is reinforced with a piece of silicon glued on the back of the cooler, small defects in the inlet/outlet region could possibly be accepted but out of an abundance of caution these were not used for installation in the experiment. The coolers are large, with just two fitting onto one wafer. For this reason the requirement that there should be no defect at all within the channel region of each cooler had a large impact on the yield during production. Not taking into account wafers discarded during production and dicing losses, approximately 69% of produced coolers were accepted for modules to be installed into LHCb. Due to the difficulty of scanning the images by eye and the dependence on the operator, image processing software was produced to make any defects easier to identify. The images were transformed to colour scale, and background noise was suppressed. This made the defects easier to spot, as shown in Fig. 13. All defects were reverified on the original SAM images before final classification.




Fig. 13. Left: Typical SAM image, with details showing examples of classifications into Grade A, B and C. Right: illustration of regions with a defect together with the image after colourisation.

4. Performance simulation

This section groups the various simulations which were done to support the microchannel development and module design choices. At the start of the project when the general principle and the choice of coolant was still open, CoBra simulations were run to gain an understanding of the thermal performance of the microchannel cooling. In a next step, FEA simulations were used to optimise the channel layout for the most efficient heat transfer. Finally, the full detector simulation was used to estimate the material contribution of the microchannel cooling solution.

4.1. CoBra simulations

The CoBra simulation programme was used to analyse the expected coolant flow in the microchannels. The results of this analysis were used to derive guidelines for the channel lengths and dimensions in the first prototypes. CoBra has been developed to predict the complex behaviour of particle detector cooling circuits. It is based on flow pattern models [23,24] which describe flow conditions and the related heat transfer and pressure drops, and include predictions of dry-out conditions. A useful figure of merit is the volumetric heat transfer conduction, $\frac{Q}{V \cdot dT}$, where Q is the absorbed heat, V is the volume of the tube, and dT is the difference in temperature between the tube and the liquid. The simulation was performed with model parameters of a tube length of 120 mm, absorbed heat of 1.5 W, a temperature at the inlet of -20 °C and a vapour quality of the outlet of 0.8. Coupling between channels was neglected and the tubes were assumed to be circular in cross-section. The result of the simulation is shown in Fig. 14. As the tubes get smaller and the relative contribution of the wall area increases, the heat transfer coefficient improves, up to the point where the coolant viscosity limits the flow. It is instructive to compare the performance of CO_2 with other coolant candidates. It can be seen that CO_2 is well matched to small channels, and has an excellent performance due to its high latent heat and low viscosity. The final dimensions of $120 \mu\text{m} \times 200 \mu\text{m}$ for the LHCb VELO main channel cross section, driven by manufacturing and pressure resistance considerations, corresponds to a tube hydraulic diameter of approximately $150 \mu\text{m}$.⁴

⁴ The hydraulic diameter of a channel with rectangular cross section is given by $D = 4A/P$ where A is the area and P the perimeter of the rectangle




Fig. 14. Result of CoBra simulation showing the heat transfer coefficient for various channel diameters and different coolants. The optimum channel diameter is at $200 \mu\text{m}$. For this simulation parameters were chosen to be close to those of the VELO; heat dissipation of 1.5 W, tube length of 12 cm and coolant temperature of -20 °C.

4.2. FEA simulations

A finite element analysis was used to quantify the heat flow through a microchannel cooler and to optimise the spacing between the microchannels. This study modelled a homogeneous $400 \mu\text{m}$ silicon cooler containing $70 \mu\text{m} \times 200 \mu\text{m}$ microchannels, connected with a $100 \mu\text{m}$ thick stycast glue layer to a thin heat source on one side which provides $3 \text{ W} / \text{ASIC}$ ($\sim 36 \text{ W}$ per module), comparable to the power dissipation expected in the final modules. The main purpose of this study was to optimise the channel dimensions and layout. It was found that a separation of $500 \mu\text{m}$ between channel edges is optimal, allowing efficient heat exchange along three sides of the channel, instead of being dominated by the side closest to the heat source which is the case

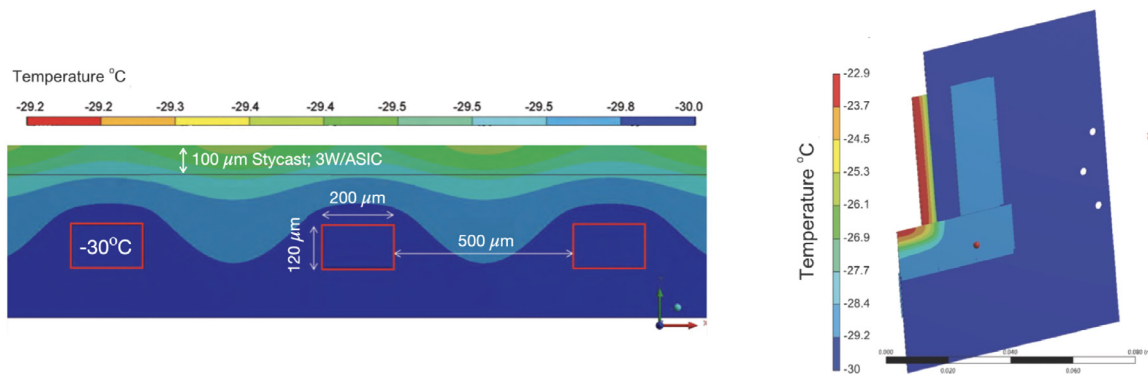


Fig. 15. (left) Thermal map from the microchannel FEA simulations, with a cross section illustrating the optimal channel separation and the very small temperature gradients to the heat source. (right) View of the full module at nominal power. The maximum ΔT to the silicon tip, from which the cooler is retracted, is estimated to be about 7 °C.

for smaller pitches. In addition the larger spacing provides maximal bonding surface. With this design the simulation shows a negligible temperature differential between the coolant and heated surface, as shown in Fig. 15 (left). The simulation was used to estimate the final ΔT that can be expected between the coolant and the tip of the silicon, with the retracted cooler. The full module under maximum power with a coolant temperature of -30 °C shows an expected ΔT to the silicon tip of about 7 °C, as illustrated in Fig. 15 (right). The final achieved performance comes very close to this ideal number, as illustrated in Fig. 27.

4.3. Full detector simulations

A description of the detector material was constructed using the basic volumes from the Geant4 toolkit [25]. This description was then used to estimate the material through which a particle travels while traversing the VELO. Fig. 16 (bottom left) illustrates the material, in terms of percentage of a radiation length, encountered when passing through a pair of VELO modules along trajectories parallel to the beam axis. The cooler, hybrids, sensors and ASICs can all be seen towards the centre of the figure, while the most prominent feature is the invar cooling connector, which is well removed from the interaction region. A particularly important parameter for the extrapolation of tracks to the secondary and primary vertices is the amount of material before the second measured point. Thanks to the thermal efficiency of the microchannel solution it is possible to minimise the contribution of the cooling to this parameter by retracting the cooler from the tip of the sensor by 5mm. The material scans shown in Fig. 16 (top) demonstrate that this reduces the material presented by the cooler to an average of 0.06% of a radiation length, or less than 2% of the total material traversed before the second measured point. Rather, the dominant contribution comes from the thin corrugated aluminium foil that separates the VELO secondary vacuum from the primary LHC machine vacuum, known as the R.F. foil. When considering the total material traversed by a particle while it travels through the VELO, the cooler contributes on average only 18% of that total, corresponding to 4.2% of a radiation length.

5. Microchannel assembly

5.1. The assembly challenge

The CO_2 is delivered to the microchannels via a connector with an integrated manifold and inlet and outlet tubes. The connector must

be attached to the silicon with a robust and leak tight seal, which is achieved via a soldering process. The cooling connector also serves as the mechanical interface with the module support, in order to avoid introducing stresses on the microchannels by constraining them with a mechanical attachment. Experience with the previously installed VELO detector has shown that where flux has been used in the cooling system construction there is a long term risk of corrosion, so the use of flux is not allowed. The connector attachment must withstand the safety limit of 186 bar, must not block the channels, and must have no large internal voids which could compromise its high pressure resistance.

The soldering process in which the cooler and its fluidic connector are attached proceeds in several steps. First, the cooler and the connector are pretinned in separate processes. The two components are then carefully aligned, and a reflow performed in order to form the connection.

5.2. Connector design

The connector head is made from Invar 36 due to its coefficient of thermal expansion being close to that of silicon while the tubes and glands are made from stainless steel 316L. The connector head dimensions are $28.1 \times 8.2 \times 3.0$ mm with inlet and outlet slits of dimensions $8.9 \times 0.8 \times 2.1$ mm. The tubes have an unbent length of 318.6 mm and external diameter of 1.59 mm with the inlet having an internal diameter of 0.57 mm while the outlet tube has an internal diameter of 0.87 mm. The dimensions of the connector are shown in Fig. 17. The connector head, tubes, and glands are vacuum brazed by CEA-CCFE. Rings of braze are placed into a custom-made jig designed to hold the connector head, tubes, and glands in the correct positions. The male nut is held in place around the tube while the tube and glands are brazed. During the brazing, the braze is above its melting point of 950 °C for 30 min.

The connectors are then tested to ensure they can withstand the safety limit of 186 bar, do not leak at the brazing points, and the inlet and outlet slits on the connector head do not exhibit large chamfers (sloping edges, as opposed to a right angle) or burrs (jagged pieces of metal that are created during the manufacturing process). First, the leak rate is measured with a helium leak test. The connector inlet and outlet slits are sealed by enclosing the connector head within two Invar blocks which are bolted together. Then the inlet tube is connected to a leak detector under vacuum and helium gas is sprayed around the brazing points. The leak rate inside the inlet tube is measured over 10 min with a precision of 0.1×10^{-10} mbar.l/s. This procedure is repeated for the outlet tube. A pressure test is then performed, where the connector inlet and outlet slits are sealed by enclosing the connector head within

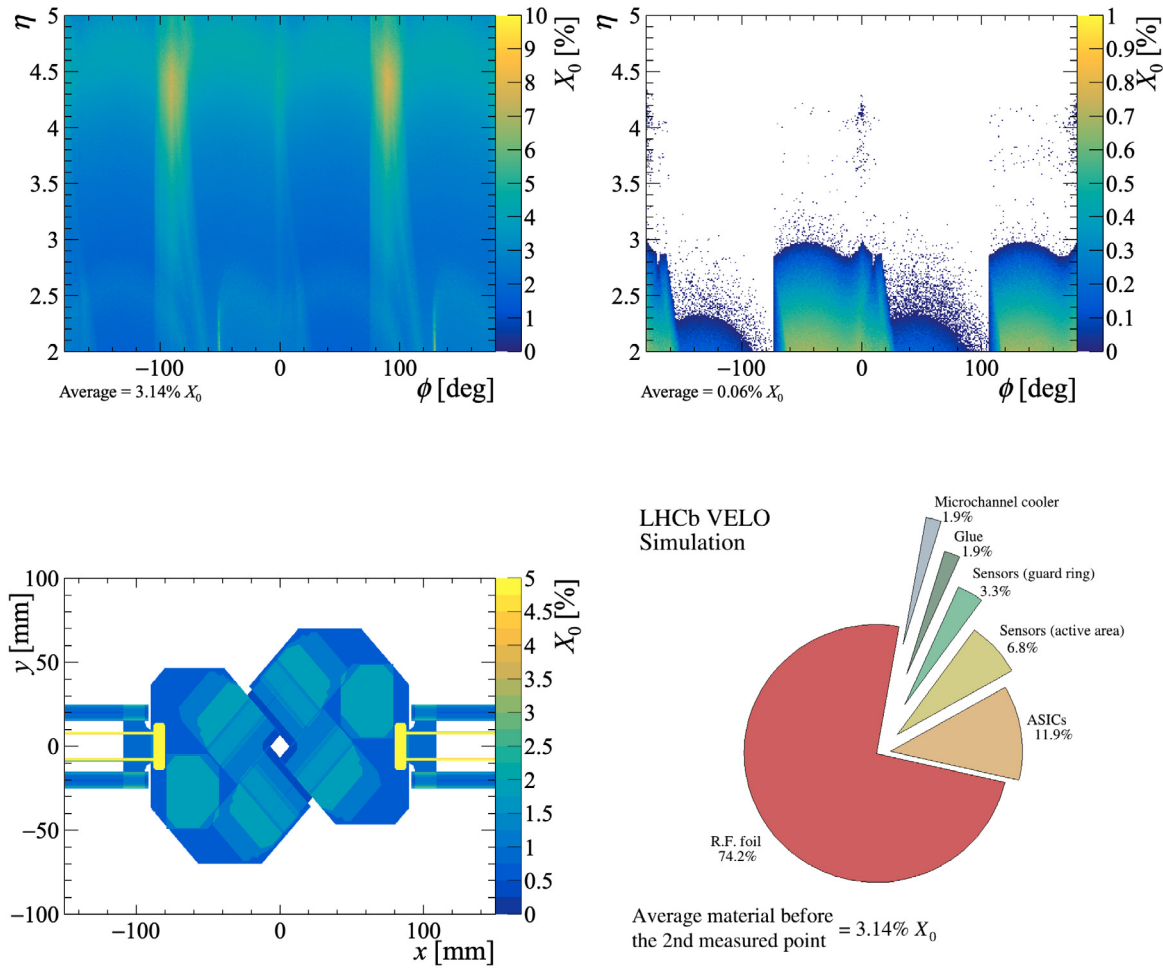


Fig. 16. (top left) The amount of material, expressed as a percentage of a radiation length (X_0), traversed by a particle before the second measured point, as a function of the pseudorapidity (η) and azimuthal angle (ϕ) of its momentum vector. (top right) Similarly, but considering only the contribution of the microchannel cooler to the material encountered before the second measured point, again as a function of η and ϕ . (bottom left) The amount of material, expressed as a percentage of a radiation length, presented by a pair of VELO modules when considering trajectories parallel to the beam axis. (bottom right) The relative contribution of each part of the detector to the material before the second measured point.

two Invar blocks which are bolted together. Then the inlet tube is connected to a water pump, where the water pressure is brought to 200 bar for two minutes. The connector and tube are observed for signs of leaking or failure due to the pressure. This test is then repeated for the outlet tube and then the connector is flushed with pressurised air to remove residual water from the tubes and connector head. Lastly, a visual inspection with a high magnification microscope is performed to search for defects on the connector surface, the inlet and outlet, and the tubes. In particular, the connector is inspected to ensure chamfers at the surface edges, including the inlet and outlet, do not exceed a length of 100 μm . Additionally, any burrs found at the inlet and outlet are required to be < 50 μm before passing inspection.

5.3. Connector and cooler preparation

Before being pre-tinned, the connector must be metallised. Experience showed that in order to ensure quality metallisation layers, a highly smooth and clean connector surface is necessary. Precise procedures were developed in order to obtain consistent results. First, the surface of the connector is polished in order to remove any ridges or depressions which might lead to inconsistencies or imperfections in the metallisation layer. The polishing is performed with increasingly fine sand paper. Visual inspection under a microscope is performed throughout the process. The connector must then be thoroughly cleaned. A




Fig. 17. Illustration of a connector head with the tubes attached and relevant measurements indicated.

sequence of ultrasound baths and rinsing with acetone, ethanol, and demineralised water is performed in a clean room environment. The connectors are then kept in a box filled with nitrogen until ready for metallisation.

The metallisation process consists of depositing a nickel layer of 5 μm followed by a gold layer of 200 nm. The nickel reacts and




Fig. 18. Illustration of assembly steps, with cross section of connector/cooler attachment.




Fig. 19. Images from the translational alignment of the connector with the cooler. Image processing tools are used to determine the position of the attachment region, and the connector is then aligned using these measurements.

mixes with the solder to create the intermetallic compounds during the reflow, while the gold is used to prevent the oxidation of the samples and dissolves in the solder. Following this process, the metal layers are removed from inside the inlet and outlet regions and the connector is once again cleaned in an ultrasound bath and rinsed with ethanol. As described in Section 3, the coolers have been previously metallised by LETI with a matching metallisation layer around the microchannel inlet/outlet region. Before pre-tinning, both the connector and the cooler undergo plasma cleaning in order to remove any possible organic contaminants. At this point, the components are ready to be pre-tinned. The metal stack is illustrated schematically in Fig. 18.

A 100 μm thick $\text{Sn}_{36}\text{Pb}_{62}\text{Ag}_2$ foil (a tin, lead, and silver alloy used as solder) is placed on the metallisation layer, and the component left in a vacuum tank in order to outgas. After a period of at least two hours (twelve hours) for the cooler (connector), the tank is filled and flushed with nitrogen gas. Formic acid vapour is introduced in the nitrogen environment at a concentration of approximately 3%. The sample is then heated to 150 $^{\circ}\text{C}$ and is left at this temperature for five minutes, in order to let the formic acid react with oxides on the foil and metallisation layer to create a salt or formate. These salts will decompose into CO_2 , water, and hydrogen when the sample is further heated above approximately 200 $^{\circ}\text{C}$ [26–29]. The heating of the sample is initially performed slowly in order to maintain approximately equal temperature over the surface and a uniform melting of the foil. At around 183 $^{\circ}\text{C}$, the foil is fully melted and the sample is heated rapidly to a temperature of 235 $^{\circ}\text{C}$. The tank is put under vacuum in order to allow trapped gas bubbles to escape and then cooled down to room temperature.

The pre-tinned coolers and connectors are visually examined to ensure no solder or other debris entered the slit regions. The components are mounted in the vacuum tank for the soldering, with the connector supported above the cooler by a stage which allows for micro adjustments of the rotation and translation with respect to the cooler. The connector is then carefully aligned over the matching pre-tinned area of the cooler using digital microscopes and cameras. Firstly, the rotational degrees of freedom are adjusted until the surface of the connector is completely parallel to the cooler. This is done with two cameras, one positioned in front and the other to the side. With the aid

of image processing tools, a precision of 10 – 20 μm in the alignment of the edges is typically achieved. An overhead camera with image processing software is then used to perform the necessary translational adjustments. The process is demonstrated in Fig. 19.

Once aligned, the connector is lowered to make contact with the solder layer of the cooler, which lies on an 8 inch silicon wafer used as a lightweight planar mechanical support capable of rapid heating in vacuum. This wafer is itself supported by ball bearings to allow for self-correction of any remaining small misalignment. The oven is heated in vacuum at a rate of 1 $^{\circ}\text{C}/\text{s}$ to a final temperature of 235 $^{\circ}\text{C}$, and the solder layers melt to form the connection. Before cooling down, the tank is pressurised with atmospheric air in order to compress any voids in the solder. The sample is then cooled down to 100 $^{\circ}\text{C}$, at which point vacuum is reintroduced, and the same procedure repeated in order to reduce the chance of any lingering voids.

The final step of the assembly consists of gluing a rectangular piece of silicon onto the cooler on the side opposite the connector. This additional thickness serves as reinforcement against the pressure of the CO_2 flowing from the connector into the cooler. A photograph of a finished microchannel assembly is shown in Fig. 20.

6. Microchannel quality assurance

The qualification procedure must ensure that the microchannel assemblies can be safely installed inside the secondary vacuum of the detector, and that there is no obstruction which could cause a reduction of CO_2 flow. Therefore modules are extensively checked for leak tightness and overall robustness. Long-term effects in the solder joint such as creep and fatigue are tested with a subset of prototypes.

6.1. Microchannel assembly quality assurance

6.1.1. X-ray qualification

The quality assurance begins with an X-ray of the solder joint. This test allows for verifying that solder did not enter the inlet or outlet region, and that there are no large voids in the joint. The absence of such features rules out the possibility of the presence of obstructions which might affect the CO_2 flow and potential weak points which might




Fig. 20. Microchannel assembly, consisting of a microchannel cooler soldered to a fluidic connector, ready to be equipped with VELO module components.

compromise the integrity of the joint. Examples of the X-rays can be seen in Fig. 21, both for the case of a failed and successful soldering. If solder access into the slit region is observed, a 3D X-ray tomography is performed in order to determine if the channels are blocked. The tomography obtained typical resolutions of $\sim 15 \mu\text{m}$.

6.1.2. Check for leak tightness

Before subjecting the sample to the rigorous robustness tests, it is first checked for leak tightness. A leak tester is attached to the fluidic connector and creates a vacuum inside the sample. Helium is sprayed all around the cooler while a gas spectrometer is used to detect any helium entering the vacuum. The procedure excludes leaks down to approximately 10^{-9} mbar l/s.

6.1.3. Check for pressure resistance

To check for robustness, the samples are subjected to high pressure tests. The system is initially pressurised to 130 bar, which corresponds to the burst disk release pressure installed in the final system. High purity nitrogen is used for the pressurisation in order to ensure internal cleanliness. After 45 min, the pressure is increased to a final value of 186 bar. This last step is done to include an additional safety buffer in the qualification, and lasts for 15 min.

6.1.4. High-pressure leak check

A final high-pressure leak test is performed. The microchannel assembly is installed inside a vacuum chamber and is pressurised to 60 bar with helium while a gas spectrometer is used to detect any leakage of helium into the vacuum chamber. The typical background level is around 10^{-7} mbar l/s and no significant leak rate increase is observed when the microchannel assembly is pressurised.

6.1.5. Planarity measurements

For the subsequent module gluing and assembly steps, the silicon surfaces must also meet flatness requirements. The planarity is measured using light reflectometry devices capable of reaching an accuracy of $\pm 3 \mu\text{m}$. As the silicon has a reflective surface, a fine layer of talc is applied in order to be able to resolve the surface. The tests are performed before and after the connector attachment, in order to confirm stress from the soldering procedure did not introduce deformations. An example of these measurements can be seen in Fig. 22. The maximum deviations from planarity never exceeded $100 \mu\text{m}$.

6.1.6. Mechanical characteristics and cleanliness

The silicon surface must also be clean and free of any surface residue. The cooler is rinsed with ethanol and a final visual inspection is performed. The shape of the connector capillaries must also conform to predetermined dimensions, and is checked using a custom 3D jig.

6.2. Creep and fatigue

Over time the solder joint may be subject to failure due to sustained strain over a long period leading to plastic deformation. Creep effects can typically be neglected if the temperature is less than T_{critical} , defined as 50% of the absolute melting temperature. For the SnPb solder used this corresponds to $T_{\text{critical}} = -46 \text{ }^\circ\text{C}$. The VELO will operate for nearly a decade, and will be maintained at $-30 \text{ }^\circ\text{C}$ for the majority of its lifetime. Therefore, creep effects are expected to be small. Using the model in Ref. [30] the strain rate was simulated for the solder joint, taking into account two scenarios, one with the stress from the CO_2 pressure spread uniformly over the joint, and one with the stress concentrated around the slit. It was found that at $-30 \text{ }^\circ\text{C}$ the time to fail for the worst case was approximately 10^6 days, dropping to approximately 2×10^4 days at room temperature. During the test, the sample pressure was kept at 60 bar to mimic the CO_2 pressure at $20 \text{ }^\circ\text{C}$ while the temperature was kept between $90\text{--}120 \text{ }^\circ\text{C}$ to accelerate the creep process. In total, 15 samples were tested during a testing campaign of approximately three months (2477 h). Two failures were observed with samples which were prepared with flux and none of the samples prepared with the fluxless procedure failed.

The potential effects of fatigue on the joint were also considered, corresponding to a weakening of the joint due to repeatedly applied loads. During the VELO lifetime a few hundred temperature cycles are expected, corresponding to powering during the LHC fills. In order to check the effects of fatigue a dedicated test stand was set up. Microchannel samples were connected with a high pressure connection and mounted on a heating/cooling plate equipped with six 40 W Peltier coolers. The temperature was cycled between $-40 \text{ }^\circ\text{C}$ and $+60 \text{ }^\circ\text{C}$ and at each point in the cycle it was possible to cycle the pressure between atmospheric pressure and 186 bar. The fatigue effect was evaluated with nine samples. Overall, approximately, 6800 temperature cycles and 40000 pressure cycles were performed over several months. The pressure cycles were done only when the sample was at high temperature to avoid accumulating ice inside the channels while the sample is below $0 \text{ }^\circ\text{C}$. One failure was observed on a sample with flux. The post-mortem microscopic photos showed that there were regions which were not properly attached due to a thin layer of flux between the connector and silicon interface. The pressure and temperature recorded during one fatigue test is illustrated in Fig. 23. It can be seen that due to the connection to the Peltier plate the extreme pressure range could be explored while maintaining the temperature steady, hence vastly extending the parameter space beyond what can be expected in normal operation and providing a robust probe of the resistance of the device.

6.3. COVID procedures

The microchannel cooler is the backbone of the detector modules, and a constant supply of the microchannel assemblies was critical for the timeliness of the detector construction. The COVID-19 pandemic began in the middle of the construction period, and presented several organisational challenges in assuring a reliable supply of coolers.

One major aspect of this was the interruption of travel, which required adaptation in transportation of materials to avoid supply chain disruptions. Travel between the different sites was originally frequent enough that components could simply be carefully packed in personal luggage. With the advent of the pandemic, it became necessary to rely on more conventional shipping methods. Particularly in the case of the assembled microchannel coolers, this required development of secure shipping methods that will be described in detail in the next section.



Fig. 21. X-ray of solder joints attaching the fluidic connector to the microchannel cooler. The left image demonstrates an unsuccessful soldering in which solder has crept into the slit regions. Voids connect to the slits are also visible. The right image demonstrates a successful soldering, with no voids or solder excess.




Fig. 22. Microchannel surface survey after the attachment of the fluidic connector. The blue(red) regions indicate the lowest(highest) regions. The overall variation in this case is around 25 μm over the surface of the cooler.




Fig. 23. The two plots on the left show recorded data during one temperature cycle, lasting around 30 min. At the warmest temperature (around 60 °C) a pressure cycle is performed. The two plots on the right show data from multiple cycles. The min–max difference is shown for the temperature and pressure respectively from left to right. It can be seen that the pressure provided to the sample decreases over time since the reservoir becomes depleted with every pressure cycle.

Compliance with social distancing guidelines required careful planning in the lab spaces. Restrictions elsewhere meant limited availability of several services such as connector metallisation, X-ray tomography, and clean room access. These complications required increased parallelisation of procedures. Originally, many parts of the process were done one-by-one, such that if a problem arose, for example with the metallisation, not many components would be affected. With the above complications, such a system was no longer possible and so more components needed to be processed simultaneously. In order to not risk losing batches of the limited supplies to unidentified problems, stricter quality control was enforced to ensure that problems were identified at an early stage.

Despite these obstacles, the above adjustments allowed for a constant and reliable production of the coolers. The production had a high yield, and did not result in any significant delays with respect to other elements of the detector.

6.4. Packaging and transport

After the qualification and final visual inspection the coolers were packed into small transport boxes as illustrated in Fig. 24. The tubes were held in place by a 3D printed jig and the fluidic connector was gently pushed down against a base plexiglass plate which is covered by

a soft clean room grade paper. Each small transport box could carry up to two microchannel assemblies and were typically transported as hand luggage either by plane or car. Later on, those boxes were integrated into a larger transport box fixed on a pallet to allow the usage of standard air freight. The large transport box was quite oversized for such small structures and had at least 9 cm of polyethylene foam around each of two small transport boxes. No microchannel assemblies were damaged during the transportation to the module assembly sites.

6.5. Assembly yield

The fluxless soldering procedure was refined and tested extensively using lower grade coolers in order to optimise the procedure. As the modules will be mounted in the secondary vacuum of the LHC, only those coolers which perfectly passed all QA procedures were accepted for installation. Firstly, the procedure was perfected with the manufacture of approximately 30 prototype coolers, which had the additional benefit of providing a plentiful supply of coolers for to exercise the module assembly steps. Next, the production phase was started. Over a period of 27 months, including the COVID lockdown period, 81 installation quality microchannel assemblies were produced with a yield of 87%. The assemblies which did not pass the full QA procedure most often had solder accumulate in the slit region or outside




Fig. 24. Microchannel assembly mounted in the transport jig.

the solder joint. Analysis shows that eight could be recovered, and these are held in reserve.

7. Cooling performance

7.1. Characterisation and operational stability testing

Two microchannel assemblies were used to measure the full fluidic characteristics, thermal performance and stability under joint running with a CO₂ cooling plant called CORA (CO₂ Research Apparatus) [31]. They were each equipped with a full set of 12 silicon heaters to simulate the ASIC load in the final VELO module and an additional silicon heater to simulate the sensor load in the region where the sensor overhangs the cooler and suffers the highest radiation. The temperature was monitored with nine temperature probes on each cooler, tracking the temperature evolution from the inlet through to the outlet manifold. In order to be as close as possible to the final configuration, the CO₂ was supplied through a capillary with a length similar to that in the VELO tank, equipped with a 150 μm gasket to control the flow as is installed for the final assembled VELO. Results from the cooling test are shown in Fig. 25. On the left, the flow as a function of the pressure drop over the cooler is shown. The accumulator CO₂ temperature set point was -35 °C and it was checked that the boiling was correctly triggered in the microchannels. It can be seen that there is a very small spread between the curves with the module set to a range of power dissipations between zero and the maximum. This gives a good indication that the system will avoid instabilities in the case of parallel operation of many modules which may have different power dissipation. At the nominal flow of 0.4 g/s the pressure drop is approximately 9 bar. The plot on the right gives an indication of the temperature evolution along the length of the cooling channels. The highest temperature is seen at the inlet. After passing the point where the boiling is triggered, the temperature then drops along the length of the channel. This is due to a small drop in pressure, due to the additional resistance created by the boiling liquid.

A stability test was also performed to check the interplay that might be expected between two VELO modules in the case of changes in fluidic resistance due to changes in power. Therefore, the initial tests were performed with the two equipped microchannel assemblies in parallel. While keeping the power constant in one assembly, the second one was cycled from unpowered to the nominal expected power several times as can be observed in Fig. 26. During this specific run, the power on the microchannel assembly B was kept at 30 W while the power on microchannel assembly A was switched on and off (from 0 W to 30 W) multiple times. Overall, a small variation of temperature was observed on the microchannel assembly B (order of 1 °C or less) and, more importantly, the boiling was maintained throughout the power cycles.

Finally a measurement was performed to check that the temperature at the very tip of the module. Here the silicon overhangs the microchannel and so is not directly cooled. This temperature difference is expected to increase as the irradiated tip draws more current at the end of lifetime. The measurement was performed as a function of coolant flow for four different power dissipations for an outlet temperature of -32 °C. As seen in Fig. 25 higher flows are associated with greater pressure drops due to the increased boiling in the channel, resulting in a temperature drop during the 2-phase cooling step. The results, displayed in Fig. 27, show that the temperature in the sensor can be kept comfortably below the target of -20 °C for the nominal power dissipation of 1 W and coolant flow of 0.4 g/s.

7.2. Construction of VELO modules using microchannel assemblies

The microchannel assemblies finally undergo a series of construction steps in order to attach the active components and services to complete the VELO modules. This involves a sequence of complex handling procedures.

In a first step the microchannel assemblies are glued, via the connector, onto a carbon fibre mid-plate, which allows the position of the cooler to be accurately defined with respect to the module base, to a level of 50 μm in the glue thickness direction (beam direction) and a nominal 20 μm in the transverse direction (module height). For this step the edges of the cooler, butted against alignment pillars on the jig, are used. Next, the hybrid pixel tiles and front end hybrids are glued to the cooler, using Stycast FT2850 with catalyst 23LV for optimum heat transfer for the tiles, and Loctite SI 5145 for the hybrids where more flexibility is required due to the CTE difference. This is done in a double sided process, with the module supported in a turnplate, attached by the foot and with additional support for the cooler, sandwiched between independent vacuum jigs supporting the tiles. The alignment marks on both sides of the cooler are useful to relate the position of the cooler and the elements glued on it with respect to the jig. The electrical connections between the hybrids and tiles are formed with wirebonds, with support provided by custom-made mechanical supports below the coolers.

Extensive metrology and visual qualification is performed after every assembly step, the final module is then electrically and thermally evaluated in vacuum. The full suite of tests is repeated again after a thermal cycling consisting of ten cycles of cooling and heating from room temperature down to -30 °C and back. More details about the module assembly steps can be found in Ref. [32].

Throughout the assembly of the first 52 modules the microchannel assemblies have been shown to be robust against these handling procedures. This is best shown on bandgap⁵ measurement from the ASICs (see Fig. 28), where, even after thermal cycling, all tiles at the end of life of the module are expected to observe a temperature difference less than 8 degrees from the coolant. The difference in ΔT between the tiles is dominated by the thickness of the glue layer under the ASICs, as obtained from the module metrology. It can be concluded that the thermal performance is excellent and is homogeneous across the module cohort.

These data can also be used to measure the Thermal Figure of Merit (TFM) for the system, by monitoring the change in temperature on each ASIC, as measured by the internal band gap, as the module power is increased from zero to 26 W. At maximum power, each tile consumes about 5.7 W, hence has a power density of 0.88 W/cm² (1.48 W/cm² is taken for the overhanging tiles due to their reduced contact area with the cooler). The data from 680 ASICs was used to calculate the TFM in bins of glue thickness. The glue thickness is derived for each tile from an average of more than 100 metrology measurements with a standard deviation of less than 5 μm for the ASICs used for this analysis. The result is shown in Fig. 28. It can be seen that the TFM

⁵ internally implemented temperature measurement in the integrated circuit

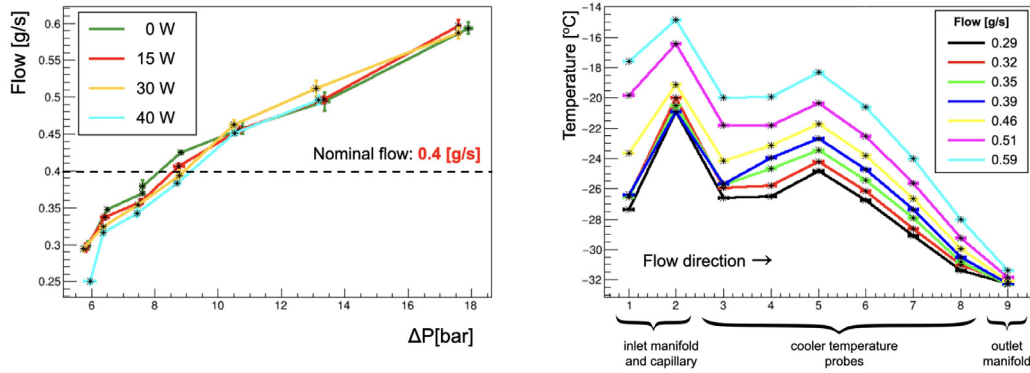


Fig. 25. Left plot: Measured flow as a function of pressure drop over the cooler. The curves are similar for a range of power dissipation, indicating good operational stability. Right plot: Temperature evolution from the inlet to the outlet for a range of power dissipations where the probes from left to right indicate increasing distance along the channels. Overall there is a drop in temperature as the coolant progresses along the channels due to the pressure drop. There are two peaks due to local changes of condition; the peak at probe number 2 is due to preheater activity on liquid CO₂ inside the inlet capillary and the peak at temperature probe number 5 is due to the overhanging position of the heater at this position.

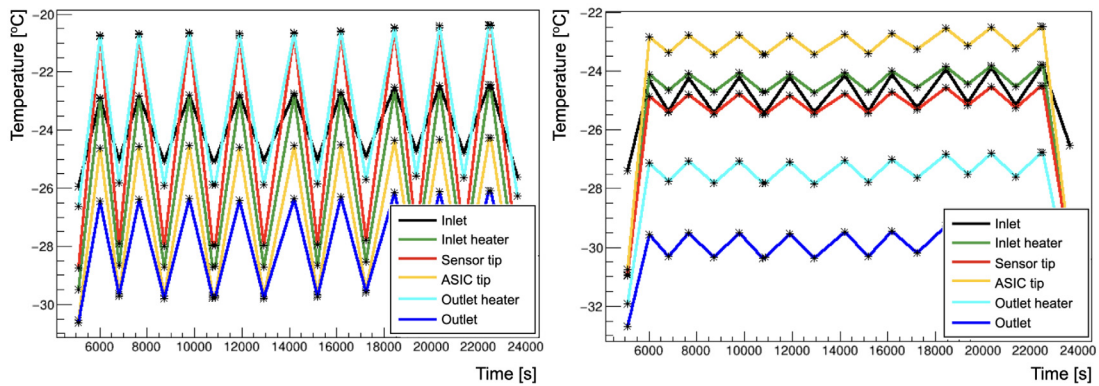


Fig. 26. While keeping the constant nominal power dissipation of 30 W on microchannel B (right plot), the power on the microchannel assembly A (left plot) was switched on and off multiple times to test the stability of the boiling. Both plots show the temperature on the coolers in different locations as a function of time.

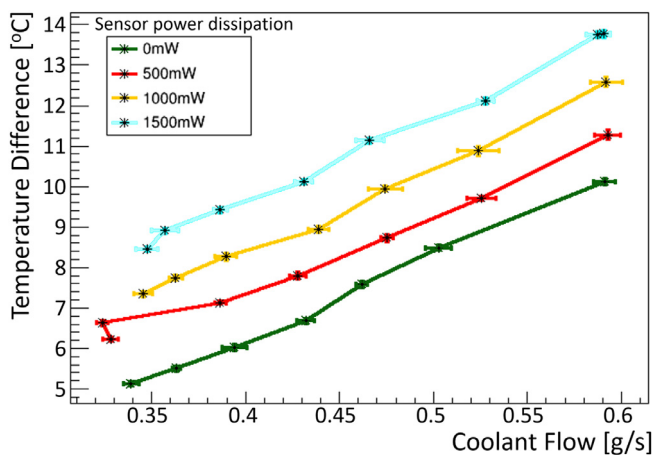


Fig. 27. Temperature difference between the outlet manifold and the tip of the sensor – which extends beyond the microchannel cooler – as a function of coolant flow. Four curves are shown. The nominal flow is 0.4 g/s and the maximum sensor power dissipation at the end of lifetime of operation at the LHC is 1 W [1].

has a dependence on glue thickness and for the thinnest glue layers TFM values of less than 2 K cm² W⁻¹ are achieved, rising to just above 3 K cm² W⁻¹ for the thickest layers. This result, measured on the final modules as constructed, showcases the high efficiency of the microchannel cooling design and confirms the expectations from the initial tests described in Section 2.1. For comparison the best measured

TFM to date on a detector system installed at the LHC is 13 K cm² W⁻¹ for the integrated steel pipe cooling system of the ATLAS IBL [33], while standalone microchannel and integrated test chips have been shown to achieve TFM values of 3 – 4 K cm² W⁻¹ [34,35]. The LHCb tile attachment uses on a star shaped glue pattern to eliminate the possibility of encapsulated voids in the vacuum system, and is not designed for thicknesses much below 50 μm. The glue thickness is an effective glue thickness as measured which may contain some empty areas. A uniform glue coverage may improve the TFM still further.

8. Summary and outlook

The microchannel assembly production for the VELO Upgrade I pixel detector has been successfully completed. The necessary R&D to bring this emergent technology from concept to production stage has been described. The microchannel coolers circulate bi-phase CO₂, allowing the VELO detector to operate at very high beam intensity, in vacuum, with ultra light weight and radiation hard cooling. The microchannel coolers also provide mechanical support for the sensors, front-end electronics, and hybrids. The cooling is efficient enough to allow the sensors to overhang at the module tip, hence there is no material budget contribution from cooling for the first few mm of track measurements. The microchannel R&D was started in 2013, the production was launched in 2016, the first installation grade soldered assembly was produced in 2017, and the production completed in 2021. Crucial to the R&D process was the availability of a facility which was able to produce small glass-pyrex demonstrators on short timescales for the prototyping phase, as well as a manufacturer for

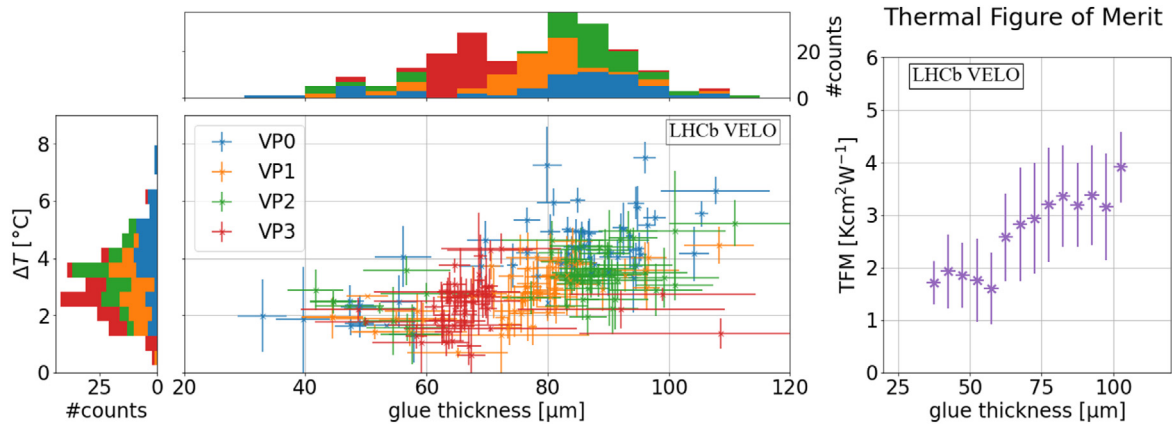


Fig. 28. Left plot: change in temperature of the four tiles on the modules (VP0 to VP3), when going from zero to full power (26 W). The values are derived from the ASIC bandgap measurements after thermal cycling and each tile is assigned the average value of the three ASICs. The x-axis gives the metrology measurement of the glue thickness for each tile. Right plot: Thermal figure of Merit (TFM) for the VELO modules, derived from the temperature and power measurements of each ASIC and plotted as a function of glue thickness. The values are binned with 5 μm thickness steps and a constraint on ASIC thickness variation of $< 5 \mu\text{m}$ is applied. All errors represent standard deviation of the measurements.

the production which was willing to be taken on the challenges and risks of this new technology. As of this date all microchannel assemblies used to produce VELO modules have shown excellent performance, exhibiting no problems due to handling, glueing, wirebonding and thermal cycling.

Challenges to be considered when deciding whether to implement microchannel technology include considerations of the cost and fragility of the devices, the size of the final coolers and the effort and skills needed to complete production. There is extensive ongoing R&D to address the outstanding issues for microchannel production and to make the process more cheap and reliable. The stringent QA demands on the quality of the bonding can have an impact on the production yield, particularly if very large coolers are produced. Reducing the size of the cooler should increase the yield, however in this case the connector should be made lower mass and the connection procedure simplified. This could be addressed by new developments of low mass plug and play hydraulic connections [36] which automatically engage between adjacent modules. R&D is ongoing and such “LEGO” connectors have been tested up to 40 bar [19]. The production of the patterns of channels could be made cheaper, faster, and more flexible with respect to design changes by the use of different techniques. Maskless laser writing could replace the time consuming and expensive lithographic steps and provide a more prototype-friendly process flow [37]. Alternatively the entire microchannel cooler could be produced with additive manufacturing [19], based on plastics (acrylic and epoxy-based photosensitive resins) or ceramics (Zirconia, Alumina, Silicon carbide, and Aluminium nitride). Such a technique was successfully pursued as a back up option for the LHCb VELO, as described in section Section 1.1. By using alternative silicon wafer bonding techniques it may be possible to integrate the cooling channels directly into the active silicon and avoid the thermal barrier of the glue layer. For this, low temperature bonding is being investigated, to enable the etched wafer to be directly bonded to the active silicon, or alternatively the technique of buried microchannels may be used, whereby channels are anisotropically etched into the back of a CMOS sensor in a post processing step [38,39]. Finally, for those applications in extremely intense environments where there may be high power consumption in tandem with a need to keep the irradiated sensors at very low temperatures, a move to an alternative coolant such as krypton may be necessary to access lower temperatures for the coolant. For a more complete discussion of future perspectives in microchannel cooling see [40].

The LHCb VELO microchannel cooled modules are currently being installed and are expected to run for the next decade. For the forthcoming Upgrade II [41,42] new cooling solutions will be needed, among

which microchannels incorporating the latest developments will be a strong contender.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We thank the technical and administrative staff at the LHCb institutes. We acknowledge support from CERN and from the national agencies: NWO (Netherlands); MNiSW and NCN (Poland); MSHE (Russia); MICINN (Spain); SNSF and SER (Switzerland); STFC (United Kingdom).

References

- [1] LHCb Collaboration, LHCb VELO Upgrade Technical Design Report, Tech. Rep., CERN, 2013, URL <https://cds.cern.ch/record/1624070>.
- [2] A.o. Alves, LHCb Collaboration Collaboration, The LHCb detector at the LHC, JINST 3 (2008) S08005, <http://dx.doi.org/10.1088/1748-0221/3/08/S08005>, Also published by CERN Geneva in 2010. URL <https://cds.cern.ch/record/1129809>.
- [3] P.R.o. Barbosa-Marinho, LHCb Collaboration Collaboration, LHCb VELO (VVertex LOcator): Technical Design Report, Technical Design Report. LHCb, CERN, Geneva, 2001, URL <http://cds.cern.ch/record/504321>.
- [4] T. Poikela, et al., The VeloPix ASIC, JINST 12 (01) (2017) C01070, <http://dx.doi.org/10.1088/1748-0221/12/01/C01070>.
- [5] K. Hennessy, Commissioning the readout for the new LHCb VELO, in: TWEPP 2019 Topical Workshop on Electronics for Particle Physics, 2019, p. 109, URL <https://cds.cern.ch/record/2681428>.
- [6] B. Verlaet, Controlling a 2-phase CO_2 loop using a 2-phase accumulator, in: ICR 2007. Refrigeration Creates the Future. Proceedings of the 22nd IIR International Congress of Refrigeration, 2007, pp. ICR07-B1-1565, <http://dx.doi.org/10.1088/1748-0221/12/01/C01070>, URL <https://iifir.org/en/fridoc/controlling-a-2-phase-co2-loop-using-a-2-phase-accumulator-24408>.
- [7] P. Petagna, B. Verlaet, A. Francescon, Two-phase thermal management of silicon detectors for high energy physics, in: Encyclopedia of Two-Phase Heat Transfer and Flow III, 2018, pp. 335–412, http://dx.doi.org/10.1142/9789813229471_0005, URL https://www.worldscientific.com/doi/abs/10.1142/9789813229471_0005.
- [8] J. Daguin, LHCb MAUVE system design, integration and construction, 2017, URL <https://www.docslides.com/embed/780601>.
- [9] B. Verlaet, E. Roeland, M. van Beuzekom, A. van Lysebetten, Design, manufacture and test results of the VTCS evaporator for the LHCb experiment at CERN, in: HEFAT2010; 7th International Conference on Heat Transfer, Fluid Mechanics and Thermodynamics 19-21 July 2010 Antalya, Turkey, 2010.
- [10] A. Van Lysebetten, B. Verlaet, M. van Beuzekom, CO_2 Cooling experience (LHCb), PoS Vertex 2007 (2008) 009, <http://dx.doi.org/10.22323/1.057.0009>.

- [11] V. Papavassiliou, BTeV Collaboration, BTeV: A proposal for a new b physics experiment at the Fermilab TEVATRON collider, in: M. Greco (Ed.), Frascati Phys. Ser. 17 (2000) 843–864.
- [12] K.E.G. P. K. Schelling, Managing heat for electronics, Mater. Today 8 (6) (2005) 30–35, [http://dx.doi.org/10.1016/S1369-7021\(05\)70935-4](http://dx.doi.org/10.1016/S1369-7021(05)70935-4), URL <https://www.sciencedirect.com/science/article/pii/S1369702105709354>.
- [13] T. Guillemet, P.-M. Geffroy, J.-M. Heintz, N. Chandra, Y. Lu, J.-F. Silvain, An innovative process to fabricate copper/diamond composite films for thermal management applications, Composites A 43 (10) (2012) 1746–1753, <http://dx.doi.org/10.1016/j.compositesa.2012.04.015>, CompTest 2011. URL <https://www.sciencedirect.com/science/article/pii/S1359835X12001443>.
- [14] J. Buytaert, LHCb VELO microchannels experience, in: Belle II VXD Open Workshop, 2019, URL <https://indico.cern.ch/event/810687>.
- [15] F. Sanders, Making modules for the LHCb VELO upgrade, 2019, URL https://lhcb.web.cern.ch/lhcb_page/lhcb_documents/presentations/Seminars/Default.htm.
- [16] D. Tuckerman, R. Pease, High-performance heat sinking for VLSI, IEEE Electron Device Lett. 2 (5) (1981) 126–129, <http://dx.doi.org/10.1109/EDL.1981.25367>.
- [17] C.J. Glassbrenner, G.A. Slack, Thermal conductivity of silicon and germanium from 3° K to the melting point, Phys. Rev. 134 (1964) A1058–A1069, <http://dx.doi.org/10.1103/PhysRev.134.A1058>, URL <https://link.aps.org/doi/10.1103/PhysRev.134.A1058>.
- [18] G. Aglieri Rinella, et al., The NA62 GigaTrackR: A low mass high intensity beam 4D tracker with 65 ps time resolution on tracks, JINST 14 (2019) P07010, <http://dx.doi.org/10.1088/1748-0221/14/07/P07010>, arXiv:1904.12837.
- [19] G.o. Aglieri, Strategic R&D Programme on Technologies for Future Experiments - Annual Report 2020, Tech. Rep., CERN, Geneva, 2020, URL <https://cds.cern.ch/record/2764386>.
- [20] A. Nomerotski, et al., Evaporative CO₂ cooling using microchannels etched in silicon for the future LHCb vertex detector, JINST 8 (2013) P04004, <http://dx.doi.org/10.1088/1748-0221/8/04/P04004>, arXiv:1211.1176.
- [21] D. Li, X. Cui, M. Du, Y. Zhou, F. Lan, Effect of combined hydrophilic activation on interface characteristics of Si/Si wafer direct bonding, Processes 9 (9) (2021) <http://dx.doi.org/10.3390/pr9091599>, URL <https://www.mdpi.com/2227-9717/9/9/1599>.
- [22] F. Fournel, H. Moriceau, R. Beneyton, Low temperature void free hydrophilic or hydrophobic silicon direct bonding, ECS Transactions 3 (6) (2006) 139–146, <http://dx.doi.org/10.1149/1.2357063>, <https://doi.org/10.1149/1.2357063>.
- [23] L. Cheng, G. Ribatski, J. Moreno Quiben, J. Thome, New prediction methods for CO₂ evaporation inside tubes: Part I - A two-phase flow pattern map and a flow pattern based phenomenological model for two-phase flow frictional pressure drops, Int. J. Heat Mass Transfer 51 (1–2) (2008) 111–124, <http://dx.doi.org/10.1016/j.ijheatmasstransfer.2007.04.002>, URL <http://infoscience.epfl.ch/record/115207>.
- [24] L. Cheng, G. Ribatski, J.R. Thome, New prediction methods for CO₂ evaporation inside tubes: Part II — An updated general flow boiling heat transfer model based on flow patterns, Int. J. Heat Mass Transfer 51 (1) (2008) 125–135, <http://dx.doi.org/10.1016/j.ijheatmasstransfer.2007.04.001>, URL <https://www.sciencedirect.com/science/article/pii/S0017931007002840>.
- [25] S. Agostinelli, et al., Geant4 collaboration Collaboration, Geant4: A simulation toolkit, Nucl. Instrum. Meth. A506 (2003) 250, [http://dx.doi.org/10.1016/S0168-9002\(03\)01368-8](http://dx.doi.org/10.1016/S0168-9002(03)01368-8).
- [26] A. Hanss, M. Hutter, J. Trodler, G. Elger, Solder process for fluxless solder paste applications, in: 2016 IEEE 66th Electronic Components and Technology Conference, ECTC, 2016, pp. 447–454, <http://dx.doi.org/10.1109/ECTC.2016.330>.
- [27] C. Lee, J. Kim, Fundamentals of fluxless soldering technology, in: Proceedings. International Symposium on Advanced Packaging Materials: Processes, Properties and Interfaces, 2005, 2005, pp. 33–38, <http://dx.doi.org/10.1109/ISAPM.2005.1432041>.
- [28] Y.-S. Lin, C.-H. Shih, W. Chang, Fluxless reflow of eutectic solder bump using formic acid, in: Proceedings of the 2009 12th International Symposium on Integrated Circuits, 2009, pp. 514–517.
- [29] W. Lin, Y. Lee, Study of fluxless soldering using formic acid vapor, IEEE Trans. Adv. Packag. 22 (4) (1999) 592–601, <http://dx.doi.org/10.1109/6040.803451>.
- [30] X. Shi, Z. Wang, W. Zhou, H. Pang, Q. Yang, A new creep constitutive model for eutectic solder alloy, J. Electr. Packag. 124 (2) (2002) 85–90, <http://dx.doi.org/10.1115/1.1462624>, arXiv:https://asmedigitalcollection.asme.org/electronicpackaging/article-pdf/124/2/85/5882916/85_1.pdf.
- [31] V. Bhanot, L. Zwalinski, J. Noite, H. Postema, J. Godlewski, T. Köttig, B. Verlaet, The CORA CO₂ cooling plant, in: 10th IIR Gustav Lorentzen Conference on Natural Refrigerants, 2012, URL <https://iifir.org/en/fridoc/the-cora-co2-cooling-plant-28456>.
- [32] P. Svihra, The design and construction of LHCb VELO upgrade modules, in: 2019 IEEE Nuclear Science Symposium (NSS) and Medical Imaging Conference, 2020, <http://dx.doi.org/10.1109/NSS/MIC42101.2019.9059693>, LHCb-PROC-2020-005, Poster-2020-1041.
- [33] M. Capeans, G. Darbo, K. Einsweiler, M. Elsing, T. Flick, M. Garcia-Sciveres, C. Gemme, H. Pernegger, O. Rohne, R. Vuillemeret, ATLAS Collaboration Collaboration, ATLAS Insertable B-Layer Technical Design Report, Tech. Rep., CERN, 2010, URL <https://cds.cern.ch/record/1291633>.
- [34] D. Hellenschmidt, M. Bomben, G. Calderini, M. Boscardin, M. Crivellari, S. Ronchin, P. Petagna, New insights on boiling carbon dioxide flow in mini- and micro-channels for optimal silicon detector cooling, Nucl. Instrum. Methods Phys. Res., A 958 (2019) <http://dx.doi.org/10.1016/j.nima.2019.162535>, 162535. 4 p, URL <https://cds.cern.ch/record/2689031>.
- [35] C. Da Via, P. Petagna, G. Romagnoli, D. Hellenschmidt, F. Munoz-Sanchez, N. Dann, Vertically integrated system with microfabricated 3D sensors and CO₂ microchannel cooling, Front. Phys. 9 (2021) 86, <http://dx.doi.org/10.3389/fphy.2021.633970>, URL <https://www.frontiersin.org/article/10.3389/fphy.2021.633970>.
- [36] A. Francescon, G. Romagnoli, A. Mapelli, P. Petagna, C. Gargiulo, L. Musa, J.R. Thome, D. Del Col, Development of interconnected silicon micro-evaporators for the on-detector electronics cooling of the future ITS detector in the ALICE experiment at LHC, Appl. Therm. Eng. 93 (2016) 10. 1367–1376, <http://dx.doi.org/10.1016/j.applthermaleng.2015.09.013>, URL <http://infoscience.epfl.ch/record/216255>.
- [37] A. Mouskeftaras, S. Beurthey, J. Cogan, G. Hallewell, O. Leroy, D. Grojo, M. Perrin-Terrin, Short-pulse laser assisted fabrication of a Si-SiO₂ Microcooling device, Micromachines 12 (9) (2021) <http://dx.doi.org/10.3390/mi12091054>, URL <https://www.mdpi.com/2072-666X/12/9/1054>.
- [38] M. Boscardin, P. Conci, M. Crivellari, S. Ronchin, S. Bettarini, F. Bosi, Silicon buried channels for pixel detector cooling, Nucl. Instrum. Methods Phys. Res. A 718 (2013) 297–298, <http://dx.doi.org/10.1016/j.nima.2012.10.014>, Proceedings of the 12th Pisa Meeting on Advanced Detectors, URL <https://www.sciencedirect.com/science/article/pii/S0168900212011461>.
- [39] P. Zellner, L. Renaghan, M. Agah, CMOS-compatible three dimensional buried channel technology (3DBCT), in: TRANSDUCERS 2009-2009 International Solid-State Sensors, Actuators and Microsystems Conference, 2009, pp. 192–195, <http://dx.doi.org/10.1109/SENSOR.2009.5285530>.
- [40] A. Mapelli, P. Petagna, M. Vos, Micro-Channel Cooling for Collider Experiments: Review and Recommendations, Tech. Rep., CERN, Geneva, 2020, URL <http://cds.cern.ch/record/2712079>.
- [41] R.o. Aaij, LHCb Collaboration Collaboration, Physics Case for an LHCb Upgrade II - Opportunities in Flavour Physics, and Beyond, in the HL-LHC Era, Tech. Rep., CERN, Geneva, 2018, 10347/15157. ISBN: 978-92-9083-494-6. URL <https://cds.cern.ch/record/2636441>, <http://arxiv.org/abs/10.17181/CERN.QZRZ.R4S6>.
- [42] R.o. Aaij, LHCb Collaboration Collaboration, Expression of Interest for a Phase-II LHCb Upgrade: Opportunities in Flavour Physics, and Beyond, in the HL-LHC Era, Tech. Rep., CERN, Geneva, 2017, URL <https://cds.cern.ch/record/2244311>.