

LHCb-DP-2024-001 April 19, 2024

# The LHCb VELO Upgrade Module Construction

K. Akiba<sup>2</sup>, M. Alexander<sup>10</sup>, C. Bertella<sup>12</sup>, A. Biolchini<sup>2</sup>, A. Bitadze<sup>12</sup>,

G. Bogdanova<sup>5</sup> , S. Borghi<sup>12</sup>, T.J.V. Bowcock<sup>11</sup>, K. Bridges<sup>11</sup>, M. Brock<sup>13</sup>,

A.T. Burke<sup>12</sup>, J. Buytaert<sup>7</sup>, W. Byczynski<sup>7</sup>, J. Carroll<sup>11</sup>, V. Coco<sup>7</sup>, P. Collins<sup>7</sup>,

A. Davis<sup>12</sup>, O. De Aguiar Francisco<sup>12</sup>, K. De Bruyn<sup>14</sup>, S. De Capua<sup>12</sup>, K. De Roo<sup>2</sup>,

F. Doherty<sup>10</sup>, L. Douglas<sup>10</sup>, L. Dufour<sup>7</sup>, R. Dumps<sup>7</sup>, D. Dutta<sup>12</sup>, L. Eklund<sup>15</sup>,

A. Elvin<sup>12</sup>, S. Farry<sup>11</sup>, A. Fernandez Prieto<sup>6</sup>, V. Franco Lima<sup>11</sup>, J. Freestone<sup>12</sup>,

C. Fuzipeg<sup>12</sup>, M.D. Galati<sup>2</sup>, A. Gallas Torreira<sup>6</sup>, R.E. Geertsema<sup>2</sup>, E. Gersabeck<sup>12</sup>,

M. Gersabeck<sup>12</sup>, F. Grant<sup>10</sup>, T. Halewood-leagas<sup>11</sup>, K. Hennessy<sup>11</sup>,

W. Hulsbergen<sup>2</sup>, D. Hutchcroft<sup>11</sup>, D. Hynds<sup>2</sup>, E. Jans<sup>2</sup>, D. John<sup>2</sup>, M. John<sup>13</sup>,

N. Jurik<sup>7</sup> , T. Ketel<sup>2</sup> , S. Klaver<sup>3</sup> , P. Kopciewicz<sup>4</sup> , I. Kostiuk<sup>2</sup> , M. Kraan<sup>2</sup>,

M. Langstaff<sup>12</sup>, T. Latham<sup>9</sup>, A. Leflat<sup>5</sup>, E. Lemos Cid<sup>6</sup>, V. Lukashenko<sup>2</sup>,

M. Merk<sup>2</sup>, M. Milovanovic<sup>7</sup>, M. Monk<sup>9</sup>, D. Murray<sup>12</sup>, I. Nasteva<sup>1</sup>,

A. Oblakowska-Mucha<sup>4</sup> , T. Pajero<sup>13</sup> , C. Parkes<sup>12</sup> , A. Pazos Alvarez<sup>6</sup>

E. Perez Trigo<sup>6</sup> (D), M. Perry<sup>12</sup> (D), F. Reiss<sup>12</sup> (D), K. Rinnert<sup>11</sup> (D), E. Rodriguez Rodriguez<sup>6</sup> (D),

J. Rovekamp<sup>2</sup>, F. Sanders<sup>2</sup>, L.G. Scantlebury Smead<sup>13</sup>, M. Schiller<sup>10</sup>, T. Shears<sup>11</sup>,

N.A. Smith<sup>11</sup>, A. Snoch<sup>2</sup>, P. Švihra<sup>12</sup>, T. Szumlak<sup>4</sup>, M. van Beuzekom<sup>2</sup>,

M. van Overbeek<sup>2</sup>, P. Vazquez Regueiro<sup>6</sup> , V. Volkov<sup>5</sup> , M. Wormald<sup>11</sup>, G. Zunica<sup>12</sup>

<sup>1</sup>Universidade Federal do Rio de Janeiro (UFRJ), Rio de Janeiro, Brazil

<sup>2</sup>Nikhef National Institute for Subatomic Physics, Amsterdam, Netherlands

<sup>3</sup>Nikhef National Institute for Subatomic Physics and VU University Amsterdam, Amsterdam, Netherlands

<sup>4</sup>AGH - University of Science and Technology, Faculty of Physics and Applied Computer Science, Kraków, Poland

 $^{5}Affiliated$  with an institute covered by a cooperation agreement with CERN

<sup>6</sup>Instituto Galego de Física de Altas Enerxías (IGFAE), Universidade de Santiago de Compostela, Santiago de Compostela, Spain

- <sup>7</sup>European Organization for Nuclear Research (CERN), Geneva, Switzerland
- <sup>8</sup>H.H. Wills Physics Laboratory, University of Bristol, Bristol, United Kingdom

<sup>10</sup>School of Physics and Astronomy, University of Glasgow, Glasgow, United Kingdom

<sup>12</sup>Department of Physics and Astronomy, University of Manchester, Manchester, United Kingdom

<sup>&</sup>lt;sup>9</sup>Department of Physics, University of Warwick, Coventry, United Kingdom

<sup>&</sup>lt;sup>11</sup>Oliver Lodge Laboratory, University of Liverpool, Liverpool, United Kingdom

<sup>13</sup>Department of Physics, University of Oxford, Oxford, United Kingdom

<sup>14</sup> Van Swinderen Institute, University of Groningen, Groningen, Netherlands, associated to <sup>2</sup>

<sup>15</sup>Department of Physics and Astronomy, Uppsala University, Uppsala, Sweden, associated to <sup>10</sup>

\* Contact author: sdecapua@cern.ch

#### Abstract

The LHCb detector has undergone a major upgrade for LHC Run 3. This Upgrade I detector facilitates operation at higher luminosity and utilises full-detector information at the LHC collision rate, critically including the use of vertex information. A new vertex locator system, the VELO Upgrade, has been constructed. The core element of the new VELO are the double-sided pixelated hybrid silicon detector modules which operate in vacuum close to the LHC beam in a high radiation environment. The construction and quality assurance tests of these modules are described in this paper. The modules incorporate  $200 \,\mu\text{m}$  thick, n-on-p silicon sensors bump-bonded to 130 nm technology ASICs. These are attached with high precision to a silicon microchannel substrate that uses evaporative  $CO_2$  cooling. The ASICs are controlled and read out with flexible printed circuits that are glued to the substrate and wire-bonded to the chips. The mechanical support of the module is given by a carbon fibre plate, two carbon fibre rods and an aluminium plate. The sensor attachment was achieved with an average precision of  $21 \,\mu m$ , more than 99.5% of all pixels are fully functional, and a thermal figure of merit of  $3 \,\mathrm{Kcm^2 W^{-1}}$ was achieved. The production of the modules was successfully completed in 2021, with the final assembly and installation completed in time for data taking in 2022.

Submitted to JINST

(c) 2024 CERN for the benefit of the LHCb collaboration. CC BY 4.0 licence.

## Contents

1	$\mathbf{Intr}$	oduction	1							
<b>2</b>	Requirements 3									
3	Det	ector design	<b>5</b>							
	3.1	General layout	5							
	3.2	Module layout	6							
		3.2.1 Microchannel plate	7							
		3.2.2 Hurdle	8							
		3.2.3 Tiles	8							
		3.2.4 Hybrids and cables	10							
		3.2.5 IV harness	11							
	33	Alternative designs and changes	12							
	0.0	3.3.1 Cooling substrate	12							
		3.3.2 Machanical support	14							
		2.2.2 Front and and control hybrids layout	14							
		2.2.4 Chus shajaa	14							
		5.5.4 Glue choices	14							
4	Mod	dule assembly procedures	16							
-	<i>A</i> 1	Database	17							
	4.1 1 2	Bare module assembly	18							
	4.2	Tiles placement	20							
	4.0	4.2.1 Clue peremeters and pettern	20							
		4.3.1 Glue parameters and pattern	20							
		4.5.2 The attachment procedure at Manchester	21							
	4 4	4.3.3 The attachment procedure at Nikner	23							
	4.4		26							
		4.4.1 Procedure at Manchester	26							
		4.4.2 Procedure at Nikhet	27							
	4.5	Wire-bonding and HV cables attachment	28							
	4.6	Cables attachment	29							
<b>5</b>	Module assembly measurements 31									
-	5.1	Bare module flatness	31							
	5.2	Tile metrology	32							
	0.2	5.2.1 Tile position	32							
	53	Clue layer thickness and tile flatness	33							
	5.0 5.4	Under layer tillexitiess and the flatness	00 94							
	0.4 5 5		04 24							
	0.0	wire-bond puil test	54							
6	Mod	dule validation	37							
	6.1	Test setup	37							
	6.2	Mechanical tests	38							
		6.2.1 Module displacement	38							

R	efere	nces		70		
9	Acknowledgements 69					
8	Con	nclusio	ns	67		
	7.3	Modul	le quality	61		
		7.2.10	Summary of I-V measurements	61		
		7.2.9	Summary of cooling performance tests	57		
		7.2.8	Summary of noise tests	57		
		7.2.7	Summary of PRBS tests	57		
		7.2.6	Summary of communication tests	57		
		7.2.5	Summary of displacement measurements	54		
		7.2.4	Summary of wire-bond pull tests	54		
		7.2.3	Summary of glue layer thickness and tile flatness measurements .	52		
		7.2.2	Summary of tile position measurements	52		
		7.2.1	Summary of bare module flatness measurements	51		
	7.2	Summ	ary plots	51		
	7.1	Produ	$  tion rate \dots $	50		
7	Ove	erview	of module production	<b>50</b>		
	6.4	'Therm	nal performance	48		
		6.3.5	Equalisation and Noise Scan	46		
		6.3.4	Pseudo random bit stream test	44		
		6.3.3	Time and fast control test	44		
		6.3.2	VeloPix communication test	43		
		6.3.1	IV curves	42		
	6.3	Electri	ical tests	42		
		6.2.2	Thermal cycles	41		

### 1 Introduction

The LHCb detector is a forward spectrometer located at the Large Hadron Collider (LHC) optimized for the reconstruction of b- and c- hadrons decays. The physics programme includes searches for effects beyond the Standard Model, mostly through precision studies of *CP* Violation and rare decays, searches for new hadronic states, electroweak precision measurements and a broad QCD phenomenology in the forward region. A major change of the experiment has been undertaken for the LHCb Upgrade I [1] which will operate for the next decade (Run 3 and Run 4). This detector incorporates a new vertex locator (VELO) system [2], which provides crucial input to the trigger with precise primary and secondary vertices reconstruction. The production and quality assurance of the modules of this system are described in this paper.

The LHCb Upgrade will allow the experiment to run at an instantaneous luminosity of  $2 \times 10^{33} \text{ cm}^{-2} \text{s}^{-1}$ , a factor five higher than in the previous runs. The experiment is expected to collect an integrated luminosity of 50 fb<sup>-1</sup> by the end of Run 4. While the previous experiment relied on a purely hardware based first level trigger, with a maximum readout rate of 1 MHz [3], a purely software trigger is now utilised. This includes the identification of heavy-flavour decays with vertex reconstruction, to expand the physics programme. The new trigger [4] requires all detector systems to be upgraded for 40 MHz readout and all front- and back-end electronics to be replaced in order to cope with the increased data rates. The Upgrade was installed during the LHC long shutdown LS2 (2019-2021), with the first data taken in 2022.



Figure 1: A CAD model of the layout of the VELO Upgrade detector, showing the RF foil in the centre (grey), the detector modules, Vacuum Feedthrough Board (light green) and the Opto-and-Power Board (purple).

Using the same mechanics as its predecessor, the upgraded VELO is separated into two retractable halves which surround the interaction point, as shown in Fig. 1. Each half is enclosed in a thin corrugated RF aluminium box, that keeps the detector vacuum separated from the main LHC vacuum, as well as shielded from beam-induced currents. During LHC beam injection, the two VELO halves are kept in a retracted position at a safe distance of 30 mm from the beams and only moved inward to enclose the beams once the conditions have been declared stable. In the closed position, the sensitive elements of the detectors are as close as 5.1 mm to the beams. Each VELO half consists of 26 VELO modules, mounted on an aluminium base which allows the required movements in both horizontal and vertical directions. The electrical signals and service voltages are routed to each module through an Opto-and-Power board (OPB) and a vacuum feedthrough board (VFB). The data and signals between the OPB and LHCb data acquisition system are carried by optical fibres.

### 2 Requirements

The VELO plays an essential role in the LHCb tracking system, by precisely reconstructing the location of *pp* collision vertices and secondary displaced vertices. This information is fed to the LHCb reconstruction algorithm and provides crucial information for the event selection. In order to fulfill its crucial role, a complete redesign and replacement of the VELO silicon sensors and electronics was necessary. The LHCb collaboration has opted for a solution based on hybrid pixel sensors. A new radiation hard ASIC (VeloPix [5]), based on the design of the Timepix/Medipix family of chips [6], was developed to cope with the higher data rates. The main changes with respect to the original detector are summarised in Table 1.

	2009 - 2018	2022 - 2031
RF-foil inner radius (minimum thickness)	$5.5 \text{ mm} (300 \mu\text{m})$	$3.5 \text{ mm} (150 \mu\text{m})$
Inner distance of active silicon detector	8.2 mm	$5.1 \mathrm{mm}$
Total fluence (silicon tip) $[n_{\rm eq}/{\rm cm}^2]^1$	$4 \times 10^{14}$	$\sim 8 \times 10^{15}$
Sensor segmentation	$r - \phi$ strips	square pixels
Pitch	$37-97~\mu m$	$55 \ \mu m$
Technology	n-on-n	n-on-p
Areas of active silicon detectors	$0.22 \text{ m}^2$	$0.12 \ {\rm m}^2$
Number of modules	42	52
Total number of channels	172 k	41 M
Readout rate [MHz]	1, analogue	40, binary
		(zero suppressed)
Whole-VELO data rate	150  Gbit/s	$\sim 2 \text{ Tbit/s}$
Total power consumption	800 W	$\sim 2 \text{ kW}$

Table 1: Specifications of the upgraded VELO compared to those of the original version.

The requirements of the VELO Upgrade system are described in [1] and some key elements are summarised here. The distinctive signature of b- and c- hadrons decaying in LHCb is given by tracks that originate from secondary vertices with a large impact parameter to all primary vertices. The vertex resolution is improved by reducing the pixel size, the amount of material traversed by the track, the distance between the active area and the interaction point, and by increasing the distance between the measured points. The spatial resolution requirements led to the 55  $\mu$ m pitch pixels, with a minimum distance of only 5.1 mm from the beams, operated in vacuum with 52 modules spread just over 1 m, each with an individual slim support structures, and arranged to close around the beams. The choice of a 55  $\mu$ m pitch led to the need of a < 30  $\mu$ m (i.e. half pixel size) precision assembly and metrology, as well as stability in the operation under the thermal and mechanical environment.

It is expected that the sensor with highest occupancy will receive on average 8.5 charged particles per bunch crossing. The bunch crossing rate in LHCb is expected to be 27 MHz on average, with a 40 MHz peak (25 ns spacing) due to variations in the LHC filling scheme. In the high occupancy regions of the detector, this leads to an expected peak pixel hit rate of  $\sim 900 \text{ MHz}/\text{ASIC}$ , with the whole VELO reaching a total data

 $<sup>^{1}1\,\</sup>mathrm{MeV}$  neutron equivalent per cm<sup>2</sup>.

rate of 2.85 Tbit/s. After 50 fb<sup>-1</sup> of integrated luminosity exposure, it is expected that some chips will have accumulated an integrated flux of  $8 \times 10^{15} n_{\rm eq}/{\rm cm}^2$ . With such a dose, leakage currents of about 200  $\mu$ A/cm<sup>2</sup>, or ~ 7 nA per 55×55  $\mu$ m<sup>2</sup> square pixel, are expected at -25 °C, with a bias voltage of 1000 V. For these reasons, both the VeloPix and the silicon sensors are designed to cope with high data rates and large non-uniform irradiation.

Such demanding front-end processing requires a significant amount of power and highly efficient on-detector cooling to protect the silicon from thermal runaway. The new VELO modules have a large heat generation in the front-end electronics (up to 30 W per module). The adopted cooling solution must ensure that all modules are kept below -20 °C. Since the operating temperature of the CO<sub>2</sub> in the cooling plant is approximately -30 °C, the temperature drop between the sensor and the coolant is required to be kept below 8 °C. The solution must also ensure a minimal amount of material and avoid leaks of the coolant in the detector vacuum. The module cooling design is fully integrated within the module. The coolant, evaporative CO<sub>2</sub>, is circulated within microchannels directly etched into thin silicon substrates that provide the mechanical support to the silicon sensors in the module. Silicon features an excellent heat transfer coefficient, has a high radiation length and matches the coefficient of thermal expansion of the readout electronics. The module design and its components are discussed in the following chapter.

### 3 Detector design

### 3.1 General layout

The new detector consists of 52 modules. Each module is identical in order to simplify the assembly process and quality control. They are grouped into two retractable halves of 26 modules, as illustrated in Fig. 2. Modules are distributed perpendicular to the beamline and arranged to fully cover the LHCb pseudorapidity acceptance  $(2 < \eta < 5)$  and ensure that at least four pixel layers are crossed by any track from the interaction region, for all azimuthal directions [7]. A number of modules are placed upstream of the interaction point in order to improve the reconstruction of primary vertices. The minimum spacing between modules is 25 mm for both halves, although in one half they are displaced by +12.5 mm in z to allow for overlap when the detector is in the closed configuration and thus provide a full azimuthal coverage.



Figure 2: Schematic cross-section at y = 0 with an illustration of the z-extent of the luminous region and the nominal LHCb acceptance.

The detector is operated in vacuum, separated from the LHC vacuum by a thin aluminium alloy RF (Radio Frequency) foil. The RF foil provides protection to the front-end electronics from beam-induced currents as well as providing beam wakefield suppression. The RF foil surface is corrugated to follow the L-shaped active area of the modules, which serves the purpose of creating a rectangular active area in the x - y plane split across both modules, as shown in Fig. 3. The silicon detectors are rotated 45° around the z axis, this minimises any risk of the detectors grazing the RF foil during installation.

A custom-made vacuum feedthrough board (VFB) is used to route the high speed data signals, the control signals, the LV power supply ( $\sim 2.5$  A) and the bias voltage (< 1000 V) inside the vacuum. Due to the high rate data requirement, a careful design of the custom transmission lines was needed [8]. The VFB is a 2 mm thick 12-layer printed circuit, which is glued to an aluminium flange with Araldite 2011 on the air side and Araldite 2020 on the vacuum side. A photograph of the vacuum feedthrough board, glued to an aluminium flange, is shown in Fig. 4. This board was also used in the module



Figure 3: Left: engineering drawing showing the clearance, in mm, between the RF foil and the modules. Right: schematic layout of the chips in the x - y plane with the VELO in the closed configuration. Half the chips are placed on one side of the module (grey) and half on the opposite side (blue). The countour in red highlights the VeloPix ASICs belonging to the same module.

quality assurance vacuum tank test setup.

The feed-through board has connectors soldered on both sides of the vacuum barrier, allowing for a fast connection to the module inside the vacuum, and to the OPB [9] in the air side. The OPB is the interface between the detector modules and the off-detector electronics. Its main purpose is to perform the electrical to optical conversion of the data links and of the control and monitoring signals, as well as to perform the DC/DC conversion of the supply voltages and distribute them to the front-end hybrid.

### 3.2 Module layout

A VELO pixel module is a double-sided detector structure with a microchannel plate at its core and two tiles arranged in an L shape attached on either side of the microchannel plate. Each tile consists of three VeloPix ASICs bump-bonded to a silicon sensor. A custom-made invar cooling connector, with  $CO_2$  supply and return capillaries, is soldered to the microchannel substrate. The invar connector is glued to a *hurdle*, which consists of a carbon fibre plate, two carbon fibre poles and an aluminium foot. On both module sides, two VeloPix chips and one GBTx control chip are hosted each on a front-end hybrid. These carry services to the chips and route their data out. Interconnect cables connect the VeloPix hybrids to the GBTx hybrid. The GBTx chip [10,11] ensures the configuration and synchronisation of the VeloPix ASICs by decoding and distributing the LHCb clock and other control signals. VeloPix data are routed through four thin PCB cables each with up to seven differential links, for a total of 20 links per double-sided module. Power is delivered via an assembly of 20 silicone-coated copper cables and a PCB transition





Figure 4: A photograph of an opto-and-power board (top) and of a vacuum feedthrough board (bottom) glued to its aluminium flange (vacuum side on the left).

bridge that reduces the thickness of the cables and thus the material budget close to the beamline. The transition bridge also provides mechanical support for the various cables. The bias voltage is delivered via thin PCB tapes which are wire-bonded to the top surface of the sensors. All parts of a VELO module are shown in Fig. 5 and described in more detail in the next sections.

### 3.2.1 Microchannel plate

Since the module is operated in vacuum, power dissipated in the electronics must be efficiently removed. The microchannel cooling technology was chosen since it provides an excellent thermal performance, minimum material budget and no mismatch of thermal expansion with respect to the tiles. Fig. 6 shows a drawing of the microchannels. The power dissipation of the electronics components on the module is removed by the evaporative  $CO_2$  coolant running through the track-like cooling channels. This technology builds on the experience of operating the first evaporative  $CO_2$  cooling system in the original VELO [12]. In total, 19 channels are embedded into a silicon plate of 500  $\mu$ m thickness. Each channel has its own input and output and the coolant distribution is done under the fluidic connector. The connector is made of invar and provides interface between the microchannels and the cooling line. An extensive R&D programme on the soldering of the connector to the microchannel plate was needed, in order to establish a reliable connection and adequate quality control procedure [13]. The initial section of the channels (roughly 4 cm long is narrower (60  $\mu \text{m} \times 60 \,\mu \text{m}$ ) to ensure an even distribution of flow and prevent instabilities. The main channels are  $120 \,\mu\text{m} \times 200 \,\mu\text{m}$  and around 26 cm long. Overall, the cooling substrate outer dimensions are around  $11.6 \text{ cm} \times 11.4 \text{ cm}$ .



Figure 5: Left: connector side and right: non-connector side of a fully-assembled VELO module.

### 3.2.2 Hurdle

The hurdle provides the microchannel plate with a mechanical support that is light and thermally insulates the microchannel from the mounting point. It consists of a carbon fibre plate and two carbon fibre rods fixed to an aluminium foot.

The module foot is machined out of aluminium 5083. It contains a precision-drilled hole for the mounting bolts and the alignment dowel, providing the module with a solid and accurate mounting point on the detector base. Two additional holes are used to insert and glue the carbon fibre rods. These are cylindrical tubes made of T700 carbon fibres, aligned along the length of the cylinder and glued together with a high performance epoxy resin. The choice of the alignment direction of the fibres ensures minimal deformation due to thermal stress. At the other end of the rods, a flattened section cut is used to glue them to a carbon fibre plate (midplate). The midplate is a thin sheet of carbon fibres aligned in the 0°/90° and  $-45^{\circ}/+45^{\circ}$  directions. The midplate is directly glued to the invar cooling connector with Araldite 2011 and it is the only point of contact between the microchannel assembly and its mechanical support. A 3D-printed clamp, made of polyether ether ketone (PEEK) and attached to the carbon fibre rods, secures the LV cables and cooling pipes and relieves them from mechanical and thermal stress.

#### 3.2.3 Tiles

A VELO tile consists of a pixelated planar silicon sensor bump-bonded to three pixelated VeloPix chips (ASICs). The VeloPix [5] contains an active matrix of  $256 \times 256$  ( $55 \times 55 \ \mu m^2$ ) pixels, for a total sensitive area of  $1.98 \text{ cm}^2$ . The chip is fabricated in 130 nm technology<sup>2</sup>,

<sup>&</sup>lt;sup>2</sup>by Taiwan Semiconductor Manufacturing Company, Hsinchu 300-096, Taiwan, R.O.C.



Figure 6: Layout of the final microchannel design, overlaid with the hybrid pixel tiles and front-end hybrids. The tiles on the front and back of the cooler are shown in brown, and the position of the front-end hybrids on the front side of the module in green.

which has proven radiation hardness above 400 MRad [14]. The sensor is a 200  $\mu$ m thick, non-p silicon sensor<sup>3</sup>, with an overall dimension of 43.470×14.980 mm<sup>2</sup> (including a 450  $\mu$ m wide inactive region for the guard rings). This builds on the experience of operating the first n-on-p silicon sensor in the original VELO [12]. The sensor is implanted with square pixels with a pitch of 55  $\mu$ m to match the pixel array of the read-out electronics. The flip chip bonding<sup>4</sup> consists of almost two hundred thousand solder bumps per VeloPix triplet. Since the sensors are attached to three ASICs in line, the pixels at the inter-chip regions are elongated to the size of 55 × 137.5  $\mu$ m to minimize the dead area. The schematic of a sensor tile is illustrated in Fig. 7.

Both the ASICs and the sensors come with etched alignment markers, which are used during the positioning and attachment of the tile to the microchannel plate, as well as during metrology measurements after attachment. Fig. 8 shows a photograph of a VELO tile and two microscopic images of the corner of the ASIC and of the sensor.

The nomenclature [15] of the tiles and the ASICs is shown in Fig. 9. On the fluidic connector side, the tiles are named CLI (VP0) and CSO (VP3), while on the non-connector side they are named NLO (VP1) and NSI (VP2). Within each tile, the ASICs are labelled with a number that ranges from 0 to 2. While CSO and NLO are fully enclosed in the microchannel plate, tile CLI overhangs the plate by 5 mm along its full length, and tile NSI by the same amount along its ASIC 0. Each VeloPix has four output links, however

<sup>&</sup>lt;sup>3</sup>made by Hamamatsu Photonics K.K., Hamamatsu, Shizuoka 435-8558, Japan.

<sup>&</sup>lt;sup>4</sup>carried out at ADVACAM Oy, Tietotie 3, FI-02150 Espoo, Finland.



Figure 7: Schematic of the sensor tile, showing the overall dimensions of the sensor and its pixel layout. An additional row of pixels, indicated in dark red, provides a connection between the ASIC ground and the innermost guard ring of the sensor.



Figure 8: Top: a photograph of a VELO tile. Bottom left: microscopic image of the ASIC corner and its alignment marker. Bottom right: microscopic image of the sensor round corner with its cross- and disk-like shaped alignment markers.

the number of links that are enabled is configurable to just one link, or two links or four links. The VeloPix chips in the hottest region of the detector (VP0-2 and VP2-0) have all four links enabled, while those located in the lower occupancy regions operate with either two or one link. In a VELO module, there are a total of 20 output links.

#### 3.2.4 Hybrids and cables

On a VELO module, the routing of the the outward-bound data and the distribution of control signals and services is provided by a set of hybrids. These are  $390 \,\mu\text{m}$  thick, four-layer, flexible printed circuits and come in two types: the front-end hybrids and the GBTx hybrids, as shown in Fig. 10. The front-end hybrids provide the electronic interface to the VeloPix chips through wire-bonds. The GBTx hybrids carries the gigabit



Figure 9: Mapping of the tiles and ASICs. The schematics also show how the 20 data links are distributed among the 12 VeloPix chips.

transceiver chip [11] and is responsible for the distribution of timing and fast controls to the ASICs, as well as providing slow controls and monitoring of the bias voltage. Each hybrid also carries a temperature probe (NTC).



Figure 10: A photograph of VeloPix and GBTx hybrids.

Each hybrid connects to two different cables: the interconnect cables, which connect the front-end hybrids to the GBTx hybrid, and the data cables, which route the outputs of each tile to the off-module electronics. Each VELO module is equipped with a set of four front-end hybrids, two GBTx hybrids, four interconnect cables and four data cables, evenly split onto the two sides of the module. The high voltage (HV) is carried to the silicon sensors by the HV tapes, which are connected to the microchannel plate with a fast-curing glue (Araldite 2012) and afterwards wire-bonded to the surface of the sensor. Data, interconnect and HV cables are shown in Fig. 11.

#### 3.2.5 LV harness

Power is delivered to the chips via an assembly of 20 silicone-coated copper cables and a PCB transition bridge. The transition bridge is glued to the carbon fibre legs where it establishes a mechanical link while changing from thinner gauge cables close to the readout, to thicker ones going back to the module foot. The LV cables are connected to the front-end and control hybrids via custom-made PCB connectors. At the other end of the LV harness, pins are inserted in the LV foot connector. A photograph of a LV harness is shown in Fig. 12.



Figure 11: A photograph of a set of interconnect cables (left), data cables (middle) and HV tapes (right).



Figure 12: A photograph of a LV harness, showing the cables, the custom-made PCB connectors, the LV bridge pieces, the LV pins and the LV foot connector.

### 3.3 Alternative designs and changes

During the VELO Upgrade R&D programme, a range of solutions were investigated for many elements of the module. Here we provide brief comments on some of the design choices that were made and changes made during the process.

### 3.3.1 Cooling substrate

The R&D focussed on the microchannel solutions as it is well suited to the needs of this project. Due to the complexity of the silicon microchannel technology, described above, two other microchannel implementations were considered: 3D printed titanium substrates and ceramic plates with embedded stainless-steel capillaries (see Fig. 13).

The 3D printed titanium option offers a cheap solution, flexible layout and fast turnaround for prototyping. Cooling pipes can be brazed or welded directly to it and the restrictions needed at the inlet to trigger the  $CO_2$  evaporation can be implemented easily in the design

of the substrate. A wall thickness of  $100\,\mu{\rm m}$  was already achievable at the time of our R&D programme.

The ceramic solution was investigated by manufacturing a plate in aluminium nitride, with precisely machined grooves to house stainless steel cooling pipes of 0.64(0.4) mm outer(inner) diameter. The inlet block is equipped with orifices of 140 µm diameter for the expansion of the coolant, which is then routed to the four cooling pipes. The ceramic offers an excellent heat transfer (about 100 W/mK at 25 °C) and a robust support for the tiles.

Although both options had the potential to be viable solutions for the VELO Upgrade, they were rejected in favour of the microchannel technology. The use of a thick ceramic plate, with embedded steel capillaries, increased the material budget in the LHCb acceptance; deviations from planarity were observed on the surface, although smaller than the expected thickness of the glue layer used to attach the ASICs; the manufacturing procedure consisted of a several-step manual process. Microscope inspections of the 3D-printed titanium substrates showed that the surface was covered with partially melted powder particles, sometimes reaching a height of tens of microns; the possible formation of these structures on the internal walls, posed a risk of perturbation to the flow of the coolant. Taking into account these considerations, the additional R&D required to achieve an optimal cooling performance from either of these alternative solutions was incompatible with the project timeline with the effort available.



Figure 13: CAD models of the 3D printed titanium substrate (left) and the aluminium nitride plate with embedded cooling pipes (right).

### 3.3.2 Mechanical support

Alternative designs of the module mechanical support were considered. The original VELO used a carbon fibre structure with an open triangular cross-section [3]. A support made from a single 1 mm thick carbon fibre sheet was considered, with the microchannel wafer glued to the carbon fibre support. This provided good stability in X-Y and by design had flexibility along Z. This flexibility led to individual module vibrational stability issues, with a carbon fibre constraint system along the VELO length to position and secure the modules. In the hurdle design an alternative mid-plate was considered, made of a silica glass mid-plate instead of carbon fibre, which was rejected after cracking occurred during production of prototypes. Alternatives were studied under stability in vibration tests, thermal studies, in vacuum operation, and in the construction of prototypes. The final design was chosen based on its low-material, ease of construction and advantages of the open geometry for routing of the cables and cooling pipes.

### 3.3.3 Front-end and control hybrids layout

A single hybrid for each side of the module was originally envisaged. The choice of separating the front-end electronics into three parts, rather than one large single hybrid, was implemented to reduce the stress on the module. The thickness of the hybrids, driven by the copper cross section required to bring sufficient power to the chips, leaves them significantly stiff. Consequently, a single stiff hybrid with a coefficient of thermal expansion different to that of silicon, would exert an unacceptable force on the microchannel plate over the 50 °C thermal change between room temperature and the operational temperature of the detector.

### 3.3.4 Glue choices

Significant studies were carried out to select and test suitable adhesives for each of the different gluing steps in the module assembly and optimise the way in which they were used. The gluing of the sensor tiles to the module, the hybrid to the module and the high voltage connection required particular attention. Thermal tests, including thermal cycling, test after irradiation and mechanical shear tests were performed. The deposition patterns were optimised as further discussed in Sect. 4.3.1.

During prototyping it was observed that glue between the ASICs increases the capacitance of the system increasing the electronic noise, so any glue excess filling the space between the ASICs must be avoided. The glue selected for the tiles' attachment was a thermally conductive epoxy encapsulant<sup>5</sup> [16]. Detailed tests showed that the mechanical connection between the tiles and the microchannel substrate was not ideal, leading to adhesion failures on test samples as well as on an early module. The cause was traced to the hygroscopic nature of the chosen catalyst, causing the accumulation of a layer of humidity on the deposited patterns that hinders the bonding properties of the glue. The moisture layer becomes visible only after thermal cycling and worsens over time, progressively weakening the adhesive strength and shortening the lifetime of the modules. For this reason, an extensive study was carried out, to test alternative epoxies for their mechanical strength, radiation hardness and thermal performance. At the end

 $<sup>^{5}</sup>$ Stycast 2850FT with catalyst 9

of this campaign, the glue chosen for the module production was the same as originally proposed, but in combination with a different catalyst<sup>6</sup>, and the additional use of a heat gun treatment (see Sect. 4.3.1).

The main requirements on the glue for the attachment of the hybrids to the substrate were elasticity and good radiation hardness. The elasticity of the adhesive is needed due to the different coefficient of thermal expansion (CTE) of the silicon substrate and the hybrid PCB (due to its copper layers). There was some evidence in prototypes of microchannel rupture when testing adhesives without flexibility. Since the cooling of these components is not as critical as for the tiles (less produced heat), an alkoxy silicone adhesive<sup>7</sup> was chosen, which gave suitable results.

The use of conductive adhesives for the attachment of the high voltage tapes to the back side of the silicon sensors was briefly investigated. No reliable candidate suitable for a heavy radiation environment was identified. Hence, the HV tapes were glued<sup>8</sup> and the electrical connection made by wire-bonding.

 $<sup>^6\</sup>mathrm{Stycast}$  2850FT with catalyst 23LV

<sup>&</sup>lt;sup>7</sup>Loctite SI 5145.

 $<sup>^{8}</sup>$ using Huntsman Araldite 2012.

### 4 Module assembly procedures

Module production took place in two assembly sites, the University of Manchester (Manchester, UK) and Nikhef (Amsterdam, NL). The assembly process requires high precision and consistency between the two sites, thus a detailed construction procedure was developed to ensure excellent quality across the whole production (see also [17–19]). In order to achieve this goal, a number of checks are conducted after each of the assembly steps. They include visual inspections, metrology, electrical tests, thermal and mechanical performance measurements in air and in vacuum. The mounting of the modules onto the VELO bases, as well as the testing of the full DAQ chain and cooling system, took place at the University of Liverpool (Liverpool, UK). From there, the two fully assembled VELO halves were transported to CERN (Geneva, CH) for final commissioning of the whole detector.

The module assembly process consists of five major steps, shown as blue blocks in the flow diagram in Fig. 14: construction of the bare module, attachment of tiles, attachment of hybrids, wire-bonding and finally cabling. After each of these steps, a number of tests is carried out and they are shown as orange blocks in the flow diagram. The information from both the assembly and testing steps is uploaded to an online database and automatically analysed, thus providing an instantaneous feedback about the quality of the performed task. The database also guarantees an off-site record of all the actions performed on a module, which will serve as future reference for the long term, and is discussed below. The progress of the module construction through the five module assembly steps is shown in Fig. 15 and these steps are described in detail in the following sections.



Figure 14: A flow chart of the assembly steps. Processes are shown in blue and measurements in orange.



Figure 15: Pictures of a VELO upgrade module after each of the five assembly steps. From left to right: bare module assembly, tiles attachment, hybrids attachment, wire-bonding and cabling.

### 4.1 Database

The database provides tracking of the module assembly as it progresses, allowing to store all relevant data. It contains information about all the components that are used to build a module, along with every process and test carried out on individual parts as well as modules.

The database uses PostgreSQL [20], accessed via Django framework [21] based on Python [22]. The Python libraries Bokeh and Holoviews [23] are used to enable the interactive control of the generated plots, such that a simplified data analysis can be performed immediately after uploading the results of a test. This provides a quick visual representation of the measurements and allows for an immediate feedback on the quality of the assembly procedure at any stage. An example of the modules page in the database is shown in Fig. 16, where the rows represent individual modules and the columns correspond to the assembly/testing steps, arranged in the same order as the tests are performed during construction.

ID	Nickname	Location	ABM	ΡI	SF	Completed	Grade	Quality	Issues	Ready
MOD_105	M105	CERN, Point 8	B (A)	A	в	•	Grade A	98.3	None	•
MOD_104	M104	CERN, Point 8	A (A)	A	A	•	Grade A	98.7	None	•
MOD_103	M103	Liverpool, Velo	A (A)	A	Α	•	Grade A	100.0	None	•
MOD_102	M102	Liverpool, Velo	A (A)	A	A	0	Grade A	99.3	None	•

Figure 16: Extract from an overview page of the construction database showing a few assembled modules. Each row corresponds to a module. A few assembly and testing steps are shown as columns: the bare module assembly (ABM), a photo inspection of the bare module (PI) and the substrate flatness measurement (SF). The table also contains a grade and a quality number.

Each assembly or test step is graded automatically by an algorithm that assigns a

letter based on the performance, defined as follows: [A] excellent; [B] very good, [C] good, [D] one missing functionality, [F] multiple missing functionalities. Grades A, B and C are assigned to fully functional modules, which meet the production quality requirements and are recommended to be suitable for installation in the VELO detector. The automatic gradings for the validation tests are based on the quality criteria described in Sect. 6.

The table shown in Fig. 16 also provides quick access to any process or test performed, by simply clicking on the corresponding cell and landing on a page with all the information concerning that particular step. In addition, the first column gives access to a more detailed module page, containing a comprehensive list of every action taken during the assembly. A small section of a module page is shown, as an example, in Fig. 17.

Summary									
ID	MOD_110 Nickname		M110						
Location	CERN, Point 8 v +								
Reason for Location Change									
	Fill if changing location; ex. "For test beam"								
Location Summary	ion Summary 2021-11-16: moved to Liverpool, Velo; reason: For installation 2021-04-26: manufactured at Manchester; all ok								
Grade	Grade A Manual grading If ticked will ignore automatic grading procedure								
Module Sign-off (fields become available after module is finished and inspected)									
	Completed available after 'final photo-inspection (finished i performed; reverts to 'false' if adjacent 'signee' 'checklist' fields not filled	Signee Stefano de Capua, Manch person responsible for signing-o			e Capua, Manchester	*	+		
Bare Module Assembly									
Bare Module	by Julian Freestone on 2021-04-26	Substrate		SUB_106					
Tile Gluing									
C side tile gluing	by Stefano de Capua on 2021-04-28	CLI tile		TIL_275		CSO tile	TIL	_277	

Figure 17: Extract from a database module page showing the summary and two of the assembly and measurements steps performed on a module.

### 4.2 Bare module assembly

The first step in the construction of a VELO module consists of assembling together the mechanical and cooling structure, forming a bare module onto which the other components will then be attached. The base of the module is an aluminium foot with a steel dowel pin on one side, used for alignment purposes. Two carbon fibre rods are glued into the foot and a carbon fibre plate is glued at the other end of the rods. This structure is then inserted in the bare module jig, shown in Fig. 18, together with the microchannel plate and a cooling clamp which relieves the microchannel plate from thermal and mechanical stress from the rest of the system.

The bare module jig ensures that the microchannel plate is placed within tolerance with respect to the dowel hole in the foot. Three alignment pins are used for this purpose (which the L-shaped edge of the microchannel plate fits against). After correctly positioning it in the jig, the microchannel plate is pushed flush against these three pins. A small amount



Figure 18: Drawing of the bare module jig. The jig is used to align the microchannel plate to the mechanical support (hurdle) and hold it in place during curing of the glue.

of glue<sup>9</sup> is deposited on the top surface of the cooling connector and the carbon fibre plate is pressed onto it and clamped.

After curing, the module is moved into a different jig, called the turn-plate, and held in place by the frame at two fixation points at its sides. The turn-plate is used to carry out the attachment of tiles and hybrids, as shown in Fig. 19. A local reference system is defined using a glass marker located at the top of the turn-plate, and visible from either side, and the dowel hole at the bottom of the turn-plate. A transformation is then performed to translate all coordinates into the LHCb reference frame, in which the origin coincides with the interaction point, x is oriented along the long side of the module, y is oriented along the short side of the module, and z along the beam line.



Figure 19: A bare module installed in the turn-plate. Both the connector and non-connector side of a module are shown.

<sup>&</sup>lt;sup>9</sup>Huntsman Araldite 2011.

### 4.3 Tiles placement

Placement of the tiles on a module is one of the assembly steps that requires the highest precision, aiming at achieving a precision better than  $\pm 30 \,\mu\text{m}$  relative to the dowel hole on the module foot. In their nominal position, the tiles attached on opposing sides of the module should have an overlap of 200  $\mu\text{m}$ . In addition, the absolute position of the tiles with respect to the module foot in the horizontal plane is also critical, given that the smallest nominal clearance between sensors and RF foil is 890  $\mu\text{m}$ .

The tiles are first placed in a dedicated alignment system, held in place by vacuum chucks (see Sect. 4.3.2 and 4.3.3). The alignment system comprises a camera for pattern recognition and motorized motion stages that allow for micron-precision positioning. The alignment process locates the markers etched on the corners of each VeloPix ASIC (see Fig. 20) and uses this information to automatically move the alignment stages to the target position. Afterwards, the two tiles are transferred to a dedicated frame system for adhesive deposition and attachment to the microchannel substrate. Fig. 20 also shows the markers present on the back of the sensor that, after attachment, are visible in the two overhanging tiles (one per side) and can be used for front-back cross-reference.

The attachment procedure followed by the two module production sites, although based on the same approach, slightly differs in terms of equipment employed and assembly sequence. Both procedures are described in the following sections.



Figure 20: Markers used by the alignment system camera for pattern recognition. Left: Marker in the corner of an ASIC. Right: markers on the back of a sensor.

#### 4.3.1 Glue parameters and pattern

The glue layer connecting the ASICs with the microchannel substrate is the only path for heat dissipation in the final detector: maximising the glue coverage, while minimising its thickness, is critical to ensure sufficient cooling performance at the overhanging sensor tip. Although a thin, uniform layer of glue underneath the tile ensures the best cooling performance, the glue layer also needs to be thick enough to absorb any mechanical stress induced by the temperature changes. In addition, the maximum size of the filler particles is about 60  $\mu$ m, which needs to be taken into account to prevent damages to the silicon substrate when pressing the tiles down. Taking into account these considerations, the targeted glue layer thickness underneath each tile was set to 80  $\mu$ m.

Furthermore, the glue layer should ideally cover as much area of the tile as possible, to maximise the heat transfer, but at the same time spillage in the gaps between adjacent ASICs should be avoided. In addition, since the detector is operated in vacuum, air pockets trapped in the glue layer may lead to damage and an appropriate glue pattern needs to be chosen in order to minimise the chances of trapping air. For these reasons, a glue dispenser robot was chosen to achieve the accuracy and control required to meet all the aforementioned requirements. The glue deposition was performed using a programmable robot<sup>10</sup> coupled to a syringe and a needle. The fluid dispenser runs on compressed air.

An extensive R&D campaign was carried out to investigate different adhesives and deposition patterns [16]. Initially, a snake pattern (see Fig 21) was proposed, which allows for the air to escape through the openings and ensures a good coverage. However, in case of accidental connections between the points near the openings, an air pocket could be created. Hence, a star shaped pattern (see Fig. 21) was evaluated and proved to be inherently better at preventing air trapping in the pattern. The star pattern has also a lower chance to cause glue spillage, as the endpoints of the star can be kept at a safe distance from the ASICs, while still providing an excellent coverage. The chances of spillage were further reduced by concavely shaping the core of the star in a hourglass shape (see Fig. 21), achieving a square shaped pattern after attachment.

The glue selected for the tiles attachment was a thermally conductive epoxy [16] (see Sect. 3.3.4). The initial step of the glueing procedure consists of heating the epoxy to  $50 \,^{\circ}$ C for one hour and then leaving it to cool down to room temperature. This procedure helps to standardise the properties of the glue before the attachment is performed, and it also ensures a proper mixing between the main resin and the aluminium oxide filler. The epoxy is then mixed to the catalyst in a planetary mixer, poured into a syringe and its viscosity is checked to ensure it lies within an acceptable range. The viscosity is determined before every attachment by measuring the amount of glue dispensed per unit time at a given force. This allows to set the line deposition speed according to the actual viscosity and hence ensures the same amount of glue is deposited in every module. The deposition is then started, bringing the tip of the needle at 200 µm from the surface of the tiles, which are held in place by vacuum chucks. After the glue pattern has been deposited, a heat gun pointing towards the tile is used to heat up the top layer of the glue to  $60 \,^{\circ}$ C for 1 minute, such that any moisture that might have accumulated on the surface of the glue during deposition, and that can lead to adhesion failures, is removed.

#### 4.3.2 Tile attachment procedure at Manchester

The alignment and attachment of the tiles in Manchester is performed using a custommade set of jigs and an alignment system (Fig. 23). The system consists of two translation

<sup>&</sup>lt;sup>10</sup>Fisnar F5200N.



Figure 21: Comparison of glue patterns deposited on tiles. Top left: snake pattern. Top right: star pattern. Bottom left: star pattern with concave inner sides. Bottom right: infrared image taken after tiles attachment, showing the resulting squared shape patterns.



Figure 22: Photograph showing the glue dispenser, the heat gun and four silicon tiles.

and one rotation stage and uses a camera for pattern recognition. A pair of tiles are placed on the stages and held by vacuum. The alignment process locates the markers on the ASICs of each tile and adjusts the x, y and  $\phi^{11}$  coordinates of each individual stage with respect to a turn-plate, which will hold the bare module during tiles attachment. The turn-plate has the same mechanical fixing to the bare module foot as in the final VELO base.



Figure 23: Photographs of the custom-made set of jigs and alignment system used for tiles attachment in Manchester. Left: the alignment stages and camera. Right: the transfer plates, installed under the glue dispenser.

After alignment, the two tiles are transferred to a dedicated transfer frame (see Fig. 24) and the transfer is validated visually by looking through holes in the transfer plate which are located in front of the ASIC markers. The same procedure is then repeated for the second pair of tiles. The two transfer-plates, carrying the tiles for both the connector and non-connector sides, are places under the glue dispenser and the glue patterns are deposited on all four tiles, on the VeloPix side. The bare module is then installed in the turn-plate and the two transfer plates are placed either side of the turn-plate and held in place with retention springs. The three jigs are left in this configuration for 24 hours while the glue cures.

Photographs of the two sides of a module, after tiles attachment, are shown in Fig. 25.

### 4.3.3 Tile attachment procedure at Nikhef

As in Manchester, the placement of the tiles on the bare module is performed using custom-built mechanical jigs, motorized motion stages and a camera. Fig. 26 shows the turn-plate, the mechanical frame used to hold the module during tiles attachment.

<sup>&</sup>lt;sup>11</sup>defined as the angle in the xy plane.



Figure 24: A CAD model showing how the transfer plates (in green and red) are brought together and attached to the turn-plate (in blue) which holds the bare module.



Figure 25: Photographs of the two sides of a module installed in the turn-plate, after tile attachment.

Two main stages allow to move the turn-plate in x and y, relative to the camera, while three smaller stages move the tile relative to the turn-plate in x, y and  $\phi$  (Fig. 27). Finally, one stage is responsible for controlling the camera focus.

Fig. 28 shows the transfer plate used for holding the tiles during attachment to the bare module. There are two types of transfer plates, one for the attachment to the connector side of the bare module and one for the non-connector side. The relative alignment between the turn-plate and the transfer plates is determined by a set of kinematic mounts that allow to reposition the two jigs with an accuracy of a few micrometers.

Initially, the turn-plate is installed in the pick and place machine with the VeloPix side facing upwards. Then, the positioning jig is placed in the turn-plate and moved to the



Figure 26: The turn-plate, shown inside the pick and place machine. The orange arrows indicate the position of the glass markers used to determine the position of the jig relative to the camera. The green arrows point to the kinematic mounts that align the transfer plate relative to the turn-plate.



Figure 27: Tile vacuum positioning jig on the  $x-y-\phi$  stack, with a dummy tile in the centre and the three pillars that hold the turn-plate in the pick and place machine.

target position according to the tile to be aligned. The tile is placed on the tile positioning jig and the jig itself is installed on the turn-plate. The camera is used to properly align the two jigs. Finally, the tile is aligned to its target position by using the two outer ASIC markers and then transferred to the transfer plate. The procedure is repeated for the second tile, such that the transfer plate carries both tiles that need to be attached to one side of the module. The procedure is repeated for the other side of the module. The attachment, including the glue curing time (24 hours), is performed one side at the time.



Figure 28: The transfer plate, shown with two dummy sensor tiles attached to the vacuum chucks.

### 4.4 Hybrid placement

The procedure for attaching the front-end and control hybrids to a module is performed similarly to that of the tiles. However, the hybrid placement does not require the same level of precision, with an alignment of  $100 \,\mu\text{m}$  with respect to the tile being enough to ensure a successful wire-bonding. The choice of the glue is mainly dictated by the need for good radiation hardness and flexibility. The cooling requirements are not as critical as for the tiles, as discussed in Sect. 3.3.4. The glue is deposited with the same robot, but using a different syringe and needle. Since the cooling requirements are less stringent, the glue pattern does not need to maximise the area covered, but rather to provide mechanical stability during the wire-bonding process (for the FE hybrids) and a reliable adhesion to the microchannel substrate. For these reasons, the chosen pattern consists of a small number of lines: five evenly spaced lines for the front-end hybrids, with one of them directly underneath the bond pads, and nine lines for the GBTx hybrids. At Nikhef, a star pattern was implemented for the GBTx hybrid, in order to potentially provide a better heat transfer to the control chip. Since the thermal performance measurements did not show a significant difference between the star and the line patterns, the star pattern was not adopted by Manchester (see Fig. 29). The front-end hybrids are also plasma cleaned before attachment, in order to ensure good cleanliness of the bonding pads and of the pins in their connectors.

### 4.4.1 Procedure at Manchester

The front-end and GBTx hybrids are positioned in dedicated alignment frames, shown in Fig. 30, where they are mechanically aligned against a set of pins, which ensure the precision required of about  $100 \,\mu\text{m}$ . Once the alignment is completed, vacuum is activated to hold the hybrids in position.

The attachment of the hybrids to the module is performed following the same strategy used for the tiles. The hybrids for the connector side of the module are transferred to the corresponding transfer plate, shown in Fig. 23, and installed under the glue dispenser where the glue is deposited. This procedure is then repeated for the non-connector side hybrids and finally the six hybrids are attached to the module on both sides simultaneously. In Manchester, the interconnect and data cables are attached at a later stage.

Photographs of both sides of a module with hybrids attached are shown in Fig. 31. On



Figure 29: The glue pattern of the FE and GBTx hybrids as implemented in Manchester (left) and Nikhef (right).



Figure 30: Hybrids alignment jigs, for the connector side (left) and non-connector side (right).

the left hand side, the connector side of a module after hybrids' attachment in Manchester. On the right hand side, the non-connector side of a module after hybrids' and cables' attachment at Nikhef.

### 4.4.2 Procedure at Nikhef

The main difference in the procedure at Nikhef is that all the flat cables are connected to the hybrids before gluing. This was motivated by concerns that the relative rigidity of the flat cables could cause difficulties in their attachment, though in practice this did not prove to be an issue. Nevertheless, at Nikhef, the placement of the hybrids and cables is performed in a single step. A vacuum pickup tool is used to hold the hybrids, and is shown in Fig. 32 left. As the alignment of the hybrid placement is not as critical as for the placement of the tiles, a simplified alignment scheme is used without the need for an



Figure 31: Photographs of both sides of a module with hybrids attached. Left: connector side of a module after hybrids' attachment in Manchester. Right: non-connector side of a module after hybrids' attachment at Nikhef.

optical camera. A mechanical jig with alignment pins against which the FE hybrids can be positioned is used (Fig. 32 right).



Figure 32: Left: the hybrid transfer jig, showing the vacuum chucks for the two FE hybrids and the GBTx hybrid. Right: the alignment jig for the connected hybrids before being lifted by the transfer jig.

The GBTx hybrid is not strongly constrained, in order to prevent additional stresses in the interconnect cables. Once the hybrids are in place, they are held by vacuum and the transfer jig placed on top. The hybrids and flat cables assembly is picked up and placed under the same automated glue dispenser described previously. Glue is then deposited in a line (star) pattern on the backside of the front-end (GBTx) hybrids. To prevent large forces on the microchannel substrate, distance screws are adjusted when the hybrids are picked up. Corresponding spacers on the jig where the module is held during gluing prevent the hybrids from pressing on the substrate.

### 4.5 Wire-bonding and HV cables attachment

After they have been attached to a module, the VeloPix front-end ASICs are wire-bonded to their readout hybrids. In order to support the microchannel substrate during this process, a jig has been manufactured in each assembly site to hold the module from the opposite side, as shown in Fig. 33. The vacuum chucks, used to hold the module in place, are positioned such as to avoid contact with the sensors on the supporting side.



Figure 33: Wire-bonding support jig used in Manchester (left) and Nikhef (right).

The first step of the procedure consists of attaching the HV cable to one side of the module, as shown in Fig. 34 left. The glue is left to cure and then the module is installed in the wire-bonding jig and fixed in place under the wire-bonding machine. The machine is aligned to the module by using image recognition and the alignment validated. The automatic bonding program is executed and all 840 wire-bonds on that side of the module. Finally, the HV tape is manually wire-bonded to the sensor surface. The same procedure is then repeated for the other side of the module.

### 4.6 Cables attachment

The final step in the assembly consists of attaching all remaining cables to the module. In Manchester, this includes the interconnect cables, data cables and LV cables. At Nikhef, since the flat cables are attached together with the FE and GBTx hybrids, this step only concerns the attachment of the LV cables.

The flat cables are thin tape-like copper-polyimide cables, equipped with high density connectors. As for the front-end hybrids, these cables are also plasma cleaned to ensure optimal connectivity of the pins in the connectors. During attachment, the module is installed in a frame and held vertically on a stand, free to rotate around the vertical axis such that the optimal angle can be reached for each individual connection, as shown in Fig. 34.

The same frame and stand are used to hold the module during attachment of the LV harness. The transition bridges, fully equipped with LV cables, are placed on top of the two plastic pillars shown in Fig. 34, such that the right distance from the module foot is set. The two parts of the LV foot connector are screwed to the aluminium foot. Then, the transition bridges are glued to the carbon fibre legs and left to cure. Finally, the LV cable connectors are mated to the FE and GBTx hybrids and the pins are inserted into the LV foot connector. The finished module is shown in Fig. 35.



Figure 34: Left: drawing of a module mounted on a free-rotating stand for cables attachment. Right: a photograph of a module after HV cables attachment.



Figure 35: Photographs of the front and rear side of a fully assembled module.

### 5 Module assembly measurements

Each individual assembly step is followed by a corresponding measurement to assess the quality achieved and assign a grade to the step. These measurements are performed at both assembly sites with similar techniques. The results are analysed with the same software and graded upon the same set of criteria, and stored in the database (see Sect. 4.1). While the two assembly sites employed different equipment, the test procedures (and their outcome) are fairly comparable. We choose to describe here the procedures used in Manchester, and point out the differences, whenever relevant.

### 5.1 Bare module flatness

Before shipment to the assembly sites, a flatness measurement of the microchannel substrate is carried out at CERN, in order to confirm that no deformations were introduced during the fluidic connector soldering procedure. The survey is performed with a 3D optical profilometer<sup>12</sup>, a 3D optical system that uses light reflectometry to reach an accuracy of  $\pm 3 \,\mu m$  [13].

The flatness of the cooling substrate after being assembled into a bare module is measured again at the assembly sites to ensure that no damage has occurred during transport, storage and assembly of the substrate. The measurement is important also to verify the presence of any mechanical stress introduced in the silicon substrate by the assembly procedure, which might cause failures in the attachment of the tiles.

Here, the flatness is measured using a 3D optical system<sup>13</sup> that includes a camera and a laser. This device can record data with micrometer precision in the xy plane and along the z axis. It is also capable of automated feature-finding methods and can locate shapes according to their different contrasts. During the survey, the module is held by the turn-plate (as shown in Fig. 19) and placed under the optical system. The turn-plate carries markers on a glass surface, such that they are visible from either side of the plate. These markers are used to define a reference system with respect to the precision hole in the module foot. The survey is performed using using the laser of the optical system, which scans the surface of the microchannel plate along x and y while recording the z values. In Manchester, only the connector side substrate flatness is measured, while both sides are surveyed at Nikhef.

The results obtained from these surveys are compared with those from CERN. They are also used to generate a substrate flatness grade. An example of a comparison between a survey performed at CERN (left) and the one carried out at an assembly site (right) is shown in Fig. 36. As illustrated by this example, the measurements performed at the assembly site after bare module assembly and those carried out at CERN on the microchannel substrate, show a significant difference at the sides of the substrate. This disagreement is due to the fact that, in the assembly sites, the substrate on the bare module is held in place by the holding frames at two fixation points at its sides, whereas at CERN the substrate is freely placed on a flat surface. However, no issue due to this constraint has been observed during module assembly.

The two-dimensional scan of z measurements across the surface of the plate is taken and the data points acquired are used to perform a fit with a least-square plane. Finally,

 $<sup>^{12}</sup>$ Keyence model VR-3200.

<sup>&</sup>lt;sup>13</sup>OGP SmartScope Flash 500.



Figure 36: Comparison of a flatness measurement performed on an unconstrained cooling substrate (left) and as a bare module (right).

the distribution of variations between the measured points and the fitted plane is used to calculate the four quartiles. The quality criteria assigned are based on the value of the first and third quartiles, as indicated in Table 2. In general, grades are assigned as being A,B,C,D or F, although some grades are not used in specific cases.

Table 2: Quality criteria and assigned grading for the microchannel substrate flatness. In this test, the grade F is not used.

grade	absolute value of quartiles $Q_1$ and $Q_3$
Α	$\max( Q_1 ,  Q_3 ) < 50 \ \mu m$
В	50 $\mu m < max( Q_1 ,  Q_3 ) < 75 \mu m$
$\mathbf{C}$	75 $\mu m < max( Q_1 ,  Q_3 ) < 100 \mu m$
D	$\max( Q_1 ,  Q_3 ) > 100 \ \mu m$
$\mathbf{F}$	_

### 5.2 Tile metrology

The tiles represent the active part of a VELO module. Their precise positioning is required to avoid mechanical interference, while a precise prior knowledge of their relative positions is a crucial starting point for the alignment of the whole detector in LHCb. The glue layer underneath each tile is also vital for a good alignment in z and for ensuring an optimal flow of the heat produced by the ASICs. For these reasons, an accurate metrology of the tiles after attachment is performed.

#### 5.2.1 Tile position

The positions of the tiles are measured with the 3D optical system. As in the case of the bare module flatness measurement, the glass markers on the turn-plate are used to define a reference system relative to the dowel hole in the module foot. Using pattern recognition features, the left-most and right-most fiducials of the two tiles facing the camera are found and their (x,y) position values are recorded. The same technique is used to measure the fiducials on the back of the sensor attached to the opposite side of the module (visible because of its overhang). The procedure is then repeated for the other side of the module.
The data are used to calculate the middle point along the tile and its rotation with respect to its nominal orientation, as illustrated in Fig. 37. On the non-connector side of the module, where the two tiles are placed at only 145  $\mu$ m from each other, their relative distance is also measured. The quality criteria applied to assign a grade are listed in Tab. 3.



Figure 37: Left: diagram showing the middle point x-y offset and angle which are measured for every tile. Right: diagram showing the closest distance between the two tiles on the non-connector side of the module.

Table 3: Quality criteria and assigned grading for the tile positions. The grade is determined by the worst result of the four parameters.

grade	middle point position	angle	distance 1 and 2
Α	$< 55 \mu m$	$< 1000 \mu rad$	$> 130 \ \mu m$
В	$< 110 \ \mu m$	$< 2000 \mu rad$	$> 100 \ \mu m$
$\mathbf{C}$	$> 110 \ \mu m$	$> 2000 \mu rad$	$> 50 \ \mu m$
D	_	_	$< 50 \ \mu m$
$\mathbf{F}$	_	_	$< 10 \ \mu m$

## 5.3 Glue layer thickness and tile flatness

This measurement is performed using the laser of the 3D optical system, one tile at a time, and is made immediately after the tile position measurement. The glue layer thickness will determine both the adhesion quality and the thermal performance, while the tile flatness checks for any irregularities.

The module is in its turn-plate and the glass markers on the turn-plate are used to define the reference system. The laser scans an area covering the tile as well as the regions on the cooling substrate that surround the tile. In this way, the glue thickness and its variation can be calculated by comparing the measurements obtained on the tile and on the substrate. The scan over the tile surface also shows the variations in the height of the tile itself, and it is used as a measurement of the tile flatness. The difference of the two-dimensional scan of z measurements across the four tiles is determined with respect to a least-square fit with a plane. An example of the glue layer thickness and tile flatness measurements are shown in Fig. 38. The quality criteria used to assign a module grading to both measurements are listed in Table 4.



Figure 38: Example measurements of the glue layer thickness (left) and tile flatness (right). These heat maps show the variation of the z measurements relative to the cooling substrate.

Table 4: Quality criteria and assigned grading for the tile flatness and glue layer thickness. The grade is determined by the worst result of the two parameters

grade	tile flatness: $ mean +2$ ·stdev	glue layer thickness: mean $\pm 2$ ·stdev
Α	$< 30 \ \mu m$	$40 - 120 \ \mu m$
$\mathbf{B}$	$< 60 \ \mu m$	$20 - 160 \ \mu m$
$\mathbf{C}$	$< 120 \ \mu m$	$0-200$ $\mu m$
D	$> 120 \ \mu m$	$> 200 \ \mu m$
$\mathbf{F}$	_	_

# 5.4 Hybrid metrology

As mentioned in Sect. 4.4, the accuracy required in the positioning of the hybrids is mainly dictated by the wire-bonding process, which requires the bond pads on the hybrids to be well aligned to those on the ASICs, and to be at a consistent distance from each other over the full length of the tile. For this reason, only the front-end hybrids are checked. After their attachment, this alignment is verified by taking high magnification photographs of the region where hybrids and tiles are facing each other. These pictures also allow for spotting any glue spillage that might affect the quality of the wire-bonding. An example of these photographs is shown in Fig. 39.

### 5.5 Wire-bond pull test

After the wire-bonding process, the quality of the wire-bonds is assessed by measuring the strength of their attachment. This is a destructive test which is performed on wire-bonds



Figure 39: High magnification photographs showing the alignment at both corners of a tile relative to its front-end hybrid.

that serve no electrical purpose and are bonded solely to serve for this indication of the strength of the wire-bonding process in that particular module. Each ASIC contains four sacrificial wire-bonds, which are pulled by a dedicated device that records the breaking force. The device uses a small hook that is lowered next to the sacrificial wire-bond to catch the wire and pull it upwards until it breaks, as shown in Fig. 40. The breaking force is then recorded and the hook is moved to the next sacrificial wire-bond. After all sacrificial bonds have been pulled, the broken wires are removed. The quality criteria used to assign a grade to this step are summarised in Table 5.



Figure 40: Photograph showing the hook of the pull-testing device catching the sacrificial wire-bond to be pulled on.

Table 5: Quality criteria and assigned grading for the wire-bond pull force.

4 g
2 g

# 6 Module validation

Once the assembly process is completed, each module is validated through a sequence of mechanical, thermal and electrical tests, such that its overall quality is assessed. The test setup and the tests performed are described in the following sections, with the results obtained discussed in Sect. 7. While differing in some details, the setup at each of the two assembly sites provided information for all the essential tests. For brevity we choose to describe the Manchester setup here.

# 6.1 Test setup

Each assembly site is equipped with a test setup that replicates the running conditions of the detector in the final experiment. The setup consists of a vacuum tank that can accommodate a single module. It is connected to a  $CO_2$  cooling system and to the full electronics chain. A schematic drawing of the setup is shown in Fig. 41.



Figure 41: A schematic drawing of the test setup used at the assembly site for module validation.

The vacuum tank consists of two parts, that can be brought together, sealed with a rubber O-ring. One part is an empty cylinder that can move on two rails and it is attached to a vacuum pump<sup>14</sup> through a bellow. The cylinder is also equipped with two germanium windows to facilitate photography with an infra-red camera. The other part of the vacuum tank houses a number of dedicated vacuum feed-throughs, in order to feed LV and HV power inside the vacuum, to enable data transfer from the module to the data acquisition system, to allow for the coolant to be circulated inside the module and to control a number of thermal and mechanical probes. This part of the tank also houses a frame for the installation of the module. The pressure inside the tank is kept below  $10^{-3}$  mbar.

In the experiment, the VELO detectors are operated in vacuum and kept cold by circulating  $CO_2$  at a flow of 0.4 g/s and evaporated at a pressure of about 12 bar. The same conditions are required at the assembly sites, so that the modules can be evaluated in realistic running conditions. The assembly sites also use evaporative  $CO_2$  systems to cool down the modules.

In the module validation setups, the electrical chain consists of the same electronics system as in the final experiment, wherever possible. In particular, the vacuum feed-through board (VFB) and the opto- and power board (OPB) are the same components

 $<sup>^{14}\</sup>mathrm{Pfeiffer}$  Vacuum D-35614 Asslar Turbo Pump.

which are employed in the experiment, as described in Sect. 3.1. The module data cables are connected, via high speed tapes, to the vacuum side of the VFB, which then carries these signals out of the vacuum to the OPB. The OPB is then connected to the data acquisition system (MiniDAQ) via fibres. The HV is provided to the module directly via the VFB, while the LV is instead provided through the OPB, which also has a connection to read out the temperature probes (NTC) in the front-end and GBTx hybrids and on the OPB itself. These NTC sensors are continuously monitored, logged and connected to a hardware interlock system. The MiniDAQ [24] is a Linux server running CentOS 7, containing a special PCIe40 card specifically adapted for the assembly sites. This card differs from the readout PCIe40s installed in the experiment as it combines the functionalities otherwise performed by three separate cards for the slow control, the timing and fast control and the readout. In the experiment, these functions are carried out by the Experiment Control System (ECS), Time & Fast Controls (TFC) and TELL40 respectively [4]. The MiniDAQ also carries the drivers, the analysis software and the SCADA [25] projects required to run the experiment.

In Manchester, the validation setup is powered using two programmable power supplies<sup>15</sup>, each of which is a dual-channel system which can supply up to 188 W. Three channels are required to power the OPB and the VeloPix digital and analogue supplies. Each channel carries a voltage of 6V to power an array of DC-DC converters on the OPB, which then step down the voltage level to that required by the front-end components.

A VELO module requires four HV lines, one for each tile. In Manchester, this is provided by a 3kV/1mA multichannel high voltage module<sup>16</sup>, which delivers the required bias voltage to the module sensors during testing. However, these boards have a current resolution of about 50 nA, which is not accurate enough when measuring the I-V characteristics of a silicon sensor and for this reason, when performing I-V measurements, the HV cables are switched to a power supply<sup>17</sup> which can deliver up to  $\pm 1100$  V and has a current resolution of 1 nA.

After assembly, the module is still held in the frame used during cable attachment; this frame can also be attached to a mounting point inside the tank. The module is first positioned against the tank flange by aligning the LV foot connector to its counterpart inside the tank and then screwed to the mounting point. The two LV connectors are then mated. An aluminium bar housing two displacement sensors is attached to the module frame, such that the displacement sensors face the non-connector side of the module. Finally, the two HV cables, the four data tapes and the two cooling capillaries are mated to the corresponding long HV, long data cables and cooling pipes installed inside the tank. Fig. 42 shows a module after installation in the vacuum tank.

# 6.2 Mechanical tests

### 6.2.1 Module displacement

A VELO module is an assembly of various parts made of different materials (silicon, copper, kapton, adhesives,...) with different CTE. It is therefore expected that the module will distort when its temperature is changed from room to operating values. Thermal

 $<sup>^{15}\</sup>mathrm{Rohde}$  & Schwarz HMP2020.

<sup>&</sup>lt;sup>16</sup>CAEN AG536N board, installed in a CAEN SY4527LC crate.

<sup>&</sup>lt;sup>17</sup>Keithley model 2410.



Figure 42: Photograph of a module installed in the validation setup, with the vacuum chamber cap retracted on the left.

distortions can put a strain on the substrate, and represent a risk for the integrity of the module. Moreover, it is vital to ensure that these deformations do not significantly reduce the clearance between the modules and the RF foil. Finally, knowing the final position of each module at its operating temperature of -30°C is critical for the alignment of the detector during operation. In addition to thermal distortions, there are other factors that can lead to module displacements: the mating of the HV tapes and data cables during installation, the connection of the cooling capillaries to the tank cooling pipes and the evacuation of the tank. Displacements are measured at two different locations on the non-connector side of the substrate, as shown in Fig. 43.



Figure 43: The displacement sensors are positioned in front of the non-connector side of the module substrate, in the two locations marked as D1 and D4.

This displacement measurement is performed by using two capacitive sensors<sup>18</sup>, housed

<sup>&</sup>lt;sup>18</sup>Micro-Epsilon CSH1-CAm1,4.

in a dedicated aluminium holder that is directly attached to the frame holding the module. These capacitive sensors have a range of about 1.4 mm and therefore are placed at about 500  $\mu$ m from the silicon sensor surface. The displacement measurements are continuously recorded during the tightening of the cooling connectors, the evacuation of the vacuum tank, and the cooling down from room temperature to -30 °C. An example of results from these measurements for a single module are shown in Fig. 44. The displacements from cooling pipe connection and from evacuation are comparatively small, while the displacement under cool down is larger and clearly correlated with the temperature of the four sensors (CLI,CSO,NLO,NSI). The displacement of each module is measured and a grade assigned according to the criteria in Tab. 6.



Figure 44: Displacements of the module measured with the two displacement sensors (D1 and D4). Top: displacement versus time during evacuation of the vacuum tank from atmospheric pressure to below  $10^{-3}$  mbar. Middle: displacement versus time when cooling down the modules. Lower: temperature versus time when cooling down the modules, for the same procedure and with the same abscissa as the middle plot.

grade	cooling pipes connection	pump-down	cool-down
$\mathbf{A}$	$< 50 \ \mu m$	$< 50 \ \mu m$	$< 100 \ \mu m$
В	$< 80 \ \mu m$	$< 80 \ \mu { m m}$	$< 150 \ \mu { m m}$
$\mathbf{C}$	$< 100 \ \mu { m m}$	$< 100 \ \mu { m m}$	$< 200 \ \mu { m m}$
D	$> 100 \ \mu m$	$> 100 \ \mu m$	$> 200 \ \mu m$
$\mathbf{F}$	—	—	—

Table 6: Quality criteria and assigned grading for the module displacement tests.

#### 6.2.2 Thermal cycles

Although during normal operation of the LHCb experiment the modules are kept cold to reduce radiation damage, it is expected that they will undergo several tens of thermal cycles between -30 and +20 °C during LHC Run 3 and 4. All modules have therefore to be able to withstand the effects of a large number of  $\Delta T=50$  °C cycles, without any significant impact on their performance. For this purpose, a set of ten thermal cycles is performed for every assembled module. All electrical tests are carried out twice, before and after thermal cycling, such that any degradation of the module performance is detected. The cycling process consists of cooling down the module to -30 °C, then switching the cooling off and allowing the system to warm up back to room temperature. Each cycle takes about 30 minutes. An example of a set of ten thermal cycles, together with the displacements observed in the module, is shown in Fig. 45 where the displacements observed are seen to be highly reproduceable between cycles.



Figure 45: Displacement (top) and temperature (bottom) of the four tiles versus time during the thermal cycling step.

# 6.3 Electrical tests

Immediately after installation of a module in the vacuum tank, a quick connection test is performed before closing the tank and starting the evacuation. This quick test checks the basic functionalities of the ASICs and ensures that all on-module and off-module cables have been connected properly. Afterwards, the vacuum tank is evacuated to a pressure below  $10^{-3}$  mbar and the CO<sub>2</sub> coolant circulated in the module to bring its temperature down to -30° C. The full chain of electrical tests is then carried out, as detailed in the following sections. The main purpose of the electrical tests is to assess the quality of the chips (VeloPix chips and GBTx's) and of the silicon sensors on the final module, but they are also useful to detect any problems with the interconnecting and data cables or with the wire-bonds between the tiles and the front-end hybrids. At the end of the first run of electrical tests, the thermal cycles are performed and then a second run of electrical tests is carried out. The results of the two runs of tests are then compared. The electrical tests are controlled using the SCADA [25].

### 6.3.1 IV curves

The I–V characteristics of each sensor tile, at room temperature, is measured by the tile manufacturer and again at CERN on delivery in order to assess the quality of the tiles before shipment to the assembly sites. The same characterization is then repeated for each module after construction, in order to ensure that no degradation has occurred during assembly. The voltage scan is performed with automated software, customized independently by the two assembly sites, in the range 0-1000 V, in steps of 5 V and with a ramp-up speed of 1 V/s. The maximum operating voltage of the modules is 1000 V. This voltage may be needed to be applied for some tiles at the end of life in the experiment because of radiation damage. At each step of the scan, a waiting time of 10 s is set before the software acquires five consecutive current measurements and outputs their average. An example of the I-V characteristics of a tile, as measured by the tile manufacturer, tile quality assurance site and module assembly site, is shown in Fig. 46. As the test setups and environmental conditions are not the same in the different tests a precise agreement in results is not expected.

In addition to this test, a ramp-up test at the nominal cooling conditions (-30° C) is performed, in order to ensure that every non-irradiated tile can reach the minimum bias voltage of 140 V, while drawing a current lower than 1  $\mu$ A.

The grading criteria for both the IV and HV tests are listed in Tab. 7. The criteria for the IV measurements are based on the measurements performed in a probe station after reception from the manufacturer, as explained in Sect. 7.2.10.

	grade	IV curve	lowest voltage before 1 $\mu {\rm A}$
-	Α	< high quality	> 700  V
	В	< good quality	< 700  V
	$\mathbf{C}$	< good quality below 140 V	_
	D	> good quality below 140 V	—
	$\mathbf{F}$	_	$< 250 {\rm ~V}$

Table 7: Quality criteria and assigned grading for the IV test and HV ramp-up test.



Figure 46: IV curves of an example VELO tile measured after bump-bonding at Advacam, for tile quality assurance at CERN, and after module assembly at Manchester. All were measured at around room temperature.

### 6.3.2 VeloPix communication test

The VeloPix chips receive their LV power through dedicated power line wire-bonds. The power is supplied by 1.3 V DC-DC converters [26], but due to losses over the OPB traces, VFB connector and LV cables, the voltage drops to a typical value of about 1.2 V at the VeloPix chips. Therefore, the purpose of this test is to ensure that communication can be established with all the chips of a module, and that all chips are supplied with the required minimum power (1.18 V). The VeloPix chips are equipped with an internal monitoring of their analogue and digital voltages, and temperature. They can also report their Chip ID, which is burnt to an eFuse on the chip itself. The test is carried out by reading back these values from each ASIC, thus verifying the integrity of the electrical chain from the VeloPix through the wire-bonds, the front-end hybrids, the interconnecting cables, the GBTx chip and the long data tapes. Typical values of the supplied voltages and internal temperatures of the VeloPix chips are shown in Fig. 47 for one module.

VeloPix	CLI-0/VP0-0	CLI-1/VP0-1	CLI-2/VP0-2	CSO-0/VP3-0	CSO-1/VP3-1	CSO-2/VP3-2
Status	Minimal/QA	Minimal/QA	Minimal/QA	Minimal/QA	Minimal/QA	Minimal/QA
Chip ID	00000D3B	00000D34	00000D35	00000D5C	00000D66	00000D48
Analogue	1.2432 Volt	1.2403 Volt	1.2447 Volt	1.2579 Volt	1.2593 Volt	1.2615 Volt
Digital	1.2571 Volt	1.2601 Volt	1.2542 Volt	1.2659 Volt	1.2703 Volt	1.2747 Volt
Band Gap	0.2803 Volt	0.2767 Volt	0.2784 Volt	0.2486 Volt	0.2513 Volt	0.2908 Volt
Temperature	-46.82 degree	-48.90 degree	-47.74 degree	-61.81 degree	-61.23 degree	-41.17 degree
Back Hybrid	1 —— I					
Back Hybrid VeloPix	NLO-0/VP1-0	NLO-1/VP1-1	NLO-2/VP1-2	NSI-0/VP2-0	NSI-1/VP2-1	NSI-2/VP2-2
Back Hybrid VeloPix Status	NLO-0/VP1-0 Minimal/QA	NLO-1/VP1-1 Minimal/QA	NLO-2/VP1-2 Minimal/QA	NSI-0/VP2-0 Minimal/QA	NSI-1/VP2-1 Minimal/QA	NSI-2/VP2-2 Minimal/QA
Back Hybrid VeloPix Status Chip ID	NLO-0/VP1-0 Minimal/QA 08001462	NLO-1/VP1-1 Minimal/QA 00000232	NLO-2/VP1-2 Minimal/QA 0000023B	NSI-0/VP2-0 Minimal/QA 0000145C	NSI-1/VP2-1 Minimal/QA 00001453	NSI-2/VP2-2 Minimal/QA 00001471
Back Hybrid VeloPix Status Chip ID Analogue	NLO-0/VP1-0 Minimal/QA 08001462 1.2601 Volt	NLO-1/VP1-1 Minimal/QA 00000232 1.2667 Volt	NLO-2/VP1-2 Minimal/QA 0000023B 1.2608 Volt	NSI-0/VP2-0 Minimal/QA 0000145C 1.2652 Volt	NSI-1/VP2-1 Minimal/QA 00001453 1.2652 Volt	NSI-2/VP2-2 Minimal/QA 00001471 1.2681 Volt
Back Hybrid VeloPix Status Chip ID Analogue Digital	NLO-0/VP1-0 Minimal/QA 08001462 1.2601 Volt 1.2637 Volt	NLO-1/VP1-1 Minimal/QA 00000232 1.2667 Volt 1.2674 Volt	NLO-2/VP1-2 Minimal/QA 0000023B 1.2608 Volt 1.2747 Volt	NSI-0/VP2-0 Minimal/QA 0000145C 1.2652 Volt 1.2740 Volt	NSI-1/VP2-1 Minimal/QA 00001453 1.2652 Volt 1.2703 Volt	NSI-2/VP2-2 Minimal/QA 00001471 1.2681 Volt 1.2755 Volt
Back Hybrid VeloPix Status Chip ID Analogue Digital Band Gap	NLO-0/VP1-0 Minimal/QA 08001462 1.2601 Volt 1.2637 Volt 0.3263 Volt	NLO-1/VP1-1 Minimal/QA 00000232 1.2667 Volt 1.2674 Volt 0.2791 Volt	NLO-2/VP1-2 Minimal/QA 0000023B 1.2608 Volt 1.2747 Volt 0.2821 Volt	NSI-0/VP2-0 Minimal/QA 0000145C 1.2652 Volt 1.2740 Volt 0.2940 Volt	NSI-1/VP2-1 Minimal/QA 00001453 1.2652 Volt 1.2703 Volt 0.2945 Volt	NSI-2/VP2-2 Minimal/QA 00001471 1.2681 Volt 1.2755 Volt 0.2635 Volt

Figure 47: Typical values of the VeloPix chips internal analogue and digital voltages and uncalibrated temperatures.

### 6.3.3 Time and fast control test

The TFC system [27] is responsible for the proper synchronization of the VELO modules with the rest of the experiment. The TFC signals reach the VeloPix chip through a single differential pair of wire-bonds. Any issue with this pair would result in the module being unusable by the experiment, as it would be operating asynchronously to the rest of the detector. The TFC signals are sent out from the PCIe40 card [28] to the GBTx chips on the module, and from there they are distributed to all VeloPix chips, where counters are incremented according to the signals received. A dedicated WinCC panel (see Fig. 48) reads these counters from a set of specific registers and displays the results for each individual VeloPix, as well as returning a green (red) flag for successful (unsuccessful) tests.

The gradings based on the internal digital and analog voltages and the TFC test are assigned according to the criteria listed in Tab. 8.

VeloPi	x TFC	Respo	nse										
Front Hybr	rid												TFC Actions
TFC Counter	VP0-0 Cnt	VP0-0 Snap	VP0-1 Cnt	VP0-1 Snap	VP0-2 Cnt	VP0-2 Snap	VP3-0 Cnt	VP3-0 Snap	VP3-1 Cnt	VP3-1 Snap	VP3-2 Cnt	VP3-2 Snap	SODIN GBTtest.Cor
BXID Reset	1315105	1311596	1294614	1285972	1315115	1311602	1318019	1314504	1294654	1286026	1318023	1314517	
FE Reset	1	1	1	1	1	1	1	1	1	. 1	1	1	Configure SODIN
Calib A	2	2	2	2	2	2	2	2	2	2	2	2	
Snapshot	93	92	91	90	93	92	93	92	91	. 90	93	92	Send Sync
SYNC	10	10	10	10	10	10	10	10	10	10	10	10	
	Reset		Reset		Reset		Reset		Reset		Reset		Send FE Reset
Back Hybr	id												Run x Orbits 2
TFC Counter	VP1-0 Cnt	VP1-0 Snap	VP1-1 Cnt	VP1-1 Snap	VP1-2 Cnt	VP1-2 Snap	VP2-0 Cnt	VP2-0 Snap	VP2-1 Cnt	VP2-1 Snap	VP2-2 Cnt	VP2-2 Snap	Stop Run
BXID Beset	1307558	1271755	1294616	1285979	1307607	1300273	1307608	1300277	1294607	1285964	1318006	1314497	
	1	1	1	1	1	1	1	1	1	1	1	1	-VoloBix Actions
FE Reset		2	2	2	2	2	2	2	2	2	2	2	VeloPix Actions
FE Reset Calib A	2		91	90	92	91	92	91	91	90	93	92	Reset All Counters
FE Reset Calib A Snapshot	2	89	21			10	10	10	10	10	10	10	
FE Reset Calib A Snapshot SYNC	2 90 10	89 10	10	10	10								

Figure 48: The WinCC panel used to carry out TFC tests. Signals are sent to the VeloPix chips and specific registers are then read back to verify that the increase in the corresponding counters matches the expectation.

Table 8: Quality criteria and assigned grading for electrical communication tests.

grade	analog and digital voltage	TFC result
Α	> 1.20 V	all signals seen
В	> 1.18 V	—
$\mathbf{C}$	$< 1.18 { m V}$	—
D	one missing	one missing
$\mathbf{F}$	multiple missing	multiple missing

### 6.3.4 Pseudo random bit stream test

The data stream coming from the VeloPix chips is routed to the DAQ directly via dedicated links in the hybrid data cables, without passing through the GBTx chip. Any error in the stream can lead to corrupted data packets and affect the performance of the final experiment. In order to verify the integrity of this path, a bit error rate test is carried out by using an industrial standard, the Pseudo-Random Bit Stream (PRBS). A stream is generated by the VeloPix using logic within the chip, and then propagated down the data links to the MiniDAQ, where it is collected and decoded. The result is then compared to the known pattern of the type of PRBS chosen, and the number of bits not matching the expected pattern recorded. This allows to assess the overall quality of the path carrying data signals from the point of PRBS generation on chip to the long data cables. However, this test is not sensitive to issues in the bump-bonds or with the silicon sensor itself, and cannot detect any damage in the pixel matrix. A typical result of a PRBS test is shown in Fig. 49.

Bi	t Error	Rate	Test				
Qua	artus Transo	eiver To	olkit ———				
Firm	ware sof-file	re/lhcb d	dag firmware md	2 v20200303-1.sof			
Outp	ut file	home/velo/modules/MOD 119/prbs_preTC tyt					
Rits I	tested	nonic/ve	w 100 12	115/pr05_prefet.txt			
Dics	lested	1	x 10 12				
N	Measure BER	]		Transceiver Toolkit			
Link	Bits Tested		Errors	BER			
0	1016303345	5664	0	0			
1	1016349057	7024	0	0			
2	1016394014	1720	0	0			
3	1016434212	2864	0	0			
4	1016477196288		3	2.9514E-12			
5	1016516435	5968	8	7.8700E-12			
6	1016557723	3648	0	0			
7	1016601411	1584	0	0			
8	1016645206	5016	0	0			
9	1016685314	1048	0	0			
10	1016694341	L632	0	0			
11	1016735948	3800	0	0			
12	1016777105	5408	0	0			
13	1016812232	2704	0	0			
14	1016855363	3584	0	0			
15	1016893497	7344	0	0			
16	1016934211	1584	0	0			
17	1016976326	5656	36	3.5399E-11			
18	1017017188	3352	0	0			
10	1017058009	9088	0	0			

Figure 49: Typical result of a PRBS test.

The PRBS test returns the number of bits tested and the corresponding number of errors for each of the twenty links in a module. The database uses this data to calculate the bit error rate (BER), which is then used to assign a grade according to the criteria summarised in Table 9.

grade	BER of all links	BER of more than half of the links
$\mathbf{A}$	$< 1 \times 10^{-11}$	$< 1 \times 10^{-12}$
В	$< 1 \times 10^{-10}$	$< 1 \times 10^{-11}$
$\mathbf{C}$	$> 1 \times 10^{-10}$	$> 1 \times 10^{-11}$
D	one link down	_
$\mathbf{F}$	multiple links down	_

Table 9: Quality criteria and assigned grading for PRBS test.

#### 6.3.5 Equalisation and Noise Scan

The noise level in the ASICs is a critical parameter that affects the performance of the detector and thus needs to be characterised. For reference, a MIP in a 200  $\mu$ m thick silicon sensor generates about 16,000 electron-hole pairs, which in the VeloPix correspond to about 1070 DAC counts (~15 e-h pairs per DAC count [29]). Moreover, because of imperfections in the silicon of the sensor and of the ASIC, each pixel may respond slightly differently to a given deposited energy and threshold. In order to correct for these differences, each pixel in the VeloPix is equipped with two configurable threshold values: the threshold value and the trim value. The threshold value is a 14-bit global threshold that is applied to all pixels on the same chip, while the trim is a 4-bit offset that is used to tune the differences between pixels in the same chip. Therefore, to equalise the response of the whole matrix, a threshold scan is performed and used to set a specific threshold trim value for each pixel. These scans are carried out automatically by a dedicated WinCC panel, which records the response of each pixel in the matrix to the minimal (trim=0) and maximal (trim=F) configurations. A typical result from these scans is shown in Fig. 50.



Figure 50: An equalisation scan showing the result for trim-0 (red), trim-F (blue) and the global equalisation (black).

The 0-trim scan is performed by setting the pixel trim value to the lowest of its 16

settings and then varying the global threshold value in a predefined range. At each step, the number of pixels that go from not firing to firing is counted, resulting in the red distribution shown in Fig. 50. This procedure is then repeated with the trim value set to its highest value of 15 (trim-F), leading to the blue curve shown in Fig. 50. Finally, the midpoint of the two trims is used to build the global equalisation threshold that is shown in black in Fig. 50. During these scans, the noise of each pixel is also measured and a noise distribution obtained for each pixel. The sigma of the distribution is then used to build a 2D noise map of each VeloPix. These maps are important not only to assess the noise level of the chip, but also to highlight any issue that might have been introduced during tiles attachment to the substrate. An example of a 2D noise map for trim-0 and trim-F is shown in Fig. Fig. 51. Finally, the average of all the values in the trim-F noise map, and their standard deviation, is used to assigned an average noise value and uncertainty to each VeloPix.



Figure 51: 2D noise maps, in DAC counts, evaluated at trim-0 and trim-F; the average noise and its uncertainty is also shown in the tile of the plot.

The noise and equalisation scans produce a large amount of data (about 15 MB/chip), specific for each VeloPix chip. The grading criteria are based on the number of masked pixels per ASIC and per tile, and on the average noise and its standard deviation, as summarised in Table 10.

Table 10: Quality criteria and assigned grading for the equalisation and noise scan.

grade	tile masked	ASIC masked	noise	noise stdev
Α	< 1%	< 5%	—	—
В	< 5%	—	_	_
$\mathbf{C}$	> 5%	—	_	_
D	—	_	> 8	> 2
$\mathbf{F}$	—	_	_	_
				1

### 6.4 Thermal performance

Assessing the cooling performance of a module is important to ensure that the heat transfer capabilities from the tiles meet the requirements discussed in Sect. 2. The cooling performance is measured using the VeloPix band-gap circuit. This circuit generates a reference voltage of around 300 mV on the periphery of the analogue read-out, independently of temperature changes. Therefore, a proportional-to-absolute-temperature (PTAT) circuit is used to monitor the temperature changes on the chip and to feed them back into the band-gap circuit, allowing the generation of the temperature-independent reference voltage. These two voltages are written into two registers of the VeloPix, labelled as band-gap and temperature, and can be used to monitor the variation of the internal temperature. In parallel, the NTC temperature sensor values are also logged and used as an independent measurement of the temperature observed near the tiles.

The cooling performance of a tile is thus determined by varying the power consumption in the three ASICs, measuring the corresponding temperature variations, and extrapolating the tile  $\Delta T$  to the end-of-life power consumption (26 W). The variation of the power consumption is achieved by enabling different functionalities in the chip (power modes). An example of the results obtained from the cooling performance test is shown in Fig. 52. The NTC temperature is typically lower than that on the ASICs due to the thermal gradient between the ASICs and hybrid.



Figure 52: Temperature variation in a tile as function of the power consumption in its three ASICs. The dashed line at 26 W indicates the estimated end-of-life power consumption.

The measurements are repeated before and after ten thermal cycles to identify any potential thermal degradation. The quality criteria applied and grading assigned are listed in Tab. 11.

Table 11: Quality criteria and assigned grading for the band-gap test and thermal performance change before and after thermal cycles are performed. The mean and maximum variations are computed over all the 12 ASICs of a module.

grade	$\Delta T$ at 26 W	mean variation	max variation
A	< 8 °C	< 1 °C	_
В	$< 10 \ ^{\circ}\mathrm{C}$	$> 1 \ ^{\circ}\mathrm{C}$	—
$\mathbf{C}$	> 10 °C	_	_
D	missing ASIC	_	>5 °C
$\mathbf{F}$	_	> 5 °C	_
	1	1	

49

# 7 Overview of module production

# 7.1 Production rate

The module production started in July 2020 in the two assembly sites, the University of Manchester and Nikhef. The production thus occurred during the period of the COVID-19 pandemic which led to significant additional organisational and logistical challenges. Despite these difficulties, the production was completed in time to allow installation for the start of LHC Run 3. The rate at which modules have been assembled and qualified is shown in Fig. 53. Only modules that have passed all quality criteria and were declared of production quality are shown in this plot. After an initial learning curve of about six months, the production rate increased to about 1.5 modules per week. However, in August 2021 the rate slowed down again, mainly due to delays in the production of a new batch of front-end and control hybrids. The target number of required modules (52) was achieved in December 2021.



Figure 53: Module production rate (only detector quality modules are shown here). The dotted grey line indicates the target of 52 detector quality modules needed for the VELO Upgrade.

In Fig. 54, the number of days needed to fully assemble and qualify each module is shown. This plot also illustrates how the amount of time spent on each module decreased over time, and reached a minimum of about two weeks per module during the central phase of production, when all procedures were well established and production was running smoothly. This completion time then increased again towards the end of the production, for the reason explained above. However, a number of steps in the module assembly and qualification could be carried in parallel, thus reducing the effective time per module: the production of the modules shown in the plot took approximately 16 months (493 days), which means that, in average, a new module was ready for the experiment every 9.5 days.



Figure 54: The time taken to fully assemble and perform the qualification tests on each module at one of the production sites. The days elapsed from the start to the completion are given for each module are shown on the y-axis. The x-axis indicates the start date of the module assembly.

# 7.2 Summary plots

During module production, the results of each individual test performed on a module were uploaded to the database, where the data were automatically processed and a grade assigned to the test. The analysis also allowed us to monitor any drift in the measured quantities such that any required action could be taken immediately to rectify the issue. A collection of summary plots for all the measured quantities is presented in the following sections, where these are shown for consistency for all modules produced at a single production site and ordered by date of completion. In addition some numbers characterising the overall module assembly accuracy and performance are given, where these are calculated for the set of 52 modules installed into the experiment.

### 7.2.1 Summary of bare module flatness measurements

The flatness of a bare module is measured as described in Sect. 5.1. The results are shown in Fig. 36. Given the large area that needs to be covered during the measurements, a number of outlying data points are present in each data set due to failures in the focusing of the optical devices. Thus, for the comparison between different modules, the  $\Delta z = z_{\text{meas}} - z_{\text{fit}}$  residual distributions are split into four ranges (quartiles), each containing 25% of the data points. The first three quartiles are then used to monitor the trend of the bare module flatness measurements, as shown in Fig. 55, while the fourth quartile is more affected by the mismeasured points. As expected, the median of the flatness distribution  $(Q_2)$  lies around  $\Delta z = 0$ , while quartiles  $Q_1$  and  $Q_3$  are fairly consistent along the whole module production. This test demonstrates that none of the substrates were damaged during transportation and that there were no irregular surface features that would prevent the attachment of the tiles.



Figure 55: The bare module flatness results as a function of module number across the production, showing the quartiles of the height distribution after the subtraction of a fit of a plane ( $\Delta z = z_{meas} - z_{fit}$ ).

#### 7.2.2 Summary of tile position measurements

As described in Sect. 5.2.1, the position of a tile after attachment is evaluated by measuring the x and y coordinates of the left fiducial in the leftmost ASIC and of the right fiducial of the rightmost ASIC. In addition, also the fiducials visible on the back of the over-hanging tiles are measured. These values can then be compared to the target position and the absolute differences used to assess the quality of the tile placement. This approach treats the measurements as independent while the tile is a rigid object. In order to better estimate the placement, the x and y coordinates of the middle point between the measured fiducials, together with the angle of rotation of the tile, is considered. The trends of the x and y offsets and the angle are shown in Fig. 56 for the four tiles. A 20  $\mu$ m drift from module 93 until 107 can be observed in the trend of the x coordinate for CLI, moving from around 20  $\mu$ m to 40  $\mu$ m. Having this careful monitoring procedure in place allowed this to be identified during production. This drift was corrected back to below 20  $\mu$ m by adjusting the alignment constants used for the positioning of the tiles from module 108 onwards.

The mean offset of the tiles from their ideal position is computed by taking, for all tiles on all installed modules, the average of the absolute distance for the centre of the tiles from its ideal position. This yields 21  $\mu$ m with a standard deviation of 10  $\mu$ m. The mean angular deviation of the tiles is 4  $\mu$ rad with a standard deviation of 243  $\mu$ rad. The tile positioning is thus far better than is needed from mechanical constraints and provides an excellent starting point for the software alignment of the detector with tracks.

On the non-connector side of a module, the distances d1 and d2 between the two tiles are also measured. The trend for those quantity is shown in Fig. 57. The distance is intended to be 145  $\mu$ m and is a key parameter in the assembly. As shown in the plot values close to this were achieved throughout the production.

#### 7.2.3 Summary of glue layer thickness and tile flatness measurements

The glue layer thickness and the tile flatness are measured using deviations from a plane fitted to measured data points. The same data set is used to both evaluate the glue thickness and the tile flatness, as discussed in Sect. 5.3. The mean and standard



Figure 56: The achieved tile positions relative to their ideal location as a function of module number across the production. The plots are for the four tiles on each module, arranged from top to bottom as CLI, CSO, NLO, NSI. In each plot the offset of the middle point of the tile in x and y and the rotation of the tile is shown.



Figure 57: The distances d1 and d2 between the two tiles on the non-connector side as a function of module number across the production.

deviation of the distributions obtained for each tile, together with those from the flatness measurements, are used to compare the quality of the attachment of the tiles during module production. The trends of these are shown in Fig. 58 for each of the four tiles as a function of the module number. Module 97 shows a larger than typical glue thickness. The hypothesis is that a dust particle was deposited along with the glue in the CSO tile which caused the microchannel to deform in the jig and led to variations in the glue thickness under all tiles. This leads to a somewhat worse thermal performance in the CSO tile (see Fig 65) but still within the acceptable range.

The mean deviation of the tile surfaces from a fitted plane for all points measured across all tiles on all installed modules is 0  $\mu$ m with a standard deviation of 4  $\mu$ m, thus showing that no significant bowing was introduced during the tile assembly or attachment. The glue layer thickness is an important parameter as it determines both the mechanical integrity and thermal performance of the module. Given the complex glue patterns, required treatment and evolving viscosity, the achieved average glue thickness of 78  $\mu$ m, with 5  $\mu$ m standard deviation, can be considered to be in excellent agreement with the target of 80  $\mu$ m.

### 7.2.4 Summary of wire-bond pull tests

As described in Sect. 5.5, after a module has been fully wire-bonded, its four sacrificial bonds are pulled and the breaking force recorded. The results for all modules are shown in Fig. 59. As the VELO modules are double-sided, obtaining sufficient support for the bonding was a concern during the design and R&D. However, the bonding quality met the specifications throughout the production without any trends in bonding performance observed.

#### 7.2.5 Summary of displacement measurements

As discussed in Sect. 6.2.1, the displacement of each module is evaluated during tightening of the cooling connectors, the evacuation of the test tank and the cooling down from room temperature to -30° C. The trends are shown in Fig. 60. While the first two displacements are usually small and fairly consistent across the whole lot of modules, the cooling down process induces larger deformations that can also vary significantly from one module



Figure 58: The glue thickness and tile flatness measurements are shown as a function of module number across the production. The points show the mean value while the error bars indicate the standard deviation. The plots are for the four tiles on each module, arranged from top to bottom as CLI, CSO, NLO, NSI. In each plot the mean glue thickness for each of the four ASICs on the module is shown (left axis) and the mean tile flatness (right axis).



Figure 59: The measured wire-bond strength, determined from pull tests on bonds on the production modules, as a function of module number across the production. Less measurements were made directly after the wire bonding on the first production modules.

to the other. Nonetheless, the displacement during cooling down is always in the same direction and very consistent when repeated on the same module. These results are a useful input to the software alignment that is performed using track reconstruction, and can provide a starting point, especially for the z position of the installed modules.



Figure 60: Displacements measured during cooling down the module (top) and evacuating the test tank (bottom), as a function of module number across the production. The results for both displacement sensors (see text) are shown.

## 7.2.6 Summary of communication tests

The purpose of measuring the analogue and digital current in each VeloPix chip is discussed in Sect. 6.3.2. These measurements are performed three times, first in air at room temperature, and then in vacuum at -30° C before and after thermal cycling. The results for the measurements taken before thermal cycling are shown in Fig. 61. For a small number of tiles the link reading the ASIC voltage was not functional, although the voltage was correctly delivered to the ASIC. It was noted that the input voltage to the VeloPix chips increases at lower temperatures, as expected. No significant effect due to the thermal cycling was observed.

## 7.2.7 Summary of PRBS tests

The PRBS test is carried out before and after thermal cycling, as all the other electrical tests. The results for all the modules are shown in Fig. 62. The majority of the links do not show any error in the  $10^{12}$  bits tested, while some specific links are noisier irrespective of the module under test. The issue seems to be due to the specific electrical test setup in which the measurement was performed.

## 7.2.8 Summary of noise tests

The noise and equalisation scans are discussed in Sect. 6.3.5. The average noise is an important parameter for the performance. For reference, a MIP in a  $200\mu$ m thick silicon sensor generates about 16,000 electron-hole pairs. The average noise in each of the three ASICs in a tile are shown in Fig. 63, as function of the module number. As illustrated by the trends, the typical average noise in a VeloPix is just below 6 DAC counts, which corresponds to about 90 e-h pairs. While this excellent performance gives a very high signal:noise ratio at the start of operations, the level of noise is expected to increase with radiation damage. The fraction of masked pixels is also shown in the plots. The fraction of masked pixels is observed after the thermal the experiment is 0.5%. No appreciable change in the results is observed after the thermal cycling.

### 7.2.9 Summary of cooling performance tests

The importance of measuring the cooling performance of each module is discussed in Sect. 6.4. The main parameters affecting the thermal behaviour of a tile are the fluidic characteristics of the microchannel substrate and the pressure, flow and temperature of the coolant. Fig. 52 shows an example of the thermal performance of a module and the extrapolation at 26 W is used to evaluate the cooling performance of each VeloPix. In order to compare the behaviour of different modules at slightly different running conditions, a correction is applied to the measured temperatures. The change in temperature of the ASIC between 0 W and that expected at 26 W is taken. This is then corrected by subtracting the temperature value of the closest NTC sensor, which corresponds to NTC<sub>CLI</sub> for tiles CLI and NLO, and to NTC<sub>NSI</sub> for tiles CSO and NSI, to make a  $\Delta$ T result that is more independent of the operational conditions of the cooling system. The location of these NTC sensors on the front-end hybrids is shown in Fig. 64.



Figure 61: The analogue and digital currents, measured by the VeloPix chips, as a function of module number across the production. The plots are for the four tiles on each module, arranged from top to bottom as CLI, CSO, NLO, NSI. In each plot the analogue and digital voltages of each of the three ASICs is shown.



Figure 62: The bit error rates (BER) during PRBS tests performed before (top) and after (bottom) thermal cycling, as a function of module number across the production. The empty bins represent links that had no errors in the  $10^{12}$  bits tested. BER values above  $10^{-7}$  are not shown in order to enhance the granularity.

The corrected cooling performance results for the modules are shown in Fig. 65. The larger temperature gradient measured on ASIC 1 of tile CLI in module 95 is due to a malfunctioning ASIC which provides a wrong reading of its bandgap voltage. The worse cooling performance shown by module 97 is instead due to its larger glue layer thickness as discussed above in Section 7.2.3.

As for all the other qualification tests, the thermal performance of a module is also evaluated twice, before and after the thermal cycling. The comparison between the two measurements provides an important check for any potential degradation of the glue layer due to thermally-induced mechanical stress, which may impact the cooling performance. The change of the thermal performance ( $\Delta\Delta T$ ) between the two tests for the modules is shown in Fig. 66. No significant change in the performance is observed.

The cooling performance is also affected by the glue layer thickness underneath the tile. Fig. 67 shows this correlation for all the modules. As expected, a thinner glue layer typically leads to a better cooling performance. The adjustment of the assembly setup was such that tile CSO typically had a thinner glue layer than the other tiles. In addition, the change of the cooling performance before and after thermal cycling is shown in Fig. 68, showing no dependence on the glue layer thickness. The missing data in the CSO trend (module 96) and NSI trend (module 101) are due to a damaged trace in the monitoring circuit of the middle ASIC, which prevents all the internal readings to reach the corresponding line in the front-end hybrid. The effectiveness of a cooling design can be quantified by the thermal figure of merit (TFM), which is defined as the temperature



Figure 63: The average noise (in DAC counts, left axis) and the percentage of masked pixels (right axis), as a function of module number across the production. The plots are for the four tiles on each module, arranged from top to bottom as CLI, CSO, NLO, NSI. In each plot the values for each of the three ASICs are shown.



Figure 64: Location of the NTC sensors used to correct the temperature gradient measured in each tile.

gradient between the hottest point on the surface of the silicon sensor and the coolant, divided by the thermal density power dissipated by the ASICs and sensors:

$$TFM = \frac{\Delta T_{\text{heater-output}}}{\text{power per unit area}}$$

For reference, conventional designs that are based on cooling blocks can achieve a TFM of 20 Kcm<sup>2</sup>W<sup>-1</sup>, which can be improved to values around 12 Kcm<sup>2</sup>W<sup>-1</sup> in designs with integrated pipe-structures [30]. For silicon microchannel solutions, the TFM is expected to reach values in the range 1.5-4.2 Kcm<sup>2</sup>W<sup>-1</sup> [31]. The TFM for the VELO modules is shown in Fig. 69 as function of the glue layer thickness. The TFM varies between 1.5 and  $3.5 \text{ Kcm}^2\text{W}^{-1}$ , as expected; for the average glue thickness of 80  $\mu$ m, the TFM is about 3 Kcm<sup>2</sup>W<sup>-1</sup>, which is the best achieved value in high energy physics experiments.

### 7.2.10 Summary of I-V measurements

The summary of all I-V measurements performed on tile NLO, as an example, is shown in Fig. 70. The I-V characterisation of a tile on the assembled module is performed before (solid lines) and after (dotted lines) thermal cycling. The black, orange and red dotted lines represent the performance limits that define the grading of the tiles, which were defined based on measurements carried out in a probe station when the tiles were received from the manufacturer. In this plot, the IV curves are coloured depending on their position relative to the performance limits (green in case they remain below them). As shown in the plot, the vast majority of the curves lie well within the high quality limit and therefore no issues are expected from increasing the bias voltage to compensate radiation damage.

## 7.3 Module quality

During module construction, each individual step described in Sect. 5 and Sect. 6 is assigned a grade based on the criteria discussed. These grades are then combined together in two ways to assess the overall quality of each module, a final grade and a quality metric.



Figure 65: The cooling performance of the modules (see text), as a function of the module number across the production. The plots are for the four tiles on each module, arranged from top to bottom as CLI, CSO, NLO, NSI. In each plot the values for each of the three ASICs are shown.



Figure 66: The change in cooling performance of the modules (see text) between the tests before and after thermal cycling, as a function of the module number across the production. The plots are for the four tiles on each module, arranged from top to bottom as CLI, CSO, NLO, NSI. In each plot the values that test for a potential degradation in cooling performance each of the three ASICs are shown.



Figure 67: Cooling performance as function of the glue layer thickness for installed modules.



Figure 68: Cooling performance degradation as function of the glue layer thickness for installed modules.

The final grade (A,B,C,D,F) is used to determine whether the module is installed in the experiment, with the quality metric (0-100) giving further information.

The final grade determines the module quality as good (grades A, B or C), acceptable (grade D) or bad (grade F). Good quality modules are assigned these three different grades (A, B or C) in order to have a finer assessment of their quality. Only modules with the good quality grades are used in the final detector. The position of the module within the VELO base is assigned using both the grade and the quality metric, such that the very highest quality modules are located in the most important regions for tracking. Modules are assigned a grade D when one functionality is missing, e.g. a data link is not functioning or a VeloPix is dead; these modules are considered to be still available for installation



Figure 69: Thermal Figure of Merit as function of the glue layer thickness for installed modules.



Figure 70: All IV curves for tile NLO. Solid (dotted) lines correspond to the measurements performed before (after) thermal cycling. The black, orange and red dotted lines define the quality grade assigned to the tile.

in the detector if needed. Modules with more than just one missing functionality are assigned the grade F and they are not considered for installation in the VELO detector.

Table 12: Number of production modules assigned each final grading.

Final grade	$ \mathbf{A} $	В	$\mathbf{C}$	D	$\mathbf{F}$
Number of modules	24	28	6	3	2

This approach provides a simple way to assess the overall quality of a module, in terms of its construction quality and performance. However, the rules listed above do not allow to distinguish between modules that have obtained the same final grade but have different results in the various tests performed during the assembly, which have different impact on the final performance of the detector. For this reason, a quality metric was also developed. The quality metric is also based on the grades of the assembly and qualification tests, but it attributes points to the grades of all the individual measurements. In addition, a weighting mechanism allows the performance results to affect the final quality of a module more than the assembly steps. This goal is achieved by splitting the production into three groups: (1) the processes, containing all the assembly steps, (2) the metrology, thermal cycling and thermal degradation tests, and (3) the electrical tests, before and after thermal cycling. The quality number is then defined in a range from 0 to 100, based on the results obtained in these three groups, weighted according to the ratio of 2:3:5 respectively, thus giving more weight to the performance measurements.

With the mechanisms described above, each module is assigned a grade and a quality number. The summary trend for all modules is shown in Fig. 71. Modules graded A, B or C are shown in green, while D and F grade modules are shown in orange and red, respectively.



Figure 71: Final grade and quality number of all the 64 production modules.

# 8 Conclusions

The LHCb VELO Upgrade is one of the most advanced pixel detectors constructed in high-energy physics and the modules have been designed to meet a number of exacting specifications. The modules operate in vacuum and in close proximity to the LHC beams. The closest elements are expected to accumulate an integrated flux of  $8 \times 10^{15} n_{eq}/cm^2$ . The VELO will reach a total data rate of 2.85 Tbit/s with 25 ns beam collision spacing. Up to 30 W per module cooling is required and is achieved by evaporative CO<sub>2</sub> circulated within microchannels that are etched into the silicon substrate that provides mechanical support to the sensors. The detection element of the modules are 200 µm thick planar n<sup>+</sup>-on-p sensors bump-bonded onto a 130 nm technology readout ASIC.

The double-sided module is composed of a significant number of components: bumpbonded ASIC and sensor tiles; microchannel cooling substrate; front-end hybrids; HV tapes; control (GBTx) hybrid; interconnect tapes; hurdle; bridge; carbon fibre legs; clamp; cooling pipes; LV cables; LV foot connector; aluminium module foot. Each component was designed, optimised and tested for its performance in the expected environment and to allow it to be incorporated during the module assembly. The module assembly procedure included a number of steps: a bare-module was formed from a foot, rods and microchannel plate; bump-bonded ASIC and sensor tiles were attached with high precision; readout and control hybrids were placed; wire-bonding was performed; and cable attachments made.

An extensive set of metrology and validation procedures were carried out, many in a realistic environment in a vacuum tank with  $CO_2$  cooling and a near-final electrical readout chain. The measurements included: module flatness; tile positions and flatness; glue layer thickness; hybrid positions; wire-bond pull-tests; module displacements with cooling; thermal cycling; IV curves; ASIC communication tests; timing and control tests; data transmission error rate tests; module threshold equalisation and noise level determination. All assembly and characterisation steps were recorded in a custom database and an overall module grading was obtained.

The assembly and quality assurance of sufficient high-quality modules for the full detector was completed in December 2021 at the University on Manchester and Nikhef. The sensor attachment was achieved with an average precision of 21  $\mu$ m, more than 99.5% of all pixels are fully functional, and a thermal figure of merit of  $3 \text{ Kcm}^2 \text{W}^{-1}$  was obtained. The modules were integrated into the two detector-halves at the University of Liverpool, with system tests of the halves performed, and then transported to CERN. Following final tests and adjustments, each half was installed in LHCb in time for the start of operation of LHC Run 3. The insertion of one of the detector halves into the VELO tank at LHCb is shown in Fig. 72. The detector is planned to operate for the next decade during Run 3 and Run 4 of the LHC.



Figure 72: Installation of an LHCb VELO detector half into the vacuum tank in the LHCb cavern.
## 9 Acknowledgements

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 and The Royal Society (United Kingdom).

## References

- LHCb collaboration, R. Aaij et al., The LHCb Upgrade I, arXiv:2305.10515, accepted by JINST.
- [2] LHCb collaboration, LHCb VELO Upgrade Technical Design Report, CERN-LHCC-2013-021, 2013.
- [3] LHCb collaboration, LHCb VELO (VErtex LOcator): Technical Design Report, CERN-LHCC-2001-011, 2001.
- [4] LHCb collaboration, LHCb Trigger and Online Upgrade Technical Design Report, CERN-LHCC-2014-016, 2014.
- [5] M. van Beuzekom et al., VeloPix ASIC development for LHCb VELO upgrade, Nucl. Instr. Meth. A 731 (2013) 92, PIXEL 2012.
- [6] V. Gromov et al., Development and applications of the Timepix3 readout chip, PoS VERTEX2011 (2011) 046.
- T. Bird, Flavour studies with LHCb: b-meson mixing, lepton-flavour violation and the velo upgrade, PhD thesis, University of Manchester, 2015, https://cds.cern. ch/record/2256124.
- [8] R. Bates et al., High speed electrical transmission line design and characterization, JINST 12 (2017) C02002.
- [9] M. van Beuzekom et al., The VELO Optical and Power Board, LHCb Public Note LHCb-PUB-2021-012 (2021).
- [10] P. Moreira, A. Marchioro, and Kloukinas, The GBT: A proposed architecture for multi-Gb/s data transmission in high energy physics, Proceedings of the Topical Workshop on Electronics for Particle Physics (2007) 332.
- [11] P. Moreira et al., The GBT-SerDes ASIC prototype, JINST 5 (2010) C11022.
- [12] R. Aaij et al., Performance of the LHCb Vertex Locator, JINST 9 (2014) P09007, arXiv:1405.7808.
- [13] O. Augusto De Aguiar Francisco et al., Microchannel cooling for the LHCb VELO Upgrade I, Nucl. Instrum. Meth. A 1039 (2022) 1.
- [14] R. Plackett et al., Measurement of Radiation Damage to 130nm Hybrid Pixel Detector Readout Chips, Proceedings of the Topical Workshop on Electronics for Particle Physics (2009) 157.
- [15] J. Back et al., Velo Upgrade Module Nomenclature, LHCb Public Note LHCb-PUB-2019-008 (2019).
- [16] S. De Capua et al., Shear and Thermal Testing of Adhesives for VELO Upgrade, LHCb Public Note LHCb-PUB-2016-026 (2016).

- [17] D. Murray, Developing silicon pixel detectors for LHCb: constructing the VELO Upgrade and designing a MAPS-based tracking detector, PhD thesis, University of Manchester, 2021, https://cds.cern.ch/record/2808484.
- [18] P. Svihra, Developing a Silicon Pixel Detector for the Next Generation LHCb Experiment, PhD thesis, University of Manchester, 2021, https://cds.cern.ch/record/ 2806219.
- [19] G. Zunica, Searches for CP violation in  $D^0 \to K_S^0 K \pi$  decays with the Energy test method, module assembly and commissioning for the LHCb VELO Upgrade, PhD thesis, University of Manchester, 2023, https://cds.cern.ch/record/2869815.
- [20] The PostgreSQL Global Development Group, PostgreSQL: The World's Most Advanced Open Source Relational Database, Webpage (2020), https://www.postgresql. org/.
- [21] Django Software Foundation, Django, Webpage (2020), https://djangoproject. com.
- [22] Python Software Foundation, Python, Webpage (2020), https://www.python.org/.
- [23] HoloViz developers, *HoloViews*, Webpage (2020), https://holoviews.org/.
- [24] L. Granado Cardoso et al., LHCb MiniDAQ control system, EPJ Web Conf. 214 (2019) 01005.
- [25] S. SIMATIC, WinCC OA SCADA V3.16, Webpage (2021), https://www.winccoa. com/.
- [26] F. Faccio et al., FEAST2: A Radiation and Magnetic Field Tolerant Point-of-Load Buck DC/DC Converter, 2014 IEEE Radiation Effects Data Workshop (REDW) (2014) 1.
- [27] F. Alessio et al., Clock and timing distribution in the LHCb upgraded detector and readout system, JINST 10 (2015) C02033.
- [28] M. Bellato et al., A PCIe Gen3 based readout for the LHCb upgrade, Journal of Physics: Conference Series 513 (2014) 012023.
- [29] T. Poikela et al., The VeloPix ASIC, Journal of Instrumentation 12 (2017) C01070.
- [30] P. Petagna, B. Verlaat, A. Francescon, Two-phase thermal management of silicon detectors for high energy physics, Encyclopedia of Two-Phase Heat Transfer and Flow III (2018) 335.
- [31] A. Nomerotski et al., Evaporative CO2 cooling using microchannels etched in silicon for the future LHCb vertex detector, JINST 8 (2013) P04004.