

The design, construction, operation and performance of the Belle II silicon vertex detector

1 Introduction

The Belle II Silicon Vertex Detector ([SVD](#)) is an important part of the upgrade of the KEK *B* factory consisting of the KEKB asymmetric-energy e^+e^- collider [1] and the Belle experiment [2] to a Super *B* factory with SuperKEKB and Belle II [3]. This upgrade significantly raises the physics requirements for the silicon tracking device. While KEKB collided 8 GeV electrons on 3.5 GeV positrons resulting in a Lorentz boost of $\beta\gamma = 0.425$, the more constrained low-emittance optics and higher beam currents of SuperKEKB require a reduction of the beam asymmetry, and collisions

of 7 GeV electrons on 4 GeV positrons imply a Lorentz boost factor of only 0.28, about two-thirds of that in KEKB. The resolution in the decay time difference Δt of the two B mesons produced in collisions at the center-of-mass energy of the $\Upsilon(4S)$ resonance is important for the study of time-dependent CP violation and should not deteriorate due to this change. As Δt is related to the Lorentz boost and the vertex separation Δz , $\Delta t \approx \Delta z/\beta\gamma c$, maintaining the same Δt resolution as in Belle implies improving the vertexing performance by about two-thirds.

At a very early stage in the vertex detector design, it became apparent that this is only achievable by reducing the 30 mm beam pipe of Belle to 15 mm or even 10 mm. Early simulations indicated that beam-induced backgrounds would increase by a factor of 15 in Belle II compared to Belle [4]. The Belle II SVD must thus be considerably faster than the Belle SVD2 to limit hit occupancy in the innermost layers to a level acceptable for the reconstruction software. This determined the choice of the APV25 as the frontend readout ASIC, which is capable of a 50 ns shaping time compared to the 800 ns of the VA1TA chip used in the Belle SVD2. Other guiding factors of the Belle II SVD design are the radiation tolerance and a drastic material reduction to avoid deterioration of the particle tracking performance due to multiple scattering in the device material. For radiation hardness, a total ionization dose (TID) of 10 Mrad corresponding to ten years of operation is assumed in the SVD design.

The result of this design and construction effort is the Belle II VerteX Detector (VXD) shown in figure 1, which consists of two layers of DEpleted P -channel Field Effect Transistor (DEPFET) pixel sensors (PXD) [5] and four layers of double-sided silicon strip sensors (SVD), arranged cylindrically around the beam pipe at radii between 14 and 135 mm. The cylindrical beam pipe has an inner radius of 10 mm and incorporates a 1 mm paraffin cooling channel defined by 0.4 mm and 0.6 mm inner and outer beryllium walls, resulting in an outer radius of 12 mm. While the placement of the innermost layer is guided by the desired vertexing performance, the outer radius is set by mechanical considerations of the support structure. Four silicon strip layers are required to perform autonomous particle trajectory finding in the SVD, in order to improve the tracking performance at the low momentum end of the particle spectrum and to be able to provide regions of interest for a reduction of the PXD data size. The VXD is mounted directly on the Belle II beam pipe and installed in the central opening of the Central Drift Chamber (CDC), the main tracking device of the Belle II detector, shown in figure 2. It operates in the 1.5 T magnetic field of the Belle II solenoid and provides inner tracking, impact parameter measurement, decay vertex reconstruction, and low-momentum particle identification for the Belle II detector [6].

This paper describes the design, construction, operation, and performance of the SVD. An initial appreciation of the overall project can be obtained by a reading of the overview paragraphs opening each of the following sections, where the main concepts are explained and the key features highlighted.

The paper is organised as follows. Section 2 contains the motivations and details of the SVD design. The construction of the detector is discussed in section 3. Section 4 deals with environmental and radiation monitoring and related safety interlocks. Section 5 is dedicated to the SVD online and control software. The different steps of SVD installation and commissioning in the Belle II detector are reviewed in section 6, while section 7 describes the operation of the detector during SuperKEKB physics data taking in the years 2019 and 2020. After an introduction to the SVD offline software (section 8), the main performance parameters are discussed in section 9. Finally, an outlook of future operations and developments is given in section 10.

A glossary and acronym list are available at end of the main text to help interpret the many




Figure 1. The Belle II vertex detector (VXD), composed of the pixel (PXD) and silicon strip (SVD) detectors.




Figure 2. The Belle II detector, surrounding the electron-positron interaction region: charged particles tracking is performed by the PXD, SVD silicon sensors and the CDC drift chamber, occupying the inner volume.

acronyms and technical terms. Active links are used when terms are used for the first time in a section to simplify browsing when reading the paper on a computer.

2 Detector design

The [SVD](#) design is guided by the major physics goals described in section [1](#), as part of the Belle II [VXD](#), operating in the challenging environment of the high-luminosity SuperKEKB electron-positron collider. The SVD contributes to inner tracking, impact parameter measurement, decay vertex reconstruction, and particle identification via energy loss measurements, with tight constraints on the material budget and strict requirements on radiation hardness and acceptable occupancy.

To fulfil these needs, four SVD *layers* with cylindrical symmetry (figure [3](#)) are inserted in the active volume between [PXD](#) and [CDC](#). The windmill-like arrangement of the sensors around the beam direction causes a partial overlap of the sensor active regions (between 2% and 8% depending on the layer), which is needed for the internal alignment of the device. The Lorentz boost, originating from the asymmetric-energy collisions, imposes a forward-backward asymmetry to Belle II detectors. To cope with the inclination of the forward tracks as well as to minimize the material traversed in the forward part, a lamp-shade geometry was adopted for the forward end of the three outer *layers* of SVD, whose detectors are inclined with respect to the cylinder axis. SVD layers are conventionally numbered from 3 to 6, taking into account the first two inner PXD layers.

The well-proven [Double-sided Silicon Strip Detector \(DSSD\)](#) technology is the basic choice for the SVD. Section [2.1](#) describes the optimization of sensor parameters for our application: shape, size, thickness, and readout strip pitch. The inner layer 3 has DSSD rectangular sensors with smaller size and pitch; the outer layers 4, 5 and 6 have larger rectangular sensors, and trapezoidal sensors for the lamp-shade part, necessitating dedicated R&D and design.

One of the major challenges in the design is the need to cope with the increased particle rates and hence with the readout channels occupancy, especially in the inner layers. The front-end readout APV25 [ASIC](#) [[7](#)], originally developed for the [CMS](#) experiment, has the required properties of reduced signal shaping time, good time resolution, on-chip integrated pipeline for dead-time free readout, and radiation tolerance.

The use of APV25 and its implications for the project are described in section [2.2](#), dedicated to the on-detector front-end electronics. Faster shaping time comes with larger noise: to keep it at an acceptable level, readout chips are placed as close to the sensors as possible. Novel solutions to achieve this goal, while keeping a low material budget, are implemented by a “chip-on-sensor” concept. Thinned-down readout chips are mounted on a flex circuit placed on top of the sensors, separated by a layer of foam for thermal and electrical isolation; flexible fan-outs are wrapped around the edge of the sensor to reach the strips on the bottom side. This scheme is named [origami scheme](#) with reference to the folding action.

The mechanical integration of individual detectors and their thermal control are also challenging. Each detector layer is organized in mechanically and electrically independent subsets called *ladders*, resembling the staves of a barrel. Although cold operation of the sensors, which operate around room temperature, is not required, active cooling based on dual-phase CO₂ is needed to remove the heat produced by the readout electronics, leading to large temperature variations (up to 40 °C) along the

ladders and when the electronics is turned off. The employed materials and the mechanical design are carefully optimized to avoid thermal stresses on the sensors and on the mechanical support.

The seven ladders of layer 3 are shorter and made of only two sensors, that can be read out by chips mounted on two traditional *hybrid boards* located outside the active volume, at the forward and backward ends respectively. Layers 4, 5, 6 are made of 10, 12, and 16 longer ladders, each containing 3, 4, and 5 sensors, respectively, for a total of 172 individual sensors. Only the first and last sensors in these ladders can be read out by the hybrid board arrangement; the central sensors require *origami boards* for the readout scheme mentioned above. In all cases flexible *pitch adapters* are used to connect sensor strips to APV25 channels.

The original solutions for these complex ladders, including insulating foam, DSSDs and support ribs are described in section 2.3. For assembling the ladders in layers 4 to 6 the relevant sub-parts of each ladder are the slanted *forward section* with one trapezoidal DSSD, and the remaining *barrel part*, made of rectangular sensors.

Figure 3 gives an overview of the SVD layout with its organization in *layers*, *ladders*, and sensors. The definition of the Belle II reference frame and of local frames for the individual sensors are recalled at the end of the present section. Sensors at different z position along the ladders, forward (FW), central, backward (BW), are labeled with increasing numbers starting with one for the FW sensor. In the rest of the paper each sensor is indicated by its position in the layer, ladder (ϕ), and z position along the ladder with the following convention: L3.2.1 stands for the position in layer 3, ladder 2, sensor 1 (FW) along the ladder, as shown in figure 3. An ‘x’ replacing a number indicates the entirety of sensors for that position; for instance L5.x.1 means all the sensors in layer 5, any ϕ position, z position 1 (that is, all FW sensors in layer 5).

In total, the SVD has an active DSSD area of about 1.1 m^2 and features 223744 channels read out by a total of 1748 APV25 chips. The material budget per DSSD layer corresponds to approximately 0.7% of the radiation length. The detector is designed to withstand an integrated radiation dose of 10 Mrad before significant performance degradation.




Figure 3. A schematic layout of the Belle II SVD: (left panel) longitudinal cut, identifying layers from 3 to 6, and individual sensors with their main properties; (right panel) transverse cut, showing a projection of layers and ladders, and the numbering scheme of the sensors. The reference frame is defined later in this section.

The entire SVD is longitudinally divided in two azimuthal halves, which are separately assembled before being mounted on the beam pipe and the PXD layers. The ladders of each half are supported and held in position at their two ends by joining the ladder aluminum *end-mounts* to stainless-steel (half) *end-rings* precisely glued on forward and backward carbon-fiber (half) *end-cones*. The innovative solutions adopted to meet the demanding requirements of position precision and stability

are discussed in section 2.4. The same section also describes the CO₂ cooling system designed to remove the heat generated by the front-end chips, using meander-shaped pipes in thermal contact with front-end ASICs, connected at their ends to cooling channels integrated in the end-rings.

Finally, the integration of the APV25 readout in the Belle II data acquisition system required a very significant investment in design, prototyping, programming and tests. The overall back-end electronics design, and the floating power supplies allowing a flexible scheme for detector bias and front-end powering, are described in section 2.5. A critical aspect for the detector performance is the electrical grounding scheme, described in section 2.5.5. The overall SVD design was validated by several cycles of prototypes and beam tests, summarized in section 2.6.

The main elements of the SVD readout chain are shown in figure 4. Both the hybrid boards and the origami flexes hosting APV25 chips are connected by 2.5 m copper cables to *junction boards* located in the forward and backward areas of Belle II, in front of the CDC end-wall, dedicated to interconnections, called **DOCKs**. These junction boards also contain the radiation-hard voltage regulators, which provide the supply voltages needed by the APV25 ASICs. Analog data are sent from the junction boards to the *Flash Analog to Digital Converter (FADC)* boards located on top of the Belle II detector. In addition to analog level translation, the **FADCs** perform first level data processing and formatting. Digital data are handed over to the *Finesse Transmitter Board (FTB)* boards, which distribute it through optical links to the *COmmon Pipelined Platform for Electronics Readout (COPPER)* boards [8] of the unified Belle II data acquisition system. Data is also sent to the *DATA CONcentrator (DATCON)* boards [9] of the PXD readout system, which use it to project candidate tracks and find regions of interest (**ROIs**) in the PXD.

The right-handed global Belle II reference frame, used in figure 3 and repeatedly throughout this report is shown in figure 5: the origin is located at the nominal **Interaction Point (IP)**; the z -axis is directed along the bisector of the angle between the two SuperKEKB beam pipes in the horizontal plane, pointing in the direction of the Lorentz boost (forward direction determined by the high-energy electron beam); the x -axis points towards the outside of the SuperKEKB rings, the y -axis points vertically upwards.

The SVD covers all azimuth angles ϕ and a polar angle θ ranging from 17° to 150°. Moving from the in- to the outside, the SVD layers 3, 4, 5 and 6 are located at radii R (measured from the z axis) of 39 mm, 80 mm, 105 mm and 135 mm. The DSSD readout strips on the N - and P -sides, as determined by the silicon diode configuration, read the z and $R\phi$ coordinates, respectively.

A local right-handed reference frame is also defined for each DSSD, with a u -axis in the sensor plane, pointing in the $R\phi$ direction, a v -axis parallel to the global z -axis, and the w -axis perpendicular to the sensor surface (figure 6). These local frames are used in the detector alignment as well as in the space point and track reconstruction. In different contexts, the detector sides are also identified as v/N -side and u/P -side.

2.1 Silicon sensors

The design of the silicon sensors is driven by the fact that multiple scattering is a key factor for the Belle II track reconstruction. Thus, the largest available sensors ought to be used to minimize dead material and the design has to be double-sided to achieve a low material budget. At the time of SVD design, typical wafer processing facilities were able to produce six-inch wafers, resulting in a maximum sensor size of approximately 12 cm length and a width of 6 cm. Rectangular barrel




Figure 4. An overview of the elements of the Belle II SVD readout chain.




Figure 5. The Belle II reference frame. Electrons circulate in the High Energy Ring (HER), while positron in the Low Energy Ring (LER).




Figure 6. Local coordinates for each DSSD.

sensors were fabricated by Hamamatsu Photonics (HPK) [82], while the production of trapezoidal sensors for the forward part was carried out by Micron Semiconductor [83].

The double-sided silicon-strip detectors from both companies are manufactured starting from an N -type bulk with high resistivity and a thickness of about $300\ \mu\text{m}$, which corresponds to the standard material available in industry. It also matches the design of the APV25 front-end ASIC [7] which is designed for capturing the signal charge from $300\ \mu\text{m}$ thick silicon sensors. The sensing strips are implanted in the N -type bulk either with acceptors or donors, depending on the sensor side. The acceptor-implanted side is the junction side, or “ P -side”, while the other side is called the ohmic side, or “ N -side”. The width of the implants is $14\ \mu\text{m}$ ($30\ \mu\text{m}$) on the P -side (N -side) of the barrel sensors and $15\ \mu\text{m}$ ($30\ \mu\text{m}$) on the P -side (N -side) of the trapezoidal sensors. The implanted strips are connected to a common bias rail using integrated bias resistors on each side with nominal values of $10\ \text{M}\Omega$. The silicon oxide contains fixed positive oxide charges, which attract electrons and thus create a conductive layer in the silicon shorting all N -type doped areas. For that reason, the N -strips on the junction side are isolated by a dedicated P -implantation surrounding the strips. This P -implantation is called P -stop.

To improve spatial resolution a floating strip is placed between two readout strips on both P - and N -sides. The charge induced in the floating strip is shared by the neighboring strips and the effective strip pitch is reduced to half of the readout pitch.

The readout metal strips are placed on top of the implanted strips, separated by a dielectric silicon oxide, hence the readout strips are AC-coupled. The width of the metallization is $20\ \mu\text{m}$ ($40\ \mu\text{m}$) on the P -side (N -side) of the barrel sensors and $25\ \mu\text{m}$ ($40\ \mu\text{m}$) on the P -side (N -side) of the trapezoidal sensors. The capacitance of the AC coupling is 15 and $30\ \text{pF/cm}$ on the P - and N -side, more than one order of magnitude larger than the typical interstrip capacitances, that are below $1\ \text{pF/cm}$. The breakdown voltage of the dielectric is typically well above $100\ \text{V}$, with all the AC capacitors tested up to $20\ \text{V}$. These figures are adequate considering that the readout ASICs are operated above the detector bias voltage (the floating power supply scheme, section 2.5.3). By using this biasing scheme, as explained in section 2.5.3, the potential across the AC decoupling capacitor is below $1\ \text{V}$, thus minimizing the probability of breakdown of the AC coupling. When a broken AC decoupling capacitor between the strip implant and the metal readout electrode is present, called a [pinhole](#), the corresponding pre-amplifier channel does not operate correctly, although some compensation are possible as explained in section 3.1.4. Pinholes can be created during sensor production or in the module assembly phases. A few additional pinholes, due to weaker capacitors, can also be generated during beam operation in case of large bursts of radiation delivered when the bias voltage is applied, as explained in section 7.4.4.

Table 1 shows the geometrical characteristics of the SVD sensors. All the sensor types carry the same kind of alignment marks (large F-marks in the corners) which can be detected automatically during the assembly process. The designs details of the forward and barrel sensors are described in the next two sections.

2.1.1 Forward sensors

Trapezoidal, $300\ \mu\text{m}$ thick sensors are used in the forward, lamp-shade part of the detector. Geometrical parameters are detailed in table 1 and figure 7, while electrical parameters are given in table 2. The strips of the P -side (junction side) feature a variable pitch, giving them a fan shape.

Table 1. Geometrical details of the double-sided silicon strip sensor types used in the SVD. All sensors have one intermediate floating strip between two readout strips.

	Small rectangular	Large rectangular	Trapezoidal
Readout strips <i>P</i> -side	768	768	768
Readout strips <i>N</i> -side	768	512	512
Readout pitch <i>P</i> -side (μm)	50	75	50 – 75
Readout pitch <i>N</i> -side (μm)	160	240	240
Sensor active area (mm^2)	122.90×38.55	122.90×57.72	$122.76 \times (38.42\text{--}57.59)$
Sensor thickness (μm)	320	320	300
Manufacturer	Hamamatsu	Hamamatsu	Micron

The strips of the *N*-side (ohmic side) are parallel to each other and perpendicular to the central strip of the *P*-side. Electrical strip separation on the ohmic side is achieved by *P*-stop blocking technique, where the so-called “atoll” variant was chosen, maximizing the signal-to-noise performance both before and after irradiation [10].




Figure 7. Dimensions of the active area of trapezoidal sensors in millimeter. Not shown are the multi-guard ring and the edge ring structures about 1.4 mm from the edge of the active area. The strip numbering for both *P*- and *N*-sides are indicated.

To achieve and maintain a good HV stability and high breakdown voltage, the sensors are designed with metal strips extending a few micrometers beyond the width of the implant strips, causing some of the field lines to end in the aluminum rather than in the implant. This reduces the electric field strength at the edges of the strip implants and takes advantage of the much higher breakdown voltage of the silicon oxide.

An additional measure to improve the reliability is the so called *edge-ring*. This is a heavily *N* type doped zone close to the cut region of the sensor which keeps the space charge region away from the imperfections created by the dicing. The edge ring is implanted on both sides and left floating. The edge rings on both sides has the *N*-side potential, which confines the bias voltage drop to a defined area.

Table 2. Electrical parameters of the trapezoidal sensor [3]

Quantity	Value
Substrate doping	<i>N</i> -type
Substrate resistivity	8 k Ω cm
Full depletion voltage	30–40 V
Breakdown voltage	> 100 V
Polysilicon bias resistor	10 \pm 5 M Ω
Coupling capacitance	100 pF (typ.)
Dark current	2 μ A (typ.), 10 μ A (max.)

In the area between the bias and the edge ring a multi-guard ring structure is realized comprising nine individual floating guard rings. These rings are used to shape the electric field and constitute a special feature of the sensors produced by Micron. A picture of this region is shown in figure 8.




Figure 8. Multi-guard ring structure of the Micron sensors for both the *N*-side (left) and the *P*-side (right). The small ruler in the top corners corresponds to 400 μ m

2.1.2 Barrel sensors

Geometrical and electrical parameters are detailed in figure 9 and tables 1 and 3. The barrel sensors are produced from 320 μ m thick, 15 cm diameter wafers. As the silicon wafers become curved in the semiconductor processes, this is the smallest thickness that achieves the flatness required for the

production at HPK. The sensors are made from *N*-type floating-zone silicon wafers with $6\text{ k}\Omega\text{ cm}$ resistivity and the expected full depletion voltage is about 65 V . As for the trapezoidal Micron sensors, electrical strip separation on the ohmic side is achieved by an atoll *P*-stop blocking structure. The aluminum trace above the implant is slightly wider than the implant to mitigate the border electric field, which causes a micro-discharge resulting in an increase of the strip current.




Figure 9. Dimensions of the layer 4 to 6 rectangular sensors. The region outside the active area is less than 1 mm wide and implements the edge ring structure. The location of the *P*-side bonding pads rows is indicated by the yellow regions.

Table 3. Electrical parameters of the barrel sensors.

Quantity	Value
Substrate doping	<i>N</i> -type
Substrate resistivity	$6\text{ k}\Omega\text{ cm}$
Full depletion voltage	$50\text{--}60\text{ V}$
Breakdown voltage	$> 200\text{ V}$
Polysilicon bias resistor	$4\text{ M}\Omega$ (min.), $10\text{ M}\Omega$ (typ.)
Coupling capacitance	$> 100\text{ pF}$ (<i>P</i> -side and <i>N</i> -side)
Dark current	$< 5\text{ }\mu\text{A}$ at $V_{\text{bias}} = 120\text{ V}$

The resistance of the polysilicon bias resistor is typically $10\text{ M}\Omega$ to keep the parallel noise negligible even given the short-shaping time of the AVP25 front-end. The bonding pads and polysilicon resistors are located inside the sensitive region as otherwise they would couple electrically to the undepleted silicon bulk just below the dielectric and cause large additional capacitance. In the barrel sensor the dead area outside active region is less than 1 mm wide (figure 10). In this area a heavily *N*-doped edge ring is implemented to prevent the extension of the depleted region to the cut edge. No multi-guard ring structures are present.




Figure 10. Detailed drawing of the *P*-side of a large rectangular sensor.

Two types of barrel sensors, one for layer 3 (124.88 mm × 40.43 mm) and one for layers 4 to 6 (124.88 mm × 59.60 mm) are produced (table 1). In addition to the wire bonding pads at the sensor edge, bonding pads are added in the middle of the sensor for the origami connection to the readout chips-on-sensor (figure 9).

2.1.3 Radiation tolerance

At the time of the Belle II TDR a preliminary estimate of the expected dose in the innermost SVD layer was based on the extrapolation from the dose received in the Belle SVD sensors, of about 90 krad/ab⁻¹. Assuming 50 ab⁻¹ integrated luminosity for the projected lifetime of Belle II, the radiation dose of the innermost SVD layer was then expected to be 4.5 Mrad.

To include some safety margin, a total dose of 10 Mrad was then initially set as requirement for the Belle II SVD sensors.

More detailed analyses of the beam background expected in the Belle II SVD, with a dedicated full simulation, were performed in the following years and those results, summarized in section 7.8, confirmed the previous assumption for DSSD radiation hardness requirements were adequate. From these background studies it was calculated that at the design luminosity of $8 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$ the expected average background level in the layer 3 sensors corresponded to 1–3% strip occupancy, 0.1–0.3 Mrad/yr integrated dose, and an equivalent neutron fluence of $0.2\text{--}0.6 \times 10^{12} \text{ n}_{\text{eq}}/\text{cm}^2/\text{yr}$. The interval quoted corresponds to the values from the initial MC background simulation result and from a more refined estimate, that uses background measured on data, up to 2020 run period, to re-scale the initial MC prediction.

Even with this more accurate background estimate, the 10 Mrad radiation hardness requirement, initially set for the SVD sensors, together with a requirement to an equivalent neutron fluence about $10^{13} \text{ n}_{\text{eq}}/\text{cm}^2$ can be then considered adequate for 10 years of operation at design luminosity. The radiation hardness of DSSD sensors up to these levels were already proven in the past.

As an example, the experience from the BaBar Silicon Vertex Tracker (SVT) DSSD sensors, produced by Micron with a design similar to the one used for our SVD sensors can be considered. The BaBar SVT was exposed to a radiation field similar to the one expected in the SVD. The SVT sensors were successfully operated for several years in BaBar reaching a total integrated dose of about 4.5 Mrad [11]. In addition during the lifetime of the BaBar experiment several irradiation campaigns were performed on the DSSD sensors to further study the effect of bulk damage, that was not negligible in that background environment. During these tests, described in [12] and [13], dedicated test structures, single DSSD sensors, and an entire module, with its sensors connected to the readout electronics, were irradiated with a 0.9 GeV electron beam, up to about 9 Mrad, corresponding to an equivalent neutron fluence of about $2.7 \times 10^{13} \text{ n}_{\text{eq}}/\text{cm}^2$. The measured electrical properties of the sensor and the charge collection efficiency of the full module proved that the DSSD detectors can be successfully operated without any problem up to the final irradiation step of 9 Mrad. No issues were found even after type-inversion, reached at about 2–3 Mrad for that sensor with initial depletion voltage of 25 V. After the final irradiation step a modest decrease of the charge collection efficiency, below 10%, was reported, together with an increase in the sensor leakage current of about $2 \mu\text{A}/\text{Mrad}/\text{cm}^2$ at 23 °C, similar to the observations in the BaBar sensors of about $1 \mu\text{A}/\text{Mrad}/\text{cm}^2$ at 20 °C.

Dedicated irradiation campaigns were also performed on the Belle II SVD rectangular and trapezoidal DSSD sensors, assembled with their readout chips in the origami and rigid hybrids. The modules were exposed to 9.5 Mrad of ^{60}Co γ radiation at SCK-CEN in Mol (Belgium) over one week [14]. Before and after this irradiation campaign the system was tested at a CERN test beam line and the response to minimum ionizing particles was measured. Although a moderate increase of the sensor leakage current was observed, the high signal-to-noise performance was maintained, and no degradation of resolution or the tracking performance was observed. In the beam test the system was operated with CO₂ cooling and the stability of the origami structure at low temperatures was also confirmed.

2.2 On-detector Electronics

Following the design principle outlined at the beginning of this section, the chip-on-sensor concept was developed, where the readout chips, thinned down to 100 μm to minimize their contribution to the material budget, are mounted on a flex circuit board that is placed on top of the sensors, separated by a layer of foam for thermal and electrical isolation. Moreover, flexible fanouts are wrapped around the edge of the sensor to reach the strips on the bottom side. In that way, it is possible to align all readout chips in a single row, cooled by a single pipe flowing two-phase CO₂. As a reference to that “folding” concept and the country where it is used, that scheme was named “[origami scheme](#)”.

2.2.1 Origami scheme

Figure 11 shows the structure of the origami scheme. The APV25 chips are aligned along Z, in the active region. The u/P strips are readout with two flexible pitch adapters (PA1 and PA2) that bend around the edge of the sensor, while the v/N strips are connected through a short pitch adapter (PA0). The PA1 and PA2 pitch adapters bend over the wirebonds of the sensor to the PA0, making the assembly quite challenging. A photograph of a completed origami module is shown in figure 12.




Figure 11. Structure of the [origami scheme](#). On the left, two views of the sensor are shown, both from the N-side. In the top sketch, the N-side strips are shown, connected to the pitch adapter PA0 and then to the APV25 chips. In the bottom sketch, the P-side strips are shown (as if the sensor was transparent), along with the PA1 and PA2 pitch adapters in a flat position. These pitch adapters are folded around the edge of the sensor, as shown on the right in a cross sectional view, so that the P-side strips can be connected to the APV25 chips.




Figure 12. Photo showing a completed [origami module](#) with P-side pitch adapters wrapped around the edge of the sensor.

2.2.2 APV25 front-end chip

The APV25 [7] is a low-noise charge sensitive amplifier chip in 250 nm CMOS technology. Its final version, APV25S1, was released in the year 2000 and about 70,000 devices were installed in the CMS Tracker. The APV25 has a shaping time of 50 ns (adjustable), a 192-cell deep analog




Figure 13. Building blocks of one of the 128 channels of the APV25 front-end readout chip. Reprinted from [7], Copyright (2001), with permission from Elsevier.

pipeline for each of its 128 channels, and a nominal operating clock of 40 MHz. The chip is known to withstand more than 100 Mrad of total ionizing dose [15].

The internal structure of the APV25 is shown in figure 13. The usual preamp/shaper architecture can be seen on the left, where an optional inverter is placed between these blocks. As the supply rails are limited to 2.5 V in the 0.25 μm technology, this switch was introduced to optimize the dynamic range depending on the polarity of the signal, i.e., reading out positive (*P*-side, inverter on) or negative (*N*-side, inverter off) detector currents. Thus, a linear range of approximately -2 to $+7$ MIPs (referred to a standard 300 μm -thick sensor) can be achieved for both polarities.

The shaper output is written to a pipeline of 192 cells at the clock frequency. In fact, the pipeline is implemented as a ring buffer memory with an additional FIFO of 32 words to label memory addresses that are requested for output by pending triggers. Until those data are read out, the tagged cells are skipped in the write cycle. Consequently, the available pipeline depth can vary between 160 and 192 cells, depending on the number of pending triggers. By multiplication with the clock period, this translates to the maximum trigger latency time.

After the pipeline, the APV25 has an analog pulse shape processor (APSP). This is, in fact, a switched capacitor filter that can perform a so-called *deconvolution* [16] or, in a different configuration, simply pass on the pipeline contents. Subsequently, the amplitude can be adjusted to some extent in the so-called MUX gain stage. Finally, the strip data are multiplexed through three hierarchical stages and sent to the differential-current-mode output.

Various bias voltages and currents as well as general parameters of the APV25 can be configured through its I2C interface. Clock and trigger signals are received by dedicated differential inputs. The trigger line also accepts special 3-bit symbols, namely 100 =trigger, 110 =calibration request, and 101 =soft reset. Consequently, two triggers must be at least 3 clock cycles apart; otherwise, they would be misinterpreted as a special symbol.

In the Belle II SVD setup, all APV25 chips used in the origami chip-on-sensor scheme are thinned down to 100 μm , while conventional hybrids located at the edges of each ladder, which reside outside the sensitive volume, use standard chips with a nominal thickness of 325 μm . The APV25 clock frequency in Belle II is 31.805 MHz, corresponding to 1/16 of the SuperKEKB RF clock of 508.887 MHz. A trigger latency of 159 clock cycles or approximately 5 μs — the maximum possible without reducing efficiency for high trigger rates — is used.

The deconvolution requires clock-synchronous collisions, which is not applicable to Belle II and therefore the non-processing *peak* mode is used or, more precisely, the *multi-peak* mode, where

three samples are recorded for each trigger. By duplicating the original trigger signal after three clock cycles, six consecutive samples are actually recorded.

A set of three or six samples along the shaped waveform allows to reconstruct both peak amplitude and hit time, provided that the timing is correct and the maximum sample is not at the edge. Especially, the peak time is important to distinguish real hits from off-time background. Initially, only 6-sample data are recorded, but a mixed 3/6 operation is in preparation to save both bandwidth and dead time at high luminosity operation. In that mixed mode, three samples are taken when precise trigger timing is available, while six samples are recorded in other cases — dynamically switching event-by-event, depending on the trigger source. Simulation has shown that this mixture can reduce the dead time fraction [17] from more than 3% (for 6-samples only) down to less than 1% at the maximum trigger rate of 30 kHz, although background rejection is negatively affected by a deterioration of the time separation between signal and background hits.

2.2.3 Hybrid boards

The sensors at the edge of each ladder, i.e. the so-called *forward* and *backward* sensors, are read out from the ends in a conventional way, using hybrid boards made of glass-reinforced epoxy laminate material (FR4), as shown in figure 14. Slightly different variants of those boards exist for *P*- and *N*-sides, forward and backward as well as layer 3. Generally, *P*-side hybrids host six APV25 chips, while *N*-side types have only four, except for layer 3 which features six chips on both sides. The numbering scheme of the APV25 chips on the hybrid boards is shown in figure 15.




Figure 14. The hybrid boards used in layers 4/5/6 (left: forward *P*-side board with 6 APVs, and right: backward *N*-side board with 4 APVs).

For ease of construction, each hybrid board is single-sided, meaning that components and APV25 chips are all attached onto the top side, while the bottom side remains empty. At a later stage, namely the construction of forward/backward modules and L3 ladders, pairs of *P*- and *N*-type hybrid boards are glued together back-to-back with a thin layer of thermally conductive, but electrically isolating, composite sheet of silicone rubber and fiberglass (“Sil-Pad”) in between. Finally, a ground wire is fed through matching vias in the hybrid boards and soldered onto both boards. This ground wire serves as a reference for the return lines of the HV bias supplies and is later attached to the metal structure of the ladder.




Figure 15. Numbering scheme of APV25 chips on origami boards (left) and conventional hybrids boards (right). Chips on the P and N-side are numbered independently. The small numbers near the chip edge are the APV25 channel numbers.

Each hybrid board is equipped with a 51 pin connector out of which 50 contacts are actually used [84]. All analog and digital signals, low voltage (LV) power, and high voltage (HV) bias are sent over this connector and the attached halogen-free fine-pitch (0.635 mm) flat cable with twisted pairs. Only in case of L3, where space is too tight for such a connector, the hybrid cables are directly soldered to the hybrid boards. All hybrid cables have the same length of 249 cm to ensure synchronous delivery of clock and trigger signals. Any cable slack is stored near the Junction Boards.

2.2.4 Origami flexible boards

For sensors in the central part (one, two or three depending on the layer), the readout circuits are realized with flexible PCBs called **origami boards**. The origami boards are a complex stack of conductive and insulating layers with a maximum thickness of 282 μm : 3 layers of 9 μm -thick copper laminated on a 25 μm -thick polyimide substrate; epoxy and insulating layers; polyimide coverlays to protect the circuitry. Unlike the hybrids, the origami boards serve both *P*- and *N*-sides and therefore have two connectors at the far end which are stacked on top of the hybrid boards at the time of module assembly. The connector region is stiffened by an H-shaped piece of FR-4 glued around the connector to prevent damage at the time of plugging/unplugging.

There are three different versions of **origami boards**, serving the (up to three) barrel sensors, labeled according to their Z location: “-Z” or “+Z” and “CE” for central. Two of them are read out on the backward side (“Origami_-Z” and “Origami_CE”), while one flex needed only in L6 is routed over the slanted trapezoidal sensor (“Origami_+Z”). Figure 16 shows a fully loaded **origami board** of the longest type (“Origami_CE”).



Figure 16. Fully loaded Origami_CE flexible PCB with dimensions $450.71 \times 57.39 \text{ mm}^2$. The 6 P-side and 4 N-side APV25 are located on the far left of the circuit.

On the **origami board**, all APV25 chips are aligned, such that they can be cooled using a single tube flowing dual-phase CO_2 . The four chips which read out the *N*-side strips (facing the flex) are

located in the center of each sensor, being complemented by two groups of three chips on each side, which read out the P -side strips to be reached by flexible pitch adapters bent around the edge (shown in figure 11). The numbering scheme of the APV25 chips on the [origami board](#) is shown in figure 15.

2.2.5 Junction boards

As shown in figure 4, the junction boards serve two purposes: 1) to connect the hybrid cables from the front-end to the data cables to the Flash ADCs ([FADCs](#)) and to the power cables from power supplies, and 2) to reduce the power supply voltages to a level accepted by the APV25 chips. The latter task is accomplished by DC/DC converter modules plugged onto the Junction Boards. Figure 17 shows one junction board, which serves four DSSDs that are powered, biased, and read out on both P - and N -sides. In consequence, there are eight (staggered) connectors for hybrid/origami cables, which are merged onto four broader connectors that carry the data of one DSSD each to the [FADCs](#). LV and HV power are supplied from the two connectors on the left side.



Figure 17. Junction Board with eight DC/DC converter modules.

The circuitry on the junction board ties together the negative side of the LV supplies to the respective HV bias potential on both P - and N -sides with an adjustable offset voltage V_{SEP} (section 3.1.4). This scheme implies that the front-end low voltages and thus also the output of the readout chips are not at ground level. The FADC boards provide electrical isolation of the chip signals to the ADC.

The outer walls of the CDC are located about 1 m in the backward direction and 1.6 m in the forward direction from the interaction point. Those regions host the [DOCK](#) boxes of PXD and SVD. One box is used for CO₂ cooling on each side. For the SVD, each [DOCK](#) box contains six junction boards and has a water cooling pipe at its bottom. Connectors for environmental sensors like the thermal NTC sensors, the diamond radiation sensors and optical fibers (section 4) are mounted onto some of the [DOCK](#) boxes.

Since the [DOCKs](#) are located in a moderate radiation zone (tens of krad) as well as in the magnetic field (1.5 T), specific design considerations must be applied. As the APV25 analog output can drive the signal through up to 25 m long Cu cables, the junction boards do not require repeater functionality, which simplifies the radiation hard design.

Local DC/DC voltage conversion was adopted to reduce the current and the voltage drop on the long cables, thus protecting the front-end chips from potentially dangerous transient over-voltages occurring in conjunction with sudden load variations. It was considered that the enhanced safety was worth the additional complexity in the junction boards design. Although off-the-shelf commercial

components operating in a radiation environment and inside a magnetic field were not available, the FEASTMP DC/DC converter modules [18], developed by CERN for the use inside LHC experiments with tolerance up to 100 Mrad and 4 T, meets the design requirements.

Each Junction Board holds a set of four converters, one each for the two voltage levels (2.5 V and 1.25 V) on the *P*- and *N*-sides. In order to account for a moderate voltage drop along the hybrid cables and [origami board](#) traces, the converters generate nominal output voltages of 2.66 V and 1.44 V, respectively, from the input of up to 12 V. In addition to that primary set, there is a second, redundant arrangement of DC/DC converters, where each device can be individually activated instead of the primary converter. All analog output voltages are monitored as well as the “Power-Good” output of each converter, so that DC/DC converter failures can readily be diagnosed.

2.3 Detector module and ladder design

The ladder design aims at keeping the material budget inside the active detector volume as low as possible (figure 18). Two carbon fiber (CF) reinforced ribs with a core made of very lightweight foam were chosen as the mechanical support structure, onto which the sensors are glued with epoxy. At their ends, the ribs are glued into two aluminum blocks ([end-mounts](#)), which are fixed to the SVD support structure. Since the readout chips dissipate heat and are actively cooled, the temperature along the ladder can vary from room temperature down to the coolant temperature ($-20\text{ }^{\circ}\text{C}$), requiring that the ladder support structure must be capable of compensating the mechanical stress induced by thermal expansion. Hence the backward fixation is designed as a solid bearing, while the forward end-mount integrates a sliding mechanism. The mechanical support and cooling are discussed in section 2.4.

2.3.1 Target mechanical and electrical requirements

The ultimate tracking accuracy is determined by the precision alignment obtained in software using actual particles (section 9.7). The requirements for the mechanical precision are rather driven by the mechanical tolerances of the support structures and the narrow space between adjacent ladders. In the mechanical design a minimum clearance of about 1 mm between adjacent components is adopted. In order to have sufficient margin, and taking into account that the achievable precision decreases with the length of the ladders, a target tolerance in the x and y directions of $\pm 100\text{ }\mu\text{m}$ for L3, L4 and L5 and $\pm 150\text{ }\mu\text{m}$ for L6 is defined. For the z coordinate a slightly relaxed tolerance of $\pm 150\text{ }\mu\text{m}$ in L3, $\pm 200\text{ }\mu\text{m}$ in L4 and L5, and $\pm 250\text{ }\mu\text{m}$ for L6 is adopted. These tolerances refer to any reference point on any sensor with respect to its nominal position, measured while the ladder is still mounted on its assembly support.

In the floating powering concept (section 2.5) the ground levels of the APV25 chips are bound to +HV on the *N*-side and -HV on the *P*-side, respectively. It is thus important to ensure that all sensors and hybrid boards are well isolated from each other and also fully disconnected from conductive components of the support structure. On the other hand, a common grounding scheme is needed in order to reduce susceptibility to noise of the system. Therefore, two ground points are defined at the forward and backward end-rings, respectively, where all electrically conductive components of the ladders as well as the HV return lines of the hybrid boards and origami boards are connected.




Figure 18. Exploded view of a layer 6 ladder. From bottom to top: rib sub-assembly, four rectangular and one trapezoidal DSSDs, pitch adapters (red), hybrid boards at both ends of the ladder, AIREX[®] foam, three versions of origami flexes.

2.3.2 Ladder design

The basic element of the SVD, a so-called ladder, consists of a mechanical support onto which the sensor modules are glued (figure 18). The ladder designs of layers 4, 5 and 6 (L4, L5 and L6) are very similar and can be divided into the barrel part with rectangular DSSDs and the slanted forward section with one trapezoidal DSSD. The ladders vary only in the slant angle of the trapezoidal sensor and the number of rectangular sensors in the barrel part. In contrast, layer 3 (L3) ladders consist of two small rectangular sensors and do not have a slanted section. Another unique feature of the L3 assembly is that the sensors face radially inwards while they face outwards in the other SVD layers. A photograph of all four kinds of ladders is shown in figure 19.

Mechanical ladder support structure. The mechanical support structure of the ladders consists of two ribs made of 3 mm thick AIREX[®] R82.60, a commercially available polymeric foam with outstanding dielectric properties, excellent stiffness to weight ratio, low water absorption, and high




Figure 19. From top to bottom: ladders of L6, L5, L4 and L3 placed on their transport frames, in the very front the forward and backward sub-assemblies are shown on their transport containers.



Figure 20. Mechanical support structure of a layer 4 ladder, consisting of a backward end-mount, 2 CF reinforced AIREX[®] ribs, a forward end-mount with sliding mechanism and two “Kokeshi” pins (refer to the text for a description of these elements).

durability in particular in radiation environments [85]. The foam core is reinforced by a 105 μm thick carbon fiber (CF) sheet on either side. In L3 a simple bar-shaped rib design is sufficient to achieve the required mechanical stiffness, while for the longer outer layers a truss design is adopted for the ribs (figure 20).

At their ends, the two ribs are glued into aluminum blocks called **end-mounts**, forming the **Rib Sub-Assembly (RSA)**. In L4, L5 and L6 precisely milled pins are inserted into holes located underneath the hybrid boards in the center of the end-mounts. With these pins the ladders are plugged into the precision holes of the end-rings and finally fixed with set screws. These set screws are designed to pull down the end-mounts onto the end-ring surface by a V-shaped notch. Since the shape of the pins looks like a traditional Japanese “Kokeshi” doll, these pins were named “**kokeshi-pins**” (see figure 25). The fixation of the backward end-mount is designed as a rigid joint and the center of its **kokeshi-pin** defines the mechanical origin of a ladder. On the forward side there is a sliding mechanism in the end-mounts to compensate for thermal expansion of the ladders: a prism made of brass is held by the kokeshi-pin, while the end-mount is pushed onto it by spring-loaded stainless steel balls and can slide along the beam axis (z) within a range of more than ± 1 mm.

Due to geometric constraints, the aluminum blocks of layer 3 have a different design and are longer than those of the outer layers. They extend from the hybrids at both ends of the ladders and are called bridges. L3 ladders are fixed with normal M3 screws, without using kokeshi-pins. At the forward bridge an oblong hole allows compensation of thermal expansion.

Pitch adapters. Pitch adapters are flexible printed circuits used to connect DSSDs with APV25 chips. They are produced in different shapes and lengths, depending on the sensor type, side, and position in the SVD layout. The electrical connection between the chips and the pitch adapters, as well as between pitch adapters and the sensor strips, is established by ultrasonic wedge wire bonding with 25 μm -thick aluminum wire. The main requirements for the pitch adapters are thus mechanical flexibility, flatness, fine pitch, good capability of performing wirebonding on the pads, and electrical quality. Mechanical flexibility is particularly important for the **origami scheme** and a single copper layer design is chosen to enable bending without peeling off after gluing. The base material is a 25 μm polyimide film with 5 μm copper. A coverlay of a 12 μm polyimide film with a 15 μm glue layer is added on both top and bottom of the base material. The bottom coverlay is required to ensure the flatness required for good gluing quality.

The bonding pad pitch is smallest on the APV side and amounts to only 44 μm . To achieve a pad pitch with enough bonding pad width for good wire bonding, the bonding pads are arranged in three rows as shown in figure 21. The minimum pattern width and pitch are 20 μm and 15 μm on the design, respectively. For good wire bonding quality, a minimum bonding pad width of 30 μm is required, and a thin displaced gold plating is applied onto a nickel plating with a thickness of 3 to 4 μm .

The name and functions of the eleven different pitch adapter geometries are summarized in (table 4). Their specific use and locations is described in the discussion of the ladder design and construction.

Layer 3. Layer 3 ladders consist of two small DSSDs, glued on the mechanical support structure, made by two aluminum bridges and two support structures called “ribs”, made by a 1 mm AIREX[®] sheet sandwiched by 0.1 mm carbon fiber reinforced plastic (CFRP) plates. The APV25 readout chips are located on hybrid boards at both ends of the ladder. There are separate boards for the N - and P -side chips on each side, glued together with an electrically isolating but thermally conductive foil (Sil-Pad[®] 800 [86]) in between. These hybrid sandwiches are then screwed onto the aluminum




Figure 21. Bonding pad arrangement in three rows for the pitch adapters (top), used to achieve a small bonding pitch ($55\ \mu\text{m}$) while keeping a sufficient bonding pad width ($45\ \mu\text{m}$). The corresponding corner of the APV25 chip is shown below, together with a few wire bond connections in red. (In reality, APV25 and pitch adapter are farther apart.)

Table 4. List of the eleven different types of flexible polyimide circuits, the pitch adapters, used in the layer 3 – 6 ladders.

Name	Layer	Location
P3F1, P3F2	3	Forward sensor, P - (1) and N -side (2)
P3B1, P3B2	3	Backward sensor, P - (1) and N -side (2)
PF1, PF2	4 – 6	Forward sensor, P - (1) and N -side (2)
PB1, PB2	4 – 6	Backward sensor, P - (1) and N -side (2)
PA0	4 – 6	N -side, on origami boards
PA1, PA2	4 – 6	Wrapped from the origami P to the N -side

bridges. The L3 pitch adapters are glued at one end to the hybrid boards in front of the readout chips and at the other end directly to the sensor surface. The heat produced by the APV25 chips is drained through the bridges, which are screw-mounted to the CO_2 cooled end-rings. To improve the thermal contact between the P -side chips and the bridges, a 0.5 mm thick soft and thermally conductive pad (Keratherm Softtherm 86/235 [87]) is inserted. The APV25 chips on the N -side hybrid board are indirectly cooled through the PCB. A schematic of the full Layer 3 ladder, with the ribs and the bridges highlighted on the u/P side, is shown in figure 22.




Figure 22. Schematic of a completed Layer 3 ladder, u/P side (top) and v/N side (bottom). The orange shapes are the pitch adapters.

Layer 4, 5 and 6. Since the design of the ladders of L4, L5 and L6 is quite similar, the L6 design is described as an example. It comprises five DSSDs and is the most complex SVD ladder design. An exploded view of an L6 ladder is shown in figure 18.

At both ends of the ladder there are the [SVD Forward module \(SFW\)](#) and [SVD Backward module \(SBW\)](#). Those are very similar to the sensor modules of the L3 ladders and consist of one DSSD, two pitch-adapters and two hybrid boards, each. The hybrid boards (section 2.2.3) are conventional PCBs with four APV25 chips on the *N*-side and six chips on the *P*-side, respectively. Two flexible polyimide pitch adapters are used to connect the sensor strips to the APV25 inputs. As in L3 these pitch adapters are glued onto the hybrid boards and the sensors and connected by wires bonds. These hybrid boards are screwed to the forward and backward end mounts, which are located outside the sensitive area of SVD.

In the central section of the ladder, the origami chip-on-sensor scheme is used. Three rectangular DSSDs are glued onto the CF-reinforced ribs with epoxy glue. A sheet of 1 mm thick AIREX[®] R82.60 foam [85] is glued on top of all sensors, including the forward and backward DSSDs. This AIREX[®] sheet acts as a thermal and electrical isolation layer between the DSSDs and the readout electronics. In order to compensate the steps along the surface of the DSSDs, caused by the pitch adapters glued on the forward and backward sensors, the foam is cut from plates of 1.2 mm thick AIREX[®] R82.60 and then thermally formed to the final shape and thickness. The readout of the central sensors is done via thin flex PCBs, the [origami boards](#), introduced in section 2.2.4, which are glued on top of the AIREX[®] sheet. The [origami boards](#) carry only the APV25 chips and a few capacitors and resistors in the acceptance region. The connectors and larger electronic components are all located at the end of the ladder and stacked on top of the hybrid sandwiches of the forward or backward modules, respectively. Therefore, the Origami_-Z is glued on top of the tail section of the Origami_CE flex. The Origami_+Z flex is attached to the second forward sensor and is routed to the forward end of the ladder. All the APV25 chips of a ladder are arranged in a row in order to allow attachment of a single thin walled pipe for cooling. To hold the cooling pipe, small clips made from PEEK [88], a high performance plastic material, are glued in between the APV25 chips.

2.4 Mechanical support and cooling

A schematic view of the VXD mechanical structure is shown in figure 23. The rigid VXD cylindrical structure is made of 1 mm thick Carbon Fiber Reinforced Plastic (CFRP) cylinder (outer cover), and the end-flanges made of aluminum on each end. The structure is divided into two parts along the yz plane (naturally defining the halves as $+x$ and $-x$) for installation of the SVD sensors. The end-flanges are divided into two halves combined by the connection rings in the forward and backward directions. The VXD components are connected to the end-flanges. The radiation shield made of tungsten alloy and the tantalum part of the beam pipe weigh 40 kg and 20 kg, respectively, and are the main source of mechanical stress to the VXD structure.




Figure 23. Main components of the VXD structure. The SVD and PXD sensors are not shown.

Inside the cylindrical structure, the end-rings, on which the ladders are mounted, are glued to the support cones, made from 3 mm thick CFRP, which are connected to the end-flanges. The beam pipe is supported near the center of gravity of the forward and backward tantalum chamber to reduce stress to the beryllium pipe. The CFRP used for the outer cover and support cones [89] has tensile elasticity 212 GPa, shearing strength 42 MPa, and thermal expansion coefficient 3 ppm/°C.

The ladders are mounted onto the end-rings by precision “Kokeshi” pins on each ladder that correspond to precision holes in the end-rings. All the support structures described in this section are located outside of the active volume and material budget is of less concern. Cooling channels integrated into the end-rings carry dual-phase CO₂ that is used to remove the heat dissipated by the APV25 chips on the hybrid boards. Figure 24 shows one half of the support structures mentioned above. The two halves are populated with ladders individually and later combined around the PXD.

The APV25 chips of the origami boards are cooled by a meander-shaped pipe, called origami pipe, which is attached onto a half-layer once all its ladders are mounted (figure 26). The circuits of layers 4 and 5 are connected in series, while layer 6 has its own circuit. The full list of cooling circuits is given in table 5. A special miniature connector, originally developed for the CMS Pixel Detector and named “Streuli” after its inventor, is used for the interconnections of cooling pipes in the active region, where space constraints prohibit the use of bulky commercial products. Ceramic



Figure 24. Support structure of the $+x$ half of the SVD, composed of carbon fiber cones, stainless steel end-rings, and aluminum end-flanges. Cooling pipes are interfaced through ceramic isolators (in beige plastic holders) at the end-flanges. The BWD (left) and FWD (right) parts are later connected by the outer-cover (not shown).

insulators are located in dedicated pockets of each end-flange and electrically separate the cooling pipes inside the SVD volume, which are grounded to the end-flange, from the outside world.

Table 5. List of SVD cooling circuits. There is a small asymmetry between $+x$ and $-x$ because of the different number of L3 ladders (4 and 3, respectively). A nominal power consumption of 0.4 W is assumed here for each APV25 chip.

Half	Name	APV25 chips	Power [W]
$+x$	BWD_end-ring	238	95.2
	FWD_end-ring	238	95.2
	L45_Origami	170	68.0
	L6_Origami	240	96.0
$-x$	BWD_end-ring	226	90.4
	FWD_end-ring	226	90.4
	L45_Origami	170	68.0
	L6_Origami	240	96.0
Total		1748	699.2

2.4.1 Cones and end-rings

The carbon fiber cones are shaped to fit the boundary of the PXD cable space underneath. Due to the lantern-like structure of the SVD, the FWD cone is significantly longer than its BWD counterpart. In total, there are four half-cones, one each for $+x$ and $-x$ halves as well as BWD and FWD, respectively.

For the end-rings, stainless steel 316L (1.4404) [90] was chosen because of its relatively low CTE (16 ppm/°C), high machinability, and small magnetic permeability that make it suitable to be used in the 1.5 T magnetic field of Belle II. The actual mount points for ladders are precision holes in the end-rings which receive the “Kokeshi-pin” of the ladders in layers 4, 5 and 6. The set

screw presses onto the V-shaped notch of those pins and push the ladder end-mounts against the end-ring surface, as shown in figure 25. The end-rings are made from two diffusion-welded stainless steel pieces with an internal cooling channel to remove the heat from the ladder end-mounts. The end-rings of layers 5 and 6 have large openings to accommodate the hybrid cables of the inner layers. Layers 3 and 4 are mounted onto a common end-ring. In contrast to the “Kokeshi” scheme, L3 ladders use straight precision pins and are screwed from top. Copper inserts are placed at the L3 mount points to supplement the limited thermal conductivity of stainless steel.




Figure 25. Left: cross section of the BWD ladder mount mechanism: end-mount (yellow) and “Kokeshi” pin (red) belong to the ladder, while end-ring (grey) and set screw (green) are support parts. The embedded cooling channel between the two diffusion-welded parts of the end-ring is clearly visible. Right: on the FWD side, a spring-loaded sliding mechanism allows relative motion along the beam axis (z) while maintaining a thermal contact.

A ground braid connects each end-ring to the end-flange which serves as the local grounding point. Two NTC temperature probes are attached to each end-ring, and three diamond radiation sensors are glued onto each cone between the L3/4 and L5 end-rings. There are two nitrogen inlets (for creating the dry detector volume) and one suction pipe (for dew point monitoring) on each half on the FWD side, and one nitrogen inlet and one suction pipe on the $-x$ BWD half. The environmental and radiation monitors are described in section 4.

2.4.2 Origami pipes

Once all ladders of a half-layer (4/5/6) are mounted, a cooling pipe (1.6 mm outer diameter, 0.1 mm wall thickness) with a meander-like shape following the surface of the ladders is attached to each ladder using custom-made clips made of PEEK. Soft and thermally conductive pads are placed between APV25 chips and the pipe to efficiently transfer the heat from a flat surface to the round pipe. The clips allow to remove the pipe at a later stage and access the individual ladders if needed.

Half-layers 5 and 6 have an even number of ladders, such that the inlet and outlet of the pipes are at the BWD end of the straight sections of first and last ladder. L4 however has an odd number, such that the pipe needs to be routed back to BWD on the last ladder, as shown in figure 26.

2.4.3 Outer-cover

The detailed structure of the outer cover is shown in figure 27. The two halves of the outer cover are identical, 898 mm long with a 152.8 mm radius. The brackets are glued to the outer cover at




Figure 26. The origami pipe is held by PEEK clips and heat is transferred from the APV25 chips to the pipe through heat conductive soft pads. Due to the odd number of ladders, the pipe needs to be returned on the last ladder of L4.

both ends, so the total length of the outer cover becomes 913.5 mm. The brackets interface the outer cover to the end-flanges to keep the cylindrical shape of the outer cover under the strong stress due to the weight of the VXD. To ensure a proper ground connection between the components of the read-out electronics, a 0.1 mm thick aluminum sheet is glued on the outer surface of the cylinder and electrically connected to the brackets, to establish an electrical connection with the forward and backward end-flanges.




Figure 27. The structure of the VXD outer cover. All dimensions are in mm.

Because of the dual-phase CO₂ cooling system, the temperature can vary from $-20\text{ }^{\circ}\text{C}$ to $+20\text{ }^{\circ}\text{C}$ (room temperature) within the volume of the VXD. The CFRP is a suitable material to reduce the thermal stress to the VXD sensors.

A finite element analysis using Ansys [91] was performed, assuming the VXD is supported at the forward and backward end-flanges. It was found that the maximum deviation of the outer cover is as small as $15\text{ }\mu\text{m}$ near the center, while the maximum stress on the outer cover is 2.1 MPa at the interface to the end-flange. For the assembly of the VXD structure, the Stycast 2850FT epoxy adhesive with the 23LV catalyst [92], originally developed for cryogenic applications, was used.

The adhesive has a large elastic modulus (6700 MPa) and small CTE (39 ppm/°C). To confirm the reliability of the adhesive, five thermal cycles between $-200\text{ }^{\circ}\text{C}$ (liquid nitrogen) and room temperature were performed. The glue samples were also irradiated up to a dose of 10 Mrad with a ^{60}Co γ source and subjected to a shear strength test. The shear strength measured after the irradiation was larger than 7 MPa, more than a factor 3 above the calculated maximum stress. The adhesive and the material were therefore judged strong enough to keep the structure robust over the entire lifetime of the Belle II experiment.

2.4.4 Cooling plant and distribution

The power consumption of 1748 APV25 chips amounts to 700 W (table 5). As the neutron flux in the SVD region is expected to be about $10^{13}\text{ n}_{\text{eq}}/\text{cm}^2$ after 10 years of operation at the design luminosity (section 7.8), the bulk damage to the DSSDs is limited and the SVD can be operated at room temperature. The two-phase CO_2 cooling system operated at $-20\text{ }^{\circ}\text{C}$ cools the APV25 chips located in the active detector region. The following considerations lead to the choice of a CO_2 cooling system:

- The liquid CO_2 has a large latent heat. When 1 g of liquid CO_2 evaporates at $-20\text{ }^{\circ}\text{C}$, it absorbs 300 J of heat. Theoretically 700 W of heat can be removed with a flow of 3 g/s. In case of the water cooling, to limit the temperature rise to $\Delta T = 10\text{ }^{\circ}\text{C}$, a flow $17\text{ cm}^3/\text{s}$ of water is needed.
- The viscosity of the liquid CO_2 is small and allows the use of small-diameter tubes for coolant circulation. For the SVD, the inner tube diameter is only 1.5 mm. The longest tube inside the VXD region is 1 m long, removing heat from 240 APV chips.
- A wall thickness of 0.05 mm provides sufficient pressure margin. As a result, the material budget for the cooling inside the Belle II acceptance is reduced to 0.003% radiation length per layer, which corresponds to about 1/7 of that of the water cooling [3].

The liquid CO_2 plant, IBelle [19], supplies cooled liquid CO_2 to the VXD system (SVD and PXD) up to a rate of 30 g/s. has been operated stably by the KEK-MPI collaboration. A detailed description of this cooling system is given in section 4.5.

2.4.5 Mechanical integration

The ladder mounting procedure for each half consists of numerous steps, and generally works from inside out, starting with an empty set of half-cone assemblies mounted on a dummy beam pipe. The cones are equipped with environmental sensors, grounding, and cooling infrastructure (figure 28).

Before the first actual ladder is mounted, mechanical interference tests are performed between PXD and L3 as well as all neighboring ladders (within and between adjacent layers) using geometrically accurate dummies. Due to the overlapping windmill-like pattern of the sensors, the ladder mounting must be done in the direction of ascending ϕ angle.

For layer 3, the procedure is completed after mounting the 4 (+x half) or 3 (-x half) ladders. All outer layers (L4/5/6) require attachment of a pre-bent [origami pipe](#) after placing all ladders of one layer. Each ladder is electrically tested after mounting (without cooling) and each half-layer is tested (with cooling) after the corresponding [origami pipe](#) is mounted. This is done inside a dry box covered with infrared transparent foil, so that the actual cooling contact between the pipe and all




Figure 28. The $+x$ FWD half-cone assembly on the ladder mount table with diamonds, NTCs, grounding braids, N_2 and sniffing pipes, cooling pipes and ceramic insulators.

APV25 chips can be confirmed with a thermal camera. On the ladder mount table in Tsukuba hall B1, the cooling was done with an open-loop CO_2 system.




Figure 29. Fully completed $+x$ half of the SVD. The visible ladders as well as their [origami pipe](#) belong to the outermost layer 6.

After completing a half (see figure 29), it is tested as a whole, followed by attaching the outer-cover and its transfer to the storage area, in order to make space for the other half on the ladder mount table. Later, the two halves are placed into a dry and dark box for cosmic testing before being attached onto beam pipe and PXD. After a combined cold testing, the completed VXD is installed in Belle II.

2.5 Off-detector electronics

The general outline of off-detector electronics is summarised in figure 4 and consists of Junction Boards, the **Flash Analog to Digital Converter (FADC)** boards, the **Finesse Transmitter Board (FTB)** boards, the **COMmon Pipelined Platform for Electronics Readout (COPPER)** boards, and the **DATA CONcentrator (DATCON)** readout modules. Their functions and characteristics are described here in more detail. Due to the lack of active electronics on the Junction Boards, the APV25 chips have to drive their output all the way up to the **FADC** boards and need to be properly terminated there. Signal distortions due to the limited bandwidth of these 15.5 m of twisted pair cables need to be addressed as well as potential reflections. Moreover, the signals are at positive and negative bias potentials and need to be translated to ground levels before they can be digitized.

As the analog circuitry needs a considerable amount of space, a 9U VME system was chosen which features large boards ($400 \times 366.7 \text{ mm}^2$). The “intelligence” of the system is packed into a powerful FPGA as the central device on each **FADC** board. It receives the digitized data streams from the front-end and sends the processed data out on a parallel bus leading to the optical “Belle2link” (Xilinx RocketIO over fiber) that transmits data to the unified **COPPER Belle II DAQ** modules [8]. Moreover, it has an interface to the VME bus and a Gigabit Ethernet port. A second optical link is used to stream the same data to the **DATCON** units of the PXD, which are used for online tracking in order to find regions of interest (RoI) on the pixel planes and reduce the PXD data transmission rate.




Figure 30. Schematic readout chain of the Belle II SVD. Reprinted from [20], Copyright (2013), with permission from Elsevier.

The **FADC** system (figure 30) is distributed over four 9U VME crates, two each on FWD and BWD sides, located on top of Belle II. It consists of one **FADC_Controller** board that interfaces to

four Buffer boards, one in each crate, and a total of 52 FADC boards, each of which sends out data to the central data acquisition (DAQ) and **DATCON** through an attached **FTB**. The communication within each crate takes place over a customized Backplane between Buffer and **FADC** boards. An asymmetric distribution of **FADC** boards (32 on BWD and 20 on FWD) reflects the fact that the majority of cables exit the **VXD** on the BWD side. One Junction Board is attached to one **FADC** board, except for L3 where the output of each Junction Board is shared between two **FADC** boards in order to reduce the bandwidth burden for the upstream DAQ system.




Figure 31. Overall Belle II DAQ scheme.

The unified DAQ [21] consists of 9U VME modules named **COPPER**, which collect the data from all subdetectors through the optical links mentioned above, as shown in figure 31. Those data are then read out by PCs, combined in the Event Builder and processed by a computer farm known as **High Level Trigger (HLT)**, which performs tracking and reconstruction and finally selects events to be written to disk. Clock, trigger and related information are distributed to the **FADC_Controller** and each **FTB** board by the **Front-end Timing Switch (FTSW)** [17].

2.5.1 FADC system

The **FADC** system has one **FADC_Controller** which interfaces to 52 **FADC** boards, distributed over four crates, through one buffer board per crate. The controller receives its input (such as a trigger) either over VME (in case of local runs) or from the **FTSW** system (global runs). Some monitoring and also spy data can be read from the **FADC** boards through the VME bus. The primary data output is through a parallel interface to **FTB** and from there optically to DAQ (and **DATCON** in parallel). The communication between the single controller and each **FADC** board is performed through Buffer boards and custom Backplanes to which every **FADC** (figure 32) is connected. Each **FADC** also features a Gigabit Ethernet (GbE) port, which is intended for fast local data taking.




Figure 32. FADC board (9U VME).

Each 9U VME crate has three 6U slots on the left side (VME controller, empty, Buffer), followed by 18 full-sized slots (FADCs). The Backplane is attached on the rear of the J2 connectors and runs from slot 3 (Buffer) to 21. It utilizes pins for internal communication which are not used by the VME bus. Broadcasts (from the FADC_Controller to the FADCs) are sent over differential LVDS lines, and a few open collector lines are available for communication in the opposite direction.

In general, one Junction Board is connected to one FADC board, and thus the latter receives data from up to four *P*- and four *N*-type hybrid boards, corresponding to up to 48 channels. To avoid readout bandwidth limitations in the FADC to COPPER connection, the readout of Junction Boards which serve L3 hybrids are split between 2 FADC boards each. Therefore, the SVD system has 48 Junction Boards in total, but 52 FADC boards.

Figure 33 shows the functional blocks of an FADC board. Data processing mainly happens within the firmware of the central FPGA, an Altera (now Intel) Stratix IV GX. It is supported by another simpler FPGA (Altera Cyclone II) that can be used to flash the firmware over VME onto the main device.

The signals to/from the Junction Boards are shown to the left in figure 33: clock (31.805 MHz), trigger and I²C control signals are sent to the APV25 front-end chips, while their analog output is




Figure 33. Schematic block diagram of the FADC. Reproduced from [22]. © IOP Publishing Ltd and Sissa Medialab. All rights reserved.

received as well as HV bias current, LV and Junction Board temperature measurements. Moreover, Enable and Power Good signals are propagated to/from the DC/DC converters on the Junction Boards. All those signals are related to the APV25 supply voltages, which reside on top of the negative (*P*-side) and positive (*N*-side) HV bias levels and therefore need to be translated to ground levels on the FADC. This is done using capacitive coupling for the signals (essentially a high-pass filter) and commercial digital insulators (based on integrated HF transformers) for the digital signals. After level translation, the analog signals are digitized with adjustable clock phase and sent to the central FPGA.

Inside the FPGA firmware (shown schematically in figure 34), each 10-bit ADC input data stream is first de-serialized and subjected to a digital **Finite Impulse Response (FIR)** filter with eight coefficients. This effectively compensates for the nonlinear transfer function of the long cables (due to bandwidth limitation) and removes moderate reflections (due to imperfect termination). In *raw mode*, these data are sent to the output. Otherwise, APV25 data frames are detected in the data stream and extracted, followed by re-ordering of the strip data — which is necessary because the order is twisted due to the 3-level output multiplexer of the APV25 chip. Only the strip data, but without further processing, are sent out in *transparent mode*. In *zero-suppressed* and *hit-time-finding modes*, the pre-loaded pedestals of each strip are subtracted.

Common mode noise is removed by subtracting the average amplitude of strips from each individual strip value (“Common Mode Correction”, CMC). This CMC can be performed either in groups of 32 (default), 64 or 128 strips, i.e. in quarters, half or full chips, and it is performed in two stages. In the first round, masked strips (typically noisy or known bad ones) are excluded to avoid abnormal distortion of the average value. In the second pass, also strips with a signal above the hit cut threshold (normally three times the RMS noise of each strip) are excluded in addition to masked strips for obtaining an unbiased average. During the CMC calculation, the strip values are stored in 10 bits, with a range of -512 to 511 . Finally, the data are zero suppressed, discarding all strip data below a certain threshold (by default three times the strip noise). The results are presented as unsigned 8 bit numbers (range $0 \dots 255$).

In terms of monitoring, the FADC measures the values of each of the eight sensor bias currents and the temperatures of the two NTC probes on the Junction Board it is connected to. Moreover, it




Figure 34. Functional blocks of the data processing in the FADC firmware. Depending on the mode setting, data is processed according to one of the four columns shown (Raw, transparent, zero-suppressed or hit-time-finding). Reproduced from [23]. 2018 IOP Publishing Ltd and Sissa Medialab. CC BY-NC-ND 4.0.

probes the APV25 supply voltage levels (2.5 V and 1.25 V nominal) near the APV25 chips on each hybrid board using sense lines. All that metering is performed with slow 24bit ADCs which are read out over VME at a rate of approximately 1 Hz.

The processed data of the [FADC](#) are passed on to the [FTB](#) boards attached to the J3 connector at the rear of each [FADC](#). A busy handshake is implemented towards the DAQ chain to avoid transient data loss. In case of persistent busy, this can lead to back pressure that eventually blocks further triggers because ultimately the APV25 internal buffers become full.

Special runs are used to determine operational parameters such as the optimum clock phase for each ADC or FIR filter coefficients. They are very similar to the ones used in the common APVDAQ test system and are described in section [3.1.4](#) and [5.2.1](#).

2.5.2 FTB boards



Figure 35. FTB board.

The [FTB](#) is a 1U board that transmits the [FADC](#) data to the [COPPER](#) board and [DATCON](#) module (figure 35). The left side of the board is connected to the J3 connector on the rear side of the [FADC](#) board, and the 32-bits data stream is received by the Xilinx Spartan-6 FPGA (XC6SLX100 T). In the FPGA, the data are duplicated and sent to the [COPPER](#) board and to the [DATCON](#) module outputs. At each data output in the FPGA, the data are serialized using 8b/10b encoding and transmitted using a Gigabit transceiver in the Spartan-6 FPGA. For the serial-data-transmission to the [COPPER](#) board, a dedicated protocol called *belle2link* is used. The serialized data rate of *belle2link* is 1.27 Gbps, which is fast enough that all the [FADC](#) data can be sent without buffering in the [FTB](#) board. The *belle2link* protocol also allows remote register access to the FTB board through the [COPPER](#) board. For the serial-data-transmission to the [DATCON](#) board, the Xilinx Aurora module is used for the data transmission. For the error detection, the CRC16 checksum of the [FADC](#) data is checked, and also the [FTB](#) board calculates another CRC16 checksum for another check downstream. The [FTB](#) board also receives the 127 MHz Belle II DAQ clock and level-1 triggers from the [FTSW](#) board through the RJ45 connector.

2.5.3 Power supplies

The power supply design is based on fully floating low and high voltage sources with a common ground connection made on the detector to avoid ground loops. The APV25 readout ICs require

two regulated positive voltages, 1.25 V and 2.5 V, while the sensors require a bias voltage, called HV, ranging from 40 V to 200 V. Bias to the sensors is provided by two independent 100 V floating voltage supplies on the two sides: HV_P with negative polarity on the *P*-side, and HV_N with positive polarity on the *N*-side. The two HV supply return lines are connected to ground at the detector side, not at the power supply side, as shown in the simplified schematic of the power supply connections in figure 36. The hybrids on which the APV25 chips are mounted are referenced to the HV potential of the sensor side (*P* or *N*) to which they are connected, thus limiting to a few volts the potential difference between the input of the APV25 chip and the sensor bias applied to the strip implant, i.e. the voltage across the strip coupling capacitor. The two APV25 supply voltages are generated by radiation-hard DC/DC converters installed on the Junction Boards (section 2.2.5) located inside the detector. These converters are powered by a regulated floating Low Voltage Power Supply (LVPS), that should be adjustable between 5 and 12 V to optimize the DC/DC converter efficiency.




Figure 36. Power supply schematics.

Since the input to the APV25 chip stays at approximately 0.8 V, in addition to the main power supply system described above, an additional floating voltage source is required, called **Separation Voltage (V_{SEP})**, slightly offsetting the APV25 reference w.r.t. the sensor bias voltage. The V_{SEP} is used to adjust the APV25 input voltage to prevent current flowing into the pre-amplifier in the case of broken AC coupling capacitors on the sensors (pinhole compensation). The V_{SEP} polarity can be reversed to allow full diagnostic on the faulty capacitors (see section 3.1.4 for details). This polarity inversion is realized with a relay in a separate hardware component, the Power Distribution Panel (not shown in figure) which also allows the matching of connectors between the cables from the detectors and the power supply front panels.

For reasons of cost and dimensions of the cable plant, hybrids are connected in parallel in groups, and powered from the same LV source, with at most four hybrids in every group. Sensors in two hybrid groups (at most eight hybrids) share the same bias HV. Hybrid groups are defined following certain requirements: same readout direction (forward or backward); same sensor side (*P*

or N); same detector layer; same sensor type (wedge and rectangular). These considerations lead to a total number of 96 hybrid groups, 60 in the forward direction and 36 in the backward direction, and therefore 96 low voltage channels, 48 HV channels, and 48 V_{SEP} channels. Considering the power consumption of the APV25 chips, the efficiency of the DC/DC converters, and the expected leakage current of the sensors, the LVPS channels should be able to deliver a power of at least 20 W between 5 V and 10 V, while the HV channels should supply 1 mA at 100 V. The V_{SEP} channels are required to deliver only a few volts, with virtually zero current. To allow sufficient setting accuracy, the actual V_{SEP} voltage is obtained through a resistive divider ($\times 10$) from the power supply voltage.

2.5.4 Power and data cables

Three types of cables lead to the Junction Boards: LV power (Leoni [93] LEHC 002054, 17 m), HV power (3M HF659/10, 17 m), and data (Amphenol [94] 125-3097-998, 13 m). The first one contains four conductors with a cross-section of 1.5 mm^2 each, plus two shielded pairs, which serve as sense wires, while the cable is originally meant for a fieldbus installation, delivering power and data. The HV cable is actually a 10-conductor flat cable rolled up and shielded. It delivers the bias voltages and the offset voltages between HV and LV (V_{SEP}) for both P - and N -sides, thus utilizing 8 of the available wires. Thin conductors (28 AWG) are sufficient here because the currents are tiny ($\ll 1 \text{ mA}$).

In order to avoid any impedance disruption, the data cable is of a similar type as the hybrid cable (see section 2.2.3), but with 68 instead of 50 conductors (because two hybrid data are merged into one cable at the Junction Board). Unlike the hybrid cable, it is a flat cable (with twisted pairs), which cannot be laid in curves. Thus, to lay this cable on the outside of the CDC end-walls, a machine was built that removes the plastic cover on the outside of the flat cable by heat, makes the cable round and wraps Kapton tape around it. The first 6 m of every data cable are converted into a round cable in this way in order to cover the entire CDC region even for the longest cable path. All cables are flame retardant and free of halogens.

2.5.5 Grounding

As shown in figure 37, the local grounding point of the SVD is the end-flange of the BWD side. It is connected to the metal structure of Belle II with a thick grounding cable. The FWD side end-flange is connected to the local ground point through the aluminum foil on the outer-cover. Each end-ring is connected to the end-flange with a copper braid, and the same applies to each [origami pipe](#). All cooling pipes are electrically separated from the outside world by ceramic insulators located at the end-flanges.

Each ladder is electrically and thermally coupled to end-rings on both BWD and FWD sides, and ground wires lead from the end-mount to each hybrid and [origami board](#), where they are connected to the HV return lines of both the P - and N -side bias voltages.

The whole SVD is inserted into the CDC, but electrically isolated from CDC, PXD and also the beam pipe. Similarly, the DOCK boxes are electrically decoupled from the CDC, but connected to the corresponding end-flange with individual ground cables. Therefore, all on- and off-detector electronics are connected (in a star topology) to the central SVD ground point.




Figure 37. SVD grounding scheme: the BWD end-flange acts as the central grounding point for the SVD. The FWD side end-flange is connected to ground through the metal coated outer-cover.

2.6 Beam test validation of design

A number of test beam experiments were performed to develop and validate the various parts of the SVD system. In table 6 the main experiments performed since 2004, the beams used, and the goal of the test are listed.

3 Detector construction

The validation of the detector design by a series of tests on prototypes (section 2.6) is the starting point for the SVD construction, presented here. This process involved several institutions and companies, and required a complex organization as well as a careful and continuous quality control to ensure the integrity of the final detector. The timeline spanning several years, the logistics across several sites tracked in a construction database, the optimization and standardization of procedures and tests for quality control and assurance are outlined in section 3.1.

The silicon sensors were produced by two manufacturers, Hamamatsu Photonics K.K. [82] and Micron Semiconductor Ltd. [83]. Section 3.2 recalls the sensor specifications and describes the tests and full characterization performed on each device. The resulting quality classification was used to select the sensors for the assembly of ladders, installed in the experiment or kept as spares.

The procurement of the components of the on-detector electronics (APV25, pitch adapters, hybrid and origami boards, junction boards: section 3.3) required similar attention to details, and remedial actions in a few problematic cases.

The assembly of sensors and front-end electronics in modules and ladders for different detector layers had to proceed in parallel at different sites, under the responsibility of several institutions. As described in section 3.4, this process required the development of common assembly techniques and procedures, standardising assembly jigs, as well as alignment, gluing, wire bonding, and electrical

Table 6. Belle II SVD test beam experiments since 2004. The “JP” modules where an early design described in [3].

Date	Site	Device under test	Comments
Aug 13–18, 2004	CERN SPS T7	Triplet (single-sided)	
Apr 6–11, 2005	KEK PS	Triplet	UV module, doubled-sided readout
Aug 10–12, 2005	PSI piE1	Triplet, MWPC	
Aug 26–29, 2006	PSI piM1	Flex	
Nov 2007	KEK Fuji hall	Triplet, JP-single, Flex	SVD2.5 readout system (NECO, MAMBO, REBO)
Jun 2008	CERN SPS H6	JP-ganged, JP-single, Micron, Triplet, Flex	
Nov 2008	KEK Fuji	JP-ganged, JP-single, Micron, Triplet, Flex	
Aug 2009	CERN SPS H6	Origami 4'', Micron	First test of origami concept
Oct 2010	CERN SPS H6	Origami 6'' module, Wedge module, Micron baby	Including irradiation at SKC-CEN Brigitte [24], first full size origami module with B2 HPK sensor
Oct 2011	CERN SPS H6	3× origami 6'' modules	First Origami flex with high density JAE connector [84]
Oct 2012	CERN SPS H6	2-DSSD origami, 2× origami, Wedge	Including irradiation at SKC-CEN Rita [24], CO ₂ cooling
Jan 2014	DESY TB24	L3 module, 3× origami 6'' (L4-L6)	Together with PXD, installation started in Dec 2013, final FADC readout
Nov 13-24, 2014	CERN SPS	L3 module, 3× origami 6'' (L4-L6), SFW993, SBW993	First test with CAEN power supply
May/June 2015	CERN SPS H6	L5.903 + FADC readout	With EMC tests in Zaragoza, first completed L5 class B ladder
Apr 2016	DESY TB24	Phase 2 cartridge (first iteration)	
May 2017	DESY TB24	Phase 2 cartridge	

test procedures, including their quality assurance, in order to reach the required mechanical tolerances and detector performance.

The implementation of the innovative mechanical design is described in section 3.5: support end-cones, end-rings and an outer cover form the rigid support structure of the detector, with an embedded system of channels and pipes for the circulation of the dual-phase CO₂ coolant. Section 3.6 describes the fabrication and tests of the off-detector electronics: FADC and power supplies.

The final step of the detector construction, the so called “ladder mount”, was performed in a dedicated clean room at KEK, with the contribution of experts from several institutions, as detailed in section 3.7. A rotating mounting support was used to precisely position and mount the ladders shipped to KEK from the different construction sites, starting with the inner layer, for each of the two detector halves. Cooling pipes were installed with a special positioning device after the completion of each half-layer. Temperature monitoring sensors were also installed at this time. Complete functional tests on each ladder confirmed its performance before and after installation.

During the several years of SVD construction, distributed over several laboratories, many challenges appeared along the way, requiring considerable effort to be overcome. Section 3.8 gives a short report on these issues and on the lessons learned during the construction.

3.1 Organization and quality control

The SVD is composed of a large number of different components, described in section 2, that were designed, fabricated, and tested at different places and times. The integration, quality assurance and control of the entire system was challenging, requiring careful coordination and a dedicated [Quality Control and Assurance Group \(QCG\)](#), as discussed in section 3.1.2. The SVD construction effort required several years. Figure 38 gives an overview of the timeline for the R&D, fabrication, and construction of the various components.

An effort was made to minimize the number of different module designs and standardize the production, so that components could be reused as much as possible in different assemblies. Individual detector modules were assembled by gluing together a sensor, pitch adapters and hybrid or origami boards, which are equipped with APV25 readout chips; subsequently the electrical interconnections are made using ultrasonic wire-bonding techniques. Four basic detector module designs exist: Layer 3 modules; forward modules used in Layers 4, 5 and 6; backward modules used in Layers 4, 5 and 6; origami modules used in Layers 4, 5 and 6 (although the origami design depends upon the location).

Detector modules were then assembled into ladders using a hierarchical structure so that rework and reuse is possible until the last gluing steps. Completed and qualified ladders were shipped to KEK in preparation for the final ladder mount phase. In parallel, the mechanical support structure, the cooling circuits, the environmental and radiation monitoring sensors (described in section 4), and the off-detector electronics were prepared.

The ladders were then mounted on the two half-SVD mechanical supports, installing the required environmental sensors and origami cooling pipes. Each completed detector half was then commissioned and stored until the coupling with the PXD detector prior to the installation in Belle II.

3.1.1 Construction sites

The SVD construction was carried out by an international team distributed across several institutions and sites. A high-level summary of the construction responsibilities for the various components is shown in figure 38. Regarding the actual assembly and mounting of the detector ladders, it was carried out at five major sites, as shown in table 7. It should be noted that export regulations forbid the shipping of the radiation-tolerant APV25 chips to India. This is a key reason why the TIFR-Mumbai team, who had the responsibility of L4 ladder assembly, relocated their operation to the Kavli-IPMU clean room, that was expanded to host them.




Figure 38. Timeline and major responsibilities of the production of the SVD components and assembly. For Layer 4 (TIFR operating at Kavli-IPMU) the timeline includes a regular break of about three weeks during the prototyping and fabrication phase owing to the travel scheduling of the team.

Table 7. SVD assembly and construction sites.

Site	Activities
INFN — Pisa	Assembly of forward-backward modules
University of Melbourne	Assembly of L3 ladders
TIFR — Mumbai ^(*)	Assembly of L4 ladders
HEPHY — Vienna	Assembly of L5 ladders
Kavli IPMU — Tokyo	Assembly of L6 modules
KEK	Assembly of mechanical structure, ladder mount, commissioning and installation

^(*)operating at Kavli IPMU — Kashiwanoha, University of Tokyo.

The assembly of ladders for each SVD layer was assigned to a different site, so that the site could specialize and limit the number of different components to process. As a consequence, components had to travel between the construction sites, that resulted in a rather complex logistical organization.

3.1.2 Quality control and assurance

The completed SVD must be of high quality for Belle II to obtain its anticipated physics performance. It is therefore of the utmost importance that all the components and ladder assembly operations maintain the highest possible quality standard. This is essential at all steps of the process, particularly because the limited number of spare components makes the replacement of faulty or under-performing ladders difficult and expensive. Given the distributed nature of the assembly of the SVD ladders and the many groups involved, it was not straightforward to guarantee that the quality control and assurance procedures were similar at all sites.

It was therefore decided to constitute the **QCG** with the charge of defining and documenting all the assembly and quality control procedures, as well as verifying that the different sites were prepared for the work they had to carry out. The **QCG** was formed by site representatives and some additional expert members from the SVD group. The **QCG** carried out numerous site visits to verify the conditions at the sites, and to ensure maximal cross-communication between the experts at different sites. Besides regular meetings, the **QCG** held a number of focused reviews that resulted in written reports, which was very important with respect to identifying the points that required action, such as the safety and detailed logging of operations, the reproducibility of the gluing or bonding steps, or the storage and analysis of the electrical measurements.

Components were graded into classes according to their quality. Class A and B are fully working and usable components, with class B identifying lower quality components, for instance with more than 1% of defective sensor strips. Class C and D are mechanical components used for both assembly-procedure optimization and training. Class C are non-working real components that can therefore be used exactly as the real components in the assembly procedures, including wirebonding. Class D are instead substitute components, with the same mechanical dimension, for instance made of aluminum instead of silicon. Assembly operations were designed and optimized starting from class D components, progressively transferring the experience gained to class C, B, and finally class A components assemblies, to avoid wasting good components during the optimization stage.

Over the years, the QCG covered many topics, ranging from very detailed operations, like gluing or wirebonding procedures, to broader activities like the ladder mount or installation activities. The group's composition was adapted to the topic being considered at a given time, including also experts from other Belle II subsystems. The most relevant reviews are, in chronological order:

- Site qualification — verifying that the sites have the proper infrastructure and personnel.
- Class B assembly qualification — review the assembly procedures on electrically working but reduced quality modules, produced with some class B components. Passing this review was a pre-requisite to use class A components
- Class A assembly qualification — review the assembly procedures on an installation quality class A module.
- Class A monitoring — review over time the assembly procedures and the quality of the modules.
- Ladder mount operation — review the suitability and safety of the ladder mount operation.
- SVD commissioning — review the safety and appropriateness of the SVD commissioning.
- SVD operation — review the safety and organization of the SVD operation.
- Environmental and radiation monitoring — review the environmental and radiation monitoring systems.

Detailed assembly manuals were prepared for each assembly site, specifying all site-specific details, while following the common procedures. They included a standardized process flow diagram, describing the sequence of all assembly steps, and a set of detailed work instructions for each task. The completeness and correctness of the manuals was reviewed during the site qualification reviews.

The main function of the QCG group was to provide a unified view of the construction process, enforcing the use of a common workflow, procedures, electrical test setup, and quality assurance methodology.

3.1.3 Logistics for the mass-production

The complexity of the SVD modules and ladders, as well as the distributed nature of the assembly process, led to a complex logistic organization that required a unified tracking tool. To keep track of the actual location, assembly status and quality of all these individual components, as well as the assemblies made from them, a construction database named HephyDB [25], a web application based on the CakePHP [95] framework and utilizing the MySQL [96] database engine, was developed and routinely used at all assembly sites.

The basic element in the database is an item, assigned to an item type, which categorizes the individual items according to their features and functionality. The item type also defines whether the item is a simple component consisting of a single piece (such as amplifier chips and sensors) or a composite item built from other items (such as electronic boards, ladder sub-assemblies and so on). For each composite item type a parts list is defined, which has to be filled with corresponding items when the component is created. Thus, it is easy to trace the items used for any given assembly. For the assembly of more complex items, like ladders, in addition to the parts list, a checklist was defined, specifying the sequence of assembly steps and all the required measurements and tests to be performed during the assembly process. The checklist allowed the assembly to be monitored and ensured continuous quality control during the production of detector components. For each item and

assembly HephyDB provided a feature to upload files like pictures and documents as well as to store measurement results, which can be displayed in an online data browser. Finally, the database was also used to track the shipment of the components and assemblies between the assembly sites and KEK, so that the actual location of every component at any given time could be traced.

3.1.4 Common electrical test setup

The Belle II SVD front-end electronics were tested with a unified setup called the “APVDAQ system” [26] at various stages of construction, i.e. from bare hybrids and origami boards, over BWD and FWD sub-assemblies, to full ladders of all flavors. The purpose of these tests was to verify the overall electrical functionality as well as to diagnose specific issues like [pinhole](#) strip defects of the silicon sensor. The results were recorded after each test, and the defects compared to the corresponding results from the manufacturer’s sensor testing, so that any new defects introduced in the assembly process could be readily identified. The same system was replicated at each assembly site and at KEK so that measurements at different sites could be directly compared.

The APVDAQ system is a modular VME-based set of electronics modules developed and built by HEPHY Vienna to read out APV25 chips connected to silicon sensors. It was designed for laboratory and beam tests and consists of 6U VME modules connected to repeater boards that can be up to 30 m away. The repeater boards provide signal amplification and implement AC coupling for both control signals and readout data, allowing the system to work with double-sided silicon sensors (DSSDs), where front-end supply voltages are tied to the bias potentials. Furthermore, the APVDAQ system accepts a hardware trigger signal, typically generated by a scintillator attached to a photomultiplier when used with real particles, or by an external clock when used with a laser. Figure 39 shows the elements of the APVDAQ system when reading out a DSSD with four APV25 chips on each side.




Figure 39. A schematic view of the APVDAQ system in a standard configuration reading out a double-sided silicon sensor.

The APVDAQ system was adapted to the requirements of the SVD construction, and in particular a small [DAC](#) board was added to each sensor side to create a programmable offset voltage in the range

± 5 V between the sensor bias and the APV25 ground potential called V_{SEP} , which is a powerful tool to detect broken strip coupling capacitors.

The results of each run are saved to disk in various formats, including an xml file suitable for subsequent uploading to the construction database (see 3.1.3). Several run types listed in table 8 are defined, and normally executed in that sequence. Since the analog signals from the APV25 are sampled at a reasonably high frequency (40 MHz in the APVDAQ setup, while 31.8 MHz is used in Belle II) a careful adjustment of the ADC clock phase is required (ADC delay scan). In addition, a digital [Finite Impulse Response \(FIR\)](#) filter is applied to reduce signal distortion and reflection effects. The coefficients for the FIR filter are optimized in a dedicated run (FIR calculation). The pedestal and noise of individual strips are calculated in the Pedestal Run, while the gain is measured by injecting a known charge through test capacitors connected to each channel during the Internal Calibration (IntCal) Scan. In the IntCal vs. V_{SEP} Scan, the response of each channel is measured varying V_{SEP} from -5 V to $+5$ V, allowing a powerful defect-finding analysis, as described below. Finally, the Normal Runs are used to record actual sensor signals, that can be generated with a laser, a radioactive source, or a particle beam. For modules with a sensor attached, the sensor current-voltage characteristic is measured by the IV scan, varying the voltage between the N- and P-side from 0 to 100 V. Based on the APVDAQ system experience, the same run types were later implemented in the Belle II SVD readout system (see 5.2.1).

Table 8. Run types of the APVDAQ system.

Run type	Purpose
ADC Delay Scan	Find the best clock phase for each ADC
FIR Calculation	Obtain FIR filter coefficients for each input
Pedestal Run	Measure pedestals and noise of each strip
Internal Calibration Scan	Obtain the gain of each strip
IntCal vs. V_{SEP} Scan	Detect pinholes and other strip defects
Normal Run	Record radioactive source or laser hits

The results of the tests performed with the APVDAQ or FADC system (see section 2.5.1) were analyzed by a dedicated software package called *aDefectFinder* developed by INFN-Pisa. The *aDefectFinder* package analyses the output of the electrical tests in order to identify the defective strips, checking the noise, the time and amplitude of the APV25 waveform in response to the injected signal, and the response from the V_{SEP} scan output. Examples of defects and their response during the electrical tests are shown in figures 40 and 41. The possible defects and their appearance during electrical tests are the following:

1. Pinhole: the AC decoupling capacitor between the implant and the aluminum of a readout strip is broken. Since the APV25 input potential is at $+0.8$ V with respect to the strip implant voltage, a pinhole leads to a non-negligible current flowing into the APV25 preamplifier, saturating it; the APV25 channel exhibits low gain and low noise that prevents it from working properly, as shown in figure 40(a). In most of the cases the pinhole can be recovered applying a $V_{\text{SEP}} = -0.8$ V to remove the voltage difference between the strip implant and the APV25 input. This compensation effect is evident in figure 40(c), where the measured APV25 amplitude

output signal, in response to the charge injection in the internal calibration circuit, is shown as a function of the V_{SEP} voltage. In very few cases on the u/P side, the AC oxide layer is broken to the N-type substrate, instead of to the implanted P strip. These are called “substrate pinholes” or “bad pinholes” because of their detrimental effect on the entire sensor. Because to the reverse bias, the natural potential in the N substrate in that region is about +2 V, so the direct connection to the APV25 input with a substrate pinhole causes not only the saturation of the corresponding APV25 channel, but also a huge sensor current, that can only be compensated with $V_{SEP} > +2$ V. Since a single V_{SEP} setting is in common to many sensors, it was decided to compensate the majority of standard pinholes with $V_{SEP} = -0.8$ V, and to disconnect the very few substrate pinholes.

2. Short: two or more strips are shorted together, therefore the signal collected on one strip is shared with one or more other strips. During the calibration run, the charge injected in one APV25 channel is also shared with the adjacent channel, resulting in a lower gain for both adjacent channels, as shown in figures 40(b) and 40(d).
3. Open: the signal collected on the strip does not reach the APV25 channel input. Since the APV25 channel is not connected to the sensor strip, its noise is very low. But since in the electrical test a common mode correction is applied (see section 5.2.1), a single disconnected channel appears to have an extremely high noise corresponding to the common mode noise of the connected channels, as shown in figure 41(a) and 41(c).
4. Noisy strip: the noise of the strip is higher than that of regular strips, although not as high as that if a short, as shown in figures 41(b) and 41(d).

If any of the criteria tuned to identify the defects described above are matched, then the program classifies the defective strip accordingly. First it checks if the strip is a pinhole, by checking if the gain of the APV25 channel changes for different V_{SEP} values, as shown in figure 40(c). If the strip is not a pinhole, but has a low gain for all the V_{SEP} values, then it is classified as a short. If the noise is very large and above a given threshold it is classified as open, while if the noise is high, but below the threshold optimized to identify the opens, it is classified as a noisy strip. Specific values of the selection and classification criteria change with sensor types and locations and a detailed quantitative description is beyond the scope of this paper. It should be noted that in some cases strips exhibit a borderline behaviour, which makes the classification uncertain. A strip is also classified as a defect if, in spite of showing a normal electrical behaviour, during the laser or source scan it records a number of hits much lower (typically less than 50%) than the average number of hits observed in the neighbouring strips. The threshold for this defect classification is adjusted depending on the specific conditions at the different construction sites.

3.2 Silicon sensor fabrication and test

The silicon sensors, whose specifications can be found in table 1, were tested by the manufacturers as well as by the SVD group in different ways. The small and large rectangular sensors produced by HPK were fully characterized by the vendor, with only verification sample-testing performed by SVD groups. For the trapezoidal sensors produced by Micron, on the other hand, only basic tests were performed by the manufacturer, while the full characterization of every sensor was performed by the SVD groups. The tests performed on all DSSDs can be grouped into two categories:




Figure 40. Examples of plots analyzed by *aDefectFinder* for different types of defects. The defective strip is shown in red, while the adjacent strips are shown in green and blue. The upper row shows the analog output of the APV25 vs time, without applying any V_{SEP} voltage, the lower row shows the height of the peak of the APV25 output, used to estimate the channel gain, as a function of V_{SEP} . The first column, (a) and (c), shows an example of a pinhole response, with gain becoming equal to normal strips when $V_{SEP} = -0.8$ V. The second column, (b) and (d), shows a shorted strip, with low gain for two adjacent channels not depending on the applied V_{SEP} .

Global measurements. Overall current and capacitance measurements (IV and CV curves) are used to judge the overall health of the sensor and to measure the depletion and breakdown voltages;

Strip scans. AC (metal) and DC (implant) strip scans are performed on both P- and N-side to characterize each individual strip. The DC strip parameters, related to the quality of the wafer, implants, and bias resistors, are the strip leakage current, strip insulation with respect to its neighbors, and the poly-silicon resistance. The AC parameters are the coupling capacitance and the capacitor leakage current, related to defective AC coupling capacitors (pinholes) or to other defects in the metal layer, such as strip shorts and interruptions.

3.2.1 Rectangular Sensors

The HPK rectangular sensor production was carried out using 15 cm wafers over the course of several years. The first prototypes of the large rectangular sensors were delivered in March 2010 and used for design verification and ladder assembly procedures development. Following the




Figure 41. Examples of plots analyzed by *aDefectFinder* for different types of defects. The defective strip is shown in red, while the adjacent strips are shown in green and blue. The upper row shows the analog output of the APV25 vs time, without applying any V_{SEP} voltage, the lower row shows the height of the peak of the APV25 output, used to estimate the channel gain, as a function of V_{SEP} . In the first column, (a) and (c), the response of an open channel is shown, with a large noise due to the common mode subtraction (see text for more details). The second column, (b) and (d), shows a regularly connected noisy strip.

full characterization of the prototype sensors [27], in summer 2010 an origami flex module was assembled and successfully tested on a minimum ionizing particles beam by the HEPHY group, allowing the start of the DSSD mass production.

The full characterization of the DSSDs, which presents distinctive quality control challenges, was performed at HPK contacting the wafers only on the front side. It included global measurements and complete DC and AC strip scans on the P-side, where sensor biasing was provided through a front-side N-type bulk contact realized on the sensor P-side periphery, effectively contacting the back side. This technique is not viable on the N-side, since it is not possible to realize a contact from the N-side surface to the opposite side P strips. In any case on the N-side only AC strip scans, not requiring sensor biasing, are required, since the N strips individual currents are not important for detector characterization. The initial diamond-saw wafer dicing method was replaced in 2012 by laser dicing, leading to improvements in the quality of the sensor edges and in the bias IV characteristics. The production of the large sensors was completed in 2013, while the small sensors were delivered in 2014 and 2015.

3.2.2 Trapezoidal sensors

In January 2013, a contract was made with Micron Semiconductor (U.K.) for the fabrication of trapezoidal sensors to be installed in the forward modules of layers 4, 5, and 6. The first sensors were delivered in August 2013 and the original plan was to have the fabrication completed by the end of that year. However, it was realized that sensor characterization was a bottleneck in the production at Micron. To avoid excessive delays, it was decided to shift the time-consuming full strip scans to the clean rooms of HEPHY Vienna and INFN Trieste. A full back side contact technique was used in these laboratories, allowing complete DC and AC strip scans on both sides of the sensors. The final batch of sensors was delivered in September 2014.

3.2.3 Quality assurance of sensors

Similar quality assurance procedures were applied to both HPK and Micron sensors, although they had some evolution during the several years of production. Specifications included limits on the depletion voltage (< 120 V), on the leakage current (< 10 μ A), and on the bias resistor (> 4 M Ω), as well as on the percentage of defective strips as identified by the AC and DC strip scans.

All coupling capacitors were tested at 20 V. The average fraction of the AC failures (pinholes) was about 0.1% (0.03%) in the P-side (N-side) for the HPK sensors, and about 0.4% (0.2%) in the P-side (N-side) for the Micron sensors. The fraction of DC failures (shorts and interruptions) was in many cases as large as 5%, but it was found that most of these strips behaved normally after the sensor was connected to the readout system. Therefore, the impact of the DC defects on the sensor quality was calculated as an effective defect fraction, obtained dividing the actual defective strip fraction by a failure-mode-dependent factor (up to 10) estimated from the sensors that had been connected to the readout system. Sensor quality classification was done using the sum of the effective DC defect fraction summed to the full AC failure fraction. Sensors with an overall fraction of defective strips less than 1% were graded as class A. The rest of sensors were graded as class B and used for the sensor and assembly studies, as well as spare ladder production.

An overall summary of sensor production, their use in completed ladders, and their installation in the experiment is shown in table 9. It should be noted that although the overall number of class A sensors would have been sufficient to avoid installing lower quality sensors, the complexity of the assembly procedures, a few incidents and unforeseen challenges (like the peel-off one described in section 3.4.3) caused losses of modules or portions of ladders, forcing the use of a few class B sensors in the installed ladders.

Table 9. Summary of the DSSD production quantities and their usage in the assembly process.

Class	Usage	HPK Small	HPK Large	Micron
A	Total	26	188	52
	Used in ladders	24	170	48
	Installed in SVD	14	114	32
B	Total	2	27	35
	Used in ladders	2	17	12
	Installed in SVD	-	6	6

3.3 On-detector Electronics Fabrication and Test

The front-end electronics were built in a modular way, such that readout units (hybrid and origami boards) could be tested before being integrated with silicon sensors, thus avoiding loss of expensive good-quality items by attaching defective electronics. Known good dies from the APV25 production were assembled onto hybrid and origami boards. The connection from the APV25 to the sensor strips is realized with pitch adapter flexible circuits that come in several different designs as discussed in section 2.3.2.

3.3.1 APV25 fabrication and thinning

The APV25 chips were produced in 0.25 μm CMOS technology on 8" wafers with a thickness of $325 \pm 25 \mu\text{m}$ by an IBM foundry and subsequently wafer tested using an automated probe card setup at Imperial College, London. The overall production yield (fraction of good chips) was 88% [28], translating to more than 300 known good dies per wafer. About 75,000 chips were installed in the CMS Tracker, while almost twice that amount of good chips were produced in total, making some excess quantities available for other experiments like Belle II.

Chips mounted on the origami boards are located in the detector active region and therefore require thinning down to 100 μm , while thick chips are mounted on the hybrid boards. Since wafer level thinning is a well established procedure, several wafers were thinned and subsequently diced. The chips were mounted on the origami boards and connected to the first batch of PA0 pitch adapters, which turned out to be defective (see section 3.3.4). When the problem was discovered, all the remaining APV25 wafers with standard thickness were already diced, and the replacement of the defective origami boards required a non-standard thinning procedure applied to individual chips. In the first attempt it was observed that microscopic cracks appeared at the edges of many chips, sometimes causing malfunction. In the second attempt with a different company, the single-die thinning was successful, as verified by optical inspection and probe card testing, yielding a sufficient number of chips to populate a second set of origami boards. The numbers of the purchased and used APV25 chips are summarized in table 10.

Table 10. Summary numbers for the APV25 production. Thick chips are $325 \pm 25 \mu\text{m}$ thick, while thin chips were thinned down to 100 μm .

	Thin	Thick	Total
Purchased			≈ 7100
Used on modules	3330	2172	5502
Installed in Belle II	820	928	1748

3.3.2 Pitch adapter fabrication and test

A contract for the final fabrication of the pitch adapters was awarded in 2015. The quality of the produced pitch adapters was ensured by performing three main checks:

- Visual inspection at the pitch adapter flex production company, Tokai Denshi [97], to check the quality of materials, pattern defects, and dimensional tolerances.

- Electrical measurement of all lines on the pitch adapters. Quality control procedures require all lines to have low resistance, no interruption, and no short circuit to adjacent lines. The resistance of each line and the resistance between neighboring lines were measured to identify open lines and shorts. The electrical measurements were performed by Daiei Electronics [98].
- Dimensional verification of the bonding pad width to ensure bondability. The widths of each pad in the innermost row at the APV25 side, which is the narrowest in the design, were measured at IPMU with an optical [Coordinate Measurement Machine \(CMM\)](#) and were required to be larger than 30 μm . Additional visual inspection was performed at each assembly site.

The number of produced pitch adapter circuits is summarized in table 11.

Table 11. Approximate numbers of produced pitch adapter boards for the different uses.

	Layer 3	FWD/BWD modules	Origami	
	P3F1/2, P3B1/2	PF1/2, PB1/2	PA1/2	PA0
Purchased	120	400	400	270
Installed in Belle II	28	152	164	82

3.3.3 Hybrid board fabrication and test

The hybrid boards are PCBs composed of six layers with an overall thickness of 0.7 mm and were manufactured in industry with gold plated surfaces to enable aluminum wire bonding to the APV25 chips. The boards were electrically tested for open- and short-circuits with a needle probe tester at the manufacturer [99], with 16 hybrid boards combined on one single panel.

The assembly of electronic components was done at HEPHY Vienna. First, solder paste was applied using a manual stencil dispenser with the proper solder mask corresponding to the PCB panel. The passive components (resistors, capacitors and the connector) were then positioned with an automatic pick and place machine, followed by vapor phase soldering. After that, conductive glue was dispensed at the locations of APV25 chips using an adapter on the CMM. Then, APV25 chips were positioned using again the pick and place machine, with some additional manual alignment. After curing of the glue, the individual hybrid boards on a panel were separated and all APV25 chips on each hybrid were electrically connected to the PCB by aluminum wire bonding. For L3 hybrid boards, the hybrid cable was directly soldered onto each board in the last step, using a shield to protect the APV25 chips and their wire bonds. All hybrid boards were electrically tested using the APVDAQ setup (see section 3.1.4).

The numbers of produced hybrid boards are summarized in table 12. In this case, class A refers to a completely flawless hybrid board, whereas class B indicates fully functional boards with some minor defects, such as: some APV25 chips are slightly displaced from their nominal position; one of the (up to four) redundant wire bonds is missing; the amount of conductive glue is slightly extending beyond the normal area, without causing unwanted connections. Class C indicates not-functional, mechanical-quality boards that can be used for testing assembly procedures. Additional dummy boards were produced initially for training purposes, but those are not included in this list. A total of

180 hybrid boards are installed in the Belle II SVD (among which are 161 of class A and 19 of class B) — each of the 45 ladders holds a pair of hybrid boards glued back-to-back on either end.

Table 12. Summary of the hybrid board production.

Grade	Number
A	310
B	51
C	24
Total	385

3.3.4 Origami board fabrication and test

The first mass production of origami boards started in 2013, after full verification of a test module. About 150 such boards were produced in the first batch. After mounting APV25 chips and passive components, the boards were tested performing calibration with test pulses. However, a major problem was identified only after the connection to the sensor, with many channels showing no signal at all. Using a microscope for visual inspection, many tiny cracks of the copper lines on the PA0s, i.e. the flexes connecting the sensor N side strips to the APV25 (see section 2.3.2), were found, localised near the bonding pads at the sensor side, where the curvature of the copper lines is largest. All the inspected origami boards (about 20) had cracks affecting about 20% of the copper lines, mostly on the bottom layer of two PA0 layers. Due to the very high crack fraction, and the unfeasibility of any repair work, a second mass production of PA0 pitch adapters and origami boards was required, as well as a new thinning campaign for the APV25 chips, as described in section 3.3.1.

After detailed investigation of the production procedure several weak points were identified: the design of the copper lines, the cutting method of PA0s, and the lack of individual lines electrical testing. In addition, the glue used between the two PA0 layers was changed from acrylic to epoxy type out of concerns of radiation tolerance, leading to an increase in stiffness.

The problematic points were solved by improving the production and testing process as well as the sign-off procedures used to define reviewed and approved specifications: the design was improved to be more tolerant to mechanical stress; the original acrylic glue, more flexible and sufficiently radiation tolerant, was used; a laser cutting method was introduced; electrical measurements of each individual line resistance and isolation from the neighboring lines were performed; fully documented specifications including all the design details, tolerances, and testing procedures were prepared, formally signed-off, and forwarded to the manufacturer.

The final mass production started in 2015 at Taiyo Industrial [100]. Electrical tests were performed by Taiyo for origamis and Daiei Electronics [98] for PA0s, selecting only circuits with no failures. Gluing of the PA0 on the origami, the assembly with surface mount passive components and APV25 chips, as well as the wirebonding of the chips digital and power lines, were performed in companies, with frequent visual inspection steps, while the final inspection and electrical test on the assembled origami boards were performed at KEK.

During the first batch of mass production it was discovered that blisters were produced on origami flexes in the automated reflow soldering process. Although there was suspicion that the

problem could be due to an excessive water content between the various origami layers, a workable solution could not be found, and hand soldering was adopted instead of reflow soldering, to avoid the blisters.

Table 13 summarizes the mass production numbers. Origami boards that passed all the inspections and electrical tests with no bad channels are selected to be class A for the ladder assembly. Of these boards 82 are installed in the experiment. Class B identifies origami boards with defects observed in the visual inspection or some minor non-compliance of electrical specifications. Class C boards have significant failures, such as dead APV25 chips or assembly defects, and are used only for mechanical tests.

Table 13. Summary of the origami board production.

Grade	Number
A	116
B	18
C	12
Total	146

3.3.5 Junction board fabrication and test

The junction board PCBs, which arrange the APV25 signals and produce the 1.25 V and 2.5 V supplies for the APV25 by using air-core DC/DC converters [18], were manufactured and assembled in industry, where they were also subjected to optical inspection, but no functional test. At HEPHY Vienna, they were equipped with the mounting/cooling bracket and the DC/DC converter modules. Then, each board was tested using a commercial cable tester (Adaptronic NT-210, see left side of figure 42). This device measures resistances and capacitances as well as diode characteristics between any combination of its 512 I/O pins and thus in two passes compares a whole junction board connected to it to a previously recorded reference board.

Before their assembly on the junction boards, each DC/DC converter module was characterized individually by recording two IV curves with a dummy load using a Source Measure Unit (SMU, Keithley 2410), once with its ENABLE signal on and once off, in order to verify their healthiness. After installing them on the junction board, the individual output voltages (nominally 2.66 V and 1.44 V) were measured using a dedicated metering board that can switch between the default and redundant converters (see right side of figure 42), allowing a complete test of the system and verifying that the correct voltage variants were installed.

3.4 Module and ladder assembly

Modules and ladders were produced with responsibilities shared among SVD collaborating institutions: **SVD Forward module (SFW)**s and **SVD Backward module (SBW)**s by INFN Pisa, Layer 3 ladders by the University of Melbourne, Layer 4 ladders by TIFR-Mumbai, Layer 5 ladders by HEPHY-Vienna. The Layer 6 ladders were produced by a consortium comprising Kavli IPMU, University of Tokyo, Science University of Tokyo, Niigata University, Nippon Dental University,




Figure 42. Left: junction board (bottom) under test with the cable tester (top). Right: voltage display (bottom) for a junction board under test (top). The switch on the bottom allowed selecting between the default and redundant converter for a complete test.

Tohoku University, KEK, Kyungpook National University, and Seoul National University. The TIFR-Mumbai team shared the clean room at Kavli IPMU because of restrictions on the export of radiation hard components to India.

The modules and ladders were assembled using the same techniques across all the production sites. In particular, since Layer 4, 5 and 6 ladders share the same ladder structure they were fabricated following very similar assembly procedures. In this section, techniques and procedures common to all assembly sites are first presented, followed by the module and ladder production procedures for the Layer 3 ladders, the forward and backward modules, and the Layer 4, 5 and 6 ladders.

3.4.1 Common techniques and procedures

Clean room setup. All the module and ladder assembly procedures were carried out in a controlled environment to avoid contamination and improve the reproducibility of the operations. The assembly areas were required to maintain an ISO8 cleanliness level, equivalent to class 100000 of US Fed. Std. 209E [101]. The environment of the clean rooms was controlled and/or monitored to ensure constant conditions, needed to ensure dimensional stability of the parts and reproducible glue deposition and curing characteristics. Typical values were (23 ± 1) °C for the temperature and (50 ± 3) % for the relative humidity. Clean rooms were equipped with site-specific anti-static protection systems including anti-static floor, mats, bracelets, boots; proper grounding of the operators and equipment was routinely verified.

Assembly jigs and alignment. A large number of assembly jigs were developed in the different sites and were either fabricated in the local workshops or outsourced to precision mechanical machining companies. Depending on the use and required precision, aluminum or stainless steel was used. All the surfaces that would have been in contact with the sensors or other delicate components were fabricated in Delrin® [102] or Teflon™ [103] to prevent scratches. Inside most of the jigs, vacuum channels or holes were realised to allow vacuum chucking of the sensors. The jigs were all

qualified with a touch-probe CMM to ensure compliance with the required tolerances, which were set between 10 and 100 μm , depending on the mechanical elements to be held and on the length of the jigs. Precision coupling between different jigs was realized using commercial pins and bushings or precision holes. Particular care was devoted to avoid any interference between the jigs and the module wire-bondings, especially on the side of the module that is lying on the jig, which is not visible and cannot be visually checked. Many assembly steps were carried out under a CMM that allowed a precision alignment of the components and the measurement of their position. Micrometer screw gauges were used in many locations to adjust the position of the various components. Silicon sensors have an “F-shaped” alignment mark in each corner which allows automatic detection and position measurement by the CMM. For the positioning of the sensors, a dedicated XYZ Θ -stage, allowing a 4-axis adjustment of the sensor position under the CMM, was developed and produced in several copies for the different ladder assembly sites.

Component gluing. The different components of the modules and ladders are for the most part glued to each other. The Araldite 2011[®] [104] epoxy glue, whose radiation hardness had been confirmed by the LHC experiments [29], is used throughout the assembly process. The position and spread of glue deposited below the pitch adapters are very tightly controlled using optimized glue deposition patterns, to ensure a sufficient glue spread below the pitch adapter that makes a solid foundation for wire-bonding pads. Under-flown glue could leave an empty region under wire-bonding pads, while over-flown glue from the pitch adapter could spoil bonding pads near the pitch-adapter edge. Both conditions could result in wire-bonding failure. The typical glue thickness under the pitch adapters was 50 μm , controlled by the position of the gluing jigs with a precision of 5 μm to 10 μm . Each assembly site developed a special setup to control the glue position and spread. The typical glue curing time was 24 hours.

Wire bonding. The sensor strips, the pitch adapters, and the APV25 chips are electrically connected with 25 μm diameter aluminum wires using the ultrasonic wedge bonding technology. An example of wire bonding is shown in figure 43.

The requirements are set to a bonding efficiency greater than 99% and a pull force greater than 5 gf (gram force), as measured with a destructive pull tester. To satisfy the requirement, each assembly site optimized the parameters of its wire-bonding machine, such as the ultrasonic bonding power and time, the wedge pressure on the pads, the shape and height of the wire loop. Table 14 summarizes the total number of wire-bonds performed for the various layers at the different sites, considering all connections between sensor, pitch adapter and APV25 chips (128×2 per chip). The wire-bonds for the connections of the chip bias (18 per chip considering redundancy), signal, clock and control lines (48 per chip) were made at the hybrid or origami production sites, and correspond to an additional 26% with respect to the numbers in table 14.

Electrical test. Ladders and modules were electrically tested at almost every stage of the assembly process, as well as before and after their transfer to KEK, using the common APVDAQ readout system described in section 3.1.4. The APVDAQ test stand was complemented with a system capable of injecting charge in the sensor, equipped with stepping motors to scan the entire sensor surface: in Pisa and Melbourne an infrared light laser emitter was used, while in HEPHY and IPMU a radioactive β source was used. The response of individual strips was analyzed using the




Figure 43. Photo of the wire bonding between a pitch adapter (PA0) and an APV25 chip.

Table 14. Number of wire-bonds for the installed modules for the readout channels. Considering all the actually produced modules a factor 1.5 to 2 should be considered.

Site	L3	L4	L5	L6	Total
Melbourne	43008				43008
Pisa		51200	61440	81920	194560
TIFR		25600			25600
HEPHY			61440		61440
IPMU				122880	122880
Total	43008	76800	122880	204800	447488

aDefectFinder software package, described in section 3.1.4, to classify defects and compare them with the list of known defects from the sensors manufacturer or from previous testing steps. Such frequent and systematic testing allowed the early detection and often repair of defects, and prevented the use of problematic modules in the successive stages of the assembly process.

Procedure documentation and tracking. An extensive set of manuals was prepared, reviewed by the QCG, and made available to all collaborators. Each step in the procedures was described in the manual, with an associated checklist to be followed. The checklist was also implemented in the construction database and all the relevant information, such as optical inspection pictures, mechanical measurements, and results of electrical tests, were inserted in the database, so that any piece of information on the module parts could be recovered easily.

Equipment at assembly sites. Although the assembly sites had the same basic capabilities, the specific equipment differed from site to site, leading to slight differences in the procedures that had to be fully understood and documented. A summary of the equipment used in the assembly sites is shown in table 15.

Table 15. Summary of the equipment used at different sites.

Site	CMM	Glue robot	Glue dispensing	Wire Bonding
Melbourne	Mitutoyo QV-PRO302	SONY CAST Pro II	EFD 1500 XL	Hesse & Knipps BJ715M
Pisa	Mitutoyo BHN506	I&J Fisnar 750	EFD	K&S 8090
TIFR (*)	Mitutoyo QV-X606P1L-C	Musashi ShotMaster 3	Musashi ML-5000XII	F&K Delvotec 6400
HEPHY	Mitutoyo Euro-C776	Customized Mitutoyo CMM	Loctite 97006	F&K Delvotec 64000 G5
IPMU	Mitutoyo QV-X606P1L-C	SONY CAST Pro; JANOME SCARA	Musashi ML-5000XII	Cho-on-pa Kogyo REBO-7

(*) some of the equipment is shared between TIFR and IPMU.

3.4.2 Layer 3 ladder assembly

Layer 3 ladders are made by two small rectangular sensors connected to APV25 chips through pitch adapters. A description of Layer 3 ladders is given in section 2.3.2. The following procedure describes the main steps of Layer 3 ladder assembly.

At first, the Layer 3 hybrid boards of u/P and v/N sides are glued together to build the Layer 3 hybrid sandwich: the P-side board is positioned on the hybrid gluing jig, where it is glued with the N-side hybrid board using two countersunk screws to find the alignment. After 24 hours of glue curing, the hybrid sandwich is optically inspected and electrically tested. Since for Layer 3 hybrid boards no connector is used and cables are soldered, handling the layer 3 hybrid sandwich requires extreme care to avoid tensions caused by the permanently attached cables.

The assembly of a Layer 3 ladder requires wire-bonding and other operations to be performed on both sides of the sensors, hence N-jigs and P-jigs were produced. They can be coupled through precision pins, allowing picking-up a sensor and turning it upside down without losing the alignment. The jigs are divided in two halves handling the forward and backward sensors of the ladder. The two silicon sensors are positioned on the N-jig, then the XYZ θ -stage is used to align the sensors under the CMM, using the F-marks of sensor corners as a reference, as shown in figure 44 a). The main reference system is determined using reference points of the assembly base, which is a support structure where the N-jig is positioned. After the alignment, the N-jig is lifted, the two P-jigs are positioned on the assembly base, then the N-jig is lowered again on top of the P-jigs, where the sensors are transferred, with their u/P side up. Figure 44 b) shows the N-jig lifted up after sensors have been transferred to P-jigs. For each P-jig, the P-side pitch adapter is glued on the sensor and

the hybrid sandwich. P-side pitch adapters are positioned on a dedicated jig, and glue is dispensed using a robot. Figure 44 c) shows the P-jig with the hybrid sandwich, the sensor and the u/P side pitch adapter already glued, and the P-side pitch adapter chuck used to perform the gluing. After glue curing and optical inspection, P-side wire bonding and electrical test is performed. Each sensor is then picked up with one half of the N-jig, and is turned upside-down to prepare it for the N-side pitch adapter gluing, which is performed after properly adjusting the height of the hybrid sandwich. Figure 44 d) shows the N-jig with one v/N side pitch adapter glued, and with the other under the N-side pitch adapter chuck, during gluing operation. After glue curing, the two halves of the ladder are optically inspected, wire bonded on the N-side and electrically tested. At this stage, two half-ladders are ready for the assembly of a full Layer 3 ladder.




Figure 44. Some steps of Layer 3 ladder assembly: a) alignment of DSSD using XYZ θ -stage; b) positioning silicon sensors on the P-jigs; c) gluing of P-side pitch adapter between DSSD and hybrid sandwich; d) gluing of N-side pitch adapter.

The two half-ladders are positioned P-side up on the basement-jig. Forward and backward aluminum bridges are positioned and fixed on the corresponding hybrid sandwiches. The ribs are aligned and glue is dispensed on the stand-offs and in the rib slots of the forward and backward bridges. The ribs are then lowered on the P-jig until the stand-offs are at $\sim 100 \mu\text{m}$ from the sensor and the two ends are inserted in the respective slots of the bridges. After 24 hours of glue curing, the ladder is completed. It is optically inspected and electrically tested, then it is finally stored in the transportation box and prepared for the shipment to KEK.

3.4.3 Forward - backward modules

The following procedure applies to both forward and backward modules. Two different gluing jigs are needed for the P-side and N-side gluing of the pitch adapters. Each jig consists of: a tower with a vacuum chuck holding the silicon sensor on its Teflon surface; a second tower where the double-sided hybrid board is fixed by screws; a transparent vacuum chuck holding the pitch adapter positioned on top of the two towers. All the operations are done under a CMM. All the manipulation and transfer of the modules are done with several different types of vacuum chucks specifically designed to match the relevant dimensions and avoid damage to wirebonds.

The silicon sensor is placed P-side up, aligned by a mechanical stop on its tower, and the reference system is defined using the CMM software. The hybrid sandwich is then placed on its tower and aligned to the reference system. The pitch adapter is positioned on its transparent vacuum chuck exploiting two reference holes, and horizontally aligned to the detector-hybrid pair by means of micrometer screws. The pitch adapter chuck can then be removed, turned upside down, and placed under a glue deposition robot, keeping the mechanical reference to the detector-hybrid pair. After the glue-line deposition on the pitch adapter, the chuck is returned to its position. The joint is realized by raising the sensor and hybrid towers to the nominal height, filling with glue the 50 μm gap between the pitch adapter and the detector-hybrid plane. The transparency of the chuck allows to see the spreading of the glue, both on the detector and the hybrid sides. An optical survey after glue curing verifies that the glue line has spread uniformly below all the bonding pads of the pitch adapter.

The module is transferred to a P-side multi-purpose chuck, which can be coupled with precision pins to an N-side multi-purpose chuck to turn the module upside-down. Vacuum is used in order to hold safely the assembly. After flipping, the module is transferred to the N-side jig, where a second pitch adapter gluing operation is repeated for the N-side. After the final survey of the glue spread, the module is wire-bonded, using dedicated P-side and N-side wire-bonding jigs. The module is then transferred to the multi-purpose chuck, which allows both vacuum chucking and mechanical clamping of the module and is used for storage, electrical tests, and shipping. A full electrical characterization is performed (see section 3.1.4): I-V curve; pedestal and internal calibration run; VSep scan. To test the complete functionality of the module a scan under an infrared laser is performed. The modules are then shipped to the various assembly sites for Layer 4, 5 and 6 ladder production.

The production consisted of 60 SFW + 56 SBW modules and was completed in 2017. With respect to the nominal 47 pairs needed to assemble all ladders, some extra spare modules were produced to cope with the possible failures during the construction of the ladders. All the built modules showed good quality:

- the achieved mechanical precision of sensor position was better than 100 μm ;
- very few defects (<0.1%) were added during assembly with respect to the initial sensor defects (at the level of $\approx 1\%$ or better);
- the typical signal-to-noise ratio was measured to be ~ 24 for the u/P side and ~ 30 for the v/N side for APV25 calibration pulses.

About 10% of the SFWs developed a high bias current after the connection to APV25 chips, which was traced to the presence of substrate pinholes, discussed in section 3.1.4. Disconnecting the faulty strip from the amplifier input (by removing the wire bond), the modules could be recovered losing only one strip.

When the module production was already well advanced, in some of the completed Layer 4 and Layer 6 ladders the P-side pitch adapter started to detach from the forward sensor. The root cause of this failure, called “peel-off issue”, was traced to a limited overlap between the forward sensor and the pitch adapter (only 1.3 mm), leading to a weak glue joint. When the SFW is rotated to its nominal angle in the ladder assembly, this glue joint is subject to a large mechanical stress that in some cases caused the detachment of the P-side pitch adapter. Although the joint robustness was tested in Layer 5 ladder assembly and no peel-off happened, it turned out that the mechanical conditions in Layer 4 and Layer 6 were significantly more severe. A delicate emergency reinforcement procedure was applied to all SFWs, consisting in the insertion of a thin glue-covered Mylar® [105] strip between the P and N-side pitch adapters near the sensor edge, resulting in a much stronger joint.




Figure 45. Backward module P-side gluing jig.

3.4.4 Layer 4, 5, and 6 ladder assembly

The procedure for the assembly of Layer 4, 5 and 6 ladders is complex and requires a large number of steps, which are summarized in figure 46. The main points, which are common to the three layers, are described in this section in some detail to allow an understanding of the entire process.

Layer 4, 5 and 6 ladder are composed of a forward module, one origami module with one to three sensors, and a backward module connected together by the rib sub-assembly, as shown schematically in figure 18. The main difference in the assembly procedures between the layers is due to the different number of sensors in the origami module, as shown in table 16.




Figure 46. Process flow of main steps of the layer 4, 5, and 6 ladder assembly procedure. Reproduced from [30]. © IOP Publishing Ltd and Sissa Medialab. All rights reserved.

Table 16. SFW/SBW and origami modules used in the ladders of different layers.

Ladder	SFW module	Origami_+Z	Origami_CE	Origami_-Z	SBW module
L4	✓			✓	✓
L5	✓		✓	✓	✓
L6	✓	✓	✓	✓	✓

The Layer 4, 5 and 6 full ladder assembly procedure can be separated into the following six steps:

1. Pre-production of [Sensor-Pitch Adapter sub-assembly \(SPA\)](#).
2. Pre-production of [Rib Sub-Assembly \(RSA\)](#).
3. Sensor alignment of one to three SPAs, one SBW, and one SFW.
4. Half-ladder assembly.
5. Origami module assembly.
6. Integration of half-ladder and origami module.

In the following, these six steps are described in some detail, but without trying to account for slight differences of procedures at the various production sites.

1. Pre-production of [Sensor-Pitch Adapter sub-assembly \(SPA\)](#)

Two pitch adapters, PA1 and PA2, are glued onto the P-side of the large rectangular silicon sensors, and then electrically connected to the sensor strips by wire-bonding. This sub-assembly is called SPA and is produced in advance, to save time in the ladder assembly process. The SPA structure is shown in figure 47. The SPA is assembled with the following steps: (a) The DSSD is placed on the DSSD-jig P-side up. The jig has a Delrin[®] plate with a vacuum circuit inside, allowing the sensor to be held by vacuum chucking. (b) PA1 and PA2 are placed, with their backside up, on the PA-jig using two alignment pins, and held by vacuum chucking. Glue is deposited on the backside of PA1 and PA2 using a programmable glue dispensing robot. (c) The PA-jig is flipped and positioned on top of the DSSD-jig, using two guide shafts to find the correct

alignment between the parts. (d) After a few minutes, the vacuum of the PA-jig is released and the PA-jig is lifted, leaving PA1 and PA2 on the surface of the sensor. A visual check is made to spot any possible glue overflow on the pitch adapters bonding pads. After waiting 24 hours for glue curing, the wire-bonding is performed. The SPA is then stored into a Gel-Box™ [106] container, which is kept in a desiccator until it is used for the ladder assembly.




Figure 47. a) SPA assembly and components. b) RSA assembly and components.

2. Pre-production of Rib Sub-Assembly (RSA).

The rib sub-assembly is the support structure of a ladder. Like the SPA, it can be produced in advance, in order to speed up the ladder assembly process. It is composed of two truss-shaped ribs, made by a 1 mm thick AIREX® sheet sandwiched by two 0.1 mm thick CFRP plates, a forward and a backward end-mount. The RSAs are different for each layer, while the end-mounts are all the same and are the mechanical interface between the ladder and the end-rings. The RSA assembly is shown in figure 47. A dedicated jig for each layer is required to produce the RSA, but the procedure is exactly the same: (a) Both end-mounts are positioned on the jig, using a pin to find the correct alignment. (b) Ribs are placed in the corresponding slots of the end-mount and aligned on the plane defined by the jig. (c) Glue is deposited in the slots of the end-mounts, and a stainless steel pin is inserted to reinforce the glue joint.

3. Sensor alignment of SPA(s), SBW and SFW.

The alignment of SPA(s), SBW and SFW is the first step of the ladder assembly. The precision of the alignment determines the mechanical quality of the full ladder. The alignment is performed under the CMM, using the XYZ θ -stage to position the sensors with a 10 μ m precision with respect to the nominal position of the sensors fiducial F-marks. Figure 48 shows the SPA(s), SBW and SFW alignment and pickup procedure: (a) The SBW/SFW and SPA(s) are placed on the assembly bench using jigs and procedures designed to avoid collisions. The assembly bench is moved under the CMM, where each sensor is precisely aligned using the XYZ θ -stage, which allows to move the sensor in two directions and to rotate it around the z-axis. (b) After the alignment, the SBW and SFW are picked up with two jigs and are prepared for the half-ladder production (described in step 4), while the SPA(s) stay on the assembly bench for the origami module assembly (described in step 5).




Figure 48. a) SPA(s), SBW and SFW placed on the assembly bench and aligned under the CMM using by the XYZθ-stage. b) SBW and SFW picked up for half-ladder assembly.

4. Half-ladder assembly.

The SBW and the SFW are glued on the RSA to make a half-ladder. The steps of the half-ladder assembly are shown in figure 49: (a) The assembly bench is lifted up and put aside. The rib-jig is placed on the basement-jig. The RSA is positioned on the rib-jig. Then height blocks and height shafts are placed. (b) Glue is manually dispensed on the ribs, where the SBW and the SFW are glued. The SFW hybrid board is bent to its design angle with respect to the sensor surface and held in this position with a clamp. Then the pick-up jigs holding the SBW and the SFW are placed on top of the rib-jig, to glue the sensors to the ribs. The pick-up jigs are coupled to the basement-jig with precision pins, so that the initial alignment is kept through the assembly steps. After glue curing, the SBW and SFW hybrid boards are fixed to the RSA end-mounts, then the vacuum is released and the SBW and SFW pick-up jigs are removed. The half-ladder is optically inspected and electrically tested, then it is stored until its integration with the origami module.




Figure 49. a) assembly bench removed from the basement jig. b) SBW/SFW glued on the RSA.

5. Origami module assembly.

The origami module assembly is shown in figure 50. (a) The assembly bench with SPAs is placed back on the basement-jig. (b) Glue is dispensed on the sensors with a programmable robot, then an AIREX[®] sheet is glued on the sensors, using a dedicated jig to find the correct alignment. The origami boards are glued on top of the AIREX[®] sheet, again using dedicated jigs to align them to the sensors. (c) Wire-bonding of the N-side is performed between each sensor and the corresponding origami board, as well as between the PA0 and the APV25 chips. (d) After an optical inspection of the N-side wire-bondings, both PA1 and PA2 of each sensor are wrapped and glued on top of the origami board. (e) After glue curing, the final step is the wire-bonding between wrapped pitch adapters and P-side APV25 chips. (f) The origami module is optically inspected and electrically tested before proceeding to the integration with the half-ladder.




Figure 50. a) The assembly bench positioned on the basement jig. b) Origami boards gluing. c) PA1/PA2 wrapping from DSSD P-side to the top of the origami boards.

6. Integration of half-ladder and origami module.

The final step, shown in figure 51, is the integration of the half-ladder and the origami module. (a) The half-ladder on the rib-jig is placed on the basement jig and glue is dispensed on the ribs. The assembly bench is then placed on two electrically controlled coupled vertical stages and slowly lowered until the gap between the sensors of the origami module and the ribs is around 100 μm . After glue curing, the tails of the origami boards are fixed on the SBW and SFW hybrid boards. (b) The ladder is picked up from the assembly bench. A mechanical survey is done under the CMM to assess the mechanical quality of the ladder, then an optical inspection and a full electrical test are performed before storing the ladder in its container. The ladder is then shipped to KEK to be mounted on the SVD support structure.

3.4.5 Summary of module and ladder production

The total numbers of produced ladders are summarized in table 17. They include the class B ladder used in the Phase 2 commissioning cartridge described in section 6.1.2 Each completed ladder was mechanically surveyed. The position of the fiducial F-marks of the silicon sensors was measured and compared to the nominal position, in a local reference frame defined for the ladder assembly, where the x-y plane corresponds to the sensor surface, with the origin in center of the backward kokeshi pin, and the z-axis is perpendicular to the sensor surface.




Figure 51. a) Integration of half-ladder and origami module. b) Completed ladder.

Table 17. Summary of the ladder production quantities and their usage in the installed SVD.

Class	Usage	L3	L4	L5	L6
A	Installed in SVD	7	10	12	16
	Spares	5	2	3	3
	Total	12	12	15	19
B	Total	2	2	2	2

Table 18. Summary of the ladder mechanical precision. Average and maximum deviations, in μm , from the nominal position of the sensor fiducial F-marks in the completed class A ladders.

Layer	Δx		Δy		Δz	
	Aver.	Max.	Aver.	Max.	Aver.	Max.
L3	13	42	28	78	80	173
L4	55	327	37	145	67	478
L5	69	254	65	317	79	342
L6	40	233	45	199	62	234

A summary of the ladders mechanical precision is reported in table 18, where the average and maximum deviations for each axis are reported. Comparing the average deviations with the target tolerances presented in section 2.3.1, it can be observed that the overall achieved precision is well within the requirements for all axes. The maximum deviations are higher than the required tolerances, but the occurrences of these deviations are very rare exceptions that do not affect the mechanical quality of the ladders. Furthermore, as long as the ladders are mechanically stable, the final position resolution is determined by the alignment with tracks, as discussed in section 9.7.

3.5 Mechanical support and cooling fabrication

The overall mechanical structure of the SVD plays an essential role in guaranteeing the mechanical precision and stability of the ladder location and ultimately the detector performance. The cooling circuits are fabricated inside the mechanical structure in the case of the end-rings, while are applied on the already mounted ladders in the case of the origami pipes. The support cones and the outer cover were fabricated and assembled by KEK. The origami pipes were fabricated by HEPHY, while the pipe assembly technique and jigs were developed by Pisa.

3.5.1 Support cones, outer cover and end-rings

The support cone design is shown in figure 52. The cylindrical regions are for the gluing of the end-flanges and the end-rings. The CFRP cones are machined to allow 0.1-0.3 mm gap for the glue.




Figure 52. The CFRP support cone for (A) backward (B) forward region.

The end-mounts of the SVD ladders are fixed to the mounting surfaces of the end-rings. The heat produced by the readout chips flows through the end-mount to the end-ring and is removed by the CO₂ coolant flowing in the channel embedded in the end-ring. The channels are made by using the diffusion bonding technique, illustrated in figure 53.




Figure 53. The production of the end-ring with diffusion bonding. (A) The surface of the stainless plates are flattened and polished. (B) Grooves are milled on the surface of both plates. (C) The two plates are stacked face to face and diffusion bonding is performed at high temperature and pressure. Metal crystals are developed at the border resulting in a strong connection between the surfaces. (D) The plates are thinned to their design thickness. (E) Machining to the final shape. (F) CO₂ tubes are brazed in a vacuum oven.

Because of the slanted sensor in Layer 4, the supporting points in Layer 3 and Layer 4 are so closely spaced that the layer 3 and layer 4 end-rings are combined into one body, as shown in figure 54.




Figure 54. The end-rings for the Layer 3 and Layer 4 in the backward direction. (A) The end-ring in its final shape, after bonding; (B) exploded view of the individual layers before bonding, showing the CO₂ channels.

Figure 55 shows the brazing scheme of the CO₂ tubes and the end-rings in the backward region. Since the stainless steel tube has an outer diameter of 1.7 mm and a wall thickness of only 0.1 mm, a special brazing jig is prepared to tightly fix the end-rings at their final distance. The full VXD assembly procedure is designed to avoid stress on the cooling tubes, maintaining the relative position of the end-rings until they are glued to the support cones.




Figure 55. The brazing of end-rings and CO₂ cooling tubes shown in the front and rear views.

The design parameters of the end-rings are shown in figure 56 and table 19. To allow the overlaps of sensors in each layer, the ladders are arranged in a wind-mill like shape. Parameter D_{offset} is the offset of the support point of the ladder from the symmetric position. An angular offset A_{offset} is introduced to optimize the clearance between sensors in neighboring layers. The end-rings and support cones are designed for a 0.2 mm gluing gap, ensured by the end-ring gluing jig that positions the parts at the proper distance during the gluing procedure.




Figure 56. Parameters describing the shape of the end-ring. N denotes the number of ladders in a layer.

Table 19. Production parameters of the end-rings. Note that L3 is mounted on the same ring as L4 and has therefore the same R_{in} .

Layer	3	4	5	6	Unit
$N_{ladders}$	7	10	12	16	
R_{in} (B)		41.2	54.2	70.2	mm
R_{in} (F)		40.2	45.2	63.2	mm
R_{mount} (B)	39.0	67.1	89.3	120.3	mm
R_{mount} (F)	38.0	48.8	66.3	85.3	mm
D_{offset}	4.16	10.63	10.97	19.1	mm
A_{offset}	5	10	8	4	°



Figure 57. Finished Layer 4, 5 and 6 (from right to left) pipes on their storing and shipping masks.

3.5.2 Origami pipes

The origami cooling pipes are made of stainless steel SUS304 (1.4301) [107] with a 1.6 mm outer diameter and a 0.1 mm wall thickness. Starting from 6 m long straight pipes, they are bent in a 3D fashion by using a dedicated jig for each left and right bending. Three different sizes of cooling pipes were prepared for Layer 4, 5 and 6. After bending, metal-seal 300 bar stainless steel connectors are brazed at each end. Figure 57 shows the bent pipes on their storage containers, also used for shipment. Before the shipment to KEK, each pipe was tested for tightness by immersing in a water bath and applying a pressure of 200 bars for 15 minutes. No leaks were ever detected. A good thermal contact between the pipe and the APV25 chips is realized by interposing a 1.0 mm or 1.5 mm-thick (depending on the location) layer of Keratherm [87], a soft foam with good thermal conductivity.

3.6 Off-detector electronics fabrication and test

SVD-specific off-detector electronics consist of the FADC data acquisition system and the power supplies. They were fabricated through a combination of external companies and in-house personnel, under the responsibility of HEPHY-Vienna for the FADC and INFN-Pisa for the power supplies.

3.6.1 FADC system fabrication and test

The FADC boards were developed by HEPHY-Vienna, and assembled commercially. The FADC boards were optically tested at this company, but not electrically. This is the reason why a test system had to be implemented to check all the components, interconnections, and external connections on these boards [31].

The testing procedure performed is partially automated. It can be performed by following a well-defined step-by-step instruction manual, without detailed knowledge of the hardware or firmware. Some hardware-testing equipment is connected to specific ports of the device under test, while some test firmware is programmed into the FPGA, controlled by a PC via the VME crate. The operator can see if the test is passed and study the plots taken during the test to identify faulty or marginal components. All measured values including the graphs are recorded for later reference. The first tests are thermal inspection, using an infrared camera, and power consumption, measured by the internal current monitors of the VME crate.




Figure 58. FADC board testing setup. In reality, this ensemble is in a VME crate with a controller, a customized backplane for P2 and other peripherals. Reprinted from [31], Copyright 2020, with permission from Elsevier.

A special hardware test set, shown in figure 58, was developed for the FADC boards together with firmware for their onboard FPGAs and software for the PC controlling the VME bus.

To test each external digital connection for conductivity and isolation against the others, the firmware sends 1024 signal pulses to one single output at different frequencies up to 40 MHz depending on the type of the interface hardware. Independent counters enumerate the signal changes on all inputs individually at the same time. The results are then read out and interpreted by the




Figure 59. Statistical distributions of the ADC output data with uncorrelated square-wave input from the FADC tester board. Eight individual channels are shown in different colors. See section 3.6.1 for explanation. Reprinted from [31], Copyright 2020, with permission from Elsevier.

software. If the input connected to the active output counted any other number than 1024, or if any other input counted anything else but zero, the test has failed, and the faulty line is reported. The same is then repeated for all other lines on these two connectors successively, and also for some digital signals of the front connectors like I²C lines or the power-good and enable signals of the DC/DC converters on the junction boards. This digital test only needs a few seconds per board.

The analog inputs of the FADC boards for the detector signals are tested in a different way. The testing hardware is connected to one of the four front connectors of the FADC board by a 68-conductor twisted-pair cable (figure 58, lower left corner). The main FPGA on the FADC board uses a dedicated testing firmware to generate several square-wave differential signals e.g. on the CLK and TRG lines, which are sent through the cable to the FADC tester board. There the signals are converted to analog levels using two voltage dividers, of which one or both can be selected through reed relays. The square-wave output signals of the voltage dividers then can be connected to any of the analog inputs of the current connector on the FADC board via a relay matrix controlled by the FPGA. These signals are distorted and slurred a little by the cable which has the same length as the signal cable in the real system, in order to get characteristic signal shapes. It is important to note that the frequency of those square-wave signals and the sampling clock of the ADCs on the FADC board are completely uncorrelated in this test. In that way, histograms of the ADC levels are recorded to obtain statistical distributions (figure 59) by performing (pseudo-) random equivalent-time sampling, as used in most spectrum analyzer devices. If the input were an ideal square wave, only two sharp peaks in the histogram for low and high states would be visible. However, as the bandwidth is limited and there is also some ringing, entries exist between (and beyond) those ideal states. These histograms are then automatically examined by the software, checking if they fit into envelope templates generated from known good boards (pass) or not (fail). Since the FADC board tester can verify one of the four front connectors of an FADC board at a time, it is connected to each of the four ports successively by the operator as instructed by the software.

The remaining lines on the front connectors are analog signal inputs for slow control measurements like voltage and temperature monitors. They are automatically verified by feeding them with voltages generated by different voltage dividers also switched by reed relays on the FADC tester board. They are converted to 24-bit digital data by sigma-delta ADCs on the FADC boards which the software reads and compares to the nominal levels. The JTAG connectors are tested using a USB blaster to flash and verify the EEPROMs of the two FPGAs (which are additionally verified by rebooting in the procedure). The onboard JTAG flasher is checked by flashing the Stratix FPGA via the Cyclone FPGA, and the Gigabit Ethernet by sending and receiving packets to and from the PC using an Ethernet cable. On-board components like the delay chips are tested using sweep signals while recording histograms that are automatically analyzed afterwards.

In the next test phase one FADC after the other gets connected to a junction board and to a spare front-end detector module. The electrical isolation of the sensor bias voltage (“HV”) on the FADC boards, the cables, and the junction boards are tested using a ± 200 V power supply with current measurement, and all channels of the ladder are read out by the DAQ PC. This procedure (but without the HV power supply) is repeated later for each channel every time before the final detector gets connected to the system.

During the initial FADC tests at HEPHY Vienna in the spring of 2018, some component assembly errors were detected, like a couple of so-called tombstone SMD resistors and capacitors (where only one side of the component is soldered properly and the other one is up in the air), a few resistors soldered in wrong positions, and several SMD connectors where some pins were without solder. One clock distributor IC was rotated by 180 degrees, and one FPGA had badly soldered BGA pins. These failures were successfully repaired, and all 60 boards behaved well afterwards. On the junction boards, some minor errors were found, for example a loose SMD resistor jammed in between two connector contacts, and they were corrected on-the-fly.

No damage was found on the retest of the individual components after the shipment to KEK, and also the DAQ tests of the individual boards performed as expected. The whole readout system was successfully tested during the commissioning of the two separate SVD halves and later on after the coupling with PXD, when the full VXD was assembled.

3.6.2 Power supplies fabrication and test

The power supply system was based on commercial components and fabricated by CAEN SpA [108]. High density fully floating switching power supply sources are used, allowing the installation of the full power supply system in just two standard rack-mount 19” crates (model SY4527). The selected modules are specified in table 20.

Table 20. Installed power supply modules types. Technical documentation on the modules is available on the CAEN website [108].

Type	Model	Specifications	Channels/ module	Installed modules	Total channels
LV	A2519A	5–15 V / 5 A (< 50 W)	8	12	96
HV	A1519B	250 V / 1 mA	12	4	48
V _{SEP}	A1510	100 V / 10 mA	12	4	48

Pre-production boards showed a large noise when channels were connected in series as foreseen in the schematics. The problem was promptly corrected by the company with additional filtering. All the boards were tested with a resistive load mimicking the load in the experiment, measuring the noise spectrum. In addition, all channels were verified on a test SVD module, performing a noise run with the full FADC system. The noise is fully comparable with what was observed with linear power supplies.

3.7 Ladder mount and mechanical survey

The assembly of the two SVD half-shells is performed in a dedicated clean room at KEK. In this section, a description of the workflow is given, describing the clean room and the setup of the ladder mount table, the procedure to mount a single ladder and to attach the cooling pipe on layers 4, 5 and 6 half-shells, the mechanical survey performed after a layer half-shell is completed. After the whole half-shell is assembled, it is stored on a dedicated table, where it is kept for the commissioning period and until the SVD is combined with PXD to assemble the full VXD.

3.7.1 Clean room setup and ladder mount table

The clean room is divided into two areas. The larger area, $4.5 \times 4.5 \text{ m}^2$, is used for the ladder mount and the tests, with the ladder mount table placed in the center of the area. The smaller area, $3.0 \times 4.0 \text{ m}^2$, is used for the inspection of the ladders just before mounting, and for the mount tool preparation and storage. The cleanliness of the room is maintained at the required ISO8 level (equivalent to class 100000 of US Fed. Std. 209E) by three HEPA fan filters. The floor of the room is covered by an anti-static mat connected to the ground. The temperature, humidity, cleanliness, O_2 and CO_2 densities in the clean room are monitored and recorded continuously. O_2 and CO_2 sensors are installed in the clean room, to monitor their levels especially during the tests of the cooling system, operated with liquid CO_2 .

The ladder mount table, shown in figure 60, is designed to position the support structure that is used to assemble the SVD. The support structure, on which the end-rings are mounted, is equipped with a worm-gear that can rotate around the z-axis to allow mounting all SVD ladders from the same side of the table, with the procedure described in section 3.7.2. The worm-gear has a reduction factor of 50 to allow the setting of the angle of the end-rings with a precision of 0.001 degrees. The table accommodates a Coordinate Measurement Machine (CMM) system used to perform mechanical measurements before and after the ladder mount, as described in section 3.7.4. The CMM is located on the same side where the ladder mount attachment tool is positioned. On the opposite side of the table, the structure for the cooling pipe attachment tool is located. During the ladder mount period, this structure holds a cover that protects the already installed ladders, while after layer 4, 5 and 6 half-shells are completed, the cover is temporarily removed to perform the cooling pipe attachment, described in section 3.7.3.

After the cooling pipe of each layer is mounted, a cooling test is performed to check that there are no leaks in the cooling circuit and to verify the thermal contact between the cooling pipe and the APV25 chips using an infrared camera. Given that the half-shell cannot be moved from the table at this point, the test is performed on the ladder mount table, using a custom design open circuit CO_2 cooling system connected to the half-shell cooling pipe through temporary connections. After checking the connections for tightness, a small box is mounted around the SVD half-shell to create a dry environment with a low dew-point to avoid water condensation when the cooling test is performed at $-20 \text{ }^\circ\text{C}$. The volume is flushed with nitrogen for the whole duration of the cooling test.




Figure 60. A CAD drawing of the Ladder mount table.

3.7.2 Single ladder mount procedures

Some preparatory steps are needed on the ladder mount table as well as on the ladder to be mounted before the installation of each ladder.

Ladder preparation. Four plexiglass covers of different lengths and radii were prepared, one for each layer, to be installed just before the assembly of each half layer, to protect the ladders already installed. The protection covers are sustained by the same support structure used for the cooling pipe attachment, described in section 3.7.3. Ladders are installed from one side of the ladder mount table, moving horizontally towards the end-rings. The protection covers are designed to leave enough space for the ladder mounting, without being moved. On the table, the end-rings are rotated so that the planes of the two end-ring supports on which the ladder is mounted are perpendicular to the surface of the table. A measurement of the [Kokeshi-pin](#) hole and of four points for each end-ring support is performed with the 3D CMM system, to verify the position and orientation of the plane on which the ladder is installed. The tolerance on the planes orientation with respect to the nominal angle is set at ± 0.15 degrees for both θ and ϕ angles. It turned out that this tight tolerance could not be reached for L3 rings after gluing the Layer 3 supports to the Layer 4 ring. Nonetheless, such a reduced tolerance on Layer 3 planes is acceptable because the larger gaps between Layer 3 modules requires less demanding accuracy to avoid mechanical interference.

Meanwhile, the ladder is optically inspected in the dedicated smaller area of the clean room. The goal of the inspection is to verify that the ladder has no mechanical issues: after removing the top and side walls of the ladder container, all wire-bonds, pitch adapters gluing points and CO₂ clips are checked. If the optical inspection is successful, Layer 4, 5 and 6 ladders, still fixed on the container supports, are brought to the Keratherm [87] attachment station, while Layer 3 ladders are ready to be mounted.

The Keratherm, used to ensure proper thermal contact with a cooling surface, has to be positioned on the surface of APV25 chips of the origami boards on which the cooling pipe is mounted. Extreme care must be taken during the installation of Keratherm strips on top of APV25

chips of the origami boards because they are placed very close to the APV25 chips wire-bonds, so any mistake in the operation can lead to wire-bonding damage, that could compromise the functionality of APV25 chips in the ladder. To properly attach Keratherm, a specific procedure was established: four Keratherm strips are cut, as in figure 61 (a), and are then positioned on a jig that holds them by vacuum chucking. The ladder is positioned as in figure 61 (c), tilting the ladder container so that the top surface of the ladder is facing the operator. Keratherm strips are aligned on the Keratherm-jig, and the protective film is removed to expose the adhesive layer. The Keratherm-jig is then positioned on a movable stage with four degrees of freedom (X - Y - Z - ϕ). The alignment of Keratherm strips with APV25 chips is verified with the help of a magnifying glass, then the movable stage is brought forward until Keratherm strips touch all APV25 chips. Finally, the vacuum is switched off, the Keratherm jig is removed and the plastic film on the other surface of Keratherm strips is removed. The procedure is repeated for all origami boards of the ladder. Once Keratherm strips have been attached, all wire-bondings of the ladder are optically inspected again to verify that no damage has occurred during Keratherm attachment. After this last optical inspection, the ladder is ready to be installed on the end-rings.




Figure 61. Main steps of Keratherm attachment on an origami board of a Layer 4 ladder. (a) the Keratherm strip is cut into four pieces; (b) the pieces are positioned and aligned on the Keratherm-jig; (c) the Keratherm-jig is positioned on the movable stage that is used to align the strips with the ladder; (d) Keratherm strips are attached to the APV25 chips of an origami board.

Ladder installation. For Layers 4, 5 and 6 ladders, hybrid cables are tested, assigned to each connector of the ladder and labeled on both sides, to easily recognize which cable is attached to

which connector. For Layer 4 ladders, cables must be fed through Layer 5-6 end-rings prior to the ladder installation. The last step of the preparation for Layers 4, 5 and 6 ladder, before their installation, is the verification of **Fiber Optic Sensor (FOS)** (see section 4.2) insertion, to be sure that the fibre can be properly inserted in the dedicated channel of the AIREX® layer.

The ladder container, without all side walls, is positioned on a table beside the ladder mount table. A ladder mount tool, shown on the left side of figure 62, is used to grab the ladder: made of aluminum, the tool has two screws and two corresponding nuts at the extremities that are used to fix the ladder on the FWD and BWD ends, two handles used by the operator to grab the tool, a knob connected to a gear that allows a rotation of the ladder around the X-axis, and a quick-release plate that is used to fix the tool on the ladder mount stage. The operator brings the ladder mount tool on top of the ladder, then two more operators fix the tool to each end of the ladder, engaging the screw and fixing the nut (figure 63 (a)). Now the pick-up screws that hold the ladder to two vertical supports of the container are released, so that the ladder can be extracted from the container. The main operator starts lifting the ladder, while another operator guides the extraction of kokeshi-pins from the container vertical supports. Once the ladder is extracted, the main operator brings and fixes the ladder mount tool on the ladder mount stage engaging the locking mechanism of the quick-release plate (figure 63 (b)). The ladder mount stage, shown on the right side of figure 62, is equipped with three micrometer screws used to move the ladder along the three axes. The whole stage block is mounted on rails on which it can slide towards the ladder mount table, with a vacuum lock mechanism that must be released by pressing a pedal. After all these operations, hybrid cables are temporarily connected to the ladder and an electrical test is performed to verify that the ladder has not been damaged during the operation.




Figure 62. On the left side, the ladder mount tools are shown. This tools are used to grab the ladders, allowing the operator to engage them on the ladder mount stage. See text for a step by step description of the procedures. On the right side, the ladder mount stage.

After the electrical test, the FWD hybrid cables are detached and the ladder is slowly moved towards the end-rings, sliding the ladder mount stage on its rails, while the distance to the end-rings is continuously monitored with the help of a ruler mounted on the rails. The stage is stopped when the ladder is 5 mm far from the end-rings, then the stage is locked on the rails. Any further movement is controlled by the micrometer screws on the ladder mount stage. A helper on the backward side takes care of pulling the still connected BWD hybrid cables while the ladder is moved

toward the end-rings, to avoid any entanglement with the end-rings and interference with any other previously mounted ladder. The main operator turns the X-axis micrometer screw bringing the BWD kokeshi-pin surface at a distance of approximately 1 mm from the end-ring; then turns the Z-axis micrometer screw to bring the ladder in the proper position in Z. The BWD kokeshi-pin head is aligned with the corresponding hole of the end-ring acting on Z and Y axes micrometer screws. Once the BWD kokeshi-pin is aligned, the position of the FWD one is adjusted using the knob on the ladder mount tool holding the ladder: the rotation axis of the ladder mount tool is close to the BWD kokeshi-pin location, so that only the FWD kokeshi-pin location changes while turning the knob. A fine adjustment of the alignment between kokeshi-pin heads and end-rings holes may be done if necessary, then the ladder is moved towards the end-rings using the X-axis micrometer screws so that the chamfered part of the kokeshi-pin heads are engaged in the end-rings holes. The main operator grabs the ladder from the FWD and BWD extremities, while two helpers on both sides release the nuts so that the ladder is released from the ladder mount tool (figure 63 (c)). The fixing screws of the tool, however, are still fixed to the ladder, so the ladder cannot fall even if it slips from the grab of the operator. The main operator inserts the ladder fully into the end-rings, then the helpers fix the set screw using a torque controlled screw driver set to 50 N·cm. The fixation screws of the ladder mount tool are released (figure 63 (d)), the ladder mount stage is brought back to its original position and the ladder mount tool is finally removed from the stage.

FWD hybrid cables are connected, then all cables on both sides are properly routed through the assigned slots of the end-flange. For Layers 4, 5 and 6 ladders, the FOS is inserted in the corresponding channel in the AIREX[®] sheet and the clamp is fixed with a set screw to the fiberglass surrounding the connectors on the BWD side. Finally, pictures of the installed ladder are taken and the reference points of each sensor are measured with the 3D coordinate measurement system. A full electrical test of the installed ladder, including a readout test of the FOS, follows. If the test is successful, the end-rings are rotated and prepared for the next ladder.

After the installation of all ladders of each SVD half-shell, the cooling pipe attachment and its test are done for Layers 4, 5 and 6 as described in the next section.

3.7.3 Cooling pipe mounting procedures

In order to mount the cooling pipes on a completed half SVD Layer 4, 5 and 6, a mechanical system, shown in figure 64, was developed. The system is positioned on the ladder mount table, on the opposite side with respect the ladder mount tool. On a thick base plate, fixed to the ladder mount table, two rails allow a coarse translation of the positioning tower (PT) toward the SVD half. Micrometer screws are used to drive the fine movements along the x, y and z axes, together with a rotational stage. The two-arms support (TAS) holds a cradle, on which the cooling pipe is fixed through plastic clamps. The TAS can be rotated to align the horizontal plane of the two arms with the plane of the ladder mount table. The cradle can be wrapped around the half SVD layer by means of a coarse mechanism and a fine micrometer screw. In the following, general procedures that are common to all layers, are summarised.

1. Mechanical System setup.

The cart is in its starting position, far from the SVD. The PT-TAS system is rotated by 180°, pointing outside of the table. The cradle is detached from the TAS.



Figure 63. Key steps of the ladder mount procedure. (a) the ladder mount tool is attached on the ladder using the fixing screws and the nuts; (b) the ladder mount tool is fixed on the ladder mount stage; (c) after the alignment of kokeshi-pin heads with the corresponding holes, the ladder is inserted in the end-rings; (d) the ladder is released from the ladder mount tool. Note that in the actual ladder mount procedures gloves were worn at all times to avoid sensor contamination. Some pictures were taken in early stages of the procedures development, with mechanical samples that required less attention.

2. Pipe pick up.

The cradle is positioned on the cooling-pipe container, where the cooling pipe is fixed with some screws. The cooling pipe is released from the container, one straight line at a time, and inserted in the corresponding clamps of the cradle. Once the cooling pipe is fully attached, the cradle is positioned on the TAS, engaging the cantilever parts of the TAS. The PT-TAS-cradle system is then rotated toward the SVD.

3. Approach the SVD.

The cart moves the PT-TAS cradle system towards the SVD, reaching a pre-determined safe position, using the SVD end-rings as a reference. This is the starting point of finer adjustments. Acting on x and y-axis micrometer screws, the cradle is moved so that the first line of the cooling pipe reaches the mounting position. If necessary, the cradle can be rotated in order to align the pipe with the ladder at the top. All movements performed by micrometer screws are carefully checked by three watchers, who make sure that the cooling pipe nor the cradle touch the SVD

ladders. The cradle is moved with the micrometer screw around the axis of the top cooling line, until it reaches its final position, where all the straight lines of the cooling pipe are at a radial distance of 1.5 mm from the clamps on the ladders.

4. Cooling pipe attachment.

Starting from the first line on top, the cooling pipe is released from the cradle clamps and engaged onto the ladder CO₂ clamps using the mounting tool, which is held by hand by an experienced operator. The mounting tool is at first engaged in the two rings of the ladder CO₂ clamp, then the cooling pipe is gently pushed down. This procedure is repeated for every clip on the ladders, following the path along the cooling pipe. Figure 65 shows the stage at which the cooling pipe is attached to the Layer 4 ladders, with the cradle still in its final position. Finally, after the cooling pipe is mounted on the SVD, the cradle is brought to its “open” position again and the PT-TAS-cradle group is moved away from the SVD using the cart.




Figure 64. Drawing of the mechanical system used to mount the cooling pipes.

The procedures were tested and tuned with mechanical prototype ladders for each layer, resulting in a cooling pipe mounting time of a few hours. Only in one case a connection of the cooling pipe was found damaged, and a detachment of the installed pipe was needed in order to replace it with a spare cooling pipe. No damage was observed on the ladders after each cooling pipe installation.

3.7.4 Mechanical contactless survey system and results

The mechanical precision of both end-rings and installed ladders is verified on the ladder mount table using a built-in CMM positioned on the side of the table where ladders are mounted, as shown in figure 66. The position of the precision holes for the kokeshi-pins has to be within 150 μm of the



Figure 65. Layer 4 cooling pipe attachment: after positioning the cradle on the Layer 4 half-shell, the pipe is inserted into the CO₂ clamps of the ladders.

nominal position and the angular deviation must be less than 1.5 mrad. The position of the sensors fiducial F-marks on each mounted ladder has to be within 400 μm of its nominal position, with a maximum rotational error of 2.5 mrad. These specifications are determined mainly by the safety requirement of avoiding interference between ladders.

The principle of operation. The CMM system consists of a **CCD** microscope camera (Schott SOD-10X Telecentric lens) with working distance of 55 mm and focusing depth of 17 μm , mounted on three-axes linear rails. Each axis is equipped with an absolute position encoder (Heidenhain LIP481, LIP581 absolute linear encoder) of < 0.1 μm precision. The end-rings are fixed to a rotating axis bar. The distance between the lens and the object is determined by using the CCD image. The rotation angle is measured with an angle encoder (Heidenhain ROC413 absolute angle encoder) with a resolution of 0.05°. The resolution of the coordinate measurement of the camera is better than 1 μm in all directions.

To cover the whole SVD structure, the length of the Z linear rail is 1 meter long. The overall measurement ambiguity of the system, mainly due to uncertainties in the mechanical setting and alignment, is estimated to be 50 μm .

When the position of an object is measured, the camera is first moved to the nominal position of the target by the CMM control program. Then the measurement operator adjusts the X, Y and Z coordinates by watching the CCD image on the PC screen. The final adjustment below 100 μm is done by a focus finding program. The resolution of the automated focusing is better than 1 μm , so the uncertainty due to different operators is avoided. After the final adjustment, the coordinates and the focus image are immediately stored to the PC.

To avoid accidents due to operation errors or hardware malfunctions, the CMM movement was constrained both by software limits and by a hardware safety system based on limit switches.




Figure 66. The 3D coordinate measurement machine system on the ladder mount table. The ladder mount arm is not shown here. (a) The CCD camera with working length of 55 mm, (b) (c) (d) X, Y, Z linear rails with position encoder, (e) End-rings, (f) Support rod, (g) Angle encoder, (h) Bearings, (i) Index plate and the worm gear for the rotation of (f), (j) (k) Limiter plate and switch for crash prevention, and, (l) A ladder.

Measurement procedure and results. Before a ladder is mounted, the rotation angle of the end-ring is set within 0.001° with respect to the nominal angle, so that the kokeshi-pins of the ladder are inserted straight into the mounting holes of the end-rings. Eight points of the end-ring surface and the center of the mounting hole are measured.

After the ladder is mounted, the coordinates of the four F-marks of each sensor are measured. The measured coordinates are compared with the nominal position to judge if the ladder is properly mounted on the support system. If the deviation is above the $400\ \mu\text{m}$ limit, the ladder mount procedure is examined, and if necessary the ladder is un-mounted. For the Layers 4, 5 and 6 ladders, the measurement is repeated after the CO_2 pipe is mounted, to check if the layer structure remains within the required tolerance.

The measured coordinates are compared to the measurements taken after the ladder assembly process, to understand if the ladder was deformed during the mounting process. Only in one case one layer 3 ladder was found deformed outside the defined tolerances, and it was replaced with another ladder. For all other ladders, no significant differences were found comparing the measurements before and after the ladder mount.

The measurements performed after the cooling pipe attachment show a slight deformation of the ladders due to the intrinsic force exerted by the cooling pipe. The biggest deformation is observed

on the slanted forward sensors and is of the order of a few hundreds of micron, still within the acceptable range and with no danger of interference with ladders of the other layers.

The coordinate measurements data are the starting point for the alignment of the SVD with respect to the other Belle II sub-detectors. The final position of the sensors in the Belle II detector are determined with real particle tracks from cosmic ray or beam collisions by using an alignment software, described in section 9.7.

3.7.5 SVD half-shell storage

After its completion, the first SVD half-shell had to be removed from the ladder mount table and stored safely in order to assemble the second half-shell. In addition, between the completion of both halves and the full VXD assembly, the commissioning of the detector had to be performed. For this purpose, two similar movable tables and two so called “pick-up tools” were used to store the two SVD half-shells.

Two 1 mm CFRP outer covers are designed to surround the two SVD half-shells, to protect the whole Vertex Detector (VXD) once it is fully assembled. Once an SVD half-shell is complete, the first step before dismounting it from the ladder mount table is the installation of the outer cover. On the inner surface of each cover, an optical fiber temperature sensor, similar to the one described in section 4.2, is attached. The outer cover is placed on top of the SVD half-shell and taper pins of 2.5 mm diameter are inserted in the corresponding holes to accurately align the end-flange and the outer cover. Screws are then used to fix the cover to the end-flange. Taper pins are finally extracted from their location.

The pick-up tool structure is shown in figure 67. It is positioned on a storage table equipped with wheels, most of it is made of aluminum, on the two side supports two rails with braking mechanisms are mounted, on which the tool can slide toward the SVD half-shell. The side supports are also equipped with six ball plungers that allow the movement of the whole tool between the tables. The storage table is brought beside the ladder mount table and raised up so that the surfaces of the two tables are at the same height. The whole pick-up tool structure is moved on the ball plungers and fixed on the ladder mount table, then the pick-up tool approaches the SVD half-shell until the two curved ends of the tool engage the SVD end-flanges on both sides. The SVD half-shell is fixed to the pick-up tool, then all cables are fixed to some supports installed on the pick-up tool, and finally the SVD is released from the support of the ladder mount table. Once the SVD half-shell is fully detached from the ladder mount table, the pick-up tool is moved back on the storage table, where it is fixed with screws. In the bottom part of the table, two containers are placed to store the cables. In order to protect the inner part of the SVD half-shell, which is exposed at this point, a protection cover is installed on the side. The table was stored into a box and constantly flushed with dry air until the commissioning started.

3.8 Issues and lessons learned during the SVD construction

The SVD construction was a complex and challenging task carried out over the course of several years. As in all large projects, many surprises, problems and difficulties appeared and were overcome, sometimes with considerable effort. In this section these issues and the relevant lessons learned during the construction process are briefly reported.



Figure 67. (a) The bare SVD pick-up tool on the storage table. (b) SVD half-shell on the storage table held by the pick-up tool, with the front cover installed.

3.8.1 Quality control of components

In several cases some components were not of sufficient quality, because of either insufficient design margins, or lack of proper quality control procedures. This required extended quality control procedures, or the fabrication of additional components, sometimes with design modifications possibly realized at very short notice.

For instance, initial pitch adapter productions had issues with the pad size and material control, leading to great difficulties in the wire bonding process and even resulting in the creation of cracks in the metal traces causing electrical connection failure. These problems were discovered in some of the first prototypes allowing an optimization of the pitch adapter design and production process.

Another example was the appearance of cracks on the APV25 chips after the thinning and dicing process. In spite of the positive electrical characterization of the devices, there was the concern that mechanical stress induced especially by large temperature variation could cause the cracks to extend leading to failures. A manual chip inspection program was hastily setup to identify the cracked dies and reject them.

On more than one occasion unexpected features of components were observed, e.g. various traces of material on top of the sensors, leading to investigations and analyses to determine whether the features were the indication of some problem, or rather a normal outcome of the production process.

The lesson learned from these problems is that quality control and assurance must be carried out and documented in the most systematic and extensive fashion possible; experience indicates that what is not checked is likely to become problematic, which can be particularly damaging in systems where replacing single failing components is hard or impossible.

3.8.2 Assembly and integration issues

Ideally all the assembly procedures and final integration should be fully defined and tested with real components before production starts. In practice, this is often impossible and one has to test the procedures only piece-wise, maybe with surrogate components, opening the possibility that bad surprises appear in the middle of the production process, requiring retro-fitting and adaptation.

One important and potentially very damaging issue appeared during the construction of the ladders: the pitch adapters of some forward modules started to detach from the sensor after the assembly of the ladder. Although the pitch adapter gluing procedures had been tested and verified, it turned out that in layer 4 and layer 6 ladders the pitch adapter bending radius is significantly smaller, leading to much higher stress on the glue joint. An emergency retrofitting program was launched on all forward modules, including the ones already installed on ladders, with the development of a reinforcement procedure which was quite challenging and risky. In this case the problem was generated by the combination of a marginal design (very small glue area between pitch adapter and sensor) and the lack of a proper prior analysis and verification of what would be the most extreme usage situation of the components.

Another example of problems discovered only very late was related to the fixing of layer 3 modules to the support cones. The screw initially foreseen in the design had the potential of causing a short on the hybrid board since the ground plane was very close to the screw hole. An insulation system was developed and implemented. This example shows how crucial it is to perform final and complete integration tests with final components to spot design imperfections early on.

Ensuring the complete reproducibility of the assembly procedures has proven to be extremely challenging, since both the laboratory conditions and the personnel undergo changes that may result in small difference in the procedures, which should be avoided as much as possible. The Quality Control and Assurance Group described in section 3.1.2 was instrumental in ensuring that all the minute details were specified in the manuals and that the instructions were consistently and uniformly followed at all sites.

3.8.3 Project management

Running a large construction project requires considerable effort and must be properly organized and staffed. In general, very marginal contingency was used in planning the SVD construction activities, resulting in an often slipping schedule. Expert person-power was limited and there were many cases of single points of failure, namely operations that could be performed only by a single person. Preliminary risk assessment should be given serious consideration, to facilitate a timely and correct completion of the project. Ensuring a correct design requires great attention to detail, as well as a continuous internal and external review process. For instance, the initial design of the SVD had some serious mechanical incompatibilities. As another example, the Layer 3 cooling design was discovered to be insufficient very late in the process, thanks to simulations and a test mockup, leading to a last minute retrofit to avoid excessive temperature gradients. As a general observation, spares were insufficient at all levels, from the individual component to the complete ladder. Although the additional cost and effort needed for spares is sometimes difficult to justify, having a significant number of spares makes the project much more resilient to problems, such that it is more likely to be completed in time.

4 Environmental and radiation monitoring

Safe operation of the vertex detector (SVD and PXD) requires continuous and reliable monitoring of its environmental parameters, in particular radiation dose rates, temperature, and humidity. If their values happen to be outside accepted limits, an interlock system, independent of the computer

network and of the slow control software, must be able to promptly shut down the power supplies as well as to notify the central interlock system of the experiment. Environmental and radiation monitoring were introduced at a late stage in the VXD design. The solutions successfully adopted are briefly described here, including the requirements, the design and construction at INFN Trieste, and the installation at KEK.

The radiation from beam losses is monitored by a system based on diamond detectors, that also generates beam-abort requests to SuperKEKB in case of excessive beam losses; this function is essential to prevent severe local radiation damage to VXD parts. This system is described in section 4.1.

Temperature is monitored by [Negative Temperature Coefficient \(NTC\)](#) thermistors at the inlets and outlets of the SVD cooling pipes and by Bragg-type [Fiber Optic Sensors \(FOSs\)](#) close to the SVD front-end [ASICs](#), as described in section 4.2.

The dew point of the residual humidity in the dry nitrogen circulating in the VXD volume is kept under control via “sniffing pipes” connected to external humidity sensors. As explained in section 4.3, this is essential to prevent condensation or ice formation on parts of the SVD cooling system, based on dual-phase CO₂ circulating in thin pipes (section 3.5).

Section 4.4 describes the [Programmable Logic Controller \(PLC\)](#)-based hardwired interlock system, which protects both SVD and PXD by disabling their power supplies when local and general Belle II environmental conditions exceed acceptable limits.

Finally, section 4.5 briefly describes the working principles and design of the CO₂-based evaporative cooling plant.

4.1 Radiation monitoring and beam abort

The radiation monitor and beam abort system is designed to protect VXD sensors and front-end electronics from damaging beam losses and to continuously monitor the radiation dose rates in the interaction region. Detailed descriptions of the design, construction, calibration, and performance of this system are given in references [32, 33].

4.1.1 Requirements

According to extrapolations from the former B-factories, KEKB and PEP-II, and simulations of the expected beam backgrounds at SuperKEKB, the system was expected to measure dose rates from a few $\mu\text{rad/s}$ to several krad/s , depending on the beam currents and accelerator conditions.

The SuperKEKB beam-abort process, triggered by a large increase in beam losses and radiation dose rates, is constrained by the 10 μs revolution time of electrons and positrons in the SuperKEKB rings. The beam abort system should react in a few μs , before a full revolution of the beams. Localised radiation damage in SVD microstrip detectors may occur when a large radiation dose, in excess of about 1 rad, is delivered in a time interval of the order of 1 ms (see discussion in section 7.4.4). For the dose-rate monitoring, archival, and integration, a recording rate of 10 Hz is required.

A configuration of radiation detectors distributed on the beam pipe and SVD support cones was initially specified; a request followed to instrument also the bellows, flexible connections between the vacuum beam pipes of the two rings, and the beam pipe section inside the VXD, located close to the [final-focus superconducting quadrupole \(QCS\)](#) magnets. The present configuration includes a




Figure 68. A diamond sensor packaged into a detector unit, as described in the text. Reprinted from [33], Copyright 2021, with permission from Elsevier.

total of 28 diamond detectors: eight on the beam pipe, twelve on the SVD support cones, eight on the bellows, as shown in figure 69.

4.1.2 Detectors design and construction

Artificially grown single-crystal CVD (sCVD) diamond sensors are chosen for their fast and temperature-independent response as well as their radiation resistance. Electrons and holes produced by ionisation in the diamond bulk drift to respective electrodes located on opposite faces, when a bias voltage is applied to them. The induced current in the external circuit is proportional to the dose rate.

The electronic-grade sCVD sensors [109] have lateral dimensions (4.5×4.5) mm² and 0.50 mm thickness; two (4×4) mm² (100 nm Ti + 120 nm Pt + 250 nm Au) electrodes are deposited [110] on the two opposite faces. Each diamond sensor is mounted on a ceramic-like [111] printed-circuit board, providing back screening and electrical contacts, as shown in figure 68. The diamond sensor is glued to a square pad, which provides connection of the back-side electrode to a soldering pad. The front-side electrode is connected to the second soldering pad by ball-bonded gold wires. The inner conductors of two miniature coaxial cables, each 2.5 m long, are soldered to the printed board pads, while conductive glue is used to mechanically fix the cables and establish a connection of their outer conductors with the outer shielding of the printed circuit, completed by a thin Aluminum cover on the front side. The cable length of 2.5 m is fixed by the space constraints within the VXD set-up.

Each detector is characterised and calibrated as detailed in reference [32]. The quality of the detector assembly is assessed by measurements of dark current I as a function of bias voltage V ($I - V$). In all cases, the measured dark current is less than a few pA at 500 V and not more than about 1 pA at 100 V. The uniformity of the transport properties of the charge carriers and of the electric field in the diamond bulk are validated using the Transient Current Technique (TCT) [34], based on the localised injection of a controlled amount of electrons and holes at a very small depth, close to one of the electrodes, by monochromatic α -particles from a ²⁴¹Am source. For each detector the radiation-induced current-voltage relation is established with a constant electron flux from a radioactive ⁹⁰Sr β -source. At least one bias polarity shows a clear plateau and stable response, leading to the choice of the operating bias voltage at 100 V for the optimal polarity.

Finally the measured current is related to the radiation dose-rate and an individual calibration factor is determined for each detector. The flux of ionizing electrons from the almost point-like radioactive ^{90}Sr β -source is controlled by varying the distance between the source and detector. The measured current I in the detector polarisation circuit is compared with the dose rate dD/dt predicted by a detailed simulation of the entire set-up; a reference silicon diode is used to constrain the source activity and reduce the systematic uncertainties in the simulation.

The measured conversion coefficients k in the linear relationship $dD/dt = kI$ correspond to about 35 (mrad/s)/nA on average. They can differ from the average by up to approximately $\pm 50\%$ for different detectors. The uncertainty of about 8% in their determination is predominantly systematic [32]. Measurements with a strong ^{60}Co γ source confirm a linearity in the response to the dose rate over at least three orders of magnitude.

4.1.3 Electronics, monitoring and beam abort

Groups of four diamond detectors are controlled and read out by a purpose-designed [Diamond Control Unit \(DCU\)](#). The DCU digital core is a board hosting an [FPGA](#), that receives commands via an Ethernet interface, drives four [HV](#) modules independently through a [DAC](#), and accepts input data from front-end modules providing amplification and analog-to-digital conversion of the input signals. The DCU is also able to deliver VXD abort requests for the [HER](#) and [LER](#) beams, and receives the SuperKEKB abort signals. The current signals from diamond detectors are amplified by trans-impedance amplifiers and digitised by 16-bit [ADCs](#) at 50 Msamples/s. The FPGA firmware stores sums of 125 ADC samples every $2.5\ \mu\text{s}$ at 400 kHz in a large revolving buffer memory.

Further sums of five million ADC values can be read out at 10 Hz. The online control software converts these values from ADC to dose-rate units after subtracting pedestals; these data are archived and used for continuous monitoring. From the archived values, accompanied by time stamps, integrated doses are computed.

Three current-measurement ranges can be selected: 36 nA, $9\ \mu\text{A}$, and 4.5 mA. The first range allows precise monitoring of relatively small beam losses, while the third range avoids saturation in the detection of large radiation spikes, and is used for beam abort. The intermediate range may be used in the future, with increasing dose rates.

For each diamond detector, two moving sums of 400 kHz data are computed by the DCU firmware. The moving sums are updated at each $2.5\ \mu\text{s}$ cycle, by subtracting the oldest added value and adding the newest one. The two moving sums, representing integrated doses, are compared with programmable thresholds. The resulting eight logical signals in a DCU can be combined using programmable logic to produce an abort request, if the DCU is enabled to do so. In the present configuration of seven DCUs connected to the 28 detectors, only one DCU, controlling four detectors located on the beam pipe, is enabled to deliver abort requests.

Communication with the SuperKEKB beam abort system is based on four signals. The diamond system delivers two abort requests, separately for the electron [High Energy Ring \(HER\)](#) and the positron [Low Energy Ring \(LER\)](#). When the abort kicker magnets complete an abort that either originated from this system or from other sources, two “SuperKEKB abort” timing signals (HER and/or LER) are broadcast and received by the DCUs. These signals stop the buffer memory writing. The 400 kHz data can then be read out and used for a “post-abort” analysis of the beam losses preceding the abort.




Figure 69. Installation of diamond detectors: (upper) on the beam pipe; (middle) on SVD cones, forward and backward; (lower) on the beam pipe bellows. Reprinted from [33], Copyright 2021, with permission from Elsevier.

4.1.4 Installation

Four diamond detectors and prototype electronics were installed prior to the accelerator commissioning in 2017 (Phase 1), as part of the [BEAST II](#) detectors dedicated to background studies [35]. During a pilot run in 2018 (Phase 2) the final version of the electronics was tested with eight detectors in their final position on the beam pipe, as discussed in section 6.1.2. An additional 20 diamond detectors were delivered and installed at KEK in 2017 on the beam pipe (4 + 4), the SVD cones (6 + 6), while the eight detectors of Phase 2 were recycled to instrument the forward and backward beam-pipe bellows, as shown in figure 69. The 2.5 m short, thin coaxial cables are linked at the [DOCKs](#) (section 2.2.5) to longer and thicker double-shielded (25 – 30 m) coaxial cables, connected to seven [DCUs](#) in the Belle II electronics hut.

4.2 Temperature monitoring

The power dissipation by the [PXD DEPFET](#) sensors and their front-end electronics amounts to about 18 W per module, 360 W in total for the 20 PXD modules. The APV25 front-end chips of the SVD dissipate about 700 W in total during full operation. This power is removed by a cooling system based on heat exchange with a dual-phase CO₂ fluid at a temperature of about –20 °C, circulating in thin pipes with good thermal contact with the front-end readout chips, as described in section 3.5.

Temperature monitoring is therefore important to ensure proper operation of the cooling system, with a required absolute accuracy of 1 °C and a resolution of 0.1 °C. It is composed of two subsystems. The first system measures input and output temperature of the CO₂ cooling pipes and is based on [NTC](#) thermistors. The second system ([FOS](#)) monitors the temperature of SVD [ladders](#), with distributed measurement points close to the front-end chips, based on Bragg-type sensors.

4.2.1 NTC system design

For the first task 12 [NTC](#) thermistors are mounted on the outer surface of the SVD support [end-rings](#) (one sensor on each of the six half-rings) and 16 sensors on the cooling pipes (one for each inlet and outlet). Doubling the number of sensors for redundancy at each measurement spot brings the total number to 56 sensors in the SVD volume. Additional sensors are installed on the CO₂ transfer lines, in the docks chiller cooling system, and at cross-reference points with the [FOS](#) sensors, for a total of 84 sensors. The location of the sensors is shown in figure 70.

[NTC](#) thermistors have accuracy better than 1 °C. They require only two wires per sensor, since their resistance (typically 100 kΩ) can be chosen to be much higher than that of connecting cables. The readout of [NTC](#) sensors is performed by a custom-designed system, built by INFN Trieste. It is based on the [Embedded Local Monitor Board \(ELMB\)](#), designed and used in large numbers at CERN [36]. Each ELMB hosts a CANbus-connected [112] processor and an ADC multiplexed to 64 input channels. To optimize accuracy, half of them are connected to reference resistors; the remaining 32 input channels measure the voltage drop across the [NTC](#) thermistors. Each ELMB is mounted on a motherboard providing the current sources and hosting comparators with trimmer-adjustable thresholds, OR-ed in groups of eight, which provide four fast interlock signal outputs for 32 input channels. These signals activate the hardware interlock system (section 4.4) when the preset threshold temperature is exceeded. Three motherboards with their ELMBs are




Figure 70. Physical layout of the NTC system: the red dot pairs represent NTC sensors placed on half end-rings and on the inlets and outlets of cooling pipes.




Figure 71. Block diagram of the readout of the NTC system: one readout unit hosts three motherboards with one ELMB card and four interlock lines each, as explained in the text (see also section 4.4).

mounted together with power supplies in one unit, providing 96 input channels, and 12 interlock outputs in total. The main features of the NTC readout system are shown in figure 71.

4.2.2 FOS system design

The FOS system measures the temperature of the SVD ladders, close to the APV25 chips. For layers from 4 to 6, optical fibres are inserted in channels prepared in the Airex foam, with several sensors along the ladder in positions roughly corresponding to the presence of the front-end ICs on the

origami boards (section 2.2.4). Each sensor on a fibre is realized by a **Fiber Bragg Grating (FBG)** with a characteristic reflection wavelength, mechanically shifted by temperature-sensitive acrylate coating. The temperature is derived by the measurement of shifts in reflection wavelength, with an accuracy better than 0.6 °C.

In total 260 sensors are distributed on 44 fibres with nominal wavelength at 20 °C between 1522 and 1573 nm, distributed between the PXD (12 sensors), SVD layers and outer cover (248 sensors in total). For all contiguous **FBG** sensors, the separation in nominal wavelength is in steps of 3 nm. For cross-calibration, a subset of eight fibres have an additional sensor located outside the Airex foam channel, thermally coupled with an NTC thermistor.

The readout of **FBG** sensors is performed by an “optical sensing interrogator” [113], featuring a high-power, low-noise swept-wavelength laser. The interrogator has 16 input channels and an internal 4 × 4 multiplexer; external multiplexing is provided by 16 1 × 4 splitter/coupler modules. The association of fibres with external multiplexers and interrogator input channels is optimized to exclude overlapping nominal **FBG** wavelengths at the same input. The overall readout scheme is sketched in figure 72.




Figure 72. Readout scheme of **FBG** sensors for SVD: each **ladder** in **layers** 4, 5, and 6 hosts one fibre with the indicated number of sensors. Groups of four fibres are connected via (4 × 1) couplers to one of the 16 inputs of the interrogator.

For calibrations and debugging the Micron Optics ENLIGHT sensing analysis software [114] provides a single suite of tools for data acquisition, computation and analysis of optical sensor networks. A custom Micron Optics protocol via Ethernet is used for readout by the SVD run and slow control software (section 5.1).

4.2.3 NTC system construction, calibration and installation

The **NTC** thermistors, with 10 kΩ resistance at 20 °C, are radiation-resistant and halogen-free [115]. Their calibration is obtained by measuring their resistance at −30 °C, 10 °C and 50 °C, and fitting the Steinhart-Hart coefficients [116]. The obtained accuracy, better than ± 1 °C, is sufficient for

our purpose. At each step of SVD cooling pipes mounting procedure (section 3.7.3), the relevant thermistors and their cables are checked to ensure they are well connected and properly working. Halogen-free flat twisted-pair cables (SVD cables: section 2.5.4) were used to connect the sensors to intermediate patch panels at the VXD DOCKs and finally to the readout unit in the electronics hut (e-hut). Two spare readout units were built and used in parallel during construction and commissioning: one for the ladder-mount setup and the other for the SVD commissioning setup.

4.2.4 FOS calibration and installation

Each FOS was calibrated in a climatic chamber by determining the coefficients of a polynomial fit to the measured temperature-wavelength relation.

During the SVD ladder mount for layers 4, 5 and 6 (section 3.7.2) one fiber was inserted in each ladder and checked for proper functioning. Particular care was required to handle the fibers and store their excess length on rolls next to the detector, when they were temporarily connected to the interrogation device, which was located outside the clean room. All tests were repeated during the assembly and commissioning of the two SVD halves. After the installation of SVD in the final position, the fibers were laid down in cable trenches and connected to the multiplexers and the interrogator in the e-hut.

4.3 Humidity monitoring

The whole volume of the VXD must be kept at low humidity by a flux of dry Nitrogen, to prevent water vapour condensation and ice formation on the cooling pipes and other cold components. The atmosphere inside the VXD volume needs to be kept at a dew point below -30°C , whereas the measured dew point temperature is about -80°C . The miniature humidity sensors available on the market are not sufficiently radiation hard and their replacement during operation is not feasible. Therefore to monitor the dew point, samples of the circulating gas are extracted from the dry volume and channeled to humidity monitor units located in the e-hut, consisting of a dew point transmitter, a mass flow meter and a pressure transmitter, as shown in figure 73.




Figure 73. Simplified block diagram of the VXD humidity monitoring system, showing the four external units, each equipped with a dew point transmitter, a mass flow meter and a pressure transmitter.

4.3.1 Design

The system collects six input pipes from the forward and backward sides of the cold and warm dry volumes of the VXD.

Four of these lines are each connected to one humidity measurement unit. Two electro-valves allow the connection of each fully equipped line to the VXD dry volumes or to a calibration gas source. The four lines merge into a vacuum pump allowing the gas suction from the detector dry volumes. The suction flow of each line can be precisely set by means of manual flow regulators using the mass flow meter information. An additional vacuum pump grants redundancy: interchange of pumps is controlled by automatic motor protection switches. An alarm LED in the front panel warns if a pump switching occurs.

The two remaining lines, not connected to the suction pump, are connected to a pressure transmitter, to monitor the dry volume pressure.

4.3.2 Construction and installation

The humidity monitoring system was designed, built, and tested at INFN Trieste. It was then shipped to KEK and installed in a rack in the electronics hut (e-hut). The required pipes were installed between VXD dry volume and the e-hut. The readout electronics, integrated in the system, was interfaced to the SVD environmental monitor server via USB connections. The system was fully tested and commissioned before the beginning of the physics run.

The system was originally equipped with miniature dew point transmitters Vaisala DMT143 [117]. This device is able to measure the dew point in the range $-80\text{ }^{\circ}\text{C}$ to $+20\text{ }^{\circ}\text{C}$ with an accuracy in air and nitrogen of $\pm 2\text{ }^{\circ}\text{C}$; it was well matched to the expected a dew point of about $-60\text{ }^{\circ}\text{C}$ to $-70\text{ }^{\circ}\text{C}$, but the level of humidity turned out to be lower, with a dew point below $-80\text{ }^{\circ}\text{C}$. As a consequence, after a few months of operation the sensors frequently went into a self-calibration state, resulting in error conditions, and were subsequently substituted with miniature dew point transmitters Vaisala DMT152 [117], with a measurement range $-100\text{ }^{\circ}\text{C}$ to $0\text{ }^{\circ}\text{C}$ and the same accuracy. These sensors need to be re-calibrated every two years. Short period re-calibration is performed by a built-in software stored in the sensors internal memory.

The DMT152 provides both analogue (4 – 20 mA signal) and digital (RS485) readout. The most important information, the dew point readings, are sent to both front panel displays and serial-over-USB lines. The analog current outputs of the Vaisala sensors are connected to loop-powered seven-segment displays of type Nokeval 302 [118], located in the front panel. Those displays compare the measured values to a programmable range and provide two relay outputs for upper and lower alarms, respectively, whenever the measured value is out of range. The alarm relays are connected to the VXD hardware interlock system. As DMT143/DMT152 communicate digitally via the RS485 bus, a 4-ports FTDI USB-COM485-Plus4 interface was used to connect the sensors to the readout server via USB interface.

In addition, a similar, extended system, with increased redundancy for safe long-term operation, was designed and built at INFN Trieste and is ready for installation. It features one additional spare measurement line, for a total of five. Each line has redundant dew point measurements with two different sensors: an Alpha Moisture Systems Model AMT [119] ($-120\text{ }^{\circ}\text{C}$ to $+20\text{ }^{\circ}\text{C}$ range) and a Michell Systems Easydew [120] ($-110\text{ }^{\circ}\text{C}$ to $+20\text{ }^{\circ}\text{C}$). Mass-flow meters and pressure sensors are

the same type as in the other system. Each line is managed by a dedicated processor board, with both USB and Ethernet readout for an improved interfacing. The system will be installed at the first shutdown opportunity.

4.4 Hardwired interlock system

A fast but orderly shutdown of the PXD and SVD power supplies should occur on a number of critical conditions, triggered by hardwired signals, independent of the slow control network and software.

4.4.1 Design

The [VXD Local Hardwired Interlock \(VLHI\)](#) must be both reliable and flexible, to accommodate different digital and analog input signals, and evolving interlock conditions. An optimal solution, adopted by several experiments, is based on industry-standard programmable logic controllers (PLCs), that are reliable, programmable, and easily expandable with both digital and analog input/output modules. For the VXD, the requirement of [EPICS](#) compatibility led to the choice of a Schneider M340 PLC [121]. The VLHI receives input signals and alarms from::

- 12 groups of NTC temperature sensors;
- 4 dew point sensors and their flowmeters and pressure sensors;
- the water cooling system of the dock chiller (water leaks, water flow);
- the CO₂ cooling system;
- the central Belle II interlock system (solenoid, water leak, [e-hut](#) power, environmental monitors).

The firmware of the PLC implements a “heartbeat counter”, which is incremented continuously. This counter has to be reset by the slow control software within a given maximum period. Exceeding this timeout is counted as a software error and an interlock is triggered.

The following interlock output signals are triggered due to a combination of the input signals:

- SVD low voltage (LV) and high voltage (HV) power supplies;
- PXD power supplies;
- dock chiller and CO₂ cooling system;
- central Belle II interlock system.

4.4.2 Construction and installation

The [VLHI](#) hardware was assembled and the PLC programmed at INFN Trieste. A test box was used to simulate all the input signals with the agreed protocols and their electrical properties. The [PLC](#) program was thoroughly tested and debugged in the laboratory, before the shipment to KEK. All tests were then repeated for each input and output, in the VXD commissioning phase before the start of physics data taking.

4.5 Cooling plant and CO₂ distribution

A schematic of the cooling system of Belle II VXD, serving both the PXD and SVD (whose design and motivations are discussed in section [2.4.4](#)), is illustrated in figure [74](#). The liquid CO₂ plant,

IBBelle, was originally designed for the IBL pixel detector for the ATLAS experiment [19]. A photo of the plant is shown in figure 75.

IBBelle is based on the 2PACL (2-Phase Accumulator Controlled Loop) [37, 38] concept, and consists of a chiller, a heat exchanger, liquid CO₂ pumps, and an accumulator. The R404A refrigerant, which can be operated below $-50\text{ }^{\circ}\text{C}$, is chosen for the chiller. The refrigerant flows to the heat exchanger as well as to the cooling coil in the accumulator. The liquid CO₂ is pressed by the membrane pumps and supplied to the Belle II detector through a CO₂ transfer tube. After absorbing heat at the detector, the CO₂ returns from the Belle II detector and goes through the heat exchanger again to start a new cycle. The accumulator is a stainless-steel reservoir tank of CO₂; it has a volume of 50 L and includes the cooling coils and electric heaters. The cooling and heating power determines the temperature, pressure and amount of liquid CO₂ in the accumulator. Consequently, the accumulator operation determines the temperature of CO₂ at the heat load. It also determines the stability of the whole system. Therefore, the accumulator requires a dedicated control and monitoring system.




Figure 74. A schematic of IBBelle and Belle II VXD cooling system. The approximate length of the pipes is 34 m from IBBelle to the distribution manifold and 9 m from the manifold to the VXD.

The connection between IBBelle and the Belle II detector is done by using a triple wall coaxial transfer tube. The CO₂ from IBBelle goes to the inner tube and the returning CO₂ flows through the intermediate tube. The outer tube is evacuated for thermal insulation. On the detector side, the CO₂ is distributed by using a manifold. There are 12 cooling circuits: eight for the SVD and four for the PXD. Before the VXD heat load a capillary is inserted, reducing the pressure of the liquid, as discussed in the next section. The CO₂ after the heat load flows back to the CO₂ plant through the vacuum isolated transfer tubes up to the manifold, where the individual lines are joined and then further to the accumulator of the IBBelle plant.

In order to avoid water condensation, the regions that face the cooled CO₂ are sealed with insulation foam and filled with nitrogen gas coming from the cold evaporator tank. The dew point is normally kept at about $-80\text{ }^{\circ}\text{C}$.

Operation of the CO₂ system. Before the cooling operation, the system is evacuated and then filled with liquid CO₂ at room temperature. At $30\text{ }^{\circ}\text{C}$ the pressure of CO₂ reaches 7 MPa. With a 1.5 safety factor, the design pressure of the system is 11 MPa and the acceptance tests were performed with even higher pressure (16 MPa). The cold operation starts when the R404A chiller is activated.



Figure 75. The IBelle cooling plant. The four chassis from left to right are for the R404A refrigerator, the CO₂ pumps and tubing, the accumulator, and the control unit. In total, the cooling plant occupies a $6 \times 1.5 \times 2 \text{ m}^3$ volume.

Figure 76 shows a pressure-enthalpy diagram of the CO₂. The region between the liquid and vapor curves is the dual-phase region, where heat can be removed from the heat load at constant temperature and pressure.




Figure 76. The phase diagram of CO₂. The horizontal axis shows the enthalpy, the vertical axis shows the pressure. The red and green lines represent constant-temperature transformations. The purple lines correspond to specific values of Z , the fraction of liquid component in the dual-phase mode.

The blue line with the red points shows the IBelle cycle after the operation is stabilized, which can be briefly described as follows, with reference to the points in figures 74 and 76.

- (a) The liquid CO₂ is first cooled with the refrigerant from the chiller, typically to $-40\text{ }^{\circ}\text{C}$.
- (b) The CO₂ is pressurized by the CO₂ pumps.
- (c) The pressure of the CO₂ is reduced by an orifice located between points (b) and (c). The volume of the CO₂ changes very little and the enthalpy is kept constant in the process.
- (d), (e), (f) The CO₂ flows in the transfer tube with constant pressure. The enthalpy and temperature are increasing because of friction with the tube. The CO₂ lines are branched with the manifold (d). From the manifold to the VXD, triple wall coaxial tubes with smaller diameter are used. While the CO₂ flows through the transfer tube, the returning CO₂ removes the heat of the going CO₂ to minimize the increase of enthalpy.
- (g) A capillary tube is inserted just before the heat load, between points (f) and (g), to reduce the CO₂ pressure and temperature to the target values of $-20\text{ }^{\circ}\text{C}$ and 2 MPa. The pressure is determined by the return line connected to the accumulator.
- (h) The coolant removes heat from the detector and the enthalpy increases, as well as the fraction of vapor, that increases to about 30% as shown in figure 76 by the purple lines. The vapor fraction should be kept below 50% for an efficient heat removal. The returning CO₂ flowing through the outer part of the triple wall coaxial transfer tube gives enthalpy to the CO₂ flowing in the opposite direction in the inner part.

After the CO₂ returns to the CO₂ plant, it goes through the heat exchanger to be cooled down before a new cycle.

The CO₂ circulation system is connected to the accumulator. The refrigerant flow rate and the heater power are controlled to stabilize the operation and total amount of CO₂ in the whole system.

5 Detector control and data acquisition

The detector control and data acquisition software is responsible to ensure proper operation of the Belle II SVD during physics data taking and local calibration runs, and acts as an interface between the SVD hardware, the Belle II run and slow control, the SuperKEKB control system as well as the operators. It provides software modules to configure, control, and monitor the power supplies, the front- and back-end data acquisition hardware, the environmental monitoring sensors, the detector safety and interlock unit, and the radiation monitoring system.

In addition to these run and slow control tasks, it allows dedicated calibration runs to be performed, which are necessary to determine operational parameters like pedestals, noise levels, amplifier gain and calibration constants, which are then used in the back-end electronics for data processing in order to reduce the amount of data sent to Belle II data acquisition. These parameters are also uploaded to the Conditions Database and finally used by the reconstruction software (see section 8) to decode and analyse the data.

Section 5.1 describes the grouping and functions of the SVD control software modules, integrated in the Belle II framework and synchronized by state machines. Efficient and error-free operation of the detector relies on well designed operator interfaces.

The readout chain and data stream are described in section 5.2, following the path from the SVD front-end APV25 chips to the back-end as well as to the Belle II central DAQ. The characteristics of local runs and details of data processing are also reported.

5.1 Detector Control Software

In the Belle II run and slow control systems the subdetectors are integrated into two frameworks, the custom Network Shared Memory 2 (NSM2) [39], and the Experimental Physics and Industrial Control System (EPICS) [122]. NSM2 is used for the global Belle II run and slow control (Belle II RC) as well as for the outer Belle II subdetectors and the central data acquisition system (Belle II DAQ). The VXD and in particular most modules of the SVD run and slow control system (SVD RC/SC) are implemented in EPICS. Only a few DAQ related modules of SVD RC/SC are within the NSM2 regime. The interface between these two platforms is realized by an NSM2-to-EPICS gateway, which translates process information and control requests between them. In EPICS, software modules are implemented as so-called input/output controllers (IOCs), which either interface to hardware components or fulfil dedicated software tasks, such as data analysis or the computation of process parameters. Every IOC stores its data in process variables (PVs), which are shared among other IOCs via an Ethernet network. Each PV in the network has a unique name and can be accessed by other clients via a sequence of UDP and TCP requests. Thus, EPICS allows for a simple and standardized data exchange between the individual software modules and moreover provides an interface to standardized services like an archiver, an alarm system, a message logger and so on.

5.1.1 SVD online software modules

The Belle II SVD online software covers all software modules required to power, configure, control, and monitor the detector. As shown in figure 77, the SVD online software is divided into four groups of I/O controllers according to their functionality:

- SVD Controller & FADC;
- Power Supplies;
- Environmental Monitors & Interlock;
- SVD EPICS Infrastructure.

The **SVD controller & FADC** group contains all core IOCs required to operate the SVD and control data taking. Those are the SVD run controller (SVD:CTRL), the IOC to configure, control, and monitor the FADC system described in section 2.5 (SVD:FADC) and the online low-level data quality monitoring (SVD:QM). SVD:FADC is also responsible to acquire spy data (as described in 2.5.1) during physics runs, which are forwarded to SVD:QM where they are processed and evaluated. Since the amount of these data is rather large, a direct TCP socket connection is established between SVD:FADC and SVD:QM to avoid large data traffic over the SVD EPICS network. The control of the low- and high-voltage power supplies is performed by the SVD:PS IOC, which is the only member of the **power supplies** group in the block diagram. SVD:PS directly communicates with the CAEN power supplies ensuring that the correct voltages are set and all currents are within limits.

The **Environmental Monitors & Interlock** group comprises all IOCs used to configure, monitor, and acquire data from the environmental monitoring sub-systems, like temperature, humidity, and radiation dose monitoring. The detailed functionality of these IOCs is listed below.

- SVD:Env: master ENV IOC that is used to collect the state of the subordinated ENV system and propagates errors and software interlock to SVD:PS and SVD:CTRL;




Figure 77. The block diagram of SVD I/O controllers. Reprinted from [40], Copyright 2020, with permission from Elsevier.

- SVD:NTC: monitors the temperatures of cooling pipes and end rings, measured by NTC resistors;
- VXD:FOS: measures the temperatures of SVD and PXD ladders with fiber optic sensors (FOS);
- VXD:Hum: dew point monitoring of the VXD volume;
- VXD:Rad: radiation dose monitoring with diamond sensors;
- VXD:VLHI: monitoring of the state of the VXD hardwired interlock system (VLHI).

Most of the ENV IOCs have prefix VXD, which means that they are responsible for systems jointly used for SVD and PXD. Those with prefix SVD are responsible for sub-systems dedicated to SVD alone. In the fourth group, the **SVD EPICS Infrastructure** group, all common services like the archiver, the alarm system, the message logger and the gateway between the SVD EPICS network and the Belle II EPICS network are collected. So far, available standard EPICS tools are used for those systems.

5.1.2 State machines

The IOCs of the SVD run and slow control are synchronized by a finite state machine (FSM), with four static states (*Idle*, *Ready*, *Running*, and *Error*) and four transitional states (*Configuring*, *Starting*, *Stopping*, and *Aborting*), which are adopted during the transition between the static states. The state diagram of the FSM is shown in figure 78. From each state there are typically very few possible requests to initiate the transition to the next state. In *Idle* the IOC does not have any information about the hardware of the sub-system, and waits for a *Configure* request to switch to the transitional state *Configuring* as well as to execute all tasks required to initialize and configure the sub-system. If this is completed successfully, the FSM changes to the state *Ready*, in which the system is configured and waits for the *Start* request in order to initiate the transition from *Ready* to *Running*, in which the system is ready for data taking. It stays there until either a *Stop* or *Abort* request is received. The exact

meaning of the states for the individual run control IOCs is listed in table 21. Whenever an error occurs, the FSM goes into *Error* state, where it remains until the confirmation by the operator. In any state only valid requests are accepted and executed, any invalid request, e.g. *Start* in state *Idle*, is ignored.




Figure 78. State Machines of SVD Run Control. Reprinted from [40], Copyright 2020, with permission from Elsevier.

Table 21. Meaning of states for each run control IOC.

State	SVD:PS	SVD:FADC	SVD:QM
Idle	LV off, HV off	off	off
Ready	LV on, HV on (stand-by voltage)	configured	configured
Running	LV on, HV on (peak voltage)	running, data taking	analysis ongoing

For the environmental monitoring IOCs a simpler FSM, as shown in the right diagram of figure 78, consisting of two stable (*Idle*, *Running*) and two transitional states (*Configuring*, *Aborting*), is chosen. The reason is that these IOCs are mainly used for monitoring purposes, staying in the *Running* state most of the time, and are not directly controlled by the Belle II RC. They provide data even when the Belle II detector and thus also the SVD is off. On the other hand, the SVD run control IOCs are controlled by the Belle II RC during global data taking and thus have to follow the run states of the whole Belle II detector.

The interfacing between Belle II RC and the SVD slow control is realized via a set of dedicated PVs. To initiate a state transition, the Belle II RC sets the so-called request PV to the requested transition, e.g. *Configure*. The SVD IOC reads *Request* PV and sets its state PV to the corresponding state, e.g. *Configuring*. Once the transition is completed the request PV is set to *Processed*, which indicates that the IOC has finished the transition and reached a new stable state. If the transition is successful, the state PV is set to the new state, e.g. *Ready*. In the case of a failure, the state PV is either set to *Error*, if the error is unrecoverable, or back to its previous state, if a safe fallback is possible. The Belle II RC regularly polls these two PVs in order to verify that the SVD slow control has finished the transition and in which state it finally arrived.

5.1.3 Operator interfaces

The operator interfaces (OPI) of the SVD run and slow control are implemented in Control System Studio (CSS) [123]. CSS is a collection of tools to monitor and operate large-scale control systems. It directly interacts with the underlying EPICS network to read and set the PV values. Apart from graphical OPIs it provides a data browser, which allows to retrieve from an archiver and generate history plots of each archived PV value. The main OPIs of the SVD RC/SC are the SVD overview




Figure 79. SVD control operator interface.

OPI and the SVD local control OPI (figure 79). The first one displays the general state of all SVD sub-systems, each SVD ladder, the power supplies, the VME crates, as well as general informations from Belle II RC like experiment and run numbers, the actual trigger configuration and so on. This OPI is foreseen as an entry point, allowing the operator to navigate to other OPIs with more detailed information of each sub-system. In addition to these SVD RC/SC user interfaces, there is a dedicated OPI for each environmental monitoring sub-system, which shows the actual temperature (figure 80), humidity, and radiation values as well as the condition of the interlock system.

5.2 Online DAQ software and data processing

The SVD readout chain is shown in figure 81. Its key components are the FADC_Controller board, the FADC boards, and the frontend APV25 readout chips, see sections 2.2.2 and 2.5. The FADC_Controller board is the main interface to the FTSW, which forwards the Belle II clock and trigger to each subdetector. The incoming clock received by the FADC_Controller is scaled down from ~ 127 MHz (one fourth of the RF clock) to ~ 31.8 MHz, the SVD system frequency that drives the FADC boards and APV25 chips. The FADC_Controller moreover converts incoming trigger signals to the APV25 trigger format. Both signals are in turn distributed with 4 Buffer boards over 52 FADC boards to the 1748 APV25 chips.

In idle mode, the APV25 outputs the so-called tick mark every 35 clock cycles. Upon receiving a trigger the next tick mark is replaced with a data frame composed of:

- three header bits: used to identify the start of the data frame;
- eight pipeline address bits: unique address of the internal analog APV25 data buffer, that is incremented every clock cycle;




Figure 80. VXD FOS temperatures OPI.




Figure 81. Data stream within the SVD FADC system.

- an error bit: indicating either an latency or buffer full error;
- 128 strip data samples in a multiplexer (MUX) output order.

The end of the output is indicated by either a tick mark or the three header bits of the next frame. The number of APV25 data frames per Belle II trigger depends on the configured APV25 operation mode as well as the generated trigger output of the FADC_Controller. During a typical run, the APV25 chips are operated in the multi-peak mode generating three frames per trigger. To acquire 6 frame data the FADC_Controller board generates two consecutive APV25 triggers for each incoming Belle II trigger. When 6 or 3 data frames are acquired, the acquisition mode is referred to as 6- or

3-sample mode, respectively. As anticipated at the end of section 2.2.2, only 6-sample event data has been acquired in the initial phase of the experiment. During specific local runs, the APV25 chips are operated in a single peak mode with one data frame output per trigger.

Besides the main data stream chain that provides the data to Belle II there is an additional *spy* channel that allocates a fraction of the data to a dedicated local readout server running the slow control system. Details on how the spy data are acquired and what information is extracted from them can be found in section 5.2.1, while section 5.2.2 describes the processing of the data that are handled to Belle II DAQ.

5.2.1 Local data acquisition and local runs

The FADC system also provides an additional side channel, the so-called SPY channel, that reads out a fraction of events from the main data stream. The SPY data are transferred over four VME connections to a dedicated local readout server running the slow control system, as discussed in section 5.1. The SVD:QM and SVD:FADC perform configuration and calibration of the FADC system. During the local calibration runs the Belle II trigger input as well as the FTB data output are masked. The system is operated in a stand-alone mode. The SVD:FADC issues the trigger generation on the FADC_Controller board over VME. The SPY data format returned by each FADC board can be set to either one of two data types:

- Raw data (RAW): raw APV25 output without any processing;
- Transparent data (TRP): strip signals, applying the FIR correction, the data frame detection, and reordering the strip signal from MUX to the APV25 channel order.

The two data types above are mainly used for debugging and local runs. The SPY data in turn are forwarded to SVD:QM over a TCP connection for processing and evaluation of the calibration. The resulting hardware parameters are propagated back to the SVD:FADC over a configuration file. The same file is in turn pushed to the offline Conditions Database, housing the parameters required for offline analysis.

Several types of local runs, described below, are implemented and routinely performed during detector operation as discussed in section 7.2.

- The **ADC delay scan** evaluates the sampling delay of the ADC chip. Therefore, the height of the APV25 tick mark is scanned while varying the delay in increments of one 64th of the clock cycle (~ 0.5 ns). An example of the ADC delay scan is shown in figure 82. The optimal delay setting of each individual APV25 depends on the layout of the front-end electronics. Since groups of six APV25 chips are connected to one ADC chip, a shared delay setting, dotted line, is selected with following conditions:
 - the common delay is at least 4 ns away from the falling edges of each chip, solid lines;
 - the signals of the tick marks are maximized within the limits of the previous condition.
- The **FIR scan** evaluates the FIR filter coefficients used to compensate the signal dispersion of the data cables. This dispersion is originated in the transfer function of the data cables between front-end electronics and FADC boards and the huge amplitude of digital data at the beginning of the APV25 data frame. The extent of the distortion can be observed in the first




Figure 82. Tick mark vs delay of APV25s connected to the same delay channel, the vertical black dotted line marks the evaluated delay setting.




Figure 83. Baseline distortion, after a tick mark with a signal height of 705 ADC (which corresponds to sample number 0 not drawn here).

seven baseline samples after a tick mark in figure 83. The tick mark resembles an impulse response, where the following seven samples are affected by the dispersion. Thus eight FIR coefficients determined by the following equation

$$\begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} d_0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ d_0 & d_1 & 0 & 0 & 0 & 0 & 0 & 0 \\ d_0 & d_1 & d_2 & 0 & 0 & 0 & 0 & 0 \\ d_0 & d_1 & d_2 & d_3 & 0 & 0 & 0 & 0 \\ d_0 & d_1 & d_2 & d_3 & d_4 & 0 & 0 & 0 \\ d_0 & d_1 & d_2 & d_3 & d_4 & d_5 & 0 & 0 \\ d_0 & d_1 & d_2 & d_3 & d_4 & d_5 & d_6 & 0 \\ d_0 & d_1 & d_2 & d_3 & d_4 & d_5 & d_6 & d_7 \end{pmatrix} \cdot \begin{pmatrix} k_7 \\ k_6 \\ k_5 \\ k_4 \\ k_3 \\ k_2 \\ k_1 \\ k_0 \end{pmatrix} \quad (5.1)$$

were chosen, with d_t representing the tick mark ($t = 0$), and seven following baseline samples after offset subtraction and normalization. The k_t is the FIR coefficient at sample t . The incoming samples s_t^* are corrected with the relation $s = \sum_{t=0}^7 k_t \cdot s_{7-t}^*$, with s representing the corrected sample.

- The **pedestal run** samples the average of signals without the sensor hit, called **pedestal** in the rest of the paper. During this stage the FIR filter coefficients are uploaded to the FADC. The frame detection and strip signal extraction are performed by the slow control software.
- The **noise run** evaluates the strip **noise** as root mean square (RMS) of the each strip signal. After the strip samples have been extracted and the pedestal subtracted, a simple **common-mode noise** correction (**CMC**) is applied. To that end, the strips connected to a single APV25 chip are divided in groups of 16, **32** (default setting), 64 or 128 channels. The common mode within a group is approximated with the truncated average over 50% of strip signals. Channels with known defects, and 50% of the strips with highest and smallest signals are excluded from the evaluation. Finally, the noise of every channel is computed as RMS. As an example the strip noise of a sensor in the innermost layer is shown in figure 84.
- The **calibration run** provides an estimate of the signal generated by a Minimum Ionizing Particle **MIP** in units of ADC counts. The APV25 chip has a so-called internal calibration circuit, which injects configurable amount of charge into the preamplifier of each channel. The absolute calibration of the injected pulse, corresponding to an MIP signal, was performed with testbeam data. Given the spread of about 15% in the relative gain among different APV25 chips due to the significant processing variations in the values of the small on-chip test capacitors [7], this absolute calibration performed with only a few chips from the testbeam data is accurate at the 15% level. The timing of the charge injection is configurable in the step of one eighths of the system clock tick. Figure 85 shows the analog outputs of all the 128 strips connected to an APV25 chip as a function of time, in response to the calibration injection pulse. The maxima correlating to the deposited charge are evaluated with an third-order polynomial fit over the full width half maximum. The two output parameters of the calibration run are the CalPeak and CalTMax (t_{peak}), the height and peaking time of the signal corresponding to an MIP. These parameters from local runs (i.e. the noise, the gain from calPeak, and t_{peak}) are uploaded on the Conditions Database (see section 8.1.4) and used to calibrate the strip charge and time, as explained later in section 8.2.1.




Figure 84. Strip noise of a layer 3 u/P side sensor. The spikes at every 128 strips (APV25 borders) are caused by crosstalk, originating from the power inputs of the chips and are further discussed in section 7.5.




Figure 85. Analog outputs of 128 strips connected to an APV25 chip.

- The V_{SEP} scan is used to detect pinholes and determine a V_{SEP} setting to be set during physics runs, to compensate for pinholes. The maxima of the injected voltage over the internal calibration circuit are sampled at various V_{SEP} settings, see also section 3.1.4.
- The IV scan produces a current vs. voltage plot for each sensor. Since radiation damage is expected to cause an increase in leakage current, the IV curve measurements are used to monitor degradation of the sensor performance with time.

5.2.2 Data processing

A single FADC board processes the output of up to 48 connected APV25 chips. The following sequence of operations on the input data is performed in each FADC board:

1. de-serialization and digitization of the analog APV25 signal with signed 10-bit precision;
2. signal conditioning: removal of reflections and compensation of nonlinear transfer function by a FIR filter;
3. re-ordering of the strip data;
4. pedestal subtraction: the strip pedestal, determined in a dedicated local run and pre-loaded to the FADC, is subtracted from the raw signal;
5. common mode correction: the first pass consists in the subtraction of the average amplitude of strips (excluding masked strips) from each individual strip value. The average is computed in groups of 32 consecutive strips. The second pass consists in the subtraction of the average amplitude of strips (excluding masked strips and strips with signal above three times the noise) from each individual strip value. The average is computed in groups of 32 consecutive strips;
6. online zero suppression: keep only strips with at least one sample above n times the noise, whereas n is an integer coefficient factor, programmed to the FADC boards during configuration. Since no clustering is performed in the firmware, the same cut is applied online for all strips, currently set to three for Belle II physics runs. Different cuts are instead applied in the offline cluster reconstruction for the cluster seed and for neighbouring strips (8.2.2).
7. strips samples are finally presented as unsigned eight bit numbers (0 . . . 255).

The zero suppressed data are forwarded to the corresponding FTBs (see figure 81) that connect to both the PXD system and the Belle II DAQ. In particular, in the Belle II DAQ chain, the COPPER boards that receive the data from FTB, encapsulate them in objects that can be read by the Belle II software [41], so that they can be processed by the *offline* software. The data contain several additional information besides the strip data (FADC board number, APV25 chip number, APV25 channel, 3 or 6 digitized samples) including:

- trigger number;
- trigger type (e.g.: subdetector that triggered the event, random trigger, delayed physics trigger, . . .);
- acquisition mode (3- or 6-acquired sample per trigger);
- several possible errors detected in the processing up to FTB.

Beside the zero-suppression acquisition mode used in normal runs and described above, there are two more acquisition modes that are used for specific studies or calibrations. In the raw mode data are sent to the output after the FIR correction; in the transparent mode data are sent to the output after the re-ordering of the strip data.

6 Detector commissioning and installation

A well-planned and thoroughly executed commissioning of the detector is a key factor in its successful and efficient operation at the beginning of the physics data taking.

The initial commissioning of SuperKEKB with collisions was carried out from April to July 2018, before the installation of the VXD. During this period, called Phase 2, the inner volume of the Belle II detector was instrumented with detectors dedicated to the characterization of beam-related backgrounds. As described in section 6.1, final PXD and SVD ladders, one per layer, were positioned at the design radii in the horizontal plane of the intersecting beams, where backgrounds were expected to be larger. These measurements were essential to provide a realistic extrapolation of the level of backgrounds from Phase 2 to Phase 3 and ensure the safety of the full VXD, which was a pre-condition for the detector installation. In addition, it confirmed the performance of the PXD and SVD sensors and their readout electronics in the SuperKEKB environment before the VXD installation.

In parallel with Phase 2, the complete SVD, fully equipped with electronics and environmental monitors, was operated and tested with cosmic muons for about two months. This commissioning phase, summarized in section 6.2, allowed a thorough verification of all hardware and software components, including cooling, calibration procedures, track reconstruction, and efficiency mapping.

The final installation and commissioning steps are described in section 6.3. In October 2018, the two halves of the SVD were mounted around the beam-pipe — PXD assembly in a clean room. The procedure required a careful alignment of the SVD. The VXD could take calibration and cosmic data in these conditions for about one month and was fully characterized before its insertion in the CDC of the Belle II detector. Special procedures and mechanical devices were developed for the insertion process, made difficult by tight tolerances and limited space for cables. The insertion of the VXD was carried out on 21 November 2018, and followed by the cabling and piping until the middle of December 2018.

6.1 Phase 2 SVD commissioning

Before the final VXD could be installed into the Belle II detector system, a commissioning operation of SuperKEKB was carried out from April to July 2018. This campaign of the SuperKEKB operation without the VXD is called “Phase 2”, while the operation after the VXD installation is called “Phase 3”. The Phase 2 operation was a major milestone for SuperKEKB to confirm the feasibility of the nano-beam scheme [3] as well as to understand the beam background on various Belle II subdetectors. During the Phase 2 operation, SuperKEKB succeeded to squeeze the beams down to $\beta_y^* = 3 \text{ mm}$ with beam currents of more than 800 mA for the both LER and HER rings. The peak luminosity achieved in Phase 2 was $5.55 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

Although the final SVD was not installed in Phase 2, this operation provided a good opportunity to test the sensors and the developed FADC system under the operating conditions of the experiment as a Phase 2 SVD commissioning. For this purpose, one SVD ladder per layer was installed in Phase 2.

6.1.1 Commissioning goals

The main goal of the Phase 2 commissioning was to evaluate the beam background condition in the VXD volume before its installation.

Being very close to the interaction point, high background levels could cause radiation damage to the SVD sensors, as well as an increased hit occupancy, which could degrade tracking performance and increase data rates in the DAQ. The most restrictive limit was related to the degradation in the tracking performance. Background levels in the SVD are measured by the hit occupancy, and the equivalent limit in the Layer 3 hit occupancy is about 3%. Another task of the Phase 2 commissioning was to check the performance and stability of the SVD. The Phase 2 commissioning was the first test in the real Belle II environment where any noise interaction to other subdetector systems could occur. Another important purpose of Phase 2 was to test the VXD installation, using the developed procedures and tools.

6.1.2 Phase 2 SVD setup

For the Phase 2 commissioning, a special VXD which had one PXD or SVD ladder per layer, with a detector configuration shown in figure 86 was designed to study the detector performance in each layer. The FANGS, CLAWS, and PLUME sensors constituted the BEAST II system for background measurement [42]. For the ease of the installation of those four SVD ladders, an “SVD cartridge” accommodating the four ladders, mechanical supports, and cooling services was developed, as shown in figure 87. These ladders had the same longitudinal direction parallel to the beam direction, and the sensor planes were facing the interaction point as shown in figure 86. The cartridge was mounted on the beam pipe flanges at the outside direction of the SuperKEKB main ring (the +X direction in the Belle II coordinate system). The radial distances of the sensors from the IP were the same as the Phase 3 design, i.e. $R = 39$ mm, 80 mm, 105 mm, and 135 mm. The outside direction was chosen because beam-background simulations [43] indicate that the outside has the largest background contribution from the storage beams. There were three CO₂ circuits in the cartridge which absorbed heat from the APV25 chips on the DSSD modules, in which flowed the dual-phase CO₂ provided by IBelle. One long circuit was attached on the Origami modules of Layer 4, 5, and 6. The remaining two were attached on the forward and backward aluminum support blocks that cooled the chips on the Forward and Backward modules via the aluminum mount blocks of the ladders. Two CFRP plates of 1 mm thickness were located in parallel on both sides of the ladders to support the cartridge structure. As seen in figure 87, all the four ladders, support aluminum blocks, and cooling circuits were accommodated in the SVD cartridge, whose installation was easily done by just attaching it on the VXD support structure. The region between the forward and backward aluminum blocks was the physics acceptance where only light materials could be located.

6.1.3 Results and conclusions

During the Phase 2 operation, several beam background studies were performed to understand the background composition and compare it with the MC expectation. The studies provided correction factors to extrapolate MC simulation results to the final SuperKEKB machine condition. In Phase 2, the measured occupancy on the innermost layer, Layer 3, was less than 0.4%, which is well below the 3% limit. The expected occupancy in early Phase 3 was similar to the one measured in Phase 2 since beam currents and machine optics were expected to be similar. From this observation and studies, the background conditions for early Phase 3 were considered safe for the installation of the full vertex detector.




Figure 86. A cross sectional view of the Phase 2 VXD setup, which contains PXD ladders, SVD ladders, FANGS sensors, CLAWS sensors, and PLUME sensors. The PLUME sensors are not shown here. The radial distances of the SVD sensors from the IP are written below the sensors.




Figure 87. The SVD cartridge setup for the Phase 2 commissioning.

Analyzing the data taken during the Phase 2 SuperKEKB operation, the performance of the SVD ladders was evaluated. The noise and gain of all sensors were found to be stable and no visible change was observed. Figure 88 shows distributions of the cluster [Signal-to-Noise Ratio \(SNR\)](#) (see the definition in section 8.2.2) for the large rectangular sensors. The most-probable-values of the distributions were over 20. These high SNR values resulted from the low noise characteristics due to the chip-on-sensor concept described in section 2.2. Efficiencies on the Layer 3 forward sensor of $99.08 \pm 0.02 \%$ and $99.44 \pm 0.02 \%$ were measured for the u/P side and v/N side, respectively. From these measurements, the excellent performance of the SVD ladders was confirmed.

6.2 Stand-alone full SVD commissioning

Two SVD halves for Phase 3 were fully assembled in the first basement floor of the KEK Tsukuba experimental hall and then moved to a storage box in the fourth basement floor (B4) in February 2018




Figure 88. Cluster Signal-to-Noise Ratio (SNR) distribution on the large rectangular sensors in the Phase 2 Layer 4, 5, 6 ladders in Exp.7 Run.783. Black and red histograms correspond to the clusters on the u/P side and v/N side, respectively.

for the first half and in July 2018 for the second half. In the storage box, before the coupling with PXD in October 2018, a stand-alone commissioning test for the SVD halves was performed using cosmic muons. This commissioning was performed in parallel with the Phase 2 commissioning, which is described in section 6.1.

6.2.1 Targets of stand-alone SVD commissioning

The first target in this commissioning test was to check the healthiness of all the DSSD sensors, front-end electronics, and environmental monitoring sensors (NTC and FOS, see section 4.2) on the SVD halves, because this was the last performance test before coupling PXD and SVD. Another target was to confirm the tracking capability of the SVD analysis software measuring cosmic muons, since it was the first opportunity to detect the particle tracks with the real geometrical configuration of the multiple layers.

Figure 89 shows the setup of the stand-alone commissioning test. For the commissioning, the two SVD halves held by the pick-up tool (see section 3.7.5) were moved into the storage box built with aluminum frames and panels. The box offered a protection for the detector, but was also a dark environment that allowed applying high voltage to the sensors. It was also flushed with nitrogen to offer a dry environment with a dew point lower than -40°C , necessary to operate the cooling system which uses CO_2 at -20°C . Scintillation counters were set on the roof of the box and under the storage tables to provide triggers for the cosmic rays. During the two months commissioning, detailed calibrations were performed for the first time with the full SVD ON and with the final configuration of the CAEN power supply system and the FADC readout system.

For the cooling of the SVD during commissioning, the MARCO (Multi-purpose Apparatus for Research on CO_2 cooling) cooling system was utilized. MARCO is a CO_2 cooling unit based on the 2PACL (2-Phase Accumulator Controlled Loop) [37, 38] concept. The cooling capacity of MARCO is in the order of 1 kW, and the operational temperature ranges from the room temperature to -40°C .




Figure 89. Stand-alone SVD commissioning setup. The two SVD halves were put in the storage box (in the photo, the front panel is open), and fixed on the floor with anchors as a protection against earthquakes. After closing the front panel of the box, they were tested in the box.

MARCO was designed for non-expert users with all its necessary operations being automated with a PLC. The CO₂ cooling design of MARCO was utilized also for IBBelle and the ATLAS IBL cooling. Figure 90 shows MARCO installed in the Tsukuba hall B4 floor for the commissioning.

In the commissioning test, the final system of FADC and FTB boards, that were described in section 2.5, were used for the readout of the detector signals. Also, the final environmental system to monitor the temperature of the detector and the CO₂ cooling tubes, that was described in section 4.2, was used. The functionality of the equipment and the firmware and software was confirmed.

6.2.2 Results and conclusions

The SVD halves ran stably during the commissioning and about 3×10^7 cosmic-muon events were recorded, confirming that all the sensors were working properly. All the noise and calibration constants were stable over the entire two-month period. The noises in the sensors were consistent with our expectations which took into account the various noise contributions significantly depending on the sensor side and location in the ladder. A summary of the noise contributions for the various sensors is shown in section 9. The number of defective strips, which had anomalous values of the noise, gain, or pedestals, was about 1% per sensor. This number was driven by the pinholes which



Figure 90. MARCO cooling system installed in the Tsukuba hall B4 floor.

were already present in the produced silicon sensors, and only a few additional bad channels, like short- or open-circuits, introduced during the ladder assembly.

It was also possible to observe and fully characterize a number of peculiar effects, for example the sensor turn ON effect, or cross talk effects, which were visible for the first time during commissioning in the cold and dry environment. Although they do not affect the performance, they deserve attention to operate SVD in a proper way. Details of these effects are described in section 7.

The flow of the CO₂ coolant in the cooling circuits was also confirmed by checking temperatures of the cooling channels measured with the NTC. Only one NTC sensor was found dysfunctional, however not affecting the ability of the environment monitor and control thanks to the redundancy of the sensors. The functionality of other NTC and all FOS for the temperature measurement was confirmed.

Tracks of the cosmic muons are reconstructed from the hits on the sensors. Using the clusters associated to the reconstructed tracks, the SNR and hit efficiency of SVD sensors are measured. Figure 91 shows the SNR distributions measured in the stand-alone commissioning. Depending on the sensor location, the incident angle distribution of the cosmic muons is found to be different, and this affects the signal charge and the cluster size of the hit on the sensor. The SNR distributions change according to the sensor location and side. By considering the effect of the incident angle, it was confirmed that the measured SNR and efficiency are consistent with that observed during the beam tests and Phase 2 commissioning.




Figure 91. SNR distribution of a Layer 5 ORIGAMI $-Z$ sensor measured in the stand-alone SVD commissioning with cosmic rays.

6.3 Installation

As the final step of the SVD construction, it needed to be coupled to the PXD to build the full VXD. Another commissioning phase was carried out for the full VXD, as a final performance check before its installation. Once the performance was confirmed, the VXD was finally installed in the Belle II detector. In the following subsections, those steps until the VXD installation are explained.

6.3.1 Coupling to PXD

The coupling of the SVD halves around the PXD to obtain the full VXD was performed at the beginning of October 2018. Before the coupling, which took place on a granite table in a clean room built on the B4 floor, the interaction point section of the beam pipe (IP beam pipe) was assembled with heavy metal shields, and the PXD, fabricated in two halves in several German sites and transported to KEK, was mounted around the IP beam pipe. In the coupling the two SVD halves were attached around the PXD from both sides. The same pick-up tool used for the SVD pick-up procedure (see section 3.7.5) as well as for the storage of the two SVD half shells, was used in the coupling operations.

During the attachment, an accurate position control of the moving SVD half was important to prevent the PXD and SVD from touching each other, since the design clearance between them was only 1.5 mm. The pick-up tool allowed us to control the vertical position and rotation angle of the SVD half, and they were precisely aligned by using alignment gauges. Also, guide pins on the beam pipe flanges prevented any accidental displacement of the SVD, which could result in a mechanical interference with the PXD. With the alignment mechanism and the guide pins, coupling to the PXD was performed successfully. Figures 92 and 93 show pictures of the first SVD half coupled to the PXD and the fully assembled VXD. The coupling of the SVD with the PXD took one day for each half, while two days were required afterwards to perform precise alignment of the two halves and connect all additional mechanical components and services.



Figure 92. A picture of the first SVD half already coupled to the PXD.

6.3.2 Combined PXD-SVD commissioning

After the coupling of the SVD to the PXD was completed, commissioning of the fully assembled VXD was performed for about a month starting from the middle of October 2018. In the combined PXD-SVD commissioning, key performances of the SVD were checked as the final step before the VXD installation. At first, the healthiness of all SVD components was checked in order to ensure that no damage was caused to the detector during the coupling with the PXD. It was confirmed that the full detector can be read out without any issue, and observed noise, gain, and pinholes were consistent with the results before the coupling. Another important task was to confirm the moderate temperature of the detector using FOS and NTC when sensors were biased, APV25 chips were powered on and configured, and CO₂ cooling was active. This test was important because it was the first temperature evaluation inside the closed SVD volume with the full heat dissipation from APV25 chips (in total, about 700 W) and the heat absorption of the CO₂ cooling using the real setup. The origami modules showed temperatures from about 4 °C to 16 °C, confirming that the produced heat was removed effectively and there was no overheating of the detector volume. The hardware, software, and firmware were also successfully tested with the full-scale FADC system.

As for the stand-alone SVD commissioning, scintillation counters on top and bottom of the VXD were installed to provide triggers for cosmic data taking. Figure 94 shows one event display of the cosmic event reconstructed with the SVD reconstruction software, that is described in section 8. Using the cosmic muon hits, the reasonable distributions of the energy loss and the performance of the alignment software were obtained.




Figure 93. A picture of the fully assembled VXD in the B4 clean room.




Figure 94. An event display of a cosmic muon track in the combined PXD-SVD commissioning. It shows inner two layers of PXD sensors (the second layer is incomplete) and outer four layers of SVD sensors. Hit strips in the SVD are indicated by blue lines, and a reconstructed track from these hits is shown as a brown line.

6.3.3 Installation in Belle II



Figure 95. Photo taken during the VXD installation in the Belle II detector: lift-up of the VXD with a crane. The cables and pipes were fixed in the cables trays on both sides.

After the VXD commissioning, the bellows pipes were attached to both ends of the IP beam pipes to connect them to the final focus beam pipes. On the bellows pipes, the diamond sensors were installed to monitor the radiation dose around the pipes (see section 4.1).

During the VXD installation preparation, the cables and pipes of the PXD and SVD were laid straight and fixed in cable trays connected to the VXD. During the transportation, the VXD and cable trays were supported by the crane tool, and the combined structure was lifted by the crane as shown in figure 95. The VXD and cable trays were lowered and put on the slide guide, that was then installed inside the Belle II detector and extended up to the forward side of the QCS. The slide guide was designed to safely insert the VXD into the dedicated location at the center of the CDC cylinder. The outer radius of VXD is 155 mm while the inner radius of the CDC is 160 mm, therefore there is a 5 mm clearance in the design. After the crane tool was detached, the VXD and cable trays were slid on the guide by hand to insert them into the CDC, as shown in figure 96.

The VXD was fixed on the CDC with alignment pins. The positions of the alignment pins on the VXD were adjusted in advance based on the results of a position survey, in which the CDC position with respect to the accelerator devices was measured with a precision of about 100 μm . With this alignment, the VXD and IP beam pipe were located at the correct position. After the



Figure 96. Photo taken during the VXD installation in the Belle II detector: installation of the VXD in the Belle II detector. The slide guide installed in the Belle II detector using the slide guide installed in the forward part of Belle II.

fixation of the VXD, the slide guide was disassembled. The part inside the Belle II detector was removed by pulling it out from the gap between the VXD and CDC.

The transportation and insertion of the VXD were completed on 21 November, 2018. After the insertion of the VXD, the cables and pipes from the PXD and SVD were laid and connected to the junction boxes, which were then connected to the readout electronics, the power supplies, and the CO₂ cooling plant IBBelle. The cabling and piping works were carried on until the middle of December, due to the large amount of cables and the limited service space for cables and pipes. As a last commissioning step, the installed VXD was tested again using cosmic rays. The results of this last test were found to be consistent with the ones taken before the installation, and confirmed an excellent performance of the SVD.

7 Detector operation

Careful operation is essential to efficiently record high-quality data for subsequent analysis, as well as to identify performance trends in the evolving high luminosity conditions of SuperKEKB. Section 7.1 gives a summary of the data taking periods in 2019 and 2020, followed by a description of the organizational efforts to progressively simplify the tasks of operators, and to relieve part of the

burden on KEK residents, by allowing remote operation from the collaborating institutes scattered around the globe.

A relevant part of the detector operation is the periodic recalibration during local runs (section 7.2), to provide updates of the calibration constants (pedestals, noise, gain), used for zero-suppression and offline analysis, and of the maps of defective channels. The experience accumulated in monitoring the evolution of calibration constants and defective channels is summarized in section 7.3. A few issues on specific sensors and on the power supply system, encountered during two years of operations, are also briefly mentioned.

A great deal of attention is devoted (section 7.4) to monitor the accumulated radiation dose, by observing the evolution of noise and leakage current in the **DSSD** sensors. A specific tool is developed to monitor the depletion voltage of the sensors, while local calibration runs are used to search for additional defective channels after high beam-loss events: the limited number of new defects confirms the robustness of the **SVD** sensors.

Further observations on the behaviour of the detector components are detailed in section 7.5, in particular cross-talk effects and baseline shifts depending on the trigger rate. In section 7.6 the implementation of a firmware feature to minimise the dead time in case of high occupancy is described.

Section 7.7 is dedicated to the operation of the beam loss monitor and beam abort system, based on diamond detectors; a brief outline of the experience on temperature and humidity monitors is also given, followed by some remarks on the hardware interlock operation. Finally, section 7.8 presents some results on machine induced-backgrounds, both in single-beam configuration and with colliding beams. These preliminary studies are essential for the extrapolation of the background levels to future higher-luminosity conditions of SuperKEKB.

7.1 Data taking periods and shift organization

The full **VXD** system was placed inside the Belle II detector in November 2018. After a series of tests made to verify that the system was working properly, Phase 3 operations started on January 21, 2019. Each year was divided into three running periods. The runs considered in this section are the following:

- **2019a**: January 21, 2019 – March 31, 2019
- **2019b**: April 1, 2019 – July 1, 2019
- **2019c**: October 15, 2019 – December 12, 2019
- **2020a**: February 25, 2020 – March 31, 2020
- **2020b**: April 1, 2020 – July 1, 2020

The first six weeks of operation of the full Belle II detector in the 2019a run were devoted to cosmic ray runs, used to test the full **DAQ** chain and the detector performance. Machine operations started on March 11, 2019 with single beams circulating in the machine, achieving the milestone of first collisions in Phase 3 on March 25. Unfortunately, a fire in a building near the Linac complex stopped machine operations for four weeks. After the restart, injection studies allowed operation of both beams with continuous injection from May 14, 2019. The machine ran stably with these conditions until the end of 2019b run. No other major issues occurred during the 2019c run. Some 2019 beam time was devoted to machine tuning and studies, during which the Belle II detector was turned off,

but the whole DAQ chain was kept running in order to test its stability. Moreover, background studies were performed periodically, and in particular every time the machine optics was changed, to understand and mitigate the background induced by the accelerator on the detector. Results for some of these studies are discussed in section 7.8.

To manage SVD operations in Phase 3, the role of an “operation coordinator” was established. A shift system, that included local and remote shifters, was developed in order to safely operate the SVD detector. During 2019 runs, the local shifters ensured constant presence at KEK, being “on-call” in case there were problems with the SVD that required actions from an expert. The remote shifters were in charge of monitoring the detector and the environmental variables, as well as of checking data quality using [Data Quality Monitoring \(DQM\)](#) plots. In 2020, with the spread of the COVID-19 pandemic and travel restrictions, the shift system was modified adding an expert remote shifter to operate the SVD detector remotely, with the local shifter role being limited to a few non-ordinary operations that required physical presence in the experimental hall.

7.2 Detector calibration

To efficiently operate the detector, it is important to monitor and constantly update the calibration constants, like pedestal, noise and gain, since they are used both for the online zero suppression during data taking, and in offline reconstruction for the signal energy conversion and hit time reconstruction as explained in detail in section 8.2.1. For the online zero suppression the strip signal, pedestal subtracted, is compared with a threshold selecting only strips that have a signal higher than three times their noise (that is the strip Signal-to-Noise Ratio is greater than 3). During reconstruction the strip signal is then converted to charge using the strip gain measured with the internal APV25 calibration circuit.

Periodic calibration runs are also useful to monitor defective channels, i.e. those with anomalous calibration constants, that need to be masked out to avoid adverse impacts on the detector performance. Periodic measurements of the characteristic IV curve of all the sensors were also performed to evaluate effects of radiation damage. All types of calibration runs, here referred to as “local runs”, are described in section 5.2.1.

Different sets of local runs were then defined to monitor all important parameters and they were taken with different frequency during data taking. This frequency was optimized based on some initial experience, which showed that in general the calibration constants were very stable; these constants were found to slowly evolve with radiation damage or due to temperature changes related to different detector settings that could change the power consumption.

Local runs were arranged in three different groups, depending on the time available as well as the calibration constants to be monitored. Local run groups are defined as follows:

- **Short:** pedestal, noise (~ 20 min). The short local run is used when there is a short no-beam time, to monitor noise and pedestals of every channel in the SVD. This is particularly important since the signal of each channel depends on pedestal and noise values, so even a small change in one of these parameters can affect the signal of each channel.
- **Intermediate:** [ADC](#) delay scan, [FIR](#), noise, pedestal, calibration (~ 55 min). This kind of local run can be done when there is a longer period without beams, or during maintenance days (every two weeks). In the intermediate set of local runs, the ADC delay and the FIR are

also evaluated to check their stability over time, while the gain of each channel is evaluated with the calibration.

- **Complete:** ADC delay scan, FIR, noise, pedestal, calibration, V_{SEP} scan, IV scan (~ 150 min). In addition to all runs of the intermediate set of local runs, V_{SEP} and IV scans are performed. The V_{SEP} scan is used to find defective strips, pinholes in particular, while the IV scan is used to evaluate the characteristic IV curve of each sensor. Typically, the complete set of local runs is taken every two weeks during maintenance days, or during the downtime after a severe beam abort due to large radiation spikes, that can potentially create damage in the SVD sensors by adding new pinholes.

The information on calibration constants is included in the configuration file used for data taking, which is generated using calibration constants measured during a specific local run. When a new local run is performed, calibration constants are compared with the ones used previously for data taking. If the differences in the calibration constants are significant, a new configuration file is generated using the calibration constants of the last local run.

7.3 Stability

During the first two years of operation, a number of small issues and effects on sensors were also observed, and are reported in this section, together with the evolution of calibration constants, which were monitored taking local runs frequently.

7.3.1 Evolution of calibration constants and defective strips

Before the accelerator started operating, calibration constants were very stable, while after the start of beam operations an evolution related to radiation damage effects was observed. The largest effect was visible in Layer 3 sensors, which are the most exposed to machine induced background.

The evolution of calibration constants from March 2019 to July 2020 is shown in figure 97 for Layer 3 sensors in the forward region. Large changes were observed on noise levels due to radiation damage: in Layer 3 mid plane, the most exposed to radiation, a noise increase up to 15% and 25% was observed on v/N and u/P sides, respectively, with some saturation effect already visible. In the external layers exposed to lower dose the increase in noise levels was around 5–10% only until July 2020. More details on the evolution of the noise are shown in section 7.4.1. Pedestals initially changed by about 2-4% and then stabilized; a small reduction of the gain around 5% was also measured.

The number of defective SVD channels, i.e. strips masked on the basis of their anomalous values of noise, pedestal, and gain from local runs, was monitored and was very stable during all 2019 runs and the first two runs of 2020. The total number of masked strips was about 1% on each sensor's side since the beginning of operations in March 2019. The number of masked strips is mainly due to the number of defective strips in the produced ladders, related to sensor defects, which had already been identified during ladder production. A few additional defects, mainly short- and open-circuits, were introduced during ladder assembly. There are also a few "hot strips" that developed during data taking; these strips become very noisy after some time, producing a significant increase in the sensor [occupancy](#) as well as distorting the distribution of the relevant quantities for the sensor performance monitoring. Hot strips need to be masked out so that they do not affect data quality monitoring. In




Figure 97. Temporal evolution of calibration constant for Layer 3 FW sensors. Top plots refer to u/P side, bottom plots refer to v/N side. From left to right, the plots show the evolution of noise, pedestals and 1/gain.

total, only 16 new hot strips developed in the entire SVD in 2019, and 3 in 2020, corresponding to 0.008% of the total SVD strips. The overall number of masked strips (including defective and hot strips) remained basically constant throughout the run: the total fraction of masked strips at the end of 2019c run was 0.96%, with 1.00% of strips masked on the u/P side and 0.86% of strips masked on the v/N side. Pinholes are not masked, because they are normally compensated by the $V_{SEP} = -0.8$ V (see 3.1.4 for the V_{SEP} definition).

7.3.2 Issues on specific sensors

During the 2019 run, some peculiar behaviour of a few sensors was observed, and is briefly described in this section. It should be noted that most of these effects did not cause any performance deterioration, except for one single chip in Layer 3 that was excluded from data taking (causing an efficiency loss in the corresponding sensor) for a few months in spring 2019 and which was later recovered during the summer shutdown.

Connection issue. From March 10, 2019, one APV25 chip of sensor L3.2.2.n was responsible for some SVD errors during data taking. At first, it was thought that ADC delay scan or FIR parameters were changed or wrongly evaluated during the last local runs, causing the errors. Other sets of local runs were taken to verify this hypothesis, but the errors were still issued even during local-run data taking. The only way to proceed with data taking without issuing errors was to disable the problematic APV25 chip, losing a fraction equal to 1/84 of the Layer 3 v/N side acceptance.

When the access to FADC crates was granted during a maintenance day, attempts to reconnect data cables on the FADC board front panel and to use different FADC channels were made without success. Once access to dock boxes was granted at the end of 2019b run, data cables were disconnected and reconnected on the junction boards. Tests performed after the reconnection showed

a recovery of the DC/DC converter voltages, a good signal output coming from the APV25 chip previously disabled, and good noise levels. The issue was considered solved and the full SVD was again available for data taking. However, at the beginning of 2020a run, DC/DC converter voltages corresponding to the same APV25 chip decreased to a level lower than normal, indicating a remaining potential issue with the connection. The lower DC/DC voltages did not affect the performance of the APV25 chip, which was subsequently included in data taking.

High sensor current. When the machine started using continuous injection, which entails keeping the beam currents almost constant by performing top-up injections for both beams, an anomalous rise of the leakage current of sensor L6.14.2 v/N side was observed. The rise was seen only when beams were continuously operated and the bias voltage was applied to the sensors, while no increase was visible with the bias voltage applied but no beams in the machine. The increase in the leakage current was around $4 \mu\text{A}$ per day, as shown in figure 98. After one week, SVD was unbiased for some days, while the machine was operating with high currents, and this caused a sudden decrease of the L6.14.2.n leakage current. As soon as the detector was biased again, the increase restarted at the same rate as before, reaching $40 \mu\text{A}$ after two weeks of operation.




Figure 98. L6.14.2 v/N side leakage current evolution from the start of continuous injection studies to the end of the run.

The hypothesis to explain this behaviour was that some positive charge had accumulated on the u/P side of sensor L6.14.2 due to both ionization created by beam background radiation and the presence of an electric field in air between the v/N side of Layer 5 ladders and the u/P side of Layer 6 ladders. The positive charge created on top of the u/P side can increase the electric field at the junction side, close to a sensor defect, and could increase the leakage current due to generation from impact ionization at that position. The hypothesis was consistent with the observation that when beams were stopped the increase in the leakage current also stopped, confirming the relation with beams. A similar effect was also observed in the past in the BaBar silicon sensors [11]. In addition, a few additional small groups of noisy strips on the two sides of the sensor were observed during local runs taken with high leakage current. On the u/P side the new noisy strips developed close to a preexisting pinhole, probably due to a local defect on the sensor that prevented the healthy growth of the oxide. This was an additional indication that the high current could be localized in that part of the sensor, at the intersection of the two groups of noisy strips on either sides of the sensor, close

to a preexisting defect, with the effect being constantly aggravated by an accumulation of positive charges on the u/P side surface of the sensor.

To verify the hypothesis of higher leakage current due to impact ionization, the sensor leakage current was measured at different temperatures (changing the set point of the cooling plant by 5 and 10 °C). All sensors in the same Layer 6 ladder doubled the leakage current with a temperature change of 7 °C, as was expected for a leakage current due to thermal generation in the bulk or surface. On the problematic sensor, the current increased by only 10% with the higher temperature setting, supporting the hypothesis of generation due to impact ionization.




Figure 99. High voltage configuration between the four layers of the SVD. On the left side the original configuration is shown, with symmetrical bias voltages on P and N sides for all layers. On the right side, the modified configuration is shown, with Layer 6 sensors' voltages changed so that the electric field between the v/N side of Layer 5 and the u/P side of Layer 6 is reduced to 70 V. Reproduced from [44]. CC BY 4.0.

In order to mitigate the problem, which could lead to the sensor reaching its breakdown point, the electric field in air between the v/N side of Layer 5 and the u/P side of Layer 6 was reduced by applying different absolute values of high voltage in Layer 6: keeping $\Delta V = 10$ V between the u/P and v/N sides of L6 ladders, but reducing the electric field between the v/N side of Layer 5 and the u/P side of Layer 6. A schematic view of the solution is shown in figure 99.

The new **High Voltage (HV)** configuration was applied in June 2019 with an HV shift of 30 V, which reduced the voltage between the v/N side of Layer 5 and the u/P side of Layer 6 to 70 V. The decrease in the electric field between Layer 5 and Layer 6 ladders was sufficient to arrest the increase in the sensor leakage current, that remained stable around 38 μ A until the end of the 2019b run. During the summer 2019, the current recovered to very low values, as expected due to the discharge, and started to rise with a smaller slope during the 2019c beam operation. The current reached a saturation value around 42 μ A and remained stable until the end of the 2019c run. The HV shift adopted as a countermeasure is still in place, effectively mitigating this effect. No other sensors developed the same symptoms.

Leakage current instability. In May 2019, sensor L6.8.3 developed an instability in the leakage current, with spontaneous jumps that were not correlated to particular interventions on either the detector or the machine. The jumps in the leakage current were symmetrical between P and N sides, with the highest recorded value around 16–17 μ A. Some tests were made to understand the behaviour, but none of them gave a clear explanation of the phenomenon, although the IV curve showed some soft breakdown toward 100 V. After some time, the leakage current went back to a low

and stable value, within 5 μA , still being above the average with occasional small fluctuations. No particular countermeasures were taken against this issue, since its origin was not clearly understood, and the effect did not cause sensor performance degradation.

7.3.3 Turn-on effect of Micron sensors

Already during commissioning, an effect was observed on some of the slanted sensors in the forward region: immediately after turning ON the high voltage, the v/N side sensor noise and therefore occupancy increased with time, while the leakage current decreased with time, with a time constant up to 12 hours. Due to this effect, a local run taken right after turning ON the high voltage would underestimate the noise level. The effect could be explained assuming that after the sensor bias is turned ON, positive charges start to accumulate on the passivation layer of the N side. This effect can increase the extension of the electron accumulation layer, always present on the N side at the Si-SiO₂ interface, close to the v/N side strip. The accumulation of negative charges extends the width of the v/N side implanted strip, therefore increasing the inter-strip capacitance, which results in the observed increase in noise and occupancy on the v/N side. This mechanism also explains the reduction in the leakage current: electron accumulation reduces the extension of the surface depletion region that contributes to the surface leakage current, which results in a smaller current. This effect was already observed [45]: the steady state conditions in the e⁻ accumulation layer at the Si-SiO₂ interface is slowly reached, due to the movement of charge on top of the passivation layer, with a time constant that depends on humidity and bias history. This effect in our double sided sensors can influence both inter-strip capacitance and leakage current, resulting in the observed anticorrelation. The behaviour is shown in figure 100, where leakage current and occupancy are




Figure 100. L5.3.1 v/N side leakage current (blue) and occupancy (green) after turning ON the sensor bias.

shown for the v/N side of L5.3.1 sensor: right after the high voltage is turned ON, the leakage current and occupancy start changing until, after some hours, the sensors settle to a condition where leakage current and occupancy have stable values similar to other sensors. There is no specific cure for the observed issue, but it is important to avoid taking local runs right after turning ON HV, especially if HV was OFF for a long time, because in this way the noise would be underestimated, with the consequence of a high occupancy observed once these sensors reach the operating condition.

7.3.4 Power supply related issues

Two issues regarding power supplies arose in 2019. The first issue was observed during the 2019a and 2019b runs, when a few channels of some **Low Voltage Power Supply (LVPS)** boards had a failure of a **MOSFET**, and were temporarily removed from the experiment to replace the defective transistor. The availability of spare boards (25% of those installed) allowed the experiment to keep running during the repair of problematic channels, reducing the unavailability of the SVD just to the time needed to diagnose the issue and replace the board, around two to three hours. As a precaution, at the end of the 2019c run, in all LV boards a systematic replacement of the problematic transistors with ones that have a higher maximum rating was made to avoid similar failures in the future.

A second issue was found that was related to the high voltage boards. It was observed, when setting $HV = 0\text{ V}$, that the leakage current of some sensors increased up to a few hundreds of μA . After tests performed with a high voltage spare board and a spare Layer 5 ladder, it turned out that under particular conditions the internal regulation circuit of the HV board caused an oscillation on the output voltage that could have resulted in a wrong sign current on the sensors. To avoid this possibility, a modification on the regulation circuit was made on one spare HV board, that was first tested with the spare Layer 5 ladder, then installed in the experiment for a couple of months to verify its functionality. After verifying that the modification had no drawbacks, at the end of the 2019c run the modification was systematically applied to all HV boards.

Both modification campaigns were performed by specialized CAEN [108] technicians. Modified boards were first tested by technicians, tested with the spare system and then reinstalled in the experiment. All the 15 LV boards and 5 HV boards were successfully modified.

7.4 Radiation damage

Exploiting the good correlation between the SVD occupancy and the diamond sensors' instantaneous dose rate, it was possible to estimate the integrated dose in the different SVD layers, which is shown in figure 101 as evaluated in 2020. The ladders that received the highest dose are those in Layer 3 that lay on the horizontal plane, as expected. The method has large uncertainty associated, and a more accurate analysis is under development, but this initial estimate was very useful to correlate the effect of radiation damage seen in SVD with the expectations.

The 1 MeV equivalent neutron fluence in Layer 3 mid-plane sensors corresponding to the integrated dose was also estimated, using two methods. In the first method, also used for the discussion in section 2.1.3, the full background simulation is performed to estimate, in the same location, the ratio between the integrated dose in SVD and the equivalent neutron fluence, which can be considered a reasonable estimate, although the absolute value of background from simulation is known to be affected by large uncertainties. Since irradiation results [46] indicate that the energy-dependent damage factors used in the simulation to calculate the 1 MeV equivalent neutron fluence may be overestimated, an independent calculation was also performed. The particle fluence, extracted from the measured occupancy, was translated to the 1 MeV equivalent neutron fluence using the **Non-Ionising Energy Loss (NIEL)** model and a damage factor of 0.01, measured in [46] for 10 MeV electrons. The dose to fluence factor in Layer 3 mid-plane, calculated from the average of the two methods described, is $1.5 \times 10^{12} \text{ neq/cm}^2$ per Mrad.




Figure 101. Integrated dose estimated for all SVD Layers as evaluated in 2020. See references in section 10 for updated radiation study results. Reproduced with permission from [47]. CC BY-NC-ND 4.0.

7.4.1 Noise

As already shown in section 7.3.1, after beam operations started, noise levels increased due to initial radiation damage on the sensor surface, with different effects depending on the sensor position.




Figure 102. Noise evolution of sensor L3.1.1 u/P side between January 2019 (before starting operations) and July 2020. On the left, the change in the noise levels for the u/P side strips for a Layer 3 mid plane sensor is shown; on the right, the evolution of sensor average noise level vs time, from January 2019 to July 2020, is shown. Reproduced with permission from [47]. CC BY-NC-ND 4.0.

An example is given in figure 102: on the left the change in the noise levels for the u/P side strips for a Layer 3 mid-plane sensor is shown; on the right, the evolution of sensor average noise level vs. time, from January 2019 until July 2020, is reported. As expected, the noise started to rise due to the radiation damage effect: on the sensor surface the higher fixed oxide charge, induced by irradiation, increases the inter-strip capacitance and therefore the strip noise. After the first months, in July 2019, the noise had increased by about 15% on both sides. On the v/N side the increase in noise was already almost saturated at that point, while on the u/P side the noise continued to rise during the 2020 runs, but with a lower slope, reaching about 25% increase in the considered Layer 3

mid-plane sensors. The noise increase is also expected to saturate in the u/P side, since the increase in the fixed oxide charge also saturates. In Layers 4, 5 and 6, exposed to lower dose, the increase in noise was at the level of 5–10% only, as of July 2020.

7.4.2 Leakage currents

Radiation damage is also responsible for the increase in the sensor leakage current. The effect is more visible in Layer 3 sensors, where the radiation absorbed is higher. Figure 103 shows the evolution of the leakage current of one backward Layer 3 sensor up to July 2020. The general trend shows a clear correlation of the increase in leakage current with beam operations. Some annealing effects are also visible when the detector was operated with no beams. In the first months of 2020, before the start of the operation with beams, a reduction of the leakage current in small periods is visible in the plot and it is correlated with the operation at lower temperature, which coincided with the PXD being off.

It has to be underlined that during the 2020 run period there were a few changes in the temperature close to Layer 3, by around 2–3 °C, related to changes in the power settings of the PXD, which sometimes caused significant changes in the Layer 3 sensor current. Although these changes were compensated for during data taking by promptly updating the calibration constants, the detailed analysis of the leakage current increase and its correlation with dose is more complicated because of the temperature dependence.




Figure 103. Evolution of leakage currents of one backward Layer 3 sensor up to July 2020.

The leakage current increase for one of the Layer 3 sensors in the mid plane (L3.2.2) is visible in figure 104, where the comparison of two IV measurements taken in January 2019 (before operations with beams started) and in July 2020 is shown. After exposure to around 100 krad integrated dose, the increase in leakage current for the mid-plane sensors was of about 8 µA, both due to surface and bulk damage, although it is difficult to distinguish between the two sources.

The observed increase is in reasonable agreement with the expectation based on the increase in leakage current measured in similar sensors of the BaBar experiment [11] of about 1 µA/Mrad/cm² at 20 °C. The leakage current increase measured in the BaBar sensors, and confirmed with irradiation campaigns to be mainly due to bulk damage from low-energy electrons is a good estimate of the effect expected in the SVD, since the radiation fields in BaBar and Belle II vertex detectors are




Figure 104. Comparison among IV measurements for one of the Layer 3 sensors in the mid plane (L3.2.2) taken in Jan 2019 (line) and July 2020 (crosses).

expected to be similar. With an estimated integrated dose of about 100 krad this increase rate corresponds to an 8 μA current increase in Layer 3 mid-plane sensors (47.5 cm^2 area), assuming a running temperature of 26°C .

A calculation of the expected current increase from NIEL scaling using the estimated neutron fluence was also tried. Considering the current related damage rate $\alpha = 4 \times 10^{-17} \text{ A/cm}$ [48] and the dose to fluence factor of $1.5 \times 10^{12} \text{ n}_{\text{eq}}/\text{cm}^2$ per Mrad, the expected increase in leakage current at 20°C in our $320 \mu\text{m}$ thick sensors would be $1.9 \mu\text{A}/\text{Mrad}/\text{cm}^2$. This is a factor two higher than the estimate calculated with the previous method based on BaBar data, but here no detailed estimation of the annealing effect with temperature and time is considered, and this can reduce the α parameter even by a factor two.

The SVD performance did not change and is not expected to deteriorate with the increase of leakage currents. In fact, even with an absorbed dose up to 10 Mrad, the strip noise contribution from leakage current, thanks to the very short APV25 shaping time, is expected to be very small with respect to the main capacitance contribution.

7.4.3 Depletion voltage monitor

An effective and relatively fast method to monitor possible changes in the depletion voltage due to bulk damage was developed. The depletion voltage for the installed sensors is evaluated with a simple set of calibrations, measuring the noise as a function of the bias voltage applied. With this method there is no need to use beam time to monitor the depletion voltage, as in other techniques based on the reconstruction of the clusters with different bias voltages. In the $\text{p}^+/\text{n}/\text{n}^+$ double-sided sensors used in SVD, the depletion region expands from the P-side, so the N-side strips become insulated only when the N-type bulk is fully depleted. When N-side strips are insulated, their noise drops to a minimum level, indicating that the full depletion voltage is reached. With over-depletion voltage applied, the noise on the N-side continues to slightly decrease: the inter-strip capacitance, which is the dominant noise contribution, is proportional to the ratio between the strip implants width and the pitch. The over-depletion bias reduces the electron accumulation layer present on the N-side surface, which on N-type bulk sensors causes a reduction of the effective width of the N-strip implant.




Figure 105. Example of the noise in a L3 sensor during the July 2020 HV scan, used to measure the depletion voltage. The indicated HV is half of the bias voltage. On the left, the N-side strip noise as a function of the strip number for increasing HV is shown. Red points correspond to full depletion. The noise dependence on the strip number is due to the variation in length (and capacitance) of the pitch adapter traces connecting the sensor strips to the APV25. The bump around strip 400 is due to a slightly lower resistivity in the wafer center, with a locally higher depletion voltage. On the right, the average N-side noise as a function of the applied HV, with the red vertical line indicating the drop at $V_{\text{depl}} = 56 \text{ V}$ ($2 \times 28 \text{ V}$). Reproduced from [44]. CC BY 4.0.

The method was proven to work, with initial scans taken in the spring/summer 2020, when still no changes in depletion voltage were expected given the low integrated fluence of about $1.5 \times 10^{11} \text{ n}_{\text{eq}}/\text{cm}^2$. Results obtained are in good agreement with the depletion voltage measured directly on the sensor before their assembly. An example of the behaviour of the N-side strip noise as a function of the bias voltage is shown in figure 105 for a Layer 3 sensor, showing a depletion voltage of 56 V. All the 172 installed sensors have an initial depletion voltage between 20 and 60 V.

7.4.4 Additional defects after large beam loss

Silicon sensors and electronics are subject to damage when a relatively large amount of radiation is absorbed in a very short time. In accelerators, this can happen when beams become unstable and in case of beam dust events. On June 9, 2019, a quench of some superconducting magnets of the final-focus system was caused by a so called “beam dust” event. During events of this kind beams became unstable, causing a sudden jump in background levels on the detector. It was estimated that in one of these events a radiation dose above 3 rad was delivered on the Belle II detector in $\sim 40 \mu\text{s}$. After this event, local runs were taken from each Belle II subdetector to evaluate any possible damage caused by the high radiation dose deposited on the detector in such a short time. For the SVD, a comparison between the two local runs taken before and after this event showed that a few new pinholes were created. This result was somehow expected as observed in other experiments [11]. In normal conditions the voltage drop across the AC decoupling capacitors is very small, thanks to the floating power supply biasing scheme (section 2.5.3). Bursts of radiation releasing a large amount of charge in the silicon can induce a rapid discharge of the sensor capacitance, causing the full bias voltage (100 V) to be applied across the AC decoupling capacitors. Although these capacitors have a typical breakdown voltage well above 100 V, and all of them were tested up to 20 V, this effect

could induce a breakdown of the oxide layer in the weaker capacitors, creating additional pinholes. The limited number of new defects confirmed the robustness of the AC coupling capacitors and of the sensors. A protocol was established in order to perform local runs and check the health of each Belle II subdetector every time a similar event occurs.

7.5 Observations on data during operation

Analyzing the ADC distributions after the pedestal and common-mode-noise subtraction during the operation, it was found that, under certain conditions, the distributions are shifted by a quantity that depends on the strip position. With further investigations, the sources of this shift were understood and classified into three cases: crosstalk from APV25 data outputs, crosstalk from APV25 power inputs, and baseline tilt at high trigger rate. Although such effects in the ADC distributions are not desirable, the effect on the performance is found to be negligible. In this section, the sources and the characteristics of these three cases are explained.

7.5.1 Crosstalk from APV25 data outputs

A part of the output lines of APV25 (see section 2.2.2) on [origami board](#) has capacitive couplings to the sensor strips, especially where they are aligned in parallel. Thus, the APV25 output can create crosstalk signals on the corresponding strips through the couplings. These crosstalk signals become prominent when the APV25 output changes rapidly at the transitions between digital-low and digital-high signals. This rapid transition happens when the APV25 transmits the tick mark for the synchronization (which is called ‘synchronization pulse’) or the data header [7]. Figure 106 shows the shifts of the ADC distributions observed at the synchronization pulse in a Layer 4 sensor. The most visible shifts are observed in APV4 (see 2.2.4 for the APV25 numbering scheme on origami boards), where the output lines in the origami boards are parallel to the u/P side strips.




Figure 106. ADC distributions for strips in L4.1.2 u/P side, after the pedestal and [common-mode noise](#) subtraction. Data are taken when the synchronization pulse is generated on the APV25 output lines. The ADC distributions are shifted due to the crosstalk from the synchronization pulse.

7.5.2 Crosstalk from APV25 power inputs

The wire-bonds and the bonding pads of the APV25 power inputs on the chip are adjacent to the strip-signal inputs of channels 0 and 127. Therefore, there are capacitive couplings between them.

The current in these power supply lines rapidly changes when APV25 processes data, as reported by the CMS collaboration [49]. Such rapid changes can cause crosstalk on channels 0,1 and 126,127 via the capacitive couplings. The strip-number dependence of the ADC distribution shifts due to the crosstalk is shown in the left plot of figure 107, in which a significant effect can be observed only on channels 0 and 127. The polarity of the spike is inverted between the two edges, reflecting the polarity of the power supply of +1.25 V close to channel 0 and -1.25 V close to channel 127. The time structure of the output in channel 0 when APV25 starts sending data is shown in the right plot of figure 107. The structure is almost identical to that of the power supply current reported by the CMS collaboration; the ADC output also sharply increases at the start of data readout (i.e. at the data header) and decreases 110 clock-cycles before. Due to this time structure, the ratio of the affected events is proportional to the trigger rate.




Figure 107. Left: mean values of the ADC distributions after the pedestal and common-mode-noise subtraction for the strips of one APV25 in L3.1.1 u/P side at the start of the APV25 data readout. The polarity of the spike is opposite between the two edges of the APV25 channels, reflecting the polarity of the power supply of ± 1.25 V. Right: time dependence of the mean values at channel 0 in the same APV25 chip. The green shaded areas correspond to the data headers and the green dashed line is the timing of the retrieval of the data samples from the pipeline which happens 110 clock-cycles before the data header transmission.

7.5.3 Baseline tilt at high trigger rate

When the trigger rate approaches 30 kHz, another shift of the distribution is observed that corresponds to a linear increase with APV25 output order, as shown in the left plot of figure 108. Since the data are read out only through AC-coupling in FADC, the circuitry tries to balance out any excursions, so that the long-term average becomes equal to the DC operating point of each channel. This has some implications depending on the trigger rate. At a low trigger rate, the digital baseline is by far the dominant signal, so it is almost equal to the operating point. In that case, an exponential decay in the analog strip signals is observed, because those are higher than the digital baseline. This is also the condition in which the pedestals are recorded. However, at a high trigger rate, the actual strip data become the dominant signal, while the digital baseline largely disappears, and then the AC coupling mandates that the average strip signal is leveled out to the operating point. In that case, the strip signals are flat, and the exponential decay is observed only in the (rarely occurring) digital baseline. This also implies that the pedestal subtraction overcompensates a slope that is not present

anymore, and therefore introduces the observed rising slope. The tilted slope is linearly increasing with the trigger rate, as shown in the right part of figure 108. The tilt cannot be mitigated by the current pedestal nor common-mode-noise subtraction. The characteristic of this baseline tilt also appears in the occupancy. If the baseline is above the average, the noise is more likely to exceed the signal threshold. Thus, the occupancy increases in half of the sensors, and decreases in the other sensors, but the effect is of the order of 10^{-4} , so the performance is barely affected.




Figure 108. Left: ADC distributions for strips of one APV25 in L4.1.2 u/P side, after the pedestal and common-mode-noise subtraction. The data are taken with a 30 kHz Poisson trigger. The circle marker represents the mean value for each column, and the bar represents the standard deviation. The baseline tilts up to around ± 1 ADC. Right: the slope of the linear-fitting in the same APV25 chip for 1, 5, 10, 20, 30 kHz trigger rate. The fitting uncertainties are smaller than the point size. There is a linear correlation between the slope and the trigger rate.

7.6 APV trigger veto

The pipeline readout of the APV25 is implemented as a ring buffer and a FIFO (see 2.2). It is notable that the chip is not equipped with a mechanism to prevent the overflow of the FIFO. Without countermeasures, 6-sample readout at a trigger rate of 30 kHz causes a FIFO overflow error in a few milliseconds, which then requires resetting the chip. To overcome this error, the “APV trigger veto” module was developed and implemented in the firmware of the most upstream [Front-end Timing Switch \(FTSW\)](#) [17]. The module monitors trigger and reset signals that are eventually sent to the APV25 chips and calculates the future FIFO occupancy. The triggers are vetoed if they would overflow the FIFO. In this way, the system prevents the FIFO overflow with minimal dead time, rather than sending busy signals from front-end electronics. According to a simulation shown in figure 109, the dead time due to the APV trigger veto is around 3% for 6-sample readout of 30 kHz triggers following the Poisson distribution. The simplest alternative, vetoing until the FIFO gets empty after every trigger, would yield 44% dead time. With the mixed 3/6-sample acquisition mode the dead time would be suppressed to less than 1% if the fraction of 3-sample triggers is more than 0.4.

The trigger veto module was integrated to the DAQ since Phase 2 and is working as expected.

7.7 Beam loss monitor and aborts, environmental monitors and interlocks

The monitoring of environmental parameters is a critical aspect of the operation and performance stability of SVD. Some key aspects of the operational experience concerning the radiation monitor




Figure 109. Dead time due to APV trigger veto versus input trigger rate in mixed 3/6-sample readout. The fraction of 3-sample triggers, R_3 , is changed.

and beam abort, the temperature and humidity monitors, and the hardware interlock are summarized here. The details of these systems are described in 4.

After the installation and commissioning of the full system of 28 diamond sensors, in 2019 the DCU firmware was modified, to allow a faster abort cycle: $2.5 \mu\text{s}$ instead of $10 \mu\text{s}$, as described in section 4.1. As a result of the experience and improvements, the full system was recently operated with 24 diamond detectors in monitor mode (in the most sensitive 36 nA range) and four beam-pipe diamond sensors dedicated to abort requests (in the 4.5 mA range), providing effective protection to both VXD and QCS magnets. The diamond system issued 305 beam-abort requests in 2019, and 655 beam-abort requests in the first semester of 2020. Two abort thresholds were recently applied in parallel:

- “very fast”: 4 mrad integrated in a moving time-window of $10 \mu\text{s}$, updated and compared with threshold every $2.5 \mu\text{s}$;
- “fast”: 40 mrad integrated in a moving time-window of 1 ms, with the same abort cycle.

The “very fast” threshold detects sudden large spikes in beam losses, while the “fast” threshold is meant to catch a slower building up of unwanted radiation doses. Most aborts were of the “very fast” type.

Data from the most sensitive “range 0” (section 4.1) diamond sensors, read out at 10 Hz, have also been used to set up “injection inhibit” algorithms, to veto continuous injection when it results in large radiation doses.

The 10 Hz data are also used for the measurement of integrated radiation doses. The integrated dose is underestimated, if range 0 is saturated during the short injection time intervals. An upgrade program to improve the dynamic range of the diamond electronics and to avoid or mitigate such saturation effects is now considered.

The temperature monitoring by NTC and FBG on FOS sensors (section 4.2) was in operation without failures since their installation. Archived monitor data provided useful correlations with the detector operation conditions. Hardware interlock signals on preset temperature thresholds behaved as expected.

The humidity monitoring system (section 4.3), after about two months of continuous operation in 2019, experienced some failures of Vaisala Dew Point Transmitter DMT143 sensors. The problem was traced to the fact that the frost point of the dry volume went lower than originally expected, typically down to about 85 °C, outside the measurement range of the instrument (−80 °C to 20 °C), causing very frequent self-calibrations and finally error conditions. The substitution of these sensors with Vaisala DMT152 (−100 °C to 0 °C) fixed the problem.

Some hang-ups also occurred in the USB read-out connection: at some point no more data were coming from one of the four Vaisala sensors. As a result, the [VXD Local Hardwired Interlock \(VLHI\)](#) system immediately shut down the VXD and set the IBBelle cooling plant (4.5) to warm operation (+15 °C), as expected when there is a failure of one of the environmental monitoring systems. The reason for the failure was found between the communication board of the humidity system and the corresponding [Input/Output Controller \(IOC\)](#): data corruption was interpreted by the IOC as a malfunction of the Vaisala sensor, causing the VXD interlock. Improving the code library used to manage the communication between the humidity system and the corresponding IOC made the communication more stable and avoided similar problems in the next runs. As noted in section 4.3, a new system, with different redundant humidity sensors and more robust Ethernet readout, was built at INFN Trieste and is ready for shipping and installation.

The VLHI hardware interlock system, described in section 4.4, was in continuous operation without failures. It protected the VXD in a few events, due to alarms from the CO₂ cooling system or from water leaks sensors, detecting condensation of water on some cryogenic transfer lines. A complete spare system is ready for shipment and installation for the long-term operation of the experiment.

A small issue that caused unexpected VXD interlocks at the beginning of 2019a run was due to the chiller used for the dock-boxes cooling system. The chiller is connected to the VLHI system so that in case of failure the SVD is automatically shut down. At the beginning of 2019a run some interlocks were issued by the chiller line, although the chiller was working properly. Looking at the output of the low-pass filter module of the chiller that sends the signal to the VLHI hardware, spikes were observed when one of the chiller's relays switched. These spikes gave fake inputs to the VLHI system, that asserted the interlock at each time. The input was temporarily disabled and subsequently a software interlock was implemented on the water flow of the chiller. The possibility of filtering the short spikes coming from the chiller was considered and implemented in 2020, to re-enable the hardwired interlock of the chiller on the VLHI system.

7.8 Background studies

The challenging design luminosity of SuperKEKB means that particular care must be taken of the machine induced background conditions, which are expected to be severe and can affect the lifetime of the detector and its performance. The main background sources in SuperKEKB can be classified in two groups: single beam backgrounds (Touschek effect, beam-gas scattering, synchrotron radiation, injection background) and luminosity related backgrounds (radiative Bhabha scattering, two-photon processes). A brief description of these background sources and results of the first background studies performed in Phase 1 are given in [35]. Before the start of Phase 2, preliminary Monte Carlo simulations were used to estimate the expected background levels at design machine parameters. Figure 110 shows the estimated SVD occupancy levels due to backgrounds, with the contributions of each background source. This initial background extrapolation at design luminosity gave on average

for Layer 3 sensors an estimate of about 1% occupancy, 0.1 Mrad/yr integrated dose, and an equivalent neutron fluence of about $0.2 \times 10^{12} \text{ n}_{\text{eq}}/\text{cm}^2/\text{yr}$. Considering the radiation hardness of the DSSD sensors, conservatively set to 10 Mrad and about $3 \times 10^{13} \text{ n}_{\text{eq}}/\text{cm}^2$ (as discussed in section 2.1.3), these numbers give enough margin to safely operate the SVD even for 10 years at design luminosity. Also the extrapolated occupancy was well below the 3% limit set to ensure good tracking performance, and it was considered promising for the operation of the VXD and the full Belle II detector.




Figure 110. Occupancy of all SVD Layers estimated by preliminary Monte Carlo simulations at design machine parameters, $L = 8 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$. Different contributions from luminosity related background (two-photon processes and radiative Bhabha, which include “RBB”, “RHWide” and “BHWide LA”, referring to different photon emission angles) and single beam related background (beam-gas scattering and Touschek effect) are shown.

In order to understand the real background levels in the VXD volume, dedicated studies were conducted during Phase 2, with the machine in its final configuration, the Belle II detector in position, but with only one slice of the Vertex Detector. The remaining volume of the VXD was instrumented with the BEAST II detector system, designed to measure the radiation field in the inner detector region. The full description of the system and the results of the studies are available in [42]. After the VXD installation and the start of Phase 3, many studies were conducted during every data taking period by machine and detector groups to carefully evaluate background levels, in order to separate background components and to optimize collimators settings, to lower the backgrounds in the interaction region preserving the beam lifetime and injection efficiency [50].

Background studies allowed a better understanding of background levels in the interaction region. Data from these studies and comparisons with simulations were used to re-scale MC backgrounds at design machine parameters to reflect the real background conditions in the interaction region. The re-scaled MC background levels, according to recent background studies based on 2020 runs data set, are shown in figure 111, again with contributions of each background source.

A comparison with figure 110 shows that the SVD background levels in the re-scaled MC are much higher by almost a factor 3: for Layer 3 sensors about 3% occupancy, 0.3 Mrad/yr integrated dose and an equivalent neutron fluence of about $0.6 \times 10^{12} \text{ n}_{\text{eq}}/\text{cm}^2/\text{yr}$ are expected. This estimation,




Figure 111. Occupancy of all SVD Layers after re-scaling MC simulation using background studies measurements based on 2020 data set. Different contributions from luminosity related background (two-photon processes and radiative Bhabha, which include “RBB”, “RHWide” and “BHWide LA”, referring to different photon emission angles) and single beam related background (Bremsstrahlung, Coulomb scattering, and Touschek effect) are shown. Design machine parameters are considered, $L=8 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$.

though still affected by a large uncertainty, is more reliable than the preliminary one shown in figure 110, because it is calculated using data taken with the actual machine and detector. In terms of integrated radiation damage, even with the re-scaled background estimate there are no concerns to safely operate the SVD for several years at design luminosity. However, the estimated occupancy level is close to the 3% limit for good tracking performance, and further mitigation strategies are now considered to increase the safety margin.

With the running conditions and machine parameters used in 2019 and 2020, background levels were quite low ($< 0.4\%$) and did not affect detector performance. To mitigate the higher background levels expected at design luminosity, some countermeasures are being evaluated and implemented. On the machine side, new beam collimators were installed, the design of additional ones is being considered, and a new beam pipe for the interaction region is being prepared, allowing the installation of better optimized heavy metal shields. On the SVD side, the 3-sample DAQ mode and the SVD hit time, described in section 9.6, can be used to reduce the required bandwidth, maintaining the tracking efficiency in high background conditions. More details of the strategy to cope with high background conditions are given in section 10.

In conclusion, operations of the SVD was smooth in the first two years of the experiment. A constant periodic monitoring has shown the expected evolution of the relevant parameters of the silicon sensors, with a few issues that are under control, well understood, and not affecting performance.

8 Simulation and reconstruction software

The SVD plays a key role in the Belle II reconstruction, both for charged particle tracking, and for low momentum particle identification. Tracks reconstructed in the CDC are first extrapolated to the SVD and then to PXD, to significantly improve the precision of the track parameters. The

SVD is also used for stand-alone pattern recognition, specifically for low-transverse-momentum tracks (approximately below 90 MeV/c) that do not reach or are not reconstructed in the CDC because of their bending in the magnetic field. In addition, using the energy loss in the material, the SVD significantly improves the particle identification capability below 1 GeV. As an example, this capability is essential to reconstruct the slow pion of the $D^{*+} \rightarrow D^0 \pi_s^+$ decay from B decays, which is crucial for the Belle II physics program.

The complete tracking chain is run online in the High Level Trigger (HLT) to determine the regions of interest (ROI) on the PXD sensors in which hits from tracks are expected, allowing to reduce by one order of magnitude the PXD data, which would otherwise violate storage and bandwidth limits. Hence, there are quite stringent limits on the execution time of the tracking chain (including SVD reconstruction) as well as on the quality of the tracking that must be achieved during data taking, i.e. before the dedicated calibration of the tracking detectors and algorithms that is usually done *using* the data that have already been collected. In fact, a sub-optimal quality of the online tracking, especially the SVD part, may result in a permanent loss of PXD signal clusters with in a significant degradation of the track quality.

The high quality and precision needed for the Belle II physics program require the SVD detector to be efficient and precise in the reconstruction of the hits produced by charged particle crossing the sensors. One of the main challenges of the SVD reconstruction is the reduction of the background hits produced by uninteresting electromagnetic processes, such as $e^+e^- \rightarrow e^+e^-e^+e^-$, occurring hundreds of nanoseconds before the triggered events. Since the bunch crossing frequency is ≈ 250 MHz, and the tail of the response of the APV is order of 400 ns long, the SVD sees particles hitting it since roughly 100 bunch crossings before the triggered event. These off-time hits populate the detector, increasing significantly its occupancy from 10^{-5} to a few percent, and making the pattern recognition much more challenging. The measurement of the hit time is particularly powerful to remove this background: a simple selection on the hit time allows to remove $\approx 40\%$ of the background 1-D hits while keeping an efficiency above 99.9%.

The main technical features of the SVD software are summarized in section 8.1: integration in the Belle II framework, reconstruction data flow, simulation workflow (section 8.1.3), and their interactions with the data store and Conditions Database (section 8.1.4).

The reconstruction (section 8.2) starts with the unpacking of the raw data and the calibration of the strips (section 8.2.1), that are then combined in clusters, representing the 1-D hit on the sensor plane. The computation of the cluster position and time is done using the strip time and charge. The cluster charge is calculated summing the strip charges and is then used for the particle identification. As described in section 8.2.2, clusters on the two sides of the sensors are finally combined in spacepoints that represent the 3-D hits in space and are the key ingredient for the track finding algorithms. At this stage, clusters are rejected based on their hit time.

An essential part of the reconstruction performance is the calibration of the detector and of the reconstruction algorithms. The strip noise and gain calibrations are measured in local runs, and stored in the Conditions Database. The calibration of the hit time is done exploiting its correlation with the time of the event with respect to the trigger arrival (section 8.2.4). Calibrations of each sensor side are also stored in the Conditions Database.

In order to monitor the quality of the data during data acquisition, the reconstruction of a fraction of the events is performed online on dedicated computer farms. The Data Quality

[Monitoring \(DQM\)](#) software (section 8.2.3) produces a set of plots that are shown live in the control room and automatically change colors in case some anomaly is detected. This is essential to immediately identify and react if something is not working as expected. A set of expert-mode plots are also available to identify issues in the shortest time possible, to maximize data taking efficiency.

The accurate modeling and simulation of the interaction of charged particles with the silicon sensors, the APV response to the induced charge, and the DAQ chain is very important. In the first place, the track finding algorithms are calibrated and optimized using simulated data: a mis-modeling of the data would result in a degraded tracking performance and would contribute to increase the systematic uncertainty for the physics analysis. Furthermore, the simulation provides an essential tool to understand the features observed in the data, and to predict the reconstruction performance for different configurations of the detector and for varying external conditions, such as an increased background.

As outlined in section 8.1.3, the simulation starts from generated MC particles, whose interaction with the material is simulated by Geant4 [51], and ends with a list of raw strips, equivalent to the output of the unpacker on the real data. The sensor response is obtained by producing a set of small groups of ionization charges along the track, that are then drifted in a simplified electrical and magnetic field towards the electrodes on both sensor sides reaching the collection implants. The charge collected on the implants is coupled to the nearby floating and readout strips, to simulate the charge sharing. No weighting field is used to model signal induction during charge drift. The APV response is then simulated, including the strip gain, while the electronic noise is taken directly from random trigger data in run-dependent simulations. Neither [common-mode noise](#) correction, nor the [FIR](#) are simulated. The calibration of the charge sharing parameters as well as of the strip gain was performed using test-beam data. This simplified model provides a fair agreement for cluster charge signal-to-noise ratio, but underestimates the size of the clusters with respect to data. A study is ongoing with a re-parameterization of the charge sharing variables and a re-calibration with data in order to improve the agreement in the cluster charge and cluster size.

8.1 Belle II SVD software

The SVD Offline Software is fully integrated into the Belle II official software with a dedicated package in the Belle II Software Analysis Framework (*basf2*) [41]. The full *basf2* source code is publicly available in [52].

The building blocks of the *basf2* processing chain are the *modules*, each having a particular task that goes from a simple task like reading the data from disk, to more complex ones like clusterization. A python steering file is used to define an ordered list of modules that are executed one after the other for each event, producing the required output. The input data of a module (e.g. list of strips of the event) are stored in a common storage, the DataStore. Each module has read and write access to the storage. Run condition information are also needed for certain modules (e.g. strip noise calibration), and are stored in the Conditions Database. As an example, in figure 112 a schematic view of the data flow of the SVD offline reconstruction with indications of modules, objects read and written to the DataStore and conditions taken from the Conditions Database is shown.

The SVD package in *basf2* contains the necessary software that allows to:




Figure 112. SVD Offline Reconstruction data flow. Modules, inputs from the DataStore and inputs from the Conditions Database are shown. For the sake of clarity, the prefix SVD is removed. The conditions in green are uploaded on the database during data taking, while the ones in blue are computed offline.

- unpack real data and simulate the strip signals recorded by the sensor in response of a charged particle traversing it, processed by the APV25, and digitized in the FADCs;
- reconstruct the hit time, position, and ionization charge from the digitized signal of the strips;
- calibrate the simulation and reconstruction algorithms, storing and retrieving the information from the Conditions Database;
- monitor the detector performance during data taking.

The simulation and calibration software run offline, while the reconstruction and monitoring software run offline for data reprocessing, and are also executed online on the High Level Trigger farm for event selection and on dedicated machines for the online data quality monitoring (ExpressReco). In particular, the SVD reconstruction (and tracking) runs on the HLT to identify regions on the PXD where the track has most probably crossed the sensor (ROI) in order to reduce the data size of the PXD, both for bandwidth and storage limitations. The required data reduction factor is 10, with an efficiency which must be close to 1 in order not to lose the precious information of the PXD hit for offline reconstruction. Since SVD reconstruction runs online, there are additional requirements on the software, especially in terms of the execution time and the memory usage. Moreover, the online reconstruction needs to work reasonably well with calibrations (e.g. for clustering and hit time computation) determined on previously recorded data, while for the offline reconstruction the calibration is done using a subset of the runs that will be re-processed. The tracking pattern recognition algorithms are trained on simulated events, therefore a good simulation of the SVD is fundamental to have a high tracking efficiency.

8.1.1 Local coordinate systems

Apart from the global coordinate system defined in section 2, each SVD sensor has a local right-handed coordinate system (u, v, w) centered at its center, with v in the direction of z , w pointing away from the IP, and u pointing along the global $r\phi$ direction. The N strips measure the local v

coordinate, while the P strips measure the local u coordinate. In the slanted sensors (forward sensors of layer 4, 5, 6), because the u/P strips are not parallel to each other but form a varying angle with the global $r\phi$ direction, the u position of the hit can be determined only if its v position is also known.

8.1.2 Time reference frames

There are two main time reference systems: the SVD reference system with $t_{\text{SVD}} = 0$ at the time of the first sample of the APV, and the global Belle II reference system with $t_{\text{glob}} = 0$ at the trigger arrival time, simultaneous for all subdetectors. The two systems are synchronized, but the APV25 clock frequency is 4 times smaller than the trigger clock. In order to switch from one reference frame to the other, the FADC registers the trigger signal arrival time within one APV clock, then stored it in the `trigger bin` variable. For example, if the trigger signal arrives during the first quarter of the APV clock, the trigger bin is 0. Therefore, the trigger bin allows us to translate the hit time measured in the SVD reference frame to the trigger reference frame, common to all detectors, and vice-versa.

8.1.3 SVD simulation

The SVD Simulation covers the part of `basf2` software that simulates SVD response to particles traversing the sensitive volumes of SVD sensors. In figure 113 a schematic view of the data flow of the SVD simulation, with indications of modules, objects written to the DataStore, and conditions taken from the Conditions Database is shown.




Figure 113. SVD simulation workflow, showing `basf2` modules and objects in the DataStore that form the SVD simulation.

Technically, detector simulation is a three-step process, with *physics event generators* simulating collisions in the accelerator and other relevant radiation-generating processes (step 1), and the Geant4 [51] engine propagating the output of physics generators through the detector (step 2). The Geant4 tracker uses callbacks to individual subdetector routines to defer the calculation of detector response from energy depositions in the sensitive subdetector volumes to the dedicated subdetector software (step 3). The last process is called *digitization* and is the main topic of this section.

There are several physics event generators supported in `basf2` as part of its *externals* software package, including `EvtGen` [53], `PYTHIA8` [54], and `KKMC` [55]. All such generators produce a list of Monte Carlo (MC) Particles, `MCParticle`. A crucial part of `basf2` software is a geometric representation of the whole Belle II detector and (parts of) the accelerator, as well as the related mechanical and civil structures. The geometry is implemented as a hierarchy of Geant4 volumes

and stores its parameters in the Belle II Conditions Database. The geometry is built subdetector-wise, with each subdetector defining its geometry builder class (`GeoSVDCreator` for SVD) and its `SensitiveDetector` Geant4 callbacks.

SVD simulation starts with a collection of `MCParticle`'s tracks going through sensitive the regions of the SVD and ends with a collection of `SVDShaperDigits` objects, which are the same objects used to encode the data on fired SVD strips during the real data acquisition process, thus producing data that imitate what is provided by the SVD hardware for reconstruction.

The SVD Sensitive Detector class is the piece of software that gets a Geant4 (MC Particle) track from the Geant4 tracker and stores its energy deposition to be used by the `SVDDigitizer`. For this, the software uses objects called `SVDSimHits`, which represent segments of a piecewise-linear approximation to the MC Particle track through the sensor. They store a detailed profile of energy deposition along the segment. The `SVDSensitiveDetector` also forms another collection of data objects called `SVDTrueHits`. An `SVDTrueHit` represents a single MC Particle passage through an SVD sensor. These objects facilitate a simple MC-truth-based tracking.

Sensor Information. Each silicon sensor has an associated `SensorInfo` instance that stores sensor-specific parameters, including position, strip electric parameters (strip capacitances and electronic noise for the simulation) and material constants. The `SensorInfo` class also has methods to calculate some simulation-related quantities for the sensor, such as the electric and magnetic fields, charge carrier mobilities, and drift velocities. The parameters and equations used are listed in table 22.

Digitization is the process of calculating strip signals from energy depositions in the sensor. The digitizer gets the information about SVD data acquisition and trigger arrival from the `SVDEventInfo` object. The latter stores the DAQ mode (6- or 3-sample acquisition), the trigger type, and the [trigger bin](#).

The `SVDSimHits` store detailed information about energy deposition used in the digitization four steps:

1. create charge carrier clouds based on the energy deposition pattern;
2. allow charge carriers to drift and diffuse in the electric and magnetic field in sensor bulk, reaching the collection implants;
3. calculate the signal that this induces on the sensor strips through capacitive coupling;
4. process the raw strip signals by taking into account strip noise, charge sharing, and readout electronics.

The *basf2* implementation of SVD digitization is implemented in the `SVDDigitizer` class. Digitization starts with energy deposition profiles in individual `SVDSimHits`, encoding energy deposition along a quasi-linear piece of a particle track through a silicon sensor. It also includes information about the time of MC Particle passage through the sensor. The piece of track is divided into several small pieces of 5 μm , and the energy deposited along each piece is converted into free charge carriers — electrons and holes — by dividing the energy deposition by the conversion energy for silicon, 3.65 eV. The number of charge carriers is randomized by the Poisson smearing. In the following, these small clouds of electrons and holes are called *chargelets*. A chargelet is characterized by its position within a sensor, the number of carriers it contains, and time of its creation.

Table 22. Simulation-related quantities provided by `SVD::SensorInfo` methods. The value of the temperature is $T = 300$ K.

Quantity	Formula	Parameters
Electric field	$E(z) = \frac{2V_{\text{depl}}}{d} \left(\frac{z - d/2}{d} \right) - \frac{V_{\text{bias}} - V_{\text{depl}}}{d}$	$V_{\text{depl}} = 40$ V $V_{\text{bias}} = 100$ V $d = 300 - 320$ μm
Carrier mobility [56]	$\mu(E(z), T) = \frac{v_m/E_c}{\left(1 + \left(\frac{E}{E_c}\right)^\beta\right)^{1/\beta}}$	Electrons: $v_m = 1.53 \cdot T^{-0.87} \cdot 10^9$ cm/s $E_c = 1.01 \cdot T^{+1.55}$ V/cm $\beta = 2.57 \cdot T^{+0.66} \cdot 10^{-2}$ Holes: $v_m = 1.62 \cdot T^{-0.52} \cdot 10^8$ cm/s $E_c = 1.24 \cdot T^{+1.68}$ V/cm $\beta = 0.46 \cdot T^{+0.17}$
Hall factor	$r_H(T) = r_H^0 + r_H^T \cdot (T - 273.15 \text{ K})$	Electrons: $r_H^0 = 1.13$ $r_H^T = 0.0008$ K ⁻¹ Holes: $r_H^0 = 0.72$ $r_H^T = 0.0005$ K ⁻¹
Drift velocity	$\mathbf{v}(\mathbf{E}, \mathbf{B}) = \frac{\mu\mathbf{E} + \mu\mu_H\mathbf{E} \times \mathbf{B} + \mu\mu_H^2\mathbf{B}(\mathbf{E} \cdot \mathbf{B})}{1 + \mu_H^2\ \mathbf{B}\ ^2}$	μ carrier mobility, $\mu_H = \mu \cdot r_H$ Hall mobility

Chargelets drift in the sensor bulk under the combined effect of the electric and magnetic fields (including the Lorentz shift effect). Mobilities and velocities of charge carriers are calculated using the formulas in table 22. Diffusion is also simulated by increasing the size of the chargelet cloud proportional to \sqrt{Dt} , where D is the diffusion coefficient.

A simplified model that does not include the effect of induced signal on the strips due to charge moving in the sensor (Ramo-Shockley theorem [57, 58]) is used in the simulation. The signal is instead produced only when the charge is *collected* on the sensor surface and it is assigned to the nearest floating or readout strip. A single chargelet can produce signal on both a readout and a floating strip, assuming a Gaussian distribution of the carriers in the chargelet a fraction of the charge corresponding to the area of the Gaussian distribution is assigned to each strip.

After this step, charge sharing among the various implants is simulated using a simplified model based on the circuit shown in figure 114, that takes into account:

- C_i : the interstrip capacitance between an implanted strip and its first neighbouring strip. The capacitance between one implant and the second neighboring strip is not considered, as it is expected to be negligible. C_i varies from 0.5 to 1 pF/cm depending on the different sensor geometry and sides. These estimated values are affected by quite a large uncertainty, depending not only on details of the sensor design, like the ratio of the width of the implant over the




Figure 114. Sensor capacitance coupling scheme implemented in simulation.

implant pitch (half the readout pitch), but also on the positive oxide charge on the sensor surface. The latter can be different from sensor to sensor, and also varies with irradiation.

- C_b : the capacitance to the back of the sensor. It can be calculated from the geometrical parameters of the sensor and it scales with the implanted pitch. C_b is about 0.1 pF/cm for u/P side, with smaller pitch, and about 0.4 pF/cm for v/N side with larger pitch.
- C_c : the decoupling capacitance to the APV input. It is more than an order of magnitude larger than C_i and C_b , with measured values of 15 (30) pF/cm on the u/P (v/N) side. With this large value, the charge collected on a readout implant is efficiently collected on the corresponding APV input.

The coupling coefficients among implanted strips calculated considering the simplified network shown in figure 114 are the following:

- k_1 , the fraction of the charge on a floating strip coupled to each neighboring readout implant:

$$k_1 = \frac{C_i}{(2C_i + C_b)}$$

Ideally it should be 50%, which is the case for small pitch, i.e. small C_b . In this case even when the charge is collected on a floating strip a full collection of the signal on the two nearby readout strips is obtained. For large pitch, when C_b becomes comparable to C_i , some fraction of the signal is instead lost. In the layer 4, 5, 6 sensors for the v/N side k_1 is 35–40%, and a total charge loss due to this effect can be 20–30%, as observed in data.

- k_2 , the fraction of the signal on a readout strip coupled to the next readout implant:

$$k_2 = \frac{1}{2} \frac{C_i}{(C_i + C_b + C_c)}$$

Since C_c is much larger than the other capacitances this coupling is very small and the signal coupled from one readout strip to the next one is in general below threshold and then lost. This effect can give rise to an additional few percent charge loss, larger on the u/P side where C_c is smaller compared to the v/N side.

After charge sharing simulation, the charge of each readout strip is smeared with a Gaussian distribution whose width includes the Fano factor [59].

Finally, the response of the APV for each signal on each readout strip is simulated by the waveform of the response as a function of time. The start time of the waveform is computed as half the drift time of the chargelet to the sensor surface. For each strip the APV responses of each chargelet is summed, and electronic noise is added on top of the signal waveform, obtaining the final APV response. This waveform is then digitized by sampling it with the APV clock frequency and converting its amplitude in 8-bit ADC counts with a conversion factor, `ADUEquivalent`. The Online zero suppression is applied to limit the number of `SVDSHaperDigits` produced in output of the simulation. In particular, only `SVDSHaperDigits` with at least one sample larger than three times the noise of the strip are stored, as done in real data acquisition.

Simulation calibration. There are several parameters in the simulation that need to be tuned separately for the two sensor sides of each sensor type (small rectangular, large rectangular, and trapezoidal or wedge):

- noise superimposed to the APV waveform;
- `ADUEquivalent` to convert electrons into ADC;
- the three capacitances for the charge sharing description. For C_c and C_b the design parameters provide already a good estimation, while C_i is much more difficult to estimate from design parameters. Therefore, only the latter has been tuned on data.

The noise in electrons was estimated using the measured noise (in ADC) converted in electrons using the measured gain. For the tuning of C_i and `ADUEquivalent`, 2017 test-beam data (see table 6) was used. In the test-beam, performed with high energy electrons crossing the sensor with perpendicular incidence, most of the clusters have a cluster size 1 or 2, depending on the impact point of the tracks. Only a small amount of clusters have higher cluster size and are not used for this tuning. Due to the coupling coefficients k_i described above, the total cluster charge for cluster size 1 and 2, Q_{cls1} and Q_{cls2} are sensitive to both parameters C_i and `ADUEquivalent`, since the other capacitances in the charge sharing model are assumed to be well known and fixed. The ratio $R = Q_{\text{cls2}}/Q_{\text{cls1}}$ is instead only sensitive to C_i and this sensitivity can be exploited to tune the interstrip capacitance C_i with data in the following way. From the full simulation the dependence of the ratio R from C_i is extracted; then the value of R measured on test-beam data is compared with simulation to extract C_i for the various sensors types and sides. After this step, the total cluster charge in data and simulation is compared to extract the `ADUEquivalent` value. A quite good agreement for the total cluster charge distributions between simulation and test-beam data is achieved with this method, as shown in figure 115 for a sensor in a Layer 5 ladder. For cluster size 2 the effect of charge loss of about 20% on the v/N side with respect to the u/P side is clearly visible, which is related to the presence of a floating strip and larger pitch on the v/N side, as described above. Although in test-beam data a fair agreement with the simulation was achieved, further tuning of the charge sharing model is now ongoing to improve the matching of the cluster properties between simulation and collision data.

Timing in simulation. In order to simulate trigger jitter, the time of all `MCParticles` of an event are randomly shifted using a Gaussian distribution centered at 0, and with a width of 10 ns; consequently, also the time of `SVDSimHits` are shifted. The simulation also handles the start of the




Figure 115. Comparison of cluster charge between test-beam data and simulation after tuning of the C_i and ADUEquivalent values for Layer 5 sensors, the u/P side on the top row and the v/N side on the bottom row; for cluster size 1 on the left column and cluster size 2 on the right column. As described in section 2, L5.x.3 means all the sensor in layer 5, any ϕ position, z position 3 (that is, all Origami_-Z sensors in layer 5).

APV sampling with respect to the arrival of the trigger. As explained in section 8.1.2, in real data taking the quarter of the APV clock period in which the trigger arrived is stored as the trigger bin. In simulation the trigger bin is chosen as a random number between 0 and 3, and the beginning of the sampling is shifted by the trigger bin times the APV period divided by 4 (≈ 7.86 ns).

Run-dependent simulation. In addition to the run-independent simulation (same conditions for all events), the run-dependent simulation is produced, to properly take into account effects of the detector that vary from run to run. Strips fired on noise and positive signals from the cross-talk effect described in section 7.5.1 are automatically, and realistically, added to the run-dependent simulation by overlaying random-triggered real events to simulated signal events. The overlay allows to sum sample-by-sample the SVDSHaperDigits from simulation to the ones in real data. The negative signals from cross-talk are not part of the simulation, but studies on efficiency and cluster position resolution do not indicate that the performance is sensitive to these negative signals. Finally, all the strips and APV chips disabled during data taking for the simulated run are masked. This information is stored on the Conditions Database.

Simulation performance. The reader should note that, at this point of the section, the reconstruction has not been discussed yet, therefore this section contains references to later parts of the text, in particular to section 9 for what concerns data.




Figure 116. SVD cluster SNR distributions for data (black) and simulated $B\bar{B}$ plus beam background (green) for the u/P side on the top row and the v/N side on the bottom row. L3 clusters on BW sensor on the left, and L6 cluster on one of the Origami sensors on the right.

The simulation reasonably describes the cluster properties for the collision data. A fair agreement was found for the cluster time, as can be seen comparing figure 130 for simulation and figure 141 for data. The collected signal and the signal-to-noise ratio are also well simulated on the u/P side, while on the v/N side some data-MC differences up to 30% were seen for the layers 4, 5, and 6 sensors with larger pitch, indicating that still some tuning of the charge loss and sharing model is needed to improve the agreement. This can be seen in figure 116. The cluster size is systematically smaller on simulation with respect to collision data, especially for the u/P sides, which have a smaller pitch, as shown in figure 117.

The disagreement in cluster size is an indication that the simplified signal formation model implemented in the simulation is not good enough to reproduce the collision data. This level of disagreement was not observed in test-beam data used to tune the simulation, in which there was a much better agreement in cluster size between data and simulation. This difference could be due to a different online zero-suppression criterion applied during test-beam ($\text{SNR}_{\text{strip}} > 5$) with respect to physics runs ($\text{SNR}_{\text{strip}} > 3$). A tighter zero-suppression criterion reduces the effect of the noise on the cluster size, but it degrades the cluster position resolution. Beside this, there is another significant difference between data and simulation regarding the strip time. In data, the strip at the edge of a cluster of size 3 has a hit time significantly smaller than that of the central strip, as shown in figure 118. The effect is more evident for the u/P side with smaller pitch, than the v/N side. Both these effects are probably due to the fact that the signal induced by moving carriers (Ramo-Shockley theorem) is not simulated: the fast APV25 may catch the signal induced by charges




Figure 117. SVD cluster size distributions for data (black) and simulated $B\bar{B}$ plus beam background (green) for the u/P side on the top row and the v/N side on the bottom row. L3 clusters on BW sensor on the left, and L6 clusters on one of the Origami sensors on the right.




Figure 118. Time distributions of the strips belonging to layer 3 clusters of size 3 for perpendicular tracks in di-muon events reconstructed in several runs of 2019 data. The time distribution of central strip is shown in red, while the overlaid time distributions of the left (right) edges strips is shown green (blue).

moving in the weighting field even if this signal integrates a null charge. Studies using an external software [60] confirm that the timing structure of the strip in the clusters is reproduced when the signal formation is simulated with the Ramo-Shockley theorem. Further studies and simulation tuning are now ongoing to improve the implemented model for signal formation in order to have a better agreement with data.

Finally, the cluster position resolution is optimistic in simulation if compared to data, as shown in figure 139. As commented in section 9.5, this effect is not a surprise given the disagreement in

cluster size and the partial inaccuracy of the charge sharing model. Preliminary simulation studies extending the time structure of the strip signals as a way of obtaining a more accurate cluster position estimate do not seem to significantly deteriorate the simulated resolution to match the measured one. However, these studies are not based on first principles, but on a phenomenological parametrization of what is observed in calibrations, and may therefore still miss the physics that would have an impact on the resolution.

8.1.4 Belle II conditions database

The SVD detector status and calibrations are stored in the Belle II Conditions Database [61]. Conditions are stored in serialized ROOT objects called payloads. Intervals of Validity (IoVs) specify the time interval for which a payload is valid. Time intervals are defined with run-level granularity. A given payload may have multiple IoVs assigned to it if it is valid for different periods of time. A Global Tag is a list of IoVs and associated payloads, and it is used to select a complete set of conditions for a given processing or reprocessing effort.

The Online Global Tag contains all the conditions that are used during data taking by the High Level Trigger as well as for the online data-quality monitoring. The SVD detector status and the results of the local calibrations, described in section 5.2.1, are promptly uploaded to the Online Global Tag. This allows an optimal data quality for the following runs that ensures an unbiased High Level Trigger selection.

Other calibrations cannot be performed directly during data taking and are performed on triggered events. This is for example the case of the hit time calibration described in section 8.2.4. The results of these calibrations are uploaded to the Global Tags that are used for the regular reprocessing of the data. These reprocessing allow the best possible data quality for physics analyses.

8.2 Reconstruction software

The goal of the SVD reconstruction is to provide the 3D position, time, and ionization charge of the charged particles crossing the SVD sensors. This information is used to reconstruct the tracks and to provide information on the particle identification. The reconstruction of charged particles with low transverse momentum completely relies on SVD reconstruction, as the CDC-only tracking efficiency is $\sim 80\%$ for $p_T \simeq 100 \text{ MeV}/c$.

The SVD reconstruction was developed to cope with an expected background hit rate of $1.5 \text{ MHz}/\text{cm}^2$ at the SuperKEKB design luminosity, $8 \times 10^{35} \text{ cm}^{-2}\text{s}^{-1}$, which corresponds to an occupancy of $\simeq 1.5\%$ in the innermost layer, layer 3. In order to reconstruct the signal tracks (the average number of tracks in an $\Upsilon(4S)$ event being 10) most of the hits from particles belonging to a different collision event, mostly due to machine background particles (off-time hits), should be identified and removed before the tracking stage. Therefore, beside the ability to correctly reconstruct the position of the hit on the sensor, the hit time must be estimated with a precision that allows to distinguish on-time from off-time hits. As of 2020, given the low machine backgrounds, all reconstructed hits are provided to the tracking, including the off-time hits. In the future, if the machine background further deteriorates, the distinction between on- and off-time hits will become crucial for the tracking performance of the experiment.

In the following, the algorithms that allow to reconstruct the clusters and space-points are described, and the reconstruction performance on simulated events are briefly shown, while in




Figure 119. u/P side (left) and v/N side (right) digitized waveforms of the APV25 in data, shown for 4 random strips, in hadronic events collected in 2020. The noise ranges between 2.5 and 4 ADC on u/P and v/N side.

section 9 the reconstruction performance on data is discussed. Performance results refer to 6-sample acquired events both for data and simulation. From the point of view of reconstruction and hit time calibration, the simulated events are treated exactly in the same way as the data.

8.2.1 Unpacking and strip reconstruction

As shown schematically in figure 112, the raw data received from Belle II DAQ are unpacked into

- an array of digits (`SVShaperDigits`), one per each strip passing the online zero suppression, and consisting of the 6 digitized values for the sampled waveform;
- a single object that contains event-wise information `SVEventInfo`: the data acquisition mode (6- or 3-sample acquisition), the trigger type, and the `trigger bin`: the arrival time of the L1 trigger relative to the APV sampling clock. The trigger bin allows to move the hit time measured in the SVD reference frame to the trigger reference frame, common to all subdetectors, and vice-versa.

The unpacker module checks the data format, and the consistency of event-wise information across the FADC, FTB and APV. It also checks for errors from FADC, FTB or APV reported in the raw data. Detected errors are written in a dedicated object, which is not stored for reconstruction but is used in the online Data Quality Monitor. If an error is detected, the unpacker does not produce any strip information in output, hence that event does not have SVD data available for reconstruction. In figure 119 examples of the digitized waveforms of strips for both sides are shown.

The `SVShaperDigits` are reconstructed by a dedicated module that computes the charge and the hit time of each strip, then stored in a dedicated array named `SVRecoDigits`. The **charge** of the strip corresponds to the highest of the six digitized samples, converted in electrons using the strip-gain as determined during the local calibration (section 5.2.1), and read from the Conditions Database. The raw **hit time** (t_{raw}) is determined as a weighted average of the sampling time (t_i) with the ADC count of the i -th sample A_i corrected by t_{peak} to remove differences in peaking times among the strips:

$$t_{\text{raw}} = \sum_{j=0}^{j=5} \frac{t_j \cdot A_j}{A_{\text{tot}}} - t_{\text{peak}}, \quad \text{with} \quad A_{\text{tot}} = \sum_{j=0}^{j=5} A_j \quad (8.1)$$

where $t_j = j \cdot 1/f_{\text{APV}}$ and $f_{\text{APV}} = 31.805$ MHz, and t_{peak} is determined for each strip from the local calibration and stored in the Conditions Database, see section 5.2.1.

As the last step of the strip reconstruction, the raw strip time is calibrated:

$$t_{\text{strip}} = f(t_{\text{raw}}) + \delta t \quad (8.2)$$

where $f(t_{\text{raw}})$ is the calibration function and δt is an event-wise time shift that allows to move the hit time from the SVD reference frame to the trigger time reference frame, as explained in section 8.2.4. The calibration function is stored in the Conditions Database, together with the value of its parameters, that depend on each sensor side.

8.2.2 Cluster reconstruction and SpacePoint creation

The cluster is reconstructed from the strips passing the online zero suppression, which rejects the signal from a strip whose maximum signal height S_i is less than three times the noise N_i of the strip. Please note the difference between the strip $\text{SNR}_{\text{strip}}$, defined as the ratio of the strip signal to its noise, and the cluster SNR, defined as ratio of the total charge of a cluster to its total noise (discussed below).

A bunch of strips, which comprises only adjacent strips with $\text{SNR}_{\text{strip}} > \text{SNR}_{\text{adj}}$, is regarded as a cluster if at least one of the strips (seed strip) in the bunch has $\text{SNR}_{\text{strip}} > \text{SNR}_{\text{seed}}$ and the bunch has $\text{SNR} > \text{SNR}^{\text{min}}$. In the current reconstruction $\text{SNR}_{\text{adj}} = 3$ (the same as the online zero-suppression criterion, therefore all strips available at the reconstruction level are also available for clustering), $\text{SNR}_{\text{seed}} = 5$, $\text{SNR}^{\text{min}} = 0$. Given the high SNR (see section 9.3), this online zero suppression cut at 3 is small enough for the resolution purpose, as verified also from testbeam data, where a good resolution was achieved even with a higher (5) online zero suppression cut. For Belle II physics runs a zero suppression cut of 3 was adopted, not to overload the data bandwidth with strips firing on noise while retaining good resolution. The reconstruction performance are evaluated on simulated events, including beam background expected at the design luminosity.

The reconstruction efficiency of clusters, defined as the fraction of `SVDTrueHits` for which a cluster was reconstructed, is close to 1 and it is shown in the left plot of figure 120 for different sensor types and layers. The cluster purity reported in the right plot of the same figure is defined as the fraction of signal strips in the clusters associated to `SVDTrueHits`. Barrel sensors show a higher purity because of the angular distribution of beam background which is smaller in the barrel region compared to the forward and backward regions. Among the barrel sensors, layer 3 has a lower purity because it is located closest to the IP where beam background is higher as it decreases rapidly going outward.

The cluster charge is computed as the sum of the charges of each strip belonging to the cluster:

$$S_{\text{CL}} = \sum_{i=0}^{i<\text{size}} S_i \quad (8.3)$$

while the noise of the cluster is computed as the quadrature sum of the noise of the single strips:

$$N_{\text{CL}} = \sqrt{\sum_{i=0}^{i<\text{size}} N_i^2} \quad (8.4)$$




Figure 120. Clustering efficiency (left) and cluster purity (right) for the u/P (red) and v/N side (blue) for different sensor types, estimated on simulated events, including beam background expected at the design luminosity.

and the signal-to-noise ratio of the cluster is defined as the ratio between its charge and noise, $SNR = S_{CL}/N_{CL}$. Distributions of cluster size, charge, and SNR are shown in figures 121, 122, and 123, respectively. The cluster size is correlated to the incident angle of the track on the sensor. Layer 3 forward sensors show a higher cluster size because of the larger incident angle and the smaller pitch on the v/N sides. The outer layers have a slanted sensor in the forward direction, therefore the cluster sizes on the two sides are similar and smaller with respect to layer 3. Also, in the cluster charge there are visible variations of the energy released, due to the different track incident angle and track length in the sensitive volume, as detailed in section 9.1. The signal peaks at around 24k electrons, as expected for a MIP crossing 320 μm silicon for sensors 3.x.2 and 6.x.4 (see section 2 for a definition of the nomenclature), with tracks mostly perpendicular. Larger MPV are visible in sensors in the forward region. The SNR distributions are more similar among layers, since the effect of the larger charge related to larger incident angle is partly removed considering the cluster noise, that includes the cluster size in its definition. Higher SNR values are visible for the v/N side, where noise is significantly lower compared to the u/P side, except for layer 3 sensors.

Cluster position. The cluster position (x_{CL}) and its uncertainty (Δx_{CL}) are computed using the charge (S_i) and position (x_i) of its strips, with different algorithms depending on the cluster size.

For clusters formed by a single strip, the position corresponds to the center of the implant of the readout strip, while the position uncertainty takes into account the fact that the adjacent strips may have a signal below the threshold:

$$x_{CL1} = x_1 \quad (8.5)$$

$$\Delta x_{CL1} = p \frac{Q_{ph}}{S_1 + Q_{ph}} \quad (8.6)$$

where p is the strip pitch and Q_{ph} is the phantom charge, i.e. the maximum charge that the strip should have not to be included in a cluster $Q_{ph} = SNR_{adj} N_i$.




Figure 121. Reconstructed cluster size for simulated events including beam background expected at design luminosity, for different layers and sensors: layer 3 forward (upper-right) and backward (upper-left) and two barrel layer 6 sensors (bottom-right and left).




Figure 122. Distributions of reconstructed charge for simulated events including beam background expected at design luminosity, for different layers and sensors: layer 3 backward (upper-left) and forward (upper-right) and two barrel layer 6 sensors (bottom-left and right)




Figure 123. Distributions of reconstructed cluster SNR for simulated events including beam background expected at design luminosity, for different layers and sensors: layer 3 backward (upper-left) and forward (upper-right) and two barrel layer 6 sensors (bottom-left and right).

For clusters formed by two strips the position is computed with the center-of-gravity algorithm:

$$x_{\text{CL2}} = \frac{x_1 S_1 + x_2 S_2}{S_1 + S_2} \quad (8.7)$$

$$\Delta x_{\text{CL2}} = p \frac{\text{SNR}_{\text{adj}}}{S_{\text{cl}}/N_1} \quad (8.8)$$

where S_{cl} is the charge of the cluster and N_1 is the noise of the first strip.

Position of clusters with more than two strips are computed with the head-tail algorithm:

$$x_{\text{CL3}} = \frac{1}{2} \left[x_{\text{head}} + x_{\text{tail}} + p \frac{S_{\text{head}} - S_{\text{tail}}}{S_{\text{center}}} \right] \quad (8.9)$$

where the subscripts head and tail indicate the position x (or charge S) of the two strips at the edge of the cluster, and $S_{\text{center}} = (S_{\text{CL}} - S_{\text{tail}} - S_{\text{head}})/(size - 2)$, with S_{CL} the total charge of the cluster. The position uncertainty is computed as:

$$\Delta x_{\text{CL3}} = \frac{p}{2} \sqrt{\left(\frac{1}{sn}\right)^2 + \frac{1}{2} \left(\frac{S_{\text{tail}}}{S_{\text{center}}}\right)^2 + \frac{1}{2} \left(\frac{S_{\text{head}}}{S_{\text{center}}}\right)^2} \quad (8.10)$$

with $sn = \frac{S_{\text{center}}}{\text{SNR}_{\text{adj}} N_{\text{CL}}}$.

Simulation studies have shown that using the calculated uncertainty the cluster position residual pulls do not have unit width as expected. To take this discrepancy into account, the cluster position

uncertainty is corrected by a scale factor that depends on the sensor type, side, and the number of strips in the cluster (1, 2, > 2), listed in table 23. These scale factors s_{err} are computed using simulation by requiring that 68% of the pull is found within 1σ with respect to the median of the distribution. The corrected uncertainty provided to tracking $\overline{\Delta x_{\text{CL}}}$ is therefore:

$$\overline{\Delta x_{\text{CL}}} = s_{\text{err}} \Delta x_{\text{CL}}.$$

Table 23. Cluster position uncertainty scale factors used in the reconstruction, estimated on simulated events.

sensor	side	s_{err} for size		
		1	2	> 2
HPK small	u/P	1.352	1.137	0.559
HPK small	v/N	1.638	1.168	0.430
HPK large	u/P	1.312	0.871	0.538
HPK large	v/N	2.338	1.418	0.468
wedge	u/P	1.728	1.209	0.662
wedge	v/N	1.766	1.481	0.433

The cluster position resolution was estimated using simulated di-muon events, based on the same analysis as for data explained in section 9.5, and it is reported as a function of the track incident angle in figure 124. The observed resolution has the expected shape as a function of the angle, showing a minimum at the incident angle for which the projection of the track along the direction perpendicular to the strips on the detector plane corresponds to two strip pitches. Given the various sensor pitches with one floating strip (table 1), the minimum is expected at 4 (7) degrees for the u/P side and at 14 (21) degrees on the v/N side, for layer 3 (4, 5, 6). The expected digital resolution for perpendicular track is $7\ \mu\text{m}$ ($11\ \mu\text{m}$) on u/P side and $23\ \mu\text{m}$ ($35\ \mu\text{m}$) on v/N side for layer 3 (4,5,6), similar to the measured resolutions. The method tends to overestimate the resolution by up to $2\ \mu\text{m}$, depending on the incident angle, as observed during the validation using the true information available for simulated events.




Figure 124. Cluster position resolution vs. track incident angle estimated on simulated events, which include the beam background expected at design luminosity, with the same analysis used for data. Results of the u/P (v/N) side is shown for the layer 3 in red and for the outer layers in black.

Since masked strips are not considered when forming a cluster, they can cause a cluster to be split in two by a masked strip in the middle. Given that the number of masked strips is stably below 1%, it is not critical to take into them account in clustering. In case an entire APV chip is masked, a *fake* cluster positioned in the middle of the region read by the disabled APV is created in order to use the information provided by the clusters on the opposite side.

Studies of the cluster position resolution on data shown in figure 139 of section 9 indicate that the resolution observed in simulated events is optimistic, and that there is still room to improve the resolution on data. Studies are ongoing to improve the resolution on data by testing alternative algorithms, as well as to improve the data-MC agreement, for example for the cluster size distribution as mentioned before.

Cluster time. The cluster time is computed as an average of the calibrated time of the strips forming the cluster, weighted with the charge of the strip. The cluster time resolution is estimated from the distribution of the difference between the cluster time and the event t_0 where event t_0 is the time of the event reconstructed by the CDC, fitting this distribution with a sum of two independent Gaussians of means μ_i and widths σ_i . The resolution is estimated as the weighted average of the widths of the Gaussians corrected by the nonzero values of the means of these Gaussians:

$$\sqrt{f\sigma_1^2 + (1-f)\sigma_2^2 + f(1-f)(\mu_1 - \mu_2)^2} \quad (8.11)$$

where f is the fraction of the first Gaussian, μ_i and σ_i are the mean and width of the i -th Gaussian. The cluster time resolution on simulated events is of the order of 2 ns, as shown in figure 125.




Figure 125. Cluster time — event t_0 for clusters related to tracks for simulated events: layer 3 for the u/P side on the left and v/N side on the right. The event t_0 is the time of the event reconstructed by the CDC. The resolution on the cluster time, estimated with equation (8.11) is 2.2 ns on the u/P side and 1.8 ns on the v/N side.

Although currently the hit time is reconstructed using all the six APV samples, alternative algorithms have been developed to compute the cluster time using only a subset of the 6 samples. The long right tail of the APV waveform, shown in figure 85, does not bring very useful information on the hit time with respect to the three samples around the peak. Therefore, first the amplitudes of each of the strips forming the cluster are summed sample-by-sample, and then the three samples corresponding to the waveform peak of each strip forming the cluster are selected, storing the time of the first selected sample (FF). The raw cluster time is then computed as a weighted average of




Figure 126. Cluster time - event t_0 for clusters related to tracks in simulated events, with the alternative algorithm using three APV samples: layer 3 for the u/P side on the left and v/N side on the right. event t_0 is the time of the event reconstructed by the CDC. The resolution on the cluster time, estimated with equation (8.11) is 2.0 ns on the u/P side and 1.3 ns on the v/N side.

the sample time with the sample amplitude, similarly to what is done to compute the strip time. The calibration is performed in the same way as the default one, exploiting the correlation with the event t_0 from the CDC as discussed in section 8.2.4, equation (8.15). However, in this case an additional shift must be considered, FF , the time of the first sample:

$$t_{cl3sample} = a + b \cdot t_{raw} + c \cdot t_{raw}^2 + d \cdot t_{raw}^3 + \delta t + FF. \quad (8.12)$$

The calibration parameters are different than the ones determined from the calibration of the default algorithm, and are therefore stored in a dedicated payload. The performance of this alternative time estimation are better by 10-30% with respect to the default algorithm, as shown in figure 126.

SpacePoint creation. The position of the 3D hits, stored in the `SVDSpacePoint`, is provided by combining all clusters on one side of the sensor with those on the opposite side. At this stage SpacePoints formed by clusters whose time (t_{CL}) or time-difference ($|t_{CL}(u) - t_{CL}(v)|$) does not satisfy a certain selection can be rejected. The software allows to choose one of the following selection criteria:

- $t_{CL} < t_{min}$: early clusters are not taken into account for the combination, where t_{min} is the minimum time of the allowed cluster;
- $|t_{CL}| > \Delta t$: in case of small jitter and correct latency, the time of the cluster is expected to be around $t_{CL} = 0$ ns, clusters with time that differ from this expectation more than Δt are rejected;
- $|t_{CL}(u) - t_{CL}(v)| > \Delta t'$: clusters of opposite sides created by the same particle are expected to have similar times, if their time differ more than $\Delta t'$ the combination is rejected.

The choice of the selection does not require to recompile the software, it is stored in the Conditions Database and can be changed by changing the payload read by the module. At the moment, none of these selection criteria are applied for data reconstruction, and all possible combinations of opposite-side clusters are provided to the tracking. This is possible because the current beam background is sufficiently low for the tracking to handle all the SpacePoints. In the future, when




Figure 127. Distribution of the difference between the u/P and the v/N cluster time belonging to the same SpacePoint, of the layer 3 (left) and the barrel sensors of the outer layers (right) for simulated events. The distributions were fit with a sum of two Gaussians (the first one with fraction f_1) with mean μ_i and width σ_i (both in ns). The widths of the distributions, estimated with equation (8.11) is 2.0 ns for layer 3 and 1.3 ns for the outer layers.

luminosity increases, it will become critical to reject SpacePoints based on the cluster time in order to have an efficient tracking with a low fake rate. As an example, for simulated events with superimposed beam backgrounds expected at design luminosity, a selection on the time difference of the clusters of a SpacePoint is applied: if the time difference is larger than 10 ns, the combination is rejected. In figure 127 the cluster time difference distribution for signal cluster from simulated events is shown. This selection reduces the number of SpacePoints by 60%, yielding a 30% improvement in the execution time of pattern recognition algorithms that involve SVD, and a reduction of the track fake rate for the SVD-only pattern recognition by 25%, down to 7.5%, with an unchanged hit finding efficiency.

8.2.3 Data quality monitoring

During data taking the quality of the data is constantly monitored by the control-room shifters, and available for detector experts. A dedicated computing server, ExpressReco, receives a fraction of the events from the EventBuilder and performs the complete official reconstruction online. Downstream reconstruction dedicated DQM modules that fill histograms and run simple analysis code to determine if the histogram has the expected shape are run. The color of the histogram canvas turns to red if the analysis code detects an unexpected distribution, orange if the distribution may be problematic, grey if there is not enough statistics, and green in the other cases. Reference histograms are superimposed to the live histograms to help in judging data quality. A list of selected of histograms is available to the control-room shifters, all other histograms are available to detector experts to debug problems and assess the quality of the run. A few DQM modules run downstream reconstruction also on the High Level Trigger (HLT) machines, in this case the check is performed on all events.

The monitoring of SVD data starts from the data received from DAQ: the data format of all events is checked by a DQM module running on HLT. The corresponding plot shown on figure 128 provides information on the fraction of events with data format issues, the involved FADC board, and the error type. On ExpressReco the occupancy due to beam backgrounds is checked by computing the average fraction of strips with $\text{SNR}_{\text{strip}} > 5$ per sensor on the u/P side. The average sensor




Figure 128. DQM plots continuously monitored by the control-room shifter: data format (top left), average u/P side occupancy computed with strips with $\text{SNR}_{\text{strip}} > 5$ (top right), average v/N side sensor efficiency (bottom left), signal-cluster time of layer 4,5,6 (bottom right). Examples taken from a run in 2020 in which no problem was observed, indeed all plots show a green canvas.

efficiency for the v/N side is monitored, as well as the v/N side cluster time distribution of layers 4, 5 and 6 cluster used to reconstruct a track in order to confirm that the acquisition window is at the correct time position. In summary, the four plots available to the control-room shifter shown in figure 128 are:

- data format;
- average u/P side occupancy computed with strips with $\text{SNR}_{\text{strip}} > 5$;
- average v/N side sensor efficiency;
- v/N side cluster time of clusters used to reconstruct a track on layers 4, 5, 6.

Many more detailed plots produced on the ExpressReco machines are made available in real time to the SVD experts.

8.2.4 SVD hit time calibration

The hit time calibration procedure uses clusters in order to minimize the effect of the electronic noise. The cluster time is calculated as an average of the strip time weighted with the strip charge, summed over the strips in the cluster. As discussed in section 8.1.2, t_{raw} is the raw cluster time measured in the SVD time reference frame, while t_0 is the time of the event estimated by the CDC measured in the trigger reference frame. The calibration, exploiting the correlation between the cluster time and the time of the event, was optimized on data. In the following, the method is explained showing results on simulated events; results on data are reported in section 9.6.

First, the time of the event expressed in the SVD reference frame as t_0^{svdRF} is computed shifting the event t_0 by a known event-dependent quantity δt , based on the trigger bin, and that gives an




Figure 129. An example of the $(t_0^{\text{svdRF}}, t_{\text{raw}})$ correlation plot between the time of the event in the SVD reference frame and the raw cluster time used for the SVD hit time calibration and superimposed calibration function (red line) for v/N size of a barrel sensor of layer 6.

additional information about the arrival of the trigger to the SVD system:

$$t_0^{\text{svdRF}} = t_0 - \delta t. \quad (8.13)$$

The goal of the calibration is to produce a calibrated hit time, corresponding to the t_0^{svdRF} , starting from the $(t_0^{\text{svdRF}}, t_{\text{raw}})$ correlation plot. In this calibration procedure the effect of flight time of the particles is neglected, and the calibration function $f(t_{\text{raw}})$ is determined empirically using a third-order polynomial. The calibration function parameters are extracted by a fit on data of the x-profile of the two-dimensional $(t_0^{\text{svdRF}}, t_{\text{raw}})$ correlation histogram:

$$f(t_{\text{raw}}) = a + b \cdot t_{\text{raw}} + c \cdot t_{\text{raw}}^2 + d \cdot t_{\text{raw}}^3 \quad (8.14)$$

In figure 129 an example of the $(t_0^{\text{svdRF}}, t_{\text{raw}})$ calibration 2D plot, with the fitted function superimposed is shown. The calibration parameters (a, b, c, d) are separately extracted for each sensor side and are stored in the Conditions Database. During reconstruction the strip time is calibrated using equation (8.14) and finally shifted in the trigger reference frame, which is common to all detectors. The calibrated strip time is then:

$$t_{\text{strip}} = a + b \cdot t_{\text{raw}} + c \cdot t_{\text{raw}}^2 + d \cdot t_{\text{raw}}^3 + \delta t. \quad (8.15)$$

The calibration for the v/N and u/P sides of each sensor is studied separately, since the waveforms provided by the two sides are different due to a number of effects:

- different capacitive load that can change the waveform shape given by the APV25;




Figure 130. Cluster Time distributions of the u/P (left) and v/N side (right) obtained on simulated events. Clusters associated to tracks are represented by the shaded region.

- the APV25 chips operate in a different configuration for P and N sides (inverter on/off), to be able to read signal of opposite charge collected on the two sides of the sensors.

Both these effects contribute to the difference observed in the shape of the APV25 signals obtained during calibration runs, injecting a known pulse in the internal APV25 calibration circuit. In particular, signals on the u/P side have a slightly longer peaking time, by about 10 ns, and larger width of about 20 ns.

In addition to these effects, already visible in calibration, the rise time of the real signal coming from the silicon sensor is also affected by the different mobility of holes and electrons. Holes, collected on the u/P side, are ~ 3 times slower than electrons, collected on the v/N side. So the final signal of the APV25 for the u/P side is slower than that on the v/N side. Figure 119 shows the differences between the waveforms for the two SVD sensor sides.

As a consequence of these effects a better resolution on the time estimated on the v/N side (1.8 ns) with respect to the u/P side (2.2 ns) is achieved, as shown previously in figure 125. In figure 130 the calibrated cluster time distribution from simulated events with beam background expected at design luminosity is shown, highlighting the distribution of the signal clusters, i.e. those used by the tracking. The corresponding distributions for data are shown in figure 141 and have a similar shape. The calibrated time of clusters associated to tracks (no SVD hit time information is used in reconstruction yet) nicely peaks at zero, around the event t_0 , while background clusters are more spread in the window. The peak on the left of the window, between -100 and -50 ns, is due to background off-time tracks hitting the sensors before the beginning of the acquisition window. The shaped signal of the strips has a long tail of few hundreds of ns and the signal is still above threshold when this tail enters the acquisition window, causing an accumulation of earlier background hits at the beginning of the window. The rest of the background hits, arriving inside the SVD acquisition window, are more uniformly distributed in time, as expected. The time separation between signal and background hits will be exploited in the future for background mitigation.

9 Performance

During the first 1.5 years of data taking (2019-2020) with the complete SVD installed, Belle II has collected approximately 70 fb^{-1} integrated luminosity at the $\Upsilon(4S)$ resonance. These data are used

to characterize the performance of the SVD, and confirm the excellent behaviour of all the 172 SVD sensors, with stable performance in time. The performance of the SVD is studied in all the aspects outlined in the description of software and reconstruction in section 8. Unless otherwise indicated, candidate hadronic events are used, defined by the requirement of at least four reconstructed tracks from the IP.

As a starting point, in section 9.1 clusters of adjacent strips associated to tracks are considered. Their total collected charge, normalized to the track length across the DSSD sensor, is compared with the expectations for the two sensor sides.

The strip noise (section 9.2) is dominated by the capacitive input load to the front-end APV25 ASICs; it gives a very good signal to noise ratio, ranging between 13 and 30 (section 9.3), confirming the success of the novel origami scheme in the SVD design (section 2).

Section 9.4 describes the evaluation of hit efficiency using reconstructed tracks. A few sensors have localized defects; in general the efficiency per sensor is above 99%, being continuously monitored by DQM software.

Position and time resolution (sections 9.5, 9.6) are the essential ingredients in track reconstruction as well as in the rejection of background hits. The measured cluster position resolution, about 15 μm (20–35 μm) in u/P (v/N), is adequate and close to the expectation. Further improvements are expected optimizing the position reconstruction algorithms directly on data. An accurate calibration of the hit time determination is complemented by the analysis of its stability.

The alignment algorithms (section 9.7) pin down the position of sensing elements to an accuracy of about 10 μm . They describe the position by six parameters for each sensor as a rigid body, plus seven parameters to describe deformations of its surface.

The main features of the performance of track finding and reconstruction are briefly outlined in section 9.8. Finally, section 9.9 describes the contribution of SVD to charged particle identification by the measurement of specific ionization (dE/dx), in particular for the low-momentum particles with limited or missing information from the CDC.

9.1 Cluster charge

After the full reconstruction is performed, signal clusters, defined in section 8.2.2, are used to measure their collected charge on data. As detailed in section 8.2.1 signals from strips and clusters are converted to energy/charge, expressed in terms of electrons, using the gain measured with the internal APV25 calibration circuit, section 5.1.1.

The cluster charge released in the SVD strongly depends on the track incident angle θ with respect to the sensors, as shown in figure 131, left. As defined in section 2 and shown in figure 3, sensors in different z positions along the ladders, are indicated with increasing index number starting with 1 for the forward sensor. According to their position, sensors see tracks coming from the IP with very different incident angles, as shown in figure 131, right, and the distribution of the collected cluster charge reflects this effect.

As an example, in figure 132 the cluster charge for sensors in specific z positions in layer 3 and 6, summed over all ϕ ladders, are shown. Sensors are numbered as defined in section 2 indicating the position in the layer, ladder (ϕ), z position along the ladder with the following convention: L3.2.1 stands for the position in layer 3, ladder 2, sensor 1 (FW) along the ladder. L3.x.1 (FW) and L3.x.2 (BW) are shown together with two other layer 6 sensors in the central z position of the ladders,




Figure 131. Left: dependence of the energy/charge released in the silicon sensor (E measured in e^-) from the track length ℓ in the sensitive volume, with $d =$ detector thickness. Right: SVD sensors layout, with sensors in the different position along the ladders (forward, central, backward) indicated with increasing index number. Reproduced from [44]. CC BY 4.0.




Figure 132. Cluster charge for sensors in specific z positions in layer 3 and 6, summed over all ϕ ladders; u/P side in red and v/N side in blue. From top to bottom and left to right: L3.x.2, L3.x.1, L6.x.4, L6.x.2. Data for a typical run with colliding beams in 2019 are shown. The vertical dashed lines correspond to the MIP signal at normal incidence.

L6.x.2 and L6.x.4. As expected, from their position with respect to the IP, and the track incident angle on them, the charge is significantly larger for sensors toward the FW position.

To remove the effects due to the track incidence angle, and verify the correct calibration of the system, the cluster charge rescaled to detector thickness, considering the track length ℓ in the sensitive volume, is evaluated. This distribution is shown for the previous four sensors in figure 133 and, as expected, it is independent of the sensor location. The most probable value (MPV) for the normalized cluster charge is on average 21000 e^- on the u/P side. Taking into account the large uncertainty ($\sim 15\%$) in the absolute APV25 gain calibration (see section 5.2.1), this measured value is in fair agreement with the expected MIP signal in a 320 μm thick silicon sensor ($\sim 24000 e^-$).

In these distributions the reconstructed cluster charge on the two sides of the sensors are similar, but a signal loss of about 10–30% is visible on the v/N side with respect to the u/P side. This effect




Figure 133. Cluster charge rescaled to detector thickness, considering the track length ℓ in the sensitive volume, for sensors in specific z positions in layer 3 and 6, summed over all ϕ ladders; u/P side in red and v/N side in blue. From top to bottom and left to right: L3.x.2, L3.x.1, L6.x.4, L6.x.2. Data for a typical run with colliding beams in 2019 are shown. The vertical dashed lines correspond to the expected most probable value of the MIP signal at normal incidence.

was expected and partly reproduced in MC simulation, as explained in section 8.1.3. It is due to the large pitch on v/N side combined with the presence of a floating strip, giving a coupling capacitance of the strips to the back plane no longer negligible with respect to the interstrip capacitance, thus reducing the signals coupled to the readout strips when tracks release charge close to the floating strip. In fact, for a track crossing the sensor close to a floating strip, as shown in detail in section 8.1.3, the signal seen on the adjacent readout strips via capacitive coupling is about $\frac{2C_i}{2C_i+C_b}$ of the total signal released, about 75% for our sensors.

9.2 Strip noise

The strip noise is measured in terms of ADC counts in the “noise local run” (see section 5.2.1), and then converted to Equivalent Noise Charge (ENC) in electrons, using the gain measured with the internal APV25 calibration circuit. The estimated accuracy on the absolute gain calibration, performed with testbeam data, is about 15%. Table 24 shows the ENC measured for the different sensor types and sides at the beginning of data taking in 2019, before the increase due to radiation damage discussed in section 7.4.1.

The strip noise is dominated by the capacitive input load to the APV25, thus the noise figure for the different sensor sides are different by almost a factor two among the longer u/P side strips, with smaller pitch and larger interstrip capacitance, and the shorter v/N side strips with larger

Table 24. Average Noise (ENC) measured for each sensor side and position in the ladders at the beginning of data taking in 2019.

Sensor position/type	u/P side ENC (e^-)	v/N side ENC (e^-)
Layer 3 (HPK small)	930	630
Layer 4/5/6 Origami (HPK large)	958	510
Layer 4/5/6 BWD (HPK large)	790	680
Layer 4/5/6 FWD (Micron wedge)	740	640

pitch. Additional small differences in noise are also visible HPK large rectangular sensors mounted in different ladder position, since in some cases the contribution from the capacitance of the pitch-adaptor circuit, used to connect the sensor strips to the chip, can be significant.

Thanks to the chip-on-sensor concept exploited in the ladder design, the total channel noise is at most about $1000 e^-$, as expected from design, allowing to achieve very good signal to noise ratio in each SVD sensor. After a year and half of beam operation, an increase of about 25(15)% in noise was observed on the u/P (v/N) side in the most exposed sensors in the layer 3 middle plane, after an estimated total dose of about 100 krad in this location. The noise increase has already shown some saturation, as expected for effects related to the increased interstrip capacitance due to the higher fixed oxide charge, as shown in section 7.3.1. In the external layers, exposed to lower dose, the increase in noise observed was of 5 to 10%.

9.3 Signal to noise ratio

The cluster SNR, defined in section 8.2.2 as the total cluster charge over the sum in quadrature of the noises of the strips belonging to the clusters, is an important figure of merit to be monitored. Due to the online zero suppression criteria, based on signal strip higher than 3 times its noise, low SNR values increase the probability to miss signal strips, deteriorating the hit efficiency and position resolution. The cluster SNR depends on collected charge, strip noise, and cluster size. As shown in previous sections, signals depend strongly on sensor position, due to the track incidence angles, that also influences the cluster size. Noise is higher in u/P side, with longer strips with respect to v/N side. This last effect dominates, then SNR is in general higher on v/N side, with quite large variation among the sensors, depending on their position.

The distribution of the SNR measured for the FW and BW Layer 3 sensors and for some Layer 6 sensors in different positions is shown in figure 134. The map of the SNR most probable values for all the SVD sensors is shown in figure 135 and clearly indicates a pattern determined by their position and the sensor side considered. All the 172 sensors have very good cluster SNR performance with MPV ranging between 13 and 30 as of July 2020. A small SNR reduction in Layer 3 u/P side of about 10% was observed, with respect to the beginning of the run in 2019, which is related to the increased noise due to radiation damage, with no effect on performance.

9.4 Hit efficiency

The efficiency of cluster reconstruction, described in section 8.2.2, was measured on an early data set with an *unbiased* technique and continuously monitored with [Data Quality Monitoring \(DQM\)](#)




Figure 134. Signal to noise ratio for sensors in specific z positions in layer 3 and 6, summed over all ϕ ladders; u/P side in red and v/N side in blue. From top to bottom and left to right: L3.x.2, L3.x.1, L6.x.4, L6.x.2. Data for a typical run with colliding beams in 2019 are shown.




Figure 135. Most probable value of the cluster SNR for all SVD sensors organized in a map. The horizontal axis identifies the ladder ϕ position (from 1 to 16 depending on the layer), while the vertical axis enumerates the different layers and the z position of the sensor in the ladder (from 1 to 5 depending on the layer). Data for a typical run with colliding beams in 2019 are shown.

tools with a *biased* technique. The unbiased technique consists in reconstructing tracks excluding the sensor under study and retaining only clusters with $SNR > 10$. The following selection criteria are applied to tracks:

- tracks must originate from the vicinity of the IP, $|d_0| < 0.1$ cm and $|z_0| < 0.2$ cm, where d_0 and z_0 are the transverse and longitudinal track impact parameters, respectively;
- tracks must have at least one SVD hit and 20 CDC hits;
- tracks must have a transverse momentum of at least 1 GeV/c.

Selected tracks are then extrapolated to the sensor under study. An example of the distribution of the track extrapolation uncertainty (i.e. the uncertainty in the fitted track extrapolation position on the sensor) for the barrel sensors of layer 4 is shown in the left plot of figure 136. The extrapolation to the u/P side shows a smaller uncertainty because of the better resolution on this direction both on SVD (smaller pitch) and CDC measurements. A fiducial area of ± 0.5 cm from the edge of the active area of the sensor was defined in order to exclude border effects; tracks extrapolated outside the fiducial area are discarded.




Figure 136. (left) Track extrapolation uncertainty (i.e. the uncertainty in the fitted track extrapolation position on the sensor) to u/P (blue) and v/N (red) sides for layer 4 barrel sensors; (right) residual distribution for layer 4 barrel sensors.

The histogram of the residuals (the distance between the extrapolated track and the cluster position) is built considering the cluster nearest the intercept. An example from layer 4 is shown in the right plot of figure 136. The analyzed data were reconstructed with a preliminary alignment and a residual misalignment is visible, especially on the v/N side. However, as shown in section 9.7 the current alignment is very much improved. The efficiency is defined as the fraction of intercepts with at least one cluster found in a window of ± 0.5 mm from the median of the residual distribution.

In figure 137 the map of the average efficiency of each sensor is shown. Only a few sensors have an efficiency below 99%. The efficiency on the v/N side is on average higher than the u/P side thanks to the higher SNR on this side.

A few sensors have localized defects slightly worsening the efficiency, for example the v/N side of sensor 4.3.3, 6.5.2, 6.6.2, and 6.7.1, and the u/P side of sensor 6.6.4. The efficiency for those sensors is still above 98%. Figure 137 also shows a slight variation of the efficiency depending on the position of the sensors. This is most probably due to the fact that the position of the sensor is correlated with the incident angle of the track on the sensor, and therefore the charge released in the sensor and the SNR change with position. Since the cluster reconstruction is independent of the sensor, the effect of clusterization may be slightly different depending on the position of the sensor.

The average efficiency per layer is above 99.3% for all layers, as listed in table 25.




Figure 137. Average efficiency for each sensor side, v/N on the left, and u/P on the right, organized in a map. The horizontal axis identifies the ladder ϕ position (from 1 to 16 depending on the layer), while the vertical axis enumerates the different layers and the z position of the sensor in the ladder (from 1 to 5 depending on the layer). The uncertainty in the efficiency varies from fraction of per-mill in the innermost sensors to fraction of per-cent in the outer sensors.

Table 25. Average sensor efficiency.

layer	$\varepsilon(u/P)(\%)$	$\varepsilon(v/N)(\%)$
3	99.83 ± 0.01	99.48 ± 0.03
4	99.69 ± 0.03	99.68 ± 0.03
5	99.66 ± 0.03	99.77 ± 0.04
6	99.31 ± 0.08	99.58 ± 0.06




Figure 138. Average efficiency, v/N in orange, and u/P in green. Data collected between March and June 2020 are shown, with the corresponding run number on the horizontal axis. Reproduced with permission from [47]. CC BY-NC-ND 4.0.

As described in section 8.2.3, the average sensor efficiency is monitored online. For this measurement a *biased* technique is used, including clusters on the sensor under study in the track reconstruction. The rest of the analysis is performed in the same way. Figure 138 shows the online average SVD efficiency in the last part of 2019 data-taking as well as in the first part of 2020 data-taking.

9.5 Position resolution

The cluster position resolution is measured with a di-muon sample ($e^+e^- \rightarrow \mu^+\mu^-$): events with only two tracks are considered, tracks with transverse momentum greater than 1 GeV/c, originating from the IP ($|z_0| < 0.5$ cm and $|d_0| < 0.4$ cm), and with at least one PXD hit, 8 SVD and 30 CDC hits are selected, and only muon pairs with an invariant mass between 10 and 11 GeV/c² are selected. After reconstruction, the muon tracks are refitted excluding the cluster on the sensor under study. Since the track extrapolation uncertainty is of the same order as the resolution that we are measuring, tracks with a track extrapolation uncertainty greater than 100 μm are excluded. A fiducial area of ± 0.5 cm from the sensor edge is defined. For each track i the unbiased residuals (Δx_i) and the unbiased track extrapolation uncertainty ($\sigma_{x,i}^{trk}$) are computed, to calculate the quantity r_i^2

$$r_i^2 = (\Delta x_i)^2 - (\sigma_{x,i}^{trk})^2 \quad (9.1)$$

The resolution σ_x is then computed as the square root of the average of r_i^2 :

$$\sigma_x = \sqrt{\frac{1}{N} \sum_i^N r_i^2} \quad (9.2)$$

with its uncertainty $\Delta\sigma_x$:

$$\Delta\sigma_x = \frac{1}{\sqrt{N}} \frac{\sigma_{r_i^2}}{2\sigma_x} \quad (9.3)$$

where $\sigma_{r_i^2}$ is the RMS of the r_i^2 distribution. Before computing the average of r_i^2 the distribution is restricted at $\pm 15\,000 \mu\text{m}^2$, removing the tails of the distribution. A value of $|r_i^2| > 15\,000 \mu\text{m}^2$, would correspond to a resolution greater than $\sim 120 \mu\text{m}$, which is much worse than the expected resolution of the SVD sensors. The analysis is performed as a function of the track incident angle, requiring at least 300 clusters for each bin. This method was validated with simulation, exploiting the true hit information, yielding a resolution within 2 μm compared to the true value.

The cluster position resolution as a function of the track incident angle is shown in figure 139 for data and simulated events. The observed resolution has the expected shape, showing a minimum at the incident angle for which the projection of the track along the direction perpendicular to the strips on the detector plane corresponds to two strip pitches. Given the various sensor pitches with one floating strip, the minimum is expected at 4 (7) degrees for the u/P side and at 14 (21) degrees on the v/N side, for layer 3 (4, 5, 6). As discussed in section 8.2.2, the expected digital resolution for perpendicular tracks is 7 μm (11 μm) on u/P side and 23 μm (35 μm) on v/N side for layer 3 (4,5,6), in very good agreement with the resolutions previously measured in testbeam data with perpendicular tracks [62], and also similar to the resolution shown for perpendicular tracks in figure 139. In particular, for the larger pitches (v/N sides) resolution measured in collision data is in very good agreement with the expected one, from previous testbeam results, while the MC resolution obtained is too optimistic. On the contrary for the small pitches (u/P sides) there is room for improvements: resolution in collision data is higher than previously measured in testbeam data, while MC results are more similar to the expected values.

The measured resolution in collision data is anyway good and adequate for the main goals of SVD tracking, though the algorithms for the determination of the cluster position, described in




Figure 139. Cluster position resolution as a function of the track incident angle for data and MC simulation. In the upper row u/P side results for layer 3 (left) and layer 456 barrel sensors (right). In the lower row v/N side results for layer 3 (left) and layer 456 barrel sensors (right).

section 8.2.2, have not been fully optimized on data yet, therefore there is still room for improvement of the cluster position resolution especially for the small pitches. The algorithms were in fact optimized on simulated events, and the resolution obtained with the same technique on simulation is much better than on data. Given the significant differences with respect to data, especially concerning the cluster size (see Simulation Performance in section 8.1.3), the difference in resolution between data and MC is not a surprise. Work both on improving the simulation of the signal formation, and on optimizing the position reconstruction algorithm directly on data is in progress. Due to the long editorial process required for this paper, some updated performance results could not be incorporated, and are included in the references in section 10.

9.6 Hit time determination

Although with present background levels the reconstruction does not require to use any background rejection based on time selection, a good hit time resolution will be crucial in future operation with higher luminosity and higher background occupancy to efficiently reject off-time background hits.

The calibration of the cluster time in data is performed with the same procedure used for simulated events (section 8.2.4), exploiting the correlation between the cluster time and t_0^{svdRF} the event t_0 expressed in the SVD reference frame. Figure 140 shows an example of the correlation plot used for the calibration of a single sensor: the calibration is performed fitting the correlation between the t_0^{svdRF} and the raw cluster time with a third-order polynomial. The parameters obtained from




Figure 140. Example of the $(t_0^{\text{svdRF}}, t_{\text{raw}})$ correlation plot between the time of the event in the SVD reference frame and the raw cluster time, used for the SVD hit time calibration. Data of one barrel sensor of layer 6 are shown, u/P side on the left and v/N side on the right, with the calibration function superimposed.

the fit are then used to calculate the calibrated time, which is finally shifted in the trigger reference frame according to equation (8.15).

The distribution of the calibrated cluster time is shown in figure 141. The time of the clusters used in tracking (performed without SVD hit time information) nicely peaks at zero, around the event t_0 , while background clusters are more spread in the window. The background peak on the left of the window, between -100 and -50 ns, is due to off-time clusters hitting the sensor before the beginning of the acquisition window. The shaped signal of the strips has a long tail of few hundreds of ns and the signal is still above threshold when the sampling starts, causing an accumulation of earlier background hits at the beginning of the window. The rest of the background hits, arriving inside the SVD acquisition window, are more uniformly distributed in time, as expected. The corresponding distributions for simulation are shown in figure 130 and have a similar shape. The time separation between signal and background hits can be exploited in the future for background mitigation.




Figure 141. Cluster time for all clusters and those related to tracks (shaded region) for the u/P side on the left and v/N side on the right.

The resolution of the cluster time can be estimated from the distribution of the difference between the cluster time, for the clusters related to tracks, and the event t_0 , the time of the event estimated with the CDC, as detailed in equation (8.11). Figure 142 reports the estimated hit time

resolution for layer 3 clusters used in tracking, that currently have the hit time reconstructed using the 6 APV samples. The hit time resolution for both sides is 4.3 ns, since the error on event t_0 is




Figure 142. Cluster time - event t_0 for clusters related to tracks on layer 3 for the u/P side on the left and v/N side on the right for data. The resolution on the cluster time, of 4.3 ns on both sides, is estimated as the weighted average of the widths of the two Gaussian functions used in the fit, corrected by the nonzero values of the means of these Gaussians (equation (8.11)).

negligible (≈ 0.7 ns). There is a small bias of order of 0.1 ns on the u/P side which has no effect on the reconstruction. The hit time resolution is also estimated using an alternative algorithm, with 3 APV samples out of 6. The results are reported in figure 143, showing a better resolution of 2.4 ns on the v/N side and 2.9 ns on the u/P side.




Figure 143. The difference (cluster time - event t_0) for clusters related to tracks in data, on layer 3 for the u/P side on the left and v/N side on the right, with the alternative algorithm that uses 3 samples out of the acquired 6. The resolution, 2.4 ns on the v/N side and 2.9 ns on the u/P side, is estimated from the weighted average of the widths of the two Gaussian functions used in the fit, corrected by the nonzero values of the means of these Gaussians (equation (8.11)).

The resolution of the cluster time can also be estimated independently of the event t_0 . Assuming that the time of all clusters connected to a track is the same, the resolution is estimated as the $\text{RMS}/\sqrt{2}$ of the time difference of two clusters on the same track. Figure 144 shows the estimated time resolution using SVD-only information for the 6 sample time estimator. They are compatible with the estimation presented in figure 142.




Figure 144. Cluster time resolution in ns, estimated as the $\text{RMS}/\sqrt{2}$ of the time difference, for couples of clusters on different layers belonging to the same track. Rows (columns) represent the outer (inner) layer number whom the cluster belong to. u/P side (left) and v/N side (right).

Calibration parameters in real data may differ for different data sets depending on changes in SVD (for example changes of APV latency), or in external conditions (for example changes of trigger latency). Consecutive runs are merged and calibrated as a single data set if the raw hit time average of layer 3 v/N side sensors differs by less than 2 ns or if the statistics of the run would be insufficient to perform a separate calibration.

Figure 145 shows the stability of the calibrated hit time versus the run number for a data taking period in 2020. The plot, automatically produced in the calibration procedure, shows the distribution over all sensors of the average of the difference between the calibrated cluster-related-to-tracks time and the event t_0 , divided by the RMS of the event t_0 for the two sides separately, for each run. The calibration is considered successful if the average of the distributions is within ± 0.5 for each run. As can be seen from the plot, the calibration is extremely stable, and the variation of the average difference between the cluster time and the event t_0 is much smaller that the event t_0 variation.




Figure 145. For each run reported on the x axis, the distribution over all sensors of the average of the difference between the calibrated cluster-related-to-tracks time and the event t_0 , divided by the RMS of the event t_0 is shown for the u/P (v/N) side sensors in blue (red). The inset shows a blown up image of the distributions for 4 runs, to better illustrate the figure.

9.7 Alignment

For a precise reconstruction of tracks and vertices the position of silicon sensors has to be determined with average precision better than $10\ \mu\text{m}$. The required precision is much higher than what was achieved during detector installation, of order of hundreds of microns. Thus, the positions of the sensors are determined by means of track-based alignment, which can achieve $1\text{--}2\ \mu\text{m}$ precision according to MC simulation studies, which do not simulate possible time-dependence. In terms of alignment, the difference between the measured (m) and predicted (p) positions of a hit i for the track j ($u_{ij}^{m,p}$) is defined as:

$$r_{ij}(\boldsymbol{\tau}_j, \mathbf{a}) = u_{ij}^m - u_{ij}^p(\boldsymbol{\tau}_j, \mathbf{a}), \quad (9.4)$$

where $\boldsymbol{\tau}_j$ is a vector of track helix parameters and \mathbf{a} is a vector of alignment parameters. Alignment parameter corrections $\delta\mathbf{a}$ are computed via minimisation of a χ^2 function linearized around initial estimates of track and alignment parameters ($\boldsymbol{\tau}_j^0, \mathbf{a}^0$):

$$\chi^2(\boldsymbol{\tau}, \mathbf{a}) = \sum_j^{\text{tracks}} \sum_i^{\text{hits}} \left(\frac{r_{ij}(\boldsymbol{\tau}_j, \mathbf{a})}{\sigma_{ij}} \right)^2 \approx \sum_j^{\text{tracks}} \sum_i^{\text{hits}} \frac{1}{\sigma_{ij}^2} (r_{ij}(\boldsymbol{\tau}_j^0, \mathbf{a}^0) + \frac{\partial r_{ij}}{\partial \mathbf{a}} \delta\mathbf{a} + \frac{\partial r_{ij}}{\partial \boldsymbol{\tau}_j} \delta\boldsymbol{\tau}_j)^2, \quad (9.5)$$

where σ_{ij} denotes the uncertainty of measurement of hit i within track j .

The solution is obtained by a simultaneous minimization with respect to all track parameters (order of millions) and alignment parameters (about three thousand for VXD alignment) using the Millepede II [63] algorithm. The simultaneous fit ensures all correlations among alignment parameters are properly taken into account. With a track sample in a broad topology (cosmic tracks combined with tracks from e^+e^- collisions), the ‘‘weak modes’’ (poorly controlled systematic detector deformations like twist) can be sufficiently constrained. With the full Belle II detector, the VXD alignment is performed simultaneously with the alignment of CDC layers.

The position and orientation of individual sensors (and ladders or half-shells) is characterized by six rigid body parameters: three shifts (u, v, w) and three rotations (α, β, γ). To accommodate sensor deformation, their surface is described using two-dimensional Legendre polynomials with three quadratic parameters (P_{02}, P_{11}, P_{20}) and four cubic parameters ($P_{03}, P_{12}, P_{21}, P_{30}$). The sensor alignment parameters are illustrated in figure 146 (see also figure 6).

The alignment procedure was performed and validated during several stages of detector construction and installation, leading to several optimizations. Cosmic data samples were studied during stand-alone SVD and combined PXD+SVD commissioning runs. Following the VXD installation, both the cosmic tracks and tracks from e^+e^- collisions were processed for alignment purposes. At the beginning of the tests, only the rigid body parameters were considered. But validation algorithms indicated substantial effects caused by deformations (figure 147) of origami sensors (middle-ladder sensors of layer 4, 5 and 6). The reconstruction framework was extended to correct for deformation effects, and the quadratic and cubic surface parameters were introduced into the alignment procedure. Their final impact on the residual distribution for one of the used sensors can be found in figure 147.

In order to verify and monitor time dependence of alignment parameters, a dedicated validation procedure was developed [64, 65]. During the spring and summer 2019 data taking periods, Belle II collected about $5.6\ \text{fb}^{-1}$ in 1200 runs. Within these periods all the alignment parameters were




Figure 146. Sensor alignment parameters: shift and rotation parameters in the local coordinates system, defined in section 8.1.1, and visualisation of the base Legendre polynomials corresponding to quadratic and cubic surface parameters. Reproduced with permission from [64].




Figure 147. Deformation along the w coordinates in μm identified by the monitoring algorithm (left), residual distributions of L4.1.2 sensor during alignment tests (center and right): residual distribution before applying alignment procedure (red), applying rigid body parameters (green), applying quadratic surface parameters (orange), applying cubic surface parameters (violet) and after the second iteration of alignment procedure (blue). Reproduced with permission from [64].

validated and monitored. An example of validation plot can be found in figure 148 — the alignment parameters fluctuate within a range of $\pm 10 \mu\text{m}$.




Figure 148. Time-dependent alignment validation of sensor 4.3.2: shift (first line), rotation (second line), quadratic surface (third line), and cubic surface (fourth line) parameters. Because of validation procedure [64] only the rotation parameters α and β (see text) can be validated. During these data taking periods a fire accident break happened, shown with a magenta vertical line. The units of alignment parameter axes are μm [66] — for cubic and quadratic parameters, this corresponds to the (deviation of) amplitude of the sensor deformation; for angles it is the displacement at sensor edge corresponding to the rotation

9.8 Track finding and reconstruction performance

The SVD information is necessary to improve the overall track finding efficiency as well as the determination of track parameters at the IP.

The impact parameter resolution is dominated by the hit closest to the IP, provided by PXD. However the SVD information plays a crucial role in extrapolating CDC tracks to the PXD sensors with the required precision to pick up the correct hit and also to perform an efficient on-line PXD data reduction, as detailed below. In fact, the uncertainty in the extrapolation of a charged particle to the PXD sensors provided by the SVD is smaller by almost one order of magnitude with respect to the one provided by CDC on the transverse plane and by two orders of magnitudes along the beam axis. Thus the SVD information is necessary to correctly identify the clusters on the PXD belonging to the particle track, that are needed to provide the ultimate resolution on the impact parameters. The SVD information also improves the measurement of the particle transverse momentum. The disadvantage of a smaller lever arm of the SVD with respect to the CDC is counterbalanced by its better spatial resolution.

Importantly, low transverse momentum tracks, unable to reach the CDC, can be reconstructed using SVD-only information. This is fundamental, for example, to reconstruct the slow pions π_s^+ from $B \rightarrow D^{*+} X$, $D^{*+} \rightarrow D^0 \pi_s^+$.

At present, the track finding strategy starts the search from the CDC. Each track found among the CDC hits is extended inward the IP. A combinatorial Kalman-filter is used to identify the best SVD clusters compatible with the extrapolated positions on the sensor. The SVD clusters are attached to the track candidate identified by the CDC to improve the determination of its transverse and longitudinal impact parameters, as well as of momentum vectors at the point of closest approach to




Figure 149. Track finding efficiency as a function of transverse momentum for the CDC-only tracking (black) and the full tracking with SVD information (red) estimated using simulated generic $B\bar{B}$ events.

the IP. The SVD is then used as a stand-alone track finding detector on the remaining clusters not associated to CDC tracks. Its four-layer structure together with its proximity to the IP is particularly valuable to identify and reconstruct particles in the low transverse momentum region $p_t < 0.3$ GeV/ c . In this p_t range the CDC stand-alone track finding efficiency is around 95% at 0.2 GeV/ c then it gradually degrades to 80% at $p_t \sim 0.1$ GeV/ c below which it quickly drops down to zero, as shown in figure 149 [67], because of the limited number of hits left by the particle in the CDC.

A stand-alone track finder [Vertex Detector Track Finding \(VXDTF2\)](#) [68, 69] had been developed to improve the overall track finding efficiency as well as to have a track finder orthogonal to the CDC finder for cross-checks and for the study of systematic effects. The SVD together with the VXDTF2 brings the track finding and fitting efficiency above 87% for $p_t \sim 0.1$ GeV/ c for all particles in $B\bar{B}$ events (see figure 149). Since the angular acceptance of the SVD exceeds the CDC one by a few degrees, a better coverage in the very forward and backward region is also obtained as an additional bonus [67].

The VXDTF2 [68] gradually builds up the track candidates starting from the SpacePoints whose reconstruction is described in section 8.2.2. Pairs of SpacePoints, whose positions are sufficiently close to the trajectory of a charged particle from the IP, are combined in a segment. Pairs of segments that have the same central SpacePoint and that are reasonably close to an helix from the IP are then combined in a triplet. The triplets having a common segment are joined using a cellular automaton [69] that provides a set of raw track candidates. At this stage, each SpacePoint can be associated to several track candidates and the trajectory of a given particle is represented by a set of several track candidates. Finally, each SpacePoint is associated to the single best track candidate identified by a figure-of-merit based on the track to clusters residual, time and energy. The VXDTF2 can exploit the cluster time information by rejecting segments and triplets whose space points are too far away in time. This feature is particularly useful to reduce the combinatorial burden produced

by random combinations of background hits. At present it is leveraged only on the reconstruction of MC simulations at full luminosity and background. The algorithm is fast enough to be used on-line at the [High Level Trigger \(HLT\)](#). The tracks found by the VXDTF2 are then extended outward into the CDC using a combinatorial Kalman-filter. A final combinatorial Kalman-filter adds the PXD clusters to the tracks found in the SVD.

An accurate extrapolation of the reconstructed tracks inside the PXD volume, made possible with the SVD, is also crucial to predict the position at which the particle crosses a PXD sensor and to define a rectangular [Region Of Interest \(ROI\)](#) around that position to perform an efficient on-line PXD data reduction. Given the long integration time of the PXD and its proximity to the IP, the data rate from PXD only, dominated by the background component, could be as high as 10 GB/s, 10 times the rate of the other Belle II subdetectors combined [70]. A PXD data reduction is mandatory at full luminosity in order to reduce to a manageable level the size of the data storage required by the PXD data stream. The background contribution of the PXD can be heavily reduced at an early stage of the data-flow by discarding PXD clusters that are outside all ROIs. The ROI are defined on the HLT and sent to the [ONline Selection Nodes \(ONSEN\)](#) subsystem, which selects the PXD clusters belonging to the ROI and sends them to the Event Builder.

The viability of the PXD data reduction had been demonstrated during several test beams at DESY [71] and was tested during several runs of Belle II data-taking, demonstrating that the required data reduction of a factor 10 is reachable. A second PXD data reduction scheme is at present under test and development [9]. This scheme exploits an FPGA-based SVD stand-alone track finder system called DATCON to define the [ROI](#) and send them to the ONSEN [72].

9.9 dE/dx performance

Particle identification plays an important role in the physics program of Belle II. For instance, a better identification of the slow pion (π_s) would improve the precision at which the lepton-flavor-universality ratio $R(D^*)$ can be measured in the decay channel $B \rightarrow D^* \ell \nu$, $D^* \rightarrow D \pi_s$. Such low-momentum particles ($p_T \approx 65 \text{ MeV}/c$) are often unable to reach the CDC, the main tracking device of the experiment, owing to their small radius trajectory. Our goal here is to exploit specific ionization (dE/dx) by these particles in the SVD towards identifying them. Even if the charged particles have a higher p_T enabling them to reach the CDC, the dE/dx values measured in the SVD can still provide complementary information towards their identification.

According to the Bethe-Bloch formula [73], the dE/dx of a charged particle traversing through detector material only depends on its $\beta\gamma$ value. Thus, if the dE/dx for various types of particles vs. their momentum ($\beta\gamma mc$) is plotted, the difference in the resulting two-dimensional distributions can be used to distinguish these particles in the SVD. Of particular importance is the charged pion-kaon separation performance of the SVD.

The self-tagging decay $D^{*+} \rightarrow D^0(\rightarrow K^- \pi^+) \pi^+$ and its charge conjugate process is used to develop the dE/dx particle identification method. Our study is based on a Belle II data sample of 3.1 fb^{-1} recorded near the $\Upsilon(4S)$ resonance and the obtained results are compared with that of an MC sample. To reconstruct $D^{*+} \rightarrow D^0(\rightarrow K^- \pi^+) \pi^+$ decay candidates, all possible combinations of three tracks in an event, requiring exactly two of them to have same charge, are considered. One of the two same-charge particles is taken to be the slow pion candidate, if its momentum is less than $1.0 \text{ GeV}/c$. The remaining two tracks are assumed to be the kaon and pion candidates from the

D^0 decay, where the latter is required to have the same charge as the “slow-pion” candidate. The reconstructed D^0 mass is required to lie between 1.85 and 1.88 GeV/c^2 , corresponding to about a $\pm 3\sigma$ window around the peak, where σ is the mass resolution.

The signal and background shapes in the D^* and D^0 mass difference (Δm) are modeled with a sum of two Gaussians and a threshold function, respectively. The latter is given by

$$f(\Delta m) = c (\Delta m - m_\pi)^a e^{-b(\Delta m - m_\pi)}, \quad (9.6)$$

where m_π is the nominal mass of the charged pion, and a , b , and c are shape parameters. Figure 150 shows results of the Δm fit to the data and MC sample. There is a qualitative data-MC agreement having a somewhat larger background contribution and broader signal width in data.




Figure 150. Results of the Δm fit performed on the data (left) and MC (right) sample. Points with error bars are data, solid blue curves are the total fit results, dashed red and dotted green curves denote the signal and background component, respectively. Reproduced from [75]. CC BY 4.0.

Various background-subtracted distributions obtained using the $sPlot$ method [74] are then studied. Slow pions are found to have an average momentum of 500 MeV/c , while the momentum spectra for D^0 daughter pions and kaons extend up to 5 GeV/c , and are therefore used for building probability density functions (PDFs) above 500 MeV/c . The dE/dx value obtained from each SVD hit for a given charged particle follows a Landau distribution [76]. The hit-level dE/dx information is then combined to reconstruct the track-level dE/dx value. It turns out that a simple arithmetic mean would be insufficient as the mean of the Landau distribution is undefined. Truncation must be applied to the individual hit-level dE/dx measurements while combining them to obtain the track-level dE/dx value, also reducing the effect of fluctuations in energy loss. When a particle traverses the entire SVD, at least eight dE/dx values are obtained: four each for the u/P and v/N side. To truncate dE/dx , the two highest energy clusters are excluded, since they are expected to come from the same wafer and to be fully correlated, making the removal of just one hit less effective. Figure 151 shows the scatter plot of truncated dE/dx values for pions and kaons as a function of their momentum in the data and MC sample.

This provides us a good motivation to construct a two-dimensional PDF \mathcal{P}_m for the $[(dE/dx), p]$ pair for a likelihood-based particle identification method, where the index m stands for π and K . A likelihood function for each particle hypothesis m is constructed: $\mathcal{L}_m(dE/dx, p) = \prod_i \mathcal{P}_m[(dE/dx)_i, p]$, with i running over all dE/dx values assigned to a track. The dE/dx calibration exercise was performed with the control sample of $D^{*+} \rightarrow D^0(\rightarrow K^-\pi^+)\pi^+$. Finally, it was verified whether the particle identification performance is improved compared to the case where the SVD dE/dx information is not used.




Figure 151. Scatter plot of truncated dE/dx value for pions and kaons as a function of their momentum in the data (left) and MC (right) sample. Reproduced from [75]. CC BY 4.0.




Figure 152. Comparison of performance for pion (left) and kaon (right) identification with and without the SVD dE/dx information. The track momentum it is required to be less than 1 GeV/ c . Reproduced from [75]. CC BY 4.0.

Figure 152 shows the performance improvement obtained for a momentum below 1 GeV/ c in the plot of identification efficiency vs. fake rate. For a fixed kaon fake rate of 12% the pion identification efficiency is found to improve from about 93% (97%) to 94.5% (98.5%) in data (MC) events. Similarly, for a fixed pion fake rate of 4% the kaon efficiency increases from 82% (89.5%) to 86% (92.5%) for data (MC) events. There is a data-MC difference in absolute terms, which can be attributed to the difference observed between data and simulation in the cluster charge distribution. It is expected that the simulated dE/dx performance will match better that on data once a more realistic simulation is available.

10 Conclusion and outlook

The SVD has operated smoothly since the start of data taking, providing well understood, high quality data. It is essential to maintain the data quality in the coming years when luminosity will be gradually increased towards the design value.

Background projections and strategies for mitigation. The current background occupancy in SVD is quite low, below 0.4% in Layer 3 and therefore no specific background rejection strategy is applied in SVD data reconstruction. In future operation with higher luminosity, the background occupancy could increase to the level of 3% in Layer 3 (see figure 111). While this is within the operational limits of SVD, it provides very little margin against a number of adverse effects on SVD performance and operation, such as an increase of fake tracks, data loss due to bandwidth limitations, or an increase in reconstruction CPU time.

In addition to machine background mitigation plans, which are in general required to allow a smooth operation of the entire detector, the very good SVD hit time resolution can be exploited to reject off-time background hits to keep the current excellent tracking performance. The SVD strategy to cope with high background conditions is under development and includes:

- reducing the sampling of the strip amplifier waveform from the standard 6 to 3 samples. The choice between 3- and 6-sample data acquisition can be made event by event on the basis of the expected trigger jitter, estimated from the calorimetric trigger signal size. This technique was shown to preserve the hit time resolution while reducing by 30% on average the required bandwidth;
- applying a selection on the absolute time and the time difference between the two sides when building a [spacepoint](#). This was shown to provide a background rejection in excess of 60% while retaining the full signal efficiency;
- using hit time differences in track pattern recognition to reduce the number of background hits combined with hits from physics events, which leads to badly reconstructed or fake tracks.

The various background mitigation strategies are tested and optimized on the current data and will be gradually introduced in the standard data taking and reconstruction.

The [CDC](#) is also subject to performance degradation in high background conditions, that can lead to tracking efficiency reduction. The SVD has the potential of becoming the central device for pattern recognition, thus finding tracks first in the SVD alone and then extrapolating to the CDC and the [PXD](#). In addition, using the precise SVD timing in the event T0 calculation can improve and speed up CDC track reconstruction.

Perspective on possible upgrades. Currently, the accelerator is operating at a peak luminosity of about $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, still a large factor below the design luminosity of $8 \times 10^{35} \text{ cm}^{-2} \text{ s}^{-1}$. In order to increase luminosity while keeping backgrounds under control, various accelerator consolidation and upgrade options are under study, including the possibility of a significant redesign of the interaction region. Although the PXD and SVD are deemed to be able to cope with high background conditions, margins are limited, and performance may be adversely affected. In addition, the interaction region redesign may require the modification of the detector envelope and mechanical interfaces, imposing the fabrication of a new vertex detector. Therefore the development of solutions for replacing either the PXD, or the SVD, or both, with lighter, more performant, and more robust devices has started. Possible options include thinner double-sided strip detectors, an improved DEPFET-based PXD, and fully pixelated 5-layer systems based on Depleted CMOS Monolithic Active Pixels or Silicon-On-Insulator technologies. Initial Expressions of Interest have

been presented inside Belle II at the beginning of 2021, and it is expected that a Conceptual Design Report of the upgrade will be prepared in the course of 2022.

Conclusion. The Silicon Vertex Detector of Belle II was designed and fabricated by a large international collaboration in the period 2012–2018 using strongly coordinated multiple assembly sites with common procedures and methodologies. Since 2019 it has been in operation providing high quality data. The small number of defective channels ($< 1\%$), the large hit efficiency ($> 99\%$), the good signal-to-noise ratio (well in excess of 10 for all sensor configurations and tracks), as well as the good control over the alignment are all essential factors to achieve good tracking efficiency and accuracy. A continuous effort is in place to monitor and improve the detector performance and to refine the matching between data and simulation, which is essential for all physics analyses.

In this paper we have tried to document all the aspects of the SVD challenges and achievements, in the spirit of providing information to the broader community and help the development of high quality detector systems, which are fundamental tools to carry out physics research.

It should be noted that such an extended paper required a long editorial process and in a few cases some more current results that could not be incorporated in the paper are available. In particular updated performance studies can be found in [77–79], and radiation damage analyses in [80, 81].

Acknowledgments

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreements No 644294 and 822070. This work is supported by MEXT, WPI, and JSPS (Japan); ARC (Australia); BMWF (Austria); MSMT (Czechia); CNRS/IN2P3 (France); AIDA-2020 (Germany); DAE and DST (India); INFN (Italy); NRF-2016K1A3A7A09005605 and RSRI (Korea); and MNiSW (Poland).

Glossary

annealing heat treatment that alters the physical and/or chemical properties of a material. In silicon sensors, it results in doping atoms to diffuse into substitutional positions in the crystal lattice, where they can become donors or acceptors. [128](#)

cluster set of adjacent strips representing a 1-D [hit](#) on the plane of sensor. [138](#), [162](#), [183](#)

common-mode noise a type of noise that causes correlated fluctuations on all the channels. [105](#), [131](#), [139](#)

DOCK space reserved for interconnections of cables and services, close to the [CDC](#) end plates. [6](#), [18](#), [89](#), [92](#)

e-hut three-floor electronics hut, located at one side of the Belle II detector. [92](#), [93](#), [94](#)

end-cone part of the [SVD](#) support structure, cone-shaped and realized in carbon fibre. [5](#), [40](#), [183](#)

end-ring part of the support structure for a given [SVD layer](#), located on an [end-cone](#). [5](#), [40](#), [89](#), [90](#), [183](#)

end-mount aluminum part of the [SVD](#) ladder, located at both ends of the ladder, and joined to the support [end-ring](#). [5](#), [19](#), [22](#)

hit generic term indicating the signal or set of signals induced by a charged particle track crossing the sensor; 1-D hit: signal on a set of adjacent strips in the sensor plane, forming a [cluster](#); 2-D hit: combined signal from 1-D hits on the two opposite sides, in the sensor plane; 3-D hit: complete information on the signal induced by a charged particle track, including space position, time and released ionization charge. [138](#), [162](#), [182](#), [183](#)

kokeshi-pin a precision pin with a groove allowing good planar coupling between two surfaces with a lateral screw. [22](#), [26](#), [74](#)

ladder mechanically and electrically independent subset of [SVD layers](#), resembling the stave of a barrel. [4](#), [5](#), [39](#), [41](#), [89](#), [90](#), [91](#), [92](#)

layer part of the [SVD](#) detector, located at approximately fixed radius. [4](#), [5](#), [39](#), [90](#), [91](#), [92](#), [183](#)

noise root mean square (RMS) of the each strip signal after pedestal subtraction and common-mode corrections. [105](#)

occupancy fraction of strips above threshold in a single triggered event. [121](#), [138](#)

origami pipe The thin stainless steel pipes used to cool the APV25 chips on the origami modules once assembled in ladders. [25](#), [28](#), [29](#), [30](#), [38](#)

origami module One of the module types that use the origami scheme. Three types exist: O-Z, O-CE, O+Z. [14](#)

origami board The flexible printed circuit board used in the origami modules to distribute power and signals to the APV25 chips. [5](#), [17](#), [18](#), [19](#), [24](#), [25](#), [38](#), [91](#), [131](#)

origami scheme Chip-on-sensor concept, obtained by connecting DSSD strips to front-end APV25 ASICs via flex circuits, bent over the DSSD edge. [4](#), [13](#), [14](#), [22](#), [162](#)

pedestal average of each strip signal. [105](#)

pinhole broken AC decoupling capacitor between the strip implant and the metal readout electrode. [8](#), [45](#)

ROI Region Of Interest on the PXD sensors where a signal cluster is probably located. [6](#), [140](#)

spacepoint combination of a u/P side cluster with a v/N side cluster to make the 3-D [hit](#). [138](#), [181](#)

substrate pinhole broken AC decoupling capacitor between the strip implant and the detector substrate. [47](#)

trigger bin the arrival time of the L1 trigger registered by the FADC with a clock frequency that is four times larger than the APV25 clock frequency. The trigger bin allows to move the hit time measured in the SVD reference frame to the trigger reference frame, common to all detectors, and vice-versa. [141](#), [142](#), [150](#), [159](#)

Acronyms

ADC Analog to Digital Converter. [87](#), [120](#)

ASIC Application-specific integrated circuit. [2](#), [4](#), [6](#), [85](#)

basf2 Belle II Analysis Software Framework. [139](#)

BEAST Beam Exorcism for A Stable Belle II experiment. [89](#)

CCD Charge-Coupled Device. [80](#)

CDC Central Drift Chamber. [2](#), [4](#), [108](#), [137](#), [159](#), [162](#), [176](#), [181](#), [182](#)

CTE Coefficient of Thermal Expansion. [26](#)

CFRP Carbon Fiber Reinforced Plastic. [22](#), [25](#), [63](#), [67](#)

CMM Coordinate Measurement Machine. [52](#), [73](#)

CMS Compact Muon Solenoid experiment. [4](#)

CMC Common Mode Correction. [105](#)

COPPER Common Pipelined Platform for Electronics Readout. [6](#), [31](#), [32](#), [33](#), [36](#)

CVD Chemical Vapour Deposition. [86](#)

DAC Digital to Analog Converter. [45](#), [87](#)

DAQ Data Acquisition. [119](#)

DATCON DATa CONcentrator. [6](#), [31](#), [32](#), [36](#)

DCU Diamond Control Unit. [87](#), [89](#)

DEPFET Depleted P-Channel Field Effect Transistor. [89](#)

DQM Data Quality Monitoring. [120](#), [138](#), [139](#), [158](#), [162](#), [165](#)

DSSD Double-sided Silicon Strip Detector. [4](#), [119](#), [162](#)

ELMB Embedded Local Monitor Board. [89](#), [90](#)

EPICS Experimental Physics and Industrial Control System. [94](#)

FADC Flash Analog to Digital Converter. [6](#), [18](#), [31](#), [32](#), [33](#), [34](#), [35](#), [36](#), [40](#), [70](#), [132](#), [140](#)

FBG Fiber Bragg Grating. [91](#), [134](#)

FIFO First In First Out. [15](#), [133](#)

FIR Finite Impulse Response. [34](#), [46](#), [103](#), [120](#), [139](#)

FOS Fiber Optic Sensor. [76](#), [85](#), [89](#), [90](#), [92](#), [134](#)

FPGA Field Programmable Gate Array. [70](#), [87](#)

FTB Finesse Transmitter Board. [6](#), [31](#), [32](#), [36](#)

FTSW Front-end Timing Switch. [32](#), [36](#), [133](#)

FTDI Future Technology Devices International. [93](#)

HEPA High Efficiency Particulate Air. [73](#)

HER High Energy Ring. [87](#)

HLT High Level Trigger. [32](#), [138](#), [140](#), [158](#), [178](#)

HV High Voltage. [87](#), [124](#)

I2C Inter-Integrated Circuit serial communication bus. [15](#)

IOC Input/Output Controller. [98](#), [135](#)

IP Interaction Point. [6](#), [109](#), [151](#), [162](#), [176](#)

LER Low Energy Ring. [87](#)

LVPS Low Voltage Power Supply. [37](#), [126](#)

ONSEN ONline SElection Nodes. [178](#)

NIEL Non-Ionising Energy Loss. [126](#)

NTC Negative Temperature Coefficient. [85](#), [89](#), [90](#), [91](#), [134](#)

MC Monte Carlo method. [109](#)

MIP Minimum Ionizing Particle. [105](#), [163](#)

MOSFET Metal-Oxide-Semiconductor Field-Effect Transistor. [126](#)

MPV Most Probable Value. [163](#)

OPI OPerator Interfaces. [101](#)

PCB Printed Circuit Board. [54](#)

PLC Programmable Logic Controller. [85](#), [94](#)

PT Positioning Tower. [77](#)

PXD Pixel Detector. [2](#), [4](#), [84](#), [89](#), [108](#), [128](#), [137](#), [140](#), [176](#), [181](#), [186](#)

QCG Quality Control and Assurance Group. [41](#), [43](#), [57](#)

QCS final-focus superconducting quadrupole. [85](#)

RC Run Control. [98](#), [100](#)

ROI Region Of Interest. [138](#), [178](#)

RF accelerator radiofrequency, 508.9 MHz. [101](#)

RSA Rib Sub-Assembly. [22](#), [62](#), [63](#)

SBW SVD Backward module. [24](#), [54](#)

SFW SVD Forward module. [24](#), [54](#)

SNR Signal-to-Noise Ratio. [110](#)

SPA Sensor-Pitch Adapter sub-assembly. [62](#)

SVD Silicon Vertex Detector. [1](#), [4](#), [39](#), [84](#), [108](#), [119](#), [161](#), [182](#), [183](#), [186](#)

TAS Two-Arms Support. [77](#)

TCT Transient Current Technique. [86](#)

VLHI [VXD](#) Local Hardwired Interlock. [94](#), [135](#)

V_{SEP} Separation Voltage. [37](#)

VXD Vertex Detector, including [PXD](#) and [SVD](#). [2](#), [4](#), [85](#), [86](#), [92](#), [94](#), [108](#), [119](#), [135](#), [186](#)

VXDTE2 Vertex Detector Track Finding. [177](#)

References

- [1] S. Kurokawa and E. Kikutani, *Overview of the KEKB accelerators*, *Nucl. Instrum. Meth. A* **499** (2003) 1.
- [2] BELLE collaboration, *The Belle Detector*, *Nucl. Instrum. Meth. A* **479** (2002) 117.
- [3] BELLE-II collaboration, *Belle II Technical Design Report*, [arXiv:1011.0352](https://arxiv.org/abs/1011.0352).
- [4] sBELLE DESIGN GROUP collaboration, *sBelle Design Study Report*, [arXiv:0810.4084](https://arxiv.org/abs/0810.4084).
- [5] J. Kemmer and G. Lutz, *New Detector Concepts*, *Nucl. Instrum. Meth. A* **253** (1987) 365.
- [6] BELLE-II collaboration, *Detectors for extreme luminosity: Belle II*, *Nucl. Instrum. Meth. A* **907** (2018) 46.
- [7] M.J. French et al., *Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker*, *Nucl. Instrum. Meth. A* **466** (2001) 359.
- [8] T. Higuchi et al., *Modular pipeline readout electronics for the SuperBelle drift chamber*, *IEEE Trans. Nucl. Sci.* **52** (2005) 1912.
- [9] F. Bernlochner, B. Deschamps, J. Dingfelder, C. Marinas and C. Wessel, *Online Data Reduction for the Belle II Experiment using DATCON*, *EPJ Web Conf.* **150** (2017) 00014 [[arXiv:1709.00612](https://arxiv.org/abs/1709.00612)].
- [10] M. Valentan, T. Bergauer, M. Dragicevic, A. Frankenberger, M. Friedl, C. Irmler et al., *Performance studies on the ohmic side of silicon microstrip sensors*, *Nucl. Instrum. Meth. A* **732** (2013) 182.
- [11] BABAR collaboration, *The BABAR Detector: Upgrades, Operation and Performance*, *Nucl. Instrum. Meth. A* **729** (2013) 615 [[arXiv:1305.3560](https://arxiv.org/abs/1305.3560)].
- [12] I. Rashevskaya, S. Bettarini, G. Rizzo, L. Bosisio, S. Dittongo and E. Quai, *Radiation damage of silicon structures with electrons of 900-MeV*, *Nucl. Instrum. Meth. A* **485** (2002) 126.
- [13] S. Bettarini, M. Bondioli, G. Calderini, F. Forti, G. Marchiori, G. Rizzo et al., *Measurement of the charge collection efficiency after heavy non-uniform irradiation in BABAR silicon detectors*, *IEEE Trans. Nucl. Sci.* **52** (2005) 1054 [[physics/0411182](https://arxiv.org/abs/physics/0411182)].
- [14] M. Valentan, *The Silicon Vertex Detector for b-tagging at Belle II*, Ph.D. thesis, Atominstitut, Technische Universität, Vienna, Austria (2013), <https://doi.org/10.34726/hss.2013.21509>.
- [15] E.A.N. Messomo, *Radiation and Temperature Effects on the APV25 Readout Chip for the CMS Tracker*, Ph.D. thesis, Imperial College, London, U.K. (2002), <http://cds.cern.ch/record/2284989>.
- [16] N. Bingefors et al., *A Novel technique for fast pulse shaping using a slow amplifier at LHC*, *Nucl. Instrum. Meth. A* **326** (1993) 112.
- [17] M. Nakao, C. Lim, M. Friedl and T. Uchida, *Minimizing dead time of the Belle II data acquisition system with pipelined trigger flow control*, in *Proceedings of the 18th Real-Time Conference*, Berkeley, U.S.A., 11–15 June 2012 [[10.1109/RTC.2012.6418145](https://arxiv.org/abs/10.1109/RTC.2012.6418145)].
- [18] F. Faccio, G. Blanchot, C. Fuentes, S. Michelis, S. Orlandi, S. Saggini et al., *FEAST2: A Radiation and Magnetic Field Tolerant Point-of-Load Buck DC/DC Converter*, in *Proceedings of the Radiation Effects Data Workshop*, Paris, France, 16 July 2014, [10.1109/REDW.2014.7004569](https://arxiv.org/abs/10.1109/REDW.2014.7004569)
- [19] B. Verlaat, M. Ostrega, L. Zwalinski, C. Bortolin, S. Vogt, J. Godlewski et al., *The ATLAS IBL CO₂ cooling system*, *2017 JINST* **12** C02064.
- [20] M. Friedl et al., *The Belle II Silicon Vertex Detector*, *Nucl. Instrum. Meth. A* **732** (2013) 83.

- [21] S. Yamada, R. Itoh, K. Nakamura, M. Nakao, S. Y. Suzuki, T. Konno et al., *Data Acquisition System for the Belle II Experiment*, *IEEE Trans. Nucl. Sci.* **62** (2015) 1175
- [22] M. Friedl, T. Bergauer, A. Frankenberger, I. Gfall, C. Irmeler and M. Valenta, *The Belle II silicon vertex detector readout chain*, *2013 JINST* **8** C02037.
- [23] BELLE-II SVD collaboration, *Electronics and Firmware of the Belle II Silicon Vertex Detector Readout System*, *PoS TWEPP-17* (2017) 109.
- [24] A. Fernandez Fernandez, B. Brichard and F. Berghmans, *Irradiation facilities at SCK-CEN for radiation tolerance assessment of space materials*, in *Materials in a Space Environment*, K. Fletcher, ed., vol. 540 of *ESA Special Publication*, pp. 627–632, 2003.
- [25] B. Würkner, *Beam test data analysis and resolution studies for the Belle II Silicon Vertex Detector*, Diploma thesis, Technische Universität, Vienna, Austria (2015), <https://doi.org/10.34726/hss.2015.29265>.
- [26] C. Irmeler, *Upgrade studies for the Belle silicon vertex detector*, MSc thesis, Technische Universität, Vienna, Austria (2008), <https://resolver.obvsg.at/urn:nbn:at:at-ubtuw:1-27472>.
- [27] P. Dolejschi, *Characterisation of interstrip parameters on silicon sensors for the Belle II vertex detector*, MSc thesis, Technische Universität, Vienna, Austria (2012), <https://docs.belle2.org/record/2832>.
- [28] M. Raymond, R. Bainbridge, M. French, G. Hall and P. Barrillon, *Final results from the APV25 production wafer testing*, in *Proceedings of the 11th Workshop on Electronics for LHC and Future Experiments*, Heidelberg, Germany, 12–16 September 2005, p. 453–457 [[10.5170/CERN-2005-011.453](https://arxiv.org/abs/10.5170/CERN-2005-011.453)].
- [29] F. Guarino, C. Hauviller and M. Tavlet, *Compilation of radiation damage test data. 4. Adhesives*, CERN Yellow Reports: Monographs. CERN, Geneva, 2001 [[10.5170/CERN-2001-006](https://arxiv.org/abs/10.5170/CERN-2001-006)].
- [30] C. Irmeler et al., *Construction and test of the first Belle II SVD ladder implementing the origami chip-on-sensor design*, *2016 JINST* **11** C01087.
- [31] BELLE-II SVD collaboration, *Series production testing and commissioning of the Belle II SVD readout system*, *Nucl. Instrum. Meth. A* **958** (2020) 162942.
- [32] G. Bassi, L. Bosisio, P. Cristaudo, M. Dorigo, A. Gabrielli, Y. Jin et al., *Calibration of diamond detectors for dosimetry in beam-loss monitoring*, *Nucl. Instrum. Meth. A* **1004** (2021) 165383 [[arXiv:2102.03273](https://arxiv.org/abs/2102.03273)].
- [33] S. Bacher et al., *Performance of the diamond-based beam-loss monitor system of Belle II*, *Nucl. Instrum. Meth. A* **997** (2021) 165157 [[arXiv:2102.04800](https://arxiv.org/abs/2102.04800)].
- [34] H. Pernegger et al., *Charge-carrier properties in synthetic single-crystal diamond measured with the transient-current technique*, *J. Appl. Phys.* **97** (2005) 073704
- [35] P.M. Lewis et al., *First Measurements of Beam Backgrounds at SuperKEKB*, *Nucl. Instrum. Meth. A* **914** (2019) 69 [[arXiv:1802.01366](https://arxiv.org/abs/1802.01366)].
- [36] B.I. Hallgren, H. Boterenbrood, H.J. Burckhart and H. Kvedalen, *The Embedded Local Monitor Board (ELMB) in the LHC Front-end I/O Control System*, in *Proceedings of the 7th Workshop on Electronics for LHC Experiments*, Stockholm, Sweden, 10–14 Sep 2001, pp. 325–330 [[10.5170/CERN-2001-005.325](https://arxiv.org/abs/10.5170/CERN-2001-005.325)].
- [37] B. Verlaat, *Controlling a Two-Phase CO₂ Loop Using a 2-Phase Accumulator*, in *Proceedings of the International Conference of Refrigeration*, Beijing, China, 21–26 August 2007, ICR07-B2-1565.

- [38] B. Verlaet, A. Van Lysebetten and M. van Beuzekom, *CO₂ Cooling for the LHCb-VELO Experiment at CERN*, in *Proceedings of the 8th IIF/IIR Gustav Lorentzen Conference on Natural Working Fluids*, Copenhagen, Denmark, 7–10 September 2008, T3-08.
- [39] M. Nakao and S. Suzuki, *Network shared memory framework for the BELLE data acquisition control system*, *IEEE Trans. Nucl. Sci.* **47** (2000) 267.
- [40] BELLE-II SVD collaboration, *Run and slow control system of the Belle II silicon vertex detector*, *Nucl. Instrum. Meth. A* **958** (2020) 162706.
- [41] T. Kuhr, C. Pulvermacher, M. Ritter, T. Hauth and N. Braun, *The Belle II Core Software*, *Comp. Softw. Big Sci.* **3** (2018) 1.
- [42] Z.J. Liptak et al., *Measurements of Beam Backgrounds in SuperKEKB Phase 2*, [arXiv:2112.14537](https://arxiv.org/abs/2112.14537).
- [43] H. Tanigawa, *A study of beam background from SuperKEKB on Belle II Silicon Vertex Detector*, MSc thesis, The University of Tokyo, Tokyo, Japan (2019), <https://docs.belle2.org/record/1794>.
- [44] G. Rizzo et al., *The Belle II Silicon Vertex Detector: Performance and Operational Experience in the First Year of Data Taking*, *JPS Conf. Proc.* **34** (2021) 010003.
- [45] A. Chilingarov, D. Campbell and G. Hughes, *Interstrip capacitance stabilization at low humidity*, *Nucl. Instrum. Meth. A* **560** (2006) 118.
- [46] A. Ritter et al., *Investigations on radiation hardness of DEPFET sensors for the Belle II detector*, *Nucl. Instrum. Meth. A* **730** (2013) 79.
- [47] BELLE-II SVD collaboration, *The Belle II Silicon Vertex Detector: Performance and Running Experience*, *PoS ICHEP2020* (2021) 718.
- [48] ROSE collaboration, *3rd RD48 status report*, Tech. Rep., CERN, Geneva (1999), [CERN-LHCC-2000-009](https://arxiv.org/abs/hep-ex/9905009).
- [49] CMS collaboration, *Studies of the CMS Tracker at High Trigger Rate*, Tech. Rep., CERN, Geneva (2009), [CMS-NOTE-2009-015](https://arxiv.org/abs/hep-ex/0903015).
- [50] BELLE-II SVD collaboration, *Beam background study for the Belle II Silicon Vertex Detector*, *Nucl. Instrum. Meth. A* **982** (2020) 164580.
- [51] GEANT4 collaboration, *GEANT4 — a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250.
- [52] BELLE II collaboration, *Belle II Analysis Software Framework (BASF2) - source code*, [10.5281/zenodo.5574115](https://zenodo.org/record/5574115).
- [53] D.J. Lange, *The EvtGen particle decay simulation package*, *Nucl. Instrum. Meth. A* **462** (2001) 152.
- [54] T. Sjöstrand, S. Ask, J.R. Christiansen, R. Corke, N. Desai, P. Ilten et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159 [[arXiv:1410.3012](https://arxiv.org/abs/1410.3012)].
- [55] S. Jadach, B.F.L. Ward and Z. Was, *The Precision Monte Carlo event generator KK for two fermion final states in e^+e^- collisions*, *Comput. Phys. Commun.* **130** (2000) 260 [[hep-ph/9912214](https://arxiv.org/abs/hep-ph/9912214)].
- [56] C. Canali, G. Majni, R. Minder and G. Ottaviani, *Electron and hole drift velocity measurements in silicon and their empirical relation to electric field and temperature*, *IEEE Trans. Electron Devices* **22** (1975) 1045.
- [57] S. Ramo, *Currents Induced by Electron Motion*, *Proc. IRE* **27** (1939) 584.
- [58] W. Shockley, *Currents to Conductors Induced by a Moving Point Charge*, *J. Appl. Phys.* **9** (1938) 635.
- [59] U. Fano, *Ionization Yield of Radiations. 2. The Fluctuations of the Number of Ions*, *Phys. Rev.* **72** (1947) 26.

- [60] F. Cenna, N. Cartiglia, M. Friedl, B. Kolbinger, H.-W. Sadrozinski, A. Seiden et al., *Weightfield2: A fast simulator for silicon and diamond solid state detector*, *Nucl. Instrum. Meth. A* **796** (2015) 149
- [61] L. Wood, M. Bracko, T. Elsethagen, K. Fox, C.F. Gamboa, T. Kuhr et al., *Performance of the Belle II Conditions Database*, *EPJ Web Conf.* **214** (2019) 04050.
- [62] BELLE-II SVD collaboration, *Performance studies of the Belle II Silicon Vertex Detector with data taken at the DESY test beam in April 2016*, *PoS Vertex2016* (2017) 057.
- [63] V. Blobel and C. Kleinwort, *A New method for the high precision alignment of track detectors*, in *Proceedings of the Conference on Advanced Statistical Techniques in Particle Physics*, Durham, U.K., 18–22 March 2002, [[hep-ex/0208021](#)].
- [64] J. Kandra, T. Bilka, L. Kapitanova, M. Uchida, H. Ozaki, T. Van Dong et al., *Calibration and alignment of the Belle II tracker*, in *Proceedings of Connecting the Dots and Workshop on Intelligent Trackers*, Valencia, Spain, 2–5 April 2019 [[arXiv:1910.06289](#)].
- [65] BELLE-II collaboration, *B lifetime and $B^0 - \bar{B}^0$ mixing results from early Belle II data*, [[arXiv:1906.08940](#)].
- [66] BELLE-II collaboration, *Alignment Studies at Belle II Vertex Detector*, *Acta Phys. Polon. B* **51** (2020) 1385.
- [67] BELLE II TRACKING GROUP collaboration, *Track finding at Belle II*, *Comput. Phys. Commun.* **259** (2021) 107610 [[arXiv:2003.12466](#)].
- [68] T. Bilka et al., *The track finding algorithm of the Belle II vertex detectors*, *EPJ Web Conf.* **150** (2017) 00007.
- [69] R. Frühwirth, R. Glattauer, J. Lettenbichler, W. Mitaroff and M. Nadler, *Track finding in silicon trackers with a small number of layers*, *Nucl. Instrum. Meth. A* **732** (2013) 95.
- [70] BELLE II PXD collaboration, *Belle II Pixel Detector Commissioning and Operational Experience*, *PoS Vertex2019* (2020) 015.
- [71] T. Bilka et al., *Demonstrator of the Belle II online tracking and pixel data reduction on the High Level Trigger system*, *IEEE Trans. Nucl. Sci.* **62** (2015) 1155 [[arXiv:1406.4955](#)].
- [72] T. Geßler, W. Kühn, J.S. Lange, Z. Liu, D. Münchow, B. Spruck et al., *The ONSEN Data Reduction System for the Belle II Pixel Detector*, *IEEE Trans. Nucl. Sci.* **62** (2015) 1149 [[arXiv:1406.4028](#)].
- [73] W.R. Leo, *Techniques for Nuclear and Particle Physics Experiments: A How to Approach*, Springer, Berlin, Germany (1987).
- [74] M. Pivk and F.R. Le Diberder, *SPlot: A Statistical tool to unfold data distributions*, *Nucl. Instrum. Meth. A* **555** (2005) 356 [[physics/0402083](#)].
- [75] S. Hazra et al., *Particle Identification in Belle II Silicon Vertex Detector*, *JPS Conf. Proc.* **34** (2021) 010018.
- [76] L. Landau, *On the energy loss of fast particles by ionization*, in *Collected Papers of L.D. Landau*, D. Ter Haar, ed., pp. 417–424, Pergamon, (1965) [[10.1016/B978-0-08-010586-4.50061-4](#)].
- [77] BELLE II SVD collaboration, *The Silicon Vertex Detector of the Belle II Experiment*, in *12th International Conference on Position Sensitive Detectors*, 11, 2021 [[arXiv:2111.13509](#)].
- [78] BELLE-II SVD collaboration, *Simulation of the Belle II silicon vertex detector*, *Nucl. Instrum. Meth. A* **1032** (2022) 166630.

- [79] BELLE-II SVD collaboration, *Measurement of the cluster position resolution of the Belle II Silicon Vertex Detector*, *Nucl. Instrum. Meth. A* **1033** (2022) 166746.
- [80] L. Massaccesi, *Performance study of the SVD detector of Belle II and future upgrades*, MSc thesis, Università di Pisa, Pisa, Italy (2021), <https://docs.belle2.org/record/2759>.
- [81] BELLE-II SVD collaboration, *The Silicon Vertex Detector of the Belle II experiment*, *Nucl. Instrum. Meth. A* **1033** (2022) 166688.
- [82] Hamamatsu Photonics, https://en.wikipedia.org/wiki/Hamamatsu_Photonics.
- [83] Micron Semiconductor Ltd, <http://www.micronsemiconductor.co.uk/>.
- [84] FI-R51S-HF Connector by Japan Aviation Electronics, <https://www.jae.com/en/>.
- [85] AIREX[®] r82.80 technical data sheet, <https://www.3accorematerials.com/en/products/airex-foam>.
- [86] BERGQUIST SIL PAD TSP 1600, https://www.henkel-adhesives.com/at/en/product/sil-pads/bergquist_sil_padtsp1600.html.
- [87] KERATHERM[®] SOFTTHERM[®] 86/125 ceramic filled foil, <https://www.kerafol.com/en/thermal-management/keratherm-produktuebersicht/keratherm-softtherm-folien>.
- [88] Polyether ether ketone, a thermoplastic used in demanding applications, https://en.wikipedia.org/wiki/Polyether_ether_ketone.
- [89] Mitsubishi Chemical Carbon fiber product, https://www.m-chemical.co.jp/en/products/departments/mcc/composite-products/product/1201233_7508.html.
- [90] 316L 1.4404: a molybdenum-containing Austenitic Stainless Steel, <https://www.aperam.com/product/316l-1-4404/>.
- [91] Ansys Engineering Design and Simulation Toolbox, <https://www.ansys.com/>.
- [92] Stycaset data sheets, https://www.henkel-adhesives.com/it/en/product/potting-compounds/loctite_stycast_2850ft.html.
- [93] LEONI wires, cables and wiring systems, <https://www.leoni.com/>.
- [94] Amphenol Spectra-Strip, https://www.datasheetarchive.com/companies/spectra-strip_com.html.
- [95] Cake Software Foundation, *CakePHP (version 2.5.9)*, 2005. <http://cakephp.org>.
- [96] ORACLE, MySQL (version 5.1.73), <http://mysql.com>, 2008
- [97] Tokai Denshi, The company is no longer active.
- [98] Daiei Electronics Co., Ltd, <https://gb.kompass.com/c/daiei-electronics-co-ltd/jp033872/>.
- [99] Leiton PCB Manufacturer, <https://www.leton.de>.
- [100] Taiyo Industrial Co., LTD, <http://www.taiyo-xelcom.com/>.
- [101] Cleanroom classification, <https://en.wikipedia.org/wiki/Cleanroom>.
- [102] Delrin[®] trade name of polyoxymethylene by Dupont, <https://www.dupont.com/brands/delrin.html>.
- [103] Teflon[™] trade name of a Polytetrafluoroethylene composite by Chemours, <https://www.teflon.com/en>.
- [104] Araldite adhesives, <https://www.huntsman.com/products/detail/333/araldite>.
- [105] Mylar[®] trade name of biaxially-oriented polyethylene terephthalate by Dupont Teijin, <https://www.tekra.com/products/brands/dupont-teijin-films/mylar>.
- [106] Gel-Box[™] is a trade name of Gel-Pak[®], <https://www.gelpak.com/gel-box-ad/>.

- [107] SUS304 1.4301: a chromium-nickel Austenitic Stainless Steel, <https://www.aperam.com/product/304-1-4301/>.
- [108] CAEN SpA, <https://www.caen.it/>.
- [109] Element Six (U.K.) Ltd, <http://www.e6.com/>.
- [110] CIVIDEC Instrumentation GmbH, <https://cividec.at/>.
- [111] Rogers Corporation, <http://www.rogerscorp.com/>.
- [112] Controller Area Network (CANbus). Specifications Version 2.0., https://en.wikipedia.org/wiki/CAN_bus.
- [113] Micron Optics Optical Sensing Interrogator sm225, <http://micronoptics.ru/uploads/library/documents/Datasheets/MicronOpticssm225.pdf>.
- [114] Micron Optics ENLIGHT Optical Sensing Analysis software, http://micronoptics.ru/sensing_software.html.
- [115] Murata thermistor NCP15XV103E03RC, <http://www.murata.com/en-sg/products/productdetail?partno=NCP15XV103E03RC>.
- [116] Steinhart-Hart coefficients for NTC thermistors, <https://www.ametherm.com/thermistor/ntc-thermistors-steinhart-and-hart-equation>.
- [117] Vaisala sensors, <https://www.vaisala.com/>.
- [118] Nokeval sensors, <https://nokeval.com/en/>.
- [119] Alpha Moisture humidity sensors, <http://www.amsystems.co.uk/>.
- [120] Michell Instruments dew point sensors, <http://www.michell.com/uk/category/dew-point-transmitters.php>.
- [121] Schneider Electric, Modicon 340 industrial process control module, <https://www.se.com/ww/en/product-range/1468-modicon-m340/>.
- [122] Experimental physics and industrial control system, <https://epics.anl.gov>.
- [123] Control System Studio, <https://controlsystemstudio.org/>.

The Belle-II SVD collaboration

K. Adamczyk,ⁱ L. Aggarwal,ⁿ H. Aihara,^z T. Aziz,^u V. Babu,^{u,1} S. Bacher,ⁱ S. Bahinipati,^d M. Bari,^w Ti. Baroncelli,^k To. Baroncelli,^k G. Bassi,^{w,2} G. Batignani,^{q,p} J. Baudot,^f A. Bauer,^c P.K. Behera,^e T. Bergauer,^c V. Bertacchi,^a S. Bettarini,^{q,p} T. Bilka,^r F. Bosi,^p L. Bosisio,^{x,w} A. Bozek,ⁱ F. Buchsteiner,^c L. Bulla,^c G. Casarosa,^{q,p} G. Cautero,^{w,3} M. Ceccanti,^p Y.Q. Chen,^y S.R. Chendvankar,^u L. Corona,^{q,p} P. Cristaudo,^w T. Czank,^{g,4} S.B. Das,^l N. Dash,^{e,5} G. de Marino,^{p,6} M. De Nuccio,^{p,7} G. De Pietro,^{p,8} S.T. Divekar,^u P. Dolejschi,^c Z. Doležal,^r G. Dujany,^f D. Dutta,^{u,9} C. Finck,^f K.D. Fischer,^c F. Forti,^{q,p,*} M. Friedl,^c R. Frühwirth,^c A. Gabrielli,^{x,w} E. Ganiev,^{x,w} I. Gfall,^c D. Giuressi,^{w,3} B. Gobbo,^w S. Halder,^u K. Hara,^{h,t} S. Hazra,^u T. Higuchi,^g S. Himori,^{v,†} T. Horiguchi,^v C. Irmler,^c A. Ishikawa,^{h,t} H.B. Jeon,^j Y. Jin,^w C. Joo,^{g,10} M. Kaleta,ⁱ A.B. Kaliyar,^u J. Kandra,^r K.H. Kang,^g P. Kapusta,ⁱ P. Kodyš,^r T. Kohriki,^h S. Koike,^{h,†} M.M. Kolwark,^u I. Komarov,^{w,7} M. Kumar,^l R. Kumar,^o P. Kvasnička,^r C. La Licata,^g K. Lalwani,^l L. Lancieri,^{x,w} K. Lautenbach,^a R. Lebourcher,^a S.C. Lee,^j J.Y. Lee,^s J. Lettenbichler,^c Y.B. Li,^b J. Libby,^e T. Lueck,^{p,11} S. Maity,^d P. Mammini,^p R. Manfredi,^w L. Martel,^f A. Martini,^{p,7} L. Massaccesi,^{q,p} S.N. Mayekar,^u G.B. Mohanty,^u S. Mohanty,^{u,5} J.A. Mora Grimaldo,^z T. Morii,^g K.R. Nakamura,^{h,t} Z. Natkaniec,ⁱ K. Negishi,^{v,12} N.K. Nisar,^{u,13} M. Oberegger,^c Y. Onuki,^z W. Ostrowicz,ⁱ F. Otani,^z A. Paladino,^{q,p} H. Palka,^{i,†} E. Paoloni,^{q,p} H. Park,^j M. Pernicka,^{c,†} F. Pilo,^p J. Pirker,^c L. Polat,^a A. Profeti,^p K.K. Rao,^u I. Rashevskaya,^{w,14} P.K. Resmi,^{e,15} I. Ripp-Baudot,^f G. Rizzo,^{q,p} N. Rout,^e M. Rozanska,ⁱ D. Sahoo,^{u,16} T. Saito,^{v,17} S. Sandilya,^{u,18} J. Sasaki,^z N. Sato,^h Y. Sato,^h S. Schultschik,^c C. Schwanda,^c Y. Seino,^m J. Serrano,^a T. Shimasaki,^z N. Shimizu,^z H. Steininger,^c J. Stypula,ⁱ J. Suzuki,^h S. Tanaka,^{h,t} H. Tanigawa,^z G.N. Taylor,^k F. Tenchini,^{q,p} R. Thalmeier,^c R. Thomas,^u R. Tiwary,^u T. Tsuboyama,^{h,t} Y. Uematsu,^z M. Valentan,^c L. Vitale,^{x,w} M. Volpi,^k K. Wan,^z Z. Wang,^z M. Watanabe,^{m,19} S. Watanuki,^{v,20} I.J. Watson,^{z,21} J. Webb,^k O. Werbycka,^w J. Wiechczynski,ⁱ S. Williams,^k B. Würkner,^c H. Yin,^c A. Zanetti,^w L. Zani^a and T. Zhang^z

^a Aix Marseille Université, CNRS/IN2P3, CPPM, 13288 Marseille, France

^b Key Laboratory of Nuclear Physics and Ion-beam Application (MOE) and Institute of Modern Physics, Fudan University, Shanghai 200443, China

^c Institute of High Energy Physics, Austrian Academy of Sciences, Vienna 1050, Austria

^d Indian Institute of Technology Bhubaneswar, Satya Nagar 751007, India

^e Indian Institute of Technology Madras, Chennai 600036, India

^f Université de Strasbourg, CNRS, IPHC, UMR 7178, 67037 Strasbourg, France

^g Kavli Institute for the Physics and Mathematics of the Universe (WPI), University of Tokyo, Kashiwa 277-8583, Japan

^h High Energy Accelerator Research Organization (KEK), Tsukuba 305-0801, Japan

ⁱ H. Niewodniczanski Institute of Nuclear Physics, Krakow 31-342, Poland

^j Kyungpook National University, Daegu 41566, South Korea

^k School of Physics, University of Melbourne, Victoria 3010, Australia

^l Malaviya National Institute of Technology Jaipur, Jaipur 302017, India

^m Niigata University, Niigata 950-2181, Japan

ⁿ Panjab University, Chandigarh 160014, India

^o Punjab Agricultural University, Ludhiana 141004, India

^p INFN Sezione di Pisa, I-56127 Pisa, Italy

^q Dipartimento di Fisica, Università di Pisa, I-56127 Pisa, Italy

^r Faculty of Mathematics and Physics, Charles University, 121 16 Prague, Czech Republic

^s Seoul National University, Seoul 08826, South Korea

^t The Graduate University for Advanced Studies (SOKENDAI), Hayama 240-0193, Japan

^u Tata Institute of Fundamental Research, Mumbai 400005, India

^v Department of Physics, Tohoku University, Sendai 980-8578, Japan

- ^w *INFN Sezione di Trieste, I-34127 Trieste, Italy*
- ^x *Dipartimento di Fisica, Università di Trieste, I-34127 Trieste, Italy*
- ^y *University of Science and Technology of China, Hefei 230026, China*
- ^z *Department of Physics, University of Tokyo, Tokyo 113-0033, Japan*
-
- ¹ *Now at Université de Strasbourg, CNRS, IPHC, UMR 7178, 67037 Strasbourg, France*
- ² *Now at Scuola Normale Superiore, I-56127 Pisa, Italy*
- ³ *Also at Elettra-Sincrotrone Trieste S.C.p.A., Basovizza 34149, Italy*
- ⁴ *Now at Tokyo Metropolitan University, Tokyo 192-0397, Japan*
- ⁵ *Now at National Institute of Science Education and Research, Homi Bhabha National Institute, 752050 Jami, India*
- ⁶ *Now at Université Paris-Saclay, CNRS/IN2P3, IJCLab, 91405 Orsay, France*
- ⁷ *Now at Deutsches Elektronen-Synchrotron, 22607 Hamburg, Germany*
- ⁸ *Now at INFN Sezione di Roma Tre, I-00146 Roma, Italy*
- ⁹ *Now at University of Manchester, Manchester, U.K.*
- ¹⁰ *Now at University of Nebraska, Lincoln, U.S.A.*
- ¹¹ *Now at Ludwig Maximilians University, 80539 Munich, Germany*
- ¹² *Now at Iwate University, Morioka, Japan*
- ¹³ *Now at Brookhaven National Laboratory, Upton, New York 11973, U.S.A.*
- ¹⁴ *Now at TIFPA-INFN, Dipartimento di Fisica, Università di Trento, I-38123 Trento, Italy*
- ¹⁵ *Now at Aix Marseille Université, CNRS/IN2P3, CPPM, 13288 Marseille, France*
- ¹⁶ *Now at Iowa State University, Ames, Iowa 50011, U.S.A.*
- ¹⁷ *Now at ICEPP, University of Tokyo, Tokyo, Japan*
- ¹⁸ *Now at Indian Institute of Technology Hyderabad, Telangana 502285, India*
- ¹⁹ *Now at Nippon Dental University, Niigata, Japan*
- ²⁰ *Now at Yonsei University, Seoul 03722, South Korea*
- ²¹ *Now at Seoul National University, Seoul 08826, South Korea*
- [†] *Deceased*
- ^{*} *Corresponding author*