Track Finding at Belle II

Belle II Tracking Group, Valerio Bertacchi[a], Tadeas Bilka[b], Nils Braun[c], Giulia Casarosa[a], Luigi Corona[a], Sam Cunliffe[d], Filippo Dattola[d], Gaetano De Marino[e], Giacomo De Pietro[f], Giulio Dujany[g], Patrick Ecker[c], Michael Eliachevitch[h], Tristan Fillinger[g], Oliver Frost[d], Rudolf Frühwirth[i], Uwe Gebauer[j], Sasha Glazov[d], Nicolas Gosling[k], Aiqiang Guo[d,l], Thomas Hauth[c], Martin Heck[c], Mateusz Kaleta[m], Jakub Kandra[b], Navid Khandan[d], Claus Kleinwort[d], Thomas Kuhr[k], Simon Kurz[d], Peter Kvasnicka[b], Jakob Lettenbichler[i], Thomas Lueck[k], Alberto Martini[f], Felix Metzner[c], Dmitrii Neverov[n], Carsten Niebuhr[d], Eugenio Paoloni[a], Sourav Patra[o], Leo Piilonen[p], Cyrille Praz[d], Markus Tobias Prim[c], Christian Pulvermacher[c], Sebastian Racs[c], Petar Rados[d], Martin Ritter[k], Armine Rostomyan[d], Bianca Scavino[q], Tobias Schlüter[r], Benjamin Schwenker[j], Stefano Spataro[s], Björn Spruck[q], Henrikas Svidras[d], Francesco Tenchini[d], Dong Thanh[t], James Webb[u], Christian Wessel[h], Laura Zani[a]

[a] *Scuola Normale Superiore and INFN Sezione di Pisa, Largo B. Pontecorvo 3, Pisa, 56127, Italy*
[b] *Charles Univ. Prague, Institute of Particle and Nuclear Physics, Ke Karlovu 3, Praha, 12116, Czechia*
[c] *Karlsruhe Institute of Technology (KIT), Institute of Experimental Particle Physics (ETP), Wolfgang-Gaede-Str. 1, Karlsruhe, D-76131, Germany*
[d] *Deutsches Elektronen-Synchrotron (DESY), Notkestrasse 85, Hamburg, D-22607, Germany*
[e] *Laboratoire de L'accelerateur Lineaire (LAL) Orsay, Batiment 200 BP34 91898 ORSAY Cedex, France*
[f] *Dipartimento di Matematica e Fisica, Università di Roma Tre, I-00146 Roma INFN Sezione di Roma Tre, Roma, I-00146, Italy*
[g] *Université de Strasbourg, CNRS, IPHC, UMR 7178, Strasbourg, 67037, France*
[h] *Rheinische Friedrich-Wilhelms-Univ. Bonn, Physikalisches Institut, Nussallee 12, Bonn, D-53115, Germany*
[i] *Institut of High Energy Physics of the Austrian Academy of Sciences, Nikolsdorfer Gasse 18, Wien, 1050, Austria*
[j] *University of Göttingen, Faculty of Physics, Friedrich-Hund-Platz 1, Göttingen, D-37077, Germany*
[k] *Ludwig Maximilians Univ. München (LMU), Faculty of Physics, Geschwister-Scholl-Platz 1, Munich, D-80539, Germany*
[l] *Institute of High Energy Physics (IHEP), 19B Yuquan Road, Beijing, 100049, China*
[m] *Institute of Nuclear Physics PAN, Division of Particle Physics and Astrophysics, Department of Leptonic Interactions, ul. Radzikowskiego 152, Krakow, 31342, Poland*

Preprint submitted to Elsevier *March 25, 2020*

[n] Nagoya Univeristy, Graduate School of Science, Furou-cho Chikusa-ku, Aichi, Nagoya-shi, 464-8602, Japan
[o] Indian Institute of Science Education and Research (IISER) Mohali, Knowledge city, Sector 81, SAS Nagar, Manauli, Punjab, 140306, India
[p] Virginia Polytechnic Institute and State Univ., 315 A Robeson Hall, Blacksburg, VA 24061-0435, USA
[q] Johannes Gutenberg Univ. of Mainz, Johann-Joachim-Becher-Weg 45, Mainz, D-55128, Germany
[r] LP-Research, Inc., Nishiazabu 1, Minato City, Tokyo, 106-0031, Japan
[s] INFN and Univ. Torino, Via P. Giuria 1 I, Torino, 10125, Italy
[t] Fudan University, Key Laboratory of Nuclear Physics and Ion-beam Application (MOE) and Institute of Modern Physics, Shanghai, 200443, China
[u] Univ. of Melbourne, School of Physics, Melbourne, Experimental Particle Physics Group, Parkville Victoria, Melbourne, 3010, Australia

Abstract

This paper describes the track-finding algorithm that is used for event reconstruction in the Belle II experiment operating at the SuperKEKB B-factory in Tsukuba, Japan. The algorithm is designed to balance the requirements of a high efficiency to find charged particles with a good track parameter resolution, a low rate of spurious tracks, and a reasonable demand on CPU resources. The software is implemented in a flexible, modular manner and employs a diverse selection of global and local track-finding algorithms to achieve an optimal performance.

1. Introduction

The SuperKEKB accelerator complex [1] located at Tsukuba, Japan is designed to achieve a world-record instantaneous luminosity for e^+e^- collisions of $8 \times 10^{35}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The collisions of $4\,\mathrm{GeV}$ positron and $7\,\mathrm{GeV}$ electron beams are recorded by the upgraded successor of the Belle detector [2], which is called Belle II [3]. The expected data sample with an integrated luminosity of $50\,\mathrm{ab}^{-1}$ will allow the Belle II experiment to study B meson decays with unprecedented accuracy.

The high instantaneous luminosity poses, however, several additional challenges. The signal and background rates are expected to increase significantly compared to those observed at Belle. The larger data samples will act to reduce statistical uncertainties, this emphasizes the need to keep systematic effects under control. The experiment therefore requires highly performing track-finding software, capable to cope with high rates and significant background, while maintaining high efficiency and resolution for particles with momenta as low as $50\,\mathrm{MeV}/c$.

The track-finding algorithms used in Belle II are built on the experiment's modular software framework [4] and can be combined for an optimal overall performance. The algorithms use both local and global track-finding methods based on cellular automaton [5, 6, 7, 8] and Legendre transformation [9], respectively, as well as combinatorial Kalman filter (CKF) approaches [10, 11, 12, 13]. A specific feature of the Belle II tracking is a heavy use of multivariate methods, based on the gradient boosted decision tree implementation provided by the FastBDT package [14], to improve background filtering and track-candidate search. The performance of the track finding is estimated using a detailed simulation of the Belle II detector using $\Upsilon(4\mathrm{S}) \to \mathrm{B}\overline{\mathrm{B}}$ events with expected background overlaid.

The paper is organized as follows. Section 2 describes the main components of the Belle II tracking devices: the silicon-based vertex detector (VXD) and the central drift chamber (CDC). Properties of signal events and background are discussed in Section 3. Section 4 describes the event simulation and methods used to gauge the tracking performance. The description of the reconstruction of hits in each of the tracking detectors is given next in Section 5. The general strategy for track reconstruction is outlined in Section 6 after which the CDC track finding is explained in Section 7. Track finding with the silicon vertex detector (SVD) using the concept of dedicated Sector Maps and a local track finding algorithm is discussed in Section 8 followed by the description of the CKF in Section 9. Section 10 presents performance studies of the Belle II track finding using simulated events and Section 11 summarizes the results.

2. Belle II Tracking System

The trajectories of the charged long-lived decay products of the B mesons are measured by the Belle II tracking detectors: the silicon based vertex detector and the central drift chamber. The origin of most of these trajectories is in the proximity of the interaction point (IP). The trajectories pass through the beam pipe which is comprised of two thin walls of beryllium enclosing a duct through which liquid paraffin flows. The inner wall of the beam pipe

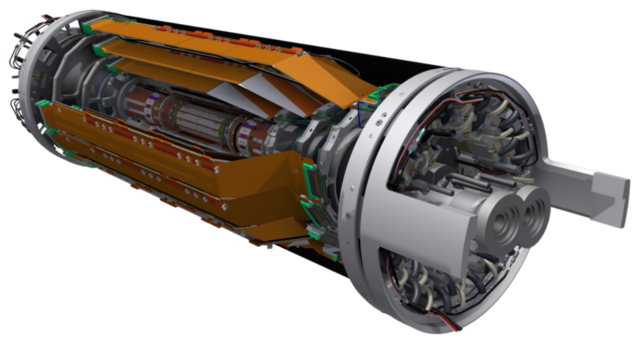

Figure 1: The Belle II Vertex Detector volume. The four outer layers are the silicon vertex detector and the pixel detector is in the center.

is sputtered with a thin layer of gold to shield the VXD from synchrotron radiation. The beam pipe radiation length for particles crossing it at a 90° angle is 0.79%. A thin superconducting solenoid provides a magnetic field of about 1.5 T directed along the nominal mechanical axis of the CDC support cylinder. A system of final focusing quadrupole and compensating solenoid magnets is situated close to the IP. The field remains fairly homogeneous and varies on the order of 1% in the entire tracking volume.

In spherical coordinates, with the z axis parallel to the CDC axis of symmetry and directed along the boost direction, the CDC covers the θ range comprised between 17° and 150° and the full ϕ range. Just outside the CDC there are additional detectors for the reconstruction of neutral particles and particle identification.

Table 1: Specifications of the Belle II PXD.

Layer	Radius (mm)	Ladders	Sensors	Pixels/Sensor $u \times v$	Pitch $u \times v$ (μm × μm)
1	14	8	16	250×768	50 × (55 to 60)
2	22	12	24	250×768	50 × (70 to 85)
Sum		20	40	7 680 000	

A rendering of the VXD is shown in Figure 1. The VXD is composed of two detectors, the pixel detector (PXD) and the SVD, which are based on DEPFET [15] and double-sided silicon strip technologies, respectively. An overview of the key figures of the PXD is shown in Table 1. The PXD consists of two approximately cylindrical layers with radii of 14 and 22 mm. The inner (outer) layer contains eight (twelve) ladders with a size of approximately 1.5 by 10 cm (1.5 by 13 cm). Each ladder is built by gluing two DEPFET modules together at their short edge. In total there are 40 PXD sensors. The ladders overlap with each other in r-ϕ (local u coordinate), while there is a 0.85 mm gap between the two sensors on each ladder in z (local v coordinate). The sensitive region of the PXD is 75 µm in thickness while the edges, which provide the mechanical stiffness to the structure and make the PXD ladder self-supporting, are 450 µm thick. One of the two long sides of each ladder is equipped with twelve switchers, six for each module. These switchers are the only PXD ASICs (Application Specific Integrated Circuits) inside the tracking volume. The other ASICs of the PXD are on the two short edges of each ladder in close contact to the cooling blocks that support the detector. The structure is extremely light with the equivalent thickness for a PXD layer of 0.2% of the radiation length. In both layers, the PXD pixel matrix is organized in rows comprising of 250 pixels with a pitch of 50 µm that run in the u direction, and columns comprising of 768 pixels with pitches varying between 55 µm and 85 µm that run along the v direction. In total the PXD comprises approximately eight million pixels.

Table 2: Specifications of the Belle II SVD.

Layer	Radius (mm)	Ladders	Sensors	Strips/Sensor u, v	Pitch u, v (µm, µm)
3	39	7	14	768, 768	50, 160
4	80	10	30	768, 512	75 to 50, 240
5	104	12	48	768, 512	75 to 50, 240
6	135	16	80	768, 512	75 to 50, 240
Sum		35	172	132 096, 91 648	

Table 2 shows the key figures of the SVD. The SVD consists of four layers of double-sided silicon strip detectors. All the layers have a barrel-shaped part with rectangular sensors. The forward section of the outermost three layers has a lamp-shade geometry made of trapezoidal sensors. This setup minimizes the amount of material for the particles originating from the IP. The radii of the four SVD layers range from 39 mm to 135 mm. The layers consist of 7 to 16 ladders, with 2 to 5 sensors per ladder, respectively. Similarly to the PXD, the SVD ladders overlap in u while there is a 2 mm gap between the sensors on each ladder in v. Each sensor of the first layer of the SVD has 768 strips per side, with readout pitches of 50 µm on the side measuring the u coordinate and 160 µm on the side measuring the v coordinate. The barrel sensors of the three outer layers have 768 strips with a readout pitch of 75 µm in u and 512 strips with a readout pitch of 240 µm in v. The slanted sensors of these layers have the same number of strips in the respective directions, and the same pitch in v. The pitch in u-direction varies from 75 µm at the back to 50 µm at the front side, due to the trapezoidal shape. The readout strips are interleaved with floating strips to improve the spatial resolution. In total, there are 172 SVD sensors with about 220 thousand read-out strips. Each SVD sensor has a thickness of 320 µm. The contribution to the overall radiation length due to mechanical support structure, electronic read-out and cooling is kept at a minimum so that the material of the outer SVD layers is equivalent to 0.6% radiation length at normal incidence.

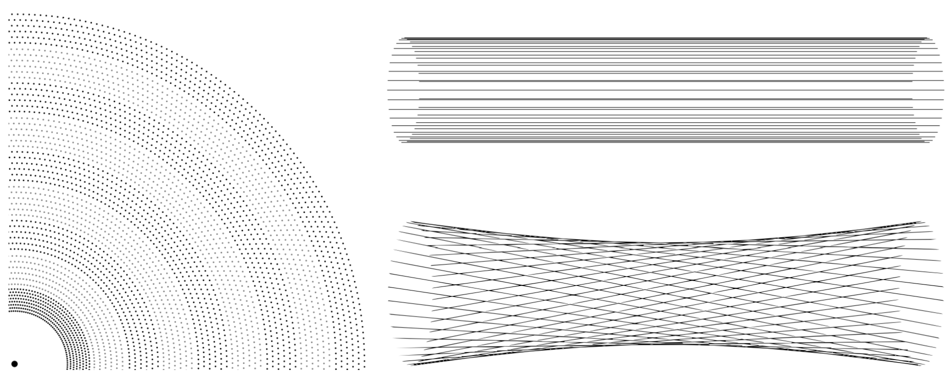

Figure 2: Left: A quadrant of a slice of the r-ϕ projection of the drift chamber. The innermost superlayer contains eight layers, all others contain six. Right: A visualization of stereo wires (bottom) relative to axial wires (top). The skew is exaggerated.

The main specifications of CDC are given in Table 3. The inner volume of the CDC contains about 50 000 sense and field wires, defining drift cells with a size of about 2 cm. The electric field in the drift cells is approximately cylindrical leading to a two-fold ambiguity with the same drift time measured for the tracks passing at the same distance on either side of the sense wire (*left-right passage* ambiguity). The sense wires are arranged in layers, where six or eight adjacent layers are combined in a superlayer, as seen in Figure 2. The outer eight superlayers consist of six layers with 160 to 384 wires. The innermost superlayer has eight layers with 160 wires in smaller

Table 3: Specification of the Belle II CDC.

Layer	Radius of Sense Wires (mm)	Number of Wires	Drift Cell Size (cm)	Average Resolution (µm)
1 to 56	168 to 1111.4	160 to 384	$\sim$1 to $\sim$2	120

(half-size) drift cells to cope with the increasing background towards smaller radii. The superlayers alternate between axial (A) orientation, aligned with the solenoidal magnetic field, and stereo (U, V) orientation. Stereo wires are skewed by an angle between 45.4 and 74 mrad in the positive and negative direction. The direction changes sign between U and V layers, with a total superlayer configuration of AUAVAUAVA. The drift distance resolution of the drift chamber is about 120 µm. By combining the information of axial and stereo wires it is possible to reconstruct a full three-dimensional trajectory.

3. Belle II Events and Background

The events recorded by the Belle II experiment can be classified according to the e^+e^- scattering process occurring at the interaction point. The main category is composed of the $\Upsilon(4S)$ events in which the annihilation of an electron-positron pair produces an $\Upsilon(4S)$ resonance. This resonance decays promptly into a quantum entangled state of two B mesons. The B meson decay vertices have an average spatial separation of $\sim$130 µm. Thus a tracking detector resolution significantly better than that is required to resolve them. This is crucial for the measurements of the time dependent CP and T violation as well as tests of the CPT symmetry in the B meson system. The decay-vertex resolution relies on the spatial resolution of the PXD sensors as well as their proximity to the IP in order to reduce the extrapolation lever arm, and thus the effects of multiple Coulomb scattering on the measurement of the impact parameters.

Studies of (semi)leptonic B decays often require the reconstruction of the missing neutrino by exploiting four-momentum conservation. Hence the tracking algorithm needs to find all of the charged final state particles. This is demanding since there are about 11 tracks per event on average. Moreover, the momentum spectrum of the particles is quite soft, ranging from a few

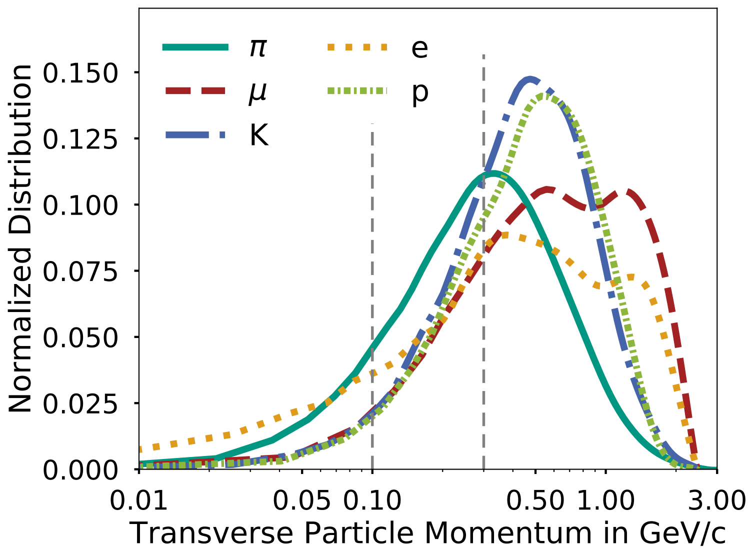

Figure 3: Transverse momentum distributions of primary charged particles as simulated for $\Upsilon(4S)$ events. A logarithmic scale is used for the x axis. The distribution of each charged particle type is normalized to the total number of tracks from the respective type. The vertical line at 100 MeV/c indicates the transverse momentum threshold below which a track can only be found by the SVD. Charged particles with transverse momenta below the value 300 MeV/c marked by the second vertical line can curl inside the CDC volume.

tens of MeV/c to a few GeV/c (Figure 3). It is also essential to keep the rate of fake and duplicate tracks as low as possible.

The reconstruction of particles with momenta below 200 MeV/c is particularly challenging since the trajectories are heavily affected by multiple Coulomb scattering and by energy loss in the material. Moreover, only the measurements of the four layers of the SVD are available to the pattern recognition algorithms for most of the tracks in this low momentum region. The soft momentum spectrum is also challenging for the CDC since particles with momenta below 300 MeV/c can loop several times in the CDC volume producing hundreds of hits. The relative abundance of the long lived charged particles produced in $\Upsilon(4S)$ decays is illustrated in Figure 4.

Other categories of events are also of importance to the experiment. Most notably τ-pair and $c\bar{c}$ events improve the existing limits and measurements on the τ lepton sector and on the charmed mesons. The experiment will also be used to search for non–Standard Model particles, i.e. dark photons, axion-

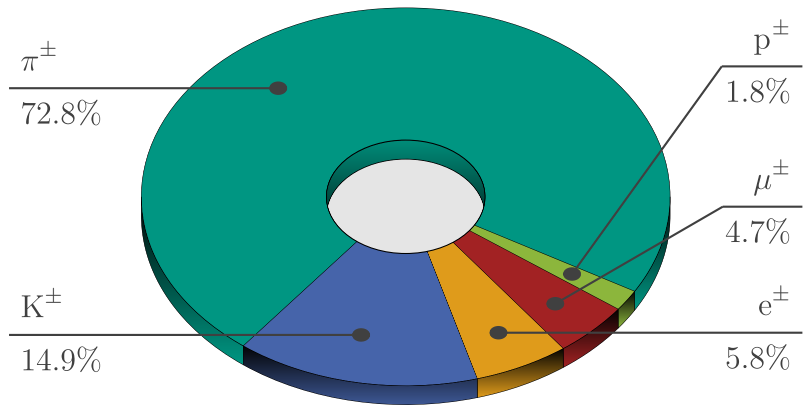

Figure 4: Fractions of charged particle types in generic $\Upsilon(4S)$ events.

like particles, or magnetic monopoles that might be produced directly in e^+e^- collisions. These events are characterized by a lower track multiplicity, a stiffer momentum spectrum, and by a less spherical event topology.

Particles lost by beam-gas and Touschek scattering, as well as due to non-linearities of the machine lattice, lead to additional hits in the detector. The occupancy due to this machine background is expected to be very high as a consequence of the high beam currents, small emittances, and large beam-beam tune shifts needed to reach the design luminosity [16]. The electromagnetic processes occurring at the interaction point, radiative Bhabha and electron-positron pair production, whose cross sections are of the order of several mb, are going to be the leading effects for the beam particle loss rate at nominal luminosity.

The VXD occupancy is expected to be largely dominated by soft electron-positron pairs produced at the IP by the process $e^+e^- \to e^+e^-e^+e^-$. The forward and backward sections of the SVD may be hit by an electromagnetic shower originating from Bhabha electrons interacting in the support structure of the final focusing magnets. The number of background hits exceeds the signal hits by two orders of magnitude resulting in a PXD inner layer pixel occupancy close to 2% and an SVD inner layer strip occupancy close to 3%.

The CDC occupancy is also expected to be dominated by the hits left by particles coming from electromagnetic showers initiated by beam particles. These interact with the material around the final focusing magnets which are well inside the CDC volume. Figure 5 shows the CDC measurements pro-

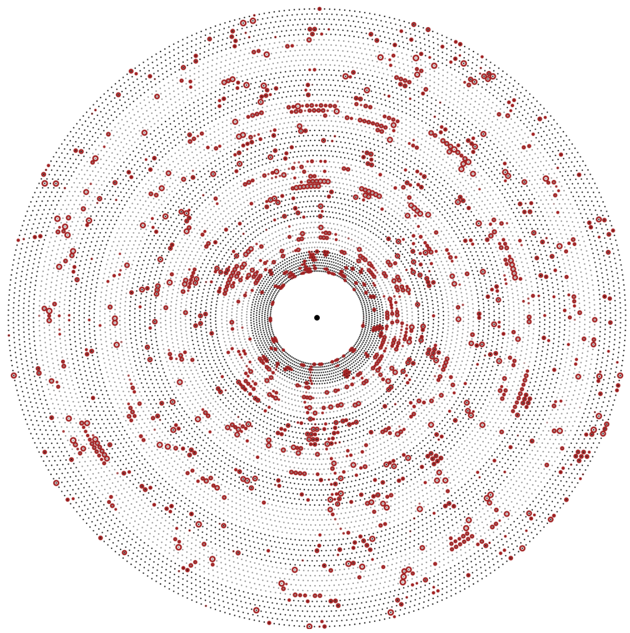

Figure 5: CDC measurements produced by simulated beam-induced background anticipated for the nominal instantaneous luminosity.

duced by simulated beam-induced background for the nominal instantaneous luminosity.

4. Simulation and Track Finding Efficiency Definition

A full simulation tool based on Geant4 [17] is used to model the detector and collider properties. Using the information from the particle generator and the Geant4 simulation of the particles traversing the detector volume, an ideal track finder, called Monte Carlo (MC) track finder, is implemented. Its performance is limited only by the detector acceptance, efficiency and resolution, and by definition cannot be surpassed. A set of figures of merit has been developed to qualify and tune the track finding algorithms. The analysis is limited to tracks identified by the MC track finder (MC-tracks from now on) having enough hits to completely determine the five parameters of the helix-like trajectory. A good track finding algorithm should behave as closely as possible to the MC track finder. In particular, each track should be assigned all of the hits of one and only one MC particle. Two figures of merit are defined for each pair of MC-track and a track found by the pattern recognition (PR-track):

- The *hit efficiency* quantifies how efficient the pattern recognition is in identifying *all* the hits belonging to a single particle. It is defined as the fraction of hits of a given MC-track contained in a given PR-track. Ideally, there should be one and only one PR-track containing all the hits of a given MC-track, thus the *hit efficiency* should be 100% for the correct pair and zero for all others.

- The *hit purity* quantifies how precise the pattern recognition is in identifying the hits belonging to *only one* particle. It is defined as the fraction of hits of a given PR-track contained in a given MC-track. Ideally, there should be one and only one MC-track to which all the hits of a given PR-track belong, thus the *hit purity* should be 100% for the correct pair and zero for all others.

A PR-track is defined as *matched* to a given MC-track if the hit purity exceeds 66% and the hit efficiency exceeds 5%. The low hit efficiency requirement accounts for low momentum tracks curling in the tracking volume which may leave several hundred of hits.

If there are two or more PR-tracks that are matched to the same MC-track, the PR-track with the highest hit purity is defined as the correctly identified match and the remaining PR-tracks are defined as *clones*. If multiple PR-tracks have the same hit purity, the hit efficiency is used in addition to the purity to identify the match. The *track finding efficiency* is defined as the fraction of matched MC-tracks over all MC-tracks.

If the PR-track fails the purity requirement, e.g. the PR-track is made up of hits from two MC-tracks, each one with a hit purity below 66% or the PR-track is made of background hits, it is defined as a *fake*.

5. Input to Tracking Algorithms

PXD Reconstruction

In order to reduce the Belle II data rate to an acceptable level, events are required to pass a software-based high-level trigger (HLT). Data from the PXD do not contribute to the HLT decision, and are therefore buffered in the readout chain. In case an event is accepted, the track information from the HLT is used to define so-called *Regions Of Interest* (ROIs) on the PXD planes. Only PXD hits within these ROIs are stored.

Table 4: Cluster position resolutions for different cluster sizes (one, two and larger than two) and strip pitch, evaluated on MC simulation. Resolutions are measured as 68% coverage of the residual distributions.

Sensor	Side	Pitch (µm)	Resolution (µm) Size = 1	Size = 2	Size > 2
Layer 3	u	50	5.2	3.7	7.6
	v	160	18.1	12.1	18.0
Slanted	u	52–75	6.8	4.5	8.6
	v	240	34.4	18.0	21.4
Barrel	u	75	7.7	5.1	8.8
	v	240	24.8	17.1	20.5

Neighboring pixels with a charge above a threshold are combined into clusters. The cluster position and charge are taken as input for the tracking algorithm.

SVD Reconstruction

The SVD reconstruction software provides in addition to cluster charge and position information also cluster time information to the tracking algorithms.

SVD clusters are formed by combining adjacent strips with a signal-over-noise ratio (SNR) above three. At least one strip in the cluster is required to have a SNR above five. The charge of the cluster is computed as the sum of the charges of the strips, while the cluster time and position are evaluated as the charge-weighted average of the strip times and positions, respectively. The cluster position resolution depends on the cluster size and strip pitch, as shown in Table 4.

The creation of space points follows the clustering and is achieved by combining all clusters on one side of a sensor with clusters on the other side. The only requirement is that the cluster time (on both sides) is greater than a minimum value. The cluster time information helps to reject the majority of the out-of-time clusters, created by beam-background particles produced before, or after the collision event of interest. The SVD cluster time resolution varies between 2 ns (for Bhabha events) and 4 ns (hadronic events), being slightly better on the v side due to a faster response of the electronics on that side. The good time resolution allows the algorithm to reject 60% of the background space points, while retaining 100% of the interesting ones.

The cluster time information provided by the SVD is also used later in the reconstruction, in the pattern recognition step, as described in Section 8.

CDC Reconstruction

The front-end read-out electronics of the CDC use a time-to-digital converter with a 1 ns resolution. This is used to measure the time between the event's trigger signal and the arrival of the drift electrons at the sense wire, the so-called *drift time*. With the x-t relation function, which is an approximation between drift time and distance parameterized in various areas of the CDC, the actual relative distance between the sense wire and a passing particle can be computed and used for track finding purposes. An additional front-end read-out provides amplitude information, sampled at 33 MHz. This information is used for the determination of the energy loss, employed by the particle identification. It can also be used to separate signal and background hits.

6. High-Level Description of the Tracking Setup

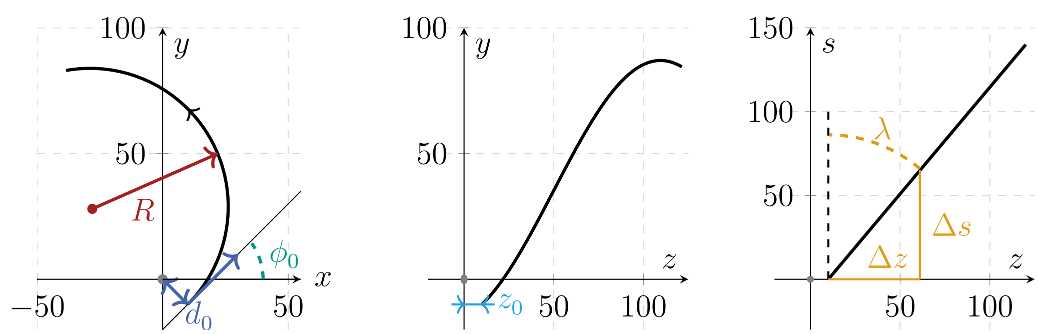

Figure 6: A schematic representation of the track's trajectory in the x-y (left), z-y (middle) and z-s (right) projections. All dimensions are in cm. The track parameters are: d_0, the signed distance of the closest approach to the z axis (POCA); ϕ_0, the angle defined by the x axis and the track transverse momentum at the POCA; z_0, the z coordinate at the POCA; and λ, the track dip angle. Also shown is the track radius R, which is the inverse of the absolute value of the track curvature ω.

The Belle II software uses data processing *modules*, written in C++, which are loosely coupled and transfer data via a common exchange container. This allows for the reconstruction task to be split into different subtasks which can be placed into a chain of independent and interchangeable

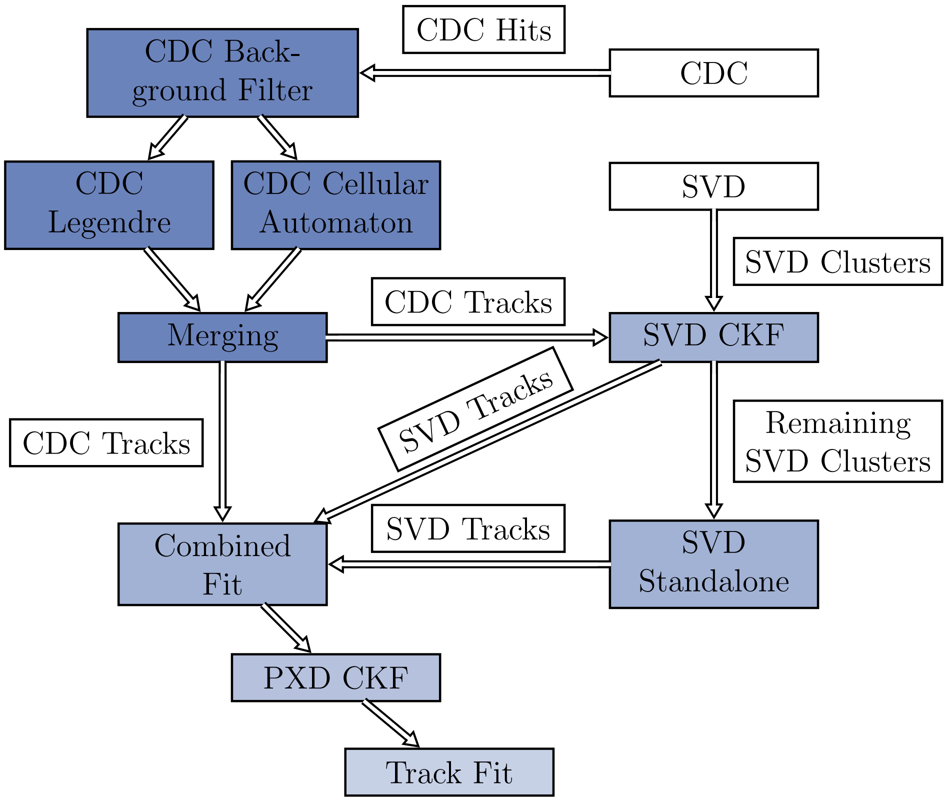

Figure 7: Overview of the steps performed for track reconstruction at Belle II. See text for more details.

modules performing the corresponding task. The `RecoTrack` class is used as a common exchange format between algorithms to transfer track candidates from the different tracking detectors and their respective hits or clusters. The final output of the track reconstruction is the `Track` class, which provides the fitted track parameters for the analysis user.

The track trajectories are represented locally using the helix parameterization, see Figure 6. The three helix parameters in the x-y plane are: the signed distance of the point of closest approach (POCA) to the z axis, d_0; the angle defined by the x axis and the track transverse momentum at the POCA , ϕ_0; and the track curvature signed with the particle charge, ω. The helix can be represented by a straight line in the s-z space, with s being the path length along the circular trajectory in the x-y projection. The two corresponding parameters are: the z coordinate at d_0, z_0; and the tangent of the dip angle $\tan\lambda$.

Figure 7 shows an overview of the steps performed for track reconstruction at Belle II. Due to the very different properties of the three tracking detectors, different algorithms are used for each of them. As a first step, the measured signals in the CDC are filtered and reconstructed by two independent algorithms: a global track finding based on the Legendre [9] algorithm and a local algorithm employing a cellular automaton. The results of both algorithms are merged and the CDC-only tracks are fitted employing a deterministic annealing filter (DAF) [18]. A combinatorial Kalman filter (CKF) is used to enrich the CDC tracks with SVD clusters. High-curvature tracks that did not produce enough hits in the CDC are reconstructed with a standalone SVD track finder using an advanced filter concept called Sector Map and a cellular automaton. The results are combined, fitted again with a DAF and extrapolated to the PXD with a second CKF. At this step the track finding stage is complete. The following sections describe these steps in more detail.

The final step after the track finding includes a track fit using the DAF provided by the GENFIT2 [18] package. For the fit, a specific particle hypothesis must be assumed to calculate the energy loss and the material effects correctly. In Belle II, all reconstructed tracks are fitted with the π, K and p hypotheses. The results of the fit is stored to be used in physics analyses.

7. CDC Algorithm

Two distinct algorithms are used for the track finding in the CDC: global, and local track finding. This enables a high track-finding efficiency while keeping the fake rate low. The global track finding searches for patterns of hits consistent with helix trajectories, even with missing hits, while the local track finding detects extended patterns of nearby hits.

Both algorithms make use of the specific geometry of the CDC and exploit the flexibility of the software framework. The software is written in a modular manner allowing for different sequences of algorithms. Currently, the global track finding is performed first, after the initial filtering of the CDC hits. Thus the global algorithm serves as the primary finding algorithm, which is followed by the local track finding algorithm. The latter helps with reconstructing displaced tracks which originate far away from the interaction point. The track candidates of both algorithms are then merged and post-processing is performed to remove falsely attached hits and, potentially, to attach additional ones. The reconstructed tracks are then passed to the DAF algorithm to be fitted.

Global CDC Track Finding

The global track finding in the CDC is based on the Legendre transformation [9]. It is first performed in the r-ϕ plane, using wire information from axial layers only. After that, it is extended to the three-dimensional space, by attaching wires from stereo layers to existing r-ϕ trajectories. The primary target of the algorithm is finding tracks originating from the vicinity of the origin in r-ϕ. It is adjusted to identify also slightly offset tracks.

In the first step of the algorithm, the position information in axial layers is approximated by drift circles. These drift circles are calculated using a calibrated x-t relation and time information corrected for particle time-of-flight and signal propagation time along the sense wire. For the time propagation correction, it is assumed that particle trajectories are straight lines from the origin, that particles travel with the speed of light, and that they cross the sense wires in the middle. These assumptions are revised when the track parameters are determined.

The reconstruction in the r-ϕ plane continues with a conformal mapping with the center at the origin. This operation transforms circular trajectories starting from the origin to straight lines while the drift circles remain circles. The track finding in the conformal space is thus reduced to the determination of straight lines tangential to a set of circles.

The equation of a tangent to a drift circle in conformal space can be represented using the two Legendre parameters ρ and θ as

$$\rho = x_0 \cos\theta + y_0 \sin\theta \pm R_{\mathrm{dr}} \,,$$

where $(x_0,\, y_0)$ and R_{dr} represent the center of the circle and its radius, respectively. Hence, each drift circle maps to a pair of sinusoids in the ρ-θ track-parameter space. The track recognition and track parameter determination correspond to finding the most populated regions in the ρ-θ space. An efficient method to localize these regions is a two-dimensional binary search algorithm, as illustrated in Figure 8a. The algorithm consists of dividing the ρ-θ space into four equally sized bins and selecting the most populated of them for further subdivision, until convergence.

The two-dimensional binary search algorithm uses a dedicated *quad tree* data structure to store intermediate search results. Each node in the quad tree is linked to four children, corresponding to four sub-bins of the node. In

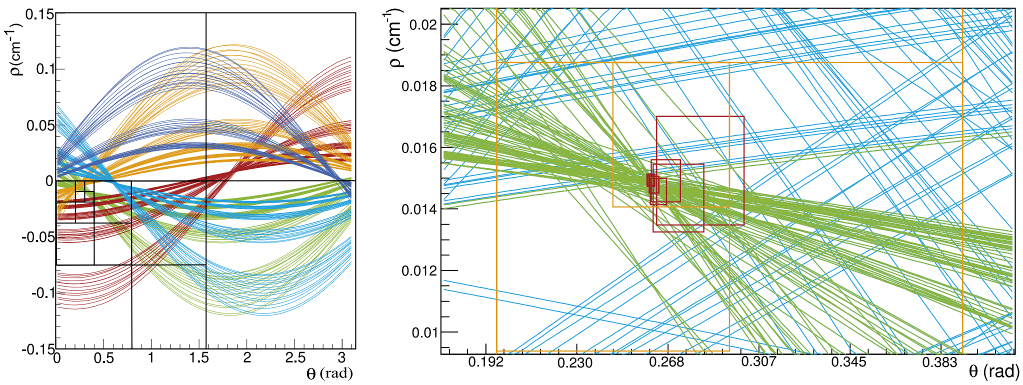

(a) Two-dim. binary search

(b) Sliding bins example

Figure 8: Examples of: (a) standard two-dimensional binary search algorithm; (b) modified algorithm with variable bin size. See text for more details.

general, the search is continued only for the sub-bin containing the most hits. However, it is possible to step back and examine other directions, without repeating the search from the beginning, which speeds up the search for multiple track candidates.

The binary search stops when the bin size becomes smaller than a resolution parameter that is taken to be dependent on ρ. This accounts for the smearing of the track parameters due to the energy loss, non-uniformity of the magnetic field, displaced IP, uncertainty of the drift circle radii and wire displacements. The resolution function is optimized using simulated events.

The introduction of the resolution function as the stopping criterion allows to extend the algorithm to non-standard bin sizes. For a track that is displaced from the origin, the crossing points in the Legendre space may be split between two bins. This effect can be reduced greatly by allowing for overlapping bins. Bins extended by 25% with respect to the exact division are used. A positive side effect of this feature is that the overlapping bins tend to *slide* towards the maximal density of intersections, as illustrated in Figure 8b.

Multiple tracks are found iteratively, using several passes over the Legendre space. At each pass a new track candidate is declared to be found when it satisfies certain quality criteria, such as the number of attached hits. These quality criteria can be varied to increase finding efficiency for different track topologies. Hits corresponding to the found track candidates are removed from further iterations. The high-momentum tracks crossing all CDC layers are searched for first, followed by curling tracks and tracks with large longitudinal momentum, which leave the chamber at smaller radii.

The r-ϕ track candidates are subjected to a post-processing step, performed in the physical r-ϕ space using the fast fitting algorithm of [19]. Firstly, the track candidates are checked to see if they can be merged, to reduce the clone rate. The merge algorithm uses a χ^2-based criterion, comparing the quality of the circular fits to the hits from the separate track-candidates to the fit to the combined set of hits. In addition, hits from the track candidates are examined to determine if they have to be removed or reassigned to other tracks. Finally, all unassigned hits are checked to determine if they can be attached to the existing track candidates.

Hits from stereo layers, containing z information, are added to the r-ϕ trajectories at the next step. The r-ϕ trajectory is used to reconstruct the position information of each stereo measurement. As the stereo wire can be approximated by a straight line and the drift circle does not have direction information, finding the position gives two solutions: either the drift circle is enclosed by the trajectory circle, or not — giving two possible position values for each hit. Stereo hits with a reconstructed z coordinate z_{rec} determined far outside the detector volume are dismissed. Given that z_{rec} depends strongly on the estimated r-position of the trajectory at the stereo wire which may be not very accurate, hits as far as twice the physical drift chamber length are retained.

The problem of track finding becomes very similar to the search in the conformal mapping of the r-ϕ space which makes it possible to use the same algorithm as described above. This time, the trajectory is straight in the s-z space and it can be described by the equation

$$z_0 = z_{\mathrm{rec}} - \tan\lambda \cdot s_{\mathrm{rec}}\,,$$

with s_{rec} being the path to the stereo-wire hit. This gives a line of possible trajectory parameters $(z_0, \tan\lambda)$ for each stereo-layer hit. The point with the most intersections of the lines in the parameter space is used to determine the track parameters. For this, an analogous implementation of the quad tree algorithm as described above is used.

The stereo-wire hits that are found are added to the r-ϕ track only in the case they are not selected for another r-ϕ track. In the latter case the hits are not added to any track which increases the purity of the hit assignment.

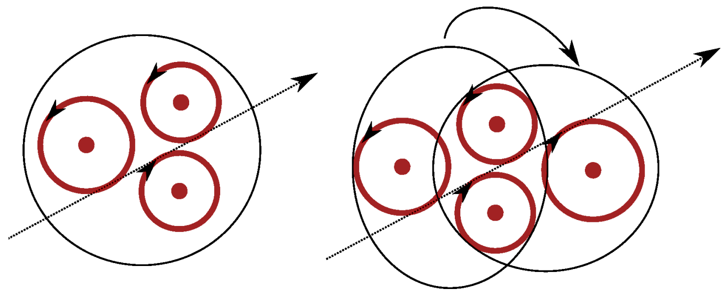

Figure 9: Combination of three neighboring wire hits to a graph vertex (left) and two triplets sharing two wire hits to a graph edge (right).

Local CDC Track Finding

To complement the global search approach, and to detect short tracks and tracks displaced from the IP with a high efficiency, the local track finder operates without any assumption on the origin of tracks. The algorithm searches for connected hits in the CDC superlayers, so called segments. This search uses the cellular automaton concept, which acts on an acyclic graph of vertices connected by edges. More specifically, a weighted cellular method is used, where the vertex i has the weight Θ_i and the edge between the vertices i and j has the weight w_{ij}. Now, track finding can be formulated as maximization of an energy function which can be formulated with

$$E_i = \sum w_{ij} + \sum \Theta_j \,,$$

where the sums are taken along a path to the vertex i. The concept of the weighted cellular automaton is employed in two different stages:

Segment building stage. Vertices (*triplets*) are formed by combining three neighboring hits and assuming the left-right passage hypotheses for a unique trajectory through these three hits[1]. A linear trajectory is then extracted from the measured drift circles by a least-squares method and the weight Θ_j is assigned based on the χ^2 value of the fit. Edges are created from neighboring triplets that share two hits and which pass loose feasibility cuts, with the weight w_{ij} determined based on the χ^2 value of the straight line fit to the four drift circles (see Figure 9). The algorithm allows for information missing from one CDC layer.

[1] In general, several triples are built for a given set of three neighboring hits, depending on the left-right passage hypothesis.

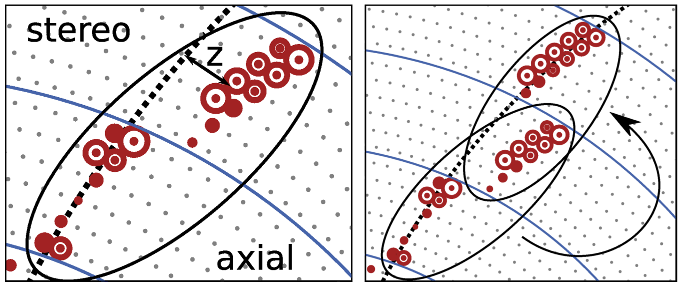

Figure 10: Combination of a pair of axial- and stereo-wire segments to one graph vertex (left) and the combination of vertices, that share one segment, to a graph edge (right).

Track building stage. This stage combines the individual segments found in the axial and stereo superlayers to longer tracks. The vertices are created from a pair of segments in neighboring axial- and stereo-wire superlayers. The weight Θ_j of each vertex is computed with a χ^2 circle fit using the Riemann method [20] and the reconstruction of the z coordinate is performed using a linear fit in the s-z space. Neighboring vertices that share one segment form the edges in the cellular automaton's graph (see Figure 10). The corresponding weight w_{ij} is computed based on the χ^2 value of the fit to hits from all segments. Additional information, such as the number of hits per segment, can be included in the weight calculation using multivariate analysis methods.

Combination of Local and Global Tracking Results

The two track finding algorithms described above are both used to find tracks from the full set of CDC hits. This is done to exploit their specific benefits, with the global track finding capable to reconstruct tracks with several missing layers and the local track finding having similar efficiency regardless of the track origin.

For combining the results of both tracking approaches, the track candidates from the global track finder are used as a baseline. Segments found by the local track finder are added to those tracks using a multivariate approach. The track-segment combination is based on FastBDT [14], which uses several variables calculated from the track and the segment (e.g. the number of common hits, helix parameters, hit-to-trajectory distances) into one single number, which classifies between correct and wrong matches. The multivariate method is trained using simulated events.

Several quality filters based on multivariate estimators are applied to the found tracks and their hits. This increases the hit purity, improves the track parameter resolution, and decreases the rate of fake and clone tracks.

8. SVD Standalone Algorithm

A dedicated standalone algorithm is employed for the task of track finding with the SVD. This algorithm reconstructs the low momentum particles with a transverse momentum of less than 100 MeV/c which deposit too few hits in the CDC. However, due to the proximity of the SVD to the beam, the algorithm has to cope with a high occupancy from beam-induced background. The original idea and implementation for this algorithm, called the VXD Track Finder (VXDTF), is described in [21]. Further improvements of the algorithm, which are partly described in [22] and [23], led to its second version, VXDTF2.

The input to the VXDTF2 algorithm is the set of the three dimensional space points created in the pre-processing steps described in Section 5 from the SVD measurements[2]. The VXDTF2 algorithm consists of three steps. In the first step, graphs of related space points are created using geometrical information. A map, called a Sector Map [24], containing the geometrical relations between different regions of the silicon detector as well as additional selection criteria, supplies the necessary input for this step. The prepared graphs are then evaluated in the second step by a cellular automaton which yields a set of paths. As third step, the final set of SVD track candidates is chosen by selecting the best paths.

Sector Map

The Sector Map is a data structure that holds information about how space points in different regions of the detector can be related by tracks. To cope with the high number of possible combinations of space points, the

[2] It is also possible to use the three dimensional measurements provided by the PXD. However, due to the combined effect of its proximity to the beam line and the lower readout frequency of its sensors, the PXD is subject to a substantially higher occupancy from beam-background induced hits. Omitting the PXD measurements from the track-finding process simplifies the combinatorial problem and leads to a purer set of track candidates produced by the standalone algorithm. The task of evaluating the additional information available in the measurements of the PXD is passed on to a dedicated algorithm described in Section 9.

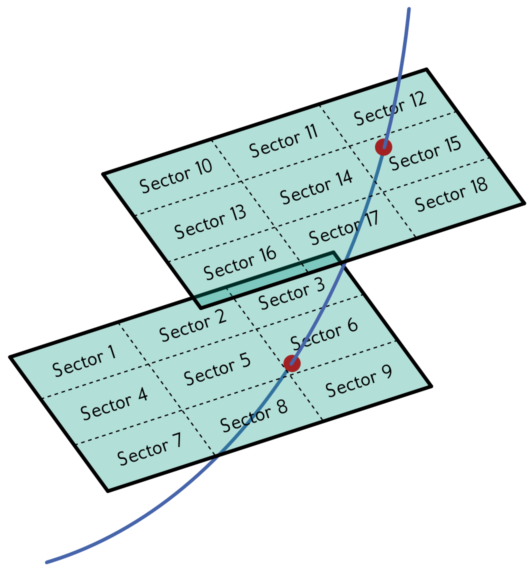

Figure 11: Illustration of the subdivision of two sensors into nine sectors each and of the relation between the sectors 6 and 15 which are traversed by the same track.

sectors on sensor concept — originally proposed in [24] — is used for track finding with the SVD. This concept consists in subdividing each sensor into smaller sections, called *sectors*. The default setup is a division of each sensor element into three parts along its width and three parts along its length, resulting in nine sectors per sensor. It is possible to adjust the number of sectors individually for each sensor, which allows the granularity to be adapted to changing detector conditions during the run time of Belle II. This representation of the SVD geometry allows to define directed relations between sectors of the detector which commonly contain measurements of the same track, as illustrated in Figure 11, where the sectors 6 and 15 are related due to the track traversing both sectors. The direction of the relations is defined by the order in which the sectors are traversed. A mapping among sectors defined by these relations allows for a significantly reduced number of combinations of space points as input to the algorithm.

In addition to the relations between sectors of the detector geometry, the Sector Map holds selection criteria to be fulfilled by combinations of space points on related sectors. These criteria are called *filters* and are defined for pairs of two as well as triplets of three space points. They provide a way to reject background hits and thereby a further reduction of the space point combinations to be evaluated per event. Each filter is a function which calculates a specific quantity, called *filter variable*, for a given space point pair or triplet, and checks if the result is within a given validity range. The validity range depends on the filter variable and on the sectors. It is stored for each individual sector combination alongside the respective relation between sectors in the Sector Map. Filter variables are mostly geometrical quantities derived from the spatial information of the space points, or are calculated from the precise timing information provided by the SVD.

The variables calculated for filters for space point pairs are simple quantities such as distances between the two space points (in one-, two-, and three-dimensions), angles in ϕ- and θ-direction defined by the two space points, or the difference in their detection times. An illustration of the combined application of a selection of such filters is depicted in Figure 12.

More complicated quantities can be evaluated for filters based on the combination of three space points. These include for example the angle enclosed by the two segments defined by the three space points, or the position of the center as well as the radius of the circle defined by the three space points in the x-y plane. Space point triplet filters based on the SVD timing information are also employed. As the SVD is composed of only four layers and a triplet of space points already provides enough degrees of freedom to unambiguously define a helix trajectory of a charged particle in a magnetic field, further filters for combinations of four and more space points are not considered.

The directed relations between sectors as well as the filter selection criteria are obtained via a training process based on Monte Carlo events. For this purpose, a dedicated sample of representative $\Upsilon(4S)$ events is generated. Track candidates are selected from this training sample using the MC track finder (see Section 4). Additional samples of high-momentum tracks such as simulated high-momentum muon events or simulated Bhabha events can be incorporated into the training process as $\Upsilon(4S)$ events don't typically produce such tracks. Based on this set of tracks, directed relations between pairs of sectors are obtained for all pairs of sectors which have been traversed subsequently by at least one track. The selection ranges for the filters are defined by the minima and the maxima or by quantiles of the distributions of the respective filter variables as observed during the training for each individual sector combination.

This training process allows the Sector Map to learn the geometry of the SVD setup. Hence, it can adapt to changing detector conditions like defects

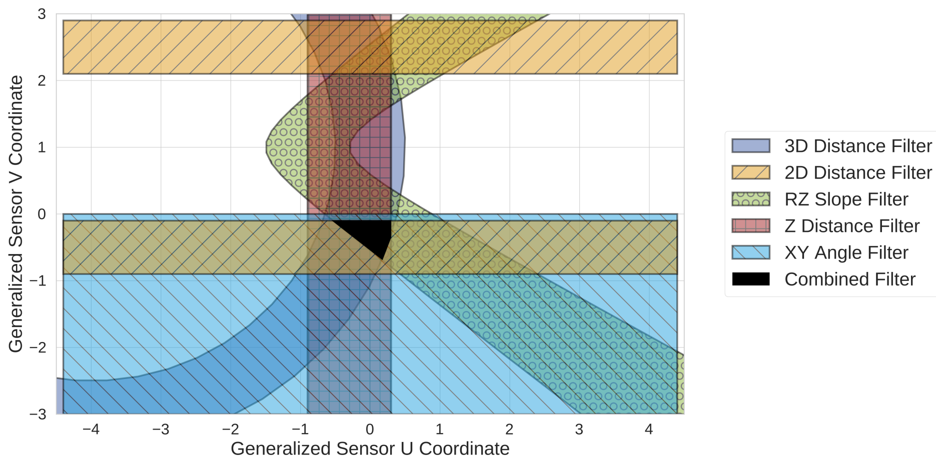

Figure 12: Illustration of the selection power of the combination of several Sector Map filters. The illustration depicts a sensor plane and the areas selected by several Sector Map filters calculated for a given space point on another sensor layer. Only space points within these areas are allowed to be combined with the given space point on the other layer. The combination of all filters reduces the allowed region to the area shown in black.

on the sensors or even the loss of complete sensors or layers, as long as these defects are modeled by the simulation. The Sector Maps produced in this manner are stored in the database of Belle II, which allows defining different Sector Maps for different run conditions.

Track Finding Algorithm

To address the high combinatorics during the process of building track candidates from the space points provided by the SVD, the Sector Map filters are used. A first directed graph, called *sector network*, is build for an event with all active sectors (sectors on which hits are detected) as nodes. The edges connecting the nodes of the sector network are given by the directed relations stored for the respective active sectors in the Sector Map.

The next step comprises the creation of a second directed graph with the space points on the active sectors of an event as nodes. The edges of this *space point network* are given by the edges of the sector network and connect space points in pairs if they pass the criteria of the Sector Map filters for space point pairs. Their directions are defined by the respective edges in the sector network. The resulting space point pairs are called *segments*. Next, pairs of such segments that have a space point in common are combined into triplets of space points, creating a third directed graph with segments as nodes, which is therefore referred to as a *segment network*. Again, the criteria given by the respective filters for space point triplets provided by the Sector Map are considered during the combination of the segments to triplets.

All paths given by the edges of the segment network with a minimal length of three space points are considered as track candidates without further restrictions from the Sector Map. A cellular automaton is used to gather the longest paths in the graph, beginning with the nodes on the outermost layers as these are least occupied by beam-background induced hits.

Gathered this way, the track candidates in an event may share SVD clusters or even space points. At this stage of the algorithm, the number of fake and clone track candidates make up more than half of the track candidates and are directly related to the overlaps among a set of tracks. To reduce the fraction of fakes and clones, the final set of track candidates for an event is required to be composed only of candidates which do not share any SVD clusters among them. As roughly 5% of all tracks in a normal $\Upsilon(4S)$ event share at least one SVD cluster with another track, this introduces a small loss of less than 1% in finding efficiency to the benefit of an increase in purity for the final set of track candidates by a factor of roughly two. Any cases where two or more track candidates share common hits are resolved based on a rating of all track candidates, followed by a greedy local selection as explained in the following paragraphs.

For the rating of each track candidate, a quality indicator determined from the goodness of a fast fit to the candidate is employed. The fit method used for this objective is an adapted version of the Triplet Fit introduced in [25]. This method is chosen because it takes into account the multiple scattering relevant for the tracks left by low momentum particles of interest to the VXD standalone track finding. The Triplet Fit is applied to each path supplied by the cellular automaton, as well as their subpaths obtained by excluding one or multiple space points from the original path. The latter allows for the exclusion of misattributed space points and results in a track candidate with higher purity. Furthermore, the inclusion of the subpaths can lead to a recovery of the efficiency loss due to overlapping true tracks. When creating the subpaths, the rule of a minimal length of three space points is still obeyed.

The Triplet Fit yields a χ^2 value for each track candidate by combining fits to all space point triplets contained within a candidate under consideration of the effect of multiple scattering. For this estimate the average radiation length of the SVD sensor material as reported in [3], as well as a first approximation of the entrance angle of the particle with respect to the sensor plane are taken into account. The p-value is calculated for each track candidate from its χ^2 value and degrees of freedom and used as a quality indicator. Based on these quality indicators, the final set of non-overlapping track candidates is chosen via a greedy selection which takes the candidate with the highest quality indicator among the ones competing for a space point.

Optionally, a multivariate method can be applied which combines the acquired quality indicator with further features, such as a particle momentum estimate, the number of space points and properties of the involved SVD clusters. This approach can yield a quality indicator with an enhanced performance. For this purpose a FastBDT is trained on Monte Carlo events obtained by applying the candidate creation steps up to the point of the Triplet Fit. In the resulting training sample, the track candidates with a purity of 100% are marked as truth target. Therefore, a FastBDT trained in this manner has learned to identify track candidates with a high purity. Enabling this auxiliary multivariate analysis method for the overlap removal increases the achieved track finding efficiency, albeit with a significant drop in hit efficiency. This option is not therefore used as a basis for the resolution of overlaps, but used to produce a refined track quality estimate in an additional step. This indicator of the track quality is stored for all tracks of the final set and can later be accessed and used in the event selection of physics analyses.

The algorithm is further optimized as it is found to perform more slowly than acceptable at the HLT for certain peculiar Bhabha events. The slowdown is understood as follows. In rare cases, highly energetic electrons scatter in the material of the final focusing magnets, thereby causing a shower of secondary particles which leave a large number of clusters in a small area of the SVD. This leads to a significant increase in the combinatorics during the candidate creation process that cannot be restricted by the Sector Map filters. To tackle this issue, two additional measures are incorporated into the candidate creation. Firstly, a limit on the number of nodes and edges in the three networks is introduced, as the problematic Bhabha events can mostly be identified based on noticeably high values for these quantities. If the limits determined on Monte Carlo simulations are exceeded, the processing of the event is aborted and the problematic event is marked as such. The limits are chosen so that the desired measurements of the $\Upsilon(4S)$ resonance are not affected. Secondly, during the path-collection step an additional selection procedure based on the segment network is applied, which evaluates overlaps already in this graph. All paths associated with a given segment are grouped together and evaluated with the Triplet Fit. Only a fixed number of best candidates from each group is considered for further processing steps. This early candidate reduction based on space point pair overlaps imposes an additional limit on the combinatorics for problematic events which slip through the aforementioned limits. By means of these additional selection steps, the problematic events can be handled by the VXDTF2 and the run time limits imposed by the requirements of the HLT are satisfied.

9. Combinatorial Kalman Filter

The combinatorial Kalman filter (CKF) is widely used in tracking in high-energy physics experiments [10, 11, 12, 13]. One of the advantages of the method is that it produces tracks with high purity also in environments of high hit densities. The CKF is an iterative local algorithm and was first described in [26]. Starting with a seed estimation of the track parameters with uncertainties, the track is extrapolated with the Runge-Kutta-Nyström method [27] into the detector volume. Hereby, non-uniform magnetic fields are included in the numerical solution of the equation of motion. A correction of the energy and the uncertainties due to energy loss and multiple scattering can optionally be included. After the extrapolation, possible hit candidates are determined based on the current position and uncertainties of the track candidate. The next hit candidate is added to the track and the procedure is repeated. If there are multiple mutually exclusive next-hit candidates, the whole track candidate is duplicated and subsequently treated as two tracks. In the end, the final track candidate is selected according to different quality criteria.

As a first step, the track candidates found by the CDC track finding algorithm are used as seeds to attach SVD space points. Hereby, low momentum tracks can have both start and end points in the inner layers of the CDC, so both points can be used as a possible seed to account for wrongly assigned charges in the CDC. The CDC seeds are fitted using a DAF algorithm assuming a pion mass hypothesis. These seeds are iteratively extrapolated to the SVD sensors and SVD space points are attached. Material effects are disregarded at this stage to increase the processing speed. Due to the complex detector layout, the different use cases, and the complex input data from the CDC track finding algorithm, the filter decisions in the CKF are taken by a FastBDT trained on simulated events including the beam-induced background. Variables such as the distance between extrapolated and measured hit position as well as the calculated χ^2 of the hit are taken into account. They are enriched with information about the track candidate, for example the number of attached space points or the estimated transverse momentum. The number of track candidates that are processed in parallel is restricted to ten to keep the computational effort on a manageable level. After a final candidate selection based on FastBDT using full-track information such as the summed χ^2 and the number of missing layers, the combined CDC-SVD track is refitted using another DAF with a full material effect handling.

Due to hit inefficiencies of the CDC algorithm, especially for the stereo layers, the track resolution can be extremely poor when $|z_0|$ is above 1 cm. Therefore, it is not possible to attach SVD space points reliably to every reconstructed CDC track. To solve this issue and to find additional low-momentum tracks, the VXDTF2 described in Section 8 is applied to the set of remaining space points in the SVD. The merging of additional SVD candidates with unmerged CDC tracks is performed by a second pass of the CKF. These unpaired CDC tracks are used as seeds and only space points found by the VXDTF2 are allowed as input. Compared to the first pass of the CKF with all space points, simpler filters are applied during the processing due to the high purity of the VXDTF2 algorithm.

All reconstructed CDC-SVD tracks are then used to extract regions of interest in the PXD during the online reconstruction. In the offline reconstruction, the PXD clusters collected in these regions of interest are used as input to the last application of the CKF and are attached to their combined CDC-SVD tracks. The implementation is based on the same general principles as the SVD CKF. It uses the same FastBDT filters which are now applied to the PXD clusters. An additional input in the BDT classification is given by the position and the shape of the PXD clusters.

The precise positions of the PXD clusters in the tracks improve the resolution on the spatial track parameters d_0 and z_0 by a factor of two and more. The efficiency of attaching SVD (PXD) hits is over 85% (89%). The purity of the attached SVD or PXD hits is above 98% and 96%, respectively.

10. Performance of the Track Finding

The following section describes the performance of the tracking algorithms presented in this paper. The performance is evaluated on an independent set of simulated $\Upsilon(4S)$ events, including beam-induced background simulated for the anticipated full instantaneous luminosity of $8 \times 10^{35}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. In the simulation, a detector setup with nominal positions is used. The results of the reconstruction are compared to the respective MC tracks. Quantities such as the purity or efficiency are calculated on this sample using the definitions from Section 4. All quoted uncertainties on these quantities are calculated using the method of bootstrapping [28].

Figure 13 shows the track-finding efficiency for different simulated transverse momenta and different levels of beam-induced background relative to the anticipated level. A distinction between final-state particles stemming from the primary e^+e^- interaction and decays of short-lived particles, produced by event generators (*primaries*), and all final-state particles including those produced by Geant4 during the travel through the detector (*secondaries*) is made. Most of the analyses rely only on the former, whereas the latter can give valuable additional information for decays in flight or for particle identification. Comparing with the momentum spectrum shown in Figure 3, the efficiency for most of the charged particles expected at typical Belle II collisions is higher than 93% for up to two times the expected beam background. Tracks with transverse momenta below $100\,\mathrm{MeV/c}$ impose complex problems to the track finding due to the small number of hits, high multiple scattering and the high level of background in the innermost layers. As a result, the efficiency decreases. The difference between the non-background and the expected beam background is small.

In Figure 14a the finding efficiency on primaries is compared for different simulated particle types. Due to the different interaction of electrons with the material, their trajectories are more likely to differ from the nominal helical path, making their reconstruction more challenging. However, the Belle II algorithms are able to achieve high efficiencies for every shown particle type for up to twice the expected beam background level.

After the final fit with the DAF provided by GENFIT2 the tracks are extrapolated to the POCA to the origin to extract their helix parameters. In the following, only the results for the pion hypothesis are shown, as most of the produced charged final-state particles are pions.

As a first result, Figure 14b shows the finding efficiency calculated only

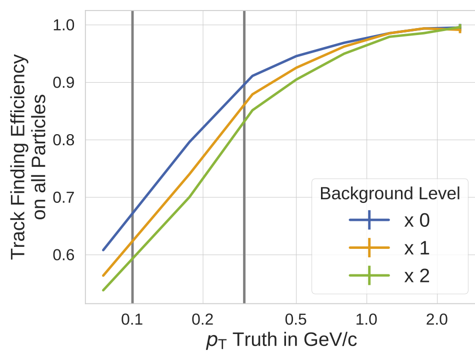

(a) On all particles.

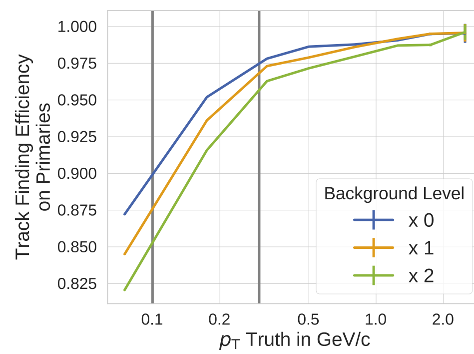

(b) Only on primaries.

Figure 13: Track finding efficiency calculated for simulated $\Upsilon(4S)$ events with different levels of beam-induced background relative to the expected level. Figure 13a is calculated on all trackable simulated particles, whereas Figure 13b only takes into account trackable particles from the primary interaction. The gray vertical lines indicate the typical transverse momentum of particles only trackable in the VXD (below left line) and with high efficiency in the CDC (above right line).

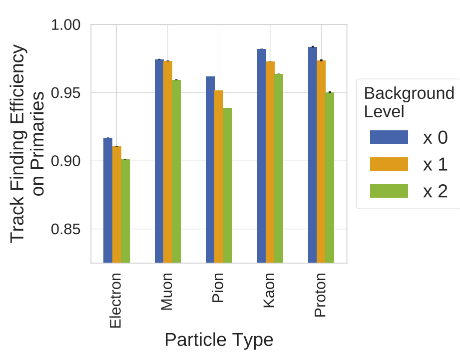

(a) Track finding efficiency extracted for the most important particle types present in the primary $B\overline{B}$-decays at Belle II dependent on the beam-induced background level. The overall finding efficiency is dominated by pions.

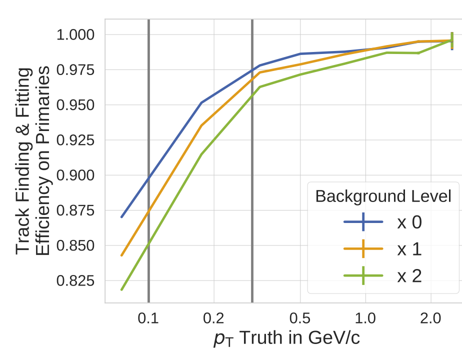

(b) Calculated finding efficiency for those tracks, which are also successfully fitted by the DAF algorithm dependent on the beam-background level. The difference with respect to Figure 13b is negligible.

Figure 14: Finding efficiency by particle type and combined finding and fitting efficiency.

with those tracks, where the GENFIT2 fit converged and the extrapolation succeeded. As expected, the difference to Figure 13b is negligible demonstrating that the track finding algorithms deliver high-quality track candidates to the fitting algorithm.

The helix parameters of the tracks at the POCA can then be compared to the MC truth values. The resolution r_x is calculated as the 68% coverage of the the residual x between reconstructed and truth value given as

$$r_x = P_{68\,\%}\left(|x - P_{50\,\%}(x)|\right) \ ,$$

where P_q calculates the q-th percentile of a distribution. For a Gaussian distribution, the 68% coverage and the standard deviation agree. For non-Gaussian distributions the coverage is more robust against outliers. As only the results calculated with the pion hypothesis are shown, only true pions from $\Upsilon(4S)$ decays are taken into account for this study.

In Figure 15, the resolution as a function of the truth transverse momentum is shown for the helix parameters d_0 and z_0, and for p_T (the transverse momentum). Both of the spatial parameters, d_0 and z_0, are mainly influenced by the precise PXD measurements. Due to the application of the CKF in the PXD and the combination of the VXDTF2 and the CKF for the SVD, a high precision (which is almost independent of the background level) is achieved. The resolution of the extracted transverse momentum follows expectation: as smaller momenta are more strongly influenced by multiple scattering and a smaller number of measurable hits in the detector, the resolution decreases with smaller transverse momenta.

Tracking is one of the most complex tasks in the reconstruction. It therefore requires a large fraction of the processing time allocated for the online reconstruction on the HLT. In Figure 16 the processing time of different components of the online reconstruction performed on one of the HLT worker nodes is shown. Due to the higher number of tracks in $\Upsilon(4S)$ events, tracking takes longer in this category. The track fitting, vertexing, and the track-based collision time (T_0) extraction are heavily influenced by the handling of the detector geometry in the software. Different techniques are planned to further optimize the time spent in the geometry navigation. This is expected to decrease the total processing time significantly which would allow to introduce additional higher level algorithms for the HLT decision. However, even with the large contribution to the total processing time from tracking, a stable reconstruction on the HLT has been achieved.

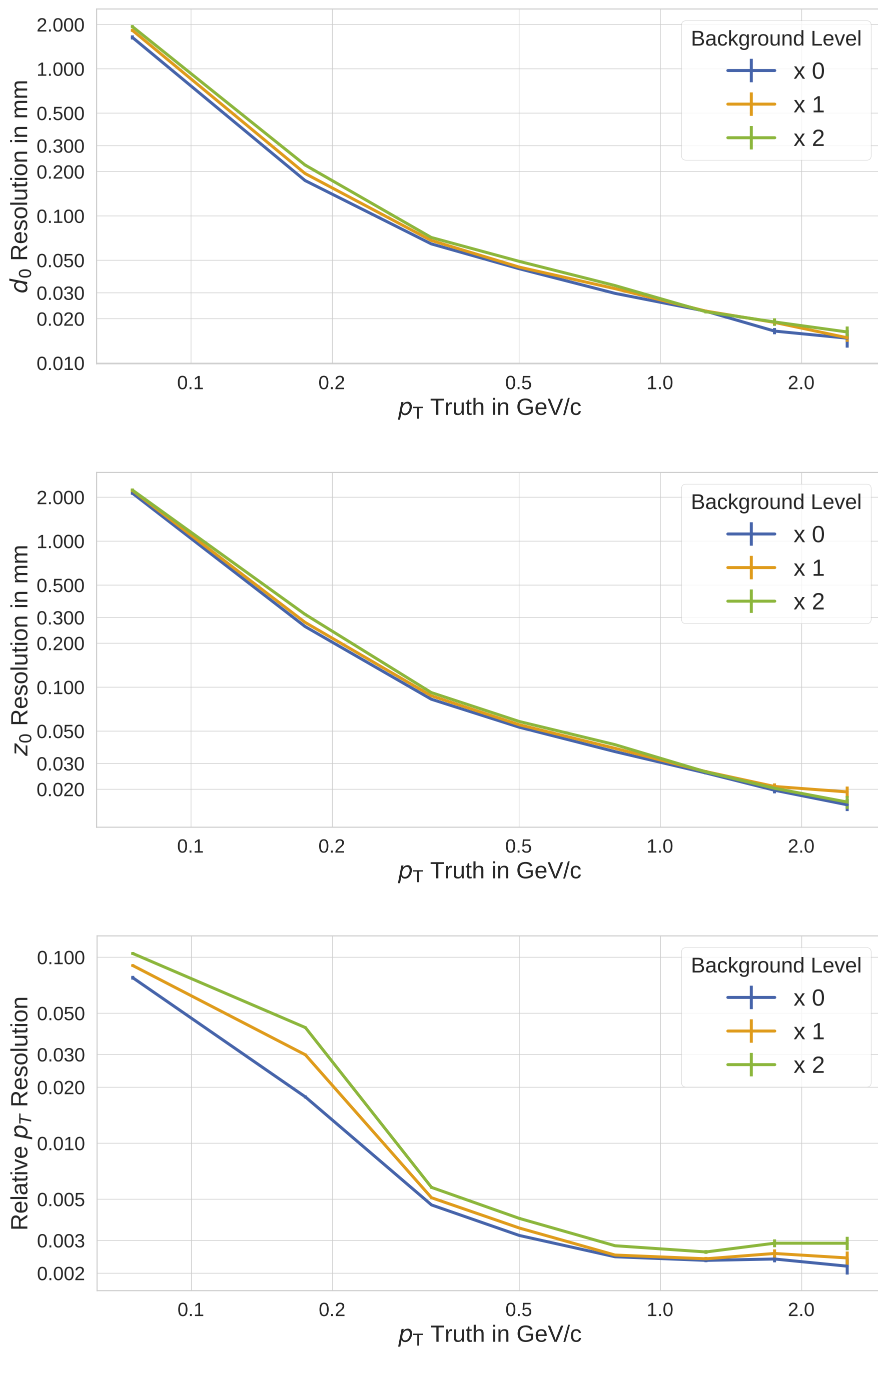

Figure 15: Resolutions for typical simulated $\Upsilon(4S)$ events with different levels of beam-induced background. As most of the simulated final state particles are pions, only the results of this fit hypothesis are shown.

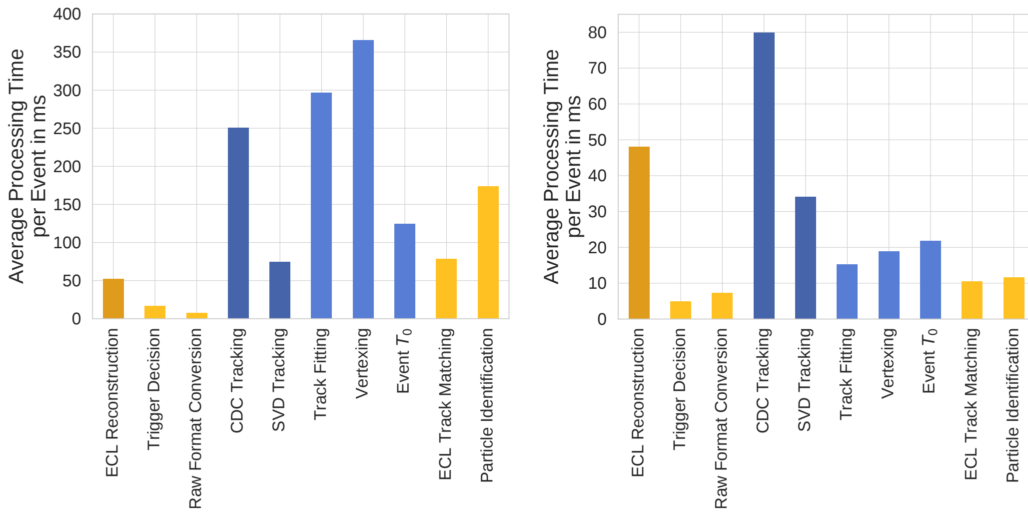

(a) For $\mathrm{B}^+\mathrm{B}^-$ events.

(b) For $\mathrm{e}^-\mathrm{e}^+$ events.

Figure 16: Processing time of a standard online reconstruction performed on the anticipated HLT worker nodes in single-processing mode for different event types. The components marked in dark (light) blue can be related to track finding (track fitting and other tracking-related tasks). The abbreviation *ECL* refers to the electromagnetic calorimeter of the Belle II experiment.

11. Summary

The Belle II track-reconstruction software consists of multiple independent algorithms to process the measurements of each tracking detector and integrate all available information into one final set of tracks available for physics analyses. This allows the use of different algorithms whose properties are especially suited for the three different tracking detectors. Ultimately, the software provides a set of tracks based on the measurements of all tracking detectors, thereby alleviating the complex task of track combination which would otherwise be forced on physics analyses.

One challenge in this approach is to perform the combination of the tracks reconstructed in each of the tracking detectors without increasing the fake and clone rate. The best method to achieve this for the Belle II experiment turned out to be a Combinatorial Kalman Filter to link hits and tracks across detector boundaries.

The upgraded SuperKEKB collider will have much higher beam background radiation and the newly developed tracking algorithms are designed to address this. Here, multiple methods like early background hit filtering using multivariate methods and a fine-grained candidate selection using the Sector Map concept are used. The studies with simulated background environments at design luminosity of the accelerator show that tracking performance remains adequate for the expected background rate.

The Belle II track-reconstruction software has been extensively studied and used for the reconstruction of simulated events. In the years 2018 and 2019, the software was also employed during the first data taking of collision data and performed well for the commissioning of the Belle II detector, studies of the background rates, and first physics results.

12. Acknowledgments

We would like to acknowledge the contributions of Karol Adamczyk, Bozek Andrzej, Kirill Chilikin, David Dossett, Torben Ferber, Stefan Ferstl, Renu Garg, Hadrien Grasland, Yinghui Guan, Yoshihito Iwasaki, Peter Kodys, Ilya Komarov, Jo-Frederik Krohn, Miriam Künzel, Stefano Lacaprara, Klemens Lautenbach, Frank Meier, Moritz Nadler, Michael De Nuccio, Benjamin Oberhof, Hitoshi Ozaki, Johannes Rauch, Giuliana Rizzo, Jonas Roetter, Michael Schnell, Alexei Sibidanov, Marko Staric, Jacek Stypula, Maiko Takahashi, Umberto Tamponi, Viktor Trusov, Makoto Uchida,

Jonas Wagner, Ian James Watson, Jarek Wiechczynski and Michael Ziegler to the development of the Belle II tracking software.

[1] Y. Ohnishi, T. Abe, T. Adachi, et al., Accelerator design at SuperKEKB, Progress of Theoretical and Experimental Physics 2013 (3) (2013) 3A011–0. doi:10.1093/ptep/pts083.
URL https://academic.oup.com/ptep/article-lookup/doi/10.1093/ptep/pts083

[2] J. Brodzicka, T. Browder, P. Chang, et al., Physics achievements from the Belle experiment, Progress of Theoretical and Experimental Physics 2012 (1) (2012) 4D001–0. arXiv:1212.5342, doi:10.1093/ptep/pts072.
URL https://academic.oup.com/ptep/article-lookup/doi/10.1093/ptep/pts072

[3] T. Abe, I. Adachi, K. Adamczyk, et al., Belle II Technical Design ReportarXiv:1011.0352.
URL http://arxiv.org/abs/1011.0352

[4] T. Kuhr, C. Pulvermacher, M. Ritter, T. Hauth, N. Braun, The Belle II Core Software, Computing and Software for Big Science 3 (1) (2019) 1. arXiv:1809.04299, doi:10.1007/s41781-018-0017-9.
URL http://link.springer.com/10.1007/s41781-018-0017-9

[5] M. Gardner, Mathematical Games – The fantastic combinations of John Conway's new solitair game life, Scientific American 223(4) (1970) 120–123. doi:10.1038/scientificamerican1070-120.

[6] A. Glazov, I. Kisel, E. Konotopskaya, G. Ososkov, Filtering tracks in discrete detectors using a cellular automaton, Nucl. Instrum. Meth. A329 (1993) 262–268.

[7] I. Abt, I. Kisel, S. Masciocchi, D. Emelyanov, CATS: A cellular automaton for tracking in silicon for the HERA-B vertex detector, Nucl. Instrum. Meth. A489 (2002) 389–405. doi:10.1016/S0168-9002(02)00790-8.

[8] D. Funke, T. Hauth, V. Innocente, et al., Parallel track reconstruction in CMS using the cellular automaton approach, J. Phys. Conf. Ser. 513 (2014) 052010. doi:10.1088/1742-6596/513/5/052010.

[9] T. Alexopoulos, M. Bachtis, E. Gazis, G. Tsipolitis, Implementation of the Legendre Transform for track segment reconstruction in drift tube chambers, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 592 (3) (2008) 456–462. doi:10.1016/j.nima.2008.04.038.
URL https://www.sciencedirect.com/science/article/pii/S0168900208005780https://linkinghub.elsevier.com/retrieve/pii/S0168900208005780

[10] R. Mankel, A. Spiridonov, The Concurrent Track Evolution algorithm: extension for track finding in the inhomogeneous magnetic field of the HERA-B spectrometer, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 426 (2-3) (1999) 268–282. doi:10.1016/S0168-9002(99)00013-3.
URL http://linkinghub.elsevier.com/retrieve/pii/S0168900299000133

[11] R. Mankel, Pattern recognition and event reconstruction in particle physics experiments, Reports on Progress in Physics 67 (4) (2004) 553–622. arXiv:0402039, doi:10.1088/0034-4885/67/4/R03.
URL http://arxiv.org/abs/physics/0402039

[12] CMS Collaboration, Description and performance of track and primary-vertex reconstruction with the CMS tracker, Journal of Instrumentation 9 (10) (2014) P10009–P10009. arXiv:1405.6569, doi:10.1088/1748-0221/9/10/P10009.
URL http://arxiv.org/abs/1405.6569

[13] ATLAS Collaboration, Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2, The European Physical Journal C 77 (10) (2017) 673. arXiv:1704.07983, doi:10.1140/epjc/s10052-017-5225-7.
URL http://arxiv.org/abs/1704.07983

[14] T. Keck, FastBDT: A Speed-Optimized Multivariate Classification Algorithm for the Belle II Experiment, Computing and Software for Big Science 1 (1) (2017) 2. arXiv:1609.06119, doi:10.1007/s41781-017-0002-8.
URL http://arxiv.org/abs/1609.06119

[15] J. Kemmer, G. Lutz, New Detector Concepts, Nuclear Instruments and Methods in Physics Research 253 (1987) 365–377.
URL https://ac.els-cdn.com/0168900287905183/1-s2.0-0168900287905183-main.pdf?_tid=9c6da4c9-b91f-464b-ae11-89016a3e2531&acdnat=1524476602_881ecbe6d098bc1af5882f526389be5a

[16] P. M. Lewis, et al., First Measurements of Beam Backgrounds at SuperKEKB, Nucl. Instrum. Meth. A914 (2019) 69–144. arXiv:1802.01366, doi:10.1016/j.nima.2018.05.071.

[17] S. Agostinelli, J. Allison, K. Amako, et al., Geant4a simulation toolkit, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 506 (3) (2003) 250–303. doi:10.1016/S0168-9002(03)01368-8.
URL https://www.sciencedirect.com/science/article/pii/S0168900203013688https://ac.els-cdn.com/S0168900203013688/1-s2.0-S0168900203013688-main.pdf?_tid=2ab2887a-7a2b-4813-9966-25cc44845930&acdnat=1524482129_7273334d05ad6f2c66d093ae3c67b132

[18] T. Bilka, N. Braun, T. Hauth, et al., Implementation of GENFIT2 as an experiment independent track-fitting frameworkarXiv:1902.04405.
URL http://arxiv.org/abs/1902.04405

[19] V. Karimäki, Effective circle fitting for particle trajectories, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 305 (1) (1991) 187–191. doi:10.1016/0168-9002(91)90533-V.
URL https://linkinghub.elsevier.com/retrieve/pii/016890029190533V

[20] A. Strandlie, J. Wroldsen, R. Frühwirth, B. Lillekjendlie, Particle tracks fitted on the Riemann sphere, Computer Physics Communications 131 (1-2) (2000) 95–108. doi:10.1016/S0010-4655(00)00086-2.
URL https://linkinghub.elsevier.com/retrieve/pii/S0010465500000862

[21] J. Lettenbichler, Real-time Pattern Recognition in the Central Tracking Detector of the Belle II Experiment, Ph.D. thesis, Technische Universität

Wien (2016).
URL http://www.ub.tuwien.ac.at

[22] J. Wagner, Track Finding with the Silicon Strip Detector of the Belle II Experiment, Master's thesis, Karlsruhe Institute of Technology (2017).
URL https://ekp-invenio.physik.uni-karlsruhe.de/record/48930/files/EKP-2017-00057.pdf

[23] S. Racs, Track Quality Estimation and Defect Sensor Studies for the Silicon Vertex Detector of the Belle II-Experiment, Master's thesis, Karlsruhe Institute of Technology, Karlsruhe (2018).
URL https://ekp-invenio.physik.uni-karlsruhe.de/record/49052

[24] R. Frühwirth, R. Glattauer, J. Lettenbichler, W. Mitaroff, M. Nadler, Track finding in silicon trackers with a small number of layers, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 732 (2013) 95–98. doi:10.1016/j.nima.2013.06.035.
URL https://linkinghub.elsevier.com/retrieve/pii/S0168900213008607

[25] N. Berger, A. Kozlinskiy, M. Kiehn, A. Schöning, A new three-dimensional track fit with multiple scattering, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment 844 (2017) 135–140. doi:10.1016/j.nima.2016.11.012.
URL https://linkinghub.elsevier.com/retrieve/pii/S016890021631138X

[26] P. Billoir, Progressive track recognition with a Kalman-like fitting procedure, Computer Physics Communications 57 (1-3) (1989) 390–394. doi:10.1016/0010-4655(89)90249-X.
URL http://linkinghub.elsevier.com/retrieve/pii/001046558990249Xhttps://linkinghub.elsevier.com/retrieve/pii/001046558990249X

[27] E. Lund, L. Bugge, I. Gavrilenko, A. Strandlie, Track parameter propagation through the application of a new adaptive Runge-Kutta-Nyström method in the ATLAS experiment, Journal of Instrumentation 4 (04)

(2009) P04001–P04001. doi:10.1088/1748-0221/4/04/P04001. URL http://stacks.iop.org/1748-0221/4/i=04/a=P04001?key=crossref.4db573cdf47b35d9171ab976705899ed

[28] B. Efron, Bootstrap Methods: Another Look at the Jackknife, The Annals of Statistics 7 (1) (1979) 1–26. doi:10.1214/aos/1176344552. URL http://projecteuclid.org/euclid.aos/1176344552