Proceedings of CTD 2020
PROC-CTD2020-40
July 20, 2020

Performance of Belle II tracking on collision data

Simon Kurz

On behalf of the Belle II Tracking Group,
Deutsches Elektronen Synchrotron (DESY)
Germany

ABSTRACT

The tracking system of Belle II consists of a silicon vertex detector (VXD) and a cylindrical drift chamber (CDC), both operating in a magnetic field created by the main solenoid of 1.5 T and final focusing magnets. The tracking algorithms employed at Belle II are based on a standalone reconstruction in SVD and CDC as well as on a combination of the two approaches, they employ a number of machine learning methods for optimal performance. The tracking reconstruction is tested on the collision data collected in 2018 and 2019. The first experience with data introduced additional challenges which are mitigated with the introduction of new algorithms such as track finding seeded by calorimeter clusters and CDC cross-talk filtering.

PRESENTED AT

Connecting the Dots Workshop (CTD 2020)
April 20-30, 2020

1 Introduction

The focus of the Belle II experiment is the study of CP-violation and rare decays of B mesons. The experiment is located at the SuperKEKB collider, which provides e^+e^- collision at a center of mass energy of 10.58 GeV. This energy is just above the mass of the $\Upsilon(4S)$ resonance, which almost exclusively decays to a pair of B mesons. Furthermore, the asymmetric beam energies of 4 and 7 GeV, respectively, lead to a forward boost of the $B\bar{B}$ system so that lifetime of the B mesons can be studied.

The peak luminosity of the SuperKEKB collider is 40 times higher than the one of the predecessor KEKB, achieving up to $L = 8 \times 10^{35}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. This is achieved by a significant reduction of the beam size ('nano beam' scheme) and increased beam currents. Over the full lifetime, SuperKEKB is expected to deliver an integrated luminosity of $50\,\mathrm{ab}^{-1}$. So far about $30\,\mathrm{fb}^{-1}$ have been recorded by the Belle II experiment. The current tasks SuperKEKB is facing is to increase the instantaneous luminosity while the beam background levels have to be reduced.

The design of the Belle II experiment shown in Figure 1 is motivated by this high instantaneous luminosity and the challenges that come with it. The tracking system is made up of three tracking detectors: the pixel detector (PXD) and the strip vertex detector (SVD), which are commonly summarized as the vertex detector (VXD), as well as the central drift chamber (CDC). The PXD is installed right outside the beam pipe and consists of two layers, based on the DEPFET Pixel Technology, but currently only the first layer and four modules of the second layer are installed. One of the advantages of this technology is a very low material budget of about 0.2% of a radiation length. The PXD allows for a very precise measurement of the impact parameter of tracks as the inner layer is installed at a radius of just about 1.4 cm. The SVD encloses the PXD and consists of four layers of double-sided silicon strip detectors, which corresponds to a thickness of about 0.7% radiation lengths, and it has an excellent hit time resolution of about 3 ns. This allows the SVD to provide robust tracking despite the background levels, also for standalone tracking. The outermost tracking detector is the CDC, which consists of 56 layers of wires grouped in alternating superlayers of axial and stereo wires. Stereo wires are skewed by an angle between 45.4 and 74 mrad in the positive and negative direction with respect to the axis of the beam and are necessary to determine the longitudinal position of tracks. Most importantly, the CDC has a very high momentum resolution as it covers the major part of the tracking volume up to $l \times r = 2.3\,\mathrm{m} \times 2.2\,\mathrm{m}$.

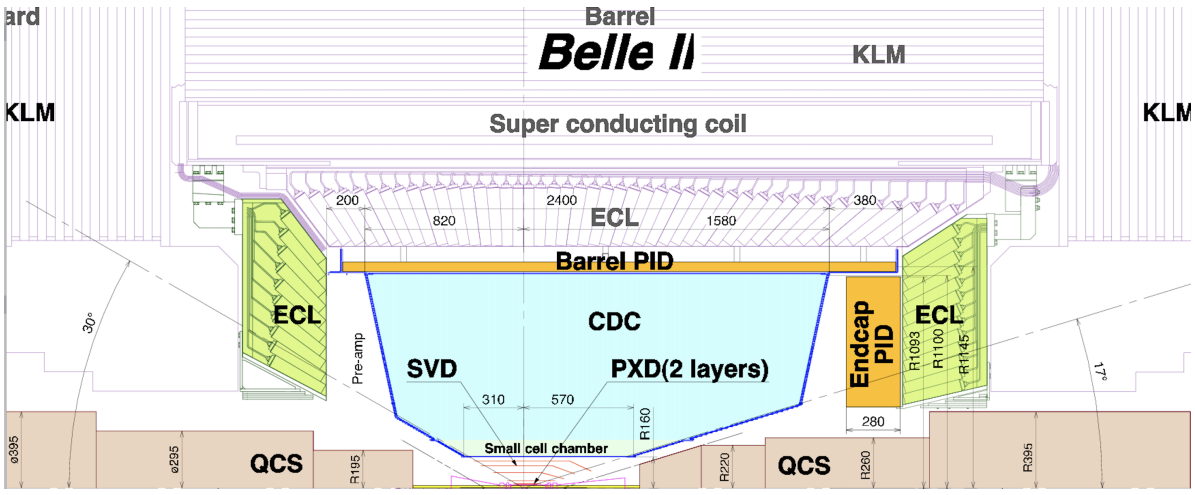

Figure 1: Technical drawing of a section of the Belle II detector [10].

The combination of these three detectors is perfectly suited for the tracking challenges of the $\Upsilon(4S)$ events that are the focus of the studies of the Belle II experiment. A typical $B\bar{B}$ event consists on average of eleven tracks with a soft momentum spectrum, most of which are tracks from charged pions, as depicted in Figure 2.

Furthermore, high backgrounds from the machine are expected, i.e. the number of of background hits

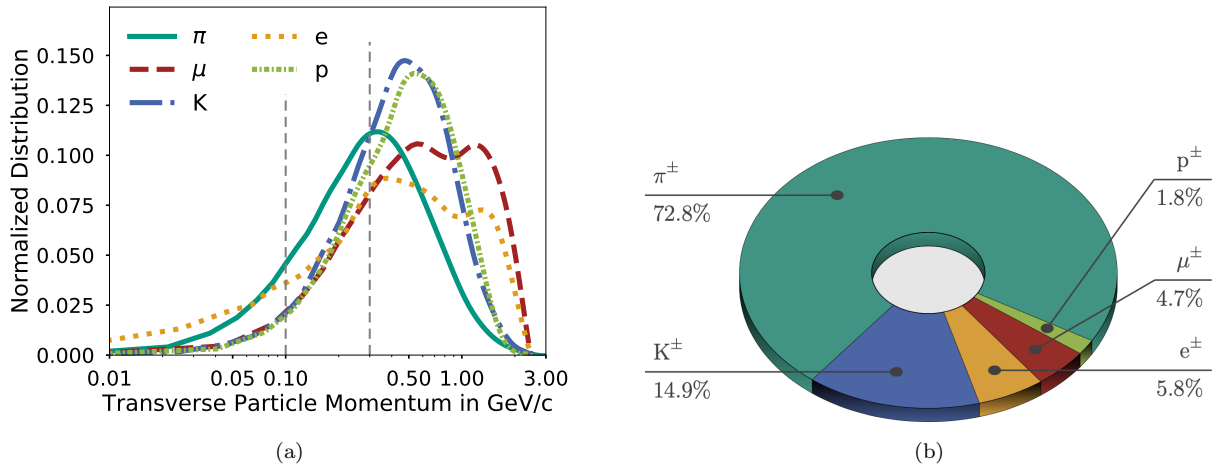

(a) (b)

Figure 2: (a) Normalized distributions (log scale) of transverse momentum of primary charged particles as simulated for $\Upsilon(4S)$ events and (b) relative fraction of charged particle types in these events [1].

is about two orders of magnitude larger than the number of signal hits. On the other side, many analyses require a precise measurement of primary and secondary vertices. Given the forward boost, the flight distance of the $B\bar{B}$ system is in the order of just 10–100 µm. Often a technique called 'Full Event Interpretation' (FEI) [3] is used, where the idea is to reconstruct the full decay chain in the event including all particles that participated. To that end, it is necessary to reconstruct all tracks down to a transverse momentum p_{T} of approximately 50 MeV. Moreover, no fake or duplicate tracks should be found.

2 Background Suppression and Track Finding

Before the actual track finding is performed it is necessary to reject as many background hits as possible, while preserving as many signal hits as possible. For each of the subdetectors designated approaches have been developed, of which the background suppression in the CDC is briefly described here as a showcase. Currently, a classical approach is employed: each CDC hit is required to have a minimum of deposited charge and time over threshold, i.e. short pulses are rejected, too. This simple selection is very efficient in rejecting hits from machine background, as well as so-called cross talk hits. Cross talk describes a feature of the ASIC read-out chips, where neighboring channels can light up if a signal hit is registered. To further reduce the contribution from cross talk hits, which typically have a small charge and peculiar time structure, a designated filter was developed targeting these exact properties.

The track finding is structured in a very modular manner, which is also reflected in the actual implementation. This modularity of the track finding sequence allows for reordering of the independent modules in case it is needed, i.e. it can be adapted according to the background condition, a potential degradation of the detector performance, etc. Furthermore, it is very convenient to add novel track finding approaches. The current sequence is summarized in Figure 3.

The CDC track finding consists of two independent algorithms and is used as the baseline. In the second step, SVD hits are added to the previously found tracks making use of a Combinatorial Kalman Filter (CKF). The third step considers the unused SVD hits and tries to reconstruct tracks out of them. These SVD standalone tracks are then extrapolated to the CDC again via another CKF. Finally all independently found tracks are merged to a single collection that is then extrapolated to the PXD in the final step by a third CKF.

In the following sections, the basic idea behind each of the steps of the track finding sequence is given. More information can be found in the recently submitted Belle II Tracking Paper [1].

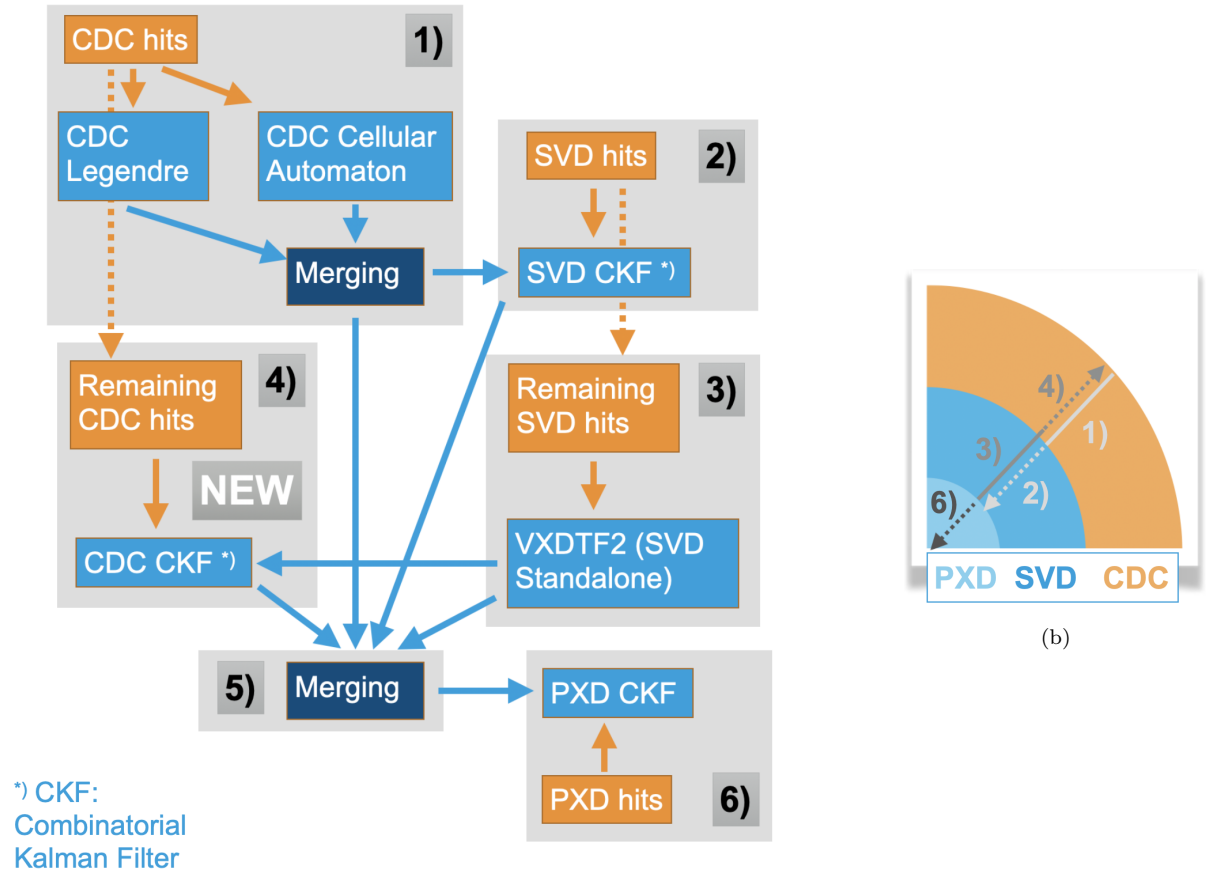

(a) (b)

Figure 3: The current sequence of the Belle II track finding algorithms (a), and the corresponding subdetectors for each one of the steps (b).

2.1 CDC Track Finding

The CDC track finding consists of two algorithms: a global approach based on Legendre Parameters and a local approach using a Cellular Automaton.

The global track finder is performed in three main steps, which are illustrated in Figure 4. These steps only consider hits in axial wires, i.e. the track finding is executed in the $r-\phi$ plane. The hits in the stereo wires are added to the tracks as a last step to reconstruct the z information. The left figure shows the first step of the global track finder in the two dimensional plane so that the track can be seen as a circular trajectory. The hit positions are approximated by so-called drift circles around the positions of the wires, where the radius is derived from the timing information of the individual hits. In the next step a conformal transformation (Hough transform) is applied, which transforms the circular trajectory to a straight line while the drift circles remain circles. This significantly simplifies the task as now a common tangent to a set of circles has to be found. In the final step, Legendre Parameters are used to describe the position (x_0, y_0) and radius R_{dr} of the drift circles as

$$\rho = x_0 \cos\theta + y_0 \sin\theta + R_{\mathrm{dr}}. \tag{1}$$

This equation corresponds to a sinusoidal curve for each of the drift circles as shown in right figure. Tracks are reconstructed by determining points of maximum density in this parameter space. This task is implemented very efficiently as a two dimensional binary search with a dedicated quadtree structure and a 'sliding bin' technique, which recenters the area of interest around the point of maximum density.

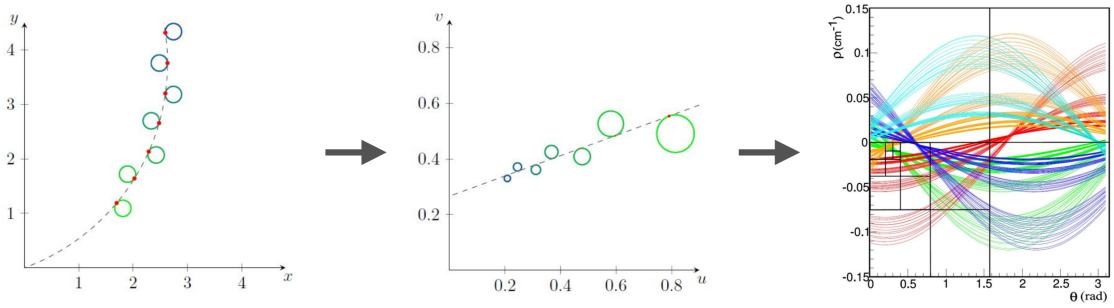

Figure 4: Main steps of the CDC Legendre track finding algorithm. The track (dashed line) is reconstructed using the information from the drift circles around the CDC wires (colored circles). The algorithm uses a conformal transformation (middle) and formulation in Legendre Parameters (right) to find tracks in this global approach.

Compared to the global track finder, which shows a high performance for tracks from the primary vertex, the local track finder focuses on short and displaced tracks. The three main steps of this algorithm are shown in Figure 5. In the first step, triplets of hits in neighboring layers are formed assuming a certain right-left-passage of the track. The second steps connects triplets that share two hits and assigns a weight based on a common fit to these segments. Finally, overlapping segments are combined and a weight is assigned based on a common fit.

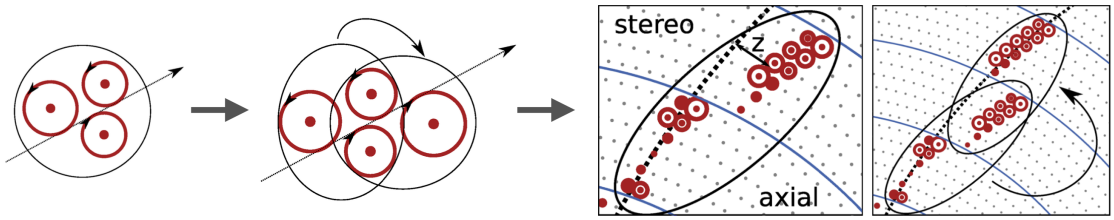

Figure 5: Main steps of the CDC cellular automaton track finding algorithm. This local approach is based on subsequent combinations of track segments.

Currently, the local track finder is not used as a standalone algorithm due to the non-negligible fake rate. Instead, reconstructed segments are added to tracks found by the global track finder using multivariate methods *. As the last step of the CDC track finding there are several further MVA-based filters to improve fake and clone rate, as well as purity and track parameter resolution. Furthermore, an additional quality indicator is currently being studied to reduce the fake rate of the local track finder so that it can be used as a standalone algorithm.

2.2 Extrapolation from CDC to SVD and SVD to PXD

The next step in the standard track reconstruction chain is to extrapolate the previously found CDC tracks inwards to the SVD. This is done via a CKF, which may skip detector layers to account for inefficiencies of the detector. The iterative procedure is as follows: all hits within a certain $\theta \times \phi$ window around the extrapolated track are considered and the best of these candidates are picked using MVA-based filters. Then, new paths are defined based on all selected candidates and the procedure is repeated for all of them. In the end, the best candidate is selected for each CDC track seed.

The implementation of the algorithm is highly configurable as it is completely template based. This has the advantage that filters for the selection of the best hit and path candidates can be modified and tuned

*The multivariate methods are typically based on FastBDT [2] throughout the Belle II software.

easily. Furthermore, basically the same code is used for the last step of the track reconstruction chain where all tracks are extrapolated to the PXD.

As a side node, material interactions are currently disabled during the CKF as many extrapolations have to be done to reduce computing time. This only has a minor influence on the reconstructed tracks as the CKF is only used to add hits to previously found tracks and the actual track fit considers material interactions. Nevertheless, studies are going on to replace the GENFIT2 tracking toolkit [4], which is used as the basis of many track definition and track reconstruction tasks throughout the software, with ACTS [5], which promises more efficient modeling of material interactions amongst other advantages.

2.3 SVD Track Finding

All SVD hits that have not been added to tracks extrapolated to the SVD in the previous step are subject to the SVD standalone track finding algorithm, which is commonly called VXDTF2. This algorithm is based on a 'sector map' that stores friendship relations between sectors. To that end, each of the SVD sensors is split to 3×3 sectors. The friendship relation between two sectors is stored in the map if consecutive hits of simulated tracks are found and if certain filter criteria are passed that are used to reject outliers. Accordingly, the sector map greatly reduces the complexity of the track finding problem by lowering the potential combinations of hits that have to be considered. Furthermore, the granularity of the sectors can be tuned to optimize the track finding for high efficiency or low fake rate.

To actually find the tracks a cellular automaton is employed that considers pairs of space points. In the next step, all combinations that do not pass certain filter requirements based on quantities like the angle between the hits or timing are rejected. In the final step, a MVA is used to reject fake and clone tracks.

This algorithm, which was developed in-house, shows a very high performance at the current background levels and is especially important to reconstruct low momentum and forward tracks, which both typically only have a low number of CDC hits and are thus hard to reconstruct by the CDC track finding algorithms.

2.4 Extrapolation from SVD to CDC

This second CKF step aims to add CDC hits to the tracks found by the standalone SVD track finder, which improves the momentum resolution of the tracks. This is illustrated in Figure 6. The left plot shows the fraction of reconstructed tracks that have CDC hits with and without this additional CKF, as a function of the forward angle of the tracks. This slope is denoted as $\tan\lambda$ with $\lambda = \pi/2 - \theta$, so that $\tan\lambda = 0$ means that the track is perpendicular to the direction of the beam. The right figure shows the invariant mass as calculated for two-track events, where both tracks are required to be in the forward detector regions as indicated by the blue dashed lines in Figure 6 (a).

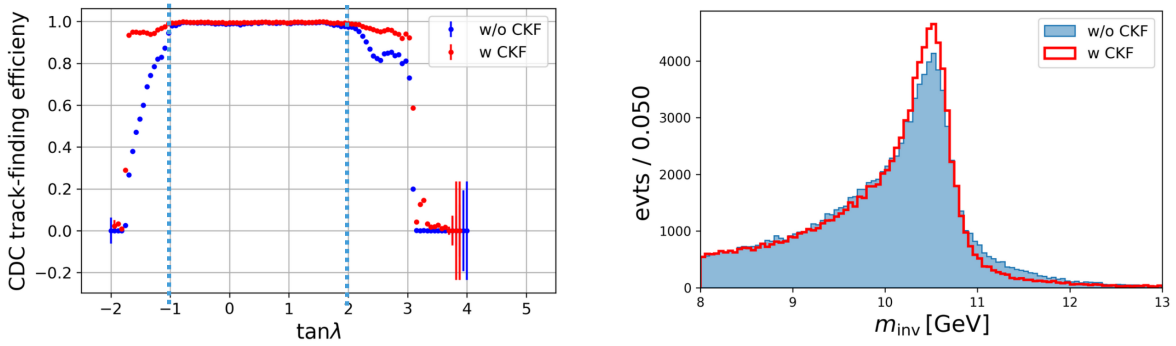

(a) (b)

Figure 6: (a) The fraction of tracks with CDC hits with and without the additional CKF, where $\tan\lambda$ denotes the slope (forward angle) of the tracks. (b) The improvement of the resolution of the invariant mass for events with forward tracks as indicated by the blue dashed lines in (a).

The implementation of this algorithm is very similar to the CKF described in Section 2.2 due to the template based nature but some modifications were necessary to take effects like drift time and left-right-passage into account.

2.5 Current developments

One of the novel tracking algorithms that are currently being developed is introduced to serve as an example. As part of the performance studies discussed in Section 3, it is observed that the electron track finding efficiency is about 5% lower due to material effects (Bremsstrahlung). One of the ideas to mitigate this effect is to use calorimeter showers as seed objects for an additional CKF, as the template-based nature of the CKF allows for easy code development.

The algorithm starts with a helix extrapolation for both charge assumptions from the position of the shower to the center of detector. The helix parameters are then used as a seed for the CKF. As only one of the charge assumptions is expected to be a valid track the candidate with the lower number of hits is rejected. The implementation of the CKF is done, so currently the parameters of CKF are being optimized and the validation on data is performed. Possible extensions of this algorithm include to add Bremsstrahlung photons along the track and to implement a procedure to find vertices from pair production.

3 Performance

This section provides a few benchmark points to illustrate the excellent performance of the Belle II track finding algorithms. Figure 7 shows the track finding efficiency as a function of the transverse momentum p_T and for different particle types for different background levels with respect to the expected beam background at nominal luminosity. These figures show that the track finding efficiency is quite robust against beam background and the performance is still acceptable at twice the expected background level. Furthermore, there is still room for improvements since no optimizations for higher than expected background have been done yet.

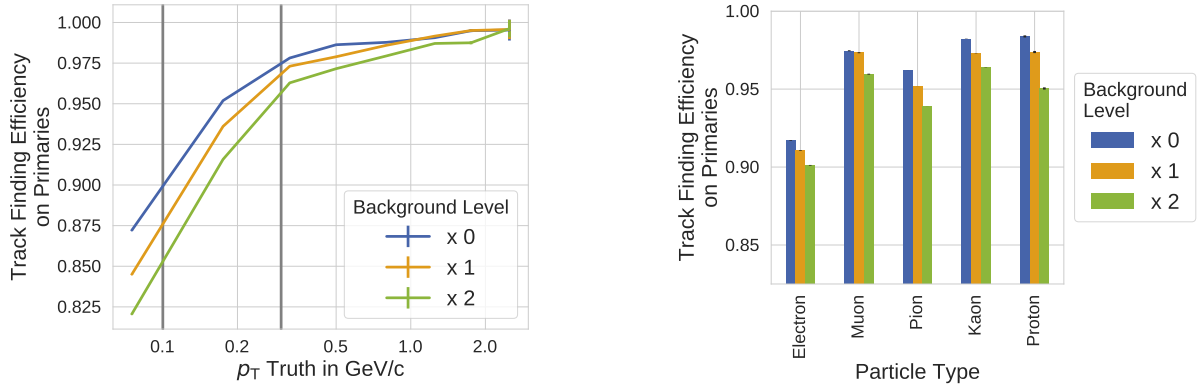

(a) (b)

Figure 7: Track finding efficiency of particles from the primary interaction calculated for simulated $\Upsilon(4S)$ events with different levels of beam-induced background relative to the expected level as a function of the transverse momentum (a) and the particle type (b) [1].

Figure 8 illustrates the excellent track parameter resolution that is achieved with the combination of all tracking algorithms discussed in the previous section. As expected, the d_0 resolution is dominated by the PXD hits and thus basically independent of the background level. In contrast to that, the p_T resolution degrades slightly for low momentum particles at higher background levels.

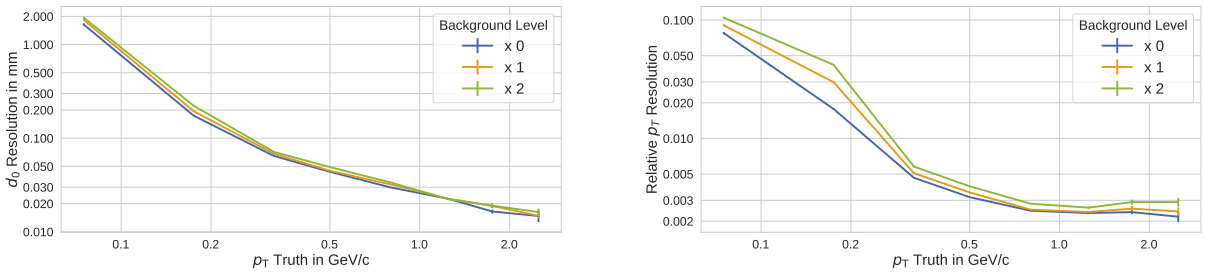

(a) (b)

Figure 8: Track parameter resolutions of (a) the transverse impact parameter d_0 and (b) the transverse momentum p_T for simulated $\Upsilon(4S)$ events with different levels of beam-induced background [1].

More generally, the tracking code runs very stable on the High Level Trigger (HLT). Due to the complexity, tracking is the most time-consuming process as illustrated in 9; all in all about 1 s per event is needed for the track reconstruction. The largest share of this goes to the CDC track finding, the track fitting and the vertexing. The current CPU consumption due to tracking is not critical at current conditions but is expected to rise at higher background levels. However, there is still room for improvement. As mentioned before, there are studies going on to replace GENFIT2 with ACTS, which provides faster fitting and modeling of material interactions, and, as mentioned before, the chain of track finding algorithms can restructured thanks to its modularity.

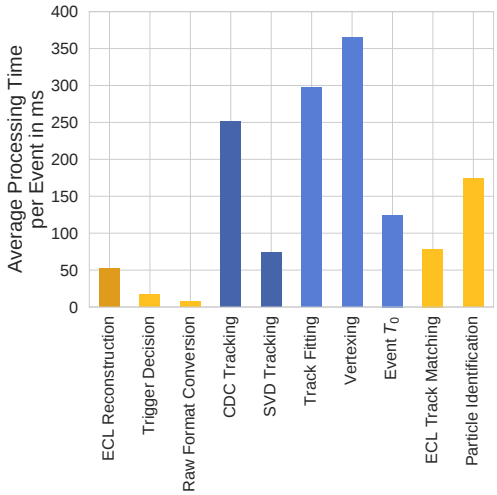

Figure 9: Processing time of a standard online reconstruction performed on s HLT worker node in single-processing mode for different event types [1]. The abbreviation *ECL* refers to the electromagnetic calorimeter of the Belle II experiment.

4 Validation

The track finding efficiency is a key performance driver of Belle II physics. As mentioned before, many physics analyses make use of the Full Event Interpretation approach for which it is especially important to achieve a high tracking performance, know where its limitations are and validate it on data. The track finding efficiency is validated via several tag-and-probe studies focusing on a variety of final states to cover a

large track momentum spectrum. Three of these final states are introduced in Figure 10 for low momentum (left, $p_{\mathrm{T}} \lesssim 200\,\mathrm{MeV}$), intermediate momentum (middle, $200\,\mathrm{MeV} \lesssim p_{\mathrm{T}} \lesssim 2\,\mathrm{GeV}$) and high momentum (right, $p_{\mathrm{T}} \gtrsim 2\,\mathrm{GeV}$).

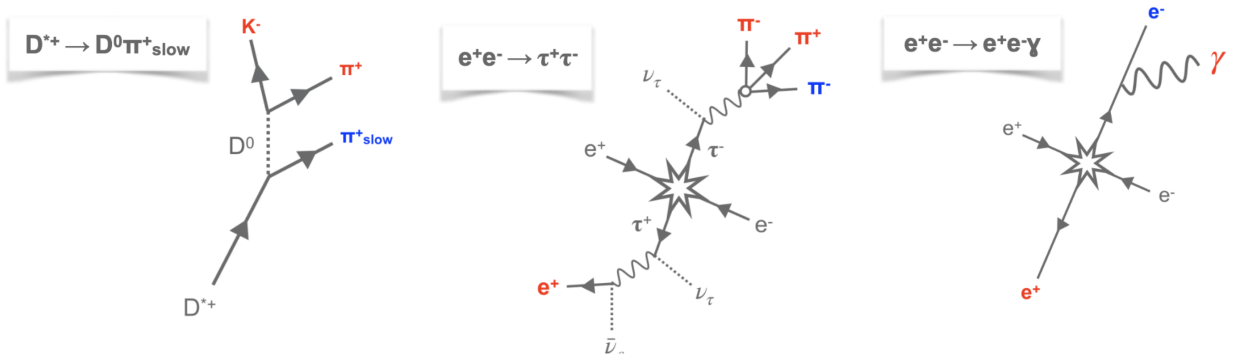

Figure 10: Examples of final states used for Tag-and-Probe studies to validate tracking efficiencies in data for low (left), intermediate (middle) and high (right) momentum tracks. The tracks used to tag the event are shown in red, while the blue track corresponds to the probe object.

The main principle of these tag-and-probe analyses can be explained quite easily focusing on the $\tau-$pair events depicted in the middle figure. To tag the events it is required that they have three good quality tracks with a total charge of ± 1 (shown in red). The probe then consists of a fourth track (blue) that passes only loose selections and conserves the charge. The track finding efficiency ϵ can then be derived using the ratio of the number of events where the probe track is found N_4 and where it is not found N_3 as

$$\epsilon \cdot A = \frac{N_4}{N_3 + N_4} \tag{2}$$

where A corresponds to the geometric acceptance of the detector. These studies show high level of agreement between data and simulation - otherwise they can serve as direct input to improvements of the simulation.

5 Beam Spot Width and Alignment

As mentioned in the introduction, SuperKEKB strongly squeezes the beams at the interaction point. This leads to the fact that the vertical beam spot size σ_y is much smaller than the resolution of the transverse impact parameter d_0 with $\sigma_y < 2\,\mu\mathrm{m}$. The horizontal beam spot size is in the order of $\sigma_x \approx 15\,\mu\mathrm{m}$. In comparison, Figure 11 shows the impact parameter d_0 as measured on data, with (red) and without (blue) the inclusion of PXD hits during the track reconstruction. Only horizontal tracks were selected so that the size of the beam spot can be neglected. Furthermore, it should be noted that for this figure data from an early experiment was used, where only parts of the vertex detector were installed. As anticipated, the PXD significantly increases the resolution down to $\sigma = 12.1\,\mu\mathrm{m}$, which is in agreement with the expected value of $\sigma \approx 10\,\mu\mathrm{m}$. It should be mentioned here that this high performance is not possible without good detector alignment. In our case, the alignment is corrections are derived using the Millepede-II framework [9].

Furthemore, the d_0 distribution can be determined as a function of the polar angle ϕ_0 as shown in Figure 12 (a), which reveals a dual peak shape coming from the asymmetric beam profile ($\sigma_y \ll \sigma_x$). The simulation is observed to underestimate the resolution slightly. Moreover, this distribution can be unfolded using the d_0 resolution discussed previously to determine the beam profile, which is shown in Figure 12 (b). Here, a good simulation of the beam profile can be seen. All in all, this illustrates the excellent performance of the track reconstruction and alignment procedures once more.

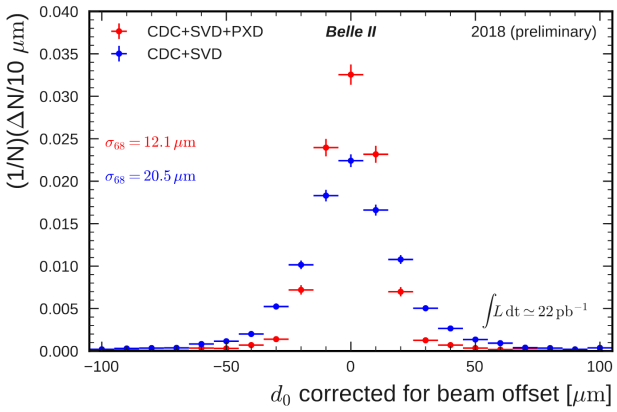

Figure 11: Distribution of transverse impact parameter d_0 as determined from events with two horizontal tracks [7].

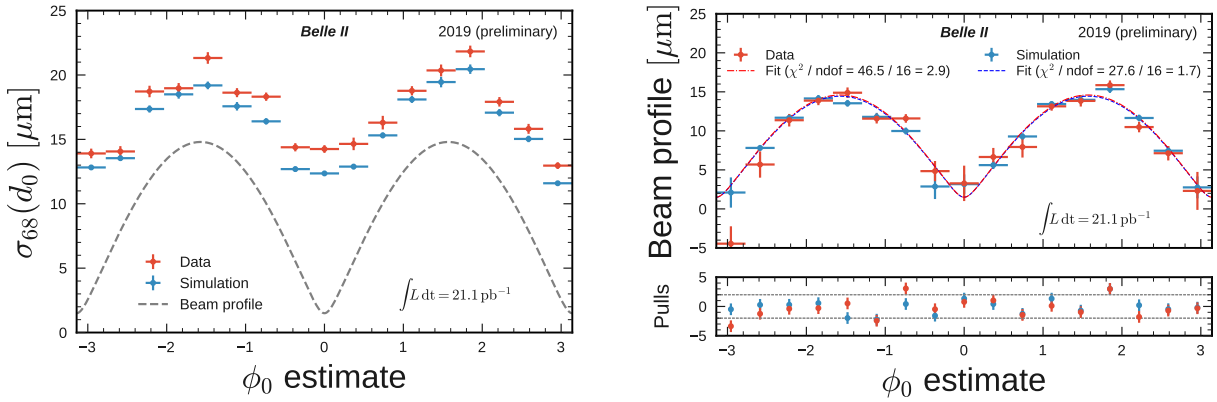

(a) (b)

Figure 12: Distribution of transverse impact parameter d_0 as measured in dimuon events (a) and unfolding of the distribution using measured the d_0 resolution to show the beam profile (b) [6].

6 Measurement of D^0 lifetime

As a highlight, one of the first highly track-based analysis is described as it is an excellent exercise to test vertexing and alignment, apart from the obvious physics aspect. The analysis is based on pure samples of $D^{*+} \to D^0\pi^+$ events, where the D^0 is reconstructed in three different final states. The largest systematic uncertainties are expected to stem from residual misalignment and beam spot measurements. As illustrated before, these effects are well understood and corrected for to best current knowledge.

One of the results of the analysis is shown in Figure 13, which shows the distribution of the lifetime of the D^0 meson in one of the channels. The clear exponential decay towards long lifetimes is especially reassuring as it confirms the excellent resolution of achieved with the Belle II tracking detectors and employed algorithms. The D^0 lifetime extracted from the fit of the very limited data is determined as 370 ± 40 fs, which is statistically compatible with the expectation.

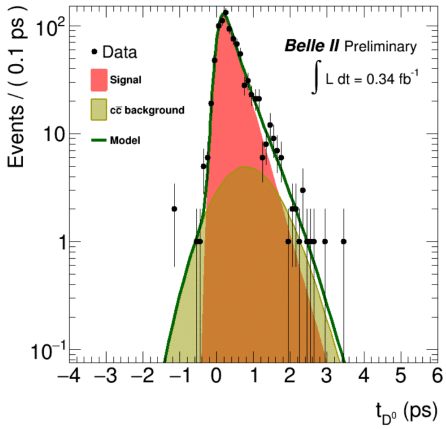

Figure 13: Determination of the D^0 lifetime [8].

7 Conclusions

The Belle II tracking detectors and algorithms show excellent performance throughout all of the presented studies. The high flexibility of the modular approach proves to be advantageous to adjust tracking to background conditions and performance requirements, as well as future extensions of algorithms. Currently, many studies are going on trying to exploit MVA-based approaches for all kinds of use cases like background rejection, track quality and many others. The high performance of Belle II tracking is also reflected in the first tracking-focused studies and analyses like the determination of the D^0 lifetime.

References

[1] Belle II Tracking Group, "Track Finding at Belle II", submitted to Computer Physics Communications, 2020, arXiv:2003.12466.

[2] Thomas Keck, "FastBDT: A speed-optimized and cache-friendly implementation of stochastic gradient-boosted decision trees for multivariate classification", Computing and Software for Big Science, 2017, arXiv:1609.06119.

[3] Thomas Keck et al., "The Full Event Interpretation – An Exclusive Tagging Algorithm for the Belle II Experiment", Computing and Software for Big Science, 2019, arXiv:1807.08680.

[4] Nils Braun, Thomas Hauth, Thomas Kuhr et al., "Implementation of GENFIT2 as an experiment independent track-fitting framework", 2019, arXiv:1902.04405.

[5] Xiaocong Ai for the ACTS developers, "Acts: A common tracking software", 2019, arXiv:1910.03128.

[6] Nils Braun, Alexander Glazov, Carsten Niebuhr, Eugenio Paoloni and Cyrille Praz, "Study of the impact parameter resolution and the beam profile in early Phase 3 data", Belle II Note, 2019, BELLE2-NOTE-PL-2019-011.

[7] Nils Braun, Alexander Glazov, Eugenio Paoloni and Cyrille Praz, "Study of the d0 track resolution", Belle II Note, 2019, BELLE2-NOTE-PL-2018-037.

[8] Gaetano de Marino and Giulia Casarosa, "D0 Lifetime On Bucket6 Phase3 Data", Belle II Note, 2019, BELLE2-NOTE-PL-2019-003.

[9] Klaus Kleinwort, Millepede II Documentation, `https://www.desy.de/~kleinwrt/MP2/doc/html`.

[10] Belle II Collaboration, "Belle II Technical Design Report", 2010, arXiv:1011.0352.