

Decay Chain Reconstruction at Belle II

Francesco Tenchini^{1,*} and Jo-Frederik Krohn²
(on behalf of the Belle II Software Group)

¹Deutsches Elektronen-Synchrotron, Hamburg, Germany

²University of Melbourne, Melbourne, Australia

Abstract. The Belle II detector underwent commissioning in 2018. Featuring a state of the art vertex detector with innovative pixel sensors, it will record collisions of e^+e^- beams from the SuperKEKB accelerator which is slated to provide luminosities 40 times higher than KEKB.

This will bring new opportunities and challenges for physicists seeking to perform precision measurements, as increasingly precise vertex position reconstruction and background rejection power is needed to take advantage of the unprecedented statistics provided.

We discuss the vertex reconstruction techniques available that make such measurements possible and present an approach to particle decay reconstruction based on the global fit of the full decay chain, implemented using a Kalman Filter. This technique is especially well suited to the treatment of processes with neutral decays. It allows for improved background rejection, access to a larger number of decay modes and a flexible framework for consistent treatment of measurement uncertainties.

1 Experimental Context

The Belle II experiment [1] is located at the SuperKEKB asymmetric electron-positron collider at KEK, Japan, which was designed to provide an instantaneous luminosity of $8 \times 10^{35} \text{ cm}^2\text{s}^{-1}$ at or near a center of mass energy of 10.58 GeV, i.e. the $Y(4S)$ bottomonium resonance. This rate, a factor of 40 higher than previous colliders of this kind[2][3], is exploited by Belle II to precisely measure the properties of the B mesons originating from the $Y(4S)$ decay. For this reason, experiments such as this and the one before it are often called *B-factories*.

A full description of the Belle II detector is available elsewhere in [1]. Succinctly, Belle II surrounds the collision point of the two charged beams with a cylindrically symmetric detector, composed of a barrel and two endcaps, immersed in a 1.5 T magnetic field. The barrel is composed of, in outward order, a silicon pixel detector (PXD), a silicon strip vertex detector (SVD), a wire drift chamber (CDC), a time-of-propagation Cherenkov counter (TOP), an electromagnetic calorimeter (ECL) and a K-long and muon detector (KLM). The endcaps contain an additional ECL each, as well as a ring-imaging Cherenkov detector (ARICH) in the forward endcap.

*e-mail: francesco.tenchini@desy.de

2 Reconstruction Paradigm

Belle II presents a relatively clean environment in which each stable decay product of the initial interaction can be reconstructed individually:

- *Charged particles* are reconstructed as helicoidal tracks in the Belle II tracking volume. Both momentum and impact parameter of these objects can be measured accurately.
- *Neutral particles* such as photons are observed as energy deposits in the Belle II electromagnetic calorimeter. Momentum is estimated using the cluster position, assuming the particle to originate from the interaction point. This is not always a safe assumption, as we will show.

Other detectors provide particle identification. A typical physics analysis is *hypothesis-based*: candidates for the desired process are built out of the available final state particles; then one attempts to use known physical properties of the decay to separate the real signal from the combinatorial component, called "background". A powerful approach here consists in fitting particle trajectories and 4-momenta together to determine whether they originate from a common vertex, and to extract the kinematics of the composite particle that produced it. This is commonly called "vertex fitting" and is a straightforward procedure when charged tracks are involved; less so in the presence of neutrals.

The traditional and most simple approach to fitting a chain of several particle decays is vertex by vertex (also known as "leaf by leaf"), where decay vertices are reconstructed and fitted individually. We will discuss the $B^0 \rightarrow K_S^0(\rightarrow \pi^+\pi^-)J/\psi(\rightarrow \mu^+\mu^-)$ decay as an example (Fig. 1). To fit this decay chain, we would first fit the $K_S^0 \rightarrow \pi^+\pi^-$ and $J/\psi \rightarrow \mu^+\mu^-$ decay vertices, then proceed to fit $B^0 \rightarrow K_S^0 J/\psi$ using the results from the fits above. In this approach, the fit is blind to information two or more stages down in the chain: only partial knowledge of the nature of final state tracks ($\pi^+\pi^-\mu^+\mu^-$) is retained in the final stage since the covariance is linearly propagated.

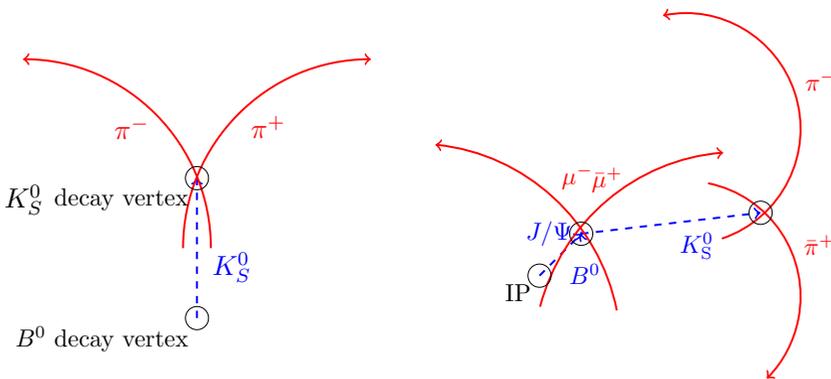


Figure 1. (Left) Simple decay vertex with two daughter tracks. (Right) Complex decay tree with a prompt decay and a displaced secondary vertex.

The Belle II analysis software framework (*basf2*) [4] provides several vertex fitting tools to the user to perform hypothesis-based analysis. The key ones are:

- *KFit*: This is the simplest fitter which was inherited from the Belle software [5]. It relies on least square minimisation and is well suited to fitting single vertices. While extended versions have been considered in the past, they do not scale well (see Section 3).

- *RAVE*: A standalone package originating from the CMS libraries [6]. Also suitable for single vertex fits, this algorithm uses a Kalman Filter [7] approach and allows for adaptive vertex fitting, where outliers not belonging to the main vertex are progressively downweighted and their contribution removed. This particular aspect is not useful in hypothesis-based fitting, but has applications in inclusive B-tagging algorithms.

Let's now consider a different example in $B^0 \rightarrow K_S^0(\rightarrow \pi^0(\rightarrow \gamma\gamma)\pi^0(\rightarrow \gamma\gamma))J/\psi(\rightarrow \mu^+\mu^-)$ (Fig. 2)

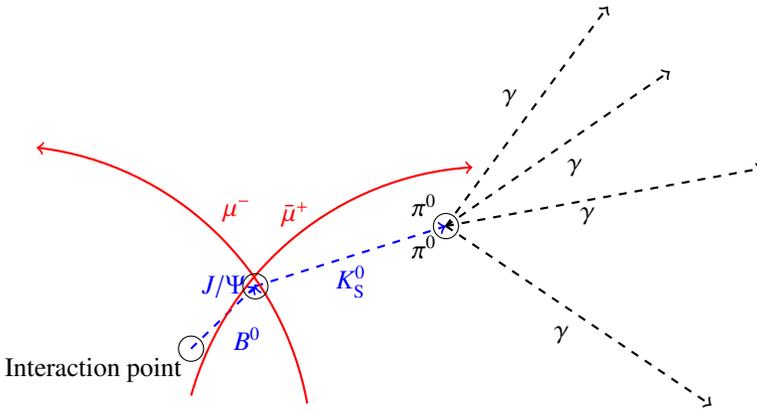


Figure 2. Variant of the Fig. 1(left) channel where the displaced K_S^0 vertex decays into a neutral-rich, trackless topology.

As mentioned, neutral particles do not have associated tracks and are only identified by their energy deposit in the ECL. Therefore their true momentum vector is unknown; reconstruction assumes it to originate from the interaction point. This is problematic as the K_S^0 vertex can be displaced from the interaction point by several centimeters; however, if fitting leaf by leaf, the fitter is blind to this information until reaching that point. The first $\pi^0 \rightarrow \gamma\gamma$ fits are performed under the naive assumption of no displacement, resulting in a biased 4-momentum estimate. When progressing upstream to the $K_S^0 \rightarrow \pi^0\pi^0$ vertex the fitter loses access to the γ , meaning this bias cannot be corrected anymore.

3 Global Fit Challenges and Implementation

To solve this issue and generalise the vertex fitting concept, we consider an alternate, global technique where the full decay tree is parametrised and fitted simultaneously. This is a more organic approach which guarantees that constraints are applied in a consistent way across the fit chain and, depending on the topology, can allow for fitting of such neutral-rich decays.

A first version of this approach was originally developed for the BaBar collaboration and is described in [8] along with the underlying algebra. We build on this concept, implement it and optimise it for the basf2 framework, under the name of *TreeFitter*.

The fundamental challenge of a global approach is of computational nature as the minimisation is limited by the speed at which large matrices can be inverted. A decay chain can be described by the free parameters of each particle it is composed of; one particle can have up to eight free parameters (E , p_x , p_y , p_z , x , y , z and flight length λ). Given a N total

parameters, an approach based on least square minimisation requires to solve a system of N equations, which is equivalent to inverting a $N \times N$ matrix. This rapidly stops being viable as the decay complexity increases, as even for optimised algorithms the computation time scales as $O(N^{2.373})$ [9].

We solve this problem by using a Kalman filter. The requirements imposed on the fit (constraints) are applied individually in iterative fashion until convergence is reached or divergence is observed. Only a subset of the state vector is affected and therefore a large matrix inversion is replaced with a series of smaller matrix operations.

We rely on the Eigen library [10] for fast treatment of matrix operations. This is the first major implementation of the library within basf2, designed for speed and robustness of linear algebra calculations. Past proof-of-concept implementations have shown up to an order of magnitude improvement in execution speed compared to older approaches [11].

4 Practical Applications

4.1 Neutral-Rich Topologies

Let us return then to the neutral-rich topology in Fig. 2. In a global approach this problem is naturally solved, as the K_S^0 origin is well defined by the short lived J/ψ vertex, allowing to iteratively correct the mass measurement by propagating the information back.

We test this on a simulated sample; the results are visible in Fig. 3. The K_S^0 mass bias present in the unfitted sample is corrected and the resolution is improved. A mass selection could then allow for the separation of combinatorial background.

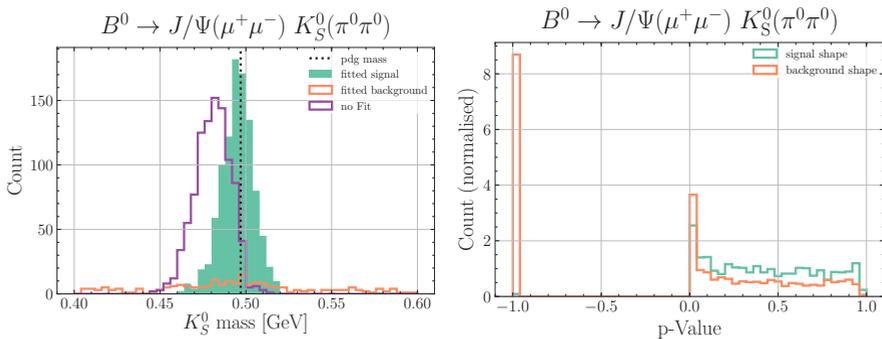


Figure 3. (Left) K_S^0 invariant mass before and after the global fit. The bias is corrected and the resolution is improved. A selection around the mass peak would allow rejection of combinatorial background (in orange) from the sample. (Right) P-values of the global fit. Convergence fails for a significant portion of the combinatorial component, represented by the negative bin; this already provides moderate background suppression.

4.2 Direct Extraction of Observables

In a global approach, quantities relating to the decay chain can be parametrised in a simple manner and saved as output for physics analysis. Thanks to TreeFitter having full access to the covariance matrix of the parameters, propagation of uncertainties is straightforward and consistent throughout. This greatly simplifies error treatment on the user side.

As an example, we extract the D^0 decay length from a simulated sample of $\bar{B}^0 \rightarrow D^{*+}(\rightarrow D^0(\rightarrow K^-\pi^+)\pi^+)\pi^-$ (see Fig. 4). The result is visible in Fig. 5 and shows good agreement with the generated values.

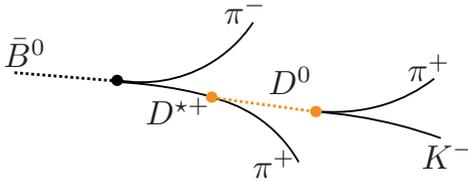


Figure 4. Sketch of the $\bar{B}^0 \rightarrow D^{*+}(\rightarrow D^0(\rightarrow K^-\pi^+)\pi^+)\pi^-$ decay. Decay lengths not to scale.

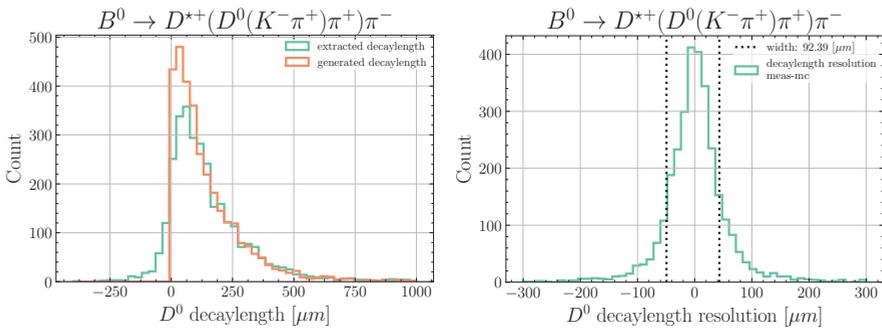


Figure 5. Decay length of the D^0 meson directly extracted from the decay chain (*left*) and corresponding deviation from generated values (*right*).

4.3 Customised Constraints

Since that each constraint is filtered independently when applied the TreeFitter could be designed in a modular way, allowing for further development and implementation of user-defined constraints based on their needs. As an example, we present the case of $\bar{B}^0 \rightarrow D^{*+}\mu^-\nu_\mu$ (see Fig. 6), with $D^{*+} \rightarrow D^0\pi^+$. This decay chain is difficult to fit due to the many free parameters introduced by the undetected neutrino.

Instead, we limit the D^{*+} production vertex (i.e. the \bar{B}^0 decay vertex) to an independently measured spatial distribution from generic B decay events and do not fit the \bar{B}^0 explicitly. The D^{*+} decay vertex resolution is improved as a consequence (Fig. 7).

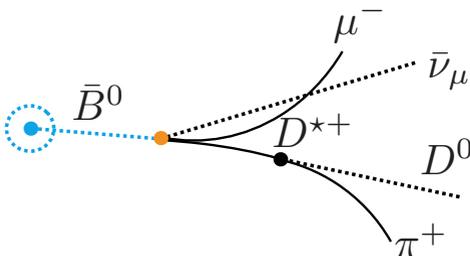


Figure 6. Sketch of the origin-constrained $\bar{B}^0 \rightarrow D^{*+}\mu^-\nu_\mu$ decay. Decay lengths not to scale.

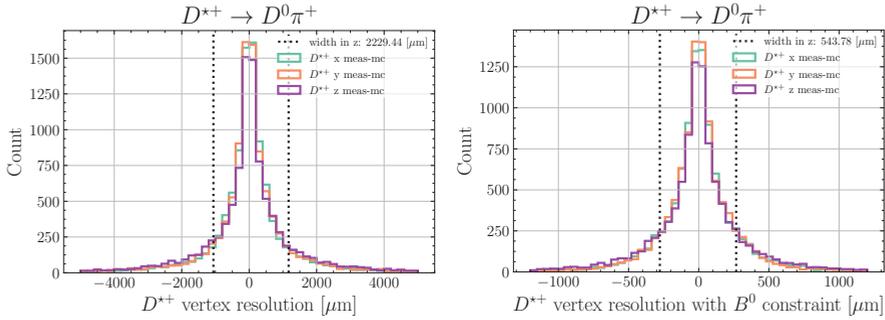


Figure 7. $D^{*+} \rightarrow D^0\pi^+$ vertex resolutions from simulated sample, before (left) and after (right) constraining the \bar{B}^0 decay vertex

5 Conclusions

We implemented a global fitting technique designed for B-factory environments which is especially suited for the treatment of neutral-rich topologies and the fit of underconstrained vertices. It will provide a powerful tool for future Belle II analyses due to its ability to suppress combinatorial backgrounds; treat constraints - and therefore uncertainties - consistently across the fit; and allowing reconstruction of hard to fit neutral decays, facilitating their inclusion in physics measurements and therefore increasing their statistical power.

An article [12] describing the implementation in greater detail is currently under internal collaboration review.

Acknowledgements

We would like to thank Wouter Hulsbergen for his original paper, the discussion and sharing his code.

References

- [1] T. Abe, et al., Belle II Technical Design Report, KEK Report 2010-1
- [2] S. Kurokawa, et al., Nucl. Instrum. Meth. A **499**, 1 (2003)
- [3] PEP-II: An Asymmetric B Factory. Conceptual Design Report., SLAC-418, LBL-5379 (1993)
- [4] T. Kuhr, et al. Comput. Softw. Big Sci. **3** 1 (2019)
- [5] J. Tanaka, Belle Note **194**
- [6] W. Waltenberger, IEEE Trans. Nucl. Sci. **58**, 434 (2011)
- [7] R. Friuwirth, Nucl. Instrum. Meth. A **262** 444 (1987)
- [8] W. D. Hulsbergen, Nucl. Instrum. Meth. A **552** 566 (2005)
- [9] D. Coppersmith, S. Winograd, Journal of Symbolic Computation, **9** (3), 251-280 (1990)
- [10] Eigen project, "Eigen" [software], version 3.3.5, 2018.
 Available from <https://bitbucket.org/eigen/eigen/> [accessed 2018-10-25]
- [11] T. Keck, "FastFit" [software].
 Available from <https://github.com/thomaskeck/FastFit> [accessed 2018-10-25]
- [12] Global Decay Chain Vertex Fitting at B-Factories, in preparation