Journal of
Instrumentation

A multi-Hough-based displaced vertex track trigger for the Belle II experiment

To cite this article: Kai Lukas Unger *et al* 2025 *JINST* **20** C02051

Published by IOP Publishing for Sissa Medialab

Received: *November 2, 2024*
Accepted: *January 30, 2025*
Published: *February 26, 2025*

Topical Workshop on Electronics for Particle Physics
University of Glasgow, Scotland, U.K.
30 September–4 October 2024

A multi-Hough-based displaced vertex track trigger for the Belle II experiment

Kai Lukas Unger,[a,*] **Jürgen Becker**,[a] **Timo Forsthofer**,[b] **Simon Hiesl**,[b] **Christian Kiesling**,[b] **Marc Neu**,[a] **Elia Schmidt**,[b] **Felix Meggendorfer**,[b,c] **Tiancheng Mi**[a] **and Zuwen Ou**[a,1]

[a] *Karlsruhe Institute of Technology (KIT), Kaiserstraße 12, 76131 Karlsruhe, Germany*
[b] *Max-Planck-Institute for Physics (MPI), Boltzmannstr. 8, 85748 Garching, Germany*
[c] *Technical University Munich (TUM), Arcisstraße 21, 80333 München, Germany*

E-mail: kai.unger@kit.edu

Abstract: The Belle II experiment at the asymmetric-energy electron positron collider SuperKEKB aims to explore physics beyond the standard model (BSM). One of the widely discussed signatures for BSM are new, long-lived neutral particles, which decay into charged mesons or lepton pairs originating from a vertex usually far from the interaction point (IP) of the colliding beams. The current level-1 (L1) track trigger system is optimized for particles created at the IP and will thus reject, with high probability, such interesting events. This makes it necessary to develop a special L1 track trigger for events with a vertex displaced from IP, the Displaced Vertex Trigger (DVT). The pipelined and deadtime-free L1 trigger system of Belle II utilizes a set of FPGA boards to make rapid decisions within 5 microseconds. The new DVT will be operating in parallel with the existing track trigger systems. The DVT identifies events, where two tracks with opposite charge originate from a common vertex away from the IP. The track finding is done by a set of Hough transformations, each one assuming a certain track origin from a grid spanning the tracking volume. The correct vertex is then determined via a shape analysis of the Hough clusters, using neural networks. To manage the large number of track origin hypotheses, a pre-selection of candidates based on the properties of the Hough map is employed, which significantly reduces the required FPGA resources.

Keywords: Trigger algorithms; Trigger concepts and systems (hardware and software)

*Corresponding author.
[1] Now at Technische Universität Dresden (TUD).

https://doi.org/10.1088/1748-0221/20/02/C02051

2025 JINST 20 C02051

Contents

1 The Belle II trigger system

One of the main tasks of the Belle II experiment at the asymmetric-energy electron-positron collider SuperKEKB (Tsukuba, Japan) is to search for physics beyond the standard model (BSM). Promising BSM event classes predict new, long-lived neutral particles [1–3] which decay into charged hadron or lepton pairs far from the interaction point (IP), thus producing event signatures with displaced vertices. However, the current level-1 (L1) track trigger [4] is optimized for tracks from the IP and therefore tends to discard track pairs from a displaced vertex as background. To address this problem, a new trigger system, named Displaced Vertex Trigger (DVT), is required that can identify two oppositely charged tracks originating far from the IP.

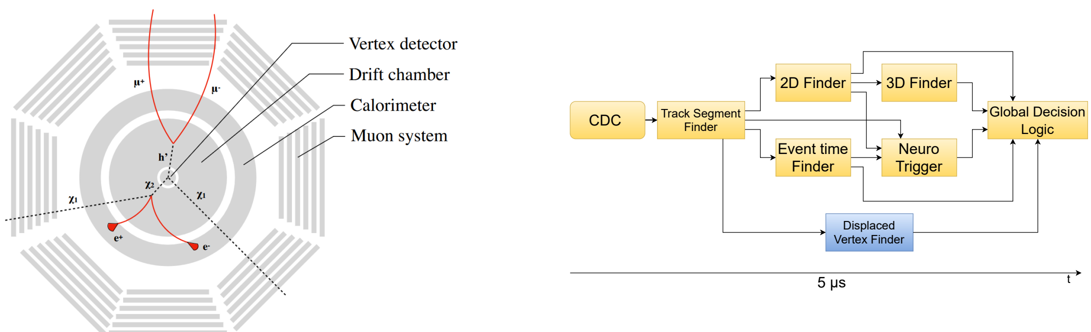

Figure 1. *Left:* the Belle II L1 Track Trigger System, including the position and the input/output connections of the DVT (blue box) to the existing L1 track trigger. *Right:* an example displaced vertex signature in the schematic 2D view of the Belle II detector. Reproduced from [2]. CC BY 4.0.

The pipelined L1 trigger system consists of custom-designed FPGA boards, called Universal Trigger Board 4 (UT4), which host Xilinx Virtex Ultrascale FPGAs. The L1 decision to keep an event

2025 JINST 20 C02051

must not take longer than 5 µs, from the detector readout to the final decision. The new Displaced Vertex Trigger (indicated by the blue box in the figure 1) will run parallel to the existing track trigger system. The hit information for the track triggers, derived from Belle II's Central Drift Chamber (CDC), is filtered beforehand in the customized Track Segment Finder (TSF) [5] and then forwarded to the DVT. The DVT should then search for two oppositely charged tracks from a common vertex that is displaced (typically more than 10 cm) from the IP.

2 Concept of the Displaced Vertex Trigger

The idea for recognizing two tracks from a common, a priori unknown, displaced vertex is to perform a set of two-dimensional Hough transformations in the transverse (x, y) plane of the CDC. Since a Hough transformation requires a vertex hypothesis, the (x, y) plane of the CDC is divided into a grid of macro cells, and each of the macro cell centers serves as vertex hypothesis for one of the 2D Hough transformations. Using a size for the macro cells of about 10 cm^2 would result in about 400 such cells, resulting in the same number of Hough transformation to be executed in parallel. This number can be reduced to about 100 if the macro cell sizes increase towards larger radii. In the outer 10–15 % of the CDC it is anyway impossible to find two tracks coming from a common vertex as there are far too few hits before the tracks leave the detector, so no macro cells are needed in this region. The optimal placement of the macro cells is the goal of further investigations. In this paper we focus on the possible hardware implementation of the DVT and the limiting constraints imposed on the FPGA side.

3 Architecture of the Displaced Vertex Trigger

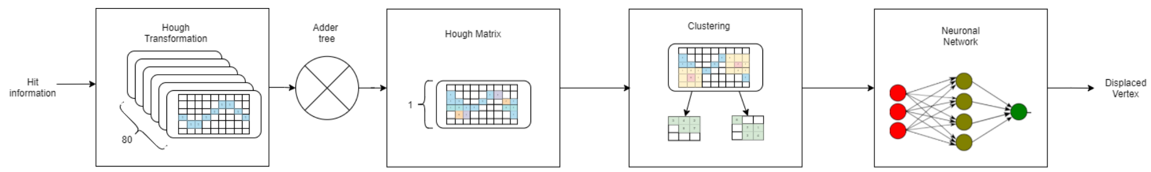

Figure 2. Schematic structure of the architecture design for the DVT.

The architecture of the DVT consists of 5 sub-modules (shown in figure 2). First, the two-dimensional Hough curves are calculated with the help of look-up tables, fully parallel in each of the 4 quadrants of the CDC. As is the case for the macro cell grid, the Hough space is also binned in the two parameter variables $\rho = 1/r$ and ϕ, where r is the radius (or transverse momentum) of the track candidate and ϕ is the azimuthal emission angle of the track at the macro cell center. The entire processing is done in a pipelined manner and is deadtime free so that no events are lost. Pre-filtering limits the maximum number of hits to 80 hits per CDC quadrant. As the FPGA is clocked 4 times faster than the data rate of the trigger electronics (32 MHz), 20 Hough curves must be generated in parallel within the ρ, ϕ plane of the Hough map. The individual Hough curves are then added up bin by bin to create a "colored" Hough map with weights in each Hough cell. In this map, the maximum bin is then searched for and a predetermined fixed shape is used to cluster the Hough cells around this maximum. The cluster parameter information and the hits contributing to the cluster are afterwards fed into a pre-trained multilayer perceptron (MLP) which performs an analysis on the cluster shapes and finds by this the correct macro cell (origin) of the displaced vertex.

2025 JINST 20 C02051

4 Implementation results

The hardware architecture presented above (except for the MLP part) is implemented on a UT4 board equipped with an UltraScale XCVU160 FPGA. Based on the estimated resource consumption, no more than two vertex hypotheses with the architecture presented above will fit on one FPGA. With approximately 100 vertex hypotheses envisaged would mean 50 FPGAs. Technically, it would be possible to integrate the architecture into the Belle II trigger system, but only with excessive development time and cost. One problem is the adder tree since each Hough curve resulting from the 80 admitted hits has to be added up. In addition, excessive memory is required to store the entire final Hough matrix until the clustering is completed. Furthermore, finding the maximum in the Hough map, clustering the region around the found maximum, and vetoing this region for the search of the second maximum is also quite resource-consuming.

Table 1. Hardware utilization of one instance of the DVT on UltraScale XCVU160 (without the neural network part).

Resource	Utilization	Available	Utilization %
LUT	355661	926400	38.39 %
Flip-Flop	196391	1852800	10.60 %
BRAM	50	3267	1.53 %
DSP	348	1560	22.31 %

5 Concept of the Hough-based pre-selection

Obviously, the number of vertex hypotheses so far considered for the cluster shape analysis is too large, so the number of candidates must be reduced by intelligent pre-selection. By studying the Hough map for different vertex hypotheses, a distinguishing pattern of the clusters in the weighted Hough map becomes visible: the further away the vertex hypothesis is from the real origin, the flatter and wider the cluster shape in the Hough map becomes. As can be seen in figure 3, the upper hypothesis is very close to the correct vertex, and the lower hypothesis is quite far away. If the vertex hypothesis is close to the real origin, two maxima representing the two particle tracks are easily recognized as sharp maxima. In addition, the Hough curves around the cluster maxima are clearly fanning out into the regions at the top and bottom (in the $\rho = 1/r$ direction) of the Hough map. The lower image shows the Hough map for a vertex hypothesis at large distance from the true origin. No sharp maximum is visible here and the two clusters in the Hough map merge into each other. In addition, no fanning out can be seen in the upper and lower regions of the Hough map.

These findings suggest to simply count the pixels in a black and white image in the upper and lower parts of the Hough map (N_{out}), and in the center part (N_{in}) and calculate the ratio $R = N_{\text{out}}/N_{\text{in}}$. The ratio R provides a simple selection criterion for the most probable origin hypotheses. Figure 3 shows the method without noise in the CDC for clarity. With the present noise level, the procedure works robustly. Its behavior projected to the design luminosity, entailing much higher background noise, is the subject of ongoing investigations.

2025 JINST 20 C02051

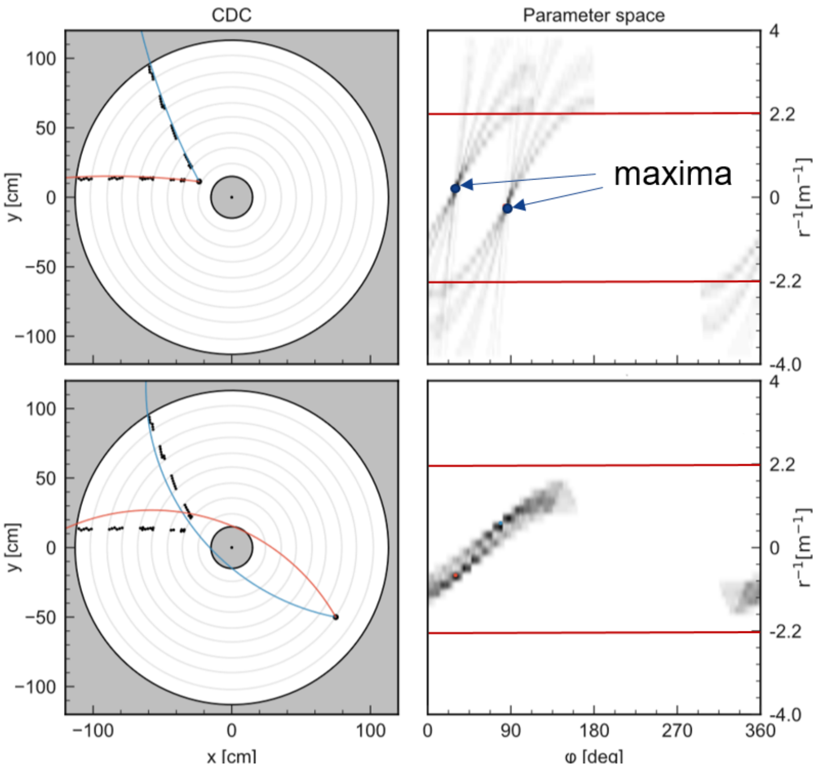

Figure 3. Illustration of two vertex hypotheses. On the left the transverse (x, y) plane of the CDC with the macro cell centers (light gray dots), the wire hits (black dots) and the corresponding reconstructed tracks are shown, on the right the corresponding Hough maps ($\rho = 1/r, \phi$ plane).

6 Architecture of the Hough-based pre-selection

The pre-selection architecture shown in figure 4 is based on the DVT architecture described above. But the difference now is that the Hough curves are generated as a 1-bit black and white image, not "colored" anymore like the actual weighted Hough cells. In order to save the hardware-intensive adder tree, the individual curves are no longer added up, but simply ORed. This means that the image memory size remains the same, but information is lost and it is no longer possible to search for maxima. Once the black and white Hough map has been created, it is split into 3 regions (see figure 3). In each region, the number of black pixels (represented as 1) is counted. The number of pixels in the upper and lower regions is then subtracted from the number of pixels in the center area. The ratio R can be used to sort the vertex hypotheses and determine a limited set, say O(▽–∞ɪ), of highly probable vertex hypotheses for a more detailed analysis using the "colored" Hough map. This detailed analysis can be done with the architecture presented in section 3 as fewer FPGAs are now required (1–3 FPGAs), depending on how reliably the set of vertex hypotheses can be determined with the pre-selection.

7 Implementation of the pre-selection

The pre-selection was again implemented on the UT4 board (UltraScale XCVU160) used in the experiment, with table 2 showing the results of the pre-selection. Compared to the DVT (table 1), digital signal processors (DSPs) are no longer required as the adder tree has been replaced by the OR

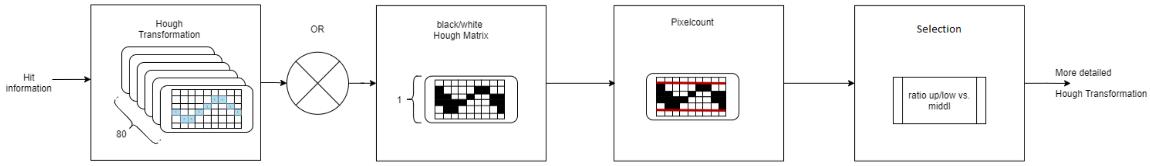

Figure 4. Schematic structure of the architecture design for the pre-selection.

function. In addition, the resource-intensive maximum finding and clustering are no longer required and replaced with the pixel counter. The BRam resources required stay the same in both cases as it is used for the pre-calculated Hough transformations.

2025 JINST 20 C02051

Table 2. Hardware utilization one instance of the preselection DVT on UltraScale XCVU160.

Resource	Utilization	Available	Utilization %
LUT	59340	926400	6.41 %
Flip-Flop	35105	1852800	1.89 %
BRAM	50	3276	1.53 %
DSP	0	1560	0 %

A first successful implementation with a set of 12 pre-selected vertex hypotheses shows that with 10 FPGAs (9 for pre-selection, and 1 for the DVT) a complete DVT with 100 original hypotheses can be realized. Compared to the first pure Hough DVT implementation (50 UT4 FPGAs) this is a saving of 40 FPGAs. In addition, there are possibilities to optimize the pre-selection algorithm and to further reduce the hardware resources. The idea is to perform the pixel count for each Hough curve directly during creation. This would mean that only the three values ($N_{\mathrm{out,upper}}$, $N_{\mathrm{out,lower}}$, and N_{in}) need to be saved and not a complete Hough map. However, an analysis of the possibly reduced efficiency and an implementation in hardware is still pending.

8 Summary

In this paper, we present prototypical implementations of a new Displaced Vertex Trigger (DVT) for the Belle II experiment at the first trigger level (L1). The DVT is intended to run in parallel to the existing L1 track trigger systems. The aim of the DVT is to recognize two oppositely charged tracks from a common displaced vertex using the method of Hough transforms in two dimensions (x, y plane of the Central Drift Chamber, CDC). Since the displaced vertices can occur anywhere in the CDC, the (x, y) plane is subdivided into about 100 macro cells, their centers serving as vertex hypotheses for the Hough transformation. The Hough curves are calculated using look-up tables and are processed in a pipeline without deadtime. The selection of the correct vertex is done by a shape analysis of the clusters around the Hough maxima using a neural network (MLP). In order to limit the number of required FPGAs, a pre-selection of the possible candidates is done, using a black-and-white Hough map. In a second step the full Hough map with weighted cells is applied to the (few) remaining candidates and the correct vertex is then found with the MLP. This architecture is successfully implemented on UT4 boards, equipped with an UltraScale XCVU160, at the moment still without the MLP part.

Acknowledgments

Funded by the German Federal Ministry of Education and Research under "Verbundprojekt 05H2021 (ErUM-FSP T09) - Belle II: Pixeldetektor, Software und erste Datenanalysen".

References

[1] T. Ferber, C. Garcia-Cely and K. Schmidt-Hoberg, *BelleII sensitivity to long-lived dark photons*, *Phys. Lett. B* **833** (2022) 137373 [arXiv:2202.03452].

[2] M. Duerr et al., *Long-lived Dark Higgs and Inelastic Dark Matter at Belle II*, *JHEP* **04** (2021) 146 [arXiv:2012.08595].

[3] M. Duerr et al., *Invisible and displaced dark matter signatures at Belle II*, *JHEP* **02** (2020) 039 [arXiv:1911.03176].

[4] S. Bähr et al., *The neural network first-level hardware track trigger of the Belle IIexperiment*, *Nucl. Instrum. Meth. A* **1073** (2025) 170279 [arXiv:2402.14962].

[5] K.L. Unger et al., *Data-driven design of the Belle II track segment finder*, 2023 *JINST* **18** C02001.