ACAT 2022



Contribution ID: 201

Type: Poster

Accelerating the DBSCAN clustering algorithm for low-latency primary vertex reconstruction

Thursday 27 October 2022 11:00 (30 minutes)

In this work we present the adaptation of the popular clustering algorithm DBSCAN to reconstruct the primary vertex (PV) at the hardware trigger level in collisions at the High-Luminosity LHC. Nominally, PV reconstruction is performed by a simple histogram-based algorithm. The main challenge in PV reconstruction is that the particle tracks need to be processed in a low-latency environment $O(1 \ \mu s)$. To achieve this an accelerated version of the DBSCAN algorithm was developed to run in a Field Programmable Gate Array (FPGA). A CPU-optimized version of DBSCAN was implemented in C++ to serve as a benchmark for comparison. The CPU version of DBSCAN resulted in an average PV reconstruction latency of 93 μs , while the FPGA firmware only had a latency of 0.73 μs resulting in a 127x speedup. The speedup is a result of running all the input tracks in parallel, which ultimately results in high resource consumption, of up to 48.6 % of the available logic. Most of the logic was attributed to the use of sorting networks that allows for the parallel processing of the input tracks. To tune the firmware for a specific latency and resource usage constraints, the firmware has been parametrized by the number of input tracks to consider at a time. The accelerated DBSCAN method yielded a higher PV reconstruction efficiency when compared to the simpler histogram-based method. As clustering applications are prominent in High Energy Physics, we modified the accelerated DBSCAN algorithm for higher-dimensional datasets.

Significance

In general the DBSCAN clustering algorithm is one of the most flexibility and accurate clustering algorithms available. This work demonstrates that DBSCAN can be utilized in a low-latency environment by using FPGA acceleration. The accelerated algorithm was used to reconstruct primary vertices in collisions at the LHC, however it can be generalized to any clustering application.

References

Experiment context, if any

CMS

Authors: TAPPER, Alex (Imperial College London); ROSE, Andrew; BORGNA, Lucas Santiago (Imperial College (GB)); BARBONE, Marco; BAINBRIDGE, Robert John (Imperial College (GB)); LUK, Wayne

Presenter: BARBONE, Marco

Session Classification: Poster session with coffee break

Track Classification: Track 2: Data Analysis - Algorithms and Tools