Imperial College London

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

Lucas Borgna, Marco Barbone, Andrew Rose, Alex Tapper, Robert Bainbridge, Wayne Luk

Density Based Spatial Clustering of Applications with Noise

DBSCAN can be used to reconstruct the primary vertex.

- Finds the z_0 of the primary vertex.
- Primary vertex defined by the cluster with largest $\sum p_T$.
- Elementary 1D application for development,
- Optimal DBSCAN configuration.
  - minPts = 2.
  - $\varepsilon = 0.15\ cm$.

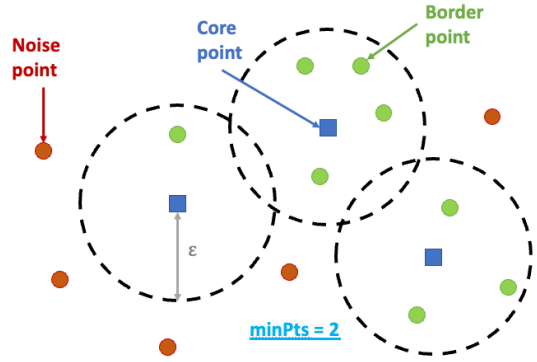

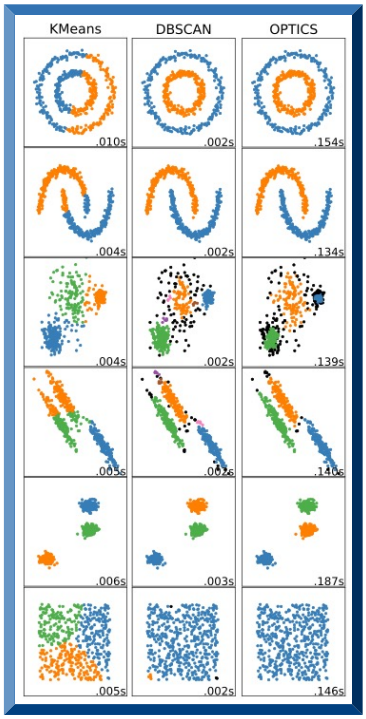

Tracks to reconstruct the primary vertex in high pileup

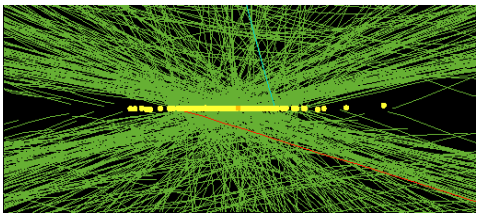

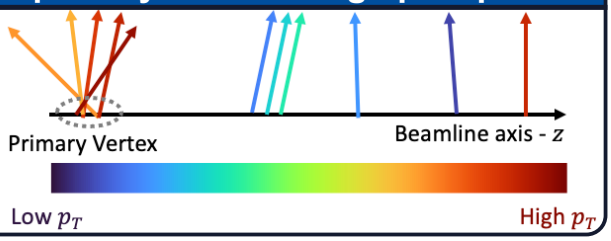

Accelerated DBSCAN

Accelerated DBSCAN for FPGA deployment.

- Sorts tracks along the the z_0 axis.
- Processes all tracks in parallel.
- Sorts boundary indices.
- Sorts cluster by p_T.
- Low latency, but high resource usage.
- Max number of tracks to process is 232.
- Compared to CPU optimized version.
- High resource usage due to sorting networks.
- Maxeler DFE with VU9P FPGA at 100 MHz and tt 200 pileup events.

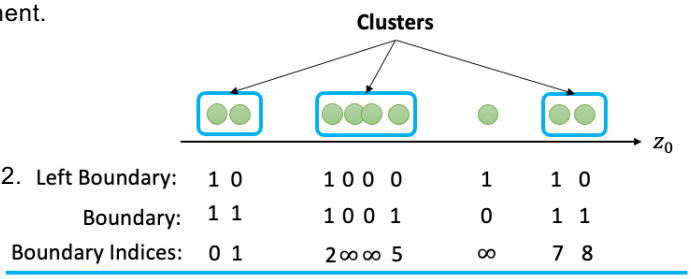

Resource	Usage
Logic	48.6 %
BRAM18	14.1 %
URAM	12.3 %

Latency
FPGA = 0.726 μs
CPU = 92.7 μs
127× Speedup!

Batched DBSCAN

Batched DBSCAN – firmware parametrized by number of tracks to process in parallel.

- Runs the cluster finding part of the Accelerated DBSCAN in fixed size batches.
- Cluster and noise points need to be merged afterwards.
- Customizable latency vs resource usage

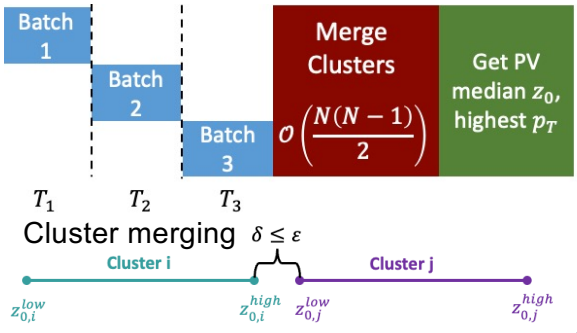

Further Work

- Further Optimizations
- Expand to 2D/3D cases
- Expand to minPts > 2.

GitHub repo