


Fast GPU Nearest Neighbors search algorithms for the CMS experiment at LHC

Alessandro Degano (Università di Torino),
Felice Pantaleo (CERN)

San Francisco 13/10/2016




	2016 conditions	2025 conditions
Luminosity	$5 \text{ nb}^{-1}\text{s}^{-1}$	$50 \text{ nb}^{-1}\text{s}^{-1}$
Pileup	50	140(200)



Phase II upgrades:

- ▶ Full tracker substitution with 4x granularity
- ▶ Calorimeters Endcaps with higher granularity
- ▶ Muon Endcaps with additional GEMs and RPC
- ▶ Level 1 trigger to comply with 750kHz rate
- ▶ High Level Trigger farm with 12x current computational power

The EndCaps will be substituted by a sampling calorimeter.



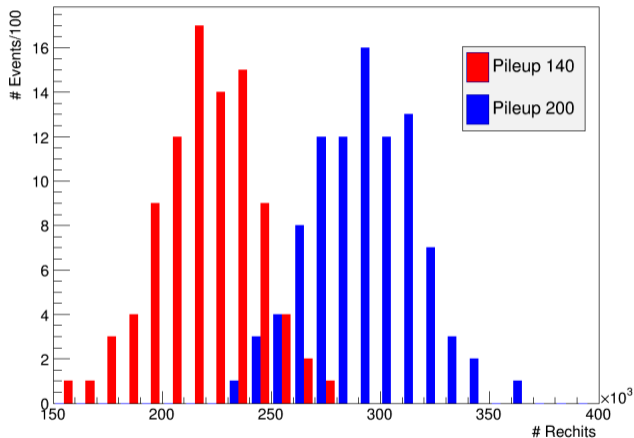
Active material:

- ▶ 0.53 cm^2 or 1.05 cm^2 silicon pad

Passive material:

- ▶ Copper plate Tungsten coated

Rechits distribution



The challenge:

find all the NN of $\sim 300k$ points in the few ms available

Based on FKDTree, a heapified KDTree implemented in TBB, OpenMP, CUDA and OpenCL. Some features critical for the optimal GPU NN search:

- ▶ NO recursion, only iteration
- ▶ Left Balanced Tree → evaluate nodes position, NO query
- ▶ GPU memory management and executions all asynchronous with CUDA Streams

Sequential Version



Asynchronous



A GPU FKDtree implementation (ctd)

FKDtree performs a parallel build using a parallel quickselect: instead of selecting the median, the point which makes the sub-tree tree a heap is selected



Search is performed in parallel and it is iterative and branchless





Simulation of Physics events with high number of electrons in a large P_t spectrum.

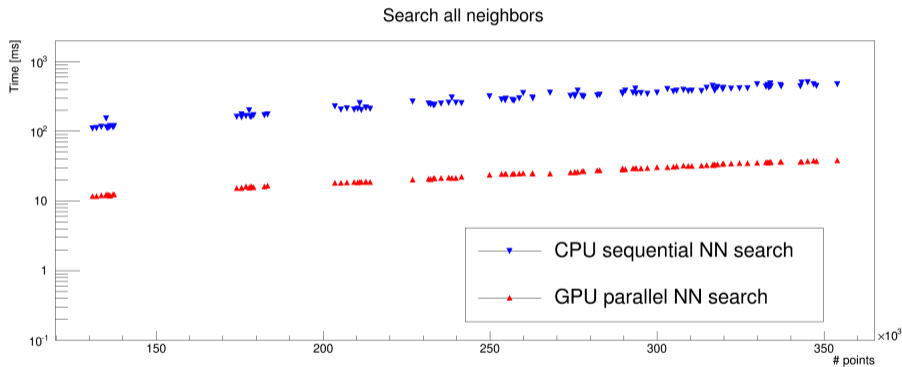
Validation:

- ▶ Compare results with *naive* method, 100% accuracy required


Performance assessment:

- ▶ Measure execution time of tree building vs number of points
- ▶ Measure execution time of sequential ANN
- ▶ Measure execution time of parallel ANN
- ▶ Calculate speedup

10 sets of 10 events, 100 electrons, $P_t:[10, 100]$ GeV/c, $\eta [1.6, 2.9]$



Points (10^3)	135	177	210	238	259	276	293	316	333	345
i7 4771k (ms)	152	199	251	302	349	378	405	444	482	503
NV K40c (ms)	13	16	19	21	25	26	29	33	35	37



Speedup between **9** and **15**



- ▶ HGCAL will gather up to 300k points per event during Phase-II
- ▶ Finding the clusters for those points is computationally intensive
- ▶ By 2023 hardware will have much better performance
- ▶ Many-cores paradigm is here to stay
- ▶ Massive parallelization non-trivial
- ▶ Benefits extend to energy consumption