ALICE TPC Tracking on GPU for Pb-Pb Run in December 2010

D. Rohr[1]

[1]Frankfurt Institute for advanced Studies, Frankfurt, Germany

GPU Tracking

A fast online tracker based on a cellular automation principle and the Kalman filter was developed for the ALICE HLT [1]. The tracker was ported to run on NVIDIA GTX295 GPUs [2]. A dynamic scheduler ensured good utilization of the GPU. On top of that, the introduction of a pipeline allowed to process the tracking on the GPU, the initialization on the CPU, and the DMA transfer in parallel. In 2010 the new Fermi GPU generation was released. For expansion and compatibility reasons the tracker was adapted and optimized for the Fermi architecture.

Commisioning

34 compute nodes equipped with GTX480 GPUs were installed in fall 2010 at the ALICE HLT. The GPU tracker was commissioned and tested during the Pb-Pb run in December. Due to reduced luminosity it was not necessary to process the entire tracking on the GPU. Instead, the run was used to test the GPU tracker in the HLT framework. All tests were successful and the tracker ran stable. Fig. 1 shows an event reconstructed using the GPU tracker. To ensure that the GPU tracker output matches the CPU version, statistics of CPU and GPU runs were compared (See Fig. 2). No difference in CPU and GPU tracking efficiency was found.

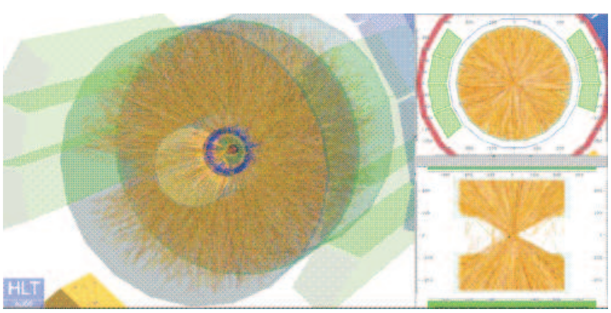

Figure 1: Snapshot of Online Event Display.

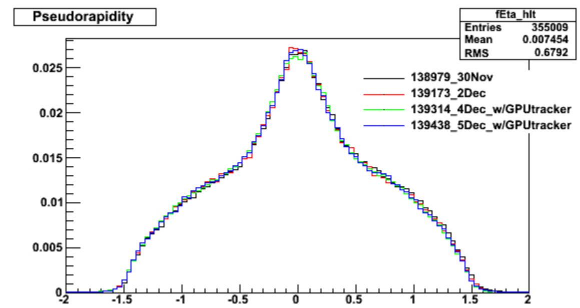

Figure 2: QA Statistiks for GPU Tracker.

Further Optimizations

The original GPU tracker was developed on quad-core Nehalem nodes. However, the GPU compute nodes employ twelve-core Magny-Cours Processors. The overall performance of the Magny-Cours is clearly better. However, when it comes to single-core performance they are inferior to the quad-cores. The GPU tracker only used a single CPU core and was slowed down by the new CPUs. The problem became even more critical due to a change in the output format which is more compute intensive. The pipeline [2] did no longer work well since the GPU was waiting for the single CPU core most of the time.

This was solved by multithreading the CPU part of the GPU-tracker. Contrary to the pipelin, other optimizations originally made for the 295 could directly be used for the GTX580, e.g. the dynamic scheduler ensures a GPU utilization of close to 70%. In late 2010 the GTX580 was released. Fig. 3 shows a performance comparison.

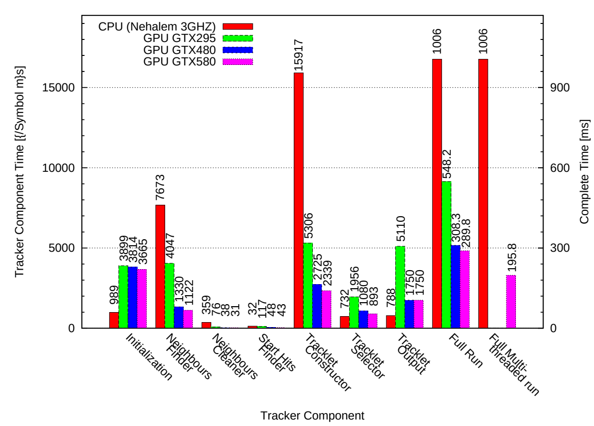

Figure 3: Tracker Performance on different architectures.

References

[1] S. Gorbunov, M. Kretz, D. Rohr, “Fast Cellular Automaton tracker for the ALICE High Level Trigger”, GSI Scientific Report 2009, http://www.gsi.de/scirep2009/

[2] D. Rohr, “ALICE TPC Online Tracking on GPGPU based on Kalman Filter“, University of Heidelberg, Diploma Thesis, http://www.kip.uni-heidelberg.de/Veroeffentlichungen/details.php?id=2035