Scalability of the CA-based track finder in the CBM experiment

I. Kisel[1] *and I. Kulakov*[2,3]

[1]GSI, Darmstadt, Germany; [2]Goethe University, Frankfurt, Germany; [3]University of Kyiv, Ukraine

The main challenge of the CBM experiment is to cope with high interaction rates up to 10^7 collisions per second. The average track density is about 1000 particles per central Au+Au collision. The tracking system of CBM is positioned within a dipole magnet and is build of double-sided silicon strip detectors. In addition to such high input rate and complicated event topology, the full event reconstruction and selection will be done at the First Level Event Selection (FLES) stage. This requires utilization of the full potential of modern many-core CPU/GPU architectures. Since developments of modern processors tend to increase a number of cores rather than a CPU frequency, a good scalability of the event reconstruction is essential.

The core of the track reconstruction for the CBM experiment is the Cellular Automaton (CA) based reconstruction [1] in the Silicon Tracking System (STS). It is the most sophisticated and time consuming part of the event reconstruction. Therefore it should have a good scalability with respect to the number of cores.

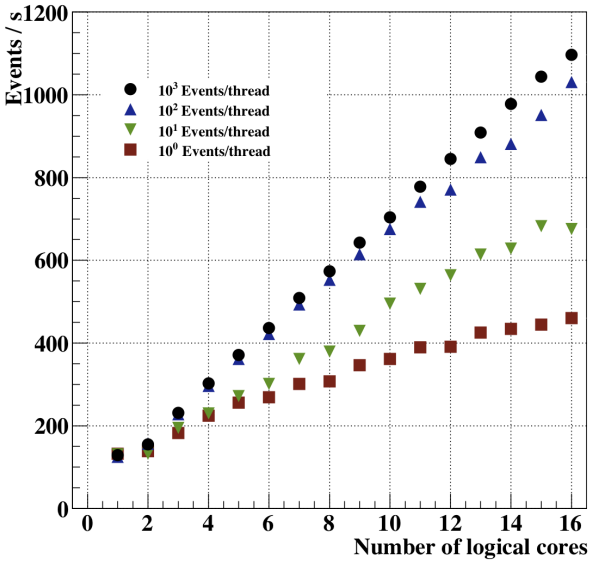

Figure 1: Scalability of the CA track finder for minimum bias events

A stand-alone package [2] was used for the investigation of the CA track reconstruction scalability. Tests were performed with a two Xeon X5550 processors computer having 8 cores in total at 2.7 GHz and with 8 MB L3 cache (GSI lxir039). Each physical core has two logical cores due to the hyper-threading technology. In order to investigate

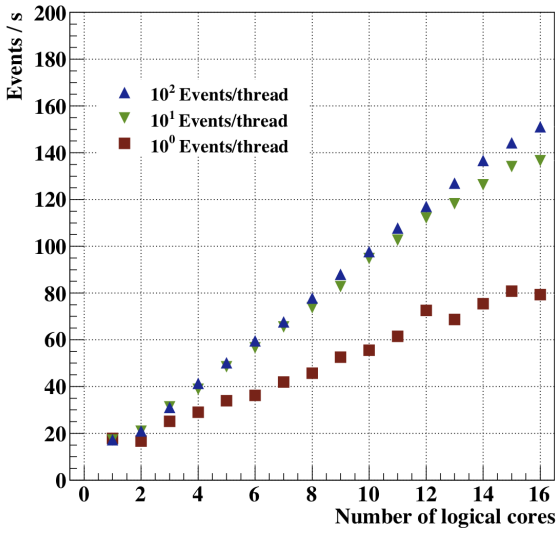

Figure 2: Scalability of the CA track finder for central events

the scalability, track finding was run on various numbers of cores. The Intel Threading Building Blocks (TBB) software package [3] was used for parallel execution between cores.

In Fig. 1 the scalability on the logical cores is shown for minimum bias events. Here, the track reconstruction was parallelized by execution of one thread per (one) logical core. Reconstruction of 1, 10, 100 and 1000 events was executed for each thread. In the same way the track reconstruction was executed for 1, 10 and 100 central events (see Figure 2). The figures show a good linear scalability for large groups of events, while for small groups of events an overhead is observed.

In conclusion, running on a computer with 8 cores the CA based track finder demonstrates the maximum throughput of 150 central or 1100 minimum bias events per second using the Intel Threading Building Blocks. The strong many-core scalability of the CA track finder makes it possible to keep the reconstruction at the event-level parallelism.

References

[1] I. Kisel, Nucl. Instrum. Meth. **A 566** (2006) 85

[2] I. Kisel et al., *CBM Progress Report 2008*, Darmstadt 2009, p. 79

[3] http://www.threadingbuildingblocks.org