



Contribution ID: 173

Type: Oral presentation to parallel session

Parallel track reconstruction in CMS using the cellular automaton approach

Tuesday, 15 October 2013 14:16 (22 minutes)

The Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) at CERN near Geneva/Switzerland is a general-purpose particle detector which led, among many other results, to the discovery of a Higgs-like particle in 2012. It comprises the largest silicon-based tracking system built to date with 75 million individual readout channels and a total surface area of 205 m^2 .

The precise reconstruction of particle tracks from this tremendous amount of input channels is a compute-intensive task and requires an elaborate set of algorithms. A large portion of the reconstruction runtime can be accounted to the Kalman filter-based, iterative track finding and fitting procedures. The combinatorial complexity of these algorithms depends on the number of particle tracks in one event, which itself scales with the number of simultaneous proton-proton collisions in the detector, named pile-up.

The foreseen LHC beam parameters for the next data taking period, starting in 2015, will result in an increase in the number of pile-up interactions. Due to the stagnating clock frequencies of individual CPU cores, new approaches to particle track reconstruction need to be evaluated in order to cope with the computational challenge of the increased number of tracks per event.

Track finding methods that are based on cellular automata (CA) offer a fast and parallelizable alternative to the well-established Kalman filter-based algorithms. The CA proceeds in three steps: i) identify compatible combinations of three energy deposits in adjacent layers of the tracking system; ii) join these fragments to complete track candidates; and iii) select the best tracks for further processing. All computations are rather simple and data local, therefore, they can be implemented in a highly data-parallel fashion. This makes CA a very promising method to be run on modern GPGPUs and hardware accelerators, as well as multi-core CPUs equipped with vector units.

We present a new cellular automaton based track reconstruction, which copes with the complex detector geometry of CMS. By combining the CA with a grid-based data structure, we enable fast access to measured positions on the silicon detectors. Furthermore, we detail the specific design choices made to allow for a high-performance computation on GPU and CPU devices, such as branch-avoidance and the use of computationally inexpensive cut criteria before employing more involved ones. To address the specifics of data parallel programming requirements, we further partition the steps of the CA into more fine-grained sub-tasks and use a two-pass approach for each of them, hence separating computation and memory write access. We conclude by evaluating the physics efficiency, as well as computational properties of our implementation on various hardware platforms.

Summary

Primary authors: FUNKE, Daniel (KIT - Karlsruhe Institute of Technology (DE)); HAUTH, Thomas (CERN); INNOCENTE, Vincenzo (CERN)

Co-authors: Mr SCHIEFERDECKER, Dennis (Karlsruher Institut für Technologie); QUASt, Gunter (KIT - Karlsruhe Institute of Technology (DE)); Prof. SANDERS, Peter (Karlsruher Institut für Technologie)

Presenter: FUNKE, Daniel (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Software Engineering, Parallelism & Multi-Core

Track Classification: Software Engineering, Parallelism & Multi-Core