Contribution ID: **210** Type: **Parallel**

Development and Evaluation of Vectorised and Multi-Core Event Reconstruction Algorithms within the CMS Software Framework

Thursday 24 May 2012 17:00 (25 minutes)

The processing of data acquired by the CMS detector at LHC is carried out with an object-oriented C++ software framework: CMSSW. With the increasing luminosity delivered by the LHC, the treatment of recorded data requires extraordinary large computing resources, also in terms of CPU usage. A possible solution to cope with this task is the exploitation of the features offered by the latest microprocessor architectures. Modern CPUs present several vector units, the capacity of which is growing steadily with the introduction of new processor generations. Moreover, an increasing number of cores per die is offered by the main vendors, even on consumer hardware. Most recent C++ compilers provide facilities to take advantage of such innovations, either by explicit statements in the programs' sources or automatically adapting the generated machine instructions to the available hardware, without the need of modifying the existing code base. Programming techniques to implement reconstruction algorithms and optimised data structures are presented, that aim to scalable vectorization and parallelization of the calculations. One of their features is the usage of new language features of the C++11 standard. Portions of the CMSSW framework are illustrated which have been found to be especially profitable for the application of vectorization and multi-threading techniques. Specific utility components have been developed to help vectorization and parallelization. They can easily become part of a larger common library. To conclude, careful measurements are described, which show the execution speedups achieved via vectorised and multi-threaded code in the context of CMSSW.

Summary

The processing of data acquired by the CMS detector at LHC is carried out with an object-oriented C++ software framework: CMSSW. With the increasing luminosity delivered by the LHC, the treatment of recorded data requires extraordinary large computing resources, also in terms of CPU usage. A possible solution to cope with this task is the exploitation of the features offered by the latest microprocessor architectures. Modern CPUs present several vector units, the capacity of which is growing steadily with the introduction of new processor generations. Moreover, an increasing number of cores per die is offered by the main vendors, even on consumer hardware. Most recent C++ compilers provide facilities to take advantage of such innovations, either by explicit statements in the programs' sources or automatically adapting the generated machine instructions to the available hardware, without the need of modifying the existing code base. Programming techniques to implement reconstruction algorithms and optimised data structures are presented, that aim to scalable vectorization and parallelization of the calculations. One of their features is the usage of new language features of the C++11 standard. Portions of the CMSSW framework are illustrated which have been found to be especially profitable for the application of vectorization and multi-threading techniques. Specific utility components have been developed to help vectorization and parallelization. They can easily become part of a larger common library. To conclude, careful measurements are described, which show the execution speedups achieved via vectorised and multi-threaded code in the context of CMSSW.

Authors: PIPARO, Danilo (CERN); Mr HAUTH, Thomas (KIT - Karlsruhe Institute of Technology (DE)); INNOCENTE, Vincenzo (CERN)

Presenter: Mr HAUTH, Thomas (KIT - Karlsruhe Institute of Technology (DE))

Session Classification: Software Engineering, Data Stores and Databases

Track Classification: Software Engineering, Data Stores and Databases (track 5)