



Contribution ID: 219

Type: Oral presentation to parallel session

Speeding up HEP experiments' software with a library of fast and autovectorisable mathematical functions

Tuesday, 15 October 2013 17:47 (20 minutes)

During the first four years of data taking at the Large Hadron Collider (LHC), the simulation and reconstruction programs of the experiments proved to be extremely resource consuming. In particular, for complex event simulation and reconstruction applications, the impact of evaluating elementary functions on the runtime is sizeable (up to one fourth of the total), with an obvious effect on the power consumption of the hardware dedicated to their execution. This situation clearly needs improvement, especially considering the even more demanding data taking scenarios after this first LHC long shut down. A possible solution to this issue is the VDT (VectorisD maTh) mathematical library. VDT provides the most common mathematical functions used in HEP in an open source product. The functions' implementations are fast, can be inlined, provide an approximate accuracy and usable in vectorised loops. Their implementation is portable across platforms: x86 and ARM processors, Xeon Phi coprocessors and GPGPUS. In this contribution, we describe the features of the VDT mathematical library, showing significant speedups with respect to the Libm library and comparable accuracies. Moreover, taking as examples simulation and reconstruction workflows ran in production by the experiments, we show the benefits of the usage of VDT in terms of runtime reduction and stability of physics output.

Primary author: PIPARO, Danilo (CERN)

Co-authors: HAUTH, Thomas (CERN); INNOCENTE, Vincenzo (CERN)

Presenter: PIPARO, Danilo (CERN)

Session Classification: Software Engineering, Parallelism & Multi-Core

Track Classification: Software Engineering, Parallelism & Multi-Core