Contribution ID: **103** Type: **Oral**

Evolution of the ATLAS Software Framework towards Concurrency

Tuesday, 2 September 2014 17:25 (25 minutes)

The ATLAS experiment has successfully used its Gaudi/Athena software framework for data taking and analysis during the first LHC run, with billions of events successfully processed. However, the design of Gaudi/Athena dates from early 2000 and the software and the physics code has been written using a single threaded, serial design. This programming model has increasing difficulty in exploiting the potential of current CPUs, which offer their best performance only through taking full advantage of multiple cores and wide vector registers. Future CPU evolution will intensify this trend, with core counts increasing and memory per core falling. Maximising performance per watt will be a key metric, so all of these cores must be used as efficiently as possible.

In order to address the deficiencies of the current framework, ATLAS has embarked upon two projects: first, a practical demonstration of the use of multi-threading in our reconstruction software, using the GaudiHive framework; second, an exercise to gather requirements for an updated framework, going back to the first principles of how event processing occurs.

In this paper we report on both these aspects of our work. For the hive based demonstrators, we discuss what changes were necessary in order to allow the serially designed ATLAS code to run, both to the framework and to the tools and algorithms used. We report on what general lessons were learned about the code patterns that had been employed in the software and which patterns were identified as particularly problematic for multi-threading. These lessons were fed into our considerations of a new framework and we present preliminary conclusions on this work. In particular we identify areas where the framework can be simplified in order to aid the implementation of a concurrent event processing scheme. Finally, we discuss the practical difficulties involved in migrating a large established code base to a multi-threaded framework and how this can be achieved for LHC Run 3.

Primary authors: WYNNE, Benjamin Michael (University of Edinburgh (GB)); LEGGETT, Charles (Lawrence Berkeley National Lab. (US)); STEWART, Graeme Andrew (University of Glasgow (GB)); JONES, Roger (Lancaster University (GB))

Presenter: JONES, Roger (Lancaster University (GB))

Session Classification: Computing Technology for Physics Research

Track Classification: Computing Technology for Physics Research