Workshop on Concurrent Frameworks in the Multi-Core Era
Discussion session notes

This is a summary of the final discussion session during the Workshop on Concurrent Frameworks in the Multi-Core Era held at FNAL on November 21st -22nd 2011. The goal was to explore the possibility that interested High Energy Physics (HEP) institutions and projects collaborate on concurrent frameworks R&D.

Multi-core computing technology is developing quickly and it is difficult to predict exactly how systems will evolve toward the exascale era; this poses a problem on deciding software development directions. We are far from even maximally using current hardware. Optimization work on current frameworks and applications would improve computing performance of both current and multi-thread implementations.

We know that software applications should be improved in the direction of more vectorization and data locality; we should investigate how to factorize the HEP software in parts to run serially and parts that can be parallelized to improve computing performance.

The LHC experiments expressed their wish to make software changes as transparent as possible to their developers and users. In addition, any modification should be done adiabatically and without affecting the physics performance. CMS is supportive of exploratory R&D and is ready to contribute to the effort. ATLAS prefers to focus on the multi-process solution. LHCb has a workshop on computing in December 2011 and this presents a good opportunity to discuss their interest and potential level of involvement. In contrast, future experiments, including those exploring the intensity frontier, are in a position to adopt quickly more radical changes to their software.

The concrete proposal is to develop demonstrators for different capabilities in a small scale with clear deliverables and metrics, and work in short cycles a few months long. Eventually, success in the small scale for different areas will bring more developers to the effort, critical mass would be reached, and a plan for a prototype of a HEP framework could emerge.

There is agreement that the interested parties should work together to develop the demonstrators and agree minimally on technology so that they can share code and compare results. It is not clear, however, that the ultimate goal of developing a common framework for all HEP experiments is achievable for running experiments. The R&D code we develop would evolve into elements to use in re-engineering of software tools such as Geant4 and ROOT, and frameworks.

Future Activities

The LHC experiments have a timeline for getting framework-related activities completed that is dominated by the window of opportunity offered by the long shutdown planned for 2013-2014. The goal of the common effort emerging from this workshop is to make rapid progress on exploratory R&D activities during 2012, find a common way forward, and aim to take advantage of the window of

opportunity offered by the long LHC shutdown. The time scale associated with the intensity frontier experiments is slightly longer, and the concentration early on will be more in the area of online, real-time filtering and processing.

There were several opinions expressed about what can be shared between the different groups participating in the workshop. There was consensus on the benefit of:

- Coming up with common goals
- Sharing knowledge
- Sharing a common concurrency model

Several of the test projects, both planned and ongoing, share a common theme. The major concepts that were similar are using task parallelism (at both the event and module level), and retaining an event-like data structure to communicate products. Different tools were used to test these ideas.

Complete freedom for implementing demonstrators may not be a good idea. Instead, it would be advantageous to come up with a common semantic interface that facilitates comparison of different approaches.

Current activities focus on building demonstrators in order to study different models, technologies, and tools. A number of demonstrators were mentioned during the workshop and include the following:

1. Object sharing (ATLAS Paolo Calafiura). Investigate moving data objects between processes.
2. Multiple process management (ATLAS Paolo Calafiura). Investigate better tools to manage processes and inter-process communication in Athena.
3. Performance tools (FNAL/CMS Lassi Tuura, Chris Jones). Report on tools for measuring multi-core application performance and debugging them. (An example is Instruments from MacOS.) Intel tools such as Thread Checker should be included, and standard options for compilation and linking determined.
4. Investigate the effect of improved data locality in a realistic case (FNAL-Ph. Canal) in the context of Geant4 detector simulation.
5. Scheduling work (FNAL/CMS Chris Jones, FNAL CET group). Further investigate use of libdispatch (GCD). Investigate use of OpenMP, TBB, and CnC for task parallelism at the event and sub-event level.
6. Parallelization within modules (FNAL CET group). This includes investigation of 2D Hough transforms and waveform compression algorithms.
7. 'Whiteboard' service implementation (CERN Pere Mato, Benedikt Hegner), including random number generators, geometry and magnetic field interpolation, and resolving service interface bottleneck issues.
8. Histograms (FNAL/CMS Lassi Tuura). Address concurrent building of histograms, and TFileService.
9. Detector simulation (FNAL/CERN R&D project - Philippe Canal, Federico Carminati). Study data locality and vector processing in geometry, particle transport etc.

10. Evaluation of framework (FNAL/CMS Chris Jones, Lassi Tuura). Define performance metrics (event throughput, scaling). Define standard set of modules timings and configurations, data samples, and set of module dependencies. Define torture tests that will help determine performance limits.
11. Multithreaded I/O (FNAL/CMS Chris Jones). Continue the investigation already initiated on performance improvement of the ROOT I/O sub-system by scheduling the different operations (disk I/O, compression, object serialization) concurrently.
12. Tracking reconstruction in a complex geometry, demonstrator using GPU. (CERN Vincenzo Innocente and student.)
13. Study solutions that make use of virtualisation technologies (CERN Predrag Buncic). Goals are to develop a test suite to continuously monitor the performance of selected LHC applications in a VM context and in longer term investigate possible use of GPGPU/MIC from VM environment.
14. Use of the 'Go' language to schedule modules (ATLAS - S. Binet). The goal is investigate task parallelism at the event and sub-event level by coding the top-level steering of the framework in Go.

Tools

There are tools for debugging (gdb does not work but Totalview does) and profiling (IgProf). Openlab may want to collaborate with the task of checking whether the multi-thread code is behaving properly. Fermilab may need to establish a relationship with Intel similar to that of CERN-Openlab with its industrial partners.

Communication

We agreed to create forums to communicate and exchange information and results. We will use the rd-multicore@cern.ch egroup list for e-mail communication and meet every two weeks, probably on Wednesdays at 17 hs CERN time to discuss progress. We would also hold face-to-face workshops every 3-4 months.