2018 IEEE 14th International Conference on e-Science

Simulating HEP Workflows on Heterogeneous Architectures

1st Charles Leggett
on behalf of the ATLAS collaboration
Lawrence Berkeley National Laboratory
Berkeley, CA 94720 USA
cgleggett@lbl.gov

2st Illya Shapoval
on behalf of the ATLAS collaboration
Lawrence Berkeley National Laboratory
Berkeley, CA 94720 USA
ishapoval@lbl.gov

***Abstract*—The next generation of supercomputing facilities, such as Oak Ridge's Summit and Lawrence Livermore's Sierra, show an increasing use of GPGPUs and other accelerators in order to achieve their high FLOP counts. This trend will only grow with exascale facilities. In general, High Energy Physics computing workflows have made little use of GPUs due to the relatively small fraction of kernels that run efficiently on GPUs, and the expense of rewriting code for rapidly evolving GPU hardware. However, the computing requirements for high-luminosity LHC are enormous, and it will become essential to be able to make use of supercomputing facilities that rely heavily on GPUs and other accelerator technologies.**

ATLAS has already developed an extension to AthenaMT, its multithreaded event processing framework, that enables the non-intrusive offloading of computations to external accelerator resources, and is developing strategies to schedule the offloading efficiently. Before investing heavily in writing many kernels, we need to better understand the performance metrics and throughput bounds of the workflows with various accelerator configurations. This can be done by simulating the workflows, using real metrics for task interdependencies and timing, as we vary fractions of offloaded tasks, latencies, data conversion speeds, memory bandwidths, and accelerator offloading parameters such as CPU/GPU ratios and speeds.

We present the results of these studies, which will be instrumental in directing effort to make the ATLAS framework, kernels and workflows run efficiently on exascale facilities.

***Index Terms*—exascale, ATLAS, simulator, accelerator, GPGPU**

Lawrence Berkeley National Lab work was supported by the Office of Science, Office of High Energy Physics, of the U.S. Department of Energy under contract DE-AC02-05CH11231

978-1-5386-9156-4/18/$31.00 ©2018 IEEE
DOI 10.1109/eScience.2018.00087