

21st International Conference on Computing in High Energy and Nuclear Physics (CHEP2015)



Contribution ID: 536

Type: **oral presentation**

Simulation of LHC events on a million threads

Thursday, April 16, 2015 11:45 AM (15 minutes)

Demand for Grid resources is expected to double during LHC Run II as compared to Run I; the capacity of the grid, however, will not double. The HEP community must consider how to bridge this computing gap. Two approaches to meeting this demand include targeting larger compute resources, and using the available compute resources as efficiently as possible. Argonne's Mira, the fifth fastest supercomputer in the world, can run roughly five times the number of parallel processes that the ATLAS experiment typically uses on the Grid. We have ported Alpgen, a serial x86 code, to run as a parallel application under MPI on the Blue Gene/Q architecture. By analysis of the Alpgen code, we reduced the memory footprint to allow running 64 threads per node, utilizing the four hardware threads available per core on the PowerPC A2 processor. Event generation and unweighting, typically run as independent serial phases, are coupled together in a single job in this scenario, reducing intermediate writes to the filesystem. By these optimizations, we have successfully run LHC proton-proton physics event generation at the scale of a million threads, filling two-thirds of Mira.

Primary authors: BENJAMIN, Doug (Duke University (US)); CHILDERS, Taylor (Argonne National Laboratory (US)); LE COMPTE, Thomas (Argonne National Laboratory (US)); URAM, Tom (urn:Google)

Presenter: CHILDERS, Taylor (Argonne National Laboratory (US))

Session Classification: Track 8 Session

Track Classification: Track8: Performance increase and optimization exploiting hardware features