

Scaling the Geneva Library Collection to large HPC clusters

Tuesday, March 15, 2016 2:00 PM (30 minutes)

The Geneva Library Collection was designed to allow parametric optimization of demanding scientific and engineering problems in distributed and parallel computing environments. It has been successfully tested with a code relevant for hadron physics and is in production use in the automotive industry. Scalability of Genevas distributed execution model may be of particular importance for deployment scenarios involving complex models and simulations. Where these feature a non-trivial quality surface and many parameters to be optimized, as well as long running evaluation functions, computing demands may rise to hundreds or even thousands of cores, used over days or weeks. Scalability and stability in this case depend on many factors, some of which must be tuned in highly distributed environments. GSI Darmstadt has teamed up with the KIT spin-off Gemfony scientific to modify the Geneva library in such a way that it allows distributed execution on the new GSI Kronos Cluster, potentially involving more than thousand cores. The presentation introduces the steps that had to be taken both on the code- and the cluster-side to make Geneva scale to the desired level, and involves a practical demonstration on the target cluster.

Primary authors: KNEDLIK, Jan (GSI Darmstadt); Dr SCHWARZ, Kilian (GSI Darmstadt); Prof. MATTHIAS, Lutz (GSI Darmstadt); Dr BERLICH, Ruediger (Gemfony scientific UG (haftungsbeschraenkt))

Co-authors: Dr GARCIA, Ariel (Gemfony scientific UG (haftungsbeschraenkt)); Dr GABRIEL, Sven (Gemfony scientific UG (haftungsbeschraenkt))

Presenters: Dr SCHWARZ, Kilian (GSI Darmstadt); Dr BERLICH, Ruediger (Gemfony scientific UG (haftungsbeschraenkt))

Session Classification: Massively Distributed Computing and Citizen Sciences Session

Track Classification: Massively Distributed Computing and Citizen Sciences