ExaLearn – GenTen Tensor Software ECP Milestone

Hemanth Kolla, Eric Phipps, Michael Wolf

July 2022

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Objective:
The objective of this milestone was to finish integrating GenTen tensor software with combustion application Pele using the Ascent in situ analysis software, partnering with the ALPINE and Pele teams. Also, to demonstrate the usage of the tensor analysis as part of a combustion simulation.

Report:
GenTen is a C++ software library for generalized tensor decompositions and provides performance portability via Kokkos. Within ExaLearn, an anomaly detection algorithm based on decomposition of cokurtosis (4^{th}-order joint moment) tensor was targeted in GenTen to provide anomaly detection capability to the Combustion-Pele AD code. In-situ deployment of the GenTen cokurtosis kernel was enabled through the Ascent software from the ALPINE project. The objective of this activity was to complete the GenTen <> Ascent <> Pele integration and demonstrate anomaly detection with a realistic combustion simulation with Pele. Prior to this effort, the cokurtosis kernel was implemented directly within Ascent, which was exposed to Pele as an Ascent "filter". With this activity, the Ascent now exposes the cokurtosis kernel to Pele from an external library (GenTen), establishing a template for AD codes invoking analyses kernels in-situ from third-party independently developed libraries.

Another objective was also to demonstrate the readiness of this integration on pre-exascale machines. To this end efforts were focused on porting Genten to heterogeneous architectures with NVIDIA (Perlmutter) and AMD (Crusher) GPUs. The implementation of the cokurtosis algorithms uses a combination of Kokkos-compatible linear algebra kernels through the KokkosKernels layer, and vendor optimized TPL kernels, i.e., CUDA/CUBLAS/CUSOLVER and ROCm/ROCBLAS/ROCSOLVER.

The implementation in GenTen, porting to Crusher and Perlmutter, and integrating with Ascent/Pele were successfully completed. Scalability of the cokurtosis kernel were studied within an ALPINE milestone (STDM16-23). A demonstration combustion simulation is also documented in that milestone report. The scalability study showed good scaling up to ~200 GPUs, before some degradation of scalability.

Related Milestones:
The GenTen development efforts are closely tied to three ALPINE milestones, and are documented in those milestone reports:

- STDM16-21
- STDM16-22
- STDM16-23

Artifacts:

- The GenTen library is available at https://gitlab.com/tensors/genten/. The cokurtosis kernel ported to, and tested on, Crusher/Summit/Perlmutter are available on the master branch, SHA 16c12de91dd77a1191575ca92509c67fac4b3ea6.
- Building instructions for compiling and building Ascent/Pele/Genten are available through the ALPINE project repository, https://github.com/Alpine-DAV/ascent/wiki/Cokurt---GenTen---Pele.
- Algorithmic innovations that improved the performance of the cokurtosis computation were presented as an accepted paper recently at The Platform for Advanced Scientific Computing

(PASC) Conference: Li, Zitong, and Kolla, Hemanth, and Phipps, Eric, "*Parallel Memory-Efficient Computation of Symmetric Higher-Order Joint Moment Tensors*."

Next Steps:
An integration KPP-3 milestone involving ALPINE <> ExaLearn <> Pele is being targeted for Q1-FY23. A KPP-3 relevant combustion simulation, running at scale on Crusher, will be targeted for demonstrating in-situ anomaly detection. Preparations are currently underway to set up a combustion simulation that is of scientific interest, while cleaning up the integration framework and streamlining the compiling and building of the coupled software libraries.