SAND2020-10629R

U.S. DEPARTMENT OF ENERGY

SANDIA NATIONAL LABORATORIES

SNL ADTM

FY20Q4 report for ATDM AD projects to ECP

September2020

Prepared by: Gabrielle Trujillo

Prepared for:

ECP Confluence updates

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-NA-0003525.

SAND2020-6640 R

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology and Engineering Solutions of Sandia, LLC.

Activities:

- Kokkos: Complete conversion of remaining Intrepid2 projection algorithms to Kokkos.
- Kokkos: A prototype for Kokkos-Graphs was developed which uses CUDA graphs under the hoods. It is currently under code review for integration into the mainline Kokkos release.
- VTK-m performance evaluation: Successfully compiled Catalyst with VTK-m enabled filters for the Vortex GPU hardware.

Accomplishments:

- Kokkos: Completed embedded error estimation in EMPIRE-Fluid, added testing of Runge-Kutta App Actions in EMPIRE-Fluid and EMPIRE-Hybrid, fixed bugs and added several more tests of Tempus features.
- Kokkos: Collaborated with NVIDIA on NVSHMEM and demonstrated several orders of magnitude improvement using aggregation and caching schemes developed by the Kokkos team for a PGAS implementation of a CGSolve. This work potentially could be used to support AMD GPUs too, since it sits on top of native ibverbs not actually using any NVIDIA libraries. But the code is based on the open source NVSHMEM implementation.
- Kokkos: Supported ATDM and IC code Kokkos adoption, debugging and optimization through in house consulting, participation in design discussions, and optimization of Kokkos capabilities.
- Kokkos: A prototype for Kokkos-Graphs was developed which uses CUDA graphs under the hoods. It is currently under code review for integration into the mainline Kokkos release.
- VTK-m: Established a "driver" program that can read CGNS and Exodus files and drive Catalyst as if a simulation were producing the data. This program as been compiled on Vortex. We have also successfully compiled Catalyst with VTK-m enabled filters for the Vortex GPU hardware. However, we are currently working through build and script issues for executing the VTK-m code.

Next Steps:

- Kokkos: Engage the C++ standards committee to further the adoption of successful Kokkos concepts into the standard and provide feedback on proposed concurrency mechanisms such as the executors proposal. In particular develop a prototype for the proposed linear algebra capabilities with executor integration, in order to allow an adoption of linear algebra into the standard by 2023.

Outreach:

- Several OS/R team members attended PathForward final reviews and the September El Capitan quarterly review.
- OS/R: MPI Forum work has progressed in several key improvements for MPI 4.0. Partitioned communication has been officially added and work continues on improvements for GPU communication. We have also contributed to multiple improvements on collective operations, feedback on dynamic sessions in MPI and work has gone into the improvements for 64-bit support for very large messages. We have also contributed to reviewing and providing extensive feedback on updates to the underlying semantics of the MPI specification and have contributed some text as well. We actively participate in multiple working groups and co-lead work on partitioned communication and advanced collectives and persistent operations. We are also actively engaged

as a key partner with the hybrid programming models group, contributing work on efficient native MPI API support for GPU architectures. OpenMP Language Committee work has focused on preparing for the release of version 5.1 of the specification. Our contributions include leading the task parallelism subcommittee and editing/revising two chapters of the specification to ensure correctness. The 5.1 public comment draft was released in August.