SAND2019-7477R

SANDIA NATIONAL LABORATORIES

SNL ADTM

FY19Q3 report for ATDM ST projects to ECP

June 27, 2019

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.

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

2.3.1.04 ATDM PMR

Accomplishments and Progress:

PERFORMANCE TO PLAN / MEETS

- The Kokkos team organized a 3-Day Kokkos Usergroup meeting at Sandia, which was attended by up to 80 developers (up to 55 in person, and typically 15-30 online). This meeting allowed Kokkos users to present their experience in using Kokkos and thus allowed the Kokkos team to collect valuable feedback, while fostering exchange of ideas between different user projects. Kokkos team collaborating with AMD via the Path Forward effort. This is advancing Kokkos support for [AMD] backends, necessary for future exascale/HPC architectures.
- Direct improvements in performance in Kokkos to accelerate Sparc and Empire codes. As Sparc was already well optimized, team achieved 10% improvement there. Team achieved 4x improvement with Empire when running on GPUs, mostly by identifying performance issues due to missing GPU offload and transfer issues with UVM memory. Focus on ATDM apps requirements remains a top priority for the team's ongoing activities. DARMA team successfully completed work on DARMA-MPI interoperability mini-apps with PIC move kernel in DARMA and solver in MPI adapted from HPCCG.
- Kokkos team continues to have significant impact on C++ standard and taking leadership position on mdspan (shorter term goal) and Executors (longer term goal). Team is now leading the effort to introduce linear algebra capabilities into the C++ standard.

EXCEEDS (these are not planned exceeds)

- Team is having significant impact in the C++ standards community, and taking leadership positions on components viewed as critical for the HPC community.

- **Q3: C++ Standards Work and Kokkos Community Growth**
  - Members of the Kokkos team attended the C++ standards meeting in March.
  - While the committee is mostly focused on finishing up C++20 no concrete progress was made on mdspan.
  - Chances of early inclusion into the C++23 draft increased significantly.
  - Committee members from the hardware vendors asked Sandia to take the lead on a proposal which will cover the needs of HPC and machine learning.
  - Status: complete

2.3.2.04 ATDM Tools

Accomplishments and Progress:

PERFORMANCE TO PLAN / MEETS

- EMPIRE and SPARC are running on TrInity and Sierra. Collaborative work with the solvers team has resulted in dramatic improvements in the scalability of the EMPIRE solver. These are reported in the solver deliverables.
- Developing Spack-based build of third-party libraries needed by Trilinos, EMPIRE, GEMMA, and SPARC.

EXCEEDS (these are not planned exceeds)

- EMPIRE and SPARC have been ported to Astra after to move to the SRN in support of next year's L1 milestone.

- **Q3: Performance Analysis studies for ATDM SPARC code for ATS-2/Sierra**
    - Kokkos Deep Dive for APRC provided the breakdown of kernel execution times.
    - Prototype Identified several kernels where performance has been slower than expected on ATS2 environments. In some cases, these have related to MPI communication issues and in others they are kernel problems. The most important fix identified here is that communication to MPI ranks ("self-communication") is incredible slow on ATS2. By changing this code to do internal copies between buffers we have been able to significantly improve performance (times down from several seconds to fractions of a second).
    - OpenMPI 4.0.1, CUDA 10.1 and GCC 7.2.0 environment for Trilinos and ATDM application testing. We identified several bugs with GCC 7.4.0 which meant that this was not acceptable for moving forward and so this has moved back to GCC 7.2.0. The results show a mixture of bugs in Trilinos, some of which are bugs in the NVIDIA compiler, that we are providing work arounds for. This will underpin the CUDA environment for ATDM Level 1 milestones in FY20.
    - Status: Complete
- **Q3: Improve Trilinos integration and multi-platform CI testing for ATDM apps**
    - Added more platforms and nightly ATDM Trilinos builds to protect and support EMPIRE, GEMMA and SPARC.
    - Added automated testing of SPARC VOTD (version of the day) with Trilinos VOTD for CUDA 9.2 on ATS-2 clone.
    - Addressed several breakages of Trilinos w.r.t. SPARC and allowed several Trilinos updates (much more frequently).
    - Addressed several issues for EMPIRE.
    - Collaborated with Si Hammond on Trilinos testing with CUDA 10.10 with GCC 7.2.0 and set up nightly testing for ATDM Trilinos configuration.
    - Status: On Schedule

2.3.3.04 ATDM Math Libraries

PERFORMANCE TO PLAN / MEETS

- Completed the forward UQ and dimensionality reduction for SPARC 1-D thermal ablation (no CFD) simulations (this work was shared at the ATDM deep dive), continuing to work towards the 2-D simulations with coupled CFD in Q4
- The KK-based solvers supporting the SPARC L2 milestone have achieved 2X speedup in strong scaling on GPUs. GPU simulations are now 3.2x faster than the CPU ones.

EXCEEDS

- The scalable solvers team is running multigrid SPARC solves on 65536 cores on Trinity, which represents the largest multigrid solves ever performed by the scalable solvers team with turbulence. [the original goal was to get up to 4000 nodes on Trinity]
- Irina Tezaur and Jerry Watkins are standing up an automated performance test using MiniEM with the small and medium blob meshes. Post-processing will be done using Jupyter notebooks. [This is unplanned work for Q3]

- **Q3: Profile and analyze performance and scalability of electromagnetics (EM) linear solver for relevant simulations on GPU architectures.**
    - Analysis of EMPIRE and MiniEM performance on GPUs at Sandia and LLNL is ongoing. Christian Glusa has been assessing performance of different multigrid smoothers.
    - Status: Complete
- **Q3: Develop, integrate and support agile components in ASC/ATDM applications.**
    - Integrate Tempus into EMPIRE-Fluids and demonstrate on problem of interest. Implement hierarchic basis support for a single cell topology type in Intrepid2. Complete and release conservative projections in Intrepid2.
    - Status: Complete
- **Q3: Provide sensitivity analysis, UQ, and calibration support to SPARC and EMPIRE teams for validation calculations relating to the FY19 L2 milestone.**
    - Developed Dakota workflow in support of the EMPIRE L2 milestone. This workflow will be used for subsequent sensitivity analysis and UQ calculations on the B-dot L2 milestone problem in Q4. Completed forward UQ and dimensionality reduction for SPARC 1-D thermal ablation (no CFD) simulations. Working towards 2-D simulations with coupled CFD in Q4. Short-listed a set of SPARC inputs (aerodynamic and heating in nature) whose uncertainties will be propagated through when validating against flight-test data. Test simulations of the flight test ongoing, as well as the process of wrapping Dakota around the flight simulation.
    - Status: Complete
- **Q3: Deliver performance portable kernels needed by multigrid methods for EMPIRE needs, especially multithreaded smoothers.**
    - We have optimized KK based solvers to support SPARC L2 milestone. The GPU based solvers have become ~2x better on the strong scaling case. The performance improvements are even better when larger problems are used. This has enabled SPARC GPU simulations to be up to 3.2x faster on than corresponding CPU based simulations. We have optimized kernels such as the Sparse Matrix Vector multiply on Volta GPUs for EMPIRE use cases. We have also started optimization of kernels for EMPIRE use cases, especially smoothers used by multigrid methods.
    - Status: Complete
- **Q3: Demonstrate scalability of electromagnetics linear solver for relevant simulations on Trinity-like architectures.**
    - Show strong and weak scaling (up to 4000-ish Trinity nodes) of MueLu for a steady-state hypersonic RANS physics for a stockpile RV. Compare strong and weak scaling of MueLu to a block tridiagonal solve.
    - Several tweaks have been incorporated into the multigrid coarse mass stabilization technique. Have successfully run some large-scale SPARC simulations with SA turbulence on 65536 cores on Trinity. We are currently running the multigrid solver for these SA runs, which would be the largest MG solves that we have performed with turbulence.
    - Status: Complete

2.3.4.04 ATDM Data and Viz

PERFORMANCE TO PLAN / MEETS

- Delivered in situ, post-processing, and artistically-rendered visualization for LSCI runs and ASC PI meeting.
- Continue to address scaling issues for third-party I/O libraries. A number of issues were identified and addressed by vendors.
- Completed burst-buffer enabled IOSS library. Testing underway. Advertising availability to ATDM and IC teams.
- ML for meshing tested graph-classification methods and found topological methods to be more accurate
- ML for system software -- project is influencing data-center collection information for new ASTRA ARM supercomputer. A novel aspect of the work is to closely couple facility-level information with the operating system running on the supercomputer, enabling two-way exchange of information and new capabilities including power forecasting and workload shifting to more effectively leverage advanced capabilities.

EXCEEDS

- ML for meshing: A director's challenge ESRF LDRD project was proposed and awarded for FY20 that leverages some of the ideas developed in this project.

- **Q3: IO-BB: Provide BB-enabled IOSS library to Sandia apps for production testing validation.**
    - BB-enabled IOSS library is available for use by the Sierra applications and Trilinos-based applications. Capability has not been advertised to users but is available and undergoing testing.
    - Status: Complete
- **Q3: TuckerMPI-Unstructured: Build functional tensor compression of unstructured meshes into TuckerMPI for in situ integration with Catalyst.**
    - Adopted a simplistic approach for compressing unstructured mesh data, interpolating onto an implicit structured mesh and then launching TuckerMPI, which is a wrapping effort around TuckerMPI (so no major integration effort needed). Regardless, we are continuing to investigate pieces of the functional tensor approach, which was our original strategy, and have submitted a paper to a KDD workshop.
    - Status: Complete
- **Q3: ML for Meshing: Use graph-classification techniques to learn features rather than relying on subject-matter experts.**
    - Completed several experiments with graph-classification techniques for characterizing geometry. Results indicated that topology-based features were more accurate and efficient.
    - Status: Complete
- **Q3: Adaptive Solvers: Develop workflow for analysis and feature importance**
    - Workflow for feature importance, including decision tree path weights and cosine similarity has been developed.
    - Status: Complete
- **Q3: Adaptive Solvers: Find minimal feature set for various datasets and get "bounding box" for features in real applications.**
    - Bounding box of features extracted for ALEGRA prototypes on serrano. Efforts on computing minimal feature set for Mk3c dataset will be completed by end of Q3.
    - Status: Complete
- **Q3: Provide visualization support for the LSCI scaling study milestone**

- Delivered Catalyst in-situ visualizations for SPARC and Nalu. Provided post-processing conversion code for Sparta for visualizing results in ParaView. Used Houdini software (commercial visual effects software) to post-process Exodus files from Empire into artistically-rendered model of EM experiments.
- Status: Complete

2.3.5.04 ATDM SW ecosystem

PERFORMANCE TO PLAN / MEETS

- SPARC container build represents, to the best of our knowledge, the first containerization of a full production ASC mission application.

EXCEEDS:

- N/A
- **Q3: ASC Application Container**
  - Demonstrate an ASC mission application in a container environment. Coordinate with ATDM DevOps and Sandia HPC support teams to develop a methodology for utilizing containers for testing and deploying mission applications. Demonstrate a production code, ideally SPARC or EMPIRE, within a container environment. Compare performance / scalability with native execution.
  - The SPARC app has compiled successfully within a custom container image. This container targets a CTS1 deployment using Intel 18 compilers and OpenMPI. The container is currently being deployed on SRN Doom machine. In the remainder of the month we will do the performance study comparing the container versus native build.
  - Status: Complete