SAND2019-3647R

SANDIA NATIONAL LABORATORIES

SNL ADTM

FY19Q2 report for ATDM AD projects to ECP

March 29, 2019

Prepared by: William Rider

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.

Performance to Plan:

- The SPARC development team will invest in changes early in FY19 that improve developer and user productivity. Planned activities include: (a) code maintenance and technical debt reduction, (b) implement better pre-push testing workflows for developers, (c) overhaul SPARC's parser in preparation for production use by aero analysts.
- Implement hybrid RANS/LES capabilities in SPARC in preparation for FY19 V&V activities and subsequent use of SPARC for simulation of reentry scenarios.
- Implement and demonstrate bow shock adaptivity for structured grids.
- Sandia's ATDM components vision has enabled apps like SPARC to build on foundational capabilities developed and deployed by other teams, providing great leverage and potential for reuse and increased impact. However, these dependencies are also risks: SPARC depends on the component teams to deliver, and that delivery has not always been as complete and timely as needed. We expect that such issues will persist, and that the dependencies that SPARC is built upon will continue to present some risk to the SPARC team's success. To mitigate, we will work closely with to the component teams to ensure delivery and/or find work-around solutions. The SPARC team is on the hook for three L2 milestones this FY.
- Development of performant GPU code has been highly challenging, due to instabilities in the advanced architecture testbeds and difficulties with CUDA compilers. CUDA-related development tasks have taken substantially longer to complete than expected as a result.
- Substantial progress has been made in EM solver performance on Trinity: up to 5 solves per second are being seen on the FY18 milestone problem (only a factor of two behind our target). Scaling problems and deep copy performance in Kokkos are currently being addressed. This progress has been made possible by dedicated effort and tight integration with the EMPIRE development team by the solver team.
- Implement segregated solve with exact Schur complement and demonstrate electromagnetic solver performance improvement in EMPIRE. This is to address the major feedback from the FY18 L2 Milestone review. This has been demonstrated in the surrogate mini-EM application and implementation in EMPIRE will be complete by the end of Q1. In scaling tests performed on Trinity, speed improvements of up to 20x have been observed, and some problems are now meeting the target of 10 linear solves/second. Committed development by the solver team, along with the algorithmic improvements defined by this deliverable, has resulted in these substantial performance gains.
- Extend and demonstrate the EMPIRE PIC/Fluid Hybrid capability on a Sierra-like GPU platform and assess its performance. This will identify focus areas for both algorithm and performance development to ensure success in the FY20 L1 milestone. The conversion of critical evaluators in EMPIRE-Fluid to use Kokkos (and hence be performant on GPU platforms) is on track to be complete by the end of Q2. This is the critical step to making EMPIRE-Hybrid functional on GPU platforms. Initial results from EMPIRE-PIC suggest that solver performance will again need substantial effort to be performant on GPU's, but the solver team is already engaged and making progress. Integration of PIC and Fluid using Tempus as the time integrator and SPIN to provide collision operators is on track for Q2 completion.
- If adequate allocations on Trinity and Sierra platforms cannot be obtained due to the number of teams needing to run milestone calculations simultaneously, then our completion of the L2 and ATS-1 milestones could be jeopardized.

Exceeds:


Sandia National Laboratories