SAND2019-7478R

SANDIA NATIONAL LABORATORIES

SNL ADTM

FY19Q3 report for ATDM AD 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.

Performance to Plan:

- SPARC assembly time has improved by over a factor of 4x since the start of the FY19 on ATS-2. This has come from many sources within the code including optimizations to SPARC, Kokkos and Tpetra
- The performance and scaling analysis and optimization for x86 and GPU architectures is essentially complete. Continued optimization and assessment on Astra (stretch goal for the L2 Milestone) are planned for Q4.
- The flight test validation is making slow but steady progress. Modeling improvements are underway to improve the aerothermal predictive capability of SPARC for this particular flight test.
- For the high-order methods work, the SPARC team has evaluated and reduced numerical dissipation in direct numerical simulations of a Mach 3.5 flat plate boundary layer.
- SQA for SPARC - report was delivered - Adequate SQA practices.
- Scaling results on Trinity have been rerun at XXL scale (1.6B elements) for the first time since the FY18 milestone and show continuation of the existing strong and weak scaling trends. Performance on HSW is essentially at the target level, and lags by approximately a factor of two on KNL. This enormous progress has been made possible by dedicated effort and tight integration with the EMPIRE development team by the solver team.
- Particle update performance has been improved on CUDA by a factor of four through Q3.
- EMPIRE-Hybrid development continues to make progress as multiple software components are combined to advance towards the capability required for the FY20 milestone. A validation problem using the CEA RKA experiment that can span PIC to hybrid capabilities has been defined to drive code improvements through FY20.
- The default particle push in EMPIRE now uses the Verlet algorithm, offering improved accuracy for verification for minor additional computational cost.
- CHEETAH: Demonstrated Woodcock & standard particle tracking and querying on a GPU, stepping towards capability of photon transport within stochastic materials on GPUs.
- As part of a performance benchmarking effort, the SPARTA code has been recently exercised on the ATS-2 Sierra platform up to 16,384 GPUs. Benchmark studies will be completed in FY19 Q4.
- An EMPIRE validation study is underway to examine modeling sensitivity and uncertainties. Code verification activities include mesh and time-step convergence studies and, where necessary, account for the stochastic aspects of the models.

Exceeds:

- SPARC has exceeded expectations for GPU performance improvements and still have opportunities for further performance gains.
- Porting of the SPARTA code to the ATS-2 Sierra platform, including the implementation and testing of Kokkos-based computational kernels, is ahead of schedule.

Milestone 6808: SPARC Credibility, Performance, and Scalability **(Q3)**

Status: [Green] On Track

Reviews scheduled:

- Initial [01/25/2019] | Mid-year [04/22/2019] | Final [8/22/2019]

Brief status:

- SPARC assembly time has improved by over a factor of 4x since the start of the FY19 on ATS-2. This has come from many sources within the code including optimizations to SPARC, Kokkos, and Tpetra.

- The performance and scaling analysis and optimization for x86 and GPU architectures is essentially complete. Continued optimization and assessment on Astra (stretch goal for the L2 Milestone) are planned for Q4.
- The flight test validation is making slow but steady progress. Modeling improvements are underway to improve the aerothermal predictive capability of SPARC for this particular flight test.

Milestone 6809: Validation of EMIPRE Cavity-SGEMP Simulation Results with ICF Experimental Data **(Q3)**
Status: [Green] On Track
Reviews scheduled:

- Initial [01/16/2019] | Mid-year [05/20/2019] | Final [08/14/2019]

Brief status:

- The four key emission models necessary for B-dot cavity physics and DSMC (Direct Simulation Monte Carlo) collision models for Argon background gas have been implemented and tested.
- A series of B-dot simulations have been performed with vacuum, gas fill, and gas fill with space-charge limited emission models. Results are being validated against NIF experimental data. These were presented at the mid-term panel review.
- The ATS-1 L2 milestone is dependent on this milestone and is similarly on track. The mid-term review occurred at SNL on 5/15/2019. Target problems at the requisite scale and of interest to this Milestone have been defined and will be run during Q4.

News notes/accomplishments:

SPARC

Sandia/SPARC Enhancement Performance for Sierra. From FY19 Q1 to Q3, Sandia developers on the SPARC (Sandia Parallel Aerodynamics and Re-Entry Code) team have significantly improved the performance on ATS-2/Sierra. The optimizations have come from essentially all parts of code, encompassing much of the assembly, solve and post-solve (i.e. gradient computation, field communication) phases of execution, with contributions coming from the SPARC, Kokkos and KokkosKernels teams at Sandia. As measured on a per-V100 basis, and for a steady-state, perfect gas simulation with a two-equation turbulence model exercising the structured finite-volume discretization, SPARC's total time step execution time now ranges between approximately 2x and 3.6x faster than ATS-1/Trinity Haswell nodes for problems over 1GB of memory per node or GPU. As measured on a per-ATS-2 node basis (4 V100s per node), SPARC's execution time ranges between approximately 8x and 10x faster than ATS-1/Haswell for problems over 1GB of memory per node or GPU. This is a significant achievement for SPARC on GPU-based systems and for production readiness of SPARC on ATS-2/Sierra. There are additional optimizations under preparation and evaluation this FY that will further increase SPARC's GPU performance.

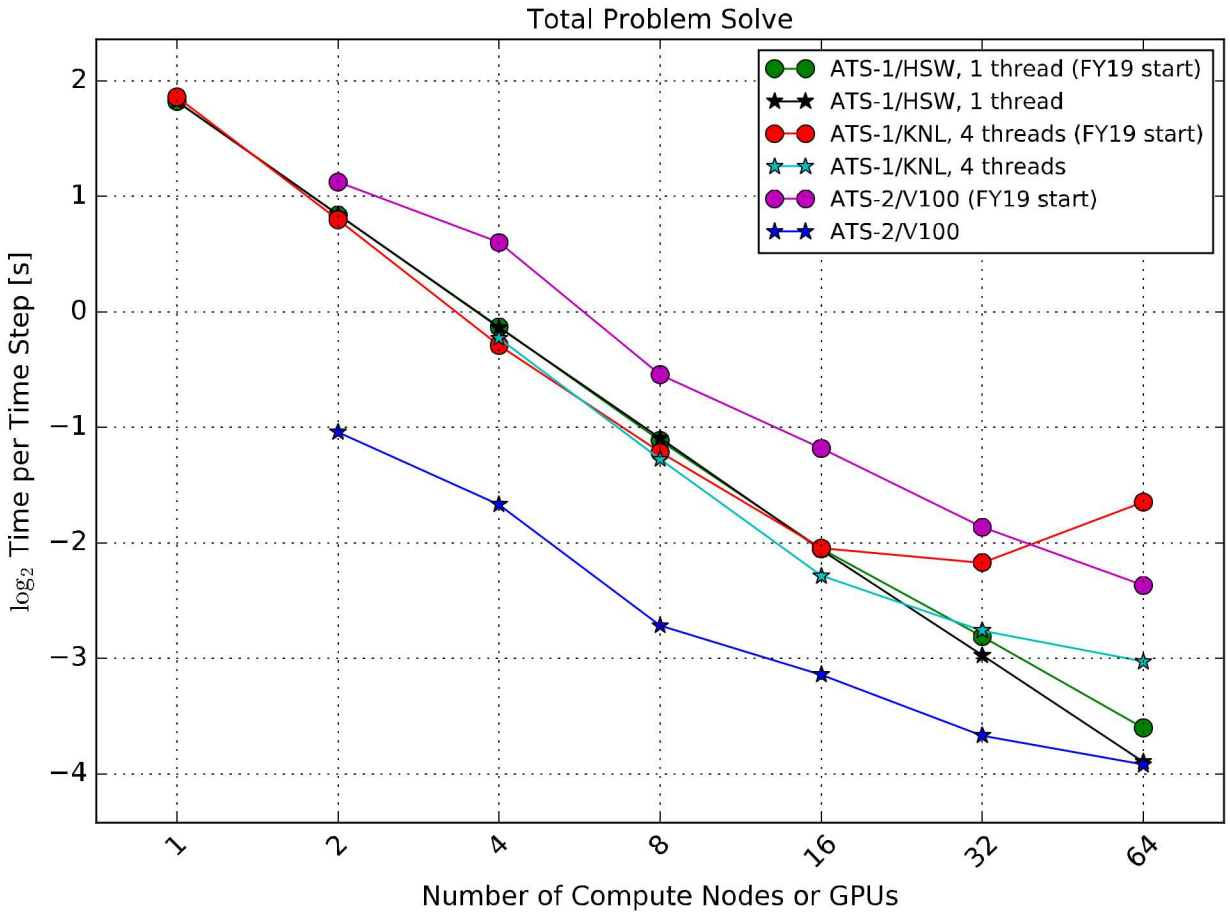

Accomplishment Summary: Strong scaling of total time step execution time on ATS-1/Haswell, ATS-1/KNL and ATS-2/V100 as compared to FY19 start and present-day SPARC. X-axis is number of ATS-1 nodes or number of V100 GPUs; Y-axis is time per timestep on a log2 scale. The use-case is a steady-state, perfect gas simulation with a two-equation turbulence model of the HIFiRE-1 hypersonic flight experiment geometry.

Sandia National Laboratories