



Sandia National Laboratories



FY20Q1 report for NNSA-ST projects to ECP

December 2019

Prepared by: Gabrielle Trujillo

Prepared for:

ECP Confluence updates

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.



Sandia National Laboratories



Activities:

Math Libraries (Kokkos Kernels):

- KK CNARS process plan created (Rajamanickam) After several discussions we have arrived at a plan for CNARS process that causes the minimal pain to Trilinos and Kokkos Core + Kokkos Kernels developers, while allowing applications a path forward to take our codes to SCN as quickly as possible.
- KK new implementation to highlight (Rajamanickam) Sparse triangular solve for triangular matrices arising from direct solvers and incomplete factorizations. We also support Third party libraries where available (targeted for 3.0 release); The new sparse triangular solver is 2-13x compared to NVIDIA's sparse triangular solver; A sparse ILU(k) for better preconditioning options on GPUs; a cluster-coloring based Gauss-Seidel preconditioner that reduces the launch overhead and reduces the number of iterations (this will assist EMPIRE as a smoother); a DotBasedGEMM in Kokkos Kernels that is tuned for the problem sizes we care about (the performance evaluation experiments against cuBLASgemm showed up to 16x speedup. This also cutdown the cost of the apply by 40% in EMPIRE. We have reported this problem to NVIDIA and they have fixed it in their codes); multiple right hand side support for Gauss-Seidel preconditioners (this helps several projects with an immediate impact to Exawind project); performance portable graph algorithms for breadth-first search and RCM ordering, balanced graph clustering (this allows developing preconditioners based on these graph algorithms); team/serial batched sorting so applications can sort within the teams.
- KK helps GEMMA via Pliris, a new LU solver package in Trilinos (Rajamanickam) New Kokkos-based Pliris (LU solver in Trilinos) runs for GEMMA on LLNL's Lassen system and investigating on-going communication overhead issues for the new Pliris; This solver will be added as a new package in Trilinos.

Programming Models & Runtimes (Kokkos):

- Supported Empire debugging effort and developed workarounds for bugs
- Conducted a 2-day Kokkos user training event for Sandia customers.
- Initial development started on AMD (HIP) and Intel (SYCL) backends, with base capabilities such as "parallel for" already working

Data & Viz (VTK-m):

- Good progress on burst-buffer support for SIERRA. Coordinating this activity with folks at NERSC and LLNL.
 - * New I/O work for Empire-Fluid (see deliverables below)
- New work for in situ analysis of SPARC data has identified a number of Catalyst bugs. These are being prioritized and addressed by Kitware.



- New work in HPC tensor analysis (Danny Dunlavy--PI) has identified a number of interesting areas for evaluation, including runtime performance from ASTRA to support long-term resiliency objectives. Joint work with 1420 and 8750.
- Completed TuckerMPI ParaView plugin (remaining work from FY19). Will be available in next release of ParaView (January 2020).
- Completed deliverable on burst buffer resilience. Developed an analytics model for "delayed resilience" that allows checkpoints to progress in an unreliable way immediately, but be improved through various resilience approaches (similar to RAID) to ensure a mix of reliable and unreliable burst buffers devices.

SW Ecosystem & Delivery (OS&ONR):

- No deliverables due this quarter, but Tech Demo Q2 deliverable completed early
- See individual deliverables below for progress on other deliverables.
- Publication: Richard Canon and Andrew Younge. "A Case for Portability and Reproducibility of HPC Containers" in Workshop on Containers and New Orchestration Paradigms for Isolated Environments in HPC (CANOPIE HPC) @ SC19

Accomplishments:

Data & Viz (VTK-m): Analytically evaluate a range of resiliency paradigms across storage backed by unreliable burst buffers.

Next Steps:

Math Libraries (Kokkos Kernels):

- Increase transparency of data flow patterns within Tempus; in particular, redesign observer / time step control objects to make inputs / outputs transparent to the user. Extend this redesign to other Tempus object as possible.
- Support the L1 milestone for SPARC by developing new linear algebra kernels and improving current ones on GPU platforms. The primary focus will be on the performance of line solves.
- Support the L1 milestone for EMPIRE by developing new linear algebra kernels and improving current ones on GPU platforms. The primary focus will be on the solver kernels identified by scalable solvers team on GPUs.



Sandia National Laboratories



SW Ecosystem & Delivery (OS&ONR): OS/R: Use the simulation framework developed for quantifying MPI resource usage to characterize ATDM workloads and examine the relationship between MPI resource usage and application performance and scalability.

Challenges:

Outreach & Communication:

Data & Viz (VTK-m): Oldfield began to meet and collaborate with the ND Data Engineering Strategy team (Brett Remund -- 2000). Strategy/Team lacked a connection to HPC and some key ASC capabilities. There is potential to collaborate to improve engagement between broader ND program and Sandia's ASC program

SW Ecosystem & Delivery (OS&ONR):

- Continuing collaboration with other labs on containers in ASC and ECP. Container BoF @ SC19 led by Andrew Younge was standing room only.
- Participating in COE and NRE activities for El Capitan, including system software and network/messaging working groups