

LLNL NNSA Software TECHNOLOGY

The NNSA supports the development of open-source software technologies that are both important to the success of national security applications and externally impactful for the rest of the ECP and the broader community. These software technologies are managed as part of a larger Advanced Simulation and Computing (ASC) portfolio, which provides resources to develop and apply these technologies to issues of importance to national security. The software technologies at LLNL span programming models and runtimes (RAJA/Umpire/CHAI), development tools (Debugging @ Scale), mathematical libraries (MFEM), productivity technologies (DevRAMP), and workflow scheduling (Flux/Power).

The RAJA/Umpire/CHAI team is providing software libraries that enable application and library developers to meet advanced architecture portability challenges. The project goals are to enable writing performance portable computational kernels and coordinate complex heterogeneous memory resources among components in a large integrated application.

The software products provided by this project are three complementary and interoperable libraries: RAJA provides software abstractions that enable C++ developers to write performance portable numerical kernels; Umpire is a portable memory resource management library that provides a unified high-level Application Programming Interface (API) in C++, C, and Fortran for resource discovery, memory provisioning, allocation, transformation, and introspection; and CHAI contains C++ “managed array” abstractions that enable transparent and automatic copying of application data to memory spaces at run time as needed based on RAJA execution contexts.

Debugging @ Scale provides an advanced debugging, code-correctness, and testing toolset to

facilitate reproducing, diagnosing, and fixing bugs within HPC applications. The current capabilities include STAT, a highly scalable lightweight debugging tool; Archer, a low-overhead OpenMP data race detector; ReMPI/NINJA, a scalable record-and-replay and smart noise injector for message passing interface (MPI); and FLiT/FPUChecker, a tool suite for checking floating-point correctness.

The MFEM library is focused on providing high-performance mathematical algorithms and finite element discretizations to next-generation, high-order applications. This effort includes the development of physics enhancements in the finite element algorithms in MFEM and the MFEM-based BLAST Arbitrary Lagrangian-Eulerian code to support ASC mission applications and the development of unique unstructured adaptive mesh refinement algorithms that focus on generality, parallel scalability, and ease of integration in unstructured mesh applications.

DevRAMP is creating tools and services that multiply the productivity of developers through automation. The capabilities include Spack, a package manager for high-performance systems that automates the process of downloading, building, and installing different versions of software packages and their dependencies, and Sonar, a software stack for performance monitoring and analysis that enables developers to understand how high-performance computers and applications interact.

Flux /Power is a next-generation resource management and scheduling software framework. The team is providing a portable, user-level scheduling solution for complex exascale workflows and a system resource manager and scheduler for exascale systems.

Progress to date

- The RAJA/Umpire/CHAI team developed support for multidimensional kernel dispatch; atomic operations on GPU accelerators; memory allocation on CPU, GPU, unified, and “pinned” memory resource and copying data between resources; high-performance memory pools; and an Umpire-based backend for CHAI that adds additional flexibility and capability.
- The Debugging @ Scale team completed the port of STAT, Archer, ReMPI/NINJA, and FLiT/FPUChecker to Sierra, deployed these tools on Sierra, and assisted those using these tools for debugging and testing.
- The MFEM team released MFEM version 3.4 with many new features including a significantly improved nonconforming unstructured adaptive mesh refinement capability, block nonlinear operators; general high-order-to-low-order refined field transfer; and specialized time integrators.
- The DevRAMP team implemented the ability for Spack to output build reports to CDash, developed syntax and a workflow for reproducible multipackage deployments, and provided a cloud-based build farm.
- The Flux/Power team enabled two major scientific workflows to complete their calculations on the Sierra pre-exascale system and released a version of their software containing all the functionalities used by these workflows.

PI: Becky Springmeyer, Lawrence Livermore National Laboratory

Collaborators: Lawrence Livermore National Laboratory