HPC Tools for Performance, Debugging, and Performance Analysis

LLNL has a long and proud history of leadership in HPC tooling across many projects
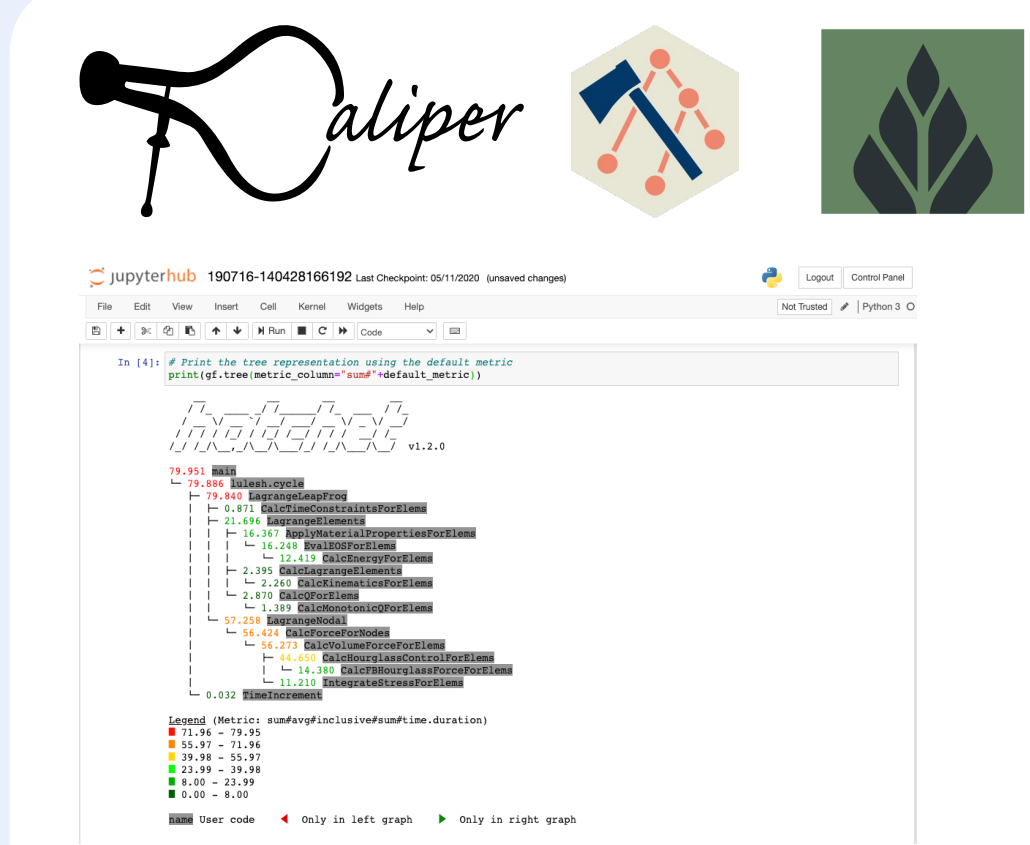
Matthew P. LeGendre (LLNL)

LLNL R&D

LLNL-led R&D tool projects identify and fill gaps based on experiences from running the world's largest computers.
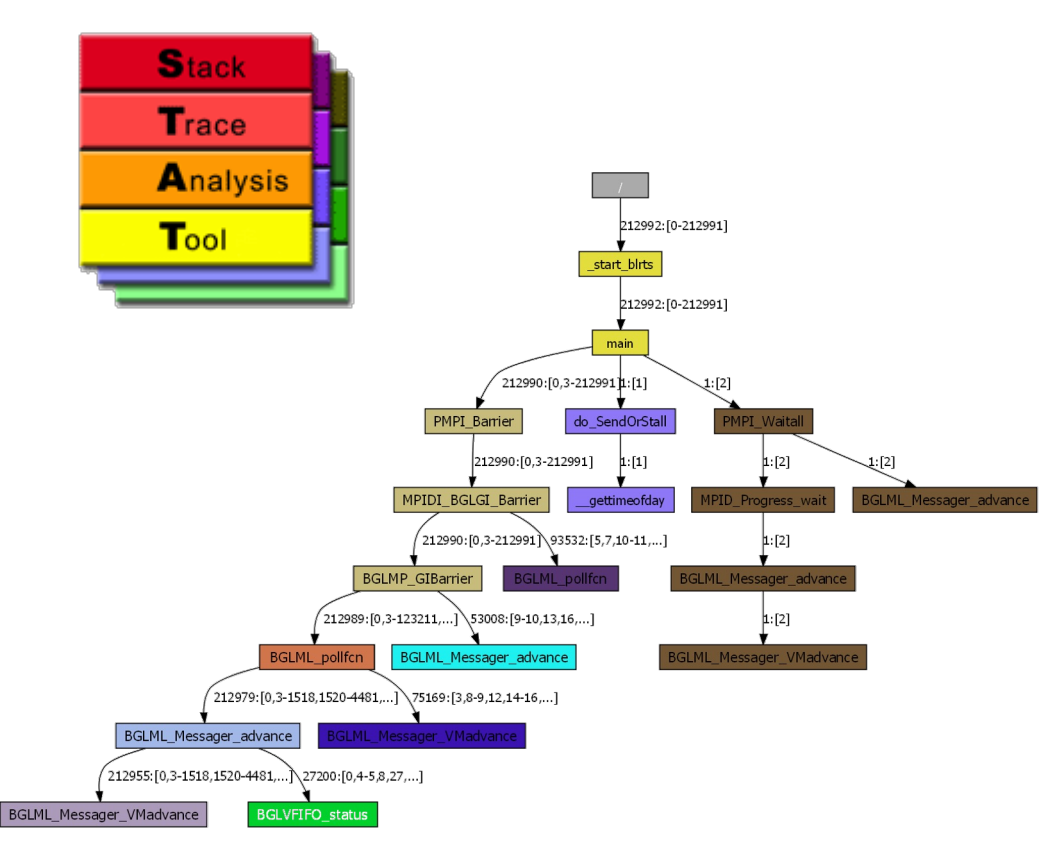
Dual focus on both innovative research and production-quality tooling.

Averages 4–5 FTEs/year.

Caliper, Hatchet, and Thicket provide **ubiquitous performance analysis**. Enabling applications to tightly integrate with performance analysis tooling.

Impact: App teams can run performance analysis tools

STAT enables debugging at the massive scales found in LLNL HPC. R&D100 Award winner that has been adopted across the HPC ecosystem.

Impact: Debugging at massive scale
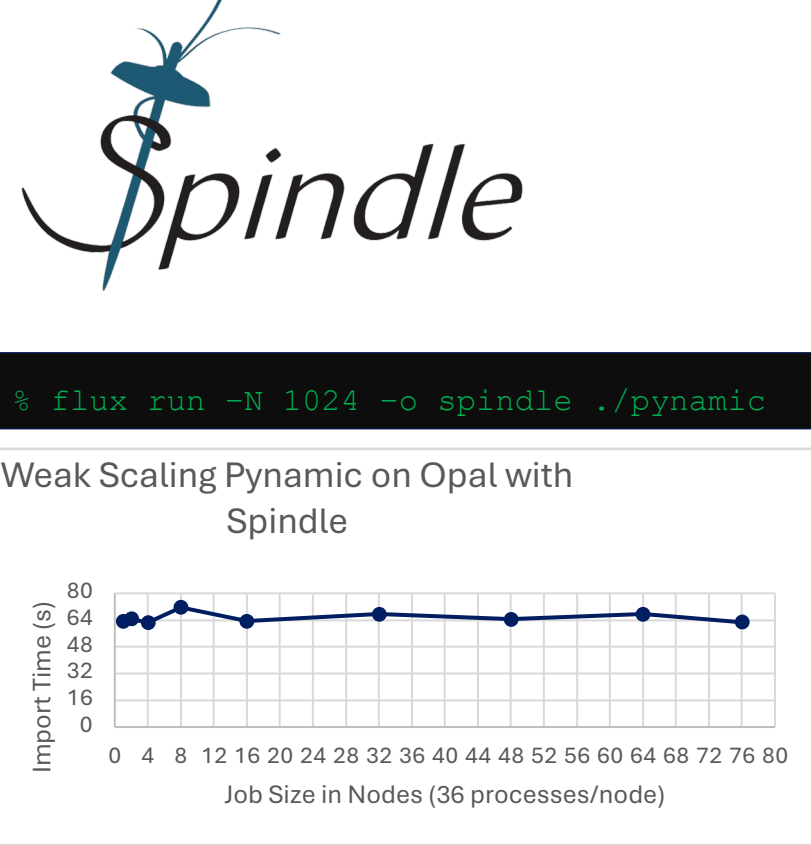
Spindle accelerates dynamic library loading on HPC. Making it possible to launch complex applications at scale.

*Impact: Job launch at massive scale*

Academic Partnerships

Partnerships with research institutions expand the HPC tools ecosystem and create employee pipelines.

Implemented through contracts, which range from production/ research focus to pure research.
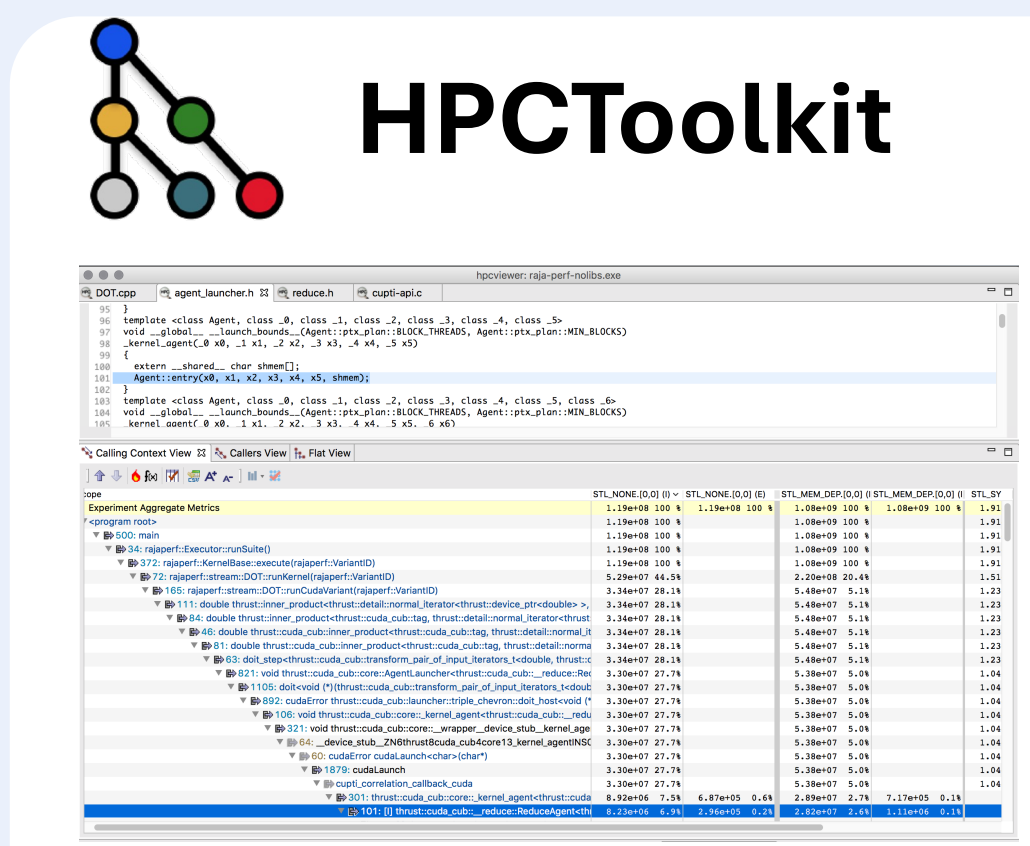
HPCToolkit

Rice University partners with the Tri-Labs to provide **HPCToolkit** as a production-quality HPC performance analysis tool.
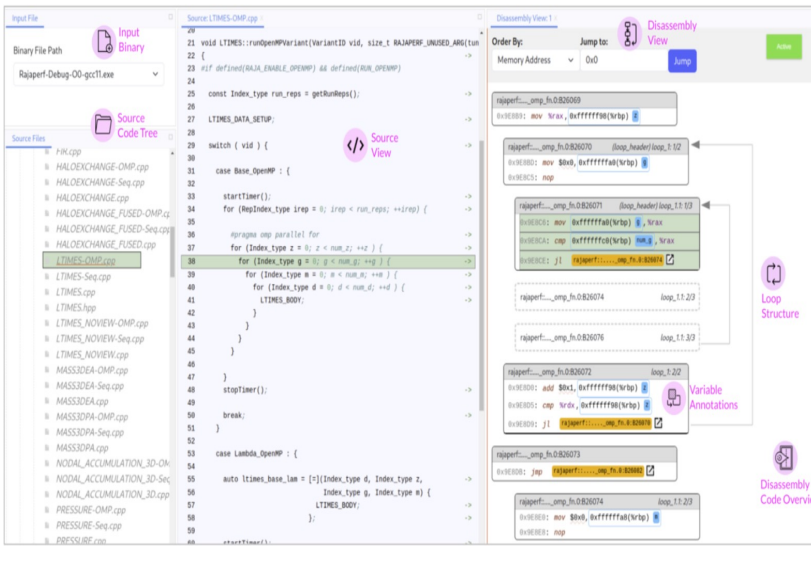
Impact: App teams have a cross-platform HPC tool
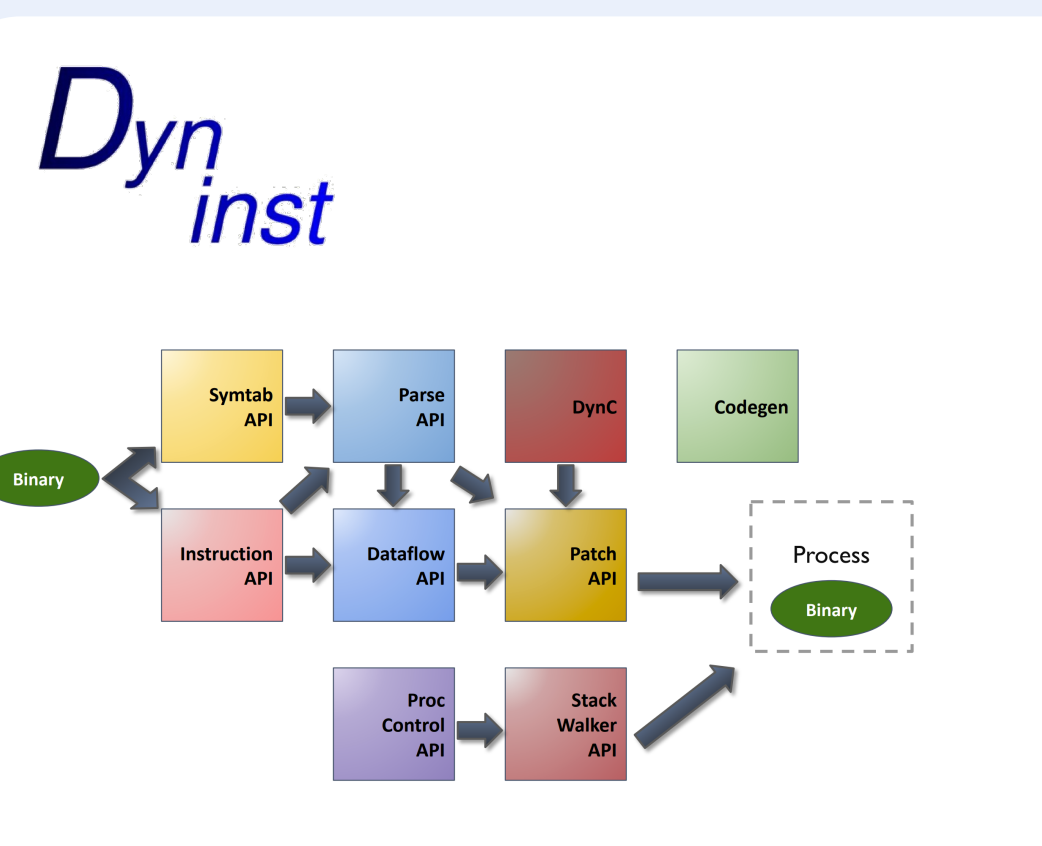
University of Wisconsin-Madison works with Tri-Labs to deliver **DyninstAPI**'s binary analysis and instrumentation libraries that underly many other tools.

Impact: Tool developers get robust infrastructure

ccNav

University of Utah's research on **ccNav** will help visualize compiler optimizations in binaries.
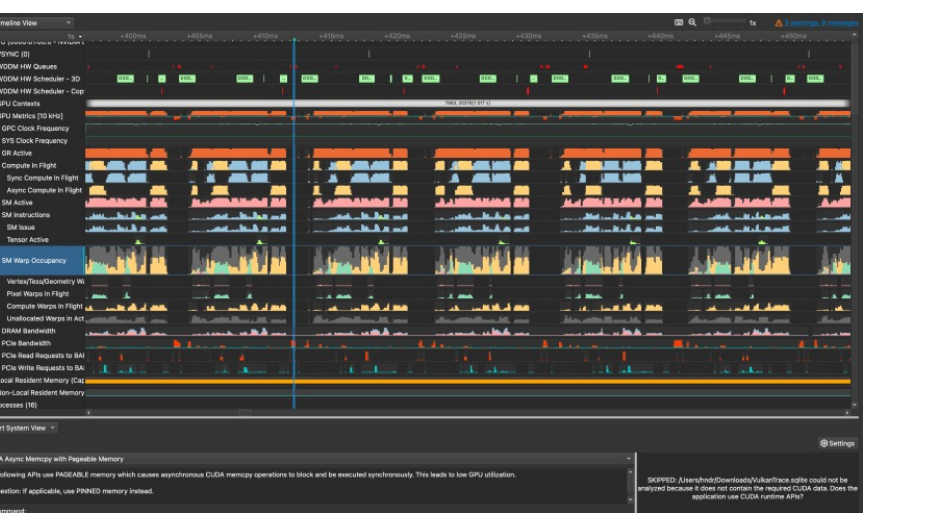
Impact: Research supports performance tuners

Industry Collaborations

Collaborations bring HPC tool expertise to system vendors.

Hardware vendors have expertise in processor and GPU details but lack in-depth knowledge in HPC workflows and applications that is needed to make their tools work on our systems.
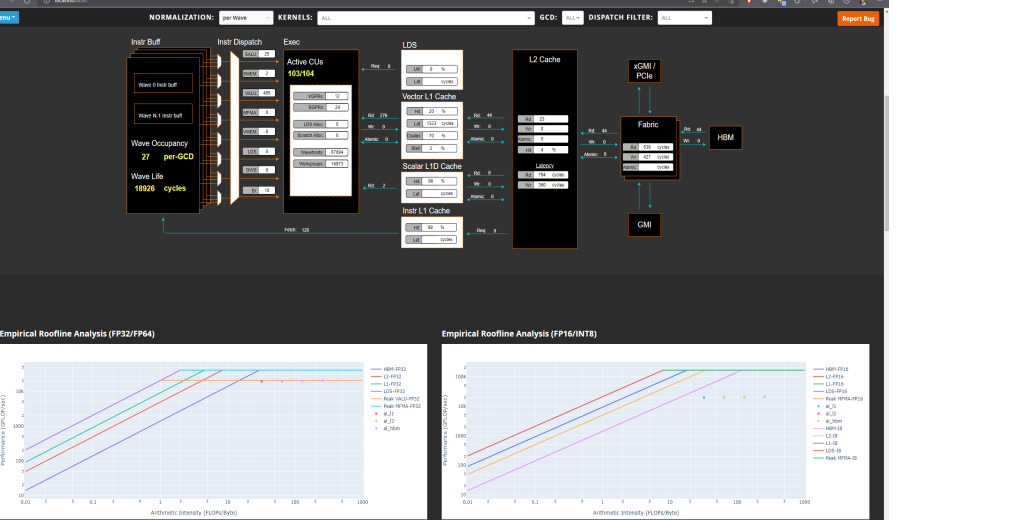
Nvidia Nsight

Worked with Nvidia during their development of **Nsight** performance analysis tool. Experiences from Sierra played an important role in scaling the tooling.

Impact: App teams can analyze Nvidia GPU kernels
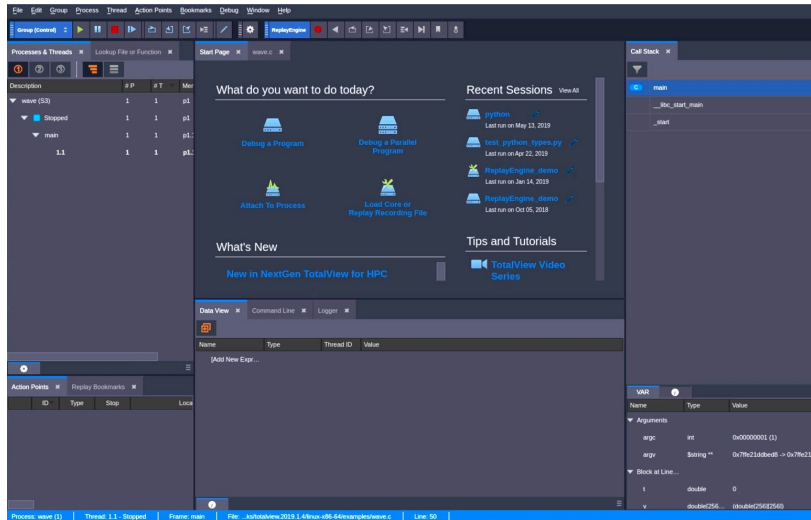
AMD Rocprofiler

Worked with AMD on their design of **Rocprofiler** for El Capitan from the start.

Focused on features like PC sampling, HPC support, counter models, and documentation.

Impact: App teams can analyze AMD GPU kernels

TotalView

Years of collaboration with **TotalView** has helped Perforce build an HPC debugger that works at large scale and with complex applications.

Impact: App teams have a full-featured debugger

Lawrence Livermore National Laboratory

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. LLNL-POST-2004558

National Nuclear Security Administration