STEP

SOFTWARE TOOLS ECOSYSTEM PROJECT

The US/DOE's Software Tools Ecosystem Project

July 7, 2025

Terry Jones, STEP Director

STEP: Software Tools Ecosystem Project

- The Importance of High Performance Computing
- The US DOE and efforts to help the Tools Ecosystem
- STEP and its 5 Initial software packages
- How to Engage with STEP and looking forward

STEP: Software Tools Ecosystem Project

- The Importance of High Performance Computing
- The US DOE and efforts to help the Tools Ecosystem
- STEP and its 5 Initial software packages
- How to Engage with STEP and looking forward

Science Drives Innovation & Innovation Drives Economies

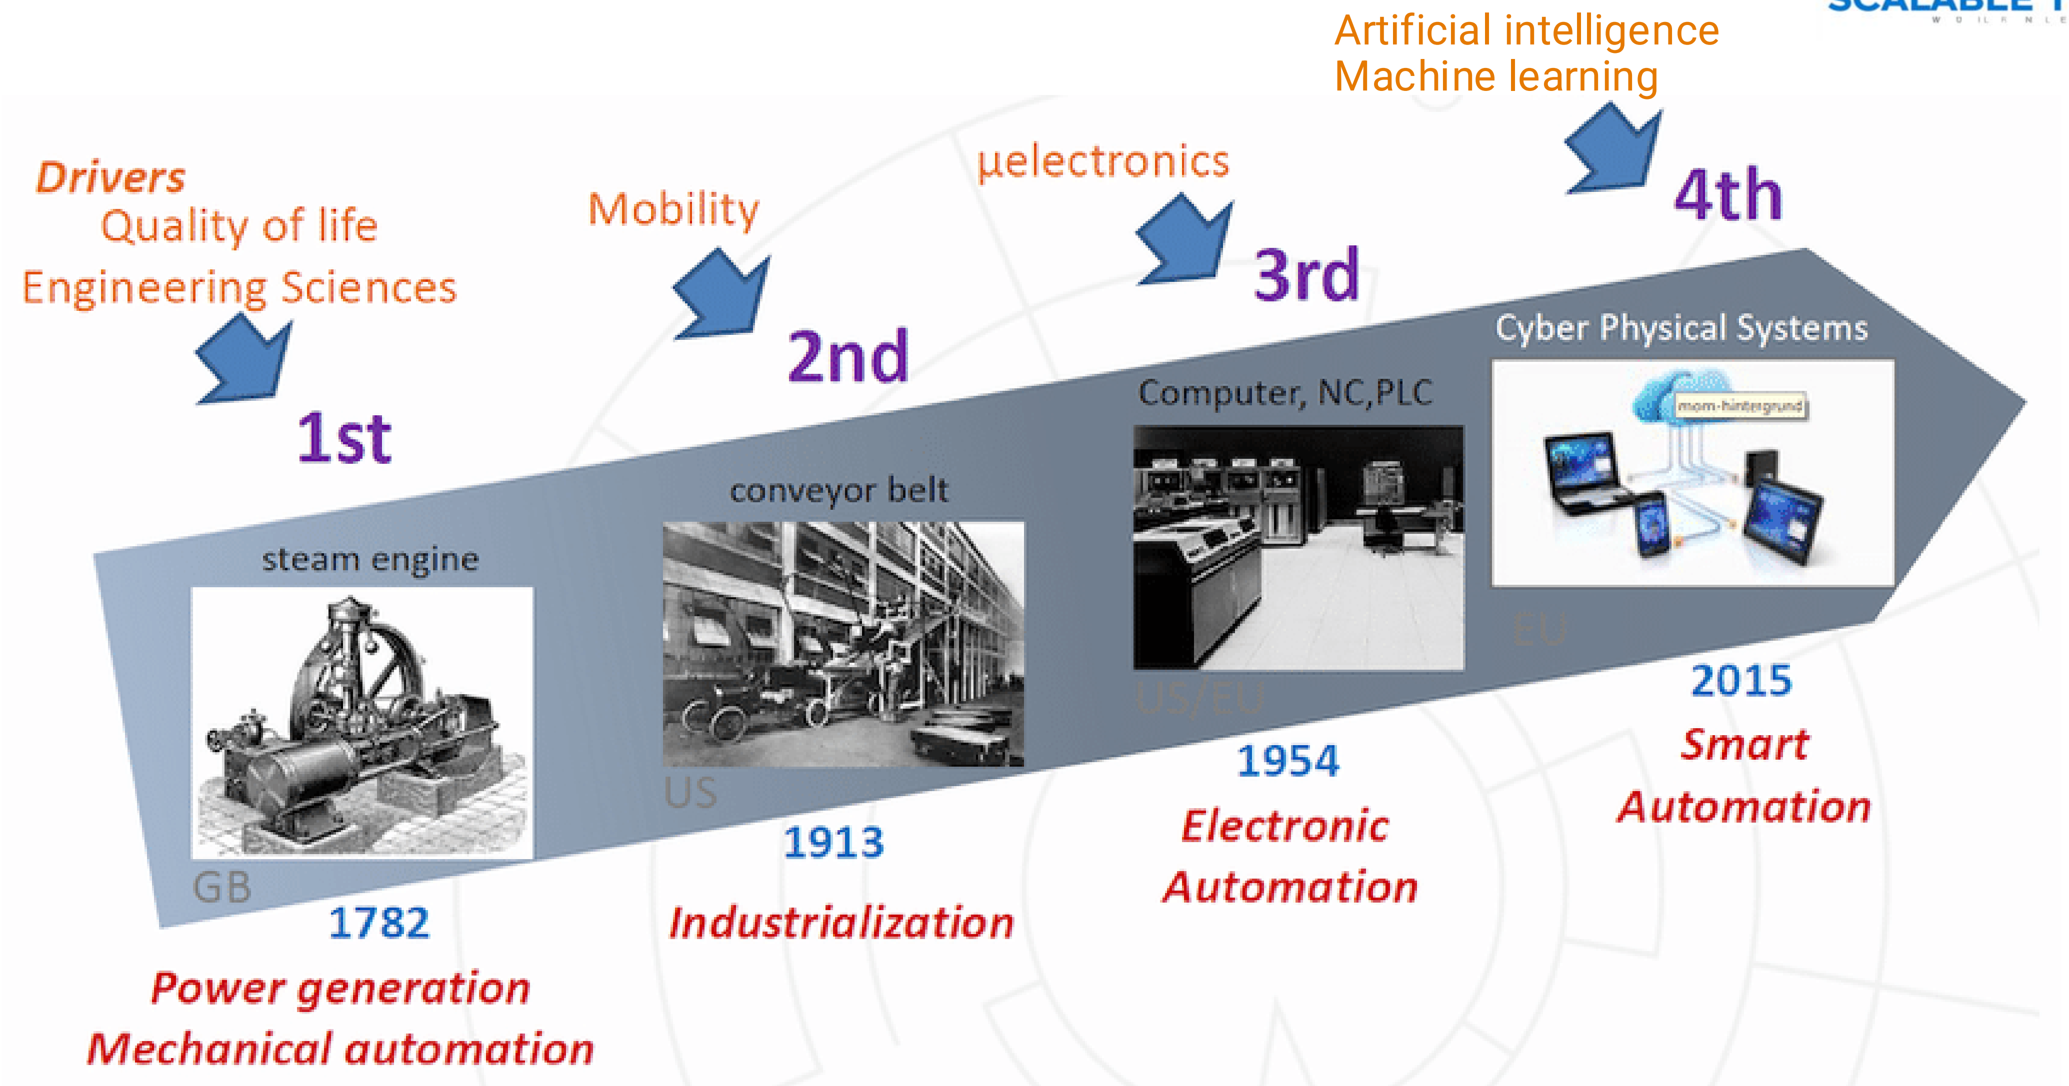

The Two Pillars of Science

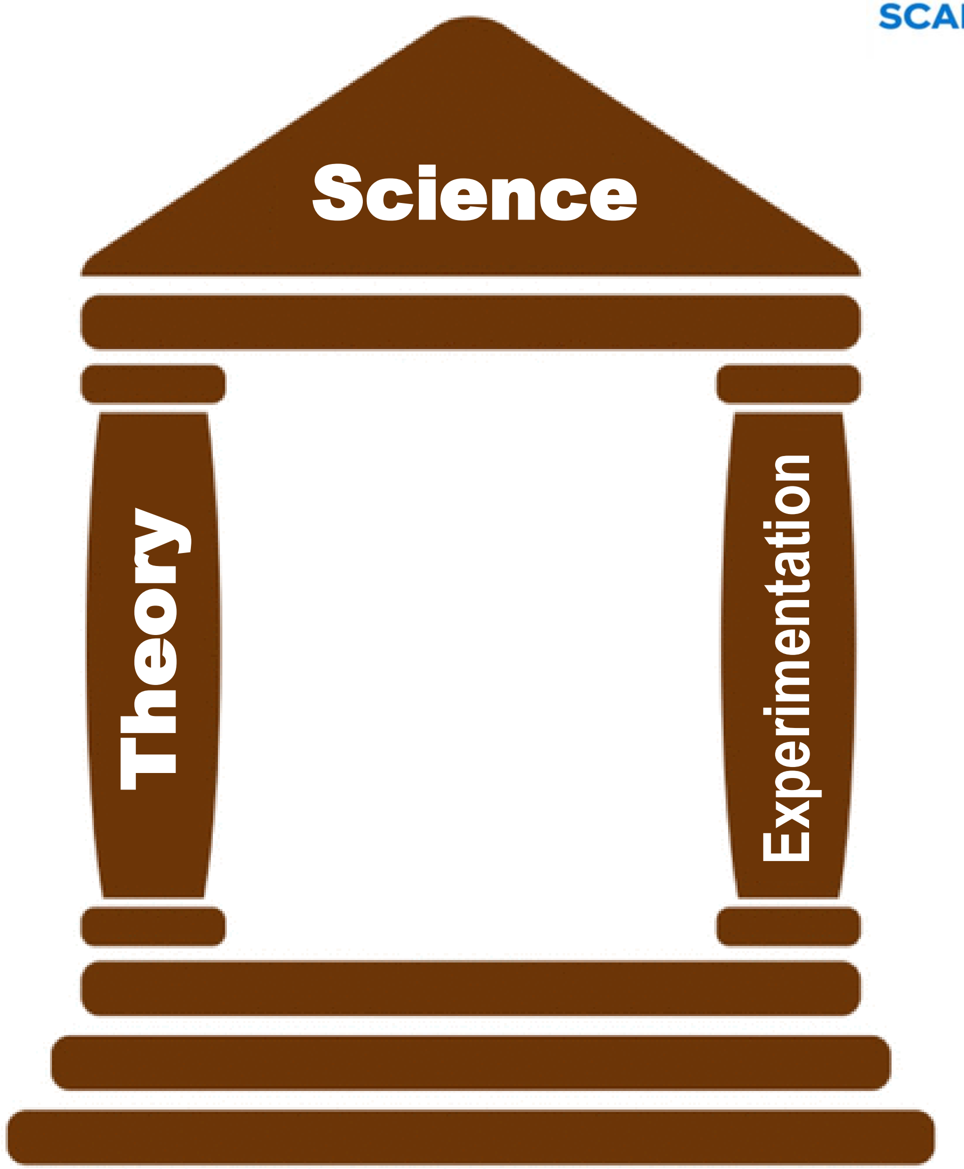

The ~~Two~~ Three Pillars of Science

Computers have changed the way we conduct experiments. Given enough computer power, we can perform accurate experiments more quickly, more cheaply, and often with greater control.

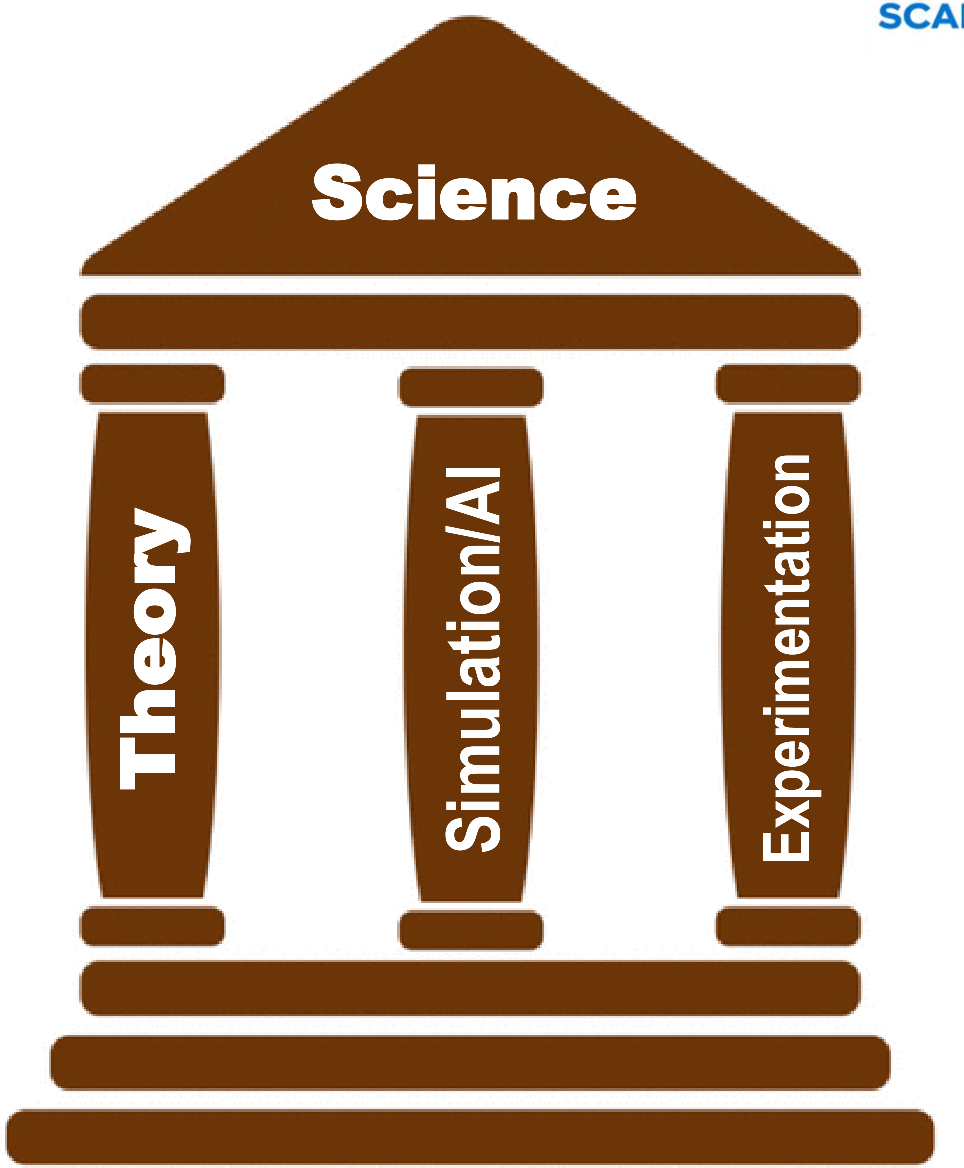

STEP
SOFTWARE TOOLS ECOSYSTEM PROJECT

Scientific discovery has long benefitted from a disciplined process…

…and Supercomputing is changing that process

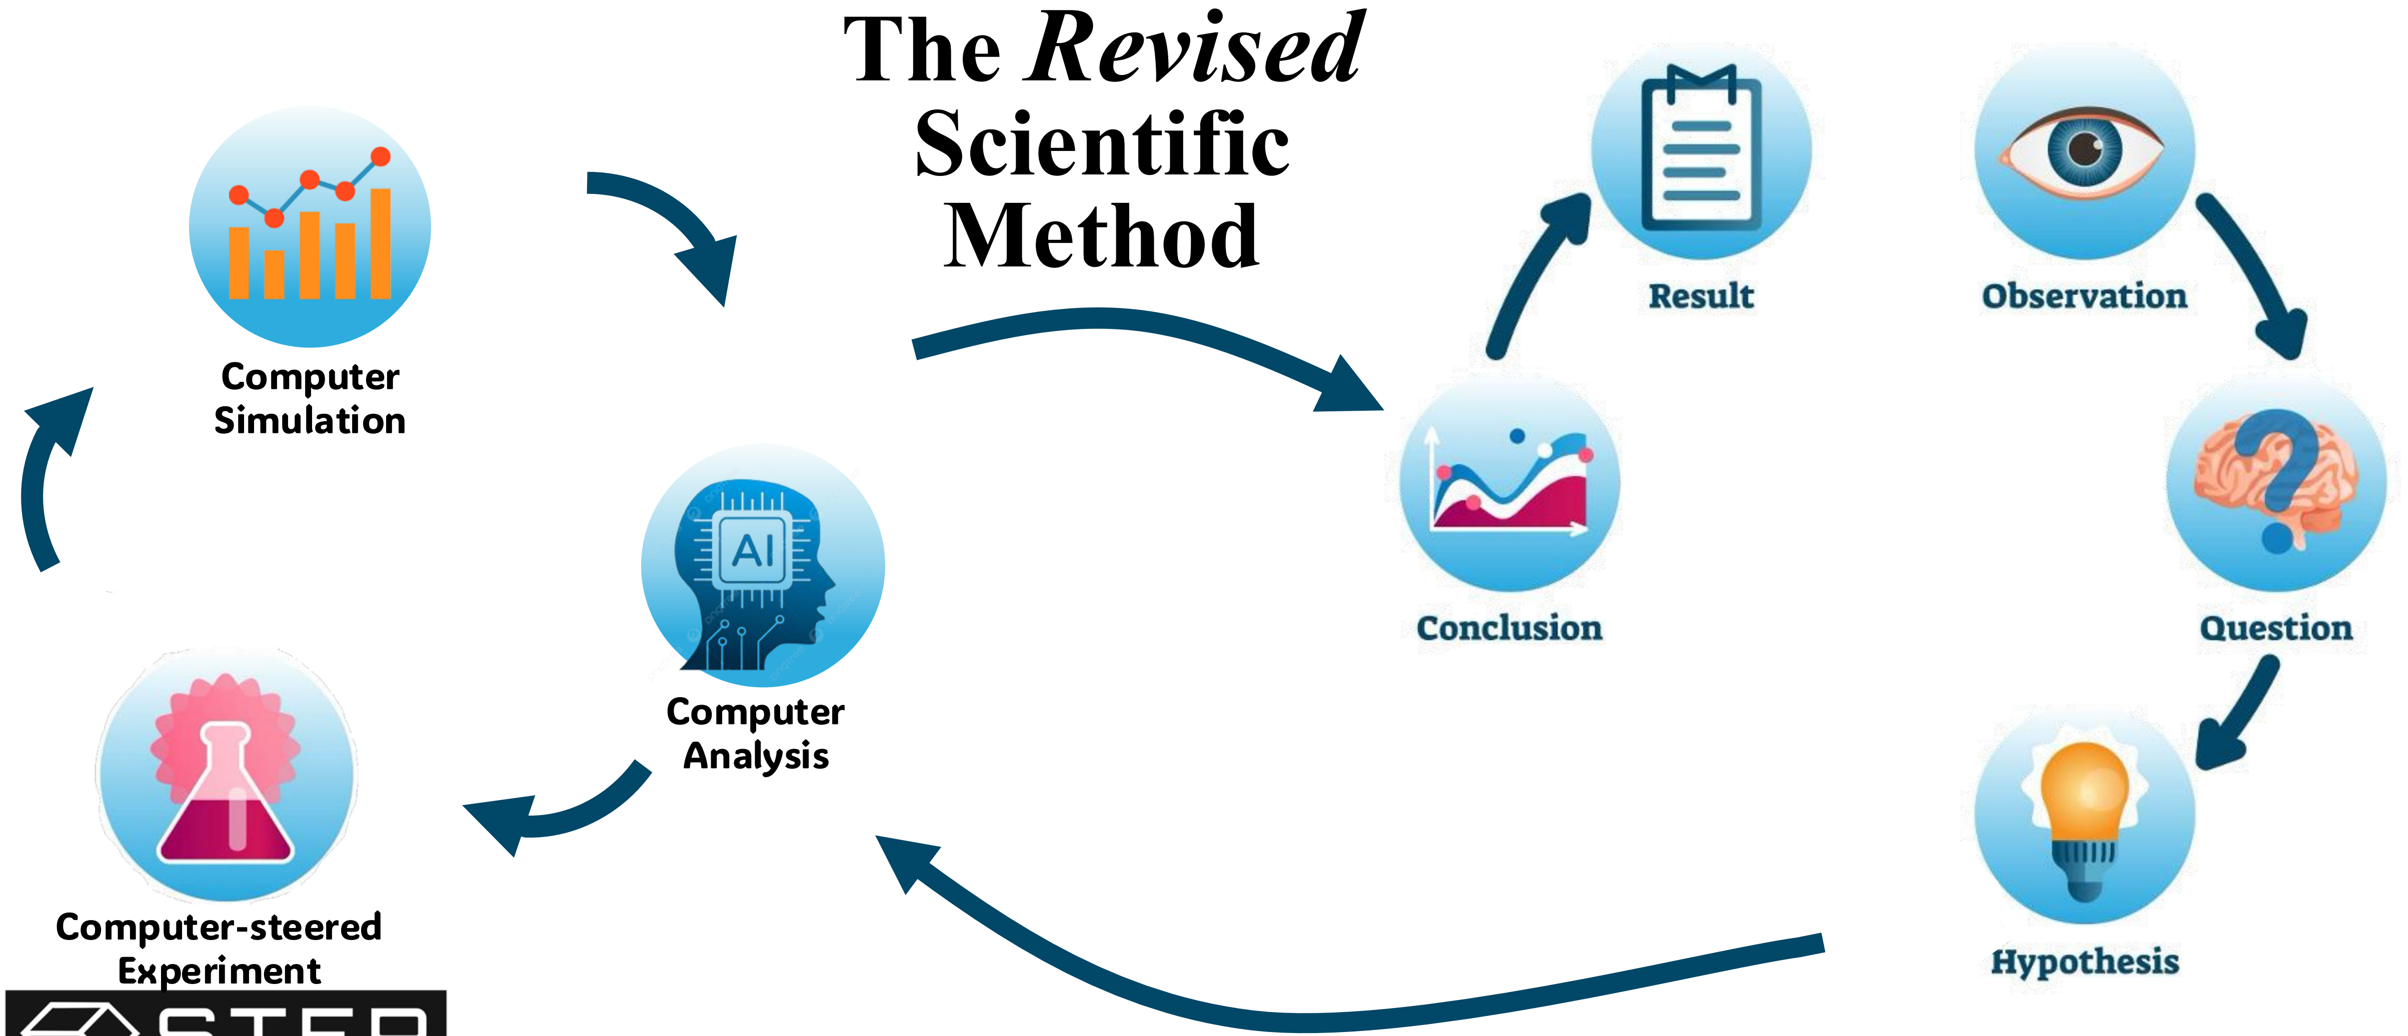

STEP: Software Tools Ecosystem Project

- The Importance of High Performance Computing
- The US DOE and efforts to help the Tools Ecosystem
- STEP and its 5 Initial software packages
- How to Engage with STEP and looking forward

US DOE National Laboratory system has a complex collection of world class facilities spread over a large geographical area

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

DOE/ASCR's Software Sustainability Efforts

Many of the scientific discoveries that our world are made possible by a sophisticated ecosystem of software technology.

This software ecosystem has evolved decades of research with input from countless stakeholders and application scientists.

Sustaining it and ensuring a strong foundation for the future requires a comprehensive and thoughtful approach to software stewardship and advancement.

DOE's Exascale Computing Program 2017 through part of 2024

Fiscal Year	Project Budget
2017	$164.0M
2018	$162.9M
2019	$232.7M
2020	$188.7M
2021	$168.9M
2022	$129.0M
2023	$77.0M
2024	$14.0M*

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

DOE/ASCR's Software Sustainability Efforts

2025 SCALABLE TOOLS

DOE's Next Generation Software Technology Program 2024 and thereafter

- NGSST Phase I: Six Funded Seedlings in FY23
  - 3 Service focused seedlings
  - 3 Functional area seedlings
- Leads to a NGSST Phase II Software Sustainability Program in FY24

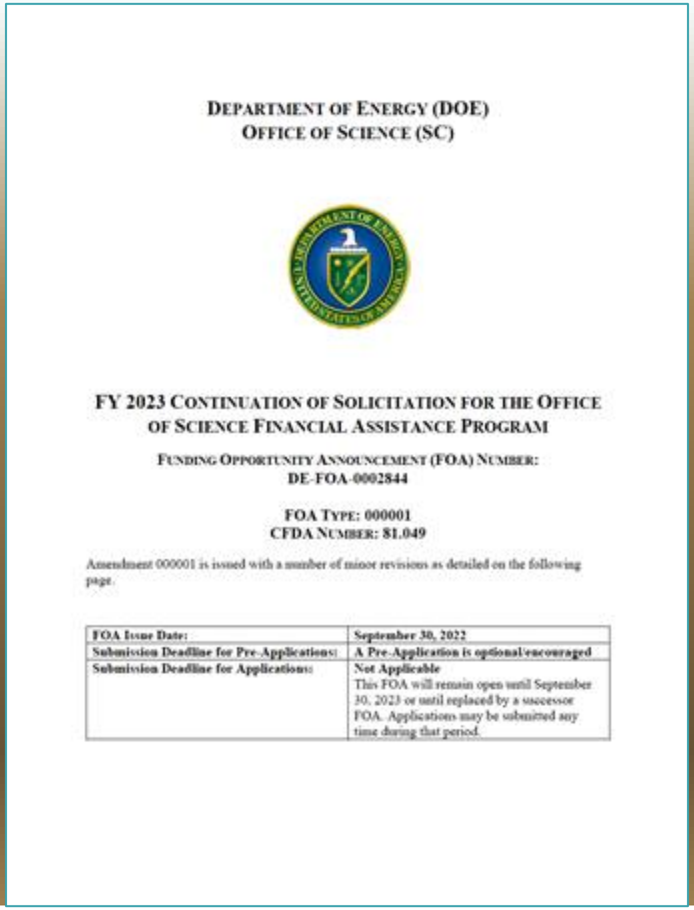

DEPARTMENT OF ENERGY (DOE)
OFFICE OF SCIENCE (SC)

FY 2023 CONTINUATION OF SOLICITATION FOR THE OFFICE OF SCIENCE FINANCIAL ASSISTANCE PROGRAM

FUNDING OPPORTUNITY ANNOUNCEMENT (FOA) NUMBER:
DE-FOA-0002844

FOA TYPE: 000001
CFDA NUMBER: 81.049

Amendment 000001 is issued with a number of minor revisions as detailed on the following page.

FOA Issue Date:	September 30, 2022
Submission Deadline for Pre-Applications:	A Pre-Application is optional/encouraged
Submission Deadline for Applications:	Not Applicable This FOA will remain open until September 30, 2023 or until replaced by a successor FOA. Applications may be submitted any time during that period.

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

STEP: Software Tools Ecosystem Project

- The Importance of High Performance Computing
- The US DOE and efforts to help the Tools Ecosystem
- STEP and its 5 Initial software packages
- How to Engage with STEP and looking forward

The US DOE's Software Tools Ecosystem Project (or STEP)

Objectives and Scope

Tools and supporting software for monitoring, analysis, and diagnosis of performance and behavior of codes on advanced computing systems.

STEP

https://ascr-step.org

Lead PI: Terry Jones
Deputy PI: Phil Carns

Program Mgr: David Rabson

Initial Funded Software Packages

Software Tool	PI
HPCToolkit	John Mellor-Crummey, Rice Univ.
PAPI	Heike Jagode, Univ. of Tennessee
Dyninst	Barton Miller, Univ. of Wisconsin
TAU	Sameer Shende, Univ. of Oregon
Darshan	Wei-Keng Liao, Northwestern

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

Clarifying STEP's definition of ***Tools***

- We define Tools to mean "*the collection of tools and utilities that can be applied to both understanding performance bottlenecks and facilitating run time mitigation of performance degrading phenomena.*"

- These tools interact with hardware, compilers, communication libraries, programming model runtimes, operating systems, middleware layers, as well as the many applications that utilize these tools.

HPCToolkit: Measurement and Analysis of Complex Apps at Scale

2025 SCALABLE TOOLS

Current Capabilities on GPU-accelerated Architectures

NVIDIA
- heterogeneous profiles
- GPU instruction-level execution and stalls using PC sampling
- heterogeneous traces

AMD
- heterogeneous profiles
- no instruction-level measurements within GPU kernels
- measure OpenMP offloading using OMPT interface
- hardware counters to measure kernels
- heterogeneous traces

Intel
- heterogeneous profiles
- GPU instruction-level measurements with instrumentation; heuristic latency attribution to instructions
- measure OpenMP offloading using OMPT interface
- heterogeneous traces

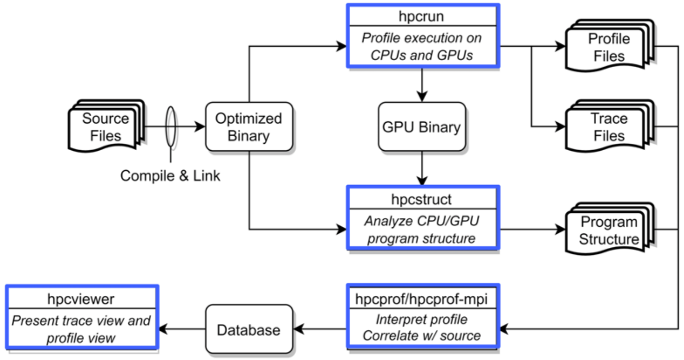

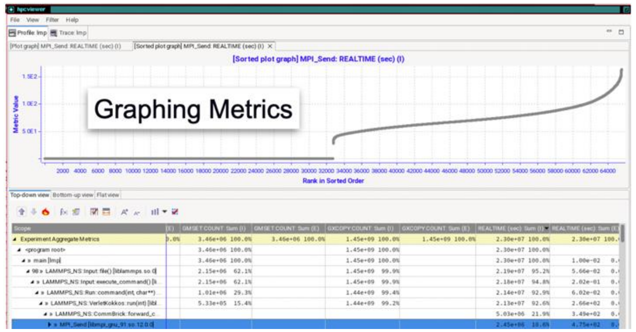

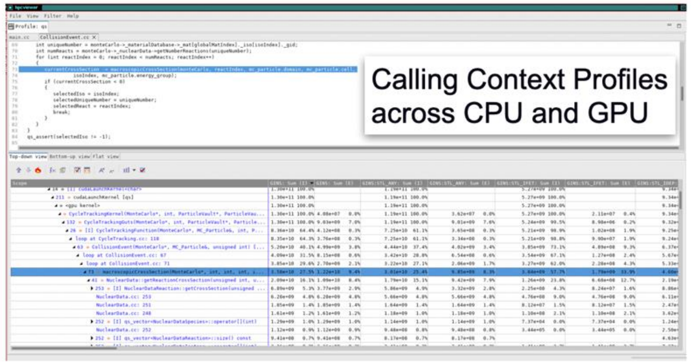

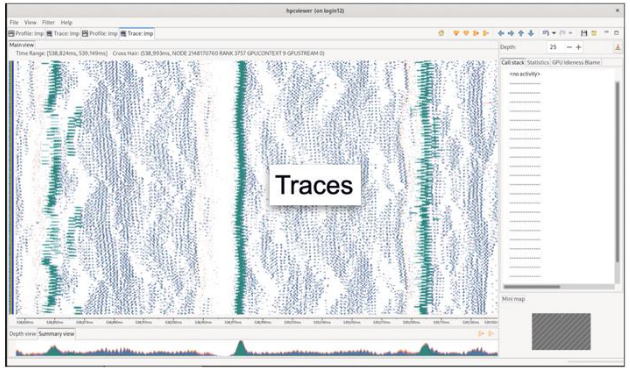

Sample Activities

- *Overhaul HPCToolkit to use emerging completely new AMD and Intel GPU APIs*
- *Extend HPCToolkit to exploit emerging support for instruction-level measurement of GPU computations using PC sampling on AMD and Intel GPUs*

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT   HPCToolkit   RICE

PAPI: Performance API

What is PAPI?

- PAPI provides a **universal interface** and methodology for **monitoring performance counters** and **managing power** across a diverse range of **hardware** and **software** components.

Monitoring Capabilities

- Robust support for **advanced hardware** resources, which includes major CPUs, GPUs, accelerators, interconnects, I/O systems, power interfaces, and virtual cloud environments.
- **Finer-grain power management** for dynamic optimization of applications under power constraints.
- **Software-Defined Events** originating from software stack (e.g., communication and math libraries, runtime systems).
- **Semantic analysis** of hardware counters so that scientists can better make sense of the ever-growing list of events.

Community Impact

- PAPI channels the monitoring capabilities of **hardware counters**, **power usage**, and **software-defined events** into a robust software package for HPC systems.
- PAPI provides the ability to monitor these metrics through a **single, portable API**.
- Without PAPI, applications must use multiple APIs to access all available metrics, which can hinder productivity.

PAPI performance and power measurements of a Kokkos application using Vampir / Score-P on NVIDIA GPUs.

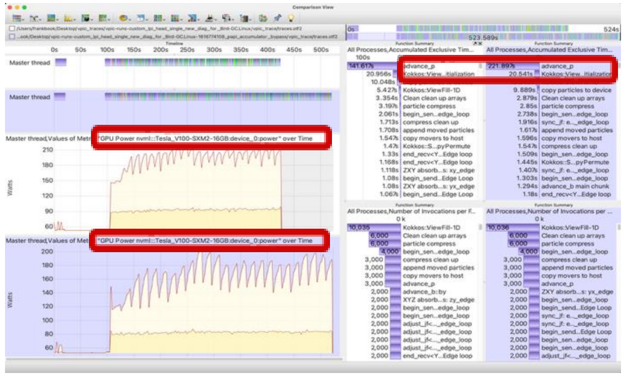

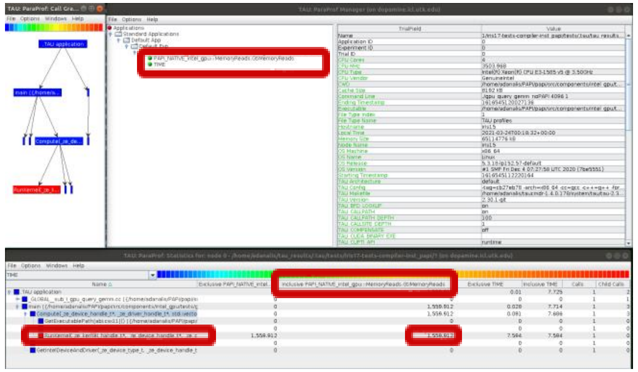

PAPI performance measurements of a math kernel using TAU on Intel GPUs.

Expected Future Impact from PAPI

- Monitoring capabilities of **new and advanced technologies** across the hardware-stack, e.g., GPU-to-GPU interconnect support for novel GPUs, rocprofiler v2 support for AMD GPUs.
- PAPI abstractions for expressing the **internal behavior of software** components and improved **interoperability** across the software-stack.

DyninstAPI: Binary Code Analysis and Instrumentation Toolkit

Current Capabilities on CPU Architectures

Includes x86, ARM and Power

Binary Code Analysis

- Abstract program representation avoids the need for architecture specific knowledge
- Full control flow analysis of binary code to identify functions, loops, and basic blocks
- Dataflow analysis to support slicing, symbolic evaluation and pointer analysis
- Operates on stripped code

Binary Code Instrumentation

- Both static (modify executables and libraries) and dynamic (modify running applications)
- Supports dynamically and statically linked code
- High level: Can instrument function entry, exit and call; loop entry, exit and top;
- Portable: DyninstAPI tools can work across any of the supported architectures.
- Fine grained: Can instrument at any instruction
- Efficient

Is the instrumentation foundation for HPCToolkit, AMD Omnitrace, TAU, Cray ATP, Red Hat SystemTap, and others

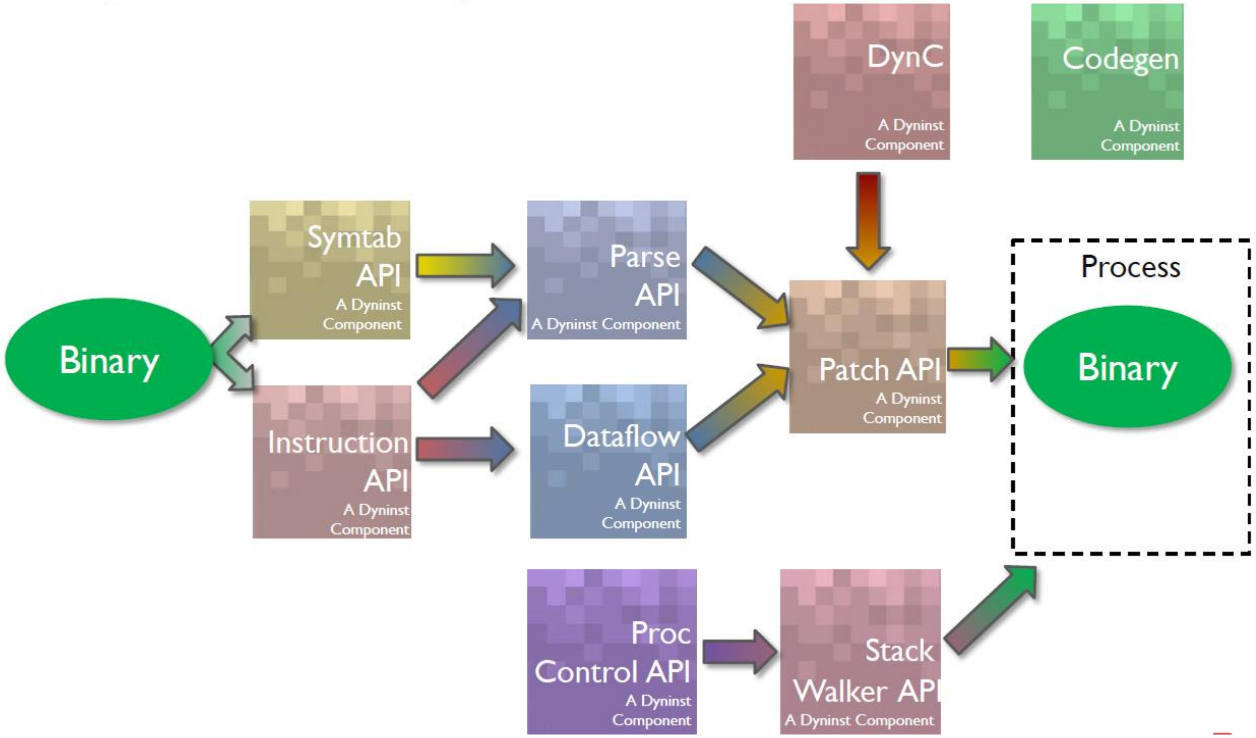

Sample Activities

- *Full Dyninst support for the newest x86 CPU families*
- *Extend Dyninst to support binary instrumentation of AMD GPU*

Slide source: Barton Miller

TAU Performance Evaluation Tool for GPU-accelerated Architectures

2025 SCALABLE TOOLS

Current Capabilities

- Integrated performance instrumentation, measurement, and analysis toolkit
- TAU is installed as a module at ALCF, LLNL, OLCF, LANL, Sandia, NERSC, and other supercomputing centers
- Supports DyninstAPI for instrumentation, PAPI for hardware performance counters based measurement, HPCToolkit and Darshan for profile display in TAU's 3D profile browser, ParaProf
- TAU supports Python instrumentation (TensorFlow, PyTorch, Horovod) and HPC runtimes (Kokkos, MPI, OpenMP)
- TAU supports tracing with OTF2 and visualization using Vampir (TU Dresden)
- tau_exec launcher requires no modification to the application binary!
- On AMD GPUs, TAU supports OMPT, ROCProfiler, ROCTracer, LLVM plugin for compiler-based instrumentation using hipcc, ROCTx
- On NVIDIA GPUs, TAU supports CUDA Profiling Tools Interface (CUPTI), NVTx
- On Intel GPUs, TAU supports OMPT, Intel Level Zero and OpenCL, Kokkos

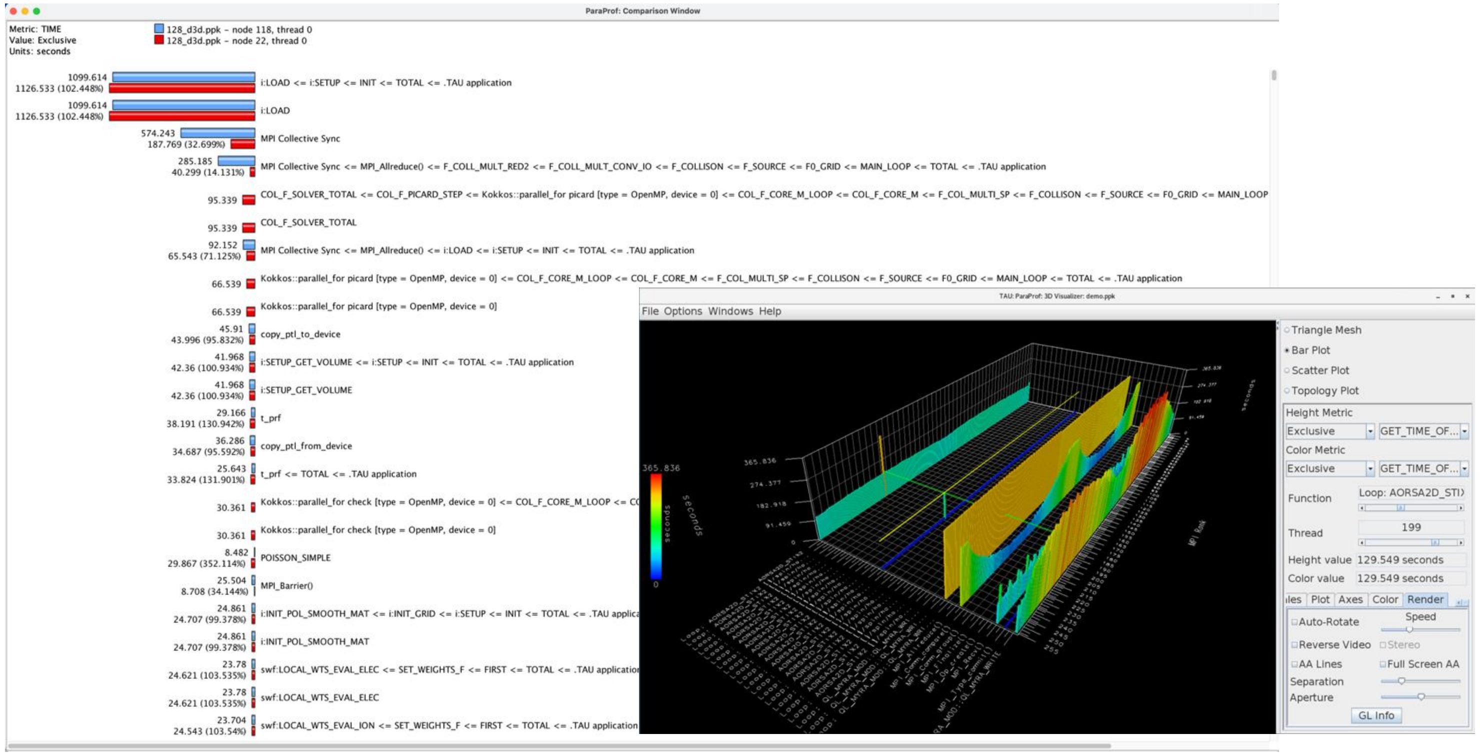

Sample Activities

- *Expand TAU's support for rewriting binaries and Dynamic Shared Objects using DyninstAPI*
- *Expand CI/CD support for TAU using GPUs from Intel, AMD, and NVIDIA on Frank at UO*

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

Slide source: Sameer Shende

Darshan: Enabling Insights into I/O Behavior of HPC Applications

Darshan is an application I/O characterization tool commonly deployed by HPC facilities

- Lightweight instrumentation methods to avoid perturbing applications
- Transparently enabled, requiring no modifications to users' code
- Generates summary reports of applications' I/O behavior

❶ Darshan runtime library

Transparently intercept application I/O calls, capture I/O characterization data, and generate log file when application terminates

Existing instrumentation modules span the HPC I/O stack: HDF5, PnetCDF, MPI-IO, POSIX, Lustre, etc.

❷ Darshan log analysis tools

New PyDarshan framework to enable extraction and analysis of Darshan log data using popular Python packages like pandas, matplotlib, etc

PyDarshan includes a job summary report tool to provide I/O overviews for users' logs

- E.g., I/O performance estimates (left) and I/O activity heatmaps (right)

files accessed	1026
bytes read	50.10 GiB
bytes written	49.30 GiB
I/O performance estimate	164.99 MiB/s (average)

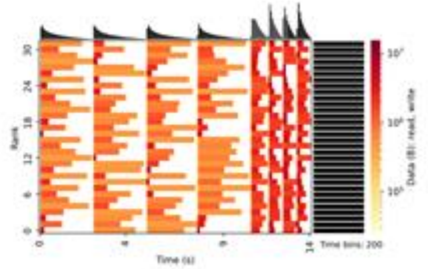

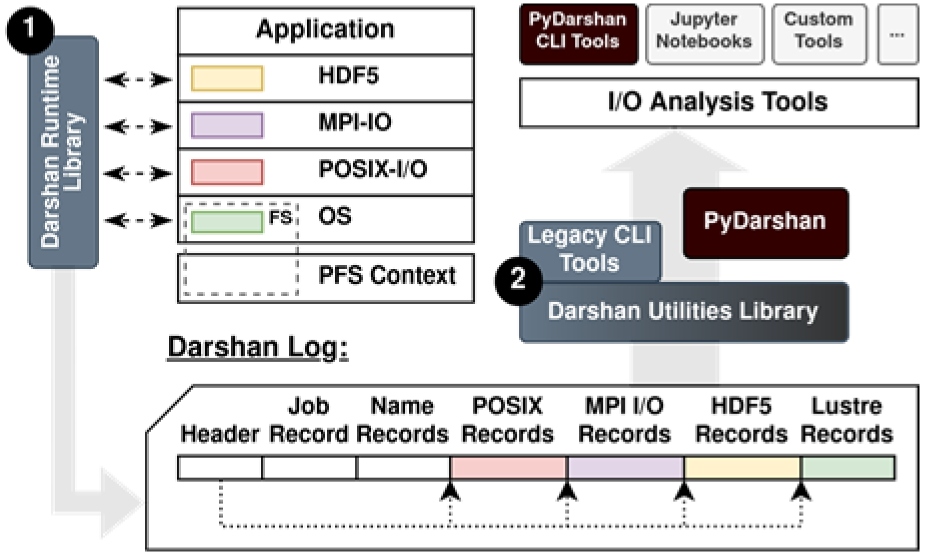

Sample Activities

- *Extend Lustre FS instrumentation to account for new progressive file layout (PFL) feature*
- *Improve design and capabilities of PyDarshan analysis tools*
- *Enhance CI testing capabilities using facility resources*

Slide source: Shane Snyder

STEP: Software Tools Ecosystem Project

- The Importance of High Performance Computing
- The US DOE and efforts to help the Tools Ecosystem
- STEP and its 5 Initial software packages
- How to Engage with STEP and looking forward

How STEP Can Be Transformative for Tools

Building An Enduring Ecosystem for Vital Tools

- "STEP will bring together a diverse community of High Performance Computing (HPC) tools developers and stakeholders to **develop plans for the sustainability of the HPC tools ecosystem.**"
- As DOE facilities update their testbeds and largest machines, STEP offers **support for vendor decisions and community-driven development directions necessary to sustain the HPC tools ecosystem."**
- "Further, many collaborative efforts lose steam when internal priority directions outweigh collaborative gain. Our approach **increases the collaborative gain by bringing together communities to address challenges caused by their dependencies."**

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

Engaging with STEP

- Web site: https://ascr-step.org
- Join CASS announcement mailing list or Participate in CASS working groups.
- Reach out to a STEP member software package for our events and news (most are open to public participation)
- Coming soon: CASS will be initiating a vendor forum to engage with vendors of hardware, software, and services.

Challenges

Technical challenges:

- Rapidly changing technologies
- Steady increase in complexity of these systems
- Software diversity in OS versions, programming languages, environments, schedulers, and libraries
- Support during growing CPU and GPU diversity / emerging opportunities with neuromorphic, quantum, ...

Non-technical challenges and project risks:

- Transitioning from ECP to NGSST funding levels
- Staff turnover (e.g., lead for Darshan (Shane Snyder) leaving Argonne for NVIDIA)
- Develop strategy for increasing use of tool packages—not just marketing/awareness but demonstrating value to scientists
- What are the best strategies to communicate with various stakeholders?
- Training is essential, but how do we accomplish it?

Opportunities

- Performance tools can be a vehicle to improve communication among vendors, facilities & application teams
- Performance tools can be a clearing-house for efficiently managing critical HPC technology
- Performance tools can be a resource for developing a skilled HPC workforce

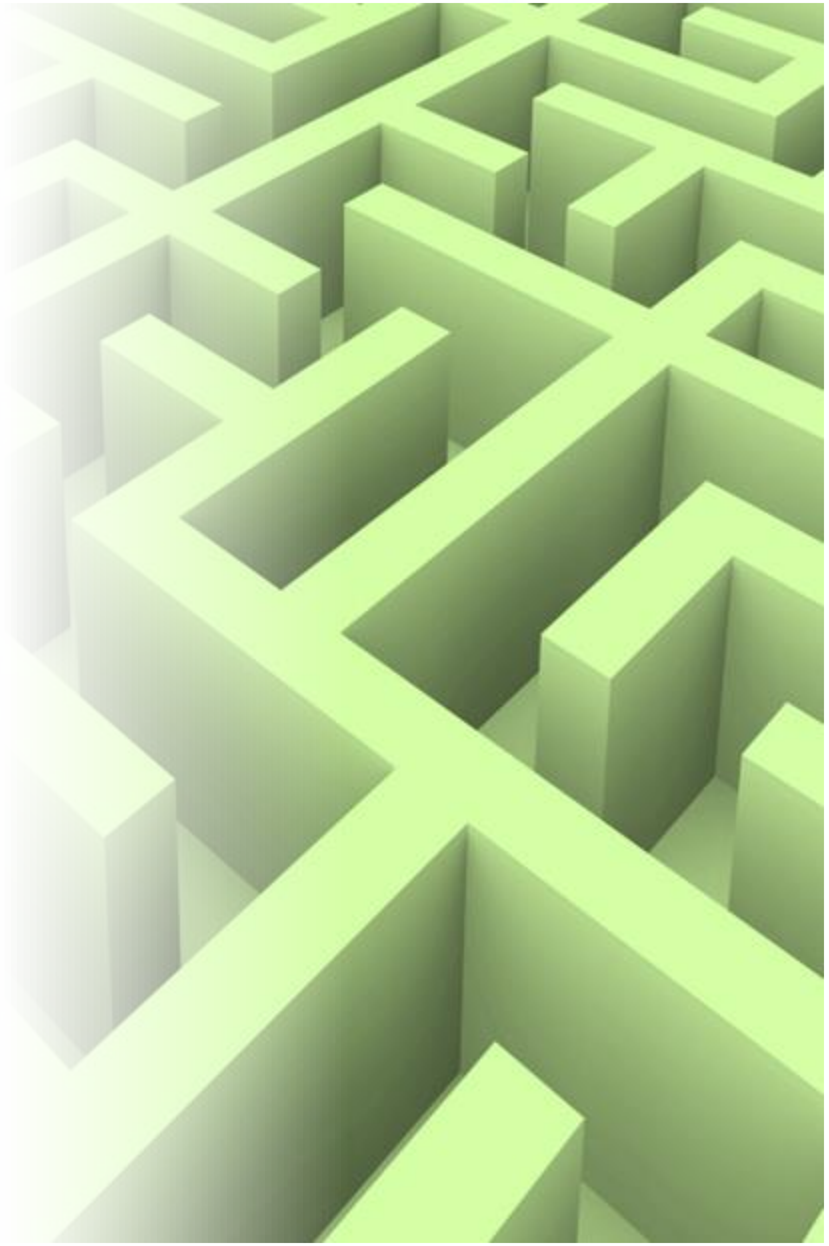

A Sampling of Current Activities

- Working together to develop a proactive stance on the rapidly evolving hardware landscape: new (completely different) tool APIs from AMD and Intel that include support for GPU PC sampling

- Working with vendors to help design their tool APIs, such as AMD's emerging rocprofiler-sdk and Intel's emerging PTI-GPU.

- Working with facilities and application teams to understand their emerging tool needs.

A Look Forward

Can Tools Transition from Reactive to Proactive?

A Look Forward

Can Tools Transition from Reactive to Proactive?

Can Tools Address the Memory Imbalance?

When We Look at Performance in Numerical Computations …

- Data movement has a big impact
- Performance comes from balancing floating point execution (Flops/sec) with memory->CPU transfer rate (Words/sec)
  - "Best" balance would be 1 flop per word-transfered
- Today's systems are close to 100 flops/sec per word-transferred
  - Imbalanced: Over provisioned for Flops

Machine Balance

Ratio of Fl Pt Ops per Data Movement over Time

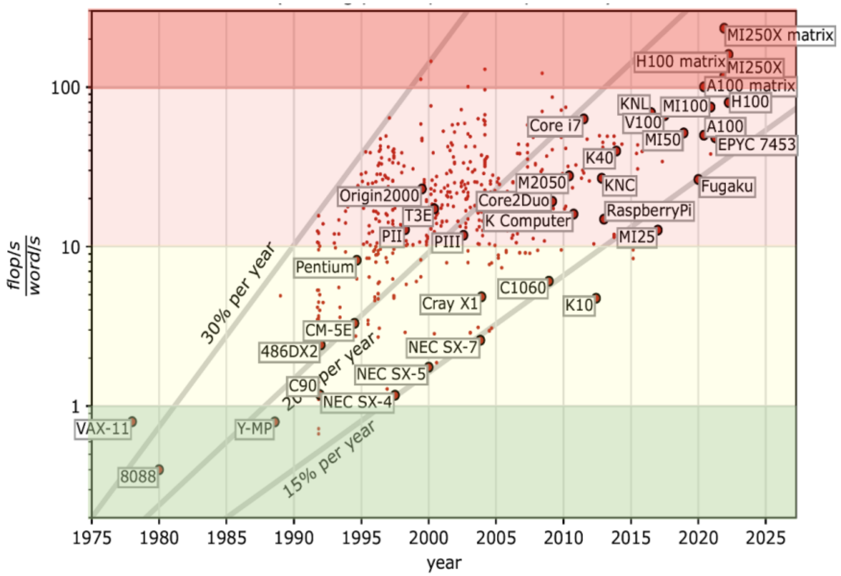

Graph from Mark Gates

Plot for 64-bit floating point data movement & operations
(Bandwidth from CPU or GPU memory to registers)

*Source: Jack Dongarra, https://new.nsf.gov/events/overview-high-performance-computing-future

A Look Forward

Can Tools Transition from Reactive to Proactive?

Can Tools Address the Memory Imbalance?

What's The Role for Tools In The Long Term? (Matt Welsh @ Harvard)

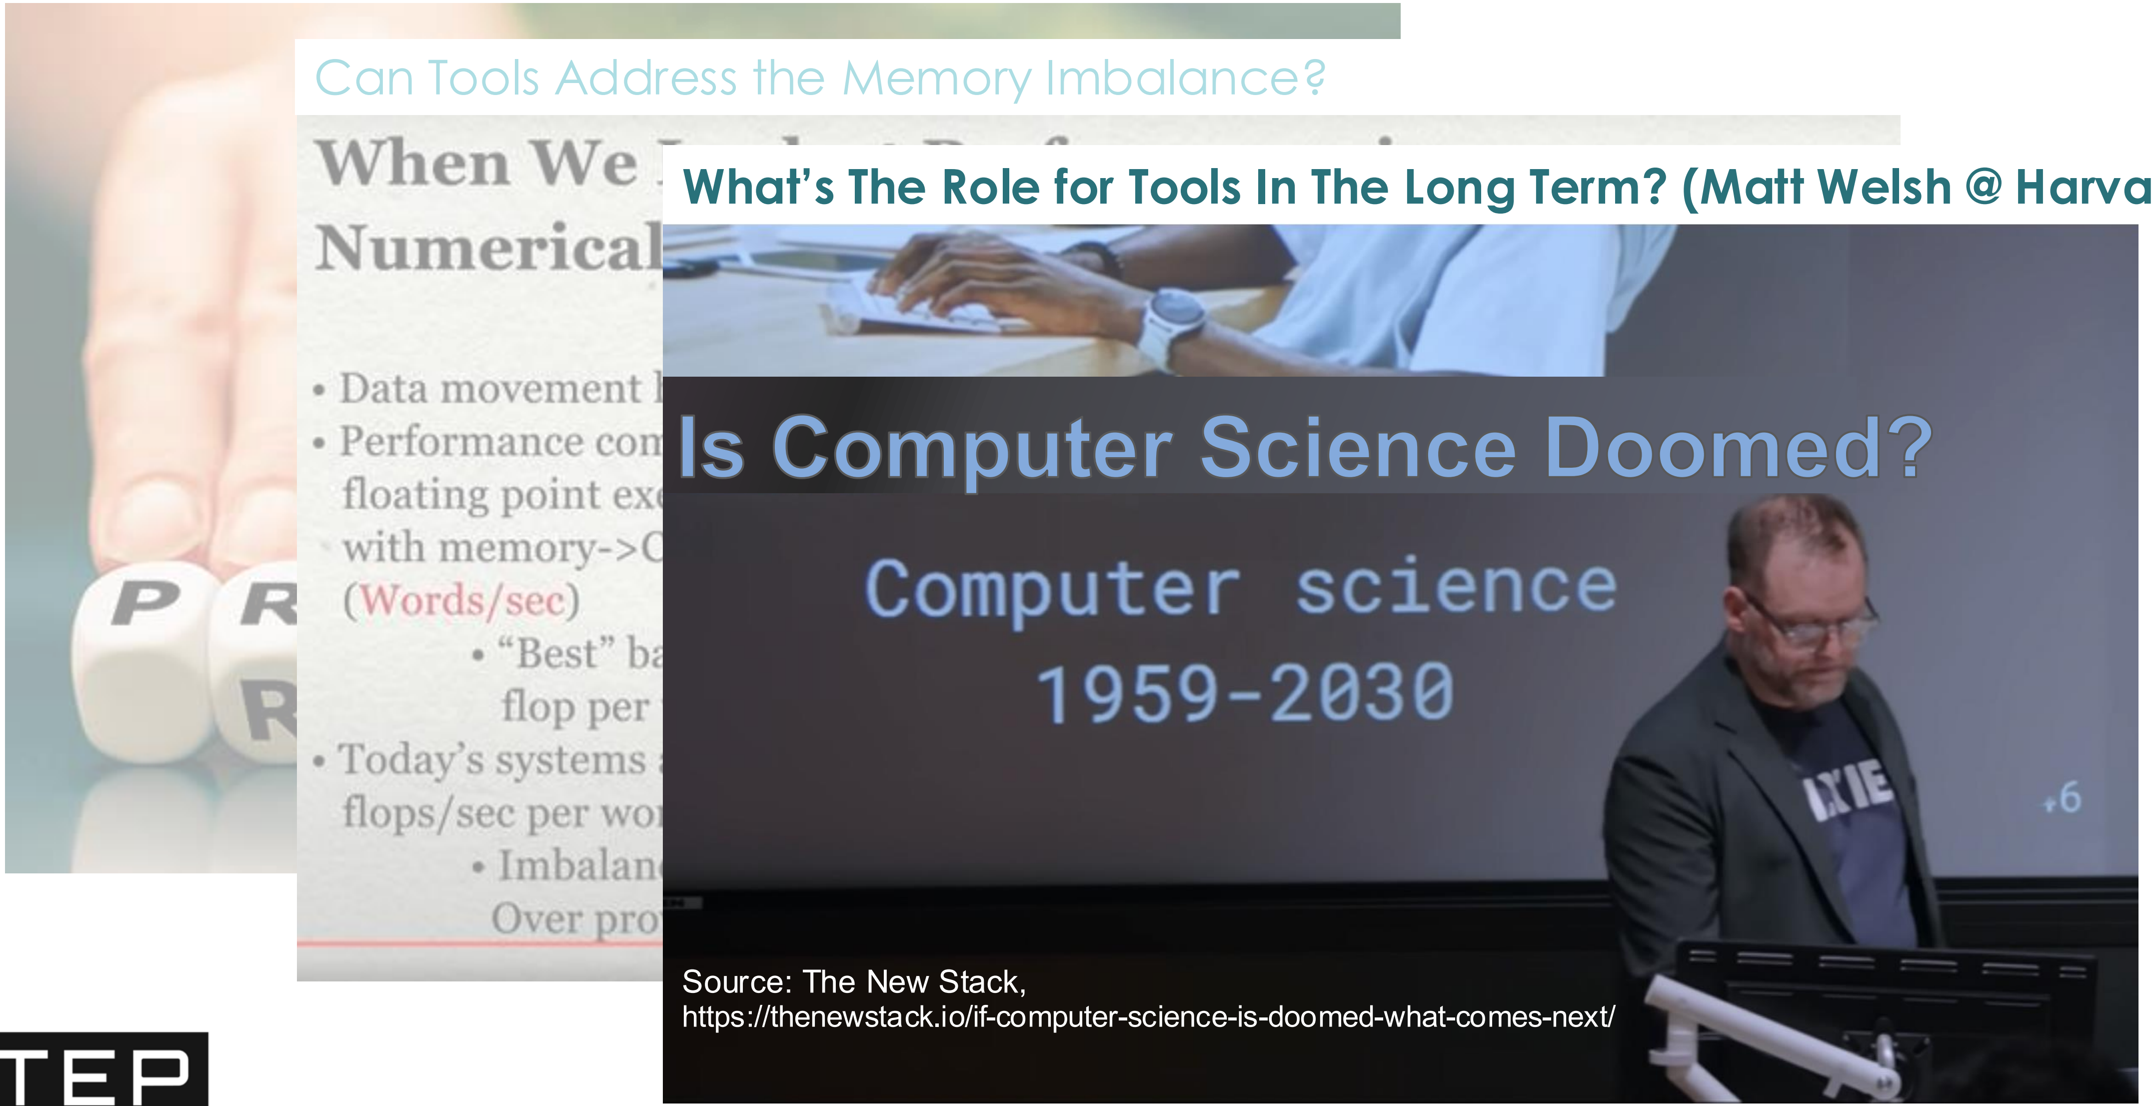

Source: The New Stack,
https://thenewstack.io/if-computer-science-is-doomed-what-comes-next/

The STEP Team

I'd Like to acknowledge the Software Tools Ecosystem Project (STEP) ***full Team*** *below* and our program manager ***David Rabson*** *of DOE's ASCR Office*

NAME	PERSPECTIVE	NAME	ORGANIZATION
Lead PI	Tools	Terry Jones	Oak Ridge National Lab
Deputy PI	Tools	Philip Carns	Argonne National Lab
Co-PI	Tools	James Brandt	Sandia National Laboratories
Co-PI	Vendors	James Custer	Hewlett Packard Enterprise
Co-PI	Tools	Ann Gentile	Sandia National Laboratories
Co-PI	Facilities	Richard Gerber	Lawrence Berkeley National Lab
Co-PI	Facilities	Kevin Harms	Argonne National Lab
Co-PI	Tools	Heike Jagode	University of Tennessee
Co-PI	Tools	Mike Jantz	University of Tennessee
Co-PI	Tools	Matthew Legendre	Lawrence Livermore Natl Lab
Co-PI	Tools	Wei-Keng Liao	Northwestern University
Co-PI	Vendors	John Linford	NVIDIA
Co-PI	Vendors	Keith Lowery	Advanced Micro Devices
Co-PI	Facilities	Verónica Melesse Vergara	Oak Ridge National Lab
Co-PI	Tools	John Mellor-Crummey	Rice University
Co-PI	Tools	Barton Miller	University of Wisconsin
Co-PI	Vendors	José Moreira	IBM
Co-PI	Applications	Erdal Mutlu	Pacific Northwest Natl Lab
Co-PI	Facilities	Suzanne Parete-Koon	Oak Ridge National Lab
Co-PI	Tools	Sameer Shende	University Oregon
Co-PI	Tools	Galen Shipman	Los Alamos National Laboratory
Co-PI	Tools	Devesh Tiwari	Northeastern University
Co-PI	Applications	Theresa Windus	Ames National Lab

STEP SOFTWARE TOOLS ECOSYSTEM PROJECT

STEP

SOFTWARE TOOLS
ECOSYSTEM PROJECT

STEP Member Organizations

AMD
The University of Wisconsin–Madison
8
Rice
Argonne National Laboratory
IBM
The University of Tennessee Knoxville
Northeastern University
Los Alamos National Laboratory
Ames National Laboratory
Hewlett Packard Enterprise
NVIDIA
Oak Ridge National Laboratory
Lawrence Livermore National Laboratory
Pacific Northwest National Laboratory
Sandia National Laboratories
University of Oregon

this presentation

DOI: https://doi.org/10.6084/m9.figshare.29483747
Speaker: Terry Jones, trj@ornl.gov
Date: July 7, 2025
Venue: 2025 Scalable Tools Workshop
Funding: *This work was supported by the U.S. Department of Energy, Office of Science, Advanced Computing Research, through the Next-Generation Scientific Software Technologies (NGSST) under contract DE-AC05-00OR22725*