DOE Exascale Computing Project
Update for the September ASCAC Meeting

EXASCALE COMPUTING PROJECT

Lori Diachin, LLNL
Project Director
Exascale Computing Project

Virtual
September 28, 2023

National Nuclear Security Administration

U.S. DEPARTMENT OF ENERGY | Office of Science

ECP is in the homestretch!

- A $1.8B project started in 2016
  - 15 DOE Laboratories
  - Dozens of university and industrial partners
  - 1000 participants
  - Three technical focus areas and 81 research projects
- ECP has reached threshold in the key performance parameters!
- Three months of technical work remain
- The ECP leadership team is increasingly focused on project closeout, communications, and outreach

We are capturing and documenting the ECP impact and legacy

- A suite of applications that will impact problems of national importance for decades to come
- Integrated software stack for GPU-accelerated computing widely available for use
- Best practices and lessons learned for thinking about how to program GPUs – moving the nation forward
- Over 1000 researchers trained and ready for accelerator-based computing
- Best practices for running a large-scale software development RD&D 413.3b project

The ECP KPP's are complete!

KPP ID	Description of Scope	Threshold KPP	Objective KPP	Threshold Completion Status (8/17/22)
KPP-1	11 selected applications demonstrate performance improvement for mission-critical problems	✓ 6 of 11 applications demonstrate Figure of Merit improvement ≥50 on their base challenge problem	All 11 selected applications demonstrate their stretch challenge problem	✓ **9 SC apps (5 complete, 4 under review)**
KPP-2	14 selected applications broaden the reach of exascale science and mission capability	✓ 5 of 10 DOE Science and Applied Energy applications *and* 2 of 4 NNSA applications demonstrate their base challenge problem	All 14 selected applications demonstrate their stretch challenge problem	✓ **6 SC apps (3 complete, 3 under review)** ✓ **3 ATDM apps under review**
KPP-3	76 software products selected to meet an aggregate capability integration score	✓ Software products achieve an aggregate capability integration score of at least 34 out of a possible score of 68 points	Software products achieve the maximum aggregate capability integration score of 68 points	✓ **40 integration points (25.5 complete, 5.5 SME approved, 9 under SME review)**
KPP-4	Delivery of 267 vendor baselined milestones in the PathForward element	✓ Vendors meet 214 out of the total possible 267 PathForward milestones	✓ Vendors meet all 267 possible PathForward milestones	✓ **267 PathForward milestones**

EXASCALE COMPUTING PROJECT

Achieving our KPPs would have been impossible without close collaborations with the ASCR HPC facilities

- Full system available to all ECP 4/2023
- All application KPPs to date run on Frontier; seeing excellent performance
- Many ST / Co-design KPPs achieved on Frontier
- Large scale jobs still seeing some stability issues; work arounds in place that leverage fast SSD capabilities for checkpoint/restart

- All ECP teams have run on Frontier: ~5M node hours
- 6 ECP teams have INCITE allocations: ~3.6M node hours

- Full system installation completed 6/2023
  - Currently in the scaling test phase; usual issues seen with serial #1 and large-scale systems
  - Lustre, interconnect, and some hardware issues being worked
- Limited access available July 2023 (ANL liaisons only); full system access expected October 2023
- ~10% of ST KPP integrations achieved on Sunspot; major test for performance portability in final quarter

- 61 ECP teams (186 users) have run on Sunspot (Aurora TDS): ~ 9K node
- 12 ECP teams (14 users) have run on Aurora: ~ 24K node hours

ExaSMR develops first-of-a-kind simulations of advanced nuclear reactors such as small modular reactors

PI: Steven Hamilton, ORNL

- **Objective**: Help DOE meet its goal of an operational small modular reactor (SMR) in 10 years, a key part of the Department's goal to develop safe, clean, and affordable nuclear power options

- **ECP accomplishments**
  - First fully coupled, fully resolved simulation of nuclear reactor core with coupled Monte Carlo neutron transport and CFD; achieved nearly 100X overall performance improvements in the science work rate
  - Computational fluid dynamics solver reached 1.4 trillion degrees of freedom per second on Frontier, an increase of 400X from the beginning of ECP
  - Refactored neutron transport algorithms from particle-based to event-based, improving performance on GPUs; tracked 1.1B particles/second on Frontier, an increase of 113X from the beginning of ECP
  - Significantly improved the accuracy of simulations for neutron transport and fluid dynamics using performance portability tools and state-of-the-art physics models

Small modular reactor; complex geometries that require fully coupled simulations to resolve physical phenomenon. Image reproduced with permission of NuScale

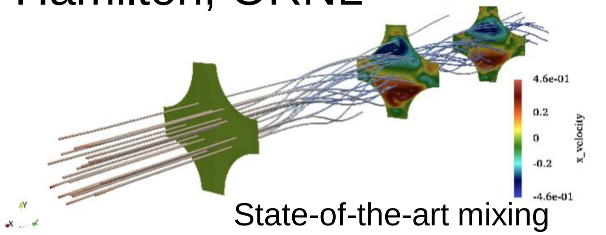

State-of-the-art mixing vane model resolved with CFD solver; used to optimized heat transfer in fuel rod bundles

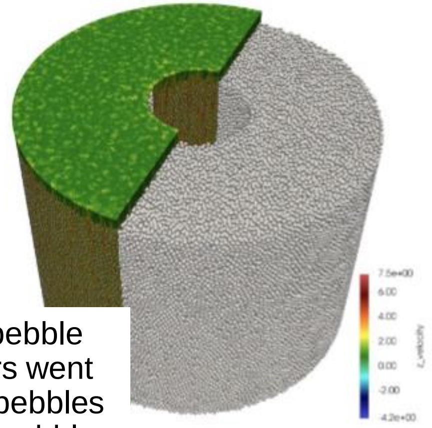

Models of pebble bed reactors went from 1000 pebbles to 350,000 pebbles through ECP advancements

Deliver experiment-quality simulations of reactor behavior to enable the design and commercialization of advanced nuclear reactors and fuels at significant savings in time (from years to months) and money

HIP on Aurora – Bringing HIP to Intel GPUs

PI: Brice Videau, ANL

Rationale and Scope

- Intel GPUs lack native support for the CUDA or HIP programming models
- A large fraction of GPU enabled HPC applications are written using CUDA or HIP
- Bring AMD's HIP and Nvidia CUDA to Aurora
- Approach: develop new HIP backend based on the new Intel GPU runtime called Level Zero and existing OpenCL APIs
- Port required math libraries and tooling
- Bring real application support to this new HIP backend

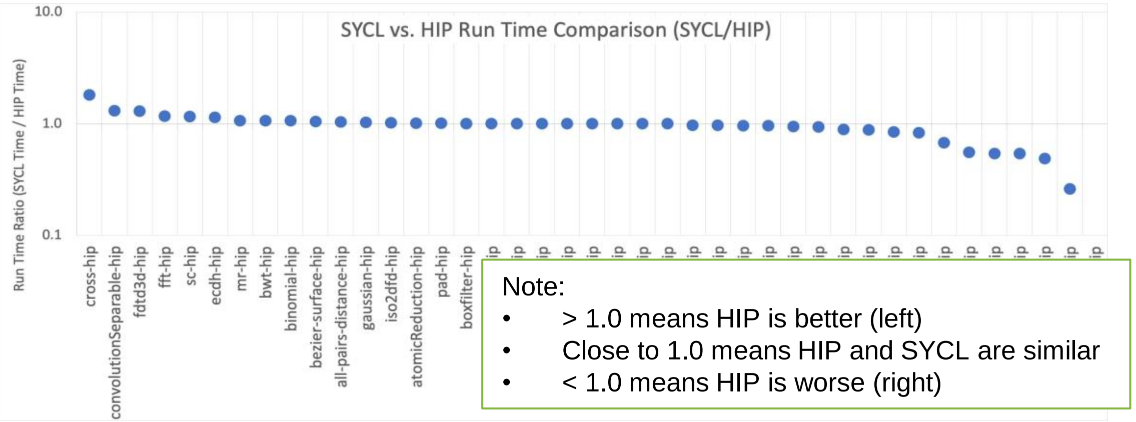

- On Aurora PVC GPUs (Credit to Kali Balleda (Intel))
- Using HecBench: HIP vs SYCL version of benchmarks with verified results

Progress

- Version 1.0 release of chipStar https://github.com/CHIP-SPV/chipStar ; provides HIP support on LLVM 16 and 15
- Supported applications: CP2K, GAMESS, libCEED
- Supported HIP libraries: HipBLAS, hipSOLVER, hipRTC
- Upcoming support for ExaBiome, hipFFT, and more...
- Experimental CUDA support

Impact

- Reduce effort for applications to move to Aurora
- Reduced porting and maintenance effort for ECP applications with HIP or CUDA backend
- Ecosystem improvements: better tooling to translate between programming models
- Post ECP: Continued partial support from ALCF, SASPT seedling, and Intel; Work with AMD to push HIP as an open standard

ECP-funded technologies and participants are leadership class

R&D 100 AWARDS

CANDLE (2023)
ZFP (2023)
Varorium (2023)
Flash-X (2022)
Mochi (2021)
SZ (2021)
Flux (2021)
Legion (2020)
Spack (2019)
UCX (2019)
Darshan (2018)
Swift/T (2018)
Nek5000 (2016)

HPCwire People to Watch

Jeff Vetter (2021)
Bronis de Supinski (2021)
Lori Diachin (2019)
Doug Kothe (2018)

acm

Gordon Bell Finalist | ExaSMR (2023)
Gordon Bell | WarpX (2022)
Gordon Bell Special Prize | CANDLE (2022)
A. M. Turing Award | Jack Dongarra (2021)
ACM Fellow | Bronis de Supinski (2022), Rob Ross (2021), Rick Stevens (2020)
IEEE CS Ken Kennedy Award | Ian Foster (2022)
IEEE CS Sidney Fernbach Award | Salman Habib (2020)

…and many more notables, including

CRA Distinguished Service Award | Paul Messina (2018)
Best Open Source Contribution (For IPDPS 2023 paper) | ExaIO (2023)
SIAM/ACM prize in Computational Science and Engineering | SUNDIALS (2023)
The Ernest Orlando Lawrence Award | ExaStar and DataLib (2021)
SIAM Fellow | Mike Heroux (2019); IEEE Fellow | Rajeev Thakur (2022)

ECP has a significant focus on training the next generation to create a pipeline for the DOE HPC workforce

Sustainable research pathways for HPC

Summer 2022:

- 61 participants at 9 labs: 13 student track, 16 faculty track (+29 students), 3 self-funded students
- 82% of overall participants represent at least
  1 element of diversity

Summer 2023:

- Multi-lab CRLC program spanning ECP and other computational and data science projects
- 189 faculty & students at 10 labs (120 funded thru ECP, 69 funded thru labs

Summer 2024:

- ECP launch via SRP Matching Workshop, Oct 31-Nov 3, 2023
- Targeting @ 100 participants


SRP-HPC students and faculty, summer 2022


SRP-HPC mentors/co-mentors

Dan Martin, LBL SRP thrust lead for ECP

Keisha Moore, SHI SRP Program Coordinator

Intro to HPC Bootcamp (New)

- Hosted at NERSC, August 7-11, 2023
- 60 students from 22 states
  - 48% first-generation scholars
- 14 project leaders & trainers (ANL, LBL, ORNL)
- 7 group projects: HPC topics related to energy justice
- Many have now applied to be part of SRP matching workshop for 2024
- **Website:** https://shinstitute.org/introduction-to-high-performance-computing-bootcamp

APTESC

- Incorporated into ECP at project inception; 2 week training program at ANL focused on late graduate career, postdocs, early career scientists
- *“There are numerous HPC-related training offerings available across the community, but ATPESC is truly the standard by which all others are measured.”*   -- Eric Neilsen, NASA Langley, Lead of FUN3D

EXASCALE COMPUTING PROJECT

We have been very active with our industry and agency council since our last update

NASA/ST deep dive

- July 18, 2023, Virtual
- 67 attendees (47 NASA, 20 ECP) including Ames, Langley, Goddard, Marshall, and HQ
- Overview of NASA software and hardware environments at NASA Ames and Goddard
- E4S technologies including Spack, I/O, Compression and use in ECP applications

IAC Quarterly Call

- August 17, 2023
- Update from ECP
- Communications and outreach strategy
- Software sustainability for ST and AD

NASA collaboration points

- Use E4S/Spack to streamline LAVA dependency management
- Explore use of Kokkos instead of CUDA and in situ visualization tools in LAVA
- Explore GPU in situ data processing and I/O of engineering outputs in FUN3D

NOAA collaboration points

- E4S use on the cloud
- Performance portability using Kokkos and OpenMP
- Adaptive mesh refinement using AMReX

ECP is moving toward to project closeout activities

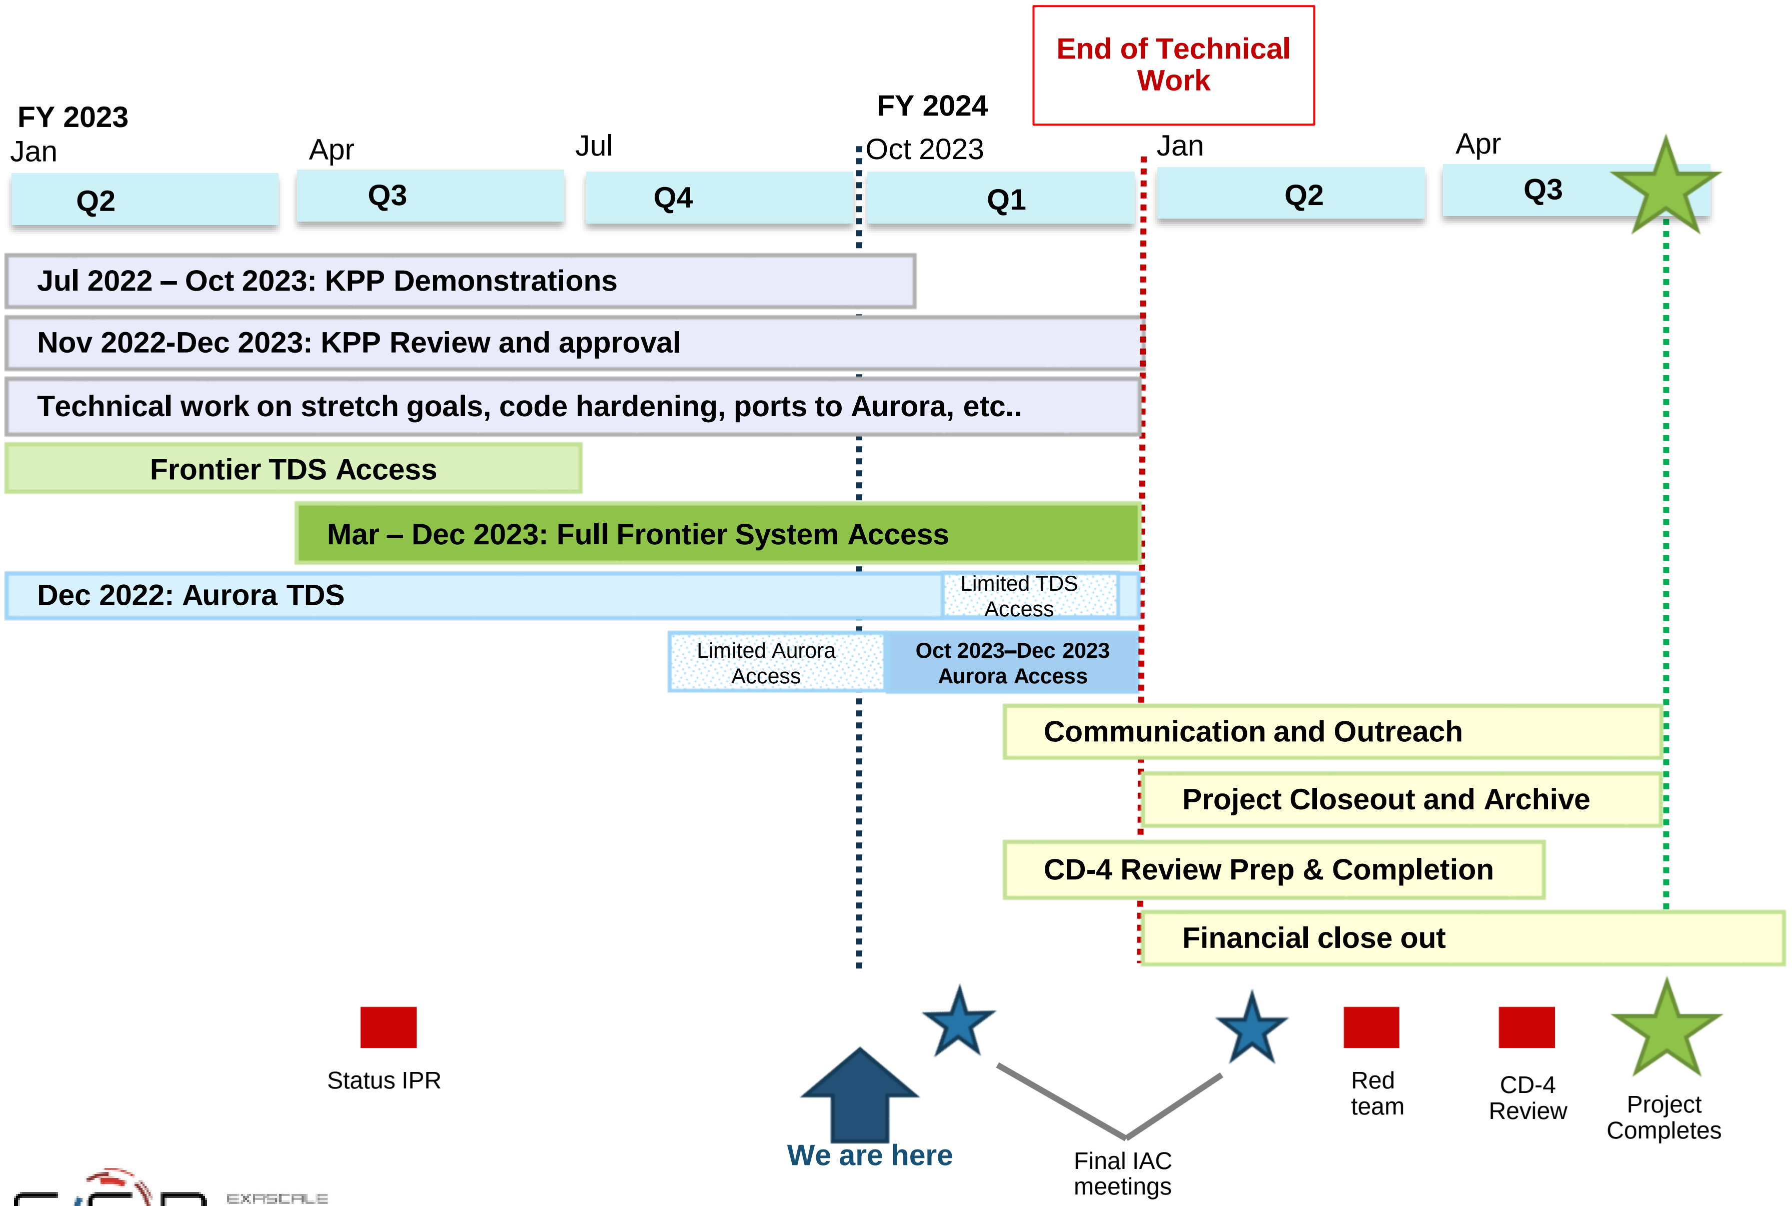

Scope for the final quarter of technical work includes:

- KPP run completions, write up and review for any teams not already done
- KPP stretch goals
- Porting to and running on Aurora
- Final milestone reports
- Outreach activities including conferences, journal articles, etc.
- Code documentation and hardening

ECP Communication Strategy

- ECP has increased staff resources to 'spread the word' to a broad audience
  - Integration case studies
  - Exascale and Industry
  - ECP and US Agencies
  - Then and Now articles
  - Application impact slides/stories
- Other planned activities
  - ECP Book
  - Significant presence at SC23
  - Stakeholder briefings
- ECP is working with the IAC on spotlight articles and other media communications

Articles recently published

Exascale Drives Industry Innovation for a Better Future

"Exascale is a massive accelerator for technology, productivity, engineering, and science."

− Dave Kepczynski, chief information officer for GE Research and co-chair of ECP's Industry and Agency Council

For the full article see https://www.exascaleproject.org/exascale-drives-industry-innovation-for-a-better-future/

Building a Capable Computing Ecosystem for Exascale and Beyond

The exascale computing era is here.

For the full article see https://www.exascaleproject.org/building-a-capable-computing-ecosystem-for-exascale-and-beyond/

ECP will have a significant presence at SC23

The Impact of Exascale and the Exascale Computing Project on Industry

Thursday, Nov 16 at 3:30 pm MST

Moderator: Fran Hill, DoD

Speakers:

- Peter Bradley, Raytheon Technologies
- Rick Arthur, GE Aerospace Research,
- Sean Dettrick, TAE Technologies, Inc.
- Mike Townsley, ExxonMobil

Lois McInnes will be an SC23 invited speaker: **Broadening Participation in HPC: Together we can make a difference**

Lori Diachin will be a speaker in the SC23 Press session on Exascale Computing

ECP also to be seen at the Gordon Bell Prize (finalists); DOE Booth; and many talks, tutorials, and posters

EXASCALE COMPUTING PROJECT

We will hold two final in-person IAC meetings and a virtual workshop

October 3, 2023	**Continuous Integration/Continuous Deployment Workshop** Nearly 70 registered participants Virtual	
Oct 24 – 25, 2023	**Fall Meeting 2023** IAC members and ECP leadership Focus on stewardship and computing in a post-ECP world Washington, DC	
January/early February 2024	**Winter Meeting 2024** Two-day meeting with plenaries and breakouts Expecting approximately 100 people Argonne Conference Center	

Altair, Ansys, Boeing, Cascade Technologies, Chevron, Cummins, EPRI, ExxonMobil, GE, gm, Goodyear, KatRisk, Lilly, P&G, Raytheon Technologies, tae Technologies, Westinghouse, Whirlpool Corporation, NOAA, NSF, NASA, NIH National Institutes of Health

ECP EXASCALE COMPUTING PROJECT

Leadership team will remain engaged through ESAAB (Q3/FY24)

- Finalizing project documentation
  - Final milestone reports
  - Financial closeout plan
  - Transition to operations plan
  - Project closeout report
  - Lessons learned
- Preparing for the CD-4 review; following up on any recommendations
- Helping archive project artifacts
- Continuing to work on post-ECP sustainability (transitioning ECP technologies to ‘operations’)

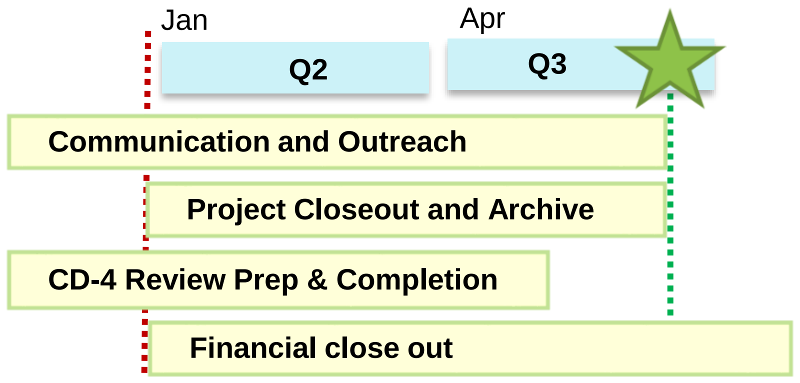

- Conducting technical outreach – successes and impact of ECP
  - Final IAC meeting
  - SIAM PP24 meeting
  - Stakeholders
- Financial close out will be a major endeavor
  - 390 active MPOs to close out
  - 150 subcontracts to universities

EXASCALE COMPUTING PROJECT

Lessons Learned

Lessons Learned: Technical

- Performance portability, programming models, strategies for increasing arithmetic intensity, refactoring code, new algorithm design, etc..
- Strategies for moving from GPU-accelerated to GPU-resident
- When facing an inflection point in the HW, S/W investment must be a first class citizen
- Node-level solutions apply at all levels of computing

Lessons Learned: Collaboration

- The value of diverse, multi-disciplinary teams
- Collaborative solutions can't be dictated but they can be incentivized
- Build integration into project structure and measures of success
- High quality software is the foundation for collaboration in scientific computing
- Open, frequent comms (good/bad/ugly) imperative with sponsors, stakeholders, staff

Lessons Learned: Project Management

- Projectizing R&D works if agile project management are aggressive change control are in place
- Empower the leadership team then hold them accountable
- Understand and manage external dependencies
- Highly functioning diverse leadership team are a must
- Good centralized project management tools do not guarantee success but bad ones can sure impede progress
- Stability and long term planning results in remarkable innovation

The community has recognized the impact of ECP

Dave Kepczynski,
GE Research

"Over the last five to six years, as a result of the nation's Exascale Computing Initiative (ECI) and the significant progress made as part of ECP, we are now **delivering on the promise of next-generation computing capabilities** that enable these significant advances in technology and breakthrough products that have real implications for the future."

Earl Joseph,
Hyperion Research

"When combined with the new exascale class computers, the ECP efforts have **enabled researchers to move their science forward more quickly** and will help them incorporate new technologies like AI, machine learning and large language models. This will make the leadership computers dramatically more productive and useful to the nation."

Addison Snell,
Intersect360 Research

"The goal wasn't to build an exascale computer just to check a box. It was to build a computer capable of addressing the next generation of scientific challenges. The role of the Exascale Computing Project (ECP) was to ensure the system did more than turn on. We had to program it. We had to power it sustainably. We had to use it. **Because of ECP, the United States' exascale computing capability is a breakthrough resource for discovery and advancement**."

RTX Pratt and Whitney

""The holistic approach taken by the Exascale Computing Project is **already paying dividends to industry and state-of-the-art technology**. As the project built the most powerful computers in history, it also built out the supporting software, technology infrastructure, and human experience necessary to hit the ground at full speed when the systems came online. The results will positively impact the lives of practically everyone on earth through advances in fields as diverse as energy efficiency, medicine, travel, climate, and physics."

ECP EXASCALE COMPUTING PROJECT

ECP Application Development Update

Andrew Siegel
Director of Application Development, Exascale Computing Project
Senior Scientist, Argonne National Laboratory

ASCAC Meeting, Sept 28, 2023

AD: Three key lessons in achieving Exascale in ECP

1. Code porting exercise
2. “Over the fence” producer-consumer enabling software
3. Do-it-yourself application groups with minimal external dependencies

1. Holistic rethinking and restructuring of implementation, algorithm, and models
2. Software/application tight co-design loop: Develop → Deploy → integrate → Accrete
3. Leverage community tools and expertise broadly – E4S, LCFs, Centers of Excellence, teams with highly diverse expertise

EXASCALE COMPUTING PROJECT

KPP Verification Workflow

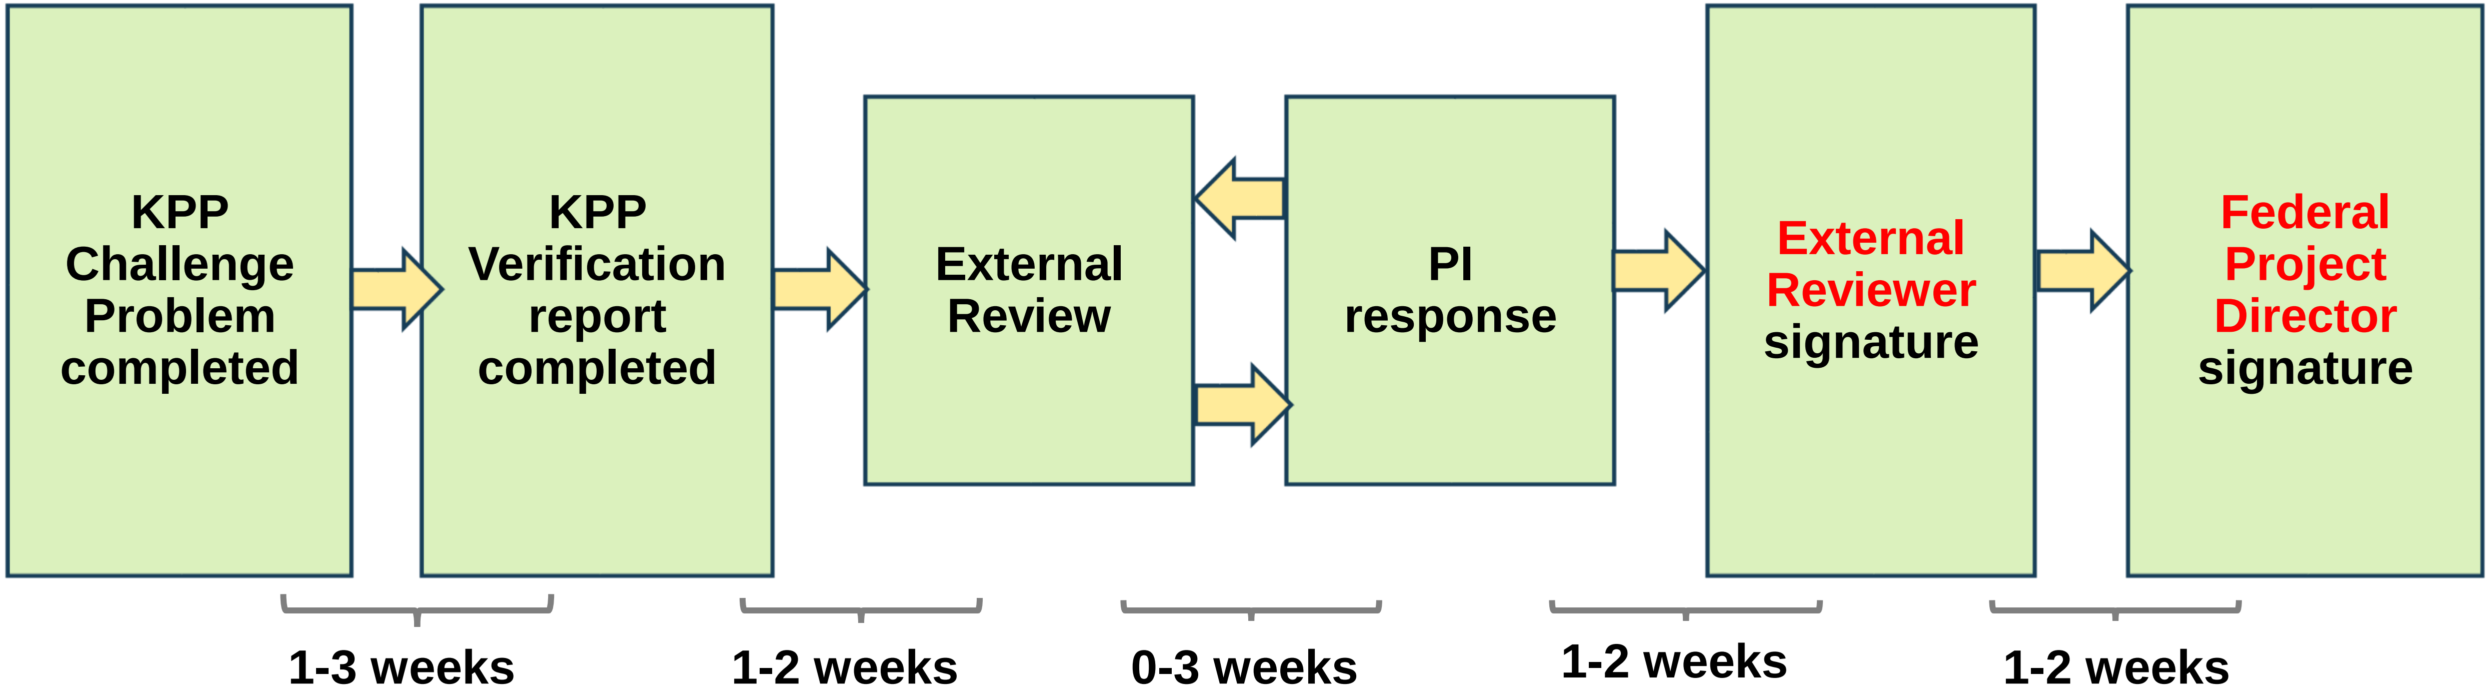

Status of KPP-1 applications on Frontier

Threshold: 6/11

KPP in progress	KPP completed	KPP under review	KPP signed (SME)	KPP signed (FPD)
NWChemEx		EXAALT →	LatticeQCD	ExaSMR
QMCPACK		E3SM-MMF	CANDLE	WDMApp
				WarpX
				ExaSky
				EQSIM

EXASCALE COMPUTING PROJECT

Status of KPP-2 applications on Frontier

Threshold: 5/10

KPP in progress	KPP completed	KPP under review	KPP signed (SME)	KPP signed (FPD)
ExaWind		GAMESS →		Combustion-PELE
ExaStar		ExaAM		MFIX-Exa
Subsurface		ExaSGD		ExaBiome
ExaFEL				

Six Remaining Projects with KPPs In Progress

Project	KPP Class	Risk	Comment
QMCPACK	KPP-1	4	OpenMP compiler bugs; latest ROCm update hopeful
NWChemEx	KPP-1	2	Steady progress: aiming to finish by mid-October
ExaFEL	KPP-2	3	Steady progress scaling up throughput. Some questions about how to demonstrate KPP without reservation since KPP measures throughput.
ExaWind	KPP-2	2	Good news: ran 150 steps of challenge problem on Frontier. Attempting longer run.
ExaStar	KPP-2	3	Stuck on Fortran compiler bug with Thornado. Elevated to highest priority. Executing CP on Summit.
Subsurface	KPP-2	2	Good news: resolved Hypre/PETSc bug; prepping for KPP run, but hampered by slow queue turnaround on Frontier

1: Near certainty
2: Likely
3: Possible
4: Unlikely
5: Would take a miracle

What Programming models were used to meet KPPs?

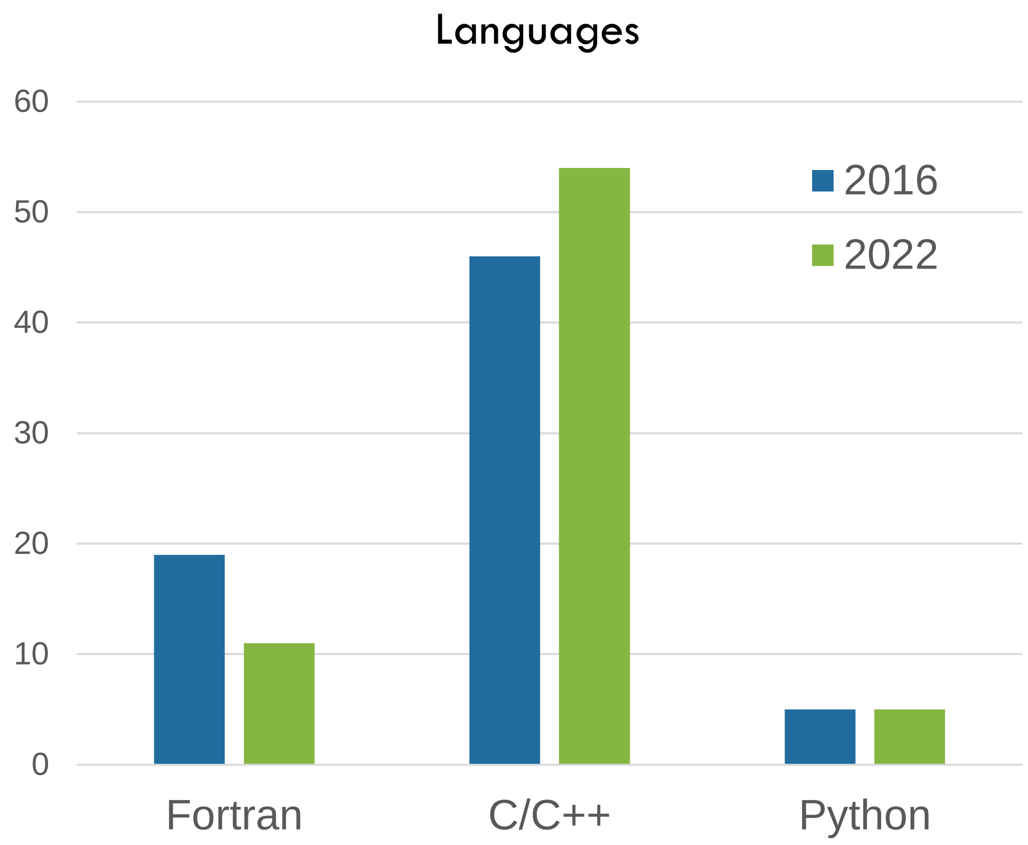

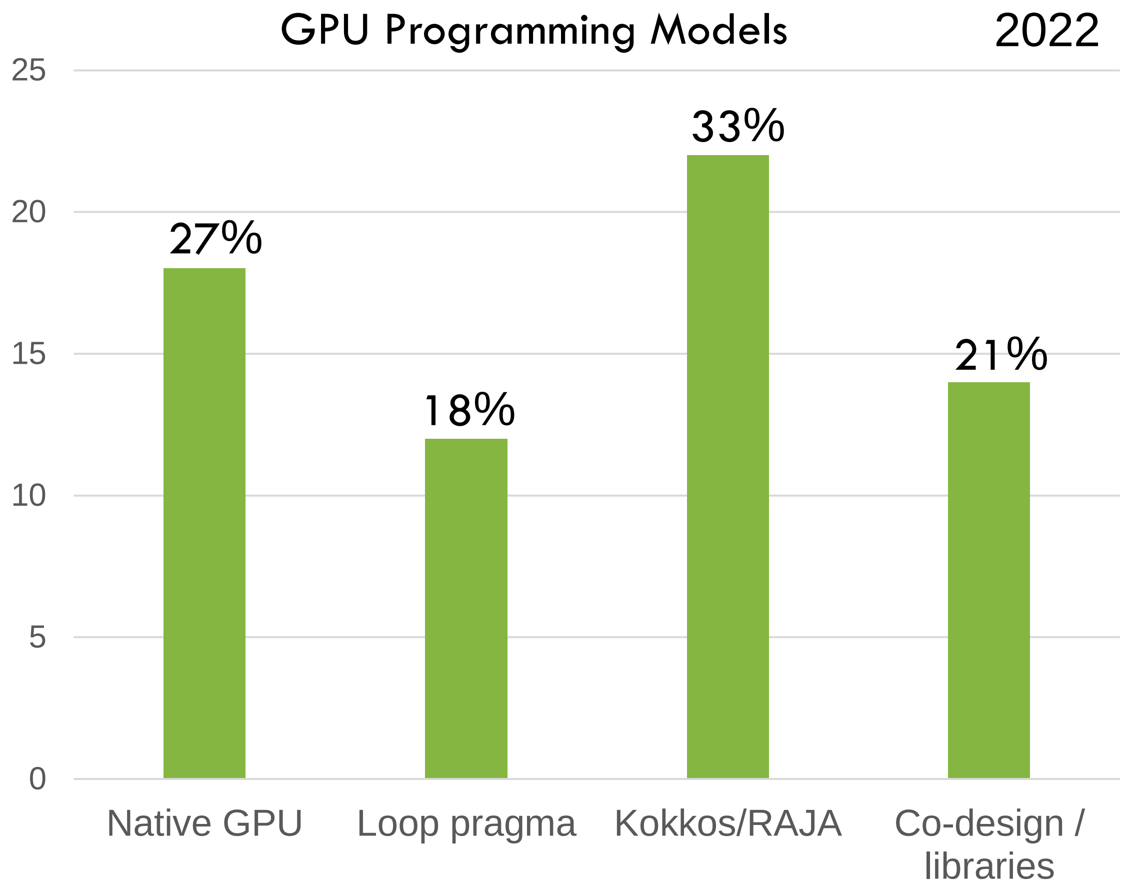

Many ECP applications started out using native GPU and loop pragma models before moving to C++ abstractions and co-design libraries

EXASCALE COMPUTING PROJECT

Evans TM, Siegel A, Draeger EW, *et al.* *"A survey of software implementations used by application codes in the Exascale Computing Project."* The International Journal of High Performance Computing Applications. 2022;36(1):5-12.

A Deeper Look at KPP-1 and KPP-2: Background and Lessons Learned

Generation 1: General Definition of KPP-1 and KPP-2

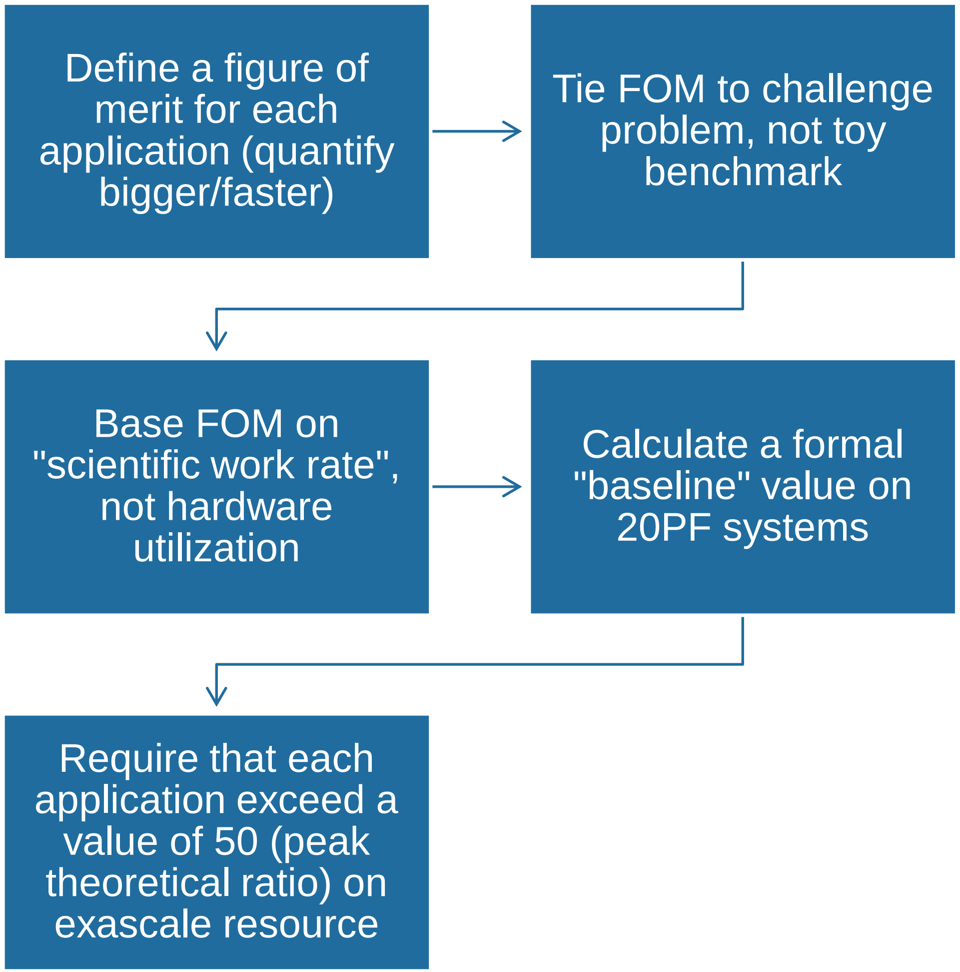

Problem: Not everything fits in the FOM model.

- New codes?
- Existing codes doing totally new things
  - ➢ No way to define a baseline
  - ➢ Want to encourage innovation

Solution: **Define a second KPP (KPP-2) based on a demonstration of capability** while making "efficient use" of the the exascale resource.

EXASCALE COMPUTING PROJECT

Generation 2: Individual FOM Formulas

The process of finalizing each application's FOM took several years

Heavily leveraged precedents and experience from LCFs

ECP was unique in focus on a science challenge problem:

- Deep questions regarding "fair" FOM definitions
- How to separate algorithmic vs. hardware improvement
- Potential paradox: more optimized codes at inception are penalized
- How to remain adaptable with new discoveries and as each domain evolves over seven year

Capability demonstration for KPP-2 applications

- how much detail?
- how to verify?

"The downside of ... benchmarks is that innovation is chiefly limited to the architecture and compiler. Better data structures, algorithms, programming languages, …cannot be used, since that would give a misleading result. The system could win because of, say, the algorithm, and not because of the hardware or the compiler. While these guidelines are understandable when the foundations of computing are relatively stable, as they were in the 1990s and the first half of this decade, they are undesirable during a programming revolution. ***For this revolution to succeed, we need to encourage innovation at all levels."***

-Hennessy and Patterson, Computer Architecture, A Quantitative Approach

Generation 3: Writing the formal KPP verification contracts

- Inclusion of SMEs in contract review
- Very difficult to strike the right balance between specificity and flexibility.
- How to formally handle cases where KPP run details deviate from contract specifications based on knowledge gained during KPP execution?
- How to quantify the “efficient use” of exascale resource for KPP-2?

➢ Avoid over-specifying details

➢ Carefully define up front role leewsay for expert judgment in reviewer role

Table 4: Comparison of minimum criteria and achieved challenge problem parameters.

Minimum criteria	Achieved
200 000 MC tally cells	213 860
6 reactions per MC cell tally	6
1×10^{10} histories per MC cycle	5.12×10^{10}
10 inactive/20 active MC eigenvalue cycles	20 inactive/20 active
200×10^{6} CFD spatial elements	1098×10^{6}
70×10^{9} CFD DOF	376×10^{9}
1000 CFD timesteps	1500

Table 7: Comparison of minimum criteria and achieved challenge problem parameters.

Minimum criteria	Achieved
100 M solver vertices	187M solver vertices
Minimum time-step 1×10^{-1} ion sound wave period	0.47×10^{-1} for most unstable KBM modes in ITER pedestal
1×10^{12} particles in XGC	1.6×10^{12} in XGC in edge region
10 ion time steps	10 ion timesteps

ECP Software Technology Update

Michael A. Heroux
Director of Software Technology, Exascale Computing Project
Senior Scientist, Sandia National Laboratories

ASCAC Meeting, September 28, 2023

ECP Impact – Portable Libraries and Tools for Accelerators

ECP Software Technology works on products that apps need now and in the future

Key themes:
- Focus: GPU node architectures and advanced memory & storage technologies
- Create: New high-concurrency, latency tolerant algorithms
- Develop: New portable (Nvidia, Intel, AMD GPUs) software product
- Enable: Access and use via standard APIs

Software categories:
- **Next generation established products:** Widely used HPC products (e.g., MPICH, OpenMPI, PETSc)
- **Robust ~~emerging~~ products:** Address key new requirements (e.g., Kokkos, RAJA, Spack)
- **New products:** Enable exploration of emerging HPC requirements (e.g., zfp, Variorum)

2023 R&D 100 WINNER

Legacy: A stack that enables performance portable application development on leadership platforms

Example Products	Engagement
MPI – Backbone of HPC apps	Explore/develop MPICH and OpenMPI new features & standards
OpenMP/OpenACC –On-node parallelism	Explore/develop new features and standards
Performance Portability Libraries	Lightweight APIs for compile-time polymorphisms
LLVM/Vendor compilers	Injecting HPC features, testing/feedback to vendors
Perf Tools - PAPI, TAU, HPCToolkit	Explore/develop new features
Math Libraries: BLAS, sparse solvers, etc.	Scalable algorithms and software, critical enabling technologies
IO: HDF5, MPI-IO, ADIOS	Standard and next-gen IO, leveraging non-volatile storage
Viz/Data Analysis	ParaView-related product development, node concurrency

Example: Addressing growing gap of ops vs bw vs memory
ZFP compressed multidimensional array primitive

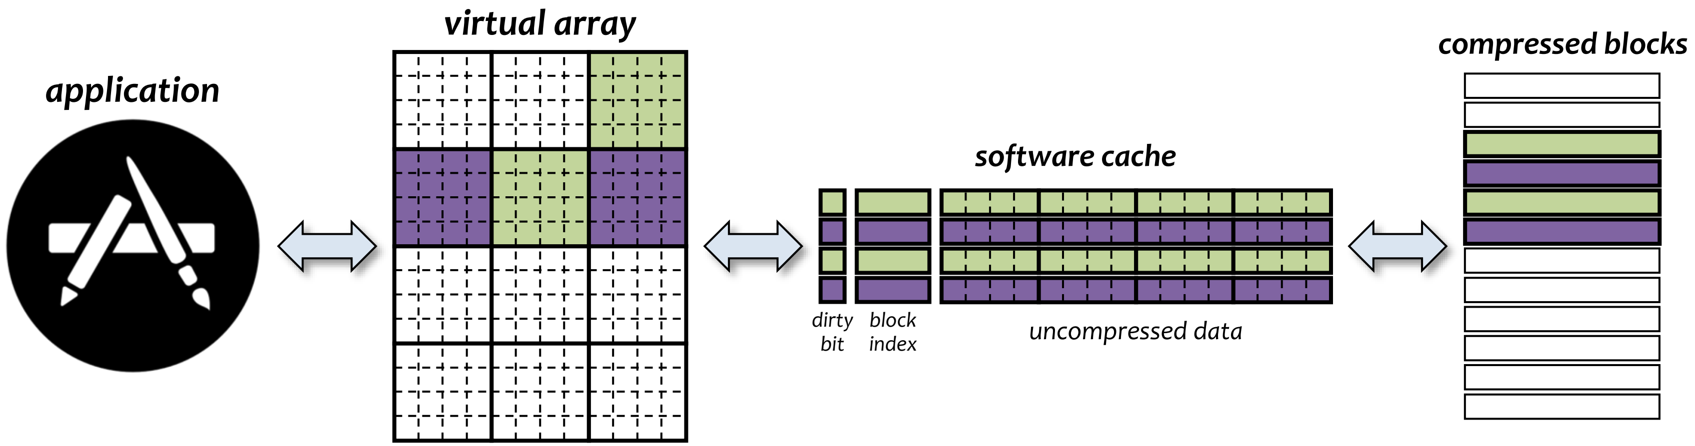

- Fixed-length compressed blocks enable fine-grained read & write **random access**
  - C++ compressed-array classes hide complexity of compression & caching from user
  - User specifies per-array storage footprint in bits/value
- Absolute and relative **error tolerances** supported for offline storage, sequential access
- Fast, hardware friendly, and parallelizable: **150 GB/s throughput** on NVIDIA Volta
- **HPC tool support**: ADIOS, CEED Exascale Discretizations, Conduit, HDF, intel, VTK-m

Inline compression boosts solution accuracy by 40x over IEEE in CEED code

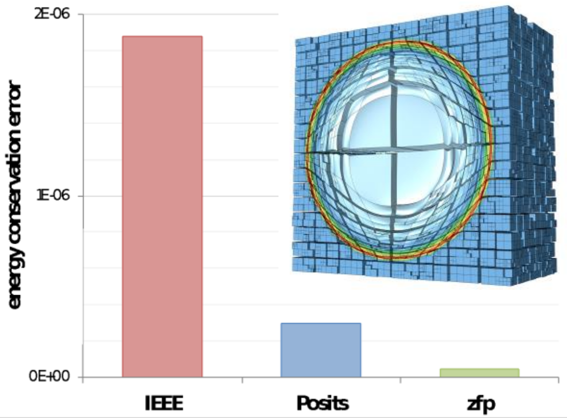

Current HPC Systems under-utilize power significantly, and this trend is expected to worsen at scale with GPU-based systems

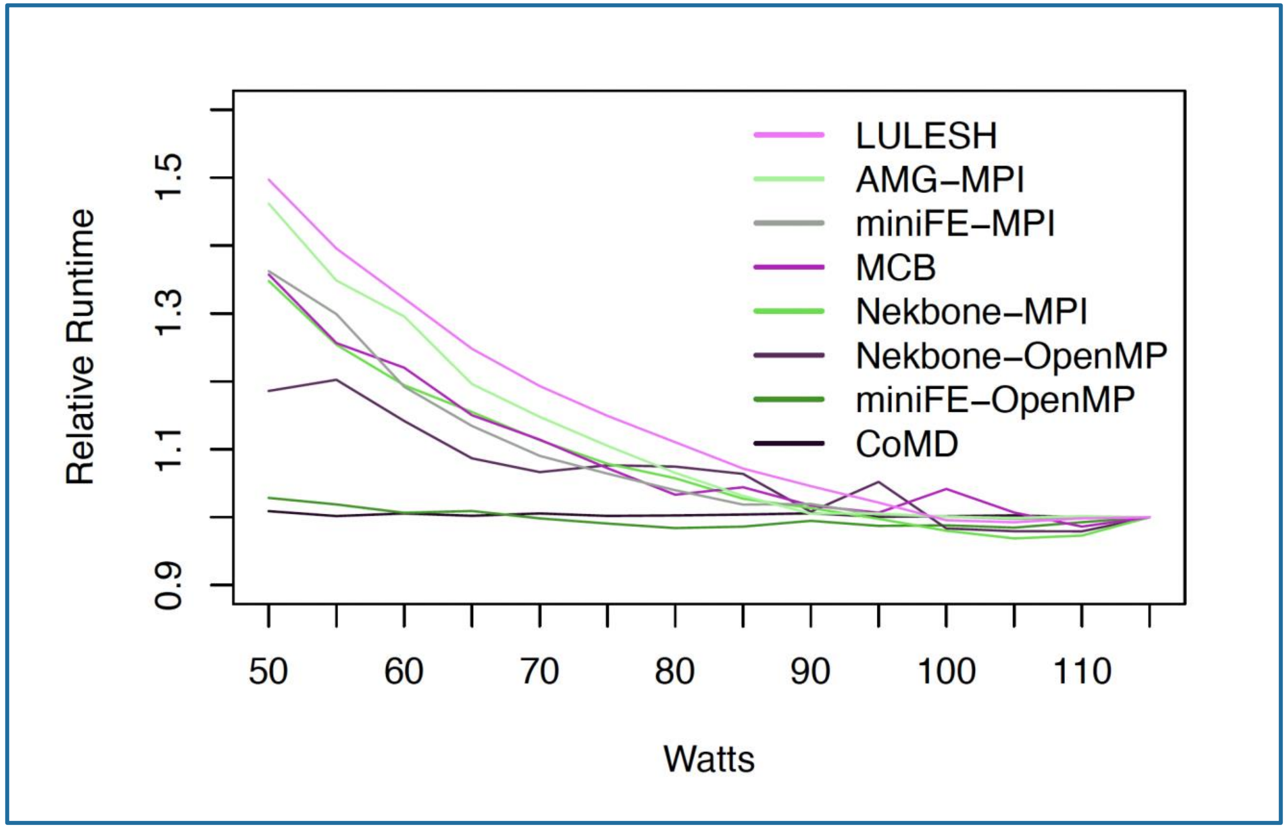

Source: Daniel A Ellsworth, Allen D Malony, Barry Rountree, and Martin Schulz. 2015. Dynamic power sharing for higher job throughput. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. ACM, 80

- Power-performance curves for HPC applications vary significantly due to their CPU, GPU and memory characteristics. Many applications do not utilize the amount of power that is procured to them. Performance of such applications is not impacted by giving them less power, as can be seen on the graph, where some applications continue to have the same runtime despite power being reduced by half (115 W to 51 W).
- Allocating power to where it is needed most results in optimal performance, energy savings, and higher scientific throughput, requiring a power stack which allows for intelligent power scheduling.

Low-level dials and interfaces for power and performance management vary significantly across vendors

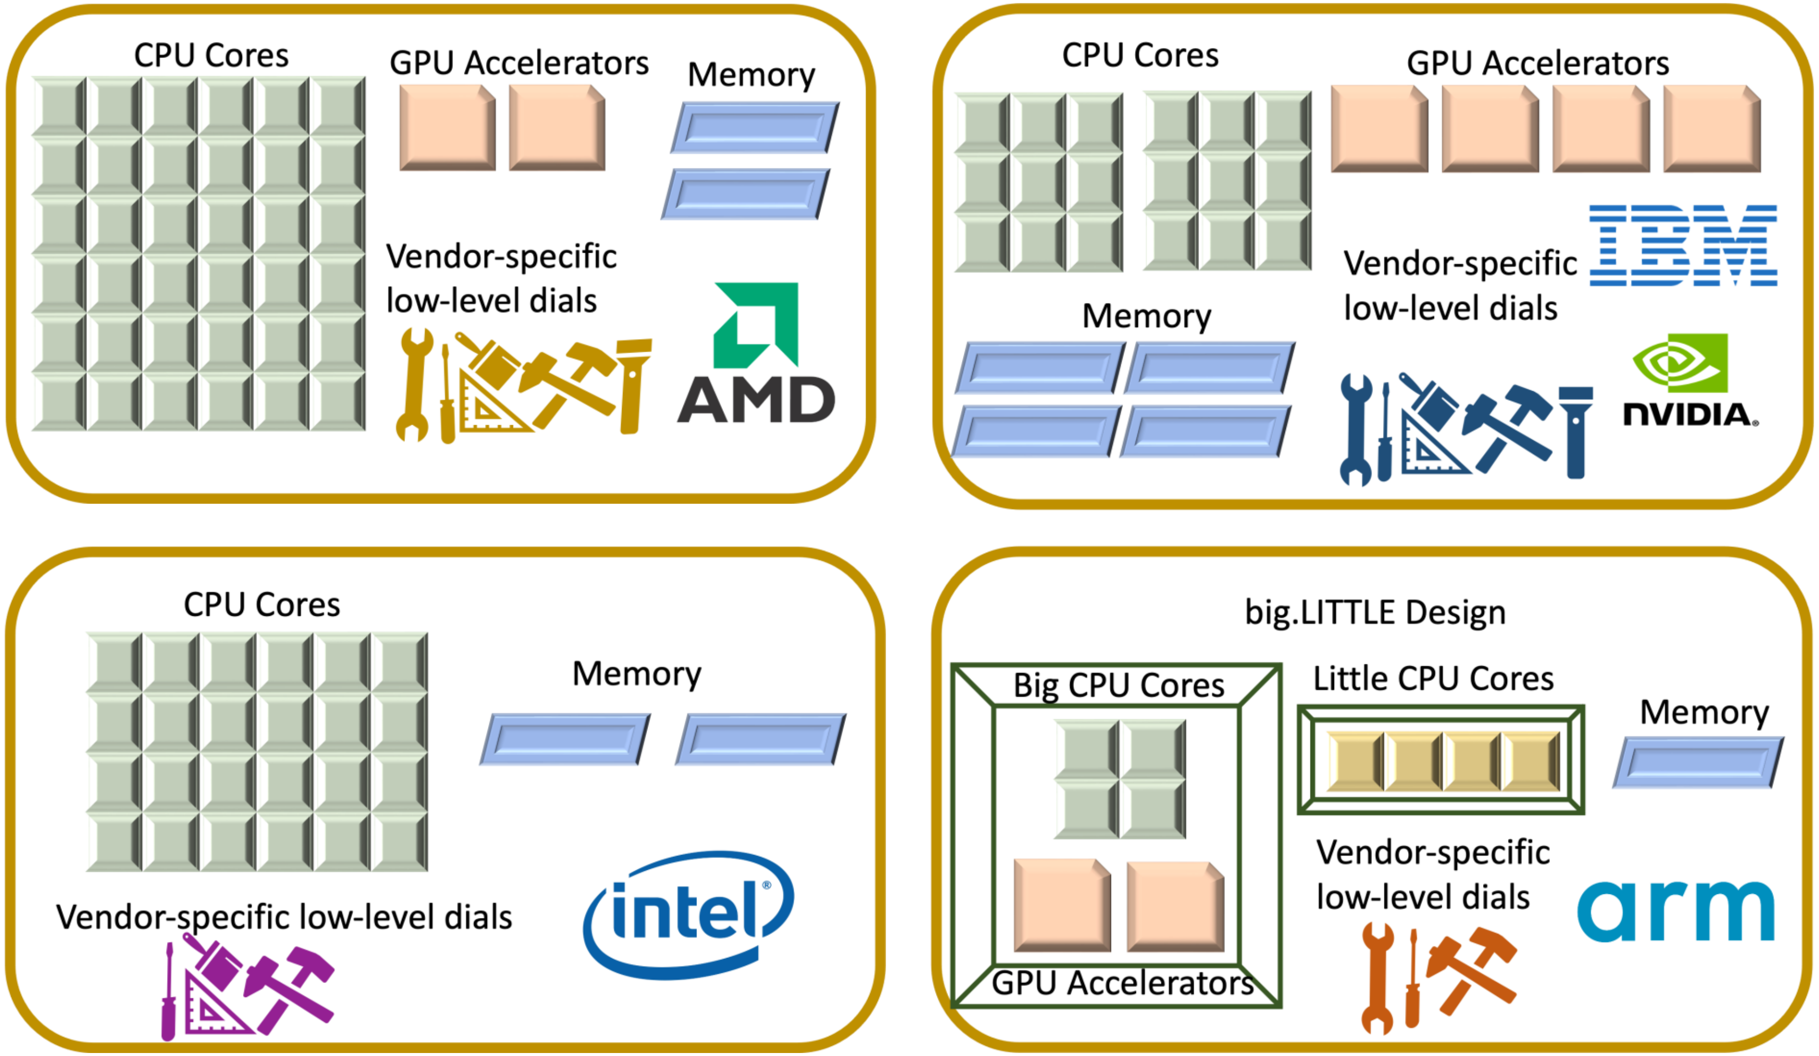

Accessing vendor-specific dials is unwieldy for domain scientists as well as experts, due to lack of uniform APIs, privileged access, and limited documentation

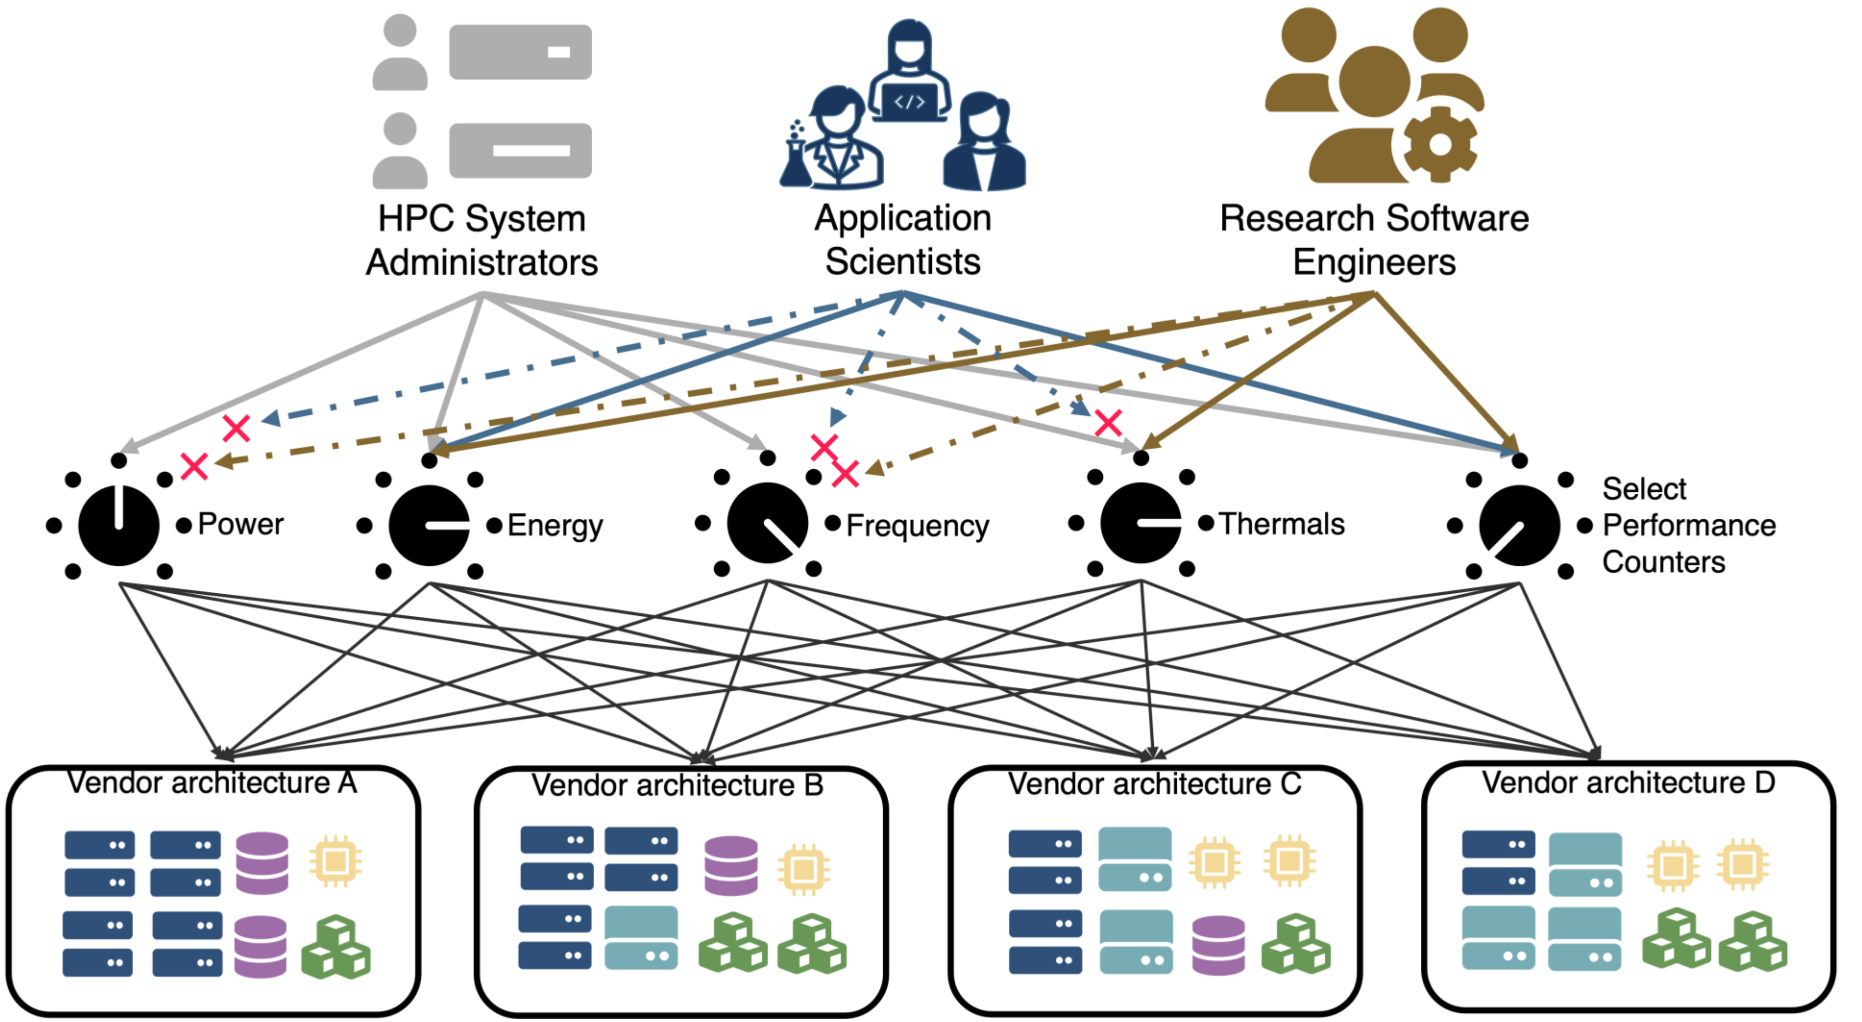

Variorum provides safe, user-space, vendor neutral access all users: administrators, application scientists and RSEs

PI: Tapasya Patki

Team: Kathleen Shoga, Stephanie Brink, Eric Green, Aniruddha Marathe, Barry Rountree

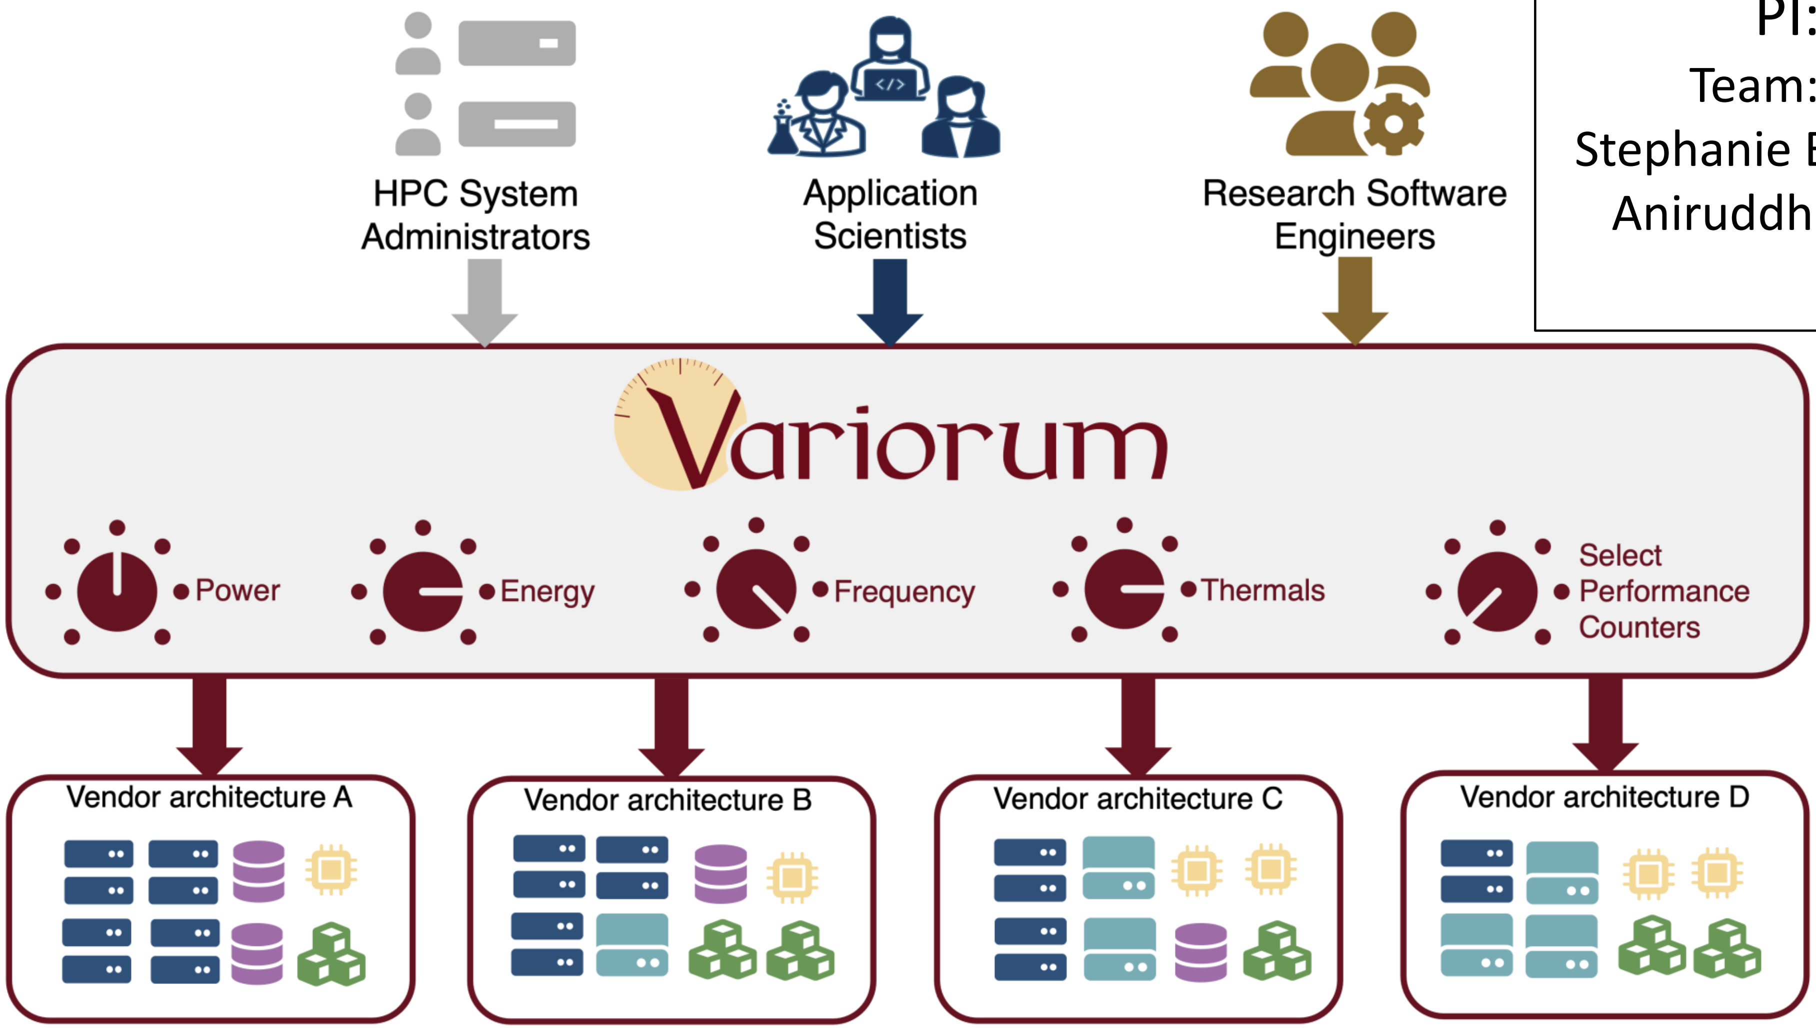

Variorum: Vendor-neutral user space library for power management

- Power management capabilities (and their interfaces, domains, latency, capabilities) widely differ from one vendor to the next, needing common interfaces
- Variorum: Platform-agnostic vendor-neutral, simple front-facing APIs
  - Evolved from *libmsr*, and designed to target several platforms and architectures
  - Abstract away tedious and chaotic details of low-level knobs
  - Implemented in C, with function pointers to specific target architecture
  - Integration with higher-level power management software through JSON
- Integrated with Flux, GEOPM, LDMS, Kokkos, Caliper and PowerAPI to enable a PowerStack
- Supported on all upcoming exascale systems (Aurora, Frontier, El Capitan) and several other supercomputers: architecture support includes CPU support for ARM, AMD, Intel, IBM; and GPU support for NVIDIA, AMD and Intel.

And Many More…

- ECP generated a
  - Collection of portable GPU-capable libraries and tools for AMD, Intel, and NVIDIA devices
  - Designed for future adaptation to next-generation highly-concurrent node architectures
  - Foundation for others who will make the transition from CPU to GPU and beyond
- ECP has contributed to the exascale readiness of 70 libraries and tools

The Extreme-scale Scientific Software Stack (E4S) Update

E4S: Extreme-scale Scientific Software Stack

- E4S is a community effort to provide open-source software packages for developing, deploying and running scientific applications on HPC platforms.
- E4S has built a comprehensive, coherent software stack that enables application developers to productively develop highly parallel applications that effectively target diverse exascale architectures.
- E4S provides a curated, Spack based software distribution of 100+ HPC, EDA (e.g., Xyce), and AI/ML packages (e.g., TensorFlow, PyTorch).
- With E4S Spack binary build caches, E4S supports both bare-metal and containerized deployment for GPU based platforms.
  - X86_64, ppc64le (IBM Power 9), aarch64 (ARM64) with support for GPUs from NVIDIA, AMD, and Intel
  - HPC and AI/ML packages are optimized for GPUs and CPUs.
- Container images on DockerHub and E4S website of pre-built binaries of ECP ST products.
- Base images and full featured containers (with GPU support) and DOE LLVM containers.
- Commercial support for E4S through ParaTools, Inc. for installation, maintaining an issue tracker, and ECP AD engagement.
  - https://dashboard.e4s.io https://e4s.io/talks/E4S_Support_July23.pdf
- E4S for commercial cloud platforms: AWS image supports MPI implementations and containers with remote desktop (DCV).
  - Intel MPI, NVHPC, MVAPICH2, MPICH, MPC, OpenMPI
- e4s-cl container launch tool allows binary distribution of applications by substituting MPI in containerized app with system MPI.
- Quarterly releases: E4S 23.08 released on Aug 31, 2023: https://e4s.io/talks/E4S_23.08.pdf

E4S 23.08: What's New?

- E4S includes support for Intel oneAPI 2023.2.0 software (BaseKit and HPCToolkit) in containers on x86_64 with support for HPC packages built with Intel compilers
- New platform: ARM64 (aarch64) with H100 (90) with CUDA 12.2 and NVHPC 23.7
- E4S includes support for CUDA architectures
  - 70 (V100), 80 (A100), and 90 (H100) under x86_64
  - 70 under ppc64, and
  - 75, 80, and 90 (H100) under aarch64
- E4S includes supports ROCm for gfx908 (MI100) and gfx90a (MI200) architectures under x86_64
- E4S includes support for DOE LLVM under x86_64, ppc64le, and aarch64
- E4S includes new applications: Ascent, Bricks, Blaspp, Camp, HDF5-vol-cache, HDF5-vol-log, Lapackpp, MGARD, pruners-ninja, Xyce (with pymi), LBANN, Quantum Espresso, LAMMPS, WARPX, Dealii, and OpenFOAM
- E4S includes support for AI/ML frameworks such as TensorFlow 2.12.0 and PyTorch 2.0.0 support for A100 as well as H100 GPUs is integrated in E4S 23.08
- E4S supports updates to 115 HPC packages on x86_64, aarch64, and ppc64le, 100K+ binaries in E4S Spack Build Cache
- New E4S tools: e4s-alc (à la carte) customizes container images, e4s-cl (container launch) replaces MPI at runtime!
- Detailed documentation for installing E4S on bare-metal and using containers

E4S: First-of-a-kind collaboration between facilities, libraries, tools

- Efforts to provide DOE-sponsored libraries and tools through the E4S stack is game changing
- Success achieved by deep and broad collaboration between libraries and tools teams and facilities
- As ECP ends, we have proven collaborations, processes, and tools that we will continue
- E4S would not be possible without dedication of facilities staff to work with us and help us improve
- Thank you

Leveraging the Potential of ECP Investments

100X

100X Demonstrated: ECP-sponsored application FOMs

Project/PI	**EXAALT: Molecular Dynamics Danny Perez**
Challenge Problem	Damaged surface of Tungsten in conditions relevant to plasma facing materials in fusion reactors • 100,000 atoms • T=1200K
FOM Speedup	**398.5**
Nodes Used	**7000**
ST/CD Tools	Used in KPP Demo: **Kokkos, CoPa**

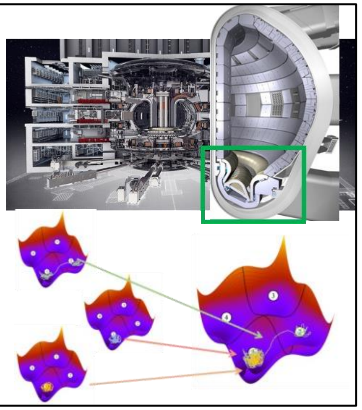

Project/PI	**ExaSMR: Small Modular Reactors Steve Hamilton**
Challenge Problem	NuScale-style Small Module Reactor (SMR) with depleted fuel and natural circulation • 213,860 Monte Carlo tally cells/6 reactions • 5.12×10^{12} particle histories/cycle, 40 cycles • 1098×10^{6} CFD spatial elements • 376×10^{9} CFD degrees of freedom • 1500 CFD timesteps
FOM Speedup	**70**
Nodes Used	**6400**
ST/CD Tools	Used in KPP Demo: **CEED** Additional: Trilinos

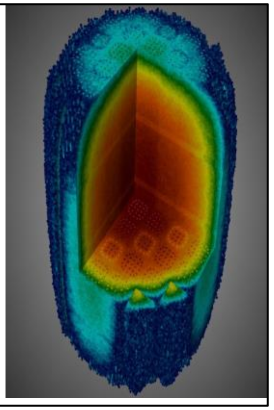

Project/PI	**ExaSky: Cosmology Salman Habib**
Challenge Problem	Two large cosmology simulations • gravity-only • hydrodynamics
FOM Speedup	**271.65**
Nodes Used	**8192**
ST/CD Tools	Used in KPP demo: **none** Additional: CoPa, VTK-m, CINEMA, HDF5.0

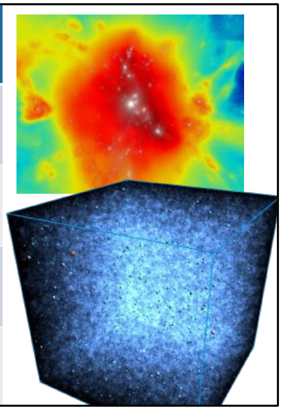

Project/PI	**WarpX: Plasma Wakefield Accelerators Jean-Luc Vay**
Challenge Problem	Wakefield plasma accelerator with a 1PW laser drive • 6.9×10^{12} grid cells • 14×10^{12} macroparticles • 1000 timesteps/1 stage
FOM Speedup	**500**
Nodes Used	**8576**
ST/CD Tools	Used in KPP Demo: **AMReX**, **libEnsemble** Additional: ADIOS, HDF5, VTK-m, ALPINE

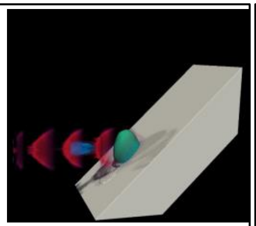

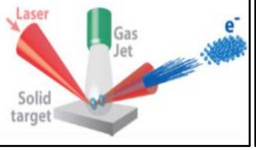

Project/PI	**WDMApp: Fusion Tokamaks Amitava Bhatacharjee**
Challenge Problem	Gyrokinetic simulation of the full ITER plasma to predict the height and width of the edge pedestal
FOM Speedup	**150**
Nodes Used	**6156**
ST/CD Tools	Used in KPP Demo: **CODAR**, **CoPA**, **PETSc**, **ADIOS** Additional: VTK-m

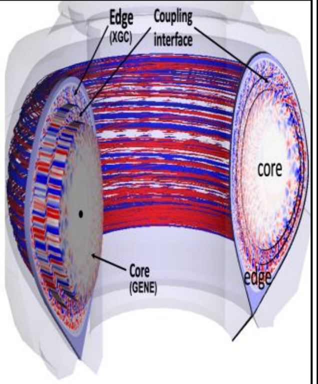

Project/PI	**EQSIM: Earthquake Modeling and Risk Dave McCallen**
Challenge Problem	Impacts of Mag 7 rupture on the Hayward Fault on the bay area.
FOM Speedup	**3467**
Nodes Used	**5088**
ST/CD Tools	Used in KPP Demo: **RAJA, HDF5**

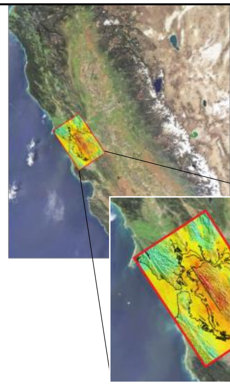

ECP investments enabled a 100X improvement in capabilities

- **7 years** building an **accelerated**, **cloud-ready** software ecosystem
- Positioned to utilize **accelerators from multiple vendors** that others cannot
- **Emphasized software quality**: testing, documentation, design, and more
- Prioritized **community engagement**: Webinars, BOFs, tutorials, and more
- DOE **portability layers** are the credible way to
  - Build codes that are sustainable **across multiple GPUs** and
  - **Avoid vendor lock-in**
  - **Avoid growing divergence** and hand tuning in your code base
- ECP software can **lower costs** and **increase performance** for **accelerated** platforms
- Outside of AI, industry has not caught up
  - DOE enables an entirely different class of applications and capabilities to use accelerated nodes
  - In addition to AI
- **ECP legacy: A path and software foundation for others to leverage**

Opportunities to realize 100X leveraging ECP investments

- Port to full use of GPUs:
  - Hotspot use of GPUs is a start but not sufficient.
  - Scalability very limited and capped for future GPU devices
- Utilize Spack ecosystem:
  - Opens ready access to hundreds of curated libraries and tools
  - Makes your code easy to consume if you publish Spack recipes for your code
  - Utilize Spack build caches (10X speedup in rebuild times)
- Utilize E4S
  - Curated libraries, tools, documentation, build caches, and more
  - Commercial support via ParaTools
  - Pre-built containers, binaries,
  - Cloud instances for AWS, Google – Permit elastic expansion, neutral collaboration for cross-agency work
- Leverage ECP team experience:
  - Engage with DOE HPC staff

100X Recipe

- Ingredients
  - A compelling science impact story
  - $$ - $$$
  - Staff
  - Computing resources, training
  - The deliverables and experience from DOE/ECP
  - Delivered via post-ECP organizations like PESO
  - And more…
- Steps
  - Translate science story to strategy and plan – leverage experience from ECP, others
  - ID node-level parallelization strategy – CUDA, HIP, DPC++, Kokkos, RAJA, OpenMP, others
  - Survey existing libraries and tools – Vendors, E4S, others
  - Explore available platforms – DOE Facilities, cloud, others
  - Leverage existing software ecosystem – containers, Spack, others
  - Leverage software communities – Product communities, communities of practice, others
  - Construct new codes within the broader ecosystem
  - Produce new science results

More than one way to leverage 100X

- 100X can be realized as exciting new science capabilities at the high end
  - Fundamental new science on leadership platform
  - New opportunities on affordable machines that fit in current data centers
- **But can also reduce costs**
- Migration to accelerated platforms can be used to
  - Migrate a problem from an HPC cluster to a deskside or laptop systems
  - Lower your AWS monthly charges – E4S is available for container/cloud
  - Keep energy costs in check while still growing computing capabilities
- **Biggest ECP impact will be accelerating GPU transition – at all levels**
- **Transitioning software stacks to GPUs is essential**
  - **CPU-based HPC system realize only modest energy efficiency improvements**
  - **Migrating to GPUs is key to improving HPC environmental impact**

ECP Software Technology Takeaways

- KPP-3 Threshold Achieved September 15, 2023
- ECP legacy:
  - A collection of portable GPU-enabled libraries and tools ready for today's platforms
    ready to be adapted to future systems
  - A de-risked path for other communities:
    - Libraries and tools for their use: Spack, E4S
    - Lessons-learned to be leveraged
  - A foundation for a 100X improvement across many scientific domains
  - A key to addressing HPC environmental impact concerns
- Next steps:
  - Continued engagement with DOE facilities on delivery of software via E4S
  - Cultivation of ECP impacts in collaboration with other agencies and industry
  - We are just getting started

Conclusion and Summary

Thank you

https://www.exascaleproject.org

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

EXASCALE COMPUTING PROJECT

ECP Director: Lori Diachin
ECP Deputy Director: Ashley Barker

Thank you to all collaborators in the ECP and broader computational science communities. **The work discussed in this presentation represents creative contributions of many people who are passionately working toward next-generation computational science.**

Looking Ahead

- **Frontier**
  - All 62 ECP teams are now on Frontier, with priority for those not yet hitting threshold targets
  - Run science challenge campaigns, harden code, strive for stretch KPPs
- **Aurora**
  - All 62 ECP teams on Aurora TDS (Sunspot) and as many as possible on Aurora with a focus on any blocking issues
  - Full access is expected in Oct 2023
  - Showcase performance portability
  - Strive for large-scale demonstrations on at least half the system for teams that may meet their threshold KPP or are otherwise likely to deliver scientific impact (e.g., KPP objective).

- **ECP KPPs**
  - Complete review of KPPs; document for project completion
  - KPP-1 Target: Expect 9-10 (of 11) applications will complete their base challenge problem with performance exceeding 50X
  - KPP-2 Target: Expect 7-8 SC applications (of 10) plus 3 (of 4) NNSA applications to meet their base challenge problem
  - KPP-3 Target: Not unrealistic to see an integration score of 55+ points (of 68)
- **Opportunities**
  - Scope post KPP threshold for each AD and ST team will target KPP stretch goals, performance portability, post ECP uptake, and buying down technical debt

At closeout, ensure expeditious closeout of all MPOs and subcontracts, successful CD-4 review, return of all uncommitted ASCR funds back to ASCR for appropriate rescoping

Questions?