Benchmarks, accelerated datacenters, parallel language standards, and regular software testing at scale

Jack Wells

Science Program Manager, NVIDIA

Challenge: Scientifically Faithful & Sustainable Benchmarks

Benchmarks should be developed that are faithful to modern science applications and workflows.

Opportunity for more open development of benchmarks

- Track science and engineering capability progress
- Realistic challenges for technology and system R&D – Influence industry
- High quality marketing and communications to stakeholders and society

Sustaining the vitality and relevance for benchmarking over time is difficult

- Increase open, community-based collaboration
- Having the application developers remain connected with benchmark curation is an advantage
- Ongoing discussion of science benchmarks outside of “business motion” should improve quality

Application Benchmarks in NVIDIA

Process

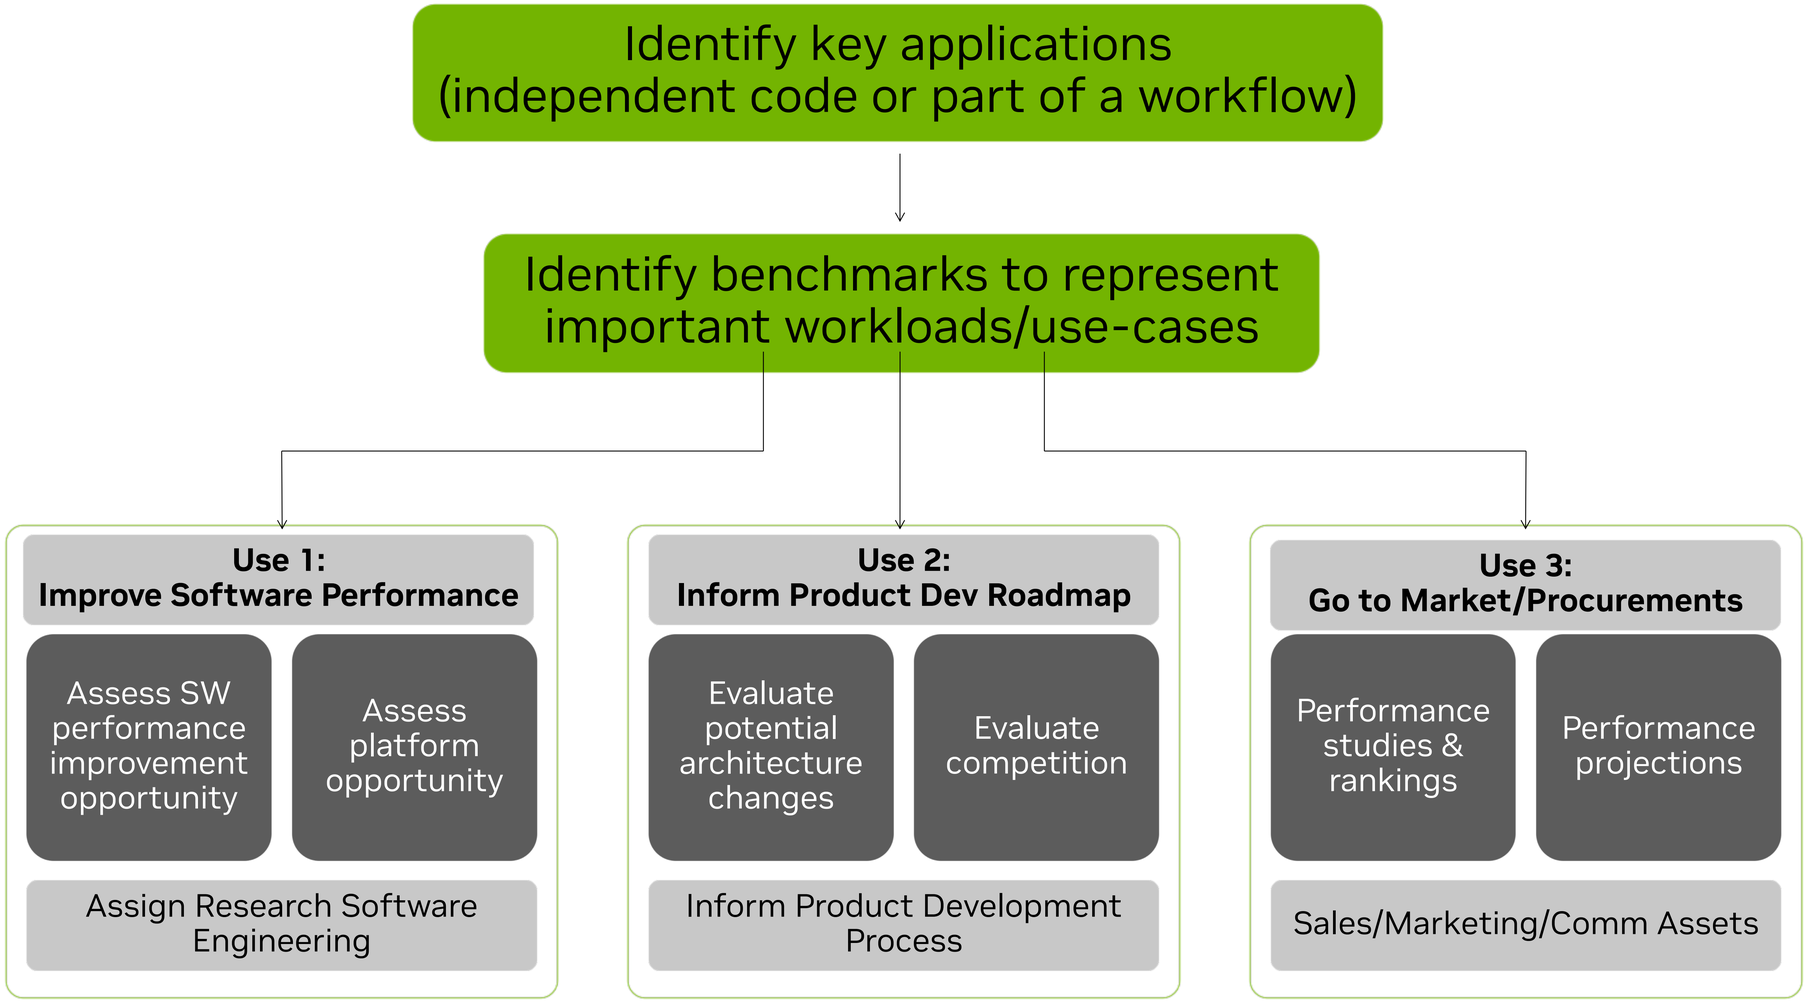

The ultimate impact depends critically on the ability to identify the important research domains on which to work

Influence through Communicating Benchmark Lists

Well curated and communicated lists of benchmarks have proven to be very influential.

- TOP500™ is a notable example – See Jack Dongarra's A.M. Turing Award Lecture

- Pay attention to the principles of effective marketing and communication

Benchmarking efforts that aspire to similar influence must not overlook the important role of marketing and communications for their work.

Various Governance Models Used for Sustainability

Two Examples

TOP500™ is the property of a small, private company

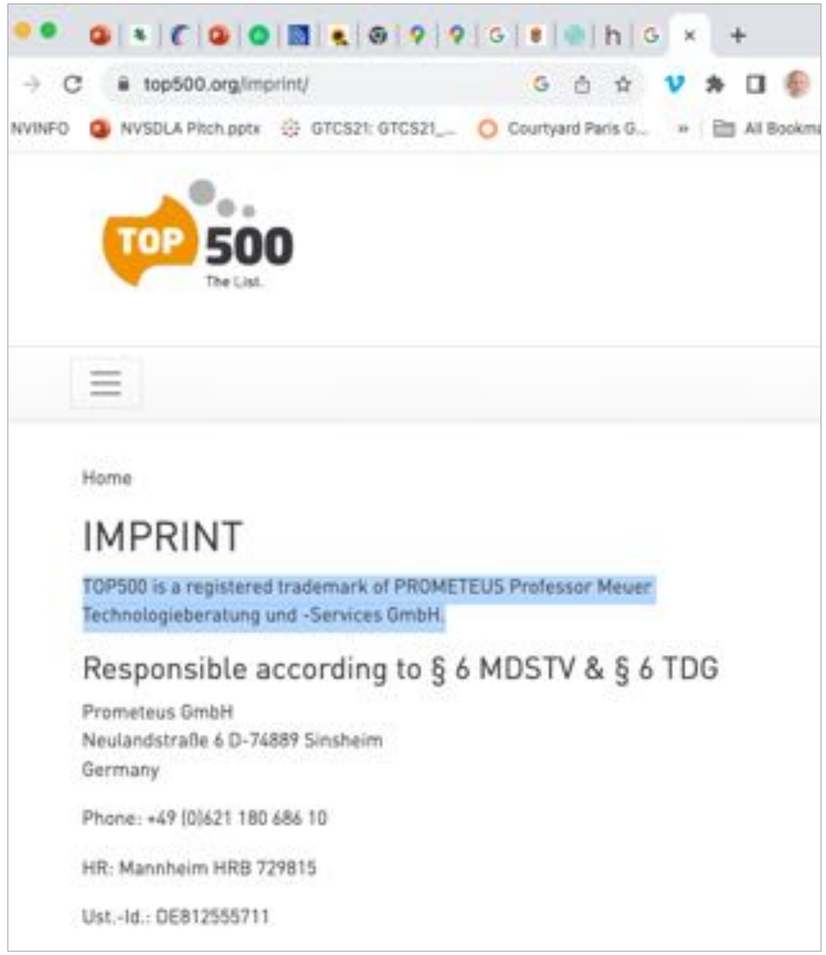

MLPerf governance is based on open standards.

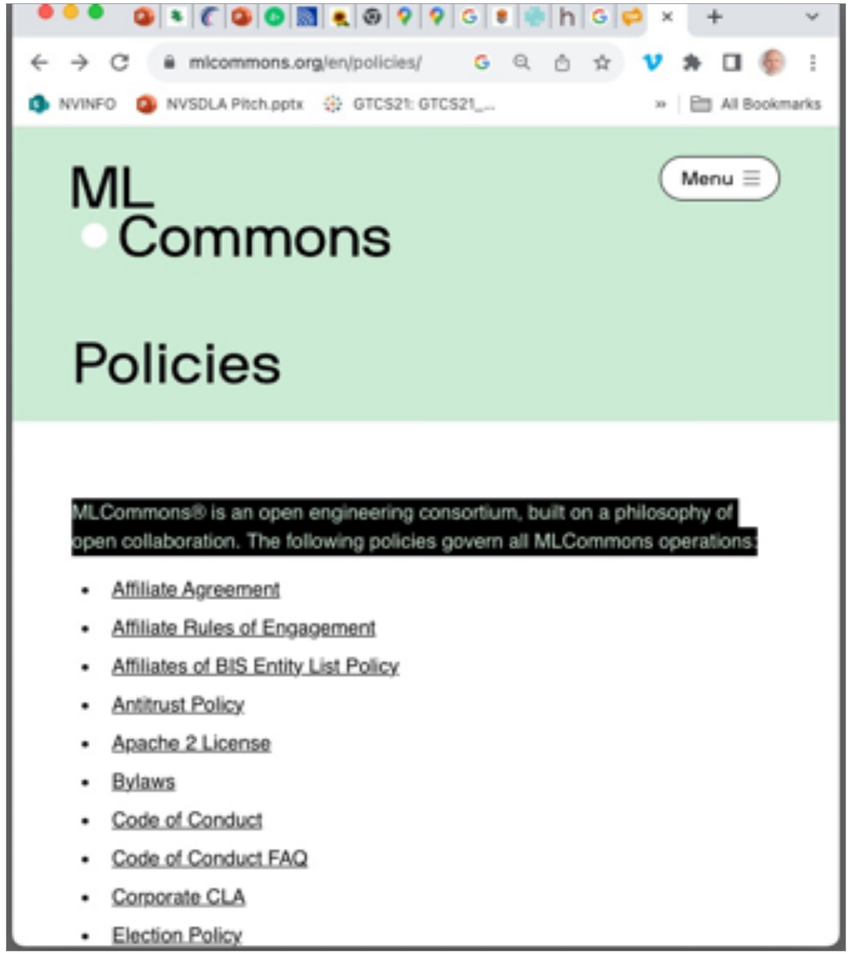

Energy efficiency has grown to a primary consideration

Need to trade off performance against energy consumption

AI and HPC public perception problem

HPC datacenters consuming lots of energy seen as negative

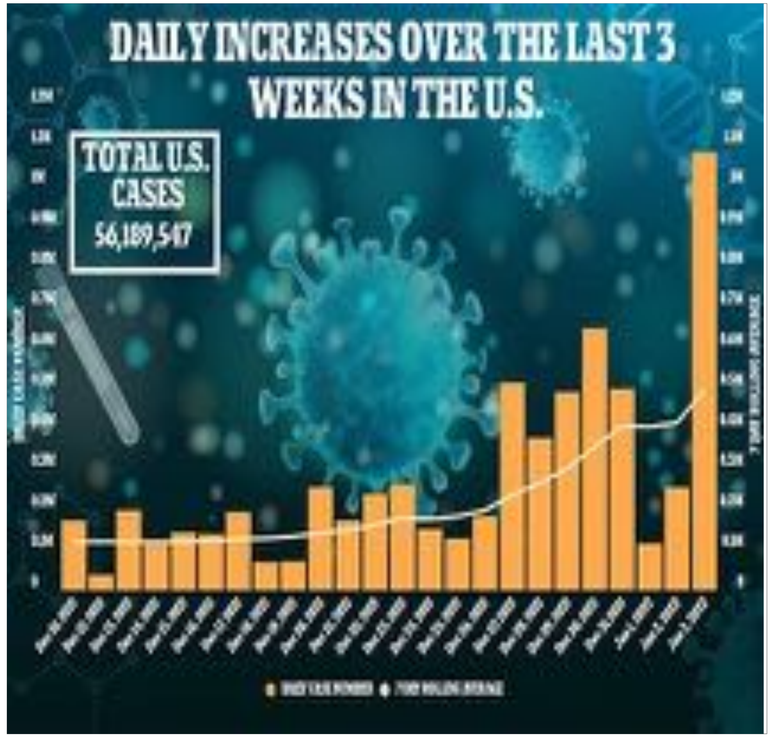

Demand for fast solutions only grows

Public understanding of science & methods poses challenges

Water and other resources important too

Nascent focus on datacenter water usage for cooling

Accelerated Computing Drives Energy Efficient Data Centers

Simulation + AI Supercharge Science

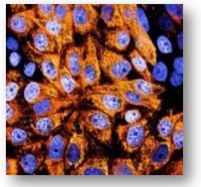

Light Sheet Microscopy

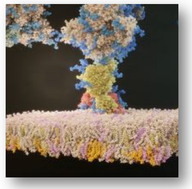

Molecular Dynamics

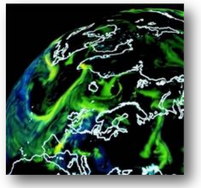

Climate Simulation

Fluid Flow Resolution

Energy Used in a Typical HPC task

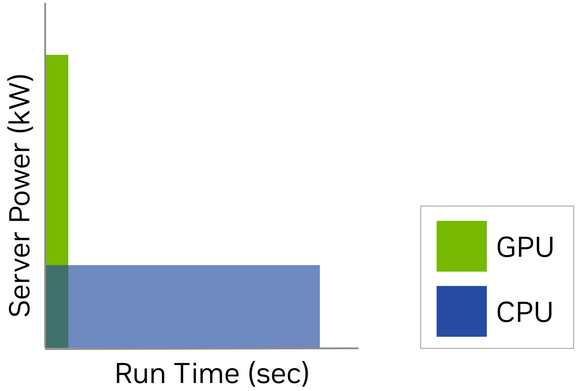

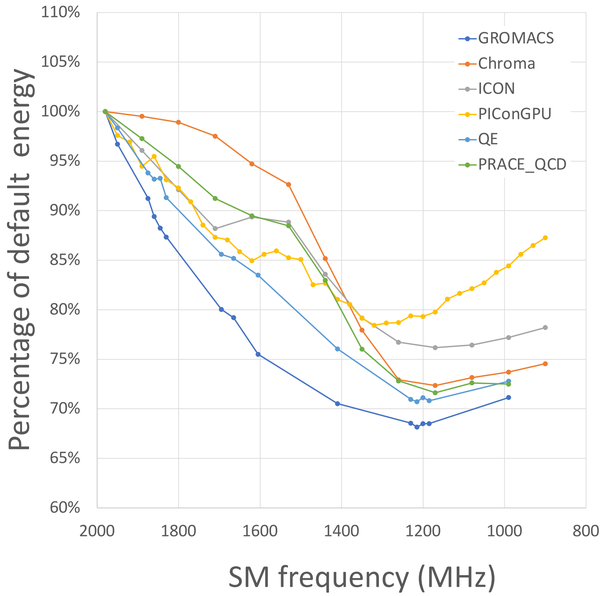

Demand for Computational Science Growing

Compute Intensive Functions Drive Energy Consumption

GPUs Use Less Energy for HPC and AI

Rapid workload completion uses less energy overall

More Opportunities for Energy Efficiency

Increasing HPC and AI efficiency doesn't stop with acceleration

Energy-to-solution and Time-to-Solution

MaxQ: Running at the Power Efficiency Sweet Spot by Dynamic Power Capping

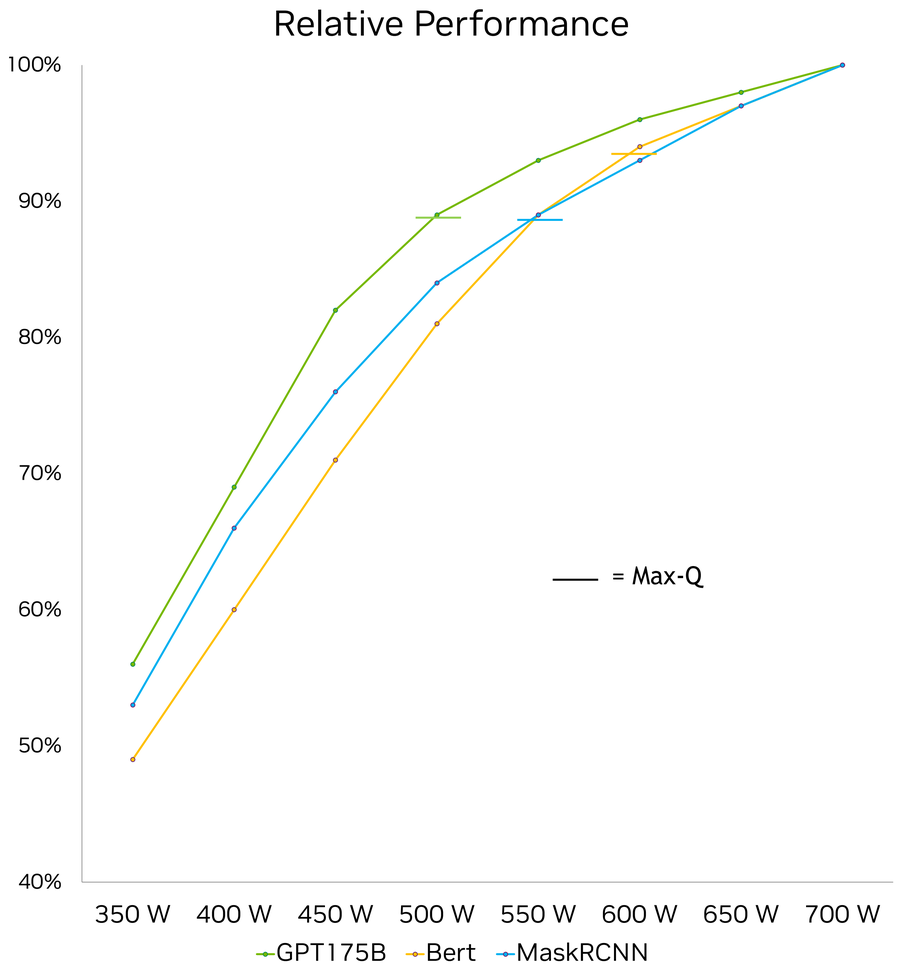

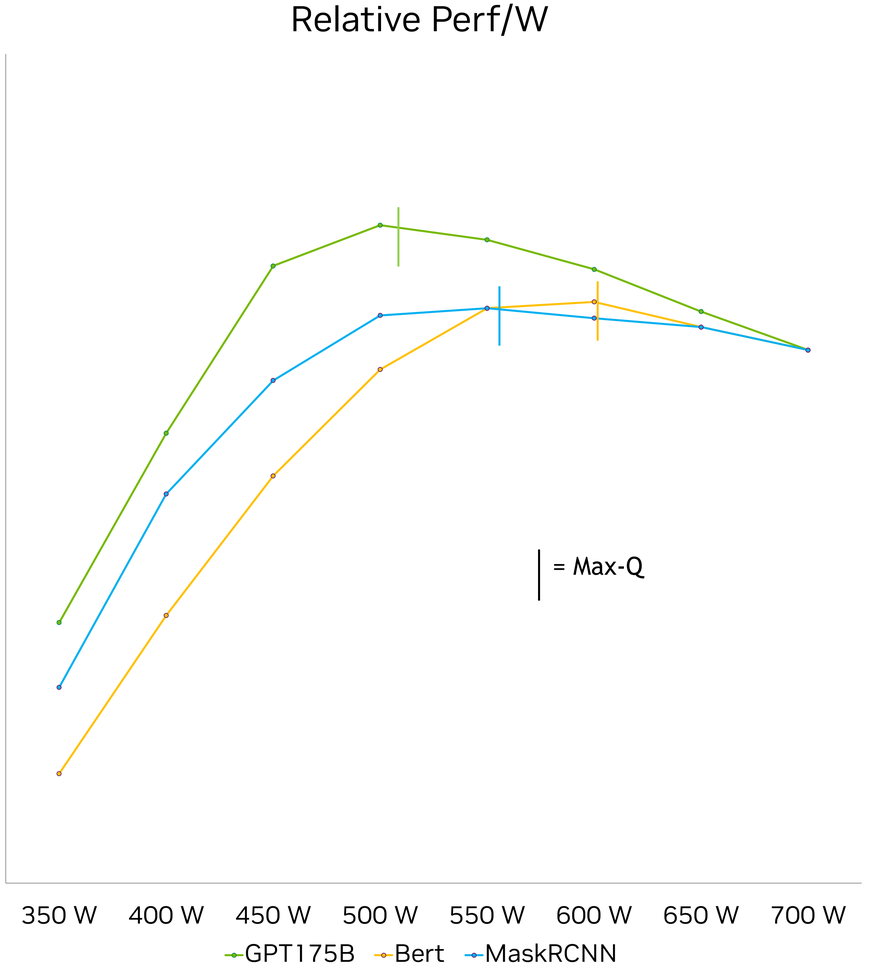

Systems: x86-hosted H100 HGX w/NVLink and IB interconnect. Workloads are training indicated models at scale: GPT175B 8192 GPUs, Bert 8 GPUs, MaskRCNN 272 GPUs

Energy-to-solution and Time-to-Solution

Need to trade off time-to-solution against resource (energy) consumption

Where on the scaling curve should we run?

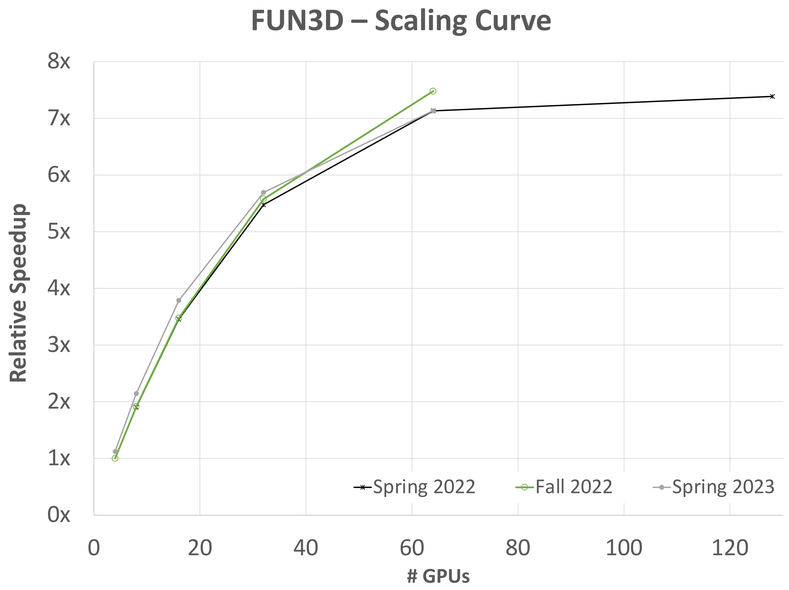

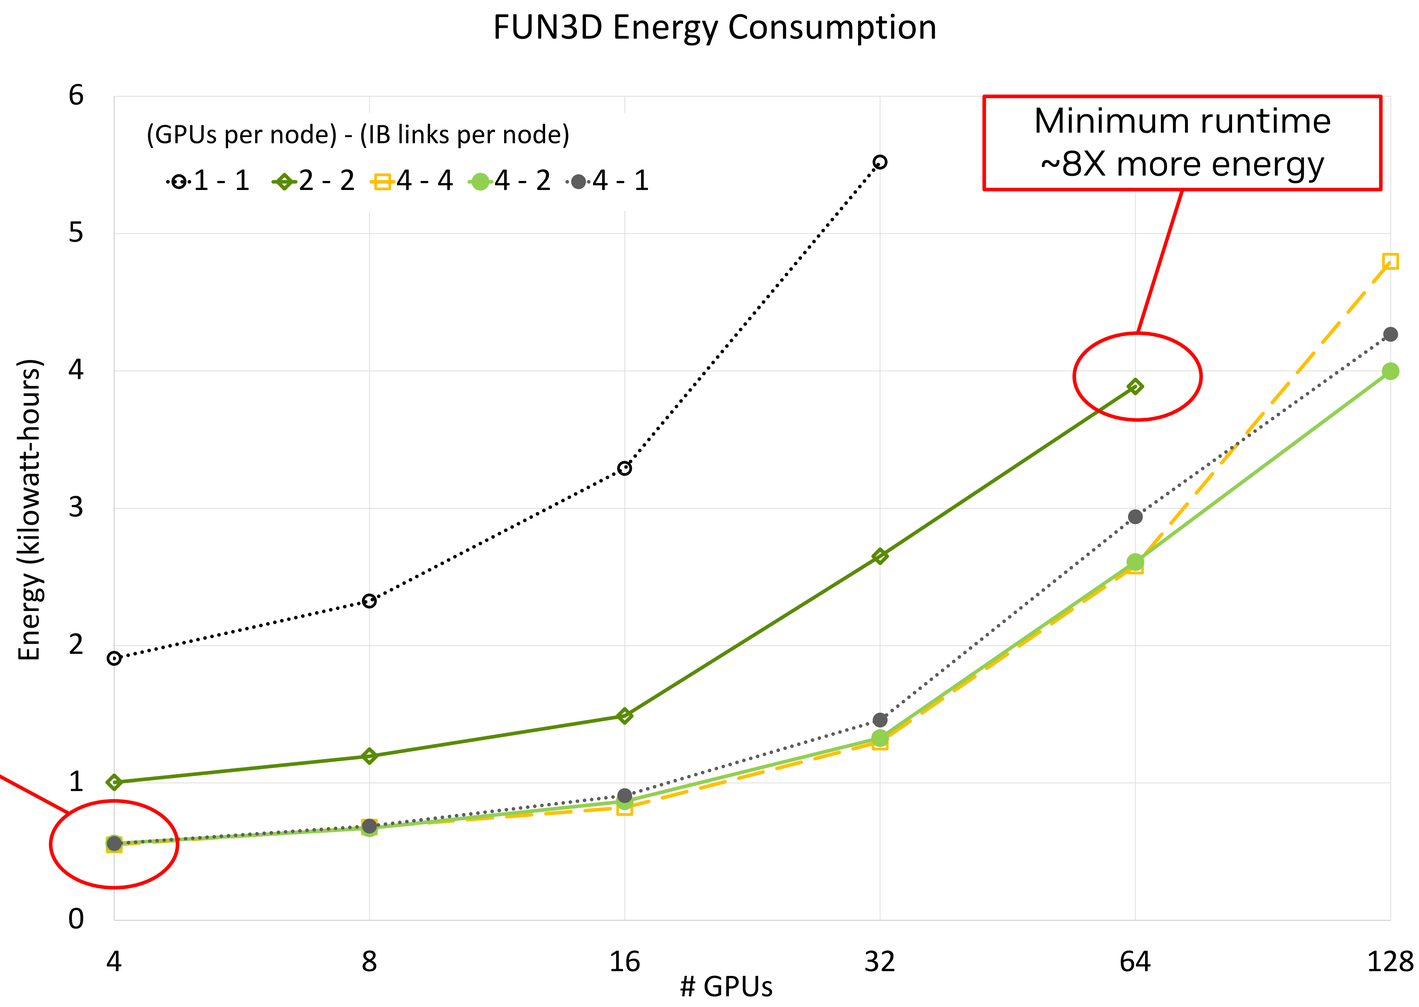

Energy Corrected Speedup

Best compromise between speed and energy consumption?

$$SPEEDUP_{energy} =: \frac{\left[\frac{runtime_1}{runtime_n}\right]}{\left[\frac{energy_n}{energy_1}\right]} = \frac{[runtime_1 * energy_1]}{[runtime_n * energy_n]}$$

$$\boxed{Minimize\ [runtime_n * energy_n]}$$

Offered as potential answers to symposium organizers questions:

- What metrics are the most valued?
- How do you characterize its performance?

Energy-to-solution and Time-to-Solution

Optimize Energy*Time

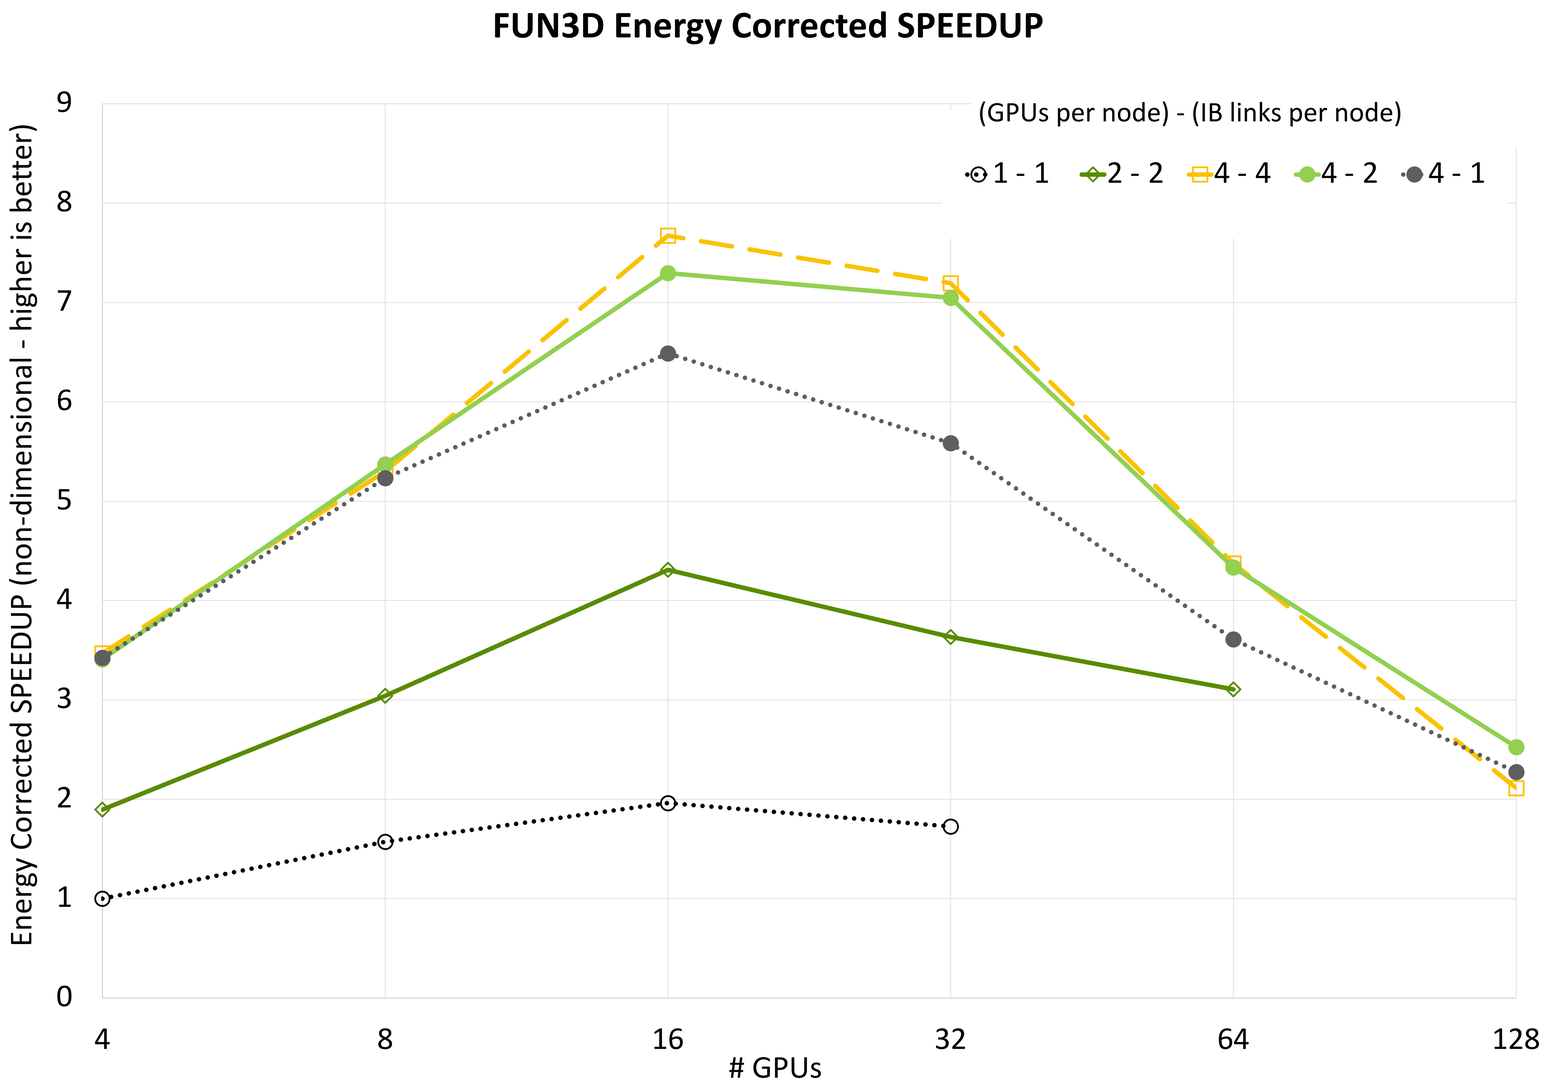

FUN3D: Version: 14.0-d03712b , Dataset: WB.C-30M

SC23 BoF: The Future of Benchmarks in Supercomputing (FoBS)

Wednesday, 15 November 2023,
5:15pm - 6:45pm MST

- The why?
  - Track science and engineering capability progress
  - Realistic challenges for technology and system design
  - High quality marketing and communications assets
- What?
  - Encourage open collaboration on HPC benchmarks
  - Think beyond LINPACK for science benchmarks
  - Energy-to-solution & Time-to-Solution
  - Benchmarking workflows.
- How?
  - Encourage community collaboration
  - Inform and initiate FoBS workshops in 2024
  - Share best practices in communicating benchmark results

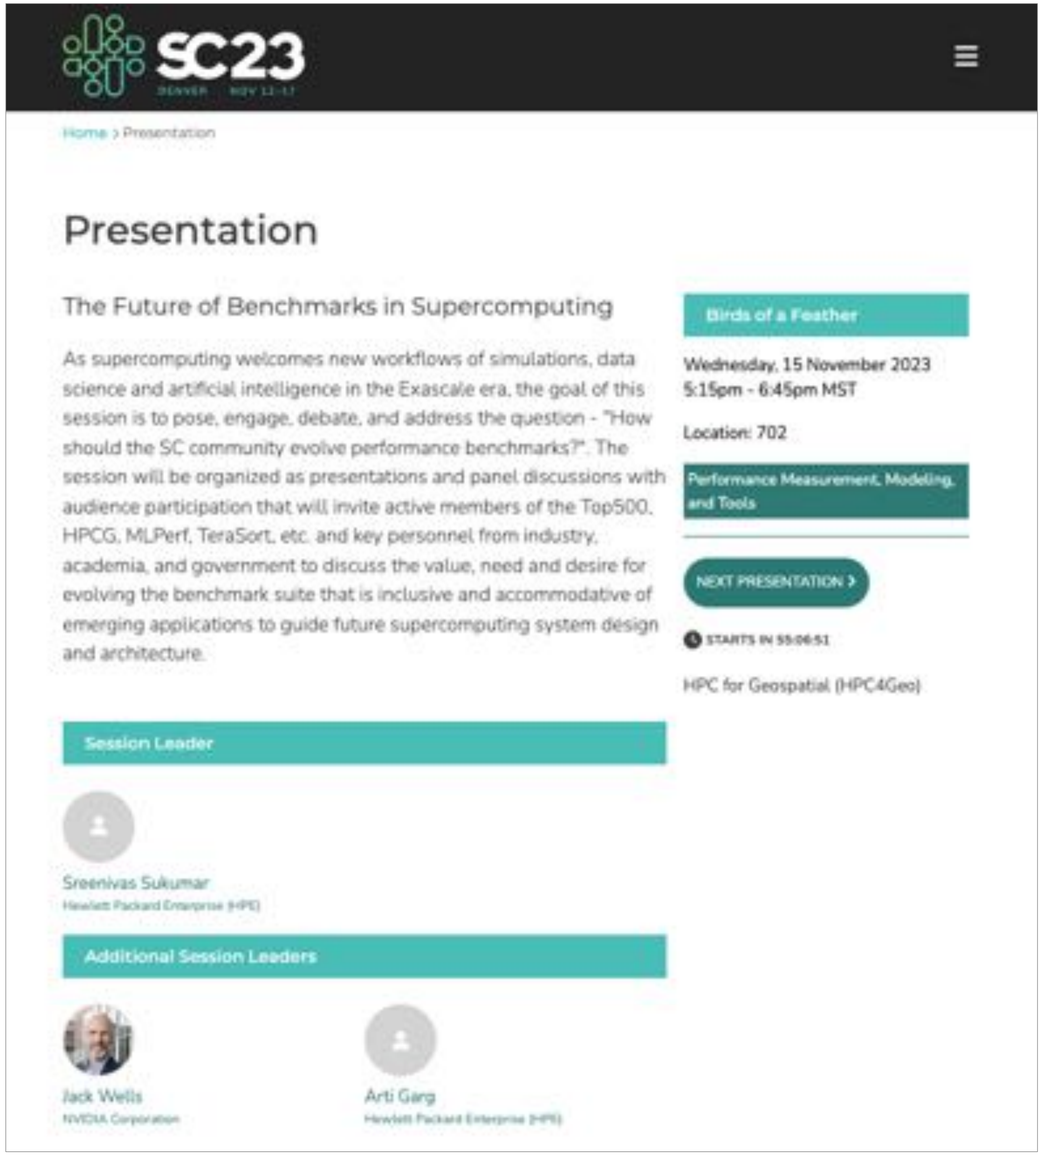

Additional Sustainability Topics

Grace CPU: Excellent Performance per Watt

OpenFOAM 2206, MotorBike 5M

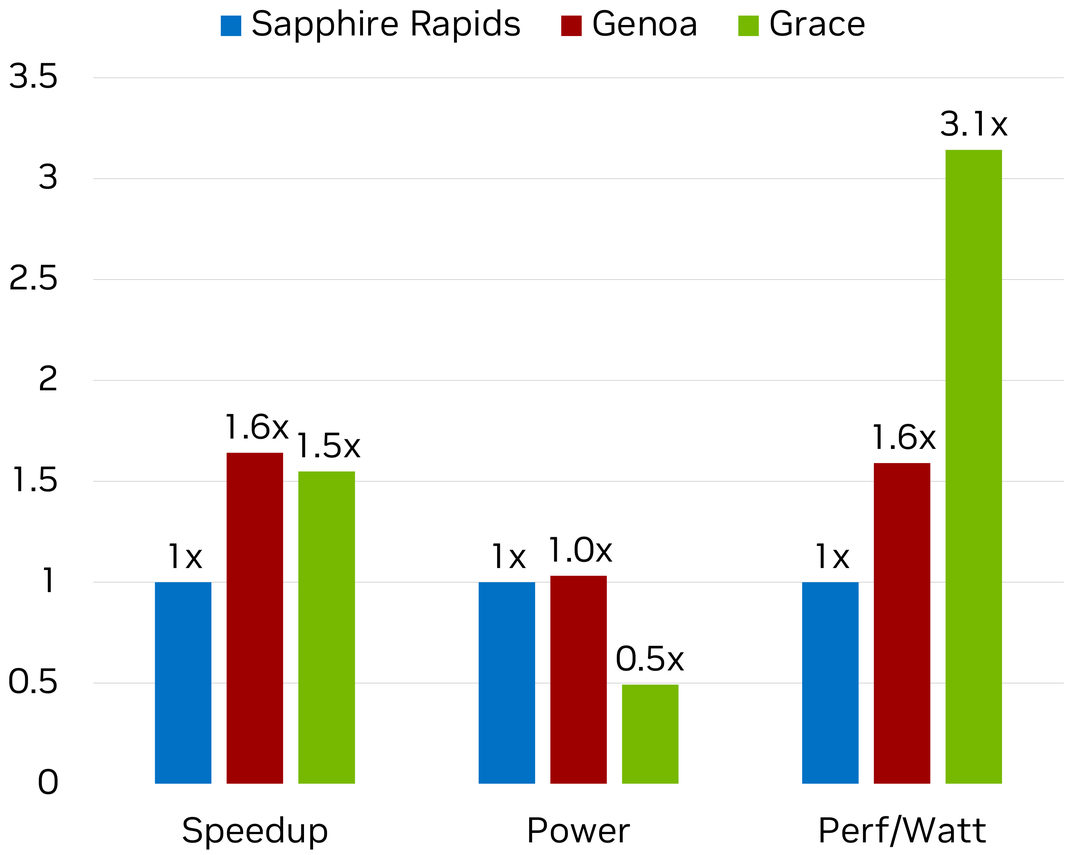

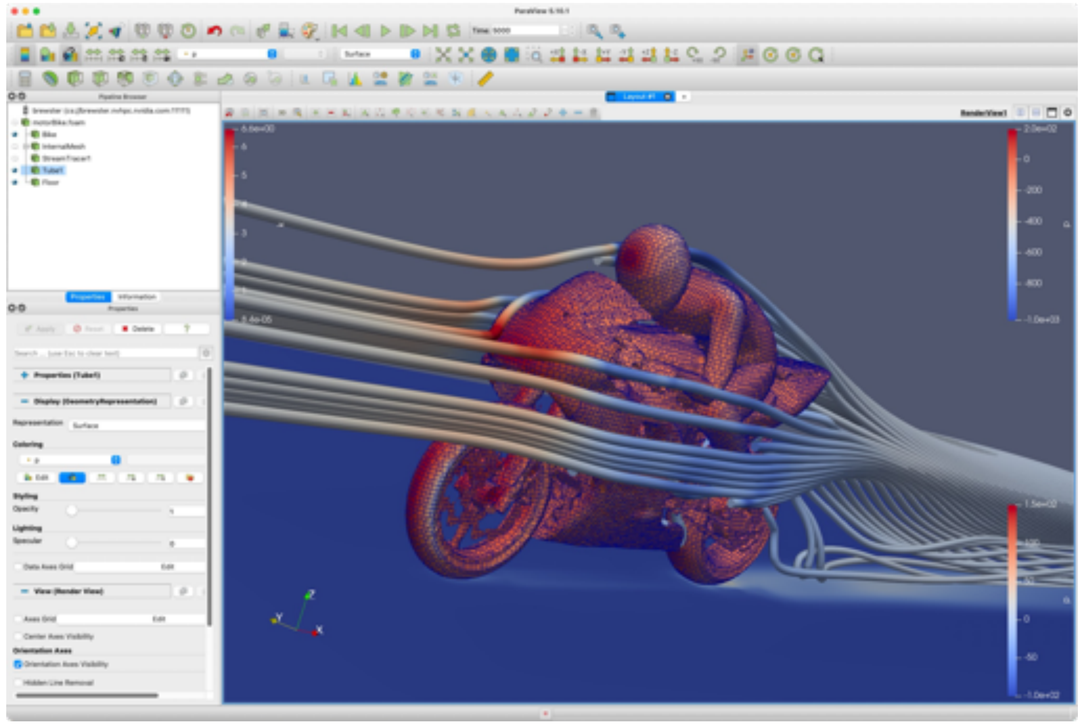

Sapphire Rapids: Intel Xeon Platinum 8470Q, 52c @ 2.1GHz - 3.8GHz
Genoa: AMD EPYC 9654, 96c @ 1.5GHz – 3.7GHz
Grace: NVIDIA Engineering Sample, 72c @ 3.2GHz
Best single socket time using best compilers (GCC, ICC, AOCC) and best rank/thread decomposition

Grace numbers based on engineering sample measurements

Grace Software Ecosystem is Built on Standards

The NVIDIA platform builds on optimized software from the broad Arm software ecosystem

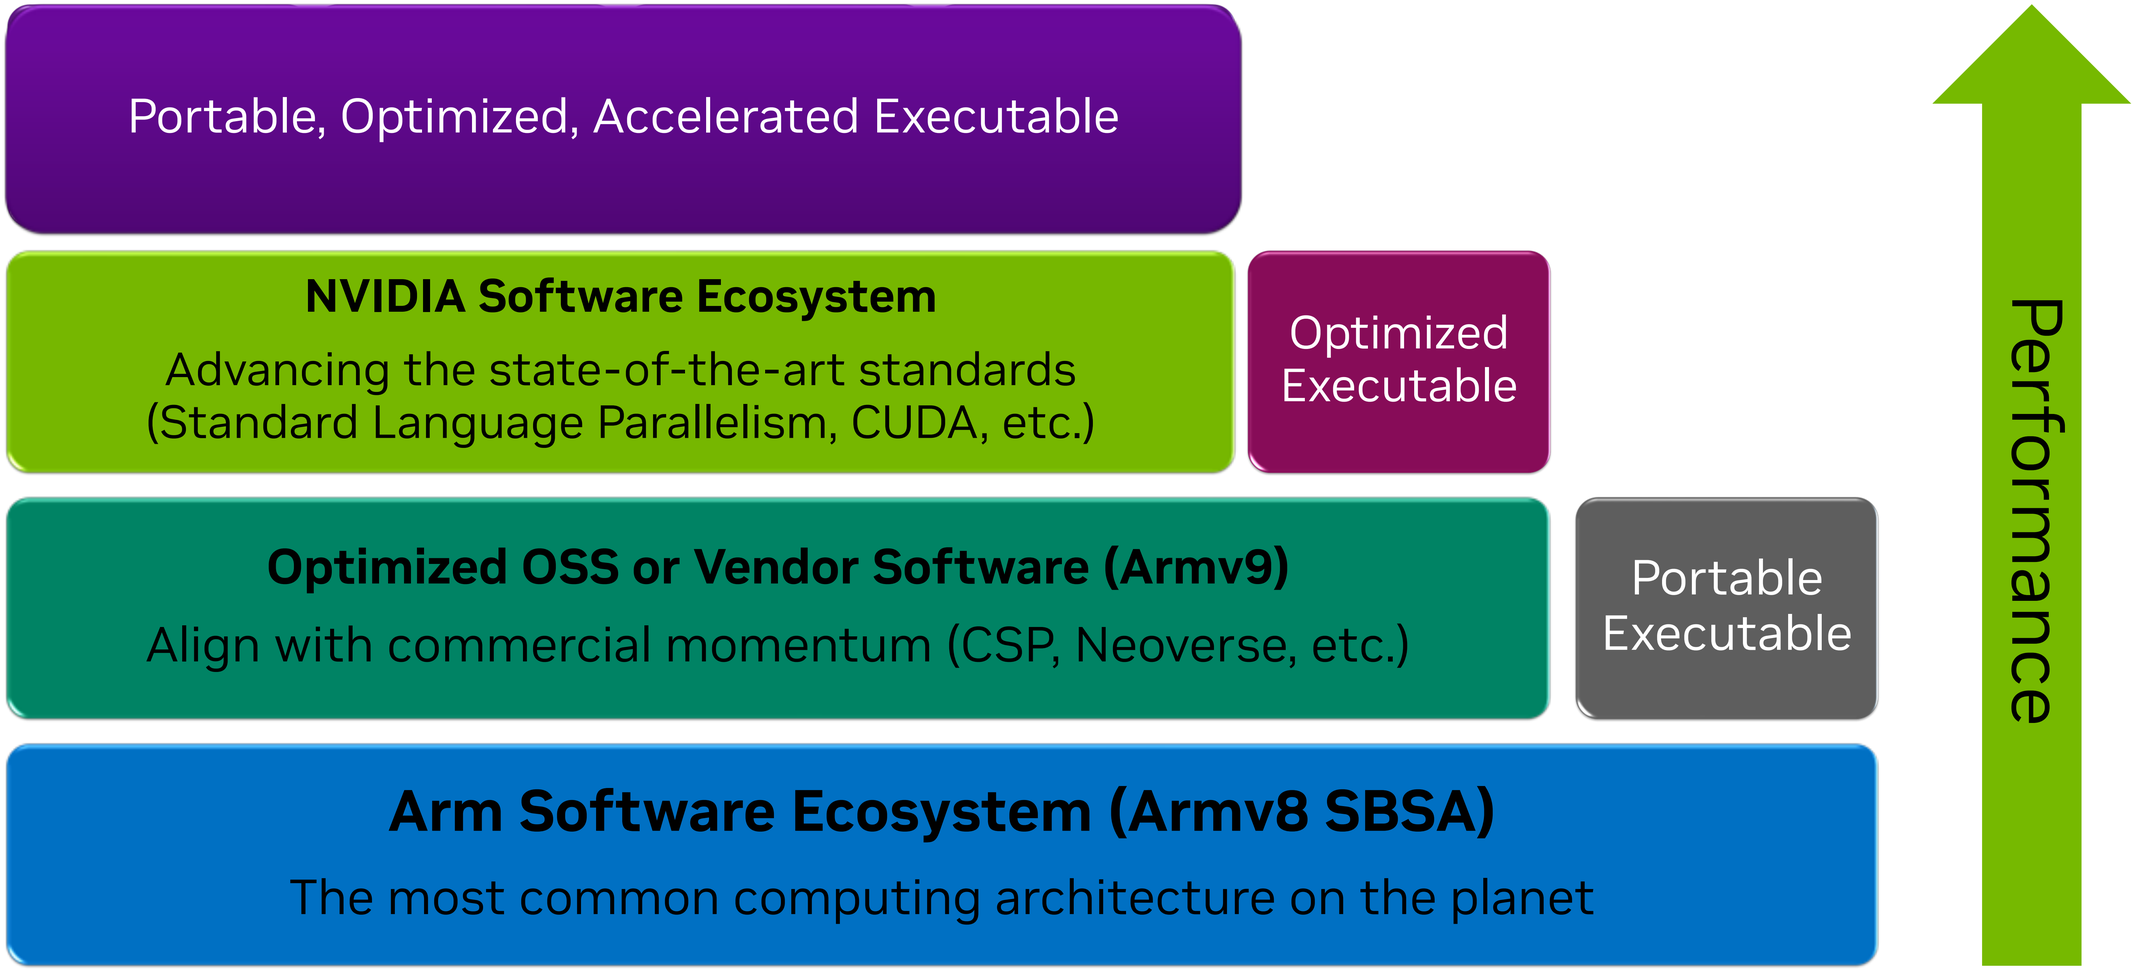

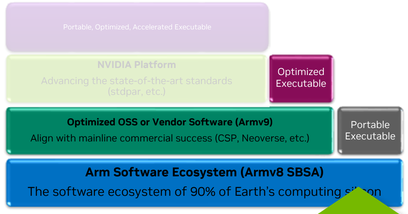

The Arm Software Ecosystem

Developed by NVIDIA, the OSS community, and other Arm partners

NVIDIA Software
Frameworks, Libraries, SDKs, Toolkits, etc.

Applications and Frameworks
- Deep Learning
- HPC
- General Compute

Developer Tools
Open Source and Commercial

Communication Libraries
Open Source and Commercial

Compilers
Open Source and Commercial

Libraries
Open Source and Commercial

Filesystems
All major parallel filesystems

Admin Tools & Packaging

OS
All Major Linux Distributions

Arm Server Ready Platform
Standard Firmware and RAS

Job Schedulers
Open Source and Commercial

Cluster Management
NVIDIA Bright, Open Source, and Commercial

Containers & Orchestration
NVIDIA Container Toolkit and All Major Frameworks and Tools

Programming the NVIDIA Platform

Open Standards Balanced with Platform Specialization

	Accelerated Standard Languages	Portability Frameworks & Directives	Platform Specialization
Languages & Programming Models	C++, python™, Fortran	kokkos, alpaka, RAJA, OpenACC More Science, Less Programming, OpenMP	NVIDIA CUDA, C++ | Fortran | Python
Libraries & Frameworks	Acceleration Libraries: Core, Math, Communication, Data Analytics, AI, DSLs		
Hardware	GPU	CPU	Interconnect

Wherever You Run

PCs | Workstations | On Prem | Cloud | At the Edge

HPC software ecosystem sustainability requires testing the software ecosystem at scale

Infrastructure for automated testing of the diverse hardware and software employed within the HPC enterprise should enhance ecosystem effectiveness.

- HPC users typically compose their software development tool chain from many sources, including DOE, GNU, and Vendor-supplied toolkits.

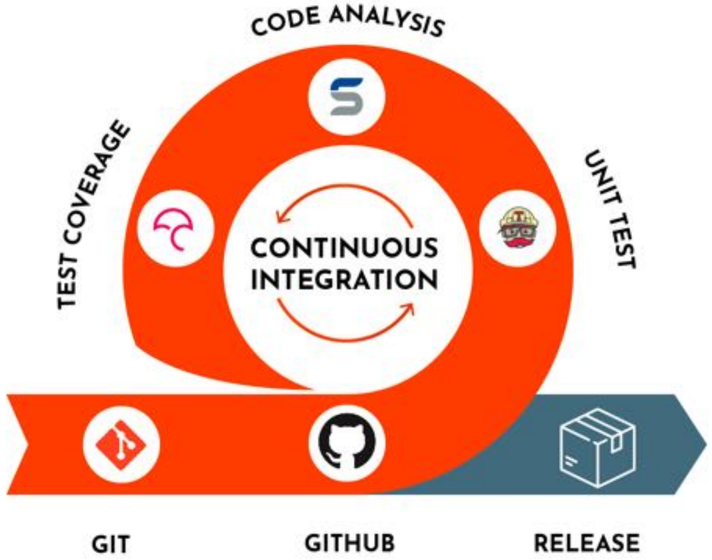

- NVIDIA maintains infrastructure for the automated testing of NVIDIA's software development kits but does not maintain infrastructure to test with the broad and varied composable HPC ecosystem.
- Ecosystem scale testing and integration services will improve the quality of the composed HPC software ecosystem.
- Large SciComp Centers have a unique opportunity to provide infrastructure and resources for software ecosystem testing.

Benchmarking and Sustainability: Conclusions

1. Benchmarks should be faithful to modern science applications
2. More open benchmark development practices would have many benefits
3. Influence of the benchmark list-makers is large
4. Energy-to-solution is a top-level priority metric moving forward
5. Sustainability in software strategy should include priority investments in ISO language parallelism (C++, Fortran), other *de-facto* standards, modern libraries, etc.
6. Software sustainability requires testing of the full, composed software ecosystem. Infrastructure to accomplish this are limited.

Acknowledgements

- Steve Oberlin – MaxQ optimization
- Chris Porter – Energy Corrected Speedup
- John Linford – Grace CPU Performance Results
- Jeff Larkin – Programming Models, Software Testing
- Rangan Sukumar, Arti Garg – SC23 Future of Benchmarks in Supercomputing BoF