NUG Monthly Meeting

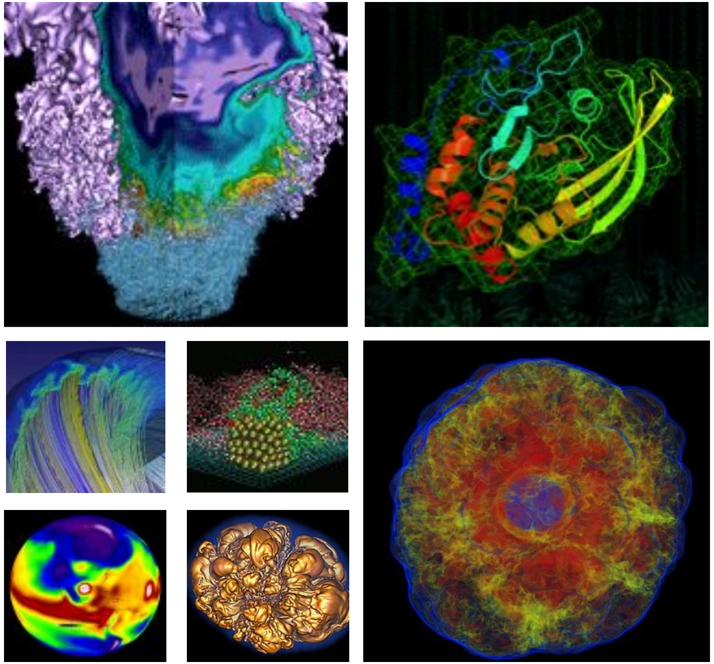

NERSC

19 November, 2020

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Today's plan

- Interactive - please participate!
  - Raise hand or just speak up
  - **NERSC User Slack** (link in chat), **#webinars** channel

- Agenda:
  - Win-of-the-month
  - Today-I-learned
  - Announcements/CFPs
  - Topic of the day: **NERSC and NERSC Users at SC20**
  - Coming meetings: topic suggestions/requests?
  - Last month's numbers

Win of the month

Show off an achievement, or shout out someone else's achievement, e.g.:

- Had a paper accepted
- Solved a bug
- A scientific achievement (maybe candidate for Science highlight, or High Impact Scientific Achievement award)
- An Innovative Use of High Performance Computing (also a candidate for an award) (https://www.nersc.gov/science/nersc-hpc-achievement-awards/)

Tell us what you did, and what was the key insight?

Win of the month

- Dilip Asthagiri had a paper published in J Physical Chemistry Letters

  Shows that temperature influences protein folding, perhaps prompts a re-thinking partitioning proteins into hydrophilic/hydrophobic groups.

  https://news.rice.edu/2020/11/11/folding-proteins-feel-the-heat-and-cold/

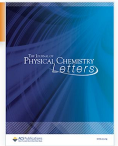

Hydrophilic Interactions Dominate the Inverse Temperature Dependence of Polypeptide Hydration Free Energies Attributed to Hydrophobicity

Dheeraj S. Tomar*, Michael E. Paulaitis*, Lawrence R. Pratt*, and Dilipkumar N. Asthagiri*

Cite this: *J. Phys. Chem. Lett.* 2020, 11, XXX, 9965–9970
Publication Date: November 10, 2020
https://doi.org/10.1021/acs.jpclett.0c02972

SUBJECTS: Free energy, Solution chemistry, Hydrophobicity, Solvents, Hydration

Abstract

We address the association of the hydrophobic driving forces in protein folding with the inverse temperature dependence of protein hydration, wherein stabilizing hydration effects strengthen with increasing temperature in a physiological range. All-atom calculations of the free energy of hydration of aqueous deca-alanine conformers, holistically including backbone and side-chain interactions together, show that attractive peptide–solvent interactions and the thermal expansion of the solvent dominate the inverse temperature signatures that have been interpreted traditionally as the hydrophobic stabilization of proteins in aqueous solution. Equivalent calculations on a methane solute are also presented as a benchmark for comparison. The present study calls for a reassessment of the forces that stabilize folded protein conformations in aqueous solutions and of the additivity of hydrophobic/hydrophilic contributions.

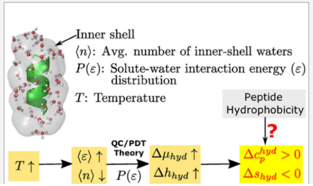

Win of the month

Show off an achievement, or shout out someone else's achievement, e.g.:

- Had a paper accepted
- Solved a bug
- A scientific achievement (maybe candidate for Science highlight, or High Impact Scientific Achievement award)
- An Innovative Use of High Performance Computing (also a candidate for an award) (https://www.nersc.gov/science/nersc-hpc-achievement-awards/)

Tell us what you did, and what was the key insight?

Today I learned

What surprised you that might benefit other users to hear about?

(and might help NERSC identify documentation improvements!)

Eg:

- Something you got stuck on, hit a dead end, or turned out to be wrong about
  - Give others the benefit of your experience!
  - Opportunity to improve NERSC documentation
- A tip for using NERSC
- Something you learned that might benefit other NERSC users

"If we knew what it was we were doing, it would not be called research, would it?" - Einstein

Announcements and CFPs

Check latest weekly email for these:

- NERSC User Survey - invitation coming this week!
  - You'll see an email from NBRI on behalf of NERSC
  - Please give feedback! The survey is short, and important for NERSC - both for reporting and for identifying what we should keep doing and what we should improve
- Maintenance yesterday: PE defaults unchanged, new 20.10 PE replaces old 20.03 one
- Upcoming training events for NVidia HPC SDK and for Totalview HPC debugger, Dec 8-10

Power maintenance Dec 15-18

The **final power upgrade** for the Perlmutter installation will take place **December 15-20.** During most of this time, power will be cut to the building where NERSC is housed. You can expect that for the duration of the outage, **all NERSC resources will not be available.**

Especially: we expect power to be out from Tuesday Dec 15 to Friday Dec 18

Changes to Premium Charging for 2021

- When to use premium
  - For a scientific emergency (i.e. not expiring allocation)
    - Unexpected scientific event, review that needs results right away
  - Infrequently
- Starting in AY2021, projects can use premium at x2 charge factor until they've spent 20% of their allocation on premium jobs
- Then charge increases to x4
  - NERSC may change this in response to premium usage
- PIs will be able to toggle permissions for users in their project to charge a job to premium in iris
  - Instructions will be posted on the web page shortly

New Julia kernel for jupyter.nersc.gov

- Added support for IJulia to jupyter.nersc.gov
  - Kernel has access to all officially-supported packages (in `/global/common/cori_cle7/software`) as well as the user packages (in `~/.julia`)
- Looking for folks to kick the tires

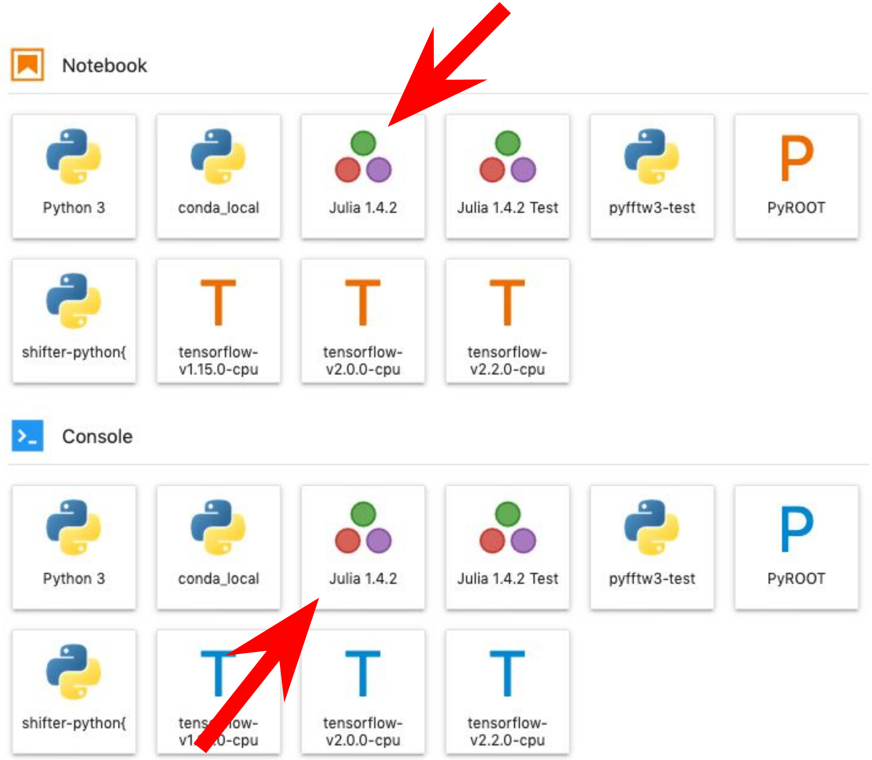

Topic of the day - NERSC at SC20

- Most NERSC users use computational science to advance research in another field
- But a number of NERSC users, and NERSC itself, contribute to the field of High Performance Computing itself
- This week and last, SC20 - the International Conference for High Performance Computing, Networking, Storage and Analysis - has been happening, and NERSC and a number of NERSC users have been involved
- So for today, we have a brief showcase of how NERSC participates in the field of HPC

SC20

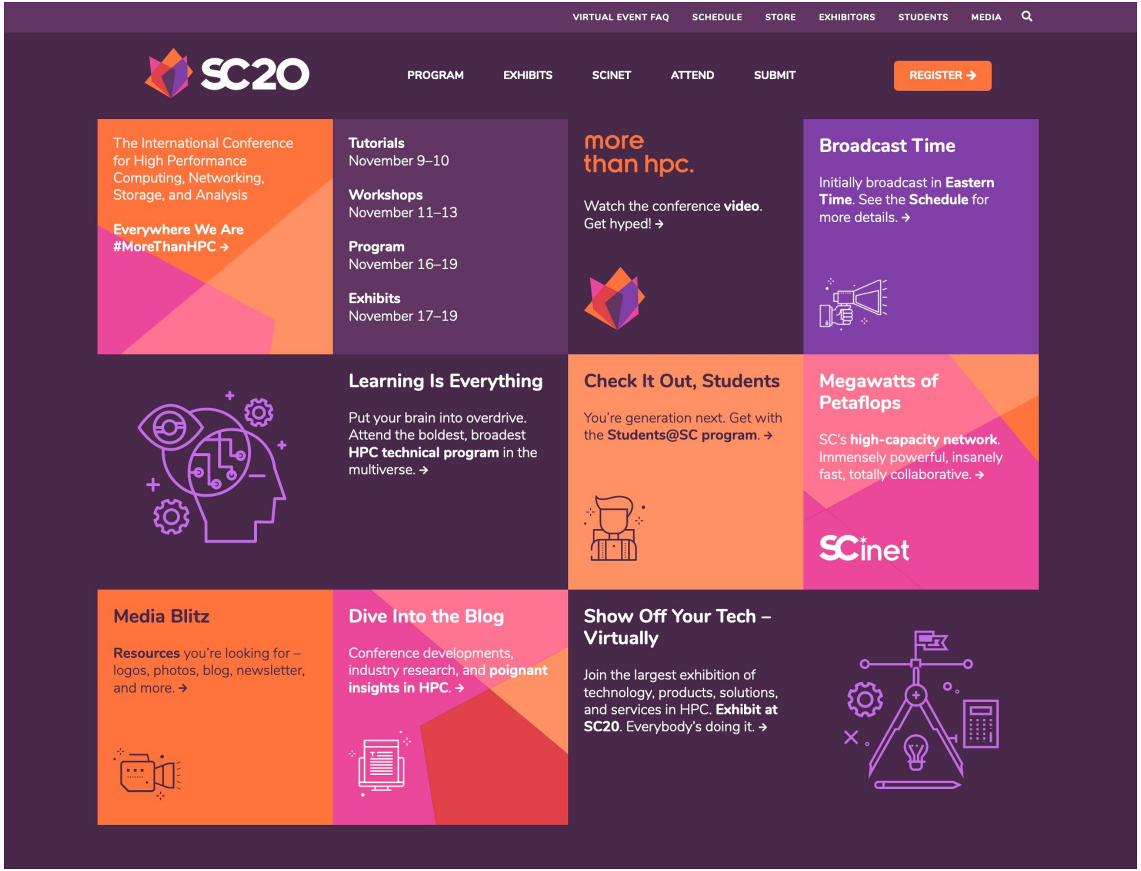
VIRTUAL EVENT FAQ
SCHEDULE
STORE
EXHIBITORS
STUDENTS
MEDIA
SC20
PROGRAM
EXHIBITS
SCINET
ATTEND
SUBMIT
REGISTER →
The International Conference for High Performance Computing, Networking, Storage, and Analysis
Everywhere We Are #MoreThanHPC →
Tutorials November 9–10
Workshops November 11–13
Program November 16–19
Exhibits November 17–19
more than hpc.
Watch the conference video. Get hyped! →
Broadcast Time
Initially broadcast in Eastern Time. See the Schedule for more details. →
Learning Is Everything
Put your brain into overdrive. Attend the boldest, broadest HPC technical program in the multiverse. →
Check It Out, Students
You're generation next. Get with the Students@SC program. →
Megawatts of Petaflops
SC's high-capacity network. Immensely powerful, insanely fast, totally collaborative. →
SCinet
Media Blitz
Resources you're looking for – logos, photos, blog, newsletter, and more. →
Dive Into the Blog
Conference developments, industry research, and poignant insights in HPC. →
Show Off Your Tech – Virtually
Join the largest exhibition of technology, products, solutions, and services in HPC. Exhibit at SC20. Everybody's doing it. →

Tutorials

NERSC staff and users presented:

- **Performance Tuning** with the Roofline Model on GPUs and CPUs
- In Situ Analysis and **Visualization** with SENSEI and Ascent
- **Parallel I/O** In Practice
- The **OpenMP** Common Core: A Hands-On Introduction
- **Deep Learning** at Scale
- **UPC++**: An Asynchronous RMA/RPC Library for Distributed C++ Applications
- Using **Containers** to Accelerate HPC

Workshops

NERSC staff and users either organized or presented in:

- **Women in HPC**
- First International Workshop on **Quantum Computing Software**
- 5th **Deep Learning** on Supercomputers Workshop
- HUST-20: 7th International Workshop on **HPC User Support Tools**
- Seventh SC Workshop on Best Practices for **HPC Training and Education**
- WORKS20: 15th Workshop on **Workflows** in Support of Large-Scale Science
- CANOPIE-HPC: **Containers** and New Orchestration Paradigms for Isolated Environments in HPC

More Workshops

NERSC staff and users either organized or presented in:

- Fifth International **Parallel Data Systems** Workshop
- The 11th International Workshop on **Performance Modeling, Benchmarking and Simulation** of High-Performance Computer Systems
- Seventh Workshop on **Accelerator Programming** Using Directives
- Fourth Workshop on **Interactive** High-Performance Computing
- 2nd Workshop on **Machine Learning** for Computing Systems
  2nd Annual Workshop on Extreme-Scale **Experiment-in-the-Loop-Computing**

Papers

NERSC

NERSC staff presented a few papers:
(plus more at workshops)

MESHFREEFLOWNET: A Physics-Constrained Deep Continuous Space-Time Super-Resolution Framework

Chiyu "Max" Jiang*† chiyu.jiang@berkeley.edu
Soheil Esmaeilzadeh*§ soes@stanford.edu
Kamyar Azizzadenesheli¶** Kamyar@purdue.edu
Karthik Kashinath‡ kkashinath@lbl.gov
Mustafa Mustafa‡ mmustafa@lbl.gov
Hamdi A. Tchelepi§ tchelepi@stanford.edu

Iris: Allocation Banking and Identity and Access Management for the Exascale Era

Gabor Torok
Mark R. Day

Tuning Floating-Point Precision Using Dynamic Program Information and Temporal Locality

Hugo Brunie*, Costin Iancu*, Khaled Z. Ibrahim*, Philip Brisk†, Brandon Cook*
*Lawrence Berkeley National Laboratory
{hbrunie,cciancu,kzibrahim,BGCook}@lbl.gov
†University of California, Riverside
philip@cs.ucr.edu

Abstract—We present a methodology for precision tuning of full applications. These techniques must select a search space composed of either variables or instructions and provide a scalable search strategy. In full application settings one cannot assume compiler support for practical reasons. Thus, an additional important challenge is enabling code refactoring. We argue for an instruction-based search space and we show: 1) how to exploit dynamic program information based on call stacks; and 2) how to exploit the iterative nature of scientific codes, combined with temporal locality. We applied the methodology to tune the implementation of scientific codes written in a combination of Python,

approach to select the appropriate implementation each time the function is called. We anticipate that this problem will become increasingly prevalent over the next decade, given current trends for hardware specialization in High Performance Computing (HPC).

We argue in favor of a generic performance-driven methodology to tune the precision of full-fledged applications. To maximize the appeal of our approach to HPC practitioners, we assume no compiler support, while providing the ability to refactor code to accommodate mixed precision. We also designed and implemented heuristics

Accelerating Large-Scale Excited-State GW Calculations on Leadership HPC Systems

Mauro Del Ben*, Charl … …ie*‡ and Jack Deslippe*
mdelben@lbl.gov cjy… …edu jrdeslippe@lbl.gov
*L… …SA
†Department of M… …5, California, USA
‡Department of Physics, University of California at Berkeley, California 94720, USA

Gordon bell finalist!

***Abstract*—Large-scale GW calculations are the state-of-the-art approach to accurately describe many-body excited-state phenomena in complex materials. This is critical for novel device design but due to their extremely high computational cost, these calculations often run at a limited scale. In this paper, we present algorithm and implementation advancements made in the materials science code BerkeleyGW to scale calculations to the order of over 10,000 electrons utilizing the entire Summit at OLCF. Excellent strong and weak scaling is observed, and a 105.9 PFLOP/s double-precision performance is achieved on 27,648 V100 GPUs, reaching 52.7% of the peak. This work for the first time demonstrates the possibility to perform GW calculations at such scale within minutes on current HPC systems, and leads the way for future efficient HPC software development in materials, physical, chemical, and engineering sciences.**

***Index Terms*—electronic structure, excited states, GW method.**

electronic engineering. Exactly solving for the properties of electrons using quantum mechanics requires a computational cost that increases exponentially with the problem size. This is intractable in practice, and approximations *must* be made.

First-principles computations aim at obtaining important materials' properties using fundamental theories without fitting parameters from experiments, and density-functional theory (DFT) [1] is one of the most common approaches. However, DFT is a ground-state theory and it presents serious errors for excited-state properties. For example, the band gap (an excited-state property) of silicon is underestimated by over 50% with the standard approximations in DFT [2].

The first-principles GW approach [2], [3] is the prevail-

Papers

NERSC Users presented even more

Improving All-to-Many Personalized Communication in Two-Phase I/O

Qiao Kang*, Robert Ross§, Robert Latham§, Sunwoo Lee*, Ankit Agrawal*, Alok Choudhary*, and Wei-keng Liao*

denoted as file domains. ...ests and data, based on ... of the processes in the ...out I/O operations with ...In general, the commu- ... and I/O aggregators is ...lementation of the MPI-

Foresight: Analysis That Matters for Data Redu...

1st Pascal Grosset, *Data Science at Scale, CCS-7*, *Los Alamos National Laboratory*, Los Alamos, USA, pascalgrosset@lanl.gov

2nd Christopher M. Biwer, *Data Science at Scale, CCS-7*, *Los Alamos National Laboratory*, Los Alamos, USA, cmbiwer@lanl.gov

3rd Jesus Pul..., *Data Science at Sca...*, *Los Alamos National ...*, Los Alamos, U..., pulido@lanl.g...

4th Arvind T. Mohan, *CCS-2: Computational Physics and Methods*, *Los Alamos National Laboratory*, Los Alamos, USA, arvindm@lanl.gov

5th Ayan Biswas, *Data Science at Scale, CCS-7*, *Los Alamos National Laboratory*, Los Alamos, USA, ayan@lanl.gov

6th John Patc..., *Data Science at Sca...*, *Los Alamos National ...*, Los Alamos, U..., patchett@lanl.

7th Terece L. Turton

8th David H. Rogers

9th Daniel Livescu

Distributed Many-to-Many Protein Sequence Alignment using Sparse Matrices

Oguz Selvitopi*[1], Saliya Ekanayake†[1], Giulia Guidi‡*, Georgios A. Pavlopoulos§, Ariful Azad¶, Aydın Buluç*‡

Computational Research Division, Lawrence Berkeley National Laboratory, USA

†*Microsoft Corporation, USA*

...erkeley, USA

...r Fleming", 34 Fleming Street, 16672, Vari, Greece

... USA

...ov

...ation of our work is the identification of protein families, ... are groups of proteins that descend from a common ...tor. It is a hard problem since the relationship between ...nce similarity and homology is imprecise, meaning that ...annot use a similarity threshold to accurately conclude ...vo proteins are homologous or belong to the same family. ...ferent algorithms attempt to infer homology directly via ...rity search [1], [2] with variable degree of success in ... of sensitivity and specificity. Alternatively, one can

A Performance-Portable Nonhydrosta... Atmospheric Dycore for the Energy Exasca... System Model Running at Cloud-Resol... Resolutions.

Luca Bertagna, Sandia National Laboratories, PO BOX 5800, 87122, Albuquerque, NM, USA, lbertag@sandia.gov

Oksana Guba, Sandia National Laboratories, PO BOX 5800, 87122, Albuquerque, NM, USA, onguba@sandia.gov

Mark A. Taylor, Sandia National Laboratories, PO BOX 5800, 87122, Albuquerque, NM, USA, mataylo@sandia.gov

Jar..., Sandia N..., PO B..., Albuqu..., jgfo...

Jeff Larkin, NVIDIA Corporation, Santa Clara, CA, USA, jlarkin@nvidia.com

Andrew M. Bradley, Sandia National Laboratories, PO BOX 5800, 87122, Albuquerque, NM, USA, ambradl@sandia.gov

Sivasankaran Rajamanickam, Sandia National Laboratories, PO BOX 5800, 87122, Albuquerque, NM, USA, srajama@sandia.gov

Andr..., Sandia N..., PO B..., Albuqu..., agsa...

Abstract

We present an effort to port the nonhydrostatic atmosphere dynamical core of the Energy Exascale Earth System Model (E3SM) to efficiently run on a variety of architectures, in-

II. Performance Attri...

Attribute title	Attribute valu...
Category achievement	Time-to-soluti...
Type of method used	Finite elemen... time-stepping

Task Bench: A Parameterized Benchmark for Evaluating Parallel Runtime Performance

Elliott Slaughter*[1], Wei Wu†[1], Yuankun Fu‡, Legend Brandenburg§, Nicolai Garcia§, Wilhem Kautz§, Emily Marx§, Kaleb S. Morris§, Qinglei Cao¶, George Bosilca¶, Seema Mirchandaney*, Wonchan Lee‖, Sean Treichler‖, Patrick McCormick†, Alex Aiken§

*SLAC National Accelerator Laboratory, †Los Alamos National Laboratory, ‡Purdue University, §Stanford University, ¶University of Tennessee, Knoxville, ‖NVIDIA

[1]Corresponding authors: eslaught@slac.stanford.edu, wwu@lanl.gov

***Abstract*—We present Task Bench, a *parameterized* benchmark designed to explore the performance of distributed programming systems under a variety of application scenarios. Task Bench dramatically lowers the barrier to benchmarking and comparing multiple programming systems by making the implementation for a given system orthogonal to the benchmarks themselves: every benchmark constructed with Task Bench runs on every Task Bench implementation. Furthermore, Task Bench's parameterization enables a wide variety of benchmark scenarios that distill the key characteristics of larger applications.**

To assess the effectiveness and overheads of the tested systems, we introduce a novel metric, *minimum effective task*

to implement m benchmarks on n systems is $\mathcal{O}(mn)$. Task Bench's design reduces this work to $\mathcal{O}(m+n)$, enabling dramatically more systems and benchmarks to be explored for the same amount of programming effort. New benchmarks created with Task Bench immediately run on all systems, and new systems that implement the Task Bench interface immediately run all benchmarks.

Benchmarks in Task Bench are based on the observation that regardless of the programming system in which an application is written, many applications can be modeled as coarse-grain

Posters

NERSC staff and users presented at least a few posters

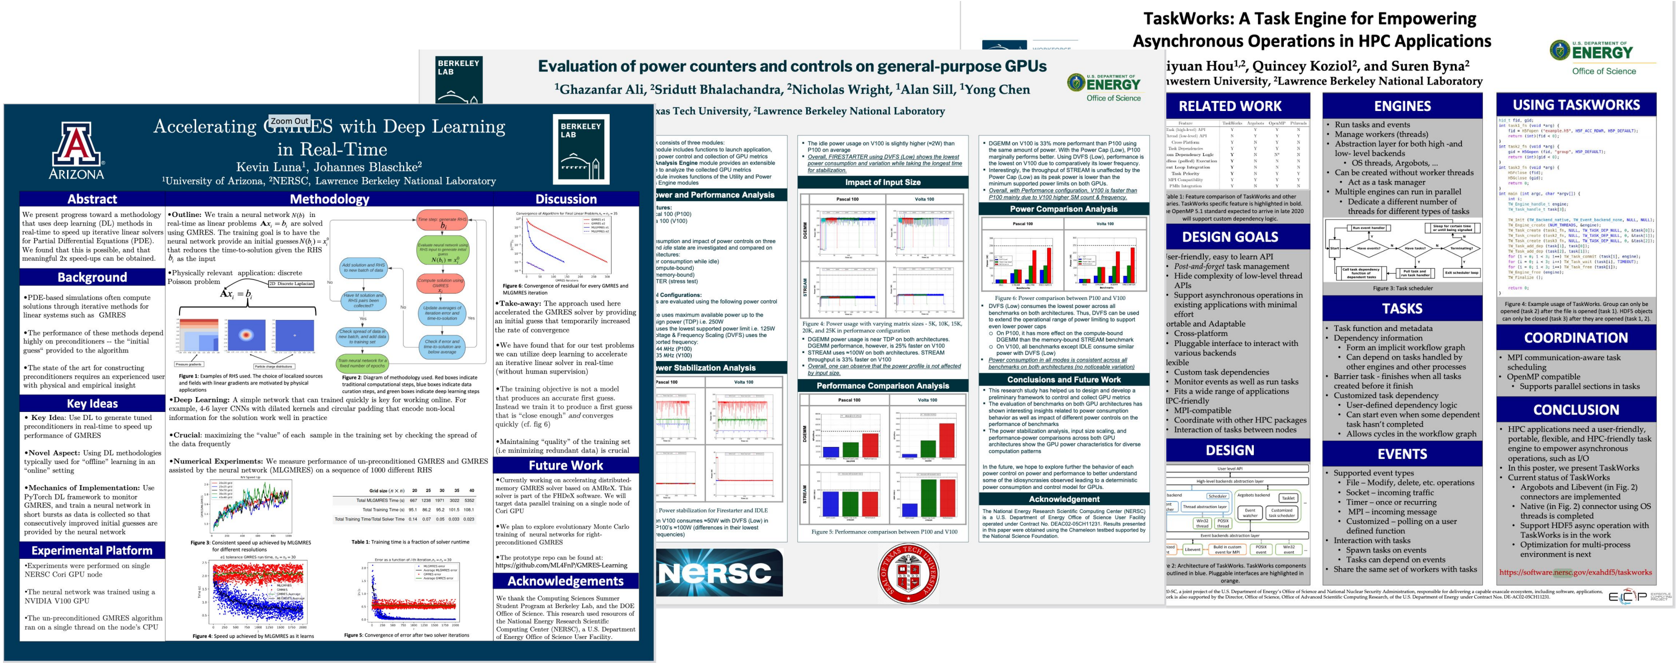

Students

- NERSC staff and users contributed to the Students@SC program
- Pic below is leaderboard for Virtual Student Cluster Competition (winner to be announced today!)

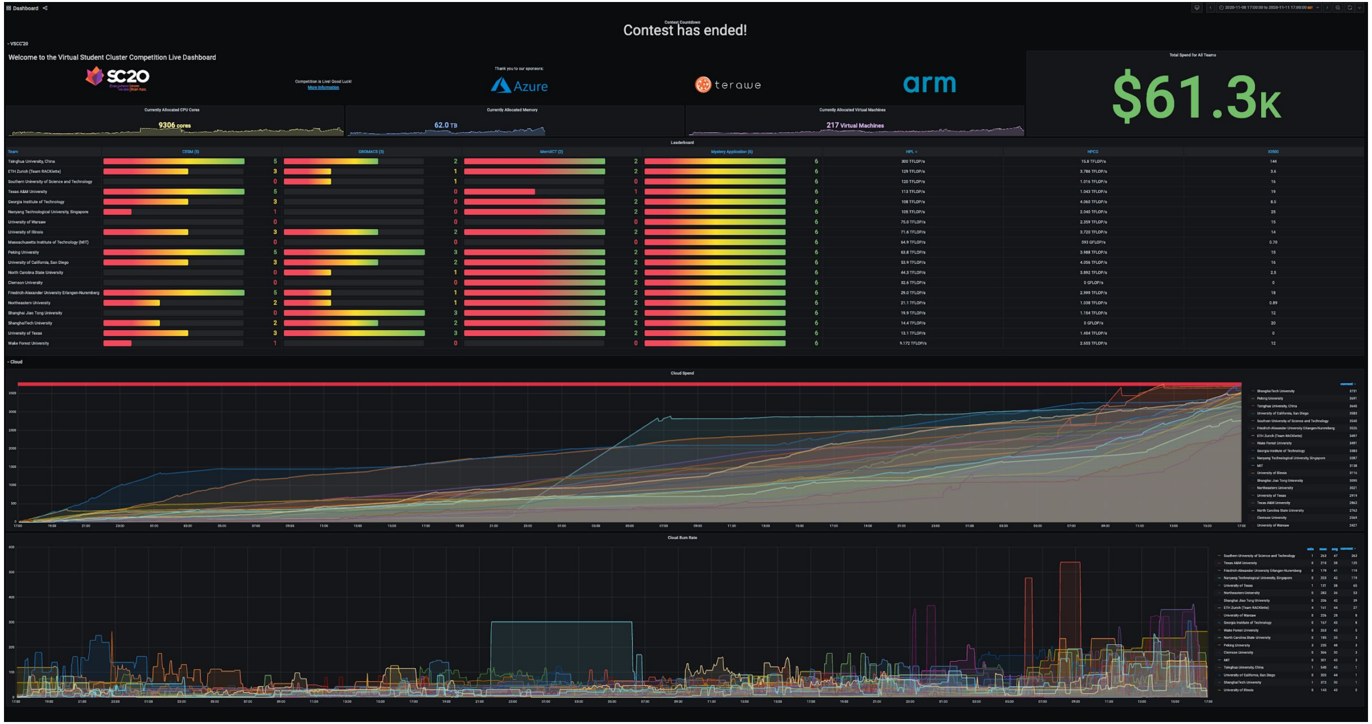

BoFs, Panels and more

NERSC staff were involved in:

- State-of-the-Practice talks
- Spack Community BoF
- Panels:
  - Exotic Computation and System Technology
  - Diverse Approaches to Tiering HPC Storage
  - The Inner Workings of SCinet as Told by Four Alumni of the Women in IT Networking at SC (WINS) Program
- SCinet
- Committees

NERSC at SC20

Key take-home here is:

NERSC is heavily involved in the development of HPC as a field

More than just a provider of compute resources

Our users are a key part of this

Coming up

Topic requests/suggestions?

- Eg:
  - Deep(ish) dive into NERSC's Slurm setup
  - HPSS tips, and stories of the move from Oakland to Berkeley

We'd love to hear some lightning talks **from NERSC users** about the research you use NERSC for!

Last month's numbers - October

Scheduled and overall availability:

	Scheduled	Overall
Cori	95.9%	94.2%
HPSS	100.00%	100.0%
CFS	100.00%	99.8%

Cori:

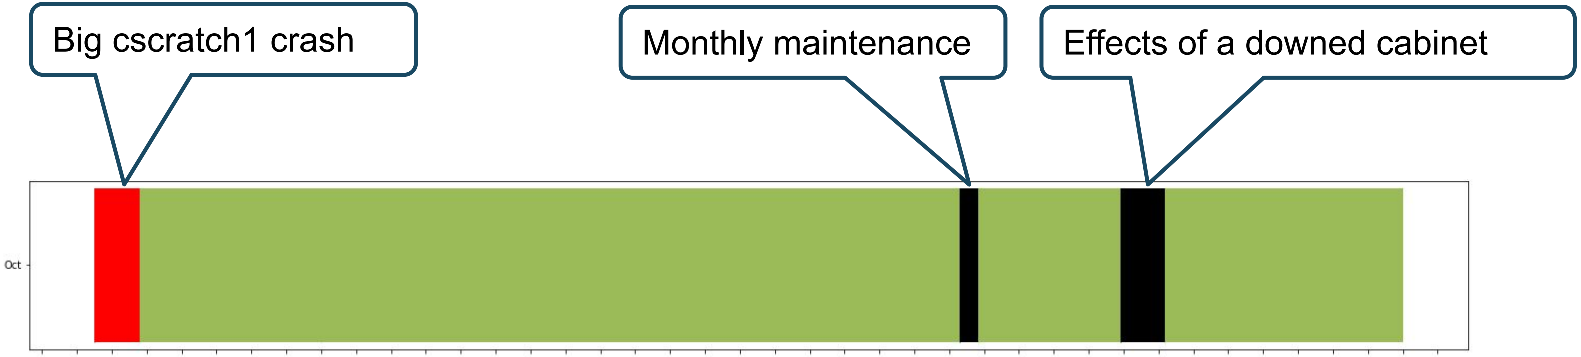

Last month's numbers - October

Cori Utilization: 91.4%

Large jobs: 28.0%

New Tickets: **669**

Closed Tickets: **711**

Backlog at 1 Oct: **594**

Thank You