



SC21 Student Cluster Competition

SCC, VSCC, and Indy SCC

Housekeeping

Q&A session at end - please use the link above slides to ask questions

We're recording this session, the recording and these slides will be posted on the webinars page

Today:

- Meet the committee
- What is a student cluster competition?
- What's it like? Reflections of past participants
- About this year's contest(s)
- How to prepare
- Introducing the IndySCC!
- Q&A

Introducing the SCC/VSCC and Indy SCC Chairs



Kathleen Shoga
SCC/VSCC Chair



Stephanie Brink
SCC/VSCC Vice-Chair



Aroua Gharbi
IndySCC Co-Chair



Darshan Sarojini
IndySCC Co-Chair

Introducing the Committee



Yasaman Ghadar
Applications Lead



Rigoberto Moreno Delgado
Infrastructure Lead



Andy Howard
Cloud Lead



Sam Coleman
Onsite

Introducing the Committee

Tapasya Patki
Power Cap Component
Co-Lead

Dan Ellsworth
Power Cap Component
Co-Lead

Le Mai Weakley
Reproducibility
Challenge Chair

Ann Gentile
Posters/Lightning Talks
Lead



Steve Leak
Webinars



Georgios
Markomanolis
Applications,
IO500



Rebecca
Hartman-Baker
Prior Year SCC
Chair

Hai Ah Nam
Prior Year SCC
Chair



Ben Schwaller
Posters/Lightning
Talks

Introducing the committee



Abhinav Thota
Applications



Amiya Maji
Applications



Paul Peltz
Infrastructure



Dan Dietz
Infrastructure

Ben Matthews
Infrastructure



Ramin Nosrat
Committee Member

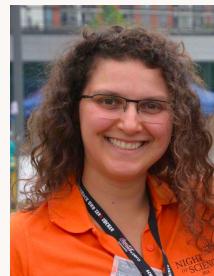


Giuseppe (Joe) Di Natale
Committee Member



Danielle Sikich
Committee Member

SCaler Volunteers

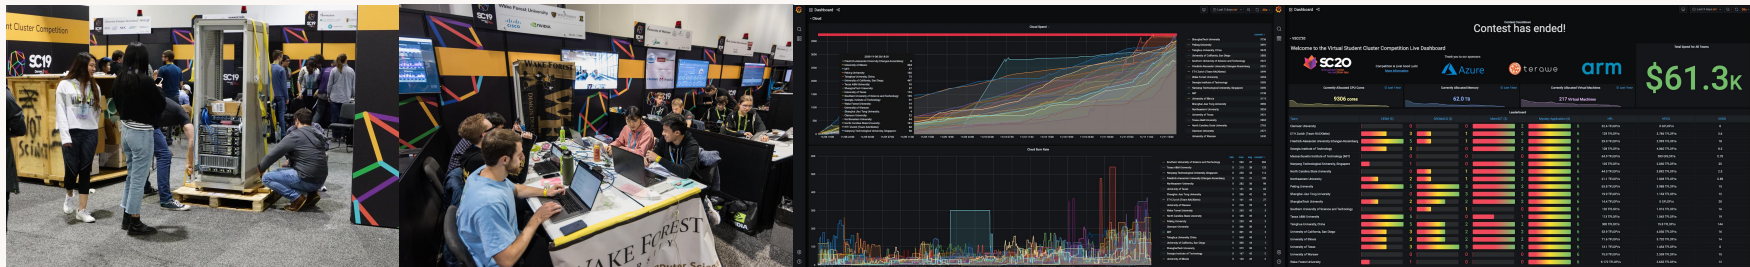


Anja Gerbes

John Ravi

The SC21 Student Cluster Competition

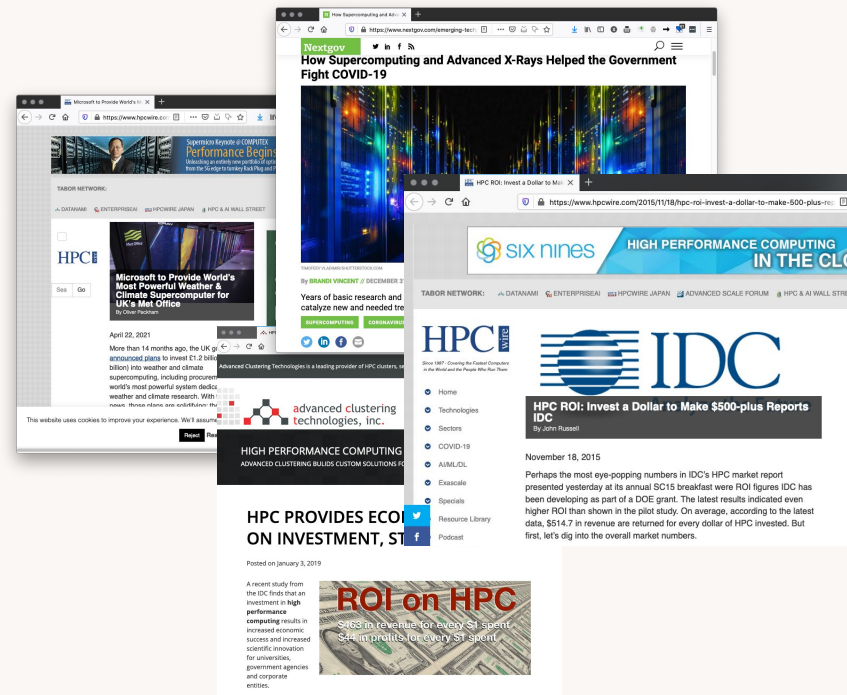
SC is The International Conference for High Performance Computing, Networking, Storage, and Analysis.



Nov 14-19
America's Center,
St Louis, MO

What is a Student Cluster Competition?

- HPC is used to solve complex problems in many areas including science and engineering
- Cloud-based HPC has increasing adoption
- The SCC provides an opportunity for students to be introduced to HPC, build skills, and connects students with HPC professionals



Screenshots from <https://www.advancedclustering.com/hpc-provides-economic-return-investment-study-finds/>,
<https://www.nextgov.com/emerging-tech/2020/12/how-supercomputing-and-advanced-x-rays-helped-government-fight-covid-19/171070/>,
<https://www.hpcwire.com/2015/11/18/hpc-roi-invest-a-dollar-to-make-500-plus-reports-idc/> and
<https://www.hpcwire.com/2021/04/22/microsoft-to-provide-worlds-most-powerful-weather-climate-supercomputer-for-uks-met-office/>

Ok, but What actually is a Student Cluster Competition?

Traditionally: A 48-hour, non-stop contest to build and run a (small) supercomputer, within a power budget and with the ability to recover from a disaster.

More recently: Also a contest to deploy and use cloud-based HPC resources for a HPC workload, within a money budget.

This year's teams compete in either a **physical** (SCC) or **cloud-based** (VSCC).

Students will:

- Design their architecture
- Deploy their clusters
- Run benchmarks
- Use HPC resources to meet the demands of real scientific computing applications
- Report on their results

What's it like?

Let's hear from some of the people who were in it last year!

Competitor Sophia and team advisor Sam from Wake Forest University, and competitors Max and Jacob and team advisor Mary from UC San Diego, chatted about their experiences in the following video



SC
COMPUTER CENTER

Hybrid SCC and a New Competition

Student Cluster Competition:

- SCC (In-Person): Teams will travel to St Louis
 - Fewer physical teams due to COVID restrictions
 - Each team designs, brings and builds their own cluster
- VSCC (Virtual): Teams compete from wherever they are, deploying instances on Microsoft Azure within a set budget
 - The exact budget is unknown until the competition begins!
 - Teams have a practice budget to prepare beforehand

New Competition: IndySCC

- A new education-focused track of the student cluster competition
- Lower entry barriers and provide continuous support

What's the process?

- Preparation before the event:
 - Teams form partnerships with institutions and vendors
 - **Design a cluster** (physical or in the cloud), **practice building and running** the applications
 - **Webinars** to help you prepare
 - In-person teams will need to consider logistics of getting to the competition
 - SCC provides conference registration for in-person and virtual teams (6 team members plus 1 advisor)
 - Also lodging, for in-person teams
 - We encourage institution and vendor partners to budget for other expenses
 - Build and install your cluster, **make posters, start on reports**

What's the process? - In Person SCC

- Friday/ Saturday
 - Teams arrive
- Saturday
 - Afternoon: Safety briefings and Student Orientations
 - After the briefings: building begins!
- Monday morning:
 - **Benchmarking** begins!
 - At the end of benchmarking, **final configuration is locked** in - no more hardware changes (including turning things off!)
- Monday evening:
 - **Mystery app, datasets announced**
 - **Main contest starts**
- Monday -> Wednesday
 - Teams run applications, gather results, present lightning talks and posters
- Wednesday **5:30pm**
 - **Competition ends!**
- Thursday
 - **Results announced!**

What's the process? - Virtual (VSCC)

- Monday morning:
 - **Benchmarking** begins!
- Monday evening:
 - **Mystery app, datasets announced**
 - **Main contest start**
- Monday -> Wednesday
 - Teams run applications, gather results, present lightning talks and posters (via Zoom or other platform)
- Wednesday **5:30pm**
 - **Competition ends!**
- Thursday
 - **Results announced!**

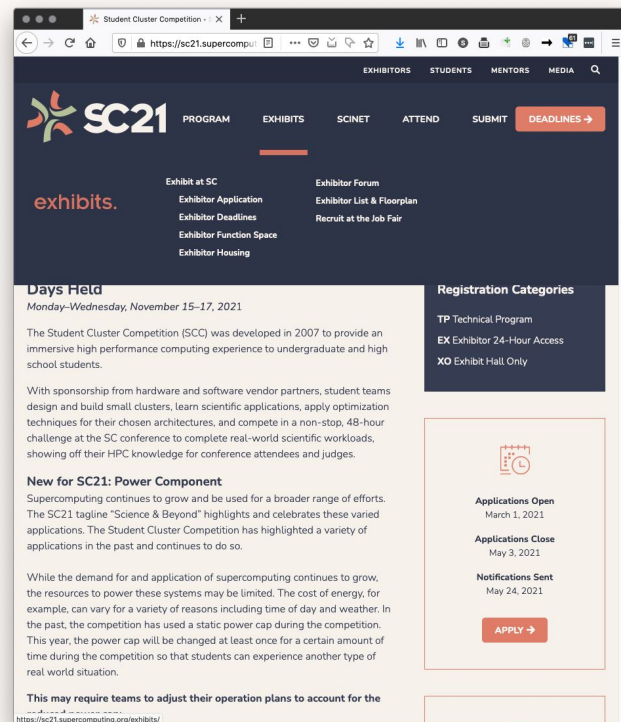
Hardware, cloud resources and software: Things to know

SCC (In person)

- New power challenges
 - Teams must adapt to power caps between 2500W and 4000W without turning off any hardware

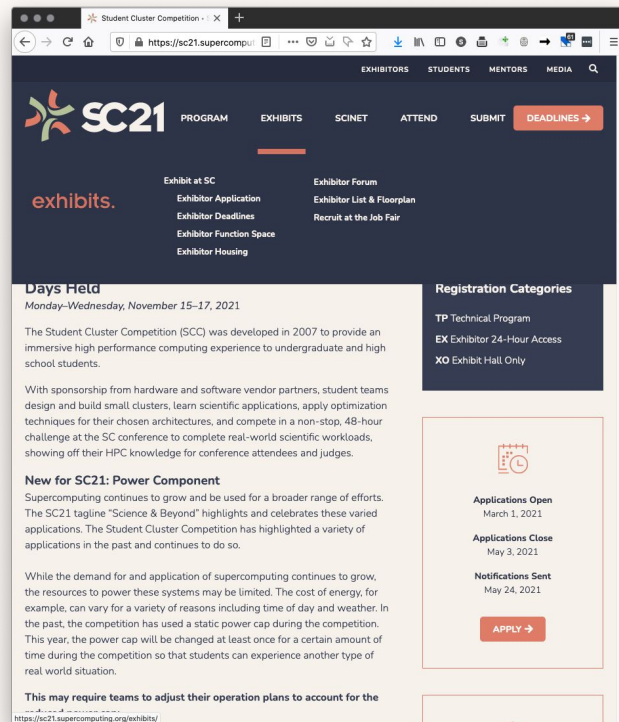
VSCC (Virtual)

- Cloud budgets
 - Separate budgets for practice and the competition
 - Competition budget will not be announced until the beginning of the competition!



Competition Components

- Benchmarks - get the highest HPL, HPCG and IO500 score
- Applications - complete tasks and report results and performance
- Posters and Lightning talks - present your team in the format of a scientific conference
- Reproducibility challenge: reproduce the results of an SC20 paper, and write your findings in a journal-quality report
- Real-world surprises - unpredictable impacts to power, budget, etc to prepare for



The benchmarks

- Linpack (HPL) (Shows peak performance)
- HPCG (A more typical computational science kernel, tougher for performance)
- IO500 (You will need storage, and a parallel filesystem)

The applications

See

<https://sc21.supercomputing.org/program/studentssc/student-cluster-competition/>

- Cardioid - cardiac simulation suite, Gordon Bell finalist
 - Parallelized using MPI, and has separate optimized loops for taking advantage of OpenMP, SIMD instruction sets for CPU architectures, and CUDA for GPU architectures
- Quantum ESPRESSO - first-principles electronic-structure calculations and materials modeling
 - Webinar Mon, June 28, 2021 from 1:00 PM – 2:00 PM PDT
- Reproducibility challenge: “[A Parallel Framework for Constraint-Based Bayesian Network Learning via Markov Blanket Discovery](#)” by Ankit Srivastava, Sriram P. Chockalingam, and Srinivas Aluru.
- Mystery App!

Each app will have a deep-dive webinar to help you prepare

Posters and Lightning Talks

Each team will present a poster and a lightning talk about their team and the design to conference attendees

- Introduction to your team and school
- Description of your cloud configuration and SKUs / hardware architecture with discussion on why you choose it
- Discussion of your deployment and management strategy of cloud resources / hardware and software resources
- Discussion of how you prepared for the competition
- Discussion of your general strategies for running and optimizing the applications
- Why your team will win!

Exact specifications for the poster and lightning talk are TBD but we hope to get them soon!

Communications (Coming Soon)

Teams will be invited to the Students@SC slack in order to ask questions from the Applications Experts and SCC Committee

- All students and the team advisor **must** join the slack
- Announcements will be made on the slack as well as through email

New this year

- A google document will be available for teams to see an FAQ-like page of previously answered questions or announcements in addition to the slack
- Google group will be a do not reply list for access to the google doc

The IndySCC

Website: <https://sc21.supercomputing.org/program/studentssc/indyscc/>

Email: indyscc@info.supercomputing.org

Timeline

IndySCC will be conducted in 4 phases from July to November. Think of it as a university course semester-long project during which you get hands-on training and support from HPC leaders.

- Regular webinars covering cloud platform, benchmark, and applications
- “Assignment” reports due periodically and count towards final competition score

Phase 1 (July 2021): Familiarizing with cloud resources and learning about architecture

Phase 2 (August 2021): HPL/HPCG benchmarks on different architectures and compiler

Phase 3 (September, October 2021): Experiment with Applications

November 6th and 7th 2021 (weekend): Compete!

- 48 hours straight
- New datasets for the applications
- Teams will run applications with maximum throughput under a power cap
- Performance and power metrics used by experts to judge the competition

Cloud Resources and Hardware

- **Cloud Vendor**
 - Chameleon Cloud: <https://www.chameleoncloud.org/>
 - Free access given to all teams
 - Login and provide each team member username to the organizers at indyscc@info.supercomputing.org
- **Hardware**
 - All teams will be given access to the same hardware to give a level playing field
 - Hardware will be a combination CPU and GPU nodes
- **Cloud budgets**
 - Manage your cloud budget! You will not get more of it.
 - You will need to reserve the nodes for Nov 6th and 7th well in-advance

Benefits and Support to the Teams

Top teams will be awarded a **free registration in SC21 that includes the full technical program and workshops** as well as a **complementary one year student membership to SIGHPC**. The winner of IndySCC 21 will receive **bonus points** on their SC22 SCC application hence improving their chance of being accepted.

- Free access to HPC resources from a cloud vendor
- Webinars on the cloud platform, benchmarks, and applications
- Office hours with HPC experts
- Feedback on reports submitted throughout the competition

The Benchmarks

- HPCG (A more typical computational science kernel, tougher for performance)

The Applications

- **GROMACS**

GROMACS is a versatile package to perform molecular dynamics, i.e. simulate the Newtonian equations of motion for systems with hundreds to millions of particles.

- **John the Ripper**

John the Ripper is an open source password cracker available for Unix, macOS, Windows, DOS, BeOS, and OpenVMS. It is mainly designed for the detection of weak Unix passwords and contains several cracking modes.

- **Mystery Application**

Will be related to scientific computing. Having team non-CS members helps!

Each app will have a deep-dive webinar to help you prepare



SC21

St. Louis, MO | science
& beyond.

Questions?