

STUDENT CLUSTER COMPETITION

A MULTI-DISCIPLINARY UNDERGRADUATE HPC EDUCATIONAL TOOL

- ❖ Stephen Lien Harrell, Purdue University
- ❖ Hai Ah Nam, Los Alamos National Laboratory
- ❖ Verónica G. Vergara Larrea, Oak Ridge National Laboratory
- ❖ Kurt Keville, Massachusetts Institute of Technology
- ❖ Dan Kamalic, GNS Healthcare

STUDENT CLUSTER COMPETITION

THE NEED

- Undergraduate programs don't provide for the interdisciplinary nature of HPC
- Computational science is, at best, offered as an elective
- Exposure to parallel computing is only attained through a computer science degree program.
- SCC created as an educational tool to immerse undergraduates in HPC
- The SCC can complement a strong parallel and distributed computing (PDC) curriculum through multi-disciplinary experiential learning and engagement
- The SCC is designed so that teams are aware of the interconnections between system architecture and applications



HPC WORKFORCE NEEDS


MULTI-DISCIPLINARY WORKFORCE

- Science application developers are typically not trained to develop massively parallel applications.
- Design choices are heavily impacted by HPC architecture complexity
- Many different kinds of staff are required for a functioning HPC center
 - Data Scientists
 - Computational Scientists
 - System Engineers
 - System Administrators
 - User Assistance Specialists
- These staff have a direct impact on the quantity and indirect impact on the quality of research done.
- Both researchers and staff benefit greatly from understanding both the science and the technology

HPC WORKFORCE NEEDS

WORKFORCE AVAILABILITY

- Whitehouse executive order for formation of National Strategic Computing Initiative
- 93% of HPC centers have difficulty hiring staff with the requisite skills.
- Current approaches to workforce development are inadequate to address today's needs.
- The pace of HPC growth in industry is outpacing government labs and academia.



STUDENT CLUSTER COMPETITION

THE DETAILS

- Teams of six undergraduates assemble small clusters at the conference and race to complete a real-world HPC workload.
- Teams also are scored on
 - Technical interviews
 - Posters
 - Presentations
- Teams can bring any architecture and hardware
 - Hardware must be able to execute the computation without going over the power limit
 - Power limit is 3120 watts
- Teams run HPL and 4 scientific and/or engineering applications
 - All applications are posted many months in advance save one
 - One mystery application is revealed at the start of the competition



STUDENT CLUSTER COMPETITION

INTENDED IMPACT

- Students experience a microcosm of an HPC center in the competition
- Encourage breadth and multidisciplinary methodology in HPC/PDC education
- Show the impact of HPC to the world we live in and inspire students to pursue HPC as a career
- Students experience the HPC community



STUDENT CLUSTER COMPETITION

STUDENTS EXPERIENCE



- Some comments from a survey sent to the teams after SC13:
 - “Our team concluded that SC12 itself was the BIG experience. That is, coming from a small liberal arts college setting [...] It was the experience of being with the HPC community that made the big difference.”
 - “The Cluster Competition was a great way to learn skills that are not taught whatsoever at school.”
 - “The SCC is a great event for undergraduate students who never have HPC experiences during their undergraduate study... This event guides them into the HPC community.”
- Many SCC alumni go into HPC related jobs and there are a growing number of SCC alumni that participate in the SC conference every year




STUDENT CLUSTER COMPETITION

GEOGRAPHIC DEMOGRAPHICS AND IMPACT TO CURRICULUM

- Trend is becoming more international
- Good representation from US and Asia
- Americas, Australia and Europe has room for growth




- 60% from returning academic institutions have adopted more formal coursework to prepare students for the HPC workforce and, by proxy, the competition.



STUDENT CLUSTER COMPETITION

STUDENT POPULATION DEMOGRAPHICS

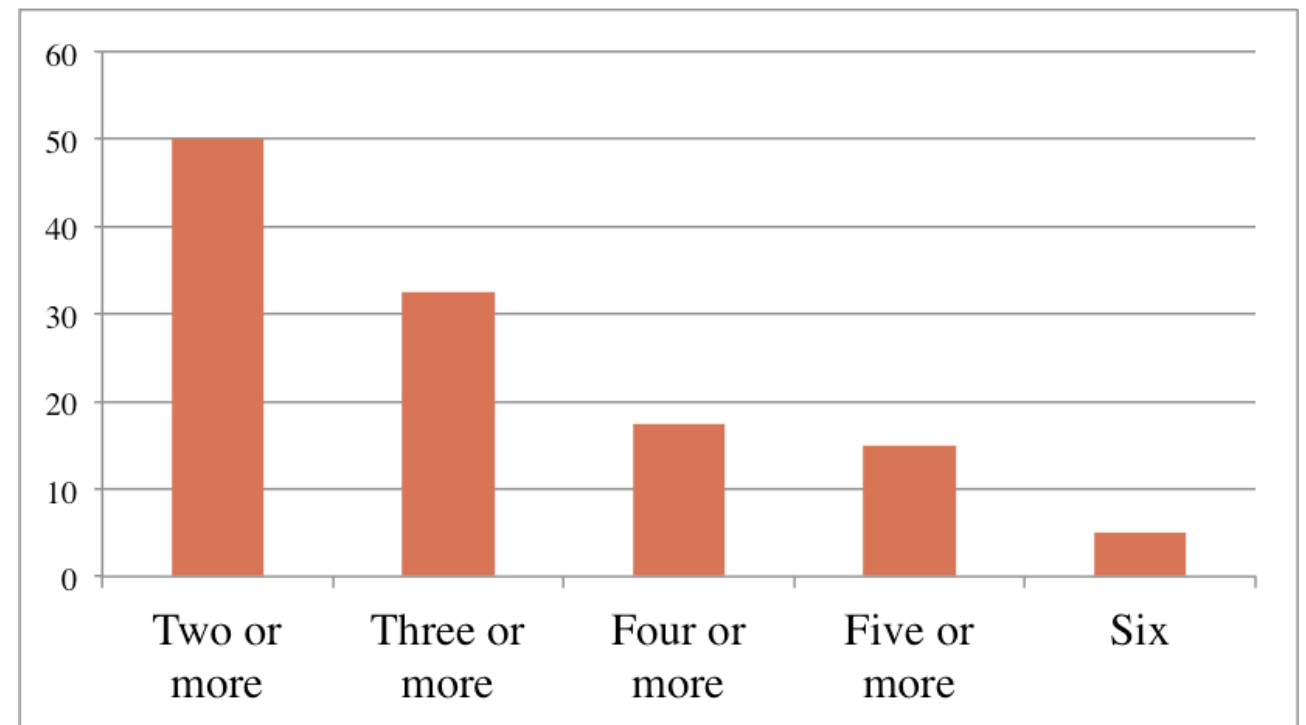
- Multi-disciplinary work is important in the HPC workforce
- SCC give students their first glance at the inter-disciplinary nature of HPC
- 1/6th of students have a non-computing major or an average of one person per team is from a domain science or other major.
- Not only are some domain scientists being exposed to HPC but computing majors are being exposed to domain scientists and their tools.



STUDENT CLUSTER COMPETITION

TEAM APPLICATION RATES


- Continued interest in the competition is a measure of usefulness to the HPC community.
- 57% of teams apply two or more times
- Schools that get a rejection on their first application rarely apply twice



STUDENT CLUSTER COMPETITION

SCHOOL SIZE DEMOGRAPHICS

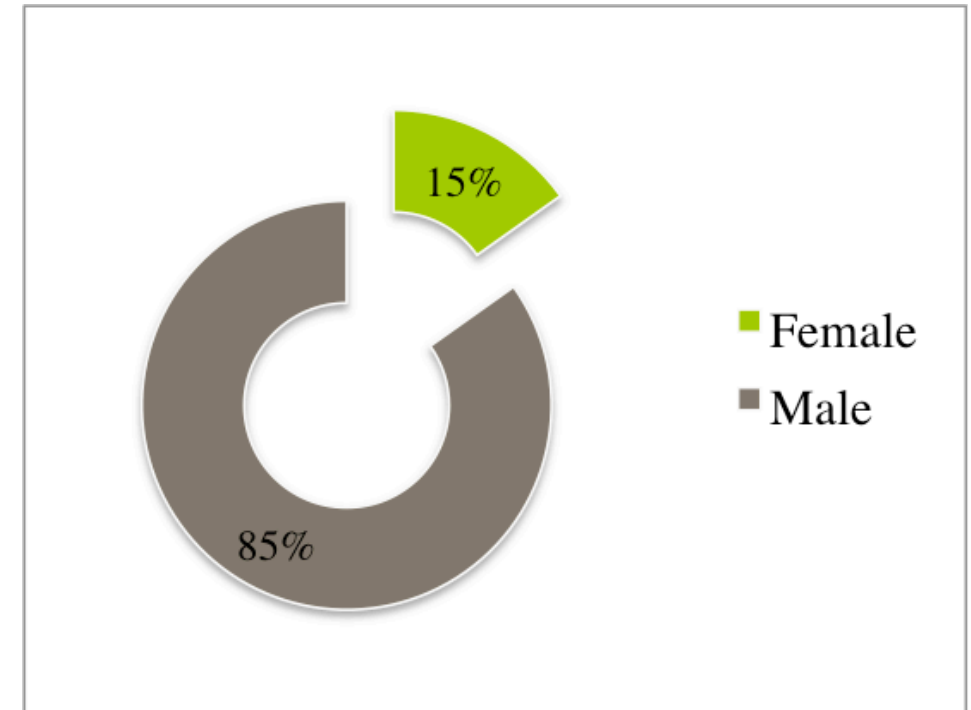
- No schools under 5k accepted
- Weak correlation with underserved institutions
- Only 8% of applications have been from small schools
- Applications were not adequate, however, feedback and support system needed.



STUDENT CLUSTER COMPETITION

WHAT WE CAN DO TO IMPROVE?

- Better advertisement to underserved groups
 - Gender ratio is very low
- Clearer expectations and feedback on judging team applications
- Better follow up with students after the competition. Reinforce key multidisciplinary themes.
- More team matchmaking for smaller schools with vendors and other mentor schools.
- More data: put in place a regular survey that quantifies the student impact long term



THANK YOU

FOR YOUR SUPPORT

- Creators of the competition
 - Brent Gorda
 - Ricky Kendall
- Funding the competition
 - Supercomputing conference series, IEEE/ACM
 - Sponsors and vendors
- This paper supported by research at:
 - Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory
 - Supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725
 - Los Alamos National Laboratory
 - Supported by the U.S. Department of Energy contract DE- FC02-06ER25750.
 - The publication has been assigned the LANL identifier LA-UR-15-27495
 - Purdue University
 - Massachusetts Institute of Technology
 - GNC Healthcare