Team ThreadShredders

Thread Shredders

Arunav Gupta, John Ge, John Li, Kaiwen He, Matthew Mikhailov, Zihao Kong. (Students)
Mary Thomas[1], Bryan Chin[2] (Mentors), Maximillian Apodaca[2], Lewis Carroll[3], Andy Goetz[1], Jacob (Xiaochen) Li[2], Paul Rodriguez[1], Manu Shantharam[1], Robert Sinkovits[1], Mahidhar Tatineni[1] , Paul Yu[4] (Advisors)
[1] San Diego Supercomputer Center, [2] UCSD, [3] AMD, [4] Microsoft Azure

SC21 St. Louis, MO | science & beyond.

Team Members

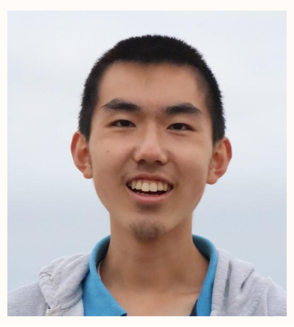

Kaiwen He
Computer Engineering
Skills: System Administration, Software Engineering, System Networking, Computer Architecture
Role: Team lead for IO500, Co-lead for HPL

Arunav Gupta
Data Science & Math
Skills : Machine Learning, Cloud Computing, Statistics
Role: Team lead for ramBLe, Co-lead for HPCG

Zihao Kong
Computer Engineering
Skills: RTL design, Parallel Programming, Deep Learning Accelerators
Role: Team lead for HPL, Co-lead for Cardioid

Matthew Mikhailov
Computer Engineering & Physics
Skills : Computer Architecture, Computational Chemistry, RTL Design
Role : Team lead for Quantum Espresso, Co-lead for IO 500

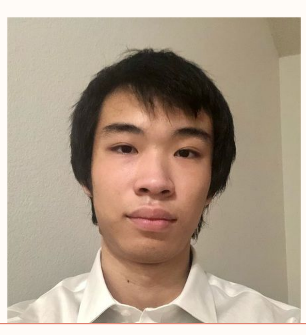

John Ge
Computer Science
Skills: Software Engineering, System Networking, Web Security
Role: Team lead for Cardioid, Co-lead for ramBLe

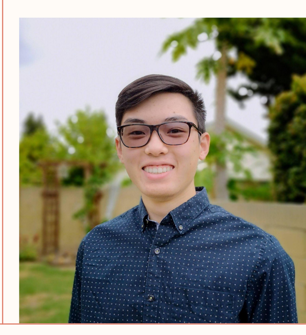

John Li
Mathematics-Computer Science
Skills: System Administration, Software Engineering, Radiation Effects Analysis, Database Design
Role: Team lead for HPCG, Co-lead for Quantum Espresso

Building our Team

While recruiting people for this competition, we reached out to several departments and student organizations to achieve diversity comparable to the UCSD Engineering pool: 2 out of 8 members (6 team + 2 alternates) are from underrepresented groups. Next year, we will reach out to more diverse organizations earlier in the application process.

As shown above, our team is self-motivated, passionate about HPC, and experienced with a broad range of academic backgrounds and technical skills. This allows every team member to learn about all applications and to be the champion of one of the applications.

Application Computational Requirements

The table below shows our planned compute usage for each benchmark and application. We plan on creating a separate Cyclecloud cluster for each benchmark/application with a dedicated environment and dependencies. Estimates are based on prior experience or tests on each application. The total estimated cost for running all applications on Azure is $2,156.

App	Architecture	Cost Analysis
HPL	• **VM Type: Azure ND96asr_v4** • Double precision arithmetic throughput • A100 GPUs and cuBLAS library optimized for dense linear algebra	$33/hr/node * 4 nodes * 3 hrs = $396
HPCG	**VM Type: Azure ND96asr_v4** • cuSPARSE library optimized for sparse linear algebra • A100 GPUs for memory bandwidth • A100 GPUs offer much greater performance here than V100s with comparable cost	$33/hr/node * 4 nodes * 1.5 hrs = $198
IO/500	**VM Type: Azure F4** • Network throughput per dollar • F series show the best performance per dollar for networking • BeeGFS as a high-performance distributed filesystem	$0.219/hr/node * 40 nodes * 4 hrs = $35

Cardioid	**VM Type: Azure ND96asr_v4**
	• Spatial diffusion phase:
	  ○ RDMA for network latency
	  ○ A100 GPUs for memory bandwidth
	• Reaction ODEs:
	  ○ A100 GPUs for computation throughput
		$33/hr/node
		* 2 nodes
		* 5 hrs
		=
		$330

Quantum Espresso	**VM Type: Azure HB120rs_v3** • Dense computation throughput • CPU-based node for better performance per dollar: using GPUs gave a 4x speed improvement, but cost 7x more	$3.96/hr/node * 12 nodes * 5 hours = $238
Bayesian Network	**VM Type: Azure HB120rs_v3** • Paper tests core counts up to 1024, on a 64-node cluster, so low $/core is essential. • We also use Xeon processors and the *-march=native* optimization flag to compile ramBLe	~$51.48/hr * 14 hours = $721
Mystery	**VM Type: Azure HB120rs_v3** • Scheduler with Slurm • Storage backend with BeeGFS • Flexibility on type of VM (Computer, Mem Network)	$3.96/hr/node * 12 nodes * 5 hours = $238

About UC San Diego and San Diego Supercomputing Center

UC San Diego is a large and diverse campus located in the border city of San Diego. The school is known for its research focus and great weather. The nearby beach is a popular attraction and so is the architecture of the campus.

The San Diego Supercomputing Center is a leader in advanced computation (hosting HPC systems such as Expanse and Voyager) and "Big Data", which includes data integration and storage, performance modeling, data mining and predictive analytics, software development, and more.

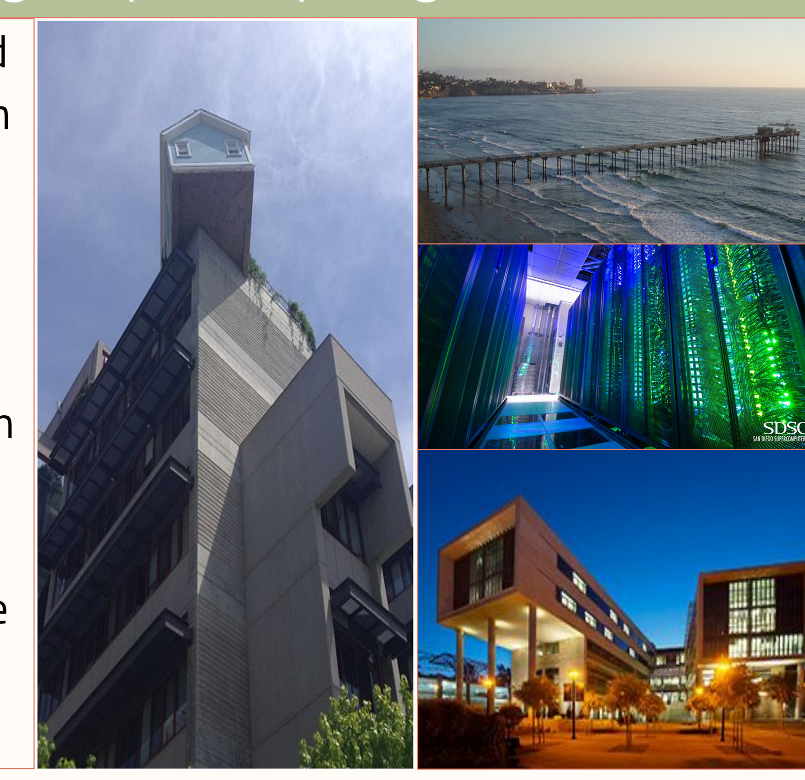

Our Tools

BeeGFS: Filesystem for IO500	CentOS: VM Operating System	Spack: Package Manager	
EPYC: Choice of CPU	InfiniBand: Networking	Chef Progress: Automation Tool	slurm workload manager: Job Scheduler

Competition Preparation and Optimization Strategy

- Attend SDSC HPC User Training classes (or self guided tutorials).
- Participate in weekly meetings with experienced mentors.
- Hands-on experience with SDSC Expanse supercomputer.
- Have a lead and a co-lead be responsible for each application and benchmark.
- Optimize the performance for each application by testing with different VMs on Azure.
- Learn to launch VMs using CycleCloud Templates, CHEF recipes, and Spack.
- Develop tools for monitoring VMs during run-time using CycleCloud, a customized job prediction spreadsheet, and near real-time spend updates.

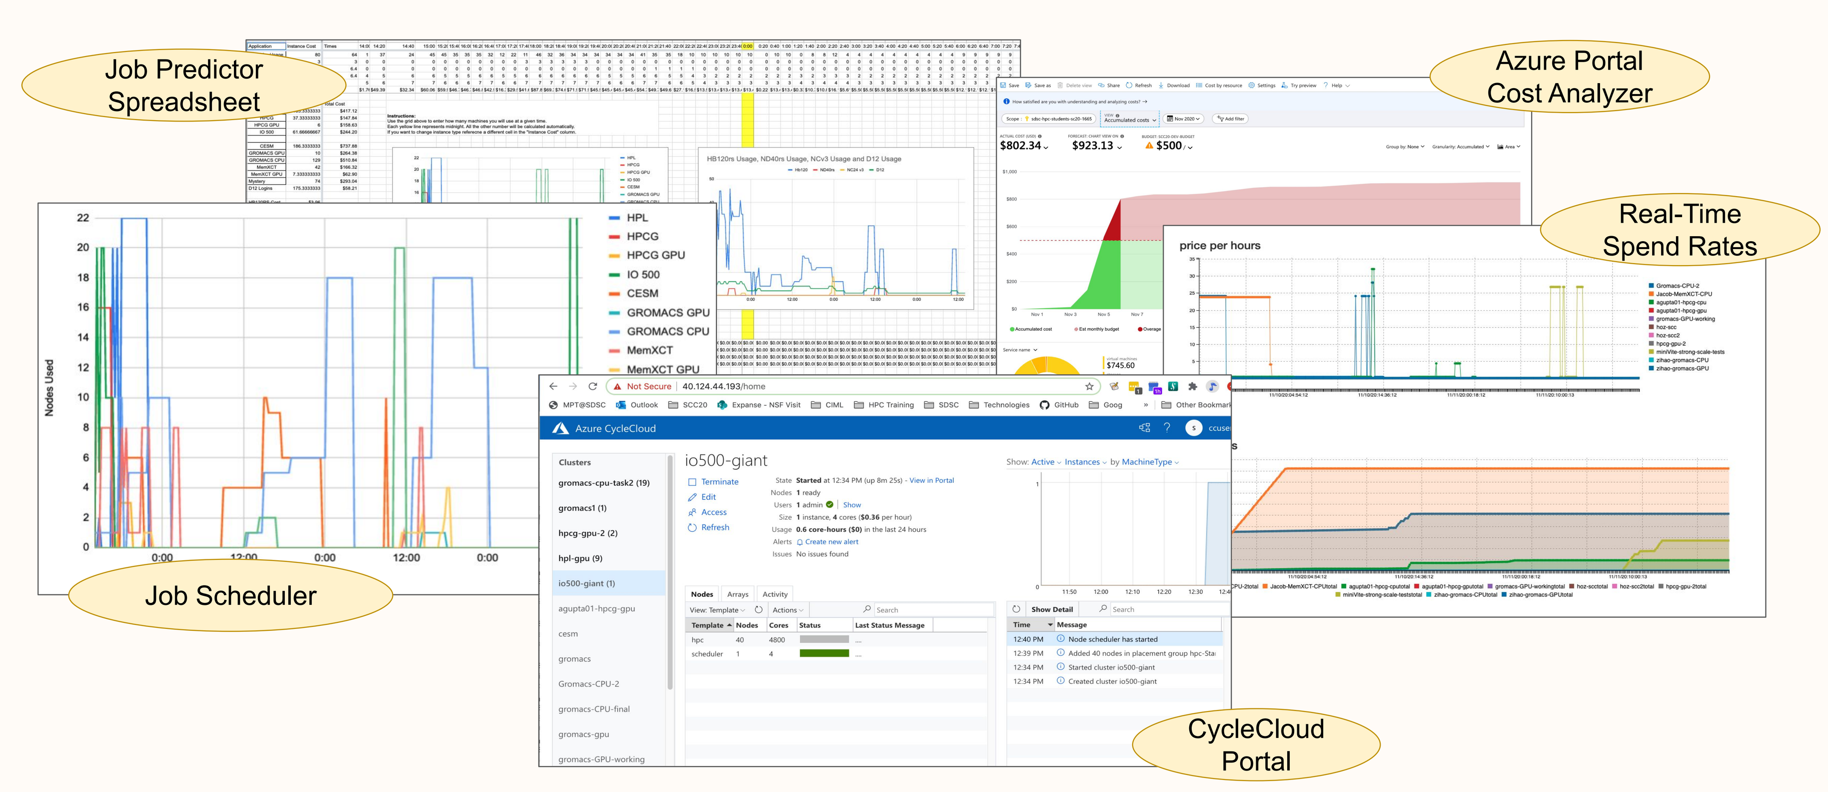

Cluster Deployment & Management

- Cluster Management:
  - Use CycleCloud to manage different clusters for different benchmarks/applications.
  - Use Slurm as our job scheduler for all clusters.
- Budget Management:
  - Obtain CycleCloud cost reports for each cluster on the CycleCloud web browser.
  - Use an Excel Sheet to make plans about budgeting during the competition.
- Cluster Deployment:
  - Develop an HPC Software stack optimized for a CentOS image, with a standard template that installs BeeGFS and SLURM.
  - Develop shell scripts and CHEF recipes to configure the environment for benchmarks and apps.
  - Reduce the VM setup time and increase efficiency in both testing and deployment.

Sponsors

UCSD · SDSC San Diego Supercomputer Center · AMD · NVIDIA · Azure · Dell · amazon · Hewlett Packard Enterprise · NSF