



Indiana University
Bloomington

Members




Lucas Snyder

- Year: Junior
- Major: Intelligent Systems Engineering
- HPL, HPCG, Mystery Application



Zachary Graber

- Year: Senior
- Major: Computer Science
- LAMMPS, Spack, Cloud



Nrushad Joshi

- Year: Junior
- Major: Intelligent Systems Engineering
- HPL, HPCG, LAMMPS



About Our Team

This is a first-time collaboration for the Purdue-IU Student Cluster Competition team; no team members have previously participated in the competition.

Diversity & Inclusion

Diverse teams not only strengthen the variety of skills and expertise at a team's disposal, but also enhance its critical thinking and problem solving capabilities. Our cross-institutional team is comprised of students from diverse educational and cultural backgrounds. One third of our constituency grew up outside of the US, and majors range from various flavors of Computer Science to Aviation Technology.

Purdue University



Members


Larkin Nickle

- Year: Senior
- Major: Computer and Information Technology
- IO500, Storage



Dhruv Sujatha

- Year: Sophomore
- Major: Computer Science & Data Science
- Cloud, PHASTA, Power Delivery, General Hardware

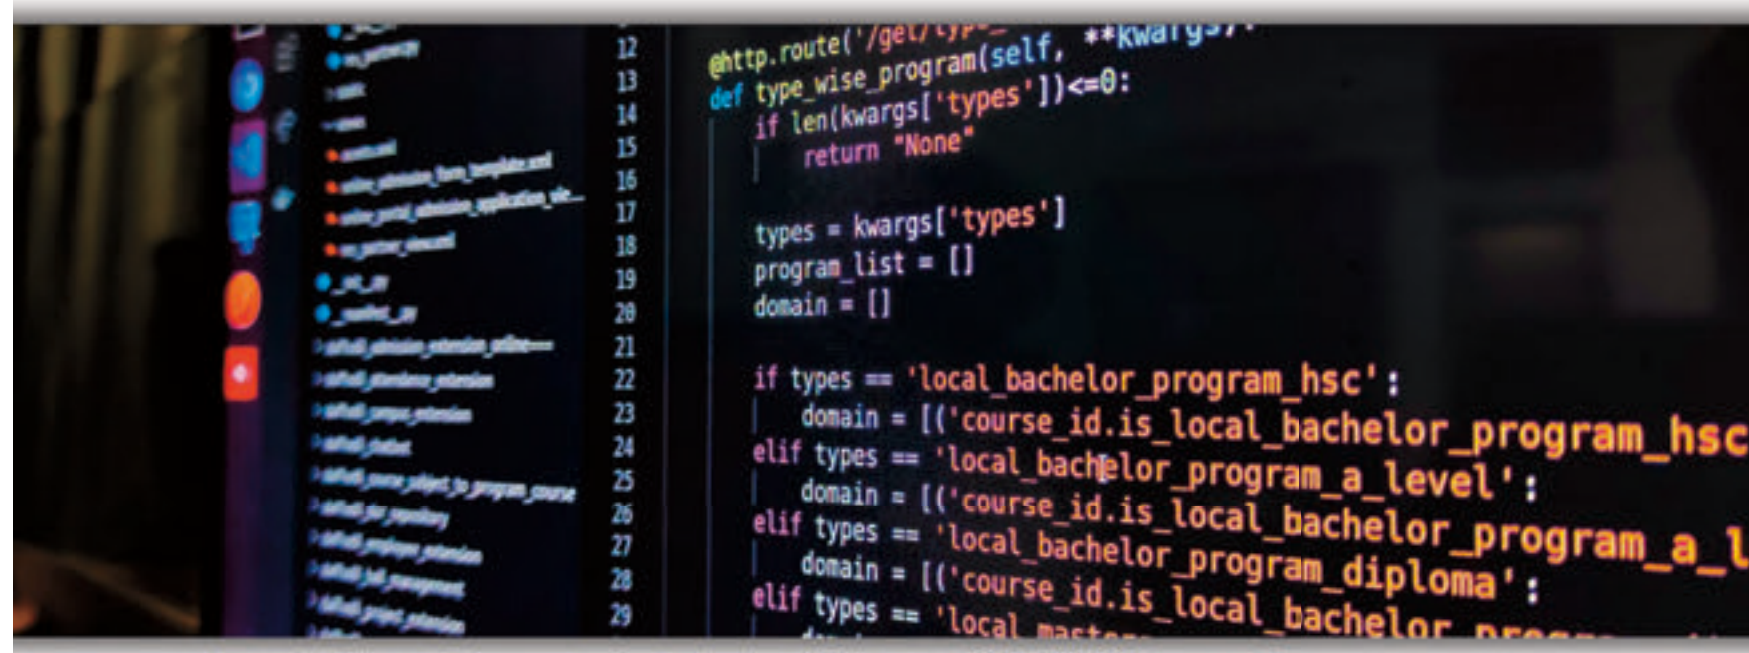


Karl Oversteyns

- Year: Junior
- Major: Unmanned Aerial Systems
- General Hardware, Power Delivery, IO500, Storage



Software



- OS: Rocky Linux 8
- Compilers: gcc or intel
- Libraries: OpenMPI/OpenMP, CUDA, NCCL
- Nvidia Containers: HPC-Benchmark 21.4, LAMMPS
- Monitoring: Netdata, Prometheus, Grafana
- Other: Python, xCAT, Slurm, and Spark

Hardware


- 4 x Dell PowerEdge 7525 chassis, with:
 - 2 x AMD EPYC 7543 @2.8GHz
 - 2 x Nvidia A100 80GB
 - 512GB RAM
 - 4 x 2TB NVMe SSD
 - HDR Infiniband, 200Gbps
 - Mellanox QM8790 Network Switch



Benchmarks & Applications

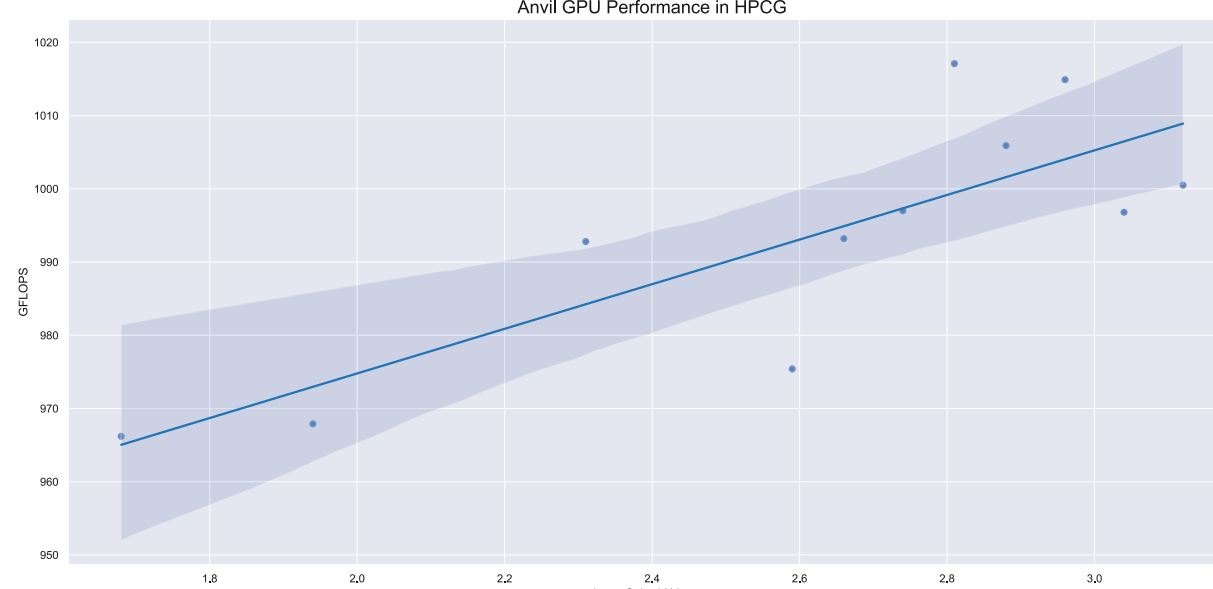
1) HPL

- Classic dense linear systems benchmark measuring
- Find optimal values of input parameters to exploit GPU utilization and its memory



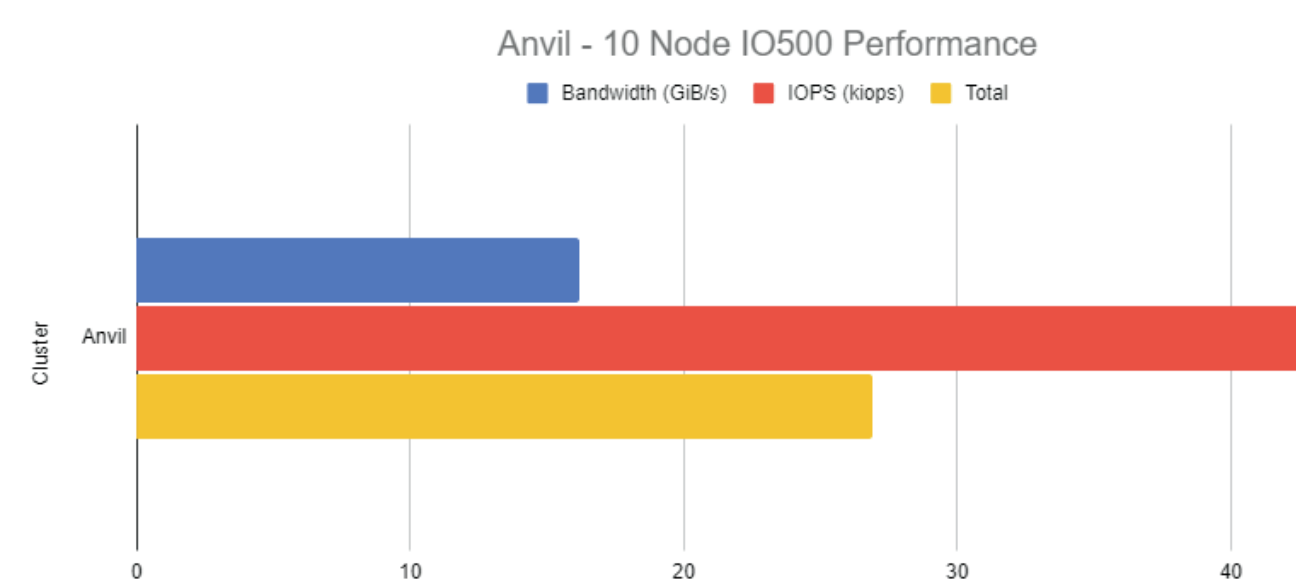
2) HPCG

- A more relevant variation of HPL, it uses the Conjugate Gradient algorithm to solve sparse linear systems.
- Find optimal values of nX, nY, nZ to exploit GPU capability




3) IO500

- Evaluates and measures the storage devices using IOR and MDtest phases
- Utilize B




4) PHASTA

- Performs higher-order, hierarchical transient analysis of compressible and incompressible Navier Stokes equations
- Leverage our 256 AMD CPU cores to accelerate performance



5) LAMMPS

- Molecular Dynamics simulator for solids, liquids, and gases
- Understand the input parameters and dependencies to optimize GPU performance



6) Mystery Application

- Run previous competition's mystery application to prepare ourselves for unexpected challenges
- Randomly select scientific applications and build them on the cluster to get valuable experience with unfamiliar applications




7) Reproducibility Challenge


- Analyze and understand the productivity, portability, performance: data-centric Python research paper
- Learn code optimization techniques in Python
- Attempt to duplicate results presented in the research paper

Teamwork & Preparation

- Self-paced learning with resources from Purdue RCAC and IU IT Training helps our team members become more familiar with topics like Unix, clusters, and the Anvil cluster, encouraging self-guided learning
- Twice-weekly virtual meetings help our team coordinate, overcoming the physical distance between campuses
- Team approach (live + asynchronous) to reproducibility challenge
- In-person meetings and bootcamps led by SMEs facilitate constructive education and deeper team bonding



Power Management



- Use Dell iDRAC power capping along with nvidia-smi GPU power capping
- Able to change power caps while system running through racadm
- Monitoring at the PDU level with SNMP
- Checking from userspace with racadm and ipmitool

Cloud Utilization



- We plan to leverage cloud computing resources in a reserve capacity for a number of uses:
- Running workloads locally and on cloud simultaneously when needed
 - Executing re-prioritization by offloading workloads to cloud, freeing up local resources

Why We Will Win

- Our team was able to prepare for SC22 using top-ranked HPC clusters at each of our institutions
- Multiple members have industry-level work experience in addition to working with the relevant benchmarks and applications, as well as real in-production clusters
- Our advisors have extensive experience from previous cluster competitions

