



FELINEOPS

INDYSCC TEAM
UNIVERSITY OF CALIFORNIA, MERCED

Students: Alex Villa², Alvin Chen², Aaron Rundberget², Aryana Neuman², Tiep Nguyen²
Advisors: Xiaoyi Lu¹, Robert Romero. Mentors: Adam Weingram¹, Yuke Li¹
PADSYS Lab¹, UC Merced², UC Merced OIT³



ABOUT US

FeLineOps UC Merced's high performance computing (HPC) team. Driven by our search for new challenges and growth in competitive environments. With experience in undergraduate research and HPC infrastructure, we are eager to push boundaries, expand our knowledge, and showcase our skills in this competition.



Dr. Xiaoyi Lu

- Associate Professor of Computer Science and Engineering (CSE) at the University of California, Merced
- Research Interests: Distributed and parallel computing, high-performance communication and I/O technologies.
- Director of Parallel & Distributed Systems Laboratory.



Alex Villa

- CSE Major, Applied Mathematics Minor
- Contributed to *FleCSI* at Los Alamos National Laboratory



Aaron Rundberget

- 3rd-year CSE major
- Competed for UC Merced in Western Regional Collegiate Cyber Defense Competition Qualifiers



Alvin Chen

- 4th-year CSE Major, Business Minor
- UCSC Hackathon Winner for Most Technical Implementation



Aryana Nenum

- 4th-year CSE major
- Research Assistant at Parallel and Distributed Systems Lab



Tiep Nguyen

- 4th-year CSE Major, Mathematics Minor
- Research Assistant for Language, Learning, Teaching Lab



Jack Berry

- 4th-year CSE Major
- Specialty: Embedded Systems, Operating Systems

WHAT MAKES US UNIQUE

Our team comprises individuals with diverse backgrounds, each bringing unique abilities and passions to the vast field. We aim to combine these talents, enhancing the potential of every student involved, and by extension, our whole team.

UC Merced is a highly diverse institution focused on providing first-rate education and opportunities to communities around California's Central Valley. UC Merced consistently ranks among the best schools in the nation for social mobility and supports top researchers within the HPC field.

Our team reflects that same spirit of inclusivity and growth. For two-thirds of our members, this will mark their first experience in working with researchers and professionals within the HPC field, embodying UC Merced's mission to expand access to advanced fields of study and foster upward mobility.

TEAM WIDE PREPARATION

As we embark on our second IndySCC journey, we have extensively planned to ensure readiness and support from various partners within our institution. Building on last year's momentum, we are reviewing and improving our scripts, workflows, and pitfalls to ensure our success in the competition.

Summer Preparation

- Sub-teams focus on application specific strategies and demos.
- Independent study of key parallel computing techniques, architectures, concepts: MPI, OpenMP, Infiniband, CUDA, parallelism strategies.
- Hands-on experience with PerfTools, AMD µProf, NVIDIA Nsight to profile bottlenecks and identify measurable speedups.

Academic Year Preparation

- Regular presentations to expand technical breadth, share results, and sharpen communication.
- Build reproducible benchmarks' workflows for automated deployments.
- Run early, large-scale tests on UC Merced's Pinnacles HPC cluster.

APPLICATION SPECIFIC STRATEGIES

In order to maximize readiness, we have each member specializing in a specific application in order to have deep understanding and skills

ECE EMULATOR


ECE is a fast statistical model that reproduces outputs of all climate simulation without running the underlying physics. It learns spatial and temporal patterns including fine-scale fluctuations then generates realistic hourly climate fields on demand.

Approach:

Identify bottlenecks in the emulator and optimize these aspects (i.e. Cholesky factorization and the Spherical Harmonic Transform) by profiling per-phase timings and tuning runtime knobs

Outlook:

Improve GPU scaling for Cholesky and SHT by balancing CPU/GPU workloads, tuning tile sizes, and overlapping computation with communication to achieve better performance on larger datasets. Adding adaptive-rank tuning to improve both speed and accuracy across future climate simulations.



HPL HERO RUN

HPL is a portable linear system solver. Achieving maximum performance on this benchmark is critical to our IndySCC success

Strategies:

Automated machine learning lets us efficiently sample runs to find the best parameters. Large scale experiments on our home cluster, Pinnacles, help us find effective strategies for both parameter and low level tuning.

$$\begin{bmatrix} Op \\ BL \end{bmatrix}^T \times \begin{bmatrix} en \\ AS \end{bmatrix}$$

OUR CLUSTER

Platform & OS

- Ubuntu 22.04 LTS.
- Persistent, Mounted Volumes For Large Datasets And Log Files.
- Manila FS
- Slurm Scheduler

Build System & System Software

- Compilers: GCC 13 And LLVM/Clang 17.
- OpenMP (GCC/Clang) And OpenMPI 4.1.2.
- Math Libs: AMD AOCL And OpenBLAS
- Package Management: Spack
- Profilers
- Perf, AMD MProf, And HPCToolkit



SST TOOLKIT

Focus areas:

Simulating CPU models, memory controllers, interconnect/NoC fabrics.

Design:

Run performance tests (sweeps) on different CPU, memory, and network architectures to compare and analyze their behavior.

Outlook:

Using SST allows us to easily repeat experiments and identify performance trade-offs between system components.



ACKNOWLEDGEMENTS

We have been directly been supported by:

- Parallel and Distributed Systems Lab @ UC Merced
 - NSF Grants: 2340982, 2321123 and 2505106
- Cyberinfrastructure and Research Technologies (CIRT) @ UC Merced
- NSF-MRI Grant: 2019144 - Pinnacles Cluster

