Indy SCC 2022 : Random Access Clemories

CLEMSON UNIVERSITY

Random Access Clemories

DELL Technologies

CLEMSON COMPUTING & INFORMATION TECHNOLOGY

Team Introductions:

High-performance computing (HPC) has long been a priority and topic of research across several departments and colleges at Clemson University. Clemson students and faculty regularly use HPC to revolutionize their field of study. Clemson has assembled a diverse and highly competitive team for this year's Student Cluster Competition under the mentorship of Dr. Jon C. Calhoun.

Cooper Sanders

Major: Computer and Electrical Engineering
Interests: HPCG
NSF REU Student and CUDA programmer. Multiple research internships in supercomputing and artificial intelligence.

Logan Durham

Major: Computer Science
Interests: Computer hardware repair IT support software development on legacy systems and languages

Moises Martinez

Major: Computer Science
Interests: Involved in IT hardware and software support for Clemson. Interested in Cybersecurity and IT management.

David Krasowska

Major: Computer Engineering
Interests: Research in HPC lossy compression with Argonne National Laboratories. Has interest in reduction in memory transfers.

Ethan Gindlesperger

Major: Computer Engineering
Interests: Research in data sampling and interning with Intel. Interested in computer architecture and embedded systems. System power management skills.

Benjamin Schlueter

Major: Computer Engineering
Interests: Involved in artificial intelligence, HPC and quantum computing with Clemson groups. Interested in high performance software optimization.

Software Description:

Software	Versions(s)
OpenMPI	3.1.6, 4.1.1
CUDA	10.2.89, 11
GCC	8.3.1
CMake	3.20.2
FFTW	3.3.8
MPICH	3.3.2
Python	3.6.8

Software Stack	
Rocky Linux	OS
ZFS	File System
OpenPBS	Scheduler

- Our software architecture resembles our Raspberry Pi clusters we built to practice for this competition and that of Clemson's Palmetto Cluster.
- This choice of software allows us to get assistance from CCIT and use systems and methods with which we are familiar.

Power Management Configuration:

- Power management tools will be controlled with **Intel PowerStat/cpufreqset** for CPUs and **nvidia-smi** for GPUs to handle the varying power supply limits. We do not expect both the CPUs and the GPUs to be 100% active at any time.
- Our plan is to schedule when different benchmarks/applications are to be run based on the power budget at the current time.

Strategizes for Running and Optimization:

Benchmarks and Applications:

LINPACK
A dense matrix multiplication benchmark that makes heavy use of GPU and memory resources.

- Highly tunable for a variety of hardware.
- The homogeneous nature of cloud resources allows simpler tuning and faster end results (better optimization of P and Q).
- Utilize high performance GPU nodes with large GPU Memory capacity to reduce data transfer costs and provide the optimal environment.
- Optimize the split of CPU and GPU to have efficient parallelization to not hinder performance.
- Test performance of optimized binaries from Nvidia and Intel against building from standard source.
- Support from CCIT who has run HPL on the Palmetto cluster for TOP500 runs.

NAMD
Dynamics simulation software that simulates large complex systems of atoms and molecules.

- Written in Charm++, the application is CPU scalable, and is memory intensive, but built to do automatic load balancing
- There are GPU implementations in CUDA of a handful of the simulations that NAMD offers
- Two 10-Core CPUs with 64 GB of memory each will give an optimization challenge

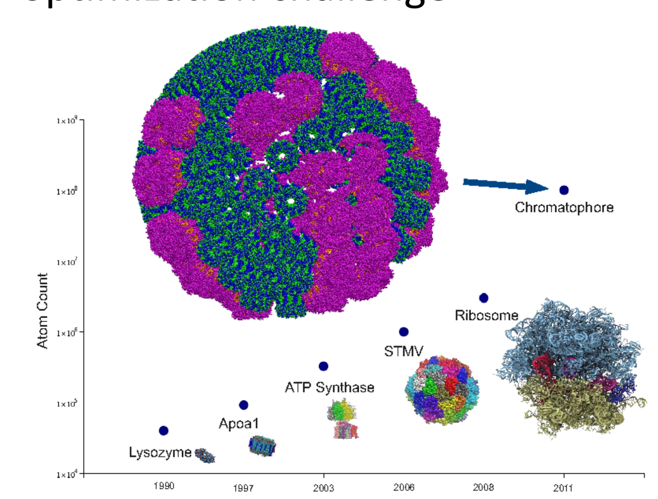

[1]

miniVite
Single phase of Louvain method in distributed memory for graph community detection.

- Written in C++, MPI, and OpenMP making the implementation scalable across CPU cores
- Uses the Louvian method to pull communities from large networks of graphs
- Memory intensive, which will require using varying problem sizes for the best results using 64GB in Chameleon

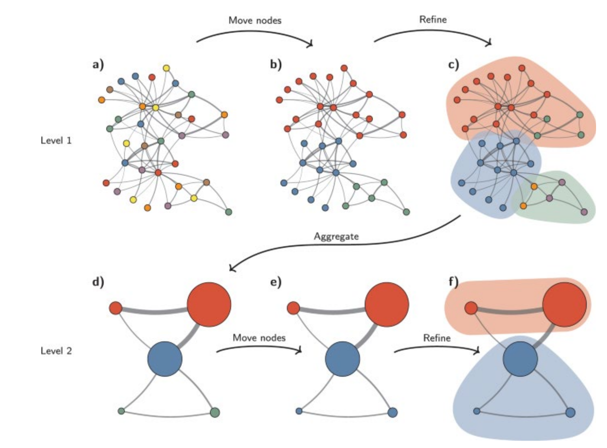

[2]

Team Preparation

- Attend weekly webinars to learn more about the applications we are running and become more familiar with the cluster.
- Test the competition applications in different configurations to obtain the best configuration for running.
- Hold weekly meetings to share our progress on application optimization and competition strategy.
- Utilize Clemson's Palmetto cluster to quickly iterate and test applications and environments.

Why Will We Win?

- We have a diverse group with complimentary experience in fields such as HPC and cloud computing.
- Our application leads have backgrounds suitable for handling the unique aspects of each competition application.
- As a first time IndySCC team, we are prepared to be "All In" (a Clemson motto) and bring our best skills to win.
- We will use our power plan to effectively control our cluster at all points in the competition.

[1] http://charm.cs.illinois.edu/research/moldyn [2] https://www.nature.com/articles/s41598-019-41695-z