

1. About Team Kent Ridge

- **1st-time participation for National University of Singapore in IndySCC.**
- Team Kent Ridge (TKR), like 'Bletchley Park', is a metonymy that references the location of our home campus and our talented team of passionate students.

2. Our Members



Way Yan is a CS and Math freshman. His experience in Math and C programming is key to understanding and optimizing HPL.



Akash, a CS senior, is skilled in Kernel/OS development, and has participated in DEFCON and the Intl. Cybersecurity Championship.



Himaya is an Electrical Engineering junior. A microcontroller expert, she designed a self-balancing robot in the Robomasters competition.



Chien Hao, a CS and Physics freshman, has previously worked on a theoretical condensed matter project at the Institute of High-Performance Computing.



Tze Han, a CS senior, is experienced with inter-node communications and has recently participated in Robonation's RobotX and Robosub competitions



Jiefan, aka **Jeff**, is a CS senior who is skilled in sysadmin and taught several machine learning and parallel programming courses.

And not forgetting our advisor, Dr. **Cristina Carbutaru**, our instructor, **Sriram Sami**, and the many other tireless individuals of the CS3210 (Parallel Computing) teaching team who have graciously supported us.



3. Our Strats

High-Performance Linpack

IC: Himaya

- Parameter (PxQ, NB) tuning to balance memory usage with message-passing.
- Parameter (N) tuning to increase computation/communication fraction, up to memory limits (Gustafson's Law).
- Tuning panel factorization, broadcast, lookahead and swap algo. parameters.
- Evaluating Nvidia HPL and AOCL-libFLAME binaries, as well as OpenBLAS and AOCL-BLIS implementations.

NANoscale Molecular Dynamics

ICs: Way Yan, Jeff

- Disable checkpointing (writing to restart, dcdfreq and xstFreq files).
- Parameter tuning (outputEnergies) to reduce extraneous energy computations.
- Charm++ SMP parameter tuning for CPU affinity (+pemap, +commap, +setcpuaffinity).
- (GPU-offload) Tuning CPU-core-to-GPU ratio.

MLPerf Inference

ICs: Akash, Chien Hao

- Evaluate use of half-precision floating point for improved efficiency.
- Evaluate C++, CUDA and PyTorch implementations, as well as different runtime backends.
- Utilize CM experiment automation script to tune hyperparameters.
- Evaluate pruned and quantized Stable Diffusion models for gains within accuracy tolerances.

ICOsahedral Nonhydrostatic Climate Model

IC: Tze Han

- Parameter tuning (nproma) to optimize contiguous memory utilization.
- Parameter tuning (mpi_oce_nodes) to balance node allocation between ocean and weather simulation.
- Parameter tuning (v0snow) to adjust the terminal velocity of snow.
- Parameter tuning (rain_n0fac) to adjust intercept parameter of raindrop size distribution.

General Optimizations

- Recompile to support vector instructions
- Configure RDMA over Converged Ethernet (RoCE).
- Disable SMT/HyperThreading.
- Disable unnecessary daemons.
- Tune MPI process affinity and placement (etc --distribution).

4. Our Stack



5. Our Prep

- Many, many weeks of meetings :)
- Attending many HPC User Group Sessions and workshops by the National SuperComputing Centre (NSCC), learning to optimize and run jobs on the Aspire2A petascale cluster.
- Training exercises and dry-runs on CS3210 and SoC GPU Compute clusters.
- Continuous literature reviews.
- Preparatory coursework: CS2106 *Introduction to Operating Systems*, CS3210 *Parallel Computing*.



6. Why We Will Win

Simply put, our collective skills, experiences, cohesion, methods, and dedication will put TKR ahead of the competition at IndySCC.