Evaluating the Power Monitoring Capabilities of Aurora

Precious Eyabi[†] , Kazutomo Yoshii[‡] , Xingfu Wu[‡] , Valerie Taylor[‡] , Jon C. Calhoun[†]

[†] Clemson University
[‡] Argonne National Lab

Background

Aurora Exascale Supercomputer

One of the world's first exascale computers.

Architecture	~9,000 nodes 2 Intel Xeon Max CPUs 6 Intel Max series GPUs
Peak Performance	1.012 exaFLOPS[1]
Power	38.7 MW[1]

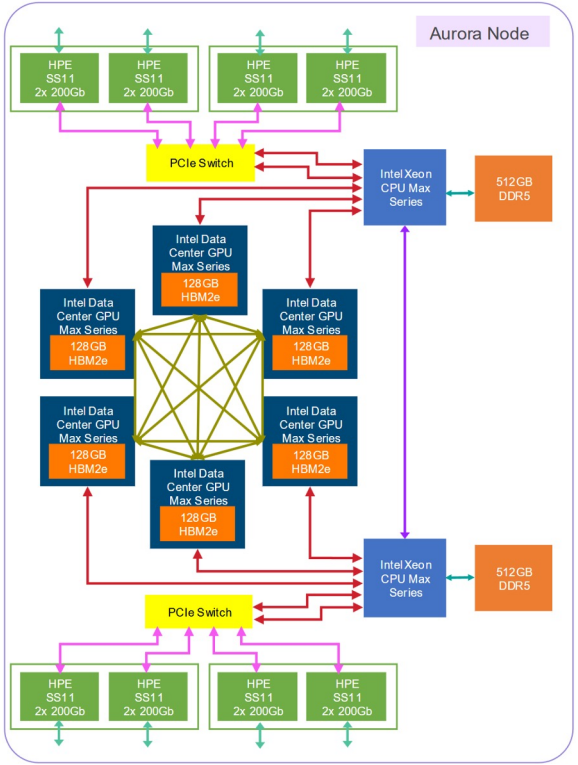

Layout of a single Aurora node[2]

[1] https://en.wikipedia.org/wiki/Aurora_(supercomputer
[2] https://www.alcf.anl.gov/sites/default/files/2024-11/Overview-of-Aurora-Oct-2024.pdf

Motivation

Motivated by Three Main Factors

Power is a limited resource

Tools

GEOPM and APMIDG

Designed for exascale systems

Utilize an Intel Running Average Power Limit (RAPL) backend

GEOPM

Global Extensible Power Manager[1]

- Open-source framework for power and performance management on HPC systems
- Provides runtime power monitoring and control across nodes and CPUs
- Uses agents to automatically adjust power limits based on workload behavior

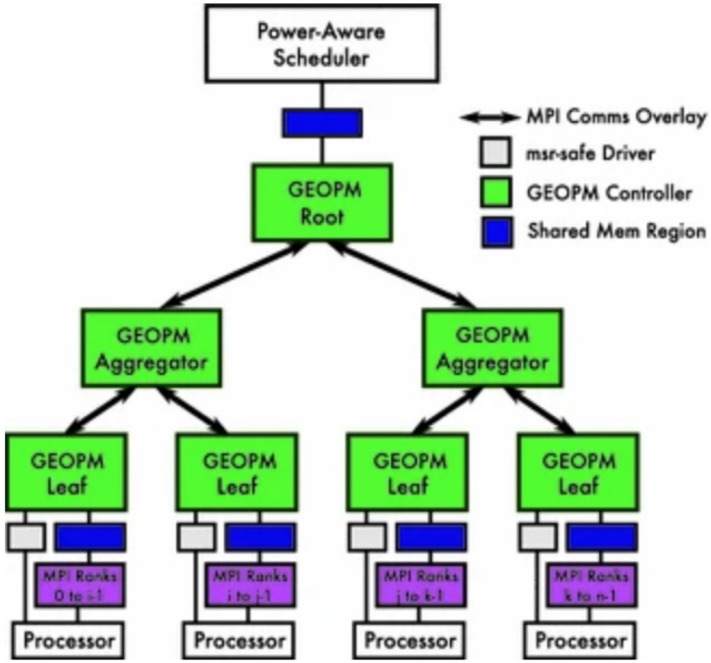

Architecture of GEOPM[1]

[1] https://link.springer.com/chapter/10.1007/978-3-319-58667-0_21

APMIDG

Argo Power Management Glue Layer for Intel Discrete GPUs[1]

- Middleware layer enabling power monitoring and control for Intel discrete GPUs
- Provides a unified interface to access GPU power, frequency, and telemetry data

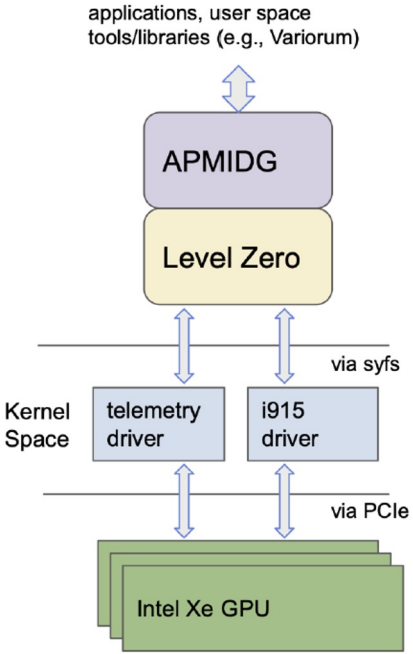

Architecture of APMIDG[1]

[1] https://github.com/anlsys/apmidg

Analysis

Idle Power Characterization

Measured the idle power of a node to set a baseline for concurrent measurements

Measured the idle temperature to examine the correlation

Idle Power Characterization

Total GPU and CPU Power Usage When Idle

- GeoPM and APMIDG report comparable average GPU power levels
- GeoPM exhibits greater variability at higher frequencies
- GPU power patterns appear irregular

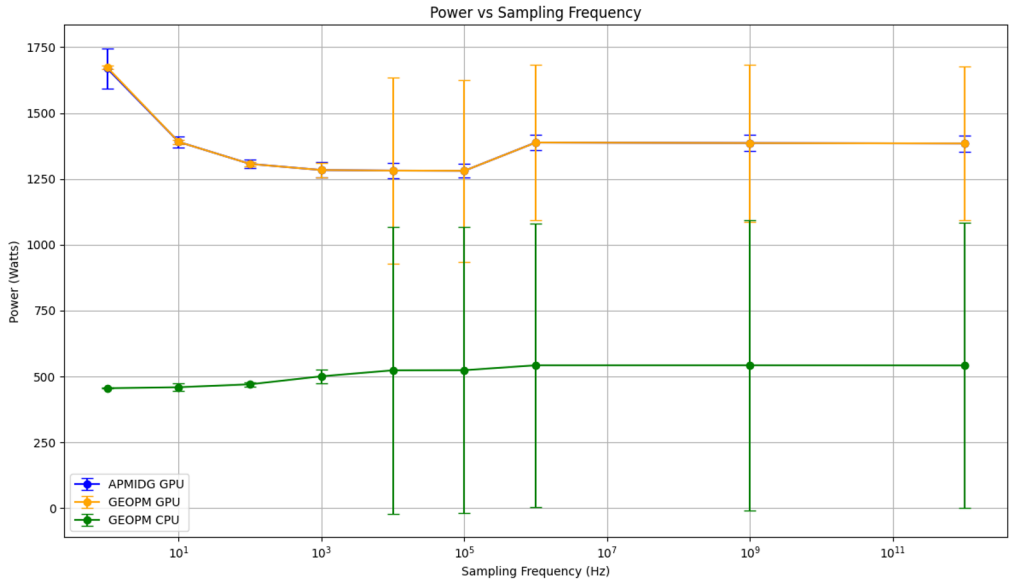

Idle Power

Individual GPU Power Usage When Idle

- Individual GPU power as the sampling frequency increases
- Possible explanation for unexpected behaviors in power reading

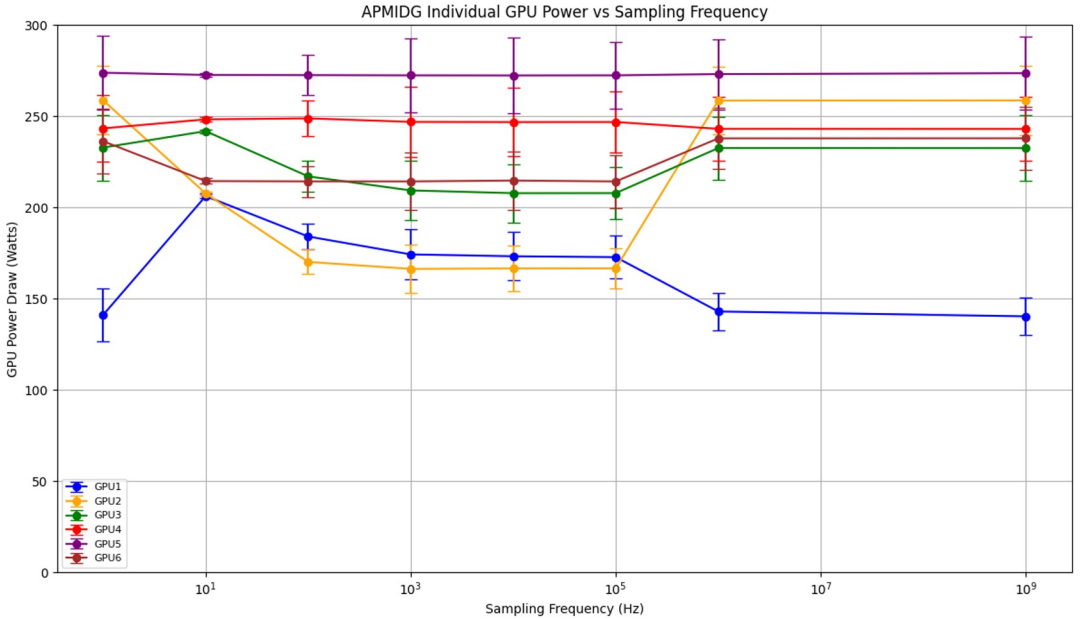

Idle Power Across Multiple Trials

Total GPU Power After Multiple Trials

- Seven trials performed per node
- Similar trend observed in initial run

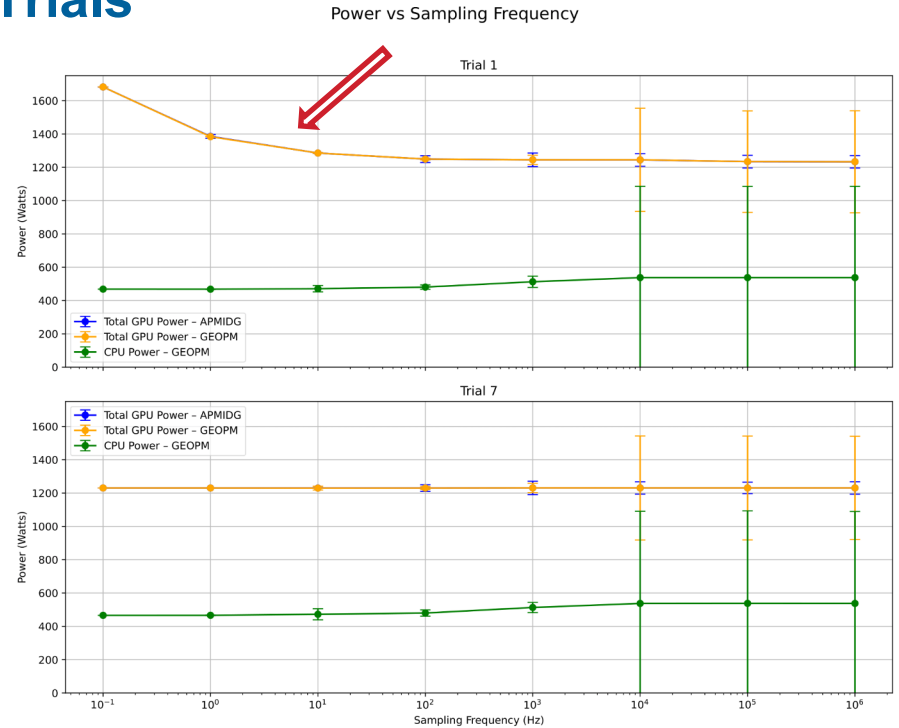

Individual GPU Idle Power Across Multiple Trials

Individual GPU Power After Multiple Trials

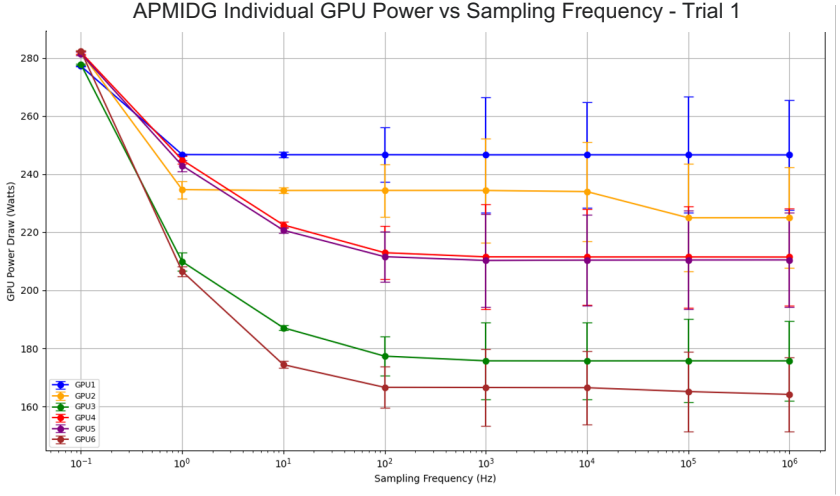

Individual GPU power after multiple trials

Idle Temperature

CPU Temperature

- Displays expected trend of increasing as workload increases

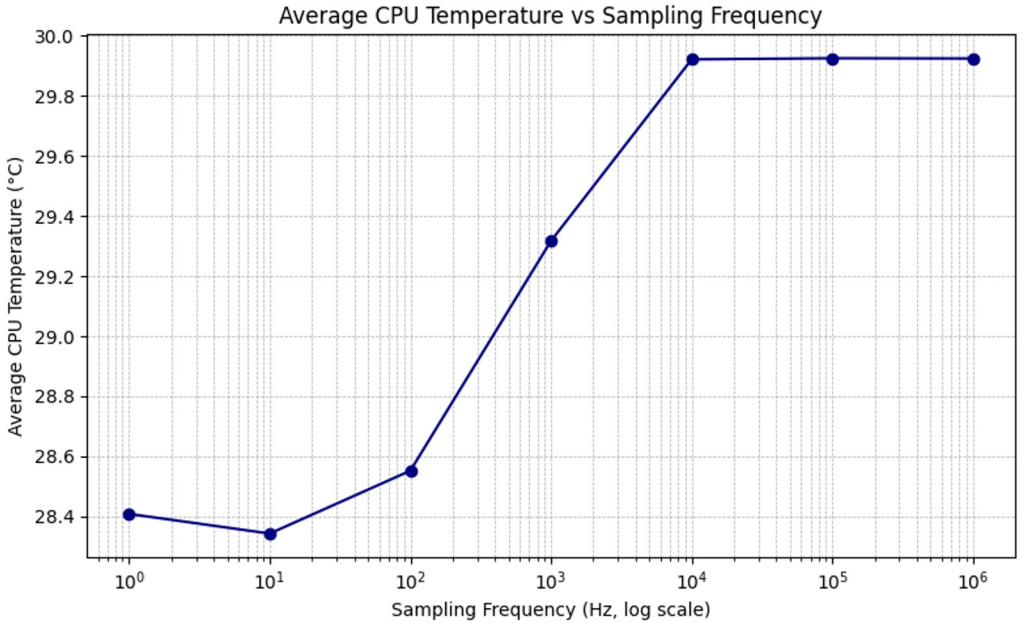

Idle Temperature

GPU Temperature

GPU temperatures as the sampling frequency increases

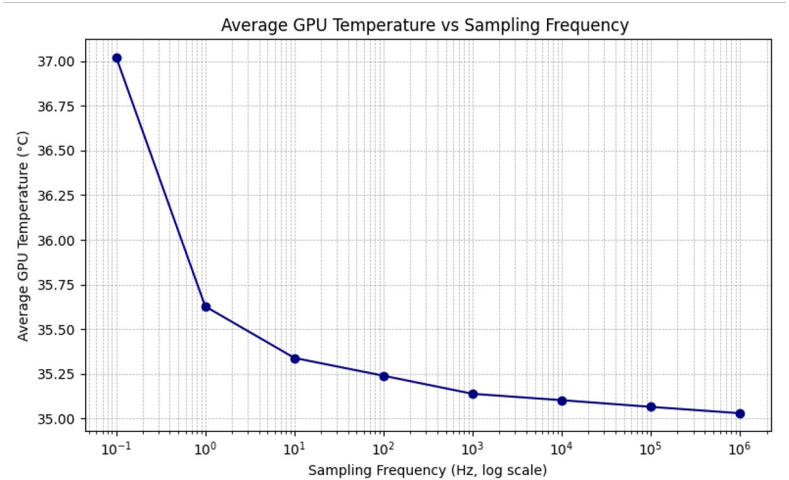

Average of all 6 GPU’s

Analysis

Accuracy Test

Applied a Fast Fourier Transform (FFT) to determine the accuracy of each tool

Accuracy Analysis

Matrix Multiplication Benchmark

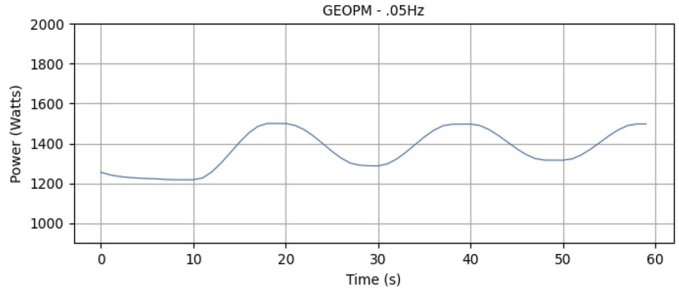

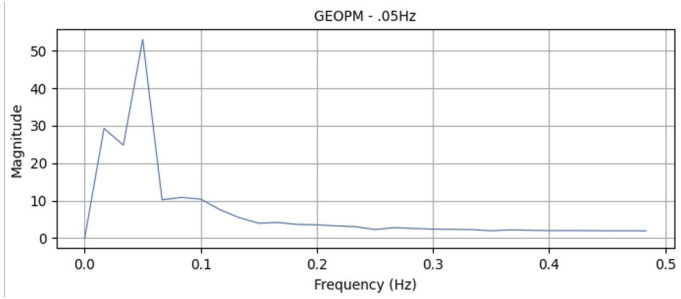

Reflects APMIDG's low-level access to the power API directly

Accuracy Analysis

Matrix Multiplication Benchmark

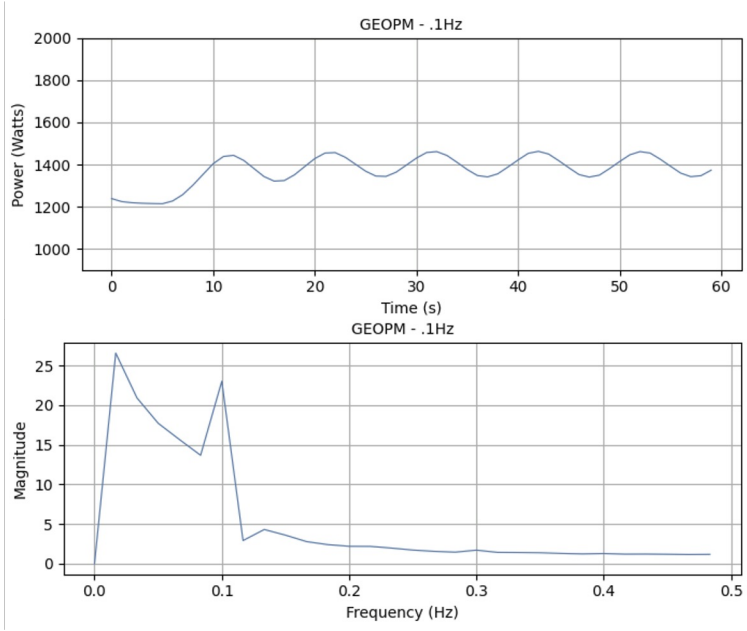

GEOPM exhibits spectral artifacts that intensify at higher frequencies

Analysis

HPC Benchmarks

Ran two GPU heavy HPC benchmarks

Benchmarks used were NekRS and XSBench

Workloads that are regularly utilized at Argonne

Benchmark

XSBench[1] - Neutron Transport Kernel

- Uses Single GPU
- Displays similar trend of internal filter with GeoPM

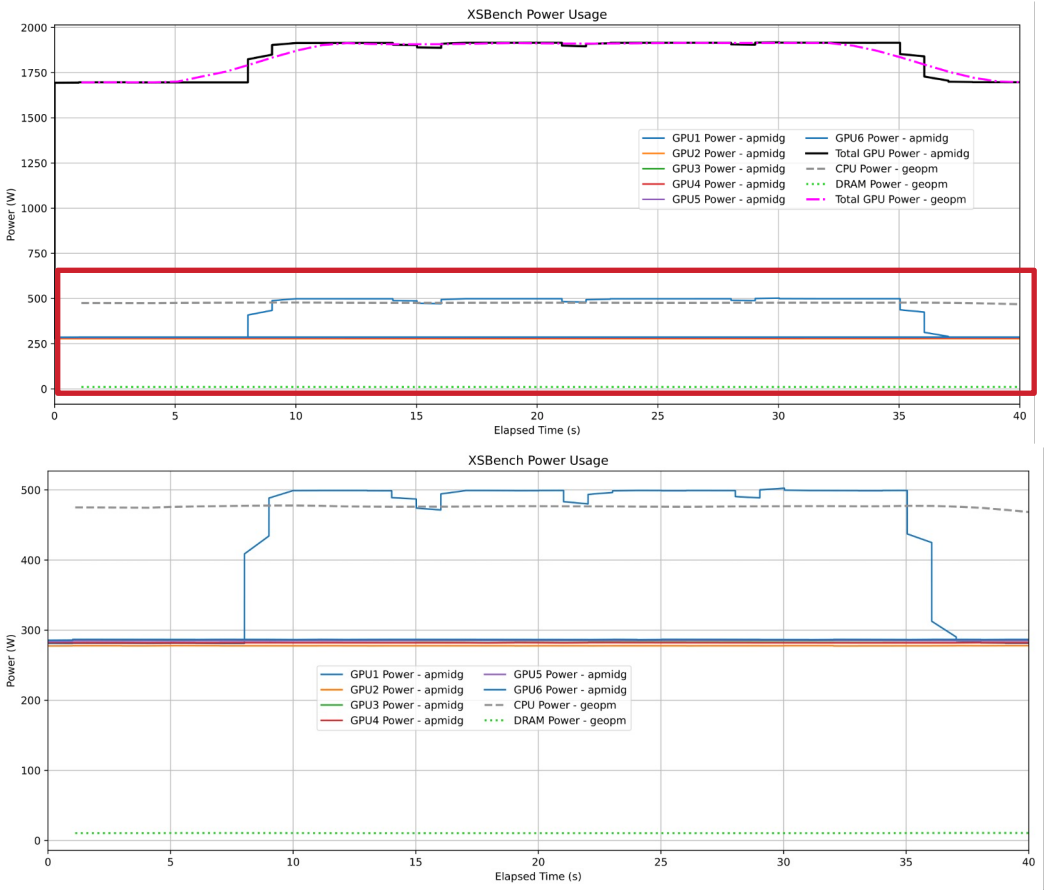

[1] https://github.com/ANL-CESAR/XSBench

Benchmark

NekRS[1] - Computational Fluid Dynamics Solver

- APMIDG can capture high frequency events
- GeoPM monitoring results contain a delay and look as if it has been low-pass filtered

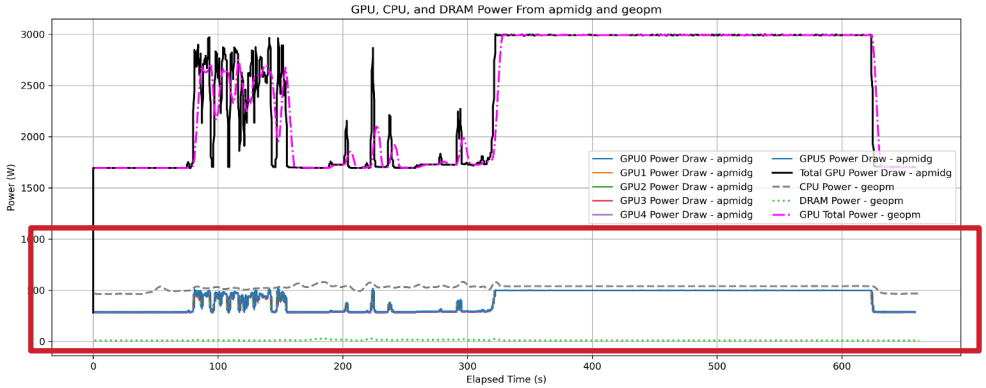

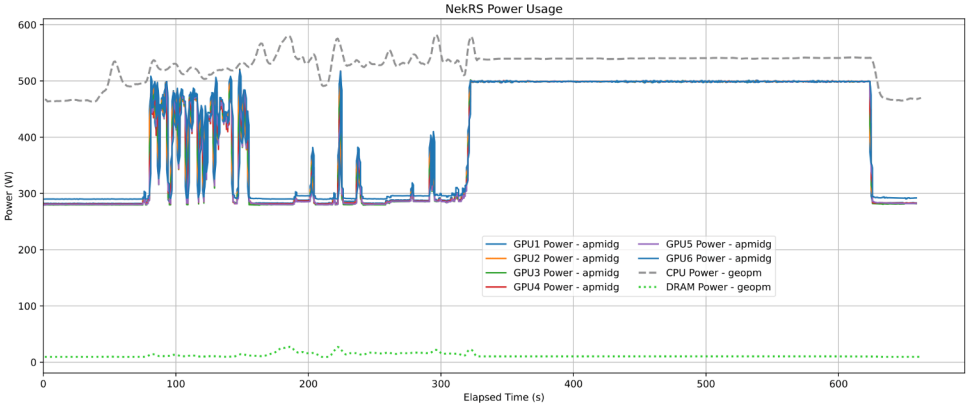

[1] https://github.com/Nek5000/nekRS

Summary

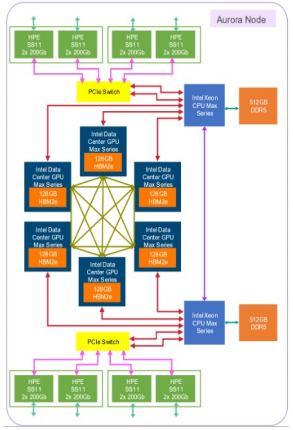

Aurora Node

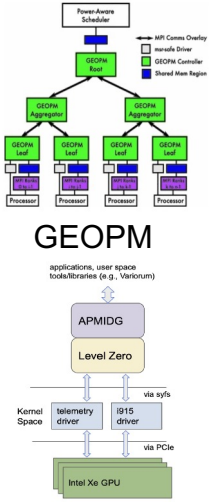

GEOPM

APMIDG

3 Stage Analysis

- Baseline Measurements
- Accuracy Test
- HPC Benchmarks

GEOPM

- Internal averaging/low pass filtering Reduce fine scale accuracy
- Better suited for system level monitoring

APMIDG

- Low level, per device access to power
- Better suited for fine grained, high resolution power profiling and debugging

Next steps

- Multi-node power monitoring
  - Examine the behavior of these tools as they monitor multiple nodes in tandem
  - Examine behaviors across different queues
- Measure multi-node applications
- Investigate the portability of these tools to other exascale systems

Questions?

Acknowledgements
This work was supported in part by DOE ASCR Base PIONEER (Program for Innovative Optimization of Next-gen Energy Efficient Resources) project, and in part by the Office of Advanced Scientific Computing Research's Computer Science Competitive Portfolios program under Contract No. DE-AC02-05CH11231.
We acknowledge the Argonne Leadership Computing Facility for use of Aurora under the project EE-ECP, and Misun Min and YuHsiang Lan for providing the NekRS baseline example on Aurora. This material is based upon work in part supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357, at Argonne National Laboratory, and in part supported by the U.S. DOE ASCR, under contract number DE-SC0026344.