OAK RIDGE National Laboratory

Fine-Grained Application Energy and Power Measurements on the Frontier Exascale System

Oscar Hernandez, **Wael Elwasif**
Oak Ridge National Laboratory

Cray User Group Meeting (CUG 2025)

U.S. DEPARTMENT OF ENERGY

ORNL IS MANAGED BY UT-BATTELLE LLC
FOR THE US DEPARTMENT OF ENERGY

Motivation

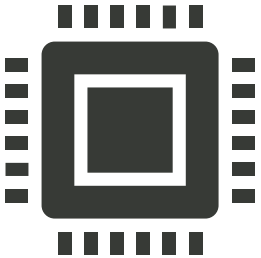

HPC Power Challenges

Increasing power consumption in modern chips (e.g., **AMD MI250X**)

Need for fine-grained power and energy optimizations

Moore’s Law slowdown and efficiency trade-offs

Why Fine-Grained Power Analysis?

Identifying power spikes and inefficiencies

Optimizing application performance and energy use

Key Contributions

- Development of an open-source toolkit for fine-grained power measurements
- Integration of Linux hwmon, Cray PM, and ROCm SMI with PAPI and Score-P
- High-resolution power sampling (1 ms)
- Mapping power usage to specific application regions

Key Toolkit Components

- Scalable Performance Measurement Infrastructure for Parallel Codes (Score-P)
  https://www.vi-hps.org/projects/score-p/
- Performance Application Programming Interface (PAPI)
  https://icl.utk.edu/papi/
- The Cray Power Management data access for users (PM)
  https://cray-hpe.github.io/docs-csm/en-10/operations/power_management/user_access_to_compute_node_power_data/
- The AMD ROCm system management interface library (ROCm-smi)
  https://github.com/ROCm/rocm_smi_lib
- Vampir performance analysis and visualization
  https://vampir.eu/

Performance Analysis Toolkit for Frontier

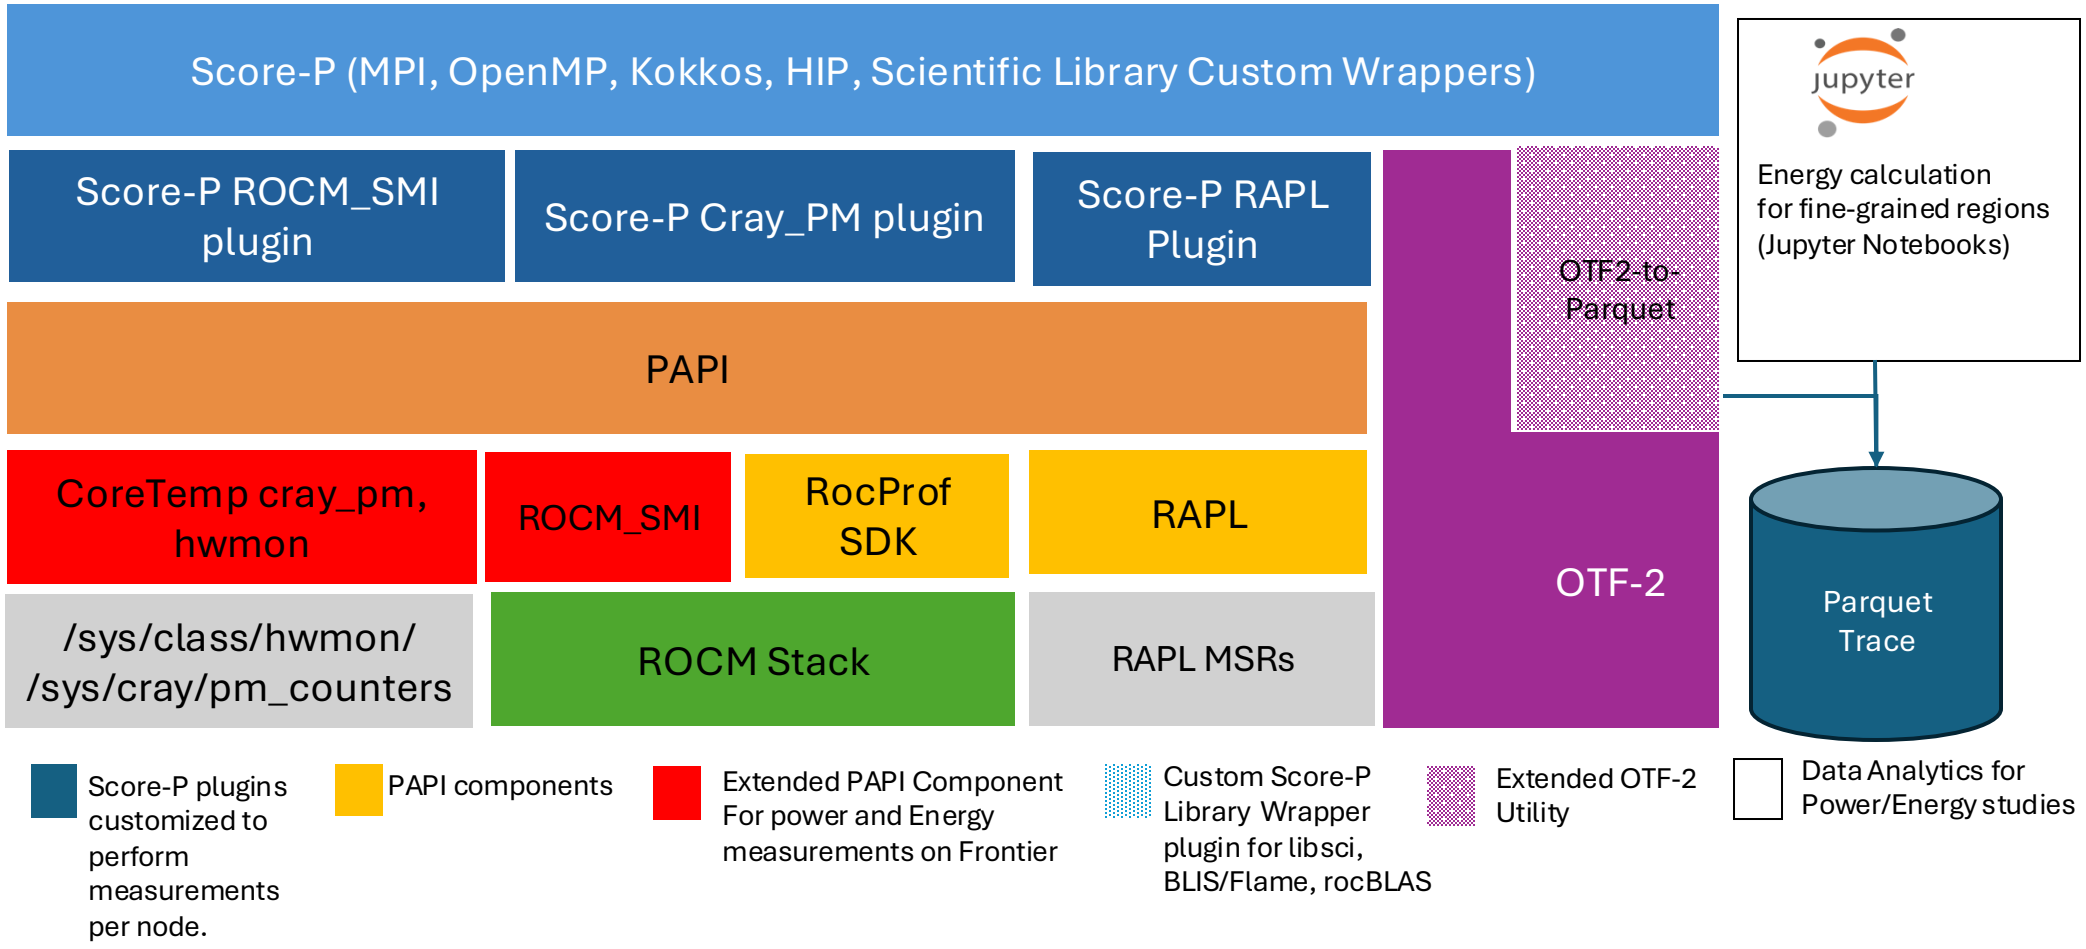

Based on tools from:

Hewlett Packard Enterprise · AMD · VAMPIR · Score-P Scalable performance measurement infrastructure for parallel codes · icl-utk-edu/papi

Performance Analysis Toolkit for GH200

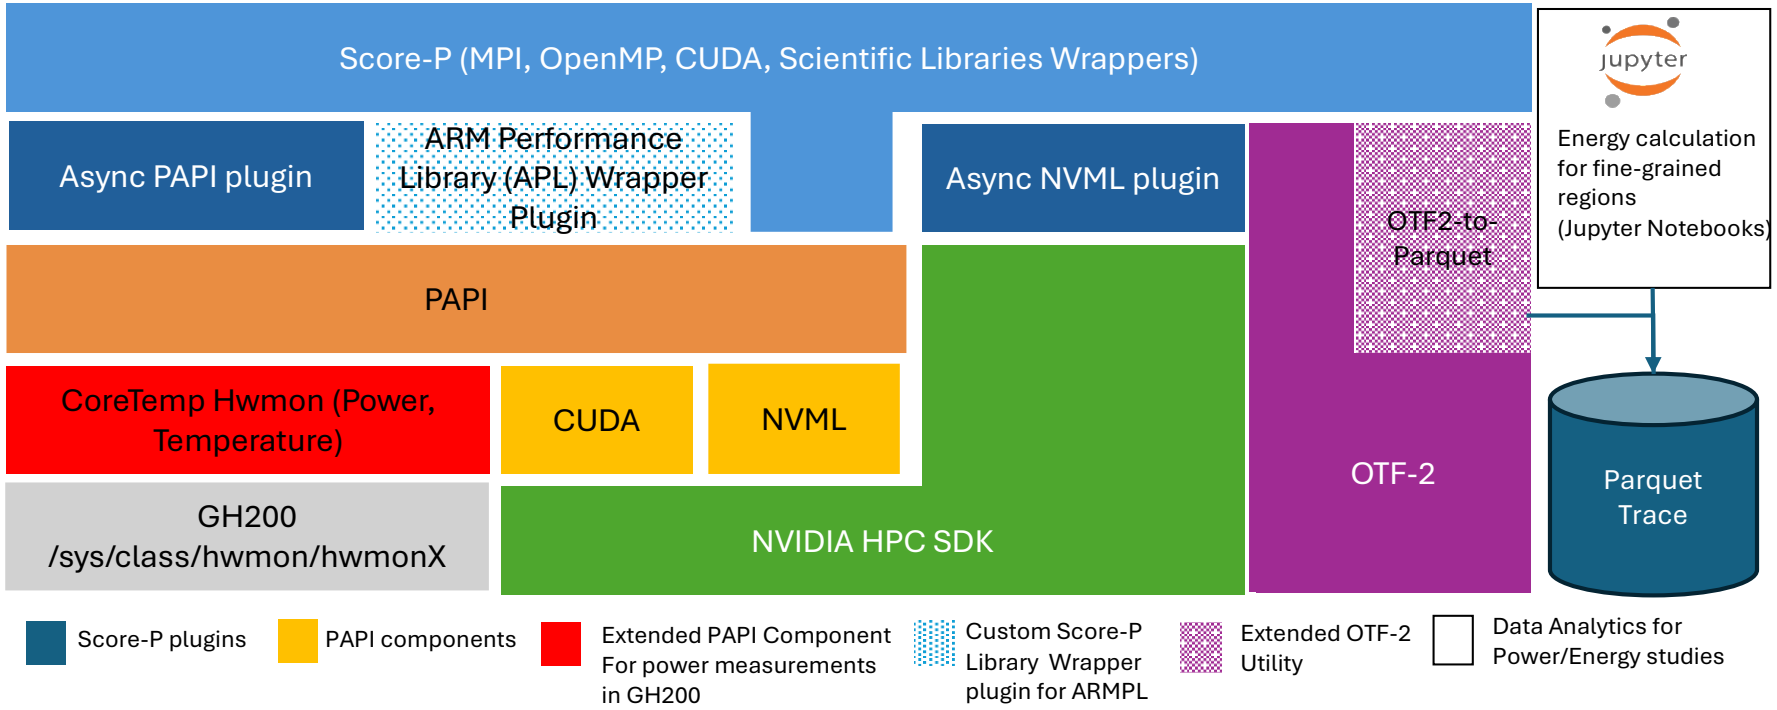

Python Data Analytics

Extended Components

- PAPI extension for hwmon to measure Superchip and Grace CPU power
- Score-P NVML and PAPI plugins readings
- OTF2-to-Parquet conversion utility
- Jupyter-based analytics for energy profiling

Supported Architectures

- GH200, potential for AMD Instinct and other platforms

Based on tools from:

VAMPIR

Score-P
Scalable performance measurement infrastructure for parallel codes

icl-utk-edu/papi

NVIDIA

arm

Performance Analysis Toolkit for GH200

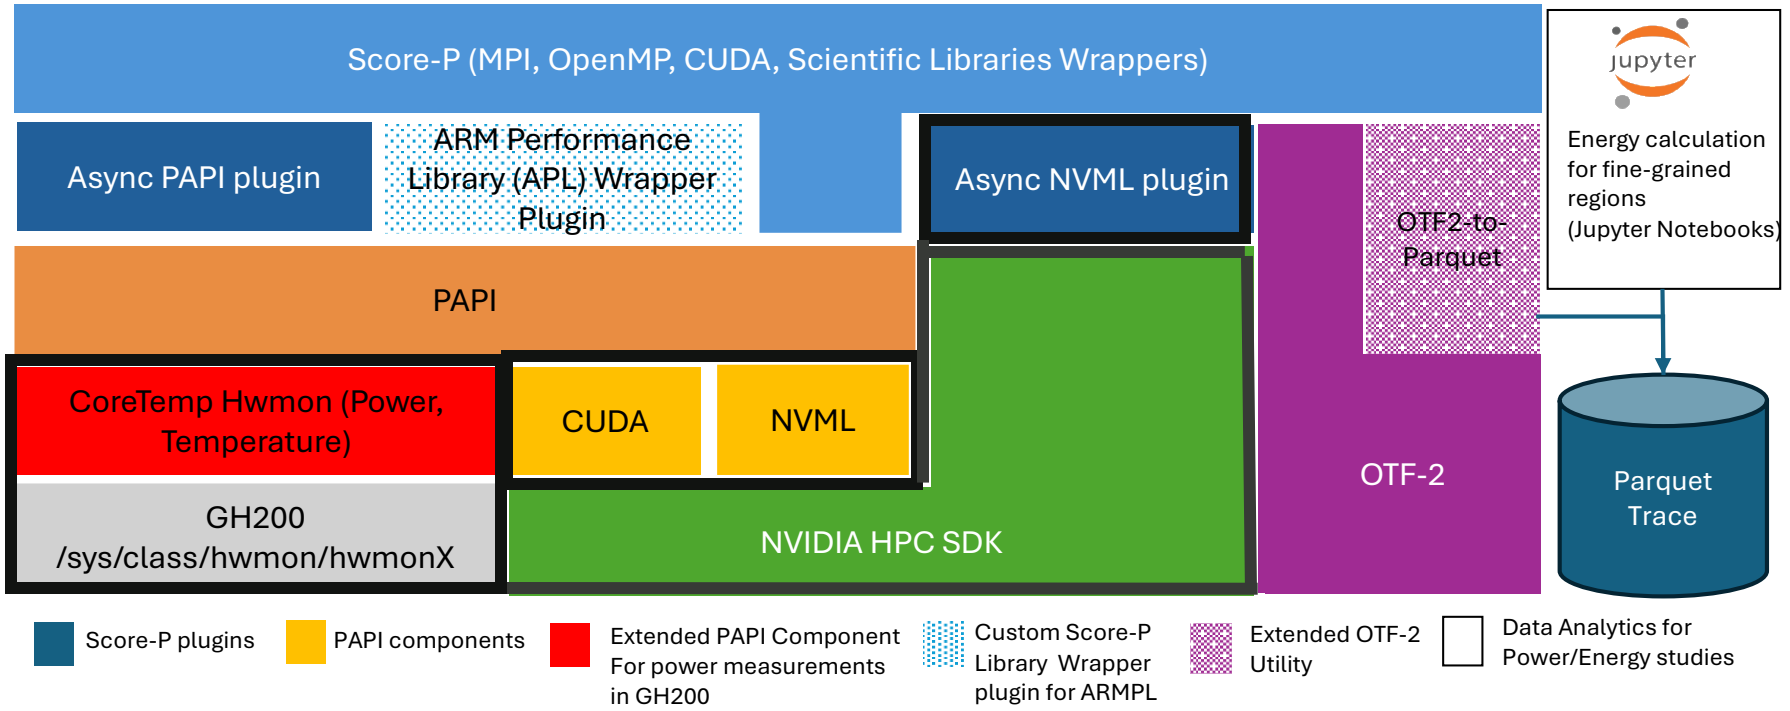

HPCAsia 2025, International Workshop on Arm Based HPC. ***Preliminary Study on Fine-Grained Power and Energy Measurements on Grace Hopper GH200 with Open-Source Performance Tools.*** ***https://doi.org/10.1145/3703001.3724383***

Extended Components

- PAPI extension for hwmon to measure Superchip and Grace CPU power
- Score-P NVML and PAPI plugins readings
- OTF2-to-Parquet conversion utility
- Jupyter-based analytics for energy profiling

Supported Architectures

- GH200, potential for AMD Instinct and other platforms

Based on tools from:

VAMPIR · Score-P Scalable performance measurement infrastructure for parallel codes · icl-utk-edu/papi · NVIDIA · arm

PAPI component improvements

Component	Improvement
ROCM_SMI	Added accumulated energy support
Coretemp	Added support for cray_pm counters
RAPL	Support for Milan and Trento CPUs

Score-P plugins customization

Plugin Name	Scope	PAPI component
ROCM_SMI	Per Node	ROCM_SMI
CRAY_PM	Per Node	Coretemp extended to read cray_pm counters
RAPL	Per Node	rapl

Key Benefits

- A unified timeline where power and energy measurements from different libraries, sub-systems, and tools are available.
- Using application trace timeline enables fine-grained analysis of power/energy characteristics of specific code regions
- Open storage format (OTF2) enables third party tools for analysis and insight (though scaling remains a challenge).

User accessible power management (PM) on Cray EX nodes

- Target telemetry and job profiling
- Collection for point-in-time power and energy data
- Made available using sysfs interface **`/sys/cray/pm_counters/XYZ`**
- Source: HPE `User Access to Compute Node Power Data`

File	Description
power	Point-in-time power (Watts).
energy	Accumulated energy, in joules.
cpu_power	Point-in-time power (Watts) used by the CPU domain.
cpu_energy	The total energy (Joules) used by the CPU domain.
cpu_temp	Temperature reading (Celsius) of the CPU domain.
memory_power	Point-in-time power (Watts) used by the memory domain.
memory_energy	The total energy (Joules) used by the memory domain.
accel_energy	Accumulated accelerator energy (Joules).
accel_power	Accelerator point-in-time power (Watts).
generation	A counter that increments each time a power cap value is changed.
startup	Startup counter.
freshness	Free-running counter that increments at a rate of approximately 10Hz.
version	Version number for power management counter support.
power_cap	Current power cap limit in Watts; 0 indicates no capping.
raw_scan_hz	The power management scanning rate for all data in pm_counters.

Accessing GPU Power/Energy on Frontier Nodes

Feature	`rsmi_dev_energy_count_get()`	`rsmi_dev_power_ave_get()`	`cray_pm power/energy`
Scope	MI250X GPU	MI250X GPU	MI250X GPU (Blade Controller)
Sampling Interval	1 millisecond	1 millisecond	100 millisecond
Units	Micro Joule	Micro Watt	Watts / Joules
Granularity	High resolution energy counter	Average power over time	Mid-resolution sensors/counters
Accessibility	rocm_smi (amd_smi)	rocm_smi (amd_smi)	`/sys/cray/pm_counters`
Sampling rate	Fixed Interval	Fixed Interval	Fixed Interval
Overhead	API (through sysfs)	API (through sysfs)	User sysfs IO access
Use Case	Fine-grained event correlation	Coarse-grained event correlation	Coarse-grained monitoring / telemetry

Recent rocm_smi library versions include **`rsmi_dev_current_socket_power()`** and **`rsmi_dev_power_get()`** which return average power on the MI250X

Support for Accumulated Energy Monitoring per GPU in ROCm-SMI PAPI component

ROCM_SMI function Integrated:

```
rsmi_status_t rsmi_dev_energy_count_get(uint32_t device, uint64_t *count, float *resolution, uint64_t *timestamp)
```

New Native Event Added:

- Format: `energy_count:device=X`
- Description: *"Accumulated GPU energy, in microjoules (μJ)"*

Usage:

```
%papi_native_avail | grep energy
| rocm_smi:::energy_count:device=0
|          Accumulated GPU energy, in microjoules (μJ).
```

Integration Points:

- Function pointer registered in dynamic symbol table
- Event handler **access_rsmi_dev_energy_count(...)** implemented to read energy in μJ
- Event appended to `event_function_table` and registered in htable

Counter access:

- Accumulated energy = `(raw_count - offset) × resolution` `int64_t` in μJ
- Offset refers to the initial value of the energy during initialization, used to start energy accumulation from zero.

Instantaneous GPU Power Estimation using rocm_smi

Methodology

- rocm_smi provides accumulated energy consumption in microjoules (uJ).
- Instantaneous power is approximated using the difference between consecutive energy samples:

$$P_{\text{inst}}^{(i)} \approx \frac{E^{(i)} - E^{(i-1)}}{\Delta t}$$

Sampling Details

- Sampling rate: 1 ms - allows mapping transient power spikes to specific GPU kernels
- Enables correlation of power surges with high-activity regions in application traces

Tradeoffs

- High temporal resolution → More accurate correlation to kernel activity
- May introduce **apparent overestimation** due to sampling noise or short-term transients

Frontier Overview

Architecture

- **9856 Node HPE Cray EX with Slingshot Interconnect**

Key Node Features

- 64-core AMD EPYC 7A53 CPU (Trento)
- 4x AMD MI250X GPU with 128GB HBM3
    - 2x64 GB GCDS
- Infinity Fabric CPU-GPU and GPU-GPU links

TDP: 500–560W (*GPU*) 225-240W (*CPU*)

Challenges

- Managing power across CPU, GPU, and I/O
- High-resolution monitoring for transient power spikes

Frontier Node

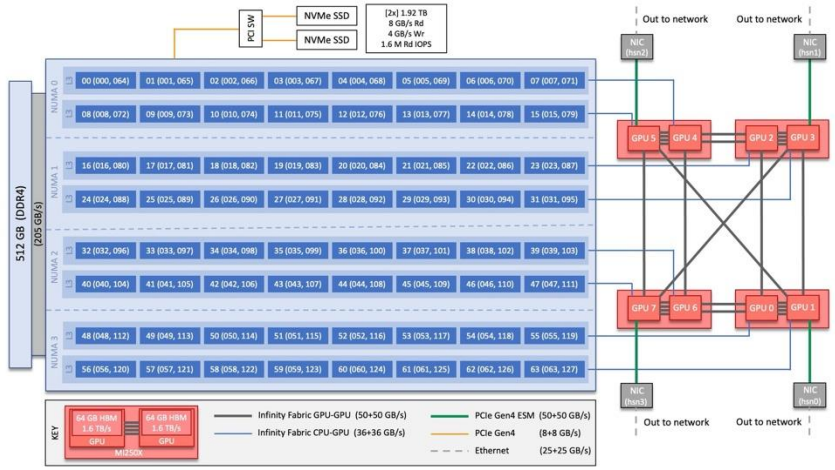

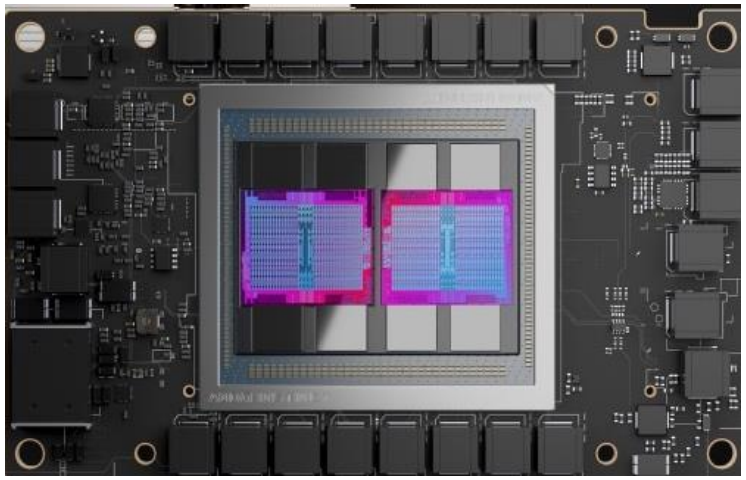

AMD MI250X

Case Study: GPU SGEMM

- **Application trace timeline**
- CPU initialization phase
- GPU Execution Phase

A

- **SMI average power measurements**
- Sampling rate : 1 ms
- Ramp up response to instantaneous change in activity

B

- **SMI instantaneous power estimate**
- Sampling rate 1 ms
- Timely capture of sudden changes
- Min/Max/Average ➔ variability

C

- **Cray PM**
- Sampling rate : 1 ms
- Delayed response to sudden changes, little/no variability

D

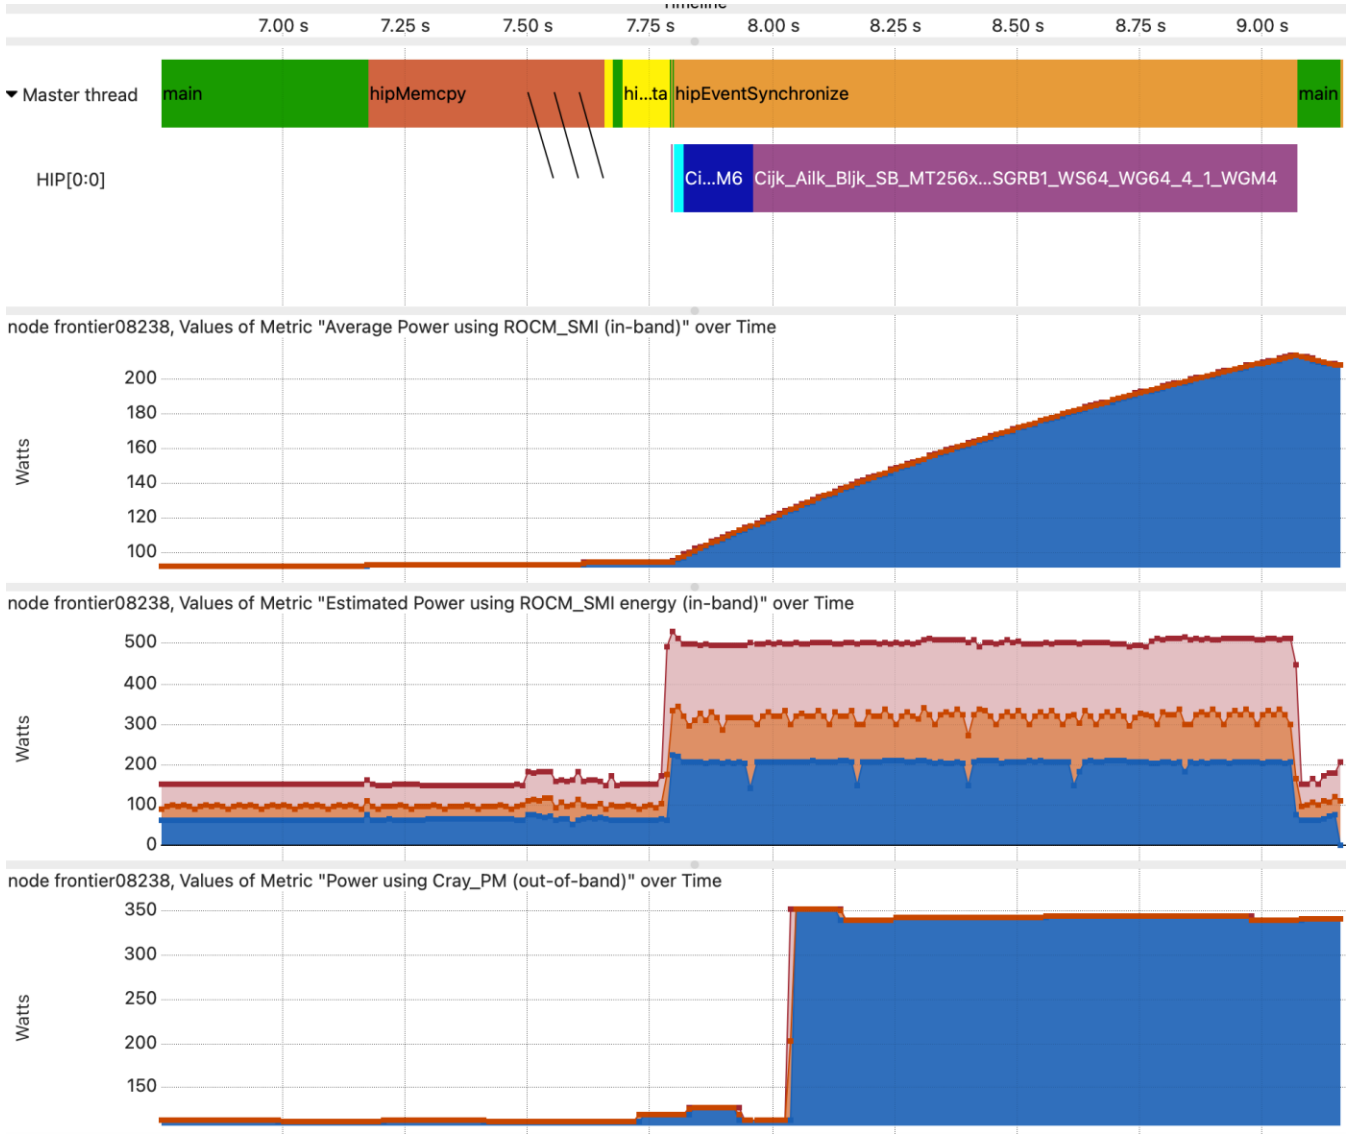

Case Study : CPU DGEMM with OpenMP

Sampling Rate : 1 ms

CPU thread activity timeline
Barriers do not consume much energy

Process Timeline

Cray PM CPU

Cray PM CPU + Memory

RAPL Package0

RAPL PP0

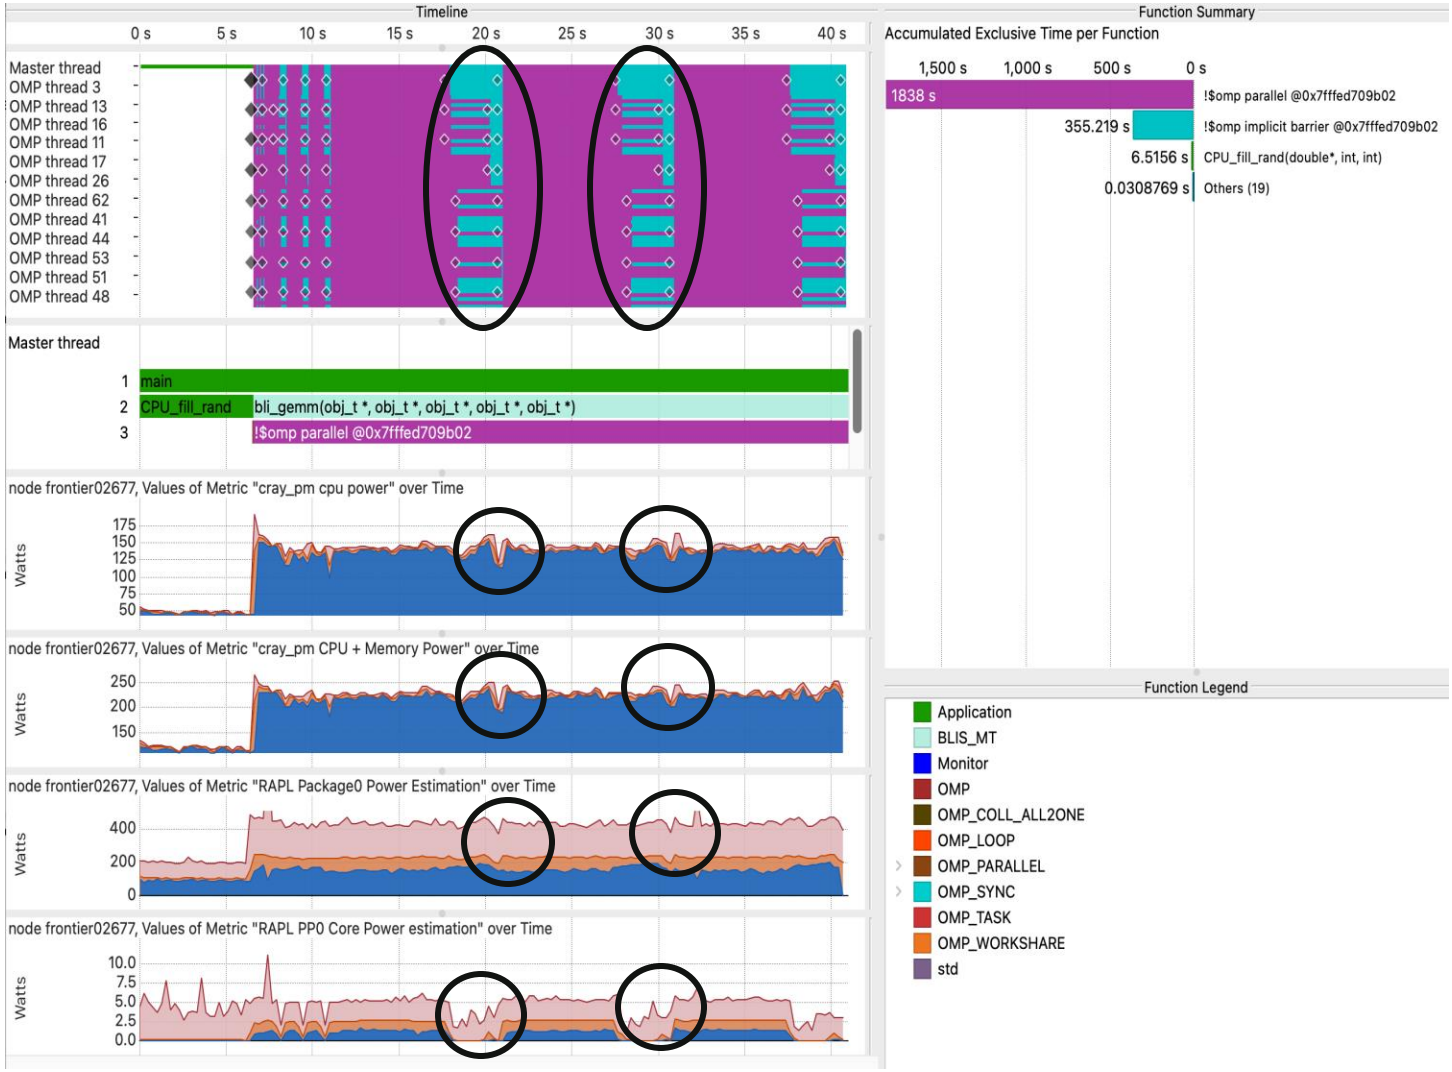

Case Study: rocHPL 6.0.0 (Single rank, 7 OpenMP threads)

Sampling Rate : 1 ms

- Process timeline
- Cray PM node power
- RAPL Package Power
- Cray PM Package (CPU + Memory)
- RAPL Package0
- SMI Instantaneous power estimate
- SMI GPU utilization
- SMI average power

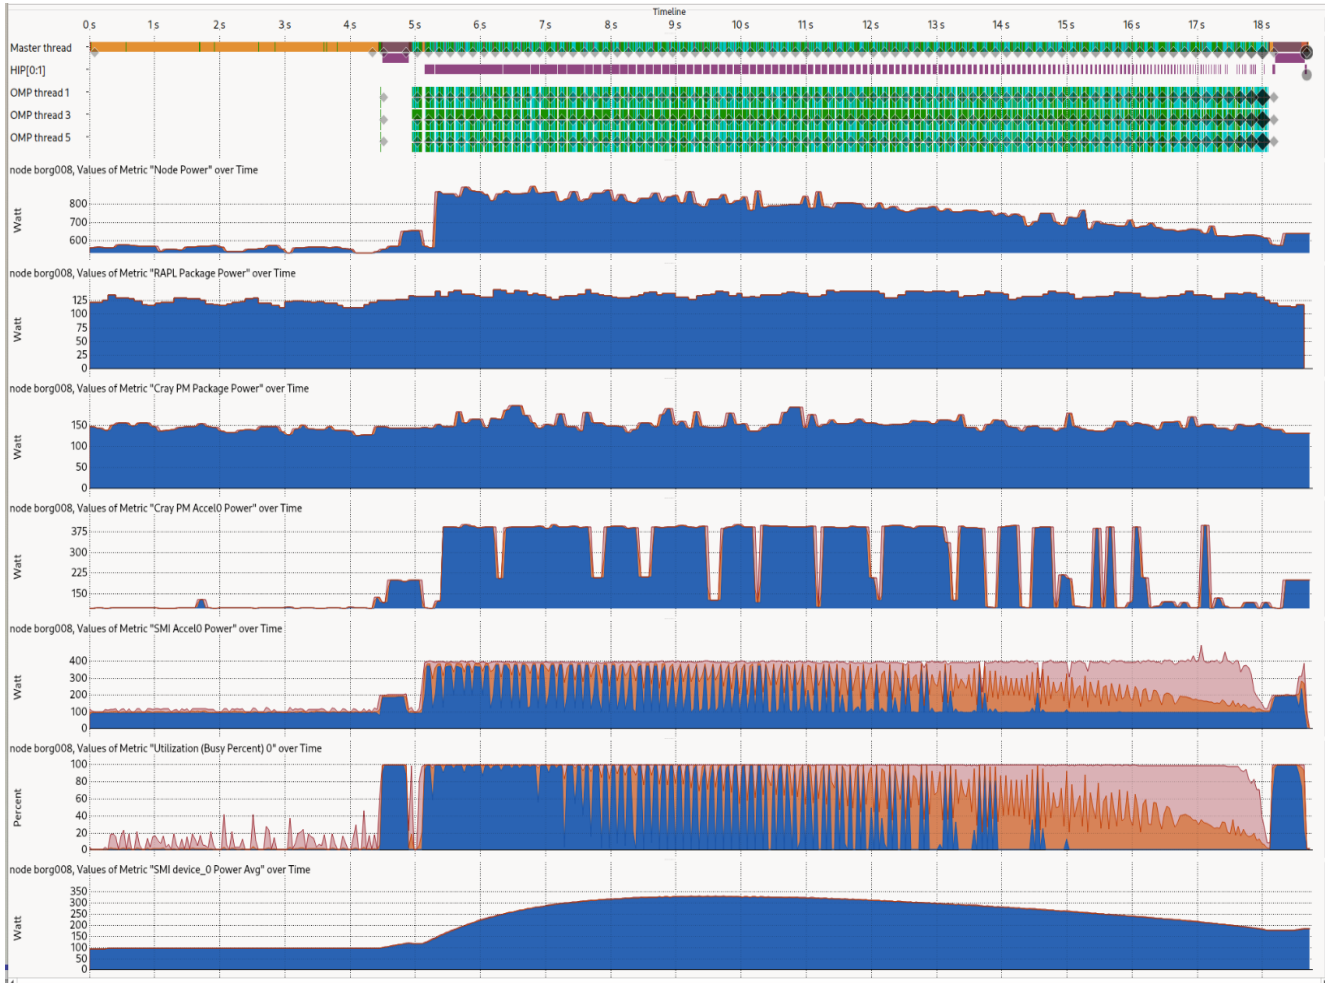

Case Study : 8 rank rocHPL x 7 OpenMP threads

- Process timeline
- Cray PM node power
- RAPL Package Power
- Cray PM Package (CPU + Memory)
- Cray PM Sum Component Power
- Sum SMI power + Cray PM memory + Cray PM CPU
- Sum SMI power + RAPL Package Power

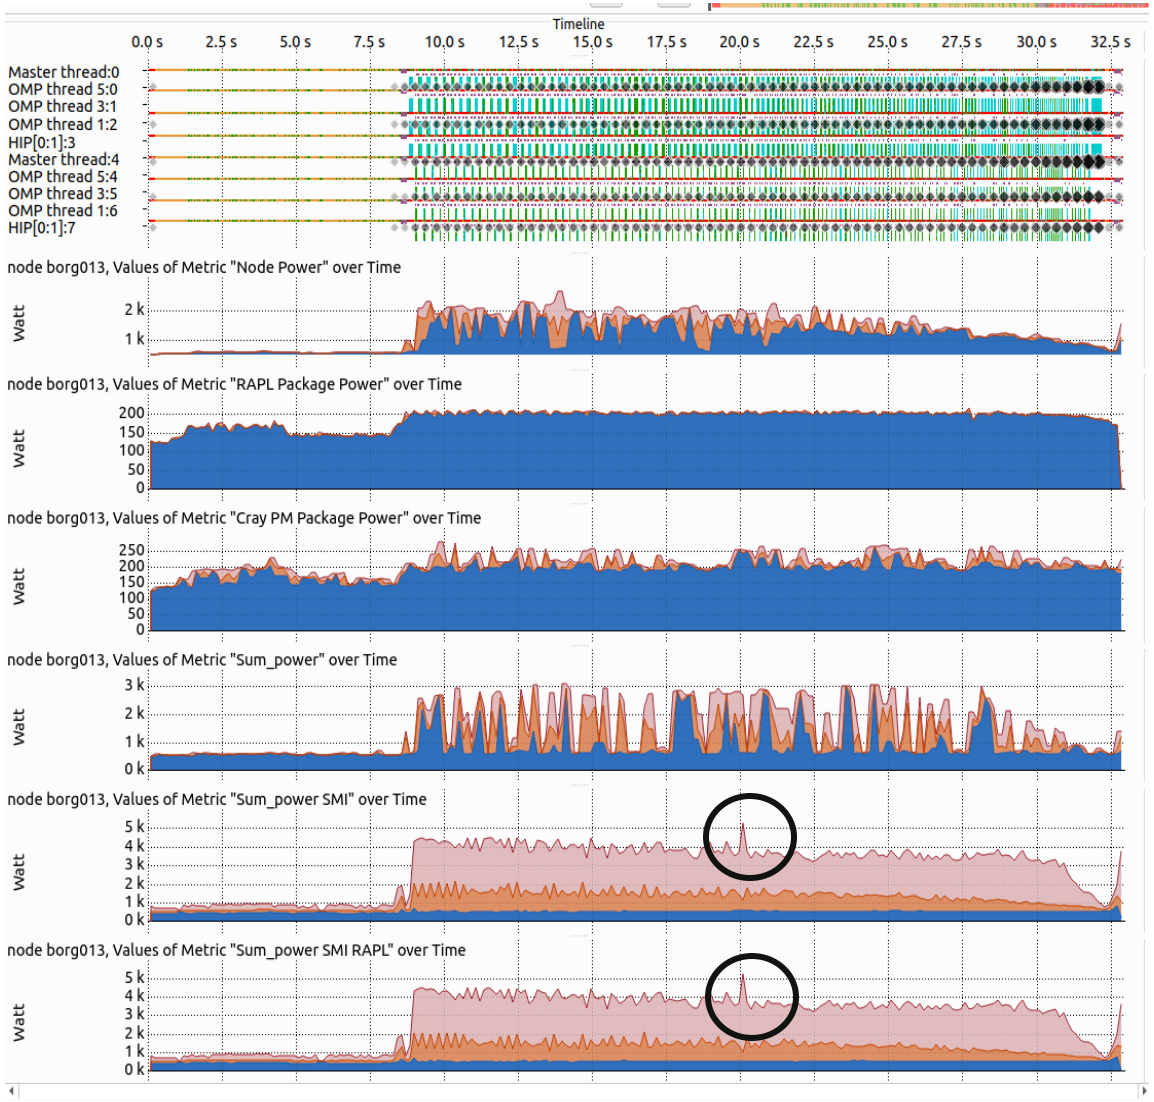

Case Study : 8 rank rocHPL x 7 OpenMP threads

- Process timeline
- Cray PM node power
- RAPL Package Power
- Cray PM Package (CPU + Memory)
- Cray PM Sum Component Power
- Sum SMI power + Cray PM memory + Cray PM CPU
- Sum SMI power + RAPL Package Power

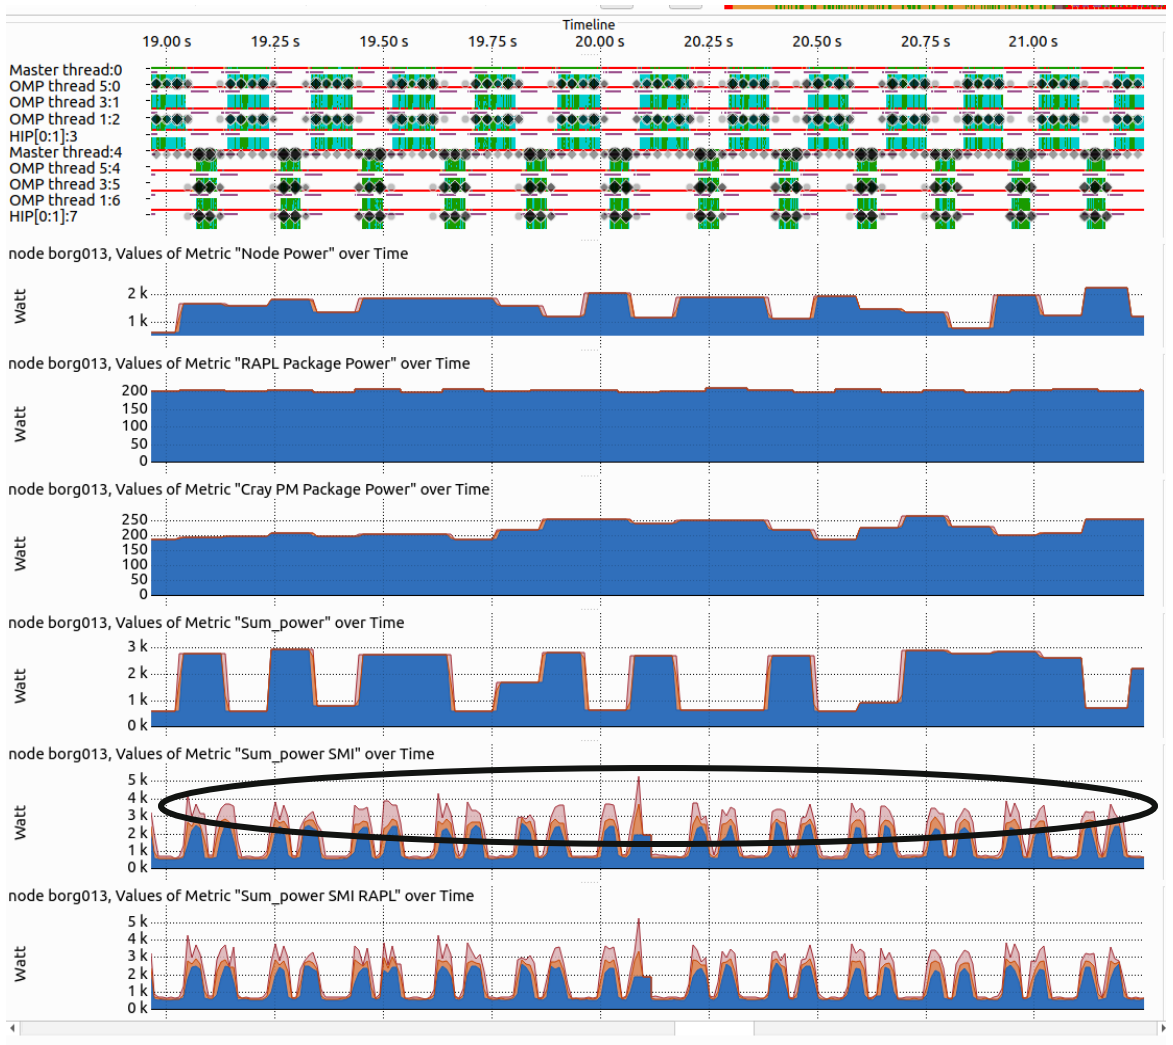

Conclusions & Further work

- Toolkit for fine grained application energy and power analysis
  - Combines Score-P, Cray PM, ROCm SMI, and PAPI
  - Single timeline for in-band and out-of-band measurements
  - Capture rapid changes in power usage, attribute to code regions
  - Foundation for detailed energy analysis and optimizations
- Further investigations and questions
  - Scalable analytics for large scale OTF2 traces
  - RAPL readings for core energy consumption
  - Investigate energy dependence on performance counters
  - Cray PM aggregate node power vs sum of individual components

Acknowledgement

This research used resources from the Oak Ridge Leadership Computing Facility, which is a US Department of Energy (DOE) Office of Science user facility supported under contract DE-AC05-00OR22725. This work was also supported by the DOE Office of Science, Advanced Scientific Computing Research, Express project “Leveraging Open- Source Simulators to Enable hw/sw Co-design of Next-generation HPC Systems” (DE-FOA-0002950).

2025 OLCF User Meeting

August 5-6, 2025

- Facility updates & expert talks
- Deep-dive discussions with OLCF staff
- Networking with top computational scientists

Register Today!

www.olcf.ornl.gov/olcf-user-meeting

OAK RIDGE National Laboratory

Questions?