

Overview of HPCToolkit

Performance Tools for GPU-accelerated Supercomputing

John Mellor-Crummey, Rice University
ALCF Hands-on HPC Workshop
Oct. 11, 2023

HPCToolkit Project Team

- **Rice University**

- HPCToolkit PI: Prof. John Mellor-Crummey
- Research staff: Laksono Adhianto, Mark Krentel, Wil Phan, Matt Barnett
- Contractor: Marty Itzkowitz
- Students: Jonathon Anderson, Dragana Grbic, Vladimir Indjic, Yumeng Liu

- **University of Wisconsin – Madison**

- Dyninst PI: Prof. Barton Miller

Outline

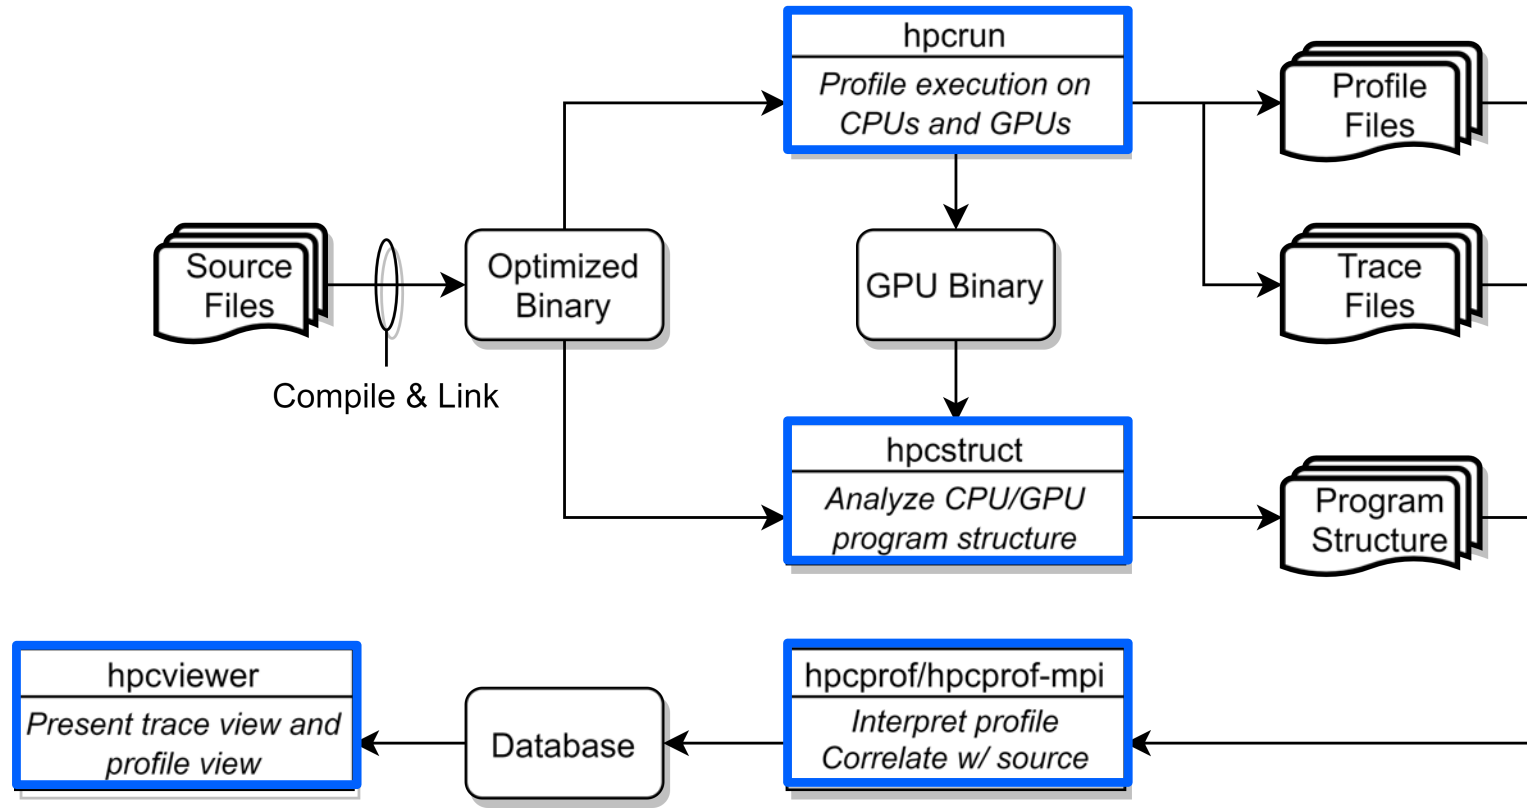
- **HPCToolkit performance tools for CPU and GPU-accelerated applications**
 - Overview of HPCToolkit components and their workflow
 - HPCToolkit's graphical user interfaces
- **Analyzing the performance of GPU-accelerated codes with HPCToolkit**
 - GAMESS
 - Quicksilver
 - LAMMPS at Exascale
- **Status**
- **Resources**
- **Hands-on directions**

Rice University's HPCToolkit Performance Tools

Measure and analyze performance of CPU and GPU-accelerated applications

- **Easy: profile unmodified application binaries**
- **Fast: low-overhead measurement**
- **Informative: understand where an application spends its time and why**
 - call path profiles associate metrics with application source code contexts
 - optional hierarchical traces to understand execution dynamics
- **Broad audience**
 - application developers
 - framework developers
 - runtime and tool developers
- **Supported platforms**
 - CPU: x86_64, Power, ARM
 - GPU: NVIDIA, AMD, Intel

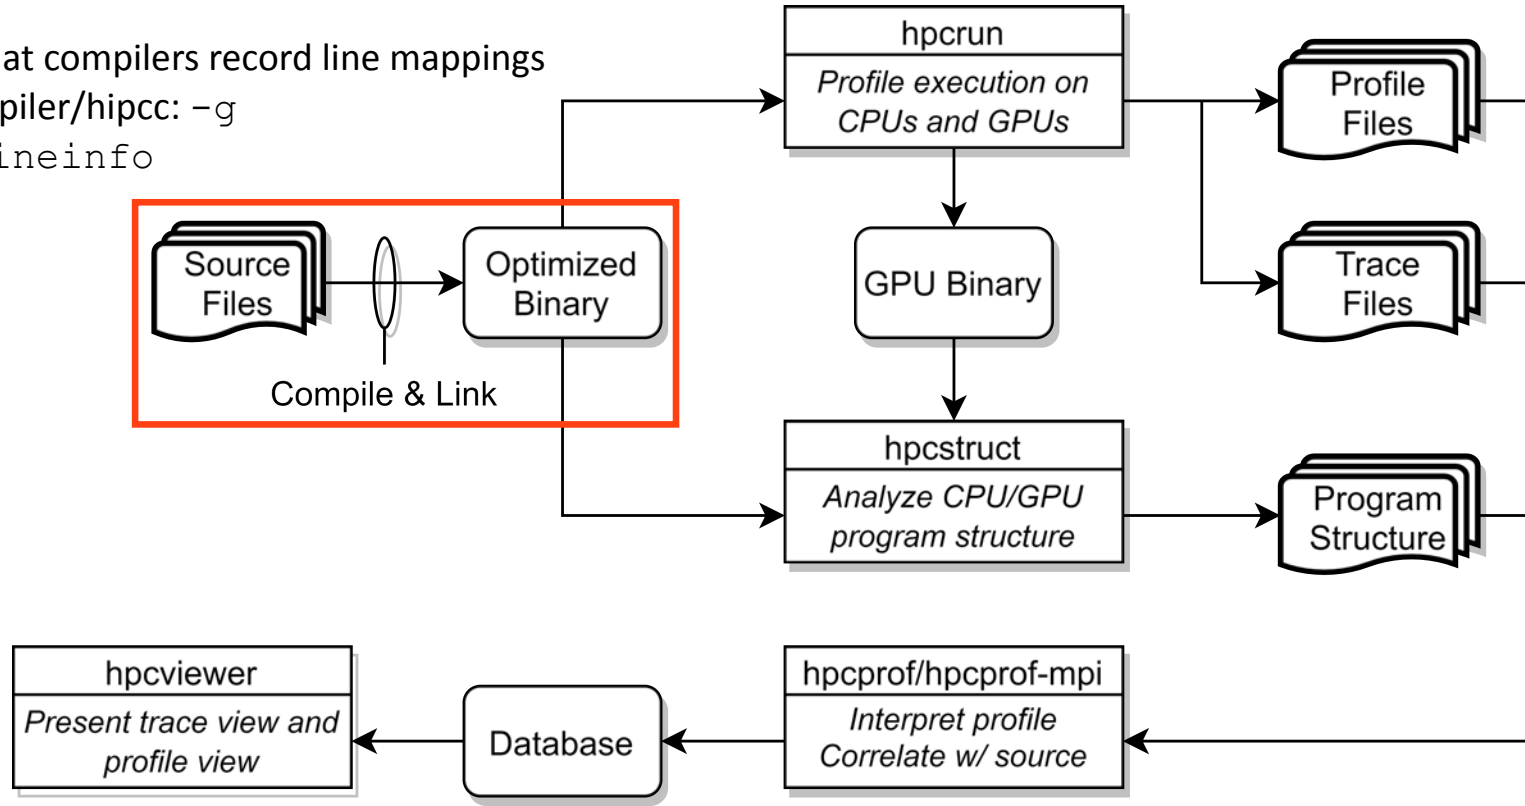
HPCToolkit's Workflow for GPU-accelerated Applications



HPCToolkit's Workflow for GPU-accelerated Applications

Step 1:

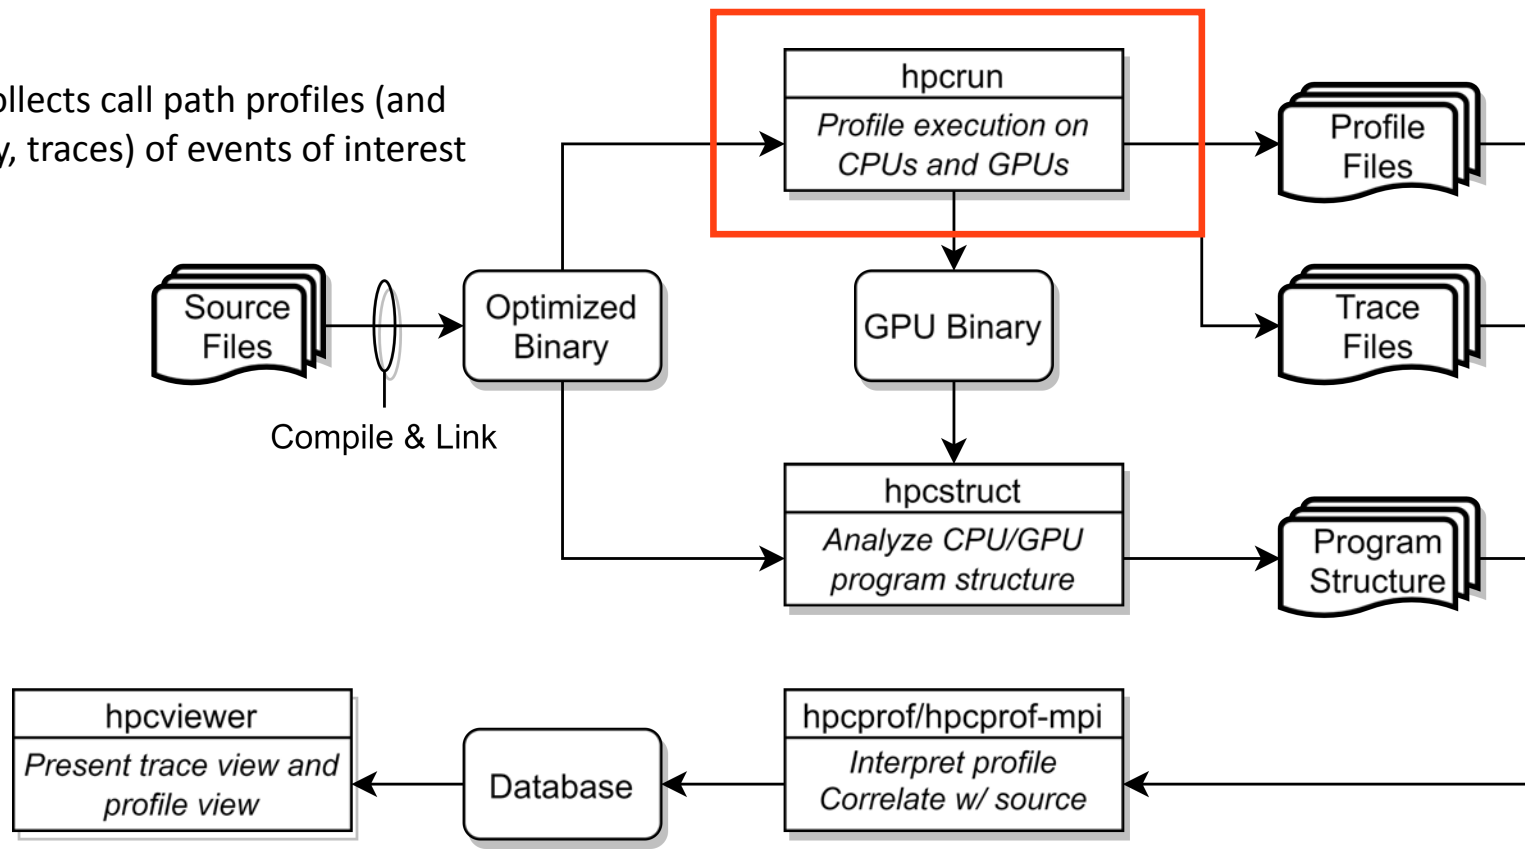
- Ensure that compilers record line mappings
- host compiler/hipcc: `-g`
- nvcc: `-lineinfo`



HPCToolkit's Workflow for GPU-accelerated Applications

Step 2:

- *hpcrun* collects call path profiles (and optionally, traces) of events of interest



Measurement of CPU and GPU-accelerated Applications

- **Sampling using Linux timers and hardware counter overflows on the CPU**
- **Callbacks when GPU operations are launched and (sometimes) completed**
- **Event stream for GPU operations**
- **Instruction-level measurements**
 - PC Samples (NVIDIA)
 - Binary instrumentation of GPU kernels (Intel)

hpcrun: Measure CPU and/or GPU activity

- **GPU profiling**

- hpcrun -e gpu=**xxx** <app> ...

xxx ∈ {*nvidia,amd,opencl,level0*}

- **GPU PC sampling (NVIDIA GPU only)**

- hpcrun -e gpu=**nvidia**,pc <app>

- **CPU and GPU Tracing (in addition to profiling)**

- hpcrun -e CPUTIME -e gpu=**xxx** **-t** <app>

- **Use hpcrun with job launchers**

- jsrun -n 32 -g 1 -a 1 hpcrun -e gpu=**xxx** <app>

Profiles

- a calling context tree per thread
- instruction level measurements

CPU traces

- trace of call stack samples

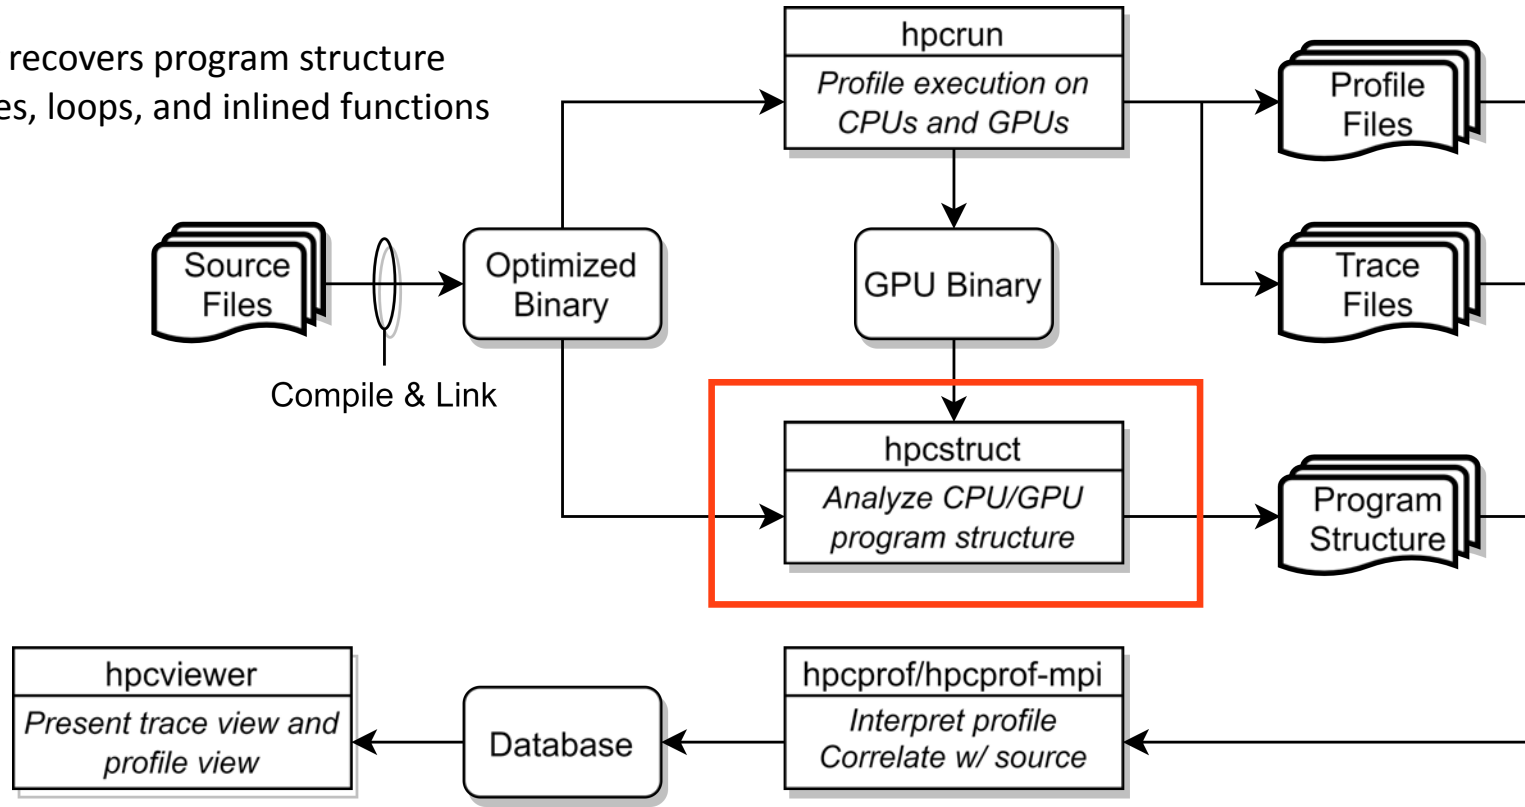
GPU traces

- trace of call stacks that initiate GPU operations

HPCToolkit's Workflow for GPU-accelerated Applications

Step 3:

- *hpcstruct* recovers program structure about lines, loops, and inlined functions



hpcstruct: Analyze CPU and GPU Binaries Using Multiple Threads

- **Usage**

```
hpcstruct [ --gpucfg yes ] <measurement-directory>
```

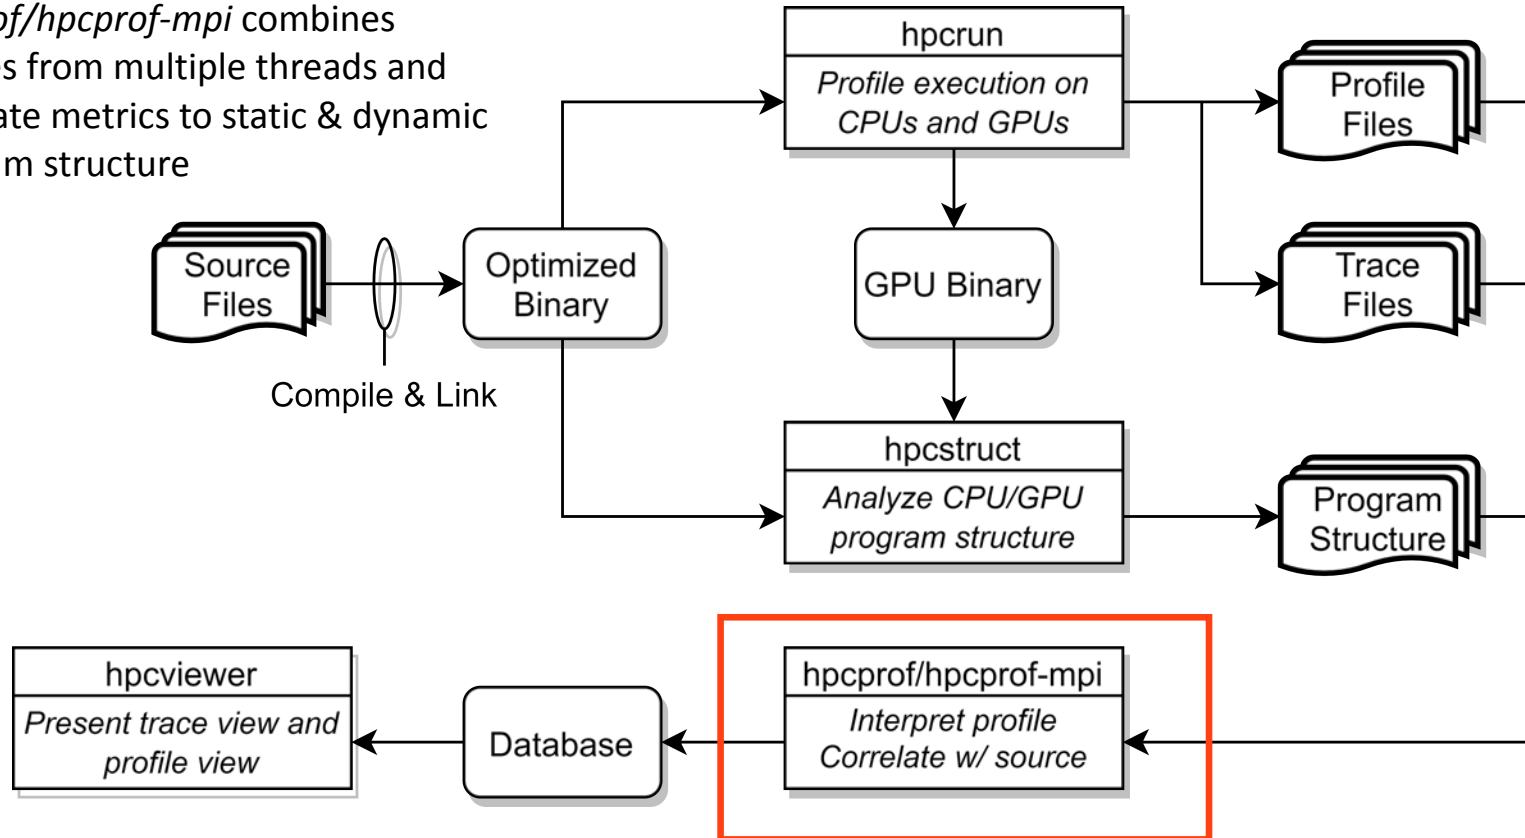
- **What it does**

- Recover program structure information
 - Files, functions, inlined templates or functions, loops, source lines
- Analyze all CPU and GPU binaries that were measured by HPCToolkit using multithreading
- Cache binary analysis results for reuse when analyzing other executions

HPCToolkit's Workflow for GPU-accelerated Applications

Step 4:

- *hpcprof/hpcprof-mpi* combines profiles from multiple threads and correlate metrics to static & dynamic program structure



hpcprof/hpcprof-mpi: Associate Measurements with Program Structure

- **Analyze data from modest executions with multithreading**

```
hpcprof <measurement-directory>
```

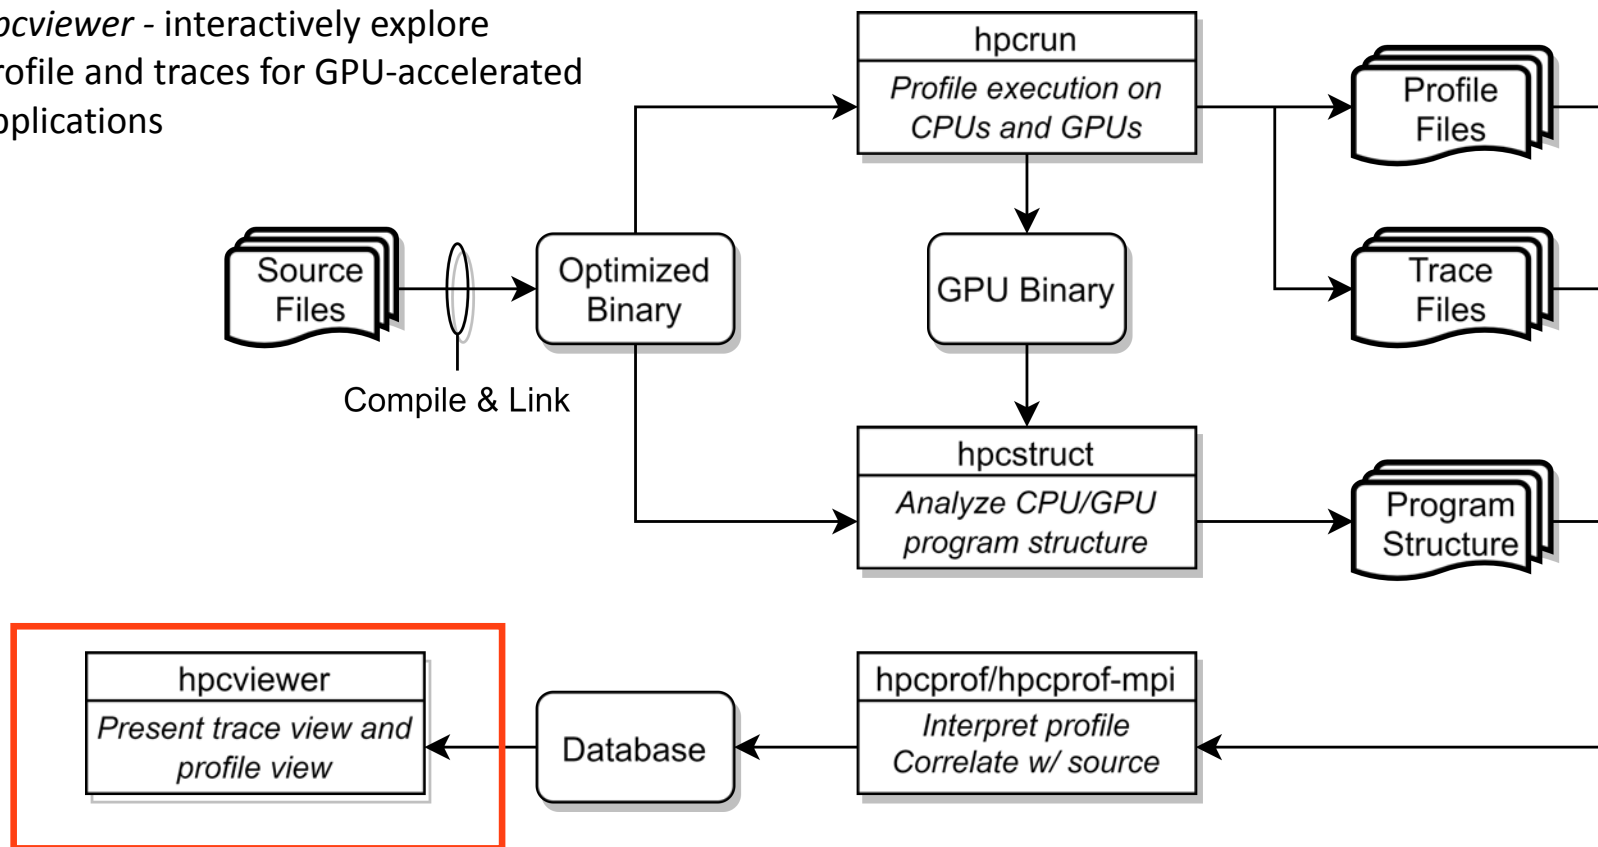
- **Analyze data from large executions with distributed-memory parallelism + multithreading**

```
jsrun -n 2 -a 1 -c 18 -b packed hpcprof-mpi <measurement-directory>
```

HPCToolkit's Workflow for GPU-accelerated Applications

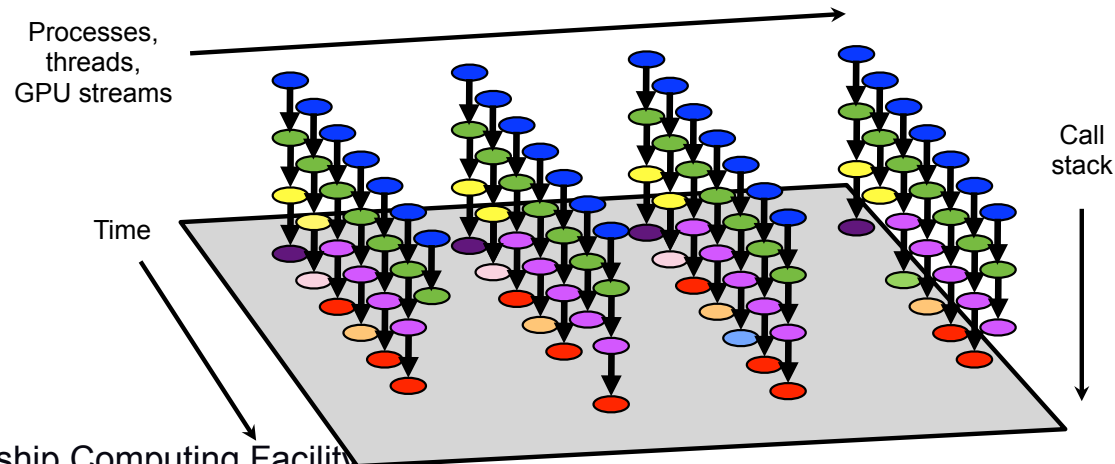
Step 4:

- *hpcviewer* - interactively explore profile and traces for GPU-accelerated applications



Understanding Temporal Behavior

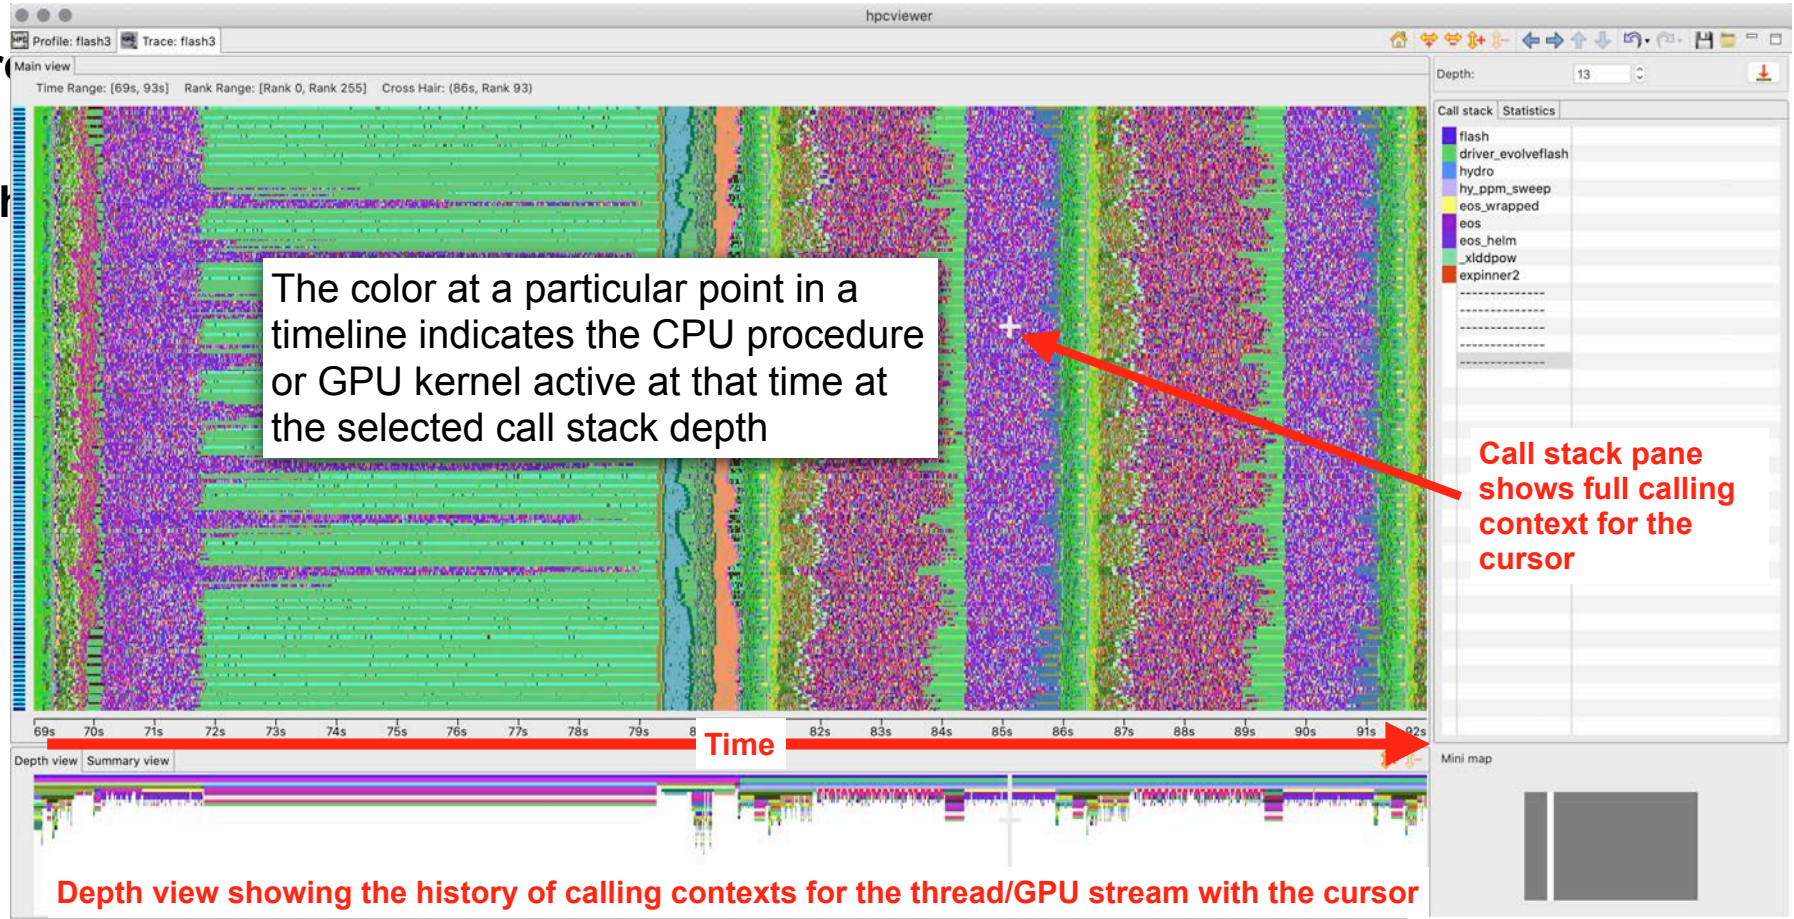
- **Profiling compresses out the temporal dimension**
 - Temporal patterns, e.g. serial sections and dynamic load imbalance are invisible in profiles
- **What can we do? Trace call path samples**
 - N times per second, take a call path sample of each thread
 - Organize the samples for each thread along a time line
 - View how the execution evolves left to right
 - What do we view? assign each procedure a color; view a depth slice of an execution



Time-centric Analysis with hpcviewer

MPI ranks, OpenMP Threads, GPU streams

- Profile
- Call stack



Minimap indicates part of execution trace shown

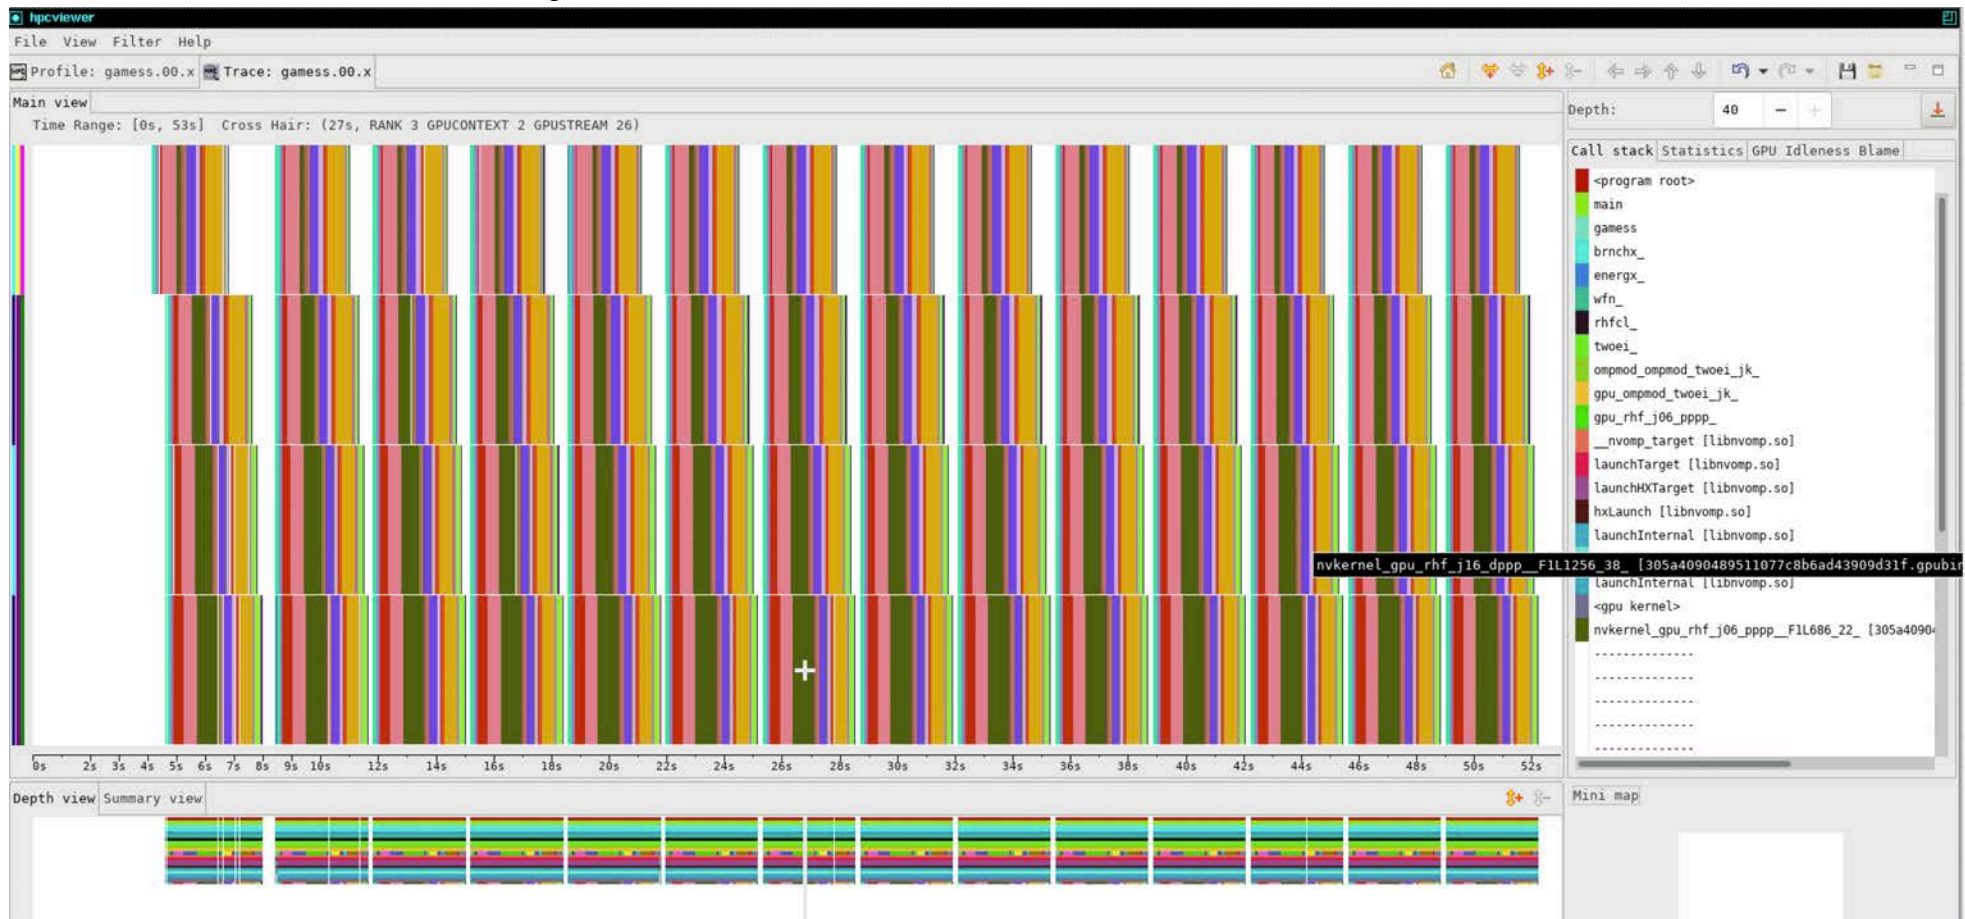
Case Studies

- **GAMESS: MPI, Fortran, OpenMP offloading**
- **Quicksilver: MPI, C++, CUDA**
- **LAMMPS: MPI, C++, Kokkos at exascale**

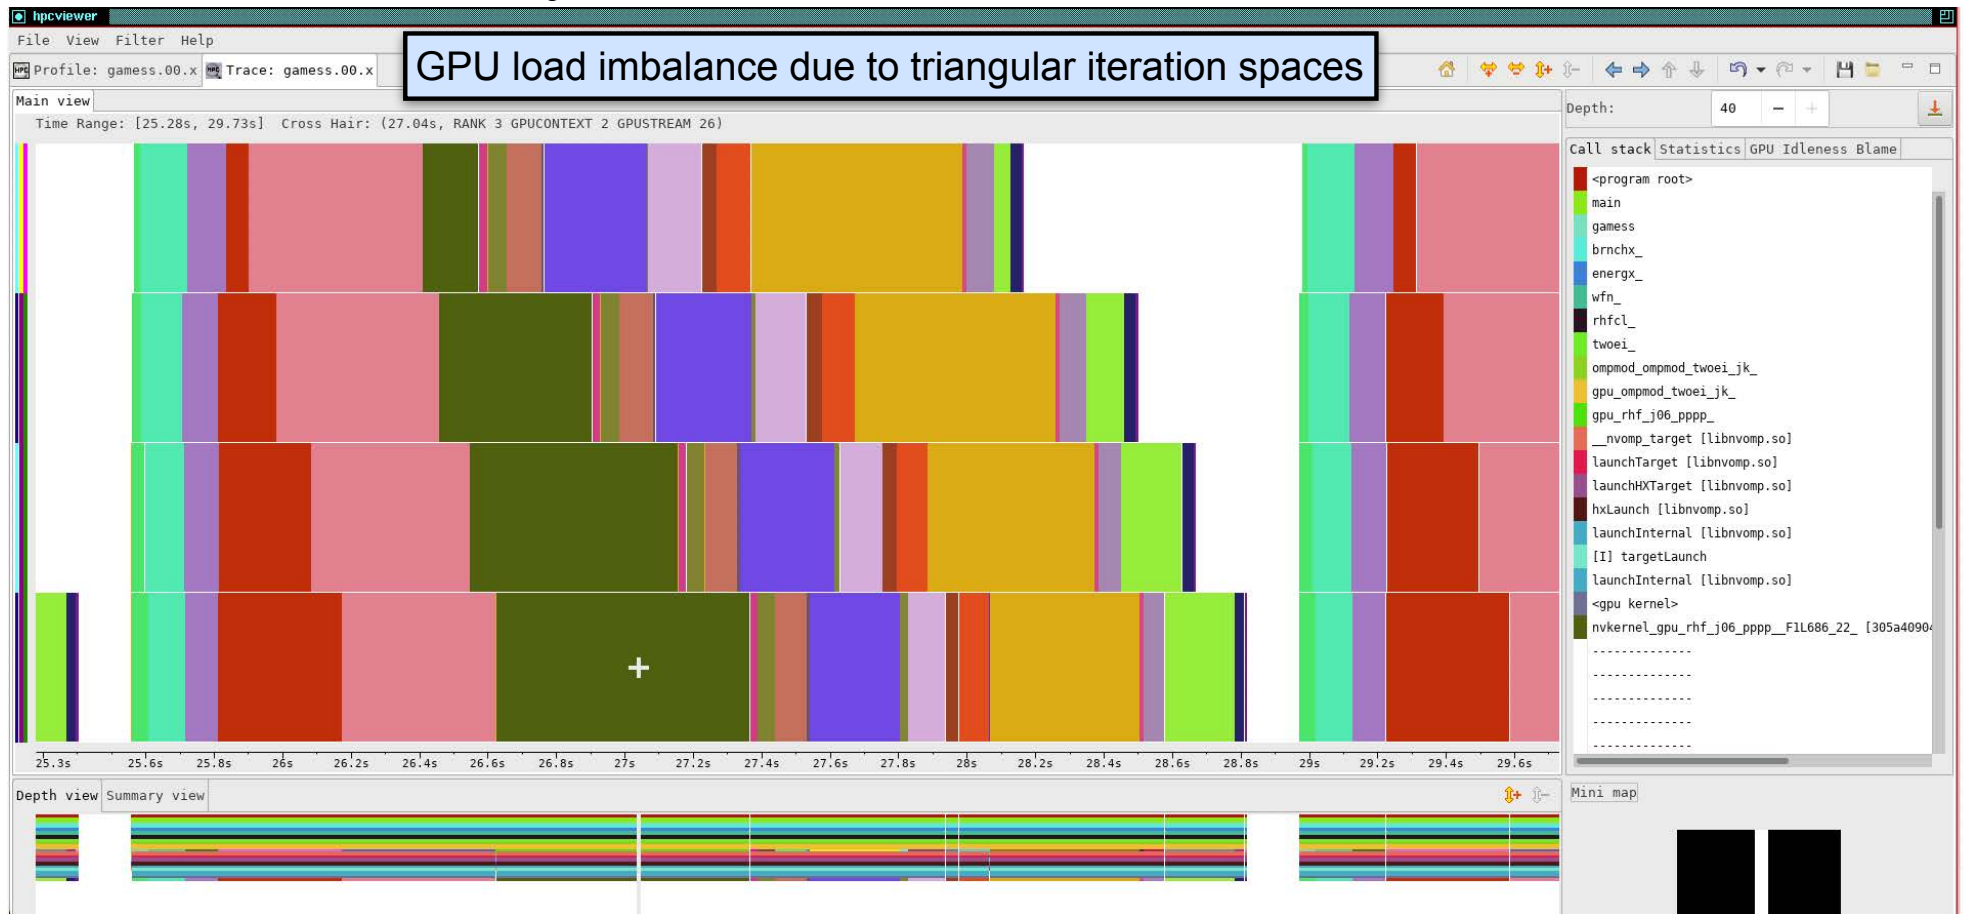
Time-centric Analysis: GAMESS 4 ranks, 4 GPUs on Perlmutter



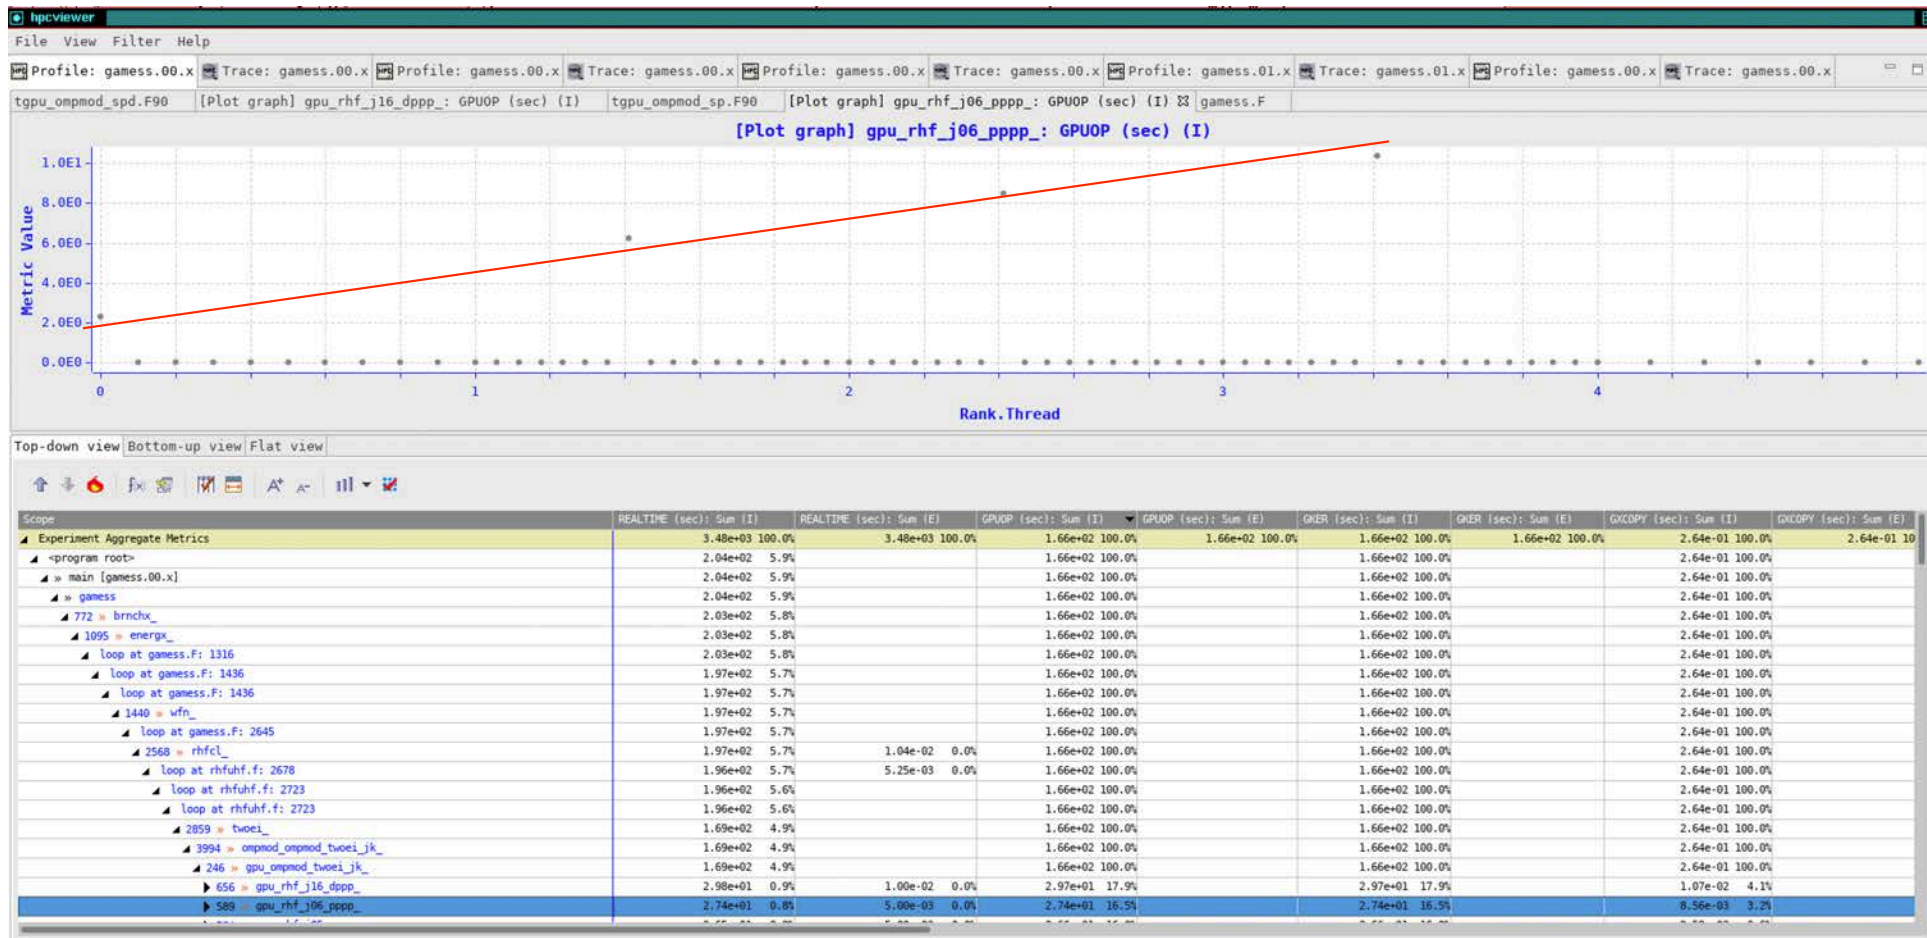
Time-centric Analysis: GAMESS 4 ranks, 4 GPUs on Perlmutter



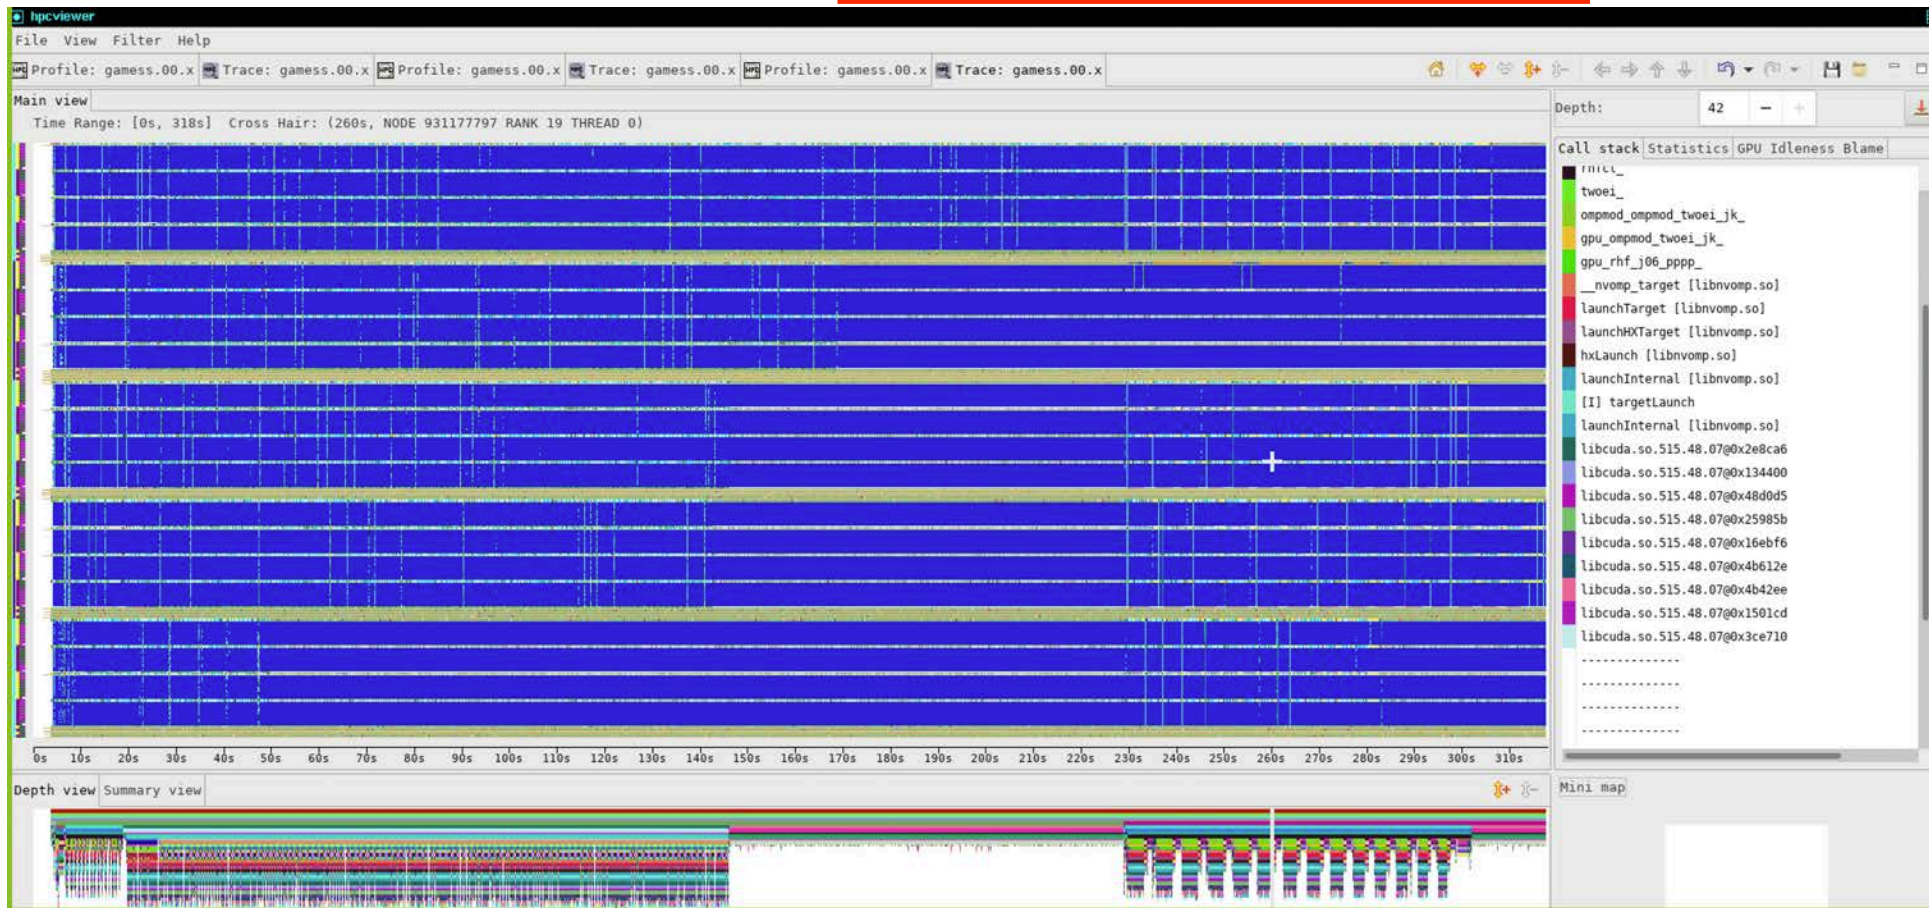
Time-centric Analysis: GAMESS 4 ranks, 4 GPUs on Perlmutter



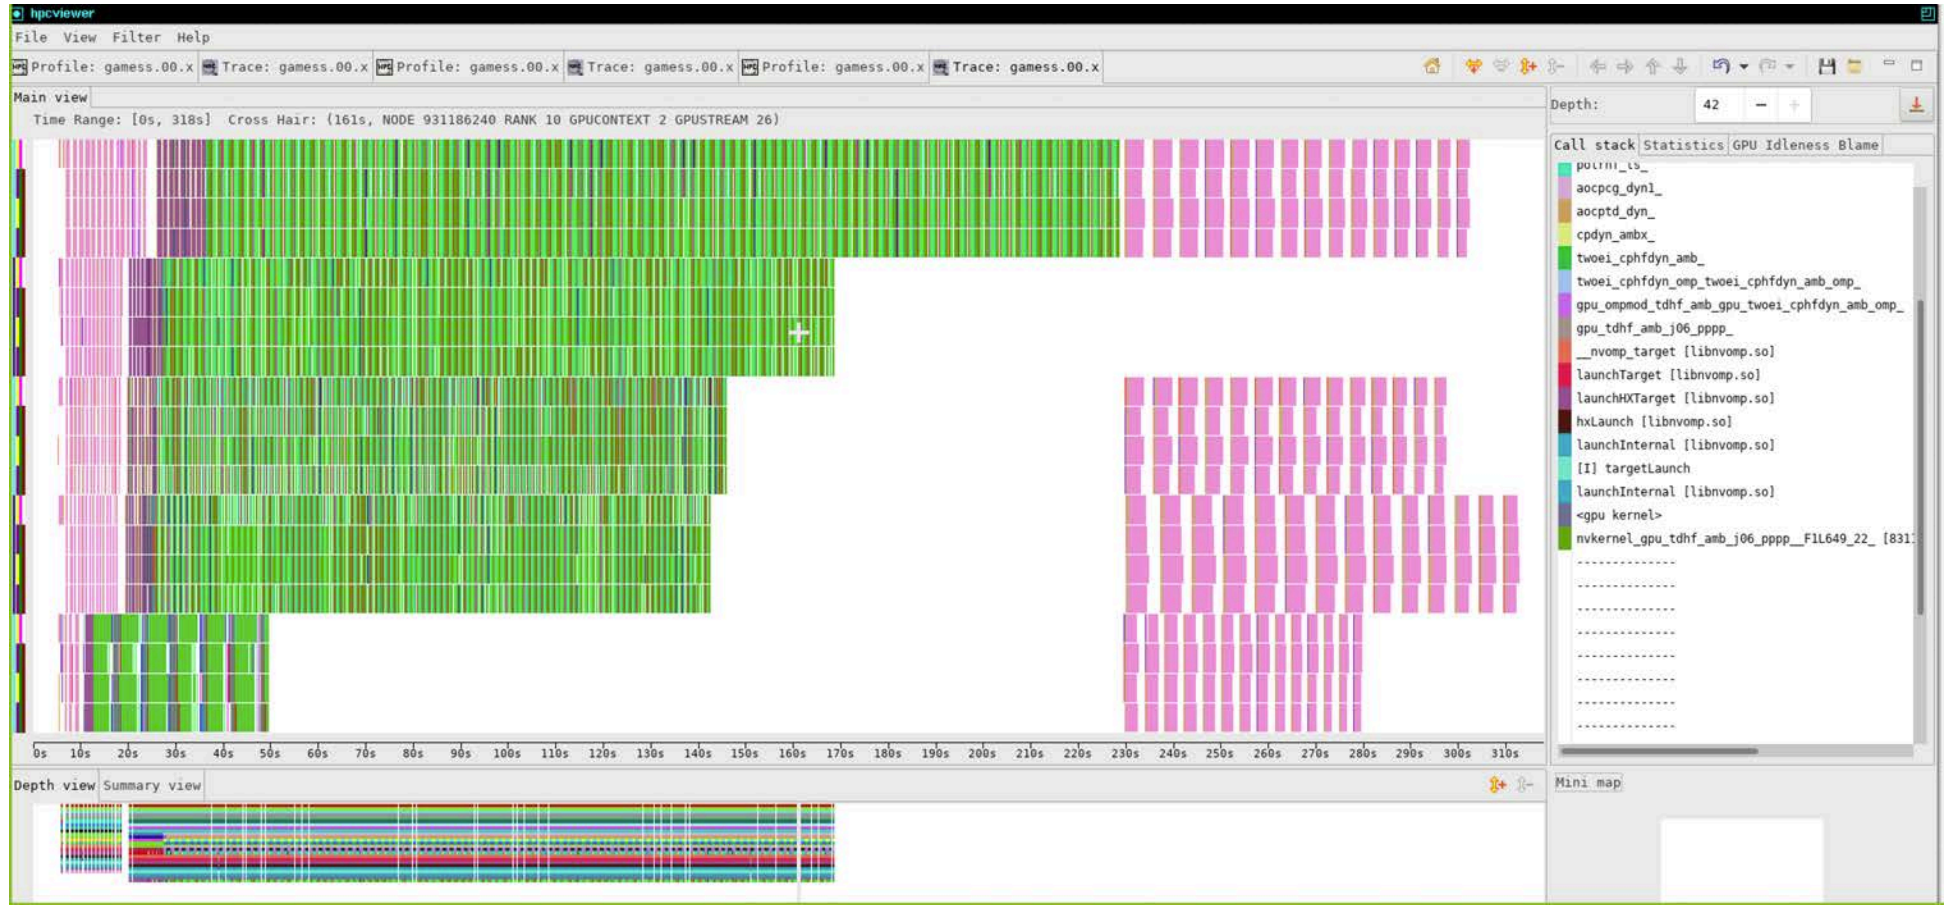
Time-centric Analysis: GAMESS 4 ranks, 4 GPUs on Perlmutter



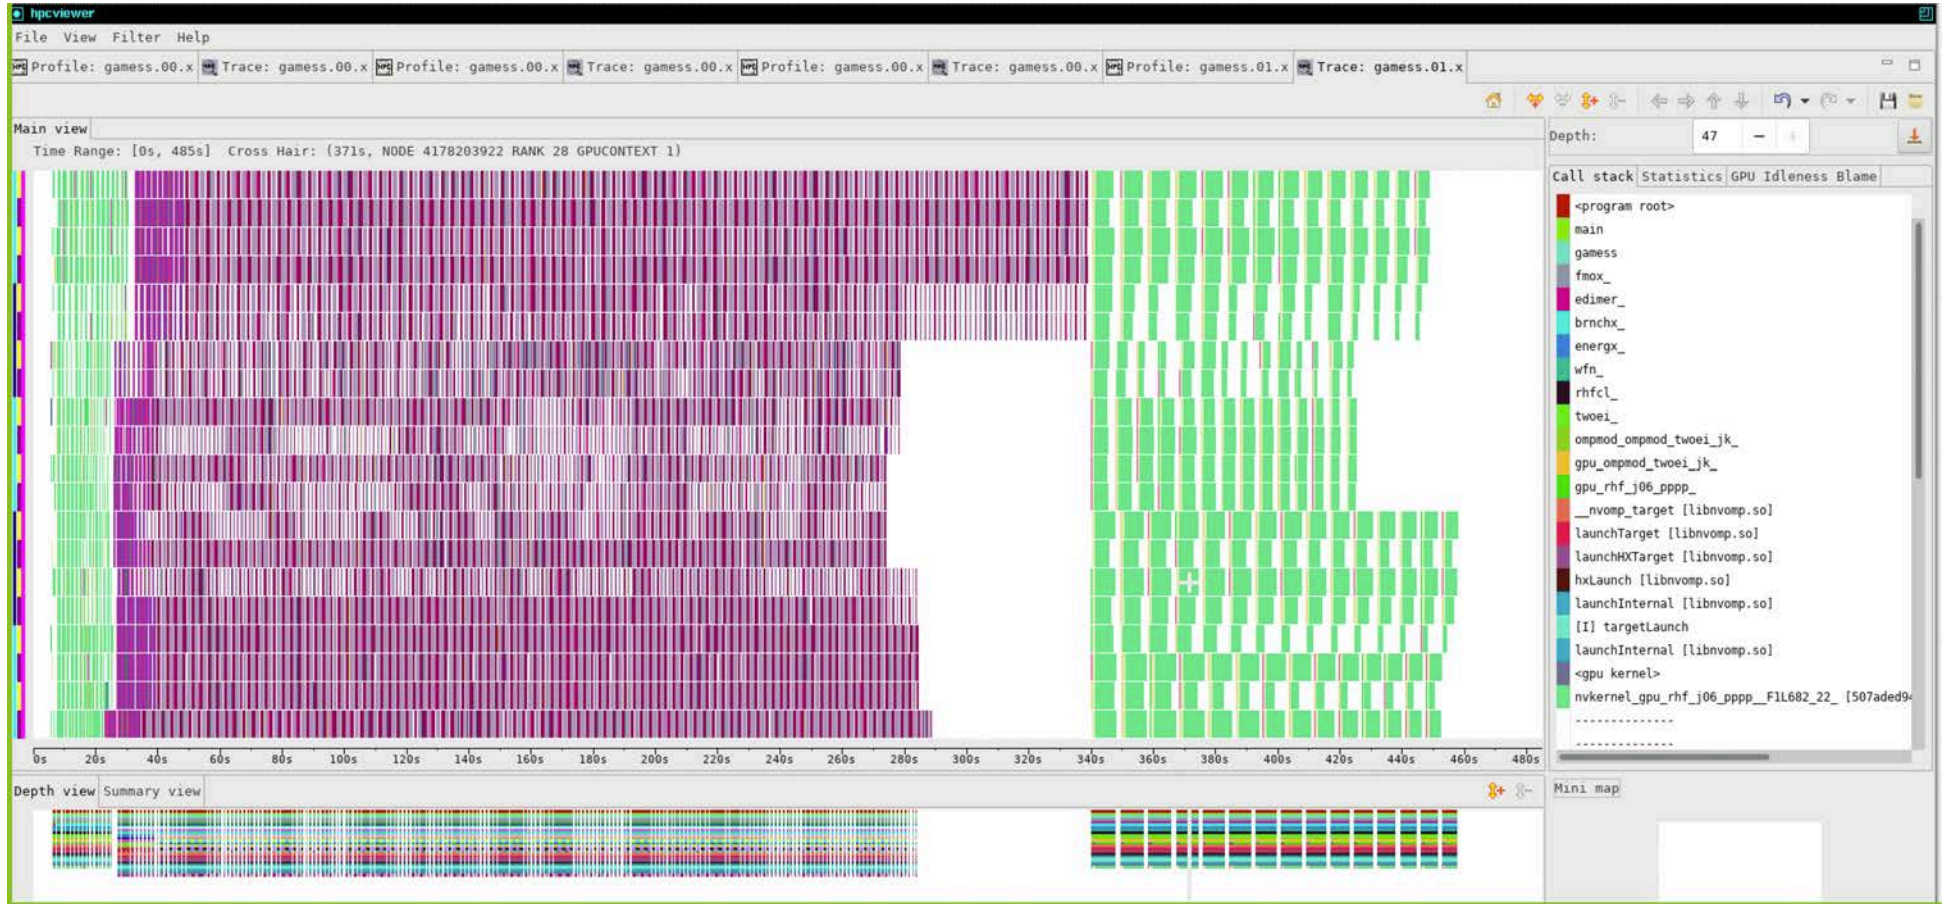
Time-centric Analysis: GAMESS 5 nodes, 40 ranks, 20 GPUs on Perlmutter



Time-centric Analysis: GAMESS 5 nodes, 40 ranks, 20 GPUs on Perlmutter



Time-centric Analysis: GAMESS 5 nodes, 40 ranks, 20 GPUs on Perlmutter



Quicksilver: Detailed analysis within a Kernel using PC Sampling

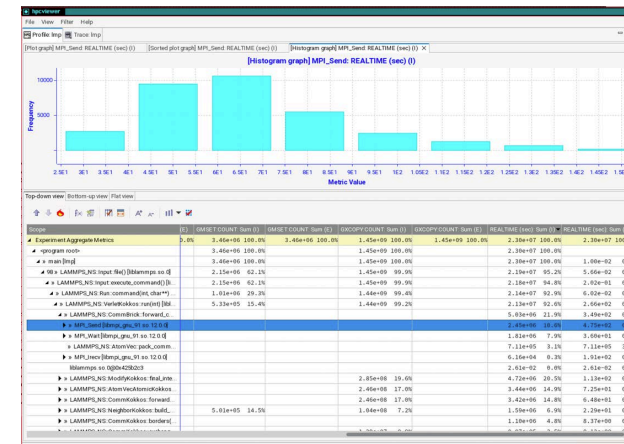
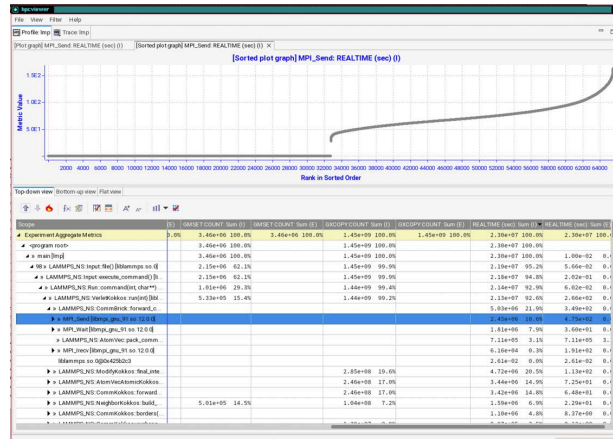
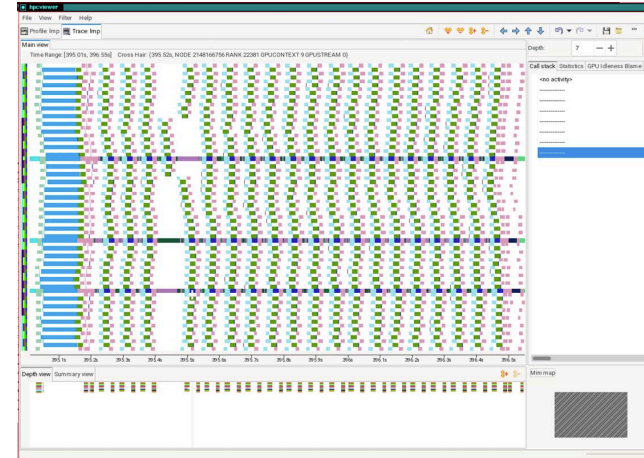
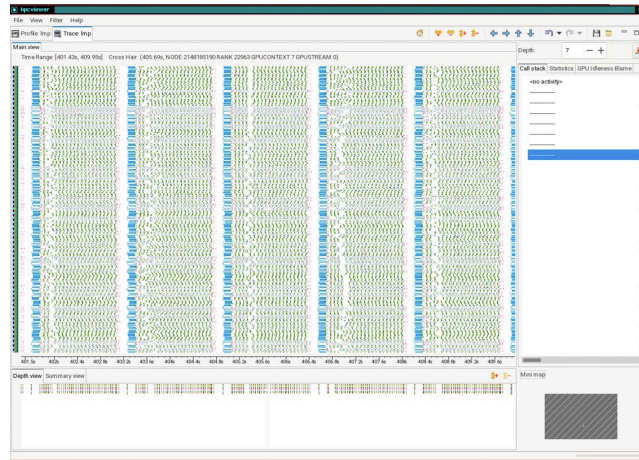
The screenshot shows the hpcviewer interface. The top pane displays C++ code from CollisionEvent.cc, with line 73 highlighted: `currentCrossSection = macroscopicCrossSection(monteCarlo, reactIndex, mc.particle.domain, mc.particle.cell, isoIndex, mc.particle.energy_group);`. The bottom pane shows a performance table with columns for various GINS metrics. A red box highlights the 'Scope' column, and a blue box highlights the row for `macroscopicCrossSection`.

Scope	GINS: Sum (I)	GINS: Sum (E)	GINS:STL_ANY: Sum (I)	GINS:STL_ANY: Sum (E)	GINS:STL_IFET: Sum (I)	GINS:STL_IFET: Sum (E)	GINS:STL_IDEP:
14 » [I] cudaLaunchKernel<char>	1.30e+11 100.0%		1.19e+11 100.0%		5.27e+09 100.0%		9.34e+
211 » cudaLaunchKernel [qs]	1.30e+11 100.0%		1.19e+11 100.0%		5.27e+09 100.0%		9.34e+
<gpu kernel>	1.30e+11 100.0%		1.19e+11 100.0%		5.27e+09 100.0%		9.34e+
CycleTrackingKernel(MonteCarlo*, int, ParticleVault*, ParticleVau...	1.30e+11 100.0%	4.08e+07 0.0%	1.19e+11 100.0%	3.62e+07 0.0%	5.27e+09 100.0%	2.11e+07 0.4%	9.34e+
132 » CycleTrackingGuts(MonteCarlo*, int, ParticleVault*, Particle...	1.30e+11 100.0%	9.03e+09 7.0%	1.19e+11 100.0%	9.01e+09 7.6%	5.24e+09 99.5%	8.98e+06 0.2%	9.32e+
26 » [I] CycleTrackingFunction(MonteCarlo*, MC_Particle&, int, P...	8.36e+10 64.4%	4.12e+08 0.3%	7.25e+10 61.1%	3.65e+08 0.3%	5.21e+09 98.9%	1.02e+08 1.9%	9.25e+
loop at CycleTracking.cc: 118	8.35e+10 64.3%	3.76e+08 0.3%	7.25e+10 61.1%	3.34e+08 0.3%	5.21e+09 98.8%	9.90e+07 1.9%	9.24e+
63 » CollisionEvent(MonteCarlo*, MC_Particle&, unsigned int) [...]	5.20e+10 40.1%	4.99e+09 3.8%	4.44e+10 37.4%	4.02e+09 3.4%	3.85e+09 73.1%	4.89e+08 9.3%	6.37e+
loop at CollisionEvent.cc: 67	4.09e+10 31.5%	8.15e+08 0.6%	3.42e+10 28.8%	6.54e+08 0.6%	3.54e+09 67.1%	1.27e+08 2.4%	5.67e+
loop at CollisionEvent.cc: 71	3.85e+10 29.6%	2.70e+09 2.1%	3.22e+10 27.1%	2.06e+09 1.7%	3.27e+09 62.0%	2.28e+08 4.3%	5.33e+
73 » macroscopicCrossSection(MonteCarlo*, int, int, int, i...	3.58e+10 27.5%	1.22e+10 9.4%	3.01e+10 25.4%	9.85e+09 8.3%	3.04e+09 57.7%	1.79e+09 33.9%	4.60e+
41 » NuclearData::getReactionCrossSection(unsigned int, u...	2.09e+10 16.1%	1.09e+10 8.4%	1.79e+10 15.1%	9.42e+09 7.9%	1.26e+09 23.8%	6.68e+08 12.7%	2.19e+
253 » [I] NuclearDataReaction::getCrossSection(unsigned ...	6.89e+09 5.3%	3.77e+09 2.9%	5.86e+09 4.9%	3.32e+09 2.8%	2.25e+08 4.3%	8.24e+07 1.6%	8.86e+
NuclearData.cc: 253	6.28e+09 4.8%	6.28e+09 4.8%	5.66e+09 4.8%	5.66e+09 4.8%	4.76e+08 9.0%	4.76e+08 9.0%	6.11e+
NuclearData.cc: 251	1.85e+09 1.4%	1.85e+09 1.4%	1.64e+09 1.4%	1.64e+09 1.4%	8.12e+07 1.5%	8.12e+07 1.5%	2.47e+
NuclearData.cc: 248	1.61e+09 1.2%	1.61e+09 1.2%	1.18e+09 1.0%	1.18e+09 1.0%	1.10e+08 2.1%	1.10e+08 2.1%	3.62e+
252 » [I] qs_vector<NuclearDataSpecies>::operator[](int)	1.29e+09 1.0%	1.29e+09 1.0%	1.14e+09 1.0%	1.14e+09 1.0%	7.37e+04 0.0%	7.37e+04 0.0%	1.24e+
NuclearData.cc: 252	1.12e+09 0.9%	1.12e+09 0.9%	9.48e+08 0.8%	9.48e+08 0.8%	3.44e+05 0.0%	3.44e+05 0.0%	2.50e+
252 » [I] qs_vector<NuclearDataReaction>::size() const	9.41e+08 0.7%	9.41e+08 0.7%	8.17e+08 0.7%	8.17e+08 0.7%			4.63e+

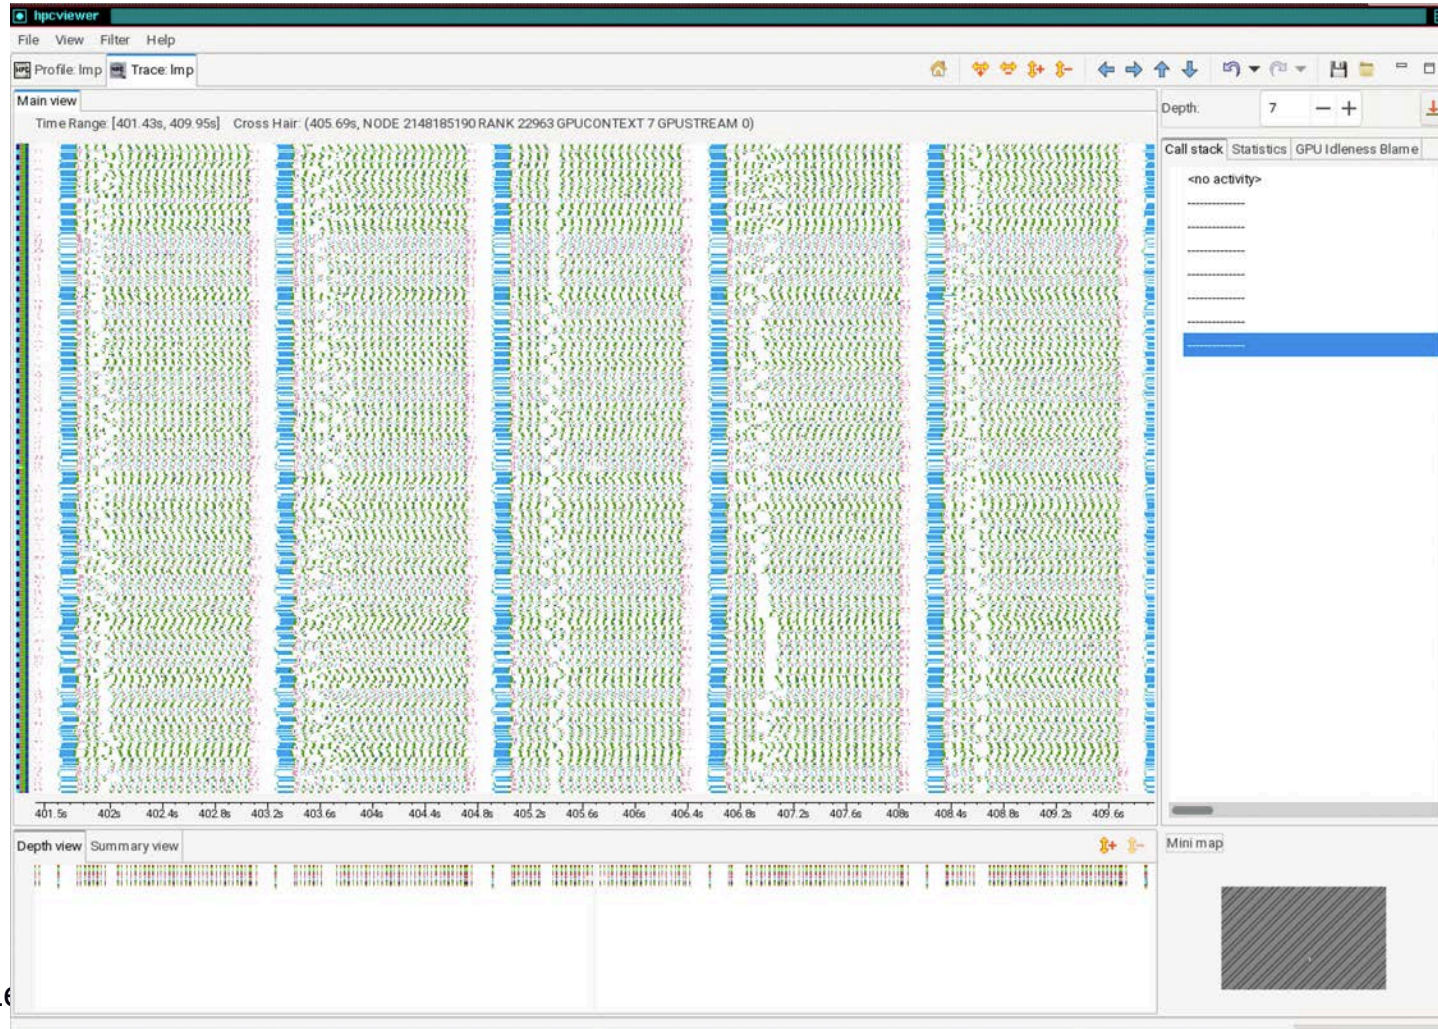
Quicksilver: Detailed analysis within a Kernel using PC Sampling

```
Scope
  ▲ 14 » [I] cudaLaunchKernel<char>
    ▲ 211 » cudaLaunchKernel [qs]
      ▲ » <gpu kernel>
        ▲ » CycleTrackingKernel(MonteCarlo*, int, ParticleVault*, ParticleVau...
          ▲ 132 » CycleTrackingGuts(MonteCarlo*, int, ParticleVault*, Particle...
            ▲ 26 » [I] CycleTrackingFunction(MonteCarlo*, MC_Particle&, int, P...
              ▲ loop at CycleTracking.cc: 118
                ▲ 63 » CollisionEvent(MonteCarlo*, MC_Particle&, unsigned int) [...
                  ▲ loop at CollisionEvent.cc: 67
                    ▲ loop at CollisionEvent.cc: 71
                      ▲ 73 » macroscopicCrossSection(MonteCarlo*, int, int, int, i...
                        ▲ 41 » NuclearData::getReactionCrossSection(unsigned int, u...
                          ▶ 253 » [I] NuclearDataReaction::getCrossSection(unsigned ...
                            NuclearData.cc: 253
                            NuclearData.cc: 251
                            NuclearData.cc: 248
                          ▶ 252 » [I] qs_vector<NuclearDataSpecies>::operator[](int)
                            NuclearData.cc: 252
                          ▶ 252 » [I] qs_vector<NuclearDataReaction>::size() const
                          ▶ 252 » [I] qs_vector<NuclearDataReaction>::operator[](int)
```

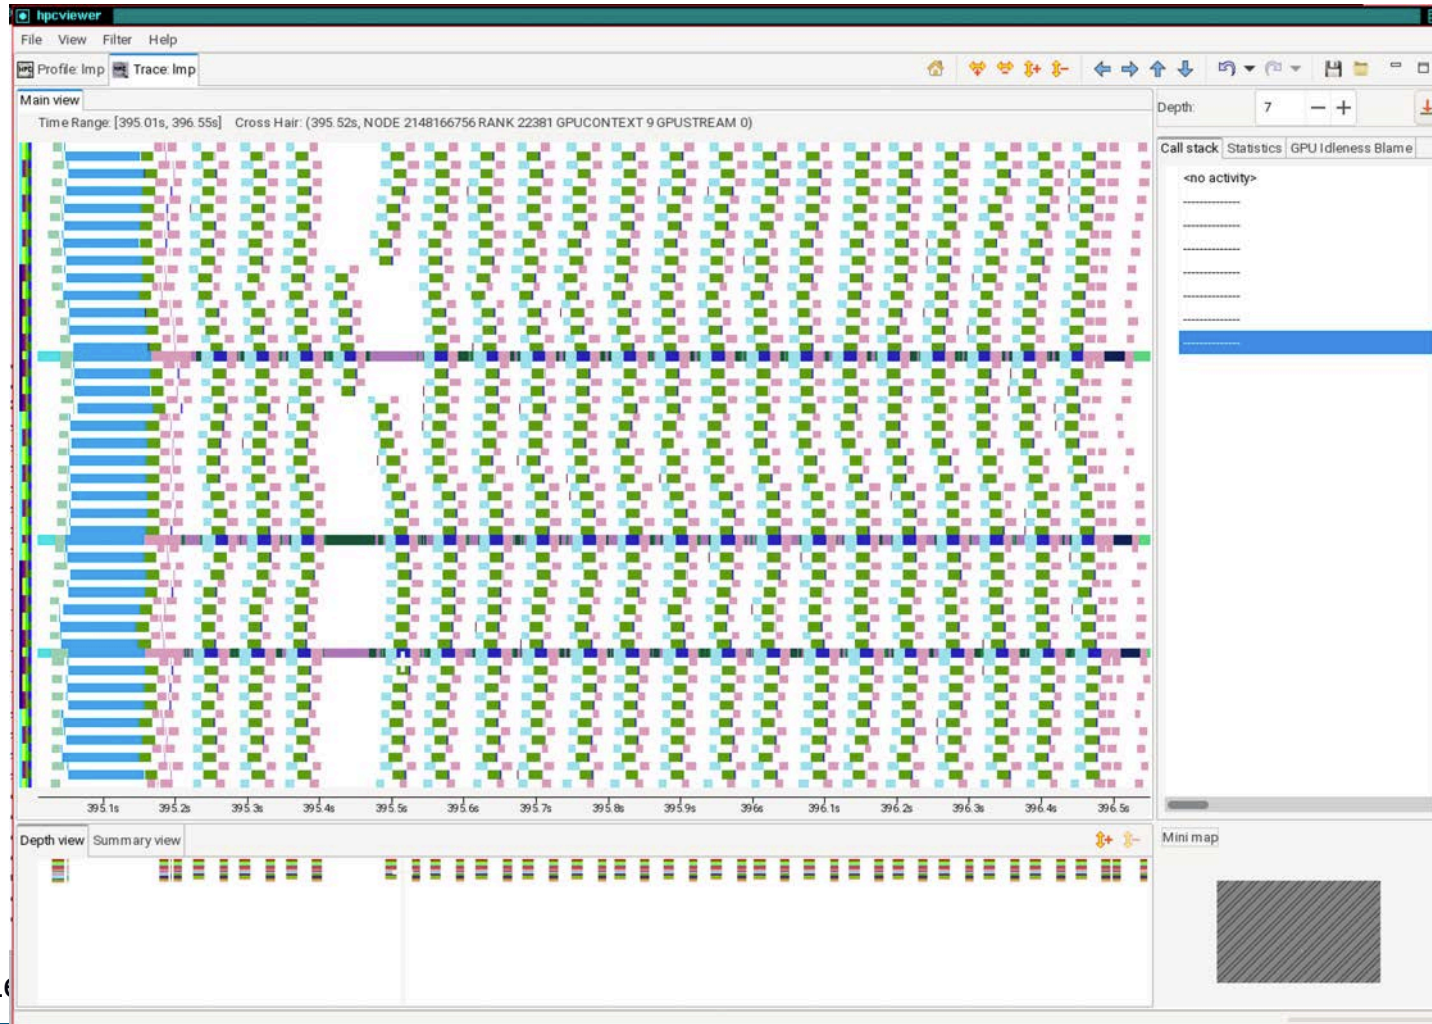

LAMMPS on Frontier: Executions with Kernel Duration of Milliseconds



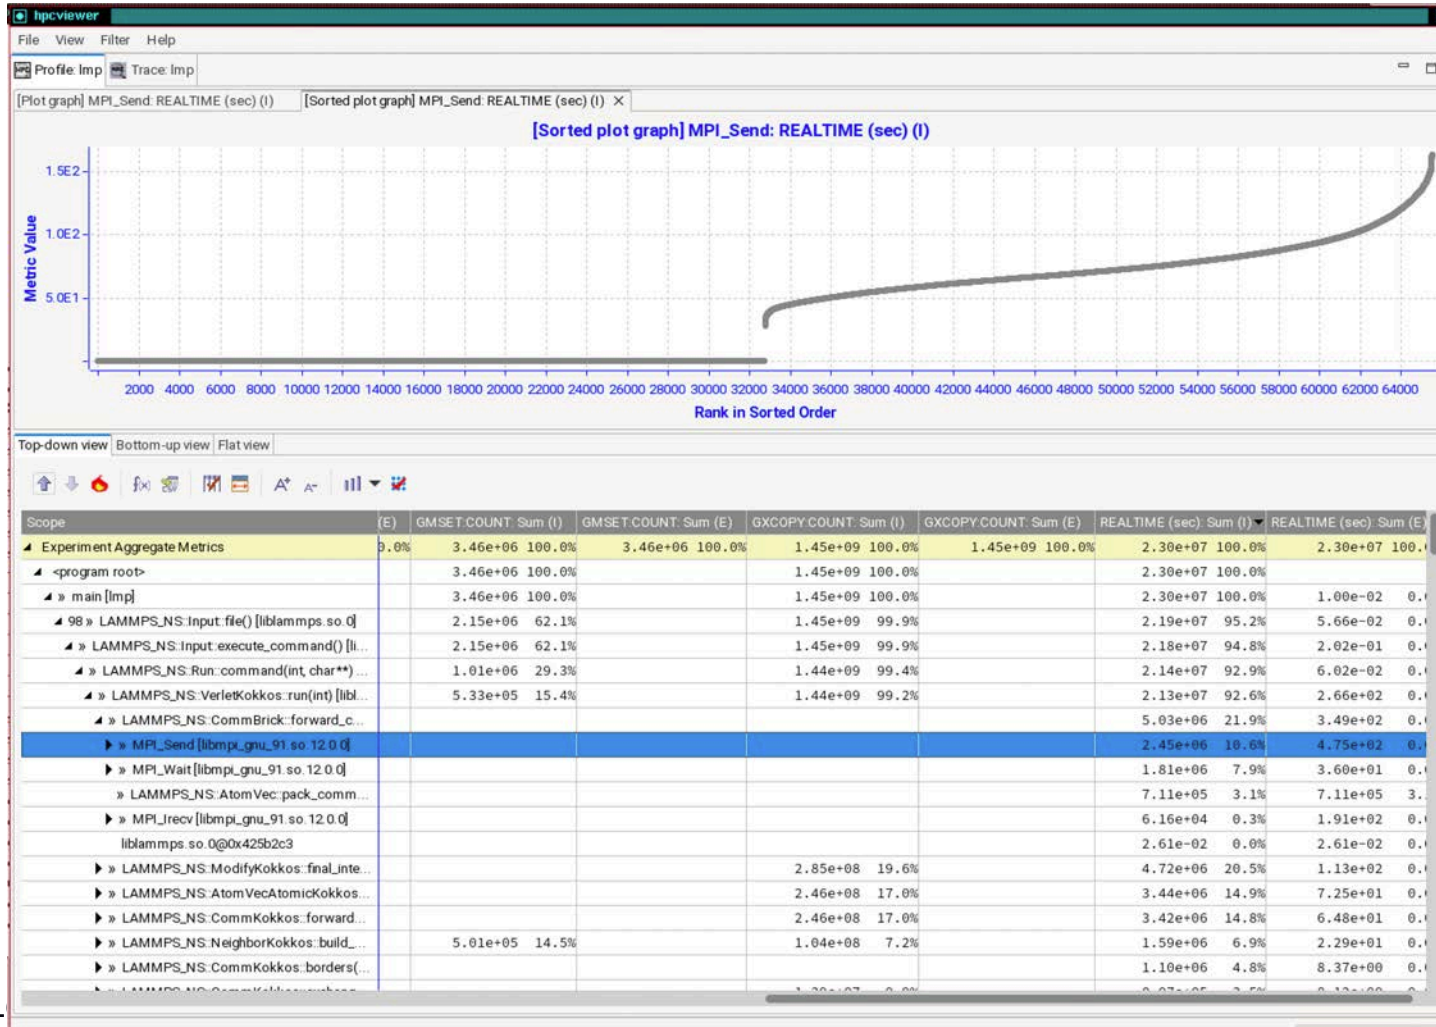
LAMMPS on Frontier: Executions with Kernel Duration of Milliseconds



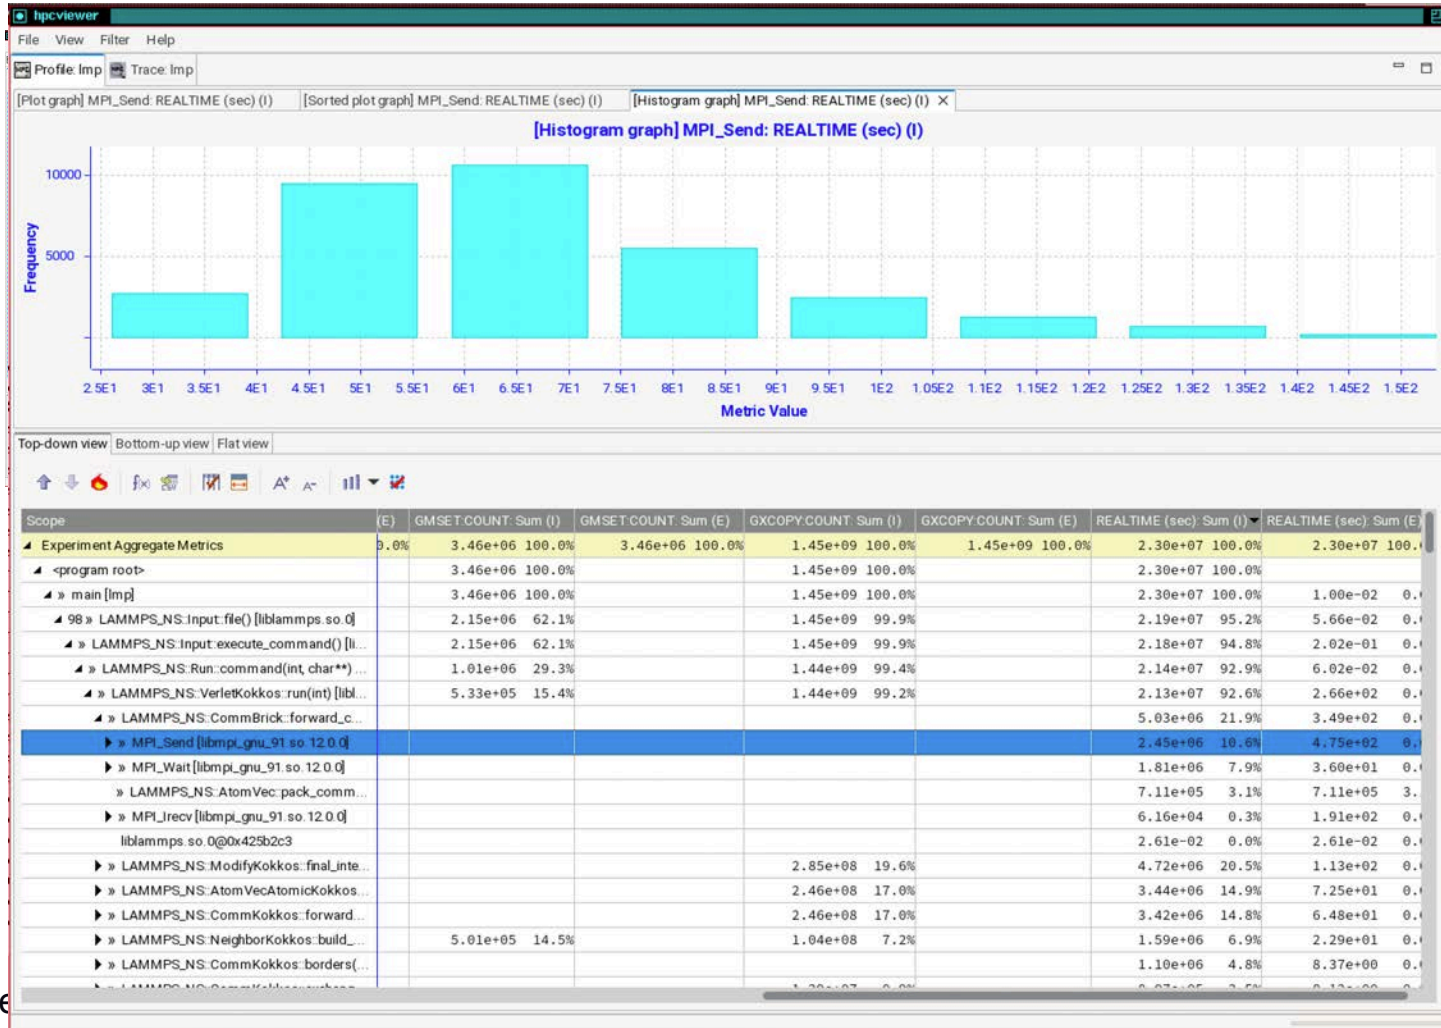
LAMMPS on Frontier: Executions with Kernel Duration of Milliseconds



LAMMPS on Frontier: Executions with Kernel Duration of Milliseconds



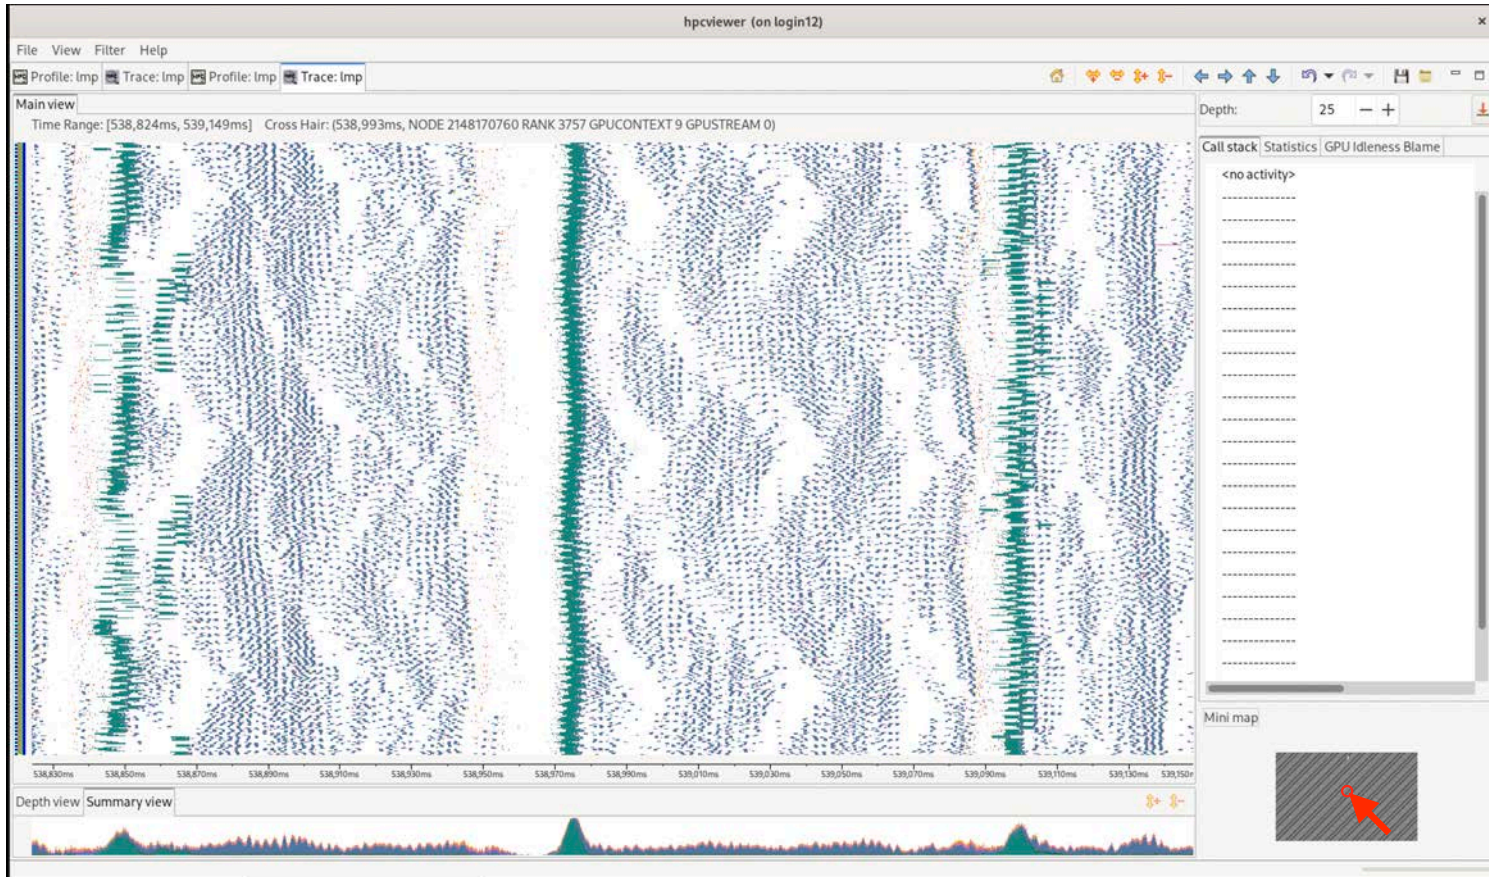
LAMMPS on Frontier: Executions with Kernel Duration of Milliseconds



LAMMPS on Frontier: 8K nodes, 64K MPI ranks + GPU times

Kernel duration of microseconds

Trace: 4.3TB



HPCToolkit Status on AMD, Intel, NVIDIA GPUs

- **Heterogeneous profiles**
- **GPU operation traces**
- **Hardware counters to measure GPU kernels**
- **Instruction-level measurement within GPU kernels**
 - NVIDIA: PC sampling
 - Intel: binary instrumentation
 - AMD: PC sampling coming soon!

HPCToolkit Documentation and Training

- **User Manual**

- <http://www.hpctoolkit.org/manual/HPCToolkit-users-manual.pdf>

- **Installing HPCToolkit's hpcviewer on your Laptop**

- <http://www.hpctoolkit.org/download.html>

- **Training slides and videos: <http://www.hpctoolkit.org/training.html>**

- Introduction to HPCToolkit [[Youtube](#) (13:22)] [[Slides](#)]
- Sampling-based Performance Analysis with HPCToolkit [[Youtube](#) (23:09)] [[Slides](#)]
- Identifying Scalability Bottlenecks with HPCToolkit [[Youtube](#) (19:27)] [[Slides](#)]
- Analyzing GPU-accelerated Applications with HPCToolkit [[Youtube](#) (23:59)] [[Slides](#)]
- Using HPCToolkit to Analyze the Performance of GPU-accelerated Applications [[Youtube](#) (1:27:35)] [[Slides](#)]
- Analyzing GPU-accelerated Applications Using HPCToolkit [[Youtube](#) (34:02)] [[Slides](#)]
- HPCToolkit Graphical User Interface [[Youtube](#) (35:03)] [[Slides](#)]
- Analyzing CPU Applications with HPCToolkit [[Youtube](#) (2:24:30)] [[Slides](#)]

- **Downloading and Installing HPCToolkit**

- <http://www.hpctoolkit.org/software-instructions.html>

Want Some Help?

- **Join our hpctoolkit-ECP Slack workspace**
 - https://join.slack.com/t/hpctoolkit-ecp/shared_invite/zt-24rtkvwma-4HNYe~TiwFwEiJpH~RqUuw
- **Send email to our mailing list**
 - hpctoolkit-forum@rice.edu

Hands-on Directions

- **Log into Polaris with X11 forwarding for hpcviewer**
 - `ssh -Y <username>@polaris.alcf.anl.gov`
- **Download example programs to measure and analyze**
 - `git clone https://github.com/hpctoolkit-tutorial-examples`
- **Load HPCToolkit into your environment**
 - `module use /soft/perftools/hpctoolkit/polaris/modulefiles`
 - `module load hpctoolkit`
- **Set up environment variables for running HPCToolkit examples at the workshop**
 - `export HPCTOOLKIT_TUTORIAL_PROJECTID=fallwkshp23`
 - `export HPCTOOLKIT_TUTORIAL_RESERVATION=fallws23single`
 - `export HPCTOOLKIT_HPCSTRUCT_CACHE=$HOME/.hpctoolkit/hpcstruct-cache`

Working with HPCToolkit's Tutorial Examples

- **Examples in** `hpctoolkit-tutorial-examples/examples/gpu`
 - Quicksilver (highly recommended)
 - LAGHOS
 - MiniQMC
 - PeleC
 - LAMMPS
- **Working with an example**
 - `cd <example-name>`
 - `source setup-env/polaris.sh`
 - `make build`
 - `make` *# enumerates commands for running jobs and inspecting their results*
- **Pro tip**
 - `watch ls -l # wait until a done file appears, e.g. log.run-pc.done` to indicate your data is ready
 - also look at the `log*.stderr` and `log*.out` logs