











DOE's GPU-Accelerated Exascale Platforms







- Frontier compute nodes (OLCF)
 - 1 AMD EPYC "Trento" CPU
 - 4 MI250X AMD Radeon Instinct GPUs
 - 4 Slingshot 11 endpoints
 - Unified memory architecture
- Aurora compute nodes (ALCF)
 - 2 Intel Xeon "Sapphire Rapids" processors
 - 6 Intel Data Center GPU Max 1500
 - 8 Slingshot 11 endpoints
 - Unified memory architecture
- El Capitan compute nodes (LLNL)
 - 4 AMD MI300 APU
 - 4 Slingshot 11 endpoints
 - Unified memory architecture

Outline

- Introduction to HPCToolkit performance tools
 - —Overview of HPCToolkit components and their workflow
 - —HPCToolkit's graphical user interfaces
- Analyzing the performance of GPU-accelerated codes with HPCToolkit
 - Slides: Exawind (AMReX)
 - —Slides: LAMMPS at Exascale (Kokkos)
 - —Demo: GAMESS (OpenMP)
 - —Hands-on: Quicksilver (CUDA)
 - —Hands-on: TeaLeaf



Linux Foundation's HPCToolkit Performance Tools

Collect profiles and traces of unmodified parallel CPU and GPU-accelerated applications

Understand where an application spends its time and why

call path profiles associate metrics with application source code contexts

analyze instruction-level performance within GPU kernels and attribute it to your source code

hierarchical traces to understand execution dynamics

Parallel programming models

across nodes: MPI, SHMEM, UPC++, ...

within nodes: OpenMP, Kokkos, RAJA, HIP, DPC++, Sycl, CUDA, OpenACC, ...

Languages

C, C++, Fortran, Python, ...

Hardware

CPU cores and GPUs within a node

CPU: x86_64, Power, ARM

GPU: NVIDIA, AMD, Intel



Why HPCToolkit?

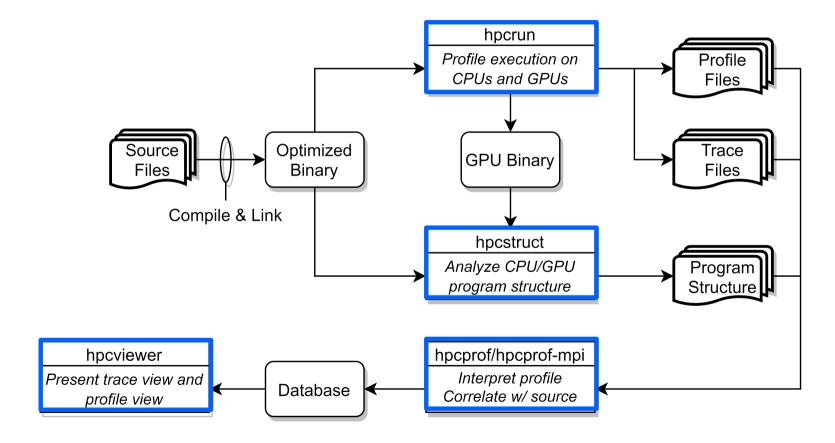
- Measure and analyze performance of CPU and GPU-accelerated applications
- Easy: profile unmodified application binaries
- Fast: low-overhead measurement
- Informative: understand where an application spends its time and why
 - —call path profiles associate metrics with application source code contexts
 - —optional hierarchical traces to understand execution dynamics
- Broad audience
 - —application developers
 - —framework developers
 - —runtime and tool developers
- Unlike vendor tools works with a wide range of CPUs and GPUs



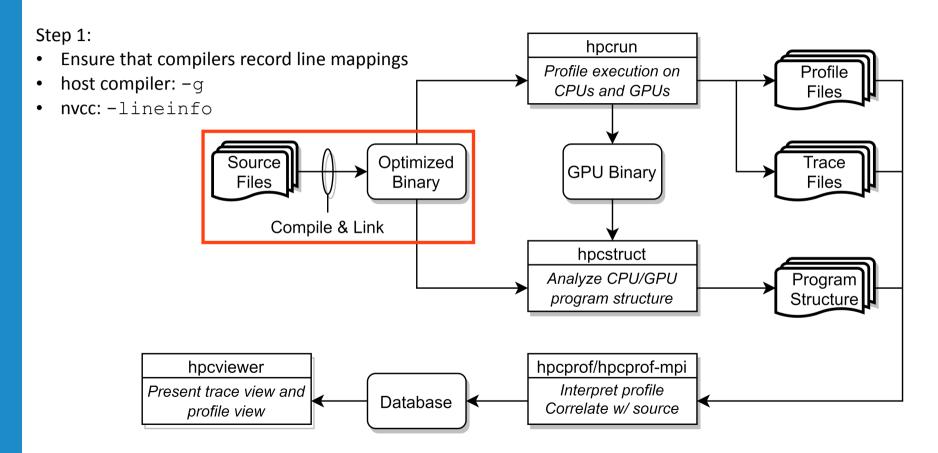
How does HPCToolkit Differ from NVIDIA's Tools?

- NVIDIA NSight Systems
 - —tracing of CPU and GPU streams
 - —analyze traces when you open them with the GUI
 - long running traces are huge and thus extremely slow to analyze, limiting scalability
 - —designed for measurement and analysis within a node
- NVIDIA NSight Compute
 - —detailed measurement of kernels with counters and execution replay
 - —very slow measurement
 - —flat display of measurements within GPU kernels
- HPCToolkit
 - —supports more scalable tracing than Nsight Systems
 - measure exascale executions across many GPUs and nodes
 - —scalable, parallel post-mortem analysis vs. non-scalable in-GUI analysis
 - —detailed reconstruction of estimates for calling context profiles within GPU kernels











Step 2: hpcrun hpcrun collects call path profiles (and Profile execution on **Profile** optionally, traces) of events of interest CPUs and GPUs Files Optimized Source Trace **GPU** Binary Binary Files **Files** Compile & Link hpcstruct Analyze CPU/GPU Program program structure Structure | hpcprof/hpcprof-mpi hpcviewer Interpret profile Present trace view and Database Correlate w/ source profile view



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
- PC Samples: NVIDIA (in progress: AMD, Intel)
- Binary instrumentation of GPU kernels on Intel GPUs for fine-grain measurement



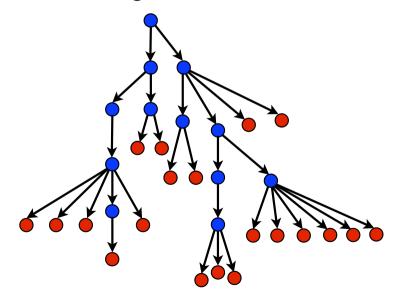
Call Stack Unwinding to Attribute Costs in Context

- Unwind when timer or hardware counter overflows
 - —measurement overhead proportional to sampling frequency rather than call frequency
- Unwind to capture context for events such as GPU kernel launches

Call path sample

return address
return address
return address
instruction pointer

Calling context tree

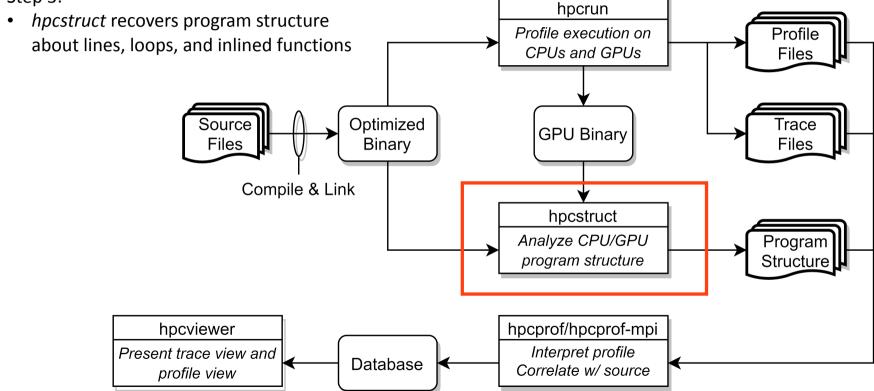




hpcrun: Measure CPU and/or GPU activity



Step 3:





hpcstruct: Analyze CPU and GPU Binaries Using Multiple Threads

Usage

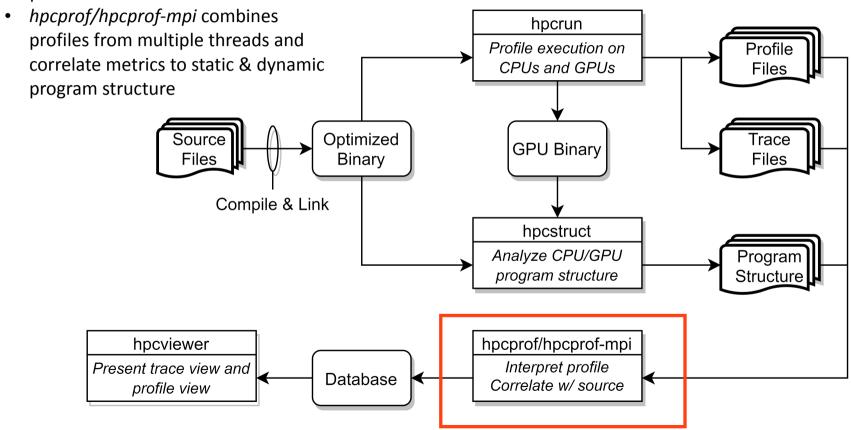
```
hpcstruct [--gpucfg yes] <measurement-directory>
```

- What it does
 - Recover program structure information
 - Files, functions, inlined templates or functions, loops, source lines
 - In parallel, analyze all CPU and GPU binaries that were measured by HPCToolkit
 - —typically analyze large application binaries with 16 threads
 - —typically analyze multiple small application binaries concurrently with 2 threads each
 - Cache binary analysis results for reuse when analyzing other executions

NOTE: --gpucfg yes needed only for analysis of GPU binaries for interpreting PC samples on NVIDIA GPUs



Step 4:





hpcprof/hpcprof-mpi: Associate Measurements with Program Structure

Analyze data from modest executions with multithreading (moderate scale)

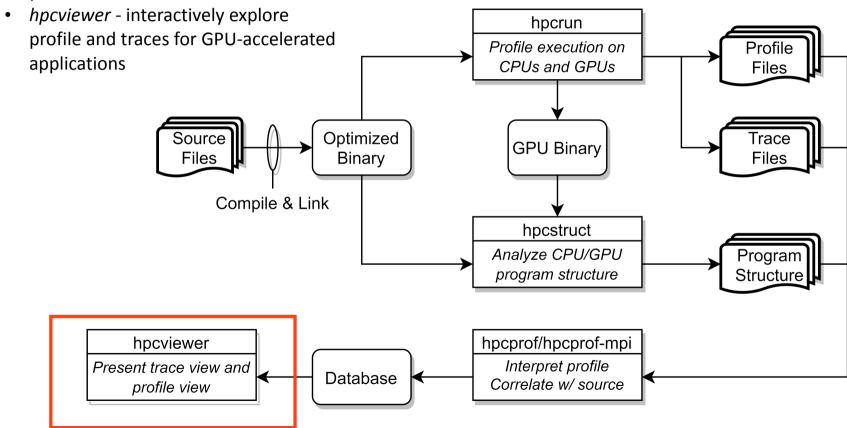
```
hpcprof <measurement-directory>
```

Analyze data from large executions with distributed-memory parallelism + multithreading (large scale)

```
mpiexec -n ${NODES} --ppn 1 -depth=128 \
    hpcprof-mpi <measurement-directory>
```

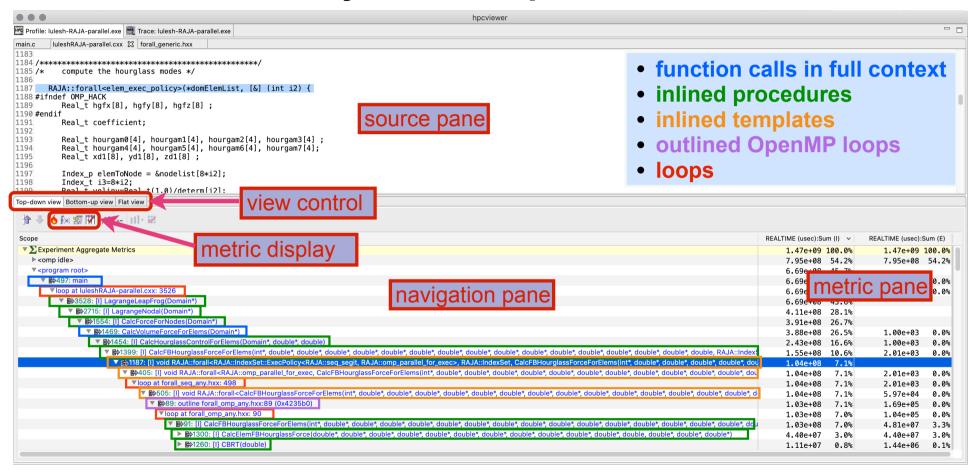


Step 4:





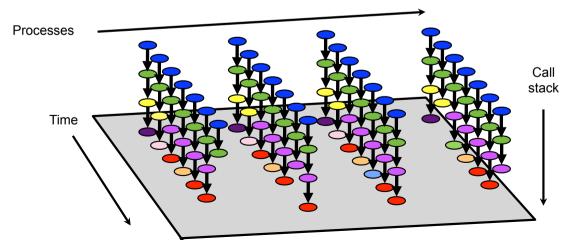
Code-centric Analysis with hpcviewer





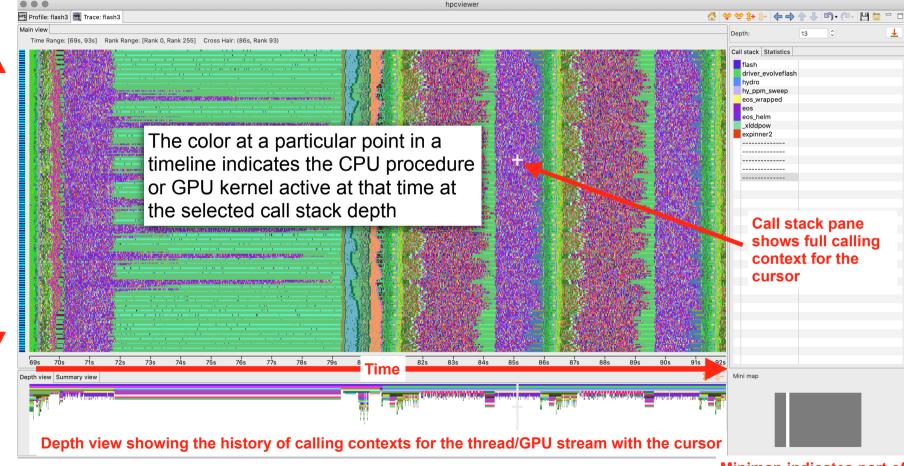
Understanding Temporal Behavior

- 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





GPU

Threads,

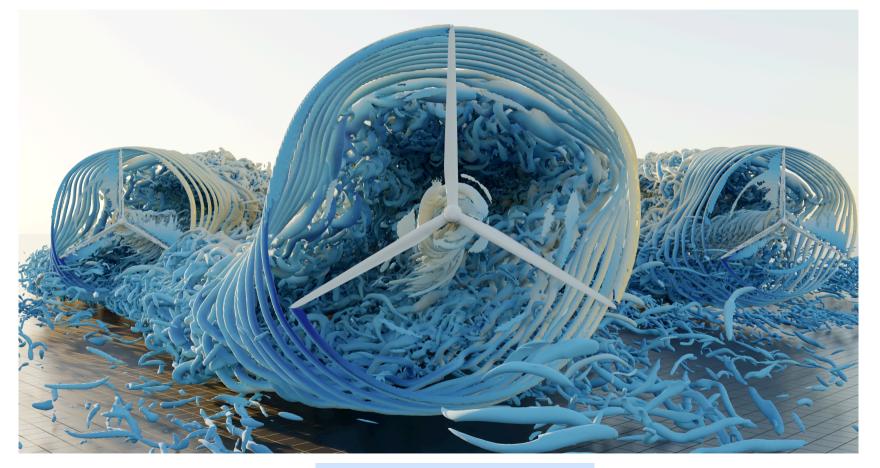
ranks, OpenMP

Case Studies

- ExaWind
- GAMESS (OpenMP)
- Quicksilver (CUDA)
- PeleC (AMReX)
- LAMMPS (Kokkos) at exascale



ExaWind: Wakes from Three Turbines over Time





ExaWind: Visualization of a Wind Farm Simulation





ExaWind: Execution Traces on Frontier Collected with HPCToolkit

Traces on roughly ~70K MPI ranks for ~17minutes

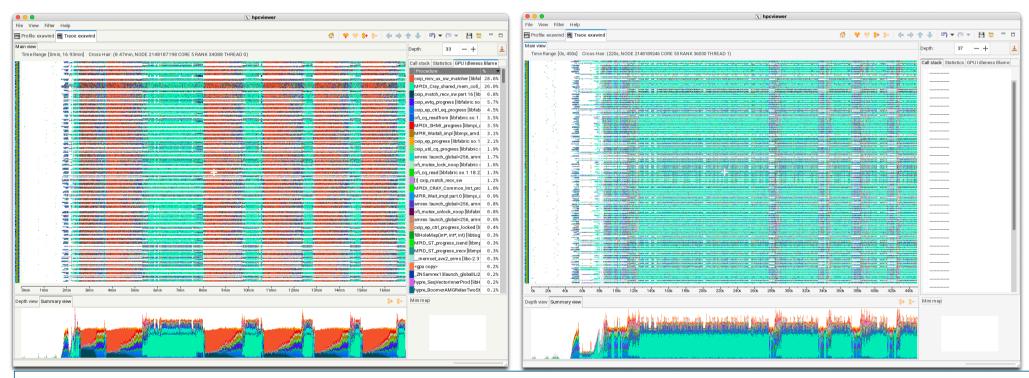


Figure credits: Jon Rood, NREL

*replaced non-blocking send/recv with ialltoallv

ExaWind Testimonials for HPCToolkit

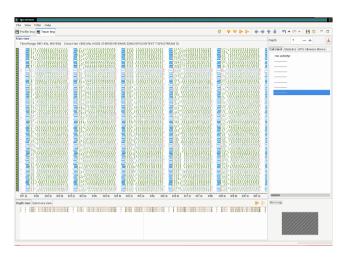
I just wanted to mention we've been using HPCToolkit a lot for our ExaWind application on Frontier, which is a hugely complicated code, and your profiler is one of the only ones we've found that really lets us easily instrument and then browse what our application is doing at runtime including GPUs. As an example, during a recent hackathon we had, we improved our large scale performance by 24x by understanding our code better with HPCToolkit and running it on 1000s of nodes while profiling. We also recently improved upon this by 10% for our total runtime.

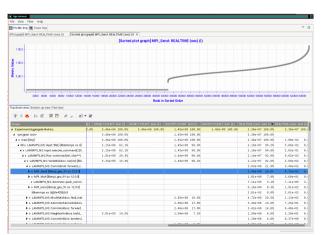
- Jon Rood NREL (5/31/2024)

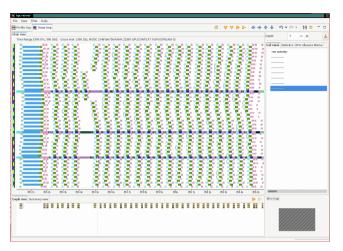
One big thing for us is that we can't overstate how complicated ExaWind is in general, and how complicated it is to build, so finding out that HPCToolkit could easily profile our entire application without a ton of instrumentation during the build process, and be able to profile it on a huge amount of Frontier with line numbers and visualizing the trace was really amazing to us.

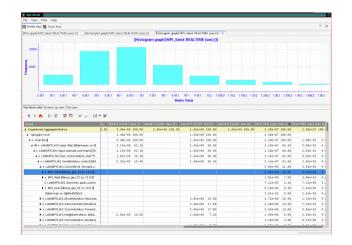
- Jon Rood NREL (6/3/2024)



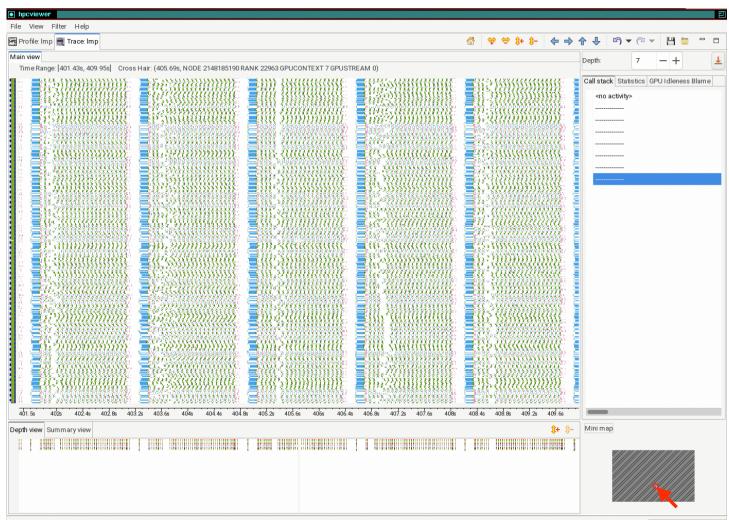




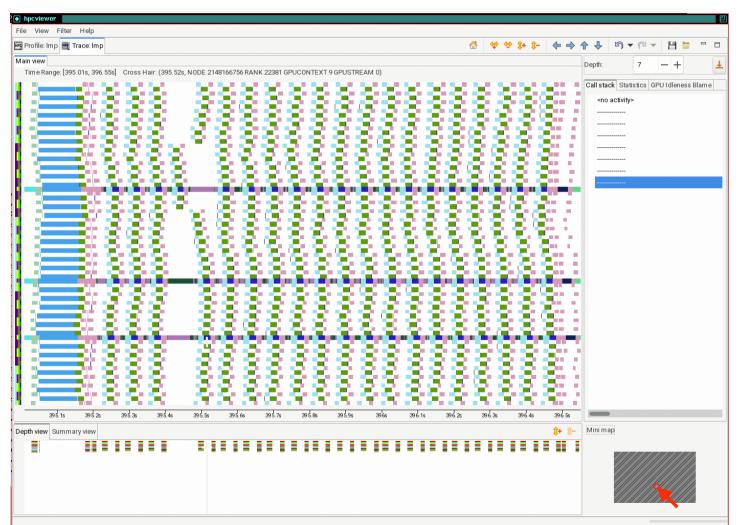




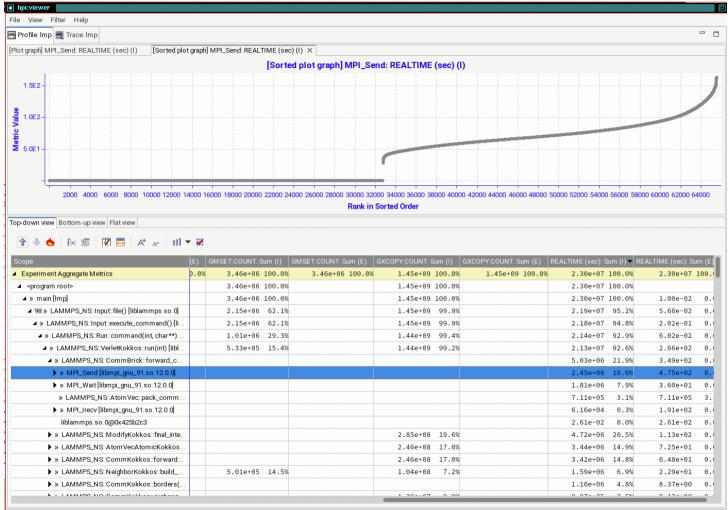




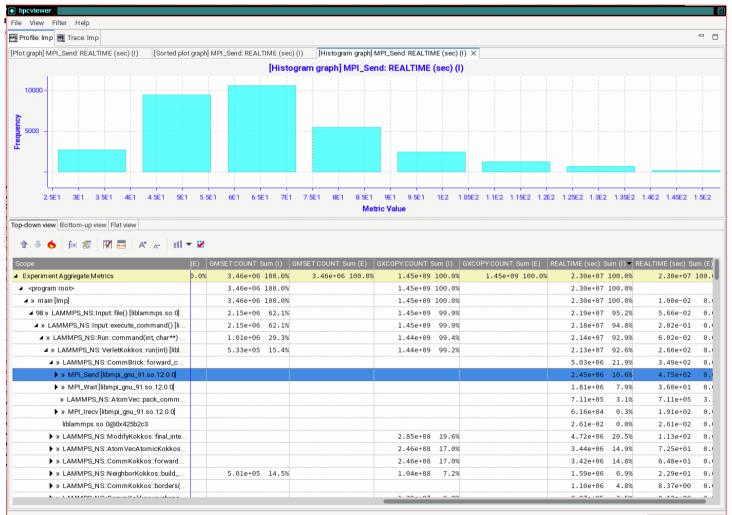








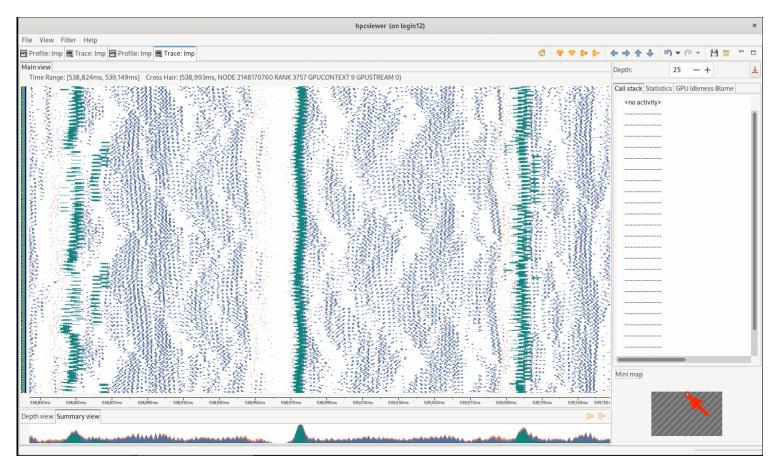






LAMMPS on Frontier: 8K nodes, 64K MPI ranks + 64K GPU tiles

Kernel duration of microseconds





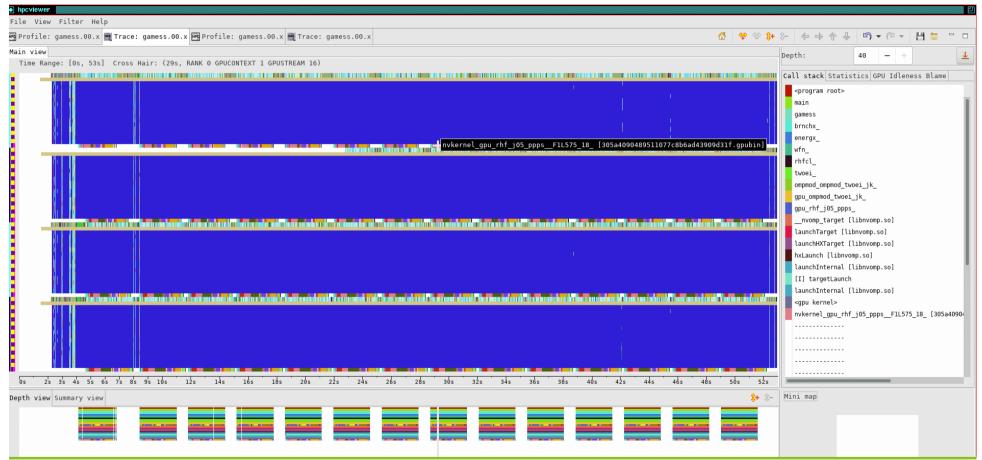
Case Study: GAMESS

- General Atomic and Molecular Electronic Structure System (GAMESS)
 - —general *ab initio* quantum chemistry package
- Calculates the energies, structures, and properties of a wide range of chemical systems
- Experiments
 - GPU-accelerated nodes at a prior Perlmutter hackathon
 - Single node with 4 GPUs
 - Five nodes with 20 GPUs

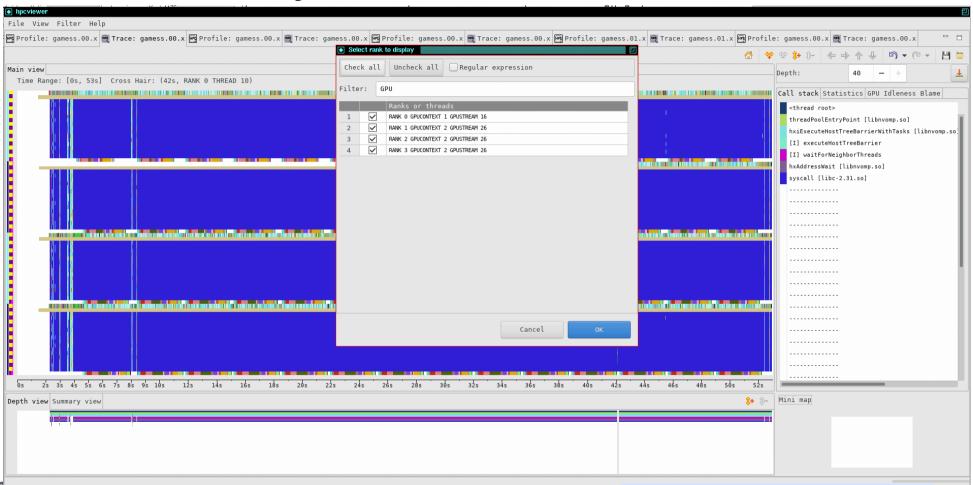
Perlmutter node at a glance

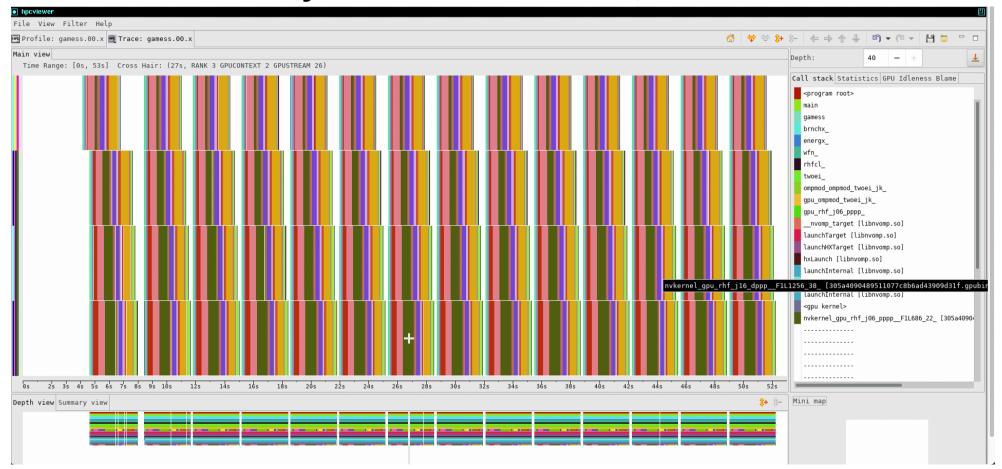
AMD Milan CPU 4 NVIDIA A100 GPUs 256 GB memory



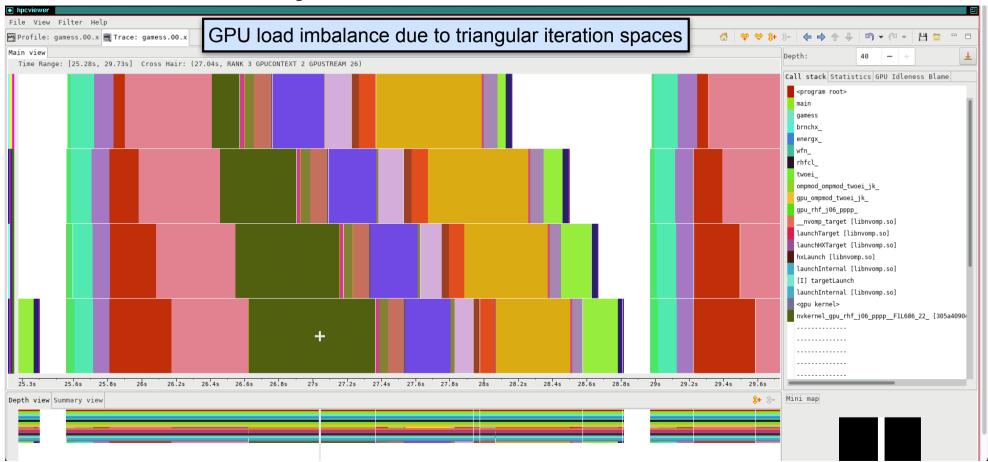






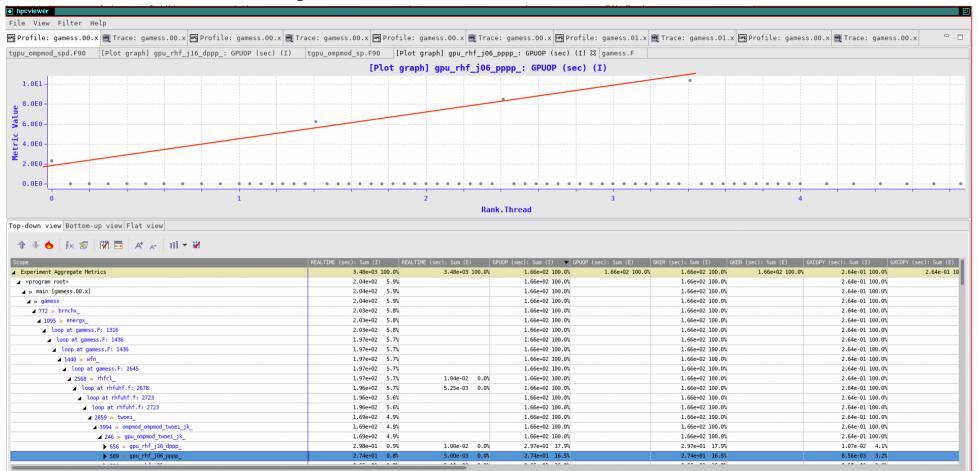




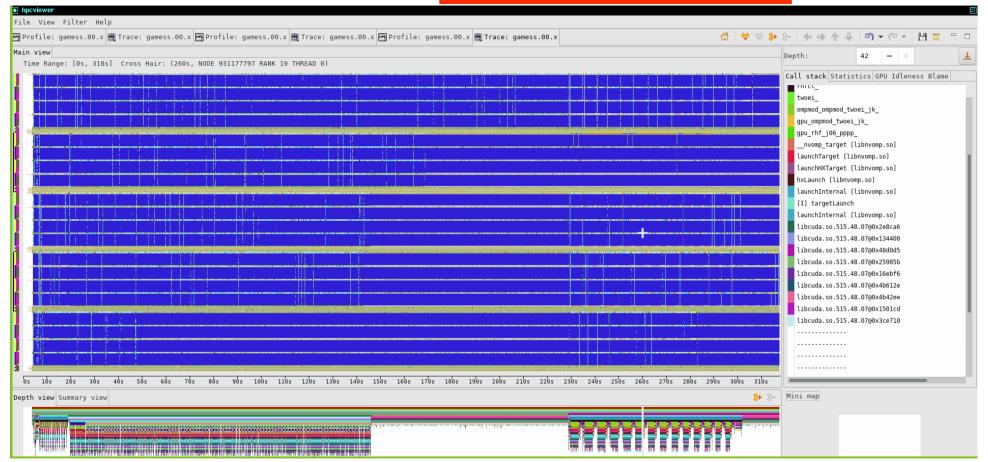




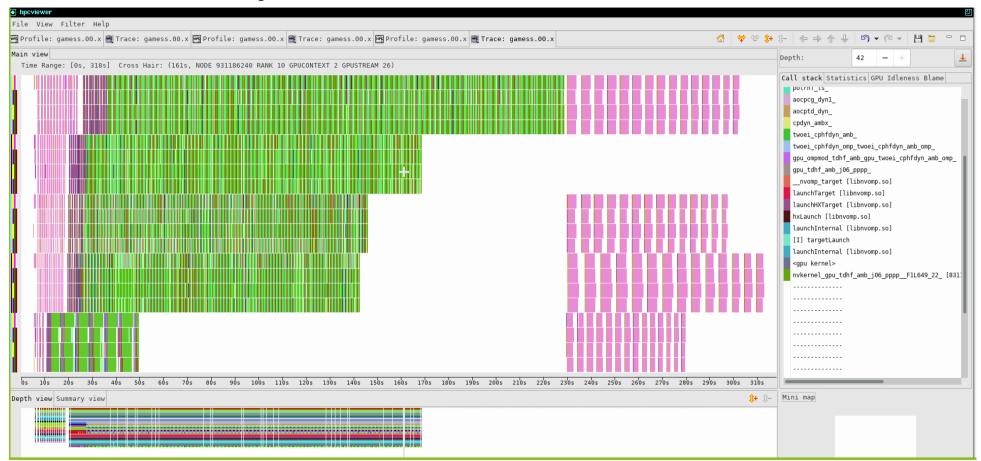
Time-centric Analysis: GAMESS 4 ranks, 4 GPUs on Perlmutter



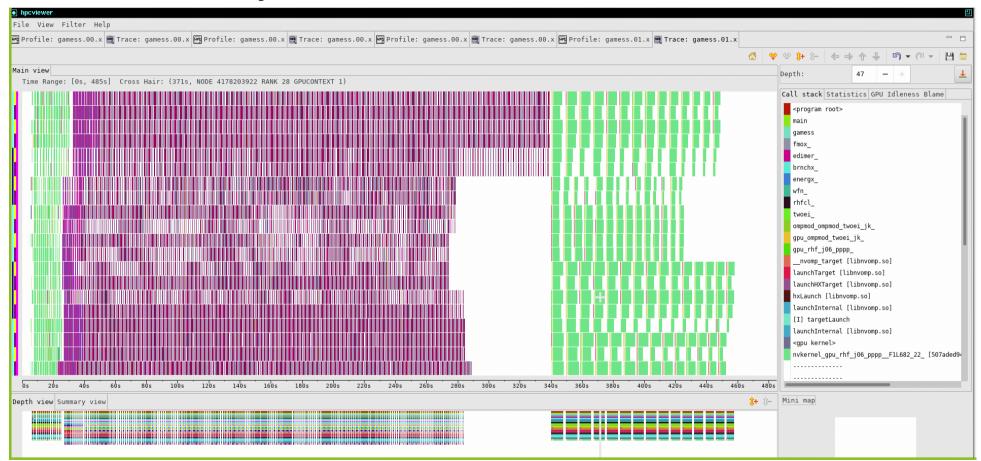




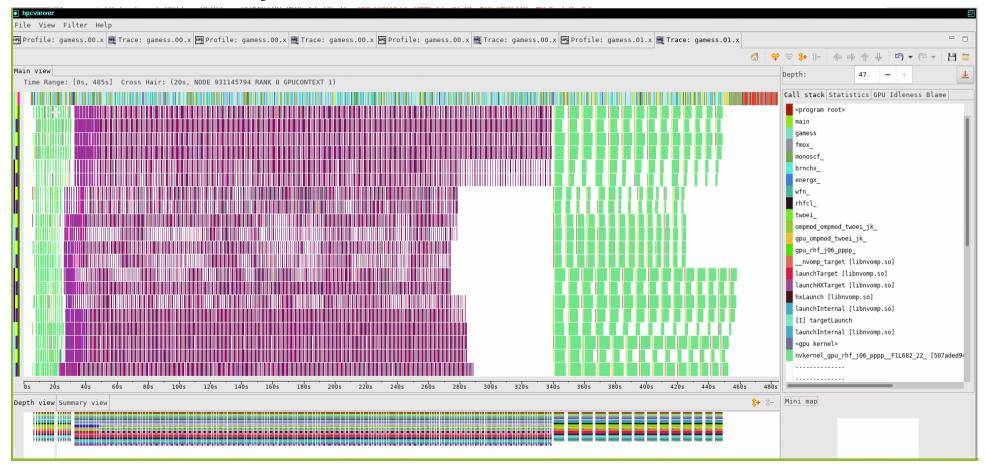




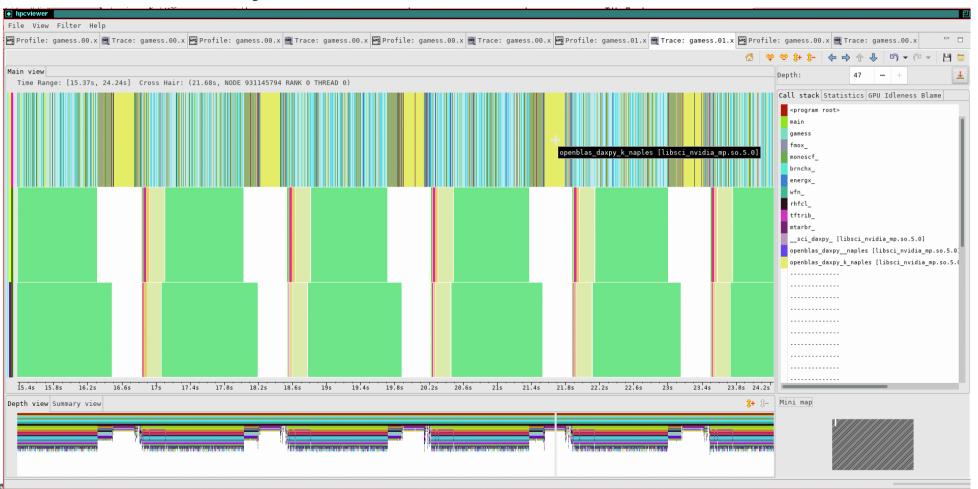












```
File View Filter Help
🚧 Profile: gamess.00.x 🖻 Trace: gamess.00.x 💌 Profile: gamess.00.x 💌 Profile: gamess.00.x 💌 Profile: gamess.00.x 💌 Trace: gamess.00.x
mthlib.f ₩
 1054 C
 L055 C*MODULE MTHLIB *DECK MTARBR
          IMPLICIT DOUBLE PRECISION(A-H, 0-Z)
 1061 C
          DIMENSION A(*),B(NA.MB),AB(NAB.MB)
 1063 C
 1064
         PARAMETER (ZERO=0.0D+00)
 1066 C* 31 OCT 1979
 1067 C*
 L068 C*FUNCTION
                - TO MULTIPLY SYMMETRIC MATRIX A
 1069 C*
                  TIMES RECTANGULAR MATRIX B AND GET RECTANGULAR MATRIX AB
 1070 C*
 1071 C*PARAMETERS
        A - THE INPUT REAL SYMMETRIC MATRIX OF ORDER NA
                 STORED IN SYMMETRIC STOAGE MODE.
                - THE INPUT REAL NA BY MB RECTANGULAR MATRIX
              - THE ORDER OF MATRIX A
              - THE COLUMN DIMENSION OF MATRICES B AND AB
               - THE OUTPUT PRODUCT NA BY MB MATRIX
          NAB - THE INPUT ROW DIMENSION OF MATRIX AB
         INCA - ADDRESS DIFFERENCE OF TWO ADJACENT ELEMENTS OF A
 1081
          INC=INCA
 1082 C
            PROCESS DIAGONAL ELEMENTS OF INPUT MATRIX A
 1084 C
          IJ=1-INC
          DO 120 I=1,NA
            IJ=IJ+I*INC
             AIJ=A(IJ)
            DO 110 K=1,MB
               AB(I,K)=AIJ*B(I,K)
 1091 110 CONTINUE
 1092 120 CONTINUE
         IF(NA.EO.1) RETURN
 1094 C
            PROCESS OFF-DIAGONAL ELEMENTS OF INPUT MATRIX A
          IJ=1-INC
          DO 150 I=2,NA
            IJ=IJ+INC
            TM1=T-1
            DO 140 J=1, IM1
               AIJ=A(IJ)
               IF(AIJ.EQ.ZERO) GO TO 140
                  CALL DAXPY(MB.AIJ.B(I.1).NA.AB(J.1).NAB)
                  CALL DAXPY(MB, AIJ, B(J, 1), NA, AB(I, 1), NAB)
 1108 150 CONTINUE
 1109
          RETURN
Top-down view Bottom-up view Flat view
```



```
File View Filter Help
🚧 Profile: gamess.00.x 🖻 Trace: gamess.00.x 💌 Profile: gamess.00.x 💌 Profile: gamess.00.x 💌 Profile: gamess.00.x 💌 Trace: gamess.00.x
mthlib.f ₩
 1054 C
 L055 C*MODULE MTHLIB *DECK MTARBR
         IMPLICIT DOUBLE PRECISION(A-H, 0-Z)
 1061 C
         DIMENSION A(*),B(NA.MB),AB(NAB.MB)
 1063 C
 1064
         PARAMETER (ZERO=0.0D+00)
 1066 C* 31 OCT 1979
 1067 C*
 L068 C*FUNCTION
                - TO MULTIPLY SYMMETRIC MATRIX A
 1069 C*
                  TIMES RECTANGULAR MATRIX B AND GET RECTANGULAR MATRIX AB
 1070 C*
 1071 C*PARAMETERS
        A - THE INPUT REAL SYMMETRIC MATRIX OF ORDER NA
                 STORED IN SYMMETRIC STOAGE MODE.
                - THE INPUT REAL NA BY MB RECTANGULAR MATRIX
              - THE ORDER OF MATRIX A
              - THE COLUMN DIMENSION OF MATRICES B AND AB
               - THE OUTPUT PRODUCT NA BY MB MATRIX
         NAB - THE INPUT ROW DIMENSION OF MATRIX AB
         INCA - ADDRESS DIFFERENCE OF TWO ADJACENT ELEMENTS OF A
         INC=INCA
 1082 C
            PROCESS DIAGONAL ELEMENTS OF INPUT MATRIX A
 1084 C
         IJ=1-INC
         DO 120 I=1,NA
            IJ=IJ+I*INC
             AIJ=A(IJ)
            DO 110 K=1,MB
               AB(I,K)=AIJ*B(I,K)
 1091 110 CONTINUE
      120 CONTINUE
         IF(NA.EO.1) RETURN
             DROCECC OFF DIACONAL ELEMENTS OF INDUT MATRIX A
         IJ=1-INC
         DO 150 I=2,NA
            IJ=IJ+INC
            TM1=T-1
            DO 140 J=1,IM1
               AIJ=A(IJ)
               IF(AIJ.EQ.ZERO) GO TO 140
                  CALL DAXPY(MB.AIJ.B(I.1),NA.AB(J.1),NAB)
                  CALL DAXPY(MB, AIJ, B(J, 1), NA, AB(I, 1), NAB)
 1108 150 CONTINUE
         RETURN
Top-down view Bottom-up view Flat view
```



```
File View Filter Help
🖷 Profile: gamess.00.x 💌 Trace: gamess.00.x 💌 Profile: gamess.00.x 💌 Trace: gamess.00.x 💌 Profile: gamess.00.x 💌 Trace: gamess.00.x
 1096 C
               IJ=1-INC
 098
               DO 150 I=2.NA
 099
                    IJ=IJ+INC
                    IM1=I-1
                    DO 140 J=1.IM1
                        IJ=IJ+INC
                        AIJ=A(IJ)
                        IF(AIJ.EQ.ZERO) GO TO 140
                            CALL DAXPY(MB.AIJ.B(I.1).NA.AB(J.1).NAB)
                            CALL DAXPY(MB,AIJ,B(J,1),NA,AB(I,1),NAB)
         140
                    CONTINUE
         150 CONTINUE
               RETURN
               END
Top-down view Bottom-up view Flat view
```

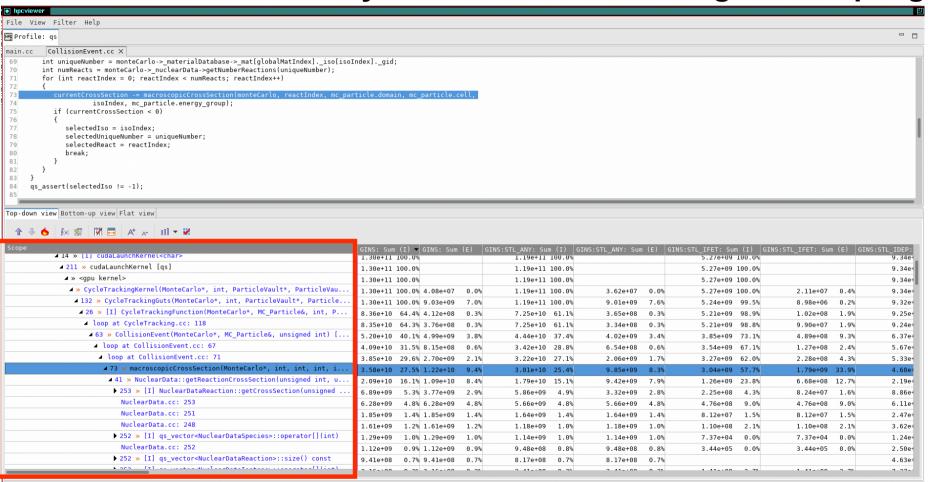


Case Study: Quicksilver

- Proxy application that represents some elements of LLNL's Mercury workload
- Solves a simplified dynamic Monte Carlo particle transport problem
 - Attempts to replicate memory access patterns, communication patterns, and branching or divergence of Mercury for problems using multigroup cross sections
- Parallelization: MPI, OpenMP, and CUDA
- Performance Issues
 - load imbalance (for canned example)
 - latency bound table look-ups
 - a highly branchy/divergent code path
 - poor vectorization potential



Quicksilver: Detailed analysis within a Kernel using PC Sampling





Quicksilver: Detailed analysis within a Kernel using PC Sampling

```
■ 14 » [1] cudaLaunchKernel<char>

■ 211 » cudaLaunchKernel [qs]

→ » <apu kernel>

   A » 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

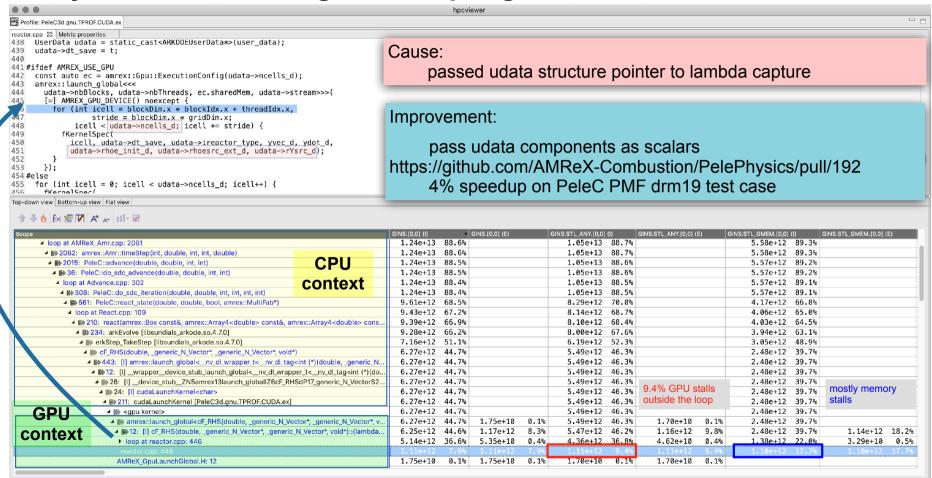
■ loop at CollisionEvent.cc: 71

■ 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
```



Analysis of PeleC using PC Sampling on an NVIDIA GPU





Key Metrics for GPU Kernels

- GPUOP: GPU operation time (kernel launch, copies, etc.)
- GXCOPY:* GPU copies of various kinds
- GKER: GPU kernel time
- GKER:FGP_ACT: fine grain parallelism actual (active warps per SM)
- GKER:FGP MAX: maximum possible fine-grain parallelism (max warps per SM)
- GKER:BLK_THR: threads per block
- GKER:BLK_SM: block shared memory
- GKER:OCC_THR: theoretical thread occupancy



Metrics for GPU Kernels with PC Samples

- GINS: GPU instructions
- GINS:STL ANY: GPU instruction stalls for any reason
- GINS:STL IFET: GPU instruction stalls for instruction fetch
- GINS:STL_GMEM: GPU instruction stalls for global memory
- GINS:STL_CMEM: GPU instruction stalls for constant memory
- GINS:STL_IDEP: GPU instruction stalls for instruction dependences
- GINS:STL_PIPE: GPU instruction pipeline stalls
- GINS:STL_MTHR: GPU instruction stalls for memory throttling
- GSAMP:EXP: expected number of samples
- GSAMP:TOT: total number of samples recorded
- GSAMP:UTIL: GPU utilization = (PC samples expected) / (PC samples total)



HPCToolkit Resources

- Documentation
 - —User manual for HPCToolkit: http://hpctoolkit.org/manual/HPCToolkit-users-manual.pdf
 - —Cheat sheet: https://gitlab.com/hpctoolkit/hpctoolkit/-/wikis/HPCToolkit-cheat-sheet
 - —User manual for hpcviewer: https://hpctoolkit.gitlab.io/hpcviewer
 - —Tutorial videos
 - http://hpctoolkit.org/training.html
 - recorded demo of GPU analysis of Quicksilver: https://youtu.be/vixa3hGDuGg
 - recorded tutorial presentation including demo with GPU analysis of GAMESS: https://vimeo.com/781264043
- Software
 - —Download hpcviewer GUI binaries for your laptop, desktop, cluster, or supercomputer
 - OS: Linux, Windows, MacOS
 - Processors: x86_64, aarch64, ppc64le
 - http://hpctoolkit.org/download.html
 - —Install HPCToolkit on your Linux desktop, cluster, or supercomputer using Spack
 - http://hpctoolkit.org/software-instructions.html



Some Hpcviewer Tips



Information for Using Hpcviewer

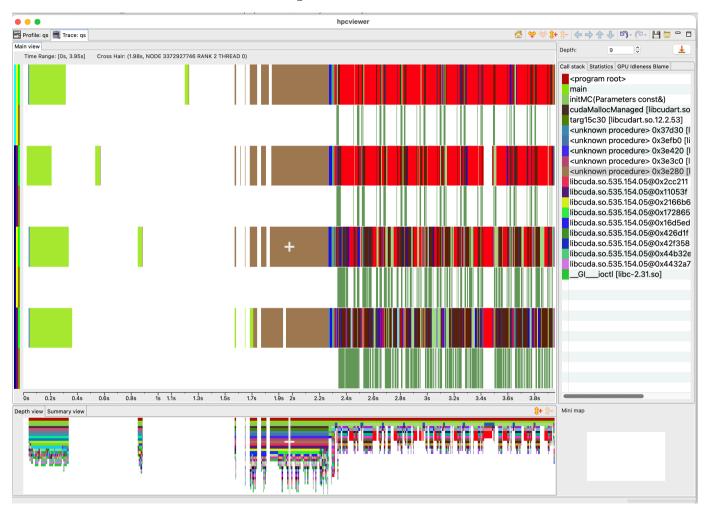
- Filtering GPU traces
 - Can use the filter menu to select what execution traces you want to see
 - · cpu only, gpu, a mix
 - type a string or a regular expression in the chooser select or unselect the new set
 - only traces that exceed a minimum number of samples
- Filtering GPU calling context tree nodes to hide clutter
 - hide individual CCT nodes: e.g. lines that have no source code mapping library@0x0f450
 - hide subtrees: MPI implementation, implementation of CUDA primitives
- When inspecting GPU activity, be aware that hpcviewer has two modes
 - expose GPU traces or not
 - means: when displaying GPU trace lines, don't just show GPU activity if the time in the middle of a pixel is in a GPU operation. instead, show the first (if any) GPU operation between the time in the middle of the pixel and the middle of the next pixel
 - why? GPU activity is so short, it may be hard to find if we don't "expose" where it is
 - · downside: makes the GPU appear more active than it is
 - can correct the statistics by turning the mode off
 - mode can be selected from <File>:<Preferences>:<Traces>



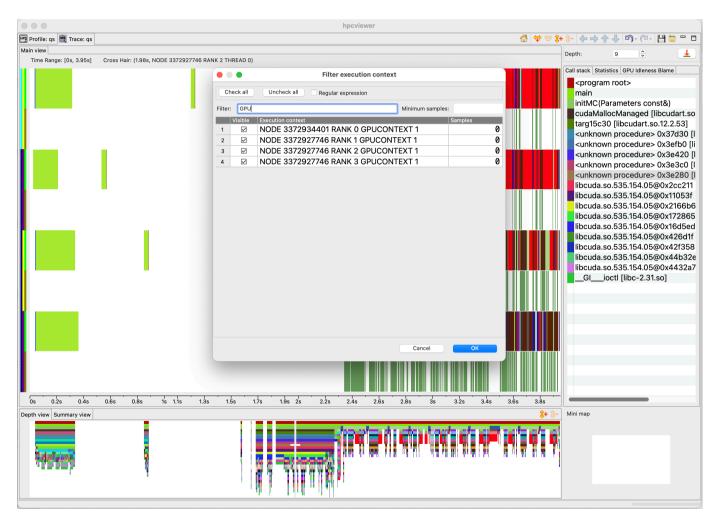
Quicksilver Hand-on Notes



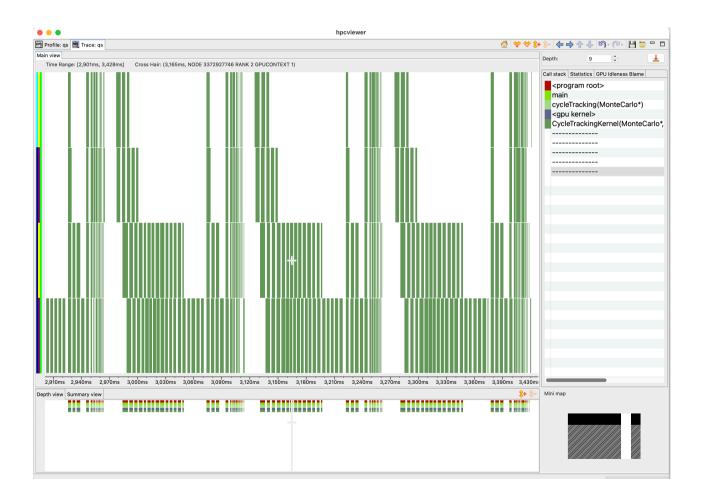
Select the Tab "Trace: qs"



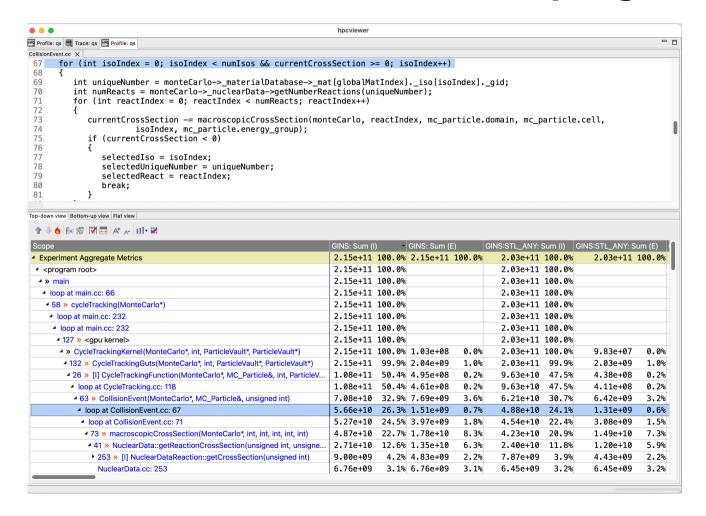
Use the Filter to "Uncheck all" and Check "GPU" streams



See Load Imbalance Across the Four GPUs



The Profile View in the other "PC Sampling" Database



Hands-on Tutorial Examples

```
%git clone https://github.com/hpctoolkit/hpctoolkit-tutorial-
examples
%cd hpctoolkit-tutorial-examples/gpu/nvidia
%ls
% arborx.kokkos lammps.kokkos quicksilver.cuda
```



Analyzing Quicksilver Traces

- Select the Trace tab "Trace: qs"
- Identifying the traces
 - Select a pixel on a trace line
 - Look at legend on the top of the display, which reports the location of the "cross hair"
 - Is this a CPU or GPU trace line?
 - Repeat this a few times to identify what each of the trace lines represents
- Notice that each time you select a colored pixel on a trace line, you will be shown the function call stack in the rightmost pane
- At the top of the pane is a "depth" indicator, that indicates what level in the call stack you are viewing. The selected level will also be highlighted
- You can change the depth of your view by using the depth up/down, typing a depth, or simply selecting a frame in the call stack at the desired depth
- You can select $\stackrel{1}{-}$ above the call stack frame to show the call stacks at the deepest depth
 - If a sample doesn't have an entry at the selected depth, its deepest frame will be shown



Analyzing Quicksilver Traces

- Zoom in on a region in a trace by selecting it in the trace display
- Use the back button to undo a zoom
- - expand or contract the pane
 - move left, right, up, or down
- Keep an eye on the minimap in the lower right corner of the display to know what part of the trace you are viewing
- Use the home button to reset the trace view to show the whole trace



Analyzing Quicksilver Traces

- Select the Trace tab "Trace: qs"
- Configure filtering
 - Use the Filter menu to select Filter Execution Contexts
 - In the filtering menu, select "Uncheck all"
 - Now, in the empty box preceded by "Filter:", type "GPU" and then click "Check all"
 - Select "OK".
 - Now, the Trace View will show only trace lines for the GPUs.
- Inspect the trace data
 - Is the work load balanced across the GPUs? How can you tell?
 - Bring up the filter menu again. Select "Uncheck all". Type in "RANK 3" in the Filter box. Select thread 0 and the GPU context. Select "OK".
 - Move the call stack to depth 2
 - What CPU function is Rank 3 thread 0 executing when the GPU is idle?
 - Does this suggest any optimization opportunities?



Analyzing the Quicksilver Summary Profile

- Select the Profile Tab "Profile: qs"
- Use the column selector to deselect and hide the two REALTIME columns
- Select the GPU OPS column, which represents time spent in all GPU operations
- Select the 6 button to show the "hot path" according to the selected column
 - the hot path of parent will continue into a child as long as the child accounts for 50% or more of the parent's cost
- The hot path will select "CycleTrackingKernel" a GPU kernel that consumes 100% of the GPU cost in this profile
- Use the profiles button to graph "GPU OPS (I)" inclusive GPU operations across the
 - Are the GPU operations balanced or not across the execution contexts (ranks)?



Analyzing the Quicksilver Summary Profile

- You will notice that for quicksilver, HPCToolkit doesn't report any data copies between the host and device
 - The quicksilver code uses "unified memory" so that all of the data movement occurs between CPU and GPU using page faults rather than explicit copies
 - Today's GPU hardware doesn't support attribution of page faults to individual instructions
 - We could profile them, but not attribute them to code



Analyzing Quicksilver PC Samples

Using a measurement database with traces that was collected *with* PC sampling enabled

Using the default top-down view of the profile

- Select the column "GINS (I)" to focus on the measurement of inclusive GPU Instructions
- Select use the flame button to look at where the instructions are executed
- In the call stack revealed, you will <gpu kernel> placeholder that separates CPU activity (above) from GPU kernel
 activity (below)
- Below the <gpu kernel> placeholder you will see the function calls, inlined functions, loops and statements in HPCToolkit's reconstruction of calling contexts within the CycleTrackingKernel
- Using the bottom-up view of the profile
 - Select the bottom-up tab of above the control pane
 - Select the GINS STL_ANY (E) column, which will sort the functions by the exclusive GPU instruction stalls within that function
 - Scroll right to see which of the types of contributing types of stalls accounts for most of the STL ANY amount
 - Select the function that has the most exclusive stalls
 - Select the the hot path to see where this function is called from.
 - Where do the calls to the costly function come from?
 - Does there appear to be an opportunity to reduce the number of calls to this function?



Filtering Tips to Hide Unwanted Implementation Details

- Filter "descendants-only" of CCT nodes with names *MPI* to hide the details of MPI implementation in profiles and traces
- Filter internal details of RAJA and SYCL templates to suppress unwanted detail using a "self-only" filter



Other Databases to Inspect



Other Performance Databases

See /tmp/hpctoolkit-databases for each of the following

- Quicksilver PC Sampling
- QMCPACK CPU prototype: a Quantum Monte Carlo code
 - inspect performance data from an early prototype on 32 MPI ranks x 32 threads per rank
- Deepware-rtm: Reverse time migration using Pytorch
 - trace data shows forward and backward phases utilizing a GPU
 - PC sampling measurements shows GPU stencil computations
- PeleC: a turbulent combustion code being used to simulate jet aircraft engines
 - large data from 2025 NERSC hackathon run on 16 CPUs + 16 GPU

