

Towards Performance Tools for Emerging GPU-Accelerated Exascale Supercomputers



John Mellor-Crummey
Rice University

18th Annual Workshop on Charm++ and Its Applications
October 21, 2020
University of Illinois at Urbana-Champaign (Virtual)



DOE's Forthcoming Heterogeneous Exascale Platforms



- **Aurora compute nodes (ALCF)**
 - 2 Intel Xeon “Sapphire Rapids” processors
 - 6 Intel Xe “Ponte Vecchio” GPUs
 - 8 Slingshot endpoints
 - Unified memory architecture



- **Frontier compute nodes (OLCF)**
 - 1 AMD EPYC CPU
 - 4 purpose-built AMD Radeon Instinct GPUs
 - multiple Slingshot endpoints
 - Unified memory architecture



- **El Capitan compute nodes (LLNL)**
 - Next-generation AMD EPYC “Genoa” CPU (5nm)
 - Next-generation AMD Radeon Instinct GPUs

Node-level Programming Models for Forthcoming Exascale Systems

- **Native programming models from platform vendors**

- Intel DPC++

C++ + SYCL specification 1.2.1 + extensions: device queues, buffers, accessors, parallel_for, single_task, parallel_for_work_group, events

- HIP: AMD's CUDA-like model

- **Directive-based models**

- OpenACC offload structured blocks and device functions, work sharing loops, data environment, data mappings
- OpenMP

- **C++ template-based models**

- RAJA parallelism loops, iteration spaces, execution policies, traversal templates, lambda functions, n-dimensional array abstractions, and lambda functions
- Kokkos

A Diverse Set of Global Programming Models

- **Message passing**
 - MPI
- **Partitioned global address space programming models**
 - languages
 - Coarray Fortran, Coarray C++, Chapel, UPC
 - librarys
 - UPC++, GASNet, OpenSHMEM , Global Arrays
- **Object-based**
 - Charm++

Source: Frontier Spec Sheet

https://www.olcf.ornl.gov/wp-content/uploads/2019/05/frontier_specsheet.pdf

Performance Analysis Challenges for GPU-accelerated Supercomputers

- **Myriad performance concerns**

- Computation performance
 - Principal concern: keep GPUs busy and computing productively
 - need extreme-scale data parallelism!
- Data movement costs within and between memory spaces
- Internode communication
- I/O

- **Many ways to hurt performance**

- insufficient parallelism, load imbalance, serialization, replicated work, parallel overheads ...

- **Hardware and execution model complexity**

- Multiple compute engines with vastly different characteristics, capabilities, and concerns
- Multiple memory spaces with different performance characteristics
 - CPU and GPU have different complex memory hierarchies
- Asynchronous execution

Measurement Challenges for GPU-accelerated Supercomputers

- **Extreme-scale parallelism**
 - Serialization within tools will disrupt parallel performance
- **Multiple measurement modalities and interfaces**
 - Sampling on the CPU
 - Callbacks when GPU operations are launched and (sometimes) completed
 - GPU event stream
- **Frequent GPU kernel launches require a low-overhead measurement substrate**
- **Importance of third-party measurement interfaces**
 - Tools can only measure what GPU hardware can monitor
 - support for fine-grain measurement will be essential to diagnose GPU inefficiencies
 - Linux perf_events for kernel measurement
 - GPU monitoring libraries from vendors

Engineering Challenges for Performance Tools

- **Complex applications**

- Compositions of programming models
- > 100 dynamic libraries
- Application binaries exceeding 5GB
- HPC libraries that intercept system calls (mmap, munmap, open, close)
- Quirky application characteristics
 - NAMD: exit initiated by a non-initial thread
 - Kull: forking non-readable helper application

- **Dynamic library loading**

- Implicit system locks on dynamic library state
- RUNPATH: library-specific library load path
- Early threads in library init constructors
- Nested dynamic library loading

- **Provisioning thread local state**

- Implicit lock when creating or destroying thread local storage

- **Process fork**

- atfork handlers trigger thread destructors

- **Interactions with vendor tool substrates**

- Libraries lack documentation of their actions, e.g. creating threads
- Callbacks for submission and completion on unspecified (and sometimes different) threads

- **Interaction between tools and software stack**

- Interaction of signals with everything
- Managing monitoring when forking

- **Lack of vendor tooling and documentation**

- Non-standard GPU binary formats that lack public documentation

GPU Performance Tools

- **Features**

- Trace view
 - A series of events that happen over time on each process, thread, and GPU stream
- Profile view
 - A correlation of performance metrics with program contexts

- **Existing GPU performance tools**

- GPU vendors
 - Nsight Systems, Nsight Compute, nvprof, ROCProfiler, Intel VTune
- Third parties
 - TAU, VampirTrace, ARM Map

Tool Shortcomings for Analyzing Complex GPU-accelerated Programs

- They lack a comprehensive profile view to analyze
 - Sophisticated CPU calling contexts
 - A GPU API (e.g., cudaMemcpy) invoked in different places
 - Sophisticated GPU calling contexts
 - OpenMP Target, Kokkos, and RAJA generate code with many small procedures
- At best, existing tools only attribute runtime cost to a flat profile view of functions executed on GPUs

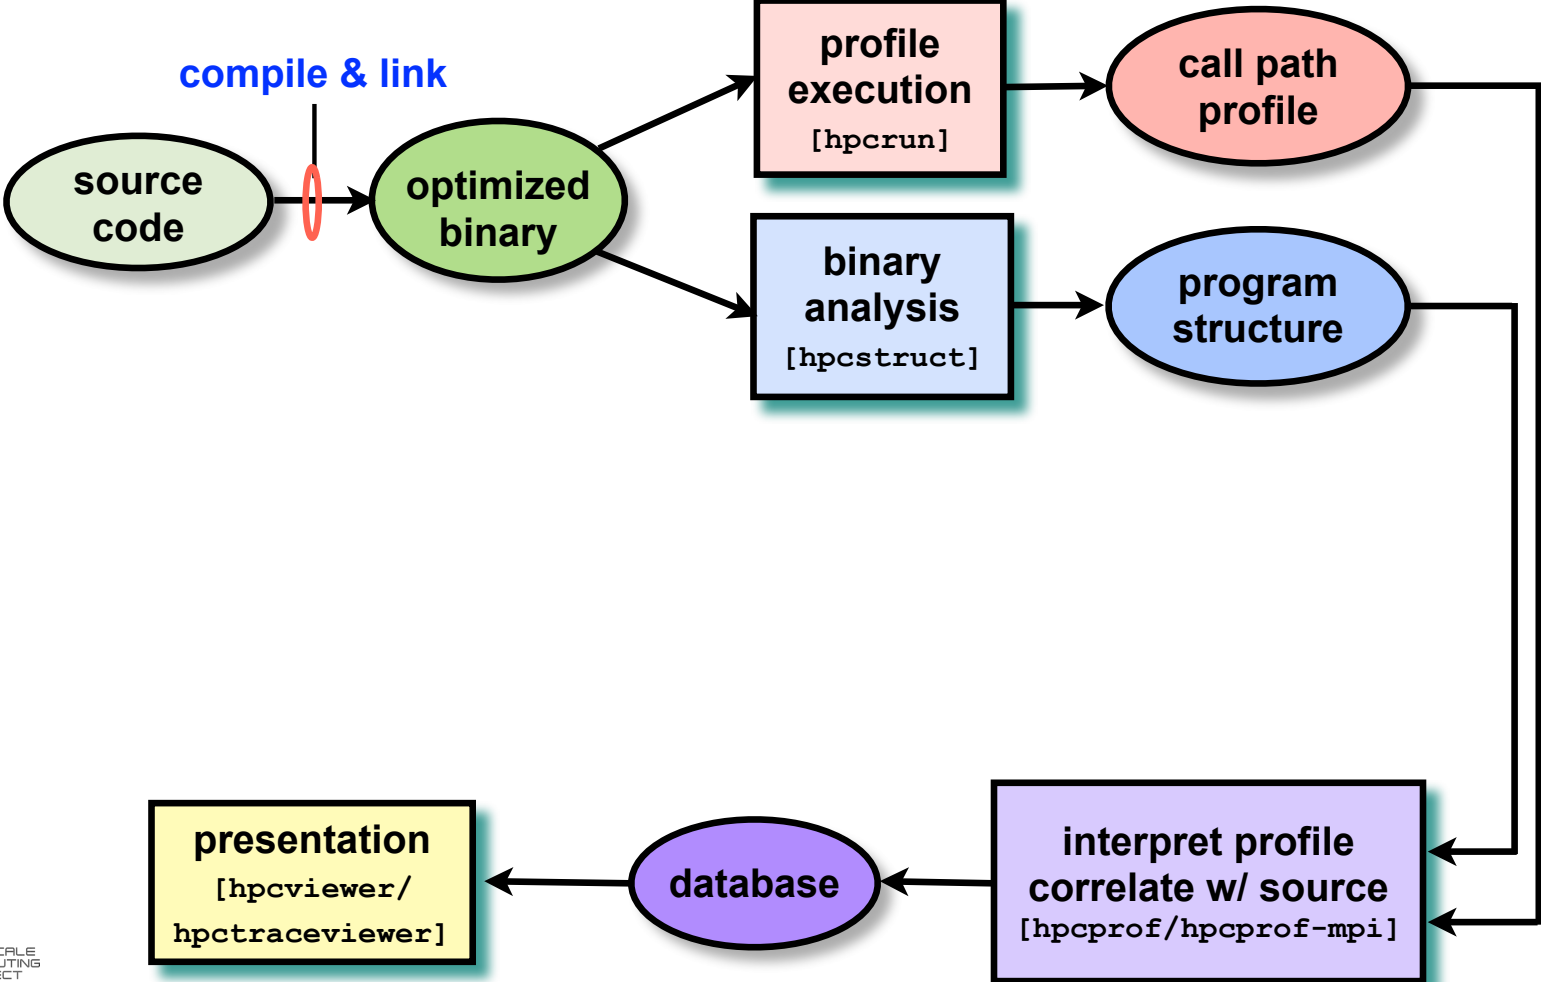
Outline

- **Performance measurement and analysis challenges for GPU-accelerated supercomputers**
- **Introduction to HPCToolkit performance tools**
 - Overview of HPCToolkit components and their workflow
 - HPCToolkit's graphical user interfaces
- **Analyzing the performance of GPU-accelerated supercomputers with HPCToolkit**
 - Overview of HPCToolkit's GPU performance measurement capabilities
 - Collecting measurements
 - Analysis and attribution
 - Scalable analysis of performance data
- **Status, ongoing work, final remarks**

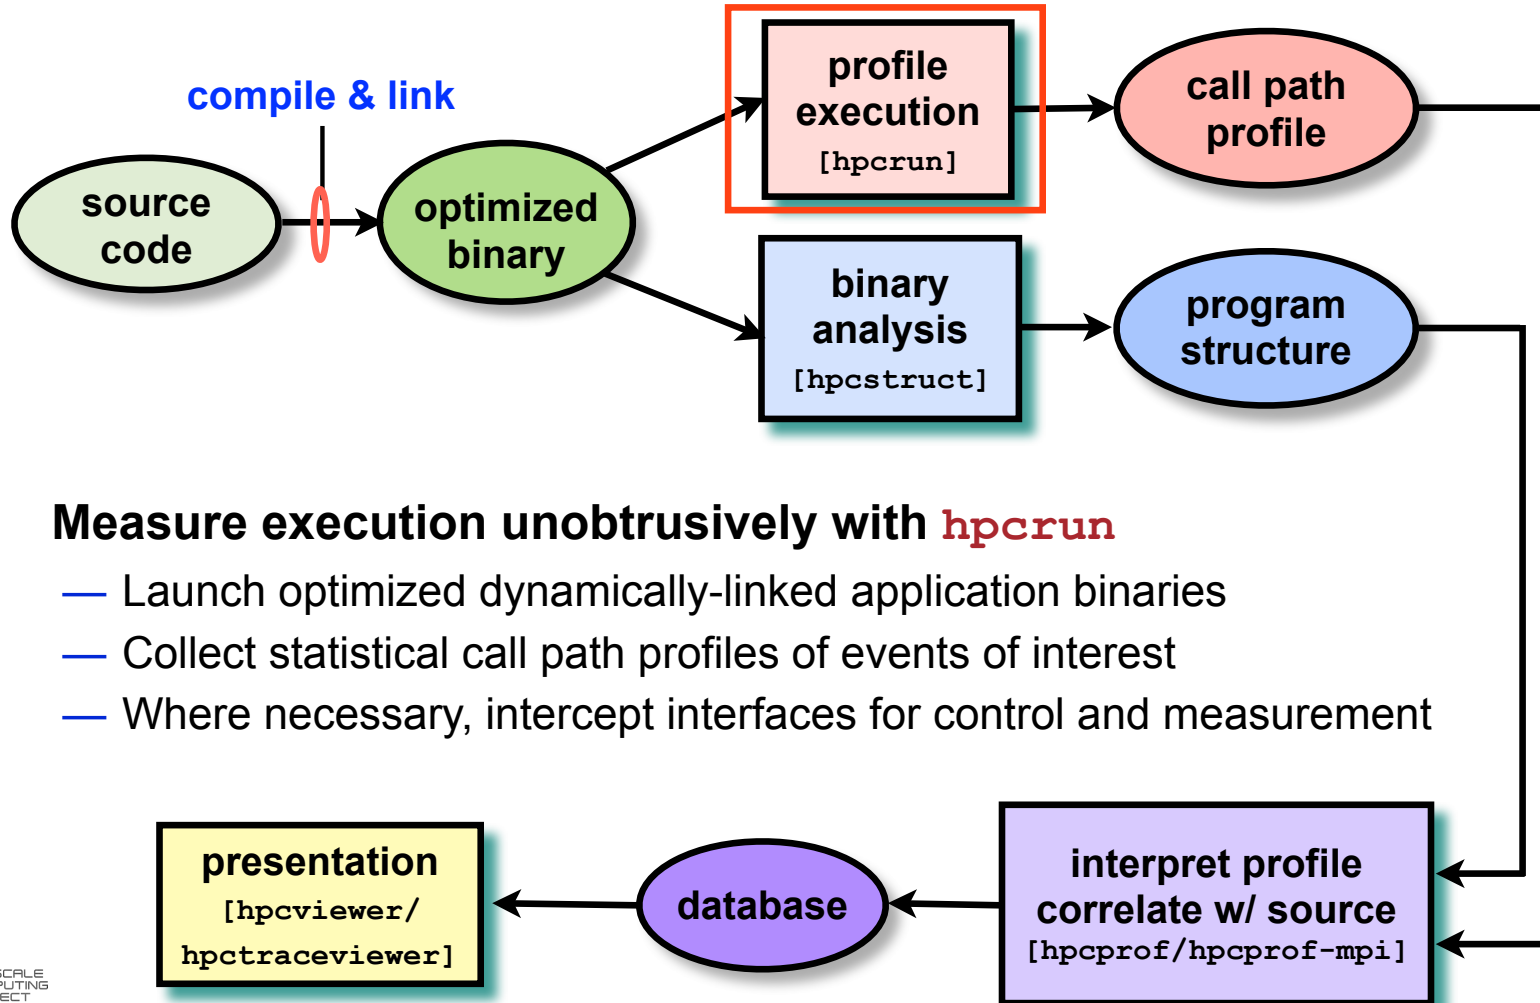
Rice University's HPCToolkit Performance Tools

- **Employs binary-level measurement and analysis**
 - Observes executions of fully optimized, dynamically-linked parallel applications
 - Supports multi-lingual codes with external binary-only libraries
- **Collects sampling-based measurements of CPU**
 - Controllable overhead
 - Minimize systematic error and avoid blind spots
 - Enable data collection for large-scale parallelism
- **GPU performance using measurement APIs provided by vendors**
 - Callbacks to monitor launch/completion of GPU operations
 - NVIDIA and AMD: activity API provides information about asynchronous operations on GPU devices
 - NVIDIA PC sampling and Intel GTPin instrumentation for fine-grain measurement
- **Associates metrics with both static and dynamic context**
 - Loop nests, procedures, inlined code, calling context on both CPU and GPU
- **Enables one to specify and compute derived CPU and GPU performance metrics of your choosing**
 - Diagnosis often requires more than one species of metric
- **Supports top-down performance analysis**
 - Identify costs of interest and drill down to causes: up and down call chains, over time

HPCToolkit Workflow



HPCToolkit Workflow




Measure execution unobtrusively with **hpcrun**

- Launch optimized dynamically-linked application binaries
- Collect statistical call path profiles of events of interest
- Where necessary, intercept interfaces for control and measurement


Call Path Profiling

- **Measure and attribute costs in context**
 - Sample timer or hardware counter overflows
 - Gather CPU calling context using stack unwinding

Call path sample

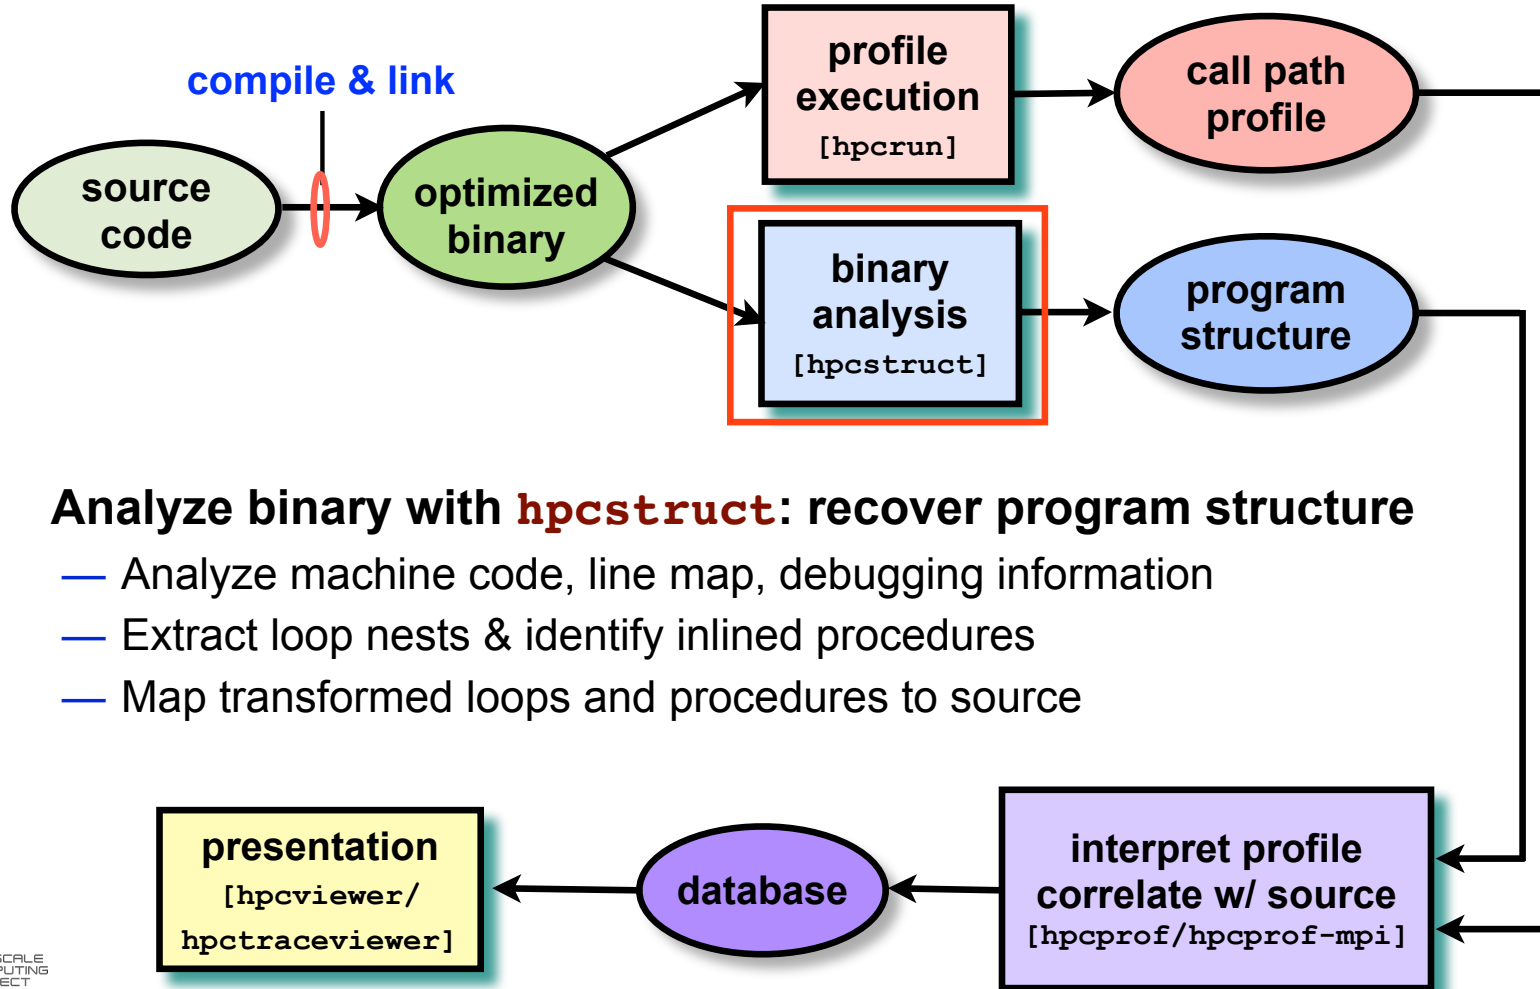


Calling context tree



Overhead proportional to sampling frequency, not call frequency

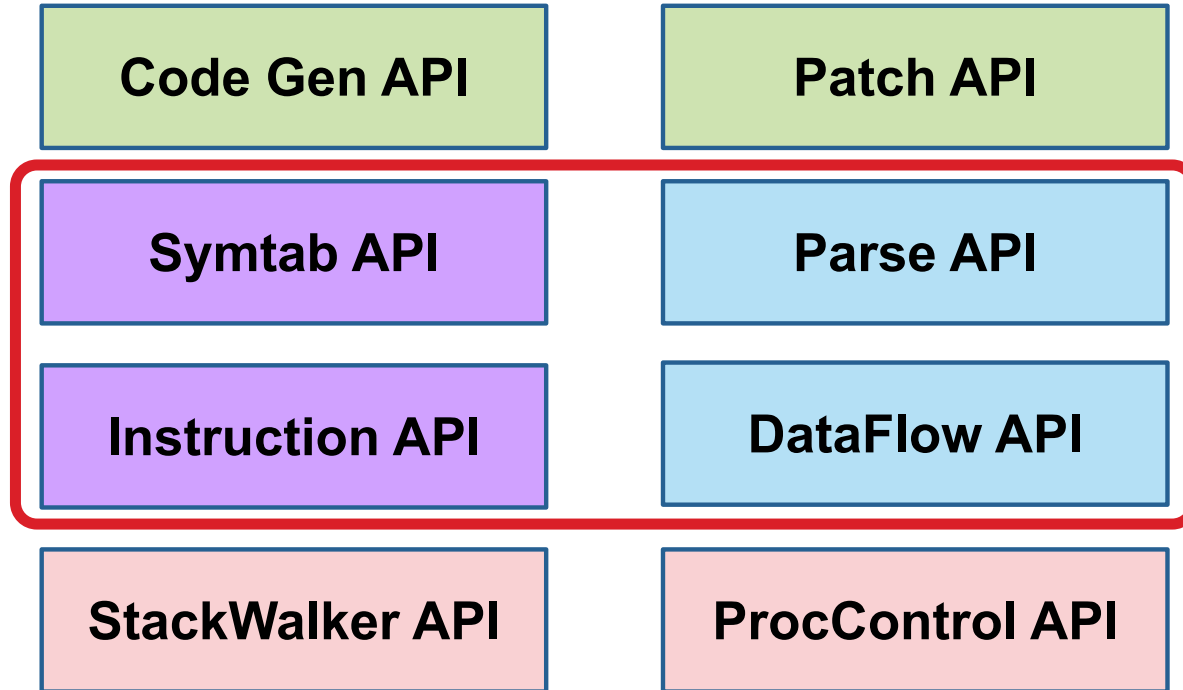
HPCToolkit Workflow



Analyze binary with **hpcstruct**: recover program structure

- Analyze machine code, line map, debugging information
- Extract loop nests & identify inlined procedures
- Map transformed loops and procedures to source

Dyninst: A Toolkit for Binary Analysis and Instrumentation



Architectures

X86_64

Power/BE

Power/LE

ARM

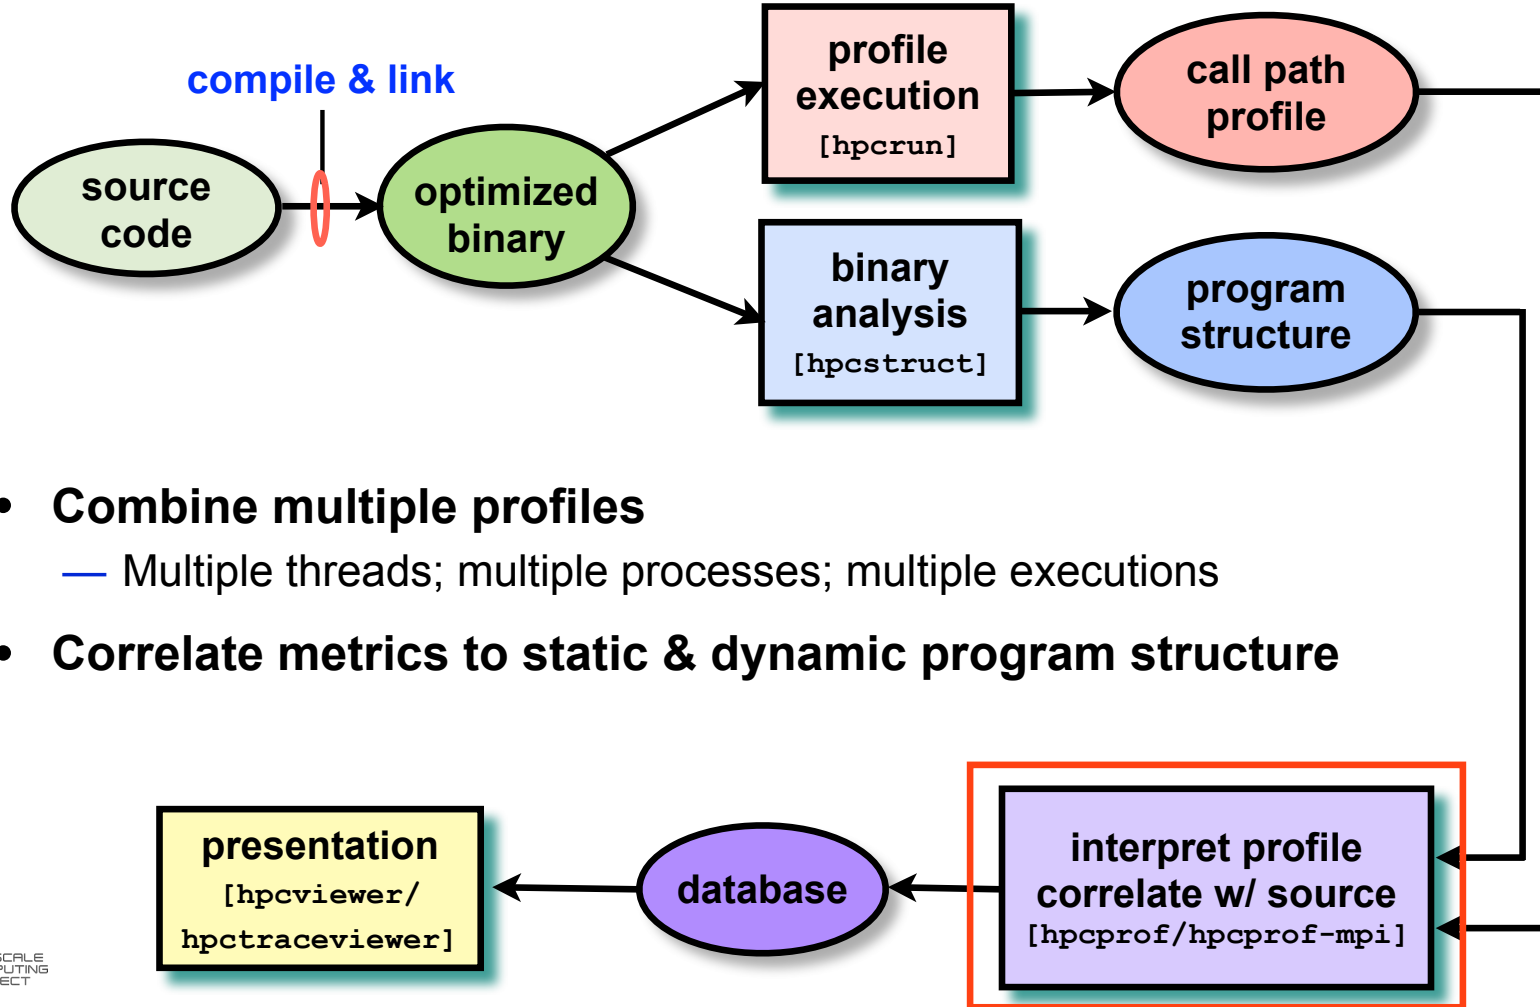
AMD Vega

CUDA

Intel GPU

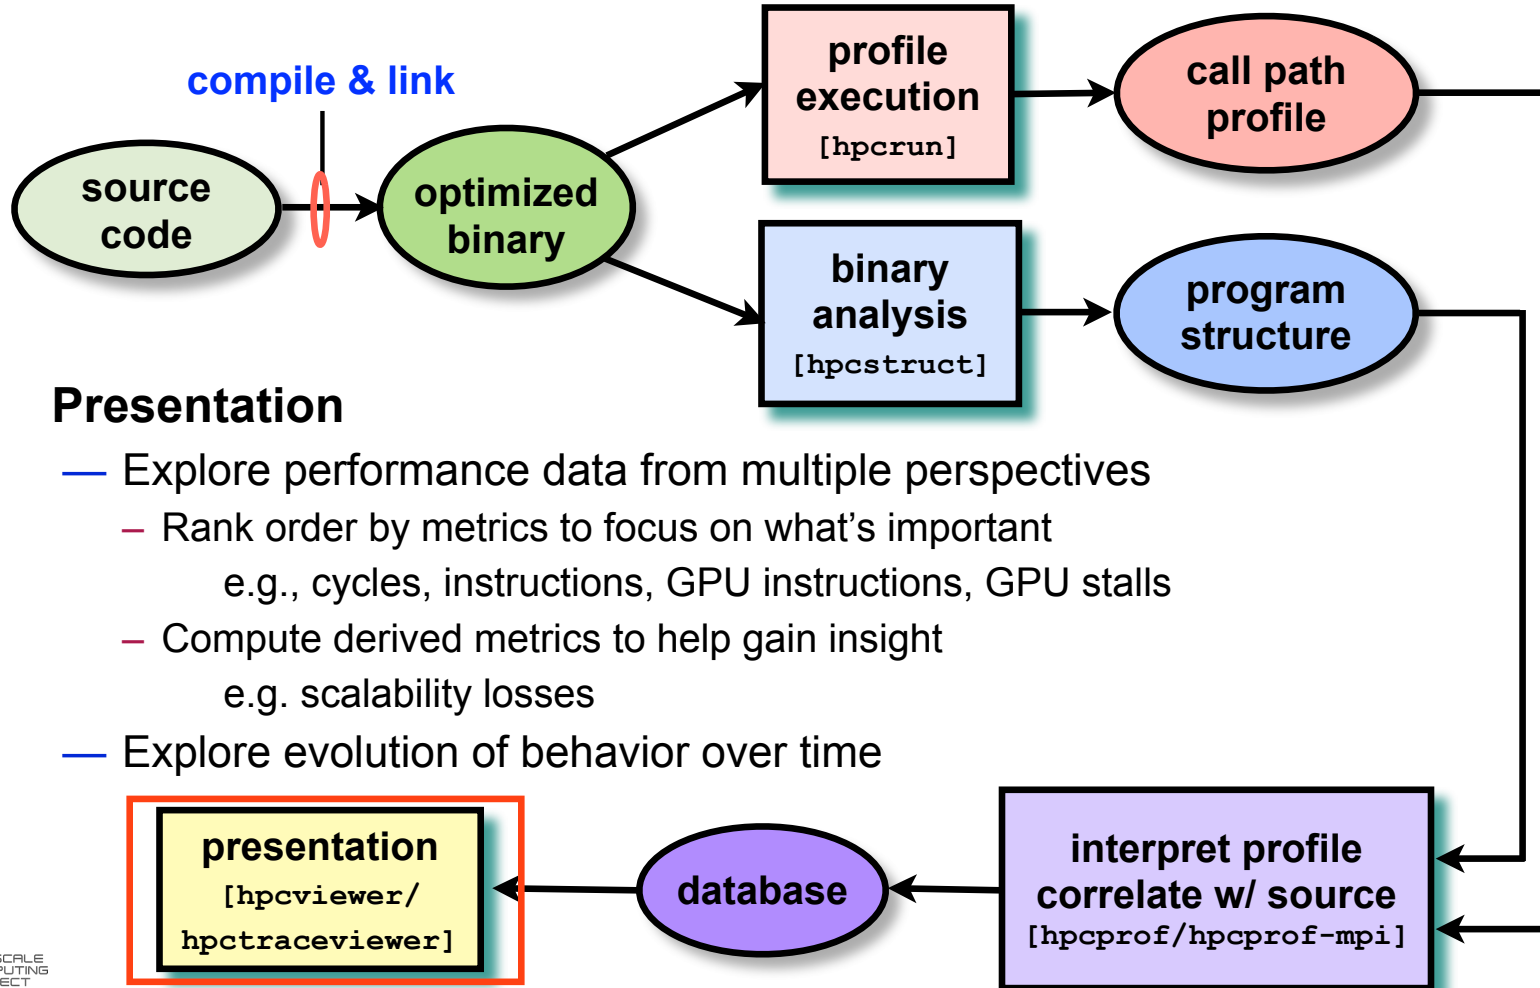
Lead Institution: University of Wisconsin – Madison

HPCToolkit Workflow



- **Combine multiple profiles**
 - Multiple threads; multiple processes; multiple executions
- **Correlate metrics to static & dynamic program structure**

HPCToolkit Workflow



Presentation

- Explore performance data from multiple perspectives
 - Rank order by metrics to focus on what's important
e.g., cycles, instructions, GPU instructions, GPU stalls
 - Compute derived metrics to help gain insight
e.g. scalability losses
- Explore evolution of behavior over time

Code-centric Analysis with hpcviewer

- function calls in full context
- inlined procedures
- inlined templates
- outlined OpenMP loops
- sequential loops


The screenshot shows the hpcviewer interface with several components highlighted by red boxes and arrows:

- source pane:** Shows the source code for `forall_generic.hxx` with lines 402-407. The code includes a `forall` function call with `EXEC_POLICY_T()`.
- view control:** A toolbar with buttons for `Calling Context View`, `Callers View`, and `Flat View`.
- metric display:** A toolbar with icons for search, zoom, and other navigation functions.
- navigation pane:** A tree view showing the execution scope, including `main`, `loop at luleshRAJA-parallel.cxx: 3526`, and various function calls like `LagrangeLeapFrog`, `CalcForceForNodes`, and `CalcVolumeForceForElems`.
- metric pane:** A table showing performance metrics for each scope.

Scope	REALTIME (usec):Sum (I)	REALTIME (usec):Sum (E)
Experiment Aggregate Metrics	2.26e+08 100 %	2.26e+08 100 %
<program root>	1.45e+08 63.9%	
497: main	1.45e+08 63.9%	6.01e+03 0.0%
loop at luleshRAJA-parallel.cxx: 3526	1.44e+08 63.8%	
3528: [] LagrangeLeapFrog(Domain*)	1.44e+08 63.8%	
2715: [] LagrangeNodal(Domain*)	8.70e+07 38.5%	
1554: [] CalcForceForNodes(Domain*)	8.30e+07 36.7%	
1469: CalcVolumeForceForElems(Domain*)	8.25e+07 36.5%	
1454: [] CalcHourglassControlForElems(Domain*, double*, double)	5.15e+07 22.8%	
1399: [] CalcFBHourglassForceForElems(int*, double*, double*, double*, double*, double)	3.10e+07 13.7%	
1187: [] void RAJA::forall<RAJA::IndexSet::ExecPolicy<RAJA::seq_segit, RAJA::omp_parallel>>(const INDEXSET_T&, LOOP_BODY loop_body)	2.43e+07 10.8%	
405: [] void RAJA::forall<RAJA::omp_parallel_for_exec, CalcFBHourglassForceForElems(int*, double*, double*, double*, double*, double)>(const INDEXSET_T&, LOOP_BODY loop_body)	2.43e+07 10.8%	
loop at forall_seq_any.hxx: 498	2.43e+07 10.8%	
505: [] void RAJA::forall<CalcFBHourglassForceForElems(int*, double*, double*, double*, double*, double)>(const INDEXSET_T&, LOOP_BODY loop_body)	2.43e+07 10.8%	1.00e+03 0.0%
89: outline forall_omp_any.hxx:89 (0x423620)	2.42e+07 10.7%	3.91e+04 0.0%
loop at forall_omp_any.hxx: 90	2.42e+07 10.7%	3.41e+04 0.0%
91: [] CalcFBHourglassForceForElems(int*, double*, double*, double*, double*, double)	2.42e+07 10.7%	9.84e+06 4.3%
1300: [] CalcElemFBHourglassForce(double*, double*, double*, double*, double*, double)	1.11e+07 4.9%	1.11e+07 4.9%
1260: [] CBRT(double)	3.27e+06 1.4%	2.00e+05 0.1%

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



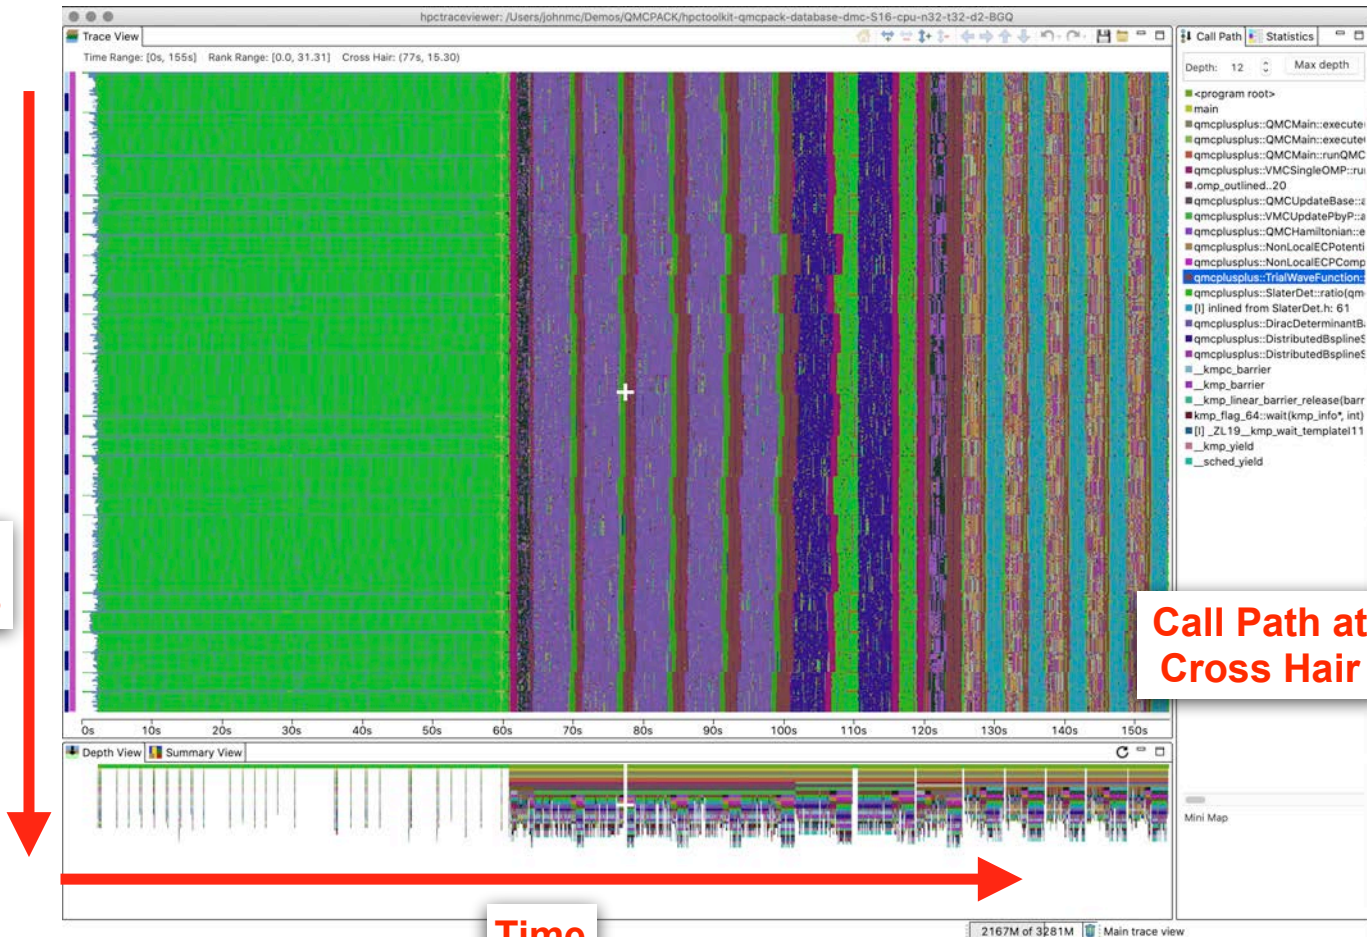
Time-centric Analysis with hpctraceviewer

Experimental version of QMCPack on Blue Gene Q

- 32 ranks
- 32 threads each

Call stack view

Ranks/
Threads



Call Path at
Cross Hair

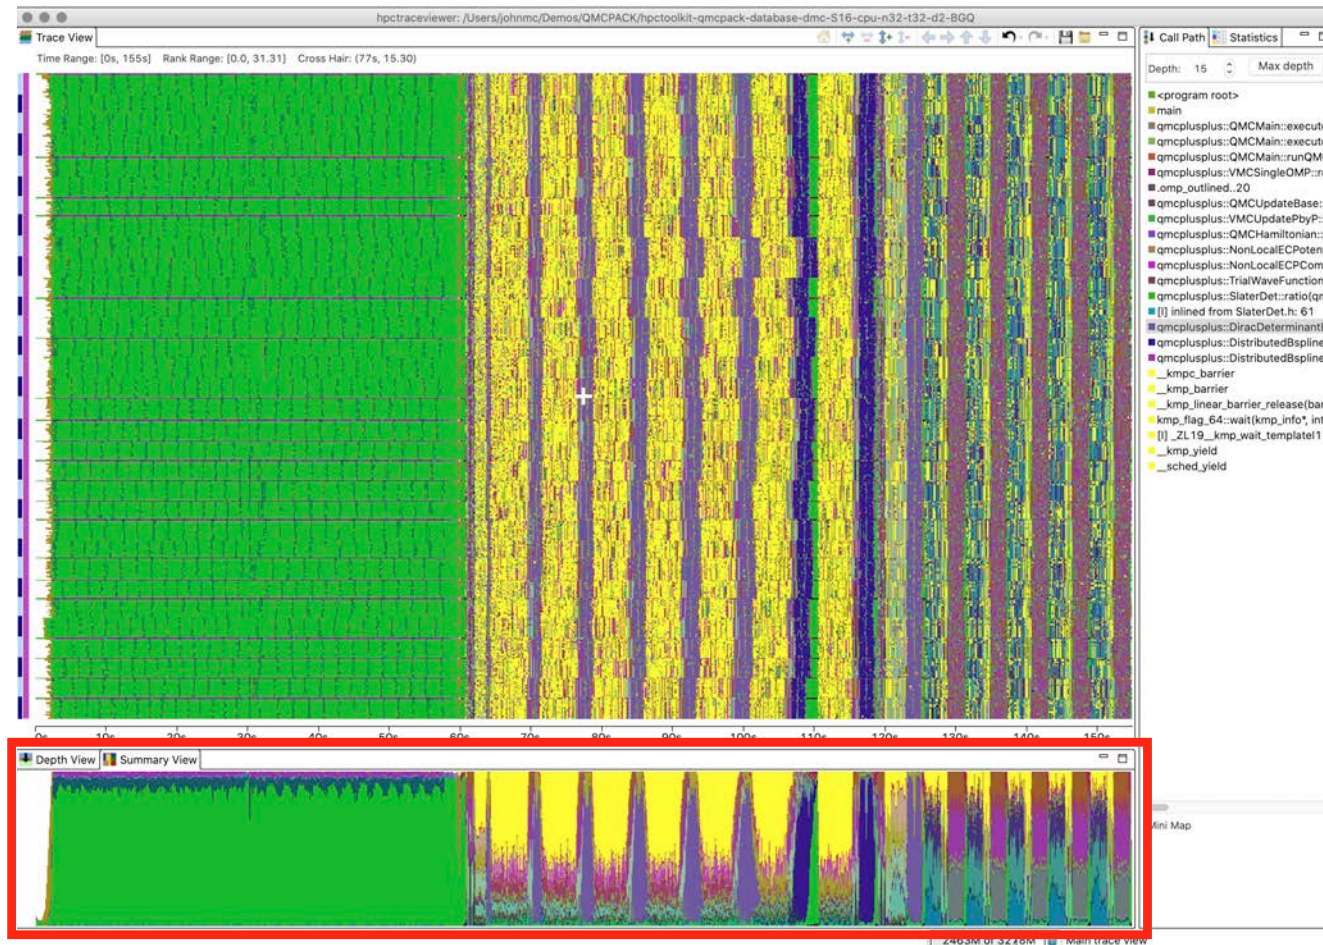
Time

Time-centric Analysis with hpctraceviewer

Experimental version of QMCPack on Blue Gene Q

- 32 ranks
- 32 threads each

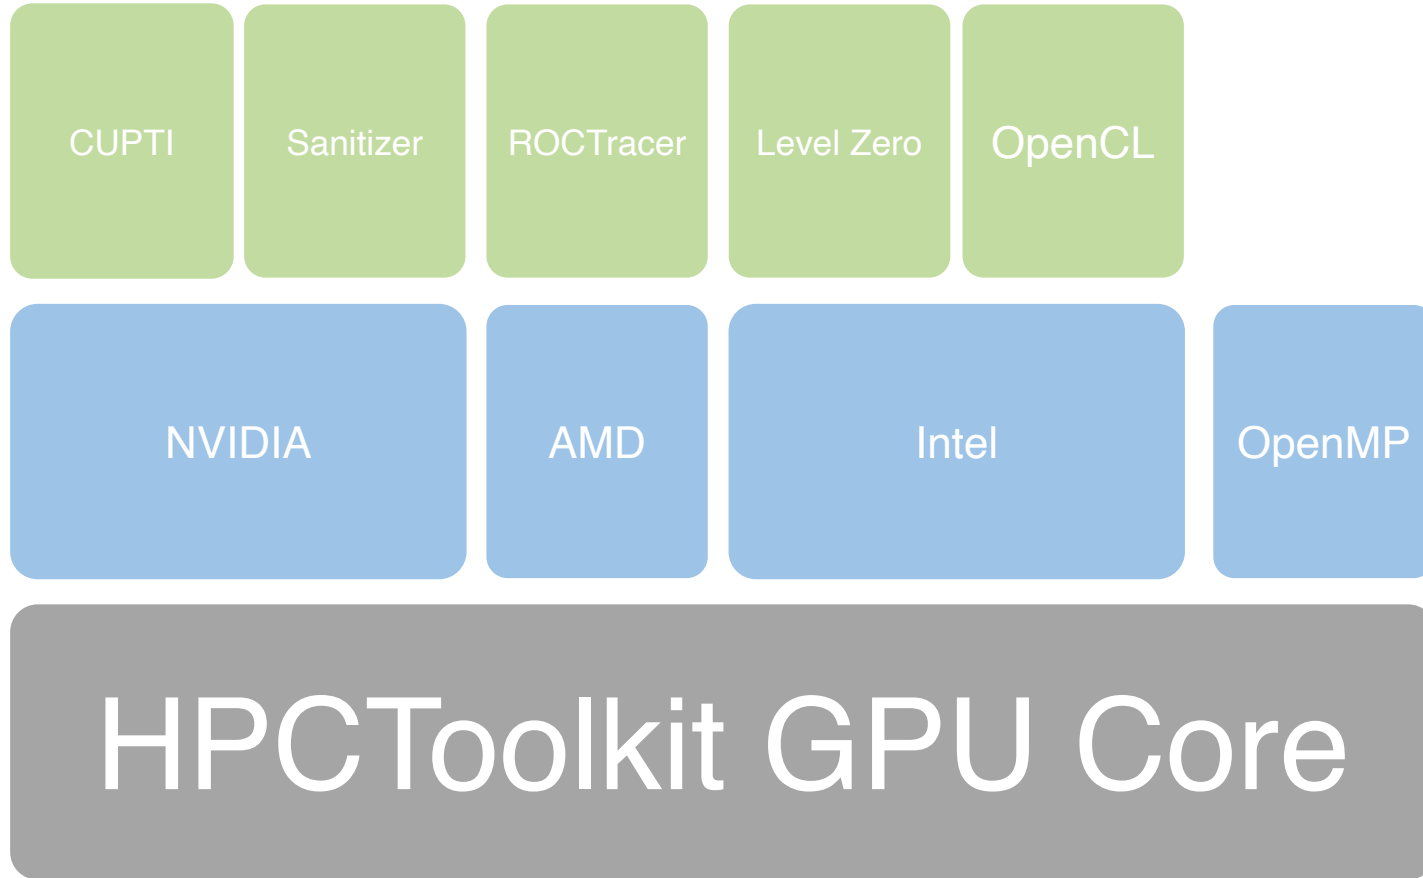
Summary view modeled after Projections



Outline

- **Performance measurement and analysis challenges for GPU-accelerated supercomputers**
- **Introduction to HPCToolkit performance tools**
 - Overview of HPCToolkit components and their workflow
 - HPCToolkit's graphical user interfaces
- **Analyzing the performance of GPU-accelerated supercomputers with HPCToolkit**
 - Overview of HPCToolkit's GPU performance measurement capabilities
 - Collecting measurements
 - Analysis and attribution
 - Scalable analysis of performance data
- **Status, ongoing work, final remarks**

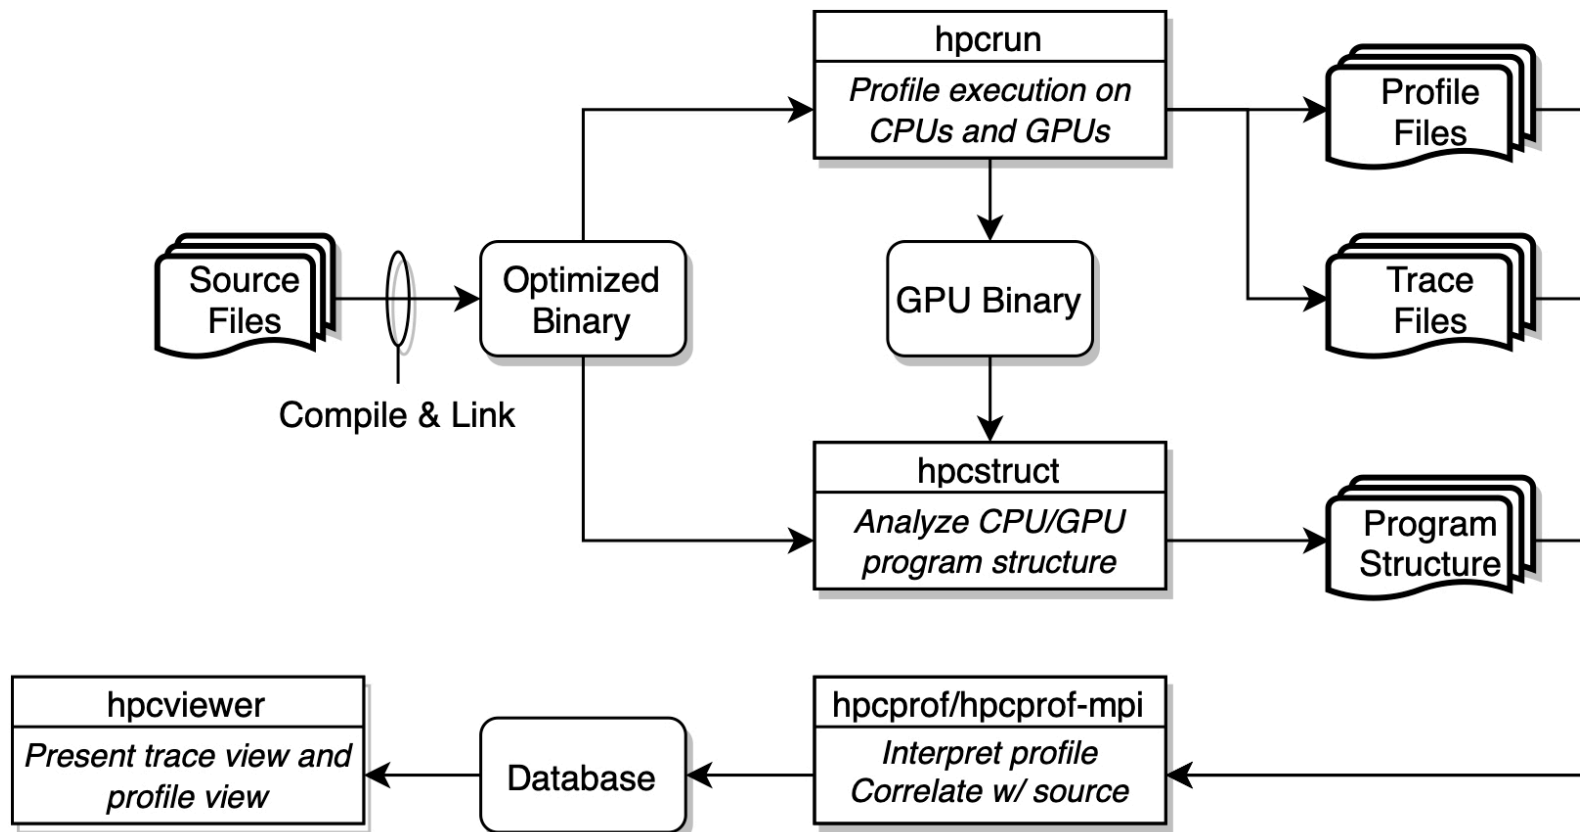
HPCToolkit for GPU-accelerated Computations



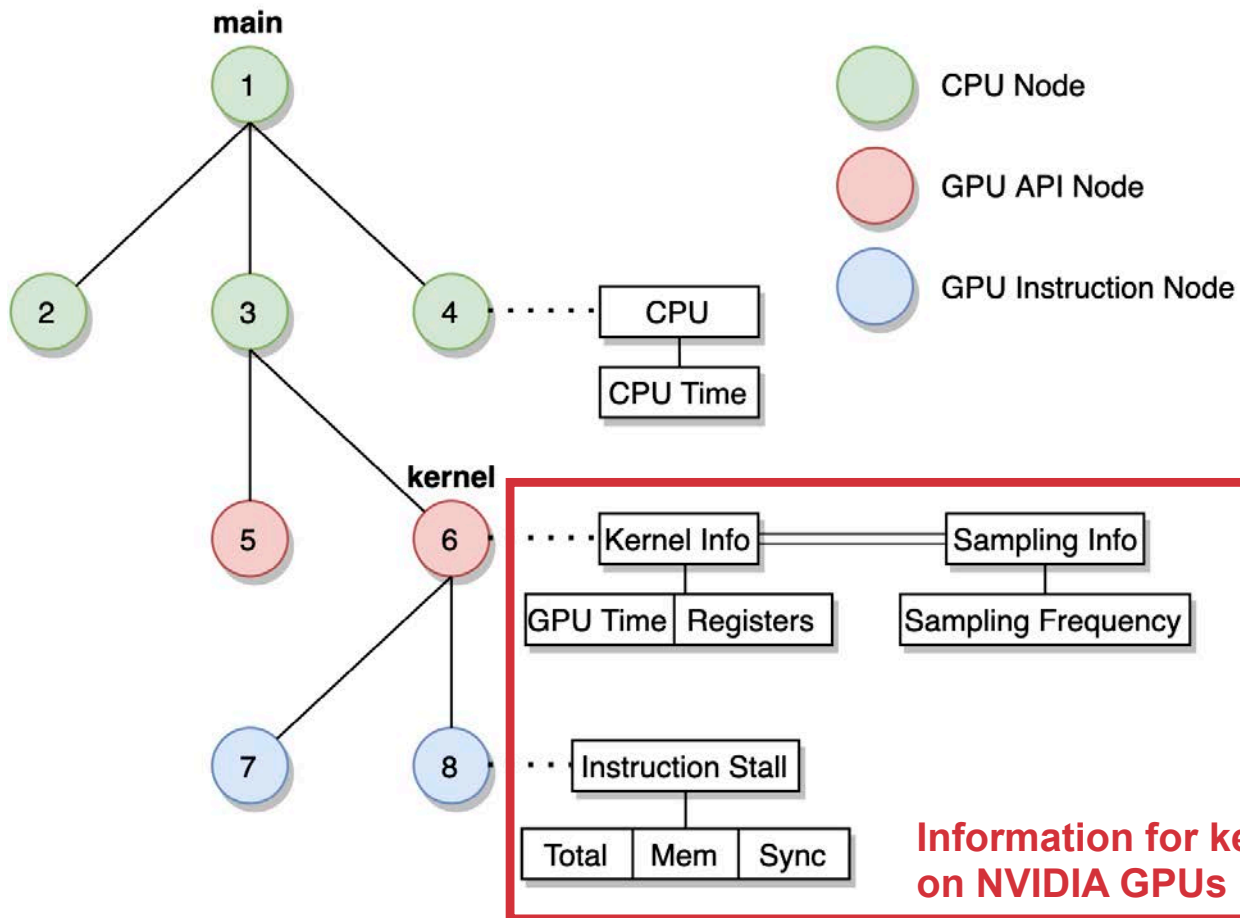
Highlights of HPCToolkit's Support for GPU-accelerated Codes

- **It unwinds the CPU call stack to identify the CPU calling context for each GPU API invocation**
- **It employs a novel and fast wait-free data structure for inter-thread communication**
- **It employs binary analysis of GPU code to attribute fine-grain performance measurements to functions, inlined functions and templates, loops, and statements**
 - NVIDIA, Intel, and AMD GPU binaries
- **It uses a novel technique to reconstruct an approximate GPU calling context tree for computations from instruction-level measurements**
- **On NVIDIA GPUs: derive a rich set of metrics from PC samples from a single execution**
- **It performs scalable analysis of sparse representations of performance measurements and produces sparse representations tailored for graphical user interfaces**

HPCToolkit's Workflow for GPU-accelerated Applications



HPCToolkit's Sparse Representation of Measurements at Run-time



HPCToolkit's Code-Centric Profiles of GPU-accelerated Code

Profile: PeleC3d.gnu.CUDA.ex

main.cpp | mechanism.cpp | **cuda_runtime.h**

```

1911  /*convert to chemkin units */
1912  for (id = 0; id < 9; ++id) {
1913      C[id] *= 1.0e-6;
1914      wdot[id] *= 1.0e-6;
1915  }
1916 }
    
```

Top-down view | Bottom-up view | Flat view

	Fine-grained Metrics				Coarse-grained Metrics				Derived Metrics
	GINS:Sum (I)		GINS:Sum (E)		GKER (sec):Sum (I)		GKER (sec):Sum (E)		GPU UTIL
↳ 1325: ParallelFor<_nv_dl_wrapper_t<_nv_dl_tag<void (PeleC::*)(double, c	3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
↳ 797: [I] amrex::launch_global<_nv_dl_wrapper_t<_nv_dl_tag<void (*)(3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
↳ 11: [I] _wrapper_device_stub_launch_global<_nv_dl_wrapper_t<_	3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
↳ 114: [I] _device_stub_ZN5amrex13launch_globalIZNS_11Paralle	3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
↳ 109: [I] cudaLaunchKernel<char>	3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
CPU Calling Context ↳ 209: cudaLaunchKernel [PeleC3d.gnu.CUDA.ex]	3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
GPU API Node ↳ <gpu kernel>	3.30e+09	1.0%			2.99e-01	16.5%			2.49 %
↳ _ZN5amrex13launch_globalIZNS_11ParallelForIZNSP	3.08e+09	1.0%	2.57e+08	0.1%	2.99e-01	16.5%	2.99e-01	16.5%	
↳ loop at AMReX_GpuLaunchFuncsG.H: 797	3.08e+09	1.0%	2.56e+08	0.1%					
↳ [I] inlined from React.H: 163	3.08e+09	1.0%	2.56e+08	0.1%					
↳ loop at React.H: 163	3.08e+09	1.0%	2.53e+08	0.1%					
↳ [I] inlined from EOS.H: 59	2.88e+09	0.9%	5.28e+07	0.0%					
↳ 168: CKWC	4.57e+08	0.1%	4.56e+08	0.1%					
GPU Calling Context mechanism.cpp: 1914	2.13e+07	0.0%	2.13e+07	0.0%					

GPU Performance Measurement

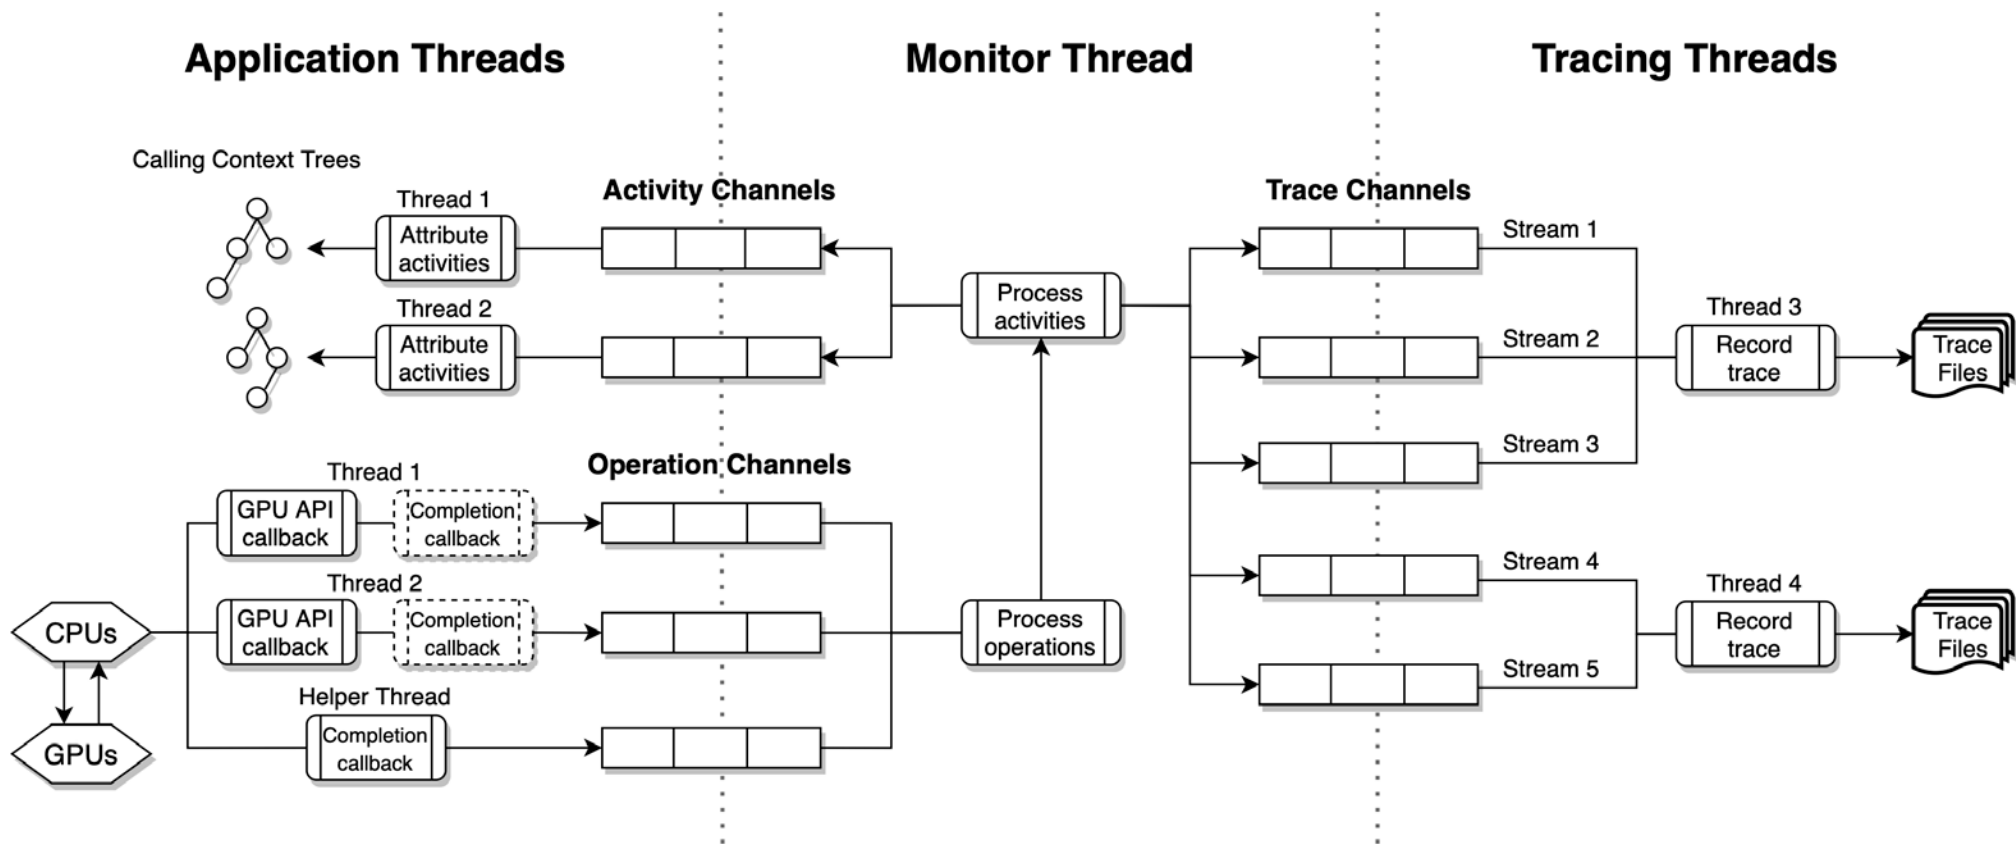
- Three categories of threads

- Application Threads (N per process)
 - Launch kernels, move data, and synchronize GPU calls
- Monitor Thread (1 per process)
 - Monitor GPU events and collect GPU measurements
- Tracing Threads (1 for every K GPU streams)

- Interactions

- **Create correlation:** An application thread T creates a correlation record when it launches a kernel and tags the kernel with a correlation ID C , notifying the monitor thread that C belongs to T
- **Attribute measurements:** The monitor thread collects measurements associated with C and communicates measurement records back to thread T
- **Record traces:** The monitor thread sends activity traces to tracing threads to record in a separate trace file per GPU stream (NVIDIA, AMD) or device queue (Intel, AMD)

HPCToolkit's Runtime Monitoring Infrastructure for OpenCL



Wait-free Channels for Inter-thread Transfer of Measurement Data

- Unidirectional wait-free channel between a pair of threads
 - implemented by a pair of stacks: one private, one shared
- Operations
 - **Producer PUSH(CAS)**: push an item on a shared stack
 - **Consumer STEAL(XCHG)**: steal the contents of the shared stack, push the contents onto a private stack
 - **Consumer POP**: pop an item from the private stack
- Wait-free because **PUSH** fails at most once when a concurrent thread **STEALS** contents of the shared stack
- Bi-directional channel: pair of wait-free unidirectional channels

Correlating GPU API Invocations with CPU Calling Context

- Unwind a call stack from each API invocation
 - Kernel launch, memory copy, and synchronization
- Initial approach
 - Identify the function enclosing each call site in the call stack using a global shared map
- Problem
 - Applications have deep call stacks and large codebase
 - Nyx: up to 60 layers and 400k calls
 - Laghos: up to 40 layers and 100k calls

Fast Unwinding

- Memoize common call path prefixes
 - Temporally-adjacent samples in often share common call path prefixes
 - Employ eager (mark bits) or lazy (trampoline) marking to identify LCA of stack unwinds
- Avoid costly access to mutable concurrent data
 - Cache unwinding recipes in a per thread hash table
- Avoid duplicate unwinds
 - Filter CUDA Driver APIs within CUDA Runtime APIs

Approximation of GPU Calling Contexts to Understand Performance

- GPU code from C++ template-based programming models is complex
- NVIDIA GPUs collect flat PC samples
- Flat profiles for instantiations of complex C++ templates are inscrutable
- HPCToolkit reconstructs approximate GPU calling contexts
 - Reconstruct call graph from machine code
 - Infer calls at call sites
 - PC samples of call instructions indicate calls
 - Use call counts to apportion costs to call sites
 - PC samples in a routine

Scope	GPU INST:Sum (l)	GPU STALL:Sum (l)	MIXINTEGRALADD3:Sum (l)
143: [l] void RAJA::forall<RAJA::policy::cuda::cuda_exec<256ul, true>, RAJA::TypedRangeSegment<long, long>, __nv_	7.28e+11	88.5%	6.46e+11 93.1%
723: [l] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<RAJA::type_traits::is_indexset_polic	7.28e+11	88.5%	6.46e+11 93.1%
370: [l] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<RAJA::type_traits::is_indexset_pc	7.28e+11	88.5%	6.46e+11 93.1%
183: [l] void RAJA::policy::cuda::forall_impl<RAJA::TypedRangeSegment<long, long>, __nv_d1_wrapper_t<__nv_	7.28e+11	88.5%	6.46e+11 93.1%
190: void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterators::numeric_iterator<long, long, lo	7.28e+11	88.5%	6.46e+11 93.1%
145: [l] __wrapper__device_stub_forall_cuda_kernel<256ul, RAJA::Iterators::numeric_iterator<long int>, __m	7.28e+11	88.5%	6.46e+11 93.1%
37: [l] __device_stub_ZNRAJA6policy4cuda4impl18forall_cuda_kernellm256ENS_9Iterators16numeric	7.28e+11	88.5%	6.46e+11 93.1%
26: [l] cudaLaunchKernel<char>	7.28e+11	88.5%	6.46e+11 93.1%
209: cudaLaunchKernel	7.28e+11	88.5%	6.46e+11 93.1%
cuda_init_placeholders	7.28e+11	88.5%	6.46e+11 93.1%
RAJA::cuda::Reduce_Data<false, RAJA::reduce::sum<double>, double>;:grid_reduce(double*)	3.92e+11	47.7%	3.59e+11 51.6%
_INTERNAL_43_tmpxft_000131b5_00000000_6_DOT_Cuda_cpp1_ii_a3c0234b:;_shfl_xor_sync	3.40e+10	4.1%	2.77e+10 4.0%
_cuda_sm20_rem_s64	3.01e+10	3.7%	2.38e+10 3.4%
_INTERNAL_43_tmpxft_000131b5_00000000_6_DOT_Cuda_cpp1_ii_a3c0234b:;_shfl_xor_sync	2.83e+10	3.4%	2.30e+10 3.3%
void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterators::numeric_iterator<long	2.43e+10	3.0%	2.01e+10 2.9%
RAJA::cuda::Reduce<false, RAJA::reduce::sum<double>, double, false>;:=Reduce()	2.17e+10	2.6%	1.99e+10 2.9%
RAJA::operators::plus<double, double, double>;:operator()(double const&, double const&)&c	1.94e+10	2.4%	1.59e+10 2.3%
_cuda_sm20_dlv_s64	1.56e+10	1.9%	1.24e+10 1.8%
_syncthreads_or	1.38e+10	1.7%	1.32e+10 1.9%
rajaperf::stream::DOT::runCudaVariant<rajaperf::VariantID>;:(lambda(long)#1);:operator()(long) c	1.36e+10	1.7%	1.17e+10 1.7%
RAJA::internal::Privatizer<rajaperf::stream::DOT::runCudaVariant<rajaperf::VariantID>;:(lambda(long)#1);:lambda(d	1.32e+10	1.6%	1.24e+10 1.8%
rajaperf::stream::DOT::runCudaVariant<rajaperf::VariantID>;:(lambda(long)#1);:~VariantID()	1.24e+10	1.5%	1.17e+10 1.7%

Scope	GPU INST:Sum (l)	GPU STALL:Sum (l)	MIXINTEGRALADD3:Sum (l)
143: [l] void RAJA::forall<RAJA::policy::cuda::cuda_exec<256ul, true>, RAJA::TypedRe	7.25e+11	88.5%	6.46e+11 93.1%
723: [l] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<R	7.25e+11	88.5%	6.46e+11 93.1%
370: [l] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<R	7.25e+11	88.5%	6.46e+11 93.1%
183: [l] void RAJA::policy::cuda::forall_impl<RAJA::TypedRangeSegment<long	7.25e+11	88.5%	6.46e+11 93.1%
190: void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterator	7.25e+11	88.5%	6.46e+11 93.1%
145: [l] __wrapper__device_stub_forall_cuda_kernel<256ul, RAJA::Iterators	7.25e+11	88.5%	6.46e+11 93.1%
37: [l] __device_stub_ZNRAJA6policy4cuda4impl18forall_cuda_kerne	7.25e+11	88.5%	6.46e+11 93.1%
26: [l] cudaLaunchKernel<char>	7.25e+11	88.5%	6.46e+11 93.1%
209: cudaLaunchKernel	7.25e+11	88.5%	6.46e+11 93.1%
<cuda kernel>	7.25e+11	88.5%	6.46e+11 93.1%
void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJ	7.25e+11	88.5%	6.46e+11 93.1%
151: RAJA::internal::Privatizer<rajaperf::stream::DOT::runC	6.10e+11	74.5%	5.48e+11 78.9%
54: rajaperf::stream::DOT::runCudaVariant<rajaperf::Varia	5.97e+11	72.9%	5.35e+11 77.1%
129: RAJA::ReduceSum<RAJA::policy::cuda::cuda_red	5.85e+11	71.4%	5.24e+11 75.4%
190: RAJA::cuda::Reduce<false, RAJA::reduce::sum	5.73e+11	69.9%	5.12e+11 73.8%
848: RAJA::cuda::Reduce<false, RAJA::reduce::su	5.61e+11	68.5%	5.01e+11 72.2%
843: RAJA::cuda::Reduce_Data<false, RAJA::re	5.39e+11	65.9%	4.81e+11 69.3%
[l] inlined from reduce.hpp: 203	4.00e+11	48.8%	3.49e+11 50.1%
reduce.hpp: 293	1.31e+11	16.0%	1.24e+11 17.9%
loop at reduce.hpp: 203	8.78e+10	10.7%	7.09e+10 10.2%
loop at reduce.hpp: 203	4.02e+10	4.9%	3.25e+10 4.7%
205: _INTERNAL_43_tmpxft_000131	1.54e+10	1.9%	1.27e+10 1.8%
reduce.hpp: 205	1.50e+10	1.8%	1.19e+10 1.7%

Approximation of GPU Calling Contexts to Understand Performance

Scope	GPU INST:Sum (l)	GPU STALL:Sum (l)
143: [l] void RAJA::forall<RAJA::policy::cuda::cuda_exec<256ul, true>, RAJA::TypedRangeSegment<long, long>, _nv_	7.28e+11 88.5%	6.46e+11 93.1%
723: [l] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<RAJA::type_traits::is_indexset_polic	7.28e+11 88.5%	6.46e+11 93.1%
370: [l] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<RAJA::type_traits::is_indexset_pc	7.28e+11 88.5%	6.46e+11 93.1%
183: [l] void RAJA::policy::cuda::forall_impl<RAJA::TypedRangeSegment<long, long>, _nv_dl_wrapper_t<_nv	7.28e+11 88.5%	6.46e+11 93.1%
190: void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::iterators::numeric_iterator<long, long, lo	7.28e+11 88.5%	6.46e+11 93.1%
145: [l] _wrapper_device_stub_forall_cuda_kernel<256ul, RAJA::iterators::numeric_iterator<long int>, _n	7.28e+11 88.5%	6.46e+11 93.1%
37: [l] _device_stub_ZN4RAJA6policy4cuda4impl18forall_cuda_kernellm256ENS_9iterators16numeric	7.28e+11 88.5%	6.46e+11 93.1%
26: [l] cudaLaunchKernel<char>	7.28e+11 88.5%	6.46e+11 93.1%
209: cudaLaunchKernel	7.28e+11 88.5%	6.46e+11 93.1%
cuda_init_placeholders	7.28e+11 88.5%	6.46e+11 93.1%
RAJA::cuda::Reduce_Data<false, RAJA::reduce::sum<double>, double>::grid_reduce(double*)	3.92e+11 47.7%	3.59e+11 51.6%
_INTERNAL_43_tmpxft_000131b5_00000000_6_DOT_Cuda_cpp1_ii_a3c0234b::_shfl_xor_sync(3.40e+10 4.1%	2.77e+10 4.0%
_cuda_sm20_rem_s64	3.01e+10 3.7%	2.38e+10 3.4%
_INTERNAL_43_tmpxft_000131b5_00000000_6_DOT_Cuda_cpp1_ii_a3c0234b::_shfl_xor_sync(2.83e+10 3.4%	2.30e+10 3.3%
void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::iterators::numeric_iterator<long	2.43e+10 3.0%	2.01e+10 2.9%
RAJA::cuda::Reduce<false, RAJA::reduce::sum<double>, double, false>::~~Reduce()	2.17e+10 2.6%	1.99e+10 2.9%
RAJA::operators::plus<double, double, double>::operator()(double const&, double const&) c	1.94e+10 2.4%	1.59e+10 2.3%
_cuda_sm20_div_s64	1.56e+10 1.9%	1.24e+10 1.8%
_syncthreads_or	1.38e+10 1.7%	1.32e+10 1.9%
rajaperf::stream::DOT::runCudaVariant(rajaperf::VariantID)::(lambda(long)#1)::operator()(long) c	1.36e+10 1.7%	1.17e+10 1.7%
RAJA::internal::Privatizer<rajaperf::stream::DOT::runCudaVariant(rajaperf::VariantID)::(lambda(l	1.32e+10 1.6%	1.24e+10 1.8%
rajaperf::stream::DOT::runCudaVariant(rajaperf::VariantID)::(lambda(long)#1)::~VariantID()	1.24e+10 1.5%	1.17e+10 1.7%

flat profile of functions called by a GPU kernel

Approximation of GPU Calling Contexts to Understand Performance

- GPU code from C++ template-based programming models is complex
- NVIDIA GPUs collect flat PC samples
- Flat profiles for instantiations of complex C++ templates are inscrutable
- HPCToolkit reconstructs approximate GPU calling contexts
 - Reconstruct call graph from machine code
 - Infer calls at call sites
 - PC samples of call instructions indicate calls
 - Use call counts to apportion costs to call sites
 - PC samples in a routine

Scope	GPU INST:Sum (l)	GPU STALL:Sum (l)
143: [] void RAJA::forall<RAJA::policy::cuda::cuda_exec<256ul, true>, RAJA::TypedRangeSegment<long, long>, __nv_	7.28e+11 88.5%	6.46e+11 93.1%
723: [] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<RAJA::type_traits::is_indexset_polic	7.28e+11 88.5%	6.46e+11 93.1%
370: [] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<RAJA::type_traits::is_indexset_pc	7.28e+11 88.5%	6.46e+11 93.1%
183: [] void RAJA::policy::cuda::forall_impl<RAJA::TypedRangeSegment<long, long>, __nv_d1_wrapper_t<__nv_	7.28e+11 88.5%	6.46e+11 93.1%
190: void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterators::numeric_iterator<long, long, lo	7.28e+11 88.5%	6.46e+11 93.1%
145: [] __wrapper__device_stub_forall_cuda_kernel<256ul, RAJA::Iterators::numeric_iterator<long int>, __m	7.28e+11 88.5%	6.46e+11 93.1%
37: [] __device_stub_ZNRAJA6policy4cuda4impl18forall_cuda_kernellm256ENS_9Iterators16numeric	7.28e+11 88.5%	6.46e+11 93.1%
26: [] cudaLaunchKernel<char>	7.28e+11 88.5%	6.46e+11 93.1%
209: cudaLaunchKernel	7.28e+11 88.5%	6.46e+11 93.1%
cuda_init_placeholders		
RAJA::cuda::Reduce_Data<false, RAJA::reduce::sum<double>, double>;:grid_reduce(double*)	3.92e+11 47.7%	3.59e+11 51.6%
_INTERNAL_43_tmpxft_000131b5_00000000_6_DOT_Cuda_cpp1_ii_a3c0234b:;_shfl_xor_sync	3.40e+10 4.1%	2.77e+10 4.0%
_cuda_sm20_rem_s64	3.01e+10 3.7%	2.38e+10 3.4%
_INTERNAL_43_tmpxft_000131b5_00000000_6_DOT_Cuda_cpp1_ii_a3c0234b:;_shfl_xor_sync	2.83e+10 3.4%	2.30e+10 3.3%
void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterators::numeric_iterator<long	2.43e+10 3.0%	2.01e+10 2.9%
RAJA::cuda::Reduce<false, RAJA::reduce::sum<double>, double, false>;>:Reduce()	2.17e+10 2.6%	1.99e+10 2.9%
RAJA::operators::plus<double, double, double>;:operator()(double const&, double const&)&c	1.94e+10 2.4%	1.59e+10 2.3%
_cuda_sm20_dlv_s64	1.56e+10 1.9%	1.24e+10 1.8%
_syncthreads_or	1.38e+10 1.7%	1.32e+10 1.9%
rajaperf::stream::DOT::runCudaVariant(rajaperf::VariantID):;(lambda(long)#1):;operator()(long) c	1.36e+10 1.7%	1.17e+10 1.7%
RAJA::internal::Privatizer<rajaperf::stream::DOT::runCudaVariant(rajaperf::VariantID):;(lambda(long)#1):;operator()(long) c	1.32e+10 1.6%	1.24e+10 1.8%
rajaperf::stream::DOT::runCudaVariant(rajaperf::VariantID):;(lambda(long)#1):;~VariantID()	1.24e+10 1.5%	1.17e+10 1.7%

Scope	GPU INST:Sum (l)	GPU STALL:Sum (l)	MIXINTEGER.ADD3:Sum (l)
143: [] void RAJA::forall<RAJA::policy::cuda::cuda_exec<256ul, true>, RAJA::TypedRe	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
723: [] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<R	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
370: [] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<R	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
183: [] void RAJA::policy::cuda::forall_impl<RAJA::TypedRangeSegment<long	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
190: void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterator	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
145: [] __wrapper__device_stub_forall_cuda_kernel<256ul, RAJA::Iterator	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
37: [] __device_stub_ZNRAJA6policy4cuda4impl18forall_cuda_kerne	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
26: [] cudaLaunchKernel<char>	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
209: cudaLaunchKernel	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
<cuda kernel>			
void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJ	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
151: RAJA::internal::Privatizer<rajaperf::stream::DOT::runC	6.10e+11 74.5%	5.48e+11 78.9%	3.53e+09 78.8%
54: rajaperf::stream::DOT::runCudaVariant(rajaperf::Varia	5.97e+11 72.9%	5.35e+11 77.1%	3.52e+09 78.6%
129: RAJA::ReduceSum<RAJA::policy::cuda::cuda_red	5.85e+11 71.4%	5.24e+11 75.4%	3.51e+09 78.4%
190: RAJA::cuda::Reduce<false, RAJA::reduce::sum	5.73e+11 69.9%	5.12e+11 73.8%	3.50e+09 78.2%
848: RAJA::cuda::Reduce<false, RAJA::reduce::su	5.61e+11 68.5%	5.01e+11 72.2%	3.50e+09 78.0%
843: RAJA::cuda::Reduce_Data<false, RAJA::re	5.39e+11 65.9%	4.81e+11 69.3%	3.47e+09 77.5%
inlined from reduce.hpp: 203	4.00e+11 48.8%	3.49e+11 50.1%	3.41e+09 76.1%
reduce.hpp: 293	1.31e+11 16.0%	1.24e+11 17.9%	7.48e+06 0.2%
loop at reduce.hpp: 203	8.78e+10 10.7%	7.09e+10 10.2%	8.15e+08 18.2%
loop at reduce.hpp: 203	4.02e+10 4.9%	3.25e+10 4.7%	4.35e+08 9.7%
205: _INTERNAL_43_tmpxft_000131	1.54e+10 1.9%	1.27e+10 1.8%	3.01e+08 6.7%
reduce.hpp: 205	1.50e+10 1.8%	1.19e+10 1.7%	1.15e+08 2.6%

Approximation of GPU Calling Contexts to Understand Performance

Scope	GPU INST:Sum (I)	GPU STALL:Sum (I)	MIX:INTEGER.ADD3:Sum (I)
↳ 143: [I] void RAJA::forall<RAJA::policy::cuda::cuda_exec<256ul, true>, RAJA::TypedRa	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 723: [I] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t<R	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 370: [I] std::enable_if<camp::concepts::all_of<camp::concepts::metalib::negate_t	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 183: [I] void RAJA::policy::cuda::forall_impl<RAJA::TypedRangeSegment<long	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 190: void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJA::Iterator	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 145: [I] __wrapper__device_stub_forall_cuda_kernel<256ul, RAJA::Iterators	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 37: [I] __device_stub_ZN4RAJA6policy4cuda4impl18forall_cuda_kerne	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 26: [I] cudaLaunchKernel<char>	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 209: cudaLaunchKernel	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
<cuda kernel>	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
void RAJA::policy::cuda::impl::forall_cuda_kernel<256ul, RAJ	7.25e+11 88.5%	6.46e+11 93.1%	4.12e+09 91.9%
↳ 151: RAJA::internal::Privatizer<rajaperf::stream::DOT::runCu	6.10e+11 74.5%	5.48e+11 78.9%	3.53e+09 78.8%
↳ 54: rajaperf::stream::DOT::runCudaVariant(rajaperf::Varia	5.97e+11 72.9%	5.35e+11 77.1%	3.52e+09 78.6%
↳ 129: RAJA::ReduceSum<RAJA::policy::cuda::cuda_red	5.85e+11 71.4%	5.24e+11 75.4%	3.51e+09 78.4%
↳ 190: RAJA::cuda::Reduce<false, RAJA::reduce::sum	5.73e+11 69.9%	5.12e+11 73.8%	3.50e+09 78.2%
↳ 848: RAJA::cuda::Reduce<false, RAJA::reduce::su	5.61e+11 68.5%	5.01e+11 72.2%	3.50e+09 78.0%
↳ 843: RAJA::cuda::Reduce_Data<false, RAJA::re	5.39e+11 65.9%	4.81e+11 69.3%	3.47e+09 77.5%
[I] inlined from reduce.hpp: 203	4.00e+11 48.8%	3.48e+11 50.1%	3.41e+09 76.1%
reduce.hpp: 293	1.31e+11 16.0%	1.24e+11 17.9%	7.48e+06 0.2%
loop at reduce.hpp: 203	8.78e+10 10.7%	7.09e+10 10.2%	8.15e+08 18.2%
loop at reduce.hpp: 203	4.02e+10 4.9%	3.25e+10 4.7%	4.35e+08 9.7%
↳ 205: _INTERNAL_43_tmpxft_000131	1.54e+10 1.9%	1.27e+10 1.8%	3.01e+08 6.7%
reduce.hpp: 205	1.50e+10 1.8%	1.19e+10 1.7%	1.15e+08 2.6%

CPU Calling Context

GPU API Node

GPU Calling Context

GPU Loops

GPU Hotspot

Approximate Reconstruction of GPU Calling Context Trees

- **Problem**

- Unwinding call stacks on GPU is costly for each GPU thread
- NVIDIA's CUPTI does not provide an unwinding API

- **Challenges**

- GPU functions may be invoked from different call sites
- Need to decide how to attribute costs to each call site

- **Solution**

- Reconstruct GPU calling context tree from flat instruction samples and static GPU call graph





Approximate Reconstruction of GPU Calling Context Trees

1. Construct a GPU static call graph based on functions and call instructions. Initialize call counts using counts or samples of call instructions.

2. For call graphs based on samples: if a function has samples and no incoming call edge has a non-zero weight, assign each of its incoming call edges a weight of 1; repeat for call edges of callers until at least one incoming call edge has samples.

3. Identify strongly connected components (SCCs) using Tarjan's algorithm. Rewire call graph, removing SCC internal structure and linking external calls to SCC.

4. Build CCT by splitting call graph. Like gprof, assume that every call to a function has equal cost. Apportion costs of each function among its call sites according to ratios of calls from each call site.







Approximate Reconstruction of GPU Calling Context Trees

1. Construct a GPU static call graph based on functions and call instructions. Initialize call counts using counts or samples of call instructions.

2. For call graphs based on samples: if a function has samples and no incoming call edge has a non-zero weight, assign each of its incoming call edges a weight of 1; repeat for call edges of callers until at least one incoming call edge has samples.



3. Identify strongly connected components (SCCs) using Tarjan's algorithm. Rewire call graph, removing SCC internal structure and linking external calls to SCC.

4. Build CCT by splitting call graph. Like gprof, assume that every call to a function has equal cost. Apportion costs of each function among its call sites according to ratios of calls from each call site.


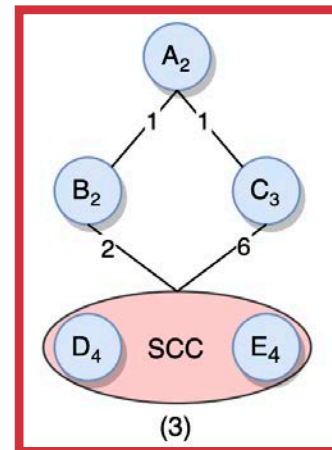


Approximate Reconstruction of GPU Calling Context Trees

1. Construct a GPU static call graph based on functions and call instructions. Initialize call counts using counts or samples of call instructions.
2. For call graphs based on samples: if a function has samples and no incoming call edge has a non-zero weight, assign each of its incoming call edges a weight of 1; repeat for call edges of callers until at least one incoming call edge has samples.



3. Identify strongly connected components (SCCs) using Tarjan's algorithm. Rewire call graph, removing SCC internal structure and linking external calls to SCC.



4. Build CCT by splitting call graph. Like gprof, assume that every call to a function has equal cost. Apportion costs of each function among its call sites according to ratios of calls from each call site.

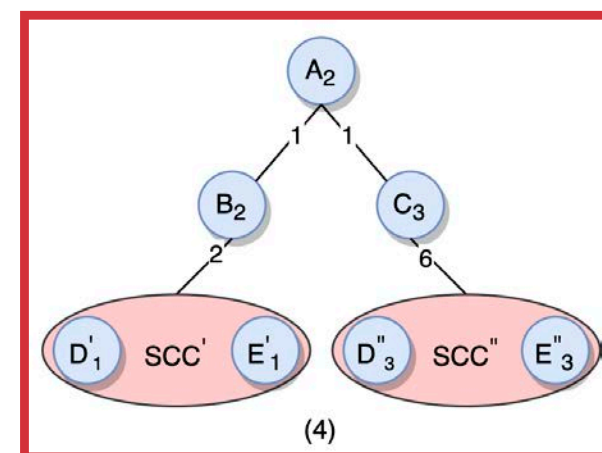



Approximate Reconstruction of GPU Calling Context Trees

1. Construct a GPU static call graph based on functions and call instructions. Initialize call counts using counts or samples of call instructions.

2. For call graphs based on samples: if a function has samples and no incoming call edge has a non-zero weight, assign each of its incoming call edges a weight of 1; repeat for call edges of callers until at least one incoming call edge has samples.

3. Identify strongly connected components (SCCs) using Tarjan's algorithm. Rewire call graph, removing SCC internal structure and linking external calls to SCC.

4. Build CCT by splitting call graph. Like gprof, assume that every call to a function has equal cost. Apportion costs of each function among its call sites according to ratios of calls from each call site.



ECP Quicksilver GPU Proxy Application: Detailed Profile View

Compute Node

- 2 Power9
- 6xNVIDIA GPU

• Optimized (-O2) compilation with nvcc

• 1 GPU stream

• Detailed measurement and attribution using PC sampling

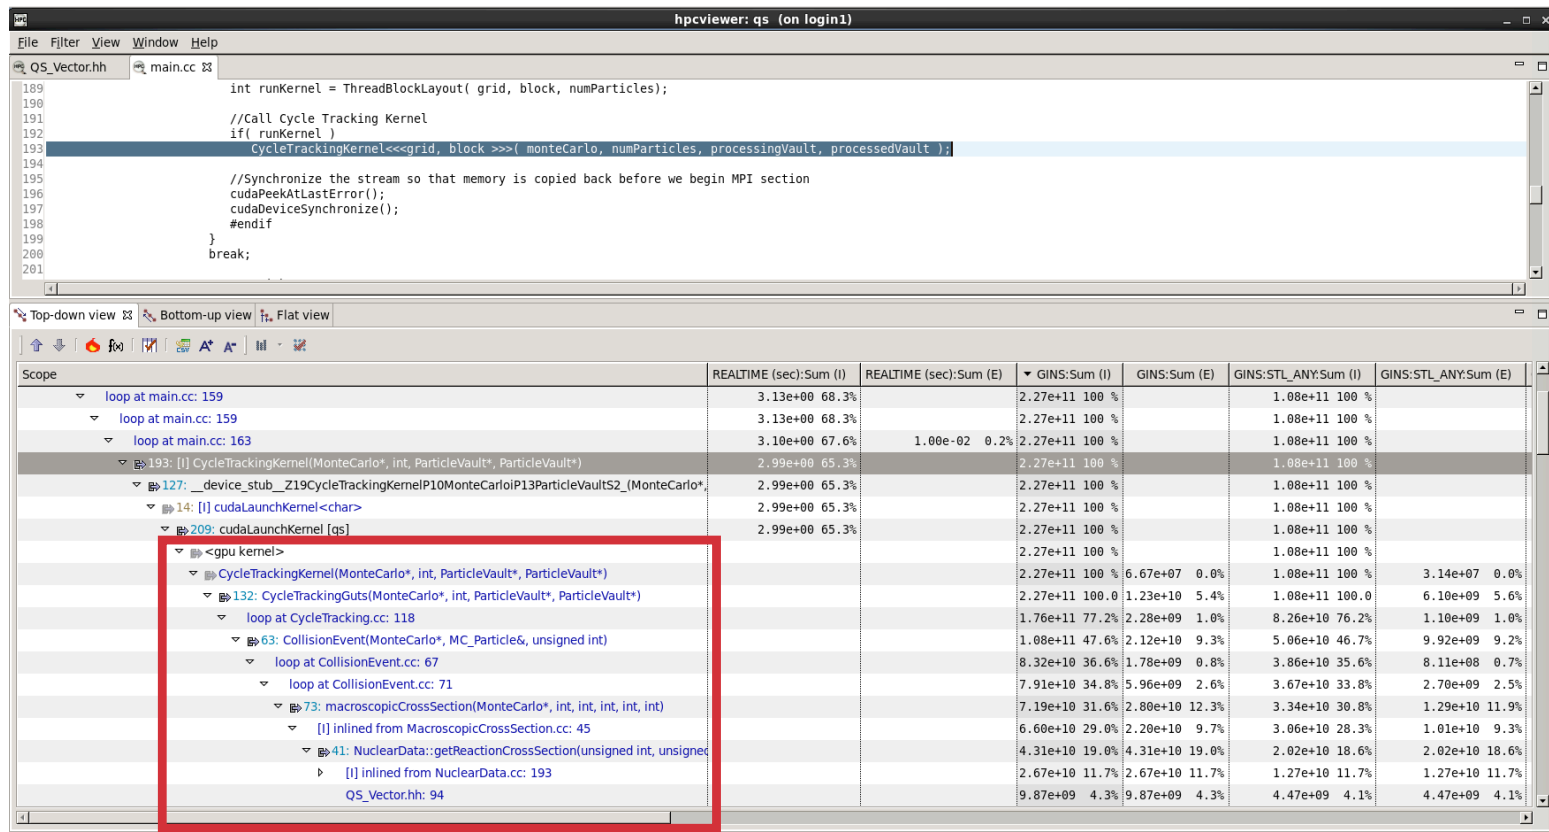
• Reconstruct approximate call graph on GPU from flat PC samples

• Attribute information to heterogeneous calling context including

- CPU calling context
- GPU kernel
- GPU calling context
- GPU loops
- GPU statements

• Metrics

- instructions executed
- instruction stalls and reasons
- GPU utilization



Quicksilver GPU CUDA Example: Detailed Profile View

Detailed Attribution on GPUs

- Optimized (-O2) compilation with nvcc
- 1 GPU stream
- Detailed measurement and attribution using PC sampling
- Reconstruct approximate call graph on GPU from flat PC samples
- Attribute information to heterogeneous calling context including
 - CPU calling context
 - GPU kernel
 - GPU calling context
 - GPU loops
 - GPU statements
- Metrics
 - instructions executed
 - instruction stalls and reasons
 - GPU utilization

```
▼ <gpu kernel>
  ▼ CycleTrackingKernel(MonteCarlo*, int, ParticleVault*, ParticleVault*)
    ▼ 132: CycleTrackingGuts(MonteCarlo*, int, ParticleVault*, ParticleVault*)
      ▼ 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, int, int)
                ▼ [I] inlined from MacroscopicCrossSection.cc: 45
                  ▼ 41: NuclearData::getReactionCrossSection(unsigned int, unsigned int)
                    ▸ [I] inlined from NuclearData.cc: 193
                      QS_Vector.hh: 94
```

Support for OpenMP TARGET

- HPCToolkit implementation of OMPT OpenMP API
 - host monitoring
 - leverages callbacks for regions, threads, tasks
 - employs OMPT API for call stack introspection
 - GPU monitoring
 - leverages callbacks for device initialization, kernel launch, data operations
 - reconstruction of user-level calling contexts
- Leverages implementation of OMPT in LLVM OpenMP and libomptarget

ECP QMCPACK Project: miniqmc using OpenMP TARGET (Power9 + NVIDIA V100)

Reconstruct full calling contexts that include

- Outlined procedures for OpenMP parallel regions
- Offloaded OpenMP TARGET computation and synchronization

Scope	CPUTIME (usec):SUM	KERNEL:TIME (us):SUM	XDMOV:TIME (us):SUM	SYNC:TIME (us):SUM
<program root>	9.06e+07 74.1%	5.63e+05 100 %	8.87e+04 100 %	1.80e+06 100 %
main	9.06e+07 74.1%	5.57e+05 99.1%	8.80e+04 99.2%	1.78e+06 99.1%
loop at miniqmc_sync_move.cpp: 432	1.75e+07 14.3%	4.81e+05 85.4%	6.57e+04 74.0%	1.49e+06 82.7%
434: omp_outlined..54	1.68e+07 13.8%	4.81e+05 85.4%	6.57e+04 74.0%	1.49e+06 82.7%
435: [i] omp_outlined_debug_.53	1.68e+07 13.8%	4.81e+05 85.4%	6.57e+04 74.0%	1.49e+06 82.7%
loop at miniqmc_sync_move.cpp: 435	1.08e+07 8.8%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
loop at miniqmc_sync_move.cpp: 459	1.08e+07 8.8%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
loop at miniqmc_sync_move.cpp: 461	8.33e+06 6.8%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
480: qmcpusplus::WaveFunction::flex_ratioGrad(std::vector<double> const&, double)	8.33e+06 6.8%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
415: qmcpusplus::WaveFunction::ratioGrad(qmcpusplus::DiracDeterminant const&, double)	7.74e+06 6.3%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
qmcpusplus::DiracDeterminant<qmcpusplus::Delay>::Delay	7.69e+06 6.3%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
97: qmcpusplus::einspline_spo_omp<double>::einspline_spo_omp	7.69e+06 6.3%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
340: qmcpusplus::einspline_spo_omp<double>::einspline_spo_omp	7.65e+06 6.2%	3.80e+05 67.6%	4.84e+04 54.6%	1.19e+06 66.1%
loop at einspline_spo_omp.cpp: 304				
311: <omp tgt kernel>		3.80e+05 67.6%		
omp_offloading_fd00_88088b_ZN: einspline_spo_omp.cpp: 316		3.80e+05 67.6%		
<cuda sync>				1.19e+06 66.1%

Support for RAJA and Kokkos C++ Template-based Models

- RAJA and Kokkos provide portability layers atop C++ template-based programming abstractions
- HPCToolkit employs binary analysis to recover information about procedures, inlined functions and templates, and loops
 - Enables both developers and users to understand complex template instantiation present with these models

ECP EXAALT Project: LAMMPS using Kokkos over CUDA (Power9 + NVIDIA V100)

The screenshot shows the hpcviewer interface with a code editor and a call stack. The code editor displays the following snippet:

```
331 // Invoke the driver function on the device
332 cuda_parallel_launch_constant_memory<DriverType>
333 <<< grid, block, shmem, cuda_instance->m_stream >>>();
334
335 #if defined( KOKKOS_ENABLE_DEBUG_BOUNDS_CHECK )
336 ...
```

The call stack on the right shows the following entries:

Scope	KERNEL:TIME (us):Sum (l)
65: LAMMPS_NS::Input::file()	2.40e+07 100.0
loop at input.cpp: 165	2.40e+07 100.0
229: LAMMPS_NS::Input::execute_command()	2.40e+07 100.0
864: void LAMMPS_NS::Input::command_creator<LAMMPS_NS::Run>(LAMMPS_NS::LAMMPS*, int, char**)	2.35e+07 97.9%
881: LAMMPS_NS::Run::command(int, char**)	2.35e+07 97.9%
182: LAMMPS_NS::VerletKokkos::run(int)	2.35e+07 97.9%
loop at verlet_kokkos.cpp: 321	2.35e+07 97.9%
477: LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>::compute(int, int)	1.66e+07 69.4%
121: s_EV_FLOAT LAMMPS_NS::pair_compute<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, void>(LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, int, int)	1.66e+07 69.4%
911: s_EV_FLOAT LAMMPS_NS::pair_compute_neighlist<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1u, void>(LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, int, int)	1.66e+07 69.4%
900: [I] void Kokkos::parallel_for<LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%
224: [I] Kokkos::Impl::ParallelFor<LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%
540: Kokkos::Impl::CudaParallelLaunch<Kokkos::Impl::ParallelFor<LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%
332: [I] cuda_parallel_launch_constant_memory<Kokkos::Impl::ParallelFor<LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%
106: [I] wrapper_device_stub_cuda_parallel_launch_constant_memory<Kokkos::Impl::ParallelFor<LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%
268: __device_stub_ZN6Kokkos4Impl36cuda_parallel_launch_constant_memoryINS0_11ParallelForINS0_11PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%
265: [I] cudaLaunchKernel<char>	1.66e+07 69.2%
209: <cuda kernel>	1.66e+07 69.2%
34: void Kokkos::Impl::cuda_parallel_launch_constant_memory<Kokkos::Impl::ParallelFor<LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>(LAMMPS_NS::PairComputeFuncor<LAMMPS_NS::PairLJCutKokkos<Kokkos::Cuda>, 1, true>, int)	1.66e+07 69.2%

Reconstruct full calling contexts that include

- Inlined Kokkos templates
- Offloaded Kokkos CUDA computation

Deriving GPU Metrics

- **Problem**

- GPU PC sampling cannot be used in the same pass with metric collection
- Nsight-compute runs nine passes to collect multiple metrics for kernels

- **Our approach**

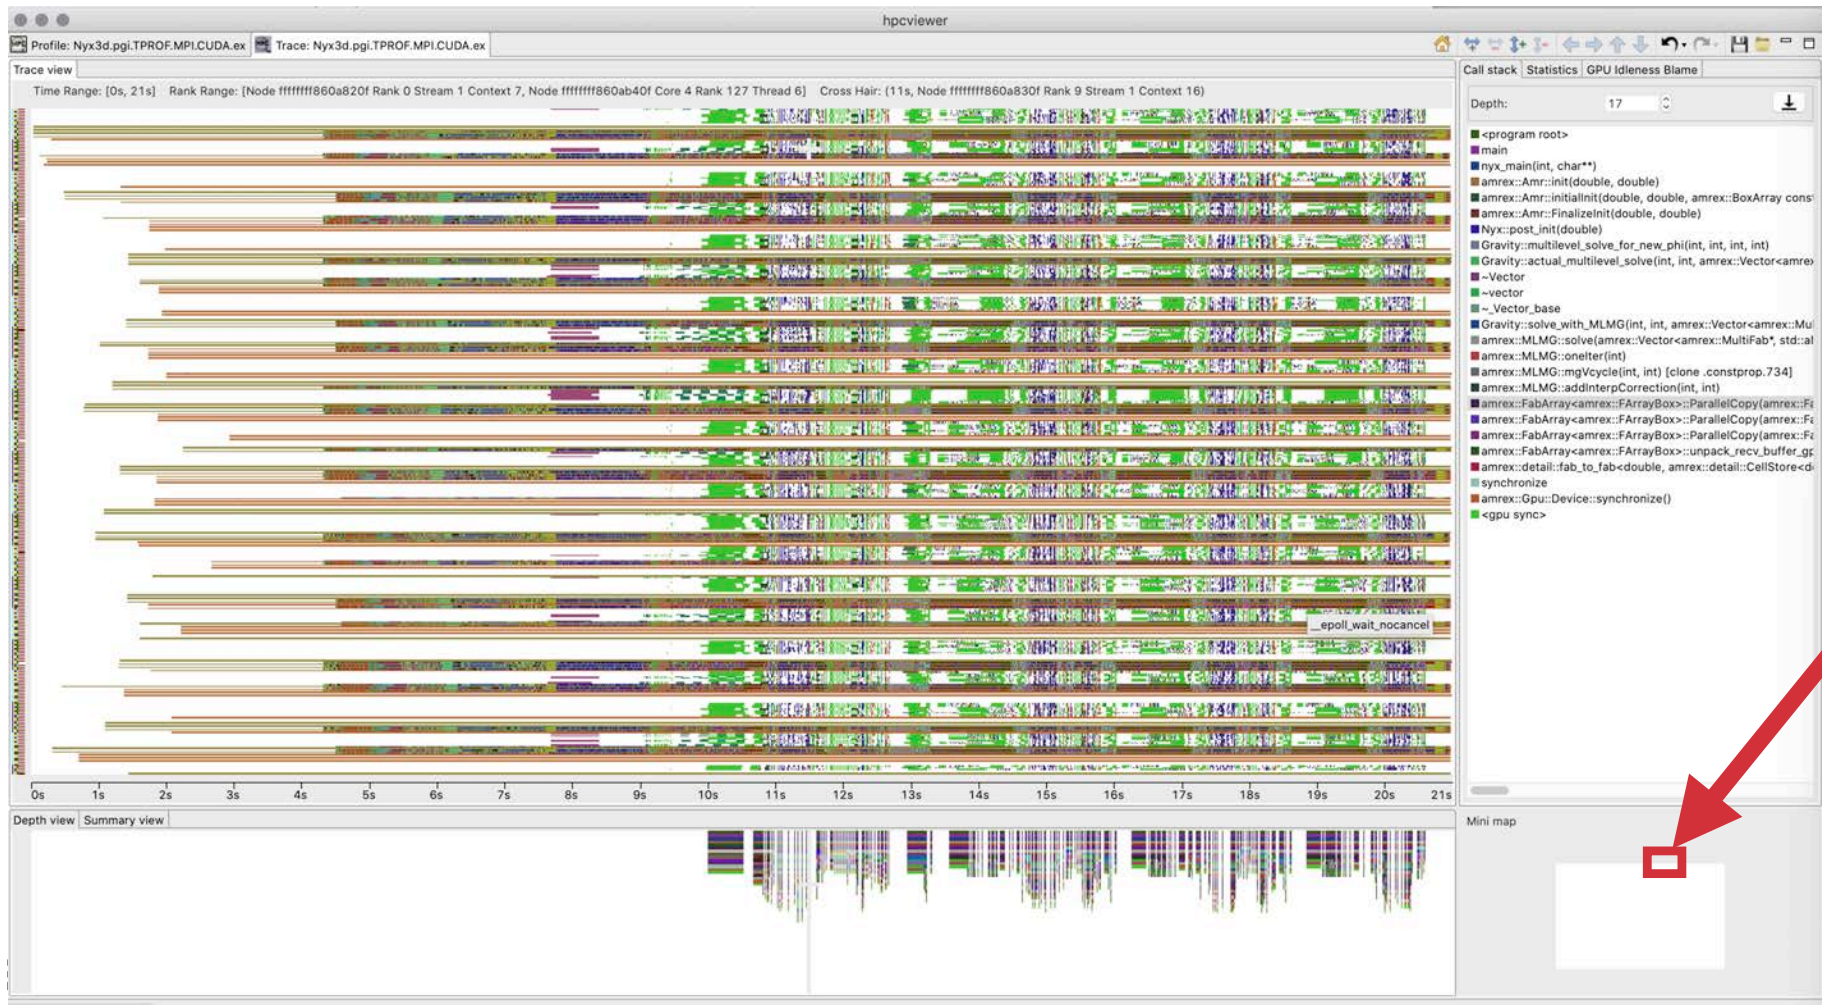
- Derive multiple metrics using PC sampling and other activity records
 - e.g., GPU SM utilization, GPU occupancy, ...

Measurement of GPU-accelerated NAMD3 using Charm++

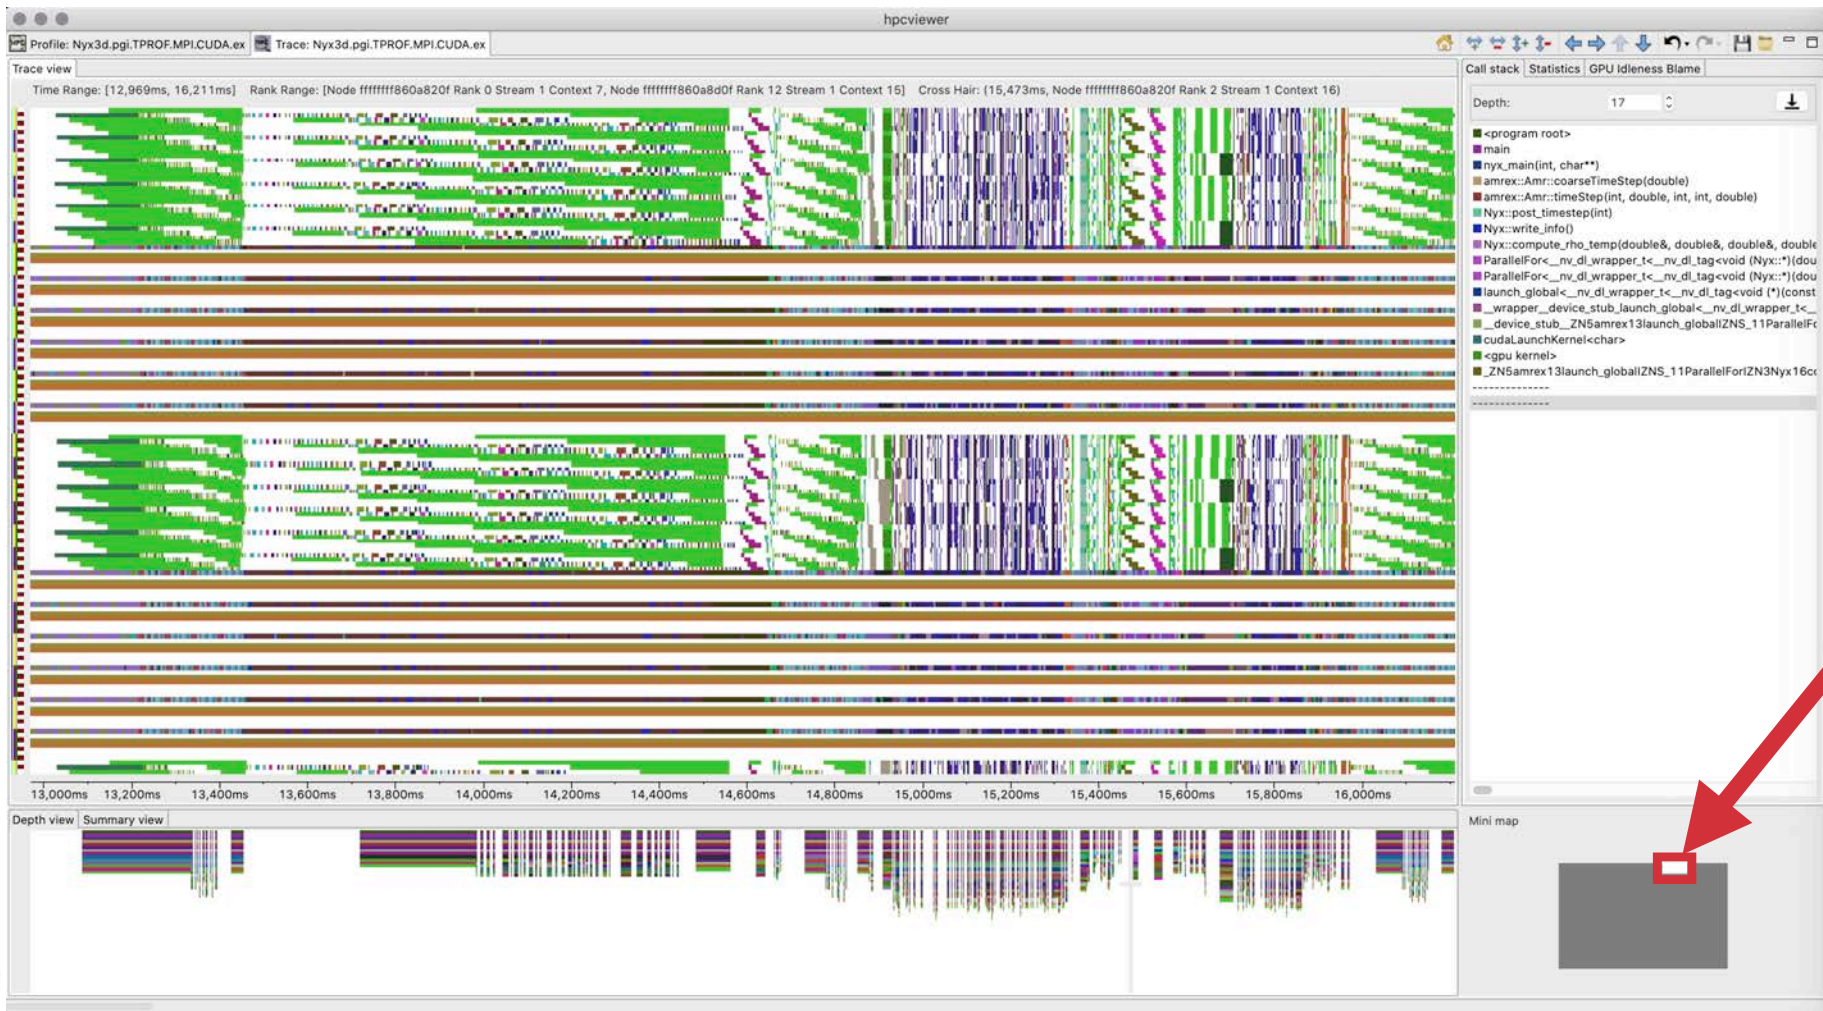
identify a costly line in context
key cost: memory stalls
full GPU utilization by the kernel
kernel characteristics

Scope	GINS:Sum (I)	GINS:STL_ANY:Sum	GINS:STL_GMEM:	GSAMP:UTIL	GKER:THR_REG:	GKER:BLK_THF	GKER:BLK_	GKER:OCC_1
Experiment Aggregate Metrics	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	300.00 %	129	896	5862	207.62 %
<partial call paths>	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	100.00 %				
make_fcontext [namd3]	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	100.00 %				
CthStartThread [namd3]	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	100.00 %				
Sequencer::algorithm()	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	100.00 %				
loop at Sequencer.C: 269	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	100.00 %				
330: Sequencer::integrate_CUDA_SOA(int)	2.39e+10 100.0%	9.78e+09 100.0%	4.82e+09 100.0%	100.00 %				
loop at Sequencer.C: 732	2.32e+10 97.0%	9.49e+09 97.0%	4.68e+09 97.1%	100.00 %				
819: Sequencer::runComputeObjectsCUDA(int, int, int)	1.81e+10 75.7%	7.03e+09 72.0%	3.91e+09 81.1%	100.00 %				
484: CudaComputeNonbonded::launchWork()	1.17e+10 49.0%	3.97e+09 40.6%	2.28e+09 47.2%	100.00 %				
1121: CudaComputeNonbonded::doForce()	1.13e+10 47.1%	3.75e+09 38.4%	2.17e+09 45.1%	100.00 %				
1347: CudaComputeNonbondedKernel::nonbondedForce(CudaTileLi	1.12e+10 46.9%	3.73e+09 38.1%	2.17e+09 44.9%	100.00 %				
loop at <unknown file> [namd3]: 0	1.12e+10 46.9%	3.73e+09 38.1%	2.17e+09 44.9%	100.00 %				
_device_stub_Z20nonbondedForceKernelLb1ELb1ELb0ELbC	8.13e+09 33.9%	2.90e+09 29.7%	1.72e+09 35.7%	100.00 %				
cudaLaunchKernel [namd3]	8.13e+09 33.9%	2.90e+09 29.7%	1.72e+09 35.7%	100.00 %				
<gpu kernel>	8.13e+09 33.9%	2.90e+09 29.7%	1.72e+09 35.7%	100.00 %				
174: nonbondedForceKernel<true, true, false, false, f	8.13e+09 33.9%	2.90e+09 29.7%	1.72e+09 35.7%	100.00 %				
loop at CudaComputeNonbondedKernel.cu: 673	8.00e+09 33.4%	2.86e+09 29.3%	1.71e+09 35.4%					
loop at CudaComputeNonbondedKernel.cu: 847	7.53e+09 31.4%	2.70e+09 27.7%	1.59e+09 33.0%					
CudaComputeNonbondedKernel.cu: 53	4.25e+09 17.7%	1.59e+09 16.1%	1.59e+09 33.0%					
CudaComputeNonbondedKernel.cu: 43	1.20e+09 5.0%	4.82e+08 4.9%						

Nyx with CUDA: Trace of Multi-rank Multi-GPU Executions



Nyx with CUDA: Trace of Multi-rank Multi-GPU Executions



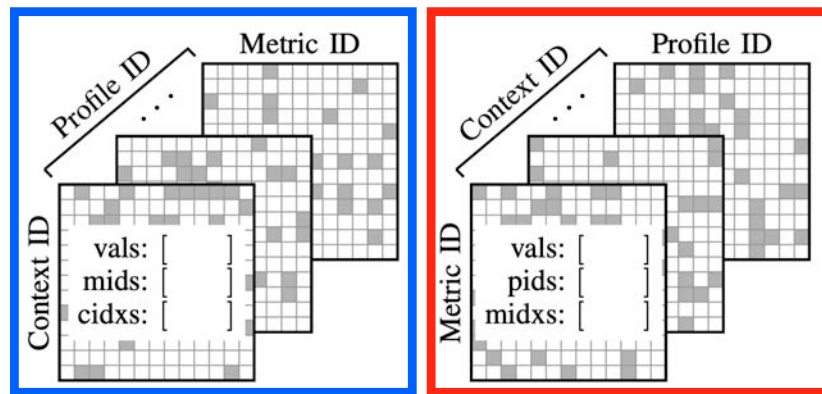
Scalable Analysis of Performance Data

- **When to reduce profile data?**

- After termination: Linux perf, NVIDIA nvvp, and Paraver record detailed traces
- At termination
 - Scalasca, Tau, Vampir use MPI to unify profile data into CUBE format
 - HPCToolkit saves separate profiles and traces per thread

- **Scalable analysis of performance data using out-of-core algorithms**

- Inspect profiles and balance across ranks by aggregate size
- Unify call stacks from all threads
- Overlay static information on calling context trees: procedures, inline functions, loops, stmts
- Generate computed statistics: aggregate and per profile
- Write out two sparse outputs
 - [profile-major-sparse database](#)
 - [calling-context-major-sparse database](#)
- Implementation: MPI + OpenMP




Scalable Analysis of Performance Data: MPI vs. OpenMP

- **Compare relative speedup of all MPI ranks vs. all OpenMP threads**

- Input: 8K profiles
- use all OpenMP threads at various scales
- use all MPI ranks at various scales

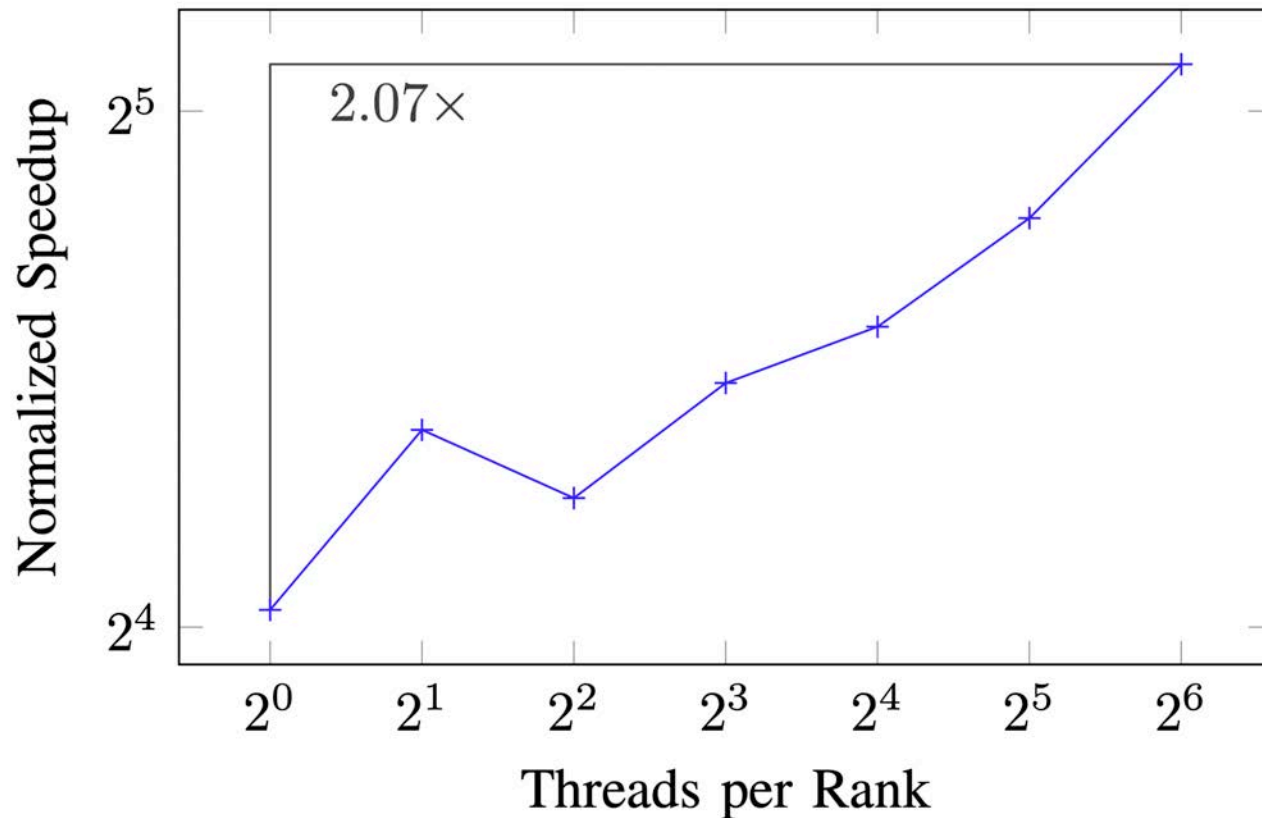
- **Findings**

- OpenMP has superior scaling and speedup
- MPI speedup degrades using more than one thread per core
- OpenMP speedup improves from 64 to 128 threads
- Note: need 2+ threads per core to use all memory B/W on KNL



Scalable Analysis of Performance Data: Threads vs. Ranks

- Use 64 cores and threads
- Compare balance of ranks vs. threads
 - 64 ranks
 - 32 ranks, 2 threads/rank
 - ...
 - 1 rank, 64 threads/rank
- Findings
 - Performance increases by trading threads for ranks



Scalable Analysis of Performance Data: Value of Sparsity

- **Assess the space savings of sparse profiles**

- AMD2006 CPU
 - 1 metric
 - 9 metrics, including some rare metrics
- Nyx GPU
- LAMMPS GPU

- **Findings**

- as much as 21x space reduction for measurements
- as much 337x reduction for output data

		Size (MiB)		
Dataset		Dense	Sparse	Ratio
AMG2006 (1)	M	659.0	911.0	0.723×
	D	7370.0	836.0	8.819×
AMG2006 (9)	M	21.7	11.1	1.956×
	D	2290.0	33.5	68.34 ×
Nyx	M	5890.0	278.0	21.14 ×
	D	130 GiB	601.0	221.5 ×
LAMMPS	M	85.5	5.23	16.35 ×
	D	8250.0	24.5	336.9 ×

Scalable Analysis of Performance Data: 64K profiles of AMG2006

Input

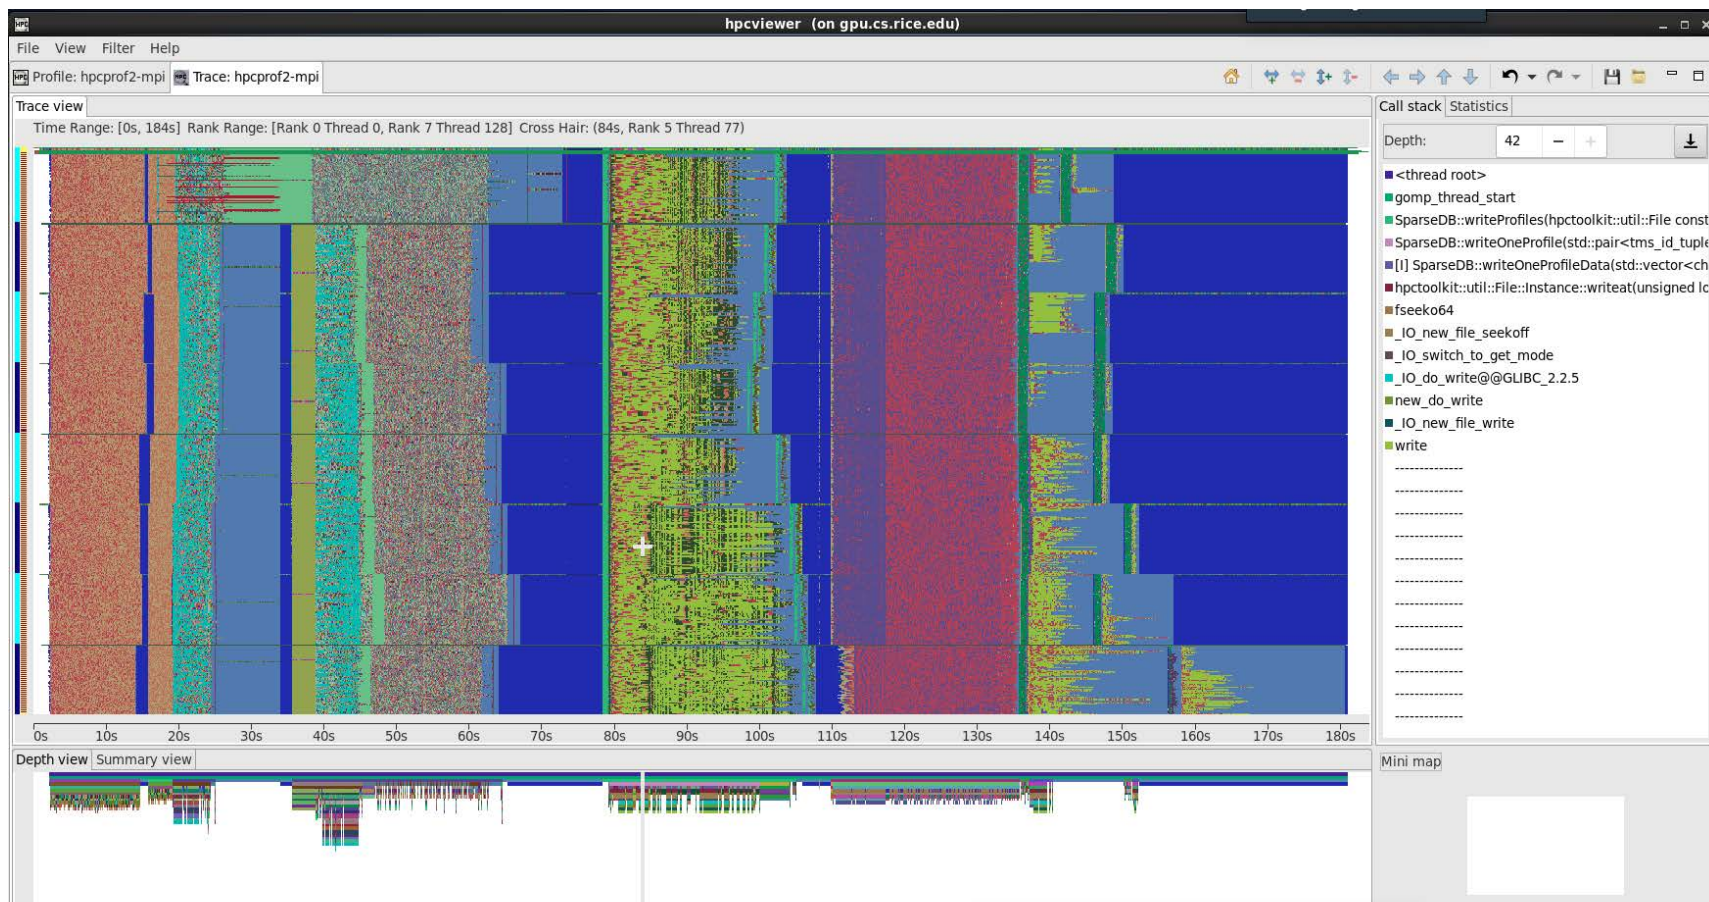
- 5GB profiles
- 225GB traces

Analysis

- 8 KNL nodes
- 1 rank / node
- 128T / rank

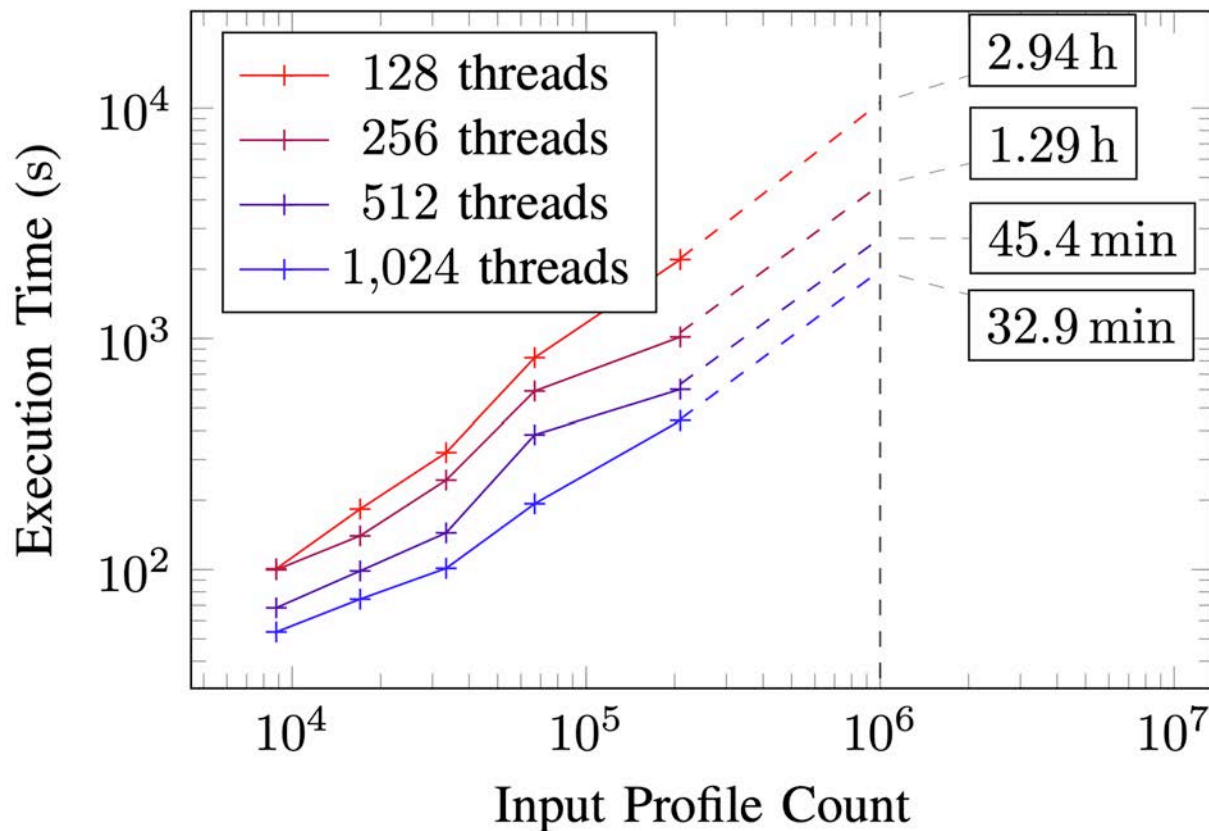
Execution time

- 184s



Scalable Analysis of Performance Data

- **Assess scalability vs. profile counts**
- **Explore scaling using 128-1024 threads**
 - 128 threads per rank
- **Extrapolate to 1M profiles**
- **Findings**
 - Good scaling with number of profiles
 - Relative efficiency drops to 67% when scaling from 128 to 1024 threads
 - 200K profiles (9GB total) is a small problem for 1024 threads



Outline

- **Performance measurement and analysis challenges for GPU-accelerated supercomputers**
- **Introduction to HPCToolkit performance tools**
 - Overview of HPCToolkit components and their workflow
 - HPCToolkit's graphical user interfaces
- **Analyzing the performance of GPU-accelerated supercomputers with HPCToolkit**
 - Overview of HPCToolkit's GPU performance measurement capabilities
 - Collecting measurements
 - Analysis and attribution
 - Scalable analysis of performance data
- **Status, ongoing work, final remarks**

Status for Various GPUs

Vendor	Coarse-grain measurement	Fine-grain measurement	Tracing	Binary analysis: loops, inlined code
NVIDIA	CUPTI	PC sampling	CUPTI	nvdiasm + dyninst
Intel	OpenCL and Level 0	GTPin instrumentation	OpenCL callbacks	IGA + dyninst
AMD	Roctracer	—	Roctracer	dyninst + emerging native decoder

Many of these capabilities are on different branches

Ongoing Work

- **Interface**

- Emerging GPU Performance Advisor tool for NVIDIA GPUs
 - attributes instruction stalls with backward slicing, analyzes code, offers advice about most promising improvements
- New integrated user interface supports both profiles and traces
 - Automated serialization analysis of CPU and GPU traces in hpctraceviewer GUI
- Performance analysis of machine learning frameworks: Pytorch, Tensorflow

- **Internals**

- Refinement of implementation atop Intel's Level 0
- Binary analysis of AMD GPU instructions
- Refining scalable aggregation

Final Remarks

- **Nice to work with national labs and have early involvement in a big procurement**
 - Amplifies our ability to affect vendor hardware and software in the near term
- **Software development challenges are myriad**
 - Developing tools for three GPU software stacks at the same time is ridiculous
 - Building capabilities ahead of current vendor hardware and software
 - AMD and Intel software is a work in progress
 - Intel: unstable with significant flaws
 - Both: API-breaking changes are common
 - Relying on vendor closed-source components is a challenge
 - standards specify only an API, but internals matter for tools that see all
 - undocumented behaviors about things that matter
 - missing capabilities, e.g. need excellent DWARF mappings for optimized GPU code
 - NVIDIA serializes kernels to facilitate measurement with PC sampling

Acknowledgments

- **Current funding**

- DOE Exascale Computing Project (Subcontract 4000151982)
- NSF Software Infrastructure for Sustained Innovation (Collaborative Agreement 1450273)
- DOE Labs: ANL (Subcontract 9F-60073), Tri-labs (LLNL Subcontract B633244)
- Industry: AMD

- **Team**

- Rice University
 - HPCToolkit PI: Prof. John Mellor-Crummey
 - Research staff: Laksono Adhianto, Mark Krentel, Xiaozhu Meng, Scott Warren
 - Contractor: Marty Itzkowitz
 - Students: Jonathon Anderson, Aaron Cherian, Dejan Grubisic, Yumeng Liu, Keren Zhou
 - Recent summer interns: Vladimir Indjic, Tijana Jovanovic, Aleksa Simovic
- University of Wisconsin – Madison
 - Dyninst PI: Prof. Barton Miller