

Preparing for Performance Analysis at Exascale



Jonathon Anderson, Yumeng Liu, John Mellor-Crummey

Rice University

STW 2022

June 20, 2022



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Performance Analysis in the Exascale Era

- Forthcoming exascale systems pose new challenges
 - Loss may only appear at the full scale of the machine, 10000s of compute nodes
 - Performance must be gathered from every application thread: **very large data**
 - Rich performance data must be gathered for a complete picture, over 130 GPU metrics
 - Some metrics are only present for GPU code regions: **very sparse data**
- Performance tools must handle these issues...
 - ...But still provide detailed, useful analysis results!

Outline

- HPCToolkit at a glance
- Exploiting natural sparsity in performance data
- Streaming aggregation for highly-parallel processing of performance data
- Evaluation
- Conclusions
- Ongoing work: exploiting distributed object storage

HPCToolkit: Fine-grain Measurement and Attribution within Kernels

```

Profile: PeleC3d.gnu.TPROF.CUDA.ex
reactor.cpp | Metric properties
438 UserData udata = static_cast<ARKODEUserData*>(user_data);
439 udata->dt_save = t;
440
441 #ifndef AMREX_USE_GPU
442 const auto ec = amrex::Gpu::ExecutionConfig(udata->ncells_d);
443 amrex::launch_global<<<
444   udata->nbBLocks, udata->nbThreads, ec.sharedMem, udata->stream>>>{
445   [=] AMREX_GPU_DEVICE() noexcept {
446     for (int icell = blockDim.x * blockIdx.x + threadIdx.x,
447          stride = blockDim.x * gridDim.x;
448          icell < udata->ncells_d; icell += stride) {
449       fKernelSpec(
450         icell, udata->dt_save, udata->ireactor_type, yvec_d, ydot_d,
451         udata->rhoe_init_d, udata->rhoesrc_ext_d, udata->Rysrc_d);
452     };
453   });
454 #else
455   for (int icell = 0; icell < udata->ncells_d; icell++) {
456     fKernelSpec(

```

Cause:
passed udata structure pointer to lambda capture

Improvement:
pass udata components as scalars
<https://github.com/AMReX-Combustion/PelePhysics/pull/192>
4% speedup on PeleC PMF drm19 test case

| Scope | GINS.[0,0] (I) | GINS.[0,0] (E) | GINS.STL_ANY.[0,0] (I) | GINS.STL_ANY.[0,0] (E) | GINS.STL_GMEM.[0,0] (I) | GINS.STL_GMEM.[0,0] (E) |
|--|----------------|----------------|------------------------|------------------------|-------------------------|-------------------------|
| loop at AMReX_Amr.cpp: 2061 | 1.24e+13 88.6% | | 1.05e+13 88.7% | | 5.58e+12 89.3% | |
| loop 2062: amrex::Amr::timeStep(int, double, int, int, double) | 1.24e+13 88.6% | | 1.05e+13 88.7% | | 5.58e+12 89.3% | |
| loop 2015: PeleC::advance(double, double, int, int) | 1.24e+13 88.5% | | 1.05e+13 88.6% | | 5.57e+12 89.2% | |
| loop 36: PeleC::do_sdc_advance(double, double, int, int) | 1.24e+13 88.5% | | 1.05e+13 88.6% | | 5.57e+12 89.2% | |
| loop at Advance.cpp: 302 | 1.24e+13 88.4% | | 1.05e+13 88.5% | | 5.57e+12 89.1% | |
| loop 308: PeleC::do_sdc_iteration(double, double, int, int, int, int) | 1.24e+13 88.4% | | 1.05e+13 88.5% | | 5.57e+12 89.1% | |
| loop 561: PeleC::react_state(double, double, bool, amrex::MultiFab*) | 9.61e+12 68.5% | | 8.29e+12 70.0% | | 4.17e+12 66.8% | |
| loop at React.cpp: 109 | 9.43e+12 67.2% | | 8.14e+12 68.7% | | 4.06e+12 65.0% | |
| loop 210: react(amrex::Box const&, amrex::Array4<double> const&, amrex::Array4<double> cons... | 9.39e+12 66.9% | | 8.10e+12 68.4% | | 4.03e+12 64.5% | |
| loop 234: arkEvolve [libsundials_arkode.so.4.7.0] | 9.28e+12 66.2% | | 8.00e+12 67.6% | | 3.94e+12 63.1% | |
| loop erkStep_TakeStep [libsundials_arkode.so.4.7.0] | 7.16e+12 51.1% | | 6.19e+12 52.3% | | 3.05e+12 48.9% | |
| loop cf_RHS(double, _generic_N_Vector*, _generic_N_Vector*, void*) | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop 443: [=] amrex::launch_global<_nv_d_wrapper_t<_nv_d_tag<int (*)>(double, _generic_N... | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop 12: [=] _wrapper_device_stub_launch_global<_nv_d_wrapper_t<_nv_d_tag<int (*)>(do... | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop 26: [=] _device_stub_ZN5amrex13launch_globalI76cf_RHSdP17_generic_N_VectorS2... | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop 24: [=] cudaLaunchKernel<char> | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop 211: cudaLaunchKernel [PeleC3d.gnu.TPROF.CUDA.ex] | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop <gpu kernel> | 6.27e+12 44.7% | | 5.49e+12 46.3% | | 2.48e+12 39.7% | |
| loop amrex::launch_global<cf_RHS(double, _generic_N_Vector*, _generic_N_Vector*, v... | 6.27e+12 44.7% | 1.75e+10 0.1% | 5.49e+12 46.3% | 1.70e+10 0.1% | 2.48e+12 39.7% | |
| loop 12: [=] cf_RHS(double, _generic_N_Vector*, _generic_N_Vector*, void*):(lambda... | 6.25e+12 44.6% | 1.17e+12 8.3% | 5.47e+12 46.2% | 1.16e+12 9.8% | 2.48e+12 39.7% | 1.14e+12 18.2% |
| loop at reactor.cpp: 446 | 5.14e+12 36.6% | 5.35e+10 0.4% | 4.36e+12 36.8% | 4.62e+10 0.4% | 1.38e+12 22.0% | 3.29e+10 0.5% |
| reactor.cpp: 446 | 1.11e+12 7.9% | 1.11e+12 7.9% | 1.11e+12 9.4% | 1.11e+12 9.4% | 1.10e+12 17.7% | 1.10e+12 17.7% |
| AMReX_GpuLaunchGlobal.H: 12 | 1.75e+10 0.1% | 1.75e+10 0.1% | 1.70e+10 0.1% | 1.70e+10 0.1% | | |

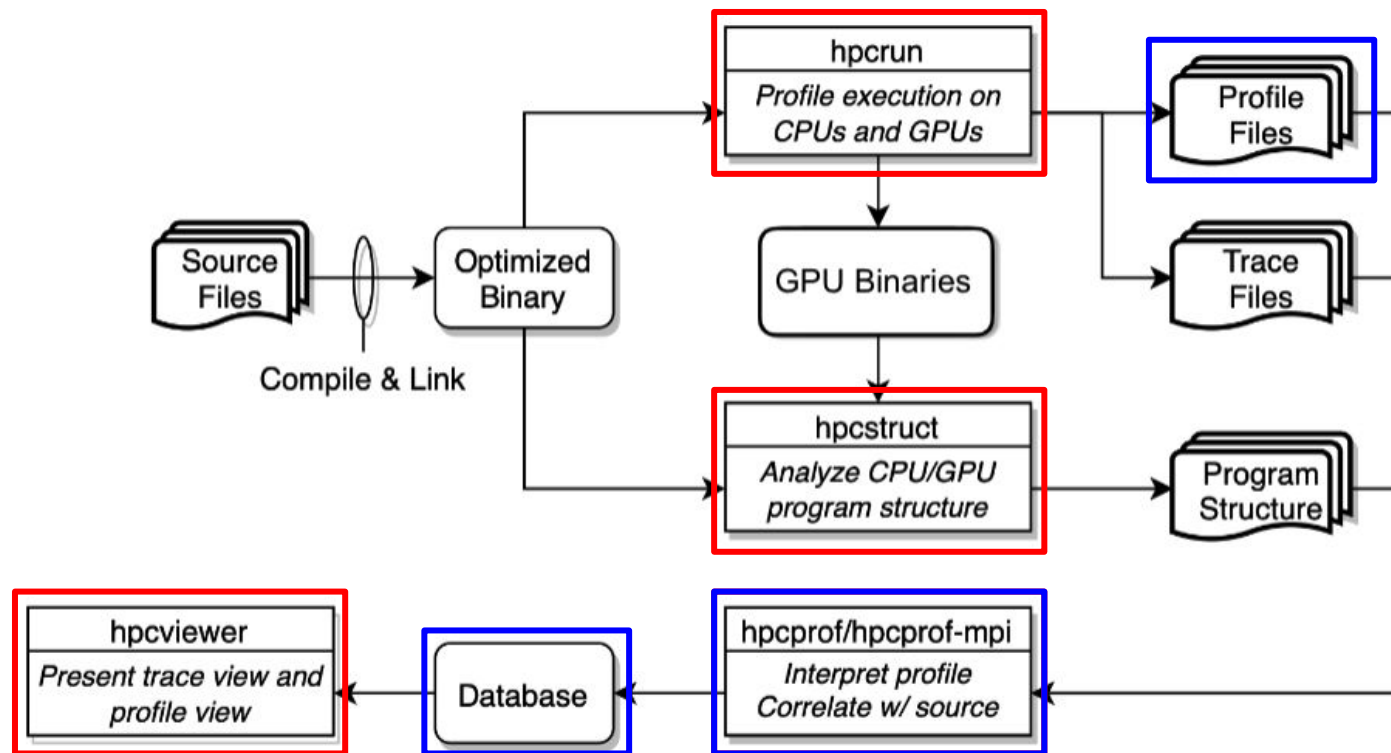
CPU context

GPU context

9.4% GPU stalls outside the loop

mostly memory stalls

HPCToolkit Our Work



Performance Data is Sparse

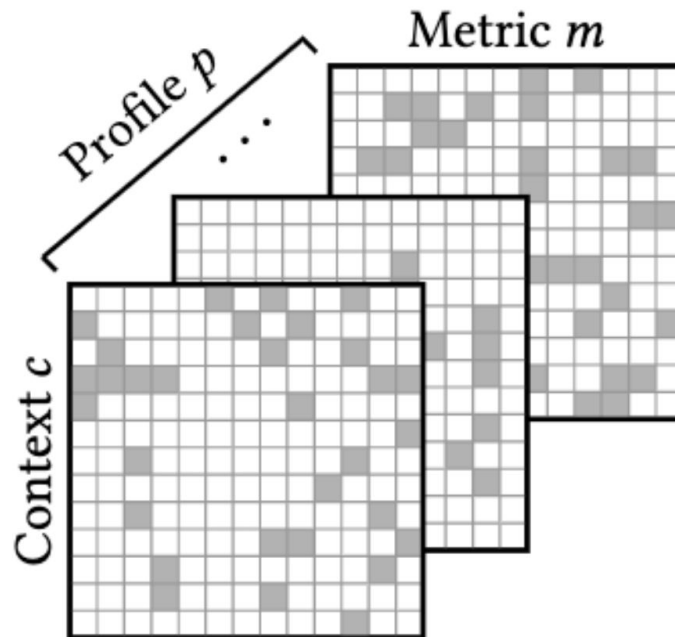
- Threads have different roles, zeros for contexts only in other threads
 - MPI helper threads, OpenMP worker threads, GPU streams...
- Metric costs accumulate in expensive leaf functions
 - Distant callers often have no exclusive metrics
- Many metrics apply only to certain instructions or code regions
 - TLB miss only on memory access; GPU cycles only in GPU code
- Some metrics only apply to very specific contexts
 - Kernel launch parameters only for GPU kernels, data motion volume for GPU copy
- All these factors contribute to sparsity!

Performance data

For each value we record, we need to know:

- which metric
- which calling context
- which thread

≈ 3-dimensional tensor



Sparsity in Practice

| Application | Density (%) | | |
|-------------------------------|---------------------|---------|-------|
| | Contexts | Metrics | |
| <i>Measurement data (In)</i> | | | |
| CPU | AMG 8K | 69.1 | 100.0 |
| | AMG 8K [†] | 22.7 | 20.7 |
| GPU | LAMMPS 1K | 17.7 | 1.8 |
| | PeleC 1K | 15.8 | 2.0 |
| <i>Analysis results (Out)</i> | | | |
| CPU | AMG 8K | 0.301 | 0.182 |
| | AMG 8K [†] | 0.059 | 0.017 |
| GPU | LAMMPS 1K | 2.360 | 1.390 |
| | PeleC 1K | 0.599 | 0.635 |

Measurement data

With 7 CPU metrics:

- 23% of the contexts have metric values
- 21% of the values for contexts are non-zero

With GPU metrics:

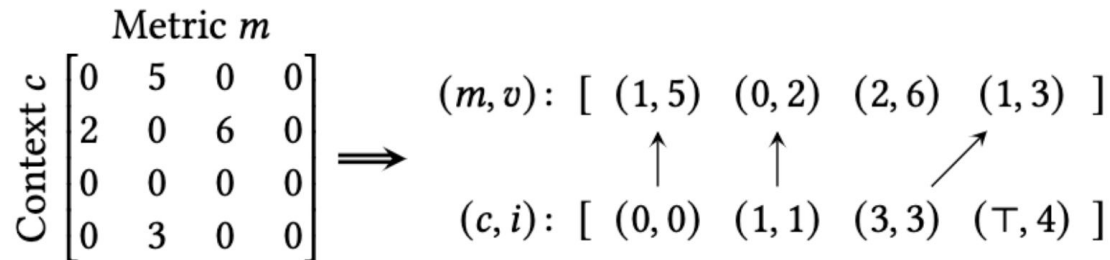
- 16% of the contexts have metric values
- 2% of the values for contexts are non-zero

Analysis results

In our experiments, at most:

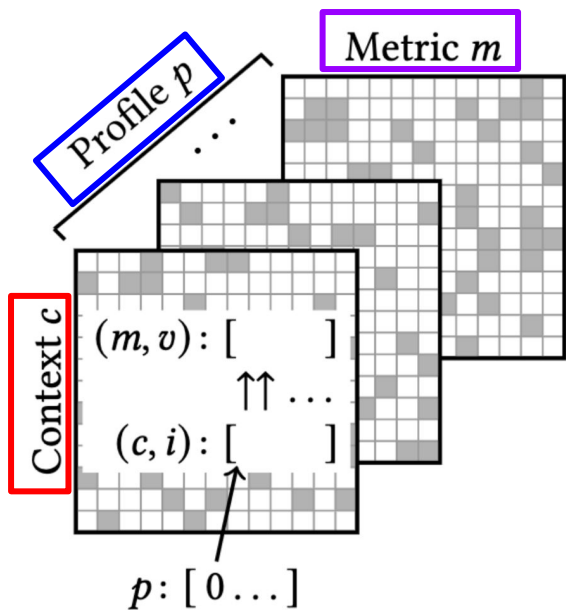
- 2.3% of the contexts have metric values
- 1.4% of the values for contexts are non-zero

Exploiting Sparsity in Measurement Data

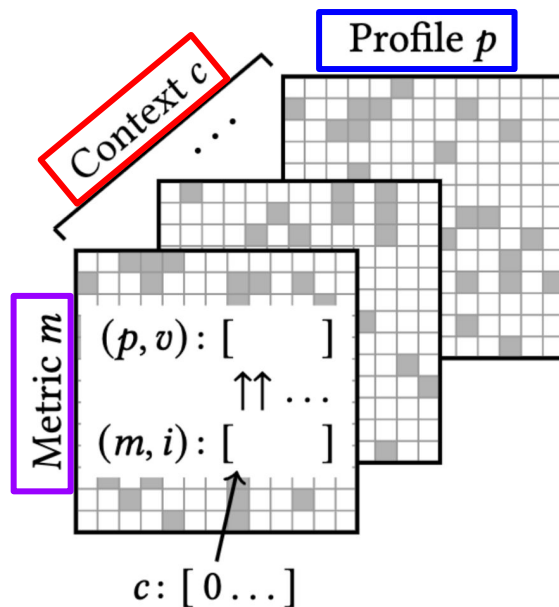


- Each value v measured for a metric m in a calling context c
 - (m, v) vector lists non-zero metric values
 - (c, i) vector maps contexts to contiguous ranges of the (m, v) vector
- Exploit sparsity in both metric and context dimensions to reduce space
- Ensure logarithmic access time to values

Exploiting Sparsity in Analysis Results



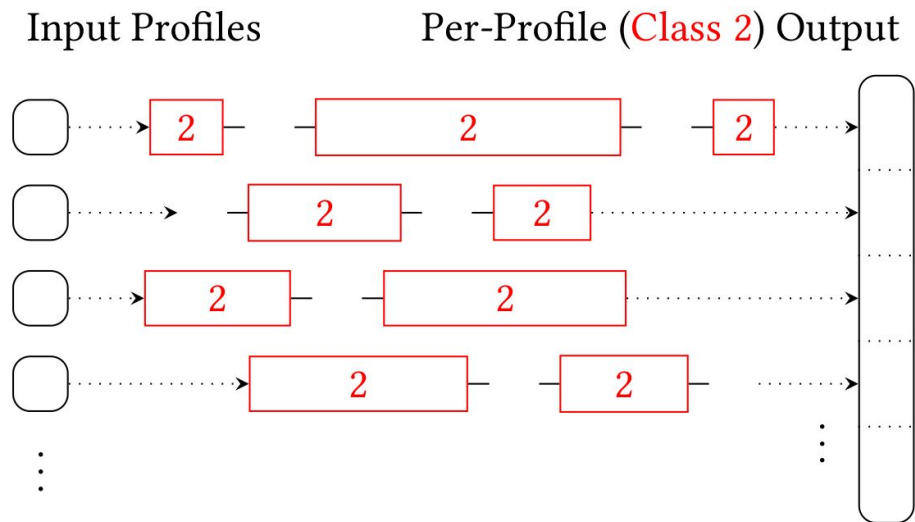
Profile-Major-Sparse (PMS):
data for each profile is contiguous



Context-Major-Sparse (CMS):
data for each context is contiguous

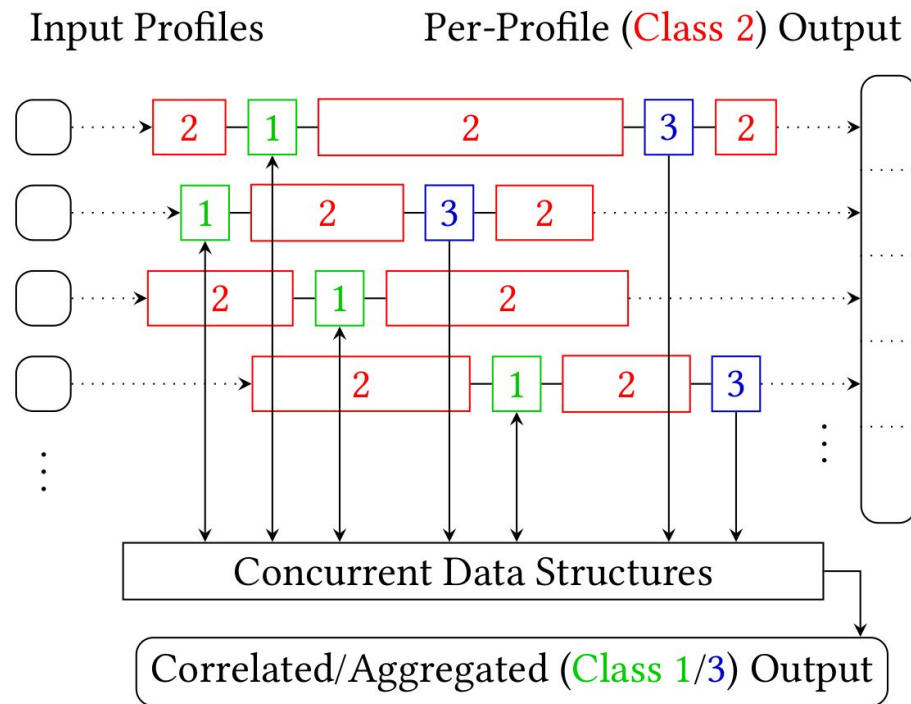
Streaming Aggregation: 3 Classes of Operations

- Data dependency based on number
 - In order: class 1, then 2, then 3
- **Class 2** operations run in parallel
 - Independent, per-profile analysis
 - e.g. calculating inclusive cost (per-thread)
 - Read inputs and write outputs directly
 - No significant synchronization
- “Streaming” approach reduces memory
 - Only in-flight profiles in memory!



Streaming Aggregation: 3 Classes of Operations

- **Class 1** operations as needed by **Class 2**
 - Associative, commutative, idempotent
 - e.g. correlate calling contexts across threads
- **Class 3** ops aggregate **Class 2** results
 - Associative and commutative
 - e.g. summary statistics for entire execution
- Update concurrent data structures
- Streaming parallel approach to “aggregation”
 - No defined order, no phases or barriers
 - Fine-grained, data-centric synchronization
 - Outperforms classic reduction approach!



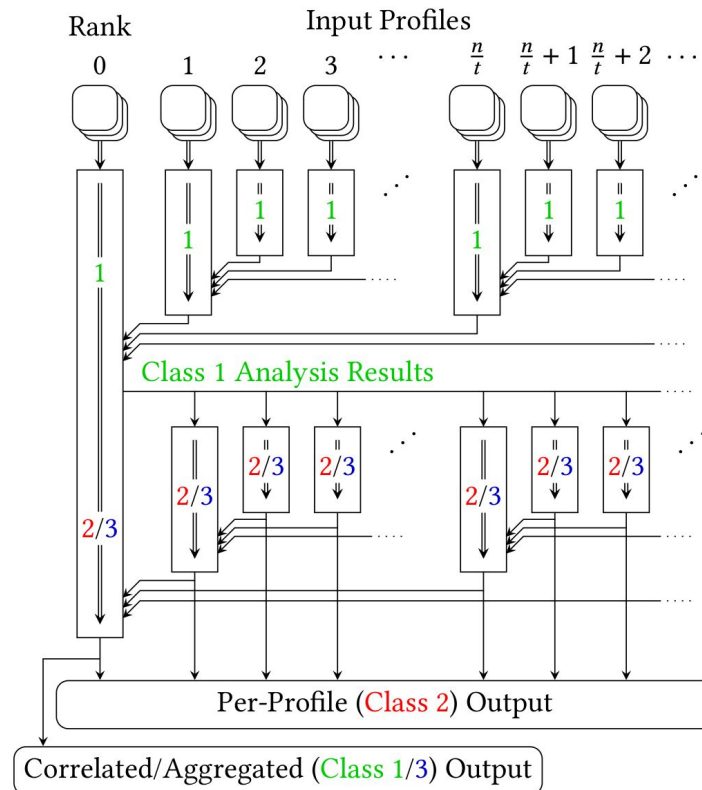
Lessons from a Proxy Application



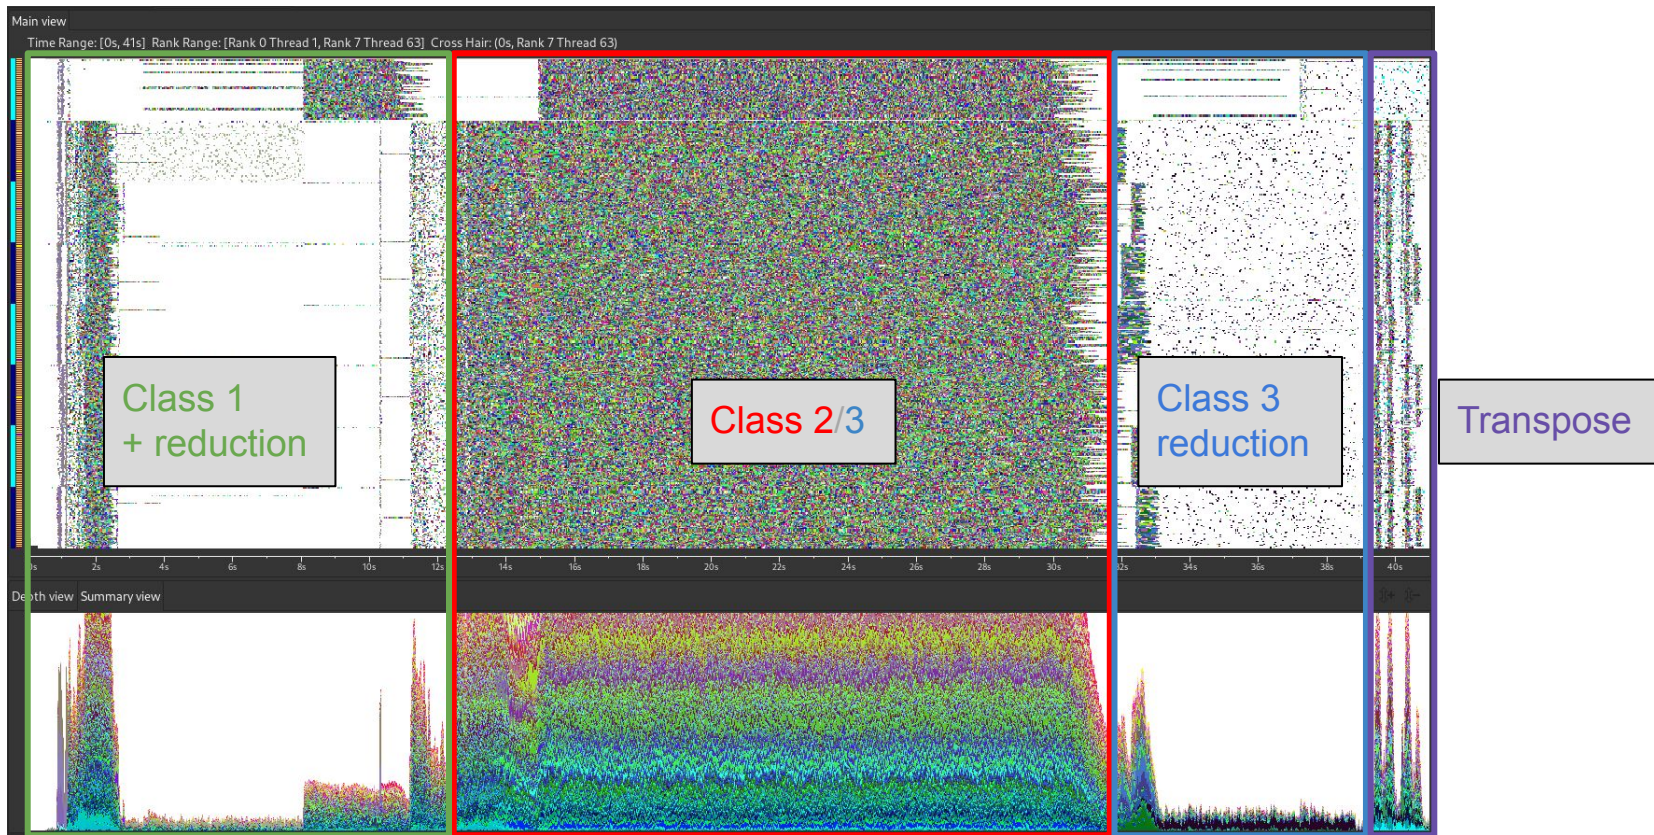
- Reduction phase causes large serialization: **phases are costly**
- Lots of allocations also takes time: **high footprint costly**
- Choice data structure affects performance: **choose/tune wisely**

Extension to Multiple Nodes

- Hybrid approach
 - MPI across nodes
 - Streaming Aggregation within
- **Class 1/3** ops use reductions across nodes
 - Shared-memory parallelism within a node
- **Class 2** ops distributed across nodes
 - Independent, no communication needed
- All nodes perform I/O individually
 - Exploit distributed file system



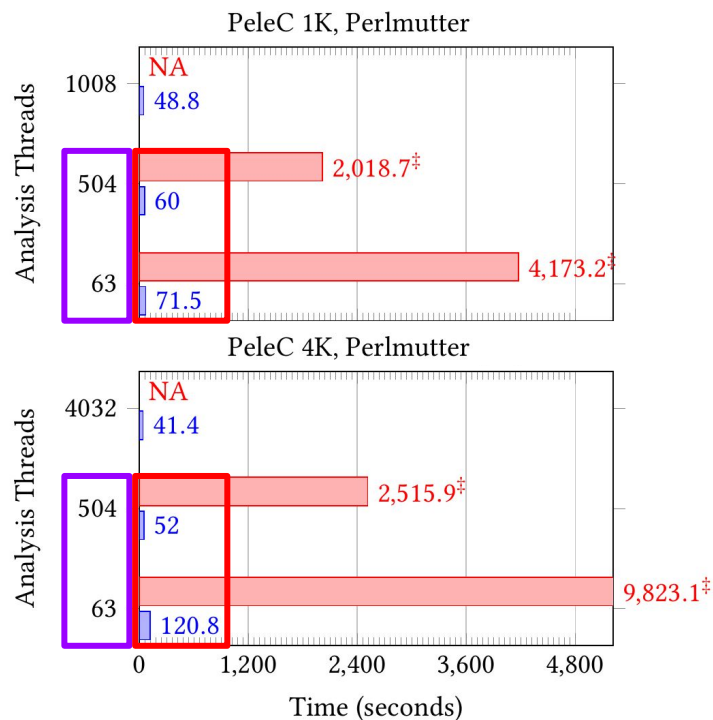
Dense Parallelism from Streaming Aggregation



Performance Improvements Analyzing GPU Measurements

- Post-processing performance from PeleC
 - 512+512 or 2K+2K (threads + GPUs)
 - Analysis uses 63 threads/node
- **30-80x performance improvement**
- **Smaller memory footprint is crucial**
 - Original `hpcprof` exceeded 256GB/node
 - Required 16x total compute resources
- **Minimal resources vs. application**
 - Ran PeleC on 128 and 512 nodes
 - Up to 1/3 of Perlmutter
 - Analyzed measurements on 1 and 8 nodes
- **Results in 1-2 mins**

See ICS paper for CPU-only results for AMG



‡Increased node count by 16x and used 4 ranks/node to avoid out-of-memory errors.

Storage Improvements

- Performance data from real-world apps
 - Measurements (In) and results (Out)
- **Significant space reductions**
 - Up to 10x compression in measurements
 - Up to 1254x compression in results
- **Results in GBs... not TBs!**

| Tool | Size (GiB) | | | |
|---------------|-------------|-------------|-------------|--------------|
| | AMG 65K | | AMG 262K | |
| | In | Out | In | Out |
| HPCToolkit | 5.88 | 195 | 36.9 | 1250 |
| HPCToolkit-SA | 6.14 | 5.55 | 38.2 | 36.6 |
| <hr/> | | | | |
| Tool | PeleC 1K | | PeleC 4K | |
| | In | Out | In | Out |
| HPCToolkit | 21.2 | 4800 | 50.7 | 14300 |
| HPCToolkit-SA | 2.85 | 5.27 | 4.8 | 11.4 |

Conclusions

- “Less is more”: Efficient use of compute resources is key
 - Exploit sparsity to reduce storage and I/O
 - Exploit multithreading to efficiently use each node’s cores and threads
 - Exploit distributed memory parallelism for scale
- Novel sparse formats for performance measurements and analysis results
 - Up to 1254x compression in analysis results for our experiments
- Novel highly-parallel streaming aggregation approach to performance analysis
 - Results for large-scale executions in minutes!
- **End result: better prepared for performance analysis at exascale!**
- Improvements actively being integrated into HPCToolkit
 - Will become widely available in a future release

Ongoing Work: Accelerating I/O

Problems

- Metadata server is slow:
 - Our approach of recording 2 files per thread may be costly
- Unnecessary overhead for maintaining page cache consistency between multiple processes on different nodes
 - Our writes don't overlap

Approach

- When available, exploit object-based solid state storage for ultimate performance

Ongoing work: **D**istributed **A**synchronous **O**bject **S**torage (DAOS)

- Designed for massively distributed NVM
- Affordable, fast, large-capacity PM
 - SCM (store metadata, latency-sensitive small data)
 - NVMe SSDs (bulk data)

Properties

- High throughput and IOPS at arbitrary alignment and size
- Fine-grained I/O operations with true zero-copy I/O to SCM
- Non-blocking data and metadata operations to allow I/O and computation to overlap
- Scalable distributed transactions with guaranteed data consistency and automated recovery

An I/O abstraction layer for HPCToolkit

- Freedom to choose any I/O option (DAOS, POSIX, Lustre ...) for any file
 - Example:
 - `hpcprof -o daos://<POOL>/<CONT>/database-dir measurement-dir`
 - We use DAOS for database and use POSIX for measurement files at the same time
- Easy integration of other I/O options in the future
 - Example: HPE Rabbit Near Node Storage (Livermore's El Capitan), Lustre ...
- Details of I/O are invisible to other parts of the software
 - Example: just call `io->write(...)` without concern for underlying I/O implementations

HPCToolkit Funding Acknowledgments

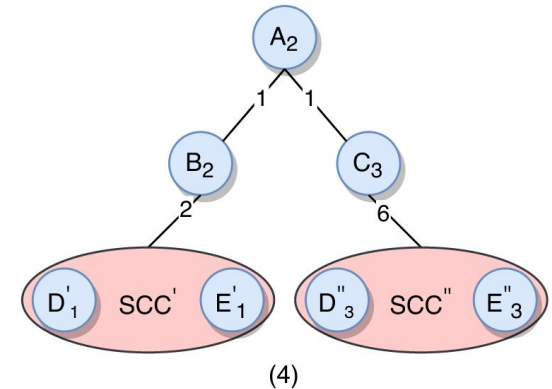
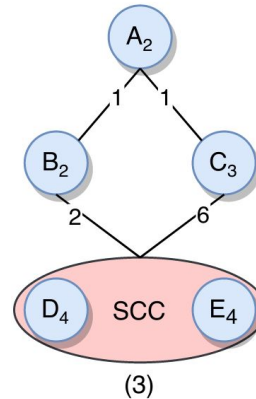
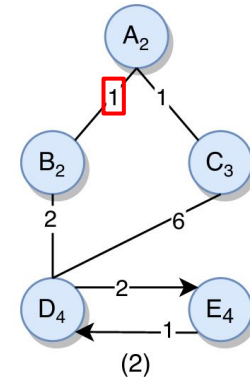
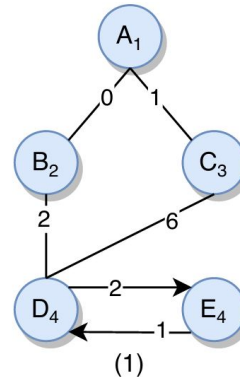
- Exascale Computing Project 17-SC-20-SC
- Lawrence Livermore National Laboratory Subcontract B645220
- Argonne National Laboratory Subcontract 9F-60073
- Advanced Micro Devices
- Intel Corporation
- TotalEnergies EP Research & Technology USA, LLC.

Backup Slides



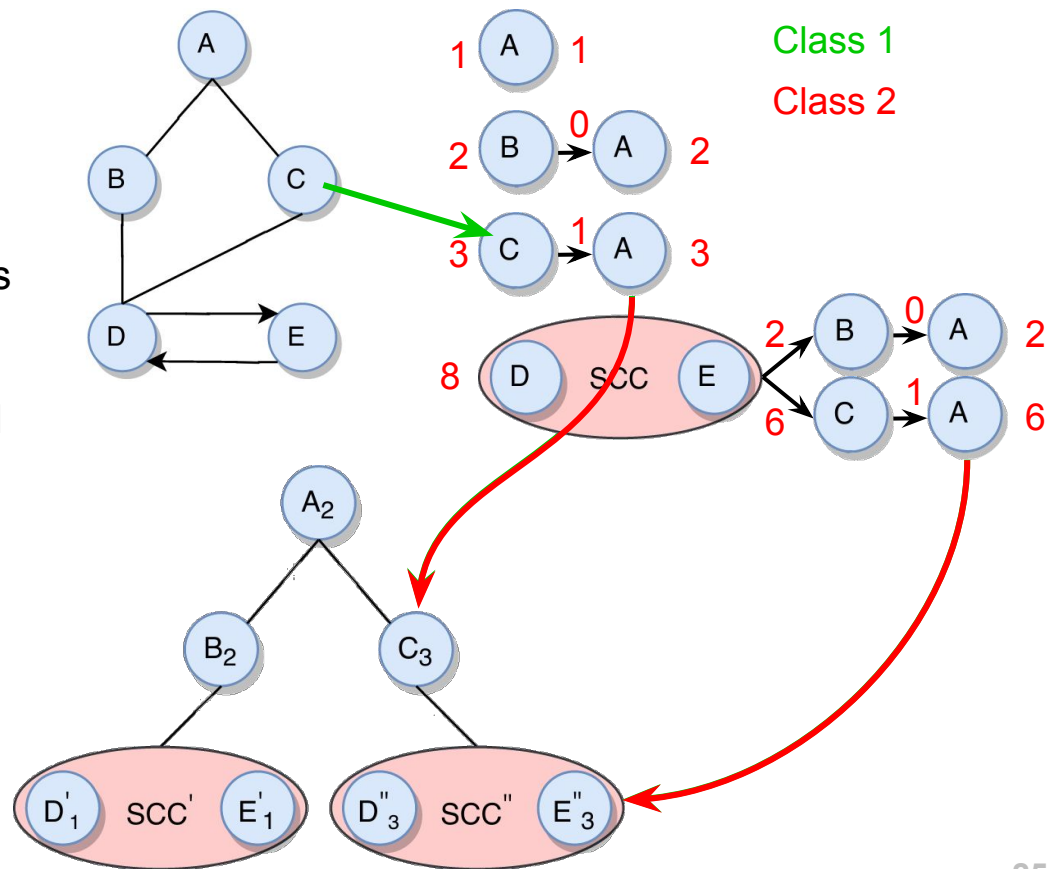
Adapting Algorithms to New Constraints

- E.g. GPU calling context reconstruction
 - GPU PC sampling is always flat
 - Reconstruct calling context based on static CFG from binary analysis
- 4-step algorithm:
 - Attribute flat samples to call graph
 - Fixup missing edge weights
 - Convert to DAG by SCC
 - Expand DAG into calling context tree
- Problem: does not obey S.A. constraints!
 - (4) must be **class 1** for final CCT...
 - ...but (1) must be **class 2** for values!



Adapting Algorithms to New Constraints

- Generate full possible CCT (**Class 1**)
 - Iterate reverse static call paths
 - Create full context for each path
- Reattribute metric values (**Class 2**)
 - Attribute flat samples to heads
 - Distribute among branching paths
 - Assign path values to CCT
- Reconstruction mixed with normal ops
 - **Generate** when a context in need of reconstruction is parsed
 - **Reattribute** just before inclusive metric propagation



DAOS Integration Experiences - I

No matching DAOS function for every POSIX function

- `fread`, `fwrite`, `fseek`...
 - `fread`: We reworked our code to avoid it by using `read_at`
 - `fwrite`: DAOS only has `write_at`, manage our own buffer and cursor
 - `fseek`: We reworked our code to avoid it by using `read_at` and `write_at`
- `write`
 - Multiple processes append to the same log file (debug info for developers)
 - With `write_at`, we need to share a buffer or a cursor between processes
 - We reworked so that each process creates its own log file
- C++ filesystem abstractions (`directory_iterator`, `exists`, `remove_all` ...)
 - We needed to reimplement these abstractions using our I/O abstraction layer

DAOS Integration Experiences - II

- Extra initialization and finalization steps
 - Each process initializes its own DAOS pool and container handles
 - We set up the output directory and begin recording trace data BEFORE MPI initialization
- Need to manage file system accesses carefully
 - Before: use POSIX anywhere without any direct coordination
 - Now: pass around I/O abstractions to access the correct file system
- Caching semantics
 - Objects created were only visible through command line after cache refreshed
 - Using `--disable-caching` can help, but may hurt performance