Overview of TAU

ALCF Handson HPC Workshop
Sameer Shende
University of Oregon and ParaTools, Inc.
1:30pm – 1:45pm CT, Wednesday, Oct. 11, 2023
2:30pm – 5pm Room 1406. Breakout session

http://tau.uoregon.edu/TAU_ALCF23.pdf

TAU Quickstart Guide on Polaris at ALCF

Setup:

- `% module load tau`

Profiling with an un-instrumented application:

- MPI: `% aprun -n 64 tau_exec -ebs ./a.out`
- CUDA+Sampling: `% aprun -n 64 tau_exec -T cupti -cupti -ebs ./a.out`
- Pthread: `% aprun -n 64 tau_exec -T mpi,pthread -ebs ./a.out`

Analysis: `% pprof -a -m | more;  % paraprof (GUI)`

Tracing:

- Vampir: MPI:
  ```
  % export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2
  % aprun -n 64 tau_exec ./a.out; vampir traces.otf2 &
  ```
- Chrome:
  ```
  % export TAU_TRACE=1; aprun -n 64 tau_exec ./a.out; tau_treemerge.pl;
  % tau_trace2json tau.trc tau.edf -chrome -ignoreatomic -o app.json
  ```
  Chrome browser: chrome://tracing (Load -> app.json) or Perfetto.dev
- Jumpshot:
  ```
  % export TAU_TRACE=1; aprun -n 64 tau_exec ./a.out; tau_treemerge.pl;
  % tau2slog2 tau.trc tau.edf -o app.slog2; jumpshot app.slog2 &
  ```

Setup: Installing TAU on Laptops

- Prerequisites: Java in your path
- Microsoft Windows
  - Install Java from Oracle.com
    - http://tau.uoregon.edu/tau.exe
    - Install, click on a ppk file to launch paraprof
- macOS (x86_64)
  - Install Java 11.0.3:
    - Download and install http://tau.uoregon.edu/java.dmg
    - If you have multiple Java installations, add to your ~/.zshrc (or ~/.bashrc as appropriate):
      - export PATH=/Library/Java/JavaVirtualMachines/jdk-11.0.3.jdk/Contents/Home/bin:$PATH
  - Download and install TAU (copy to /Applications from dmg):
    - http://tau.uoregon.edu/tau.dmg
    - export PATH=/Applications/TAU/tau/apple/bin:$PATH
    - paraprof app.ppk &
- macOS (arm64, Apple Silicon M1/M2)
    - http://tau.uoregon.edu/java_arm64.dmg
    - http://tau.uoregon.edu/tau_arm64.dmg
- Linux (http://tau.uoregon.edu/tau.tgz)
  - ./configure; make install; export PATH=<taudir>/x86_64/bin:$PATH; paraprof app.ppk &

TAU Performance System®

- Parallel performance framework and toolkit
  - —Supports all HPC platforms, compilers, runtime system
  - —Provides portable instrumentation, measurement, analysis

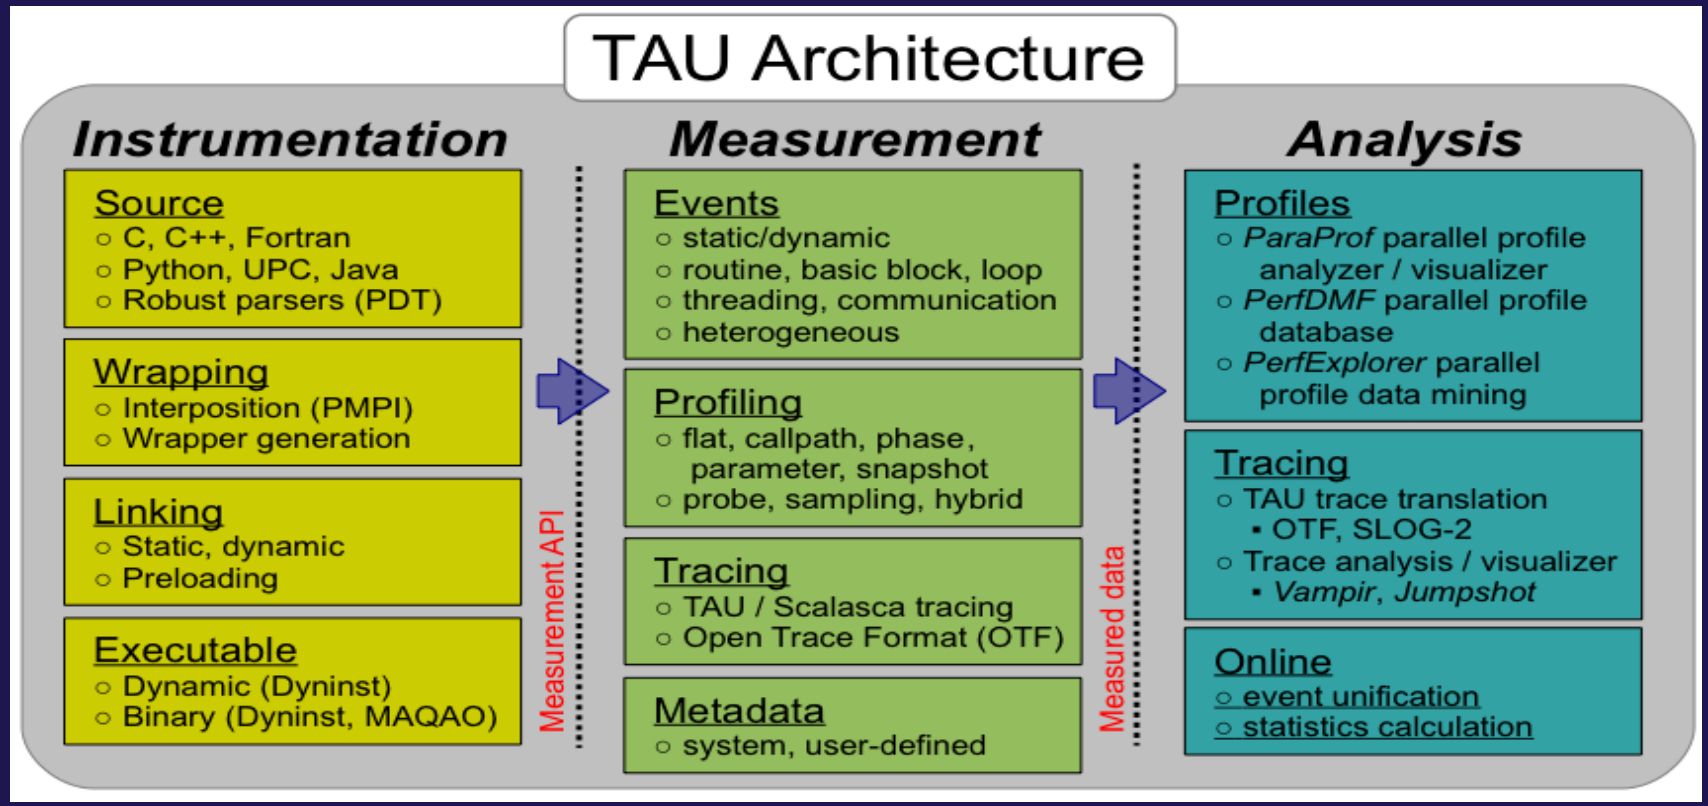

TAU Performance System®

- Instrumentation
  - Fortran, C++, C, UPC, Java, Python, Chapel, Spark
  - Automatic instrumentation
  - Map manual instrumentation APIs from other tools to TAU
    - NVTX, ROCTx
    - CAMTimers, PerfStubs, PETSc, Caliper, Kokkos API
- Measurement and analysis support
  - MPI, OpenSHMEM, ARMCI, PGAS, DMAPP
  - pthreads, OpenMP, OMPT interface, hybrid, other thread models
  - GPU: Intel oneAPI DPC++/SYCL, AMD ROCm (RocProfiler and RocTracer), CUDA, OpenCL, OpenACC, Kokkos
  - Parallel profiling and tracing
- Analysis
  - Parallel profile analysis (ParaProf), data mining (PerfExplorer)
  - Performance database technology (TAUdb)
  - 3D profile browser

Application Performance Engineering using TAU

- How much time is spent in each application routine and outer *loops*? Within loops, what is the contribution of each *statement*? What is the time spent in OpenMP loops? In kernels on GPUs. How long did it take to transfer data between host and device (GPU)?
- How many instructions are executed in these code regions?
  Floating point, Level 1 and 2 *data cache misses*, hits, branches taken? What is the extent of vectorization for loops?
- How much time did my application spend waiting at a barrier in MPI collective operations?
- How can I use my app multi-node GPU systems? With unmodified binary on all 3 vendor GPUs?
- What is the memory usage of the code? When and where is memory allocated/de-allocated? Are there any memory leaks? What is the memory footprint of the application? What is the memory high water mark?
- How much energy does the application use in Joules? What is the peak power usage?
- What are the I/O characteristics of the code? What is the peak read and write *bandwidth* of individual calls, total volume?
- How does the application *scale*? What is the efficiency, runtime breakdown of performance across different core counts?

Instrumentation

- Add hooks in the code to perform measurements
- **Source instrumentation using a preprocessor**
  - Add timer start/stop calls in a copy of the source code.
  - Use Program Database Toolkit (PDT) for parsing source code.
  - Requires recompiling the code using TAU shell scripts (tau_cc.sh, tau_f90.sh)
  - Selective instrumentation (filter file) can reduce runtime overhead and narrow instrumentation focus.
- **Compiler-based instrumentation**
  - Use system compiler to add a special flag to insert hooks at routine entry/exit.
  - Requires recompiling using TAU compiler scripts (tau_cc.sh, tau_f90.sh…)
- **Runtime preloading of TAU's Dynamic Shared Object (DSO)**
  - No need to recompile code! Use **aprun tau_exec ./app** with options.

Profiling and Tracing

Profiling

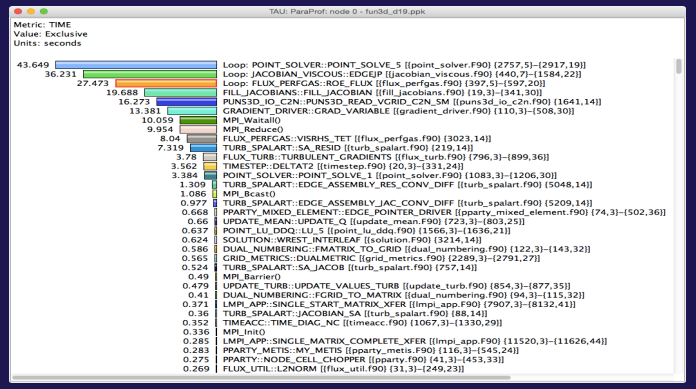

- **Profiling** shows you **how much** (total) time was spent in each routine

Tracing

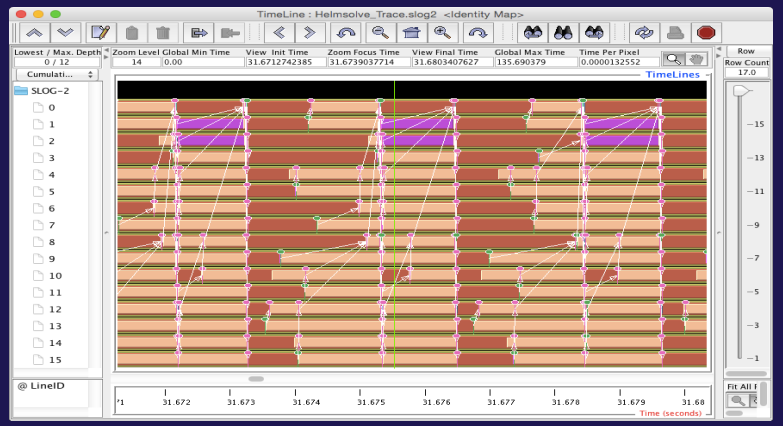

- Tracing shows you when the events take place on a timeline

- Profiling and tracing

  Profiling shows you **how much** (total) time was spent in each routine

  Tracing shows you **when** the events take place on a timeline

Instrumentation

- Direct and indirect performance observation
- Instrumentation invokes performance measurement
- Direct measurement with *probes*
- Indirect measurement with periodic sampling or hardware performance counter overflow interrupts
- Events measure performance data, metadata, context, etc.

- User-defined events
  - ***Interval*** (start/stop) events to measure exclusive & inclusive duration
  - ***Atomic events*** take measurements at a single point
    - Measures total, samples, min/max/mean/std. deviation statistics
  - ***Context events*** are atomic events with executing context
    - Measures above statistics for a given calling path

Inclusive vs. Exclusive values

- Inclusive
  - Information of all sub-elements aggregated into single value
- Exclusive
  - Information cannot be subdivided further

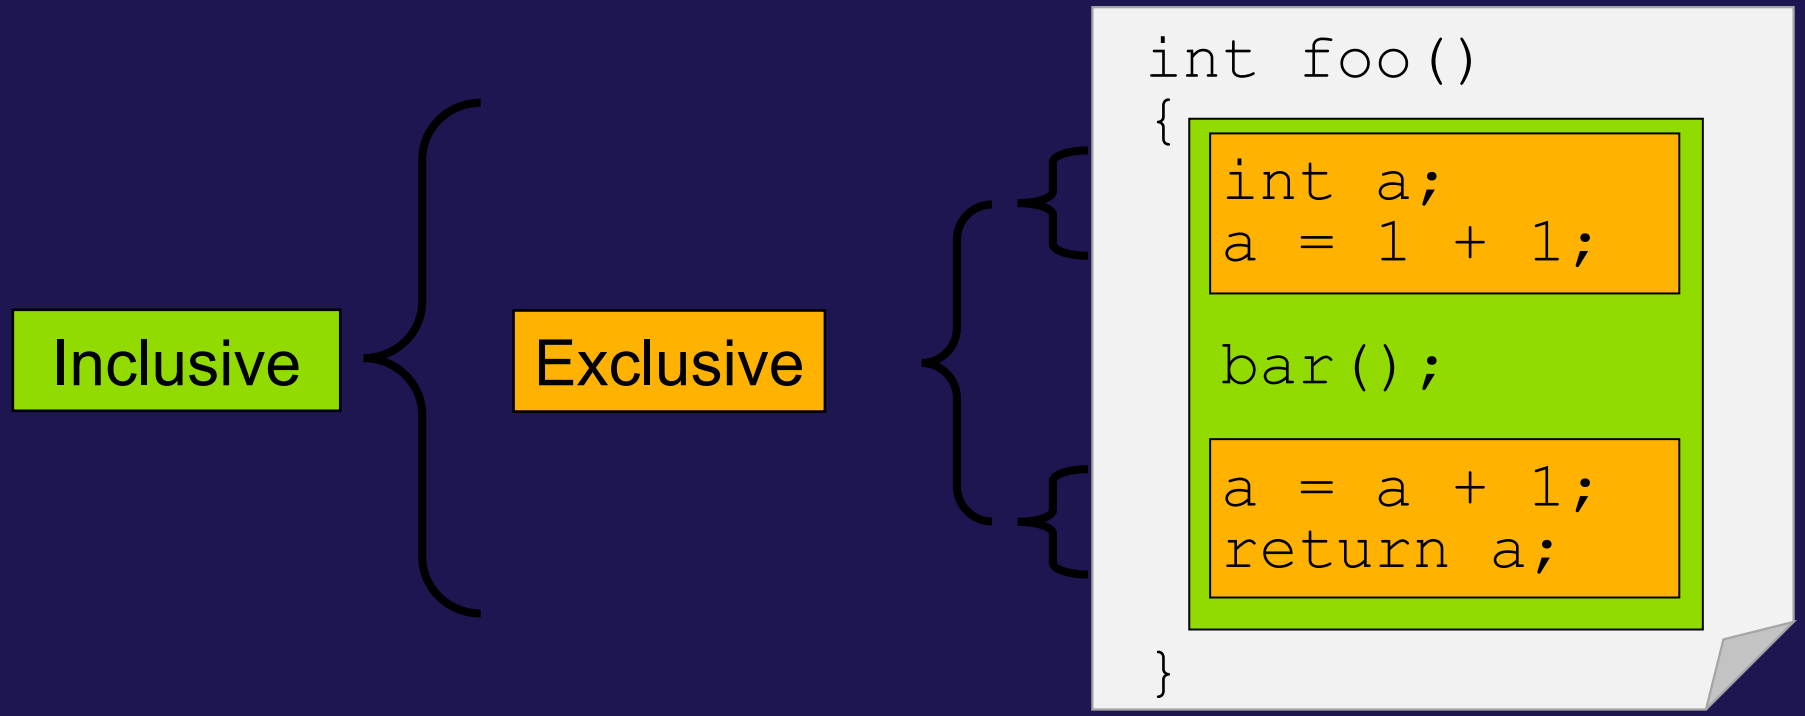

Inclusive Measurements

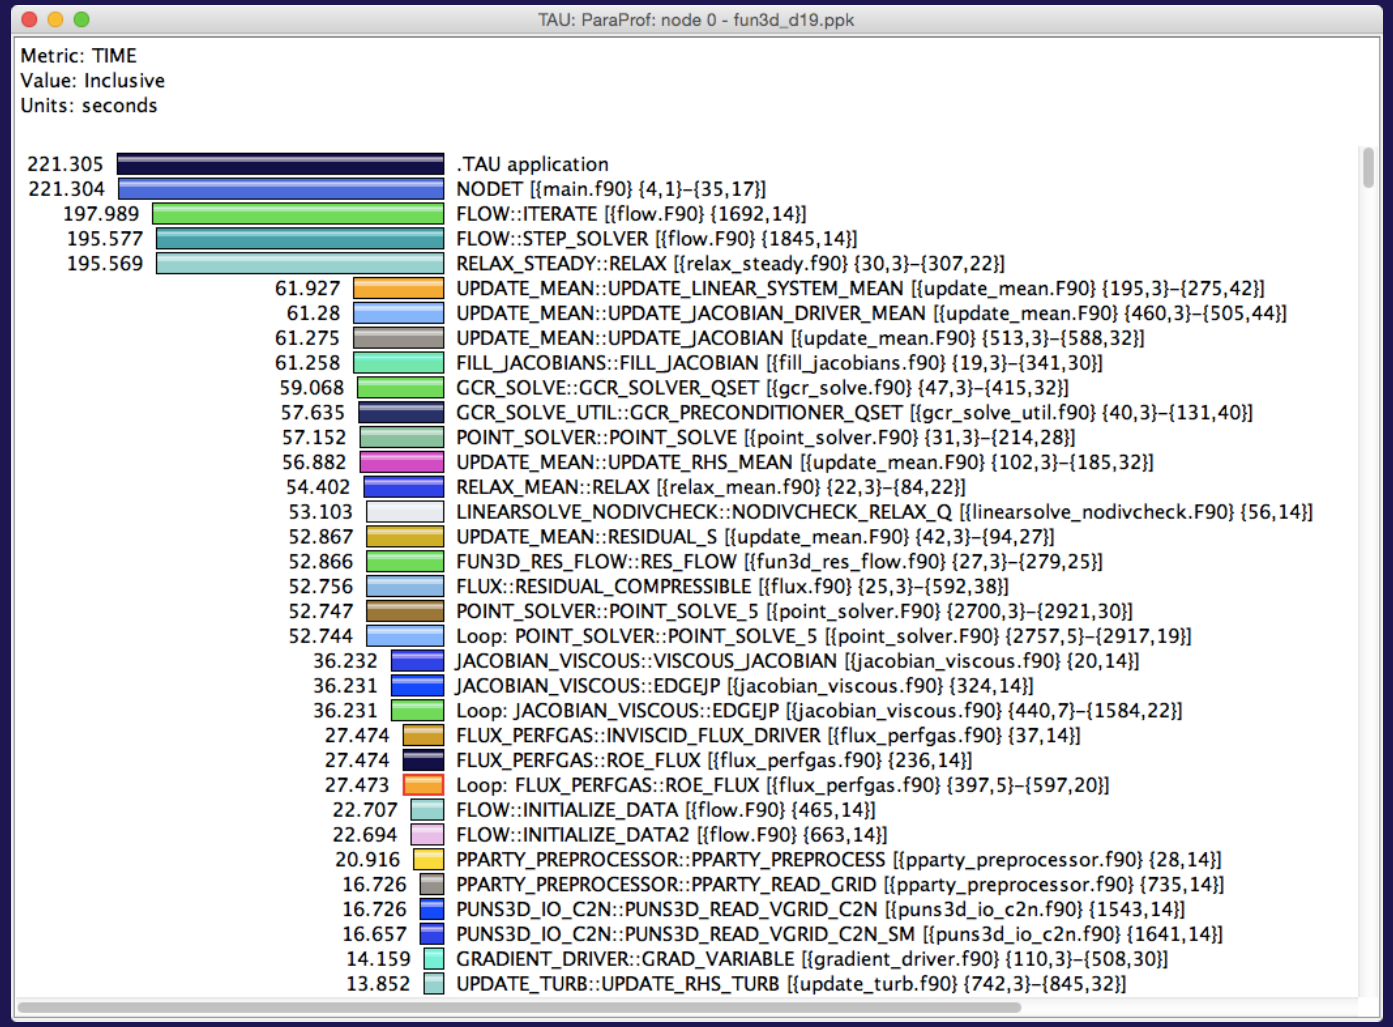

Exclusive Time

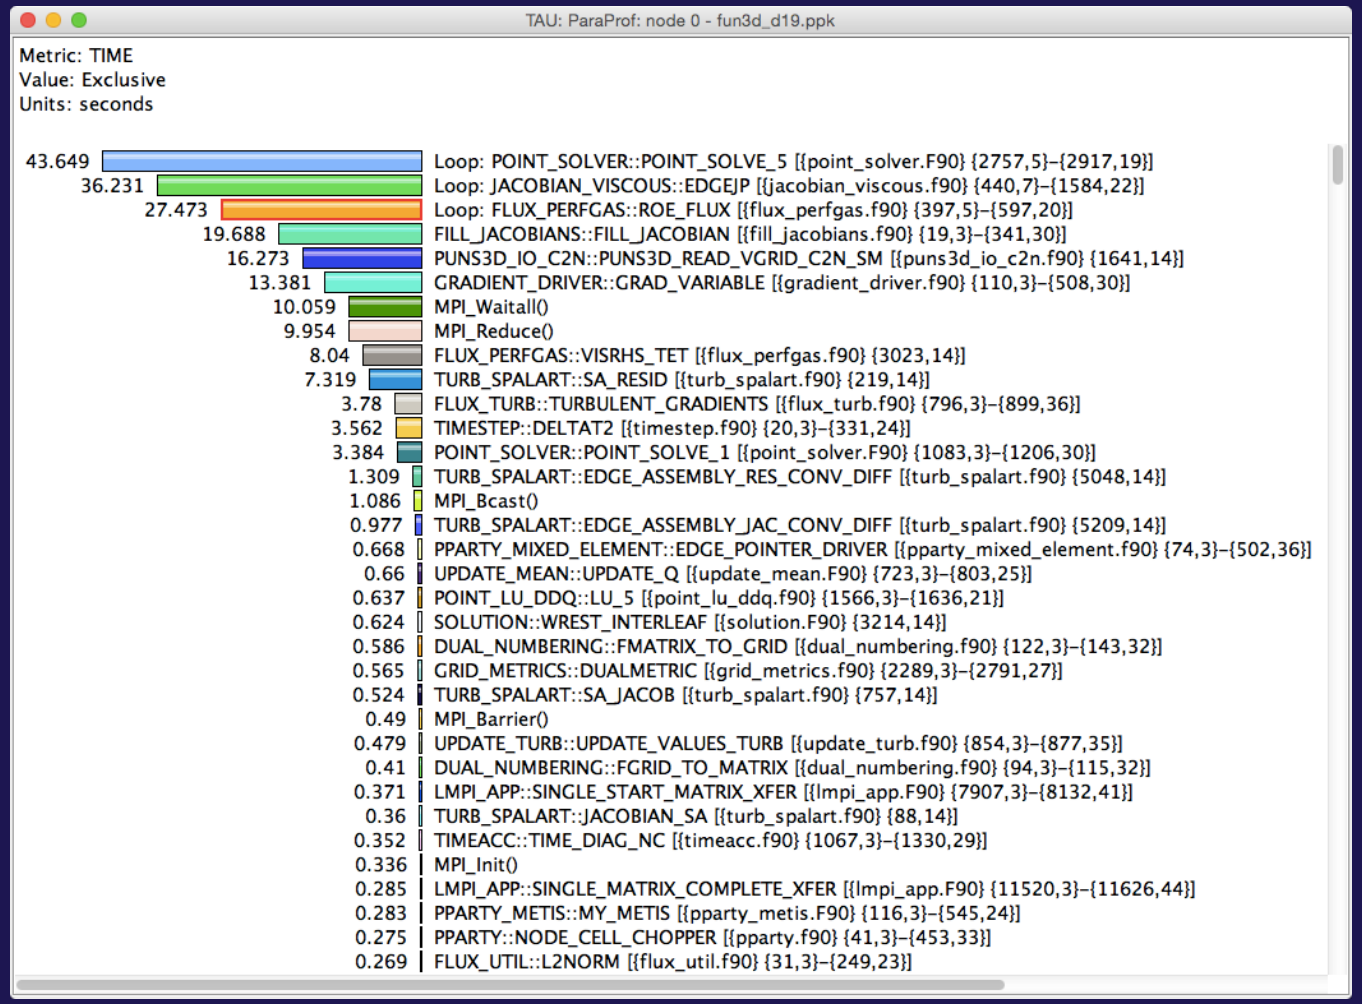

How much data do you want?

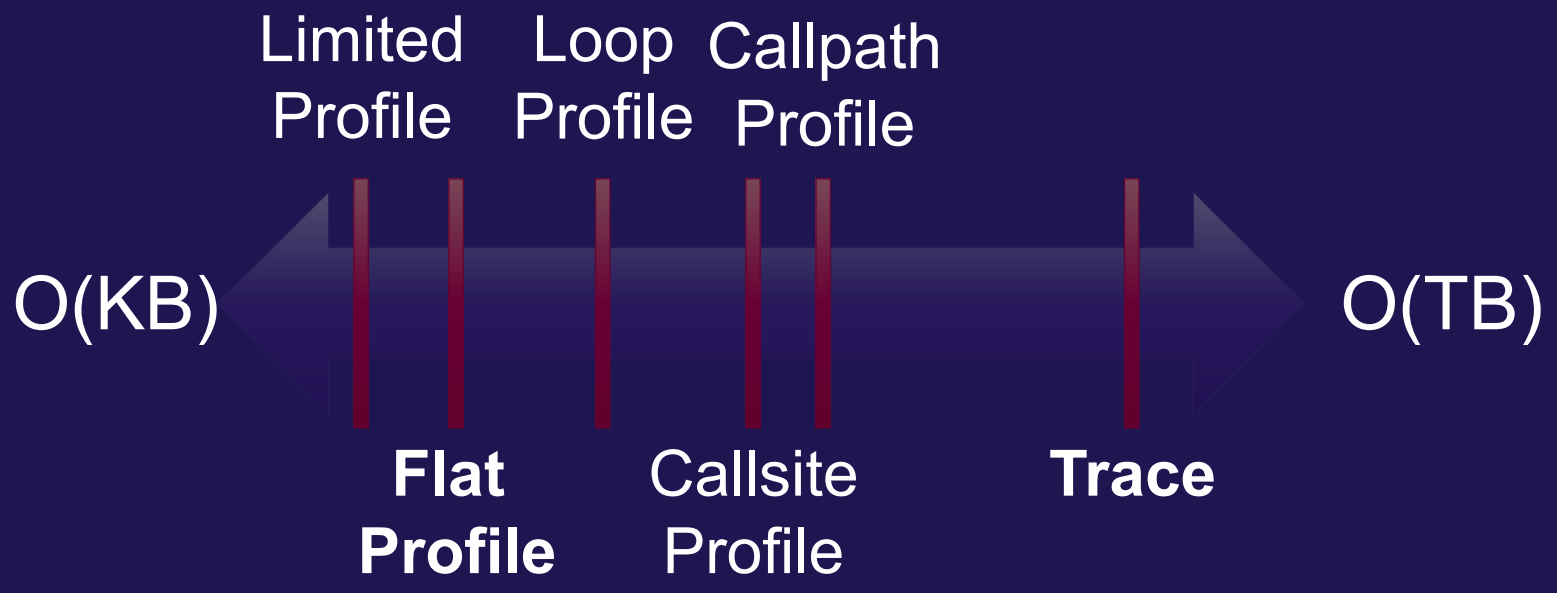

Performance Data Measurement

Direct via Probes

```
Call
START('potential')
// code
Call
STOP('potential')
```

- Exact measurement
- Fine-grain control
- Calls inserted into code

Indirect via Sampling

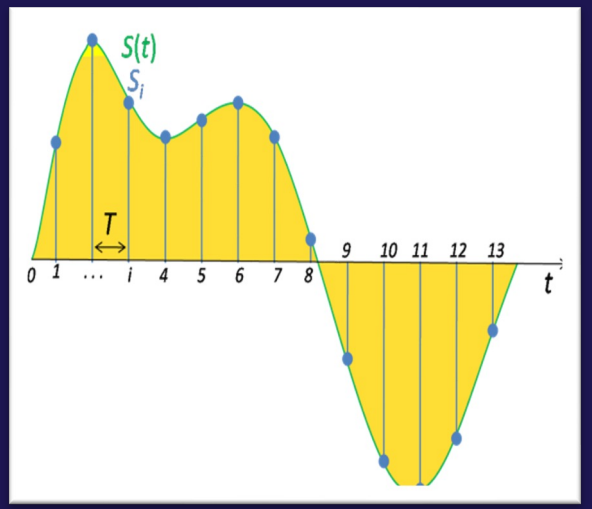

- No code modification
- Minimal effort
- Relies on debug symbols (-g)

Event-Based Sampling (EBS)

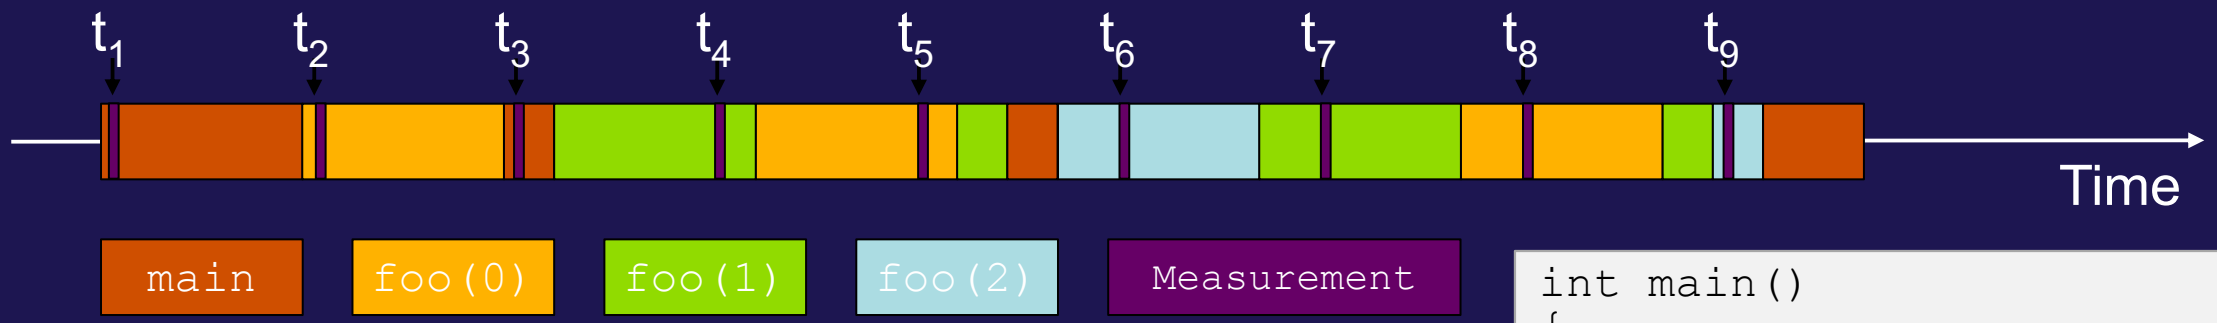

- Running program is periodically interrupted to take measurement
  - Timer interrupt, OS signal, or HWC overflow
  - Service routine examines return-address stack
  - Addresses are mapped to routines using symbol table information
- Statistical inference of program behavior
  - Not very detailed information on highly volatile metrics
  - Requires long-running applications
- Works with unmodified executables (tau_exec –ebs)

```
int main()
{
  int i;

  for (i=0; i < 3; i++)
    foo(i);

  return 0;
}

void foo(int i)
{

  if (i > 0)
    foo(i - 1);

}
```

Instrumentation

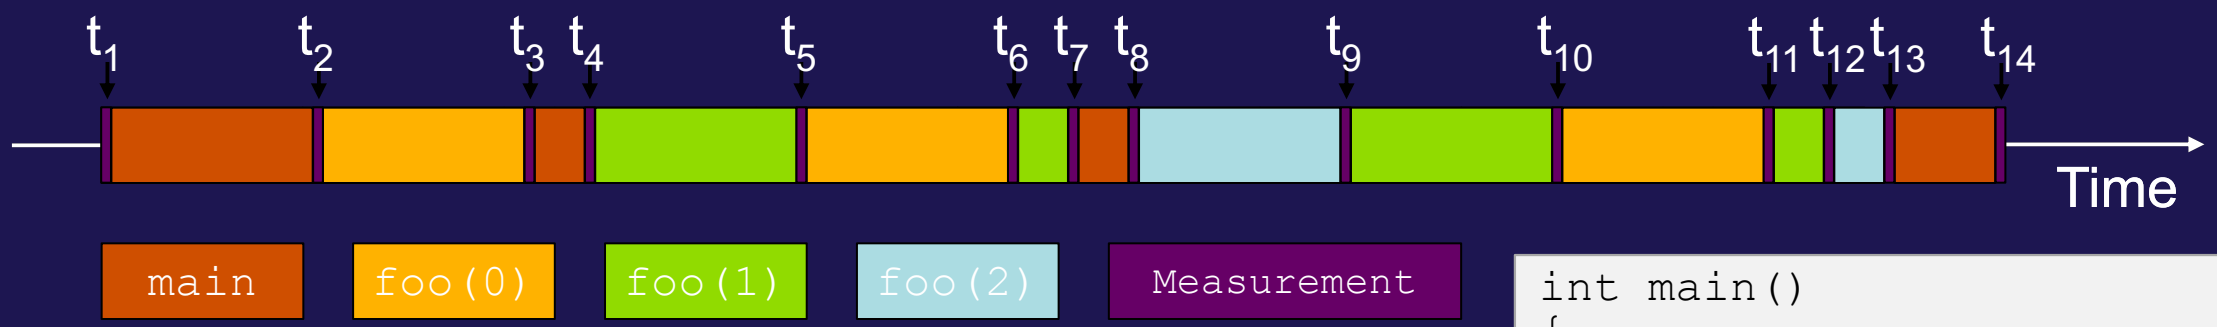

- Measurement code is inserted such that every event of interest is captured directly
  - Can be done in various ways
- Advantage:
  - Much more detailed information
- Disadvantage:
  - Processing of source-code / executable necessary
  - Large relative overheads for small functions

```
int main()
{
  int i;
  TAU_START("main");
  for (i=0; i < 3; i++)
    foo(i);
  TAU_STOP("main");
  return 0;
}

void foo(int i)
{
  TAU_START("foo");
  if (i > 0)
    foo(i - 1);
  TAU_STOP("foo");
}
```

Using TAU's Runtime Preloading Tool: tau_exec

- Preload a wrapper that intercepts the runtime system call and substitutes with another
  - **MPI**
  - **OpenMP**
  - **POSIX I/O**
  - **Memory allocation/deallocation routines**
  - **Wrapper library for an external package**
- No modification to the binary executable!
- Enable other TAU options (communication matrix, OTF2, event-based sampling)

TAU Execution Command (tau_exec)

- Uninstrumented execution
  - % aprun -n 256 ./a.out
- Track GPU operations
  - % aprun -n 256 tau_exec -T rocprofiler –rocm ./a.out
  - % aprun -n 256 tau_exec –l0 ./a.out
  - % aprun -n 256 tau_exec –cupti ./a.out
  - % aprun -n 256 tau_exec –opencl ./a.out
  - % aprun -n 256 tau_exec –openacc ./a.out
- Track MPI performance
  - % aprun -n 256 tau_exec ./a.out
- Track I/O, and MPI performance (MPI enabled by default)
  - % aprun -n 256 tau_exec -io ./a.out
- Track OpenMP and MPI execution (using OMPT)
  - % export TAU_OMPT_SUPPORT_LEVEL=full;
  - % aprun -n 256 tau_exec –T ompt,mpi -ompt ./a.out
- Track memory operations
  - % export TAU_TRACK_MEMORY_LEAKS=1
  - % aprun -n 256 tau_exec –memory_debug ./a.out (bounds check)
- Use event based sampling (compile with –g)
  - % aprun -n 256 tau_exec –ebs ./a.out
  - Also -ebs_source=<PAPI_COUNTER> -ebs_period=<overflow_count> -ebs_resolution=<file | function | line>

Configuring TAU and choosing a configuration in tau_exec

```
% cd /soft/perftools/tau/tau-2.32; cat .all_configs
 ./configure -ompt -mpi -bfd=download -unwind=download -iowrapper -dwarf=download
      -papi=<dir> -pdt=<dir> -pdt_c++=g++ -otf=download
% make install
% module load tau
% ls $TAU/Makefile*
/soft/perftools/tau/tau-2.32/craycnl/lib/Makefile.tau-gnu-mpi-cupti-pdt
/soft/perftools/tau/tau-2.32/craycnl/lib/Makefile.tau-gnu-papi-mpi-pdt
/soft/perftools/tau/tau-2.32/craycnl/lib/tau-gnu-papi-mpi-pthread-cupti-pdt

% aprun -n 4 tau_exec -T cupti -cupti -ebs ./a.out
Will preload libTAU.so from
/soft/perftools/tau/tau-2.32/craycnl/lib/shared-gnu-mpi-cupti-pdt/

Corresponding to
/soft/perftools/tau/tau-2.32/craycnl/lib/Makefile.tau-gnu-mpi-cupti-pdt

-T mpi is chosen by default. Please use -T serial for non-mpi cases.
```

RUNTIME PRELOADING

- Injects TAU DSO in the executing application
- Requires dynamic executables
- We must compile with –dynamic –g
- Use tau_exec while launching the application

ParaProf Profile Browser

```
% paraprof
```

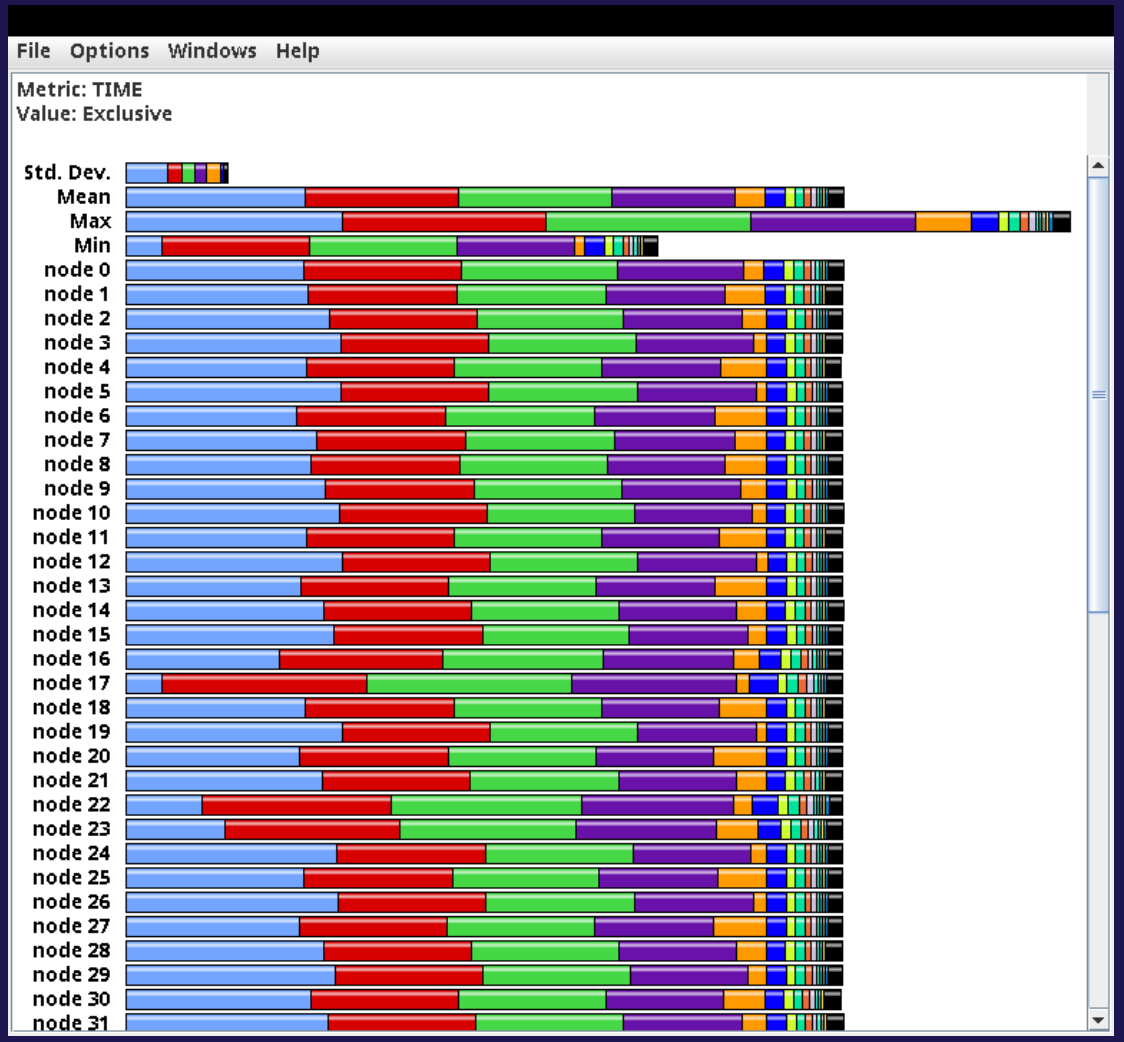

ParaProf Profile Browser

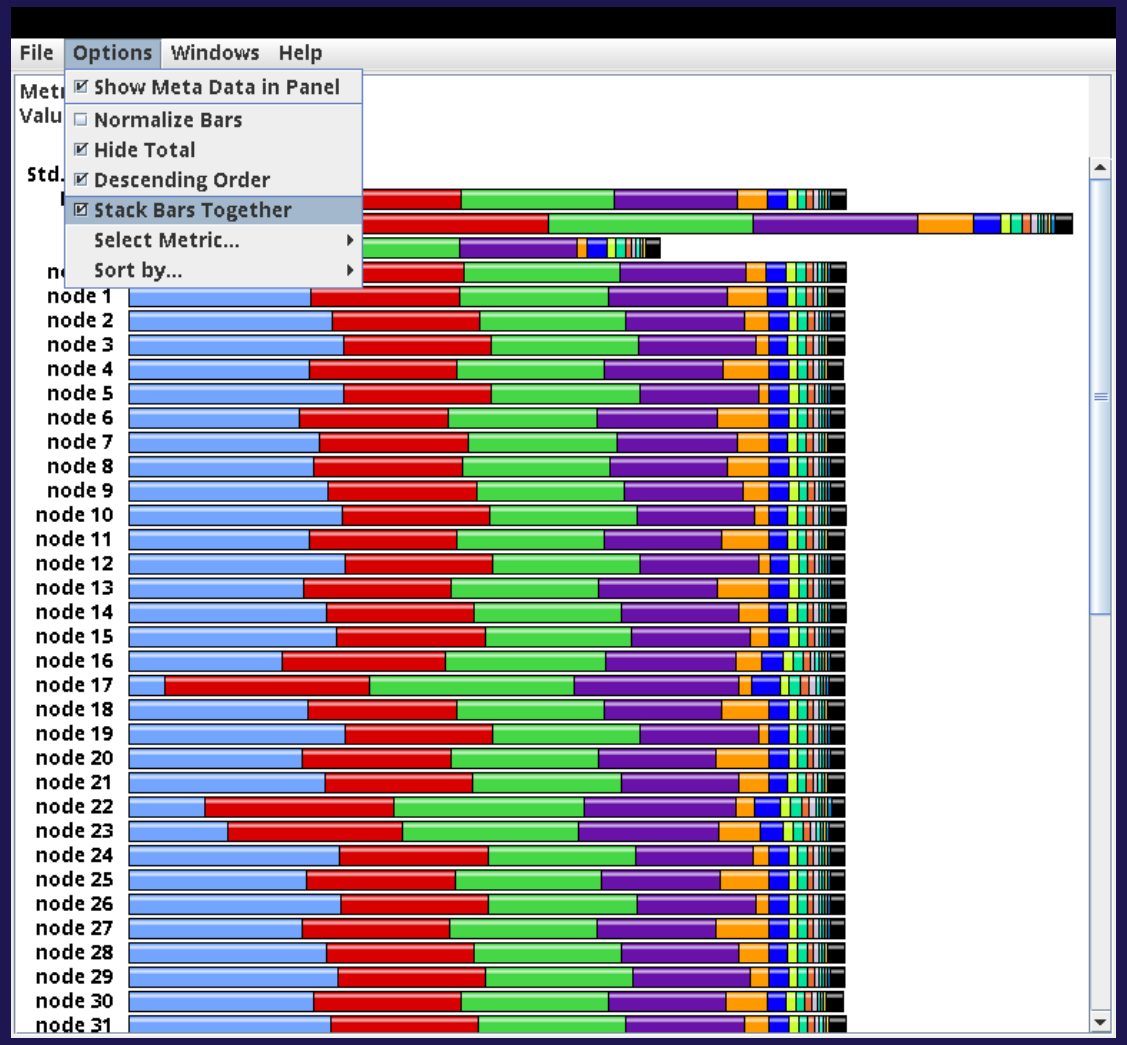

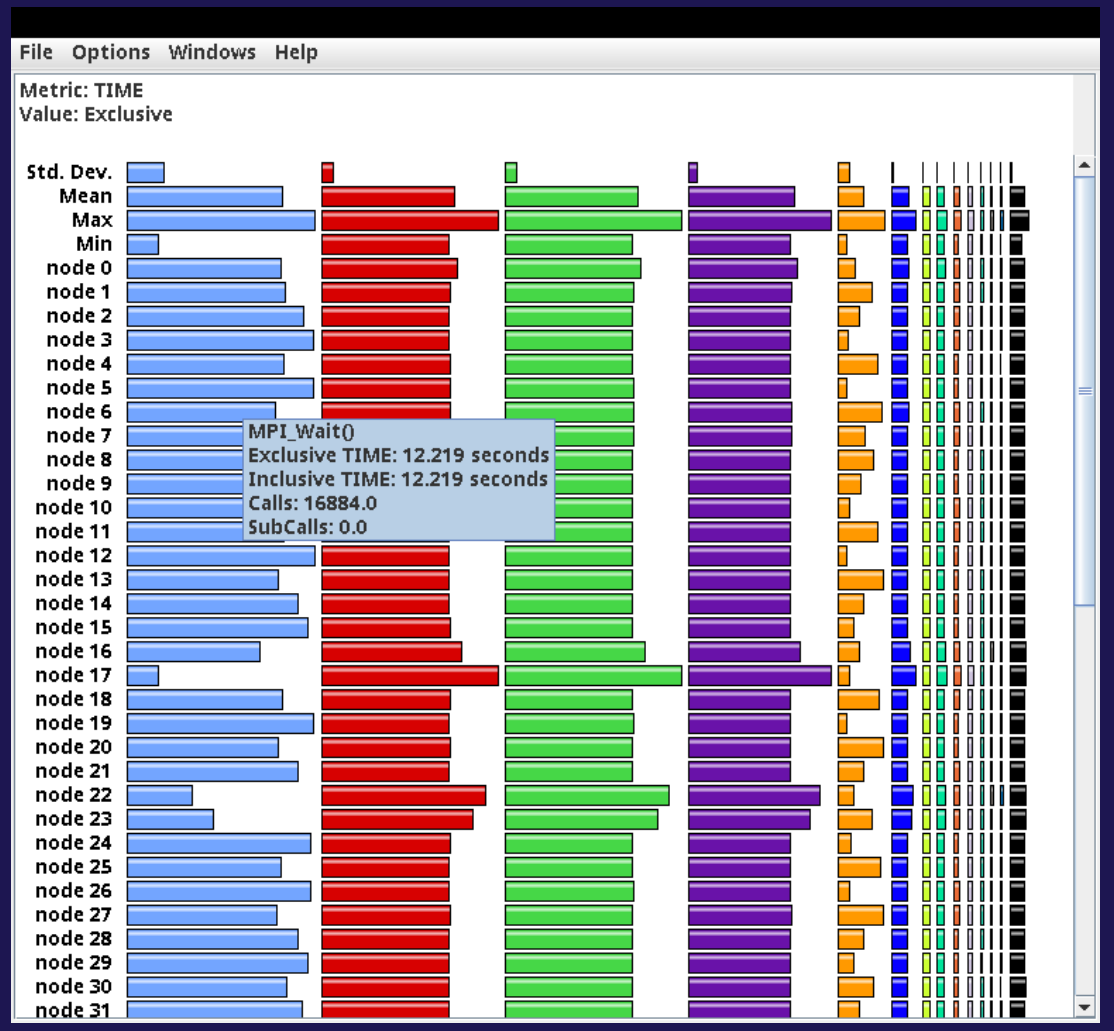

ParaProf Profile Browser

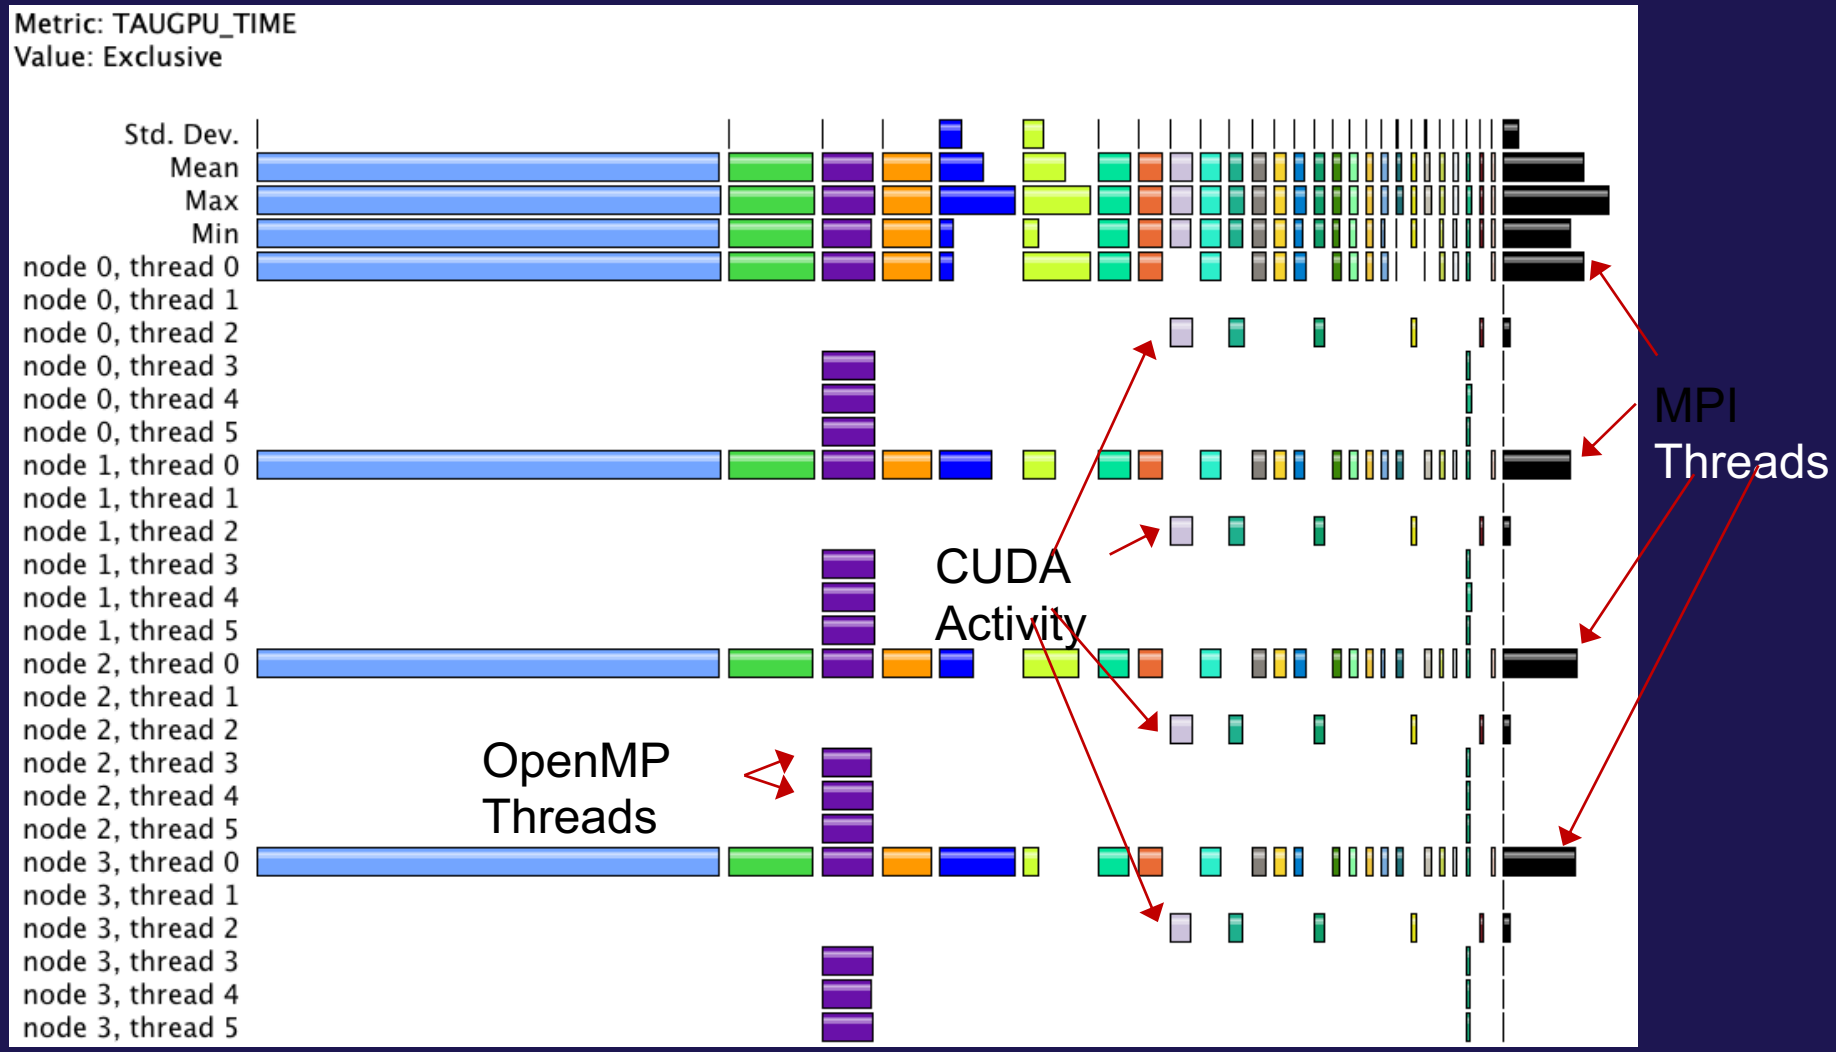

ParaProf Profile Browser: Choose Thread Statistics Window

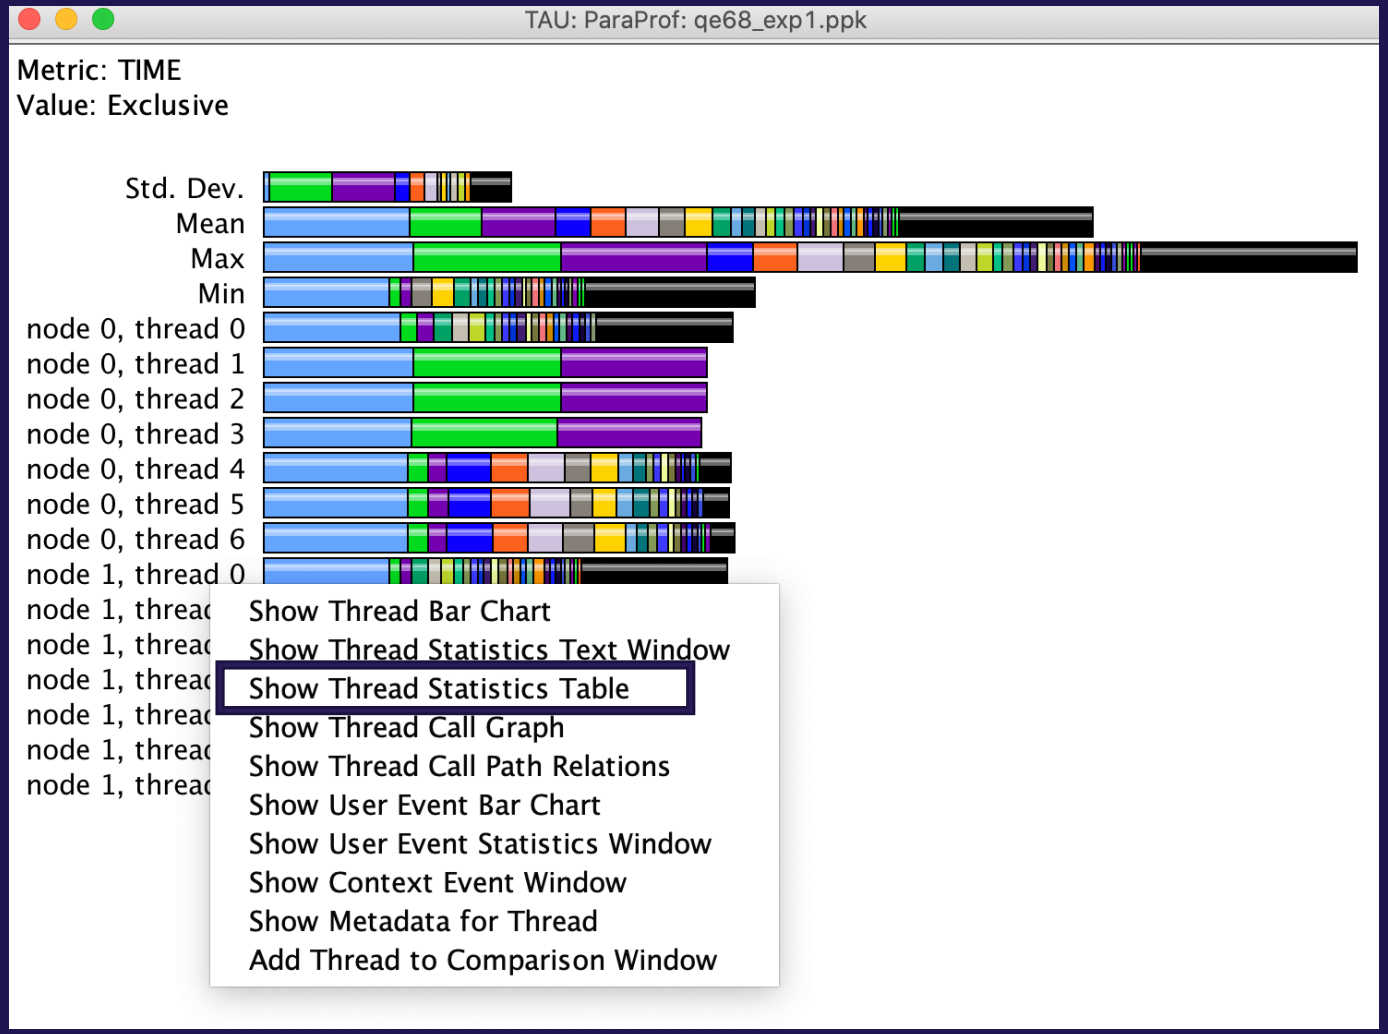

Right click

ParaProf Thread Statistics Table

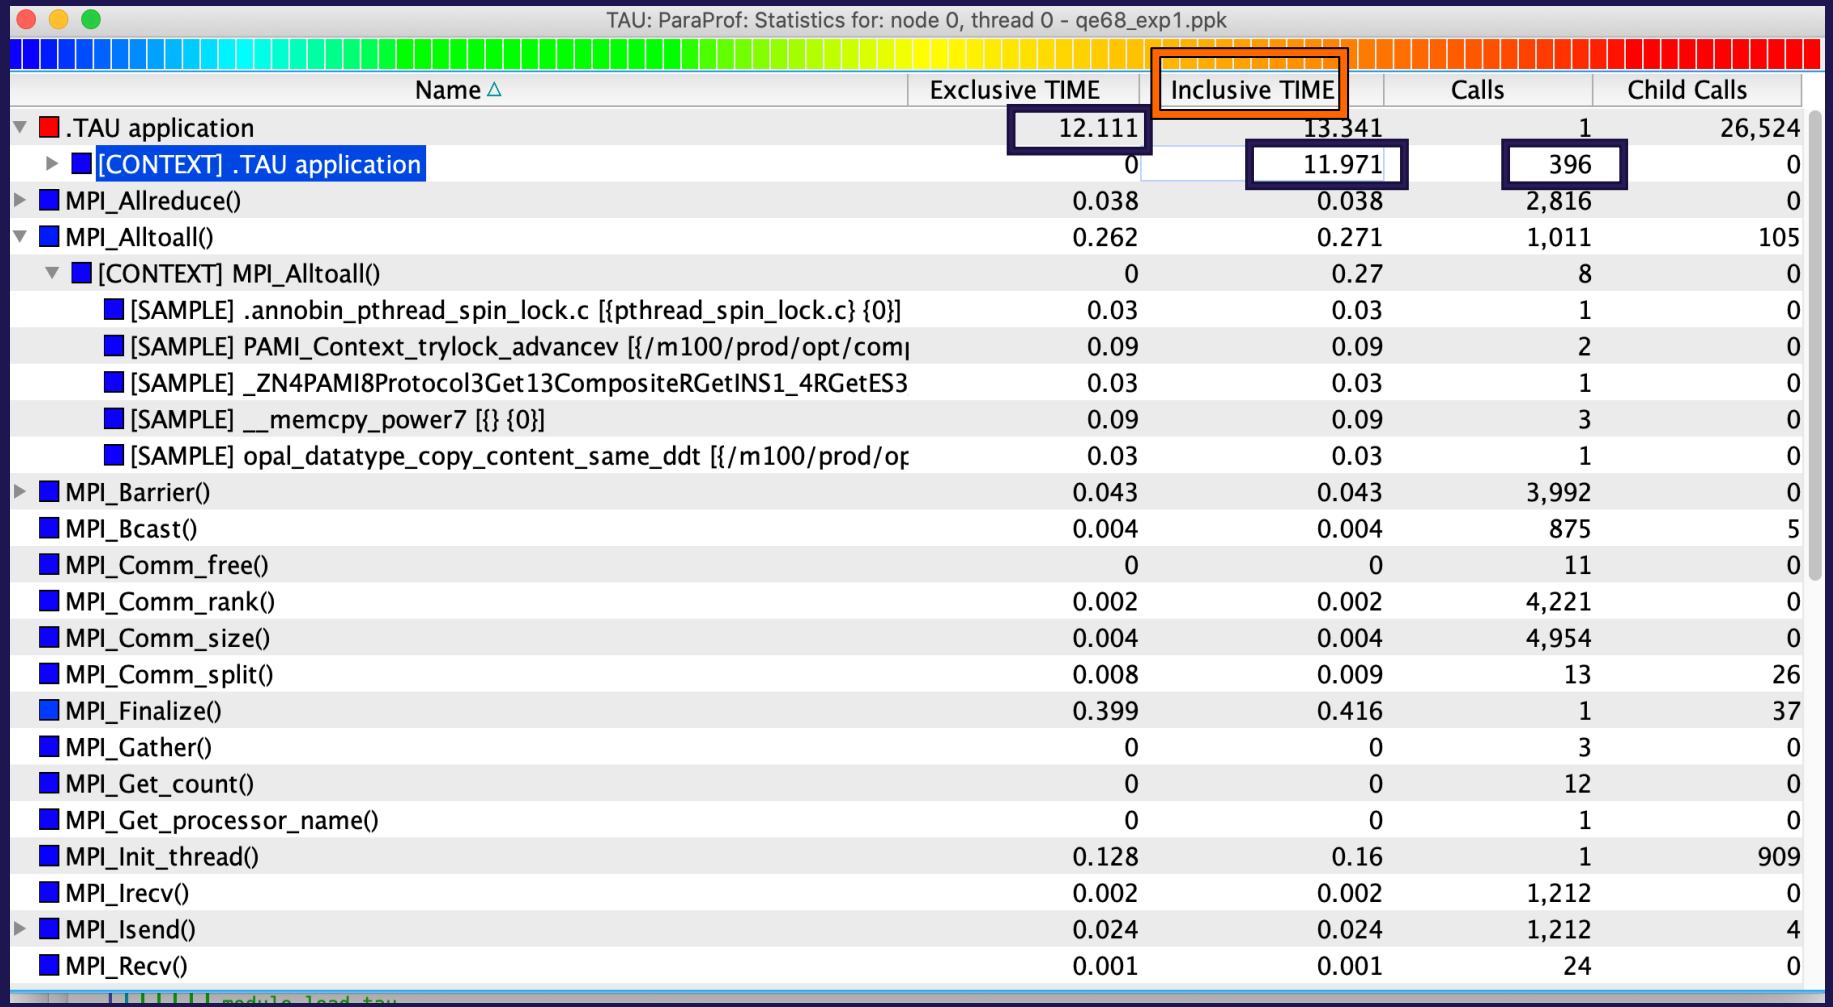
TAU: ParaProf: Statistics for: node 0, thread 0 - qe68_exp1.ppk

Name	Exclusive TIME	Inclusive TIME	Calls	Child Calls
.TAU application	12.111	13.341	1	26,524
[CONTEXT] .TAU application	0	11.971	396	0
MPI_Allreduce()	0.038	0.038	2,816	0
MPI_Alltoall()	0.262	0.271	1,011	105
[CONTEXT] MPI_Alltoall()	0	0.27	8	0
[SAMPLE] .annobin_pthread_spin_lock.c [{pthread_spin_lock.c} {0}]	0.03	0.03	1	0
[SAMPLE] PAMI_Context_trylock_advancev [{/m100/prod/opt/com	0.09	0.09	2	0
[SAMPLE] _ZN4PAMI8Protocol3Get13CompositeRGetINS1_4RGetES3	0.03	0.03	1	0
[SAMPLE] __memcpy_power7 [{} {0}]	0.09	0.09	3	0
[SAMPLE] opal_datatype_copy_content_same_ddt [{/m100/prod/op	0.03	0.03	1	0
MPI_Barrier()	0.043	0.043	3,992	0
MPI_Bcast()	0.004	0.004	875	5
MPI_Comm_free()	0	0	11	0
MPI_Comm_rank()	0.002	0.002	4,221	0
MPI_Comm_size()	0.004	0.004	4,954	0
MPI_Comm_split()	0.008	0.009	13	26
MPI_Finalize()	0.399	0.416	1	37
MPI_Gather()	0	0	3	0
MPI_Get_count()	0	0	12	0
MPI_Get_processor_name()	0	0	1	0
MPI_Init_thread()	0.128	0.16	1	909
MPI_Irecv()	0.002	0.002	1,212	0
MPI_Isend()	0.024	0.024	1,212	4
MPI_Recv()	0.001	0.001	24	0

Using sampling, TAU can explain 11.971 seconds out of 12.111 seconds using 396 samples.

ParaProf Thread Statistics Table

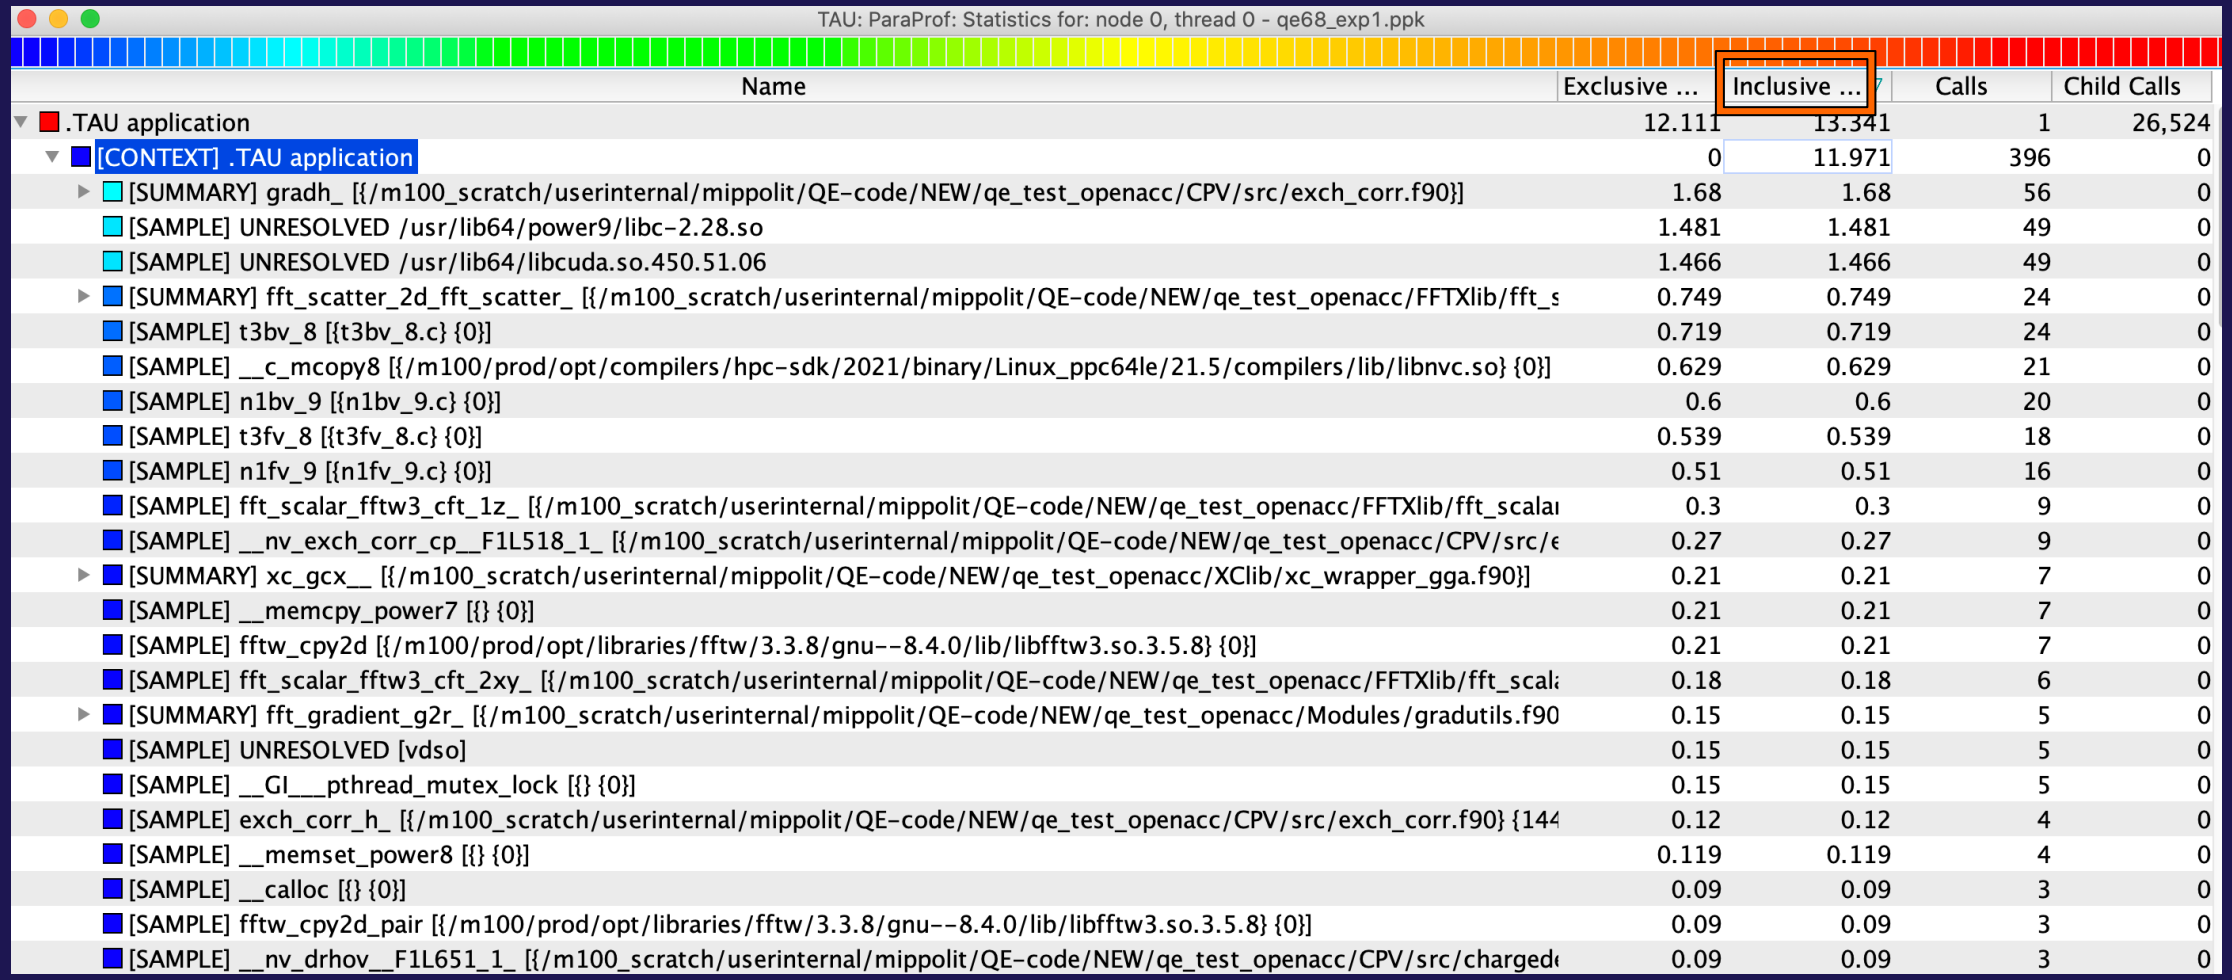

Name	Exclusive ...	Inclusive ...	Calls	Child Calls
.TAU application	12.111	13.341	1	26,524
[CONTEXT] .TAU application	0	11.971	396	0
[SUMMARY] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90}]	1.68	1.68	56	0
[SAMPLE] UNRESOLVED /usr/lib64/power9/libc-2.28.so	1.481	1.481	49	0
[SAMPLE] UNRESOLVED /usr/lib64/libcuda.so.450.51.06	1.466	1.466	49	0
[SUMMARY] fft_scatter_2d_fft_scatter_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/FFTXlib/fft_s	0.749	0.749	24	0
[SAMPLE] t3bv_8 [{t3bv_8.c} {0}]	0.719	0.719	24	0
[SAMPLE] __c_mcopy8 [{/m100/prod/opt/compilers/hpc-sdk/2021/binary/Linux_ppc64le/21.5/compilers/lib/libnvc.so} {0}]	0.629	0.629	21	0
[SAMPLE] n1bv_9 [{n1bv_9.c} {0}]	0.6	0.6	20	0
[SAMPLE] t3fv_8 [{t3fv_8.c} {0}]	0.539	0.539	18	0
[SAMPLE] n1fv_9 [{n1fv_9.c} {0}]	0.51	0.51	16	0
[SAMPLE] fft_scalar_fftw3_cft_1z_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/FFTXlib/fft_scalar	0.3	0.3	9	0
[SAMPLE] __nv_exch_corr_cp__F1L518_1_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/e	0.27	0.27	9	0
[SUMMARY] xc_gcx__ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/XClib/xc_wrapper_gga.f90}]	0.21	0.21	7	0
[SAMPLE] __memcpy_power7 [{} {0}]	0.21	0.21	7	0
[SAMPLE] fftw_cpy2d [{/m100/prod/opt/libraries/fftw/3.3.8/gnu--8.4.0/lib/libfftw3.so.3.5.8} {0}]	0.21	0.21	7	0
[SAMPLE] fft_scalar_fftw3_cft_2xy_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/FFTXlib/fft_scala	0.18	0.18	6	0
[SUMMARY] fft_gradient_g2r_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/Modules/gradutils.f90	0.15	0.15	5	0
[SAMPLE] UNRESOLVED [vdso]	0.15	0.15	5	0
[SAMPLE] __GI___pthread_mutex_lock [{} {0}]	0.15	0.15	5	0
[SAMPLE] exch_corr_h_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {144	0.12	0.12	4	0
[SAMPLE] __memset_power8 [{} {0}]	0.119	0.119	4	0
[SAMPLE] __calloc [{} {0}]	0.09	0.09	3	0
[SAMPLE] fftw_cpy2d_pair [{/m100/prod/opt/libraries/fftw/3.3.8/gnu--8.4.0/lib/libfftw3.so.3.5.8} {0}]	0.09	0.09	3	0
[SAMPLE] __nv_drhov__F1L651_1_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/charged	0.09	0.09	3	0

ParaProf Thread Statistics Table

Name	Exclusive ...	Inclusive ... ▽	Calls	Child Calls
.TAU application	12.111	13.341	1	26,524
[CONTEXT] .TAU application	0	11.971	396	0
[SUMMARY] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90}]	1.68	1.68	56	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {339}]	0.66	0.66	22	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {315}]	0.6	0.6	20	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {344}]	0.18	0.18	6	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {308}]	0.06	0.06	2	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {353}]	0.06	0.06	2	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {320}]	0.03	0.03	1	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {325}]	0.03	0.03	1	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {331}]	0.03	0.03	1	0
[SAMPLE] gradh_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/exch_corr.f90} {303}]	0.03	0.03	1	0
[SAMPLE] UNRESOLVED /usr/lib64/power9/libc-2.28.so	1.481	1.481	49	0
[SAMPLE] UNRESOLVED /usr/lib64/libcuda.so.450.51.06	1.466	1.466	49	0
[SUMMARY] fft_scatter_2d_fft_scatter_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/FFTXlib/fft_s	0.749	0.749	24	0
[SAMPLE] t3bv_8 [{t3bv_8.c} {0}]	0.719	0.719	24	0
[SAMPLE] __c_mcopy8 [{/m100/prod/opt/compilers/hpc-sdk/2021/binary/Linux_ppc64le/21.5/compilers/lib/libnvc.so} {0}]	0.629	0.629	21	0
[SAMPLE] n1bv_9 [{n1bv_9.c} {0}]	0.6	0.6	20	0
[SAMPLE] t3fv_8 [{t3fv_8.c} {0}]	0.539	0.539	18	0
[SAMPLE] n1fv_9 [{n1fv_9.c} {0}]	0.51	0.51	16	0
[SAMPLE] fft_scalar_fftw3_cft_1z_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/FFTXlib/fft_scalar	0.3	0.3	9	0
[SAMPLE] __nv_exch_corr_cp__F1L518_1_ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/CPV/src/e	0.27	0.27	9	0
[SUMMARY] xc_gcx__ [{/m100_scratch/userinternal/mippolit/QE-code/NEW/qe_test_openacc/XClib/xc_wrapper_gga.f90}]	0.21	0.21	7	0
[SAMPLE] __memcpy_power7 [{} {0}]	0.21	0.21	7	0
[SAMPLE] fftw_cpy2d [{/m100/prod/opt/libraries/fftw/3.3.8/gnu--8.4.0/lib/libfftw3.so.3.5.8} {0}]	0.21	0.21	7	0

TAU supports Python, MPI, and CUDA

- Without any modification to the source code or DSOs or interpreter, it instruments and samples the application using Python, MPI, and CUDA instrumentation. TAU needs to be built with the same Python as the application.

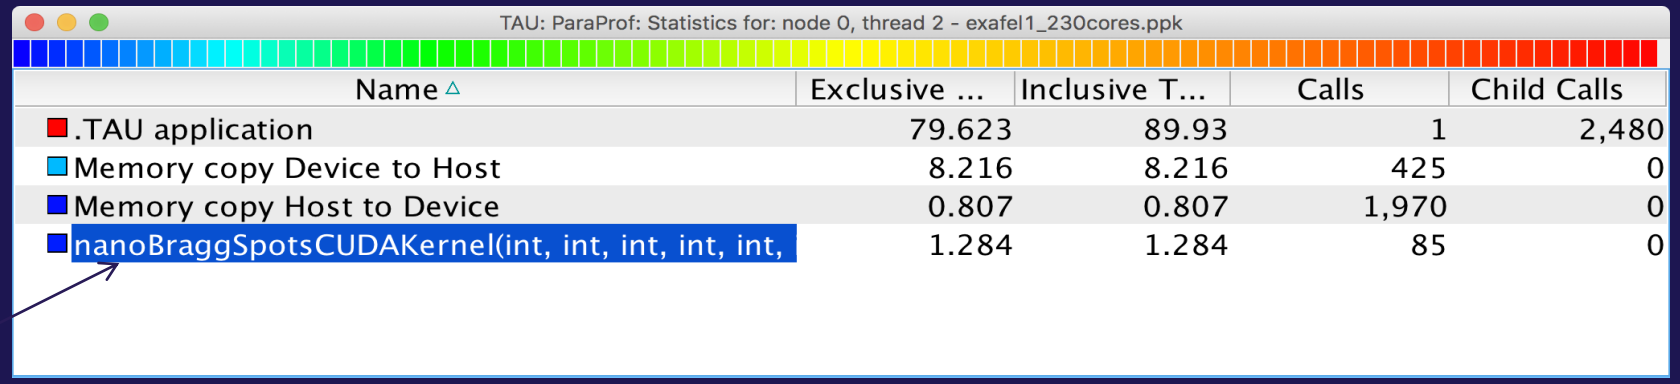

Kernel on GPU

% aprun –np 230 tau_python –T cupti,mpi,pdt –ebs –cupti ./exafel.py
Instead of:
% aprun –np 230 python ./exafel.py

TAU Thread Statistics Table

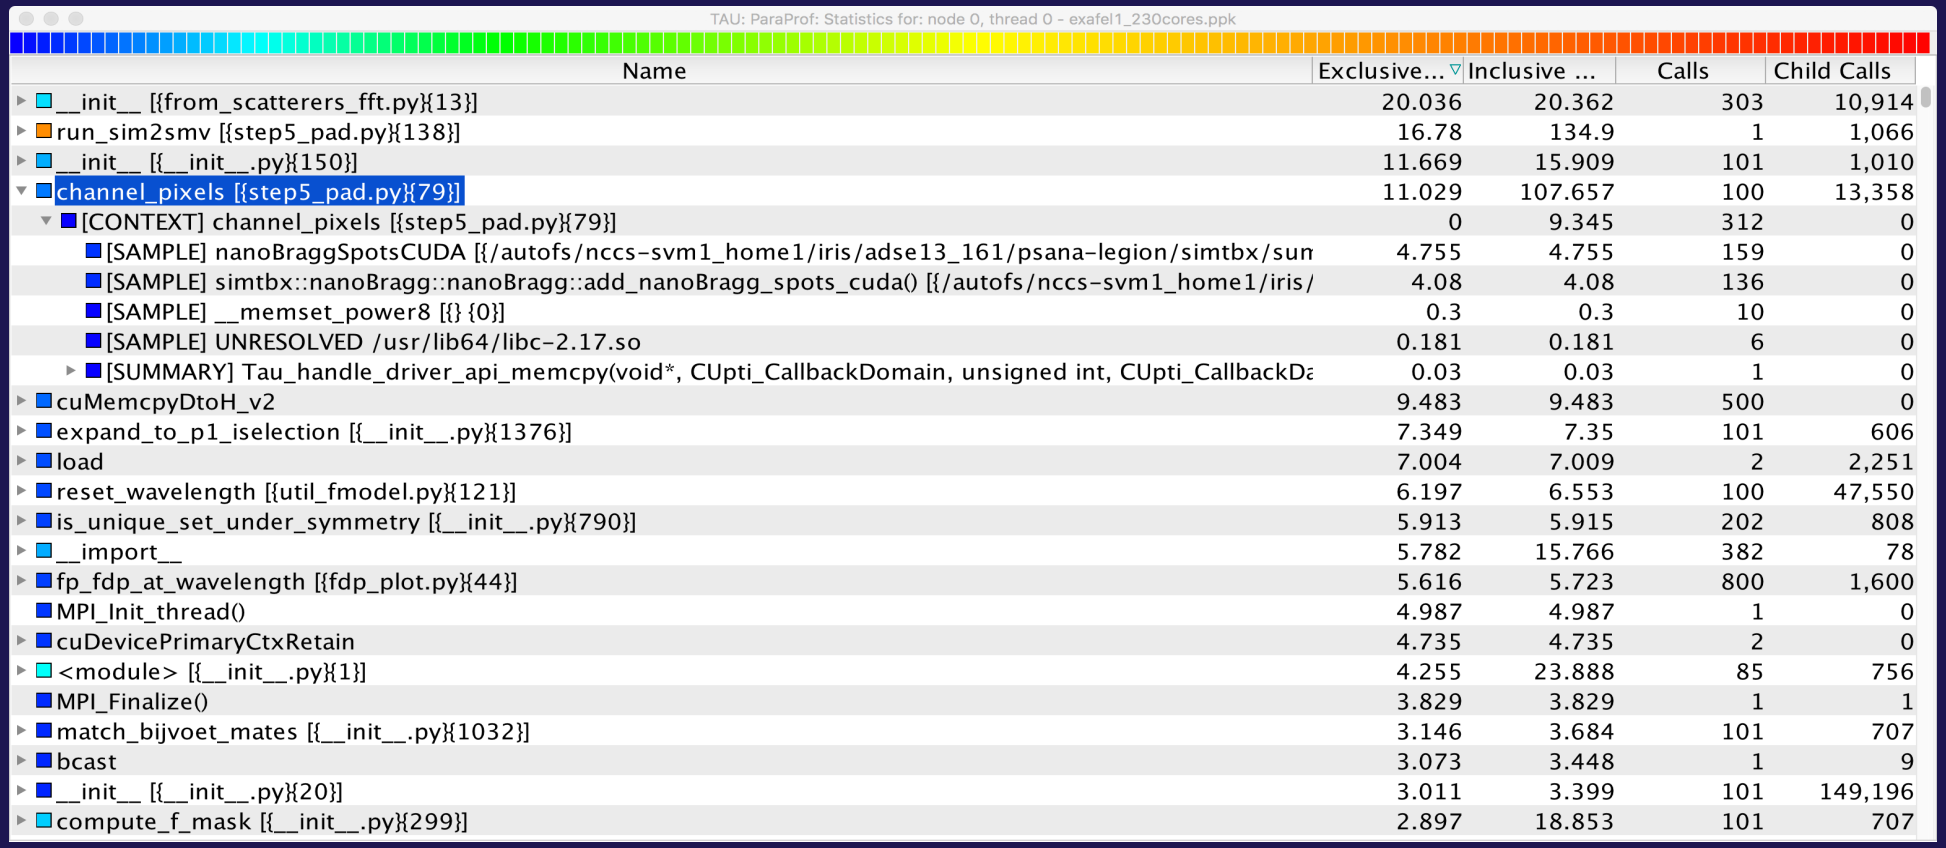

Name	Exclusive...	Inclusive ...	Calls	Child Calls
__init__ [{from_scatterers_fft.py}{13}]	20.036	20.362	303	10,914
run_sim2smv [{step5_pad.py}{138}]	16.78	134.9	1	1,066
__init__ [{__init__.py}{150}]	11.669	15.909	101	1,010
channel_pixels [{step5_pad.py}{79}]	11.029	107.657	100	13,358
[CONTEXT] channel_pixels [{step5_pad.py}{79}]	0	9.345	312	0
[SAMPLE] nanoBraggSpotsCUDA [{/autofs/nccs-svm1_home1/iris/adse13_161/psana-legion/simtbx/sun	4.755	4.755	159	0
[SAMPLE] simtbx::nanoBragg::nanoBragg::add_nanoBragg_spots_cuda() [{/autofs/nccs-svm1_home1/iris/	4.08	4.08	136	0
[SAMPLE] __memset_power8 [{} {0}]	0.3	0.3	10	0
[SAMPLE] UNRESOLVED /usr/lib64/libc-2.17.so	0.181	0.181	6	0
[SUMMARY] Tau_handle_driver_api_memcpy(void*, CUpti_CallbackDomain, unsigned int, CUpti_CallbackDa	0.03	0.03	1	0
cuMemcpyDtoH_v2	9.483	9.483	500	0
expand_to_p1_iselection [{__init__.py}{1376}]	7.349	7.35	101	606
load	7.004	7.009	2	2,251
reset_wavelength [{util_fmodel.py}{121}]	6.197	6.553	100	47,550
is_unique_set_under_symmetry [{__init__.py}{790}]	5.913	5.915	202	808
__import__	5.782	15.766	382	78
fp_fdp_at_wavelength [{fdp_plot.py}{44}]	5.616	5.723	800	1,600
MPI_Init_thread()	4.987	4.987	1	0
cuDevicePrimaryCtxRetain	4.735	4.735	2	0
<module> [{__init__.py}{1}]	4.255	23.888	85	756
MPI_Finalize()	3.829	3.829	1	1
match_bijvoet_mates [{__init__.py}{1032}]	3.146	3.684	101	707
bcast	3.073	3.448	1	9
__init__ [{__init__.py}{20}]	3.011	3.399	101	149,196
compute_f_mask [{__init__.py}{299}]	2.897	18.853	101	707

Python, MPI, CUDA, and samples from DSOs are all integrated in a single view

ParaProf

- Click on Columns:
- to sort by incl time

- Open binvcrhs
- Click on Sample

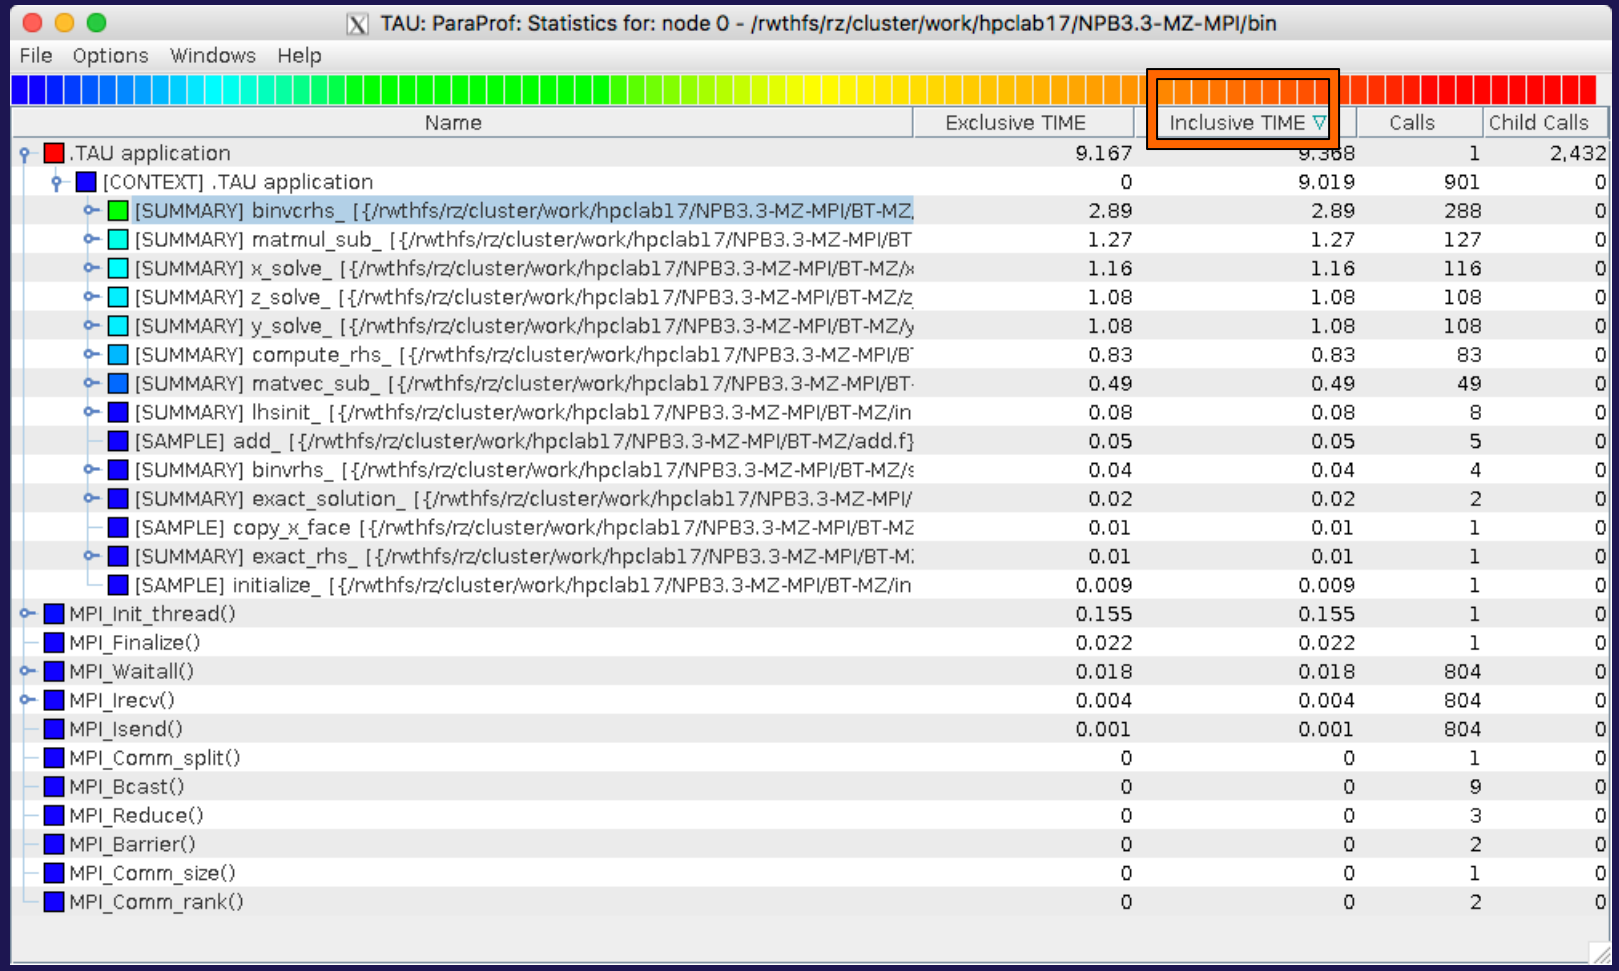

ParaProf

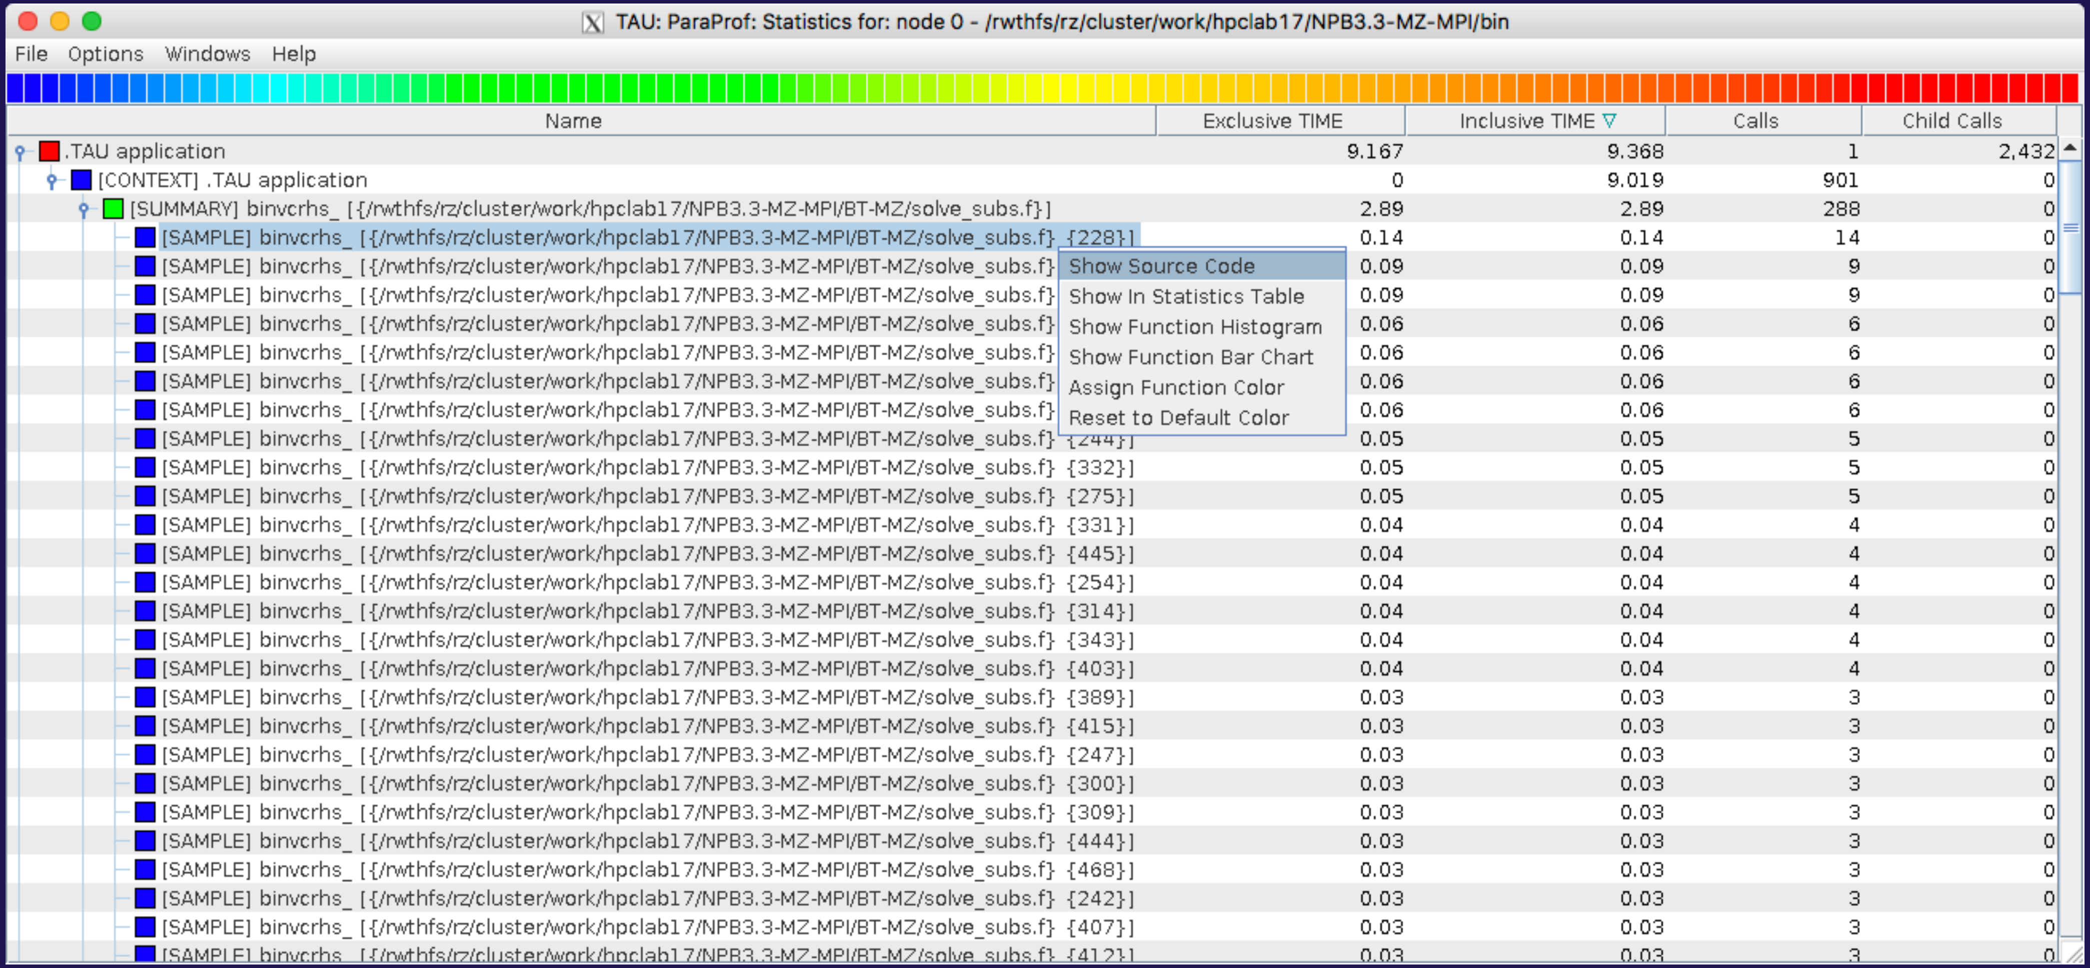

Right click

TAU Context Event Window

TAU: ParaProf: Context Events for: node 0, thread 0 - exafel1_230cores.ppk

Name	Total	NumSamples	MaxValue	MinValue	MeanValue	Std. Dev.
<module> [{step5_batch.py}{1}]						
tst_one [{step5_batch.py}{23}]						
run_sim2smv [{step5_pad.py}{138}]						
channel_pixels [{step5_pad.py}{79}]						
cudaMemcpy						
Bytes copied from Device to Host	15,300,000,000	500	36,000,000	9,000,000	30,600,000	10,800,000
Bytes copied from Host to Device	15,423,816,000	2,300	36,000,000	8	6,706,006.957	13,564,989.185
cuMemcpyHtoD_v2						
Bytes copied from Host to Device	15,423,816,000	2,300	36,000,000	8	6,706,006.957	13,564,989.185
cuMemcpyDtoH_v2						
Bytes copied from Device to Host	15,300,000,000	500	36,000,000	9,000,000	30,600,000	10,800,000
Bytes copied from Device to Host	30,600,000,000	1,000	36,000,000	9,000,000	30,600,000	10,800,000
Bytes copied from Host to Device	30,847,632,000	4,600	36,000,000	8	6,706,006.957	13,564,989.185
Message size for broadcast	827,971,798	2	827,971,794	4	413,985,899	413,985,895

TAU tracks the data transfers between the host and the GPU.

TAU's tracking of Python and MPI

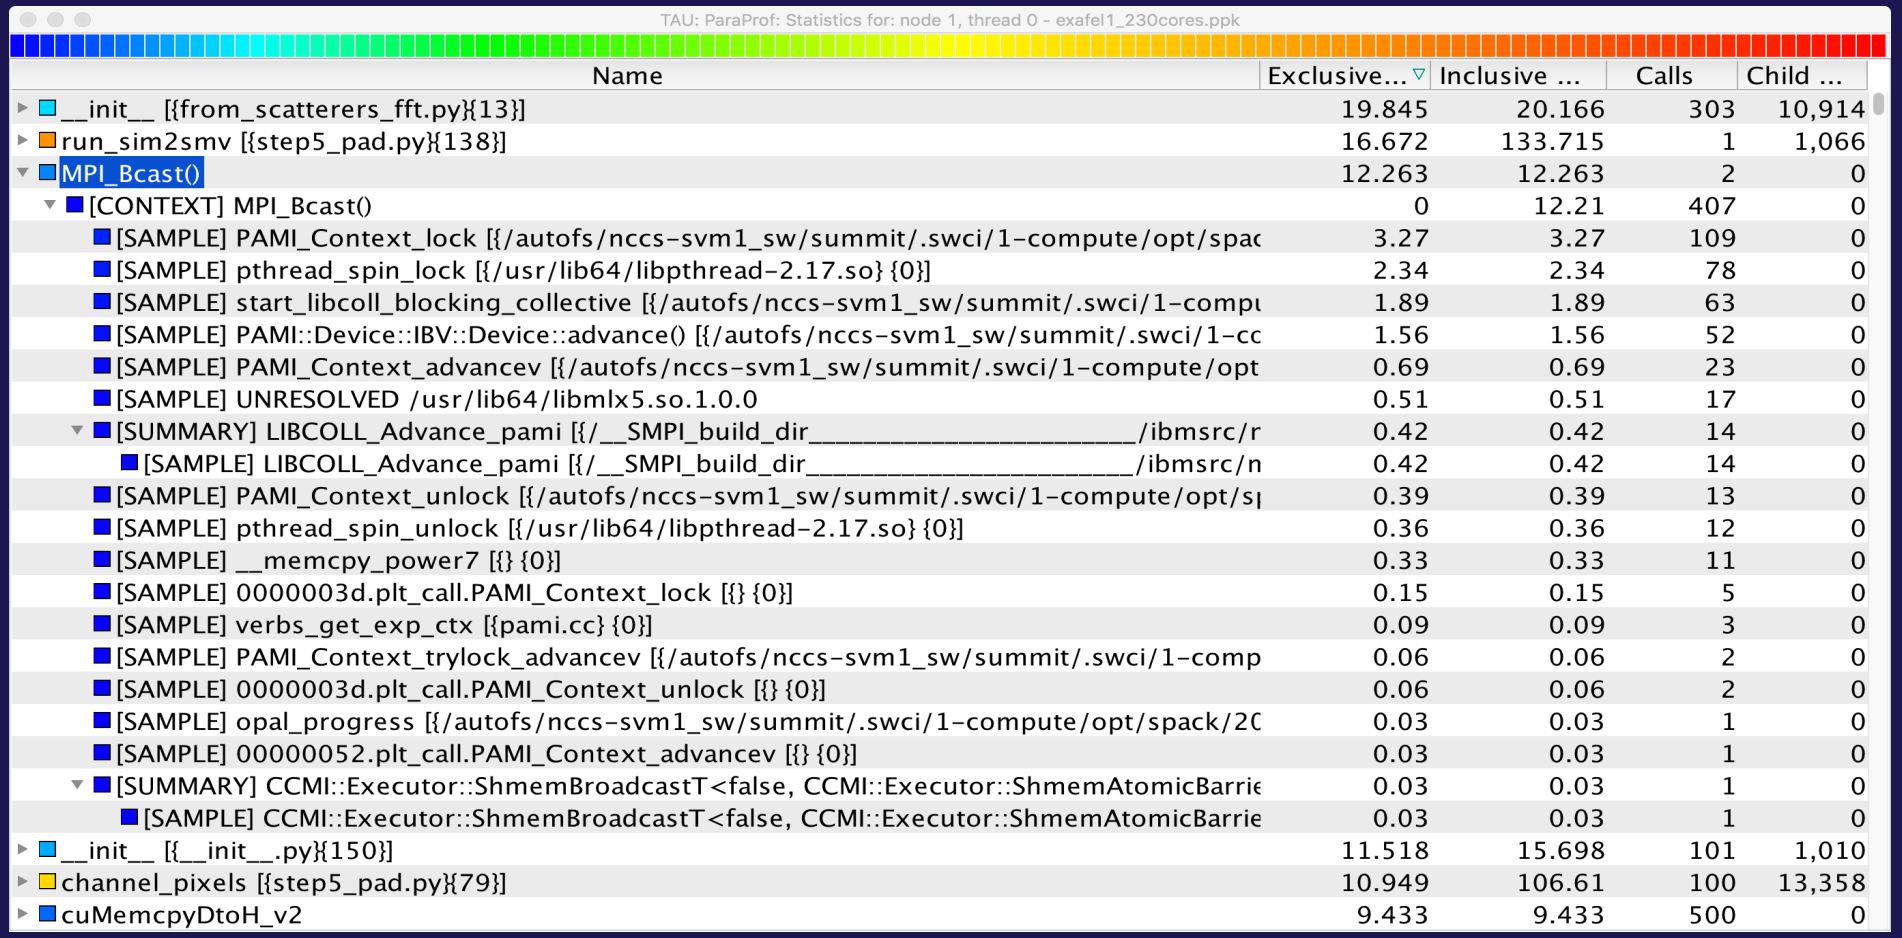

Name	Exclusive...	Inclusive ...	Calls	Child ...
__init__ [{from_scatterers_fft.py}{13}]	19.845	20.166	303	10,914
run_sim2smv [{step5_pad.py}{138}]	16.672	133.715	1	1,066
MPI_Bcast()	12.263	12.263	2	0
[CONTEXT] MPI_Bcast()	0	12.21	407	0
[SAMPLE] PAMI_Context_lock [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt/spac	3.27	3.27	109	0
[SAMPLE] pthread_spin_lock [{/usr/lib64/libpthread-2.17.so} {0}]	2.34	2.34	78	0
[SAMPLE] start_libcoll_blocking_collective [{/autofs/nccs-svm1_sw/summit/.swci/1-compu	1.89	1.89	63	0
[SAMPLE] PAMI::Device::IBV::Device::advance() [{/autofs/nccs-svm1_sw/summit/.swci/1-cc	1.56	1.56	52	0
[SAMPLE] PAMI_Context_advancev [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt	0.69	0.69	23	0
[SAMPLE] UNRESOLVED /usr/lib64/libmlx5.so.1.0.0	0.51	0.51	17	0
[SUMMARY] LIBCOLL_Advance_pami [{/__SMPI_build_dir_____/ibmsrc/r	0.42	0.42	14	0
[SAMPLE] LIBCOLL_Advance_pami [{/__SMPI_build_dir_____/ibmsrc/n	0.42	0.42	14	0
[SAMPLE] PAMI_Context_unlock [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt/s	0.39	0.39	13	0
[SAMPLE] pthread_spin_unlock [{/usr/lib64/libpthread-2.17.so} {0}]	0.36	0.36	12	0
[SAMPLE] __memcpy_power7 [{} {0}]	0.33	0.33	11	0
[SAMPLE] 0000003d.plt_call.PAMI_Context_lock [{} {0}]	0.15	0.15	5	0
[SAMPLE] verbs_get_exp_ctx [{pami.cc} {0}]	0.09	0.09	3	0
[SAMPLE] PAMI_Context_trylock_advancev [{/autofs/nccs-svm1_sw/summit/.swci/1-comp	0.06	0.06	2	0
[SAMPLE] 0000003d.plt_call.PAMI_Context_unlock [{} {0}]	0.06	0.06	2	0
[SAMPLE] opal_progress [{/autofs/nccs-svm1_sw/summit/.swci/1-compute/opt/spack/2(	0.03	0.03	1	0
[SAMPLE] 00000052.plt_call.PAMI_Context_advancev [{} {0}]	0.03	0.03	1	0
[SUMMARY] CCMI::Executor::ShmemBroadcastT<false, CCMI::Executor::ShmemAtomicBarrie	0.03	0.03	1	0
[SAMPLE] CCMI::Executor::ShmemBroadcastT<false, CCMI::Executor::ShmemAtomicBarrie	0.03	0.03	1	0
__init__ [{__init__.py}{150}]	11.518	15.698	101	1,010
channel_pixels [{step5_pad.py}{79}]	10.949	106.61	100	13,358
cuMemcpyDtoH_v2	9.433	9.433	500	0

TAU can observe events in closed-source vendor libraries (e.g., in MPI_Bcast)!

Callstack Sampling in TAU

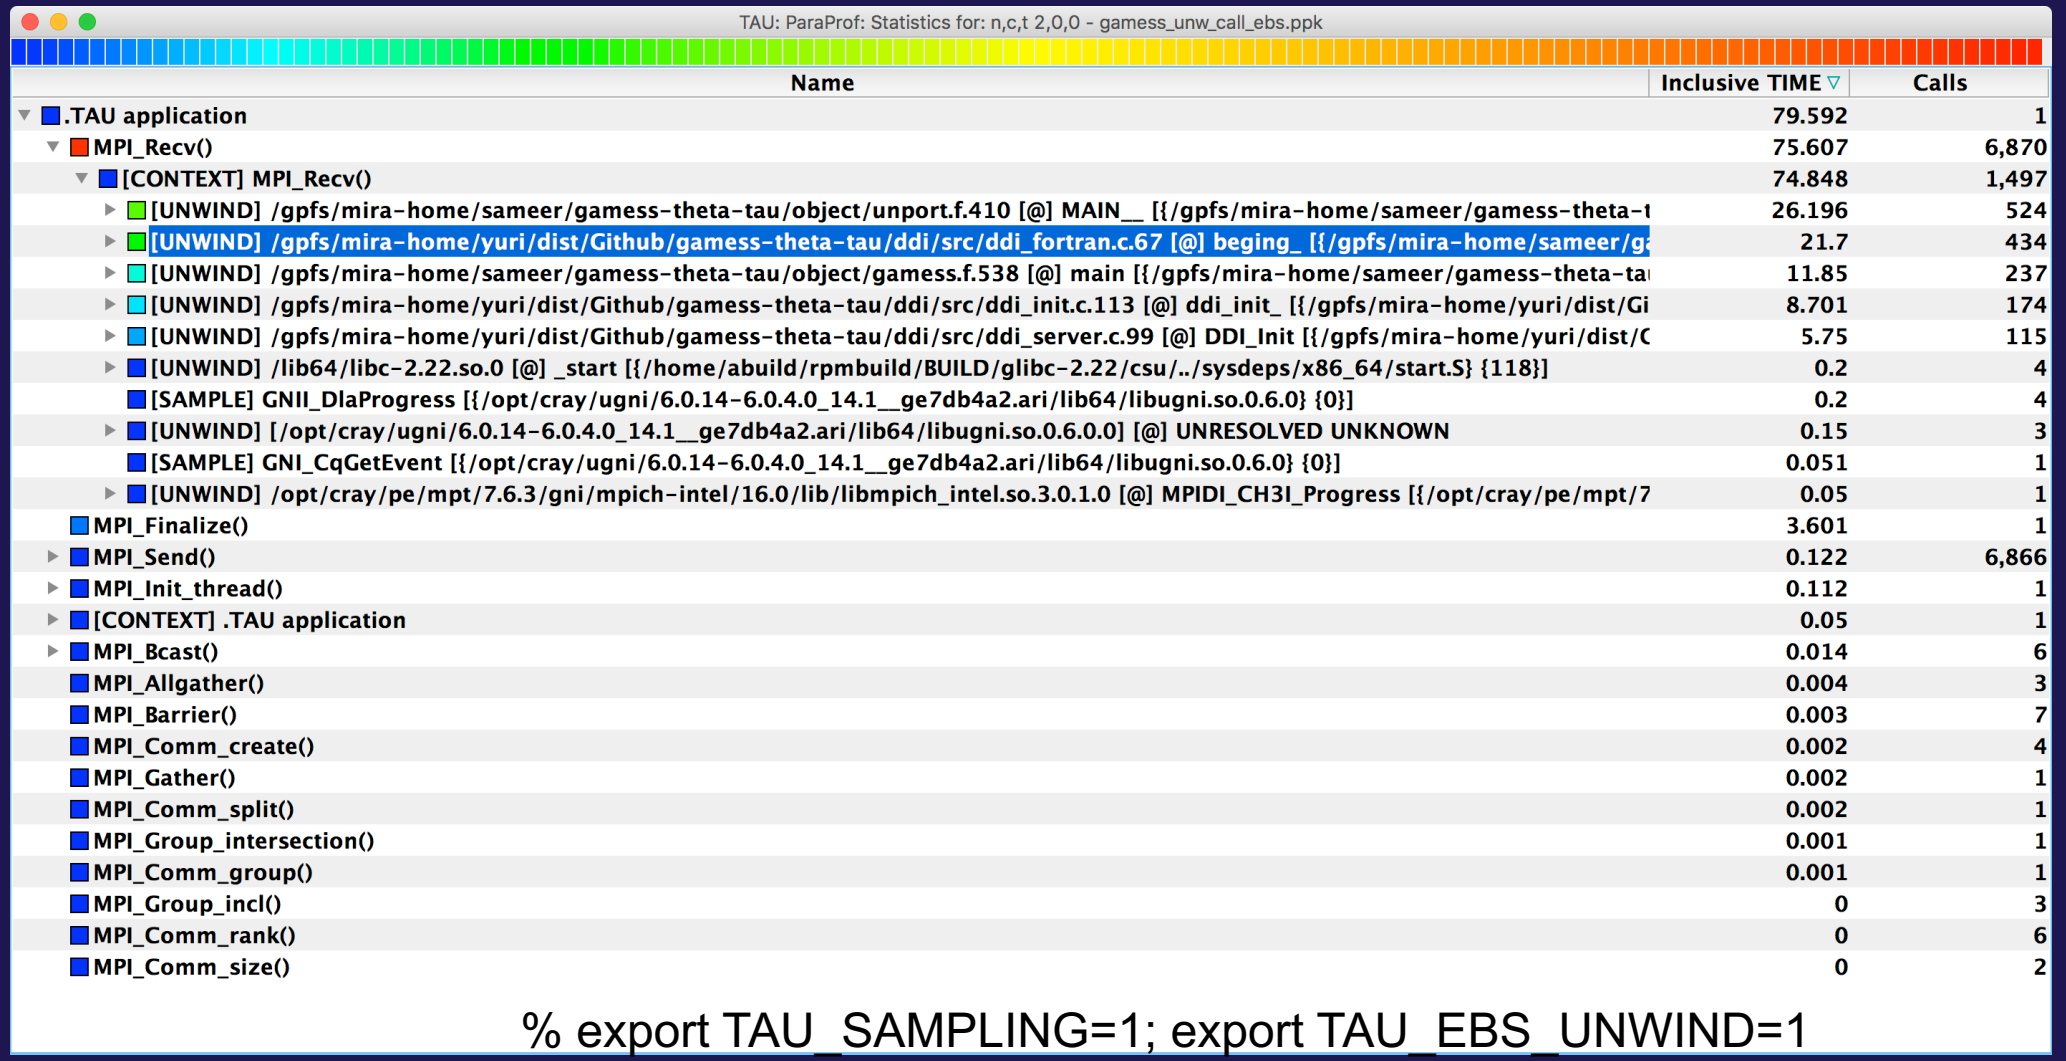

Name	Inclusive TIME ▽	Calls
.TAU application	79.592	1
MPI_Recv()	75.607	6,870
[CONTEXT] MPI_Recv()	74.848	1,497
[UNWIND] /gpfs/mira-home/sameer/gamess-theta-tau/object/unport.f.410 [@] MAIN__ [{/gpfs/mira-home/sameer/gamess-theta-t	26.196	524
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_fortran.c.67 [@] beging_ [{/gpfs/mira-home/sameer/g	21.7	434
[UNWIND] /gpfs/mira-home/sameer/gamess-theta-tau/object/gamess.f.538 [@] main [{/gpfs/mira-home/sameer/gamess-theta-ta	11.85	237
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_init.c.113 [@] ddi_init_ [{/gpfs/mira-home/yuri/dist/Gi	8.701	174
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_server.c.99 [@] DDI_Init [{/gpfs/mira-home/yuri/dist/C	5.75	115
[UNWIND] /lib64/libc-2.22.so.0 [@] _start [{/home/abuild/rpmbuild/BUILD/glibc-2.22/csu/../sysdeps/x86_64/start.S} {118}]	0.2	4
[SAMPLE] GNII_DlaProgress [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	0.2	4
[UNWIND] [/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0.0] [@] UNRESOLVED UNKNOWN	0.15	3
[SAMPLE] GNI_CqGetEvent [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	0.051	1
[UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPIDI_CH3I_Progress [{/opt/cray/pe/mpt/7	0.05	1
MPI_Finalize()	3.601	1
MPI_Send()	0.122	6,866
MPI_Init_thread()	0.112	1
[CONTEXT] .TAU application	0.05	1
MPI_Bcast()	0.014	6
MPI_Allgather()	0.004	3
MPI_Barrier()	0.003	7
MPI_Comm_create()	0.002	4
MPI_Gather()	0.002	1
MPI_Comm_split()	0.002	1
MPI_Group_intersection()	0.001	1
MPI_Comm_group()	0.001	1
MPI_Group_incl()	0	3
MPI_Comm_rank()	0	6
MPI_Comm_size()	0	2

% export TAU_SAMPLING=1; export TAU_EBS_UNWIND=1

UNWINDING CALLSTACKS

Name	Inclusive TIME	Calls
.TAU application	79.592	1
MPI_Recv()	75.607	6,870
[CONTEXT] MPI_Recv()	74.848	1,497
[UNWIND] /gpfs/mira-home/sameer/gamess-theta-tau/object/unport.f.410 [@] MAIN__ [{/gpfs/mira-home/sameer/gamess-theta-	26.196	524
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_fortran.c.67 [@] beging_ [{/gpfs/mira-home/sameer/g	21.7	434
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_init.c.113 [@] ddi_init_ [{/gpfs/mira-home/yuri/dist	21.7	434
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_server.c.99 [@] DDI_Init [{/gpfs/mira-home/yuri/	21.7	434
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_recv.c.65 [@] DDI_Server [{/gpfs/mira-home/y	21.7	434
[UNWIND] /lus/theta-fs0/software/perftools/tau/tau-2.26.3/src/Profile/TauMpi.c.2371 [@] DDI_Recv_request [{/gpfs/mira	21.7	434
[UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPI_Recv [{/lus/theta-fs0/sof	21.7	434
[UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] PMPI_Recv [{/opt/cray/pe/n	21.7	434
[UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPIDI_CH3I_Progress [{/(	21.45	429
[UNWIND] /opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0.0 [@] MPID_nem_gni_poll [{/	15.95	319
[SAMPLE] GNI_SmsgGetNextWTag [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0}	10.349	207
[SAMPLE] GNI_CqGetEvent [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	5.6	112
[UNWIND] gni_poll.c.0 [@] MPID_nem_gni_poll [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_inte	5.25	105
[UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPID_nem_gni_poll [{/	0.25	5
[UNWIND] UNRESOLVED [@] MPIDI_CH3I_Progress [{/opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_int	0.25	5
[UNWIND] /gpfs/mira-home/sameer/gamess-theta-tau/object/gamess.f.538 [@] main [{/gpfs/mira-home/sameer/gamess-theta-ta	11.85	237
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_init.c.113 [@] ddi_init_ [{/gpfs/mira-home/yuri/dist/G	8.701	174
[UNWIND] /gpfs/mira-home/yuri/dist/Github/gamess-theta-tau/ddi/src/ddi_server.c.99 [@] DDI_Init [{/gpfs/mira-home/yuri/dist/(	5.75	115
[UNWIND] /lib64/libc-2.22.so.0 [@] _start [{/home/abuild/rpmbuild/BUILD/glibc-2.22/csu/../sysdeps/x86_64/start.S} {118}]	0.2	4
[SAMPLE] GNII_DlaProgress [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	0.2	4
[UNWIND] [/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0.0] [@] UNRESOLVED UNKNOWN	0.15	3
[SAMPLE] GNI_CqGetEvent [{/opt/cray/ugni/6.0.14-6.0.4.0_14.1__ge7db4a2.ari/lib64/libugni.so.0.6.0} {0}]	0.051	1
[UNWIND] /opt/cray/pe/mpt/7.6.3/gni/mpich-intel/16.0/lib/libmpich_intel.so.3.0.1.0 [@] MPIDI_CH3I_Progress [{/opt/cray/pe/mpt/	0.05	1
MPI_Finalize()	3.601	1
MPI_Send()	0.122	6,866
MPI_Init_thread()	0.112	1
[CONTEXT] .TAU application	0.05	1

% export TAU_SAMPLING=1; export TAU_EBS_UNWIND=1

Event-Based Sampling (EBS)

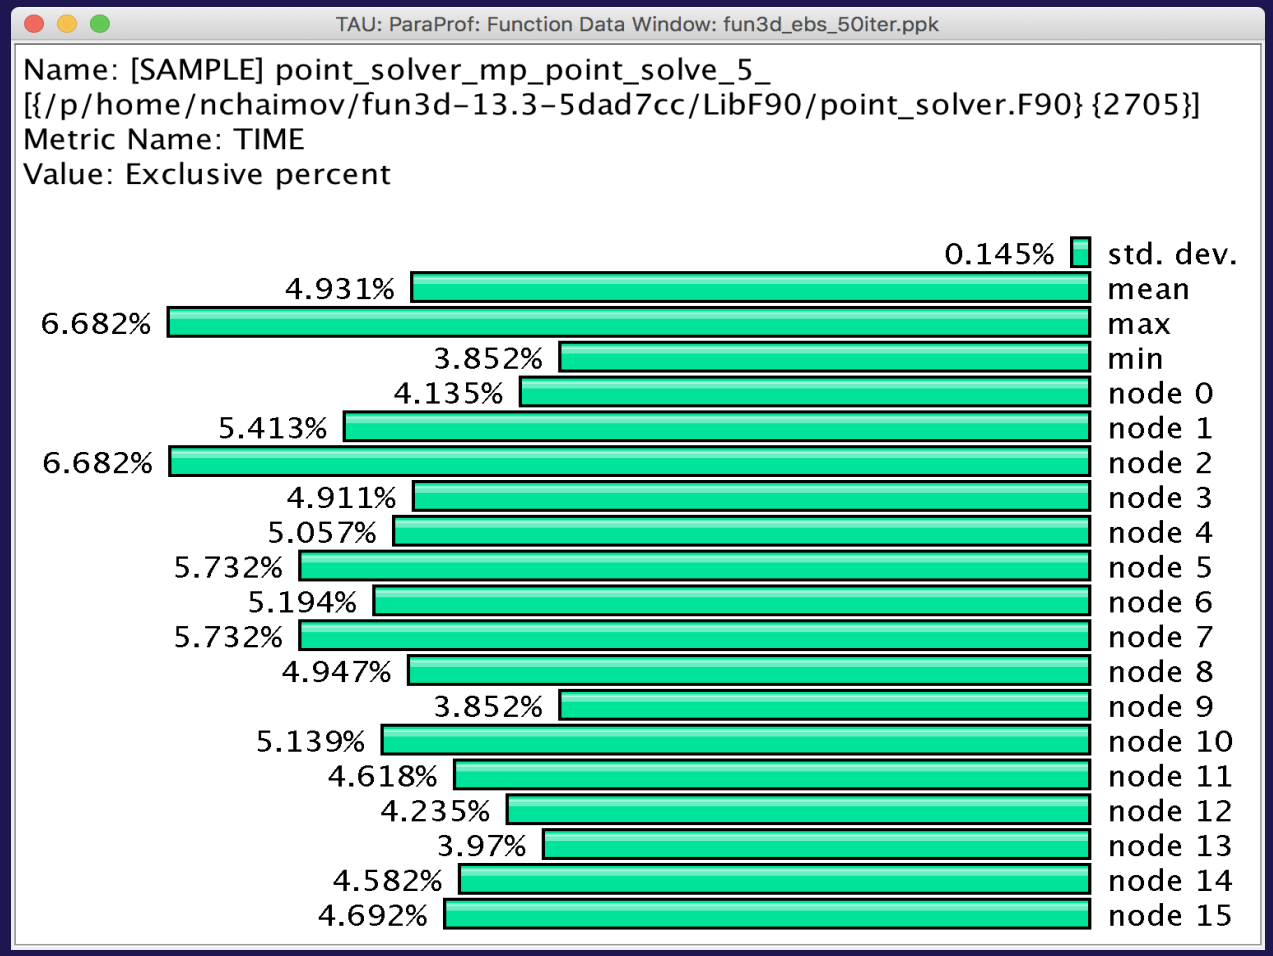


```
% aprun -n 16 tau_exec -ebs a.out
```

Callsite Profiling and Tracing

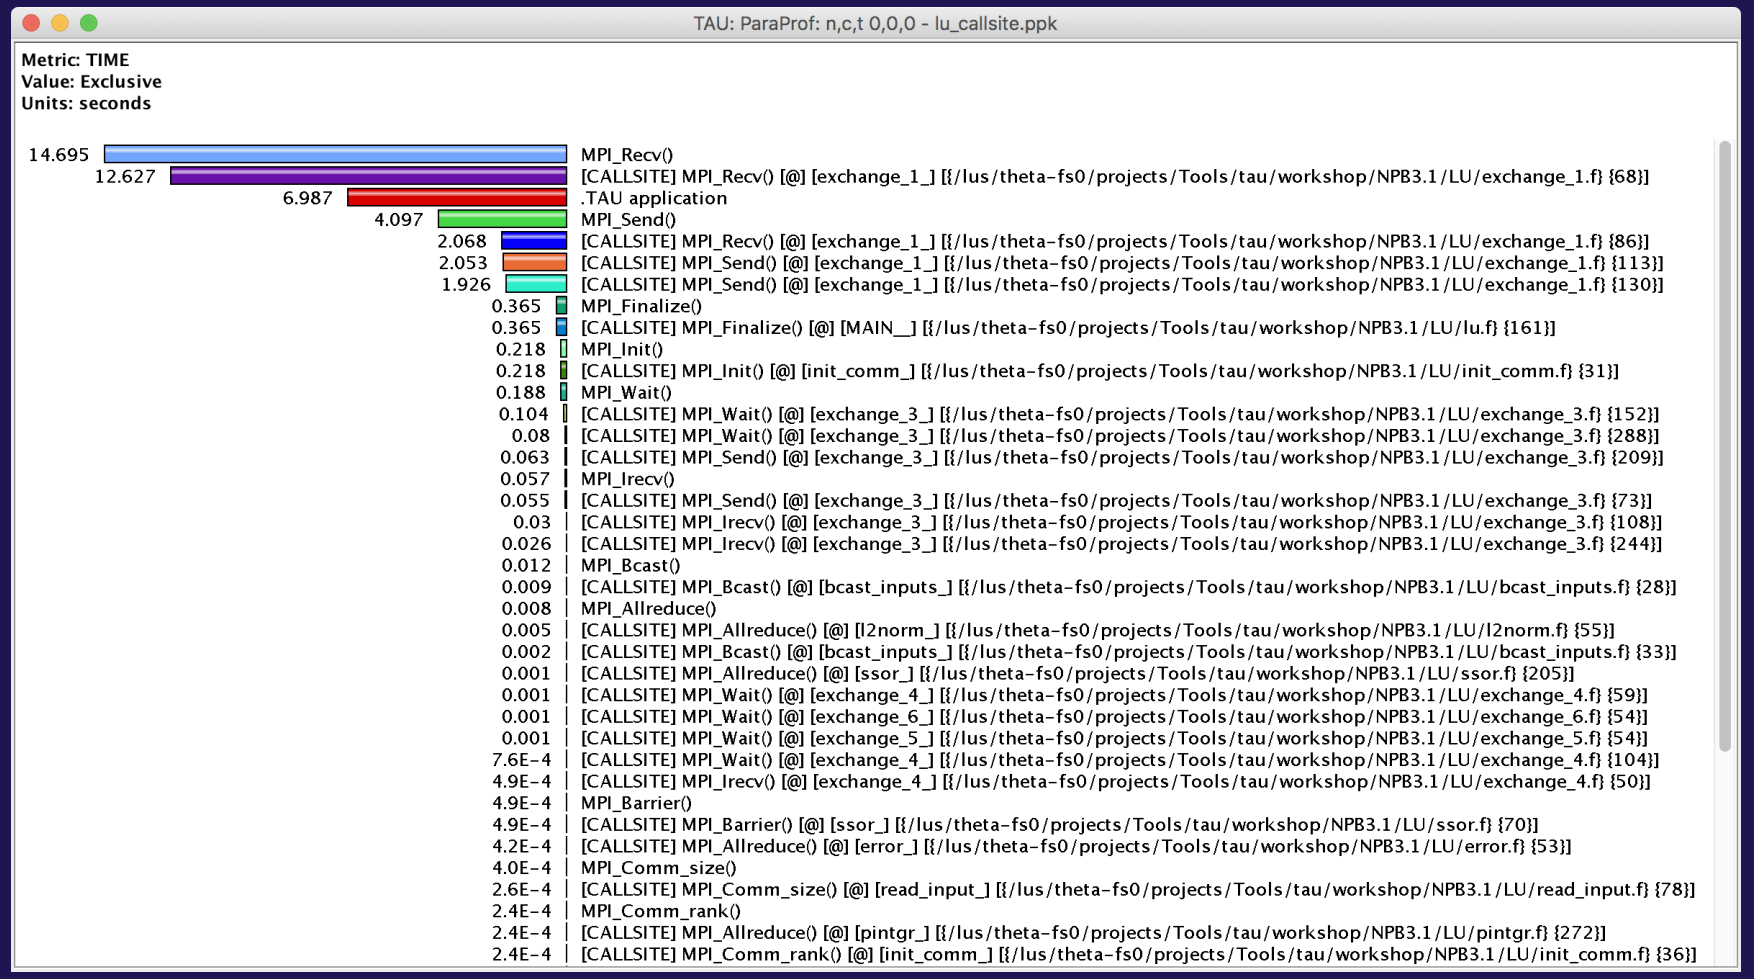

% export TAU_CALLSITE=1

Identifying Collective Wait States: Thread Callpath Relations Window

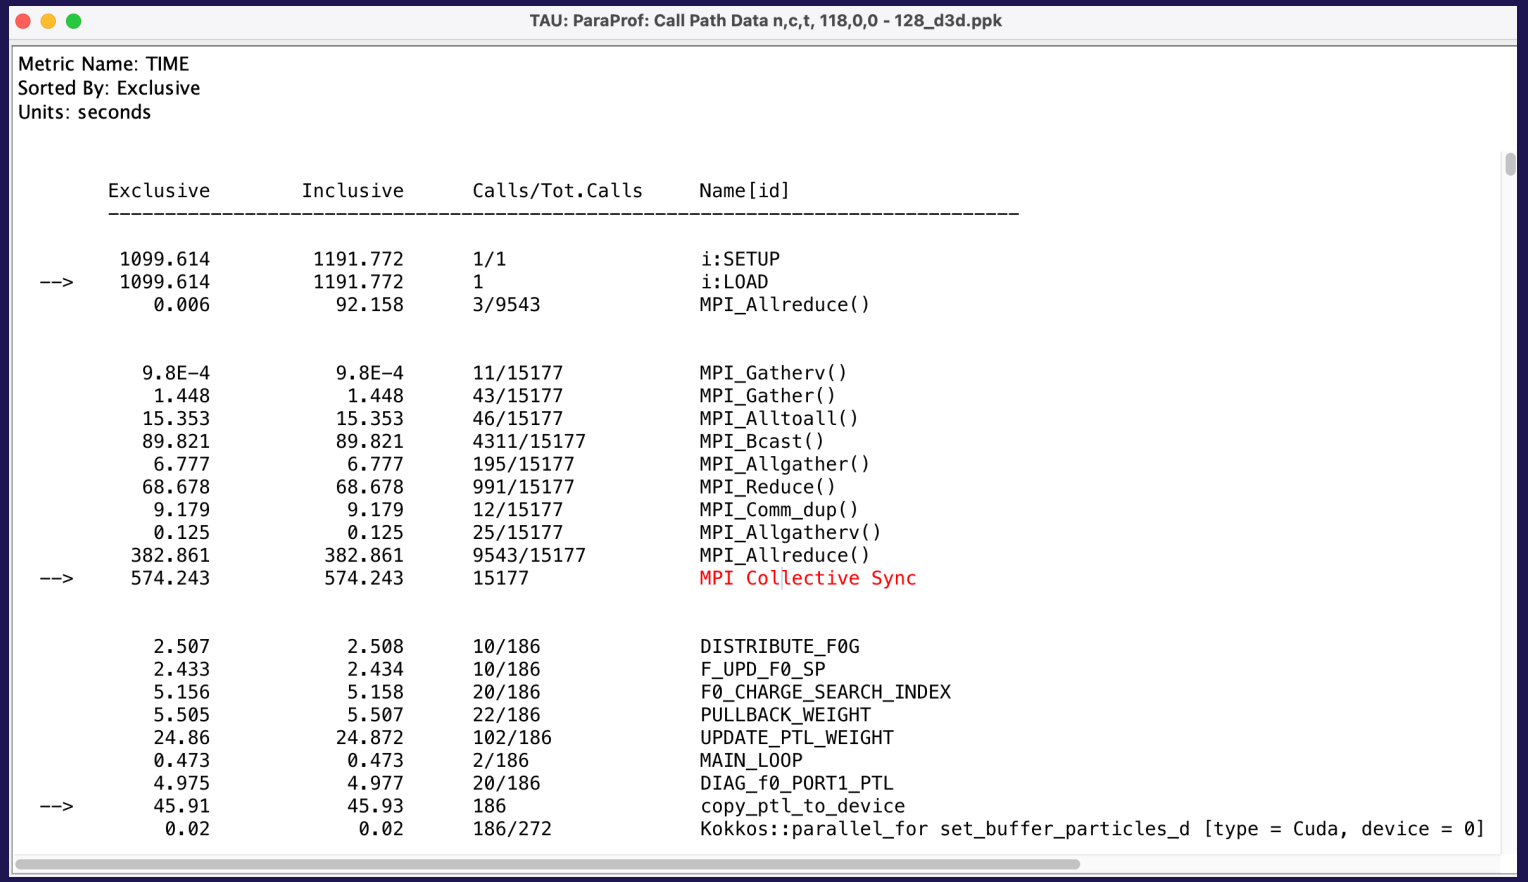

MPI Collective Sync is the time spent in a barrier operation inside a collective

ParaProf Thread Comparison Window

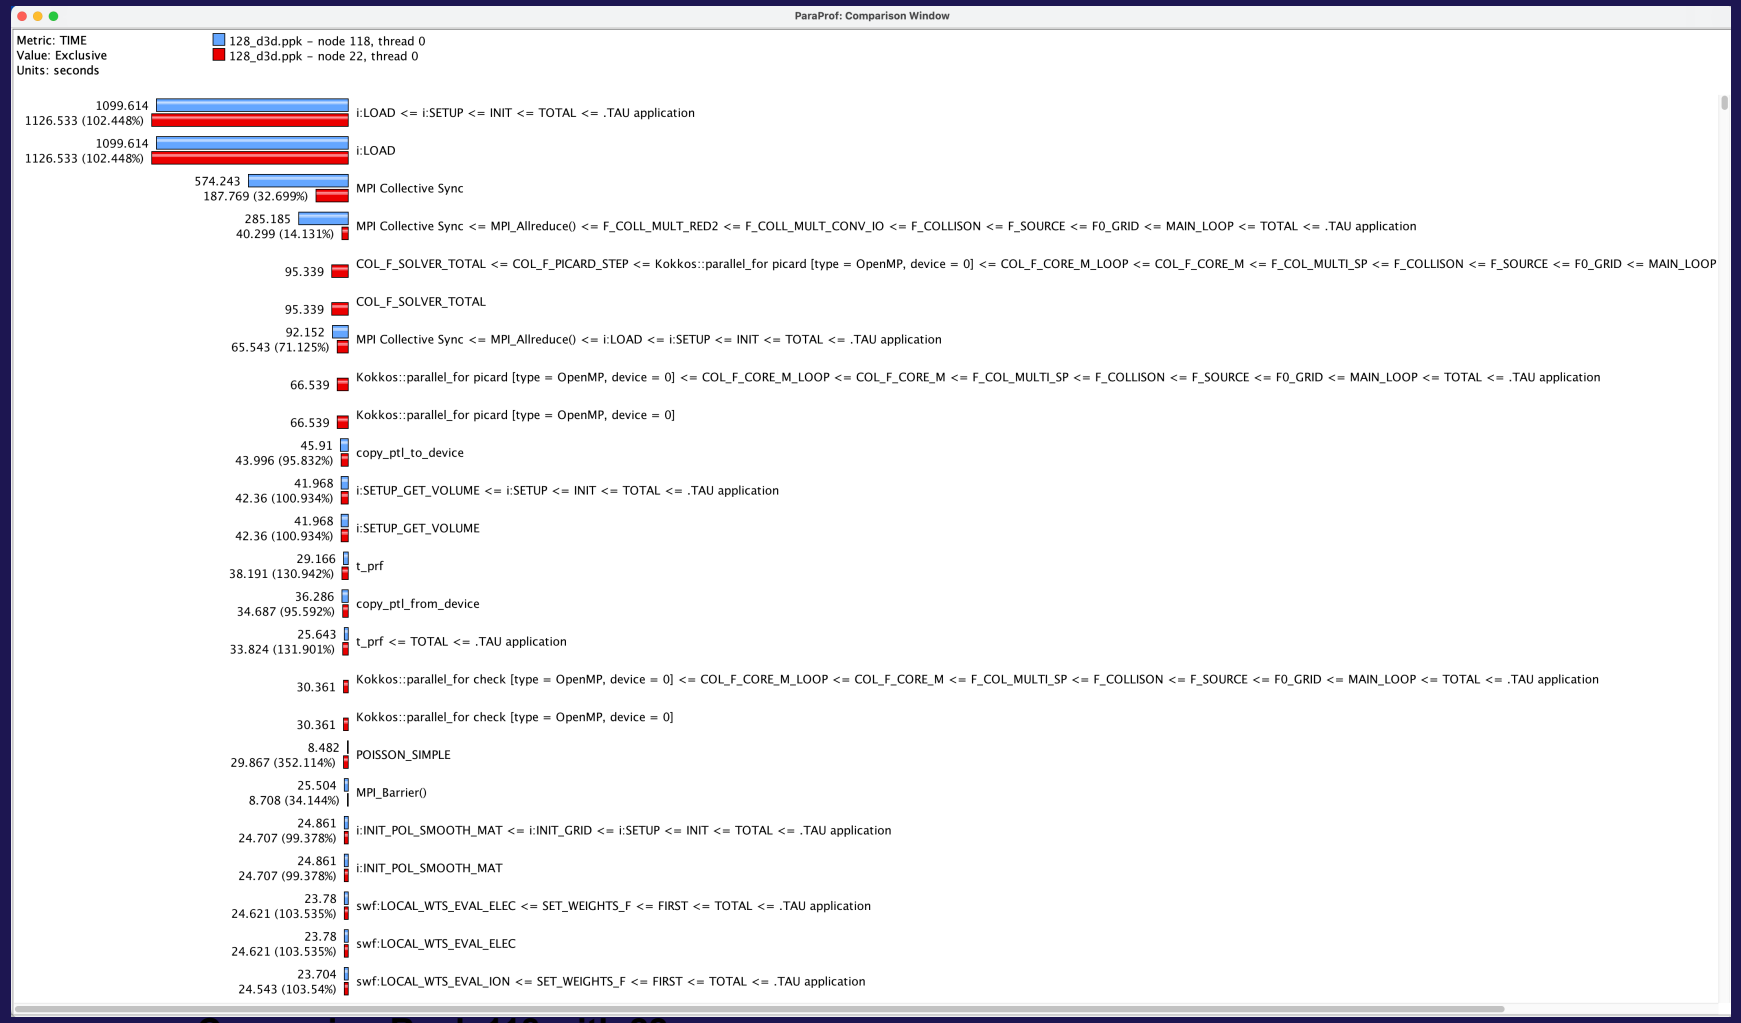

Comparing Rank 118 with 22.

Right click on “node 118” -> Add node to comparison window

TAU – Context Events

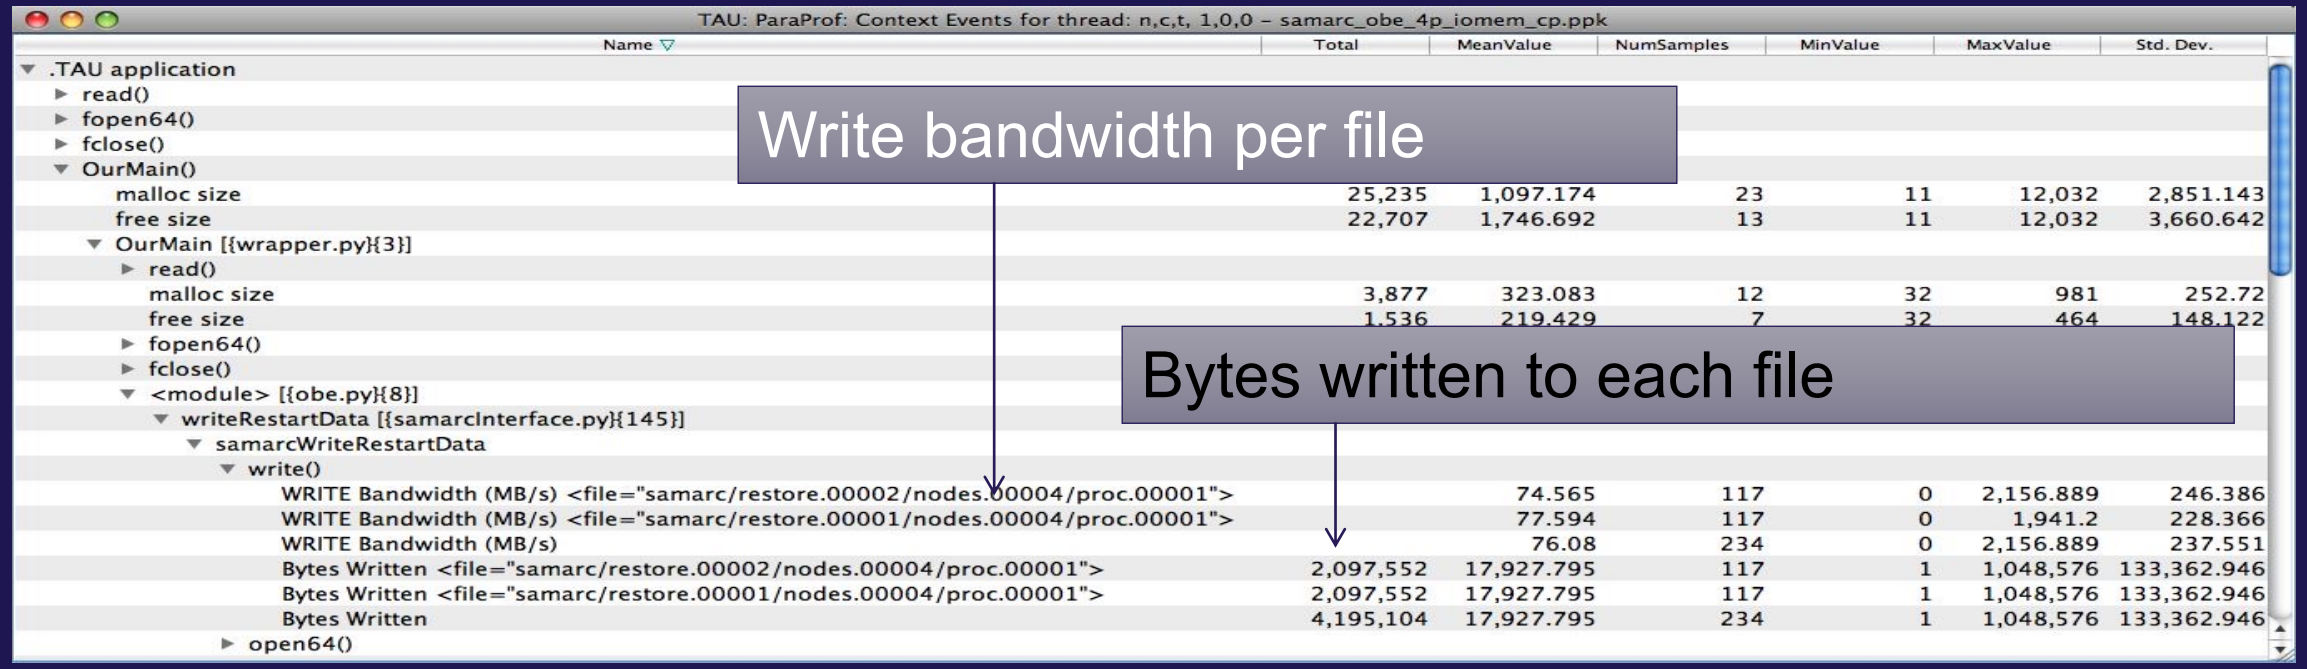

Name	Total	MeanValue	NumSamples	MinValue	MaxValue	Std. Dev.
.TAU application						
read()						
fopen64()						
fclose()						
OurMain()						
malloc size	25,235	1,097.174	23	11	12,032	2,851.143
free size	22,707	1,746.692	13	11	12,032	3,660.642
OurMain [{wrapper.py}{3}]						
read()						
malloc size	3,877	323.083	12	32	981	252.72
free size	1,536	219.429	7	32	464	148.122
fopen64()						
fclose()						
<module> [{obe.py}{8}]						
writeRestartData [{samarcInterface.py}{145}]						
samarcWriteRestartData						
write()						
WRITE Bandwidth (MB/s) <file="samarc/restore.00002/nodes.00004/proc.00001">		74.565	117	0	2,156.889	246.386
WRITE Bandwidth (MB/s) <file="samarc/restore.00001/nodes.00004/proc.00001">		77.594	117	0	1,941.2	228.366
WRITE Bandwidth (MB/s)		76.08	234	0	2,156.889	237.551
Bytes Written <file="samarc/restore.00002/nodes.00004/proc.00001">	2,097,552	17,927.795	117	1	1,048,576	133,362.946
Bytes Written <file="samarc/restore.00001/nodes.00004/proc.00001">	2,097,552	17,927.795	117	1	1,048,576	133,362.946
Bytes Written	4,195,104	17,927.795	234	1	1,048,576	133,362.946
open64()						

```
% tau_exec -io ./a.out
```

ParaProf 3D Profile Browser: Triangle Mesh

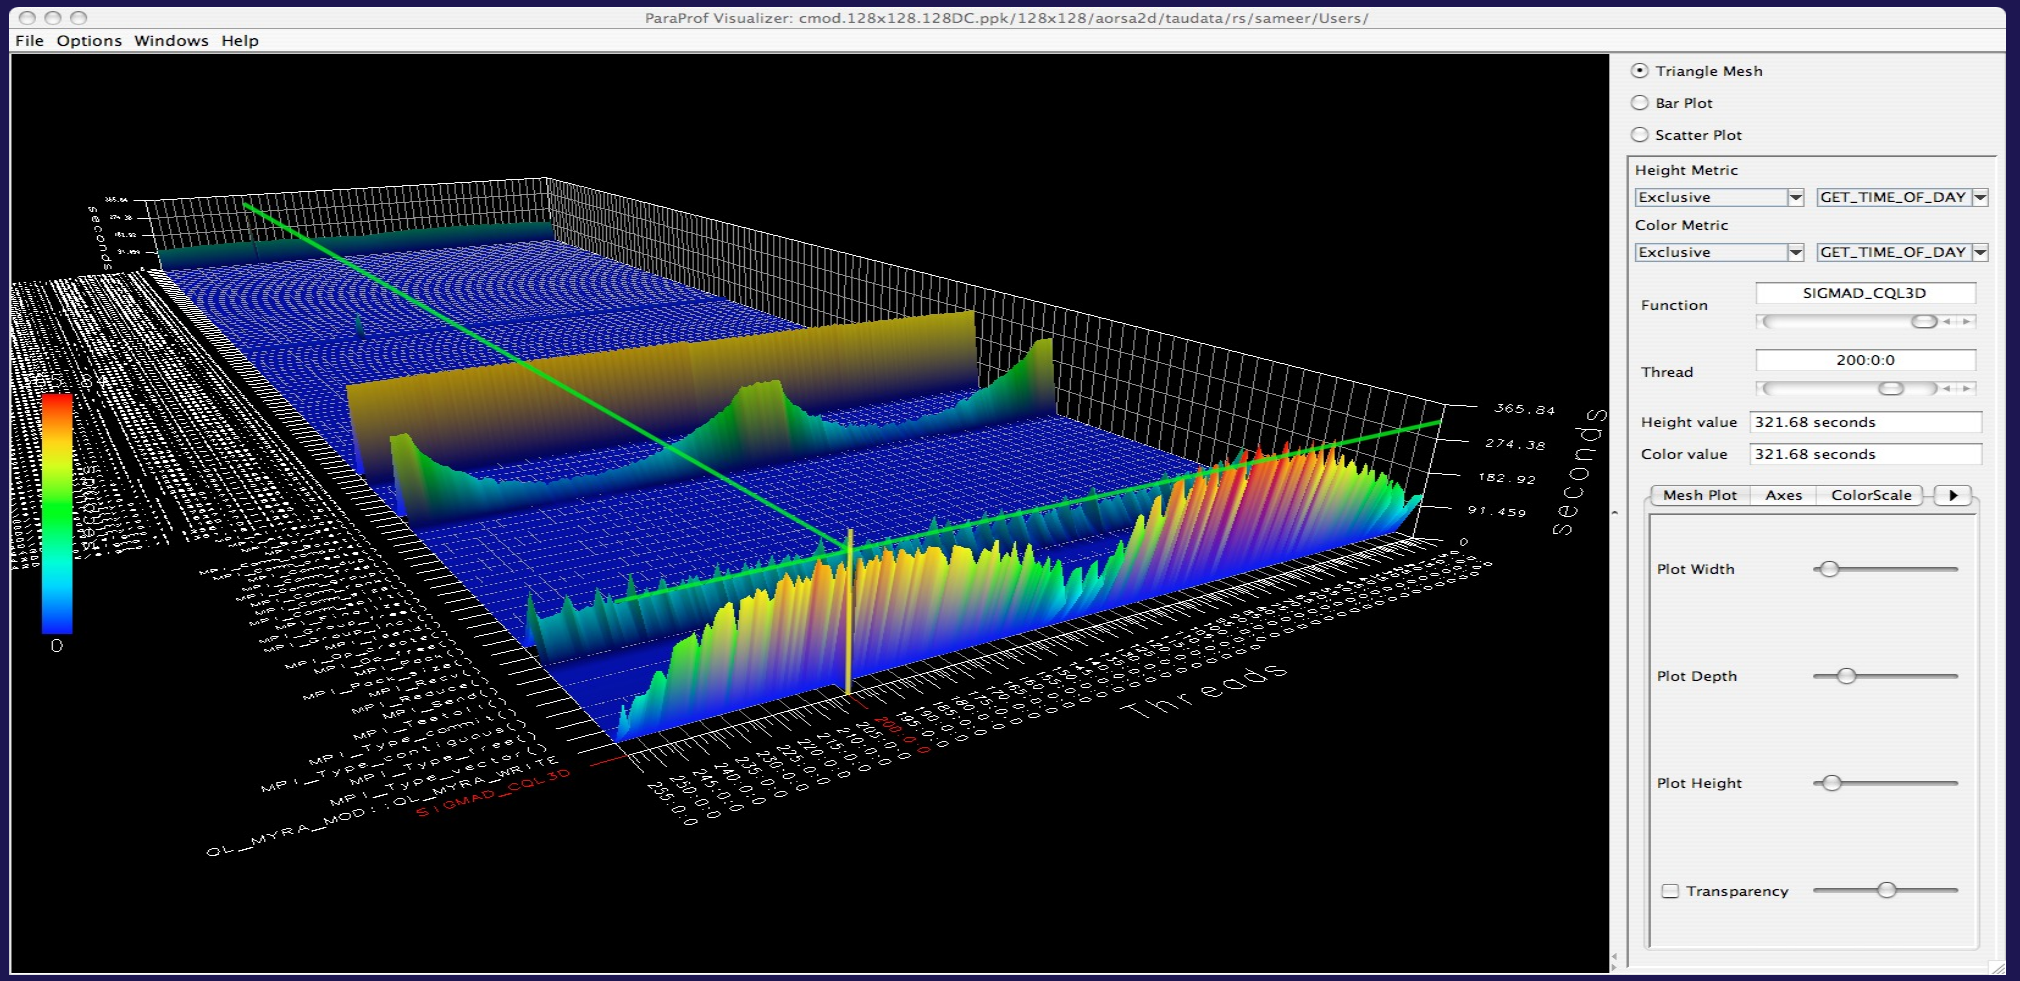

Windows -> 3D Visualization

ParaProf 3D Profile Browser: Bar Plot

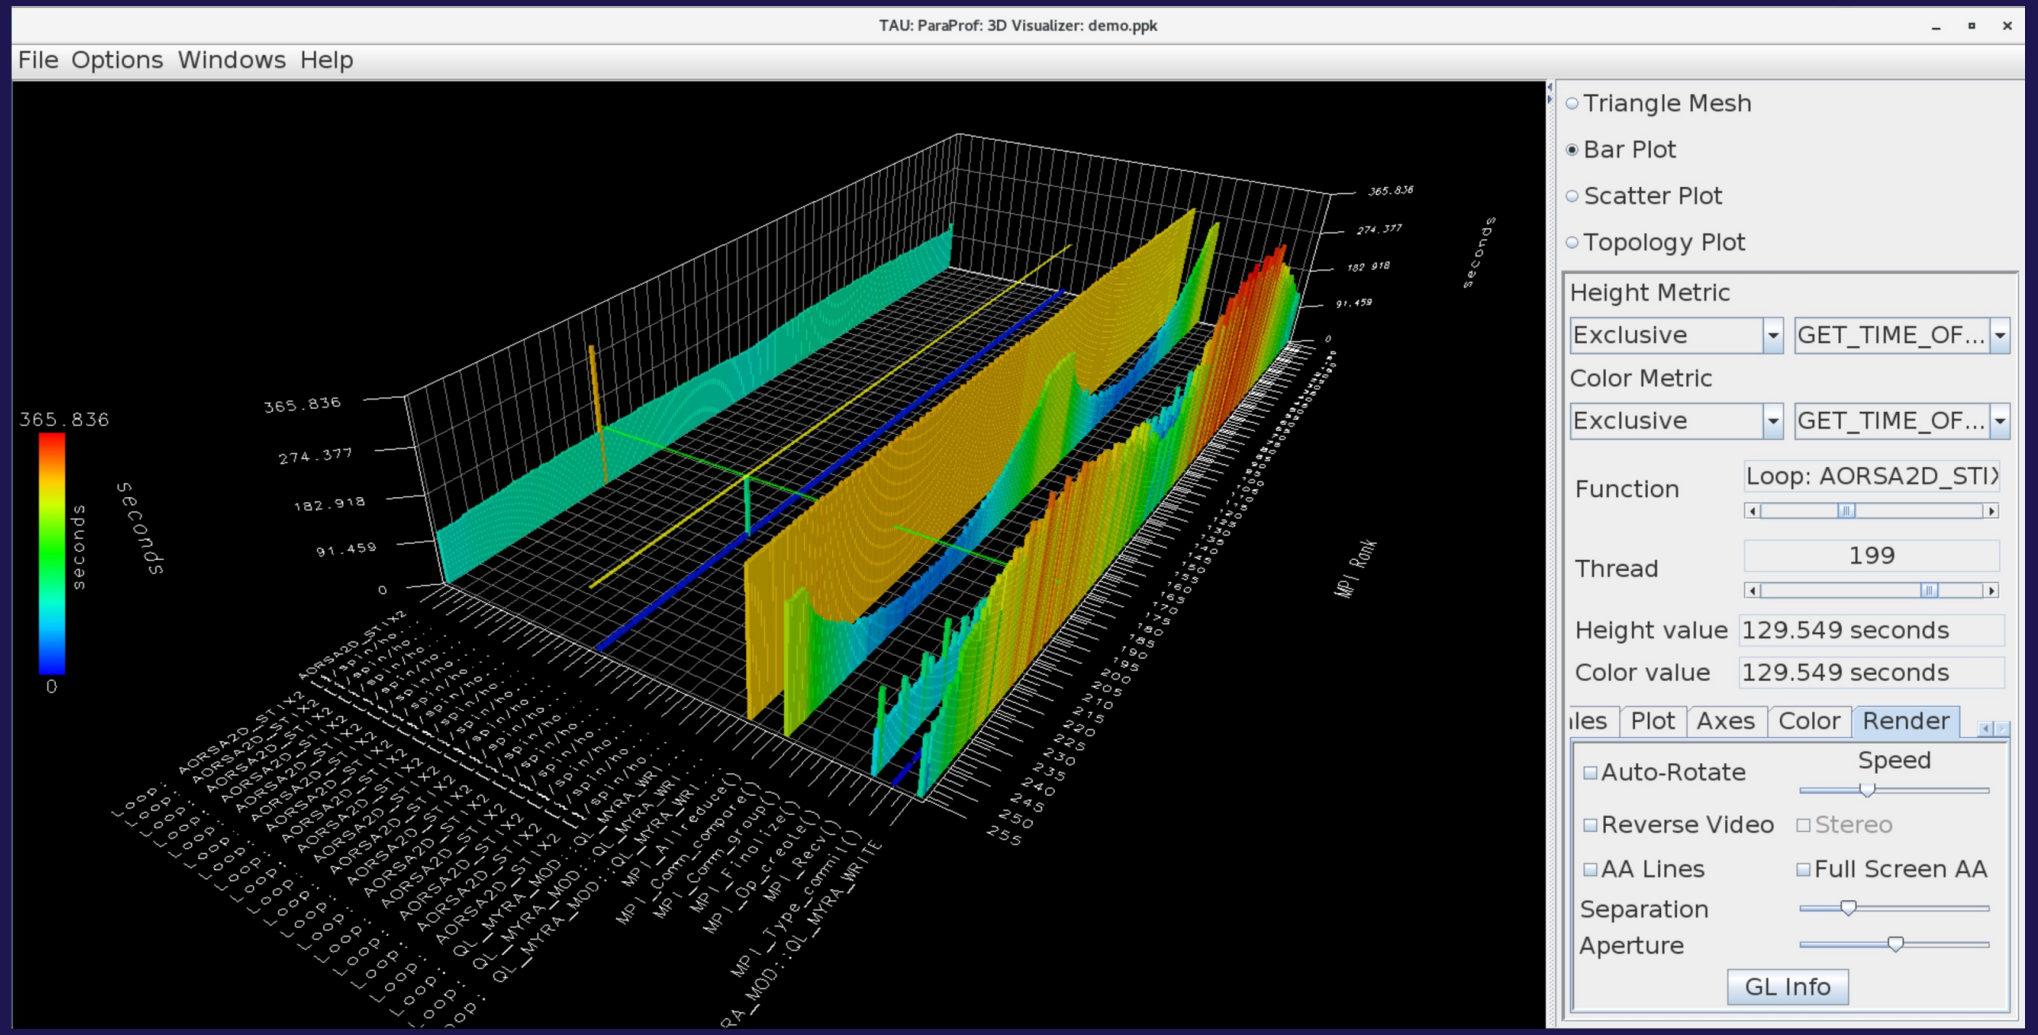

TAU – ParaProf 3D Visualization: Bar Plot using cross-hairs to zoom into a location (function, thread)

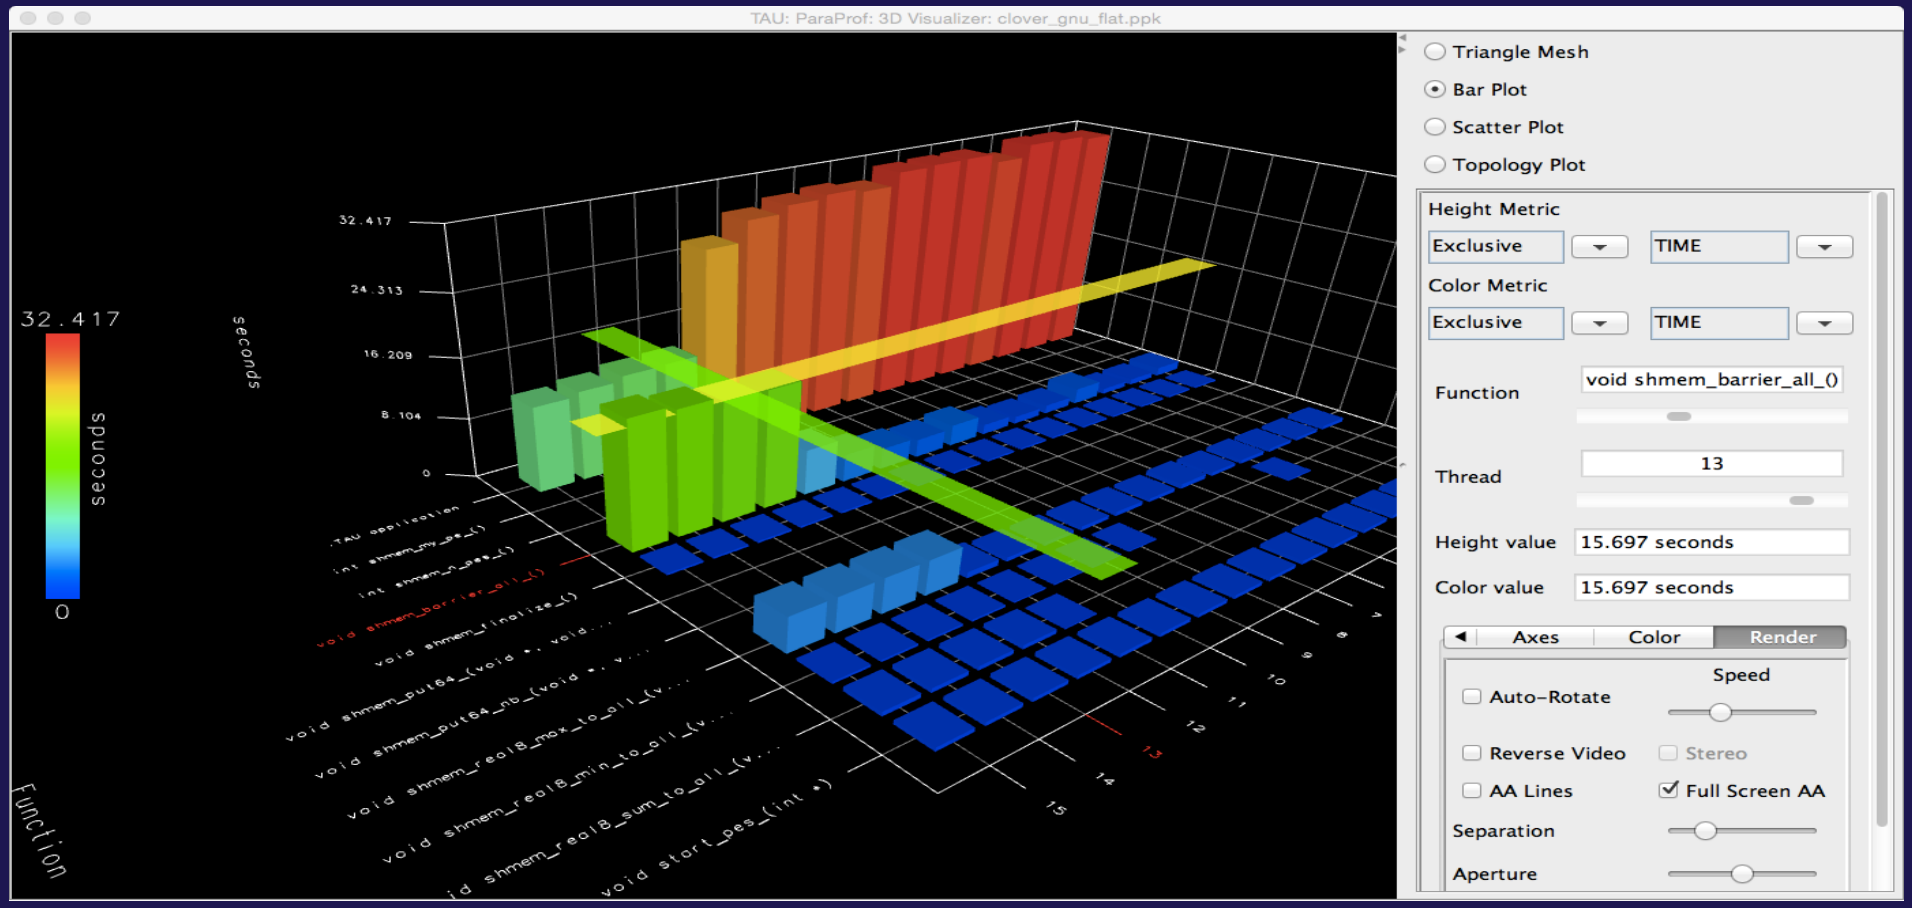

% paraprof app.ppk
Windows -> 3D Visualization -> Bar Plot (right pane)

TAU: ParaProf Topology Plot Window

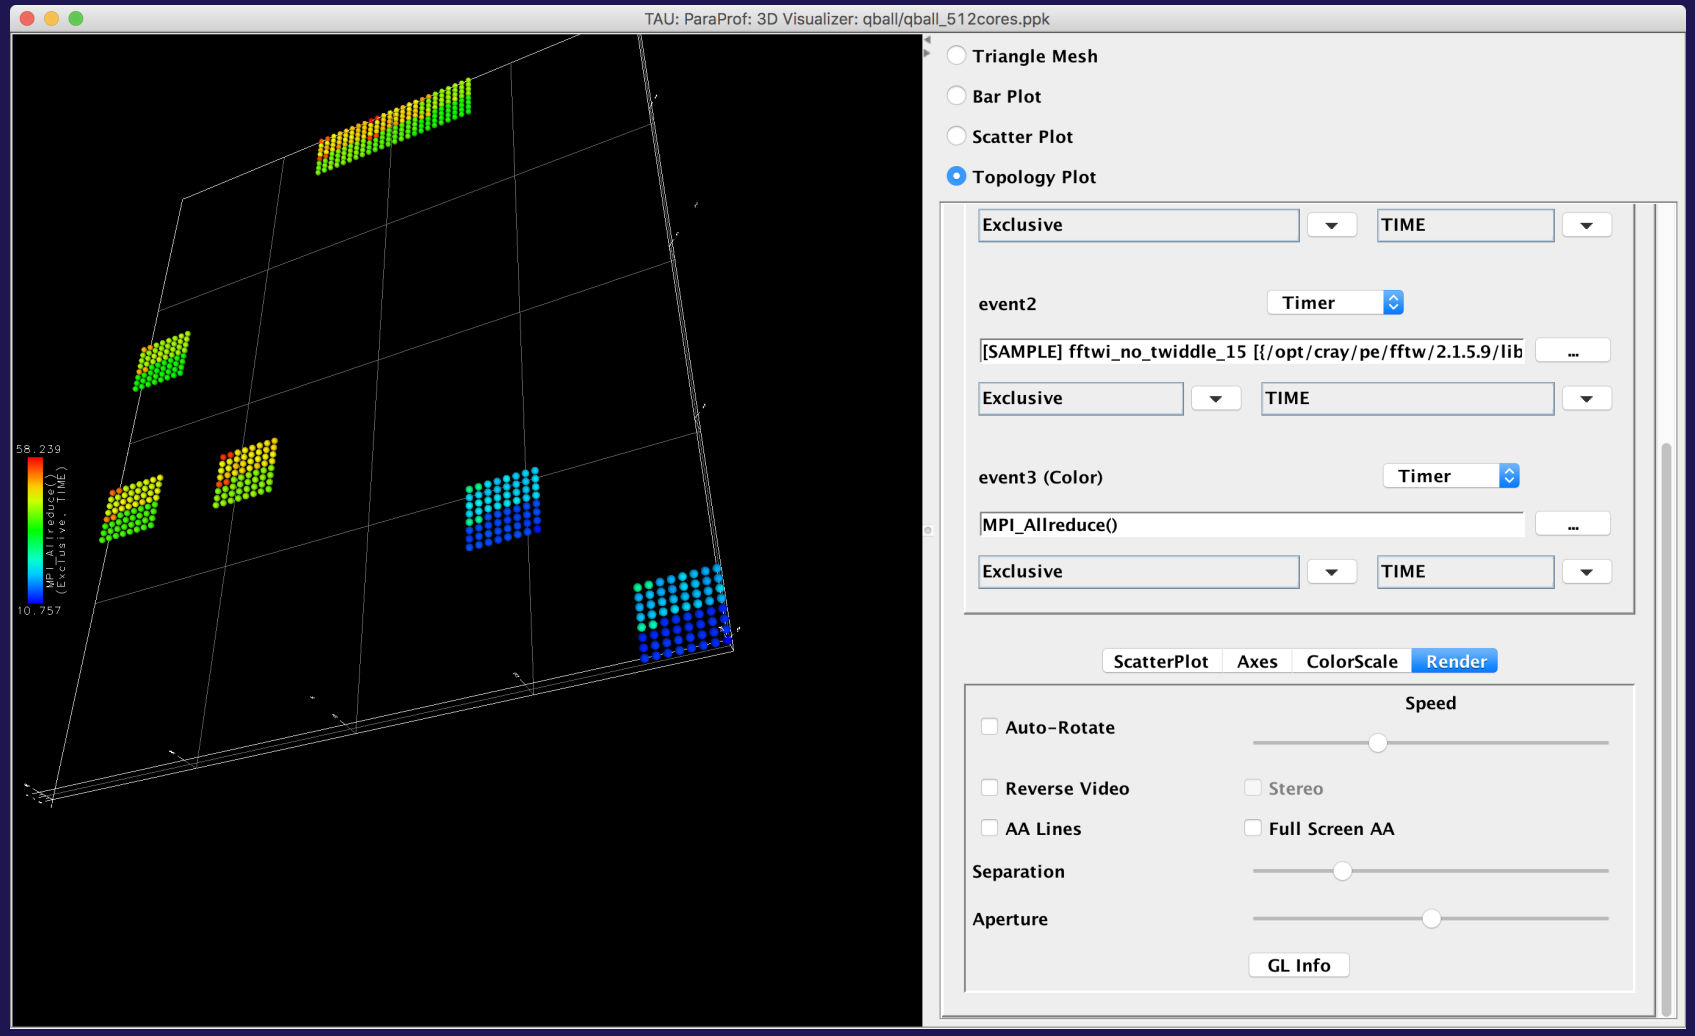

TAU – 3D Communication Window

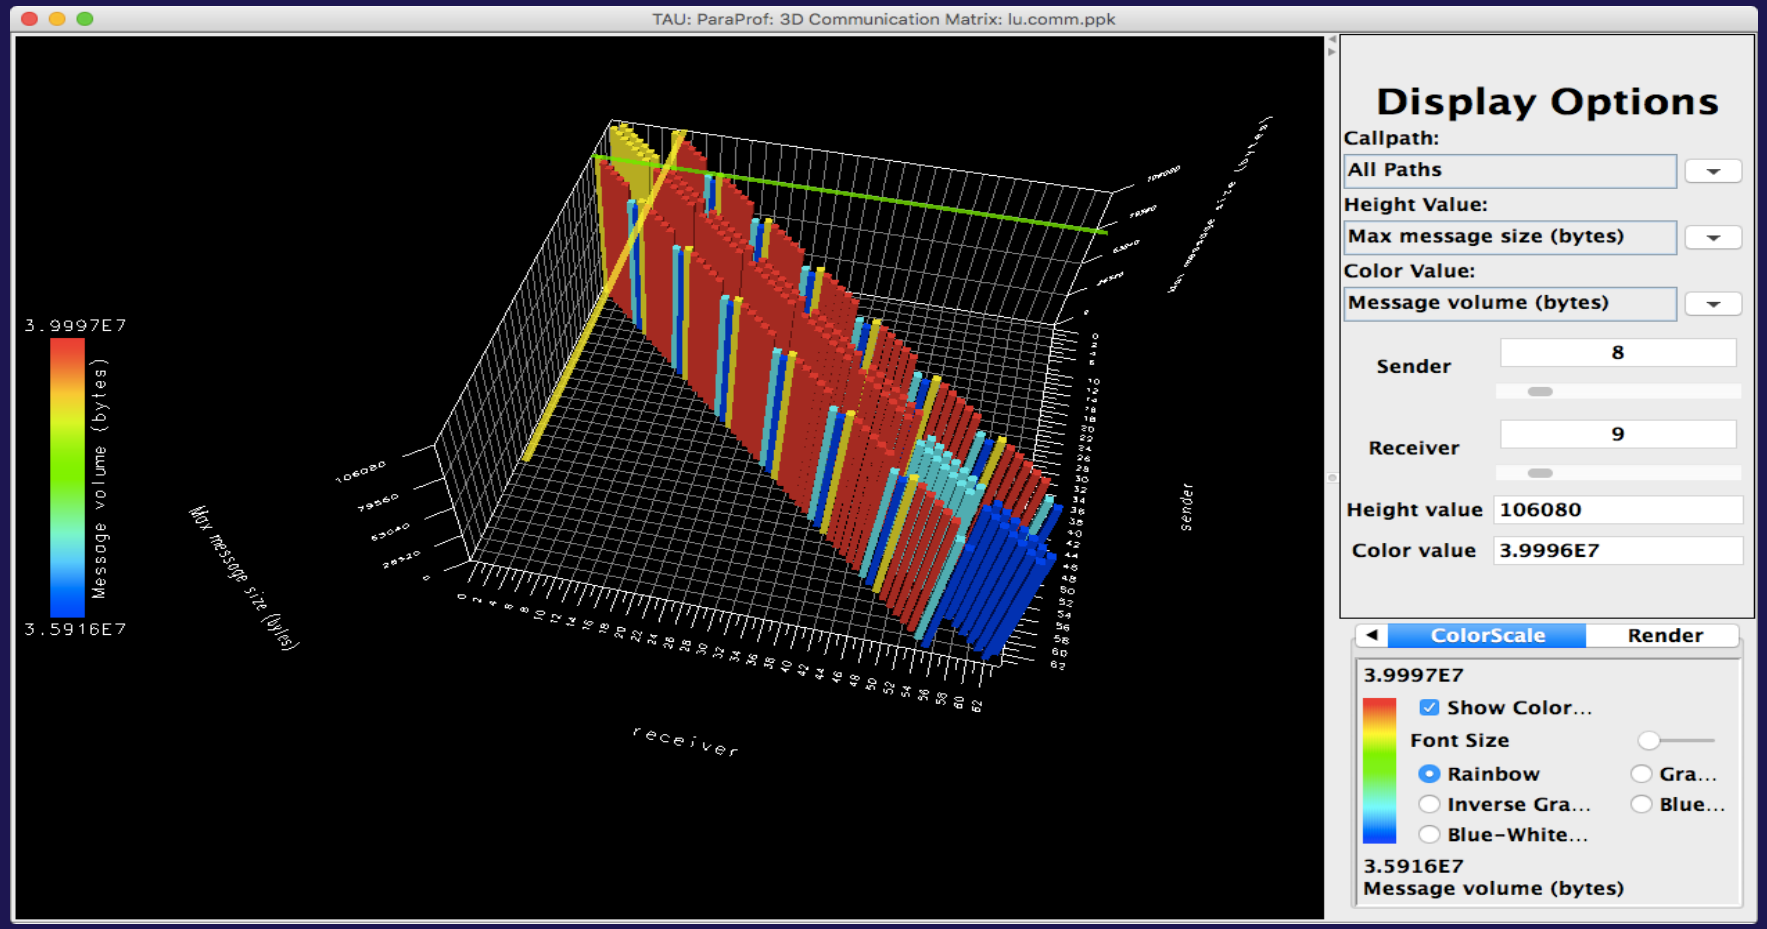

% export TAU_COMM_MATRIX=1; aprun … tau_exec ./a.out
% paraprof ; Windows -> 3D Communication Matrix

Using TAU on GPUs

TAU: Intel oneAPI DPC++ on an Intel Gen12LP or DG1 GPU

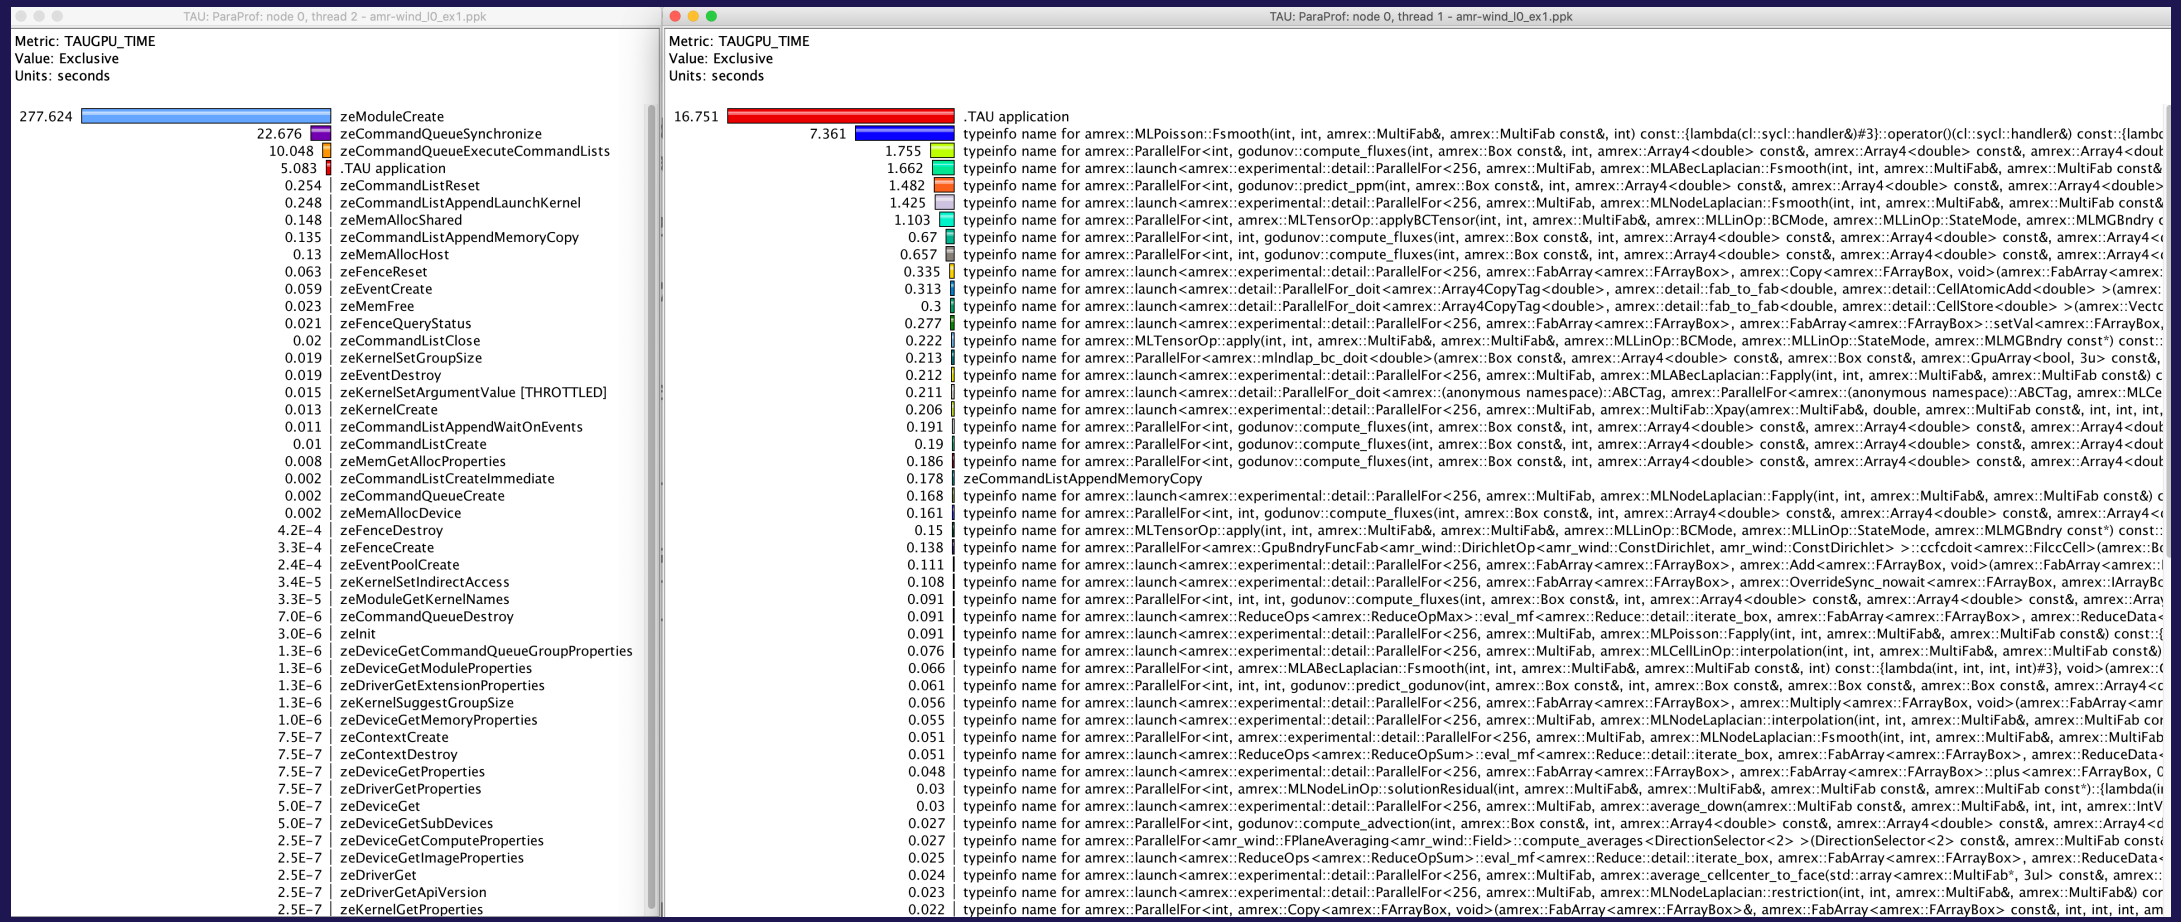

% tau_exec –T level_zero,serial –l0 ./a.out

TAU: Intel oneAPI DPC++ on an Intel Gen12LP or DG1 GPU

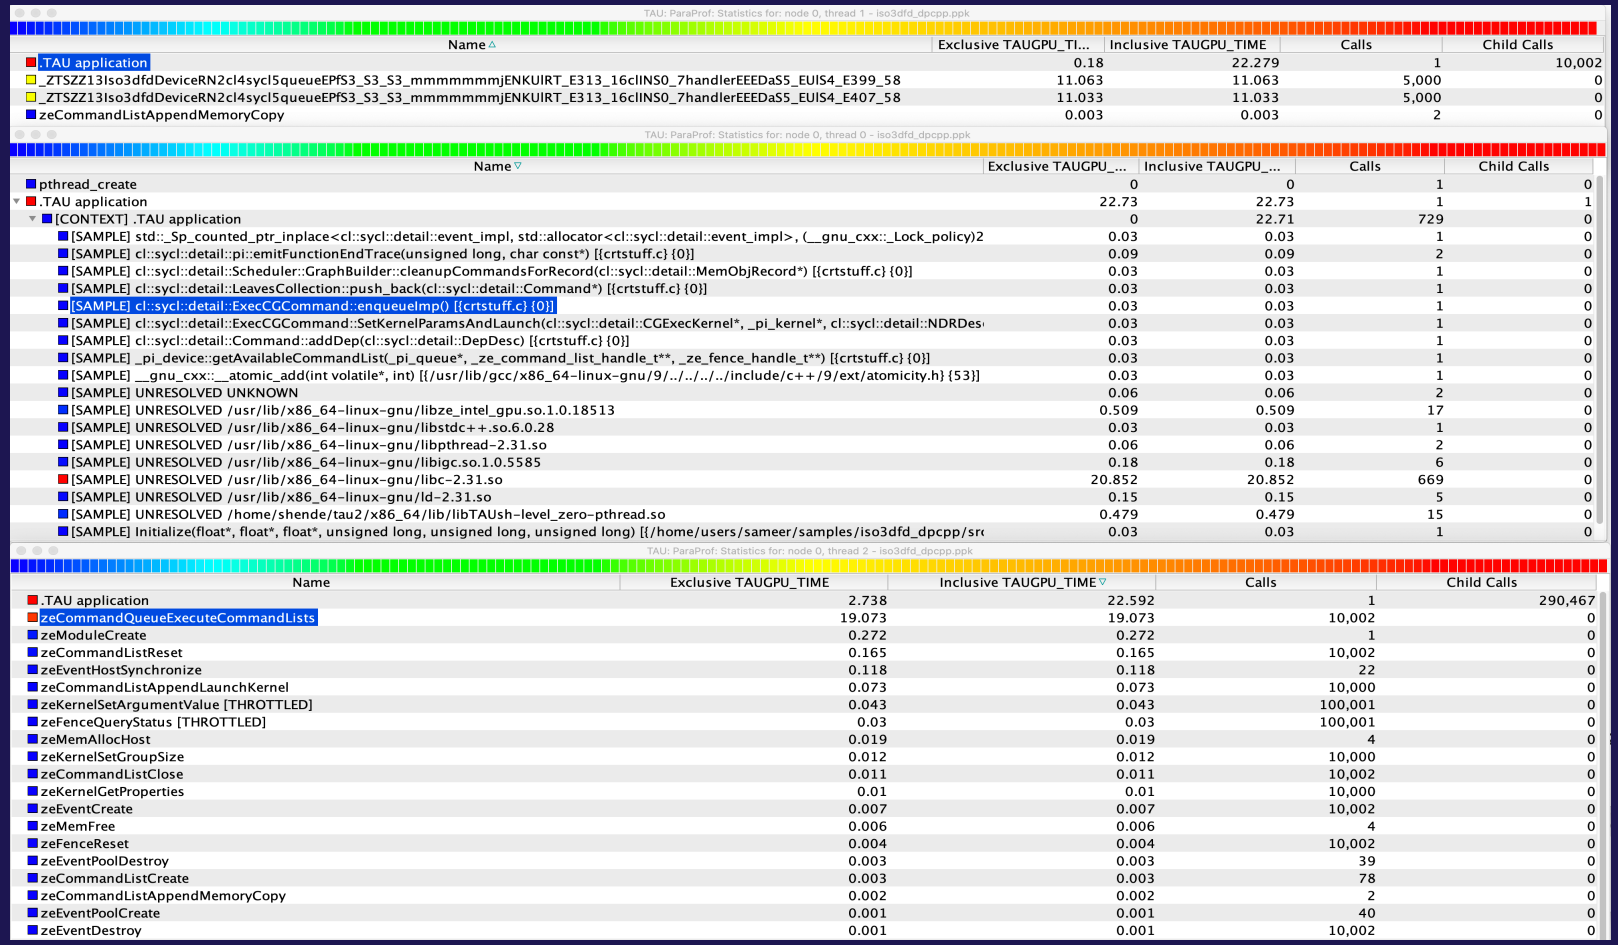

TAU: ParaProf: Statistics for: node 0, thread 1 - iso3dfd_dpcpp.ppk

Name	Exclusive TAUGPU_TI...	Inclusive TAUGPU_TIME	Calls	Child Calls
.TAU application	0.18	22.279	1	10,002
_ZTSZZ13Iso3dfdDeviceRN2cl4sycl5queueEPfS3_S3_S3_mmmmmmmjENKUlRT_E313_16clINS0_7handlerEEEDaS5_EUlS4_E399_58	11.063	11.063	5,000	0
_ZTSZZ13Iso3dfdDeviceRN2cl4sycl5queueEPfS3_S3_S3_mmmmmmmjENKUlRT_E313_16clINS0_7handlerEEEDaS5_EUlS4_E407_58	11.033	11.033	5,000	0
zeCommandListAppendMemoryCopy	0.003	0.003	2	0

TAU: ParaProf: Statistics for: node 0, thread 0 - iso3dfd_dpcpp.ppk

Name	Exclusive TAUGPU_...	Inclusive TAUGPU_...	Calls	Child Calls
pthread_create	0	0	1	0
.TAU application	22.73	22.73	1	1
[CONTEXT] .TAU application	0	22.71	729	0
[SAMPLE] std::_Sp_counted_ptr_inplace<cl::sycl::detail::event_impl, std::allocator<cl::sycl::detail::event_impl>, (__gnu_cxx::_Lock_policy)2	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::pi::emitFunctionEndTrace(unsigned long, char const*) [{crtstuff.c} {0}]	0.09	0.09	2	0
[SAMPLE] cl::sycl::detail::Scheduler::GraphBuilder::cleanupCommandsForRecord(cl::sycl::detail::MemObjRecord*) [{crtstuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::LeavesCollection::push_back(cl::sycl::detail::Command*) [{crtstuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::ExecCGCommand::enqueueImp() [{crtstuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::ExecCGCommand::SetKernelParamsAndLaunch(cl::sycl::detail::CGExecKernel*, _pi_kernel*, cl::sycl::detail::NDRDes	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::Command::addDep(cl::sycl::detail::DepDesc) [{crtstuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] _pi_device::getAvailableCommandList(_pi_queue*, _ze_command_list_handle_t**, _ze_fence_handle_t**) [{crtstuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] __gnu_cxx::__atomic_add(int volatile*, int) [{/usr/lib/gcc/x86_64-linux-gnu/9/../../../../include/c++/9/ext/atomicity.h} {53}]	0.03	0.03	1	0
[SAMPLE] UNRESOLVED UNKNOWN	0.06	0.06	2	0
[SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/libze_intel_gpu.so.1.0.18513	0.509	0.509	17	0
[SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/libstdc++.so.6.0.28	0.03	0.03	1	0
[SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/libpthread-2.31.so	0.06	0.06	2	0
[SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/libigc.so.1.0.5585	0.18	0.18	6	0
[SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/libc-2.31.so	20.852	20.852	669	0
[SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/ld-2.31.so	0.15	0.15	5	0
[SAMPLE] UNRESOLVED /home/shende/tau2/x86_64/lib/libTAUsh-level_zero-pthread.so	0.479	0.479	15	0
[SAMPLE] Initialize(float*, float*, float*, unsigned long, unsigned long, unsigned long) [{/home/users/sameer/samples/iso3dfd_dpcpp/sr	0.03	0.03	1	0

TAU: ParaProf: Statistics for: node 0, thread 2 - iso3dfd_dpcpp.ppk

Name	Exclusive TAUGPU_TIME	Inclusive TAUGPU_TIME	Calls	Child Calls
.TAU application	2.738	22.592	1	290,467
zeCommandQueueExecuteCommandLists	19.073	19.073	10,002	0
zeModuleCreate	0.272	0.272	1	0
zeCommandListReset	0.165	0.165	10,002	0
zeEventHostSynchronize	0.118	0.118	22	0
zeCommandListAppendLaunchKernel	0.073	0.073	10,000	0
zeKernelSetArgumentValue [THROTTLED]	0.043	0.043	100,001	0
zeFenceQueryStatus [THROTTLED]	0.03	0.03	100,001	0
zeMemAllocHost	0.019	0.019	4	0
zeKernelSetGroupSize	0.012	0.012	10,000	0
zeCommandListClose	0.011	0.011	10,002	0
zeKernelGetProperties	0.01	0.01	10,000	0
zeEventCreate	0.007	0.007	10,002	0
zeMemFree	0.006	0.006	4	0
zeFenceReset	0.004	0.004	10,002	0
zeEventPoolDestroy	0.003	0.003	39	0
zeCommandListCreate	0.003	0.003	78	0
zeCommandListAppendMemoryCopy	0.002	0.002	2	0
zeEventPoolCreate	0.001	0.001	40	0
zeEventDestroy	0.001	0.001	10,002	0

% tau_exec –T level_zero,serial –l0 ./a.out

Intel Level Zero (TigerLake Gen12LP integrated CPUs or DG1)

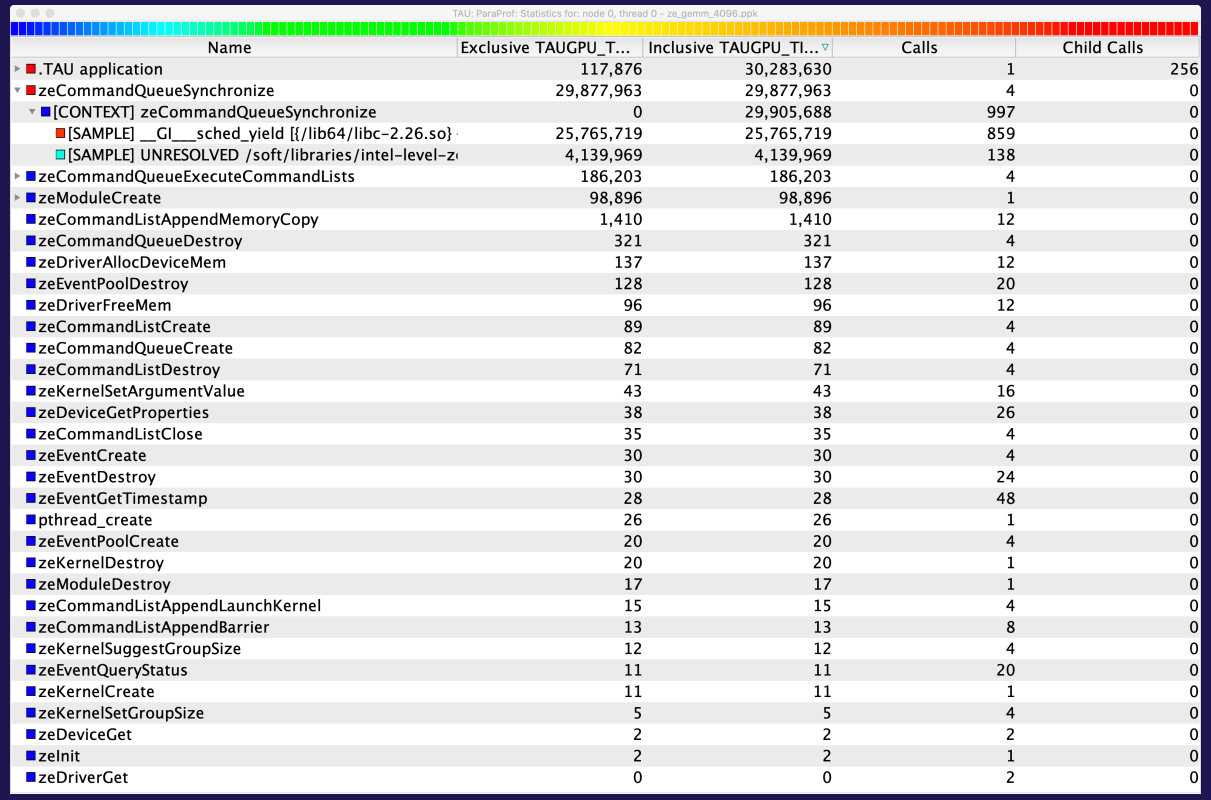

Name	Exclusive TAUGPU_T...	Inclusive TAUGPU_TI...	Calls	Child Calls
.TAU application	117,876	30,283,630	1	256
zeCommandQueueSynchronize	29,877,963	29,877,963	4	0
[CONTEXT] zeCommandQueueSynchronize	0	29,905,688	997	0
[SAMPLE] __GI___sched_yield [{/lib64/libc-2.26.so}	25,765,719	25,765,719	859	0
[SAMPLE] UNRESOLVED /soft/libraries/intel-level-z	4,139,969	4,139,969	138	0
zeCommandQueueExecuteCommandLists	186,203	186,203	4	0
zeModuleCreate	98,896	98,896	1	0
zeCommandListAppendMemoryCopy	1,410	1,410	12	0
zeCommandQueueDestroy	321	321	4	0
zeDriverAllocDeviceMem	137	137	12	0
zeEventPoolDestroy	128	128	20	0
zeDriverFreeMem	96	96	12	0
zeCommandListCreate	89	89	4	0
zeCommandQueueCreate	82	82	4	0
zeCommandListDestroy	71	71	4	0
zeKernelSetArgumentValue	43	43	16	0
zeDeviceGetProperties	38	38	26	0
zeCommandListClose	35	35	4	0
zeEventCreate	30	30	4	0
zeEventDestroy	30	30	24	0
zeEventGetTimestamp	28	28	48	0
pthread_create	26	26	1	0
zeEventPoolCreate	20	20	4	0
zeKernelDestroy	20	20	1	0
zeModuleDestroy	17	17	1	0
zeCommandListAppendLaunchKernel	15	15	4	0
zeCommandListAppendBarrier	13	13	8	0
zeKernelSuggestGroupSize	12	12	4	0
zeEventQueryStatus	11	11	20	0
zeKernelCreate	11	11	1	0
zeKernelSetGroupSize	5	5	4	0
zeDeviceGet	2	2	2	0
zeInit	2	2	1	0
zeDriverGet	0	0	2	0

Units: microseconds

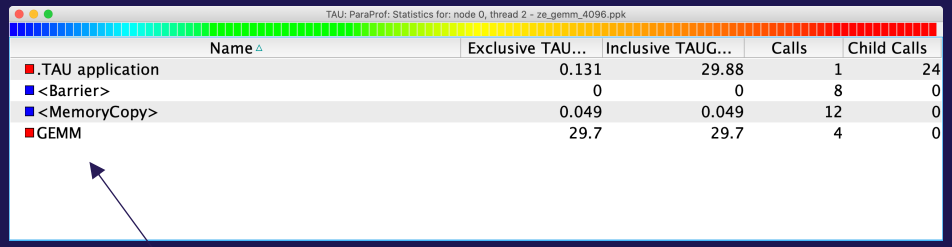

Name	Exclusive TAU...	Inclusive TAUG...	Calls	Child Calls
.TAU application	0.131	29.88	1	24
<Barrier>	0	0	8	0
<MemoryCopy>	0.049	0.049	12	0
GEMM	29.7	29.7	4	0

Time spent in GEMM kernel

Units: seconds

% mpirun –np 64 tau_exec –l0 ./a.out

TAU and Vampir [TU Dresden]: Intel oneAPI with MPI

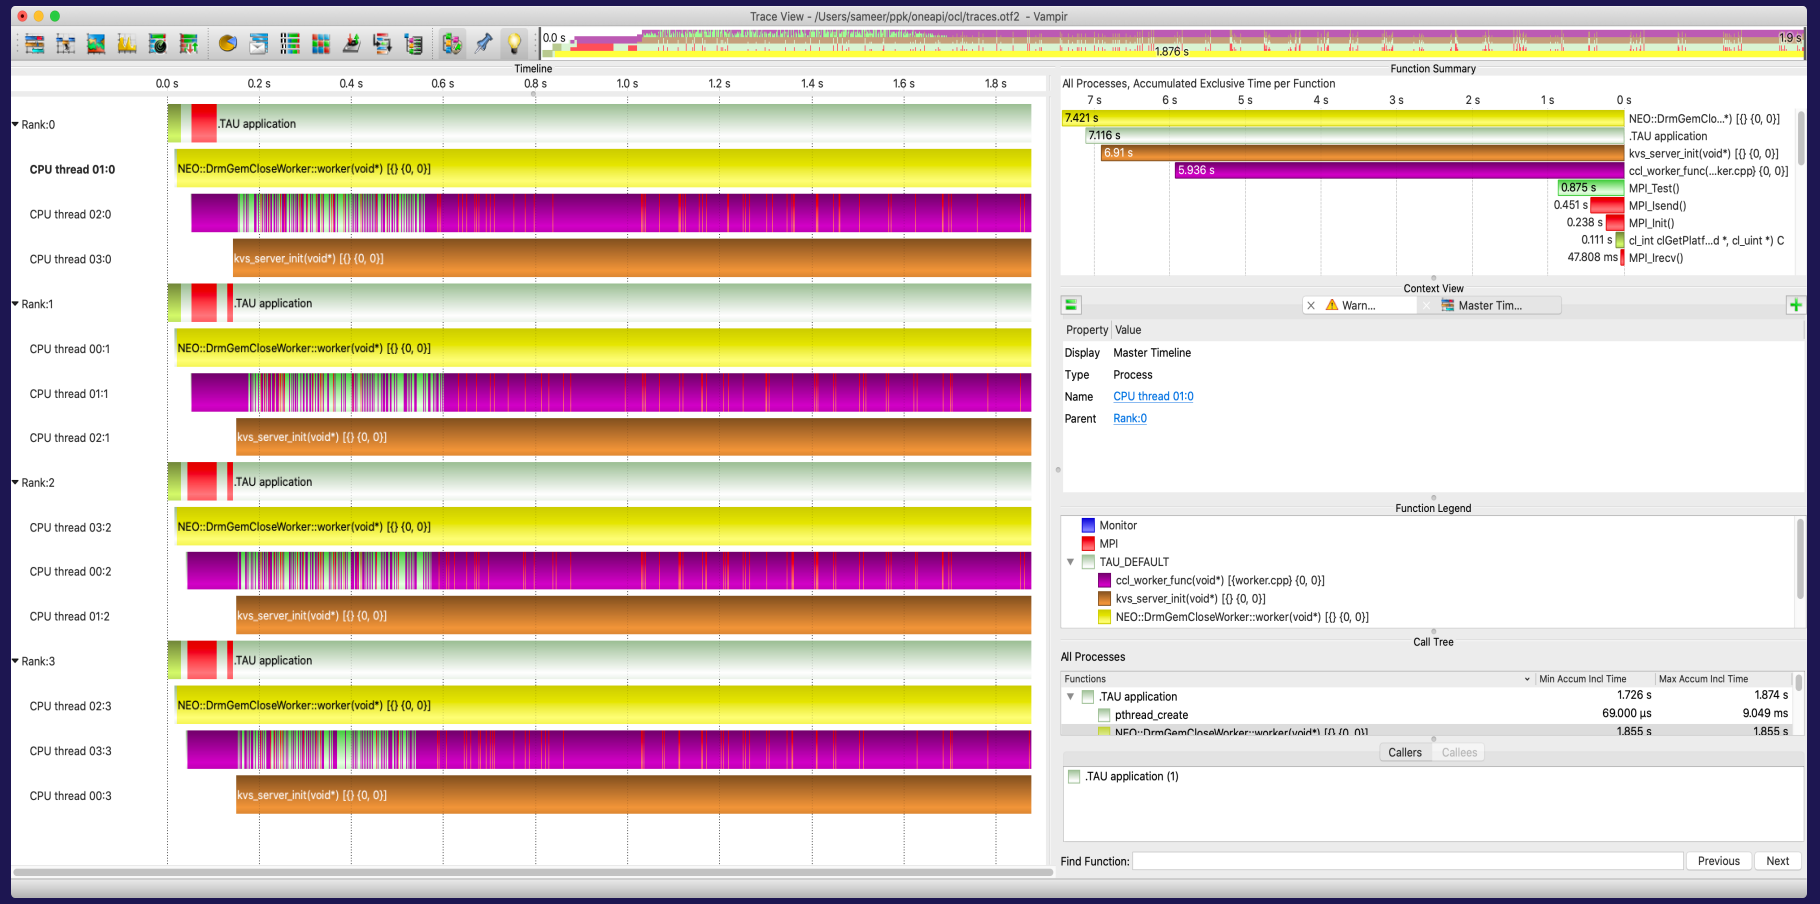


```
% export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2
% mpirun –np 4  tau_exec –T level_zero –opencl ./a.out
```

AMD GPU Tracing support uses RocTracer

```
$ TAU_TRACE=1 TAU_TRACE_FORMAT=otf2 tau_exec -T serial,roctracer ./gpu-stream-hip
```

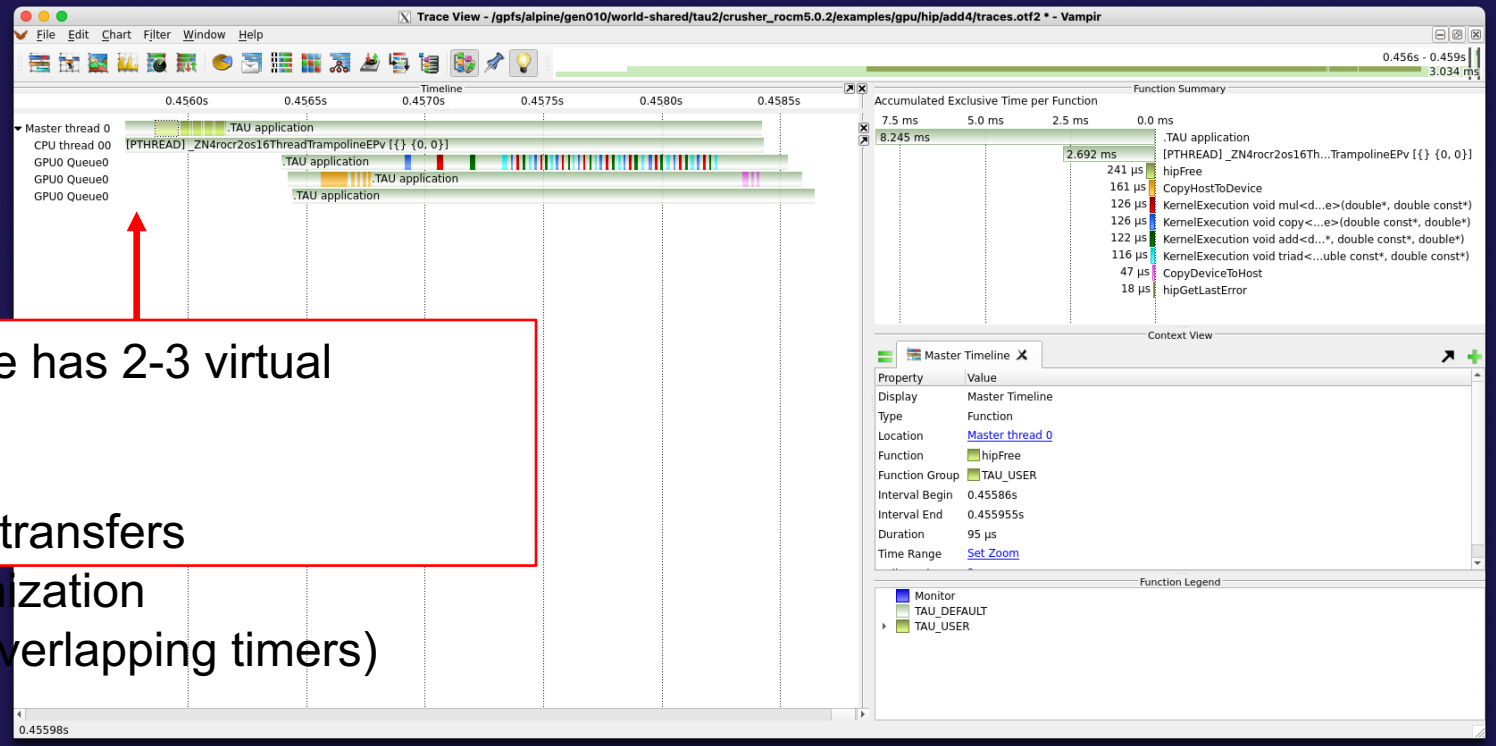

Each device has 2-3 virtual threads:
1) kernels,
2) memory transfers
3) synchronization
(prevents overlapping timers)

TAU output shown in Vampir

Tracing: Jumpshot [ANL] (ships with TAU)

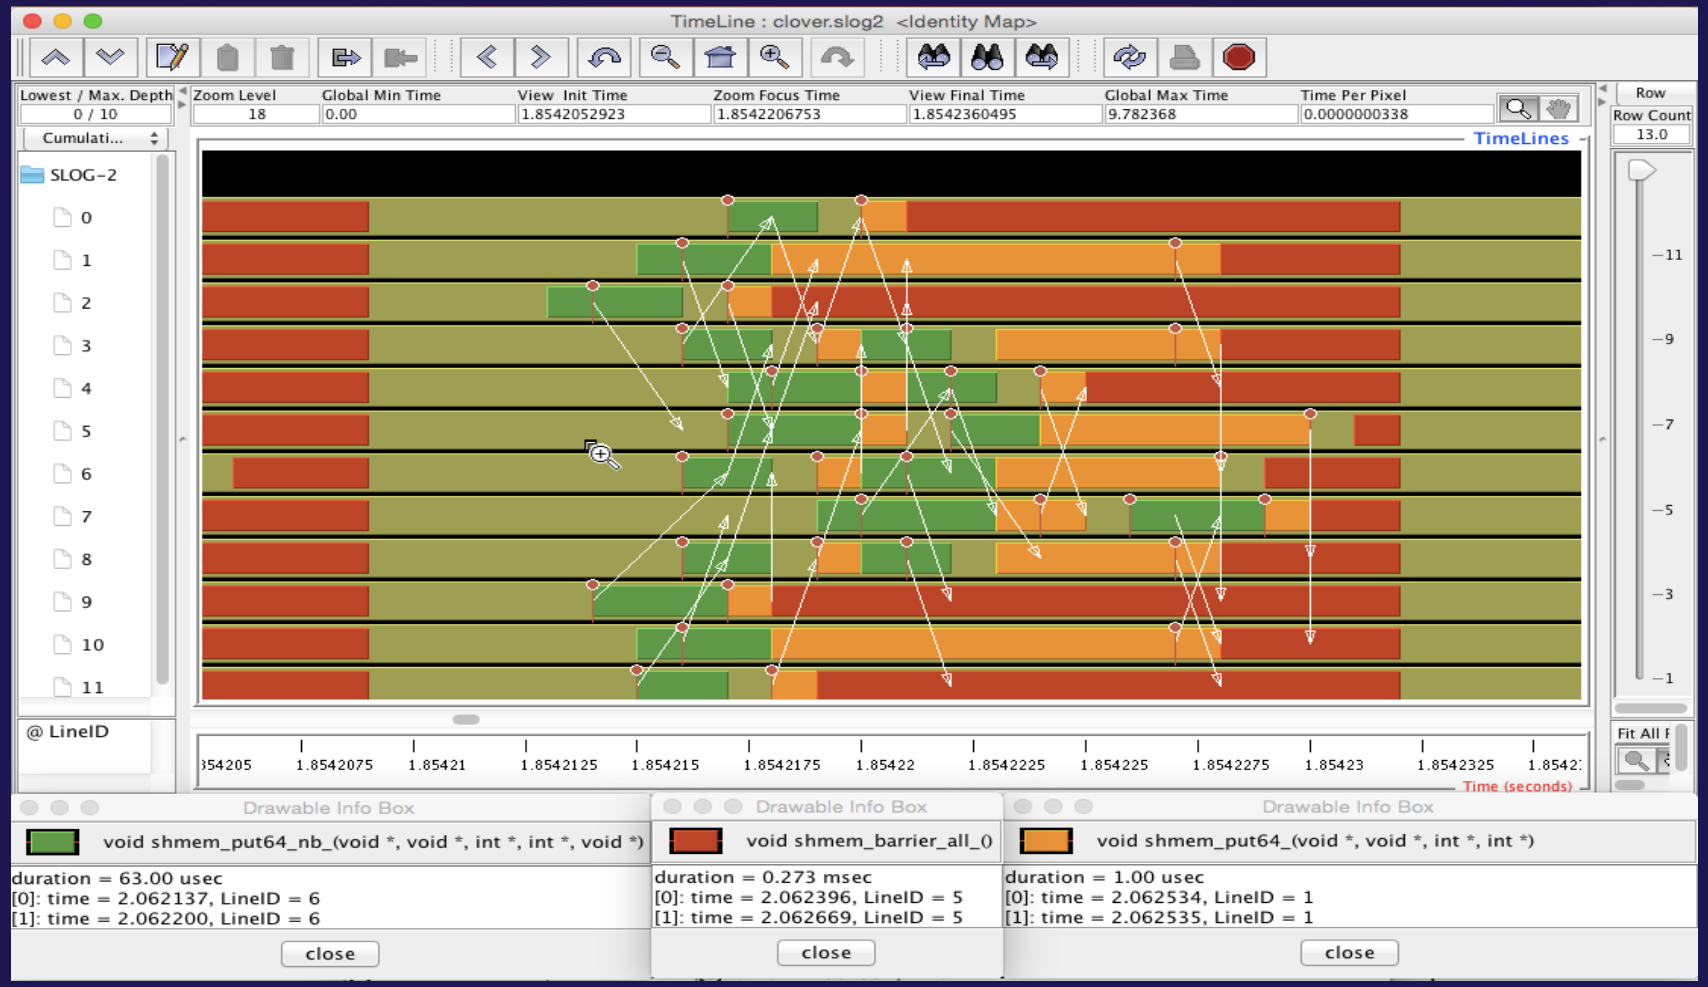

Tracing: Chrome Browser

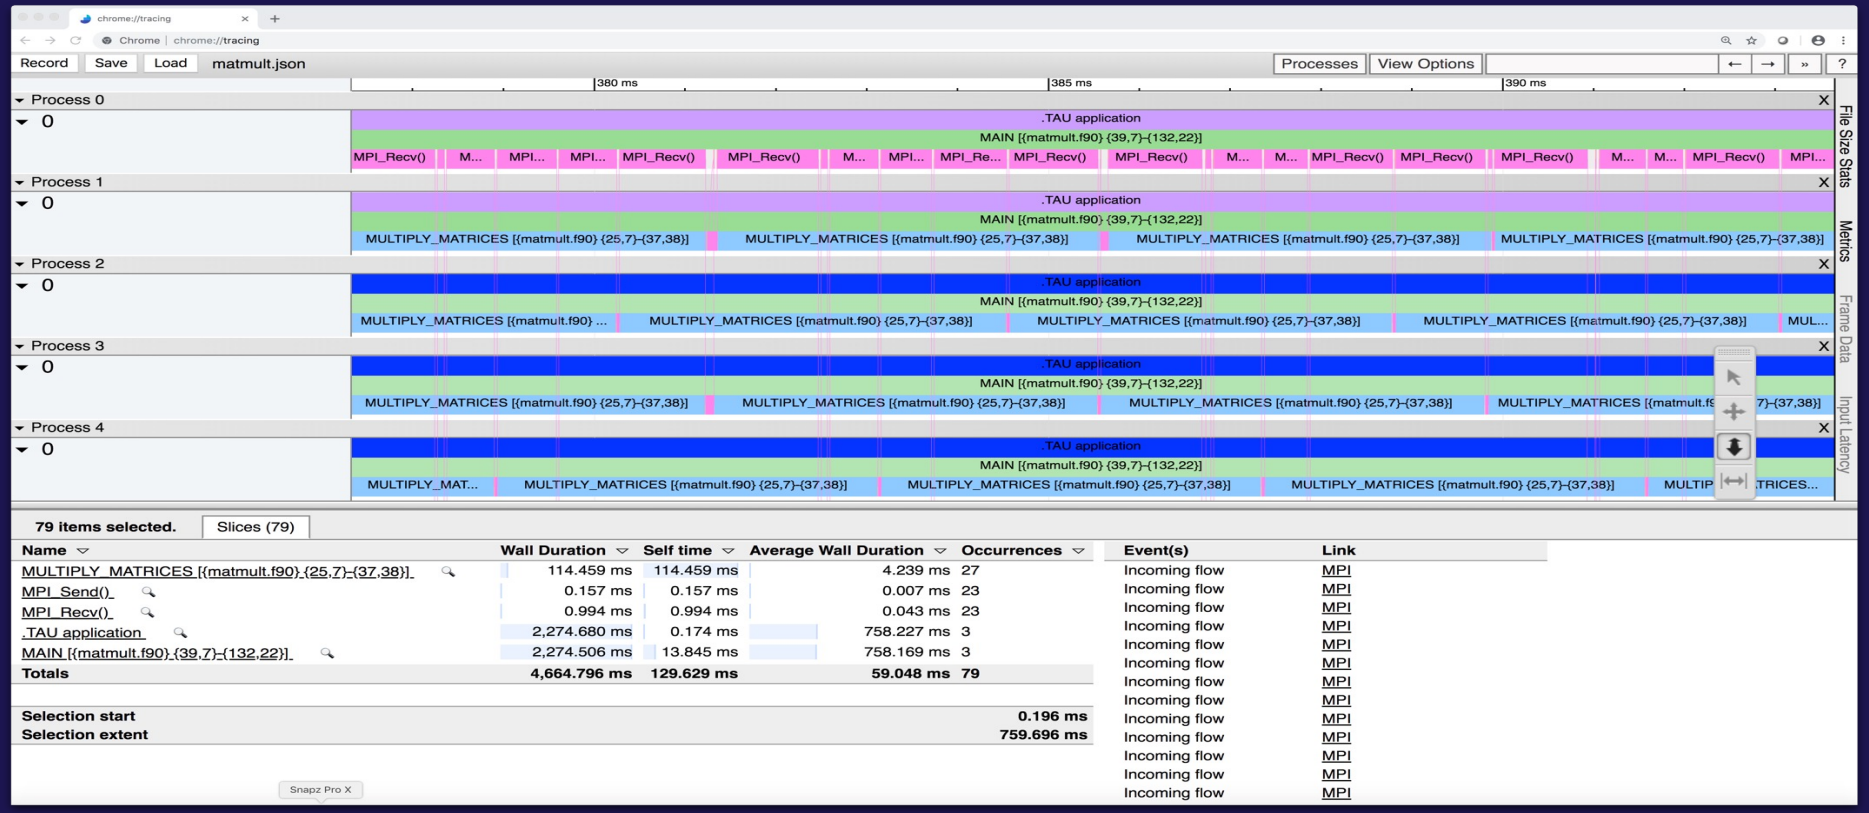


```
% export TAU_TRACE=1
% mpirun –np 256 tau_exec ./a.out
% tau_treemerge.pl; tau_trace2json tau.trc tau.edf –chrome –ignoreatomic –o app.json
```

Chrome browser: chrome://tracing (Load -> app.json)

Perfetto.dev Trace Browser: Kokkos Example

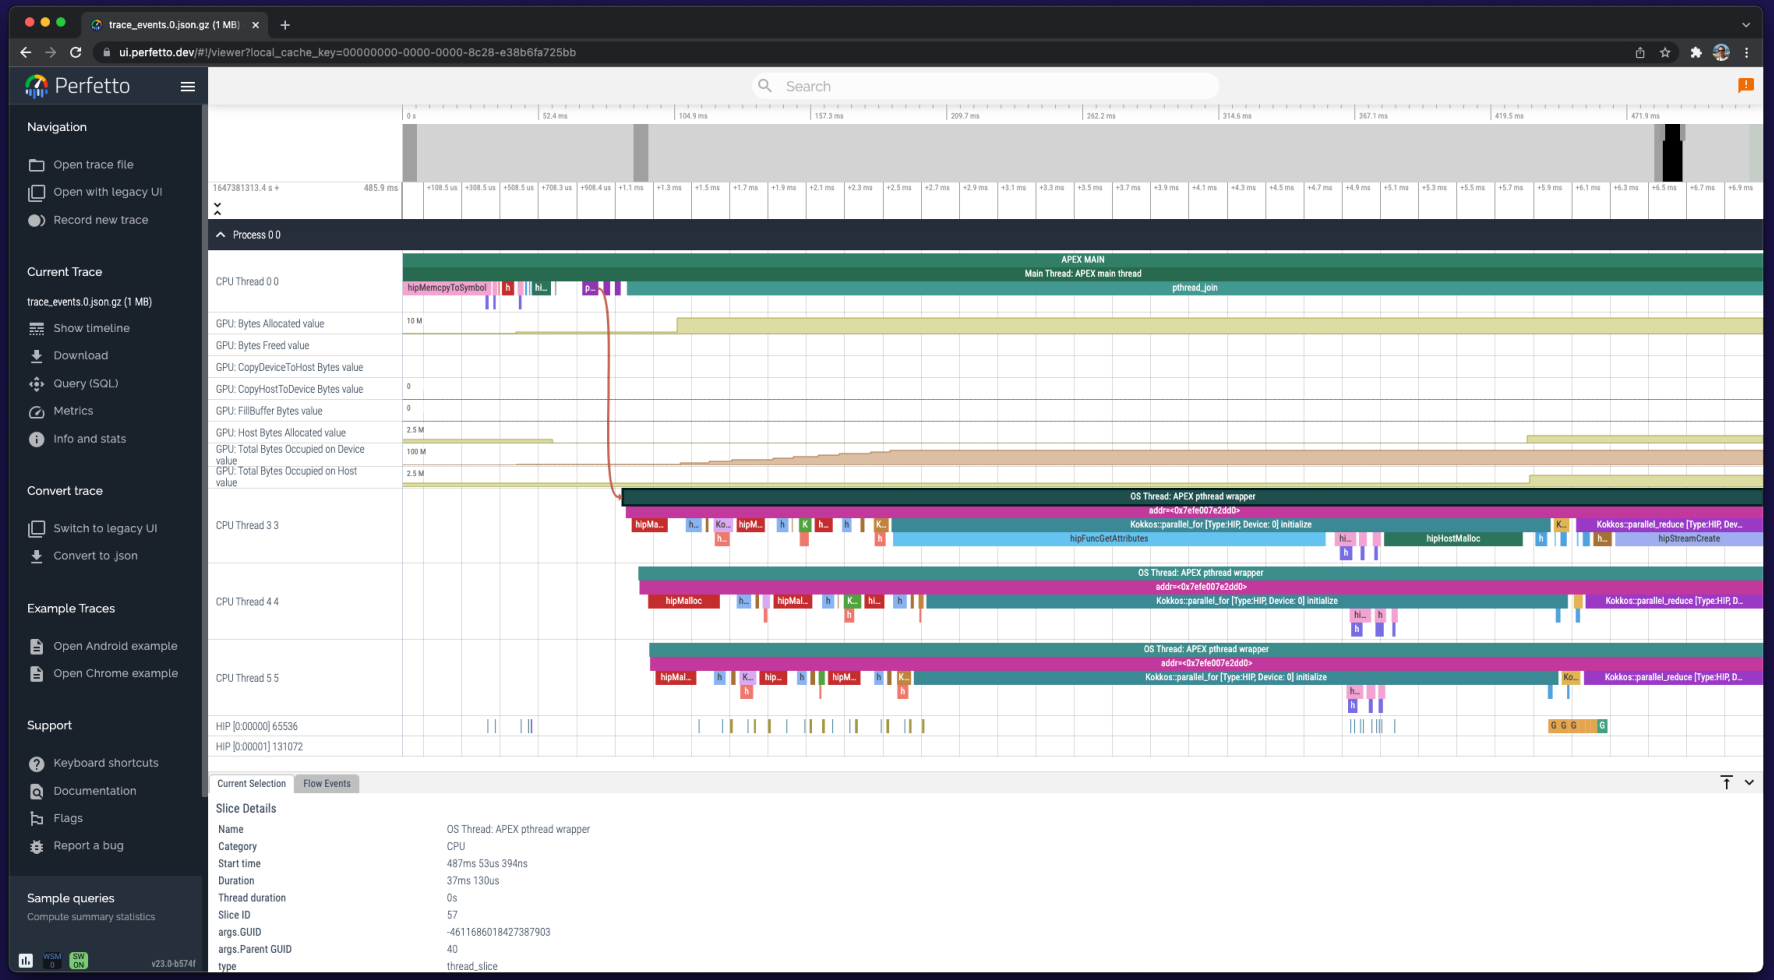

Perfetto.dev Trace Browser

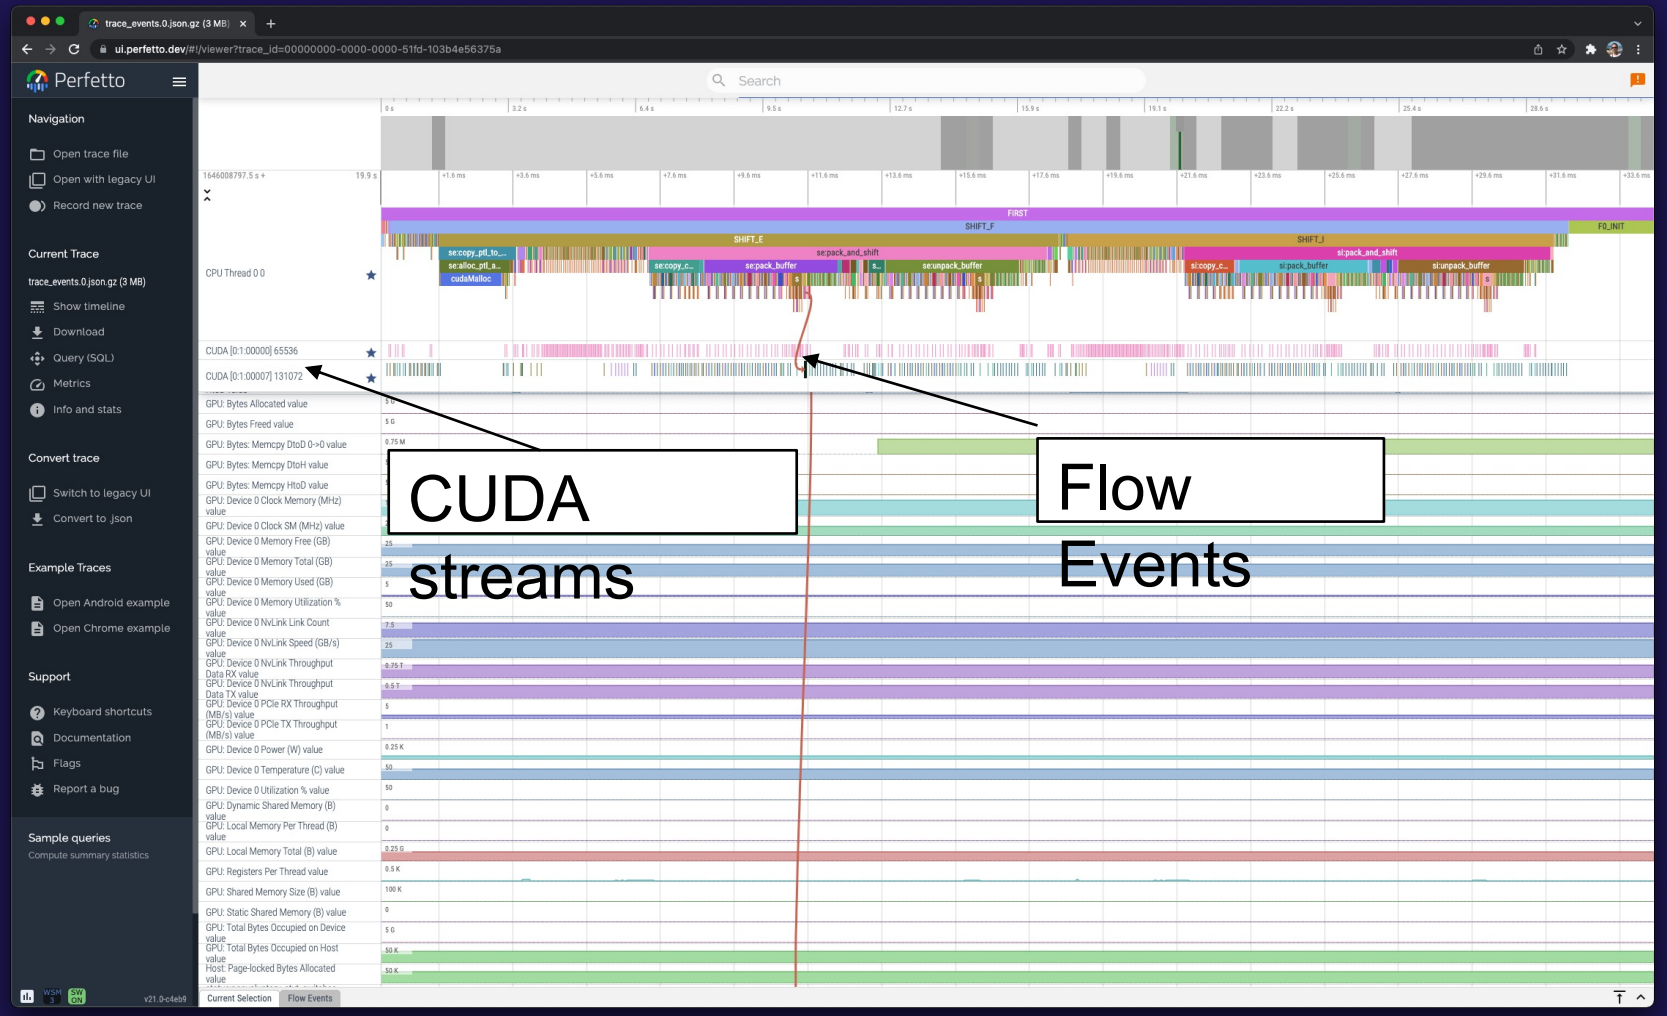

Vampir [TU Dresden] Timeline: Kokkos

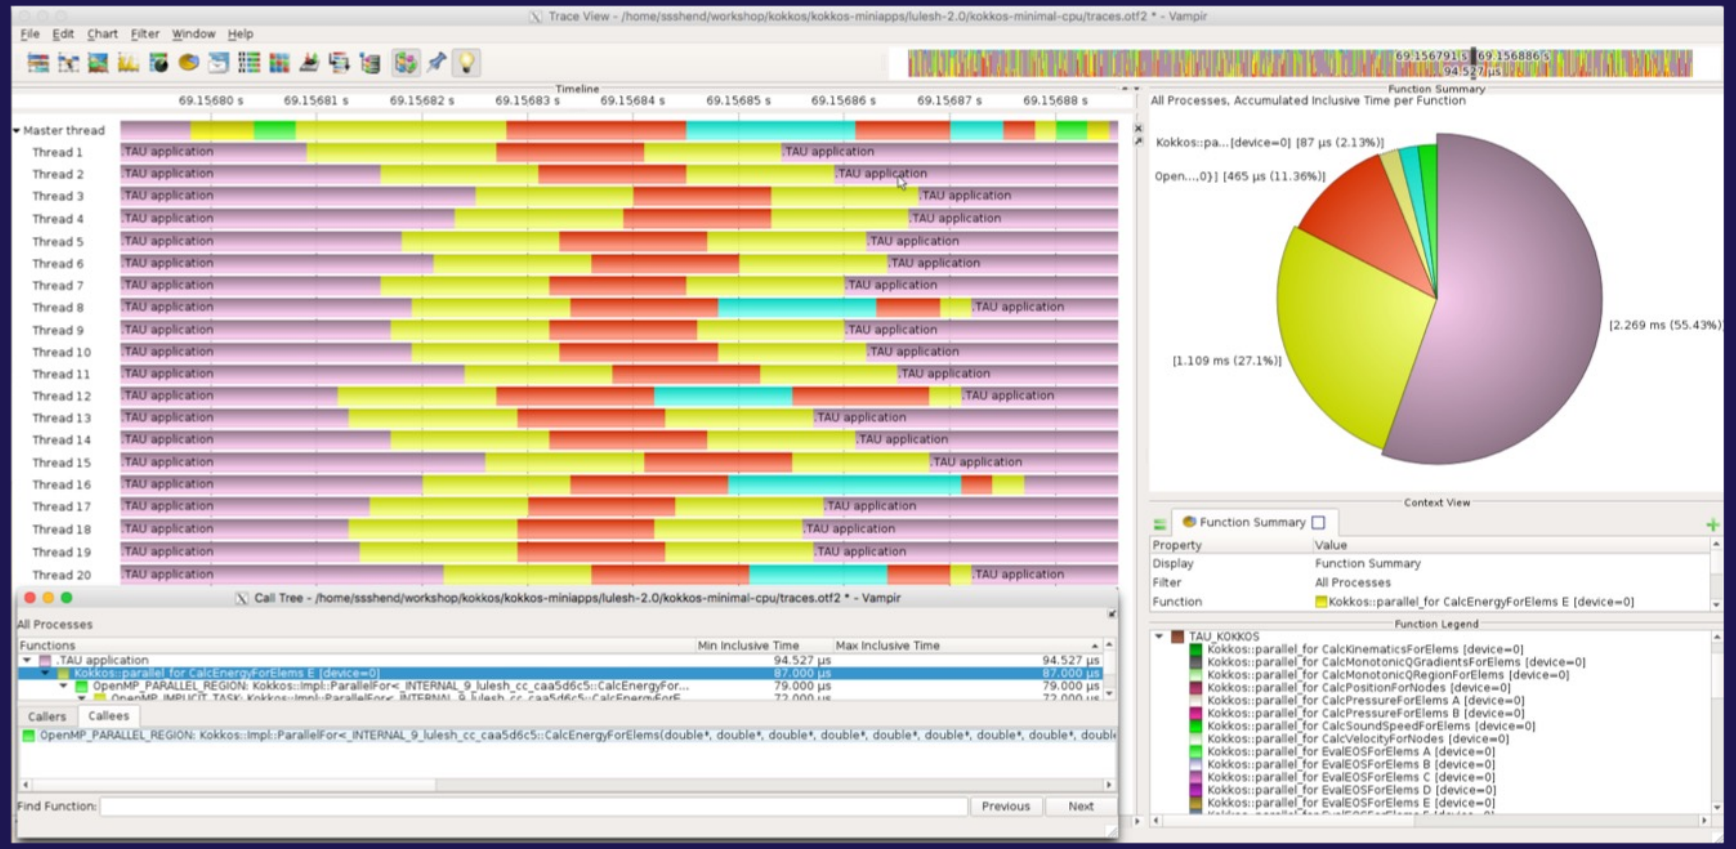


```
% export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2
% tau_exec –T serial,ompt -ompt  ./a.out
% vampir traces.otf2 &
```

Kokkos

- Provides abstractions for node level parallelism (X in MPI+X)
- Productive, portable, and performant shared-memory programming model
- Helps you create single source performance portable codes
- Provides data abstractions
- C++ API for expressing parallelism in your program
- Aggressive compiler transformations using C++ templates
- Low level code targets backends such as OpenMP, Pthread, CUDA
- Creates a problem for performance evaluation tools
- Gap: performance data and higher-level abstractions
- Solution: Kokkos profiling API for mapping performance data

TAU's Support for Runtime Systems

- *MPI*
  - PMPI profiling interface
  - MPI_T tools interface using performance and control variables
- *Pthread*
  - Captures time spent in routines per thread of execution
- *OpenMP*
  - OMPT tools interface to track salient OpenMP runtime events
  - Opari source rewriter
  - Preloading wrapper OpenMP runtime library when OMPT is not supported
- *OpenACC*
  - OpenACC instrumentation API
  - Track data transfers between host and device (per-variable)
  - Track time spent in kernels

TAU's Support for Runtime Systems (contd.)

- *OpenCL*
  - —OpenCL profiling interface
  - —Track timings of kernels
- *CUDA*
  - —Cuda Profiling Tools Interface (CUPTI)
  - —Track data transfers between host and GPU
  - —Track access to uniform shared memory between host and GPU
- *ROCm*
  - —Rocprofiler and Roctracer instrumentation interfaces
  - —Track data transfers and kernel execution between host and GPU
- *Kokkos*
  - —Kokkos profiling API
  - —Push/pop interface for region, kernel execution interface
- *Python*
  - —Python interpreter instrumentation API
  - —Tracks Python routine transitions as well as Python to C transitions`

Examples of Multi-Level Instrumentation

- *MPI + OpenMP*
  - —MPI_T + PMPI + OMPT may be used to track MPI and OpenMP
- *MPI + CUDA*
  - —PMPI + CUPTI interfaces
- *OpenCL + ROCm*
  - —Rocprofiler + OpenCL instrumentation interfaces
- *Kokkos + OpenMP*
  - —Kokkos profiling API + OMPT to transparently track events
- *Kokkos + pthread + MPI*
  - —Kokkos + pthread wrapper interposition library + PMPI layer
- *Python + CUDA + MPI*
  - —Python + CUPTI + pthread profiling interfaces (e.g., Tensorflow, PyTorch) + MPI
- *MPI + OpenCL*
  - —PMPI + OpenCL profiling interfaces

TAU Execution Command (tau_exec)

- Uninstrumented execution
  - % aprun -n 256 ./a.out
- Track GPU operations
  - % aprun –np 256 tau_exec –rocm ./a.out
  - % aprun –np 256 tau_exec –cupti ./a.out
  - % aprun –np 256 tau_exec –opencl ./a.out
  - % aprun -np 256 tau_exec –l0 ./a.out
  - % aprun –np 256 tau_exec –openacc ./a.out
- Track MPI performance
  - % aprun -n 256 tau_exec ./a.out
- Track I/O, and MPI performance (MPI enabled by default)
  - % aprun -n 256 tau_exec -io ./a.out
- Track OpenMP and MPI execution (using OMPT for Intel v19+ or Clang 8+)
  - % export TAU_OMPT_SUPPORT_LEVEL=full;
  - % aprun –np 256 tau_exec –T ompt,intel,mpi -ompt ./a.out
- Track memory operations
  - % export TAU_TRACK_MEMORY_LEAKS=1
  - % aprun –np 256 tau_exec –memory_debug ./a.out (bounds check)
- Use event based sampling (compile with –g)
  - % aprun –np 256 tau_exec –ebs ./a.out
  - Also -ebs_source=<PAPI_COUNTER> -ebs_period=<overflow_count> -ebs_resolution=<file | function | line>

AMD HIPCC: OMPT Target Offload Support in TAU

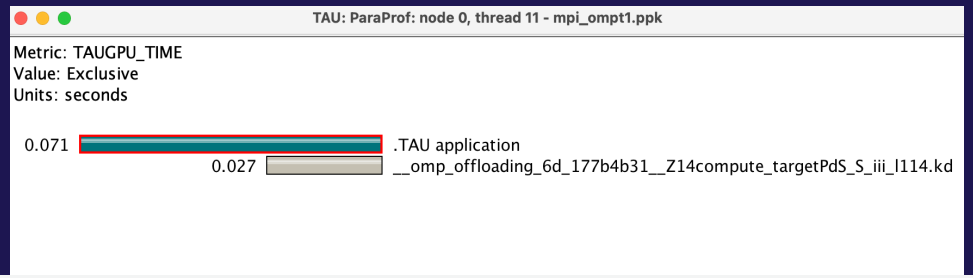

TAU: ParaProf: Statistics for: node 0, thread 0 - mpi_ompt1.ppk

Name	Exclusive...	Inclusive ...	Calls	Child Calls
.TAU application[{0}-{0xffffffff}]	1.566	1.829	1	124
MPI_Comm_rank()	0	0	1	0
MPI_Finalize()	0.01	0.01	1	1
MPI_Init()	0.118	0.118	1	0
OpenMP_Parallel_Region do_work() [{/autofs/nccs-svm1_home1/sameer/public/tau2/examples/openmp/mpi_target/matmult} {0, 0}]	0	0.025	20	47
OpenMP_Parallel_Region do_work() [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {148, 0}]	0	0.007	10	20
OpenMP_Sync_Region_Barrier_Implicit do_work() [{/autofs/nccs-svm1_home1/sameer/public/tau2/examples/openmp/mpi_target/matmult} {0, 0}]	0.023	0.023	20	0
OpenMP_Sync_Region_Barrier_Implicit do_work() [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {148, 0}]	0.006	0.006	10	0
OpenMP_Target allocateMatrix(int, int) [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {38, 0}]	0.007	0.007	30	1
OpenMP_Target compute_target(double*, double*, double*, int, int, int) [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {113, 0}]	0	0	20	0
OpenMP_Target compute_target(double*, double*, double*, int, int, int) [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {114, 0}]	0.094	0.094	10	0
OpenMP_Target freeMatrix(double*, int, int) [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {58, 0}]				
OpenMP_Work_Loop .omp_outlined. [{/ccs/home/sameer/tau2/examples/openmp/mpi_target/matmult.cpp} {48, 0}]				
pthread_create				
pthread_join				

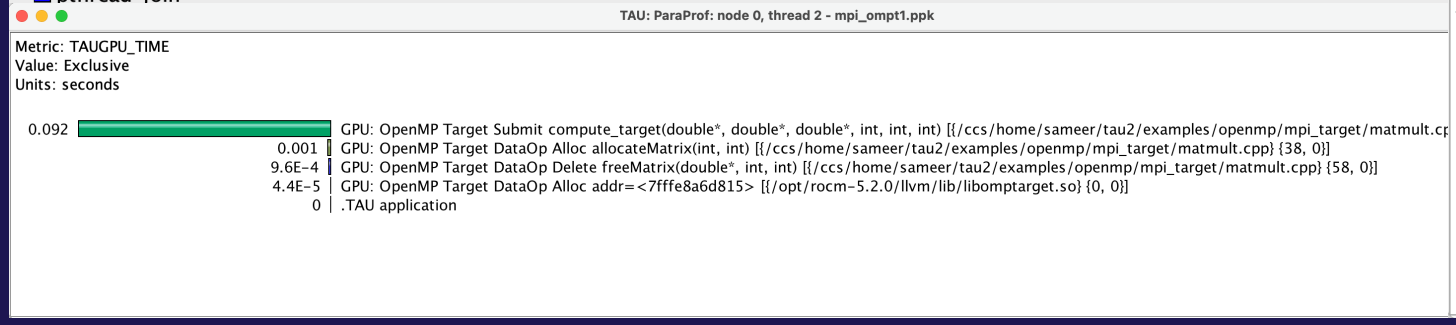

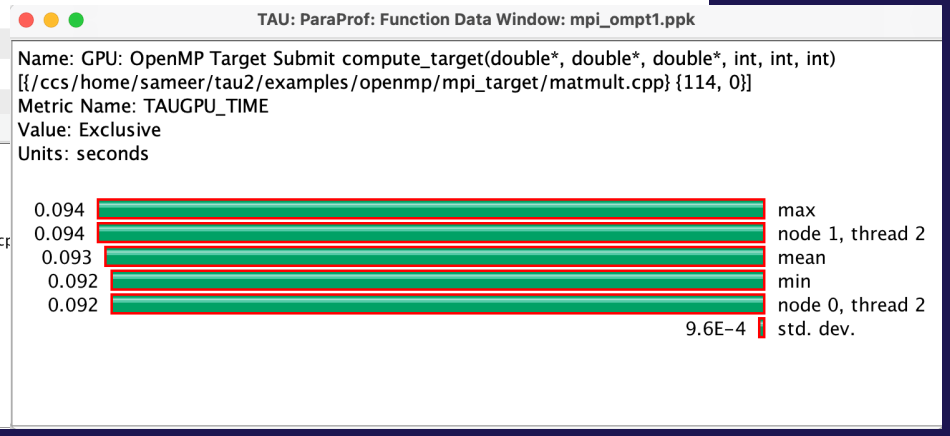

TAU's Runtime Environment Variables

Environment Variable	Default	Description
TAU_TRACE	0	Setting to 1 turns on tracing
TAU_CALLPATH	0	Setting to 1 turns on callpath profiling
TAU_TRACK_MEMORY_FOOTPRINT	0	Setting to 1 turns on tracking memory usage by sampling periodically the resident set size and high water mark of memory usage
TAU_TRACK_POWER	0	Tracks power usage by sampling periodically.
TAU_CALLPATH_DEPTH	2	Specifies depth of callpath. Setting to 0 generates no callpath or routine information, setting to 1 generates flat profile and context events have just parent information (e.g., Heap Entry: foo)
TAU_SAMPLING	1	Setting to 1 enables event-based sampling.
TAU_TRACK_SIGNALS	0	Setting to 1 generate debugging callstack info when a program crashes
TAU_COMM_MATRIX	0	Setting to 1 generates communication matrix display using context events
TAU_THROTTLE	1	Setting to 0 turns off throttling. Throttles instrumentation in lightweight routines that are called frequently
TAU_THROTTLE_NUMCALLS	100000	Specifies the number of calls before testing for throttling
TAU_THROTTLE_PERCALL	10	Specifies value in microseconds. Throttle a routine if it is called over 100000 times and takes less than 10 usec of inclusive time per call
TAU_CALLSITE	0	Setting to 1 enables callsite profiling that shows where an instrumented function was called. Also compatible with tracing.
TAU_PROFILE_FORMAT	Profile	Setting to "merged" generates a single file. "snapshot" generates xml format
TAU_METRICS	TIME	Setting to a comma separated list generates other metrics. (e.g., ENERGY,TIME,P_VIRTUAL_TIME,PAPI_FP_INS,PAPI_NATIVE_<event>:<subevent>)

Runtime Environment Variables

Environment Variable	Default	Description
TAU_TRACE	0	Setting to 1 turns on tracing
TAU_TRACE_FORMAT	Default	Setting to “otf2” turns on TAU’s native OTF2 trace generation (configure with –otf=download)
TAU_EBS_UNWIND	0	Setting to 1 turns on unwinding the callstack during sampling (use with tau_exec –ebs or TAU_SAMPLING=1)
TAU_EBS_RESOLUTION	line	Setting to “function” or “file” changes the sampling resolution to function or file level respectively.
TAU_TRACK_LOAD	0	Setting to 1 tracks system load on the node
TAU_SELECT_FILE	Default	Setting to a file name, enables selective instrumentation based on exclude/include lists specified in the file.
TAU_OMPT_SUPPORT_LEVEL	basic	Setting to “full” improves resolution of OMPT TR6 regions on threads 1.. N-1. Also, “lowoverhead” option is available.
TAU_OMPT_RESOLVE_ADDRESS_EAGERLY	1	Setting to 1 is necessary for event based sampling to resolve addresses with OMPT. Setting to 0 allows the user to do offline address translation.

Runtime Environment Variables

Environment Variable	Default	Description
TAU_TRACK_MEMORY_LEAKS	0	Tracks allocates that were not de-allocated (needs –optMemDbg or tau_exec –memory)
TAU_EBS_SOURCE	TIME	Allows using PAPI hardware counters for periodic interrupts for EBS (e.g., TAU_EBS_SOURCE=PAPI_TOT_INS when TAU_SAMPLING=1)
TAU_EBS_PERIOD	100000	Specifies the overflow count for interrupts
TAU_MEMDBG_ALLOC_MIN/MAX	0	Byte size minimum and maximum subject to bounds checking (used with TAU_MEMDBG_PROTECT_*)
TAU_MEMDBG_OVERHEAD	0	Specifies the number of bytes for TAU's memory overhead for memory debugging.
TAU_MEMDBG_PROTECT_BELOW/ABOVE	0	Setting to 1 enables tracking runtime bounds checking below or above the array bounds (requires –optMemDbg while building or tau_exec –memory)
TAU_MEMDBG_ZERO_MALLOC	0	Setting to 1 enables tracking zero byte allocations as invalid memory allocations.
TAU_MEMDBG_PROTECT_FREE	0	Setting to 1 detects invalid accesses to deallocated memory that should not be referenced until it is reallocated (requires –optMemDbg or tau_exec –memory)
TAU_MEMDBG_ATTEMPT_CONTINUE	0	Setting to 1 allows TAU to record and continue execution when a memory error occurs at runtime.
TAU_MEMDBG_FILL_GAP	Undefined	Initial value for gap bytes
TAU_MEMDBG_ALINGMENT	Sizeof(int)	Byte alignment for memory allocations
TAU_EVENT_THRESHOLD	0.5	Define a threshold value (e.g., .25 is 25%) to trigger marker events for min/max

TAU: Key takeaways

- There is no need to modify your application source code, build system, or the binary
- TAU supports GPUs (Intel, AMD, NVIDIA) as well as CPUs
- Simply launch the application using tau_exec [options]
- Launch paraprof on Polaris or bring the ppk file to your laptop and launch paraprof
- You may also use Cooley for a VNC session

Download TAU from U. Oregon

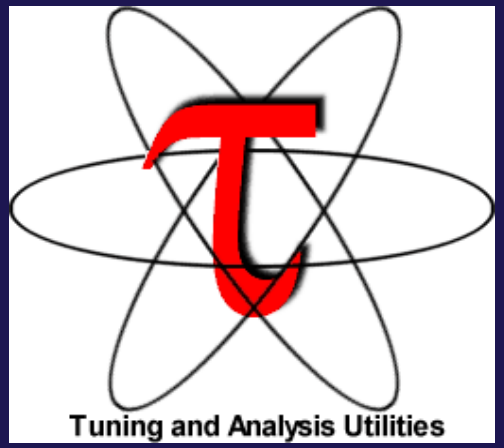

http://tau.uoregon.edu

for more information

Free download, open source, BSD license

Performance Research Laboratory, University of Oregon, Eugene

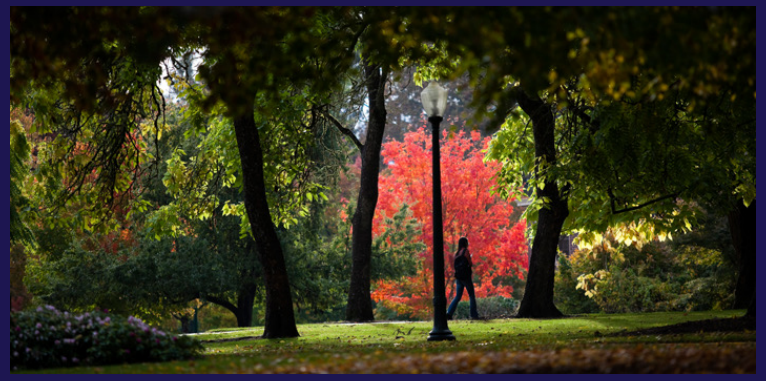

www.uoregon.edu

Support Acknowledgements

- US Department of Energy (DOE)
  - — ANL
  - — Office of Science contracts, ECP
  - — SciDAC, LBL contracts
  - — LLNL-LANL-SNL ASC/NNSA contract
  - — Battelle, PNNL and ORNL contract
- Department of Defense (DoD)
  - — PETTT, HPCMP
- National Science Foundation (NSF)
  - — SI2-SSI, Glassbox
- NASA
- CEA, France
- Partners:
  - —University of Oregon
  - —The Ohio State University
  - —ParaTools, Inc.
  - —University of Tennessee, Knoxville
  - —T.U. Dresden, GWT
  - —Jülich Supercomputing Center

Acknowledgment

“This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation’s exascale computing imperative.”

Hands-on session: TAU

Using TAU on Polaris natively

- Setup preferred program environment compilers (check instructions)

```
% ssh -Y <login>@polaris.alcf.anl.gov
% module load tau
% tar zxf /soft/perftools/tau/tar/workshop.tgz; cd workshop
% paraprof demo.ppk &

If you are on a Mac with Xquartz, you may need:
% paraprof -fix-xquartz demo.ppk &
In the directory where profile.* files are created. Xquartz 2.7.4 works well without this.
Please do not use paraprof on the compute nodes. You may also use Cooley (VNC) or
install TAU locally on your laptop.
```

TAU Breakout Session – CUDA with MPI on Polaris

- Setup preferred program environment compilers (check instructions)

```
% ssh -Y <login>@polaris.alcf.anl.gov
% module load tau
% tar zxf /soft/perftools/tau/tar/workshop.tgz
% cd workshop/TeaLeaf_CUDA;
% make clean
% make; cd bin
% qsub -I -q fallws23single -t 60 -n 1 -A fallwkshp23
% ./run.sh
% pprof -a | more
% paraprof --pack app.ppk
You may use paraprof --dump app.ppk to write out the profile.* files.
Bring ppk file to your desktop:
% paraprof app.ppk &
```

Setup: Installing TAU on Laptops

- Prerequisites: Java in your path
- Microsoft Windows
  - Install Java from Oracle.com
    - http://tau.uoregon.edu/tau.exe
    - Install, click on a ppk file to launch paraprof
- macOS (x86_64)
  - Install Java 11.0.3:
    - Download and install http://tau.uoregon.edu/java.dmg
    - If you have multiple Java installations, add to your ~/.zshrc (or ~/.bashrc as appropriate):
      - export PATH=/Library/Java/JavaVirtualMachines/jdk-11.0.3.jdk/Contents/Home/bin:$PATH
  - Download and install TAU (copy to /Applications from dmg):
    - http://tau.uoregon.edu/tau.dmg
    - export PATH=/Applications/TAU/tau/apple/bin:$PATH
    - paraprof app.ppk &
- macOS (arm64, Apple Silicon M1/M2)
    - http://tau.uoregon.edu/java_arm64.dmg
    - http://tau.uoregon.edu/tau_arm64.dmg
- Linux (http://tau.uoregon.edu/tau.tgz)
  - ./configure; make install; export PATH=<taudir>/x86_64/bin:$PATH; paraprof app.ppk &

Using VNC on Cooley to use a remote desktop

- These instructions are also in README.Cooley in /soft/perftools/tau/tar/workshop.tgz

```
% Terminal 1
ssh cooley.alcf.anl.gov
Add to ~/.soft.cooley
+tau
+java
@default

Then, launch:
vncpasswd
(set the VNC password and say no to saving view only password)
qsub -I -n 1 -t 50 -A <ACCOUNT>
see which host (e.g., cc054 or ccXX - using ccXX for the example. Please use correct hostname below instead of ccXX.)
x0vncserver --display=:0.0 --NeverShared=1 --geometry=1400x800+0+0 --PasswordFile=$HOME/.vnc/passwd --MaxProcessorUsage=100

Terminal 2
ssh -L 5900:ccXX:5900 cooley.alcf.anl.gov
ssh ccXX "export DISPLAY=:0.0; ~/.vnc/xstartup"

Open XVNC viewer
localhost:5900
Open a terminal
launch the terminal window.
paraprof app.ppk
```

TAU Breakout Session – PETSc and CUDA with MPI on Polaris

- Setup preferred program environment compilers (check instructions)

```
% ssh -Y <login>@polaris.alcf.anl.gov
% module load tau
% tar zxf /soft/perftools/tau/tar/workshop.tgz
% cd workshop/petsc-tau
% ./compile.sh
% qsub -I -q fallws23single -t 60 -n 1 -A fallwkshp23
% module load tau
% ./run.sh
% pprof -a | more
% exit
% paraprof &   (To run this on a login node on Polaris)
% paraprof --pack app.ppk
You may use paraprof --dump app.ppk to write out the profile.* files.
Bring ppk file to your desktop or use with VNC on Cooley:
% paraprof app.ppk &
```