

TAU Performance System®

IWOCL 2022


Prof. Sameer Shende
Research Associate Professor and Director,
Performance Research Laboratory, OACISS, University of Oregon
President and Director, ParaTools, Inc. and ParaTools, SAS

http://tau.uoregon.edu/TAU_IWOCL22.pdf



UNIVERSITY
OF OREGON

ParaTools



U.S. DEPARTMENT OF
ENERGY

Office of
Science

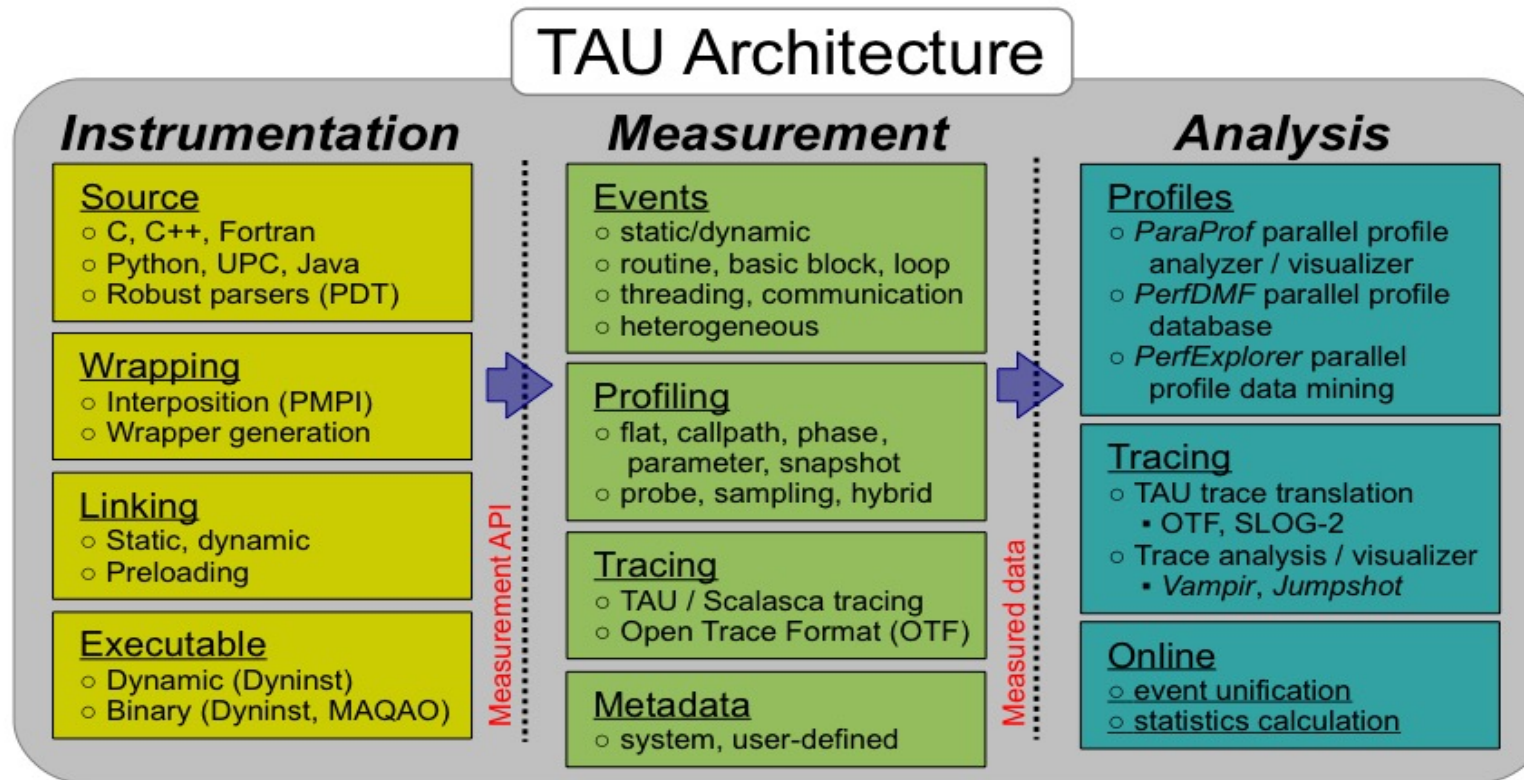

Motivation and Challenges

- With growing hardware complexity, it is getting harder to accurately measure and optimize the performance of our HPC and AI/ML workloads.
- TAU Performance System®:
 - Deliver a scalable, portable, performance evaluation toolkit for HPC and AI/ML workloads.
 - Supports OpenCL programming model across GPUs from Intel, AMD, and NVIDIA.
 - <http://tau.uoregon.edu>

TAU Performance System[®]

Parallel performance framework and toolkit

- Aims to support all HPC platforms, compilers, runtime system
- Provides portable instrumentation, measurement, analysis



TAU Performance System

Instrumentation

- Python, C++, C, Fortran, UPC, Java, Chapel, Spark
- Automatic instrumentation

Measurement and analysis support

- MPI, OpenSHMEM, ARMCI, PGAS
- pthreads, OpenMP, OMPT interface, hybrid, other thread models
- GPU: OpenCL, DPC++/SYCL, ROCm, CUDA, OpenACC
- 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?
- 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 `mpirun tau_exec ./app` with options.

Profiling and Tracing

Profiling




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

Tracing




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

ParaProf Profile Browser

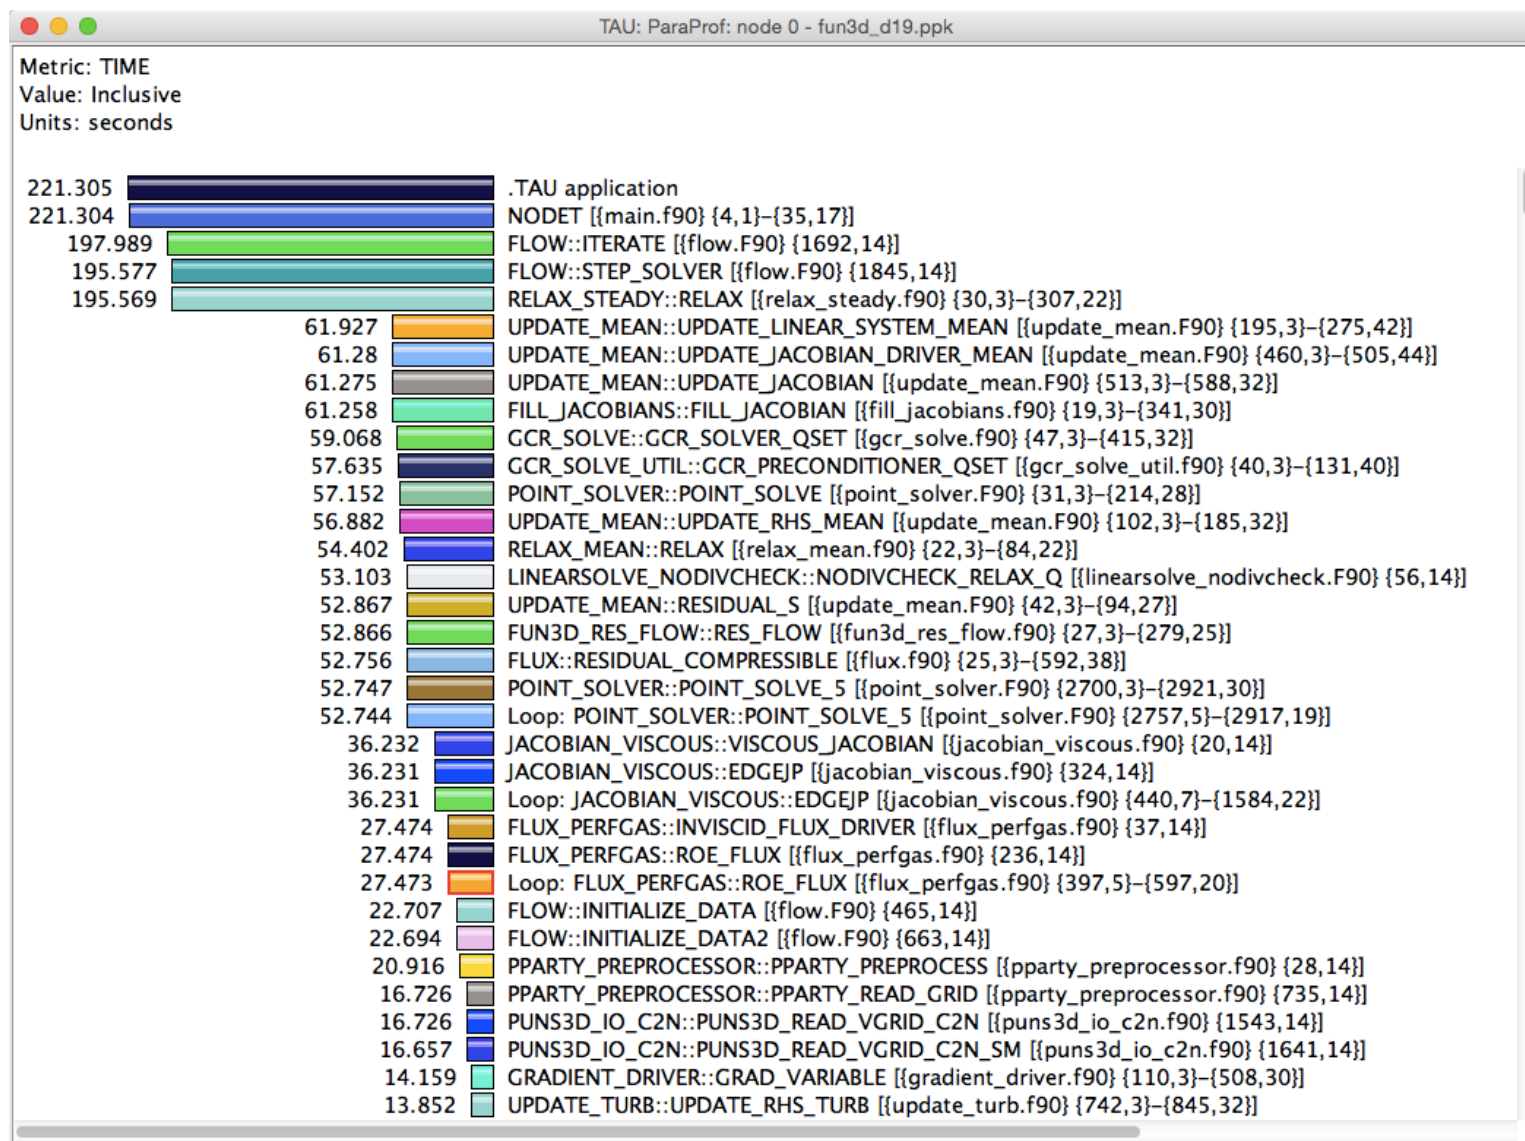


% paraprof

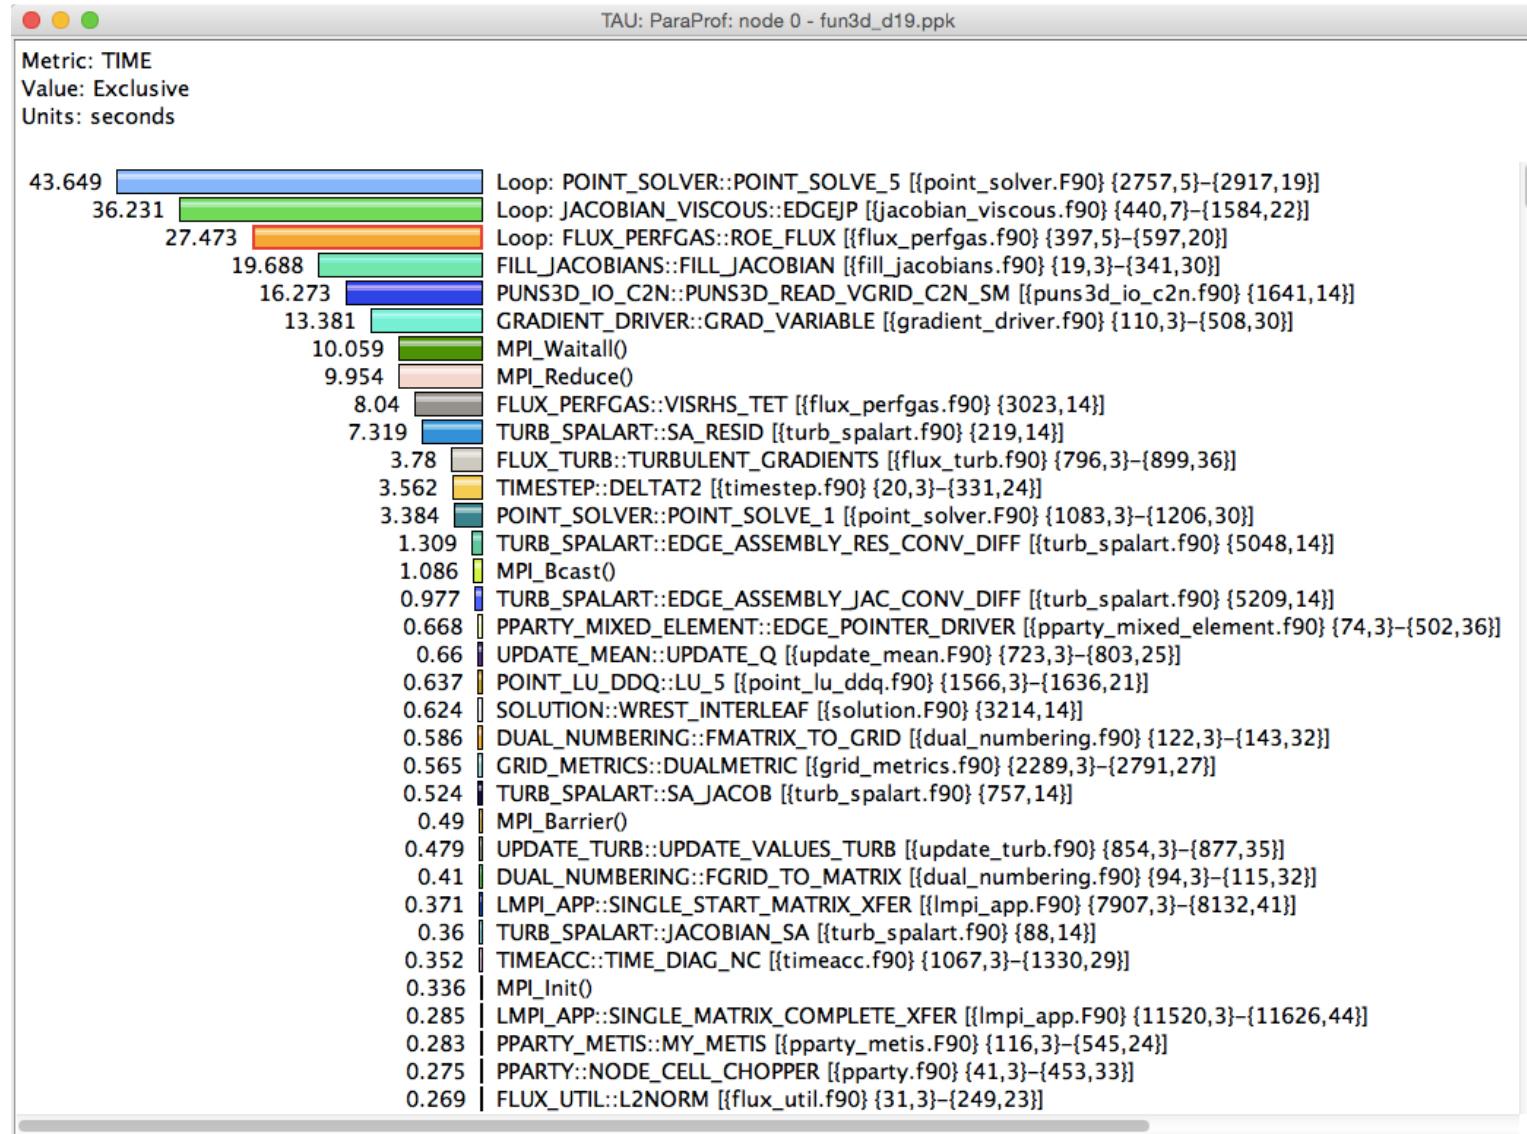
TAU's ParaProf Profile Browser




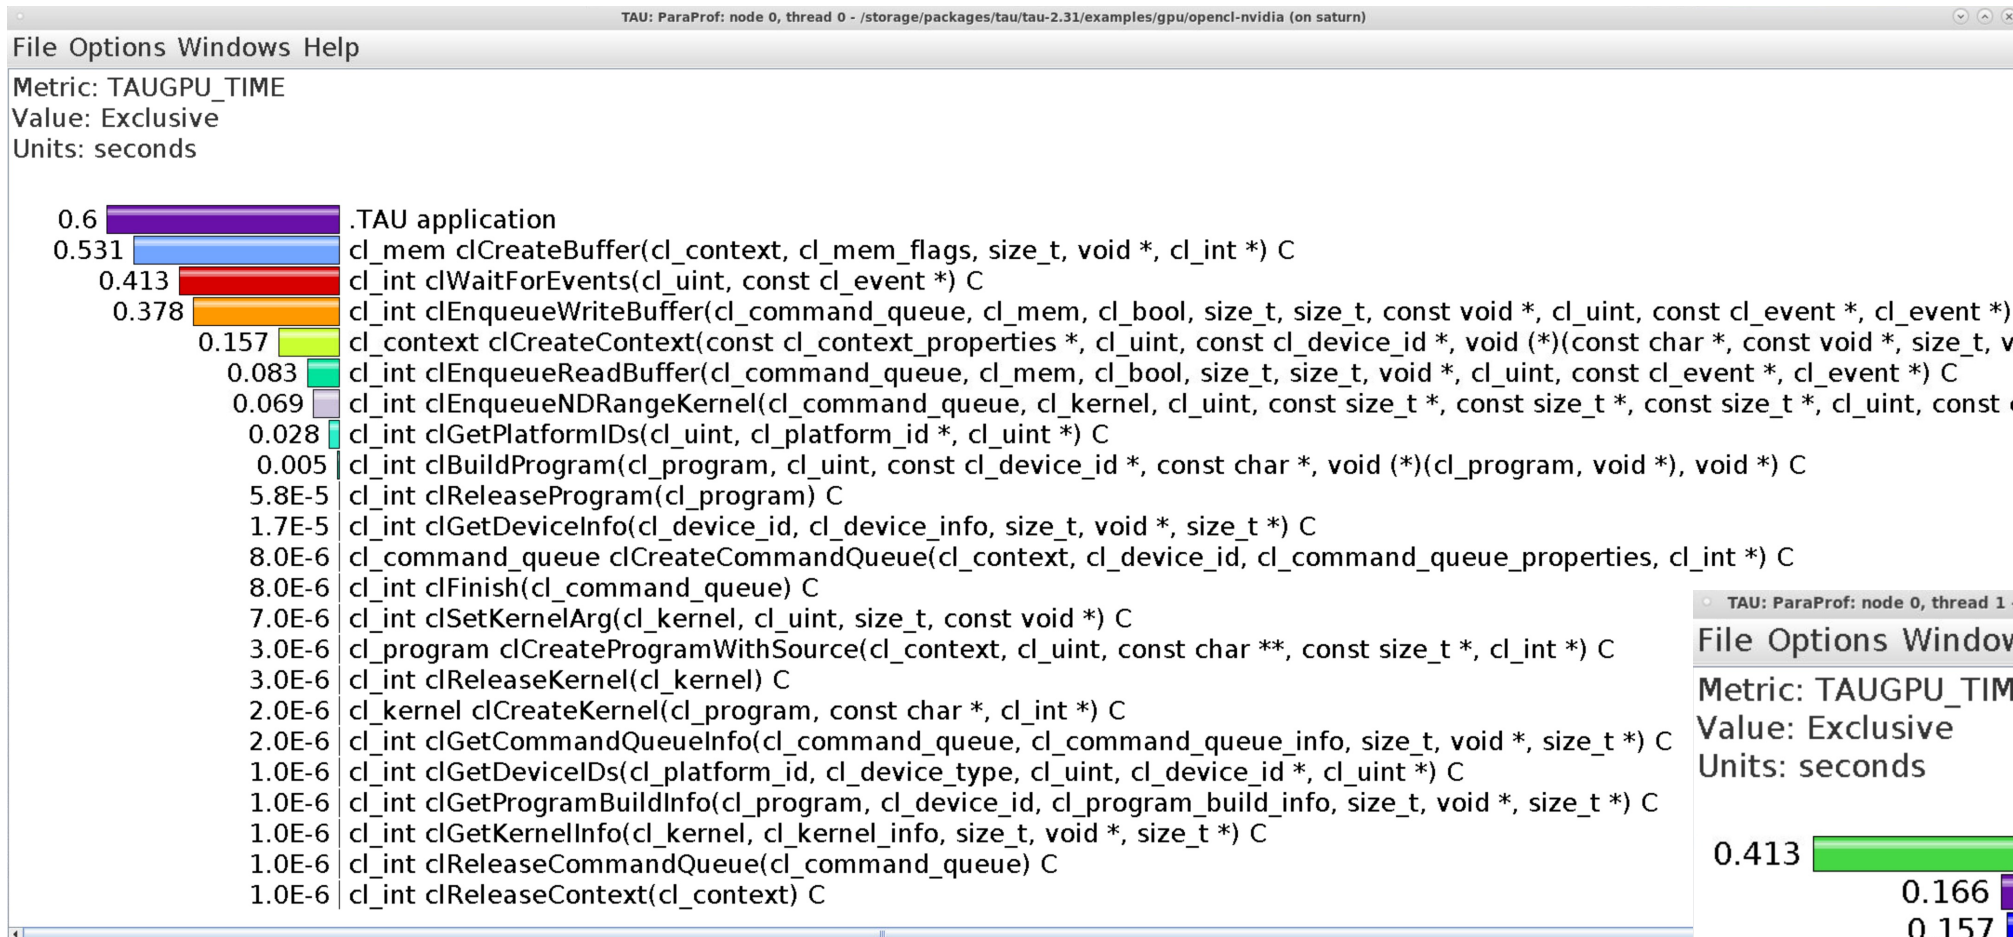
Inclusive Measurements



Exclusive Time



TAU: OpenCL profiling on NVIDIA A100 GPU



% tau_exec -T cupti,serial -opencl ./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_TIME	Inclusive TAUGPU_TIME	Calls	Child Calls
.TAU application	0.18	22.279	1	10,002
_ZTSZZ13iso3dfdDeviceRN2c14sycl5queueEPfS3_S3_mmmmmmmjENKUIRT_E313_16clINSO_7handlerEEEDa55_EUIS4_E399_58	11.063	11.063	5,000	0
_ZTSZZ13iso3dfdDeviceRN2c14sycl5queueEPfS3_S3_mmmmmmmjENKUIRT_E313_16clINSO_7handlerEEEDa55_EUIS4_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_TIME	Inclusive TAUGPU_TIME	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*) [{crststuff.c} {0}]	0.09	0.09	2	0
[SAMPLE] cl::sycl::detail::Scheduler::GraphBuilder::cleanupCommandsForRecord(cl::sycl::detail::MemObjRecord*) [{crststuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::LeavesCollection::push_back(cl::sycl::detail::Command*) [{crststuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] cl::sycl::detail::ExecCGCommand::enqueueImp0 [{crststuff.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) [{crststuff.c} {0}]	0.03	0.03	1	0
[SAMPLE] _pi_device::getAvailableCommandList(_pi_queue*, _ze_command_list_handle_t**, _ze_fence_handle_t**) [{crststuff.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/libglibc.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/lib-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/src	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)

TAU: ParaProf: Statistics for: node 0, thread 0 - ze_gemm_4096.ppk

Name	Exclusive TAUGPU_T...	Inclusive TAUGPU_Tl...	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-zero	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

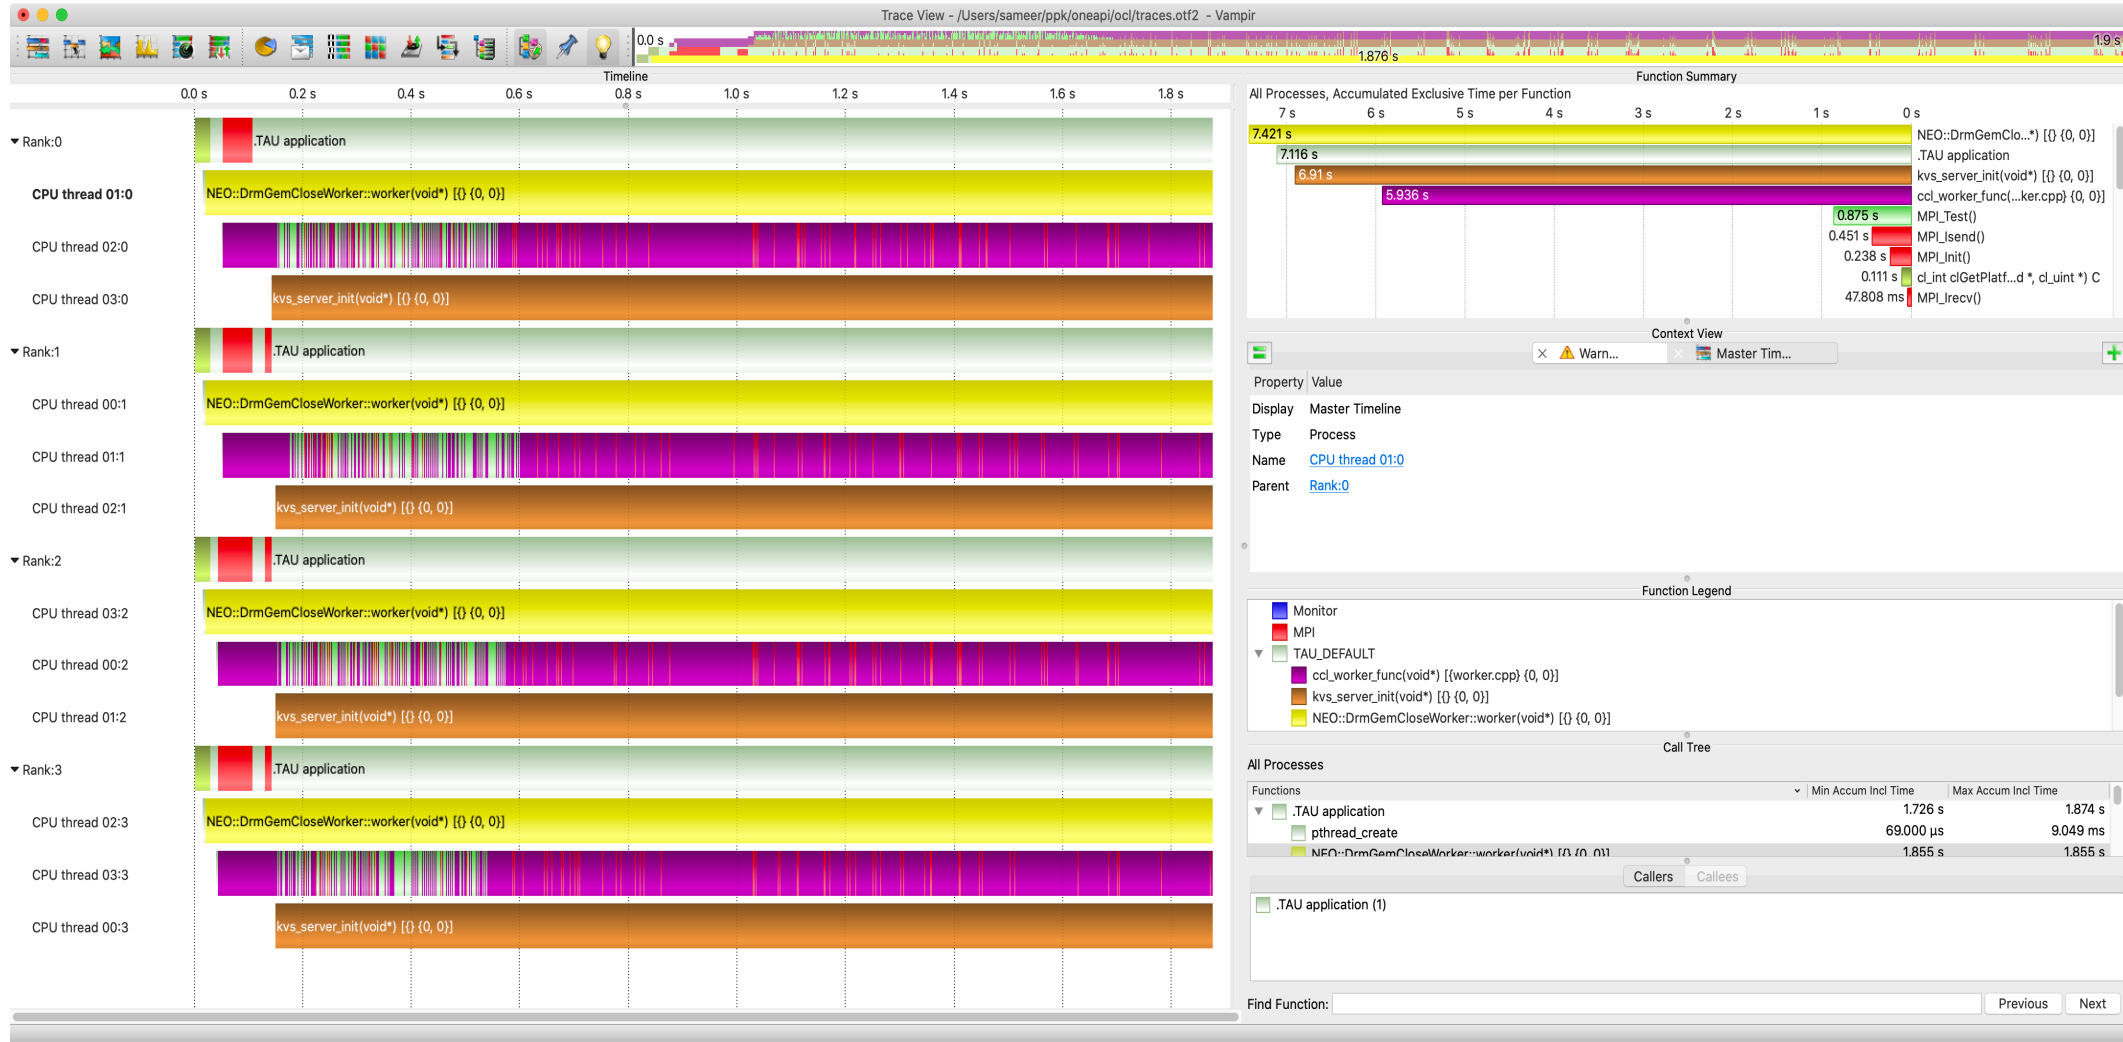
TAU: ParaProf: Statistics for: node 0, thread 2 - ze_gemm_4096.ppk

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

Units: seconds

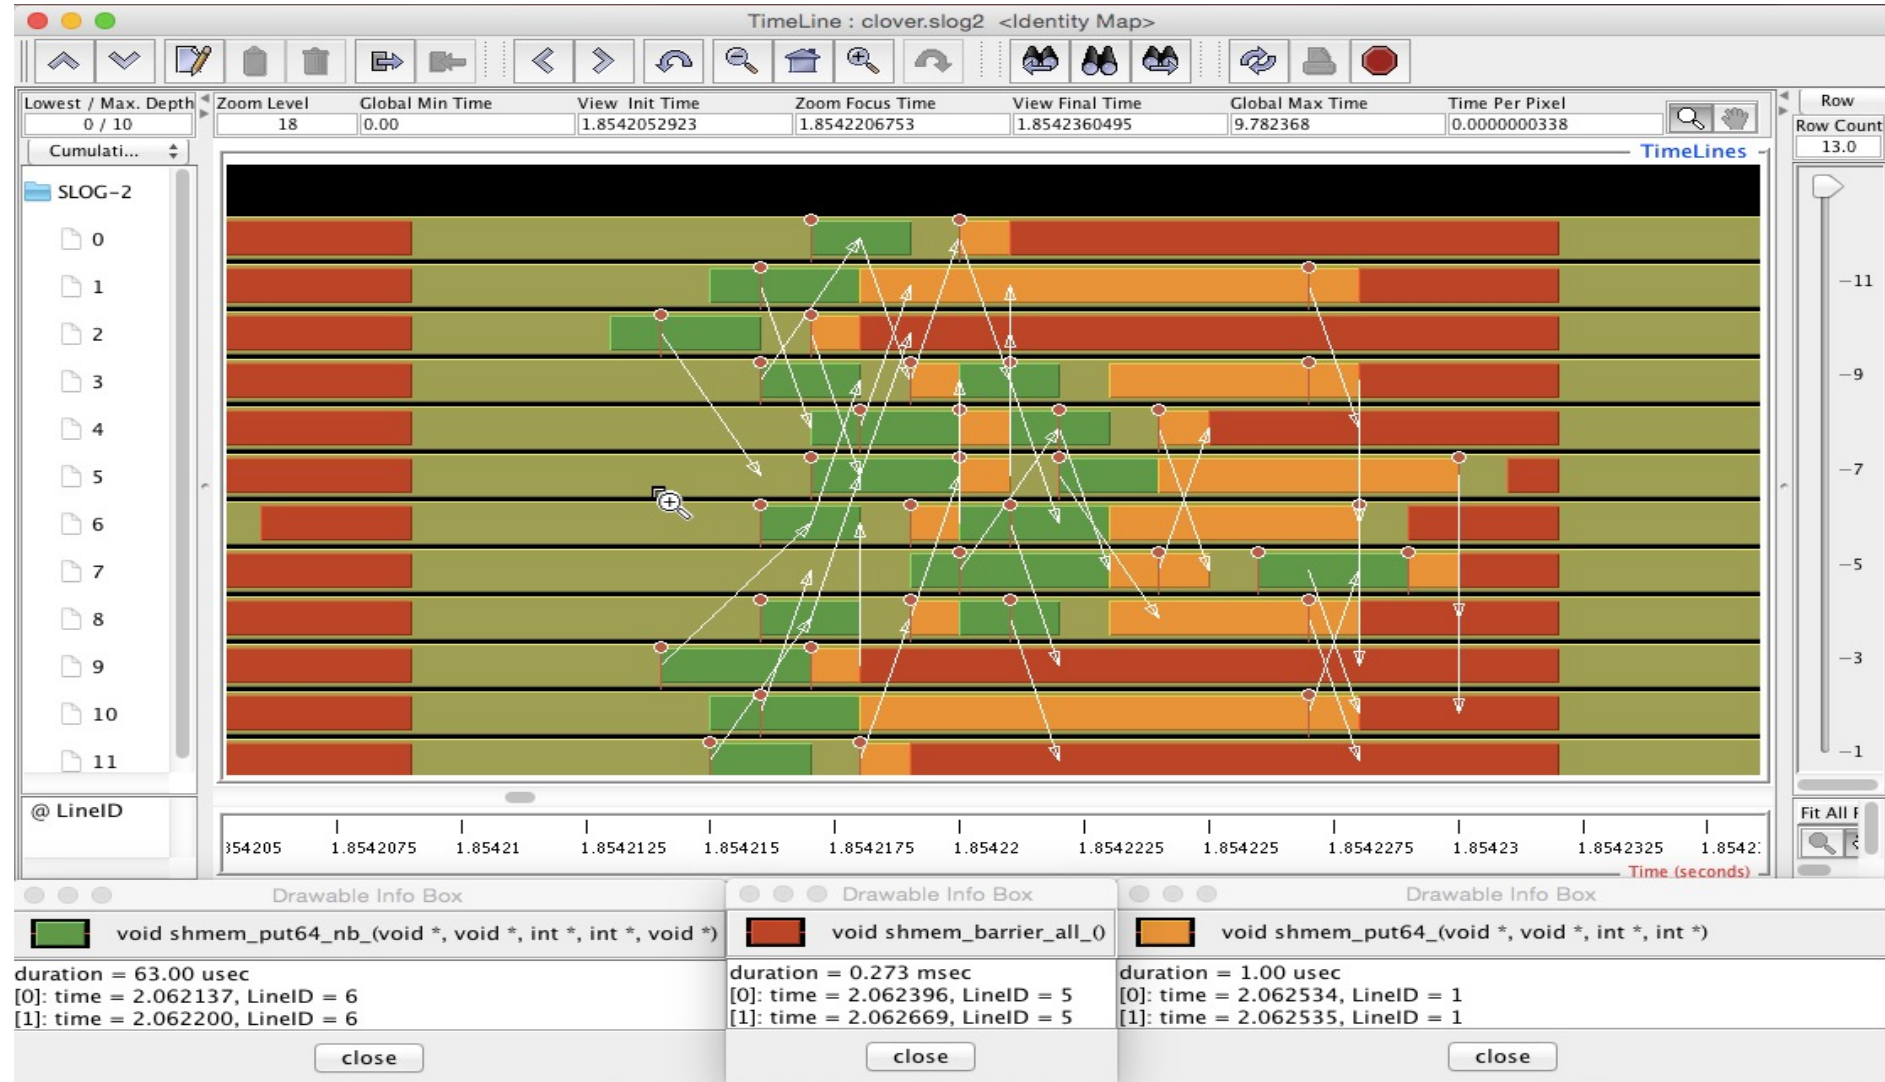
Time spent in GEMM kernel

TAU and Vampir [TU Dresden]: Intel oneAPI OpenCL

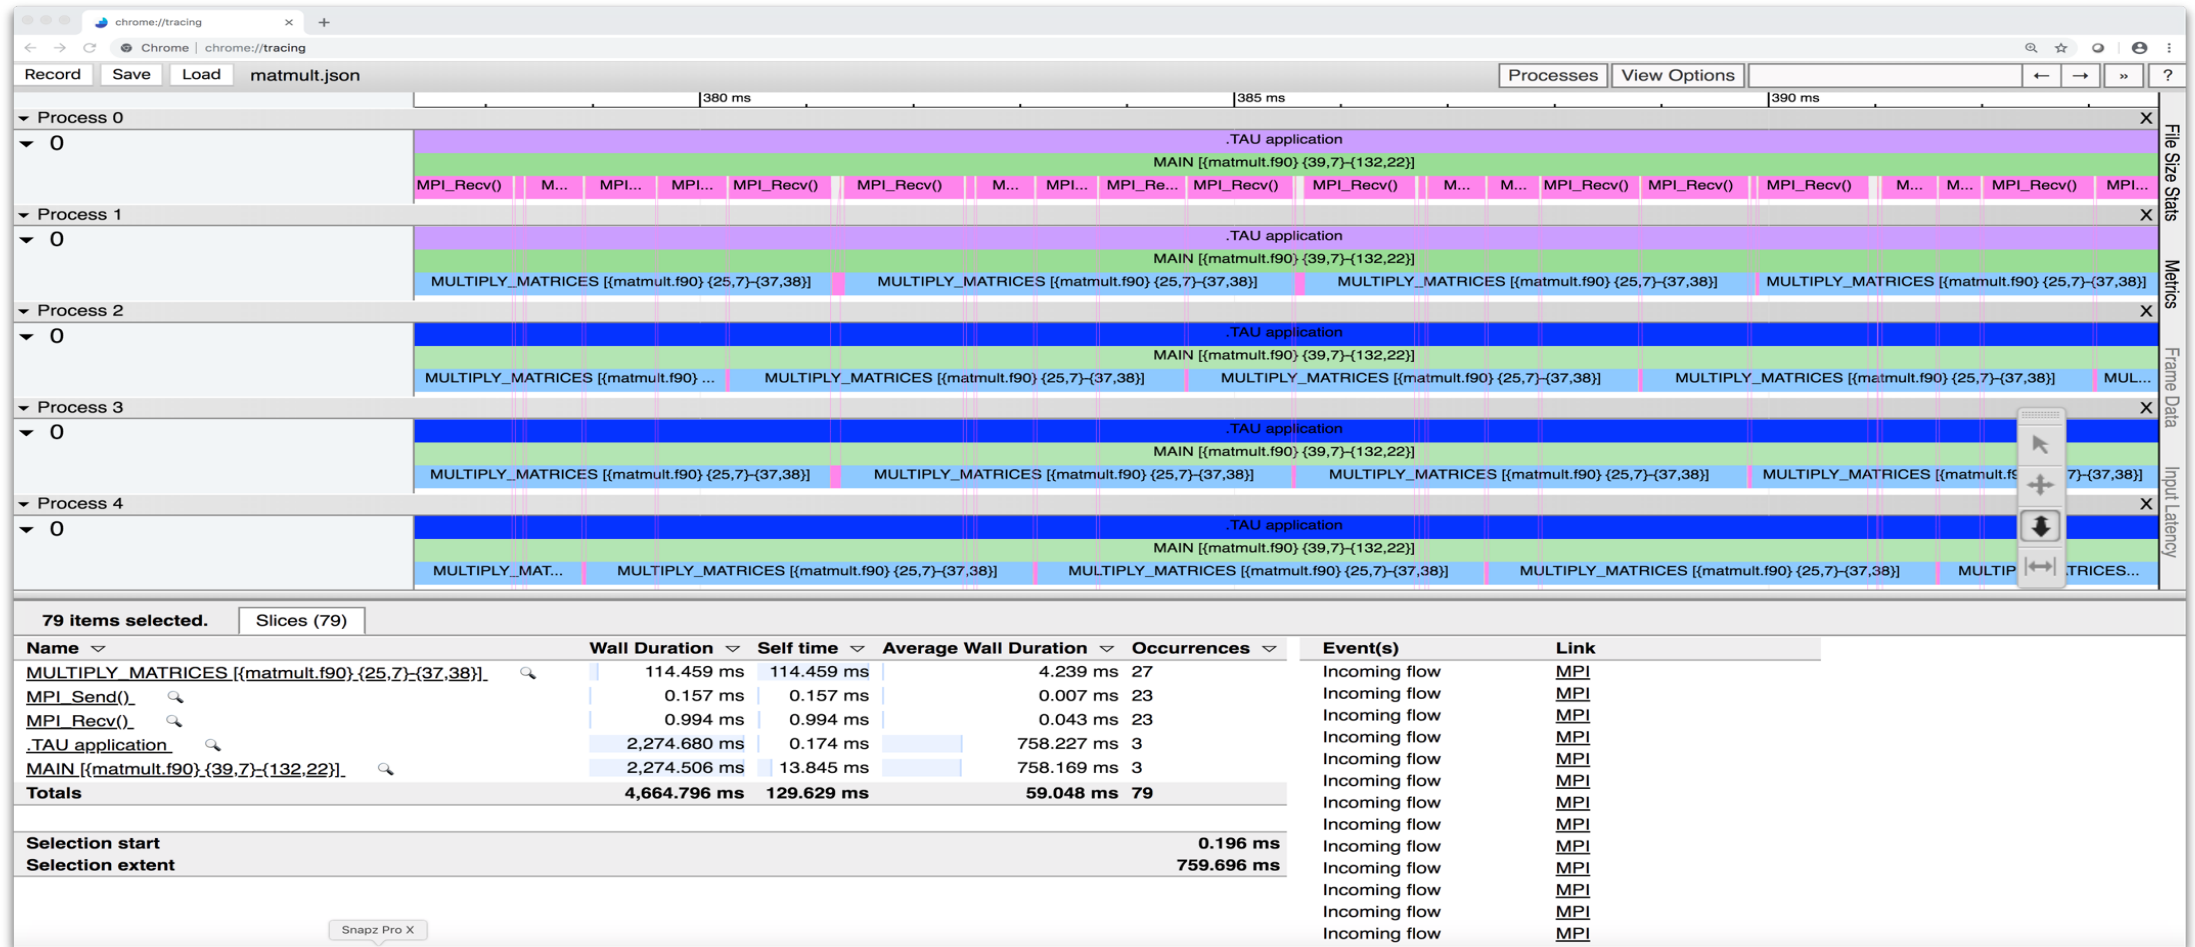


```
% export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2
% mpirun -np 4 tau_exec -T level_zero -opencl ./a.out
```


Tracing: Jumpshot (ships with TAU)



Tracing: Chrome Browser or Perfetto.dev



% export TAU_TRACE=1

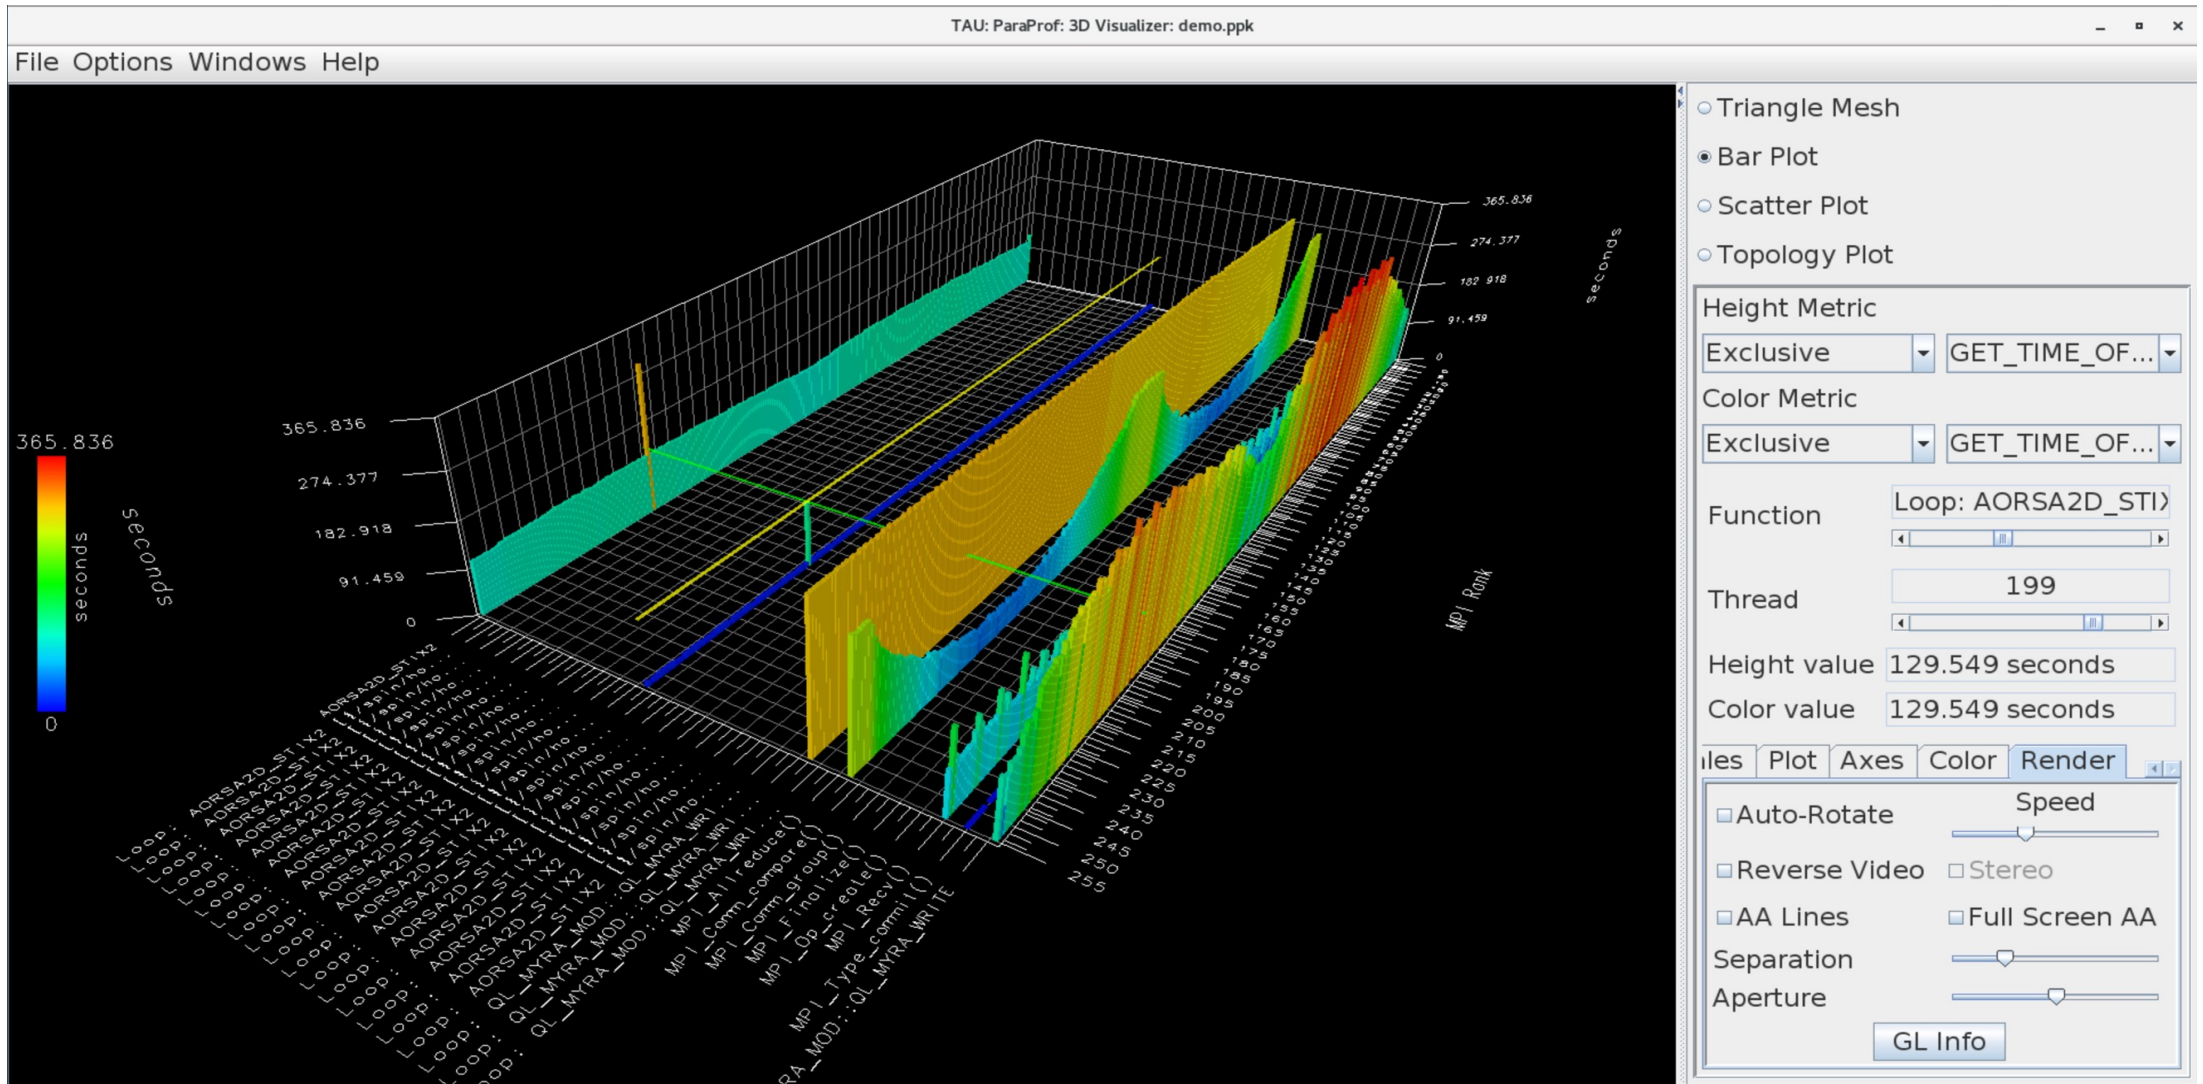
% mpirun -np 256 tau_exec ./a.out

% tau_treemerge.pl; tau_trace2json tau.trc tau.edf -chrome -ignoreatomic -o app.json

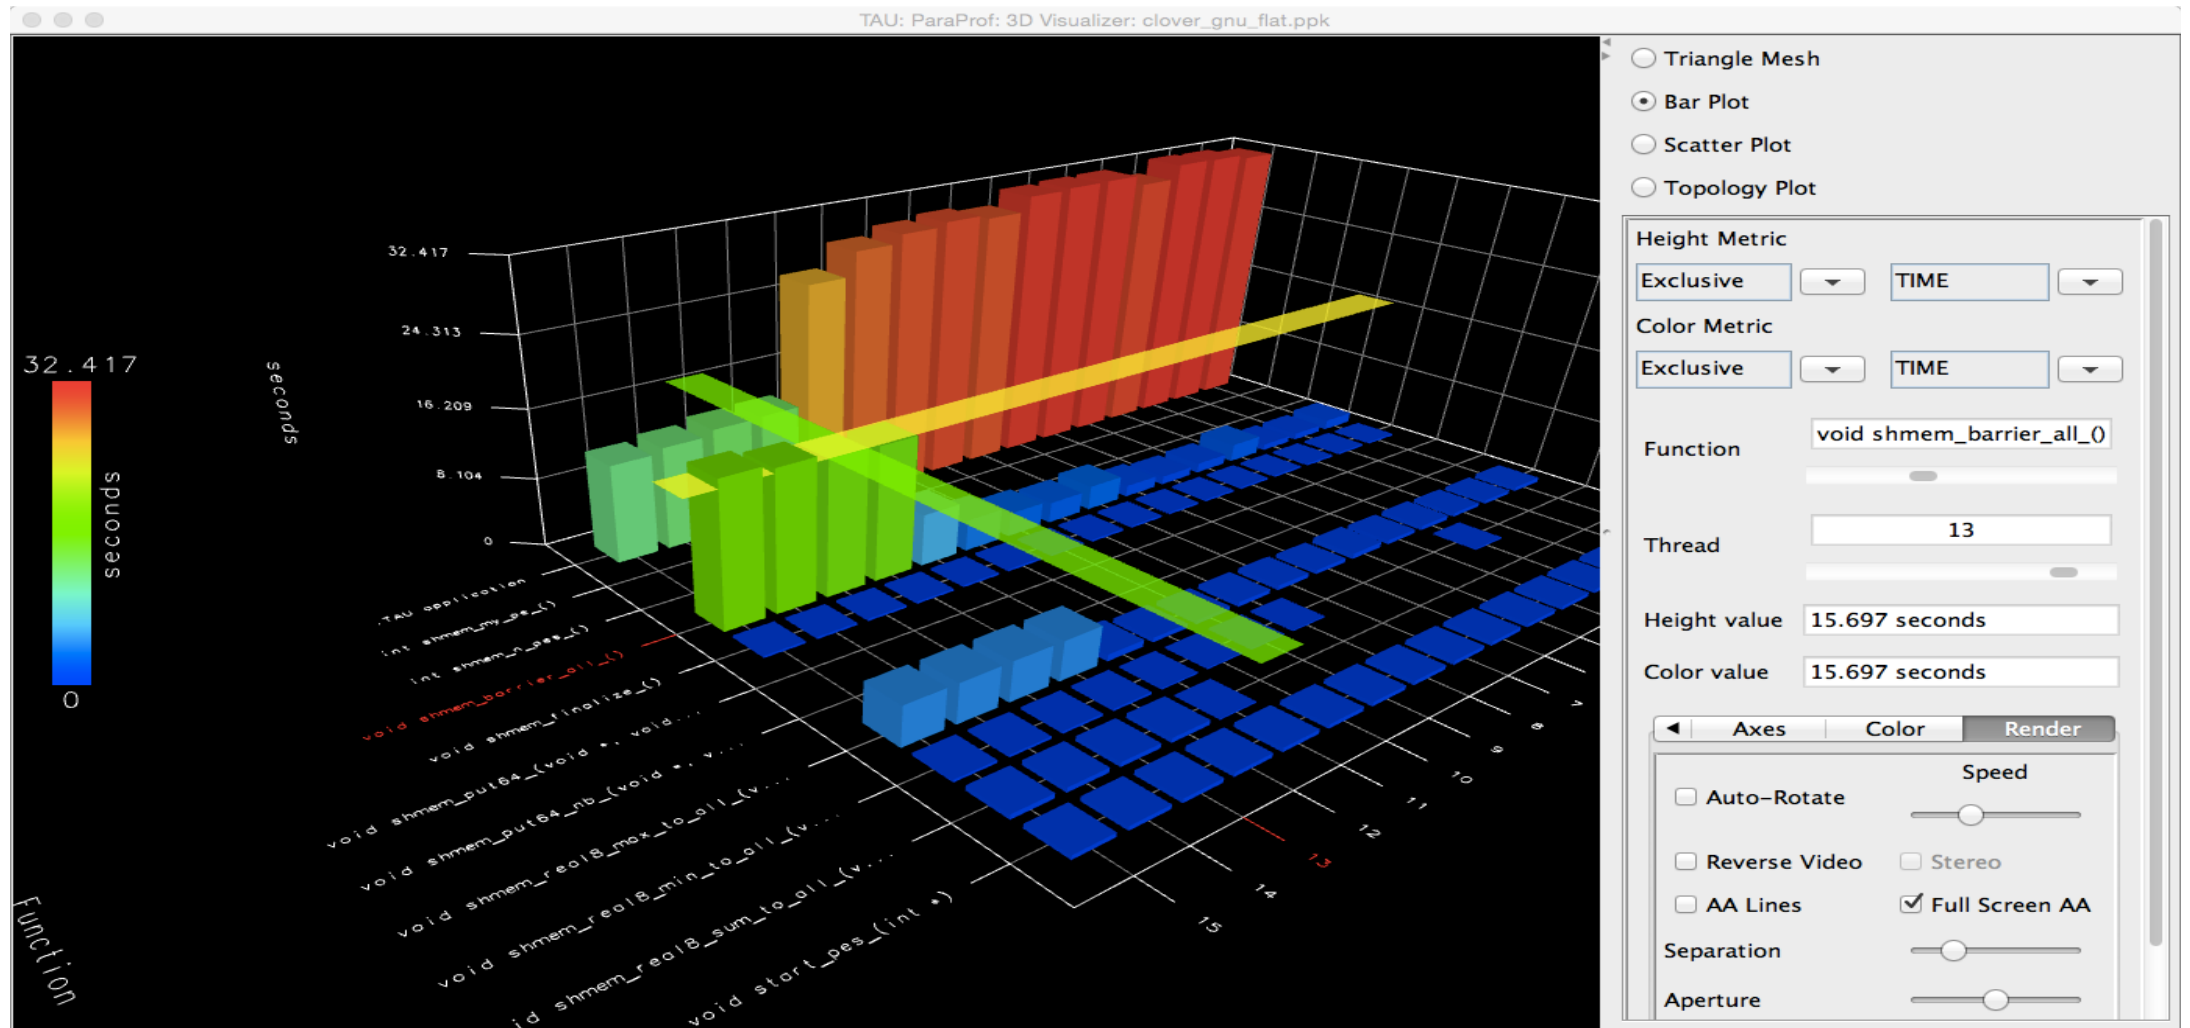
Chrome browser: chrome://tracing (Load -> app.json)

Perfetto.dev (open the UI)

ParaProf 3D Profile Browser




TAU – ParaProf 3D Visualization



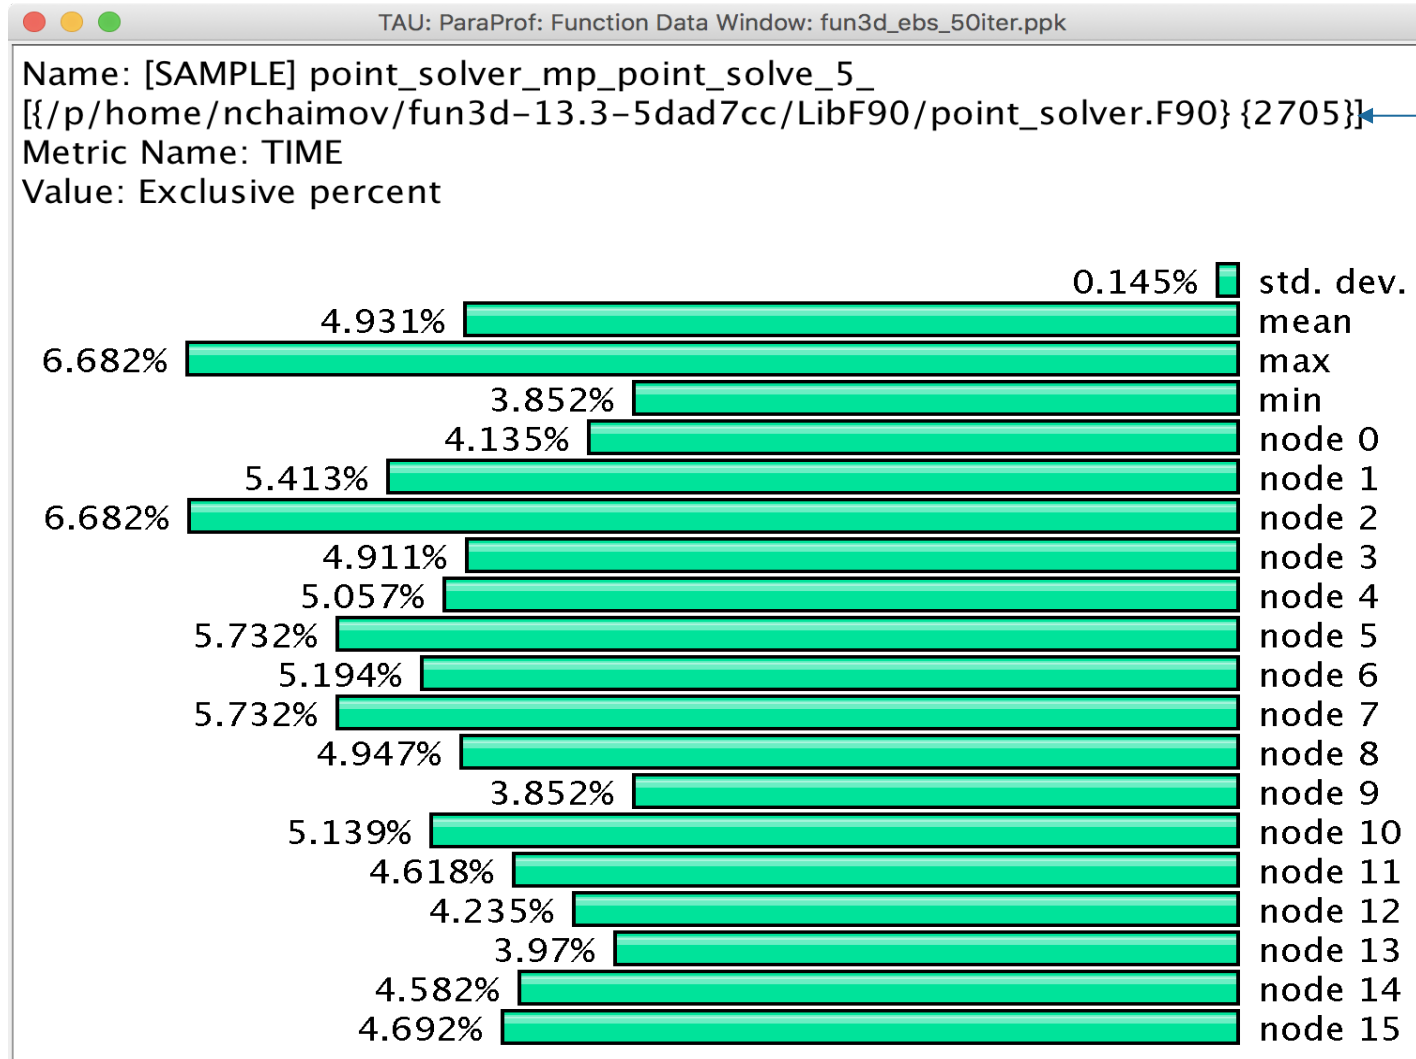
% paraprof app.ppk
Windows -> 3D Visualization -> Bar Plot (right pane)

TAU – 3D Communication Window



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


Event Based Sampling (EBS)

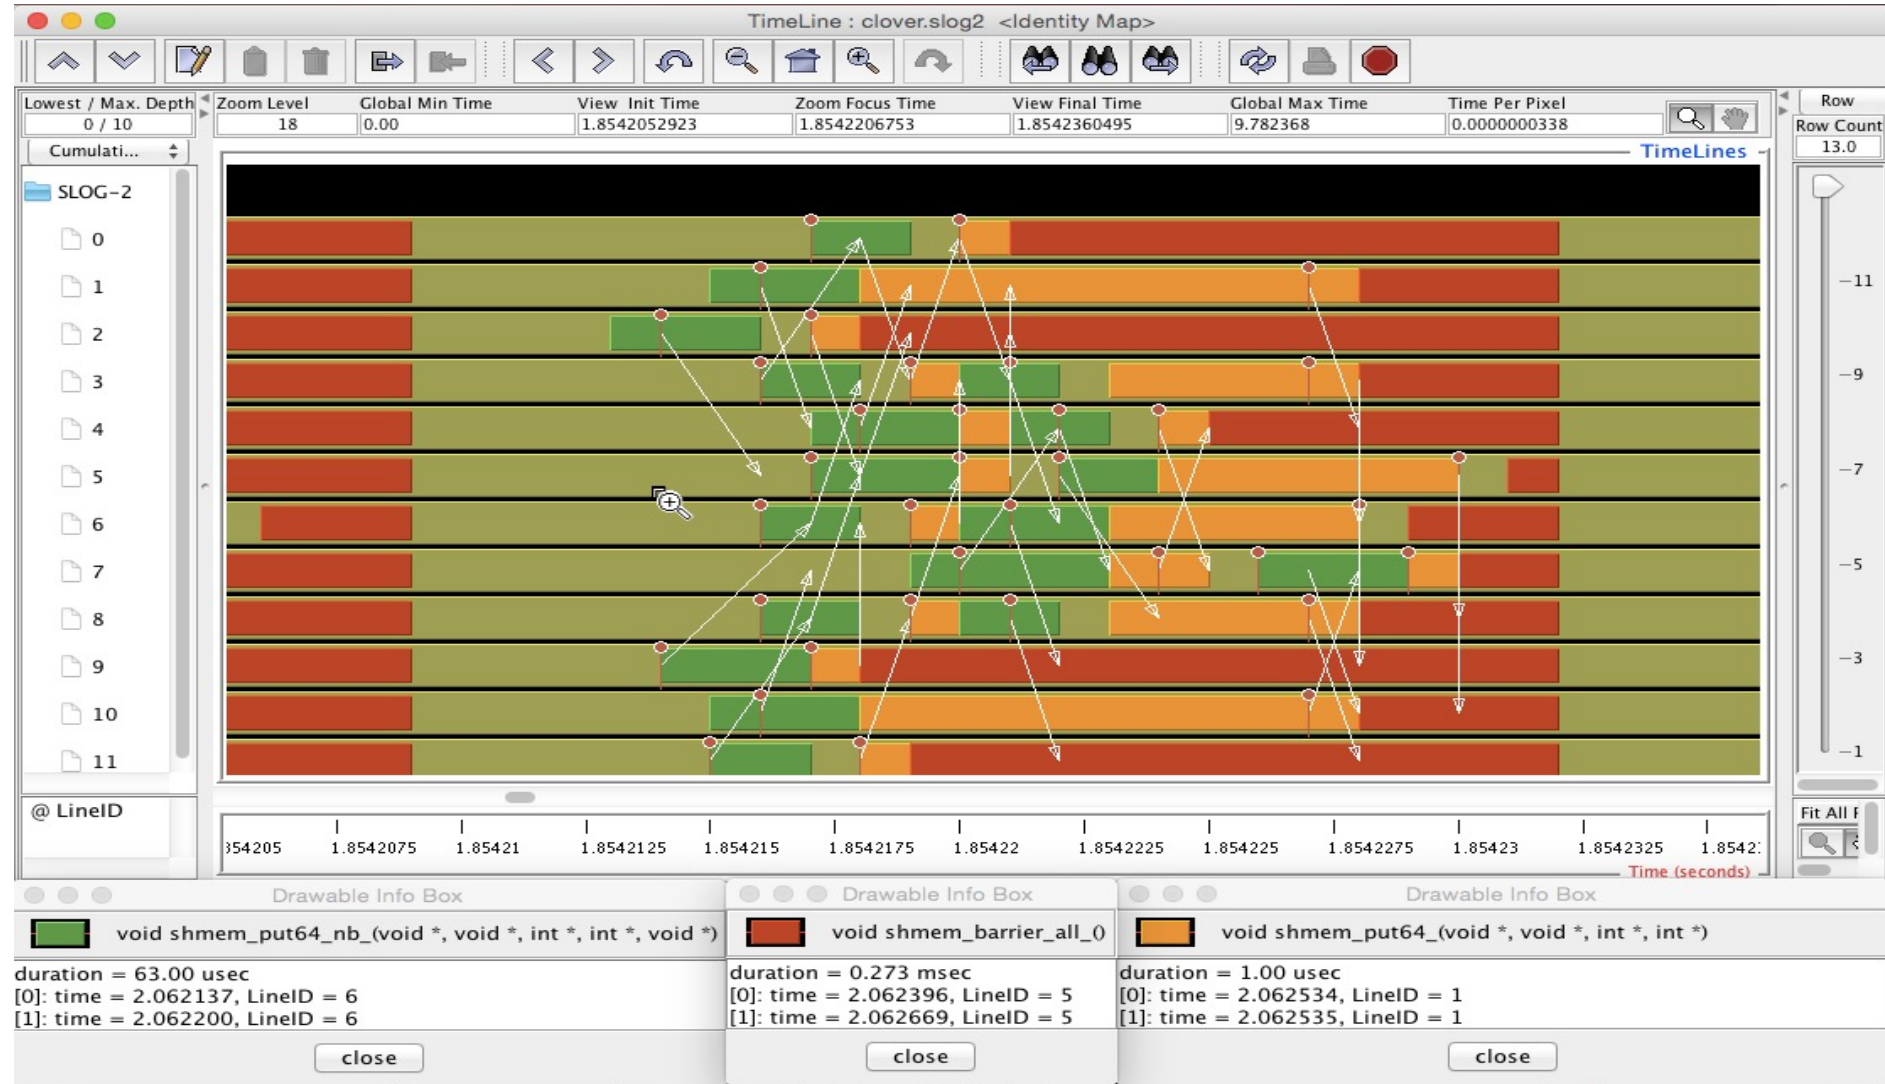


File: point_solver.F90
Line: 2705


Uninstrumented!

% mpirun -n 16 tau_exec -ebs a.out

Tracing: Jumpshot (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)

Vampir [TU Dresden] Timeline: Kokkos



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


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

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

Level Zero (DPC++/SYCL Intel oneAPI)

- Track execution of kernels on GPU
- Track time spent in level zero runtime system calls

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

Level zero + MPI

Intel oneAPI level zero (DPC++) + PMPI MPI interface

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

Installing TAU with support for OpenCL for NVIDIA, Intel, AMD GPUs

```
% wget http://tau.uoregon.edu/tau.tgz; tar xf tau.tgz; cd tau-<version>;  
% ./configure -bfd=download -cuda=/packages/cuda/11.4.1 -tag=nvidia ; make install -j  
% ./configure -bfd=download -opencl -tag=intel; make install -j  
% ./configure -bfd=download -opencl=/opt/rocm-5.1.0/opencl -tag=amd; make install -j  
% ls x86_64/lib/Makefile*  
Makefile.tau-amd  
Makefile.tau-intel  
Makefile.tau-nvidia-cupti  
  
% cd examples/gpu/opencl; make;  
% tau_exec -T serial,[amd,intel,nvidia] -opencl ./matmult  
% pprof -a | more  
% paraprof &
```


Using TAU's Runtime Preloading Tool: tau_exec

- Preload a wrapper that intercepts the runtime system call and substitutes with another
 - **MPI**
 - **OpenMP**
 - **OpenCL**
 - **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

```
% mpirun -np 256 ./a.out
```

Track GPU operations

```
% mpirun -np 256 tau_exec -cupti ./a.out
```

```
% mpirun -np 256 tau_exec -opencl ./a.out
```

```
% mpirun -np 256 tau_exec -rocm ./a.out
```

```
% mpirun -np 256 tau_exec -openacc ./a.out
```

Track MPI performance

```
% mpirun -np 256 tau_exec ./a.out
```

Track I/O, and MPI performance (MPI enabled by default)

```
% mpirun -np 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;
```

```
% mpirun -np 256 tau_exec -T ompt,v5,mpi -ompt ./a.out
```

Track memory operations

```
% export TAU_TRACK_MEMORY_LEAKS=1
```

```
% mpirun -np 256 tau_exec -memory_debug ./a.out (bounds check)
```

Use event based sampling (compile with -g)

```
% mpirun -np 256 tau_exec -ebs ./a.out
```

```
Also -ebs_source=<PAPI_COUNTER> -ebs_period=<overflow_count> -ebs_resolution=<file | function | line>
```


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 <code>-optMemDbg</code> or <code>tau_exec -memory</code>)
TAU_EBS_SOURCE	TIME	Allows using PAPI hardware counters for periodic interrupts for EBS (e.g., <code>TAU_EBS_SOURCE=PAPI_TOT_INS</code> when <code>TAU_SAMPLING=1</code>)
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 <code>TAU_MEMDBG_PROTECT_*</code>)
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 <code>-optMemDbg</code> while building or <code>tau_exec -memory</code>)
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 <code>-optMemDbg</code> or <code>tau_exec -memory</code>)
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


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
- Industry: AMD, ARM, Intel, IBM, NVIDIA
- Partners:
 - University of Oregon
 - The Ohio State University
 - ParaTools, Inc.
 - University of Tennessee, Knoxville
 - T.U. Dresden, GWT
 - Jülich Supercomputing Center



UNIVERSITY
OF OREGON

THE OHIO STATE
UNIVERSITY



ParaTools

Thank you

<https://www.exascaleproject.org>

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.



Thank you to all collaborators in the ECP and broader computational science communities. The work discussed in this presentation represents creative contributions of many people who are passionately working toward next-generation computational science.

