

Multi-Level Performance Instrumentation for Kokkos Applications using TAU

Sameer Shende, Nicholas Chaimov, Allen D. Malony
ParaTools, Inc.
{sameer,nchaimov,malony}@paratools.com

Neena Imam
Oak Ridge National Laboratory
imamm@ornl.gov

ProTools Workshop, SC19, Denver, CO
Sunday, November 17, 2019, 4pm – 4:15pm
Room 704-706

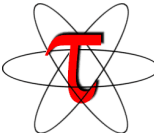
Slides:
http://tau.uoregon.edu/TAU_Kokkos_SC19.pdf

Motivation: Kokkos

<https://github.com/kokkos/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
- This talk: experience extending TAU to support Kokkos

TAU Performance System[®]



- **Tuning and Analysis Utilities (20+ year project)**
- **Comprehensive performance profiling and tracing**
 - Integrated, scalable, flexible, portable
 - Targets all parallel programming/execution paradigms
- **Integrated performance toolkit**
 - Instrumentation, measurement, analysis, visualization
 - Widely-ported performance profiling / tracing system
 - Performance data management and data mining
 - Open source (BSD-style license)
- **Easy to integrate in application frameworks**
- *<http://tau.uoregon.edu>*

Understanding Application Performance using TAU

- **How much time** is spent in each application routine and outer *loops*? Within loops, what is the contribution of each *statement*?
- **How many instructions** are executed in these code regions? Floating point, Level 1 and 2 *data cache misses*, hits, branches taken?
- **What is the memory usage** of the code? When and where is memory allocated/de-allocated? Are there any memory leaks?
- **What are the I/O characteristics** of the code? What is the peak read and write *bandwidth* of individual calls, total volume?
- **What is the extent of data transfer** between host and a GPU? In an Kokkos, OpenMP, OpenCL program.
- **What is the contribution of each *phase*** of the program? What is the time wasted/spent waiting for collectives, and I/O operations in Initialization, Computation, I/O phases?
- **How does the application scale?** What is the efficiency, runtime breakdown of performance across different core counts?

Types of Performance Profiles

Flat profiles

- Metric (e.g., time) spent in an event
- Exclusive/inclusive, # of calls, child calls, ...

Callpath profiles

- Time spent along a calling path (edges in callgraph)
- “*main=> f1 => f2 => MPI_Send*”
- Set the **TAU_CALLPATH** and **TAU_CALLPATH_DEPTH** environment variables

Callsite profiles

- Time spent along in an event at a given source location
- Set the **TAU_CALLSITE** environment variable

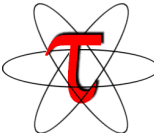
Phase profiles

- Flat profiles under a phase (nested phases allowed)
- Default “main” phase
- Supports static or dynamic (e.g. per-iteration) phases

Kokkos Profiling Interface

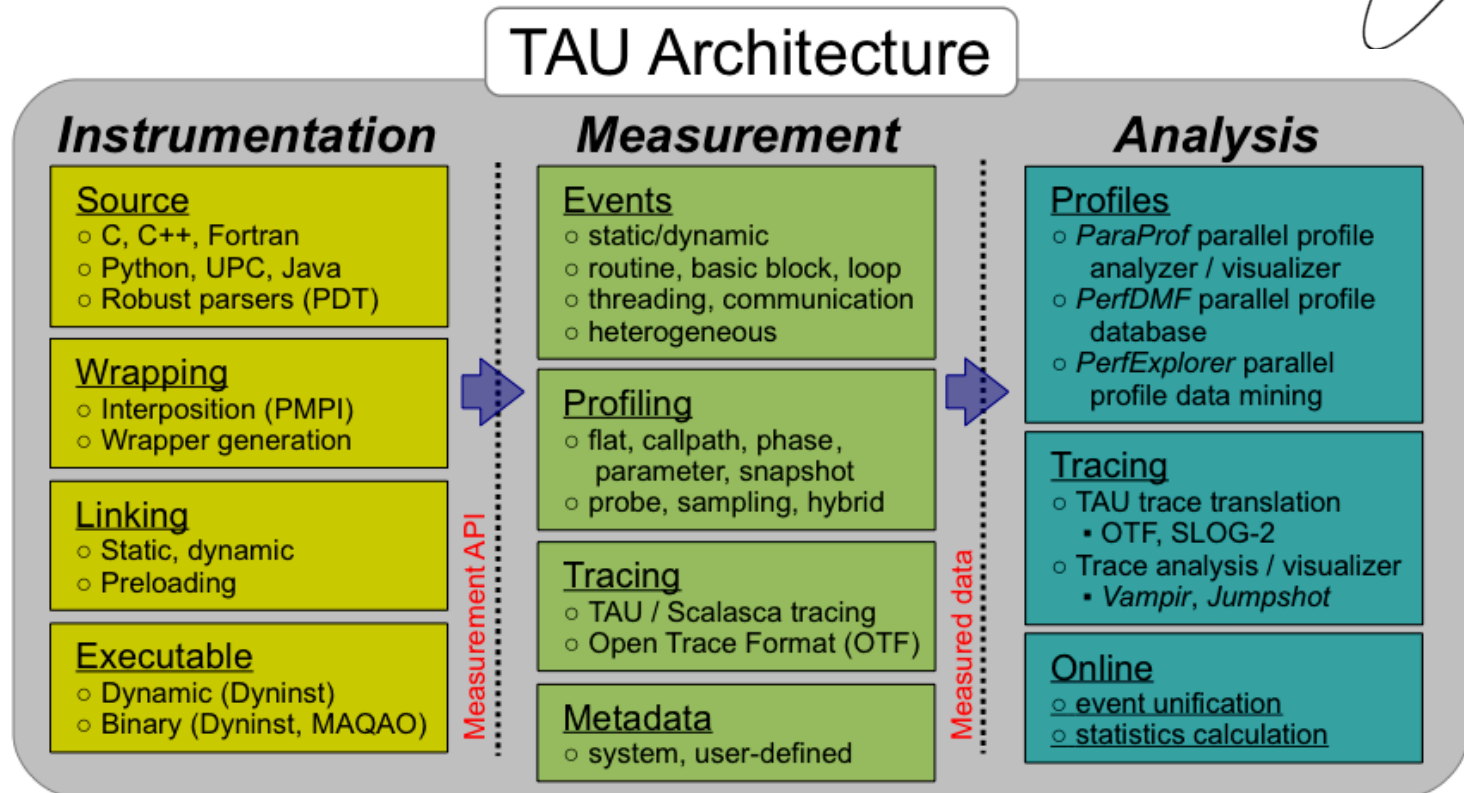
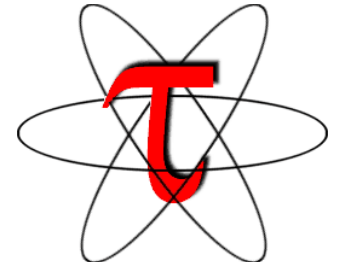
```
extern "C" void kokkosp_init_library(...);  
extern "C" void kokkosp_finalize_library();  
extern "C" void kokkosp_begin_parallel_for(const char* name,...);  
extern "C" void kokkosp_begin_parallel_reduce(const char* name, ...);  
extern "C" void kokkosp_begin_parallel_scan(const char* name, ...);  
/* corresponding end parallel constructs */  
extern "C" void kokkosp_push_profile_region(const char* name);  
extern "C" void kokkosp_pop_profile_region();  
... /* other APIs for sections, data transfers, memory allocation, ...*/
```

TAU



- Tracks kernel names specified as the first parameter in parallel API
- When name is not specified, TAU uses the template instantiation
- TAU needs to demangle mangled names of C++ entities
- TAU maps Kokkos profiling regions to TAU phases
- In a TAU phase, all functions called directly/indirectly are flattened into a flat profile under the phase
- Other runtime system calls (CUDA, Pthread, OpenMP) are also tracked alongside Kokkos calls
- Multi-level instrumentation support in TAU can help us slice through multiple runtime layers

TAU Architecture and Workflow



TAU's Support for Runtime Systems

MPI

- PMPI profiling interface
- MPI_T tools interface using performance and control variables

Kokkos

- Kokkos profiling API
- Push/pop interface for region, kernel execution interface

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

TAU's Support for Runtime Systems (contd.)

OpenCL

- OpenCL profiling interface
- Track timings of kernels

OpenACC

- OpenACC instrumentation API
- Track data transfers between host and device (per-variable)
- Track time spent in 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

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

- Python + CUPTI + pthread profiling interfaces (e.g., Tensorflow, PyTorch)

MPI + OpenCL

- PMPI + OpenCL profiling interfaces

Simplifying the use of TAU!

Uninstrumented code:

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

With TAU using event based sampling (EBS):

- `% mpirun -np 256 tau_exec -ebs ./lu.B.64`
- `% paraprof` (GUI)
- `% pprof -a | more`

NOTE:

- Requires dynamic executables (-dynamic link flag on Cray XC systems).
- Source code should be compiled with -g for access to symbol table.
- Kokkos support is on by default in tau_exec

TAU Execution Command (tau_exec)

Uninstrumented execution

- % mpirun -np 256 ./a.out

Track GPU operations

- % mpirun -np 256 tau_exec -rocm ./a.out
- % mpirun -np 256 tau_exec -cupti ./a.out
- % mpirun -np 256 tau_exec -opencl ./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)

- % export TAU_OMPT_SUPPORT_LEVEL=full;
% mpirun -np 256 tau_exec -T openmp,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>

Kokkos API use in ExaMiniMD

```
20. sameer@pegasus:~/pkgs/ORNLD/DEMO/BUILD/ExaMiniMD-pthread/ExaMiniMD/src/comm_types (ssh)
void CommMPI::update_halo() {

    Kokkos::Profiling::pushRegion("Comm::update_halo"); ← pushRegion("Comm::update_halo")

    N_ghost = 0;
    s=*system;

    pack_buffer_update = t_buffer_update((T_X_FLOAT*)pack_buffer.data(),pack_indicies_all.extent(1));
    unpack_buffer_update = t_buffer_update((T_X_FLOAT*)unpack_buffer.data(),pack_indicies_all.extent(1));

    for(phase = 0; phase<6; phase++) {
        pack_indicies = Kokkos::subview(pack_indicies_all,phase,Kokkos::ALL());
        if(proc_grid[phase/2]>1) {

            Kokkos::parallel_for("CommMPI::halo_update_pack",
                Kokkos::RangePolicy<TagHaloUpdatePack, Kokkos::IndexType<T_INT> >(0,proc_num_send[phase]),
                *this);
            MPI_Request request;
            MPI_Status status;
            MPI_Irecv(unpack_buffer.data(),proc_num_recv[phase]*sizeof(T_X_FLOAT)*3/sizeof(int),MPI_INT, proc_neighbors_recv[phase],100002,MPI_COMM_WORLD,&request);
            MPI_Send (pack_buffer.data(),proc_num_send[phase]*sizeof(T_X_FLOAT)*3/sizeof(int),MPI_INT, proc_neighbors_send[phase],100002,MPI_COMM_WORLD);
            s = *system;
            MPI_Wait(&request,&status);
            const int count = proc_num_recv[phase];
            if(unpack_buffer_update.extent(0)<count) {
                unpack_buffer_update = t_buffer_update((T_X_FLOAT*)unpack_buffer.data(),count);
            }
            Kokkos::parallel_for("CommMPI::halo_update_unpack", ← Kokkos::parallel_for
                Kokkos::RangePolicy<TagHaloUpdateUnpack, Kokkos::IndexType<T_INT> >(0,proc_num_recv[phase]),
                *this);

        } else {
            //printf("HaloUpdateCopy: %i %i %i\n",phase,proc_num_send[phase],pack_indicies.extent(0));
            Kokkos::parallel_for("CommMPI::halo_update_self",
                Kokkos::RangePolicy<TagHaloUpdateSelf, Kokkos::IndexType<T_INT> >(0,proc_num_send[phase]),
                *this);
        }
        N_ghost += proc_num_recv[phase]; ← popRegion
    }

    Kokkos::Profiling::popRegion();
};
```

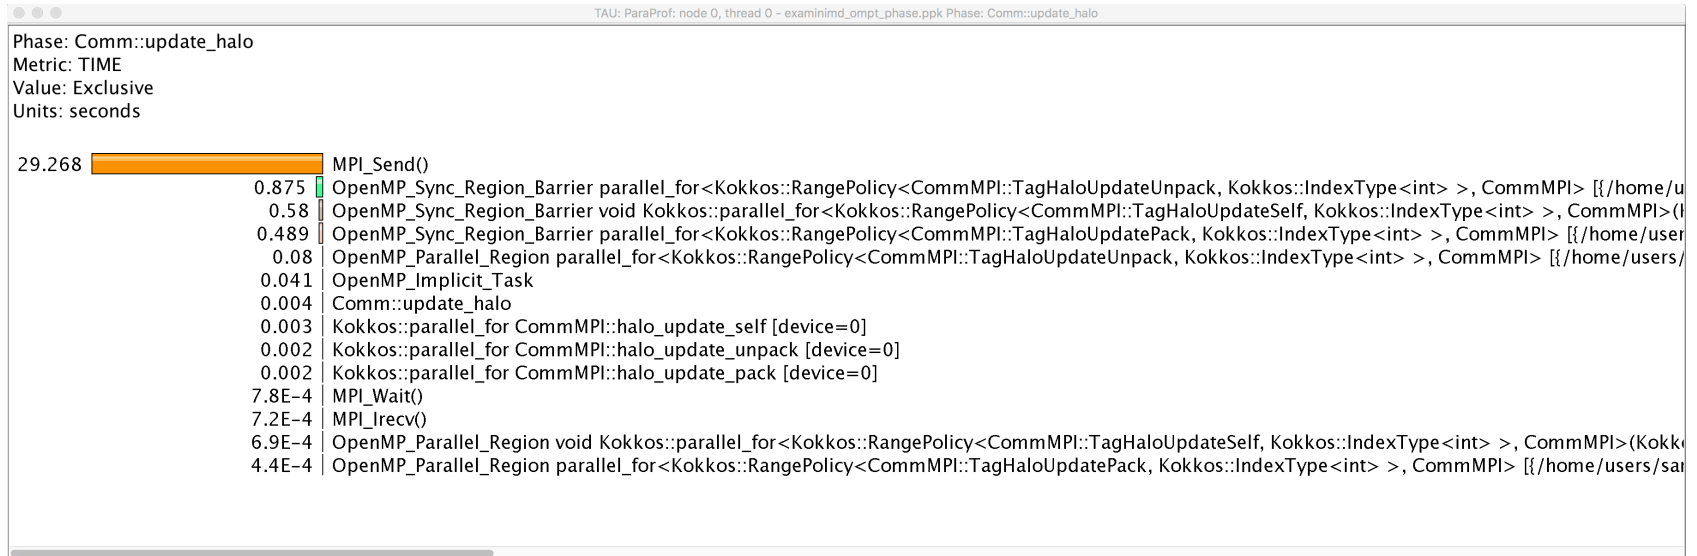
ExaMiniMD: TAU Phase

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

Name	Exclusive TIME	Inclusive TIME	Calls	Child Calls
TAU application	0.143	96.743	1	832
Comm::exchange	0.001	0.967	6	142
Comm::exchange_halo	0.001	4.702	6	184
Comm::update_halo	0.004	31.347	95	1,330
Kokkos::parallel_for CommMPI::halo_update_pack [device=0]	0.002	0.506	190	190
Kokkos::parallel_for CommMPI::halo_update_self [device=0]	0.003	0.597	380	380
Kokkos::parallel_for CommMPI::halo_update_unpack [device=0]	0.002	0.97	190	190
MPI_Irecv()	0.001	0.001	190	0
MPI_Send()	29.268	29.268	190	0
MPI_Wait()	0.001	0.001	190	0
OpenMP_Implicit_Task	0.041	1.985	760	760
OpenMP_Parallel_Region parallel_for<Kokkos::RangePolicy<CommMPI::Ta	0	0.504	190	190
OpenMP_Parallel_Region parallel_for<Kokkos::RangePolicy<CommMPI::Ta	0.08	0.968	190	190
OpenMP_Parallel_Region void Kokkos::parallel_for<Kokkos::RangePolicy<t	0.001	0.594	380	380
OpenMP_Sync_Region_Barrier parallel_for<Kokkos::RangePolicy<CommMF	0.489	0.489	190	0
OpenMP_Sync_Region_Barrier parallel_for<Kokkos::RangePolicy<CommMF	0.875	0.875	190	0
OpenMP_Sync_Region_Barrier void Kokkos::parallel_for<Kokkos::RangePol	0.58	0.58	380	0

Comm::update_halo phase in TAU ParaProf's Thread Statistics Table

ExaMiniMD: ParaProf Node Window



Event-based Sampling (EBS): CabanaMD

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

Name	Exclusive...	Inclusive...	Calls	Child Calls
TAU application	0.655	5.132	1	2,424
Comm::update_halo	0.129	1.634	95	21,755
[CONTEXT] Comm::update_halo	0	0.12	3	0
[SAMPLE] __strlen_power8 [{ } {0}]	0.09	0.09	2	0
[SAMPLE] Kokkos::Impl::SharedAllocationRecord<void, void>::increment(Kokkos::Impl::SharedAllocationRecord<void, void>*) [{/g/g20/reeve5/bin/CabanaMD}]	0.03	0.03	1	0
cudaDeviceSynchronize	0.991	0.991	3,043	0
[CONTEXT] TAU application	0	0.54	18	0
[SUMMARY] LAMMPS_RandomVelocityGeom::reset(int, double*) [{/g/g20/reeve5/pr/CabanaMD/src/input.h}]	0.27	0.27	9	0
[SAMPLE] LAMMPS_RandomVelocityGeom::reset(int, double*) [{/g/g20/reeve5/pr/CabanaMD/src/input.h} {128}]	0.09	0.09	3	0
[SAMPLE] LAMMPS_RandomVelocityGeom::reset(int, double*) [{/g/g20/reeve5/pr/CabanaMD/src/input.h} {129}]	0.09	0.09	3	0
[SAMPLE] LAMMPS_RandomVelocityGeom::reset(int, double*) [{/g/g20/reeve5/pr/CabanaMD/src/input.h} {130}]	0.06	0.06	2	0
[SAMPLE] LAMMPS_RandomVelocityGeom::reset(int, double*) [{/g/g20/reeve5/pr/CabanaMD/src/input.h} {140}]	0.03	0.03	1	0
[SUMMARY] Input::create_lattice(Comm*) [{/g/g20/reeve5/pr/CabanaMD/src/input.cpp}]	0.15	0.15	5	0
[SAMPLE] Input::create_lattice(Comm*) [{/g/g20/reeve5/pr/CabanaMD/src/input.cpp} {745}]	0.03	0.03	1	0
[SAMPLE] Input::create_lattice(Comm*) [{/g/g20/reeve5/pr/CabanaMD/src/input.cpp} {665}]	0.03	0.03	1	0
[SAMPLE] Input::create_lattice(Comm*) [{/g/g20/reeve5/pr/CabanaMD/src/input.cpp} {721}]	0.03	0.03	1	0
[SAMPLE] Input::create_lattice(Comm*) [{/g/g20/reeve5/pr/CabanaMD/src/input.cpp} {713}]	0.03	0.03	1	0
[SAMPLE] Input::create_lattice(Comm*) [{/g/g20/reeve5/pr/CabanaMD/src/input.cpp} {714}]	0.03	0.03	1	0
[SAMPLE] reference<unsigned int, unsigned int, unsigned int> [{/g/g20/reeve5/build_v100/install/kokkos/include/impl/Kokkos_ViewMapping.hpp} {2740}]	0.06	0.06	2	0
[SAMPLE] unsigned long Kokkos::Impl::ViewOffset<Kokkos::Impl::ViewDimension<0ul, 16ul, 3ul>, Kokkos::LayoutCabanaSlice<176, 16, 3, 0, 0, 0, 0>, void>::	0.03	0.03	1	0
[SUMMARY] LAMMPS_RandomVelocityGeom::uniform() [{/g/g20/reeve5/pr/CabanaMD/src/input.h}]	0.03	0.03	1	0
[SAMPLE] LAMMPS_RandomVelocityGeom::uniform() [{/g/g20/reeve5/pr/CabanaMD/src/input.h} {93}]	0.03	0.03	1	0
Comm::exchange	0.024	0.392	6	3,371
MPI_Finalize()	0.367	0.369	1	68
Comm::exchange_halo	0.026	0.351	6	4,772
MPI_Init()	0.323	0.323	1	0
Cabana::Verlet	0.004	0.256	6	438
Kokkos::parallel_for ForceLJCabanaNeigh::compute [device=0]	0.002	0.164	101	606
MPI_Allreduce()	0.082	0.082	39	0
[CONTEXT] MPI_Allreduce()	0	0.09	3	0
[SAMPLE] __GI__sched_yield [{ } {0}]	0.03	0.03	1	0
[SAMPLE] pthread_spin_unlock [{/usr/lib64/libpthread-2.17.so} {0}]	0.03	0.03	1	0
[SAMPLE] pthread_spin_lock [{/usr/lib64/libpthread-2.17.so} {0}]	0.03	0.03	1	0
Kokkos::parallel_for Kokkos::View::initialization [device=0]	0.001	0.072	35	170
Kokkos::parallel_for Kokkos::ViewFill-3D [device=0]	0.001	0.047	101	303
Kokkos::parallel_reduce ForceLJCabanaNeigh::compute_energy [device=0]	0	0.042	11	77
cudaLaunchKernel	0.015	0.028	527	1,581

Kokkos sample within Comm::update_halo

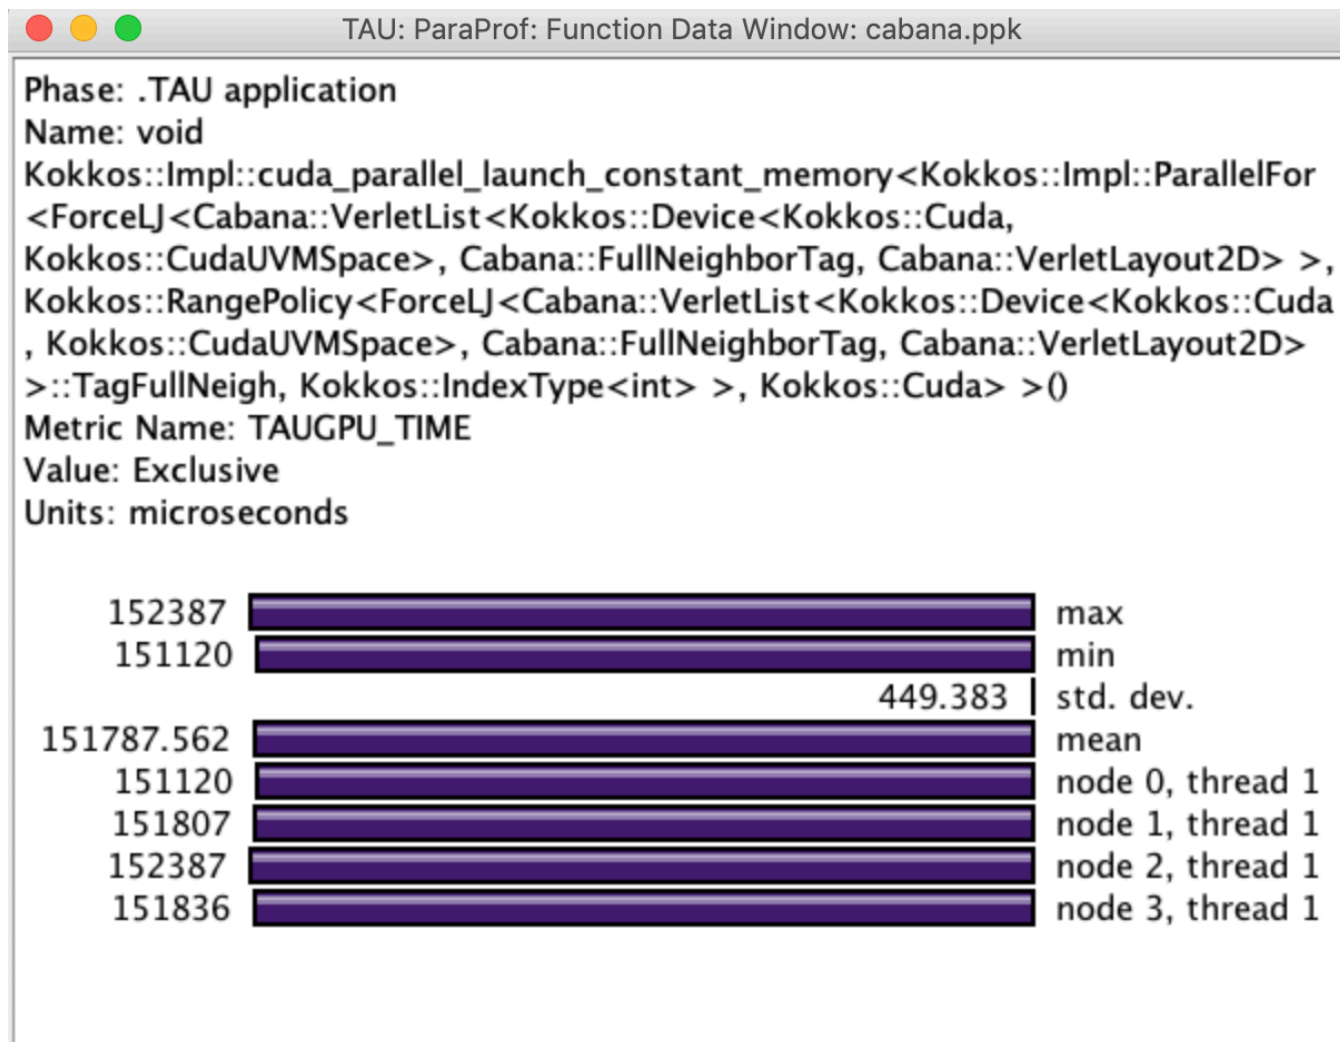
Kokkos sample within top-level application code

Instrumented Kokkos::parallel_for

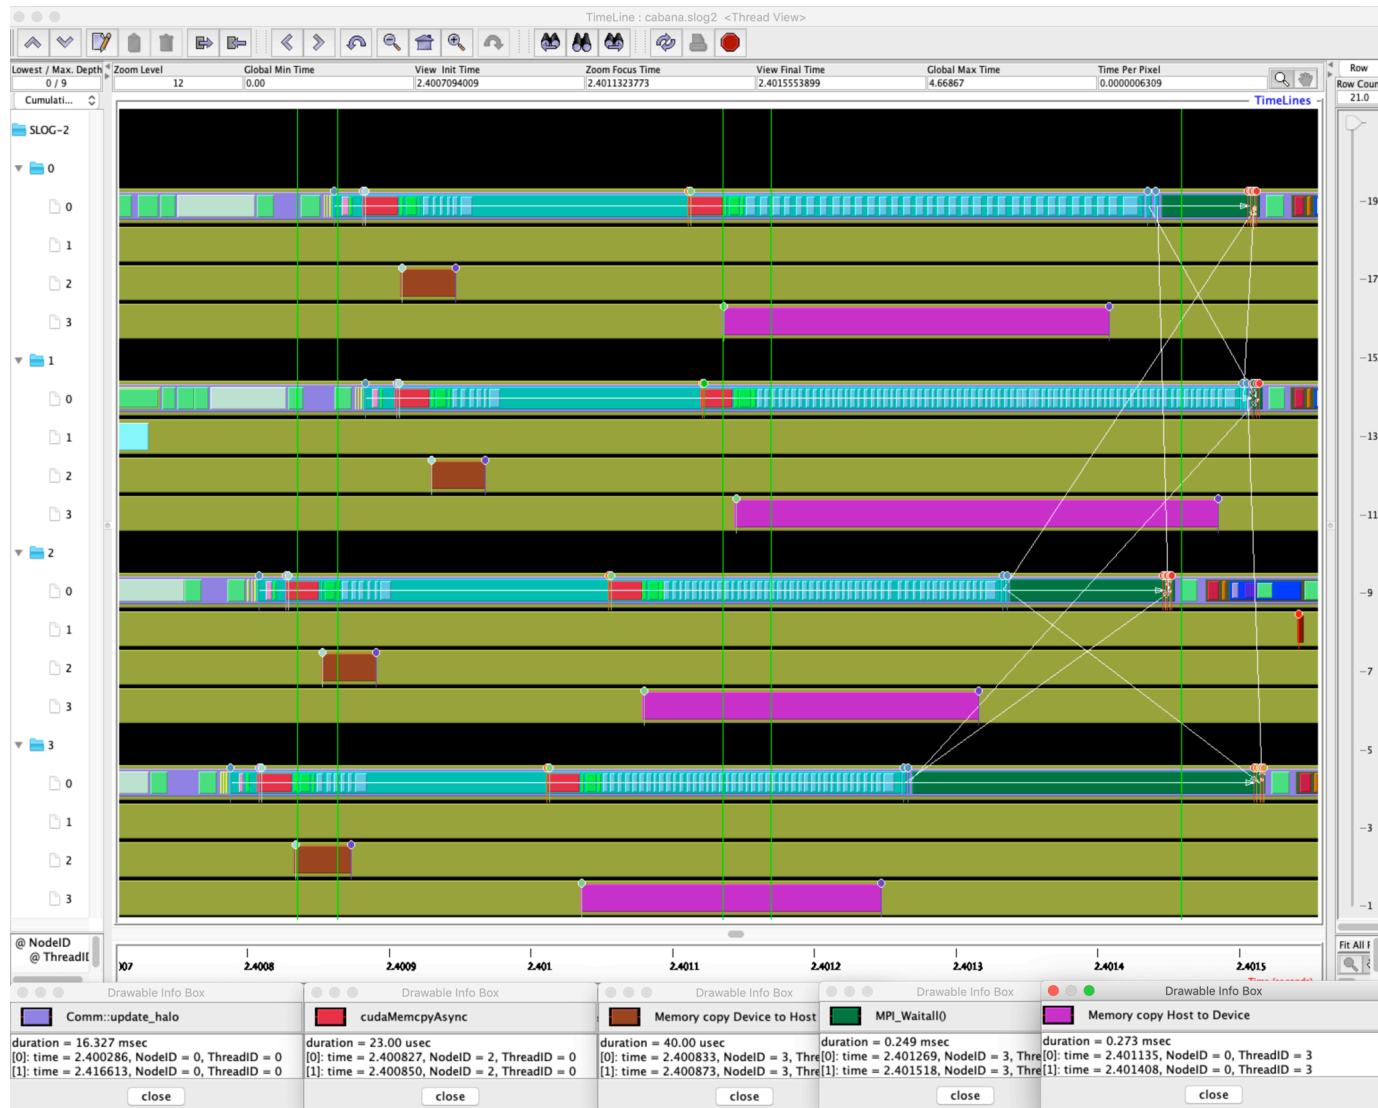
Instrumented Kokkos::parallel_reduce

EBS with Kokkos API

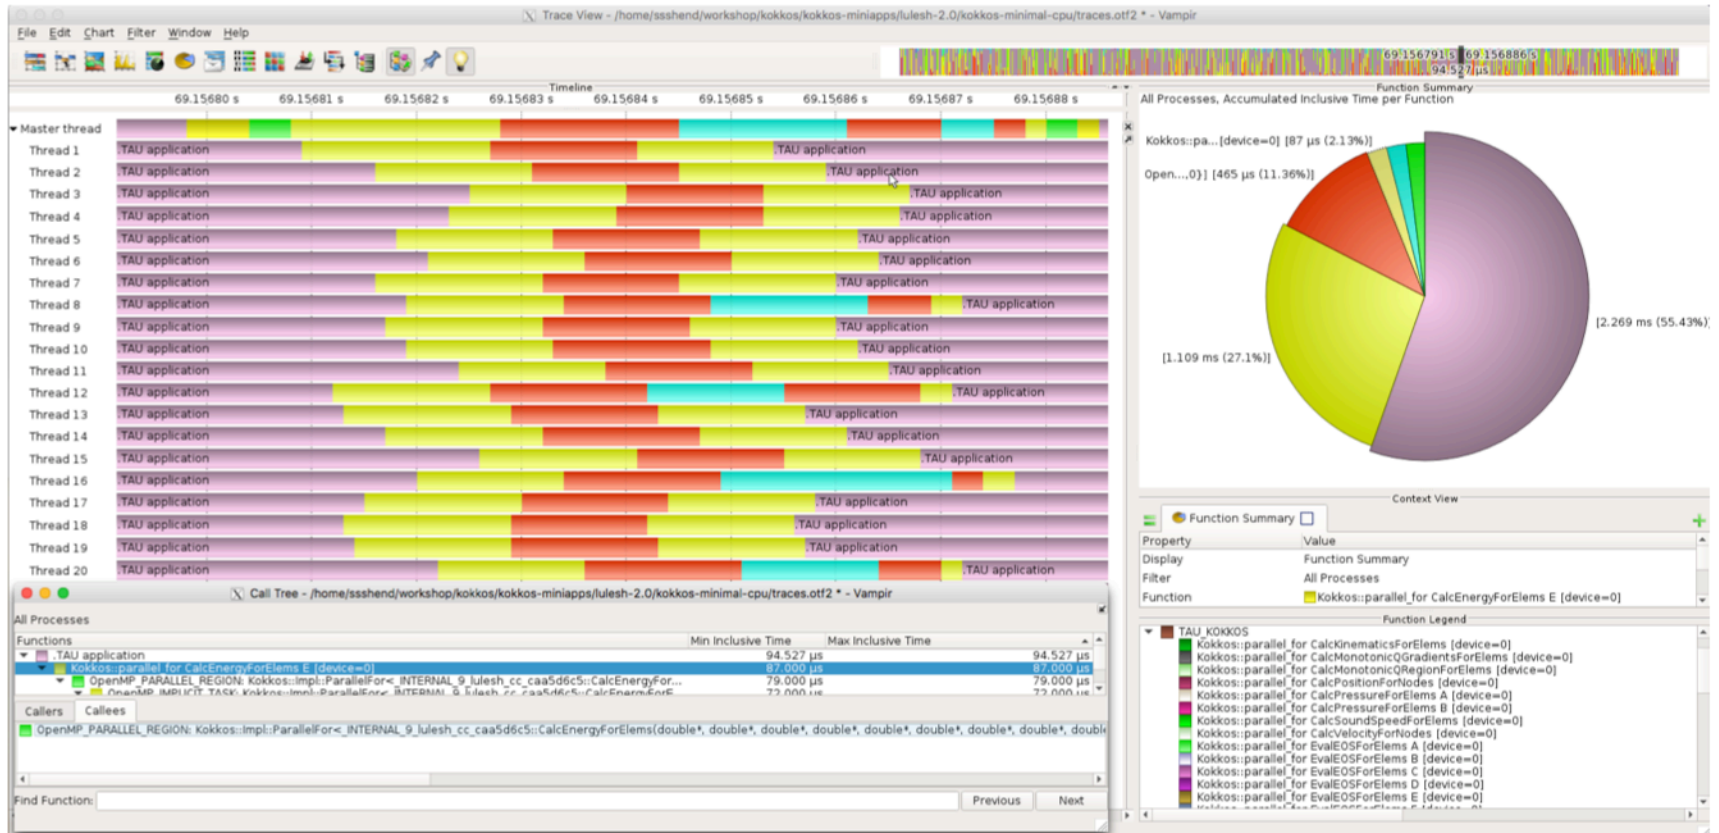
CabanaMD: CUDA Events



Jumpshot Trace Visualizer



Vampir [TU Dresden] Timeline Display



```
% export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2
% mpirun -np 16 tau_exec ./a.out
% vampir traces.otf2
```

An API model for other runtimes?

- Kokkos profiling interface is very elegant
- Kokkos API calls are defined by the tools
- Kokkos maintains compatibility between its releases
- No need for a header file, no #defines, type names
- Additional calls may be added to the API, but are not necessary for an older version of the tool to support a new version of Kokkos
- No need for tool to be compiled with a given version of Kokkos
- At startup, if KOKKOS_PROFILE_LIBRARY environment variable is defined, it loads the library, if not, profiling calls are disabled
- Profiling hooks are activated by loading the agent library
- Hooks may be disabled during configuration, enabled by default
- Is this a good model for other runtimes to adopt?
- Is there a need for a single tool to provide access to events across different runtimes/programming models?

Acknowledgment

“This work was supported by the United States Department of Defense (DoD) and used resources of the Computational Research and Development Programs, the Oak Ridge Leadership Computing Facility (OLCF) at Oak Ridge National Laboratory, and the Performance Research Laboratory at the University of Oregon. This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. This work benefited from access to the University of Oregon high performance computer, Talapas. The authors would like to thank Sam Reeve (LLNL), for his assistance with CabanaMD.”

Support Acknowledgments

US Department of Energy (DOE)

- ORNL
- Office of Science contracts, ECP
- SciDAC, LBL contracts
- LLNL-LANL-SNL ASC/NNSA contract
- Battelle, PNNL and ANL contract



US Department of Defense (DoD)

- HPCMP, ORNL

National Science Foundation (NSF)

- SI2-SSI, CSSI



NASA

CEA, France



Partners:

- University of Oregon
- The Ohio State University
- ParaTools, Inc.
- University of Tennessee, Knoxville



UNIVERSITY
OF OREGON



THE OHIO STATE
UNIVERSITY

ParaTools



ParaTools



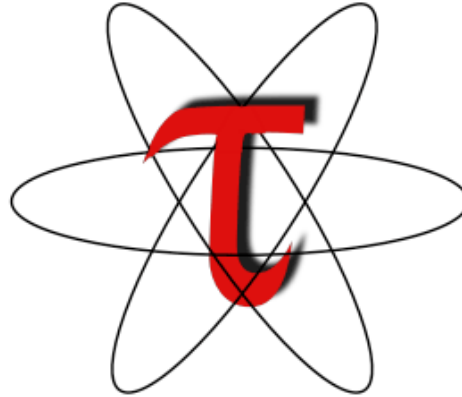
Acknowledgment



“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.”

<http://exascaleproject.org>

Download TAU



<http://tau.uoregon.edu>

<http://taucommander.com>

<https://e4s.io>

[E4S: Extreme-Scale Scientific Software Stack]

For more info on E4S: E4S BoF, Tues, 12:15pm, Room #405-407

Free download, open source, BSD license

Reference

Installing and Configuring TAU

•Installing PDT:

- `wget tau.uoregon.edu/pdt_lite.tgz`
- `./configure --prefix=<dir>; make ; make install`

•Installing TAU:

- `wget tau.uoregon.edu/tau.tgz; tar xzf tau.tgz; cd tau-2.<ver>`
- `wget http://tau.uoregon.edu/ext.tgz ; tar xf ext.tgz`
- `./configure -bfd=download -pdt=<dir>`
`-iowrapper -mpi -dwarf=download -unwind=download -`
`otf=download -papi=<dir>`
`make install`

•Using TAU:

- `export TAU_MAKEFILE=<taudir>/<arch>/lib/Makefile.tau-<TAGS>`
- `make CC=tau_cc.sh CXX=tau_cxx.sh F90=tau_f90.sh`

Compile-Time Options

Optional parameters for the TAU_OPTIONS environment variable:

% tau_compiler.sh

-optVerbose	Turn on verbose debugging messages
-optComplnst	Use compiler based instrumentation
-optNoComplnst	Do not revert to compiler instrumentation if source instrumentation fails.
-optTrackIO	Wrap POSIX I/O call and calculates vol/bw of I/O operations (Requires TAU to be configured with <i>-iowrapper</i>)
-optTrackGOMP	Enable tracking GNU OpenMP runtime layer (used without <i>-opari</i>)
-optMemDbg	Enable runtime bounds checking (see TAU_MEMDBG_* env vars)
-optKeepFiles	Does not remove intermediate .pdb and .inst.* files
-optPreProcess	Preprocess sources (OpenMP, Fortran) before instrumentation
-optTauSelectFile="<file>"	Specify selective instrumentation file for <i>tau_instrumentor</i>
-optTauWrapFile="<file>"	Specify path to <i>link_options.tau</i> generated by <i>tau_gen_wrapper</i>
-optHeaderInst	Enable Instrumentation of headers
-optTrackUPCR	Track UPC runtime layer routines (used with tau_upc.sh)
-optLinking=""	Options passed to the linker. Typically \$(TAU_MPI_FLIBS) \$(TAU_LIBS) \$(TAU_CXXLIBS)
-optCompile=""	Options passed to the compiler. Typically \$(TAU_MPI_INCLUDE) \$(TAU_INCLUDE) \$(TAU_DEFS)
-optPdtF95Opts=""	Add options for Fortran parser in PDT (f95parse/gfparse) ...

Compile-Time Options (contd.)

Optional parameters for the TAU_OPTIONS environment variable:

% tau_compiler.sh

-optShared	Use TAU's shared library (libTAU.so) instead of static library (default)
-optPdtCxxOpts=""	Options for C++ parser in PDT (cxxparse).
-optPdtF90Parser=""	Specify a different Fortran parser
-optPdtCleanscapeParser	Specify the Cleanscape Fortran parser instead of GNU gfpaser
-optTau=""	Specify options to the tau_instrumentor
-optTrackDMAPP	Enable instrumentation of low-level DMAPP API calls on Cray
-optTrackPthread	Enable instrumentation of pthread calls

See tau_compiler.sh for a full list of TAU_OPTIONS.

...

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