


The TAU Performance System[®]

Kevin Huck, Sameer Shende, Allen Malony

khuck@cs.uoregon.edu


<http://tau.uoregon.edu>



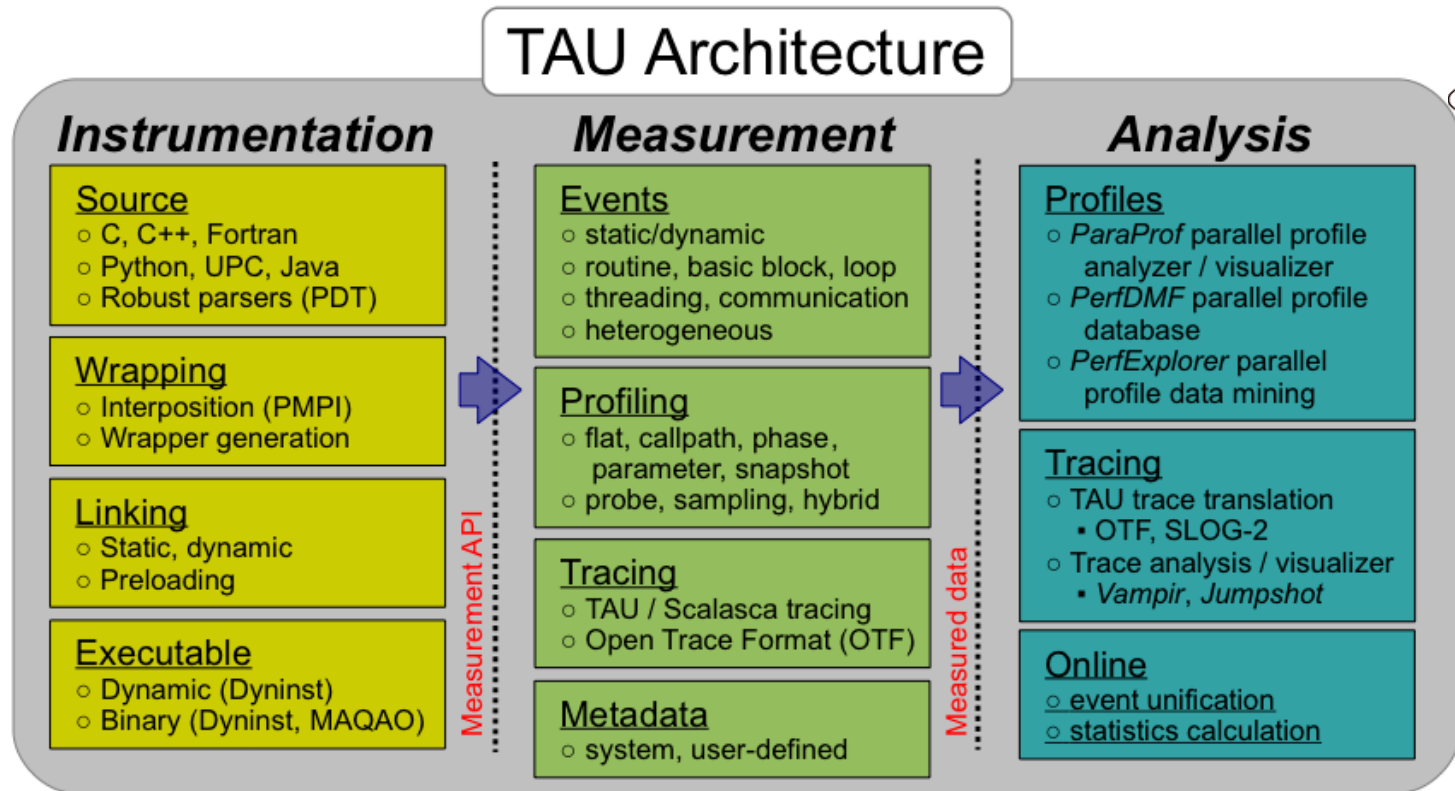

UNIVERSITY OF OREGON

TAU : brief overview

- Tuning and Analysis Utilities (28+ year project)
- Integrated performance toolkit:
 - Multi-level performance instrumentation
 - Highly configurable
 - Widely ported performance profiling / tracing system
 - Portable (java, python) visualization / exploration / analysis tools
- Supports all major HPC programming models
 - MPI/SHMEM, OpenMP/ACC, CUDA, HIP, OneAPI, Kokkos...

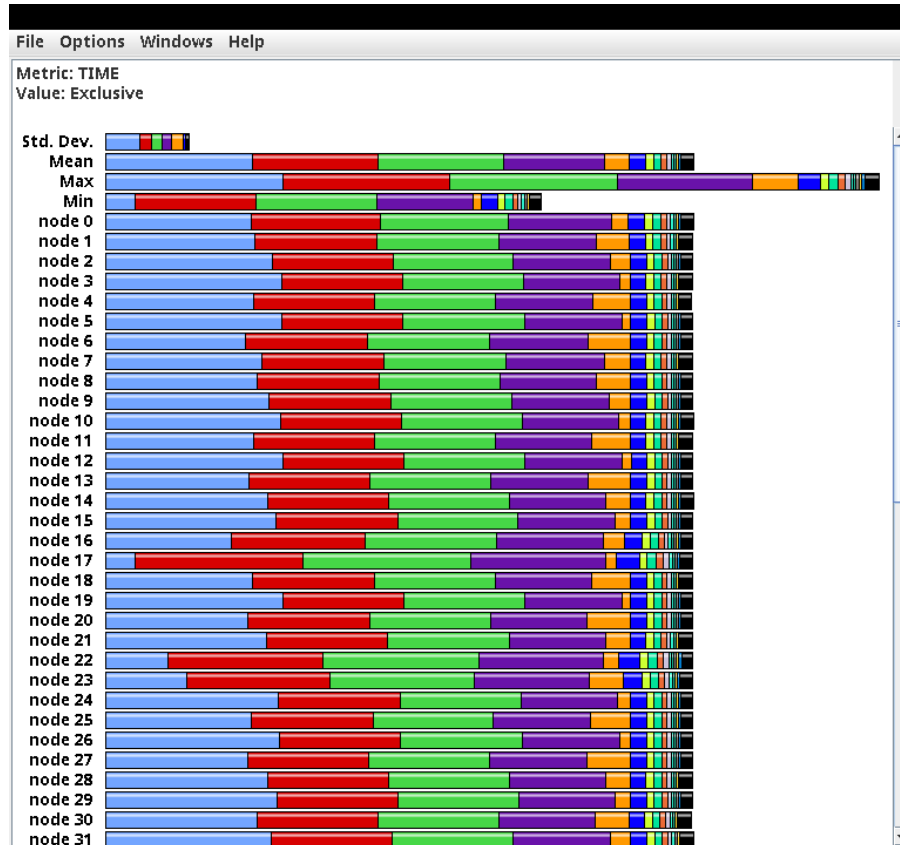


TAU : brief overview



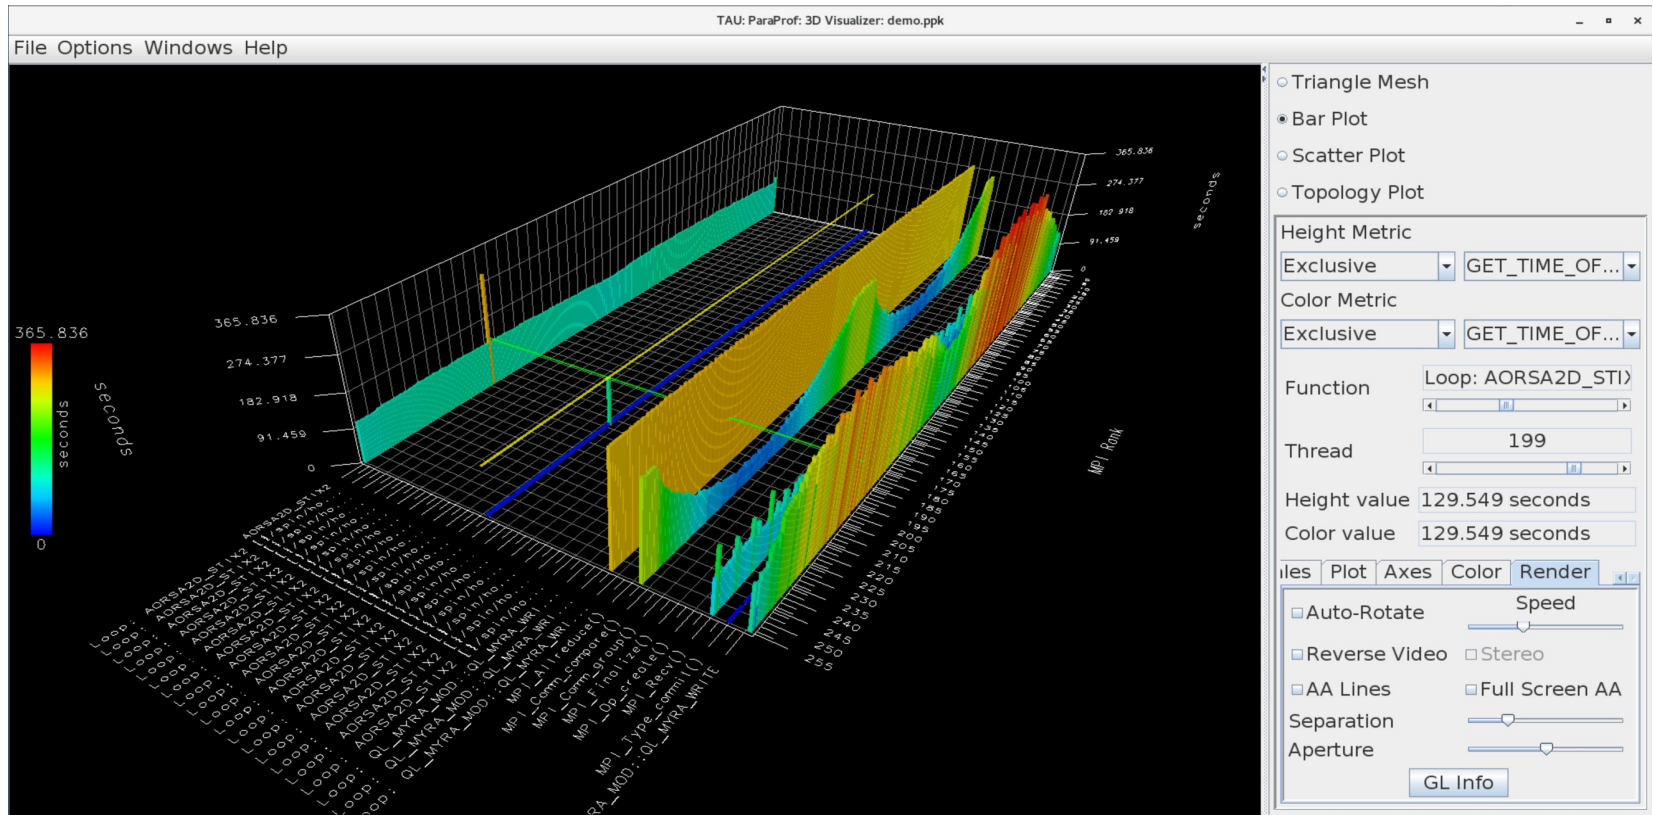
ParaProf Profile Browser

% paraprof

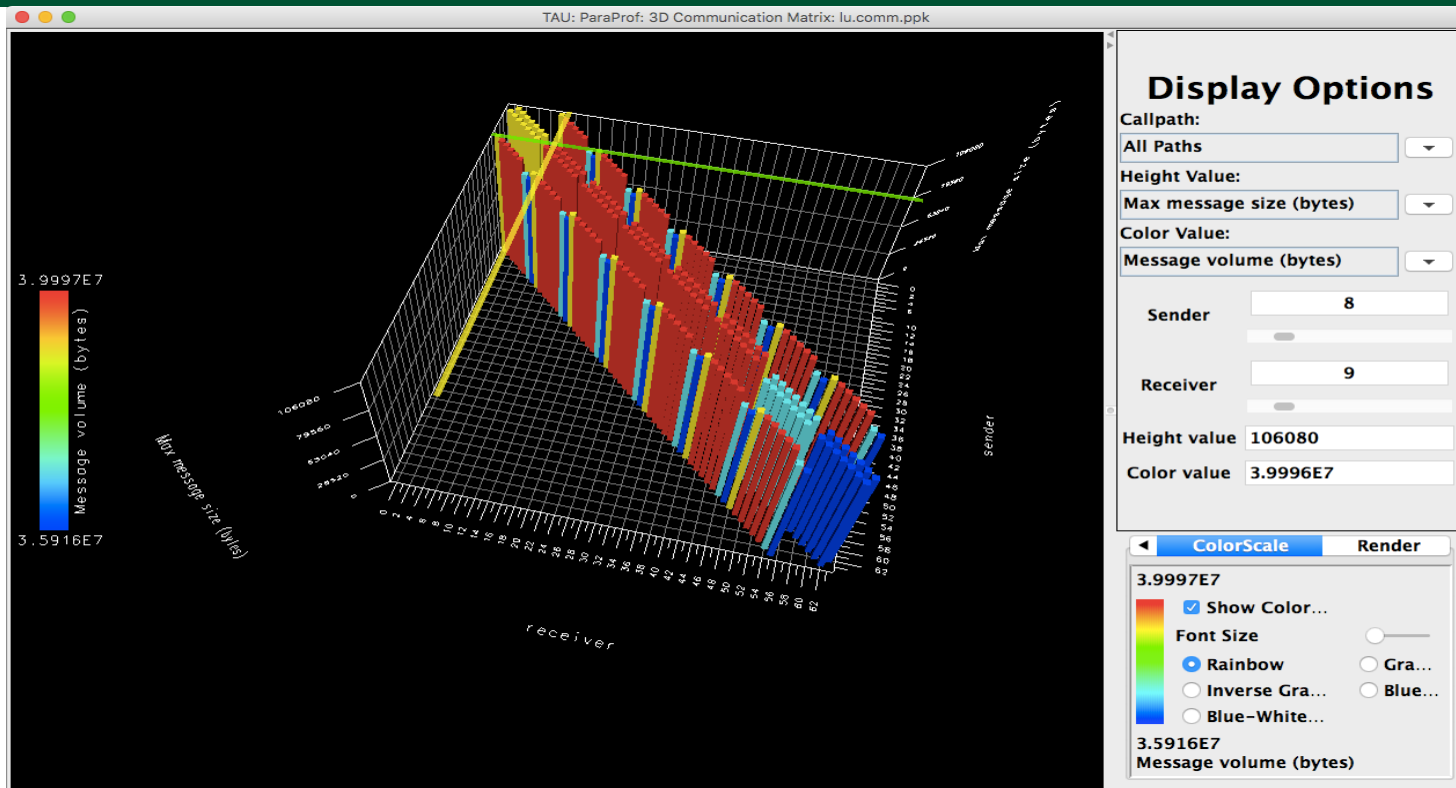


Each line is a different process/thread of execution, each color is a different function

ParaProf 3D Profile Browser

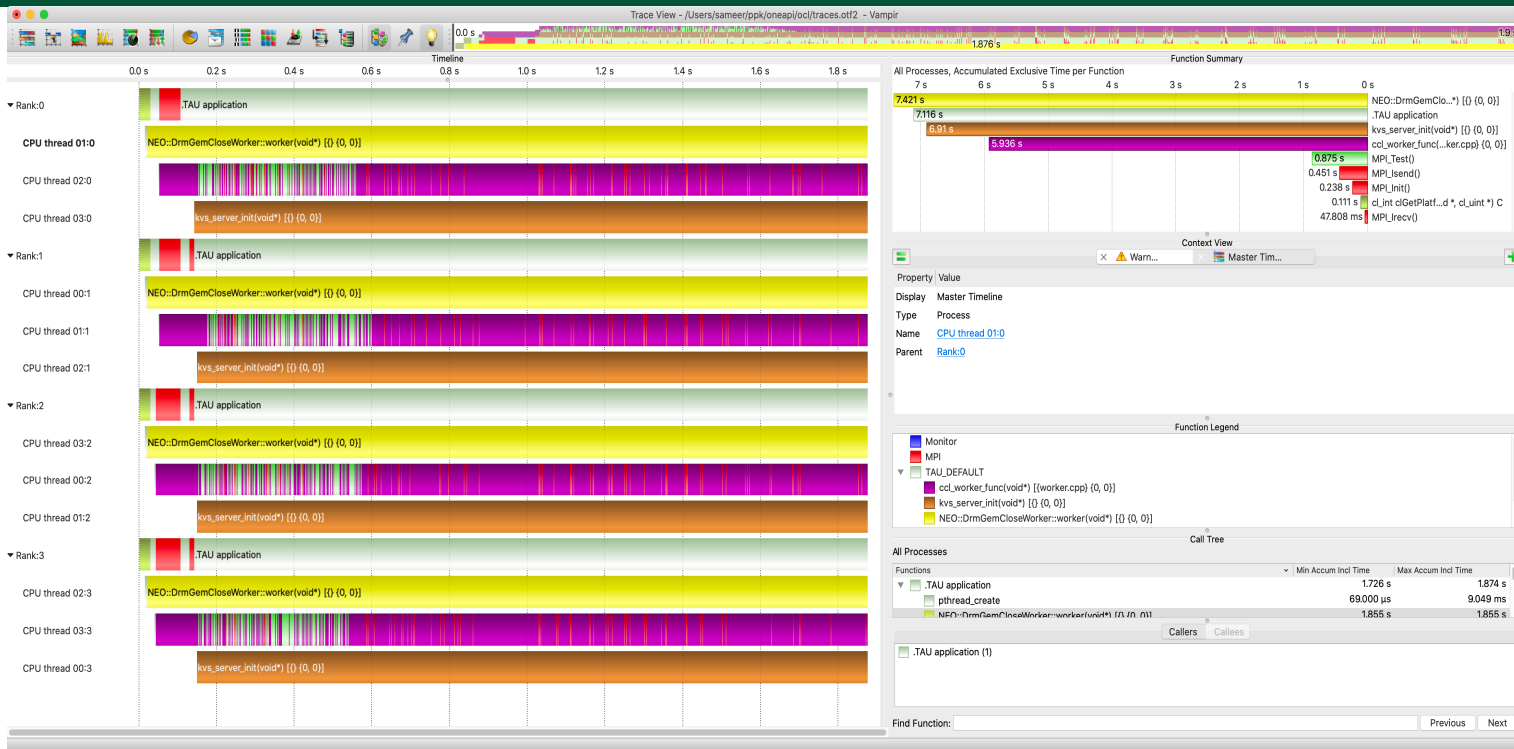


TAU – 3D Communication Window



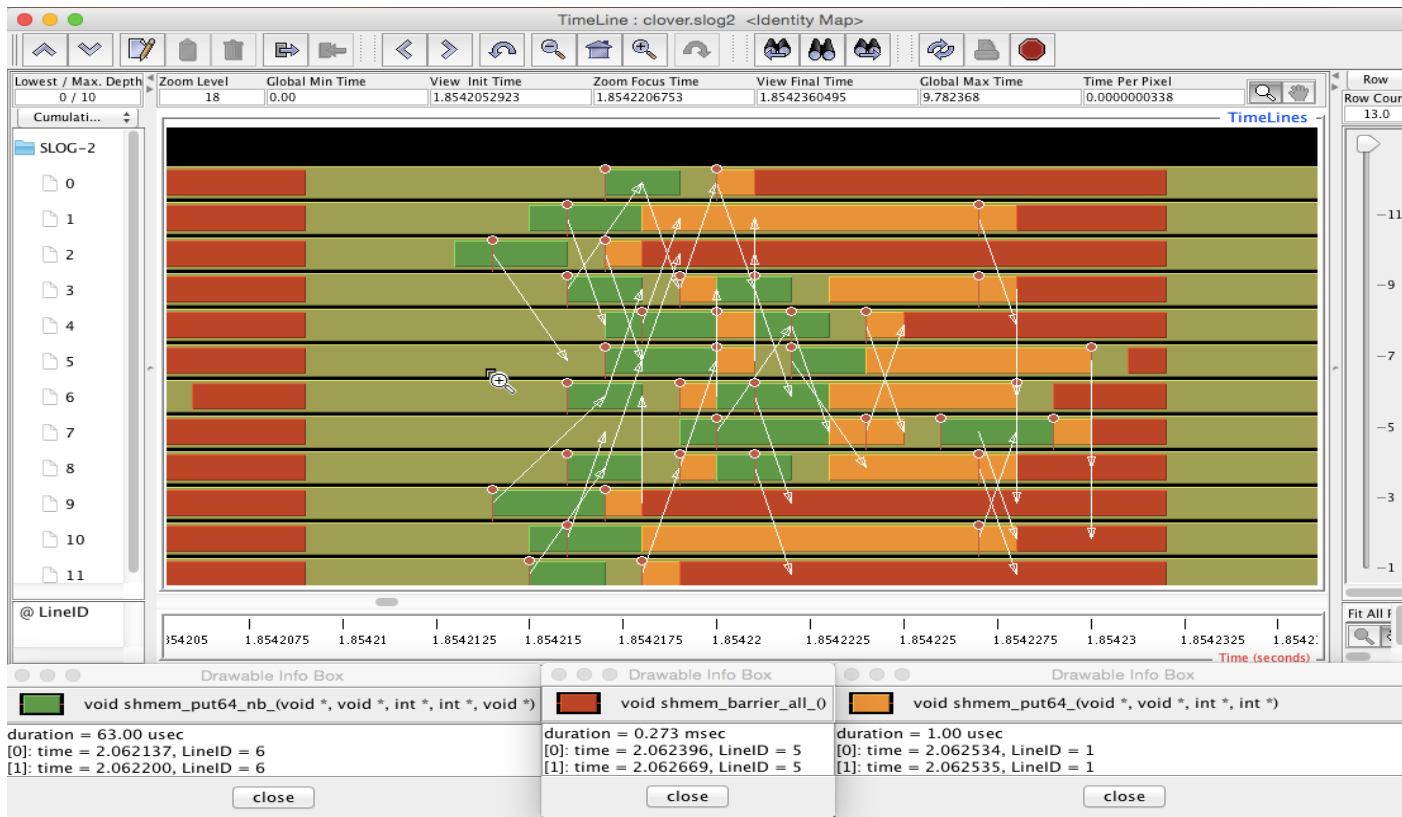
```
% export TAU_COMM_MATRIX=1; mpirun ... tau_exec ./a.out  
% paraprof ; Windows -> 3D Communication Matrix
```

TAU and Vampir [TU Dresden]: Intel oneAPI OpenCL

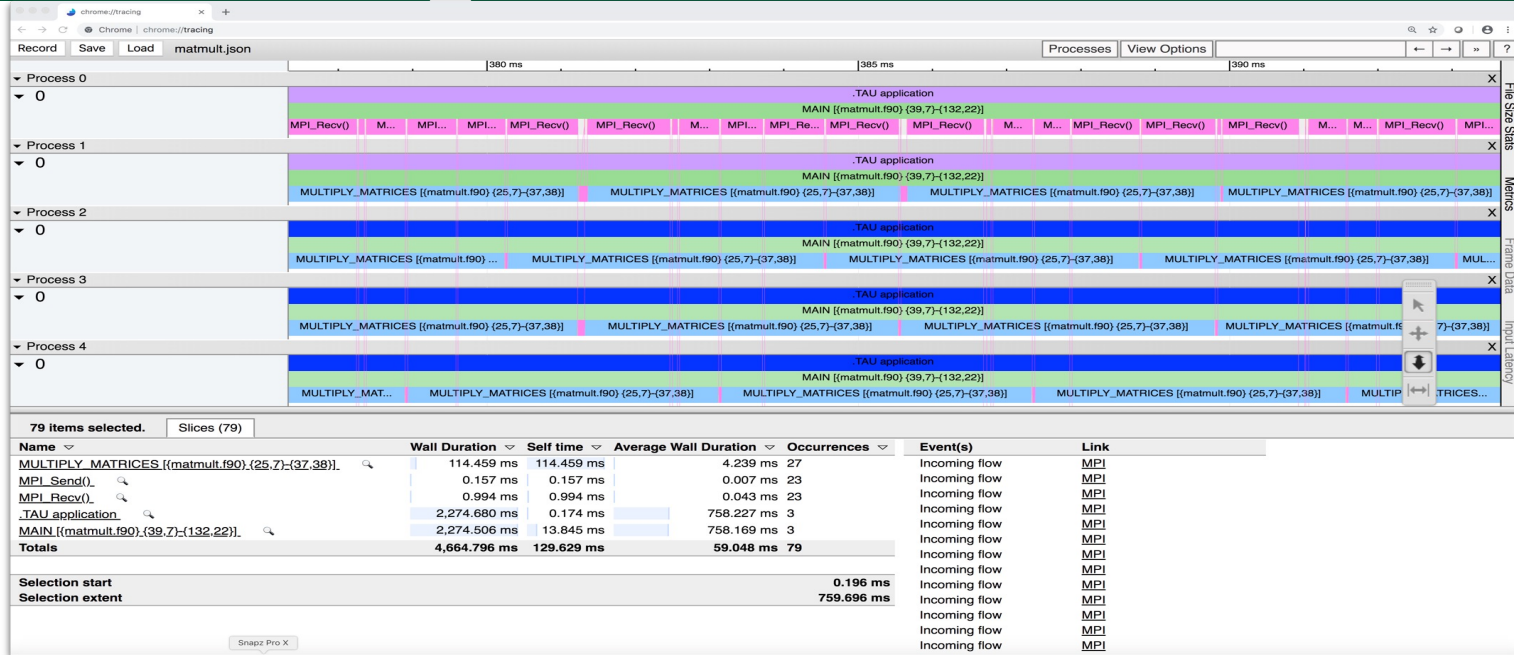


```
% export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2  
% mpirun -np 4 tau_exec -T level_zero -opencl ./a.out
```

Tracing: Jumpshot (ships with TAU)



Tracing: Chrome Browser



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

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

Or visit <https://ui.perfetto.dev/> to use Perfetto

Performance Measurement

- **Timers**

- Requires instrumentation of some kind
 - Manual, automated
 - Source, compiler provided, binary
 - Library callbacks, API wrappers, weak symbol replacement
- Simple to implement

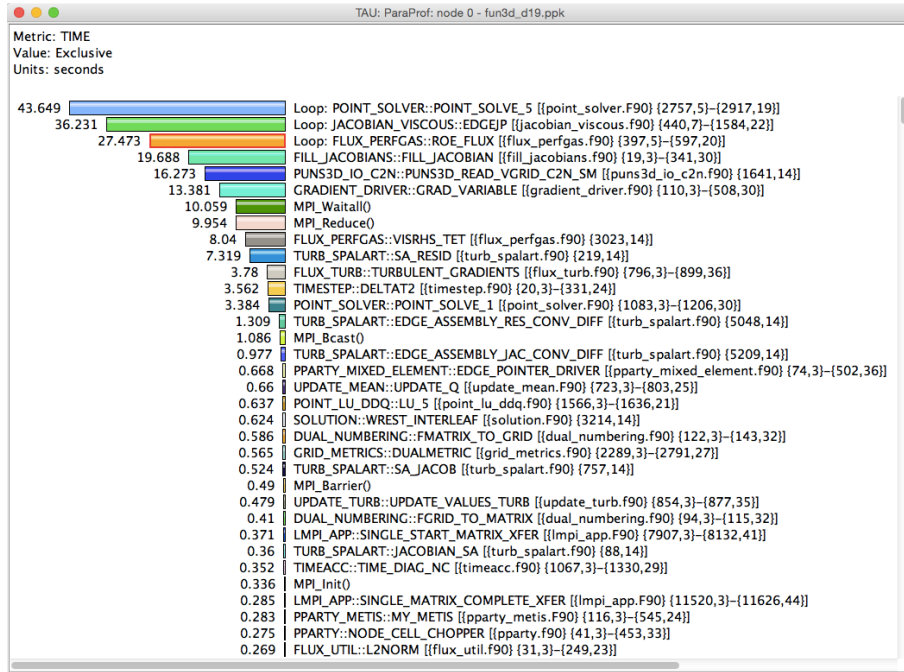
- **Sampling**

- Requires specialized system libraries / support
 - Periodic signals, signal handler
- No modification to executable/library needed
- Potential to interfere with system support (signal handlers)

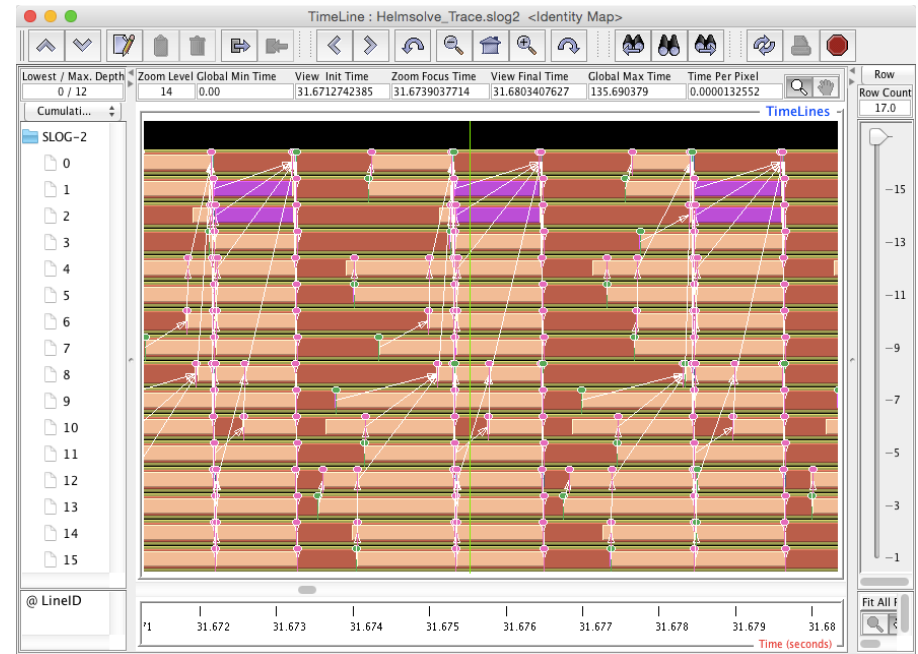
Profiling and Tracing

- **Profiling:** how much time was spent in each measured function on each thread in each process?
 - Collapses the time axis
 - No ordering or causal event information
 - Small summary per thread/process, regardless of execution time – only grows with number of timers & threads/processes
- **Tracing:** record all function entry & exit events on a timeline
 - Detailed view of what happened
 - The longer the program runs, the bigger the trace

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

Integrating TAU

Compile Time

- Compile with TAU compiler wrappers (see next slides)
- Link with TAU libraries

Runtime

- Execute with `tau_exec`
- Preloads the TAU shared object library and instantiates measurement support for different models
- More later...

Instrumentation Approaches

- Manual
 - Add TAU API calls to the code by hand:
<https://www.cs.uoregon.edu/research/tau/docs/newguide/bk05rn01.html>
- Automated:
 - PDT – optional TAU configuration
 - Compiler based instrumentation – comes with TAU
 - LLVM plugin – comes with TAU
 - Binary instrumentation - using Dyninst, PIN, or MAQAO
 - Optional TAU configuration, not covered in this tutorial
- PerfStubs API: <https://github.com/UO-OACISS/perfstubs>

TAU compiler wrappers with PDT

- `tau_cc.sh`, `tau_cxx.sh`, `tau_f90.sh`
- Usually does 3 passes to compile:
 - PDT parses the source file, writes a `.pdb` file
 - `tau_instrumentor` reads the source file, the `pdb` file, writes an instrumented source file
 - The instrumented source file is passed to the regular compiler
 - The instrumented source file is deleted
- Selective instrumentation by file, function (include/exclude)
- At link time, `tau_*.sh` will add the TAU libraries to the link

PDT Example

```
khuck@instinct:~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -g -O2 -no-pie -c simple.c -o simple.o

Debug: Parsing with PDT Parser. 1
Executing> /usr/local/packages/pdtoolkit/3.25.1-instinct/x86_64/bin/cparse simple.c -I/home/users/khuck/src/tau2/include -DPROFILING_ON -DTAU_GNU -DTAU_DOT_H_LESS_HEADERS -DTAU_LINUX_TIMERS -DTAU_LARGEFILE -D_LARGEFILE64_SOURCE -DTAU_BFD -DHAVE_GNU_DEMANGLE -DHAVE_TR1_HASH_MAP -DTAU_SS_ALLOC_SUPPORT -DEBS_CLOCK_RES=1 -DTAU_STRSIGNAL_OK -DTAU_TRACK_LD_LOADER -DTAU_ELF_BFD -DTAU_DWARF -I/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/include -I/home/users/khuck/src/tau2/include

Debug: Instrumenting with TAU 2
Executing> /home/users/khuck/src/tau2/x86_64/bin/tau_instrumentor simple.pdb simple.c -o simple.inst.c -c

Debug: Compiling with Instrumented Code 3
Executing> /usr/bin/gcc -g -O2 -no-pie -c simple.inst.c -DPROFILING_ON -DTAU_GNU -DTAU_DOT_H_LESS_HEADERS -DTAU_LINUX_TIMERS -DTAU_LARGEFILE -D_LARGEFILE64_SOURCE -DTAU_BFD -DHAVE_GNU_DEMANGLE -DHAVE_TR1_HASH_MAP -DTAU_SS_ALLOC_SUPPORT -DEBS_CLOCK_RES=1 -DTAU_STRSIGNAL_OK -DTAU_TRACK_LD_LOADER -DTAU_ELF_BFD -DTAU_DWARF -I/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/include -I/home/users/khuck/src/tau2/include -o simple.o
Looking for file: simple.o

Debug: cleaning inst file
Executing> /bin/rm -f simple.inst.c

Debug: cleaning PDB file
Executing> /bin/rm -f simple.pdb

link

TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -g -O2 -no-pie simple.o -o simple
Debug: Moving these libraries to the end of the link line:

Debug: Linking with TAU Options
Executing> /usr/bin/gcc -g -O2 -no-pie simple.o -o simple -L/home/users/khuck/src/tau2/x86_64/lib/shared-pdt -lTAU -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/lib/shared-pdt -L/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib -L/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib64 -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib64 -lbfd -liberty -lz -ldl -Wl,--export-dynamic -lrt -lm -L/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/lib -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/lib -ldwarf -lz -lelf -L/usr/lib/gcc/x86_64-linux-gnu/9/ -lstdc++ -lgcc_s -L/home/users/khuck/src/tau2/x86_64/lib/static-pdt -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/lib/shared-pdt -ldl

khuck@instinct:~/src/tau2/examples/simple$
```


PDT Instrumentation Example

```
tau_cc.sh -optKeepFiles -g -O2 -c simple.c -o simple.o
```

```
48
49 int main (int argc, char *argv[])
50 {
51     double **a = allocateMatrix(SIZE,
52     double **b = allocateMatrix(SIZE,
53     double **c = allocateMatrix(SIZE,
54     init(a);
55     init(b);
56     compute(a, b, c);
57     /* use the result */
58     printf("c[9][9] = %f\n", c[9][9]);
59     free(a);
60     free(b);
61     free(c);
62     return 0;
63 }
64
```

```
128
129 int main (int argc, char *argv[])
130 {
131     #line 50
132     TAU_PROFILE_TIMER(tautimer, "int main(int, char **) C [{simple.c} {49,1}-{63,1}]", " ", TAU_DEFAULT);
133     TAU_INIT(&argc, &argv);
134     #ifndef TAU_MPI
135     #ifndef TAU_SHMEM
136     TAU_PROFILE_SET_NODE(0);
137     #endif /* TAU_SHMEM */
138     #endif /* TAU_MPI */
139     TAU_PROFILE_START(tautimer);
140     #line 50
141     {
142         double **a = allocateMatrix(SIZE, SIZE);
143         double **b = allocateMatrix(SIZE, SIZE);
144         double **c = allocateMatrix(SIZE, SIZE);
145         init(a);
146         init(b);
147         compute(a, b, c);
148         /* use the result */
149         printf("c[9][9] = %f\n", c[9][9]);
150         free(a);
151         free(b);
152         free(c);
153     #line 62
154     { int tau_ret_val = 0; TAU_PROFILE_STOP(tautimer); return (tau_ret_val); }
155     #line 62
156     #line 63
157     }
158     TAU_PROFILE_STOP(tautimer);
159     #line 63
160     }
161
```

TAU compiler wrappers without PDT

- Same `tau_cc.sh`, `tau_cxx.sh`, `tau_f90.sh`
- Usually does 1 pass to compile:
 - Extra flags are added to the compiler:
 - Compiler based instrumentation (`-finstrument-functions`)
 - Tool has to implement two special functions:
 - » `void __cyg_profile_func_enter (void *this_fn, void *call_site);`
 - » `void __cyg_profile_func_exit (void *this_fn, void *call_site);`
 - Call a compiler plugin (LLVM only)
 - `tau_instrumentor` adds TAU API calls directly
- At link time, `tau_*.sh` will add the TAU libraries to the link
- Can be forced with `-optComplnst` flag

Compiler Based Example

```
khuck@instinct: ~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -optCompInst -g -O2 -no-pie -c simple.c -o simple.o
Debug: Using compiler-based instrumentation

Debug: Compiling with Instrumented Code
Executing> /usr/bin/gcc -g -O2 -no-pie -c simple.c -g -DPROFILING_ON -DTAU_GNU -DTAU_DOT_H_LESS_HEADERS -DTAU_LINUX_TIMERS -DTAU_LARGEFILE -D LARGEFILE64_SOURCE -DTAU_BFD -DHAVE_GNU_DEMANGLE -DHAVE_TR1_HASH_MAP -DTAU_SS_ALLOC_SUPPORT -DEBS_CLOCK_RES=1 -DTAU_STRSIGNAL_OK -DTAU_TRACK_LD_LOADER -DTAU_ELF_BFD -DTAU_DWARF -I/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/include -I/home/users/khuck/src/tau2/include -g -finstrument-functions -finstrument-functions-exclude-file-list=include,.hpp -o simple.o
Looking for file: simple.o

TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -optCompInst -g -O2 -no-pie simple.o -o simple
Debug: Using compiler-based instrumentation
Debug: Moving these libraries to the end of the link line:

Debug: Linking with TAU Options
Executing> /usr/bin/gcc -g -O2 -no-pie simple.o -o simple -L/home/users/khuck/src/tau2/x86_64/lib/shared-pdt -lTAU -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/lib/shared-pdt -L/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib -L/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib64 -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/binutils-2.36/lib64 -lbfd -liberty -lz -ldl -Wl,--export-dynamic -lrt -lm -L/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/lib -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/libdwarf-gcc/lib -ldwarf -lz -lelf -L/usr/lib/gcc/x86_64-linux-gnu/9/ -lstdc++ -lgcc_s -L/home/users/khuck/src/tau2/x86_64/lib/static-pdt -Wl,-rpath,/home/users/khuck/src/tau2/x86_64/lib/shared-pdt -ldl -g

khuck@instinct:~/src/tau2/examples/simple$
```

Simple example

```
1  /*****
2  *   Matrix Multiply - C Version
3  *   Demonstrates a matrix multiply using OpenMP.
4  *****/
5  #include <stdio.h>
6  #include <stdlib.h>
7
8  #define SIZE 1024
9
10 double** allocateMatrix(int rows, int cols) {
11     int i;
12     double **matrix = (double**)calloc(rows, (sizeof(double*)));
13     for (i=0; i<rows; i++) {
14         matrix[i] = (double*)calloc(cols, (sizeof(double)));
15     }
16     return matrix;
17 }
18
19 void freeMatrix(double** matrix, int rows, int cols) {
20     int i;
21     for (i=0; i<rows; i++) {
22         free(matrix[i]);
23     }
24     free(matrix);
25 }
26
27 /* Initialize the matrix to something other than zero */
28 void init(double **a) {
29     int i,j;
30     for (i=0; i < SIZE; i++) {
31         for (j=0; j < SIZE; j++) {
32             a[i][j] = i+j+1;
33         }
34     }
35 }
```

```
36
37 /* Perform matrix multiply */
38 void compute(double **a, double **b, double **c) {
39     int i,j,k;
40     for (i=0; i < SIZE; i++) {
41         for (j=0; j < SIZE; j++) {
42             for (k=0; k < SIZE; k++) {
43                 c[i][j] += a[i][k] * b[k][j];
44             }
45         }
46     }
47 }
48
49 int main (int argc, char *argv[])
50 {
51     double **a = allocateMatrix(SIZE, SIZE);
52     double **b = allocateMatrix(SIZE, SIZE);
53     double **c = allocateMatrix(SIZE, SIZE);
54     init(a);
55     init(b);
56     compute(a, b, c);
57     /* use the result */
58     printf("c[9][9] = %f\n", c[9][9]);
59     freeMatrix(a, SIZE);
60     freeMatrix(b, SIZE);
61     freeMatrix(c, SIZE);
62     return 0;
63 }
```

PDT Instrumentation

```
khuck@instinct:~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -optQuiet -g -O2 -no-pie -c simple.c -o simple.o
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -optQuiet -g -O2 -no-pie simple.o -o simple
khuck@instinct:~/src/tau2/examples/simple$ ./simple
c[9][9] = 367967744.000000
khuck@instinct:~/src/tau2/examples/simple$ pprof -a
Reading Profile files in profile.*

NODE 0;CONTEXT 0;THREAD 0:
-----
%Time   Exclusive   Inclusive   #Call   #Subrs   Inclusive Name
      msec     total msec                usec/call
-----
100.0         21       1,538         1         1 1538547 .TAU application
 98.6       0.078     1,517         1         9 1517236 int main(int, char **) C [{simple.c} {49,1}-{63,1}]
 97.1     1,494     1,494         1         0 1494336 void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]
  1.3         20         20          3         0   6813 double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]
  0.1          2          2          2         0   1129 void init(double **) C [{simple.c} {28,1}-{35,1}]
  0.0       0.125     0.125         3         0     42 void freeMatrix(double **, int) C [{simple.c} {19,1}-{25,1}]
-----
```

- Timer names have full signatures, start & end lines/columns
- all information available from parsing original source file with PDT

Compiler Instrumentation

```
khuck@instinct:~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -optQuiet -optCompInst -g -O2 -no-pie -c simple.c -o simple.o
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-pdt tau_cc.sh -optShared -optQuiet -optCompInst -g -O2 -no-pie simple.o -o simple
khuck@instinct:~/src/tau2/examples/simple$ ./simple
c[9][9] = 367967744.000000
khuck@instinct:~/src/tau2/examples/simple$ pprof -a
Reading Profile files in profile.*
```

```
NODE 0;CONTEXT 0;THREAD 0:
```

```
-----
%Time   Exclusive   Inclusive   #Call   #Subrs   Inclusive Name
      msec     total msec
-----
100.0         22       1,546       1         1   1546532 .TAU application
98.5         0.135     1,523       1         9   1523732 main [{/home/users/khuck/src/tau2/examples/simple/simple.c} {50,0}]
97.1         1,501     1,501       1         0   1501601 compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {38,0}]
1.3          19        19         3         0     6540 allocateMatrix [{/home/users/khuck/src/tau2/examples/simple/simple.c} {10,0}]
0.1          2         2          2         0     1128 init [{/home/users/khuck/src/tau2/examples/simple/simple.c} {28,0}]
0.0         0.118     0.118       3         0      39 freeMatrix [{/home/users/khuck/src/tau2/examples/simple/simple.c} {19,0}]
-----
```

Timer names have name only, start line only

- function entry/exit callback only has address of function and return address
- all source information retrieved during program execution with binutils (libbfd)

LLVM Plugin Instrumentation

Different TAU configuration with clang++/clang/flang

```
khuck@instinct:~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-rocm-rocprofiler-clang-ompt-pthread-pdt-openmp tau_cc.sh -optShared -optQuiet -optCompInst -g -O2
-c simple.c -o simple.o
Using selective instrumentation for LLVM
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-rocm-rocprofiler-clang-ompt-pthread-pdt-openmp tau_cc.sh -optShared -optQuiet -optCompInst -g -O2
simple.o -o simple
Using selective instrumentation for LLVM
khuck@instinct:~/src/tau2/examples/simple$ ./simple
c[9][9] = 367967744.000000
khuck@instinct:~/src/tau2/examples/simple$ pprof -a
Reading Profile files in profile.*

NODE 0;CONTEXT 0;THREAD 0:
-----
%Time   Exclusive   Inclusive   #Call   #Subrs   Inclusive   Name
        msec     total msec
-----
100.0    21          1,754      1       1       1754797    .TAU application
 98.8    18          1,733      1       1       1733520    main [{/home/users/khuck/src/tau2/examples/simple/simple.c} {49}]
 97.7    1,714      1,714      1       0       1714671    compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {38}]
-----
```

Where are the other timers?...

TAU_COMPILER_MIN_INSTRUCTION_COUNT defaults to 50, so they were filtered out

LLVM Plugin Instrumentation

Different TAU configuration with clang++/clang/flang

```
khuck@instinct:~/src/tau2/examples/simple$ export TAU_COMPILER_MIN_INSTRUCTION_COUNT=0 # show all timers
khuck@instinct:~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-rocm-rocprofiler-clang-ompt-pthread-pdt-openmp tau_cc.sh -optShared -optQuiet -optCompInst -g -O2
-c simple.c -o simple.o
Using selective instrumentation for LLVM
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-rocm-rocprofiler-clang-ompt-pthread-pdt-openmp tau_cc.sh -optShared -optQuiet -optCompInst -g -O2
simple.o -o simple
Using selective instrumentation for LLVM
khuck@instinct:~/src/tau2/examples/simple$ ./simple
c[9][9] = 367967744.000000
khuck@instinct:~/src/tau2/examples/simple$ pprof -a
Reading Profile files in profile.*
```

```
NODE 0;CONTEXT 0;THREAD 0:
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	20	1,542	1	1	1542909 .TAU application
98.7	0.083	1,522	1	9	1522755 main [{/home/users/khuck/src/tau2/examples/simple/simple.c} {49}]
97.3	1,500	1,500	1	0	1500818 compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {38}]
1.3	19	19	3	0	6636 allocateMatrix [{/home/users/khuck/src/tau2/examples/simple/simple.c} {10}]
0.1	1	1	2	0	858 init [{/home/users/khuck/src/tau2/examples/simple/simple.c} {28}]
0.0	0.232	0.232	3	0	77 freeMatrix [{/home/users/khuck/src/tau2/examples/simple/simple.c} {19}]

TAU_COMPILER_MIN_INSTRUCTION_COUNT=0 disables filtering

Instrumentation pros/cons

Pros

- Simple to implement
- Works universally
- Instruments everything – no blind spots
- Selective instrumentation available (by file or function name, or instruction count)

Cons

- Instruments *everything*
- Potentially high overhead – especially with C++
- Changes program behavior
- Potentially interferes with optimizations

Sampling

- Run the application with `tau_exec -ebs`
 - Preloads the TAU library, instantiates a signal handler and periodic interrupt to process the signal
 - The signal handler will record the current instruction pointer, all requested metrics, and optionally unwind the callstack
 - At the end of execution, all addresses are resolved to symbols in the application using `binutils/libdwarf`
- Some things that help:
 - For best support, build application with debug (-g) - all other optimizations are fine

Using TAU's Runtime Preloading Tool: tau_exec

- `<mpirun> tau_exec -T <config> <options> <executable>`
- `tau-config --list-matching <mpi/serial>` will show available configs
- Preload a wrapper that intercepts the runtime calls and substitutes with another (using `dlsym()` or weak symbol replacement)
 - 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)

Sampled Measurement

Previous instrumentation example:

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	21	1,538	1	1	1538547 .TAU application
98.6	0.078	1,517	1	9	1517236 int main(int, char **) C [{simple.c} {49,1}-{63,1}]
97.1	1,494	1,494	1	0	1494336 void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]
1.3	20	20	3	0	6813 double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]
0.1	2	2	2	0	1129 void init(double **) C [{simple.c} {28,1}-{35,1}]
0.0	0.125	0.125	3	0	42 void freeMatrix(double **, int) C [{simple.c} {19,1}-{25,1}]

Sampling example:

```
[khuck@instinct:~/src/tau2/examples/simple$ make
gcc -g -O2 -no-pie -c simple.c -o simple.o
gcc -g -O2 -no-pie simple.o -o simple
[khuck@instinct:~/src/tau2/examples/simple$ tau_exec -T serial -ebs ./simple
c[9][9] = 367967744.000000
[khuck@instinct:~/src/tau2/examples/simple$ pprof -a | grep -v CONTEXT
Reading Profile files in profile.*
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	1,730	1,730	1	0	1730983 .TAU application
84.9	1,469	1,469	49	0	30000 [SAMPLE] compute [(/home/users/khuck/src/tau2/examples/simple/simple.c) {43}]
13.9	240	240	8	0	30003 [SAMPLE] compute [(/home/users/khuck/src/tau2/examples/simple/simple.c) {42}]

Both more and less information at the same time...

Sampled Measurement

```
[khuck@instinct:~/src/tau2/examples/simple$ make
gcc -g -O2 -no-pie -c simple.c -o simple.o
gcc -g -O2 -no-pie simple.o -o simple
[khuck@instinct:~/src/tau2/examples/simple$ tau_exec -T serial -ebs ./simple
c[9][9] = 367967744.000000
[khuck@instinct:~/src/tau2/examples/simple$ pprof -a | grep -v CONTEXT
Reading Profile files in profile.*
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	1,730	1,730	1	0	1730983 .TAU application
84.9	1,469	1,469	49	0	30000 [SAMPLE] compute [(/home/users/khuck/src/tau2/examples/simple/simple.c) {43}]
13.9	240	240	8	0	30003 [SAMPLE] compute [(/home/users/khuck/src/tau2/examples/simple/simple.c) {42}]

```
37 /* Perform matrix multiply */
38 void compute(double **a, double **b, double **c) {
39     int i,j,k;
40     for (i=0; i < SIZE; i++) {
41         for(j=0; j < SIZE; j++) {
42             for (k=0; k < SIZE; k++) {
43                 c[i][j] += a[i][k] * b[k][j];
44             }
45         }
46     }
47 }
```

14% spent here

85% spent here

Simple Transformation – loop inversion

```
37 /* Perform matrix multiply */
38 void compute(double **a, double **b, double **c) {
39     int i,j,k;
40     for (i=0; i < SIZE; i++) {
41         for(j=0; j < SIZE; j++) {
42             for (k=0; k < SIZE; k++) {
43                 c[i][j] += a[i][k] * b[k][j];
44             }
45         }
46     }
47 }
```

```
37 /* Perform matrix multiply */
38 void compute(double **a, double **b, double **c) {
39     int i,j,k;
40     for (i=0; i < SIZE; i++) {
41         for (k=0; k < SIZE; k++) {
42             for(j=0; j < SIZE; j++) {
43                 c[i][j] += a[i][k] * b[k][j];
44             }
45         }
46     }
47 }
```

Reduced from 1.73 seconds

```
[khuck@instinct:~/src/tau2/examples/simple$ make
gcc -g -O2 -no-pie -c simple.c -o simple.o
gcc -g -O2 -no-pie simple.o -o simple
[khuck@instinct:~/src/tau2/examples/simple$ tau_exec -T serial -ebs ./simple
c[9][9] = 367967744.000000
[khuck@instinct:~/src/tau2/examples/simple$ pprof -a | grep -v CONTEXT
Reading Profile files in profile.*
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	976	976	1	0	976578 .TAU application
70.7	690	690	23	0	30001 [SAMPLE] compute [~/home/users/khuck/src/tau2/examples/simple/simple.c] {43}
27.6	269	269	9	0	29997 [SAMPLE] compute [~/home/users/khuck/src/tau2/examples/simple/simple.c] {42}

Both together!

Timers

```
khuck@instinct:~/src/tau2/examples/simple$ tau_exec -T serial -ebs ./simple
c[9][9] = 367967744.000000
khuck@instinct:~/src/tau2/examples/simple$ pprof -a | grep -v CONTEXT
Reading Profile files in profile.*
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	18	1,611	1	1	1611900 .TAU application
98.9	0.116	1,593	1	6	1593593 int main(int, char **) C [{simple.c} {49,1}-{63,1}]
97.6	1,572	1,572	1	0	1572560 void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]
67.0	1,079	1,079	36	0	30000 [SAMPLE] compute [/{home/users/khuck/src/tau2/examples/simple/simple.c} {43}]
29.8	480	480	16	0	30000 [SAMPLE] compute [/{home/users/khuck/src/tau2/examples/simple/simple.c} {42}]
1.9	30	30	1	0	30010 [SAMPLE] UNRESOLVED /usr/lib/x86_64-linux-gnu/libc-2.31.so
1.2	18	18	3	0	6198 double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]
0.1	2	2	2	0	1161 void init(double **) C [{simple.c} {28,1}-{35,1}]

Samples

...with callpath profiling

```
khuck@instinct:~/src/tau2/examples/simple$ TAU_CALLPATH=1 TAU_CALLPATH_DEPTH=100 tau_exec -T serial -ebs ./simple
```

```
c[9][9] = 367967744.000000
```

```
khuck@instinct:~/src/tau2/examples/simple$ pprof -a
```


```
Reading Profile files in profile.*
```

```
NODE 0;CONTEXT 0;THREAD 0:
```

```
-----  
%Time   Exclusive   Inclusive   #Call    #Subrs   Inclusive Name  
      msec     total msec  
-----  
100.0      1         2,040      1         1   2040179 .TAU application  
100.0      0         2,040      68        0   30000 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}] => void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}] => [CONTEXT] void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]  
100.0      0         2,040      68        0   30000 [CONTEXT] void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]  
99.9      0.15      2,038      1         6   2038709 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}]  
99.9      0.15      2,038      1         6   2038709 int main(int, char **) C [{simple.c} {49,1}-{63,1}]  
99.0      2,019     2,019      1         0   2019510 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}] => void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]  
99.0      2,019     2,019      1         0   2019510 void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}]  
86.8      1,769     1,769      59        0   30000 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}] => void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}] => [CONTEXT] void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}] => [SAMPLE] compute [/{home/users/khuck/src/tau2/examples/simple/simple.c} {43}]  
86.8      1,769     1,769      59        0   30000 [SAMPLE] compute [/{home/users/khuck/src/tau2/examples/simple/simple.c} {43}]  
13.2      270       270        9         0   30001 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}] => void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}] => [CONTEXT] void compute(double **, double **, double **) C [{simple.c} {38,1}-{47,1}] => [SAMPLE] compute [/{home/users/khuck/src/tau2/examples/simple/simple.c} {42}]  
13.2      270       270        9         0   30001 [SAMPLE] compute [/{home/users/khuck/src/tau2/examples/simple/simple.c} {42}]  
0.8        16        16         3         0   5562 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}] => double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]  
0.8        16        16         3         0   5562 double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]  
0.1         2         2          2         0   1182 .TAU application => int main(int, char **) C [{simple.c} {49,1}-{63,1}] => void init(double **) C [{simple.c} {28,1}-{35,1}]  
0.1         2         2          2         0   1182 void init(double **) C [{simple.c} {28,1}-{35,1}]
```


...easier to view in ParaProf

TAU: ParaProf: Statistics for: node 0 - /Users/khuck/tutorial



Name	Exclusive TI...	Inclusive TIME ▾	Calls	Child Calls
▾ .TAU application	0.001	2.04	1	1
▾ int main(int, char **) C [simple.c] {49,1}–{63,1}	0	2.039	1	6
▾ void compute(double **, double **, double **) C [simple.c] {38,1}–{47,1}	2.02	2.02	1	0
▾ [CONTEXT] void compute(double **, double **, double **) C [simple.c] {38,1}–{47,1}	0	2.04	68	0
[SAMPLE] compute [simple.c] {43}	1.77	1.77	59	0
[SAMPLE] compute [simple.c] {42}	0.27	0.27	9	0
double **allocateMatrix(int, int) C [simple.c] {10,1}–{17,1}	0.017	0.017	3	0
void init(double **) C [simple.c] {28,1}–{35,1}	0.002	0.002	2	0

Other measurement support

- Many programming models provide “hooks” for tools
- Often, instrumentation isn’t necessary!
 - **MPI**, SHMEM, Charm++
 - Pthreads, **OpenMP**, Kokkos
 - CUDA, **HIP/ROCm**, OneAPI, OpenACC, OpenCL, OpenMP offload
 - Python
 - Wrappers: POSIX, Chapel, UPC, memory, ARMCI, GASNet...
 - Java

Other TAU features

- Binary instrumentation
 - Dyninst, MAQAO, PIN
- Hardware counter support
 - PAPI, LIKWID
- Tracing support (native or converters)
 - Vampir (OTF2), Perfetto (JSON), Jumpshot (SLOG2), ...
- Plugins
 - OS/HW monitoring, ADIOS2, SOS, Mochi, SQLite3, ...

OpenMP

- <https://www.openmp.org>
- Pragma-based language extension to facilitate threading
- OpenMP 5.0 standard includes OpenMP Tools (OMPT/OMPD) specification for providing callbacks from the runtime to performance/debugging tools
- Provided by Intel, LLVM, IBM compilers
- GCC can use drop-in replacement (LLVM 8.0 runtime)
- TAU provides OPARI legacy support (when using PDT)

Adding OpenMP

```
37 /* Perform matrix multiply */
38 void compute(double **a, double **b, double **c) {
39     int i,j,k;
40     #pragma omp parallel for
41     for (i=0; i < SIZE; i++) {
42         for(j=0; j < SIZE; j++) {
43             for (k=0; k < SIZE; k++) {
44                 c[i][j] += a[i][k] * b[k][j];
45             }
46         }
47     }
48 }
```

If OMP_NUM_THREADS=4, SIZE=1024, then iteration space will be split into 4 of chunk size 256 each – 4x speedup

Compiling, Running, Reporting

```
khuck@instinct:~/src/tau2/examples/simple$ make
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-ompt-pdt-openmp tau_cc.sh -optShared -optQuiet -g -O2 -fopenmp -no-pie -c simple.c -o simple.o
TAU_MAKEFILE=/storage/users/khuck/src/tau2/x86_64/lib/Makefile.tau-ompt-pdt-openmp tau_cc.sh -optShared -optQuiet -g -O2 -fopenmp -no-pie simple.o -o simple
khuck@instinct:~/src/tau2/examples/simple$ export OMP_NUM_THREADS=2
khuck@instinct:~/src/tau2/examples/simple$ ./simple
c[9][9] = 367967744.000000
khuck@instinct:~/src/tau2/examples/simple$ pprof -s -a
Reading Profile files in profile.*
```

Compiler flag to Enable OpenMP

FUNCTION SUMMARY (total):

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	18	1,787	2	2	893950 .TAU application
97.2	1,738	1,738	2	1	869314 OpenMP_Implicit_Task
50.3	0.078	899	1	9	899006 int main(int, char **) C [{simple.c} {50,1}-{64,1}]
49.0	6	875	1	1	875235 void compute(double **, double **, double **) C [{simple.c} {38,1}-{48,1}]
48.7	0.012	870	1	1	870224 OpenMP_Thread_Type_ompt_thread_worker
48.6	0.131	868	1	1	868546 OpenMP_Parallel_Region compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {51, 0}]
1.2	21	21	3	0	7012 double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]
0.1	2	2	2	0	1071 void init(double **) C [{simple.c} {28,1}-{35,1}]
0.0	0.516	0.516	3	0	172 void freeMatrix(double **, int) C [{simple.c} {19,1}-{25,1}]
0.0	0.012	0.012	1	0	12 OpenMP_Sync_Region_Barrier compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {51, 0}]

Thread lifetime

Worker lifetime

Region

FUNCTION SUMMARY (mean):

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	9	893	1	1	893950 .TAU application
97.2	869	869	1	0.5	869314 OpenMP_Implicit_Task
50.3	0.039	449	0.5	4.5	899006 int main(int, char **) C [{simple.c} {50,1}-{64,1}]
49.0	3	437	0.5	0.5	875235 void compute(double **, double **, double **) C [{simple.c} {38,1}-{48,1}]
48.7	0.006	435	0.5	0.5	870224 OpenMP_Thread_Type_ompt_thread_worker
48.6	0.0655	434	0.5	0.5	868546 OpenMP_Parallel_Region compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {51, 0}]
1.2	10	10	1.5	0	7012 double **allocateMatrix(int, int) C [{simple.c} {10,1}-{17,1}]
0.1	1	1	1	0	1071 void init(double **) C [{simple.c} {28,1}-{35,1}]
0.0	0.258	0.258	1.5	0	172 void freeMatrix(double **, int) C [{simple.c} {19,1}-{25,1}]
0.0	0.006	0.006	0.5	0	12 OpenMP_Sync_Region_Barrier compute [{/home/users/khuck/src/tau2/examples/simple/simple.c} {51, 0}]

Synchronization

MPI Support

- MPI standard includes tool support
 - MPI_* functions are thin, weak wrappers around PMPI_* API
 - Tools create their own wrappers to replace them and intercept MPI calls
 - Tool library is preloaded or linked ahead of MPI library(ies)
 - Example:

```
int MPI_Barrier(MPI_Comm comm) {
    int returnVal;


    TAU_PROFILE_TIMER(tautimer, "MPI_Barrier()", " ", TAU_MESSAGE);
    TAU_PROFILE_START(tautimer);

    returnVal = PMPI_Barrier( comm );

    TAU_PROFILE_STOP(tautimer);
    return returnVal;
}
```

MPI example – Lulesh

- Lulesh 2.0.3 <https://asc.llnl.gov/codes/proxy-apps/lulesh>
- “The Shock Hydrodynamics Challenge Problem was originally defined and implemented by LLNL as one of five challenge problems in the DARPA UHPC program and has since become a widely studied proxy application in DOE co-design efforts for exascale.”
- C++, Serial, OpenMP, MPI
- CUDA, OpenACC, OpenCL, other models



Lulesh Profile - ParaProf

TAU: ParaProf Manager

TrialField	Value
Name	lulesh-tutorial.ppk
Application ID	0
Experiment ID	0
Trial ID	0
CPU Cores	14

TAU: ParaProf: Mean - lulesh-tutorial.ppk

Metric	Value	Units
Metric: TIME	4.096	seconds
Value: Exclusive	2.157	
Units: seconds	1.798	
	1.754	
	1.022	
	0.828	
	0.623	
	0.185	
	0.148	
	0.145	
	0.143	
	0.122	
	0.062	
	0.054	

Main window

Mean profile

TAU: ParaProf: Mean Statistics - lulesh-tutorial-callpath.ppk

Name	Exclusive TIME	Inclusive TIME	Calls	Child C...
main [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]	1.78	12.465	1	1,012,125
CalcQForElems(Domain&, double*) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]	4.082	5.077	200	600
CommSend(Domain&, int, int, double& (Domain:**)(int), int, int, in	0.233	0.951	200	1,000
MPI_Waitall()	0.714	0.714	200	0
MPI_Isend()	0.003	0.003	600	0
MPI_Comm_rank()	0	0	200	0
CommMonoQ(Domain&) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]	0.037	0.039	200	800
CommRecv(Domain&, int, int, int, int, bool, bool) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]	0.003	0.005	200	800
MPI_Irecv()	0.002	0.002	600	0
MPI_Comm_rank()	0.001	0.001	200	0
LagrangeNodal(Domain&) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]	0.002	2.632	200	800
CommSyncPosVel(Domain&) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]	0.991	1.075	200	900
MPI_Wait()	0.095	0.095	800	0
MPI_Comm_rank()	0	0	200	0
CommSend(Domain&, int, int, double& (Domain:**)(int), int, int, in	0.28	0.948	200	1,100
MPI_Isend()	0.666	0.666	200	0
MPI_Waitall()			800	0
MPI_Comm_rank()			200	0
CalcForceForElems(Domain&, double const*, double const*, double const*			200	800
CommRecv(Domain&, int, int, int, int, bool, bool) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			200	1,800
CommSbnc(Domain&, int, int, double& (Domain:**)(int), int, int, in			200	1,600
CalcVolumeForceForElems(Domain&, double const*, double const*, double const*			200	400
CommRecv(Domain&, int, int, int, int, bool, bool) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			200	1,600
CommRecv(Domain&, int, int, int, int, bool, bool) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			200	900
Domain:Domain(int, int, int, int, int, int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			1	100,010
CalcElemVolume(double const*, double const*, double const*) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			100,001	100,001
Domain:AllocateElemPersistent(int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			1	0
Domain:AllocateNodePersistent(int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			1	0
Domain:CreateRegionIndexSets(int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			1	1
Domain:BuildMesh(int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]			1	0
Domain:SetupElemElems(int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]				
Domain:SetupBound(int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]				
Domain:SetupComm(int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]				
Domain:SetupSymm(int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]				
Domain:SetupThrea(int, int, int) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]				
Timelncrment(Domain&, double const*, double const*, double const*) [/{home/users/khuck/src/lulesh-2.0.3/lulesh.cc} [2735]]				
CommSend(Domain&, int, int, double& (Domain:**)(int), int, int, in				

Treetable of callpath data

TAU: ParaProf: Main Profile Window

Metric: TIME
Value: Exclusive

Std. Dev. [Bar chart]
Mean [Bar chart]
Max [Bar chart]
Min [Bar chart]

node 3 [Bar chart]
node 4 [Bar chart]
node 5 [Bar chart]
node 6 [Bar chart]
node 7 [Bar chart]

Main Profile Window

Profile of one timer

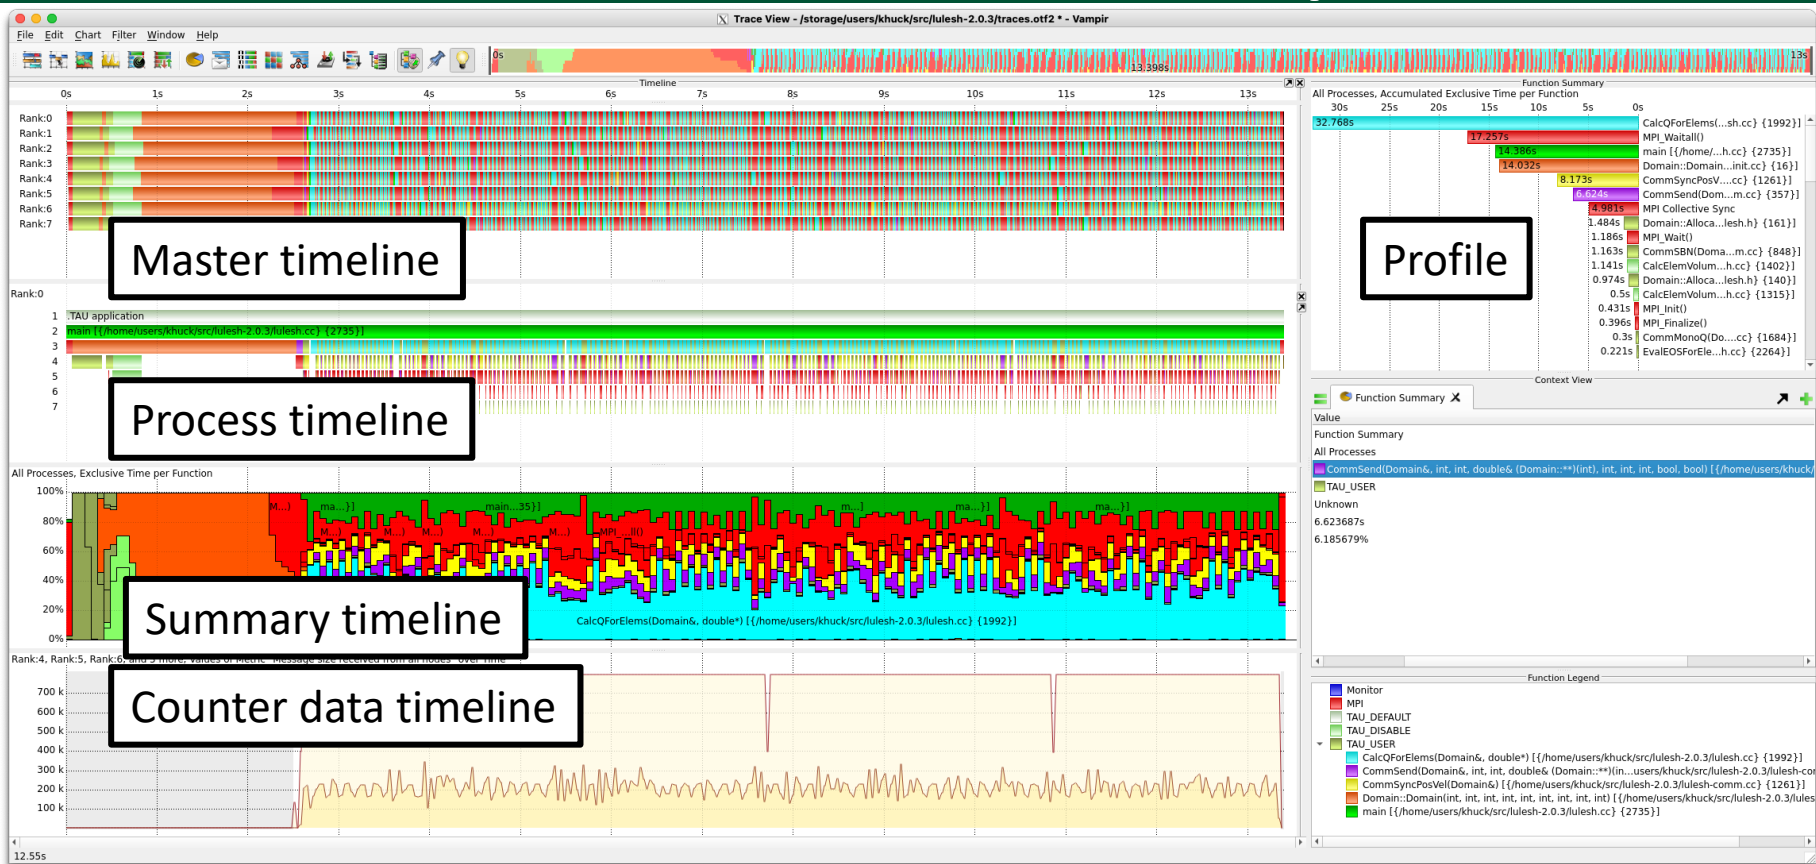
TAU: ParaProf: Profile of one timer

Name: MPI_Waitall()
Metric Name: TIME
Value: Exclusive
Units: seconds

3.102 [Bar chart]
2.927 [Bar chart]
2.985 [Bar chart]
3.009 [Bar chart]
3.102 [Bar chart]

node 0 [Bar chart]
node 1 [Bar chart]
node 2 [Bar chart]
node 3 [Bar chart]
node 4 [Bar chart]
node 5 [Bar chart]
node 6 [Bar chart]
node 7 [Bar chart]

Lulesh Trace – Vampir



Measuring HIP kernel performance

- Hip-stream – small program with 4+ kernels

```
[[khuck@login2.crusher add4]$ ./gpu-stream-hip
GPU-STREAM
Version: 1.0
Implementation: HIP
GridSize: 26214400 work-items
GroupSize: 1024 work-items
Operations/Work-item: 1
Precision: double

Running kernels 10 times
Array size: 200.0 MB (=0.2 GB) 0 bytes padding
Total size: 1000.0 MB (=1.0 GB)
Using HIP device (compute_units=110)
Driver: 50013601
d_a=0x7f0cb0000000
d_b=0x7f0ca0000000
d_c=0x7f0c90000000
d_d=0x7f0c80000000
d_e=0x7f0c70000000
Function      MBytes/sec  Min (sec)   Max          Average
Copy          1331503.944 0.00032     0.00032     0.00032
Mul           1332392.192 0.00031     0.00032     0.00032
Add4          1196944.446 0.00088     0.00089     0.00088
Triad         1256501.941 0.00050     0.00051     0.00050
GEOMEAN       1278064.946
```

Program output

```
*-----*
* Copyright 2015: Tom Deakin, Simon McIntosh-Smith, University of Bristol HPC
* Based on John D. McCalpin's original STREAM benchmark for CPUs
*-----*
```

```
template <typename T> __global__ void copy(const T * a, T * c) {
    const int i = hipBlockDim_x * hipBlockIdx_x + hipThreadIdx_x;
    c[i] = a[i];
}

template <typename T> __global__ void mul(T * b, const T * c) {
    const T scalar = 3.0;
    const int i = hipBlockDim_x * hipBlockIdx_x + hipThreadIdx_x;
    b[i] = scalar * c[i];
}

template <typename T> __global__ void
add(const T * a, const T * b, const T *d, const T *e, T * c) {
    const int i = hipBlockDim_x * hipBlockIdx_x + hipThreadIdx_x;
    c[i] = a[i] + b[i] + d[i] + e[i];
}

template <typename T> __global__ void
triad(T * a, const T * b, const T * c) {
    const T scalar = 3.0;
    const int i = hipBlockDim_x * hipBlockIdx_x + hipThreadIdx_x;
    a[i] = b[i] + scalar * c[i];
}
```

HIP kernels

Measuring HIP kernel performance

- Just add **tau_exec** and arguments to the command (between `srun/mpirun` and application when applicable)
- **tau-config** shows available configs

“use serial,rocprofiler configuration with HIP/ROCm support enabled”

```
[[khuck@login2.crusher add4]$ tau_exec -T serial,rocprofiler -rocm ./gpu-stream-hip
GPU-STREAM
Version: 1.0
Implementation: HIP
GridSize: 26214400 work-items
GroupSize: 1024 work-items
Operations/Work-item: 1
Precision: double

Running kernels 10 times
Array size: 200.0 MB (=0.2 GB) 0 bytes padding
Total size: 1000.0 MB (=1.0 GB)
Using HIP device (compute_units=110)
Driver: 50013601
d_a=0x7f48e0000000
d_b=0x7f48d0000000
d_c=0x7f47b0000000
d_d=0x7f47a0000000
d_e=0x7f4790000000
Function      MBytes/sec  Min (sec)   Max         Average
Copy          1320624.685 0.00032     0.00032     0.00032
Mul           1321623.393 0.00032     0.00032     0.00032
Add4          1217965.813 0.00086     0.00088     0.00087
Triad         1254042.504 0.00050     0.00051     0.00051
GEOMEAN      1277787.457
```

Pprof output, timers

```
[khuck@login2.crusher add4]$ pprof
Reading Profile files in profile.*
```

```
NODE 0;CONTEXT 0;THREAD 0:
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	1,031	1,031	1	1	1031765 .TAU application
0.0	0.03	0.03	1	0	30 pthread_create

Main Thread

```
NODE 0;CONTEXT 0;THREAD 1:
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	1	479	1	1	479989 .TAU application
99.6	478	478	1	0	478248 [PTHREAD] _ZN4rocr2os16ThreadTrampolineEPv

ROCm Thread

```
NODE 0;CONTEXT 0;THREAD 2:
```

%Time	Exclusive msec	Inclusive total msec	#Call	#Subrs	Inclusive Name usec/call
100.0	0.638	20	1	40	20298 .TAU application
42.4	8	8	10	0	861 void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
24.2	4	4	10	0	492 void triad<double>(double*, double const*, double const*) [clone .kd]
15.1	3	3	10	0	307 void copy<double>(double const*, double*) [clone .kd]
15.1	3	3	10	0	306 void mul<double>(double*, double const*) [clone .kd]

Device activity

Pprof output, counters

USER EVENTS Profile :NODE 0, CONTEXT 0, THREAD 2

NumSamples	MaxValue	MinValue	MeanValue	Std. Dev.	Event Name
10	2.621E+07	2.621E+07	2.621E+07	0	Grid Size : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	2.621E+07	2.621E+07	2.621E+07	0	Grid Size : void copy<double>(double const*, double*) [clone .kd]
10	2.621E+07	2.621E+07	2.621E+07	0	Grid Size : void mul<double>(double*, double const*) [clone .kd]
10	2.621E+07	2.621E+07	2.621E+07	0	Grid Size : void triad<double>(double*, double const*, double const*) [clone .kd]
10	0	0	0	0	LDS Memory Size : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	0	0	0	0	LDS Memory Size : void copy<double>(double const*, double*) [clone .kd]
10	0	0	0	0	LDS Memory Size : void mul<double>(double*, double const*) [clone .kd]
10	0	0	0	0	LDS Memory Size : void triad<double>(double*, double const*, double const*) [clone .kd]
10	32	32	32	0	Scalar Register Size (SGPR) : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	24	24	24	0	Scalar Register Size (SGPR) : void copy<double>(double const*, double*) [clone .kd]
10	24	24	24	0	Scalar Register Size (SGPR) : void mul<double>(double*, double const*) [clone .kd]
10	24	24	24	0	Scalar Register Size (SGPR) : void triad<double>(double*, double const*, double const*) [clone .kd]
10	0	0	0	0	Scratch Memory Size : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	0	0	0	0	Scratch Memory Size : void copy<double>(double const*, double*) [clone .kd]
10	0	0	0	0	Scratch Memory Size : void mul<double>(double*, double const*) [clone .kd]
10	0	0	0	0	Scratch Memory Size : void triad<double>(double*, double const*, double const*) [clone .kd]
10	8	8	8	0	Vector Register Size (VGPR) : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	4	4	4	0	Vector Register Size (VGPR) : void copy<double>(double const*, double*) [clone .kd]
10	4	4	4	0	Vector Register Size (VGPR) : void mul<double>(double*, double const*) [clone .kd]
10	4	4	4	0	Vector Register Size (VGPR) : void triad<double>(double*, double const*, double const*) [clone .kd]
10	1024	1024	1024	0	Work Group Size : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	1024	1024	1024	0	Work Group Size : void copy<double>(double const*, double*) [clone .kd]
10	1024	1024	1024	0	Work Group Size : void mul<double>(double*, double const*) [clone .kd]
10	1024	1024	1024	0	Work Group Size : void triad<double>(double*, double const*, double const*) [clone .kd]
10	5952	5952	5952	0	fbarrier count : void add<double>(double const*, double const*, double const*, double const*, double*) [clone .kd]
10	3392	3392	3392	0	fbarrier count : void copy<double>(double const*, double*) [clone .kd]
10	4672	4672	4672	0	fbarrier count : void mul<double>(double*, double const*) [clone .kd]
10	0	0	0	0	fbarrier count : void triad<double>(double*, double const*, double const*) [clone .kd]

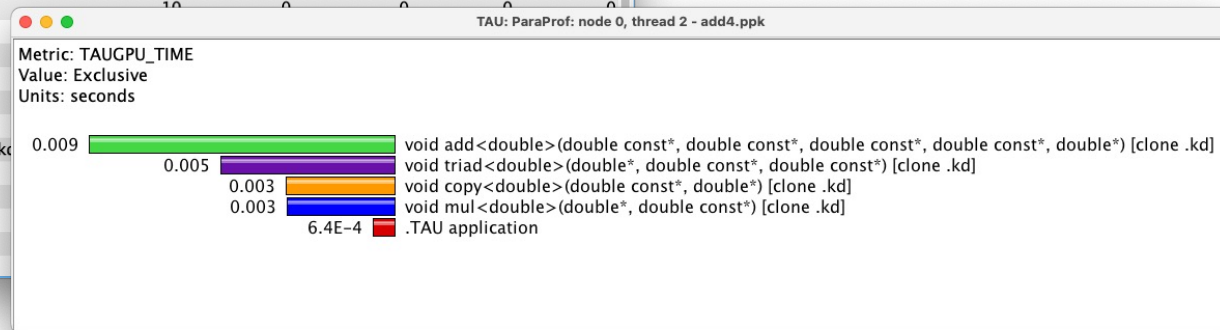
Counters for measuring register pressure and occupancy

ParaProf view of same data

TAU: ParaProf: Context Events for: node 0, thread 2 - add4.ppk

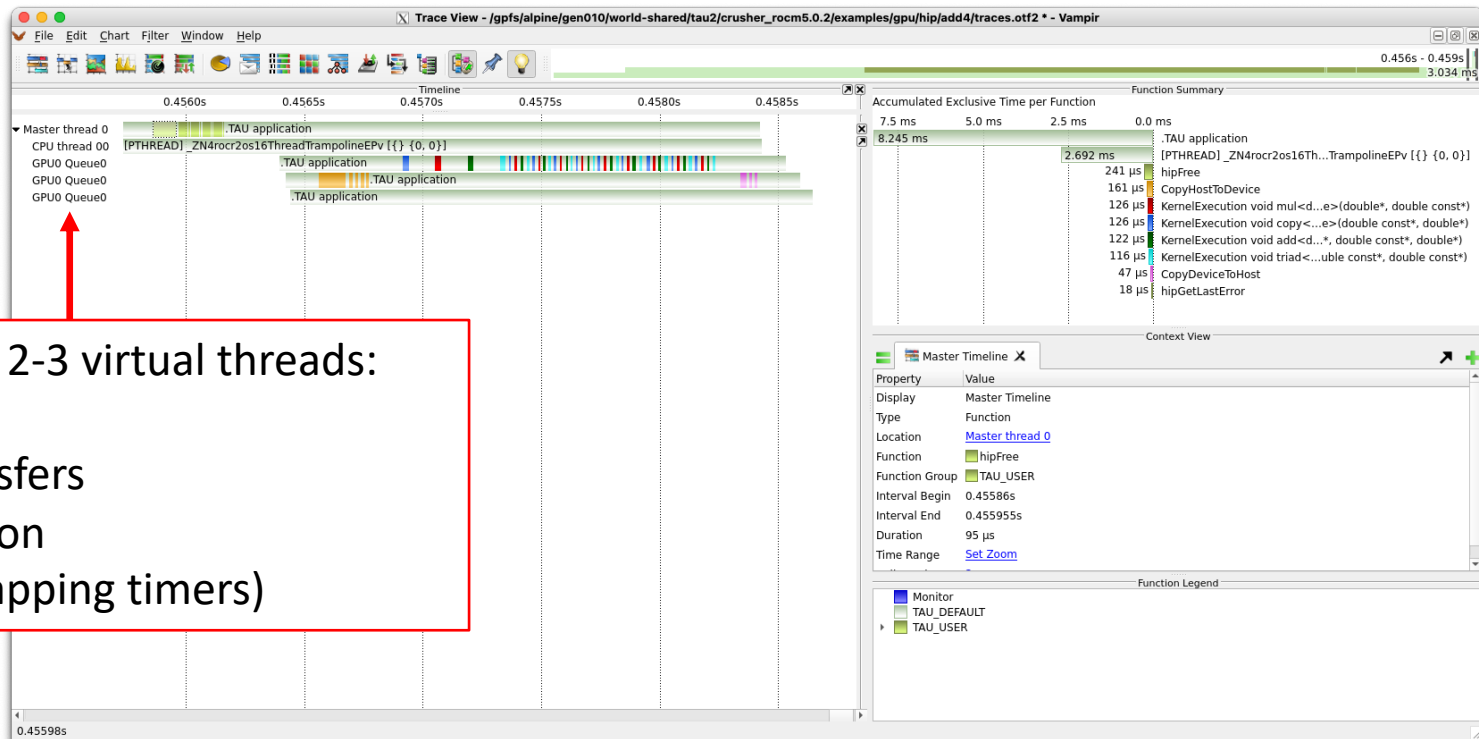
Name	NumSamp...	MaxValue	MinValue	MeanVal...	Std. Dev.
void add<double>(double const*, double const*, double const*, double cor					
Grid Size	10	26,214,400	26,214,400	26,214,400	0
LDS Memory Size	10	0	0	0	0
Scalar Register Size (SGPR)	10	32	32	32	0
Scratch Memory Size	10	0	0	0	0
Vector Register Size (VGPR)	10	8	8	8	0
Work Group Size	10	1,024	1,024	1,024	0
fbarrier count	10	5,952	5,952	5,952	0
void copy<double>(double const*, double*) [clone .kd]					
Grid Size	10	26,214,400	26,214,400	26,214,400	0
LDS Memory Size	10	0	0	0	0
Scalar Register Size (SGPR)	10	24	24	24	0
Scratch Memory Size	10	0	0	0	0
Vector Register Size (VGPR)	10	4	4	4	0
Work Group Size	10	1,024	1,024	1,024	0
fbarrier count	10	3,392	3,392	3,392	0
void mul<double>(double*, double const*) [clone .kd]					
Grid Size	10	26,214,400	26,214,400	26,214,400	0
LDS Memory Size	10	0	0	0	0
Scalar Register Size (SGPR)	10	0	0	0	0
Scratch Memory Size	10	0	0	0	0
Vector Register Size (VGPR)	10	0	0	0	0
void triad<double>(double*, double const*, double const*) [clone .kd]					
Grid Size	10	26,214,400	26,214,400	26,214,400	0
LDS Memory Size	10	0	0	0	0
Scalar Register Size (SGPR)	10	0	0	0	0
Scratch Memory Size	10	0	0	0	0
Vector Register Size (VGPR)	10	0	0	0	0

VERY helpful for understanding **register pressure** and **occupancy**



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_exec command reference

- Uninstrumented execution
 - % mpirun -np 256 ./a.out
- Track GPU operations
 - % mpirun -np 256 tau_exec -l0 ./a.out
 - % mpirun -np 256 tau_exec -opencl ./a.out
 - % mpirun -np 256 tau_exec -openacc ./a.out
 - % mpirun -np 256 tau_exec -cupti ./a.out
 - % mpirun -np 256 tau_exec -rocm ./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,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 export TAU_METRICS=TIME,PAPI_L1_DCM... -ebs_resolution=<file | function | line>
- Non-MPI execution: use -T serial
 - % tau_exec -T serial,level_zero -l0 -ebs ./a.out

TAU 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>)

TAU 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.

TAU 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

For more info...

- <https://tau.uoregon.edu>
- <https://github.com/UO-OACISS/tau2>
- <https://github.com/UO-OACISS/tau2/wiki>
- <https://github.com/UO-OACISS/tau2/wiki/Frequently-Asked-Questions-%28FAQ%29>
- Email tau-bugs@cs.uoregon.edu

Acknowledgements

Parts of 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.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.



Current/Previous Acknowledgements

