

Parallel Performance Evaluation With TAU

CSCADS 2009

Wyatt Spear

wspear@cs.uoregon.edu

<http://tau.uoregon.edu>



TAU Performance System[®] Project

- ***T*uning and *A*nalysis *U*tilities (15+ year project effort)**
- **Performance system framework for HPC systems**
 - Integrated, scalable, and flexible
 - Target parallel programming paradigms
- **Integrated toolkit for performance problem solving**
 - Instrumentation, measurement, analysis, and visualization
 - Portable performance profiling and tracing facility
 - Performance data management and data mining
- ***Partners***
 - LLNL, ANL, LANL
 - Research Centre Jülich, TU Dresden



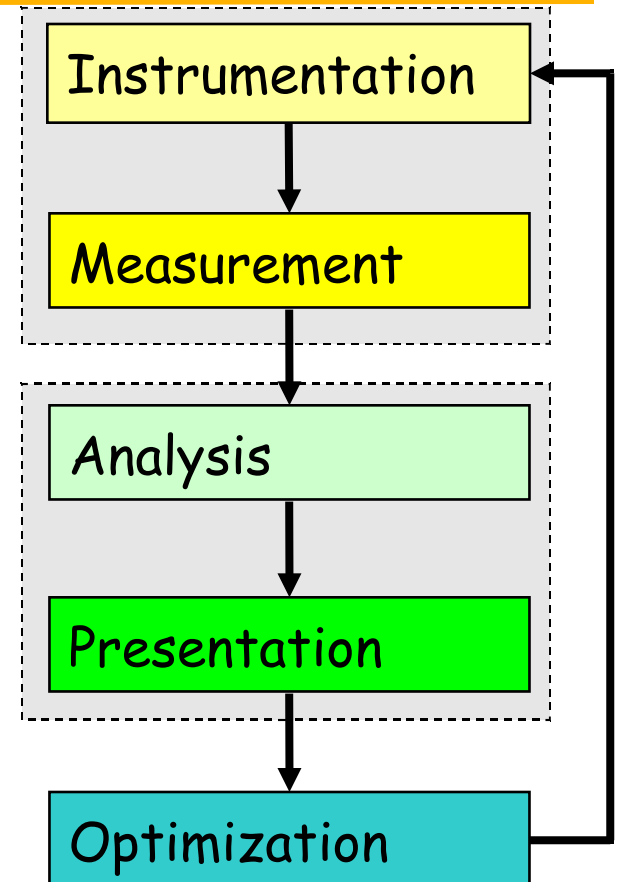
What is TAU?

- TAU is a performance evaluation tool
- It supports parallel profiling and tracing
- Profiling shows you how much (total) time was spent in each routine
- Tracing shows you *when* the events take place in each process along a timeline
- TAU uses a package called PDT for automatic instrumentation of the source code
- Profiling and tracing can measure time as well as hardware performance counters from your CPU
- TAU can automatically instrument your source code (routines, loops, I/O, memory, phases, etc.)
- TAU runs on all HPC platforms and it is free (BSD style license)
- TAU has instrumentation, measurement and analysis tools
 - paraprof is TAU's 3D profile browser
- **To use TAU, you need to set a couple of environment variables and substitute the name of your compiler with a TAU shell script**



Performance Optimization Cycle

- Expose factors
- Collect performance data
- Calculate metrics
- Analyze results
- Visualize results
- Identify problems
- Tune performance



Steps of Performance Evaluation

- Collect basic routine-level timing profile to determine where most time is being spent
- Collect routine-level hardware counter data to determine types of performance problems
- Collect callpath profiles to determine sequence of events causing performance problems
- Conduct finer-grained profiling and/or tracing to pinpoint performance bottlenecks
 - Loop-level profiling with hardware counters
 - Tracing of communication operations



Parallel Performance Properties

- Parallel code performance is influenced by both sequential and parallel factors?
- Sequential factors
 - Computation and memory use
 - Input / output
- Parallel factors
 - Thread / process interactions
 - Communication and synchronization



Performance Analysis Questions

- How does performance vary with different compilers?
- Is poor performance correlated with certain OS features?
- Has a recent change caused unanticipated performance?
- How does performance vary with MPI variants?
- Why is one application version faster than another?
- What is the reason for the observed scaling behavior?
- Did two runs exhibit similar performance?
- How are performance data related to application events?
- Which machines will run my code the fastest and why?
- Which benchmarks predict my code performance best?



TAU Parallel Performance System Goals

- **Portable (open source) parallel performance system**
 - Computer system architectures and operating systems
 - Different programming languages and compilers
- Multi-level, multi-language performance instrumentation
- **Flexible and configurable performance measurement**
- Support for multiple parallel programming paradigms
 - Multi-threading, message passing, mixed-mode, hybrid, object oriented (generic), component-based
- Support for performance mapping
- Integration of leading performance technology
- **Scalable (very large) parallel performance analysis**



Using TAU: A brief Introduction

- TAU supports several measurement options (profiling, tracing, profiling with hardware counters, etc.)
- Each measurement configuration of TAU corresponds to a unique stub makefile that is generated when you configure it
- To instrument source code using PDT
 - Choose an appropriate TAU stub makefile in <arch>/lib:
`% export TAU_MAKEFILE=/projects/tau/tau_latest/x86_64/lib/Makefile.tau-mpi-pdt`
`% export TAU_OPTIONS='-optVerbose ...'` (see `tau_compiler.sh -help`)
And use `tau_f90.sh`, `tau_cxx.sh` or `tau_cc.sh` as Fortran, C++ or C compilers:
`% mpif90 foo.f90`
changes to
`% tau_f90.sh foo.f90`
- Execute application and analyze performance data:
 - At runtime, if more than one metric is measured
 - `export TAU_METRICS=TIME:PAPI_FP_INS:PAPI_NATIVE_<native_event_name>`
 - Use `papi_native_avail`, `papi_avail`, and `papi_event_chooser` to select these preset and native event names



Using TAU

- Configuration
- Instrumentation
 - Manual
 - MPI – Wrapper interposition library
 - PDT- Source rewriting for C,C++, F77/90/95
 - Compiler-based instrumentation for C, C++, F90
 - OpenMP – Directive rewriting
 - Component based instrumentation – Proxy components
 - Binary Instrumentation
 - DyninstAPI – Runtime Instrumentation/Rewriting binary
 - Java – Runtime instrumentation
 - Python – Runtime instrumentation
- Measurement
- Performance Analysis

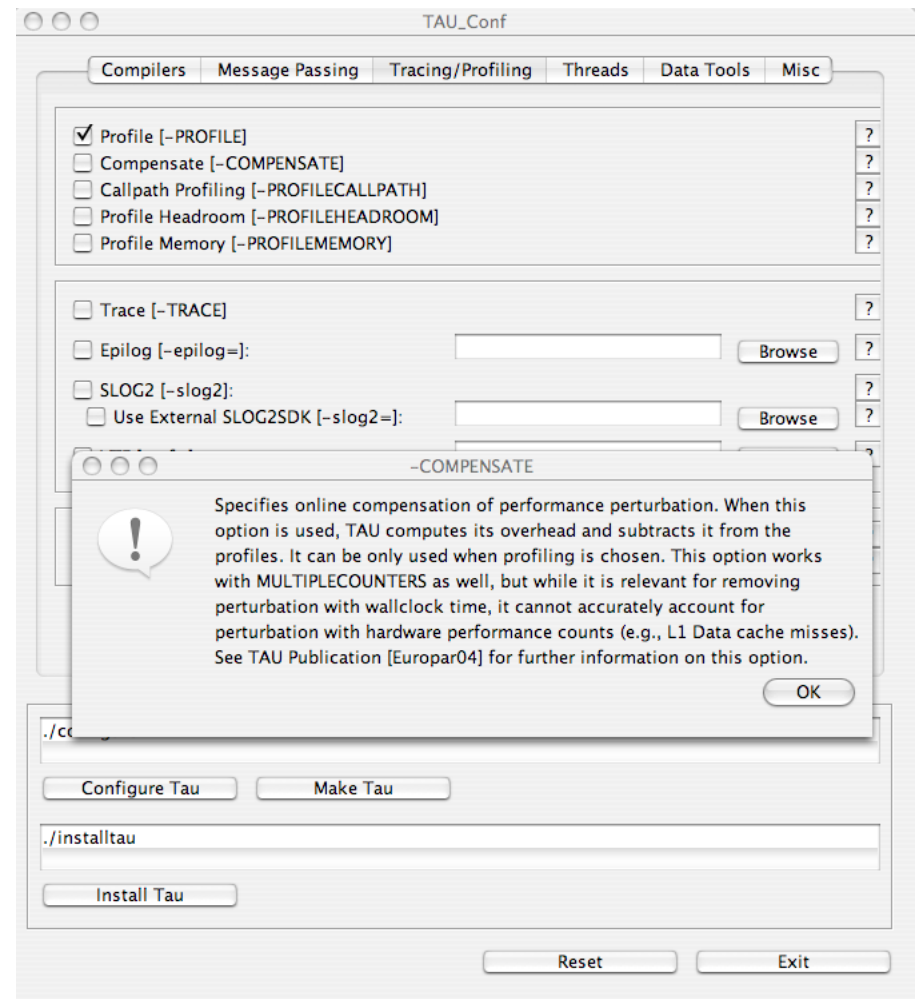
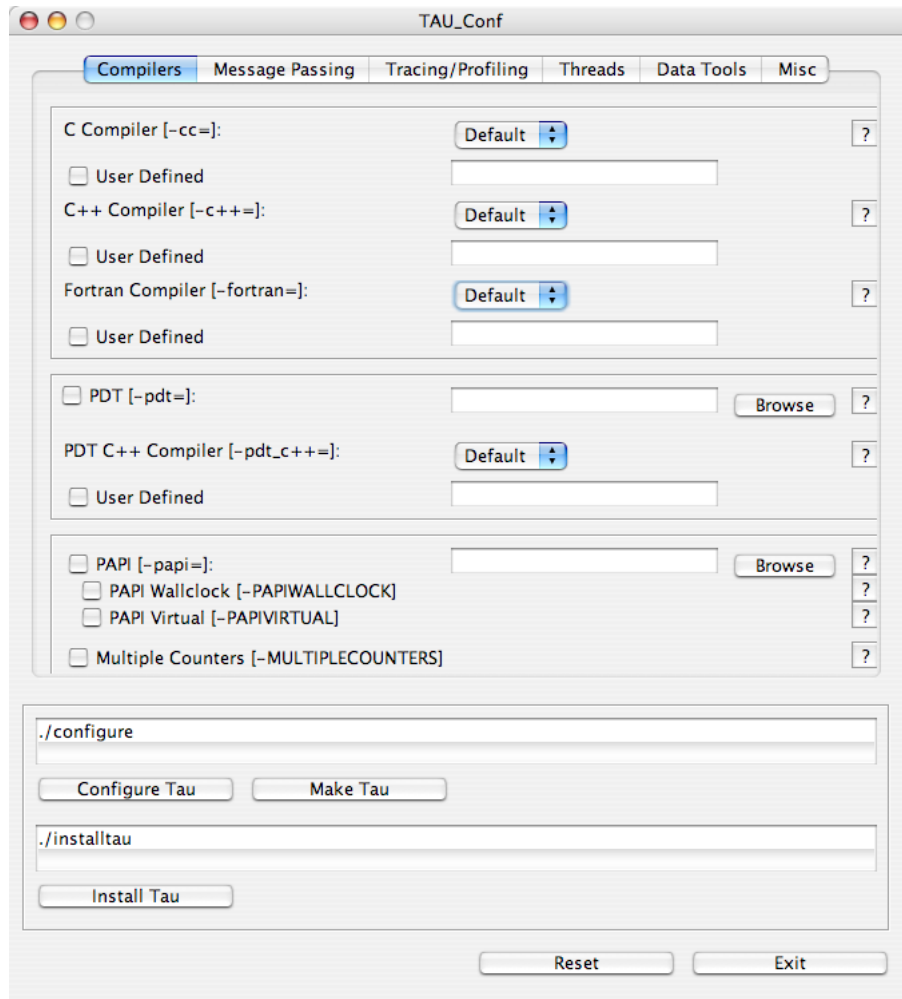


TAU Measurement Configuration – Examples

- `./configure -arch=x86_64 -pdt=/projects/tau/pdtoolkit-3.14 -mpi`
Configure using PDT and MPI
- `./configure -arch=x86_64 -papi=/projects/tau/papi-3.6.2`
`-pdt=<dir> -mpi ; make clean install`
 - Use PAPI counters (one or more) with C/C++/F90 automatic instrumentation. Also instrument the MPI library.
- Typically configure multiple measurement libraries
- Each configuration creates a unique `<arch>/lib/Makefile.tau<options>` stub makefile. It corresponds to the configuration options used. e.g.,
 - `$(PET_HOME)/tau/x86_64/lib/Makefile.tau-mpi-pdt`
 - `$(PET_HOME)/tau/x86_64/lib/Makefile.tau-mpi-papi-pdt`



TAU_SETUP: A GUI for Installing TAU



TAU Measurement Configuration – Examples

```
% cd $(PET_HOME)/tau/x86_64/lib; ls Makefile.*pgi
```

Makefile.tau-pdt

Makefile.tau-mpi-pdt

Makefile.tau-callpath-mpi-pdt

Makefile.tau-mpi-pdt-trace

Makefile.tau-mpi-compensate-pdt

Makefile.tau-mpi-papi-pdt

Makefile.tau-mpi-papi-pdt-trace

Makefile.tau-mpi-papi-pdt-epilog-scalasca-trace

Makefile.tau-pdt...

- For an MPI+F90 application, you may want to start with:

Makefile.tau-mpi-pdt

- Supports MPI instrumentation & PDT for automatic source instrumentation for PGI compilers



TAU's MPI Wrapper Interposition Library

- Uses standard MPI Profiling Interface
 - Provides name shifted interface
 - MPI_Send = PMPI_Send
 - Weak bindings
- Interpose TAU's MPI wrapper library between MPI and TAU
 - -lmpi replaced by -lTauMpi -lpmi -lmpi
- No change to the source code!
 - Just re-link the application to generate performance data
 - export TAU_MAKEFILE=<dir>/<arch>/lib/Makefile.tau-mpi -[options]
 - Use tau_cxx.sh, tau_f90.sh and tau_cc.sh as compilers



Runtime MPI Shared Library Instrumentation

- We can now interpose the MPI wrapper library for applications that have already been compiled
 - No re-compilation or re-linking necessary!
- Uses LD_PRELOAD for Linux
- On AIX, TAU uses MPI_EUILIB / MPI_EUILIBPATH
- Simply compile TAU with MPI support and prefix your MPI program with tauex

```
% mpirun -np 4 tauex a.out
```
- Requires shared library MPI - does not work on XT3
- Approach will work with other shared libraries



-PROFILE Configuration Option

- Generates flat profiles (one for each MPI process)
 - It is the default option.
- Uses wallclock time (gettimeofday() sys call)
- Calculates exclusive, inclusive time spent in each timer and number of calls

% pprof

```
emacs@neutron.cs.uoregon.edu
Buffers Files Tools Edit Search Mule Help
Reading Profile files in profile.*
NODE 0;CONTEXT 0;THREAD 0:
-----
%Time   Exclusive   Inclusive   #Call   #Subrs   Inclusive Name
      msec     total msec
-----
100.0   1           3:11.293   1       15      191293269 applu
99.6    3,667      3:10.463   3       37517   63487925 bcast_inputs
67.1    491        2:08.326   37200   37200   3450 exchange_1
44.5    6,461      1:25.159   9300    18600   9157 buts
41.0    1:18.436   1:18.436   18600   0       4217 MPI_Recv()
29.5    6,778      56,407     9300    18600   6065 blts
26.2    50,142     50,142     19204   0       2611 MPI_Send()
16.2    24,451     31,031     301     602     103096 rhs
3.9     7,501      7,501     9300    0       807 jacld
3.4     838        6,594     604     1812    10918 exchange_3
3.4     6,590      6,590     9300    0       709 jacu
2.6     4,989      4,989     608     0       8206 MPI_Wait()
0.2     0.44       400       1       4       400081 init_comm
0.2     398        399       1       39      399634 MPI_Init()
0.1     140        247       1       47616   247086 setiv
0.1     131        131       57252   0       2 exact
0.1     89         103       1       2       103168 erhs
0.1     0.966      96         1       2       96458 read_input
0.0     95         95        9       0       10603 MPI_Bcast()
0.0     26         44        1       7937   44878 error
0.0     24         24        608     0       40 MPI_Irecv()
0.0     15         15        1       5       15630 MPI_Finalize()
0.0     4          12        1       1700   12335 setbv
0.0     7          8         3       3       2893 l2norm
0.0     3          3         8       0       491 MPI_Allreduce()
0.0     1          3         1       6       3874 pinter
0.0     1          1         1       0       1007 MPI_Barrier()
0.0     0.116     0.837     1       4       837 exchange_4
0.0     0.512     0.512     1       0       512 MPI_Keyval_create()
0.0     0.121     0.353     1       2       353 exchange_5
0.0     0.024     0.191     1       2       191 exchange_6
0.0     0.103     0.103     6       0       17 MPI_Type_contiguous()
-----
--:-- NPB_LU.out (Fundamental)--L8--Top-----
```



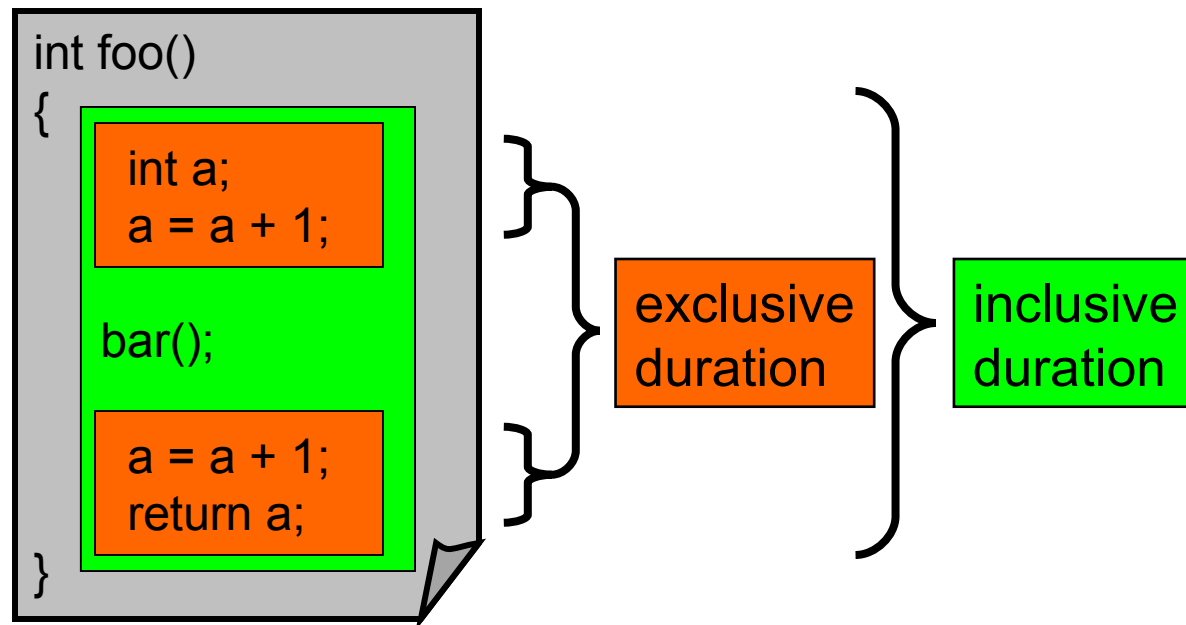
Profiling

- Recording of aggregated information
 - Counts, time, ...
- ... about program and system entities
 - Functions, loops, basic blocks, ...
 - Processes, threads
- Methods
 - Event-based sampling (indirect, statistical)
 - Direct measurement (deterministic)



Inclusive and Exclusive Profiles

- Performance with respect to code regions
- Exclusive measurements for region only
- Inclusive measurements includes child regions



ParaProf Main Window

The screenshot shows the ParaProf Manager interface. On the left is a file browser. The main window displays a thread bar chart with a context menu open over it. The context menu includes options like 'Show Thread Bar Chart', 'Show Thread Statistics Text Window', 'Show Thread Statistics Table', 'Show Thread Call Graph', 'Show Thread Call Path Relations', 'Show User Event Bar Chart', 'Show User Event Statistics Window', and 'Add Thread to Comparison Window'. A red circle highlights the bar chart, and an arrow points from the text 'click left mouse button' to the statistics table window. Another arrow points from the text 'click right mouse button' to the context menu.

click left mouse button

click right mouse button

Metric	Value	Thread
std. dev.	1339496	MPI_Send()
n,c,t 0,0,0	2096225	MPI_Init()
n,c,t 1,0,0	76679	MPI_Bcast()
n,c,t 2,0,0	74890	MAIN [matmult.f90] {39,15}
n,c,t 3,0,0	60711	MPI_Recv()
	19583	INITIALIZE [matmult.f90] {4,18}
	6313	MPI_Finalize()
	4	MPI_Comm_rank()
	3	MPI_Comm_size()



`% paraprof matmult.ppk`

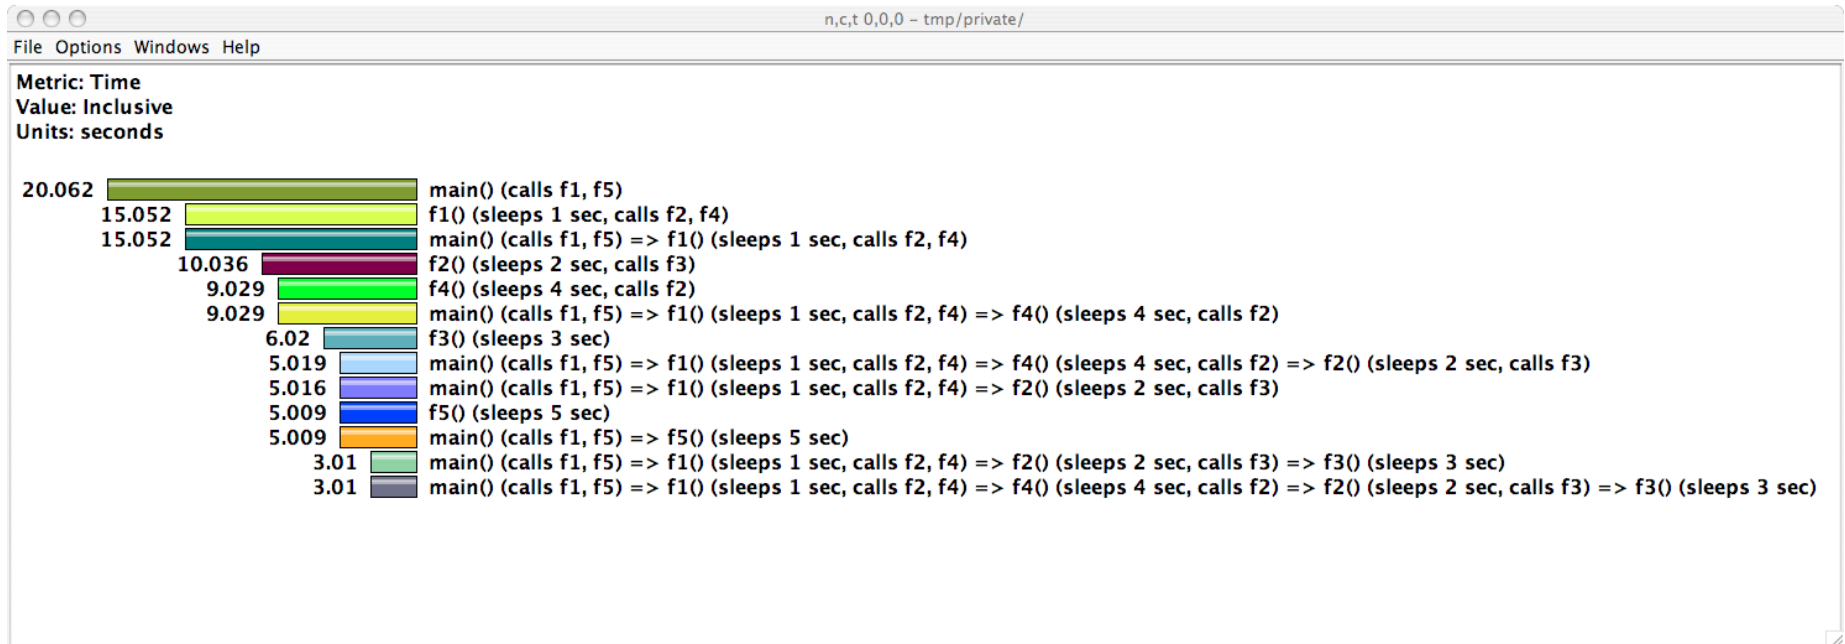
-PAPI Configuration Option

- Instead of one metric, profile or trace with more than one metric
 - % export TAU_METRICS=TIME:PAPI_L2_DCM:PAPI_FP_OPS...
- When used with `–TRACE` option, the first counter **must** be TIME
 - % export TAU_METRICS=TIME:...
 - Provides a globally synchronized real time clock for tracing
- `-papi` appears in the name of the stub Makefile
- Often used with `–papi=<dir>` to measure hardware performance counters and time
- `papi_native_avail` and `papi_avail` are two useful tools



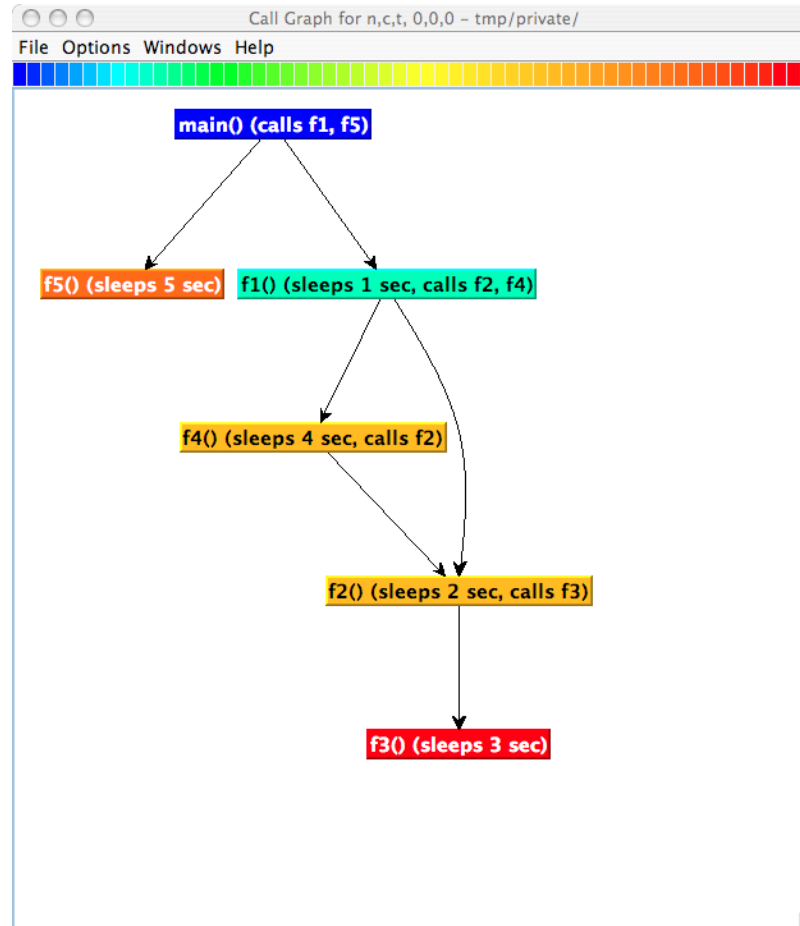
-PROFILECALLPATH Configuration Option

- Generates profiles that show the calling order (edges & nodes in callgraph)
 - A=>B=>C shows the time spent in C when it was called by B and B was called by A
 - Control the depth of callpath using `TAU_CALLPATH_DEPTH` env. Variable
 - `-callpath` in the name of the stub Makefile name
 - In TAU 2.18.2+, any executable can generate callpath profiles using
 - `% export TAU_CALLPATH=1`



-PROFILECALLPATH Configuration Option

- Generates program callgraph



Profile Measurement – Three Flavors

- **Flat profiles**
 - Time (or counts) spent in each routine (nodes in callgraph).
 - Exclusive/inclusive time, no. of calls, child calls
 - E.g.,: MPI_Send, foo, ...
- **Callpath Profiles**
 - Flat profiles, **plus**
 - Sequence of actions that led to poor performance
 - Time spent along a calling path (edges in callgraph)
 - E.g., “main=> f1 => f2 => MPI_Send” shows the time spent in MPI_Send when called by f2, when f2 is called by f1, when it is called by main. Depth of this callpath = 4 (TAU_CALLPATH_DEPTH environment variable)
- **Phase based profiles**
 - Flat profiles, **plus**
 - Flat profiles under a phase (nested phases are allowed)
 - Default “main” phase has all phases and routines invoked outside phases
 - Supports static or dynamic (per-iteration) phases
 - E.g., “IO => MPI_Send” is time spent in MPI_Send in IO phase



-DEPTHLIMIT Configuration Option

- Allows users to enable instrumentation at runtime based on the depth of a calling routine on a callstack.
 - Disables instrumentation in all routines a certain depth away from the root in a callgraph
- TAU_DEPTH_LIMIT environment variable specifies depth
 - % export TAU_DEPTH_LIMIT=1
enables instrumentation in only “main”
 - % export TAU_DEPTH_LIMIT=2
enables instrumentation in main and routines that are directly called by main
- Stub makefile has -depthlimit in its name:
export TAU_MAKEFILE=<taudir>/<arch>/lib/Makefile.tau-mpi-depthlimit-pdt



-COMPENSATE Configuration Option

- Specifies online compensation of performance perturbation
- TAU computes its timer overhead and subtracts it from the profiles
- Works well with time or instructions based metrics
- Does not work with level 1/2 data cache misses
- `export TAU_COMPENSATE=1` (in TAU v2.18.2+)




-TRACE Configuration Option

- Generates event-trace logs, rather than summary profiles
- Traces show when and where an event occurred in terms of location and the process that executed it
- Traces from multiple processes are merged:
% tau_treemerge.pl
 - generates tau.trc and tau.edf as merged trace and event definition file
- TAU traces can be converted to Vampir's OTF/VTF3, Jumpshot SLOG2, Paraver trace formats:
% tau2otf tau.trc tau.edf app.otf
% tau2vtf tau.trc tau.edf app.vpt.gz
% tau2slog2 tau.trc tau.edf -o app.slog2
% tau_convert -paraver tau.trc tau.edf app.prv
- Activated by environment variable
% export TAU_TRACE=1




Tracing Analysis and Visualization

1	master
2	worker
3	...



...			
58	A	ENTER	1
60	B	ENTER	2
62	A	SEND	B
64	A	EXIT	1
68	B	RECV	A
69	B	EXIT	2
...			



Trace Formats

- Different tools produce different formats
 - Differ by event types supported
 - Differ by ASCII and binary representations
 - Vampir Trace Format (VTF)
 - KOJAK (EPILOG)
 - Jumpshot (SLOG-2)
 - Paraver
- Open Trace Format (OTF)
 - Supports interoperation between tracing tools



-PROFILEPARAM Configuration Option

- Idea: partition performance data for individual functions based on runtime parameters
- Enable by configuring with **-PROFILEPARAM**
- TAU call: TAU_PROFILE_PARAM1L (value, "name")
- Simple example:

```
void foo(long input) {  
    TAU_PROFILE("foo", "", TAU_DEFAULT);  
    TAU_PROFILE_PARAM1L(input, "input");  
    ... }  
}
```

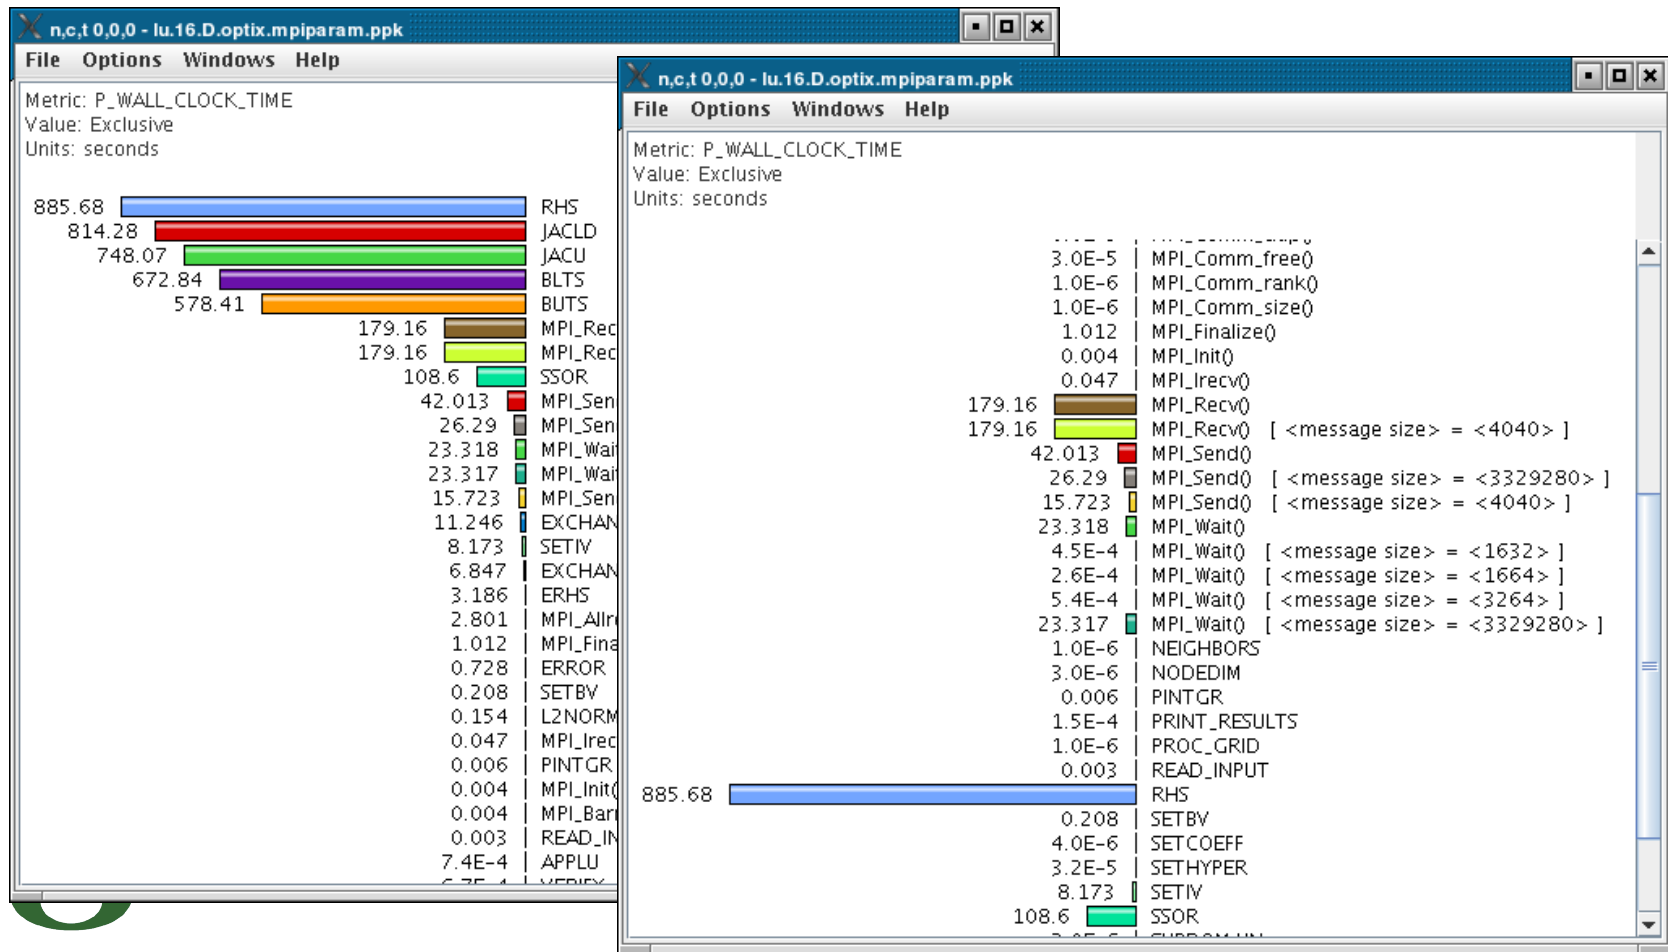
Workload Characterization

- 5 seconds spent in function “foo” becomes
 - 2 seconds for “foo [<input> = <25>]”
 - 1 seconds for “foo [<input> = <5>]”
 - ...
- Currently used in MPI wrapper library
 - Allows for partitioning of time spent in MPI routines based on parameters (message size, message tag, destination node)
 - Can be extrapolated to infer specifics about the MPI subsystem and system as a whole



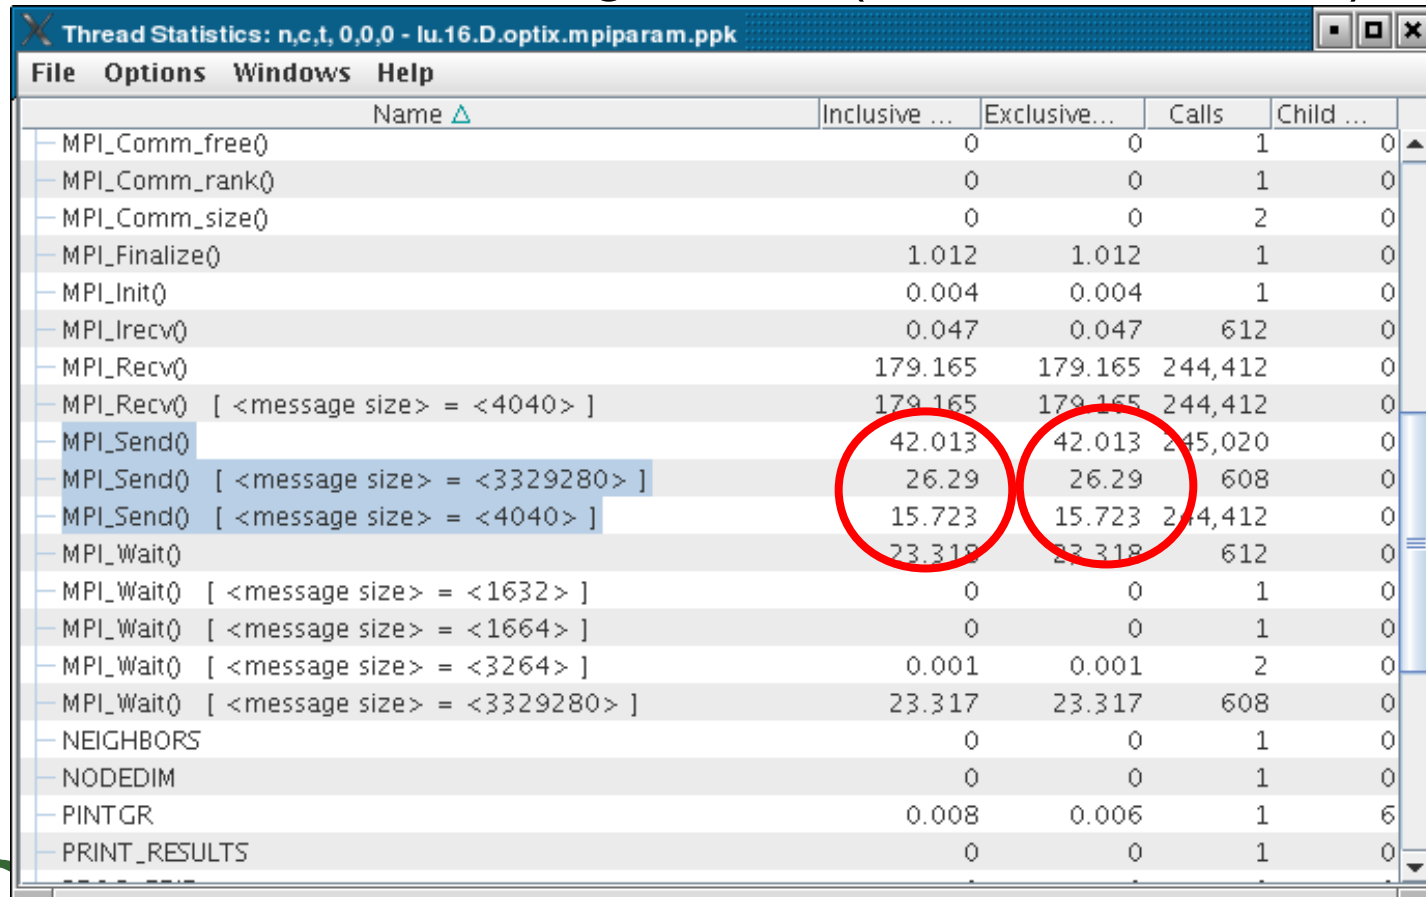
Workload Characterization

- MPI Results (NAS Parallel Benchmark 3.1, LU class D on



Workload Characterization

- Two different message sizes (~3.3MB and ~4K)



Thread Statistics: n,c,t, 0,0,0 - lu.16.D.optix.mpiparam.ppk

Name	Inclusive ...	Exclusive...	Calls	Child ...
MPI_Comm_free()	0	0	1	0
MPI_Comm_rank()	0	0	1	0
MPI_Comm_size()	0	0	2	0
MPI_Finalize()	1.012	1.012	1	0
MPI_Init()	0.004	0.004	1	0
MPI_Irecv()	0.047	0.047	612	0
MPI_Recv()	179.165	179.165	244,412	0
MPI_Recv() [<message size> = <4040>]	179.165	179.165	244,412	0
MPI_Send()	42.013	42.013	245,020	0
MPI_Send() [<message size> = <3329280>]	26.29	26.29	608	0
MPI_Send() [<message size> = <4040>]	15.723	15.723	244,412	0
MPI_Wait()	23.318	23.318	612	0
MPI_Wait() [<message size> = <1632>]	0	0	1	0
MPI_Wait() [<message size> = <1664>]	0	0	1	0
MPI_Wait() [<message size> = <3264>]	0.001	0.001	2	0
MPI_Wait() [<message size> = <3329280>]	23.317	23.317	608	0
NEIGHBORS	0	0	1	0
NODEDIM	0	0	1	0
PINTGR	0.008	0.006	1	6
PRINT_RESULTS	0	0	1	0

Memory Profiling in TAU

- Configuration option **-PROFILEMEMORY**
 - Records global heap memory utilization for each function
 - Takes one sample at beginning of each function and associates the sample with **function name**
- Configuration option **-PROFILEHEADROOM**
 - Records headroom (amount of free memory to grow) for each function
 - Takes one sample at beginning of each function and associates it with the **callstack** [TAU_CALLPATH_DEPTH env variable]
 - Useful for debugging memory usage on IBM BG/L.
- Independent of instrumentation/measurement options selected
- No need to insert macros/calls in the source code
- User defined atomic events appear in profiles/traces



Memory Profiling in TAU (Atomic events)

Sorted By: number of userEvents

NumSamples	Max	Min	Mean	Std. Dev	Name
252032	2022.7	1181.2	1534.3	410.04	MODULEHYDRO_1D::HYDRO_1D - Heap Memory (KB)
252032	2022.8	1181.7	1534.3	410.04	MODULEINTRFC::INTRFC - Heap Memory (KB)
104559	2023.2	331.13	1526.6	409.54	MODULEEOS3D::EOS3D - Heap Memory (KB)
63008	2022.7	1182	1534.3	410.01	MODULEUPDATE_SOLN::UPDATE_SOLN - Heap Memory (KB)
55545	2023.3	333.07	1514.2	408.31	DBASETREE::DBASENEIGHBORBLOCKLIST - Heap Memory (KB)
51374	2023	1179.4	1497.7	402.53	AMR_PROLONG_GEN_UNK_FUN - Heap Memory (KB)
42120	2022.7	1187.5	1533.5	409.83	ABUNDANCE_RESTRICT - Heap Memory (KB)
41958	2023	346.12	1514.9	408.39	AMR_RESTRICT_UNK_FUN - Heap Memory (KB)
31832	2022.8	1187.4	1534.1	409.91	AMR_RESTRICT_RED - Heap Memory (KB)
31504	2022.7	1181.8	1534.3	410.04	DIFFUSE - Heap Memory (KB)
26042	2023	1179.2	1501.9	403.61	AMR_PROLONG_UNK_FUN - Heap Memory (KB)

Flash2 code profile (-PROFILEMEMORY) on IBM BlueGene/L [MPI rank 0]



Detecting Memory Leaks in C/C++

- TAU wrapper library for malloc/realloc/free
- During instrumentation, specify
 - optDetectMemoryLeaks option to TAU_COMPILER
 - % export TAU_OPTIONS='-optVerbose -optDetectMemoryLeaks'
 - % export TAU_MAKEFILE=<taudir>/<arch>/lib/Makefile.tau-mpi-pdt...
 - % tau_cxx.sh foo.cpp ...
- Tracks each memory allocation/de-allocation in parsed files
- Correlates each memory event with the executing callstack
- At the end of execution, TAU detects memory leaks
- TAU reports leaks based on allocations and the executing callstack
- Set **TAU_CALLPATH_DEPTH** environment variable to limit callpath data
 - default is 2
- Future work
 - Support for C++ new/delete planned
 - Support for Fortran 90/95 allocate/deallocate planned



Memory Leak Detection

The screenshot displays three windows from a memory leak detection tool. The top window shows a list of memory events with colored bars representing memory sizes. The second window is a detailed view of a 'MEMORY LEAK!' event, showing a mean value of 52 and a standard deviation of 0. The bottom window is a table of sorted memory events.

Event Window 1:

- 180 malloc size <file=simple.inst.cpp, line=26>
- 180 malloc size <file=simple.inst.cpp, line=26> : int main(int, char **) => int foo(int) => int bar(int)
- 180 free size <file=simple.inst.cpp, line=28>
- 180 free size <file=simple.inst.cpp, line=28> : int main(int, char **) => int foo(int) => int bar(int)
- 80 malloc size <file=simple.inst.cpp, line=18>
- 80 malloc size <file=simple.inst.cpp, line=18> : int main(int, char **) => int foo(int) => int g(int) => int bar(int)
- 80 free size <file=simple.inst.cpp, line=21>
- 80 free size <file=simple.inst.cpp, line=21> : int main(int, char **) => int foo(int) => int g(int) => int bar(int)
- 52 MEMORY LEAK! malloc size <file=simple.inst.cpp, line=18> : int main(int, char **) => int foo(int) => int g(int) => int bar(int)

Event Window 2:

Name: MEMORY LEAK! malloc size <file=simple.inst.cpp, line=18> : int main(int, char **) => int foo(int) => int g(int) => int bar(int)

Value Type: Max Value

52 Mean
52 n,c,t 0,0,0
0 Std. Dev.

Sorted List:

NumSamples	Max	Min	Mean	Std. Dev	Name
3	80	48	60	14.236	malloc size <file=simple.inst.cpp, line=18>
3	80	48	60	14.236	malloc size <file=simple.inst.cpp, line=18> : int main(int, char **
2	52	48	50	2	MEMORY LEAK! malloc size <file=simple.inst.cpp, line=18> : int ma
1	80	80	80	0	free size <file=simple.inst.cpp, line=21>
1	80	80	80	0	free size <file=simple.inst.cpp, line=21> : int main(int, char **
1	180	180	180	0	malloc size <file=simple.inst.cpp, line=26>
1	180	180	180	0	malloc size <file=simple.inst.cpp, line=26> : int main(int, char
1	180	180	180	0	free size <file=simple.inst.cpp, line=28>
1	180	180	180	0	free size <file=simple.inst.cpp, line=28> : int main(int, char **

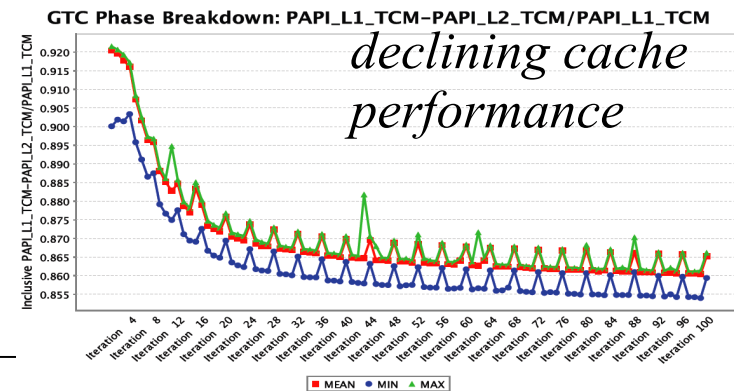

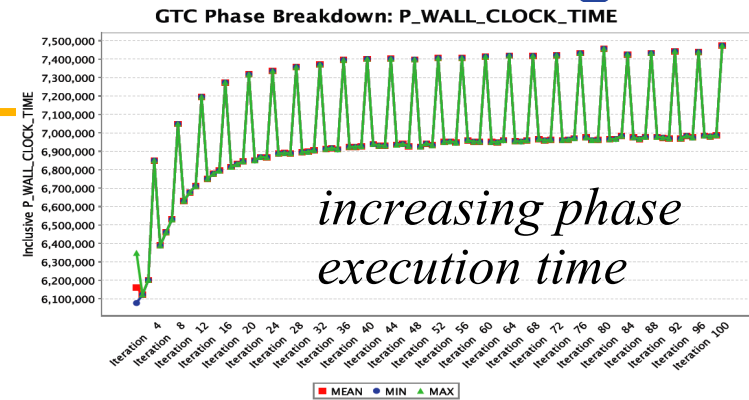
TAU Timers and Phases

- **Static timer**
 - Shows time spent in all invocations of a routine (foo)
 - E.g., “foo()” 100 secs, 100 calls
- **Dynamic timer**
 - Shows time spent in each invocation of a routine
 - E.g., “foo() 3” 4.5 secs, “foo 10” 2 secs (invocations 3 and 10 respectively)
- **Static phase**
 - Shows time spent in all routines called (directly/indirectly) by a given routine (foo)
 - E.g., “foo() => MPI_Send()” 100 secs, 10 calls shows that a total of 100 secs were spent in MPI_Send() when it was called by foo.
- **Dynamic phase**
 - Shows time spent in all routines called by a given invocation of a routine.
 - E.g., “foo() 4 => MPI_Send()” 12 secs, shows that 12 secs were spent in MPI_Send when it was called by the 4th invocation of foo.



Performance Dynamics: Phase-Based Profiling

- Profile phases capture performance with respect to application-defined 'phases' of execution
 - Separate full profile produced for each phase
- GTC particle-in-cell simulation of fusion turbulence
- Phases assigned to iterations
- Data change affects cache



TAU_COMPILER Commandline Options

- See `<taudir>/<arch>/bin/tau_compiler.sh -help`

- Compilation:

```
% mpx1f90 -c foo.f90
```

Changes to

```
% f95parse foo.f90 $(OPT1)
```

```
% tau_instrumentor foo.pdb foo.f90 -o foo.inst.f90 $(OPT2)
```

```
% ftn -c foo.inst.f90 $(OPT3)
```

- Linking:

```
% ftn foo.o bar.o -o app
```

Changes to

```
% ftn foo.o bar.o -o app $(OPT4)
```

- Where options `OPT[1-4]` default values may be overridden by the user:
`F90 = tau_f90.sh`



TAU_COMPILER Options

- Optional parameters for $\$(TAU_COMPILER)$: [tau_compiler.sh -help]
 - optVerbose Turn on verbose debugging messages
 - optCompInst Use compiler based instrumentation
 - optDetectMemoryLeaks Turn on debugging memory allocations/
de-allocations to track leaks
 - optKeepFiles Does not remove intermediate .pdb and .inst.* files
 - optPreProcess Preprocess Fortran sources before instrumentation
 - optTauSelectFile="" Specify selective instrumentation file for tau_instrumentor
 - 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)
 - optPdtF95Reset="" Reset options for Fortran parser in PDT (f95parse/gfparse)
 - optPdtCOpts="" Options for C parser in PDT (cparse). Typically
 $\$(TAU_MPI_INCLUDE)$ $\$(TAU_INCLUDE)$ $\$(TAU_DEFS)$
 - optPdtCxxOpts="" Options for C++ parser in PDT (cxxparse). Typically
 $\$(TAU_MPI_INCLUDE)$ $\$(TAU_INCLUDE)$ $\$(TAU_DEFS)$
 - ...



Compiling Codes with TAU

- If your Fortran code uses free format in .f files (fixed is default for .f), you may use:
`% export TAU_OPTIONS='-optPdtF95Opts="-R free" -optVerbose '`
- To use the compiler based instrumentation instead of PDT (source-based):
`% export TAU_OPTIONS='-optComplnst -optVerbose'`
- If your Fortran code uses C preprocessor directives (#include, #ifdef, #endif):
`% export TAU_OPTIONS='-optPreProcess -optVerbose -optDetectMemoryLeaks'`
- To use an instrumentation specification file:
`% export TAU_OPTIONS='-optTauSelectFile=mycmd.tau -optVerbose -optPreProcess'`
`% cat mycmd.tau`
BEGIN_INSTRUMENT_SECTION
memory file="foo.f90" routine="#"
instruments all allocate/deallocate statements in all routines in foo.f90
loops file="*" routine="#"
io file="abc.f90" routine="FOO"
END_INSTRUMENT_SECTION



Optimization of Program Instrumentation

- Need to eliminate instrumentation in frequently executing lightweight routines
- Throttling of events at runtime (default in tau-2.17.2+):

```
% export TAU_THROTTLE=1
```

Turns off instrumentation in routines that execute over 100000 times (TAU_THROTTLE_NUMCALLS) and take less than 10 microseconds of inclusive time per call (TAU_THROTTLE_PERCALL). Use TAU_THROTTLE=0 to disable.
- Selective instrumentation file to filter events

```
% tau_instrumentor [options] -f <file> OR  
% export TAU_OPTIONS=' -optTauSelectFile=tau.txt'
```
- Compensation of local instrumentation overhead

```
% configure -COMPENSATE  
or  
% export TAU_COMPENSATE=1 (in tau-2.18.2+)
```

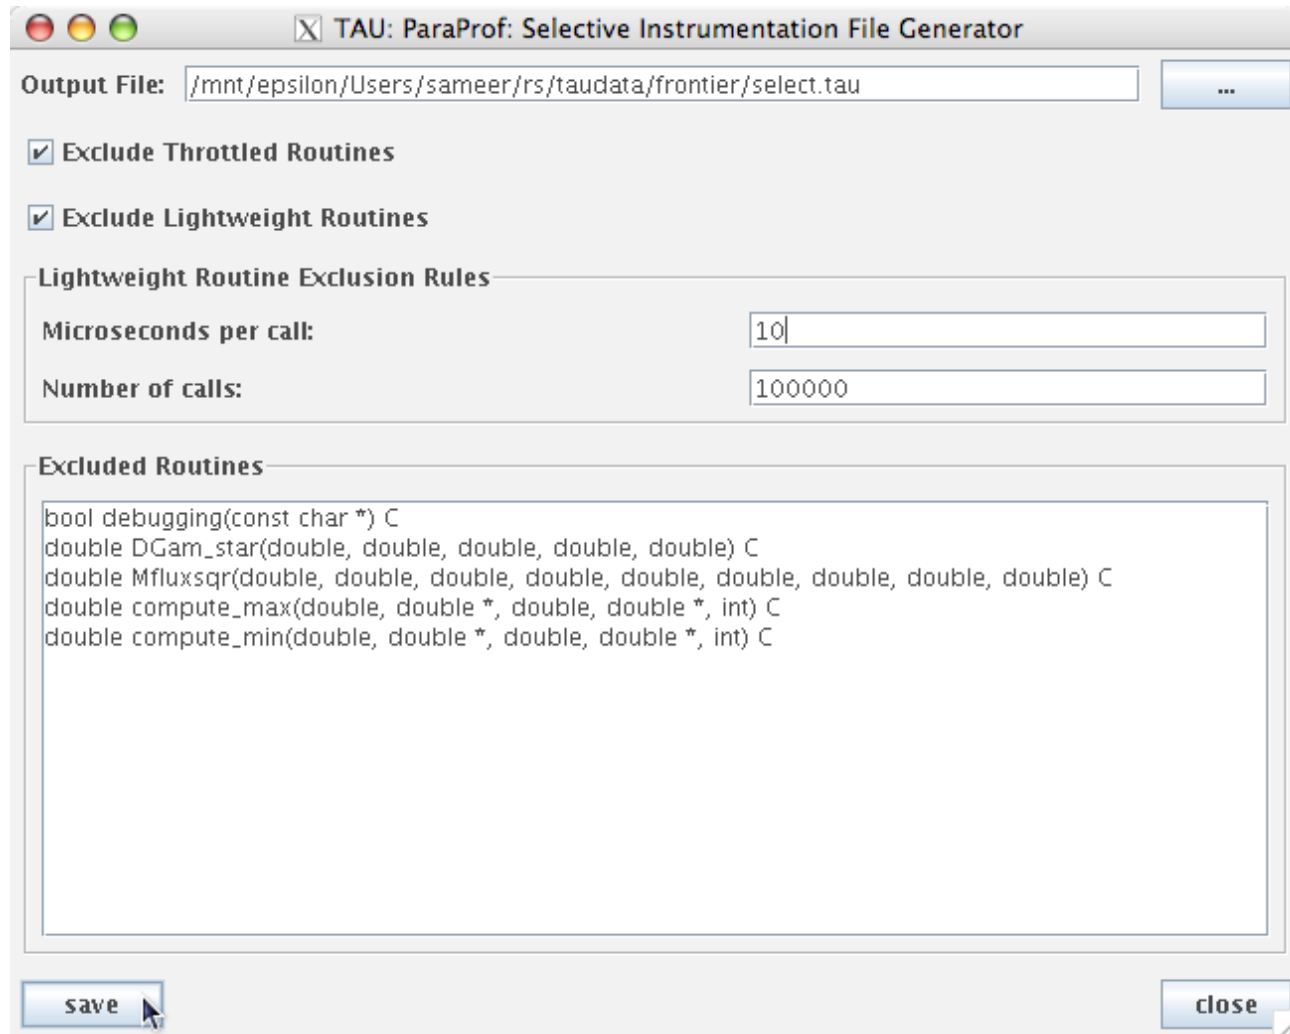


ParaProf: Creating Selective Instrumentation File

TrialField	Value
Name	200m4_p256.ppk
Application ID	0
Experiment ID	0
Trial ID	0
BGP Coords	(7,3,7)
BGP DDRSize (MB)	2048
BGP Location	R00-M1-N15-J32
BGP Node Mode	Coprocessor (22270944)
BGP Processor ID	0
BGP Size	(8,4,8)
BGP isTorus	(0,0,0)
BGP numNodesInPset	1
BGP numPsets	256
BGP psetNum	3
BGP rankInPset	24
CPU Type	450 Blue Gene/P DD2
CWD	/gpfs/home/kaman/FronTier/src/gas
Executable	/sbin.rd/ioproxy
Hostname	ion-16
Local Time	2008-08-22T12:50:33-05:00
MPI Processor Name	Rank 256 of 256 <7,3,7,0> R00-M1-N15-J32
Memory Size	1816608 kB
Node Name	ion-16
OS Machine	BGP
OS Name	CNK
OS Release	2.6.19.2
OS Version	1
Starting Timestamp	1219427292054274
TAU Architecture	bgp
TAU Config	-arch=bgp -pdt=/soft/apps/tau/pdtoolkit-3.12 -...
TAU Version	2.17.1
Timestamp	1219427456121879
UTC Time	2008-08-22T17:50:33Z
pid	355



Choosing Rules for Excluding Routines



The screenshot shows a dialog box titled "TAU: ParaProf: Selective Instrumentation File Generator". It has a standard macOS-style title bar with red, yellow, and green window control buttons. The "Output File:" field contains the path "/mnt/epsilon/Users/sameer/rs/taudata/frontier/select.tau" and has a "..." button to its right. Below this are two checked checkboxes: "Exclude Throttled Routines" and "Exclude Lightweight Routines". A section titled "Lightweight Routine Exclusion Rules" contains two input fields: "Microseconds per call:" with the value "10" and "Number of calls:" with the value "100000". At the bottom of this section is a text area labeled "Excluded Routines" containing the following C function signatures:

```
bool debugging(const char *) C
double DGam_star(double, double, double, double, double) C
double Mfluxsqr(double, double, double, double, double, double, double, double, double) C
double compute_max(double, double *, double, double *, int) C
double compute_min(double, double *, double, double *, int) C
```

At the bottom of the dialog are "save" and "close" buttons.



Selective Instrumentation File

- Specify a list of routines to exclude or include (case sensitive)
- # is a wildcard in a routine name. It cannot appear in the first column.

```
BEGIN_EXCLUDE_LIST  
Foo  
Bar  
D#EMM  
END_EXCLUDE_LIST
```

- Specify a list of routines to include for instrumentation

```
BEGIN_INCLUDE_LIST  
int main(int, char **)  
F1  
F3  
END_INCLUDE_LIST
```

- Specify either an include list or an exclude list!



Selective Instrumentation File

- Optionally specify a list of files to exclude or include (case sensitive)

- * and ? may be used as wildcard characters in a file name

```
BEGIN_FILE_EXCLUDE_LIST  
f*.f90  
Foo?.cpp  
END_FILE_EXCLUDE_LIST
```

- Specify a list of routines to include for instrumentation

```
BEGIN_FILE_INCLUDE_LIST  
main.cpp  
foo.f90  
END_FILE_INCLUDE_LIST
```



Selective Instrumentation File

- User instrumentation commands are placed in INSTRUMENT section
- ? and * used as wildcard characters for file name, # for routine name
- \ as escape character for quotes
- Routine entry/exit, arbitrary code insertion
- Outer-loop level instrumentation

```
BEGIN_INSTRUMENT_SECTION
loops file="foo.f90" routine="matrix#"
memory file="foo.f90" routine="#"
io routine="matrix#"
[static/dynamic] phase routine="MULTIPLY"
dynamic [phase/timer] name="foo" file="foo.cpp" line=22 to line=35
file="foo.f90" line = 123 code = " print *, \" Inside foo\""
exit routine = "int foo()" code = "cout <<\"exiting foo\"<<endl;"
END_INSTRUMENT_SECTION
```



Instrumentation of OpenMP Constructs

- OpenMP **P**ragma **A**nd **R**egion **I**nstrumentor [UTK, FZJ]
- Source-to-Source translator to insert **POMP** calls around OpenMP constructs and API functions
- **Done:** Supports
 - Fortran77 and Fortran90, OpenMP 2.0
 - C and C++, OpenMP 1.0
 - POMP Extensions
 - EPILOG and TAU POMP implementations
 - Preserves source code information (`#line line file`)
- **tau_ompcheck**
 - Balances OpenMP constructs (DO/END DO) and detects errors
 - Invoked by `tau_compiler.sh` prior to invoking Opari
- KOJAK Project website <http://icl.cs.utk.edu/kojak>



OpenMP API Instrumentation

- Transform

- `omp_#_lock()` → `pomp_#_lock()`
- `omp_##_nest_lock()` → `pomp_##_nest_lock()`

[# = init | destroy | set | unset | test]

- POMP version

- Calls omp version internally
- Can do extra stuff before and after call



Dynamic Instrumentation

- TAU uses DyninstAPI for runtime code patching
- Developed by U. Wisconsin and U. Maryland
- <http://www.dyninst.org>
- ***tau_run*** (mutator) loads measurement library
- Instruments mutatee
- MPI issues:
 - one mutator per executable image [TAU, DynaProf]
 - one mutator for several executables [Paradyn, DPCL]



Virtual Machine Performance Instrumentation

- **Integrate performance system with VM**
 - Captures robust performance data (e.g., thread events)
 - Maintain features of environment
 - portability, concurrency, extensibility, interoperation
 - Allow use in optimization methods
- **JVM Profiling Interface (JVMPi)**
 - Generation of JVM events and hooks into JVM
 - Profiler agent (TAU) loaded as shared object
 - registers events of interest and address of callback routine
 - Access to information on dynamically loaded classes
 - **No need to modify Java source, bytecode, or JVM**

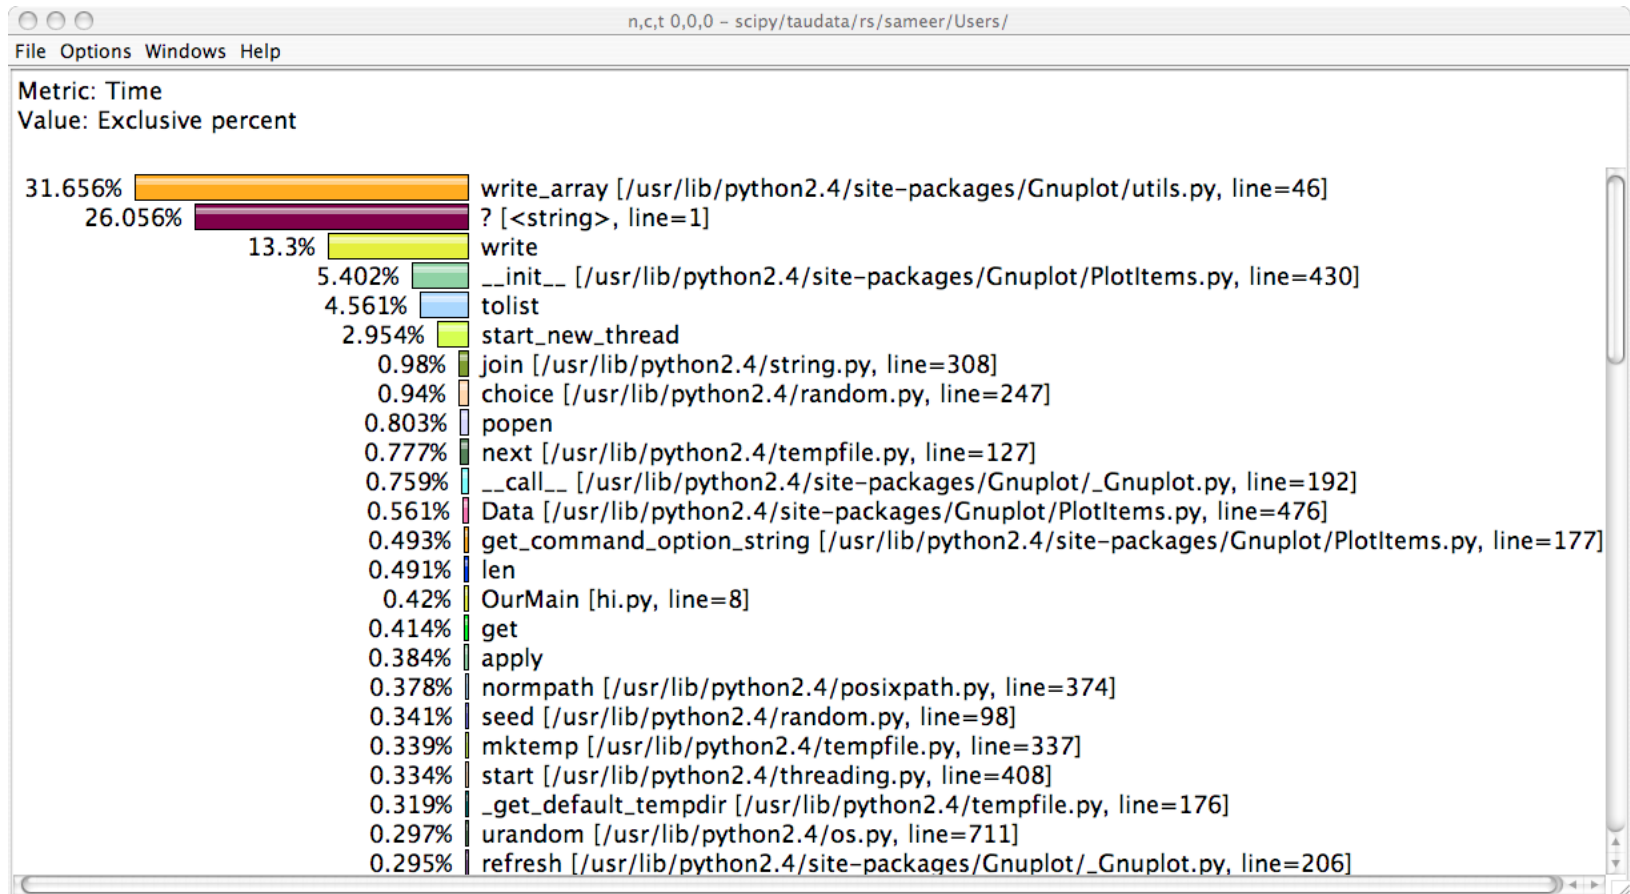


Generate a Python Profile

```
% export TAU_MAKEFILE=/projects/tau/tau_latest/ibm64
    /lib/Makefile.tau-python-pdt
% set path=(/projects/tau/tau_latest/ibm64/bin $path)
% cat wrapper.py
import tau
def OurMain():
    import foo
    tau.run('OurMain()')
Uninstrumented:
% ./foo.py
Instrumented:
% export PYTHONPATH= <taudir>/ibm64/<lib>/bindings-python-pdt
(same options string as TAU_MAKEFILE)
% export LD_LIBRARY_PATH=<taudir>/x86_64/lib/bindings-python-pdt:
$LD_LIBRARY_PATH
% ./wrapper.py

Wrapper invokes foo and generates performance data
% pprof/paraprof
```

Python Instrumentation: SciPy

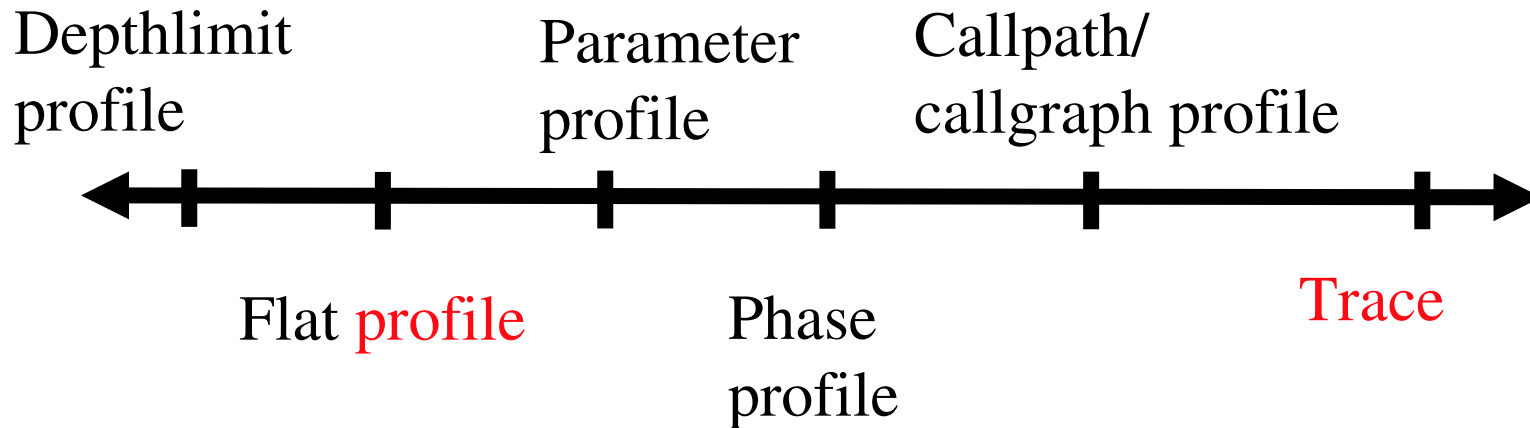


Critical issues

- Accuracy
 - Timing and counting accuracy depends on resolution
 - Any performance measurement generates overhead
 - Execution on performance measurement code
 - Measurement overhead can lead to intrusion
 - Intrusion can cause perturbation
 - alters program behavior
- Granularity
 - How many measurements are made
 - How much overhead per measurement
- Tradeoff (general wisdom)
 - Accuracy is inversely correlated with granularity



Performance Evaluation Alternatives



Each alternative has:

- one metric/counter
- multiple counters

Volume of performance data



Profiling / Tracing Comparison

- Profiling
 - ☺ Finite, bounded performance data size
 - ☺ Applicable to both direct and indirect methods
 - ☹ Loses time dimension (not entirely)
 - ☹ Lacks ability to fully describe process interaction
- Tracing
 - ☺ Temporal and spatial dimension to performance data
 - ☺ Capture parallel dynamics and process interaction
 - ☹ Some inconsistencies with indirect methods
 - ☹ Unbounded performance data size (large)
 - ☹ Complex event buffering and clock synchronization



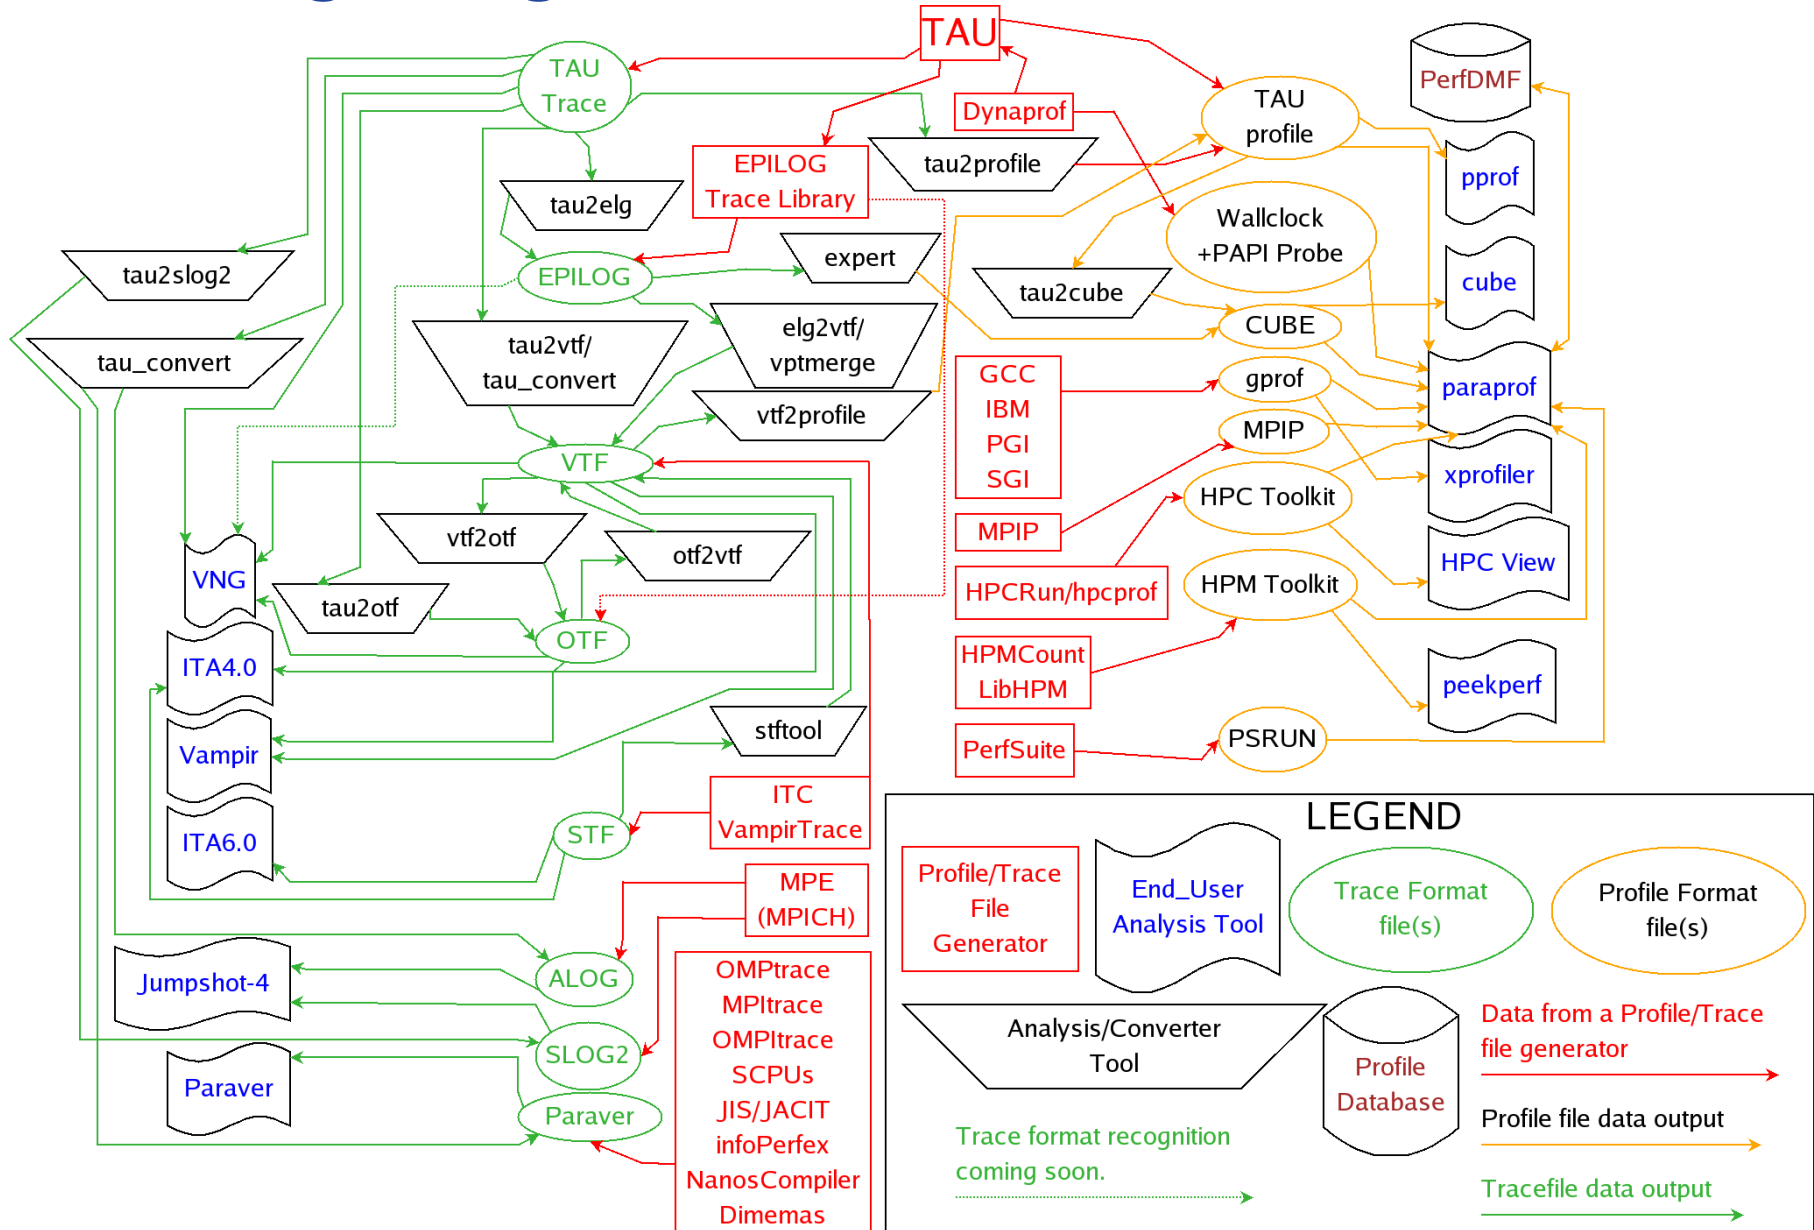
TAU Performance System Interfaces

- PDT [U. Oregon, LANL, FZJ] for instrumentation of C++, C99, F95 source code
- PAPI [UTK] for accessing hardware performance counters data
- DyninstAPI [U. Maryland, U. Wisconsin] for runtime instrumentation
- KOJAK [FZJ, UTK]
 - Epilog trace generation library
 - CUBE callgraph visualizer
 - Opari OpenMP directive rewriting tool
- Vampir/VNG Trace Analyzer [TU Dresden]
- VTF3/OTF trace generation library [TU Dresden] (available from TAU website)
- Paraver trace visualizer [CEPBA]
- Jumpshot-4 trace visualizer [MPICH, ANL]
- JVMPI from JDK for Java program instrumentation [Sun]
- Paraprof profile browser/PerfDMF database supports:
 - TAU format
 - Gprof [GNU]
 - HPM Toolkit [IBM]
 - MpiP [ORNL, LLNL]
 - Dynaprof [UTK]
 - PSR in INC.SAI





Building Bridges to Other Tools: TAU





ParaProf – Manager Window

The screenshot shows the TAU: ParaProf Manager interface. On the left is a hierarchical file tree under 'Applications'. On the right is a table with columns 'Field' and 'Value'.

Field	Value
Name	64 CPU
Application ID	4
Experiment ID	26
Trial ID	85
DATE	
COLLECTORID	
NODE_COUNT	64
CONTEXTS_PER_NODE	1
THREADS_PER_CONTEXT	1

performance database

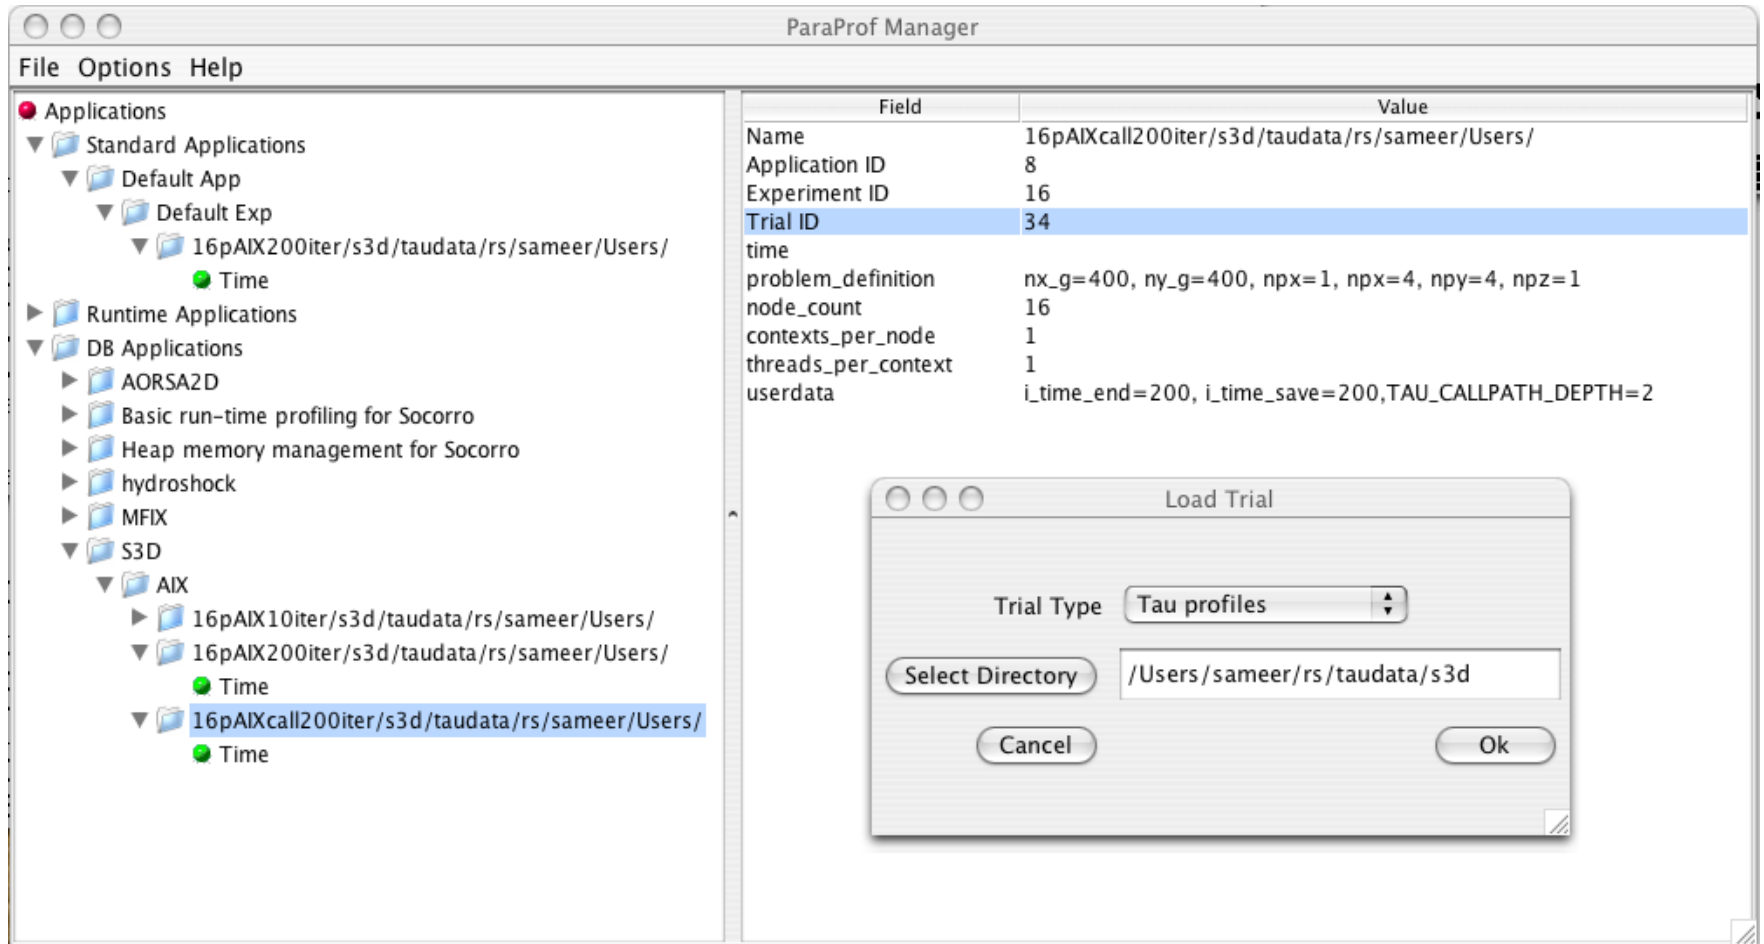
metadata

The 'Load Trial' dialog box is open, showing a dropdown menu for 'Trial Type'. The menu is expanded to show the following options:

- Tau profiles
- Tau pprof.dat
- Dynaprof
- MpiP
- HPMToolkit
- Gprof
- PSRun
- ParaProf Packed Profile
- Cube
- HPCToolkit



Performance Database: Storage of MetaData



ParaProf Manager

File Options Help

- Applications
 - Standard Applications
 - Default App
 - Default Exp
 - 16pAIX200iter/s3d/taudata/rs/sameer/Users/
 - Time
- Runtime Applications
- DB Applications
 - AORSA2D
 - Basic run-time profiling for Socorro
 - Heap memory management for Socorro
 - hydroshock
 - MFIX
 - S3D
 - AIX
 - 16pAIX10iter/s3d/taudata/rs/sameer/Users/
 - 16pAIX200iter/s3d/taudata/rs/sameer/Users/
 - Time
 - 16pAIXcall200iter/s3d/taudata/rs/sameer/Users/
 - Time

Field	Value
Name	16pAIXcall200iter/s3d/taudata/rs/sameer/Users/
Application ID	8
Experiment ID	16
Trial ID	34
time	
problem_definition	nx_g=400, ny_g=400, npx=1, npx=4, npy=4, npz=1
node_count	16
contexts_per_node	1
threads_per_context	1
userdata	i_time_end=200, i_time_save=200,TAU_CALLPATH_DEPTH=2

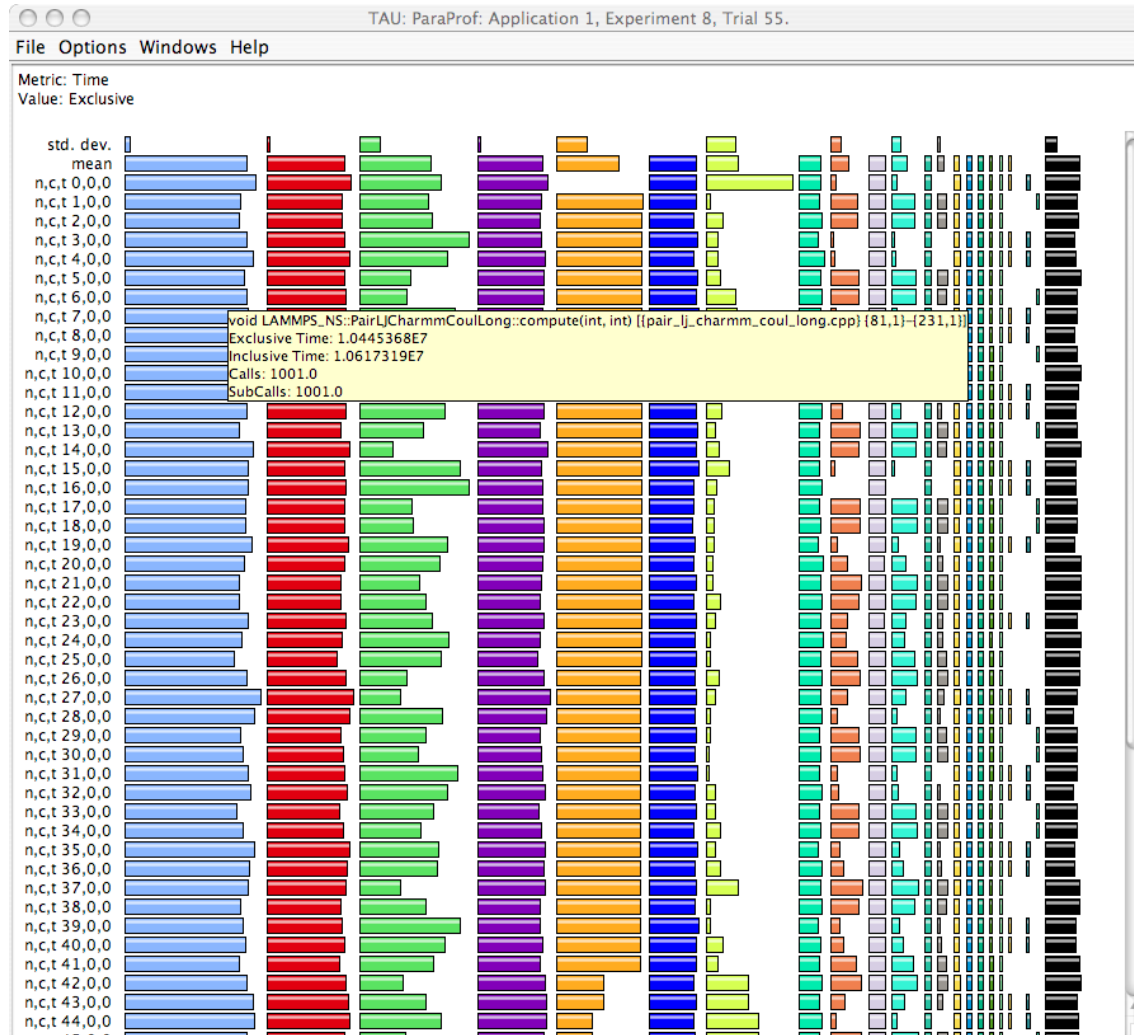
Load Trial

Trial Type: Tau profiles

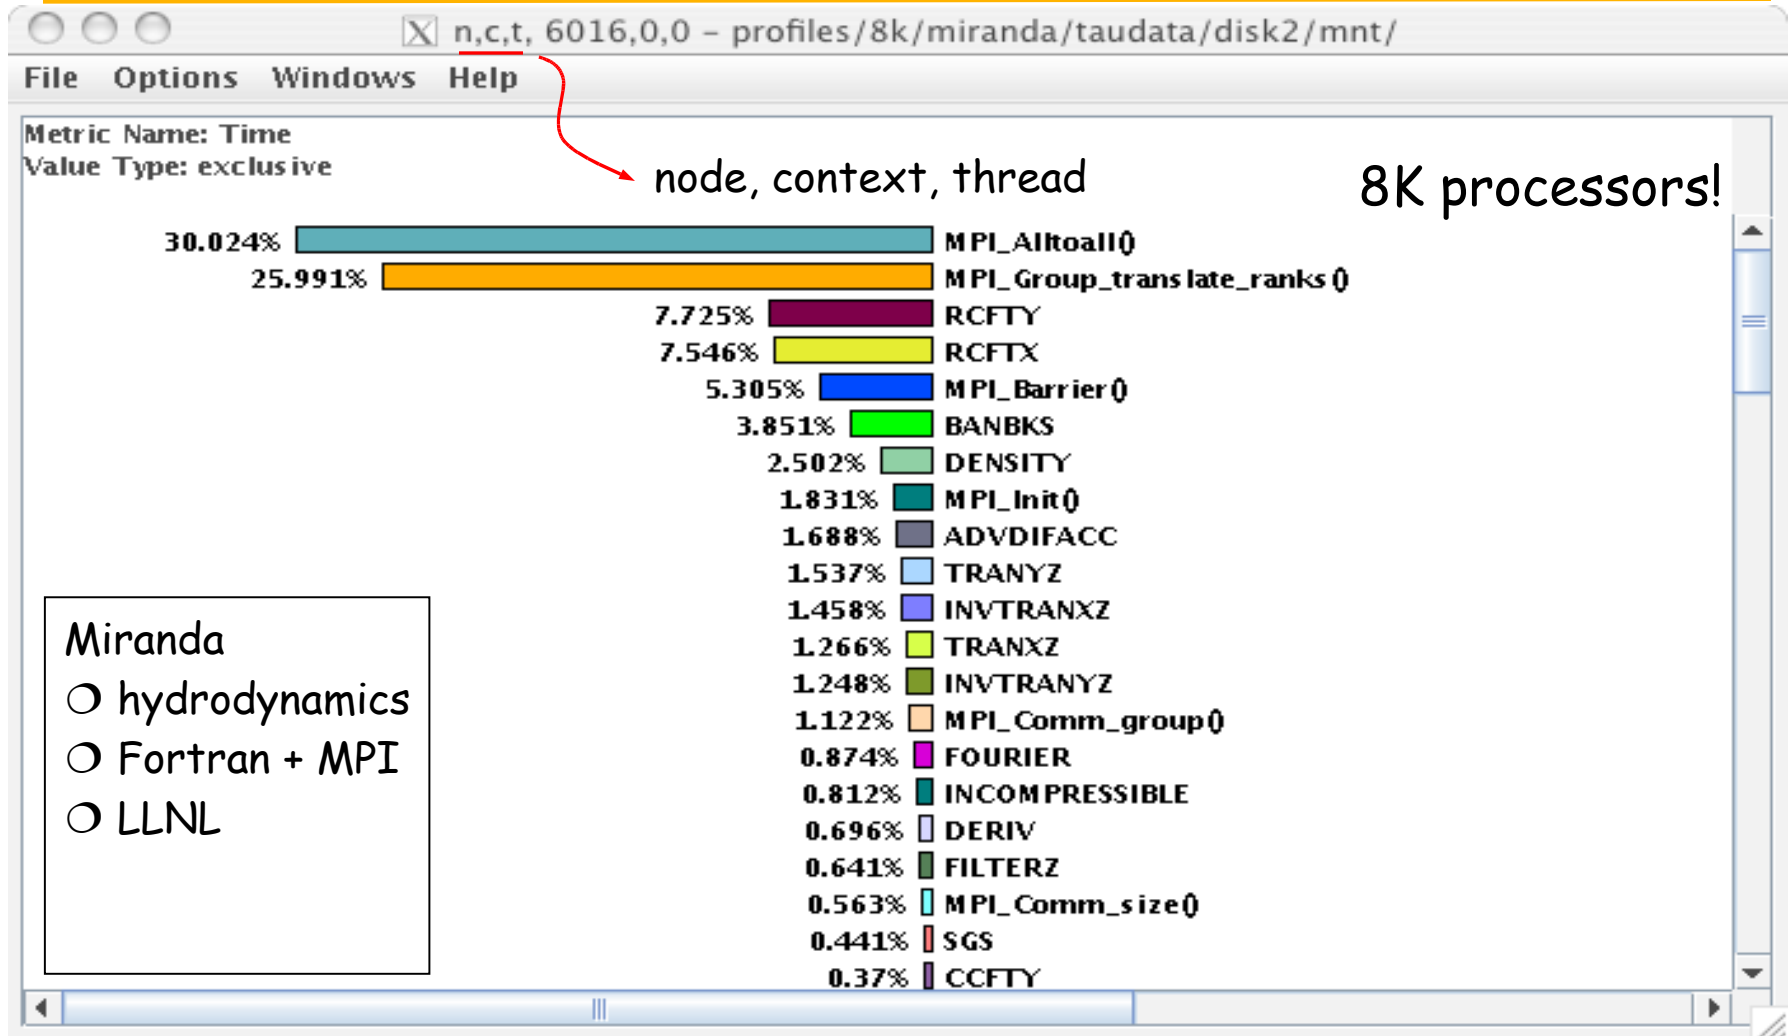
Select Directory: /Users/sameer/rs/taudata/s3d

Cancel Ok

ParaProf Main Window (Lammps)

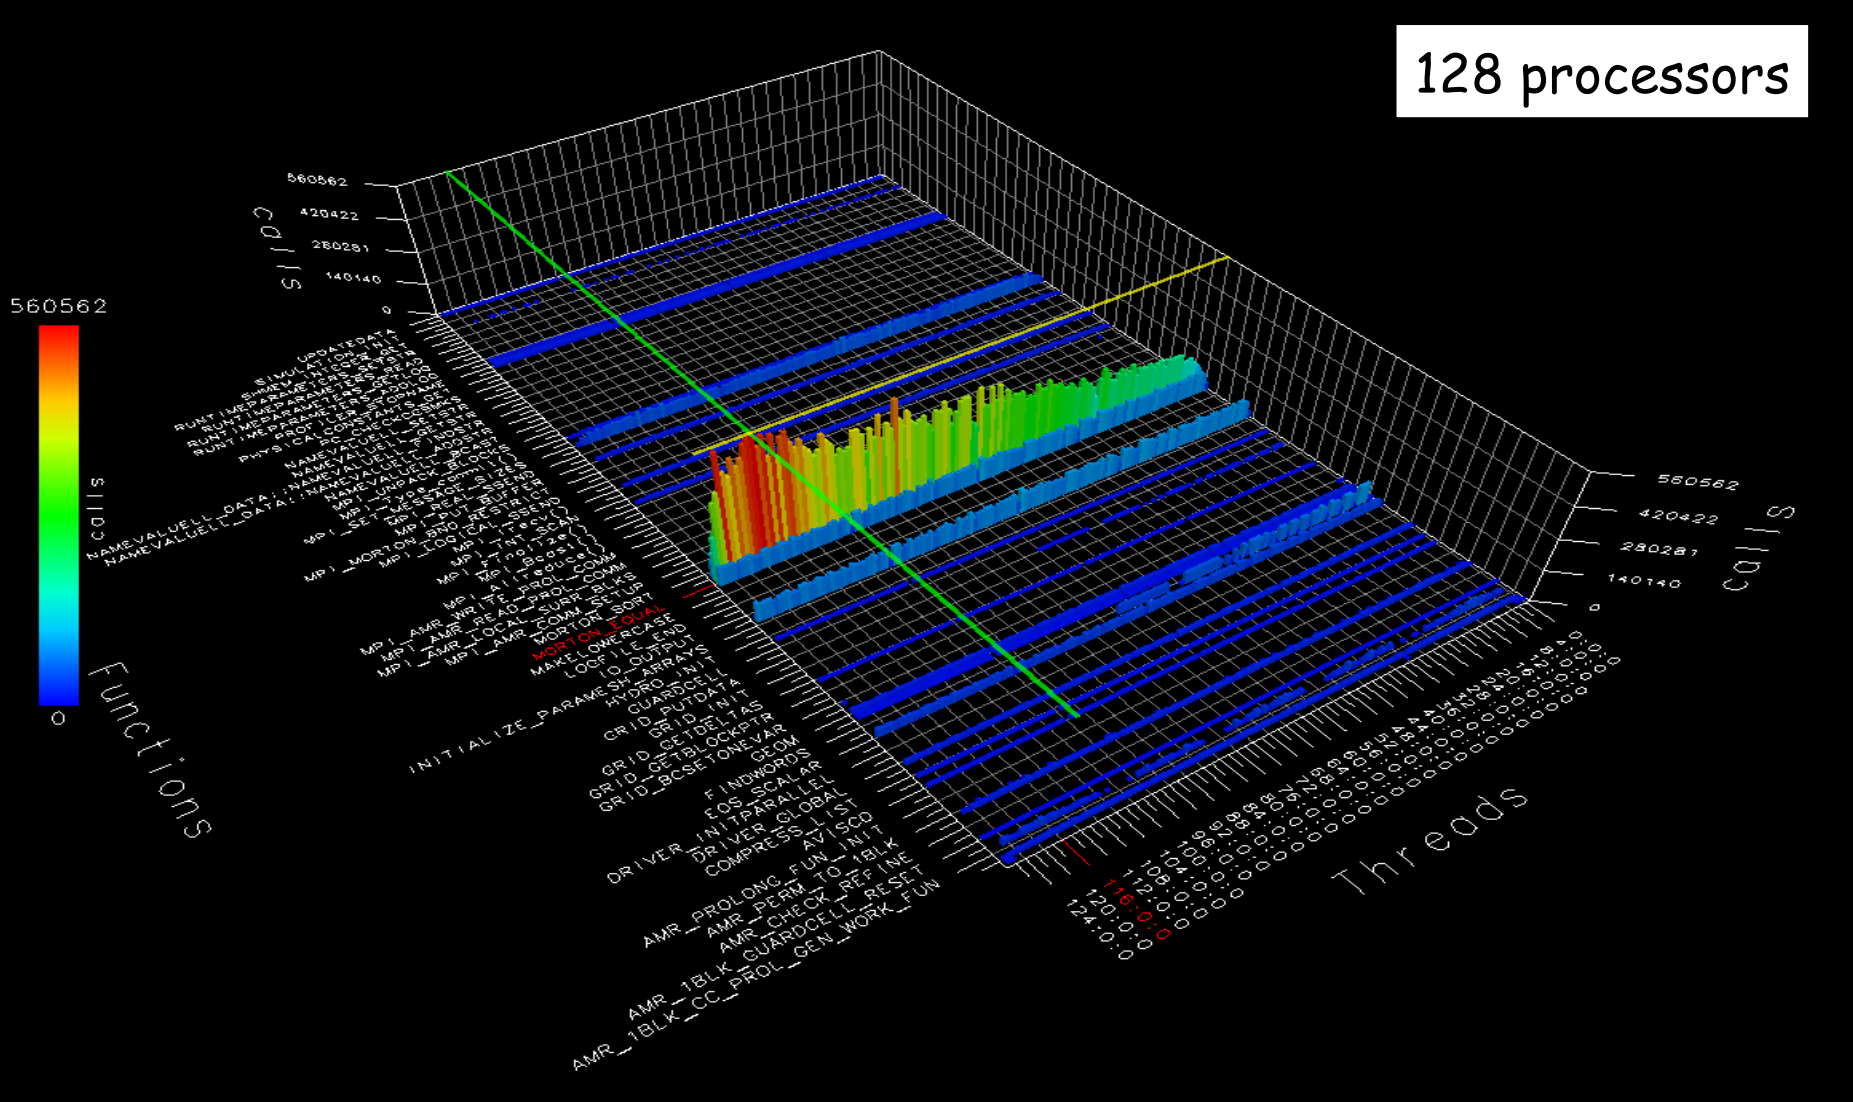


ParaProf – Flat Profile (Miranda)

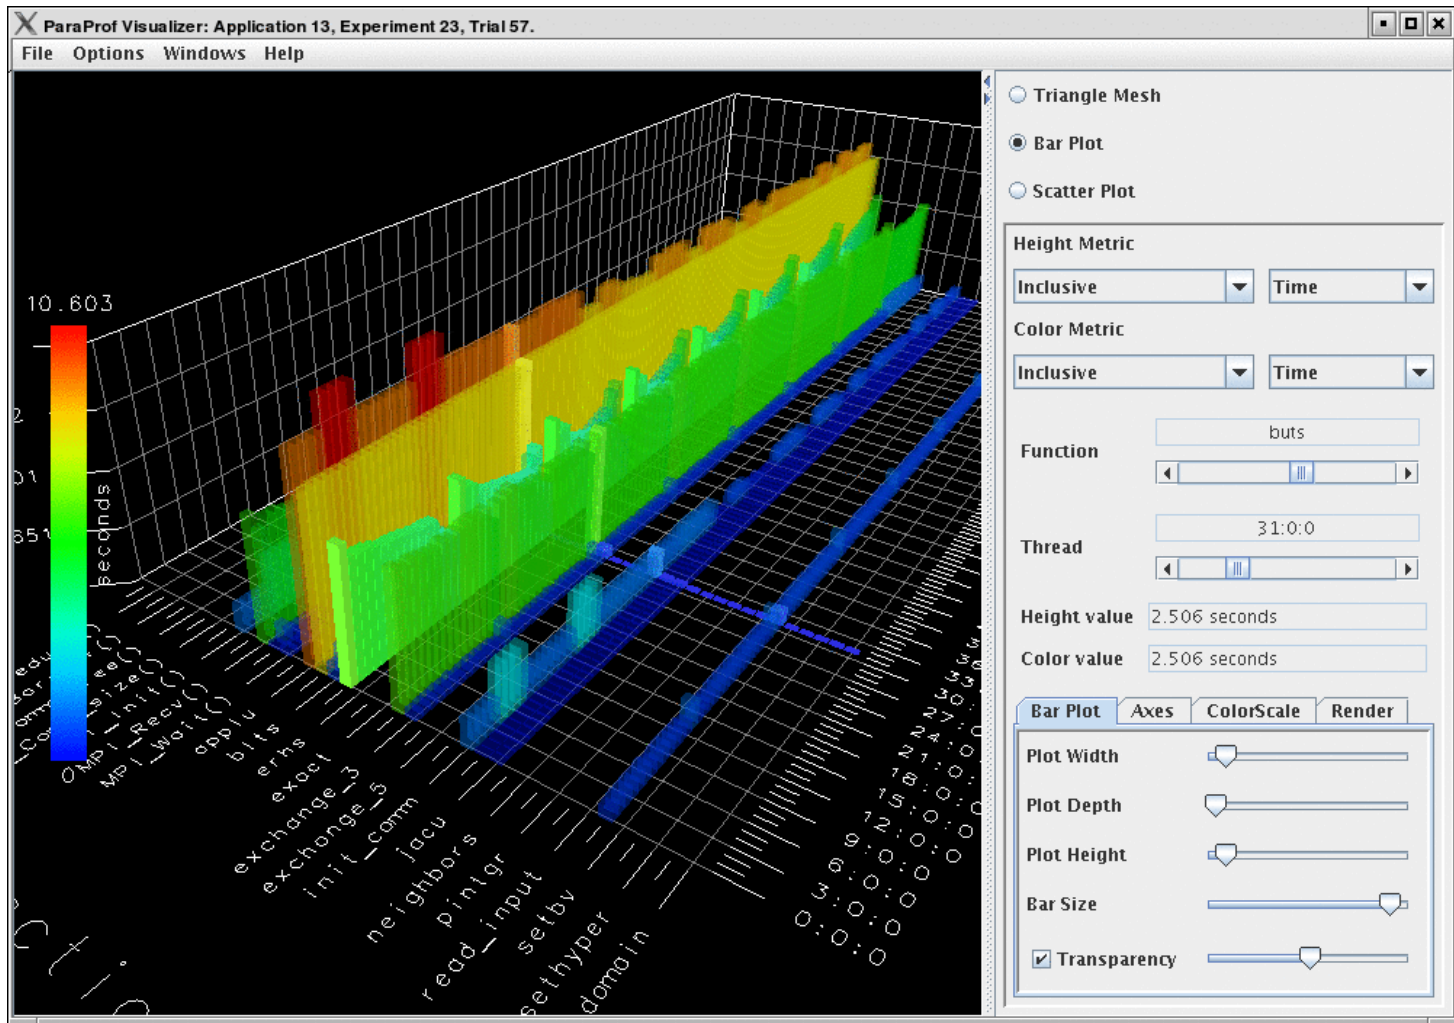


ParaProf – 3D Full Profile Bar Plot (Flash)

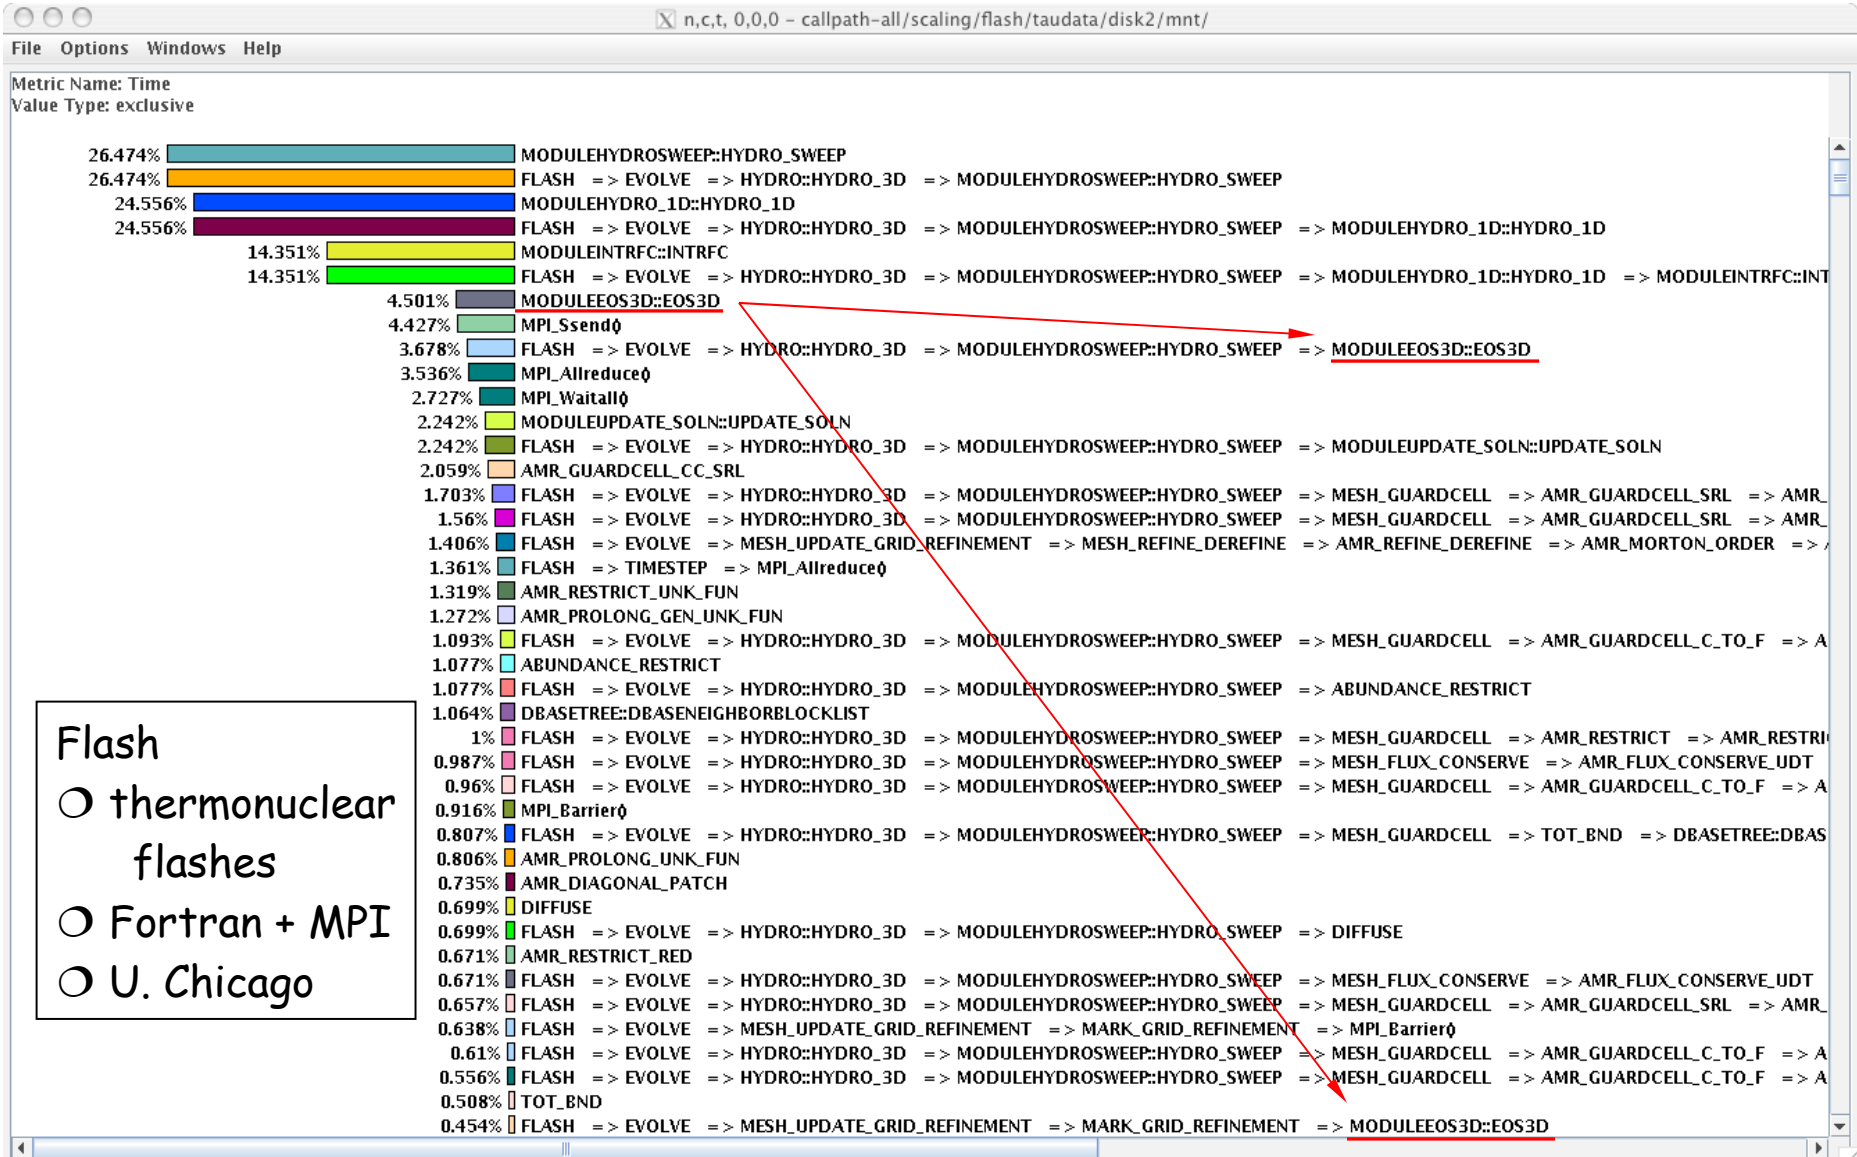
128 processors



ParaProf Bar Plot (Zoom in/out +/-)



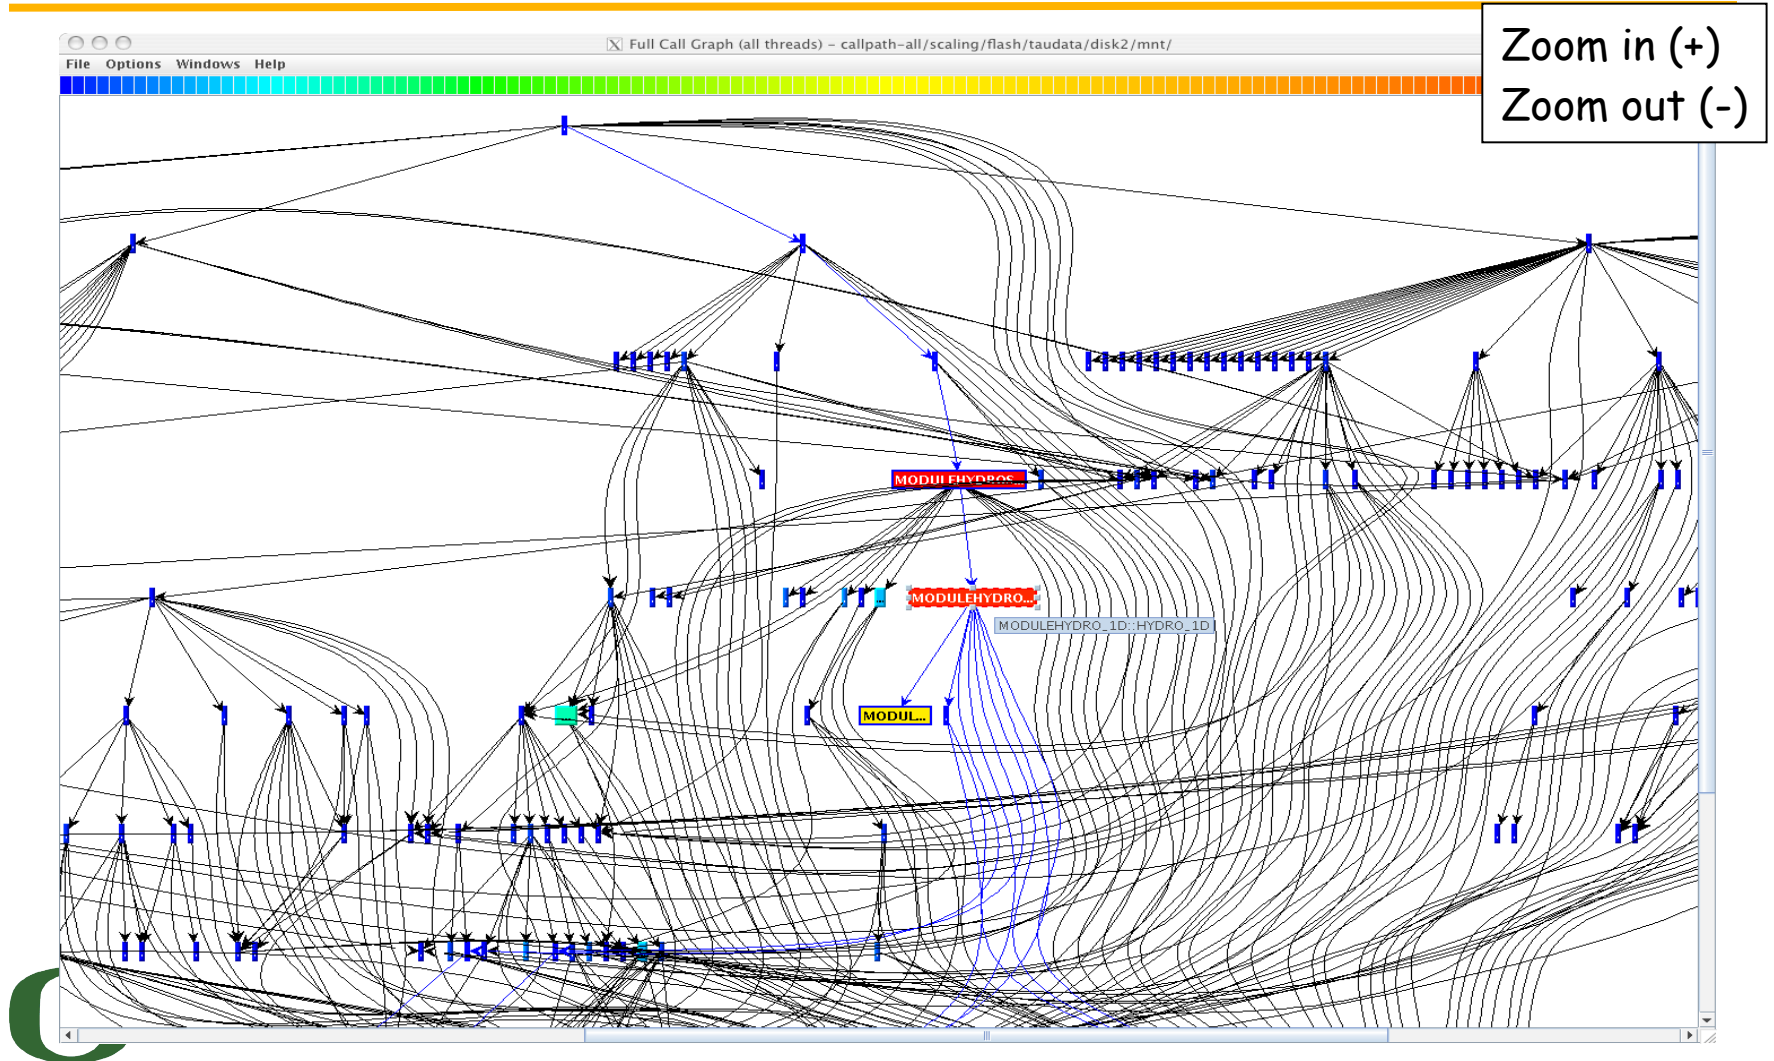
ParaProf – Callpath Profile (Flash)



Flash

- thermonuclear flashes
- Fortran + MPI
- U. Chicago

ParaProf – Callgraph Zoomed (Flash)



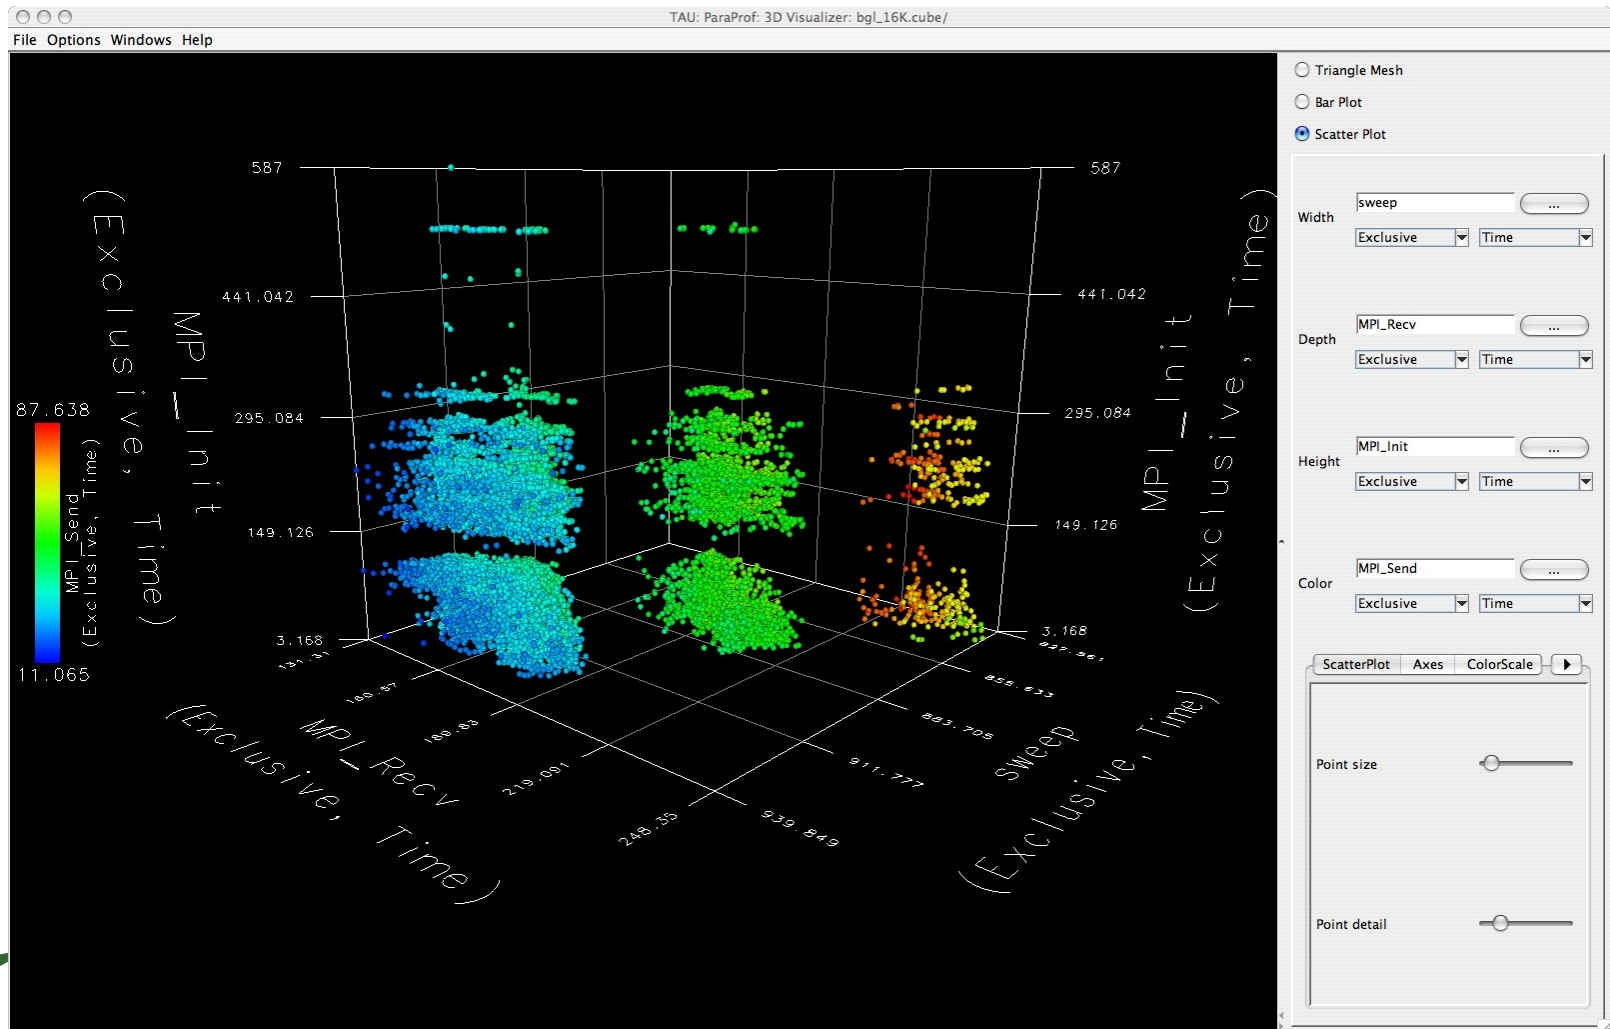
ParaProf - Thread Statistics Table (GSI)

Thread Statistics: n,c,t, 0,0,0 - comp.ppk/


File Options Windows Help

Name	Inclusive Time	Exclusive Time	Calls	Child Calls
▼ GSI	5,223.564	0.098	1	30
SPECMOD::INIT_SPEC_VARS	0.26	0.26	1	0
▶ MPI_Init()	0.056	0.054	1	1
▼ GSI SUB	5,223.094	0.012	1	13
▶ RADINFO::RADINFO_READ	0.103	0.101	1	1,196
PCPINFO::PCPINFO_READ	0.042	0.042	1	0
▼ GLBSOI	5,212.171	0.024	1	12
MPI_Finalize()	1.004	1.004	1	0
▶ OBS_PARA	3.635	0.181	1	56
JFUNC::CREATE_JFUNC	0.142	0.142	1	0
GUESS_GRIDS::CREATE_GES_BIAS_GRIDS	0.059	0.059	1	0
▶ READ_GUESS	1,406.412	0.023	1	8
▼ READ_OBS	3,770.188	0.016	1	6
MPI_Allreduce()	3,725.802	3,725.802	3	0
▶ READ_BUFRTOVS	44.369	0.254	1	871,535
SATTHIN::MAKEGVALS	0	0	1	0
▶ W3FS21	0	0	1	1
BINARY_FILE_UTILITY::OPEN_BINARY_FILE	0.025	0.012	1	3
INITIALIZE::INITIALIZE_RTM	0.099	0.001	1	2
GUESS_GRIDS::CREATE_SFC_GRIDS	0	0	1	0
▶ M_FVANAGRID::ALLGETLIST_	30.582	0	1	10
ERROR_HANDLER::DISPLAY_MESSAGE	0	0	1	0
JFUNC::SET_POINTER	0	0	1	0
OZINFO::OZINFO_READ	0.016	0.016	1	0
DETER_SUBDOMAIN	0.008	0.008	1	0
GRIDMOD::CREATE_MAPPING	0.005	0.005	1	0
INIT_COMMVARS	0.004	0.004	1	0
▶ M_FVANAGRID::ALLGETLIST_	10.711	0	1	1
GRIDMOD::CREATE_GRID_VARS	0	0	1	0

ParaProf – 3D Scatterplot (SWEEP3D CUBE)



Vampir – Trace Zoomed (S3D)

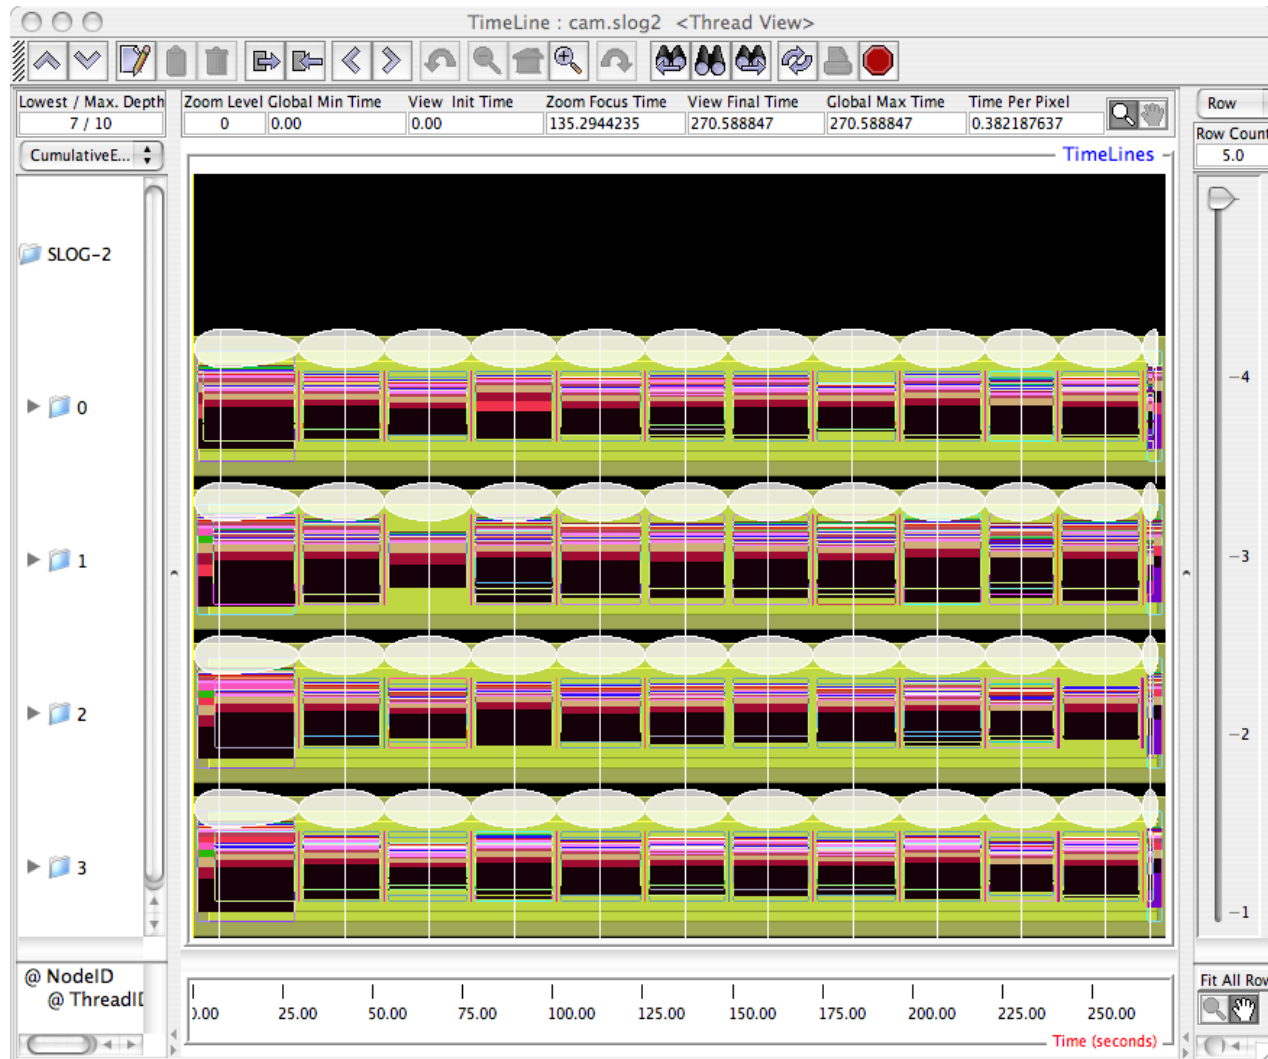


Jumpshot

- <http://www-unix.mcs.anl.gov/perfvis/software/viewers/index.htm>
- Developed at Argonne National Laboratory as part of the MPICH project
 - Also works with other MPI implementations
 - Installed on NAVO IBM and ERDC XT3/4
 - Jumpshot is bundled with the TAU package
- Java-based tracefile visualization tool for postmortem performance analysis of MPI programs
- Latest version is Jumpshot-4 for SLOG-2 format
 - Scalable level of detail support
 - Timeline and histogram views
 - Scrolling and zooming
 - Search/scan facility



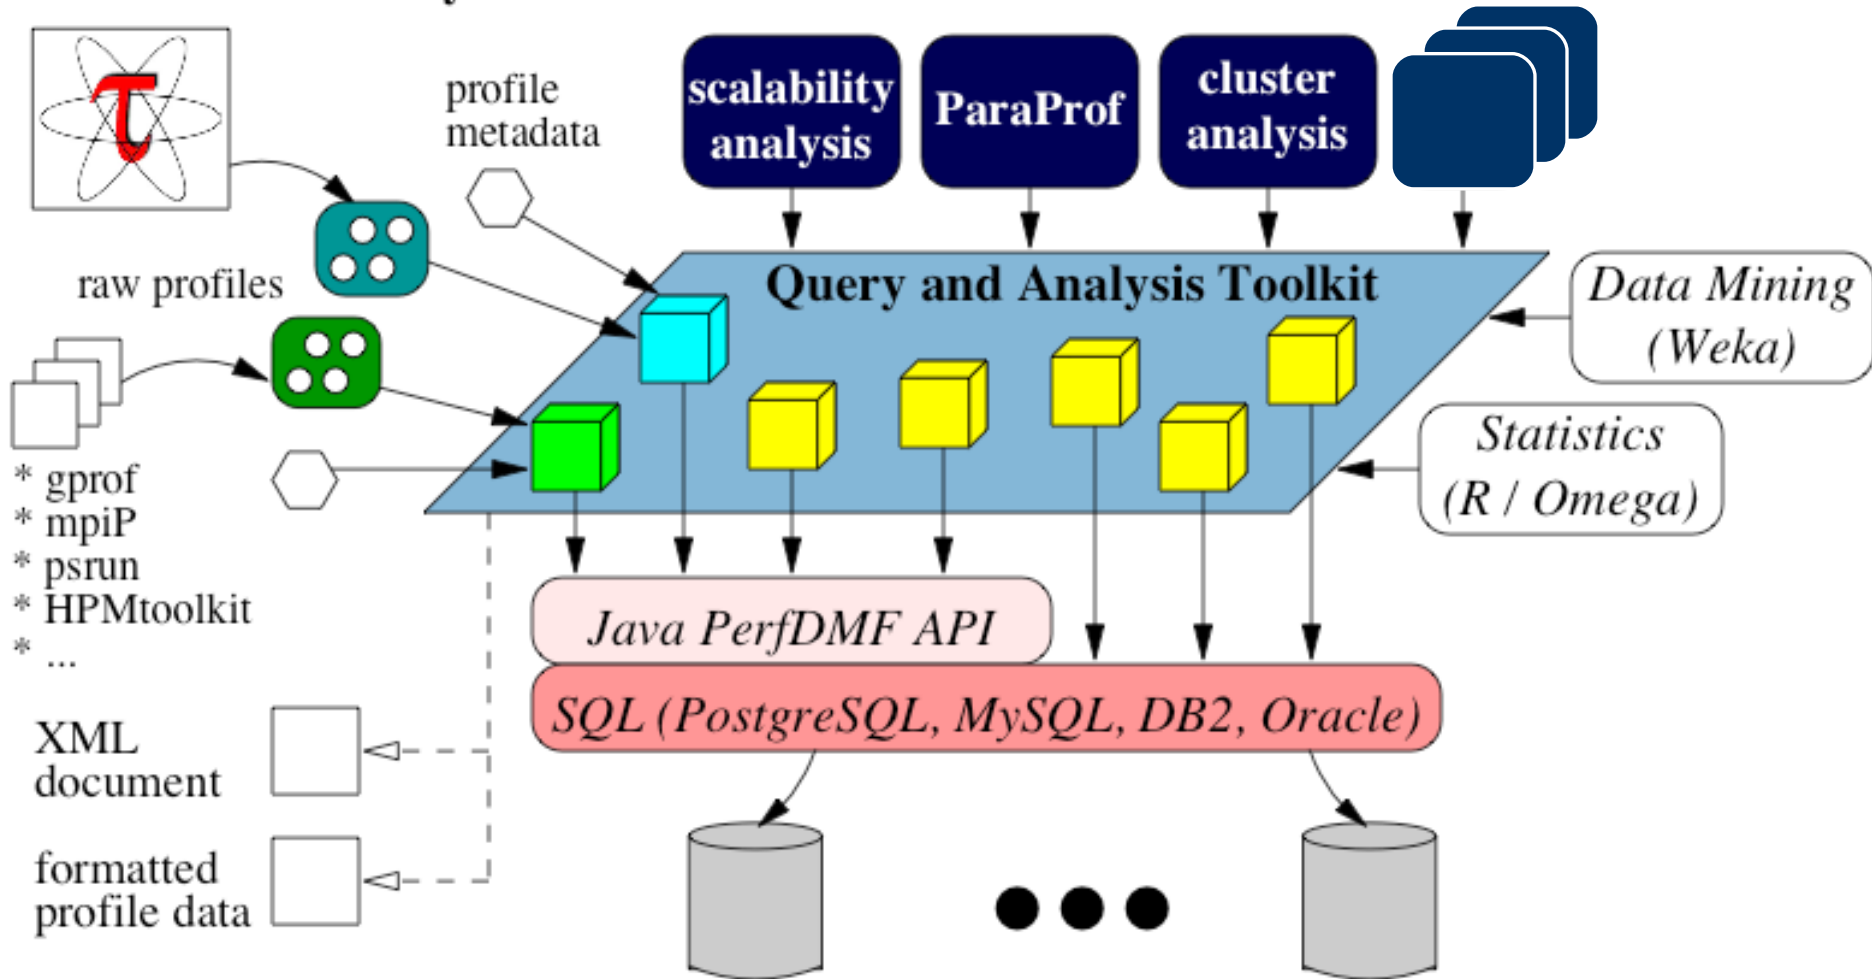
Jumpshot



PerfDMF: Performance Data Mgmt. Framework

TAU Performance System

Performance Analysis Programs



Using Performance Database (PerfDMF)

- **Configure PerfDMF (Done by each user)**
 - % perfdmf_configure --create-default
 - Choose derby, PostgreSQL, MySQL, Oracle or DB2
 - Hostname
 - Username
 - Password
 - Say yes to downloading required drivers (we are not allowed to distribute these)
 - Stores parameters in your ~/.ParaProf/perfdmf.cfg file
- **Configure PerfExplorer (Done by each user)**
 - % perfexplorer_configure
- **Execute PerfExplorer**
 - % perfexplorer



- Development of the TAU portal
 - Common repository for collaborative data sharing
 - Profile uploading, downloading, user management
 - Paraprof, PerfExplorer can be launched from the portal using Java Web Start (no TAU installation required)
- Portal URL
<http://tau.nic.uoregon.edu>



Performance Data Mining (Objectives)

- Conduct parallel performance analysis process
 - In a systematic, collaborative and reusable manner
 - Manage performance complexity
 - Discover performance relationship and properties
 - Automate process
- Multi-experiment performance analysis
- Large-scale performance data reduction
 - Summarize characteristics of large processor runs
- Implement extensible analysis framework
 - Abstraction / automation of data mining operations
 - Interface to existing analysis and data mining tools



Performance Data Mining (PerfExplorer)

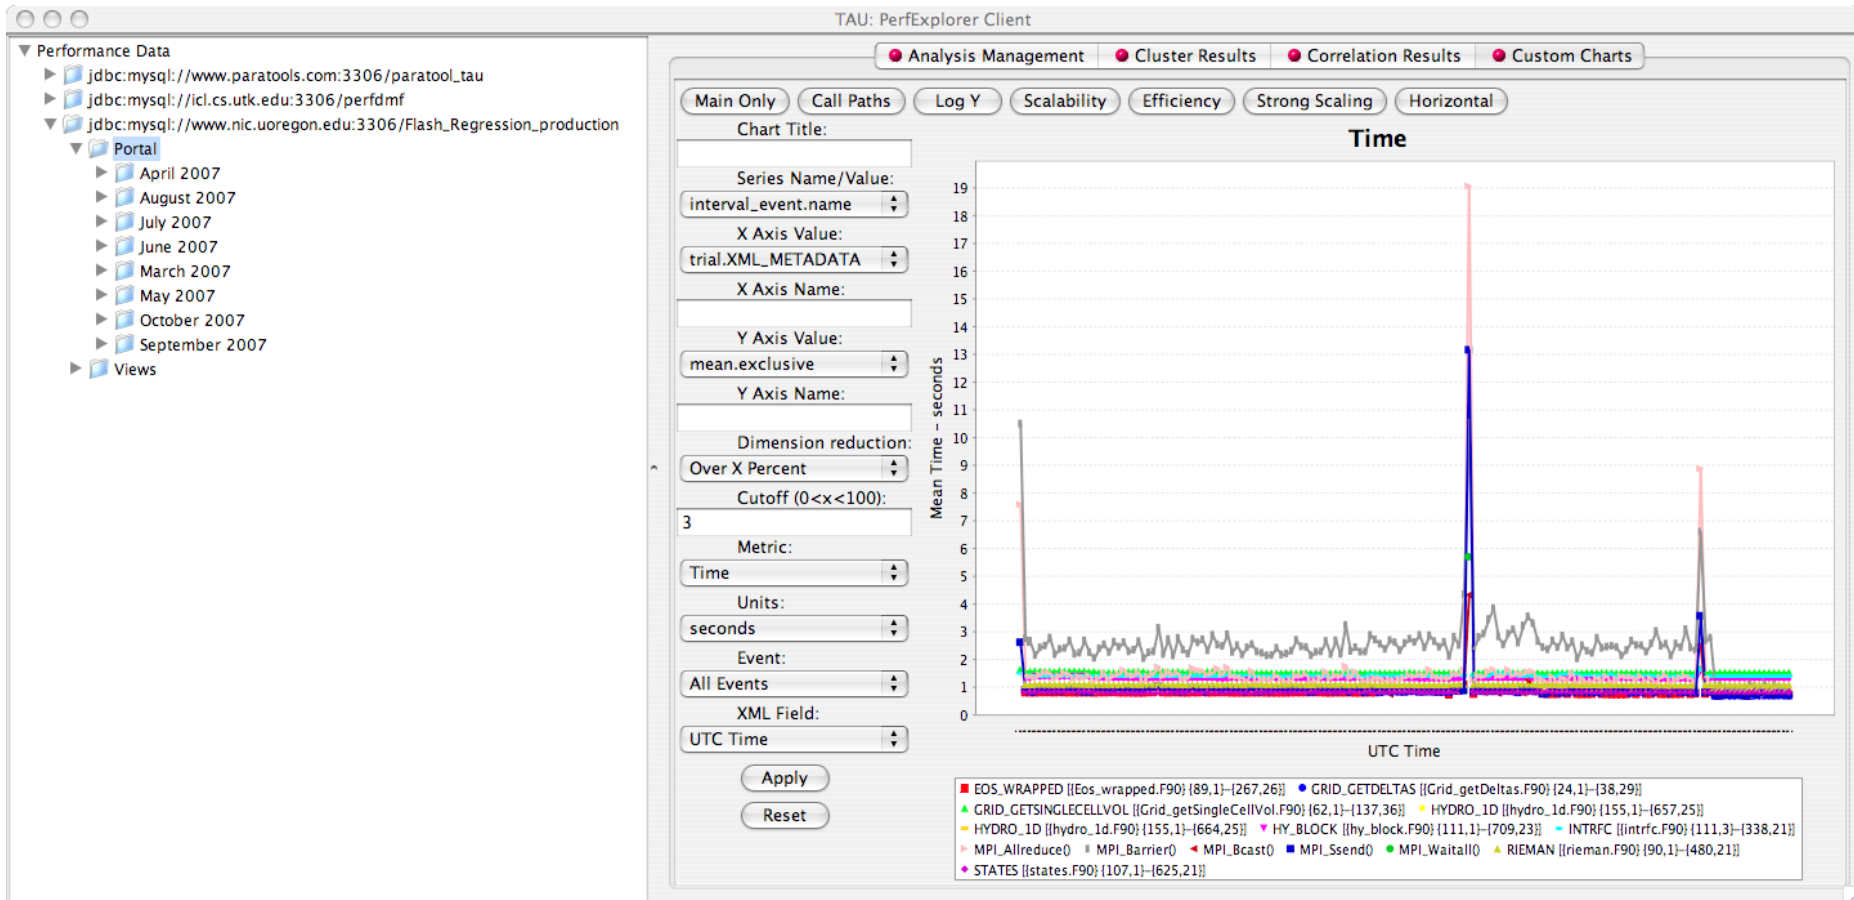
- Performance knowledge discovery framework
 - Data mining analysis applied to parallel performance data
 - comparative, clustering, correlation, dimension reduction, ...
 - Use the existing TAU infrastructure
 - TAU performance profiles, PerfDMF
 - Client-server based system architecture
- Technology integration
 - Java API and toolkit for portability
 - PerfDMF
 - WEKA data mining package
 - JFreeChart for visualization, vector output (EPS, SVG)




PerfExplorer: Regression Testing



PerfExplorer: Exclusive Time for Events (2007)



PerfExplorer: Limiting Events (> 3%), Oct 2007

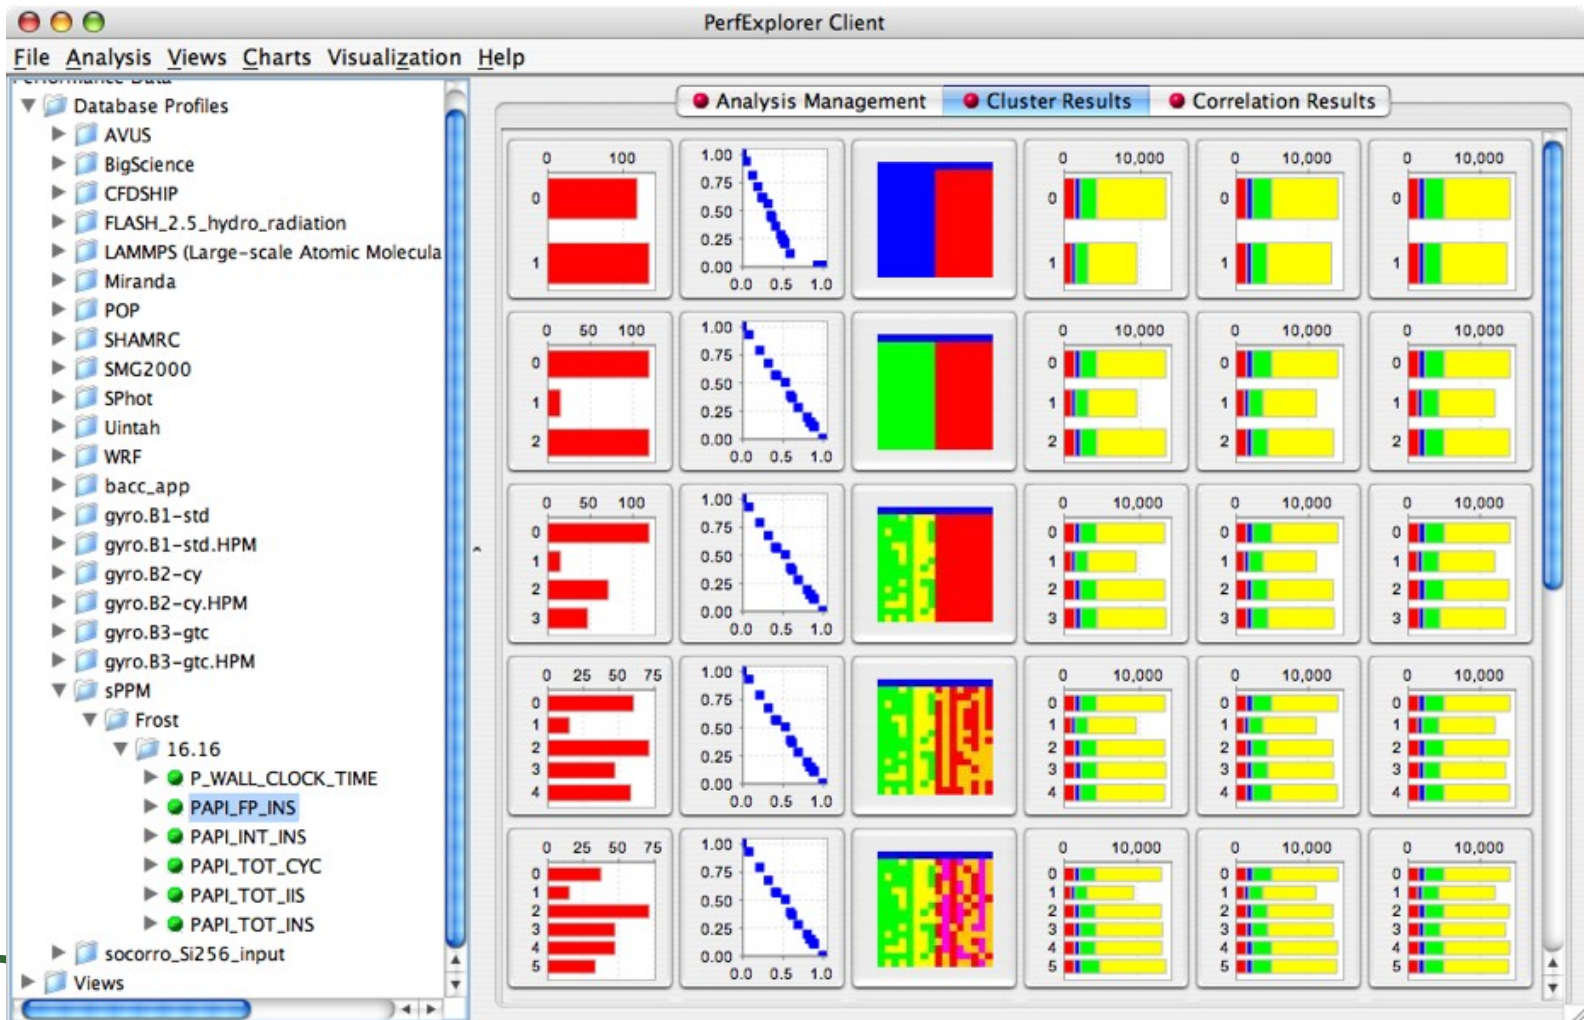


PerfExplorer - Analysis Methods

- Data summaries, distributions, scatter plots
- Clustering
 - *k*-means
 - Hierarchical
- Correlation analysis
- Dimension reduction
 - PCA
 - Random linear projection
 - Thresholds
- Comparative analysis
- Data management views




PerfExplorer - Cluster Analysis (sPPM)



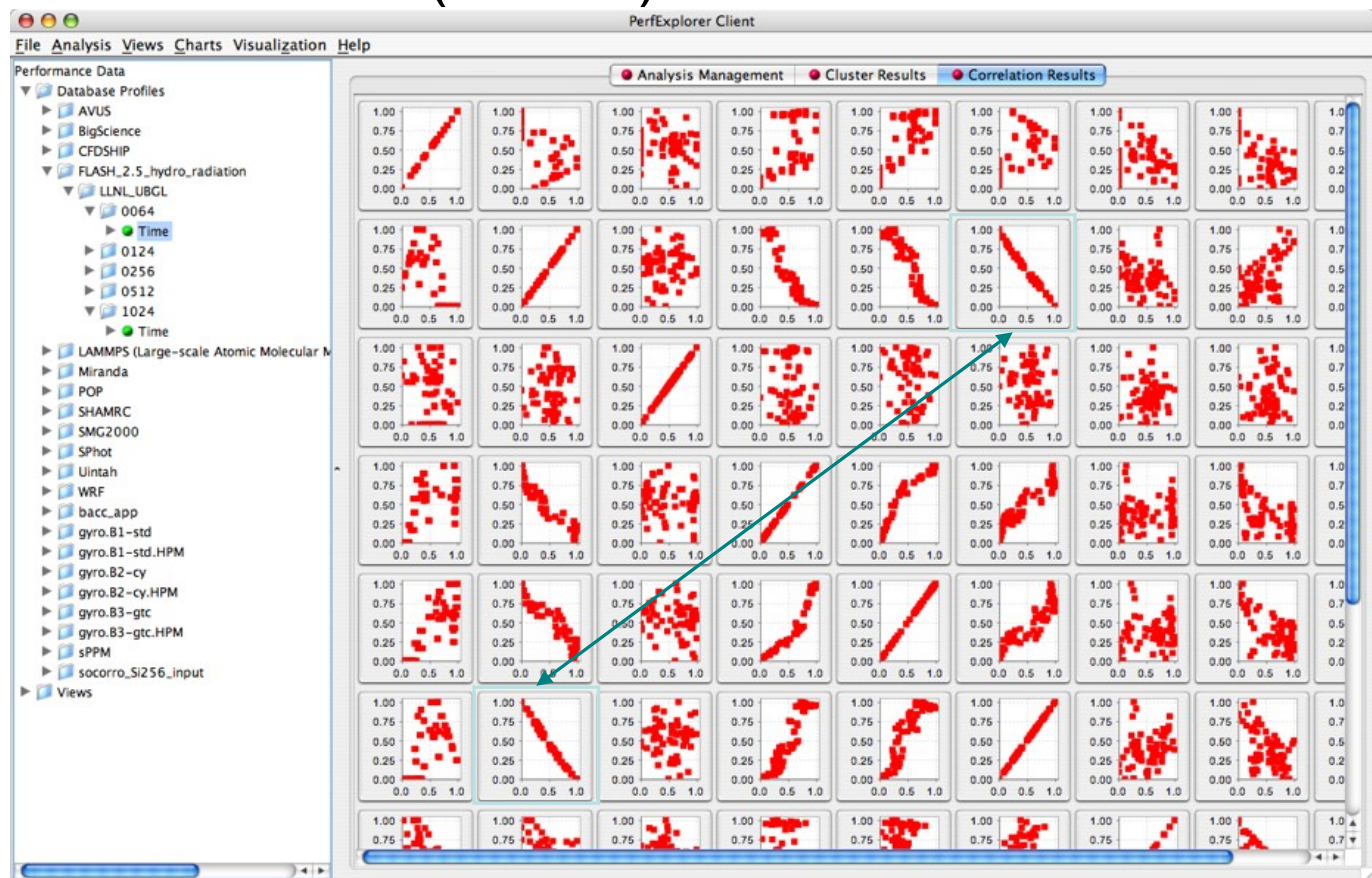
PerfExplorer - Cluster Analysis

- Four significant events automatically selected (from 16K processors)
- Clusters and correlations are visible




PerfExplorer - Correlation Analysis (Flash)

- Describes strength and direction of a linear relationship between two variables (events) in the data



PerfExplorer - Correlation Analysis (Flash)

- -0.995 indicates strong, negative relationship
- As `CALC_CUT_BLOCK_CONTRIBUTIONS()` increases in execution time, `MPI_Barrier()` decreases

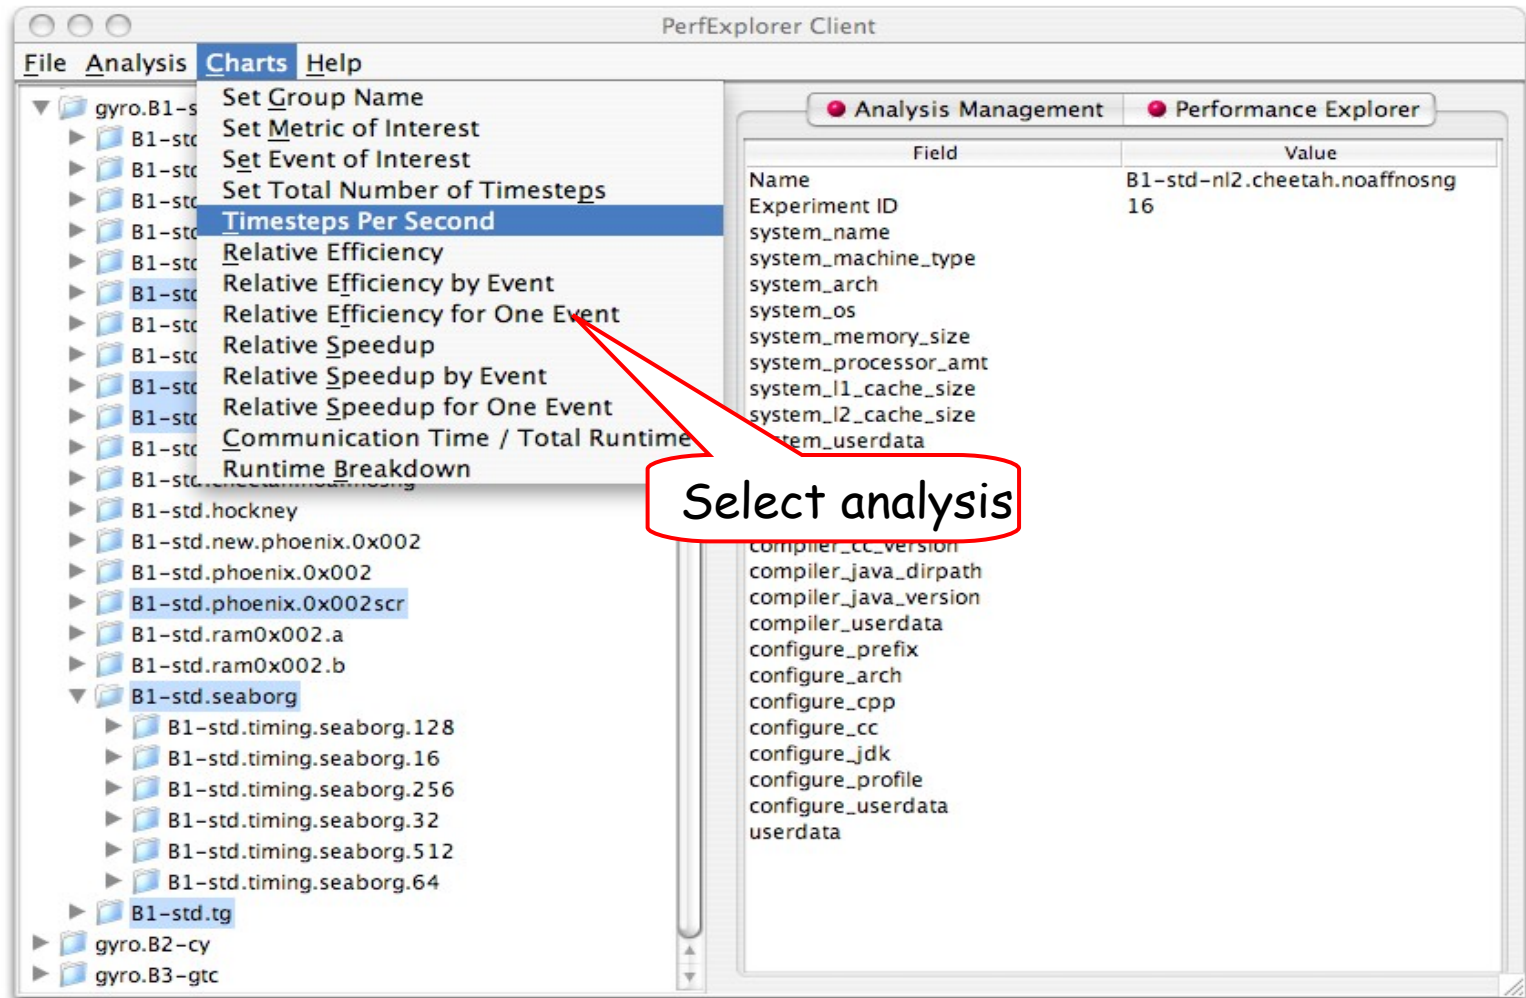


PerfExplorer - Comparative Analysis


- Relative speedup, efficiency
 - total runtime, by event, one event, by phase
- Breakdown of total runtime
- Group fraction of total runtime
- Correlating events to total runtime
- Timesteps per second
- Performance Evaluation Research Center (PERC)
 - PERC tools study (led by ORNL, Pat Worley)
 - In-depth performance analysis of select applications
 - Evaluation performance analysis requirements
 - Test tool functionality and ease of use




PerfExplorer - Interface




PerfExplorer - Relative Efficiency Plots




PerfExplorer - Relative Efficiency by Routine




PerfExplorer - Relative Speedup




PerfExplorer - Timesteps Per Second




PerfExplorer - Relative Efficiency



PerfExplorer - Relative Speedup by Event




PerfExplorer - Runtime Breakdown




PerfExplorer - Timesteps per Second for GYRO

- Cray X1 is the fastest to solution
 - In all 3 tests
- FFT (nl2) improves time
 - B3-gtc only
- TeraGrid faster than p690
 - For B1-std?
- All plots generated automatically


B1-std



B2-cy






B3-gtc

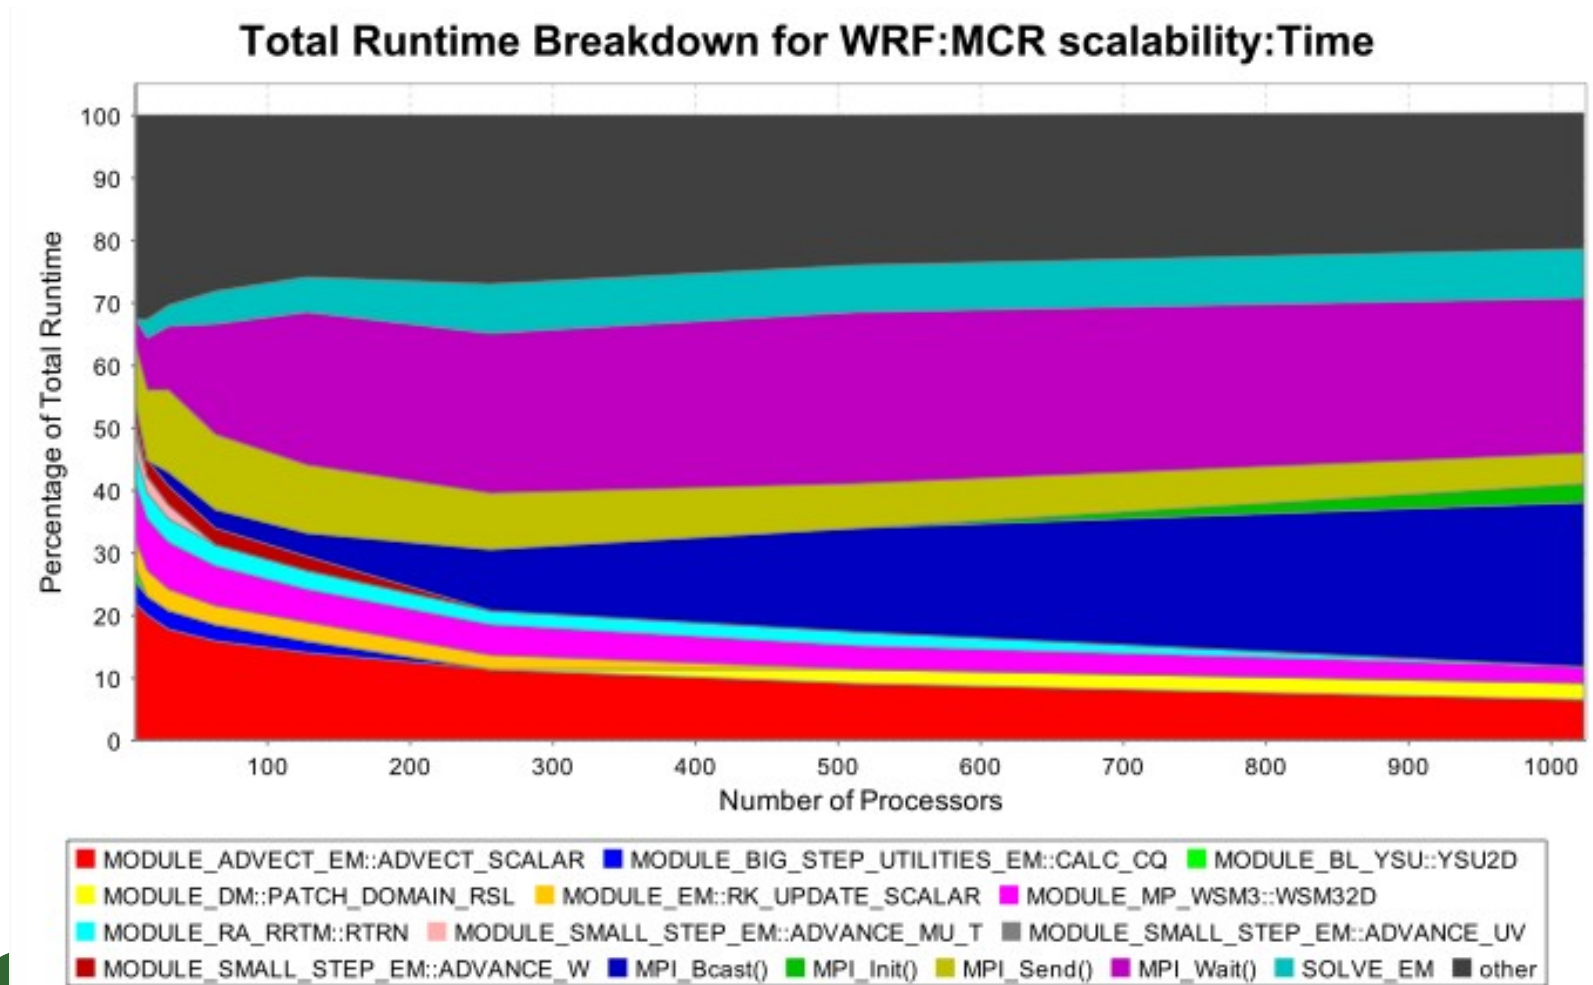


PerfExplorer - Relative Efficiency (B1-std)

- By experiment (B1-std)
 - Total runtime (Cheetah (red))
- By event for one experiment
 - Coll_tr (blue) is significant
- By experiment for one event
 - Shows how Coll_tr behaves for all experiments

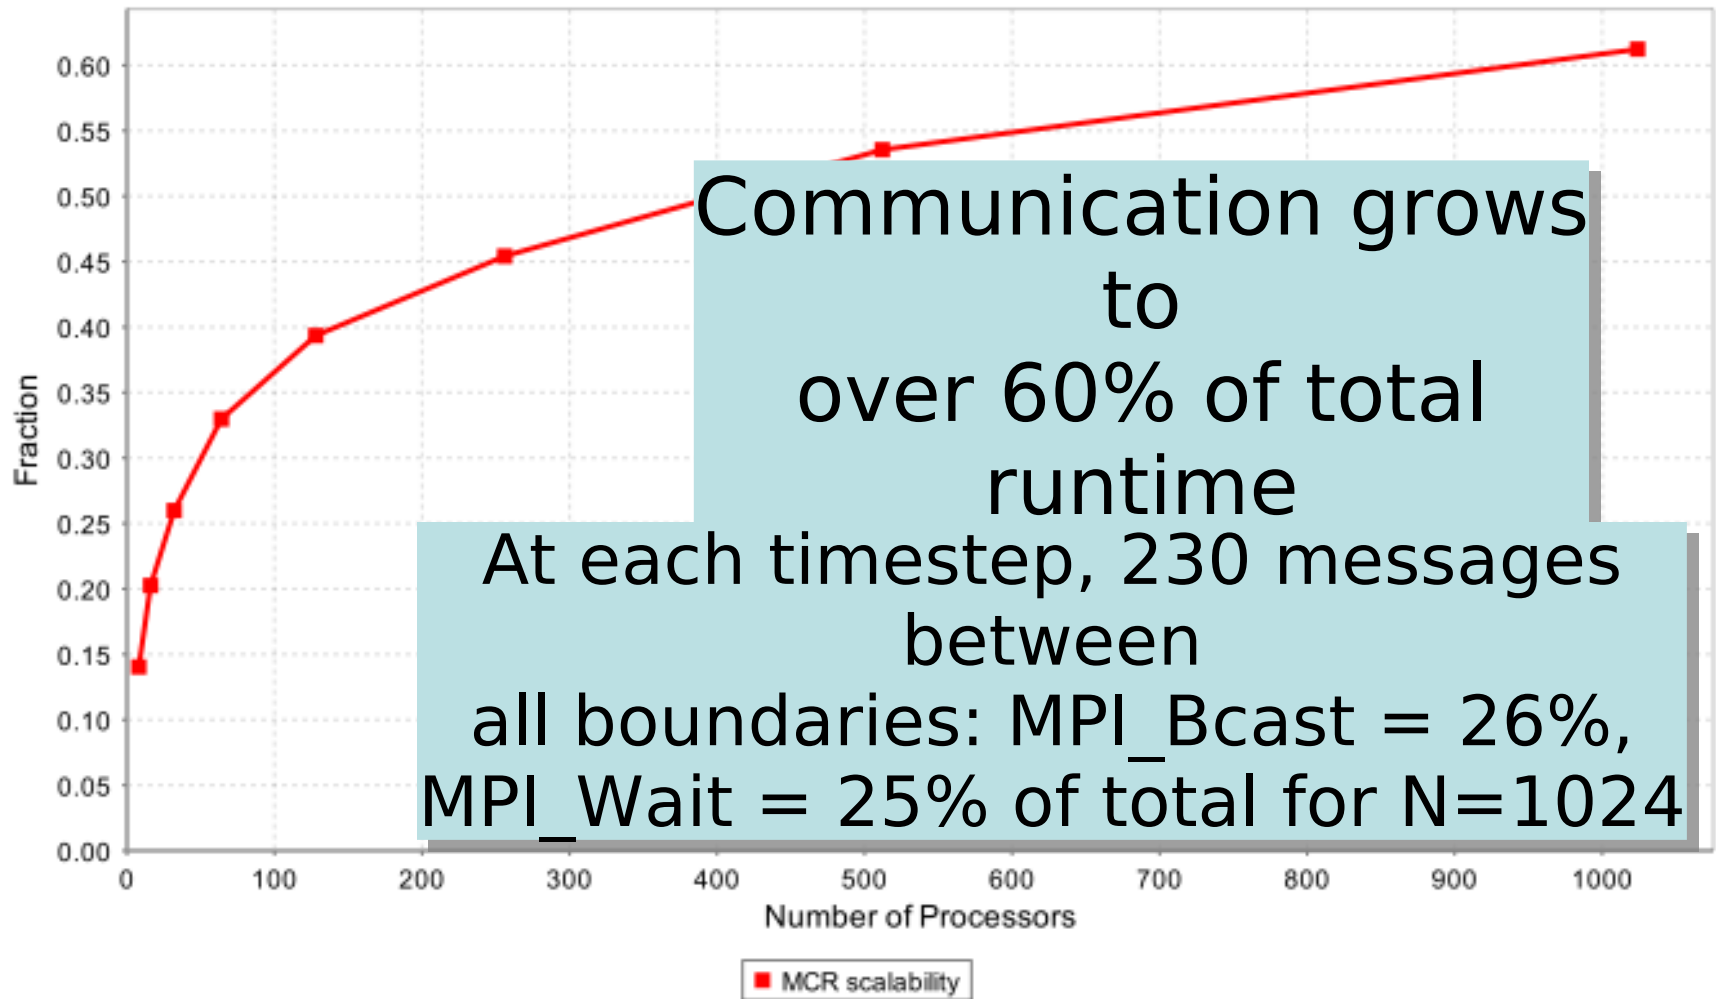


PerfExplorer - Runtime Breakdown



Group % of Total

MPI Time / Total Runtime - WRF:MCR scalability:Time



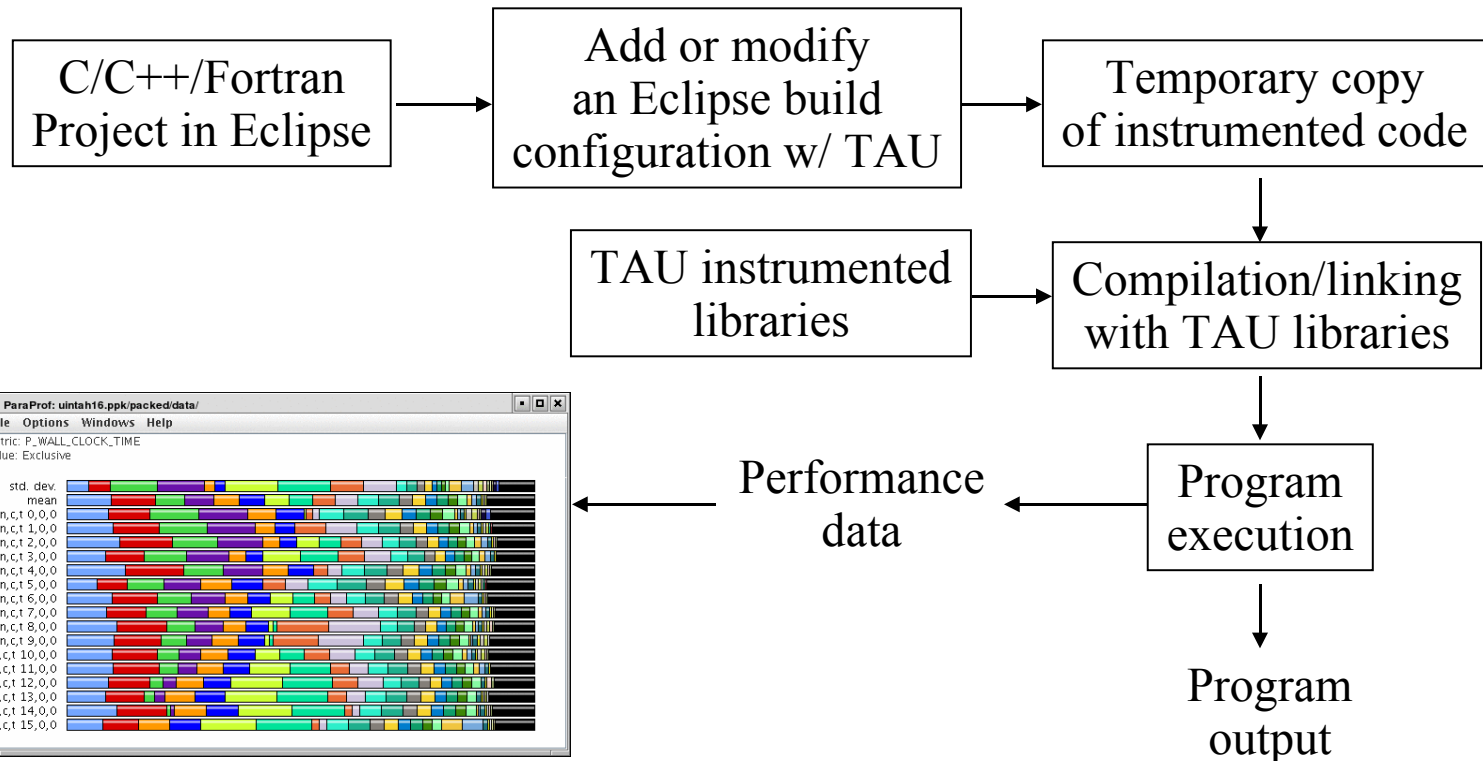
TAU Integration with IDEs

- High performance software development environments
 - Tools may be complicated to use
 - Interfaces and mechanisms differ between platforms / OS
- Integrated development environments
 - Consistent development environment
 - Numerous enhancements to development process
 - Standard in industrial software development
- Integrated performance analysis
 - Tools limited to single platform or programming language
 - Rarely compatible with 3rd party analysis tools
 - Little or no support for parallel projects



TAU and Eclipse

- Provide an interface for configuring TAU's automatic instrumentation within Eclipse's build system
- Manage runtime configuration settings and environment variables for execution of TAU instrumented programs



TAU and Eclipse

The screenshot displays the Eclipse IDE interface for a Fortran project named 'matmult.f90'. The main editor shows the following code:

```
matmult.f90 - simple matrix multiply implementation
subroutine initialize(a, b, n)
  double precision a(n,n)
  double precision b(n,n)
  integer n

! first initialize the A matrix
do i = 1, n
  do j = 1, n
    a(j,i) = i
  end do
end do

! then initialize the B matrix
do i = 1, n
  do j = 1, n
    b(j,i) = i
  end do
end do

end subroutine initialize

subroutine multiply_matrices(answer, buffer, b, matsize)
  double precision buffer(matsize), answer(matsize)
  double precision b(matsize, matsize)
  integer i, j

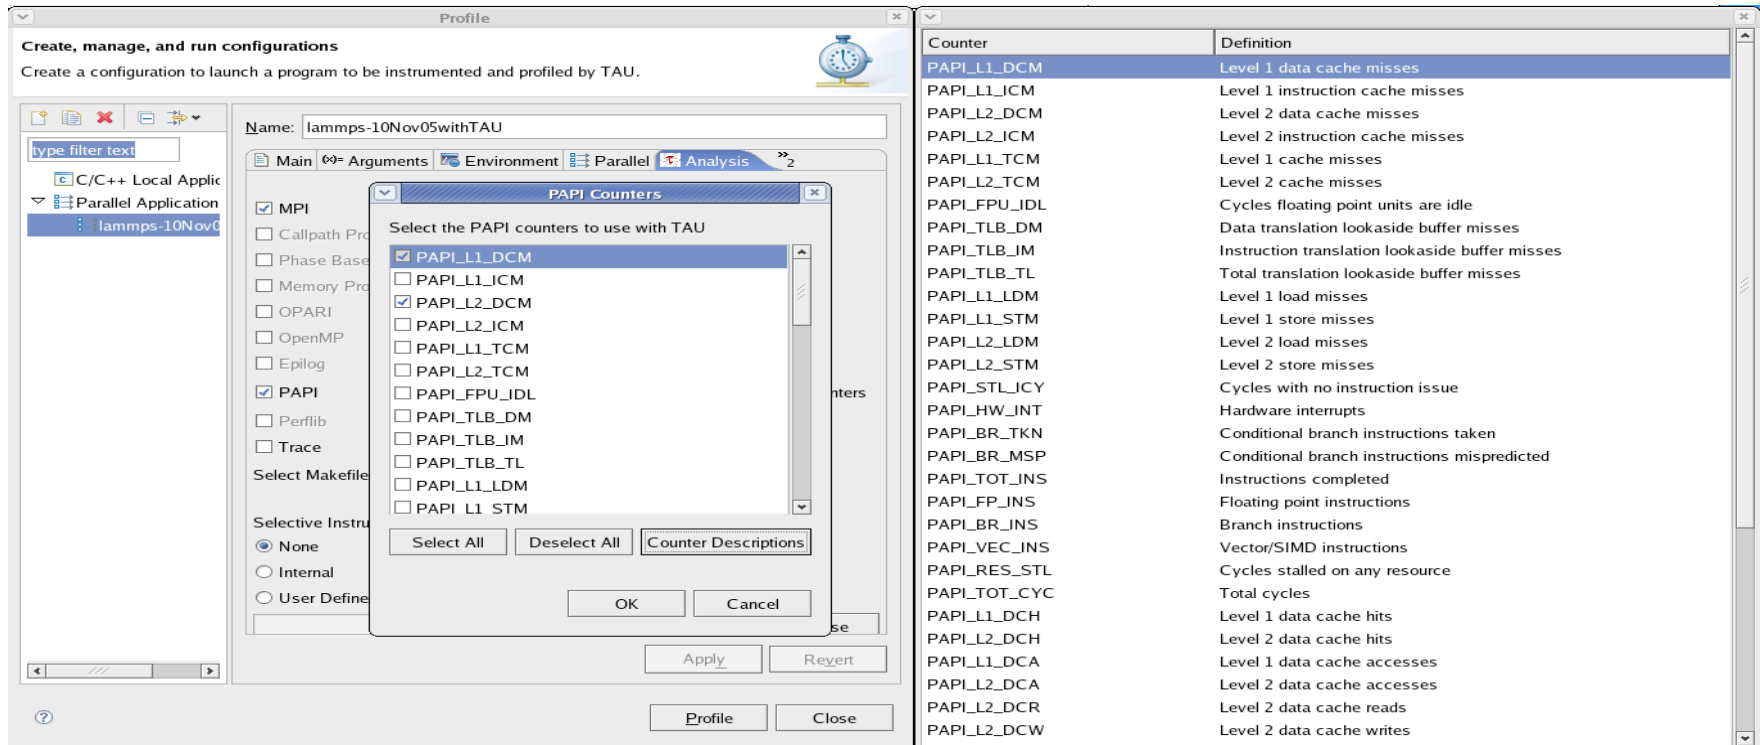
! multiply the row with the column
```

The Outline view on the right lists the subroutines: initialize, multiply_matrices, and main. The Console view at the bottom shows performance data from the Performance Data Manager, with a highlighted entry: 'The New Trial: 2006-12-02 20:36:59'. An arrow points from the text 'PerfDMF' to this entry.

PerfDMF



Choosing PAPI Counters with TAU in Eclipse



The screenshot shows the Eclipse IDE's 'Profile' window. The 'Name' field is set to 'lammps-10Nov05withTAU'. The 'Analysis' tab is active, and the 'PAPI' checkbox is checked. A dialog box titled 'PAPI Counters' is open, allowing the user to select counters to use with TAU. The dialog lists various counters, with 'PAPI_L1_DCM' and 'PAPI_L2_DCM' selected. Below the dialog, a table lists the definitions for these and other counters.

Counter	Definition
PAPI_L1_DCM	Level 1 data cache misses
PAPI_L1_ICM	Level 1 instruction cache misses
PAPI_L2_DCM	Level 2 data cache misses
PAPI_L2_ICM	Level 2 instruction cache misses
PAPI_L1_TCM	Level 1 cache misses
PAPI_L2_TCM	Level 2 cache misses
PAPI_FPU_IDL	Cycles floating point units are idle
PAPI_TLB_DM	Data translation lookaside buffer misses
PAPI_TLB_IM	Instruction translation lookaside buffer misses
PAPI_TLB_TL	Total translation lookaside buffer misses
PAPI_L1_LDM	Level 1 load misses
PAPI_L1_STM	Level 1 store misses
PAPI_L2_LDM	Level 2 load misses
PAPI_L2_STM	Level 2 store misses
PAPI_STL_ICY	Cycles with no instruction issue
PAPI_HW_INT	Hardware interrupts
PAPI_BR_TKN	Conditional branch instructions taken
PAPI_BR_MSP	Conditional branch instructions mispredicted
PAPI_TOT_INS	Instructions completed
PAPI_FP_INS	Floating point instructions
PAPI_BR_INS	Branch instructions
PAPI_VEC_INS	Vector/SIMD instructions
PAPI_RES_STL	Cycles stalled on any resource
PAPI_TOT_CYC	Total cycles
PAPI_L1_DCH	Level 1 data cache hits
PAPI_L2_DCH	Level 2 data cache hits
PAPI_L1_DCA	Level 1 data cache accesses
PAPI_L2_DCA	Level 2 data cache accesses
PAPI_L2_DCR	Level 2 data cache reads
PAPI_L2_DCW	Level 2 data cache writes

% /projects/tau/eclipse/eclipse



TAU Performance System Status

- Computing platforms (selected)
 - IBM SP/pSeries/BGL/Cell PPE, SGI Altix/Origin, Cray T3E/SV-1/X1/XT3, HP (Compaq) SC (Tru64), Sun, Linux clusters (IA-32/64, Alpha, PPC, PA-RISC, Power, Opteron), Apple (G4/5, OS X), Hitachi SR8000, NEC SX Series, Windows ...
- Programming languages
 - C, C++, Fortran 77/90/95, HPF, Java, Python
- Thread libraries (selected)
 - pthreads, OpenMP, SGI sproc, Java, Windows, Charm++
- Compilers (selected)
 - Intel, PGI, GNU, Fujitsu, Sun, PathScale, SGI, Cray, IBM, HP, NEC, Absoft, Lahey, Nagware, ...



More Information

- PAPI References:
 - PAPI documentation page available from the PAPI website:
<http://icl.cs.utk.edu/papi/>
- TAU References:
 - TAU Users Guide and papers available from the TAU website: <http://tau.uoregon.edu/>
- VAMPIR References
 - VAMPIR-NG website
<http://www.vampir-ng.de/>
- Scalasca/KOJAK References
 - Scalasca documentation page
<http://www.scalasca.org/>
- Eclipse PTP References
 - Documentation available from the Eclipse PTP website:
<http://www.eclipse.org/ptp/>



Acknowledgements

- HPCMP DoD PET Program
- Department of Energy
 - Office of Science
 - NNSA/ASC Trilabs (SNL, LLNL, LANL)
- National Science Foundation
- University of Tennessee
 - David Cronk, Shirley Moore
 - Daniel Terpstra
 - Joseph Thomas
- University of Oregon
 - Allen D. Malony, A. Morris, K. Huck, W. Spear
- TU Dresden
 - Holger Brunst, Andreas Knupfer
 - Wolfgang Nagel
- Research Centre Juelich, Germany
 - Bernd Mohr
 - Felix Wolf



UNIVERSITY
OF OREGON

