


VI-HPS



Profile Analysis with ParaProf

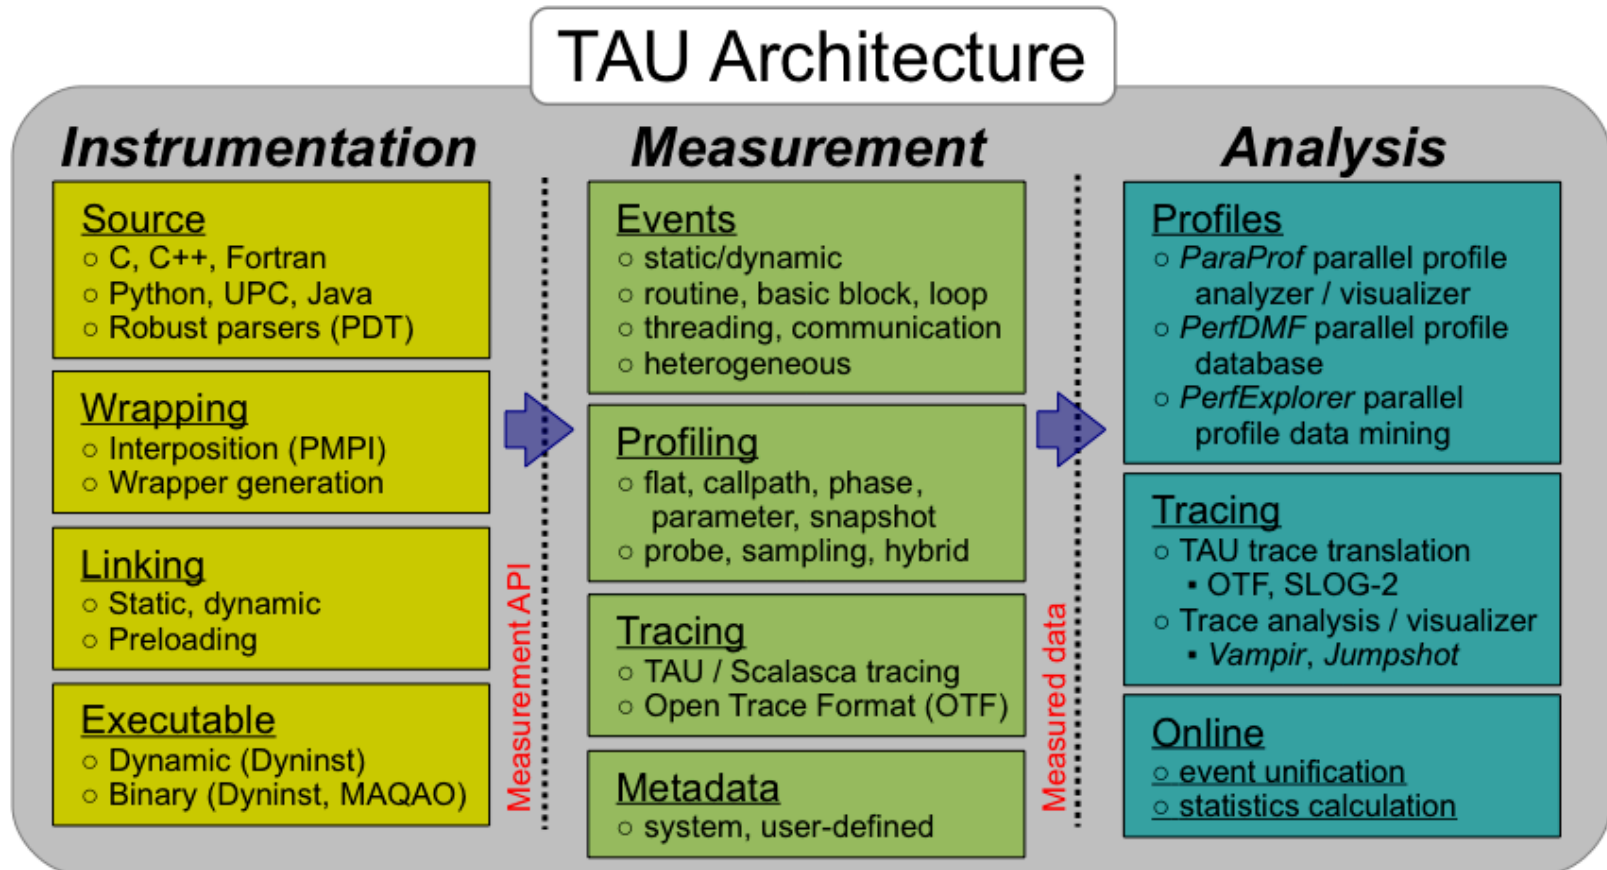




Sameer Shende

Performance Research Lab, University of Oregon

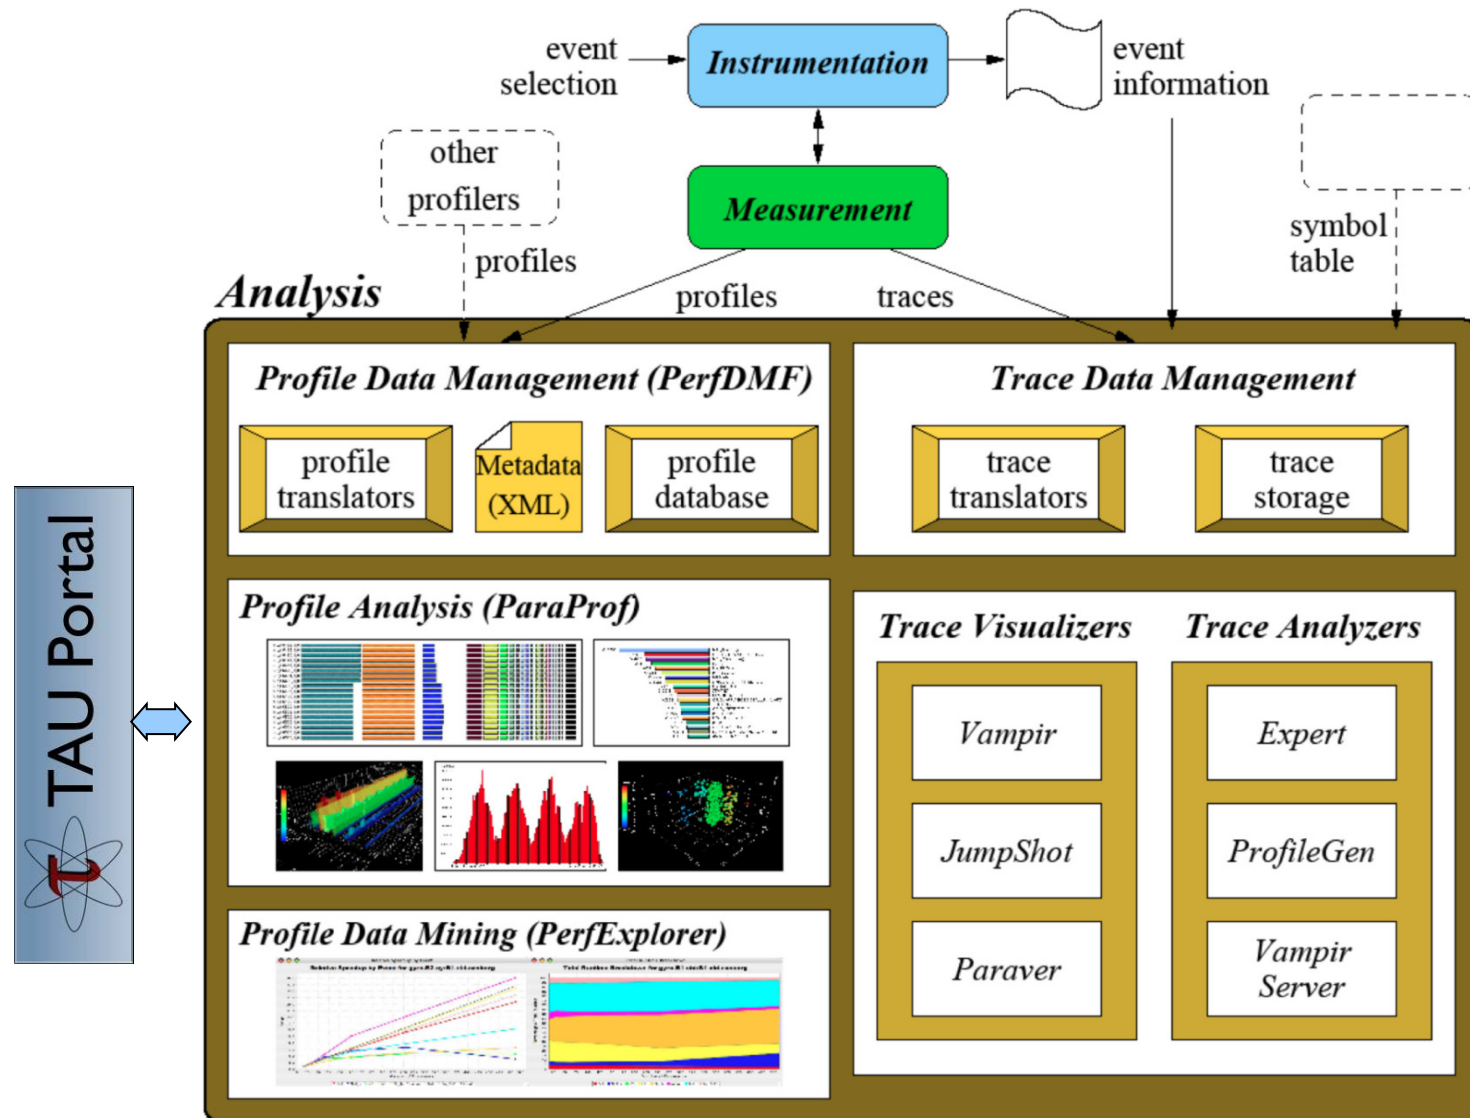
<http://TAU.uoregon.edu>

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



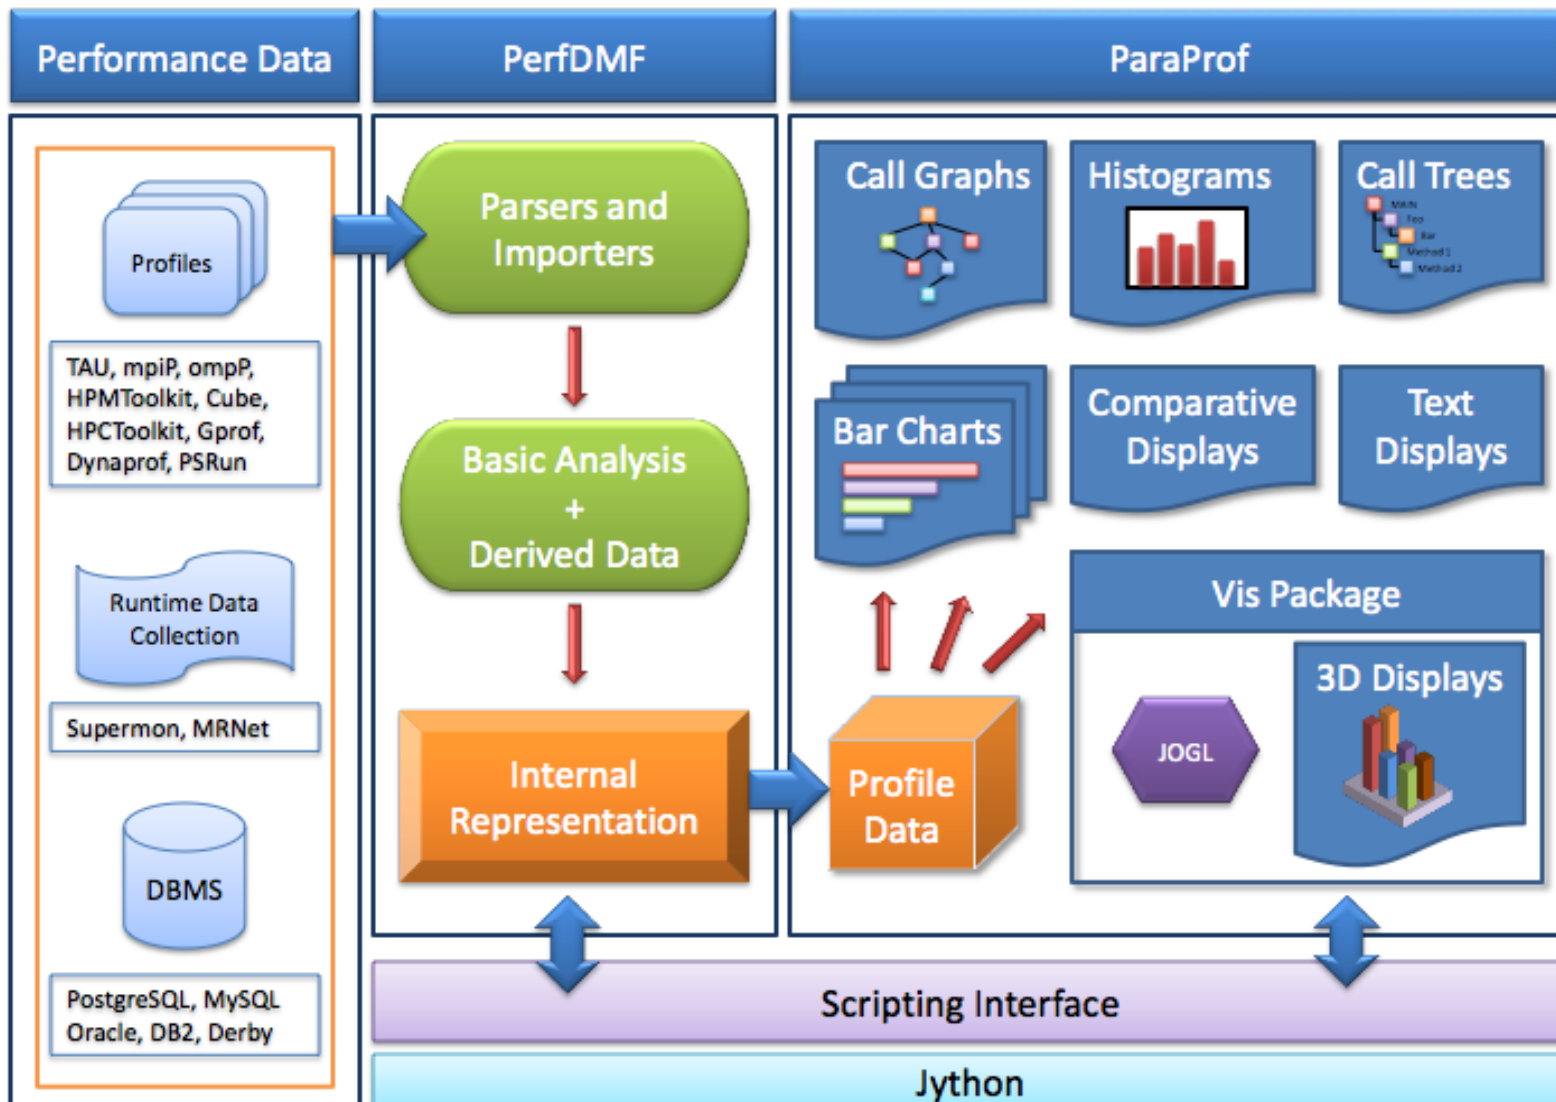


- Instrumentation
 - Fortran, C++, C, UPC, Java, Python, Chapel
 - Automatic instrumentation
- Measurement and analysis support
 - MPI, OpenSHMEM, ARMCI, PGAS, DMAPP
 - pthreads, OpenMP, hybrid, other thread models
 - GPU, CUDA, OpenCL, OpenACC
 - Parallel profiling and tracing
 - Use of Score-P for native OTF2 and CUBEX generation
 - Efficient callpath profiles and trace generation using Score-P
- Analysis
 - Parallel profile analysis (ParaProf), data mining (PerfExplorer)
 - Performance database technology (PerfDMF, TAUdb)
 - 3D profile browser




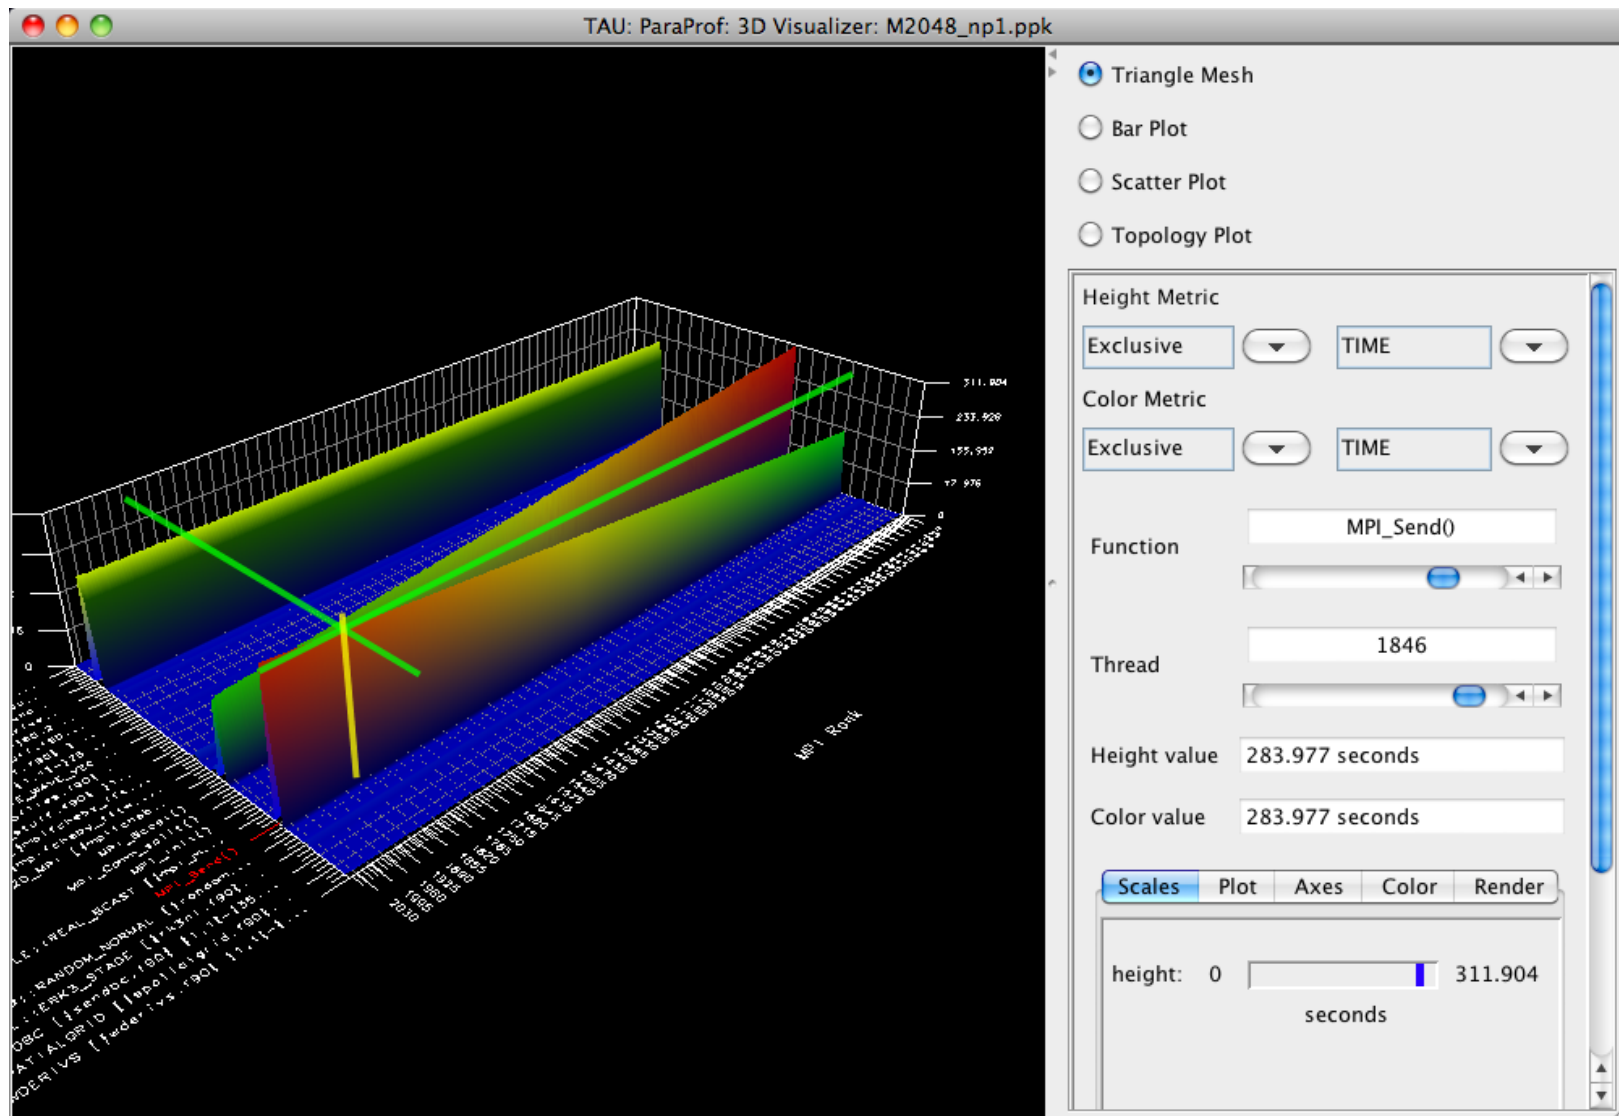
ParaProf Profile Analysis Framework

VI-HPS




VI-HPS





ParaProf: 3D Communication Matrix



- The Live-DVD contains Score-P experiments of BT-MZ
 - class “B”, 4 processes with 4 OpenMP threads each
 - collected on a dedicated node of the SuperMUC HPC system at Leibniz Rechenzentrum (LRZ), Munich, Germany

```
% cd
% cd workshop-vihps/supermuc_expts
% ls
periscope-1.5
README
run.out
scorep-20120913_1740_557443655223384
scorep_bt-mz_B_4x4_sum
scorep_bt-mz_B_4x4_sum+mets
scorep_bt-mz_B_4x4_trace
```

- Start TAU’s paraprof GUI with default profile report

```
% paraprof scorep-20120913_1740_557443655223384/profile.cubex
OR
% paraprof scorep_bt-mz_B_4x4_trace/scout.cubex
```


TAU: ParaProf Manager


File Options Help

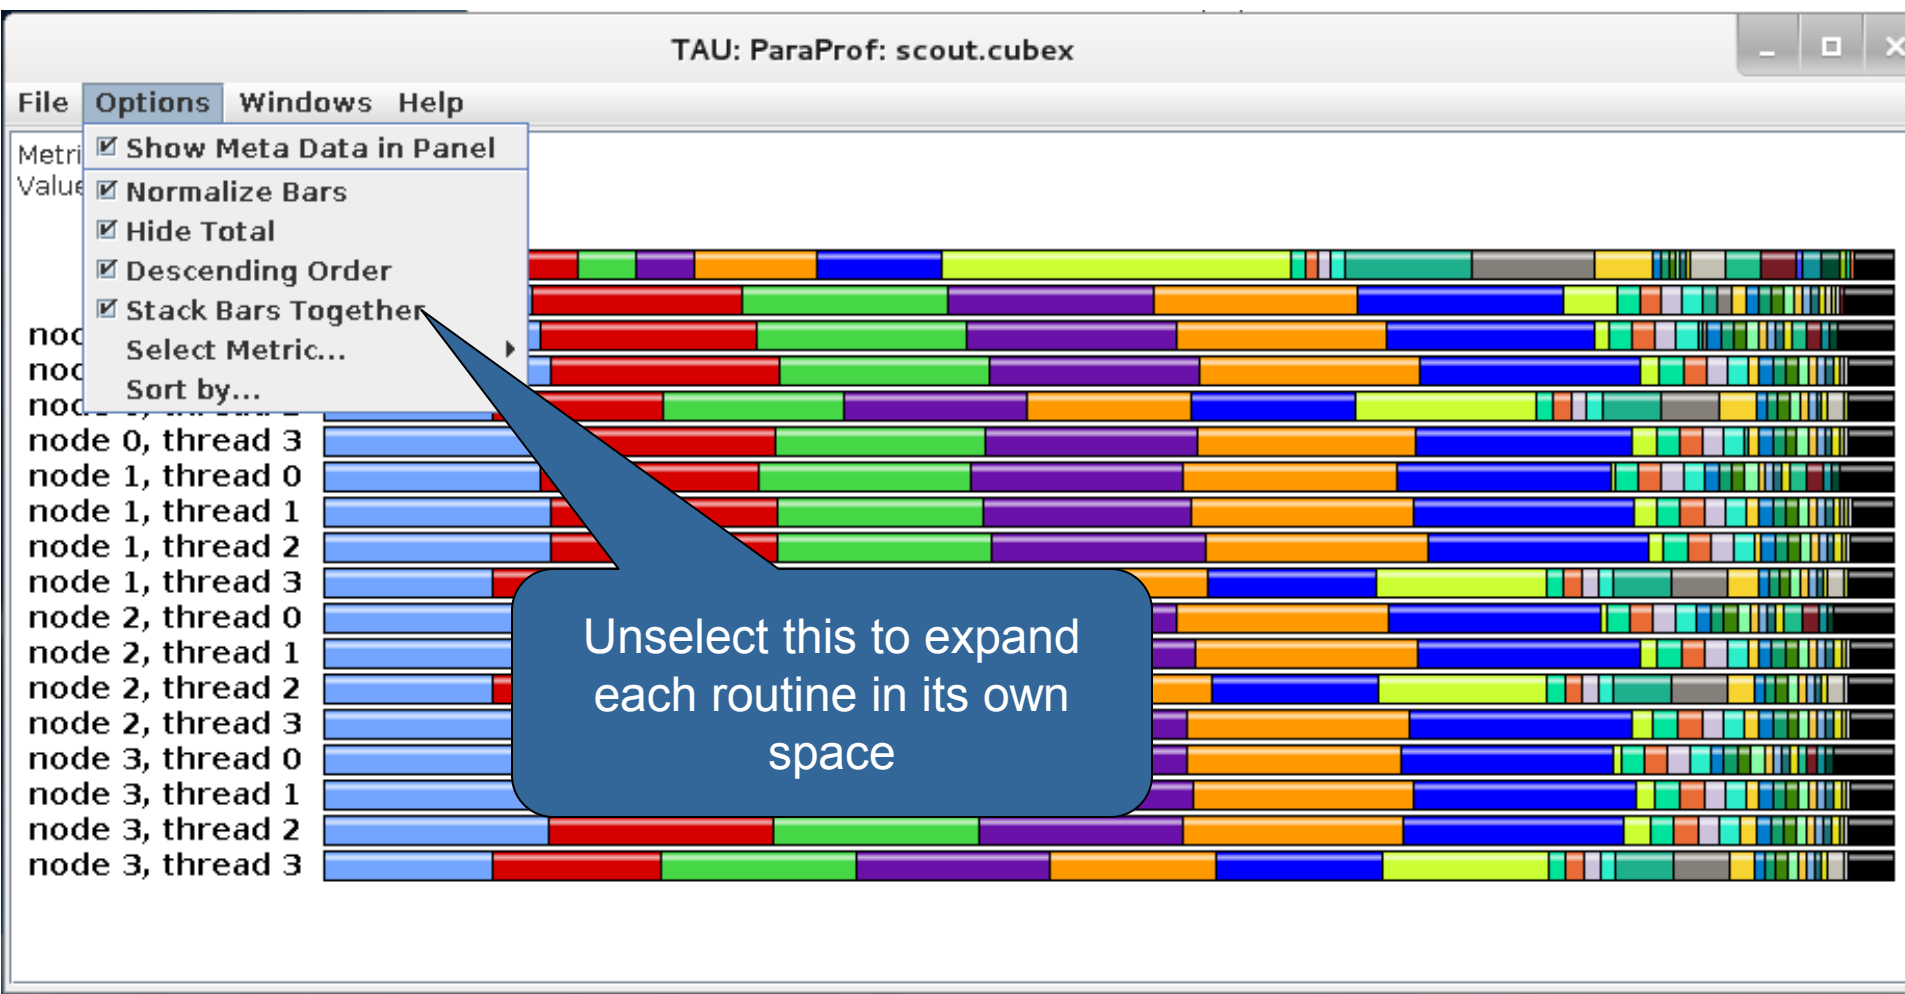
Applications

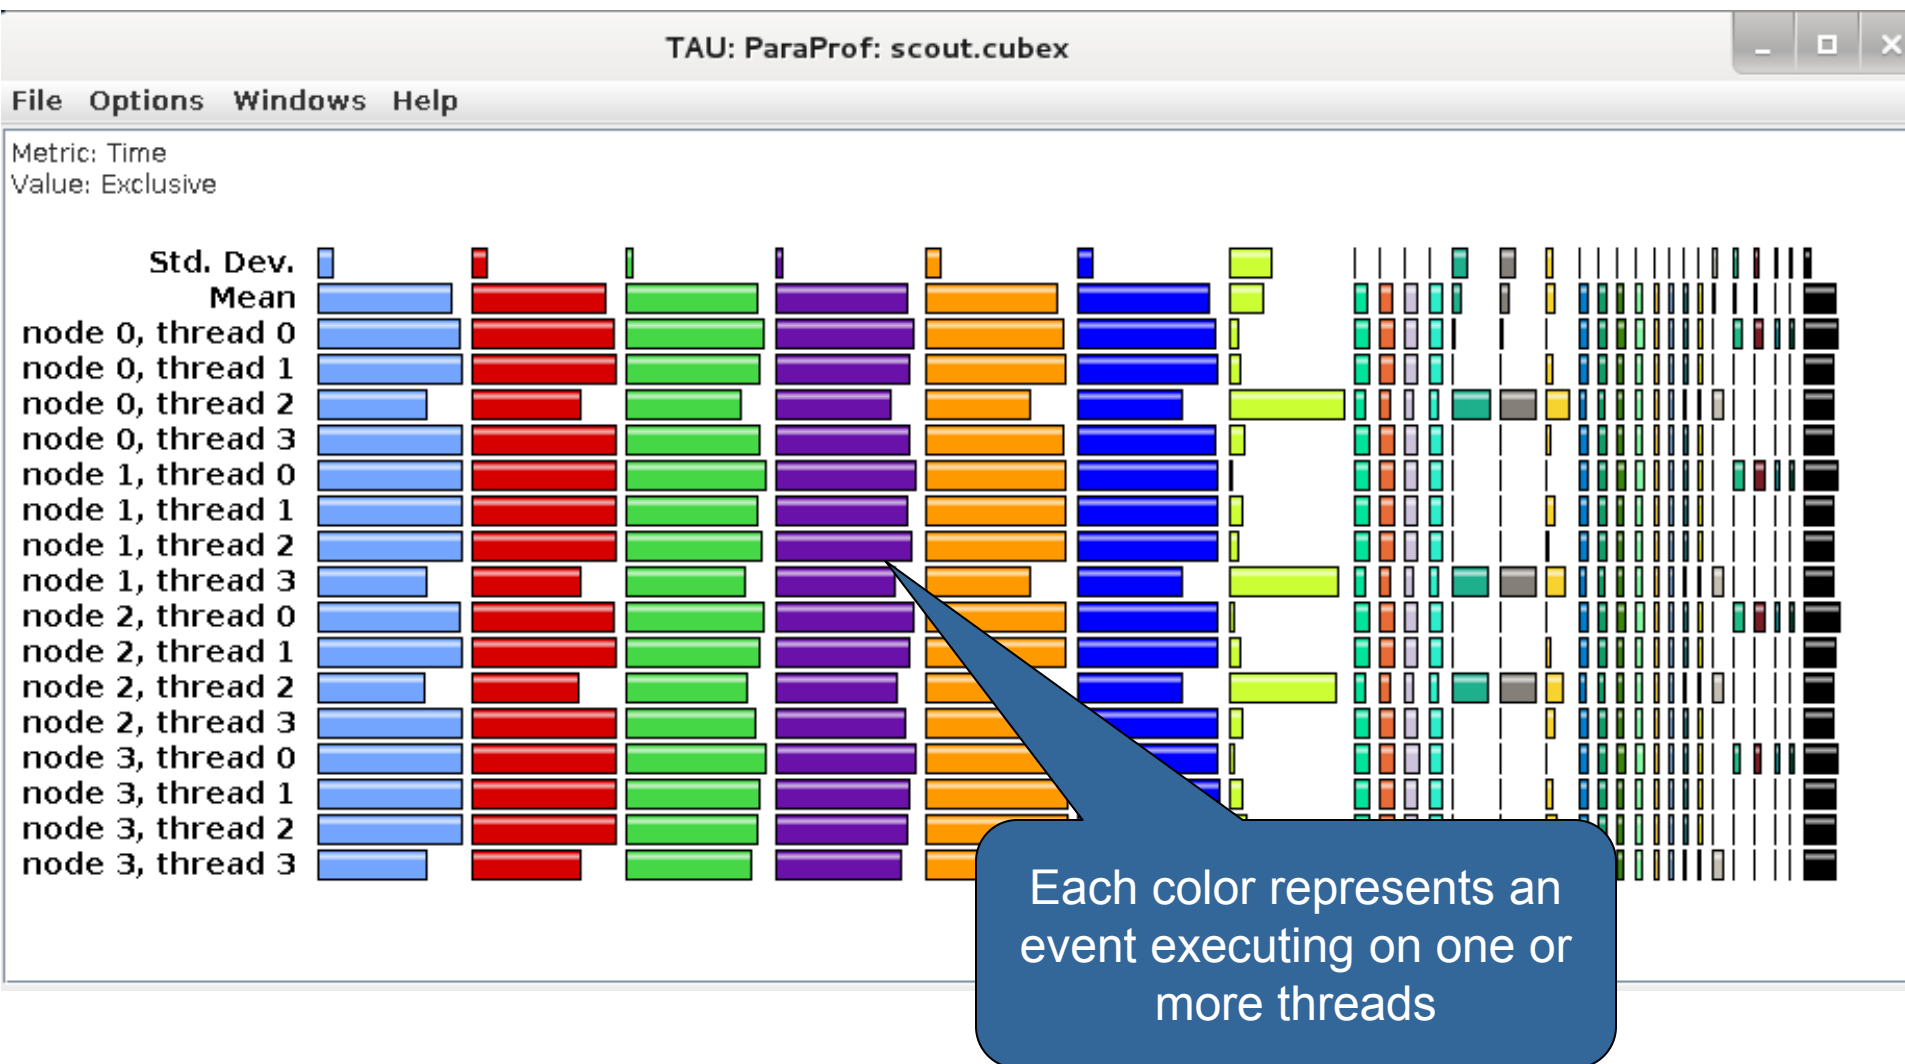
- Standard Applications
 - Default App
 - Default Exp
 - scout.cubex
 - Time
 - Wait at Barrier
 - Barrier Completion
 - Late Sender
 - Late Sender => Messages in Wrong Order
 - Late Sender => Messages in Wrong Order => Messages from different sources
 - Late Sender => Messages in Wrong Order => Messages from same source
 - Late Receiver
 - Early Reduce
 - Early Scan
 - Late Broadcast
 - Wait at N x N
 - N x N Completion
 - Management
 - Management => Fork
 - P2P send synchronizations
 - P2P send synchronizations => Late Receivers
 - P2P rcv synchronizations
 - P2P rcv synchronizations => Late Senders
 - P2P rcv synchronizations => Late Senders => Messages in Wrong Order
 - Collective synchronizations
 - P2P send communications
 - P2P send communications => Late Receivers
 - P2P rcv communications
 - P2P rcv communications => Late Senders
 - P2P rcv communications => Late Senders => Messages in Wrong Order
 - Collective exchange communications
 - Collective communications as source
 - Collective communications as destination
 - P2P bytes sent
 - P2P bytes received
 - Collective bytes outgoing
 - Collective bytes incoming
 - RMA bytes received
 - RMA bytes put

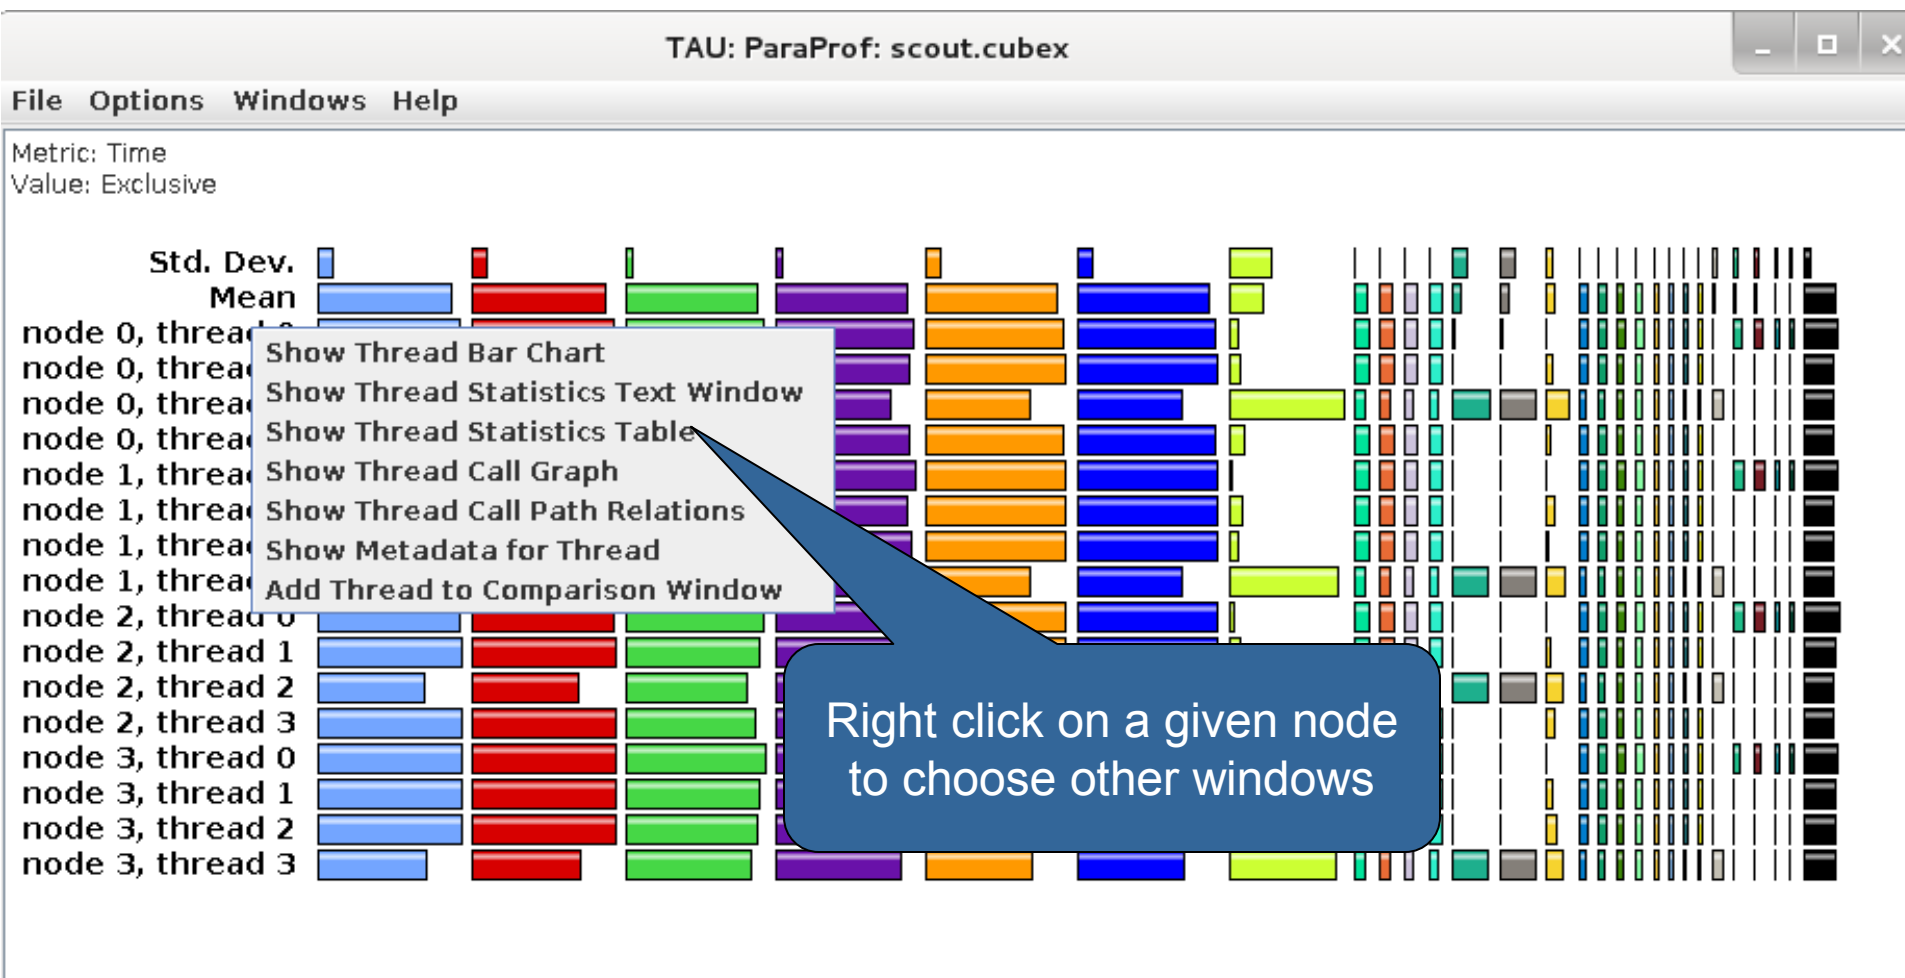
TrialField	Value
Name	scout.cubex
Application ID	0
Experiment ID	0
Trial ID	0
File Type Index	9
File Type Name	Cube

Metrics in the profile









ParaProf: Thread Statistics Table

TAU: ParaProf: Statistics for: node 0, thread 0 - scout.cubex

File Options Windows Help

Time

Name	Exclusive Time ▾	Inclusive Time	Calls	Child Calls
!\$omp do @y_solve.f:52	5.817	5.817	3,216	0
!\$omp do @z_solve.f:52	5.657	5.657	3,216	0
!\$omp do @x_solve.f:54	5.609	5.609	3,216	0
!\$omp do @rhs.f:191	0.609	0.609	3,232	0
!\$omp do @rhs.f:80	0.583	0.583	3,232	0
MPI_Waitall	0.402	0.402	1	0
!\$omp implicit barrier	0.402	0.402	1	0
!\$omp do @rhs.f:301	0.36	0.36	1	0
!\$omp implicit barrier	0.026	0.026	1	0
!\$omp implicit barrier	0	0	1	0
!\$omp do @rhs.f:37	0.343	0.343	1	0
!\$omp do @rhs.f:62	0.225	0.228	3,232	3,232
!\$omp implicit barrier	0.004	0.004	3,216	0
!\$omp implicit barrier	0	0	16	0
MPI_Init_thread	0.218	0.218	1	0
!\$omp do @rhs.f:384	0.199	0.199	3,232	0
!\$omp parallel do @add.f:22	0.099	0.111	3,216	3,216
!\$omp do @rhs.f:428	0.069	0.069	3,232	0
MPI_Isend	0.043	0.043	603	0
!\$omp do @initialize.f:50	0.04	0.04	32	0
!\$omp parallel @rhs.f:28	0.03	2.536	3,232	51,712
!\$omp parallel do @exch_qbc.f:215	0.021	0.029	6,432	6,432
!\$omp parallel do @exch_qbc.f:255	0.02	0.033	6,432	6,432
!\$omp parallel @exch_qbc.f:255	0.02	0.053	6,432	6,432
!\$omp parallel @exch_qbc.f:244				

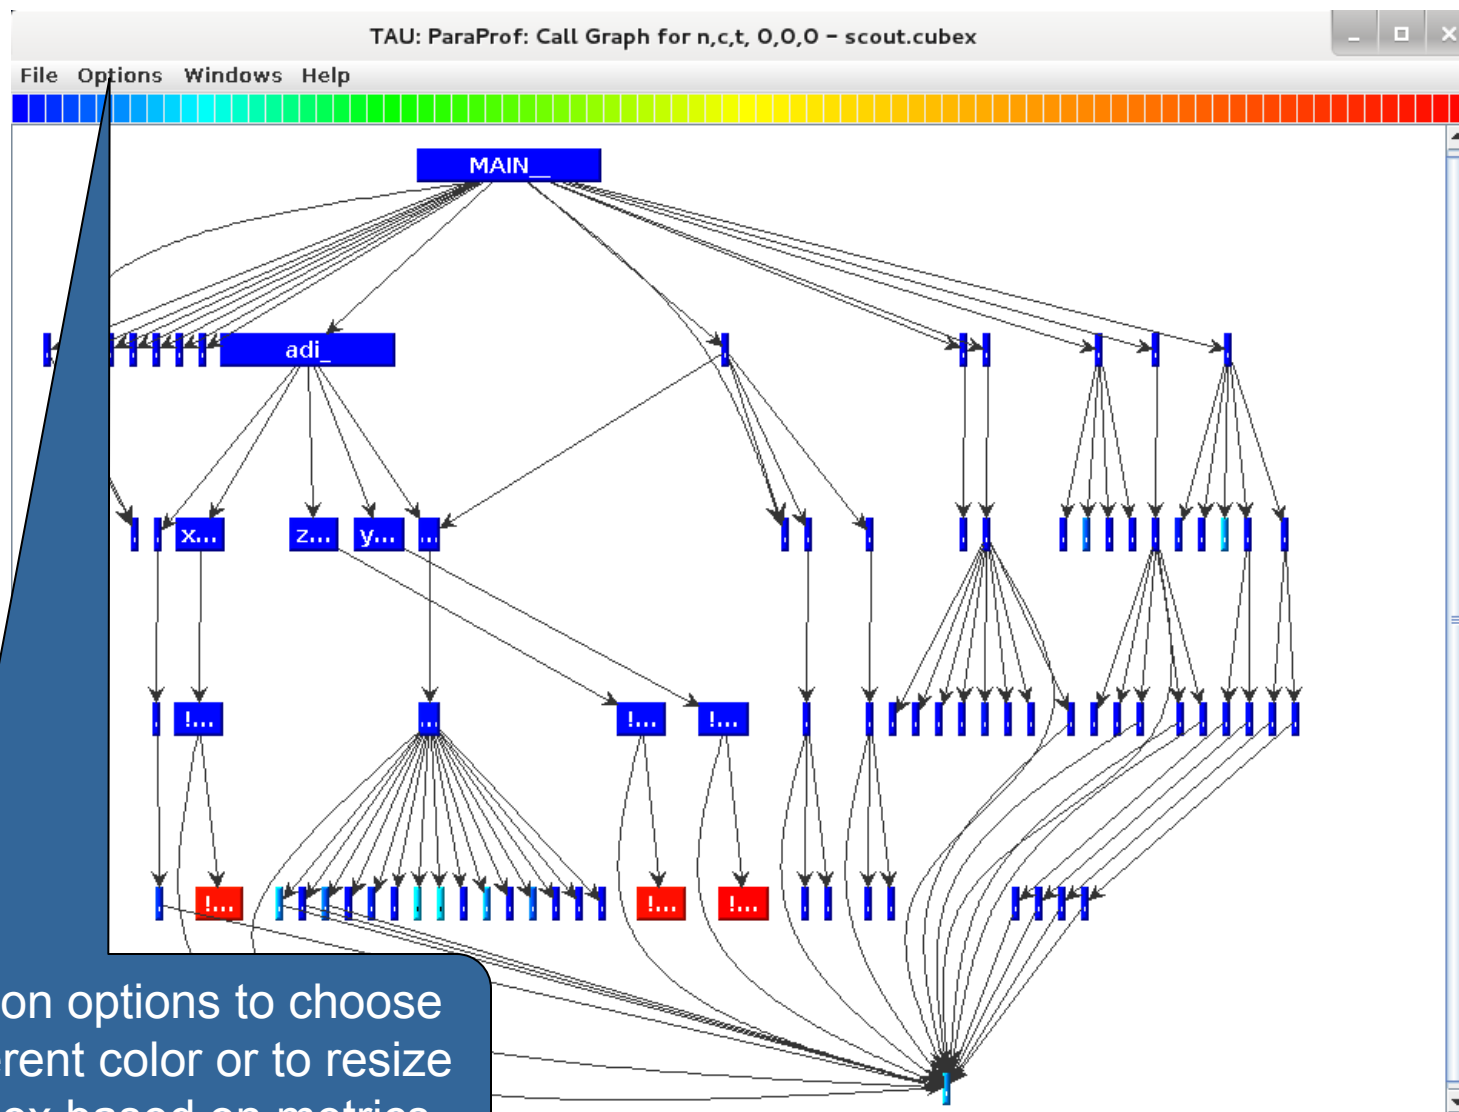
Click to sort by a given metric, drag and move to rearrange columns

Example: Score-P with TAU (LU NPB)

TAU: ParaProf: Statistics for: node 0, thread 0 - profile.cubex

File Options Windows Help

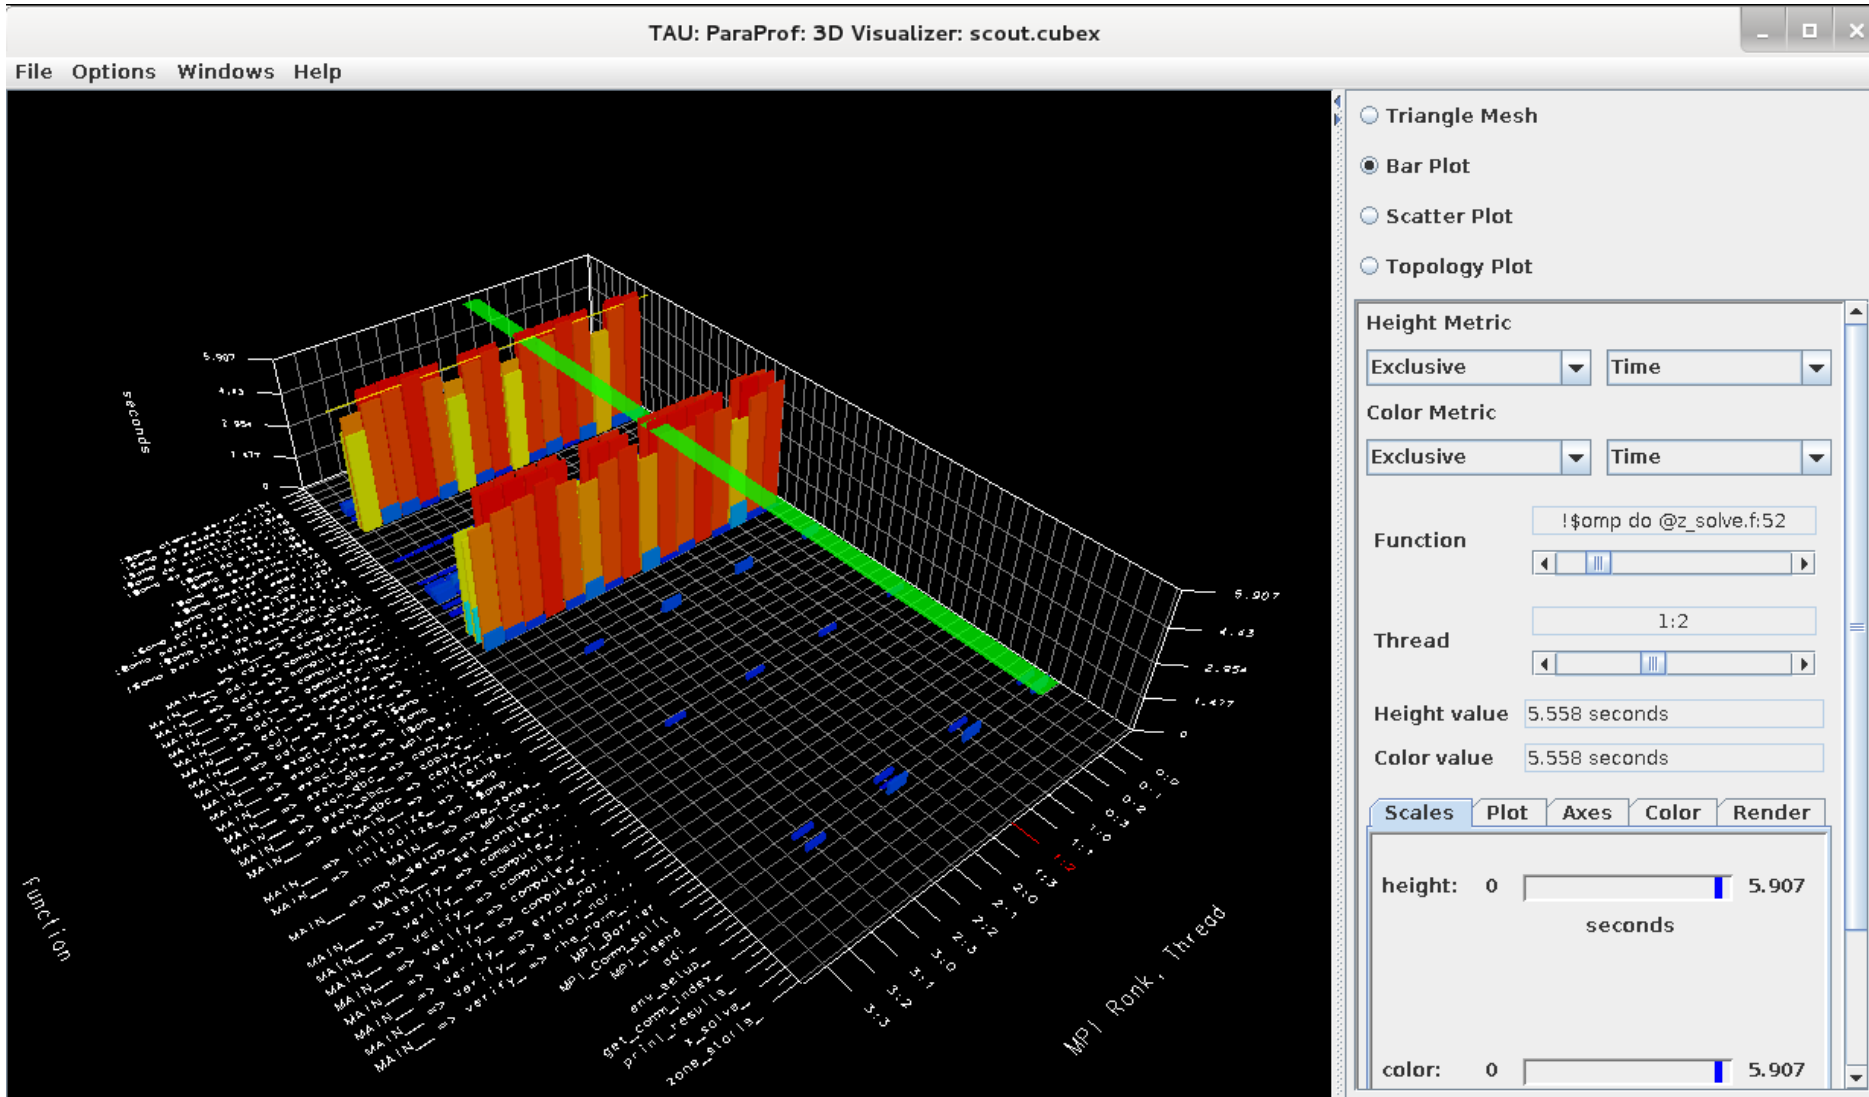
Name	Exclusive Time ▾	Inclusive Time	Calls	Child Calls
APPLU [{lu.f} {46,7}-{162,9}]	0	8.035	1	19
SSOR [{ssor.f} {4,7}-{241,9}]	0.064	6.225	2	37,643
RHS [{rhs.f} {5,7}-{504,9}]	0.743	2.524	303	606
BLTS [{blts.f} {4,7}-{259,9}]	0.613	0.658	9,331	18,662
BUTS [{buts.f} {4,7}-{259,9}]	0.612	1.871	9,331	18,662
EXCHANGE_1 [{exchange_1.f} {5,7}-{177,9}]	0.024	1.259	18,662	18,662
MPI_Recv	1.235	1.235	18,662	0
MPI_Send	0	0	0	0
JACU [{jacu.f} {5,7}-{384,9}]	0.532	0.532	9,331	0
JACLD [{jacld.f} {5,7}-{384,9}]	0.522	0.522	9,331	0
MPI_Allreduce	0.018	0.018	2	0
L2NORM [{l2norm.f} {4,7}-{68,9}]	0	0.035	4	4
MPI_Barrier	0	0	2	0
TIMER_START [{timers.f} {23,7}-{37,9}]	0	0	2	0
TIMER_STOP [{timers.f} {43,7}-{59,9}]	0	0	2	0
TIMER_CLEAR [{timers.f} {4,7}-{17,9}]	0	0	2	0
TIMER_READ [{timers.f} {65,7}-{77,9}]	0	0	2	0
SETIV [{setiv.f} {4,7}-{67,9}]	0.043	0.111	2	95,232
PROC_GRID [{proc_grid.f} {5,7}-{34,9}]	0.011	0.011	1	0
ERHS [{erhs.f} {4,7}-{536,9}]	0.004	0.108	1	2
ERROR [{error.f} {4,7}-{81,9}]	0.004	0.009	1	7,937
SETBV [{setbv.f} {5,7}-{79,9}]	0.002	0.004	2	3,400
READ_INPUT [{read_input.f} {5,7}-{125,9}]	0	0.001	1	2
VERIFY [{verify.f} {5,9}-{403,11}]	0	0	1	0
PRINT_RESULTS [{print_results.f} {2,7}-{115,12}]	0	0	1	0
PINTGR [{pintgr.f} {5,7}-{288,9}]	0	0	1	6
INIT_COMM [{init_comm.f} {5,7}-{57,9}]	0	1.565	1	4
MPI_Finalize	0	0	1	0
SETHYPER [{sethyper.f} {5,7}-{94,9}]	0	0	1	0
NEIGHBORS [{neighbors.f} {5,7}-{48,9}]	0	0	1	0
SETCOEFF [{setcoeff.f} {5,7}-{157,9}]	0	0	1	0



ParaProf: Callpath Thread Relations Window

TAU: ParaProf: Call Path Data n,c,t, 0,0,0 - scout.cubex				
File Options Windows Help				
Metric Name: Time				
Sorted By: Exclusive				
Units: seconds				
-->	0.04	0.04	32/32	!\$omp parallel @initialize.f:28
	0.04	0.04	32	!\$omp do @initialize.f:50
-->	0.03	2.536	3232/3232	compute_rhs_
	0.03	2.536	3232	!\$omp parallel @rhs.f:28
	9.8E-4	9.8E-4	3232/3232	!\$omp master @rhs.f:424
	0.225	0.228	3232/3232	!\$omp do @rhs.f:62
	0.002	0.002	3232/3232	!\$omp master @rhs.f:74
	0.002	0.002	3232/3232	!\$omp master @rhs.f:293
	0.199	0.199	3232/3232	!\$omp do @rhs.f:384
	0.002	0.002	3232/3232	!\$omp master @rhs.f:183
	0.343	0.343	3232/3232	!\$omp do @rhs.f:37
	0.016	0.016	3232/3232	!\$omp do @rhs.f:372
	0.014	0.027	3232/3232	!\$omp do @rhs.f:413
	0.609	0.609	3232/3232	!\$omp do @rhs.f:191
	0.36	0.386	3232/3232	!\$omp do @rhs.f:301
	0.583	0.583	3232/3232	!\$omp do @rhs.f:80
	0.019	0.019	3232/3232	!\$omp do @rhs.f:400
	0.006	0.006	3232/51680	!\$omp implicit barrier
	0.069	0.069	3232/3232	!\$omp do @rhs.f:428
	0.015	0.015	3232/3232	!\$omp do @rhs.f:359
-->	0.021	0.029	6432/6432	!\$omp parallel @exch_qbc.f:215
	0.021	0.029	6432	!\$omp parallel do @exch_qbc.f:215
	0.007	0.007	6432/51680	!\$omp implicit barrier
-->	0.02	0.033	6432/6432	!\$omp parallel @exch_qbc.f:255
	0.02	0.033	6432	!\$omp parallel do @exch_qbc.f:255
	0.013	0.013	6432/51680	!\$omp implicit barrier


ParaProf:Windows -> 3D Visualization -> Bar Plot



ParaProf: 3D Scatter Plot

TAU: ParaProf: 3D Visualizer: scout.cubex

File Options Windows Help



- ☐ Triangle Mesh
- ☐ Bar Plot
- ☒ Scatter Plot
- ☐ Topology Plot

Width: !\$omp do @y_solve.f:52 ...
Exclusive Time


Depth: !\$omp do @z_solve.f:52 ...
Exclusive Time

Height: !\$omp do @x_solve.f:54 ...
Exclusive Time

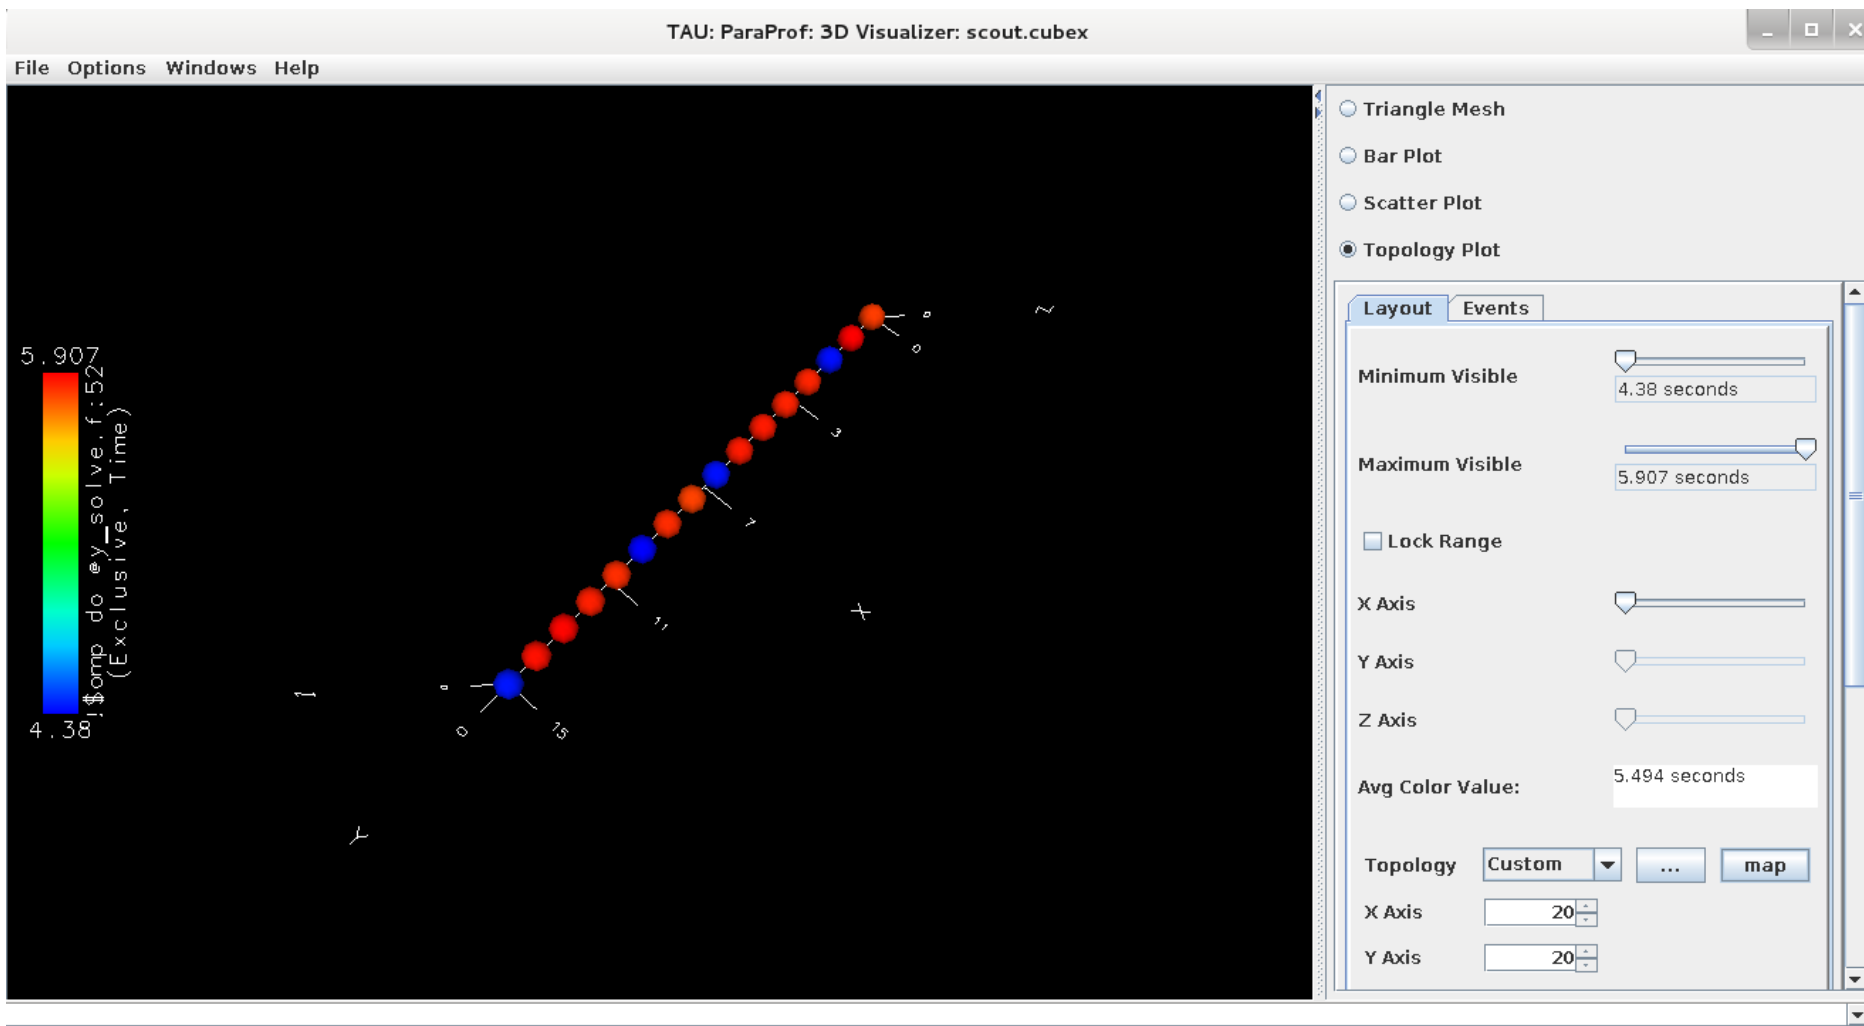
Color: !\$omp implicit barrier ...
Exclusive Time

ColorScale Render
ScatterPlot Axes

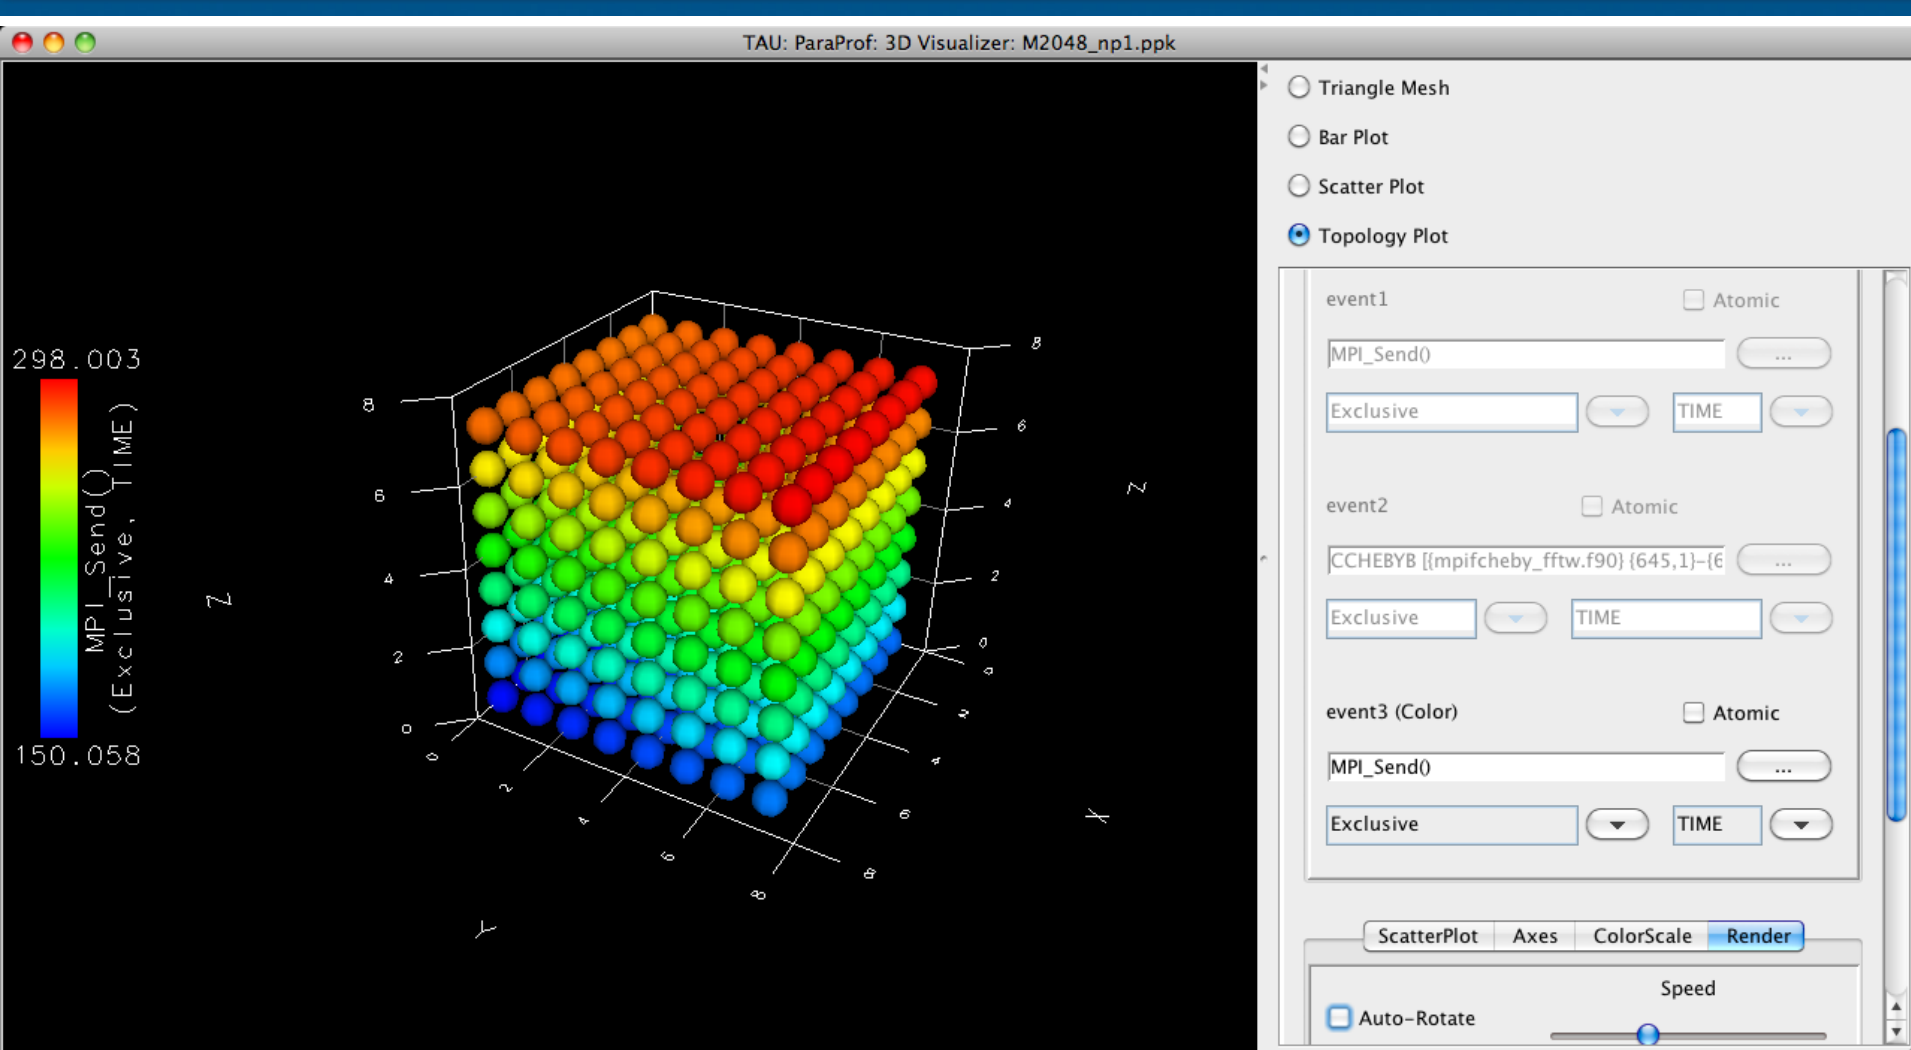
Auto-Rotate Speed
Reverse Video Stereo



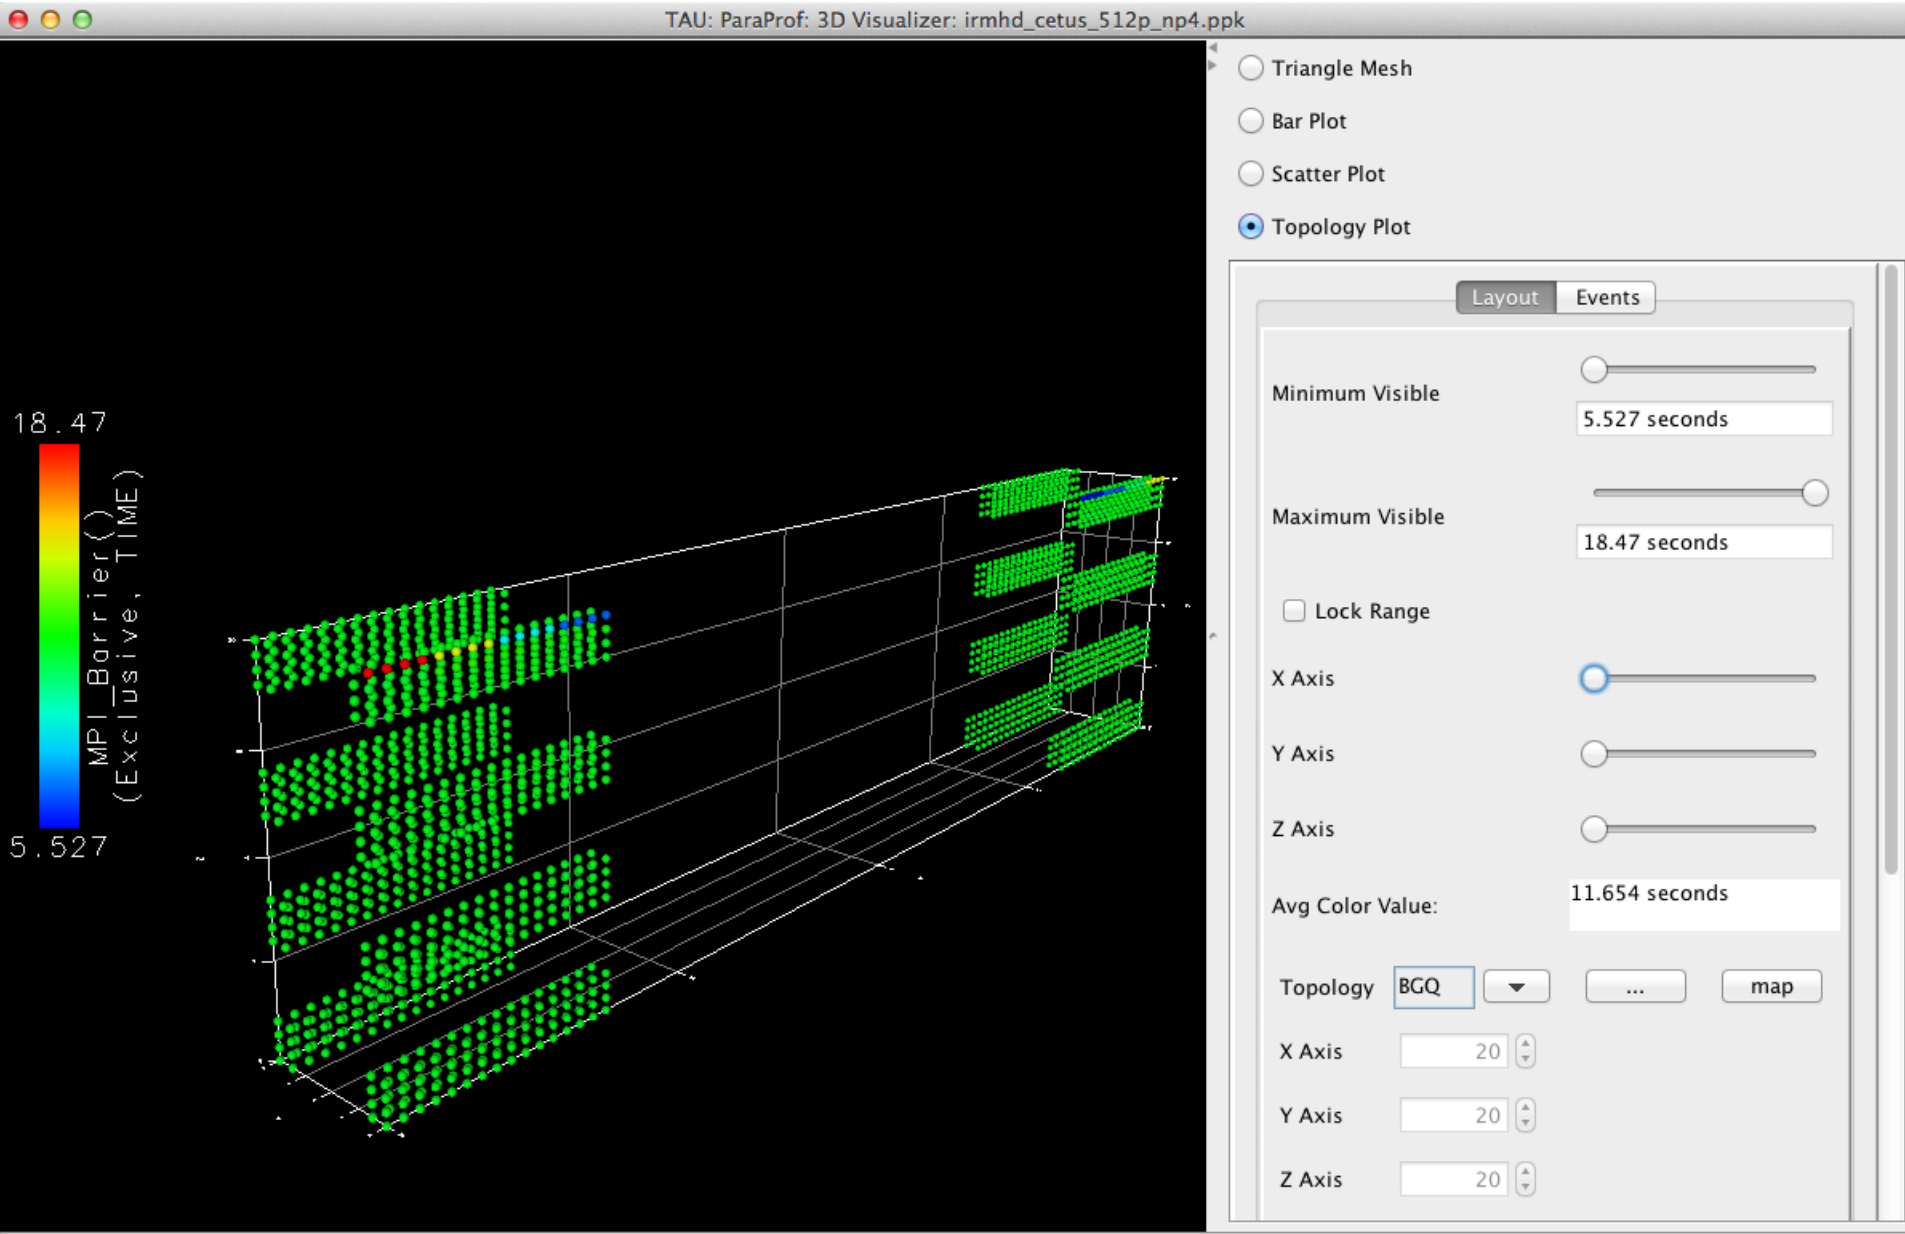
ParaProf: 3D Topology View for a Routine

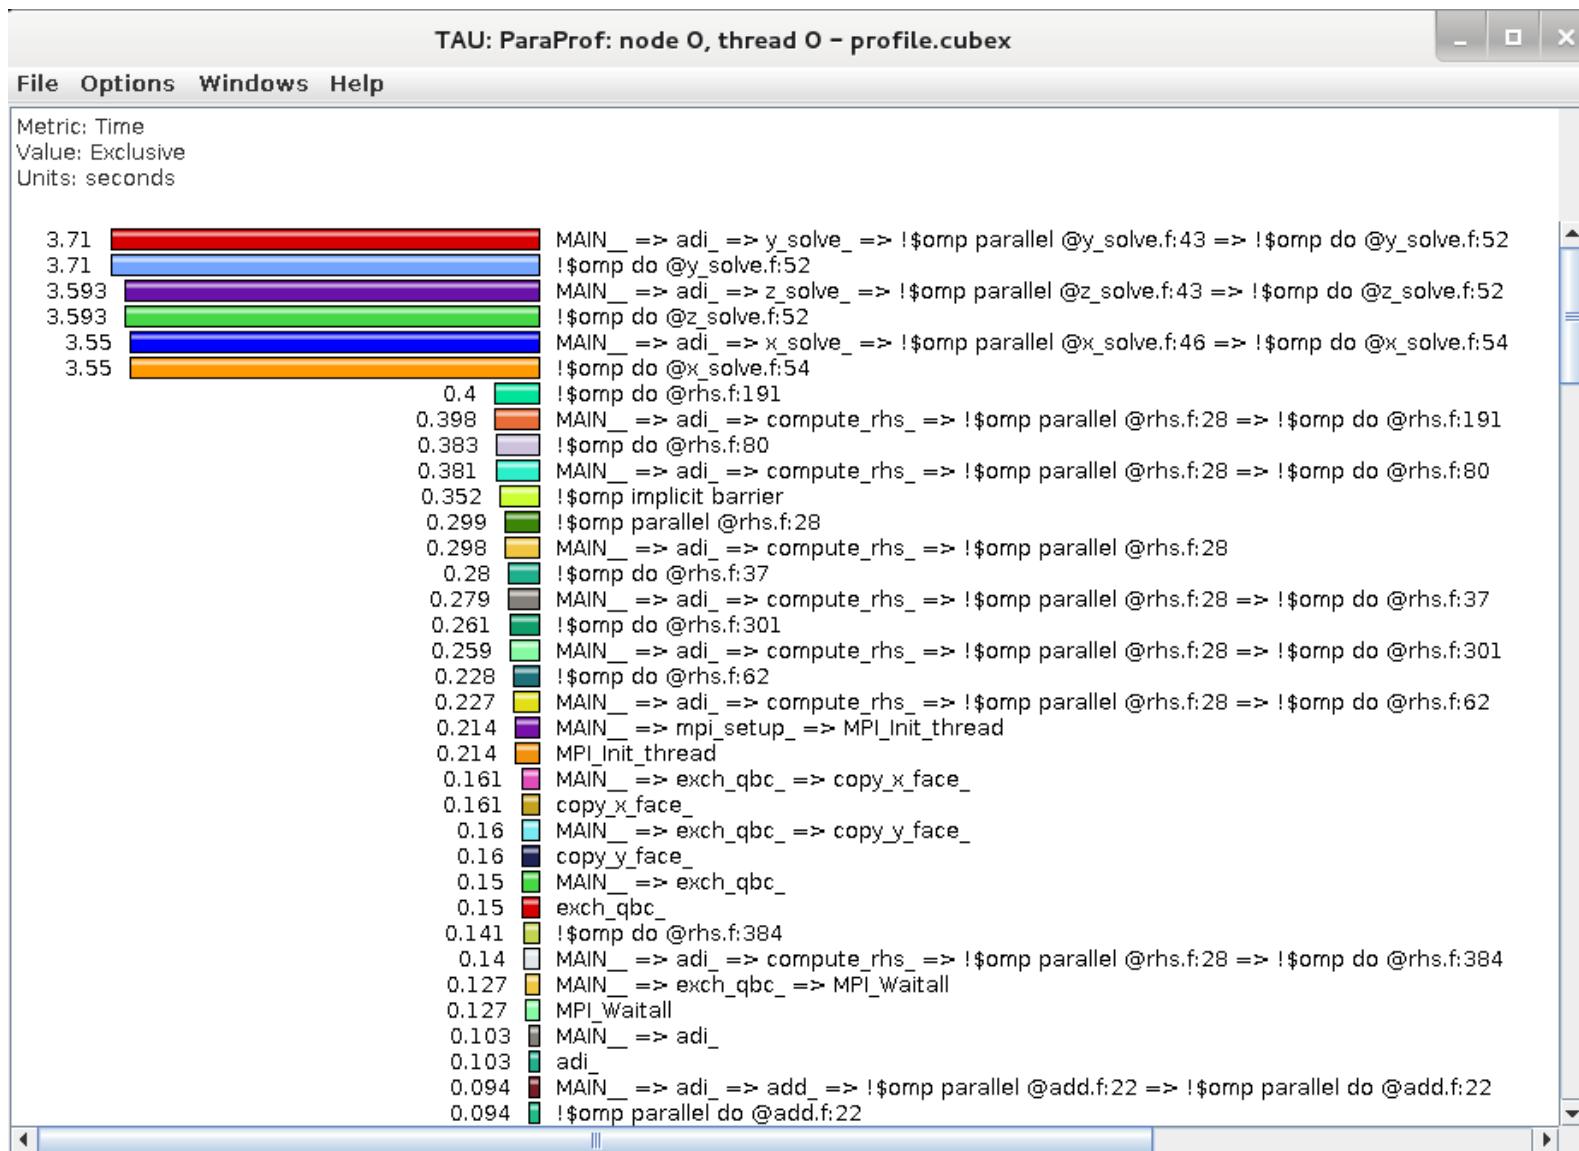


ParaProf: Topology View 3D Torus (IBM BG/P)

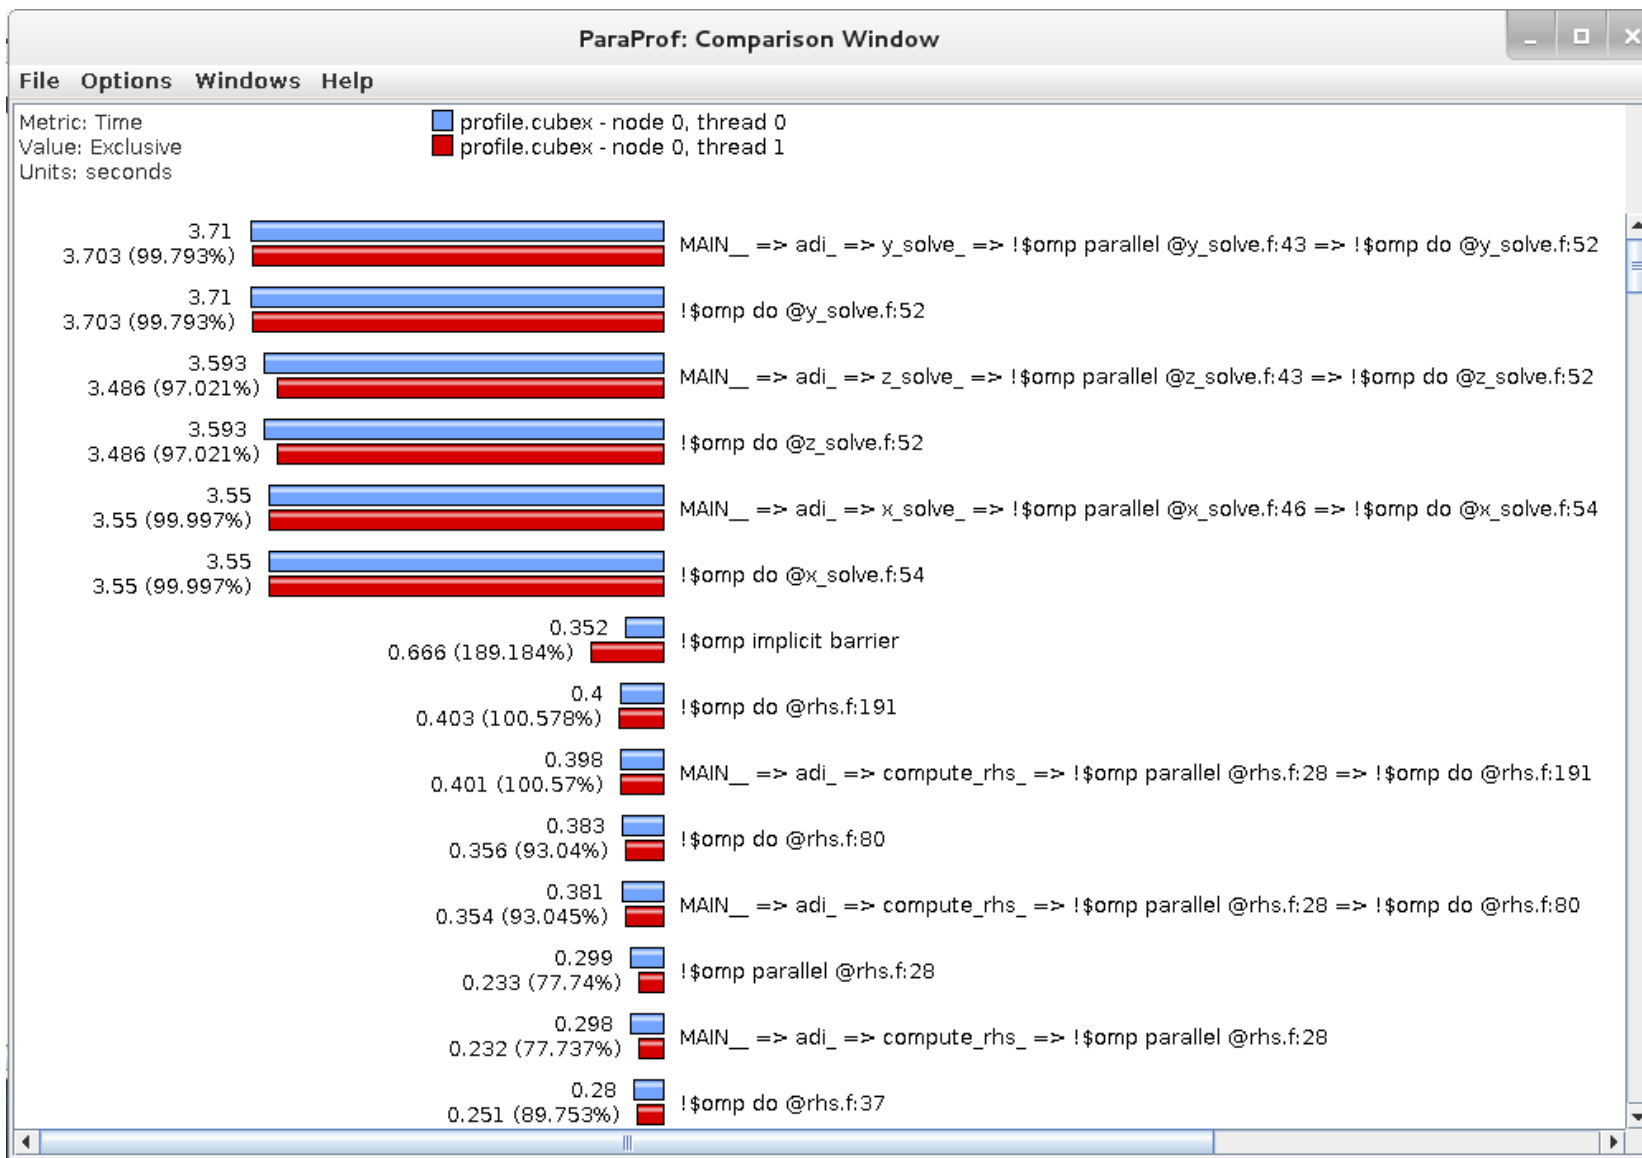


ParaProf:Topology View (6D Torus Coordinates BG/Q)





ParaProf: Add Thread to Comparison Window



ParaProf: Score-P Profile Files, Database



TAU: ParaProf Manager

File Options Help

Applications

- Standard Applications
 - Default App
 - Default Exp
 - profile.cubex
 - Time
 - Minimum Inclusive Time
 - Maximum Inclusive Time
 - PAPI_TOT_CYC
 - PAPI_TOT_INS
 - PAPI_FP_INS
 - ru_utime
 - ru_stime
 - ru_maxrss
 - ru_ixrss
 - ru_idrss
 - ru_isrss
 - ru_minflt
 - ru_majflt
 - ru_nswap
 - ru_inblock
 - ru_oublock
 - ru_msgsnd
 - ru_msgrcv
 - ru_nsignals
 - ru_nvcsw
 - ru_nivcsw
 - bytes_sent
 - bytes_received
- Default (jdbc:h2:/home/livetau/.ParaProf/perfdmf;AUTO_SERVER=TRUE)
- perfexplorer_working (jdbc:h2:/home/livetau/.ParaProf/perfexplorer_wo

Add Application
Add Experiment
Add Trial

TrialField	Value
Name	profile.cubex
Application ID	0
Experiment ID	0
Trial ID	0
File Type Index	9
File Type Name	Cube

ParaProf: File -> Preferences

ParaProf Preferences

File

Font

SansSerif

☐ Bold

Size

☐ Italic

0 10 20 30 40

Window defaults

Units

Seconds

☐ Show Values as Percent

Settings

☐ Show Path Title in Reverse

☐ Reverse Call Paths

☒ Interpret threads that do not call a given function as a 0 value for statistics computation

☐ Generate data for reverse calltree
(requires lots of memory)
(does not apply to currently loaded profiles)

☒ Show Source Locations

☒ Auto label node/context/threads

Restore Defaults

Apply Cancel

n,c,t 0,0,0

n,c,t 0,0,1

n,c,t 0,0,2

ParaProf: Group Changer Window

TAU: ParaProf: Group Changer: profile.cubex

Region	Current	Available
<p>filter: <input type="text"/></p> <ul style="list-style-type: none">!\$omp atomic @error.f:104!\$omp atomic @error.f:51!\$omp do @error.f:33!\$omp do @error.f:91!\$omp do @exact_rhs.f:147!\$omp do @exact_rhs.f:247!\$omp do @exact_rhs.f:31!\$omp do @exact_rhs.f:346!\$omp do @exact_rhs.f:46!\$omp do @initialize.f:100!\$omp do @initialize.f:119!\$omp do @initialize.f:137!\$omp do @initialize.f:156!\$omp do @initialize.f:174!\$omp do @initialize.f:192!\$omp do @initialize.f:31	<p>CUBE_DEFAULT</p>	<p><input type="text"/> <input type="button" value="new group"/></p> <p>CUBE_CALLPATH</p>

Navigation buttons: and

ParaProf: Options -> Derived Metric Panel



TAU: ParaProf Manager

File Options Help

Applications

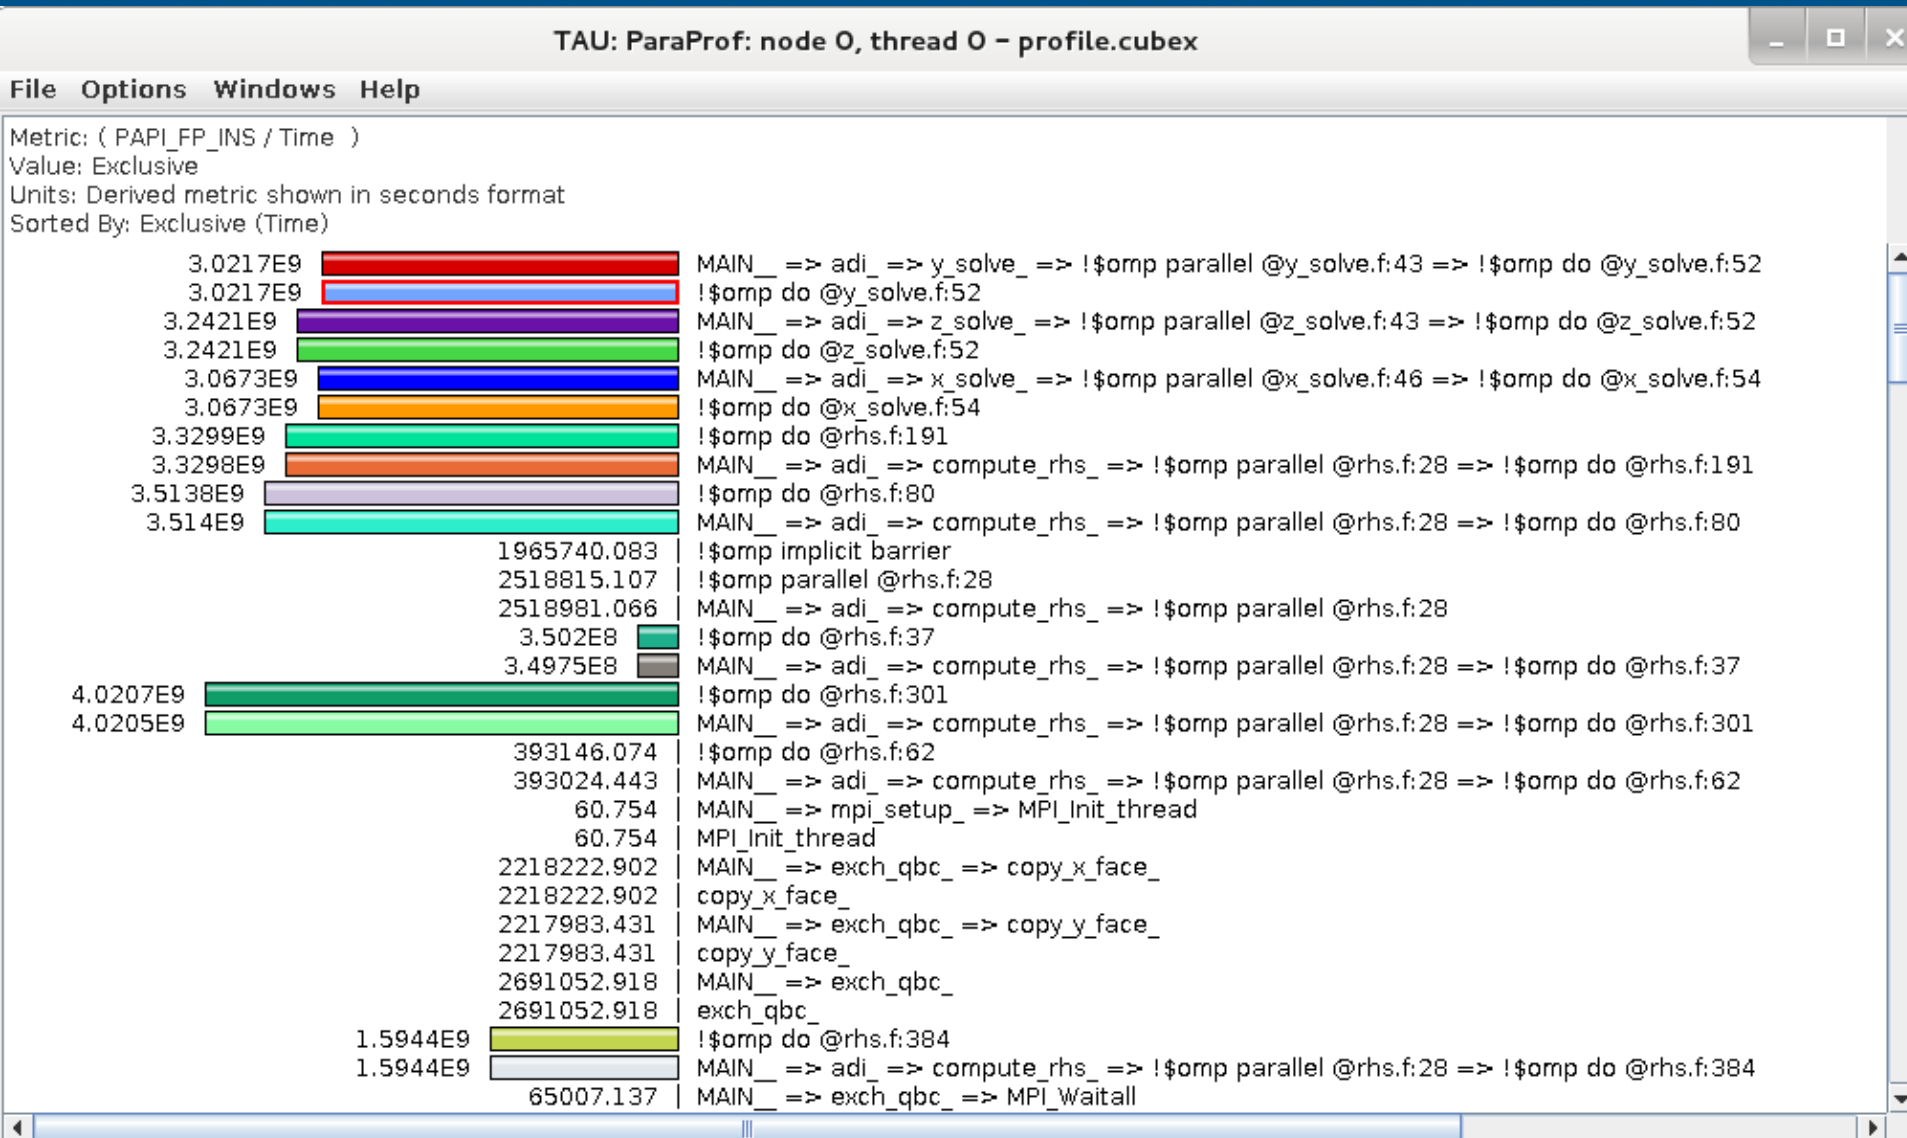
- Standard Applications
 - Default App
 - Default Exp
 - profile.cubex
 - Time**
 - Minimum Inclusive Time
 - Maximum Inclusive Time
 - PAPI_TOT_CYC
 - PAPI_TOT_INS
 - PAPI_FP_INS
 - ru_utime
 - ru_stime
 - ru_maxrss
 - ru_ixrss
 - ru_idrss
 - ru_isrss
 - ru_minflt
 - ru_majflt
 - ru_nswap
 - ru_inblock
 - ru_oublock
 - ru_msgsnd
 - ru_msgrcv
 - ru_nsignals
 - ru_nvcsw

MetricField	Value
Name	Time
Application ID	0
Experiment ID	0
Trial ID	0
Metric ID	0

Expression: "PAPI_FP_INS"/"Time"

+ - * / = { } Apply Clear

Sorting Derived Flops Metric by Exclusive Time



- U.S. Department of Energy (DOE)
 - Office of Science
 - ASC/NNSA, Tri-labs (LLNL, LANL, SNL)
- U.S. Department of Defense (DoD)
 - HPC Modernization Office (HPCMO)
- NSF Software Development for Cyberinfrastructure (SDCI)
- Juelich Supercomputing Center, NIC
- Argonne National Laboratory
- Technical University Dresden
- ParaTools, Inc.
- NVIDIA



ParaTools