VI-HPS

SOFTWARE

FAST SOLUTIONS

PRODUCTIVITY

Profile Analysis with ParaProf

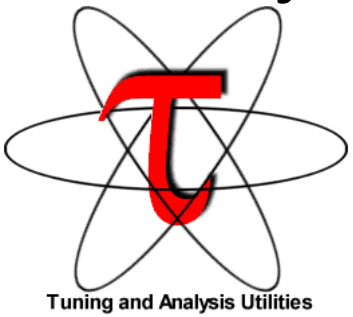

Sameer Shende
Performance Research Lab, University of Oregon
http://TAU.uoregon.edu

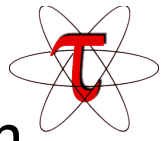

- Parallel performance framework and toolkit
  - Supports all HPC platforms, compilers, runtime system
  - Provides portable instrumentation, measurement, analysis

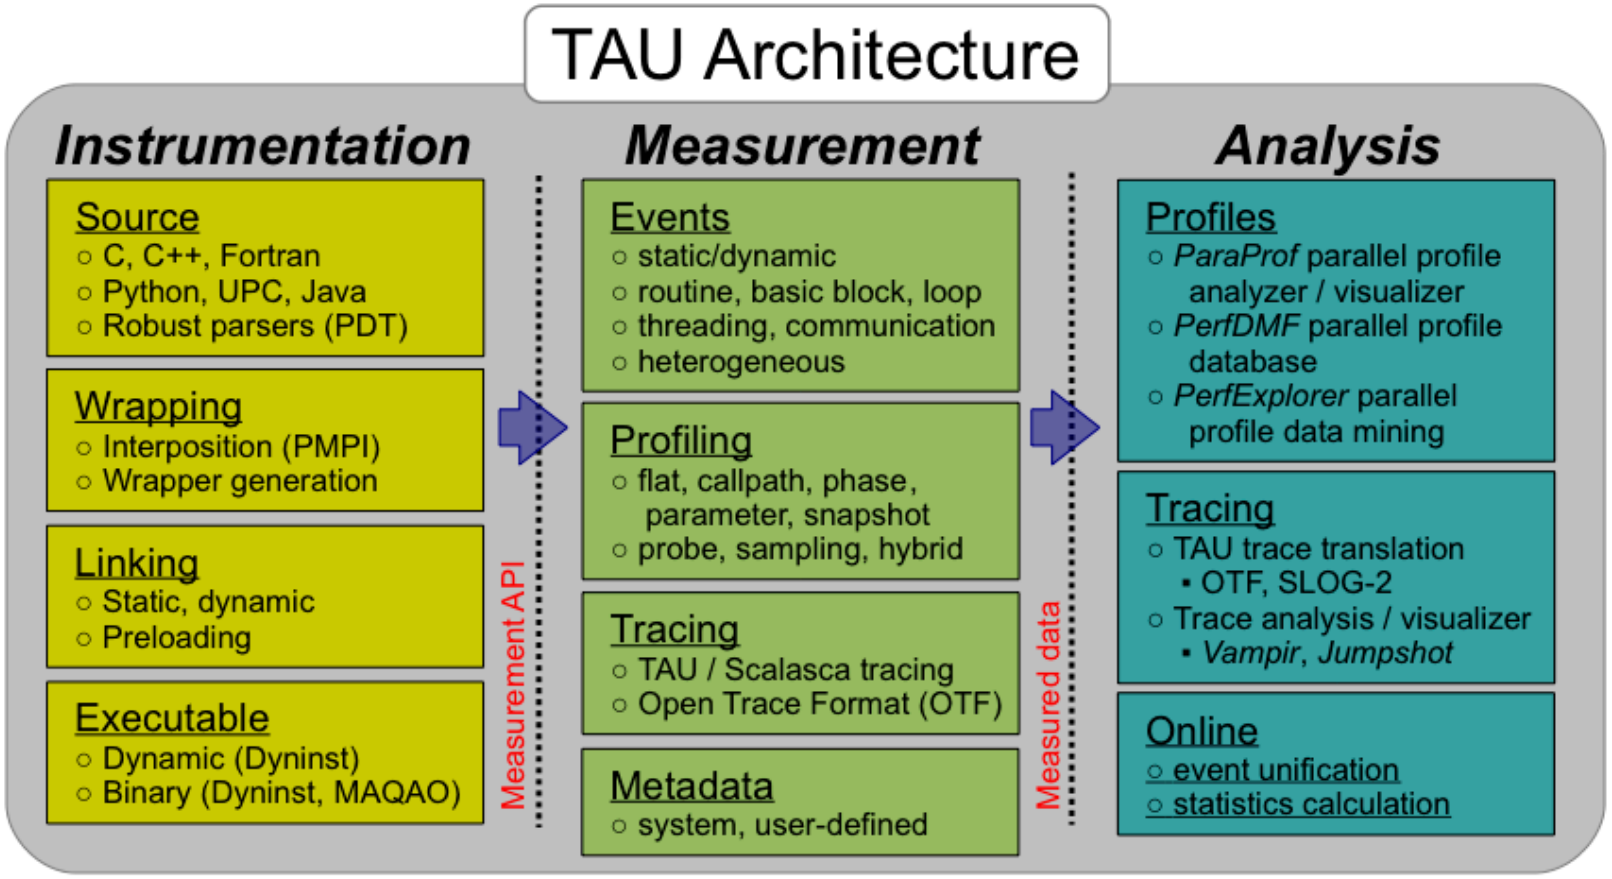

TAU Performance System®

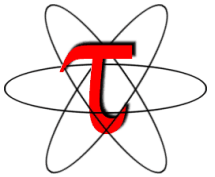

- Instrumentation
  - Fortran, C++, C, UPC, Java, Python, Chapel
  - Automatic instrumentation
- Measurement and analysis support
  - MPI, OpenSHMEM, ARMCI, PGAS, DMAPP
  - pthreads, OpenMP, hybrid, other thread models
  - GPU, CUDA, OpenCL, OpenACC
  - Parallel profiling and tracing
  - Use of Score-P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

- TAU supports both sampling and direct instrumentation
- Memory debugging as well as I/O performance evaluation
- Profiling as well as tracing
- Interfaces with Score-P for more efficient measurements
- TAU's instrumentation covers:
  - Runtime library interposition (tau_exec)
  - Compiler-based instrumentation
  - PDT based Source level instrumentation: routine & loop
  - Event based sampling (TAU_SAMPLING=1)
  - Callstack unwinding with sampling (TAU_EBS_UNWIND=1)
  - OpenMP Tools Interface (OMPT, tau_exec –T ompt)
  - CUDA CUPTI, OpenCL (tau_exec -T cupti -cupti)

LABS on Poincare: paraprof

```
module use /gpfslocal/pub/vihps/UNITE/local
module load UNITE VI-HPS-TW
cd tutorial/NPB3.3-MZ-MPI
make suite
cd bin; cp ../jobscript/mds/run.tau.ll
Uncomment the first, then second run block:
# Case 2: MPI with OpenMP (OpenMP Tools Interface (OMPT))
#mpirun -np ${LOADL_TOTAL_TASKS} tau_exec -T ompt ./bt-
mz_B.4

llsubmit run.tau.ll
Wait and then launch after the job finishes:
paraprof  (Right Click on node 0 or 1, Show Thread
Statistics Table. Show Source Code on an OMPT source
location. Also use paraprof on Score-P *.cubex files.)
```

TAU Analysis

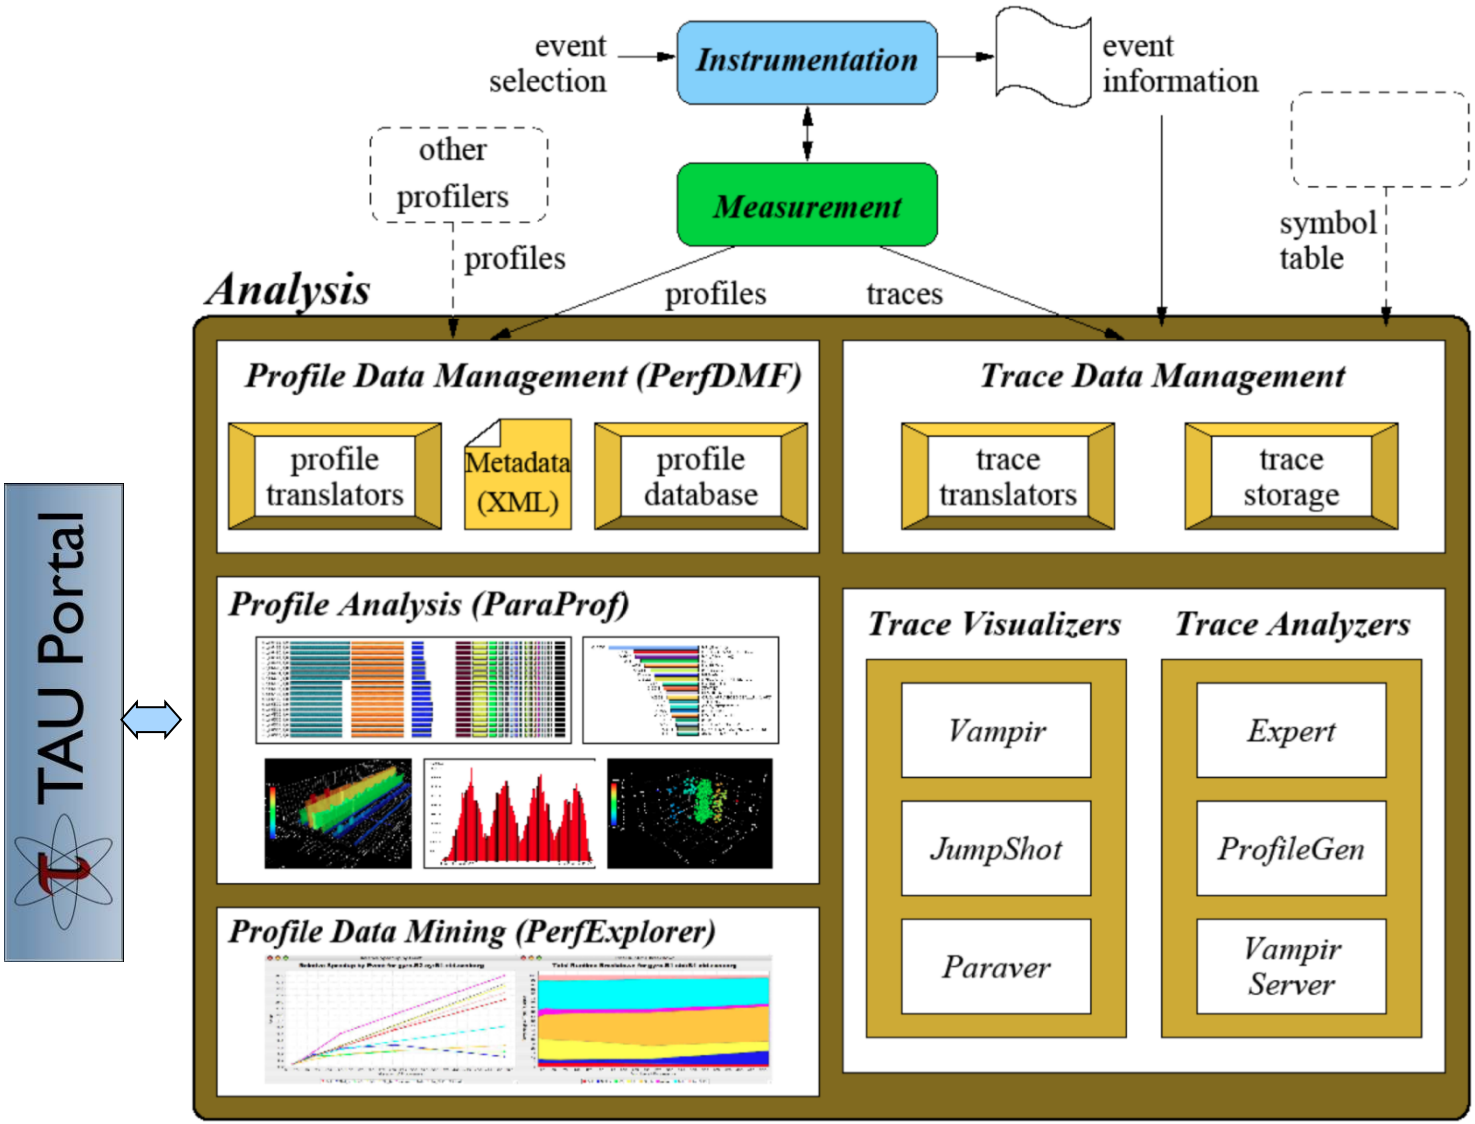

ParaProf Profile Analysis Framework

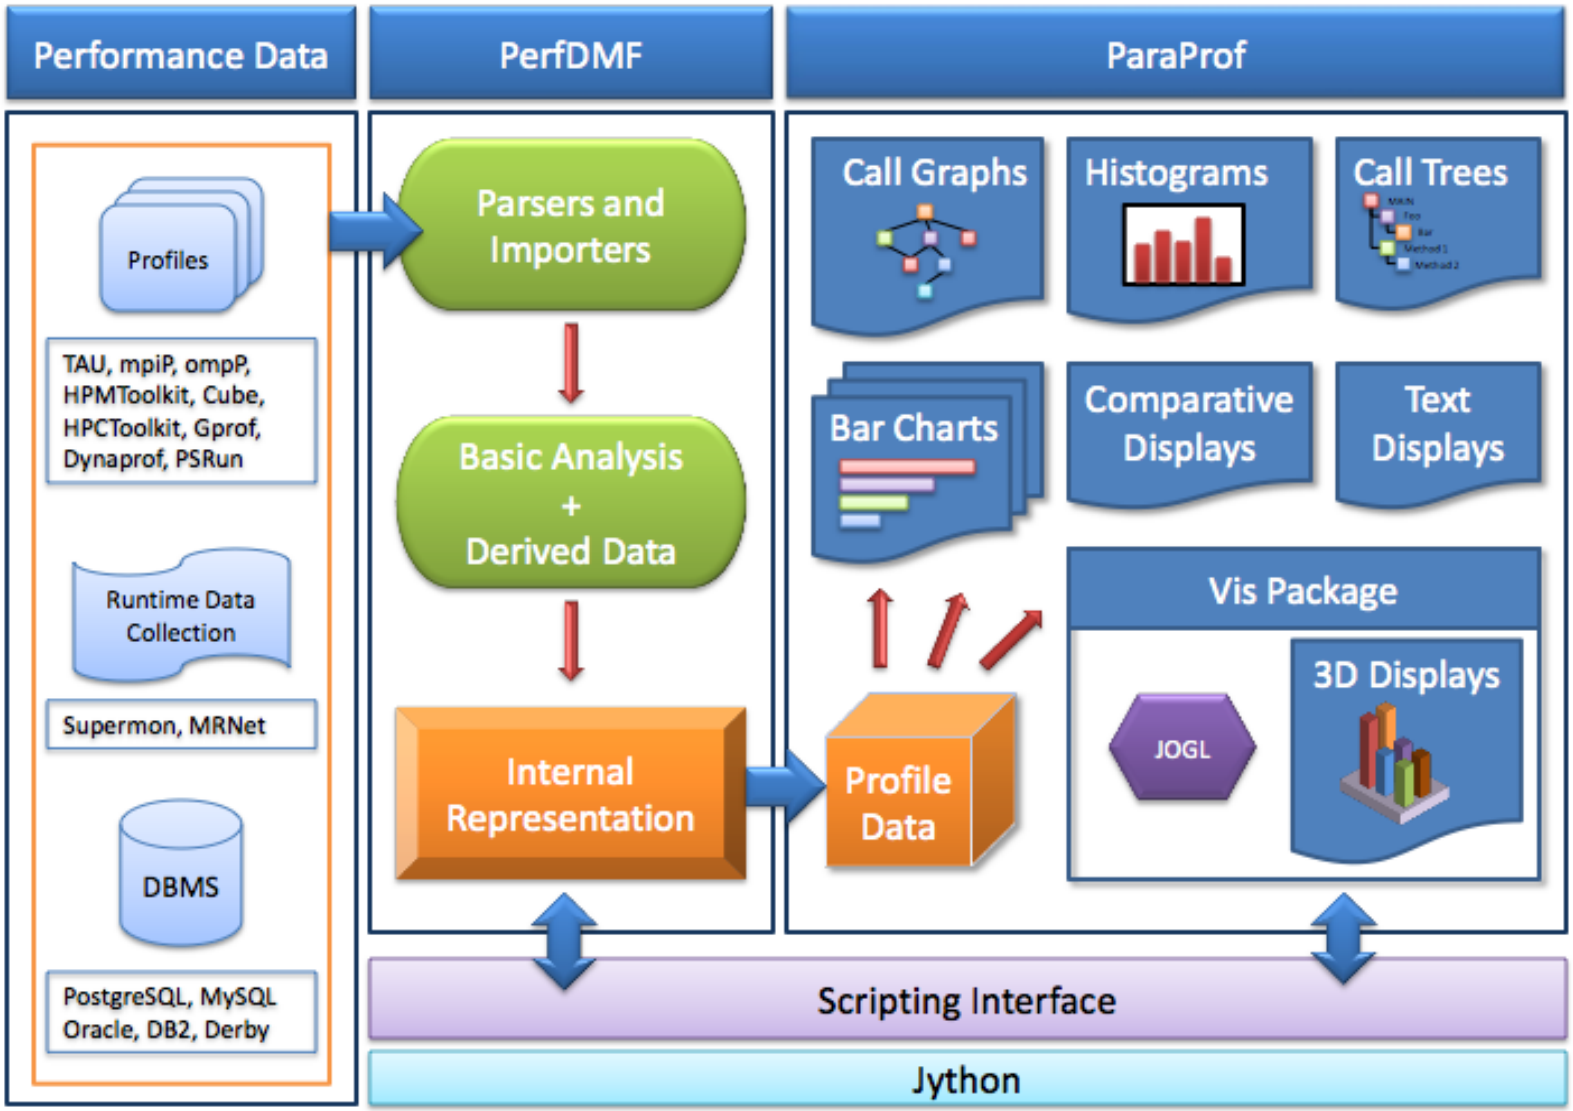

Parallel Profile Visualization: ParaProf

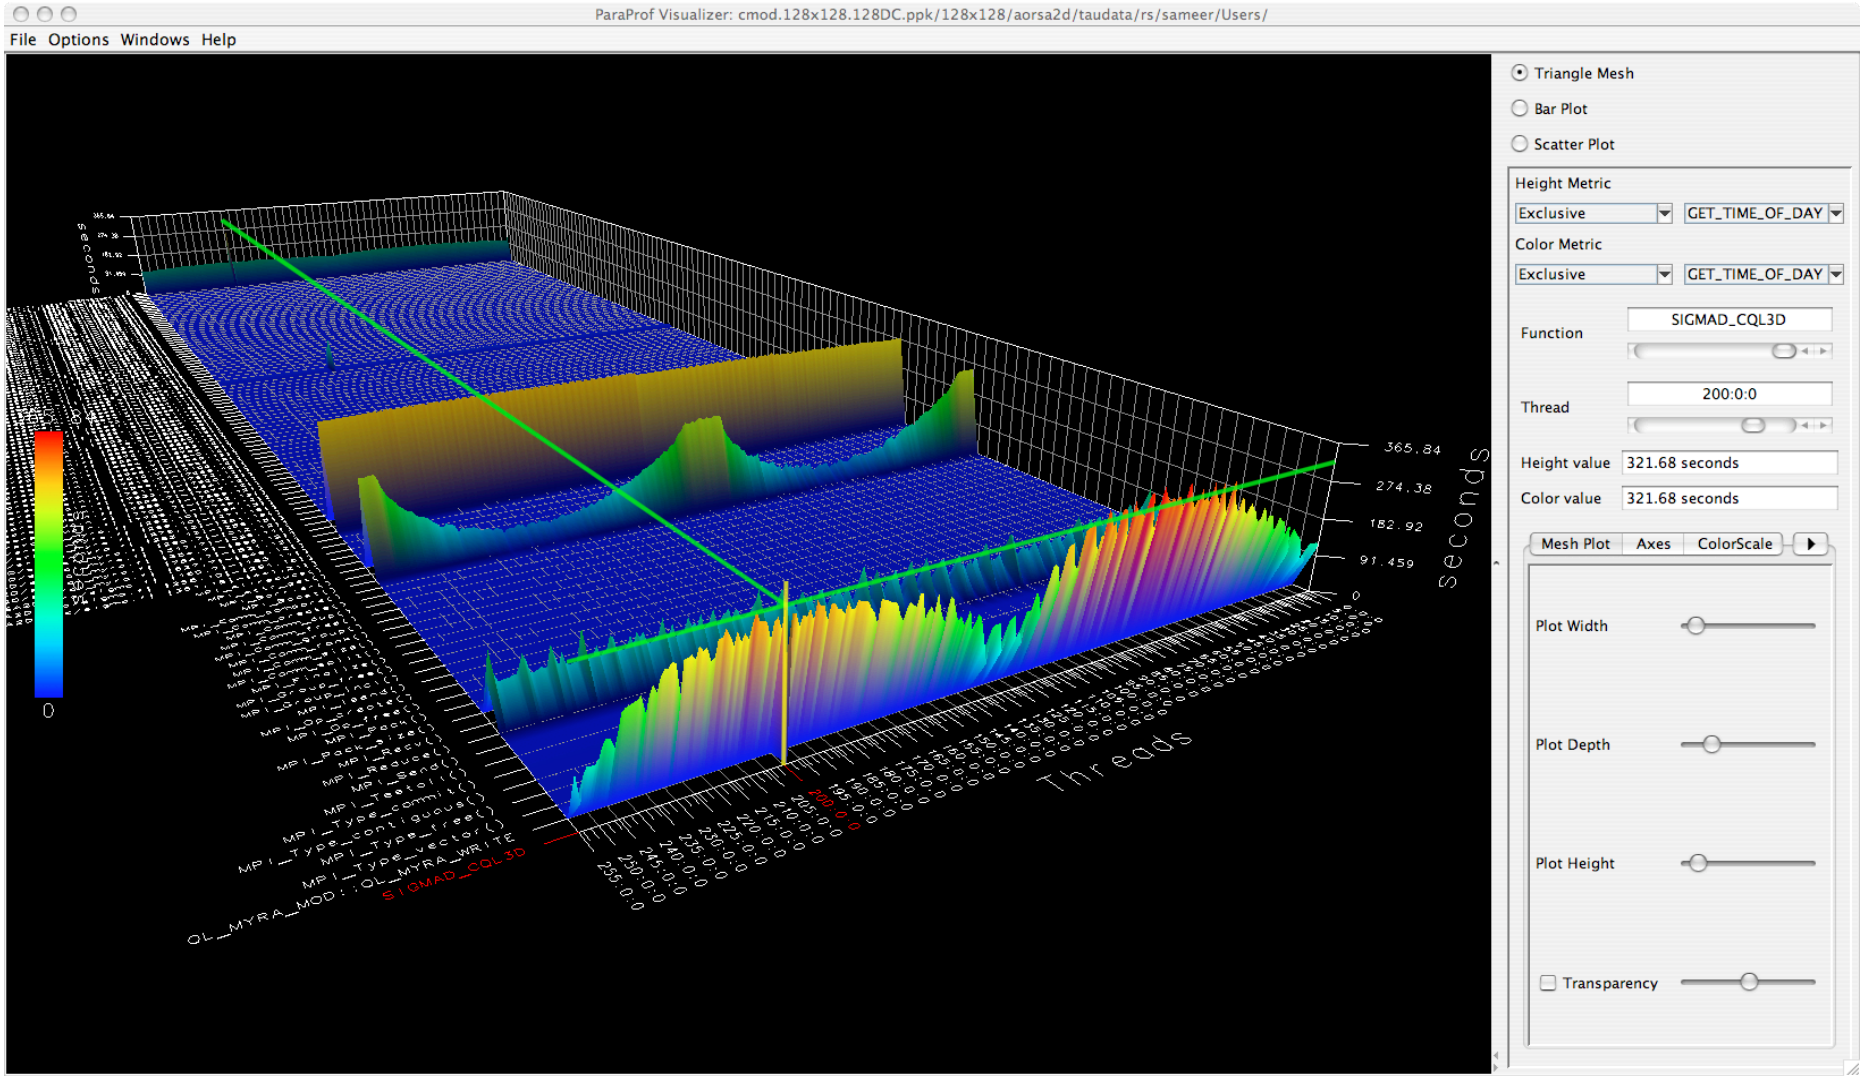

Parallel Profile Visualization: ParaProf

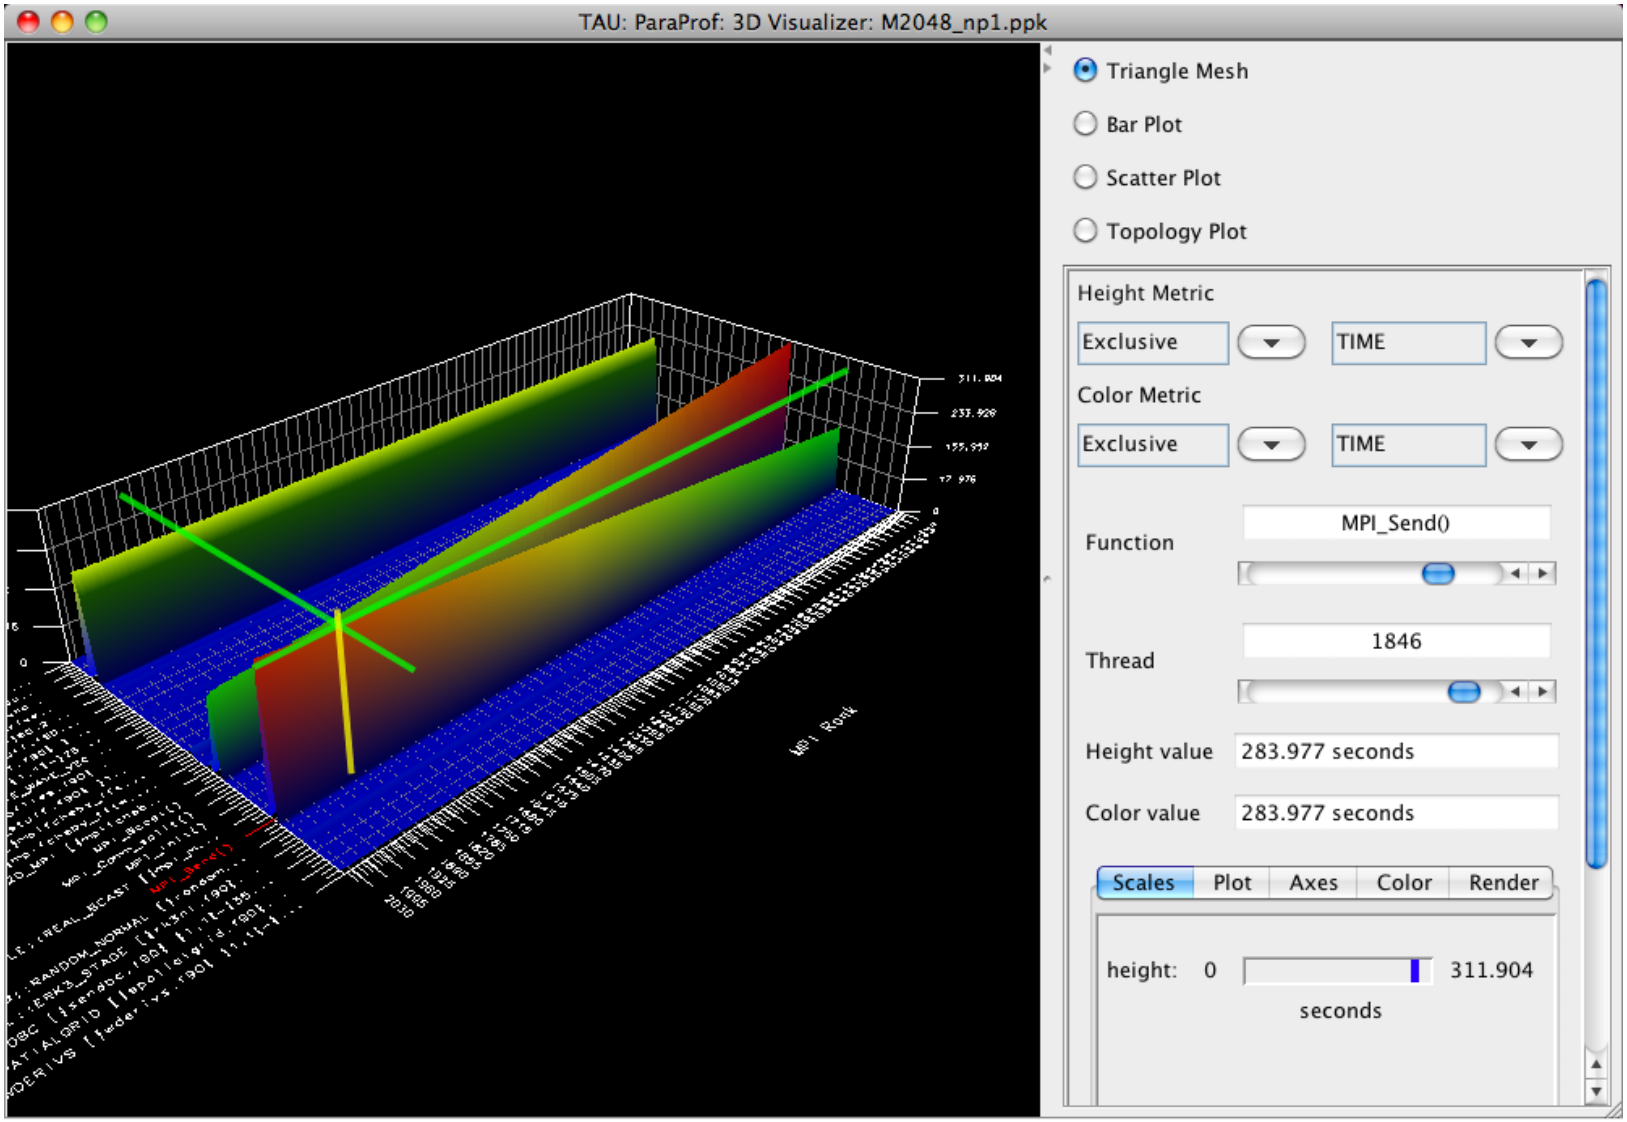

ParaProf: 3D Communication Matrix

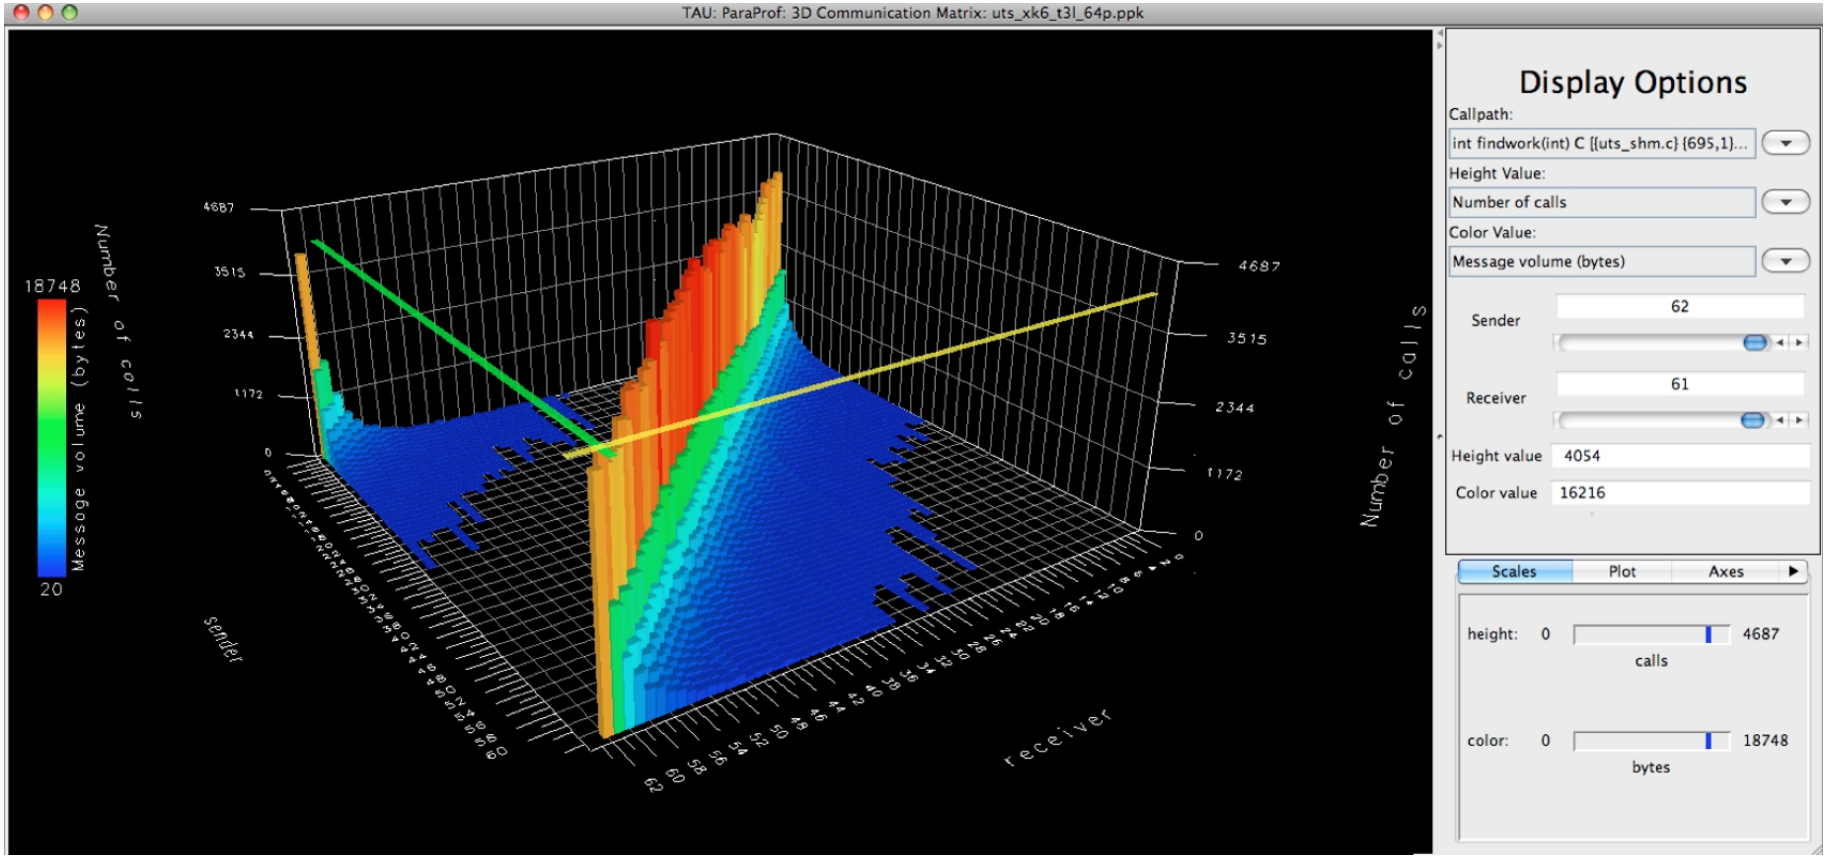

Hands-on: Profile report exploration

- The Tutorial 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
% ls
periscope-1.5                                scorep_bt-mz_B_4x4_sum
README                                       scorep_bt-mz_B_4x4_sum+mets
run.out                                      scorep_bt-mz_B_4x4_trace
scorep-20120913_1740_557443655223384
```

- 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
```

ParaProf: Manager Window: scout.cubex

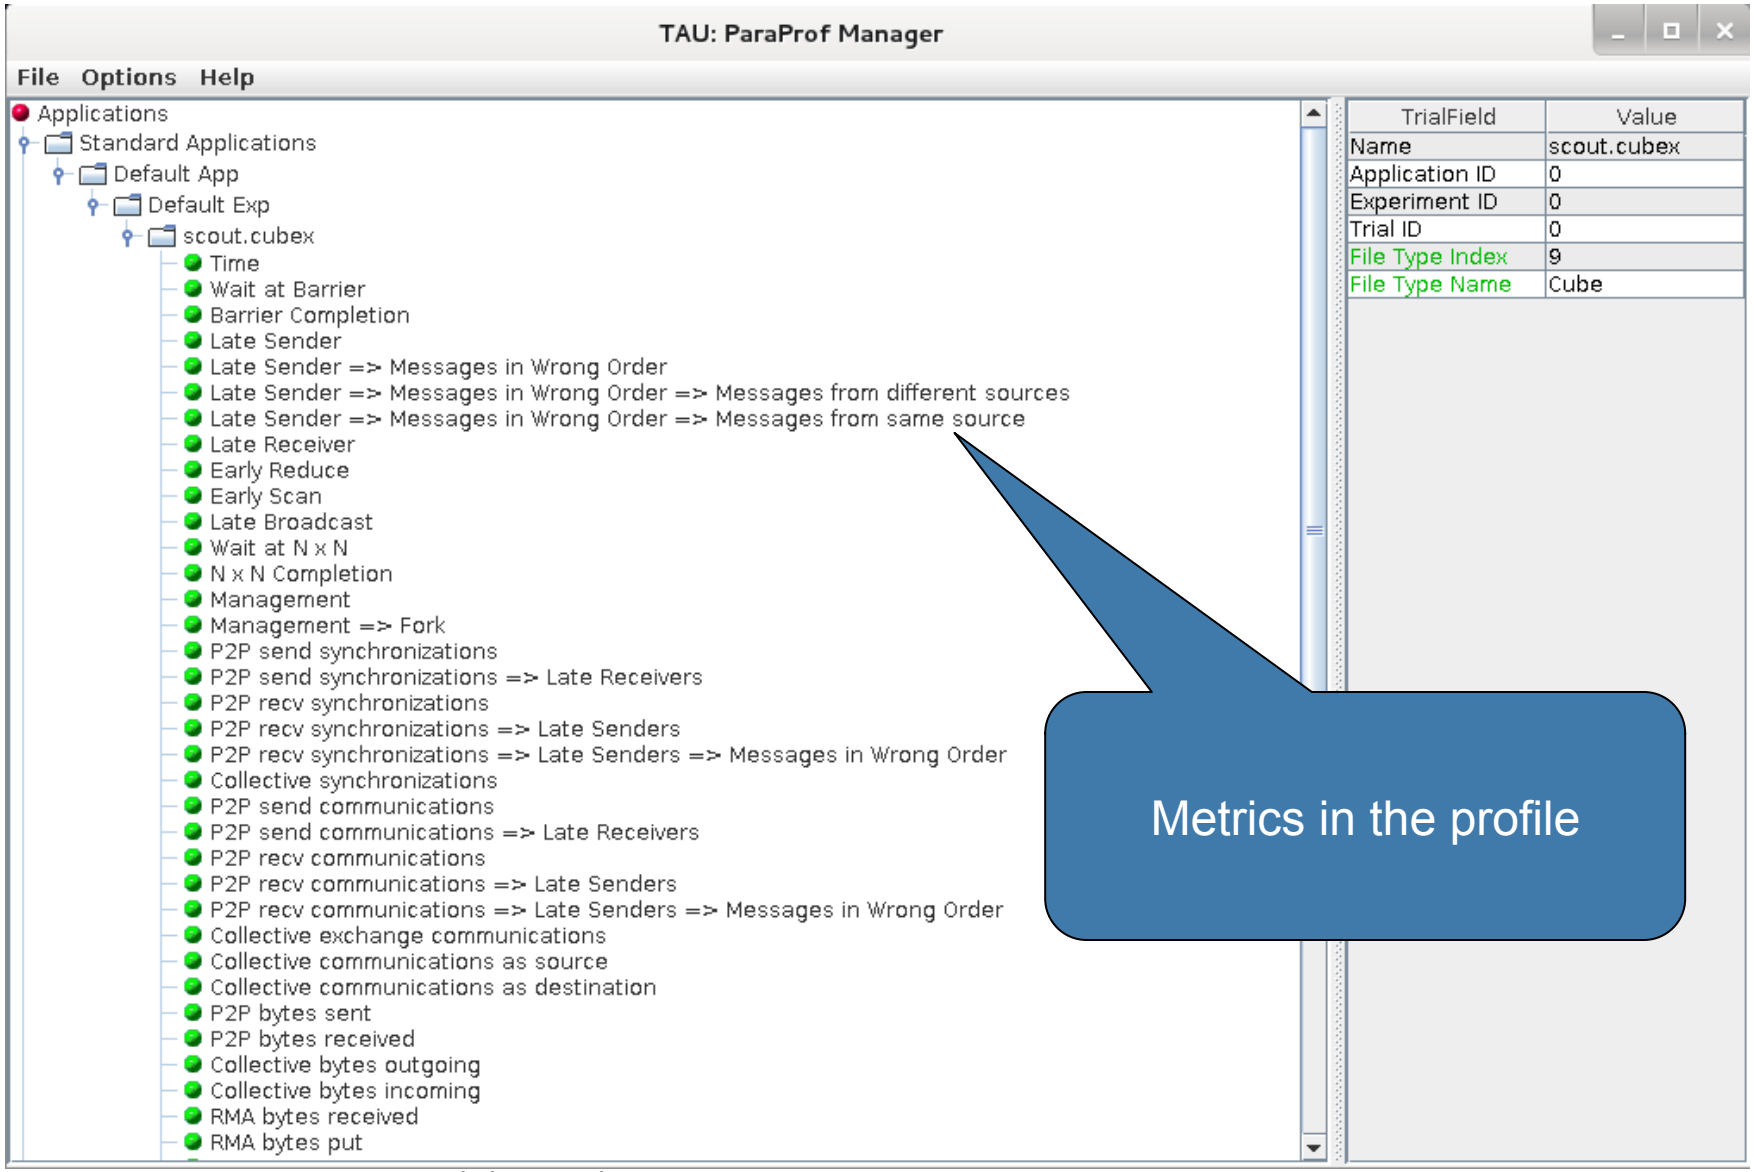

ParaProf: Main window

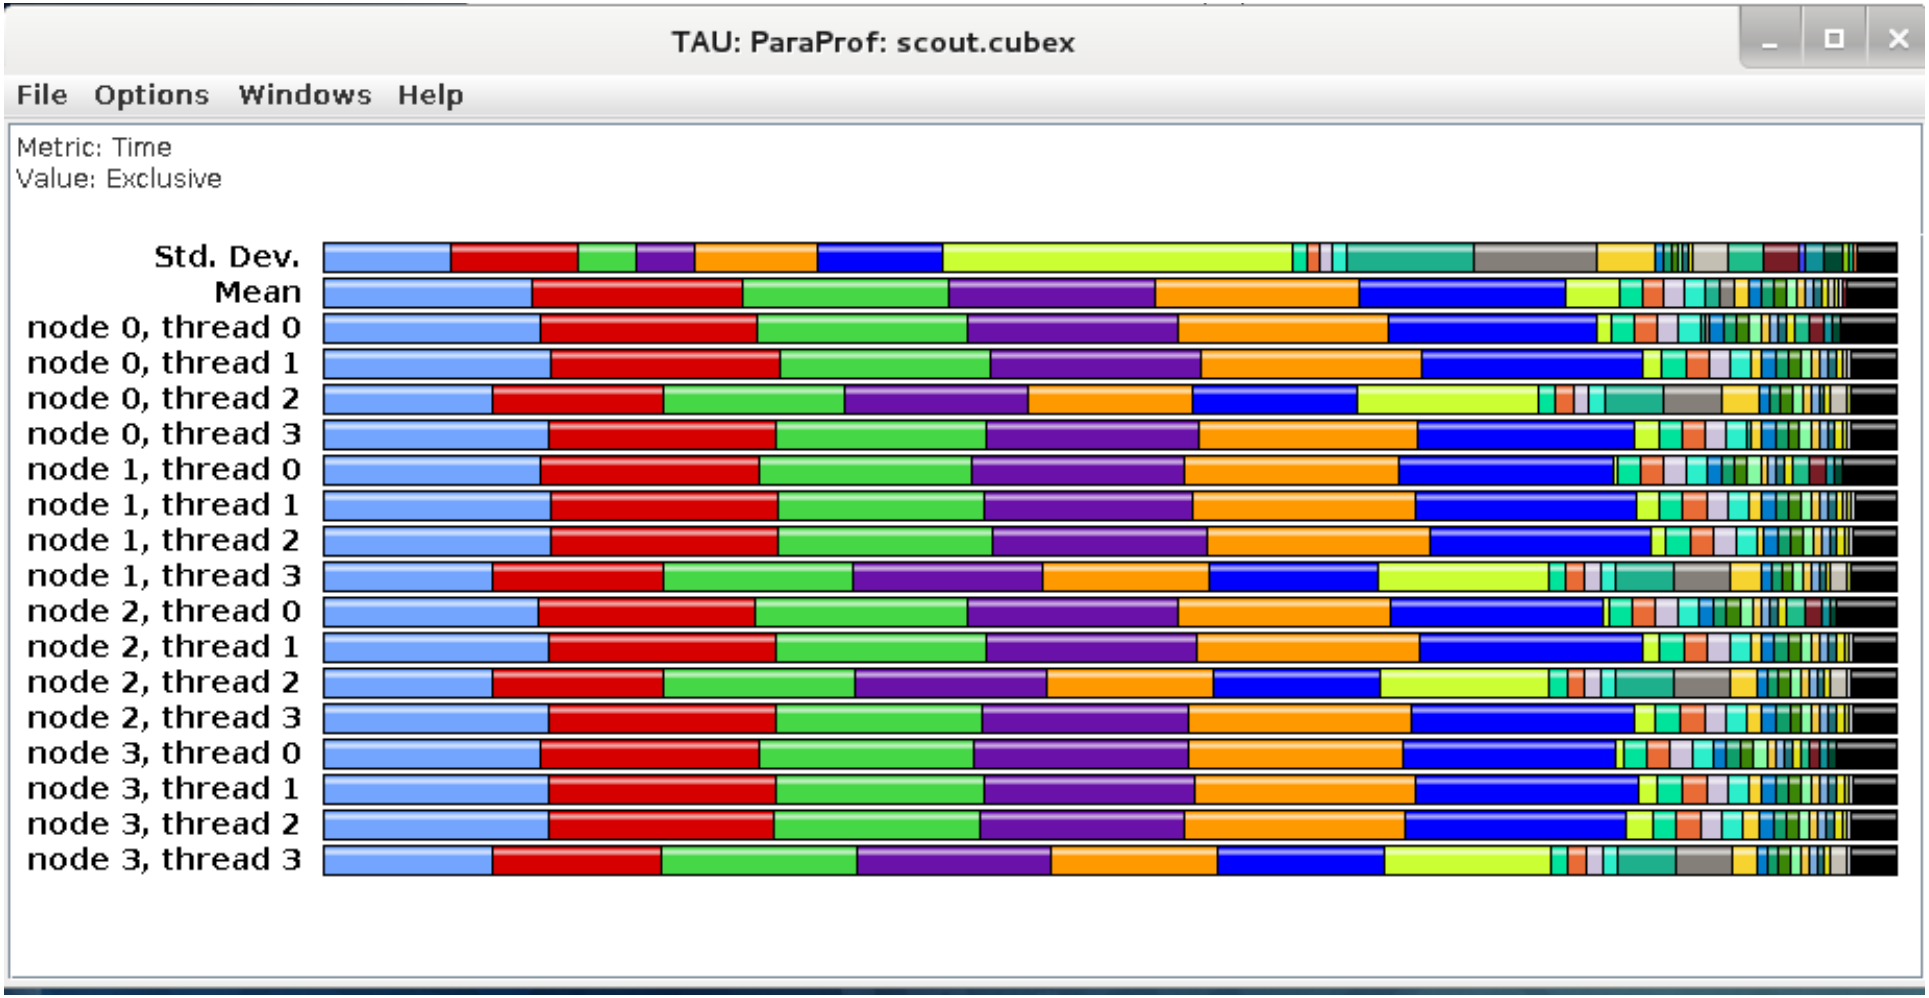

ParaProf: Options

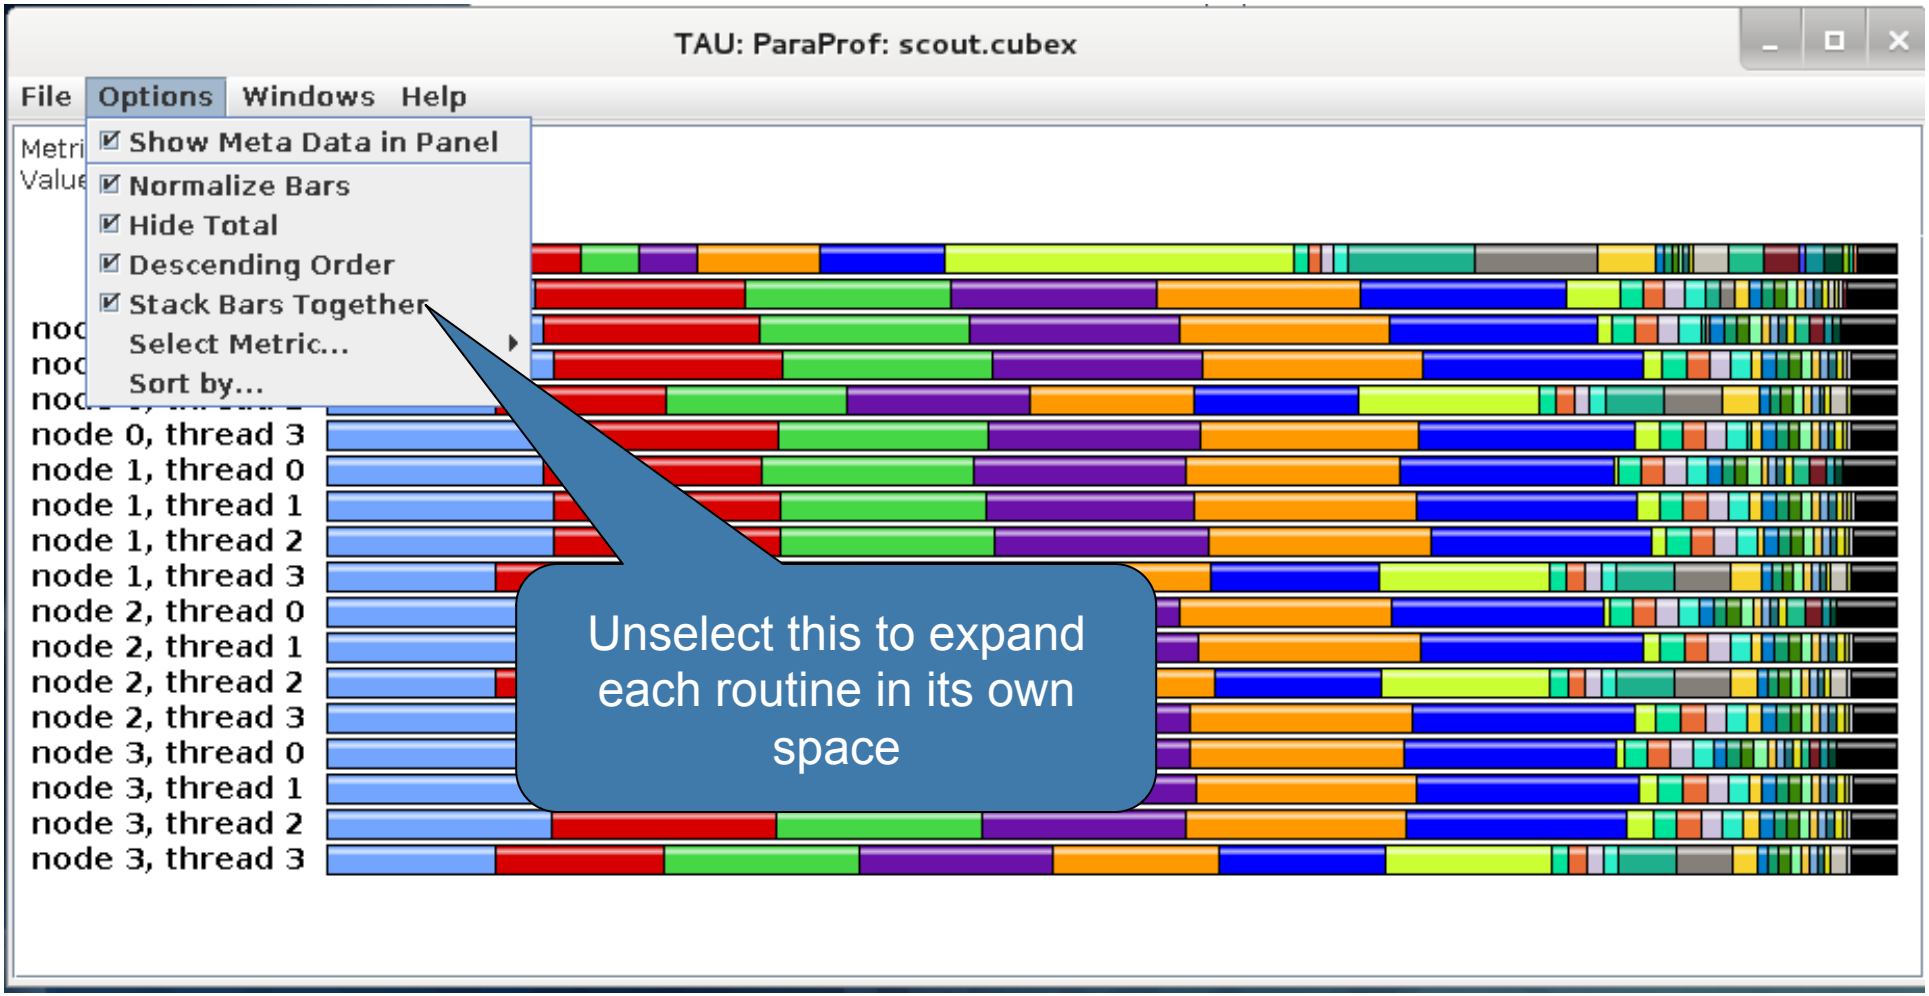

ParaProf:

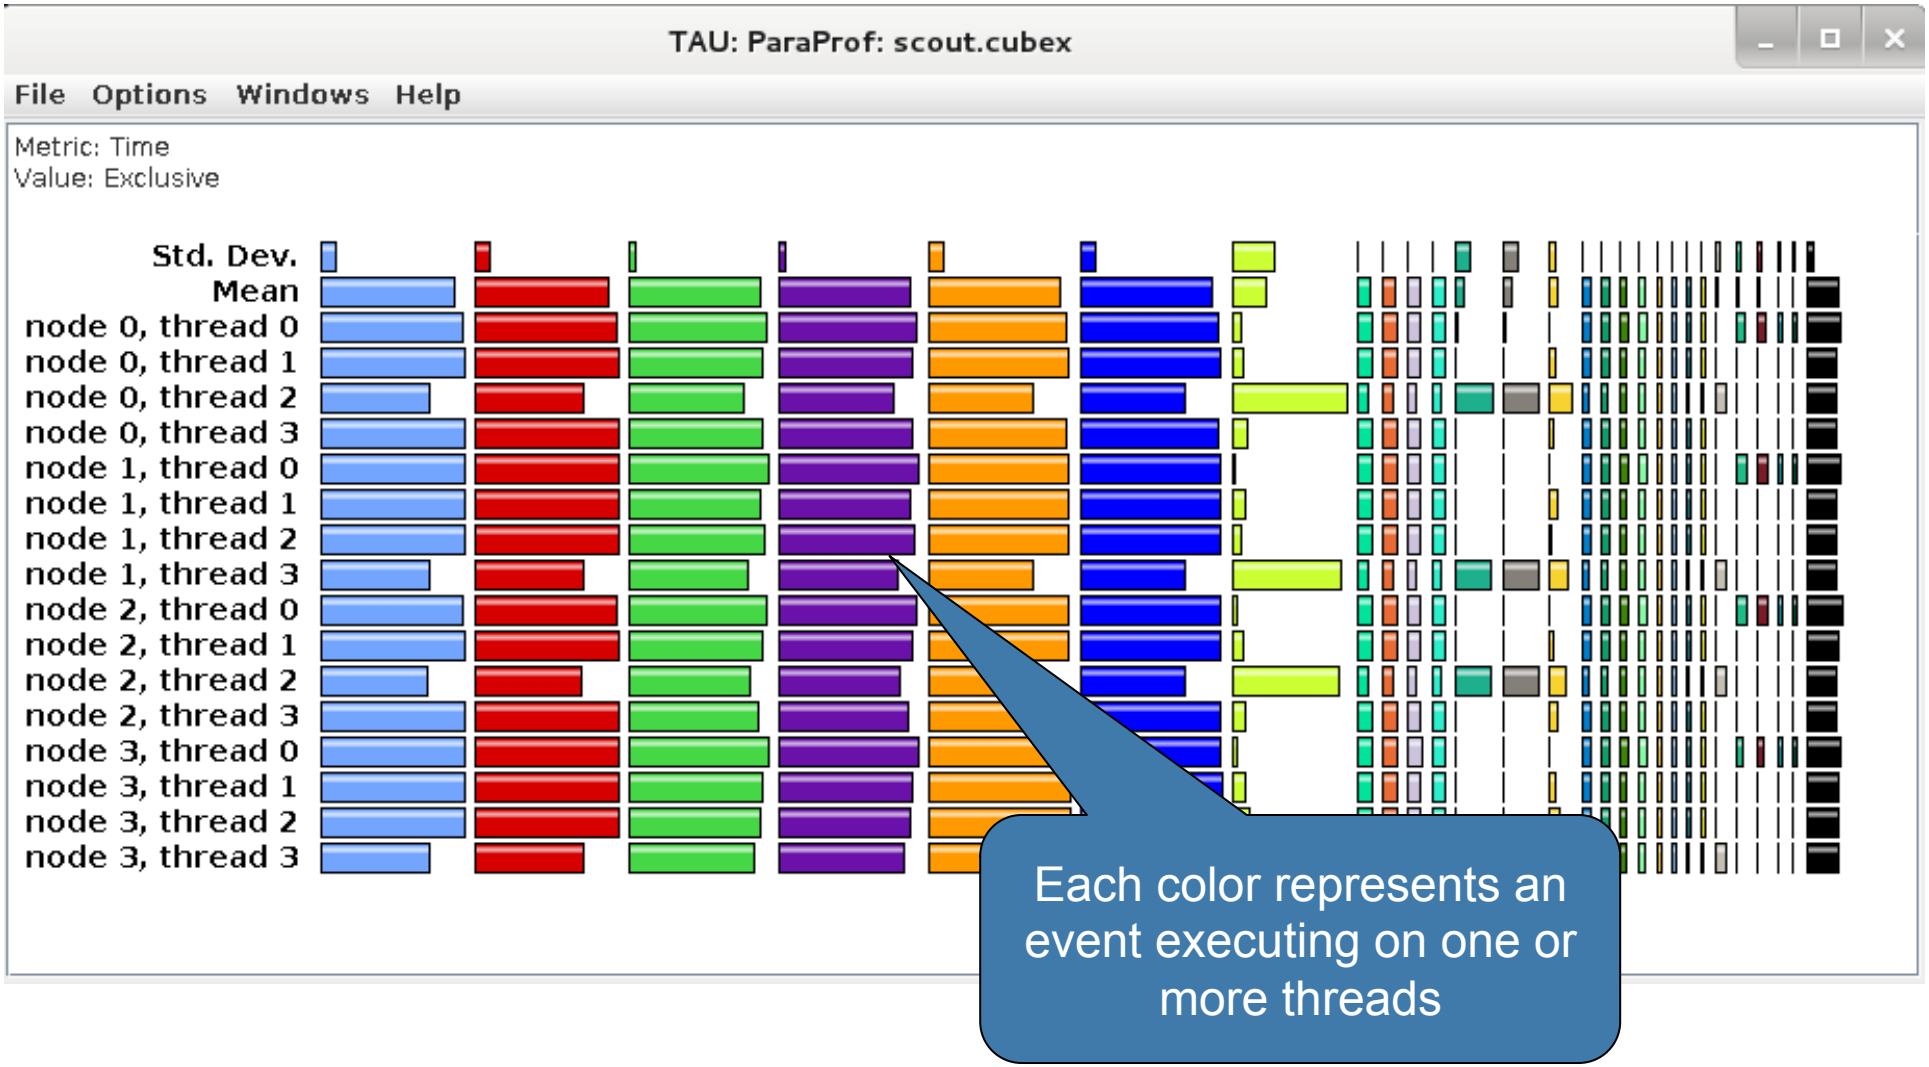
TAU: ParaProf: scout.cubex
File Options Windows Help
Metric: Time
Value: Exclusive
Std. Dev.
Mean
node 0, thread 0
node 0, thread 1
node 0, thread 2
node 0, thread 3
node 1, thread 0
node 1, thread 1
node 1, thread 2
node 1, thread 3
node 2, thread 0
node 2, thread 1
node 2, thread 2
node 2, thread 3
node 3, thread 0
node 3, thread 1
node 3, thread 2
node 3, thread 3
Each color represents an event executing on one or more threads

ParaProf: Windows

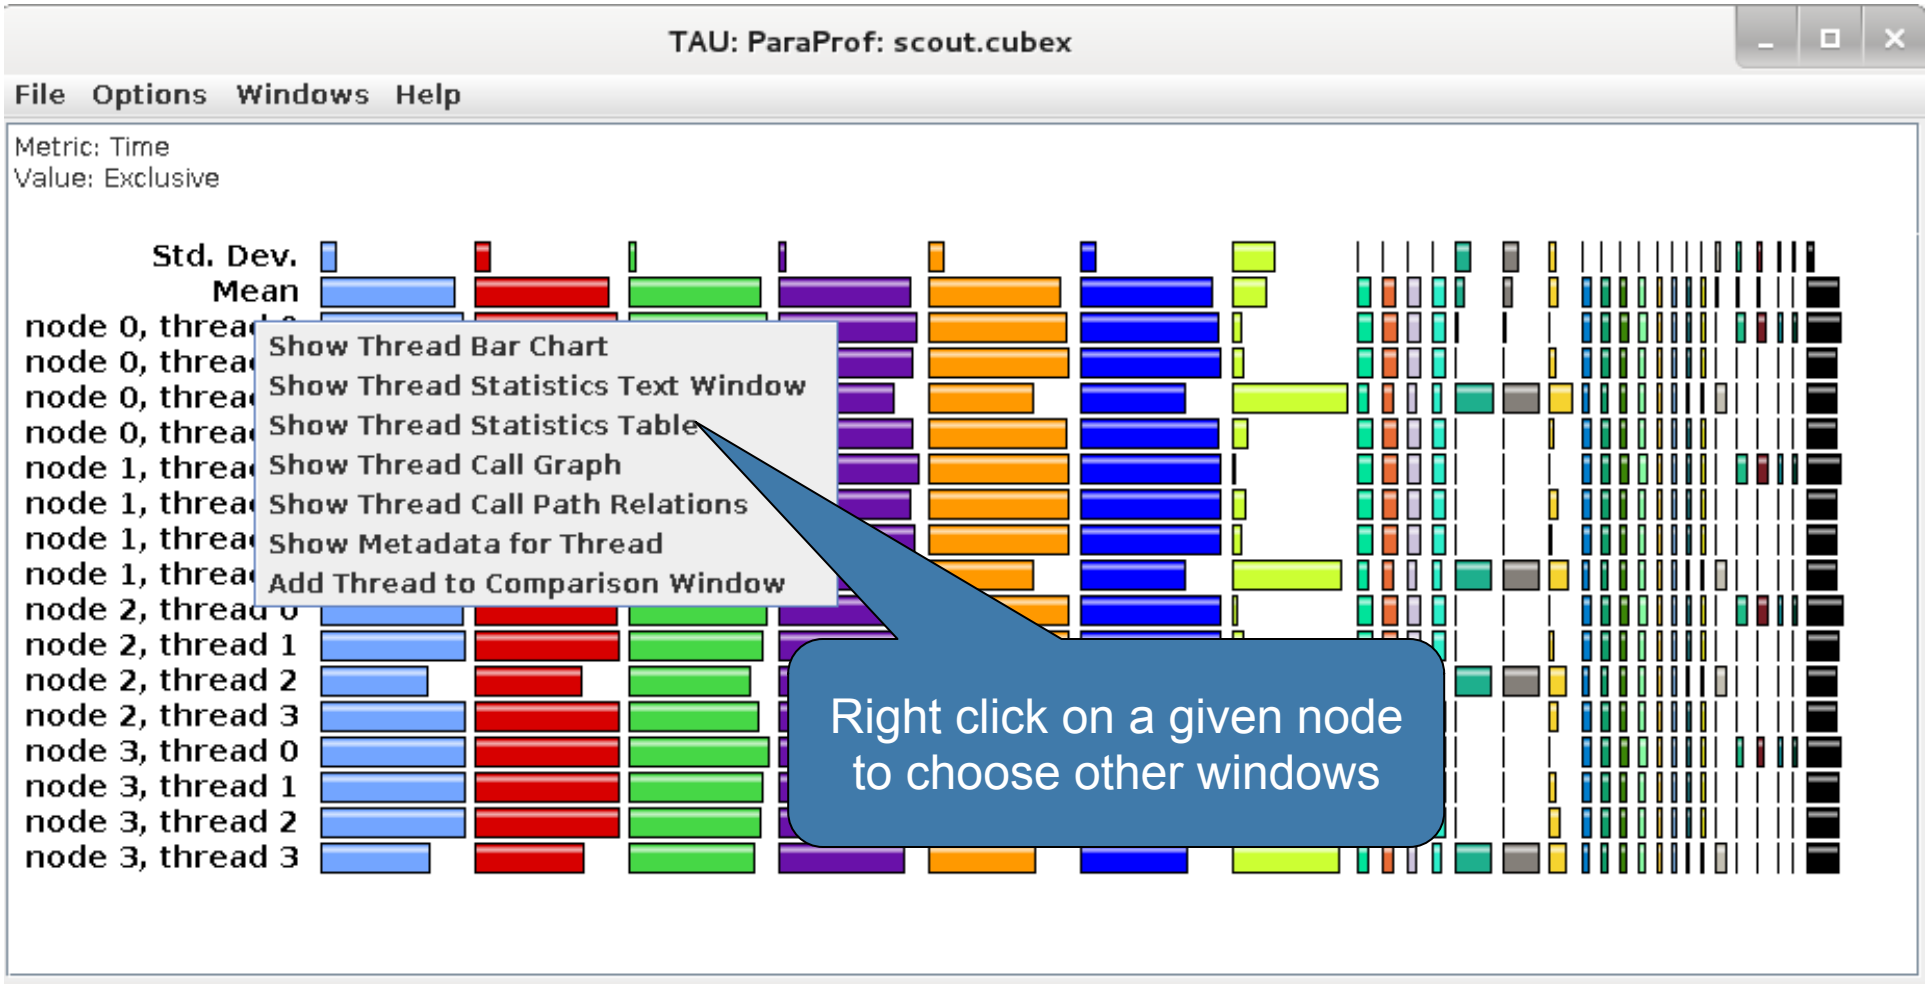

ParaProf: Thread Statistics Table

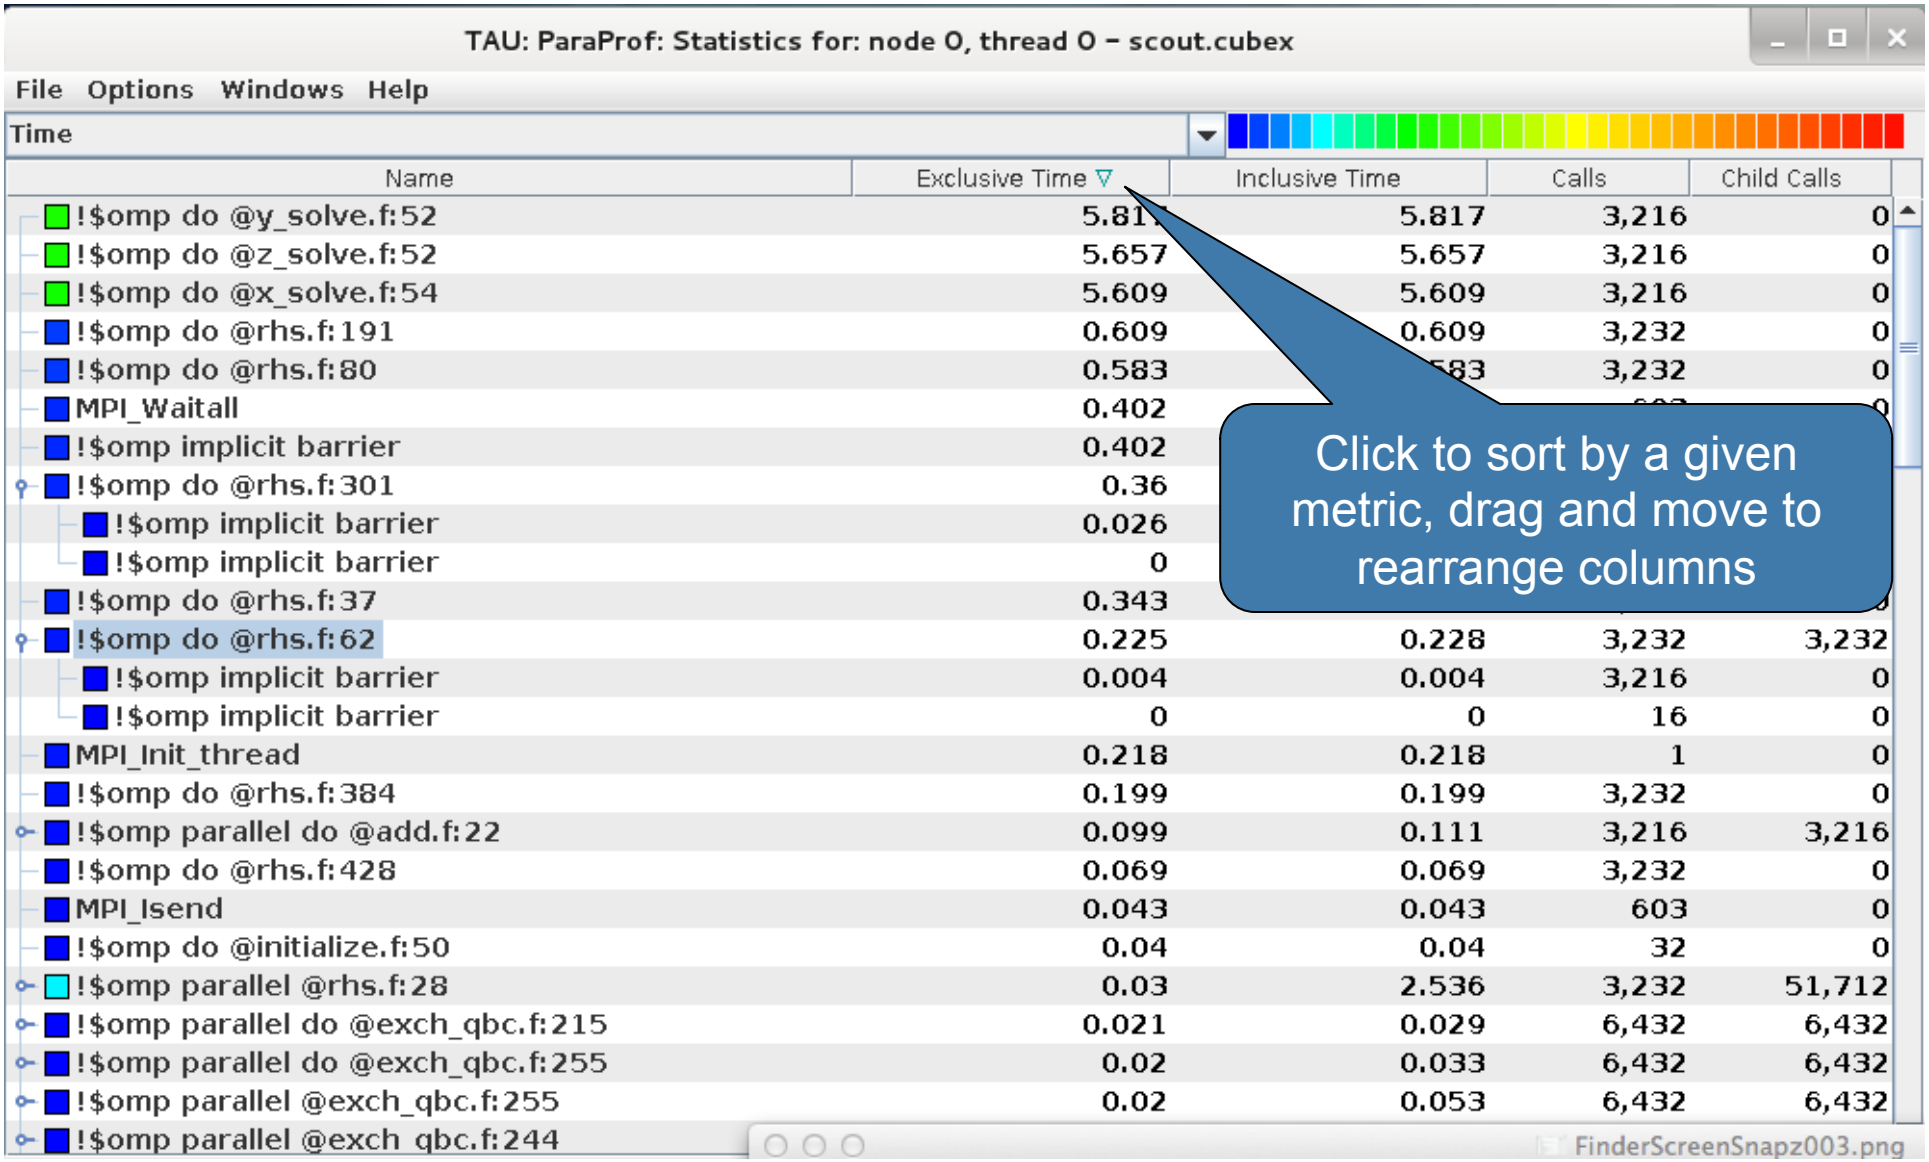

Example: Score-P with TAU (LU NPB)

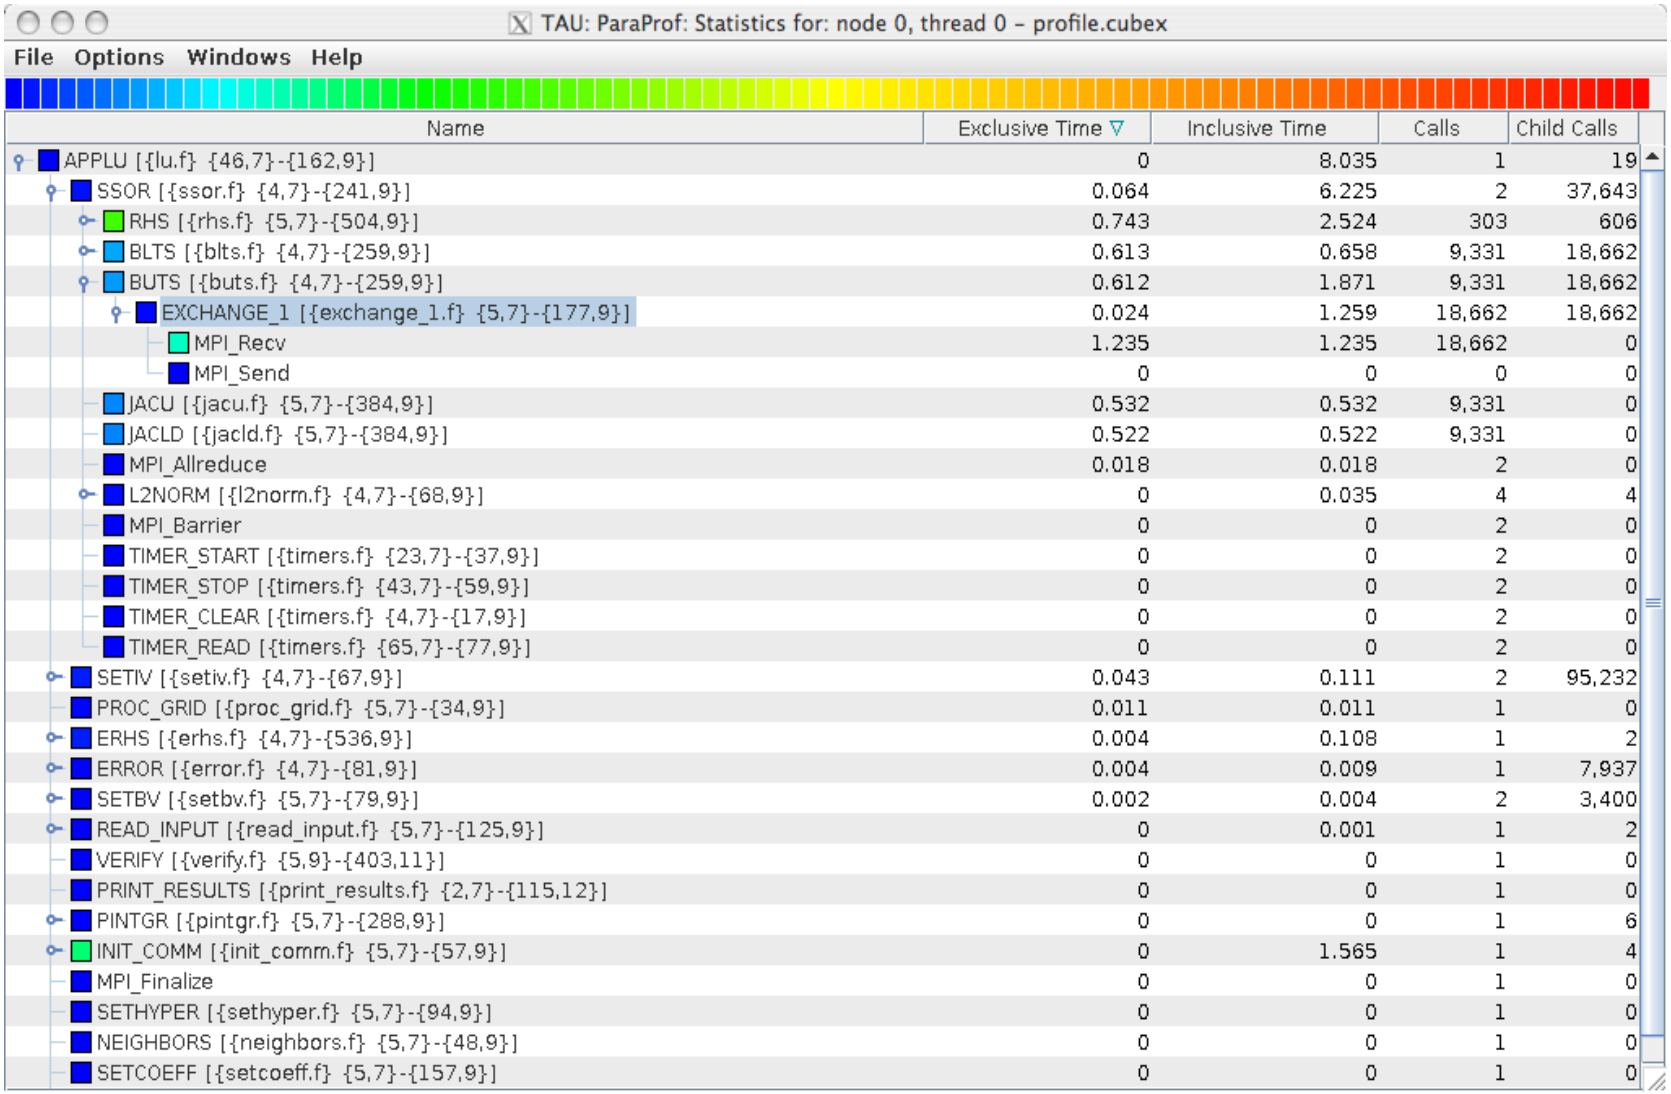

Name	Exclusive Time	Inclusive Time	Calls	Child Calls
APPLU [{lu.f} {46,7}-{162,9}]	0	8.035	1	19
SSOR [{ssor.f} {4,7}-{241,9}]	0.064	6.225	2	37,643
RHS [{rhs.f} {5,7}-{504,9}]	0.743	2.524	303	606
BLTS [{blts.f} {4,7}-{259,9}]	0.613	0.658	9,331	18,662
BUTS [{buts.f} {4,7}-{259,9}]	0.612	1.871	9,331	18,662
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: Thread Callgraph Window

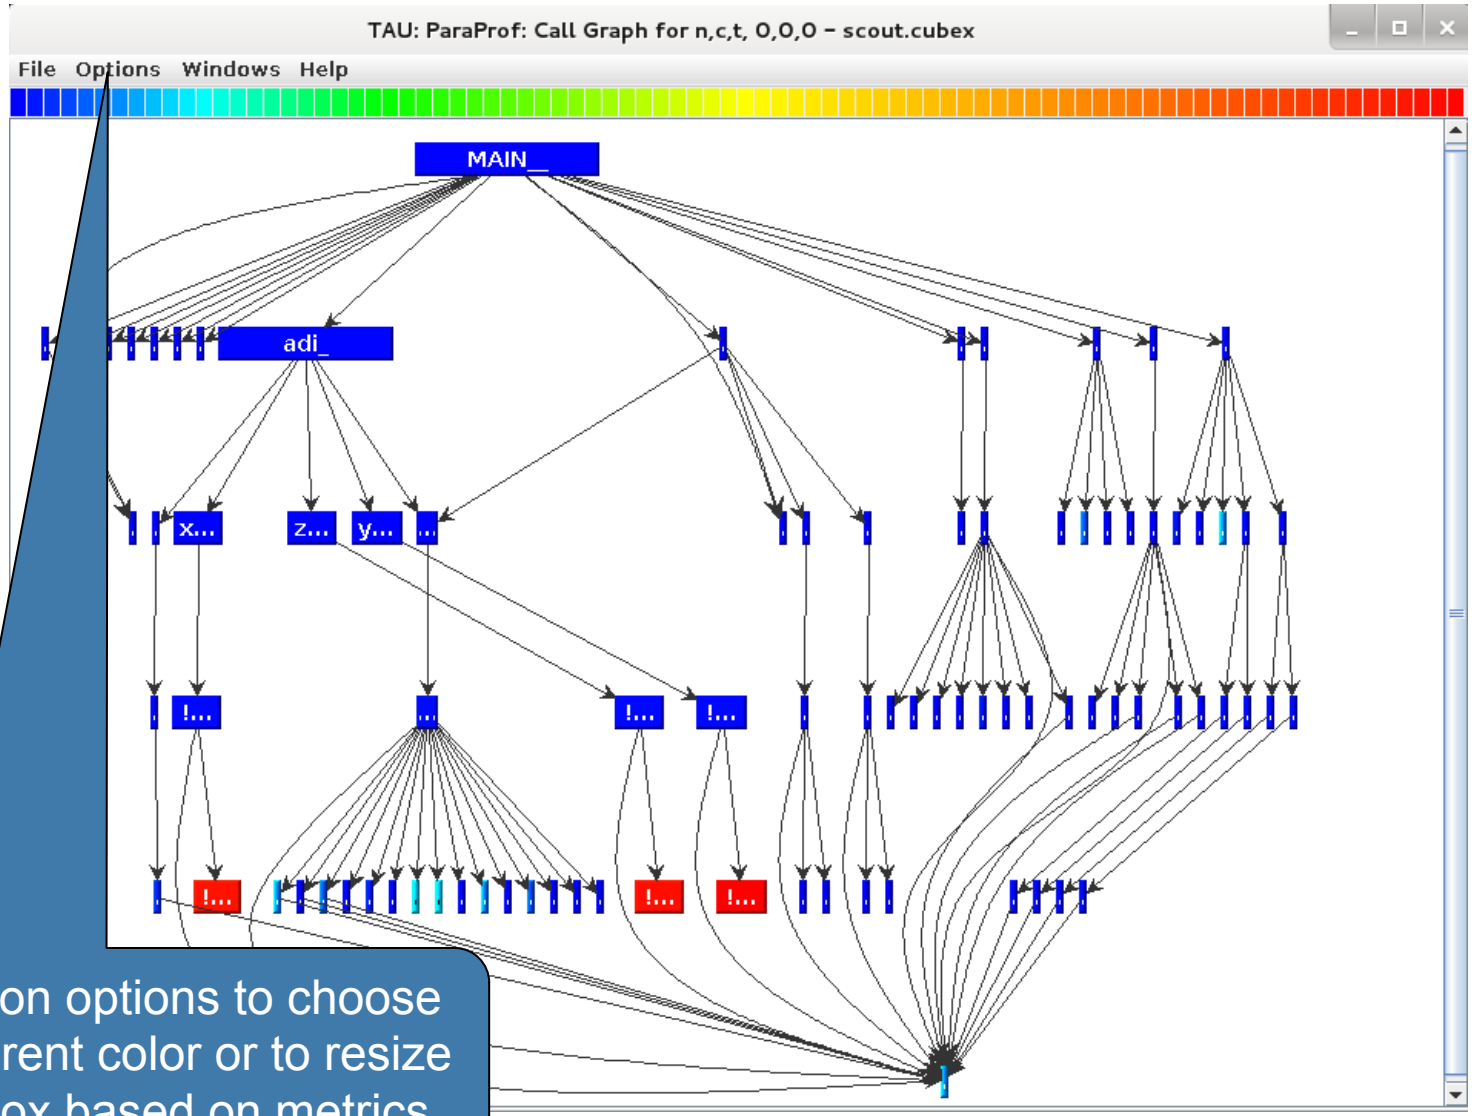

Click on options to choose a different color or to resize the box based on metrics

ParaProf: Callpath Thread Relations Window

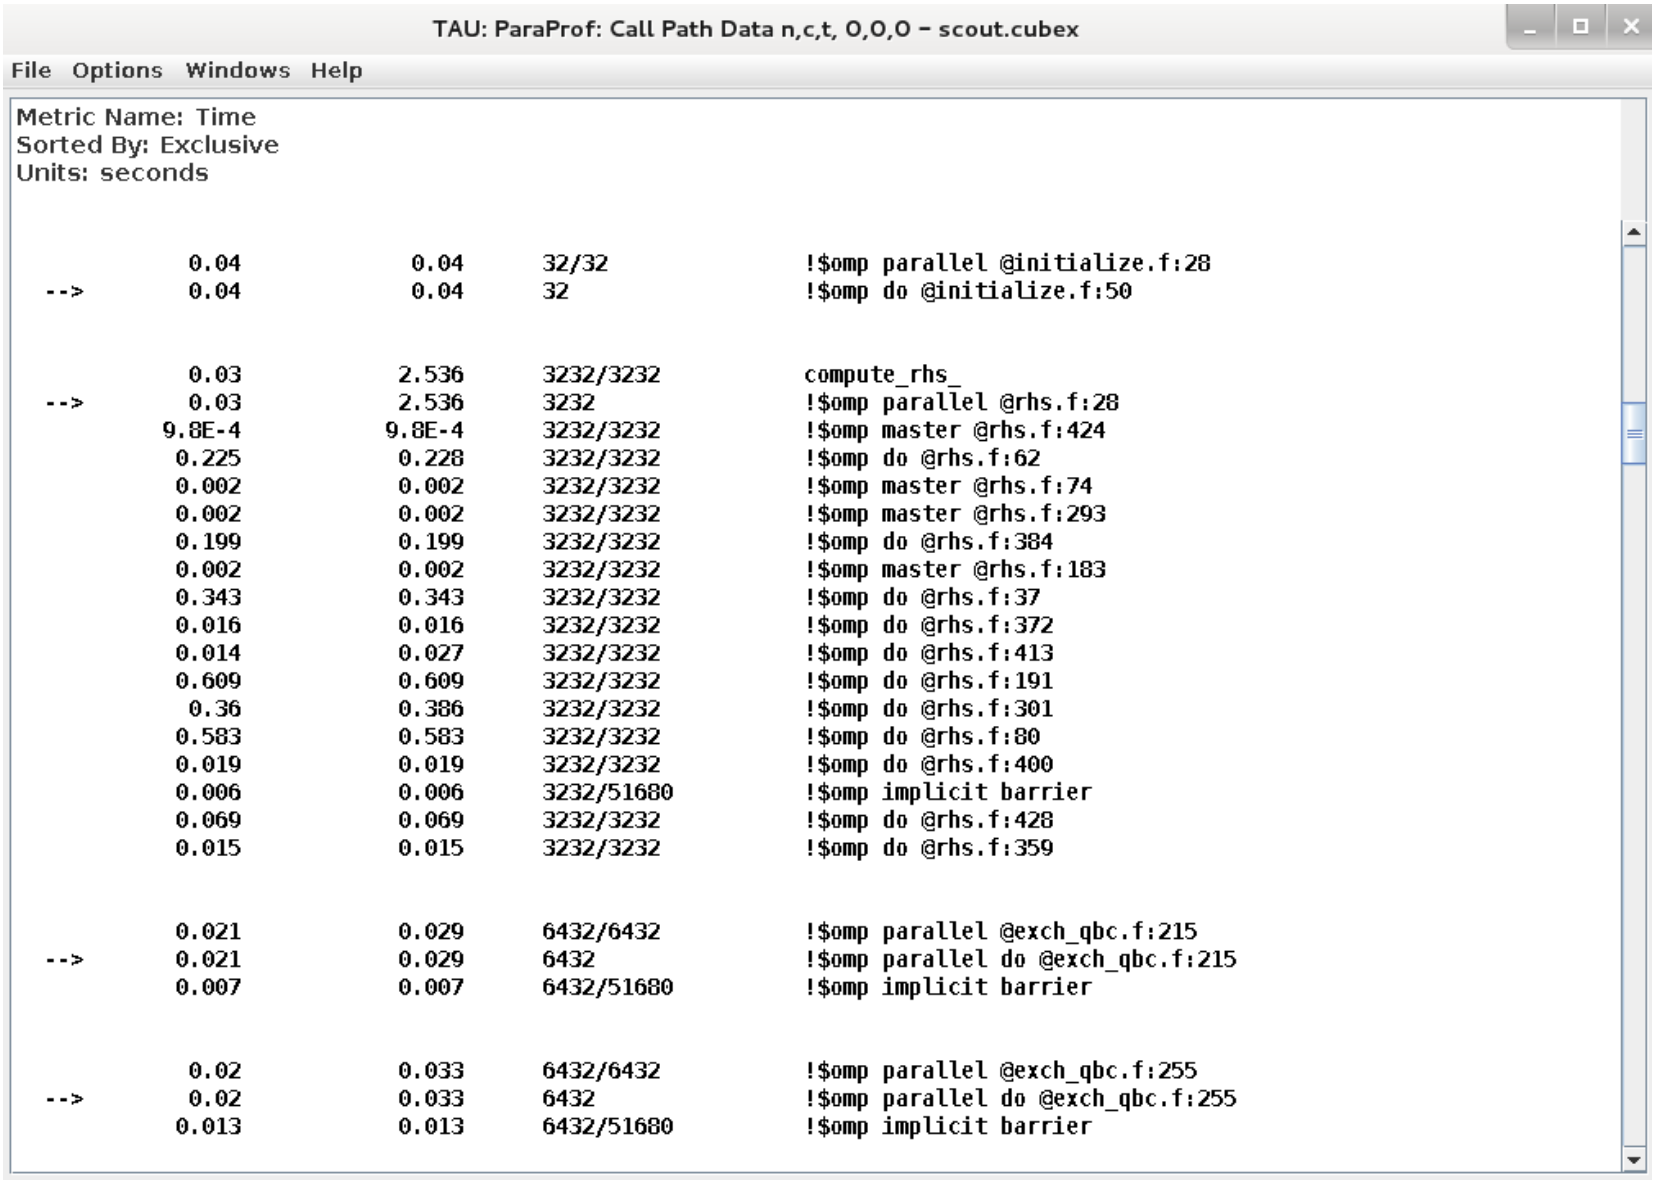

ParaProf:Windows -> 3D Visualization -> Bar Plot

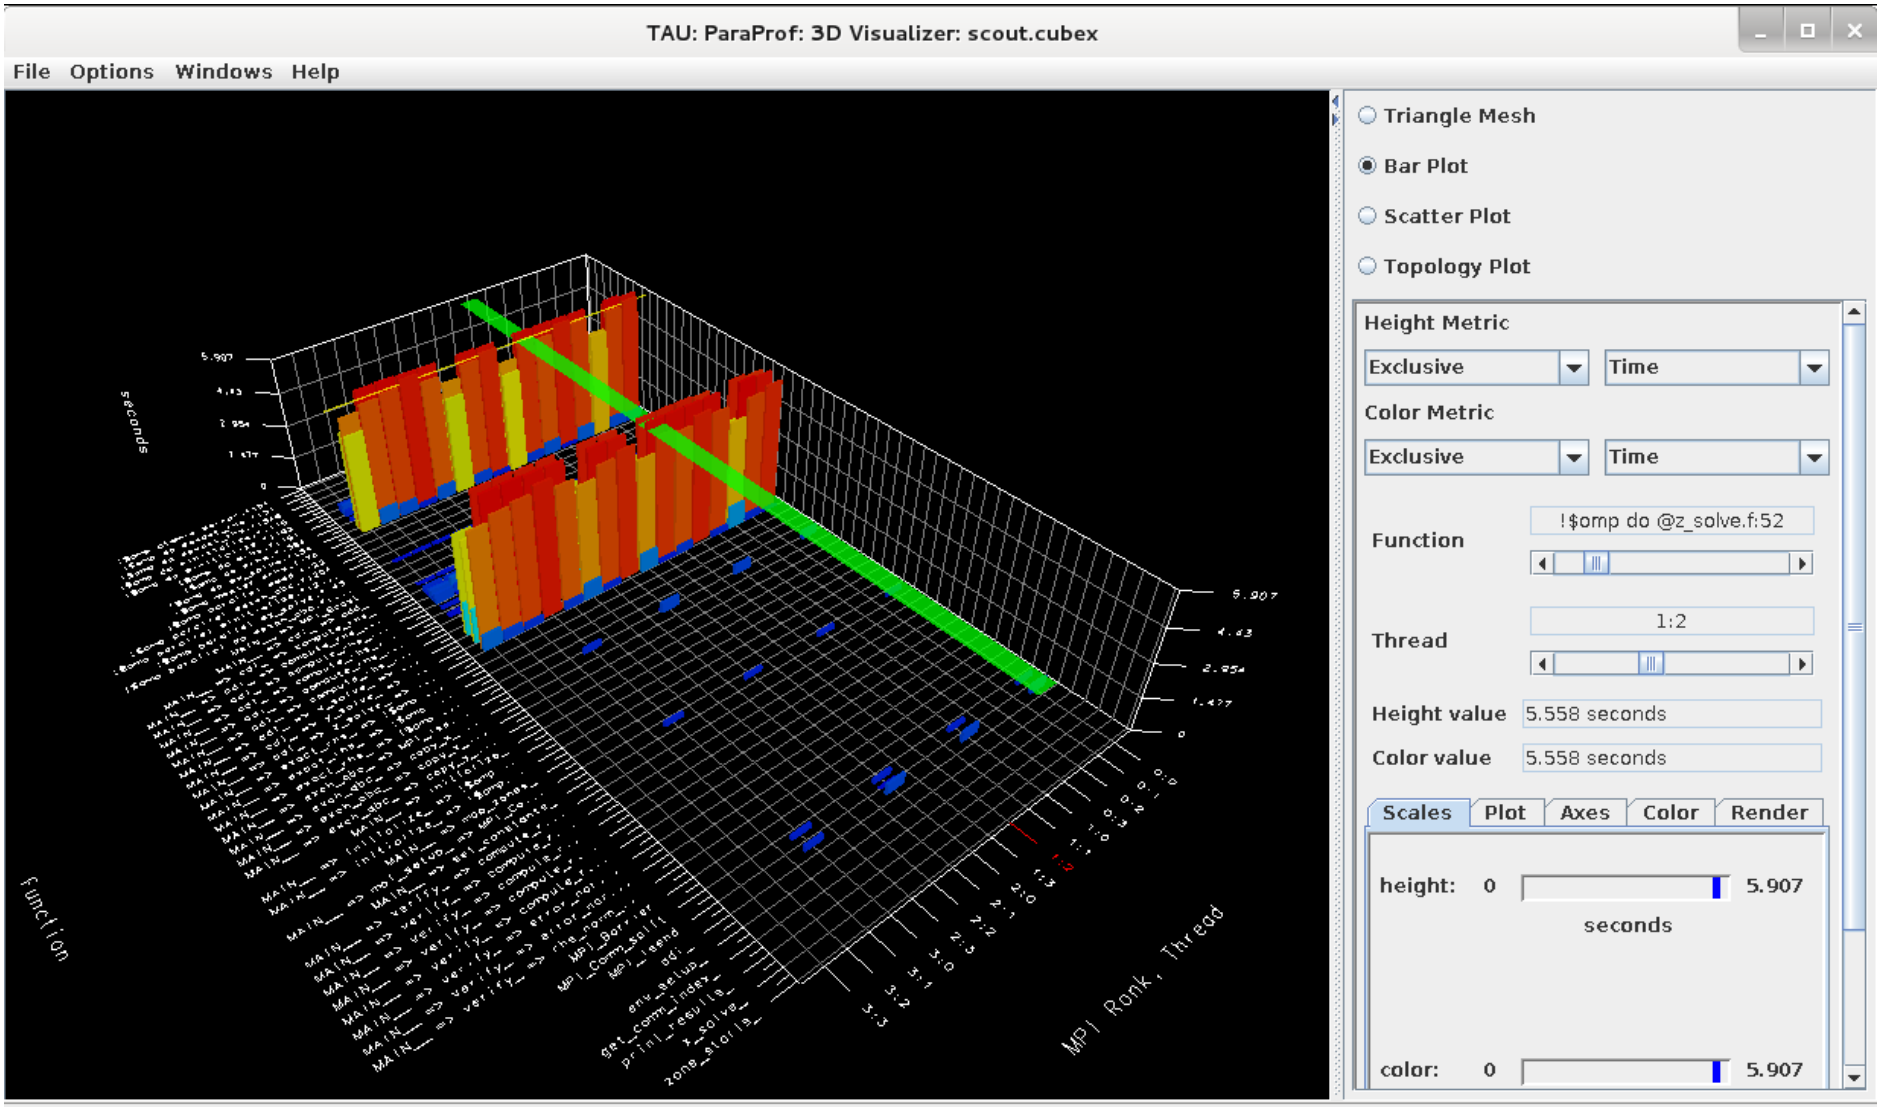

ParaProf: 3D Scatter Plot

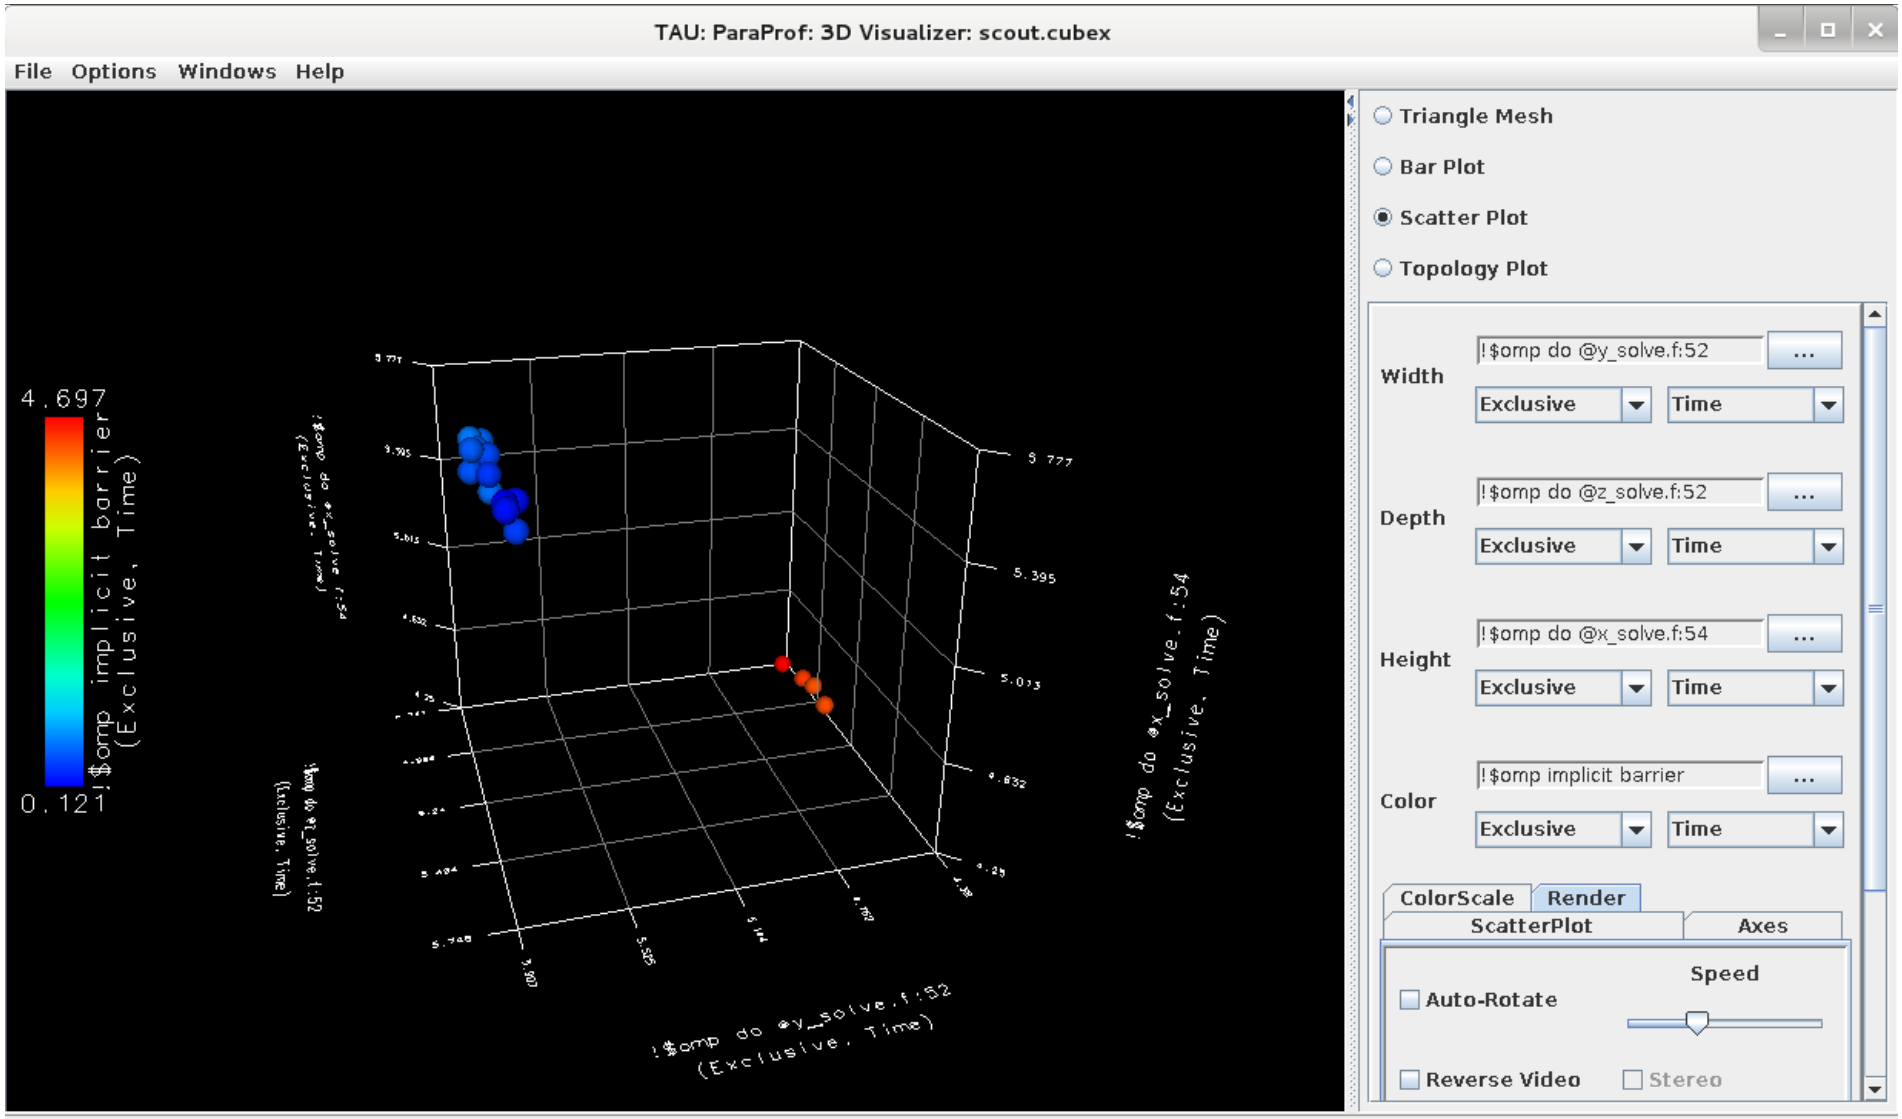

ParaProf: Scatter Plot

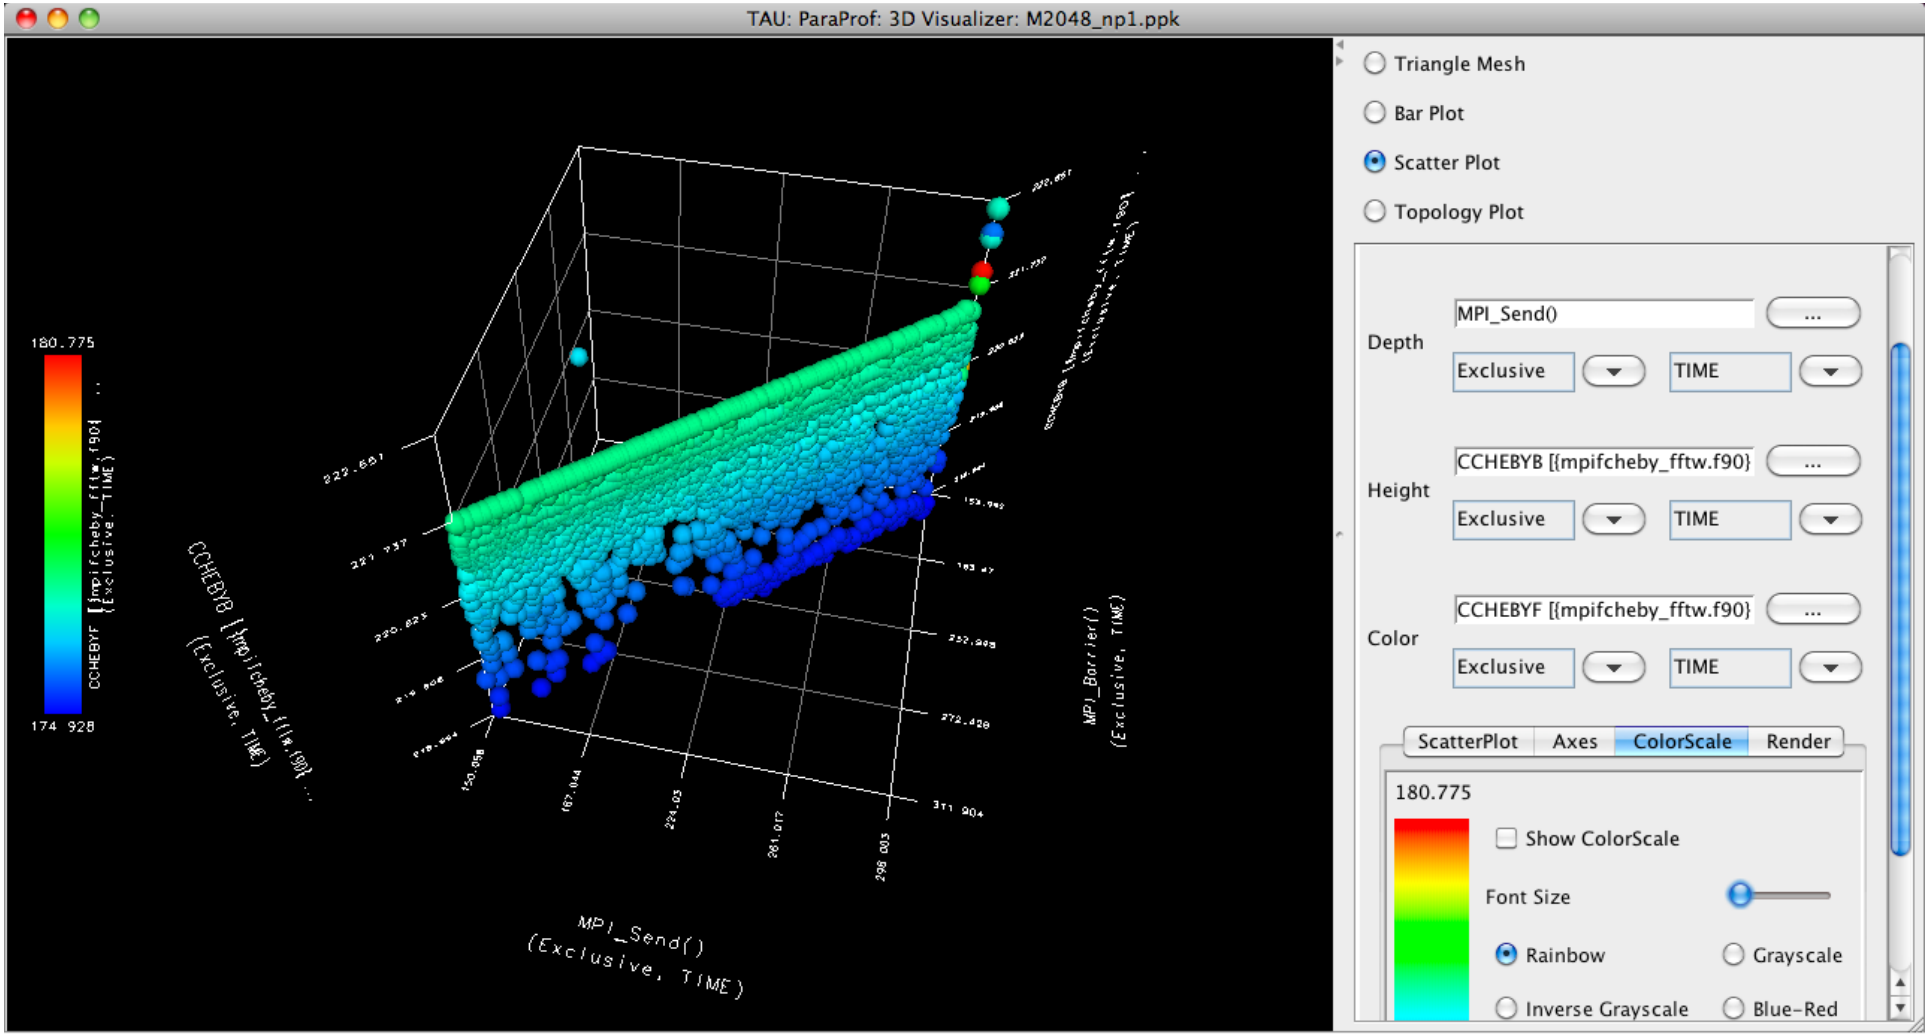

ParaProf: 3D Topology View for a Routine

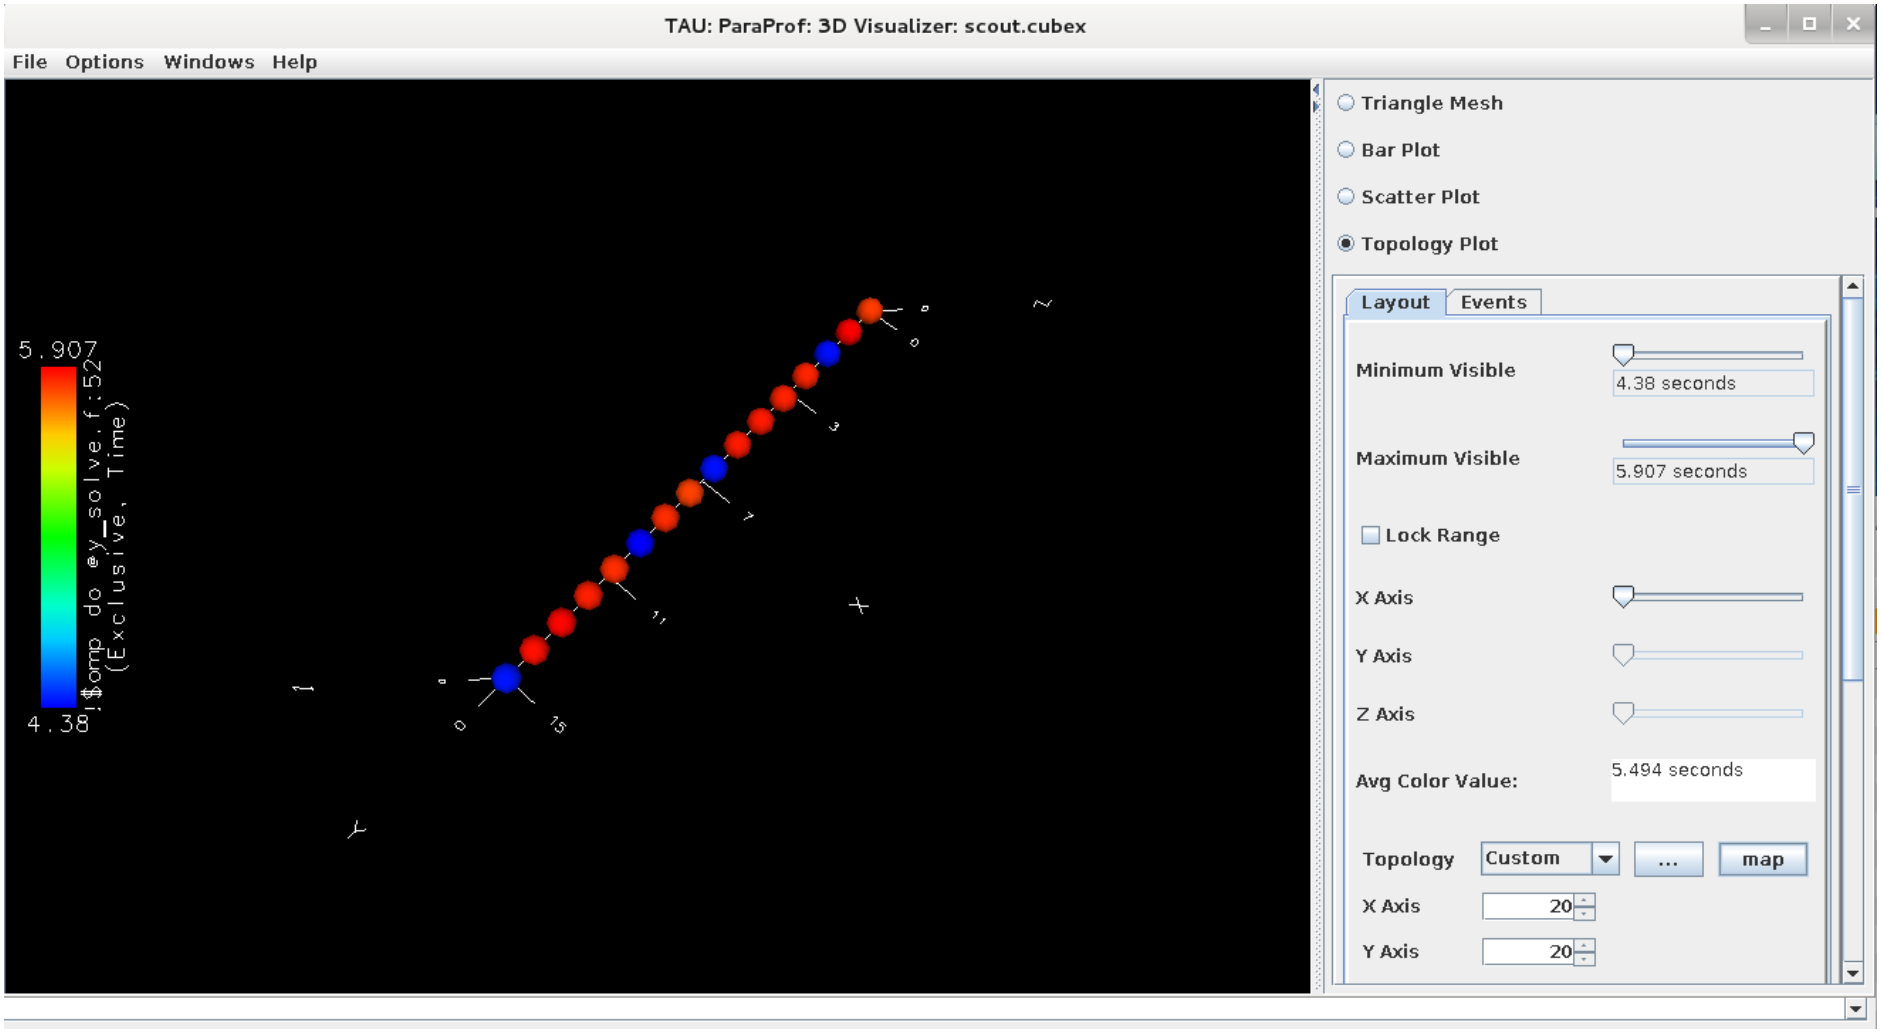

ParaProf: Topology View 3D Torus (IBM BG/P)

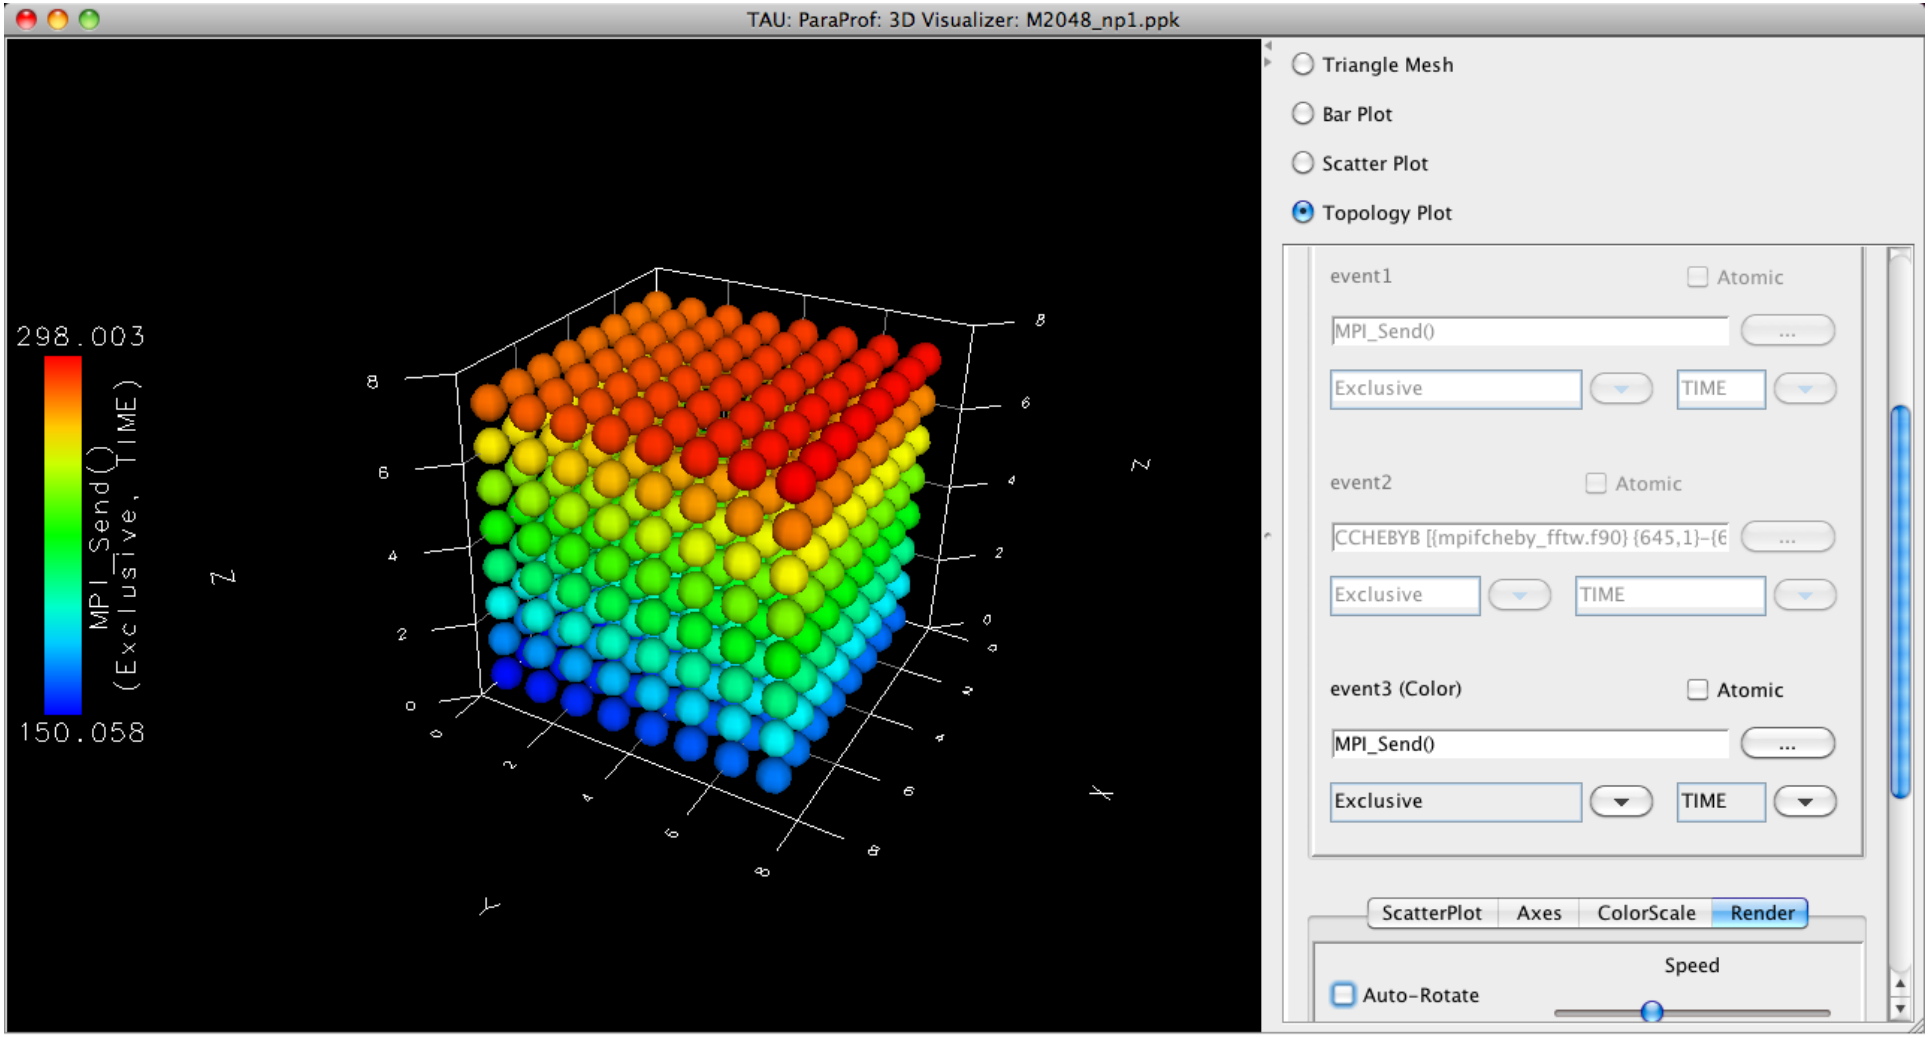

ParaProf:Topology View (6D Torus Coordinates BG/Q)

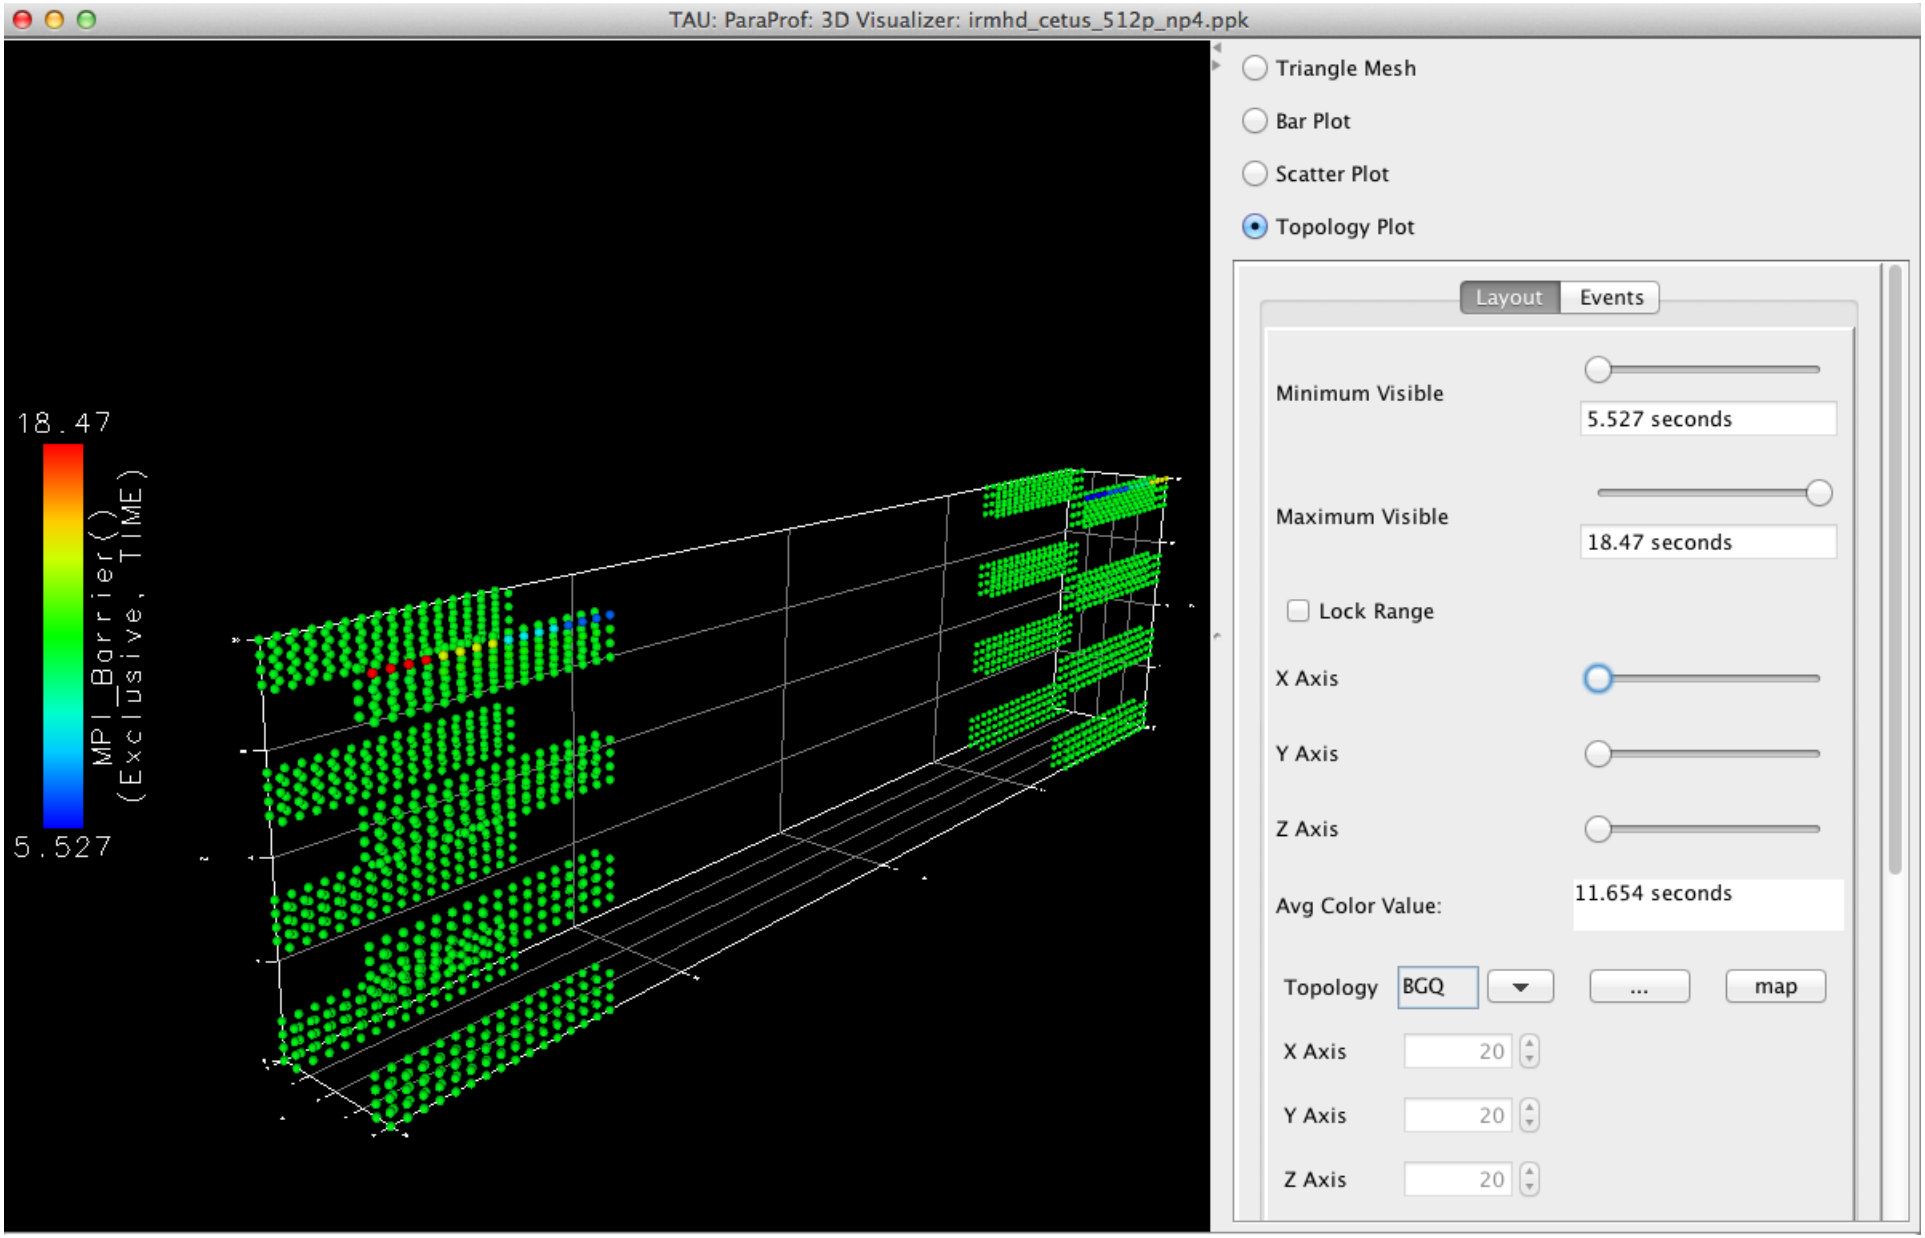

ParaProf: Node View

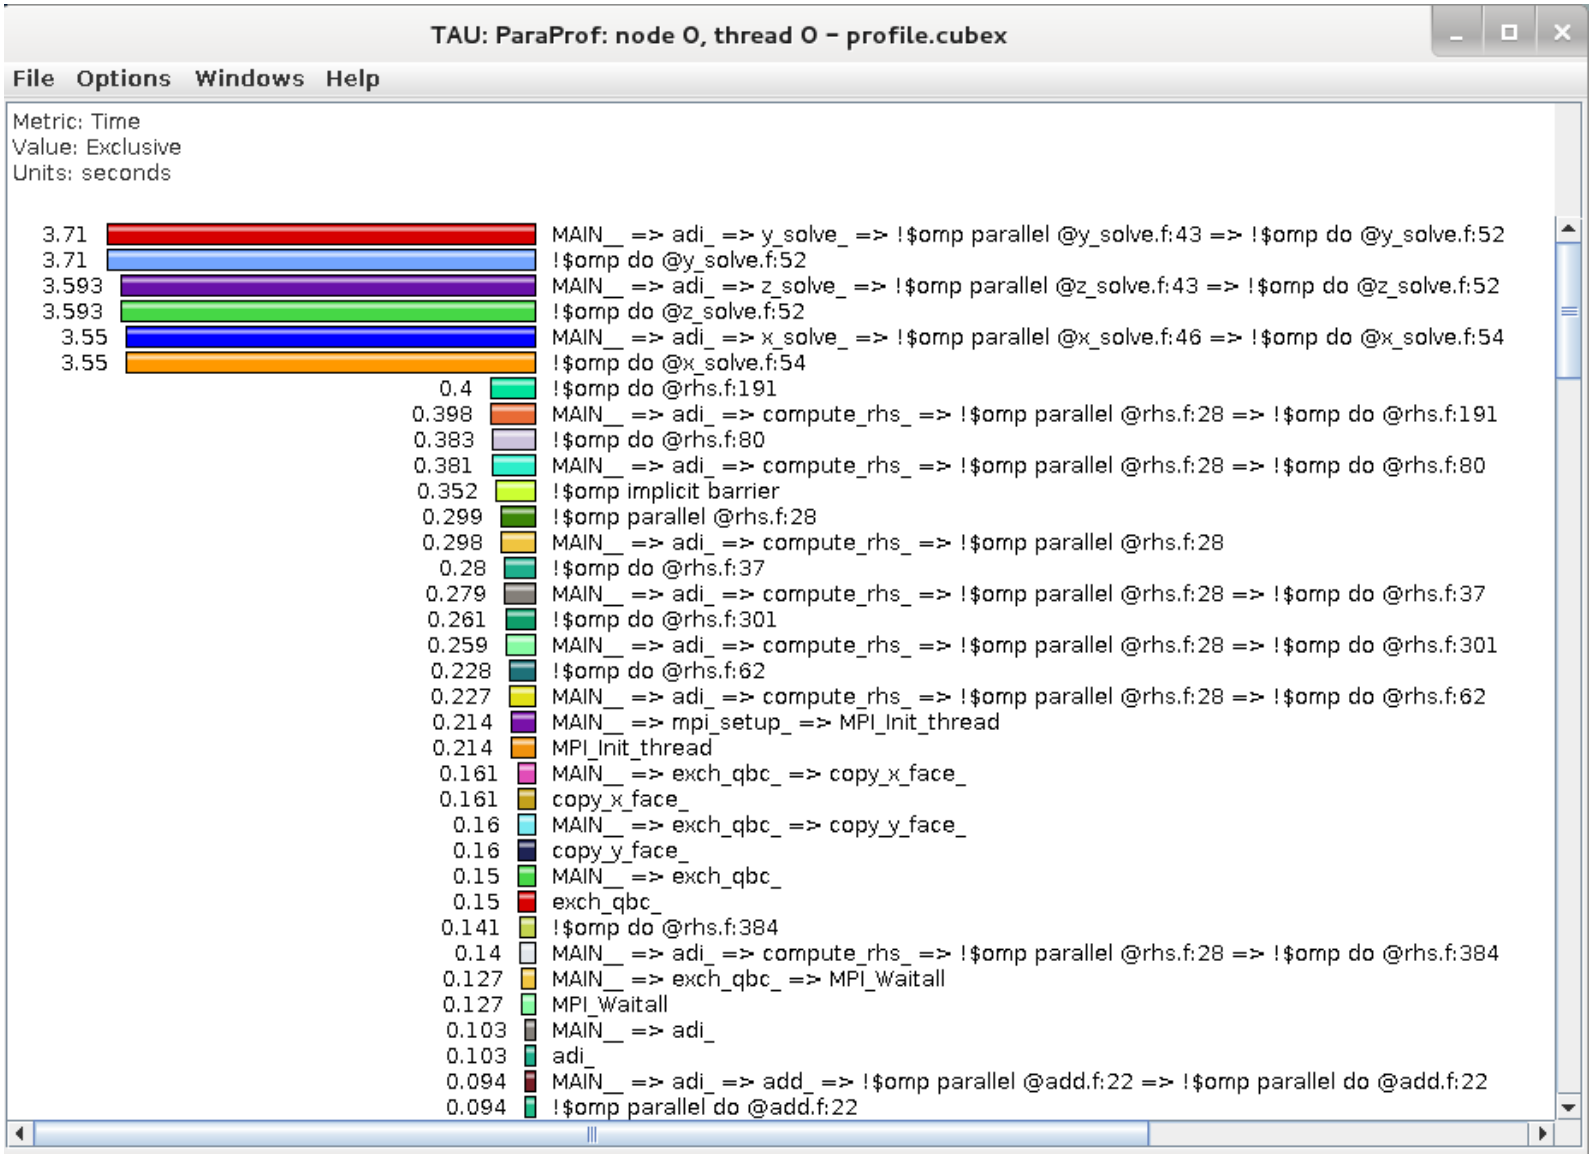

ParaProf: Add Thread to Comparison Window

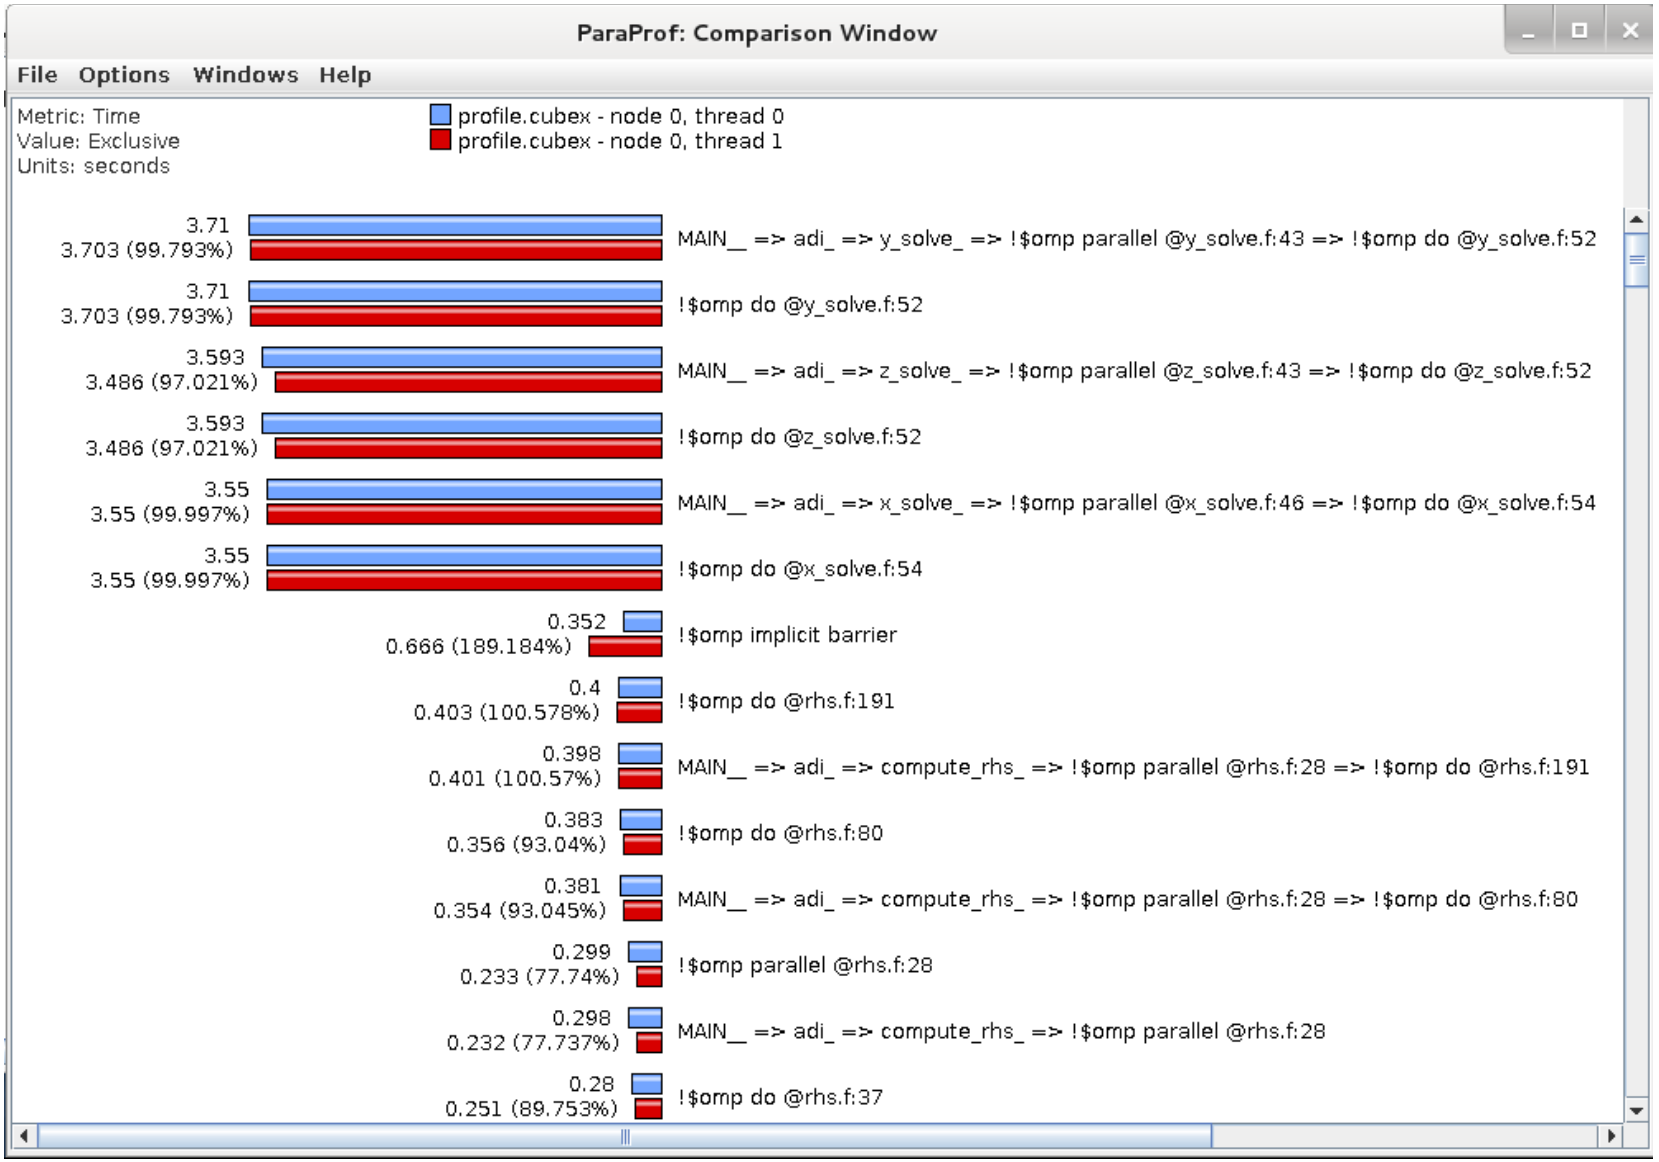

ParaProf: Score-P Profile Files, Database

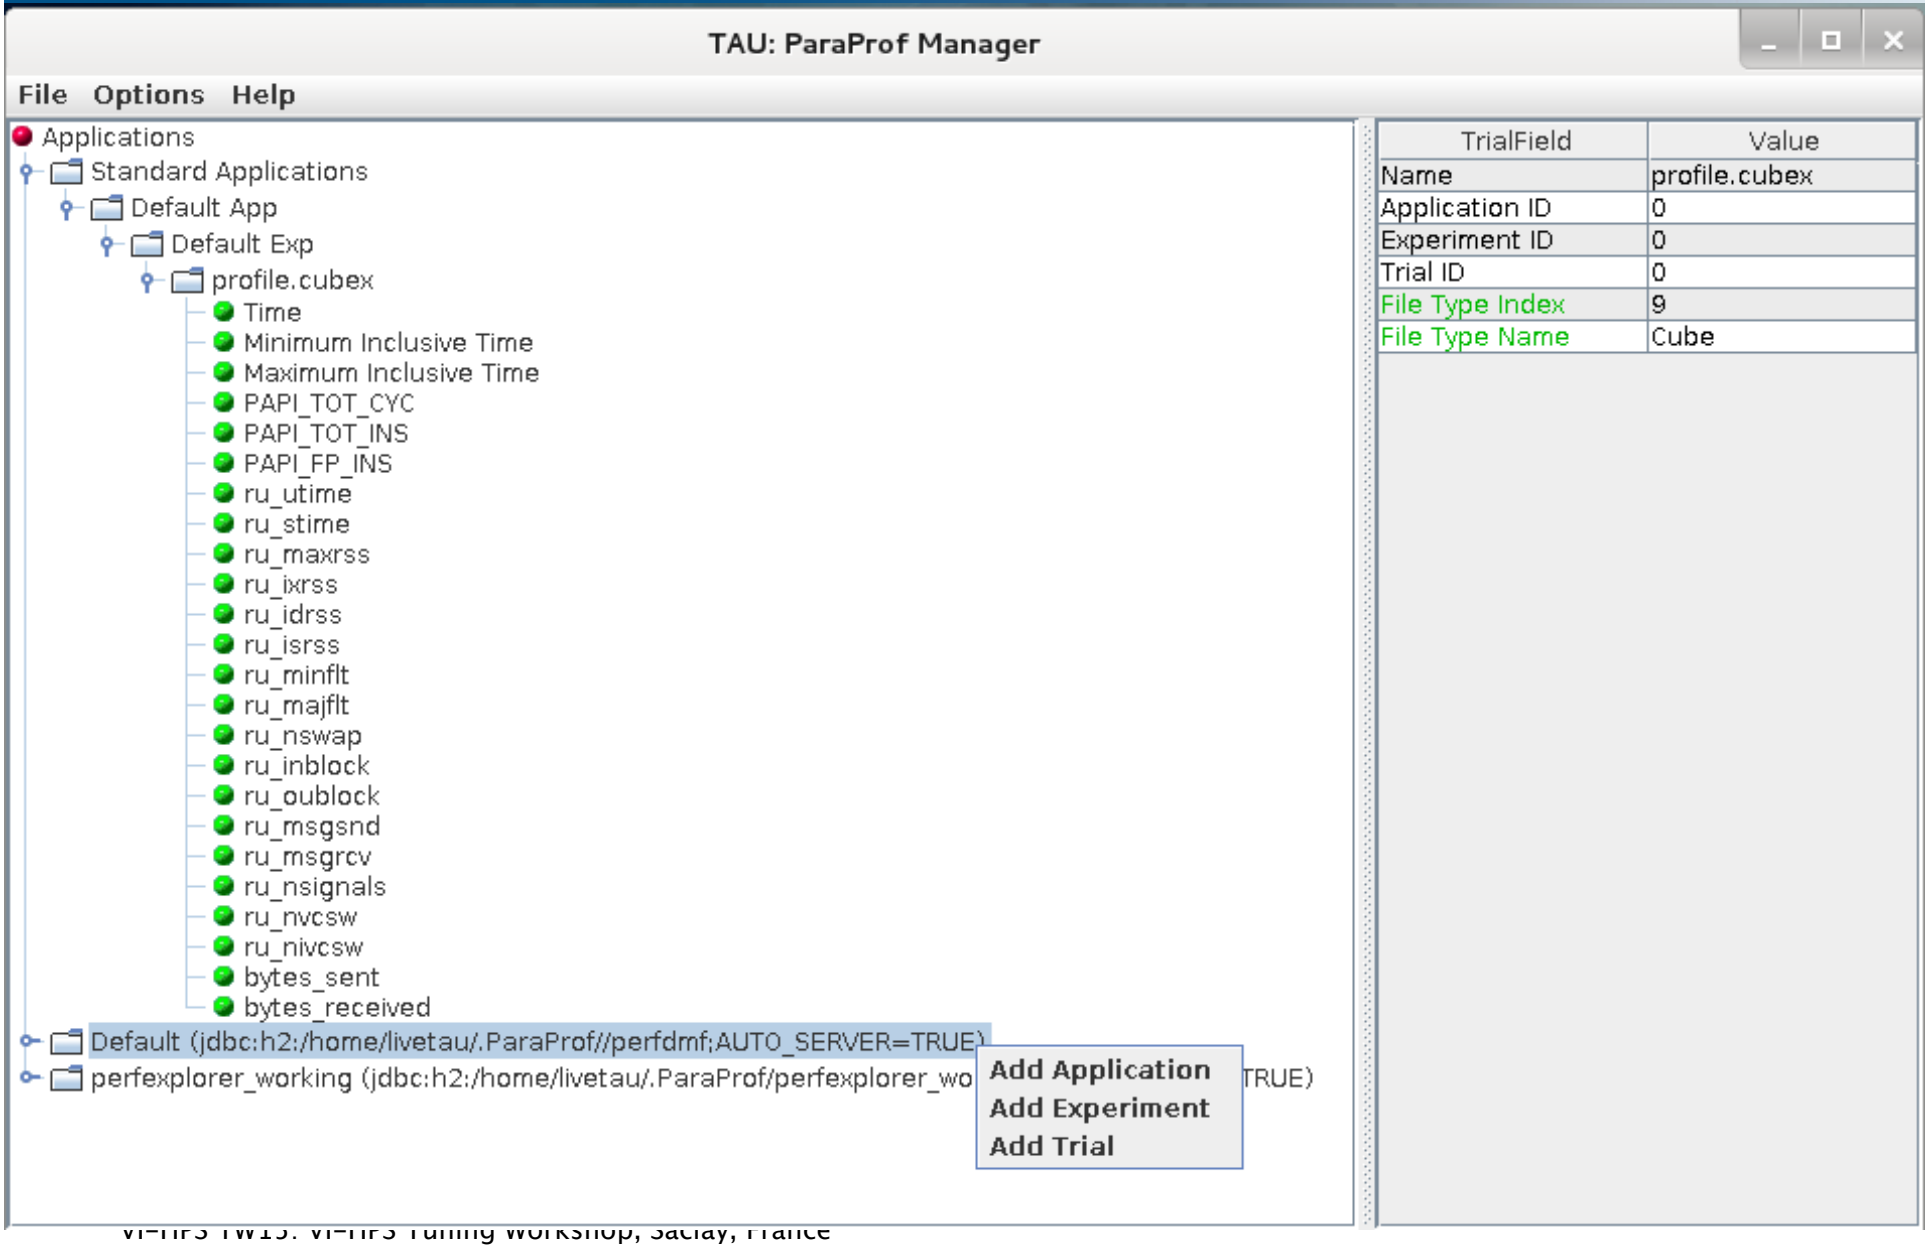

ParaProf: File -> Preferences

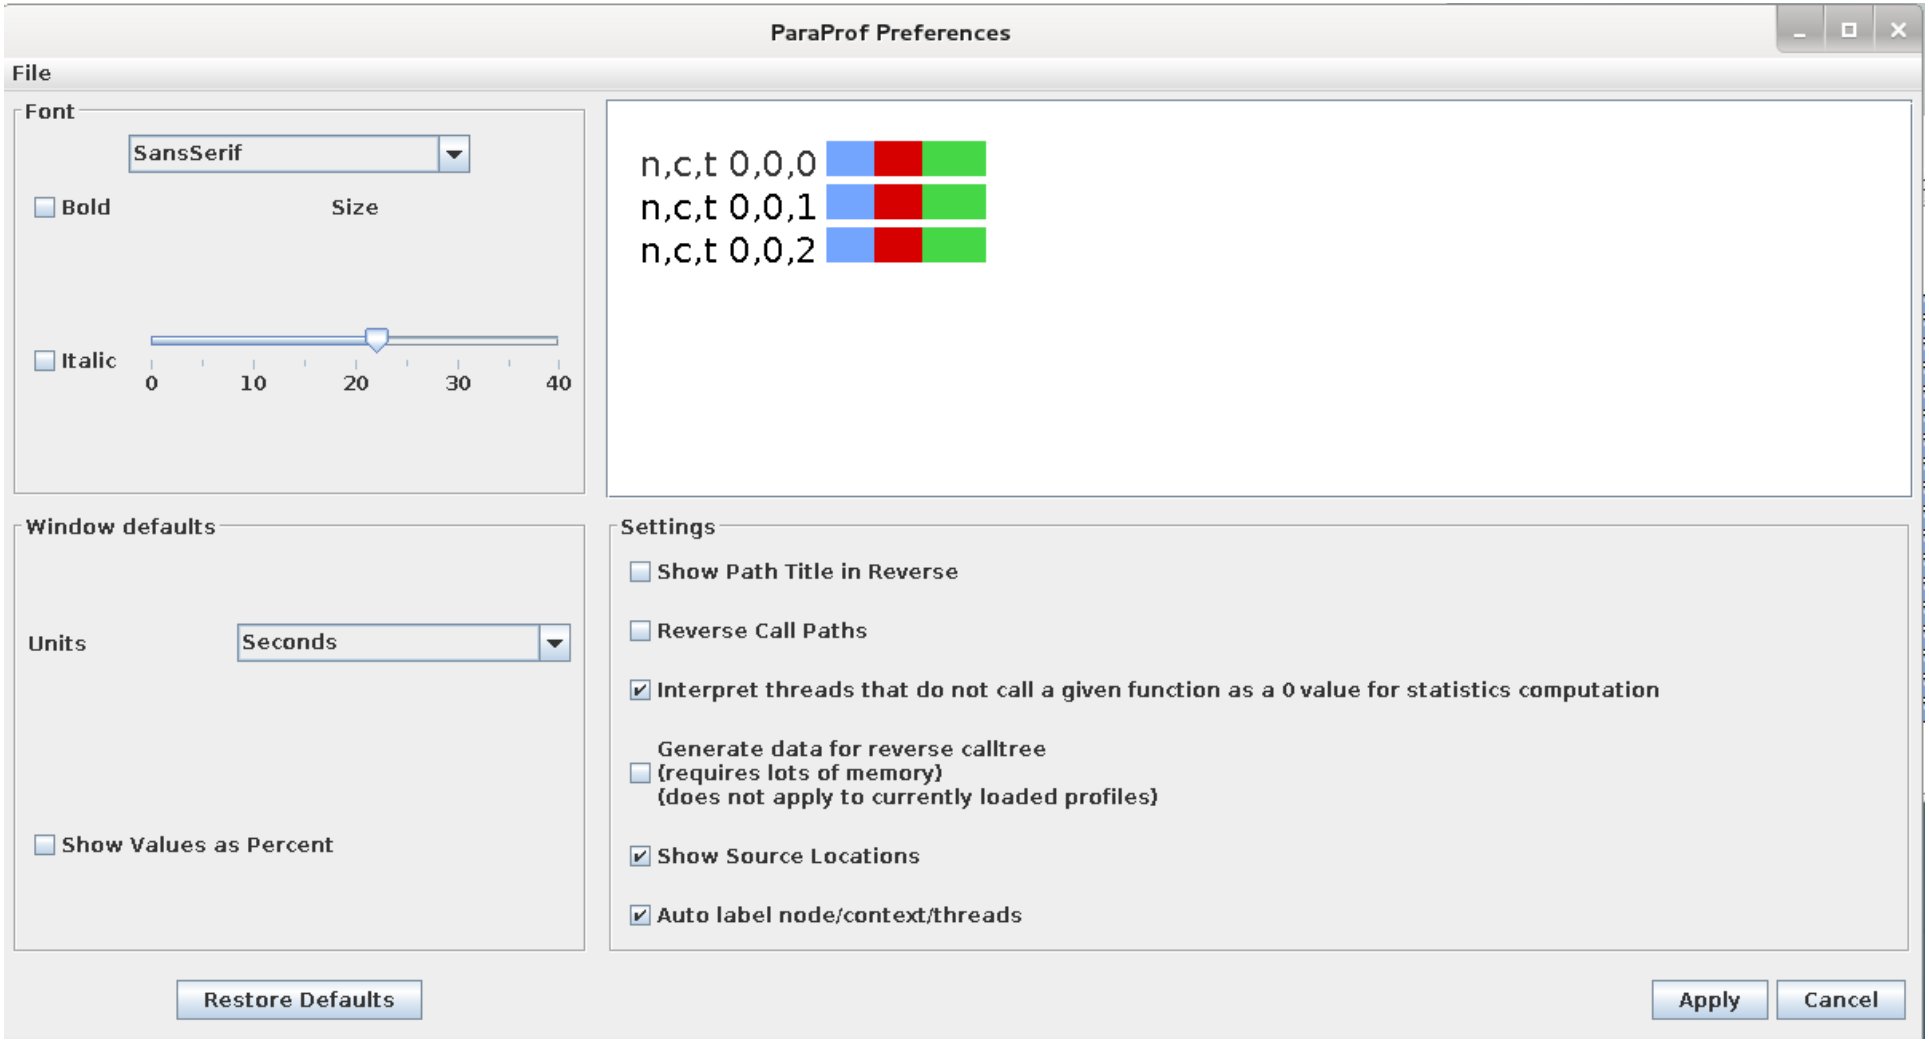

ParaProf: Group Changer Window

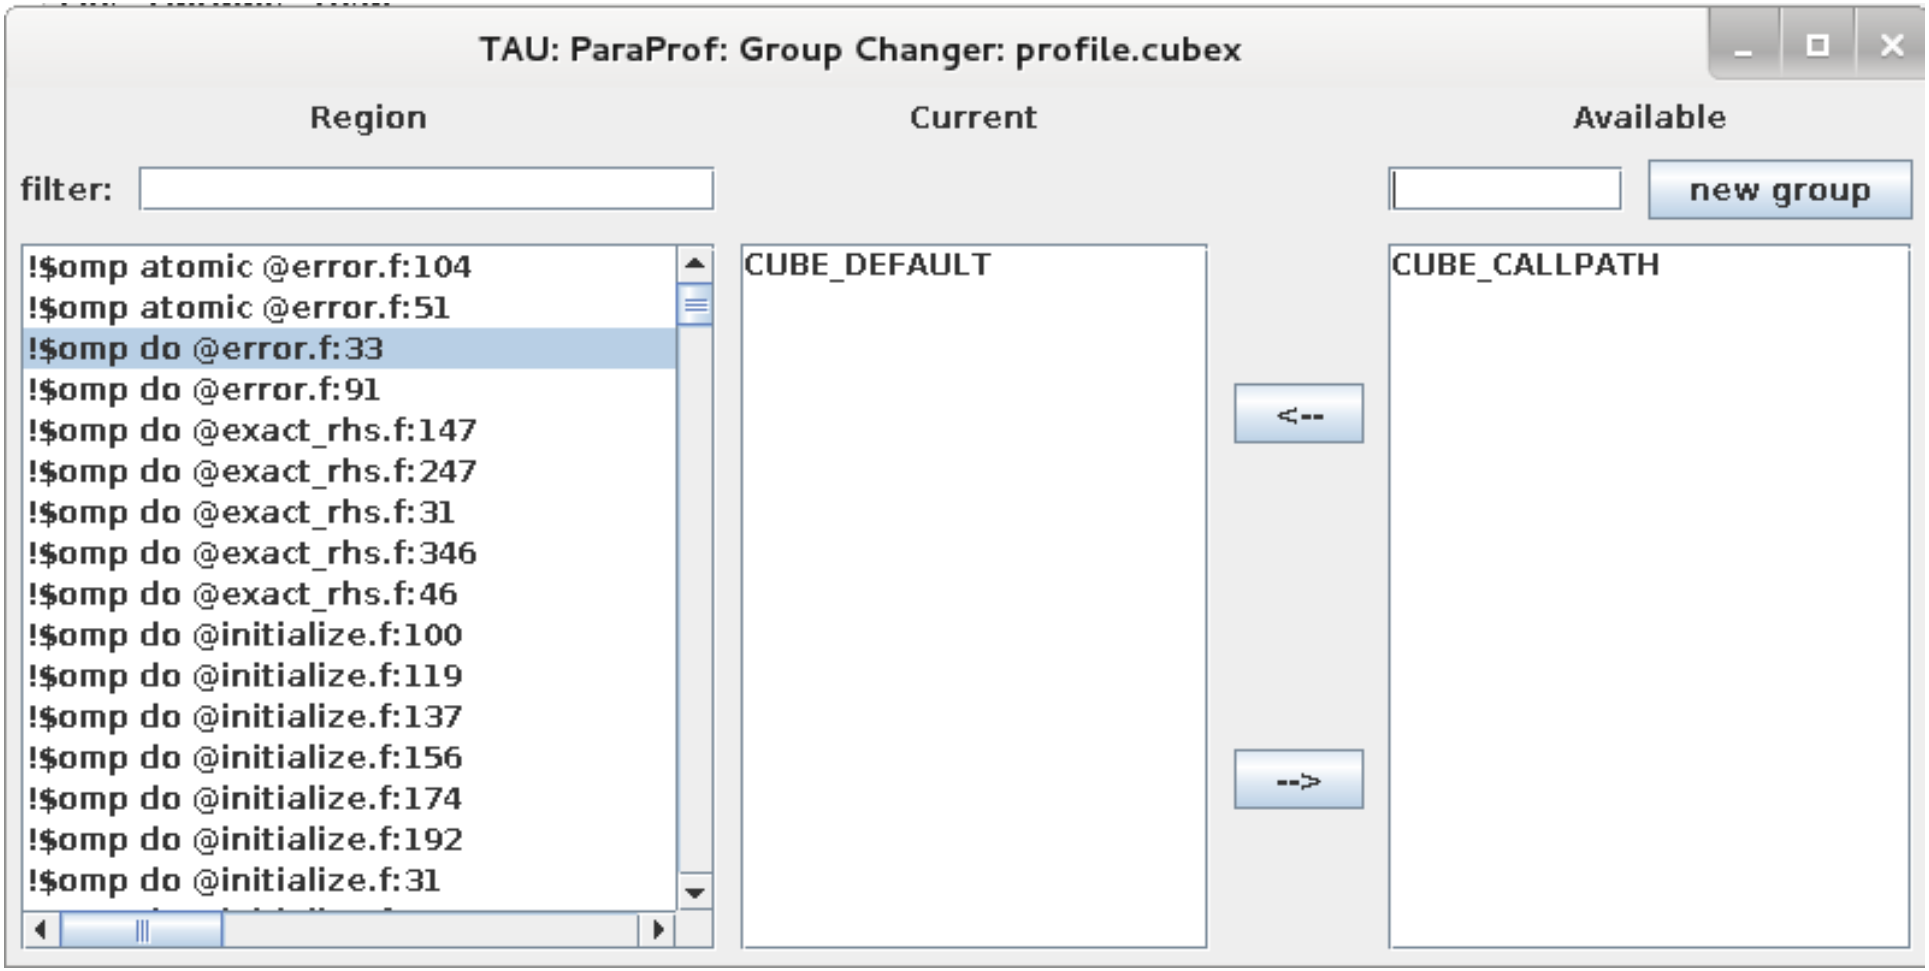

ParaProf: Options -> Derived Metric Panel

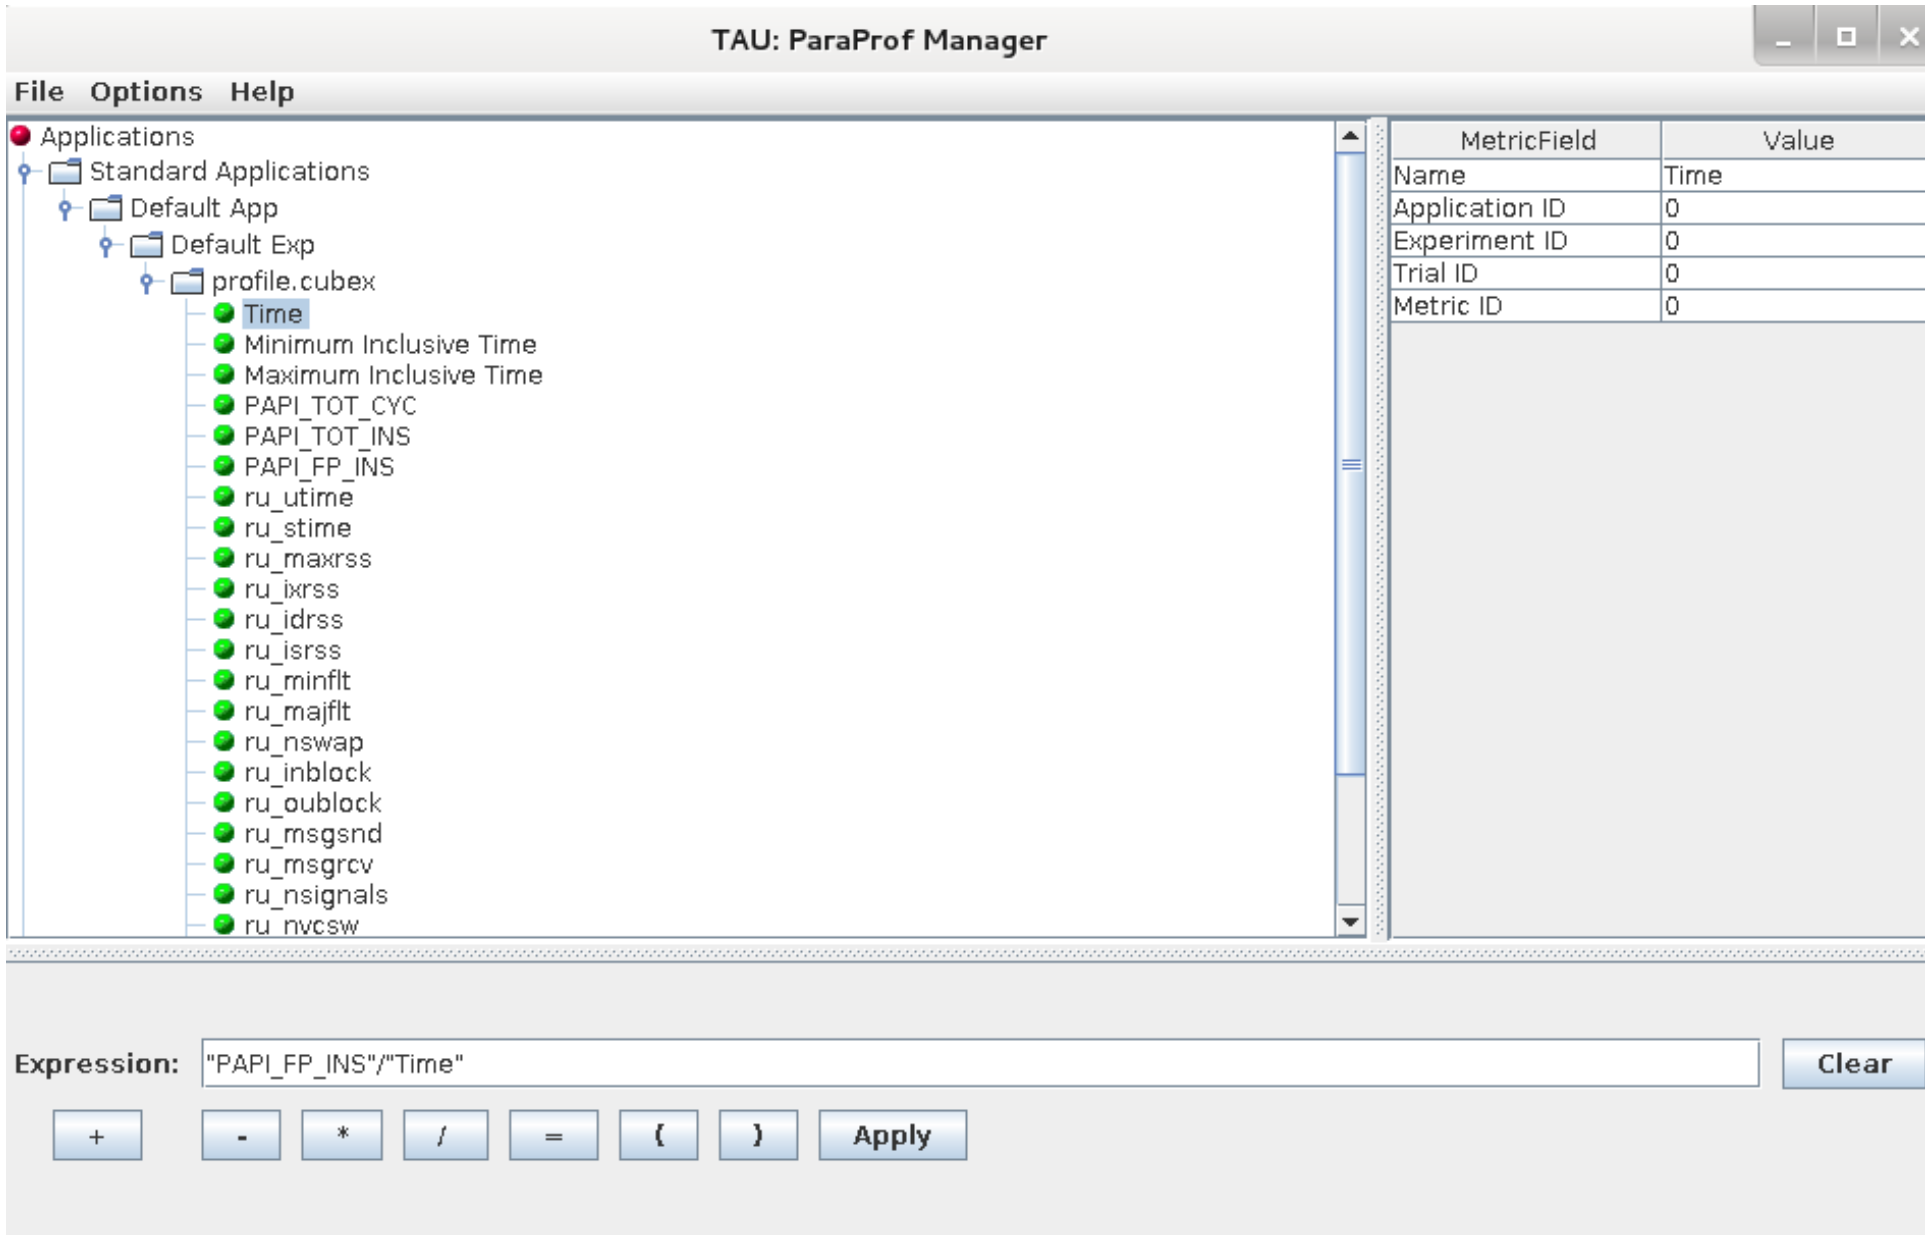

Sorting Derived Flops Metric by Exclusive Time