TAU Performance System® Hands-On

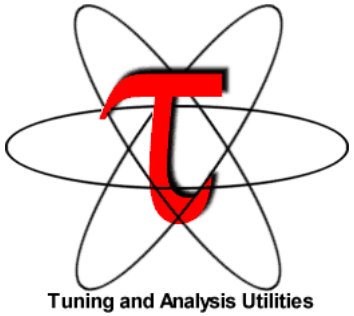

Sameer Shende
Research Professor
sameer@cs.uoregon.edu
University of Oregon

http://tau.uoregon.edu/TAU_TW44_Handson.pptx

TAU: Quickstart Guide

Profiling:

```
MPI: % mpirun -np 16 tau_exec -ebs ./a.out
```

- `Pthread:  % mpirun -np 16 tau_exec -T mpi,pthread -ebs ./a.out`
- `CUDA: % mpirun -np 16 tau_exec -T cupti,mpi -cupti -ebs ./a.out`
- `Score-P: % mpirun -np 16 tau_exec -T scorep,mpi ./a.out`

```
Analysis: % pprof -a -m | more;  % paraprof (GUI)
```

Tracing:

- `Vampir: MPI: % export TAU_TRACE=1; export TAU_TRACE_FORMAT=otf2`
  `% mpirun -np 16 tau_exec ./a.out; vampir traces.otf2 &`
- `Chrome/Jumpshot: % export TAU_TRACE=1; mpirun -np 64 tau_exec ./a.out`
  `% tau_treemerge.pl;`

```
Chrome: % tau_trace2json tau.trc tau.edf -chrome -ignoreatomic -o app.json
   Chrome browser: chrome://tracing   (Load -> app.json) or Perfetto.dev
```

- `Jumpshot: tau2slog2 tau.trc tau.edf -o app.slog2; jumpshot app.slog2`

TAU Hands-On

Hands-On Exercises on Barnard, TU Dresden

```
% module load intel TAU Java
% wget http://tau.uoregon.edu/workshop.tgz
% tar workshop.tgz; cd workshop
% cat README
```

Hands-On Exercises on CLAIX, RWTH

```
% . /home/hpc/vihps-tw44/setup.sh
% module load TAU
% wget http://tau.uoregon.edu/workshop.tgz
% tar workshop.tgz; cd workshop
% cat README
```

Exercise 1: CoMD

Hands-On Exercises: Using TAU with an un-instrumented MPI app

```
cd workshop; cat handson.txt
cd src-mpi; make
cd ../bin
sbatch tau.sbatch.tud
Or
sbatch tau.sbatch.rwth

pprof -a | more
paraprof &
```

Exercise 2: CoMD with TAU and Score-P

Using PDT to instrument source code with TAU and Score-P

Edit srun command to include tau_exec –T scorep

```
srun tau_exec -T scorep ./CoMD-mpi --xproc 4 --yproc 4 --zproc 1 \
--nx 80 --ny 80 --nz 20
```

```
cd scorep-<dir>
paraprof profile.cubex
```

More Score-P examples:
http://tau.uoregon.edu/cubex_ex.tgz

Exercise 3: NPB3.3-MZ-MPI

Hands-On Exercises with OpenMP Tools Interface (OMPT):

```
cat workshop/handson.txt
# Barnard, TUD:
salloc -N 1 --ntasks-per-node=16

# CLAIX, RWTH
salloc -N 1 --ntasks-per-node=16 -p c23test
cd workshop/NPB3.3-MZ-MPI
make clean; make suite
cd bin
cat r.sh
./r.sh
cat rt.sh
./rt.sh
```

ParaProf

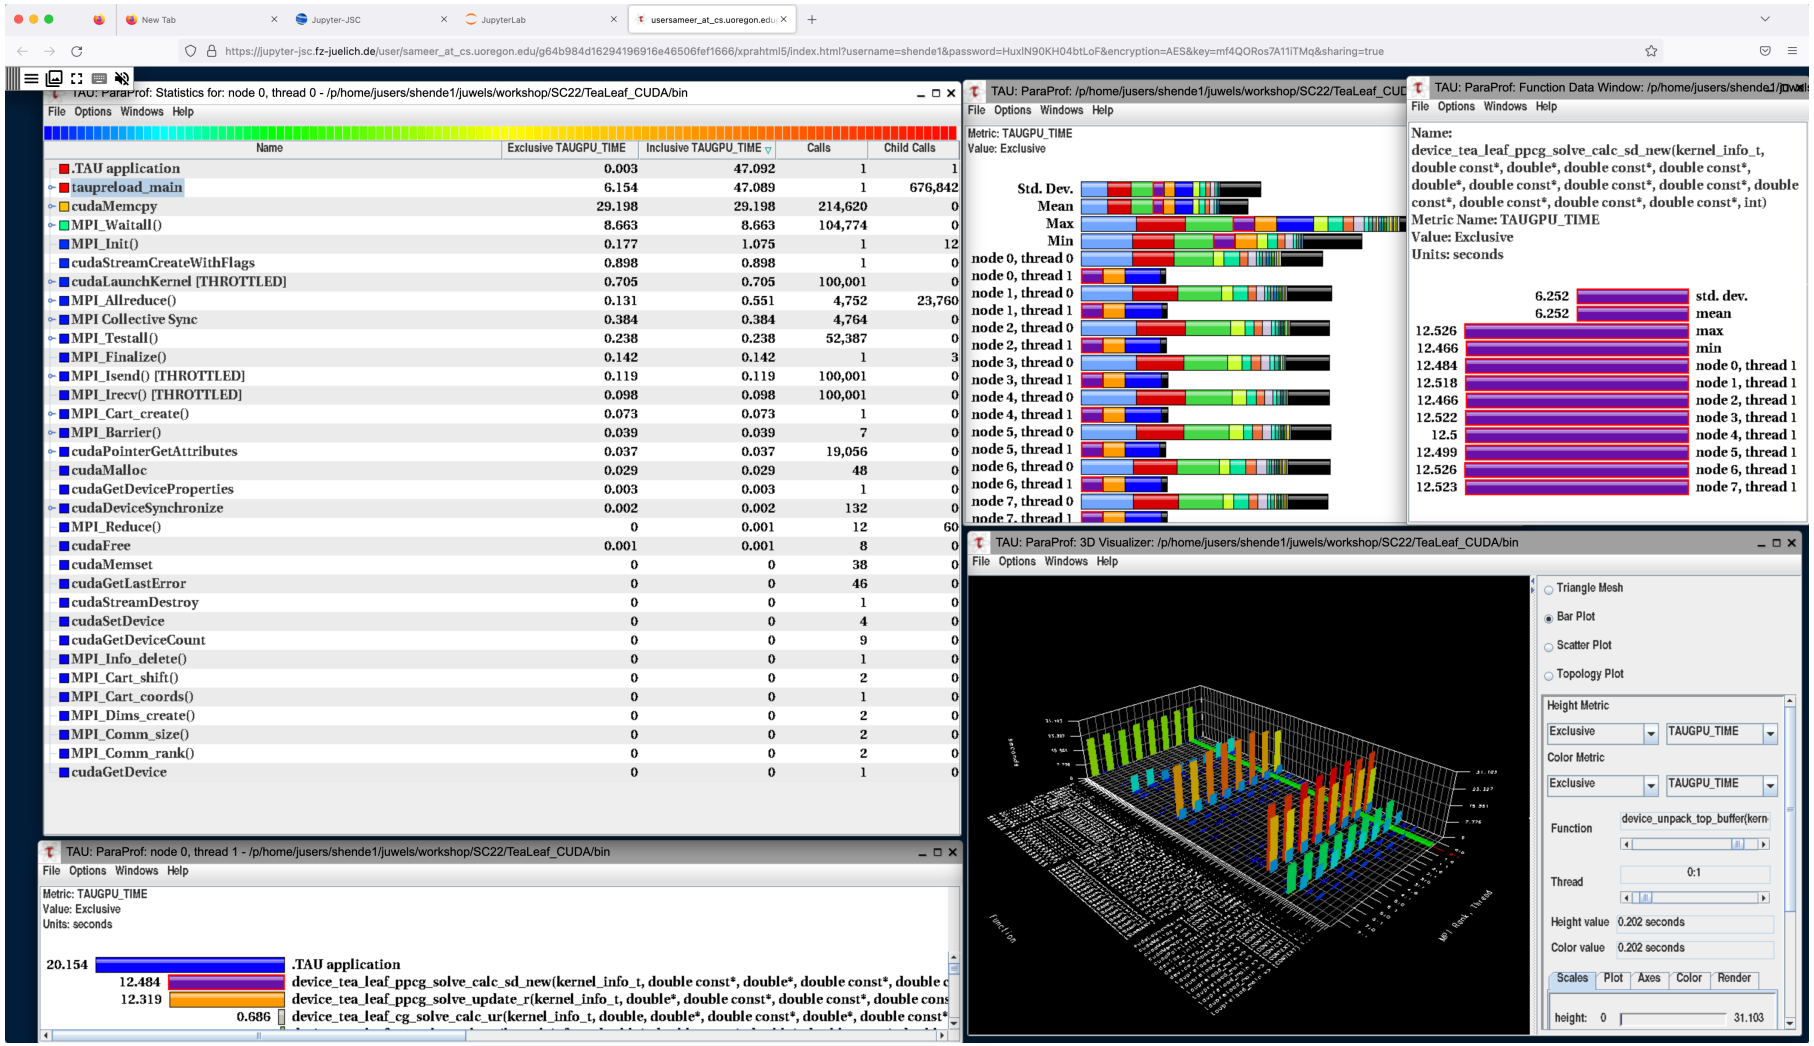
TAU: ParaProf: Statistics for: node 0, thread 0 - /p/home/jusers/shende1/juwels/workshop/SC22/TeaLeaf_CUDA/bin
File Options Windows Help
Name | Exclusive TAUGPU_TIME | Inclusive TAUGPU_TIME | Calls | Child Calls
.TAU application 0.003 47.092 1 1
taupreload_main 6.154 47.089 1 676,842
cudaMemcpy 29.198 29.198 214,620 0
MPI_Waitall() 8.663 8.663 104,774 0
MPI_Init() 0.177 1.075 1 12
cudaStreamCreateWithFlags 0.898 0.898 1 0
cudaLaunchKernel [THROTTLED] 0.705 0.705 100,001 0
MPI_Allreduce() 0.131 0.551 4,752 23,760
MPI Collective Sync 0.384 0.384 4,764 0
MPI_Testall() 0.238 0.238 52,387 0
MPI_Finalize() 0.142 0.142 1 3
MPI_Isend() [THROTTLED] 0.119 0.119 100,001 0
MPI_Irecv() [THROTTLED] 0.098 0.098 100,001 0
MPI_Cart_create() 0.073 0.073 1 0
MPI_Barrier() 0.039 0.039 7 0
cudaPointerGetAttributes 0.037 0.037 19,056 0
cudaMalloc 0.029 0.029 48 0
cudaGetDeviceProperties 0.003 0.003 1 0
cudaDeviceSynchronize 0.002 0.002 132 0
MPI_Reduce() 0 0.001 12 60
cudaFree 0.001 0.001 8 0
cudaMemset 0 0 38 0
cudaGetLastError 0 0 46 0
cudaStreamDestroy 0 0 1 0
cudaSetDevice 0 0 4 0
cudaGetDeviceCount 0 0 9 0
MPI_Info_delete() 0 0 1 0
MPI_Cart_shift() 0 0 2 0
MPI_Cart_coords() 0 0 1 0
MPI_Dims_create() 0 0 2 0
MPI_Comm_size() 0 0 2 0
MPI_Comm_rank() 0 0 2 0
cudaGetDevice 0 0 1 0
TAU: ParaProf: node 0, thread 1 - /p/home/jusers/shende1/juwels/workshop/SC22/TeaLeaf_CUDA/bin
File Options Windows Help
Metric: TAUGPU_TIME
Value: Exclusive
Units: seconds
20.154 .TAU application
12.484 device_tea_leaf_ppcg_solve_calc_sd_new(kernel_info_t, double const*, double*, double const*, double
12.319 device_tea_leaf_ppcg_solve_update_r(kernel_info_t, double*, double const*, double const*, double cons
0.686 device_tea_leaf_cg_solve_calc_ur(kernel_info_t, double, double*, double const*, double*, double const*
TAU: ParaProf: /p/home/jusers/shende1/juwels/workshop/SC22/TeaLeaf_CUD
File Options Windows Help
Metric: TAUGPU_TIME
Value: Exclusive
Std. Dev.
Mean
Max
Min
node 0, thread 0
node 0, thread 1
node 1, thread 0
node 1, thread 1
node 2, thread 0
node 2, thread 1
node 3, thread 0
node 3, thread 1
node 4, thread 0
node 4, thread 1
node 5, thread 0
node 5, thread 1
node 6, thread 0
node 6, thread 1
node 7, thread 0
node 7, thread 1
TAU: ParaProf: Function Data Window: /p/home/jusers/shende1/juwels
File Options Windows Help
Name:
device_tea_leaf_ppcg_solve_calc_sd_new(kernel_info_t, double const*, double*, double const*, double const*, double*, double const*, double const*, double const*, double const*, double const*, int)
Metric Name: TAUGPU_TIME
Value: Exclusive
Units: seconds
6.252 std. dev.
6.252 mean
12.526 max
12.466 min
12.484 node 0, thread 1
12.518 node 1, thread 1
12.466 node 2, thread 1
12.522 node 3, thread 1
12.5 node 4, thread 1
12.499 node 5, thread 1
12.526 node 6, thread 1
12.523 node 7, thread 1
TAU: ParaProf: 3D Visualizer: /p/home/jusers/shende1/juwels/workshop/SC22/TeaLeaf_CUDA/bin
File Options Windows Help
Triangle Mesh
Bar Plot
Scatter Plot
Topology Plot
Height Metric
Exclusive
TAUGPU_TIME
Color Metric
Exclusive
TAUGPU_TIME
Function
device_unpack_top_buffer(kern
Thread
0:1
Height value 0.202 seconds
Color value 0.202 seconds
Scales Plot Axes Color Render
height: 0 31.103

Performance Research Lab, University of Oregon, Eugene, USA

Support Acknowledgments

- US Department of Energy (DOE)
  - Office of Science contracts, ECP
  - SciDAC, LBL contracts
  - LLNL-LANL-SNL ASC/NNSA contract
  - Battelle, PNNL contract
  - ANL, ORNL contract
- Department of Defense (DoD)
  - PETTT, HPCMP
- National Science Foundation (NSF)
  - Glassbox, SI-2
- NASA
- CEA, France
- Partners:
  - University of Oregon
  - ParaTools, Inc., ParaTools, SAS
  - The Ohio State University
  - University of Tennessee, Knoxville
  - T.U. Dresden, GWT
  - Juelich Supercomputing Center



















Acknowledgement

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation's exascale computing imperative.

Download TAU from U. Oregon

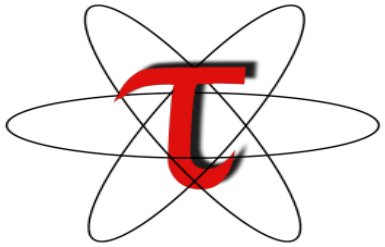

http://tau.uoregon.edu

http://www.hpclinux.com [LiveDVD, OVA]

https://e4s.io [Containers for Extreme-Scale Scientific Software Stack]

Free download, open source, BSD license