

#### **TAU Performance System®**



Sameer Shende <u>sameer@cs.uoregon.edu</u>

#### University of Oregon

http://tau.uoregon.edu/TAU\_TW41\_handson.pdf



# TAU hands-on exercises

41ST VI-HPS TUNING WORKSHOP (JSC/RWTH AACHEN - ONLINE, 7TH - 11TH FEB 2022)

## **TAU tutorial exercise objectives**

- Familiarise with usage of TAU tools
  - complementary tools' capabilities & interoperability
- Prepare to apply tools productively to your applications(s)
- Exercise is based on a small portable benchmark code
  - unlikely to have significant optimisation opportunities
- Optional (recommended) exercise extensions
  - analyse performance of alternative configurations
  - investigate effectiveness of system-specific compiler/MPI optimisations and/or placement/binding/affinity capabilities
  - investigate scalability and analyse scalability limiters
  - compare performance on different HPC platforms

• ...

## Installing TAU on your laptop for paraprof (GUI)

#### Microsoft Windows

- Install Java from Oracle.com
- http://tau.uoregon.edu/tau.exe
- Install, click on a ppk file to launch paraprof
- ■macOS (x86\_64)
  - Install Java 11.0.3:
    - Download and install <u>http://tau.uoregon.edu/java.dmg</u>
    - If you have multiple Java installations, add to your ~/.zshrc (or ~/.bashrc as appropriate):
    - export PATH=/Library/Java/JavaVirtualMachines/jdk-11.0.3.jdk/Contents/Home/bin:\$PATH
    - java -version
  - Download and install TAU (copy to /Applications from dmg):
    - http://tau.uoregon.edu/tau.dmg
    - export PATH=/Applications/TAU/tau/apple/bin:\$PATH
    - paraprof app.ppk &
- macOS (arm64, M1)
  - <u>http://tau.uoregon.edu/java\_arm64.dmg</u>
  - http://tau.uoregon.edu/tau\_arm64.dmg
- Linux (http://tau.uoregon.edu/tau.tgz)
  - ./configure; make install; export PATH=<taudir>/x86\_64/bin:\$PATH; paraprof app.ppk &

## Using TAU on JUWELS Booster at JSC, FZJ, Germany

- Setup preferred program environment compilers using PGI compilers with MPI and CUDA
  - % ssh -Y juwels-booster.fz-juelich.de -l <USER>
  - % source \$PROJECT\_training2123/setup.sh
- Load the module for TAU

% module load TAU

Copy the workshop tarball to your WORK directory

```
% cd $SCRATCH_training/$USER
% tar xf $PROJECT_training2123/examples/workshop.tgz
% cd workshop; cat README; cat handson.txt
% cd CoMD/src-mpi; make ; cd ../bin; sbatch tau.sbatch
% pprof -a | more; paraprof &
```

### Using TAU on JUWELS Booster at JSC, FZJ, Germany

#### • Using TAU with CUDA and MPI

```
% cd $SCRATCH_training2123/workshop
% cd TeaLeaf_CUDA; make; cd bin; sbatch tau.sbatch
% pprof -a | more
% paraprof --pack tealeaf_ex1.ppk
<copy to your laptop>
% paraprof tealeaf_ex1.ppk
```

### **TAU's 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>)</subevent></event>	

#### **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	0	Setting to 1 is necessary for event based sampling to resolve addresses with OMPT TR6 (-ompt=download-tr6)

#### **Runtime Environment Variables**

Environment Variable	Default	Description
TAU_TRACK_MEMORY_LEAKS	0	Tracks allocates that were not de-allocated (needs –optMemDbg or tau_exec –memory)
TAU_EBS_SOURCE	TIME	Allows using PAPI hardware counters for periodic interrupts for EBS (e.g., TAU_EBS_SOURCE=PAPI_TOT_INS when TAU_SAMPLING=1)
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 TAU_MEMDBG_PROTECT_*)
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 – optMemDbg while building or tau_exec –memory)
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 –optMemDbg or tau_exec –memory)
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

#### Performance Research Lab, University of Oregon, Eugene, USA



## **Support Acknowledgments**





## 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