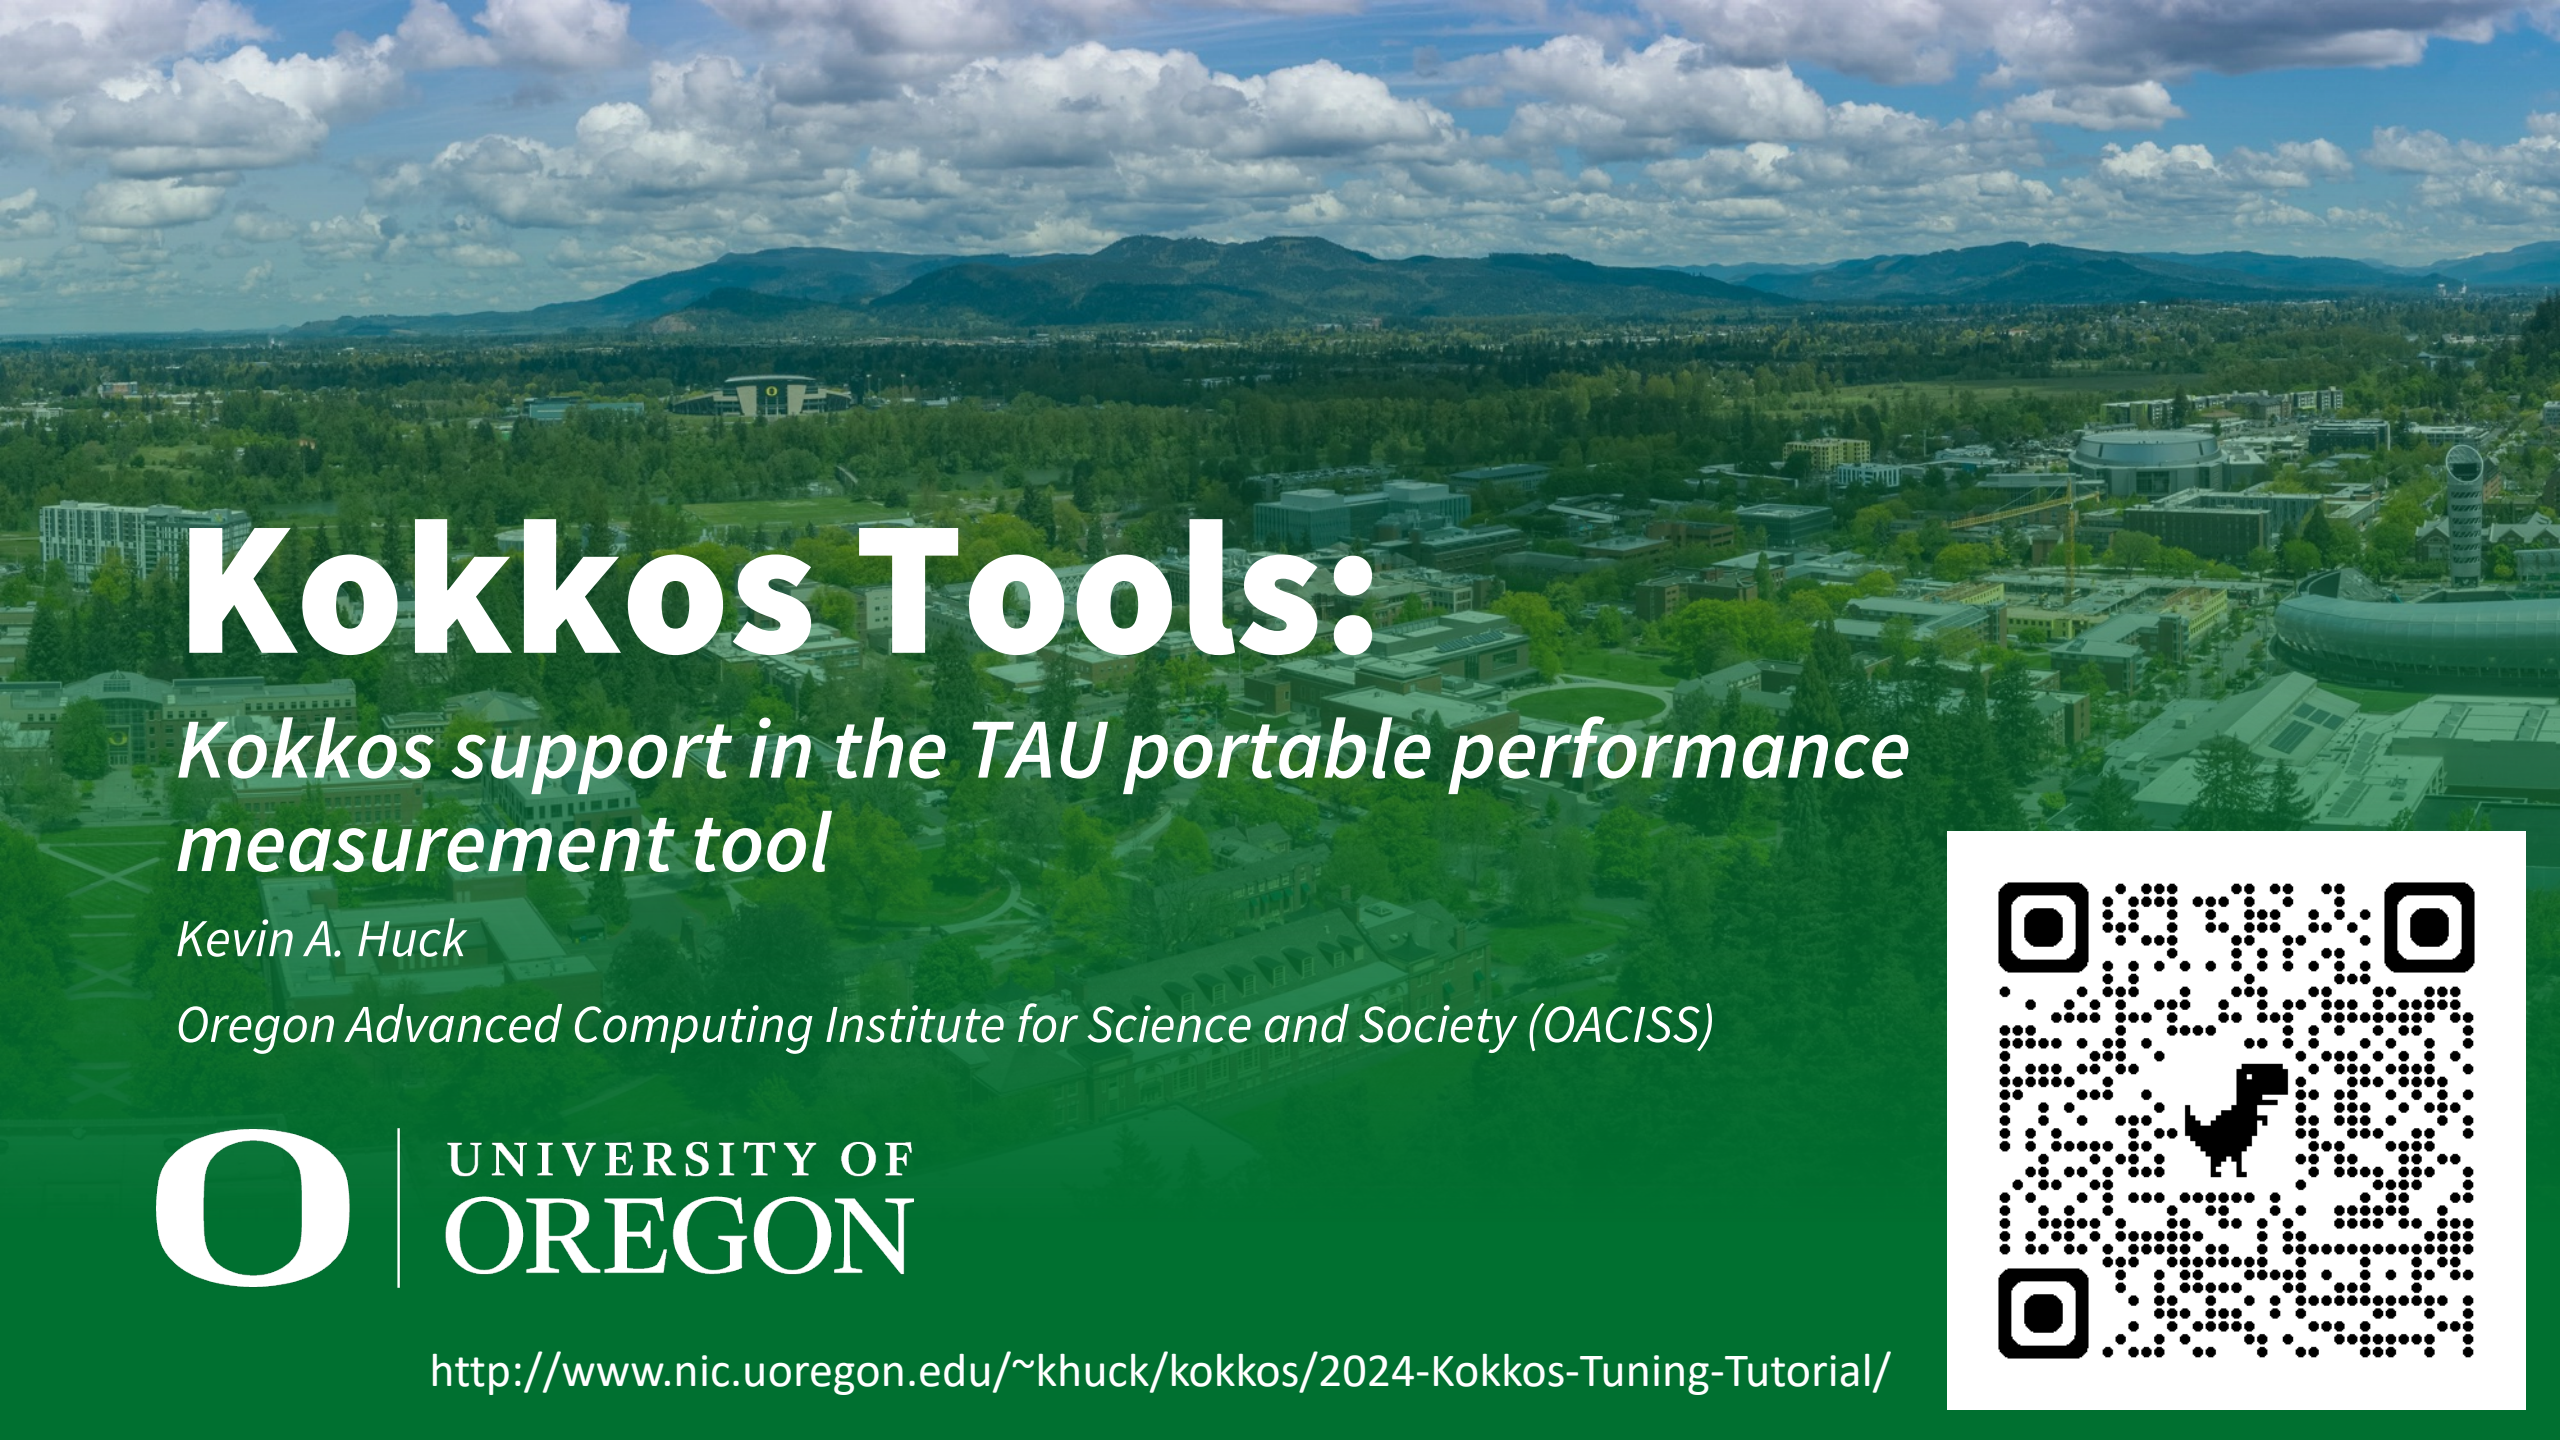

Kokkos Tools:

Kokkos support in the TAU portable performance measurement tool

Kevin A. Huck

Oregon Advanced Computing Institute for Science and Society (OACISS)

UNIVERSITY OF OREGON

http://www.nic.uoregon.edu/~khuck/kokkos/2024-Kokkos-Tuning-Tutorial/

TAU Performance System

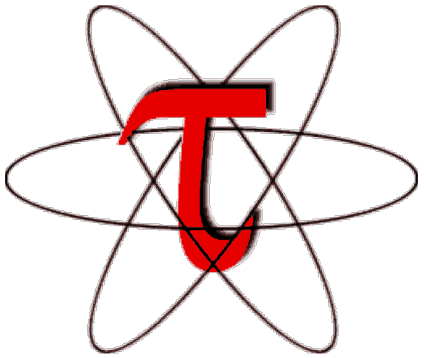

- **T**uning and **A**nalysis **U**tilities (29+ year project)
- Integrated performance toolkit:
  - Multi-level performance instrumentation
  - Highly configurable
  - Widely ported performance profiling / tracing system
  - Portable (java, python) visualization / exploration / analysis tools
- Supports all major HPC programming models
- MPI/SHMEM, OpenMP, OpenACC, CUDA, HIP, SYCL/OneAPI, **Kokkos**...
- Support for ML/AI frameworks: TensorFlow, pyTorch, Horovod
- Integrated with PAPI, LIKWID for hardware counter support
- https://tau.uoregon.edu or https://github.com/UO-OACISS/tau2 (public mirror)

Performance Measurement

- **Timers**
  - Requires instrumentation of some kind
    - Manual, automated
    - Source, compiler provided, binary
    - **Library callbacks**, API wrappers, weak symbol replacement
  - Simple to implement
- **Sampling**
  - Requires specialized system libraries / support
    - Periodic signals, signal handler
    - Call stack unwinding
  - No modification to executable/library needed
  - Potential to interfere with system support (signal handlers)
  - Can mix with timers to generate a hybrid profile

Profiling and Tracing

- **Profiling**: how much time was spent in each measured function on each thread in each process?
  - Collapses the time axis
  - No ordering or causal event information
  - Small summary per thread/process, regardless of execution time – only grows with number of timers & threads/processes

ParaProf profile viewer

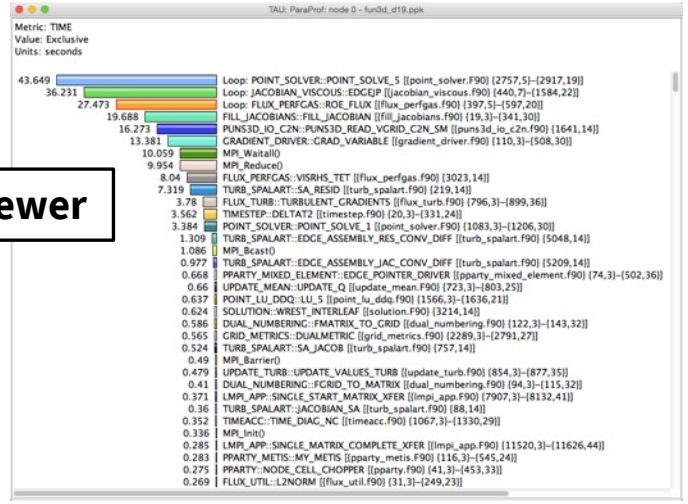

- **Tracing**: record all function entry & exit events on a timeline
  - Detailed view of what happened
  - The longer the program runs, the bigger the trace

JumpShot trace viewer

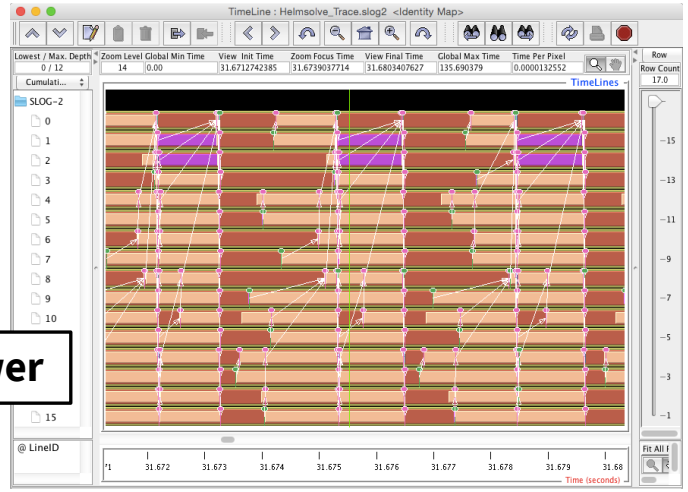

TAU Analysis Tools: ParaProf, Vampir

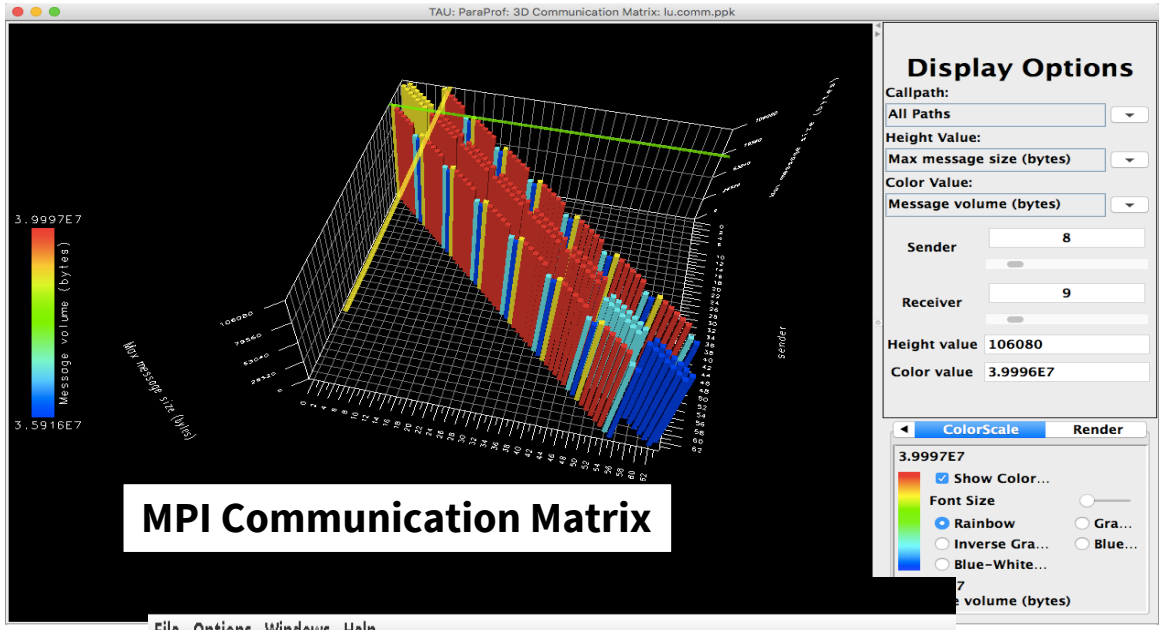

MPI Communication Matrix

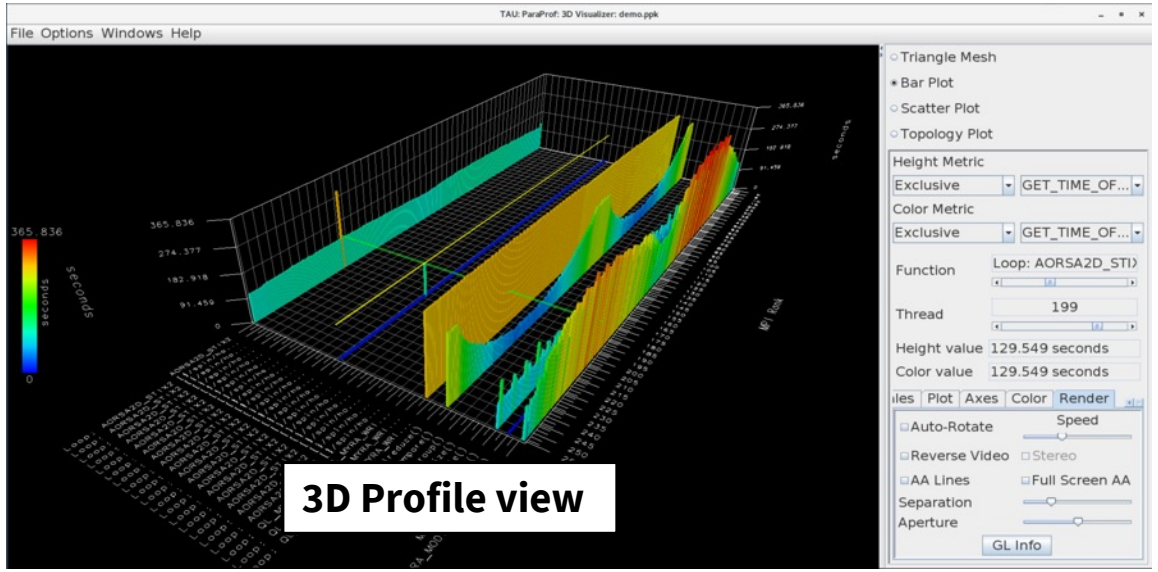

3D Profile view

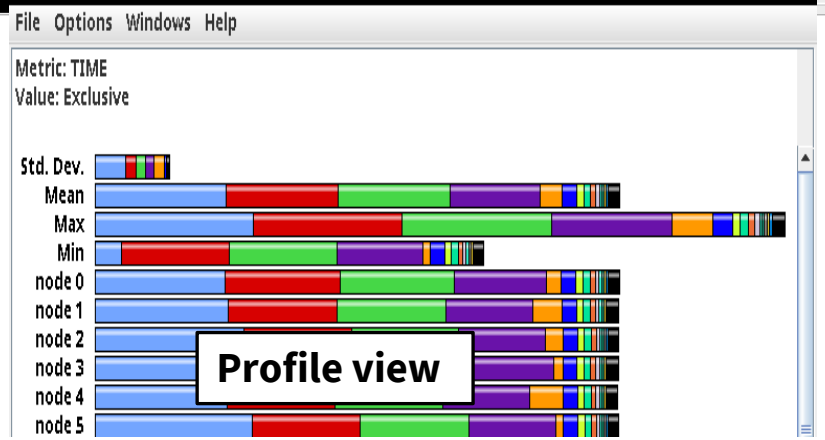

Profile view

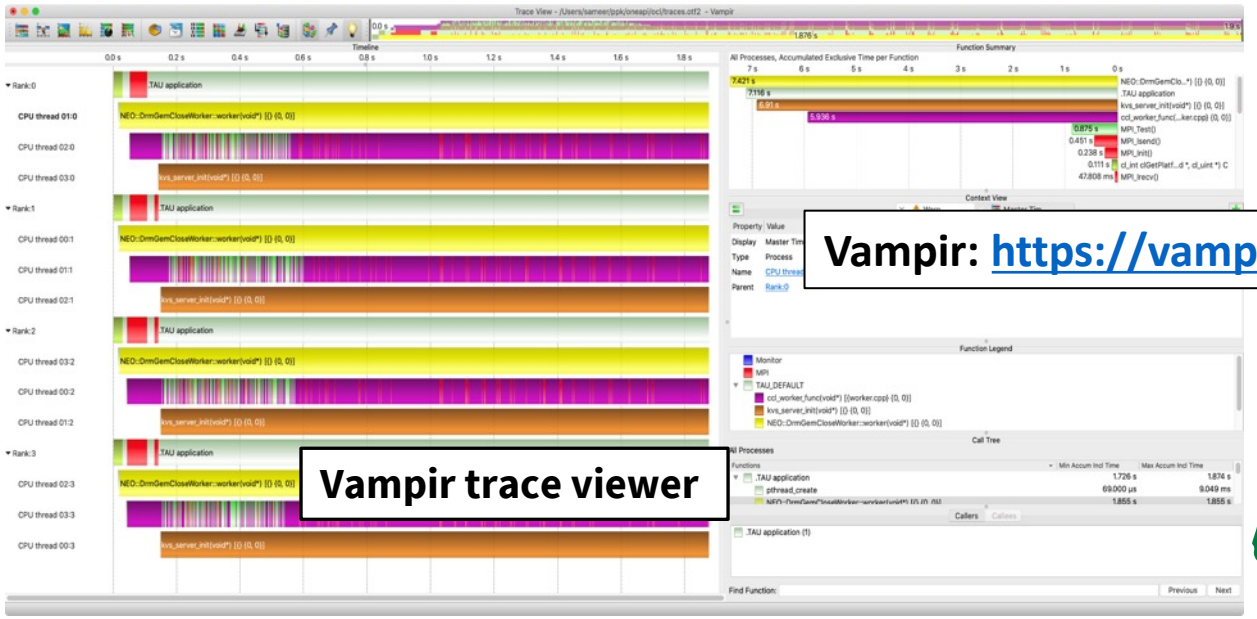

Vampir trace viewer

Vampir: https://vampir.eu

Kokkos support in TAU – since February, 2017

- TAU implements the Kokkos Profiling API (`Kokkos_Profiling_C_Interface.h`)
- TAU sets an environment variable KOKKOS_TOOLS_LIBS to tell Kokkos that it should enable profiling and enable function callbacks to the TAU implementations
- TAU implements
  - `kokkosp_[init|finalize]_library`
  - `kokkosp_[begin|end]_parallel_[for|scan|reduce]`
  - `kokkosp_[push|pop]_profile_region`
- Names for regions are passed to the tools to provide intelligent labels
- In addition, TAU also implements support for native Pthreads, OpenMP, OpenACC, CUDA, HIP, SYCL back-end measurement – no code changes necessary
- Fun fact: if you have a Raja application, and Raja is configured with `-DRAJA_ENABLE_RUNTIME_PLUGINS`, Raja implements the same callback API!

TAU Example – Kokkos Lulesh (from kokkos-miniapps)

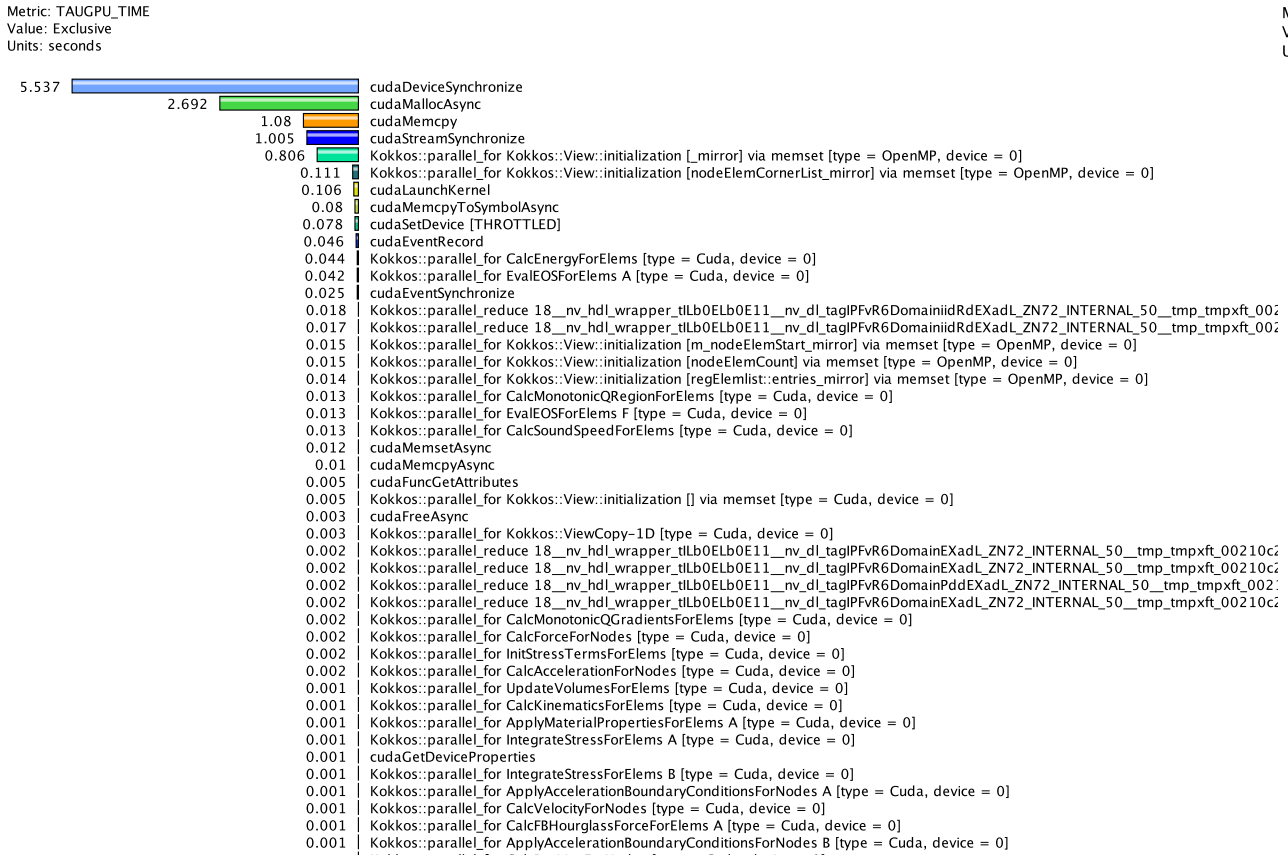

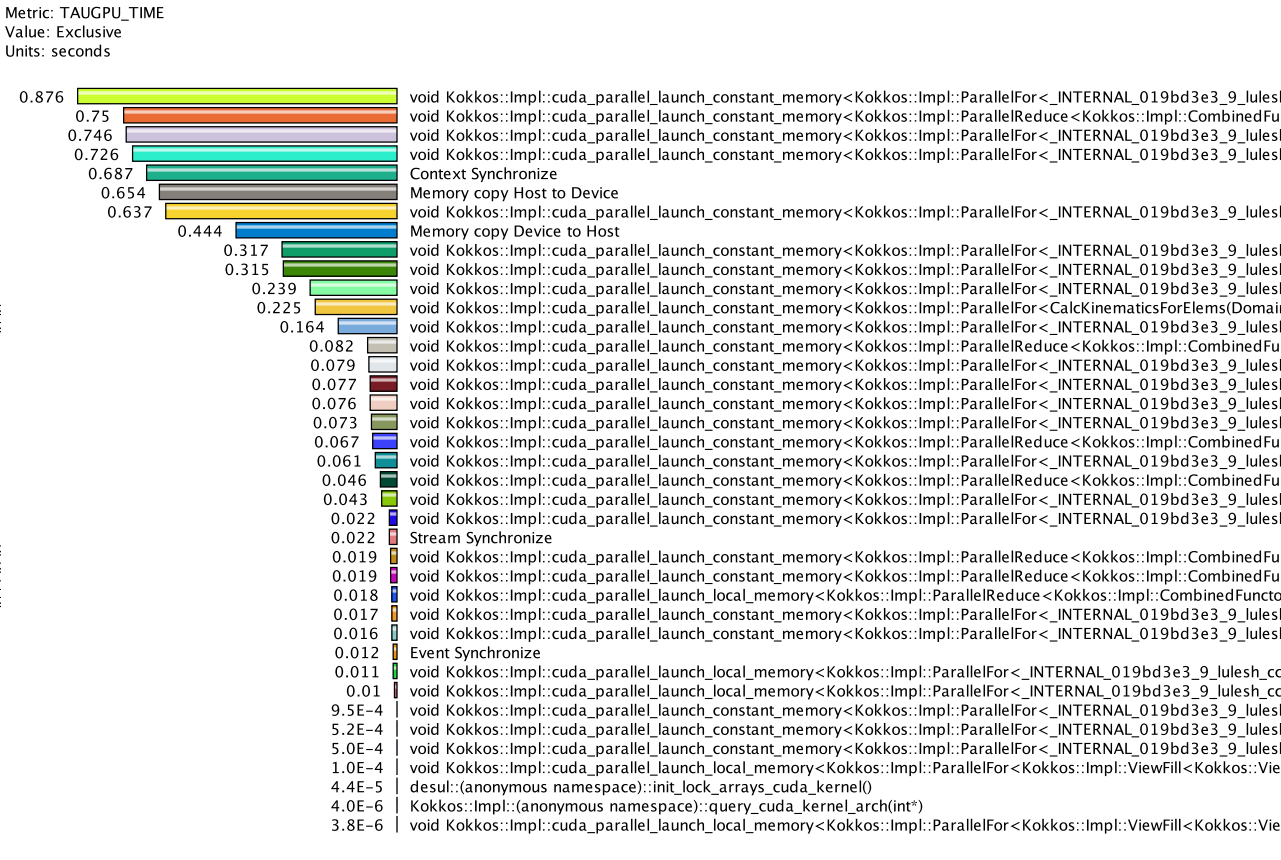

Main thread launching kernels

Virtual thread with CUDA activity

UNIVERSITY OF OREGON

Intel Xeon system with NVIDIA A100, size 256, 100 iterations, one rank

PerfStubs side note…

- PerfStubs is a “frictionless” instrumentation library
  - https://github.com/UO-OACISS/perfstubs
  - One source file, three headers
  - Provides a plugin interface for performance tools
  - Can be compiled away if desired
- Integrated into several libraries (so far) as a git submodule
  - CAMTIMERS
  - PETSc
  - Ginkgo
  - ADIOS2
  - Others?
- Provides runtime integration with TAU & APEX

Boehme, Huck, Madsen, Weidendorfer,
“The Case for a Common Instrumentation Interface for HPC Codes”
https://doi.org/10.1109/ProTools49597.2019.00010, 2019

Acknowledgements

Parts of this research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

UNIVERSITY OF OREGON