

Trilinos Users Group Data Services Update

November 30, 2021

Karen Devine

with contributions from

Greg Sjaardema (SEACAS/IOSS/Exodus team)

Siva Rajamanickam (KokkosKernels team)

Erik Boman (Zoltan2 team)

Alan Williams (STK team)

SEACAS / IOSS / Exodus

(POC: Greg Sjaardema)

- New Features
 - Assemblies – hierarchical groups of blocks/sets/assemblies
 - Blobs -- store arbitrarily-sized objects in an exodus file
 - Entity Attributes -- “provenance” or annotation data on entities and fields
 - Aprepro – Arrays, Exodus integration
 - Exodus.py – python3, improved capabilities, testing
- New Integrations – FAODEL, Catalyst2, ADIOS2, TextMesh
- In progress:
 - Discontinuous Galerkin Fields
 - HDF5 VOL
 - Compression (lossy and lossless)
- Others: windows, scalability, code quality

Kokkos Core 3.4

(POC: Christian Trott)

- HIP Backend works in previous release
 - Performance Improvements for HIP backend
- SYCL Backend Almost Feature Complete
 - Trilinos Tpetra and LAMMPS work now
 - Ongoing issues: Fine grained tasking, and arbitrary sized atomics.
- OpenMPTarget Backend Almost Feature Complete
 - OpenMPTarget still has issues with some custom reduction stuff in addition.
- Require CMake 3.16 or newer

See Brian Kelley's talk at 11am:
Trilinos Support on AMD and Intel GPUs

Kokkos Kernels 3.4

(POC: Siva Rajamanickam)

- Added full support for HIP
 - Integration with Trilinos is complete and tested on Spock
- Initial support for SYCL
 - Most features functional (coloring algorithms and some batched algorithm still disabled)
 - Integration in Trilinos has started
- Improved support for half precision
- ArmPL library support in specific kernels
- Distance 2 Maximal Independent Set (MIS-2)
- Partial coloring as input for coloring algorithms
- Improved two-stage Gauss-Seidel performance (added damping)
- Supernodal SpTRSV improvements (use batched algorithms internally)

See Brian Kelley's talk at 11am:
Trilinos Support on AMD and Intel GPUs

Zoltan2

(POC: Erik Boman)

- First multi-GPU graph partitioner for distributed-memory systems: Sphynx
 - Sphynx: Spectral Partitioning for HYbrid aNd aXelerator-based systems
 - Uses several Trilinos packages (Tpetra, Anasazi, Ifpack2, MueLu)
 - Uses Kokkos for performance portability
 - Compared to ParMETIS, Sphynx is faster on irregular graphs and obtains similar quality partitions on regular graphs
- New Multi-GPU graph coloring:
 - Distance-1, Distance-2 and Partial Distance-2 coloring
 - Uses Teuchos::Comm and Kokkos to run with MPI+GPU
 - Zoltan2 TpetraCrsColorer interface accepts Tpetra matrix as input

STK

(POC: Alan Williams)

- GPU: Improving the performance of synchronizing Fields between CPU and GPU memory spaces.
 - Primarily for Sierra SM
- AMD/HIP: stk-mesh unit-tests now build and run on AMD platforms, using ROCM 4.3.
 - Primarily for Exawind
- STK Balance: improving work-flow and performance of Balance and BalanceM2N coming soon.