What's New in Tpetra

2025 Trilinos Users Group meeting

Jonathan Hu, for the Tpetra/Performance team

SAND2025-XXXYYY

FY25 Tpetra Team & Collaborators

- Luc Berger-Vergiat (welcome!)

- Tim Fuller

- Jonathan Hu

- Brian Kelley

- Steve Kennon (welcome!)

- Curt Ober (welcome!)

- Carl Pearson

- Chris Siefert (lead)

Updates since TUG`23

Performance monitoring has been improved/extended. See Chris Siefert's talk @ 10:15a on Thursday.

Identified/mitigated performance build issues and performance regressions after package snapshots.

Performance of many-to-few communication (github issue #11803)

      `Igatherv` communication path implemented in Tpetra

      Evaluation underway with Frosch

SpGEMM (#13052)

      More robust handling of CRS matrices, depending on row pointer type

      Improves portability across GPU architectures and vendor TPLs

Improved SpMV performance for HIP (#12852)

Multivector pool, Belos-specific (#13469)

GPU-aware MPI detection logic changed (#12468)

Improved robustness of handling runtime environment variables that start with `Tpetra_` (#12561)

FY25 Plans

Still identifying specific deliverables

Known high-priority items/areas:
- Address identified performance dashboard regressions
- Support ASC codes running on El Capitan / El Dorado and Hops
  - Resolving unified memory issues (MI300a and H100)
- Optimizing Tpetra kernels of interest to ORNL application (Kennon)
- Support LANL's Venado (Grace-Hopper)
- Support Xyce and Alegra's transition to Tpetra
- Kokkos integration performance testing