Meshes, Geometry and Load Balancing Capability Area (Karen Devine, POC)

SAND2014-19229C

Sandia National Laboratories

- Mesh generation
- Mesh management
- Mesh quality improvement
- Load balancing
- Task placement
- Graph coloring
- Matrix ordering

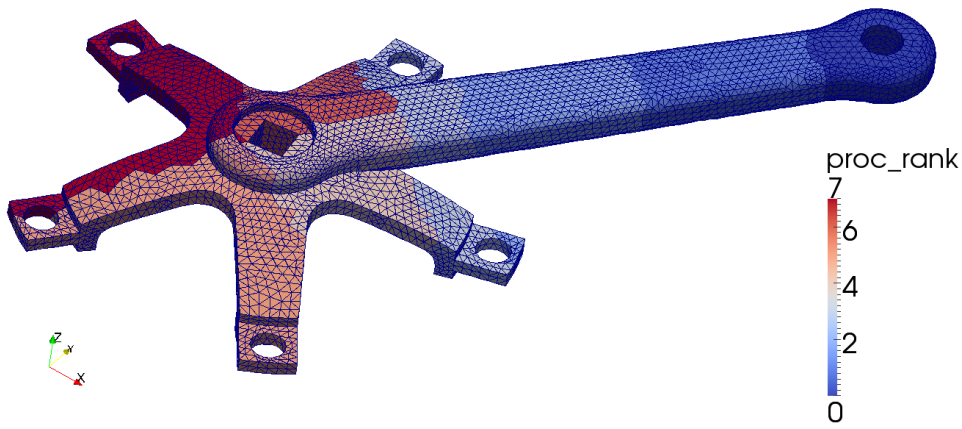

Mesh Generation: PamGen

Sandia National Laboratories

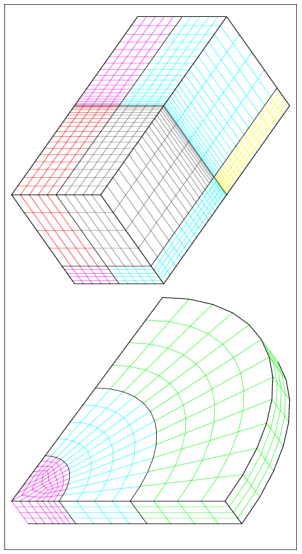

- In-line parallel finite element mesh generation
- POC: David Hensinger, 1443
- Features:
  - Regular meshes on simple geometries (e.g, boxes, cylinders)
  - Meshes generated in parallel
  - Application accesses mesh through Exodus-like interface
  - Great for weak-scaling studies in Trilinos: Generate sequence of meshes on-the-fly; no files needed
- Status: Developed as-needed for ALEGRA
  - Multi-block mesh generation: multiple generated "submeshes" accessed as a single large multi-block mesh
  - Block suppression simplifies specification of complicated geometries

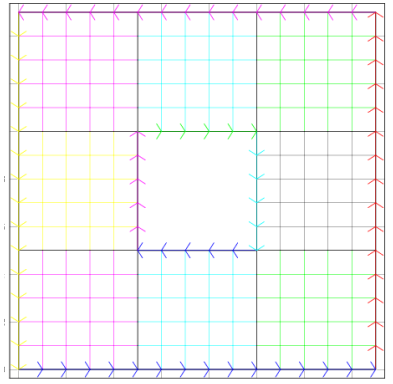

Mesh generation: Meshing Genie

Sandia National Laboratories

- 2D Voronoi Mesh Generation
- POC: Mohamed Ebeida, 1441
- Features:
  - Specify domain with line-segment boundaries
  - Interior holes and fractures are OK.
  - Generates point cloud with user-provided distance function
  - Returns point cloud and mesh information

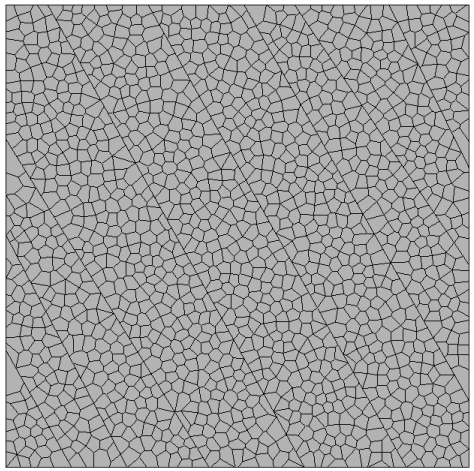

- Status:
  - On hold since 2012
  - ASCR proposal under review to renew development efforts

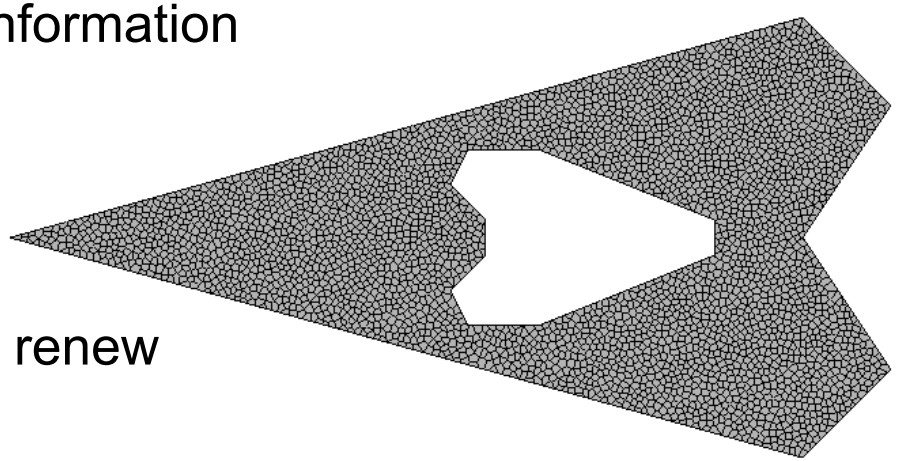

Mesh management: STK

Sandia National Laboratories

- Unstructured mesh toolkit from SIERRA
- POC: Alan Williams, 1543
- Features:
  - Parallel management of unstructured, heterogeneous, dynamic meshes
  - Algorithm support, adaptivity, rebalancing, search, transfer, I/O, refinement, enrichment
- Status:
  - Extensive refactoring and reintegration over past 1+ year
  - *See Alan Williams' talks at 3:30pm Tuesday.*

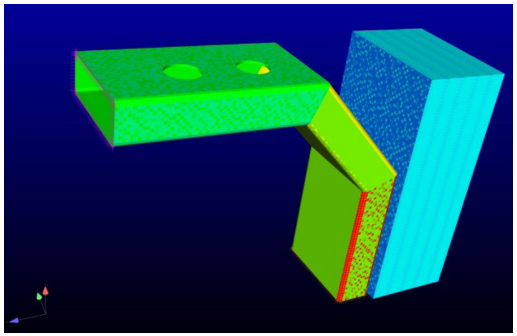

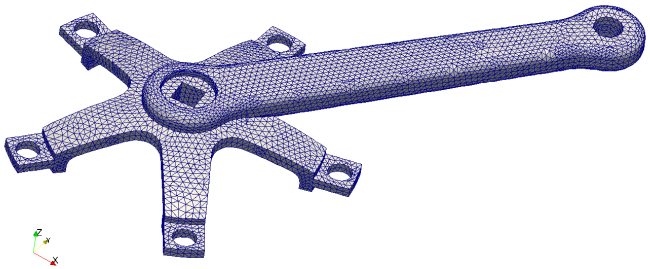

Mesh Manipulation: Mesquite

Sandia National Laboratories

- Mesh Quality Improvement
- POC: Lori Diachin, LLNL
- Features:
  - Optimize the quality of 2D & 3D meshes
  - Shape quality improvement, mesh untangling, alignment with fields, r-adaptivity, boundary mesh smoothing, high-order node smoothing
  - Supports triangles, quads, hexes, tets, pyramids
- Status: Mature; no new development

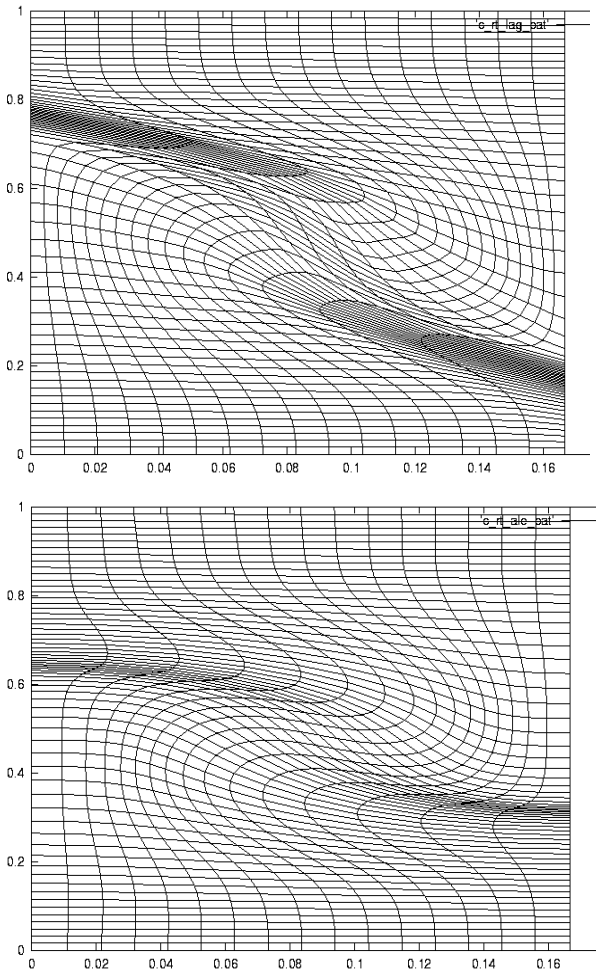

Partitioning, Ordering, Coloring

Sandia National Laboratories

- Zoltan: Parallel combinatorial toolkit
  - POC: Karen Devine, 1426
  - Partitioning and dynamic load balancing
  - Matrix/graph ordering; locality-preserving geometric ordering
  - Graph coloring (distance-1 and -2)

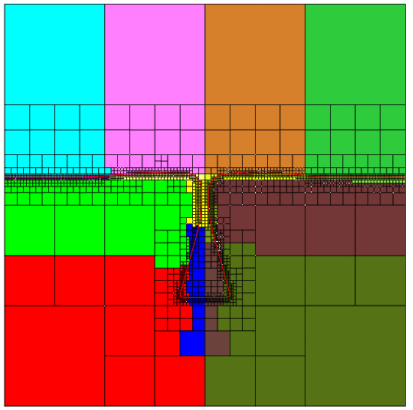

- Isorropia: Combinatorial tools for Epetra data
  - POC: Erik Boman, 1426
  - Epetra-based interface to Zoltan functionality
  - Multi-threaded graph matching, probing, 2D matrix partitioning

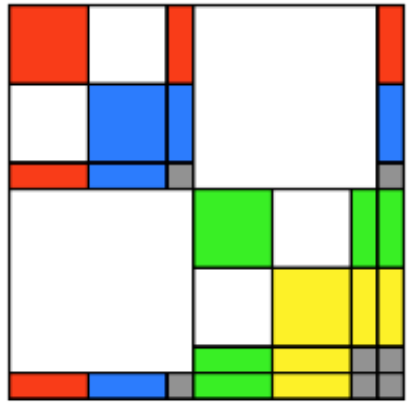

- Status: Mature; maintenance only

Zoltan2:
Next-generation combinatorial toolkit

Sandia National Laboratories

- Team: Boman, Devine, Leung, Rajamanickam, Wolf, Deveci (OSU), Riesen (Ret.)
- Software features:
  - Tighter integration with Trilinos' templated software stack
  - Application-centric user interface
    - Tpetra and Epetra matrix and multivector adapters
    - PamGen mesh adapter
    - Prototype STK mesh adapter
- Status: Platform for new architecture-aware algorithms
  - MPI+X partitioning to load-balance hybrid applications
  - Architecture-aware task placement to reduce network congestion, application communication costs
  - Local graph coloring to detect on-node parallelism
  - Matrix ordering to reduce fill, improve locality

Zoltan2 Partitioning

Sandia National Laboratories

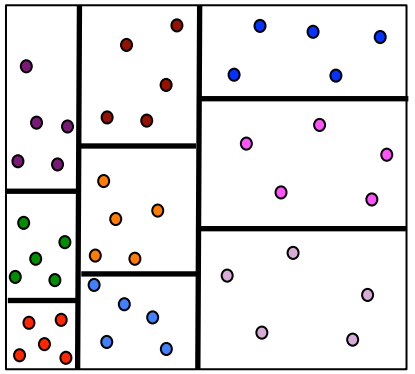

A nine-part 3x3 MJ partition

- Multijagged Geometric partitioning (MJ)
  - Generalization of Recursive Coordinate Bisection
  - Lower data movement ➔better scalability than RCB
  - MPI+OpenMP implementation
  - Used in MueLu to repartition coarse matrices
    - Demonstrated up to 78.4M coordinates, 78K parts on 524K cores (Lin, et al.)
- Graph partitioning interfaces to
  - ParMETIS 4.0.3 (U. Minnesota)
  - PT-Scotch 6.0.3 (parallel, threaded Scotch; U. Bordeaux, INRIA)
- Prototype integration with STK allows new parallel graph-based partitioning in D-to-A toolchain (Bhardwaj et al.)
  - *8B elements on 4K processors in < 5 minutes (including IO)*

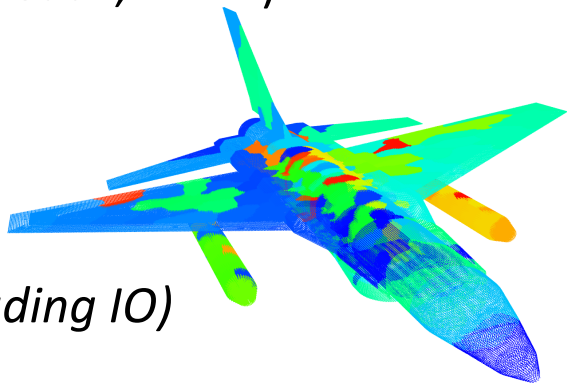

Zoltan2
Architecture-Aware Task Placement

Sandia National Laboratories

- Goal: Map MPI ranks (tasks) to allocated cores so that network congestion and communication costs are low
- Geometric task placement: Uses MJ partitioner to assign interdependent tasks to "nearby" cores in an allocation
- Adopted by MueLu to reduce multigrid solve time
  - Prokopenko, Boman
  - Geometric task mapping applied at fine-matrix level
  - Bipartite graph matching reduces communication between fine operator (on all cores) and coarse operator (on subset of cores)

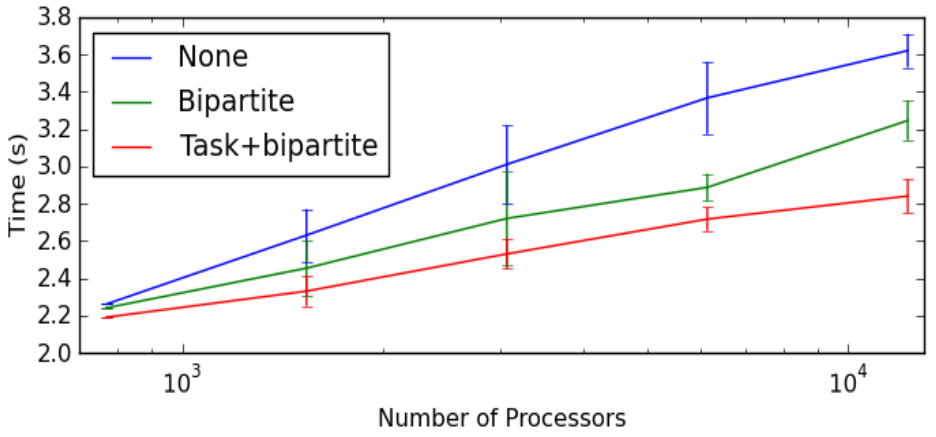

Weak scaling of MueLu on NERSC Hopper
Time for one multigrid solve
(Courtesy of A. Prokopenko)

Zoltan2 Graph Algorithms

Sandia National Laboratories

- Local graph coloring
  - Assign colors to vertices so that adjacent vertices have different colors
  - Can identify independent vertices for in-node parallelism
  - Serial implementation available; Kokkos-based multithreaded implementation in development
  - Coming soon: "Balanced" coloring for use on GPUs; coloring with roughly equal number of vertices per color.
    - *Looking for friendly first users; talk to Erik Boman*

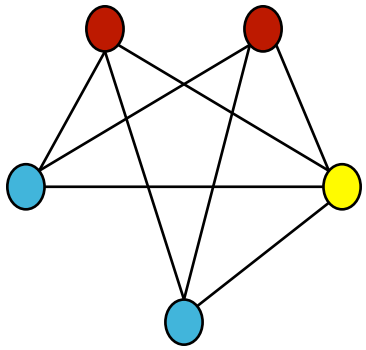

- Local matrix/graph ordering
  - Reorders data for reduced fill, increased locality
  - RCM, Sorted-degree, Random, AMD-interface

Package Summary

Sandia National Laboratories

- Mesh Generation
  - PamGen (Dave Hensinger, 1443)
  - Meshing Genie (Mohamed Ebeida, 1441)
- Mesh storage and manipulation
  - STK *(Alan Williams, 1542; Tuesday 3:30pm)*
  - Mesquite (Lori Diachin, LLNL)
- Combinatorial algorithms
  - Zoltan/Zoltan2 (Karen Devine, 1426)
  - Isorropia (Erik Boman, 1426)

- Related talks:
  - Exodus, IOSS, and SEACAS *(Greg Sjaardema, 1543; Tuesday 3:10pm)*
  - Mesh Adaptivity Using Percept and STK *(Brian Carnes, 1544; Tuesday 4pm)*