Unstructured Meshing Techniques

FASTMath Team Members: B. Ball, B. Bihari, K. Devine, L. Diachin, D. Ibanez, K.E. Jansen, J. Jiao, V. Leung, Q. Lu, V.S. Mahadevan, O. Sahni, E.S. Seol, M.S. Shephard, C.W. Smith and P. Wilson

Unstructured meshes, often adaptively defined, can yield required levels of accuracy using many fewer degrees of freedom at the cost of more complex parallel data structures and algorithms. FASTMath is providing the parallel unstructured mesh data structures and services needed by developers of PDE solution procedures targeted for exascale computers

Unstructured Meshes for Exascale Computers

Key unstructured mesh technologies needed by applications

- Effective parallel mesh representations
- Base parallel mesh functions
  - Partitioned mesh control and modification
  - Read only copies of types needed
  - Associated data, grouping, etc.
- Key services
  - Load balancing (separate poster)
  - Mesh-to-mesh solution transfer
  - Mesh optimization and adaptation

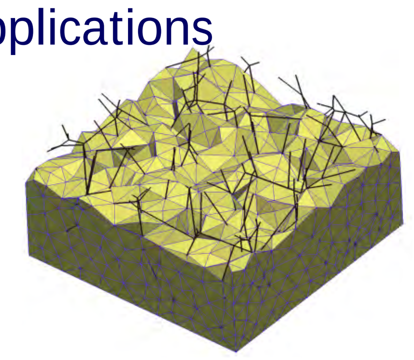

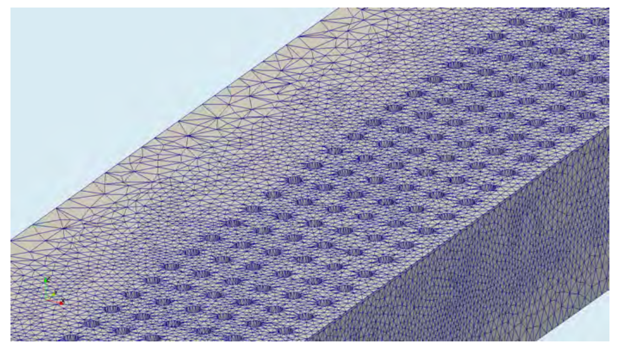

Poster highlights recent FASTMath developments

Parallel Mesh Infrastructures

Infrastructures to address the needs of SciDAC applications

SIGMA(CGM/MOAB/Lasso/MeshKit) developments support

- Both open-source (OCC) and commercial (ACIS) geometry modeling engines
- Scalable mesh (data) usage in applications on >32K cores through efficient array-based memory access
- Parallel HDF5-based I/O, visualization
- Unified interfaces to external meshing packages (MeshKit)
- Components for discretization and mesh-to-mesh coupling

PUMI developments

- Support for adaptively evolving mixed meshes
- Combined use of MPI and threads
- Meshes to 92B elements on ¾ million parts
- Support of model dimension adaptation of M3D-C1 (underway)
- Array-based mesh representation to improve memory efficiency and better support new architectures (underway)
  - 4x memory savings, 250 bytes per tetrahedron for a full rep.

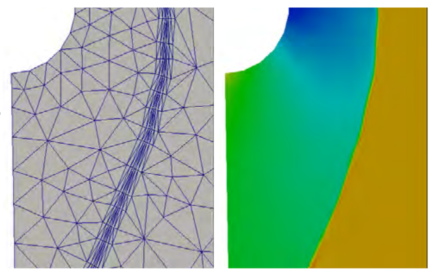

Attached Parallel Fields (APF) development underway

- Effective storage of solution fields on meshes
- Supports operations on the fields
- Support adaptive expansion of Fields from 2D to 3D in M3D-C1 (under development)
- History-dependent integration point fields for Albany plasticity models

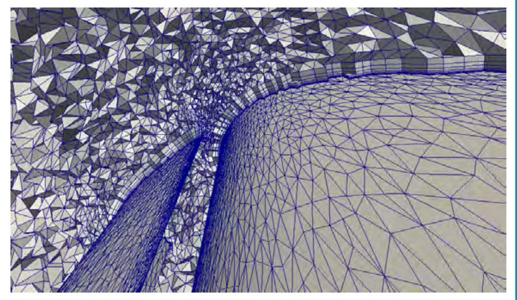

Mesh Improvement and Mesh Adaptation

Key packages

- Parallel anisotropic adaptation for a broad range of applications (MeshAdapt)
- Optimizes element shapes using target-matrix quality metrics (Mesquite)

Recent advances

- Mixed mesh adaptation for boundary layers
- Adapting meshes to 92 billion elements
- Boundary layer thickness adaptation
- Scaling of mesh quality improvement (Mesquite) to 125,000 cores
- In-memory integration of MeshAdapt into multiple simulation workflows
- Higher than quadratic curved mesh geometry with G1 surface continuity (partly complete)

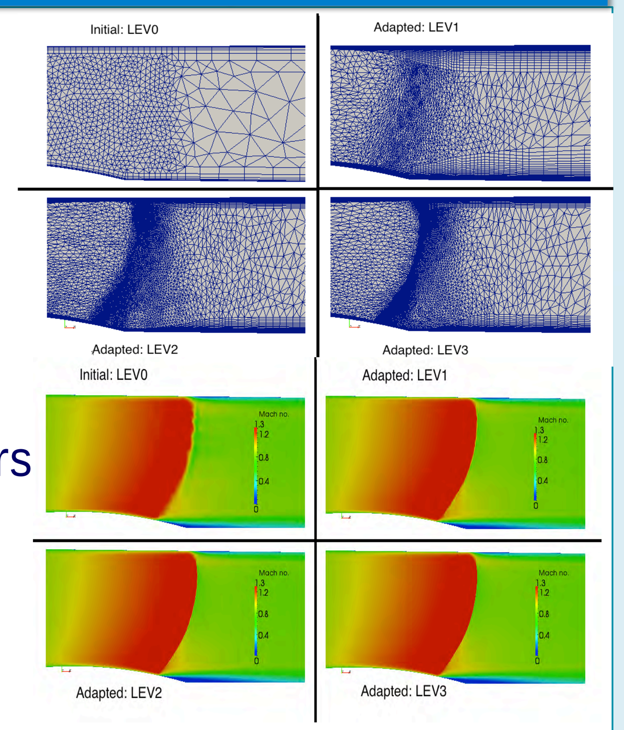

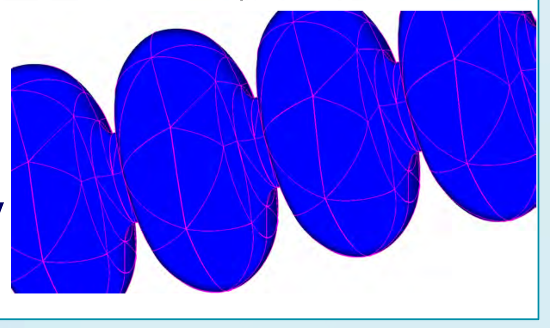

Architecture Aware Developments

Goal: provide tools and methods for operating on unstructured meshes that effectively use high core-count, hybrid parallel compute nodes

- A parallel control utility (PCU) that supports hybrid operation on unstructured meshes
- Hybrid partitioned meshes (PUMI)

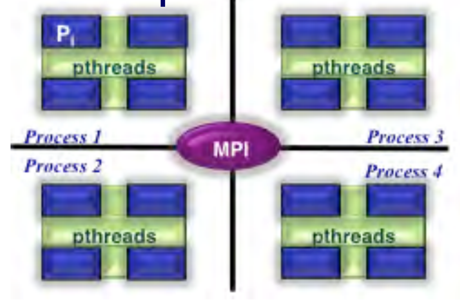

Results

- 16 threads per process on BG/Q saves 1GB of memory
  - Critical for many-core nodes where memory/core is limited
- Initial testing on TACC Stampede with running natively on Intel Phi
  - 25M element mesh
  - Equal number of BGQ and Stampede nodes
  - 2048 -> 4096 partitioning on Stampede is 8% slower than BGQ
  - 1024->2048 partitioning is 40% faster on Stampede!
- There is ongoing work in MOAB to support multi-threaded mesh entity traversals for efficient FEM kernel assemblies

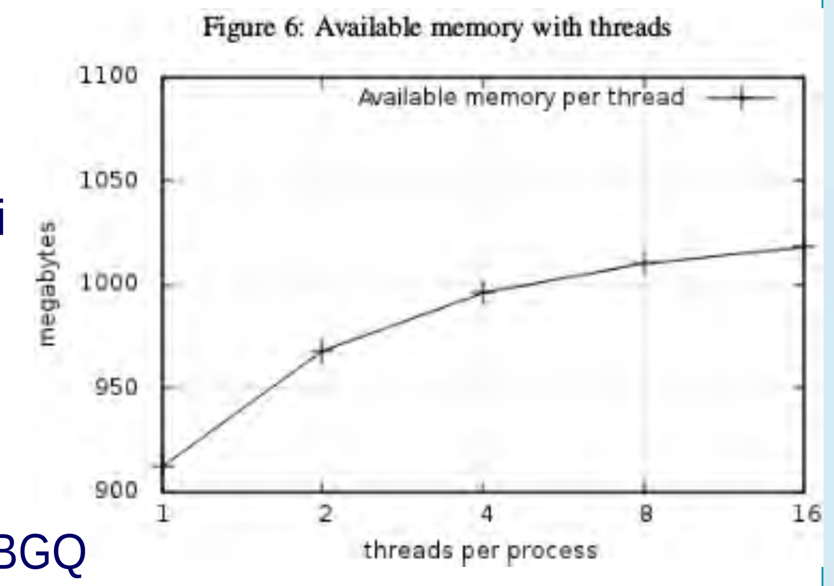

Parallel Point Location for Solution Transfer

Goal: Simplify geometry searches and unify discretization kernels

- Geometry search: parallel point-in-element query for various element topologies (edge,tri/quad/polygon,tet/hex/prism/pyramid)
- Discretization: support transformations, different higher-order basis functions (lagrange, spectral) for optimized local FE/FV

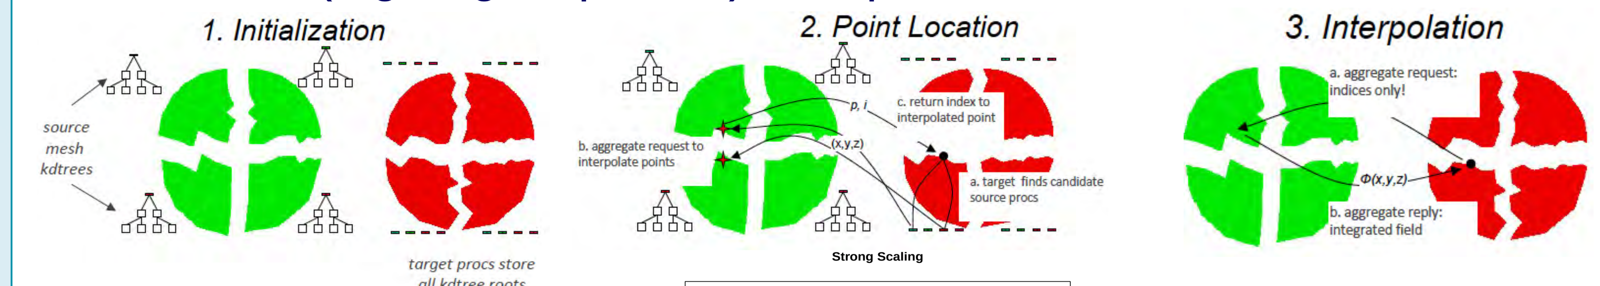

4. Global normalization

5. Subset Conservation

Bottleneck: kd-tree scales as O(nlog(n)). Consider BVH/BIH tree data-structures to attain O(log(n)) time complexity.

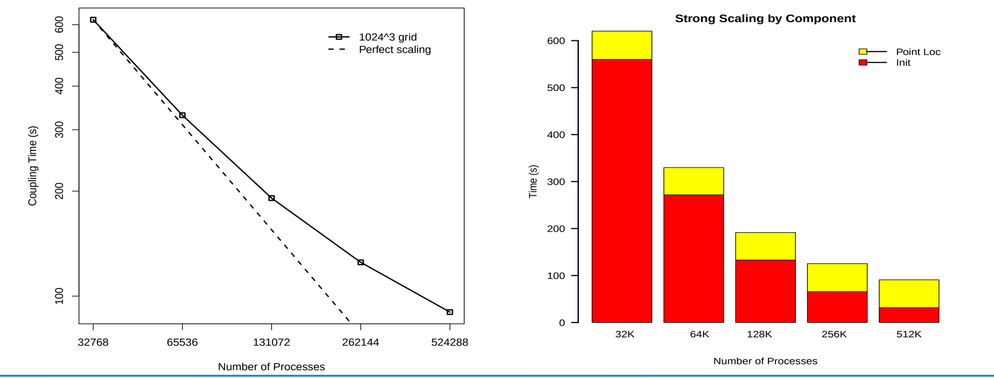

Applications List

- Accelerator Modeling (ACE3P)
- Climate data analysis (Par-NCL)
- Conservative multi-tracer transport (MBCSLAM)
- FE-based neutron transport (PROTEUS)
- Fluid/Structure interaction (AthenaVMS)
- Fusion Edge Physics (XGC)
- Fusion first wall chemistry & dynamics (XOLOTL)
- Fusion Plasmas (M3DC1)
- High-order CFD on (Nektar++)
- High-speed viscous flows (FUN3D)
- Monte Carlo neutron transport (DAG-MCNP)
- Mortar element Structural Mechanics (Diablo)
- Multiphase reactor flows (PHASTA)
- SEM-based CFD (Nek5000)
- Solid Mechanics (Albany)

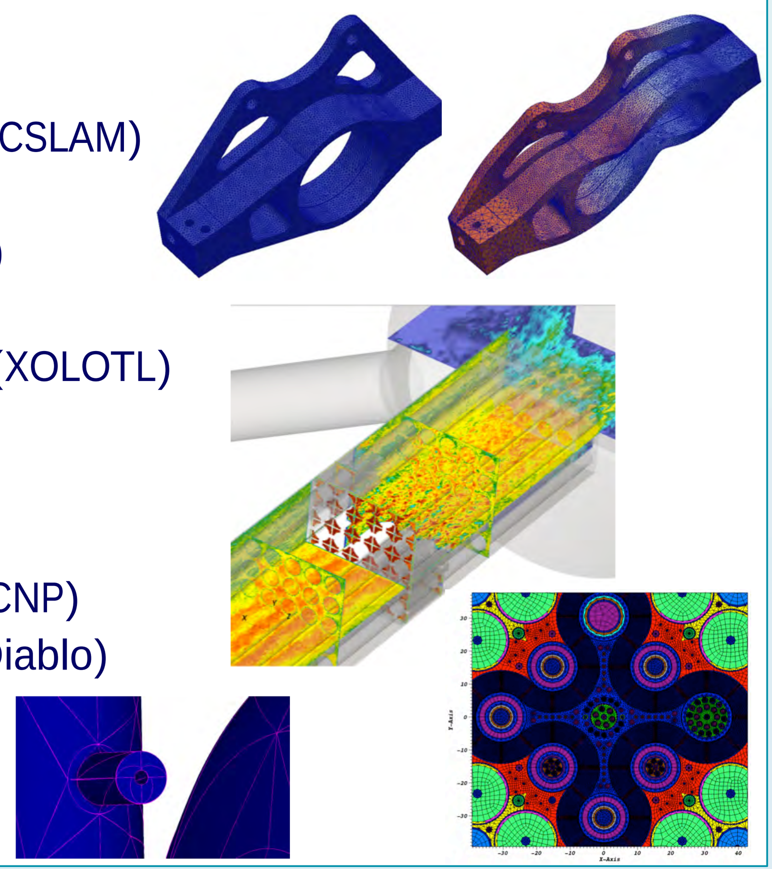

Future Plans

- Continue to improve mesh representations for memory use
- Software thread affinity control for many-core nodes
- Support combined unstructured mesh/PIC methods
- Extend and expand services to meet the needs of applications
- Complete mesh representation for higher than quadratic elements
- Continued development of fields and solution transfer methods

More Information: http://www.fastmath-scidac.org or contact Lori Diachin, LLNL, diachin2@llnl.gov, 925-422-7130