Interoperable Mesh Tools for Petascale Applications

Lori Diachin, LLNL

Representing the ITAPS Team

ITAPS focuses on interoperable meshing and geometry services for SciDAC

- ITAPS Goal
  - Improve SciDAC applications’ ability to take advantage of state-of-the-art meshing and geometry tools
  - Develop the next generation of meshing and geometry tools for petascale computing
- Technology Focus Areas
  - Complex geometry
  - High quality meshes and adaptivity
  - Coupled phenomenon
  - Dynamic mesh calculations
  - Tera/Petascale computing

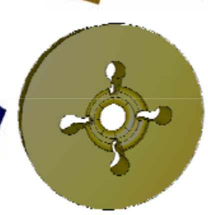

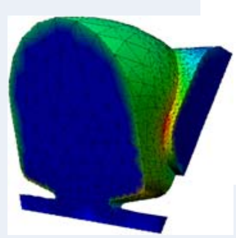

Accomplishing the ITAPS interoperability goal requires a strong team with diverse expertise

Our senior personnel are experts in complex geometry tools, mesh generation, mesh quality improvement, front tracking, partitioning, mesh refinement, PDE solvers, and working with application scientists

SciDAC Scientific Discovery through Advanced Computing

3

ITAPS

The Challenge

- Pre-existing ITAPS tools all meet particular needs, but
  - They do not interoperate to form high level services
  - They cannot be easily interchanged in an application
- In general the technology requires too much software expertise from application scientists
  - Difficult to improve existing codes
  - Difficult to design and implement new codes

The ITAPS center recognizes this gap and is addressing the technical and human barriers preventing the creation and use of advanced, interoperable mesh and geometry services

SciDAC Scientific Discovery through Advanced Computing

4

ITAPS

ITAPS will achieve its long term technology goal using a hierarchical approach

- Build on successes with SciDAC-1 applications and explore new opportunities with SciDAC-2 application teams
- Develop and deploy key mesh, geometry and field manipulation *component services* needed for petascale computing applications
- Develop advanced functionality *integrated services* to support SciDAC application needs
  - Combine component services together
  - Unify tools with *common interfaces* to enable interoperability

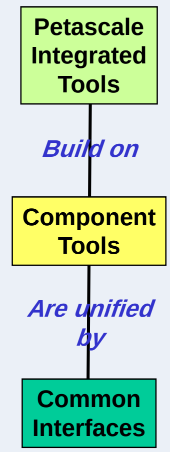

SciDAC Scientific Discovery through Advanced Computing

ITAPS

The ITAPS team successfully deployed new mesh and geometry tools in SciDAC-1

- Prototype integrated services
  - AMR-Front tracking
  - Design optimization
- Component services
  - FronTier front tracking
  - Mesquite mesh improvement
  - Adaptive loop infrastructure
- Mesh and geometry access through common interfaces
  - Mesh: MOAB, AOMD, NWGrid, GRUMMP
  - Geometry: CGM

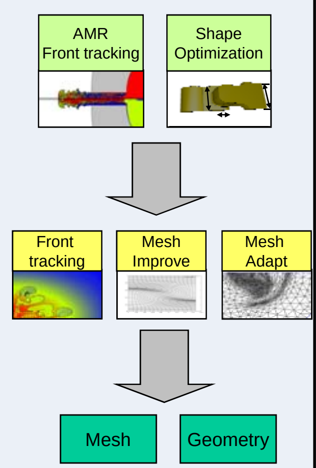

SciDAC Scientific Discovery through Advanced Computing

ITAPS

Status of the ITAPS services and their use of the TSTT interfaces

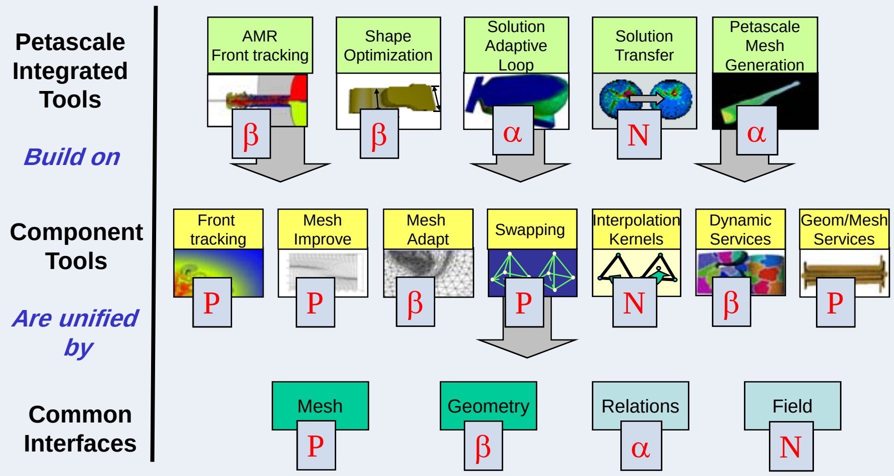

SciDAC
Scientific Discovery through Advanced Computing

7

The ITAPS interoperability goal requires abstracting the data model and information flow

- Information flows from geometrical representation of the domain to the mesh to the solvers and post-processing tools
- Adaptive loops and design optimization requires a loop

- A *data model* that encompasses a broad spectrum of mesh types and usage scenarios
- A set of *common interfaces*
  - Implementation and data structure neutral
  - Small enough to encourage adoption
  - Flexible enough to support a broad range of functionality

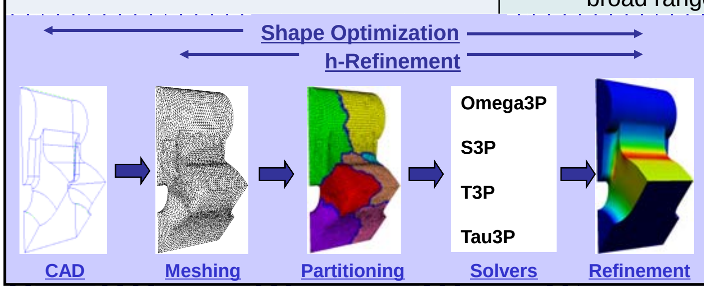

The ITAPS data model abstracts PDE-simulation data hierarchy

- Core Data Types
  - 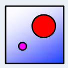 *Geometric Data*: provides a high leve boundaries of the computational dom mesh data
  -  *Mesh Data*: provides the geometric a associated with the discrete represer
  - 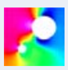 *Field Data*: provides access to the tir variables associated with application solution. These can be scalars, vectors, tensors, and associated with any mesh entity.

Each core data type has an ITAPS interface
- Mesh: iMesh
- Geometry: iGeom
- Fields: iField
- Relations: iRel

- Data Relation Managers
  - Provides control of the relationships among the core data types
  - Resolves cross references between entities in different groups
  - Provides functionality that depends on multiple core data types

The building blocks of the ITAPS interfaces consists of entities, entity sets, and tags

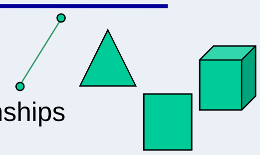

- Entity Definition
  - Unique type and topology
  - Canonical ordering defines adjacency relationships
- Entity Set Definition
  - Arbitrary grouping of TSTT mesh entities
  - There is a single “Root Set”
  - Relationships among entity sets
    - Contained-in
    - Hierarchical

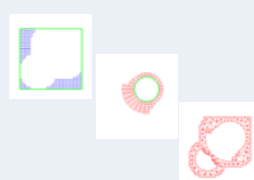

- Tags allow user-defined data association with entities and entity sets
- Blend of abstract concepts and familiar mesh/geometry specifics

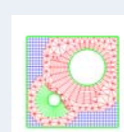

An overarching philosophy guides the interface definition efforts

- Maintain data structure neutrality
- Create a small set of interfaces that existing packages can support
  - Small set of 'core' functions that must be implemented
  - Larger set of functions supported by reference implementations
- Balance performance and flexibility
- Work with a large tool provider and application community to ensure applicability
- Target both specific uses of ITAPS services and broad dissemination as CCA components

Lowers the burden for adoption of the interface

Performance is critical for kernel computations involving mesh access; flexibility is critical for covering a broad usage spectrum

CCA provides infrastructure and guidance for domain-specific interface definition efforts

SciDAC Scientific Discovery through Advanced Computing

11

ITAPS

Interface implementations are well underway

- TSTT Mesh 0.7 Interface complete
- Geometry and relations interfaces well on their way
- Preliminary field data interface
- Implementations
  - Mesh: FMDB, GRUMMP, NWGrid, MOAB, Frontier
  - Geometry: CGM
  - Relations: Lasso
- C, C++, and Fortran language interoperability through a C and SIDL/Babel (CCA) interfaces
  - Analyzing performance ramifications of SIDL/Babel language interoperability tools (joint with the CCA)
- Interfaces stable enough to build services upon them and test interoperability

SciDAC Scientific Discovery through Advanced Computing

12

ITAPS

Achieving good performance requires understanding the memory/time tradeoffs

- Implementations of iMesh vary on speed vs. memory performance
  - Create, v-E, E-v query, square all-hex mesh
  - Entity- vs. Array-based access
- Compare iMesh (C, SIDL), Native (MOAB), Native Scd (MOAB), CUBIT
  - Ent-, Arr-based access
  - All-hexahedral square structured mesh

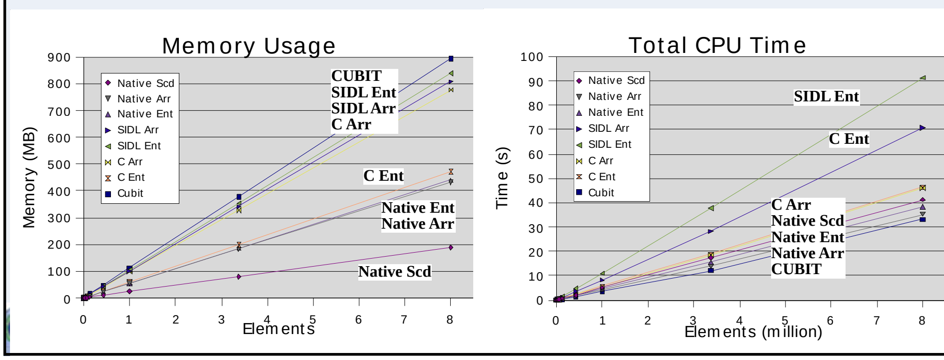

ITAPS services build on the interface definitions

- Mesh and Geometry Services
- Mesquite Mesh Quality Improvement
- FronTier Front Tracking
- Zoltan partitioning
- Adaptive Loops

SciDAC Scientific Discovery through Advanced Computing

14

ITAPS

Mesh/Geom services support operating on data from multiple sources

- Goals: Support manipulation of meshes by applications
  - Mesh topology
  - Mesh associations and groupings
  - Meshes across parallel computers
  - Interface to multiple mesh generation/management technologies
    - FMDB, NWGrid, MOAB,GRUMMP
- Current work
  - Parallel search and sort for mesh-to-mesh transfer
  - Advanced access/relations functionality

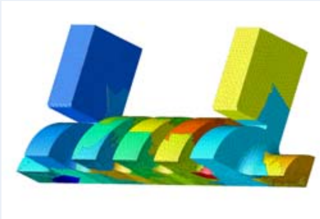

SciDAC Scientific Discovery through Advanced Computing 15 ITAPS

Mesquite provides mesh quality improvement and r-adaptive services

- Goals: to provide a portable, stand-alone software toolkit to address mesh quality improvement
  - A priori shape, size, alignment improvement
  - A posteriori solution feature capturing or error reduction
- Used in many SciDAC applications including accelerator design, climate, mesh generation
- Current work
  - Parallelization
  - R-adaptivity
  - Combination with swapping
  - Design optimization

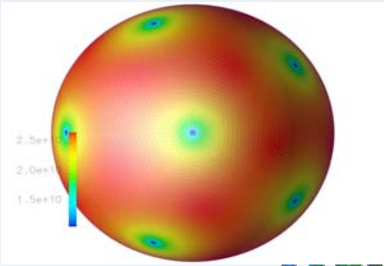

SciDAC Scientific Discovery through Advanced Computing 16 ITAPS

FronTier provides software for tracking sharply defined interfaces

- Goals of FT-Lite
  - Provide an easy-to-use front tracking software library
  - Hide geometrical and topological operations from users
- Three levels of libraries
  - Static interface library
  - Dynamic front library (*FT-Lite)*
  - Physics dependent libraries (gas, res, solid etc)
- Application to SciDAC problems
  - Fluid flow instabilities
  - Groundwater plume tracking
  - Pellet injection for tokamak design

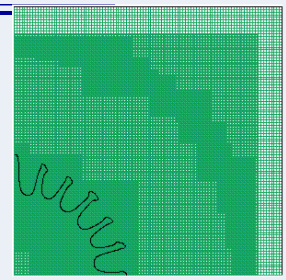

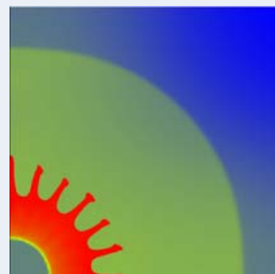

Dynamic load balancing and partitioning via the Zoltan toolkit (w/ CSCAPES)

- Goal: Reduce total execution time by
  - Distributing work evenly among processors
  - Reducing applications' interprocessor communication
  - *Keeping data movement costs low*
- Important in many SciDAC technologies including adaptive mesh refinement, parallel remeshing
- Current Work
  - Target emerging LCF platforms
    - Expand scalability studies to petascale
    - Work in application context
  - Consider a variety of partitioners
  - Develop highly efficient support technologies
    - Data migration from old to new partitions
    - Data mapping from one mesh to another

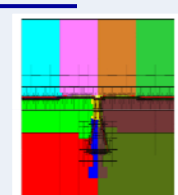

Adaptive Mesh Refinement

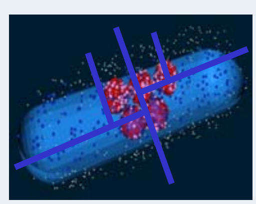

Particle Simulations

ITAPS is developing stand-alone adaptive loops for SciDAC applications

- Goal: To increase the availability of adaptive loops for simulations originally designed based on fixed grids
- Implementation Options
  - *Tightly coupled* using a single set of structures
    - Advantage is optimal and efficient if done well
    - Disadvantage is complex algorithm and code development
  - *Loosely coupled* building on existing components
    - Takes advantage of existing analysis s and adaptive tools
    - Disadvantage is the overhead of multiple structures and data conversion
- Strong need for this capability in SciDAC applications
  - Loosely coupled procedure used to improve SLAC design code
  - Tightly coupled procedure being implemented for fusion MHD

SciDAC Scientific Discovery through Advanced Computing

19

ITAPS

Adaptive unstructured mesh methods have both advantages and disadvantages

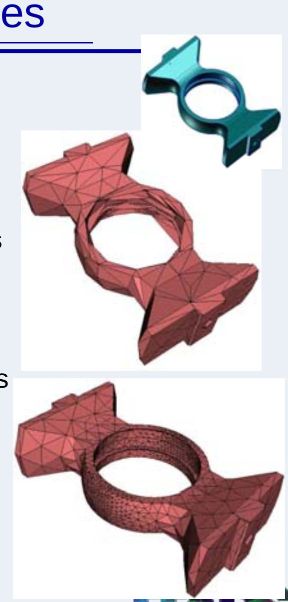

- Advantages
  - Meshes of mixed topologies and order easy
  - Commonly used spatial decomposition for finite element discretizations
  - Mesh adaptation can account for curved domains
  - General mesh anisotropy can be obtained
  - Easy to create strong mesh gradations without special numerical techniques
  - Alignment with multiple curved geometric features
- Disadvantages
  - Data structures larger and more complex (memory/coding time)
  - Solution algorithms can be more complex (CPU time)

SciDAC Scientific Discovery through Advanced Computing

20

ITAPS

Anisotropic adaptation techniques show promise for orders of magnitude improvement

- Area of active research
  - Methods to determine optimal anisotropic meshes
  - General mesh modification based on anisotropic mesh size field
  - Extensions for anisotropic adaptation of mixed meshes
- Results to date demonstrate
  - Isotropic mesh adaptation typically provides order of magnitude reduction in mesh size
  - Anisotropic adaptation typically provides another order of magnitude reduction
    - Isotropic adaptation selects $h=\min(h_1, h_2, h_3)$,
    - Unstructured anisotropic adaptation matches h_i and uses one element where isotropic uses $h_3/h_1 * h_2/h_1$ elements
    - Navier Stokes: Aspect ratios vary from $O(10^2)$ to $O(10^6)$,
    - Plasma Physics: Similar aspect ratios in ITER class plasma physics

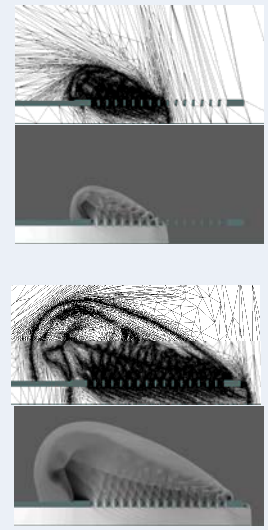

Steps to Petascale Adaptive Simulations on Unstructured Meshes

- Steps
  1. Be sure the fixed mesh solver scales to 100,000's of processors
  2. Provide parallel distributed support for mesh adaptation
  3. Construct adaptive loops in which all components run on petascale machines
- Status Summary
  - Good progress on 1. with an implicit FE flow code
  - A preliminary set of tools for supporting parallel mesh adaptation including dynamic load balancing
  - Constructing initial parallel adaptive loops

Step 1: Taking a Solver to Petascale: Progress with PHASTA

- PHASTA - Parallel Hierarchic Adaptive Stabilized Transient Analysis
  - Parallel finite element NS flow solver
  - Compressible and incompressible laminar and turbulent flow
  - **Implicit time integration -** requires solution of very large linear algebra systems at each time step using iterative solvers
  - Employs a domain decomposition scheme that breaks the total domain into subdomains with roughly the same number of elements (3% variation typical)
- PHASTA's work can be characterized as requiring:
  - Organized, substantial, and regular communication between partitions that touch each other
  - A specific number of ALL-REDUCE communications are also required

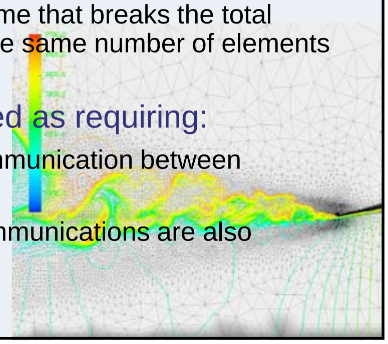

SciDAC Scientific Discovery through Advanced Computing

PHASTA performance on Blue Gene L shows excellent strong scaling

- Excellent strong scaling on Blue Gene's at IBM and Rensselaer (# 7 on June 07 top 500 list)
- Blue Gene communication fabric critical to obtaining these results

5M vertex mesh

# Proc.	t (sec)	scale
8192	16.6	0.952
4096	32.3	0.978
2048	64	0.988
1024	126	1.0
512	252	1.0

18.5M vertex mesh

# Proc.	t (sec)	scale
16384	60.6	1.04
8192	131.7	0.957
4096	241.6	1.04
2048	502.3	1.00
1024	1008.7	1.00

SciDAC Scientific Discovery through Advanced Computing

ITAPS

Parallel adaptive simulation loop under development at RPI

- Required Operations
  - Links across processor boundaries
  - Dynamic parallel partitioning
  - Mesh migration
  - Predictive balancing
  - Parallel mesh adaptation
- Parallelization of refinement
  - Operations on processor
  - Synchronize the partition data
- Parallelization of coarsening and swapping
  - On processor - performed
  - Off processor - held
  - Migration to get operation on processor

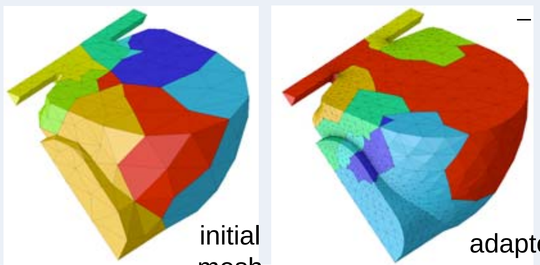

initial mesh

adapted mesh

SciDAC Scientific Discovery through Advanced Computing

25

ITAPS

ITAPS technologies impact SciDAC applications in three ways

1. Direct use of ITAPS technology in applications
   - Geometry tools, mesh generation and optimization for accelerators and fusion
   - Mesh adaptivity for accelerators and fusion
   - Front tracking for astrophysics and groundwater
   - Partitioning techniques for accelerators and fusion
2. Technology advancement through demonstration and insertion of key new technology areas
   - Design optimization loop for accelerators (w/ TOPS)
   - Petascale mesh generation for accelerators
3. Enabling future applications with ITAPS services and interfaces
   - Parallel mesh-to-mesh transfer for multi-scale, multi-physics applications
   - Dynamic mesh services for adaptive computations

SciDAC Scientific Discovery through Advanced Computing

26

ITAPS

Applications can access ITAPS services in two ways

1. Implement ITAPS interfaces on top of application data structures
2. Use a reference implementation of the interfaces to provide access to ITAPS services at the cost of a data copy

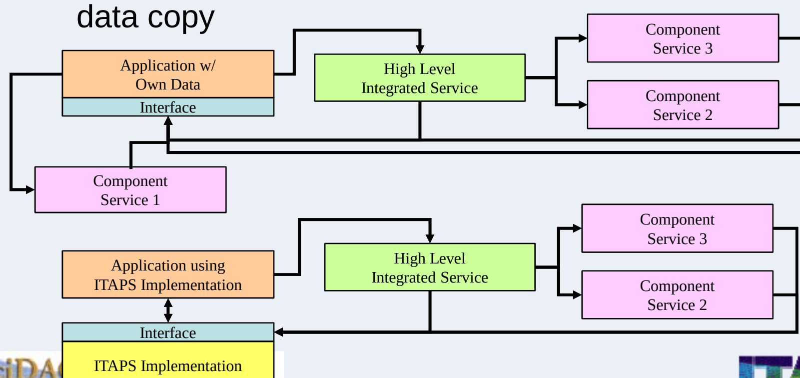

Example 1. Inserted ITAPS technology to enhance accelerator design

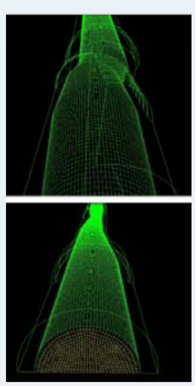

- Custom-built meshes with attention to quality requirements
  - Resulted in recent decision to use Tau3P as vehicle for further PEP-II IR design studies
  - Enabled the first wakefield analysis of a Damped Detuned Structure and direct verification of DDS wakefield suppression by end-to-end simulation
- Created an adaptive mesh simulation capability
  - Accurately predict field quantities that influence wall losses.
  - Used 1/3 the number of unknowns and were found to give the most accurate results
- Working with SLAC/TOPS to build a shape optimization capability

Components in adaptive loop for SLAC

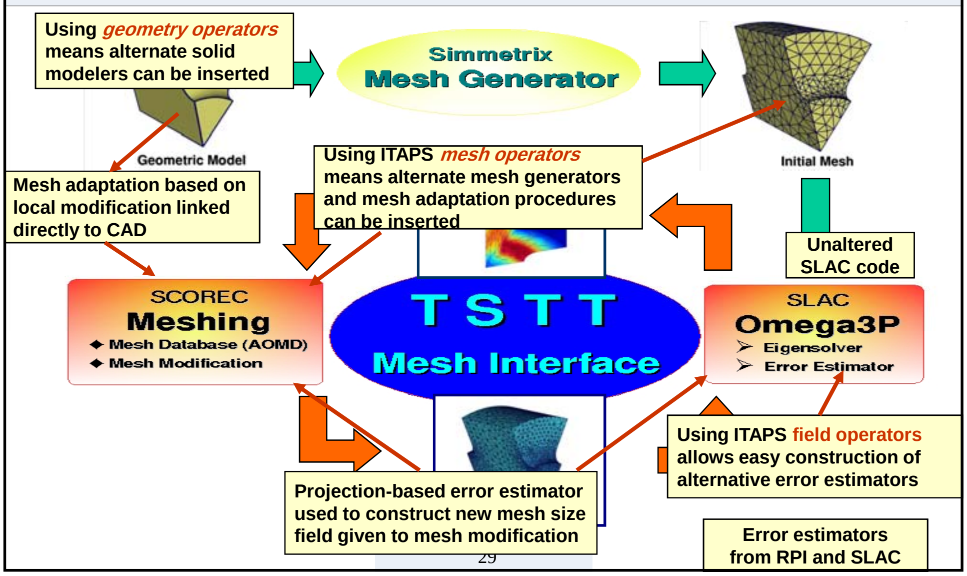

Currently working to provide parallel adaptivity for the M3D-C¹ Extended MHD Code

- Modified M3D-C¹ to use ITAPS unstructured meshes
  - Interface to unstructured mesh search algorithm
  - Efficient use of C¹ continuous shape functions for unstructured meshes
  - Reduced interprocessor communications
- Created interface with ITAPS mesh adapt software
  - Error indicator for C¹ element developed
  - Local field transfer based on interpolation
  - Constrained mesh adaptation to take into account mesh matching on periodic boundaries

SciDAC Scientific Discovery through Advanced Computing

30

Preliminary results for 2D domains are promising

Toroidal equilibrium problem that converges to steady state

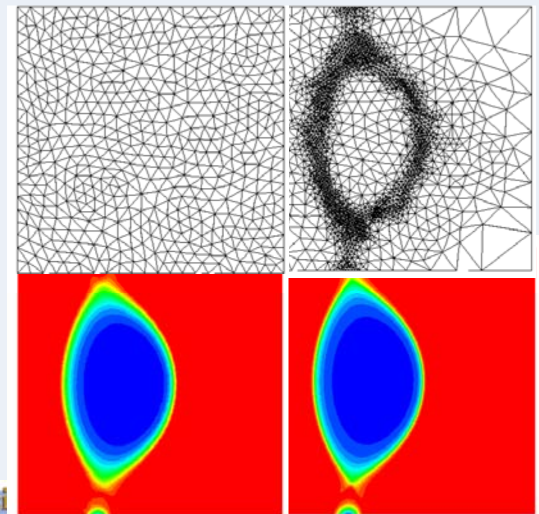

Currently working on anisotropic adaptivity capabilities

Example 2. Demonstrated optimal mesh generation and adaptive methods for climate

- *Goal:* Given an initial isotropic or anisotropic planar or surface mesh and a solution field with large gradient mountain heights, use solution based r-adaptation to minimize solution error
- *Progress*:
  - Several different meshing strategies developed for structured and unstructured surfaces
  - Proof of principle of meshing technologies demonstrated, integrated in next generation climate codes
  - Improving the prediction of rainfall, snowfall and cloud cover in regional weather models

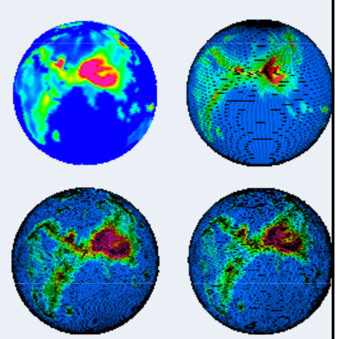

Orography field showing high altitude over the Himalayas and alps; Structured adapted spherical mesh, hybrid geodesic mesh, and unstructured mesh based on orographic field data

Example 3. ITAPS technologies used to build next generation biology tools

- *Goal*: To understand the behavior of Shewanella microbe flocs in oxygen rich environments
- Collaborated on the development of the Virtual Microbial Cell Simulator
  - Floc geometry built using image reconstruction techniques from a stack of confocal images
  - Solve reaction-diffusion equations to find the concentration of oxygen in the floc
  - Used many ITAPS technologies in concert, e.g, NWGrid, Mesquite, Frontier

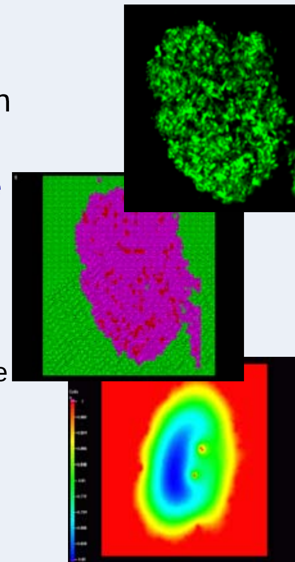

SciDAC Scientific Discovery through Advanced Computing

33

ITAPS will continue its tradition of strong partnerships with application teams

- Accelerator Modeling
  - Continued work on parallel adaptive refinement loops
  - High order mesh generation and fix up
  - Shape determination and optimization (w/ TOPs) including petascale mesh generation and improvement
  - High order embedded boundary methods for Maxwell's equations in complex geometry
- Fusion
  - Parallel adaptive loops and high order finite elements
  - Front tracking technologies for subgrid models of external sources in tokomaks (pellet fueling and plasma disruption)
- Groundwater/Subsurface Flow
  - Explore mesh generation, mesh quality improvement, and AMR- front tracking techniques

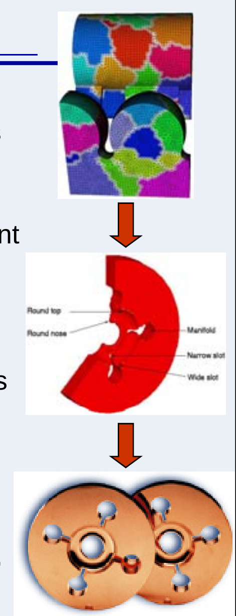

SciDAC Scientific Discovery through Advanced Computing

34

ITAPS Software Web Pages

http://www.itaps-scidac.org/software/

- Provides help getting started
- Usage strategies
- Data model description
- Access to interface specifications, documentation, implementations
- Access to compatible services software

Interface Software Access

- Links to the interface user guides and man pages where available
- Links to the C-binding and SIDL-binding files
- Links to implementations for iMesh, iGeom, iRel
  - Version 0.7 compatible software
  - Links to the home pages for more information
- Simple examples, compliance testing tools and build skeletons coming soon

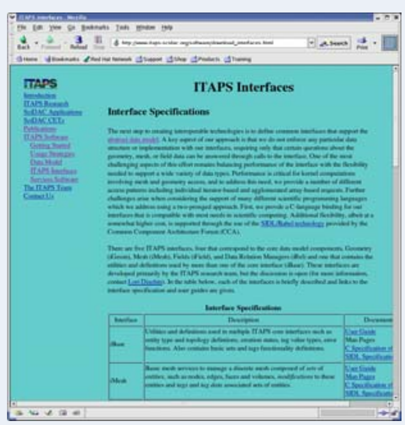

Services Software Access

- Links to the services built on the ITAPS interfaces
- Currently or very soon to be available
  - Mesquite (C, SIDL)
  - Zoltan (C, SIDL)
  - Swapping (SIDL)
  - Frontier (SIDL)
  - VisIt Plug In (C, SIDL)
- Links to home pages for more information
- Instructions for build and links to supporting software

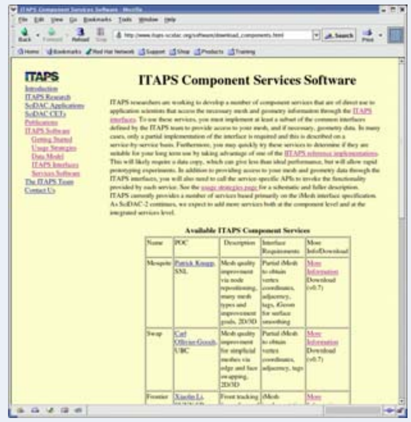

SciDAC Scientific Discovery through Advanced Computing

37

Where we're going…

- Complete interface definition work
  - Multiple implementations of ITAPS interfaces
  - Evaluate design choices in the data model; revise as appropriate
  - Continue to evaluate performance of language interoperability tools (w/ CCA)
- Component and integrated service development and interoperability
  - PDE-based applications (partitioning, mesh-to-mesh transfer, adaptive loops, error estimation, geometry modification, VACET visualization tools)
  - Continued development of ITAPS tools such as Mesquite, Frontier, Zoltan
  - Release CCA-compliant ITAPS services
- Demonstrate interoperability in SciDAC applications
  - SLAC use of adaptive meshing technologies on complex geometries
  - Adaptive mesh refinement loops for fusion
  - Shape optimization in accelerator design
  - Front tracking and AMR in supernovae simulations
  - Mesh quality improvement, front tracking and AMR in groundwater applications

SciDAC Scientific Discovery through Advanced Computing

38

Contact Information

ITAPS

- Web Site: www.itaps-scidac.org
- Email: itaps-mgmt@lists.llnl.gov
- Lori Diachin: diachin2@llnl.gov

We actively seek and welcome your input!

SciDAC Scientific Discovery through Advanced Computing

39

LLNL Disclaimer and Auspices

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

SciDAC Scientific Discovery through Advanced Computing

40