# Advanced SIMULATION Technologies for APPLICATION Scientists

Advanced simulation techniques lie at the heart of many of the nation's most pressing scientific challenges, including understanding our changing climate, designing safe and efficient energy sources, and managing the nation's nuclear stockpile. For example, in the area of designing next-generation nuclear reactors, many new or modified designs can be evaluated using computer simulations before the best designs are chosen for further study. These simulations are built using advanced mathematical models that describe the underlying physical phenomena, and sophisticated software tools that allow scientists to examine solutions for many different scenarios.

To build these simulations, research scientists first devise a mathematical model of the physical process they would like to study. This results in one or more equations that approximate physical processes, along with a description of what is occurring on the boundary (boundary conditions) and at the beginning of the simulation (initial conditions). In addition, scientists must develop a computer representation of the computational domain. The geometry of the domain can be as simple as a rectangular box or sphere or as complex as one can imagine when studying advanced scientific devices. In most cases, the mathematical equations describing the physical phenomenon cannot be solved analytically on the computational domain of interest. In such cases, the domain is decomposed into a collection of simpler geometries—a mesh—typically comprising triangles or quadrilaterals in two dimensions and tetrahedrons or hexahedrons in three dimensions. Once the mesh has been generated, the mathematical equations are approximated on that mesh resulting in a system of algebraic equations that is easier to solve on a computer than the original equation is. Once the solution of this system of equations is obtained, it is extensively analyzed and, when possible, validated against experiments to ensure the solution is correct. This process is repeated with adjustments made to the mathematical model, the computational domain, the mesh, or the numerical solution process until the scientific goal in question is achieved (sidebar "Representing Geometry and Generating a Mesh"). The Interoperable Technologies for Advanced Petascale Simulations (ITAPS) center is primarily interested in the steps of the solution process associated with modeling the computational geometry, generating the computational mesh, and refining or adjusting that mesh to ensure it optimally meets the needs of the simulation scientist.

The ITAPS center is primarily interested in the steps of the solution process associated with modeling the computational geometry, generating the computational mesh, and refining or adjusting that mesh to ensure it optimally meets the needs of the simulation scientist.

We note that the entirety of this process is significantly complicated by the computer architectures available to scientists today. These machines often have large numbers of processors, currently in the hundreds of thousands, and moving toward one million and beyond. Because the memory for these machines is distributed across the processors, the computa-

## Representing Geometry and Generating a Mesh

Representing realistic computational geometries for advanced numerical simulation can be a time-consuming process. One first starts with a high-level description of the computational domain. See figure1(a) for a typical example from accelerator modeling. This description is often generated from engineering specifications in a computer-aided design (CAD) package, and provides a functional representation of the boundaries of the domain using splines. Given the boundary representation of the computational domain provided by CAD packages, automatic mesh generation tools can divide both the surface and interior of the domain into a union of simpler shapes needed for numerical simulation. In this case, in figure 1(b) the mesh consists of triangular elements on the surface and tetrahedral elements in the interior of the geometry. For this geometry, note that some areas of high curvature have a finer mesh than other areas because we expect more interesting solution behavior in those areas. Once the mesh is generated it must be partitioned and distributed to the processors of a parallel computer. Figure1(c) shows the mesh divided into eight partitions of equal size. The final image, figure1(d), shows the numerical solution obtained on this mesh. The solution quality could potentially be improved by further refining and repartitioning the mesh or modifications to the original geometry could be considered to improve the performance of the accelerator part. In the latter case, the mesh would need to also be modified or regenerated before the numerical simulation proceeds. The ITAPS center is working to create tools that enable application scientists to easily experiment with different mesh types, refinement procedures, or modifications to the computational geometry.

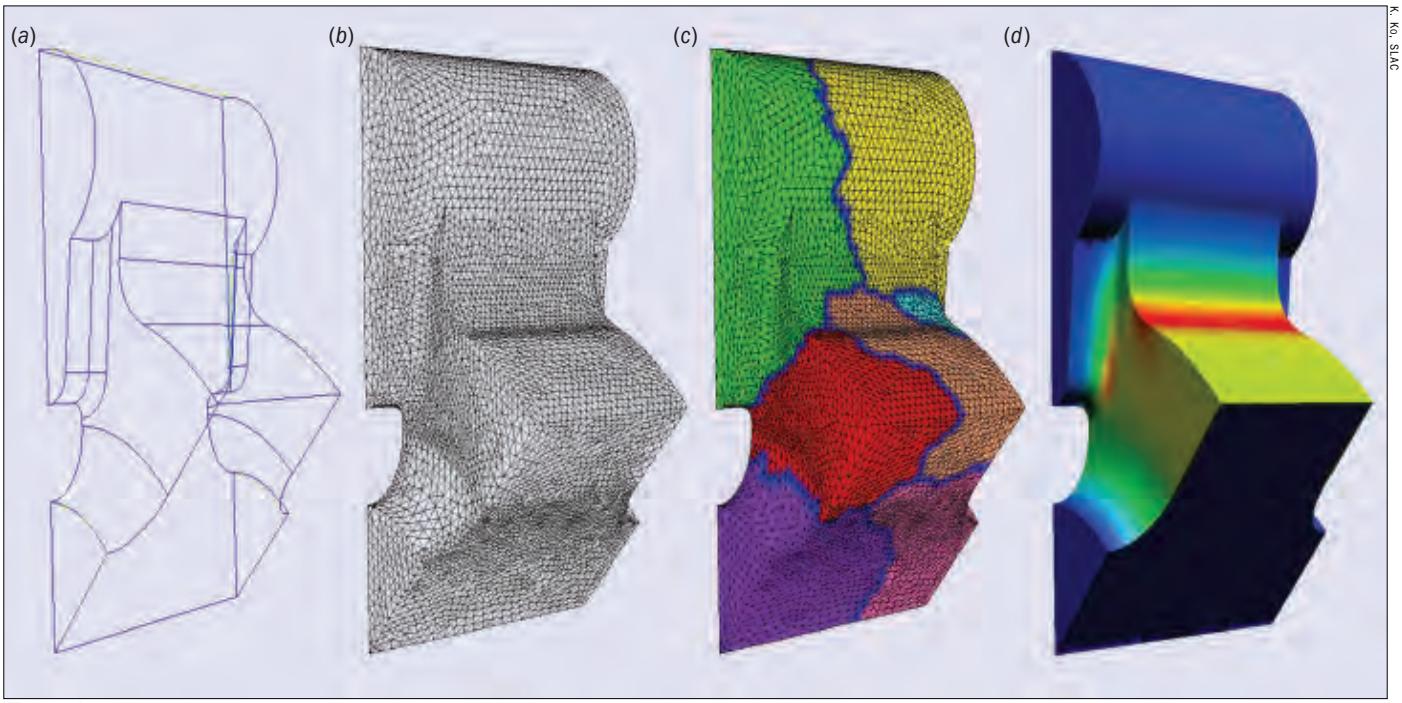


K. Ko, SLAC

**Figure 1.** *The basic steps of generating a mesh on a complex computational domain and using that mesh to solve complex physics problems in parallel.*

tional domain and mesh must also be distributed, and communication between processors must take place to ensure correct execution of the simulation. Some of the challenges we face when using this type of computing environment include ensuring that equal amounts of work are given to each processor to "balance the load" and that the overhead associated with communication costs are minimized. We must also ensure that the computations are coordinated so that the data structures and solution variables are consistent across processors as the simulation proceeds.

The ITAPS SciDAC project focuses on all of these issues, and this article highlights the work being done to represent the computational domain and to generate high-quality meshes for the simulation process on massively parallel computers.

### Computational Meshes

Optimally solving physics-based simulations involves achieving the highest possible accuracy in the most efficient way on today's and tomorrow's computer architectures. The key factors in achieving this goal that the ITAPS project focuses on are

## Meshes: Structured versus Unstructured

In figure 2 we show both a structured and unstructured mesh that has been refined to better represent the orographic (height) fields on the globe. In figure 2(a) we show in red the mountainous regions that we would like to model in more detail. Figure 2(b) shows a structured grid in which the points have been moved to focus more of them in the regions of interest. This approach allows the grid to be highly efficient both in terms of memory and computational cost, but the quality of some elements is affected. We note that other methods that use structured grids are also possible including the use of block-structured grids which use patches of potentially different grid spacing to represent features of interest in the computational domain. Such techniques preserve the quality of grid elements and the efficiency of the methods, but introduce the need for special consideration at the boundary between patches of different grid size. The rightmost image, figure 2(c), shows an unstructured grid for the same computational domain. Because the connectivity among elements can vary, it is relatively easy to generate meshes that conform to features of interest. However, these meshes have higher costs in terms of both storage and computation.

A. Khamesyeh, ORNL

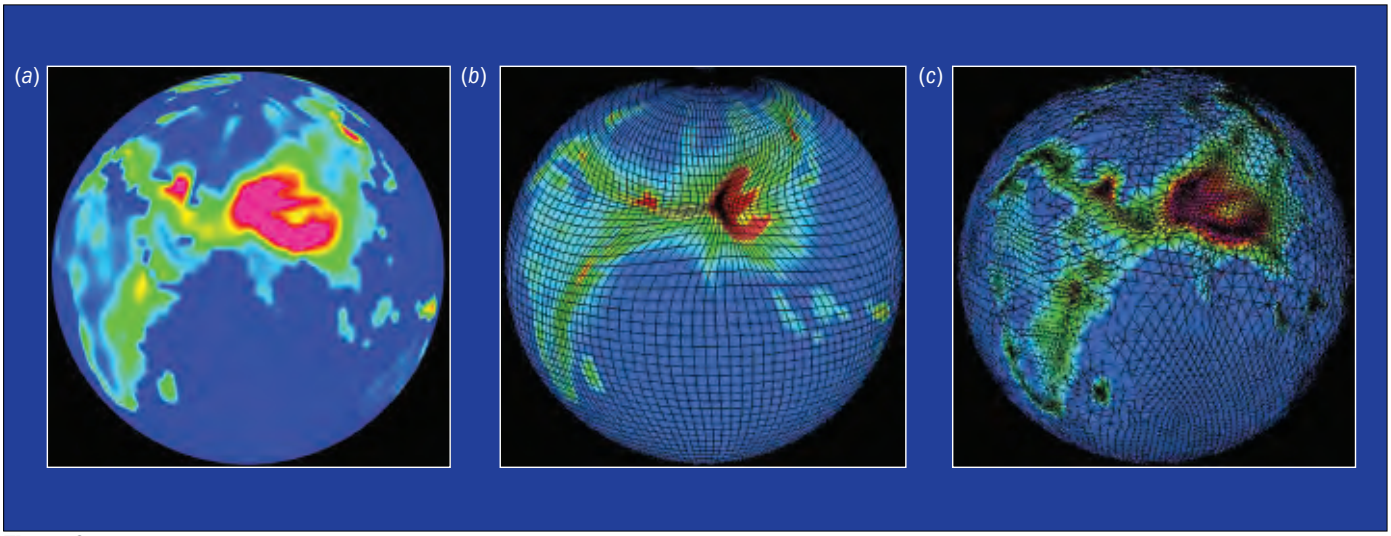


**Figure 2.** *The orographic (height) fields on a globe (a) and the structured (b) and unstructured (c) meshes, used to concentrate grid points in regions of interest.*

- Accurately representing the computational geometry; that is, maintaining fidelity to curved surfaces and complex, interconnected domains

- Generating and using a mesh that is ideally suited for the simulation; that is, it contains high-quality elements that are focused on regions of interest or special structures that follow moving interfaces or boundaries

- Balancing the load across the processors of a parallel computer

Accurately representing the geometry is important as subtleties in shape that look like small errors in representation of the boundary can introduce large errors in simulation results.

Accurately representing the geometry is important as subtleties in shape that look like small errors in representation of the boundary can introduce large errors in simulation results. To achieve the highest possible accuracy, the geometric model of computational domain should use shapes as close to reality as is possible. For example, rather than using a series of line segments to represent what in reality is a curved line, it would be more accurate to use a higher-order representation that is itself curved. Many tools exist that allow engineers and scientists to create complex geometrical models and use them in simulations. However, they are often ill-suited for use on massively parallel computers or not flexible enough to meet the advanced needs of some simulations.

Similarly, the subdivided representation of the computational domain, the mesh, can be configured in many different ways to increase solution efficiency and accuracy. Generally speaking, meshes used in numerical simulations can be either structured or unstructured (sidebar "Meshes: Structured versus Unstructured"). Structured meshes are laid out in a very regular pattern, often in a Cartesian coordinate system, and are very efficient in terms of their memory consumption and use of computational resources. However, it is generally more difficult to represent complex geometrical domains with

## Adaptive Mesh Refinement Reduces Costs by Orders of Magnitude

For problems on complex three-dimensional domains and in cases where there is substantial anisotropy in the physical solution, unstructured adaptive mesh techniques have been shown to offer distinct advantages in providing cost-effective, reliable numerical simulation. ITAPS team members have developed a number of serial and parallel adaptive simulation procedures that can be and have been used with SciDAC application codes. We collaborate closely with application scientists to define error indication procedures that drive the interaction between the application simulation procedures and ITAPS mesh adaptation service. These procedures are typically built on a set of standard error indication methodologies and determine the size and shape of elements needed to gain the desired level of accuracy.

The ITAPS mesh adaptation service is being used by the Stanford Linear Accelerator Center (SLAC) with their high-order, finite-element methods for calculating electromagnetic fields. Figure 3 shows an initial mesh of 1,595 elements and the final adapted mesh of 23,082,517 elements for one accelerator component that they model and solve for in parallel. To achieve the same level of accuracy using a uniform mesh would add at least two orders of magnitude more elements.

Additional reductions in the numbers of elements can be realized by using the ITAPS mesh adaptation service to adapt to an anisotropic mesh size field. That is, the mesh elements are stretched and oriented to align with solution features. Figure 4 (p26) shows an internal view of initial and adapted meshes for which both isotropic and anisotropic adapted meshes were constructed. In this example, the anisotropically adapted mesh has an order of magnitude fewer elements than the isotropically adapted one. Such capabilities are currently being explored with fusion application scientists modeling confined magnetic plasmas in tokamaks; such phenomena also display high degrees of anisotropy.

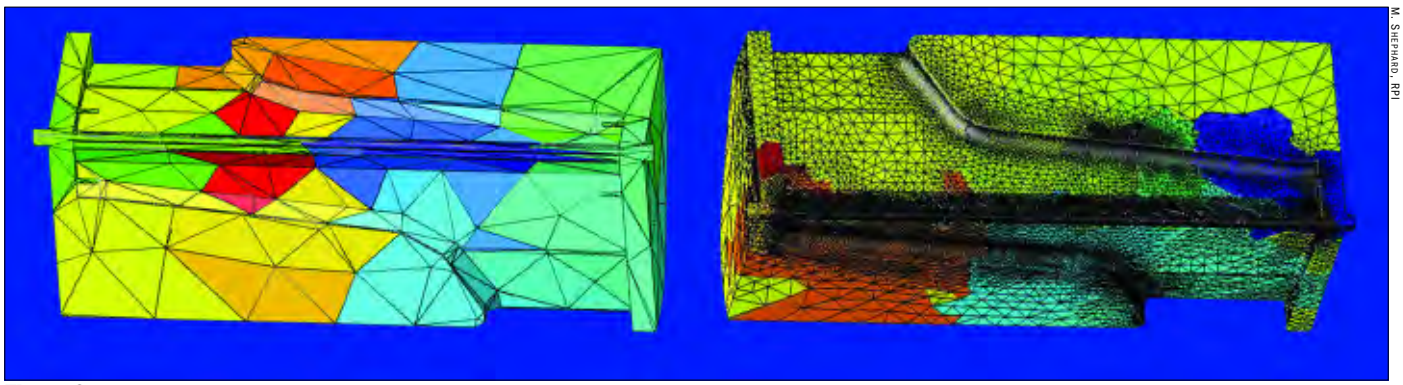

M. Shephard, RPI

**Figure 3.** *Initial and adapted mesh of an accelerator component. Colors indicate the mesh partitions in this example.*

curved surfaces using these mesh types. Various methods are being explored to overcome this limitation, including the use of overlapping grids and embedded boundary methods. While these techniques have had success in many application areas, special techniques must be used on or near the boundaries of the computational domain, and there is still much ongoing research in this area. In contrast, unstructured grids are well-suited for accurately representing the boundaries of complex geometries. Depending on the geometry and mesh type desired, they can be more time consuming and labor intensive to generate. In addition, they require more computer memory than structured grids as the connectivity between grid points must be explicitly stored, and it can also be more computationally expensive to obtain the numerical solution using these grids. However, they are a popular choice for many different engineering and scientific applications due to their ability to accurately and flexibly represent complex computational domains.

In both the structured and unstructured grid cases, one can significantly improve the accuracy and computational cost of many calculations using a variety of techniques. One of the primary mechanisms for accomplishing this is to change the resolution of the mesh as the computation proceeds to better capture the physics of interest. That is, if you use more grid points in areas where the solution is changing rapidly or other features of interest are occurring, and take away grid points in areas where the solution is relatively uninteresting, you can save a considerable amount of effort compared with using the same density of grid points everywhere. This technique is called adaptive mesh refinement (sidebar "Adaptive Mesh Refinement Reduces Costs by Orders of Magnitude") and is a powerful tool available to application scientists. Other variants of this idea can be used to increase accuracy or improve computational cost. For example, rather than increasing the number of elements in an area of interest, one

Many tools exist that allow engineers and scientists to create complex geometrical models and use them in simulations. However, they are often ill-suited for use on massively parallel computers.

M. Shephard, RPI

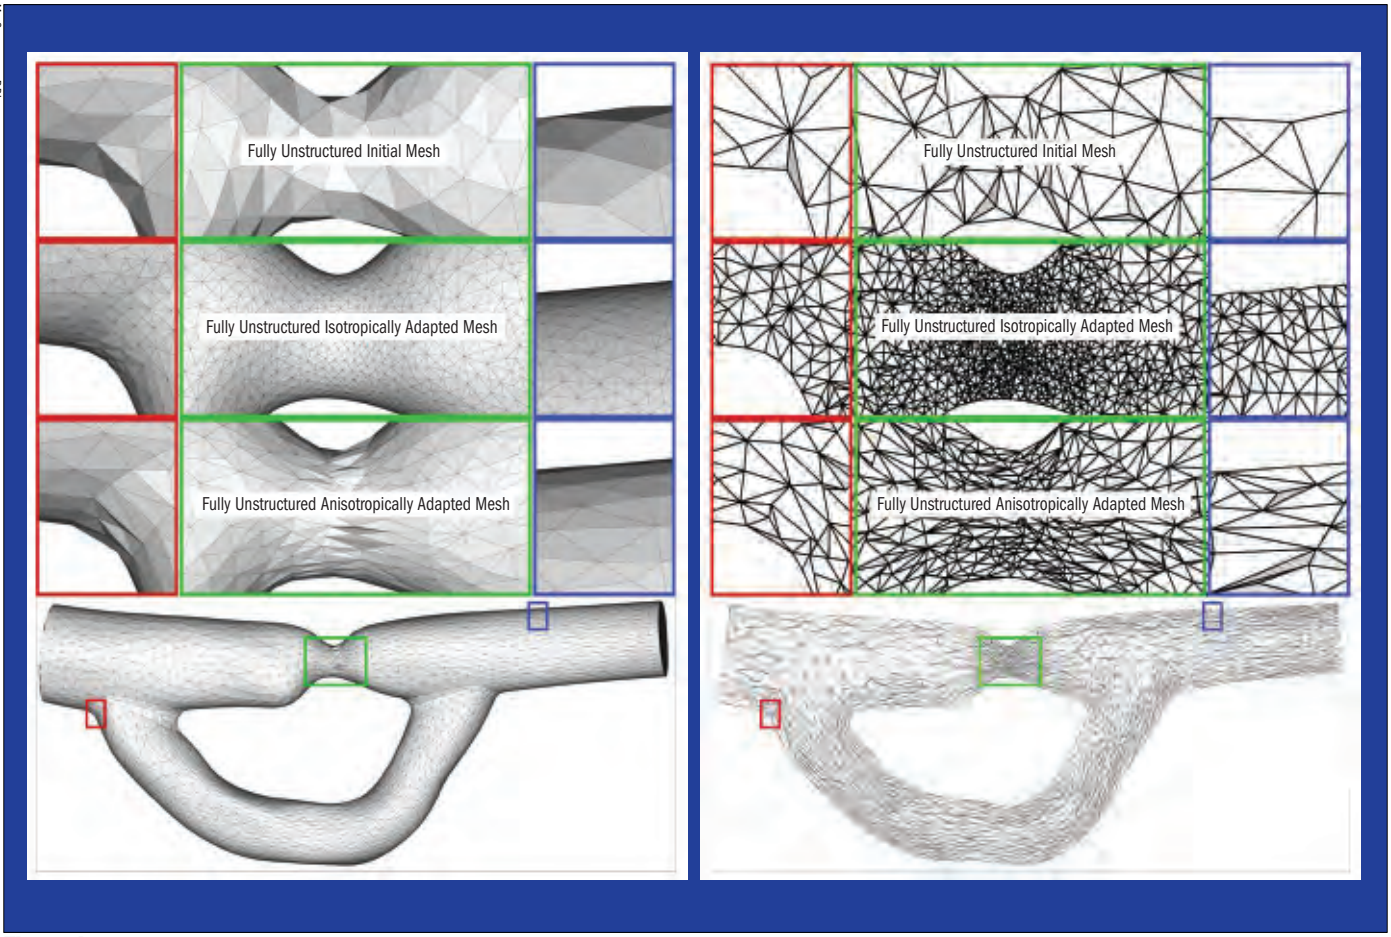


**Figure 4.** *Meshes of an arterial by-pass show that one can reduce the number of mesh elements needed by using anisotropic mesh refinement.*

could instead use more accurate numerical approximations in a local area, or cluster the grid points without adding any to minimize the solution error in the simulation. Another technique is to explicitly follow moving fronts or boundaries between two different materials in a simulation using a special mesh that tracks the interface. These techniques have proven beneficial in many application settings and have saved orders of magnitude in the cost of computation for simulations in fluid dynamics, materials modeling, astrophysics, and many others.

Unfortunately, these advanced techniques can be difficult to implement and use because they significantly complicate the underlying data structures and algorithms. Furthermore, the dynamic nature of many of these methods implies that the amount of work on each processor of a parallel computer changes over the course of the simulation. This in turn requires that the mesh be re-distributed periodically to preserve a consistent workload across processors. These issues present a considerable distraction for domain scientists and can take them away from their primary scientific enterprise. Thus, advanced mesh and geometry techniques, which have proven extremely beneficial in a wide array of applications, are not being leveraged as much as they could be in SciDAC applications. Correcting this situation will provide a significant advance in the amount of science that can be done in the SciDAC program.

Better leveraging of advanced mesh and geometry techniques will provide a significant advance in the amount of science that can be done in the SciDAC program.

## The ITAPS Project

The ITAPS project is addressing this problem by developing an interoperable infrastructure that allows scientists to more easily use sophisticated mesh, geometry, and field manipulation tools developed by the computational mathematics and computer science communities. The key new innovation we are developing as a center is the interoperability of existing tools. By creating an environment that allows scientists to easily try different tools that perform similar tasks we encourage experimentation and the use of the technology scientifically best-suited for an application, not just the technology that is convenient or easy to use. Moreover, by bringing together experts from across the DOE community in mesh and geometry technologies, we can more easily build advanced, higher-level tools that combine one or more capabilities together.

## Partitioning is Critical to Efficient Computations

One of the most critical steps in obtaining high performance on massively parallel computers is partitioning the problem across tens or hundreds of thousands of processors. This task requires careful attention to ensure that an equal amount of work is distributed to each processor, that communication overhead costs are minimized, and, if occurring repeatedly throughout the simulation, that the cost of partitioning itself does not dominate the costs in the computation. The Zoltan partitioning toolkit provides access to many different types of partitioners, including those that divide the mesh based on geometric information and those based on mesh entity connectivity information (figure 5). Geometric techniques tend to be easy to use because they require only vertex coordinate information, and they are inexpensive to execute, so they are suitable for adaptive mesh refinement. However, the partition quality can be mediocre and there is no explicit control of communication costs. Moreover, these types of partitioners can generate disconnected subdomains on complex geometries. In contrast, graph-based and other connectivity-based partitioners divide the mesh with respect to the entities' data dependencies. That is, they assign entities that depend on each other (such as two vertices that are connected by a mesh edge) to the same processor. These partitioners have proven highly successful for mesh-based PDE problems because they allow for explicit control of the communication volume resulting in higher partition quality than the geometric-based techniques (figure 6, p28). However, they are more expensive and difficult to use than their geometric counterparts. Zoltan allows applications to experiment with all of these different partitioning types through the ITAPS interfaces, giving application scientists an opportunity to determine which methods are best suited for their simulation needs.

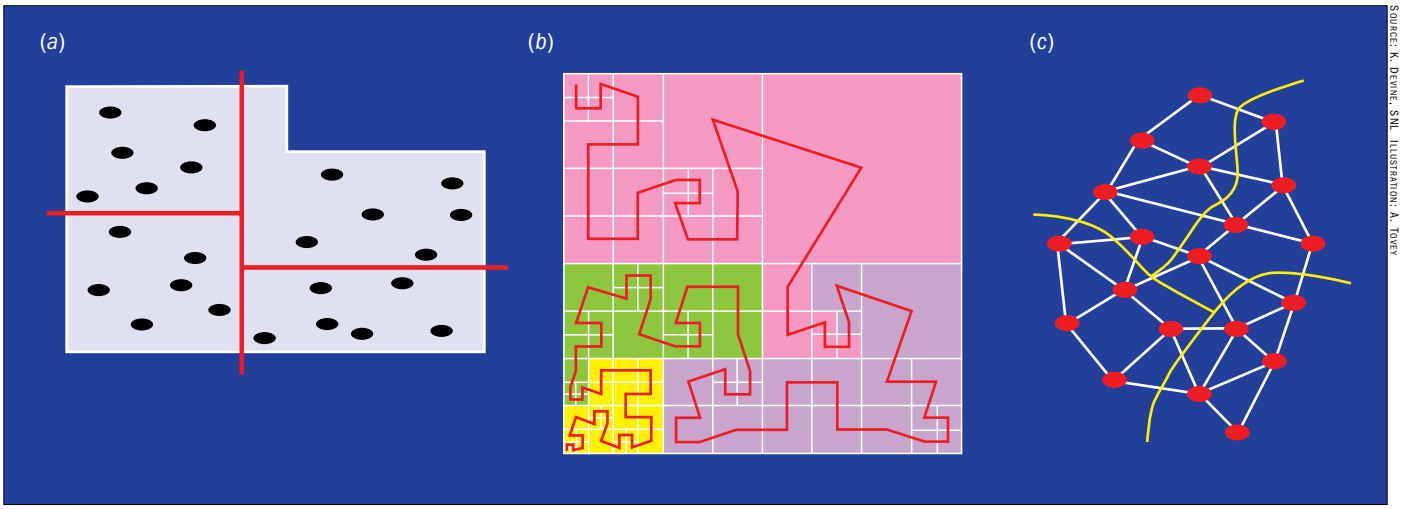


SOURCE: K. DEVINE, SNL ILLUSTRATION: A. TOVEY

**Figure 5.** *Geometric (coordinate-based) partitioning methods, (a) and (b), use geometric information to divide the mesh, whereas graph-based methods, (c), use dependencies in the mesh to improve partition quality.*

Such technologies address the more sophisticated needs of both current and next-generation application codes.

The challenges associated with accomplishing this goal are both technical and sociological. Technically speaking, the infrastructure we provide must balance the often-conflicting goals of flexibility to support many different tools, data structures, and usage scenarios, with ease of use and efficiency. If the infrastructure we provide cannot use existing application code data structures or be easily experimented with, no one will use it. On the sociological side, the ITAPS task is complicated by the fact that the mesh and geometry data structures serve as the foundation upon which the rest of the simulation is typically built. Application scientists are understandably reluctant to make changes to these data structures to adopt new tools without knowing in advance that there will be a significant payoff. Moreover, they must trust that any new tool adopted will exist and be supported for as long as it is used and needed.

The ITAPS approach to solving these challenging technical and sociological problems is to define common interfaces, along with general data model abstractions, for the types of data most often used in simulation applications. ITAPS has defined interfaces for geometry, mesh, fields and the relationship among these data types. These interfaces must be suitable for a wide variety of underlying tools and use case scenarios. We also provide access to advanced mesh and geometry services such as independent software libraries that can be called directly from an application simulation code. To date, we have

The ITAPS project is developing an interoperable infrastructure that allows scientists to more easily use sophisticated mesh, geometry, and field manipulation tools developed by the computational mathematics and computer science communities.

K. Devine, SNL

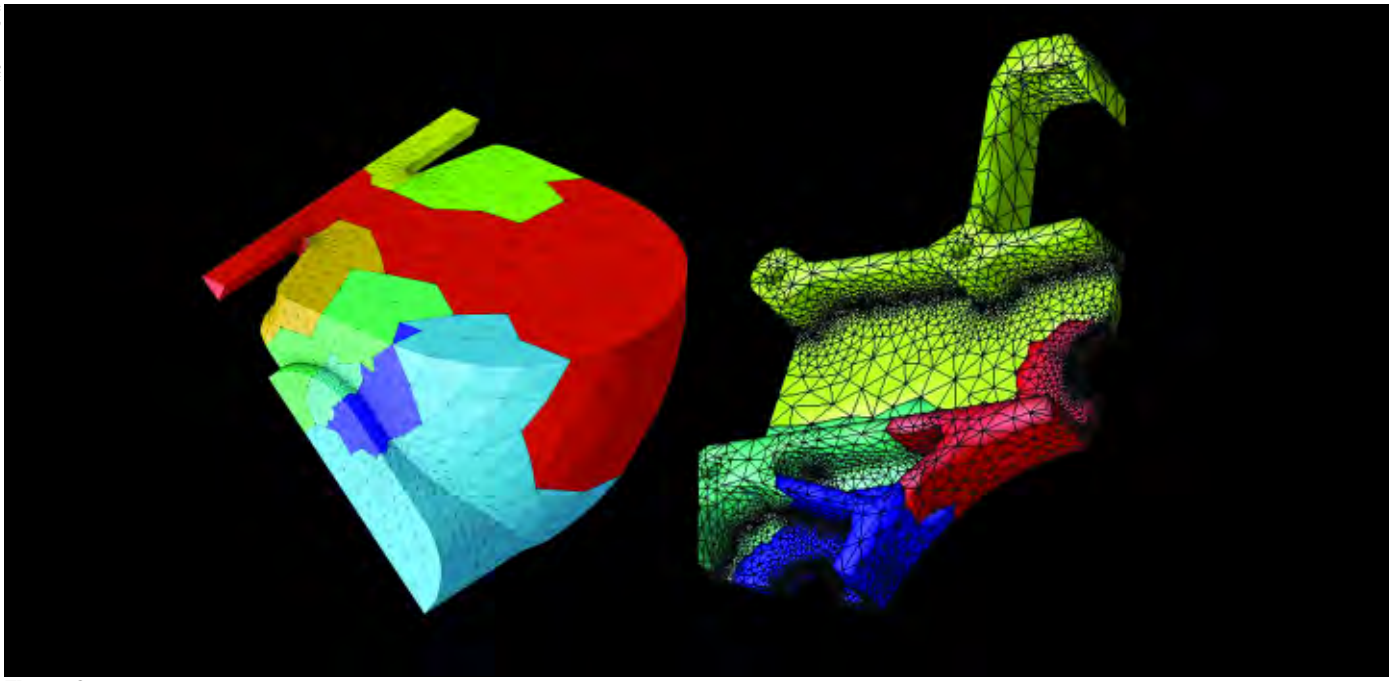

**Figure 6.** *Explicit control of communication costs enabled by graph-based partitioners enables efficient computation by reducing the number and size of messages that are communicated between processors.*

Source: X. Li, SUNY SB Illustration: A. Tovey

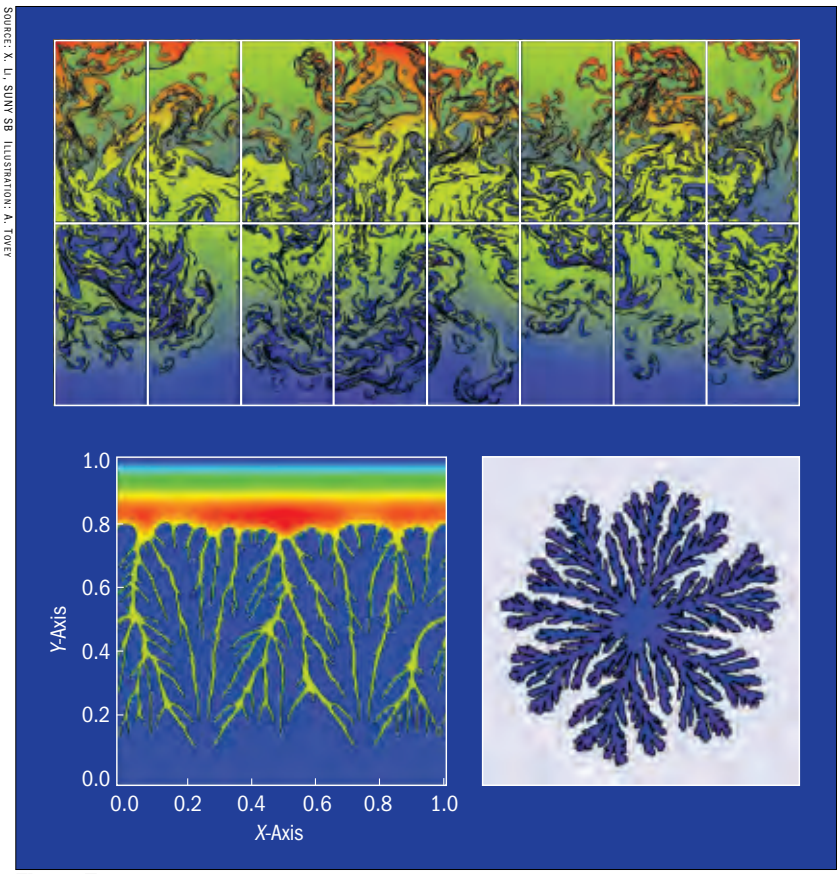


**Figure 7.** *The FronTier library has been used for a number of applications of interest to DOE including Rayleigh–Taylor instabilities which are critical in understanding fluid mixing (top) and crystal precipitation and growth in subsurface flows (bottom). The latter simulation was performed as part of a new collaboration between ITAPS and groundwater scientists at PNNL who are studying reactive transport and mineral precipitation in fractured and porous media.*

focused on interfaces for core data types, such as geometrical models and unstructured meshes. Our primary goal has been to define key functionalities that provide data access, associate user data to various entities, promote relationships among entities, and to modify the mesh and geometry. While it is important to keep the abstractions for the mesh and geometry separate so that application scientists can use only those pieces that are necessary for their application, it is also critical to provide infrastructure for tracking the relationships between these core data types. Most recently, we have also focused considerable attention on developing the interfaces needed for parallel computation.

The ITAPS team now has a considerable collection of tools that use these interfaces that are useful in application development. In particular, there are several mesh and geometry databases, mesh quality improvement tools, partitioning tools (sidebar "Partitioning is Critical to Efficient Computations" p27), and mesh adaptation tools, including mesh refinement and front tracking (sidebar "Front Tracking Techniques Lead to Unexpected Results"). These tools all work in parallel and new work has enabled extremely large simulations to be run on tens of thousands of processors (sidebar "Scaling to 100,000 Processors and Beyond" p31). In addition, the common interface allows the ITAPS team to develop "high-level services" that were not previously available. For example, adaptive mesh refinement can be combined with techniques to

# Front Tracking Techniques Lead to Unexpected Results

Front tracking techniques have been developed to explicitly track the motion of a moving interface in a simulation, for example, the interface between two different materials. The front tracking capabilities in ITAPS are provided as a stand-alone library called FronTier-Lite that has been used with a wide variety of applications of interest to DOE. These applications range from the study of fluid mixing instabilities, astrophysical supernova simulations, crystal precipitation and growth in subsurface flows, the simulation of bio-fuel jets for efficient and renewable energy, and the simulation of pellet ablation fueling techniques for ITER (figure 7), an international research collaboration that aims to demonstrate the scientific and technical feasibility of fusion power

The techniques developed by the ITAPS team involve the use of marker particles which represent the interface moving through a computational mesh. The particles are connected to each other to form a triangulated mesh or piecewise linear segments representing the interface (figure 8). The FronTier-Lite library is fully parallelized and interoperable with the ITAPS interfaces. The library hides the most complicated operations of mesh redistribution and topological bifurcation from the user and provides a smooth and accurate computation of the moving front.

In a recent collaboration with fusion scientists at General Atomics, we used the FronTier-Lite front tracking library to study the physics of pellet ablation for tokamak fueling. Obtaining a deep understanding of this phenomenon is necessary for the successful operation of ITER (figure 9, p30). Plans call for fueling ITER by injection of frozen deuterium pellets which ablate when interacting with the hot fusion plasmas; the lower the quasi-steady state ablation rate, the higher the fueling efficiency. However the physical processes associated with this ablation are not clearly understood, and advanced modeling and simulation tools are being used to shed light on these processes (figure 9, bottom).

Using the front tracking technology developed by the ITAPS Center, we have developed novel mathematical models and computational software for the numerical simulation of the pellet ablation for tokamak fueling. Using this software, the pellet ablation rate and lifetime in magnetic fields were systematically studied for the first time and compared with theory and experimental databases. Simulations revealed several new features of the pellet ablation, such as (contrary to expectations) the ablation rate depending strongly on the plasma pedestal width and the magnetic field strength (figure 10, p30). This new feature implies that pellets traversing strong plasma gradients, as in the edge pedestal region of the ITER plasma, could have significantly lower ablation rates (higher fueling efficiency) if injected at higher velocity. Simulations also demonstrated that the ablation cloud rotates with supersonic velocity about its main axis, a phenomenon that significantly influences ablation rate.

explicitly track interfaces or with high-accuracy representations of geometry to increase both the efficiency and the accuracy of numerical simulations. Previously, such tools were difficult to combine; the new ITAPS infrastructure makes this significantly easier. Descriptions of the ITAPS interfaces, along with the tools that use them, are available for download from the ITAPS website (see Further Reading, p35).

For an application scientist to experiment with these tools, they need only write a small amount of wrapper code around their internal data structures to become compliant with the ITAPS interfaces. He or she can then easily experiment with the broad array of tools that we provide through the interface. We also provide reference implementations of the ITAPS interfaces on which new applications can be constructed; these implementations can also be used alongside an application's data in cases where the cost of a data copy is acceptable. If they determine the tool is meeting their need they can incrementally implement the interfaces as needed to improve performance.

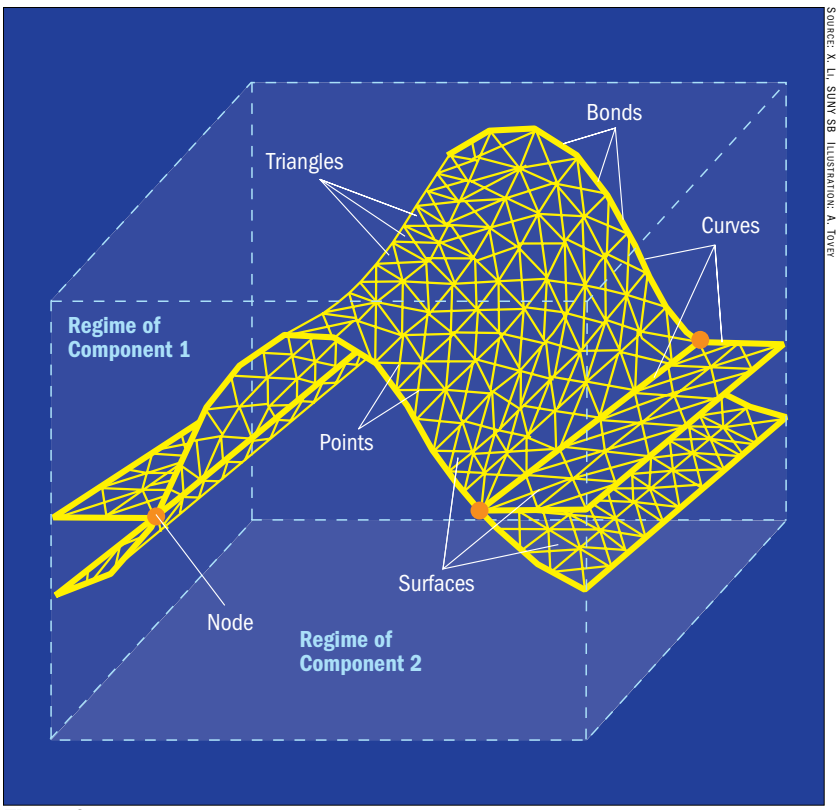


SOURCE: X. LI, SUNY SB ILLUSTRATION: A. TOVEY

**Figure 8.** *A triangular mesh representing the interface between two different materials. This mesh is moved through the computational domain to follow the motion of the interface.*

## Use in Applications

The ITAPS team is working extensively with scientists from many DOE mission-relevant application

The ITAPS team is working extensively with scientists from many DOE mission-relevant application areas, analyzing their needs for advanced geometry and mesh technologies, and collaborating with them to demonstrate the promise of such techniques in their scientific domains.

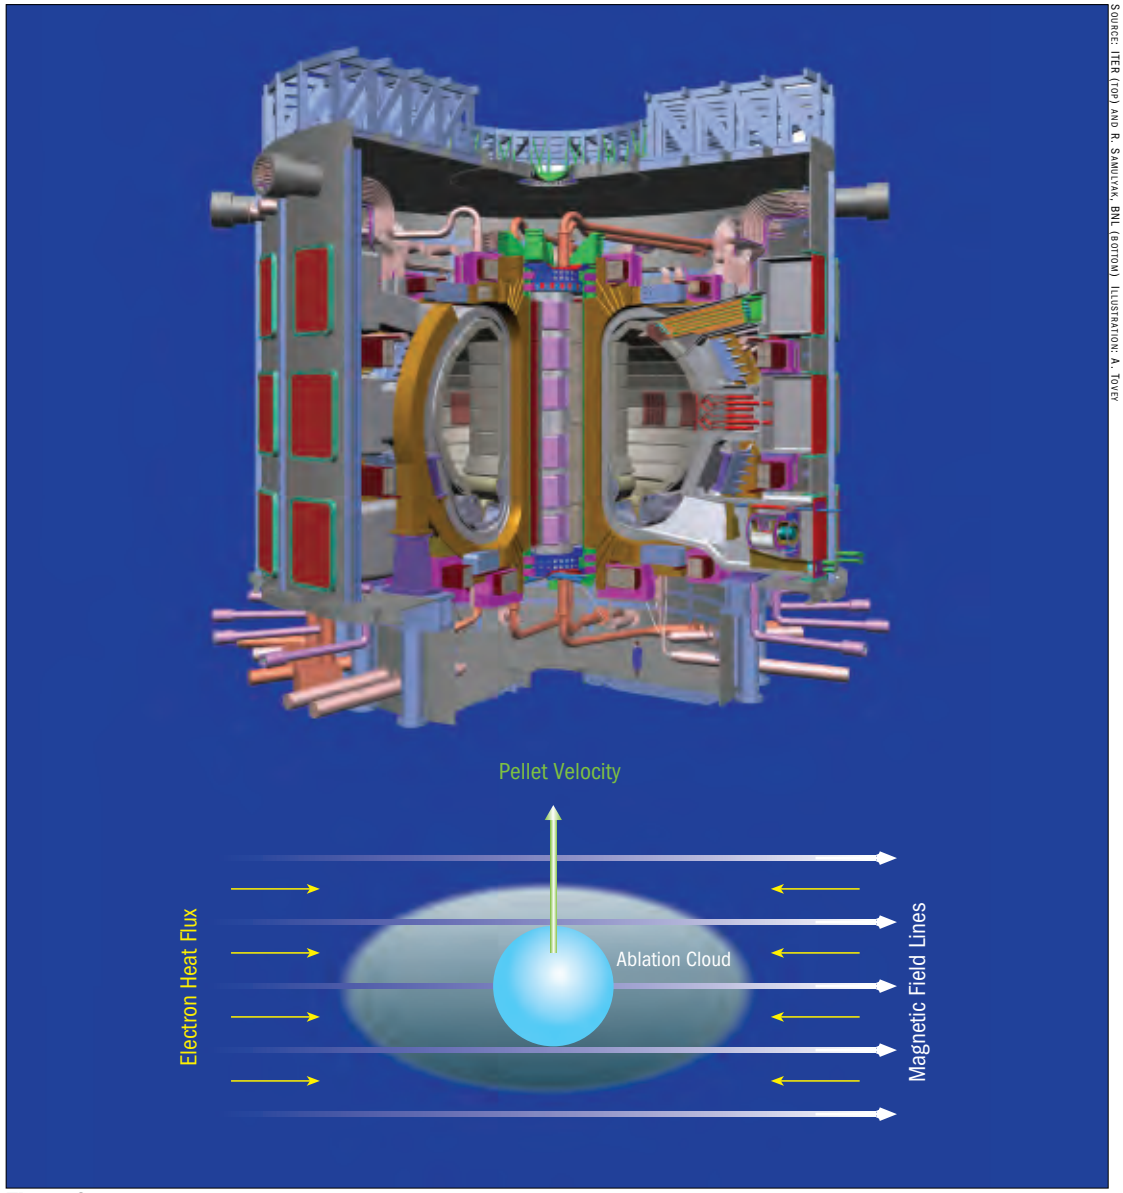


SOURCE: ITER (TOP) AND R. SAMULYAK, BNL (BOTTOM). ILLUSTRATION: A. TOVEY

**Figure 9.** *A schematic of ITER (top) and of the physical processes associated with cryogenic deuterium pellet ablation in a tokamak magnetic field (bottom). Hot electrons traveling along the magnetic field lines hit the pellet surface causing a rapid ablation. A cold, dense cloud forms around the pellet and shields it from incoming hot electrons. The most important processes determining pellet ablation from that point forward occur in the cloud.*

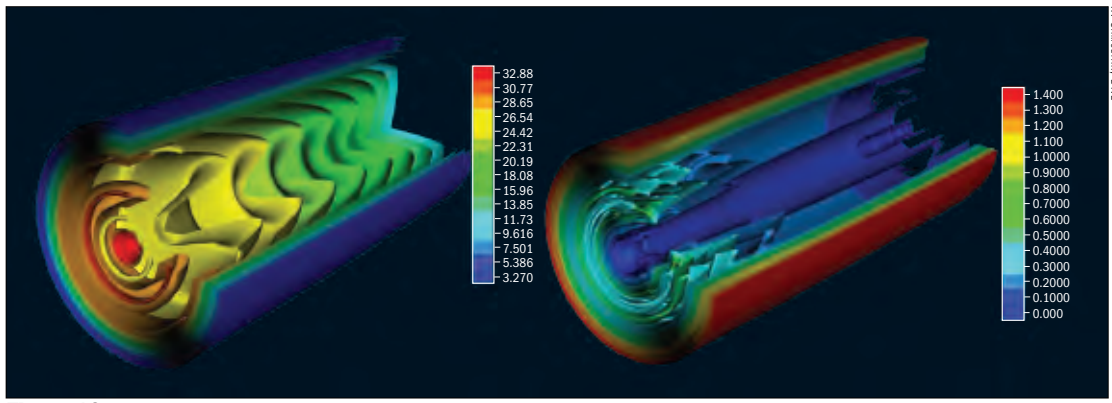


R. SAMULYAK, BNL

**Figure 10.** *The isosurfaces of pressure (left) and rotational Mach number (right) for the steady state ablation cloud. Numerical simulations of these quantities revealed new properties for pellet ablation and led to key insights for ITER fueling techniques.*

## Scaling to 100,000 Processors and Beyond

Efficiently operating on today's leadership-class computing facilities requires being able to operate on tens to hundreds of thousands of compute cores. This is a significant challenge for simulation codes, and scalability for the entire process will only be achieved when the individual steps are also scalable. The ITAPS team has successfully demonstrated strong scaling for several key solution steps to tens of thousands of processors. In particular, we have worked with high-order methods on general unstructured grids using the PHASTA simulation code. Such problems are often most effectively solved using implicit methods where large systems of simultaneous equations must be solved. The computational work is carried out in two key stages. The first stage requires formulating the equations, and each part of the mesh must communicate with processor neighbors that share mesh faces, edges, or vertices. The second stage uses an iterative equation solution method where there is neighbor-to-neighbor communication along with global communication across all processors. Because dynamic load-balancing and partitioning procedures such as those found in Zoltan can provide well-balanced mesh partitions, the scaling is typically not affected by geometric domain complexity or mesh anisotropy. PHASTA has shown strong scaling (meaning that the problem size stays the same as the number of compute nodes increases) to 32,000 processing cores of the IBM Blue Gene/L computer on adaptively-defined meshes over general geometries. For example, figure 11 shows the strong scaling results up to 128,000 processors for an adaptively-defined mesh containing one billion elements that modeled the flow in an abnormal aortic aneurysm.

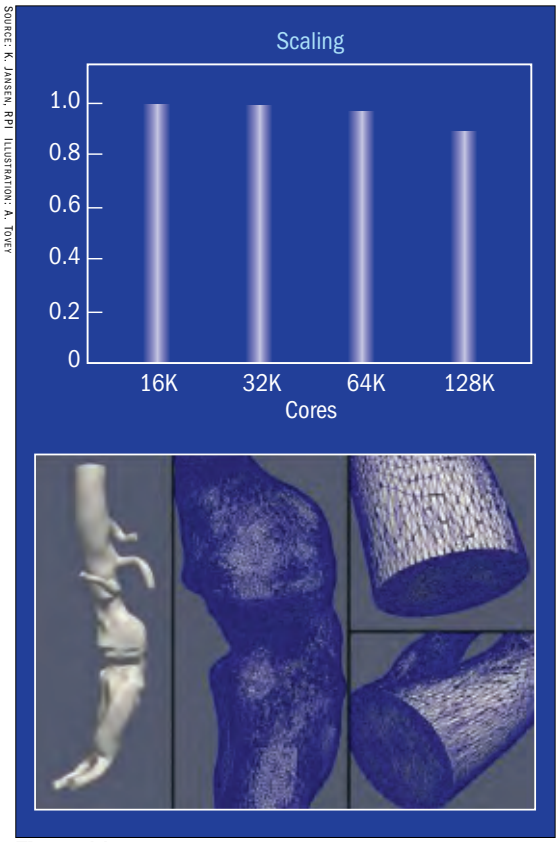


SOURCE: K. JANSEN, RPI ILLUSTRATION: A. TOVEY

**Figure 11.** *PHASTA strong scaling results on up to 128K processing cores of Blue Gene/L for the simulation of an aoritic aneurysm.*

areas, analyzing their needs for advanced geometry and mesh technologies, and collaborating with them to demonstrate the promise of such techniques in their scientific domains. As a result we have used the ITAPS tools in a large number of applications that span the DOE mission space, and highlight here a few key results in accelerator modeling, fusion simulation, biological modeling, and the development of next-generation nuclear reactor modeling codes.

*Accelerator Modeling*

Advanced accelerators have hundreds of thousands of components that need to be designed and engineered to obtain the highest-energy acceleration and beam properties. Numerical simulation is a key component in the design of the next generation of accelerator devices, and their problems are characterized by extremely complex geometries and the need for very high levels of accuracy (sidebar "Simulation Tools for Modeling Next-Generation Accelerators" p32). The ITAPS team has worked closely with scientists from the accelerator modeling community on many different aspects of their problem.

By providing high-quality meshes for complex geometries associated with the PEP-II device, we enabled the first-ever transit beam simulation using the Tau3P software, which supported a 15% increase in beam current in the upgraded device. Similar mesh generation efforts for the advanced Damped Detuned Structure (DDS) resulted in the first wakefield analysis of an actual DDS prototype and the direct verification of DDS wakefield suppression by end-to-end simulation

ITAPS researchers provided an adaptive mesh refinement capability to Stanford Linear Accelerator Center (SLAC) scientists to improve the accuracy of predicted field quantities that influence wall losses in the Rare Isotope Accelerator device by an order of magnitude

We have used the ITAPS tools in a large number of applications that span the DOE mission space.

# Simulation Tools for Modeling Next-Generation Accelerators

Particle accelerators are a significant part of the DOE science goals, accounting for seven of the top twenty facilities priorities over the next two decades. Simulation plays an increasingly important role in the design of these accelerators because of its impact on the performance improvements and operating cost reductions required to make new facilities successful. Accelerator performance is sensitive to geometric shape because of the high-frequency operating regime of beam cavities in large-scale accelerators and placement of waveguide elements in reduced-scale accelerators. This sensitivity drives the need for sophisticated geometric modeling and body-fitted mesh generation services from ITAPS. We are working extensively with the Stanford Linear Accelerator Center (SLAC) to improve the processes for the design and optimization of accelerator cavities used in the International Linear Collider (ILC), the Continuous Electron Beam Accelerator Facility (CEBAF) upgrade, Spallation Neutron Source (SNS), and other near- and mid-term priority accelerator facilities.

The electromagnetic simulations performed by SLAC are characterized by extremely complex geometries and the use of higher-order methods. Because the domains are curved and high-order methods are being used, the meshes must also be curved to provide a sufficiently higher-order geometric approximation to effectively achieve the desired level of accuracy. Standard mesh generation tools have difficulty creating such meshes and often contain inverted elements which pose a significant challenge for the simulation software. The ITAPS team has developed a mesh curve correction tool for this problem that automatically identifies the inverted elements and performs a series of mesh quality improvement operations to correct the inversion (figure 12). The mesh curving procedures allowed SLAC to perform more accurate simulations that were also more computationally efficient, resulting in up to a 30% saving in CPU time due to a better conditioned system.

The short-range wakefield simulations also require adaptive mesh refinement around the beam to resolve high frequency, while the rest of the domain can have a large mesh size. As an example, an ILC coupler has a beam pipe radius of 39 mm, but the beam region in a short-range wakefield simulation is only 300 microns. If the beam size is used to generate a uniform mesh, it will contain over 100 million tetrahedral elements, which is computationally infeasible. The ITAPS mesh adaptation service is being used to provide SLAC with a "moving mesh" refinement procedure that allows the computational scientists to focus elements in the areas that the beam is moving through and keep a coarse mesh everywhere else in the domain. Figure 13 shows this refinement region moving with the particles through the curved domains to achieve high accuracy at an acceptable level of computational efficiency. Considering that the domains are curved and high-order finite elements are used, the refined meshes must also be curved to provide a sufficiently high-order geometric approximation to ensure convergence of the solution. Using such techniques has resulted in a tenfold reduction in the computational cost of these simulations.

Future accelerator designs are currently done using manual design evaluations in which a simulation is conducted to analyze the performance of a given design. Given the results of simulation, the geometry of the accelerator cavity is modified slightly, a new mesh is generated, and the simulation is rerun to evaluate the new shape. This process currently requires significant human interaction and is impractical or impossible to perform efficiently. Advanced optimization techniques will allow accelerator scientists to explore the design space in a more automated way. The ITAPS center is working to provide tools that automatically adjust the geometry and mesh based on the optimization procedure. Such a technique requires significant interactions among many different ITAPS technologies including geometry modules, mesh databases, and mesh quality improvement tools. Figure 14 shows the steps of this procedure, as first the geometry and then the mesh are modified at each design step.

ITAPS researchers provide services for varying design geometry, quickly generating high-quality meshes for each new geometry, and automatically computing sensitivities of the mesh with respect to design parameters.

We are now working with SLAC researchers to provide tools for automatic tuning of accelerator geometries that will significantly increase the speed and decrease the cost at which new accelerators can be designed. ITAPS researchers provide services for varying design geometry, quickly generating high-quality meshes for each new geometry, and automatically computing sensitivities of the mesh with respect to design parameters.

*Fusion Modeling*

Fusion energy has the promise of providing a clean source of electrical power in the future. The fusion energy community is focused on designing and understanding large-scale experimental facilities, such as ITER, and makes extensive use of simulation capabilities. Some of the key characteristics of the problems that ITAPS is focusing on are the fact that the physical processes are highly nonlinear and anisotropic in nature, requiring adaptive mesh refinement and high-quality meshes. Toward this end, ITAPS researchers are contributing to a new effort at Princeton Plasma Physics Lab (PPPL) to develop an adaptive, high-order accurate method for studying the behavior of magnetically confined plasmas. We have found that high-order methodologies can significantly decrease the solution time needed to obtain a given level of accuracy in highly-anisotropic cases compared to lower-order accurate methods. We have also collaborated with scientists at General Atomics to apply adaptivity and front

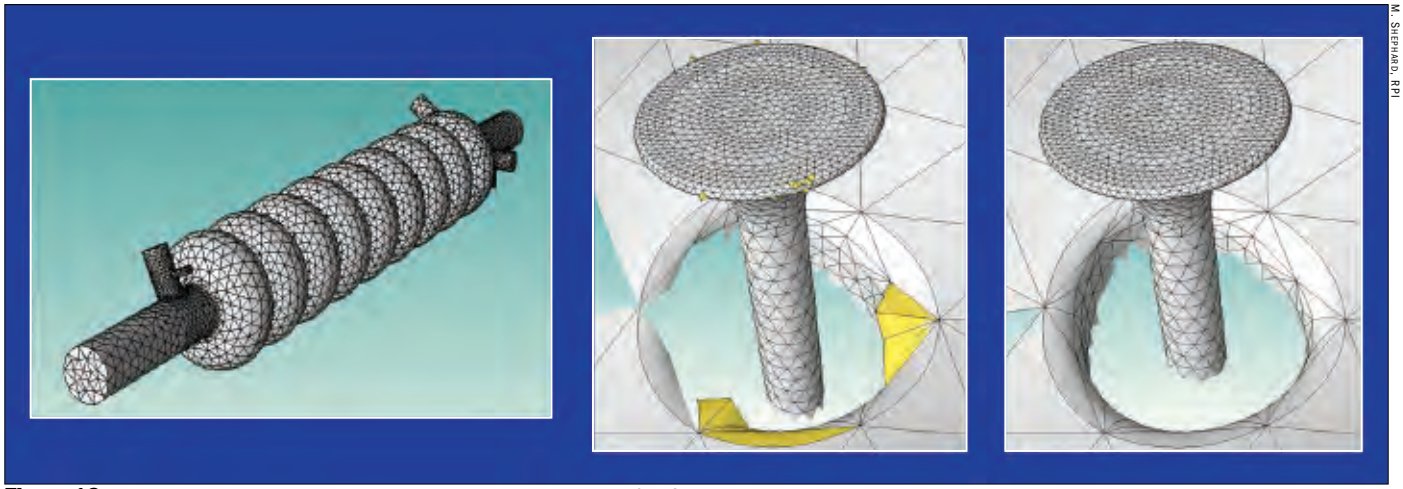

M. SHEPHARD, RPI

**Figure 12.** *The curved geometries and high-order methods used by SLAC to model next-generation accelerators require that the meshes also be curved to achieve high accuracy. We have developed automatic procedures to locate and correct inverted elements which significantly improves simulation accuracy and efficiency.*

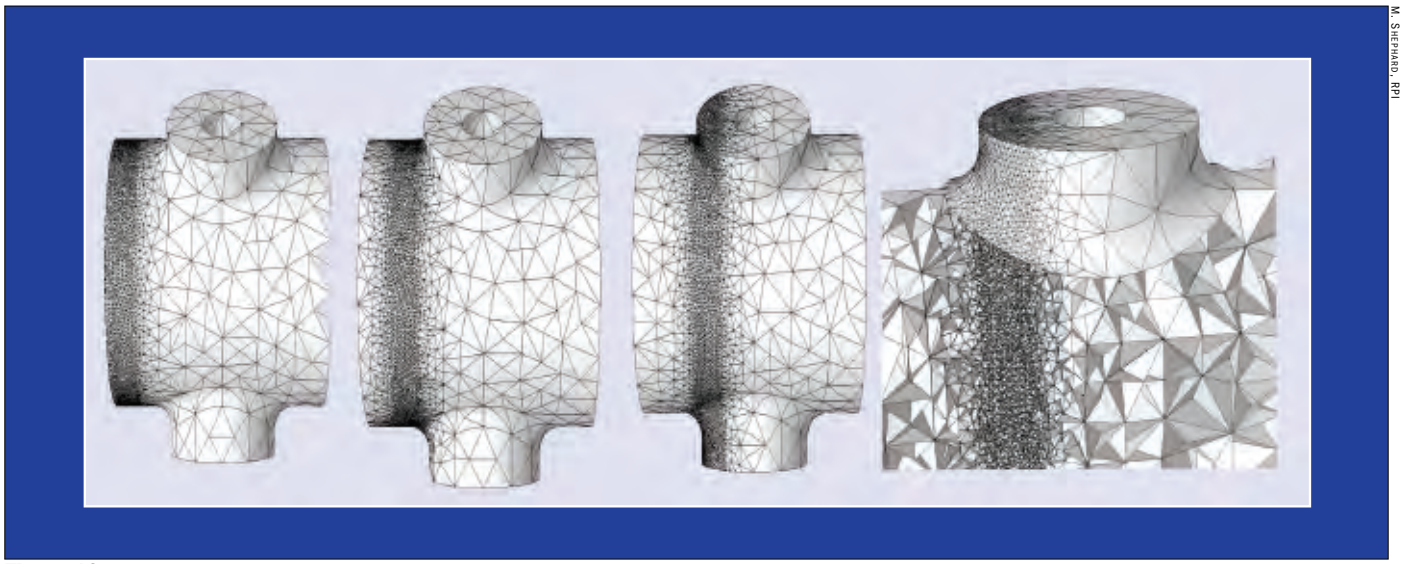

M. SHEPHARD, RPI

**Figure 13.** *The left three images show the adapted mesh for following particles in an accelerator cavity at three points in time. The right-hand image shows a cut on the interior of the mesh.*

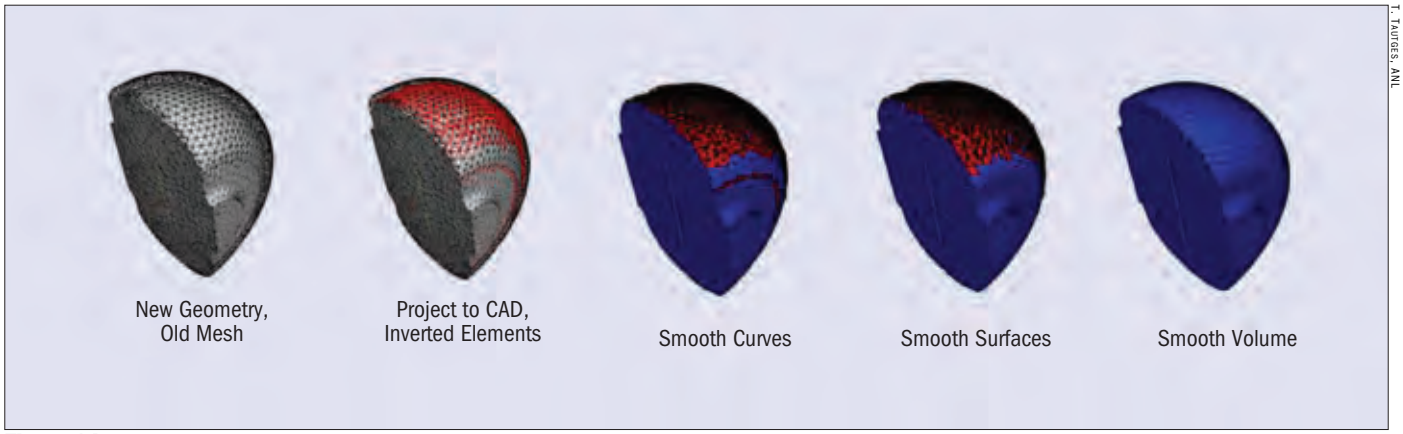


T. TAUTGES, ANL

**Figure 14.** *The steps needed to modify the geometry and mesh as part of the design optimization procedure. First, the new geometry is determined and the boundary nodes of the mesh are projected to the geometrical surface. This can result in a poor quality mesh, so the curves, surface, and volume of the mesh are then improved using smoothing techniques. The derivative information associated with the mesh motion is then computed for use in the optimization technique.*

# Next-Generation Nuclear Reactor Modeling

Nuclear energy can play a key role in securing the energy independence of the United States by offering a clean, safe alternative to traditional fossil fuel energy production. Simulation and modeling will play a key role in the development of next-generation nuclear reactors to improve their design, implementation, and operation. The physics associated with these systems operates in very complex geometrical domains and will benefit from the tools and technologies developed by ITAPS. We are involved in several collaborations whose goals are to develop and deploy high-performance computing tools for coupled multi-scale simulations of the sodium fast reactor.

The geometries of reactor cores can be quite complex and require scalable geometry and mesh generation tools. For example, the 217 pin fuel assembly shown in figure 15 uses a conformal hexahedral mesh for the 1,520 geometrical volumes. In the future, more accurate models will require the resolution of the helical wire wrap, significantly complicating the mesh generation process. Once the mesh is generated, the physics that is modeled is a combination of thermal hydraulics, structural mechanics, and neutronics. A different simulation tool is used for each of these physical processes, but there is a strong need for them to interact with and exchange data with each other. The SHARP framework being developed at Argonne National Laboratory provides the infrastructure necessary to do this coupling on large-scale computer architectures (figure 16). ITAPS technology is used as the basis of the SHARP reactor simulation project and the common ITAPS interfaces simplify the coupling between thermal/hydraulics and neutronics physics modules. These interfaces, and the ITAPS approach in general, also serve as a means for collaborating with other institutions and for coupling with commercial codes.

SOURCE: T. TAUTGES, ANL ILLUSTRATION: A. TOVEY

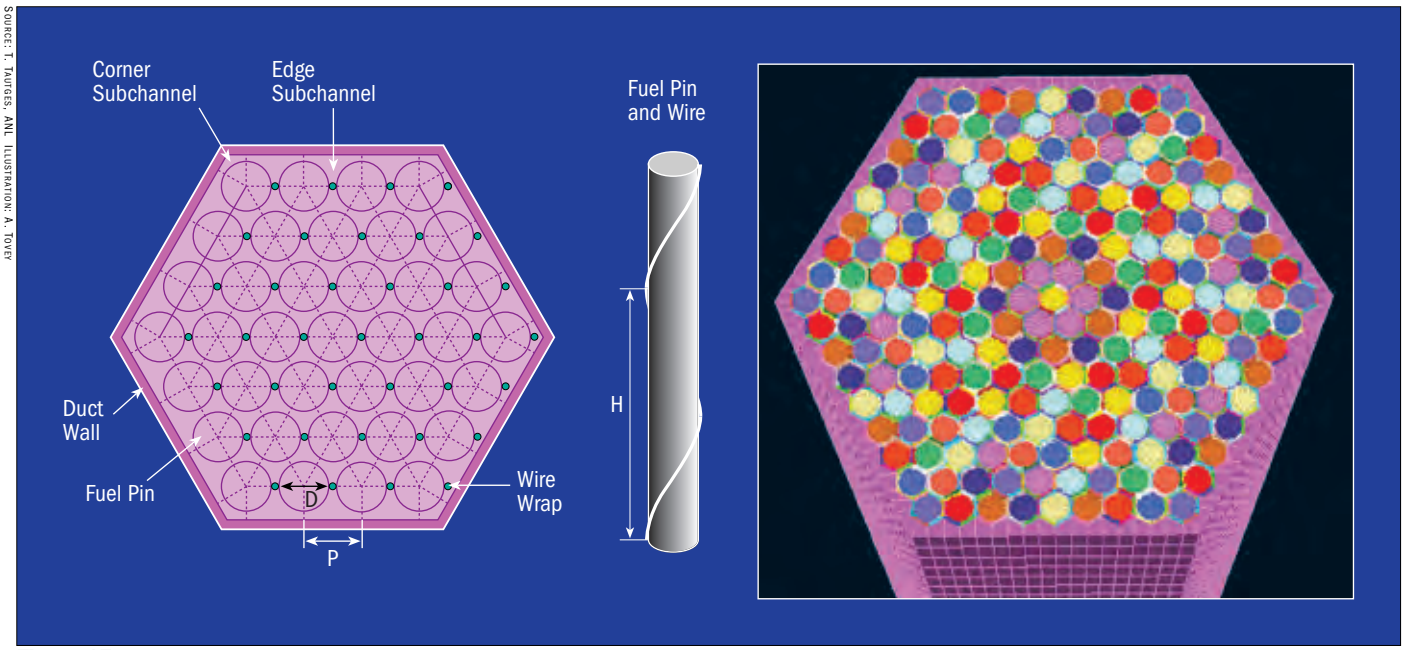


**Figure 15.** *The geometrical domain and a conformal hexahedral mesh for modeling the reactor core of a sodium fast reactor.*

tracking technologies to study the fueling process for ITER. A systematic study using these technologies provided new insight into pellet ablation, which is a key technology fueling ITER.

*Biological Modeling*

The ITAPS team has contributed to the development of the Virtual Microbial Cell Simulator (VMCS). ITAPS mesh generation and discretization technologies were used in simulations that provided new scientific insight into the flocculation behavior of Shewanella microbes in oxygen-rich environments by confirming that there is an oxygen gradient from the edges of the floc into the center. This collaboration with PNNL computational biologists is targeting DOE bio remediation problems in heavy metal waste.

*Nuclear Energy Simulation Codes*

There is a renaissance associated with the nuclear energy application area (sidebar "Next-Generation Nuclear Reactor Modeling"). Critical to the design of next-generation nuclear reactors are accurate simulations that can quantify the per-

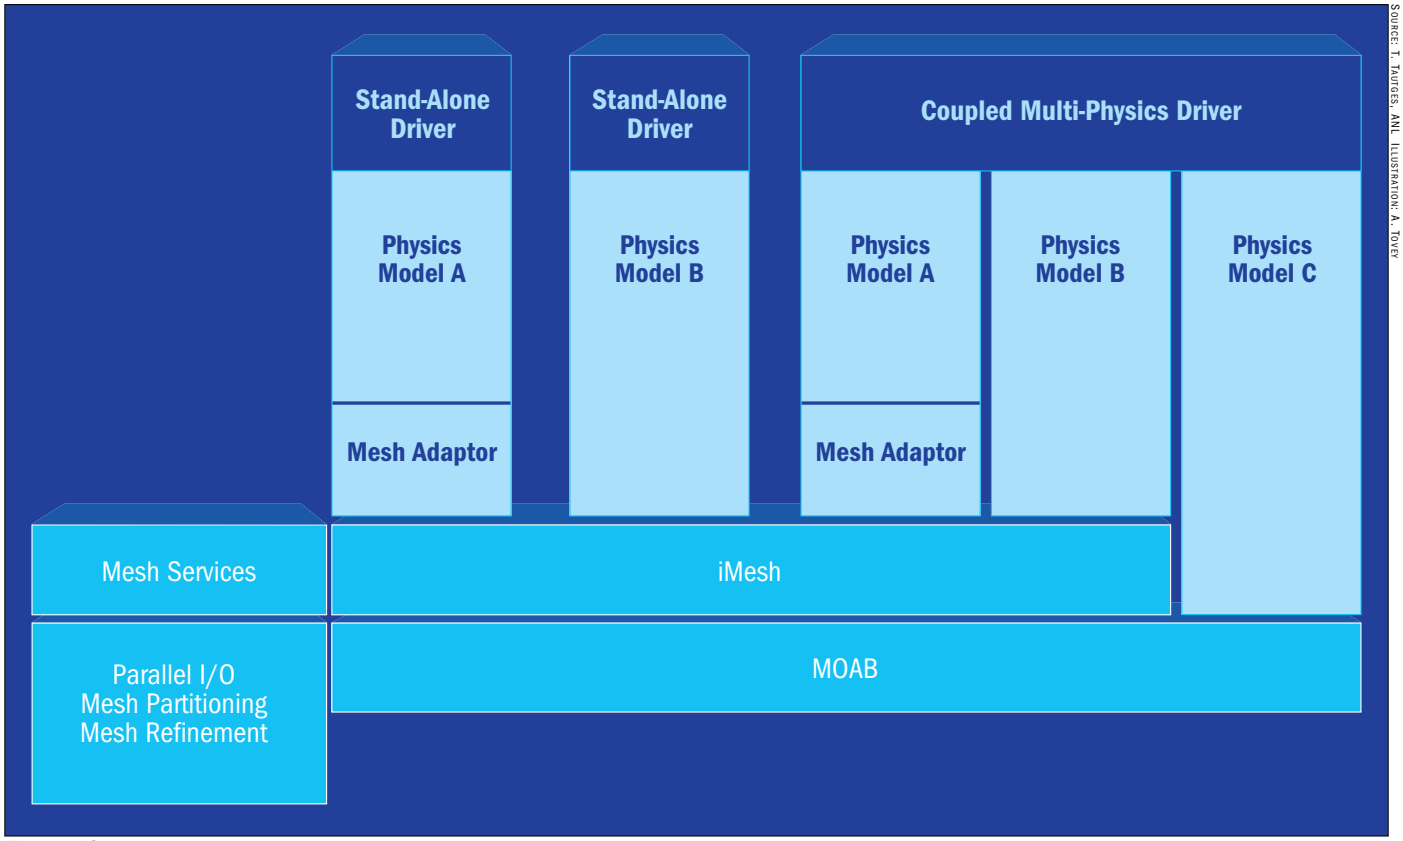


SOURCE: T. TAUGES, ANL ILLUSTRATION: A. TOVEY

**Figure 16.** *The SHARP framework developed at Argonne National Laboratory uses the ITAPS interfaces and tools as a key mechanism for coupling different physics modules together to perform high-fidelity simulations of nuclear reactors.*

formance of nuclear reactor components, characterize fuel components, and help quantify the design margins for safe, efficient operation. Simulating the processes of a nuclear reactor involve complex geometries, and coupled, multi-physics simulations. ITAPS researchers are providing mesh generation, partitioning, and mesh-to-mesh transfer tools in a next-generation reactor core modeling framework being developed at Argonne National Laboratory. We are also providing front tracking and adaptive technologies for the study of phase transitions and material relocation during hypothetical nuclear fuel disruptive accidents in Generation IV power plants.

## Outreach

The ITAPS team is proactively engaging the broader DOE community through coordinated team presentations and tutorials at major conferences along with peer-reviewed journal publications. We are actively encouraging and pursuing collaborations with new application areas, such as subsurface flow modeling, and expanding our collaborations with the accelerator and fusion modeling communities. More information regarding our research efforts and software tools can be found on the ITAPS website (see Further Reading).

## Summary

Simulation has become critical to reaching scientific goals in a number of application areas of importance to the DOE mission. A major component of this process is defining the computational domain, representing the resulting geometry, creating a high-quality initial mesh, adapting that mesh to capture the physics, and ensuring that this all works well on next-generation computers. The ITAPS team is providing an interoperable infrastructure that allows easy exploration of state-of-the-art tools that address these challenges. We are working closely with domain scientists in a number of application areas and have already had significant impact by enabling simulations that could not previously be performed, dramatically reducing total time to solution through the use of adaptive mesh refinement techniques, and increasing accuracy through front tracking methods and high-fidelity representations of the computational geometry.

**Contributors** Dr. Lori Diachin, Lawrence Livermore National Laboratory

## Further Reading

http://www.itaps-scidac.org