# Seeing the Unseeable

The SciDAC Visualization and Analytics Center for Enabling Technologies (VACET) is a highly productive effort combining the forces of leading visualization researchers from five different institutions to solve some of the most challenging data understanding problems in modern science. The VACET technology portfolio is diverse, spanning all typical visual data analysis use models and effectively balancing forward-looking research with focused software architecture and engineering resulting in a production-quality software infrastructure. One of the key elements in VACET's success is a rich set of projects that are collaborations with science stakeholders: these efforts focus on identifying and overcoming obstacles to scientific knowledge discovery in modern, large, and complex scientific datasets.

Galileo Galilei's improvements to early telescope design first opened up the heavens, the satellites of Jupiter, sunspots, and even the rotation of the Sun. He proved the Copernican heliocentric model of the solar system: that it is the Sun, rather than the Earth, which is the center of the solar system. Thus, the telescope became the first device to make the unseeable seeable.

Today, scientific visualization plays an equally significant role in contemporary science. Such visualization transforms abstract data into readily comprehensible images (figure 1) and has become an indispensable part of scientific discovery. The landmark report of 1987 by Thomas DeFanti and others first showed the important role of visualization in scientific discovery.

Visualization research tends to fall into one of three distinct categories: exploration visualization, analytical visualization, and presentation visualization. With exploration visualization, one has no idea what to look for. This model is typically the most challenging as success relies on interactive, "random-access" exploration of large and complex datasets. Analytical visualization requires that one knows what to look for in the data and is often employed after exploratory visualization has reduced the dataset. Presentation visualization conveys a specific concept to others. These visualizations are typically part of computational science presentations.

The U.S. Department of Energy's (DOE) investment in a broad range of scientific disciplines under the SciDAC program enables researchers to study scientific phenomena via simulation on some of the world's largest computer systems. These simulations are carried out on fractional-petascale-sized machines and generate vast amounts of output data. Managing and understanding such data is widely perceived as one of the bottlenecks in contemporary science, resulting in DOE's SciDAC program which addresses data management and knowledge discovery to complement the computational science.

These efforts are being driven by the DOE SciDAC Visualization and Analytics Center for Enabling Technologies (VACET), which is working on the daunting task of enabling discovery through visualization and analytics on some of

Today, scientific visualization plays a significant role in contemporary science. Such visualization transforms abstract data into readily comprehensible images and has become an indispensable part of scientific discovery.

DATA: W. WASHINGTON, NCAR; J. DRAKE AND F. HOFFMAN, ORNL IMAGE: J. DANIEL, ORNL

**Figure 1.** *Visualization offers the ability to "see the unseeable." The image shows the component of the atmospheric $CO_2$ concentration that results from the net ecosystem exchange (NEE), which is shown on the land surface. This "green $CO_2$" is the flux due to the respiration of vegetation, respiration of soil microbes, and fire, minus that taken up by ecosystem production.*

the world's largest and most complex datasets and computational platforms. As a Center for Enabling Technology, VACET's mission is production-quality visualization and knowledge discovery software to run on the large, parallel computer systems at DOE's open computing facilities in order to improve visual data exploration and knowledge discovery in modern science, especially in the DOE science community. This article endeavors to sum up a small part of the center's first-year accomplishments across a broad section of VACET's portfolio and reflects an effective balance between forward-looking, award-winning research and the software engineering needed for production-quality visual data.

## Forward-Looking Research

As in other areas of computational science, visualization too is strained by the ever-increasing amounts of data produced by simulations and experiments. Algorithms and implementations, effective with data of modest size and complexity, may not prove adequate for future datasets. Visualization is more than just the means to present scientific results in an appealing format. Its real potential lies in the integration of interactive visual exploration and analysis with exceptional computational capacity and the spectacular capabilities of the human mind.

The research summaries that follow represent interrelated challenges in contemporary science. The first presents recent work in the area of flow field visualization focusing on computing and displaying structure. Such techniques represent an important advance in the ability to extract and convey knowledge about complex data. This work has general applicability to a diverse range of science projects. The second presents recent work using a rigorous mathematical foundation as the basis for quantitative and comparative analysis of complex data. The idea is to characterize, identify, and analyze features using topology. This approach has the potential to accelerate scientific discovery as the technique that supplants a common but imprecise comparative analysis method known as "chi-by-eye." The third addresses an orthogonal set of challenges, namely

VACET's mission is production-quality visualization and knowledge discovery software to run on the large, parallel computer systems at DOE's open computing facilities in order to improve visual data exploration and knowledge discovery in modern science.

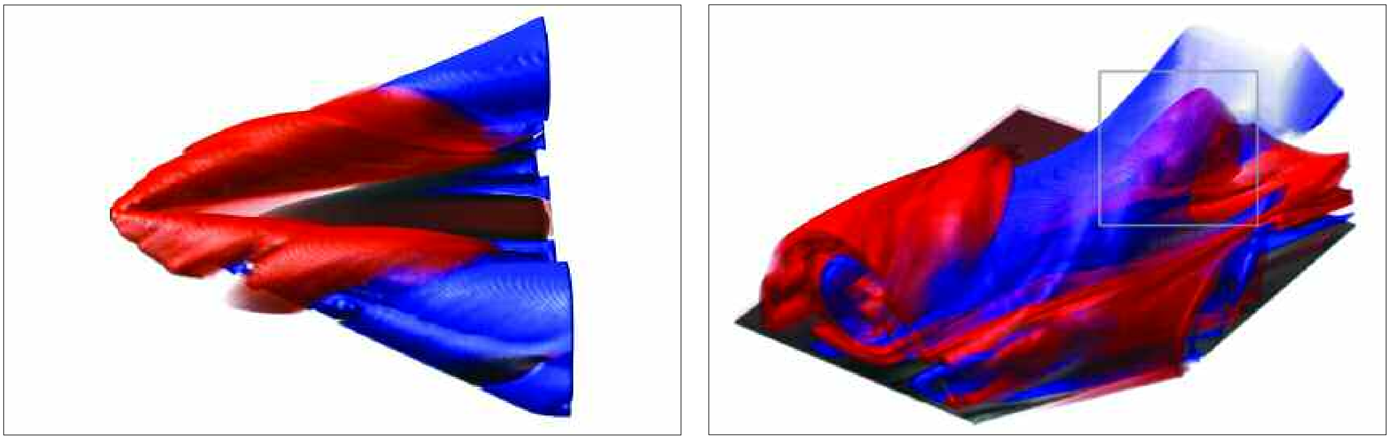

**Figure 2.** *Volume rendering of regions of high forward (red) and backward (blue) FTLE above a delta wing. Shown are the wing edge separation and the primary attachment layer (left), with inner structures occluded. The interplay of separation and attachment structures is visible on the front face (right). The gray box highlights the separation structure that characterizes a vortex breakdown bubble.*

Data: M. Rutten, German Aerospace Center Image: C. Garth, UC-Davis

Many scientific problems ranging from fluid dynamics and magnetohydrodynamics to climate and combustion research require the effective analysis of simulated flow data.

## Finite-Time Lyapunov Exponent

In dynamical systems, the Finite-Time Lyapunov Exponent (FTLE) measures the rate of separation of infinitesimally close trajectories. In the specific context of fluid flows, the FTLE value at any position can be computed from the spatial derivative of the flow map. This map indicates the new position reached by a massless particle after being transported along the flow (advection) during a finite time interval. Since the FTLE quantifies the dispersion of particles along the flow, that quantity can be computed both for forward and backward advection. High values of forward FTLE are found along repelling manifolds embedded in the flow, while high values of backward FTLE indicate the presence of an attracting manifold. Both types of manifolds determine the local behavior of the flow and are called Lagrangian Coherent Structures (LCS).

LCS are manifolds embedded in the flow that dynamically act on surrounding particles as separating surfaces. Hence, they provide natural boundaries to the coherent structures present in the flow. From a practical standpoint, their visualization requires the computation of a high-resolution, three-dimensional FTLE field in each direction. While high values of forward FTLE characterize the highly diverging behavior of neighboring particles under the action of a repelling LCS, high values of backward FTLE are indicative of a strongly converging pattern among particles that originated at widely distant positions.

simplifying creation, modification, and reuse of visualization technologies. This award-winning work takes the first steps toward reducing the amount of labor and specialized knowledge required for creating visualization "pipelines."

### Efficient Computation of Coherent Structures in Fluid Flow Applications

Many scientific problems ranging from fluid dynamics and magnetohydrodynamics to climate and combustion research require the effective analysis of simulated flow data. As the size and complexity of the corresponding vector fields grow, the efficient extraction of their salient structures becomes essential. The notion of Finite-Time Lyapunov Exponents (FTLE; sidebar "Finite-Time Lyapunov Exponent") provides a sound theoretical framework to characterize Coherent Lagrangian Structures in transient flows. Despite its conceptual simplicity, the associated computational cost is prohibitive. VACET has developed a novel approach for adaptive computation of FTLE fields in two and three dimensions that significantly reduces the computational cost. Meaningful results for three-dimensional flows can be obtained by restricting the analysis to a well-chosen plane. Moreover, the examination of some of the visualization aspects of FTLE and the introduction of several new methods have improved the analysis of specific aspects of challenging datasets.

To address the challenge raised by the size and the qualitative complexity of flow vector fields resulting from modern computational fluid dynamics (CFD) computations, scientific visualization research has explored different approaches that characterize, extract, and visually represent salient flow structures across spatial and temporal scales (figure 2). These methods are mainly divided into topological

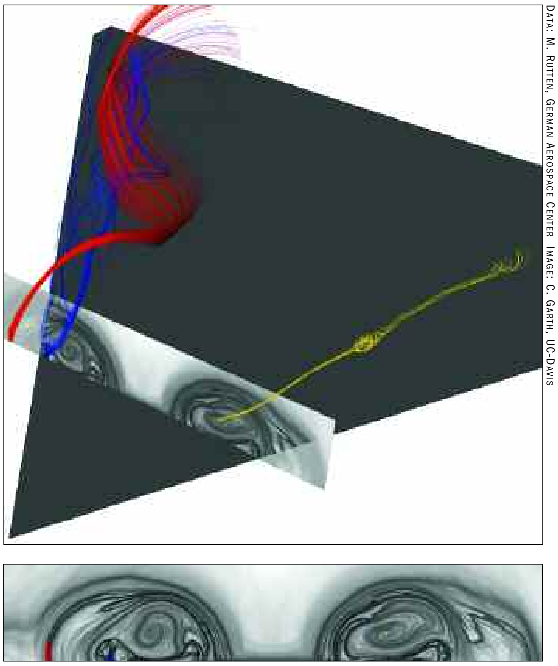

Data: M. Rutten, German Aerospace Center  Image: C. Garth, UC-Davis

**Figure 3.** *Pathlines are seeded according to FTLE strength, FTLE ridge lines, or via a user-guided probability density function (PDF). The top image shows visualization of primary (red) and secondary (blue) separation structures; pathlines are seeded according to the PDF shown in the lower image, which shows the planar FTLE visualization on a section plane perpendicular to the main flow direction. Darker regions correspond to regions of high FTLE. Colored regions indicate PDFs used to see the pathlines in the upper right image.*

and feature-based approaches. While the former leverages a sound mathematical framework and allows for an objective and fully automatic post-processing, the latter explicitly integrates significant flow structures into the analysis at the cost of ambiguous definitions and *ad hoc* methods.

In this context, the notion of Lagrangian Coherent Structures (LCS) and its quantitative assessment using the FTLE provide a promising alternative that combines a well-articulated theoretical basis with physical intuition. Specifically, coherency in steady and transient flows can be characterized in terms of repelling and attracting manifolds. Despite the versatility and consistency of this approach, its practical application is fundamentally hampered by a prohibitive computational cost associated with the required advection of a dense set of particles across the spatio-temporal flow domain.

VACET's work has made three significant contributions to the field of visual data analysis. First, VACET has achieved a lower computational cost by significantly reducing the number of particle advections required to perform visualization and analysis based on FTLE and LCS. It has developed an incremental, data-driven refinement algorithm which exploits the coherence of neighboring particle paths to generate smooth approximations of the so-called flow map from which the FTLE is computed. This approach enables high-resolution analysis of complex 4D flows and permits the construct of insightful visualization for accurate assessment of coherence. Second, VACET has proven that it is often not necessary to perform a full 3D analysis: given limited problem-specific knowledge about the flow field it is often sufficient and in some cases even beneficial to consider FTLE on 2D subsections, further reducing compute time (figure 3). Third, VACET has demon- strated several new visualization methods based upon these new techniques with data from large-scale CFD simulations.

VACET's work has made three significant contributions to the field of visual data analysis.

## Topologically Based Feature Detection, Tracking, and Quantitative Analysis

When a heavy fluid is placed above a light fluid, tiny vertical perturbations in the interface create a characteristic structure of rising bubbles and falling spikes known as Rayleigh–Taylor instability. Rayleigh–Taylor instabilities have received much attention over the past half-century because of their importance in understanding many natural and man-made phenomena, ranging from the rate of formation of heavy elements in supernovae to the design of capsules for inertial confinement fusion. VACET has developed a new, robust method for quantitative analysis of Rayleigh–Taylor instabilities whereby a hierarchical segmentation of the mixing envelope surface is extracted to identify bubbles and analyze analogous segmentations of fields on the original interface plane. This approach is based on a family of robust topological techniques that enable multiscale segmentation of scientific data for feature extraction and error-bounded quantitative analysis.

To overcome the challenge of analyzing the complex topology of the Rayleigh–Taylor mixing layer, VACET has developed a novel approach based on robust Morse theoretical techniques. This approach systematically segments the envelope of the mixing interface into bubble structures (figure 4, p28) and represents them with a new multi-resolution model allowing for the first time a multi-scale quantitative analysis of the rate of mixing based on bubble count. The analysis highlighted and provided precise measures for four fundamental stages in the turbulent mixing process that scientists previously could only observe qualitatively. This approach has led to new insights and a deeper understanding of this fundamental phenomenon (figure 5, p28).

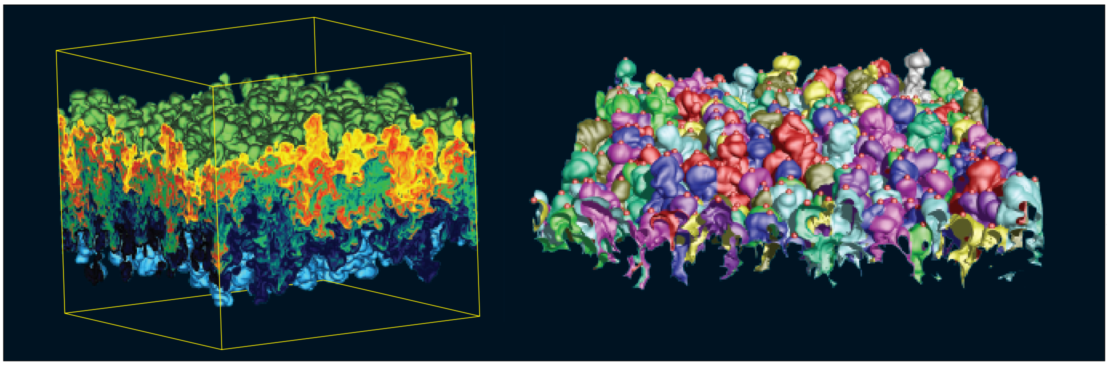

**Figure 4.** *Turbulent mixing interface of a Rayleigh–Taylor instability simulation (left). Topological segmentation of the upper envelope highlighting in different colors the bubbles that rise during the mixing process (right).*

Data: A. Cook and W. Cabot, LLNL, SciDAC Science Application "Simulation of Turbulent Flows Hit with Strong Shockwaves" Image: V. Pascucci, LLNL

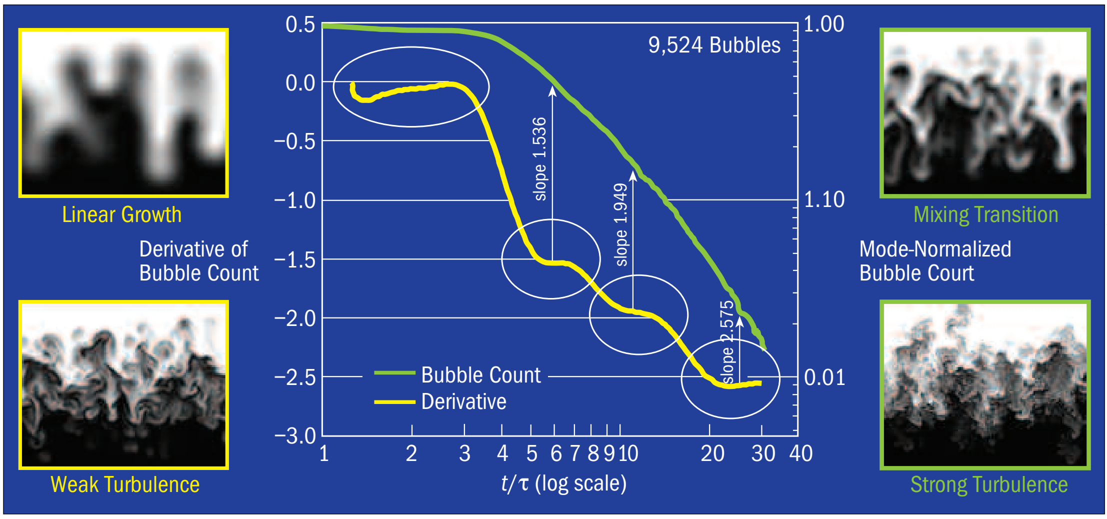

**Figure 5.** *Time analysis of the bubble structures in the Rayleigh–Taylor mixing interface. Our approach both highlights qualitatively the four main stages of the process and quantifies the mixing rates characterizing each stage.*

Data: A. Cook and W. Cabot, LLNL, SciDAC Science Application "Simulation of Turbulent Flows Hit with Strong Shockwaves" Image: V. Pascucci, LLNL

This work has been documented in a paper, which won the Best Application Paper award at IEEE Visualization 2006 and was later presented at the International Workshop on the Physics of Compressible Turbulent Mixing. Follow-up work enabled the first-ever direct comparison of two simulations based on different physics models and run with different initial conditions: the first run with one billion nodes over 758 time steps, the second run with 27 billion nodes over 220 time steps. Although comparison by superposition (for example, "chi-by-eye") of the two simulations could not yield any meaningful result, the topological approach provided a quantitative multi-scale, feature-based comparison highlighting fundamental similarities (figure 6), which validated the lower-resolution large-eddy simulation (LES) with respect to the higher-resolution direct numerical simulation (DNS).

## Querying and Creating Visualizations by Analogy

While there have been advances in visualization systems, particularly in multi-view visualizations and visual exploration, the process of building visualizations remains a major bottleneck in data exploration. A useful paradigm for building visualization applications is the dataflow model. A dataflow is a directed graph where nodes represent computations and edges represent streams of data: each node or module corresponds to a procedure that is applied on the input data and generates some output data as a result. The flow of data in the graph determines the order in which the processing nodes are executed. In visualization, it is common to refer to a dataflow network as a "visualization pipeline."

A useful paradigm for building visualization applications is the dataflow model.

VACET's work has been able to show that provenance metadata collected during the creation of pipelines can be reused to suggest similarities in related visualizations and guide semi-automated changes. To enable the effective reuse of computational (visualization) pipelines, VACET has introduced the idea of query-by-example in the context of an ensemble of visualizations, and the use of analogies as first class operations in a system to guide scalable interactions. This work, which is part of VACET's forward-looking research portfolio, received the prestigious Best Paper award at the IEEE Visualization 2007 conference.

Most visualization dataflow-based systems (such as AVS, SCIRun) have sophisticated user interfaces with visual programming capabilities that ease the creation of visualizations. Nonetheless, the path from "data to insight" requires a laborious trial-and-error process, where users successively assemble, modify, and execute pipelines. In the course of exploratory studies, users often build large collections of visualizations, each of which helps in the understanding of a different aspect of their data. A scientist working on a CFD application might need different visualizations such as 3D isosurface plots, 2D plots, and direct volume-rendering images.

Although in general each of these visualizations is implemented in a separate dataflow, there is a certain overlap, meaning they manipulate the same input datasets. Furthermore, for a particular class of visualizations, scientists generate several different versions of each individual dataflow while fine tuning visualization parameters or experimenting with different datasets.

VisTrails (see Further Reading, p33) is an open-source system which implements a provenance model that uniformly captures changes to pipeline and parameter values during the course of data exploration. This detailed history, combined with a multi-view visualization interface, streamlines the exploration process. It allows users to navigate a large number of visualizations, giving them the ability to return to previous versions, compare different pipelines and their results, and then resume their explorations. This provenance information can also be used to simplify and partially automate the construction of new visualizations.

VACET has proposed a new framework to reuse this knowledge to better assist users in performing data exploration through visualization. The process of applying pipeline differences (like a patch) to derive new pipelines can be automated in VisTrails in a process called "visualization creation by analogy." The framework consists of two key components: an intuitive interface for querying dataflows and a novel mechanism for semi-automatically creating and refining visualizations by analogy. The query engine is employed through a query-by-example interface whereby users query dataflows through the same familiar interface they used to create the dataflows (figure 7, p30). This approach allows for searching a large number of visualizations and identifing pipelines that satisfy the user-defined criteria.

While the query interface allows users to identify pipelines relevant to a particular task, the visualization by analogy component provides a mechanism for reusing these pipelines to construct new visualizations in a semi-automated manner—without requiring the user to manipulate the specifications of the pipeline. To apply an analogy, VisTrails first determines the difference between a source pair of analogous visualizations, and then transfers this difference to a third visualization (figure 8, p30). The user is not required to know the exact details of the dataflows in order to modify them. In addition, the analogy mechanism provides the basis for scalable updates: analogies can be automatically applied to many pipelines simultaneously. Together, the abilities to query visualization pipelines by example and to refine them by analogy are a significant step toward scalable pipeline development in visualization systems.

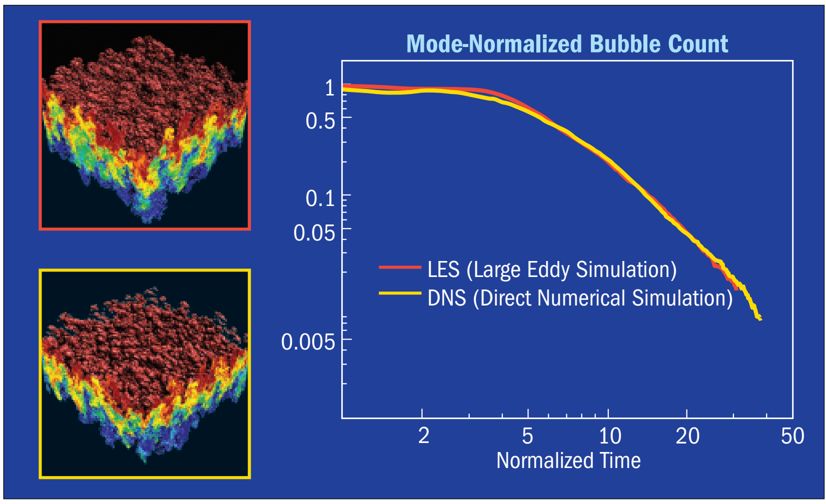

**Figure 6.** *Feature-based comparison of two Rayleigh–Taylor instability calculations based upon different models and run with different initial conditions. The results shown here validate the large-eddy simulation (top) with respect to the direct numerical simulation (bottom).*

DATA: A. COOK AND W. CABOT, LLNL, SCIDAC SCIENCE APPLICATION "SIMULATION OF TURBULENT FLOWS HIT WITH STRONG SHOCKWAVES" IMAGE: D. LANEY, LLNL

## Production Software: Building on Proven Technology

VACET has forged long-term relationships with science stakeholders who are able to define what kinds of information they hope to mine from their massive datasets as well as their methodology for hypothesis testing. With that, VACET can identify the technologies needed to achieve such capability. In some cases, VACET can adapt or extend existing technology, while in others, it must conceive of new technology. It is those ideas that have to be translated into practice—the production quality, petascale capable visual data analysis software.

Most software developers would agree that this objective represents a formidable amount of software engineering. VACET has adopted a low-risk, fast time-to-solution approach, built upon proven technology. Its team uses two primary delivery vehicles described below. Both are visualization

VACET has forged long-term relationships with science stakeholders who are able to define what kinds of information they hope to mine from their massive datasets as well as their methodology for hypothesis testing.

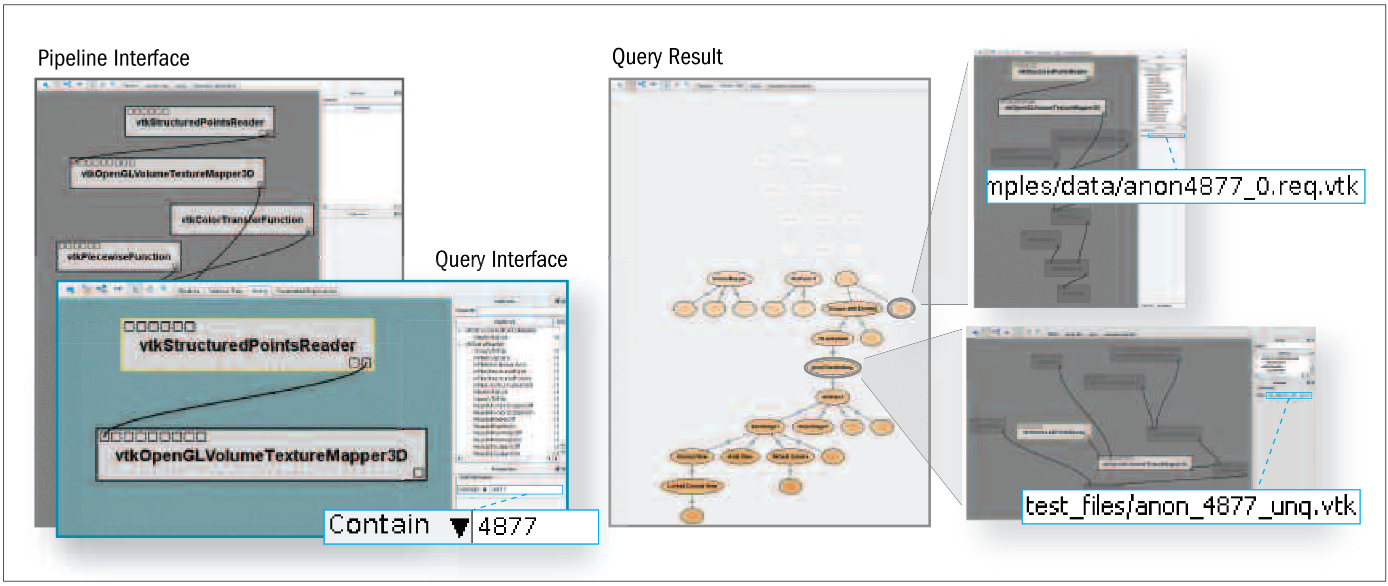

C. Scheidegger and C. Silva, University of Utah

**Figure 7.** *Querying by example. The interface for building a query over an ensemble of pipelines is the same as the one for constructing and updating pipelines. In fact, they work together: portions of a pipeline can become query templates by directly pasting them onto the Query Canvas. In this figure, the user is looking for a volume-rendered image of a file whose name contains the string 4877. The system highlights the matches both at the visualization level (version tree, shown in the middle) and at the module level (shown in the right insets).*

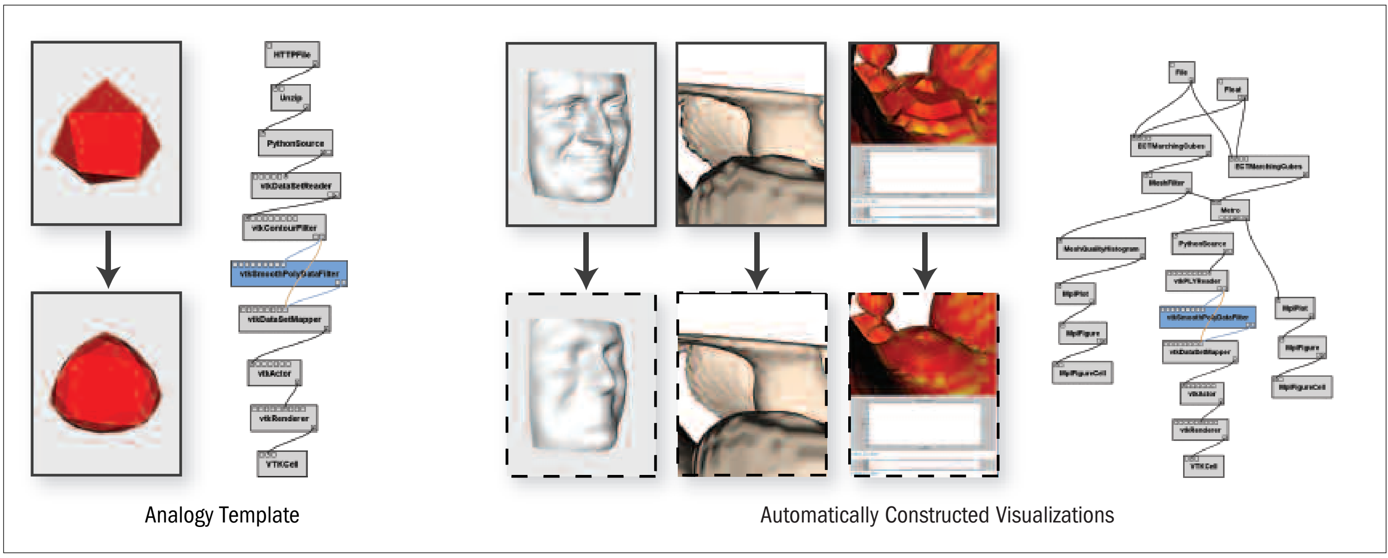

C. Scheidegger and C. Silva, University of Utah

**Figure 8.** *Visualization by analogy. The user chooses a pair of visualizations to serve as an analogy template. In this case, the pair represents a change where a file downloaded from the WWW is smoothed. Then, the user chooses a set of other visualizations that will be used to derive new visualizations. These new visualizations are derived automatically. The pipeline on the left reflects the original changes, and the one on the right reflects the changes when translated to the last visualization on the right. The pipeline pieces to be removed are portrayed in orange, and the ones to be added in blue. Note that the surrounding modules do not match exactly: the system figures out the most likely match.*

VACET can quickly add new capabilities to infrastructure that is production-quality and petascale capable.

applications, the result of decades of research and development. In this way, VACET can quickly add new capabilities to infrastructure that is production-quality and petascale capable. This strategy has proven effective and has delivered production-quality visual data analysis for use on Adaptive Mesh Refinement (AMR) data, as described below.

### VisIt Provides Key to Massive Datasets

VisIt is an open-source, turnkey application for large-scale simulated and experimental datasets. Its charter goes beyond pretty pictures; the application is an infrastructure for parallelized, general post-processing of extremely massive datasets. Target-use cases include data exploration, comparative analysis, visual debugging, quantitative analysis, and presentation graphics.

Many software developers from different areas worked together to make VisIt into a single-package product. VisIt leverages several third-party libraries: the Qt widget library for its user interface, the Python programming language for a command line interpreter, and the Visualization ToolKit (VTK) library for its data model and many of its visualization algorithms. Additionally, 50 man-years worth of effort have been devoted to the development of VisIt itself. The VisIt-specific effort has largely been focused on parallelization for large datasets, user interface, implementing custom data analysis routines, addressing non-standard data models (such as AMR and mixed materials zones), and creating a robust overall product. VisIt consists of over one-and-a-half million lines of code, and its third-party libraries have an additional one million lines of code. It has been ported to Windows, Mac, and many UNIX variants, including AIX, IRIX, Solaris, Tru64, and, of course, Linux, including ports for SGI's Altix, Cray's XT4, and many commodity clusters.

The basic design is a client–server model, where the server is parallelized. The client–server aspect allows effective visualization in a remote setting, while the parallelization of the server allows the largest datasets to be processed interactively. This tool has been used to visualize many large datasets, including a 27 billion data point structured grid (figure 9), a one billion point particle simulation, and curvilinear, unstructured, and AMR meshes with hundreds of millions to billions of elements. The most common form of the server is as a standalone process that reads in data from files. However, an alternative form exists where a simulation code can link in "lib-VisIt" and itself become the server, allowing for *in situ* visualization and analysis.

VisIt follows a dataflow network paradigm where interoperable modules are connected to perform custom analysis. The modules come from VisIt's five primary user interface abstractions, each with its own multitude of variations: 21 "plots" (ways to render data), 42 "operators" (ways to manipulate data), 85 file format readers, over 50 "queries" (ways to extract quantitative information), and over 100 "expressions" (ways to create derived quantities). Further, a plugin capability allows for dynamic incorporation of new plot, operator, and database modules. These plugins can be partially code generated, even including automatic generation of Qt and Python user interfaces.

The VisIt project originated at Lawrence Livermore National Laboratory (LLNL) as part of the Advanced Simulation and Computing (ASC) program of DOE's National Nuclear Security Agency (NNSA), but it has become a distributed project being further developed by several groups. Major hubs for the project come from VACET, ASC, and the Global Nuclear Energy Partnership (GNEP) from DOE's Office of Nuclear Energy. The VisIt project has 20 developers from many organizations and universities, including five DOE laboratories. VisIt received an R&D 100 Award in 2005 and is downloaded approximately 25,000 times per year.

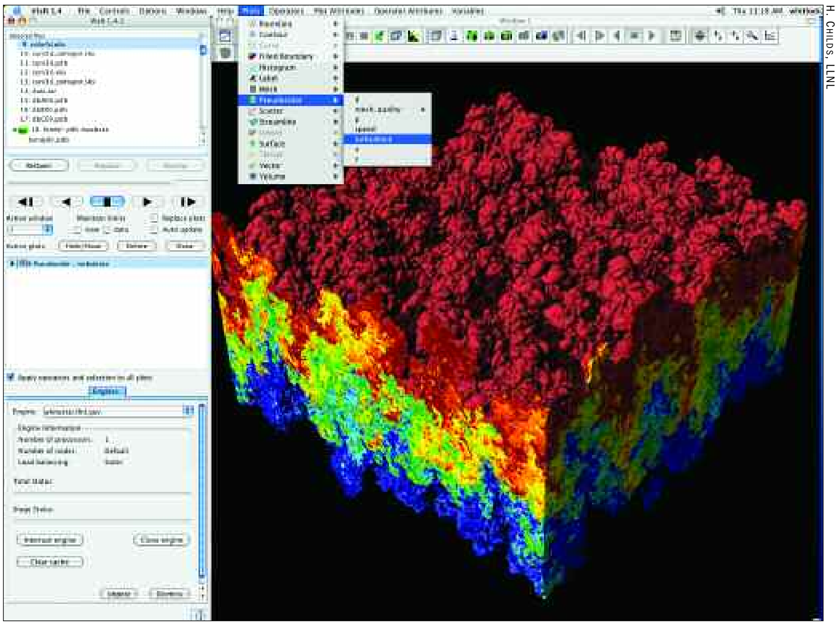

**Figure 9.** *VisIt was used to visualize this 27 billion element Rayleigh–Taylor instability, which simulates the mixing of heavy and light fluids. This simulation was run by the MIRANDA code on the NNSA Blue Gene/L machine.*

H. Childs, LLNL

## SCIRun Facilitates Large-Scale Computation, Visualization

SCIRun is a scientific problem-solving environment (PSE) that allows interactive construction and steering of large-scale scientific computations. A scientific application is constructed by connecting computational elements or modules to form a program or network. The program may contain several computational elements as well as several visualization elements, all of which work together to orchestrate a solution to a scientific problem. SCIRun is designed to facilitate large-scale scientific computation and visualization on a wide range of architectures from the desktop to large supercomputers. Geometric inputs and computational parameters may be changed interactively, and the interface provides immediate feedback to the investigator.

SCIRun is used to support the efforts of the SciDAC Center for Extended Magnetohydrodynamic Modeling in their analysis of the instabilities of magnetic fields that confine the burning plasma in fusion devices. Within SCIRun, tools have been developed to rapidly create and analyze Poincaré plots that show the behavior of the magnetic field-lines which have a periodic or quasi-periodic behavior as shown in figure 10 (p32). SCIRun is also employed to produce query-driven

VisIt received an R&D 100 Award in 2005 and is downloaded approximately 25,000 times per year.

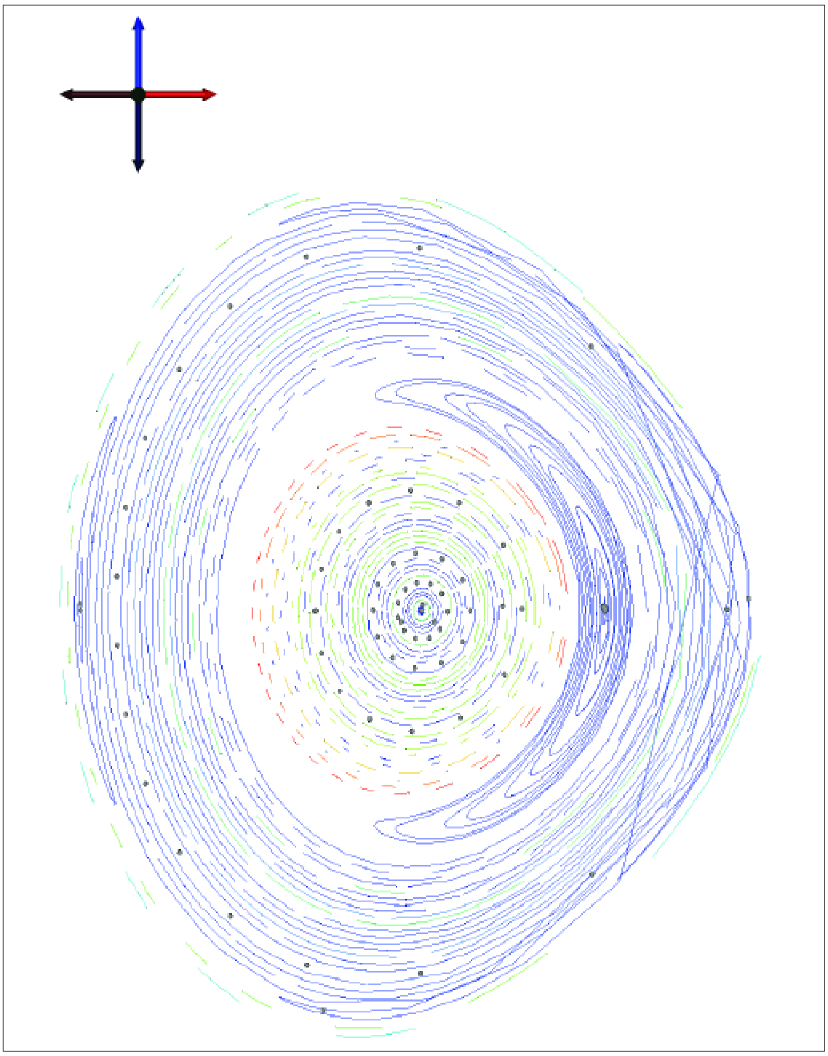

**Figure 10.** *A Poincaré plot of magnetic field instabilities from a MHD simulation of a tokamak simulation. The instabilities appear in this image as a large "banana-shaped" region located in the middle of this cross section.*

Data: S. Kruger, Tech-X Corporation and the SciDAC Center for Extended Magnetohydrodynamic Modeling Image: Scientific Computing and Imaging Institute, University of Utah

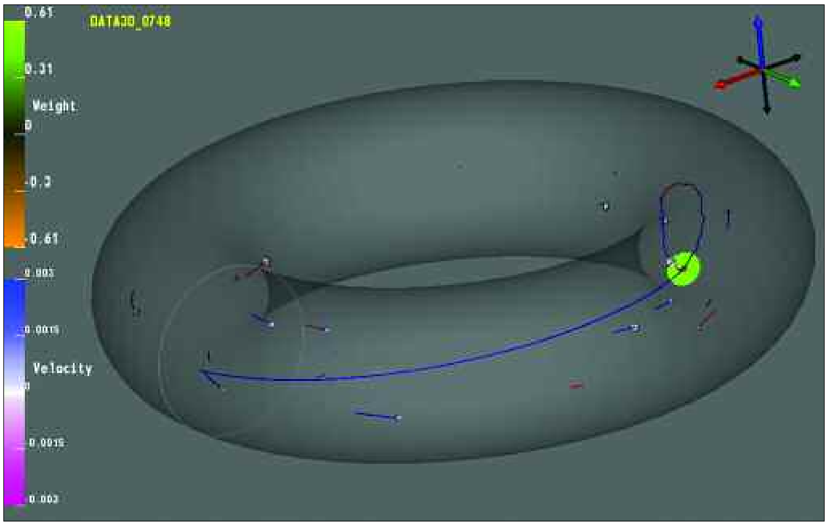

**Figure 11.** *Only 22 of the 400 million particles in a particle-in-cell simulation are being displayed based on the number of times that they are magnetically trapped (red line) and de-trapped (blue line) in relation to the externally imposed magnetic field. One particle has been highlighted by the electric potential that surrounds it throughout the simulation due to its extreme trajectory.*

Data: S. Ethier, Princeton Plasma Physics Laboratory, SciDAC Center for Gyrokinetic Particle Simulations of Turbulent Transport in Burning Plasmas Image: Scientific Computing and Imaging Institute, University of Utah

visualization of particle-in-cell simulations that are part of the SciDAC Center for Gyrokinetic Particle Simulations of Turbulent Transport in Burning Plasmas. Here, physicists are interested in analyzing just a few out of millions of particles that contribute to turbulent transport. The query-driven aspects of SCIRun allow physicists to isolate and visualize these "trapped" particles over hundreds of time steps as shown in figure 11.

### AMR Visual Data Analysis

AMR is a highly effective simulation method for spanning a large range of spatio-temporal scales, such as astrophysical simulations that must accommodate ranges from interstellar to sub-planetary. Most mainstream visualization tools lack the necessary support for AMR as a first-class data type and AMR code teams use custom-built applications for AMR visualization. VACET has provided significant enhancements to one of its technology pillars—VisIt—to provide the kind of production-quality, parallel-capable AMR visual data analysis infrastructure needed by SciDAC scientists who use AMR-based simulations.

As a result, at least one SciDAC team, the Applied Partial Differential Equations Center (APDEC; "APDEC: Algorithms and Software for Discovery," *SciDAC Review*, Summer 2007, p22), has migrated to this new platform for most of its day-to-day work, thereby realizing a substantial cost savings: they no longer expend their own effort toward developing and maintaining AMR-capable visual data analysis software.

AMR techniques combine the compact, implicitly specified structure of regular, rectilinear grids with the ability to adapt to changes in scale of unstructured grids. Handling AMR data for visualization is challenging, since coarser information in regions covered by finer patches is superseded and replaced with information from these finer patches. During visualization, it becomes necessary to manage selection of which resolutions are being used for any given operation. Furthermore, it is difficult to avoid discontinuities at level boundaries, which, if not properly handled, lead to visible artifacts in visualizations. Because of these difficulties, support for AMR as a first-class data type in production visualization tools has been lacking despite the growing popularity and usefulness of AMR simulations.

VisIt, however, accommodates AMR as a first-class data type. It handles AMR data as a special case of "ghost data," that is, data that are used to make computations more efficient, but that are not considered to be part of the simulation result. VisIt tags cells in coarse patches that are available at finer resolution as "ghost" cells, allowing AMR patches to retain their highly efficient native

Data: J. Bell and M. Day, Center for Computational Sciences and Engineering, LBNL Images: G. Weber, LBNL

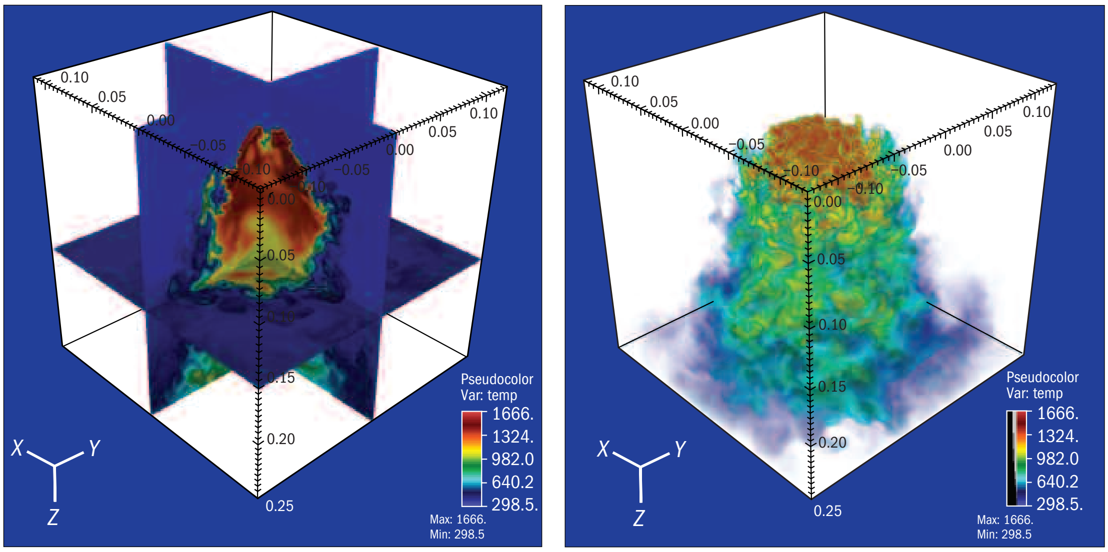

**Figure 12.** *Production-quality visualization of an AMR simulation of a hydrogen flame. The left panel shows a pseudocolor plot restricted to three axis-perpendicular slices. The right panel shows a volume-rendered image of the same data.*

VisIt offers a rich set of production-quality functions, like pseudocolor and volume-rendering plots, for visualization and analysis of complex datasets on parallel platforms, making it an ideal candidate to replace specialized AMR visualization tools.

format as rectilinear grids. It offers a rich set of production-quality functions, like pseudocolor and volume-rendering plots (figure 12), for visualization and analysis of complex datasets on parallel platforms, making it an ideal candidate to replace specialized AMR visualization tools.

Most of the work focused on implementing a set of essential debugging features offered by ChomboVis in VisIt to improve VisIt's handling of AMR data, both in terms of interface and performance.

ChomboVis provides spreadsheet "plots" that support direct viewing of numerical values on a particular slice of a patch. This function is essential for debugging and it is used by AMR code development teams on a daily basis. VACET added these spreadsheets to VisIt as shown in figure 13 and connected them to VisIt's "pick cell" feature allowing users to "link" them to other plots. We further added a capability to dynamically create new buttons in the VisIt interface to perform custom actions. This matches a capability that APDEC users valued in ChomboVis and allows new users to quickly navigate the tool. VACET also modified the VisIt selection routines to better support AMR data, allowing users to specify selections in terms of cell indices in a particular AMR level.

VACET has optimized the handling of AMR grids in VisIt. These optimizations save on memory by a factor of ten and also support more efficient rendering. Additional performance and memory optimizations improve efficiency for rendering patch boundaries. VisIt previously used a very general algorithm that was unnecessarily slow. The new, specialized algorithm is an order of magnitude faster and more memory efficient. ●

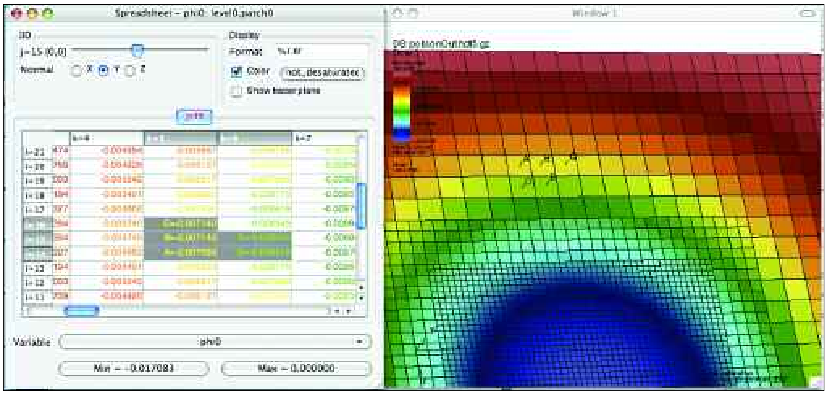

Variable phi0

Min = −0.017083

Max = 0.000000

Data: P. Colella and B. van Straalen, LBNL Image: G. Weber, LBNL

**Figure 13.** *Spreadsheet plots are an important tool for debugging AMR codes. They support direct viewing of numerical data in patch cells. VisIt labels selected cells in both spreadsheet and 3D visualizations allowing users to recognize correspondences quickly and effectively.*

**Contributors:** E. Wes Bethel at Lawrence Berkeley National Laboratory (LBNL) and Dr. Chris Johnson at the Scientific Computing and Imaging Institute, University of Utah; Dr. Charles Hansen, Dr. Claudio Silva, Dr. Steven Parker, Dr. Allen Sanderson, Lee Myers, and Martin Cole at the Scientific Computing and Imaging Institute, University of Utah; Dr. Xavier Tricoche at Purdue University; Sean Ahern, Dr. George Ostrouchov, Dr. Dave Pugmire, Jamison Daniel, and Jeremy Meredith at Oak Ridge National Laboratory; Dr. Valerio Pascucci, Dr. Hank Childs, Dr. Peer-Timo Bremer, and Dr. Ajith Mascarenhas at Lawrence Livermore National Laboratory; Dr. Ken Joy, Dr. Bernd Hamann, and Dr. Christoph Garth at the University of California–Davis; Dr. Cecilia Aragon, Dr. Gunther Weber, and Prabhat at LBNL

## Further Reading

VACET
www.vacet.org

VisTrails
www.vistrails.org