Editors: Claudio Silva, csilva@cs.utah.edu
Joel E. Tohline, tohline@rouge.phys.lsu.edu

VISUALIZING MULTIVARIATE VOLUME DATA FROM TURBULENT COMBUSTION SIMULATIONS

By Hiroshi Akiba, Kwan-Liu Ma, Jacqueline H. Chen, and Evatt R. Hawkes

To understand dynamic mechanisms, scientists need intuitive and convenient ways to validate known relationships and reveal hidden ones among multiple variables.

Recent developments in numerical methodology for combustion simulations that effectively harness modern high-performance parallel computers can simulate reacting flows by using high-fidelity models for the underlying complex processes. However, a single run of the simulation can produce multiple terabytes of raw data that are vast in the spatial (near a billion grid points), temporal (100,000 time steps), and variable (tens of variables) domains, creating a formidable challenge for subsequent analysis and interpretation. In addition to the data set's sheer size, the difficulty of knowledge extraction is compounded by the complexity of the turbulent flow fields and the phenomena under study, as well as by the different data types (particle and field data). To understand the dynamic mechanisms of extinction and reignition in turbulent flames, for example, scientists need intuitive and convenient ways to validate known relationships and reveal hidden ones between multiple variables.

In a collaboration between the University of California, Davis and Sandia National Laboratories in Livermore, California, our team developed interactive visualization techniques and an interface design that enable validation and improved understanding of turbulent combustion simulations, letting scientists study potentially complex relationships between multiple variables. Using a high-fidelity *direct numerical simulation* (DNS) of a turbulent nonpremixed flame with the detailed chemistry performed to date, we show how the resulting visualizations present information previously unseen, providing new understanding and validation of such simulations.

Turbulent Combustion Simulations

DNS is a mature and productive research tool in combustion science used to provide high-fidelity, computer-based observations of the microphysics found in turbulent, reacting flows. DNS is also a unique tool for developing and validating the reduced model descriptions used in macroscale simulations of engineering-level systems. DNS can fully resolve a turbulent combustion process's finest spatial and temporal details, focusing on a particular physical aspect of the process by adopting simple "building block" flows, such as homogeneous turbulence or free shear layers.

The Sandia 3D Direct Numerical Solver (S3D) is an existing Sandia DNS compressible Navier-Stokes solver coupled with an integrator for detailed chemistry.[1] Presently, S3D produces 10 Tbytes of raw data per run in a parametric study of several runs, presenting a significant challenge for subsequent analysis and interpretation. The current approach is to store the raw data at prescribed time intervals, then either analyze it on a supercomputer or access it via a local analysis and visualization cluster at Sandia. The raw data is archived so that various modeling groups in the combustion community can interrogate it to test model assumptions. To understand the correlation of scalar fields such as temperature, mixing rates, and species concentrations in turbulent flames, scientists must be able to visualize two or more scalars simultaneously. Conventional visualization tools don't directly support such a capability. Scientists often must make side-by-side comparisons of images of different variables by hand, which is tedious and time-consuming. Furthermore, the information that scientists can derive by looking at separate images is quite limited. Thus, they need effective methods for simultaneously visualizing multiple time-varying variables from large data sets in an interactive fashion.

Simultaneous Visualization

In many fields not limited to combustion, the capability to simultaneously visualize different variables that describe the same spatial domain and to determine their correlations is very desirable. In a multidisciplinary comput-

 Copublished by the IEEE CS and the AIP 1521-9615/07/$25.00 © 2007 IEEE

ing environment, for example, several engineering analysis programs—such as structural and flow solvers—run concurrently and cooperatively to perform a multidisciplinary design. The goal might be to identify the relevant design variables that can explain a particular phenomenon's causes, such as vortices in a flow field.

Another example is medical research and practice, in which many imaging modalities, including X-ray imaging, Doppler ultrasound, computed tomography scans, magnetic resonance imaging (MRI) scans, and so on, are available for making better diagnoses and surgical plans. Physicians generally look at these images from different modalities individually or side by side, but it would help if they could see two types of data sets simultaneously by superimposing them spatially. An example would be to visualize functional data overlaid on anatomical data.

Volume rendering has been accepted as an effective method for visualizing physical phenomena or structures defined in 3D space. Consumer graphics hardware support for 3D textures has enabled interactive volume rendering, which is very attractive to many disciplines. However, previous volume visualization algorithms were mostly designed for looking at one volume at a time or for animating time-varying volume data. Such data from different modalities almost always has either different resolutions or various degrees of distortions, thus the user must first do a registration step to align the data points. Similarly, in many scientific or engineering studies, researchers must look at simulation data and experimental data of different resolutions or on different kinds of computational mesh structures together, and then resample them to match one resolution to the other.

To visualize data from S3D, however, we're concerned only with rendering scalar volume data on the same mesh, which is a simpler problem than multimodality volume visualization in medical imaging or comparative visualization of simulated scientific data from different codes.[2]

The techniques we developed can generate effective visualizations that reveal interaction and cause–effect between different scalar properties of the same flow domain. Our approach to the simultaneous visualization problem is to use hardware-accelerated volume rendering, user-controlled data fusion, and mixed rendering styles. Interactive rendering lets users freely change rendering and visualization parameters, as well as data-fusion schemes. Rendering different volumes with different styles, if done appropriately, can enhance perception of shape, structure, and spatial relationship. The data-fusion problem here is to determine how to display multiple data values defined at the same spatial location.

A few approaches prove effective for simultaneously visualizing an S3D simulation of an H_2/CO turbulent jet flame undergoing local extinction and reignition.[3] To understand the strong coupling between turbulence, the turbulent mixing of scalars such as temperature and species concentrations, and reaction, it's helpful to render the stoichiometric mixture fraction isoline overlaid on the scalar dissipation rate (that is, local mixing intensity) image and on the hydroxyl radical concentration image. Doing so will identify the actively burning flame surface relative to these other quantities.

Data Fusion

When rendering multivariate volume data, we must decide how to treat the multiple values defined at the same spatial position. This data-fusion problem has three basic approaches:

- use one value for each color channel;
- use one of the values based on some criterion; or
- compute a weighted sum of all the values.

The first approach is probably the simplest to implement and verify. Although it's limited to visualizing three or fewer volumes, in practice, we hardly need to see more than three volumes simultaneously. The problem really limiting this approach, however, is the difficulty viewers will have interpreting the resulting color. Peter Hastreiter and Thomas Ertl give an example of the second approach, using alternating sampling for rendering two volumes,[4] which seems to work well for medical imaging but not for fluid flow visualization. The third approach gives us more freedom. We could use the opacity transfer function for one volume to enhance or deenhance some aspects of the other volume, for example. This is similar to the common practice of volume visualization in which we use gradient magnitude to enhance boundary surface. One of the authors and his colleagues have visualized multimodality volume data by computing values weighted by a function of some or all other values.[5] Furthermore, we might use the weighted sum with scaling that reflects a desired property, such as distance from the viewer.[6]

In our application, we used a simple linear function for most of the cases to blend colors from multiple variables based on user-specified weights. We can freely and interactively adjust the weights and obtain immediate visual feedback. Given that we use hardware texture volume rendering, this fusing

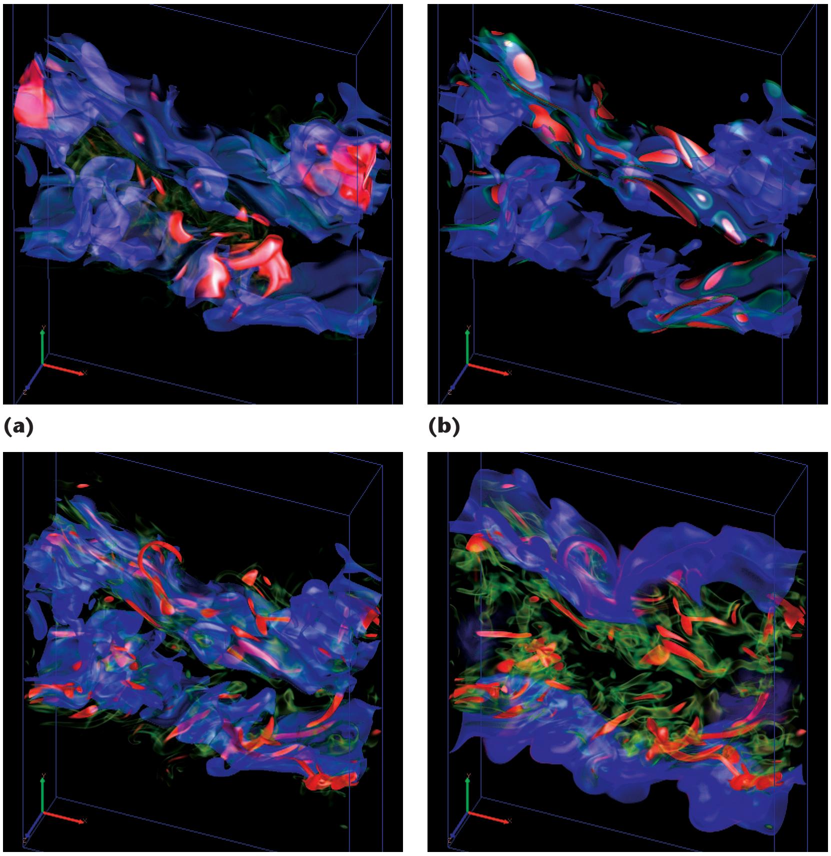

Figure 1. Simultaneous visualization. (a) The mass fraction of the hydroxyl radical, OH (red), and stoichiometric mixture fraction isosurface (blue); (b) χ (red/yellow) and a mixture fraction (blue); (c) vorticity (red/yellow) and a mixture fraction (blue); and (d) vorticity (red/yellow) and heat release (blue). All the images are for time step 41. Except for the mixture fraction isosurface using isovalue = 0.5, all others emphasize high values.

step corresponds to the determination of a fragment color based on multiple scalar values along with the associated weight and 1D transfer function for each variable. We obtain the final image by compositing the resulting fragments in the frame buffer.

Figure 1 shows four simultaneous visualization examples using this data-fusion approach. In nonpremixed combustion, in which fuel and oxidizer mix and react simultaneously, the flame is typically located where the fuel and oxidizer are in stoichiometric proportions. The mixture fraction variable represents the fraction of mass at the local position that originates in the fuel stream. It has a value of unity for pure fuel and zero for pure oxidizer. For the present mixture conditions, the stoichiometric mixture fraction is 0.42. For a fully burning flame, where chemistry timescales are fast relative to turbulence timescales, this isosurface will be a good representation of the flame, and high reaction rates, temperature, and radical concentrations will all coincide with this region. However, when strong finite-rate chemistry effects exist—in which the local turbulent mixing intensity exceeds the governing chemical rates—local extinction can occur. Some areas of the stoichiometric surface might be only weakly burning or completely extinguished, characterized by lower radical levels, temperature, and heat-release rate. Later, these regions might reignite independently or with assistance from heat conduction and radical diffusion from neighboring fully burning flame elements. Hence, the mixture fraction alone is insufficient for characterizing the local combustion, and simultaneous visualization of multiple, time-varying scalars is required.

Figure 1a shows the mass fraction of the hydroxyl radical (OH) in red and the mixture fraction 0.42 isosurface. We can see that due to the effects of turbulence, the OH field is far from uniform on the mixture fraction isosurface, which we would expect from fast chemistry limits. High OH regions are found where the mixture fraction isosurface is highly convoluted, and low ones are in regions where the mixture fraction isosurface is stretched or unwrinkled. We can explain these observations by looking at Figure 1b, which shows the simultaneous visualization of the mixture fraction 0.42 isosurface and the scalar dissipation χ. The scalar dissipation, which is closely related to the strain rate, is essentially the rate of mixing between fuel and oxidizer, and is a very important quantity in nearly all models of nonpremixed combustion. Some mixing is desirable for efficient combustion, but very large mixing rates lead to heat and radical losses that are so great that the chemical reaction is no longer self sustaining. Figure 1b shows that the high χ regions occur in low-volume, pancake-like structures and are aligned according to principal strain-rate directions of the shear flow. Comparing Figures 1a and 1b shows that high χ and low OH regions coincide. Figure 1c shows vorticity magnitude and a mixture fraction, whereas Figure 1d shows vorticity magnitude and the heat-release rate.

These images give insight into how the turbulent flow field interacts with the flame. In this configuration, vortex stretching induced by the mean shear results in a high density of fine-scale eddies that cause significant mixing.

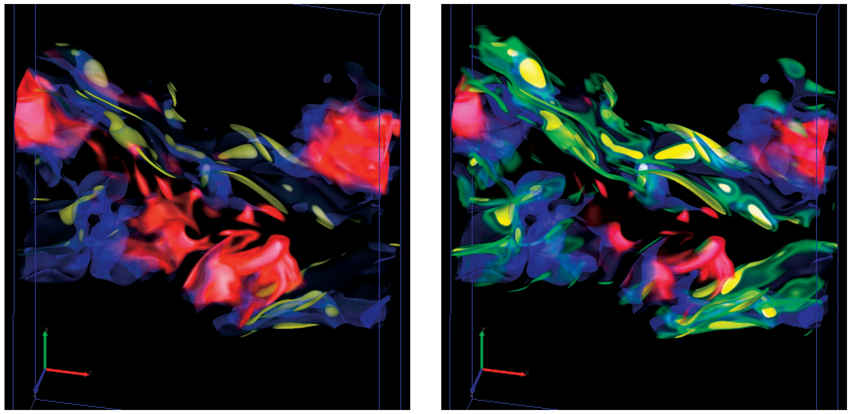

(a) (b)

Figure 2. Simultaneous visualization. With three variables—χ, OH, and a mixture fraction—one visualization emphasizes (a) χ, whereas the other emphasizes (b) OH.

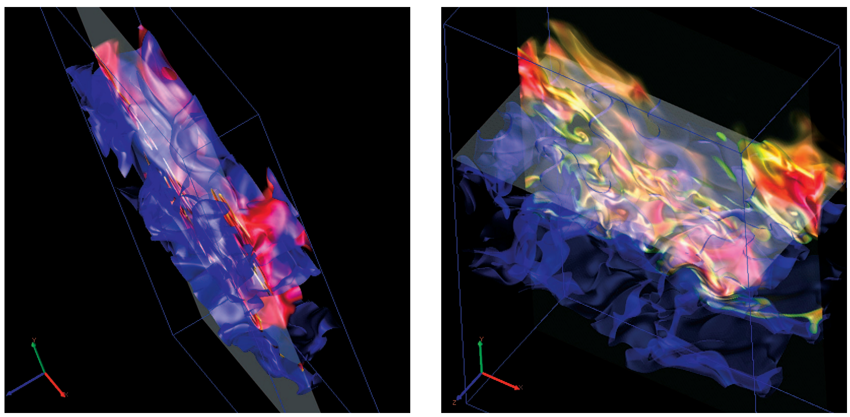

(a) (b)

Figure 3. Cut-away. For (a) interactive univariate cutting, we applied the vertical cutting plane to OH. For (b) interactive multivariate cutting, we applied the vertical cutting plane to OH and the other plane to a mixture fraction.

Interactive Fusing and Cutting

When we use data fusion, it's frequently difficult to determine the spatial relationship between two or more variables. Our system lets users interactively vary one variable's weight (or transparency). In addition to varying the view and the cut-away capability we discuss next, users can also remove most of the ambiguity. In Figure 2, two visualizations show the result of varying blending weight. We had to compare Figures 1a and 1b to understand the combustion dynamics, but the three-variable visualization in Figure 2 highlights much more readily that high OH and χ regions are exclusive due to rapid diffusive losses in the high χ regions.

A common problem in 3D visualization is occlusion, which is more severe when the visualization consists of rendering two or more scalar quantities. Cut-away (or slicing) is a widely employed method for visualizing 3D data to remove occlusion (or see inner structures). Our system supports cut-away in a variable-dependent fashion—that is, users can interactively cut away one variable while keeping the other's integrity. Users can also specify multiple cutting planes to reveal a particular region of interest. In our system, cutting is implemented with the fragment program: given a variable to be cut and cutting-plane equations, the variable's opacity is set to zero if the current fragment position was inside the cut-away region the plane equations specified. In this way, the polygon slices and the volume data stay intact. This fragment-based method is easy to implement and can achieve interactive frame rates during cutting.

Figure 3a shows visualizations using univariate cut-away. Compared with Figure 1a, which obscures them, the cut-away lets scientists observe low OH regions, characteristic of weakly burning or extinguished flames. In particular, occlusion is a problem at the edges of low OH regions, which are often obscured by high OH regions. These regions are of particular interest in the reignition scenario when diffusion from the strongly burning regions (high OH) causes reignition near these edges. Figure 3b shows multivariate cut-away, which we can use to further relieve the occlusion problem.

Illustration and Enhancement

To disambiguate the complex relationships between multiple variables in the 3D and 4D domains, we should also employ some nonphotorealistic rendering (NPR) methods, which involve applying techniques used by artists for computer-generated imagery (CGI). Scientific visualizations created with

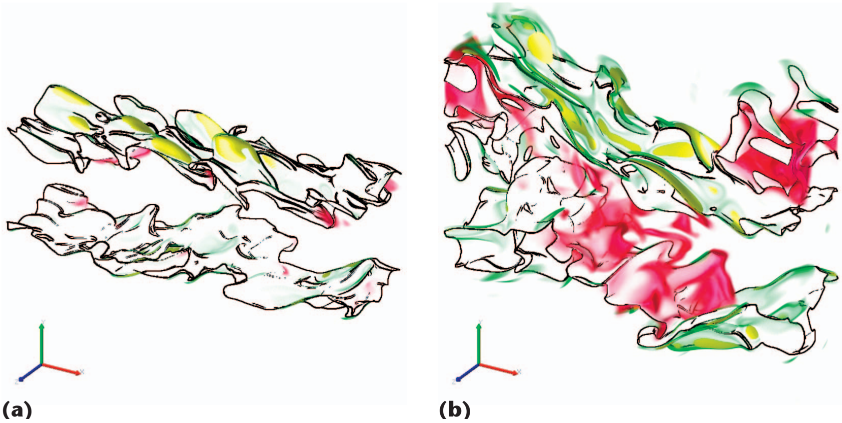

Figure 4. Nonphotorealistic rendering of three variables. At (a) time step 11, and (b) time step 41, we see different views of OH (red), χ (yellow and green), and a mixture fraction isosurface.

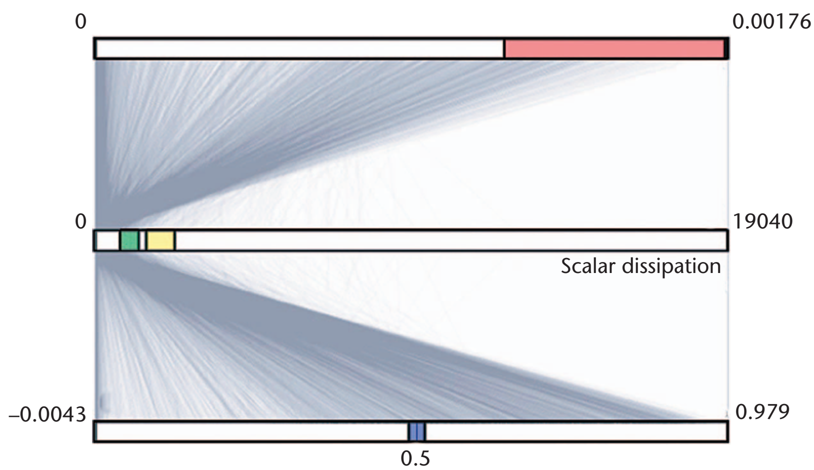

Figure 5. A parallel coordinate interface for making a multivariate data visualization. The polylines show the relationship between χ, OH, and a mixture fraction.

these techniques can be more meaningful than those generated with more traditional photorealistic ones. NPR is usually associated with surface representation, but it has also been applied to direct volume rendering.[6,7] By applying several NPR techniques, we can accentuate key features in a volume while de-emphasizing structures that might obstruct or detract from those features. Figure 4 illustrates the mixture fraction isosurface using silhouette rendering displayed simultaneously with volume rendering of OH and χ. Silhouette rendering consists of adding dark lines around an object and can be very effective in enhancing fine structures. It can also aid in depth perception when viewing overlapping structures of similar color. The spatial relation between these two variables is much clearer in this image.

Our system lets users color and light each variable independently, which helps distinguish them. However, these variables often share spatial locations. Thus, highlighting those locations can be helpful—for example, users can enhance the variables with color and use NPR to improve clarity further. These images are particularly useful for studying extinction and reignition dynamics, in which the focus is on areas of the stoichiometric surface that have low OH values. Using the enhancement technique, we can visually identify these regions without confusion regarding the two volume fields' relative spatial locations. After identifying the low OH regions, we can observe their temporal evolution to understand the reignition dynamics as these regions interact with burning regions.

Visualization Interfaces

The ability to simultaneously visualize multiple variables is very powerful. However, the relationship between these variables is often nontrivial, requiring repeated exploration in the data's temporal, spatial, and variable domains. An easy-to-use, interactive interface for specifying different combinations of multivariate data visualization is therefore desirable. We've developed an interface based on parallel coordinates[8] and time histograms[9] that effectively supports time-varying multivariate feature extraction.

Figure 5 shows the parallel coordinate interface, in which selected points in the three-variable space are represented as polylines with vertices on the parallel line segments. The parallel-coordinate interface can thus show the relationships between two or more variables or time steps. Figure 6 shows the time-histogram interface, which displays each variable's temporal characteristic and helps users identify time steps of interest. Together with volume, surface, and cutting-plane visualizations, these different views of the data are tightly coupled to facilitate trispace data exploration.

Figure 7 presents a simultaneous visualization example for selected time steps specified over the parallel coordinate interface in Figure 5. The red region in the images covers the high OH values, whereas the green and yellow regions cover low χ values. The parallel-coordinate display also suggests such a combination.

Detecting and tracking time-varying flow features can be very tedious, especially when we can't define the feature of interest analytically. The new visualization interface lets us define a feature by linking a set of transfer functions for rendering multiple variables. Note that as we manipulate in the parallel-coordinate space as well as the transfer-function space, interactive visual feedback is critical so that we can refine the feature accordingly. One feature of interest, for example, is isolating local extinction followed by reignition. The surface identified as the flame corresponds to a mixture fraction isolevel with a value of 0.42. Just because a region has a dissipation rate higher than the steady laminar critical value doesn't necessarily mean it will extinguish in a turbulent environment. However, it would be useful to track those regions together, say, with the temperature or OH fields at a prescribed threshold value (the threshold corresponds to one half of an OH concentration at the steady state extinction limit of a laminar strained diffusion flame). Visually, extinction should appear as a hole in the OH, O, H, or temperature fields. We want to observe the transition from burning to extinguished regions, as well as track the edge of the hole—to see if it's growing or retreating locally and determine the speed at which it's moving.

Our interface design lets us easily select quantities from a large parameter space, whereas the interface gives us an overview of the selected variables as well as the relationships among them. To locate the feature of interest described earlier, we tried to define it by linking the relevant variables through the parallel-coordinate interface, as Figure 8 shows, and to verify it over the temporal domain by looking at the time histogram. Figure 9 shows the visualization results we got for four selected time steps. As anticipated, holes (in red) appear on the surfaces. With this interactive interface, we can specify the relationships that we think we understand and immediately observe if the simulation supports our hypothesis. The interface thus not only lets us freely explore our simulation data but also gives us an easy way to validate our simulations.

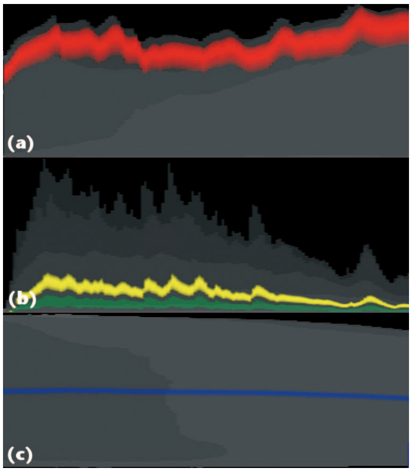

Figure 6. Time-histogram interface showing data distributions and trends in the temporal domain. The *x*-axis is time and the *y*-axis is data value. The shade of each pixel indicates the number of occurrences for the corresponding value at the corresponding time step: (a) OH, (b) χ, and (c) a mixture fraction. The colored regions correspond to the visible flow features in the visualization.

Discussion

We and others have made a lot of progress in advancing visualization technologies, but it can take quite a long time before such technologies are available to scientists. Although there is a perceptual limit on how many variables users can view simultaneously, our current system serves as a testbed we can use to study other approaches. We believe our multivariate exploration capability, coupled with adaptive feature extraction and tracking methods, will let us glean considerable insight from our simulation data. A typical time step in a 3D production run generates at least 25 Gbytes of data and represents roughly 2 to 4 nanoseconds of physical combustion. Ideally, we'd want to save data generated from all time steps to provide accurate analysis in real time or postprocessing, particularly for transient, intermittent events such as ignition and extinction that are impossible to determine a priori. Unfortunately, it takes both time and space to concurrently store data from a large number of processors on a supercomputer. As such, production runs today are constrained to save data at predetermined intervals, causing some adverse effects, including a potential loss of critical data (such as extrema, sharp quantitative changes, and so on), more complicated data analysis, and difficulties in real-time code steering. We must thus develop a new paradigm for extracting salient knowledge from terabyte data sets generated on supercomputers. One approach couples parallel feature detection and visualization to provide a unified solution for knowledge extraction. It's also promising to design the feature-detection process based on topological and level-set methods that we can apply to subset the combustion data, which exhibit large dynamic ranges.

Additionally, our next modeling phase will track Lagrangian fluid particles and flame element positions in space and time. Tracking these elements will help us understand the cu-

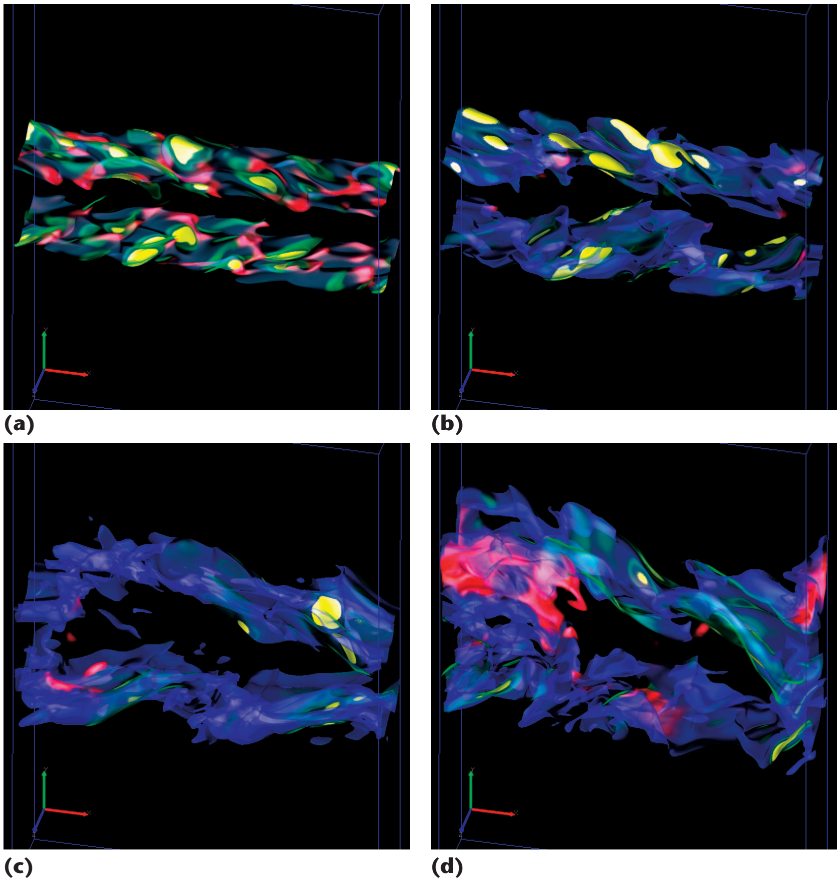

Figure 7. Simultaneous visualization of three variables at selected time steps: (a) 4, (b) 11, (c) 21, and (d) 31.

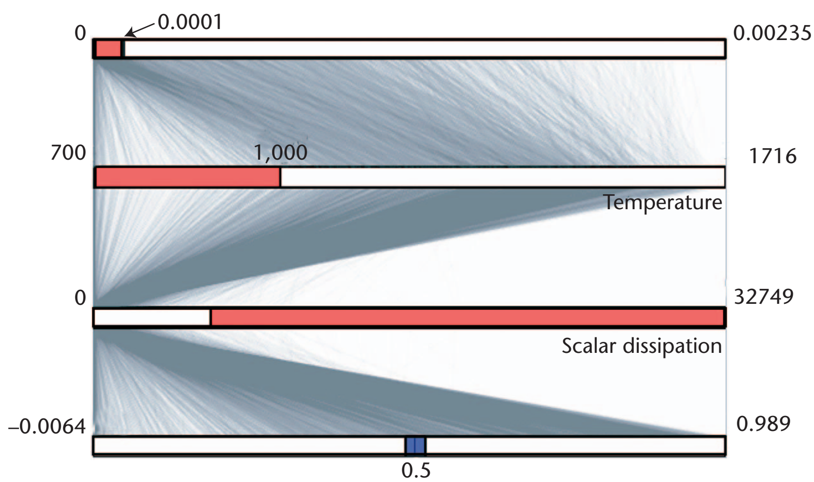

Figure 8. The multivariate data visualization interface. We use this interface to extract a feature involving four variables: OH, temperature, χ, and a mixture fraction. Figure 9 shows the resulting visualizations.

mulative unsteadiness effects that a portion of the flame encounters as the turbulence modulates the flame topology and the local burning intensity. We will then want to visualize the evolution of derived quantities interpolated onto the particle and element locations, and thus better understand dynamic localized events such as extinction and reignition. For this, we'll be able to apply the particle tracking and visualization technique we previously developed for accelerator simulations.[10]

To design cleaner, more efficient, and safer combustion devices, we need a greatly improved understanding of basic combustion processes that we can use directly during design or incorporate into engineering-level computational fluid dynamics models. High-fidelity simulation approaches can provide such information, but there's currently a bottleneck in the translation of vast quantities of raw data representing complex turbulence-chemistry interactions into useful knowledge. These interactions aren't fully understood precisely because they're so complex, thus multiscale and new science-driven visualization techniques are required to resolve this complexity. We foresee that visualization's major roles in the DNS of combustion will be to provide qualitative understanding of the underlying combustion phenomena and subsequently guide combustion scientists toward regions requiring more quantitative analysis.

These new interactive, multivariable visualization capabilities will provide combustion scientists with never-before-seen views of the underlying phenomena and are destined to lead to new understanding. The ability to simultaneously visualize multiple variables is essential because combustion is inherently a multiscalar problem with intricate coupling between reactive scalars and the flow field. The techniques we described here are already proving useful in guiding qualitative understanding, and we expect they will also prove very useful in guiding further quantitative analysis. Without such techniques, it would be nearly impossible to extract salient knowledge from combustion data of this magnitude.

Acknowledgments
The work is sponsored in part by the US Department of Energy's Division of Chemical Sciences, Geosciences, and Biosciences, the Office of Basic Energy Sciences, the SciDAC program, and the US National Science Foundation's ITR program.

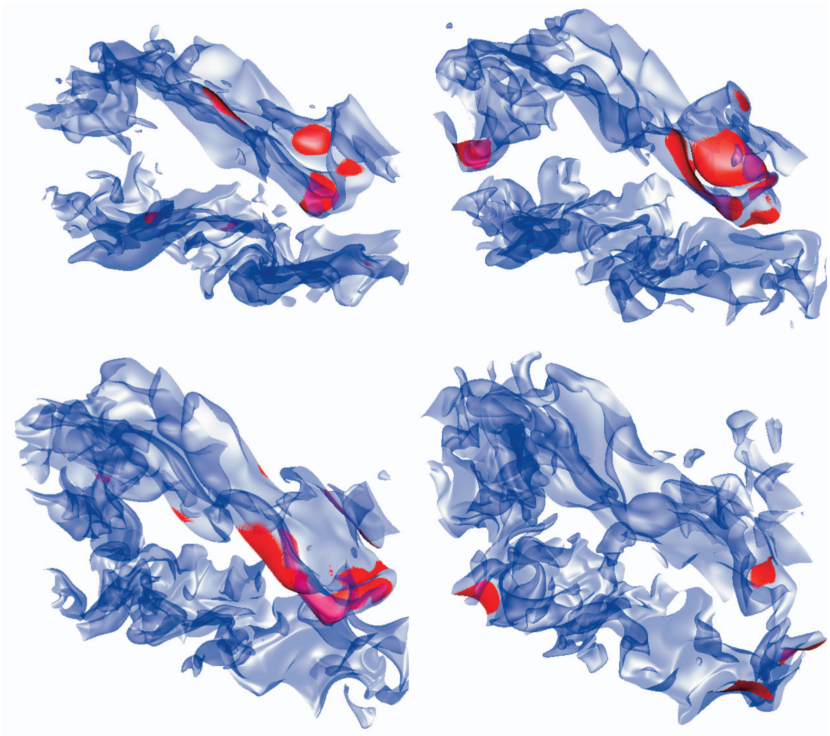

Figure 9. Feature extraction using the new multivariate data visualization capability. The interface shows the feature of interest, extinction, as holes on the isosurface of a mixture fraction. From top to bottom, left to right, the time steps are 21, 25, 29, and 35.

References

1. J.H. Chen et al., *Sandia Internal Report*, tech. report, Sandia Nat'l Laboratories, 2005.
2. J. Ahrens et al., "Quantitative and Comparative Visualization Applied to Cosmological Simulations," *J. Physics: Conference Series*, vol. 46, 2006, pp. 526–534.
3. E.R. Hawkes et al., "Scalar Mixing in Direct Numerical Simulations of Temporally-Evolving Plane Jet Flames with Detailed CO/H2 Kinetics," *Proc. Combustion Institute*, Elsevier, 2006; doi:10.1016/j.proci.2006.08.079.
4. P. Hastreiter and T. Ertl, "Integrated Registration and Visualization of Medical Image Data," *Proc. Computer Graphics Int'l*, IEEE CS Press, 1998, pp. 78–85.
5. B. Wilson, E. Lum, and K-L. Ma, "Interactive Multi-Volume Visualization," *Proc. Int'l Workshop Computer Graphics and Geometric Modeling*, Springer-Verlag, 2002, pp. 102–110.
6. E.B. Lum and K-L. Ma, "Hardware-Accelerated Parallel Non-Photorealistic Volume Rendering," *Proc. Int'l Symp. Non-Photorealistic Animation and Rendering* (NPAR 02), ACM Press, 2002, pp. 67–74.
7. P. Rheingans and D. Ebert, "Volume Illustration: Nonphotorealistic Rendering of Volume Models," *IEEE Trans. Visualization and Computer Graphics*, vol. 7, no. 3, 2001, pp. 253–264.
8. A. Inselberg, "The Plane with Parallel Coordinates," *The Visual Comput*er, vol. 1, no. 2, 1985, pp. 69–91.
9. H. Akiba, N. Fout, and K.-L. Ma, "Simultaneous Classification of Time-Varying Volume Data Based on the Time Histogram," *Proc. Eurographics Visualization Symp.*, Eurographics, 2006, 171–178.
10. K-L. Ma et al., "Advanced Visualization Technology for Terascale Particle Accelerator Simulations," *Proc. IEEE/ACM SC06 Conf.*, IEEE CS Press, 2002; www.sc-2002.org/paperpdfs/pap.pap224.pdf.

Hiroshi Akiba is a PhD candidate in computer science at the University of California, Davis. His research focuses on time-varying data visualization. Contact him at akiba@cs.ucdavis.edu.

Kwan-Liu Ma is a professor of computer science at the University of California, Davis and the director of the US Department of Energy's SciDAC Institute for Ultra Scale Visualization. His research interests include visualization, computer graphics, and high-performance computing with particular efforts aimed at advancing the state of the art in data visualization technology to address the challenges presented by data that grows in size and complexity. Ma has a PhD in computer science from the University of Utah. Contact him at ma@cs.ucdavis.edu.

Jacqueline H. Chen is a Distinguished Member of Technical Staff at Sandia National Laboratories and a principal investigator in the DOE Basic Energy Sciences Chemical Sciences Program. Her research interests and expertise are in terascale direct numerical simulations of turbulent combustion focusing on fundamental understanding of turbulence-chemistry interactions and validation of models. Chen has a PhD in mechanical engineering from Stanford University. Contact her at jhchen@sandia.gov.

Evatt R. Hawkes is a postdoctoral combustion researcher at Sandia National Laboratories. His research spans several high-fidelity simulation approaches, including direct numerical simulation, large-eddy simulation, and the development and testing of turbulent combustion models. Hawkes has a PhD in engineering sciences from Cambridge University. Contact him at erhawke@sandia.gov.