ANALYZING INDUSTRIAL FURNACE EFFICIENCY USING COMPARATIVE VISUALIZATION IN A VIRTUAL REALITY ENVIRONMENT

Lori Freitag
Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL 60439

Tim Urness
Scientific Visualization Laboratory
1003 Cowley Hall
University of Wisconsin - La Crosse
La Crosse, WI 54601

Address Queries to:

Lori Freitag
MCS 221/C223
Argonne National Laboratory
Argonne, IL 60439
(630) 252-7246
freitag@mcs.anl.gov

ABSTRACT

We describe an interactive toolkit used to perform comparative analysis of two or more data sets arising from numerical simulations. Several techniques have been incorporated into this toolkit, including (1) successive visualization of individual data sets, (2) data comparison techniques such as computation and visualization of the differences between data sets, and (3) image comparison methods such as scalar field height profiles plotted in a common coordinate system. We describe each technique in detail and show example usage in an industrial application aimed at designing an efficient, low-NOx burner for industrial furnaces. Critical insights are obtained by interactively adjusted color maps, data culling, and data manipulation. New paradigms for scaling small values in the data comparison technique are described. The display device used for this application was the CAVE virtual reality theater, and we describe the user interface to the visualization toolkit and the benefits of immersive 3D visualization for comparative analysis.

1   INTRODUCTION

Designing a new manufacturing system often involves using numerical simulations to analyze different options before actual prototypes are built. Critical insights can be obtained by examining both the results of a particular simulation and the differences among several solution sets. For example, when designing a system prototype, an engineer may wish to examine the effects of varying an input parameter on the final solution, isolate the regions and physical quantities most affected by the change, and visualize and analyze multiple solution data sets simultaneously. Unfortunately, in most cases, visualization tools are limited to importing and displaying a single data set; multiple data sets must be loaded in successive order and displayed individually. To compare simulation results in this paradigm, the user must either remember the characteristics of interest from one data set to the next or initiate multiple visualization sessions and display the windows side by side. In both cases it is difficult, if not impossible, to isolate the differences among the data sets, particularly if the differences are significantly smaller than the physical features of the numerical simulation.

To address this problem, researchers have recently begun developing comparative visualization techniques for analyzing multiple data sets simultaneously. These methods fall into two primary categories: image comparison and data comparison. Image comparisons are performed either by displaying data set images side by side in the same visualization coordinate system, or, more directly, by computing the difference in the *image* produced by the visualization tools [5, 6]. Shortcomings of this technique include the fact that subtle differences are difficult to discern, particularly in side by side comparisons. In addition, different visualization techniques operating on the same data set can produce spurious differences unrelated to the underlying data set [5]. The second approach to comparative visualization, data comparison, is performed by creating an intermediate data set by combining information from two or more data sources and visualizing the resulting "differenced" data set. This technique has been used successfully to compare simulation and experimental data [8, 6, 7] and also to compare the data produced by different visualization algorithms, such as volume rendering techniques [5].

In this paper, we describe a toolkit that incorporates a combination of comparative visualization techniques for analyzing multiple data sets arising from numerical simulations. This toolkit, called the ALICE Differencing Engine (ADE), was motivated by a joint project between the U.S. Department of Energy (DOE) and Air Products and Chemicals, Inc. to analyze fuel efficiency for a new burner targeted for use in industrial furnaces. We describe the application in some detail in Section 2. The techniques developed for comparative visualization and analysis, including successive

individual data set visualization, data comparison techniques, and image comparison techniques, are described in Section 3. As each technique is described, its use is illustrated in the context of the aluminum smelting furnace application. The display device used with this toolkit is the CAVE immersive virtual reality system [3], and we briefly describe the interactive environment and user interface in Section 4. Finally, in Section 5 we offer concluding remarks and directions for future development of the ADE toolkit.

2   A MOTIVATING APPLICATION: ALUMINUM SMELTING FURNACE EFFICIENCY

Recently, the U.S. Department of Energy and Air Products and Chemicals, Inc. entered into a joint project to design, build, and test a new burner nozzle for industrial furnaces. As a test case, an aluminum smelting furnace (shown in Figure 1) was chosen for initial evaluation of the nozzle. In typical furnace units, an operator continuously supplies aluminum source to an external charging bin, where it is mixed with already molten aluminum. A pump in an adjacent external tank causes the molten aluminum to circulate through the body of the furnace where combustion occurs and the molten aluminum temperature is maintained. Over the course of several hours, the volume of molten aluminum increases to a critical level, the ratio of metals in the alloy is checked for correctness, and the furnace is tapped to form the alloy ingot. Better production rates of aluminum ingot can be achieved with higher combustion temperatures; in particular, a fuel consisting of pure O_2 will burn much hotter than an air fuel, but at the cost of higher operating expenses and higher NO_x emissions.

Figure 1: An aluminum smelting furnace

The overall goal of the joint DOE/Air Products project is to find a nozzle and fuel combination that achieves high efficiency and low NO_x emissions at low cost to the furnace operator. Complex numerical simulations that coupled the combustion process in the furnace to the molten aluminum were performed for three different fuel choices: air, pure O_2, and air enriched with O_2. To determine which fuel is best, a number of factors must be examined, including the initial cost of the fuel, the level of NO_x emissions (which corresponds directly to the amount of O_2) the amount and kind of contaminates in the molten aluminum alloy caused by the fuel, and the efficiency of the furnace. Because the new nozzle and fuel combination will be installed directly into a production furnace

(no prototype furnace environments are available), the process engineers rely on the numerical simulations to quantify these tradeoffs before installation, to minimize the impact on production.

The numerical simulations were performed with the commercial computational fluid dynamics package FLUENT [2]. The computational mesh used for each simulation consisted of a static, logically regular grid containing roughly 90,000 cells. In this case the mesh used for each simulation was identical; only the boundary conditions were changed to model the three different fuel types. Numerous vector and scalar field values were computed at each cell center, but in this case, the following simulation data of most interest:

- the velocity fields and temperature distributions in both the combustion gases and molten aluminum, in particular, the mixing and uniformity of the temperature distributions of combustion gases and the circulation characteristics and stagnation points in the molten aluminum; and

- the distribution of O_2 throughout the computational domain, because it is a prime indicator of the level of NO_x emissions produced during operation of the unit.

Direct comparisons among the data sets that focused on the differences in temperature distributions, O_2 levels in furnace, and mixing characteristics of each fuel type would enhance the ability of process engineers to quickly determine the best fuel nozzle combination to use. Unfortunately, such direct comparisons are not possible using the FLUENT visualization environment. Instead, separate, side by side comparisons of the data sets were used to obtain an approximate measure of the differences between data sets.

3 VISUALIZATION TOOLKIT FOR COMPARATIVE ANALYSIS

To allow the engineers to quickly and intuitively make these comparisons, we developed the ALICE Differencing Engine (ADE), a toolkit that can be used in an immersive virtual reality environment for comparative analysis. As shown in Figure 2, the toolkit provides three primary functionalities: the visualization of individual data sets, data comparison via the visualization of the differences between data sets, and image comparison via the simultaneous visualization of multiple data sets. For each functionality, one or more standard visualization techniques, such as vector field glyphs, animated streamlines and flow fields, cutting planes, and scalar field height profiles, are used to gain insight into the data. Additional insights are gained through the dynamic selection of the color maps and through data manipulation techniques such as magnification, culling, and exaggeration. We now describe the visualization techniques, color mappings, and manipulation strategies for each of the three functionalities.

3.1 Visualization of Individual Data Sets

In Figure 3 we show the options available to the user to analyze an individual data set. Standard visualization techniques such as tetrahedral dart vector glyphs are used to represent vector and scalar fields simultaneously; their direction and length correspond to the direction and magnitude of the vector field at that data point. The color, C, assigned to each vector glyph is determined by interpolating the field value of interest, s, over some range, $[m_r, M_r]$, which can be dynamically selected by the user. That is,

$$C = \frac{M_r - s}{M_r - m_r}.$$

Data Manipulation: Culling, Magnification, Focus, Projection
Color Map Options: Vector, Scalar, Source, Data Subsets
Visualization Techniques: Vector Glyphs, Cutting Planes, Streamlines

Individual Data Sets	Data Comparison	Image Comparison

Raw Data: Mesh   Data Set 1   Data Set 2   ....   Data Set 10

Figure 2: Overview of the ALICE Differencing Engine (ADE) toolkit for comparative analysis of multiple data sets

If $s = m_r$, the color assigned to the vector glyph is blue; if $s = M_r$, the color assigned to the vector glyph is red. The extrema values, m_r and M_r, can be chosen in a number of different ways: by the individual data set extrema, m_d and M_d; by the global extrema over all the data sets, m_g and M_g, or by user-defined bounds, m_u and M_u. In addition, the user may manipulate the vector field glyphs by culling vectors whose corresponding scalar value falls outside the range $[m_u, M_u]$ or by uniformly magnifying them for easier viewing.

Cull		Magnify	
Scalar	Data Set	Global	User-Defined

Vector Glyphs	Streamlines	Cutting Plane

Individual Data Set Analysis (Data Set I or J)

Figure 3: The options available in the ADE toolkit for analyzing individual data sets

In Figure 4 we show the use of the vector field glyphs for the three different fuel types under consideration: air, an air/O_2 mix, and pure O_2 (left to right, respectively). The color map is determined by the global range $[m_g, M_g]$ for temperature. Blue indicates the coldest temperature among the three data sets, approximately 273° F in the molten aluminum in the air data set, and red indicates the hottest temperature, approximately 3050° F in the combustion space of the O_2 data set. The four areas of high-speed flows are in the vicinity of the furnace burners. This series of images clearly shows that air fuel yields the coldest temperatures in the combustion space, averaging only about $1600\,^\circ$F, and therefore yields the poorest furnace efficiency. The temperatures corresponding to the O_2 fuel are the hottest, with the air/O_2 mix falling in between.

In addition to the static vector field glyphs, the user may choose to visualize the fluid flow in individual data sets by using animated streamlines, dynamic flow field movies, or interactively defined cutting planes. The streamlines and flow field movies may be colored by the scalar field of choice, but the range is constrained to $[m_d, M_d]$. The interactively defined cutting planes have the most flexible color maps, and the user may choose the ranges corresponding to any of the bounds defined above, $\left[m_{\{d,g,u\}},\ M_{\{d,g,u\}}\right]$, or the cutting plane extrema, $[m_c, M_c]$. In addition, if the range $[m_u, M_u]$ is used, data culling in the cutting plane also occurs; that is, data that falls outside the bounds are not drawn and appear as holes in the cutting plane.
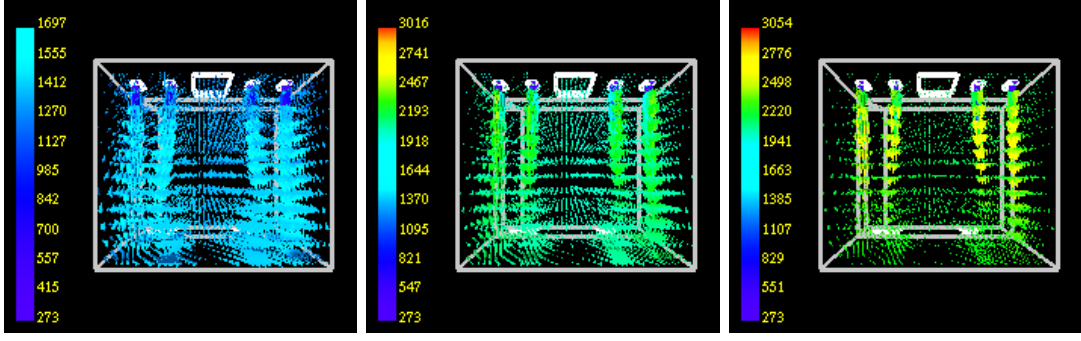
Figure 4: This series of images shows the flue gas flows in an aluminum furnace for three different types of fuel: air, air/O$_2$ mix, and pure O$_2$. Each tetrahedral dart's magnitude and direction give the magnitude and direction of the flow at that point. The color is mapped to temperature, where blue indicates the coldest temperature among the three data sets and red indicates the hottest temperature.

3.2 Data Comparison Techniques

To perform comparative analysis using the differences between data sets, the user interactively selects the two data sets to be compared, for example, data sets i and j. We currently assume that the computational grid is identical for all the data sets, thereby eliminating concern regarding the error associated with data interpolation to common coordinate locations for the purposes of comparison. This is a potential problem when the interpolation error is the same order of magnitude as the difference between the data sets. We compute pointwise differences between the two data sets as follows: pointwise scalar field differences are computed directly as

$$S_{diff} = |S_i - S_j|,$$

and pointwise differences in vector fields are computed by

$$\mathbf{u}_{diff} = [u_i - u_j, \ v_i - v_j, \ w_i - w_j].$$

Figure 5: The options available in the ADE toolkit for analyzing data differences

The options available to the user for visualizing the resulting "differenced" data set is summarized in Figure 5. The primary visualization technique used for this option is the vector field glyphs described in the preceding subsection. In this case the color map corresponds to the range of values in S_{diff} or to the range of magnitudes of $\mathbf{u}_{diff}$. In addition, the user may interactively select a subset range of values in S_{diff} to both determine the color map and cull the vector glyphs.

The differenced data set is more likely that the original data to contain small data values that are difficult to discern. To obtain additional insight, the user may manipulate the vector field glyphs in a variety of ways. The vectors may be culled to focus on regions of interest, for example, the maximum or minimum values in S_{diff}, and/or magnified uniformly for easier viewing. In addition, the user may exaggerate the differences in the scalar field above a user-defined threshold by using the formula

$$Scale\ Factor = \left(\frac{s}{f}\right)^{\alpha} \times M,$$

where s is the scalar value at the current data point, f is the threshold value, α ranges from 0 to 5, and M is the magnification constant. The resulting $Scale\ Factor$ value is multiplied by the vector glyph magnitude to determine the final size. As exp increases, vectors whose corresponding scalar values s are greater than f grow exponentially, whereas those with scalar values less than f decrease, resulting in an exaggeration of the differences above the threshold value. If $\alpha = 0$, changes in f result in no changes to the vector glyph appearance.
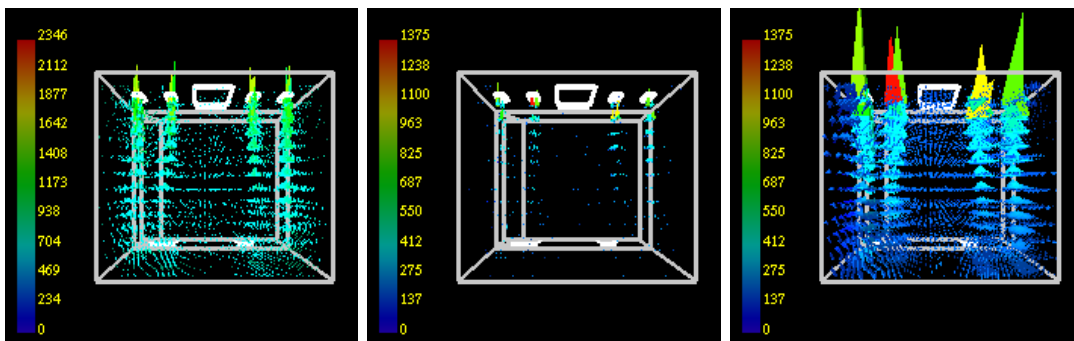

Figure 6: Example use of the data difference technique in the aluminum smelting furnace. The first two images from the left show the temperature difference between the O_2 and air fuels and the air/O_2 and O_2 fuels. The last image shows a magnified view of the second image.

Examples of the data comparison technique applied to the aluminum smelting furnace data are shown in Figure 6. The first tow images in this series from left to right show the difference between the pure O_2 and air cases and the pure O_2 and air/O_2 cases. Each point's color gives the difference in temperature at that point; blue corresponds to the smallest difference and red to the largest difference. In this case, the maximum differences are seen between the pure O_2 and air cases, roughly 1500° F difference and the minimum differences are seen between the pure O_2 and air/O_2 mix, roughly 180° F difference. Because the colors range from blue to red for each case, quantitative feedback regarding the values of S_{diff} must be provided in the form of a colorbar as illustrated in each image. Because the differences between O_2 and air/O_2 can be difficult to see, the third image shows a magnified view of the differences between them.

3.3   Image Comparison Techniques

Comparative analysis using image techniques can be as simple as plotting simulation data from different experiments in the same one- or two–dimensional graphs and coordinate spaces. We provide the analogous mechanism for this type of analysis for scalar fields in three-dimensional data sets. An overview of the image comparison technique provided in the ADE toolkit is shown in Figure 7.

Figure 7: The options available in the ADE toolkit for image differencing

Cull			Computational Grid		
Source	Scalar	Plane	Data Set	Global	User
Scalar Field Height Profiles					
Image Comparisons (Data Set I and J)					

The user may select any number of data sets and display the scalar fields from an interactively defined cutting plane using height profiles. Each data set is identified by a uniquely colored outline that corresponds to the control panel button color. Additional insights in this case are obtained through manipulation of the color map. The scalar field color maps can be bounded by any of the ranges discussed above, $\left[m_{\{u,d,g,c\}}, M_{\{u,d,g,c\}} \right]$. In addition, the differences in the heights of the cutting plane are more clearly delineated if the user selects the option to color by data source. In this case each data set is assigned a unique color, and subtle differences between the height profiles are easier to discern.

In Figure 8 we show the use of the image comparison technique for the aluminum smelting furnace data. The cutting plane from which the data is displayed is shown by the flat, gray surface; the height of the data surfaces above the cutting plane corresponds to the temperature at that point. The colored markers and outline on the boundary of each data set indicate the source of the data: red markers indicate the air case, green markers indicate the air/O_2 mix, and blue markers indicate the pure O_2 case. In the three images we show different options for displaying the height profiles: in the left image, the color of the data surface corresponds to temperature; in the middle image, the color of the surface corresponds to data source and the height field is displayed with the computational grid; the rightmost view shows the advantage of immersive display to gain additional insight from within the data set.
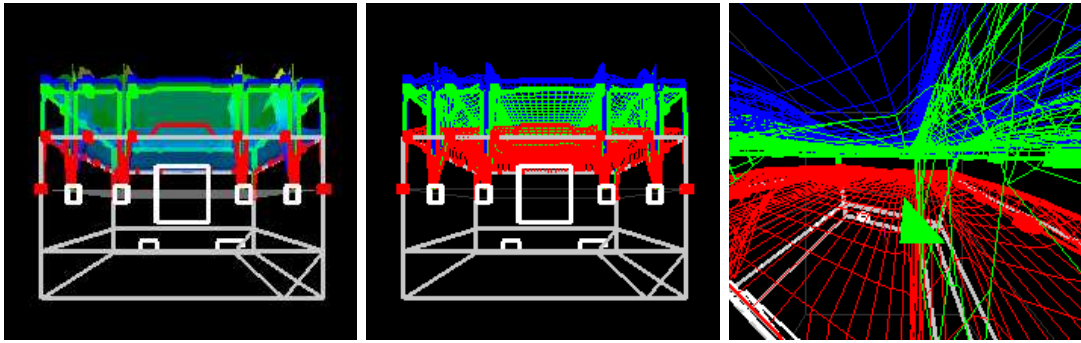

Figure 8: Data displayed from a cross-sectional cutting plane in the aluminum furnace. The location of the cutting plane is given by the gray surface. Image comparison is accomplished by plotting height profiles of scalar fields of interest, in this case temperature. The colored markers outlining each data set indicate the source of the data: red markers indicate the air case, green markers indicate the air/O_2 mix, and blue markers indicate the O_2 case. The height profiles are colored by temperature in the leftmost image and by source in the middle and rightmost image.

4   THE INTERACTIVE VIRTUAL ENVIRONMENT

The display device used with this application is based on the CAVE technology developed at the Electronic Visualization Laboratory at the University of Illinois at Chicago [3]. Users are immersed in the virtual environment by stepping into a ten-foot cube that has stereo images projected onto three walls and the floor. Several users may be immersed simultaneously in the same virtual environment and interact with the same computational model. One user is tracked by an electromagnetic tracking system, and the image orientation is calculated with respect to the head position of that user. Objects in the CAVE are manipulated by the user by means of a hand-held wand, a three-dimensional analogue of the mouse on current computer workstations.

The immersive, 3D display provided by the CAVE allowed the user to easily navigate through the various data sets and examine spatial regions of interest from several different vantage points. This was particularly useful when performing the image comparision; navigating through the scalar field height profiles provided more insight into the differences in the data than viewing them externally was able to provide (see Figure 8.

In working with engineers from industry, particularly those who are not familiar with computational techniques, we have found that it is critical to provide a realistic frame of reference for the visualization and comparison of simulation data. Toward this end, we have texture-mapped digital images of an existing aluminum smelting furnace to the boundaries of the computational domain, which in this case is a large rectangular box. In addition, we have augmented the computational domain with structures such as charging bins, exhaust stacks, and burners, which are not included in the numerical simulation but provide important visual clues for orientation (see Figure 1).

In addition, to create a natural interaction paradigm with the various toolkit options, we are using a "control panel" GUI interface developed at Argonne National Laboratory [4]. This GUI interface is written in openInventor and C++ and provides a simple API for creating toggle buttons, single and multi-selectable button rows, and sliders. Multiple control panels can be created and are selected by clicking on a tab identifier.

We have used this infrastructure to create a number of separate control panels for interacting with the ADE toolkit. In particular, each data set has its own control panel for individual data set analysis. Separate control panels exist for the data and image comparison techniques. Using the data comparison control panel, the user can interactively request two data sets and compute the differences between them. Sliders exist to magnify, cull, or exaggerate the differences as the user desires. The user may interactively choose the color map to correspond to different scalar fields. Using the image comparison control panel, the user may select any number of data sets and visualize the scalar field height profiles at any elevation in the furnace. Sliders exist to change the elevation and translucency of the cutting planes and to cull data in order to isolate regions of interest. Button toggles exist for changing the data color maps for a scalar field, changing the scalar fields visualized, projecting the cutting planes, and enabling the option to visualize the computational grid. A control panel also was created to allow the user to set various options that select the view (top, side, back, or oblique), the wireframe option verses texture-mapped view of the furnace, and the resolution of the data sets. Because there are two distinct fluid regimes in this application—the combustion gases and the molten aluminum—the user also has the ability to choose to show data from either regime, or both. Figure 9 shows the control panel developed to interact with the data comparison techniques.
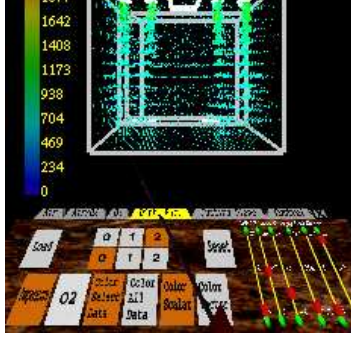
Figure 9: The control panel developed to interact with with the data comparison techniques

5 CONCLUDING REMARKS

In this article we described a new visualization toolkit that can be used to directly compare simulation data from a number of different runs in an immersive environment. The three primary functionalities of the toolkit include successive visualization of individual data sets, data differencing techniques, and image differencing techniques. Each technique added valuable insights into the simulation results; in particular, we found the following aspects of comparative visualization to be the most useful.

- The ability to interactively define the ranges for computing color maps was useful in a number of different ways: access to the global, data set, and cutting plane extrema values in the scalar field allowed the user to explore the details of individual data sets and simultaneously obtain a sense of the magnitude of the scalar field relative to the other data sets. Being able to color by source or by scalar enabled the user to discern small differences in the scalar field height profiles.

- Having the ability to interactively define ranges $[m_u, M_u]$ both to cull data and to define the color map allowed the user not only to isolate regimes of interest and examine them in detail but also to successively visualize the portions of each data set that exist in that regime.

- Both the data and image comparison tools provided significant advantages over traditional side-by-side or successive visualization paradigms for comparing multiple data sets. The data comparison tool was very useful in obtaining a global, three-dimensional image of the temperature distributions differences among the data sets. When examining the difference between the O_2 and air/O_2 fuels, magnification and exaggeration where used to highlight the subtle differences between the data sets.

Future work includes expanding the ADE toolkit to include a number of new functionalities. First, we will expand the data comparison functionality to support more general operations, particularly general linear combinations of two or more data sets. We are also planning to incorporate the communication mechanisms provided by the ALICE Memory Snooper [1], so that the ADE toolkit will be able to dynamically retrieve data from multiple time steps of an ongoing simulation, allowing the application scientist to monitor the progress of the application's solution and to investigate the differences between time steps. Finally, for the very large data sets of interest to many application scientists, we must incorporate the ability to perform general data reduction for interactive visualization and manipulation. Achieving this task will involve understanding the

errors associated with coarsening the data sets and their relation to the differences between data sets.

ACKNOWLEDGMENTS

REFERENCES

[1] Ibrahima Ba, Christopher Malon, and Barry Smith. Design of the ALICE Memory Snooper, http://www.mcs.anl.gov/ams, 1998.

[2] FLUENT Corporation, http://www.fluent.com, 1999.

[3] C. Cruz-Neira, D. J. Sandin, and T. A. DeFanti. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *ACM SIGGRAPH 93 Proceedings*, pages 135–142. ACM, 1993.

[4] Joe Insley. Control Panel User Interface for the CAVE, Private Communcation, Argonne National Laboratory, 1998.

[5] Kwansik Kim and Alex Pang. Projection-based data level comparisons of direct volume rendering algoriths. Technical Report UCSC-CRL-97-16, University of California at Santa Cruz, 1997.

[6] Hans-Georg Pagendarm and Frits Post. Comparative visualization – approaches and examples. In *Visualization in Scientific Computing*, pages 95–108. Springer, Wien, 1995.

[7] Hans-Georg Pagendarm and Frits Post. Comparative visualization of flow features. In *Scientific Visualization: Overviews, Methodologies, and Techiques*, pages 211–227. CS Press, 1997.

[8] Jens Trapp and Hans-Georg Pagendarm. Data level comparative visualizaton in aircraft design. In *Proceedings of Visualization '96*, pages 393–396. IEEE, 1996.