12TH INTERNATIONAL SYMPOSIUM ON FLOW VISUALIZATION
September 10-14, 2006, German Aerospace Center (DLR), Göttingen, Germany

INTERACTIVE VISUAL ANALYSIS AND EXPLORATION OF COMPLEX FLOW SIMULATION DATA WITH SIMVIS

Helmut Doleisch, Helwig Hauser
VRVis Research Center, Vienna, Austria
email:{Doleisch|Hauser}@VRVis.at, http://www.SimVis.at/

Keywords: *interactive visual analysis, visualization of flow simulation data, focus-plus-context visualization, multiple linked views*

ABSTRACT

SimVis is a novel technology for the interactive visual analysis of flow data which results from Computational Fluid Dynamics (CFD) simulation. The new technology which has been researched and developed over the last years here at VRVis Research Center in Vienna, introduces a new approach for interactive graphical exploration and analysis of time-dependent data (computed on large three-dimensional grids, and resulting in a multitude of different scalar/vector values for each cell of these grids). In this paper we present the major new technological concepts of our research.

1 INTRODUCTION

Visualization of complex results from computational fluid dynamics (CFD) simulation has become a very active field of research. With increasing computational power of computing systems both the frequency of CFD simulation being used as well as the complexity of the calculations (and consequently also the complexity of the results) has increased. The complexity is especially caused by the multi-dimensional, as well as also by the time-dependent character of the simulation. Data sets, that contain millions of data points with dozens of simulated flow attributes and possibly hundreds of time steps are common nowadays and require specialized tools to be handled. Visualization can be used to support the exploration and analysis of these data sets.

Visualization of flows was done long before the advent of the computer, for example with the injection of dyes and bubbles into flows. The relatively recent developments in computational flow visualization (over the last ten to twenty years) have led to visualization and interaction methods that now enable the user to not only see, but also to analyze and interact with the data graphically. Thereby it becomes possible, to quickly assess the results of a simulation (or measurement), and to ask and answer questions which are related to the respective problem.

SimVis is a system for the graphical analysis of simulation data, built on a new, cutting-edge technological approach for interactive visual analysis of large, multi-dimensional, and time-dependent data sets resulting from CFD simulation. The different components of the SimVis framework are discussed in more detail further below. The here presented new approach allows the user to view as well as to query the data, and thus to gain new insights into the simulation results, based on the engineers knowledge. In contrast to automatic feature detection systems, SimVis allows more direct and flexible access to the data, and makes it also easier to understand the results of the analysis process through the tight integration of the engineer.

After this introduction, section 2 explains different visualization approaches and the visualization process in general. In section 3 details about the basic components of the SimVis framework and the new technology behind this system are discussed. Afterwards, in section 4 the analysis of four application examples is presented, showing the flexibility and intuitiveness of this new visualization technology. The paper is rounded up by giving some hints on performance issues and presenting a few concluding remarks.

2 COMPUTATIONAL FLOW VISUALIZATION – DIFFERENT APPROACHES

With the growing use of CFD simulations in many application fields, also the demand for exploration, analysis, validation and presentation of the simulation results increases. These steps can be supported by computational visualization and related graphics techniques and are mainly targeted throughout this paper. Computational visualization (or postprocessing, as it is also called in engineering terms) is useful for more than just viewing the computed flow field. It can help with understanding the nature of the problem, with identifying and recognizing interesting relationships of different variables in the simulation output, and also with debugging the simulation process itself.

In computational flow visualization research different approaches have been in the focus of research over the last years. A few detailed state of the art reports capture most of these works and discuss advantages and disadvantages of the different approaches [17, 15, 18]. The new technology presented by us in this paper is in some sense related to previous feature-based flow visualization approaches, although we employ an interactive setup for the specification of features as opposed to classical feature-based approaches, where the selection of features mostly is (semi-)automatic, and thus hidden from the expert.

2.1 Visualization Process

As already mentioned above, computational visualization aims at supporting the tasks of exploration, analysis, and presentation of (typically) large amounts of data through graphical representations (images, videos, interactive representations), and therefore, visualization approaches can be differentiated according to how well they fit into the three stages of the visualization process [1].

- *Visualization for Exploration:* exploration is usually the first step in data investigation. Before the user can analyze a data set, exploration is carried out, supporting the user to find out certain characteristics about the data set. Examples include which dimensions are likely to play a major role during the following analysis steps, or which structures are of interest in the given data. Another goal of exploration can be to check whether the data appears to be valid, and to find and remove any obvious problems, e.g., wrong or non-converging results of a simulation process. For these purposes visualization tools supporting exploration typically should provide maximum flexibility.

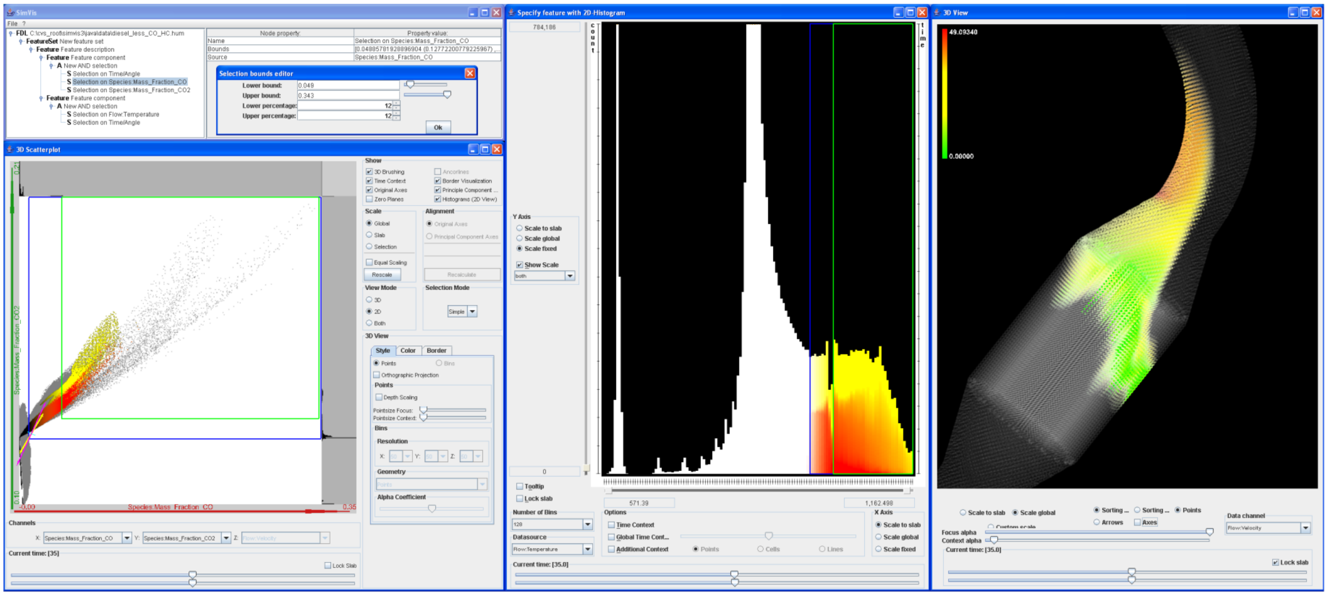

FDL C:\cvs_root\simvis3\java\data\diesel_less_CO_HC.hum
- FeatureSet New feature set
  - Feature Feature description
    - Feature Feature component
      - A New AND selection
        - S Selection on Time/Angle
        - S Selection on Species:Mass_Fraction_CO
        - S Selection on Species:Mass_Fraction_CO2
    - Feature Feature component
      - A New AND selection
        - S Selection on Flow:Temperature
        - S Selection on Time/Angle

Node property:	Property value:
Name	Selection on Species:Mass_Fraction_CO
Bounds	[0.04885781928896904 (0.12772200779225967) ,...
Source	Species:Mass_Fraction_CO

Lower bound: 0.049
Upper bound: 0.343
Lower percentage: 12
Upper percentage: 12

Show: ☑ 3D Brushing ☐ Ancorlines ☑ Time Context ☑ Border Visualization ☑ Original Axes ☑ Principle Component ... ☐ Zero Planes ☑ Histograms (2D View)
Scale: ● Global ○ Slab ○ Selection ☐ Equal Scaling
Alignment: ● Original Axes ○ Principal Component Axes
View Mode: ○ 3D ● 2D ○ Both
Selection Mode: Simple
3D View: ● Points ○ Bins ☐ Orthographic Projection
Points: ☐ Depth Scaling; Pointsize Focus:; Pointsize Context:
Bins: Resolution X: 50 Y: 50 Z: 50
Geometry: Points
Alpha Coefficient
Channels: X: Species:Mass_Fraction_CO Y: Species:Mass_Fraction_CO2 Z: Flow:Velocity
Current time: [35] ☐ Lock Slab

784,186
0
Y Axis: ○ Scale to slab ○ Scale global ● Scale fixed
☑ Show Scale: both
☐ Tooltip
☐ Lock slab
571.39
1,162.498
Number of Bins: 128
Datasource: Flow:Temperature
Options: ☐ Time Context ☐ Global Time Cont... ☐ Additional Context ● Points ○ Cells ○ Lines
X Axis: ● Scale to slab ○ Scale global ○ Scale fixed
Current time: [35.0]

○ Scale to slab ● Scale global ○ Custom scale
● Sorting ... ○ Sorting ... ● Points ○ Arrows ☐ Axes
Focus alpha
Context alpha
Data channel: Flow:Velocity
Current time: [35.0] ☑ Lock slab

Fig. 1 A sample SimVis scenario: simulated flow through a diesel particle filter (DPF) is visualized – the flow is shown at the time of 35 secs. after simulation started. The user has reflected his interest in flow regions of heavy oxidation by interactively brushing data items which exhibit a lot of carbon-oxides in the scatterplot (lower left) and then refining this specification to only apply to hot regions (in the histogram, middle). The 3D view on the right shows a focus+context visualization of the DPF with the brushed data items highlighted in color (color shows velocity magnitudes).

- *Visualization for Analysis:* based on hypotheses which emerged during the exploration phase (or also independently from exploration) the data is now visually analyzed. The final goal is to provide a thorough analysis of all the structures or processes of interest in the data. Verification or falsification of hypotheses can lead to new questions, which also are investigated and analyzed. Tools which allow flexible analysis of the data must provide some sort of querying possibilities. Therefore, interactivity is a key feature of visualization as part of the analysis process.

- *Visualization for Presentation:* the results and findings gained during the analysis of a data set eventually need to be presented and communicated to others. Visualization for presentation usually reduces the information to be shown to what is absolutely necessary to show a certain connection or fact, and therefore often needs to be relatively simple but at the same time effective. Here not the interactivity of an investigation is a primary goal, but rather a high visual quality of the representation of the results.

In practice, the three phases of visualization cannot be separated easily, and many tasks appear in at least two phases of this classification. When compared to traditional computational visualization approaches for CFD simulation data sets, the here presented SimVis approach is mainly targeted towards interactive analysis and exploration tasks, and not so well suited for presentation issues.

3 THE SIMVIS SYSTEM

SimVis is a multiple-views technology for the interactive and feature-based analysis of large and high-dimensional data from CFD simulation [1, 2, 3, 4]. SimVis is the result of several years of cutting-edge computational visualization research at the VRVis Research Center in Vienna, Austria.

As the prime goal of our visualization research was to support the tasks of exploration and analysis of very large and complex data sets resulting from CFD simulation, interactivity is a key element in our approach. To enable interactivity, several previously known concepts from classical computational visualization research have been combined and adopted accordingly in the SimVis framework. In the following we shortly describe the key concepts of SimVis, for in-depth details a number of related publications can be consulted [1, 2, 3, 4].

3.1 Multiple Linked Views, Brushing

An important concept in computational visualization, which has also been used in our approach, is the use of *multiple, linked views*. Different views can be used to show different aspects (e.g., dimensions) of the data, or they can also show the same dimensions using different representations. In SimVis, one view usually is used to show the three-dimensional layout of the data set (possibly also animated over time), while many other views mainly show simulated flow attributes or similar values (see figure 1 for a typical session setup). These additional views include, for example, simple 2D scatterplots (figure 1, left lower view), time-dependent histograms [13], combined 2D/3D scatterplots [16], or parallel coordinates [12, 11]. An example of a 3-views layout is shown in figure 1, where on the right hand side a 3D view is used to show the spatial layout of the data set (a diesel particle filter application for an exhaust system), in the middle a historgram is used to show the temperature distribution in the data, and on the left side a 2D scatterplot view plots data values of two data attributes (carbon-oxides) against each other.

Another very important concept, often used in conjunction with the usage of multiple views, is *linking and brushing*. Brushing of data means marking data points interactively in the graphical display, thereby selecting the respective data items, e.g. for emphasizing them in the visualization. Formally, brushing assigns a *degree-of-interest (DOI)* value $DOI_j \in [0,1]$ to each data point (compare to Furnas [8]). A DOI value of 1 means that the data item is brushed, a value of 0 stands for a not-brushed data item. Brushed data items are called to belong to the focus, non-brushed items belong to the context of a so-called *focus+context (F+C) visualization*. In a F+C visualization, the interactively brushed data items are emphasized, whereas the rest of the data (the context parts) are shown in a less prominent style for orientation in the whole data domain. Brushing ist often combined with linking, and also in our setup we have built on this combined approach. By linking multiple different views, the brushing information is propagated to all these linked views. Thus, marking data points in one view effects the visual representation in all the linked views.

One of the first examples of linking and brushing different visualization approaches in different views is a system called WEAVE [9], which was used to interactively analyze and visualize simulated data of a human heart application using a focus+context style. Figure 1 illustrates the linking and brushing mechanism as employed in our approach. In this session, the user has reflected his interest in flow regions of heavy oxidation by interactively brushing data items which exhibit a lot of carbon oxides in the scatterplot. This brushed information has been propagated to the histogram and also the 3D view. Then, in a second step, the brushed information has been refined by performing a second brushing operation in the histogram, selecting all data items exhibiting high temperature values. The two resulting DOI values of both brushes are logically AND-combined and the resulting areas are emphasized in the F+C visualization in the 3D view.

3.2 Fuzzy Classification

In most applications of visualization there are only binary or discrete classifications used to establish a semantic layer on top of the originally unlabeled data [18]. In medical visualization, for example, object segmentation plays an important role, and usually discrete object maps are used to label voxels of either being part of one object or another. Similarly, in flow visualization, also usually a sharp feature extraction process is used to discretely partition the flow domain into portions which represent certain flow features, e.g., a vortex or a recirculation zone.

In SimVis, fuzzy classifications (according to the terminology of fuzzy logic [20]) are used to assign probabilities of class containment. This happens interactively via *smooth brushing* [3] in the data distribution views with respect to what is currently of interest for the user. Fuzzy logic operations are used to establish a calculus, which is based on fuzzy DOI values. This is an important feature of SimVis as it is often not possible to sharply delimit flow portions of interest from all the rest- usually between a certain region of full, i.e., 100% user interest and completely uninteresting portions of the flow (DOI-values of 0) a border region exists for which a gradual change of DOI-values is assumed.

3.3 Iterative and Interactive Feature Specification

SimVis utilizes an iterative and interactive approach to feature-based visualization of large and complex data. In SimVis, the setup of synthetic DOI attributes is usually started by a simple selection of data items in one view (for example through brushing data values in a scatterplot). After investigating the visual response of this first step (e.g., in the 3D view), iterative refinement is performed to furthermore detail the feature specification in any of the other views (as done in the histogram of figure 1 as a second step, for example). The result of such an iterative process is a complex feature specification which is of hierarchical structure and usually involves a set of different data dimensions [2].

Additionally, users often also want to refine certain parts of a feature specification numerically, especially when certain thresholds carry a specific meaning, e.g., the boiling temperature of water at $100°C$. The SimVis approach incorporates an explicit representation of all this information, i.e., the hierarchical feature specification with related parameters, called *feature definition language*, together with a separate user interface which enables the direct manipulation of feature specifications [2] (compare to the upper left parts of figure 1). Such an explicit representation of the feature specification process also immediately enables load- and save-functionality which consequently results in additional advantages such as the opportunity to compare data sets by setting up an analysis for one data set, then saving it, and re-applying the same analysis to another data set.

3.4 Focus+Context Visualization

3D viewing is an essential component of the SimVis approach – users usually want to be spatially oriented. Up to now, SimVis was successful with rather simple, glyph-based 3D viewing. For every data item in the flow data a glyph (e.g. a small 3D sphere or arrow) is drawn with a certain opacity and color. The size of the glyphs is adjusted locally through a transfer function in dependence on a DOI value and globally through a user-defined scaling factor. One important fact why this approach works so well is the focus+context extension of 3D viewing [10]: according to the feature specification process, where data items are attributed with degree-of-interest values, glyphs are drawn in an emphasized (rather opaque and colored for data items in focus) or in a reduced (rather transparent and in shades of gray) style (right view in figure 1). Data items, which are associated with fractional DOI values (resulting from smooth brushing), are represented in an interpolated way (interpolated opacity and color). This focus+context rendering of data, which is laid out in 3D space, enhances the perception on the user side as an efficient part of occlusion control.

3.5 Attribute Derivation and Advanced Brushing

Brushing is an intuitive and very effective, but still very simple concept to indicate user interest. One way to characterize this kind of data classification is to understand it as a shallow, broad-band approach. Major advantages of the brushing approach are:

- that the user is not bound to specific extraction procedures, which usually are linked to specific mathematical formulae, but can select whatever is interesting (especially useful through data exploration),
- that the feature extraction process is easily comprehended by the user (since selections are formulated in explicit terms of the data, no ”magic” is going on behind feature extraction), and
- that interactive brushing perfectly fits with an iterative refinement approach where features can be formulated step by step which eases to track down interested regions of the data.

One disadvantage of brushing is that it indeed does not enable the extraction of really complex relations within the flow (which nevertheless often are of great interest). SimVis incorporates two approaches to deal with this limitation of simple brushing [4]: On the one hand, SimVis offers *advanced brushing* mechanisms such as angular brushing [11] which are useful to "dig out" relations within the data which are more complex as compared to the brushing standards. On the other hand, SimVis offers opportunities to interactively *derive* additional synthetic *data attributes* on the basis of comprehensible mathematical formulae such as gradient derivation, data smoothing, similarity measures, etc. Through the combination of these two approaches (attribute derivation and advanced brushing) it becomes possible to extract rather complex features from the flow data, which are comparable to sophisticated feature extraction processes.

Another extension, that has been recently added was the integration of *time-dependent feature specification* [4]. We consider *time-dependent features* to be those flow features which are inherently dependent on time – these features cannot be extracted from singular time steps of unsteady CFD data. In an informal user study with application engineers we identified different types of important time-dependent features, including for example features based on attribute derivates, relative feature specification or features based on local temporal extrema. As a detailed discussion would go beyond the scope of this paper and we therefore refer to the description in a separate report [4].

4 ANALYSIS EXAMPLES

SimVis has been applied in several case studies and application examples from different fields [5, 6, 7, 14, 19], four of which are described shortly in the following. These analysis examples are shown to demonstrate that the SimVis approach is generally applicable. Our original goal was to provide an

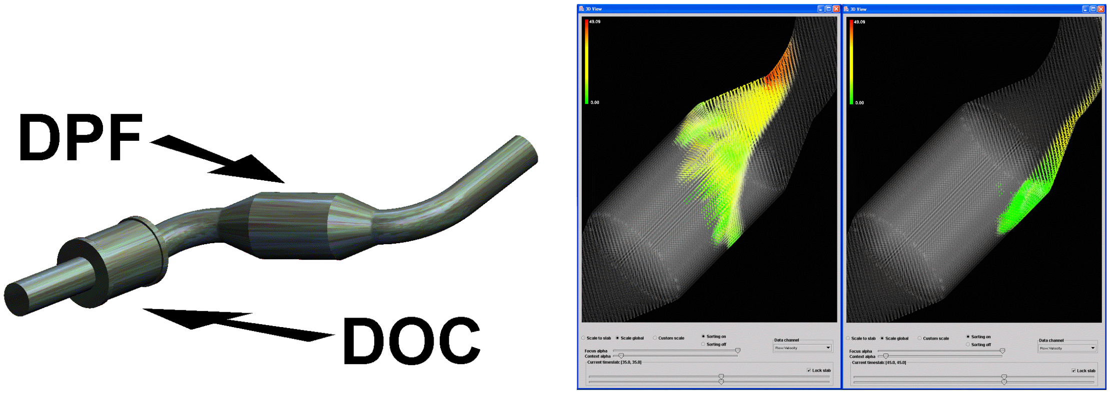

Fig. 2 Left: layout of a diesel exhaust system consisting of a diesel oxidation catalyst (DOC) and a diesel particulate filter (DPF). Right: velocity in cells of the DPF with high CO and CO_2 mass fraction and high temperature after 35 sec. (middle) and 45 sec. (right) of the simulation.

analysis and exploration technology for results from CFD simulation, especially from the automotive industry. Thus three of the following examples are out of the automotive field. But recently we also employed SimVis technology for the analysis of data from other application fields, e.g. simulated or measured data from the meteorological or medical domain, and thereby proved the generic nature of our approach. In the following we give a few very short overviews about some application examples. More examples, as well as also more detailed descriptions of the here dicussed analysis examples are available at http://www.simvis.at/.

4.1 Visual Analysis of a Diesel Exhaust System

This first analysis example is part of a case study on improving a diesel exhaust system for passenger cars powered by diesel engines [6]. Figure 2, left, shows the layout of the diesel exhaust system's geometry, consisting of a diesel oxidation catalyst (DOC), used to reduce hydrocarbons and CO emissions, and a diesel particulate filter (DPF), trapping the diesel particulates (soot) of the exhaust gas in its filter material. Over time the collected particulates block the filter, which would negatively affect the engine operation. Therefore a periodically applied filter regeneration mechanism has been developed: the collected particulates are oxidized at high temperatures to gaseous products, primarily CO_2.

In our detailed case study on this exhaust system results of different settings for the DPF regeneration have been analyzed. Four different application questions have been investigated in tight collaboration with engineers at AVL List GmbH, a partner company of our research center. One of these application questions is now discussed in detail, a description of the full case study is available at http://www.VRVis.at/via/research/diesel-case/.

Application Question: Where and how fast does the soot oxidize?

Due to soot oxidation during the regeneration phase, CO and CO_2 are generated. Regions of high CO and CO_2 concentration, where also the temperature is relatively high denote regions of current soot oxidation in the regeneration process. To display these areas, it is necessary to specify a complex feature. First in a scatterplot the mass fraction of CO_2 (Y-axis) is plotted against the mass fraction of CO (X-axis). A

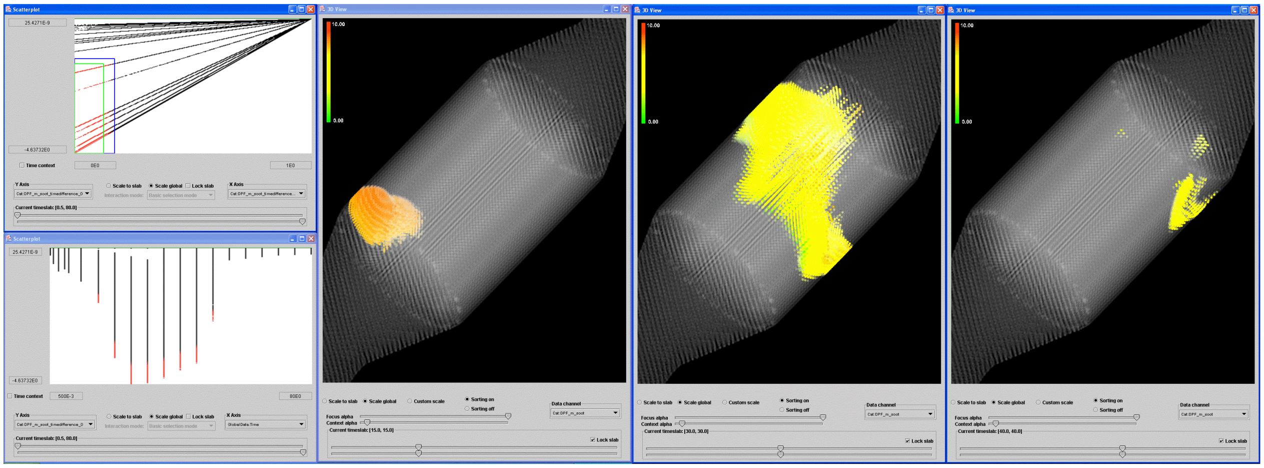

Fig. 3 Soot mass in cells with high oxidation rate after 15, 30, and 40 sec. (from left to right).

smooth brush selecting higher values of both attributes is applied (see scatterplot view in figure 1). Then, in a second step, to refine the first 2D brush, a histogram showing the data distribution of temperature values is used (see histogram view of figure 1). Here, high values of fluid temperature are brushed, leading to a 3D, composite brush. In figure 2, a F+C visualization of the hereby specified regions in the DPF in the 3D view for two different timesteps of the simulation is shown. Note, that velocity values are color mapped, giving a hint on the flow speed in the respective regions. This visualization shows, that CO and CO_2 production are not symmetric. As can be seen in the right 3D view, the velocities in this region are relatively small (colored green), which seems to be the reason for this asymmetric temporal behavior of the oxidation process.

For a better understanding, the oxidation progress has to be analyzed. This can be done by displaying only cells with a high mass soot gradient over time. As the amplitude of the gradient values changes over time, we need a method to select relatively high gradient values with respect to the maximum gradient value for each time step, to get interesting cells for each time step separately. Therefore differences of the soot mass values over time are calculated and normalized. With normalized differences we can easily select the upper 20% of difference values per time step, for example.

After calculating the differences of the soot mass values and normalizing them, a scatterplot showing the normalized mass soot gradients (X-axis) and the mass soot gradients (Y-axis) is used. Low mass soot gradients and low normalized mass soot gradients are brushed in this scatterplot (see figure 3, upper left). Note that central differences are used for the approximation of gradients, which in the case of mass soot oxidation are negative, thus the lower gradient values are selected. For better illustration and better exploration a second scatterplot showing the mass soot gradients (Y-axis) and the time domain (X-axis) is shown in figure 3. Here the changing range of mass soot gradients becomes visible very intuitively. The vertical cluttering of data items results from the discretization of time on the X-axis, each streak denoting one time step in the temporal domain. For color mapping in the three 3D views, the mass soot values are taken into account. From figure 3 it is clear, that the region with the highest soot oxidation rate after 15 seconds is not symmetric (left 3D view). The highest oxidation rate at the beginning of the soot oxidation is in the area with the highest temperatures (see full case study report [6]). After further investigations of the oxidation behaviour, it was detected, that there is not enough O_2 for oxidation at

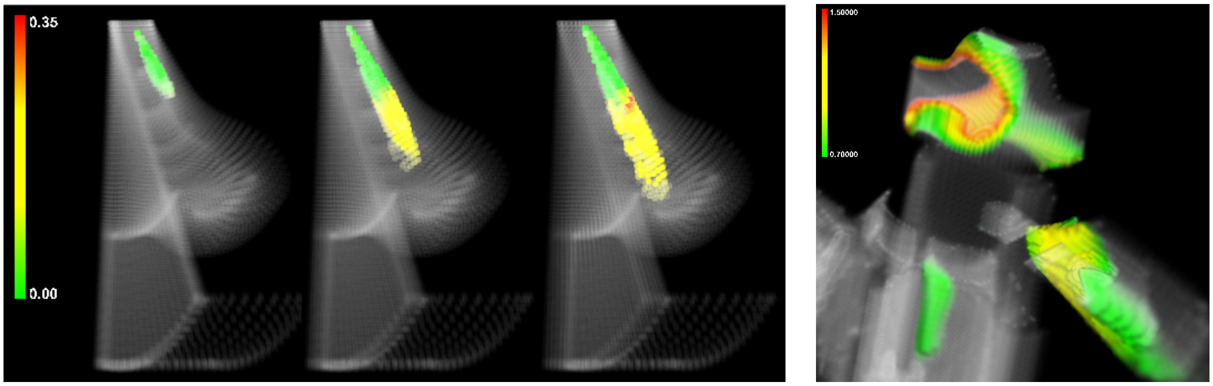

Fig. 4 Left: inspecting the injection process in a diesel engine for a heavy truck. Right: visualizing and analyzing regions containing a optimal burnable amount of fuel-air mixture.

the region displayed in the right most 3D view of figure 3. This is caused by an asymmetric asymmetric flow field, due to the bend of the geometry between the DOC and the DPF (see detailed report for this case [6].

4.2 Visual Analysis of Fluid Dynamics in Internal Combustion Engines

We performed several case studies on investigating and analyzing fluid flow dynamics in internal combustion engines [5, 19]. For these applications a concept of handling and analyzing data on time-varying grids (moving meshes) has been realized. In one application example, the analysis of the fuel injection and combustion process in a diesel engine for a heavy truck has lead to an improved injection design [5] (see figure 4, left, for a series of timesteps showing the fuel injection region). In another, very recent application example data coming from a full 720° crank angle CFD simulation of a 2-stroke engine with high pressure GDI injection has been analyzed together with experts from the internal combustion engines research field [19] (see figure 4, right side, for the visualization of regions containing a burnable amount of fuel-air mixture).

4.3 Visual Analysis of Simulated Hurricanes

In this application example, a visual analysis of the simulated hurricane Isabel, which struck the US eastern coast in 2003, was performed [7]. With our approach of the visual analysis of this application example we won the 1st price at the prestigious international IEEE Visualization Contest 2004 in the US. To get a first overview of the hurricane data set, we investigated the eye of the hurricane (characterized by comparably low pressure and relatively low wind speeds and interactively specified with a brush on a scatterplot of pressure and velocity values, see figure 5, left for one timestep of the resulting visualization) as well as the cloud structures (brushing relatively high values of the simulated cloud-attribute, see figure 5, middle left). Visualizing clouds helps to get a very intuitive picture of the hurricane, fast clouds around the eye of the storm as well as comparably slow clouds (thunderstorms at three fronts) give a good overview of the data.

Apart from the eye of the hurricane and cloud structures, we investigated a dozen of other properties of this data set as described in much more detail on the related webpage http://www.simvis.at/Isabel/.

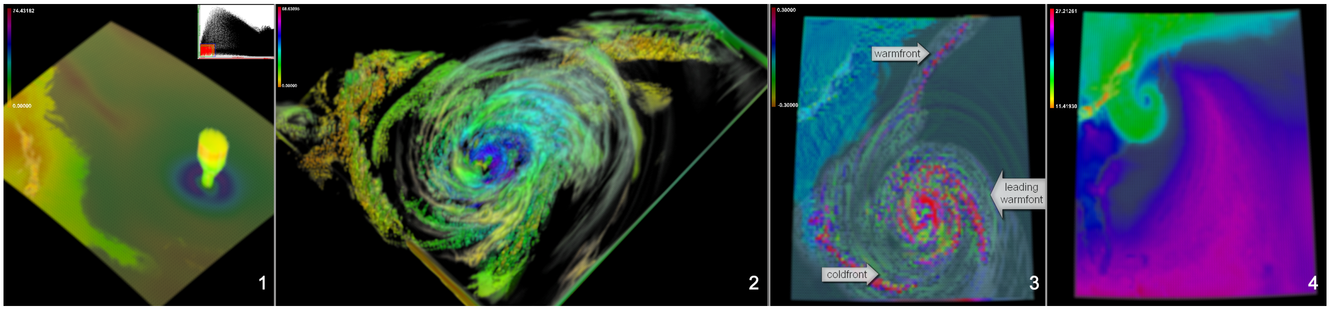

Fig. 5 Example results from the visual analysis of a simulated hurricane: (left) selecting the eye of the storm and coloring wind velocities over the sea off-shore the US east coast; (middle-left) showing fast moving clouds; (middle-right) extraction of the different fronts reveals a warm sector - regions exhibiting high precipitation are high-lighted; (right) emphasizing the influence of the warm sector by coloring according to temperature values.

Through the linking-and-brushing methodology of the SimVis approach, hurricane structures such as fronts (see figure 5, middle right), the temperature distribution (see figure 5, right), amounts of precipitation, classified wind velocities, supercoold cloud parts, convection areas, etc., can be examined, revealing interesting information about the simulated two days of this storm.

More recently we have also had the chance to investigate simulation data provided by the National Center of Atmospheric Research (NCAR) for the prediction of the desastrous hurricane Katrina, which struck New Orleans in August 2005. Detailed results including many videos showing the interactive visual analysis process are available at http://www.simvis.at/Katrina/.

4.4 Visual Analysis of Flow through a Cooling Jacket

A very recent application example includes the analysis of fluid flow through a cooling jacket [14]. This engine component has an important role in transferring heat away from the engine block. In this case study we interactively investigated important questions for engineers, two of which are: *(1) are there any areas where the flow is moving in the wrong direction?* and *(2) where in the cooling jacket are the areas of stagnant and hot flow?* More details about the study, including a discussion of further application questions, are available in a seperate paper [14].

As mentioned above, the function of a cooling jacket is to transfer heat away from the engine as efficiently as possible. Engineers are interested in learning where and how the flow deviates from the ideal, leading to less effective heat transfer. Figure 6 shows two example analysis results for the questions asked above. In the left view, all areas are highlighted, where the flow direction deviates from the main flow direction following the shortest path through the geomtry from the inlet to the outlet of the cooling jacket. The shown feature has been interactively specified by brushing negative flow components in transversal and longitudinal direction and logical AND-combination of the resulting DOI values. In the right view, areas exhibiting nearly stagnant flow and hot temperatures above the threshold of $364°K$ are shown. From these results it becomes obvious, that the resulting cooling jacket design has only a few small critical regions, and thus is close to the optimum, with respect to these two questions.

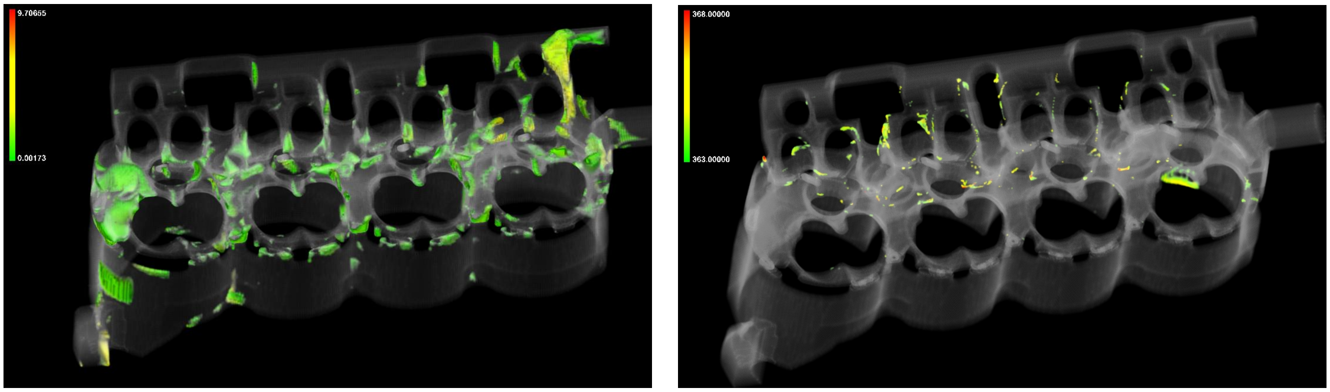

Fig. 6 Analyzing a cooling jacket for diesel engines: (left) the result of selecting all regions of reverse-longitudinal flow *and* regions of reverse-transversal flow; (right) areas of temperature $t > 364^\circ K$ and velocity $|\mathbf{v}| < 0.1 m/s$ are interactively-specified by the user and rendered in focus.

5 PERFORMANCE ISSUES

The SimVis technology is designed to work on regular PCs. Apart from certain size limitations, which restrict the amount of data to less than 3GB which can be held in memory currently, interactive visual analysis and exploration is very well possible even without very expensive hardware. The above presented case studies and application examples have been carried out in realtime on a PC-system consisting of the following components: Intel Pentium4 2.8GHz, 2GB RAM and a Nvidia GeForceFX 5950 graphics card (with 256MB of RAM). The data sets analyzed consist of between 260.000 cells (DPF-example) and 3.000.000 cells (hurricane-example), results for dozens of timesteps and attributes have been analyzed interactively.

6 CONCLUSIONS

In this paper we have presented the SimVis system, a framework for interactive visual exploration and analysis of large, time-dependent, and high-dimensional data sets resulting from CFD simulation. SimVis provides a flexible fusion of different visualization technologies to enable the interactive visual access to spatiotemporal data properties as well as to all the other simulated data attributes. Interactive feature specification is intelligently linked to focus+context visualization in 3D – color and opacity are made dependent on multiple data attributes in an intuitive way. Immediate visual feedback allows to efficiently explore and analyze the data. The flexibility of our framework allows to interactively analyze data from many different applications of CFD simulation. Besides CFD simulation data, also the analysis of any other spatiotemporal data, which has a high-dimensional attribute space assigned to each point in space and time, fits perfectly to our approach. We have, for example, also proven the effectiveness of our analysis approach on investigating measured radar data for severe weather forcasting, recently.

The proposed methodology of the SimVis approach adds new opportunities for exploration and analysis of simulation results. The results of our work are not considered as a replacement for existing technologies, nor to be the exclusive solution for the tasks aimed for. Nevertheless, by using the proposed concepts of our framework, additional help for these tasks can be identified. Especially the flexibility offered by an interactive approach is beneficial for engineers, who are running and supervising CFD simulation cases. The interactive approach as presented here allows the users to specify their interest during an investigation of the resulting data in a very intuitive way.

ACKNOWLEDGEMENTS

The authors would like to thank all their colleagues, who helped during the development process of the SimVis approach. Special gratitude goes to Philipp Muigg, Martin Gasser, and Harald Piringer. Furthermore the authors would also like to thank Michael Mayer from AVL List GmbH and Stephan Schmidt from the Institute of Internal Combustion Engines and Thermodynamics, Technical University of Graz, Austria, for their collaboration in the application examples. Most of the data sets presented in the application example are courtesy of AVL List GmbH, the hurricane Isabel data set is courtesy of NCAR, the 2-stroke engine data set has been provided by our collaboration partners from the combustion engines research group at the Technical University of Graz. Parts of the presented work have been done in the scope of the basic research on visualization (http://www.VRVis.at/via/) at the VRVis Research Center in Vienna, Austria, which is funded by the Austrian research program Kplus.

REFERENCES

[1] H. Doleisch. *Visual Analysis of Complex Simulation Data using Multiple Heterogenous Views*. PhD thesis, Technical University of Vienna, see also http://www.VRVis.at/via/resources/diss-HD/, October 2004.

[2] H. Doleisch, M. Gasser, and H. Hauser. Interactive feature specification for focus+context visualization of complex simulation data. In *Proc. of the 5th Joint IEEE TCVG - EUROGRAPHICS Symposium on Visualization (VisSym 2003)*, pages 239–248. Springer-Verlag, 2003.

[3] H. Doleisch and H. Hauser. Smooth Brushing for Focus+Context Visualization of Simulation Data in 3D. In *Journal of WSCG*, volume 10, pages 147–154, Plzen, 2002.

[4] H. Doleisch, H. Hauser, M. Gasser, and R. Kosara. Interactive focus+context analysis of large, time-dependent flow simulation data. Technical Report TR-VRVis-2004-024, currently in submission, VRVis Research Center, 2004.

[5] H. Doleisch, M. Mayer, M. Gasser, P. Priesching, and H. Hauser. Interactive feature specification for simulation data on time-varying grids. In *Proc. of the Conference Simulation and Visualization (SimVis 2005)*, pages 291–304, Magdeburg, Germany, March 2005.

[6] H. Doleisch, M. Mayer, M. Gasser, R. Wanker, and H. Hauser. Case Study: Visual Analysis of Complex, Time-Dependent Simulation Results of a Diesel Exhaust System. In *Proc. of the 6th Joint IEEE TCVG - EUROGRAPHICS Symposium on Visualization (VisSym 2004)*, pages 91–96. Springer-Verlag, 2004.

[7] H. Doleisch, P. Muigg, and H. Hauser. Interactive visual analysis of hurricane isabel. Technical Report TR-VRVis-2004-058, VRVis Research Center, see also http://www.simvis.at/Isabel/, 2004.

[8] G. Furnas. Generalized fisheye views. In *Proc. of the ACM CHI '86 Conf. on Human Factors in Computing Systems*, pages 16–23, 1986.

[9] D. Gresh, B. Rogowitz, R. Winslow, D. Scollan, and C. Yung. WEAVE: A system for visually linking 3-D and statistical visualizations, applied to cardiac simulation and measurement data. In *Proceedings IEEE Visualization 2000*, pages 489–492, 2000.

[10] H. Hauser. Generalizing focus+context visualization. In *Scientific Visualization: The Visual Extraction of Knowledge from Data (Proceedings of the Dagstuhl 2003 Seminar on Scientific Visualization)*, pages 305–327, 2005.

[11] H. Hauser, F. Ledermann, and H. Doleisch. Angular brushing of extended parallel coordinates. In *Proc. IEEE Symposium on Information Visualization 2002 (InfoVis 2002)*, pages 127–130, 2002.

[12] A. Inselberg and B. Dimsdale. Parallel coordinates: a tool for visualizing multidimensional geometry. In *Proceedings IEEE Visualization '90*, pages 361–378, 1990.

[13] R. Kosara, F. Bendix, and H. Hauser. TimeHistograms for large, time-dependent data. In *Proc. of the 6th Joint IEEE TCVG - EUROGRAPHICS Symposium on Visualization (VisSym 2004)*, pages 45–54, 2004.

[14] R. Laramee, C. Garth., H. Doleisch, J. Schneider, H. Hauser, and H. Hagen. Visual analysis and exploration of fluid flow in a cooling jacket. In *Proceedings IEEE Visualization 2005*, pages 623–630, 2005.

[15] R. Laramee, H. Hauser, H. Doleisch, B. Vrolijk, F. Post, and D. Weiskopf. The state of the art in flow visualization: Dense and texture-based techniques. *Computer Graphics Forum*, 23(2), 2004.

[16] H. Piringer, R. Kosara, and H. Hauser. Interactive focus+context visualization with linked 2D/3D scatterplots. In *Proc. of the Intl. Conference on Coordinated & Multiple Views in Exploratory Visualization (CMV 2004)*, pages 49–60, 2004.

[17] F. Post, B. Vrolijk, H. Hauser, R. Laramee, and H. Doleisch. Feature extraction and visualization of flow fields. In *Eurographics 2002 State-of-the-Art Reports*, pages 69–100, Saarbrücken Germany, 2002.

[18] F. Post, B. Vrolijk, H. Hauser, R. Laramee, and H. Doleisch. The state of the art in flow visualization: Feature extraction and tracking. *Computer Graphics Forum*, 22(2):775–792, 2003.

[19] S. Schmidt, O. Schögl, R. Kirchberger, H. Doleisch, P. Muigg, H. Hauser, M. Grabner, A. Bornik, and D. Schmalstieg. Novel visualization and interaction techniques for gaining insight into fluid dynamics in internal combustion engines. In *Proc. of the NAFEMS World Congress Conference 2005*, 2005.

[20] L. Zadeh. Fuzzy sets. *Information and Control*, 8:338–353, 1965.