Case Study: Visual Analysis of Complex, Time-Dependent Simulation Results of a Diesel Exhaust System

Helmut Doleisch[1], Michael Mayer[1], Martin Gasser[1], Roland Wanker[2], Helwig Hauser[1]

[1] VRVis Research Center, Vienna, Austria, {Doleisch,Mayer,Gasser,Hauser}@VRVis.at http://www.VRVis.at/vis/
[2] AVL List GmbH, Graz, Austria, Roland.Wanker@avl.com, http://www.avl.com/

Abstract
In previous work we have presented visualization techniques that provide engineers with a high degree of inter-activity and flexibility for analyzing large, time-dependent, and high-dimensional data sets resulting from CFD (computational fluid dynamics) simulations. In this case study we apply our techniques in the fields of the automotive engineering industry and demonstrate how users benefit from using them during their routine analysis, as well as for exploring new phenomena. For coping with some of the special requirements in this application, we adapted and extended parts of the system. A comparison of two related cases of a diesel exhaust system is presented, and some important questions about these cases are addressed.

1. Introduction

The analysis of data which results from computational simulation is challenging and complex, but also important to speed up simulation-cycles, which leads to shortening design and development times of new products. One interesting application field for computational simulations is in simulating catalyst processes in a diesel exhaust system in the automotive industry.

Data resulting from CFD (computational fluid dynamics) simulations has a few special properties, including, for example, large amounts of data or complex geometrical models used for the simulations. The large amounts of data (typically in the range of billions of data values for each simulation cycle) result mainly from three different sources: detailed geometrical models (in terms of numbers of cells), time-dependency and multi-dimensionality (many attributes) for each cell of the model. Complex geometrical models as employed, e.g. in simulating diesel exhaust systems, pose problems in understanding relationships between different data items, if not presented in a visual form.

In this paper we present a case study, showing how visualization effectively helps to analyze complex, multi-dimensional, and time-dependent simulation results. For these types of applications, suitable visualization tools must fulfill certain requirements. One key requirement is interactivity, i.e., allowing the user to change visualization parameters interactively and interact with parts of the data. Then, providing high flexibility for exploring multi-dimensional data, allowing complex specifications of what parts of the data are of current interest, is also crucial for the users. Finally, users know their data very well, thus successful visualization tools must enable the users to access their data using a data-based approach.

When visualizing and analyzing large, 3D, and time-dependent data resulting from flow simulation, often (semi-)automatic feature extraction methods are used, as discussed in an overview by Post et al. [10]. Besides these feature extraction methods, only a small amount of work on interactive feature specification has been presented so far. Henze developed a technique to visualize and analyze time-varying CFD data by using multiple, linked views in a system called Linked Derived Spaces [8]. Gresh et al. presented a system called WEAVE [7], which combines information visualization (InfoVis) and 3D scientific visualization (SciVis) and allows investigations in the 3D context based on marking data according to different attributes.

In previous work [4, 5] we have developed visualization techniques, that focus on fulfilling the requirements discussed above, in a system called SimVis (see also section 3). We have shown that combining methods from SciVis and InfoVis provides powerful possibilities for visualizing multi-dimensional simulation results computed on 3D grids. In contrast to our system, common commercial systems for post processing and visualization of simulation results usually provide only views of 2D slices or surfaces and very limited interaction possibilities. Three different types of tools are mainly used: planar cuts, views of surfaces or iso-surfaces onto which specific attributes are color-mapped.

Figure 1: *A diesel exhaust system consisting of a diesel oxidation catalyst (DOC) and a diesel particulate filter (DPF).*

Figure 2: *Emission reductions in the DOC [9]*

In this paper, we present a recent case study that we carried out at the VRVis Research Center together with our partner AVL List. We use SimVis to demonstrate how the use of InfoVis and SciVis views can extend and improve the results of analysis of CFD data.

2. Application Scenario

The application presented in this case study is a diesel exhaust system for passenger cars powered by diesel engines. We first shortly describe such a diesel exhaust system, and afterwards discuss several important questions of engineers who work on the development of such a system.

2.1. Diesel Exhaust System

For passenger cars powered by diesel engines, layout and design of the *diesel oxidation catalysts* (DOC), used to reduce hydrocarbons and CO emissions, are standard applications. Because of decreasing legal limits for particulates, *diesel particular filters* (DPF) are in the focus of recent research [9]. In this case study a diesel exhaust system (see figure 1), consisting of a diesel oxidation catalyst and a diesel particulate filter, is analyzed.

The DOC is still the only catalyst technology that demonstrates the required robustness and durability and has been commercially established in a number of light- and heavy-duty applications. At sufficiently high exhaust temperatures, diesel oxidation catalysts can provide very effective control of hydrocarbons and CO emissions, with reduction efficiencies in the order of 90%.

The DPF traps 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, diesel filter systems have to provide a way of removing particulates from the filter periodically to restore its soot collection capacity (filter regeneration). This is attained by oxidizing the collected particulates at high temperatures to gaseous products, primarily CO_2. To achieve this high temperature in the DPF, the engine operates on rich conditions, which means that there is more fuel injected than necessary for a complete combustion. Due to the incomplete combustion, the amount of CO and hydrocarbons increases in the exhaust gas, which oxidize in the DOC and thereby heat the gas. This raises the temperature of the DPF itself and the oxidation of the particulates starts.

All here analyzed results have been simulated using SWIFT [1], AVL's computational fluid dynamics (CFD) software. These CFD simulations generate large amounts of data, including many different flow attributes like velocity, pressure, turbulence kinetic energy, temperature, etc. For the diesel exhaust system additionally a catalyst module is used, which simulates the emission reductions occuring in the DOC as well as the soot oxidation in the DPF. With this module the following data is available: the mass fractions of CO, CO_2, H_2O, C_3H_6, O_2, N_2, NO, and NO_2, and the mass of soot. Note that for simplicity of simulations, C_3H_6 is taken to represent all hydrocarbons in the fuel. Furthermore, the heat transfer is calculated and, as a result, the temperature of the solid part of the DPF is available per cell.

2.2. Application Questions

During real testing of a diesel exhaust system only some questions can be answered, like the degree of conversion of pollutants and the temperature of the exhaust gases. But it is hard to analyze why something is happening. Therefore, computational simulation can help to find more answers.

The main goal of the regeneration of the DPF is a complete oxidation of soot. As during the regeneration phase of the DPF fuel consumption and emission of pollutants increases, the duration of this phase should be as short as possible. Therefore, the oxidation should be fast, which usually causes high thermal stresses, especially in the DPF.

From this goal (complete oxidation of soot) and the requirement of short regenertion times, the following application questions are of high interest to engineers developing such DPF systems:

1. *Does the soot oxidize completely?*
   And if not, what are the reasons?
2. *Where and how fast does the soot oxidize?*
   What influences the actual behaviour of the oxidation?
3. *How high is the thermal stress of the DPF?*
   Where do high thermal stresses occur and when?

To get an overview about the behaviour of the DPF during the regeneration phase, different boundary conditions for the simulation are chosen. We here present two different cases with differing boundary conditions, for the second case thereby a 30% lower mass fraction of CO and C_3H_6 has been defined for the inlet boundary.

Figure 3: *Visualization of soot mass in the DPF after 30 sec. (middle) and after 50 sec. (right). Histograms (left) and 3D views for both cases are shown.*

3. SimVis – Interactive Multi-View Visualization of Data from Flow Simulation

In previous work we presented a multiple views system for visualization of simulation results, called *SimVis* [4]. It provides users with visualization techniques that are very well suited for the analysis of data similar to the application described here. The basic idea behind SimVis is that multiple, different views are used to view and interact with the underlying multi-dimensional data. By using both, InfoVis and SciVis views, the complex structure of simulation data is explored interactively. InfoVis views are mainly used for viewing different data attributes and relations between multiple attributes, whereas SciVis views are primarily used for 3D-viewing of the spatial layout of the data. All views support the powerful concept of linking and brushing [5, 4] (see figure 5 for a typical analysis session, two scatterplots and two 3D views are used to view and interact with the data).

To allow fast and flexible exploration and analysis of large data sets, SimVis employs *focus-plus-context* (F+C) visualization [3, 6]. F+C techniques are used to show some of the data in detail, and (at the same time) the rest of the data, at a lower resolution, as context for better orientation. Visual discrimination of data in focus and context is done interactively by the user through brushing data items [2, 11] in the different InfoVis views. In addition to standard brushing, SimVis enables the use of non-binary degree-of-interest (DOI) functions for smooth brushing [5] for interactive feature specification. The specified features are assumed to be in focus during further visualization.

Besides smooth brushing, SimVis allows also the definition of complex (i.e. high-dimensional and composite) feature specifications via an XML-based *feature definition language* (FDL) [4]. The FDL enables viewing all features in all associated views, interactive changing of parameters for all feature specification parts, and storing and resuming feature specification sessions at any point of time. Thereby, also easy porting of feature specifications to other data sets for comparing results, as done in this case study, is possible.

Analyzing data with SimVis enables multiple different options of how to support the engineer in gaining more insight into large, high-dimensional data sets. First, through brushing data according to one or more data attributes and coloring data items in other InfoVis views according to the assigned DOI values, relations of different data attributes become apparent. Another task requested very often is fast feature localization in the 3D spatial layout of the data. It is especially important in the cases where engineers quickly want to analyze different behaviour patterns during debugging of simulations. SimVis allows fast feature localization by linking the DOI values for all data items also to the SciVis views. The 3D visualization is then adjusted according to the respective DOI value for each cell. We use a cell-based rendering approach, making it possible to identify each cell value separately. To achieve a F+C visualization of the flow features defined, we modulate transfer functions for colors and opacities, as well as for sizes of singular visualization objects per cell [5] (see colorplate figures for examples of such F+C visualizations in 3D views).

4. Visual Analysis of a Diesel Exhaust System

To analyze the simulation results of a regeneration phase of the DPF in a diesel exhaust system, we employ SimVis for a visual analysis of the before identified user questions (see section 2). Results of this visual analysis are now presented, and extensions which had to be added to our previously presented work [4] are discussed. For larger versions

of the images as presented here, as well as for some more results and additional animation sequences, and a longer version of this case study as technical report, please refer to *http://www.VRVis.at/vis/research/diesel-case/*.

One general extension to our previous work was the extension to handle time-dependent data, both for visualization and also for feature specification. Each view can either show data of one timestep, or accumulate the data of many successive timesteps (building a so-called *timeslab*). Another extension was added to the F+C idea. When visualizing time-dependent data, two levels of context can be distinguished: a 2^{nd} level context representing data of all timesteps which are currently not active, and the standard F+C discrimination for all data in the currently active timesteps of the chosen timeslab. The data items belonging to the 2^{nd} level context are visualized additionally using a less prominant representation, e.g. grey-colored points in a scatterplot (see figure 4).

4.1. Does the soot oxidize completely?

Complete soot oxidation is the main goal of the simulation of the DPF regeneration phase. In the case of incomplete soot oxidation, analysis of sources for this behaviour are of highest interest. To analyze the soot oxidation process, first soot mass values are investigated. Therefore, in two histograms, the soot mass values are displayed for all timesteps of both cases. A brush is applied in the upper histogram view (see figure 3), selecting all data items of case one having a soot mass value greater than 0 (selected data is colored red in SimVis InfoVis views). Note that we always show the views for the first simulation case (with more CO and C_3H_6 at the inlet) above the views of the second case for comparison of results throughout this section. To ensure correct comparison and analysis of the results of both cases, all feature specifications performed in the first case are saved and loaded again for the second case (using the FDL of SimVis), thus the same parameters for the brushes are guaranteed. In figure 3 two different timesteps are shown in the 3D views in the middle and right, where soot mass values are color-mapped in the focus parts of the above described brushed regions.

As clearly seen, in the first case there is almost no soot left after 50 seconds, but in the second case the soot does not oxidize completely, there are still fractions of the soot in the front part of the DPF after 50 seconds. To explore the reason for that, the temperature values in the DPF are investigated, as shown in Colorplate 1, left side, for 20 seconds after the simulation start. It can be seen that the temperature in the first case is much higher than in the second one. The reason for that is, that more hydrocarbons and CO oxidize in the DOC in the first case due to higher C_3H_6 and CO mass fraction at the inlet. The oxidation generates more heat, which results in a higher fluid temperature, heating the DPF and the oxidation of soot starts quicker. As a result the effectiveness of the soot oxidation increases in the rear, and the soot oxidizes completely in this region, too. To sum up the first

question, a high mass fraction of C_3H_6 and CO is needed for a complete soot oxidation.

4.2. Where and how fast does the soot oxidize?

For this question only the second case is of interest because of the incomplete soot oxidation. Due to soot oxidation, CO_2 and CO are generated in regions of high temperature. To display these areas, it is necessary to specify a complex feature. In a first scatterplot the mass fraction of CO_2 (Y-axis) is plotted against the mass fraction of CO (X-axis). A smooth brush selecting higher values of both attributes is applied (see Colorplate 1, right, upper scatterplot view). Then, in a second step, to refine the first 2D brush, a second scatterplot showing general fluid temperature (Y-axis) vs. the spatial X-coordinates of the data set (X-axis) is used. Note that the spatial X-axis of this data set correlates very well with the main flow direction through the whole data set (from left to right in figure 1). Here, high values of fluid temperature in the region of the DPF are brushed, leading to a 3D, composite brush. The two scatterplots as well as a 3D view are shown in Colorplate 1 on the right. In the 3D view, velocity values are color mapped. Visualization shows, that the CO and the CO_2 production are not symmetric.

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. As the amplitude of the gradient values changes over time, we need a method to select relativly high gradient values with respect to the maximum gradient value for each timestep, to get interesting cells for each timestep 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 timestep, for example. And here another extension of SimVis was added. We included a data derivation tool, where derived data (time-dependent) can be calculated. Standard derived data calculation in SimVis includes calculation of differences over time (central, forward or backward differences), smoothing of data over time, calculation of spatial regions exhibiting low change rates, etc. After such a derived data calculation, all of these derived data attributes can be accessed and used for visualization and analysis in any of the views of the system. We also included a tool for local (with respect to time) normalization of data attributes to cope with changing attribute value ranges.

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 Colorplate 2, 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 also shown below. Here the
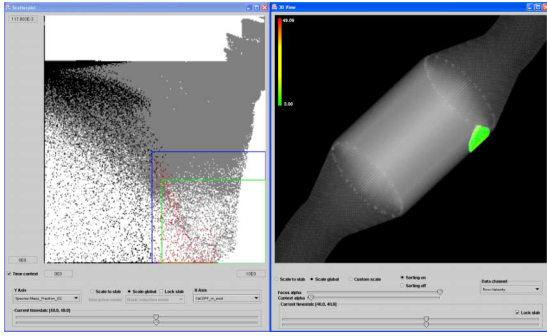
Figure 4: *Velocity values in cells with high soot mass values and low mass fraction of* O_2 *after 40 sec.*

**Figure 5:** *High solid catalyst temperature after 30 sec.*

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 timestep. For color mapping in the three 3D views, the mass soot values are taken into account.

From Colorplate 2 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 also Colorplate 1, left). The asymmetric temperature distribution originates in the asymmetric flow field, due to the bend of the geometry between the DOC and the DPF. The amplitude of the oxidation rate is rather low at the beginning, which can be seen in the lower scatterplot of Colorplate 2. The asymmetric behaviour is continued also at later timesteps (see 3D views middle to right). That means, that in the region displayed in the right most 3D view the oxidation starts later, although there seems to be enough heat for a fast oxidation. A possible reason for low oxidation rate is that there is not enough O_2 for oxidation. To check this, a scatterplot showing mass fraction values of O_2 (Y-axis) vs. the soot mass values (X-axis) is used (see figure 4). High soot mass values and low mass fraction values of O_2 are brushed, the color mapping of selected regions in the 3D view shows velocity values.

The area of high soot mass values with low mass fraction of O_2 is the same as the area where the oxidation starts later. So the low mass fraction of O_2 causes the later soot oxidation start. The reason for less O_2 is, that the velocity is low and thus only little oxygen can flow into this region. To answer the second question, the soot oxidation is executed asymmetrically, because of the asymmetric fluid flow.

4.3. How high is the thermal stress of the DPF?

Here we again compare the two cases as described in section 2. One measure for thermal stress is the temperature amplitude of the DPF. For analyzing this, a scatterplot with the solid catalyst temperature (Y-axis) and the spatial X-coordinates of the data set (X-axis) is used. High values

of solid catalyst temperature in the region of the DPF are brushed (see figure 5). As can be seen, in the first case the area with high solid catalyst temperature is bigger and the maximum value is higher than in the second case, thus the thermal stresses in the DPF are higher in the first case.

Additionally thermal stresses are caused by temperature fluctuation. Thus we search for regions with high heating or high cooling properties. To analyze this we open a new scatterplot, showing differences of solid catalyst temperature values (Y-axis) against temperature values of the flow (X-axis). Here we define two new brushes: high positive difference values of solid catalyst temperature give high heating factors, whereas high negative difference values of solid catalyst temperature represent strong cooling capacities. In the 3D views of Colorplate 3, the green-colored regions represent areas exhibiting cooling gradients, whereas the red-colored regions show areas with heating gradients. As can be seen from the two shown timesteps, gradients of both, heating and cooling, are higher for the first case, thus thermal stresses are also higher in this case.

The analysis of this question has shown, that the thermal stresses in the first case are higher, due to higher heating gradients and higher solid catalyst temperatures in the DOC, especially in the beginning of the regeneration phase.

Comparison

Analyzing the presented questions with common visualization software would be quite tedious and often even impossible, because of the three-dimensional behaviour of the fluid flow, as well as the complexity of needed data selection tools. Typical common visualization tools usually only provide 2D cuts, or displaying data attributes on surfaces, whereas SimVis allows cell-based rendering in the 3D views, showing fully three-dimensional impressions of data attribute distributions. Additionally, interactive analysis through brushing (with immediate feedback in the 3D view) and working on multiple dimensions simultaneously,

improve and speed up the analysis process compared to common visualization tools. Another advantage of SimVis is, that it is easy to define gradients or calculate other derived data attributes and normalize them afterwards.

Performance Issues

The presented case study was carried out on a system consisting of the following components: The hardware used was a standard PC-based system (Intel Pentium4, 2.8GHz, 2GB RAM, 2x 40GB HD, RAID 0 (Striped Set)) with a NVidia GeForceFX 5900 graphics card. The software is our Java-C++ based SimVis visualization environment.

The two data sets investigated during this case study each consisted of approximately 260,000 cells, where 37 attributes per cell are available. The temporal dimension of the results stored from the simulation is 0 to 80 seconds after the start of the DPF clearing process started. We used 20 different timesteps out of this interval for visualization and analysis. With these settings, loading all available timesteps into a scatterplot, for example, requires drawing approximately 5.2 million points. Fortunately, due to algorithmic optimizations, and a simple but effective data handling framework, we are able to handle analysis sessions with multiple views for data sets even a magnitude larger than the ones presented here, while still providing full interactivity.

5. Conclusions

In this paper, we are demonstrating how interactive focus+context visualization of multi-dimensional and time-dependent flow data effectively supports the exploration and analysis of results from computational flow simulation (in our case the gas flow through a diesel exhaust system). Views, which are usually known from information visualization, such as scatterplots, histograms, etc., are used to visualize the multi-dimensional attribute space of flow simulation data and to provide insight in the distribution and intra-relation of data items with respect to certain attributes of special interest. Once the user detects interesting structures in the data through exploration in multiple views, he or she is able to interactively brush data subsets of special interest directly in the InfoVis views. As one result, all the different views are updated instantaneously to jointly reflect the user selection based on visual linking – color is used to consistently highlight the selected data subsets in all views.

In the course of our case study, we are showing how different features of this visualization approach are used to effectively bring forward the exploration and analysis of a diesel exhaust system. Complex application questions are investigated, providing new insight into the data, which would not have been possible with standard technology (at least not as efficiently).The here described system (SimVis) is especially useful when multiple data attributes are to be considered during the analysis of a specific question (e.g., why oxidation of soot is not symmetric). To actually work out this case study, SimVis needed to be extended by several means.

Lessons learned from this case study are that interactive visualization with the opportunity to interactively drill down into certain aspects of the data (through brushing) effectively supports the investigation of simulation data. Also the clear formulation of features in terms of the feature definition language proves to be very useful, not at the least when it comes to comparing different data sets side by side and porting of feature specifications from one data set to the other becomes handy. In comparison to standard analysis tools for simulation data, the visual linking between the visualization of attribute space (through InfoVis views) and physical space (through 3D SciVis) also proves to be very useful.
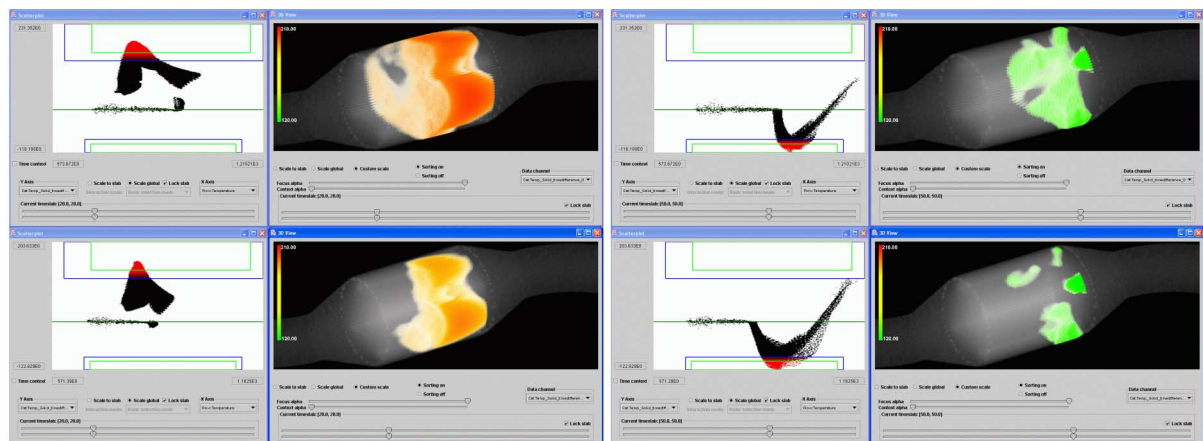
References

[1] AVL SWIFT Manual Version 3.2, July 2003. 2

[2] R. Becker and W. Cleveland. Brushing scatterplots. *Technometrics*, 29(2):127–142, 1987. 3

[3] S. Card, J. MacKinlay, and B. Shneiderman. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, 1998. 3

[4] 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, Grenoble, France, May 2003. 1, 3

[5] 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. 1, 3

[6] G. W. Furnas. Generalized fisheye views. In *Proc. of the ACM CHI '86 Conf. on Human Factors in Computing Systems*, pages 16–23, 1986. 3

[7] D. L. Gresh, B. E. Rogowitz, R. L. Winslow, D. F. Scollan, and C. K. Yung. WEAVE: A system for visually linking 3-D and statistical visualizations, applied to cardiac simulation and measurement data. In *Proc. of IEEE Visualization 2000*, pages 489–492. 1

[8] C. Henze. Feature detection in linked derived spaces. In *Proc. of IEEE Visualization '98*, pages 87–94. 1

[9] Online information service on clean diesel engines and diesel emissions. See URL http://www.dieselnet.com/. 2

[10] F.H. Post, B. Vrolijk, H. Hauser, R.S. Laramee, and H. Doleisch. Feature extraction and visualization of flow fields. In *Eurographics State of the Art Reports*, 2002. 1

[11] M. O. Ward. XmdvTool: Integrating multiple methods for visualizing multivariate data. In *Proc. of IEEE Visualization '94*, pages 326–336. 3

Colorplate 1: *left:* temperature values in the DPF are color mapped to the features (regions of non-zero soot mass) after 20 seconds, visualizations from case 1 (upper view) and case 2 (lower view) are compared; *right:* velocity in cells with high CO and CO_2 mass fraction and high temp. after 35 seconds for data from case 2.

Colorplate 2: Soot mass in cells with high oxidation rate after 15, 30, and 40 sec. (from left to right)

Colorplate 3: Time differences for solid catalyst temperature after 20 sec. (left) and after 50 sec. (right)