Opening the Can of Worms: An Exploration Tool for Vortical Flows

Simon Stegmaier*
Institute for Visualization and Interactive Systems

Ulrich Rist†
Institute for Aerodynamics and Gasdynamics

Thomas Ertl*
Institute for Visualization and Interactive Systems

University of Stuttgart, Stuttgart, Germany

ABSTRACT

Gaining a comprehensive understanding of turbulent flows still poses one of the great challenges in fluid dynamics. A well-established approach to advance this research is the analysis of the vortex structures contained in the flow. In order to be able to perform this analysis efficiently, supporting visualization tools with clearly defined requirements are needed. In this paper, we present a visualization system which matches these requirements to a large extent. The system consists of two components. The first component analyzes the flow by means of a novel combination of vortex core line detection and the λ_2 method. The second component is a vortex browser which allows for an interactive exploration and manipulation of the vortices detected and separated during the first phase.

Our system improves the reliability and applicability of existing vortex detection methods and allows for a more efficient study of vortical flows which is demonstrated in an evaluation performed by experts.

CR Categories: I.3.3 [Computer Graphics]: Interactive Rendering—Flow Visualization;

Keywords: Flow Features, Vortex Detection, Interactive Manipulation, 3D Vector Field Visualization

1 INTRODUCTION

Numerical flow simulation using supercomputers and experimental techniques like PIV (Particle Image Velocimetry) and LDA (Laser Doppler Anemometry) are valuable instruments for the development of new products in the car manufacturing and aerospace industry and other research areas where understanding of gaseous or liquid flows is required. In general, the value of these numerical and experimental flow simulations depends on the resolution of the computational grid and the number of velocity vectors that could be measured, respectively. Accordingly, much effort has been put into improving these simulation and measurement techniques, thereby pushing the size of the resulting data sets into regions where human cognition no longer suffices as a data analysis tool for the raw data. Some sort of flow visualization is, therefore, required for converting the raw data into a more meaningful representation. In the following, the term *flow visualization* is used synonymously to numerical flow visualization using streamlines, isosurfaces, etc. in contrast to physical flow visualization, the process of visualizing physical flows of gases and liquids.

Visualizing the original untransformed data, however, is often insufficient for obtaining the desired insight. Fig. 1, bottom, shows a turbulent flow visualized with the commercial flow visualization package *PowerVIZ* [4]. For this data set, streamlines become almost useless due to a dominant velocity component along the longitudinal axis and the researcher is forced to analyze the flow based on the limited information revealed by the cutting planes and the isosurface computed on velocity magnitude.

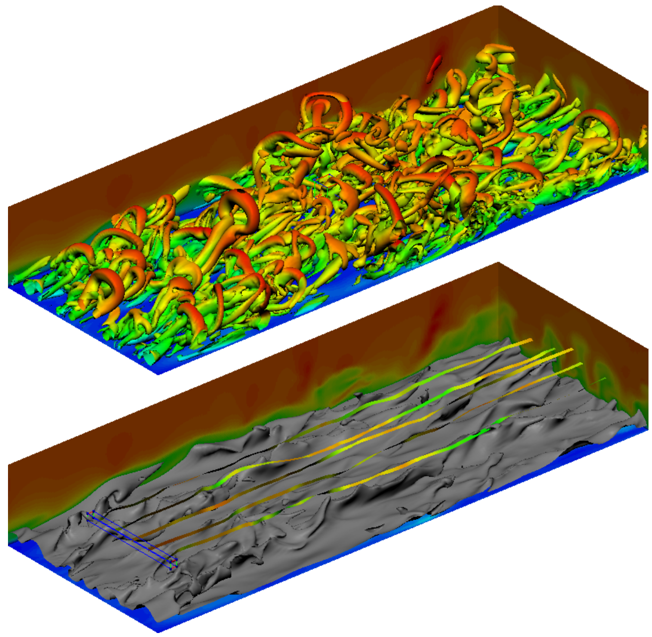

Figure 1: Advanced flow visualization by λ_2 vortices (colored with velocity magnitude, top) and conventional flow visualization using isosurfaces, stream ribbons, and cutting planes. Despite the additional insight gained from vortex visualization, analyzing the flow is nevertheless difficult since the structures are neither clearly separated nor is it clear how they interact and evolve.

As a result, recent flow visualization techniques often resort to visualizing meta data, i.e., data that is obtained by transforming the raw data—here: vector field—into a different and often more compact representation. Critical points and flow topology [7, 6], separation lines and surfaces [11], shock waves [12, 15], and vortices [9, 19, 22] are some well-established examples of this category of advanced flow visualization which use data reduction as a means for eased flow field analysis.

Of special importance is the detection and visualization of vortices since they carry most of the energy of vortical flows and, thus, contribute considerably to the flow evolution in future time steps. The method that is currently considered to be most effective for detecting vortical structures in incompressible flows (maybe except for flows with strong axial stretching, see [25]) is the λ_2 method proposed by Jeong and Hussain [9]. This λ_2 method can also be im-

*e-mail: (stegmaier|ertl)@vis.uni-stuttgart.de
†e-mail: rist@iag.uni-stuttgart.de

plemented very efficiently—even on programmable graphics hardware as was recently demonstrated [23]—and is, thus, favored by many researchers working on fluid dynamics.

Unfortunately, in contrast to methods for detecting vortex cores which automatically obtain a segmented visualization, the λ_2 method only produces a scalar field usually visualized by an isosurface. As a result, analyzing visualizations of λ_2 vortices is often more difficult than analyzing the results of less reliable but automatically segmenting core-detection approaches. In fact, for highly-turbulent flows a "can of worms" is obtained which hardly reveals the important information sought by fluid dynamics engineers (Fig. 1, top).

Thus, what is actually needed is an effective method for separating λ_2 vortices and, furthermore, a tool—a *vortex browser*—which allows for an interactive exploration of the flow by manipulating and examining individual vortices. In this paper we present such a tool and an algorithm for vortex separation to accomplish this task. The development of both the algorithm and the custom exploration tool was driven by actual requirements of fluid dynamics engineers cooperating with computer scientists in a nation-wide project dedicated to the fundamental research and comparison of experimental and numerical methods for the analysis of turbulent flows. The resulting tool was successfully evaluated by this group of experts.

Our exposition is organized according to the components of the system. Sec. 2 gives related work; Sec. 3 elaborates on the separation algorithm, followed by Sec. 4 which gives a description of the vortex browser. An evaluation of the system demonstrating its effectiveness is given in Sec. 5. The paper concludes in Sec. 6.

2 Related Work

Most similar to what is accomplished in this work is the visiometric approach to visualization proposed by Silver and Zabusky [21] and Fernandez et al. [5]. In both works, systems are presented to visualize, recognize, identify, classify, and automatically track observable flow field features to give researchers new insight into what is happening in the data and to allow for the formulation or verification of theories describing the visualized process. Object identification in these papers is done by a region-growing approach, e.g., thresholding on vorticity magnitude. This approach was later applied to a λ_2 scalar field by Rist et al. [17] to identify and separate vortices.

None of these approaches take advantage of domain knowledge and all of them exclusively depend on scalar values—for the clustering algorithm it makes no difference whether the scalar field has been derived from MRI images or from the λ_2 method to detect vortices. In any case, the same structures will be extracted. From a researcher's point of view it makes a difference since structures in the flow field might actually be connected while medical structures might be hindered from merging by tissue. Thus, the accuracy of feature extraction may be reduced if domain knowledge is not taken into account.

However, while identifying features and providing means to understand these features are both important, the effectiveness of the latter depends on the sophistication of the former. In contrast to Silver and Fernandez who employ rather basic object identification algorithms and who lay their foci on the understanding issue, this paper strives to improve the feature identification using a combination of the λ_2 method (to which the visiometric approach is only insufficiently applicable) and the algorithm proposed by Banks and Singer [2, 22]. In addition, we provide a basic accompanying flow exploration tool for illustration purposes. And this differentiates our work from the plethora of publications on vortex identification.

3 Vortex Separation

3.1 Basic Idea

As motivated in Sec. 1, the λ_2 method is popular due to its reliability in detecting vortices and the simplicity of the operations involved in the computation. This section will elaborate on this issue and point out the strengths and weaknesses of the algorithm. For this purpose, it is helpful to classify the λ_2 method according to the criteria defined by Jiang et al. [10]. The taxonomies used are:

1. The number of grid cells involved in the computations required for detecting a single vortex.
2. The algorithm's ability to work reliably even when a constant delta is added to the velocities of the original vector field, also referred to as *Galilean invariance*.
3. The vortex definition underlying the detection algorithm.

Whether a detection method is well-accepted or not depends largely on how the algorithm matches the above criteria. For example, an algorithm that needs to examine grid points distributed all over the volume in order to detect a single vortex will exhibit a large memory foot print and thus be computationally more expensive than algorithms which only have to examine neighboring grid cells. And an algorithm which will no longer work when a delta is added to the velocities will be unable to detect vortices in time-varying data where swirling motion is only seen from within a moving reference frame. Finally, an algorithm which only detects vortical regions cannot generate a flow field representation as compact as that produced by a method returning a list of vortex cores and thus might have to be declined when visualizing large data sets.

To classify the λ_2 method, one needs to examine the algorithm in detail. Let

$$\mathbf{u}(\mathbf{x}) = \begin{pmatrix} u_1 \\ u_2 \\ u_3 \end{pmatrix}$$

be a 3D velocity field from which vortices are to be extracted. For every grid point of this vector field the Jacobian or velocity gradient tensor $J = \nabla\mathbf{u}$ is then computed and decomposed into a symmetric part $S = (J + J^T)/2$ and an antisymmetric part $\Omega = (J - J^T)/2$:

$$S_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right), \qquad \Omega_{ij} = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i}\right).$$

Seen from a physical point of view S denotes the strain-rate tensor and Ω the rotation tensor. Both S and Ω are squared and added to obtain a new 3×3 matrix. This matrix is real and symmetric and thus has exactly three real eigenvalues. These eigenvalues are computed and sorted in decreasing order: $\lambda_1 \geq \lambda_2 \geq \lambda_3$. A vortex is then

---

foreach remaining seed point $\mathbf{p}_0$
  if $\mathbf{p}_0$ is not in any previous vortex
    while the vortex skeleton continues
      (1) determine vorticity ω_i at position $\mathbf{p}_i$
      (2) integrate vorticity ω_i to predict new position $\overline{\mathbf{p}}_{i+1}$
      (3) determine vorticity $\overline{\omega}_{i+1}$ at prediction $\overline{\mathbf{p}}_{i+1}$
      (4) $\mathbf{p}_{i+1}$ is the point of minimum pressure
        in the plane $P \perp \overline{\omega}_{i+1}$ (correction step)
      (5) $i \leftarrow i+1$

Figure 2: The predictor-corrector vortex detection algorithm proposed by Banks and Singer [2, 22].

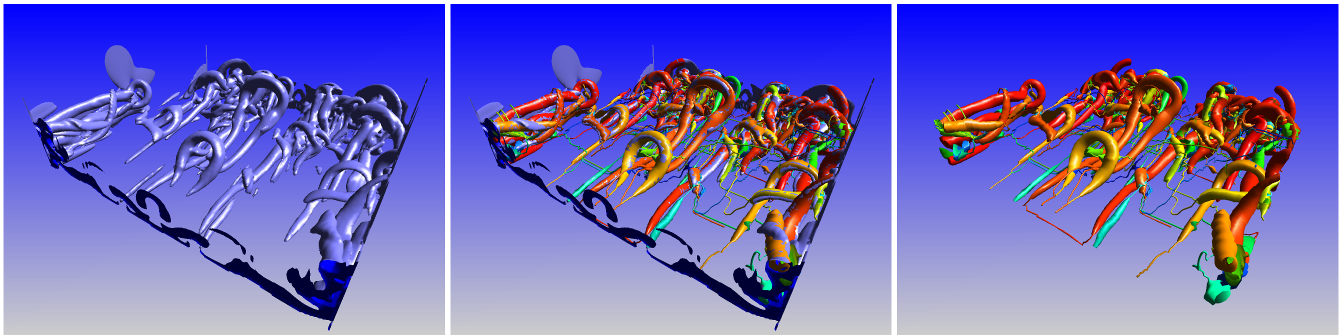

Figure 3: Comparison of λ_2 isosurface and vortex structures detected and separated with the proposed approach. The left image shows the λ_2 isosurface, the right image the separated vortices. As seen in the middle image (showing both the isosurface and the separated vortices), significant differences between the two visualizations are only seen in the boundary region. The shape of these structures, however, indicates that they correctly have not been classified as vortices.

defined as a connected region where two of the eigenvalues are negative. For visualization—as the name of the method implies—the second largest eigenvalue is picked. Since λ_2 is a scalar and since for every grid point a corresponding λ_2 value can be computed, the λ_2 method transforms the original vector field into a scalar volume which can then be visualized by any volume visualization technique, most commonly isosurfaces (Fig. 1, top). The more negative the chosen isovalue, the stronger the vortices that will be seen. A zero isovalue captures all vortices contained in the data set.

Obviously, the computation of λ_2 values requires only the Jacobian which can be easily computed by central differences from the six direct neighbors[1]; thus, the algorithm is local and computationally cheap. Furthermore, velocity is not directly used to compute the scalar field but rather its derivative (the velocity gradient tensor) so constant deltas when added to the vector field cancel each other, i.e., the method is also Galilean invariant. And finally, the output of the λ_2 method is a scalar field rather than a list of core lines so that the method sees vortices as regions.

The latter is both good and bad: on the one hand a vortex has no defined extent so discarding the concept of a vortex core is valid; but on the other hand, vortical regions smoothly merge into each other such that vortices in the proximity of other vortices may be misleadingly seen as a single structure. It becomes clear that the major drawback of the λ_2 method is its inability to separate vortices. And it also becomes clear that there is a well-justified demand for an improved vortex detection method combining the benefits of the λ_2 method with the benefits of line-oriented vortex detection methods.

Our approach to alleviate this drawback is to combine the λ_2 method with a core line detector. We chose the predictor-corrector approach proposed by Banks and Singer [2, 22]. This method is based on the assumption that pressure minima at vortex cores define seed points for tracing core lines and that vorticity defines the corresponding vortex core direction in these locations. Accordingly, by integrating along the vorticity vector a prediction for the next vertex along the core can be made and the guess corrected by minimizing pressure in the plane perpendicular to the vorticity vector at the predicted position (Fig. 2). In order to obtain tubular structures, cross-sections can be determined by radially sampling the pressure field in the planes perpendicular to the vortex core. Overall, this vortex detection approach is robust (due to its self-correcting nature) and very intuitive. But it also has a drawback: It does not work at low Reynolds numbers [9].

On the contrary, the λ_2 criterion lacks the capability to produce a list of separated vortices but it is able to capture the pressure minimum in a plane perpendicular to the vortex axis for both low and high Reynolds number flows. The λ_2 method, therefore, perfectly complements the algorithm proposed by Banks and Singer if pressure is being replaced by the λ_2 scalar field. And since pressure is hard to compute from a velocity field—or even unavailable at all in many experimentally obtained data sets—while λ_2 is not, this combination not only improves the reliability of existing vortex detection, it also widens the applicability of the methods.

3.2 Initialization

The algorithm proceeds by growing core lines from a set of selected seed points. In order to obtain a countable number of seed points, these initial seed points must lie on the grid; thus, the initial set contains as many seed points as there are grid points—for typical flow data sets millions or even tens of millions elements. Checking all of them would be extremely expensive. However, with the Banks/Singer method used in conjunction with λ_2 values, the initial set can be significantly reduced. This reduction defines the initialization phase.

There are three classes of grid points: Grid points with positive λ_2 values, grid points with negative λ_2 values which are not local minima, and grid points to which none of the previous applies.

The first group can readily be discarded since positive λ_2 values indicate non-vortical regions so no vortices will be found there. These grid points are assigned to the class `NO_VORTEX`.

Grid points which do *not* define local minima can also be discarded since they will be processed when tracing the core (Sec. 3.4). They are assigned to the class `UNDECIDED`.

Thus, what remains are grid points with negative λ_2 values defining local minima. Whether these grid points will serve as starting points for vortex cores is unclear since again they might be eliminated if enclosed in a vortex tube. These grid points, therefore, define the class `POTENTIAL_SEED`.

The $229\times116\times250$ data set shown in Fig. 3 has a total of 6,641,000 grid points. Of these, only 1,335 qualify as potential seed points. 273,571 grid points fall into the class `UNDECIDED`. The remaining vertices—95.8 percent—can be discarded.

3.3 Growing the Skeleton

The process of growing the vortex core line is almost identical to the original implementation by Banks and Singer and differs only in implementation details. The description is, therefore, kept short.

For each core line, a seed point has to be picked first. We choose the seed point with the smallest λ_2 value. Since more negative λ_2 values indicate stronger swirling motion this strategy in general results in stronger and larger vortices and will immediately generate

[1] We assume uniform grids for our discussion.

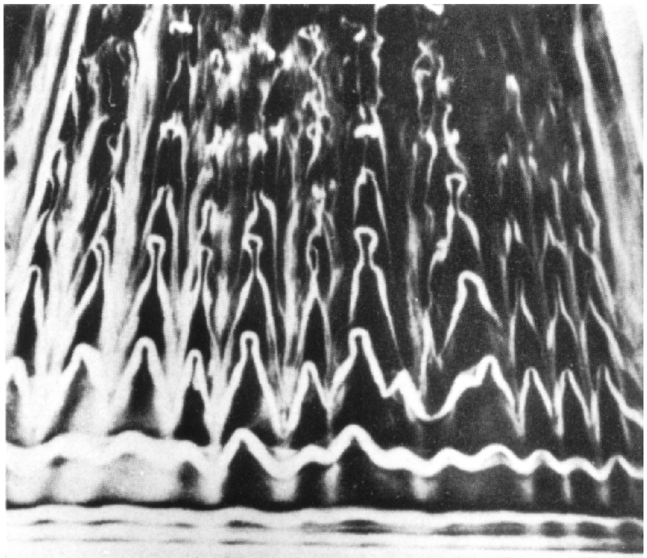

Figure 4: Photograph of actual K-type transition with physical vortex visualization using smoke inserted into the boundary layer [16]. In the upstream region of the flow near the trip wire (bottom), the Λ-vortices with the characteristic Ω-shaped heads are clearly visible.

a list of vortices sorted by strength. The initial seed point is then refined by minimizing λ_2 on the plane defined by the vorticity vector at the original seed point. Starting from the (refined) seed point, the skeleton is then grown in both directions.

We first take a step along the vorticity vector (or its negative, depending on the traced direction). For the new position we have to check, first, whether the position lies inside any of the vortex tubes already found and, second, whether the position lies inside the vortex tube currently built. In either case, the skeleton is terminated and the vortex just found is added to the vortex list. Similar to the Banks/Singer method we next determine the core direction at the predicted position and again minimize the λ_2 value in the plane defined by this direction. Since in general the new position will not lie on a grid point, the quality of the structure extraction is significantly influenced by the quality of the interpolated λ_2 values and vorticity vectors. Therefore, a four-point Lagrange interpolation is used for this purpose. The λ_2 minimization problem in turn is solved with a *direct search* approach [8] since this relieves us from the burden of having to compute gradients using computationally expensive Lagrange interpolations.

If the angle between the predicted vortex core direction and the direction at the λ_2 minimum does not exceed a given threshold the correction is accepted. In either case, the vortex cross section at the core line vertex is then determined by radially sending out rays lying in the plane perpendicular to the vortex core and sampling the λ_2 volume until the user-specified λ_2 isovalue is exceeded. As proposed by Banks and Singer the distances to the surface points are stored in a radii table for efficient encoding. Since this table represents a periodic function it can be approximated by a Fourier series which is an efficient way for storing the cross sections without loosing the ability to reconstruct them to any desired accuracy for visualization.

If the cross section area falls below a given threshold (very close to zero) the vortex tube is considered to be closed and the skeleton growth terminates. The growth is also terminated if the vortex core length exceeds a pre-defined value to enforce core line termination where the core has run into a spiral.

3.4 Seed Point Elimination

A vortex core will rarely directly hit a grid point. It is, therefore, necessary to eliminate potential seed points in a post-processing step once a new vortex has been detected. This is accomplished by iterating over the individual slabs composing the vortex tube and performing sign tests of potential seed points with respect to the triangles comprising the slab surface. By computing the slab bounding box prior to performing the test, the number of potential seed points requiring testing is reduced.

3.5 Results

Fig. 3, right, shows the result of the vortex detection applied to a DNS data set of K-type transition experiments [1]. For comparison, Fig. 4 shows a photograph of this kind of transition[2]. The photo was shot during an experiment in which smoke was inserted into the boundary layer to physically visualize vortices [16].

[2] Actual measurements of physical flows take several months or even years. Experimental data of a flow recorded under conditions equivalent to those used for the numerical simulation that provided the data for our system evaluation are not yet available.

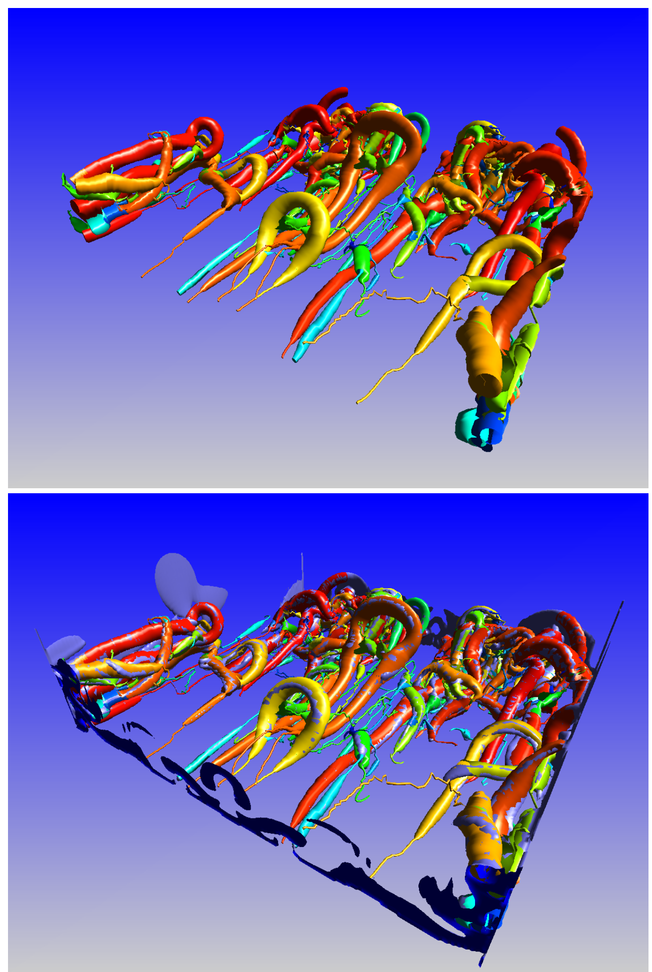

Figure 5: Visualizations obtained using an eigenvector-based estimation of the core line direction and λ_2 values for determining core vertices. The lower image again shows a comparison of the λ_2 isosurface and the vortex structures detected with the eigenvector approach.

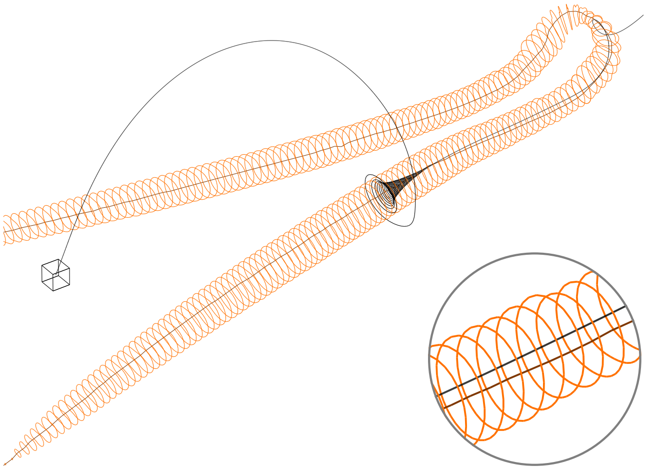

Figure 6: Particle trace computed on the velocity field. While in general the trace encloses the vortex core, in some regions it travels parallel to the core line (see close-ups).

The computations for this 229×116×250 data set took 138 seconds on a PC equipped with an AMD 1.2 GHz processor and 512 MB memory. The program detected a total of 331 vortices. Red color was assigned to the vortex structure found first which—due to the seed point selection looking for the most negative λ_2 value—is also the strongest vortex. The weaker a vortex, the greater the hue value of the respective color in HSV space.

When comparing the result to a standard λ_2 isosurface computed for the same isovalue that was chosen for determining cross sections, it becomes apparent that the visualizations match very closely. There are only two differences. First, some structures near the volume boundary are missing in the separated visualization and, second, some additional very fine structures evolve from the vortex tails.

The structures referred to in the first case are caused by inaccurate derivatives at the volume boundary. However, as the shape of these structures already implies, it is very improbable that in these regions "instantaneous streamlines mapped onto a plane normal to the vortex core exhibit a roughly circular or spiral pattern", as one popular definition of vortices demands [18]. Therefore, the structures should not be classified as vortices. And they are not—as is seen in the image—so errors made during the λ_2 computation do not find its way into the final visualization.

In the second case, the additional fine structures are vortex core lines that have been obtained by setting the cross section area to a larger value, allowing for skeleton growth also beyond vortex tails. This is another benefit when applying the Banks/Singer vortex detection method to a λ_2 scalar field since it enables the researcher to identify connected structures hitherto visualized as separate vortices.

If a higher accuracy is required, the system can be configured to better approximate the vortex tube cross sections and, accordingly, to better match the isosurface, thereby eliminating isolated bluish spots in the comparative visualizations showing both the isosurface and the vortex tubes. For the given images, 10 rays were used for sampling the λ_2 field, five coefficients where saved for the Fourier series, and 16 samples were used for displaying the vortex geometry.

Sujudi and Haimes propose to use the eigenvector corresponding to the real eigenvalue of the Jacobian instead of vorticity for estimating core directions [24]. The eigenvector, however, is not uniquely defined with respect to sign and thus cannot be used as-is

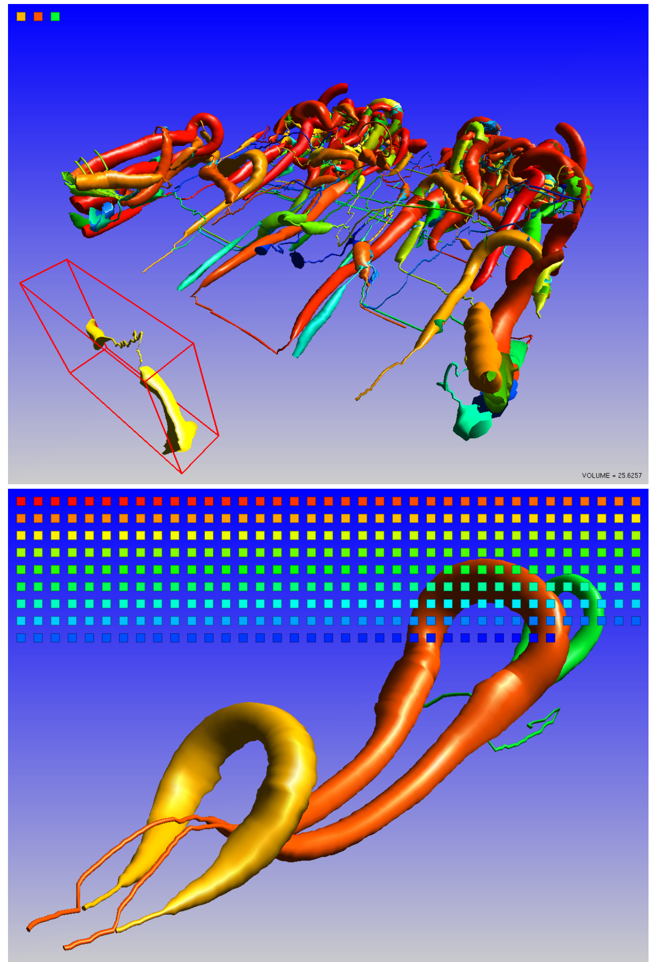

Figure 7: Manipulating the vortex set. Top: Hiding the three central Λ-vortices and selecting a single vortex for closer examination. Bottom: Inverting the selection. Using the colored buttons individual hidden vortices can be temporarily or permanently re-inserted.

for following the vortex core. We propose a simple solution to this problem in that we re-orient the eigenvector based on the sign of the dot product of the eigenvector and the vorticity vector. Fig. 5 shows visualizations obtained with this approach. The result hardly differs from what is obtained with a vorticity-based estimator and there is only a small discrepancy in the number of detected vortices: 309 vs. 331. In consideration of the high complexity involved in the computation of eigenvectors, using vorticity still seems a valid alternative.

To further verify the quality of the presented approach, Fig. 6 shows a single Λ-vortex extracted from the K-type transition data. For the image, a particle probe was placed near the detected vortex core and a trace computed by integrating the velocity field. Ideally, the resulting trace should exactly enclose and follow the core line detected during the vortex segmentation process. In the example, this requirement is essentially fulfilled and inaccuracies are only seen on a short core segment where the particle travels parallel to the vortex core (see close-up). The particle leaves the core through the appendix shown in the upper right corner near the Ω-shaped vortex segment. This is correct behavior and must be accounted to the strong velocity component along the longitudinal axis. As for the interpolations involved in tracing the core line and determining the vortex tube cross sections, a four-point Lagrange interpolation

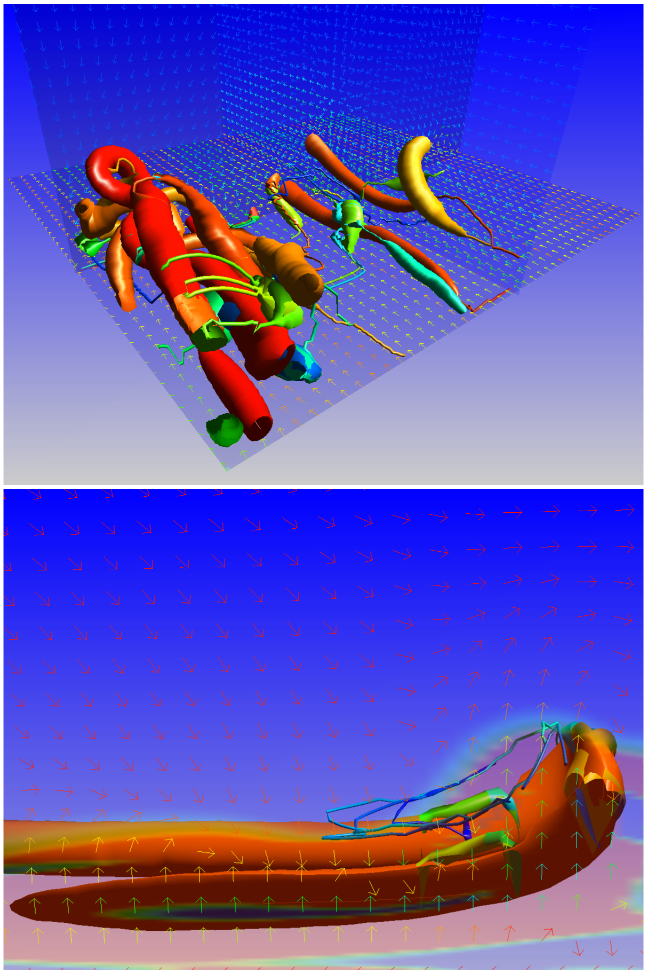

Figure 8: Use of slices. Slices can be used to display vector plots of the velocity field and to clip unwanted regions (top) and to obtain dense representations of associated scalar fields like velocity or vorticity magnitude, λ_2 values, or—as shown in the image—shear stress.

(in contrast to a less costly trilinear interpolation) is required in order to obtain accurate results.

4 The Vortex Browser

The vortex set shown in Fig. 3 comprises 331 vortices. This includes weak and strong vortices, vortices with short and long cores, and vortices which cover small and large regions of the simulation volume. Obviously, some filtering must be applied before the data can be thoroughly and efficiently analyzed. This is accomplished with the second component of our system, the vortex browser.

4.1 Aims

The research of fluid dynamics engineers is directed towards an increased understanding of fluid dynamics. This understanding refers to both fundamental questions as well as to specific applications where the flow around an actual device (automobile, airplane wing, etc.) has to be examined in order to, e.g., improve efficiency or to reduce noise generation. To gain this insight, fluid dynamics engineers like to study the interactions of different fluid flow entities and to view and take into consideration meta data. Typical questions arising in practice are: What is the volume of a certain structure? What is its swirling strength, circulation, and vorticity? How about induced velocity, vortex stretching, and dissipation?

4.2 Vortex Manipulation

None of the above questions can be answered when considering the vortex set in its entirety; thus, means are required to manipulate the vortex. For this purpose, the individual vortices are inserted into a custom scenegraph which allows for simple manipulations like picking and transformations. This, however, is insufficient since researchers either want to see a certain vortex or do *not* want to see certain vortices being irrelevant for the special application or obscuring the region of interest. Our system, therefore, provides hiding and inversion functionality (Fig. 7). In the top image, the central Λ-vortices have been hidden. Another vortex—one that would have been visualized as two separate structures with standard λ_2 isosurfaces—has been dragged out of the flow and is shown (with a selection box) in the front. Its volume is shown in the lower right corner. In the bottom image the selection has been inverted in order to allow for a closer examination of exactly the hidden vortices. In either case, for each hidden vortex a button colored with the respective vortex color is inserted. By moving the mouse cursor over any of these buttons, the corresponding vortex is shown at its original location. By clicking the buttons, the corresponding hidden vortex is permanently re-inserted into the visualized vortex set. Of course, the buttons can be disabled if desired.

Manually defining the vortex set may become tedious if there are dozens or even several hundreds of vortices. However, since the flow dynamics is dominated by the largest vortices, this task can be partially automated. In the system, this functionality is provided by means of a length filter which can be used to eliminate smaller and weaker vortices.

4.3 Visualization Techniques

From an information visualization point of view the above functionality defines the focus. The vortices, however, have been extracted from a flow field and this context should not be lost. One software feature to provide context information has already been shown in Fig. 6. The user can freely position a particle probe inside the simulation volume and integrate both along the velocity field and the

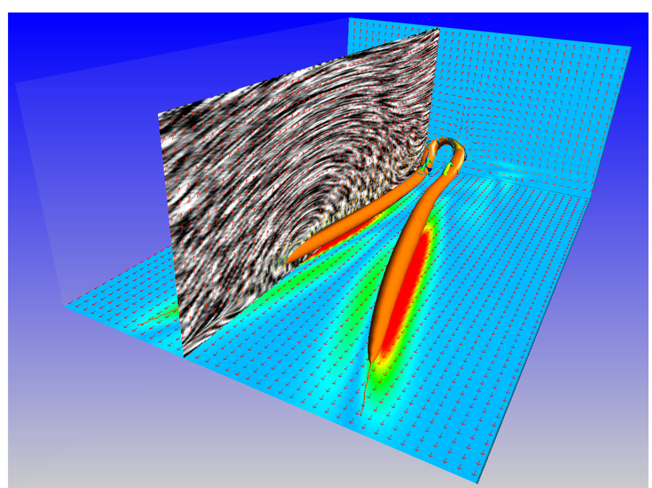

Figure 9: Combined visualization of separated vortices and slices showing both scalar field data (shear stress) and a LIC representation of the velocity field.

vorticity field.

Regarding the former one, it should be noted that no real particle will actually follow the present streamlines because the velocity field changes with time. However, the deviation of instantaneous streamlines from path lines of real particles depends on the temporal rate of change of the flow. The latter is smaller near the "legs" of the Λ-vortex and higher where many "tangled" vortices interact. Therefore, the interactive streamline tool is still useful in the first part of the data field to get a qualitative understanding of the flow but dangerous or even misleading in the second. Thus, a more reliable probe like the vorticity line probe is needed.

So-called "vorticity lines", i.e. field lines of the instantaneous vorticity field, are (by definition) everywhere parallel to the vorticity vector. This means that they should follow the vortex axis and shear layers everywhere in the flow. However, using vorticity lines alone as a means of visualization would not be advantageous, because they have the disadvantage that no distinction is made between vortices and shear layers, and that regions with small vorticity cannot be distinguished from regions with large vorticity. On the other hand, when used in an interactive manner together with the extracted vortices, they provide a useful tool to check the physical accuracy of our vortex identification and extraction. In contrast to field lines of the instantaneous velocity field (streamlines) vorticity lines are Galilean invariant such that they can be used with more confidence in an unsteady flow situation, as the one discussed here. Therefore, using the particle probe with the velocity field can, at best, show regions of strong attraction (like the strong vortex tubes) in a qualitative manner in contrast to an integration of the vorticity field which is quantitatively accurate.

Since particle traces give only meaningful information when positioned carefully—which, in particular, also applies to vorticity lines [14]—a more intuitive technique is needed to visualize the velocity and vorticity fields. For this purpose, movable slices with vector plots of variable resolutions have been integrated into the system. Although this visualization is simple and dated, it is nevertheless the most often used technique in physical flow visualization. Fig. 8, top, shows a screenshot of the visualization technique applied to K-type transition data. As seen in the images, the slices can also be used to clip unwanted parts of the visualization and to obtain a dense representation of any scalar field associated with the vector field, like, e.g., shear stress. When vector plots are insufficient, dense representations of the original vector field are often more useful. Therefore, line-integral–convolution (LIC) visualizations of the flow projected onto any movable slice can be shown [3]. Fig. 9 gives another example of movable slices and the use of LIC.

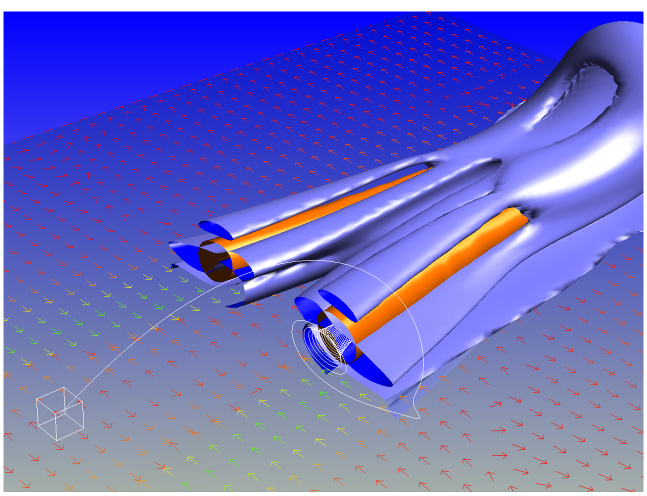

Figure 10: Combined visualization of separated vortices, particle probe integrating the velocity field, semi-transparent slice showing a vector plot of the velocity field, and shear layers.

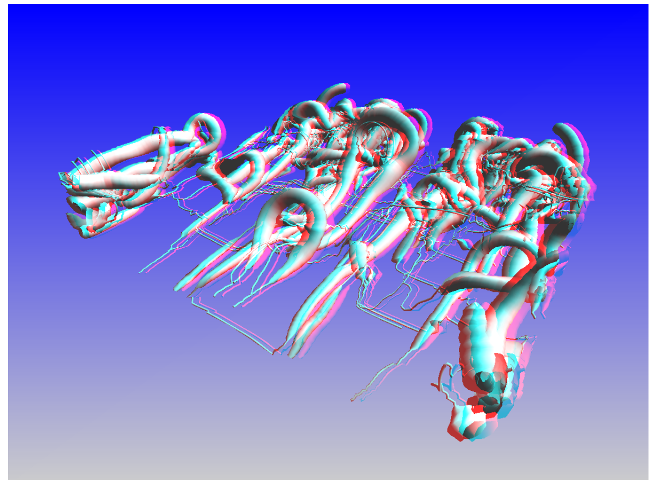

Figure 11: Stereo visualization of vortex set to help in understanding and identifying complex interwoven structures. To be viewed with red/cyan anaglyph glasses.

Shear stress is of special importance for the understanding of vortices. A fluid volume flowing along an object boundary is non-uniformly decelerated by friction with the wall. This in turn induces shear stress and finally swirling motion. Once the rotating fluid volume separates from the wall, a new vortex is born. The system allows the researcher to mix the visualization of separated vortices with an isosurface visualization of this important phenomenon. An example is shown in Fig. 10. The required shear layer computations are based on the second invariant I_2 of the strain-rate tensor S [13]:

$$
\begin{aligned}
I_2 &= \frac{1}{2}(S_{ii}S_{jj} - S_{ij}S_{ij}) \\
&= S_{11}S_{22} + S_{22}S_{33} + S_{11}S_{33} - S_{12}^2 - S_{13}^2 - S_{23}^2 \,.
\end{aligned}
$$

Since the strain-rate tensor is required for computing the λ_2 scalar field, anyway, computing the shear layer is virtually a by-product which comes for free.

As the visualizations in Figs. 3 might suggest, visualizing vortical flows results in very complex geometries. Depending on the chosen isovalue, the mesh of interwoven structures may be hard to analyze. Dissecting the set using the tools described above can alleviate the situation but it fails to ease the understanding at the very beginning. Therefore, a red/cyan anaglyph mode can be enabled making it easier for the researcher to mentally visualize the flow field and to perform his work. Fig. 11 gives a screenshot.

5 Evaluation

The present work was developed in close cooperation with fluid dynamics engineers and has been evaluated by this group of experts and practitioners. Both benefits and problems of the system were found.

The main benefit was found to be the novel combination of reliable vortex detection (showing also connected components) and selective visualization, a combination allowing for a great complexity reduction necessary for analyzing current simulation data. Once vortices have been extracted, the structures can be picked, hidden, extracted, and freely moved and rotated. Of course, the latter is standard functionality also found in more general visualization packages used so far (Tecplot, COVISE) and, indeed, the vortex detection part can be decoupled from the system to provide a set of vortices to be visualized with standard tools. However, in contrast to standard tools our system is problem– and meta-data–oriented; thus, analyzing flows with our tool was found to be significantly more efficient than with general tools, both from a usability point-of-view (intuitive, adapted user interface) and a technical point-of-view (speed).

Furthermore, the program does not require advanced graphics hardware features like programmable vertex or fragment processors or specific operating systems and can be used with virtually any workstation—a property not to be underrated in environments where most computations are performed on supercomputers and local computing hardware is predominantly acquired for computationally cheap post-processing. Accordingly, the possibility to view the visualizations in stereo using inexpensive red/cyan glasses was preferred over more sophisticated stereo technology requiring more expensive and less ubiquitous technology.

On the other hand, albeit the system can help to make the work of fluid engineers more efficient, it was also found that more meta data was needed and that the amount of data and structures was still too large for a straightforward analysis.

6 Conclusion

We have presented a novel algorithm for detecting and separating vortices from 3D vector fields. The resulting vortex set can be filtered to concentrate on the most relevant structures and the extracted vortices can be manipulated individually. Various visualization techniques were combined in order to improve the analysis possibilities. The system has been found by a group of experts to be of great value for analyzing vortical flows but it fails to reduce the original data to a level making processing trivial.

Since unsteady flows must currently be analyzed by processing the different time steps separately, incorporating automatic feature tracking as proposed by Silver and Wang [20] will be an important step towards a more useful tool. Furthermore, additional visualization techniques should be evaluated—like, e.g., texture advection which, however, currently fails to produce satisfying visualizations when applied to 3D data—as well as additional quantification functionality to distract the researchers from fancy images and let them concentrate again on the physics.

Acknowledgments

We are grateful to the German Research Council (DFG) for financing this work as part of SPP 1147.

References

[1] S. Bake, D. G. W. Meyer, and U. Rist. Turbulence mechanism in Klebanoff-transition. A quantitative comparison of experiment and direct numerical simulation. *J. Fluid Mech.*, pages 217–243, 459 2002.

[2] D. C. Banks and B. A. Singer. Vortex tubes in turbulent flows: identification, representation, reconstruction. In *VIS '94: Proceedings of the conference on Visualization '94*, pages 132–139, Los Alamitos, CA, USA, 1994. IEEE Computer Society Press.

[3] B. Cabral and L. Leedom. Imaging vector fields using line integral convolution. In *Proc. of SIGGRAPH-93: Computer Graphics*, pages 263–270, Anaheim, CA, 1993.

[4] Exa Corporation. PowerVIZ Specifications, 2001. http://www.exa.com/pdf/PowerVIZscreen.pdf.

[5] V.M. Fernandez, N.J. Zabusky, S. Bhat, D. Silver, and S.-Y. Chen. Visualization and Feature Extraction in Isotropic Navier-Stokes Turbulence. In *Proceedings of the AVS95 conference*, 1995.

[6] J. L. Helman and L. Hesselink. Representation and Display of Vector Field Topology in Fluid Flow Data Sets. *Computer*, 22(8):27–36, 1989.

[7] J. L. Helman and L. Hesselink. Surface Representations of Two- and Three-dimensional Fluid Flow Topology. In *VIS '90: Proceedings of the 1st conference on Visualization '90*, pages 6–13, Los Alamitos, CA, USA, 1990. IEEE Computer Society Press.

[8] R. Hooke and T. A. Jeeves. "Direct Search" Solution of Numerical and Statistical Problems. *J. ACM*, 8(2):212–229, 1961.

[9] J. Jeong and F. Hussain. On the Identification of a Vortex. *Journal of Fluid Mechanics*, pages 69–94, 285 1995.

[10] M. Jiang, R. Machiraju, and D. Thompson. Detection and Visualization of Vortices. In C. Johnson and C. Hansen, editors, *Visualization Handbook*. Academic Press, 2004.

[11] D. N. Kenwright. Automatic detection of open and closed separation and attachment lines. In *VIS '98: Proceedings of the conference on Visualization '98*, pages 151–158, Los Alamitos, CA, USA, 1998. IEEE Computer Society Press.

[12] D. Lovely and R. Haimes. Shock Detection from Computational Fluid Dynamics Results, 1999. AIAA 14th Computational Fluid Dynamics Conference, Paper 99-3285.

[13] D. Meyer. Direkte numerische Simulation nichtlinearer Transitionsmechnismen in der Strömungsgrenzschicht einer ebenen Platte, 2003. PhD Thesis, University of Stuttgart.

[14] P. Moin and J. Kim. The Structure of the Vorticity Field in Turbulent Channel Flow. Part 1. Analysis of Instantaneous Fields and Statistical Correlations. *Journal of Fluid Mechanics*, page 441, 155 1985.

[15] H.-G. Pagendarm and B. Seitz. An Algorithm for Detection and Visualization of Discontinuities in Scientific Data Fields Applied to Flow Data with Shock Waves. In P. Palamidese, editor, *Scientific Visualization: Advanced Software Techniques*, pages 161–177, Chichester, UK, 1993. Ellis Horwood Ltd.

[16] A. E. Perry, T. T. Lim, and E. W. Teh. A Visual Study of Turbulent Spots. *Journal of Fluid Mechanics*, 104:387–405, March 1981.

[17] U. Rist, K. Müller, and S. Wagner. Visualization of Late-Stage Transitional Structures in Numerical Data using Vortex Identification and Feature Extraction. In *Proc. 8th International Symposium on Flow Visualization*. paper no. 103.

[18] S. K. Robinson. Coherent Motions in the Turbulent Boundary Layer. *Annu. Rev. Fluid Mech*, 23:601–639, 1991.

[19] M. Roth and R. Peikert. A Higher-Order Method For Finding Vortex Core Lines. In D. Ebert, H. Hagen, and H. Rushmeier, editors, *Proceedings of IEEE Visualization '98*, pages 143–150. IEEE, October 1998.

[20] D. Silver and X. Wang. Tracking Scalar Features in Unstructured Datasets. In *Proceedings of IEEE Visualization '98*, pages 79–86. ACM Press, 1998.

[21] D. Silver and N. Zabusky. Quantifying Visualizations for Reduced Modeling in Nonlinear Science: Extracting Structures from Data Sets. *J. Visual Comm. and Image Representation*, 4(1):46–61, 1993.

[22] B. A. Singer and D. C. Banks. A Predictor-Corrector Scheme For Vortex Identification. Technical report, 1994.

[23] S. Stegmaier and T. Ertl. A Graphics Hardware-based Vortex Detection and Visualization System. In *Proceedings of IEEE Visualization '04*, pages 195–202, 2004.

[24] D. Sujudi and R. Haimes. Identification of swirling flow in 3D vector fields, 1995. AIAA 12th Computational Fluid Dynamics Conference, Paper 95-1715.

[25] J.-Z. Wu, A.-K. Xiong, and Y.-T. Yang. Axial Stretching and Vortex Definition. *Phys. Fluids*, 17(3), 2005.