On the role of domain-specific knowledge in the visualization of technical flows

Christoph Garth* Xavier Tricoche† Alexander Wiebel‡
Kenneth I. Joy§

Abstract

In this paper, we present an overview of a number of existing flow visualization methods, developed by the authors in the recent past, that are specifically aimed at integrating and leveraging domain-specific knowledge into the visualization process. These methods transcend the traditional divide between interactive exploration and feature-based schemes and allow a visualization user to benefit from the abstraction properties of feature extraction and topological methods while retaining intuitive and interactive control over the visual analysis process, as we demonstrate on a number of examples.

1 Introduction

The visualization of numerical fluid flow datasets is essential to the engineering processes that require their investigation through computational simulations. To address the need for visual representations that convey insight and enable a precise understanding of flow structures, the discipline of Flow Visualization has devised many methods and efficient implementations that support a variety of visualization tasks. Nonetheless, the ever increasing complexity of modern flow simulations puts an enormous demand on these methods, and significant limitations remain.

To this date, two major paradigms can be seen as the basis of most methods developed in flow visualization research. The first approach makes use of high-bandwidth and interactive display to present users with as much information as possible, providing maximum control in the selection and refined visual inspection of subsets that he deems interesting. Methods falling in this category typically rely on geometric primitives derived from numerical integration to interactively explore the dataset, using e.g. streamlines, stream surfaces, and pathlines. In contrast, the second paradigm is built upon the notion of data abstraction. To this end, a processing layer is added before the rendering stage to preselect significant aspects of the data and reduce the volume information that is ultimately conveyed visually. Prominent examples of this latter class of methods are topological methods and feature extraction schemes.

*Institute for Data Analysis and Visualization, University of California, Davis (cgarth@ucdavis.edu)
†Department of Computer Science, Purdue University (xmt@cs.purdue.edu)
‡Abteilung Bild- und Signalverarbeitung, Universität Leipzig (wiebel@informatik.uni-leipzig.de)
§Institute for Data Analysis and Visualization, University of California, Davis (kijoy@ucdavis.edu)

Both approaches have their specific weaknesses. Interactive exploration of datasets is based on the assumption that displaying sparse and arbitrary subsets of the data can provide an effective graphical representation that allows the user to make sense of the underlying phenomenon. This visual exploration is often hampered both by the intrinsic complexity of the data and by the geometric intricacy of corresponding visualizations. On the other hand, abstraction methods explicitly aim at solving these particular problems by de-correlating the amount of information shown in the visual representation from the sheer size of the input dataset. However, they are faced with the challenge of correctly identifying the characteristics of the data that are indeed most meaningful to an analyst. Therefore, the difficulty in adopting either of these two paradigms is that the responsibility for meaningful visualization is either completely left in the hands of users or taken away from them entirely. In the former scenario, users must make sense of a deluge of information that by all accounts overwhelms human processing ability. In the latter, they must trust that the machinery running behind the scene is properly capturing the high-level meaning of the data to offer a reliable picture of the considered phenomenon.

The intent of this paper is to provide and overview of several methods introduced by the authors in recent years that strive to reach a middle ground by allowing the user to interact with the data abstraction process and to leverage the resulting data abstraction as a high-level contextual information for the interactive exploration of a flow. Our general approach couples interaction and abstraction in two different ways. After presenting some of the basic principles and prior work in Section 2, we discuss the section plane methodology (Section 3). It is geared toward the topological analysis of certain flow structures as well as of user-prescribed subsets of a given dataset. By doing so we provide a user with the ability to leverage pre-existing domain-specific knowledge of the dataset under consideration and the structures it contains. Furthermore, we demonstrate a number of approaches that allow the user to explore subsets of a dataset on the basis of a pre-processing abstraction step, thereby retaining the benefits of interactive exploration while simultaneously limiting its complexity and providing required context. Section 4 demonstrates this idea based on Lagrangian analysis, and in Section 5, we show how topological analysis of boundary flow can serve as a building block for interactive visualization of volume flows. Both Section 4 and Section 5 focus on the application domain-specific knowledge in the form of typical structure interactions to simplify the visualization process. Finally, we conclude on the presented material in Section 6.

Remark that due to a limitation in available space, we deliberately illustrate the general approach that we wish to advocate through methods that have been developed by the authors (and previously published in [GTS04, TGK$^+$04, GGTH07, WTS07a, WTS$^+$07b]), even if our discussion is meant to be general.

2 Background

2.1 Feature Extraction

The goal of feature-based visualization methods is to generate images that restrict the depiction of complex flow data to a limited set of points, lines, and volumes representing features of particular interest for the considered application. This yields fairly abstract pictures that convey significant flow properties in a concise and compact form. The most prominent examples of features in CFD applications include vortices, separation and attachment manifolds, shock waves and recirculation zones. The loose, empiric nature of the definition of those feature explains the variety of algorithms available to locate, identify, and visualize them and requires the user to determine experimentally which method is best suited for the needs of his particular application. Further restrictions on the type of method can be imposed by the size or the structure of the data.

Vortices The extraction of vortical structures has been a major topic in visualization for quite some time. Although a vortex is most intuitively conceived as the superposition of a flow along an axis and a flow around this axis, there are few satisfying definitions that exists for this flow pattern [Lug96, Hal05]. As a result, vortex extraction methods are essentially characterized by the type of vortex criterion they are built on.

This are either region-based criteria (identifying regions of vortical flow behavior) or a line-type description (focusing on the vortical axis or vortex core line). Region definitions include high vorticity, helicity, low pressure. Most often used in engineering applications is the λ_2 definition by Jeong and Hussain [JH95]. The major limitation of λ_2, however, lies in its incapacity to isolate individual structures.

Among the line-type definitions, the approach of Sujudi and Haimes [SH95] is most widely used. The idea here is to perform on a cell-wise basis the pattern matching of a rotation motion on the vector field and to extract locally sections of the rotation axis that can be patched together to approximate the vortex core line. Because of the linear nature of the sought pattern, the method has issues with vortex core lines that are strongly curved. Roth and Peikert therefore proposed a higher-order scheme that can extract curved core lines reliably [RP98]. They also showed in a subsequent paper that this and other similar methods can be formulated in a unified framework involving their *parallel vector fields operator* [PR00]. Most recently, Weinkauf et al. extended this approach to the extraction of swirling particle motion in unsteady flows [WSTH07].

While the above methods invoke a preprocessing step and present results in the form of line segments or surfaces, there is another class of methods that try to identify a vortical or swirling flow behavior by examining the trajectories of particles [JMT02]. Recently, Garth et al. presented a more general stream surface-based approach [GTSS04] that allows for the extraction of a vortex core line approximated as the medial axis of a stream surface that undergoes vortical motion.

Separation and Attachment Lines Separation and attachment lines are another major feature type. They are defined as the lines along which the flow attaches or separates from the surface of a flow-embedded object. The analysis is focused on the non-zero, tangential shear-stress vector field defined over the surface that exhibits the same flow patterns as nearby located streamlines. In particular, flow separation and attachment induce the creation of curves of asymptotic streamline convergence in the shear stress vector field. The corresponding three-dimensional flow pattern is characterized by the presence of a stream

surface starting or ending along the feature line and swirling around a nearby located vortex. As a consequence, flow separation and vortex genesis are closely related phenomena (cf. [Dal85]).
Following the original idea of Sujudi and Haimes for vortex core lines, Kenwright et al. proposed a simple and fast method for the extraction of separation and attachment lines [KHL94]. Their basic observation is that these feature lines are present in two linear patterns, namely saddle points and nodes, where they are aligned with an eigenvector of the flow field Jacobian. The original method works on a cell-wise basis to extract this pattern within each triangle. Hence it results in disconnected line segments, caused by the discontinuity of the Jacobian. Yet, applying the parallel operator leads to the reformulation of the features in terms of lines of zero curvature and yields connected lines. However this definition is quite restrictive because it assumes that separation resp. attachment lines always have zero curvature. Moreover, since it requires derivative computation it is very sensitive to noise. Consequently strong pre-smoothing of the data is often necessary which in turn can deform and shift the features. Another approach was proposed earlier by Okada and Kao [OK97] who extend the classical Line Integral Convolution (LIC) algorithm [CL93] by color coding the flow direction so as to highlight the fast changes in flow direction that occur as streamlines approach separation resp. attachment lines. The weakness of this approach lies in the heavy computation associated with LIC on one hand, and in the fact that the geometry of the feature lines is not extracted. Instead, the method computes a density function that indicates the proximity / likelihood of these feature lines. Using a releated approach, Tricoche et al. recently proposed a scheme [TGS03] designed to overcome the restrictions imposed by the purely local analysis used in the algorithms mentioned previously. Their method is built upon monitoring the convergence of streamlines in the shear stress vector field and to determine streamlines of maximal local convergence or divergence as separation or attachment lines.

2.2 Topological Flow Analysis

Vector field topology is a powerful approach for the visualization of flows. Topology-based methods leverage basic results of the qualitative theory of dynamical systems to generate effective depictions characterized by a high level of abstraction and an accurate segmentation of the domain in regions where the flow exhibits a uniform behavior. Unfortunately, the application of this methodology to three-dimensional problems has not so far demonstrated the same usefulness in visualization applications as it does for two-dimensional flow fields. Two independent problems can be named to explain this discrepancy. First, the topology of volume flows involve stream surfaces that are plagued by self-occlusion and visual complexity. Secondly, a lack of connection between topological structures and major features of interest in fluid dynamics problems, as described in the previous section, make topological methods less useful for 3D flow visualization. Neither vortices nor separation lines can be, in general, reliably located as elements of flow topology. However, work in this direction is underway [SJH07].
In the following we provide a short introduction to essential notions of vector field topology. Our presentation is driven by the needs of visualization algorithms discussed in this paper. For a more complete survey of existing methods in topology-based flow visualization, we

refer the reader to [ST05].

Topological Graph The topology of a vector field is the decomposition of its phase portrait into regions where all streamlines have the same limit sets. The notion of *phase portrait* considers all points located along the same streamline as a single equivalence class. The *critical points* of a steady vector field are the locations where the field magnitude vanishes, and in the non-degenerate, linear case, the nature of a critical point can be classified by the eigenvalues of the vector field Jacobian matrix. For the work presented here, we are mostly interested in spiral-type critical points that induce a spiraling behavior in the vector field trajectories. Besides critical points, *cycles* are closed streamlines that correspond to periodic trajectories.
These *limit sets* present in a vector field induce a segmentation of the domain into regions where all streamlines share the same limit sets for forward and backward integration. The boundaries between such regions are called separatrices. The *topological graph* then consists of all limit sets and the separatrices connecting them.
Computationally, many algorithms exist to extract the topological graph of a vector field in both two- and three-dimensions (e.g. [GLL91]). They mostly differ by the vector field representation they are based on.

Parametric Topology and Bifurcations In the case where a vector field depends on a parameter (e.g. time), changes in the parameter value induce changes to its topology. These transformations are called bifurcations and exist in an infinite variety. Their common property nonetheless is to replace a stable structural configuration by another stable configuration through an instantaneous, unstable pattern. In this context, stability is defined with respect to the ability of a given structure to remain qualitatively unchanged after a small but arbitrary modification of the vector field. Bifurcations are either local or global depending on the extent of the region they impact. Basic classification of such bifurcations is possible in both two- and three-dimensions, however, we will not go into detail here and instead refer the reader to [GH83]. In between such bifurcation events, the critical points may move through the domain of the vector field. Together with bifurcations, the paths of critical points form the *structural graph* of a parametric vector field.
An algorithmic solution to track the continuous evolution of the topology and detect the associated bifurcations was proposed by Tricoche et al. in [TWSH02]. The method was designed for two-dimensional parametric flows represented on a piecewise linear mesh, and an extension of this scheme to three-dimensional vector fields was described in [GTS04]. Another approach called *Feature Flow Field* exploits a different computational representation of the considered vector field (cf. [TS03]).

2.3 Lagrangian Analysis

The *finite-time Lyapunov exponent* (FTLE) is a computational tool that can be used to define and extract coherent structures in transient flows studied in a Lagrangian framework (i.e. based on particle trajectories). It has been the object of a growing interest in fluid dynamics research over the last few years and has been successfully applied to a variety of

fluid dynamics problems. Its derivation and application to aperiodic time-dependent flows, has been recently described by Haller [Hal01a]. We give a brief overview of this approach in the following, with our presentation being voluntarily informal, and refer the reader to [LSM06] for a comprehensive discussion.
Essentially, Lagrangian analysis methods are based on observing the behavior of particle trajectories. In the specific case of the FTLE, one is interested in the exponential separation rate of particle trajectories that start closely together and move apart after a finite time. The regions of locally maximal exponential separation rate, i.e. the ridges of the scalar FTLE field, are then called *Lagrangian Coherent Structures*, and one can show that they provide a basic skeleton of an unsteady flow field [Hal01b, SLM05, LSM06].
While the computation of FTLE fields for a given flow volume is computationally tedious due to the large number of particle advections required for the FTLE computation, several approaches for an accelerated computation have recently been presented. Garth et al. [GGTH07] make use of large-scale coherence of particle trajectories to adaptively approximate FTLE fields, while Sadlo et al. [SP07] focus more specifically on the direct approximation of LCS by starting out with a coarse flow field sampling and then refining in the vicinity of detected LCS.

3 Moving Section Planes

The topological analysis of a vector field enables qualitative understanding of the dynamical system generated by the field in terms of the phase portrait. However, a graphical representation of the topological graph is necessary to convey the structural connections. While such a representation is straightforward for a two-dimensional domain of definition, with its topological graph consisting of points and lines, higher dimensions are much more difficult to treat. There, the limit sets that constitute the topological graph may be of dimension greater than one. In three dimensions, for example, there are separation surfaces and closed invariant tori. Even in simple cases, a straightforward depiction of the topological graph is difficult as these surfaces, as opposed to curves, occlude each other. Furthermore, a straightforward topological analysis of three-dimensional flow field datasets has proven elusive. Here, we discuss a different approach and apply the notion of *topology tracking* as a basis for the topological visualization of three-dimensional flow structures based on the concept of moving section planes (cf. [TGB$^+$04]).
In engineering practice, it is quite common to avoid occlusion problems in the visualization of three-dimensional datasets by employing planar sections. Typically, color mapping is used to depict scalar quantities of interest on the section plane. The central idea of the vector field visualization method described in this section is to extend these basic and widely used planar sections as tool for exploring flow volumes of stationary flows. The planes smoothly travel in a continuous way along curves that can be either obtained automatically by standard feature extraction schemes or directly provided by the user to explore a particular region. For each point of the plane trajectory, the vector field is resampled and projected, resulting in a two-dimensional vector field on the plane. As the plane moves, the changing projected vector field can be expressed as a parameter-dependent two-dimensional vector field and is hence amenable to topology tracking methods (cf. Section 2.2). The resulting

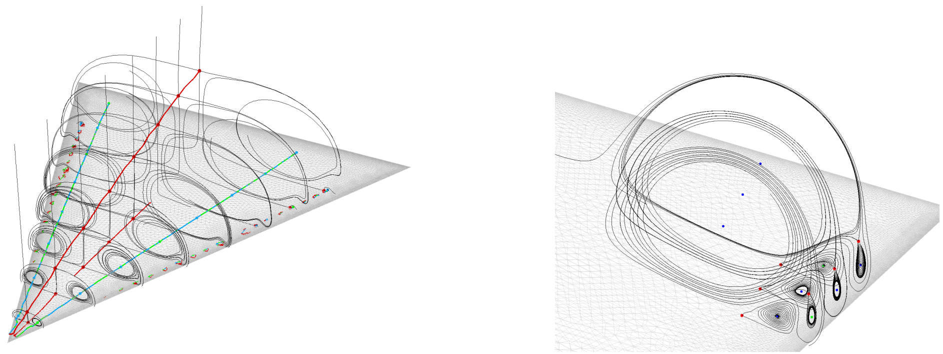

Figure 1: Flow above a delta wing: parametric topology extracts the primary vortices. The planes travel along the symmetry axis of the wing. Spiral-type critical points indicate rotation centers (left). Automatic plane orientation shows not one but three vortices above the wing (right). The interactions in the vortex system are clearly visible.

structural graph includes the paths of critical points, which are then mapped to the moving section plane coordinate system for visualization purposes.

During our visualization experiments (described below), we have made use of several types of trajectories that are based on specific knowledge about a dataset under consideration and the structures contained in it, and accomodate different visualization aspects. Using these or other trajectories, the section plane approach enables the user to introduce existing knowledge of flow structures into the visualization process.

Straight Line A first type of trajectory is a user-selected straight line. This approach has two major applications. First, if the focus is on large-scale vortical structures, the mean flow direction in the corresponding region can be selected along with a convenient start position. This technique can be used to explore datasets where application-specific knowledge predetermines the location of vortices and vortex systems. The second application arises when one is interested in structures that deviate from an overall dominating flow component, such as e.g. the flow surrounding an embedded object. By using the section plane approach and orienting the planes roughly orthogonal to the dominating flow direction, the latter is effectively discarded and visualization is focused on deviating structures. Overall, this trajectory type reflects interactive exploration of a dataset with minimal prior knowledge. In contrast to a fully three-dimensional approach, the resulting complexity is much reduced. An examples is provided in Figure 1.

Vortex Core Lines A second idea is to select a vortex core line to serve as trajectory. These feature lines are the center of the swirling flow and are therefore natural candidates to capture the local symmetry of the flow. Essentially, if the section plane is parallel to the plane of rotation, the component of the velocity field along the vortex core line is discarded, and the rotational motion is captured in the form of a spiral-type critical point. In this way,

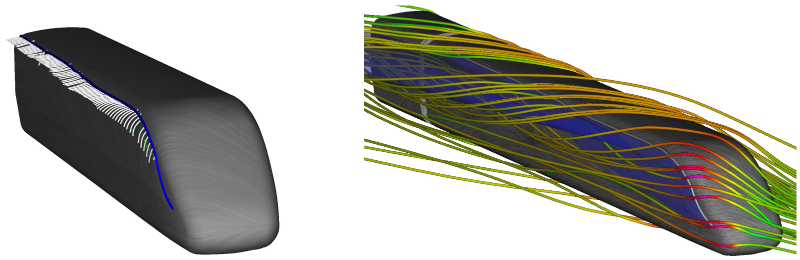

Figure 2: Section planes applied to attachment line on a high-speed train wall. The identification of the correct separation angle (left) allows the seeding of the corresponding separation surface (right).

interactions between different vortices may be visualized as they are captured by the plane topology.

Separation/Attachment Lines Similar to vortex core lines, surface feature lines such as separation and attachment lines can serve as trajectories of moving section planes in order to visualize the structure of flow separation over the surface. Furthermore, they can serve as a tool to seed separation surfaces in a correct fashion [WTS07a], thereby linking knowledge about surface features to volume flow visualization (see Figure 2).

Recirculation Bubble Axis A last type of trajectory is directly fitted to the visualization of so-called *recirculation zones*. These structures essentially represent closed vortex core lines and are typically difficult to extract reliably using vortex detection schemes. Such structures can be visualized by rotating the section plane around an approximate symmetry axis of the recirculation (cf. Figure 3).
In order to guarantee meaningful results in the purely exploratory case, the orientation of the cutting plane along the path is either fully determined or must be chosen according to the local flow structure. The first situation occurs e.g. during the investigation of flow structures with a distorted geometry or when following a straight line to investigate large-scale features. In the latter case the trajectory could provide the plane normal. However, when dealing with e.g. with a vortex core line, a normal that does not coincide with the axis of rotation may lead to an observed shift in the location of the planar rotation center. To guarantee good results in these cases, the quality of a normal is evaluated with respect to the amount flow crossing the corresponding plane, and the plane is oriented such that the crossing flow is maximized.

3.1 In-Context Visualization using Volume Rendering

The sparse visual representation provided by the moving section plane topology is ideally suited to a combination with other, dense representations. Such dense visualizations include

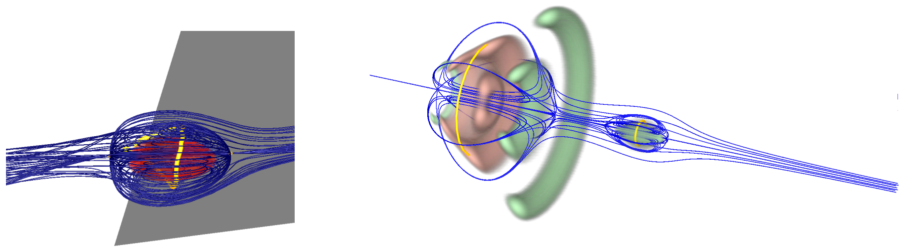

Figure 3: Section plane topology of vortex breakdown bubbles. Section planes travel around the rotational symmetry axis of the breakdown bubbles. Volume rendering of scalar flow quantities provides context and allows correlation of the breakdown geometry with the surrounding flow.

volume rendering of scalar fields. The context provided through this approach allows a better interpretation of the true three-dimensional geometry of the flow structures encountered by the moving section planes. Furthermore, the framework of multi-dimensional transfer functions in volume rendering applications is extremely powerful in allowing for the simultaneous and coherent representation of complementary derived quantities [KKH02], and intuitive techniques for the manipulation of such transfer functions exist. Overall, these methods are ideal companions for topological visualization in an interactive context, since they allow a further incorporation of domain-specific knowledge of physically relevant scalar flow quantities such as density, pressure, vorticity magnitude etc. [TGK$^+$04]. An example is given in Figure 3.

4 Lagrangian Analysis on Section Planes

Straightforward Lagrangian analysis and visualization based on the Finite-Time Lyapunov Exponent provides a strong reduction in the complexity of (time-dependent) volumetric flow analysis by essentially reducing a flow field to a scalar field that provides information about the converging or diverging nature of pathlines. A further abstraction can be achieved by computing the ridge lines of this scalar field, resulting in a skeleton of Lagrangian Coherent Structures. However, the complexity of typical flow structures makes it difficult to isolate individual structures by means of automated methods (cf. [GGTH07, SP07]). We instead propose to limit the computation of an FTLE field to a single section plane or to multiple section planes.

Visualizing a Finite-Time Lyapunov Exponent field on a section plane essentially provides a means to characterize the coherence of the particles that intersect the corresponding plane. Compared to the three-dimensional case it reduces both the visual complexity and the typically great computational complexity of the associated computation, and is therefore much more accessible in an interactive setting. Despite a lack of theoretical foundation, FTLE fields on section plane are still a reliable tool for unsteady flow visualization since the La-

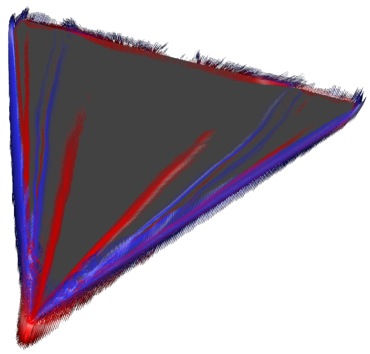

(a) Separation and attachment structures on the wing surface. Pathlines were seeded from FTLE ridges on a plane parallel to the wing.

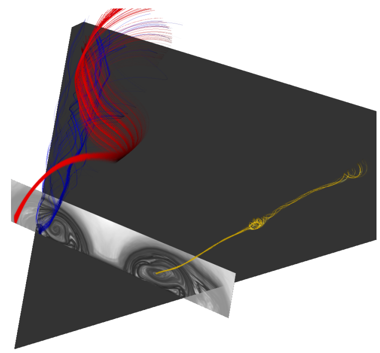

(b) Pathline visualization supported by section plane FTLE. Pathlines seeded according to PDF in (c).

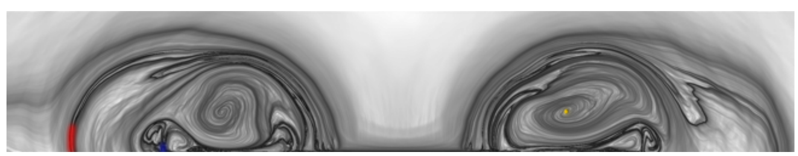

(c) Planar FTLE section perpendicular to the main flow direction. Darker regions correspond to regions of high FTLE. Colored regions indicate user-guided PDFs (see (b)).

Figure 4: Fuzzy pathline visualization of flow structures above a delta wing. Pathlines are seeded according to FTLE strength, FTLE ridge lines or through a user-guided Probability Density Function.

grangian Coherent Structures they indicate are essentially preserved as cross-sections if the section plane is approximately orthogonal to the flow direction along these structures (more details are provided in [GGTH07]).

From a visualization perspective, the two-dimensional scalar field that results from the section plane FTLE computation can further be made accessible to abstraction and interaction (see examples provided in Figure 4).

Ridge Lines In the original three-dimensional setting, Lagrangian Coherent Structures appear as the ridges of the volumetric FTLE field. Therefore, using ridge line extraction to isolate their corresponding sections on the plane allows a further abstraction step. These ridges are most useful when used to seed integral surfaces and pathlines (Figure 4(a)).

Stochastic Particle Seeding An FTLE field may be interpreted as a Probability Density Function (PDF) for particle seed distribution. The rejection rule applied to this data is then used to determine stochastically a set of seed points that overall emphasize regions

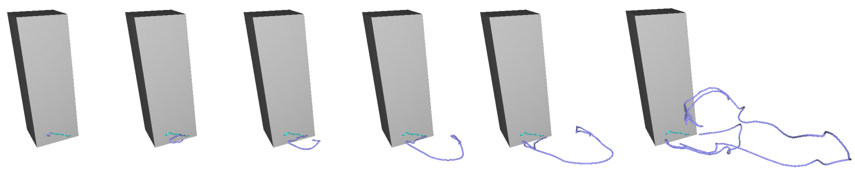

Figure 5: Generalized streak lines in the flow past a cuboid. The streak line (blue) starts from the moving swirl-type critical point on the cuboid boundary. As the critical point reaches a bifurcation, the streak line is terminated and separates from the boundary.

Figure 6: Generalized streak line emanating from a swirl-type critical point in the flow past an ellipsoid, in combination with volume rendering of λ_2-criterion (left) and vorticity (right).

of high FTLE values and as such of high structural coherence. The fuzziness of the resulting representation matches naturally the fuzziness of flow coherence and the uncertainty involved in their computational characterization.

Image-based User Interface Planar FTLE sections provide the user with a look-up map over which interesting regions that may be difficult or impossible to extract automatically can be manually and selectively identified by simple brushing to provide a PDF similar to the one mentioned previously but this time geared towards the specific focus of the analysis. This reduces the visual complexity of the final image, to emphasize most prominent aspects in the data, and it provides an intuitive interface to do so. Hence it is an effective means to further support interactive exploration of a dataset by hinting at interesting seed regions on the section plane. Furthermore, it enables to user to confirm an interpretation of specific flow structures by visualizing the corresponding integral curves. Conversely, it allows a straightforward and precise selection of those integral curves that visualize a specific pattern or feature visible on the section plane (Figures 4(b) and 4(c)).

5 Shear Stress Topology and Generalized Streak Lines

As discussed in Section 2, the topology of the wall shear stress vector field of an object embedded in a surrounding flow can be used to infer flow structures that are close to the boundary or originate on it. It was shown recently, that the combination of volumetric flow visualization methods with those that visualize boundary flow can be beneficial in understanding these interactions (cf. [GLT$^+$06]). In this section, we discuss an approach to facilitate the direct application of the knowledge of boundary structures to visualize volumetric flow structures in the case of vortices.

In [WTS$^+$07b], two of the authors proposed a method to track the paths of critical points on object boundaries over time in an unsteady dataset. This former is achieved by introducing a specific parameterization of an object boundary that simplifies the task of tracing the critical points on a curved surface to the planar setting, where efficient algorithms are readily available. The information obtained thereby is then used to facilitate three-dimensional visualization of the vortices corresponding to spiral-type critical points in the following ways.

Pathline Seeding A straightforward approach is to use information about critical points on the boundary to allow automated seeding of pathlines in close proximity. Through this, the tedious task of manual selection of such pathlines is greatly simplified. See Figure 7.

Generalized Streak Lines The paths of the tracked critical points can serve as exact loci of particle placement. By continuously injecting particles into the flow as time evolves, a *generalized streak line* is obtained, which visualizes the interrelation between wall shear stress and three-dimensional flow. Figure 5 provides an example. This method proves especially effective by combining it with volume rendering techniques (cf. Figure 6 to depict scalar flow quantities (as described in Section 3.1), thus providing necessary context and allowing an application of domain-specific knowledge to understand the evolution of a vortex as it progresses away from the surface.

6 Conclusion

In this paper, we have given an overview of a number of recently proposed visualization techniques that exemplify the idea of bridging the traditional gap between visualization techniques that focus on either interactive exploration or data abstraction. We have shown

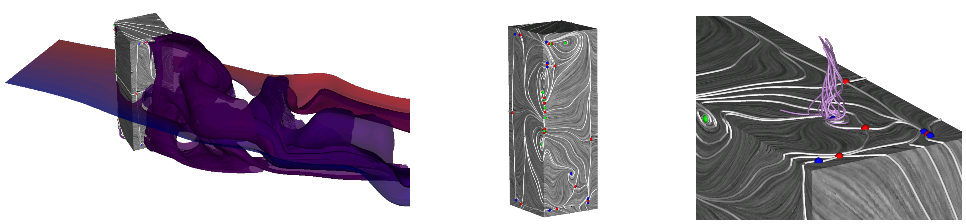

Figure 7: Depictions of the flow on the boundary of an embedded cuboid. Stream surface showing turbulent behavior behind cuboid (left). LIC and topological structures of shear stress field on cuboid (middle). Swirling streamlines seeded close to a spiral critical point indicate a vortex in the flow volume (right).

that by combining their respective properties, a user is enabled to leverage domain-specific knowledge about the visualization problem at hand. This approach can considerably simplify the visualization task for practical datasets. Future research along these lines appears promising as the path from data to knowledge becomes more and more challenging.

Acknowledgements The authors would like to thank the *FAnToM* development group at the University of Leipzig. Alexander Wiebel received supported by the Deutsche Forschungsgemeinschaft through grant SCHE 663/3-7 during the course of this work. Christoph Garth was in part supported by the DFG International Research Training Group 1131 based at the University of Kaiserslautern, Germany.

References

[CL93] B. Cabral and L. C. Leedom. Imaging vector fields using line integral convolution. In *Proc. 20th Annual Conference on Computer Graphics and Interactive Techniques*, 1993.

[Dal85] U. Dallmann. On the formation of three-dimensional vortex flow structures. Technical Report 221-85 A 13, Deutsche Forschungs- und Versuchsanstalt fuer Luft- und Raumfahrt, 1985.

[GGTH07] C. Garth, F. Gerhard, X. Tricoche, and H. Hagen. Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1464–1471, 2007.

[GH83] J. Guckenheimer and P. Holmes. *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, 1983.

[GLL91] A. Globus, C. Levit, and T. Lasinski. A tool for visualizing the topology if three-dimensional vector fields. In *IEEE Visualization Proceedings*, pages 33 – 40, October 1991.

[GLT+06] C. Garth, R.S. Laramee, X. Tricoche, J. Schneider, and H. Hagen. Extraction and visualization of swirl and tumble motion from engine simulation data. In *Proceedings of The Topology-Based Methods in Visualization Workshop*, 2006.

[GTS04] C. Garth, X. Tricoche, and Gerik Scheuermann. Tracking of Vector Field Singularities in Unstructured 3D Time-Dependent Datasets. In *Proceedings of IEEE Visualization*, pages 329–336, October 2004.

[GTSS04] C. Garth, X. Tricoche, T. Salzbrunn, and G. Scheuermann. Surface Techniques for Vortex Visualization. In *Proceedings Eurographics - IEEE TCVG Symposium on Visualization*, 2004.

[Hal01a] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional flows. *Physica D*, 149:248–277, 2001.

[Hal01b] G. Haller. Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence. *Physics of Fluids*, 13(11), 2001.

[Hal05] G. Haller. An objective definition of a vortex. *Journal of Fluid Mechanics*, 525:1–26, 2005.

[JH95] J. Jeong and F. Hussain. On the Identification of a Vortex. *Journal of Fluid Mechanics*, 285:69 – 94, 1995.

[JMT02] M. Jiang, R. Machiraju, and D. Thompson. A Novel Approach to Vortex Core Detection. In *Data Visualization 2002*, pages 217 – 226, Aire-la-Ville, Switzerland, 2002. Eurographics Association.

[KHL94] D. Kenwright, C. Henze, and C. Levit. Feature extraction of separation and attachment lines. *IEEE Transactions on Visualization and Computer Graphics*, 5(2):135–144, 1994.

[KKH02] J. Kniss, G. Kindlmann, and C. Hansen. Multi-dimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):270–285, July-September 2002.

[LSM06] F. Lekien, S. C. Shadden, and J. E. Marsden. N-dimensional lagrangian coherent structures: Definition and properties. *Physica D*, 2006.

[Lug96] H. J. Lugt. *Introduction to Vortex Theory*. Vortex Flow Press, Inc., 1996.

[OK97] A. Okada and D. L. Kao. Enhanced line integral convolution with flow feature detection. In *Proceedings of IS&T/SPIE Electronic Imaging*, pages 206–217, 1997.

[PR00] R. Peikert and M. Roth. The "parallel vectors" operator - a vector field visualization primitive. In *IEEE Visualization Proceedings '00*, pages 263 – 270, 2000.

[RP98] M. Roth and R. Peikert. A higher-order method for finding vortex core lines. In *IEEE Visualization Proceedings '98*, pages 143 – 150, 1998.

[SH95] D. Sujudi and R. Haimes. Identification of Swirling Flow in 3D Vector Fields. Technical Report AIAA Paper 95–1715, American Institute of Aeronautics and Astronautics, 1995.

[SJH07] A. Surana, G. B. Jacobs, and G. Haller. Extraction of separation and reattachment surfaces from 3d steady shear flow. *AIAA Journal*, 45:1290 – 1302, 2007.

[SLM05] S. C. Shadden, F. Lekien, and J. E. Marsden. Definition and properties of lagrangian coherent structures from finite-time lyapunov exponents in two-dimensional aperiodic flows. *Physica D*, 212:271–304, 2005.

[SP07] F. Sadlo and R. Peikert. Efficient visualization of lagrangian coherent structures by filtered amr ridge extraction. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1456–1463, 2007.

[ST05] G. Scheuermann and X. Tricoche. Topological methods for flow visualization. In C.D. Hansen and C.R. Johnson, editors, *The Visualization Handbook*, pages 341–356. Elsevier, 2005.

[TGB+04] X. Tricoche, C. Garth, T. Bobach, G. Scheuermann, and M. Ruetten. Accurate and efficient visualization of flow structures in a delta wing simulation. In *Proceedings of 34th AIAA Fluid Dynamics Conference and Exhibit, AIAA Paper 2004-2153*, 2004.

[TGK+04] X. Tricoche, C. Garth, G. Kindlmann, E. Deines, G. Scheuermann, M. Rütten, and C. D. Hansen. Vizualization of Intricate Flow Structures for Vortex Breakdown Analysis. In *Proceedings of IEEE Visualization*, pages 187–194, October 2004.

[TGS03] X. Tricoche, C. Garth, and Gerik Scheuermann. Fast and robust extraction of separation line features. In *Proceedings of the Dagstuhl Scientific Visualization Seminar*, 2003. to appear.

[TS03] H. Theisel and H.-P. Seidel. Feature flow fields. In *VISSYM '03: Proceedings of the symposium on Data visualisation 2003*, pages 141–148, Aire-la-Ville, Switzerland, Switzerland, 2003. Eurographics Association.

[TWSH02] X. Tricoche, T. Wischgoll, G. Scheuermann, and H. Hagen. Topology Tracking for the Visualization of Time-Dependent Two-Dimensional Flows. *Computers & Graphics*, 26(2):249 – 257, 2002.

[WSTH07] T. Weinkauf, J. Sahner, H. Theisel, and H.-C. Hege. Cores of swirling particle motion in unsteady flows. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1759–1766, 2007.

[WTS07a] A. Wiebel, X. Tricoche, and G. Scheuermann. Extraction of separation manifolds using topological structures in flow cross sections. In *Proceedings of the Topology-Based Methods in Visualization Workshop*, 2007. to appear.

[WTS+07b] A. Wiebel, X. Tricoche, D. Schneider, H. Jänicke, and G. Scheuermann. Generalized streak lines: Analysis and visualization of boundary induced vortices. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1735–1742, 2007.