

Lagrangian Visualization of Flow-Embedded Surface Structures

Christoph Garth¹, Alexander Wiebel², Xavier Tricoche³, Ken Joy¹ and Gerik Scheuermann²

¹Institute for Data Analysis and Visualization, University of California, Davis, USA

²Abteilung Bild- und Signalverarbeitung, Institut für Informatik, Universität Leipzig, Germany

³Computer Science Department, Purdue University, USA

Abstract

The powerful yet conceptually simple notions of Finite-Time Lyapunov Exponent (FTLE) and Lagrangian Coherent Structures provide a powerful framework for the analysis and visualization of complex technical flows. Their definition is simple and intuitive, and they are built on a deep theoretical foundation. We apply these concepts to enable the analysis of flows in the immediate vicinity of the boundaries of flow-embedded objects by limiting the Lagrangian analysis to surfaces closely neighboring these boundaries. To this purpose, we present an approach to approximate FTLE fields over such surfaces. Furthermore, we achieve an effective depiction of boundary-related flow structures such as separation and attachment over object boundaries and specific insight into the surrounding flow using several specifically chosen visualization techniques. We document the viability of our methods by presenting a number of application examples.

Categories and Subject Descriptors (according to ACM CCS): I.6.6 [Simulation and Modeling]: Simulation Output Analysis; J.2 [Physical Sciences and Engineering]: Engineering, Physics

1. Introduction

The study of fluid flows is an essential topic in a broad range of engineering, bio-medical and scientific applications and is aimed at understanding and control of technical flows. Examples include important application areas such as combustion, turbo-machinery, automotive engineering, and aeronautics. Modern simulation techniques facilitate computer models that allow a comfortable study of many application scenarios. A key role is played by flow visualization techniques by providing tools that allow increased insight into and understanding of numerical flow datasets.

The complexity and size of modern CFD simulations however complicates the visualization task greatly. One possibility to reduce its complexity is to limit the analysis of a volumetric flow to the study of its interaction with objects boundaries. This approach is feasible because flow structures of interest often result from the interplay of flows with embedded objects. Typical examples include vortex shedding and shear layers, caused by creation of vorticity at boundaries and its successive transport into the flow volume. Such

flow processes reflect on object boundaries in the form of separation and attachment flow patterns.

Recently, the notion of Lagrangian flow analysis using the notions of Finite-Time Lyapunov Exponent and Lagrangian Coherent Structures has emerged as a powerful framework for the visualization and analysis of both steady and unsteady flows. However, the applicability of methods founded on these concepts has been limited to two-dimensional flow domains or three-dimensional flow volumes. In this paper, we propose a definition of Finite-Time Lyapunov Exponents on surfaces (Section 3) in the vicinity of flow boundaries that allows a comprehensive analysis of boundary flow regions and gives rise to a number of visualization techniques that are aimed at extending this analysis away from the boundaries into the volume flow. The presented concept is closely analogous to the established Lagrangian methods, benefits from their conceptual simplicity and retains many of their desirable properties. Furthermore, we describe our implementation of this concept in Section 4 and discuss a number of visualization approaches based on surface FTLE measures and focus our presentation on both automated,

non-interactive and user-assisted scenarios (Section 5). To demonstrate the applicability of the presented work in the context of modern CFD datasets, we provide several examples and discuss specific aspects of our methods in Section 6. Finally, we conclude with a discussion of the presented ideas and future work (Section 7).

2. Previous Work

The concept of Finite-Time Lyapunov Exponents (FTLE) as a means to describe and analyze coherent Lagrangian structures in transient flows was introduced by Haller [Hal01a] in 2001. He presented FTLE as a geometric approach, aimed at characterizing coherent structures in terms of preservation of certain stability types of the velocity gradient along the path of a particle. This initial research generated a significant interest in FTLE and its applications to the structural analysis of transient flows. The same author also proved a study of the robustness of the structures characterized by FTLE [Hal02] and showed that even under approximation errors, they remain intact. Shadden et al. provided a more formal discussion of the theory of FTLE [SLM05] in two dimensions, and an extension to arbitrary dimensions was discussed in [LSM06]. These tools have been applied to the study of turbulent flows [Hal01b, GRH07, MHP*07] and used in the analysis of vortex ring flows [SDM06]. However, the visualizations presented in these papers were chosen on a case-by-case basis, and no systematic investigation of different visualization types was considered. As of recently, the topic has also been treated in the visualization community. Garth et al. focused on accelerated computation of FTLE fields and their visualization possibilities in both 2D [GLT*07] and 3D [GGTH07], while Sadlo and Peikert compared Lagrangian approaches to vector field topology [SP07b] and proposed a method for the adaptive computation of Coherent Structures and examined several variants of FTLE [SP07a].

The analysis of boundary flow has a long and rich tradition in fluid dynamics. A large body of both theoretical [Leg56, Dal83, CPC90, SGH06] and experimental work [Dél01] has been dedicated to the elucidation of the complex and critical flow structures that result from the interaction of an embedded object with its surrounding flow. Commonplace experimental flow visualization modalities like a thin film of oil applied to the surface of design provide a powerful means to facilitate the visual inspection of shear stress vector field patterns. Their expressive power inspired the scientific visualization community to develop techniques that possess a similar visual flavor [dLPPW95].

More recently, the research in flow visualization has been focusing on a structural characterization of the flow behavior. In particular, critical point theory and topology have received a significant attention [WTHS04, GLT*06, WTS*07b], while features of interest like flow separation and attachment manifolds were the object of a variety of

techniques [KHL94, TGS06, WTS07a]. Common to most of these approaches is their reliance on the shear flow vector field defined on the boundary surface itself. In many cases this quantity is not directly provided by the CFD simulation and must therefore be computed in post-processing, which is usually subject to some numerical difficulty.

In the following section, we proceed to introduce the basic idea of surface FTLE and discuss its application to surface and volume feature analysis.

3. Concepts

In the following, we will briefly discuss the basic concepts behind the *Finite-Time Lyapunov Exponent* and its computation and use in visualization. We keep the following presentation voluntarily informal and refer the interested reader to the publications listed in the previous section for a more comprehensive overview of this subject.

3.1. Finite-Time Lyapunov Exponents

The definition of the FTLE relies on concepts from the theory of dynamical systems. There, the *Lyapunov exponent* is defined to characterize the rate of separation of infinitesimally close trajectories as time approaches infinity. The idea behind FTLE is to apply this concept in the context of finite-time flow fields and to define a measure of coherence in terms of the trajectories of closely seeded particles.

Considering a (possibly time-dependent) vector field v , the position φ of a particle starting at position x_0 at time t_0 after advection along the vector field can be formulated as the *flow map* $\varphi(t; t_0, x_0)$ satisfying $\varphi(t_0; t_0, x_0) = x_0$ and $\left. \frac{\partial \varphi}{\partial t} \right|_t = \dot{\varphi}(t; t_0, x_0) = v(t, \varphi(t; t_0, x_0))$, where the dot denotes derivative with respect to the first parameter. Specifically, with a fixed initial time t_0 and a fixed time interval τ , one defines $t = t_0 + \tau$.

A linearization of the local variation of the flow map $\varphi(t, t_0, \cdot)$ around the seed position x_0 is given by its spatial gradient $J_\varphi(t, t_0, x_0) := \nabla_{x_0} \varphi(t, t_0, x_0)$ at x_0 . This gradient can be used to determine the maximal dispersion after time τ of particles in a neighborhood of x_0 at time t_0 as a function of the direction d_{t_0} along which we move away from x_0 : $d_t = J_\varphi(t, t_0, x_0) d_{t_0}$. Maximizing the norm $|d_t|$ over all possible unit directions d_{t_0} corresponds to computing the spectral norm of $J_\varphi(t, t_0, x_0)$. Therefore, maximizing the dispersion of particles around x_0 at t_0 over the space of possible directions around x_0 is equivalent to evaluating

$$\sigma_\tau(t_0, x_0) := \sqrt{\lambda_{\max}(J_\varphi(t, t_0, x_0)^T J_\varphi(t, t_0, x_0))}$$

where λ_{\max} is the maximum eigenvalue. To obtain the average exponential separation rate $\lambda(t, t_0, x_0)$, the logarithm is applied and the result is normalized by advection time τ to

obtain

$$\lambda(t, t_0, x_0) = \frac{1}{|\tau|} \log \sqrt{\lambda_{\max}(J_{\varphi}(t, t_0, x_0)^T J_{\varphi}(t, t_0, x_0))}.$$

This rate is then called the Finite-Time Lyapunov Exponent, and can be evaluated for both forward and backward advection. Large values of λ for forward advection correspond to stable manifolds while large FTLE values for backward advection correspond to unstable manifolds.

In the case of a fluid flow vector field, Shadden et al. [SLM05] showed that ridge lines in these fields correspond to so-called *Lagrangian Coherent Structures* that form a skeleton of the transient flow in terms of so-called *hyperbolic* material lines which dominate the overall flow structure. Therefore, analysis and visualization of these structures is a useful tool in flow visualization, and various techniques have been proposed for their direct or indirect visualization in the recent past (cf. [GLT*07, SP07b, GGTH07, SP07a]).

3.2. FTLE over Surfaces

The setting we examine in this paper is slightly different. Instead of analyzing a flow vector field in its entirety, we are especially interested in interactions of a flow with object boundaries embedded within it, with an emphasis on separation and attachment structures in the immediate vicinity of the surface. To this purpose, we adapt the concept of FTLE in the following way.

Considering an object embedded in a flow field, we denote its boundary surface by B . Assuming that B is a smooth surface in the differential geometric sense, we can define the offset surface B^ε as the set of points

$$B^\varepsilon := \{x + \varepsilon n(x), x \in B\},$$

where $n(x)$ is the surface normal to B at x . If $\varepsilon > 0$ is chosen small enough, B^ε is again a smooth surface [Her98], and B^ε is contained in the domain of definition of v . We denote the restriction of the flow map φ by φ^ε . In analogy to the definition of σ_τ above, we write

$$\sigma_\tau^\varepsilon(t_0, x_0) := \sqrt{\lambda_{\max}(J_{\varphi^\varepsilon}^\varepsilon(t, t_0, x_0)^T J_{\varphi^\varepsilon}^\varepsilon(t, t_0, x_0))},$$

with the difference that x_0 only varies in B^ε . The relation between the definitions of σ_τ^ε and σ_τ is clarified by observing that

$$J_{\varphi^\varepsilon}^\varepsilon = J_\varphi \cdot E_\varphi,$$

where the matrix E_φ projects to the local tangent space of B^ε at φ . In other words, σ_τ^ε only measures the exponential separation of integral curves originating from B^ε as the starting point varies on B^ε . Note that while the J_ε depends on the choice of local tangent basis reflected in E , σ_τ^ε does not as eigenvalues are invariant under a basis change. In the next section, we will focus on the interpretation of the values of σ_τ^ε .

Remark that while our construction is described on smooth surfaces, any surface that is only piecewise smooth must be decomposed and each of the pieces treated separately. This is a consequence of the fact that if the surface does not admit a smoothly varying tangent space, the computation of σ_τ^ε across tangent space discontinuities is neither possible nor meaningful.

3.3. Interpretation

The interpretation of σ_τ^ε is basically straightforward. Regions where this quantity is high hint at exponential separation of integral curves crossing the surface there. There are two typical cases that produce such regions on the offset surface:

Surface Separation implies locally increased values of σ_τ^ε for $\tau < 0$, since a separation line induces a local saddle-type flow structure that is intersected by the offset surface and has repelling properties in negative time (cf. Figure 1(c)). This flow pattern typically is indicative of vortex shedding, for example.

Surface Attachment induces locally maximal values of σ_τ^ε for $\tau > 0$. The saddle-type pattern on the surface in the vicinity of the attachment line is repelling in positive time (see Figure 1(b)). Most commonly, this flow pattern is the result of the flow being separated as it passes around an object.

Hence, observing the values of σ_τ^ε on an offset surface close to an object boundary can provide strong hints of these classes of flow structures. While we do not propose a sound fluid dynamical justification for this interpretation, it is confirmed by our experiments and allows reliable insight into flow structures close to or shed from object boundaries (see Section 6). In the remainder of this paper, we will abbreviate σ_τ^ε by FTLE^+ if $\tau > 0$, and FTLE^- if $\tau < 0$.

4. Implementation

In practice, a straightforward approach to computing FTLE fields consists of sampling the flow map $\varphi(t, t_0, \cdot)$ on a discretized domain and approximate its gradient numerically. In the following, we will describe the steps that we perform to this purpose for a given dataset.

Offset Surface We assume that object boundaries can be extracted from the dataset and brought into the form of a triangular surface mesh. This is true for most CFD application datasets, since the object boundaries are part of the necessarily discrete computational mesh and are usually readily extracted. Given such a mesh, we segment it along sharp features and treat each segment separately (cf. Section 3.1). We then compute surface normals and offset each of the original vertices along the normal direction according to the given ε to obtain an offset surface mesh. Necessarily, ε must be chosen such that the offset surface is not degenerate, i.e. it is bounded by the maximum curvature of the original surface.




Figure 1: Flow patterns close to the surface are reflected in integral curve convergence or divergence and hence in surface FTLE measures.

(Optional) Subdivision If the generated offset mesh does not have sufficient resolution to guarantee a good piecewise approximation of the flow map, it can optionally be subdivided to yield a higher mesh resolution. Care must be taken at this step, however, to preserve both smoothness of the surface and conformance to the original offset surface. We have selected a modified Butterfly scheme (as presented in [ZSS96]) as it is both interpolating and has a C^1 -limit surface. The simpler piecewise linear subdivision is not an option as it introduces discontinuities in the limit that reflect numerically even after a small number of subdivision steps.

Flow Map The flow map is then computed by numerical integration in the volumetric flow field with one sample (i.e. pathline) per vertex of the offset surface, according to the selected integration time τ .

FTLE Computation The resulting piecewise linear approximation of the flow map gives rise to a piecewise Jacobian per triangle of the offset mesh. Computation of the maximal eigenvalue and normalization is then straightforward, as the specific local orientation of triangles does not influence the end result of the computation (cf. Section 3.1).

The resulting algorithm is straightforward and simple to implement. We have incorporated it into the existing visualization system FAnToM, developed at the University of Leipzig, thereby both leveraging existing infrastructure (e.g. numerical integration in CFD meshes) and providing interactive access to our method to potential users.

Remark: Complex object boundaries may consist of millions of surfaces, therefore, flow map integration can represent a significant computational cost. We have performed various experiments with adaptive flow map approximation similar to the technique presented in [GGTH07]. Essentially, the flow map is approximated on a sequence of successively refined meshes, and in areas in which the flow map does not vary much, interpolation is used instead of evaluation. However, we did not obtain a significant performance gain (typically less than 10%), for the following reasons. Either the initial mesh was already fine enough to represent the flow

map in adequate detail and refinement was not necessary, or otherwise the flow map variation was quite uniform over the entire grid, making selective refinement superfluous. Similar results were obtained for adaptive approximation along the temporal axis if a sequence of FTLE fields is desired in the unsteady case. This is a direct consequence of the adaptive computational meshes that are used in most modern CFD computations. Furthermore, as discussed in Section 6, all our test cases required seconds to minutes to compute, further obviating the need for computational acceleration.

5. Visualization Techniques

Here, we will present a number of visualization techniques based on surface FTLE fields that are aimed at both automatic and user-assisted treatment of CFD datasets.

Large time-varying datasets often encompass hundreds to thousands of timesteps and represent an enormous amount of data. Processing each timestep manually is a tedious if not impossible task. Therefore, automated methods that require little or no user interaction are preferable if such datasets are to be visualized in their entirety. Furthermore, such methods guarantee reproducibility and objectivity. For example, the latter is important in the case where multiple datasets are to be compared.

Direct Visualization The $FTLE^{+,-}$ fields obtained from our approximation algorithm are readily accessible to visualization. Contrary to volumetric FTLE-based visualization, their surface nature does not suffer from visibility issues. A straightforward approach consists of depicting the object boundary in combination with a two-dimensional color map of $FTLE^{+,-}$. Before applying the color map, both quantities are normalized to a uniform range; here we first discard negative values (as they correspond to converging flow behavior which we are not interested in), and then normalize the remaining values to the interval $[0, 1]$. We have adopted the typically used color scheme depicted in Figure 1(a).

Stochastic Seeding Similar to the approach first presented

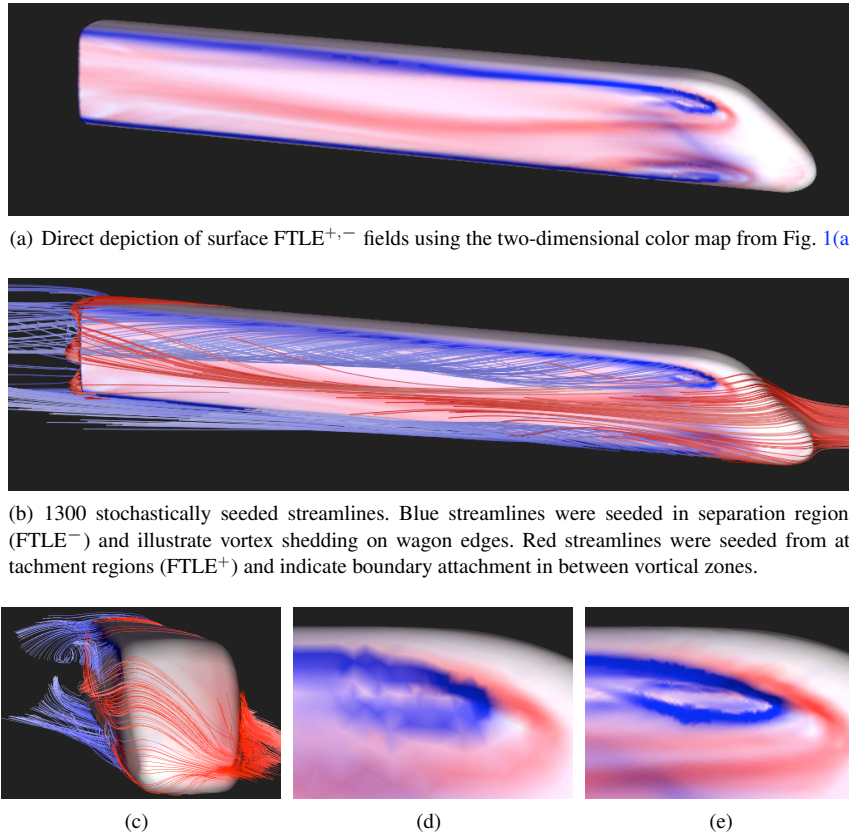


Figure 2: Direct surface FTLE visualization and stochastic streamline seeding in a high-speed train dataset.

in [GGTH07], we apply surface FTLE fields as *Probability Density Functions* (PDF) to seed integral curves in close vicinity of the boundary. Every triangle of the offset surface mesh is assigned a probability density that is proportional to its FTLE value and its area. A fixed number of seed triangles is then randomly selected according to the PDF, and a corresponding number of integral curves of prescribed length is computed and depicted. We chose positive time integration for $FTLE^-$ -based seeding, and negative time integration for $FTLE^+$ -based seeding. Obviously, this technique is not limited to FTLE fields,




Figure 4: Volumetric FTLE fields around the high-speed train nose show a strong agreement with the results of **Figure 2(a)**.

but is applicable to any scalar field defined over object boundaries. However, using FTLE fields yields especially interesting visualizations, as integral curves close to separation and attachment structures and their corresponding manifolds are typically within the selected seed set. Moreover, the uncertainty of the respective feature location as indicated by the sharpness of FTLE ridges is expressed by the fuzzy nature of randomly seeded integral curves (see e.g. Fig. 3).

While the automatic methods presented above allow computation in an offline preprocessing step and are therefore ideally suited for large datasets, there is often a need to direct the visualization to specific regions or features of interest to examine them in more detail. Furthermore, a user is uniquely able to incorporate domain-specific knowledge into the visualization process. In the following, we discuss a seeding method aimed at leveraging user feedback into surface FTLE visualization.

User-guided Seeding Curve Once surface FTLE fields have been computed, we provide an interactive tool to describe a curve on the surface (and in time in the unsteady case) by letting the user place a number of points on the

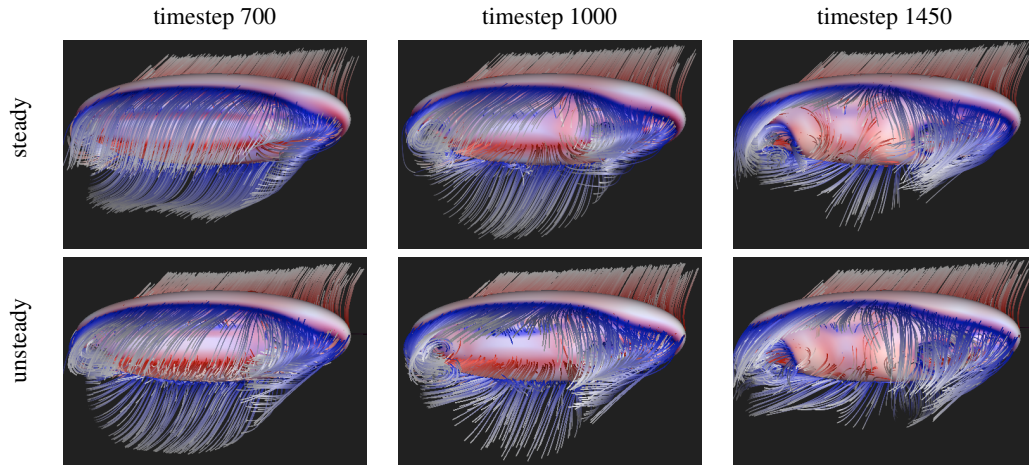


Figure 3: Comparison steady vs. unsteady surface FTLE fields and resulting flow structures illustrated by stochastic streamline/pathline seeding.

offset surface mesh on which FTLE values are depicted through a color map. The curve is then fitted through these points, and integral curves are then seeded from the curve. To increase flexibility, we allow both uniform and random seeding (optionally with a user-determined random offset) and enable a combination of the latter with the stochastic seeding described above. In this fashion, the user is able to visually identify regions of interest and examine the flow structures emanating from them. Figure 5 provides several examples.

Technically, the seeding curve is not computed through the surface points in three-space, as the resulting curve would not be constrained to the offset surface or might even leave the flow volume, and make it hard to control in general. Rather, we make use of piecewise surface parameterizations obtained using the *Least Squares Conformal Mapping* (LSCM) technique [LPRM02]. To obtain a curve on the offset surface, we map the selected points into the two-dimensional parameter space, perform curve fitting and discretization, and map the resulting points back to the offset surface. This guarantees that the curve and consequently all seed points are embedded in the offset surface. The mapping between surface and parameter space introduces only minimal distortion in the resulting curve by virtue of the LSCM parameterization.

6. Examples

In this section, we will provide a number of examples that demonstrate the discussed visualization methods on several CFD application datasets.

Fig. 2 depicts several images of a high-speed train. The corresponding stationary dataset consists of $2.7 \cdot 10^6$ unstructured elements and models the traveling train while it is

hit by strong wind from the side, causing vortices to be shed on the top on bottom edges of the wagon on the side facing away from the wind. Figures 2(a) through 2(c) show both the obtained FTLE measures and the results of stochastic streamline seeding with 1300 streamlines in total. The shedded vortices are clearly visible and are identified to emanate from separation structures (indicated by ridges of $FTLE^-$, blue). The boundary between the top and bottom vortex regions is indicated by a weak separation structure ($FTLE^+$, red). (Remark: these images were obtained without manual intervention). Figures 2(d) and 2(e) illustrate the effect of subdivision refinement on the obtained FTLE fields. Clearly, the refined offset mesh results in a much improved FTLE depiction. Computation times range from two minutes for both FTLE fields over a refined offset surface to 25 seconds for the stochastically seeded streamlines. A comparison of the obtained results with volumetric FTLE analysis of the same dataset (Figure 4, cf. [GGTH07]) shows a strong similarity of the obtained structures, indicating that our method can deliver qualitatively similar results as the much more expensive volumetric analysis.

We also applied both methods to an unsteady dataset describing a flow-embedded ellipsoid with the explicit goal of comparing steady and unsteady FTLE fields. The dataset consists of 400 timesteps over a grid of $2.6 \cdot 10^6$ elements. Figure 3 shows the steady and unsteady FTLE on the surface and 2000 streamlines/pathlines for three different time steps of the flow around the ellipsoid. Due to the slowly changing flow behavior the FTLE depictions look quite similar. The difference of streamlines and pathlines integrated for a sufficiently long time, however, reveals the unsteady character of the flow. Seeding the lines takes less than one minute and even the computation of the FTLE reaches its peak for the unsteady case with six minutes of computation time for one step.

The steady BMW dataset (Figure 5) results from a simulation of the air flow around a car and illustrates the application of manual seeding; the surface geometry is quite complex and feature sizes vary strongly in magnitude, complicating an application of the stochastic seeding approach (cf. Figure 5(a)). However, using manually determined seeding curves, we were able to quickly locate a number of interesting features in this dataset by using surface FTLE depiction as a visual guide. Figure 5(c) shows such a seeding curve (green) above the front bumper. The corresponding streamlines (Figure 5(d)) seeded from the curve confirm the vortical flow behavior hinted at by the neighboring regions of strong $FTLE^+$ and $FTLE^-$. Figures 5(e) and 5(b) provide further examples of successful identification of volume flow patterns through manual seeding. The computation of the $FTLE^{+,-}$ fields took 1.3 and 1.2 minutes respectively.

For the datasets examined here, the choice of the parameters τ and ϵ turned out to not influence the result strongly, as long as they were chosen within reasonable bounds. We have generally chosen ϵ such that the subdivided offset surface follows the boundary surface closely while still remaining entirely within the flow volume. This choice was quite simple to determine in practice by trial and error. The choice of τ is generally more involved, and no immediate principle for its choice exists (see also the discussion in [GLT*07, SP07b]). However, we found that above a certain threshold that allows the streamlines or pathlines to significantly advance from the offset surface, the resulting FTLE fields did not change much aside from slowly gaining more contrast with increased integration length.

7. Discussion and Future Work

In this paper, we have applied the concept of Finite-Time Lyapunov Exponents to flow boundaries, resulting in a both conceptually and computationally feasible characterization of flow behavior in the immediate vicinity of object boundaries. Furthermore, we have provided a number of visualization concepts that are based on these surface FTLE fields and leverage integral curves to depict flow structures that interact with or are generated at flow boundaries. We have demonstrated our methods on a number of examples and documented their usefulness and robustness with respect to state-of-the-art CFD application datasets, and we have determined that our method allows qualitatively similar conclusions to fully volumetric FTLE analysis.

Some open questions remain, nevertheless. Although surface FTLE fields admit no direct characterization of Lagrangian Coherent Structures as in the plane and volumetric cases, we would nevertheless like to examine the interaction of such LCS with flow boundaries further. Technically, an extraction of ridge lines of surface FTLE fields seems feasible, however, the current piecewise-linear representation of surface FTLE fields precludes this. Furthermore, ridges might prove an ideal candidate for integral curve seeding.

Similarly, streamlines or pathlines provide intuitive flow visualization, more advanced primitives such as stream surfaces or streak lines might be seeded using our methods. Furthermore, a more comprehensive study of the parameters of our method should be attempted, and we would like to conduct a more thorough investigation into the fluid dynamical background of our approach.

Acknowledgments

The authors would like to thank Markus Rütten from German Aerospace Center (DLR) in Göttingen, as well as VA Tech Hydro and Ronald Peikert and Filip Sadlo from ETH Zürich for providing datasets and discussion. Thanks also go to 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

- [CPC90] CHONG M., PERRY A., CANTWELL B.: A general classification of three-dimensional flow field. *Physics of Fluids 2*, A (1990), 765–777.
- [Dal83] DALLMANN U.: Topological structures of three-dimensional flow separation. *AIAA Paper 83-1735* (1983).
- [Dél01] DÉLERY J.: Robert Legendre and Henri Werlé: Toward the elucidation of three-dimensional separation. *Ann. Rev. Fluid Mechanics 33* (2001), 129–154.
- [dLPPW95] DE LEEUW W. C., PAGENDARM H.-G., POST F. H., WALTER B.: *Visual simulation of experimental oil-flow visualization by spot noise images from numerical flow simulation*. Tech. Rep. DUT-TWI-95-19, Delft, The Netherlands, 1995.
- [GGTH07] GARTH C., GERHARD F., TRICOCHÉ X., HAGEN H.: Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Transactions on Visualization and Computer Graphics 13*, 6 (2007), 1464–1471.
- [GLT*06] GARTH C., LARAMEE R., TRICOCHÉ X., SCHNEIDER J., HAGEN H.: Extraction and visualization of swirl and tumble motion from engine simulation data. In *Proceedings of The Topology-Based Methods in Visualization Workshop* (2006).
- [GLT*07] GARTH C., LI G. S., TRICOCHÉ X., HANSEN C. D., HAGEN H.: Visualization of coherent structures in transient 2d flows. In *Topology-Based Methods in Visualization 2007, to appear* (2007), Mathematics + Visualization, Springer.
- [GRH07] GREEN M., ROWLEY C., HALLER G.: Detection of lagrangian coherent structures in 3d turbulence. *J. Fluid Mech. 527* (2007), 111–120.
- [Hal01a] HALLER G.: Distinguished material surfaces and coherent structures in three-dimensional flows. *Physica D 149* (2001), 248–277.
- [Hal01b] HALLER G.: Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence. *Physics of Fluids 13*, 11 (2001).
- [Hal02] HALLER G.: Lagrangian coherent structures from ap-

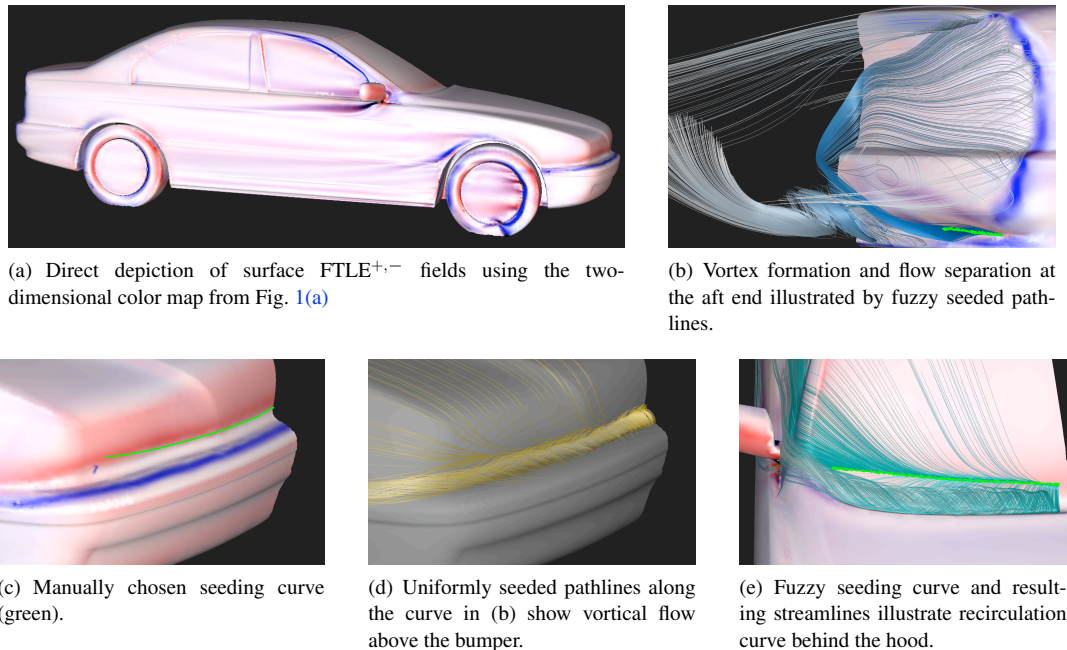


Figure 5: Manual streamline seeding on the surface of the BMW dataset.

- proximate velocity data. *Physics of Fluids* 14, 6 (June 2002), 1851–1861.
- [Her98] HERMANN T.: On the smoothness of offset surfaces. *Computer Aided Geometric Design* 15, 5 (May 1998), 529–533.
- [KHL94] KENWRIGHT D., HENZE C., LEVIT C.: Feature extraction of separation and attachment lines. *IEEE Transactions on Visualization and Computer Graphics* 5, 2 (1994), 135–144.
- [Leg56] LEGENDRE R.: Séparation de l'écoulement laminaire tridimensionnel. *La Rech. Aéro* 54 (1956), 3–8.
- [LPRM02] LÉVY B., PETITJEAN S., RAY N., MAILLOT J.: Least Squares Conformal Maps for Automatic Texture Atlas Generation. In *SIGGRAPH '02* (New York, NY, USA, 2002), ACM Press, pp. 362–371.
- [LSM06] LEKIEN F., SHADDEN S. C., MARSDEN J. E.: N-dimensional lagrangian coherent structures: Definition and properties. *Physica D* (2006).
- [MHP*07] MATHUR M., HALLER G., PEACOCK T., RUPPERT-FELSOT J. E., SWINNEY H. L.: Uncovering the lagrangian skeleton of turbulence. *Physical Review Letters* 98, 14 (2007), 144502.
- [PC86] PERRY A., CHONG M.: A series expansion study of the navier-stokes equations with applications to three-dimensional separation patterns. *Journal of Fluid Mechanics* 173 (1986), 207–223.
- [SDM06] SHADDEN S., DABIRI J., MARSDEN J.: Lagrangian analysis of fluid transport in empirical vortex rings. *Physics of Fluids* 18, 4 (2006), 047105.
- [SGH06] SURANA A., GRUNBERG O., HALLER G.: Exact theory of three-dimensional flow separation. part i: Steady separation. *J. Fluid Mech.*, in press, available at <http://web.mit.edu/ghaller> (2006).
- [SLM05] SHADDEN S. C., LEKIEN F., MARSDEN J. E.: Definition and properties of lagrangian coherent structures from finite-time lyapunov exponents in two-dimensional aperiodic flows. *Physica D* 212 (2005), 271–304.
- [SP07a] SADLO F., PEIKERT R.: Efficient visualization of lagrangian coherent structures by filtered AMR ridge extraction. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1456–1463.
- [SP07b] SADLO F., PEIKERT R.: Visualizing lagrangian coherent structures: A comparison to vector field topology. In *Topology-Based Methods in Visualization 2007*, to appear (2007).
- [TGS06] TRICOCHÉ X., GARTH C., SCHEUERMANN G.: Fast and robust extraction of separation line features. In *Scientific Visualization: The Visual Extraction of Knowledge from Data* (2006), Bonneau G.-P., Ertl T., Nielson G. M., (Eds.), Mathematics + Visualization, Springer.
- [WTHS04] WEINKAUF T., THEISEL H., HEGE H.-C., SEIDEL H.-P.: Boundary switch connectors for topological visualization of complex 3d vector fields. In *VisSym 2004 : Joint Eurographics/IEEE Symposium on Visualization* (Konstanz, Germany, 2004), Deussen O., Hansen C., Keim D. A., Saupe D., (Eds.), Eurographics, pp. 183–192.
- [WTS07a] WIEBEL A., TRICOCHÉ X., SCHEUERMANN G.: 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] WIEBEL A., TRICOCHÉ X., SCHNEIDER D., JANICKE H., SCHEUERMANN G.: Generalized streak lines: Analysis and visualization of boundary induced vortices. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1735–1742.

[ZSS96] ZORIN D., SCHRÖDER P., SWELDENS W.: Interpolating subdivision for meshes with arbitrary topology. In *SIG-GRAPH '96* (New York, NY, USA, 1996), ACM, pp. 189–192.