EUROGRAPHICS - IEEE VGTC Symposium on Visualization (2005)
K. W. Brodlie, D. J. Duke, K. I. Joy (Editors)

Localized Flow Analysis of 2D and 3D Vector Fields

Alexander Wiebel[1], Christoph Garth[2], Gerik Scheuermann[1]

[1]Computer Science Institute, University of Leipzig
[2]Computer Graphics Group, University of Kaiserslautern

Abstract

In this paper we present an approach to the analysis of the contribution of a small subregion in a dataset to the global flow. To this purpose, we subtract the potential flow that is induced by the boundary of the sub-domain from the original flow. Since the potential flow is free of both divergence and rotation, the localized flow field retains the original features. In contrast to similar approaches, by making explicit use of the boundary flow of the subregion, we manage to isolate the region-specific flow that contains exactly the local contribution of the considered subdomain to the global flow. In the remainder of the paper, we describe an implementation on unstructured grids in both two and three dimensions. We discuss the application of several widely used feature extraction methods on the localized flow, with an emphasis on topological schemes.

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

1. Introduction

In the usual analysis and visualization of CFD flows, topological methods have an established role and are in wide use for typical flow analysis tasks (e.g. [BTPS04]). Under certain conditions, however, this class of methods fails to provide an accurate structural picture of non-trivial flows. The notion of topology of a vector field is built on the occurrence of critical points, i.e. zeros of the field. If, for example, the flow is dominated by a large near-constant component, as is common in the flow past a stationary object, critical points do often not occur at all. Moreover, the definition of vector field zeros implies that they are not invariant with respect to a moving frame of reference. Due to practical reasons, the actual frame of reference of a given flow vector field does not necessarily match the one of interest. This is especially true for physical measurements that are impossible to obtain for certain geometries (e.g. a moving helicopter rotor), but also for CFD computations where practical reasons mandate that moving objects are simulated on a fixed grid. While topology is a prime example of a flow visualization method that is dependent on the underlying frame of reference, other schemes suffer from this limitation to a varying degree.

The ideas presented in this paper center around the notion of *localized flow analysis*, i.e. the analysis of the contribution in a subregion to the global flow of a given dataset. To this purpose, we construct a so-called *potential flow* that matches the original field on the boundary of the sub-domain but is otherwise simple in the sense that it has vanishing divergence and curl. In other words, it represents the laminar flow in the sub-domain. Through subtraction of this field from the original flow, we are left with a *localized flow* that is confined to the sub-domain under consideration and contains the local contribution to the global flow. As an important side-effect, the influence of the frame of reference that is reflected on the boundary of the sub-domain is removed. In this way, topological analysis and frame of reference are separated. Visualization methods that are based on divergence or rotation of the flow (both local in nature) are unaffected by this approach since the localized flow retains the original rotation and divergence. The method presented here works well for both two-dimensional and three-dimensional flow fields. The choice of subregion is arbitrary up to the condition that it is a simple domain with piecewise smooth boundary. We describe an algorithm that reflects the given ideas on unstructured triangular or tetrahedral meshes using a standard finite-element approach. Although the computation of the potential flow is a complicated numerical procedure, our algorithm works well even on large CFD datasets with millions of cells.

Our work can be seen as having similarities to what others

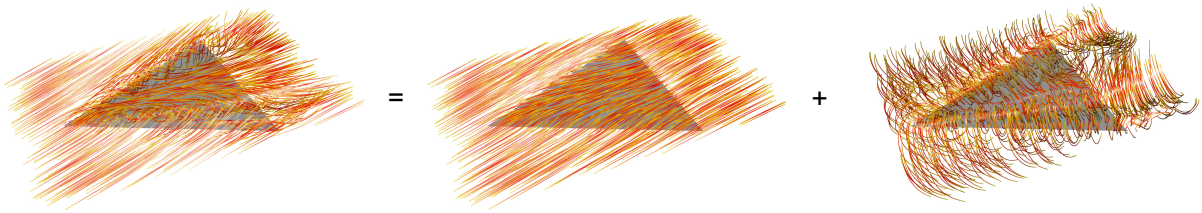

Figure 1: *Illustration of the different components of the flow around the delta wing.* ***Left****: Original flow field from the CFD simulation,* ***Middle****: Laplacian field computed by the original field's flow normal to the boundary. Note that the flow is simple but not constant.* ***Right****: Localized or region-specific flow obtained by subtracting the Laplacian field from the original field.*

have published before (cf. [PP00, PP03, TLHD03], therefore we describe some essential differences to the work presented here in Section 2 as well as other work that is related to this paper. In Section 3, we sketch out the mathematical concepts that the region-specific flow is based on and give a detailed discussion on its usefulness in dataset analysis. The implementation on triangular and tetrahedral grid is the topic of Section 4. Finally, we discuss the application of topological methods on some examples in Section 5, as well as other particular visualization and feature extraction methods. Section 6 concludes on the presented work.

2. Related Work

The notion of localized flow analysis under preservation of the original characteristics of the field (i.e. divergence and rotation) is in part related to work published by Polthier and Preuss [PP00, PP03] (in 2D) and Tong *et al.* [TLHD03] (in 3D). These authors employ the Hodge decomposition theorem from vector analysis, stating that any vector field can be decomposed into three fields containing the divergence, rotation and harmonic parts. The decomposition is given in terms of potentials for the divergence- and rotation-components:

$$v = \operatorname{grad} u + \operatorname{curl} w + h,$$

which are explicitly computed. Analysis is then attempted by locating features as extremal points of the first two components. Although these approaches seem quite similar to what we describe here, our motives and technique are different. It is our aim to analyze the localized flow with conventional flow visualization techniques, as opposed to making use of the potentials for that purpose. Moreover, in spite of the superficial similarity between the potential flow and the harmonic field h from above, we believe that our approach is better suited to the localized analysis of flow since we use specific boundary conditions to guarantee that the potential flow contains the part of the flow that does not originate in the considered domain. No such condition is imposed on h. In addition, little is known as to how existing visualization methods are affected by the Hodge decomposition, a topic we discuss and examine in detail in Sections 3.3 and 5. Last but not least, the computation of the potential flow is conceptually simpler than that of u and w, as only one potential and this only of scalar nature has to be computed. We therefore feel that our contribution is significantly different from [PP00, PP03, TLHD03].

Concerning topological analysis and feature extraction of vector fields, there is a large body of literature available. [PVH*03] provides a good overview. Of special interest in this paper are especially topological methods as treated by many authors, e.g. Helman and Hesselink [HH89], Globus [GLL91], Scheuermann *et al.* [SKMR98], Tricoche *et al.* [TWSH02] and Theisel *et al.* [TWHS03] to name just a few. We are also concerned with more general feature extraction methods, such as the vortex core line extraction method of Sujudi and Haimes [SH95] and the region-based λ_2-criterion by Jeong and Hussain [JH95] that we discuss in the context of the localized flow.

3. Localized Flow Analysis

In the following, let $v : \mathbb{R}^d \to \mathbb{R}^d, d = 2,3$ be a continuous (flow) vector field. Let $\Omega \in \mathbb{R}^d$ be an open, bounded and connected domain and n the outward normal field on $\partial\Omega$.

In order to analyze the specific contribution of the flow in Ω to the global flow field, we define the *region-specific flow* $v_R : \Omega \to \mathbb{R}^d$ by requiring two essential conditions:

1. it retains the essential behavior of the flow in terms of rotation and divergence, i.e.

$$\operatorname{div} v_R = \operatorname{div} v \quad \text{and} \quad \operatorname{curl} v_R = \operatorname{curl} v \qquad \text{on } \Omega.$$

2. it is isolated from the global flow on the boundary of the sub-domain, i.e. the region-specific flow through the boundary vanishes:

$$v_R \cdot n = 0 \quad \text{on } \partial\Omega.$$

The suitability of these conditions is discussed in more detail in 3.3. The difference of global and region-specific flow

is then given by

$$v_P := v - v_R.$$

Owing to the linearity of divergence and curl, v_P must satisfy

$$\operatorname{div} v_P = 0 \quad \text{and} \quad \operatorname{curl} v_P = 0 \quad \text{on } \Omega, \tag{1}$$

and we find that

$$v_P \cdot n = v \cdot n \quad \text{on } \partial\Omega. \tag{2}$$

We next look at how the construction of v_P can be achieved by a simple mathematical procedure.

3.1. A Special Neumann Problem

Let us assume that v_P is given as the gradient of a function $u : \Omega \to \mathbb{R}$ (then v_P is called *potential flow*). It is immediate that

$$\operatorname{curl} v_P = \operatorname{curl} \operatorname{grad} u = 0 \quad \text{on } \Omega.$$

Requiring that v_P has vanishing divergence, we compute

$$0 = \operatorname{div} v_P = \operatorname{div} \operatorname{grad} u = \Delta u \quad \text{on } \Omega,$$

where Δ denotes the Laplace operator on scalar functions. Rewriting Eq. (2) in terms of u, it turns into

$$n \cdot \operatorname{grad} u = v \cdot n \quad \text{on } \partial\Omega.$$

Hence, for $v_P := \operatorname{grad} u$ to fulfill the conditions (1) and (2), u must solve

$$\Delta u = 0 \qquad \text{on } \Omega \tag{3}$$

$$n \cdot \operatorname{grad} u = v \cdot n \qquad \text{on } \partial\Omega \tag{4}$$

This class of problem is called a Neumann-Laplace problem for u, and it is solvable up to a constant under the condition that

$$\int_{\partial\Omega} v \cdot n = 0, \tag{5}$$

i.e. the total flow through the boundary must vanish (for details on this the reader is referred to [Hac92]). From this construction, we are able to determine v_P by solving for u. The fact that u is only determined up to a constant does not matter here, since the constant vanishes after taking the gradient. However, to guarantee uniqueness of u and make the problem well defined, a condition of the form

$$\int_{\Omega} u = 0$$

can be imposed. Finally, the region-specific flow is then given via

$$v_R := v - \operatorname{grad} u.$$

The solvability condition (5) is by definition fulfilled for incompressible flows (e.g. liquid flow), since by Stokes' theorem

$$\int_{\partial\Omega} v \cdot n = \int_{\Omega} \operatorname{div} v = 0.$$

In the next section, we detail a modification of the Neumann problem for the case of compressible flows.

3.2. A Modification for Compressible Flows

When considering compressible flows, e.g. those arising as solutions of the full Navier-Stokes equations, the solvability condition (5), by the divergence theorem, needs not hold. However, compressible flows satisfy the conservation of momentum law which reads

$$\operatorname{div} \rho v = 0, \tag{6}$$

where $\rho > 0$ denotes a material density that may vary spatially. Based on this, we are able to enhance our approach from above to guarantee results for compressible flows. The modified Neumann problem then reads

$$\Delta u = 0 \qquad \text{on } \Omega \tag{7}$$

$$n \cdot \operatorname{grad} u = (\rho v) \cdot n \qquad \text{on } \partial\Omega \tag{8}$$

The solvability condition for this system coincides with Eq. (6) and is hence fulfilled. Then, v_P is again divergence- and curl-free and by setting

$$\rho v_R := \rho v - \operatorname{grad} u$$

it follows

$$n \cdot \rho v_R = n \cdot (\rho v - \operatorname{grad} u) = 0 \quad \text{on } \partial\Omega.$$

Dividing by ρ we find

$$n \cdot v_R = 0 \quad \text{on } \partial\Omega,$$

i.e. the region-specific flow is again confined to Ω and inherits the original flow's characteristics.

3.3. Interpretation of the Region-Specific Flow

Until now we have only considered the mathematical construction of the region-specific flow. Along the way, some conditions were imposed to guarantee solvability of the problem. We will now dedicate some thoughts as to how these conditions affect applicability of our method to the general localized analysis of flows. More specific results and application examples are given in Section 5.

From a purely physical point of view it does not seem feasible at first glance to manipulate a flow field in order to further its analysis. It is known practice, however, to subtract a constant vector field to reveal structures that are not visible in the original field (heuristically, the average (boundary) flow is subtracted). This is justified by the principle of Galilean invariance which states that the properties of flow have to be the same for a constantly moving and for a resting observer. However, in most cases, this approach is not

appropriate as it does not preserve boundary conditions. For example, in the flow around a stationary object, subtraction of a constant vector field yields streamlines that are not tangential to the enclosed surface. Since the boundary conditions are an integral part of the region-specific flow (via conditions (1) and (2)), it does not suffer these problems.

Furthermore, both vorticity and divergence of the original flow are preserved in the region-specific approach. Therefore, feature definitions that build on these quantities and consequently algorithms that extract these features are naturally unaffected. Recently, work has been published treating analysis and visualization of three-dimensional vector fields based on vorticity and vorticity lines (see [SPP04]). While the streamlines of the velocity are naturally different in the region-specific flow, the invariance of vorticity lines and hence the non-changing vorticity transport in the flow imply that all vortical structures are kept. This confirms our approach to be meaningful and to contain the information for the important features present in the original field. In summary, the region-specific flow contains exactly the local domain-specific contribution to the global flow.

Concerning the influence of the frame of reference on flow analysis, the region-specific flow delivers a natural abstraction. The *correct*frame of reference might not be known a-priori or hard to construct, as is often the case in the analysis of measurement datasets involving geometries moving relative to the observer. For the common case that features in the flow are obscured by a dominating through-flow, the influence of the latter is *caught*in the boundary conditions of the potential flow, even if it is non-constant. It is subsequently subtracted from the original flow and does not show up in the region-specific flow. Through this, for the case of topological methods, critical points such as sinks, sources and spirals relating to extremal divergence and vorticity are much more likely to occur than in the original flow, enabling the use of such methods in a broader context of flow analysis.

Figure 2 exemplifies some of these considerations. The 2D vector field shown represents the incompressible flow passing around a cylinder. On the downstream side of the cylinder, the well known Karman vortex street should develop. However, it cannot be observed in the original flow. Removing the average flow reveals some but not all of the features present and yields a strong diagonal flow component that has no physical interpretation. The constructed potential flow is very uniform except in the vicinity of the cylinder where it reflects the flow around it. Subtracting the computed flow from the original flow reveals all the downstream vortical structures through a topological analysis.

3.4. Constraints

The computed potential flow is simple in the sense that it is irrotational and solenoidal. Although it is still possible for saddle points to occur, this does not play a role in the analysis of the region-specific flow.

a) original flow

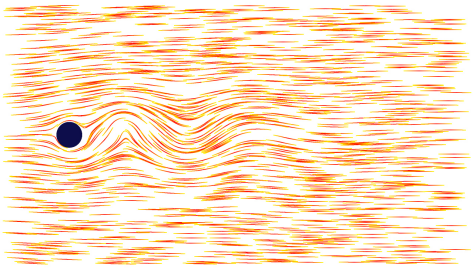

b) original flow minus average flow

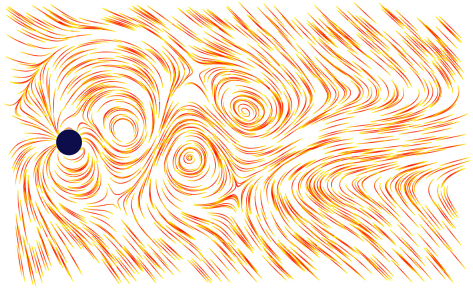

c) potential flow

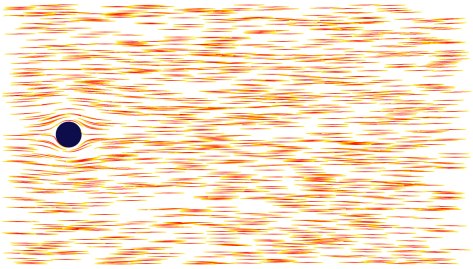

d) original flow minus potential flow

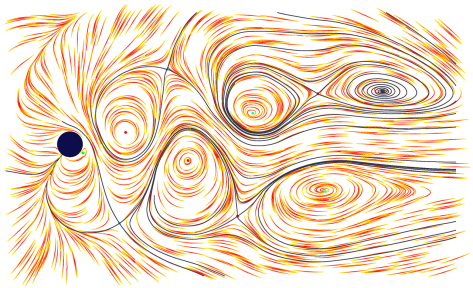

Figure 2: *Comparison of different fields obtained from cylinder dataset with Karman vortex street.* ***a)*** *Streamlines in the original flow. Only sinuous line structures give hints on the vortices.* ***b)*** *Three vortices revealed by removing average flow.* ***c)*** *Potential flow induced by the flow on the boundary of the considered region. Note how the flow attaches to the cylinder and does not seem to cross it as it would be the case for constant flow.* ***d)*** *Subtracting the Laplacian field reveals all five vortices present in the considered region by use of topology.*

In the preceeding discussion, we did not reflect on the choice of the subdomain Ω that is subject to the localized analysis. From the mathematics of solving the special Neumann Problem 3, it is only required that Ω be open, bounded and connected. These requirements are easily fulfilled and do not constrain the choice of region much. Regarding the numerical schemes we apply in the application of our ideas in the next section, a convex domain with piecewise smooth boundary is greatly beneficial in terms of convergence.

4. Implementation

The construction of the region-specific flow for discrete datasets is straightforward. In the following, we will revisit the steps of Section 3 and discuss how to implement them.

We assume that the discrete flow field is given on the vertices of a triangular or tetrahedral grid. This is a form that most modern CFD datasets are either supplied in or easily brought into. It should be noted here again that the derivation of the potential flow is basically the same in any spatial dimension. In the implementation, however, differences show up since the method works on triangles in the two-dimensional case and on tetrahedra in three dimensions. By formulation in the context of finite element methods, a unified numerical formalism can be achieved nevertheless.

The region Ω is easily discretized as a subset of the grid simplices that forms a connected set. The Neumann problem is then discretized on this set by the application of a Galerkin-type finite element method. The basic idea is simple (cf. [ACF99] for a very concise presentation) and results in a linear system of equations $Ax = b$ with a sparse symmetric positive-definite matrix A. The system can then be solved using a standard iterative technique such as the preconditioned conjugate gradient scheme, yielding a discrete potential u on the vertices of the grid. Owing to the finite element approach, the dimension of the system matrix A is equal to the number of grid points n. Therefore, A is best represented in a sparse storage format. This allows us to treat grids with millions of cells without resorting to out-of-core or cluster techniques which makes the implementation straightforward. Computational complexity is two-fold: the assembly of the system matrix A is relatively costly since the complexity is linear in the number of grid cells. The complexity of the successive matrix inversion is then again a function of n, depending on the dimension.

The implementation of the boundary condition deserves a more specific description: since the outer normal field of the grid is (depending on the dimension) edge- or face-based, we use averaging of adjacent edges or faces to compute the appropriate boundary condition. Having obtained the discrete potential u, taking its gradient gives a cell-wise constant vector field. Again, we use weighted averaging of neighboring simplices to compute the vector field values of v_P on the vertices of the grid. Finally, v_R is obtained by subtracting v_P from v at the grid vertices.

Concerning the solvability condition (5) of the Neumann problem, in our experiments we found that due to discretization errors of the underlying CFD computation, the condition is sometimes not fulfilled exactly, hindering convergence of the matrix inversion. In our experience, it is best to settle for a least squares solution to the linear system in these cases. While these do only approximate the requirements of Section 3 they can still be used to perform localized flow analysis. The validity of Eq. (5) is easily checked by summing up the orthogonal velocity components over the boundary.

In discretizing a continuous problem, there is always the question of which grid is suitable for approximation. In our case, the grid is given as a subset of the computational grid of the CFD simulation, therefore we are limited to its precision. A remeshing is prohibitively complicated and expensive, except for very simple geometries. In our experiments, however, we found the adaptive-resolution grids typical for CFD datasets well suited to our purpose.

Overall, the algorithm is easily implemented with the help of broadly available algorithms or libraries (cf. the Matrix Template Library with its companion Iterative Template Library).

5. Results and Examples

In this section we present some examples in the form of CFD application datasets that we have applied the localized flow analysis to. All datasets were treated on a consumer-type PC workstation with 3 gigabytes of RAM. Two of the datasets (the cylinder and the vortex generator) are rather academic in nature, but the delta wing configuration and the HART II datasets stem from actual application projects.

5.1. Delta Wing

As non-artificial example we studied a CFD simulation of airflow around a delta wing. The simulation was performed over 1000 time steps that show the genesis and burst of the two main vortices (see fig. 1) over time. We picked out one time step to test our method. The field consists of roughly 3 million vertices and 18.1 million unstructured tetrahedra. The region for the localized flow analysis was chosen as a box around the wing, since the dataset contains no other interesting areas. As is typically for CFD data, the resolution is much higher in highly turbulent regions, in our example near the wing. Unfortunately, our box thus contains still more than 17.3 million tetrahedra.

The results of subtracting the potential flow from this dataset can be seen in figure 1. To check the benefits of our approach, we also removed an average flow field for tests. This yielded no interesting results. The reason for this is in part to be blamed upon the large vortices above the wing that disturb the average computation for the full dataset. Subtracting only the average of the boundary yielded much

better results. But as described in Section 3.3 removing a constant component does not respect the boundary conditions of the original flow. This was striking on the surface of the delta wing. While the vortices formed above the wing in the region-specific flow, they seemed to form around and through the wing for the flow reduced only by a constant field.

Studying the localized flow, we observed in some regions around the wing a similarity between its streamlines and streamlines in the gradient field of the pressure. This is especially emphasized at the tip of the delta wing. Looking at the Navier-Stokes equations, these similarities lead to the idea that a major part of the local flow is induced by the pressure distribution, i.e. by the acceleration produced by the pressure gradient. While this interpretation applies to the flow near the tip, it is not convincing in the rear part of the wing where we found only weak similarities. Considering the Navier-Stokes equations again yields the solution. Apart from the pressure gradient there are additional forces resulting in the stress induced by the no slip boundary on the surface of the wing. At the tip of the wing the flow has not yet passed a large part of the surface, the influence of the stress is small there and thus the pressure gradient is the driving force in the local flow. Moving along the surface to the back of the wing the effect of the pressure gradient mixes with the growing stress influence and the similarities disappear.

The topology of the flow appears to be simple but meaningful. While the original vector field contains no interesting critical points, there appear two singular points in the localized flow where the main vortices cross the borders of the sub-domain. The position of these two 3D-saddles are marked in the right image of figure 5. The two one-dimensional separatrices starting from the two saddles are also shown in figure 5. Both separatrices almost exactly match the vortex core lines of the main vortices in the original flow. This means that, in the localized flow, we are able to detect the vortex core lines just by extracting the topology. This should be true for other datasets as well.

Applying standard vortex detection methods on the region-specific flow yields results that are nearly the same as in the original flow. As can be verified in the left image in figure 5 all six vortices, the two main vortices and the secondary and tertiary vortices, are present in the localized flow. This confirms the argument that these feature extraction methods built on divergence and rotation are not affected by localizing the flow.

5.2. HART II

Our second realistic dataset stems from the HART II test [vdWBY*04] for measuring helicopter rotor wakes. The aim of the test is to improve the knowledge about the evolution of the vortex generated by a moving rotor blade. This is important because the vortex location relative to the following blade is crucial for reducing rotor noise which, in fact, is created by interaction of a wake and a following blade hitting the wake. PIV (Particle Image Velocimetry) was used to obtain instantaneous flow field data in a large observation area and in a smaller close-up view of the vortex core. We consider a 2D vector field from such a close-up view. The field has 8316 quadratic cells consisting of 8500 vertices.

We tested our method for this dataset because the correct frame of reference is not known to us. As can be seen in the left image of figure 3 without working on the field no vortical structure can be found at all. A small part of the vortical structure can already be revealed removing the average of the field, but is much clearer when removing the potential flow induced by the boundary (see right image of fig. 3).

5.3. Vortex Generator

The Vortex generator dataset clearly is an academic example. It has two blade-shaped obstacles placed in the dataset to generate vortices that are very explicit and easy to handle. There appear no problems in detecting the vortex core lines. We present this dataset nevertheless, as it nicely shows how the localized flow emphasizes the structure of the flow in the considered region. Figure 4 shows extremely clear where the vortices have their start point and how the obstacles are the origin for the important parts of the local flow structure.

5.4. Performance

For most of our datasets the whole procedure of cutting, solving, computing the derivatives and subtracting the obtained potential flow field took only a few minutes. The subdomain of the delta wing, with its over 17 million cells was the exception. Computing the solution in this case took 75 minutes. Much of this time was spent for the assembly of the matrix and for incorporating the boundary conditions.

5.5. Choice of Region

Choosing the region in a good way lies in the hand of the user. It is up to him to decide which parts may be interesting without the flow induced by the boundaries. However, as the resulting field is forced to be parallel to the boundary, some care has to be taken. The region should not be to small, for example its boundary should not cut through a feature the user wants to emphasize. Cutting such a feature results in biasing of the feature in the localized flow. Keeping some distance from interesting features to the boundary can enhance the quality of the localized flow.

A good example showing that different choices can be useful is the delta wing. We chose the whole wing to get an overall picture, but it is also interesting to consider a cylinder around one of the vortices or a region at the wing apex. Indeed, choosing a box at the apex revealed the coincidence of the localized flow with the pressure gradient, mentioned earlier in this section.

6. Conclusion and Outlook

We have introduced a method to isolate the flow in sub-domains of flow datasets from the flow in the neighborhood by constructing a Laplacian field from the flow at the boundary of the sub-domain and subtracting it from the original field. The method retains the original features of the flow and is thus ready for the application of many standard methods for feature and topology analysis. We discussed the differential equation that has to be solved to obtain the Laplacian fields and described our implementation. Applied to a large dataset from CFD simulations and to a measurement dataset our method proved to be scalable and robust.

Finally, it is left mentioning that discussions with engineers showed that our concept is simple to understand for them since potential flow fields are well known.

As we have just begun to study the flow resulting from our method, there are many avenues of future research:

- The ideas given here are strongly tied to the study of flow fields. We would like to investigate in how far they can be applied to the study and visualization of other vector fields.
- By now we have only studied the change of the most common flow field quantities (e.g. vorticity) when subtracting potential flow. How other quantities as helicity or the derivatives of the flow are influenced are still open questions.
- Having no sinks or sources in the region-specific flow the streamlines have to be closed. Since the vectors at the boundary have to be tangential to it, all streamlines have to stay in the region. This leads to the idea to study how bundles of closed streamlines form knots and loops and how they are knotted.
- In more complicated cases where the potential flow field has saddles it may be interesting to not only study the localized flow but also the subtracted flow.

Acknowledgments

The authors wish to thank Markus Rütten, Hugues Richard and Berend van der Wall from German Aerospace Center in Göttingen and Braunschweig for providing the datasets. We would especially like to thank Xavier Tricoche for many constructive remarks and interesting discussions on the topic. Thanks also go to all members of the FAnToM development team for their programming efforts. This work was partly supported by DFG grants HA 1491/15-4 and HA 1491/15-5.

References

[ACF99] ALBERTY J., CARSTENSEN C., FUNKEN S. A.: Remarks around 50 lines of Matlab: short finite element implementation. *Numerical Algorithms 20* (1999), 117 – 137. 5

[BTPS04] BROWN J. L., TOBAK M., PRABHU D. K., SANDSTROM T. A.: Flow Topology About An Orbiter Leading Edge Cavity At STS-107 Reentry Conditions. *Proceedings of the 37th AIAA Thermophysics Conference* (2004). 1

[GLL91] GLOBUS A., LEVIT C., LASINSKI T.: A Tool for Visualizing the Topology of Three-Dimensional Vector Fields. In *IEEE Visualization Proceedings* (October 1991), pp. 33 – 40. 2

[Hac92] HACKBUSCH W.: *Elliptic Differential Equations : Theory and Numerical Treatment*. No. 18 in Springer Series in Computational Mathematics. Springer, 1992. 3

[HH89] HELMAN J., HESSELINK L.: Representation and Display of Vector Field Topology in Fluid Flow Data Sets. *IEEE Computer 22*, 8 (August 1989), 27 – 36. 2

[JH95] JEONG J., HUSSAIN F.: On the Identification of a Vortex. *Journal of Fluid Mechanics 285* (1995), 69 – 94. 2

[PP00] POLTHIER K., PREUSS E.: Variational Approach to Vector Field Decomposition. In *Proc. of Eurographics Workshop on Scientific Visualization* (2000), van Liere R., Post F.,, et. al., (Eds.), Springer Verlag. 2

[PP03] POLTHIER K., PREUSS E.: Identifying Vector Field Singularities Using a Discrete Hodge Decomposition. In *Visualization and Mathematics III*, Hege H. C., Polthier K., (Eds.). Springer Verlag, 2003, pp. 113 – 134. 2

[PVH*03] POST F. H., VROLIJK B., HAUSER H., LARAMEE R. S., DOLEISCH H.: The State of the Art in Flow Visualization: Feature Extraction and Tracking. *Computer Graphics Forum (Blackwell CGF) 22*, 4 (2003), 775 – 792. 2

[SH95] SUJUDI D., HAIMES R.: Identification of swirling flow in 3-d vector fields. In *12th AIAA CFD Conference* (San Diego CA, June 1995). 2

[SKMR98] SCHEUERMANN G., KRÜGER H., MENZEL M., ROCKWOOD A. P.: Visualizing nonlinear vector field topology. *IEEE Trans. Vis. Comput. Graph. 4*, 2 (1998), 109–116. 2

[SPP04] SADLO F., PEIKERT R., PARKINSON E.: Vorticity Based Flow Analysis and Visualization for Pelton Turbine Design Optimization. In *Proceedings of IEEE Visualization '04* (October 2004), Turk G., van Wijk J. J.,, Moorhead R., (Eds.), IEEE, pp. 179 – 186. 4

[TLHD03] TONG Y., LOMBEYDA S., HIRANI A., DESBRUN M.: Discrete Multiscale Vector Field Decomposition. *ACM Transactions on Graphics (TOG) 22*, 3 (July 2003), 445 – 452. 2

[TWHS03] THEISEL H., WEINKAUF T., HEGE H.-C., SEIDEL H.-P.: Saddle Connectors - An Approach to Visualizing the Topological Skeleton of Complex 3D Vector Fields. In *IEEE Visualization '03* (2003), Rushmeier H., Turk G.,, van Wijk J. J., (Eds.), pp. 225 – 232. 2

[TWSH02] TRICOCHE X., WISCHGOLL T., SCHEUERMANN G., HAGEN H.: Topology Tracking for the Visualization of Time-Dependent Two-Dimensional Flows. *Computers & Graphics 26*, 2 (2002), 249 – 257. 2

[vdWBY*04] VAN DER WALL B. G., BURLEY C. L., YU Y., RICHARD H., PENGEL K., BEAUMIER P.: The HART II test - measurement of helicopter rotor wakes. *Aerospace Science and Technology*, 8 (2004), 273 – 284. 6

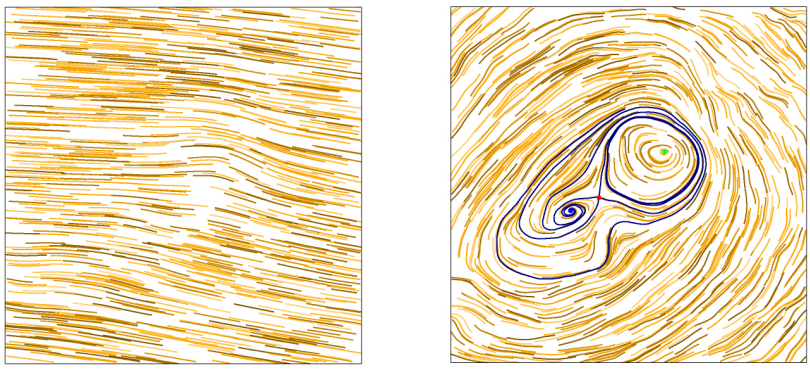

Figure 3: *HART II dataset consisting of PIV measurements. Left image shows the original measured flow with no visible features. On the right, the topology of the region-specific flow reveals the vortices present in the correct frame of reference*

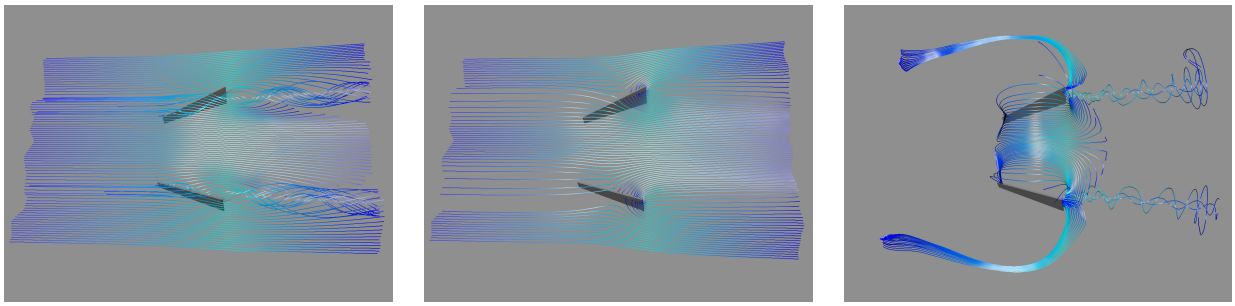

original flow | potential flow | original flow minus potential flow

Figure 4: *Original, potential and region-specific flow around two obstacles for vortex generation. The local contributions to the flow and the locations where the vortices begin are easily identified in the right image.*

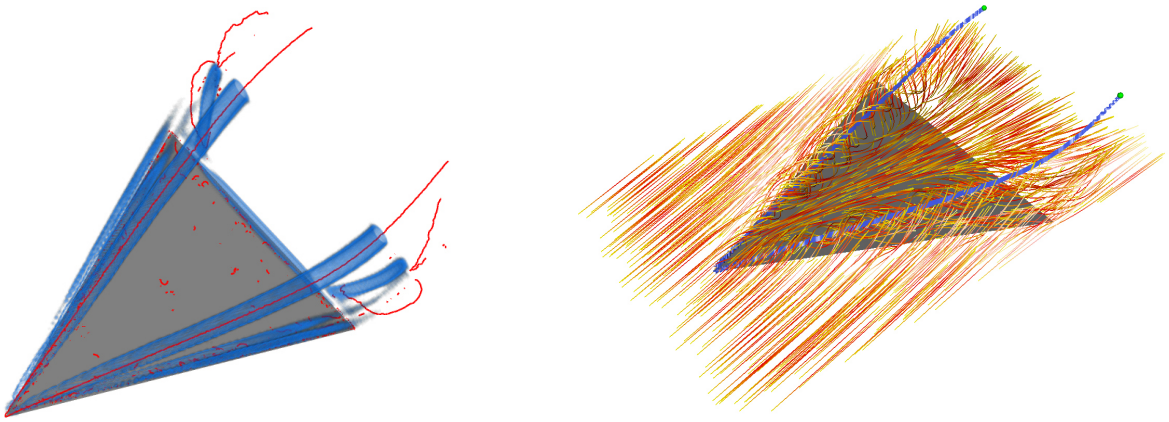

Figure 5: ***Left:****Vortex core lines and volume rendering of the* λ_2*-criterion, both computed for the localized flow around the delta wing. The vortex core lines were extracted using the algorithm of Sujudi and Haimes.* ***Right:*** *Streamlines and topology in the localized field of the delta wing dataset. Two 3D-saddle points are shown as small spheres. 1D separatrices are integrated from the saddles.*