

# Topology-based Smoothing of 2D Scalar Fields with $C^1$ -Continuity

Tino Weinkauff, Yotam Gingold and Olga Sorkine

Courant Institute of Mathematical Sciences, New York University, USA

---

## Abstract

*Data sets coming from simulations or sampling of real-world phenomena often contain noise that hinders their processing and analysis. Automatic filtering and denoising can be challenging: when the nature of the noise is unknown, it is difficult to distinguish between noise and actual data features; in addition, the filtering process itself may introduce “artificial” features into the data set that were not originally present. In this paper, we propose a smoothing method for 2D scalar fields that gives the user explicit control over the data features. We define features as critical points of the given scalar function, and the topological structure they induce (i.e., the Morse-Smale complex). Feature significance is rated according to topological persistence. Our method allows filtering out spurious features that arise due to noise by means of topological simplification, providing the user with a simple interface that defines the significance threshold, coupled with immediate visual feedback of the remaining data features. In contrast to previous work, our smoothing method guarantees a  $C^1$ -continuous output scalar field with the exact specified features and topological structures.*

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Geometric algorithms, languages, and systems

---

## 1. Introduction

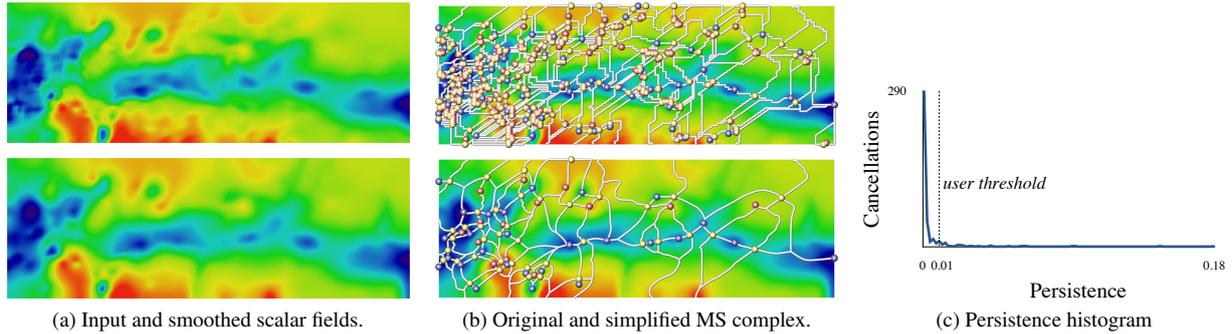
Data sets encountered in scientific computing, computer graphics, medicine and many other fields often contain noise, originating from errors in the acquisition process of real-world phenomena or from numerical instabilities in simulations. The presence of noise hinders the analysis and further processing of the data. For example, an isocontour visualization of a noisy scalar field contains a large number of connected components which clutter the visualization and mask the true features in the data. Smoothing of the data set is a commonly accepted approach to deal with noise. The common goal of all smoothing methods is to remove “small” features (noise) and keep “large” features. However, most methods handle the features in an implicit manner and do not provide guarantees regarding the removal of noise or the preservation of features.

Common examples of features include vortices in flow data sets (indicating, for instance, hurricane locations in atmospheric data), spots of high temperature in a physical simulation, or topographic highlights of a map. The ability to discriminate between noise and true features is important to most processing and visualization tasks, since features reveal significant insights about the structure and meaning of the data, while noise causes spurious features that may lead to poor visualizations, erroneous processing results and general misinterpretation of the data. It is desirable to remove

the noise in a way that leaves the true features as close as possible to their original values. For example, a temperature peak should retain its value after denoising to allow precise interpretation of the smoothed data.

In this work we explore a smoothing approach that allows explicit control over the features in the data. We focus on discrete 2D scalar fields that represent irregular and noisy samplings of phenomena that have a (piecewise) smooth ground truth. We use a general characterization of features commonly used in scientific visualization, namely the critical points (minima, maxima and saddles) of the scalar field, as well as the global topological structure they induce, i.e., the Morse-Smale complex (MS complex). Topology not only provides an abstract characterization of a scalar function, but also describes many domain-specific features such as the already mentioned vortices, temperature peaks, or topographic highlights, which can all be identified as critical points.

Our goal is to devise a data denoising and smoothing method that allows the user to explicitly control the fate of the features. The general idea is to capture *all* features (“small” and “large”), sort them according to a well-defined importance measure and allow the user to identify the threshold in this sorted list that discriminates between noise and true features. Based on this decision we guarantee the removal of the noise and the preservation of the true features. Furthermore, we guarantee that no additional features are in-



**Figure 1:** The Pressure data set. The top row shows the noisy input scalar field, and the bottom row our results. The original topology contains spurious critical points which have low persistence values and are removed by the simplification process. The persistence histogram on the right shows the number of topological simplification operations (cancellations) as function of their persistence; the obvious drop in the histogram around a persistence value of 0.01 hints the user at the appropriate noise threshold. The peak signal-to-noise ratio (PSNR) between noisy input and smoothed result is 30.8 dB.

roduced by the method. Note that standard smoothing methods generally cannot guarantee feature preservation and may even introduce new features in the output that were not part of the original data. For example, it is well known that Laplacian smoothing is prone to form singularities and new critical points in the data [GZ06].

The algorithmic core for the noise/feature discrimination of our technique consists of discrete Morse theory [For98] and topological simplification based on persistence [EHZ03] or similar importance measures (Section 4). We supply the user with a simple and intuitive interface to select the importance threshold in the form of a slider over the persistence histogram (see Figure 1c). Features whose importance is below the threshold are considered noise and are removed by means of topological simplification. The result is still a valid Morse-Smale complex, but with fewer critical points. The system provides immediate visual feedback of the remaining features and a preliminary  $C^0$  scalar field obeying the topological structure as a faithful preview of the end result (Section 5).

Our main contribution is the construction of a smooth function from the simplified Morse-Smale complex. After fairing the discrete embedding of the separatrices we run a constrained bi-Laplacian optimization of the entire domain, where the constraints are deduced from the monotonicity encoded in the simplified Morse-Smale complex (Section 6). The result is a  $C^1$ -continuous scalar field which closely fits the input data and obeys the desired topology. The positions and values of the remaining critical points are preserved exactly as they have been found in the input data. Thanks to the rigorous theoretical foundations, we are able to provide sound guarantees regarding the structure of the output, up to numerical precision. To summarize, the technical contributions of this work are:

- an algorithm that couples topological simplification and bi-Laplacian optimization to smooth scalar functions;
- real-time faithful preview of the end result;

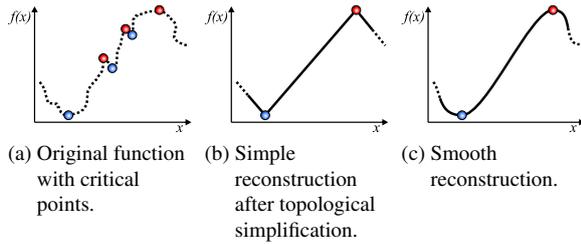
- constrained bi-Laplacian smoothing algorithm that produces  $C^1$  output, interpolates the desired features and guarantees monotonicity within each Morse cell.

We demonstrate the effectiveness of our approach on simulated and experimentally measured data sets (Section 7).

## 2. Related work

Topology-based smoothing of scalar fields is a common approach in scientific visualization. Carr et al. [CSvdP04] simplify the contour tree of the input data set, thereby suppressing small topological features. Gingold and Zorin [GZ06] modify existing filters, such as Laplacian smoothing or anisotropic diffusion, to disallow “illegal” changes in the isocontour topology (i.e., creation of new features). This approach enables one to remove some topological noise but does not provide explicit control over the resulting topological structure.

A setup similar to ours is given by Bremer et al. [BEHP04] in the sense that the Morse-Smale complex is simplified and a corresponding scalar field is obtained. There are two major differences between our method and [BEHP04]: first, they provide only  $C^0$ -continuity between Morse-Smale cells of the resulting scalar field, whereas we provide  $C^1$ -continuity there. This makes our method better suited for applications that require derivatives of the smoothed scalar field. We also describe a  $C^0$  method that we use as a real-time preview, which is comparable to the result of [BEHP04], but obtained much faster thanks to the second major difference between [BEHP04] and our approach: they smooth the function in the interior of a Morse-Smale cell *after each cancellation step* in order to obey the changed topological structure. The smoothing step itself is a complex, time-consuming procedure including 1D gradient smoothing and 2D iterative Laplacian smoothing. Since usually a large number of cancellations is necessary to remove noise, this iterative approach leads to very long computation times. In contrast, we construct a topologically valid function *after all cancellation steps* have



**Figure 2:** Overview of our algorithm using a 1D example.

been carried out. This is significantly faster and allows us to provide an interactive response for our  $C^0$  method. Note that our  $C^0$  method is comparable to the result of [BEHP04] since their 2D iterative Laplacian smoothing leads after several steps to the same harmonic function that we construct in Section 5 in a single step. In summary, both  $C^0$  methods provide comparable results, but ours is much faster to compute. More importantly, we devise a novel scheme to provide  $C^1$ -continuity.

It is worth noting that harmonic functions have been used for constructing Morse functions on 2-manifold domains [NGH04]. In that work, minima and maxima of the scalar field were prescribed; we use a similar machinery for our interactive  $C^0$  preview but we prescribe the topology fully, including saddle points and separatrices.

The topology of scalar fields is a subset of the topology of vector fields. Theisel and Weinkauff et al. construct a 2D or 3D vector field based on a given topological skeleton [The02, WTHS04]. These methods construct piecewise-linear vector fields with  $C^0$ -continuity and cannot avoid the appearance of additional critical points (i.e., in addition to the ones that have been prescribed by the user). Chen et al. [CML\*07] use Morse decomposition to edit the topology of vector fields on surfaces.

Feature-preserving smoothing is an active research area in image and geometry processing, where one strives to denoise the data without smoothing out features such as salient edges of a triangle mesh. Examples are anisotropic diffusion [HP04] and bilateral filtering [TM98]. Traditional low-pass filtering techniques such as Laplacian smoothing can also be constrained to interpolate prescribed feature points [Tau95]. However, none of these methods can generally guarantee that the features are exactly preserved and no new features are created in the smoothing process: for example, Laplacian or mean-curvature flows may have asymptotic singularities and may create novel features [GZ06, EPT\*07]. Depending on the application domain, other types of guarantees on the smoothing result have been explored, such as volume preservation [DMSB99, EPT\*07] or bounded distance from input data [HP07].

### 3. Overview

Our smoothing algorithm consists of three main steps: (i) *Persistence-based topological simplification of the*

*Morse-Smale complex.* The noisy input function contains a large amount of critical points (Figure 2a). Certain pairs of them can be assigned a well-accepted importance measure called persistence [ELZ02]. Pairs of critical points whose persistence is lower than a user-defined threshold are removed in a way that maintains a valid topological structure (Section 4). (ii) *Creation of a quick preview reconstruction of the scalar field.* Left with fewer critical points (Figure 2b), we create a preview reconstruction of the smoothed scalar field by computing a harmonic function within each cell of the Morse-Smale complex (Section 5). The harmonic function is guaranteed to obey the monotonicity requirement within each cell, but it is only  $C^0$  along the cell boundaries. (iii) *Final smooth scalar field construction by constrained bi-Laplacian optimization.* We compute a  $C^1$  function by constrained bi-Laplacian smoothing of the input scalar field, where the constraints are inequalities which force the scalar field to be monotonic within each Morse cell (Section 6 and Figure 2c). We convert the problem into an unconstrained nonlinear optimization problem by change of variables and solve it numerically using a standard truncated Newton minimization package.

### 4. Filtering Features

We identify the features of the input scalar field using topology. The topology of a scalar function consists of *critical points* as well as their relationship to each other. Critical points are found at locations where the gradient of the function vanishes. They are closely related to the behavior of the connected components of an isocontour when considering an increasing isovalue. There are three types of critical points: *minima*, where components emerge; *saddle points*, where they merge or split; and *maxima*, where components disappear. Four *separatrices* emanate from every saddle, which are forwards/backwards integrated gradient curves connecting the saddle with two maxima and two minima. It may happen that either both minima or both maxima coincide. Critical points and separatrices form a graph structure which gives rise to a domain decomposition called the *Morse-Smale complex*. It contains *2-dimensional Morse-Smale cells* that are bordered by one minimum, one maximum, and one or two saddles (Figure 3a). The combination of all Morse-Smale cells around a single maximum is called a *Morse cell* as shown in Figure 3b. Note how the border of a Morse cell consists of minima, saddle points, and the separatrices connecting them. The scalar function behaves monotonically inside a Morse cell with respect to the gradient: the function values are always monotonically increasing along a gradient curve, and any gradient curve started within the Morse cell ends at its sole maximum. Also, the Morse cells around all maxima provide a complete decomposition of the domain. Similar statements hold for Morse cells around minima.

#### 4.1. Computing the Morse-Smale Complex

We have chosen Forman's discrete Morse theory [For98, For02] as the basis for the computation of the Morse-Smale

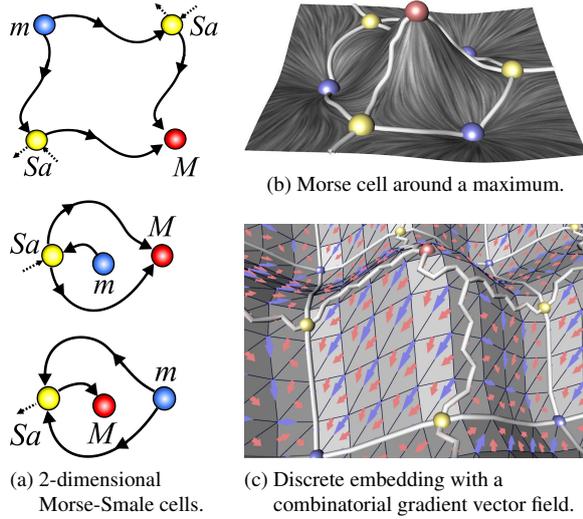


Figure 3: Topological concepts.

complex, since it allows purely combinatorial and thus consistent computations, whereas a numerical scheme based on root finding and gradient curve integration would be impaired by the noise of the underlying function. Discrete Morse theory utilizes the notion of a cell complex, which in our case of a triangular mesh consists of vertices (0-cells), edges (1-cells), and triangles (2-cells). In this discrete space, minima live on vertices, saddles on edges, and maxima on triangles. A separatrix connecting a saddle with a minimum is an alternating sequence of edges and vertices. Likewise, a separatrix from a saddle to a maximum is an alternating sequence of edges and triangles. Figure 3c illustrates this. This way of embedding the topology into the cell complex provides guarantees regarding the graph structure of the Morse-Smale complex. For example, two maxima (i.e., triangles) cannot exist next to each other without having a saddle (i.e., an edge) between them, and similarly for two minima. Furthermore, saddles can be connected to at most two maxima and two minima. The so-called monkey saddles (connected to more than two minima/maxima) are not possible. A consistent topological structure is highly important since we can then rely on valid topology for which the simplification operations are well defined, and a function that realizes the topology is guaranteed to exist.

The basis for discrete Morse theory is a *combinatorial vector field* – in our case a combinatorial *gradient* vector field. The value space of such a field is discrete. It is defined on the cell complex such that an  $n$ -cell points to a single incident  $(n + 1)$ -cell, i.e., vertex  $\rightarrow$  edge and edge  $\rightarrow$  triangle. Furthermore, it is required that a cell does not point to any of its neighbors if it is already being pointed to by one of them. Critical points are cells which neither point to a neighbor nor does a neighbor point to them (zero gradient). To compute the combinatorial gradient vector field, we follow the spanning-tree approach of Lewiner [Lew02] and

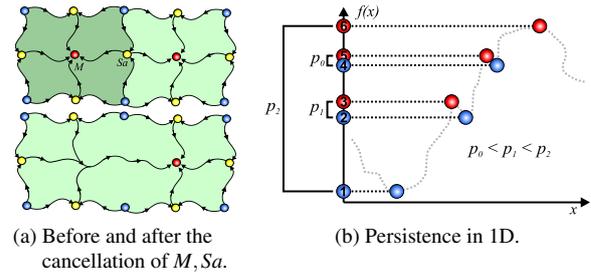


Figure 4: Persistence-based simplification by removing pairs of critical points.

Cazals et al. [CCL03]. It is a watershed-like algorithm that sweeps through the data in ascending order to establish the links between 0-cells and 1-cells, and in descending order for the links between 1-cells and 2-cells. Details of the implementation can be found in [WG09]. Figure 3c shows an example of a combinatorial vector field (note that it is computed from the negative gradient). To compute the Morse-Smale complex of a 3D scalar field, the method of Gyulassy et al. [GBHP08] provides a practical approach.

#### 4.2. Persistence-Based Topological Simplification

Topological simplification is the process of successively removing critical points from the Morse-Smale complex under the condition that the Morse-Smale complex is in a topologically consistent state after each simplification step. Consistency is only maintained if a saddle-minimum or a saddle-maximum pair is removed, which corresponds to merging two neighboring Morse cells (see Figure 4a). This is called a *cancellation*. In our implementation, we identify the order of cancellations using persistence as introduced by Edelsbrunner et al. [ELZ02, EHZ03], which is a measure of importance: critical points with lower persistence will be removed before ones with higher persistence. Loosely speaking, persistence measures the lifetime of connected components of an isocontour in a 2D scalar field when considering an increasing isovalue. More precisely, it measures the difference in function value between acts of creation (minima and splits at saddles) and acts of destruction (maxima and merges at saddles). Figure 4b illustrates the idea behind persistence using a simple 1D example: sweeping through the data in an ascending manner collects the minima (blue) and maxima (red) in the shown order. Every maximum (act of destruction) is then paired with a preceding minimum (act of creation) and the function value difference between these points is their persistence. Note how the global minimum and maximum have been assigned the highest persistence. A detailed algorithm for computing persistence for Morse-Smale complexes of 2D scalar fields can be found in [EHZ03]. This includes the pairs of critical points to be removed together. It has to be noted that paired critical points are not necessarily adjacent to each other in the initial Morse-Smale complex, but they will be adjacent right before their cancellation, as asserted by the Adjacency Lemma proven in [EHZ03].

A revealing insight into the noise level of a scalar field is often given by a simple histogram showing the number of cancellations over persistence, as depicted in Figure 1. Usually, a very high percentage of cancellations takes place at very low persistence levels. This makes it rather easy to identify a threshold discriminating between noise and true features of a data set. We use this histogram as a guidance for the user to determine the threshold.

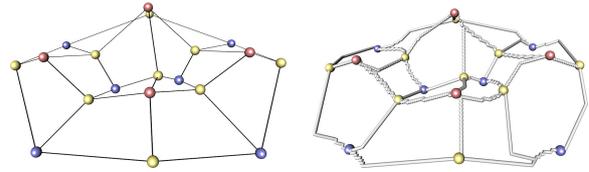
Several alternatives to persistence have been proposed in the literature that take geometric properties [CSvdP04] or application-specific aspects into account. Since our main contributions (Sections 5 and 6) are independent of the particular choice of a simplification metric, we restrict ourselves to persistence for the rest of the paper and leave a study of alternative metrics to future work.

### 4.3. From the Cell Complex to a Vertex-Based Representation

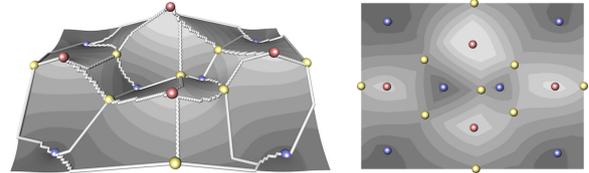
The goal of this work is to create a smooth version of the input scalar field containing only the simplified Morse-Smale complex. The necessary reconstruction of the function (Sections 5 and 6) is easier to handle if it is only carried out on the vertices of the mesh, but the discrete Morse-Smale complex lives partly on edges and triangles as described earlier. To accommodate these structures in a vertex-based representation, we insert new vertices and triangles around maxima, saddles, and the separatrices connecting them. To do so, we follow each separatrix from the saddle to the maximum, which is an alternating sequence of edges and triangles. We insert a new vertex in the center of every edge and triangle along the path and triangulate accordingly. An example of this can be seen in Figure 6. Note that this allows an unambiguous description of the Morse-Smale complex on the vertices of the mesh and is only needed for the simplified topology, i.e., only for a relatively small number of critical points and separatrices.

### 5. Reconstruction I: Interactive Preview

Given a simplified Morse-Smale complex, we create a fast “preview” version of the scalar field, which provides the user with a faithful view of the end result. It is “faithful” in the sense that it has exactly the topology prescribed by the simplified Morse-Smale complex and interpolates the values of the critical points. To do this, we first compute a valid value assignment for the separatrix curves (recall that after simplification, the values of the original function along the new separatrices are no longer monotonic). For each separatrix, we fix the critical points at its ends to their original values, and all the intermediate vertices are assigned a linearly interpolated value according to the arc-length parameter. Note that in the discrete setting, when tracing a separatrix from the starting saddle point, it is possible that we reach a junction where another separatrix with *already assigned* values joins; in this case we stop and take the junction vertex value as the end point for the linear interpolation. Figure 5a shows the resulting “linearized” separatrices.



(a) CP with their connectivity and linearized separatrices.



(b) Harmonic reconstruction as heightfield and in the 2D domain.

**Figure 5:** Constructing a faithful preview of the smoothed scalar field. The abstract connectivity of the simplified Morse-Smale complex (top left) is supplied with valid values by linearly interpolating between the endpoints of each separatrix (top right). A harmonic function is then computed with the separatrices as boundary conditions (bottom). Refer to Figure 9a for the input scalar field.

Once valid values along the separatrices have been set, we solve the Laplace equation  $\Delta f = 0$  on the domain, with the values along the separatrices and the critical points as Dirichlet boundary conditions. The solution  $f$  is a harmonic function that has no interior critical points in the domain: since the boundary conditions are monotonic, we obtain a valid scalar field that obeys the prescribed topology. The harmonic function  $f$  is  $C^\infty$  in the interior of the domain and  $C^0$  at the boundary curves, i.e., the separatrices (Figure 5b).

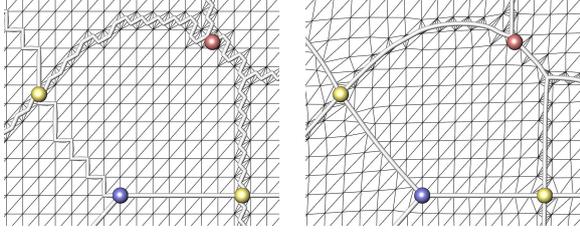
To solve the Laplace equation, we use the standard vertex-based linear FEM discretization of the Laplace operator on our parametric domain mesh, i.e., the cotangent weights matrix  $L_{\text{cot}}$  [PP93, DMSB99], or the mean-value weights matrix  $L_{\text{mean}}$  [Flo03] if the cotangent weights are negative (positive weights guarantee that the solution of the discrete Laplace equation has no interior critical points). We therefore solve the following linear system of equations for the vertex values  $\mathbf{f} = (f_1, \dots, f_n)$ :

$$L\mathbf{f} = 0, \quad \text{s.t. } f_i = \hat{f}_i \quad \forall i \in S, \quad (1)$$

where  $S$  is the set of mesh vertices that belong to the separatrices (including the critical points) and  $\hat{f}_i$  are their fixed values. The Laplacian matrix  $L$  is sparse (7 nonzeros per row on average), symmetric and positive definite; we therefore use sparse Cholesky factorization to solve the system [Tol03], which allows us to obtain the preview scalar field at interactive frame rates (some timing statistics are given in Section 7).

### 6. Reconstruction II: $C^1$ -smooth Function

The scalar field obtained in the previous section is only piecewise-smooth. In the following, we describe the recon-



**Figure 6:** The parametric embedding of the simplified separatrices may be jaggy (left) and is smoothed by Laplacian mesh optimization (right).

struction of a smooth scalar field that obeys the prescribed Morse–Smale complex and interpolates the values of the critical points. The smooth reconstruction involves three steps: first, we optimize the parametric domain so as to smooth the embedding of the separatrices (this facilitates the reconstruction of the smooth scalar field later). We then optimize a smoothing objective functional (specifically, the bi-Laplacian) in a constrained setting which guarantees monotonicity of the scalar field within each Morse cell, such that the result is a  $C^1$ -smooth function that has the prescribed topology.

### 6.1. Laplacian-Based Fairing of the Discrete Embedding of Separatrices

After the topological simplification, depending on the amount of noise in the input field, we might be left with long separatrices whose shape meanders through the parameter domain (see Figure 6, left). Such separatrices inhibit the subsequent smooth reconstruction of the scalar field and cause visual clutter, so we would like to smooth them in the parametric domain first. Since all our discrete topology information resides in the elements of the mesh, we would like to keep the mesh intact and only slightly shift the vertices in the parametric domain. We therefore employ a particular form of 2D Laplacian mesh optimization [NISA06], which does not alter the mesh connectivity.

Each separatrix  $s_i$  is a chain of vertices in the parametric domain, starting and ending with a critical point:  $(\mathbf{v}_1^i, \mathbf{v}_2^i, \dots, \mathbf{v}_{n_i}^i)$ ,  $\mathbf{v}_j^i \in \mathbb{R}^2$ . We smooth the spatial embedding of each  $s_i$  and regularize the spacing of the vertices along it, so for each inner vertex of each separatrix we have the term:

$$\left\| \mathbf{v}_j^i - 0.5 \left( \mathbf{v}_{j-1}^i + \mathbf{v}_{j+1}^i \right) \right\|^2 \quad j = 2, \dots, n_i - 1. \quad (2)$$

We also fair the entire mesh to prevent parametric non-smoothness near the separatrices, by adding an energy term for each vertex  $\mathbf{v}_k$  not on a separatrix. Fairness is defined via the cotangent Laplacian operator [NISA06]:

$$\left\| \sum_{j:(k,j) \in \mathcal{E}} w_{kj} (\mathbf{v}_k - \mathbf{v}_j) \right\|^2, \quad (3)$$

where  $w_{kj}$  are the cotangent weights and  $\mathcal{E}$  is the set of mesh

edges. To prevent the mesh vertices from moving too much, we put a soft positional constraint on each  $\mathbf{v}_k$ :

$$\omega \|\mathbf{v}_k - \mathbf{u}_k\|^2, \quad (4)$$

where  $\mathbf{u}_i$  is the original location of  $\mathbf{v}_k$  and  $\omega$  is a small positive weight ( $\omega = 0.01$  in our experiments). We sum up all the energy terms in Eqs. (2), (3) and (4) and minimize the total energy under the hard constraints of fixing the locations of all the critical points and the domain boundary vertices. Solving the resulting sparse linear optimization is performed efficiently by TAUCS [Tol03]. Figure 6 demonstrates the resulting fairing effect.

### 6.2. Monotonic Reconstruction of a Smooth Function

Our goal is to reconstruct a function  $f$  that is (i) smooth, (ii) passes close to the original input values, (iii) interpolates the prescribed critical point values and (iv) has exactly the prescribed topology. In the following, we describe the general idea behind reconstructing a function according to the four requirements, and then show how we do this specifically for the 1D separatrices and the 2D scalar field. We will first reconstruct the values along the separatrices, which will then serve as boundary constraints in the 2D optimization. The two tasks are solved in essentially the same manner, up to the particular encoding of the topological requirements.

Smoothness and closeness to the original data are addressed by defining an objective functional that our function  $f$  should minimize in a constrained setting:

$$E(f) = \int_{\Omega} |\Delta f|^2 + \omega_d |f - \hat{f}|^2 dA, \quad (5)$$

where  $\Omega$  is our parametric domain,  $\hat{f}$  is the original function and  $\omega_d > 0$  is the weight of the data term. Larger values for  $\omega_d$  lead to results which are closer to the original data. In our experiments, we used values between  $10^4$  and  $10^6$  for  $\omega_d$ . Minimizing  $E(f)$  corresponds to bi-Laplacian smoothing and produces functions that are  $C^1$  at the boundary constraints and  $C^\infty$  everywhere else (up to certain discretization and convergence conditions, see [GGRZ06]). In our discrete setting, we use the linear FEM discretization of the Laplace operator, either on a curve or a 2D mesh domain (refer to Section 5), and formulate the energy in terms of the function values  $\mathbf{f} = (f_1, \dots, f_n)$  at the vertices:

$$E(\mathbf{f}) = \|\mathbf{L}\mathbf{f}\|^2 + \omega_d \|\mathbf{f} - \hat{\mathbf{f}}\|^2 \rightarrow \min. \quad (6)$$

The interpolation of critical point values is translated simply into Dirichlet boundary conditions  $f_i = \hat{f}_i$  where  $i$  are the vertex indices of the critical points.

In order to satisfy the topology requirements, we introduce monotonicity constraints: the function should be monotonic “in-between” adjacent critical points. For 2D domains, “adjacency” of critical points is defined via the Morse complex; in each Morse cell around a maximum, the function should decrease from the maximum towards the boundary of the cell (cf. Figure 3b). We encode monotonicity as a directed, acyclic graph (the *monotonicity graph*) on the ver-

tices of the mesh with sources at maxima and sinks at minima. A path in the monotonicity graph then represents an integral line. We use the monotonicity graph also to represent the network of separatrices – each chain path in the graph then corresponds to a separatrix, i.e., the function is decreasing from a maximum to a saddle and increasing from a minimum to a saddle. Note that the monotonicity graph for the separatrices is well-defined.

Assume we have a monotonicity graph defined for our problem; each interior vertex  $i$  then has an incoming edge in the graph; the origin vertex of that edge is denoted by  $parent(i)$ . We assume that any vertex that has more than one parent (or is a critical point) has a fixed value. For each free vertex  $i$ , the topological constraints can then be formulated in terms of inequalities:

$$f_i < f_{parent(i)}. \quad (7)$$

We have now theoretically defined the complete problem setup for the reconstruction of a function that satisfies all the four requirements: one needs to minimize the energy  $E(\mathbf{f})$  in Eq. (6) under the Dirichlet boundary conditions and the inequality constraints (7). While the quadratic minimization of the bi-harmonic functional under Dirichlet boundary conditions amounts to merely solving a sparse linear system, the additional inequality constraints turn our problem into quadratic programming, which may be quite complex to solve in practice. Instead of using general tools for quadratic programming, we utilize the monotonicity graph to convert the inequalities into (nonlinear) equalities. We then end-up with a nonlinear system to solve, but since all our formulations are analytic, it is straightforward to solve by standard numerical methods for nonlinear optimization.

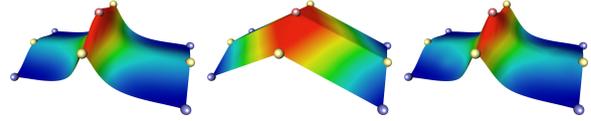
Since the function values along a path in the graph are decreasing, the value  $f_i$  should be between the value of its parent and the largest value at the end of all paths passing through  $f_i$ , denoted by  $f_{m_i}$ . We can thus write  $f_i$  as an interpolation of the two:

$$f_i = f_{m_i} + t_i (f_{parent(i)} - f_{m_i}), \quad 0 \leq t_i \leq 1. \quad (8)$$

We perform a simple change of variables in order to get rid of the inequalities  $0 \leq t_i \leq 1$ : for any free vertex  $i$  we attach a new variable  $\theta_i$ , such that

$$t_i = t_i(\theta_i) = 0.5 + 0.5 \cos(\theta_i). \quad (9)$$

The inequalities in (8) are now satisfied by construction, and the expressions for the function values  $f_i$  become dependent on the new variables  $\theta_i$ ; these expressions are then plugged into the smooth reconstruction functional (6). The energy is nonlinear in  $\theta_i$ ; typical numerical optimization methods such as Gauss-Newton require expressions for the gradient and the Hessian of the energy, i.e., first and second partial derivatives w.r.t.  $\theta_i$ 's. Fortunately, these are easy to derive analytically, and we provide the formulas in the supplemental material.



**Figure 7:** Original function (left) has been reconstructed based on the topology shown using the  $C^0$  preview scheme (middle, Section 5) and the  $C^1$  method (right, Section 6).

### 6.3. Reconstructing the Scalar Field

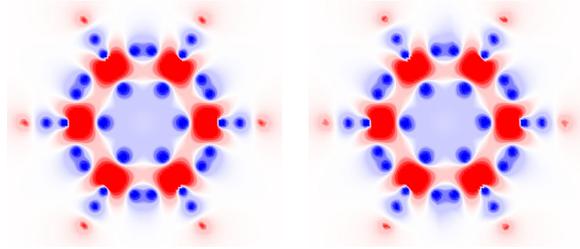
We reconstruct the smoothed scalar field using a monotonicity graph describing the Morse decomposition with sources at maxima and sinks at minima. Each Morse cell contains a single maximum and its boundary is composed of saddles and minima connected by separatrices (cf. Figure 3b). The corresponding monotonicity graph is built by performing a watershed-like algorithm on a function that has the prescribed Morse complex – we use the same harmonic function that we computed for the interactive preview (Section 5).

As mentioned in the previous section, all free vertices in the optimization should have a single parent, and we fix all the vertices that have more than one incoming edge in the monotonicity graph. Therefore we need to fix the values along the boundaries of the Morse cells, namely the separatrices that run from saddles to minima (as each such separatrix has incoming paths that descend from the maximum, as well as the separatrix path itself). The values on these boundary separatrices are unknown (recall that we only have fixed values for the critical points), therefore we first run the optimization process on the collection of these separatrices. For this curve-network optimization we discretize the 1D Laplace operator (i.e., the second derivative) in (6) using uniform weights. Once the values on the boundary separatrices have been computed, we run the optimization of the entire domain while fixing the critical points and the boundary separatrices, this time using the FEM discretization of the 2D Laplace operator (as in Section 5). Figure 7 compares the result of the  $C^1$  optimization with the  $C^0$  preview reconstruction for a single Morse cell. It clearly shows that our nonlinear optimization is able to recover the characteristics of the original function while obeying the prescribed topology. Also note the smoothness across the (omitted) separatrices. This is in contrast to the method of [BEHP04], where iterative Laplacian smoothing inside a Morse-Smale cell eventually leads to a result similar to our  $C^0$  method.

Note that the same discrete Morse decomposition of a 2D mesh domain has many different possible monotonicity graphs that are topologically valid; each monotonicity graph implies a function with different gradient lines, so that the structure of the graph has significant impact on the shape of the resulting reconstruction. In our reconstruction process, we prefer monotonicity graphs with as few edges as possible to prevent overconstraining the optimization. We have chosen to use the harmonic function as the basis for our monotonicity graph, observing that it has low variation, but in future work, it would be interesting to explore the possible monotonicity encodings further.

Data set	#Vertices	In/Out #CP	Simpl. time	Rec. I time	Max/Avg path	Rec. II time
Fig. 1, Pressure	11266	499/81	0.01	0.11	175/34.5	719
Fig. 8, Benzene	13377	85/85	–	0.04	173/38.9	981
Fig. 9, random noise	2711	481/19	0.001	0.01	88/25.7	580
Fig. 9, logo overlay	4825	43/19	0.001	0.05	95/28.7	603
Fig. 10, $P = 5.0$	20126	7733/167	0.04	0.1	344/61.3	2603
Fig. 10, $P = 18.6$	14758	7733/13	0.05	0.1	432/91.5	3027

**Table 1:** *Experimental statistics.* In/Out #CP: number of critical points; Simpl. time: topological simplification time (Section 4.2); Rec. I time: simple harmonic reconstruction time (Section 5); Rec. II time: constrained bi-Laplacian optimization time (Section 6); Max/Avg path: maximal and average path length within the monotonicity graph. All timings are given in seconds.



**Figure 8:** *Electrostatic field of the Benzene molecule.* The original data set (left) has been reconstructed (right) without topological simplification. The PSNR is 50.8 dB.

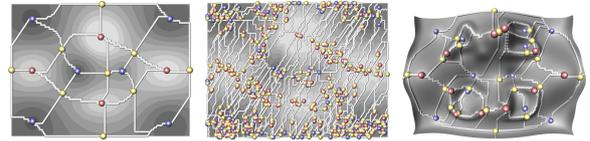
## 7. Examples

We have implemented our technique in C++ and tested it on current hardware (Intel Core 2 Duo 3 GHz with 3 GB memory). Table 1 summarizes some statistics about the performance of our system. By design, the algorithm produces interactive preview results: computation of the persistence values and subsequent topological simplification required no more than 0.01 seconds. The reconstruction of the preview scalar field (the harmonic function) is fast as well, in the order of 0.05-0.1 seconds, so it can be carried out interactively while the user is dragging the persistence slider. The final smooth reconstruction by constrained bi-Laplacian optimization can be done offline once the user is satisfied with the persistence threshold and the resulting topology. Our experiments have shown that the optimization time mainly depends on the average and maximal path length in the monotonicity graph, and less so on the total number of vertices (Table 1).

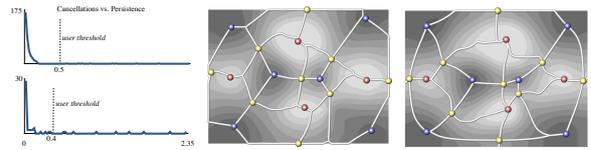
In our first experiment we test how well the method reconstructs the scalar field from the original topology, i.e., the persistence threshold is set to zero and no topological simplification takes place. Figure 8 shows the original and reconstructed electrostatic field of the benzene molecule: the features and the overall structure have been very well preserved by our optimization stage. This is also backed by the very high peak signal-to-noise ratio (PSNR) of 50.8 dB (typical PSNR values for lossy image compression are between 30 and 50 dB). Subtle differences are noticeable around the small outer H-atoms (red) where the resolution of the mesh is rather low compared to the feature size such that the discrete nature of our approach becomes visible.



(a) The original synthetic function (left) is corrupted by additive noise (middle) and by an overlay displacement field (right, displayed as a heightfield).



(b) The topological structures of the ground-truth and the test cases.

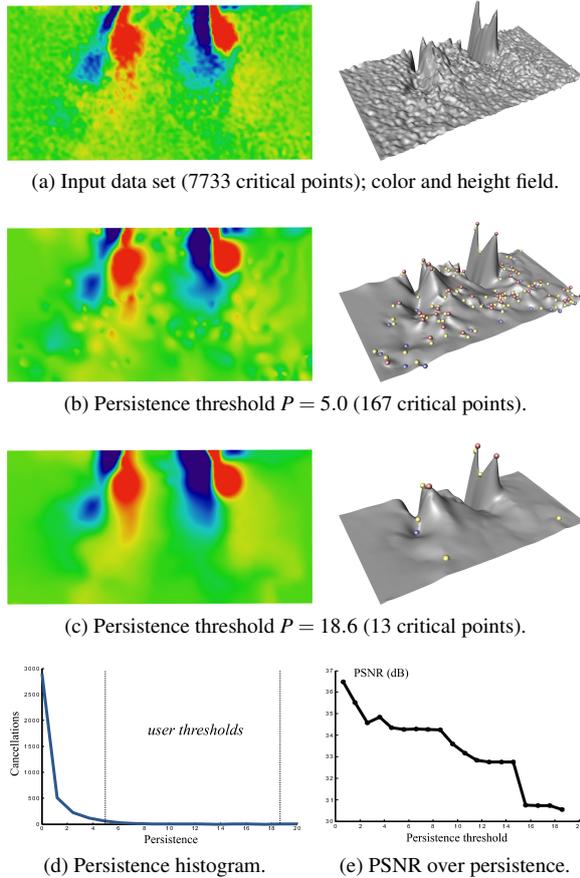


(c) Persistence histograms of the two test cases with the user-specified thresholds (left); the reconstructed results for both the noisy (middle) and the overlay (right) tests.

**Figure 9:** *Testing our algorithm on a synthetic example where the ground truth is a known smooth function.* We tested two types of corruption: random noise and overlay of another non-random function. In both cases the algorithm reconstructs a high-quality smooth scalar field with the correct topological structure.

Figure 9 shows a synthetic 2D function, which we corrupted in two different ways: by random additive noise and by adding another (weaker) overlay signal. The underlying assumption is that the persistence of the noise is lower than that of the underlying true data. One can observe that our method successfully reconstructs the features of the original scalar field in both cases; the appropriate persistence threshold is quite evident from the persistence histogram in Figure 9c and is therefore not difficult to tune.

Figure 1 shows an example from a fluid dynamics simulation of a turbulent flow behind a bluff body. This data set has been computed by Erik Wassen (Technical University Berlin) using a Large-Eddy simulation scheme at a Reynolds



**Figure 10:** Experimentally measured data set showing the instantaneous vorticity at the outlet of a combustor (flow direction from top to bottom). The size and overall shape of the major features at the burner outlet (top) are preserved while the noise has been successfully removed.

number of 500000 based on model length and incoming velocity [WT07]. The data set shows the pressure of the flow, which indicates vortex activity: regions of low pressure (blue colors) denote a vortex. We focus on the region directly behind the bluff body, where one can see a number of vortices. However, a significant amount of noise can be observed in the input pressure field (Figure 1a). This is also evident from the topological analysis (Figure 1b). To a large degree, this noise comes from the simulation scheme itself and is therefore not desired. The spurious features have very low persistence values (refer to Figure 1c) and are removed in the topological simplification process. Note how the prominent critical points are better emphasized in the optimized result.

Figure 10 shows an experimentally measured data set from fluid dynamics: a flow field obtained using Particle Image Velocimetry (PIV) in a slice at the outlet of a swirl-inducing combustor. The data set is courtesy of Arnaud Lacarelle (Technical University Berlin) and the experimental setup is described in [LFG\*09]. Shown is the instantane-

ous vorticity of the flow which indicates vortex activity. Vortices steer the mixing of fresh and hot gas in a burning chamber and are therefore essential for an efficient burning process. The setup of this PIV measurement is inherently 2D (the flow is only measured in a slice) and prone to noise as indicated by the 7733 critical points in the original data set (Figure 10a). The most dominant features in this data set are the major vortices at the top where the intake of the combustor (burner outlet) is located. Figures 10b–c show two different simplification levels obtained with our method containing only 167 and 13 critical points, respectively. Note how the noise has been filtered while the size and the overall shape of the major features has been preserved: the major peaks and pits in the height fields of Figures 10b–c are exactly at the same location and have the same height as in the input data set. In Figure 10e we plotted the PSNR values of several simplification levels in the persistence interval  $[0.6, 18.6]$ . Note that even the PSNR value for the highest chosen persistence threshold does not drop below 30 dB.

## 8. Conclusions

We have presented a smoothing method for 2D scalar fields that focuses on the topological features and enables to explicitly control the topology of the output. The interactive control mechanism provided by our method, coupled with real-time faithful preview of the end result, allows for an effective and intuitive processing workflow. The central algorithmic components of our method are discrete topological simplification and constrained bi-Laplacian optimization; the combination enables to provide guarantees regarding the topological structure of the output.

Persistence has a limitation with respect to strong outliers: their persistence is likely to be high and they will not be filtered out. Hence, in future work we will incorporate different persistence measures that help identify outliers into our framework. Another future research direction is to explore how different monotonicity descriptions within a Morse cell affect the results of the constrained optimization; in particular, it would be interesting to come up with an optimization of the monotonicity graph itself.

Finally, we believe that direct control over topological structures may have wide applications beyond data visualization and analysis. In the future we would like to expand the method to other domains, such as scalar fields on surfaces or 3D data, and explore the possibilities of modeling and editing using topologically-guaranteed methods.

## Acknowledgments

We are grateful to Denis Zorin and Ashish Myles for their valuable comments. Tino Weinkauff is supported by a Feodor Lynen research fellowship of the Alexander von Humboldt foundation. This work was supported in part by an ADVANCE Research Challenge Grant funded by the NSF ADVANCE-PAID award HRD-0820202 and by an NSF award IIS-0905502. All visualizations in this paper have been created using AMIRA – a system for advanced visual data analysis (see <http://amira.zib.de/>) [SWH05].

## References

- [BEHP04] BREMER P.-T., EDELSBRUNNER H., HAMANN B., PASCUCCI V.: A topological hierarchy for functions on triangulated surfaces. *IEEE TVCG 10*, 4 (2004), 385–396.
- [CCL03] CAZALS F., CHAZAL F., LEWINER T.: Molecular shape analysis based upon the Morse-Smale complex and the Connolly function. In *Proc. SoCG* (2003), pp. 351–360.
- [CML\*07] CHEN G., MISCHAIKOW K., LARAMEE R. S., PILARCZYK P., ZHANG E.: Vector field editing and periodic orbit extraction using Morse decomposition. *IEEE TVCG 13*, 4 (2007), 769–785.
- [CSvdP04] CARR H., SNOEYINK J., VAN DE PANNE M.: Simplifying flexible isosurfaces using local geometric measures. In *Proc. Visualization* (2004), pp. 497–504.
- [DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. ACM SIGGRAPH* (1999), pp. 317–324.
- [EHZ03] EDELSBRUNNER H., HARER J., ZOMORODIAN A.: Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete and Comp. Geometry 30*, 1 (2003), 87–107.
- [ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete and Computational Geometry 28*, 4 (2002), 511–533.
- [EPT\*07] ECKSTEIN I., PONS J.-P., TONG Y., KUO C.-C. J., DESBRUN M.: Generalized surface flows for mesh processing. In *Proc. SGP* (2007), pp. 183–192.
- [Flo03] FLOATER M. S.: Mean value coordinates. *Computer Aided Geometric Design 20*, 1 (2003), 19–27.
- [For98] FORMAN R.: Morse theory for cell-complexes. *Advances in Mathematics 134*, 1 (1998), 90–145.
- [For02] FORMAN R.: A user's guide to discrete morse theory. Séminaire Lotharingien de Combinatoire, B48c, 2002.
- [GBHP08] GYULASSY A., BREMER P.-T., HAMANN B., PASCUCCI V.: A practical approach to Morse-Smale complex computation: Scalability and generality. *IEEE Transactions on Visualization and Computer Graphics 14* (2008), 1619–1626.
- [GGRZ06] GRINSPUN E., GINGOLD Y., REISMAN J., ZORIN D.: Computing discrete shape operators on general meshes. *Computer Graphics Forum 25*, 3 (2006), 547–556.
- [GZ06] GINGOLD Y. I., ZORIN D.: Controlled-topology filtering. In *Proc. ACM SPM* (2006), pp. 53–61.
- [HP04] HILDEBRANDT K., POLTHIER K.: Anisotropic filtering of non-linear surface features. *Computer Graphics Forum 23*, 3 (2004), 391–400.
- [HP07] HILDEBRANDT K., POLTHIER K.: Constraint-based fairing of surface meshes. In *Proc. Eurographics Symposium on Geometry Processing* (2007), pp. 203–212.
- [Lew02] LEWINER T.: *Constructing discrete Morse functions*. Master's thesis, Department of Mathematics, PUC-Rio, 2002.
- [LFG\*09] LACARELLE A., FAUSTMANN T., GREENBLATT D., PASCHEREIT C., LEHMANN O., LUCHTENBURG D., NOACK B.: Spatiotemporal Characterization of a Conical Swirler Flow Field Under Strong Forcing. *Journal of Engineering for Gas Turbines and Power 131* (2009), 031504.
- [NGH04] NI X., GARLAND M., HART J.: Fair Morse functions for extracting the topological structure of a surface mesh. In *Proc. SIGGRAPH* (2004), pp. 613–622.
- [NISA06] NEALEN A., IGARASHI T., SORKINE O., ALEXA M.: Laplacian mesh optimization. In *Proc. ACM GRAPHITE* (2006), pp. 381–389.
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Exp. Math.* 2, 1 (1993), 15–36.
- [SWH05] STALLING D., WESTERHOFF M., HEGE H.-C.: Amira: A highly interactive system for visual data analysis. *The Visualization Handbook* (2005), 749–767.
- [Tau95] TAUBIN G.: A signal processing approach to fair surface design. In *Proc. ACM SIGGRAPH* (1995), pp. 351–358.
- [The02] THEISEL H.: Designing 2D vector fields of arbitrary topology. *Computer Graphics Forum 21*, 3 (2002), 595–604.
- [TM98] TOMASI C., MANDUCHI R.: Bilateral filtering for gray and color images. In *Proc. ICCV* (1998), pp. 839–846.
- [Tol03] TOLEDO S.: TAUCS: A Library of Sparse Linear Solvers, version 2.2. Tel-Aviv University, Available online at <http://www.tau.ac.il/~stoledo/taucs/>, Sept. 2003.
- [WG09] WEINKAUF T., GÜNTHER D.: Separatrix persistence: Extraction of salient edges on surfaces using topological methods. *Computer Graphics Forum (Proc. SGP '09) 28*, 5 (July 2009), 1519–1528.
- [WT07] WASSEN E., THIELE F.: LES of wake control for a generic fastback vehicle. In *37th AIAA Fluid Dynamics Conference and Exhibit* (Miami, U.S.A., 2007). AIAA-2007-4504.
- [WTHS04] WEINKAUF T., THEISEL H., HEGE H.-C., SEIDEL H.-P.: Topological construction and visualization of higher order 3D vector fields. *Computer Graphics Forum (Proc. Eurographics 2004) 23*, 3 (2004), 469–478.