SCI INSTITUTE
TECHNICAL REPORT

Edge Flows: Stratified Morse Theory for Simple, Correct Isosurface Extraction

Carlos Scheidegger, Tiago Etiene, L. Gustavo Nonato, and Cláudio T. Silva, Senior Member, IEEE

UUSCI-2010-002

Scientific Computing and Imaging Institute
University of Utah
Salt Lake City, UT 84112 USA

August 5, 2010

Abstract:

We present a method to characterize the topology of the level sets of trilinearly interpolated scalar fields. Our characterization is based on Morse theory, and in particular a variant called Stratified Morse theory capable of treating the piecewise-smooth aspect of trilinear interpolation. Algorithms such as Marching Cubes generate approximations to these level sets to a varying degree of fidelity. It is now folklore that the standard Marching Cubes algorithm has inconsistencies, and while corrected versions are well-established, it is still the case that versions which strive for homeomorphism have complicated case tables. Our characterization explains exactly what is the source of topological problems in Marching Cubes, suggests simpler algorithms that generate isosurfaces homeomorphic to the level sets, and allows short, complete proofs of correctness that can also be used to prove homeomorphism for a subset of the outputs of Marching Cubes itself. We provide an open-source implementation of this algorithm and report on the results of its verification.

THE UNIVERSITY OF UTAH

Edge Flows: Stratified Morse Theory for Simple, Correct Isosurface Extraction

Carlos Scheidegger, Tiago Etiene, L. Gustavo Nonato, and C. Silva, *Senior Member, IEEE*

Abstract— We present a method to characterize the topology of the level sets of trilinearly interpolated scalar fields. Our characterization is based on Morse theory, and in particular a variant called Stratified Morse theory capable of treating the piecewise-smooth aspect of trilinear interpolation. Algorithms such as Marching Cubes generate approximations to these level sets to a varying degree of fidelity. It is now folklore that the standard Marching Cubes algorithm has inconsistencies, and while corrected versions are well-established, it is still the case that versions which strive for homeomorphism have complicated case tables. Our characterization explains exactly what is the source of topological problems in Marching Cubes, suggests simpler algorithms that generate isosurfaces homeomorphic to the level sets, and allows short, complete proofs of correctness that can also be used to prove homeomorphism for a subset of the outputs of Marching Cubes itself. We provide an open-source implementation of this algorithm and report on the results of its verification.

Index Terms—Stratified Morse Theory, Marching Cubes, Isosurface, Topology.

1 Introduction

Isosurface extraction is a popular technique in scientific visualization and computational science. Although Marching Cubes is now about 25 years old [18], the basic idea behind it has not only proven very long-lived, but also a rich source of mathematical and computational problems [20]. In this paper, we are concerned with topological properties of the level sets of trilinearly interpolated scalar fields, and how faithfully they are reproduced in triangle meshes generated by algorithms such as Marching Cubes.

If we want to generate triangle meshes for scientific visualizations, and, more importantly, for computational simulations that depend on the properties of the generated mesh, it is important that we understand the mathematical phenomena that arise from trilinear interpolation. In this paper, we will focus on topological properties of these surfaces, in enough detail to be able to show that our algorithms generate triangular meshes respecting the topology of the level sets.

We use a variant of Morse theory to isolate the topological changes of the level sets of linear interpolation. The main complication in our scenario is that the scalar fields generated by trilinear interpolation are only piecewise smooth (as are the level sets). Consider, for example, the surfaces shown in Figure 2. Although standard Morse theory is not well-equipped to make statements about the surface (since the gradient does not exist on many points), it is intuitively clear that some analog of Morse Theory should apply. To analyze the topological behavior of these more general level sets, then, we turn to Stratified Morse Theory, or SMT. We give a criteria that tells us exactly in which situations Marching Cubes generates surfaces with the correct topology (we will make these notions precise in Section 3). We then propose a simple divide-and-conquer algorithm that uses Marching Cubes as a base case, together with an appropriately defined gluing strategy. This algorithm generates triangle meshes that are provably homeomorphic to the level sets, and is simpler than the complicated tables of currently published Marching Cubes-like algorithms for extracting such isosurfaces.

- *Scheidegger is with AT&T Labs – Research. email: cscheid@research.att.com*
- *Etiene and Silva are with the School of Computing and Scientific Computing and Imaging Institute, University of Utah. email: {tetiene,csilva}@cs.utah.edu.*
- *Nonato is with Universidade de São Paulo, Brazil. email: gnonato@icmc.usp.br.*

Manuscript received 31 March 2010; accepted 1 August 2010; posted online 24 October 2010; mailed on 16 October 2010.
For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

We point out that we do not claim here to have the first analysis of the topology of the trilinear interpolant. We do claim that our analysis is qualitatively different from the ones previously presented in the literature. Specifically, our use of Morse theory for determining the presence or absence of face and body saddles eliminates much of the manual labor of case analysis. In addition, the only case table we require is computed programmatically via a simple group theory argument of Banks et al. [3]. Furthermore, the same analysis gives a novel, simple way to compute the Euler characteristic of the level set being extracted.

Specifically, our contributions are:

- a topological characterization of the ambiguities that arise when using Marching Cubes to extract isosurfaces of trilinearly interpolated scalar fields. Based on stratified Morse theory (which, as far as we know, is the first use of SMT in the scientific visualization community), it provides a complementary approach to previous work on the topological correctness of isosurfacing algorithms;
- A straightforward proof of correctness of the resulting algorithm. While the resulting algorithm is in essence the same as the one implied by previously published subdivision schemes for the trilinear interpolant [7, 5, 6], our arguments are less tied to the actual interpolant formulation, and seem more amenable to generalization. We also believe that the resulting algorithm is simpler, and so we provide
- a reference open-source implementation. Notably, this implementation handles in a simple way the cracks generated by the nonuniform subdivision, which has been pointed out as a real hurdle for the algorithm's practicality [6]. As evidence of its effectivenes, we provide a report of the implementation's verification of both geometrical and topological properties, in the spirit of Etiene et al.'s recent work [11].

2 Related Work

Soon after the publication of the original Marching Cubes paper [18], researchers noted that some sign configurations generate topologically inconsistent isosurfaces. In a landmark paper, Nielson and Hamann pointed out these inconsistencies can happen when a face contains a *saddle point* of the bilinear interpolation [22], and suggested an extended rule that ensures no cracks appear in the final surface. This line of work eventually led to algorithms that generate isosurfaces which are homeomorphic to the trilinear level set [8, 16, 21]. However, all

such algorithms involve to the best of our knowledge large case analyses, complicating the task of verifying the implementations. In addition, there are issues with some of the stated topological correctness proofs, as we will discuss in Section 6.

Topological correctness in meshing is widely studied. The seminal work is Edelsbrunner and Shah's work showing a condition where the triangulation is homeomorphic to the original manifold [10]. This led to a flurry of work, of which Amenta and Bern's *crust* and ε-sampling idea is the classic reference [2]. Stander and Hart presented one of the earliest papers connecting isosurface extraction and Morse theory, with an algorithm which performs surgery on the resulting triangle set as it sweeps through critical points [24]. Plantinga and Vegter present a criterion for isotopic isosurfacing that can be used to show that a subset of the edge flow cases that we will present in Section 3.5 is a correct reconstruction [23]. Their proof is significantly more complicated than ours, but, more importantly, does not seem to apply straightforwardly to the piecewise-smooth setting. Their analysis of Marching Cubes cases is similar to the ones of Van Gelder and Wilhelms, and Stelldinger et al. [27, 25]. Varadhan et al. use a couple of sampling criteria to ensure that geometrical computations produce results with the right topology [28]. In a nutshell, all of these techniques try to show that the resulting isosurface component inside each cell is sufficiently simple (typically a disk). Our edge flow diagrams, which we will present next, make an explicit connection of these notions with Morse theory. By requiring that the scalar field to be as simple as possible inside each cell, we also get topological disks. We believe this is a cleaner argument that leads to simpler algorithms and verification tools.

The topology of the trilinear interpolant has been previously analyzed before. In particular, Carr and Snoeyink [7] introduce the idea of creating a simplicial decomposition of the cubic cell and placing the vertices so that critical points are isolated to vertices of the mesh. They use this decomposition to translate the isosurface sweeping into a state machine that can be efficiently implemented. Carr and Max then use this state machine to perform a full analysis of the isosurface cases for the trilinear interpolant [5, 6]. We believe our argument to be an attractive alternative because it eliminates the need for much of the error-prone case analysis. While Carr's analysis uses Nielson's manually-crafted case table, we use a simpler table mechanically computed from orbits of the cube's symmetry group. Our proof of correctness of the isosurfacing algorithm then only requires a reference to the simple cases and to an appropriate gluing procedure. We provide experimental evidence for the effectiveness of our analysis via an open-source reference implementation of the algorithm and its verification.

Morse theory has provided a very fruitful approach to understand and analyse the topology of level sets of scalar fields [19]. In particular, Morse theory is attractive because it has a clean combinatorial formulation that is intimately related to piecewise-linear reconstructions and very well suited for robust implementations [13]. We point out that in this paper we do not use the PL variant of Morse theory, because we are interested in understanding the topology of the trilinearly interpolated scalar fields. Since these are neither piecewise linear or entirely smooth, none of the well-established variants of Morse theory in scientific visualization are directly applicable. There has been some previous work in solid modeling which suggests data structures based on some form of stratification in order to deal with manifold boundaries [17]. In this paper, we are interested in the application of stratification to reconcile Morse theory with piecewise-smooth functions, in particular in the context of isosurface extraction.

3 The topology of trilinear interpolation

In this section we provide a characterization of the topological behavior of level sets of trilinearly interpolated scalar fields. One way to think about Morse Theory is that it gives a well-defined way to *localize* the changes in the topology of level sets with different scalar values. Morse theory tells us that for most scalar fields, its regular isosurfaces can only attain a finite number of different topological configurations, and that if we think of the isosurfaces as sweeping through the scalar values, the topology of these isosurfaces only change on a finite number of scalar values, known as *critical values*. In the (standard) case of smooth Morse theory, critical values are the values for which there exist points (the *critical points*) where the gradient of the scalar field vanishes. Just as importantly, it is possible to characterize the *type of topological change* by examining the behavior of the scalar field around the critical point. In the smooth case, this is given by the number of negative eigenvalues of the Hessian matrix evaluated at the critical points. The goal of this section is to show a similar classification for the case of piecewise-smooth manifolds. While this setting is not new, to the best of our knowledge it has not been presented in the context of scientific visualization. We review it here for completeness.

3.1 Stratified Morse Theory

The basic tool that we will use is *stratification*. Intuitively, we partition the piecewise-smooth manifold such that each of the subsets (called *strata*) is either zero-dimensional (that is, with a finite number of points) or has smooth structure. There are additional technical requirements that are irrelevant here, and we refer the reader to the work of Goresky and Macpherson for the full definition of what are known as *Whitney-stratified manifolds* [12]. In our case of trilinearly interpolated scalar fields, the stratification will always be the same. We partition the manifold in four sets, where each set is a (possibly disconnected) manifold of different dimensions. Set 0 contains all vertices of the grid, set 1 contains all edge interiors, set 2 contains all face interiors, and set 3 contains all cube interiors. We illustrate this in Figure 1 for the 2D case, but the 3D case is very much the same. More importantly, the level sets of the scalar field restricted to each of the subsets will be smooth (or not have a differential structure at all, in the case of the zero-dimensional stratum). The strategy of SMT is, intuitively, to use standard Morse theory on each of the strata and then combine the results in the correct way. Before we get there, we need some definitions.

3.2 Preliminaries

A *scalar field* f is a function that assigns a real value to each point in a space (in our case, a manifold representing the entire volume reconstructed by the trilinear interpolation). A stratification is a partition of the manifold into subsets (each subset is called a *stratum*, pl. *strata*). In our case, the stratification will be such that the restriction of f to each stratum is differentiable (or the stratum is zero-dimensional). We will restrict ourselves to fields where the gradient of the scalar field is zero in only a finite subset of the manifold, where the gradient is taken on the differential structure of the strata. These are the *Morse functions* of SMT. The *critical points* in SMT are the set of points with zero gradient, together with *all* points in the zero-dimensional stratum. As we will see shortly, not all critical points in stratified Morse theory induce topological changes (unlike what happens in the smooth case); see Figure 2 for an illustration. For Morse functions, the critical points will be isolated: they can always be separated by disjoint neighborhoods in the topology of M. If there exists a scalar value k for which some critical point x satisfies $f(x) = k$, then k will be called a *critical value* of f. The set of points x in M that satisfy $f(x) = k$ will be called an isosurface of M. If k is not a critical value, it is called a *regular value*, and in that case the set of points will be a submanifold of M. Let a and b be two regular values. If there are no critical values c such that $a < c < b$, then the isosurfaces given by a and b are homeomorphic. It follows that if we think of a changing isosurface as it "sweeps" through the scalar values of f, the only situation where this isosurface can change topology is as it goes through a critical value.

Another feature of Morse theory is the ability to predict exactly the type of topological change a critical point induces. One way to look at this in the smooth setting is to consider the behavior of the isosurfaces around a critical point and value. Consider a closed neighborhood around each critical point small enough such that each neighborhood only contains a single critical point, and call it $\varepsilon(p)$. We will call the boundary of $\varepsilon(p)$ the *link* of a point, and denote it by $L(p)$. We call the subset of the link with scalar value less (resp. greater) than $f(p)$ the *lower (upper) link* of p, and denote it by $L^-(p)$ ($L^+(p)$). We then define the *lower (upper) Morse data* to be the pair created from

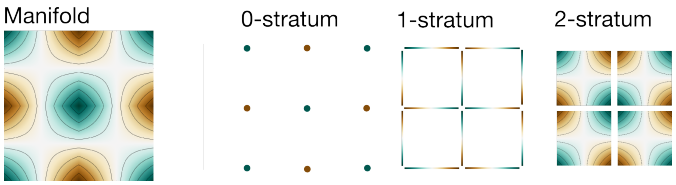

Fig. 1. An illustration of a piecewise-smooth immersed manifold. The colormap illustrates the height of each point scalar field. Notice that although the manifold itself is not everywhere differentiable, each strata is itself an open manifold that is differentiable.

the neighborhood and its appropriate link: $M^-(p) = (\varepsilon(p), L^-(p))$ and $M^+(p) = (\varepsilon(p), L^+(p))$. The topology of the lower and upper links give exactly the behavior as the isosurfaces go through the critical point p. Intuitively, the changes induced by the critical point are topologically equivalent to removing the lower link from the isosurface, and attaching the upper link. This also gives a good intuition as to why regular values do not change the topology of an isosurface: around a point, their "change" can be characterized by removing a disk and then simply adding back another disk. In the case of smooth Morse theory, every critical value changes the topology of the isosurface. As we will see, some critical points in SMT actually have a local structure that is topologically equivalent to the one of regular points, and will *not* change the topology of the level set. For any Morse data (A, B), it is always the case that $B \subset A$. While including A might seem unnecessary here, it will be of great calculational use in the stratified setting.

3.3 Characterizing critical points

In SMT, critical points have similar notions of upper and lower Morse data and links, which induce analogous topological changes to the isosurfaces. Their definition is somewhat more involved, and this is what we now turn to. This subsection follows the presentation in the monograph by Goresky and MacPherson [12].

We first state an (intuitive) property of stratified manifolds that we will use without proof. For an n-dimensional Whitney stratified manifold M and a point p in a d-dimensional stratum of M, it is possible to find a $(n-d)$-dimensional submanifold of M (which might straddle many strata) that meets transversally at p, and whose intersection consists of only p. It is, in a sense, the "topological orthogonal complement" to the neighborhood of p restricted to the stratum where p resides. Examples can be seen in Figure 3. Define the lower and upper *tangential* Morse data $T^-(p) = (\varepsilon_T(p), L^-(p))$ to be the smooth lower and upper Morse data of p, where $\varepsilon_T(p)$ is in the submanifold defined by the stratum where p resides. Similarly, define lower and upper *normal* Morse data $N^-(p) = (\varepsilon_N(p), L^-(p))$ to be the analogous notion, but the lower and upper links are taken to be subsets of $\varepsilon_N(p)$, itself a subset of the $(n-d)$-dimensional submanifold transversal to the stratum of p going through p. We are now ready to give the definition of the stratified Morse data. We omit here the signs that indicate lower and upper Morse data, since the definition is the same for both cases. Given $T(p) = (A, B)$ and $N(p) = (A', B')$, the stratified Morse data is given by the topological product of the two pairs, $M(p) = T(p) \times N(p)$, which in our case works out to be

$$M(p) = (A \times A', A \times B' \cup A' \times B), \tag{1}$$

where $\times$ denotes the topological product between two spaces. Since this definition does not give much intuition for its behavior, we now present a set of examples.

Zero-dimensional strata If a critical point p lies on a zero-dimensional stratum, then $T(p) = (\cdot, \emptyset)$. In this case, the normal Morse data $N(p) = (A', B')$ is simply the Morse data for the entire manifold, and so is the stratified Morse data, since

$$M(p) = (\cdot \times A', \cdot \times B' \cup A' \times \emptyset) = (A', B').$$

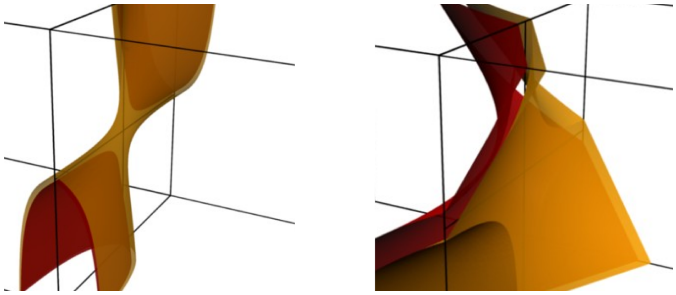

Fig. 2. Two cases of stratified Morse data illustrated. The left case is topologically equivalent to the saddles in smooth Morse theory, even though the critical point is not differentiable. The right example is a *virtual* critical point, which only happens in stratified Morse theory: a critical point that does not actually change the topology of the isosurface.

Partially non-smooth critical points Consider a point p in an internal face of the trilinear interpolation grid such that the gradient in p evaluated on the stratum is zero. In this case, $T^-(p) = (\square, \shortmid\shortmid)$, since we know that critical points in bilinear scalar fields are necessarily saddles. Let's examine the three (non-degenerate) possibilities for the normal Morse data.

- Case 1: p is a minimum along the transversal 1-disk, and so $N^-(p) = (-, \emptyset)$. In this case,

$$\begin{aligned} M^-(p) &= (\square \times -, (\square \times \emptyset) \cup (- \times \shortmid\shortmid)) \\ &= (\square, \emptyset \cup \square) = (\square, \square) \end{aligned}$$

- Case 2: p is a maximum along the transversal 1-disk, and so $N^-(p) = (-, \cdot)$. In this case,

$$\begin{aligned} M^-(p) &= (\square \times -, (\square \times \cdot) \cup (- \times \shortmid\shortmid)) \\ &= (\square, \square \cup \square) = (\square, \square) \end{aligned}$$

- Case 3: p is neither a minimum or a maximum along the transversal 1-disk, and so $N^-(p) = (-, \cdot)$. In this case,

$$\begin{aligned} M^-(p) &= (\square, (\square \times \cdot) \cup (- \times \shortmid\shortmid)) \\ &= (\square, \square) \end{aligned}$$

We illustrate these cases in Figure 2. Note that case 1 is homeomorphic to a saddle of type 1 in smooth Morse theory. Case 2 is, correspondingly, a saddle of type 2. Case 3, on the other hand, is more interesting. It would correspond to computing the Morse data for a regular point in smooth Morse theory. If we were to compute the stratified upper Morse data for case 3, we would find that the second element of the pair is again a disk. From the interpretation given in Section 3.2, we see that although p is strictly speaking a critical point, it does not actually induce a change in the topology of the manifold. In this paper, we will refer to such critical points as *virtual* critical points. In the next section we will see how to determine exactly which critical points are virtual.

Sequential stratification Our final example illustrates a technique that will be critically important for the characterization of Section 3.4. As we have explained, the idea of SMT is to compute the Morse data in parts, by combining the tangential Morse data with the normal Morse data. In the calculations we just did, the tangential Morse data has been given by smooth Morse theory, and the normal Morse data has been one-dimensional. This does not need to be the case. In fact, both tangential and normal Morse data might have been calculated as stratified Morse data, as illustrated in Figure 3. This allows us, in particular, to explicitly calculate the topology of the Morse data of critical points in the zero-dimensional stratum.

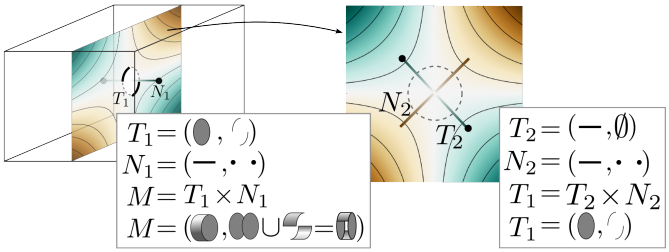

Fig. 3. Sequential stratification allows us to break complicated Morse data into simple ones which can be combined with Equation 1. The full Morse data M is given by $M = T_1 \times N_1 = (T_2 \times N_2) \times N_1$ T_2, N_2 and N_1 are all one-dimensional Morse data, with zero-dimensional, discrete, links, whose Euler characteristic can be trivially computed.

$\chi(L)$
1 $value \leftarrow \text{HEAD}(L)$
2 $rest \leftarrow \text{TAIL}(L)$
3 **return** $\begin{cases} value, & \text{if } rest = [] \\ \chi(rest), & \text{if } value = 0 \\ 1, & \text{if } value = 1, \\ 2 - \chi(rest), & \text{if } value = 2 \end{cases}$

Fig. 4. A simple algorithm to compute the Euler characteristic of a link. L is the signature of the link: the cardinalities of the one-dimensional links in the sequential stratification.

3.4 Computing the change in χ

Ultimately, we want to use Stratified Morse theory to make topological statements about the isosurfaces of piecewise smooth scalar fields such as the ones generated by trilinear interpolation. In this section, we provide a novel calculation that determines the change in the Euler characteristic of a level set, as it goes through a critical point. Because we can identify all critical points and their isovalues, this formula can be used to easily compute the Euler characteristic of a level set of the scalar field. The Euler characteristic, denoted as the function $\chi : M \to Z$, is a well-known invariant of topological spaces. In the case of compact, connected, orientable 2-manifolds without boundary, it gives an exact classification: two such surfaces are homeomorphic if and only if they have the same Euler characteristic. Although the level sets of our scalar fields are orientable, they are not necessarily connected or without boundary. The Euler characteristic, then, does not completely characterize the topology of the level sets in our case.

We now show how to compute the change in Euler characteristic induced by a single critical point, as a function of its sequential stratification. First, we recall that, for compact spaces, χ obeys the inclusion-exclusion principle

$$\chi(X \cup Y) = \chi(X) + \chi(Y) - \chi(X \cap Y). \tag{2}$$

Now, notice that the change in χ between an isosurface $S^- = \{x : f(x) = k - \varepsilon\}$ and $S+ = \{x : f(x) = k + \varepsilon\}$ for critical point p, associated critical value k and a sufficiently small $\varepsilon > 0$, denoted by $\Delta\chi(p)$, is given by

$$\Delta\chi(p) = \chi(L^+(p)) - \chi(L^-(p)). \tag{3}$$

This is because one can partition S^- and S^+ in two sets each, such that one of the sets in each partition is completely inside the small $\varepsilon-$ball. Since the sets outside this ball have no critical point, they are homeomorphic, and so share the same χ.

Let the critical point have a stratification such that either the lower or upper Morse data $M(p)$ be given by the tangential pair (T_ε, T_L) and the normal pair (N_ε, N_L). The following simple calculation gives $\chi(T_\varepsilon \times N_L \cup N_\varepsilon \times T_L)$, which is the Euler characterist of the lower (upper) link requited in 3 (see Equation (1)) . First, from Equation 2, $\chi(T_\varepsilon \times N_L \cup N_\varepsilon \times T_L) = \chi(T_\varepsilon \times N_L) + \chi(N_\varepsilon \times T_L) - \chi(T_\varepsilon \times N_L \cap N_\varepsilon \times T_L)$. Then, notice that for the manifolds we are interested in, T_ε and N_ε are always disks. We know that if d is a disk, $\chi(d) = 1$. In addition, if A and B are topological spaces with well-defined χ, $\chi(A \times B) = \chi(A)\chi(B)$. Then,

$$\chi(T_\varepsilon \times N_L \cup N_\varepsilon \times T_L) = \chi(N_L) + \chi(T_L) - \chi(T_\varepsilon \times N_L \cap N_\varepsilon \times T_L)$$

Now, partition both T_ε and N_ε in two subsets such that T_L and N_L are elements of the partition. That is, define $T_L \cup T_r = T_\varepsilon, T_L \cap T_r = \emptyset$, and similarly for N_ε and N_L.

Then, expand the partitions and products, and distribute the intersections around the unions, noticing all but one of intersections will be empty:

$$\begin{aligned} T_\varepsilon \times N_L \cap N_\varepsilon \times T_L &= ((T_r \cup T_L) \times N_L) \cap ((N_r \cup N_L) \times T_L) \\ &= ((T_r \times N_L) \cup (T_L \times N_L)) \cap \\ &\quad ((N_r \times T_L) \cup (N_L \times T_L)) \\ &= N_L \times T_L \\ \chi(T \times N_L \cap N \times T_L) &= \chi(N_L \times T_L) \\ &= \chi(N_L)\chi(T_L) \end{aligned}$$

which gives the final result

$$\chi(T_\varepsilon \times N_L \cup N_\varepsilon \times T_L) = \chi(N_L) + \chi(T_L) - \chi(N_L)\chi(T_L) \tag{4}$$

To compute $\chi(M(p))$, we apply Equation 4 recursively on a sequential stratification of p. If the scalar function we are considering is Morse, we can always find one such stratification. In that case, the simplest Morse data will be the ones such that the first element of the pair is a 1-disk. In that case, the lower and upper links will always be discrete spaces with zero, one or two points, for which computing χ is trivial (it is equal to the cardinality of the set). For internal points, the sum of the cardinality of lower and upper links in the one-dimensional strata is 2. For boundary points, the sum is 1. We then collect these cardinalities of the upper and lower one-dimensional Morse links into a tuple, and call this tuple the *signature* of a link. The upper and lower signatures contain much of the topological information about the critical point, and in particular make the computation of the Euler characteristic of the links trivial, as shown in Figure 4.

One attractive feature of this technique is that it can determine the change in χ for links in arbitrary dimensions. In addition, the same formula works for points in the boundary of the manifold. We present the possible kinds of critical points along with their changes in χ as computed by our technique in Figure 5. Let us work through an example. Consider the critical point depicted in the middle column of Figure 5. We look at its stratification as a combination of a saddle point along the 2D face (call these dimensions x and y), and a minimum in the remaining dimension, z. Notice that along the z dimension, the critical point lies in the boundary of the Morse data, so the links will have either zero or one points. On the 2D face, the saddle is such that it is the combination of a minimum in x and a maximum in y. This means that the first two elements of the signature of L^- are 0 and 2, while the signature of L^+ starts with 2, 0. Plugging these two sequences in Algorithm 4 gives that $\chi(L^-) = 2 - \chi(L_z^-)$ and $\chi(L^+) = 2 - \chi(L_z^+)$, where L_z^+ and L_z^- denote the one-dimensional links along the z dimension. Since the cardinality of L_z^- is 0, $\chi(L^-) = 2 - 0 = 0$. Similarly, from the cardinality of L_z^+ being 1, we have that $\chi(L^+) = 2 - 1 = 1$ and, $\Delta\chi = 1 - 2 = -1$.

With the machinery we have presented, we can identify and characterize all critical points belonging to the zero-dimensional stratum of M. We now turn to the other strata. For Morse functions such as the ones we are interested, all points in the one-dimensional stratum will be regular, since the scalar values on vertices have to be different in order to isolate vertex critical points. In the next sections, we will identify and characterize critical points on the two and three-dimensional

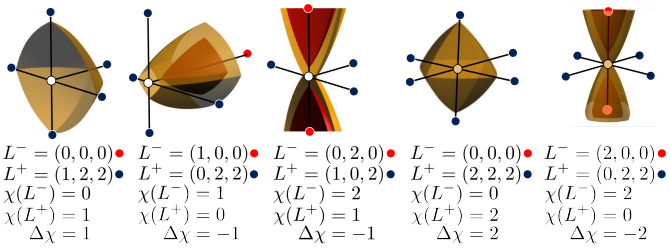

Fig. 5. All possible critical point types for piecewise smooth level sets of a stratified 3-manifold. The leftmost three types of critical points occur in the boundary of the manifold, and affect the topology of boundary loops of the level set. The rightmost two types are well-known types of internal critical points. Each of these has a negative counterpart which switches lower and upper links.

strata, and just as importantly, we will relate their presence to the topological correctness of Marching Cubes cases. Finally, once we can compute the change in χ for every critical point, we can get the Euler characteristic of a regular isosurface $s(k) = \{x | f(x) = k\}$ in the volume by summing all $\Delta\chi$ values of critical points p with isovalues $v(p)$ less than k:

$$\chi(s(k)) = \sum_{v(p)<k} \Delta\chi(p)$$

3.5 The level-set topology of trilinearly interpolated scalar fields

In this section we will present a topological characterization of the topology of level sets in a scalar field representing the simplest non-trivial three-dimensional regular grid: a single cube. Given scalar values at the cube's vertices and an isovalue, each Marching Cubes case is characterized by the relative sign of the vertices (below or above the isovalue), or equivalently by its *active edges*: the set of edges that contain the isovalue.

Definition 1 *A Marching Cubes case is* ambiguous *if, for a given isovalue, two different sets of scalar values generate the case, but non-homeomorphic level sets.*

Our formulation pinpoints exactly which trilinearly interpolated scalar fields will have no critical points other than one single minimum and one single maximum, both at vertices of the cube. In these cases, SMT tells us that the only possible topology for a level set is a disk. In addition, together with the argument of Section 4, we can then say that if no cube is ambiguous, Marching Cubes will generate triangular meshes homeomorphic to the original surface.

Since there are many different variants of Marching Cubes, we clarify by picking a specific one. We choose to use Bhaniramka et al's technique [4]. Since their tables are based on a simple algorithmic construction around the convex hull, it is straightforward to see that the cases it generates for active cells with simple edge flows are, in fact, homeomorphic to disks. While all Marching Cubes algorithms which rely solely on active edges for its decision will have ambiguous cases, every MC table that we are aware of appears to, in fact, handle unambiguous cases correctly. Bhaniramka et al's technique is very attractive in this case because its elegant formulation allows us to skip individual case-based arguments. This greatly reduces the effort for proving the correctness of the full algorithm.

Given values v_{ijk}, where $i,j,k \in \{0,1\}$, a scalar field $f(x,y,z) : [0,1]^3 \to R$ representing a trilinear interpolation is given by

$$f(x,y,z) = axyz + bxy + cxz + dyz + ex + fy + gz + h,$$

where $a = -v_{000} + v_{001} + v_{010} - v_{011} + v_{100} - v_{101} - v_{110} + v_{111}$, $b = v_{000} - v_{001} - v_{010} + v_{011}$, $c = v_{000} - v_{001} + v_{100} - v_{101}$, $d = v_{000} - v_{010} - v_{100} + v_{110}$, $e = -v_{000} + v_{001}$, $f = -v_{000} + v_{010}$, $g = -v_{000} + v_{100}$, $h = v_{000}$. Our main tool in this paper is what we call an *edge flow*. It succinctly represents much of the behavior of the level sets of the cube. An edge flow G is obtained by interpreting the vertices and edges of the boundary of the cube as a directed graph, such that if $f(a) < f(b)$ for adjacent vertices a and b, G has an edge from a to b. For now we assume vertices have different associated scalar values; we later show how the same analysis easily extends to the more general case. We call a vertex a *minimum* if the edge flow at that vertex has in-degree zero (that is, no edge points to it). Similarly, we call a vertex a *maximum* if the edge flow has out-degree zero. If an edge flow has exactly one minimum and one maximum, we call it a *simple flow*. These have direct interpretations as critical points in Stratified Morse Theory which behave in the analogous way to minima and maxima in the smooth case. Our main theorem of the section is:

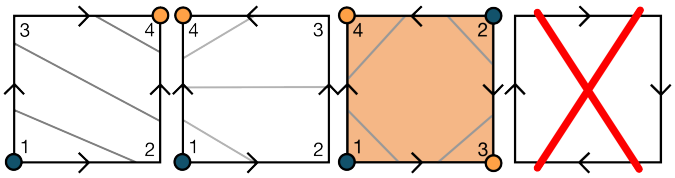

Fig. 6. The three possible edge flows for Marching Squares. The numbers next to the vertices illustrate vertex values that would induce those flows (other vertex values are also clearly possible). The internal lines represent approximations of some of the isosurfaces. The crossed-out figure to the far right represents an assignment of edge directions that cannot happen for any scalar fields.

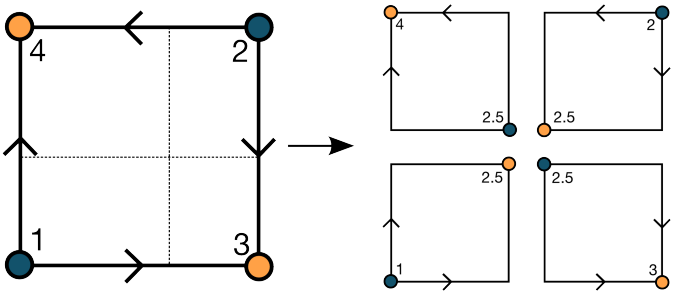

Fig. 7. Splitting an edge flow. In the two-dimensional case, splitting an ambiguous edge flow along the asymptotes creates four unambiguous subcases.

Theorem 1 *Let G be an edge flow of a scalar field f induced by a trilinear interpolation with given values at vertices. Marching Cubes will generate homeomorphic triangle meshes for every isosurface in a cube if and only if its edge flow is simple.*

In particular, we will break the main proof in two lemmata:

Lemma 1 *If G is simple, then f will have no face saddles.*

Lemma 2 *If G is simple, then f will have no internal critical points.*

An edge flow is associated with the entire scalar field, and not with a particular isovalue or active edge configuration. While some of the outputs of Marching Cubes might be correct for non-simple edge flows, for at least one isovalue there will be ambiguous configurations. Before we set out to prove the 3D case, we illustrate it with the much simpler 2D case, since Marching Squares has much of the same structure as Marching Cubes, and its ambiguities can be identified in a similar way. We use it to introduce some concepts we will need later.

3.6 A simplified scenario: Marching Squares

Marching Squares is the folklore 2D analog of Marching Cubes, and provides a simple setting to explore the ambiguity issues in the level sets of linear interpolation. We can define an edge flow for 2D cells in the same way as the three-dimensional case. The ambiguity criterion is also similar: a 2D cell will cause an ambiguous case in Marching Squares if and only if it is not simple. This is illustrated in Figure 6, which includes all possible edge flow cases modulo symmetries of the hyperoctahedral group.

The first feature of edge flows we highlight is that not all edge direction assignments represent valid edge flows. In particular, the flow must be acyclic, since such a cycle would violate Stokes's theorem applied to the closed curve representing the cycle (intuitively, it would represent an "infinite climbing" scenario that no scalar field can satisfy). We call this the *cycle restriction*. Another observation is that for all simple 2D edge flows, every nonempty regular isosurface has the topology of a closed disk. If we think of the scalar field sweeping through the square, Morse theory tells us that the topology of a level set can only change as it passes through *critical points*. In addition, minima represent the creation of isosurface pieces and maxima represent their destruction. It is no coincidence then that our criterion requires exactly one maximum and one minimum. As we illustrate below, in these cases there will be no internal critical points.

Consider the necessary conditions for a saddle point to appear inside a square (assuming distinct vertex values). First, note that partial derivatives of linear scalar fields are themselves multilinear scalar fields, in addition constant along the coordinate of the taken derivative. This means that any axis-aligned line will be monotonic with respect to the field. Now consider one of the partial derivatives around the saddle point. Since it is zero at the saddle and itself a linear scalar field (hence monotonic), it will change signs when crossing the zero-point. This means that the two edges corresponding to the flow must have opposing directions, since otherwise the partial could not be zero in the interior. The same argument applies for both edge pairs, and so it must be that an internal critical point (the saddle) happens only if the corresponding edge flow contains two minima and two maxima; the only other possible configuration with opposing edge directions violates the cycle restriction. In addition, in these cases there will be no more than one internal critical point.

This is a complete characterization of the topology of level sets of bilinearly interpolated grids. For edge flows 0 and 1, monotonicty implies that at least one of the partial derivatives will be non-zero inside the square, and so there cannot be an internal critical point. For edge flow 2, there must be exactly one internal critical point. This characterization is combinatorial: it pinpoints the presence of an internal critical point solely based on the configuration of the edge flow. Alas, the three-dimensional case will not be as simple: there are configurations with the same edge flow but with a differing number of internal critical points. Still, the main argument holds: if the edge flow is simple, there will be no saddles or internal critical points.

We introduce one final concept that will be useful in the three-dimensional case, that of *splitting* edge flows. Linear interpolations have the following *reproducing* property: it is possible to create a finer grid with new scalar values such that the union of the scalar fields over the finer grid gives exactly the same values as the original interpolation [21], and the new vertex values are given by the coarser scalar field evaluated at the vertices. We use this to define new edge flows from old ones, by splitting them in multiple squares, as illustrated in Figure 7.

Now consider what happens when we split an edge flow exactly through an asymptote. First, notice that one edge has vertices with the same scalar values. In these cases, we will consider all connected vertices with the same value as one single entities, and collapse the edges of the flow with the same value. The second, more important result, is that the resulting subflows will *always be simple*. For the case of Marching Squares, if we come across an ambiguous case, it suffices to split the edge flow along the asymptotes, generate the unambiguous Marching Squares solutions, and merge them together. Since simple edge flows always indicate Marching Squares cases that get the correct topology, it will necessarily be the case that the final result will be one with the correct topology as well. This splitting in simpler cases is the central notion of our algorithm in Section 4. It relieves us from having to consider all possible complicated cases of the interpolant. All that is required is the ability to identify ambiguous cases, to split the complicated cases into simpler ones, and to glue them back appropriately.

Edge flows also provide an arguably simpler interpretation to the behavior of the asymptotic decider [22]. One can think of the asymptotic decider as performing a data-dependent subdivision of the interpolation grid into a subgrid where the isosurface decision becomes trivial. The particular Marching Squares case will be determined by the scalar value at the intersection of the asymptotes, which is exactly the criterion employed by the asymptotic decider. The idea of considering the direction of the gradient flow across edges of a cubic cell has, to the best of our knowledge, been originally proposed by Carr [5]. We believe our application of the idea in the arguments which follow is simpler, and, in particular, it shows more directly that the absence of face saddles forces the simplicity of the flow in the cube.

3.7 Detecting critical points in the three-dimensional case

With the intuition of the Marching Squares case in our belt, we turn to the full three-dimensional case. We want to determine which configurations generate what kinds of critical points in the two- and three-dimensional strata of M. We start with a full enumeration of possible edge flows in 3D. For that, we use the technique described by Banks et al., originally developed in the context of counting cases in Marching Cubes-like algorithms [3]. Specifically, we model each edge flow as a string of twelve binary digits, one for each edge, where "0" indicates flow in one direction, and "1" indicating flow in the other direction. The starting set of edge flows is computed by picking all possible total orders on the vertices of the cube ($8! = 40320$ configurations) and computing the associated edge flows. We then determine the action of the hyperoctal group on the edge flows, and find the orbits under this group action, also using GAP [1]. The result is a total of 54 cases, shown in Figure 8.

It is worthwhile to inspect that table carefully. Cases 18 through 53 all have face saddles (the faces are shaded green in the figure), which correspond to an internal critical point in a Marching Squares case. This immediately means that the corresponding Marching Cubes case must be ambiguous, since the boundary manifold for the cube has at least two different configurations. Notice as well that cases 18 through 53 are exactly the cases that fail to have exactly one minimum and one maximum. This means that having exactly one minimum and one maximum is a necessary condition for Marching Cubes to be unambiguous on that edge flow. Finally, note that this table gives a trivial computational proof of Lemma 1:

Proof The only simple edge flow cases are 0–17. None of these cases have faces with saddles. QED.

To prove Lemma 2, we again inspect cases 0–17 and notice the following property. Partition the edges in an edge flow with respect to their directions. In all but case 4, at least one partition has every edge pointing in the same direction. Remember that the partial derivative of a trilinearly interpolated scalar field is also a trilinearly interpolated scalar field. If the edges all point in the same direction, the partials along the edges all have the same sign. By monotonicity, that partial derivative must be nonzero everywhere in the interior.

We are left with one final case. The proof that case 4 cannot have an internal critical point proceeds by contradiction, and is illustrated in Figure 9. We assume that there exists an internal critical point in a configuration, and create a splitting edge flow such that the internal vertex lies exactly on the critical point. By definition, all partial derivatives vanish, which is identical to saying that the six edges that touch the internal vertex (shown as dashed lines) have no direction. The partial derivatives are themselves linearly interpolated scalar fields (and so are monotonic along any fixed line), and this means that *edge pairs that are parallel and in opposite sides of the dashed lines must have opposite directions*. In addition, monotonicity also implies that if two parallel external edges in a flow splitting have the same direction, then

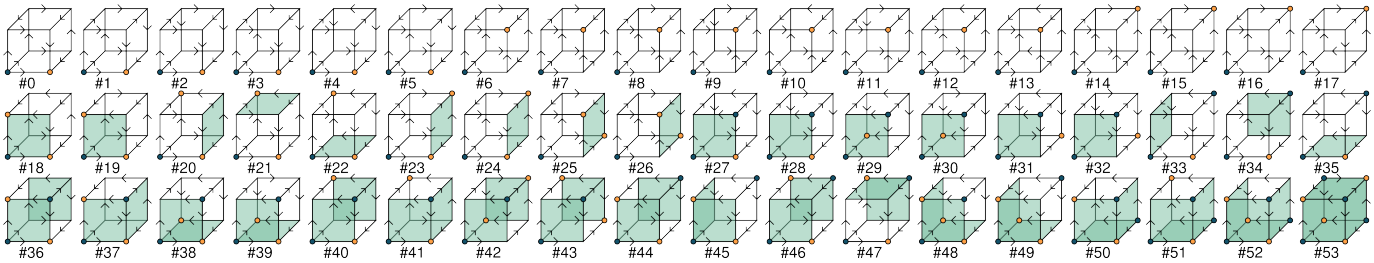

Fig. 8. All possible 54 discrete edge flow cases for trilinear interpolation. Cases 0 through 17 have one single minimum and one single maximum. As we show in the text, in this case the trilinear interpolant has no internal critical points. In addition, notice that cases 18 through 53 all have face saddles, denoted here by shaded faces. This means, then, that internal critical points require the presence of at least one face critical point.

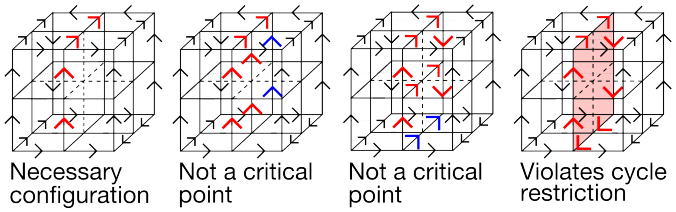

Fig. 9. An illustration of part of the proof of Lemma 2.

so does the internal edge. These two considerations together with the cycle restriction form a cascade of edge direction assignments, which illustrate in Figure 9. Of the twelve internal pairs of edges, six of them are determined by the monotonicity requirement, only two of which we need to show a contradiction: these are the red arrowheads in the figure. Now notice what happens when we ponder the possible directions if the edges opposite the red arrowheads. In all but one configuration monotonicity tells us that one of the central edges must have a well-defined direction, and hence a non-zero partial derivative. However, the remaining configuration violates the cycle restriction and so is impossible. With this, we can prove Lemma 2:

Proof Cases 0–3 and 5–17 cannot possibly have an internal critical point since at least one of the partial derivatives is nonzero. Splitting the edge flow of Case 4 using the critical point as the central vertex shows that an internal critical point would violate the cycle restriction, creating a contradiction, QED.

These give us a proof of Theorem 1:

Proof That it is necessary for the flow to be simple in order for Marching Cubes to be unambiguous follows from the contrapositive of Lemma 1, as follows. If there is a face saddle, then there exist two nonhomeomorphic surfaces with the same MC case, but face saddles only occur in non-simple edge flows, and so unambiguous MC cases can only appear if the flow is simple. Lemmas 1 and 2 together characterize all critical points on strata other than the zero-dimensional one. Since simple flows have a single minimum and a single maximum, by Morse theory it follows that the isosurface inside a cube with simple flow must be topologically a disk. In addition, for those cases, it is easy to see that Bhaniramka et al.'s MC table [4] will generate a disk, and so Marching Cubes will return the correct topology.

3.8 Classifying face and internal critical points in the three-dimensional case

The above proof already suggests a simple algorithm for extracting an isosurface that is homeomorphic to the trilinearly interpolated regular grid, by splitting an ambiguous case along the center of the cell, until all the active cells have simple edge flows, and then gluing them all back together correctly. The algorithm we describe in Section 4 works in a very similar fashion. However, to limit the size of the final mesh, we want to keep the number of necessary subdivisions to a minimum, and so we will try to find exactly where the critical points are, since splitting flows on those vertices will necessarily produce simple flows.

Face critical points will happen whenever a face contains a saddle configuration. Their exact position can be found by setting the gradient of f, namely

$$\nabla f(x,y,z) = \begin{pmatrix} e+by+cz+ayz \\ f+bx+dz+axz \\ g+cx+dy+axy \end{pmatrix},$$

to zero, and solving the resulting equation in the cases $\{x,y,z\} = \{0,1\}$. The critical points and values at faces are easily solved analytically by setting the partial to zero and solving. To determine the kind of critical point, we appeal to SMT and note that the lower and upper tangential Morse data will be a pair of lines. The lower link signature will be $(0,2)$, and the upper link signature will be $(2,0)$ (or vice-versa). The signature of the normal Morse data can be found by evaluating the trilinear interpolation in a line orthogonal to the face, and then the output of algorithm of Figure 4 can be used to get which type of saddle the critical point is.

Internal critical points will happen when all three partials are simultaneously zero in the interior of the cube. Let $\Delta_x = bc-ae$, $\Delta_y = bd-af$, and $\Delta_z = cd-ag$. We find those points by directly solving the resulting quadratic equation in x, y and z, and get two critical points:

$$\begin{aligned} x &= \frac{d\Delta_x \pm \sqrt{\Delta_x\Delta_y\Delta_z}}{a\Delta_x} \\ y &= \frac{c\Delta_y \pm \sqrt{\Delta_x\Delta_y\Delta_z}}{a\Delta_y} \\ z &= \frac{b\Delta_z \pm \sqrt{\Delta_x\Delta_y\Delta_z}}{a\Delta_z}. \end{aligned}$$

If any of Δ_x, Δ_y, or Δ_z are zero, the critical point will be degenerate, and neither smooth or stratified Morse theory are be able to characterize the critical point. If $a = 0$, then there is a single solution

$$\begin{aligned} x &= -\frac{-de+cf+bg}{2bc} \\ y &= -\frac{de-cf+bg}{2bd} \\ z &= -\frac{de+cf-bg}{2cd}. \end{aligned}$$

To classify these internal critical points, we use the standard Morse theory technique of looking at the signs of the eigenvalues of the Hessian evaluated at the critical point. We now describe how to compute the eigenvalue signs without actually solving for the eigenvalues, which reduces the risk of numerical errors.

coefficient signs			eigenvalue signs		
-	-	-	-	+	+
-	-	+	-	-	+
-	+	+	-	+	+
+	+	+	-	-	+

Table 1. The signs of the Hessian eigenvalues as a function of the signs of the coefficient matrix.

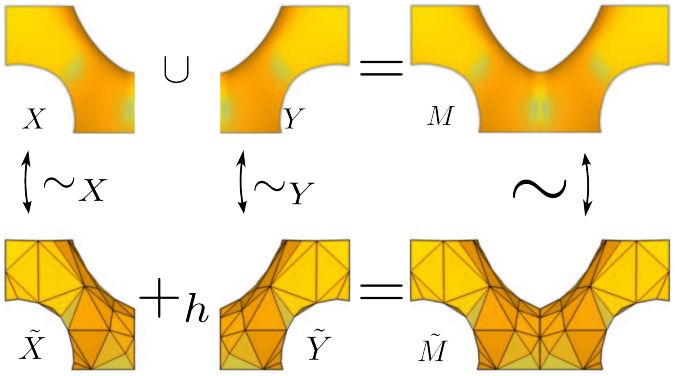

Fig. 10. An illustration of Lemma 3, which gives us a way to build a large homeomorphism by small parts.

3.9 The index of a critical point, robustly

The number of negative signs of eigenvalues gives us the critical point's index. Let us first look at the simpler case of the Hessian when $a = 0$. We will denote the Hessian evaluated at the appropriate critical point by H:

$$H = \begin{pmatrix} 0 & b & c \\ b & 0 & d \\ c & d & 0 \end{pmatrix}$$

Notice that the matrix is symmetric. In addition, because the determinant is $|H| = 2bcd$, it must also be non-singular, since if any of b, c and d were zero the solution would not be well-defined. This means that all eigenvalues are nonzero; and since the matrix is traceless, not all eigenvalues can have the same sign. From Morse theory, the critical point will be a saddle. However, the kind of saddle crucially depends on the eigenvalue signs. To that effect, note that roots are continuous functions of the polynomial coefficients (ie. eigenvalues are continuous functions of the characteristic polynomial). We know that all roots are non-zero, so in order for a root to become positive, it has to go through zero, which means that one of b, c and d must also change signs, since there will be a zero eigenvalue iff $|H| = 0$. So the configuration of eigenvalue signs is completely determined by the configuration of signs of b, c and d, as given by Table 1. If either b, c or d are zero, then again f is not a Morse function and we cannot use Morse theory to predict the topology of the isosurfaces. The case where $a \neq 0$ is very similar. In this situation, the Hessian is

$$H = \begin{pmatrix} 0 & \pm\frac{\sqrt{\Delta_x\Delta_y\Delta_z}}{\Delta_z} & \pm\frac{\sqrt{\Delta_x\Delta_y\Delta_z}}{\Delta_y} \\ \pm\frac{\sqrt{\Delta_x\Delta_y\Delta_z}}{\Delta_z} & 0 & \pm\frac{\sqrt{\Delta_x\Delta_y\Delta_z}}{\Delta_x} \\ \pm\frac{\sqrt{\Delta_x\Delta_y\Delta_z}}{\Delta_y} & \pm\frac{\sqrt{\Delta_x\Delta_y\Delta_z}}{\Delta_x} & 0 \end{pmatrix},$$

and we can determine the signs of each term by looking at the signs of Δ_x, Δ_y and Δ_z and use Table 1 as well.

4 Algorithm

We now describe an algorithm that, given a regular grid with associated scalar values, and a regular isovalue (according to the definition given in Section 3.1), returns a manifold triangle mesh that is homeomorphic to the level set of the trilinearly interpolated grid. The algorithm is straightforward, and so is our proof of correctness. We will need a simple lemma which says, essentially, that if we split a manifold, we can recover the original topology if we glue the pieces back together correctly:

```
EXTRACTISOSURFACE(F, isovalue)
1  if F has a simple flow
2     then return MARCHINGCUBES(F, isovalue)
3     else lF, rF, commonface ← SPLIT(F)
4          lMesh ← EXTRACTISOSURFACE(lF, isovalue)
5          rMesh ← EXTRACTISOSURFACE(rF, isovalue)
6          return GLUE(lMesh, rMesh, commonface)
```

Fig. 11. A simple algorithm to extract an isosurface of a trilinearly interpolated scalar field.

Lemma 3 *Let M be a topological space, and let X and Y be such that $X \cup Y = M$. Let $\tilde{X}$ be homeomorphic to X via a homeomorphism defined by $\sim_X : X \to \tilde{X}$, and similarly with $\tilde{Y}$. Let $X \cap Y = I$, and $f : I \to \tilde{X}$ be the restriction of $\sim_X$ onto I, and $g : I \to \tilde{Y}$ be the restriction of $\sim_Y$ onto I. Let A be the image of f, and $h : A \to \tilde{Y}$ be such that $h(f(x)) = g(x)$. Let $\tilde{M}$ be the adjunction space $\tilde{X} +_h \tilde{Y}$ (that is, let $\tilde{M}$ be the space defined by "gluing $\tilde{X}$ and $\tilde{Y}$ together along h"). Then, M and $\tilde{M}$ are homeomorphic.*

Proof Let $\alpha : M \to \tilde{X} +_h \tilde{Y}$ be the function such that

$$\alpha(x) = \begin{cases} \sim_X (x), \text{if } x \in X - Y \\ \sim_Y (x), \text{if } x \in Y \end{cases}$$

α is clearly continuous and invertible. If $x \notin I$, then $\alpha(x)$ is continuous and has a continuous inverse there, from the assumptions on $\sim_X$ and $\sim_Y$. If $x \in I$, then since open sets of the quotient topology are the union of the open sets under the equivalence relation, the neighborhoods in I will include open sets from both $\tilde{X}$ and $\tilde{Y}$, and so $\alpha^{-1}(\alpha(x))$ will be continuous as well, giving the homeomorphism.

The argument is illustrated in Figure 10. With this in hand, we give the algorithm for generating a triangle mesh homeomorphic to the level set in Figure 11. The proof that this algorithm generates the correct topology is by a trivial induction on the possible cases. If the edge flow is simple, the base case of the recursion gives the correct topology. If the edge flow is not simple, SPLIT subdivides the scalar field, and GLUE implements the attaching map construction by zippering the two pieces of boundary together appropriately. Notice that the divide-and-conquer algorithm given above works for regularly interpolated three-dimensional scalar grids of any size. In a real-life implementation, one would want to use an acceleration structure such as interval trees to ensure fast runtimes [9]. In the source code we include as extra material, we perform a brute-force scan through the volume. Although this is clearly inefficient, the point of our reference implementation is to provide experimental evidence of the effectiveness with respect to correctness.

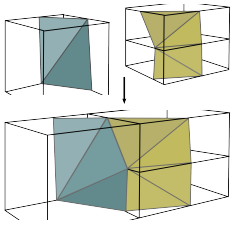

The procedure GLUE is illustrated on the right. It is quite simple: the idea is that every cell boundary of a trilinearly interpolated scalar field is homeomorphic to a set of at most two line segments. In addition, the observation by Nielson that trilinear functions can be represented as height functions on any axis [21] can be used to sort the vertices of any boundary in the correct zippering order. By splitting the edges

along the zippering, we can glue both of the surface patches. Edge splitting preserves the topology of a PL-manifold, and so the original homeomorphism is preserved. The attaching map then satisfies Lemma 3, because the sorting of the vertex ordering on both faces is the same.

4.1 Cases where Marching Cubes succeeds

We can use the previous observations to show a situation where Marching Cubes does, in fact, give a homeomorphic isosurface. We combine Theorem 1 with Lemma 3. Lemma 3 gives a condition for how the gluing between cells needs to be in order to generate a homeomorphic surface. Notice that in the case of simple edge flows, the faces, considered as flows themselves are also simple. This means that the intersection of the level set with the face will be a disc, and using the nomenclature of Lemma 3, $+_h$ is the function that glues one edge of the MC triangulation on one side to the edge of the MC triangulation on the other side. This is exactly the gluing that MC implementations perform; we have then shown that in such cases there is a full homeomorphism between the MC reconstruction and the level set.

5 Verification

We now describe the verification of our implementation. We use the word verification in the sense of verifiable visualization, described in detail by Kirby and Silva [14].

Geometric properties The analysis of the convergence of geometric properties presented by Etiene et al. applies directly to our algorithm, and so we would should expect quadratic order of accuracy for the Hausdorff distance to their manufactured solutions [11]. In addition, we should also expect linear accuracy on the normal directions, and asymptotic convergence of some order of surface area. Using that same analysis and the same manufactured solutions, we have observed an order of accuracy of 1.98 for the Hausdorff distance, an order of accuracy of 0.70 for the normal directions and area convergence of 2.02. The reason for the somewhat lower normal convergence remains to be investigated. We speculate that it might be arising out of numerical issues of small cells in the subdivision. The convergence plots are illustrated in Figure 12.

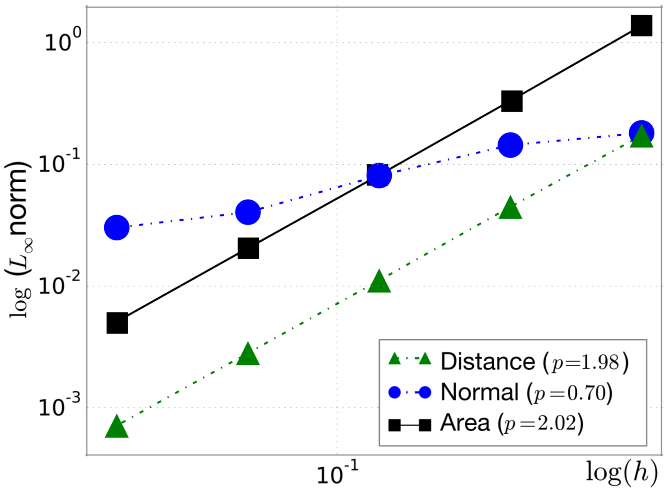

Fig. 12. Results of the analysis of the convergence of geometric properties using the method of manufactured solutions.

Topological properties In the same spirit of Etiene et al's analysis of geometric convergence properties, we believe it is important to perform a thorough investigation of the topological properties attained or not by isosurface extraction algorithms. We report here the results of verifying our reference implementation with three methods proposed in a separate report [26]. In a nutshell, these verification procedures generate a large number of scalar fields for which topological invariants of the level sets are easily computable. By checking whether these topological invariants are satisfied by the output of the algorithm, it is possible to quickly pinpoint subtle bugs in the implementation.

We employed this verification methodology throughout the development of the prototype, and we found it to be invaluable for spotting corner cases while debugging. The supplied implementation passed a large number of tests: we have checked on the order of 50000 isosurface outputs. These tests include checking if the isosurface is a piecewise-linear manifold, and if the result's Betti numbers (the ranks of the homology groups) and Euler characteristics match automatically generated test cases. For a full discussion of topological verification issues, we refer the reader to the additional material in the submission [26].

6 Discussion

Plantinga and Vegter's criterion for subdivision essentially guarantees that all gradient vectors point in the same direction, which allows them to (abstractly) build parameterizations of the isosurfaces and show an ambient isotopy [23]. Our criteria are much simpler, and although we believe will require in general fewer domain subdivisions to obtain provably disk-like isosurfaces, a full experimental evaluation is still necessary.

The phenomena of virtual critical points of stratified Morse theory has already been observed in the literature. In particular, Weber et al.'s Lemma 4 [29] provides an informal account of similar results as the ones given by our calculation. We feel that the technique of stratification presented here is more formal. In addition, it allows us to easily calculate the type of critical point, regardless or not it is in the boundary of the manifold (a case itself not handled in the above).

Both Stelldinger et al. and Van Gelder and Wilhelms present similar arguments for the presence of ambiguity in the cases in Marching Cubes [27, 25]. They show that if either of the negative or positive vertex sets are not connected (using the edges as the connectivity information), then the case will be ambiguous for that field and isovalue. As we have shown in Section 4.1, It is easy to show that if every grid cell is simple scalar fields for which every isovalue generates connected positive and negative sets have simple edge flows, and edge flows provide the added advantage of guiding the splitting of those grids so that Marching Cubes will generate correct results.

Lewiner et al. claim that it is enough to show a homeomorphism between the isosurface in two adjacent cubes and the reconstructed cases [16]. However, this misses the possibility that the gluing between the two cases is performed incorrectly (for example, it is clear that the disjoint union of both pieces will rarely give the right topology). Our Lemma 3 provides a condition on the gluing. In particular, we note that the authors of that implementation have recently uncovered a bug in the orientation of some pieces of the isosurface [15] exactly due to incorrectly gluing. Because of these subtleties, we believe that there is value in being exceedingly careful about the definitions and properties.

It would be interesting to search for a proof that does not refer to the table of possible edge flows, in particular since it gives us hope that such a proof might work for isosurfacing algorithms in any dimension. This would be an exciting development, since together with the automated case table generation of Bhaniramka et al [4] and Lemma 3, the appropriate generalization Theorem 1 would give us an almost complete description of an isosurfacing algorithm in any dimension with strong topological guarantees, together with a verification methodology for the implementation. However, we note that the $\Delta\chi$ method might not provide much information for isosurfaces of even dimension, since all isosurfaces of a scalar field of even dimension without boundary have Euler characteristic zero (via Poincaré duality). Although, implementing the GLUE procedure for boundaries of more than one dimension is not trivial, it might be possible to do it one dimension at a time: since Lemma 3 lets us build a high-dimensional homeomorphism by combining lower-dimensional constructions, it might be possible to construct the higher-dimensional map via simpler lower-dimensional ones.

7 Conclusions and Future Work

We have shown how to compute topological properties of the level sets of trilinear interpolants, and how to develop simple algorithms that extract topologically faithful isosurfaces. Stratified Morse theory both gives a framework to precisely state the property of these isosurfaces, and tools to verify the implementations of our algorithms (via the $\Delta\chi$ method). Our approach helps unify many different arguments for topological correctness in different published algorithms, and it can be seen as a way to Morse theory to explain the importance of the decomposition into disks. Practical verification of topological properties for more general interpolants remains an extremely important problem that is essentially open. One exciting avenue for future work is to use the ideas presented here to design such methods. In general, we believe that there is interesting work to be done in the interplay between designing reconstruction kernels on a grid and developing effective algorithms for extracting isosurfaces of that grid.

References

[1] Gap: Groups, algorithms, programming. a system for computational discrete algebra. `http://www.gap-system.org/`.

[2] N. Amenta and M. Bern. Surface reconstruction by voronoi filtering. In *SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry*, pages 39–48, New York, NY, USA, 1998. ACM.

[3] D. C. Banks, S. A. Linton, and P. K. Stockmeyer. Counting cases in substitope algorithms. *IEEE Trans. Vis. Comp. Graph.*, 10(4):371–384, 2004.

[4] P. Bhaniramka, R. Wenger, and R. Crawfis. Isosurface construction in any dimension using convex hulls. *IEEE Transactions on Visualization and Computer Graphics*, 10(2):130–141, 2004.

[5] H. Carr. No more marching cubes. In *Volume Graphics*, 2007.

[6] H. Carr and N. Max. Subdivision analysis of the trilinear interpolant. *IEEE Transactions on Visualization and Computer Graphics*, 16:533–547, 2010.

[7] H. Carr and J. Snoeyink. Representing interpolant topology for contour tree computation. In *TopoInVis 2007*, 2007.

[8] E. V. Chernyaev. Marching cubes 33: Construction of topologically correct isosurfaces. Technical report, Institute for High Energy Physics, Russia, 1995.

[9] P. Cignoni, P. Marino, C. Montani, E. Puppo, and R. Scopigno. Speeding up isosurface extraction using interval trees. *IEEE Transactions on Visualization and Computer Graphics*, 3(2):158–170, 1997.

[10] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. In *SCG '94: Proceedings of the tenth annual symposium on Computational geometry*, pages 285–292, New York, NY, USA, 1994. ACM.

[11] T. Etiene, C. Scheidegger, L. G. Nonato, R. M. Kirby, and C. Silva. Verifiable visualization for isosurface extraction. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1227–1234, 2009.

[12] M. Goresky and R. Macpherson. *Stratified Morse Theory*. Springer, 1988.

[13] A. Gyulassy, V. Natarajan, V. Pascucci, P. timo Bremer, and B. Hamann. Topology-based simplification for feature extraction from 3d scalar fields. In *Proceedings of IEEE Visualization*, pages 535–542, 2005.

[14] R. M. Kirby and C. T. Silva. The need for verifiable visualization. *IEEE Comput. Graph. Appl.*, 28(5):78–83, 2008.

[15] T. Lewiner. Private communication, Mar 2010.

[16] T. Lewiner, H. Lopes, A. W. Vieira, and G. Tavares. Efficient implementation of marching cubes' cases with topological guarantees. *journal of graphics, gpu, and game tools*, 8(2):1–15, 2003.

[17] H. Lopes and G. Tavares. Structural operators for modeling 3-manifolds. In *SMA '97: Proceedings of the fourth ACM symposium on Solid modeling and applications*, pages 10–18, New York, NY, USA, 1997. ACM.

[18] W. E. Lorensen and H. E. Cline. Marching Cubes: A High Resolution 3D Surface Construction Algorithm. volume 21, pages 163–169. ACM, 1987.

[19] Y. Matsumoto. *An Introduction to Morse Theory*. American Mathematical Society, 2001.

[20] T. S. Newman and H. Yi. A survey of the marching cubes algorithm. *Computers & Graphics*, 30(5):854 – 879, 2006.

[21] G. M. Nielson. On marching cubes. *IEEE Trans. Vis. Comp. Graph.*, 9(3):283–297, 2003.

[22] G. M. Nielson and B. Hamann. The asymptotic decider: resolving the ambiguity in marching cubes. In *Proceedings of the IEEE Visualization*, pages 83–91, Los Alamitos, CA, USA, 1991. IEEE Computer Society Press.

[23] S. Plantinga and G. Vegter. Isotopic approximation of implicit curves and surfaces. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 245–254, New York, NY, USA, 2004. ACM.

[24] B. T. Stander and J. C. Hart. Guaranteeing the topology of an implicit surface polygonization for interactive modeling. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 279–286, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[25] P. Stelldinger, L. J. Latecki, and M. Siqueira. Topological equivalence between a 3d object and the reconstruction of its digital image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(1):126–140, 2007.

[26] (suppressed). Topology verification for isosurface extraction. Technical report, (suppressed), 2010.

[27] A. van Gelder and J. Wilhelms. Topological considerations in isosurface generation. *ACM Trans. Graph.*, 13(4):337–375, 1994.

[28] G. Varadhan, S. Krishnan, T. Sriram, and D. Manocha. Topology preserving surface extraction using adaptive subdivision. In *SGP '04: Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing*, pages 235–244, New York, NY, USA, 2004. ACM.

[29] G. H. Weber, G. Scheuermann, H. Hagen, and B. Hamann. Exploring scalar fields using critical isovalues. In *VIS '02: Proceedings of the conference on Visualization '02*, pages 171–178, Washington, DC, USA, 2002. IEEE Computer Society.