Discrete Comput Geom (2008) 39: 138–157
DOI 10.1007/s00454-007-9011-4

Isotopic Implicit Surface Meshing

Jean-Daniel Boissonnat · David Cohen-Steiner · Gert Vegter

Received: 7 December 2005 / Revised: 2 March 2007 /
Published online: 13 September 2007

Abstract This paper addresses the problem of piecewise linear approximation of implicit surfaces. We first give a criterion ensuring that the zero-set of a smooth function and the one of a piecewise linear approximation of it are isotopic. Then, we deduce from this criterion an implicit surface meshing algorithm certifying that the output mesh is isotopic to the actual implicit surface. This is the first algorithm achieving this goal in a provably correct way.

Keywords Topology · Triangulations · Morse theory · Algorithms

1 Introduction

Implicit equations are a popular way to encode geometric objects; See, e.g., [4, 25]. Typical examples are CSG models, where objects are defined as results of boolean operations on simple geometric primitives. Given an implicit surface, associated geometric objects of interest, such as contour generators, are also defined by implicit equations. Another advantage of implicit representations is that they allow for efficient blending of surfaces, with obvious applications in CAD or metamorphosis. Finally, this type of representation is also relevant to other scientific fields, such as level set methods or density estimation [8].

However, most graphical algorithms, and especially those implemented in hardware, cannot process implicit surfaces directly, and require that a piecewise linear approximation of the considered surface has been computed beforehand. As a consequence, polygonization of implicit surfaces has been widely studied in the literature.

J.-D. Boissonnat · D. Cohen-Steiner (✉)
Projet Géométrica, INRIA Sophia-Antipolis, Nice, France
e-mail: david.cohen-steiner@sophia.inria.fr

G. Vegter
Institute for Mathematics and Computing Science, RUG, Amsterdam, Netherlands

There are two general classes of methods devoted to this problem: continuation methods and adaptive enumeration methods. A *continuation algorithm* is surface based in the sense that it starts from a seed point on the surface, and computes successive vertices of the mesh while following the surface in some tangent direction. None of the algorithms in this category comes with topological guarantees: they might miss some connected components, or merge different components into a single one. *Adaptive enumeration methods*, also called *extrinsic polygonization methods* [25], are grid based, or, more generally, based on a tessellation of the ambient 3D space. They consist of two steps: first build a tessellation of space, and then analyze the intersection of the considered surface with each cell of the tessellation to construct the approximation. The celebrated marching cube algorithm [16] belongs to this category. The goal of an implicit surface polygonizer is twofold: its output should be geometrically close to the original surface, and have the same topology. While the former is achieved by several polygonization schemes [26], the latter has been barely addressed up to now.

Some algorithms achieve topological consistency, that is, ensure that the result is indeed a manifold, by taking more or less arbitrary decisions when a topologically ambiguous configuration is encountered. This implies that their output might have a topology different from the one of the original surface, except in very specific cases [15]. The problem of topologically correct polygonization of implicit curves in the plane is treated by Snyder in [24], who uses an adaptive enumeration method. His algorithm combines interval arithmetic with a quadtree tessellation of the domain of interest. It seems hard to generalize this method to implicit surfaces in three-space. Moreover, this algorithm seems to have high complexity due to the large number of calls to the interval version of Newton's method.

When the conference version of the present paper was published (Proceedings of STOC'04), there was only one paper devoted to the problem of homeomorphic polygonization of surfaces [19]. Since then there has been several papers [5, 7, 18] that solve the same problem as ours, or a related one. The main theoretical tool used in [19] is Morse theory. The authors first find a level set of the considered function that can be easily polygonized. This initial polygonization is then progressively transformed into the desired one, by computing intermediate level sets. This requires in particular to perform topological changes when critical points are encountered. This algorithm has an intuitive justification and seems to work on simple cases. Unfortunately, the authors do not give any proof of its correctness, and it is not clear to us whether it can deal with complex shapes in a robust way. In particular, the method does not guarantee that the mesh produced are self-intersection free.

In this paper, we give the first certified algorithm for the more difficult problem of *isotopic* implicit surface polygonization. This means that our output can be continuously deformed into the actual implicit surface without introducing self-intersections [14]. For instance, if the original implicit surface is knotted, then our output is guaranteed to be knotted in the same way, which would not be guaranteed by an algorithm ensuring only homeomorphic polygonization. Moreover, the whole algorithm can be implemented in the setting of interval analysis. We only assume that the considered isosurface is smooth, that is, does not contain any critical point. By Sard's theorem [22], this is a generic condition. Our polygonization is the zero-set of the linear interpolation of the implicit function on a mesh of $\mathbb{R}^3$. We first exhibit a set of conditions on the mesh used for interpolation that ensure the topological correctness (Sect. 2). Then, we describe an algorithm for building a mesh satisfying these conditions, thereby leading to a provably correct isotopic polygonization algorithm (Sect. 3).

We note that since the publication of the conference version of the present paper, another method appeared that solves exactly the same problem as ours [18]. One difference between the two methods is that [18] uses octrees instead of triangulations. A more important difference is in the refinement stopping criterion: in [18], cells are subdivided until the intersection of the implicit surface with each cell is sufficiently flat. By contrast, we stop refinement as soon as a certain global criterion ensuring topological correctness is met. Hence, we may expect that our method is faster than [18]. This remains to be proved though, since we did not implement our method.

2 A Condition for Isotopic Meshing

Let f be a C^2 function from $\mathbb{R}^3$ to $\mathbb{R}$, and M be its zero-set. We assume that M, the surface we want to polygonize, is compact (condition a1). In what follows, T denotes a triangulation of a domain $\Omega \subset \mathbb{R}^3$ containing M and $\hat{f}$ the function that coincides with f at the vertices of T and that is linearly interpolated on the simplices of T. A vertex v will be said to be *larger* (resp. *smaller*) than a vertex u if $f(v)$ is *larger* (resp. *smaller*) than $f(u)$; the sign of f at a vertex will be referred to as the sign of that vertex. We set $\hat{M} = \hat{f}^{-1}(0)$.

2.1 Topological Background

Collapses Loosely speaking, a collapse [20] is an operation which consists of removing cells from a simplicial complex without changing its connectivity. More precisely:

Definition 1 If L is a simplicial complex and K a subcomplex of L, one says that there is an elementary collapse from L to K if there is a p-simplex s of L and a $(p-1)$-face t of s such that:

- s is not a face of any simplex of L.
- t is not a face of any simplex of L other than s.
- L is the union of K, s, and all the faces of s.
- $\partial s \setminus K$ is the relative interior of t.

Definition 2 If L is a simplicial complex and K a subset of L, one says that L collapses to K if there is a subdivision L' of L such that a subdivision of K can be obtained from L' by a sequence of elementary collapses.

Definition 2 is illustrated in Fig. 2. In Fig. 2, the complexes in the middle and on the right do not collapse to the bold curve because they would need to be "torn" in order to do so.

Fig. 1 Elementary collapse

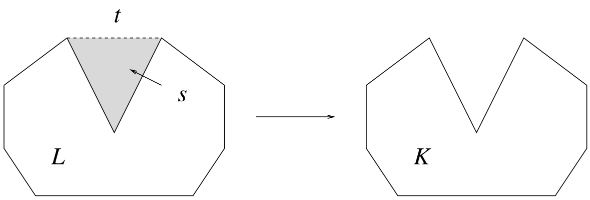

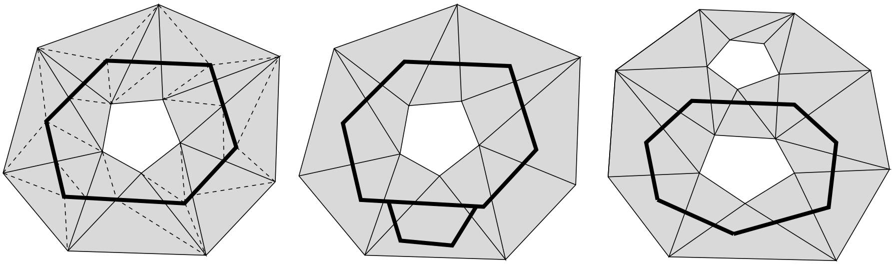

Fig. 2 The *grey complex L* on the *left collapses* to the *bold curve K* (*dashed edges* represent the subdivision L'). This is not true for the two other complexes

Smooth Morse theory The topology of implicit surfaces is usually investigated through Morse theory [17]. Given a real function f defined on a manifold, Morse theory studies the topological changes in the sets $f^{-1}(]-\infty, a])$ (*lower level-sets*) when a varies. In our case, as f is defined on $\mathbb{R}^3$, this amounts to studying how the topology of the part of the graph of f lying below a horizontal hyperplane changes as this hyperplane sweeps $\mathbb{R}^4$. Classical Morse theory assumes that f is of class C^2. In this case, as is well known, these topological changes are related to the *critical points* of f, that is, the points where the gradient ∇f of f vanishes. More precisely, the only topological changes occur when $f^{-1}(a)$ passes through a critical point p. The value a is then called a *critical value*. Generically, in the 2-dimensional case, the topology of $f^{-1}(]-\infty, a])$ can change in three possible ways, according to the type of the critical point p (see Fig. 3).

In Fig. 3, the sets $f^{-1}(]-\infty, a])$ are displayed as light grey regions. The leftmost column depicts the situation where p is a local maximum, that is, when the Hessian of f at p is positive. In this case, $f^{-1}(]-\infty, a+\varepsilon])$ is obtained from $f^{-1}(]-\infty, a-\varepsilon])$ by gluing a topological disk along its boundary. In the case of a saddle point (i.e. the Hessian has critical values of both signs), passing a critical value amounts to gluing a thickened topological line segment (in grey) along its "thickened" boundary (in bold). Finally, passing through a local minimum (negative Hessian) just amounts to adding a disk disconnected from $f^{-1}(]-\infty, a-\varepsilon])$. If p does not fall in any of these categories, that is, if the Hessian at p is degenerate, then classical Morse theory cannot be applied. C^2 functions the critical points of which all have non-degenerate Hessian are called *Morse functions*. From now on, we will assume that f is a Morse function (condition a2). Also, we require that 0 is not a critical value of f (condition a3), which implies that M is a manifold.

The number n of negative eigenvalues of the Hessian at p is classically called the index of p. However, for consistency reasons that will appear later, we call the *index*

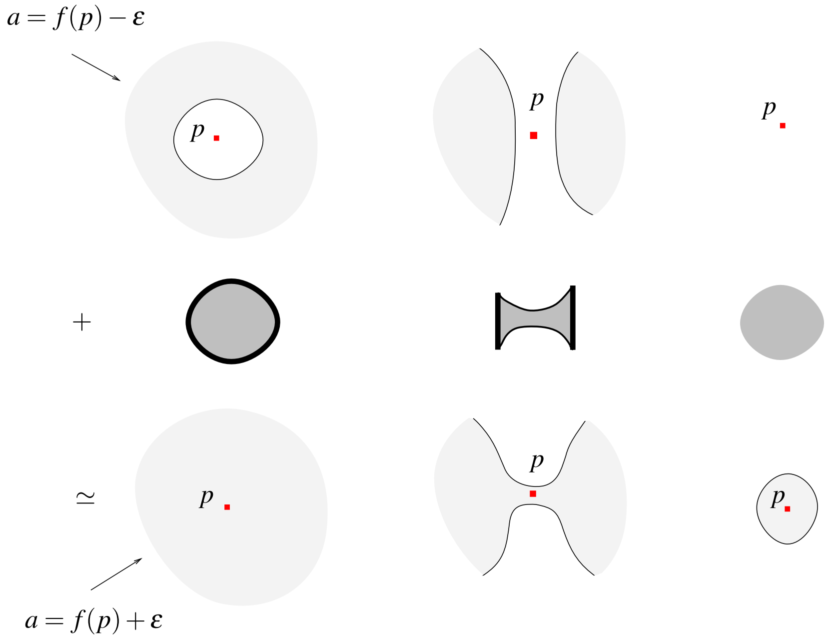

Fig. 3 Smooth Morse theory in 2D

of p the integer $(-1)^n$. The *index* of f on a region V is the sum of the indices of all critical points of f lying in V. The index satisfies the following important theorem:

Theorem 1 (Poincaré-Hopf index theorem) *The index of f on one of its lower level-sets is the Euler characteristic of that lower level-set.*

PL Morse theory Morse theory has been extended to a broad class of non-smooth functions by Goresky and McPherson [11]. We now outline the special case of PL functions, that is, we consider the case of $\hat{f}$. We assume from now on that no two neighboring vertices map to the same value under f, and that no vertex of T maps to 0 under f (conditions b1 and b2), which guarantees that $\hat{M}$ is a manifold. We refer to these assumptions as *genericity assumptions*. Let us first recall some well-known definitions [10, 11]:

Definition 1 The *star* of a vertex is the union of all simplices[1] containing this vertex. The *link* of a vertex is the boundary of its star.

Definition 2 The *lower star* $St^-(v)$ of $\hat{f}$ at a vertex v is the union of all simplices incident on v whose vertices other than v are smaller than v. The *lower link* $\mathrm{Lk}^-(v)$ of $\hat{f}$ at a vertex v is the union of all simplices of the link of v all vertices of which are smaller than v. The upper star $\mathrm{St}^+(v)$ and the upper link $\mathrm{Lk}^+(v)$ are defined similarly.

[1] By simplex we mean a closed cell of T of any dimension.

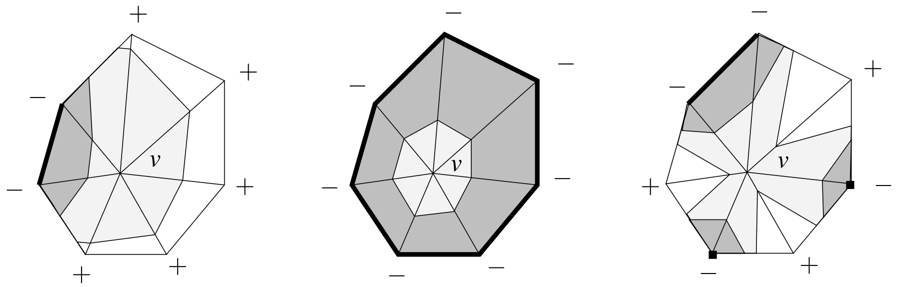

Fig. 4 Morse theory for PL functions in 2D. *Plus* and *minus signs* indicate whether neighbors of v are larger or smaller than v. Lower links are displayed in *bold*, sets $\hat{f}^{-1}(]-\infty, f(v)-\varepsilon])$ in *grey*, and sets $\hat{f}^{-1}([f(v)-\varepsilon, f(v)+\varepsilon])$ in *light grey*

Figure 4 shows that—for small ε—the topological changes between lower level-sets $\hat{f}^{-1}(]-\infty, f(v)-\varepsilon])$ and $\hat{f}^{-1}(]-\infty, f(v)+\varepsilon])$ are determined by the topology of $\mathrm{Lk}^-(v)$. In particular, in 2D, topological changes occur whenever $\mathrm{Lk}^-(v)$ is not connected or equals the link of v (right and middle cases in Fig. 4). This is what motivates the next definition in the higher dimensional case:

Definition 3 A *critical point* of $\hat{f}$ is a vertex whose lower link is not collapsible.[2] A vertex that is not a critical point of $\hat{f}$ will be called *regular*.

With this definition, topological changes in lower level-sets occur exactly at critical points, which is consistent with smooth Morse theory. The *index* of a vertex v is defined to be 1 minus the Euler characteristic of $\mathrm{Lk}^-(v)$ [2]. In particular, regular points all have index 0. The converse is not true however in dimension at least 3. Also, checking if a vertex is regular is easy for PL functions defined on three-dimensional meshes: it is sufficient to check that the lower link and the upper link are both non-empty and connected.[3] Define the index of $\hat{f}$ on a region V to be the sum of the indices of all critical points of $\hat{f}$ lying in V. Again, this definition is consistent with the smooth case, since the PL index can be shown to also satisfy the Poincaré-Hopf index theorem [2]. The following lemma will be used later:

Lemma 2 *If the gradients of $\hat{f}$ on tetrahedra incident to a vertex v all have a positive inner product with some vector, then v is regular.*

Proof By Proposition 1.2 page 450 in [1], $\hat{f}^{-1}(]-\infty, f(v)+\varepsilon])$ retracts by deformation on $\hat{f}^{-1}(]-\infty, f(v)-\varepsilon])$ for sufficiently small ε. Hence $\mathrm{Lk}^-(v)$ has the homology groups of a point, implying that it is collapsible since it is a subcomplex of the 2-sphere. □

[2] A complex is collapsible if it collapses to a point.

[3] This follows from Alexander duality together with the fact that contractible subcomplexes of the 2-sphere are collapsible.

2.2 Main Result

We assume throughout the paper that f and T satisfy conditions a1, a2, a3, b1, b2. That is, M is compact, f is a Morse function, 0 is not a critical value of f, no vertex of T map to 0 by f, and no two neighboring vertices of T map to the same value by f. Additionally, we assume that the following condition holds:

0. f does not vanish on any tetrahedron of T containing a critical point of f.

Theorem 3 *Let W be a subcomplex of T satisfying the following conditions*:

1. *f does not vanish on ∂W.*
2. *W contains no critical point of f.*
2′. *W contains no critical point of $\hat{f}$.*
3. *W collapses to $\hat{M}$.*
4. *f and $\hat{f}$ have the same index on each bounded component of $\Omega \setminus W$.*

Then M and $\hat{M}$ are isotopic in W. Moreover, the Hausdorff distance between M and $\hat{M}$ is smaller than the "width" of W, that is, the maximum over the components V of W of the Hausdorff distance between the subset of ∂V where f is positive and the one where f is negative.

Here, isotopic in W means that M can be continuously deformed into $\hat{M}$ while remaining a manifold embedded in W, so that M could not be a knotted torus if $\hat{M}$ is an unknotted one, for instance. We first prove that under the conditions of the theorem, M and $\hat{M}$ are homeomorphic. Under the assumptions of the theorem, the fact that they actually are isotopic will be a direct consequence of a result obtained in [6]. Before proving the theorem, we first show by some examples that none of its assumptions can be removed. In the three following pictures, (local) minima of f are represented by min, (local) maxima by max, and saddle points by s. Critical points of $\hat{f}$ are represented similarly but with a caret. The sign preceding a critical point symbol indicates the sign of the considered function (f or $\hat{f}$) at the critical point.

Figure 5 shows that condition 0 cannot be removed even in the 2D case. By allowing for critical points of f inside a triangle of T with positive vertices, one can build an example where M has an extra component with respect to $\hat{M}$ without violating

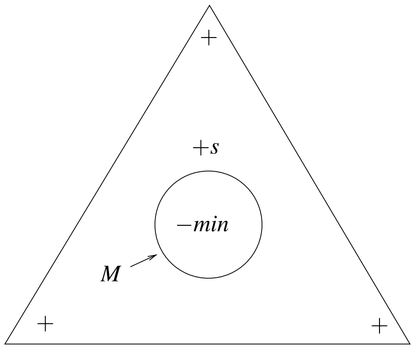

Fig. 5 Condition 0 is necessary

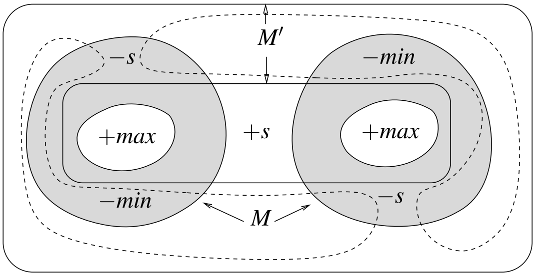

Fig. 6 Critical points do not determine the topology of level-sets

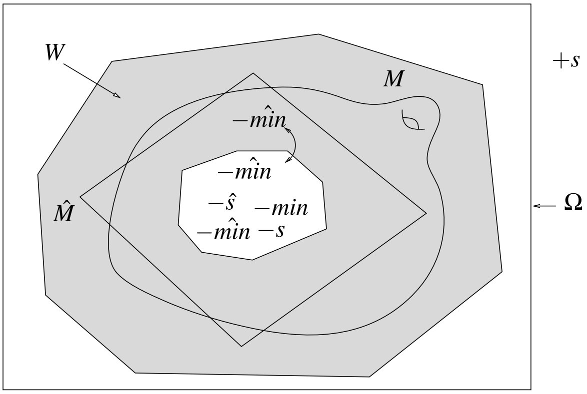

Fig. 7 Condition 2′ and 4 are necessary

conditions involving critical points and their indices. Indeed, in Fig. 5, f has index 0 on the triangle, since minima have index 1 and saddle points have index -1.

The situation in Fig. 6 is a 2D example of two zero-sets M (boundary of the grey region) and M' which are not homeomorphic, though their defining functions have the same critical points, with the same indices. The dashed curve represents a negative level-set of the function defining M'. Such an example can also be built such that $M' = \hat{M}$ for some mesh T. This shows the importance of the set W in the theorem. In particular, conditions 1 and 3 cannot be removed. Indeed, if one drops 1, taking for W any set satisfying 2 and 3 makes the theorem fail. On the other hand, if one drops 3, any W satisfying 2 and 1 also makes the theorem fail.

Figure 7 shows a 3D example where M is a torus whereas $\hat{M}$ is a sphere. This is because $\hat{f}$ has an extra negative minimum inside $\hat{f}^{-1}(]-\infty, 0])$ whereas f has an index 1 saddle point outside the bounding box Ω. Depending on whether this extra minimum lies in W or not (see the circle arc with arrows at both ends in Fig. 7), one obtains counterexamples to the theorem if assumptions $2'$ or 4 are dropped. One can build similar examples showing that condition 2 is also needed.

We now return to the proof of Theorem 3.

2.3 Proof of the Theorem

Lemma 4 *Let S and T be two subsets of a topological space X that meet* (*i.e.* $S \cap T \neq \emptyset$). *Assume the boundary of S, as well as T and $X \setminus T$, are connected. If*

$X \setminus S$ and $X \setminus T$ meet but their boundaries do not, then S is contained in the interior of T or the other way around.

Proof The boundary of S is the disjoint union of $\partial S \cap \mathrm{int}(T)$ and $\partial S \cap \mathrm{int}(X \setminus T)$ since $\partial S \cap \partial T$ is empty. So we have a partition of ∂S in two relatively open sets. As it is connected, one has to be empty. If $\partial S \cap \mathrm{int}(T)$ is empty then $\partial S \subset \mathrm{int}(X \setminus T)$ that is, $T \cap \partial S$ is empty. As a consequence, T is included in $\mathrm{int}(S)$ or in $\mathrm{int}(X \setminus S)$ by connectedness. Since S and T meet, we have that $T \subset \mathrm{int}(S)$.
Now if $\partial S \cap \mathrm{int}(X \setminus T)$ is empty then $X \setminus T$ is contained in $\mathrm{int}(S)$ or in $\mathrm{int}(X \setminus S)$ by connectedness again. Similarly as above it has to be contained in $\mathrm{int}(X \setminus S)$, which implies that $S \subset T$. Thus $\mathrm{int}(S) \subset \mathrm{int}(T)$ so $\partial S \supset S \setminus \mathrm{int}(T) = S \cap \partial T$. If S would meet ∂T, then ∂S and ∂T would meet, which is impossible. Hence, S is included in the interior of T. □

Lemma 5 *Let V be a connected component of W. $M \cap V$ is a connected smooth compact manifold without boundary.*

Proof Condition 3 implies easily that V collapses to $\hat{M} \cap V$. Therefore V contains a simplex having positive and negative vertices. As a consequence, f vanishes on V. Since f does not vanish on ∂W (condition 1), M intersects V. Also, M does not meet the boundary of V (condition 1), so $M \cap V$ is a smooth compact manifold without boundary.

Because V, which is connected, collapses to $\hat{M} \cap V$, $\hat{M} \cap V$ is a connected closed surface. Therefore, the complement of $\hat{M} \cap V$ has exactly two components, one of which is bounded. Because V collapses to $\hat{M} \cap V$, $\mathbb{R}^3 \setminus V$ also has exactly one bounded component which we denote by A and one unbounded component we denote by B (see Fig. 8). The complement of A, which is $B \cup V$, is connected, because B and V are connected. For the same reason, $A \cup V$ is also connected. Moreover, the complement of $A \cup V$, being equal to B, is also connected. In summary, A is connected as well as its complement, and the same is true for $A \cup V$.

Call now M_i, $i = 1, \ldots, n$ the connected components of $M \cap V$ (see Fig. 8). For each i, let N_i be the bounded component of $\mathbb{R}^3 \setminus M_i$. $M_i = \partial N_i$ does not meet $\partial(A \cup V) \subset \partial W$ (1), and $A \cup V$ is connected as is its complement. So N_i is included in $A \cup V$ thanks to Lemma 4. Now N_i contains at least one critical point of f. But as $N_i \subset A \cup V$, such a point has to lie in A, by 2. So N_i meets A, but since $\partial N_i = M_i$ does not meet $\partial A \subset \bar{W}$, N_i contains A by Lemma 4 again. Suppose $M \cap V$ is not connected. Then N_1 and N_2 both contain A so they intersect. Because M is smooth, their boundaries do not intersect. So one has w.l.o.g. $N_2 \subset N_1$. Now f vanishes on $\partial(N_1 \setminus N_2) = \partial N_1 \cup \partial N_2$, and therefore has an extremum in $N_1 \setminus N_2$, which is impossible by 2 because $N_1 \setminus N_2 \subset V$. □

So $M \cap V$ and $\hat{M} \cap V$ are connected compact surfaces without boundary. As seen in the preceding proof, A contains all critical points of f enclosed by $M \cap V$. Also, A contains all critical points of $\hat{f}$ enclosed by $\hat{M} \cap V$ by $2'$. From condition 4, we deduce that the volumes enclosed by $M \cap V$ and by $\hat{M} \cap V$ have the same Euler characteristic, since the Euler characteristic of a lower level set is the index of the considered function on that lower level set (Theorem 1). So $M \cap V$ and $\hat{M} \cap V$ have

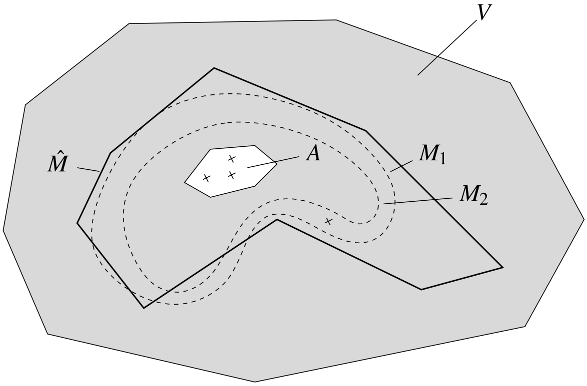

Fig. 8 Proof of Lemma 5

the same genus and are thus homeomorphic. To complete the proof that M and $\hat{M}$ are homeomorphic, it remains to check that:

Lemma 6 *M is included in W.*

Proof Let D be some component of $\Omega \setminus W$. We claim that $M \cap D$ is empty. First $\hat{M} \cap D$ is empty by condition 1 so w.l.o.g. vertices lying in the closure of D are all positive. If $M \cap D$ is not empty then some component E of $f^{-1}(]-\infty, 0])$ meets D. Moreover, ∂D does not meet E. Indeed, f is positive at vertices of ∂D, and does not vanish on $\partial D \subset \partial W \cup \partial\Omega$ by condition 1. So E, being connected, is included in the interior of D. But then E is compact and thus f reaches its minimum on E, implying that E contains a (negative) critical point of f. This is impossible since the tetrahedron containing this critical point would have negative vertices by condition 0, though being included in D. □

The proof of the bound on the Hausdorff distance between M and $\hat{M}$ is not difficult. Pick any point p in $\hat{M}$ and let V be the component of W containing it. Assume w.l.o.g. that $f(p) > 0$ and let p' be the closest point of p on the component of ∂V where f is negative. By the intermediate value theorem, the line segment pp' meets M at a point q. The distance between p and q is smaller than the distance between p and p' which is smaller than the Hausdorff distance between the two components of ∂V. This shows one part of the bound. The other part can be proved in a similar way.

Now that we know that M and $\hat{M}$ are homeomorphic, the fact that they are isotopic is a consequence of Proposition 7, which is proved in [6].

Proposition 7 *Let $\hat{M}$ be an orientable compact surface without boundary and let M be a surface such that*

- *$\hat{M}$ is homeomorphic to M,*
- *M separates the sides of a topological thickening[4] $\tilde{W}$ of $\hat{M}$.*

Then M is isotopic to $\hat{M}$ in $\tilde{W}$.

[4]This means that there is a homeomorphism $\Phi : \tilde{W} \to \hat{M} \times [0, 1]$ mapping $\hat{M}$ to $\hat{M} \times \{1/2\}$.

Indeed, considering a regular neighborhood of W [20] yields the desired topological thickening $\tilde{W}$, as can be seen from the uniqueness theorem for regular neighborhoods from piecewise-linear topology [20].

3 Algorithm

In the algorithm, we take as W a set that is related to the notion of watershed from topography. This set satisfies properties $2'$ and 3 by construction. In Sect. 3.1, we give its definition, basic properties, and construction algorithms. Section 3.2 describes the meshing algorithm itself, which ensures that W fulfills also conditions 0, 1, 2, and 4, and proves its correctness.

3.1 PL Watersheds

We first assume that the mesh T conforms to $\hat{M}$, i.e. $\hat{M}$ is contained in a union of triangles of T. We will see later how to remove this assumption, which is in contradiction with the genericity assumptions. Define W^+ as the result of the following procedure:

Positive Watershed Algorithm

```
set W+ = M̂.
mark all vertices of M̂.
while there is a positive regular unmarked vertex v of T
   such that the vertices of Lk−(v) are marked
do
   set W+ = W+ ∪ St−(v).
   mark v.
end while
return W+
```

W^- is defined as the result of the same algorithm applied to $-f$. We set $W = W^+ \cup W^-$. Note that W contains no critical point of $\hat{f}$. Also, positive marked vertices are exactly the vertices of W^+.

Lemma 8 *W collapses to $\hat{M}$.*

Proof It is sufficient to show the result for W^+. Let W_i^+ be the state of W^+ after i steps of the algorithm, and let v_i be the i-th marked vertex. As $W_0^+ = \hat{M}$, the only thing we have to show is that W_{i+1}^+ collapses to W_i^+ for all i. Let us first show that $\mathrm{Lk}^-(v_i)$ is included in W_i^+. If it is not the case, let u be the largest vertex of some simplex s of $\mathrm{Lk}^-(v_i)$ that is not in W_i^+. Simplex s is in $\mathrm{St}^-(u)$ which is therefore not included in W_i^+. This is a contradiction since v_i is marked. Therefore $\mathrm{Lk}^-(v_i) \subset W_i^+$. Now since v_i is regular, $\mathrm{Lk}^-(v_i)$ is collapsible. Consider a sequence of elementary collapses allowing to collapse $\mathrm{Lk}^-(v_i)$ to p and let $s_j \subset \mathrm{Lk}^-(v_i)$, $j = 1, \ldots, n$

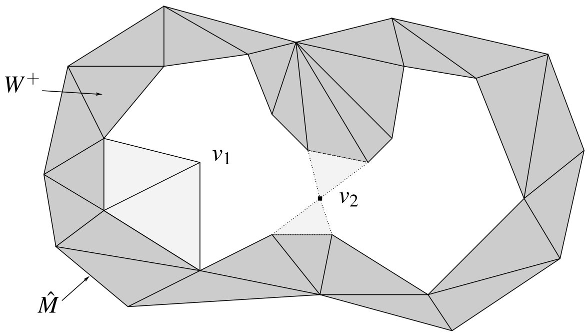

Fig. 9 Construction of W^+: lower stars of regular vertices (such as v_1) are added one by one. Lower stars of critical vertices (v_2) are discarded

be the sequence of simplices defining these elementary collapses. The simplices $\operatorname{conv}(s_j \cup v_i)$, $j = 1, \ldots, n$ and the edge pv_i define a valid sequence of elementary collapses allowing to collapse $W^+_{i+1} = W^+_i \cup \mathrm{St}^-(v_i)$ to W^+_i, which concludes the proof. □

One may prefer a more intrinsic definition of W^+. In the same spirit as in [9], one can define a partial order on the vertices of T by the closure of the acyclic relation $\prec$ defined by $u \prec v$ if $u \in \mathrm{Lk}^-(v)$ or $u = v$. We will denote this order $\prec$ again and say that v *flows* into u whenever $u \prec v$. The next lemma shows that the vertices of W^+ do not depend on the order in which the vertices are considered in the construction.

Lemma 9 *The vertices of W^+ are exactly the positive vertices that do not flow into any positive critical point of $\hat{f}$.*

Proof The vertices of W^+ have this property by construction. Let p be a positive vertex not belonging to W^+ and assume p does not flow into any positive critical point. In particular, p is regular by reflexivity. Hence, as $p \notin W^+$, the lower link of p, which is not empty, has to contain an unmarked vertex. It cannot contain a critical point because as T conforms to $\hat{M}$, vertices in $\mathrm{Lk}^-(p)$ are all non-negative, and so p would flow into a positive critical point. There is thus an unmarked vertex in $\mathrm{Lk}^-(p)$. If we can choose an unmarked positive vertex p_1 in $\mathrm{Lk}^-(p)$, then p_1 does not belong to W^+, and flows into a positive critical point. Repeating this process with p replaced by p_1, we find a strictly decreasing sequence of positive vertices, that thus has to end. Let p_k be its last term. The lower link $\mathrm{Lk}^-(p_k)$ contains no positive unmarked vertices. But as T conforms to $\hat{M}$, vertices in $\mathrm{Lk}^-(p_k)$ are all non-negative. Since vertices of $\hat{M}$ are marked, we get a contradiction. □

Note that W is the union of simplices with all their vertices in W. As a result, we get an intrinsic definition of W, and not only of its vertices. From an algorithmic point of view, it may be efficient to examine the vertices in increasing order in the construction of W^+. One can for instance maintain the ordered list of vertices neighboring W, always consider the first element of this list for marking, and discard it if it cannot be marked. Indeed, with this strategy, a vertex that cannot be marked at some point will never be marked.

Another consequence of Lemma 9, which will be useful later, goes as follows. Let c be the minimum of $|\hat{f}(v)|$, and hence the minimum of $|f(v)|$ over all critical points v of $\hat{f}$.

Lemma 10 *W contains all vertices the image of which under $|f|$ is smaller than c.*

Proof Let p be such that $|f(p)| < c$. Without loss of generality, assume that p is positive. Any critical point v into which p flows satisfies $f(v) < f(p)$. So it cannot be positive by definition of c: by Lemma 9, p lies in W^+. □

Non conforming case We now drop the assumption that T conforms to $\hat{M}$ and assume genericity again. From T and $\hat{M}$ one can build a mesh S that is finer than T, conforms to $\hat{M}$, and has all its extra vertices on $\hat{M}$. Indeed, it suffices to triangulate the overlay of $\hat{M}$ and T without adding extra vertices except those of $\hat{M} \cap T$. This can be done as the cells of the overlay are convex. The construction of W described above can then be applied to S. A positive vertex of T has its lower link in S containing only vertices of $\hat{M}$ if and only if its lower link in T contains only negative vertices. Thus, in order to find the positive vertices of $W \cap T$, one can apply the positive watershed algorithm described above to T, if at the initialization step one marks all negative vertices having a positive neighbor instead of those of $\hat{M}$. Still, note that if a negative critical point has a positive neighbor, then this neighbor will not be marked by this modified algorithm, whereas it could have been marked by the standard algorithm applied to S. However, if we assume that vertices having a neighbor of opposite sign are regular (condition c), then this does not happen and the result W' of the modified algorithm is equal to W. The negative vertices of $W \cap T$ are determined similarly. In our meshing algorithm, we will not build the mesh S, but rather make sure condition c holds, and apply the modified algorithm.

Updating W' The intrinsic definition of W—or W'—given above yields an efficient way of updating W when T undergoes local transformations. It is sufficient to describe the algorithm for updating the vertices of W^+. Let T_1 be a mesh obtained from T by removing some set of tetrahedra E and remeshing the void left by E. Call A the set of positive critical points of the linear interpolation of f on T_1 that lie in E. Then the vertex set of the positive watershed W_1^+ associated with T_1 can be computed from the vertex set of W^+ by performing the following two operations. To begin with, the set of vertices of T_1 that flow into A must be removed from W^+ (Lemma 9), which amounts to a graph traversal. The remaining vertices of T_1 all belong to W_1^+. Then, mark these vertices and apply the positive watershed algorithm loop to get the other vertices of W_1^+.

Remark The watershed we compute is in general strictly included in the 'true watershed'. The 'true watershed' seems hard to compute, though, and can intersect a triangle in a very complicated way. There might be interesting intermediate definitions between ours and the true one, for instance based on the PL analog of the Morse complex introduced in [10].

3.2 Main Algorithm

Theorem 3 enables us to build a mesh isotopic to M using two simple predicates, *vanish* and *vanish'*. The predicate *vanish* (resp. *vanish'*) takes a triangle or a box and return true if f (resp. ∇f) vanishes on that triangle or that box. We actually do not even need predicates, but rather filters. More precisely, *vanish* (or *vanish'*) may return true even if f does not vanish on the considered element, but not the other way around. Still, we require that *vanish* returns the correct answer if the input triangle or box is sufficiently small. Such filters can be designed using interval analysis.

Our algorithm also requires to build a refinable triangulation of space such that $\hat{f}$ (resp. $\nabla \hat{f}$) converges to f (resp. ∇f) when the size of the elements tends to 0. As noticed by Shewchuk [23], this is guaranteed provided all tetrahedra have dihedral and planar angles bounded away from π. In [3], Bern, Eppstein and Gilbert described an octree-based algorithm yielding meshes the angles of which are bounded away from 0. In our case, which is much easier, the desired triangulation can simply be obtained by adding a vertex at the center of each square and each cube of the octree, triangulating the squares radially from their center, and doing the same with the cubes. Indeed, resulting planar and dihedral angles are all bounded away from π. One can expect that this scheme does not produce too many elements upon refinement, because the size of elements is allowed to change rapidly as we do not require that these have a bounded aspect ratio (see Fig. 10). The main algorithm uses an octree O, the associated triangulation T, and the watershed W'. We will say that two (closed) boxes of O are neighbors if they intersect. O is initialized to a bounding box Ω of M. Such a bounding box can be found by computing the critical points of the coordinate functions restricted to M, if possible, or by using interval analysis. Besides, we maintain five sets of boxes ordered by decreasing size. *Critical1* is a certain set of boxes obtained by interval analysis (see below). This set has the property that the union of its boxes, which we call the *critical set*, encloses all critical points of f but does not intersect M. *Critical2* contains all boxes containing a critical point of $\hat{f}$ that is not in a box belonging to *Critical1*. *Index* contains all boxes neighboring a box b in *Critical1* such that f and $\hat{f}$ have different indices on the connected component of the critical set that contains b. We defer the description of a method that computes the

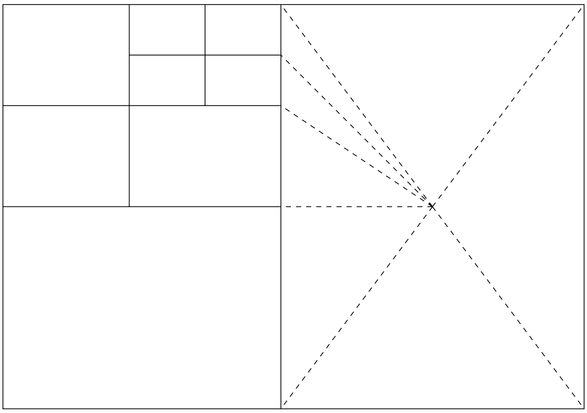

Fig. 10 Octree and triangulation used in the algorithm. In this 2D example, only the edges of the triangulation of the box on the right are shown (*dashed*)

index of f on a box in a certified way to the appendix. *Boundary1* contains all boxes containing two neighboring vertices of opposite signs one of which is critical for $\hat{f}$ (condition c, see paragraph **Non conforming case**). *Boundary2* contains all boxes that are not included in W', and that contain a triangle t of $\partial W'$ such that *vanish*(t) is true. Finally, for our algorithm to work, we need to introduce a slight modification of the watershed W', which we call W''. The modification consists of taking as W''^{+} vertices—and the same for W''^{-}—the positive vertices that do not flow into positive critical points of $\hat{f}$ nor into vertices lying in a box containing a positive critical point of f. With this modification, Lemma 8 still holds and Lemma 10 holds if one replaces c by the minimum c' of c and the minimum of $|f|$ on the boxes containing a critical point of f. Also, c' is positive as f does not vanish on these boxes.

Main Algorithm

Initialization Refine O until all boxes b satisfy either *vanish*(b) is false or *vanish'*(b) is false. Insert all boxes b such that *vanish'*(b) is true in *Critical1*. compute T and W'', and the four sets *Critical2*, *Boundary1*, *Boundary2*, and *Index*.
while (**true**) **do**
  update T, W'', and the four sets.
  if *Critical2* $\neq \emptyset$ **then**
    split its first element.
  else if *Boundary1* $\neq \emptyset$ **then**
    split its first element.
  else if *Boundary2* $\neq \emptyset$ **then**
    split its first element.
  else if f and $\hat{f}$ have different indices on some component of the critical set **then**
    split the first element of *Index*.
  else
    return $\hat{M}$
  end if
end while

Thanks to Theorem 3 applied to W'', the correctness of this algorithm amounts to its termination. We now show that the main algorithm terminates. First note that after the initialization step, no box containing a critical point of f is split, because such boxes belong to *Critical1*. The magnitude of ∇f is thus larger than a certain constant $g_{\min}$ on the complement C of the union of these boxes. Let us show that the size of the boxes of *Critical2* that are split at some point is bounded from below. As $\nabla\hat{f}$ converges to ∇f, there is a number s_1 such that for each tetrahedron with diameter smaller than s_1, $\|\nabla f - \nabla\hat{f}\|$ is smaller than $g_{\min}/2$ on the interior of that tetrahedron. If the tetrahedron is included in C, $\|\nabla f\| > g_{\min}$, which implies that $\nabla\hat{f}$ and ∇f make an angle smaller than $\pi/6$.

Lemma 11 *Let $A \subset \mathbb{R}^3$ be such that ∂A is a manifold included in C and containing no vertex of T. Suppose that all boxes meeting ∂A are smaller than s_1. Then f and $\hat{f}$ have the same index on A.*

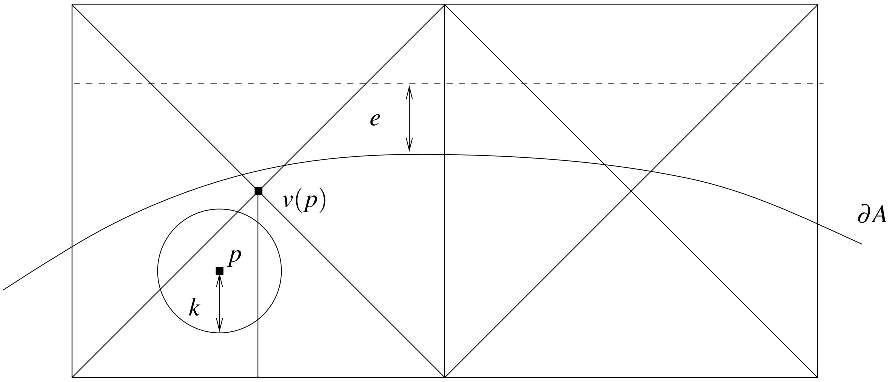

Fig. 11 Proof of Lemma 11

The proof of Lemma 11 resorts to stratified Morse theory, which is an extension of both the smooth and PL Morse theory to the case of piecewise smooth functions. We refer to [11] for a complete exposition of this subject.

Proof For $p \in \partial A$, let $d(p)$ denote the largest number such that the simplices of T that meet the open ball centered at p of radius $d(p)$ all share a vertex, $v(p)$. The quantity $d(p)$ is the 3-dimensional analog of the local feature size function introduced by Ruppert [21]. We call $d_{\min}$ the minimum of d, which is known to be positive, and set k equal to the minimum of $d_{\min}$ and e, where e is half the distance from ∂A to the closest box that does not meet ∂A.

Let us now consider a smooth nonnegative function $\phi : \mathbb{R}^3 \to \mathbb{R}$ with support included in the open ball centered at 0 of radius k. The convolution of $\hat{f}$ and ϕ is a smooth function $\tilde{f}$. Let p be a point at distance less than e from ∂A. The gradient of $\tilde{f}$ at p is a weighted average of the gradients of $\hat{f}$ at points lying in the open ball centered at p and with radius k. All gradients involved in this average are gradients of $\hat{f}$ on tetrahedra incident on $v(p)$. Moreover, the size of these tetrahedra is smaller than s_1 because $k \leq e$. As a consequence, all gradients considered make an angle smaller than $\pi/6$ with the gradient of f at $v(p)$. As the weights in the average are nonnegative, we have that the angle between $\nabla\tilde{f}(p)$ and $\nabla f(v(p))$ is smaller than $\pi/6$. Also, the angle between $\nabla f(v(p))$ and $\nabla f(p)$ is less than $\pi/3$ since both vectors make an angle smaller than $\pi/6$ with the gradient of $\hat{f}$ on some tetrahedron containing p and $v(p)$. Finally, we get that $\nabla\tilde{f}(p)$ and $\nabla f(p)$ have a positive inner product.

Let now U_1 be a neighborhood of ∂A whose closure does not contain any vertex of T and let U_2 be an open set such that $U_1 \cup U_2 = \mathbb{R}^3$. We also require that the Hausdorff distance between U_1 and ∂A is smaller than e and that $U_2 \cap \partial A = \emptyset$. Denote by $\{u_1, u_2\}$ a partition of unity subordinate to the covering $\{U_1, U_2\}$. This means that u_1 and u_2 are nonnegative smooth function defined on $\mathbb{R}^3$, with support in U_1 and U_2 respectively, and such that $u_1 + u_2 = 1$. In particular, u_2 equals 1 on the complement of U_1, and u_1 equals 1 on the complement of U_2. So the function $g = u_2\hat{f} + u_1\tilde{f}$ coincide with $\hat{f}$ on $\mathbb{R}^3 \setminus U_1$ and with $\tilde{f}$ on $\mathbb{R}^3 \setminus U_2$. Now recall that $\nabla\tilde{f}$ and ∇f have a positive inner product on ∂A, which is contained in the complement of U_2. Hence the linear homotopy between both vector fields does not vanish on ∂A: by normalization, one gets a homotopy between $\nabla\tilde{f}/\|\nabla\tilde{f}\|$ and $\nabla f/\|\nabla f\|$, considered as maps from ∂A to the unit sphere. Because the degree (see [13] p. 134 for a definition) is invariant under homotopy, we deduce that these maps have the same degree, which shows that f and $\tilde{f}$ have the same index on A. Now as g and $\tilde{f}$ coincide in a neighborhood of ∂A, f and g have the same index on A. To complete the proof, it thus suffices to show that g and $\hat{f}$ also have the same index on A. Now the critical points of $\hat{f}$ are critical for g, with the same index, as U_1 contains no such point. Potential other critical points of g can only lie in U_1. But the gradient of g at any point p of U_1 where it is defined is a convex combination of $\nabla\tilde{f}(p)$ and $\nabla\hat{f}(p)$: it thus has a positive inner product with $\nabla f(p)$. By the result of [1] which we mentioned when we stated Lemma 2, this implies that the index of p is 0. We thus proved the announced claim. □

Suppose that some box b of *Critical2* of size smaller than s_1 is split. Let v be a critical point of $\hat{f}$ included in b. All the boxes containing v are in *Critical2* and their size is smaller than s_1 since we consider boxes in decreasing order. Now the gradients of $\hat{f}$ on tetrahedra incident on v all have a positive inner product with $\nabla f(v)$ (recall ∇f and $\nabla\hat{f}$ make an angle less than $\pi/6$), which is a contradicts Lemma 2, implying that v is not critical. So the conclusion is that *Critical2* becomes—at least temporarily—empty after a finite number of consecutive splittings of boxes in *Critical2*.

Now if the algorithm splits a box b in *Boundary1*, then b contains a critical point of $\hat{f}$. This critical point, which we assume to be positive, belongs to a box containing a critical point of f as *Critical2* is empty. So the maximum of $|f|$ on b is larger than the minimum of $|f|$ on the boxes containing a critical point of f (i.e. c'). On the other hand, f vanishes on b since b contains a negative vertex. This cannot happen if the size of b is below a certain value, so that boxes in *Boundary1* cannot be split indefinitely.

Suppose that the algorithm splits arbitrarily small boxes in *Boundary2*. If a small enough box b is split, then b contains a triangle t of W'' on which f vanishes. So, if the size of b is small enough, the maximum of $|f|$ on b will be smaller than c'. By Lemma 10, all vertices of b will then belong to W'' so $b \subset W''$ which is a contradiction. Thus the size of split boxes in *Boundary2* is also bounded from below.

To complete the proof of termination, we need to prove that *Index* does not contain boxes that are too small. This is true by applying Lemma 11 to smooth neighborhoods of each connected component of the critical set. Finally:

Theorem 12 *The main algorithm returns an isotopic piecewise linear approximation of M.*

If one wishes to guarantee in addition that the Hausdorff distance between M and its approximation is less than say ε, by Theorem 3 it is sufficient to modify the positive watershed algorithm so as to control that the width of W is smaller than ε.

4 Conclusion

We have given an algorithm that approximates regular level sets of a given function with piecewise linear manifolds having the same topology. Though no implementation has been carried out, we believe that it should be rather efficient due to the simplicity of the involved predicates and the relative coarseness of the required space decomposition.

Appendix

We now briefly explain how to compute the index of a generic smooth function $f : \mathbb{R}^3 \to \mathbb{R}$ on a box $B \subset \mathbb{R}^3$ in a certified way. Without loss of generality, we assume that $B = [0, 1]^3$. Our approach is based on a recursive definition of the index of a vector field introduced in [12]. The central formula in this work is the following (see Fig. 12). If V denotes a vector field (in our case, $V = \nabla f$) defined on a compact smooth n-manifold M and not vanishing on ∂M, then the index of V satisfies:

$$\mathrm{Ind}(V) = \chi(M) - \mathrm{Ind}(\partial_- V).$$

Here $\partial_- V$ is a vector field defined on $\partial_- M$, which is the set of boundary points where V points inwards. On $\partial_- M$, $\partial_- V$ coincides with the projection of V on the tangent space of ∂M. Now suppose we can find a $(n-1)$-submanifold $M_1 \subset \partial_- M$ that contains all zeroes of $\partial_- V$. Then, to compute the index of V on M, it is sufficient to compute the index of $\partial_- V$ on M_1 (and the Euler characteristic of M_1). By repeated application of this principle, we can express the index of V as a sum of Euler characteristics and indices of vector fields defined over 1-manifolds, which are trivial to compute.

To apply this strategy to our case, in which $M = B$ has edges and corners, we conceptually consider offsets of M, which are smooth, and let the offset parameter go to 0. Almost by definition, in this setting the zeros of $\partial_- V$ are the points where V belongs to the normal cone and points inwards. Using interval analysis, it is not difficult to find a subset B_1 of $\partial_- B$ that contains all such points, and such that $\partial_- V$

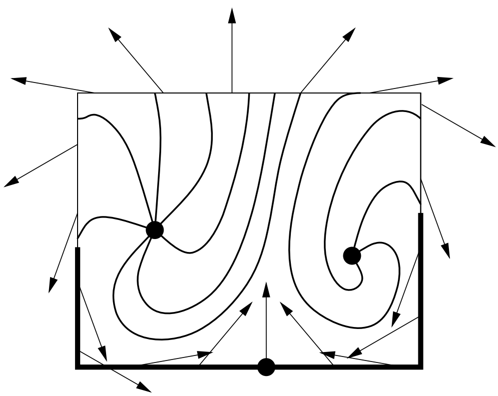

Fig. 12 An index 2 vector field V on a *square* C represented by a *few flow lines*. $\partial_- C$ is in *bold*. The *dot* on $\partial_- C$ represents the unique zero of $\partial_- V$, which has index -1

does not vanish on ∂B_1. To do this, we recursively subdivide the faces of the cube until all cells satisfy one of the two following conditions: either the cell does not contain a zero of $\partial_- V$, or it is included in $\partial_- B$. The union of the cells of the latter type will then provide a suitable B_1. For a square C lying on the face supported by, say, the plane $z = 1$, sufficient conditions ensuring that C does not contain any zero of $\partial_- V$ are

$$(V_z(C) > 0) \quad \text{or} \quad (0 \notin V_x(C)) \text{ or } (0 \notin V_y(C)).$$

Here $V_z(C) > 0$ for instance means that the z-coordinate of V is positive on C. The condition under which C is included in $\partial_- B$ is obviously $V_z(C) < 0$. Edges of the cube might also have to be subdivided. Without loss of generality we assume that edge E is supported by the line with equation $x = y = 1$. Then sufficient conditions under which E cannot contain a zero are as follows:

$$(V_x(E) > 0) \quad \text{or} \quad (V_y(E) > 0) \text{ or } (0 \notin V_z(E)).$$

Also, the condition under which E is included in $\partial_- B$ is $(V_x(E) < 0)$ and $(V_y(E) < 0)$. It can be checked that this subdivision process terminates if V has no zeroes on the surface of the cube, which is a generic condition. Upon termination of the subdivision process, we obtain a set B_1 to which the formula can be applied. It thus remains to recursively subdivide the boundary edges of B_1 in a similar way as above to complete the computation of the index of V.

References

1. Agrachev, A.A., Pallaschke, D., Scholtes, S.: On Morse theory for piecewise smooth functions. J. Dyn. Control Syst. **3**, 449–469 (1997)
2. Banchoff, Th.: Critical points and curvature for embedded polyhedra. J. Differ. Geom. **1**, 245–256 (1967)
3. Bern, M., Eppstein, D., Gilbert, J.: Provably good mesh generation. J. Comput. Syst. Sci. **48**, 384–409 (1994)
4. Bloomenthal, J.: Introduction to Implicit Surfaces. Morgan Kaufmann Series in Computer Graphics and Geometric Modeling. Morgan Kaufmann, San Mateo (1997)
5. Boissonnat, J.D., Oudot, S.: Provably good sampling and meshing of surfaces. Graph. Models **67**, 405–451 (2005)
6. Chazal, F., Cohen-Steiner, D.: A condition for isotopic approximation. In: Proc. ACM Symp. Solid Modeling and Applications, 2004
7. Cheng, S.-W., Dey, T.K., Ramos, E., Ray, T.: Sampling and meshing a surface with guaranteed topology and geometry. In: Proc. 20th Sympos. Comput. Geom., pp. 280–289, 2004
8. Dobkin, D.P., Levy, S.V.F., Thurston, W.P., Wilks, A.R.: Contour tracing by piecewise linear approximations. ACM Trans. Graph. **9**(4), 389–423 (1990)
9. Edelsbrunner, H.: Surface reconstruction by wrapping finite point sets in space. In: Aronov, B., Basu, S., Pach, J., Sharir, M. (eds.) Discrete and Computational Geometry. The Goodman–Pollack Festschrift, pp. 379–404. Springer, Berlin (2003)
10. Edelsbrunner, H., Harer, J., Zomorodian, A.: Hierarchical Morse complexes for piecewise linear 2-manifolds. In: Proc. 17th Annu. ACM Sympos. Comput. Geom., pp. 70–79, 2001
11. Goresky, M., MacPherson, R.: Stratified Morse Theory. Springer, Berlin (1988)
12. Gottlieb, D., Samaranayake, G.: The index of discontinuous vector fields. N. Y. J. Math. **1**, 130–148 (1995)
13. Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2002)
14. Hirsch, M.: Differential Topology. Springer, Berlin (1976)

15. Lopez, A., Brodlie, K.: Improving the robustness and accuracy of the marching cubes algorithm for isosurfacing. IEEE Trans. Vis. Comput. Graph. 9(1), 2003
16. Lorensen, W.E., Cline, H.E.: Marching Cubes: a high resolution 3D surface construction algorithm. Comput. Graph. **21**(4), 163–169 (1987)
17. Milnor, J.: Morse Theory. Ann. of Math. Studies, vol. 51. Princeton University Press, Princeton (1963)
18. Plantinga, S., Vegter, G.: Isotopic approximation of implicit curves and surfaces. In: Proceeding Symposium on Geometry Processing, pp. 251–260. Nice, France (2004)
19. Stander, B.T., Hart, J.C.: Guaranteeing the Topology of an Implicit Surface Polygonizer for Interactive Modeling. In: Proceedings of SIGGRAPH 97, pp. 279–286
20. Rourke, C.P., Sanderson, B.J.: Introduction to Piecewise-Linear Topology. Springer, Berlin (1982)
21. Ruppert, J.: A Delaunay refinement algorithm for quality 2-dimensional mesh generation. J. Algorithms **18**, 548–585 (1995)
22. Sard, A.: The measure of the critical values of differentiable maps. Bull. Am. Math. Soc. **48**, 883–890 (1942)
23. Shewchuk, J.R.: What is a good linear finite element? Interpolation, conditioning, and quality measures. In: Eleventh International Meshing Roundtable, Ithaca, New York, pp. 115–126, Sandia National Laboratories, September 2002
24. Snyder, J.M.: Interval analysis for computer graphics. In: Proceedings of SIGGRAPH 92, pp. 121–130
25. Velho, L., Gomes, J., de Figueiredo, L.H.: Implicit Objects in Computer Graphics. Springer, Berlin (2002)
26. Velho, L.: Simple and efficient polygonization of implicit surfaces. J. Graph. Tools **1**(2), 5–24 (1996). ISSN 1086-7651