Isotopic Implicit Surface Meshing

Jean-Daniel Boissonnat
Projet Géometrica, INRIA
Sophia-Antipolis, France
Jean-Daniel.Boissonnat
@sophia.inria.fr

David Cohen-Steiner
Computer Science
Department, Duke University,
USA
david@cs.duke.edu

Gert Vegter
Institute for Mathematics and
Computing Science, RUG,
Netherlands
gert@cs.rug.nl

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.

Categories and Subject Descriptors

G.1.2 [**Mathematics of Computing**]: Numerical Analysis—*Approximation of Surfaces and Contours*

General Terms

Algorithms, Theory

Keywords

Meshing, Computational Topology, Morse theory

1. INTRODUCTION

Implicit equations are a popular way to encode geometric objects; See, e.g., [4] and [23]. 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 [6].

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. 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* [23], are grid based, or, more generally, based on a tesselation of the ambient 3D space. They consist of two steps : first build a tesselation of space, and then analyze the intersection of the considered surface with each cell of the tesselation to construct the approximation. The celebrated marching cube algorithm [15] 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 [24], 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 theoriginal surface, except in very specific cases [14]. The problem of topologically correct polygonization of implicit curves in the plane is treated by Snyder in [22], who uses an adaptive enumeration method. His algorithm combines interval arithmetic with a quadtree tesselation 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.

To the best of our knowledge, there is only one paper devoted to the problem of homeomorphic polygonization of surfaces [17]. The main theoretical tool used in [17] 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 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 [13]. 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 iso-surface is smooth, that is does not contain any critical point, which is generic by Sard's theorem [20]. We also show that our algorithm can be easily modified so as to ensure geometric closeness between the implicit surface and its approximation. 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 (section 2). Then, we describe an algorithm for building a mesh satisfying these conditions, thereby leading to a provably correct isotopic polygonization algorithm (section 3).

2. A CONDITION FOR ISOTOPIC MESHING

Let f be a C^2 function from $\mathbb{R}^3$ to $\mathbb{R}$. We assume that $M = f^{-1}(0)$, the surface we want to polygonize, is compact. In what follows, T denotes a triangulation of a domain $\Omega \subset \mathbb{R}^3$ containing M and $\hat{f}$ the function coinciding 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

2.1.0.1 Collapses.
Loosely speaking, a collapse [18] is an operation which consists in removing cells from a simplicial complex whithout 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.
- $L = K \cup s$.
*- $\partial s \setminus K$ is the relative interior of t.*

Figure 1: Elementary collapse.

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 K can be obtained from L' by a sequence of elementary collapses.*

Definition 2 is illustrated in figure 2. In figure 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.

2.1.0.2 Smooth Morse theory.
The topology of implicit surfaces is usually investigated through Morse theory [16]. Given a real function f defined on a manifold,


Figure 2: The grey complex on the left collapses to the bold curve (dashed edges represent the subdivision). This is not true for the two other complexes.

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. 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 critical point p (see figure 3).


Figure 3: Smooth Morse theory in 2D.

In figure 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 signature $(1,1)$), 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 whose critical points all have non-degenerate Hessian are called *Morse functions*. From now on, we will assume that f is a Morse function. Also, we require that 0 is not a critical value of f, which implies that M is a manifold.

The number n of minus signs in the signature of the Hessian at p is classically called the index of p. However, for consistency reasons that will appear later, we call the *index* of p the integer -1 to the n-th power. The *index* of f on a region V is the sum of indices of all critical points of f lying in V. The index satisfies the following important theorem :

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

2.1.0.3 PL Morse theory.

Morse theory has been extended to a broad class of non-smooth functions by Goresky and McPherson [10]. 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 by f, and that no vertex of T maps to 0 by f, which guarantees that $\hat{M}$ is a manifold. We refer to these assumptions as *genericity assumptions*. Let us first recall some well-known definitions [8, 10] :

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 all vertices of which but v are smaller than v. The* lower link $Lk^-(v)$ *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.*


Figure 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 $Lk^-(v)$. In particular, in 2D, topological changes occur whenever $Lk^-(v)$ is not connected or equals the link of v (right and middle cases in figure 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 $Lk^-(v)$ [2]. In particular, regular points all have

---

[1]By simplex we mean a closed cell of T of any dimension.
[2]A complex is collapsible if it collapses to a point.

index 0. The converse is not true however in dimension at least 3. Also, checking if a vertex is regular is easy in 3D : it is sufficient to check that the lower link and the upper link —similarly defined— are both non-empty and connected. Define the index of $\hat{f}$ on a region V to be the sum of 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 [2] to also satisfy the Poincare-Hopf index theorem. The following lemma, which is a consequence of a result proved in [1] will be used later :

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

2.2 Main result

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

THEOREM 3. *Let W be a subcomplex of T.*
If W satisfies the following conditions :
1. *f does not vanish on ∂W.*
2. *W contains no tetrahedron of T containing a 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.

In the conclusion of the theorem, 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 [5]. 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: Condition 0. is needed.

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

$+\hat{min}$
$-\hat{min}$
$+\hat{max}$
$-\hat{max}$
$+\hat{s}$
$-\hat{s}$
$+min$

T with positive vertices, one can build an example where M has an extra component w.r.t. $\hat{M}$ without violating conditions involving critical points and their indices. Indeed, in figure 5, f has index 0 on the triangle, since minima have index 1 and saddle points have index -1.


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

Figure 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.

$+\hat{min}$

Figure 7: Condition 2'. and 4. are needed.

Figure 7 is 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 figure 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 (ie $S \cap T \neq \emptyset$).*
Assume the boundary of S is connected, as well as T and $X \setminus T$.
If the $X \setminus S$ 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. Let S and T be two such sets. ∂S is the disjoint union of $\partial S \cap \text{int}(T)$ and $\partial S \cap \text{int}(X \setminus T)$ 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 $\partial S \cap \text{int}(T)$ is empty then $\partial S \subset \text{int}(X \setminus T)$ that is $T \cap \partial S$ is empty. As a consequence, T is included in $\text{int}(S)$ or in $\text{int}(X \setminus S)$ by connectedness. Since S and T meet, we have that $T \subset \text{int}(S)$.
Now if $\partial S \cap \text{int}(X \setminus T)$ is empty then $X \setminus T$ is contained in $\text{int}(S)$ or in $\text{int}(X \setminus S)$ by connectedness again. Similarly as above it has to be contained in $\text{int}(X \setminus S)$, which means that $S \subset T$. Thus $\text{int}(S) \subset \text{int}(T)$ so $\partial S \supset S \setminus \text{int}(T) = S \cap \partial T$. If S would meet ∂T, then ∂S and ∂T would meet, which is impossible : S is included in the interior of T. $\square$

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

PROOF. Hypothesis **3** implies easily that V collapses to $\hat{M} \cap V$. Thus V contains a simplex having positive and negative vertices. As a consequence, f vanishes on V. Since f does not vanish on ∂W (**1**), M intersects V. Also, M does not meet the boundary of V (**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 $\hat{M}$, $\mathbb{R}^3 \setminus V$ also has exactly one bounded component which we denote by A and one unbounded component we denote by B. 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, since the complement of $A \cup V$ is B, it 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, $i = 1..n$ the connected components of $M \cap V$. 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 because $N_1 \setminus N_2 \subset V$. $\square$

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

Figure 8: Proof of lemma 5.

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 **3** 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, by condition **1**, ∂D does not meet E. Indeed, f is positive at vertices of ∂D and does not vanish on $\partial D \subset \partial W \cup \partial \Omega$. So E, being connected, is included in the interior of D. But then E is compact and thus f reaches its minimum on E : 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 half of the bound. The other half 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 [5].

PROPOSITION 7. *Let $\hat{M}$ be a 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[3] $\tilde{W}$ of $\hat{M}$.*

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

Indeed, considering a regular neighborhood of W [18] yields the desired topological thickening $\tilde{W}$. The proof of this fact, which we omit, relies on the uniqueness theorem for regular neighborhoods from piecewise-linear topology [18].

---

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

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 section 3.1, we give its definition, basic properties, and construction algorithms. Section 3.2 describes the meshing algorithm itself, which ensures that V 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 $W^+ = \hat{M}$.
mark all vertices of $\hat{M}$.
while there is a positive regular unmarked vertex v of T
     such that the vertices of $Lk^-(v)$ are marked
do
  set $W^+ = W^+ \cup St^-(v)$.
  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^+.


Figure 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.

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 $Lk^-(v_i)$ is included in W_i^+. If it is not the case, let u be the largest vertex of some simplex s of $Lk^-(v_i)$ that is not in W_i^+. s is in $St^-(u)$ which is hence not included in W_i^+. So u is either critical or not marked yet, which is a contradiction since v_i is marked. So $Lk^-(v_i) \subset W_i^+$. Now since v_i is regular, $Lk^-(v_i)$ is collapsible. Consider a sequence of elementary collapses allowing to collapse $Lk^-(v_i)$ to p and let

305

$s_j \subset Lk^-(v_i), j = 1..n$ be the sequence of simplices defining these elementary collapses. The simplices $conv(s_j \cup v_i)$ and the edge pv_i define a valid sequence of elementary collapses allowing to collapse $W_{i+1}^+ = W_i^+ \cup St^-(v_i)$ 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 $u \in St^-(v)$ or $u = v$. We will note 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 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.*

PROOF. The vertices of W^+ have this property by construction. Let $p \notin W^+$ be a positive vertex and assume p does not flow into any positive critical point. In particular, p is regular. Hence, as $p \notin W^+$, the lower link of p, which is not empty, has to contain either a critical vertex or an unmarked one. It cannot contain a critical point because as T conforms to $\hat{M}$, vertices in $Lk^-(p)$ are all non-negative, and so p would flow into a positive critical point. There is thus an unmarked vertex p_1 in $Lk^-(p)$. If p_1 can be chosen positive, then p_1 satisfies the same assumptions as p so one can define p_2 in a similar way. By going on, one obtains a strictly decreasing sequence of positive vertices, that thus has to end. Let p_k its last term. $Lk^-(p_k)$ contains no positive unmarked vertices. But as T conforms to $\hat{M}$, vertices in $Lk^-(p_k)$ are all non-negative. As 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. Call c the minimum of $|\hat{f}(v)| = |f(v)|$ over all critical points v of $\hat{f}$.

LEMMA 10. *W contains all vertices whose image by $|f|$ 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 in 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 say 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 **a**), then this does not happen and the result W' of the modified algorithm is equal to W. In our meshing algorithm, we will not build the mesh S, but rather make sure condition **a** 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. Remaining vertices all belong to W_1^+. Then, mark these vertices and apply the positive watershed algorithm loop to get the vertex set of W_1^+.

Remark. The presented definition of a watershed seems quite well-behaved and leads to an easy construction algorithm, but it is not fully satisfactory. In particular, 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 [8].

3.2 Main algorithm

Assume the critical points of f are given. Theorem 3 enables us to build a mesh isotopic to M using only one simple predicate, *vanish*. *vanish* takes a triangle or a box and returns true if f vanishes on that triangle or that box. We actually not even need a predicate, but rather a filter. More precisely, *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 elements tends to 0. As noticed by Shewchuk [21], 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 whose angles 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 180°. 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 figure 10). The main algorithm uses an octree O, the associated triangulation T, 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 using interval analysis. Besides, we maintain four sets of boxes ordered by decreasing size. *Critical* contains all boxes containing a critical point of $\hat{f}$ that is not in a

Figure 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).

box containing a critical point of f. *Index* contains all boxes neighboring a box b containing a critical point of f and such that f and $\hat{f}$ have different indices on b. *Boundary1* contains all boxes containing two neighbors -in T- of opposite signs one of which is critical for $\hat{f}$ (condition **a**, see paragraph **Non conforming case**). Finally, *Boundary2* contains all boxes that contain a triangle t of $\partial W'$ such that *vanish*(t) is true and that are not included in W'.

Main Algorithm

Initialization Refine O until *vanish*(b) is false for all boxes containing at least one critical point of f.
compute T and W', and the four sets *Critical*, *Boundary1*, *Boundary2*, and *Index*.
while (**true**) **do**
  update T, W', and the four sets.
  if *Critical*$\neq \emptyset$ **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 bounded component of $\Omega \setminus W'$ **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 almost amounts to its termination. The only problem is that W' might contain some critical point of f, thereby violating condition **2.**. It thus seems that the definition of W' needs to be slightly modified. The modification consists in 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.

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. 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 *Critical* 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, this 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.

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 [10] for a complete exposition of this subject.

PROOF. Let $p \in \partial A$ and $d(p)$ denote the local feature size of p with respect to the 2-skeleton of T, as defined -in 2D- by Ruppert [19]. Simplices of T that meet the open ball centered at p of radius $d(p)$ all share a vertex $v(p)$ -by definition, $d(p)$ is the largest number such that this holds. 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, the half of the distance from ∂A to the closest box that does not meet ∂A.

**Figure 11: Proof of lemma 11.**

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, considered gradients all 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 $\nabla f(p)$ make a positive dot product.

Let now U_1 be a neighborhood of ∂A whose closure does not contain any vertex of T and 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 for $i = 1..2$, u_i is a non negative smooth function defined on $\mathbb{R}^3$, with support in U_i, and such that $u_1 + u_2$ is identically 1. In particular, u_2 equals 1 on the complement of U_1, and vice versa. 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 $\mathbb{R}^3 \setminus U_2 \supset \partial A$. Now recall that $\nabla \tilde{f}$ and ∇f make a positive dot product on ∂A. 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 is invariant under homotopy [12], 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 makes a positive dot 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. $\square$

Suppose that some box b of *Critical* 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 *Critical* 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 make a positive dot product with ∇f which is a contradiction with lemma 2 which implies that v is not critical. So the conclusion is that *Critical* becomes -at least temporarily- empty after a finite number of consecutive splittings of boxes in *Critical*.

Now if the algorithm splits a box b in *Boundary1*, then b contains a say positive critical point of $\hat{f}$, which belongs to a box containing a critical point of f as *Critical* 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 eternally.

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 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 too small boxes. This is true by applying lemma 11 to small offsets of the boxes containing critical points of f. Finally :

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

Furthermore, if one wishes to guarantee that the Hausdorff distance between M and its approximation is less than say ε, it suffices to modify the positive watershed algorithm so as to control that the width of W is smaller than ε, thanks to theorem 3.

4. CONCLUSION

We have given an algorithm that approximates regular level sets of a given function with piecewise linear manifolds having the same topology. Moreover, our algorithm can be modified so as to ensure geometric closeness in the Hausdorff sense. Though no implementation has been carried out yet, 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.

The main drawback of our algorithm is that it requires, as is, the knowledge of the critical points of the considered function. A

closer look shows that we almost only need to find a set of boxes containing all the critical points, and on which the function does not vanish. This task, corresponding to the initialization step in the main algorithm, can be done in a certified way using interval analysis. Once this is done, the only remaining problem is to compute the index of the function on these boxes in a robust way. In a forthcoming version of the paper, we will show how this can be done in the framework of interval analysis, thereby giving a complete solution to the problem. Also, we plan to adapt the algorithm to the case of surfaces with boundaries, which is useful for instance when one wants to study the considered level set inside a user-specified bounding box.

5. REFERENCES

[1] A.A. Agrachev, D. Pallaschke, and S. Scholtes, *On Morse Theory for Piecewise Smooth Functions*, Journal of Dynamical and Control Systems 3, pp. 449-469, 1997.

[2] Th. Banchoff, *Critical points and curvature for embedded polyhedra,* J. Diff. Geom 1 (1967) $245 - 256$.

[3] M. Bern, D. Eppstein, and J. Gilbert, *Provably Good Mesh Generation*, Journal of Computer and System Sciences 48 (1994) 384-409.

[4] J. Bloomenthal, *Introduction to Implicit Surfaces*, Morgan Kaufmann Series in Computer Graphics and Geometric Modeling, Morgan Kaufmann Publishers, Inc., 1997.

[5] F. Chazal and D. Cohen-Steiner, *A condition for isotopic approximation*, submitted. Available at : ftp://ftp-sop.inria.fr/prisme/dcohen/Papers/isotopie.ps

[6] D.P. Dobkin, Silvio V. F. Levy, William P. Thurston and Allan R. Wilks, *Contour Tracing by Piecewise Linear Approximations*, ACM Transactions on Graphics, Vol.9, No.4, October 1990, Pages 389-423.

[7] A. Dold, *Lectures on algebraic topology*, Springer 1972.

[8] H. Edelsbrunner, J. Harer, and A. Zomorodian. Hierarchical Morse complexes for piecewise linear 2-manifolds. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 70-79, 2001.

[9] H. Edelsbrunner, *Surface reconstruction by wrapping finite point sets in space*, Ricky Pollack and Eli Goodman Festschrift, ed. B. Aronov, S. Basu, J. Pach and M. Sharir, Springer-Verlag, to appear.

[10] M. Goresky and R. MacPherson, *Stratified Morse Theory*, Springer Verlag, 1988.

[11] D. Gottlieb and G. Samaranayake, *The index of discontinuous vector fields*, New-York Journal of Mathematics, 1, pp 130-148, 1995.

[12] A. Hatcher, *Algebraic topology*, Cambridge University Press, 2002. http://www.math.cornell.edu/ hatcher.

[13] M. Hirsch, *Differential Topology*, Springer Verlag, 1976.

[14] A. Lopez and K. Brodlie, *Improving the Robustness and Accuracy of the Marching Cubes Algorithm for Isosurfacing*, IEEE Transactions on Visualization and Computer Graphics, Vol.9, No.1, January-March 2003.

[15] W.E. Lorensen and H.E. Cline, *Marching Cubes : A high resolution 3D surface construction algorithm*, Computer Graphics, 21(4):163-169, 1987.

[16] J. Milnor, *Morse theory*, Ann. of Math. studies 51, Princeton University Press, Princeton N.J. 1963.

[17] B.T. Stander and J.C. Hart, *Guaranteeing the Topology of an Implicit Surface Polygonizer for Interactive Modeling*, Proceedings of SIGGRAPH 97, pages 279-286.

[18] C.P. Rourke and B.J. Sanderson, *Introduction to Piecewise-Linear Topology*, Springer-Verlag, 1982.

[19] J. Ruppert, *A Delaunay refinement algorithm for quality 2-dimensional mesh generation*, Journal of Algorithms, 18:548–585, 1995.

[20] A. Sard, *The measure of the critical values of differentiable maps*, Bull. Amer. Math. Soc. 48 (1942), 883-890.

[21] J.R. Shewchuk, *What Is a Good Linear Finite Element? Interpolation, Conditioning, Anisotropy, and Quality Measures*, in preparation.

[22] J.M. Snyder, *Interval analysis for computer graphics*, Proceedings of SIGGRAPH 92, pages 121–130.

[23] L. Velho, J. Gomes, L.H. de Figueiredo, *Implicit Objects in Computer Graphics*, Springer-Verlag, 2002.

[24] L. Velho. Simple and efficient polygonization of implicit surfaces. Journal of Graphics Tools, 1(2):5-24, 1996. ISSN 1086-7651.