Optimal discrete Morse functions for 2-manifolds

THOMAS LEWINER, HÉLIO LOPES AND GEOVAN TAVARES

Department of Mathematics — Pontifícia Universidade Católica — Rio de Janeiro — Brazil
{tomlew, lopes, tavares}@mat.puc--rio.br.

Abstract. Morse theory is a powerful tool in its applications to computational topology, computer graphics and geometric modeling. It was originally formulated for smooth manifolds. Recently, Robin Forman formulated a version of this theory for discrete structures such as cell complexes. It opens up several categories of interesting objects (particularly meshes) to applications of Morse theory.

Once a Morse function has been defined on a manifold, then information about its topology can be deduced from its critical elements. The main objective of this paper is to introduce a linear algorithm to define optimal discrete Morse functions on discrete 2-manifolds, where optimality entails having the least number of critical elements. The algorithm presented is also extended to general finite cell complexes of dimension at most 2, with no guarantee of optimality.
Keywords: ***Morse Theory. Forman Theory. Computational Topology. Computational Geometry and Object Modeling.***

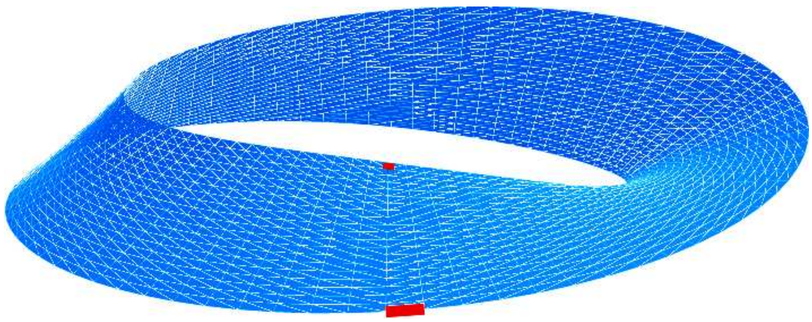

Figure 1: *An optimal discrete Morse function on a Möbius strip, with 2 critical cells.*

1 Introduction

Applications of computational topology [6, 23] in computational science and engineering are many and growing. These include meshing, morphing, feature extraction, data compression, surface coding and more, in areas such as computer graphics, solid modeling, computational medicine and astrophysics. In applications to computer graphics, being able to detect topological singularities helps designing more robust and efficient algorithms in time and space.

Morse theory [19] is a fundamental tool for investigating the topology of smooth manifolds. Particularly for computer graphics, many applications have been induced [8, 21, 14] from the smooth case.

Morse proved that the topology of a manifold is very closely related to the critical points of a real smooth map defined on it (i.e., the points where the gradient vanishes). The simplest example of this relationship is the fact that if the manifold is compact, then any continuous function must have a maximum and a minimum. Morse theory provides a significant refinement of this observation.

Forman's discrete Morse theory. Recent insights in Morse theory by Forman [10, 11] extend several aspects of this fundamental tool to discrete structures. Its combinatorial aspect allows computation completely independent of a geometric realization: the algorithm does not require any coordinate or floating-point calculation. Forman proves several results and provides many applications of his theory [12, 13], and new ones have appeared recently [2].

Once a Morse function has been defined on a smooth manifold, then informations about its topology can be deduced from its critical points. Similarly to the smooth case, Forman proved that the topology of a cell complex can be partly read out of the critical cells of a discrete Morse function defined on it.

Preprint MAT. 03/02, communicated on December 3^{rd}, 2001 to the Department of Mathematics, Pontifícia Universidade Católica — Rio de Janeiro, Brazil. The corresponding work was published in Computational Geometry: Theory and Applications, volume 26, numer 3, pp 221–233. Elsevier, 2003.

Main results. We provide here a linear algorithm to compute *optimal* discrete Morse functions on 2–manifolds, where optimality entails having the least number of critical elements. We also extend this linear algorithm to build valid discrete Morse functions for general finite cell complexes of dimension 2 (non–manifolds). Reaching optimality in the general case is here proven to be NP-hard with no polynomial approximation (see section 3(c) *The optimality problem and its complexity*). However, the experiments of our algorithm on non-manifolds lead to optimal results in most of the cases (see section 5 *Results*).

Prior work. As far as we know, there have been no results yet in computing explicitly a discrete Morse function with optimality requirements. The work of Babson and Hersh [1] gives a construction and some interpretation of Forman functions on cell complexes out of lexicographic orders. The construction of such lexicographic orders is not mentioned in [1] and the optimality of the resulting function depends on it.

Outline. In section 3 *Elements of discrete Morse theory*, we give a short introduction to Forman's theory. The basics elements of topology we will need for the proof of the optimality are grouped in section 2 *Basic concepts of combinatorial topology*. The algorithm, its proof and its extension are described in section 4 *Computing optimal discrete Morse functions* and section 4(c) *Extension to general 2–cell complexes*. We finally show some of the results in section 5 *Results*.

2 Basic concepts of combinatorial topology

(a) Finite cell complex

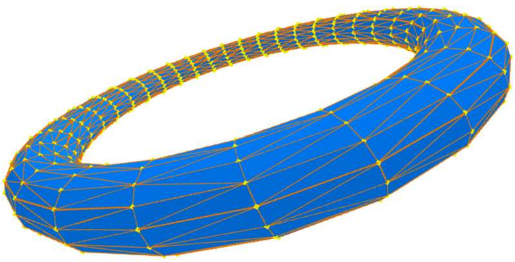

Figure 2: *A triangulated torus model.*

A cell complex is, roughly speaking, a generalization of the structures used to represent solid models: it is a consistent collection of cells (vertices, edges, faces…). Figure 2 gives an example of such a structure. A complete introduction to cell complexes can be found in [18].

More formally, a *cell* $\alpha^{(p)}$ of dimension p is a set homeomorphic to the open p–ball $\{x \in \mathbb{R}^p : \|x\| < 1\}$. When the dimension p of the cell is obvious, we will simply denote α instead of $\alpha^{(p)}$.

A *cell complex* K is built by starting off with a discrete collection of 0–cells (vertices) called K^0, then attaching 1–cells (edges) to K^0 along their boundaries, obtaining K^1, then attaching 2–cells (faces) to K^1 along their boundaries, and so on, giving spaces K^n for each n.

A cell complex will be said to be finite when it is built out of a finite number of cells. In this work, we will consider only finite (and thus regular) cell complexes.

A p–cell $\alpha^{(p)}$ is a sub–face of a q–cell $\beta^{(q)}$ $(p < q)$ if $\alpha \subset \mathit{closure}(\beta)$. If $q = p + 1$, we will use the notation $\alpha^{(p)} \prec \beta^{(q)}$., and say that α and β are *incident*.

(b) Homology groups and Betti numbers

Choice of $\mathbb{Z}_2$ homology The topology of a 2–manifold is completely described by its orientability and its Betti numbers, which are the rank of its homology groups. We will use the Betti numbers together with the weak Morse inequalities (section 3(b) *Morse inequalities*) to provide a lower bound to the number of critical cells of a discrete Morse function.

In the weak Morse inequalities, we are free to choose any field for the coefficient of the homology groups. We will choose here the field $\mathbb{Z}_2$, as this will generate higher values for the Betti numbers (in the presence of torsion, $\mathbb{Z}_2$ homology "counts" some torsion as cycles). For the reader interested only in computing the Betti numbers, a classical algorithm is described in [5].

We would like to suggest [4] for a basic introduction to homology theory, and [7] for a concise presentation of $\mathbb{Z}_2$ homology.

The chain group and the boundary operator Let K be a cell complex. A p–chain $c^{(p)}$ is a subset of p–cells in K.

$$c^{(p)} = \sum_{\sigma^{(p)} \in K} c_\sigma . \sigma^{(p)}$$

The coefficients $c_\sigma \in \mathbb{Z}_2$ only counts whether the cell belongs to the chain or not. The addition of two p–chains is trivially defined element–wise on each cell. In other words, the addition of two p–chains is the symmetric difference of the two sets. The group C_p of all p–chains is called the *chain group*. The empty set is the zero element of C_p.

The *boundary* $\partial_p(\sigma^{(p)})$ of a p–cell $\sigma^{(p)}$ is the collection of its $(p-1)$–dimensional faces, which is a $(p-1)$–chain. The boundary operator ∂_p is extended to p–chain by linearity:

$$\partial_p \left(\sum_{\sigma^{(p)} \in K} c_\sigma . \sigma^{(p)} \right) = \sum_{\sigma^{(p)} \in K} c_\sigma . \partial_p \left(\sigma^{(p)} \right)$$

A p–cycle $z^{(p)}$ is a p–chain whose boundary is null, $\partial_p\left(z^{(p)}\right) = 0$; and a p–boundary $b^{(p)}$ is the boundary of a (p+1)–chain, $b^{(p)} = \partial_{p+1}\left(c^{(p+1)}\right)$.

As the boundary operators ∂_p preserve the addition from C_p to C_{p-1}, the set of the p–boundaries $\mathrm{im}\partial_{p+1}$ and the set of the p–cycles $\ker\partial_p$ are subgroups of C_p.

Homology groups and Betti numbers An essential property of the boundary operators is that the boundary of every boundary is empty ($\partial_p \circ \partial_{p+1} = 0$). Therefore, every p–boundary is a p–cycle and $\mathrm{Im}\partial_{p+1} \subseteq \ker\partial_p$.

For each p, the group $H_p = \ker\partial_p \,/\, \mathrm{Im}\partial_{p+1}$ is called the *p–th homology group* (with coefficients in $\mathbb{Z}_2$).

The corresponding work was published in Computational Geometry: Theory and Applications, volume 26, numer 3, pp 221–233. Elsevier, 2003.

Homology groups are commutative and finitely–generated (as the cell complex is finite). Thus, they can be written as $H_p \cong \mathbb{Z}_2^{\beta_p}$, where β_p is called the p–th *Betti number*. The basic interpretation for Betti numbers is a way of counting "holes" in a given complex: β_0 counts the number of connected component, β_1 is the dimension of the vector–space of 1–cycles in a surface, β_2 the voids of a solid, and so forth.

(c) Combinatorial manifold

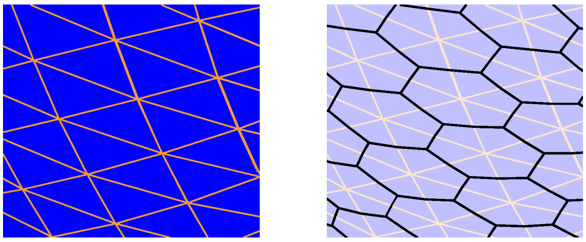

Figure 3: *A part of a triangulation (on the left) and its dual pseudograph (on the right)*

An n*–manifold* is a topological space where each point has a neighborhood homeomorphic to either $\mathbb{R}^n$ or $\mathbb{R}_+ \times \mathbb{R}^{n-1}$. The set of points whose neighborhood is $\mathbb{R}_+ \times \mathbb{R}^{n-1}$ is called the boundary of the manifold. It can be shown [18] that if a finite cell complex is an n–manifold, then each (n-1)–cell is the sub–face of either one or two n–cells.

Thus, the (n-1)–cells of a manifold can be thought as links of a pseudograph (i.e., a non–simple graph in which both loops and multiple edges are permitted) whose nodes are the n–cells of the manifold. This pseudograph will be called the *dual pseudograph* of the manifold. In particular, the cells of the boundary will be represented by loops in the dual pseudograph. For example, Figure 3 shows a part of a triangulation and its dual pseudograph (which is here a simple graph).

The classification theorem The classification theorem for surfaces [4] completely characterizes the topology of 2–manifolds in terms of their Euler characteristic and their orientability. The following theorem is a simple consequence of it:

Theorem 1 *For a connected 2–manifold K, orientable or not:*

$$H_0(K) \cong \mathbb{Z}_2 \quad \text{and}$$
$$H_2(K) \cong \begin{cases} \mathbb{Z}_2 & \textit{if } K \textit{ has no boundary} \\ 0 & \textit{if } K \textit{ has a boundary} \end{cases}$$

In this paper, we will consider 2–manifolds with possibly many connected component. As our algorithm processes on each component separately, we will be able to use this theorem to guarantee the optimality of the result in the case of 2–manifolds.

3 Elements of discrete Morse theory

Forman's discrete Morse theory relates the topology of a cell complex to the *critical cells* of a *discrete Morse function*. For a complete introduction, see Forman's presentations [12, 13] and Chari's works [2, 3]. The focus of this paper is to provide an *optimal* construction of discrete Morse functions – optimal in the sense that the function has the minimum possible number of critical cells in each dimension. We will introduce some basics of Forman's theory in the next paragraphs, and discuss the *optimality problem* in the last one.

(a) Discrete Morse function

Definition 2 (Discrete Morse function [10]) *A function f mapping each cell of a cell complex K to a real value is a discrete Morse function if it satisfies, for every cell $\sigma^{(p)} \in K$:*

$$\#\left\{\tau^{(p+1)} \succ \sigma^{(p)} : f(\tau) \leq f(\sigma)\right\} \leq 1$$
$$\#\left\{\upsilon^{(p-1)} \prec \sigma^{(p)} : f(\upsilon) \geq f(\sigma)\right\} \leq 1$$

So there is at most one "counterbalancing" sub–face $\tau^{(p+1)}$ of codimension 1 and one "counterbalancing" bounding cell $\upsilon^{(p-1)}$ for every cell $\sigma^{(p)}$. It is easy to show that a cell cannot have both of them. A cell that has none of them will be called *critical*:

Definition 3 (Critical Cell [10]) *A cell $\sigma^{(p)}$ is a critical cell of f if:*

$$\#\left\{\tau^{(p+1)} \succ \sigma^{(p)} : f(\tau) \leq f(\sigma)\right\} = 0$$
$$\#\left\{\upsilon^{(p-1)} \prec \sigma^{(p)} : f(\upsilon) \geq f(\sigma)\right\} = 0$$

We will denote by $\mathbf{m}_p(f)$ the number of critical cells of dimension p.

Examples One can define a trivial discrete Morse function by $f\left(\sigma^{(p)}\right) = p$, for which every cell is critical (see Figure 4(a)). Of course, not all functions are valid as discrete Morse function: on Figure 4(b) for example, the face (with value 4) and the edge with value 0 are assigned values invalid for definition 2. The critical cells of Figure 4(c) are assigned values 0 and 5.

(b) Morse inequalities

We aim to construct optimal Morse functions, i.e., functions with the minimum number of critical cells. To ensure we reached the optimality, we need a lower bound to the number of critical cells. The *weak Morse inequalities* provide such a bound in terms of the Betti numbers. Those inequalities are valid whatever the field is chosen to calculate the Betti numbers [11].

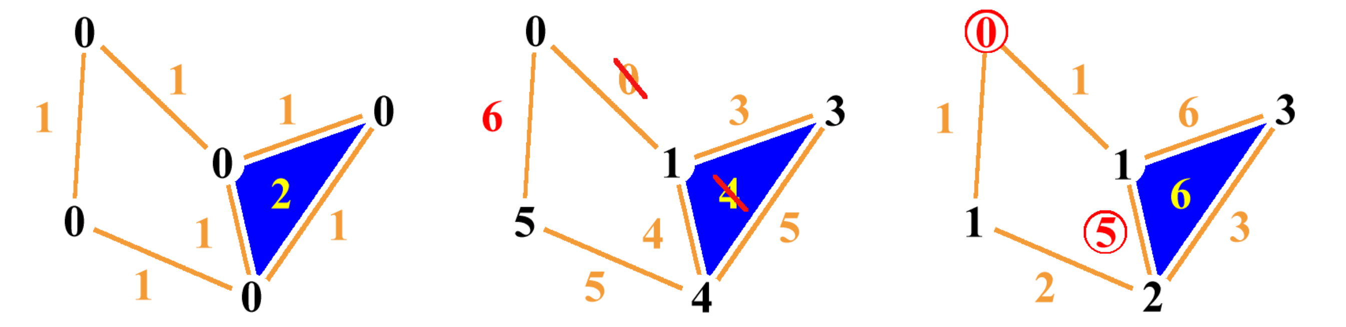

(a) trivial (valid) (b) invalid (c) valid (optimal)

Figure 4: *Examples of trivial, invalid and optimal discrete Morse functions on a small cell complex with 5 vertices, 6 edges and 1 face.*

Theorem 4 (Weak Morse inequalities [11]) *For a given finite cell complex K, any discrete Morse function f defined on it satisfies:*

$$\forall p, \quad m_p(f) \geq \beta_p(K)$$
$$\chi(K) = \sum_p (-1)^p \cdot \beta_p(K) = \sum_p (-1)^p \cdot m_p(f).$$

(c) The optimality problem and its complexity

Optimality problem	
Instance:	A pair (K, n), where K is a finite cell complex of dimension at least 2 and n is a non-negative integer
Question:	Does it exist a discrete Morse function on K with at most n critical cells?

The Morse optimality problem reduces, in the general case, to the Collapsibility problem, which is proven in [9] to be a *MAX–SNP hard problem*, i.e., an NP-hard problem for which any polynomial approximation algorithm can lead to a result arbitrary far from the optimum:

Collapsibility problem	
Instance:	A pair (K, n), where K is a finite simplicial complex of dimension 2 and n is a non-negative integer
Question:	Does K contain a subset F of 2–simplicies of cardinality at most n such that $K \setminus F$ collapses?

Let's consider a simplicial complex K of dimension at most 2. Suppose we find a discrete Morse function f on K having the minimum possible number of critical cells. The number of critical vertices is the number of connected components. From the weak Morse inequalities, we deduce that the discrete Morse function has the minimum number of critical faces. Then, [11, Theorem 3.3] proves that the subset F of the critical 2–cells is a solution to the Collapsibility problem.

However for 2–manifolds, we provide an algorithm that reaches optimality in linear time (see section 4 *Computing optimal discrete Morse functions*). For the case of general finite cell complexes of dimension 2 (see section 4(c) *Extension to general 2–cell complexes*), our algorithm is still linear and generates a valid discrete Morse function, which is optimal in most of the cases we have experimented (see section 5 *Results*).

4 Computing optimal discrete Morse functions

Our construction defines first the discrete Morse function on a spanning tree of the dual pseudograph, and then on its complement graph. The optimality of the resulting function relies on the weak Morse inequalities.

(a) Algorithm outline

Given a finite cell complex K that has the topology of a 2–manifold, the algorithm proceeds on each connected component in 4 steps:

1. Construct a spanning tree T on the dual pseudograph of K.
2. If K has a boundary, add one boundary edge of K to T.
3. Define the discrete Morse function on T.
4. Define the discrete Morse function on the complement of T.

First step: Construction on a face–spanning tree. The face–spanning tree T can be constructed by any of the standard algorithms [22]. In particular, we can use some mesh compression's strategies. For example, Figure 6 shows a spanning tree constructed by the Edgebreaker's compression algorithm [17].

Second step: Addition of one edge of the boundary. We test whether the manifold has a boundary during the first step. If we found a boundary edge, we add it to T. This edge will be a loop in the dual pseudograph, so T becomes a pseudograph. For example on Figure 5, the boundary edge (represented by a loop) has been added to T.

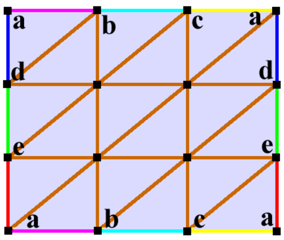

(a) A triangulated torus (with the identified vertices and edges)

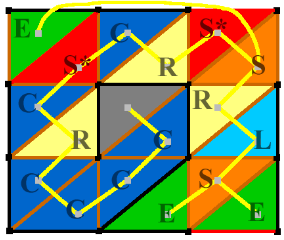

(b) Edgebreaker mesh compression traversal.

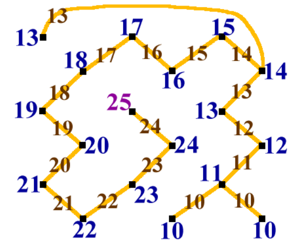

(c) The discrete Morse function on the spanning tree T. The root is the only critical face, having value 25.

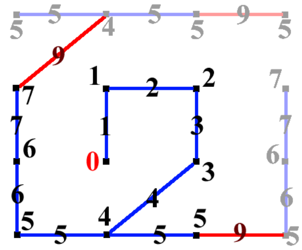

(d) The complement graph G and its discrete Morse function: 1 critical vertex (0) and 2 critical edges (9).

Figure 6: *Triangulated torus. The Edgebreaker mesh compression traversal can be used to define the spanning tree T of step 1: the critical edges correspond to the handles of [17].*

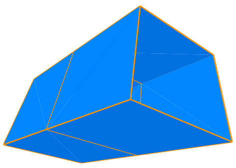

(a) The cylinder model with 8 square faces.

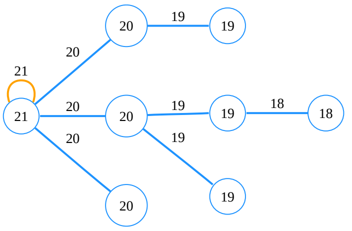

(b) The face–spanning tree T (the loop represents the boundary edge added at step 2).

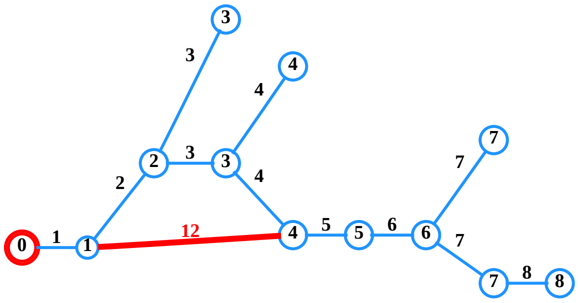

(c) The discrete Morse function on the complement graph G.

Figure 5: *Algorithm's execution on a small cylinder model: 1 critical vertex and 1 critical edge with values 0 and 12.*

Third step: Definition of the function on T. We select a root of T, and we assign to every node of T (i.e. 2–cells of K) its height in the tree plus $v + 1$, where v is the number of 0–cells of K. We assign to every edge of T (i.e. 1–cells of K) the minimum value of its two end nodes (see Figure 5).

Fourth step: Definition of the function on the complement of T. We will now consider G, the complement of T: G is a simple graph whose nodes are the vertices of K, and whose edges are the edges of K that are not represented in T. We build another spanning tree U on G. We assign to every node of G its edge distance to a selected root of U, and to every edge of U the maximum value of its two end nodes. This can be done at the same time as building the spanning tree U. We finally assign the value v to each edge of $G \setminus U$ (see the critical edge of Figure 5 with value 12).

(b) Proof and analysis of the algorithm

Valid discrete Morse function. From the construction on the trees T (step 3) and U (step 4), the resulting function f respects the inequalities of Definition 2 (for each connected component K). Moreover, there is exactly one critical vertex: the root of U. If K has no boundary, the root of T will be the unique critical face. In the other case, there is no critical face in K.

We now just need to check that the edges of K are assigned valid Morse values. From the value of the constant v, the critical edges are those of $G \setminus U$, which are assigned a value greater than the value of any vertex, and inferior to the value of any face (es). The inequalities of Definition 2 there are v nodes in G, so at most v different regular valuare obvious for each cell represented in the trees T and U. From the value of the v, every cell of T has a value greater than any cell of G. Thus, those inequalities are strictly respected between the edges of T and the vertices of G, and between the edges of U and the faces in T.

Thus, our construction yields a valid discrete Morse function with exactly 1 critical vertex ($m_0 (f) = 1$), possibly many critical edges, and 1 critical face ($m_2 (f) = 1$) only if K is a manifold without boundary ($m_2 (f) = 0$ otherwise).

Optimal discrete Morse function. From Theorem 1, we obtained the same values $m_0 (f)$ and $m_2 (f)$ as the values of $\beta_0 (K)$ and $\beta_2 (K)$ as the Betti numbers in the same cases. From the second weak Morse inequality (Theorem 4), we deduce $m_1 (f) = \beta_1 (K)$. Therefore, we reached our lower bound: the function built by the algorithm is a valid, optimal discrete Morse function.

Complexity. Once the spanning trees are built, the algorithm visits each node and edge at most once. Thus, steps 2, 3 and the second part of step 4 are of linear complexity. Building a spanning tree can be linear with a simple greedy algorithm. Therefore, the whole algorithm is linear in time.

(c) Extension to general 2–cell complexes

If the given complex does not have the topology of a 2–manifold, some edges can be incident to 3 faces and the above proof does not state anymore. However, the algorithm still produces a valid discrete Morse function, which is optimal in several cases. In fact, a cell complex of dimension 2 is not a manifold if it combines some of the following 3 elements (see Figure 7):

(1) Dangling edge: an edge not incident to any face.

(2) Singular vertex: a vertex that, when removed, disconnects incident faces.

(3) Non–regular edge: an edge incident to 3 or more face.

(1) Dangling edge. That case reduces to a graph glued to a complex. For a graph, β_1 is the number of edges that cannot be included in a spanning tree. This graph will be processed in step four of the algorithm. Thus the number of critical edges is still β_1 : the algorithm still reaches the optimality in that case.

(2) Corner vertex. This case corresponds to several cell complexes glued at a vertex. For steps one, two and three of the algorithm, each of those cell complexes is processed as distinct connected component. During step four, the algorithm will generate only one critical vertex. Therefore, the algorithm still reaches the optimality in that case.

(3) Non–regular edge. That case is the most difficult one. We will only give a heuristic that always build valid discrete Morse function, but we know from section 3(c) *The optimality problem and its complexity* that, for some very particular cases, the resulting functions can be arbitrary far from the optimum.

In that case, we will first remove from the dual pseudograph the edges that are incident to 3 nodes or more (i.e. the non–regular edges). The algorithm then runs normally through steps one to four, and the non–regular edges that cannot be included in the spanning tree of step four will be critical.

5 Results

(a) Reaching optimality

We tested our algorithm on more than 150 models from various types: triangulations, quadrangulations and general polygons; manifolds and non-manifolds; models with a consistent connectivity and topology and raw scans or VRML importation with deficient topology. The algorithm always builds a valid discrete Morse function. For all the manifolds cases, the resulting function was optimal. For the non-manifolds complexes (in particular for the examples of Moriyama and Takeuchi [20]), the function had at most 4 redundant critical cells. The experimental results on a Pentium III, 550 MHz, confirm the linear complexity (Figure 10(a)) and the independence of the complexity from topology (Figure 10(a)).

(b) Mixing with geometry

Our construction is completely independent of the geometry. This advantage of Forman's discrete Morse theory appears here on two points. First, the whole algorithm is done without any floating–point operation. Second, it is possible to add some external constraints, for example geometrical ones, without loosing optimality. There are different constraints we can add on our discrete function:

- The face–spanning tree T can be chosen to be a minimal spanning tree. This leads to a time complexity in $O(\#K \cdot \log \#K)$.
- The loop added at step two can minimize the same function, in order to have the root of the face–spanning tree at a minimal position.
- The roots of the face–spanning trees T of step one and U of step four, can also be at a minimal position.

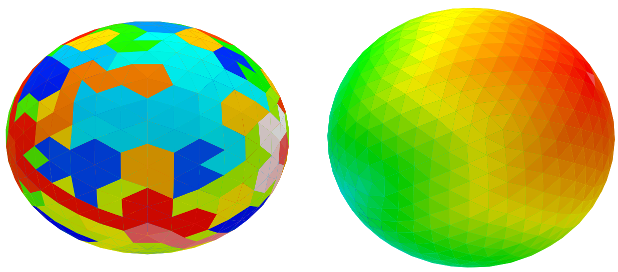

(a) Without geometrical constraint (b) With geometrical constraint

Figure 8: *Two discrete gradient vector fields on a 2–sphere, both with 2 critical cells.*

The way we include geometric constraint does not change the optimality of the resulting function. For example, Figure 8 shows two discrete gradient vector field on a 2–sphere, both with 2 critical cells. This can be used to build a discrete Morse function with localized critical points, as on Figure 9. Moreover, Forman proved in [11] that for a given cell complex K and a discrete Morse function f defined on it, the cell complex L built out of the only critical cells of f is homotopy equivalent to K. This corresponds in the smooth case to the handlebody decomposition [16].

6 Future works

We introduced a scheme for constructing discrete Morse function on finite cell complexes of dimension 2. This construction is linear in time in all cases, and is proven to be optimal in the case of 2-manifolds.

This algorithm has been extended for arbitrary finite dimensions in [15], without proof of optimality. However, the experimental results showed our algorithm gave an optimal result in most of the cases. This opens the question of which conditions on the cell complex would ensure the optimality of the resulting function.

Forman's theory seems to have still a huge potential of applications, particularly in computer graphics. More specifically, we are interested in solid mesh compression.

The corresponding work was published in Computational Geometry: Theory and Applications, volume 26, numer 3, pp 221–233. Elsevier, 2003.

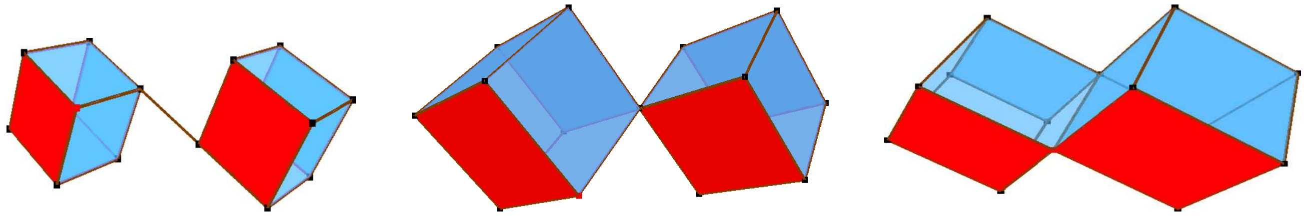

(a) dangling edge. (b) singular vertex. (c) non–regular edge.

Figure 7: *3 non–manifold cases: 1 critical vertex and 2 critical faces (squares).*

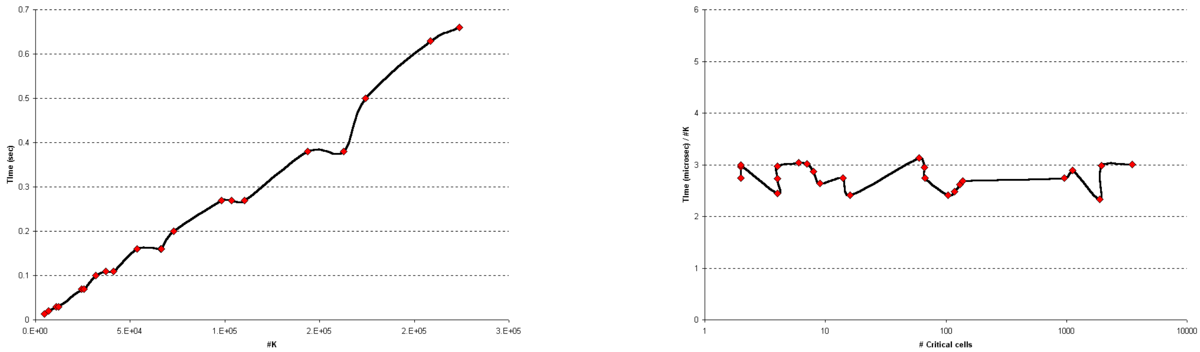

(a) Execution time versus the number of cells of the cell complex K: the complexity is linear.

(b) Ratio execution time / size of K versus the number of critical cells of K: independence.

Figure 10: *The critical steps of the extension of a torus.*

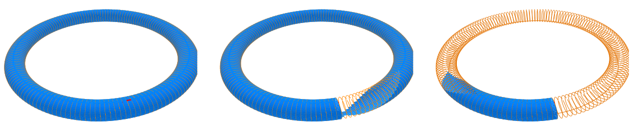

(a) Critical face (b) Collapses (c) More collapses

(d) Second critical edge (e) First critical edge (f) Critical vertex

Figure 9: *Steps of the retraction of a torus with geometrical constraint, with 1 critical vertex, 2 critical edges and 1 critical face.*

Acknowledgments

We are grateful to the referees for their patience and their contributions to this paper. We would like to thank Jean-Marie Nicolas and Anne Germa from the ENST and Marcos da Silveira from the PUC-Rio who made this cooperation possible.

References

[1] E. Babson and P. Hersh. Discrete Morse functions from lexicographic orders. *Transactions of the American Mathematical Society*, 357:509–534, 2005.

[2] M. K. Chari. On discrete Morse functions and combinatorial decompositions. *Discrete Mathematics*, 217:101–113, 2000.

[3] M. K. Chari and M. Joswig. Discrete Morse complexes. In *ACCOTA: Workshop on Combinatorial and Computational Aspects of Optimization, Topology and Algebra*, San Crist, Dec. 2002.

[4] F. H. Croom. *Basic concepts of algebraic topology*. *Undergraduate text in Mathematics*. Springer, New York, 1978.

[5] C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for Betti numbers of simplicial complexes. In *Symposium on Computational Geometry*, pages 232–239. ACM, 1993.

[6] T. K. Dey, H. Edelsbrunner and S. Guha. Computational topology. In B. Chazelle, J. E. Goodman and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary mathematics*, pages 109–143. American Mathematical Society, Providence, 1999.

[7] H. Edelsbrunner, D. Letscher and A. Zomorodian. Topological persistence and simplification. In *Symposium on Foundations of Computer Science*, pages 454–463, 2000.

[8] H. Edelsbrunner, J. L. Harer and A. Zomorodian. Hierarchical Morse complexes for piecewise linear 2–manifolds. In *Symposium on Computational Geometry*, pages 70–79. ACM, 2001.

[9] Ö. Eğecioğlu and T. F. Gonzalez. A computationally intractable problem on simplicial complexes. *Computational Geometry: Theory and Applications*, 6:85–98, 1996.

[10] R. Forman. A discrete Morse theory for cell complexes. In S. T. Yau, editor, *Geometry, Topology and Physics for Raoul Bott*. International Press, 1995.

[11] R. Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134:90–145, 1998.

[12] R. Forman. Some applications of combinatorial differential topology. In *Sullivan Fest*, 2001.

[13] R. Forman. A user guide to discrete Morse theory. In *Séminaire Lotharingien de Combinatoire*, volume 48, 2001.

[14] J. C. Hart. Morse theory for implicit surface modeling. In H.-C. Hege and K. Polthier, editors, *Visualization and Mathematics*, pages 257–268, Heidelberg, 1998. Springer.

[15] T. Lewiner, H. Lopes and G. Tavares. Visualizing Forman's discrete vector field. In H.-C. Hege and K. Polthier, editors, *Visualization and Mathematics III*, pages 95–112, Heidelberg, 2003. Springer.

[16] H. Lopes. Algorithm to build and unbuild 2 and 3 dimensional manifolds. PhD thesis, *Department of Mathematics, PUC–Rio*, 1996.

[17] H. Lopes, J. Rossignac, A. Safonova, A. Szymczak and G. Tavares. Edgebreaker: a simple compression for surfaces with handles. In C. Hoffman and W. Bronsvort, editors, *Solid Modeling and Applications*, pages 289–296, Saarbrücken, Germany, 2002. ACM.

[18] A. T. Lundell and S. Weingram. *The topology of CW–complexes*. Van Nostrand Reinhold, New York, 1969.

[19] J. W. Milnor. *Morse theory*. Number 51 in *Annals of Mathematics Study*. Princeton University Press, 1963.

[20] S. Moriyama and F. Takeuchi. Incremental construction properties in dimension two— shellability, extendable shellability and vertex decomposability. In *Canadian Conference on Computational Geometry*, pages 65–72, Fredericton, 2000.

[21] Y. Shinagawa, T. L. Kunii and Y. L. Kergosien. Surface coding based on Morse theory. *Computer Graphics and Applications*, 11:66–78, 1991.

[22] R. E. Tarjan. *Data structures and network algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, 1983.

[23] G. Vegter. Computational topology. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, pages 517–536. CRC Press, 1997.

The corresponding work was published in Computational Geometry: Theory and Applications, volume 26, numer 3, pp 221–233. Elsevier, 2003.