Dimension-Independent Simplification and Multi-Resolution Representation of Morse Complexes

Lidija Čomić*
University of Novi Sad

Leila De Floriani†
University of Genova

Federico Iuricich‡
University of Genova

Abstract—We have defined and implemented atomic and dimension-independent simplification operators on a graph-based representation of Morse complexes, and we have defined and implemented a multi-resolution model for Morse complexes built through such simplification operators.

1 Introduction

Morse theory [9] offers a natural and intuitive way of analyzing the structure of a scalar field f and of compactly representing it through decompositions of the domain M of f into meaningful regions associated with its critical points, called *Morse and Morse-Smale complexes*. A descending (ascending, resp.) Morse complex Γ_d (Γ_a) partitions manifold M into cells defined by integral lines converging to (originating at, resp.) critical points of f. For a survey of algorithms for computing Morse and Morse-Smale complexes, see [1, 4]. In our work, we use the algorithm in [10] which is based on discrete Morse theory and has been recently used in the community [5]. Our implementation starts from a decomposition of the domain of the scalar field f into a simplicial complex.

Morse and Morse-Smale complexes built on current large-size data sets can be quite large, and the noise in the data often produces over-segmentations. Simplification of these complexes can be achieved by applying the *cancellation* operator [8]. Cancellation has been used to build hierarchical models for representing Morse complexes in 2D and 3D [7].

2 Simplification Operators on Morse Complexes

In [2], we have defined atomic and dimension-independent simplification and refinement operators on Morse complexes, which form a minimally complete set of operators for creating and updating Morse and Morse-Smale complexes. The *remove* simplification operator collapses two saddles of consecutive index that are connected through a unique integral line, and such that one of them is connected to at most one saddle of the same index as the other one. Operator $remove_{i,i+1}(q,p,p')$ applies when i-saddle q is connected to exactly one other $(i+1)$-saddle p' different from p. It collapses i-saddle q and $(i+1)$-saddle p into $(i+1)$-saddle p'. In the descending complex Γ_d, it collapses i-cell q and $(i+1)$-cell p into a unique $(i+1)$-cell p' incident in q and different from p. An example of $remove_{1,2}(q,p,p')$ on a 2D descending Morse complex is illustrated in Figure 1: 1-cell q is deleted and 2-cell p is merged into 2-cell p'. Operator $remove_{i,i-1}(q,p,p')$ is completely dual.

We encode the topology of the descending and ascending Morse complexes Γ_a and Γ_d, respectively, through the *Morse Incidence Graph (MIG)* [3] (see Figure 1 for a 2D example). An *MIG* is a dimension-independent multigraph $G=(N,A)$ such that (i) there is a one-to-one correspondence between the nodes in N and the i-cells of Γ_d, (and thus the $(n-i)$-cells of Γ_a) (called *i-nodes*), and (ii) there are k arcs joining an i-node p with an $(i+1)$-node q if and only if i-cell p appears k times on the boundary of $(i+1)$-cell q in Γ_d.

*E-mail: comic@uns.ac.rs
†E-mail: deflo@disi.unige.it
‡E-mail:federico.iuricich@unige.it

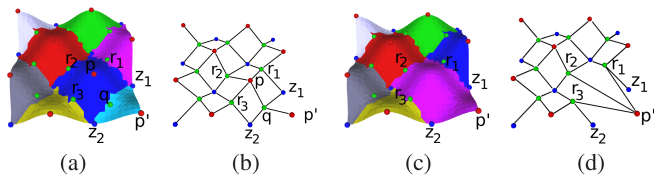

Fig. 1. The descending 2D Morse complex Γ_d for function $f(x,y)=\cos x \cos y$ (a) and the corresponding *MIG* (b). An example of $remove_{1,2}(q,p,p')$ on Γ_d (c), and on the *MIG* (d). After $remove_{1,2}(q,p,p')$, nodes p and q and the incident arcs are deleted. Arcs (p,r_j) are replaced with arcs (p',r_j), $1\le j\le 3$.

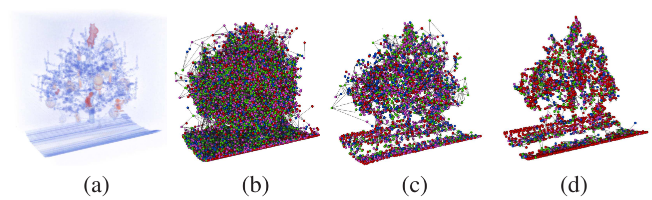

Fig. 2. The field behavior for the *xMaxTree* data set (a). The *MIG* at full resolution (b), the *MIG* after 10K (c) and after 16K simplifications (d).

A *remove* operator on an *MIG* $G=(N,A)$ eliminates the two nodes corresponding to the two critical points p and q. A $remove_{i,i+1}(q,p,p')$ is feasible on G if there is a unique arc (p,q) in A, and there is a unique $(i+1)$-node $p'\in N$ different from $(i+1)$-node p and connected to i-node q. The effect of the *remove* operator on the *MIG* is completely local, since it deletes the arcs incident in q and the arcs incident in p and in $(i+2)$-nodes in the *MIG*, and it modifies the arcs incident in p and in i-nodes in the *MIG* which become incident in p'.

The *cancellation* operator [7, 8] eliminates an i-saddle q and an $(i+1)$-saddle p that are connected through a unique integral line. There is no limit on the number of saddles connected to q and of the same index as the other saddle (i.e., p' is not unique). Let T be the set of $(i+1)$-nodes connected to q. The effect of the *cancellation* of p and q on the *MIG* consists of deleting nodes p and q and all arcs incident in either of them, and creating arcs connecting each i-node previously connected by an arc to $(i+1)$-node p with each node in T.

Thus, a *cancellation* deletes two nodes from the *MIG*, but, unlike *remove*, it may increase the number of arcs. Some strategies have been proposed to bound the number of arcs introduced by a *cancellation* [6].

Name	N Simpl	Nodes	Arcs	Cost	Time
Aneurism	-	24082	156098	1.1	-
cancellation	2000	20082	125524	1.03	762.0
remove	2000	20082	97684	0.82	204.17
Bucky	-	2645	9412	0.06	-
cancellation	350	1945	6494	0.05	0.85
remove	350	1945	3586	0.03	0.47
F16	-	32866	146549	1.04	-
cancellation	4000	24866	102847	0.9	300.08
remove	4000	24866	71646	0.6	196.82
Sf2c	-	8826	22778	0.14	-
cancellation	1500	5826	12645	0.12	3.69
remove	1500	5826	9365	0.09	3.18
Sheep	-	83089	684487	4.8	-
cancellation	10000	63089	502714	3.8	18477.3
remove	10000	63089	379491	3.1	5742.32
VisMale	-	16766	64695	0.38	-
cancellation	2000	12766	40501	0.36	36.86
remove	2000	12766	31549	0.29	31.31

Table 1. Comparison between *cancellation* and *remove*.

3 Experimental Results

We have developed a simplification algorithm for Morse complexes including operators on the dimension-independent *MIG* based on the *remove* operators. We associate a value $P(p,q)$ to any feasible remove operator, deleting two critical points p and q, as the absolute difference in function values at p and q, and we perform simplifications in order of increasing $P(p,q)$ values. We have performed experiments on the simplification of Morse complexes by using six data sets describing 2D scalar fields, and eight data sets describing 3D scalar fields on a 3.2GHz processor with 2 GBytes of memory. We have used different threshold values on $P(p,q)$: 1% of the maximum $P(p,q)$ value for light noise removal, 10% for stronger noise removal, and 20% or greater for consistently reducing the complexity of the *MIG*. The storage cost of the simplified *MIG* using these three different thresholds is equal to 95%, 65% and 35% of the cost of the *MIG* at full resolution.

In Table 1, we show the results obtained by comparing the *remove* operator with the *cancellation* operator. For each data set, we show in the first row the number of nodes and arcs in the full resolution *MIG*. In the second and third rows, we show the statistics related to *cancellation* and *remove* operators, respectively: the number of simplifications applied, the number of nodes and arcs in the simplified *MIG*, the cost of the data structure encoding the *MIG* (in Mb), and the time (in sec) needed to perform the simplifications. The number of arcs in the graph simplified with *cancellation* always exceeds the number of arcs in the graph simplified with the same number of *remove*. Such behavior influences the efficiency of the whole algorithm, doubling the time needed to manage and enqueue a larger number of arcs (and thus, a greater number of possible simplifications) for large data sets. The cost of the *MIG* is reduced by 10% $\sim$ 20% by using *remove* instead of *cancellation*.

We have applied our simplification algorithm on a 3D *xMasTree* data set (see Figure 2 (a)), which represents a tomography of a Christmas tree. The *MIG*s at full-resolution, and after 10K and 16K simplifications are shown in Figures 2 (b), (c) and (d), respectively.

4 The Multi-Resolution Morse Incidence Graph

We have defined the inverse *insert* refinement operator which undoes the effect of the *remove* simplification operator, and we have designed and implemented a multi-resolution model for the *MIG* representation of the Morse complexes, that we call the *Multi-Resolution Morse Incidence Graph (MMIG)* [3]. An *MMIG* is built from the full-resolution *MIG* by applying a sequence of *remove* operators. It consists of the base *MIG* G_B, representing the Morse complexes at the coarsest resolution, the set of *insert* refinement modifications, inverse to the *remove* applied in the simplification phase, and a dependency relation between refinement modifications, which has been shown to be a partial order relation. From an *MMIG*, it is possible to extract representations of the Morse complexes at any intermediate (uniform or variable) resolution by considering a subset of refinements that is closed with respect to the partial order.

We have implemented two data structures for storing the *MMIG*, that we call the *implicit* and the *explicit* data structures, respectively. In the *implicit* data structure, the direct dependencies of a modification μ are encoded in two arrays, *ancestors* and *descendants*, and a vector *pq-ancestors* of two bit elements of the same size as *ancestors*. *pq-ancestors*$[j]$ encodes whether a modification μ depends on node p, on node q or on both in *ancestors*$[j]$. Thus, the sets of nodes that will be connected to q or p are implicitly encoded. In the *explicit* data structure, three lists are created for each modification μ, storing the actual pointers to the above mentioned sets of nodes. The *implicit* data structure occupies 3 to 10 times less space than the *explicit* one. The extraction algorithm is from 33% to 50% faster when using the *explicit* data structure.

Acknowledgments

This work has been partially supported by the Italian Ministry of Education and Research under the PRIN 2009 program, and by the National Science Foundation under grant number IIS-1116747.

References

[1] S. Biasotti, L. D. Floriani, B. Falcidieno, P. Frosini, D. Giorgi, C. Landi, L. Papaleo, and M. Spagnuolo. Describing shapes by geometrical-topological properties of real functions. *ACM Comput. Surv.*, 40:Article 12, 2008.

[2] L. Čomić and L. De Floriani. Dimension-Independent Simplification and Refinement of Morse Complexes. *Graphical Models*, 73(5):261–285, 2011.

[3] L. Čomić, L. De Floriani, and F. Iuricich. Dimension-Independent Multi-Resolution Morse Complexes. *Computers & Graphics*, 36(5):541–547, 2012.

[4] L. Čomić, L. De Floriani, and F. Iuricich. Modeling Three-Dimensional Morse and Morse-Smale Complexes. In *Dagstuhl Seminar 11142 - Innovations for Shape Analysis: Models and Algorithms*. to appear.

[5] D. Günther, J. Reininghaus, H. Wagner, and I. Hotz. Efficient Computation of 3D Morse-Smale Complexes and Persistent Homology Using Discrete Morse Theory. *The Visual Computer*, pages 1–11, 2012.

[6] A. Gyulassy, P.-T. Bremer, B. Hamann, and V. Pascucci. Practical Considerations in Morse-Smale Complex Computation. In *Topological Methods in Data Analysis and Visualization: Theory, Algorithms, and Applications*, pages 67–78. Springer Verlag, 2011.

[7] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, and B. Hamann. A Topological Approach to Simplification of Three-Dimensional Scalar Functions. *IEEE TVCG*, 12(4):474–484, 2006.

[8] Y. Matsumoto. *An Introduction to Morse Theory*, volume 208 of *Translations of Mathematical Monographs*. AMS, 2002.

[9] J. Milnor. *Morse Theory*. Princeton University Press, New Jersey, 1963.

[10] V. Robins, P. J. Wood, and A. P. Sheppard. Theory and Algorithms for Constructing Discrete Morse Complexes from Grayscale Digital Images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(8):1646–1658, 2011.