Computing discrete Morse complexes from simplicial complexes

Ulderico Fugacci^a, Federico Iuricich^b, Leila De Floriani^c

^aGraz University of Technology, Graz, Austria ^bClemson University, Clemson, SC, USA ^cUniversity of Maryland, College Park, MD, USA

Abstract

We consider the problem of efficiently computing a discrete Morse complex on simplicial complexes of arbitrary dimension and very large size. Based on a common graph-based formalism, we analyze existing data structures for simplicial complexes, and we define an efficient encoding for the discrete Morse gradient on the most compact of such representations. We theoretically compare methods based on reductions and coreductions for computing a discrete Morse gradient, proving that the combination of reductions and coreductions produces new mutually equivalent approaches. We design and implement a new algorithm for computing a discrete Morse complex on simplicial complexes. We show that our approach scales very well with the size and the dimension of the simplicial complex also through comparisons with the only existing public-domain algorithm for discrete Morse complex computation. We discuss applications to the computation of multi-parameter persistent homology and of extrema graphs for visualization of time-varying 3D scalar fields.

Keywords: Shape analysis, topological data analysis, discrete Morse theory, homology, persistent homology, shape understanding, scientific data visualization.

1 1. Introduction

In recent years, computational topology has 2 become a fundamental tool for the analysis and 3 visualization of scientific data. In particular, 4 the efficient development of software tools for 5 extracting topological features from data has 6 led to an increasing number of applications of 7 topology-based approaches in shape analysis and 8 understanding, and in particular in the analysis of 9 sensor [1] and social [2] networks, in chemistry [3], 10 in astrophysics [4], in medicine [5]. Several math-11 ematical tools have been studied for computing a 12 compact, topologically-equivalent object starting 13 from a simplicial complex of large size. Examples 14 of these tools are the discrete Morse complex [6], 15 16 the size graph [7], and the tidy set [8].

17

¹⁸ Discrete Morse Theory (DMT) [6] is a power-¹⁹ ful theory defined in a completely combinatorial ²⁰ setting, that aims at the construction of a discrete representation of a given simplicial complex, 21 based on a discrete Morse gradient (also called 22 Forman gradient or discrete gradient field) from 23 which a homology-equivalent chain complex, the 24 discrete Morse complex is built. The Forman gra-25 dient and the associated discrete Morse complex 26 have been used both for the analysis and visualiza-27 tion of scalar fields [9], and for computing standard 28 and persistent homology [10, 11, 12]. 29

In very recent research areas, like the analysis of higher dimensional scalar fields [13] or in the analysis of shapes based on multi-parameter persistent homology [14, 15], there is a need for efficient methods capable of encoding a Forman gradient on higher dimensional simplicial complexes.

In this work, we introduce the first complete study for implementing a Forman gradient on high dimensional simplicial complexes. We start from a theoretical evaluation of the various methods used for building a Forman gradient, which are generally called reduction-based or coreduction-based. 41

30

31

32

33

34

⁴² We describe a third method initially formulated in

⁴³ [16], obtained by interleaving reductions and core-

44 ductions, and we prove the equivalence of all three

techniques. This equivalence will provide us the
freedom to implement the method that best fits
any given data structure.

To this aim, we undertake a theoretical and ex-48 perimental evaluation of the three most common 49 data structures for encoding simplicial complexes. 50 Here, we focused on data structures with avail-51 able public-domain implementations. Our exper-52 iments clearly show that the Generalized Indexed 53 data structure with Adjacencies (IA^*) [17], a data 54 structure encoding only the vertices and a subset 55 of the simplices of the complex, is the only one that 56 can suitably scale to higher dimensions without be-57 ing affected by the exponential growth in the num-58 ber of simplices. We propose a solution to com-59 pactly encode a Forman gradient attached to the 60

 $_{61}$ IA^{*} data structure.

Based on the latter encoding, we have de-62 fined and implemented an efficient, dimension-63 independent, algorithm for computing a Forman 64 gradient and for retrieving the discrete Morse com-65 plex defined by it, which is fundamental for com-66 puting, among others, homology and persistent ho-67 mology. We compare our approach to the one 68 developed in [11] and implemented in the soft-69 ware library Perseus [18] which computes a discrete 70 Morse complex using a data structure implement-71 ing the Hasse diagram of the complex, the Inci-72 dence Graph. Our experiments show that our ap-73 proach is more efficient and it is also easy to paral-74 lelize. 75

The remainder of the paper is organized as fol-76 lows. Section 2 introduces some preliminary no-77 tions about simplicial complexes, simplicial and 78 persistent homology, and discrete Morse theory. 79 Section 3 reviews some classical topological data 80 structures for simplicial complexes as well as al-81 gorithms for computing a Forman gradient and a 82 discrete Morse complex. In Section 4, we intro-83 duce, evaluate and compare the data structures for 84 compactly encoding a simplicial complex and we 85 discuss a new compact encoding for the Forman 86 gradient. In Section 5, we present the reduction 87 and coreduction-based algorithms used for comput-88 ing a Forman gradient. Section 6 is devoted to the 89 formal proof of the theoretical equivalence of these 90 approaches, while, in Section 7, we introduce a new 91

approach based on the interleaving of the two. In 92 Section 8, we describe a coreduction-based algo-93 rithm for building a discrete Morse complex based 94 on the IA^* data structure and on the compact rep-95 resentation of the Forman gradient. In Section 9, 96 we evaluate the performances of our algorithm on 97 a variety of input complexes. Finally, in Section 98 10, we draw some concluding remarks and discuss 99 applications of our approach to single and multi-100 parameter persistent homology computation and to 101 the analysis and visualization of time-varying 3D 102 scalar fields. 103

2. Background

In this section, we introduce some notions which are at the basis of our work. We briefly define and discuss simplicial complexes, simplicial homology and persistent simplicial homology, as well as discrete Morse theory.

104

110

120

121

2.1. Simplicial complexes

A k-simplex σ is the convex hull of k+1 affinely 111 independent points in the Euclidean space. For in-112 stance, a 0-simplex is a single point, a 1-simplex 113 an edge, a 2-simplex a triangle, and a 3-simplex a 114 tetrahedron. Given a k-simplex σ , the dimension 115 of σ is defined to be k, and denoted as $dim(\sigma)$. Any 116 simplex σ' , which is the convex hull of a non-empty 117 subset of the points generating σ , is called a *face* 118 of σ . Conversely, σ is called a *coface* of σ' . 119

A simplicial complex Σ is a finite set of simplices such that:

- each face of a simplex in Σ belongs to Σ ; 122
- each non-empty intersection of any two simplices in Σ is a face of both. 124

We define the *dimension* of a simplicial complex 125 Σ , denoted as $dim(\Sigma)$, as the largest dimension of 126 its simplices. Given a simplex σ of Σ , we define 127 the star of σ as the set of the cofaces of σ in Σ . A 128 simplex σ is called a *top simplex* if its star consists 129 only of σ itself. Given a simplex σ face/coface of 130 σ', σ and σ' are said to be *incident*. For k > 0, 131 two k-simplices in Σ are said to be *adjacent* if 132 they share a face of dimension k - 1, while two 133 0-simplices u and v in Σ are called *adjacent* if they 134 are both faces of the same 1-simplex. 135 136 137

Queries on a simplicial complex are often expressed in terms of the topological relations defined
by the adjacencies and incidences among its simplices.

142 • Boundary relations: given a q-simplex τ and 143 a k-simplex σ with q > k, we say that σ is in 144 boundary (q, k)-relation with τ if σ is a face of 145 τ . We denote as $bd_{q,k}(\tau)$ the set of simplices 146 in boundary (q, k)-relation with τ .

- Coboundary relations: given a q-simplex τ and a k-simplex σ with q > k, we say that τ is in coboundary (k,q)-relation with σ if τ is a coface of σ . We denote as $cbd_{k,q}(\sigma)$ the set of simplices in coboundary (k,q)-relation with σ .
- Adjacency relations: given two k-simplices σ and σ' , we say that σ is in adjacency (k, k)relation with σ' if σ is adjacent to σ' . We denote as $adj_{k,k}(\sigma)$ (or, simply $adj(\sigma)$) the set of simplices in adjacency (k, k)-relation with σ .

In the following, we will call *immediate* bound-157 ary and coboundary relations those boundary and 158 coboundary relations involving simplices of con-159 secutive dimensions. In the following, we will of-160 ten refer to them as $bd(\cdot)$ and $cbd(\cdot)$ or, when 161 we need to explicit the complex Σ with respect 162 to these relations are considered, as $bd_{\Sigma}(\cdot)$ and 163 Figure 1 illustrates the topological re $cbd_{\Sigma}(\cdot).$ 164 lations of a 1-simplex (edge) σ_0 in a simplicial 165 complex Σ . Simplices in the immediate bound-166 ary, immediate coboundary and adjacency relations 167 are depicted in blue, red, and green, respectively. 168 Specifically, $bd_{1,0}(\sigma_0) = \{v_1, v_3\}, cbd_{1,2}(\sigma_0) = \{\tau\},\$ 169 and $adj_{1,1}(\sigma_0) = \{\sigma_1, \sigma_2, \sigma_3, \sigma_4\}.$ 170

Simplicial complexes are a subclass of the more 171 general class of cell complexes [19]. They are ex-172 tensively used because of their combinatorial prop-173 erties and of the possibility of representing collec-174 tions of unorganized sets of points, usually called 175 point clouds. Alpha-shapes [20], Delaunay trian-176 gulations [21], Cech complexes [22], Vietoris-Rips 177 complexes [23], witness complexes [24, 25, 26] and 178 graph-induced complexes [27] are different ways for 179 endowing a point cloud with a simplicial structure. 180 *Cech complexes* are the most classical way to build 181 a simplicial complex starting from a point cloud, 182



Figure 1: Topological relations of edge σ_0 . Immediate boundary relation $bd_{1,0}(\sigma_0)$ consists of the two blue vertices v_1, v_3 . Immediate coboundary relation $cbd_{1,2}(\sigma_0)$ consists of the red triangle τ . Adjacency relation $adj_{1,1}(\sigma_0)$ consists of the four green edges $\sigma_1, \sigma_2, \sigma_3, \sigma_4$.

but their construction requires exponential time in the number of the input points [23].

183

184

205

Vietoris-Rips (VR) complexes [23] represent a 185 compromise between Čech complexes and the ap-186 proximations based on subsampling adopted by 187 witness [24, 25, 26] and graph-induced [27] com-188 plexes. Let G = (N, A) be a graph, a *clique* in G 189 is defined as a complete subgraph of G. The flag 190 complex of G, denoted as Flaq(G), is the simplicial 191 complex whose simplices correspond to the cliques 192 of G. Given a finite set of points P in a metric 193 space (such as the Euclidean space) and a positive 194 real number ϵ , the Vietoris-Rips (VR) complex is 195 the flag complex of the graph whose set of nodes 196 coincides with P and having an arc for each pair 197 of points in P whose distance is at most ϵ . Figure 198 2(a) shows, for each point of a set P, the neighbor-199 ing points at a distance less or equal to ϵ . Figure 200 2(b) shows the edges connecting points in P whose 201 mutual distance is less or equal ϵ . Figure 2(c) shows 202 the cliques computed on graph G and the resulting 203 VR complex Σ . 204

2.2. Simplicial and persistent homology

Simplicial homology provides invariants for shape description and characterization. Given a simplicial complex Σ , we define the *chain complex* associated with Σ as the pair $C_*(\Sigma) := (C_k(\Sigma), \partial_k)_{k \in \mathbb{Z}}$, where:

- $C_k(\Sigma)$ is the free Abelian group whose elements, called *k*-chains, are linear combinations with integer coefficients of the *k*-simplices of Σ ; 213
- $\partial_k : C_k(\Sigma) \to C_{k-1}(\Sigma)$ is the homomorphism 214 encoding the boundary relations between the 215 *k*-simplices and those (k-1)-simplices of Σ 216 such that $\partial^2 = 0$. 217



Figure 2: Construction of a VR complex Σ : given a finite set of points P, the disks of radius ϵ are computed (a); an edge is created for each pair of points at distance less than ϵ (b); the VR complex is retrieved by adding a simplex for each clique of the obtained graph (c).

Given $C_*(\Sigma)$, we denote as $Z_k(\Sigma) := \ker \partial_k$ 218 the group of the k-cycles of Σ , and as 219 $B_k(\Sigma) := \operatorname{Im} \partial_{k+1}$ the group of the k-boundaries 220 of Σ . The k^{th} homology group of Σ is defined as 221 $H_k(\Sigma) := H_k(C_*(\Sigma)) = Z_k(\Sigma)/B_k(\Sigma)$. Intuitively, 222 homology groups reveal the presence of "holes" 223 in a simplicial complex Σ . The non-null elements 224 of each homology group are cycles, which do 225 not represent the boundary of any collection of 226 simplices of Σ . The rank β_k of the k^{th} homology 227 group of a simplicial complex Σ is called the k^{th} 228 Betti number of Σ . In particular, β_0 counts the 229 number of connected components of Σ , β_1 its 230 tunnels and holes, and β_2 the shells surrounding 231 voids or cavities. 232

233

Persistent homology [28, 29, 30] aims at overcom-234 ing intrinsic limitations of standard homology by 235 allowing for a multi-scale approach defined through 236 a filtration. Let Σ be a simplicial complex, a *filtra*-237 tion F of Σ is a finite sequence of subcomplexes 238 $\{\Sigma^m | 0 \leq m \leq M\}$ of Σ such that $\emptyset = \Sigma^0 \subseteq \Sigma^1 \subseteq \cdots \subseteq \Sigma^M = \Sigma$. The *p*-persistent k^{th} ho-239 240 mology group $H_k^p(\Sigma^m)$ of Σ^m consists of the k-241 cycles included from $C_k(\Sigma^m)$ into $C_k(\Sigma^{m+p})$ mod-242 ulo boundaries. 243

Figure 3 shows an example of a filtration of a sim-244 plicial complex Σ . Persistent homology detects the 245 changes in the homology of Σ and it allows distin-246 guishing between relevant homology classes, such 247 as the 1-cycle in Σ^1 which is born at step (a) and 248 persists until the end of the filtration, and negligi-249 ble homology classes like, for instance, the 1-cycle 250 which is born at step (b) and immediately dies at 251



Figure 3: A filtration of a simplicial complex Σ . In (a), Σ^1 consists of two different connected components and one non-boundary 1-cycle (a); in (b), Σ^2 gains a non-boundary 1-cycle while it becomes connected; finally, in (c), the 1cycle created at step (b) becomes the boundary of the unique triangle in Σ^3 and its contribution in homology vanishes.

2.3. Discrete Morse theory

Discrete Morse theory due to Forman [6, 31] provides a powerful tool for analyzing the topology of an object. It has been defined for cell complexes but, for the sake of simplicity, we will review discrete Morse theory in the context of simplicial complexes. 259

252

253

260

261

262

A simplicial complex Σ is endowed with a function $f: \Sigma \to \mathbb{R}$, called a *discrete Morse function* if, for every simplex σ in Σ ,

•
$$c^+(\sigma) := \#\{\tau \in cbd(\sigma) \mid f(\tau) \le f(\sigma)\} \le 1,$$
 263

•
$$c^{-}(\sigma) := \#\{\rho \in bd(\sigma) \mid f(\rho) \ge f(\sigma)\} \le 1.$$
 264

It is easy to show (see [6], Lemma 2.5) that, for discrete Morse function, $c^+(\sigma)$ and $c^-(\sigma)$ cannot be simultaneously equal to 1. A k-simplex σ in Σ corrected by the simultaneously equal to 1.



Figure 4: (a) A discrete Morse function on a simplicial complex and (b) the corresponding Forman gradient (red simplices are critical simplices).

is called *critical simplex of index k* (or, k-saddle) 268 if $c^+(\sigma) = c^-(\sigma) = 0$. A critical simplex of index 269 0 is called a *minimum*, while a critical simplex of 270 index $d = dim(\Sigma)$ a maximum. Figure 4(a) shows 271 a discrete Morse function f defined on a simplicial 272 complex. Each simplex is labeled with the corre-273 sponding value of function f. Vertex 1 is critical 274 (minimum), since f has a higher value on all edges 275 incident to it. Edge 5 is critical (saddle), since f276 has a higher value on the incident triangle 7, and 277 lower values on its vertices. 278

A discrete vector field V on Σ is a collection of 279 pairs of simplices $(\sigma, \tau) \in \Sigma \times \Sigma$ such that $\sigma \in bd(\tau)$ 280 and each simplex of Σ is in at most one pair of 281 V. A discrete Morse function $f: \Sigma \to \mathbb{R}$ induces 282 a discrete vector field $V = \{(\sigma, \tau) \in \Sigma \times \Sigma \mid \sigma \in$ 283 $bd(\tau)$ and $f(\sigma) \geq f(\tau)$, called a Forman gradi-284 ent (or, equivalently, gradient vector field) of f on 285 Σ . A pair $(\sigma, \tau) \in V$ can be depicted as an arrow 286 from σ to τ . Given a discrete vector field V, a V-287 path (or, equivalently, a gradient path) is a sequence 288 $[(\sigma_1, \tau_1), (\sigma_2, \tau_2), \dots, (\sigma_r, \tau_r)]$ of pairs of k-simplices 289 σ_i and (k+1)-simplices τ_i , such that $(\sigma_i, \tau_i) \in V$, 290 σ_{i+1} is a face of τ_i , and $\sigma_i \neq \sigma_{i+1}$. A V-path is a 291 closed path if σ_1 is a face of τ_r different from σ_r . 292 It has been proven that a discrete vector field V is 293 the Forman gradient of a discrete Morse function if 294 and only if V is free of closed paths [6]. 295

Given a Forman gradient V on a simplicial complex Σ , the *discrete Morse complex* associated with Σ is a chain complex $\mathcal{M}_* := (\mathcal{M}_k, \tilde{\partial}_k)_{k \in \mathbb{Z}}$, where:

• groups \mathcal{M}_k are generated by the critical ksimplices;

• the boundary maps $\tilde{\partial}_k$ are obtained by following the gradient paths of V (see Subsection 8.2 for a detailed description).



Figure 5: (a) A Forman gradient computed on a simplicial complex and (b) the graph structure formed by the gradient paths (boundary maps $\tilde{\partial}_k$) connecting the critical simplices (groups \mathcal{M}_k).

A discrete Morse complex \mathcal{M}_* , associated with 304 a simplicial complex Σ , provides a homologically 305 equivalent representation of Σ (see [6], Theorem. 306 8.2). If we consider a simplicial complex Σ , and 307 we compute a Forman gradient V on it (see Figure 308 5(a), we obtain a discrete Morse complex \mathcal{M}_* hav-300 ing its cells in one-to-one correspondence with the 310 critical simplices of V. \mathcal{M}_* can be described as a 311 graph having nodes in correspondence of the criti-312 cal simplices of V, and having the arcs in one-to-one 313 correspondence with the gradient paths connecting 314 such nodes (see Figure 5(b)). 315

Since Σ and \mathcal{M}_* are homologically equivalent, 316 computing the homology on \mathcal{M}_* is preferable due 317 to the fact that the cells in \mathcal{M}_* are generally fewer 318 than the simplices in Σ . As shown in [32], the ho-319 mological equivalence between a simplicial complex 320 Σ and a discrete Morse complex \mathcal{M}_* associated 321 with Σ can be generalized to persistent homology 322 by requiring that the Forman gradient V is filtered 323 with respect to the filtration F considered. For-324 mally, given a filtration $F = \{\Sigma^m \mid 0 \le m \le M\}$ 325 of a simplicial complex Σ , a Forman gradient V 326 of Σ is *filtered* with respect to F if, for each pair 327 $(\sigma, \tau) \in V$, there exists $m \in \{1, \ldots, M\}$ such that $\sigma, \tau \in \Sigma^m$ and $\sigma, \tau \notin \Sigma^{m-1}$. 328 329

3. Related work

In this section, we review the state-of-the-art on data structures for encoding simplicial complexes and on algorithms for computing a discrete Morse complex. 334

335 3.1. Topological data structures for simplicial com 336 plexes

Several topological data structures for encoding a simplicial complex have been proposed in the literature, mainly for simplicial complexes in low dimensions, and focusing on triangle and tetrahedral meshes (see [33] for a survey). We consider here data structures specific for simplicial complexes in arbitrary dimensions.

The most general dimension-independent data 344 structure for cell and simplicial complexes is the 345 Incidence Graph. An Incidence Graph (IG) [34] 346 is a topological incidence-based representation of a 347 simplicial complex which encodes all the simplices 348 as nodes of a graph and their immediate boundary 349 and coboundary relations as its arcs. The Simpli-350 fied Incidence Graph (SIG) [35] and the Incidence 351 Simplicial (IS) data structures [36] are simplified 352 representations of the IG. A comparison among 353 IG, SIG and IS is presented in [37], while an im-354 plementation of all these data structures is included 355 in the Mangrove Topological Data Structure library 356 available in the public domain [38]. 357

In the case of triangle and tetrahedral meshes, 358 adjacency-based data structures are the most 359 widely used thanks to their compactness and ef-360 ficiency. The Generalized Indexed data structure 361 with Adjacencies (IA^*) [17] generalizes such repre-362 sentations and is capable of encoding non-manifold 363 simplicial complexes of any dimension. Recently, a 364 topological data structure has been proposed for 365 simplicial complexes embedded in the Euclidean 366 space in [39], where topological relations can be ef-367 ficiently extracted in parallel on different portions 368 of the domain. 369

In recent years, new data structures have been 370 developed well suited to perform specific tasks. The 371 Simplex Tree (ST) [40] has been defined to effi-372 ciently extract boundary relations for computing 373 persistent homology. The Simplex Tree encodes all 374 simplices in the complex and tends to be more ver-375 bose than the IA^* data structure, as shown in [39]. 376 An implementation of the ST is available in the 377 Gudhi public domain library [41]. The Maximal 378 Simplex Tree (MST) and the Simplex Array List 379 (SAL) [42] are optimized versions of the ST. To 380 the extent of our knowledge, there are no imple-381 mentations of these data structures. The skeleton 382 blocker data structure [43] has been created specif-383 ically to perform edge contraction on a simplicial 384

complex, but it can be efficiently initialized only when working with flag complexes. An implementation of the latter is provided in the *Gudhi* library [41].

389

3.2. Computing a discrete Morse complex

The process of building a discrete Morse com-390 plex from a simplicial complex typically consists 391 of two steps: (i) computing the Forman gradient 392 and identifying the critical simplices, and (ii) ex-393 tracting the boundary maps. We can classify al-394 gorithms for computing a Forman gradient as: un-395 constrained [44, 45, 46, 47, 11, 48] and constrained 396 algorithms [49, 50, 51, 10, 52, 53, 54, 55]. Uncon-397 strained algorithms compute a Forman gradient on 398 a cell/simplicial complex when no scalar value is 399 provided. The aim is to create a homologically 400 equivalent representation of the input complex hav-401 ing as few critical cells as possible. Constrained 402 algorithms start from a cell/simplicial complex en-403 dowed with a scalar function F_0 defined on its ver-404 tices, and aim at constructing a Forman gradient 405 that best fits F_0 [50, 51, 10, 52]. The discrete Morse 406 complex is used in the analysis and visualization of 407 scalar fields as a compact representation of the field 408 behavior. The aim is to obtain a decomposition of 409 the dataset in regions of influence for each critical 410 simplex. Ascending and descending traversal tech-411 niques [55, 56, 57] for the V-paths of the Forman 412 gradient have been developed for reconstructing the 413 ascending and descending Morse cells, respectively. 414 We refer to [9] for an in-depth description of these 415 methods. When computing persistent homology on 416 a simplicial complex Σ [58, 59, 60], the aim is to ob-417 tain a complex which is a compact version of Σ and 418 has the same persistent homology [10, 32, 61]. To 419 this aim, the gradient V-paths need to be visited by 420 starting from the critical simplices and by travers-421 ing the paths in a descending manner. A detailed 422 description of this process is provided in Subsection 423 8.2. 424

4. Encoding a simplicial complex endowed 425 with a Forman gradient 426

In this section, we consider the problem of encoding a simplicial complex endowed with a discrete vector field in a compact way. We start with

an analysis of existing data structures for simpli-430 cial complexes. Then, driven by the need to iden-431 tify the most efficient data structure to adopt in 432 our algorithm, we perform an experimental com-433 parison among them. Finally, we show how we can 434 encode a Forman gradient efficiently using a com-435 pact data structure which represents only vertices 436 and top simplices. This is particularly challenging 437 since a representation for a Forman gradient V on a 438 complex Σ requires encoding the pairings between 439 two simplices of consecutive dimension for all sim-440 plices in Σ . 441

442 4.1. Encoding a simplicial complex

We analyze here three data structures for en-443 coding a simplicial complex, namely the Incidence 444 Graph (IG) [34], the Simplex Tree (ST) [40, 62], 445 and the Generalized Indexed data structure with 446 Adjacencies $(IA^* \text{ data structure})$ [17]. The IG is 447 the most widely-used data structure for simplicial 448 complexes, the ST has been used in topological 449 data analysis applications, being implemented in 450 the *Gudhi* library, the IA^* data structure is a com-451 pact representation for simplicial complexes encod-452 ing only vertices and top simplices. Implementa-453 tions in the public domain exist for all of them, and 454 on such implementations we base our experimental 455 comparisons. 456

⁴⁵⁷ The Incidence Graph (IG) for complex Σ de-⁴⁵⁸ scribes its Hasse diagram [63], i.e., the graphical ⁴⁵⁹ representation of the partially ordered set gener-⁴⁶⁰ ated by all the simplices of Σ and their incidence ⁴⁶¹ relations. The IG can be viewed as a directed graph ⁴⁶² $G_{IG} = (N_{IG}, B_{IG} \cup C_{IG})$ in which:

- the nodes in N_{IG} are in one-to-one correspondence with the simplices of Σ ; with abuse of notation, we will indicate with σ both a node, and its corresponding simplex;
- a directed arc in B_{IG} (boundary arc) connects two nodes (τ, σ) in N_{IG} with $dim(\tau) = dim(\sigma) + 1$ if σ is a face of τ ;
- a directed arc (coboundary arc) in C_{IG} connects two nodes (σ, τ) in N_{IG} with $dim(\tau) = dim(\sigma) + 1$ if τ is a coface of σ .

In Figure 6(b), the *IG* representing the simplicial complex Σ depicted in Figure 6(a) is shown. Nodes are colored according to the dimension of the simplex they represent. Note that, for simplicity, 476 we have shown only one undirected arc for each 477 pair of mutual incident nodes, since if a directed 478 arc exists from τ to σ in B_{IG} , arc (σ, τ) must exist 479 in C_{IG} By storing the incidence relations between 480 simplices of consecutive dimension, the IG is 481 efficient in the retrieval of topological relations, 482 but the large amount of information encoded 483 makes it unsuitable for complexes of large size and 484 of high dimensions [37]. 485

The Simplex Tree (ST) [40] encodes also all the 487 simplices of a simplicial complex Σ as the IG, but 488 only a subset of the incidence relations encoded in 489 the IG. The ST is based on a total order selected 490 on the vertices of Σ . Let I(v) be the position in the 491 total order of a vertex $v \in \Sigma$. Given a k-simplex 492 $\sigma = \{v_0, ..., v_k\}$ in Σ , $max_v(\sigma) = max(I(v_i))$ is the 493 latest vertex of σ in the total order. The ST can 494 be viewed as a graph $G_{ST} = (N_{ST}, A_{ST})$ in which: 495

486

- the nodes in N_{ST} are in one-to-one correspondence with the simplices of Σ , and a node $\sigma \in N_{ST}$ is labeled with $I(max_v(\sigma))$; 498
- a directed arc $(\sigma, \tau) \in A_{ST}$ connects two nodes in N_{ST} , if σ is in the immediate boundary of τ , and $I(max_v(\tau)) > I(max_v(\sigma))$.

Nodes corresponding to the vertices of Σ are con-502 nected to the root of the Simplex Tree. If we select 503 a path from the root to a node $\sigma = \{v_0, ..., v_k\},\$ 504 we have that: (i) labels $\{l_0, ..., l_k\}$ are encountered 505 sorted by increasing order along the path and each 506 label appears exactly once; (ii) each label corre-507 sponds to a vertex of σ , more precisely $l_i = I(v_i)$, 508 for each i = 0, ..., k. 509

In Figure 6(c), we show the Simplex Tree rep-510 resentation of the simplicial complex depicted in 511 Figure 6(a). The order of the vertices is indicated 512 by the numbers depicted in blue, while the remain-513 ing numbers indicate the labels of the nodes corre-514 sponding to the k-simplices, with k > 0. For the 515 sake of clarity, we are not showing the connections 516 between the vertices and the root. Note that the 517 Simplex Tree is order dependent, in the sense that 518 we can have different STs for the same complex. 519 For example, Figure 6(d) shows the ST obtained 520 for Σ by using a different order for its vertices. 521

From the two graph representations, we see that $_{522}$ $N_{IG} = N_{ST}$ and that $A_{ST} \subset A_{CB}$, since it contains $_{523}$



Figure 6: A simplicial complex Σ (a) and its representation through an Incidence Graph (b) and through a Simplex Tree (c); a Simplex Tree using a different ordering for the vertices (d). Blue dots are associated with the vertices of Σ , green dots with its edges and red dots with its triangles.

all those arcs $(\sigma, \tau) \in A_{CB}$ for which $I(max_v(\tau)) >$ 524 $I(max_v(\sigma))$. The Simplex Tree has been designed 525 with the task of efficiently performing only bound-526 ary queries. In order to be able to perform also 527 coboundary queries, an extended version of the 528 Simplex Tree has been proposed in [40]. This ex-529 tended version contains a circular list linking all the 530 nodes having the same label and the same dimen-531 sion and an arc from a node to its parent. This ver-532 sion is not implemented in the Simplex Tree in the 533 public domain library Gudhi [41]. In [42], two com-534 pressed optimization of the latter have been pre-535 sented, namely the Maximal Simplex Tree and the 536 Simplex Array List, sharing the same functionali-537 ties but reducing the number of nodes encoded. To 538 the best of our knowledge, no implementations are 539 provided for these latter. 540

As mentioned before, more compact representa-541 tions for a simplicial complex can be obtained by 542 encoding only the vertices and top simplices. To 543 be able to extract boundary, coboundary and ad-544 jacency relations efficiently, the simplest represen-545 tation would encode: (i) for each top k-simplex σ , 546 its boundary defined by the references to its k+1547 vertices, and its adjacencies defined by references 548 to the simplices adjacent to σ along a (k-1)-face; 549 (ii) for each vertex v, its star, defined by the the 550 list of all top simplices incident in v. It can be 551 noticed that storing the entire star of a vertex v552 is not necessary, since the star can be efficiently 553 reconstructed by navigating the top simplices inci-554 dent in v through the encoded adjacencies. This 555 constitutes the basis for the Generalized Indexed 556 data structure with Adjacencies (IA^*) [17]. 557

⁵⁵⁸ We can describe the IA^* data structure as a ⁵⁵⁹ graph $G_{IA} = (N_{IA}, A_{IA})$ in which $N_{IA} = N_0 \cup$ N_{top} , with set N_0 corresponding to the vertices of Σ , and set N_{top} corresponding to the top simplices of Σ . The set of arcs in A_{IA} is the disjoint union of three subsets $A_{(t,0)}, A_{(t,t)}, A_{(0,t)}$ defined as follows: 563

- $A_{(t,0)}$ (boundary arcs): a directed arc (σ, v) , 564 where σ is in N_{top} and v in N_0 , belongs to 565 $A_{(t,0)}$ if v is a vertex of σ ; 566
- $A_{(t,t)}$ (adjacency arcs): an undirected arc 567 (σ, τ), where σ and τ are k-simplices in N_{top} , 568 belongs to $A_{(t,t)}$ if σ and τ share a (k-1)-face; 569
- $A_{(0,t)}$ (coboundary arcs): a subset of the arcs (v, σ), where v in N_0 and σ is in N_{top} , such that v is on the coboundary of σ , as defined below. 573

Given a vertex v, we consider the subgraph $_{574}$ $G_{IA}(v) = (N_{IA}(v), A_{IA}(v))$ of G_{IA} where: $_{575}$

- $N_{IA}(v)$ consists of all nodes $\sigma \in N_{top}$ such that v is a vertex of σ ; 577
- $A_{IA}(v)$ consists of all arcs in $A_{(t,t)}$ connecting pair of nodes in $N_{IA}(v)$. 579

Thus, an oriented arc (v, σ) is encoded in $A_{(0,t)}$ 580 for each connected component in $G_{IA}(v)$, where 581 σ is any top simplex in $N_{IA}(v)$ belonging to such 582 component. 583

584

In Figure 7, we show the nodes and the arcs encoded in the IA^* data structure (see Figure 7(b)) for the simplicial complex in Figure 7(a). Blue nodes denote vertices, while green and red nodes denote top edge and triangles, respectively. Undirected arcs represent adjacency relations among top simplices, i.e., arcs (τ_1, τ_2) and (σ_1, σ_2) . Boundary



Figure 7: A simplicial complex Σ (a) and its representation through the IA^* data structure (b). Blue dots correspond to vertices, green dots correspond to top edges and red dots to top triangles.

arcs are denoted by arrows, while coboundary arcsby dotted arrows.

The space required by the IA^* data structure 594 depends on the structure on the complex, i.e., the 595 number of arcs in $A_{(t,t)}$ and in $A_{(0,t)}$ depends on the 596 connectivity of the top simplices. If we restrict our 597 consideration to an important subclass of simpli-598 cial complexes, that of simplicial pseudomanifolds, 599 we can get some insights for comparing the space 600 required by the IA^* data structure to that of the 601 IG. Recall that a simplicial d-pseudomanifold is 602 a (d-1)-connected simplicial d-complex such that 603 any (d-1)-simplex is on the boundary of either one 604 or two *d*-simplices. 605

If Σ is a *d*-pseudomanifold, we have that the 606 number of arcs in $A_{(t,0)}$ originating from a top k-607 simplex σ is equal to k + 1. The number of arcs 608 in $A_{(t,t)}$ originating from a top k-simplex σ is also 609 equal to k+1. Thus, $|A_{(t,t)}|$ is equal to $|A_{(t,0)}|$ and, 610 thus, the total cost of storing the boundary and the 611 adjacency arcs in the IA^* data structure is equal 612 to $2|A_{(t,0)}|$. We can observe that $c = 2|A_{(t,0)}|$ is ex-613 actly the cost of storing in the IG all the boundary 614 arcs connecting a d-simplex to a (d-1)-simplex plus 615 all the dual coboundary arcs connecting a (d-1)-616 simplex to a *d*-simplex. In the IA^* data structure, 617 besides c, we have the cost c_{st} of storing some top 618 simplices in the star of the vertices. For each vertex 619 v, c_{st} is equal to the number of connected compo-620 nents of $G_{IA}(v)$. In the worst case this might be 621 equal to the number of d-simplices having v on their 622 boundary. However, in the IG we need to take into 623 account the cost of encoding the other boundary 624 and coboundary arcs which connect k- and (k-1)-625 simplices (with k < d), which will be clearly much 626

higher than c_{st} .

4.2. Experimental evaluation

This subsection provides an experimental com-629 parison among the IG, the ST and the IA^* data 630 structure. In our experiments, we have used three 631 kinds of data sets. The first data sets are volume 632 data that have been tetrahedralized. Each ver-633 tex of the dataset has an associated scalar value. 634 The DTI-SCAN is a Diffusion Tensor MRI Scan 635 of a human brain, the VISMALE dataset is a CT-636 scan of a man's head and the ACKLEY dataset is 637 a synthetic function discretizing Ackley's function 638 [64]. The datasets in the second group are networks 639 obtained from real data on which cliques have 640 been computed. Two of these datasets (AMAZON1, 641 AMAZON2) are graphs representing the "Customers 642 Who Bought This Item Also Bought" feature of 643 the Amazon website. If a product i is frequently 644 co-purchased with product j, the graph contains a 645 directed edge from i to j (notice, we are considering 646 the graph undirected). The third graph represents 647 a road network in California where intersections 648 and endpoints are described by nodes and the roads 649 connecting these intersections or road endpoints are 650 described by undirected edges (ROADNET). The 651 datasets in the third group are point clouds ex-652 tracted from a 2-sphere on which a Vietoris-Rips 653 complex has been computed (datasets S1.0, S1.2, 654 S1.3). 655

In our comparisons, we use the Simplex Tree 656 (ST) implementation in the *Gudhi* library [41], the 657 Incidence Graph (IG) implemented in *Perseus* [18], 658 which is a public domain tool for computing the 659 discrete Morse complex, and the IA^* data struc-660 ture implemented in [65]. Table 1 summarizes the 661 characteristics of the datasets we used and their 662 storage costs using the three data structures. For 663 each dataset, we provide the dimension of the re-664 sulting simplicial complex (column d), the number 665 of its vertices (column $|\Sigma_0|$) and of its top simplices 666 (column $|\Sigma_{top}|$), the size of the complex (column 667 $|\Sigma|$), and the storage cost required by the three data 668 structures, expressed in gigabytes. 669

We can observe that the storage cost of the IG and of the ST increases based on the total GTI number of simplices. The IG implemented in the Perseus library often runs out of memory, while GTI the ST has much higher limits. The storage cost GTI of the IA^* data structure depends on the number GTI

Deteret	4			$ \Sigma $	Storage Cost		
Dataset		乙 0	$ \Delta top $	니스	IA^*	IG	ST
DTI-SCAN	3	0.9 M	5.5M	24M	0.97	11.9	2.4
VISMALE	3	4.6M	26M	118M	4.7	-	9.7
Ackley4	4	1.5M	32M	204M	6.8	-	12.8
Amazon01	6	0.2M	0.4M	2.2M	0.12	1.6	0.3
Amazon02	7	0.4M	1.0M	$18.4 \mathrm{M}$	0.28	9.8	1.5
Roadnet	3	1.9M	2.5M	4.8M	0.8	3.3	1.0
Sphere-1.0	16	100	224	0.6M	0.003	0.9	0.04
Sphere-1.2	21	100	285	26M	0.0032	-	1.5
Sphere-1.3	23	100	382	$197 \mathrm{M}$	0.0034	-	11.01

Table 1: Datasets used in the experiments and storage costs for encoding the corresponding simplicial complex with the three data structures IA^* , IG and with the ST. The storage costs are expressed in gigabytes.

of top simplices. This means that simplicial 676 complexes in low dimensions (like ROADNET or the 677 volumetric datasets) may require comparatively 678 more memory than, for example, SPHERE-1.3 679 (being a 23-simplicial complex composed by less 680 than 400 top simplices). It is clear that the IA^* 681 data structure is always more compact than the 682 ST. The ratio between the storage costs of the 683 two data structures roughly depends on the ratio 684 between the number of top simplices and the size 685 of the complex. The worst-case scenario occurs for 686 (ROADNET dataset) where the IA^* data structure 687 requires 20% less memory than the ST, while in 688 the case of SPHERE-1.3 the storage cost for the 689 IA^* data structure is negligible with respect to the 690 11 gigabytes required by the ST. 691 692

693 4.3. Encoding a Forman gradient

In this subsection, we describe how to encode a discrete gradient field, like the Forman gradient V, on the data structures encoding a simplicial complex Σ .

If we consider the Incidence Graph G represent-698 ing Σ , we see that the arcs of G describe all the pos-699 sible pairings that can be defined on Σ by consider-700 ing two simplices of consecutive dimension. A For-701 man gradient can be encoded on the IG by adding 702 one bit flag to each arc a in C_{IG} indicating whether 703 the nodes incident in a are also a valid pair in V. 704 Because of this reason, the IG has been selected 705 in the *Perseus* tool [18]. This encoding cannot be 706 extended to the Simplex Tree since this latter en-707 codes only a subset of the coboundary arcs of the 708 IG.709



Figure 8: Gradient pairs encoded in a triangle. Pair between vertex v_0 and edge (v_0, v_2) is identified by moving on the first bit reserved for the 1-simplices (3 positions). We move forward of one position for each edge preceding (v_0, v_2) on the triangle (2 positions). We do not have to move forward since v_0 has position 0 on the edge.

We describe here a new representation which al-710 lows for a compact encoding of a Forman gradi-711 ent on the IA^* data structure and, in general, for 712 any data structure which encodes vertices plus top 713 simplices. In this case, the encoding for the gradi-714 ent pairs needs to be attached to the top simplices 715 only. The representation that we have defined en-716 codes, for each top k-simplex τ , a bit-vector of length $\sum_{i=1}^{k} {k+1 \choose i+1}(i+1)$ representing all the pos-717 718 sible pairings on its boundary. The first k + 1 bits 719 encode the pairing between τ and one of its (k-1)-720 faces. Then, recursively, for each *i*-face of τ , i + 1721 bits are stored until, for each 1-face, 2 bits are en-722 coded storing the pairings with one of its vertices. 723 For example, considering a 2-simplex (triangle), 3 724 bits are reserved for encoding the pairings with the 725 boundary edges. Then, for each of them, 2 bits are 726 reserved for encoding the pairings with the bound-727 ary vertices (see Figure 8). 728

If two paired simplices ρ and σ are both on the boundary of τ , the resulting pair will be encoded in the bit-vector of τ . Let j and l (with j + 1 = l) be the dimensions of ρ and σ , respectively, we check the bit associated with the corresponding pair computing:

- the position $\sum_{i=l+1}^{k} {\binom{k+1}{i+1}(i+1)}$ of the first bit reserved for *l*-simplices in τ ; 735
- the position of σ on the boundary of τ obtained enumerating the faces of τ ; 737
- the position of the vertex in σ that is not in ρ . 739

For example, in Figure 8, we consider the pairing 740 between the 0-simplex v_0 and the 1-simplex v_0v_2 . 741 The bits reserved for the 1-simplices start at position 3. The position of v_0v_2 on the boundary of the triangle is 1, so we discard the first two bits. Vertex 744 ⁷⁴⁵ v_2 is missing in v_0 and its position is 1. Then, the ⁷⁴⁶ bit representing their pairing relation is at position ⁷⁴⁷ $3 + (2 \cdot 1) + 1$.

We have implemented a prototype of the gradi-748 ent encoding based on the dynamic_bitset provided 749 by the Boost C++ library. With such encoding, 750 we have been able to represent the gradient frame 751 representation up to 40-dimensional simplicial com-752 plexes. Using more involved libraries and architec-753 tures could overcome the current limitations, but it 754 might greatly affect computation times. 755

756 5. Reductions and coreductions for discrete 757 Morse complexes

Reduction and coreduction operators [45] are two 758 homology-preserving operators used for reducing 759 the size of a simplicial complex without affecting 760 its homology. For this reason, reduction and core-761 duction operators can be used in a preprocessing 762 approach to compute homology, or persistent 763 homology of a simplicial complex [32, 45, 46, 47]. 764 Reduction and coreduction pairs can be fruitfully 765 used also in the context of discrete Morse theory 766 in order to define a Forman gradient. In this 767 section, we present the two methods based on such 768 operators, and we propose a new strategy, while 769 providing also a theoretical comparison of all these 770 techniques. 771

772

A reduction on a simplicial complex Σ corre-773 sponds to a deformation retraction of a simplex 774 which is the face of only one other simplex in the 775 complex. The problem is that, in most situations, 776 available reductions are quickly exhausted. In or-777 der to overcome this issue, coreductions have been 778 introduced [45], where a coreduction can be viewed 779 as the dual operation with respect to a reduction. A 780 coreduction is not feasible on a simplicial complex. 781 while it is available in the context of S-complexes 782 [45]. For the sake of simplicity, we consider an S-783 complex as a simplicial complex in which some sim-784 plices may be not present even if their cofaces are 785 in the complex. For instance, all the complexes de-786 picted in Figure 9 are S-complexes. In particular, 787 the complexes obtained after performing a coreduc-788 tion operator are examples of S-complexes which 789 are not simplicial complexes. 790

⁷⁹¹ Given an S-complex Σ , a pair (σ, τ) of elements of ⁷⁹² Σ , such that the coefficient of σ in $\partial \tau$ is ± 1 , is



Figure 9: (a) Removal of the reduction/coreduction pair (σ, τ) , and (b) corresponding pairing of simplices σ and τ in the gradient.

called a reduction pair if $cbd_{\Sigma}(\sigma) = \{\tau\}$, a coreduction pair if $bd_{\Sigma}(\tau) = \{\sigma\}$.

793

794

795

796

797

798

790

800

801

802

When simplifying a simplicial complex Σ , the effect of a reduction/coreduction is that of changing the structure of Σ , by removing a pair of simplices without affecting its homology (see Figure 9(a)). When building a Forman gradient V, the same pair is not removed from Σ , but added as a pair to V (see Figure 9(b)).

A coreduction-based algorithm builds a Forman 803 gradient using coreduction pairs and free simplices 804 [11], where a *free simplex* is a simplex with an 805 empty boundary. The algorithm works on two sets 806 of simplices: the set of paired simplices V, initial-807 ized as empty, and the set of non-excised simplices 808 Σ' , initialized as Σ . While Σ' admits a coreduction 809 pair, the algorithm excises a coreduction pair (σ, τ) 810 from Σ' and adds it to V. When no more core-811 duction is feasible, a free simplex is excised from 812 the complex and labeled as critical. The algorithm 813 repeats these steps until Σ' is empty. Since no 814 simplicial complex admits a coreduction pair, any 815 coreduction-based algorithm performs as its first 816 step the excision of an arbitrary vertex v, which is a 817 free simplex by definition, and declares it as critical. 818 The removal of v turns Σ' into an S-complex and 819 unlocks the possibility of pairing through a core-820 duction any vertex u adjacent to v. 821

A reduction-based approach performs reductions and removals of top simplices [48]. We recall that a top simplex is a simplex with an empty coboundary. The algorithm works on two sets of simplices: the set of paired simplices V, initialized as empty, and the set of non-excised simplices Σ' , initialized as Σ . While the set of non-excised simplices Σ' argument and the set of non-excised simplices Σ'

admits a reduction pair, the algorithm excises a re-829 duction pair from Σ' and adds it to V. When no 830 more reduction is feasible, a top simplex is excised 831 from the complex and labeled as critical. The al-832 gorithm stops when Σ' is empty. Differently from 833 a coreduction-based algorithm, whose first step is 834 necessarily the removal of a vertex, the initial step 835 in a reduction-based approach can involve the ex-836 cision of a feasible reduction pair or the removal 837 of a top simplex. Similarly to the previous case, 838 if no reduction pair is available, the approach has 839 to label an arbitrary top simplex as critical and to 840 remove it from Σ' . After such a removal, the situ-841 ation is analogous to the starting one and, so, the 842 same strategy can be applied. 843

In order to minimize the size of the discrete 844 Morse complex, in both approaches the creation of 845 a critical simplex is performed only if no more core-846 duction, or reduction is feasible. Actually, even if 847 this condition is not satisfied, the acyclicity of the 848 gradient paths is still guaranteed. In the following, 849 we refer to this two approaches, also in the case 850 in which critical simplices can be created when it 851 is not strictly necessary, as coreduction-based algo-852 rithm and reduction-based algorithm, respectively. 853

Equivalence of reduction and coreduction sequences

In this section, we prove the equivalence between 856 the use of reduction and coreduction operators in 857 the construction of a (filtered) Forman gradient and 858 we introduce another class of methods which could 859 operate reductions and coreductions in an inter-860 leaved way. The equivalence among these three 861 methods will give us the freedom to choose the one 862 that best fits our data structure. 863

In order to better understand how the removal of a coreduction, or of a reduction pair affects the coboundary and the boundary of the simplices of a simplicial complex, we first discuss some preliminary results.

Remark 1. Let τ be a simplex and let σ be one of its faces, then there exists $\dim(\tau) - \dim(\sigma)$ faces of τ in $cbd_{\tau}(\sigma)$.

Lemma 1. In a coreduction-based algorithm, each
removal operation does not modify the coboundary
of the remaining simplices.

Proof. Let Σ be a simplicial complex on which the 875 coreduction-based algorithm is executed. Clearly, 876 the removal of a free simplex does not modify 877 the coboundary of any remaining simplex. Let us 878 consider only removals of coreduction pairs. Let 879 (σ, τ) be a feasible coreduction pair in the set 880 of non-removed simplices Σ' . The only simplices 881 whose coboundary can be modified by the core-882 duction pair are those belonging to $bd_{\Sigma'}(\tau)$ and to 883 $bd_{\Sigma'}(\sigma)$. Since, for the feasible coreduction pair 884 $(\sigma, \tau), bd_{\Sigma'}(\tau) = \{\sigma\}$, the thesis is obtained by 885 proving that, before performing the coreduction, 886 $bd_{\Sigma'}(\sigma) = \emptyset$. Suppose that there exists $\nu \in bd_{\Sigma'}(\sigma)$. 887 By Remark 1, there exists in Σ a simplex $\sigma' \neq \sigma$ 888 such that $\sigma' \in bd_{\Sigma}(\tau)$ and $\nu \in bd_{\Sigma}(\sigma')$. Since (σ, τ) 889 is a feasible coreduction pair in Σ' , simplex σ' must 890 have been already removed, i.e., $\sigma' \notin \Sigma'$. Let us 891 proceed by induction. If (σ, τ) is the first coreduc-892 tion pair performed in the coreduction-based algo-893 rithm on complex Σ , then σ' has been removed as 894 a free simplex, but, since $\nu \in bd_{\Sigma}(\sigma')$ and $\nu \in \Sigma'$, 895 this leads to a contradiction. 896

Assume now that, for any removal of a core-897 duction pair performed before (σ, τ) , the simplex 898 of smaller dimension of the pair is free. Since 899 $\nu \in bd_{\Sigma}(\sigma')$ and $\nu \in \Sigma', \sigma'$ cannot be removed 900 as a free simplex, or by a coreduction pair removal 901 of the kind (ν', σ') . So, σ' has been removed by 902 operating a coreduction pair removal of the kind 903 (σ', τ') , which leads to a contradiction of the in-904 ductive hypothesis. 905

Lemma 2. In a reduction-based algorithm, each removal operation does not modify the boundary of the remaining simplices.

Proof. Let Σ be a simplicial complex on which 909 the reduction-based algorithm is executed. Clearly, 910 the removal of a top simplex does not modify the 911 boundary of any remaining simplex. Let us con-912 sider only removals of reduction pairs. Let (σ, τ) be 913 a feasible reduction pair in the set of non-removed 914 simplices Σ' . Similarly to Lemma 1, proving that, 915 before performing the coreduction, $cbd_{\Sigma'}(\tau) = \emptyset$ 916 is sufficient. If there exists $\nu \in cbd_{\Sigma'}(\tau)$, then, 917 by Remark 1, there exist $dim(\nu) - dim(\sigma) > 2$ 918 faces of ν in $cbd_{\Sigma'}(\sigma)$. But this leads to a con-919 tradiction, because (σ, τ) is a reduction and, thus, 920 $#cbd_{\Sigma'}(\sigma) = 1.$ 921

We are now ready to formalize and to prove 922

the equivalence between the coreduction-based andreduction-based algorithms.

Proposition 1. Given a simplicial complex Σ and the Forman gradient V produced by a reductionbased algorithm, it is always possible to obtain the same Forman gradient through a coreduction-based algorithm. The reverse is also true.

⁹³⁰ *Proof.* For the sake of brevity, we only prove that the Forman gradient produced by a reduction-⁹³² based algorithm on Σ can be obtained with a ⁹³³ coreduction-based algorithm. The proof of the re-⁹³⁴ verse is entirely similar (by using Lemma 1). Let ⁹³⁵ Σ be a simplicial complex and let

$$R_1^1, R_2^1, \dots, R_{i_1}^1, R_1^2, R_2^2, \dots, R_{i_2}^2, \dots, R_1^n, R_2^n, \dots, R_{i_r}^n$$

be the ordered sequence of reduction pairs and top simplices removed during the execution of a reduction-based algorithm, where, for $1 \leq l \leq n$ and $1 \leq j \leq i_l - 1$, R_j^l represents a reduction pair and, for each $1 \leq l \leq n$, $R_{i_l}^l$ represents a top simplex.

According to the notation adopted in (1), Fig-942 ure 10(a) depicts the ordered sequence of reduction 943 pairs and top simplices removed during the exe-944 cution of a reduction-based algorithm. We want 945 to prove that, by using the same removals, it is 946 possible to obtain a sequence of coreduction pairs 947 and free simplices compatible with a coreduction-948 based algorithm producing the same Forman gra-949 dient. Figure 10(b), for example, shows a sequence 950 of coreduction pairs and free simplices compatible 951 with a coreduction-based algorithm obtained by re-952 versing the reduction-based sequence depicted in 953 Figure 10(a) and producing the same Forman gra-954 dient. 955

We consider the following sequence obtained taking sequence (1) in reverse order:

$$R_{i_n}^n, R_{i_n-1}^n, \dots, R_1^n, R_{i_{n-1}}^{n-1}, \dots, R_{i_1}^1, \dots, R_2^1, R_1^1$$

⁹⁵⁸ Consider (2) as an ordered list of removal operations performed on Σ . The following properties ⁹⁶⁰ hold:

- 1. For each $1 \le l \le n$ and $1 \le j \le i_l 1$, R_j^l is a feasible coreduction pair.
- 963 2. For each $1 \le l \le n$, $R_{i_l}^l$ is a free simplex.

To prove the two properties, we denote with:

- Σ_{j}^{l} the simplicial complex obtained in (1) after performing all the removal operations up to R_{j}^{l} 966 included; 967
- S_j^l the S-complex obtained in (2) after performing all the removal operations up to R_j^l 969 excluded. 970

We have that, for each value of l and j,

$$\Sigma_j^l \sqcup S_j^l = \Sigma \tag{3}$$

964

971

972

1. Let $R_j^l = (\sigma, \tau)$ with $1 \leq l \leq n$ and $1 \leq j \leq i_l - 1$. We have to prove that it represents a coreduction in the sequence (2), i.e., $bd_{S_j^l}(\tau) = \{\sigma\}$. By Lemma 2, in (1), τ cannot be removed before the simplices in $bd_{\Sigma}(\tau)$. So, all the simplices in $bd_{\Sigma}(\tau) \setminus \{\sigma\}$ belong to Σ_j^l . Then, by (3), $bd_{S_j^l}(\tau) = \{\sigma\}$ and, thus, (σ, τ) is a feasible coreduction in S_j^l .

2. Let $R_{i_l}^l$ be the simplex σ . We have to prove that it represents a free simplex in the sequence (2), i.e., $bd_{S_{i_l}^l}(\sigma) = \emptyset$. Analogously to 1., by Lemma 2, in (1), all the simplices belonging to $bd_{\Sigma}(\sigma)$ are in $\Sigma_{i_l}^l$. Then, by (3), $bd_{S_{i_l}^l}(\sigma) = \emptyset$ and, thus, σ is a free simplex in $S_{i_l}^l$.

Sequence (2) satisfies properties 1. and 2. So, it represents a sequence of removals compatible with a correduction-based algorithm producing on Σ the same Forman gradient of (1).

It is interesting to understand if the equivalence 990 between reduction-based and coreduction-based al-991 gorithms still holds with the further condition that 992 allows for the introduction of a critical simplex 993 only if no reduction [coreduction] pair is available. 994 Proposition 1 ensures that, given a reduction [core-995 duction] sequence produced on a simplicial complex 996 Σ by an algorithm requiring such a condition, it 997 is always possible to find a coreduction [reduction] 998 sequence inducing the Forman gradient on Σ . In 999 spite of this, Proposition 1 does not guarantee that 1000 a sequence produced by an algorithm satisfying the 1001 condition mentioned above exists. Figure 11 shows 1002 that, in general, this does not hold. The Forman 1003 gradient depicted in Figure 11 can be considered as 1004 produced by a reduction-based algorithm starting 1005



Figure 10: (a) A sequence of reduction pairs (green arrows) and top simplex removals (red simplices) produced by a reduction-based algorithm on a simplicial complex and (b) the sequence of coreduction pairs resulting in the same gradient than (a).



Figure 11: A Forman gradient on a simplicial complex that cannot be produced by a coreduction-based algorithm in which critical simplices are introduced only when no more coreduction pair is feasible.

with the removal of the top simplex τ and introduc-1006 ing critical simplices only when it is strictly neces-1007 sary. This Forman gradient cannot be produced 1008 by a coreduction-based algorithm in which critical 1009 simplices are introduced only when no more core-1010 duction pair is feasible because such an algorithm 1011 applied to this simplicial complex necessarily pro-1012 duces a Forman gradient with just one critical sim-1013 plex of dimension 0 and two critical simplices of 1014 dimension 1. 1015

1016 7. Interleaving reductions and coreductions

A new method to build a gradient field V on a 1017 simplicial complex is to execute removals of reduc-1018 tion and coreduction pairs in an interleaved way. 1019 We denote as interleaved-based algorithm an algo-1020 rithm producing a discrete vector field by using re-1021 movals of reduction and coreduction pairs, of top 1022 simplices and of free simplices. Given a simplicial 1023 complex Σ , pairs of simplices are excised from Σ by 1024 arbitrarily choosing between reduction or coreduc-1025 tion pairs. When no more pairs can be removed, 1026

a free simplex or a top simplex is excised from the 1027complex and labeled as critical. The algorithm repeats these steps until Σ is empty. 1028

Here, we prove that such an algorithm actually 1030 produces a Forman gradient and that all interleaved 1031 methods are equivalent. 1032

Proposition 2. Given a simplicial complex Σ , the 1033 discrete vector field V produced by any interleavedbased algorithm is a Forman gradient. 1036

Proof. Given two pairs (σ, τ) , (σ', τ') in V, we de-1036 fine $(\sigma, \tau) \leq (\sigma', \tau')$ if there exists a V-path start-1037 ing with (σ, τ) and ending with (σ', τ') . In order to 1038 prove the thesis, i.e., that V is free of closed V-path, 1039 it is enough to prove that \leq define a partial order 1040 on V. Consider set V as built in any intermediate 1041 step of the proposed algorithm and let (σ, τ) be the 1042 last pair inserted in V. The following properties 1043 allow to achieve the thesis: 1044

- 1. (σ, τ) is a minimal element with respect to the elements already inserted in V originating from a coreduction pair; 1047
- 2. (σ, τ) is a maximal element with respect to the elements already inserted in V originating from a reduction pair.

Suppose that condition 1 does not hold. Then, 1051 there must exist an already performed coreduction 1052 pair (σ', τ') such that $\sigma \in bd(\tau')$. This implies that, 1053 at the step in which (σ', τ') has been performed, 1054 $\sigma, \sigma' \in bd(\tau')$. But this is impossible, otherwise the 1055 coreduction pair (σ', τ') could not have been performed. 1057 Suppose that condition 2 does not hold. Then, there must exist an already performed reduction pair (σ', τ') such that $\sigma' \in bd(\tau)$ and this implies that, at the step in which (σ', τ') has been performed, $\tau, \tau' \in cbd(\sigma')$. But this is impossible, otherwise the reduction pair (σ', τ') could not have been performed.

Having proven that any possible interleaved method leads to a Forman gradient, we are now interested in understanding if these different approaches could produce equivalent results or not. As an immediate consequence of Lemma 1 and Lemma 2, we can claim the following result.

Remark 2. In each interleaved-based algorithm,
each coreduction pair and free simplex removal cannot make a reduction pair feasible; each reduction
pair and top simplex removal cannot make a coreduction pair feasible.

¹⁰⁷⁶ Finally, we can prove that all interleaved meth-¹⁰⁷⁷ ods are equivalent.

Proposition 3. Given a simplicial complex Σ and the Forman gradient V on it produced by an interleaved-based algorithm, it is always possible to obtain the same Forman gradient with a reduction-based algorithm or, equivalently, with a coreduction-based algorithm.

Proof. We prove that the sequence of removals pro-1084 duced by an interleaved-based algorithm on a sim-1085 plicial complex can be also obtained with a se-1086 quence of coreduction pairs and free simplex re-1087 movals. By Remark 2, we can suitably order such a 1088 sequence, moving all the coreduction pairs and the 1089 free simplices at the beginning, thus creating a new 1090 sequence equivalent to the previous one. We apply 1091 to the last part, composed only of reduction pairs 1092 and top simplices, of this new sequence the same 1093 sorting strategy proposed in Proposition 1 to trans-1094 form a reduction-based sequence to a coreduction-1095 based sequence, and in this way, we obtain the the-1096 sis. 1097

From both an application and a theoretical point of view, it is interesting to find a method to build a Forman gradient which minimizes the number of critical simplices. It is known that, in general, this problem is NP-hard [66]. The previous results show that, from a theoretical point of view, the use of different simplification operators (such as reduction and coreduction pairs), or the combination of more than one, does not actually affect the number of resulting critical simplices.

For the sake of completeness, let us note that the 1108 results proven in this section still hold when the 1109 above-described approaches are applied to build a 1110 filtered Forman gradient. This is due to the fact 1111 that the satisfaction of the condition required to 1112 guarantee that V is a filtered Forman gradient with 1113 respect to a filtration F does not take into account 1114 if the pairs of V have been created thanks to a 1115 reduction or a coreduction operator. 1116

8. A coreduction-based algorithm for computing a discrete Morse complex

8.1. Construction of a (filtered) Forman gradient 1126

The theoretical equivalences proven in Section 6 1127 and Section 7 tell us that there is no preferable 1128 homology-preserving operator for computing a For-1129 man gradient. Here, we introduce a new dimension-1130 independent algorithm, that can also runs in par-1131 allel, which uses a representation of the simplicial 1132 complex as an IA^* data structure and the encoding 1133 of the Forman gradient discussed in Subsection 4. 1134

The basic underlying approach is the 1135 coreduction-based algorithm, introduced in 1136 [11] and implemented there only for regular grids. 1137 We summarize it for simplicial complexes. When 1138 considering simplicial complexes, the coreduction-1139 based algorithm computes a Forman gradient by 1140 using coreduction pairs starting from the simplices 1141 of lowest dimension. The set of k-simplices of 1142 complex Σ is considered by increasing values of k, 1143 starting from k = 0. As long as a coreduction pair 1144 exists between a k-simplex σ and a (k+1)-simplex 1145 τ , pair (σ, τ) is added to the Forman gradient V. 1146 When no k-simplex can be paired, one simplex is 1147 randomly chosen and declared critical. When all 1148 k-simplices have been paired or denoted as critical, 1149

the working dimension k is increased by one. Since 1150 no coreduction pair is feasible on a simplicial 1151 complex, at the first step, an arbitrary vertex v is 1152 denoted as critical in V to trigger coreductions. 1153 In [32], the coreduction-based approach is used 1154 for persistent homology computation, and thus by 1155 considering a filtration of the original complex. If 1156 each simplex is paired only with another simplex 1157 belonging to the same filtration value, the result-1158 ing discrete Morse complex will have the same 1159 persistent homology of the original complex. 1160

The dimension-independent coreduction-based 1161 algorithm proposed here, unlike previous ones, uses 1162 a local approach that allows us to work on the stars 1163 of the vertices independently, which makes it par-1164 ticularly suitable for a parallel implementation. We 1165 define an indexing on the vertices of the input sim-1166 plicial complex Σ , and we extend the indexing to all 1167 the simplices in such a way that each simplex in Σ 1168 has an index equal to the maximum of the indexes 1169 of its vertices. With such indexing, the coreduction 1170 pairs can be computed locally to the lower star of 1171 each vertex. Given a vertex v, a simplex σ belongs 1172 to the lower star of v (denoted as $St^{-}(v)$) if: (i) σ 1173 is a coface of v, and (ii) v has lowest index value 1174 among the vertices of σ . Algorithm 1 illustrates the 1175 process for computing a Forman gradient on a sim-1176 plicial complex Σ having an indexing F_0 defined on 1177 its vertices. 1178

The algorithm iterates on the vertices of Σ , ex-1179 tracting first the top simplices in the lower star of 1180 v, denoted as LT_v , which are encoded in a list. For 1181 each vertex v, the algorithm iterates on the dimen-1182 sion of the simplices in the lower star of v. The 1183 algorithm works, for each dimension, with two sets 1184 of simplices: the set of k-simplices that can be de-1185 clared critical, denoted as CR_v (row 12), and the 1186 set of (k + 1)-simplices to pair (row 14), denoted 1187 as ST_v . CR_v and ST_v have a maximum size equal 1188 to the maximum, by varying k, of the number of 1189 k-simplices in the lower star of a vertex in Σ , and 1190 they are encoded as balanced binary search trees. 1191 A candidate simplex is extracted from set ST_v (row 1192 16) and paired with its unique unpaired face (row 1193 18). Recall that a simplex τ can be paired with 1194 another simplex σ by coreduction if σ is the only 1195 unpaired face of τ . If there are no coreductions 1196 available (row 21), a new critical simplex is taken 1197 from CR_v . Every time a simplex is paired or set 1198 as critical, it is also removed from ST_v or CR_v , re-1199

Algorithm 1 - FormanGradient(Σ, F_0)

- 1: INPUT: Σ , d-dimensional simplicial complex
- 2: INPUT: F_0 , indexing of vertices of Σ
- OUTPUT: V, Forman gradient; C, set of crit-3: ical simplices
- 4: $\Sigma_0 :=$ vertices of Σ
- 5: $V := \emptyset$
- 6: $C := \emptyset$

12:

13:

14:

16:

17:

18:

19:

20:

21:

22:

23:

24:

- 7: for $v \in \Sigma_0$ do
- k := 08:
- $ST_v := \{v\}$ 9:
- $LT_v := LowerTop(v, \Sigma, F_0)$ 10:
- while $k \ll d$ do 11:

else

end if

 $CR_v := ST_v$ k := k + 1 $ST_v := LowerStar(v, \Sigma, F_0, LT_v, k)$

 $Remove(\sigma, CR_v)$

 $addCritical(\sigma, C)$

 $remove(\sigma, CR_v)$

- while $CR_v \neq \emptyset$ do 15:
 - $(\sigma, \tau) := getNextPair(v, \Sigma, ST_v, CR_v)$ if $(\sigma, \tau) \neq \emptyset$ then $addPair(\sigma, \tau, V)$ $Remove(\tau, ST_v)$

 $\sigma = getFirstCritical(CR_v)$

25: end while 26:end while 27:28: end for

spectively. When set CR_v is empty, the working 1200 dimension is increased. The algorithm terminates 1201 when all the simplices in the lower star of each ver-1202 tex v have been paired, or set as critical. 1203

The procedures and the functions, on which Al-1204 gorithm 1 is based, are:

- Function $LowerTop(v, \Sigma, F_0)$: computes all 1206 the top simplices of Σ belonging to the lower 1207 star of v and encodes such simplices in list LT_v . 1208 This is performed by navigating the star of ver-1209 tex v through the adjacency arcs in the IA^* 1210 data structure. Thus, it works in time $O(t_v)$, 1211 where t_v denotes the number of top simplices 1212 in the star of v. 1213
- Function LowerStar $(v, \Sigma, F_0, LT_v, k)$: ex-1214 tracts all the k-simplices belonging to the lower 1215

star of v from LT_v and encodes such sim-1216 plices in ST_v . This operation is performed 1217 by cycling on the elements of LT_v and collect-1218 ing the k-faces of each top simplex that are 1219 also incident in v. The extraction of the k-1220 simplices of a top simplex of dimension i is 1221 performed in $O(\binom{i+1}{k+1})$. If we denote as $t_{v,i}$ 1222 the number of top simplices of dimension i1223 incident in v, the total number N_k of simplices extracted is $N_k = \sum_{i=1}^d t_{v,i} {i+1 \choose k+1}$ in the 1224 1225 worst case, since some simplices are contained 1226 within the boundary of more than one top sim-1227 plex. Since each of such simplices is inserted in 1228 ST_v , $LowerStar(v, \Sigma, F_0, LT_v, k)$ may require 1229 $O(N_k \log N_k)$ time in the worst case. 1230

- Procedure $addPair(\sigma, \tau, V)$: adds a new pair 1231 to V. Since the gradient pairs are encoded on 1232 the top simplices only, we have to find the top 1233 simplices incident in both σ and τ . This is 1234 done by examining all the top simplices in the 1235 star of a vertex of σ and detecting all those 1236 having τ on their boundary. For each of these 1237 latter, we update the corresponding bit-vector. 1238 The operation requires $O(t_w)$, where t_w de-1239 notes the number of top simplices incident in 1240 a vertex w of σ . 1241
- Function $getNextPair(v, \Sigma, ST_v, CR_v)$: iter-1242 ates on the set of unpaired simplices ST_v se-1243 lecting the first simplex available for a core-1244 duction. For each simplex τ in ST_v , the sim-1245 plices on the boundary of τ containing v are 1246 extracted, and then, for each of such bound-1247 ary simplices σ , the membership of σ to CR_n 1248 is checked. In the worst case, we will need to 1249 check the membership of all the elements in 1250 CR_v . This leads to a worst-case complexity of 1251 $O(k|ST_v||CR_v|\log|CR_v|).$ 1252
- Function $getFirstCritical(CR_v)$: returns the first simplex in the set of candidate critical simplices CR_v . Since CR_v is implemented as a balanced binary search tree, the worst-case time complexity is $O(\log |CR_v|)$.
- Procedure $Remove(\sigma, CR_v)$: eliminates a simplex from CR_v (or ST_v). Since both CR_v and ST_v are implemented as a balanced binary search tree, the worst-case time complexity is $O(\log |CR_v|)$.

For each dimension, the computation cost is 1263 dominated by the cost of executing Function 1264 $getNextPair(ST_v, CR_v, \Sigma)$. If we denote as cr_m 1265 and as st_m the maximum size of CR_v and ST_v , re-1266 spectively over all dimensions, the time complexity 1267 for a single vertex v is $O((d-1)st_mcr_m\log(cr_m))$. 1268 Note that both cr_m and st_m can be of the order of 1269 the number of k-simplices incident in v. Since the 1270 algorithm computes the Forman gradient locally to 1271 the lower star of each vertex, the approach is easy 1272 to parallelize by running Algorithm 1 on multiple 1273 vertices at a time. Results are shown in Section 9. 1274

We prove the correctness of Algorithm 1 by showing that it is a coreduction-based algorithm ensuring that the generated discrete vector field V is a filtered Forman gradient.

Proposition 4. Let Σ be a simplicial complex, 1279 $F_0: \Sigma_0 \to \mathbb{R}$ be an injective function and F be 1280 the filtration of Σ naturally induced by F_0 . Given 1281 Σ and F_0 as input, Algorithm 1 returns a filtered 1282 Forman gradient with respect to F. 1283

Proof. Algorithm 1 processes the lower stars of the 1284 vertices of Σ independently. Without loss of gener-1285 ality, we can assume that the lower stars are pro-1286 cessed in a sequence ordered by ascending values 1287 of function F_0 . In this way, we obtain an ordered 1288 sequence of simplices added to the gradient V and 1289 to the set of critical simplices C. We prove that 1290 this sequence, denoted as S, actually represents a 1291 feasible sequence of coreduction pairs and free sim-1292 plices for Σ . Let us consider a pair of simplices 1293 (σ, τ) declared as a pair of V during the processing 1294 of the lower star $St^{-}(v)$ of v. Let σ' be a simplex 1295 in $\operatorname{bd}_{\Sigma} \tau$ different from σ . If $\sigma' \in St^{-}(v)$, then σ' 1296 has to be already added to V or to C. Otherwise, 1297 if $\sigma' \notin St^{-}(v)$, then there exists a vertex w of Σ 1298 such that $\sigma' \in St^-(w)$ and $F_0(w) < F_0(v)$. So, 1299 σ' has to be already added to V or to C during 1300 the processing of $St^{-}(w)$. In both cases, (σ, τ) can 1301 be considered as a feasible coreduction pair in the 1302 sequence S. Similarly, any simplex σ added to C 1303 during the processing of a lower star can be con-1304 sidered as a free simplex in the sequence S. So, 1305 Algorithm 1 is a coreduction-based algorithm and 1306 then, thanks to Proposition 2, it returns a Forman 1307 gradient. Moreover, since Algorithm 1 pairs only 1308 simplices belonging to the same lower star and, by 1309 the definition of F, these simplices have the same 1310 filtration value. Thus, the returned Forman gradiunit V is necessarily filtered with respect to F.

¹³¹³ 8.2. Extracting the discrete Morse complex

The discrete Morse complex \mathcal{M}_* associated with 1314 a (filtered) Forman gradient V on Σ is retrieved 1315 by navigating the paths of V. The output consists 1316 of the boundary maps $\partial_k : \mathcal{M}_k \to \mathcal{M}_{k-1}$. These 1317 latter can be seen as the arcs of a graph in which 1318 the nodes correspond to the critical simplices and 1319 each arc has a multiplicity which corresponds to a 1320 gradient path between two critical simplices. 1321

Extracting the boundary maps by visiting the 1322 paths of V may cause simplices to be visited more 1323 than once, as discussed in [10]. In the worst case, 1324 a critical k-simplex may be connected by V-paths 1325 to all the k-simplices of Σ (this set is denoted as 1326 Σ_k). Moreover, each k-simplex of this set can be 1327 visited, via multiple V-paths, more than once; in 1328 the worst case each simplex will be visited $O(|\Sigma_k|)$ 1329 times. The resulting worst-case complexity for re-1330 trieving the boundary maps of a single critical k-1331 simplex can be quadratic in the number $|\Sigma_k|$ of k-1332 simplices of Σ . 1333

Even if this is a very rare case, some solutions 1334 have been proposed to guarantee lower complex-1335 ity bounds by either using a Boolean function for 1336 marking the visited simplices [67, 57], or by us-1337 ing a priority queue [55] for limiting the number of 1338 simplices visited more than once. Both approaches 1339 have limitations however. The approach in [67] is 1340 useful for reconstructing a combinatorial represen-1341 tation for the connectivity of the critical simplices, 1342 but it does not visit all the possibile paths, which is 1343 necessary for retrieving the correct boundary maps 1344 in \mathbb{Z} . The approach in [55] can successfully retrieve 1345 the correct boundary maps, but it requires a input 1346 scalar function to be defined all over the simplices 1347 of Σ . 1348

The algorithm presented here is based on the general approach outlined in [10].

Algorithm 2 illustrates the steps required for 1351 traversing the gradient paths in a descending fash-1352 ion. Starting from a critical k-simplex τ , a breadth-1353 first traversal is performed by navigating from τ 1354 to its adjacent k-simplices passing through their 1355 shared (k-1)-simplices. The breadth-first traver-1356 sal is supported by a queue Q. Given a k-simplex 1357 τ_0 extracted from the queue Q (row 8), we examine 1358 all the (k-1)-simplices σ in the boundary of τ_0 1359

Algorithm 2 - BoundaryMaps (Σ, τ, V)

- 1: INPUT: Σ , d-dimensional simplicial complex
- 2: INPUT: τ , critical k-simplex
- 3: INPUT: V, Forman gradient
- 4: OUTPUT: M, boundary maps as collections of arcs

```
5: Q := \emptyset
```

```
6: Q.enqueue(\tau)
```

```
7: while Q \neq \emptyset do
```

- 8: $\tau_0 := Q.dequeue()$
- 9: for $\sigma_1 \in getBoundary(\tau_0, \Sigma)$ do
- 10: **if** $isPaired(\sigma_1, V, \tau_1)$ **then**

11: $Q.enqueue(\tau_1)$

12: else

- 13: $Add(M, \tau_1, \sigma_1)$
- 14: end if
- 15: end for

(row 11). For each (k-1)-simplex σ , if σ is paired 1360 with a k-simplex τ_1 (row 12), τ_1 is added to the 1361 queue (rows 15 and 16). If σ is a critical simplex, 1362 then σ is stored as on the boundary of τ . 1363

In Figure 12(a), we show an example of the de-1364 scending traversal performed by starting from crit-1365 ical triangle τ . For each edge on the boundary of 1366 τ , the paired triangle is visited and enqueued (indi-1367 cated in red in Figure 12(b)). The process contin-1368 ues recursively for each new triangle (Figure 12(c)) 1369 until the entire region associated with τ has been 1370 covered. When a critical edge σ is encountered, 1371 the relation with τ is stored in the boundary maps 1372 (Figure 12(d)). 1373

The procedures and the functions, on which Algorithm 2 is based, are: 1375

- Function getBoundary(τ, Σ): returns the immediate boundary of k-simplex τ , i.e., its (k – 1377 1)-faces. Extracting the immediate boundary 1378 is performed by taking all the combinations of 1379 the k vertices of τ , and it is a linear process in 1380 the number of vertices of τ .
- Function $isPaired(\sigma, V, \tau_1)$: returns the value $True \text{ if } (k-1)\text{-simplex } \sigma \text{ is paired with a } k$ simplex in V and the value False otherwise. Inthe former case it returns the paired simplex τ_1 . This is done by considering all the top simplices in the star of a vertex w of σ and visiting 1387



Figure 12: Descending traversal starting from τ . Expanding the gradient V-paths the critical edge σ is encountered and stored as connected to τ .

the gradient encoding of those top simplices which are incident in σ . The time complexity is $O(t_w)$ in the worst case, where t_w is the number of top simplices incident in vertex w.

Algorithm 2 is executed for each critical sim-1392 plex in the Forman gradient. For each k-simplex 1393 τ popped from Q, the for loop is performed up 1394 to k times. For each simplex σ on the boundary 1395 of τ we check whether it is paired or not $O(t_w)$. 1396 Then, we can conclude that each iteration of the 1397 while loop takes $O(kt_m)$, where t_m is the maxi-1398 mum of the number of top simplices t_k considered 1399 in *isPaired* at the varies of σ . The algorithm has 1400 $O(qkt_m)$ worst-case time complexity, where q is the 1401 number (counted with multiplicity) of k-simplices 1402 of Σ inserted in the queue Q. 1403

¹⁴⁰⁴ 9. Experimental results

In this section, we evaluate the performances of 1405 the coreduction-based algorithm for Forman gra-1406 dient computation and of the algorithm for com-1407 puting the boundary maps that give a Morse com-1408 plex, described in Section 2.3, which are based on 1409 the encoding of the original simplicial complex as 1410 an IA^* data structure. As described in Section 1411 8, computing the Forman gradient focusing on the 1412 lower star of each vertex is an operation well suited 1413 for distributed, or parallel implementation. To test 1414 the gain in performances of such an approach, we 1415 have implemented also a parallel version of our gra-1416 dient computation algorithm based on OpenMP. 1417 We compare our two implementations (sequential 1418 and parallel) with the implementation provided by 1419 *Perseus* which computes the Morse complex using 1420

Dataset	$ \Sigma $	C	IA^*	IA_p^*	IG
DTI-SCAN	24M	0.14M (171x)	$3.1\mathrm{m}$	$0.7 \mathrm{m}$	77.3h
VISMALE	118M	$0.94M_{(125x)}$	29.2m	6.5m	-
Ackley4	204M	$0.01M_{\rm \ (10^4x)}$	1.1h	$19.7 \mathrm{m}$	-
Amazon1	2.2M	$0.16M_{\rm \ (13.7x)}$	14.5s	3.7s	20.9h
Amazon2	$18.4 \mathrm{M}$	$0.37M_{\rm (49.7x)}$	281.9s	68.3s	>200h
Roadnet	4.8M	$0.75M_{(6.4x)}$	15.8s	6.06s	>200h
S1.0	0.6M	$16 \ (10^5)x$	56.8s	22.1s	61.7s
S1.2	26M	$12 \ (10^7)x$	4.2h	1.8h	-
S1.3	$197 \mathrm{M}$	$7_{\ (10^8)x}$	173h	74.3h	-

Table 2: Compression factor achieved by using the discrete Morse complex instead of the original simplicial complex. Column |C| indicates the number of critical simplices, as opposed to the number of simplices $|\Sigma|$, for each dataset. Columns IA^* , IA_p^* and IG indicate the timings required for computing the discrete Morse complex with our sequential implementation, the multi-thread implementation, and the Perseus tool, respectively.

an IG for encoding the input simplicial complex. ¹⁴²¹ To the extent of our knowledge, there are no implementations of the discrete Morse complex on a ¹⁴²³ Simplex Tree. ¹⁴²⁴

In our experiments, we consider both real and 1425 synthetic datasets. The hardware configuration 1426 used is an Intel i7 3930K CPU at 3.20Ghz with 1427 64GB of RAM. The data sets used in our experi-1428 ments are described in Table 1. There are tetrahe-1429 dralized volume data sets, and data sets obtained 1430 from networks and point clouds. Networks and 1431 point clouds have no filtration provided as input. 1432

In Table 2, we show first information about the size of the obtained discrete Morse complex (i.e., the number of cells), with respect to the original simplicial complex. The compression factor depends on the homological changes in the filtration of a dataset and on the dataset. Volumetric datasets benefit from a compression of about two orders of magnitude, network datasets are compressed by a factor of ten, while higher-dimensional
complexes are compressed by five to eight orders of
magnitude. This shows the advantage of using the
Morse complex instead of the original one for computing homological information.

By comparing the timings, we see that our ap-1446 proach (based on the IA^* data structure) always 1447 outperforms Perseus (based on the IG). When the 1448 number of simplices is low (dataset SPHERE-1.0), 1449 the two implementations require a similar amount 1450 of time but, as soon as the number of simplices 1451 increases, our approach is faster by two or three 1452 orders of magnitude. With the increasing of the di-1453 mension of the complex, we see that the complexity 1454 of computing the discrete Morse complex reaches 1455 its limits taking also 173 hours to complete for 1456 dataset Sphere-1.3. In our multi-threaded imple-1457 mentation, we have been able to use eight threads 1458 on our machine configuration, processing 8 vertices 1459 at a time. The speed up gained varies between a 1460 2x and a 5x. 1461

Figure 13 shows three evaluations. In the first 1462 graph, we are evaluating the memory used for rep-1463 resenting the simplicial complex Σ (in blue) and 1464 the Forman gradient V (in red). We can notice 1465 that for the first three data sets, the complex is 1466 the entity requiring the highest amount of mem-1467 ory. For the remaining data sets, we notice that 1468 when the dimension increases, the storage cost de-1469 creases. For example, when comparing ACKELEY4 1470 and S1.0, the total number of simplices is almost 1471 the same (see Table 2, column $|\Sigma|$), while memory 1472 consumption is dramatically reduced, being dataset 1473 S1.3 stored with less than 3.4MB compared to the 1474 7.9GB required by ACKELEY4. This is again due 1475 to the use of the IA^{*} data structure and to the en-1476 coding for the Forman gradient based on the top 1477 simplices. 1478

While extracting the lower star in Algorithm 1, 1479 the k-simplices are recursively extracted from the 1480 IA^* data structure and explicitly represented. This 1481 operation causes the main increase in the memory 1482 consumption at runtime. This is documented in the 1483 remaining graphs of Figure 13. We are indicating 1484 in green the static overhead required for storing the 1485 simplicial complex and the Forman gradient and in 1486 purple the amount of memory used at runtime. 1487

¹⁴⁸⁸ As we can notice (column IA^*), the difference ¹⁴⁸⁹ between the static overhead and the dynamic overhead is larger when working on datasets in higher 1490 dimensions, while it becomes negligible when work-1491 ing in two or three dimensions. This fact is intrinsi-1492 cally related to the dimension d of the original sim-1493 plicial complex. When d is small, the lower star of 1494 each vertex is also small. When working on higher 1495 dimensional complexes, the number of simplices in 1496 the lower star grows exponentially, since the num-1497 ber of simplices on the boundary of any k-simplex 1498 is exponential in k. In the worst-case scenario of 1499 our experiments (S1.3), the encoding of the star 1500 occupies 1.8GB at runtime, while storing the sim-1501 plicial complex and the gradient requires less than 1502 100MB. 1503

The implementation in Perseus, based on the IG, 1504 does not present a difference between static and dy-1505 namic overhead, since all the simplices are already 1506 represented at the beginning and progressively sim-1507 plified during the computation. Thus, the max-1508 imum peak is reached before starting the reduc-1509 tion algorithm. As a result, the IG presents seri-1510 ous limitations when the dimension of the complex 1511 increases. 1512

If we considering our parallel implementation 1513 (column IA_n^*), we see that the maximum peak of 1514 memory is higher, since all threads run on the same 1515 machine. Looking at the graphs in the first column 1516 (DTI-SCAN, VISMALE, ACKLEY4), we recognize 1517 that the runtime overhead of this version is still 1518 comparable to the one of the single-thread imple-1519 mentation. This is an expected result since these 1520 are low dimensional data sets with a fairly small 1521 lower star for each vertex. With the increasing in 1522 the data set dimension (second column), the over-1523 head required by the parallel implementation starts 1524 to be relevant. In the worst case, we have expe-1525 rienced a memory overhead up to 6 times larger 1526 than the single thread implementation (third col-1527 umn data set S1.3). These results suggest that 1528 the whole framework is promising for a distributed 1529 environment, where each process has its dedicated 1530 amount of memory. 1531

10. Concluding remarks

1532

We have studied different strategies to endow a ¹⁵³³ simplicial complex with a Forman gradient through ¹⁵³⁴ the use of homology-preserving operators and to ¹⁵³⁵ extract the corresponding discrete Morse complex. ¹⁵³⁶



Figure 13: Storage cost required by computing and storing the Forman gradient and the discrete Morse complex. The first graph on the left indicates the amount of memory in GB required for storing the simplicial complex (blue bars) and the Forman gradient (red bars). The remaining graphs indicate, for each dataset, the amount of memory in GB used for storing the complex and the Forman gradient (green bars), and the overhead required at runtime for computing the gradient (purple bars). Results are presented comparing the IA^* data structure, the IG, and the parallel implementation based on the IA^* data structure (indicated as IA_p^*). Missing columns represent experiments that exceeded the maximum amount of memory available.

We have formally proven the theoretical equiva-1537 lence of such methods which allow for reducing 1538 the complexity of the computation through reduc-1539 tions and coreductions. We have developed and 1540 implemented algorithms to efficiently build a dis-1541 crete Morse complex based on coreductions, on a 1542 space-efficient representation of the simplicial com-1543 plex and on a compact encoding of the Forman gra-1544 dient, also implementing a parallel version of the 1545 latter. 1546

Based on the results obtained from the parallel
implementation, we are currently working on a distributed version of Algorithm 1. Since the process
is localized within the star of each vertex, by distributing the computation on different machines,
we expect to get a boost on timings without affecting memory consumption.

We are also considering the application of this work in single-parameter and multi-parameter persistent homology computation, as the basis for tools for shape understanding and retrieval, and in segmentation of time-varying 3D scalar fields in the context of scientific data visualization.

¹⁵⁶⁰ The best implementation currently available in ¹⁵⁶¹ the literature for computing *persistent homology*

[28] on high-dimensional complexes is based on an-1562 notations and on the Simplex Tree and it represents 1563 all the simplices of the simplicial complex explicitly 1564 [62]. This is also the case for any persistent ho-1565 mology computation algorithms based on boundary 1566 map reduction, because of the need to represent all 1567 the simplices explicitly. This puts practical limita-1568 tions when working on large complexes. In these 1569 cases, our approach is particularly useful since the 1570 Morse complex is a simpler structure sharing the 1571 same persistent homology as the original simplicial 1572 complex. 1573

Multi-parameter persistent homology (also called 1574 multi-dimensional persistent homology) is an ex-1575 tension of persistent homology for data character-1576 ized by multiple parameters, like multi-field data 1577 sets. In this case, not a single filtration but mul-1578 tiple filtrations are considered. To date, the ap-1579 proaches proposed in the literature for computing 1580 multi-parameter persistent homology are at a pi-1581 oneering level and are not able to deal with the 1582 complexity and the size of real datasets. Recently, 1583 an interesting connection between multi-parameter 1584 persistent homology and discrete Morse theory has 1585 been pointed out in [14]. A formal proof is given 1586

of the equivalence between the multi-parameter 1587 persistent homology of the Morse complex defined 1588 by a Forman gradient compatible with the multi-1589 filtration and that of the underlying simplicial com-1590 plex is provided. An algorithm has been proposed 1591 by Allili et al. [15] for computing a Forman gradi-1592 ent on a vector-valued function, but its implemen-1593 tation is limited to triangle meshes of very small 1594 size. Based on the dimension-independent encod-1595 ing for the Forman gradient described in this pa-1596 per, we are planning to develop a new algorithm 1597 that works independently of the dimension of the 1598 domain (the underlying simplicial complex) and of 1599 the codomain (the number of filtrations provided). 1600 A parallel implementation will be also at the center 1601 of future studies for empowering the computation 1602 of multi-parameter persistent homology. 1603

In scientific visualization, *extremum graphs* have 1604 been defined as topological tools to understand and 1605 visualize the structure of 3D scalar fields, i.e., scalar 1606 fields defined at points in the three-dimensional Eu-1607 clidean space [13]. The extremum graph is a sub-1608 graph of the graph representing the boundary maps 1609 of the Morse complex. We recall that the boundary 1610 maps encode all the incidence relations between a 1611 critical k-simplex and a critical (k-1)-simplex, for 1612 $1 \leq k \leq d = dim(\Sigma)$. The extremum graph only 1613 represents the boundary maps between critical d-1614 simplices and (d-1)-simplices and between critical 1615 1-simplices and 0-simplices. For each pair of criti-1616 cal simplices, it also encodes the chain of simplices 1617 that connects the two critical ones. In [13], a vi-1618 sualization technique, called *topological spines*, has 1619 been developed specifically for extremum graphs 1620 not only of static 3D scalar fields, but also of time-1621 varying fields, which can be regarded as 4D scalar 1622 fields. The algorithms described in Section 8 can 1623 be suitably adapted to efficiently compute the ex-1624 tremum graphs of a scalar field. Using the scalar 1625 function as a filtering function, we can compute the 1626 Forman gradient V using Algorithm 1. The gradi-1627 ent paths of V now describe the behavior of the in-1628 put scalar field. Using Algorithm 2, we can extract 1629 the incidence relations between critical simplices 1630 by starting the descending traversal from critical 1631 d-simplices and from critical 1-simplices. Our ap-1632 proach will make the computation of extremum 1633 graphs [68] (and topological spines) feasible for 4D 1634 fields, but also for 3D fields defined on a tetrahe-1635 dral mesh (as needed for complex 3D domains), 1636

while the current approach [13] works only works 1637 on scalar fields defined on cubic grids. 1638

Acknowledgments

This work has been partially supported by the 1640 US National Science Foundation under grant number IIS-1116747. The authors wish to thank Davide 1642 Bolognini, Emanuela De Negri and Maria Evelina 1643 Rossi for their helpful comments and suggestions. 1644

References

- V. De Silva, R. Ghrist, Homological sensor networks, Notices of the American Mathematical Society 54.
- [2] R. Fellegara, U. Fugacci, F. Iuricich, L. De Floriani, Analysis of geolocalized social networks based on simplicial complexes, in: 9th
 ACM SIGSPATIAL International Workshop on Location-Based Social Networks (LSBN),
 ACM, 2016.
- S. Martin, A. Thompson, E. A. Coutsias, J.-P. 1655
 Watson, Topology of cyclo-octane energy landscape, Journal of Chemical Physics 132 (23) 1657 (2010) 234115. 1658
- [4] R. van de Weygaert, G. Vegter, H. Edelsbrunner, B. J.-T. Jones, P. Pranav, C. Park, W. A. Hellwing, B. Eldering, N. Kruithof, E. Bos, 1661
 et al., Alpha, Betti and the Megaparsec universe: on the topology of the cosmic web, in: 1663
 Transactions on Computational Science XIV, 1664
 Springer, 2011, pp. 60–101. 1665
- M. K. Chung, P. Bubenik, P. T. Kim, Persistence diagrams of cortical surface data, in: Information Processing in Medical Imaging, Springer, 2009, pp. 386–397.
- [6] R. Forman, Morse theory for cell complexes, 1670
 Advances in Mathematics 134 (1) (1998) 90- 1671
 145. 1672
- [7] P. Frosini, M. Pittore, New methods for reducing size graphs, International Journal of Computer Mathematics 70 (3) (1999) 505-517.

- [8] A. J. Zomorodian, The tidy set: a minimal simplicial set for computing homology of clique complexes, in: Proceedings of the 2010 Annual Symposium on Computational Geometry, ACM, 2010, pp. 257–266.
- [9] L. De Floriani, U. Fugacci, F. Iuricich, P. Magillo, Morse complexes for shape segmentation and homological analysis: discrete models and algorithms, Computer Graphics Forum 34 (2) (2015) 761–785.
- [10] V. Robins, P. J. Wood, A. P. Sheppard, Theory and algorithms for constructing discrete Morse complexes from grayscale digital images, IEEE Transactions on Pattern Analysis and Machine Intelligence 33 (8) (2011) 1646– 1691 1658.
- 1692 [11] S. Harker, K. Mischaikow, M. Mrozek,
 1693 V. Nanda, Discrete Morse theoretic algorithms
 1694 for computing homology of complexes and
 1695 maps, Foundations of Computational Mathe1696 matics 14 (1) (2014) 151–184.
- [12] S. Harker, K. Mischaikow, M. Mrozek, 1697 V. Nanda, H. Wagner, M. Juda, P. Dłotko, 1698 The efficiency of a homology algorithm based 1699 on discrete Morse theory and coreductions, 1700 in: Proceedings 3rd International Workshop 1701 on Computational Topology in Image Context 1702 (CTIC 2010). Image A, Vol. 1, 2010, pp. 41-1703 47. 1704
- [13] C. Correa, P. Lindstrom, P.-T. Bremer, Topological spines: a structure-preserving visual representation of scalar fields, IEEE Transactions on Visualization and Computer Graphics 1709 17 (12) (2011) 1842–1851.
- 1710 [14] M. Allili, T. Kaczynski, C. Landi, Reducing
 1711 complexes in multidimensional persistent ho1712 mology theory, Journal of Symbolic Computa1713 tion 78 (2017) 61 75.
- 1714 [15] M. Allili, T. Kaczynski, C. Landi, F. Masoni, Algorithmic construction of acyclic partial matchings for multidimensional persistence, Springer, 2017, pp. 375–387.
- 1718 [16] U. Fugacci, F. Iuricich, L. De Floriani, Ef1719 ficient computation of simplicial homology
 1720 through acyclic matching, in: Symbolic and

Numeric Algorithms for Scientific Computing 1721 (SYNASC), 2014 16th International Symposium on, 2014, pp. 587–593. 1723

- [17] D. Canino, L. De Floriani, K. Weiss, IA^* : 1724 an adjacency-based representation for nonmanifold simplicial shapes in arbitrary dimensions, Computers & Graphics 35 (3) (2011) 1727 747-753. 1728
- [18] V. Nanda, The Perseus software project for 1729 rapid computation of persistent homology. 1730 URL http://www.math.rutgers.edu/ 1731 ~vidit/perseus/index.html 1732
- [19] A. T. Lundell, S. Weingram, The topology of CW complexes, Van Nostrand Reinhold Company, 1969.
 1733
- H. Edelsbrunner, D. G. Kirkpatrick, R. Sei- del, On the shape of a set of points in the plane, IEEE Transactions on Information The-ory 29 (4) (1983) 551–559.
- [21] E. W. Chambers, V. De Silva, J. Erickson, 1740
 R. Ghrist, Vietoris-Rips complexes of planar point sets, Discrete & Computational Geometry 44 (1) (2010) 75–90.
- [22] A. Hatcher, Algebraic topology, Cambridge 1744 University Press, 2002.
 1745
- [23] A. J. Zomorodian, Fast construction of the Vietoris-Rips complex, Computer and Graphics (2010) 263-271.
- [24] V. De Silva, G. Carlsson, Topological estimation using witness complexes, in: Proceedings of the First Eurographics Conference on Point-Based Graphics, 2004, pp. 157–166.
- [25] V. De Silva, A weak definition of Delaunay triangulation, arXiv preprint cs/0310031.
- [26] L. J. Guibas, S. Y. Oudot, Reconstruction using witness complexes, Discrete & Computational geometry 40 (3) (2008) 325–356.
- T. K. Dey, F. Fan, Y. Wang, Graph induced 1758 complex on point data, in: Proceedings of the 1759 Twenty-ninth Annual Symposium on Computational Geometry, SoCG '13, 2013, pp. 107–1761 116.

- 1763 [28] H. Edelsbrunner, J. Harer, Persistent homol1764 ogy a survey, Contemporary Mathematics
 1765 453 (2008) 257–282.
- [29] A. J. Zomorodian, Topology for computing,
 Cambridge University Press, 2005.
- [30] R. Ghrist, Barcodes: the persistent topology
 of data, Bulletin of the American Mathematical Society 45 (1) (2008) 61–75.
- 1771 [31] R. Forman, A user's guide to discrete Morse
 1772 theory, Séminaire Lotharingien de Combinatoire 48 (2002) 35.
- 1774 [32] K. Mischaikow, V. Nanda, Morse theory for
 1775 filtrations and efficient computation of persis1776 tent homology, Discrete & Computational Ge1777 ometry 50 (2) (2013) 330–353.
- 1778 [33] L. De Floriani, A. Hui, Data structures for simplicial complexes: an analysis and a comparison, in: M. Desbrun, H. Pottmann (Eds.),
 1780 Proc. 3rd Eurographics Symposium on Geometry Processing, 2005, pp. 119–128.
- [34] H. Edelsbrunner, Algorithms in combinatorial
 geometry, Springer, 1987.
- [35] L. De Floriani, D. Greenfieldboyce, A. Hui, A
 data structure for non-manifold simplicial dcomplexes, in: Proceedings of the 2004 Eurographics/ACM SIGGRAPH Symposium on
 Geometry processing, ACM, 2004, pp. 83–92.
- Information [36] L. De Floriani, A. Hui, D. Panozzo, D. Canino,
 A dimension-independent data structure for
 simplicial complexes, Proceedings of the
 19th International Meshing Roundtable (2010)
 403–420.
- 1795 [37] D. Canino, L. De Floriani, Representing
 1796 simplicial complexes with Mangrove, Pro1797 ceedings of the 22nd Iinternational Meshing
 1798 Roundtable (2013) 465–483.
- [38] D. Canino, The Mangrove TDS Library: a
 C++ tool for fast prototyping of topological data structures (2012).
- ¹⁸⁰² URL http://mangrovetds.sourceforge.¹⁸⁰³ net

- [39] R. Fellegara, K. Weiss, L. De Floriani, The
 Stellar tree: a compact representation for simplicial complexes and beyond, arXiv preprint
 arXiv:1707.02211.
- [40] J.-D. Boissonnat, C. Maria, The Simplex Tree:
 an efficient data structure for general simplicial complexes, Algorithmica 70 (3) (2014)
 406–427.
- [41]С. Maria, J.-D. Boissonnat, М. Glisse, 1812 M. Yvinec, The GUDHI library: simpli-1813 cial complexes and persistent homology, in: 1814 H. Hong, C. Yap (Eds.), Mathematical Soft-1815 ware ICMS 2014, Springer, 2014, pp. 167-174. 1816
- [42] J.-D. Boissonnat, K. C. S., S. Tavenas, Building efficient and compact data structures for simplicial complexes, in: L. Arge, J. Pach (Eds.), 31st International Symposium on Computational Geometry, 2015, pp. 642– 656.
- [43] D. Attali, A. Lieutier, D. Salinas, Efficient
 data structure for representing and simplify ing simplicial complexes in high dimensions,
 in: Proceedings of the 27th ACM Symposium
 on Computational Geometry, 2011, pp. 501–
 509.
- [44] T. Lewiner, H. Lopes, G. Tavares, Optimal discrete Morse functions for 2-manifolds, Computational Geometry 26 (3) (2003) 221 – 233.
- [45] M. Mrozek, B. Batko, Coreduction homology algorithm, Discrete & Computational Geometry 41 (1) (2009) 96–118.
- [46] M. Mrozek, T. Wanner, Coreduction homology algorithm for inclusions and persistent homology, Comput. Math. Appl. 60 (10) (2010) 1837 2812–2833.
- [47] P. Dłotko, T. Kaczynski, M. Mrozek, T. Wanner, Coreduction homology algorithm for regular CW-complexes, Discrete & Computational Geometry 46 (2) (2011) 361–388.
- [48] B. Benedetti, F. H. Lutz, Random discrete 1843
 Morse theory and a new library of triangulations, Experimental Mathematics 23 (1) 1845 (2014) 66–94.

- [49] F. Cazals, F. Chazal, T. Lewiner, Molecular
 shape analysis based upon the Morse-Smale
 complex and the Connolly function, in: Proc.
 9th Annual Symposium on Computational Geometry, 2003, pp. 351–360.
- 1852 [50] H. King, K. Knudson, N. Mramor, Generating discrete Morse functions from point data,
 1854 Experimental Mathematics 14 (4) (2005) 435– 444.
- [51] A. Gyulassy, P.-T. Bremer, B. Hamann, 1856 V. Pascucci, Practical considerations in 1857 Morse-Smale computation, complex in: 1858 V. Pascucci, X. Tricoche, H. Hagen, J. Tierny 1859 (Eds.), Topological Methods in Data Analysis 1860 and Visualization: Theory, Algorithms, and 1861 Applications, Mathematics and Visualization, 1862 Springer, 2011, pp. 67–78. 1863
- 1864[52]A. Gyulassy, P.-T. Bremer, V. Pascucci, Computing Morse-Smale complexes with accurate1865geometry, IEEE Transactions on Visualization1867and Computer Graphics 18 (12) (2012) 2014–18682022.
- [53] D. Günther, J. Reininghaus, H. Wagner,
 I. Hotz, Efficient computation of 3D MorseSmale complexes and persistent homology using discrete Morse theory, The Visual Computer 28 (10) (2012) 959–969.
- 1874 [54] N. Shivashankar, S. Maadasamy, V. Natarajan, Parallel computation of 2D Morse-Smale
 1876 complexes, IEEE Transactions on Visualization and Computer Graphics 18 (10) (2012)
 1878 1757-1770.
- 1879 [55] N. Shivashankar, V. Natarajan, Parallel computation of 3D Morse-Smale complexes, Computer Graphics Forum 31 (3) (2012) 965–974.
- [56] R. Fellegara, F. luricich, L. De Floriani, 1882 K. Weiss, Efficient computation and simpli-1883 fication of discrete Morse decompositions on 1884 triangulated terrains, in: Proceedings of the 1885 22Nd ACM SIGSPATIAL International Con-1886 ference on Advances in Geographic Informa-1887 tion Systems, SIGSPATIAL '14, 2014, pp. 1888 223 - 232.1889
- [57] K. Weiss, F. Iuricich, R. Fellegara, L. De Flori ani, A primal/dual representation for discrete

Morse complexes on tetrahedral meshes, Computer Graphics Forum 32 (3) (2013) 361–370. 1893

- [58] A. J. Zomorodian, G. Carlsson, Computing persistent homology, Discrete & Computational Geometry 33 (2) (2005) 249–274.
- [59] U. Bauer, M. Kerber, J. Reininghaus, H. Wagner, PHAT Persistent Homology Algorithms
 Toolbox, in: H. Hong, C. Yap (Eds.), Mathematical Software ICMS 2014, Vol. 8592 of
 Lecture Notes in Computer Science, Springer, 2014, pp. 137–143.
- [60] U. Bauer, M. Kerber, J. Reininghaus, Distributed computation of persistent homology, 1903
 in: Proceedings of the Meeting on Algorithm 1905
 Engineering & Experiments, 2014, pp. 31– 38.
- [61] P. Dłotko, H. Wagner, et al., Simplification 1908 of complexes of persistent homology computations, Homology, Homotopy and Applications 1910 16 (1) (2014) 49–63.
- [62] J.-D. Boissonnat, T. K. Dey, C. Maria, The 1912
 compressed annotation matrix: an efficient data structure for computing persistent coho-1914
 mology, Algorithmica 73 (3) (2015) 607–619. 1915
- [63] S. Pemmaraju, S. Skiena, Computational Discrete Mathematics: combinatorics and graph theory with Mathematica, Cambridge University Press, 2003.
- [64] D. H. Ackley, A connectionist machine for genetic hillclimbing, Kluwer Academic Publishers, 1987.
- [65] F. Iuricich, The IA*, an indexed-based data structure with adjacencies for encoding simplicial complexes.
 URL https://github.com/IuricichF/ 1926
 IAstar 1927
- [66] M. Joswig, M. E. Pfetsch, Computing optimal Morse matchings, SIAM J. Discret. Math. 1929 20 (1) (2006) 11–25.
- [67] D. Günther, J. Reininghaus, I. Hotz, H. Wagner, Memory-efficient computation of persistent homology for 3D images using discrete
 Morse theory, in: 24th SIBGRAPI Conference
 1933

1935	on Graphics, Patterns and Images, 2011, pp
1936	25-32.

1937	[68]	V. Narayanan, D. M. Thomas, V. Natarajan,
1938		Distance between extremum graphs, in: 2015
1939		IEEE Pacific Visualization Symposium (Paci-
1940		ficVis), 2015, pp. 263–270.