Simplifying Complicated Simplicial Complexes: Discrete Morse Theory and its Applications

Rachel Elana Zax
rzax@fas.harvard.edu

Advisor: Michael Hopkins

Submitted to the Department of Mathematics in partial fulfillment
of the requirements for the degree of A.B. in Mathematics

Harvard University
Cambridge, Massachusetts
March 19, 2012

Acknowledgements

Many thanks to my thesis advisor, Professor Michael Hopkins, for providing invaluable guidance, advice, and candy throughout the course of this project. I am also very grateful to Zachary Abel for his help throughout the year, and to those who read and commented on my drafts: Professor Peter Kronheimer, Ernie Fontes, Ariel Rabkin, and Yair Rosenberg. Thanks to my family for their support, both within the realm of mathematics and outside of it; in particular, thank you to my mother Marcia for passing along the genes for passing the AP Calculus exam, to my father David for looking at my thesis, to my brother Ben for reading my thesis, and to my brother Aryeh for reading at least a few pages of my thesis and providing a comment – more than anyone might expect from a 15-year-old. I am deeply indebted to my instructors in mathematics: my high school teachers; Professors Anil Nerode and Allen Hatcher at Cornell; and my instructors in college: Professors Noam Elkies (also my advisor for the past three years), Lydia Bieri, Gerald Sacks, Barry Mazur, Jacob Lurie, Antti Knowles, and (again) Michael Hopkins. I am likewise grateful to my instructors in theoretical computer science, who instilled in me not only an appreciation for computer science but a deeper understanding of mathematics as well: Professors Salil Vadhan, Michael Mitzenmacher, and Michael Rabin. Finally, my thanks go out to Randall Munroe of xkcd, who has, throughout my college career, provided an incentive for me to work on math until late in the evenings by posting a new comic strip online at midnight three nights a week.

1 Introduction

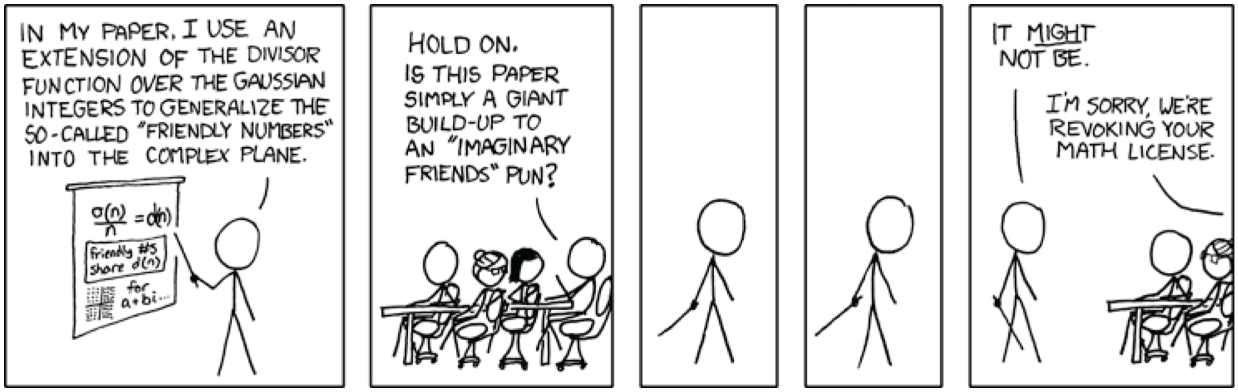

xkcd.com/410

Discrete Morse theory is a tool for determining equivalences between topological spaces arising from discrete mathematical structures. The mathematical objects we will consider are those that can be built from points, lines, triangles, tetrahedrons, and their higher-dimensional analogs; these can represent structures such as vector spaces and their subspaces, sets and their subsets, graphs comprising vertices and edges, and many others. Our notion of equivalence captures the idea of transforming one space into another by bending, shrinking, or expanding; but without tearing, creating holes, or filling holes. This theory was developed by Robin Forman in the 1990s as a combinatorial analog to Morse theory, developed by Marston Morse in the 1920s. The original theory deals with analyzing such equivalences for general topological spaces, while discrete Morse theory provides similar methods of analysis for topological spaces endowed with additional, discrete structure. For these structures, applications of the discrete theory are often more natural, as well as simpler and more straightforward to apply.

Discrete Morse theory has applications throughout many fields of pure and applied mathematics. Within pure mathematics, for example, the theory has been widely applied to problems in geometry [Gr11], topology [Be10a], and knot theory [Va93]; and within computer science, the theory has been used to evaluate data compression algorithms [Le04] and to bound the complexity of algorithms that determine whether graphs have certain properties – for example, whether all components of a graph are connected [Fo00]. If we wish to know whether a given property holds for a certain topological space, our question can often be reduced to the question of whether the space is equivalent to another space for which the property holds. For example, as we will see in Section 6 (as based on [Fo00]), whether a simple algorithm exists for determining if a graph is connected depends on whether the structure that represents the space of not-connected graphs can be shrunken to a point. Alas, it cannot, so any algorithm for testing graph connectedness must, at least in some cases, conduct an exhaustive search. This result has real-world implications: for example, it means that if we want to test a communications system – say, immediately after a disaster – to determine whether it is still connected, there is no guaranteed way of finding the answer without testing every component individually.

This thesis is an exposition of discrete Morse theory, with new applications of the theory to standard topological examples. Our presentation of the theory combines elements from the expositions of Forman in [Fo02] and [Fo04] and Jonsson in [Jo08], and also emphasizes standard results from graph theory upon which discrete Morse theory relies. We also briefly survey the literature on Morse's original Morse theory and the parallels that have been drawn between the continuous and discrete theories. We then analyze some of the standard examples of structures within topology using discrete Morse theory. Although these structures are already well understood, these proofs are new: we demonstrate how discrete Morse theory can be applied to produce the standard results, and how the intuition behind the theory can be leveraged to understand these results. We note the similarities of our technique in these examples to other applications of discrete Morse theory throughout the literature, and highlight commonalities that suggest a more general approach to applying discrete Morse theory to structures with analogous symmetries. We also present an application originally by Forman [Fo00], of discrete Morse theory to a generalization of the abovementioned problem in graph theory.

In our next section, we will define and discuss the structures that we will examine. The third section develops the theory itself, introducing the relevant ideas, exploring the mechanics of the theory, and demonstrating a few simple examples. Our fourth section links the ideas we have been developing to Morse's original theory and discusses the parallels between the two theories. The fifth section apply discrete Morse theory to standard topological examples, and the sixth section discusses the application mentioned above.

2 Topological basics

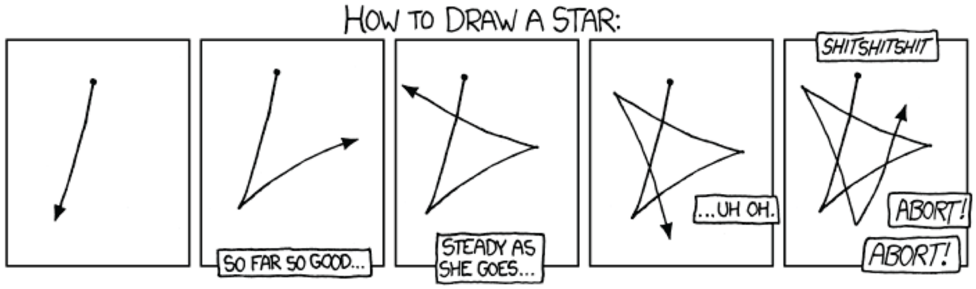

xkcd.com/1029

The situation we examine in this thesis is as follows: we have a topological space, endowed with the structure of a simplicial complex X, and we desire a simpler decomposition of this space requiring fewer cells. Often, this simplified decomposition is a standard form that is more easily recognizable. We accomplish this goal by collapsing our original structure as dictated by a discrete Morse function $f : X \to R$. Our final structure is usually no longer a simplicial complex, but a more general form called a cell complex. Discrete Morse theory in fact applies not only to simplicial complexes but to more general regular cell complexes, but for simplicity we restrict our attention in this paper to simplicial complexes.

Our main theorem will state that a simplicial complex X endowed with a discrete Morse function f is homotopy equivalent to a cell complex consisting of just the critical simplices of f (as will be defined below). In a number of special cases, this information is sufficient to entirely determine the homotopy type of X. For example, this is true when a number of spheres of the same dimension all intersect in a single point, as will happen often in our examples: as we will demonstrate later in this section, any cell complex consisting of a single 0-cell and k d-cells is homotopy equivalent to this structure.

However, in many cases attaching the cells in different ways gives rise to different homotopy types. For example, although the torus and the Klein bottle can both be constructed from one 0-cell, two 1-cells, and one 2-cell, these spaces have different homotopy types, as we will show below. With greater care, we can in fact determine the homotopy type rather than just the number of cells of each dimension, by considering the series of collapses that the discrete Morse function induces.

Before proceeding further, we define some of the aforementioned terms more precisely and justify the above statements.

Definition 2.1. An *n-simplex* is a closed convex polyhedron of dimension n created by joining $(n+1)$ vertices.

Definition 2.2. An *n-cell* is a closed region of dimension n, homeomorphic to an n-simplex, but as opposed to a simplex may have any number of vertices.

Definition 2.3. Let X be a topological space and c_d a cell of dimension d, with boundary ∂c_d. Let $a : \partial c_d \to X$ be a continuous *attaching map*. We *attach* c_d to X by taking the disjoint union of X and c_d and then identifying $x \in \partial c_d$ with $a(x) \in X$, for all $x \in \partial c_d$. We denote the resulting space by $X \cup_a c_d$.

There are many equivalent ways to define a simplicial complex; we follow Hatcher's definition in [Ha01]:

Definition 2.4. A *simplicial complex* X is a complex built from simplices attached via identification of their faces such that any simplex is uniquely determined by its vertices. We can describe X as a collection of vertices X_0 and sets X_n of n-simplices, i.e. $(n+1)$-element subsets of X_0, such that for all k-element subsets $\sigma_k \in X_k$, every j-element subset $\sigma_j \subset \sigma_k$ is an element of X_j, for all $j < k$.

Similarly, a cell complex can be defined in various ways; we follow Forman's definition in [Fo02]:

Definition 2.5. A (finite) *cell complex* X is a space for which there is a sequence

$$\emptyset \subset X_0 \subset ... \subset X_n = X$$

such that for $i = 1, ..., n$, $X_i = X_{i-1} \cup c_d$ for a cell of some dimension d. According to this definition, we require that X_0 be a 0-cell, since it is attached to the empty set.

We note that in the literature, Definition 2.5 is often given as the definition of a CW complex, while cell complexes are defined similarly but without restrictions on the order in which cells are attached. However, we follow Forman in referring to the above structure as a cell complex.

A cell complex that we will encounter often throughout this thesis is the wedge of spheres:

Definition 2.6. A *wedge of k spheres of dimension d* is a cell complex consisting of k d-cells along with a single 0-cell α, such that all of the n-cells intersect at α and nowhere else. We write this structure as $\bigvee_k S^d$.

And in fact, any cell complex built from this set of cells has the same homotopy type as a wedge of spheres. Many of our examples will require this result, so we prove it here. First, we require the following standard lemma from topology:

Lemma 2.7. *Let X be a topological space and c_d a d-cell, and $a, a' : \partial c_d \to X$ two continuous maps. If a is homotopic to a' then $X \cup_a c_d$ is homotopy equivalent to $X \cup_{a'} c_d$.*

Proposition 2.8. *A cell complex consisting of one 0-cell and k d-cells is homotopy equivalent to $\bigvee_k S^d$.*

Proof. We attach cells one at a time to create the complex. As noted above, we must begin by attaching the 0-cell. For each step of our construction thereafter, our only choice is to attach a d-cell. The first d-cell can only be attached by the constant map, so X_1 is a d-sphere. At the next step, any continuous map from the boundary of our new cell to X_1 is homotopic to the constant map, so our new structure X_2 is a wedge of 2 d-spheres. Continuing inductively, if X_i is homotopy equivalent to a wedge of i d-spheres, any continuous map from the new cell's boundary to X_i must be homotopic to the constant map and thus X_{i+1} is homotopy equivalent to a wedge of $(i+1)$ d-spheres. We continue until we reach $X_k = X$. $\square$

However, in general, the numbers of cells of each dimension is insufficient information to determine the homotopy type of a space. For example, although the torus T and the Klein bottle K can be built from the same set of cells, they have different homotopy types:

Example 2.9.

$$H_0(T) = \mathbb{Z}, H_1(T) = \mathbb{Z} \oplus \mathbb{Z}, H_2(T) = \mathbb{Z}$$

$$H_0(K) = \mathbb{Z}, H_1(K) = \mathbb{Z} \times \mathbb{Z}_2, H_2(K) = 0$$

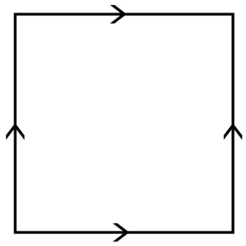

T

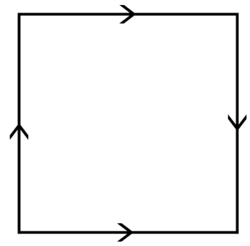

K

3 Discrete Morse theory

3.1 Definitions

Throughout this section, let X be a simplicial complex, with α, β simplices (or more generally, cells) of X.

Definition 3.1. A cell α is a *face* of β if $\alpha \subsetneq \beta$. We will denote "α is a face of β" by $\alpha \in \partial\beta$.

We will often speak about *codimension-1* faces, i.e. faces $\alpha \in \partial\beta$ such that $\dim \alpha = \dim \beta - 1$.

Definition 3.2. A cell β is a *coface* of α if $\alpha \subsetneq \beta$.

The essential components of Forman's discrete Morse theory are the discrete Morse function (Definition 3.3) and critical simplices (Definition 3.5). A discrete Morse function assigns values to simplices such that higher-dimensional simplices usually have higher values than lower-dimensional simplices, with some exceptions. A critical simplex is one near which no such exception occurs. Formally:

Definition 3.3. A *discrete Morse function* f on a simplicial complex X is a real-valued function on the set of simplices of X such that for each simplex β both conditions 1 and 2 hold:

1. At most one face α of β has a value greater than or equal to that of β.
2. At most one coface γ of β has a value less than or equal to that of β.

As we will demonstrate later (see Lemma 3.9), we cannot have both an α satisfying condition 1 as well as a γ satisfying condition 2 for a given β.

Example 3.4. Let X be the standard 2-simplex. The first diagram below gives the trivial discrete Morse function in which all simplices are mapped to their dimensions. The second gives a more complicated Morse function. The third gives a function that is not a Morse function.

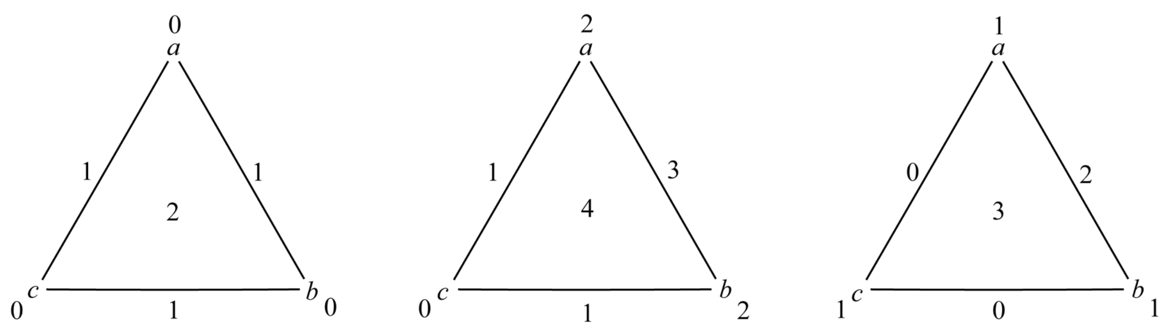

Definition 3.5. A *critical simplex* β of a discrete Morse function f on X is a simplex for which:

1. No face α of β has a value greater than or equal to that of β.
2. No coface γ of β has a value less than or equal to that of β.

Example 3.6. In our example above, the critical simplices of f on X are, in the first case, all simplices: $[a], [b], [c], [ab], [ac], [bc], [abc]$; and in the second case, the simplices $[c], [ab], [abc]$.

Note that we have here denoted simplices by their vertices, and we will use this convention throughout: if a simplex α has vertices $v_1, ..., v_n$, we may where convenient write it as $[v_1 ... v_n]$.

3.2 The main theorem

We can now state the main theorem of discrete Morse theory, as originally proven in [Fo98], although we defer the proof until the end of this section.

Theorem 3.7. *Let X be a simplicial complex with a discrete Morse function f. Then X is homotopy equivalent to a cell complex containing the same number of cells of a given dimension as there are critical simplices of f of that dimension.*

For example, for the 2-simplex of Example 3.4, we can reduce the original structure to just one 0-cell, one 1-cell, and one 2-cell, as dictated by the second function given. For the first function all cells are critical, so the theorem trivially applies: the original structure is homotopy equivalent to itself.

As another example, we consider the standard simplicial decomposition of the torus, and demonstrate how discrete Morse theory allows us to remove all but one 0-cell, two 1-cells, and one 2-cell, leaving us with the standard cell decomposition.

Example 3.8. The minimal simplicial complex on the torus consists of seven 0-simplices, twenty-one 1-simplices, and fourteen 2-simplices. We label the vertices a through g:

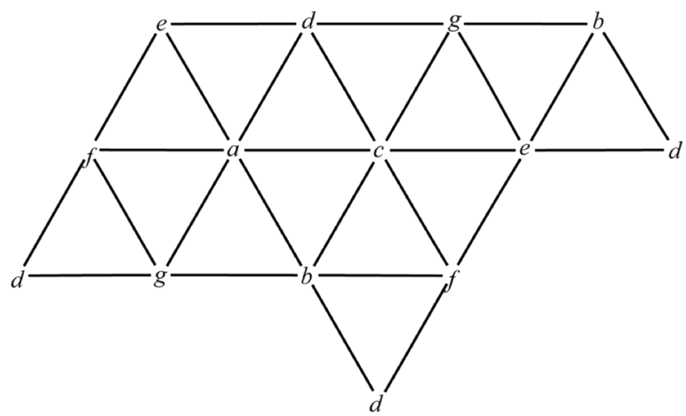

We can mark all non-critical simplices with arrows pointing into or out of them; these arrows are derived from a discrete Morse function, as we will explain, and each points from a simplex α to a simplex β such that $\dim \beta = \dim \alpha + 1$. These arrows can be construed in various ways: as a pairing of simplices, which is the approach we will generally take; as a discrete vector field, as Forman decribes in [Fo98]; and as indications of simplicial collapses, as we will describe in the proof of our main theorem. The critical simplices here are those labeled [d], [bg], [ef], and [bcf]:

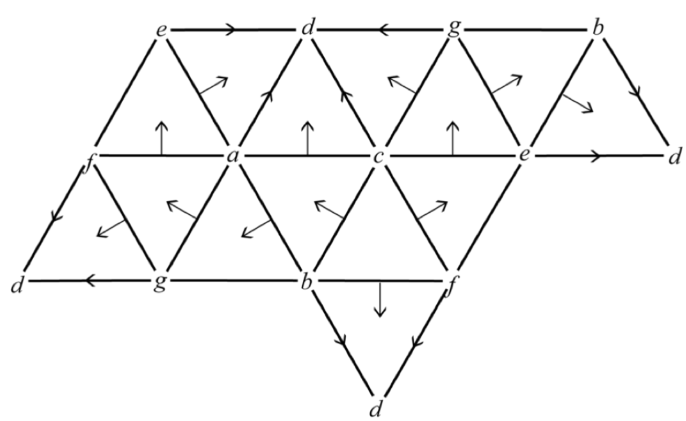

3.3 Pairings and the Hasse diagram

It is often useful to describe the discrete Morse function f not as a function but as a set of pairings of simplices P_f, or simply P when the choice of f is implied, where $\{\alpha, \beta\} \in P_f$ iff $\alpha \in \partial\beta$ and $f(\alpha) \geq f(\beta)$. This preserves the essential information regarding critical simplices without requiring us to know the actual values of f on every simplex. Thus we can work with the simplicial complex locally rather than globally, since we need only decide how to pair simplices with their neighbors, rather than assigning values over the entire complex such that the conditions are simultaneously satisfied at each simplex. The critical simplices as we have defined them above are exactly those that do not appear in the pairing P.

Instead of finding a discrete Morse function f for a complex X, it is often more convenient to find a pairing P and check that it corresponds to some discrete Morse function f on X. Clearly, every f has an associated P_f; but not every P corresponds to some f. We thus need to determine when a pairing P is equal to P_f for some f. We note that when such an f exists, it is not unique; each P_f actually corresponds to an equivalence class of discrete Morse functions, where f and f' are equivalent if they have the same critical simplices and the same pairs of simplices α, β such that $\alpha \in \partial\beta$ but $f(\alpha) \geq f(\beta)$. However, we are concerned only with the existence of a corresponding f, and not its uniqueness.

First, we note that for a pairing P associated with a discrete Morse function f, each simplex of X can be involved in at most one pair:

Lemma 3.9. *For all $\gamma \in X$, there exists at most one pair $\{\alpha, \beta\} \in P$ s.t. $\gamma = \alpha$ or $\gamma = \beta$.*

Proof. First, by condition 1 of Definition 3.3, there exists at most one α such that $\{\alpha, \gamma\} \in P$. Similarly, by condition 2, there is at most one β such that $\{\gamma, \beta\} \in P$.

Now assume that there exists a k-simplex γ such that $\{\alpha, \gamma\} \in P$ for some α and $\{\gamma, \beta\} \in P$ for some β, so that $f(\gamma) \leq f(\alpha)$ and $f(\beta) \leq f(\gamma)$. Since α is a $(k-1)$-simplex, we can write $\alpha = [v_{i_1} ... v_{i_k}]$ where the v_{i_j} are 0-simplices. Then since α is a face of γ, we have $\gamma = [v_{i_1} ... v_{i_k} v_{i_{k+1}}]$ for some 0-simplex $v_{i_{k+1}}$; and γ is a face of β, so $\beta = [v_{i_1} ... v_{i_k} v_{i_{k+1}} v_{i_{k+2}}]$ for some 0-simplex $v_{i_{k+2}}$. We consider the k-simplex $\gamma' = [v_{i_1} ... v_{i_k} v_{i_{k+2}}]$, so that $\alpha \in \partial\gamma$ and $\gamma \in \partial\beta$. We know from above that $\{\alpha, \gamma'\} \notin P$ and $\{\gamma', \beta\} \notin P$, so $f(\alpha) < f(\gamma')$ and $f(\gamma') < f(\beta)$.

Combining these inequalities with those above, we find that

$$f(\alpha) < f(\gamma') < f(\beta) \leq f(\gamma) \leq f(\alpha)$$

i.e. $f(\alpha) < f(\alpha)$, a contradiction. Thus we cannot have both $\{\alpha, \gamma\} \in P$ and $\{\gamma, \beta\} \in P$. $\square$

We illustrate the above lemma for the case $k = 1$. Here, we have $\alpha = [v_3], \gamma = [v_1 v_3], \beta = [v_1 v_2 v_3], \gamma' = [v_2 v_3]$.

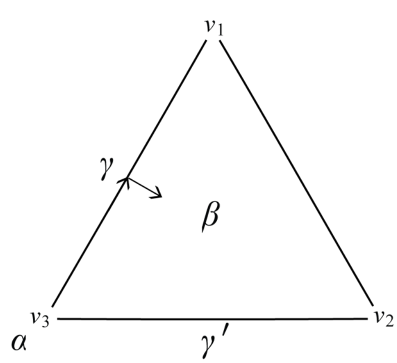

Lemma 3.9 gives us a necessary but not sufficient condition for the existence of an f corresponding to P. To derive the second condition, it is helpful to consider a graphical representation of the simplicial complex X and pairing P: the modified Hasse diagram. Hasse diagrams – which are not specific to discrete Morse theory, and are used throughout various fields of mathematics to represent finite partially ordered sets – are directed acyclic graphs, with which we represent a partial ordering via its transitive reduction, i.e. the minimal relation whose transitive closure gives the full ordering.

Definition 3.10. The **Hasse diagram** of a simplicial complex X is a directed acyclic graph on the partially ordered set of simplices of X, ordered by the relation $\alpha < \beta$ iff $\alpha \in \partial\beta$.

The transitive reduction is given by including an edge from β to α for all pairs of simplices $\{\beta, \alpha\}$ where $\alpha \in \partial\beta$ and $\dim\alpha = \dim\beta - 1$. We typically write all simplices of a given dimension in one row, and draw the diagram as a series of rows of decreasing dimension; thus we can refer to a specific level of the graph to mean the set of all simplices of a given dimension.

Example 3.11. The Hasse diagram of the standard 2-simplex:

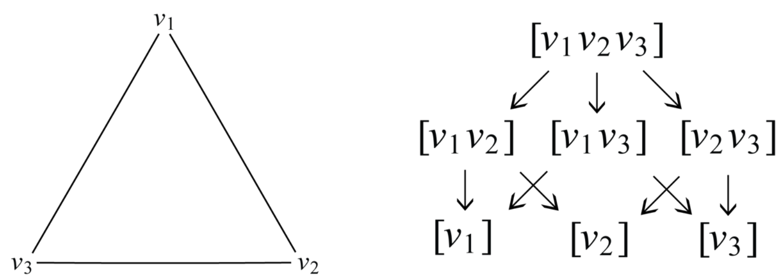

Given the Hasse diagram H of X and a pairing of simplices P, we can construct a modified directed graph H' by flipping the edge between any two simplices that appear as a pair in P. Note that a pairing P consists only of pairs $\{\alpha, \beta\}$ such that $\dim(\alpha)$ and $\dim(\beta)$ differ by one, so there is always a corresponding arrow in the graph for each possible pair. When the graph concerned is clear from context, we may refer to H' as the "Hasse diagram" rather than the "modified Hasse diagram."

Example 3.12. The modified Hasse diagram of the standard 2-simplex for a given pairing:

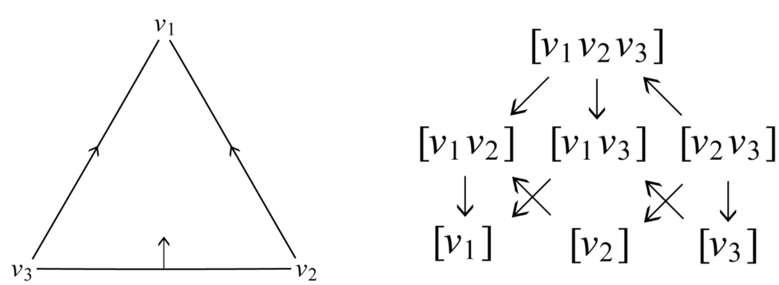

The additional condition we require for a pairing P to correspond to some f has to do with whether P – and thus the associated modified Hasse diagram H' – contains a cycle.

Definition 3.13. A pairing P is *cyclic* if there exists a sequence of simplices $\alpha_1, \beta_1, ..., \alpha_n, \beta_n, \alpha_{n+1}$ where $n > 1$ such that

- $\dim(\alpha_i) = d$ and $\dim(\beta_i) = d + 1$ for some d, for all $i = 0, ..., n$,

- both $\alpha_i, \alpha_{i+1} \in \partial\beta_i$, but $\alpha_i \neq \alpha_{i+1}$, and $\{\alpha_i, \beta_i\} \in P$ for all $i = 0, ..., n$, and
- $\alpha_1 = \alpha_{n+1}$.

We call the set $\{\alpha_1, \beta_1, ..., \alpha_n, \beta_n, \alpha_{n+1}\}$ a *cycle*. If a pairing P is not cyclic, then it is *acyclic*.

Theorem 3.14. *P is acyclic if and only if the corresponding Hasse diagram H' is acyclic.*

Proof. First, assume that P contains a cycle $\{\alpha_1, \beta_1, ..., \alpha_n, \beta_n, \alpha_{n+1}\}$. We have that $\alpha_i, \alpha_{i+1} \in \partial\beta_i$ for all i. Each pair $\{\alpha_i, \beta_i\}$ is in the pairing P, so there is an arrow in H' from α_i to β_i. Then since $\alpha_{i+1} \in \partial\beta_i$ and $\{\alpha_{i+1}, \beta_i\}$ is *not* a pair in P, there is an arrow in H' from β_i to α_{i+1}. We can follow these arrows to form a path from α_1 to $\alpha_{n+1} = \alpha_1$; thus we have a cycle in H'.

Now, assume there is a cycle in H'. We claim that this cycle can include only simplices of dimension d and $(d+1)$, for some d. Assume not; then since at every step in the path we must move up or down in dimension, at some point in the cycle we must move upward twice from a simplex α of dimension $(d-1)$ to a simplex γ of dimension $(d+1)$; this can only be achieved by a sequence of two consecutive upward arrows, from α to β and β to γ. But then $\{\alpha, \beta\}$ and $\{\beta, \gamma\}$ must have both been in P. So there is a cycle among simplices of dimension d and $(d+1)$; this cycle must alternate between downward and upward arrows, where each upward arrow indicates a pair $\{\alpha_i, \beta_i\} \in P$ and each downward arrow indicates an $\alpha_{i+1} \neq \alpha_i \in \partial\beta_i$, leading from α_1 back to $\alpha_{n+1} = \alpha_1$. This sequences gives us a cycle in P, as desired. ■

We claim that a pairing P has an associated discrete Morse function f if and only if every simplex is a member of at most one pair in P and P is acyclic; this was originally proven by Forman in [Fo98], using instead the language of discrete vector fields. With the help of the following results from graph theory, we will prove that these two conditions together are necessary and sufficient.

Lemma 3.15. *If G is a finite acyclic graph then there exists a vertex in G having no inward-oriented arrows.*

Proof. Assume every vertex v_i has an inward-oriented arrow. We choose some v_{i_1} and backtrack along one of its inward-oriented arrows to find v_{i_2}. Continuing inductively, we create a chain, always choosing a v_{i_j} that we have not already chosen. If at some point there is no such v_{i_j}, we have found a cycle. If we reach the final vertex v_{i_n} without having created a cycle, backtracking along the final inward arrow will force us to choose a vertex already in the chain, thus creating a cycle. □

This fact allows us to perform a *topological sort* on any acyclic graph: we choose one such vertex with no inward-oriented arrow, add it to the front of a list, and then remove the vertex and its associated outward arrows from the graph. The remaining subgraph is still acyclic, so we can continue inductively, removing vertices with no inward arrows one by one and appending them to the list in this order, until no vertices remain. We use this construction in the following lemma:

Lemma 3.16. *Let G be a finite directed graph. G is acyclic if and only if there is a real-valued function f on the vertices such that $f(v_i) > f(v_j)$ whenever there is an edge from v_i to v_j.*

Proof. First, assume that G contains a cycle of vertices $\{v_0, v_1, ..., v_n, v_{n+1} = v_0\}$. Then for any such function f, we require $f(v_0) > f(v_1) > ... > f(v_n) > f(v_{n+1}) = f(v_0)$, i.e. $f(v_0) > f(v_0)$, which is impossible.

Now, assume that G contains no cycles. We can perform a topological sort on G as described above, giving us a list of vertices $\{v_1, ..., v_m\}$. We assign values in decreasing order: let $f(v_i) = m - i$. Whenever there was an edge from v_i to v_j, our procedure added v_i to the list before v_j, so $f(v_i) > f(v_j)$ as desired. □

It follows that such an f is strictly decreasing along each directed path in G.

Theorem 3.17. *A pairing P on a simplicial complex X corresponds to a discrete Morse function f if and only if every simplex of X appears in at most one pair of P and P is acyclic.*

Proof. For the forward direction, if we have a discrete Morse function f, we can construct the corresponding pairing P; by Lemma 3.9 every simplex appears in at most one pair in P. If P were cyclic with cycle $\{\alpha_1, \beta_1, ..., \alpha_n, \beta_n, \alpha_{n+1}\}$ then we would have $f(\alpha_1) > f(\beta_1) > ... > f(\alpha_n) > f(\beta_n) > f(\alpha_{n+1})$ i.e. $f(\alpha_1) > f(\alpha_1)$. Thus P is acyclic.

Conversely, by Theorem 3.14, P is acyclic if and only if the corresponding modified Hasse diagram is acyclic. By Lemma 3.16, this occurs exactly when there is a real-valued function of the vertices f such that $f(\alpha) > f(\beta)$ whenever there is an edge from α to β. By the construction of the modified Hasse diagram, f assigns higher values to higher-dimensional simplices, except in cases where $\{\alpha, \beta\} \in P$. Since no simplex appears in more than one pair in P, for every simplex α this function f can, at most, assign a higher value than $f(\alpha)$ to one face of α, or assign a lower value to one coface of α. Thus, we take f to be our discrete Morse function. ■

3.4 Collapsing simplices

The final idea we introduce before proving our main theorem is that of simplicial collapse. As noted earlier, although we are working with spaces that have the structure of simplicial complexes, the process of collapsing may convert a simplicial complex into a more general cell structure. In this section, we will thus speak about both simplicial complexes and cell complexes.

Definition 3.18. Let α be a face of X. If there is no $\beta \in X$ such that $\alpha \in \partial\beta$, then α is a *maximal face* in X.

Definition 3.19. Let α, β be faces of X, with $\alpha \in \partial\beta$. If α is not a face of any other cell $\beta' \in X$, then α is a *free face* of β.

Note that if α is a free face of β, then β is maximal in X (otherwise, if $\beta \in \partial\gamma$, then $\alpha \in \partial\gamma$ as well). In addition, only a codimension-1 face can be a free face: if $\dim\alpha < \dim\beta - 1$, there would exist $\gamma \in \partial\beta$ with $\dim\alpha < \dim\gamma < \dim\beta$ such that $\alpha \in \partial\gamma$.

Theorem 3.20. *If α is a free face of β, then X is homotopy equivalent to $X\backslash(int(\alpha) \cup int(\beta))$. We refer to the deformation retraction $X\backslash(int(\alpha) \cup int(\beta))$ as the simplicial collapse of X with respect to α and β.*

In fact, when X and X' are related by such a simplicial collapse, they are *simple homotopy equivalent*. Simple homotopy theory was developed by Whitehead, and described by Cohen in [Co73]; as they demonstrate, simple homotopy equivalence is a refinement of the idea of homotopy equivalence, and thus implies homotopy equivalence. We will not concern ourselves with the distinction here.

Discrete Morse theory can be viewed as a generalization of the theory of simplicial collapse, as Jonsson formulates it in [Jo08]. We consider a simplicial complex X and a discrete Morse function f, inducing a pairing P. As we will see in the proof of Theorem 3.7, we will remove unmatched faces before performing collapses involving their neighbors, so we can drop the condition that α be a free face and instead require only that α be a face of no other matched simplices in P. In particular, our goal will be to create a series of pairs $\{\alpha_i, \beta_i\}$ such that for all i, $\alpha_i \notin \partial\beta_j$ for any $j > i$.

Theorem 3.21. *Any acyclic pairing P on X admits an ordering of pairs $\{\alpha_1, \beta_1\}, ..., \{\alpha_n, \beta_n\}$ such that for all i, $\alpha_i \notin \partial\beta_j$ for any $j > i$.*

Proof. Let $H' \subset H$ be the subgraph consisting of only the matched faces in H. An arrow oriented inward to a simplex α from a simplex γ means either a) $\alpha \in \partial\gamma$ and $\{\alpha, \gamma\} \notin P$ or b) $\gamma \in \partial\alpha$ and $\{\alpha, \gamma\} \in P$. If α has no inward-oriented arrows, neither a) nor b) is satisfied, i.e. α is not paired with any of its faces, and the coface which α is paired is the only matched face of X containing α. Since $H' \subset H$ is acyclic, by Lemma 3.15, there is a simplex $\alpha_1 \in H'$ that has no inward-oriented arrows. Thus α_1 is a face of exactly one matched simplex in H', say β_1, with which it is paired. Thus we can perform the simplicial collapse $X \to X \backslash \{\alpha_1, \beta_1\}$ to find a homotopy-equivalent complex; call it X_1. We now consider the subgraph $H'_1 := H' \backslash \{\alpha_1, \beta_1\}$ (also removing the corresponding edges). Since $H'_1 \subset H'$, it is acyclic, so we can repeat the process to find a new pair $\{\alpha_2, \beta_2\}$. Continuing inductively, at the ith step we can find another simplex that is the face of exactly one matched simplex in X_{i-1} and is thus a valid candidate for a simplicial collapse. We continue until we have gone through all the matched pairs and our final subgraph H'_n is empty.

The sequence $\{\alpha_1, \beta_1\}, ..., \{\alpha_n, \beta_n\}$ is now the desired ordering of pairs such that for all i, $\alpha_i \notin \partial\beta_j$ for any $j > i$. ■

The above theorem gives a sequence of collapses that cancels out, pair by pair, the portion of X consisting of non-critical simplices. This is a major idea of discrete Morse theory – that all of these non-critical simplices have no effect on the homotopy type of X, and thus can be removed. The proof of the main theorem of discrete Morse theory is essentially the above, together with the mechanics of addressing unmatched simplices as they arise.

Proof of Theorem 3.7: Let H be the modified Hasse diagram associated with f on X. Since H is acyclic, there is a simplex α that has no directed edges oriented into it. Either α is unmatched, in which case it has no cofaces, i.e. it is maximal and critical; or it is matched, in which case as above, α is a free face of β with $\{\alpha, \beta\} \in P$.

If α is maximal and critical, then let $X' := X \backslash int(\alpha)$. Clearly, X is the result of attaching the cell α to the cell complex X'.

If α is a free face of β with $\{\alpha, \beta\} \in P$, let $X' := X \backslash (int(\alpha) \cup int(\beta))$. Then, by Theorem 3.20, X is homotopy equivalent to X'.

We now consider $X' \subset X$: the corresponding pairing P' remains acyclic, so we can continue by induction, choosing another simplex α' in H' having no inward-oriented edges. At each step, we either find and remove another critical simplex, or perform another step in constructing the ordering of pairs of Theorem 3.21. At each step we remove some number of cells, so this protocol eventually terminates.

Retracing these steps in reverse, we successively perform one of two operations: we either attach a pair of non-critical cells, without changing the homotopy type; or we attach a critical cell. The steps in which we attach two non-critical cells have no effect on the homotopy type of our final complex, so we can construct a complex equivalent to X by leaving out these steps and just attaching the critical cells one by one. Thus X is equivalent to a complex containing a cell of dimension k for every critical simplex of dimension k, as desired. ■

3.5 Example: the torus

We demonstrate the technique on the torus, our example from above.

Example 3.22. We begin with the traditional simplicial complex containing seven 0-simplices, twenty-one 1-simplices, and fourteen 2-simplices, and end with a cell complex consisting of one

0-cell, two 1-cells, and one 2-cell. To conserve space, rather than depicting each collapse one by one, we will at each step perform all of the valid collapses at once:

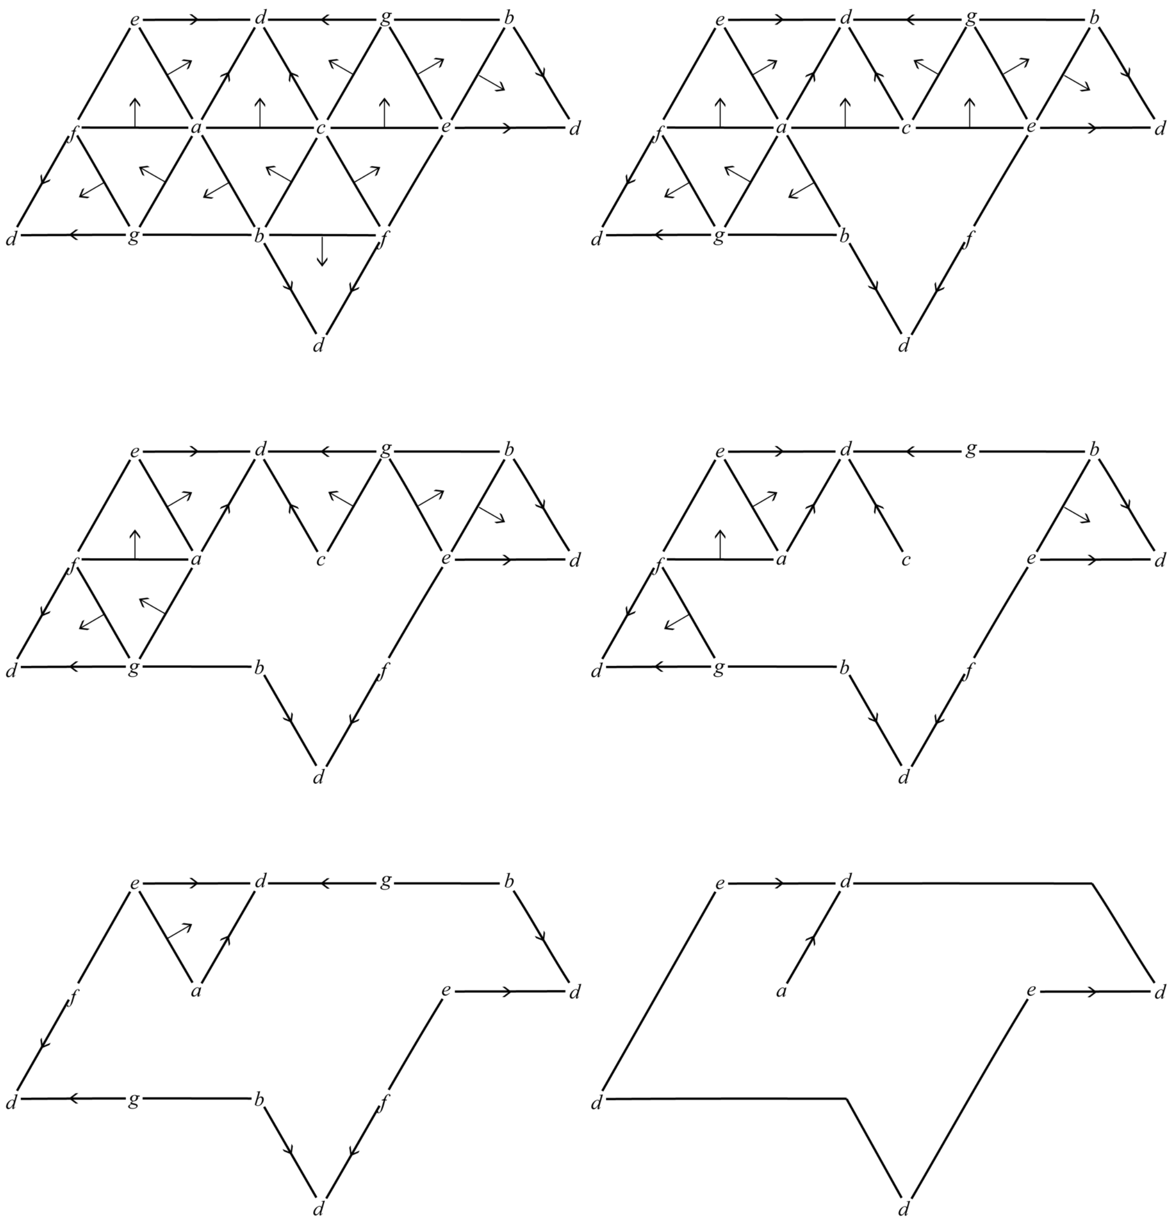

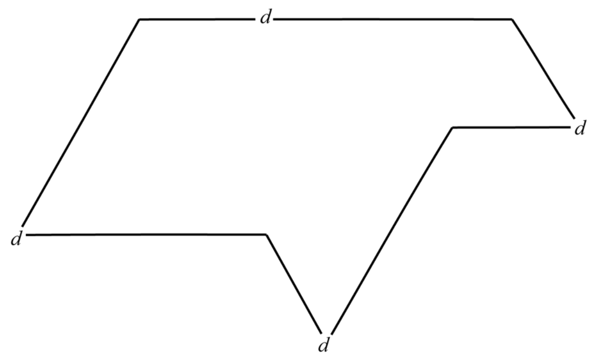

3.6 Canceling critical points

Finally, we might wonder how to construct optimal discrete Morse functions, having as few critical simplices as possible. In general this is a difficult problem – for example, as Forman remarks in [Fo02], it contains the Poincaré conjecture – and algorithmically, the problem is NP-hard [JP06]. Those constructing discrete Morse functions in practice often rely upon heuristic techniques [Le02].

However, there is a simple condition that is sufficient to cancel to critical simplices; as Forman notes in [Fo02], the continuous version appears as a key step in a more general proof by Morse in [Mo34]. Forman formulates the analogous condition for discrete Morse functions, which he presents in terms of discrete vector fields in [Fo98]. We recast this idea in the language of pairings; the theorem is as follows:

Theorem 3.23. *Let f be a discrete Morse function on X such that α, β are critical, with* $\dim \alpha = \dim \beta - 1$*, and let there be exactly one sequence of distinct simplices*

$$\{\beta =: \beta_0, \alpha_0, \beta_1, \alpha_1, ..., \beta_n, \alpha_n := \alpha\}$$

such that for $i = 0, ..., n$, $\alpha_i \in \partial\beta_i$ and for $i = 0, ..., n-1$, $\alpha_i \in \partial\beta_{i+1}$ and $\{\alpha_i, \beta_{i+1}\} \in P$. Then there is another discrete Morse function f' with the same critical simplices as f, except that α, β are no longer critical.

Proof. We take the pairing P associated with f and remove all pairs of the form $\{\alpha_i, \beta_{i+1}\}$ as described above to create a modified pairing P'. We then add to P' all pairs of the form $\{\alpha_i, \beta_i\}$ for $i \in [0, ..., n]$.

No simplex is involved in more than one pair in P', since P' just contains all the simplices paired in P, plus α and β, which were previously unpaired. And P is acyclic, since the only new path we have introduced is $\alpha_n, \beta_n, \alpha_{n-1}, ..., \alpha_0, \beta_0$ (and its subpaths), and any cycle involving this path would require an additional path from $\beta_0 = \beta$ to $\alpha_n = \alpha$; but we assumed that the sequence we just reversed was, until then, the unique path from β to α.

The unpaired simplices in P' are exactly those unpaired in P, less α and β. Thus P' corresponds to a discrete Morse function f' with the same critical simplices as f, less α and β, as desired. ∎

We illustrate an example of the above theorem: the path in the first figure can be reversed to create the path including α and β in the second.

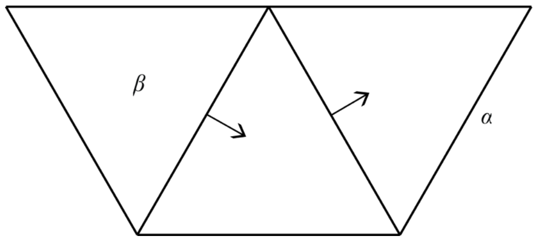
β
α
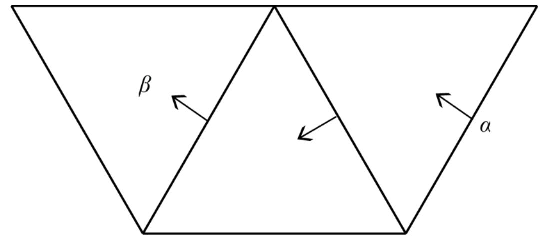
β
α

4 (Continuous) Morse theory

Discrete Morse theory is an adaption of a continuous version, simply called Morse theory, in which the object of study is a smooth manifold rather than a simplicial complex. Morse theory was introduced by Marston Morse in the 1920s [Mo34], and Milnor's book [Mi63] is now the standard reference. The theory relies on the observation that a differentiable function on a manifold will generally reflect the manifold's topology: by analyzing the function's critical points, it is possible to construct a cell structure for the manifold. In this section, we explore the parallels between discrete Morse theory and the original theory. First, we introduce the basic components of Morse theory: the Morse function and its critical points, which are analogous to the discrete versions with which we began. In this section, let X be a smooth manifold.

Definition 4.1. Let $f : X \to \mathbb{R}$ be a smooth function. A point $x \in X$ is a *critical point* if

$$\frac{\partial f}{\partial x_1}(x) = ... = \frac{\partial f}{\partial x_n}(x) = 0$$

Definition 4.2. Let $x \in X$ be a non-degenerate critical point of f (i.e. the Hessian $\frac{\partial^2 f}{\partial x_i \partial x_j}(x) \neq 0$). The number of linearly independent axes along which f is decreasing at x is the *index* of x.

The main result of Morse theory, as proven in [Mo34], is as follows:

Theorem 4.3. *Let $f : X \to \mathbb{R}$ be a smooth function. For all non-degenerate critical points x of f, X contains a cell of dimension d_x, where d_x is the index of x.*

We will explore this theory with a few examples. A Morse function that is simple to analyze is given by the height function on the manifold, which is valid whenever no two critical points occur at the same height. If they do, we can "tilt" the manifold to avoid this problem:

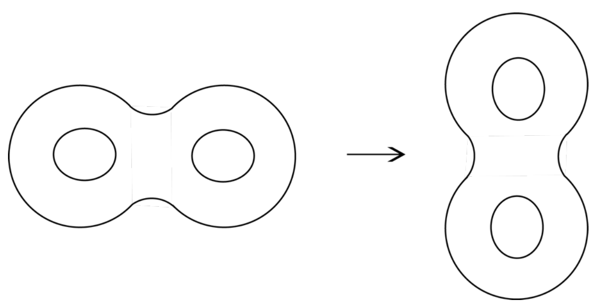

We can then understand the Morse function as giving us a method from which to build the simplicial complex from the bottom up. This interpretation of the Morse function arises from the following theorems of Morse, as demonstrated by Milnor in [Mi63]:

Definition 4.4. Let X be a smooth manifold. Then we define $X(a) \subset X$ to be the submanifold

$$\{x \in X | f(x) \leq a\}$$

Theorem 4.5. *If there is no critical point $x \in X$ such that $a < f(x) \leq b$, then $X(a)$ and $X(b)$ are homotopy equivalent.*

Theorem 4.6. *If there is exactly one critical point x such that $a < f(x) \leq b$, let d be the index of x. Then $X(b)$ is homotopy equivalent to $X(a) \cup c_d$, where c_d is a cell of dimension d.*

These theorems are actually entirely analogous to the discrete case. We define $X(a) \subset X$ in the discrete case analogously to Definition 4.4. We can then restate Theorem 3.7 to parallel Theorems 4.5 and 4.6, and the proofs follow from our proof of Theorem 3.7:

Theorem 4.7. *Let X be a simplicial complex with discrete Morse function f. If there is no critical simplex $\alpha \in X$ such that $a < f(\alpha) \leq b$, then $X(b)$ and $X(a)$ are homotopy equivalent.*

Theorem 4.8. *If there is exactly one critical simplex α such that $a < f(\alpha) \leq b$, let d be the dimension of α. Then $X(b)$ is homotopy equivalent to $X(a) \cup_{\partial\alpha} c_d$, where c_d is a cell of dimension d.*

Theorem 4.6 makes clear why having two critical points at the same height pose a problem. However, as long as the critical points are isolated, we can tilt the manifold as demonstrated above to eliminate the problem. We examine as an example the two-holed torus above, with its corresponding height function when tilted vertically:

Example 4.9. The two-holed torus:

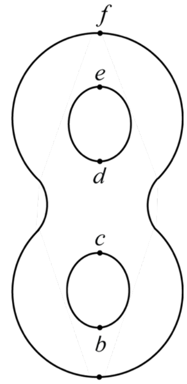

The labeled points are the critical points of the height function. At point a, the height is everywhere increasing, so this corresponds to a 0-cell. At points b, c, d, and e, the height is decreasing along one axis, so these each correspond to a 1-cell. Finally, at f, the height is decreasing along both axes, giving us a 2-cell. Thus we find that such a manifold is homotopy equivalent to a cell complex with one 0-cell, four 1-cells, and one 2-cell. We can envision building the manifold from the bottom up, attaching a cell of the appropriate dimension each time we cross a critical point. In this case we begin with the 0-cell, attach the four 1-cells in succession, and finally attach the 2-cell last.

We will examine the circle as a simple example, demonstrating how discrete and continuous Morse theory allow us to analyze the structure in parallel ways. Many structures admit analysis via both discrete Morse theory and the original Morse theory, although parallel functions are not always as simple to construct as they will be in this example.

Example 4.10. The fully reduced cell complex on the circle consists of a single 0-cell and a single 1-cell:

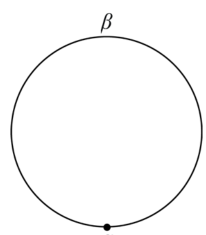

Any simplicial complex on the circle with an associated discrete Morse function will have as critical simplices this 0-cell and 1-cell. We can impose a different structure than the one show above by adding an equal number of non-critical 0-cells and 1-cells; we can then use discrete Morse theory to analyze the resulting structure. As the number of extra cells in the structure approaches infinity, we can view the analysis as a question for continuous Morse theory.

Let α be our critical 0-cell and β our critical 1-cell, and without loss of generality let $f(\alpha) = 0, f(\beta) = 1$. A decomposition with $(n+1)$ 0-cells and $(n+1)$ 1-cells each can be thought of as having the critical 0-cell at the bottom, the critical 1-cell at the top, and $\frac{n}{2}$ 0-cells and $\frac{n}{2}$ 1-cells between them on either side. Then we can assign values of f as follows: let i_γ indicate the distance, in terms of number of cells, between a given cell γ and the critical 0-cell. Let $f(\gamma) = \frac{i_\gamma}{n}$; so we have an increasing function with its minimum at α and its maximum at β.

This is a valid discrete Morse function, since every 0-cell except α has f value greater than the 1-cell below it but less than the 1-cell above it, and similarly every 1-cell except β has f value less than the 0-cell above it but greater than the 0-cell below it. This function pairs each 0-cell with the 1-cell directly below it, giving rise to a series of simplicial collapses that result in the reduced cell complex pictured above.

As $n \to \infty$, our discrete Morse function f turns into a smooth, monotonically increasing function between α and β – essentially a height function (though possibly increasing at a rate different from the height, depending on the lengths we assign to each successive extra pair of simplices). The critical points of this height function are α, where the height is nowhere decreasing, thus representing a 0-cell; and β, where the height is decreasing along one axis, thus representing a 1-cell.

Given a structure that admits analysis by both theories, we might ask which theory produces the sharpest result. First, we must introduce a standard by which to measure our functions:

Definition 4.11. The ith *Betti number* of a manifold X is $b_i := \dim H_i(X)$.

As proven in [Mi63], for a manifold X with discrete Morse function f, if n_i denotes the number of critical simplices of dimension n (or f a Morse function with n_i critical points of index n), then the following inequalities hold:

Theorem 4.12. *The weak Morse inequalities:*

1. *For* $i = 0, ..., \dim X$, $n_i \geq b_i$
2. $$\sum_{i=0}^{\dim X} (-1)^i n_i = \sum_{i=0}^{\dim X} (-1)^i b_i$$

In the second, we note that the right-hand side is the Euler characteristic of X, $\chi(X)$. This observation will be important to much of our later analysis. Intuitively, this equation comes from the fact that when we perform a simplicial collapse, we get rid of two simplices whose dimensions differ by one, so their contributions to the sum added up to 0: one of them contributed +1 while the other contributed -1.

As the name implies, these inequalities can be strengthened to the strong Morse inequalities:

Theorem 4.13. *The strong Morse inequalities: For all* $d = 0, ..., \dim X$,

$$\sum_{i=0}^{d} (-1)^{d-i} n_i \geq \sum_{i=0}^{d} (-1)^{d-i} b_i$$

These inequalities set bounds for how few critical cells we might possibly find in a Morse function; thus we can judge Morse functions by how closely they approach these bounds.

Definition 4.14. If X is a manifold with Morse function f, if $n_i = b_i$ for all i, f is *perfect*.

For example, the torus has Betti numbers $b_0 = 1, b_1 = 2, b_2 = 1$, so the weak Morse inequalities tell us that the decomposition we found in Example 3.22 is the best possible – no decomposition with fewer cells can exist.

As proven by Benedetti in [Be10a], given a fixed triangulation of a manifold X, smooth Morse theory generally approaches closer to the above bounds than discrete Morse theory applied to the given triangulation. However, as Benedetti and others demonstrate in later works, if we are allowed to refine a given triangulation, discrete Morse theory reaches as close to these bounds as Morse theory [Be10b], and may even reach closer [AB11].

We continue with a number of examples demonstrating the use of discrete Morse theory, in which we focus on constructing acyclic partial matchings of the Hasse diagram to determine simplified cell structures.

5 Topological examples

xkcd.com/982

5.1 The k-skeleton of the n-simplex

Definition 5.1. Let X be a simplicial complex of dimension n. The *k-skeleton* $X^k \subset X$ is the subcomplex consisting of all simplices up to and including those of dimension k.

Theorem 5.2. *The k-skeleton of the n-simplex is homotopy equivalent to $\bigvee_{\binom{n}{k+1}} S^k$.*

Proof. The k-skeleton of the n-simplex on vertices $v_0, ..., v_n$ consists of the $(n+1)$ 0-simplices $[v_i]$, the $\binom{n+1}{2}$ 1-simplices $[v_{i_1}v_{i_2}]$, ... , up to the $\binom{n+1}{k+1}$ k-simplices $[v_{i_1}v_{i_2}...v_{i_{k+1}}]$. We will assume without loss of generality that each simplex is written with its vertices labeled in ascending order.

We construct a matching as follows: Choose a 0-simplex to be left out of the matching, say v_0. Then match each remaining 0-simplex with the 1-simplex that connects it with v_0; i.e. for each $i \neq 0$, we create the pair $\{[v_i], [v_0v_i]\}$.

Then all 1-simplices of the form $[v_0v_i]$ are paired, leaving us only with 1-simplices of the form $[v_iv_j], i, j \neq 0$. Thus we can pair each such $[v_iv_j]$ with the 2-simplex $[v_0v_iv_j]$.

Now all remaining unpaired 2-simplices are of the form $[v_iv_jv_k], i, j, k \neq 0$, so we can pair each such simplex with the 3-simplex $[v_0v_iv_jv_k]$. Continuing inductively, we go on until we have paired all $(k-1)$-simplices of the form $[v_{i_1}...v_{i_k}], i_1, ..., i_k \neq 0$ with k-simplices $[v_0v_{i_1}...v_{i_k}]$. We have now paired all simplices in the graph exactly once except for the 0-simplex $[v_0]$ and the k-simplices $[v_{i_1}...v_{i_{k+1}}]$ where $i_1, ..., i_{k+1} \neq 0$.

This matching is acyclic. As demonstrated in Theorem 3.14, any cycle must be contained within two consecutive levels of the graph; so we assume there exists a cycle including the sequence $\alpha_1, \beta_1, \alpha_2$ where $\dim \beta_1 = d = \dim \alpha_i + 1$ for some d. Since $\{\alpha_1, \beta_1\} \in P$, α_1 does not include v_0, so we can write $\alpha_1 = [v_{i_1}...v_{i_d}], i_1...i_d \neq 0$. Then $\beta_1 = [v_0v_{i_1}...v_{i_d}]$, and $\alpha_2 = [v_0v_{i_1}...v_{i_{j-1}}v_{i_{j+1}}...v_{i_d}]$ for some j. But α_2 includes v_0, so it has no upward arrows emanating from it, and we cannot continue on to any β_2. Thus the matching is acyclic.

How many k-simplices were left unpaired? The unpaired simplices are all of the form $[v_{i_1}...v_{i_{k+1}}]$, $i_1, ..., i_{k+1} \neq 0$; that is, they are given by $(k+1)$ of the n vertices $v_1, ..., v_n$. Thus there are $\binom{n}{k+1}$ unpaired k-simplices, along with the 0-simplex v_0.

So the k-skeleton of the n-simplex is homotopy equivalent to a complex consisting of one 0-simplex and $\binom{n}{k+1}$ k-simplices; by Proposition 2.8, it is thus homotopy equivalent to $\bigvee_{\binom{n}{k+1}} S^k$. ■

As we can see from the weak Morse inequalities, the discrete Morse function we have just found is perfect – no other function could give us fewer critical simplices.

In the above pairing, we originally left out one 0-simplex, thus we had

$$\binom{n+1}{1} - 1$$

0-simplices to pair with 1-simplices, leaving

$$\binom{n+1}{2} - \left(\binom{n+1}{1} - 1\right)$$

1-simplices to pair with 2-simplices, etc. Thus the number of leftover k-simplices is

$$\binom{n+1}{k} - \binom{n+1}{k-1} + ... + (-1)^k \binom{n+1}{0}$$

which by combinatorics equals $\binom{n}{k+1}$.

As we can see, $\binom{n}{k+1}$ is the reduced Euler characteristic (denoted $\overline{\chi}$) of the k-skeleton of the n-simplex (actually, it is $(-1)^{k+1} \cdot \overline{\chi}$). This result follows from Theorem 4.12, which stated that

$$\sum_{i=0}^{\dim X} (-1)^i n_i = \sum_{i=0}^{\dim X} (-1)^i b_i$$

In fact, whenever we are able to reduce a complex to a wedge of spheres, so that we have some number of cells of dimension n and a single cell of dimension 0, their alternating sum must equal $(-1)^n \chi(X)$. Thus the number of cells of dimension n is $\overline{\chi}(X)$. Intuitively, we began with some triangulation of X, from which we can find $\chi(X)$. But $\chi(X)$ is invariant with respect to the structure we place on X, and indeed every simplicial collapse eliminates two cells of consecutive dimension, leaving $\chi(X)$ unchanged. Thus in our final structure, the entire value of $\chi(X)$, less one 0-simplex, is found among cells of dimension n.

The fact that we have $\overline{\chi}$ rather than χ arises from a subtlety that we have neglected until this point. We have chosen the convention that $\emptyset$ is not a subset of a simplicial complex X and cannot be included in a pairing; but if we consider $\emptyset$ to be a (-1)-simplex of every simplicial complex, then instead of choosing a 0-simplex to leave out of the pairing, we can instead pair that 0-simplex with $\emptyset$. This removes a single, previously critical 0-simplex from our structure, so our new result tells us about the reduced homology of the structure rather than its homology, and thus in the Morse inequalities we must replace χ with $\overline{\chi}$.

5.2 The flag complex on $\mathbb{F}_2^3$

Definition 5.3. Let V be an n-dimensional vector space. A *k-flag* of V, denoted f_k, is a set of vector spaces (sometimes referred to as a *chain*)

$$\{v_{i_1}, ..., v_{i_k} | 0 \subsetneq v_{i_1} \subsetneq ... \subsetneq v_{i_k} \subsetneq V\}$$

where v_{i_j} has dimension i_j.

For shorthand, we will often leave off the 0- and n-dimensional vector spaces, which are the same for all flags, and write just

$$f_k = v_{i_1} \subsetneq ... \subsetneq v_{i_k}$$

In the case of a vector space V over a finite field, the flag complex of V consisting of all flags, denoted $F(V)$, can be represented as a finite simplicial complex: we consider each k-flag as a $(k-1)$-simplex, and define the face relation by $f_j \in \partial f_k$ if $f_j \subsetneq f_k$.

We can analyze the homotopy type of the flag complex using discrete Morse theory. For purposes of illustration, we will first explore the flag complex on the 3-dimensional vector field over $\mathbb{F}_2$, and then extrapolate our method to the more general case. Although these results are known, the following proof is original.

Theorem 5.4. *Let $V = \mathbb{F}_2^3$. Then $F(V)$ is homotopy equivalent to $\bigvee_8 S^1$.*

Proof. We first establish some notation. V is spanned by the three standard basis vectors

$$e_1 := [1,0,0], e_2 := [0,1,0], e_3 := [0,0,1]$$

We note that there are seven lines contained in V:

$$e_1, e_2, e_3, e_1+e_2, e_1+e_3, e_2+e_3, e_1+e_2+e_3$$

Each plane is determined by two lines, and each line is contained in three planes, so there are $\binom{7}{2}/3 = 7$ planes:

$$span\{e_1, e_2\}$$

$$span\{e_1, e_3\}$$

$$span\{e_1, e_2+e_3\}$$

$$span\{e_2, e_3\}$$

$$span\{e_2, e_1+e_3\}$$

$$span\{e_3, e_1+e_2\}$$

$$span\{e_1+e_2, e_1+e_3\}$$

The 0-simplices of our complex are given by $f_1 = v_i$ $(i = 1, 2)$, where v_i is one of the 7 lines or 7 planes, and the 1-simplices are given by $f_2 = v_1 \subset v_2$. The faces of the 1-simplex $v_1 \subset v_2$ are the 0-simplices v_1 and v_2.

The seven lines and seven planes give us fourteen 1-flags. Since each line is contained in three planes (or equivalently, each plane contains three lines), there are a total of twenty-one 2-flags, namely:

$$e_1 \subset span\{e_1, e_2\}$$

$$e_2 \subset span\{e_1, e_2\}$$

$$e_1+e_2 \subset span\{e_1, e_2\}$$

$$\dots$$

$$e_1+e_2 \subset span\{e_1+e_2, e_1+e_3\}$$

$$e_1+e_3 \subset span\{e_1+e_2, e_1+e_3\}$$

$$e_2+e_3 \subset span\{e_1+e_2, e_1+e_3\}$$

This structure can be represented by the Heawood graph, as given below, with circles representing 1-flags and lines representing 2-flags. We use black circles for 1-flags corresponding to lines, labeled accordingly, and white circles for 1-flags corresponding to planes, consisting of the lines adjacent to (and thus included in) them:

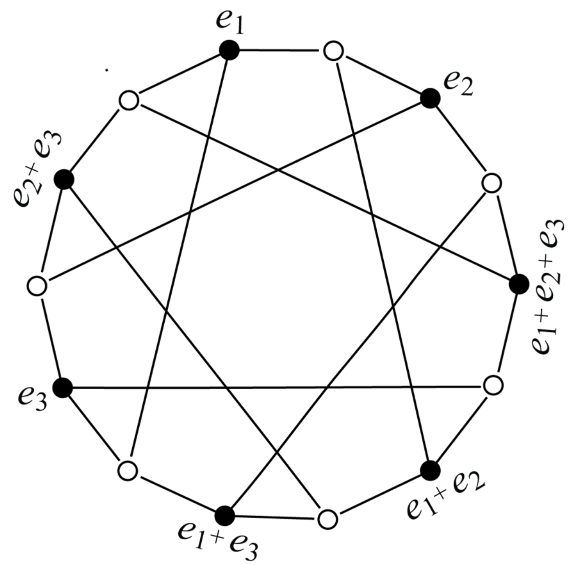

We want to construct a pairing P with no duplicates such that the resulting modified Hasse diagram is acyclic. Our approach is motivated by Brown's analysis of the flag complex in terms of *buildings* [Br89], which we will summarize briefly: Brown first sets aside a maximal flag, i.e. one containing a vector space of each dimension $i = 1, ..., n-1$:

$$f_n^0 := u_1 \subset ... \subset u_{n-1}$$

We will use the natural ordering of our basis vectors, by which $e_1 < e_2 < e_3$ and $e_1 + e_2 < e_3$, and choose our u_i to be as small as possible, so for $\mathbb{F}_2^3$ we take $u_1 = e_1, u_2 = span\{e_1, e_2\}$. Brown defines a relation *opposite* on maximal flags $v_1 \subset ... \subset v_{n-1}$:

$$v_1 \subset ... \subset v_{n-1} \text{ is opposite } v_1' \subset ... \subset v_{n-1}' \text{ if } v_i \oplus v_{n-i}' = V \text{ for all } i = 1, ..., n-1$$

Let $S := \{f_n | f_n \text{ not opposite } f_n^0\}$ be the set of all maximal flags *not* opposite to f_n^0. Then we can define the subcomplex $F' \subset F$

$$F' = \bigcup_S \bigcup_{f_k \subset f_n} f_k$$

i.e. the subcomplex consisting of all maximal flags that are not opposite to f_n^0 as well as all of their faces. (Note that $f_k \subset f_n$ includes all faces of f_n along with f_n itself.) Brown shows that F' is contractible; our goal will be to show this result using discrete Morse theory.

Thus, we want to construct a pairing P that pairs all flags $f_n \in S$ and all faces $f_k \in \partial f_n$ for all $f_n \in S$, leaving as critical simplices just the flags $f_n' \notin S$, as well as the 1-flag u_1 (to which F' will contract).

We can break the 2-flags in S into three categories, based on their "distance" from f_2^0: Either

a) $v_1 = u_1$ and $v_2 = u_2$

b) $v_1 = u_1$ or $v_2 = u_2$ (but not both)

c) $v_2 \supset u_1$ or $v_1 \subset u_2$ (but $v_1 \neq u_1, v_2 \neq u_2$)

The 2-flags not in S, which are opposite to f_2^0, fall into none of these categories. As suggested above, these categories distinguish simplices based on their distance from f_2^0, where simplices of category a) are equal to f_2^0; simplices of category b) are adjacent to those of category a), in that they share a codimension-1 face; and simplices of category c) are adjacent to those of category b). The remaining, opposite, simplices are adjacent to those of category c).

We pair the 2-flags of the above categories according to the following rules:

a) Since u_1 must be left unpaired, pair $\{v_2, v_1 \subset v_2\}$.

b) Let $i \in \{1, 2\}$ be such that $v_i \neq u_i$. Pair $\{v_i, v_1 \subset v_2\}$.

c) Let $i, j \in \{1, 2\}$ (with $i \neq j$) be such that $v_i \not\subset u_j$ and $v_i \not\supset u_j$; equivalently, $v_i \oplus u_j = V$. Pair $\{v_i, v_1 \subset v_2\}$.

We want this pairing to be a bijection between the $f_2 \in S$ and the $f_1 \neq u_1$; i.e. it includes all of the simplices we desire to include, with no duplicates. We can check this by exhibiting the inverse of the above rules, matching the desired 1-flags to unique 2-flags. (For this direction, we enumerate the possibilities according to different, more convenient categories.)

- For the 1-flag $v_2 = u_2$, pair $\{v_2, u_1 \subset v_2\}$.
- For a 1-flag $v_1 \neq u_1$, pair v_1 with the smallest $v_2 \supset v_1$, i.e. $\{v_1, v_1 \subset \{v_1, e_1, v_1 + e_1\}\}$
- For a 1-flag $v_2 \neq u_2$, pair v_2 with the smallest $v_1 \subset v_2$, i.e. either e_1, e_2, or $e_1 + e_2$. (Each $v_2 \neq u_2$ must contain one of these, since at most two lines in any plane can have an e_3 component; and no $v_2 \neq u_2$ can contain more than one, since the sum of any two is the third.)

By inspection, these two formulations of the pairing are indeed inverses. Now, we want to show that the modified Hasse diagram resulting from this pairing is acyclic. Any cycle must contain a sequence of flags f_1, f_2, f_1', f_2', where $\{f_1, f_2\}, \{f_1', f_2'\} \in P$ and $f_1' \neq f_1 \in \partial f_2$. We analyze the three possible cases:

- If $f_2 = v_1 \subset v_2$ is of category a), then the only downward arrow takes us to $v_1 = u_1$, which is unpaired and thus has no upward arrows.
- If $f_2 = v_1 \subset v_2$ is of category b), then the only downward arrow takes us to $v_i = u_i$. If $v_1 = u_1$, again this is unpaired and has no upward arrows. If $v_2 = u_2$, it is paired with $u_1 \subset u_2$, i.e. a 2-flag of category a).
- If $f_2 = v_1 \subset v_2$ is of category c), then the only downward arrow takes us to v_i, either $v_1 \subset u_2$ or $v_2 \supset u_1$. Since u_2 and u_1, respectively, are then the smallest 1-flags such that $v_i \subset v_2$ or $v_1 \subset v_i$, v_i is paired with a 2-flag of category b).

Thus no matter the category of f_2, we know that f_2' must be of an earlier category (or no valid f_2' exists). Since this holds true at all points within a path, no path can be a cycle.

Thus we have made a valid pairing, and we are left with eight unpaired 2-flags:

$$e_1 + e_2 + e_3 \subset span\{e_2, e_1 + e_3\}$$

$$e_1 + e_2 + e_3 \subset span\{e_1 + e_2, e_3\}$$

$$e_2 + e_3 \subset span\{e_2, e_3\}$$

$$e_2 + e_3 \subset span\{e_1 + e_2, e_1 + e_3\}$$

$$e_1 + e_3 \subset span\{e_2, e_1 + e_3\}$$
$$e_1 + e_3 \subset span\{e_1 + e_2, e_1 + e_3\}$$
$$e_3 \subset span\{e_2, e_3\}$$
$$e_3 \subset span\{e_1 + e_2, e_3\}$$

We return to our representation via the Heawood graph:

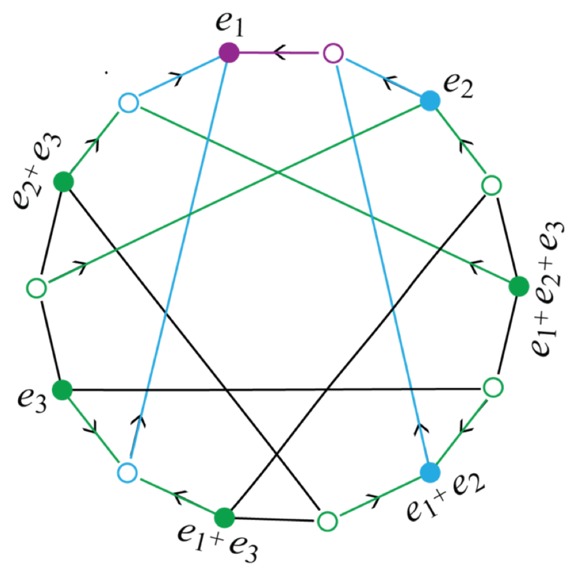

As usual, arrows indicate pairings. We can consider our pairing as being performed in steps, category by category. Simplices in purple are considered first, and are paired (or in the case of e_1, left aside as unpaired) when considering category a). Simplices in blue are adjacent to purple simplices, and are paired next, in category b); note that the arrows from blue simplices take us directly to purple simplices. Simplices in green are adjacent to blue simplices, and are finally paired when considering category c); these arrows point to blue simplices. We are now left with only 1-simplices – those opposite to f_2^0, in Brown's description – which remain unpaired, and are shown in black. The colored simplices form the subcomplex F' that is contractible, leaving behind eight 1-simplices that all meet at e_1, to which the rest of the complex has contracted. Thus the flag complex is indeed homotopy equivalent to a wedge of eight 1-spheres, as desired. ■

As in our previous example, the resulting number of spheres is in fact equal to $(-1)^n \cdot \overline{\chi}(F(V))$: 8 = 21-14+1. We note that this discrete Morse function is also perfect.

5.3 The general flag complex and lexicographic discrete Morse functions

We now examine a generalization of the above technique to the flag complex on an n-dimensional vector space over a finite field $\mathbb{F}_q$. In the general case, the order in which we pair maximal flags and their subsets is less obvious, so we turn to Babson and Hersh's lexicographic discrete Morse functions, based on lexicographic orders, as described in [BH04].

The following lemma, essential to lexicographic discrete Morse functions, allows us to construct individual pairings P_i on each subset X_i of X and combine them to create a pairing P on all of X corresponding to a discrete Morse function. This result is widely used and has been discovered independently by several mathematicians; in particular, it appears in Hersh [He05] as well as Jonsson [Jo08]. We present an adaptation of Hersh's formulation and proof:

Lemma 5.5. *Let X be a cell complex that decomposes into collections X_i of cells, indexed by the elements i in a partial order I that has a unique minimal element i_0. Assume that every cell belongs to exactly one X_i, and that for each $i \in I$, $\cup_{j \leq i} X_j$ is a subcomplex of X. For each $i \in I$, let P_i be an acyclic pairing on the cells of X_i. Then $\cup_{i \in I} P_i$ is an acyclic pairing on X.*

Proof. Let H be the modified Hasse diagram associated with $\cup_{i \in I} P_i$. By construction, no cell $\alpha \in X_i$ is paired with a cell β outside of X_i, so H contains no upward-oriented edges between cells in different components. Suppose we have a downward-oriented edge from a cell $\beta \in X_i$ to a cell $\alpha \in X_j$, $i \neq j$. Since $\cup_{j \leq i} X_j$ is a subcomplex of X, this implies $j < i$. Since this is true of every downward arrow between different components at every step of the directed path, the path can never return to X_i, so our pairing is acyclic. □

Theorem 5.6. *$F(\mathbb{F}_q^n)$ is homotopy equivalent to a wedge of $\overline{\chi}(F(\mathbb{F}_q^n))$ spheres of dimension $(n-2)$.*

Proof. First, we create a lexicographic order; the following order was suggested by Hersh [He12]. For the vector space $\mathbb{F}_q^n$, we take the standard basis $e_1, ..., e_n$. We can represent each maximal flag $f_{n-1} = 0 \subset v_1 \subset ... \subset v_{n-1} \subset V$ by an $n \times n$ matrix, constructed as follows: The first i columns of the matrix give a basis for v_i. We choose these columns, left to right, so that for each column the highest 1 (i.e. the 1 with highest row-index, furthest down in the matrix) has as high a row-index as possible. Each such leading 1 has only 0s on its right; subject to this constraint, we pick the next column so that its highest 1 has the highest possible row-index. We label a matrix by listing the index of this highest-indexed 1 in the ith column for $i = 1, ..., n$.

For example, in the case of $\mathbb{F}_2^3$, we would represent the flag $f_2 = e_1 \subset span\{e_1, e_2\}$ by the matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

giving rise to the labeling 123; the flag $f_2 = e_1 + e_3 \subset span\{e_1 + e_2, e_1 + e_3\}$ would be represented by the matrix

$$\begin{bmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}$$

giving rise to the labeling 321.

We can label each maximal flag by the label for its matrix, giving us a lexicographic ordering on the maximal flags. Note that this order has a unique minimal element. The smallest label is the only one that is strictly increasing, and a strictly increasing label can correspond only to the identity matrix: our construction of the label requires that all values below the diagonal be 0, and the condition that all leading 1s have only 0s to their right requires that all values above the diagonal be 0 as well. However, in general, different flags may have the same label, so this labeling defines a partial order: let $l = l_1...l_n, l' = l'_1...l'_n$; then $l < l'$ if, at the first index i where l and l' differ, $l_i < l'_i$. We will index these labels by a set I.

For shorthand, let $X := F(\mathbb{F}_q^n)$. For every maximal flag f_{n-1}^i, where $i \in I$ indicates the flag's index in our partial order, we create a subset X_i including f_{n-1}^i and all of its faces except those belonging to other maximal flags f_{n-1}^j with $j < i$. This construction mirrors that of Babson and Hersh in [BH04]. Formally, we have

$$X_i := f_{n-1}^i \setminus \cup_{j<i} f_{n-1}^j$$

By construction, X is the disjoint union of these sets X_i: every simplex in X is contained in some maximal face, and thus contained in some X_i, and we avoid overlap by constructing these sets in the given order and never including a face in X_i if it was already included in some previous X_j.

We note that for all $i \in I$, $\cup_{j \leq i} X_j$ includes all faces of each X_j, since X_j itself is missing only those faces that are included already in a lower-labeled simplex. Thus each such $\cup_{j \leq i} X_j$ is a subcomplex of X. Thus our complex X and partial order I fulfill the conditions of Lemma 5.5, so we just need to find an acyclic pairing on each X_i. We do so as follows:

First, if X_i contains no flags f^i_{n-2} – i.e. codimension-1 faces of the maximal flag f^i_{n-1} – then it contains no flags of any lower dimension either, since all faces of lower dimension are contained in some codimension-1 face. Thus, in this case, $X_i = \{f^i_{n-1}\}$, and our pairing P_i is empty, giving us f^i_{n-1} as a critical simplex.

Otherwise, X_i contains some codimension-1 face of f^i_{n-1}, a flag f^i_{n-2}. Let v_i be the vector space such that $v_i \in f^i_{n-1}$ but $v_i \notin f^i_{n-2}$. We create a pairing P_i by including the pair $\{f_k, f_k \cup \{v_i\}\}$ for all $f_k \in X_i$ such that $f_k \not\ni v_i$.

Firstly, $f_k \cup \{v_i\}$ is a valid $(k+1)$-flag: since f^i_{n-1} is isomorphic to an $(n-1)$-simplex, if we take any face of f^i_{n-1} and add another vertex, we get another face of f^i_{n-1}. In addition, every face of f^i_{n-1} is included in some such pair, because if we remove the vertex v_i from any face we get another valid face of f^i_{n-1}. The exception is the 1-flag v_i itself, which this rule would pair with $\emptyset$, but this is not a concern: if any codimension-1 face of f^i_{n-1} is paired already, v_i is included in this face and thus is already paired. If no codimension-1 face of f^i_{n-1} is already paired, it must be that $i = i_0$, in which case we let v_i be our single unpaired 1-flag in X.

We need to show that we never try to pair an element of X_i with a face of f^i_{n-1} that is already included in X_j for some $j < i$. Consider a flag $f_k \in \partial f^i_{n-1}$. Assume $f_k \in X_i$, i.e. $f_k \notin \partial f^j_{n-1}$ for any $j < i$, and consider the two cases: $v_i \in f_k$ or $v_i \notin f_k$. In the first case, since $f_k \notin \partial f^j_{n-1}$, neither is $f_k / \{v_i\}$, since $f_k / \{v_i\} \subset f_k$. In the second case, if $f_k \notin \partial f^j_{n-1}$, neither is $f_k \cup \{v_i\}$: if it were, since $f_k \subset f_k \cup \{v_i\}$ this would force $f_k \in \partial f^j_{n-1}$ as well. Thus for every potential pair $\{f_k, f_k \cup \{v_i\}\}$, one member of the pair is in X_i if and only if the other is as well.

Finally, P_i is acyclic: The complete flag f^i_{n-1} is isomorphic to the standard $(n-2)$-simplex, so that X_i is isomorphic to a subset of the simplex; and the pairing we have given, in which all faces collapse toward a single vertex v_i, is isomorphic to a subset of the pairing we constructed in Example 5.2. Since this original pairing was acyclic, so is P_i.

Thus, by Lemma 5.5, $P := \cup P_i$ is an acyclic pairing on X. In addition, it includes all cells except v_{i_0} and some number of maximal faces f^i_{n-1}, i.e. $(n-2)$-simplices. So, as determined above, this number of maximal faces is $\overline{\chi}(\mathbb{F}_q^n)$, and so $F(\mathbb{F}_q^n)$ is homotopy equivalent to a wedge of $\overline{\chi}(F(\mathbb{F}_q^n))$ spheres of dimension $(n-2)$.

■

As Hersh has shown [He12], the critical cells are contributed by those flags corresponding to the minimal element and maximal element of I. The single flag corresponding to i_0 contributes a 0-cell, as noted in our construction above, and the flags corresponding to the maximal element are exactly those that contribute a maximal cell. Thus we need only count the number of flags corresponding to the maximal element of I, i.e. the strictly decreasing label. Any matrix with such a label has 1s all along the minor diagonal; by our construction of the matrix, all entries below the minor diagonal must be 0. There are no restrictions on the entries above the minor diagonal, so each of these entries may take on any of q values; there are $\frac{n(n-1)}{2}$ such entries, so the total number of matrices with this label is $q^{n(n-1)/2}$. Thus, $F(\mathbb{F}_q^n)$ is in fact homotopy equivalent to a wedge of

$q^{n(n-1)/2}$ spheres of dimension $(n-2)$.

5.4 Further examples

We can combine the previous examples to examine the k-skeleton of $F(\mathbb{F}_q^n)$:

Corollary 5.7. *The k-skeleton of $F(\mathbb{F}_q^n)$ is homotopy equivalent to a wedge of*

$$\sum_{i=0}^{k}(-1)^i c_i$$

k-spheres.

Proof. We begin with the pairing P constructed above, which pairs all simplices except one 0-simplex and some number of $(n-2)$-simplices. Then for $i = n-2, ..., k+1$ we remove all pairs that include simplices of dimension i from P. After the final stage, we have a pairing on only the simplices of dimension 0 through k, which, as above, leaves one 0-simplex unpaired along with $\sum_{i=0}^{k}(-1)^i c_i$ k-simplices. □

The matchings that we have created in our examples in this section share a similar flavor, due to a common feature of the structures we have encountered: they can all be regarded as partially ordered sets whose groups of automorphisms act transitively on the elements. Thus, we can employ a single general rule for all simplices of a certain form, and these rules behave the same way regardless of our choice of "special" simplices.

Other common mathematical structures share similar underlying structures, as revealed by their homotopy types. For example, Wachs demonstrates in [Wa05] that the partially ordered set of partitions of $\{1, ..., n\}$, ordered by merging blocks, also has the homotopy type of $\bigvee_{(n-1)!} S^{n-3}$. Another example with such symmetry is the space of not-connected graphs on n vertices, as we will see in the next section.

6 Evasiveness: an application of discrete Morse theory

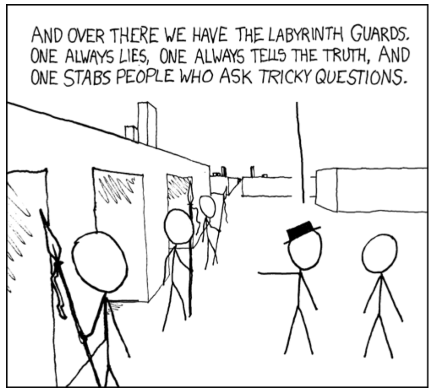

xkcd.com/246

Lest the reader think that discrete Morse theory is of interest only to mathematicians, we present a final example demonstrating its relevance to the general question of determining global information from local data. Imagine a game with two players, a hider and a guesser, and an n-dimensional simplex X with a subcomplex $Y \subset X$, where Y is known to both players. The hider picks a face $\sigma \in X$ (we consider $\emptyset \in X$ to be a valid face), and the guesser's goal is to determine whether $\sigma \in Y$, by asking questions of the form "is v_i in σ?" one by one for different vertices v_i. The guesser's algorithm can depend on responses to previous questions, but it must be deterministic. The guesser wins if she can determine whether $\sigma \in Y$ in fewer than $(n+1)$ questions – that is, without asking about every vertex in X – for every simplex σ.

Definition 6.1. A subcomplex Y is *evasive* if no winning strategy exists – i.e. for any guessing algorithm A there exists some $\sigma \in X$ such that $(n+1)$ questions are required to determine whether $\sigma \in Y$. If a winning strategy does exist, Y is *nonevasive*.

This game describes a rather general search problem, with applications to graph theory, complexity theory, and combinatorics. We will explore one such application below, returning to our example of the space of not-connected graphs.

Definition 6.2. An *evader* of a guessing algorithm A on an evasive subcomplex Y is a face $\sigma \in X$ such that when questions are asked as dictated by A, all $(n+1)$ questions are required to determine whether $\sigma \in Y$.

We note that evaders occur in pairs: if the guessing algorithm A requires $(n+1)$ questions to determine whether $\sigma \in Y$, this means that after the nth question the algorithm has narrowed down the possibilities to two, either $\sigma = \alpha \in Y$ or $\sigma = \beta \notin Y$, where β is a simplex containing all the vertices of α plus an additional v_i, so that α is a codimension-1 face of β. Thus both α, β are evaders of A.

We can take these pairs of evaders to form a pairing P, noting the technicality that if the empty set is an evader of A, then it will be paired with some 0-simplex v_j. According to the conventions we established previously, a matching cannot include the empty set. For this example, it will be convenient to deviate from this convention, pairing the empty set with a 0-simplex that, according to our earlier convention, would have been set aside as unpaired. This deviation is another way of arriving at the reduced Euler characteristic, and reduced homology, that we have encountered previously.

Theorem 6.3. *The pairing P constructed above corresponds to a discrete Morse function on X.*

Proof. Firstly, by our construction, the pairs are clearly disjoint, since the $(n+1)$st question cannot distinguish between more than two distinct simplices.

Now we need that P is acyclic. Let A be a guessing algorithm for a subcomplex $Y \in X$. For every simplex $\sigma \in X$, let the sequence $s(\sigma) = \{i_0, ... i_m\}$ give the indices of those questions dictated by A to which the answer is affirmative, in increasing order.

We impose a reverse-lexicographic ordering on these sequences, which hence defines an ordering on the simplices: let $s(\alpha) = \{i_0, ... i_m\}, s(\beta) = \{j_0, ... j_{m'}\}$. If $s(\beta)$ is an extension of $s(\alpha)$, i.e. $m' > m$ and for all $k \leq m$ we have $s(\alpha) = s(\beta)$, then we say $\beta < \alpha$. If neither $s(\beta)$ nor $s(\alpha)$ is an extension of the other, then we again order them reverse-lexicographically: if for some $k < m, m'$, we have $i_l = j_l$ for all $l < k$ and $i_k < j_k$, then $\beta < \alpha$.

We consider a potential cycle containing the sequence $\{\alpha_0, \beta_0, \alpha_1\}$. Since $\{\alpha_0, \beta_0\} \in P$, if $s(\alpha_0) = \{i_0, ... i_m\}$, then $s(\beta_0) = \{i_0, ... i_m, n+1\}$: these simplices are undistinguished until the final question, at which point A asks about the single vertex v_i such that $v_i \in \beta$ but $v_i \notin \alpha$. Since $\alpha_1 \in \partial\beta_0$, $s(\alpha_1)$ is a subset of $s(\beta_0)$, missing some single index $i_j \neq n+1$ (since $\alpha_1 \neq \alpha_0$). Thus $s(\alpha_1) = \{i_0, ..., i_{j-1}, i_{j+1}, ... i_m, n+1\}$. But $i_{j+1} > i_j$, so by our ordering $\alpha_1 < \alpha_0$. Since this holds for any α_0, α_1, every step in a path is decreasing, and thus no path can be a cycle. Thus P is acyclic, as required. ∎

The main result of this section, as demonstrated by Forman [Fo00] and others, is as follows:

Theorem 6.4. *If Y is nonevasive then Y is homotopy equivalent to a point.*

Proof. Continuing with our P constructed above, we take the pairing $P' \subset P$ on Y consisting of all pairs $\{\alpha, \beta\} \in P$ such that $\alpha, \beta \in Y$. Since P is acyclic, so is P'; therefore P' corresponds to a discrete Morse function on Y. Recall that X has no critical simplices (v_j being paired with $\emptyset$), so the unpaired – i.e. critical – simplices of Y are those paired with simplices in $X \backslash Y$, i.e. those pairs of evaders of A.

In particular, this implies that if Y is nonevasive, having an algorithm A with no evaders, then there exists a pairing P that includes all simplices of Y. Thus, adding back one 0-simplex to our count of critical simplices, we find that Y is homotopy equivalent to a point. ∎

The paradigm of evasiveness can be applied to a family of problems in graph theory: determining whether a graph satisfies a given *monotone decreasing graph property* by checking edges one by one. A graph property is a subset of G_n, the space of all graphs on n vertices, such that membership in the subset depends only on the isomorphism type of the graph, i.e. the labeling of vertices is irrelevant. A graph property is monotone decreasing if whenever a given graph has the property, so do all of its subgraphs. Formally:

Definition 6.5. A subset $P \subset G_n$ is a *graph property* if for all $G, G' \in G_n$ such that G, G' are isomorphic, then $G \in P$ if and only if $G' \in P$.

Definition 6.6. A graph property P is *monotone decreasing* if $P(G)$ implies $P(G')$ for all $G' \subset G$.

A conjecture attributed to Karp suggests that all non-trivial monotone graph properties are evasive [Fo00]. We offer as an example one such monotone decreasing graph property: the set of not-connected graphs on n vertices, N_n.

We provide a sketch of Forman's analysis of the homotopy type of N_n, as given in [Fo02]. The space N_n is constructed by creating a 0-simplex corresponding to each of the $\binom{n}{2}$ possible edges in a graph on n vertices, and then including all possible i-simplices to represent each of the possible graphs with $i+1$ edges, up to $i = \binom{n}{2} - 1$.

Theorem 6.7. *N_n is homotopy equivalent to $\bigvee_{(n-1)!} S^{n-3}$.*

Outline of Proof: The 0-simplices are the graphs composed of a single edge, denoted $\{i, j\}$ for the two vertices i, j that the edge connects. We first consider the 0-simplex $\{1, 2\}$. Let this be our unpaired 0-simplex, and then for all pairs of graphs $G, G + \{1, 2\}$ such that $G, G + \{1, 2\}$ are both unconnected, pair them (excluding the pair $\{\emptyset, \{1, 2\}\}$). For any graph G that remains unpaired, we next consider the 0-simplices $\{1, 3\}$ and $\{2, 3\}$: at most one of these can be added to G with a result that is also unconnected and not yet paired, so we add this pair to P. At the next step, we consider all remaining unpaired graphs G combined with one of $\{1, 4\}, \{2, 4\}$, and $\{3, 4\}$, and add the appropriate pair to P. We continue inductively until we have reached vertex n. Forman demonstrates that this pairing is acyclic.

The unpaired graphs that are left are those unconnected graphs that are the union of two connected trees, with the vertices 1 and 2 in separate components, and such that the vertex numbers increase along every path emanating from the vertices 1 and 2. There are a total of $(n-1)!$ such graphs, each comprising $n-2$ edges, so they are $(n-3)$-simplices. ■

Example 6.8. N_n has the homotopy type of $\bigvee_{(n-1)!} S^{n-3}$. Thus N_n is evasive. We illustrate the case for $n = 4$:

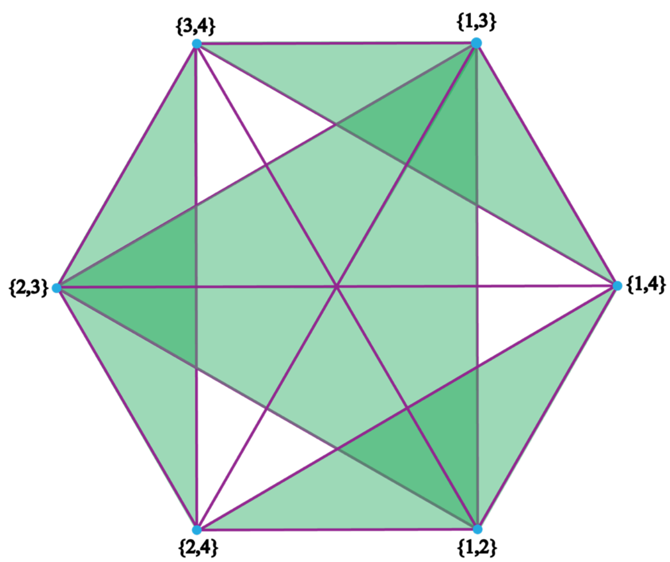

We can represent the space of all graphs on 4 vertices by a simplex on 6 vertices, taking the edges of our graph, of which there are $\binom{4}{2} = 6$, to be the vertices of our simplex. The space of not-connected graphs is a subcomplex, as given by the colored simplices above: the one-edged graphs that are not connected (i.e. all of them) are colored blue, the two-edged graphs that are not connected (again, all of them) are colored purple, and the three-edged not-connected graphs (this time, only 4 of the 20) are colored green. All graphs with four or more edges are connected (luckily for us, since they'd be difficult to color in two dimensions). As we saw earlier, this space is homotopy equivalent to a wedge of six 1-spheres, and in particular is not homotopy equivalent to a point.

The property of being not connected is thus evasive, and so any algorithm must, in at least some cases, check each of $\binom{n}{2}$ possible edges in order to decide whether a given graph is connected. This example has real-life implications: Forman in [Fo04] gives the example of trying to determine, immediately after a disaster, whether a network of communication lines remains connected, when one can only test the individual components. This analysis demonstrates that there is no guaranteed way of finding the answer without testing every individual component.

We might wonder whether it would be more fruitful to pose the question as "Is G connected?" rather than "Is G not connected?" – perhaps the space of connected graphs is homotopy equivalent to a point, even though the space of not-connected graphs is not. But the space of connected graphs is not a simplicial complex, since it does not contain all of the faces of all of its simplices: in terms of graphs, when we remove an edge from a connected graph, the result is not necessarily still connected. Thus the theorem would not apply.

This application is explored in more detail in [Fo00]. Forman demonstrates that we can determine not only whether a subcomplex is evasive, but how evasive it is – how many faces may evade even the best guessing algorithm – based on the subcomplex's homotopy type.

7 Summary

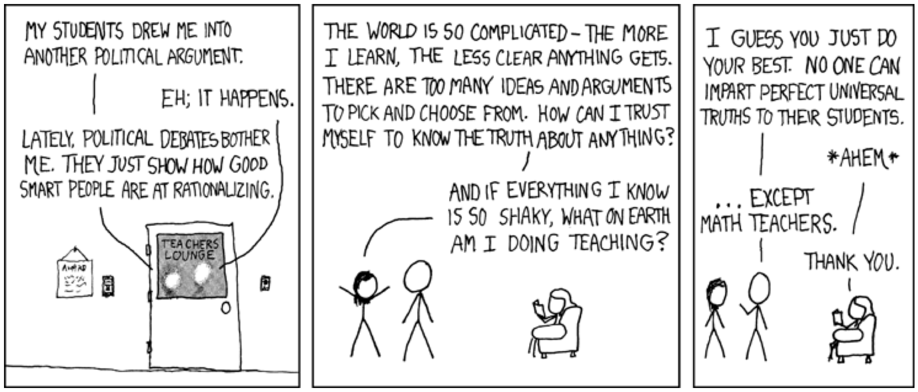

xkcd.com/263

Mathematics is often seen as a reliable but prosaic tool for finding answers. Perhaps the most fundamental question in mathematics is “How can one solve this problem?” But equally important, and more easily overlooked, is the related question “Has this problem already been solved?” Many mathematical questions appear in different guises, and recognizing an old problem in a new form is an essential – and more imaginative – task for mathematical inquiry.

Discrete Morse theory provides an approach to this latter question within the realm of topological spaces endowed with discrete structure. The theory provides a natural language with which to examine equivalences between such structures, describing many problems within algebraic topology and encompassing a great number of applications. As we saw in our study of evasiveness, questions about a space can often be reduced to questions about its topology, and so determining the space’s homotopy type reveals the answer to our original question – even when the question seemingly has little to do with topology, arising from a highly applicable, real-world problem.

In this work, the applications we have noted focus mainly on computer science and graph theory, and much interesting work in this realm has been (and continues to be) done: for example, see Shareshian’s work on 2-connected graphs [Sh01]; Babson et al. on not-i-connected graphs [BB96]; Wachs on partially ordered sets of partitions, graphs, and trees [Wa05]; and Jonsson on 3-connected and Hamiltonian graphs [Jo03]. Applications to other fields abound as well: Green applies the theory to polygons and their higher-dimensional analogs, polytopes, in [Gr11]; Vassiliev links the ideas to the study of knots in [Va93]; and Benedetti introduces a version for manifolds with boundary in [Be10a]. With such widespread realms of application, discrete Morse theory has powerful implications for our understanding of structures throughout discrete mathematics.

References

[AB11] K. Adiprasito and B. Benedetti, *Metric Geometry and Collapsibility*, ArXiv e-prints (2011).

[BB96] E. Babson, A. Björner, S. Linusson, J. Shareshian, and V. Welker, *Complexes of Not i-Connected Graphs*, Topology **38** (1996), pp. 271-299.

[BH04] E. Babson and P. Hersh, *Discrete Morse Functions from Lexicographic Orders*, Trans. Amer. Math. Soc. **357** (2005), pp. 509-534.

[Be10a] B. Benedetti, *Discrete Morse Theory for Manifolds with Boundary*, ArXiv e-prints (2010).

[Be10b] ___, *Discrete Morse Theory Is At Least As Perfect As Morse Theory*, ArXiv e-prints (2010).

[Br89] K. Brown, *Buildings.* Springer-Verlag: New York, 1989.

[Co73] M. Cohen, *A Course in Simple-Homotopy Theory.* Springer-Verlag: New York, 1973.

[Fo98] R. Forman, *Morse Theory for Cell Complexes*, Adv. Math. **134** (1998), pp. 90-145.

[Fo00] ___, *Morse Theory and Evasiveness*, Combinatorica **20** (4) (2000), pp. 489-504.

[Fo02] ___, *A User's Guide to Discrete Morse Theory*, Sém. Lothar. Combin. **48** (2002).

[Fo04] ___, *Topics in Combinatorial Differential Topology and Geometry*, IAS/Park City Math. Ser. **13** (2003), pp. 133-205.

[Gr11] R. Green, *Morse matchings on polytopes and their subcomplexes.* University of Colorado Boulder. 11 Nov 2011.

[Ha01] A. Hatcher, *Algebraic Topology.* Cambridge University Press: Cambridge, 2001.

[He05] P. Hersh, *On Optimizing Discrete Morse Functions*, Adv. in Appl. Math. **35** (3) (2005), pp. 294-322.

[He12] ___. Email to M. Hopkins. 16 Mar. 2012.

[Jo03] J. Jonsson, *On the topology of simplicial complexes related to 3-connected and Hamiltonian graphs*, J. Combin. Theory Ser. A **104** (1) (2003), pp. 169-199.

[Jo08] ___, *Simplicial Complexes of Graphs*, Lecture Notes in Mathematics **1928** (2008), pp. 51-66.

[JP06] M. Joswig and M. Pfetsch. *Computing Optimal Discrete Morse Functions*, Discrete Math **20** (2006), pp. 11-25.

[Le02] T. Lewiner, *Constructing Discrete Morse Functions.* 10 Jul. 2002. Unpublished manuscript, Pontifical Catholic University of Rio de Janeiro.

[Le04] T. Lewiner, H. Lopes, and G. Tavares. *Applications of Forman's discrete Morse theory to topology visualization and mesh compression.* Visualization and Computer Graphics **10** (5) (2004), pp. 499-508.

[Mi63] J. Milnor, *Morse Theory.* Princeton University Press: Princeton, 1963.

[Mo34] M. Morse, *The Calculus of Variations in the Large*, Amer. Math. Soc. Colloquium Pub. **18**, Amer. Math. Soc., Providence R.I., (1934).

[Sh01] J. Shareshian, *Discrete Morse Theory for complexes of 2-connected graphs*, Topology **40** (4) (2001), pp. 681-701.

[Va93] V. Vassiliev, *Complexes of Connected Graphs*, in: The Gelfand Mathematical Seminars, 1990-1992, pp. 223-235. Birkhäuser Boston: Boston, 1993.

[Wa05] M. Wachs, *Topology of posets of partitions, graphs, and trees*. San Francisco State University. 10 Aug 2005.