Bachelor's thesis

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

Nils Mork Stene

An Introduction to Discrete Morse Theory

Bachelor's thesis in Mathematical Sciences
Supervisor: Marius Thaule
May 2022

NTNU
Norwegian University of
Science and Technology

Nils Mork Stene

An Introduction to Discrete Morse Theory

Bachelor's thesis in Mathematical Sciences
Supervisor: Marius Thaule
May 2022

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

NTNU
Kunnskap for en bedre verden

Contents

1 Introduction

Imagine a regular topographic map of some real world landscape. The contour lines on the map mark the points that are on the same height for some discrete heights, for example at 100, 150 and 200 meters, but we can imagine that the map has contour lines for any height. If we look at the region that lies below any given height, we can see that it has to be separated from the region that is above the same height by a contour line. Now, intuitively, if we disregard any notion of distance or size, as we generally do in topology, these contour lines only appear, disappear or merge at points where the gradient is zero, i.e., a point where the ground is flat. That includes peaks, basins and passes, also called maxima, minima and saddle points, respectively. The topology of the region that is below some height only changes when its boundary changes, which we know only happens at the peaks, basins and passes. Morse theory, named after Marston Morse and developed in his paper [7] from 1928, describes a rigid topological analogue of this intuition. By assigning a height function to some topological space, we can use its critical points, i.e., its maxima, minima and saddle points, to analyze its topology.

Discrete Morse theory is a discrete analogue of Morse theory, with discrete analogues to many of its theorems. It was first introduced by Robin Forman in his article [4] from 1995. Due to its discrete nature, discrete Morse theory lends itself well to computations. If we can find a discrete representation, for example a triangulation, of a topological space, discrete Morse theory can simplify the representation and tell us about its homology. Or, if we have a question about the real world and can find a way to translate the information into relations between discrete points, discrete Morse theory might answer the question. One example of this is detecting potential areas without cell phone signal coverage and is detailed in [10, pp. 2–6]. It also finds applications in reconstructing road networks from GPS traces and satellite images and reconstructing neural networks from image data as described in [3, pp. 255–258]. Another application is persistent homology, which is one of the main tools in topological data analysis. More information on the subject can be found in [10, pp. 117–148].

We start by defining simplicial complexes, the structures we will be working with, and discrete Morse functions, functions from a simplicial complex to $\mathbb{R}$. These are a bit unwieldy in practice, but are great tools for proving statements. We then introduce induced gradient vector fields. They are easier to understand and deal with, but harder to prove things with. Luckily, they have a strong connection to discrete Morse functions, meaning that we can prove statements about induced gradient vector fields using discrete Morse functions. Finally, we look at some important results in discrete Morse theory.

2 Simplicial complexes

Simplicial complexes are combinatorial data structures that says what relations exist between several elements. In the following definitions, simplices represent the relations between its elements, called vertices, in that they are all related to each other. We also define some terminology related to simplicial complexes.

Definition 2.1 (Simplex)**.** A *simplex* α is a collection of elements called vertices. Its dimension p is one less than the number of vertices. A simplex α of dimension p is denoted $\alpha^{(p)}$.

Example 2.2. Let $\alpha = \{v_0, v_1, v_2\}$. Then α has dimension $\dim(\alpha) = 3 - 1 = 2$.

We take "k-simplex" to mean "a simplex of dimension k." We can draw representations of simplices up to dimension two as shown in Figure 1.

Figure 1: A 0-, 1- and 2-simplex, respectively.

We often denote the vertices by $v_0, v_1, v_2, \dots$ and, assuming there is no confusion, the simplex $\{v_a, v_b, v_c\}$ by v_{abc}. For this reason, v_0 could refer to either the vertex v_0 or the simplex $\{v_0\}$.

Now, we introduce terminology that makes it easier to talk about simplices that are part of another simplex.

Definition 2.3 (Face, coface, codimension)**.** Let $p, q \in \mathbb{R}$ with $p < q$. A simplex $\alpha^{(p)}$ is a *face* of another simplex $\beta^{(q)}$ if $\alpha \subset \beta$. A simplex $\beta^{(q)}$ is a *coface* of another simplex $\alpha^{(p)}$ if $\beta \supset \alpha$. The *codimension* of a face or coface is the difference in dimension, i.e., $|q - p|$.

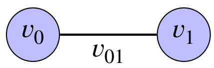

Figure 2: A collection of three simplices.

Example 2.4. In the collection of simplices represented by Figure 2, v_0 and v_1 are faces of v_{01}, as $\{v_0\} \subset \{v_0, v_1\} = v_{01}$ and $\{v_1\} \subset \{v_0, v_1\} = v_{01}$. This also means that v_{01} is a coface of both v_0 and v_1.

We now introduce the main structure we will be working with.

Definition 2.5 (Simplicial complex)**.** An *abstract simplicial complex*, which we will refer to as a *simplicial complex*, K on a collection v of elements called vertices is a collection of simplices such that

(i) every vertex is in some simplex in K, and

(ii) if a simplex $\alpha \in K$ and $\beta \subset \alpha$ then $\beta \in K$.

The first condition ensures that every vertex is part of some simplex. The second condition ensures that every face of every simplex in a simplicial complex is also in the simplicial complex. This is an important property because it means there is some connection between simplices of higher and lower dimension. In the combinatorial data structure, this means that if any amount of vertices are related, then all subsets are also related. It follows from this definition that every simplex is a subset of the vertex set v and that the collection of vertices is equivalent to the collection of 0-simplices in that for every vertex v_a there is a simplex $\{v_a\}$ and vice versa. We will therefore rarely mention the vertex set, taking it to be the collection of 0-simplices.

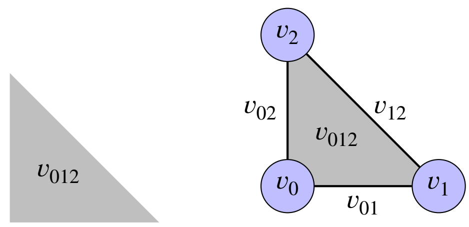

(a) Not a simplicial complex. (b) A simplicial complex.

Figure 3: Collections of simplices.

Example 2.6. The collection of simplices represented by Figure 3(a) consists of only the 2-simplex, without its faces and thus is not a simplicial complex. The collection of simplices represented by Figure 3(b) contains the 1-simplices, or lines, that are the faces of the 2-simplex, as well as the 0-simplices, or points, that are the faces of the lines. Therefore, it is a simplicial complex.

It is useful to have an upper bound on the dimension of each simplex in any given simplicial complex. Therefore, we introduce the dimension of a simplicial complex.

Definition 2.7 (Dimension)**.** The *dimension* of a simplicial complex K, $\dim(K)$, is the dimension of its largest simplex i.e., $\max_{\alpha \in K}(\dim(\alpha))$.

Example 2.8. The collection represented by Figure 2 is a simplicial complex with dimension 1, as it contains a 1-simplex and its two 0-simplex faces.

We also need to know how many simplices we have of each dimension.

Definition 2.9 (C-vector)**.** The *c-vector* of a simplicial complex states how many simplices of each dimension is in the simplicial complex, i.e., $\vec{c} = (c_0, c_1, \cdots, c_{\dim(K)})$, where $c_i = |\{\alpha^{(i)} \in K\}|$.

At this point, it might be a good idea to tie multiple concepts together in a single example.

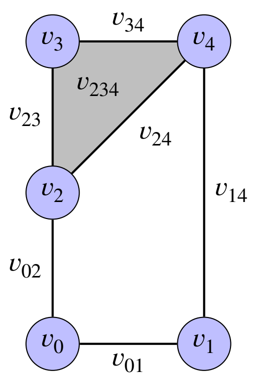

Figure 4: A simplicial complex K_5 of dimension 2.

Example 2.10. The vertices v of the simplicial complex K_5 represented by Figure 4 are $\{v_0, v_1, v_2, v_3, v_4\}$. Let

$$K_5 = \{v_0, v_1, v_2, v_3, v_4, v_{01}, v_{02}, v_{14}, v_{23}, v_{24}, v_{34}, v_{234}\}.$$

We check that K_5 is a simplicial complex. First, every vertex is also a simplex, so the first condition is satisfied. Then, we check that K_5 also contains v_{23}, v_{24} and v_{34}, the faces of the 2-simplex v_{234}. We also check for each of the 1-simplices, v_{01}, v_{02}, v_{14}, v_{23}, v_{24} and v_{34}, that their 0-dimensional faces are also in K_5. But these are the vertices that we already know are there, so the second condition is also satisfied. Thus, K_5 is a simplicial complex.

The simplex of K_5 with largest dimension is v_{234}, which has dimension 2, so the dimension of K_5 is 2. Finally, K_5 contains five 0-simplices, six 1-simplices and one 2-simplex and so, K_5 has c-vector $\vec{c} = (5, 6, 1)$.

3 Discrete Morse functions

Recall the topographic map in the introduction. Rivers, as we know, flow downhill and are often drawn as blue lines, while their direction are inferred from the contour lines. If the ground is flat, i.e., the gradient is 0, a drop of water on the ground would stay still. There are exceptions, particularly at passes, but generally, if we have a river, it must flow downhill, meaning that there are no peaks, basins or passes along the river. Now, in order to more easily calculate the simplicial homology, we want to have something like the rivers on the simplicial complexes too, but there is no intuitive height function. So, we simply give simplices a real number, obeying a couple of rules. This kind of function is called a discrete Morse function. These functions give us something called an induced gradient vector field, which has arrows that resemble flow lines, which is the direction a drop of water would go if dropped on the ground. These flow lines can be accumulated into V-paths, the discrete Morse theory equivalent of rivers, simultaneously giving us the direction of flow. These flow lines let us find critical simplices, which are the discrete Morse theory equivalent of peaks, basins and passes. While we here start by talking about the height function and work our way to the vector field, we could also have started with a vector field and found a function from it.

Definition 3.1 (Discrete Morse function)**.** Let K be a simplicial complex. A function $f : K \to \mathbb{R}$ is a *discrete Morse function* if for every simplex $\alpha^{(p)} \in K$, both

$$|\{\beta^{(p+1)} \supset \alpha^{(p)} \mid f(\beta) \leq f(\alpha)\}| \leq 1 \qquad \text{and} \qquad |\{\gamma^{(p-1)} \subset \alpha^{(p)} \mid f(\gamma) \geq f(\alpha)\}| \leq 1.$$

The two conditions are respectively saying that every simplex has no cofaces with codimension 1 of lower or equal value with at most one exception and no faces with codimension 1 of greater or equal value with at most one exception.

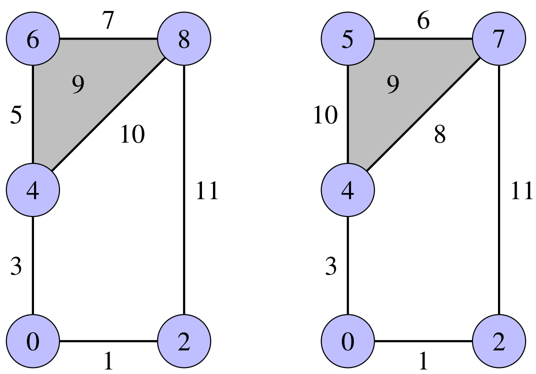

(a) $f : K_5 \to \mathbb{R}$.　　(b) $g : K_5 \to \mathbb{R}$.

Figure 5: Two discrete Morse functions on K_5.

Example 3.2. Let the numbers in Figure 5(a) represent the numbers a function f assigns to the simplices of K_5, meaning that $f(v_0) = 0$, $f(v_4) = 8$ $f(v_{23}) = 5$, $f(v_{234}) = 9$ and so on. We can check for each simplex that its faces have lower values and its cofaces have higher values with at most one exception. The 2-simplex v_{234} has three 1-simplex faces: v_{23}, v_{34} and v_{24}. These have the values $f(v_{23}) = 5$, $f(v_{34}) = 7$ and $f(v_{24}) = 10$, only one of which has higher value than $f(v_{234}) = 9$. The 1-simplex v_{34} has value $f(v_{34}) = 7$. It has two 1-simplex faces, v_3 and v_4. Only v_4, with value $f(v_4) = 8$, has a greater value than v_{34}. Its coface v_{234} also has greater value. We can check this for every simplex and conclude that f is a discrete Morse function. We can check in the same way that g, represented by the numbers in Figure 5(b), is also a discrete Morse function on K_5.

Remark 3.3. Though most examples of discrete Morse functions in this paper are injective, they do not have to be. As long as it satisfies the definition, a discrete Morse function may take the same value for any number of simplices.

(a) A discrete Morse function.　　(b) Not a discrete Morse function.

Figure 6: Two functions on K_2.

Example 3.4. Look at Figure 6. The two subfigures represent the same simplicial complex $K_2 = \{v_0, v_1, v_{01}\}$, but with two different functions. One of them is a discrete Morse function, and the other is not. The difference is that the values of the leftmost 0-simplex and the 1-simplex has been exchanged. In Figure 6(a), only the rightmost 0-simplex has a higher value than the 1-simplex, meaning that we have a discrete Morse function. In Figure 6(b), both 0-simplices have higher values than the 1-simplex, and thus the function is not a discrete Morse function.

Definition 3.5 (Critical, regular)**.** A simplex α of a simplicial complex K with a discrete Morse function f is *critical* if it has no exceptions to the rules stated in Definition 3.1 i.e., both

$$|\{\beta^{(p+1)} \supset \alpha^{(p)} \mid f(\beta) \leq f(\alpha)\}| = 0 \qquad \text{and} \qquad |\{\gamma^{(p-1)} \subset \alpha^{(p)} \mid f(\gamma) \geq f(\alpha)\}| = 0.$$

Otherwise, if it has an exception to at least one of the rules, it is *regular*. The value of a critical simplex is called a *critical value*. If a value is not critical, it is called a *regular value*.

This is the same as saying that a simplex is critical if it obeys both of the following rules:

(i) It has no faces with a higher value than itself, and

(ii) it has no cofaces with a lower value than itself.

(a) One critical value. (b) Only critical values.

Figure 7: Critical values of two discrete Morse functions on K_2.

Example 3.6. Consider the simplicial complex K_2, introduced in Example 3.4. We assign two different discrete Morse functions as shown in Figure 7. We look first at the function f, represented by Figure 7(a). The 0-simplex on the left has only the 1-simplex as a coface, which has higher value than the 0-simplex. Therefore, the 0-simplex is critical and 0 is a critical value. The 0-simplex on the right also has only the 1-simplex as a coface, but it has lower value than the 0-simplex. Therefore, the 0-simplex is regular and 2 is a regular value. The 1-simplex has two faces, one with higher value and one with lower value than the 1-simplex, and thus it is also regular and 1 is a regular value.

Similarly, we can check that K_2 with discrete Morse function g, as represented by Figure 7(b), has only critical simplices and its critical values are 0, 1 and 2.

Example 3.7. We check that K_5 with f as represented by Figure 5(a) has two critical simplices: v_0 and v_{14}. The simplex v_0 has two cofaces, v_{01} and v_{02}, where their values, $f(v_{01}) = 1$ and $f(v_{02}) = 3$, are both greater than $f(v_0) = 0$. Thus v_0 is critical, which also makes 0 a critical value. Similarly, v_{14} has two faces, v_1 and v_4, but as their values, $f(v_1) = 2$ and $f(v_4) = 8$, both are less than $f(v_{14}) = 11$, v_{14} is critical and 11 is a critical value. We verified that some of the simplices had an exception to the rules stated in Definition 3.1 in Example 3.2 and we can do this for the rest, concluding that in fact, all simplices except v_0 and v_{14} are regular. Thus all values except 0 and 11 are regular.

In the same way, we can see that K_5 with g as represented by Figure 5(b) has critical simplices v_0, v_3, v_{24} and v_{14} as well as critical values 0, 5, 8 and 11. All the other simplices and values are regular.

We often need to know if a simplex is critical or regular and what exceptions it has to the rules stated in Definition 3.1. We now present and prove a result that lets us know that any simplex can only break one of the rules. This result will prove useful when talking about induced gradient vector fields.

Lemma 3.8 (Exclusion lemma)**.** *Let $f : K \to \mathbb{R}$ be a discrete Morse function. Then no simplex can have an exception to both of the rules stated in Definition 3.1 i.e., no simplex $\alpha^{(p)} \in K$ can satisfy both*

$$|\{\beta^{(p+1)} \supset \alpha^{(p)} \mid f(\beta) \leq f(\alpha)\}| = 1 \qquad \textit{and} \qquad |\{\gamma^{(p-1)} \subset \alpha^{(p)} \mid f(\gamma) \geq f(\alpha)\}| = 1.$$

Proof. We assume that $\alpha^{(p)}$ has a coface $\beta^{(p+1)}$ such that $f(\beta) \leq f(\alpha)$ and a face $\gamma^{(p-1)}$ such that $f(\gamma) \geq f(\alpha)$. Then $\beta = \alpha \cup \{v_a\} = \gamma \cup \{v_a, v_b\}$ for some vertices v_a, v_b. But then there must exist some simplex $\delta^{(p)} = \gamma \cup \{v_a\} = \beta \setminus \{v_b\}$ i.e., δ is a face of β and a coface of γ. As f is a discrete Morse function, $\gamma \subset \delta, \alpha$ and $f(\gamma) \geq f(\alpha)$, we cannot have $f(\gamma) \geq f(\delta)$, as that would imply that γ has two exceptions to the rule that cofaces must have higher value. Therefore, we must have

$$f(\delta) > f(\gamma).$$

Similarly, since $\delta, \alpha \subset \beta$ and $f(\beta) \leq f(\gamma)$, meaning that for β, $f(\delta) \geq f(\beta)$ would be the second exception to the rule that faces must have lower value. Therefore, we must have

$$f(\delta) < f(\beta).$$

It follows that $f(\delta) > f(\gamma) \geq f(\alpha) \geq f(\beta) > f(\delta)$, which is a contradiction. Thus, no simplex can satisfy both equations. □

The following definition introduces the flow lines, as described at the start of this section.

Definition 3.9 (Induced gradient vector field)**.** Let $f : K \to \mathbb{R}$ be a discrete Morse function. The *induced gradient vector field* V_f is the collection of pairs of simplices $\{(\alpha^{(p)}, \beta^{(p+1)}) \mid \alpha \subset \beta, f(\alpha) \geq f(\beta)\}$.

We can draw these pairs as arrows from the p-simplex, the tail, to the $(p+1)$-simplex, the head, as visualised in the following example.

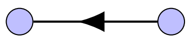

Figure 8: The induced gradient vector field on K_1 as shown in Figure 6(a).

Example 3.10. In the simplicial complex represented by Figure 6(a), the rightmost 0-simplex has higher value than the 1-simplex. Therefore, the arrow in Figure 8 goes from the rightmost 0-simplex to the 1-simplex.

A simplex being the head of an arrow means that the tail of the arrow is an exception to the first rule stated in Definition 3.1. Similarly, a simplex being the tail of an arrow means that the head of the arrow is an exception to the second rule stated in Definition 3.1. The exclusion lemma, Lemma 3.8, tells us that no simplex can be both a head and a tail, meaning that simplices are either

(i) the head of exactly one arrow, i.e., regular,

(ii) the tail of exactly one arrow, i.e., regular, or

(iii) neither a head nor a tail of any arrow, i.e., critical.

Example 3.11. Let K_5 be the simplicial complex introduced in Example 2.10 and $f : K_5 \to \mathbb{R}$ be the discrete Morse function introduced in Example 3.2. The induced gradient vector field on f, V_f is displayed in Figure 9.

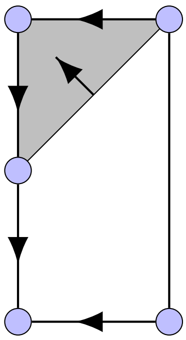

Figure 9: An induced gradient vector field on K_5.

Definition 3.12 (Discrete vector field)**.** A *discrete vector field* V on a simplicial complex K is $\{(\alpha^{(p)}, \beta^{(p+1)}) \mid \alpha \subset \beta$, each simplex in K is in at most one pair.$\}$

The discrete vector field is similar to the induced gradient vector field, as they both pair simplices together. In fact, every induced gradient vector field is a discrete vector field. This follows from the exclusion lemma, Lemma 3.8, and the definition of a discrete Morse function, Definition 3.1. The converse is not true, not every discrete vector field is an induced gradient vector field. See the counterexample in the simplicial complex represented by Figure 10.

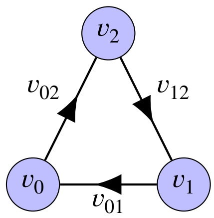

Figure 10: A discrete vector field that is not an induced gradient vector field.

No simplex is in more than one pair, so this is a discrete vector field. But if we imagine a discrete Morse function f inducing this induced gradient vector field, we find a contradiction. From the pair (v_0, v_{02}) we can deduce that $f(v_0) \geq f(v_{02})$. Since we assumed that the function is a discrete Morse function, we must have $f(v_{02}) > f(v_2)$. Thus we have

$$f(v_0) \geq f(v_{02}) > f(v_2).$$

In the same manner, we get

$$f(v_2) \geq f(v_{12}) > f(v_1) \qquad \text{and} \qquad f(v_1) \geq f(v_{01}) > f(v_0).$$

Combining these, we get $f(v_0) > f(v_0)$. This is a contradiction, meaning that the assumption that f is a discrete Morse function must be wrong. Thus Figure 10 cannot represent an induced gradient vector field. Observe that the contradiction arises from the circular chain of arrows, $\{v_0, v_{02}\}, \{v_2, v_{12}\}, \{v_1, v_{01}\}$. We formalize the idea of a chain of arrows in the following definition.

Definition 3.13 (V-path)**.** Let K be a simplicial complex and V a discrete vector field on K. A *V-path* is a sequence of simplices where each pair of simplices are one of the arrows of V, no simplex appears twice except possibly the last, and they are connected in the sense that every p-simplex is a face of the following $(p+1)$-simplex, i.e.,

$$\alpha_0^{(p)}, \quad \beta_0^{(p+1)}, \quad \alpha_1^{(p)}, \quad \beta_1^{(p+1)}, \quad \dots \quad , \quad \alpha_r^{(p)}, \quad \beta_r^{(p+1)}, \quad \alpha_{r+1}^{(p)},$$

where for every $0 \leq i, s \leq r$, $s \neq i$,

$$\{\alpha_i, \beta_i\} \in V, \quad \alpha_i \neq \alpha_r \quad \text{and} \quad \alpha_{i+1} \subset \beta_i.$$

If the sequence contains more than one simplex and starts and ends with the same simplex i.e., $r \neq -1$, $\alpha_0 = \alpha_{r+1}$, we call the sequence a *non-trivial closed V-path.*

We have already stated that every induced gradient vector field is a discrete vector field and noted a discrete vector field containing a non-trivial closed V-path that cannot be an induced gradient vector field, as shown in Figure 10. Any non-trivial closed V-path in any discrete vector field will lead to a contradiction in the same way as in the counterexample, i.e., for the i-th arrow, we have the inequality

$$f(\alpha_i) \geq f(\beta_i) > f(\alpha_{i+1})$$

and when collecting these inequalities for all arrows, we get

$$f(\alpha_0) \geq f(\beta_0) > f(\alpha_1) \geq f(\beta_1) > \cdots \geq f(\beta_r) > f(\alpha_{r+1}) = f(\alpha_0).$$

Therefore, an induced gradient vector field on K must also be a discrete vector field on K containing no non-trivial closed V-paths.

A V-path is in fact analogous to a longer river in the map analogy, and as we know, a river cannot loop back on itself. Therefore, it makes sense that an induced gradient vector field cannot contain a non-trivial closed V-path. The following theorem formalizes this.

Theorem 3.14. *Let K be a simplicial complex. An induced gradient vector field on K is equivalent to a discrete vector field on K that contains no non-trivial closed V-paths.*

It is possible to prove that a discrete vector field on K that contains no non-trivial closed V-paths is also an induced gradient vector field on K using the following result about directed graphs:

Theorem 3.15. *Given a directed graph, there exists a strictly decreasing real-valued function along each directed path if and only if there are no directed loops.*

We content ourselves with noticing the striking similarity between the two theorems: A discrete Morse function will decrease along a V-path, and a directed loop is similar to a non-trivial closed V-path. For a proof of Theorem 3.14 see, e.g., [5, pp. 21–23] or [10, pp. 68–69]. For a proof of Theorem 3.15, see [2, p. 13–14].

We can use Theorem 3.14 to find an induced gradient vector field and a discrete Morse function on any given simplicial complex by simply finding a discrete vector field without any non-trivial closed V-paths and then defining the discrete Morse function by starting at a critical 0-simplex and giving increasing values along every V-path in reverse order, making sure that before assigning a value to any simplex, all of its faces that are not part of the V-path already has a value. We allow both simplices that are part of an arrow to have the same value, but require strictly increasing values elsewhere.

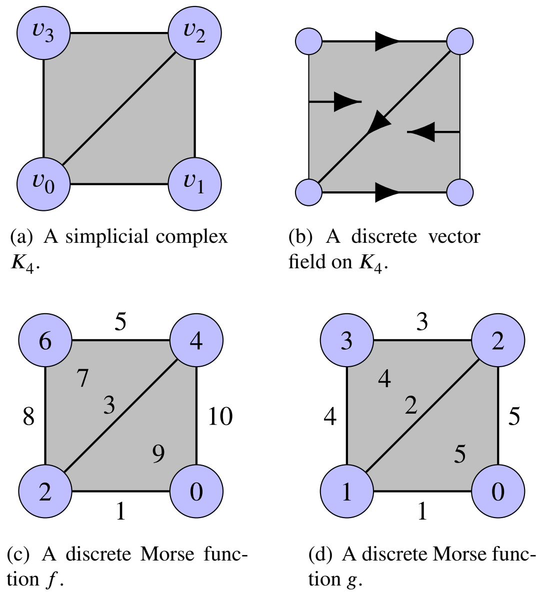

(a) A simplicial complex K_4.

(b) A discrete vector field on K_4.

(c) A discrete Morse function f.

(d) A discrete Morse function g.

Figure 11: An application of Theorem 3.14.

Example 3.16. If we take the simplicial complex K_4 as represented by Figure 11(a), we can choose a discrete vector field as represented by Figure 11(b). It is easy too see that this has no non-trivial closed V-paths and is therefore the induced gradient vector field of some discrete Morse function f. We will demonstrate two slightly different methods giving two slightly different discrete Morse functions f and g. We notice that v_1 is neither the head nor tail of any arrow and thus critical. We start with the V-path

$$v_1, v_{01}, v_0, v_{02}, v_2, v_{23}, v_3.$$

We assign values in increasing order like so:

$$f(v_1) = 0, \quad f(v_{01}) = 1, \quad f(v_0) = 2, \quad f(v_{02}) = 3,$$
$$f(v_2) = 4, \quad f(v_{23}) = 5, \quad f(v_3) = 6.$$

We could also have given the same values to each of the two simplices in each arrow like so:

$$g(v_1) = 0, \quad g(v_{01}) = 1, \quad g(v_0) = 1, \quad g(v_{02}) = 2,$$
$$g(v_2) = 2, \quad g(v_{23}) = 3, \quad g(v_3) = 3.$$

We notice that two pairs of simplices, $\{v_{023}, v_{03}\}$ and $\{v_{012}, v_{12}\}$, still need to be assigned values, so we let

$$f(v_{023}) = 7, \quad f(v_{03}) = 8, \quad f(v_{012}) = 9, \quad f(v_{12}) = 10,$$
$$g(v_{023}) = 4, \quad g(v_{03}) = 4, \quad g(v_{012}) = 5, \quad g(v_{12}) = 5.$$

We have now defined the discrete Morse functions f and g, represented by Figure 11(c) and Figure 11(d) respectively.

4 Elementary collapses and expansions

Elementary collapses and expansions are a way to modify the simplicial complex without modifying its topology and analogous to a deformation retract in topology. They let us say that two different simplicial complexes are equivalent or to simplify the complex we are working on. In our map analogy, the closest equivalent might be different amounts of detail. It might be hard to immediately see that two different maps are of the same area, but if one can be constructed from the other, they have to be. The rougher version of the map also contains less noise and superfluous information, making it easier to find your way. Similarly, simplifying simplicial complexes reduces the complexity of calculations and may therefore be useful in computations.

Definition 4.1 (Elementary collapse and expansion, free pair)**.** Given a simplicial complex K that contains a pair of simplices $\{\alpha^{(n)}, \beta^{(n+1)}\}$ such that $\alpha \subset \beta$ and α has no other cofaces, the collection $K \setminus \{\alpha, \beta\}$ is called an *elementary collapse* of K, denoted by $K \searrow K \setminus \{\alpha, \beta\}$.

Similarly, given a simplicial complex K that does not contain a pair of simplices $\{\alpha^{(n)}, \beta^{(n+1)}\}$ such that $\alpha \subset \beta$ and all other faces of β are in K, the collection $K \cup \{\alpha, \beta\}$ is called an *elementary expansion* of K, denoted by $K \nearrow K \cup \{\alpha, \beta\}$.

In both the case of the elementary collapse and the elementary expansion, the pair $\{\alpha^p, \beta^{(p+1)}\}$ is called a *free pair*.

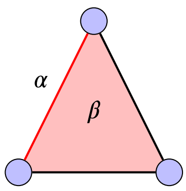

Figure 12: A simplicial complex K with a free pair $\{\alpha, \beta\}$.

Example 4.2. If the red simplices in the simplicial complex K represented by Figure 12 are part of K, then we can check that α has no other cofaces in K than β. Thus, it is a free pair and we can collapse K as shown in Figure 13.

If the red simplices are not part of K, then we can check that all the faces of β except α are in K. Thus, it is a free pair and we can expand K as shown in Figure 14.

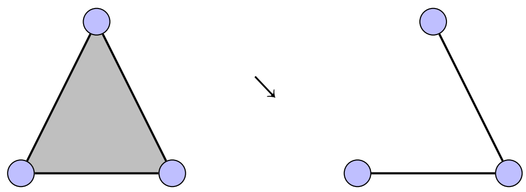

Figure 13: An elementary collapse.

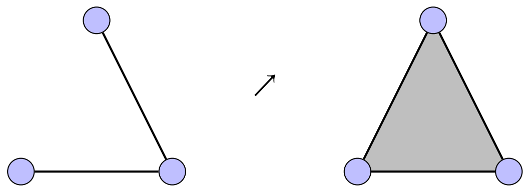

Figure 14: An elementary expansion.

In both the case of the elementary collapse and the elementary expansion, the new collection is a simplicial complex, as the following theorem states.

Theorem 4.3. *Let K be a simplicial complex. If $\{\alpha, \beta\} \subset K$ is a free pair, then $K \setminus \{\alpha, \beta\}$ is also a simplicial complex. If $\{\alpha, \beta\}$ is a free pair where $\alpha, \beta \notin K$, then $K \cup \{\alpha, \beta\}$ is also a simplicial complex.*

Proof. Let K be a simplicial complex containing a free pair $\{\alpha^{(n)}, \beta^{(n+1)}\}$. We modify the set of vertices if needed. We show that $K \setminus \{\alpha, \beta\}$ is a simplicial complex by arguing that neither α nor β is a face of any simplex $\gamma \in K \setminus \{\alpha, \beta\}$. As β is a coface of α, any coface of β would also be a coface of α i.e.,

$$\alpha \subset \beta, \beta \subset \gamma \implies \alpha \subset \gamma,$$

meaning that if β has some coface γ, a contradiction with our assumption that α has no other cofaces than β arises. Thus, $K \setminus \{\alpha, \beta\}$ is a simplicial complex.

Let K be a simplicial complex not containing a free pair $\{\alpha^{(n)}, \beta^{(n+1)}\}$. We show that $K \cup \{\alpha, \beta\}$ is a simplicial complex by arguing that all faces of both α and β are contained in $K \cup \{\alpha, \beta\}$. By assumption, all faces of β except α are in K. But as α is a face of β, all faces of α are also faces of β, meaning that they are also in K. Thus, $K \cup \{\alpha, \beta\}$ is a simplicial complex. □

Definition 4.4 (Collapsible)**.** If a simplicial complex K is *collapsible* if there exists a sequence of elementary collapses from K to a single vertex:

$$K = K_0 \searrow K_1 \searrow \cdots \searrow K_n = \{v\}. \tag{1}$$

We will write $K \searrow S$ when there exists a sequence of elementary collapses from K to S. This means that (1) could be rewritten as $K \searrow \{v\}$.

Example 4.5. The simplicial complex represented by Figure 12 is in fact collapsible, due to the sequence of elementary collapses shown in Figure 15.

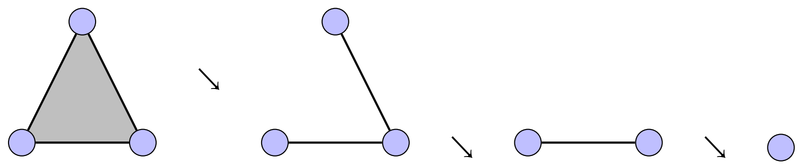

Figure 15: A collapsible simplicial complex.

We now introduce a way to look at subsets of a simplicial complex, arranged primarily by height, letting us construct the simplicial complex by adding simplices in sequence.

Definition 4.6 (Level subcomplex)**.** Given a discrete Morse function f on a simplicial complex K and any $c \in \mathbb{R}$, the *level subcomplex* $K(c)$ is the subcomplex of K containing simplices α with $f(\alpha) \leq c$ as well as their faces. That is,

$$K(c) = \bigcup_{f(\alpha)\leq c} \bigcup_{\beta\subset\alpha} \beta.$$

In other words, take all simplices with value up to c and include faces until you have a simplicial complex.

Example 4.7. We look at K_5 with f as represented by Figure 5(a). The level subcomplex $K_5(1)$ must contain the simplices marked 0 and 1, because they have values less than or equal to 1, but it must also include the simplex marked 2, as it is a face of the simplex marked 1. Observe that $K_5(1) = K_5(2)$. Similarly, $K_5(9)$ must contain all simplices that are assigned values less than or equal to 9, but even though v_{24} has value 10, it is a face of v_{234}, which has value 9, and must therefore be part of $K_5(9)$.

Remark 4.8 (Geometric realization)**.** While a simplicial complex is simply combinatorial data, the geometric realization of a simplicial complex is a topological analogue of the combinatorial data. For the purposes of this paper and at most 2-dimensional simplicial complexes, the intuitive approach of looking at the figure that represents the simplicial complex as a topological space is good enough. A discussion of geometric realizations can be found in [6, pp. 154, 167].

The i-th Betti number of a simplicial complex K is a topological invariant of the geometric realization of K, defined to be the rank of the i-th simplicial homology group of K. We will always compute simplicial homology groups with coefficients in $\mathbb{F}_2$, the integers modulo 2. Therefore, the rank of the simplicial homology group is equal to the dimension of the simplicial homology group seen as a vector space over $\mathbb{F}_2$, giving us the following definition.

Definition 4.9 (Betti number)**.** The i-th *Betti number* b_i of a simplicial complex K is given by

$$b_i = \dim(H_i(K; \mathbb{F}_2)),$$

where $H_i(K; \mathbb{F}_2)$ is the i-th simplicial homology group of K with coefficients in $\mathbb{F}_2$.

Let K be a simplicial complex. Then b_0 is the number of connected components, and for $i \geq 1$, b_i is the number of i-dimensional holes in the geometric realization of K. For example, K_5 as represented by Figure 4 has one connected component and a single

1-dimensional hole, meaning that $b_0 = 1, b_1 = 1, b_k = 0$ for $k \geq 2$. Note that the Betti numbers are independent of the discrete Morse function.

The Euler characteristic is often defined as the alternating sum of the number of simplices of each dimension. In fact, this is equivalent to the alternating sum of the Betti-numbers as shown in [9, p. 146].

Definition 4.10 (Euler characteristic)**.** The *Euler characteristic* χ of a simplicial complex K is given by

$$\chi(K) = \sum_{i=0}^{\dim(K)} (-1)^i c_i = c_0 - c_1 + c_2 - c_3 + \cdots + (-1)^{\dim(K)} c_{\dim(K)}$$

or, equivalently,

$$\chi(K) = \sum_{i=0}^{\dim(K)} (-1)^i b_i = b_0 - b_1 + b_2 - b_3 + \cdots + (-1)^{\dim(K)} b_{\dim(K)}.$$

Example 4.11. As the c-vector of K_5 as represented by Figure 4 is $(5, 6, 1)$, the Euler characteristic is

$$\chi(K_5) = c_0 - c_1 + c_2 = 5 - 6 + 1 = 0.$$

5 Main theorems

We finish with some of the important results from discrete Morse theory. The following theorem gives us that we can simplify a simplicial complex without removing critical values by collapses, or, in our map analogy, that we can remove details without losing the information we care about as long as peaks, basins and passes are still visible.

Theorem 5.1 (Collapse theorem)**.** *Let f be a discrete Morse function on a simplicial complex K. If the interval $(a, b] \subset \mathbb{R}$ contains no critical values, then $K(b) = K(a)$ or $K(b)$ collapses to $K(a)$, i.e, $K(b) \searrow K(a)$.*

Proof. For simplicity, we assume that f takes only integer values. If $K(b) = K(a)$, we are finished. We now assume $K(b) \neq K(a)$. We show that for any regular value b, we have either $K(b) = K(b-1)$ or $K(b) \searrow K(b-1)$.

Let δ be a simplex with $f(\delta) = b$. Then, δ must be regular, and thus either the head or tail of an arrow in the induced gradient vector field V_f. Renaming δ as appropriate, we call the arrow (α, β). As $(\alpha, \beta) \in V_f$, we know that $\alpha \subset \beta$ and $f(\alpha) \geq f(\beta)$. If δ is the tail, $f(\alpha) = b$. If δ is the head, $f(\alpha) \geq f(\delta) = b$. In any case, $f(\alpha) \geq b$.

We now assume that (α, β) is not a free pair. That means that α must have some coface γ in $K(b)$ with $f(\gamma) \leq b$. But then $f(\gamma) \leq b \leq f(\alpha)$, meaning that both β and γ are cofaces of α with lower value, which contradicts the assumption that f is a discrete Morse function.

Thus, every arrow in $K(b)$ where at least one of the simplices has value b must be a free pair. This means that every simplex with value b can be removed through elementary collapses on its own. But we must still argue that they can be removed in sequence. We know that no simplex is in more than one pair and that no elementary collapse, which removes simplices, can add a coface to the head of any arrow. Thus, we know that we can remove the simplices independently and therefore in sequence, meaning that $K(b) \searrow K(b-1)$.

Repeating the process, we get that for every $c \in (a, b]$, $K(c)$ is either equal to or collapses to $K(c-1)$, meaning that $K(b) \searrow K(a)$. □

Example 5.2. We verify that the collapse theorem applies to the simplicial complex represented by Figure 5(a). As noted earlier, all values are regular except for 0 and 11. Thus (0, 10] contains no critical values. The theorem states that we should be able to collapse $K_5(10)$ to $K_5(0)$, as evidenced by the sequence of collapses shown in Figure 16.

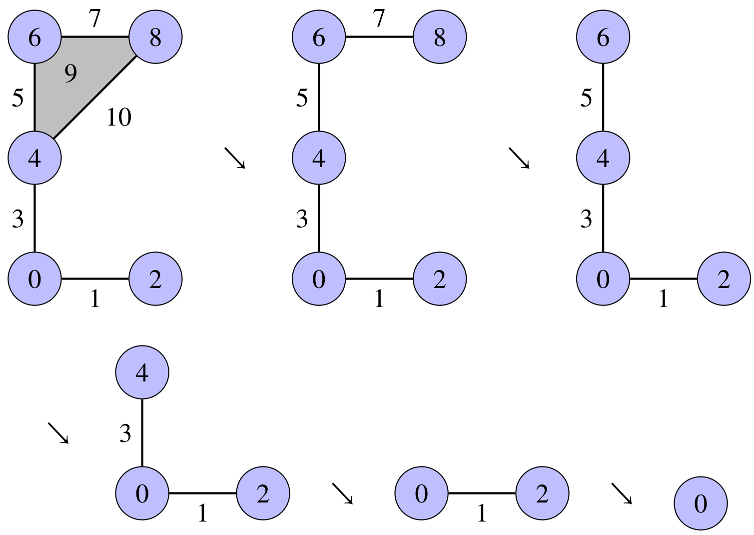

Figure 16: A sequence of collapses from $K(10)$ to $K(0)$.

We will now give a rough outline of the theory of CW-complexes, which as we will see in the next theorem, have a strong connection to simplicial complexes. A CW-complex, also called a cell complex, is a topological space that can be constructed by attaching unit balls of increasing dimension to a set of points. We call a unit ball of dimension n an $n - cell$. For example, the 1-dimensional open ball is the unit interval (0, 1). We can continuously deform this line to a half circle and attach the endpoints to two different points, giving the CW-complex represented by Figure 17. The interested reader can find a complete definition in [1, pp. 149–150].

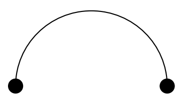

Figure 17: Two points and a 1-cell.

The following theorem can be considered the main result in discrete Morse theory, as it gives a strong connection between simplicial complexes and topological spaces.

Theorem 5.3. *Let K be a simplicial complex with a discrete Morse function f and m_i critical i-simplices. Then the geometric realization of K is homotopy equivalent to a CW-complex with m_i cells of dimension i.*

A proof of this theorem can be found in [4, p. 107].

Example 5.4. We can verify this theorem by applying it to K_5 with f, as represented by Figure 5(a).We first note that K_5 should be homotopy equivalent to a circle. As K_5

has one critical 0-simplex and one critical 1-simplex, the theorem says that it should be homotopy equivalent to a CW-complex with one cell of dimension 0 and one cell of dimension 1. That means we have to attach the boundary of a line to a single point, which yields the CW-complex represented by Figure 18. As expected, this CW-complex is a circle.

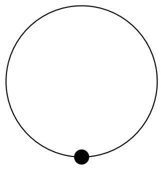

Figure 18: The circle as a CW-complex.

If we instead used g, we would have $m_0 = m_1 = 2$. In that case, we would know that K_5 must be homotopy equivalent to at least one of the four CW-complexes that can be created from two points and two 1-cells, displayed in Figure 19, but we would not immediately know which.

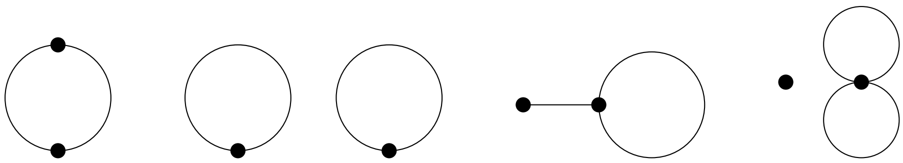

(a) (b) (c) (d)

Figure 19: Four CW-complexes constructed from two points and two 1-cells.

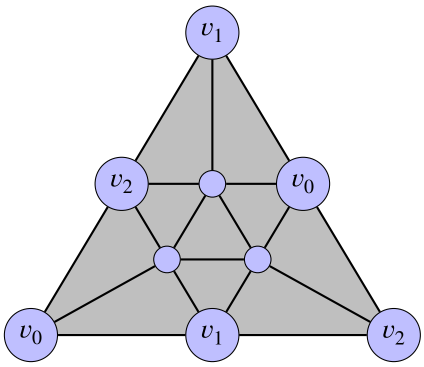

(a) A triangulation of the real projective plane.

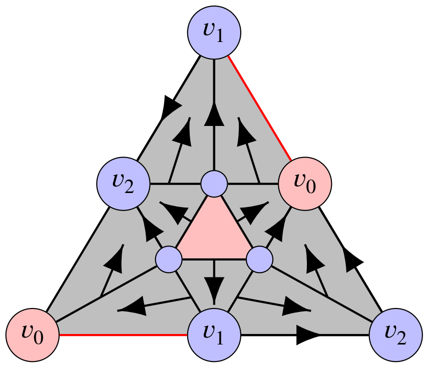

(b) An induced gradient vector field on the real projective plane.

Figure 20: Applying Theorem 5.3 to the real projective plane.

Example 5.5. We borrow an example from [5, p. 16]. The real projective plane, $\mathbb{R}P^2$, can be triangulated as represented by Figure 20(a). As the real projective plane is not easily represented in two dimensions, or even in three dimensions, we have to draw some vertices multiple times. Now, there is a discrete vector field with no non-trivial closed V-paths i.e., an induced gradient vector field, as demonstrated in Figure 20(b). The simplices marked in red are neither the head nor the tail of an arrow and therefore critical. There is one 0-simplex, one 1-simplex and one 2-simplex, so by Theorem 5.3, the real projective plane is homotopy equivalent to a CW-complex consisting of a 0-cell, a 1-cell and a 2-cell.

The following inequalities represent relations between the number of critical simplices of a simplicial complex and its Betti numbers.

Theorem 5.6 (Strong Morse inequalities)**.** *Let $f\colon K \to \mathbb{R}$ be a discrete Morse function, m_i be the number of critical i-simplices and b_i be the i-th Betti number as defined earlier. For each n such that $0 \leq n \leq \dim(K) + 1$, the following inequality holds:*

$$\sum_{i=0}^{n} (-1)^{n-i} b_i \leq \sum_{i=0}^{n} (-1)^{n-i} m_i.$$

We will not prove the strong Morse inequalities here, as the proof relies on quite technical results that are beyond the scope of this paper. A proof of the strong Morse inequalities can be found in [10, pp. 103–104]. We will, however, show that the strong Morse inequalities imply the weak Morse inequalities.

Theorem 5.7 (Weak Morse inequalities)**.** *Let $f\colon K \to \mathbb{R}$ be a discrete Morse function, m_n be the number of critical n-simplices, b_n be the n-th Betti number as defined earlier and χ the Euler characteristic. Then the following hold:*

(i) $b_n \leq m_n$, *for* $0 \leq n \leq \dim(K)$

(ii) $\sum_{i=0}^{\dim(K)} (-1)^i m_i = \chi(K)$

Proof. We first show that $b_i \leq m_i$ for $0 \leq i \leq \dim(K)$. For $n = 0$, the strong Morse inequality immediately gives $b_n \leq m_n$. For any n such that $0 < n \leq \dim(K)$, we have the strong Morse inequalities

$$\sum_{i=0}^{n} (-1)^{n-i} b_i \leq \sum_{i=0}^{n} (-1)^{n-i} m_i \tag{2}$$

and

$$\sum_{i=0}^{n-1} (-1)^{n-1-i} b_i \leq \sum_{i=0}^{n-1} (-1)^{n-1-i} m_i. \tag{3}$$

Adding (2) and (3) yields

$$\sum_{i=0}^{n} (-1)^{n-i} b_i + \sum_{i=0}^{n-1} (-1)^{n-1-i} b_i \leq \sum_{i=0}^{n} (-1)^{n-i} m_i + \sum_{i=0}^{n-1} (-1)^{n-1-i} m_i$$

$$b_n + \sum_{i=0}^{n-1} ((-1)^{n-i} b_i + (-1)^{n-1-i} b_i) \leq m_n + \sum_{i=0}^{n-1} ((-1)^{n-i} m_i + (-1)^{n-1-i} m_i)$$

$$b_n + \sum_{i=0}^{n-1} ((-1+1)^{n-i} b_i) \leq m_n + \sum_{i=0}^{n-1} ((-1+1)^{n-i} m_i)$$

$$b_n \leq m_n.$$

We now show that

$$\sum_{i=0}^{\dim(K)} (-1)^i m_i = \chi(K).$$

The strong Morse inequalities for $\dim(K)$ and $\dim(K)+1$ are

$$\sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} b_i \leq \sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} m_i \tag{4}$$

and

$$\sum_{i=0}^{\dim(K)+1} (-1)^{\dim(K)+1-i} b_i \leq \sum_{i=0}^{\dim(K)+1} (-1)^{\dim(K)+1-i} m_i. \tag{5}$$

We know that K has no simplices of higher dimension than $\dim(K)$, meaning that $m_{\dim(K)+1} = 0$, and K has no holes of higher dimension than $\dim(K)$, meaning that $b_{\dim(K)+1} = 0$. Thus (5) becomes

$$\sum_{i=0}^{\dim(K)} (-1)^{\dim(K)+1-i} b_i \leq \sum_{i=0}^{\dim(K)} (-1)^{\dim(K)+1-i} m_i$$

$$-\sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} b_i \leq -\sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} m_i$$

$$\sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} b_i \geq \sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} m_i. \tag{6}$$

Combining (4) and (6) we get that

$$\sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} b_i = \sum_{i=0}^{\dim(K)} (-1)^{\dim(K)-i} m_i,$$

which is equivalent to

$$\chi(K) = \sum_{i=0}^{\dim(K)} (-1)^i b_i = \sum_{i=0}^{\dim(K)} (-1)^i m_i. \qquad \square$$

Example 5.8. Given any simplicial complex K and discrete Morse function $f : K \to \mathbb{R}$, we can use the Morse inequality $b_i \leq m_i$ to obtain an upper bound on the Betti numbers of K. The result depends on the choice of function, and may give very loose or very tight upper bounds, possibly even the Betti numbers themselves. We use K_5 as introduced in Example 2.10 as an example. The function g, represented by Figure 5(b) has four critical simplices, two of dimension 0 and two of dimension 1. Thus we know that $b_0 \leq m_0 = 2$ and $b_1 \leq m_1 = 2$. If we instead chose the function f, represented by Figure 5(a), we would get one critical simplex of dimension 0 and one of dimension 1. That would give us $b_0 \leq m_0 = 1$ and $b_1 \leq m_1 = 1$, which in fact are the Betti numbers of K_5.

We will now discuss in detail an easily imaginable real life application from [10, pp. 2–6]. Most people have a cell phone, which communicates with the rest of the world via cell phone towers. But most of us have probably also experienced lacking cell phone service, as the cell phone towers have a limited range or for any other reason. We can use discrete Morse theory to find if there are any holes in the coverage. We assume that the towers have a set radius and that there are no obstructions blocking the signals. We also assume that the towers communicate with each other and know whether their ranges overlap. Let Figure 21 represent the grid.

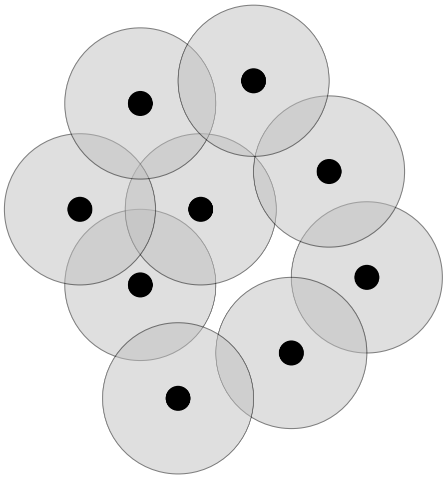

Figure 21: A grid of cell phone towers, showing their covered area.

We know which towers have overlapping coverage and can represent this data by a simplicial complex where the 0-simplices are the towers, the 1-simplices connecting two towers represent that their ranges overlap and the 2-simplices connecting three towers that they all their ranges overlap. Then we get the simplicial complex K_g represented by Figure 22.

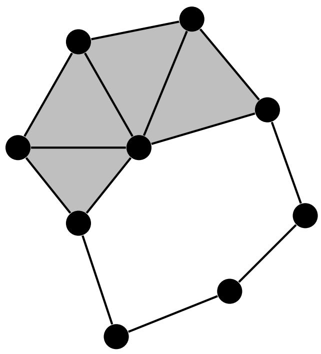

Figure 22: The connectivity of the towers represented by a simplicial complex K_g.

Now, we can use discrete Morse theory to find whether there is a hole in the simplicial complex. We find a discrete vector field on K_g with no non-trivial closed V-paths, as represented by Figure 23. Then the two red simplices are the two critical simplices. Using the weak Morse inequalities, Theorem 5.7, we learn that $b_0 \leq 1$, $b_1 \leq 1$ and $b_i = 0$ for $i > 1$. Thus, $b_0 = b_1 = 1$, and Theorem 5.3 gives us that the geometric realization of K_g is homotopy equivalent to the circle, meaning that there is a hole in the coverage. We could also have seen directly from $b_1 = 1$ that there is a hole.

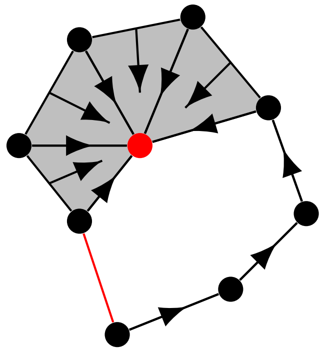

Figure 23: An induced gradient vector field on the grid.

6 Summary

We have now presented an introduction to discrete Morse theory, starting with defining simplicial complexes and moving on to discrete Morse functions and induced gradient vector fields. We now know how to construct one from the other and how they help us learn about the simplicial homology of a simplicial complex as well as being able to compare different simplicial complexes. Just as discrete Morse theory is comparable to a topographic map, we hope this introduction can be comparable to a road map on the readers way to understanding the subject. The next step, if one wishes to learn more about discrete Morse theory, might be to read [10] for a more thorough discussion or [8] for an introduction to Morse theory in the smooth case.

References

[1] Marcelo Aguilar, Samuel Gitler, and Carlos Prieto. *Algebraic topology from a homotopical viewpoint*. Universitext. Springer-Verlag, New York, 2002.

[2] Jørgen Bang-Jensen and Gregory Gutin. *Digraphs: Theory, Algorithms and Applications*. Springer Monographs in Mathematics. Springer, second edition, 2009.

[3] Tamal Dey and Yusu Wang. *Computational Topology for Data Analysis*. Cambridge University Press, Cambridge, 2022.

[4] Robin Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134(1):90–145, 1998.

[5] Robin Forman. A user's guide to discrete Morse theory. *Sém. Lothar. Combin.*, 48, 12 2001.

[6] John M. Lee. *Introduction to topological manifolds*, volume 202 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2000.

[7] Marston Morse. The foundations of a theory in the calculus of variations in the large. *Transactions of the American Mathematical Society*, 30(2):213–274, 1928.

[8] Liviu I. Nicolaescu. *An invitation to Morse theory*. Universitext. Springer, New York, 2007.

[9] Joseph J. Rotman. *An introduction to algebraic topology*, volume 119 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1988.

[10] Nicholas A. Scoville. *Discrete Morse theory*, volume 90 of *Student Mathematical Library*. American Mathematical Society, Providence, RI, 2019.

NTNU
Norwegian University of
Science and Technology