Discrete Morse Theory on Simplicial Complexes

August 27, 2009

ALEX ZORN

ABSTRACT: Following [2] and [3], we introduce a combinatorial analog of topological Morse theory, and show how the introduction of a discrete Morse function on a simplicial complex gives rise to a discrete vector field. We then move to from the combinatorial setting to the topological setting, and interpret our work in the language of homotopy classes of CW complexes. We conclude by showing the power of discrete Morse theory in analyzing a complicated simplicial complex.

CONTENTS

PART 1: SIMPLICIAL COMPLEXES

We begin by introducing the notion of an abstract simplicial complex. It should be noted that, while the following definition is purely combinatorial, it carries a natural geometric interpretation fundamental to the study of such objects.

Definition 1.1 An *abstract simplicial complex* is a finite set V and a collection K of subsets such that if $A \in K$ and $B \subset A$, $B \in K$.

The elements of V are referred to as "vertices," while elements of K are denoted "simplicies." Specifically, an element of K containing $d + 1$ vertices is called a "d-simplex." By slight abuse of notation, we can refer to the simplicial complex simply using the letter K. (Note that, in all but one case, $\emptyset \in K$. For the rest of this paper, unless explicitly stated, whenever we write "simplex" we mean "nonempty simplex").

Geometrically, we can view a simplicial complex as a structure made up of points, lines, triangles, etc. Specifically, a 1-simplex can be viewed as a line connecting 2 points, a 2-simplex is a triangle with 3 points as vertices, a 3-simplex is a tetrahedron spanned by its 4 points, and so on. With this in mind, the number d (in $d-$simplex) is called the dimension of a simplex. Below are two straightforward examples of simplicial complexes, including the sets V and K and the geometric interpretation.

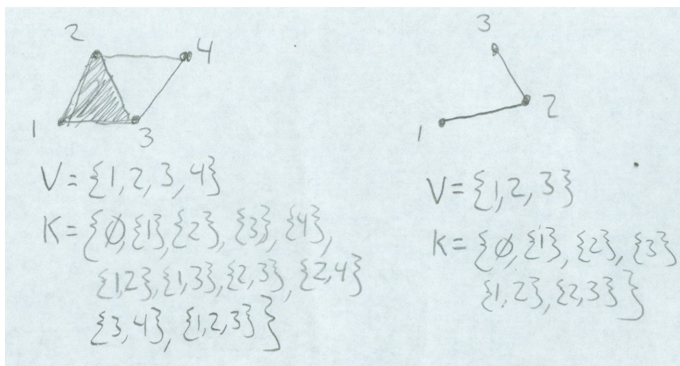

Another definition before we proceed:

Definition 1.2 If K is a simplicial complex, with $\alpha, \beta \in K$, we say α is a *face* of β if $\alpha \subset \beta$. α is a *free face* of β if α is a face of β, but not a face of any other simplex in K.

We next define some basic combinatorial "moves" or "operations" on simplices:

Definition 1.2 If K is a simplicial complex, an *elementary simplicial removal* is the act of removing a maximal simplex (a simplex which is not a face of any other simplex in K. More specifically, an *elementary d-removal* is the act of removing a $d-$simplex from K.

Definition 1.3 If K is a simplicial complex, an *elementary simplicial collapse* is the act removing a maximal simplex β and another simplex α such that α is a free face of β.

If $K_2 \subset K_1$ are simplicial complexes, then we write $K_1 \searrow K_2$ if K_2 is the result of an elementary simplicial collapse on K_1. We can equivalently write $K_2 \nearrow K_1$, and say that K_1 is the result of an elementary simplicial expansion on K_2.

For a straightforward example, consider the simplicial collapse below obtained from deleting $\alpha = \{2, 3\}$ and $\beta = \{1, 2, 3\}$ from the simplicial complex K_1:

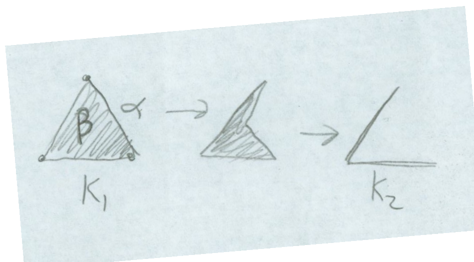

The topic of elementary collapse deserves a little more attention. In fact, as we will see later, elementary collapse is related to the topological notion of homotopy. This idea motivates the following definition:

Definition 1.4 *Simple homotopy equivalence* is the equivalence relation generated by elementary collapse. Stated otherwise, we have $K \sim K'$ if we can obtain K' from K by a series of elementary collapses and expansions.

We have one final definition:

Definition 1.5 A *complete reduction* of a simplicial complex K is a sequence of elementary collapses and elementary removals that transform K into the null simplicial complex (The complex containing only $\emptyset$).

It turns out that, as we will see in section 4, we can potentially learn a lot about K by examining the elementary removals in a complete reduction. In particular, the simplices we remove indicate what K is really "built" of, while the elementary collapses just induce homotopy equivalence (as we mentioned above). Of course, one can undergo a complete reduction solely with elementary removals by removing each simplex one by one, but in the following sections we show how to refine this process so that it gives us meaningful results about our space.

PART 2: DISCRETE MORSE FUNCTIONS

Let K be a simplicial complex and consider a function $f : K \to \mathbb{N}$. (In addition, we let $f(\emptyset) = 0$ and ignore it for the majority of the paper). We might hope that if $\alpha, \beta \in K$ and α is a face of β then $f(\alpha) < f(\beta)$. (We can accomplish this result easily by setting $f(\alpha) = d$ for α a $d - simplex$.) A discrete Morse function is a function that *almost* satisfies this property, in particular it has only one "mistake" locally. This notion is formalized in the following definition (to simplify notation from here forward, we will write a $d-$simplex α as $\alpha^{(d)}$.

Definition 2.1 K a simplicial complex. $f : K \to \mathbb{N}$ is a *discrete Morse function* if for all $\alpha^{(d)} \in K$:

- $f(\beta^{(d+1)}) \leq f(\alpha^{(d)})$ for at most one $\beta^{(d+1)} \supset \alpha^{(d)}$.
- $f(\beta^{(d-1)}) \geq f(\alpha^{(d)})$ for at most one $\beta^{(d-1)} \subset \alpha^{(d)}$.

As a simple example, consider the picture below, where the points (0-simplicies) and lines (1-simplicies) are labeled (by their value under a discrete Morse function f. Notice that, at the top line and the bottom vertex there are no "mistakes." However there is a line labeled 3 containing a point labeled 4 and a line labeled 1 containing a point labeled 2.

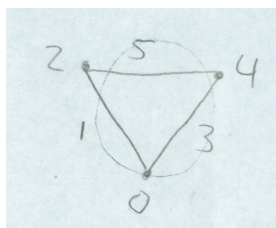

One more definition will serve us well:

Definition 2.2 For a discrete Morse function f, $\alpha^{(d)} \in K$ is a *critical point* of f if:

- For all $\beta^{(d+1)} \supset \alpha^{(d)}$, $f(\beta^{(d+1)}) > f(\alpha^{(d)})$
- For all $\beta^{(d-1)} \subset \alpha^{(d)}$, $f(\beta^{(d-1)}) < f(\alpha^{(d)})$.

Definition 2.3 If $f : K \to \mathbb{N}$ is a discrete Morse function we can *filter* K by $\mathbb{N}$ as follows: Let $K(n) = \{\alpha \in K | \exists \beta \in K$ with $\alpha \subseteq \beta$ and $f(\beta) \leq n\}$. In other words, take all simplices that are mapped to a value $\leq$ n, and then take all subsets of those simplices.

Note that $K(n)$ is a simplicial complex. Furthermore, $\{\emptyset\} = K(0) \subseteq K(1) \subseteq \cdots$, and for n big enough $K(n) = K$. We proceed with a lemma:

Lemma 2.3 If $\alpha^{(d)} \in K$ is a simplex, there do not exist simplices $\gamma^{(d-1)} \subset \alpha^{(d)} \subset \beta^{(d+1)}$ such that $f(\gamma^{(d-1)}) \geq f(\alpha^{(d)}) \geq f(\beta^{(d+1)})$.

Proof: Assume for the sake of contradiction such simplices exist. Then there is a vertex $v \in \beta^{(d+1)}$ such that $v \notin \alpha^{(d)}$. Let $\chi^{(d)} = \gamma^{(d-1)} \cup \{v\}$. If $f(\chi^{(d)}) \leq f(\gamma^{(d-1)})$, then $\chi^{(d)}$ and $\alpha^{(d)}$ are distinct $d-$ *simplices* containing

$\gamma^{(d-1)}$ which are mapped to equal or smaller values than $\gamma^{(d-1)}$, a contradiction. If $f(\chi^{(d)}) > f(\gamma^{(d-1)}) \geq f(\beta^{(d+1)})$, then $\chi^{(d)}$ and $\alpha^{(d)}$ are distinct $d-simplices$ contained in $\beta^{(d+1)}$ which are mapped to equal or larger values than $\beta^{(d+1)}$, also a contradiction.

Lemma 2.4 If $\alpha^{(d)} \in K$ is critical and $\alpha^{(d)} \subset \beta^{(d+r)}$, $f(\alpha^{(d)}) < \beta^{(d+r)}$.

Proof. We proceed by induction on r. The case $r = 1$ is given by the definition. Now suppose it is true for $r = k$, and suppose $\alpha^{(d)} \subset \beta^{(d+k+1)}$. Then there exist at least two simplices $\gamma^{(d+k)}, \chi^{(d+k)}$ such that $\alpha^{(d)} \subset \gamma^{(d+k)}, \chi^{(d+k)} \subset \beta^{(d+k+1)}$ (simply remove a vertex from $\beta^{(d+k+1)}$ that is not in $\alpha^{(d)}$). It follows that, since f is a Morse function , either $f(\beta^{(d+k+1)}) > f(\gamma^{(d+k)})$ or $f(\beta^{(d+k+1)}) > f(\chi^{(d+k)})$. WLOG we assume the former. Then by induction we have $f(\beta^{(d+k+1)}) > f(\gamma^{(d+k)}) > f(\alpha^{(d)})$. $\square$

To conclude this section, we prove two big (and very closely related) theorems regarding discrete Morse functions and critical points. The following proofs are involved, but worthwhile.

Theorem 2.5 If there do not exist critical simplices $\alpha \in K$ such that $f(\alpha) = n$, $K(n) \searrow K(n-1)$.

Proof. First notice that $K(n)$ satisfies three properties:

1) For all $\alpha \in K(n)$, there exists $\beta \in K(n)$ such that $\alpha \subseteq \beta$ and $f(\beta) \leq n$.

2) $K(n-1) \subseteq K(n)$.

3) If $f(\alpha) = n$, α is not critical in $K(n)$.

The first and second properties follow immediately from the definition. There is a subtlety in the third property: We know that if $f(\alpha) = n$ α is not critical *in K*, but it might be *in K(n)*. We show that it isn't: If $\alpha^{(d)} \in K(n)$ with $f(\alpha^{(d)}) = n$, either there exists $\beta^{(d-1)} \subset \alpha^{(d)}$ with $f(\beta^{(d-1)}) \geq \alpha^{(d)}$ or $\beta^{(d+1)} \supset \alpha^{(d)}$ with $f(\gamma^{(d+1)}) \leq \alpha^{(d)}$. In the first case, since $K(n)$ is a simplex we must have $\beta^{(d-1)} \in K(n)$. In the second case, $f(\gamma^{(d+1)}) \leq n$, so $\gamma^{(d+1)} \in K(n)$.

Now suppose M is any simplicial complex that satisfies those three properties. If $M \neq K(n-1)$, we will construct a simplicial complex M' that also satisfies those three properties, and is a simplicial collapse of M. Then, starting from K and repeating this process, we must eventually reach a stopping point. At this time we must have $K(n-1)$. Hence once we construct the above simplicial collapse we are done.

The construction is simple: Out of all the maximal simplicies in M, pick the one, say $\alpha^{(d)}$, such that $f(\alpha^{(d)})$ is maximized. If $f(\alpha^{(d)}) \leq n-1$ then we must have $M = K(n-1)$, by property 2 and the fact that every simplex in M is a face of some maximal simplex of M. If $f(\alpha^{(d)}) > n$ then $\alpha^{(d)}$ violates property 1. Hence $f(\alpha^{(d)}) = n$, so by property 3 it is not critical in M.

Then there exists $\beta^{(d-1)} \subset \alpha^{(d)}$ such that $f(\beta^{(d-1)}) \geq f(\alpha^{(d)})$. $\beta^{(d-1)}$ is a free face of $\alpha^{(d)}$: if not, pick another face $\gamma \supset \beta^{(d-1)}$. Then there exists $\chi \supseteq \gamma$ such that $f(\chi) \leq n$. But, by lemma 2.3, if we delete simplex $\alpha^{(d)}$, $\beta^{(d-1)}$ becomes a critical simplex. Hence lemma 2.4 gives us the following contradictory chain of inequalities: $n \geq f(\chi) > f(\beta^{(d-1)}) \geq f(\alpha^{(d)}) = n$. So we can delete $\alpha^{(d)}$ and $\beta^{(d-1)}$ to get a simplicial collapse to a new simplicial complex M'.

It remains to show that M' satisfies the three properties above. For property one, the only potentially problematic simplices are faces of $\alpha^{(d)}$ or $\beta^{(d-1)}$, but each of those are a face of (or equal to) some other $(d-1)-$simplex $\gamma^{(d-1)} \subset \alpha^{(d)}$. Since f is a Morse function, $f(\gamma^{(d-1)}) < f(\alpha^{(d)}) = n$. For property 2, note that $\alpha^{(d)}$ and $\beta^{(d-1)}$ were not faces of (or equal to) any simplex γ with $f(\gamma) = n-1$, hence neither of them are in $K(n-1)$. For property 3, our only potential problems are faces $\gamma^{(d-1)} \subset \alpha^{(d)}$ with $f(\gamma^{(d-1)}) = n$ or $\chi^{(d-2)} \in \beta^{(d-1)}$ with $f(\chi^{(d-2)}) = n$. However the first case is impossible because f is a Morse function. Inn the second case, if $\chi^{(d-2)}$ was critical in M', we would have $f(\chi^{(d-2)}) \geq f(\beta^{(d-1)})$, contradicting lemma 2.3. □

Theorem 2.6 $K(n)$ can be transformed into $K(n-1)$ by a series of elementary collapses and elementary removals, with exactly one elementary $d-$removal for each critical $d-$simplex $\alpha^{(d)}$ with $f(\alpha^{(d)}) = n$.

Proof. Suppose $f(\alpha) = n$, and α is critical in K. Then α is maximal in $K(n)$. To see why, assume there exists $\beta \in K(n)$ with $\beta \supset \alpha$. Then there exists $\gamma \supseteq \beta$ such that $f(\gamma) \leq n$. Then lemma 2.4 gives us the contradictory chain of inequalities $n = f(\alpha) < f(\gamma) \leq n$. This means we can take out each critical simplex α with $f(\alpha) = n$ using a series of elementary removals, one per simplex. Call the resultant simplicial complex M. We are done if we can show that M yields $K(n-1)$ through a series of simplicial collapses. Then we only need to show that M satisfies the three properties from the previous proof (remember that $K(n)$ satisfies all three).

1) For all $\alpha \in M$, there exists $\beta \in M$ such that $\alpha \subseteq \beta$ and $f(\beta) \leq n$. This still holds after removing the critical simplices. Every simplex that is a face of one of the removed simplexes $\alpha^{(d)}$ is also a face (or equal to) some simplex $\beta^{(d-1)} \subseteq \alpha^{(d)}$. But since $\alpha^{(d)}$ is critical, $f(\beta^{(d-1)}) < f(\alpha^{(d)}) = n$.

2) $K(n-1) \subseteq M$. This also still holds after removing the critical simplices, because all of them were maximal in $K(n)$, so none of them were in $K(n-1)$.

3) If $f(\alpha) = n$, α is not critical in M. We explicitly removed all the critical faces α with $f(\alpha) = n$ when we constructed M. The only possible faces whose criticality/non-criticality we might have changed in the removing are those faces $\beta^{(d-1)} \subset \alpha^{(d)}$ where $\alpha^{(d)}$ is a removed face. But then $f(\beta^{(d-1)}) < n$, so there is no problem. $\square$

Corollary 2.7 If K is a simplicial complex and $f : K \to \mathbb{N}$ is a discrete Morse function, then there exists a complete reduction of K with exactly one elementary $d-$removal for each critical $d-$simplex in K.

The above corollary is a very powerful result about discrete Morse functions, and it will play a big role in directing the flow of the paper. Part 3 will provide a tool for constructing discrete Morse functions, parts 4 and 5 work to interpret corollary 2.7 from a topological viewpoint, and part 6 uses the techniques of part 3 and the interpretation of parts 4 and 5 to establish a rather interesting result. Through all of this, however, corollary 2.7 is absolutely fundamental.

PART 3: DISCRETE VECTOR FIELDS

Writing a discrete Morse function is not exceptionally difficult (again we have the easy example of $f(\alpha^{(d)}) = d$). Writing a discrete Morse function that gives us interesting results proves to be more challenging. Specifically, corollary 2.7 suggests that we can learn more about our simplicial complex when we have less critical simplices. Note that non-critical simplices occur in pairs, $\alpha^{(d)} \subset \beta^{(d+1)}$ where $f(\alpha^{(d)}) \geq f(\beta^{(d+1)})$. We can illustrate by drawing in an "arrow" from $\alpha^{(d)}$ to $\beta^{(d+1)}$. This is illustrated below for the discrete Morse function from the previous part:

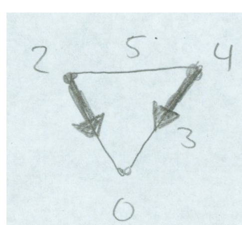

Note that by Lemma 2.3, no simplex can be at both the head and the tail of an arrow. Also, the points that are not the heads or tails of any arrow are exactly the critical simplices. Now suppose we have a simplicial complex, and draw in arrows. One would hope that we could find a discrete Morse function which gives us those arrows, in other words that Morse functions and sets of arrows are effectively the same thing. In order to properly address this question we introduce the definition of a "discrete vector field."

Definition 3.1 A *discrete vector field* on a simplicial complex K is a set of pairs $\{\alpha^{(d)} \subset \beta^{(d+1)}\}$ such that each simplex is in at most one pair.

The pairs $\{\alpha^{(d)} \subset \beta^{(d+1)}\}$ are thought of as "arrows" or "vectors" with a tail at $\alpha^{(d)}$ and a head a $\beta^{(d+1)}$. Thus the discrete vector field for the diagram above would be: $\{f^{-1}(2) \subset f^{-1}(1)\}, \{f^{-1}(4) \subset f^{-1}(3)\}$. Some discrete vector fields are shown pictorially below:

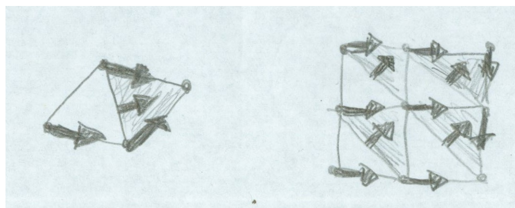

By an unfortunate notation standard, we shall refer to a discrete vector field as V (the same letter for the vertices of K). However, this will hopefully not cause confusion.

Now if $f : K \to \mathbb{N}$ is a discrete Morse function, we can construct a discrete vector field as above: Put in the pairs $\{\alpha^{(d)} \subset \beta^{(d+1)}\}$ when $f(\alpha^{(d)}) \geq f(\beta^{(d+1)})$. We call this the *gradient vector field* of f. Then we can re-ask the question above using our new terminology: *"When is a discrete vector field on a simplicial complex K the gradient vector field of some discrete Morse function on K?"* In order to properly answer this question, we introduce the notion of a $V-$path:

Definition 3.2 If V is a discrete vector field, a $V-$*path* is a set of simplices $\alpha_0^{(d)}, \beta_0^{(d+1)}, \alpha_1^{(d)}, \beta_1^{(d+1)}, \ldots, \beta_r^{(d+1)}, \alpha_{r+1}^{(d)}$ such that $\{\alpha_i^{(d)} \subset \beta_i^{(d+1)}\} \in V$ for $0 \leq i \leq r$ and $\beta_i \supset \alpha_{i+1} \neq \alpha_i$.

An example of a $V-$path is illustrated below: (α_1,α_2,α_3 are not labeled, to save space, but they are the intermediate line segments).

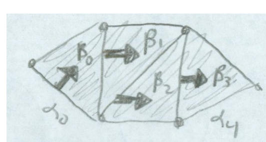

We say a $V-$path is *closed* if $\alpha_0 = \alpha_{r+1}$. In particular, a *non-trivial closed* $V-$*path* is a closed $V-$path with $r > 0$ (a trivial closed $V-$path would be just the simplex α_0). The reason for this definition is the following result, which is the main theorem of this section:

Theorem 3.3 If V is a discrete vector field on a simplicial complex K, V is the gradient vector field of some discrete Morse function on K if and only if there are no non-trivial closed $V-$paths (We say V is *acyclic*).

Proof. We prove "only if" here, and postpone the proof of "if" to the end of the section. Assume that V is the gradient vector field of a discrete Morse function $f : K \to \mathbb{N}$, and $\alpha_0^{(d)}, \beta_0^{(d+1)}, \ldots, \alpha_{r+1}^{(d)}$ is a non-trivial $V-$path. Then, since $\{\alpha_i^{(d)} \subset \beta_i^{(d+1)}\} \in V$ we have $f(\alpha_i^{(d)}) \geq f(\beta_i^{(d+1)})$. Then, since f is a discrete Morse function and $\alpha_{i+1}^{(d)} \neq \alpha_i^{(d)}$ we must have $f(\alpha_{i+1}^{(d)}) < f(\beta_i^{(d+1)})$. So we get the chain of inequalities:

$$f(\alpha_0^{(d)}) \geq f(\beta_0^{(d+1)}) > f(\alpha_1^{(d)}) \geq f(\beta_1^{(d+1)}) > \cdots > f(\alpha_{r+1}^{(d)})$$

From which we conclude $f(\alpha_0^{(d)}) > f(\alpha_{r+1}^{(d)})$, hence $\alpha_0^{(d)} \neq \alpha_{r+1}^{(d)}$. So no non-trivial $V-$path is closed. □

Note that in the proof we discovered something stronger: If V is the gradient vector field of a discrete Morse function on K, the $V-$paths are "continuous" paths along which f is decreasing.

The remainder of this section will be devoted to exploring a different way of looking at discrete vector fields which will give us the tools to finish the proof of theorem 4.3.

First, consider a simplicial complex K. We can "draw" K by representing each simplex as a point, and drawing an arrow from $\beta^{(d+1)}$ to $\alpha^{(d)}$ when $\beta^{(d+1)} \supset \alpha^{(d)}$. This is called the *Hasse diagram* of K. This process is illustrated below for a triangle:

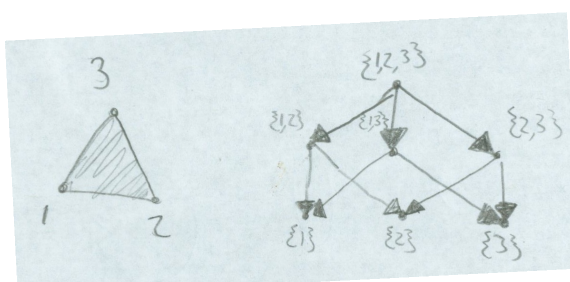

Now suppose V is a discrete vector field on K. We modify the Hasse diagram by reversing the direction of the arrow from $\beta^{(d+1)}$ to $\alpha^{(d)}$ when $\{\alpha^{(d)} \subset \beta^{(d+1)}\} \in V$. Again, we show the process:

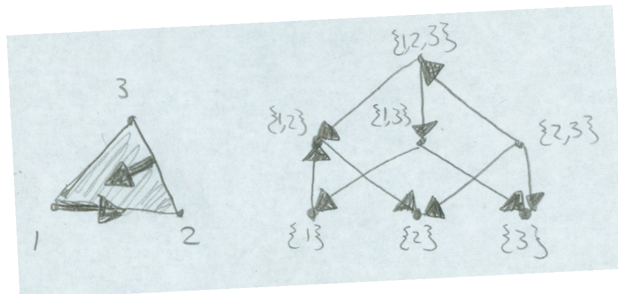

So for every simplicial complex K with a discrete vector field V, we have a corresponding Hasse diagram, which is a finite *directed graph.* Ie it is a finite amount of points with arrows between certain pairs of points. Then we have the notion of a "cycle" in the directed graph, which is simply a set of vertices in the graph $a_1, a_2, \ldots, a_n$ such that there is an arrow from a_i to a_{i+1} and from a_n to a_1. We have the following result:

Theorem 3.4 If V is a discrete vector field on a simplicial complex K, if there are no closed $V-$paths then the corresponding directed graph has no cycles.

Proof. We will show that if we have a cycle in the directed graph, we have a closed $V-$path. Note first that in a cycle in the directed graph, the dimension of adjacent simplices differs by 1, hence we know that the cycle must have an even number of elements. Let's call them $\alpha_0, \beta_0, \ldots, \alpha_r, \beta_r$. Without loss of generality we can assume $\alpha_0^{(d)} \subset \beta_0^{(d+1)}$, since if not we can just shift the cycle. (The dimension must increase at some point if we want to get back to where we started).

Now suppose we are at some simplex a_i in the cycle (I use "a" to differentiate from α and β). Then a_{i+2} cannot have a bigger dimension, otherwise there would be an arrow from a_i to a_{i+1} increasing dimension, and from a_{i+1} to a_{i+2} increasing dimension which contradicts Lemma 2.3. But since we have a cycle, this implies that a_{i+2} cannot have a smaller dimension either (otherwise there would be no way to get back to a_i when we come around). This means all the α simplices and all the β simplices have the same dimension, in particular the αs are $d-$simplices and βs are $d+1-$simplices. So our cycle gives us a $V-$path. $\square$

This is the key result connecting discrete vector fields to directed graphs. We can now use the following result from graph theory:

Theorem 3.5 (Topological Sort) If G is a finite directed graph, there exists a function f on the vertices of G (mapping into $\mathbb{N}$) such that $f(a) > f(b)$ whenever there is an arrow from a to b if and only if G has no cycle.

Proof. "Only if" is clear, since if we had a cycle such that we decrease along each arrow it would follow that the starting point must have a label less than its label, a clear contradiction.

For "if", we have the following method: Call each point that is not the tail of any arrow a node. Then, since G is finite and acyclic, a path starting from any vertex must eventually hit a node. In particular, there are a non-zero, finite number of such paths. Then for a vertex a, let $f(a)$ be the greatest number of vertices in a path from a to a node. (Note: this means $f(a) = 1$ iff a is a node). Then if there is an arrow from a to b, any path from a to a node can be appended to the arrow from a to b to make that path one vertex longer. Hence $f(a) > f(b)$ as desired. $\square$

The topological sort is illustrated for the directed graph below. Notice that the top left and bottom right vertices have more than one possible path to the upper right node point, so we have to choose the longest one. In addition, no paths to the bottom left node point are maximal.

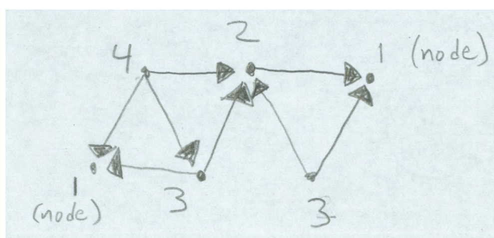

This result gives us what we need to finish theorem 3.3. We proceed in the natural way: Given a discrete vector field V with no $V-$paths, the corresponding directed graph has no cycles. Then we can make a function f such that $f(a) > f(b)$ whenever there is an arrow from a to b, which is easily seen to be a Morse function whose gradient vector field is V. This gives us the final result of this section, which is the appropriate analog of corollary 2.7 using our new language:

Theorem 3.6 Suppose K is a simplicial complex and V is an acyclic discrete vector field on K. Then there exists a complete reduction of K consisting of exactly one elementary $d-$removal for each unpaired $d-$simplex in V.

Proof: This follows immediately from corollary 2.7 and theorem 3.3.

PART 4: GEOMETRIC REALIZATION OF SIMPLICES

We now attempt to move our discussion of simplicial complexes into the topological realm. Here we return to the standard geometric picture of a $d-$simplex as a $d-$dimensional "triangle". We formalize this notion as follows:

Definition 4.1 The *standard d-simplex* Δ^d is the set $\{(t_0, t_1, \ldots, t_d) \in \mathbb{R}^{d+1} | \sum_i t_i = 1 \text{ and } t_i \geq 0 \; \forall i\}$.

Now we have written Δ^d as a subset of $\mathbb{R}^{d+1}$, and it inherits a topology from $\mathbb{R}^{d+1}$. We can also talk about the boundary, $\delta\Delta^d$ as the subspace of Δ^d such that at least one coordinate is 0. Now we proceed to the main definition of this section:

Definition 4.2 If K is a simplicial complex over the vertex set $V = \{v_0, v_2, \ldots, v_n\}$, then the *geometric realization* of K is a topological space $[K] \subseteq \mathbb{R}^{n+1}$. Explicitly, $[K] = \Delta^{n+1} \cap \{(t_0, t_1, \ldots, t_n) | \{v_i | t_i \neq 0\} \in K\}$.

This definition effectively identifies the *ith* coordinate with the *ith* vertex v_i. The definition states that , if a set of coordinates is nonzero, those vertices must be a simplex in K. With this formalism, we can state the two main results of this section:

Theorem 4.3 If K_1 and K_2 are simplicial complexes with $K_1 \searrow K_2$, then $[K_1]$ is homotopy equivalent to $[K_2]$.

Theorem 4.4 If K_1 and K_2 are simplicial complexes with $K_2 = K_1 - \alpha^{(d)}$, then $[K_1]$ is homeomorphic to a gluing of $[K_2]$ to Δ^d along $\delta\Delta^d$.

The remainder of this section is devoted to proving these two theorems.

We begin with theorem 4.3. First, let $V = \{v_0, v_1, \ldots, v_n\}$, K_1 be a simplicial complex over V and K_2 be an elementary collapse of K_1 obtained by removing simplices $\alpha^{(d)}$ and $\beta^{(d-1)}$, where $\beta^{(d-1)}$ is a free face of $\alpha^{(d)}$ (note that this implies $\alpha^{(d)}$ is maximal). Without loss of generality, we can assume $\alpha^{(d)} = \{v_0, v_1, \ldots, v_d\}$ and $\beta^{(d-1)} = \{v_0, v_1, \ldots, v_{d-1}\}$. Note that $[K_2] \subseteq [K_1] \subseteq \Delta^n$.

Now consider the vector $\mathbf{v} = (-1, \ldots, -1, d, 0, \ldots, 0)$, where the first d coordinates equal -1. We will collapse $[K_1]$ to $[K_2]$ along lines parallel to $\mathbf{v}$. For each $\mathbf{t} \in [K_1]$, let ϕ be the projection map defined by $\phi(\mathbf{t}) = \mathbf{t} + m_t\mathbf{v}$, where m_t is the minimum of the first d coordinates of $\mathbf{t}$. Now we are done if we can show the following three things:

1. ϕ maps $[K_1]$ into $[K_2]$

2. ϕ restricts to the identity on $[K_2]$

3. ϕ is homotopic to the identity on $[K_1]$

For item 1: First notice that, since the sum of the coordinates in $\mathbf{v}$ is zero, the sum of the coordinates in $\phi(\mathbf{t})$ is 1. In addition, only the first $d+1$ coordinates of $\mathbf{t}$ change. If $\mathbf{t} = (t_0, \ldots, t_{d-1}, t_d, t_{d+1} \ldots, t_n)$, then $\phi(\mathbf{t}) = (t_0 - m_t, \ldots, t_{d-1} - m_t, t_d + m_t d, t_{d+1}, \ldots, t_n)$. We see immediately that all coordinates must be nonnegative, so $\phi(\mathbf{t}) \in \Delta^n$.

Now, if $m_t = 0$ then $\phi(\mathbf{t}) = \mathbf{t}$, and $\mathbf{t} \in [K_2]$. In particular, the set $\{v_i | t_i \neq 0\}$ must be a simplex $\gamma \in K_1$, and since $m_t = 0 \Rightarrow$ one of the first d coordinates is 0 we have $\gamma \neq \alpha, \beta$, ie γ was not removed. If $m_t \neq 0$ we must have all of the first d coordinates nonzero, ie $\beta^{(d-1)} \subseteq \{v_i | t_i \neq 0\}$. But this implies $\{v_i | t_i \neq 0\} = \alpha^{(d)}$ or $\beta^{(d-1)}$. Then, $\phi(\mathbf{t})$ must have one of the first d coordinates equal to zero (the smallest one, which is equal to m_t), hence $\{v_i | \phi(\mathbf{t})_i \neq 0\}$ is some simplex $\gamma \subset \alpha^{(d)}$ not equal to $\alpha^{(d)}$ or $\beta^{(d-1)}$, hence $\gamma \in [K_2]$ so $\phi(\mathbf{t}) \in [K_2]$.

For item 2: If $\mathbf{t} \in [K_2]$, it means that $\beta \nsubseteq \{v_i | t_i \neq 0\}$, ie one of the first d coordinates must be 0. Then $m_t = 0$, so $\phi(\mathbf{t}) = \mathbf{t}$.

For item 3: For $(\mathbf{t}, \theta) \in [K_1] \times [0,1]$, let $\Gamma(\mathbf{t}, \theta) = \mathbf{t} + \theta m_t \mathbf{v}$. Then, since in an ε neighborhood of $\mathbf{t}$ m_t cannot vary by more than ε, Γ is continuous. In addition, $\Gamma(\mathbf{t}, \theta) \in [K_1]$, since $m_t \neq 0 \Rightarrow \{v_i | t_i \neq 0\} = \beta^{(d-1)}$ or $\alpha^{(d)}$, hence its image under Γ must also have nonzero coordinates from either $\beta^{(d-1)}$ or $\alpha^{(d)}$. Lastly, Γ equals the identity at $\theta = 0$ and ϕ at $\theta = 1$, hence ϕ is homotopic to the identity. This completes the proof.

Now we proceed to Theorem 4.4 and the notion of gluing. Informally, if we have two topological spaces A and B and a continuous function f from a subset $B_0 \subseteq B$ to A, we can glue B to A by identifying each point in B_0 with its image under f. To state this formally, we build up a set of elementary topological ideas:

Definition 4.5 If X is a topological space and $A \subseteq X$, there exists an *inclusion map* $i_A : A \to X$, taking the set A to itself in X. The *subspace topology* on A is as follows: U is open in A iff there exists V open in X such that $i_A^{-1}(V) = U$.

Definition 4.6 If A and B are topological spaces, their *disjoint union* $A \sqcup B$ is a topological space X which is the (setwise) disjoint union of the sets A and B. There is a *canonical injection* $\varphi_A : A \to X$, taking A to itself in X (resp. B). Then U is open in X iff $\varphi_A^{-1}(U)$ is open in A (resp. B).

Note that it is not totally correct to say $A \subset A \sqcup B$ (there is a problem if A and B are not disjoint initially), it is better to say A is equivalent to a subset of $A \sqcup B$, or $\varphi_A(A) \subset A \sqcup B$. This is the reason why we have differentiated the terms "inclusion map" and "canonical injection", and it will prevent the following proofs from seeming tautological.

Definition 4.7 If X is a topological space and $\sim$ is an equivalence relation on X, the *quotient space* $X/\sim$ is the set of equivalence classes in X. There exists a quotient map $q : X \to X/\sim$ which takes $x \in X$ to its equivalence class. A set $U \in X/\sim$ is open iff $q^{-1}(U)$ is open in X.

Remark: Inclusion maps, canonical injections, and quotient maps are continuous, as easily seen by the definition. Now we are ready for our main definition:

Definition 4.8 If A and B are topological spaces, a *gluing map* is a continuous function $f : B_0 \to A$ where $B_0 \subseteq B$. We then say $A \cup_f B = A \sqcup B/\sim$, where each element of B_0 is glued to its image. Explicitly, if $f(b) = a$, we have $\varphi_B(b) \sim \varphi_A(a)$.

Note that, while these definitions seem technical, they all do exactly what they "should do." Ie, the disjoint union of two spaces is exactly those two spaces with their separate topologies. A quotient space just takes a bunch of points and makes them "the same." The following lemma shows that subspace topologies are consistent:

Lemma 4.9 If X is a topological space and $A \subseteq B \subseteq X$, then the subspace topology induced on A is the same regardless of whether A is considered a subspace of X or of B (with the subspace topology).

Proof. Let $i_A : A \to X, i_B : B \to X, j_A : A \to B$ be the inclusion maps. Note $i_A = i_B \circ j_A$. Let τ_1 be the subspace topology on A as a subspace of X, and τ_2 be the subspace topology on A as a subspace of B. Then $U \in \tau_1 \Rightarrow \exists V$ open in X such that $i_A^{-1}(V) = U \Rightarrow i_B^{-1}(V)$ open in $B \Rightarrow j_A^{-1}(i_B^{-1}(V)) = i_A^{-1}(V) = U \in \tau_2$. Similarly, $U \in \tau_2 \Rightarrow \exists V$ open in B such that $j_A^{-1}(V) = U \Rightarrow \exists V'$ open in X such that $i_B^{-1}(V') = V \Rightarrow i_A^{-1}(V') = j_A^{-1}(i_B^{-1}(V')) = U \in \tau_1$. $\square$

Now we introduce another lemma, whose truth seems obvious, but whose proof is worth going through:

Lemma 4.10 If $A, B \subseteq X$ are topological spaces (A, B closed) let $i : A \cap B \to A$ be the inclusion map. Then X is homeomorphic to $A \cup_i B$.

Note that, if i is a gluing map, then $A \cap B$ is regarded as a subspace of B. However by lemma 4.9 this is irrelevant, hence i is continuous (being an inclusion map).

Now let $q : A \sqcup B \to A \cup_i B$ be the quotient map, $\varphi_A : A \to A \sqcup B$ be the canonical injection and $i_A : A \to X$ be the inclusion map (resp φ_B, i_B). For $[x] \in A \cup_i B$, define f as follows:

$$f([x]) = \begin{cases} i_A(y) & y \in A \text{ and } \varphi_A(y)) \in [x] \\ i_B(y) & y \in B \text{ and } \varphi_B(y)) \in [x] \end{cases}$$

Note that when $[x]$ is a singleton set, there is exactly one y that meets these criteria, since the images of φ_A and φ_B are disjoint and cover $A \sqcup B$. If $[x]$ contains two points, $\varphi_B(b)$ and $\varphi_A(a)$ such that $i(b) = a$, then we have $a = b$ since i is an inclusion map. This means $i_A(a) = i_B(b)$, so f is well-defined. Furthermore, f is an injection, since i_A and i_B are injections, so if $[x] \neq [x']$ and $f([x]) = f([x'])$ we would have $y \in A$, $\phi_A(y) \in [x]$ and $y' \in B$, $\phi_B(y') \in [x']$, and $i_A(y) = i_B(y')$. But this means $y = y' \in A \cap B$, hence $i(y') = y$, so $[x] = [x']$, a contradiction.

Now consider f as a map of sets. For $U \in A \cup_i B$, we can write $f(U) = f_A(U) \cup f_B(U)$ where $f_A(U) = (i_A \circ \varphi_A^{-1} \circ q^{-1})(U)$ and $f_B(U) = (i_B \circ \varphi_A^{-1} \circ q^{-1})(U)$. Note that: $f_A(A \cup_i B) = A$, $f_B(A \cup_i B) = B$, so $f(A \cup_i B) = A \cup B = X$, hence f is surjective. Also, for any $U \subset A \cup_i B$, $f_A(U) = f(U) \cap A$ and $f_B(U) = f(U) \cap B$.

So we have: $U \in A \cup_i B$ open $\Leftrightarrow q^{-1}(U)$ open in $A \sqcup B \Leftrightarrow \varphi_A^{-1}(q^{-1}(U))$ is open in A (resp. B) $\Leftrightarrow \exists V_A$ open in X such that $i_A^{-1}(V_A) = \varphi_A^{-1}(q^{-1}(U))$, equivalently $f_A(U) = V_A \cap A$ (resp. B). Now if $f(U)$ is open in X, we can simply take $V_A = V_B = f(U)$. On the other hand, if U is open in $A \cup_i B$, we have $f(U) \cap A = V_A \cap A$ and $f(U) \cap B = V_B \cap B$, hence $f(U) = (V_A - B) \cup (V_B - A) \cup (V_A \cap V_B)$, which is open. This means f is a homeomorphism.

SECTION 5: CW COMPLEXES

In this section, we present an array of results (mostly without proof) concerning cell complexes. This section rewrites results of the previous section into statements which are perhaps more aesthetically pleasing.

To begin, we define the *$n-$ball* $B^n = \{x \in \mathbb{R}^n | ||x|| \leq 1\}$ where $||\cdot||$ denotes the Euclidean norm. Then, the *$n-$sphere* is $S^n = \{x \in \mathbb{R}^{n+1} | ||x|| = 1\}$. Note $S^n \subseteq B^{n+1}$.

Now, the construction of a cell complex is as follows: One begins with a finite number of points (B^0). Then one adds higher dimensional balls B^n by "gluing" along the boundary. For example, suppose we begin with two points, a and b. Then notice the boundary of B^1 is S^0: 2 points. Hence we can glue on a B^1 (which is simply a line segment) by attaching one endpoint to a and the other to b. Then, if we do this again, we will have two segments between a and b which form a sort of loop (that looks a lot like S^1). Then we can glue a B^2 along that loop. The process is illustrated below:

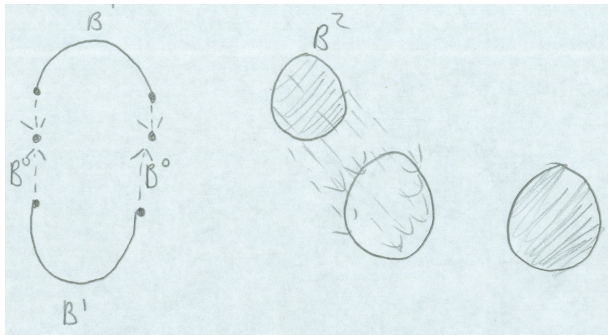

For a different sort of example, suppose we just start with one point a. If we glue on a B^1 with both endpoints at a we get a circle. If instead we glue on a B^2, we are forced to glue the entire boundary (S^1) to one point. It is not hard to see that the result of this is S^2.

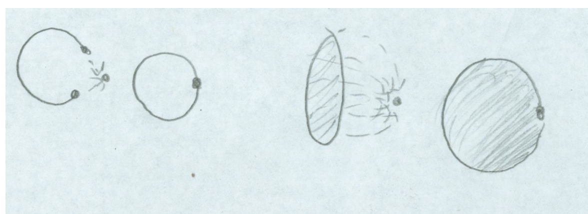

We have the following definition:

Definition 5.1 X is a *CW complex* if we have a finite chain of topological spaces $\emptyset \subset X_0 \subset \cdots \subset X_n = X$, such that $X_{i+1} = X_i \cup_f B^n$ for some n, where $f : S^{n-1} \to X_{i+1}$.

We now have the following important results:

Theorem 5.3 B^n is homeomorphic to Δ^n, with S^{n-1} mapped to $\delta\Delta^n$.

Proof Sketch: Embed Δ^n in B^n, then expand Δ^n to B^n. This is illustrated below with $n = 2$.

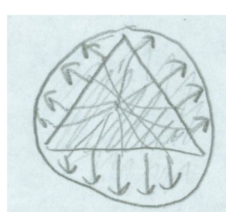

Theorem 5.4 Let $h : X \to X'$ denote a homotopy equivalence, and $f_1 : S^{n-1} \to X$ and $f_2 : S^{n-1} \to X'$ continuous maps. If $h \circ f_1$ is homotopy equivalent to f_2, then $X \cup_{f_1} B^n$ is homotopy equivalent to $X' \cup_{f_2} B^n$.

Corollary 5.6 If K is a simplicial complex with a complete reduction, then $[K]$ is homotopy equivalent to a CW complex involving one $d-$ball for each elementary $d-$removal in the reduction.

Proof. We construct a CW complex by reading the reduction in reverse. In particular, the null simplex space is homotopy equivalent to the empty topological space. Then, by theorems 4.4 and 5.3, encountering a $d-$removal prompts the gluing of a $d-$simplex, which is homeomorphic to B^n. By theorem 4.3, encountering an elementary collapse signifies a homotopy equivalence, and by theorem 5.4 this won't affect the homotopy class of later gluings. □

We are now in a position to state corollary 2.7 in its most refined form:

Corollary 5.7 Suppose K is a simplicial complex, and V an acyclic discrete vector field on K. Then $[K]$ is homotopy equivalent to a CW complex involving one B^d for each unpaired $d-$simplex in V.

Proof. This follows immediately from corollary 5.6 and theorem 3.6. □

We now proceed to demonstrate a special case of corollary 5.7.

Theorem 5.8 The CW complex $B^0 \cup_f B^n$, where $f : S^{n-1} \to B^0$ is the constant map (the only possible choice), is homeomorphic to S^n.

Proof Sketch: Show $B^n - S^{n-1}$ is homeomorphic to $\mathbb{R}^n$ is homeomorphic to $S^n - B_0$, using in the first case the map $x \mapsto x/(1 - ||x||)$ and in the second case a projection from the removed point $(1, 0, 0, \ldots, 0)$. Then, since all of S^{n-1} is glued to a point, adding S^{n-1} to $B^n - S^{n-1}$ becomes adding B_0 to $S^n - B_0$, which is S^n.

This implies that, if our only unpaired simplices are a 0-simplex and an $n-$simplex, our simplicial complex is homotopy equivalent to an $n-$sphere. We can extend this result further to a *wedge* of $n-$spheres, ie a collection of $n-$spheres glued together at one common point.

Theorem 5.9 If X is a finite wedge of $n-$spheres, any continuous function $f : S^{n-1} \to X$ is nullhomotopic.

Corollary 5.10 If K is a simplical complex, V is an acyclic discrete vector field on K, and the only unpaired simplices are a $0-$simplex and k $d-$simplices, then K is homotopy equivalent to the wedge of k $d-$spheres.

Proof. By theorems 5.9 and 5.4, gluing B^n onto a wedge of k $n-$spheres results in a space homotopy equivalent to a wedge of $k+1$ $n-$spheres, no matter what the gluing map is. Then we apply corollary 5.7. $\square$

This last result is the "aesthetically pleasing" statement we were looking for. We can now proceed to the final section.

PART 6: THE COMPLEX OF DISCONNECTED GRAPHS

In this section we use the results established in this paper to find the homotopy type of a class of simplicial complexes. To begin, we have the following definition:

Definition 6.1 Consider all graphs on n vertices. A *monotone graph property* is some collection P of those graphs such that if $G \in P$ and $G' \subset G$, we have $G' \in P$. That is, we can remove edges without losing the "property."

For example, the collection of disconnected graphs on n vertices is a monotone graph property, since if we remove an edge from a disconnected graph it will remain disconnected. Other examples are: 2-colorable, sub i-regular (each vertex has $\leq i$ edges incident to it), acyclic.

From a monotone graph property P we can make a simplicial complex in the following way: let V, the vertex set, be the set of all $\binom{n}{2}$ edges on our n vertices. Then if some collection of edges form a graph $G \in P$, the set of those edges is a simplex in K, our simplicial complex. We can now ask questions about the topology of a simplicial complex of a monotone graph property.

Since the sets V and K are fairly large, it may seem like an unweildy simplicial complex to analyze. DIscrete Morse theory proves to be a useful tool in studying such spaces. We will illustrate this with the simplicial complex of disconnected graphs: If a set of edges form a disconnected graph on n vertices, that set is a simplex. Our goal is to construct an acyclic discrete vector field on this simplicial complex. We begin as follows: Take every nonempty simplex $\alpha \in K$, such that α does not contain the edge $(1,2)$. Then, if $\beta = \alpha + (1,2)$ is also a simplex, draw an arrow $\{\alpha \subset \beta\}$. Call the resulting discrete vector field V_{12}.

Now we ask: which simplices are unpaired? First, we have $\{(1,2)\}$ itself (since we don't draw arrows from the empty simplex). Suppose some other simplex α is unpaired. If $(1,2) \in \alpha$, then α is paired with $\alpha - (1,2)$. If $(1,2) \notin \alpha$, then it is possible that $\alpha + (1,2)$ is connected. This is true when the graph of α has two connected components, one containing the vertex 1 and the other containing the vertex 2.

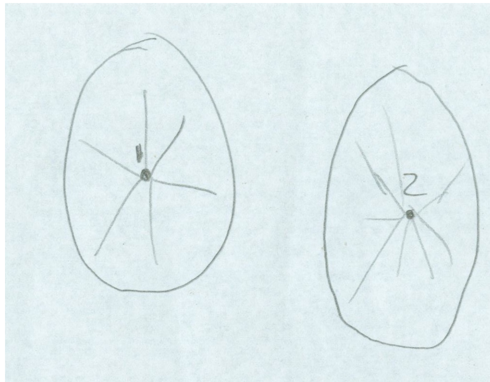

Now, suppose α is some unpaired simplex in V_{12}. Then suppose the vertex 3 is contained in the connected component of vertex 1. If α does not have the edge $(1,3)$, then we can draw an arrow from α to $\alpha + (1,3)$. If α does have the edge $(1,3)$, and $\alpha - (1,3)$ is unpaired in V_{12}, draw an arrow from $\alpha - (1,3)$ to α. It may be the case that $\alpha - (1,3)$ is paired in V_{12}, this happens when $\alpha - (1,3)$ is the union of three connected components, one containing vertex 1, one with vertex 2, and the last with vertex 3.

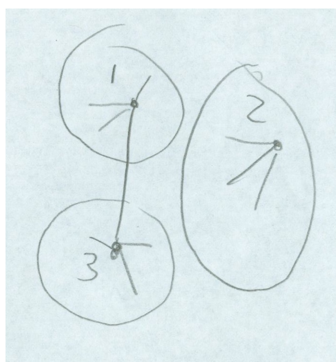

If the vertex 3 is instead in the connected component of vertex 2, we can do an analogous treatment. Call the resulting discrete vector field V_3. Then the only (nonempty) unpaired simplices in V_3 are $\{(1,2)\}$ and any simplex whose graph G satisfies $G - (1,3)$ has three connected components or $G - (2,3)$ has three connected components.

Now we look at the vertex 4. If α is unpaired in V_3, assume $(1,3) \in \alpha$, so $\alpha - (1,3)$ has three connected components. Then the vertex 4 is in the component with 1, 2 or 3. Then we can do the same analysis with the edge $(1,4)$, $(2,4)$ or $(3,4)$, whichever is applicable. Once we have exhausted these possibilities, we are left with V_4, and the unpaired simplices have graphs that are formed of four connected components, then edges added in between $1,2,3,4$ to form a tree with those vertices. Two possibilities are below (circles are connected components):

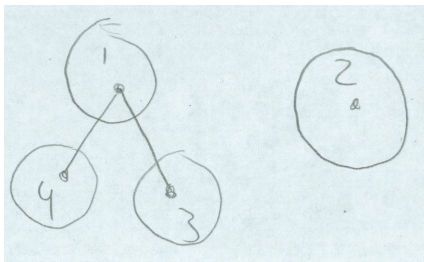

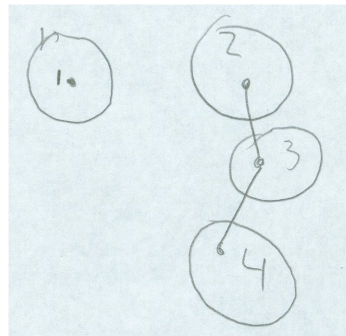

Repeating this process, we will eventually have a discrete vector field V on the simplicial complex such that the only unpaired simplices are the point $\{(1,2)\}$ and the union of two connected trees, one tree having root 1, the other having root 2, and labels increasing along rays. There are $(n-1)!$ such trees, and each has $(n-2)$ edges (points), so they correspond to $(n-3)$-*simplices.*

Lastly, we have to check that our discrete vector field V is acyclic. Suppose $\alpha_0, \beta_0, \alpha_1, \beta_1, \ldots, \alpha_{r+1}$ is a closed $V-$path. Then there is an arrow from α_0 to β_0. Suppose they are paired for the first time in V_i, $i \geq 3$. Then α_1 must either be the head of an arrow in V_i, or the tail of an arrow in V_{i-1}, and since the $V-$path continues it must be the latter. This means α_{r+1} is the tail of an arrow in $V_{i-(r+1)}$, hence it cannot be equal to α.

This means we can define a discrete Morse function on the simplex of disconnected graphs such that the only critical simplices are a point ant $(n-1)!$ $(n-3)$-simplices. Hence:

Theorem 6.2 The simplex of disconnected graphs is homotopy equivalent to the wedge of $(n-1)!$ $(n-3)-$spheres.

Using discrete Morse theory, many other similar results are possible. This is discussed in depth in [4]. To conclude this paper, we present some of the results below. Discrete Morse theory plays a prominent role in the discovery of the homotopy classes of these graph properties.

Graph Property	Homotopy Type
Forest	$\bigvee S^{n-2}$
Bipartite	$\bigvee S^{n-2}$
Disconnected	$\bigvee_{(n-1)!} S^{n-3}$
Not 2-connected	$\bigvee_{n-2)!} S^{2n-5}$
Not 3-connected	$\bigvee_{(n-3)\frac{(n-2)!}{2}} S^{2n-4}$

This concludes our investigation of discrete Morse theory.

ACKNOWLEDGMENTS

I'd like to thank my mentor Jim Fowler for sticking with me throughout this whole process, and suggesting a handful of immeasurably useful directions and resources.

REFERENCES

[1] M. M. Cohen. *A Course in Simple-Homotopy Theory.* Springer, 1973.
[2] R. Forman. *A User's Guide to Discrete Morse Theory.* Seminaire Lotharingien de Combinatoire. p. 48, 2002, article B48c.
[3] R. Forman. *Topics in Combinatorial Differential Topology and Geometry.* IAS/Park City Mathematics Series. Volume 14, 2004.
[4] J. Jonsson. *Simplicial Complexes of Graphs.* Springer-Verlag, Berlin, 2008.