On Discrete Morse Theory

Brais Macknik-Conde

July 12, 2026

A very large simplicial complex may contain many simplices, making it's topology difficult to study. Discrete Morse theory provides a combinatorial method for replacing such a structure with a simpler analogue, while still retaining the essential homology of the original simplicial complex. This paper is an introductory exposition into discrete Morse theory. The main theorem of discrete Morse theory states that an acyclic Morse matching on a simplicial complex produces a homotopy equivalent cell complex with one cell for each critical complex. It builds the theory starting from elementary definitions and is meant for all readers, even those without background in any topological field. This expository paper also does not require any knowledge of linear algebra or calculus, simply a level of mathematical maturity enough to read the text while absorbing the information as it is provided, as the later content of the paper is built almost entirely from the definitions and theorems from the beginning. Throughout the paper are many examples to aid in understanding of the text, but I recommend that the reader experiment with the definitions on paper as they read, as a deeper understanding can only be found through working with the definitions by hand. There are three main sections to the paper. The first are just definitions on and surrounding simplicial complexes, to build the groundwork that will be used extensively in the later sections of the paper. The second section is on elementary collapses, the Euler characteristic, and Betti numbers. Elementary collapses provide our first way to simplify simplicial complexes and turn them into simpler structures. The Euler characteristic and Betti number give us ways to study and measure the topology of simplicial complexes without doing any elementary collapses. Finally, the last section is on discrete Morse theory. It covers the method discrete Morse theory uses to simplify and find the topology of simplicial complexes. Throughout the last section we work on an example simplicial complex, where we look at it's Morse matching, construct it's cell complex, simplify it's cell complex, and find it's topological structure.

1 On Simplicial Complexes

1.1 What are simplices and simplicial complexes?

A simplicial complex is a structure composed of simplices. To understand the workings of a simplicial complex, we need to understand what a simplex is. A simplex is a building block of a simplicial complex. A simplex is simply defined as a set of its vertices:

$$[v_n] := \{v_0, v_1, \ldots, v_n\}$$

Above is a set of $n+1$ vertices, which creates an n-simplex. An n-simplex can be thought of intuitively as a triangle extrapolated to a dimension n. For examples, a 0-simplex (a simplex containing just 1 vertex) is a point, a 1-simplex is a line, a 2-simplex is a triangle, a 3-simplex is a tetrahedron, and so on. In this paper, we will use ω and ϕ to denote a simplex. Realize that, by definition, the number of vertices a simplex contains is always one more than the dimension of the simplex (notated $\dim(\omega)$).

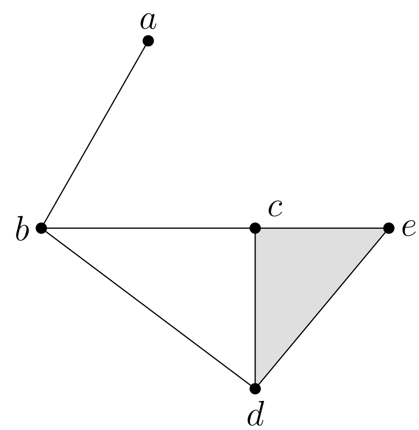

Figure 1. A simplicial complex.

Importantly, two simplices in a simplicial complex can share points with each other, like in the above example of a simplicial complex (fig. 1). We can clearly see that different 1-simplices (lines) are sharing common points. We can also see a difference between two types of triangles. One is filled in and one is not. This is because the unfilled triangle is actually three 1-simplices connected, and the filled triangle is three 1-simplices and a 2-simplex. The distinction between the two will become very important.

Notating simplices by naming each one with a new symbol and then defining what vertices the simplex contains can be very cumbersome when working with simplicial complexes with many simplices. To simplify simplex notation, we just name each simplex by what 0-simplices (vertices) it contains. For example, some of the simplices in fig. 1 are as follows:

$$a = \{a\},$$
$$ce = \{c, e\},$$
$$cde = \{c, d, e\}.$$

With this new notation, we can create notation for the structure of a simplicial complex. We can denote a simplicial complex by a set of its simplices. For the rest of the paper, L will be used as a general name for simplicial complexes. Let's denote fig. 1 using this method:

$$L = \{a, b, c, d, e, ab, bc, bd, cd, ce, de, cde\}$$

An intuitive way to think of a simplicial complex is as a structure that simply denotes how different 0-simplices are related. In Dr. Schoville's book on discrete Morse theory, he uses the example that each 0-simplex is a person, and a simplex connecting different 0-simplices notates that those people have had dinner together. Let's examine this intuition with fig. 1. We can see that b and c have had dinner, b and d have had dinner, and c and d have had dinner. However, there is no 2-simplex connecting the three of them, so they have not all had dinner together. On the other hand, c, d, and e have jointly had dinner together. Importantly, this necessarily means that c and d have had dinner, c and e have had dinner, and d and e have had dinner. This leads us into our next section.

1.2 Rules and definitions on simplicial complexes and simplices

An abstract simplicial complex L is defined as follows:

$$\omega \in L, \varnothing \neq \rho \subseteq \omega \implies \rho \in K$$

In words, all simplices ϕ that can be constructed out of the vertices in a simplex ω in the simplicial complex L must also be in L. In practice, this means that there cannot be a situation where there is a 2-simplex in a simplicial complex and one of the sides of the triangles is not in the simplicial complex, or generally, an n-simplex in a simplicial complex where the simplicial complex does not contain one of the simplices made from a subset of the vertices contained in the n-simplex.

Some examples of structures we can examine are a filled triangle and a filled triangle missing an edge.

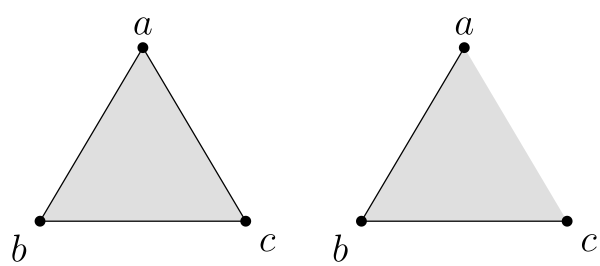

Figure 2. A filled triangle and an incomplete filled triangle.

The filled triangle on the left has abc as one of it's simplices, so because of the rule above, the simplices a, b, c, ab, ac, and bc (all of the subsets of the points of simplex abc) must also be in the simplicial complex of the filled triangle. Because this set of simplices follows the rule set on simplicial complexes, it is indeed a simplicial complex. On the other hand, the triangle on the right also has the simplex abc, but the complex doesn't contain a simplex for every subset of the points of abc; it is missing the simplex ac. Because of this it cannot follow the rule on simplicial complexes, and therefore is not a simplicial complex.

Now we will create important definitions needed for the rest of the paper.

For a simplex ω,

$$\dim(\omega) = |\omega| - 1.$$

So, for simplices named by their vertices,

$$\dim(a) = 0, \dim(ab) = 1, \dim(abc) = 2.$$

For a complex L,

$$\dim(L) = \max\{\dim(\omega) : \omega \in L\}.$$

This means that the dimension of a simplicial complex is the dimension of it's highest dimension simplex.

For simplices ϕ and ω, ϕ is a face of ω and ω is a coface of ϕ if

$$\phi \subseteq \omega.$$

This means that a simplex ϕ is a face of a simplex ω if ϕ is created from a subset of vertices from (or is a 'part' of) ω, and, for the same reason, ω is a coface of ϕ.

We can take a simplex and all of it's faces and create a simplicial complex out of them. (They create a simplicial complex by definition; the faces are all of the simplices that can be created out of the vertices of the starting simplex.) We can call this simplicial complex the complex generated by that simplex. For example, the filled triangle in fig. 2 is the complex 'generated' by abc. We can notate this simplicial complex as follows:

$$\langle abc \rangle$$

In discrete morse theory, the angled bracket notation denotes the simplest simplicial complex that contains the enclosed simplices. For example, $\langle abc \rangle = \langle abc, bc \rangle$ is a filled triangle, as the filled triangle is the simplest simplicial complex that contains the enclosed simplices in both sets of brackets. (I say simplest as a way to make the function of the angled brackets more intuitive. Really what the angled brackets do is take the union of the sets of simplices generated by each simplex within the brackets.)

We can also define a proper face. A simplex ϕ is a proper face of simplex ω if

$$\dim(\phi) = \dim(\omega) - 1$$

Notice the difference, ϕ is a face of ω if ϕ is a part of ω, but ϕ is only a proper face of ω if ϕ's dimension is only one less than ω's. For example, bc is a proper face of abc, but a is not a proper face of abc. However, a would be a proper face of ab. We can notate a proper face as follows:

$$\phi < \omega$$

If $\phi \subseteq \omega$, then the codimension of ϕ in ω is

$$\dim(\omega) - \dim(\phi)$$

The codimension just measures the difference in dimensions of a face and it's coface. For example, an edge of a triangle has codimension 1 to the triangle, and a vertex has a codimension 2 to the triangle.

The boundary of a simplex is defined as all of it's faces. For example, the boundary of the filled triangle in fig. 2 is denoted as

$$\partial(abc) = \{a, b, c, ab, ac, bc\}$$

For clarity, the notation uses a partial derivative symbol to denote a boundary, even though it has nothing to do with partial derivatives. For shorthand, the boundary of a simplex is sometimes notated as just the codimension 1 faces of it:

$$\partial(abc) = \{ab, ac, bc\}$$

This boundary is the same as the boundary before, this is just shorthand notation. Sometimes the codimension 1 faces are called the boundary faces of the simplex.

The i-skeleton of a simplicial complex L is defined as

$$L^{(i)} = \{\omega \in L : \dim(\omega) \leq i\}.$$

In words, the i-skeleton of a simplicial complex is the simplicial complex created from the simplexes of dimension less than or equal to i. For example, below is a filled triangle and it's 1-skeleton:

Figure 3. A filled triangle and it's 1-skeleton.

A facet is a simplex in a simplicial complex that is not a proper face of any other simplex in the simplicial complex. For example, in fig. 1, ab, bc, bd, cd, and cde are all facets. Notice that we make the distinction that a facet must be a proper face of a simplex because if we said that a facet must be a face of a simplex, then every simplex would be a facet, as all simplices are faces of themselves.

Facets allow us to create one of the most concise notations for simplicial complexes. We can notate any simplicial complex by taking it's facets and putting them in angled brackets, denoting the simplicial complex generated by those simplices. We can see why every simplicial complex may be uniquely generated by it's facets as follows:

Let L be a finite simplicial complex with facets $f_1, f_2, \ldots, f_n$. Because every f_1 is a simplex of L, all simplices generated by the facets must also be in L:

$$\{\omega : \omega \subseteq f_i \, \forall i \leq n\} \subseteq L.$$

Let ω be some simplex in L. Because L is finite, if ω is not a facet, we may keep enlarging ω by finding simplices properly containing it until we find a facet. This will stop once we find a facet f_j such that

$$\omega \subseteq f_j.$$

We now have the conclusion that every simplex in L is a facet or is contained in (generated by) a facet. We also have the conclusion that every simplex generated by a facet must also be in L. Therefore the simplicial complex generated by the facets of a simplicial complex L includes only the simplices within L. ■

The c-vector of a simplicial complex is a vector that counts how many of each dimension of simplex there is. For example, the c-vector of the simplicial complex L in fig. 1 is denoted as

$$\vec{c}(L) = \{5, 6, 1\}.$$

This is because there are 5 0-simplices (vertices), 6 1-simplices (lines), and 1 2-simplex (filled triangle).

Finally, let's go over some standard, good to know simplicial complexes. These have their own notation because of how commonly they are used.

The simplicial complex generated by a simplex of dimension n is referred to as $\Delta^n = \langle \omega^{(n)} \rangle$. For example, the filled triangle in fig. 2 is Δ^2.

The simplicial complex created from the boundary of some Δ^n is denoted as $S^{n-1} = \partial \Delta^n$. This can be intuitively thought of as the hollow triangle from fig. 2 (S^1) but extrapolated to whatever dimension the boundary is.

A graph is simply a simplicial complex with a maximum dimension of 1. These are commonly used as examples for advanced topics, as it's harder to denote simplicial complexes of higher dimensions on 2d paper. The hollow triangle from fig. 2 is a graph.

Lastly, if L is a simplicial complex and v is a vertex not in L, CL is a simplicial complex on L (called the cone on L) such that

$$CL = L \cup \{v\} \cup \{\omega \cup \{v\}\} : \omega \in L.$$

In words, this means that CK is the simplicial complex containing all the simplices in L, containing v, and containing all the simplices one gets by adding v as a vertex in every simplex in L. For example, below is a simplicial complex and it's cone (I chose this example as it would be much harder to show a more complicated simplicial complex's cone without introducing ambiguity):

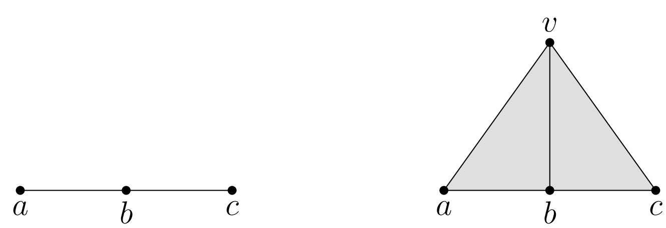

Figure 4. A simplicial complex and it's cone.

2 On Elementary Collapses and Expansions

2.1 What are elementary collapses and expansions?

This section will introduce the most important operations to know before starting discrete Morse theory. The elementary collapse and elementary expansions are two operations that take a simplicial complex and remove or add some part to it such that the fundamental structure / shape is the same. We will learn exactly what it means for the structure to be the same, but right now, we will just define what the elementary operations are.

Let L be a simplicial complex where $\phi, \omega \in L$ such that ϕ is a codimension 1 face of ω. We can call the pair of simplices

$$(\phi, \omega)$$

a "free pair" if ϕ's only non-trivial coface (coface that isn't ϕ itself) is ω. We can then remove the free pair (ϕ, ω) as an "elementary collapse:"

$$L \searrow L - \{\phi, \omega\}$$

The $\searrow$ notation indicates that on the left of the $\searrow$ is a simplicial complex before and elementary collapse and on the right of the $\searrow$ is the simplicial complex after the collapse. We can see that the simplicial complex after the collapse is the same as the simplicial complex before, just with the free pair taken out. For example, we can see how we can apply elementary collapses to a filled triangle as follows:

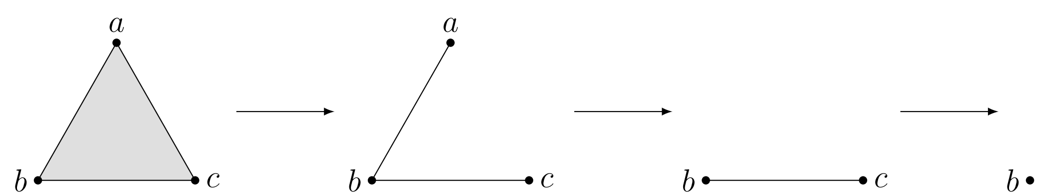

Figure 5. A collapse sequence on a filled triangle.

We can see in the above example that in the original filled triangle, $(ab, abc), (ac, abc)$, and (bc, abc) are all of the free pairs that we could remove. Then in the second simplicial complex, (a, ab) and (c, cb) are the only free pairs. Then in the third simplicial complex we only have (b, bc) and (c, cb) as free pairs. Finally, in the final complex, there isn't a free pair to choose, as there is no non-trivial coface of b for b to pair with. Also notice that in some stages of the collapse there are simplices that weren't free pairs that later became free pairs. For example, in the initial filled triangle, (a, ab) are not free pairs, as ab is not the only non-trivial coface of a (abc is also a coface of a). However, once the first collapse happened, (a, ab) became a free pair, as the other coface of a (abc) was collapsed away. Let's look at a more complex collapse sequence:

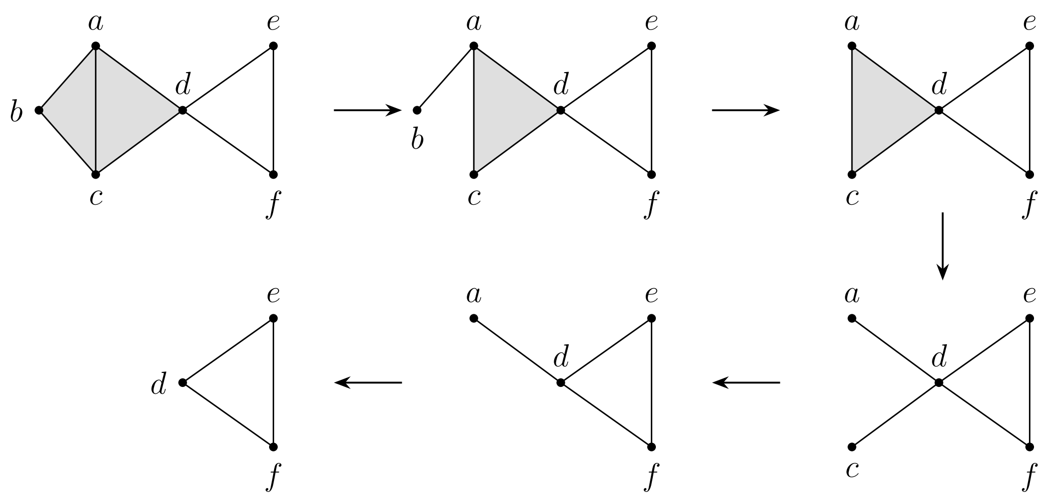

Figure 6. Collapse sequence of a simplicial complex.

The important part to notice is what part of the simplicial complex stays and what is collapsed away. The parts of the simplicial complex that elementary collapses keep, the parts crucial to the structure (again, what exactly these parts are will be discussed later) are kept. We can get an intuition that loops (like hollow triangles) are important and general filled space (like filled triangles) are not inherently important to the structure we are interested in.

Another important operation is an elementary expansion. An elementary collapse simply adds two simplices to a simplicial complex, as long as if the operation was done in reverse it would be a valid elementary collapse. There are other ways to define elementary expansions without basing it on the definition of an elementary collapse, but doing this allows us to later prove things regarding the elementary collapse and know that it hold for the elementary expansion, as we based it on the collapses.

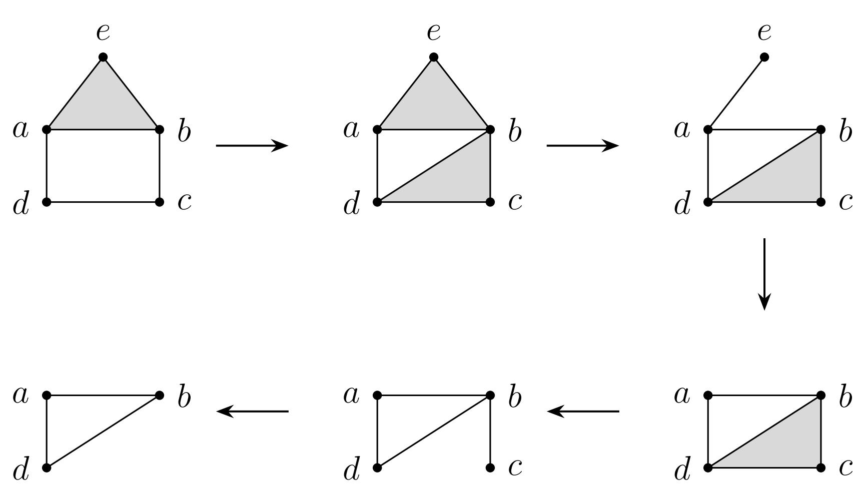

Figure 7. Example of an elementary expansion in a collapse sequence.

Notice above that the first elementary operation is an elementary expansion. This expansion allowed us to turn what was a 4 sided 'hole' into a 3 sided hole. Without using expansions, realize that it would have been impossible to collapse the 4 sided hole into a simpler 3 sided hole. In this manner, when we want to take a larger simplicial complex and collapse it into a simpler, easier to manage simplicial complex, we may need to do so using a sequence of both elementary collapses and expansions. However, it may not always be clear when a simplicial complex is fully contracted, and these elementary operations aren't able to tell us. For example, see the next example:

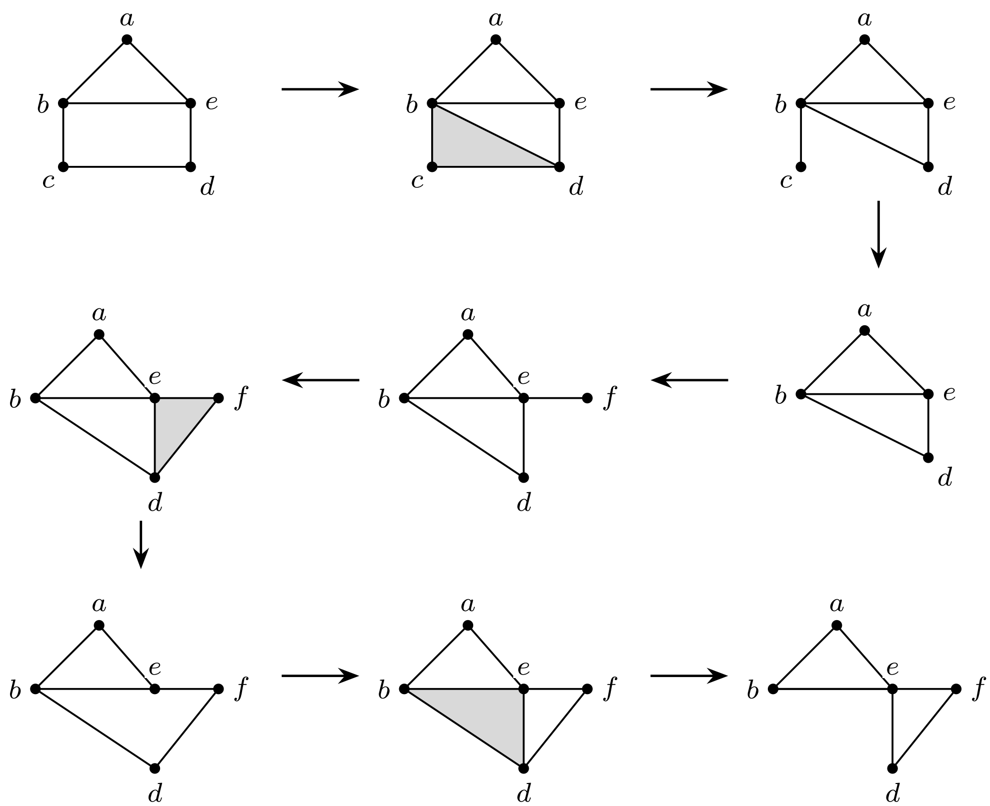

Figure 8. Collapse sequence of a simplicial complex.

This is an example of a collapse sequence where, even though we made completely valid collapses and expansions, we ended up with a simplicial complex that, while it may seem simpler as we have two triangular holes, is just as easy to analyze the structure of as the first. The first simplicical complex had two connected holes and so did the last, and the sequence of collapses helped us make nearly zero progress towards a simpler simplicial complex to examine. Even the c-vector of the two simplicial complexes is the same. In order to understand what's happening and how we can use elementary operations to truly make studying simplicial complexes easier, we need to know more about what how elementary operations preserve the structure of simplicial complexes.

2.2 Why do elementary operations preserve structure?

Let's start by making it clear that, from the knowledge we currently have, we cannot prove why elementary operations preserve structure. This is because we don't even know how to define the structure of a simplicial complex. For the rest of the paper, a main motivation is going to be to define how to even define the structure of a simplicial complex and then how to prove that our elementary operations preserve that definition of structure.

We already have defined one form of structure, and that is simply that a collection of simplices is a simplicial complex. So let's start by proving that the collection of simplices remaining after an elementary collapse is a simplicial complex. The claim we want to prove is that if we have a free pair (ϕ, ω) in simplicial complex L, then $L - \{\phi, \omega\}$ is also a simplicial complex.

This claim is surprisingly easy to prove. Let's start with how we defined a simplicial complex. We said a simplicial complex was defined as a set of simplices such that, for any simplex, if we could construct a new simplex from a set of the vectors contained within the original simplex, then that new simplex must also be a part of the simplicial complex. We had to use this language because we didn't know the definition of a face of a simplex. However, now we can rewrite our rule to be that any face of a simplex in our simplicial complex must also be in our simplicial complex. Let L be a simplicial complex such that (ϕ, θ) is a free pair. Let ρ be some simplex in L. Let

$$L' = L - \phi, \omega.$$

The only reason L' could not be a simplicial complex is if one of it's simplices was missing a face. The only simplices that were removed in L' were ϕ and ω. So we simply need to prove that ϕ and ω are not faces of any of the simplices in L'. Let's start with ω. ω cannot be a face of any ρ in L' because then ϕ would have another coface other than ω, which would make (ϕ, θ) not a free pair. And ϕ cannot be a face of ρ because then ϕ would have another coface other than ω.

And elementary expansions also result in a valid simplicial complex, as a criteria for a valid elementary expansion is that is results in a valid simplicial complex.

2.3 The Euler Characteristic

Our second measurement for the structure of a simplicial complex will be the Euler characteristic. Let's see how to calculate the Euler characteristic of a simplicial complex:

The Euler characteristic of a simplicial complex L ($\chi(L)$) is calculated from L's c-vector:

$$\chi(L) = \sum_{i=0}^{n} (-1)^i c_i = c_0 - c_1 + c_2 - \cdots + (-1)^n c_n.$$

This creates an alternating sum of the number of each dimension of simplex in L. For example, the Euler characteristic of a filled triangle is 1, while the Euler characteristic of the boundary of a 2-simplex is 0:

c_0 (vertices) $= 3$
c_1 (edges) $= 3$
c_2 (filled triangles) $= 1$
$3 - 3 + 1 = 1$

$c_0 = 3$
$c_1 = 3$
$3 - 3 = 0$

Figure 9. Euler characteristic of a filled and boundary triangle.

However, the Euler characteristic does not completely determine the structure of a simplicial complex; two simplicial complexes with different structures may share the same Euler characteristic. For example:

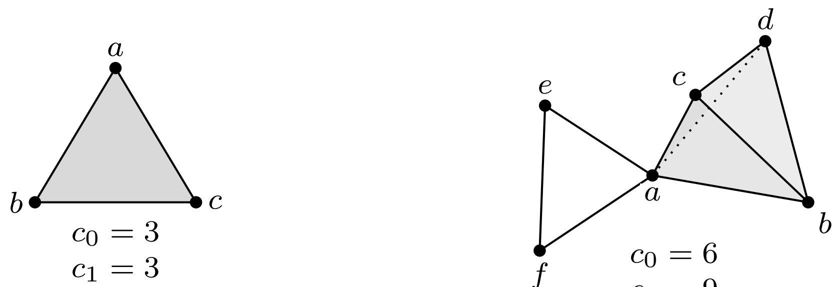

$c_0 = 3$
$c_1 = 3$
$c_2 = 1$
$\chi = 3 - 3 + 1 = 1$

$c_0 = 6$
$c_1 = 9$
$c_2 = 4$
$\chi = 6 - 9 + 4 = 1$

Figure 10. Euler characteristic of a filled triangle and of a boundary triangle connected to a boundary tetrahedron at the vertex.

In the above example, on the left we have an ordinary filled triangle, and on the right we have S^2 with S^1 connected to it at a vertex. Notice in the above example that there is no way to do elementary operations to turn one of these simplicial complexes into the other, despite the two shapes sharing an Euler characteristic. To be able to continue measuring criteria on simplicial complexes and to be able to distinguish between examples like these where the Euler characteristic fails, we will later in the paper learn more criteria on simplicial complexes.

Let's now show that elementary operations preserve the Euler characteristic:

Let (ϕ, ω) be a free pair in a simplicial complex L. Let

$$\dim(\omega) = n$$

and

$$\dim(\phi) = n - 1.$$

(The above dimensions of simplices is true for any free pair by the definition of a free pair.) This means in a free pair one simplex of dimension p and one of dimension p-1 will be removed. The Euler characteristic before and after looks like this (the c-vector shown is $\vec{c}(L)$ for both examples below):

$$\chi(L) = c_0 - c_1 + c_2 - \cdots + (-1)^{p-1}c_{p-1} + (-1)^p c_p + \cdots + (-1)^n c_n.$$

$$\chi(L - \{\phi, \omega\}) = c_0 - c_1 + c_2 - \cdots + (-1)^{p-1}(c_{p-1} - 1) + (-1)^p(c_p - 1) + \cdots + (-1)^n c_n.$$

The elementary collapse changes the Euler characteristic by

$$-(-1)^{p-1} - (-1)^p.$$

Clearly,

$$-(-1)^{p-1} - (-1)^p = 0.$$

So elementary operations do not change the Euler characteristic. Let's see an example:

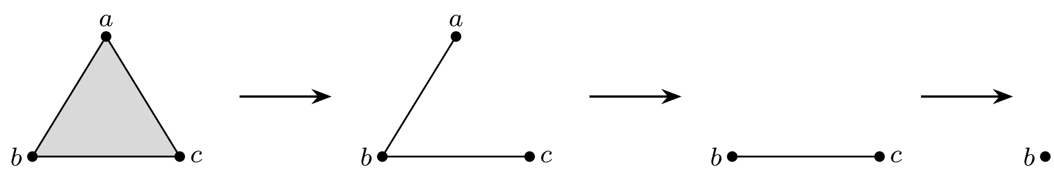

$f(L) = (3, 3, 1)$, $\chi(L) = 3 - 3 + 1 = 1$; $f(L_1) = (3, 2)$, $\chi(L_1) = 3 - 2 = 1$; $f(L_2) = (2, 1)$, $\chi(L_2) = 2 - 1 = 1$; $f(L_3) = (1)$, $\chi(L_3) = 1$

Figure 11. Euler characteristic of a filled triangle and it's collapse sequence.

Look at the above example in reverse for an example showing the Euler characteristic being preserved for elementary expansions.

Now that we have proved that elementary operations preserve the Euler characteristic, we can intuit a useful corollary that can help us see why the Euler characteristic may be useful. Because the Euler characteristic never changes, no matter the elementary operations done, realize that a simplicial complex of a certain Euler characteristic can never be turned into a simplicial complex of a different Euler characteristic with a collapse sequence. So if two simplicial complexes have different Euler characteristics, they are definitively different simplicial complexes with different structures. This doesn't help us to distinguish whether two simplicial complexes are the same or not when they share the same Euler characteristic, but it's a powerful tool when the Euler characteristics are indeed different. One can see the usefulness for telling two simplicial complexes' structures apart with the Euler characteristic, because if we had two very large simplicial complexes, there would be no practical way to try every collapse sequence to see if we could turn one into the other to see if they shared the same structure. Now we can just calculate the Euler characteristic, which is as easy as counting how many of each dimension of simplex there is, and if the Euler characteristic is different, then the structures are for sure different.

One type of simplicial complex we haven't looked at yet is a disjoint simplicial complex. It's important to mention, as it is one the types of simplicial complexes that the Euler characteristic "confuses" with others the most. A disjoint simplicial complex is just a simplicial complex where some parts of it are completely disconnected from others. See below an example of the Euler characteristic of a disjoint simplicial complex and a non-disjoint simplicial complex:

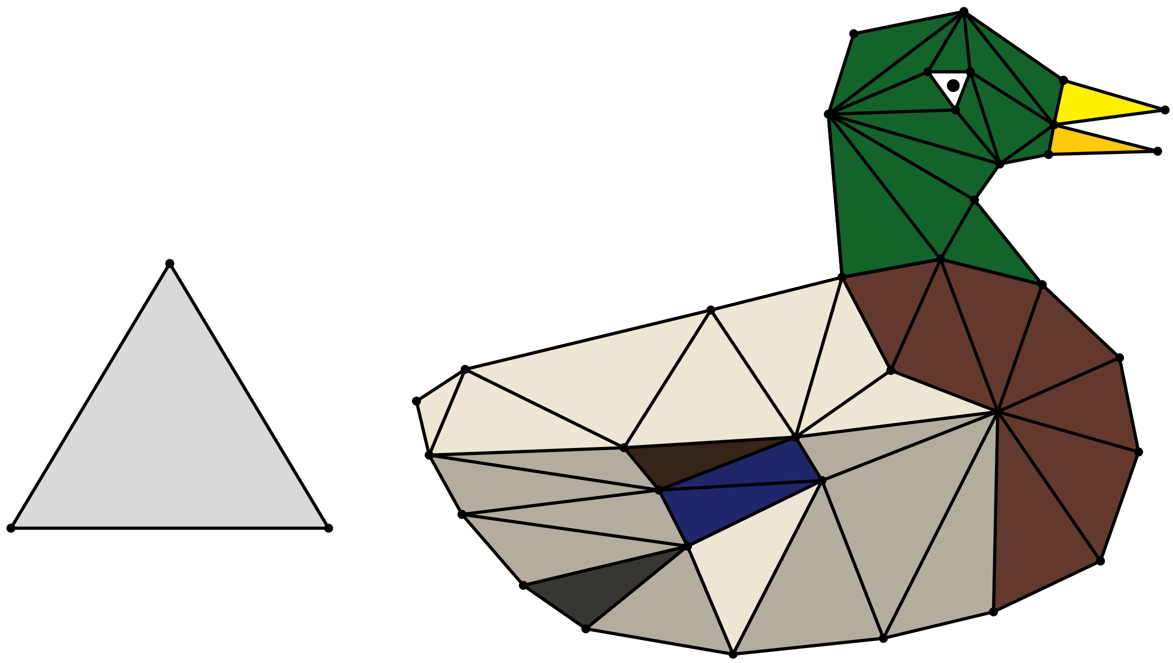

$$f_0 = 3 \quad f_1 = 3 \quad f_2 = 1 \quad \chi = 3 - 3 + 1 = 1$$

$$f_0 = 37 \quad f_1 = 80 \quad f_2 = 44 \quad \chi = 37 - 80 + 44 = 1$$

Figure 12. Euler characteristic of a filled triangle and of a disjoint simplicial complex consisting of a duck (collapses to a hollow triangle) and its pupil (a vertex).

Because the Euler characteristic counts all even dimension simplices the same and all odd dimensions simplices the same, if we allow disjoint simplicial complexes (which are just as valid as non-disjoint simplicial complexes) then we can quite easily create simplicial complexes that share the same Euler characteristic with other structures. Here in this example we see a filled triangle and a hollow triangle plus a vertex share the same Euler characteristic; I took the 2-simplex of the filled triangle and turned that into a disjoint 0-simplex.

2.4 Betti Numbers

We've seen that, even though the Euler characteristic is clearly useful (not just because it can distinguish between spaces, but also because of how easy it is to calculate it), it has limitations as seen in fig. 10 and fig. 12. A characteristic that, while more difficult to calculate, can better distinguish between different simplicial complex structures are Betti numbers.

Counting Betti numbers without doing it by eye requires a complicated process involving boundary maps and linear algebra, so we will not be covering that in this paper. Instead we will go over what we can see counting them by hand.

Before we learn what Betti numbers are, let's go over some quick notation. The Betti vector of a simplicial complex is similar to the f-vector of a simplicial complex. We simply put the 0^{th} Betti number in the first slot in the vector and so on. We count specific Betti numbers with the following notation. For example, let's say that our simplicial complex L has one 0^{th} Betti number. Then we would say:

$$\beta_0(L) = 1$$

The Betti vector is notated as follows:

$$\vec{\beta}(L) = \{\beta_0, \beta_1, \ldots, \beta_n\}$$

Now that we understand the notation, let's go into what the Betti numbers really are. We won't be going into their true definition which is based on homology, as that topic is too complex for this paper. However, we can get a strong intuition for what Betti numbers measure. Betti numbers measure the amount of "cycles" of connected simplices in the simplicial complex at specific dimensions. For example, β_0 (the 0^{th} Betti number) measures the number of cycles of 0-simplices. We can think of a cycle of vertices as just vertices that are connected. So β_0 measures how many portions of a simplicial complex are there that don't have 0-simplices connected in some way. Intuitively, β_0 measures how many disjoint sections there are in a simplicial complex. Let's continue with what the Betti numbers mean in the intuitive sense, it doesn't make sense to try to make formal definitions while missing the homology needed. β_1 represents the number of loops of 1-simplices found, so a hollow triangle would have

$$\beta_1 = 1.$$

Betti numbers continue like this from here, where β_n would count the number of n dimensional "holes." Let's see an example where the Euler characteristic fails to help us distinguish between different simplicial complexes, but that the Betti numbers can help us:

$c_0 = 3, \quad c_1 = 3$
$\chi = 3 - 3 = 0$
$\beta_0 = 1, \quad \beta_1 = 1$

$c_0 = 6, \quad c_1 = 6$
$\chi = 6 - 6 = 0$
$\beta_0 = 2, \quad \beta_1 = 2$

Figure 13. Euler characteristic and Betti vector of a hollow triangle and of a disjoint simplicial complex consisting of two hollow triangles connected at a vertex and a vertex.

So in this figure we can see where the Euler characteristic couldn't help us find a distinction between the two simplicial complexes, but by comparing the Betti numbers we can see that they are indeed different. Let's see some of our other diagrams we struggled to differentiate using the Euler characteristic earlier. For example, in fig. 12, the Betti vector for the filled triangle is simply

$$\vec{\beta} = \{1\},$$

while the Betti vector for the duck is

$$\vec{\beta} = \{2, 1\},$$

We can see this because the filled triangle is one section (no disjoint parts) and has no holes. The duck on the other hand has a hole at the eye and two disjoint sections, the pupil, and the rest of the body. We had one more set of simplicial complexes we couldn't distinguish between using the Euler characteristic in fig. 10. Neither example is disjoint, so

$$\beta_0 = 1$$

for both simplicial complexes. The left example has no holes, so the filled triangle, as before, has a Betti vector of

$$\vec{\beta} = \{1\}.$$

However, the right simplicial complex has a 1-dimensional hole connected to a 2-dimensional hole. That means it's Betti vector is

$$\vec{\beta} = \{1, 1, 1\}.$$

So now we have distinguished between every example where the Euler characteristic failed so far. However there are simplicial where the structures are different, yet the Betti numbers and Euler characteristic are the same between them. There isn't enough space to show these examples, but I encourage the reader to look at a filled triangle and a triangulation of the projective plane, which share Euler characteristic 1 and Betti vector $\{1\}$, even though they are fundamentally different structures homologically.

An interesting fact of Betti numbers that can be proved with homology (we will not be covering the proof) is that you can find the Euler characteristic of a simplicial complex simply from it's Betti vector. Simply take the alternating sum of the Betti numbers, just like we did with the f-vector, and this will retrieve the Euler characteristic of the simplicial complex:

$$\chi = \sum_{i=0}^{n} (-1)^i \beta_i = \beta_0 - \beta_1 + \beta_2 - \cdots + (-1)^n \beta_n.$$

The Betti number also does not change with elementary operations, which can also be proved using homology.

3 On Discrete Morse Theory

We have been focusing on Elementary operations until now. We have seen that they are operations for changing the shape of a simplicial complex without changing the structure. However, elementary collapses only act locally; they pair two codimension 1 simplices as a free pair and remove them. With discrete Morse theory we get a broader way to do this process. Instead of finding one free pair at a time, discrete Morse theory allows us to pair multiple simplices at the same time throughout the entire simplicial complex using similar but different rules to elementary collapses. The simplices that are paired are considered free and removable, though we don't actually take them out, we just note which ones are paired together. Once we have have added as many pairs as we can to a simplicial complex, the remaining unpaired simplices are labeled as critical simplices. We take these critical simplices and add their equivalents to a cell complex (which follow different rules than simplicial complexes). Then according to gradient paths (how those critical simplexes were related within the original simplicial complex) those cells are then connected in the cell complex. The resulting cell complex can be compared to the cell complex of any other simplicial complex, and if they match, then those two simplicial complexes are the same homologically / structurally. If they don't match then they aren't. This is the kind of distinction we have been searching for.

3.1 Morse Matchings and Critical Simplices

A discrete vector field (or a Morse matching) on a simplicial complex L is a collection of Morse pairs

$$(\phi^{(p)}, \omega^{(p+1)})$$

such that $\phi < \omega$ and every simplex of L appears in at most one pair. In words, a morse matching contains a collection of

Let's say we have a filled triangle:

$$L = \langle abc \rangle.$$

Then of course

$$L = \{a, b, c, ab, bc, ac, abc\}.$$

With what rules of creating Morse matchings we know so far (which we will promptly add to) we could do the following Morse matching on L (fig. 14):

$$(a, ab), \qquad (b, bc), \qquad (c, ac).$$

All of the above simplices are "matched." The only unmatched simplex is *abc*. So *abc* must be critical right?

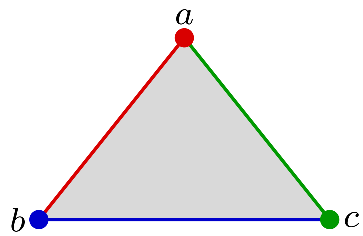

Figure 14. Incorrect Morse matching on a filled triangle, where Morse pairs are color matched.

3.2 Gradient Paths and Acyclicity

Not so fast, we still have one more rule of Morse matchings to add. This is the rule of acyclity. However, before we can define this rule, we need to define gradient paths. A gradient path is a sequence of simplices in a simplicial complex

$$\phi_0^{(p)} < \omega_0^{(p+1)} > \phi_1^{(p)} < \omega_1^{(p+1)} > \cdots > \phi_r^{(p)}$$

such that (ϕ_i, ω_1) is a matched pair and ϕ_{i+1} is another codimension-1 face of ω_i. In words, a gradient path follows a path from the lower dimension simplex of a pairing to the higher dimension one, where then we choose a different codimension 1 simplex of the higher dimension simplex. Then we go to that lower dimension simplex's pair, and so on. Let's see an example of a morse matching (which follows the rule of acyclity) and we can analyze the gradient paths we see:

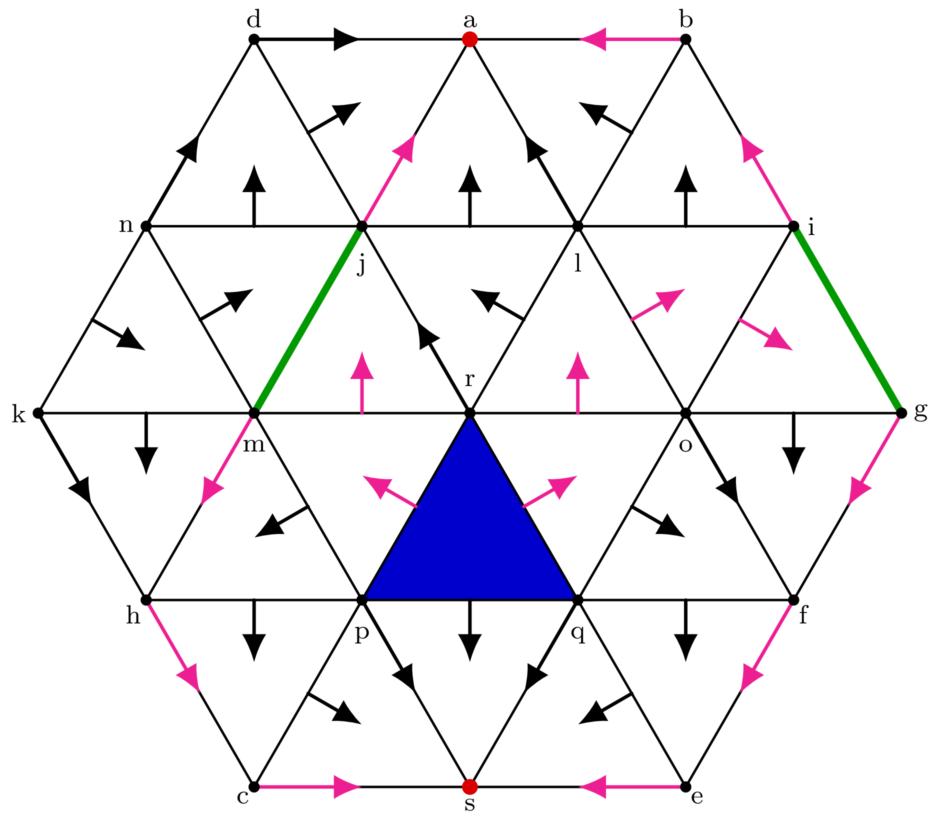

Figure 15. Large Morse matching with highlighted critical vertices and magenta arrows showing the gradient paths between critical simplices.

Let's start with what the arrows mean. The arrows represent a morse matching (ignore the color for a minute). For example, the arrow pointing from vertex c to vertex s represents that (c, cs) is a Morse matching. Or the arrow pointing from edge hp to face chp represents that (hp, chp) is a Morse matching. The start of the arrow represents the lower dimension simplex of a Morse matching, and the big end of the arrow is placed on top of the higher dimension simplex of said Morse matching. This is sometimes called "arrow notation." With this diagram we can quite clearly see the gradient paths on the simplicial complex. Let's try picking a simplex (excluding the critical vertices highlighted in red, green, and blue). If it has an arrow starting on it, follow the arrow to it's end. If an arrow doesn't start on it do nothing. Starting on either scenario one should now be on either a 1-simplex or a 2-simplex. From here choose a codimension-1 simplex that is not paired with this simplex. Then continue this pattern of following arrows and choosing codimension-1 faces. Eventually one reaches a critical simplex, where the pattern of following arrows and choosing codimension-1 faces stops. Either there are no codimension-1 faces, or to continue down a gradient path one would need to choose codimension-1 faces twice in a row. These gradient paths between critical simplices will become very important later. They are highlighted in magenta on the simplicial complex. Now that we understand how gradient paths work, we can move on to the rule of acyclity.

All of our previous rules acted quite locally around the pair we were trying to make. If we were studying a pair of simplices to see if they would follow the rules, a broken rule could usually be found quite quickly looking locally around the simplices we were planning to pair, if there was a broken rule to be found. Instead the rule of acyclity goes as follows:

For any single possible gradient path

$$\phi_0^{(p)} < \omega_0^{(p+1)} > \phi_1^{(p)} < \omega_1^{(p+1)} > \cdots > \phi_r^{(p)}$$

in a Morse matching, any simplex may not appear twice.

This means that a gradient path cannot cycle, hence the rule of acyclity. With this rule in mind, let's revisit why fig. 14 is an invalid Morse matching. Let's create arrow notation for fig. 14 below:

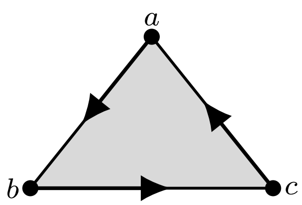

Figure 16. An incorrect Morse matching on a filled triangle with arrows showing pairings.

If we follow the gradient path, we can see the following:

$$a < ab > b < bc > c < ca > a.$$

A cycle in a gradient path. This breaks the rule of acyclity. So a different, correct Morse matching on a filled triangle could be

$$(ab, abc), \qquad (b, bc), \qquad (c, ac).$$

As seen in the arrow notation in , there are no gradient paths that cycle.

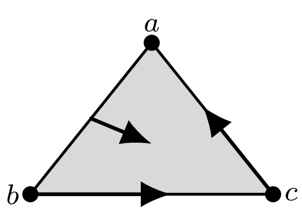

Figure 17. A correct Morse matching on a filled triangle with arrows showing pairings.

3.3 The Main Theorem and Cell Complexes

[Discrete Morse Theorem] Let L be a finite simplicial complex with an acyclic Morse matching. If m_p is the number of critical p-simplices, then L is homotopy equivalent to a cell complex with exactly m_p cells of dimension p.

Now that we understand the rule of Acyclity, we can get to what lets us actually find the inherit structure of our simplicial complexes, and what role Discrete Morse theory has in that. The main theorem of discrete Morse theory states that an acyclic Morse matching on a simplicial complex produces a homotopy equivalent cell complex with one cell for each critical complex. That means that we need to our critical simplices and put them in a cell complex, then connect them according to the gradient paths, and then analyzing the structure that takes shape within the cell complex. Before we can do any of that, we need to understand what a cell complex is. A cell complex is a structure similar to a simplicial complex. Instead of being constructed out of simplices, a cell complex is constructed out of cells. A cell in a cell complex is similar to a simplex in a simplicial complex. A 0-simplex is basically the same as a 0-cell. They are both vertices in their respective structures. A 1-cell is a line, like how a 1-simplex is a line in a simplicial complex. However, a line in a cell complex follows slightly different rules. While, like in a simplicial complex, a line can connect two vertices, in a cell complex a line can also connect a vertex to itself:

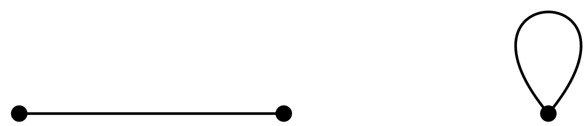

Figure 18. Two cell complexes. The left cell complex shows a 1-cell connecting two 0-cells. The right cell complex shows a 1-cell connecting a 0-cell to itself.

From there onward cells follow roughly the same rules, just applied at their respective dimensions. An n-cell fills an n-dimensional hole. For example, a 2-cell fills a 2-dimensional hole:

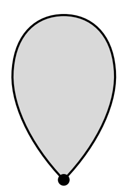

Figure 19. A cell complex of a filled loop, created from putting a 2-cell within the hole created by a 1-cell connecting a 0-cell to itself.

3.4 Constructing and Simplifying the Morse Cell Complex

Now that we understand how cells work in a cell complex, let's see how we can take our critical simplices and place them within cell complexes. For example, take a look at fig. 15. fig. 15 has 5 critical simplices. There is one critical 2-simplex, two critical 1-simplices, and two critical 0-simplices. We take these critical simplices and take their equivalent cells and create a cell complex out of them. So we need to make a cell complex out of one critical 2-cell, two 1-cells, and two 0-cells. The problem we face now is we don't know how to connect these cells. With that arrangement of cells there are many different cell complexes we could construct:

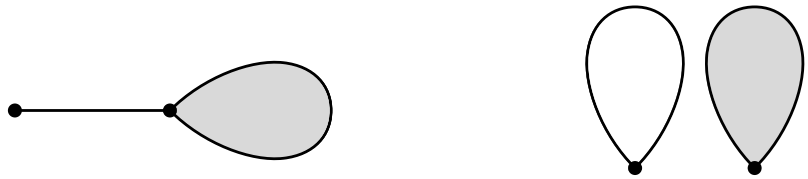

Figure 20. Two distinct cell complexes, both made from just two 0-cells, two 1-cells, and one 2-cell.

To determine how to construct the cell complex out of our cells, we need to take a look at the gradient paths. The gradient paths between critical simplices indicate how they should be connected together. For example, we can see that both critical 1-simplices in fig. 15 have gradient paths to both critical 0-simplices (shown as the arrows in magenta). Additionally, the 2-simplex has gradient paths connecting it to the two 1-simplices. With this in mind, we can construct the correct cell complex given by our Morse matching of fig. 15:

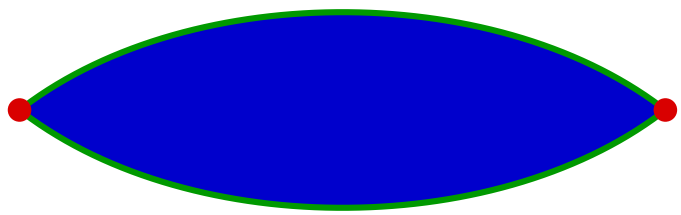

Figure 21. Cell complex for fig. 15, with the cells color coded according to which critical simplex they represent in the original simplicial complex.

However, this is not where our story ends. We have more we can do with this cell complex. Now we can do elementary collapses within the cell complex to simplify the cell complex. An elementary collapse in a cell complex follows the same rules as an elementary collapse in a simplicial complex. If we have a pair of cells

$$(\phi^{(k)}, \omega^{(k+1)})$$

such that ϕ is not a face of any other cell, then we can perform an elementary collapse on ϕ and ω (this also works in reverse like for elementary expansions, like they did in simplicial complexes). We can simplify our fig. 15 as follows:

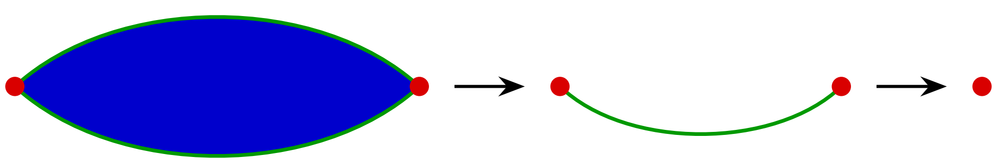

Figure 22. Collapsing of cell complex for fig. 15. The cell complex collapses to a single point.

So once we are satisfied with the cell complex we end up with out chosen simplicial complex, we can now compare this cell complex directly with cell complexes of other simplicial complexes. Cell complexes created in the long winded manner we just did represent the full homology of of the simplicial complex, so there is no need to use Betti numbers or the Euler characteristic if we have computed the cell complex for a simplicial complex. For our chosen example, we can see that fig. 15 is topologically equivalent to a point because it's cell complex is a vertex, which is also the cell complex a simplicial complex containing just a point makes.

3.5 Different Morse Matchings and Critical-Cell Bounds

Importantly, note that a Morse matching doesn't need to include as many matches as possible. For example, this is a valid Morse matching on a filled triangle:

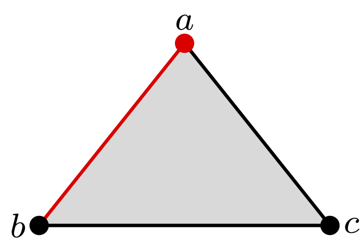

Figure 23. Morse matching on a filled triangle, where just one pair is matched in the entire simplicial complex.

With this example, all of these unpaired simplices would be critical, and following gradient paths, we would construct a cell complex as shown in fig. 21, which would collapse to a point as shown in fig. 22, still showing that the filled triangle is topologically equivalent to a point.

Additionally, note that there can be multiple Morse matchings for a single simplicial complex, and a Morse matching doesn't need to include as many matches as possible. For example, a different Morse matching for fig. 15 is as follows:

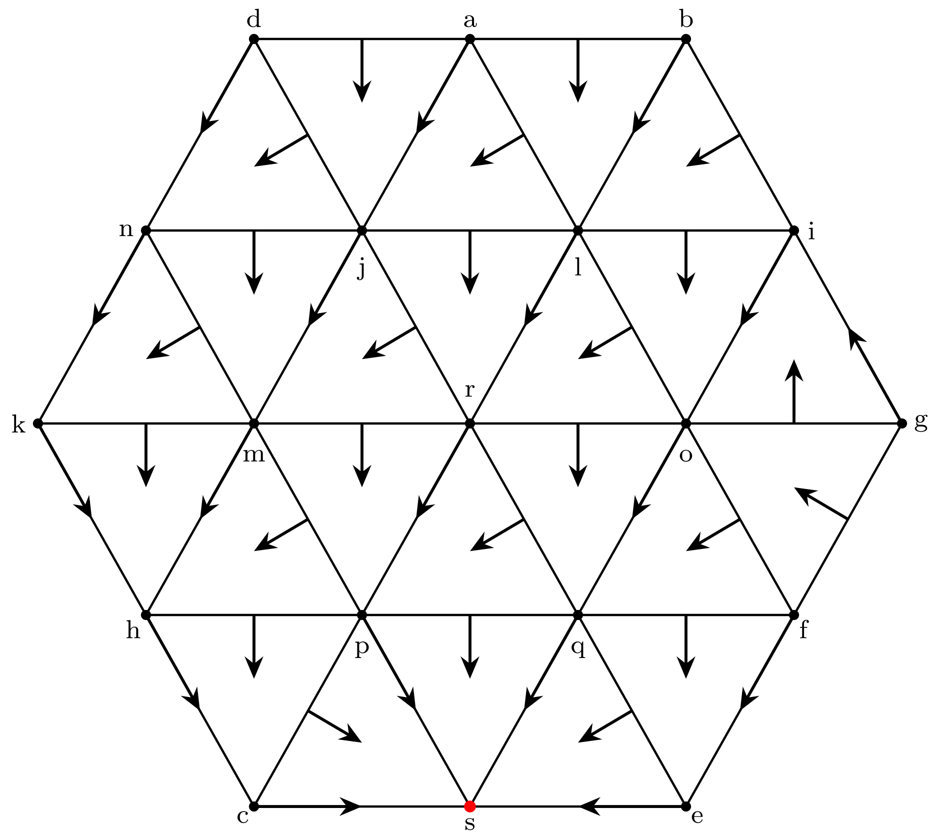

Figure 24. Different Morse matching for fig. 15.

This Morse matching leaves just a critical vertex. This would have saved us all of the trouble of simplifying the cell complex from before. The fact that different Morse matching can be made means that we can reduce the complexity of our cell structures just by finding “better” Morse matchings. But how can we know if we have the best Morse matching we can get or not? Well, using Betti numbers we can define a lower bound on the number of critical simplices of each dimension there must be. (Sometimes it is impossible to make a Morse matching on some simplicial complexes that meet this lower bound, but if one manages to reach this lower bound with a Morse matching on a simplicial complex, they can be confident that the cell complex the critical simplexes generate is fully simplified down.) The lower bound on the number of k-dimensional critical simplices is simply β_k, the k-dimensional Betti number.

References

[1] Robin Forman, *Morse Theory for Cell Complexes*, Advances in Mathematics, 134(1), 90–145, 1998.

[2] Robin Forman, *A User's Guide to Discrete Morse Theory*, Séminaire Lotharingien de Combinatoire, 48, 2002.

[3] Nicholas A. Scoville, *Discrete Morse Theory*, Student Mathematical Library, vol. 90, American Mathematical Society, Providence, Rhode Island, 2019.

[4] Jan Philipp Bullenkamp, Theresa Kaiser, Florian Linsel, Susanne Krömker, and Hubert Mara, "Discrete Morse Theory Segmentation on High-Resolution 3D Lithic Artifacts," *it – Information Technology*, vol. 66, no. 6, pp. 191–205, 2024. doi: `https://doi.org/10.1515/itit-2023-0027`.