Discrete Morse Theory and Persistent Homology

Chung-Ping Lai

Advisor: Christine Escher

Oregon State University
College of Science
Department of Mathematics

December, 2020

Abstract

In this paper we explain at the theoretical level how discrete Morse theory can provide us a more efficient approach to compute persistent homologies. In achieving so we also provide a framework for discrete Morse theory to be applied to persistent homology for other purposes.

Contents

1 Introduction

In the novel field of Topological Data Analysis (TDA), simplicial complexes are among the central objects of study. In particular, persistent homology emerged as a powerful tool in data analysis. This naturally gives rise to the need to efficiently compute the homology of various complexes.

On the other hand, Morse Theory is a powerful tool in differential topology and differential geometry which has produced important results foundational to modern geometry and topology. This inspired Robin Forman along with others to take similar approaches to study CW complexes and PL manifolds. Forman eventually produced a discrete version of the Morse inequality (see [8]) and developed what we now know as the discrete Morse theory (see [3]).

Among the ideas carried over from Morse theory is the notion of Morse homology, a homology theory whose homology groups are isomorphic to standard simplicial homology groups, see [12] for details. As part of his work, Forman formulated Morse homology discretely on CW complexes.

This particular result leads to an approach to efficiently compute homology groups. In general, the complexity of computing homology groups is related to the number of cells. Meanwhile the computation of Morse homology only uses the "critical" cells, and thus can be significantly faster to compute than the simplicial homology while still giving the same result.

The purpose of this expository Master's paper is to introduce at a theoretical level how Milschaikow and Nanda applied the above idea to optimize the computation of persistent homology [4], and in the process provide a framework for other applications of discrete Morse theory in persistent homology.

We will first define some of the algebraic and topological notions in Section 2. In Section 3 we will introduce the theory of abstract cell complexes, which tackles the issue that TDA mainly works with abstract simplicial complexes whereas discrete Morse theory was originally introduced on CW complexes. Then in Section 4, we will establish a variant of discrete Morse theory with a focus on Morse complexes, the analogue of Morse homology in discrete Morse Theory. We will briefly introduce the persistent homology of a filtration of abstract cell complexes in Section 5. Finally in Section 6 we will construct the Morse Filtration which involves fewer cells but has isomorphic persistent homology groups as the original filtration.

2 Background

In this section we will briefly establish notations and some useful results from algebraic topology.

Notation 2.1. Modules

Let R be a ring. Let $R(Y)$ denote the free R module generated by the (finite) set Y.

Moreover, for such $R(Y)$, we assign by default a pairing structure $\langle , \rangle$, so that for $Y = \{y_1, y_2, ..., y_n\}$, and all r_i, r_i''s in R, $\langle \sum_{i=1}^{n} r_i y_i, \sum_{i=1}^{n} r_i' y_i \rangle := \sum_{i=1}^{n} r_i r_i'$.

Definition 2.2. Chain Complex

A chain complex $\mathcal{C} = \{(C_d, \partial_d)\}$ *is a sequence of abelian groups C_d, called the chain groups connected by homomorphisms ∂, called the boundary operators and commonly denoted $\partial_d : C_d \to C_{d-1}$.*

$$... \xrightarrow{\partial_{d+2}} C_{d+1} \xrightarrow{\partial_{d+1}} C_d \xrightarrow{\partial_d} C_{d-1} \xrightarrow{\partial_{d-1}} ...$$

Moreover, we require $\partial_d \partial_{d-1} = 0$ for every index d.

The index d assigned to each group is generally called the dimension as elements of the groups C_d usually correspond to topological objects of $\dim d$. It is a common practice in topology to suppress the index of the boundary operators and denote them by ∂ since the dimension is usually clear from context.

For the chain complexes that will appear in this paper, all the groups C_d will be free R-modules for a fixed ring R. Moreover this ring R is uniform across all chain complexes presented. This extra property is prevalent in topology but is especially important for this paper.

Remark 2.3. *Given a chain complex $\mathcal{C}$ with chain of groups C_d. For convenience we sometimes ignore the dimensions completely and consider the chain complex C as one module, the direct sum of all the modules C_d. In that case the ∂'s can altogether be considered as one single function $\mathcal{C} \to \mathcal{C}$.*

A CW complex is another topological structure deeply related to our topic. We use the definition from Lundell and Weingram[1] among the many equivalent ones as our reference.

Definition 2.4. CW Complex

A (finite) CW complex *is a set $\bar{X}$ with a (finite) collection of functions $X := \{\alpha^{(d)} : D^d \to \bar{X} \mid d \in \mathbb{N}\}$, where D^d is the closed unit disc of dimension d with boundary S^{d-1}, satisfying the following conditions.*

- *Each $\alpha^{(d)} \in X$ is injective on $D^d - S^{d-1}$.*
- *$\bar{X}$ is the disjoint union of $\{\alpha^{(d)}(D^d - S^{d-1}) \mid \alpha \in X\}$.*
- *For any $\alpha^{(d)} \in X$, $\alpha^{(d)}(S^{d-1}) \subseteq \bar{X}^d$, where $\bar{X}^d := \bigcup\{\beta^{(d')}(D^{d'}) \mid \beta^{(d')} \in X, d' < d\}$ and is called the d-skeleton of the complex.*

We usually denote $\alpha^{(d)}, \beta^{(d')} \in X$ just as α, β, suppressing the dimension indicator d. Furthermore we abuse notation and let them also denote their images, which are the cells of X. Note that $\bar{X}$ has a topological structure induced by the cells and that is the topological structure of the CW complex.

Finally we state the definition of an abstract simplicial complex, another very important structure in TDA. Here we reference the definition of Munkres [10].

Definition 2.5. Abstract Simplicial Complex

A (finite) abstract simplicial complex *is a (finite) collection of finite sets X. Each $\alpha \in X$ is called a cell of X and has the property that $\beta \subset \alpha \in X \implies \beta \in X$. We define the* dimension of *α as its cardinality minus* 1*. And we define the* vertex set *$\bar{X}$* of *X as the union of all $\alpha \in X$.*

In addition, in this paper we always index the finite vertex set $\bar{X}$ as $\{x_0, x_1, ..., x_n\}$. Then we can orient each cell $\alpha = \{x_{i_0}, x_{i_1}, ..., x_{i_d}\}$ by the parity of the permutation of $i_0, i_1, ..., i_d$.

3 Abstract Cell Complexes

In order to work in the combinatorial setting of computers, discrete constructions are strongly preferred, with abstract simplicial complexes being the prime candidates in Topological Data Analysis (TDA). However, in order to apply ideas from discrete Morse theory, which will be introduced later in this paper, we often have to change the shape of the cells drastically, which could not be done in the theory of simplicial complexes where all cells must be simplicies. In the original discrete Morse theory developed by Forman in 1998[3], Forman built his theory on CW complexes and took advantage of the continuous properties to bend the cells. This approach is powerful but the continuous nature also makes it difficult to use for computers.

For this reason, before we describe discrete Morse theory, we will introduce a different cell complex that shares the flexibility provided by CW complexes which made discrete Morse theory possible, while also being discrete so that it is computer-friendly.

Amazingly, Tucker [7] already constructed such a complex many decades ago in 1936. With some refinement, Mischaikow and Nanda(2013) [4] were able to make Tucker's theory the groundwork of an alternative formulation of Forman's discrete Morse theory. And this is the theory we will introduce in this section.

3.1 Abstract Cell Complexes

In this particular theory, the definition of a complex, which we call an abstract cell complex, relies on defining a boundary incidence function. This has the advantage of giving rise to a natural boundary operator, and in general is able to encapsulate rich information of boundary behaviors. The trade off is that defining meaningful functions between complexes becomes exceedingly difficult, if not impossible. Moreover, so far exploration based on this approach is mostly limited to complexes with finitely many cells.

Luckily, we will see in the next subsection every abstract cell complex has an associated chain complex. The close relation between the two allows us to apply many properties of the well-studied chain complexes to build up our theory, most notably the homology theories.

Definition 3.1.1. Abstract Cell Complex

An abstract cell complex, denoted by (X, b), consists of a set $X := \bigsqcup_{d\in\mathbf{Z}} X_d$ with each X_d being a finite set, and a boundary incidence function b. Elements in X_d will be called abstract cells of dimension d. Sometimes we also write $\alpha^{(d)}$ to denote that α has dimension d.

The boundary incidence function b is defined by $b : X \times X \longrightarrow R$, where R is an integral domain with units $U(R)$, and is called the coefficient ring *of the complex, and b satisfies the following conditions.*

(i) For any $\alpha, \beta \in X$, $b(\alpha, \beta) \neq 0 \implies \dim(\alpha) = \dim(\beta) + 1$;

(ii) For any $\alpha, \gamma \in X$, $\sum_{\beta\in X} b(\alpha, \beta)b(\beta, \gamma) = 0$.

In this paper, we won't specify the coefficient rings of different abstract cell complexes since we never put together abstract cell complexes with different coefficient rings.

The intuition of the definition will be more apparent, as we see later in this section how b can be derived from operations we are very familiar with in the theory of homology, which motivated condition (ii) of its definition.

For now, to better illustrate the idea, we will work on an example.

Example 3.1.2. *Let* $R = \mathbb{Z}$, $X_2 = \{\alpha\}$, $X_1 = \{\beta_1, \beta_2, \beta_3\}$, $X_0 = \{\gamma_1, \gamma_2, \gamma_3\}$. *Then define* b *as* 0 *except for the following cases.*

$$b(\alpha, \beta_1) := 1 \qquad b(\alpha, \beta_2) := 1 \qquad b(\alpha, \beta_3) := 1$$
$$b(\beta_1, \gamma_2) := -1 \qquad b(\beta_1, \gamma_3) := 1 \qquad b(\beta_2, \gamma_3) := -1$$
$$b(\beta_2, \gamma_1) := 1 \qquad b(\beta_3, \gamma_2) := -1 \qquad b(\beta_3, \gamma_2) := 1$$

It is easy to see that condition (i) of Definition 3.1.1 is satisfied. For condition (ii), note that condition (i) implies $b(\alpha', \beta')b(\beta', \gamma') \neq 0$ only when $\dim \alpha' = \dim \gamma' + 2$. Hence we only have to check the pairings of $(\alpha, \gamma_1), (\alpha, \gamma_2), (\alpha, \gamma_3)$.

For (α, γ_1), $\sum_{\beta \in X} b(\alpha, \beta)b(\beta, \gamma_1) = b(\alpha, \beta_2)b(\beta_2, \gamma_1) + b(\alpha, \beta_3)b(\beta_3, \gamma_1)$
$= (1)(1) + (1)(-1) = 0$. The computation for other pairings is similar.

As we will in the next subsection, there is a deep relation between the boundary incidence function b and the boundary operator we know from the theory of CW complexes or simplicial complexes. In fact, much like an abstract simplicial complex, despite relying only on set theory and algebra in the definition, intuitively it is very much still a theory of topology and is applied to solve topological problems. To display that, we show how we can use a CW-complex with oriented cells to describe Example 3.1.2.

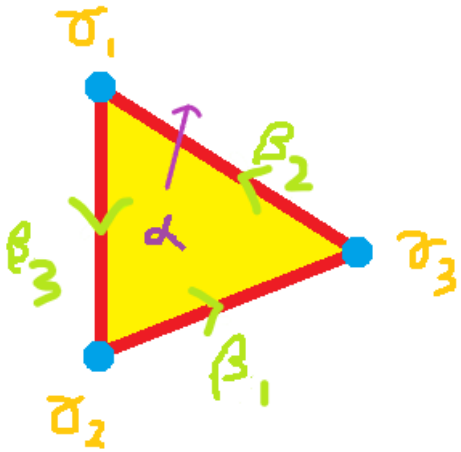

Assuming the ring R for the abstract cell complex is $\mathbb{Z}$. The idea is that X_0, X_1, X_2 corresponds to the set of vertices, edges, and surfaces, and $X_3, X_4, ...$ corresponds to the higher dimension cells. On the other hand the boundary incidence function b assigns the degree of the attaching map from the boundary of a dimension $d+1$ cell α' (in the sense of a CW complex) to a dimension d cell β', to the pairing (α', β'). This particular interpretation motivated condition (i) so that b is only ever non-zero on pairings of the form $X_{d+1} \times X_d$. For pairings where the dimensions don't match, we still assign zeros to them for coding convenience, and also for when we later apply

ideas of Remark 2.3.

More specific to Example 3.1.2 we now explain how we arrive at the triangle image above. By inspecting the elements of X_0, X_1, X_2, $\gamma_1, \gamma_2, \gamma_3$ are the vertices, $\beta_1, \beta_2, \beta_3$ are the edges, while α is the face. Then we figure out how the pieces stick together by inspecting b.

We begin with the three vertices and start adding edges to it. $b(\beta_1, \gamma_2) = -1$ is interpreted as the edge β_1 leaving γ_2 once, while $b(\beta_1, \gamma_3) = 1$ is interpreted as β_1 arriving γ_3 once, akin to the adjacency matrix of directed graphs in graph theory. β_2 and β_3 are treated similarly.

Finally, for α, as a face its boundary would be a loop. Then $b(\alpha, \beta_1) = b(\alpha, \beta_2) = b(\alpha, \beta_3) = 1$ would be interpreted as the loop traversing each edge once and with the same orientation as the edge as defined above. Choosing an order of $\beta_1 \to \beta_2 \to \beta_3 \to \beta_1$ and we form an anti-clockwise loop for α, using the right-hand rule we have α orienting "upward" as a face.

Moreover, we can see that all finite abstract simplicial complexes also admit an abstract cell complex structure.

Example 3.1.3. *Finite Abstract Simplicial complex*

Let X be a finite abstract simplicial complex with index set $\bar{X} := \{x_0, ..., x_n\}$.

Let R be any ring with 1. Let $X_d := \{\beta \in X \mid \dim \beta = d\}$

Define the boundary incidence function b by $b(\alpha, \beta) = (-1)^k$ for $\alpha = \{x_{i_0}, ..., x_{i_d}\}$ and $\beta = \{y_{i_0}, ..., \hat{y_{i_k}}, ..., y_{i_d}\}$ (the i's are ordered so that $i_0 < i_1 < ... < i_d$), and 0 for everything else.

Note that letting $b(\alpha, \beta) = (-1)^k$ is what allows condition (ii) of Definition 3.1.1 to be fulfilled.

Together with the next remark, we see that the structure of an abstract cell complex is indeed what we are looking for, a bridge between CW complexes and abstract simplicial complexes that translates to the discrete setting.

Remark 3.1.4. *Note that we can also reverse the process following Example 3.1.2 to construct an abstract cell complex from a CW-complex. More specifically, we will see in Theorem 3.2.4 how to rigorously construct an abstract cell complex from the cellular chain complex of CW-complexes.*

As a final note while CW complexes and simplicial complexes only make up a small portion of all abstract cell complexes, these are also the cases that have been of interest in the field of TDA so far.

With the definition settled for now we look at some features immediately derived from it.

Much like all other theories of complexes, the face relation between cells of the same complex play an important role in the theory. Similar to CW complexes however, certain face relations have better behavior than others, which leads to the definition of regular faces.

Definition 3.1.5. Faces and Regular Faces

Given an abstract cell complex (X, b), and any cells $\alpha, \beta \in X$.

When $b(\alpha, \beta) \neq 0$, β is called a face of α and α is called a coface of β. The relation is also denoted by $\beta \prec_b \alpha$.

When $b(\alpha, \beta) \in U(R)$, β is called a regular face of α and α is called a regular coface of β. The relation is also denoted by $\beta \preceq_b \alpha$.

In practice, no more than one face relation will be employed at a time, thus we will just use $\prec$ and $\preceq$ and assume it is based on b.

This is an analogue to faces and regular faces in CW complexes. It is worth noting however, when viewing CW complexes as abstract cell complexes, every face of a cell in the abstract cell complex always corresponds to a face of the corresponding cell, but the opposite is not necessarily true. We can see this in the simple example below. As a CW complex γ_0 is a face of β_0, but as an abstract cell complex, $b(\beta_0, \gamma_0) = 0$ since the 1-cell β_0 leaves γ_0 once and then returns to γ_0 on the other end. So the attaching map on the boundary of β_0 has degree 0. Thus γ_0 is not a face of β_0.

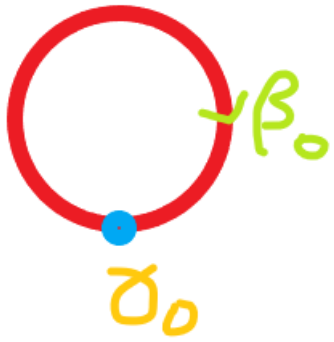

$X_0 = \{\gamma_0\}$,
$X_1 = \{\beta_0\}$.

Similar to faces, the recurring idea of subcomplexes also finds room in this theory. In fact, we will see that the idea of subcomplexes is especially important in the theory's applications in Section 5 and 6.

Definition 3.1.6. Subcomplex

Given an abstract cell complex (X, b). We say $(X', b\mid_{X' \times X'})$, abbreviated (X', b), is a subcomplex of X if $X' \subset X$, and given any $\alpha \in X'$, for every $\beta \in X$, $\beta \prec \alpha$ implies $\beta \in X'$. Moreover, for every cell $\alpha \in X'$, its dimension in X' is the same as its dimension in X.

In other words, a subcomplex X' of X is a subset of X that preserves all boundary relationships and contains all the faces of its cells.

We will then establish that a subcomplex of an abstract cell complex is also an abstract cell complex.

Proposition 3.1.7. *Given an abstract cell complex (X, b) and (X', b) a subcomplex of (X, b). Then (X', b) is also an abstract cell complex.*

Proof. We check that (X', b) satisfies the conditions of Definition 3.1.1.

(i) Let $X'_d := X_d \cap X'$ denote the set of cells in X' of dimension d. Then b automatically satisfies the dimension condition.

(ii) Given $\alpha, \gamma \in X'$. By definition $\beta \prec \alpha$ implies $\beta \in X'$. By contrapositive, $\beta \in X \setminus X'$ implies $b(\alpha, \beta) = 0$.

Hence , $\sum_{\beta \in X} b(\alpha, \beta) b(\beta, \gamma) = \sum_{\beta \in X'} b(\alpha, \beta) b(\beta, \gamma) + \sum_{\beta \in X \setminus X'} b(\alpha, \beta) b(\beta, \gamma)$
$= \sum_{\beta \in X'} b(\alpha, \beta) b(\beta, \gamma) + \sum_{\beta \in X \setminus X'} 0 \cdot b(\beta, \gamma) = \sum_{\beta \in X'} b(\alpha, \beta) b(\beta, \gamma).$

But $\sum_{\beta \in X} b(\alpha, \beta) b(\beta, \gamma) = 0$ since (X, b) is an abstract cell complex. □

Lastly, we establish a useful lemma. If two cells in a complex have an intermediate face, then there must be an additional intermediate face.

Lemma 3.1.8. *Given an abstract complex (X, b). If $\gamma \prec \beta \prec \alpha$, then there exists $\beta' \neq \beta$ such that $\gamma \prec \beta' \prec \alpha$.*

Proof. Assume no such cells exist. Then for any $\beta' \in X$ such that $\beta' \neq \beta$, by definition either $b(\alpha, \beta') = 0$ or $b(\beta', \gamma) = 0$, i.e. $b(\alpha, \beta') b(\beta', \gamma) = 0$.

Meanwhile the hypothesis $\gamma \prec \beta \prec \alpha$ implies that $b(\alpha, \beta) \neq 0$ and $b(\beta, \gamma) \neq 0$.

So $\sum_{\beta' \in X} b(\alpha, \beta') b(\beta', \gamma) = b(\alpha, \beta) b(\beta, \gamma) \neq 0$, which contradicts the definition of a complex.

□

3.2 Associated Chain Complex

We will see in this subsection that every abstract cell complex has an associated chain complex, and it is the deep relation between the two that gives the theory of abstract cell complexes power to prove some of the very useful results. In fact, Mischaikow and Nanda's version of discrete Morse theory is partially inspired by another version developed by Kozlov [5] for chain complexes.

The construction of associated chain complexes from abstract cell complexes mirrors the construction of cellular chain complexes from CW complexes, leading to homology with R coefficients.

Definition 3.2.1. Associated Chain Complex

Let $C_d(X) := R(X_d)$.

Define $\partial_d : C_d(X) \to C_{d-1}(X)$ on a d-cell by $\partial_d(\alpha^{(d)}) = \sum_{\beta^{(d-1)} \in X_{d-1}} b(\alpha, \beta)\beta$ and then extend the definition linearly to all of $C_d(X)$.

Thus we obtain the associated chain complex of (X, b), denoted by $C(X)$:

$... \longrightarrow C_d(X) \xrightarrow{\partial_d} C_{d-1}(X) \xrightarrow{\partial_{d-1}} C_{d-2}(X) \longrightarrow ...$

Remark 3.2.2. *The homology of the associated chain complex of abstract cell complexes is akin to the cellular homology of CW complexes with R coefficients, and thus also defines homology groups with R coefficients.*

We then establish that an associated chain complex is indeed a chain complex.

Proposition 3.2.3. *Given the associated chain complex $C(X)$ of an abstract cell complex (X, b). Then $C(X)$ is a chain complex i.e. $\partial_{d-1} \partial_d = 0$ for any d.*

Proof. We omit the dimension d in ∂_d. By definition of ∂ as homomorphisms, it suffices to verify $\partial^2 = 0$ for the basis elements of each $C_d(X)$ in X_d. For any $\alpha \in X_d$, we obtain

$$\partial^2(\alpha) = \partial\Big(\sum_{\beta^{(d-1)}} b(\alpha, \beta)\beta\Big) = \sum_{\beta^{(d-1)}} b(\alpha, \beta)\Big(\sum_{\gamma^{(d-2)}} b(\beta, \gamma)\gamma\Big)$$

$$= \sum_{\gamma^{(d-2)}} \Big(\sum_{\beta^{(d-1)}} b(\alpha, \beta) b(\beta, \gamma)\Big)\gamma = 0$$

The last equality follows from condition (ii) of Definition 3.1.1.

□

Hence every abstract cell complex corresponds to a chain complex of R-modules, and next we will see that this relation works in both directions. In fact, abstract cell complexes can be considered as an alternative definition of certain classes of chain complexes.

Theorem 3.2.4. *Given a fixed integral domain R, a chain complex of finitely-generated free R-modules has a canonical abstract cell complex structure.*

Moreover, the associated chain complex of the abstract cell complex is isomorphic to the original chain complex.

Proof.

Let the following sequence, denoted by C, be a chain complex, such that each C_d is a finitely-generated free R-module and the connecting maps are R-module homomorphism.

$$... \to C_{d+1} \xrightarrow{\partial_{d+1}} C_d \xrightarrow{\partial_d} C_{d-1} \xrightarrow{\partial_{d-1}} ...$$

For each d, let X_d be the basis of the dth module in the chain. Moreover define the boundary incidence function b as $b(\alpha, \beta) :=< \partial_d \alpha, \beta >$ for any $\alpha \in X_d$ and $\beta \in X_{d-1}$. Otherwise $b(\alpha, \beta) := 0$. Let $X = \bigcup_d X_d$ and we claim (X, b) is an abstract cell complex.

Condition (i) of Definition 3.1.1 follows immediately from the definition.

For condition (ii), note that by (i) we know that $\sum_{\beta \in X} b(\alpha, \beta) b(\beta, \gamma)$ can only be non-zero when $\dim(\alpha) = \dim(\gamma) + 2$. However, from $\partial^2 = 0$, we know that for any $\alpha \in X_d$,

$0 = \partial^2(\alpha) = \partial(\sum_{\beta^{(d-1)}} \langle \partial\alpha, \beta \rangle \beta) = \sum_{\beta^{(d-1)}} b(\alpha, \beta)(\sum_{\gamma^{(d-2)}} \langle \partial\beta, \gamma \rangle \gamma)$

$= \sum_{\gamma^{(d-2)}} (\sum_{\beta^{(d-1)}} b(\alpha, \beta) b(\beta, \gamma)) \gamma.$

So for any $\gamma^{(d-2)}$, $\sum_{\beta^{(d-1)}} b(\alpha, \beta) b(\beta, \gamma) = 0$.

Finally, we look at the associated chain complex of (X, b). For each d, $C_d(X) = R(X_d) \cong C_d$ since X_d is the basis of the free R-module C_d. It remains to show that the boundary maps $\partial_{C_d(X)}$ and ∂_{C_d} are the same for each d and we do so by checking the map on the basis.

For each d, given $\alpha \in X_d$, following Definition 3.2.1, $\partial_{C_d(X)}(\alpha) = \sum_{\beta^{(d-1)}} b(\alpha, \beta)\beta = \sum_{\beta^{(d-1)}} \langle \partial_{C_d} \alpha, \beta \rangle \beta$. On the other hand since $\partial_{C_d}(\alpha)$ has dimension $d - 1$ and is thus spanned by X_{d-1}, i.e. $\partial_{C_d}(\alpha) = \sum_{\beta^{(d-1)}} \langle \partial_{C_d}(\alpha), \beta \rangle \beta$.

□

In fact, mapping α, β to $\langle \partial\alpha, \beta \rangle$ is precisely how we derived the definition of $b(\alpha, \beta)$ for abstract cell complexes.

4 Discrete Morse Theory

When Forman formulated discrete Morse theory in 1998, see [3], it was done on CW-complexes with heavy reference to traditional Morse theory, in particular by providing an actual Morse function. His ideas have since then been distilled and transferred to various other contexts with modifications. Here we will apply his ideas in the context of abstract cell complexes, as was done by Mischaiwkow and Nanda [4]. This version of discrete Morse theory is derived from Kozlov's earlier version based on chain complexes, see [5]. Recall from subsection 3.2 that chain complexes are in fact closely related to abstract cell complexes. A key quality inherited by this lineage is the omission of a Morse function. Instead, a construction known as acyclic partial matching allows us to go directly to the Morse complex, a key tool in the field.

It should be noted though, that the idea of acyclic partial matchings was already known to Forman from the very beginning. Moreover, Forman was able to reveal the connection between acyclic partial matchings and discrete Morse functions. Hence it is interesting to first introduce Forman's Morse function and show its relationship to acyclic partial matchings. We will be working in the context of abstract cell complexes which we know by Remark 3.1.4 also includes CW complexes, but the key ideas remain the same as those of Forman.

4.1 Discrete Morse Function

For the most part smooth Morse theory focuses on studying the critical points of a smooth Morse function. An important insight is that it only takes a few points on the entire manifold to capture much useful information. The analogue in this discrete Morse theory is that in CW complexes and also in abstract cell complexes there are a few critical cells among the large complexes that capture a large amount of information, and a discrete Morse function can help us find and analyze the critical cells.

In smooth Morse theory, for a given manifold $\mathcal{M}$, Morse functions are smooth functions $\mathcal{M} \to \mathbb{R}$ with no degenerate critical points. When it comes to defining a discrete Morse function, a similar procedure is followed. For a given abstract cell complex (X, b), we look at functions $X \to \mathbb{R}$ and first identify the critical cells.

Definition 4.1.1. Critical Cell

Given an abstract cell complex (X, b) *and a function* $f : X \to \mathbb{R}$. *Then* $\beta^{(d)} \in X$ *is a* critical cell *if the following two conditions hold*

- $U_f(\beta) := |\{\alpha \in X_{d+1} \mid \beta \prec \alpha \text{ and } f(\alpha) \leq f(\beta)\}| = 0$;
- $L_f(\beta) := |\{\gamma \in X_{d-1} \mid \gamma \prec \beta \text{ and } f(\gamma) \geq f(\beta)\}| = 0$;

The critical cells are defined by an order-preserving property between the face relation and the order of the outputs of f. More specifically, a critical cell β of f is defined by the property that $\gamma \prec \beta \prec \alpha$ if and only if $f(\gamma) < f(\beta) < f(\alpha)$. The functions $U_f(\beta)$ and $L_f(\beta)$ count the number of $\alpha^{(d+1)}$ and $\gamma^{(d-1)}$ that violates this property, and hence they must both be 0 for any critical cells $\beta^{(d)}$.

And just as not all smooth functions are smooth Morse functions, we ask for additional properties of the $f : X \to \mathbb{R}$ that are useful for our purpose.

Definition 4.1.2. Discrete Morse function

Given an abstract cell complex (X, b)*, a* discrete Morse function *is a function* $f : X \to \mathbb{R}$ *such that for each* $\beta \in X$*, the following hold:*

- $L_f(\beta) \leq 1, U_f(\beta) \leq 1$.
- *If* $\beta \prec \alpha$ *and* $f(\beta) \geq f(\alpha)$*, then* $\beta \preceq \alpha$.

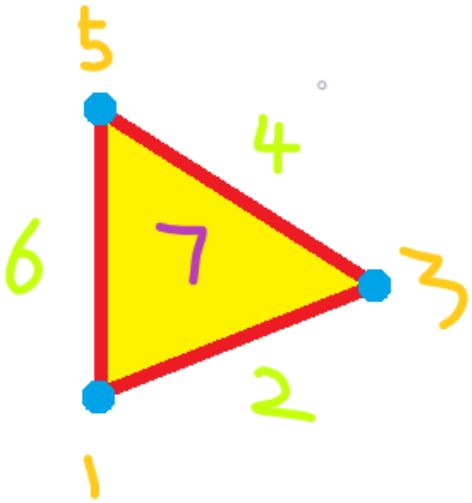

Figure 1: A discrete Morse function on Example 3.1.2.

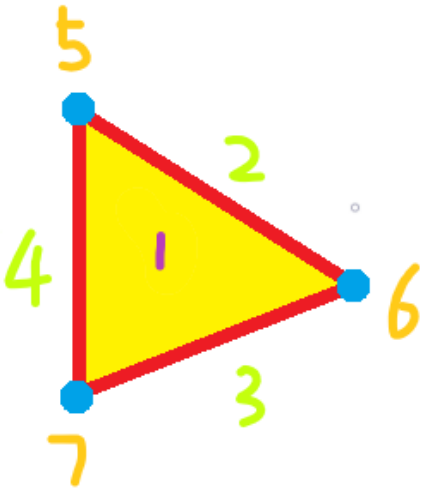

Figure 2: Not a discrete Morse function.

The first condition will allow us to eventually put all the non-critical cells into pairs (β, α) so that $\beta \prec \alpha$ and $f(\beta) > f(\alpha)$. The second condition forces $\beta \preceq \alpha$, i.e. $b(\alpha, \beta) \in U(R)$, which will allow us to divide by $b(\alpha, \beta)$.

We end the subsection by establishing a key property of a discrete Morse function that allows said pairings to happen, as we will see in Proposition 4.2.4.

Proposition 4.1.3. *Given an abstract cell complex* (X, b) *with a discrete Morse function* f*. For any* $\beta \in X$*,* $L_f(\beta)$ *and* $U_f(\beta)$ *cannot both be* 1.

Proof. Assume there exists $\beta^{(d)}$ such that $L_f(\beta) = U_f(\beta) = 1$. Then we can find $\alpha^{(d+1)}$ and $\gamma^{(d-1)}$ such that $f(\gamma) \geq f(\beta) \geq f(\alpha)$ with $\gamma \prec \beta \prec \alpha$. But by Lemma 3.1.8 there exists $\beta'^{(d)} \neq \beta$ such that $\gamma \prec \beta' \prec \alpha$. Since f is a Morse function, $L_f(\alpha) \leq 1$ and so β is the unique face of α such that $f(\beta) \geq f(\alpha)$, hence $f(\beta') < f(\alpha)$. Similarly, $U_f(\gamma) \leq 1$ and so β is the unique coface of γ such that $f(\gamma) \geq f(\beta)$, hence $f(\gamma) < f(\beta')$.

Putting everything together we get $f(\alpha) \leq f(\beta) \leq f(\gamma) < f(\beta') < f(\alpha)$. This gives a contradiction. □

4.2 Partial Matchings

An acyclic partial matching is a partial matching with the acyclic property, both of which will be defined in this section. After that we will show in Proposition 4.2.4 how a discrete Morse

function induces an acyclic partial matching. In addition we will also see that each acyclic partial matching conversely induces a discrete Morse function.

It is worth noting that Forman called a partial matching a discrete vector field and considered an acyclic partial matching a gradient vector field induced by a discrete Morse function, see [3] for details. That different names have been given to the same concept reflects the change of approach towards this theory as the field of TDA progressed, see for example [4, 5].

We begin with defining a partial matching.

Definition 4.2.1. Partial Matching (Discrete Vector Field)

For an abstract cell complex (X, b) *a* partial matching *is a partition of* X *into three subsets* A,Q *and* K*, with a bijection* $w : Q \to K$*. Moreover we require that for each* $q \in Q$*,* $q \preceq w(q)$*, i.e.* $b(w(q), q) \in U(R)$*.*

A partial matching is usually denoted by $(A, w : Q \to K)$*.*

Sometimes we also call a partial matching just as a matching.

As an example we provide a partial matching on Example 3.1.2. Let $A := \{\gamma_2, \beta_3, \alpha\}$, $Q := \{\gamma_1, \gamma_3\}$, $K := \{\beta_1, \beta_2\}$, w maps $\gamma_1 \mapsto \beta_2, \gamma_3 \mapsto \beta_2$.

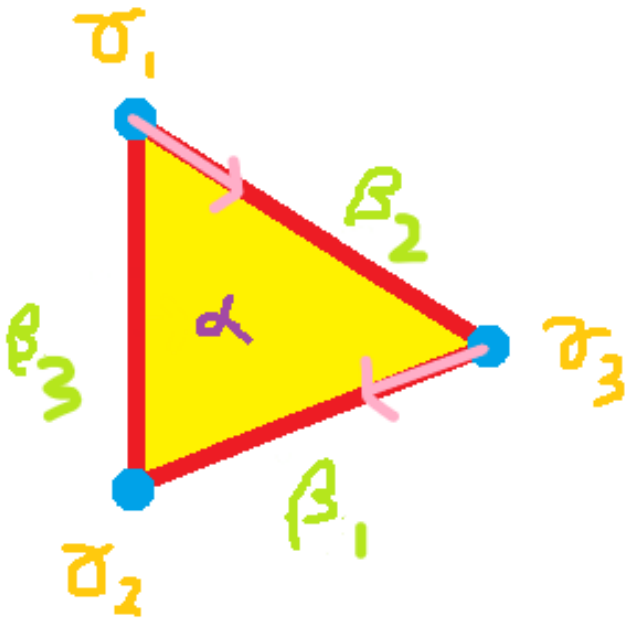

Figure 3

Graphically in Figure 3, w is represented by pink arrows. We omit the "orientation" in the figure since the orientation only shows the signs $\pm$ of $b(\alpha, \beta)$'s, however for our purpose we are only concerned with $b(\alpha, \beta) = 0$ and with $b(\alpha, \beta) \in U(R)$, but both are unaffected by signs when $R = \mathbb{Z}$.

We then define a relation on Q induced by a partial matching. And it will be acyclicity of this relation that defines the acyclic partial matchings among other partial matchings.

Definition 4.2.2. Gradient of a Partial Matching

Given a partial matching $(A, w : Q \to K)$, we define the relation $\lessdot$ on Q by defining $\beta' \lessdot \beta$ if $\beta' \prec w(\beta)$. The relation $\lhd$ on Q is defined as the transitive closure of $\lessdot$. We call $(Q, \lessdot, \lhd)$ the gradient *of $(A, w : Q \to K)$.*

Aside from defining acyclicity, the definition of gradient will remain useful in the later sections. The notation and definition of both $\lessdot$ and the transitive closure $\lhd$ are both retained as they can serve separate purposes.

Definition 4.2.3. Acyclic Partial Matching

Given a partial matching $(A, w : Q \to K)$ with gradient $(Q, \lessdot, \lhd)$. It is an acyclic partial matching if there do not exist any $\beta_1, ..., \beta_n \in Q$ such that $\beta_1 \lessdot \beta_2 \lessdot ... \lessdot \beta_n \lessdot \beta_1$.

Note that Definition 4.2.3 is equivalent to $(Q, \lhd)$ being a partially ordered set.

Take the partial matching defined in Figure 3 as an example. The only gradient relation induced in this matching is $\gamma_3 \lessdot \gamma_1$. Thus it is an acyclic matching.

Alternatively, the matching in Figure 4 induces gradient relationship $\gamma_1 \lessdot \gamma_2 \lessdot \gamma_3 \lessdot \gamma_1$ and is thus not acyclic. And graphically, we can see the arrows representing w form a "cycle".

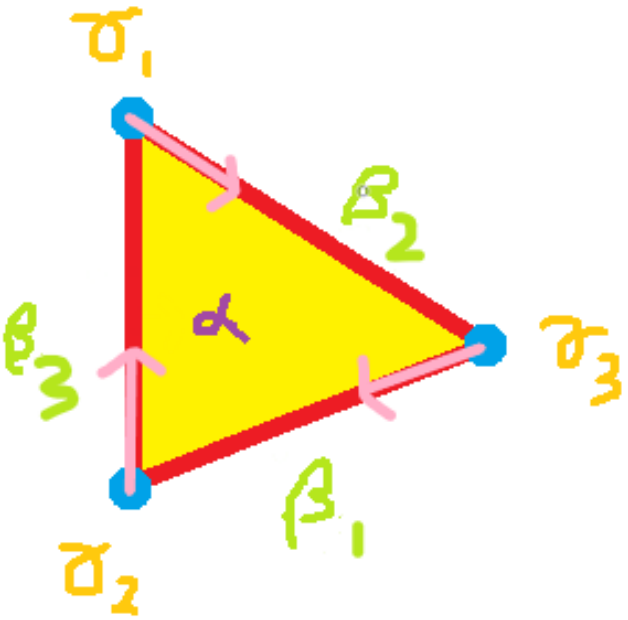

Figure 4: Another partial matching on Example 3.1.2. $A := \{\alpha\}$, $Q := \{\gamma_1, \gamma_2, \gamma_3\}$, $K := \{\beta_1, \beta_2, \beta_3\}$. w maps $\gamma_1 \mapsto \beta_2, \gamma_2 \mapsto \beta_3$, $\gamma_3 \mapsto \beta_1$.

Similar to the practice in standard Morse theory, while there can be many choices of acyclic partial matchings on a complex, most of the time we will only assign one acyclic partial matching to each complex.

In fact, the process of assigning an acyclic partial matching is very much the same as assigning a discrete Morse function. In particular, it is not too difficult to see that a discrete Morse function induces an acyclic partial matching. And this acyclic partial matching is basically what Forman first constructed as a gradient vector field, only to eventually be simplified into Definition 4.2.3, [4].

Proposition 4.2.4. *A discrete Morse function induces an acyclic partial matching.*

Proof. Given a discrete Morse function $f : X \to \mathbb{R}$. Let A denote the set of critical cells.

Let $Q := \{\beta \in X \mid L_f(\beta) = 1\}$, and let $K := \{\beta \in X \mid U_f(\beta) = 1\}$.

By Proposition 4.1.3 we know that A, Q, K are partitions of X.

For all $\beta \in Q$, by the definition of Q and K there is a unique $\alpha \in K$ such that $\alpha \succ \beta$ and $f(\alpha) < f(\beta)$. Thus we can define $w : Q \to K, \beta \mapsto \alpha$. Moreover, w is bijective, since for any $\alpha \in K$ there is a unique $\beta \in Q$ such that $\beta \prec \alpha$, which means $w(\beta) = \alpha$. Note this means that $f(\beta) > f(w(\beta))$ always holds.

It remains to check that for each $\beta \in Q$, β is a regular face of $w(\beta)$. This follows directly from Definition 4.1.2 and the fact that $\beta \prec w(\beta)$ but $f(\beta) > f(w(\beta))$.

Hence $(A, w : Q \to K)$ induces a partial matching on (X, b), and we let $\lhd$ be its gradient. We now show that this matching is acyclic.

Assume it is not. We can then find some cycle $\beta_1 \lessdot .... \lessdot \beta_n \lessdot \beta_1$ in Q where all β_i's are distinct. By the definition of gradient this means $\beta_2 \prec w(\beta_1), ..., \beta_1 \prec w(\beta_n)$. Note that for each $w(\beta_i)$, β_i is by definition of our w the unique element such that $\beta_i \prec w(\beta_i)$ but $f(\beta_i) > f(w(\beta_i))$. So the above chain of face relations implies that $f(\beta_2) < f(w(\beta_1)), ..., f(\beta_1) < f(w(\beta_n))$. Putting everything together we obtain $f(\beta_1) > f(w(\beta_1)) > f(\beta_2) > f(w(\beta_2)) > ... > f(\beta_n) > f(w(\beta_n)) > f(\beta_1)$ which is a contradiction. □

For example, the discrete Morse function on Figure 1 induces the acyclic partial matching on Figure 3 by the process above while the function on Figure 2 can't do the same.

Surprisingly, the converse of Theorem 4.2.4 is also true [3, Theorem 9.3; 6, Theorem 3.5]. And in simpler examples it is easy to imagine how we reproduce Figure 1 from Figure 3, and how Figure 4 probably isn't induced by any discrete Morse function. Since a detailed proof will deviate deeply into graph theory and is not tied to the main topic of this paper, we will just state the Theorem.

Theorem 4.2.5. [3, 6] *Any acyclic partial matching is always induced by some discrete Morse function.*

4.3 Morse Complex

In this section, we construct the Morse complex of an abstract cell complex and prove that a Morse complex has the same homology as the original complex (Theorem 4.3.7), which is one of the most important theorems in discrete Morse theory. The concept of a Morse complex is parallel to Morse homology in standard Morse theory. Both are induced by the Morse functions and both preserve the homology groups of the underlying object. While Morse homology uses critical points of Morse functions as basis of the chain complex, a Morse complex of an abstract cell complex uses critical cells.

Recall that we have discussed in Section 3 that the idea of an abstract cell complex is mostly a reformulation of the idea of a chain complex. As we will see the main advantage of this reformulation is that we can now directly construct in Definition 4.3.5 a Morse complex from the original complex. On the other hand, the reformulation also means we have to greatly modify the proof of Theorem 4.3.7.

The definition of a Morse complex begins with defining a special kind of sequence that begins and ends at A, but traverses along Q.

For notational simplicity, without changing any definitions we first expand on a few notations introduced in the previous sections.

Notation 4.3.1. *Given a partial matching $(A, w : Q \to K)$ with gradient $(Q, \lessdot, \lhd)$. For any $q \in Q, \alpha, \beta \in A$, we expand our notations as follows.*

- *$q \gtrdot \beta$ denotes $w(q) \succ \beta$; $\alpha \gtrdot q$ denotes $\alpha \succ q$.*
- *$q \rhd \beta$ (and resp. $\alpha \rhd q$) denotes that there exists $q' \in Q$ such that $q \rhd q' \gtrdot \beta$ (and resp. $\alpha \gtrdot q' \rhd q$).*
- *$\alpha \gtrdot \beta$ denotes $\alpha \succ \beta$; $\alpha \rhd \beta$ denotes that there exists $q' \in Q$ such that $\alpha \rhd q' \rhd \beta$.*
- *$w(\alpha) = \alpha$ for all $\alpha \in A$.*

Definition 4.3.2. Q-sequence and Multiplicity

Given an abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K)$. A Q-sequence is any sequence $\rho := (\alpha = q_0, q_1, ..., q_{l+1} = \beta)$ where q_i are distinct elements in Q for $1 \leq i \leq l$, while $\alpha, \beta \in A$.

We also write ρ as $\alpha \overset{\rho}{\rightsquigarrow} \beta$ and call l the stepcount of ρ, denoted as $l(\rho)$.

Then we define the multiplicity of a Q-sequence ρ as:

$$\mu_b(\rho) := \frac{\prod_{i=0}^{l(\rho)} b(w(q_i), q_{i+1})}{\prod_{i=1}^{l(\rho)} -b(w(q_i), q_i)}.$$

We then zoom into a special kind of Q-sequence which we call gradient paths.

Definition 4.3.3. (Un)saturated Gradient Paths

Given an acyclic partial matching $(A, w : Q \to K)$ on an abstract cell complex (X, b) and let its gradient be $(Q, \lessdot, \lhd)$. Fix $\alpha, \beta \in A$. A (Saturated) Gradient Path *from α to β in $(X, \lhd)$ is a Q-sequence $\alpha \overset{\rho}{\rightsquigarrow} \beta$ with the property that*
$\alpha = q_0 \gtrdot q_1 \gtrdot q_2 \gtrdot ... \gtrdot q_l \gtrdot q_{l+1} = \beta$.

An unsaturated gradient *path is any subsequence of a gradient path that preserves $\alpha = q_0$ at the beginning and $\beta = q_{l+1}$ at the end.*

Lastly we denote the set of all saturated (and resp. unsaturated) gradient paths from α to β as $\nabla_{\gtrdot}(\alpha, \beta)$ (and resp. $\nabla_{\rhd}(\alpha, \beta)$).

Note that the definition of an unsaturated gradient path is equivalent to requiring $\alpha = q_0 \rhd q_1 \rhd q_2 \rhd ... \rhd q_l \rhd q_{l+1} = \beta$.

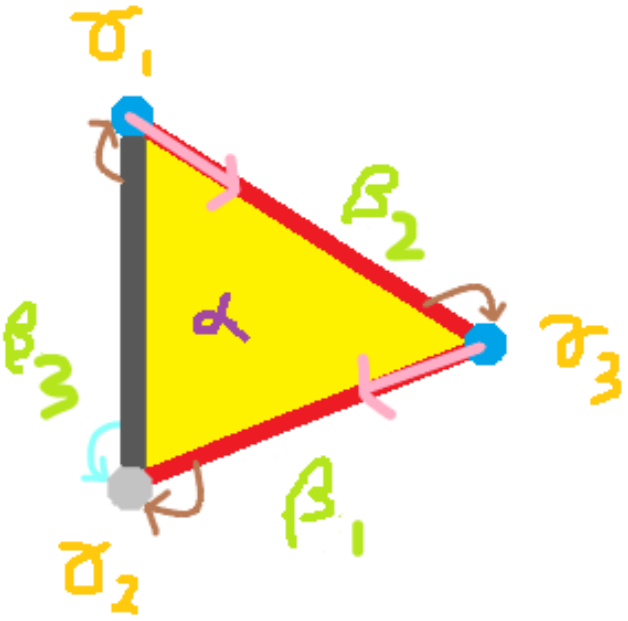

Figure 5: Two saturated gradient paths on the partial matching of Figure 3 from β_3 to γ_2, the brown path $(\beta_3, \gamma_1, \gamma_3, \gamma_2)$, and the cyan path (β_3, γ_2).

As an example Figure 5 describes two of the saturated gradient paths induced by the partial matching of Figure 3. Recall that $A = \{\gamma_2, \beta_3, \alpha\}$, $Q = \{\gamma_1, \gamma_3\}$, $K = \{\beta_1, \beta_2\}$, w maps $\gamma_1 \mapsto \beta_2, \gamma_3 \mapsto \beta_2$. So $\beta_3 \succcurlyeq \gamma_1 \succcurlyeq \gamma_3 \succcurlyeq \gamma_2$ since $\beta_3 \in A$ and $\beta_3 \succ \gamma_1$, $w(\gamma_1) = \beta_2 \succ \gamma_3$, $w(\gamma_3) = \beta_1 \succ \gamma_2$. i.e. $(\beta_3, \gamma_1, \gamma_3, \gamma_2)$ is a saturated gradient path with step count 2. The multiplicity of this gradient path is $\dfrac{b(\beta_3, \gamma_1)b(\beta_2, \gamma_3)b(\beta_1, \gamma_2)}{b(\beta_2, \gamma_1)b(\beta_1, \gamma_3)} = \dfrac{(1)(-1)(-1)}{(1)(1)} = 1$

On the other hand $\beta_3 \succ \gamma_2$ so (β_3, γ_2) is trivially a gradient path with step count 0. The multiplicity of this gradient path is simply $b(\beta_3, \gamma_2) = -1$

Before going to the main definition, we establish a small proposition on gradient paths and the multiplicity of a Q-sequence.

Proposition 4.3.4. *Given a Q-sequence $\alpha \overset{\rho}{\rightsquigarrow} \beta$ in an abstract cell complex (X, b), $\mu_b(\rho) \neq 0$ if and only if ρ is a saturated gradient path.*

Proof. Let $\rho := (\alpha = q_0, q_1, ..., q_{l+1} = \beta)$ be a Q-sequence. Then $\mu_b(\rho) \neq 0$ if and only if $b(w(q_i), q_{i+1}) \neq 0$ for every i, which is true if and only if $\alpha \succ q_1$ and $q_i \succcurlyeq q_{i+1}$ for every $i \geq 1$, i.e. ρ is a saturated gradient path.

□

The implication of the proposition is that ρ has zero multiplicity unless it is a saturated gradient path. This allows us to conveniently add or remove Q-sequences that aren't saturated gradient paths when we are working with multiplicity. This flexibility is helpful in dealing with some of the technicalities in the proof of Theorem 4.3.7.

We now have the necessary prerequisites to define the Morse complex. As stated before, an explicit construction of a Morse complex is the key motivation behind approaching our problem

with the theory of abstract cell complexes, as opposed to in the theory of chain complexes where a Morse complex can only be inductively defined.

Definition 4.3.5. Morse Complex

Given a abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K)$. Its associated Morse complex is (A, b'), where $b'(\alpha, \beta)$ is defined as the sum of multiplicities of all unsaturated gradient paths from α to β:

$$b'(\alpha, \beta) := \sum_{\rho \in \nabla_{\triangleright}(\alpha,\beta)} \mu_b(\rho) = \sum_{\rho \in \nabla_{\gg}(\alpha,\beta)} \mu_b(\rho).$$

Remark 4.3.6. *Note that Proposition 4.3.4 implies that in computing $b'(\alpha, \beta)$ it is indifferent whether we sum all ρ in $\nabla_{\triangleright}(\alpha, \beta)$, or only those in $\nabla_{\gg}(\alpha, \beta)$. In fact we can sum up all Q-sequences as well since they also have zero multiplicity.*

Figure 6: The Morse Complex induced by the matching on Figure 3.

The example on Figure 6 gives us some idea on what the Morse Complex intuitively looks like. The Morse Complex corresponding to Figure 3 consists of $A = \{\gamma_2, \beta_3, \alpha\}$, $b'(\alpha, \beta_3) = 1$, $b'(\beta_3, \gamma_2) = 0$.

In particular $b'(\beta_3, \gamma_2)$ is sum of the multiplicities of the two gradient paths described on Figure 5. Meanwhile $b'(\alpha, \beta_3) = b(\alpha, \beta_3)$ since only the trivial path connects the two.

Graphically, to construct the Morse complex we remove the non-critical cells and then, when appropriate, "stretch" each critical cell (only β_3 in this case) along the gradient paths (those on Figure 5) and through the "void" left behind, until they contact other critical cells (γ_2) at the end of the paths. This is also why information like $b(\alpha, \beta_1)$ is simply ignored, since the place β_1 occupied has been "taken over" by other critical cells.

Intuitively, the multiplicity of the gradient path records how the contact happens. The definition of $b'(\alpha, \beta)$ then aggregates all these instances of touches to find the net number of times the boundary of a critical cell α is covering another critical cell β in new Morse Complex, which is also what we want b' to represent.

From the definition we can see clearly that the Morse complex has a lot less cells than the original complex. We will now see in Theorem 4.3.7 that despite being a "simpler" complex, the homology, the key property, is preserved. This opens up a path to greatly simplify the computation of homology which will be discussed in Section 6.

Theorem 4.3.7. *Given an abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K)$). Then the associated Morse Complex (A, b') is an abstract cell complex. Moreover (A, b') has the same homology as (X, b), that is $H_*(X, R) \cong H_*(A, R)$.*

The proof of this theorem revolves around inductively removing q and $w(q)$ pair by pair from X until Q and K are empty while preserving the homology. At each step we will construct what we call a q-reduced complex which will be established in the following definition.

Definition 4.3.8. q-reduced Complex and q-reduced Matching

Given an abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K))$. For $q \in Q$, let $X_q = X \setminus \{q, w(q)\}$, $Q_q = Q \setminus \{q\}$, and $K_q = K \setminus \{w(q)\}$. Then define $b_q : X_q \times X_q \to R, (\alpha, \beta) \mapsto b(\alpha, \beta) - \dfrac{b(\alpha, q)b(w(q), \beta)}{b(w(q), q)}$.

We call (X_q, b_q) a q-reduced complex *of (X, b). Denoting $w\,|_{Q_q}$ by w, we call $(A, w : Q_q \to K_q)$ the* q-reduced matching.

We first establish a fact from this definition, we will prove then the reduced complexes and reduced matchings are exactly what they are called. Note that b_q is well-defined because $b(w(q), q) \in U(R)$ as required by Definition 4.2.1.

Proposition 4.3.9. *Given an abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K))$ and gradient $(Q, \lessdot, \lhd)$, and any of its q-reduced complexes (X_q, b_q). For any $\alpha, \beta \in X_q$, $b(\alpha, \beta) \neq b_q(\alpha, \beta)$ if and only if $\beta \prec w(q)$ and $q \prec \alpha$ in (X, b).*

In particular, for any $q' \in Q$ such that $q' \neq q$, $b(w(q'), q') = b_q(w(q'), q')$.

Proof. Note that $b(\alpha, \beta) - b_q(\alpha, \beta) = \dfrac{b(\alpha, q)b(w(q), \beta)}{b(w(q), q)}$ by definition of $b_q(\alpha, \beta)$.

We already know that $b(w(q), q) \in U(R)$. So the difference is non-zero if and only if both $b(\alpha, q)$ and $b(w(q), \beta)$ are non-zero, which by definition of $\prec$ is true if and only if $q \prec \alpha$ and $\beta \prec w(q)$.

Hence, if $b(w(q'), q') \neq b_q(w(q'), q')$, then $q \prec w(q')$ and $q' \prec w(q)$. It follows that $q \lhd q' \lhd q$, contradicting the acyclic property. □

We now show that a q-reduced complex is indeed an abstract cell complex.

Proposition 4.3.10. *Given an abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K))$. For any $q \in Q$, the q-reduced complex of (X, b), (X_q, b_q) is an abstract cell complex.*

Proof. Since X_q is a subset of X, we only need to show b_q is a boundary incidence function for X_q. Hence we show that b_q satisfies properties (i) and (ii) of Definition 3.1.1.

(i) Assume that $b_q(\alpha, \beta) \neq 0$. Then $b(\alpha, \beta)$ and $b(\alpha, \beta) - b_q(\alpha, \beta)$ must not both be zero.

If $b(\alpha, \beta) \neq 0$, then by (X, b) being a complex we know $\dim(\beta) = \dim(\alpha) - 1$.

If $b(\alpha,\beta)-b_q(\alpha,\beta)\neq 0$, then Proposition 4.3.9 implies $q\prec\alpha$ and $\beta\prec w(q)$. Also recall that $b(w(q),q)\in U(R)$, so $b(w(q),q)\neq 0$. And again $\dim(\beta)=\dim(w(q))-1=\dim(q)=\dim(\alpha)-1$, and condition (i) is satisfied.

(ii) For any $\alpha,\gamma\in X$,

$$\begin{aligned}\sum_{\beta\in X_q} b_q(\alpha,\beta)b_q(\beta,\gamma) &= \sum_{\beta\in X_q}(b(\alpha,\beta)b(\beta,\gamma)-b(\alpha,\beta)\frac{b(\beta,q)b(w(q),\gamma)}{b(w(q),q)}\\ &\quad - b(\beta,\gamma)\frac{b(\alpha,q)b(w(q),\beta)}{b(w(q),q)}+\frac{b(\alpha,q)b(w(q),\beta)}{b(w(q),q)}\frac{b(\beta,q)b(w(q),\gamma)}{b(w(q),q)})\\ &= \sum_{\beta\in X_q} b(\alpha,\beta)b(\beta,\gamma)-\frac{b(w(q),\gamma)}{b(w(q),q)}\sum_{\beta\in X_q} b(\alpha,\beta)b(\beta,q)\\ &\quad -\frac{b(\alpha,q)}{b(w(q),q)}\sum_{\beta\in X_q} b(w(q),\beta)b(\beta,\gamma)\\ &\quad +\frac{b(\alpha,q)b(w(q),\gamma)}{b(w(q),q)^2}\sum_{\beta\in X_q} b(w(q),\beta)b(\beta,q) \qquad (*)\end{aligned}$$

We claim $b(w(q),\beta)\cdot b(\beta,q)=0$. Assume to the contrary that $b(w(q),\beta)\cdot b(\beta,q)\neq 0$. Then $\dim(w(q))=\dim(\beta)+1=\dim(q)+2$ which contradicts the partial matching. It follows that the last summand of $(*)$ vanishes.

Recall that for any $\alpha',\gamma'\in X$, $\sum_{\beta\in X} b(\alpha',\beta)b(\beta,\gamma')=0$ by definition of the boundary incidence function b. Since $X_q=X\setminus\{q,w(q)\}$, we obtain $\sum_{\beta\in X_q} b(\alpha',\beta)b(\beta,\gamma')=-b(\alpha',q)b(q,\gamma')-b(\alpha',w(q))b(w(q),\gamma')$.

For the first summand of $(*)$, $\alpha'=\alpha$, $\gamma'=\gamma$.

Then, recall $b(\beta',\beta')=0$ for any $\beta'\in X$.

For the second summand of $(*)$, $\alpha'=\alpha$ and $\gamma'=q$, in particular $-b(\alpha',q)b(q,\gamma')=-b(\alpha,q)b(q,q)$ vanishes.

For the third summand of $(*)$, $\alpha'=w(q)$ and $\gamma'=\gamma$, in particular $-b(\alpha',w(q))b(w(q),\gamma')=-b(w(q),w(q))b(w(q),\gamma)$ vanishes.

Putting everything together, we can rewrite the right hand side of $(*)$ as

$$\begin{aligned}&-b(\alpha,q)b(q,\gamma)-b(\alpha,w(q))b(w(q),\gamma)-\frac{b(w(q),\gamma)}{b(w(q),q)}(-b(\alpha,w(q))b(w(q),q))\\ &-\frac{b(\alpha,q)}{b(w(q),q)}(-b(w(q),q)b(q,\gamma))\\ &=0\end{aligned}$$

Thus (X_q,b_q) is an abstract cell complex.

□

Alternatively, we can see the q-reduced complex as a special kind of Morse complex.

Remark 4.3.11. *The q-reduced complex is also the Morse complex of (X,b) induced by a simple acyclic partial matching $(A',w':Q'\to K')$ defined by $A'=X\setminus\{q,w(q)\}$, $Q'=\{q\}$, and $K'=\{w(q)\}$.*

The observation is that for such a matching every pair of critical cells α, β has at most two (unsaturated) gradient paths (α, β) and (α, q, β). Let b' the boundary incidence function of this Morse complex, then $b'(\alpha, \beta)$ is the sum of the multiplicities of the two paths. But the two multiplicities are precisely $b(\alpha, \beta)$ and $-\dfrac{b(\alpha, q)b(w(q), \beta)}{b(w(q), q)}$ and their sum is by definition $b_q(\alpha, \beta)$.

Meanwhile, a q-reduced matching is an acyclic partial matching on the q-reduced complex.

Lemma 4.3.12. *Given an abstract cell complex (X, b) with acyclic partial matching $(A, w : Q \to K)$. For any $q \in Q$, the q-reduced matching $(A, w : Q_q \to K_q)$ is an acyclic partial matching of the q-reduced complex (X_q, b_q).*

Moreover, any gradient path $\alpha \overset{\rho}{\rightsquigarrow} \beta$ of $(A, w : Q_q \to K_q)$ is at least an unsaturated gradient path of $(A, w : Q \to K)$.

Proof. For any $q' \in Q_q$, by Proposition 4.3.9, $b_q(w(q'), q') = b(w(q'), q') \in U(R)$, so $(A, w : Q_q \to K_q)$ is a partial matching of (X_q, b_q).

To prove acyclicity, we first let $(Q, \lessdot, \lhd)$ be the gradient of $(A, w : Q \to K)$ and $(Q_q, \lessdot_q, \lhd_q)$ be the gradient of $(A, w : Q_q \to K_q)$.

Now, we claim that for any $q', q'' \in Q_q \cup A$, $q' \lessdot_q q''$ implies $q' \lhd q''$. Note that we are using the expanded notation of Notation 4.3.1 for the case of q' or q'' being in A.

To see the claim, note that $q' \lessdot_q q''$ implies that $b_q(w(q''), q') \neq 0$. Hence $b(w(q''), q')$ and $\dfrac{b(w(q''), q)b(w(q), q')}{b(w(q), q)}$ must not both be zero. The former being non-zero means $q' \lessdot q''$ while the later being non-zero means $q' \lessdot q \lessdot q''$. In either case $q' \lhd q''$.

Hence, if we have a cycle for $\lessdot_q$, say $q_1 \lessdot_q q_2 \lessdot_q ... \lessdot_q q_n \lessdot_q q_1$, then we would have a cycle $q_1 \lhd q_2 \lhd ... \lhd q_n \lhd q_1$ for $\lhd$, which we know is impossible as $(A, w : Q \to K)$ is an acyclic matching of (X, b). This gives acyclicity of $\lessdot_q$ and $\lhd_q$.

Finally, given $\rho = (\alpha, q_1, ..., \beta)$ a gradient path of $(A, w : Q_q \to K_q)$, i.e. $\alpha \gtrdot_q q_1 \gtrdot_q ... \gtrdot_q \beta$, then $\alpha \rhd q_1 \rhd ... \rhd \beta$. So ρ is indeed at least an unsaturated gradient path of $(A, w : Q \to K)$. □

Then we consider the corresponding chain complexes of both X and X_q, $C(X)$ and $C(X_q)$, to show that X and X_q have the same homology by introducing two chain maps in the following proposition. The chain maps will be represented as a single map between $C(X)$ and $C(X_q)$, which can be considered two big R-modules as was described in Remark 2.3.

Proposition 4.3.13. *Given an abstract cell complex (X, b) with q-reduced complex (X_q, b_q). We define functions $\phi : C(X) \to C(X_q)$, $\psi : C(X_q) \to C(X)$ between the associated chain complexes by linearly extending $\phi : X \to C(X_q)$ and $\psi : X_q \to C(X)$, which are defined as follows.*

$$\phi(\alpha) := \begin{cases} 0 & \alpha = w(q) \\ -\sum_{\beta \in X_q} \dfrac{b(w(q), \beta)}{b(w(q), q)} \beta & \alpha = q \\ \alpha & \alpha \in X_q \end{cases}$$

$$\psi(\alpha) := \alpha - \frac{b(\alpha, q)}{b(w(q), q)} w(q)$$

Then ϕ and ψ are chain maps. Moreover, $\phi \circ \psi$ and $\psi \circ \phi$ are chain homotopic to the identity maps, i.e. $(\phi\circ\psi)_* = (id_{C(X_q)})_*$ and $(\psi\circ\phi)_* = (id_{C(X)})_*$. Thus $\phi_* = (\psi_*)^{-1}$ gives an isomorphism between the homology groups of (X, b) and (X_q, b_q), that is for each d, $H_d(X; R) \cong H_d(X_q; R)$.

Remark 4.3.14. *It is helpful to know that in each dimension the modules of $C_d(X_q)$ are in fact submodules of $C_d(X)$. In this context, $\phi(q)$ can be alternatively defined as $q - \dfrac{\partial_{C(X)}(w(q))}{b(w(q), q)}$ since $-\sum_{\beta\in X_q} \dfrac{b(w(q),\beta)}{b(w(q),q)}\beta = q - \sum_{\beta\in X} \dfrac{b(w(q),\beta)}{b(w(q),q)}\beta = q - \dfrac{\partial_{C(X)}(w(q))}{b(w(q),q)}$.*

Proof of Proposition 4.3.13. We will abbreviate all the boundary operators as simply ∂ where the context is clear. Recall that $X^q = X \setminus \{q, w(q)\}$.

Since ψ and ϕ are defined by linear extension it suffices to check the condition on the basis, i.e. the cells. For ϕ, we consider the three cases of q, $w(q)$, and cells in X_q, which we denote by α. We also let $\dim(q) = d$.

$\alpha = w(q)$:

$\phi \circ \partial(w(q)) = \phi(\sum_{\beta^{(d)}\in X} b(w(q),\beta)\beta) = \sum_{\beta^{(d)}\in X_q} b(w(q),\beta)\phi(\beta) + b(w(q),q)\phi(q)$

$= \sum_{\beta^{(d)}\in X_q} b(w(q),\beta)\beta + b(w(q),q)(-\sum_{\beta^{(d)}\in X_q} \dfrac{b(w(q),\beta)}{b(w(q),q)}\beta) = 0$

Note that $b(w(q),\beta) \neq 0$ implies that $\dim\beta = d$, and thus it makes sense to only look at β's of this dimension. Then the only dimension d element in $X \setminus X_q$ is q.

$\partial \circ \phi(w(q)) = \partial(0) = 0$

$\alpha = q$:

$\phi \circ \partial_{C(X)}(q) = \phi(\sum_{\gamma^{(d-1)}\in X} b(q,\gamma)\gamma) = \sum_{\gamma^{(d-1)}\in X} b(q,\gamma)\phi(\gamma) = \sum_{\gamma^{(d-1)}\in X_q} b(q,\gamma)\gamma$

Note that as above, $b(w(q),\beta) \neq 0$ implies $\dim\beta = d$. In this case there are no dimension $d-1$ elements in $X \setminus X_q$.

$\partial \circ \phi(q) = \partial_{C(X_q)}(-\sum_{\beta^{(d)}\in X_q} \dfrac{b(w(q),\beta)}{b(w(q),q)}\beta) = -\sum_{\beta^{(d)}\in X_q}\sum_{\gamma^{(d-1)}\in X_q} \dfrac{b(w(q),\beta)b_q(\beta,\gamma)}{b(w(q),q)}\gamma$

$= -\sum_{\gamma^{(d-1)}\in X_q}(\sum_{\beta\in X} \dfrac{b(w(q),\beta)b(\beta,\gamma)}{b(w(q),q)}\gamma - \dfrac{b(w(q),q)b(q,\gamma)}{b(w(q),q)}\gamma) = \sum_{\gamma^{(d-1)}\in X_q} b(q,\gamma)\gamma$

Note that we used the fact that for any $\beta^{(d)}, \gamma^{(d-1)} \in X_q$, $b_q(\beta,\gamma) = b(\beta,\gamma)$. Otherwise by Proposition 4.3.9, $\gamma \prec w(q)$ and $q \prec \beta$, which would violate the dimension requirement between faces. Note also that by the definition of the boundary function $\sum_{\beta\in X} b(w(q),\beta)b(\beta,\gamma) = 0$.

$\alpha \in X_q$:

$\phi \circ \partial(\alpha) = \phi(\sum_{\beta\in X_q} b(\alpha,\beta)\,\beta + b(\alpha,q)q + b(\alpha,w(q))w(q))$

$= \sum_{\beta\in X_q} b(\alpha,\beta)\beta - \sum_{\beta\in X_q} \dfrac{b(\alpha,q)b(w(q),\beta)}{b(w(q),q)}\beta = \sum_{\beta\in X_q} b_q(\alpha,\beta)\beta$

$\partial \circ \phi(\alpha) = \partial_{C(X_q)}(\alpha) = \sum_{\beta\in X_q} b_q(\alpha,\beta)\beta$

Clearly everything matches and so both ϕ and ψ are chain maps. We then look at the two compositions $\phi \circ \psi$ and $\psi \circ \phi$.

For any $\alpha \in X_q$, $\phi \circ \psi(\alpha) = \phi(\alpha - \dfrac{b(\alpha,q)}{b(w(q),q)}w(q)) = \alpha$. So $\phi \circ \psi$ is in fact the identity map already.

While for $\psi \circ \phi$, we first write down the function explicitly for the basis X of $C(X)$. Again we divide the calculation into three cases, $w(q)$, q and every other $\alpha \in X_q$.

$\alpha = w(q)$:

$\psi \circ \phi(w(q)) = \psi(0) = 0$

$\alpha = q$:

$\psi \circ \phi(q) = \psi\Big(- \sum_{\beta \in X_q} \frac{b(w(q),\beta)}{b(w(q),q)}\beta\Big) = - \sum_{\beta \in X_q} \frac{b(w(q),\beta)}{b(w(q),q)}\Big(\beta - \frac{b(\beta,q)}{b(w(q),q)}w(q)\Big)$

$= - \sum_{\beta \in X_q} \frac{b(w(q),\beta)}{b(w(q),q)}\beta + \sum_{\beta \in X_q} \frac{b(w(q),\beta)b(\beta,q)}{b(w(q),q)^2}w(q) = - \sum_{\beta \in X_q} \frac{b(w(q),\beta)}{b(w(q),q)}\beta$

Note that the second summand vanishes since $b(w(q),\beta)$ and $b(\beta,q)$ can never both be non-zero.

$\alpha \in X_q$:

$\psi \circ \phi(\alpha) = \psi(\alpha) = \alpha - \frac{b(\alpha,q)}{b(w(q),q)}w(q)$

In summary, $\psi \circ \phi(\alpha) = \begin{cases} 0 & \alpha = w(q) \\ - \sum_{\beta \in X_q} \frac{b(w(q),\beta)}{b(w(q),q)}\beta & \alpha = q \\ \alpha - \frac{b(\alpha,q)}{b(w(q),q)}w(q) & \alpha \in X_q \end{cases}$

Now we claim the following functions $\{\theta_d : C_d(X) \to C_{d+1}(X)\}$, which we simply call θ, give us the chain homotopy between $\psi \circ \phi$ and id on $C(X)$. In homology theory this would mean $(\psi \circ \phi)_* = id_* = id_{H_*(X)}$ as desired.

Again we only show the function on the basis X of $C(X)$.

$\theta_d(\beta) = \begin{cases} \frac{1}{b(w(q),q)}w(q) & \beta = q \\ 0 & \beta \neq q \end{cases}$

It remains to show that the two functions, $\theta_d \circ \partial + \partial \circ \theta_{d+1}$ and $id_{C(X)} - \psi \circ \phi$, coincide. Again it suffices to compare $(\theta_d \circ \partial + \partial \circ \theta_{d+1})(\alpha)$ and $(id_{C(X)} - \psi \circ \phi)(\alpha)$ for $\alpha = q$, $\alpha = w(q)$, and $\alpha \in X_q$.

Assume $\dim(q) = d'$

$\alpha = w(q)$:

$(\theta_{d'} \circ \partial + \partial \circ \theta_{d'+1})(w(q)) = \theta_{d'}(\partial(w(q)) + \partial_{C(X_q)}(0) = \theta_{d'}(b(w(q),q)q) = w(q)$

$(id - \psi \circ \phi)(w(q)) = w(q) - 0 = w(q)$

$\alpha = q$:

$(\theta_{d'-1} \circ \partial + \partial \circ \theta_{d'})(q) = 0 + \partial(\frac{1}{b(w(q),q)}w(q)) = \frac{\partial(w(q))}{b(w(q),q)}$

$(id - \psi \circ \phi)(q) = q - (- \sum_{\beta \in X_q} \frac{b(w(q),\beta)}{b(w(q),q)}\beta) = q - \phi(q) = \frac{\partial(w(q))}{b(w(q),q)}$

Note that last equation follows from Remark 4.3.14.

$\alpha \in X_q$:

$(\theta_d \circ \partial + \partial \circ \theta_{d+1})(\alpha) = \theta_d(\sum_{\beta \in X} b(\alpha,\beta)\beta) = \frac{b(\alpha,q)}{b(w(q),q)}w(q) = (id - \psi \circ \phi)(\alpha)$

So we indeed have $\theta_d \circ \partial + \partial \circ \theta_{d+1} = id_{C(X)} - \psi \circ \phi$.

□

While Proposition 4.3.13 is sufficient for us to reduce X all the way to A while preserving the homology, we can't tell yet if A has the same abstract cell complex structure as the Morse complex of X does. To do so, we introduce a lemma to show that X and its q-reduced X_q have the same Morse Complex.

Lemma 4.3.15. *Let (X, b) be an abstract cell complex with acyclic partial matching $(A, w : Q \to K)$. For any $q \in Q$, the acyclic partial matching $(A, w : Q_q \to K_q)$ on the q-reduced complex (X_q, b_q) induces the same Morse Complex as $(A, w : Q \to K)$, i.e. (A, b') and (A, b'_q) are the same complex, or $b' = b'_q$.*

Proof. We compare b' and b'_q for each pair of $\alpha, \beta \in A$. For simplicity we write $\nabla_{\rhd_q}(\alpha, \beta)$ as ∇_q and $\nabla_{\rhd}(\alpha, \beta)$ as ∇. The goal is to show

$$b'(\alpha, \beta) = \sum_{\rho \in \nabla} \mu_b(\rho) = \sum_{\rho \in \nabla_q} \mu_{b_q}(\rho) = b'_q(\alpha, \beta)$$

.

We already know from Lemma 4.3.12 that $\nabla_q \subset \nabla$. By considering any Q-sequences ρ as ordered sets, we define $\nabla^{\notin} := \{\rho \in \nabla, q \notin \rho\}$ and $\nabla^{\in} :=$
$\{\rho \in \nabla, q \in \rho\}$ which is a partition of ∇. Clearly, $\nabla_q \subset \nabla^{\notin}$.

From Proposition 4.3.4, we know every $\rho \in \nabla^{\notin} - \nabla_q$ has $\mu_{b_q}(\rho) = 0$, and so we will instead prove that $\sum_{\rho \in \nabla} \mu_b(\rho) = \sum_{\rho \in \nabla^{\notin}} \mu_{b_q}(\rho)$, or rather,

$$\sum_{\rho \in \nabla^{\in}} \mu_b(\rho) = \sum_{\rho \in \nabla^{\notin}} \big(\mu_{b_q}(\rho) - \mu_b(\rho)\big)$$

.

To do so we will first remove another class of Q-sequences from $\nabla^{\notin}$.

Given $\rho = (q_0, ..., q_{l+1}) \in \nabla^{\notin}$. Recall the definition of $\nabla^{\notin}$, $q_0 \rhd q_1 \rhd ... \rhd q_{l+1}$ and none of the q_i's are q. Now we look at those ρ's where for any $0 \le m \le l(\rho)$, $q_m \rhd q \rhd q_{m+1}$ does NOT hold. In such a case, we know for sure that $q \prec w(q_m)$ and $q_{m+1} \prec w(q)$ must not both be true. By Proposition 4.3.9, this implies $b(w(q_m), q_{m+1}) = b_q(w(q_m), q_{m+1})$ for any $0 \le m \le l(\rho)$.

Recall that $\mu_{b_q}(\rho) = \dfrac{\prod_{i=0}^{l} b_q(w(q_i), q_{i+1})}{\prod_{i=1}^{l} -b_q(w(q_i), q_i)}$ and $\mu_b(\rho) = \dfrac{\prod_{i=0}^{l} b(w(q_i), q_{i+1})}{\prod_{i=1}^{l} -b(w(q_i), q_i)}$ by definition. So for the class of Q-sequences ρ's we just stated, $\mu_{b_q}(\rho)$ and $\mu_b(\rho)$ have the same numerator. Also by Proposition 4.3.9 we see that they have the same denominator. Hence this class of Q-sequences has the property that $\mu_{b_q}(\rho) - \mu_b(\rho) = 0$ and can be ignored for our purpose.

Let $\nabla^{\#}$ be the set of the rest of the Q-sequences ρ's in $\nabla^{\notin}$, i.e. Q-sequences where $q_m \rhd q \rhd q_{m+1}$ holds for at least one index m, and we further reduced our claim to proving

$$\sum_{\rho \in \nabla^{\in}} \mu_b(\rho) = \sum_{\rho \in \nabla^{\#}} \big(\mu_{b_q}(\rho) - \mu_b(\rho)\big)$$

For $\rho \in \nabla^{\#}$, we first note that the index m for which we have $q_m \rhd q \rhd q_{m+1}$ must be unique. If we instead had some index m' with $m < m'$ satisfying this property, then $q \rhd q_{m+1} \rhd ... \rhd q_{m'} \rhd q$ contradicting acyclicity of $(A, w : Q \to K)$. Hence we have $\alpha = q_0 \rhd q_1 ... \rhd q_m \rhd q \rhd q_{m+1} \rhd ... \rhd q_{l+1} \rhd \beta$ where all q_i's are distinct, so $\rho^+ := (\alpha, ..., q_m, q, q_{m+1}, ..., \beta) \in \nabla^{\in}$.

Hence for every i other then $i = m$, $q_m \rhd q \rhd q_{m+1}$ does NOT hold, and so we can again

apply Proposition 4.3.9 to get $b(w(q_i), q_{i+1}) = b_q(w(q_i), q_{i+1})$. Therefore we have the following.

$$
\begin{aligned}
&\mu_{b_q}(\rho) \\
&= \frac{\prod_{i=0}^{m-1} b_q(w(q_i), q_{i+1})}{\prod_{i=1}^{m} -b_q(w(q_i), q_i)} \Big(b_q(w(q_m), q_{m+1})\Big) \frac{\prod_{i=m+1}^{l(\rho)} b_q(w(q_i), q_{i+1})}{\prod_{i=m+1}^{l(\rho)} -b_q(w(q_i), q_i)} \\
&= \frac{\prod_{i=0}^{m-1} b(w(q_i), q_{i+1})}{\prod_{i=1}^{m} -b(w(q_i), q_i)} \Big(b(w(q_m), q_{m+1}) - \frac{b(w(q_m), q)b(w(q), q_{m+1})}{b(w(q), q)}\Big) \frac{\prod_{i=m+1}^{l(\rho)} b(w(q_i), q_{i+1})}{\prod_{i=m+1}^{l(\rho)} -b(w(q_i), q_i)} \\
&= \frac{\prod_{i=0}^{l(\rho)} b(w(q_i), q_{i+1})}{\prod_{i=1}^{l(\rho)} -b(w(q_i), q_i)} + \frac{\prod_{i=0}^{m-1} b(w(q_i), q_{i+1})}{\prod_{i=1}^{m} -b(w(q_i), q_i)} \Big(\frac{b(w(q_m), q)b(w(q), q_{m+1})}{-b(w(q), q)}\Big) \frac{\prod_{i=m+1}^{l(\rho)} b(w(q_i), q_{i+1})}{\prod_{i=m+1}^{l(\rho)} -b(w(q_i), q_i)} \\
&= \mu_b(\rho) + \mu_b(\rho^+)
\end{aligned}
$$

Or $\mu_{b_q}(\rho) - \mu_b(\rho) = \mu_b(\rho^+)$.

The proof is finished if there exists bijective function $+_q : \nabla^\# \to \nabla^\in, \rho \mapsto \rho^+$ where ρ^+ is constructed from ρ as above.

Firstly, since $q \in \rho^+$, clearly $\rho^+ \in \nabla - \nabla_q = \nabla^\in$. So $+_q$ indeed maps elements in $\nabla^\#$ to $\nabla^\in$. The well-definedness of $+_q$ follows from the uniqueness of the choice of m.

For injectivity, note that since $+_q(\rho)$ is ρ with q added into the sequence, ρ is the merely $+_q(\rho)$ with q removed. So $+_q(\rho) = +_q(\rho')$ implies ρ and ρ' are results of the same sequence with the same element (q) removed, and thus ρ and ρ' must be the same sequence.

Lastly, for any $\rho' \in \nabla^\in$, $q \in \rho'$. So we can consider the subsequence $\rho := \rho' - \{q\}$. Obviously ρ as a subsequence of an extended gradient path is also an extended gradient path, so $\rho \in \nabla$. It is also clear that there is a unique position in ρ where we can reinstate q to get ρ' back, so $\rho \in \nabla^\#$ and $+_q(\rho) = \rho'$.

Hence $\sum_{\rho \in \nabla^\#} \mu_{b_q}(\rho) - \mu_b(\rho) = \sum_{\rho \in \nabla^\#} \mu_b(+_q(\rho)) = \sum_{\rho' \in \nabla^\in} \mu_b(\rho')$ as desired.

□

Finally, we have all the tools to prove the main theorem of this section.

Proof of Theorem 4.3.7. Let (X, b) be an abstract cell complex with acyclic partial matching $(A, w : Q \to K)$). Let (A, b') be the Morse Complex.

Let $Q = \{q_1, q_2, ..., q_n\}$. Define inductively abstract cell complexes (X_0, b_0), (X_1, b_1), ... , (X_{n-1}, b_{n-1}) so that $(X_0, b_0) = (X, b)$ and for every i, $(X_{i+1}, b_{i+1}) := ((X_i)_{q_{i+1}}, (b_i)_{q_{i+1}})$ is

the q_{i+1}-reduced complex of (X_i, b_i). Also, inductively construct the q_{i+1}-reduced matchings $(A, w : Q_{i+1} \to K_{i+1})$ from $(A, w : Q_i \to K_i)$.

Note in particular that $X_n = A$ and the Morse Complex (A, b'_n) of (X_n, b_n) is merely itself. By applying Lemma 4.3.15 inductively to all X_i, all of the (X_i, b_i) share the same Morse Complex.

It remains then to apply Proposition 4.3.13 inductively to each (X_i, b_i) and see that all (X_i, b_i)'s have the same homology. In particular (X, b) and $(X_n, b_n) = (A, b')$ have the same homology groups, i.e. $H_*(X, R) = H_*(A, R)$, and we are done.

□

5 Persistent Homology

Persistent homology, introduced by Gunnar Carlson and others, has since then been highly successful at presenting topological information of data sets and is the flagship of applying algebraic topology in data analysis, see [2] for a survey.

Applying tools like the Vietoris-Rips Complex and the Cech Complex, we can obtain from data in the form of some discrete sets in $\mathbb{R}^n$, a "filtration" of abstract cell complexes. The general idea of persistent homology revolves around processing this "filtration" and track how basis elements of the homology groups "persist" over the series. The end result is usually represented in a graphical form which is called a "barcode".

We will provide definitions based on abstract cell complexes. However, in general the ideas presented in this section can be easily replicated in various contexts like geometric simplicial complexes and abstract simplicial complexes.

5.1 Filtration

In general set theory a filtration refers to a family of sets indexed by a totally ordered set, so that the order also preserves the subset relation. In particular this means the family of sets is totally ordered with respect to the subset relation. In the context of abstract cell complexes, we further restrict the sets to be abstract cell complexes and the subset relation to be the subcomplex relation.

In practice, the family of abstract cell complexes we work with is always finite. Thus in this article we use the natural numbers as our index.

Definition 5.1.1. Filtration of Abstract Cell Complexes

A filtration of abstract cell complexes*, denoted by* $\mathcal{F} = \{X^p, b\}$*, is a sequence of complexes* $.... \subset X^{p-1} \subset X^p \subset X^{p+1} \subset ...$ *such that each* (X^p, b) *is a subcomplex of* (X^{p+1}, b)*.*

Each X^p *will be called the* p-th frame *of* $\mathcal{F}$*. The last frame of the filtration, for which every frame is subcomplex of, will be referred to as* X*.*

The subcomplex relation is particularly important as we can construct chain maps between the associated chain complexes based on it. Moreover, we can see from Definition 3.1.6 that the subcomplex relation allows every frame X^p to share the same boundary incidence function b in the sense that the boundary incidence function of X^p is $b\,|_{X^p}$.

Definition 5.1.2. Inclusion Map

Let (X^p, b) *and* (X^{p+1}, b) *be abstract cell complexes such that* X^p *is a subcomplex of* X^{p+1}*. For each* p *and each dimension* d*, the inclusion map* $i_d^p : C_d(X^p) \to C_d(X^{p+1})$ *is defined as the linear extension of the inclusion map* $i_d^p : X_d^p \to X_d^{p+1}, \beta \mapsto \beta$*. The collection of maps* $\{i_d^p\}$ *is denoted by* $i^p : C(X^p) \to C(X^{p+1})$*.*

Moreover, we let $i^{p\to p'} : C(X^p) \to C(X^{p'})$ *denote the collection of maps* $\{i_d^{p\to p'} : C_d(X^p) \to C_d(X^{p\prime})\}$ *where* $i_d^{p\to p'} = i_d^{p'-1} \circ ... \circ i_d^p$*.*

Note that we can alternatively define $i_d^{p\to p'}$ by considering X^p as a subcomplex of $X^{p\prime}$.

We will then prove that this inclusion map is indeed a chain map.

Proposition 5.1.3. *The inclusion map i^p is a chain map from $C(X^p)$ to $C(X^{p+1})$, and thus so is $i^{p\to p'}$ from $C(X^p)$ to $C(X^{p\prime})$.*

Proof. We look at each dimension d. Then for any $\alpha \in X^{(d)p}$

$i^p_{d-1} \circ \partial^p_d(\alpha) = i^p_{d-1}(\sum_{\beta\in X^p} b(\alpha,\beta)\beta) = \sum_{\beta\in X^p} b(\alpha,\beta)\beta.$

Since X^p is a subcomplex of X^{p+1}, by Definition 3.1.6, $\alpha \in X^p$ and $\beta \notin X^p$ implies that $b(\alpha,\beta) = 0$.

So $i^p_{d-1} \circ \partial^p_d(\alpha) = \sum_{\beta\in X^{p+1}} b(\alpha,\beta)\beta = \partial^{p+1}_d(\alpha) = \partial^{p+1}_d \circ i^p_d(\alpha)$ as desired.

For $i^{p\to p'}$, just recall that it is the composition of $i^p, i^{p+1}, ..., i^{p'-1}$.

□

Thus each series i^p of i^p_d are effectively a series of embeddings, which embed the module $C_d(X^p)$ to the module $C_d(X^{p+1})$. From the perspective of considering an entire chain complex as one module as was done in Remark 2.3, i^p embeds $C(X^p)$ into $C(X^{p+1})$.

Remark 5.1.4. *While not covered in this paper, the theory still works when we use real numbers as our index set.*

Similar to the treatment of ∂_d, we often omit the index and write the inclusion functions i_p merely as i. So the chain map condition is reduced to $i\partial = \partial i$.

With the above established we can define the persistent homology groups.

Definition 5.1.5. Persistent Homology Group

The dth (p,p')-persistent homology group of $\mathcal{F}$ is defined by

$$H^{p\to p'}_d(\mathcal{F}) = \frac{i^{p\to p'}_d(Z_d(X^p))}{i^{p\to p'}_d(Z_d(X^p)) \cap B_d(X^{p'})}$$

where $Z_d(X^p) = \ker(\partial^p_d)$, $B_d(X^{p'}) = \mathrm{Im}(\partial^{p'}_d)$

Remark 5.1.6. *An element of $H^{p\to p'}_d(\mathcal{F})$ is usually denoted as $\bar{x}$, the equivalence class of $x \in i^{p\to p'}_d(Z_d(X^p)) \subset Z_d(X^{p'})$. Note that $i^{p\to p'}_d(Z_d(X^p)) \subset Z_d(X^{p'})$ holds since $i^{p\to p'}$ is a composition of embeddings.*

The following proposition shows that we can also define $H^{p\to p'}_d(\mathcal{F})$ as $\mathrm{Im}((i^{p\to p'}_d)_*)$ where $(i^{p\to p'}_d)_*$ is the homomorphism $H_d(X^p) \to H_d(X^{p'})$ induced by $i^{p\to p'}_d$ as a chain map. $H_d(X^p)$ and $H_d(X^{p'})$ are the R coefficient homology groups stated in Remark 3.2.2

Proposition 5.1.7. *$H^{p\to p'}_d(\mathcal{F}) \cong (i^{p\to p'}_d)_*(H_d(X^p))$, where $(i^{p\to p'}_d)_*$ is the map $H_d(X^p) \to H_d(X^{p'})$ induced by $i^{p\to p'}_d$.*

Proof. Recall $H^{p\to p'}_d(\mathcal{F})$ is a quotient group of $i^{p\to p'}_d(Z_d(X^p))$. Thus for any $\bar{x} \in H^{p\to p'}_d(\mathcal{F})$, any of its representatives x would be in $i^{p\to p'}_d(Z_d(X^p))$ and thus represents an element in $(i^{p\to p'}_d)_*(H_d(X^p))$. In fact given any $x \in Z_d(X^{p'})$,

$\bar{x} \in H^{p\to p'}_d(\mathcal{F}) \iff x \in i^{p\to p'}_d(Z_d(X^p)) \iff \bar{x} \in (i^{p\to p'}_d)_*(H_d(X^p))$ $\quad (*)$

Thus we can construct a function from $H^{p\to p'}_d(\mathcal{F})$ to $(i^{p\to p'}_d)_*(H_d(X^p))$ in the following way. Given any $\bar{x} \in H^{p\to p'}_d(\mathcal{F})$, choose a representative x of $\bar{x}$ which must be in $i^{p\to p'}_d(Z_d(X^p))$, hence

x must also represent an element $\bar{x}$ in the subgroup $(i_d^{p\to p'})_*(H_d(X^p))$ of the homology group $H_d(X^{p'})$.

Clearly, the function preserves addition, hence if it maps zero to zero it must be well-defined and injective. If $\bar{x} = 0 \in H_d^{p\to p'}(\mathcal{F})$, then $x \in i_d^{p\to p'}(Z_d(X^p)) \cap B_d(X^{p'})$ and so $\bar{x} = 0$ in $H_d(X^{p'})$ as well as its subset $\mathrm{Im}((i_d^{p\to p'})_*)$.

Lastly, note that the implications in $(*)$ work both ways, so the function must also be surjective, and is in fact an isomorphism from $H_d^{p\to p'}(\mathcal{F})$ to $(i_d^{p\to p'})_*(H_d(X^p))$.

□

The mathematical idea of a persistent homology group $H_d^{p\to p'}(\mathcal{F})$ is to track how many d-dimension "holes" survive from the p-th frame to the p'-frame.

In practice, we usually choose some field to be the coefficient ring of the abstract cell complexes, which makes the persistent homology groups vectors spaces. This allows the information of all the dth persistent homology groups, usually denoted together as H_d, to be represented in a barcode diagram for every d. Here, we represent the dimension of $H_d^{p\to p'}(\mathcal{F})$ by the number of bars passing both p and p', see the example below.

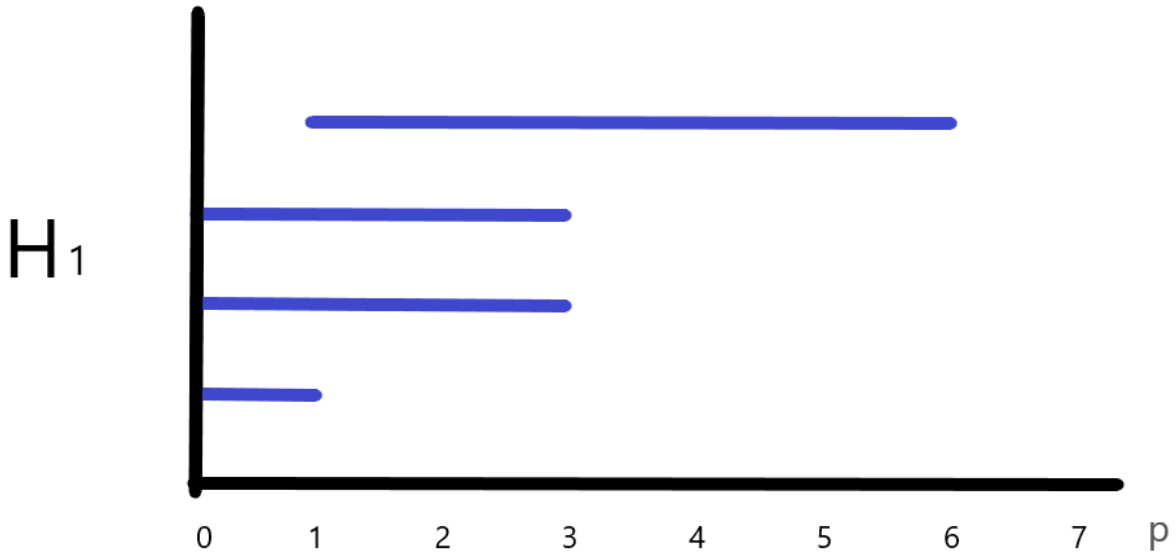

Our particular example is a the H_1 barcode of a filtration as follows.

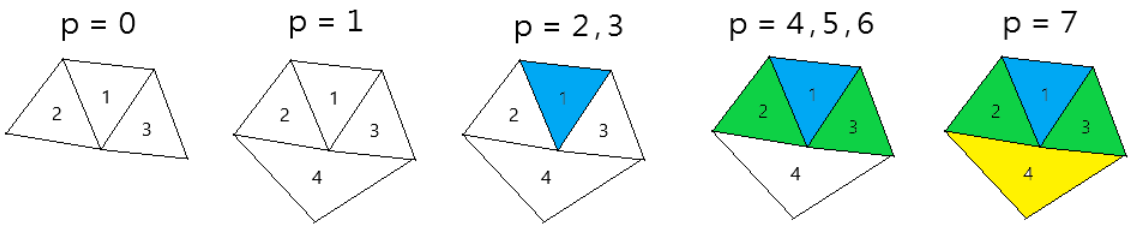

Each bars tracks one of the triangles in the complexes, starting at the value p where the triangle appears and ends at the value p where the triangle is filled.

To see how this relates to the persistent homology groups, note that $H_1^{1\to 3}(\mathcal{F})$ has dimension 3, since if we compare X^1 and X^3, triangles 2, 3, and 4 "persist" through the inclusion. But $H_1^{1\to 4}(\mathcal{F})$ only has dimension 1, since triangles 2 and 3 are already filled at X^4.

On the barcode diagram this is shown by having three bars "persist" from $p = 1$ to $p = 3$, but two of the bars, which would represent triangles 2/3, do not persist to $p = 4$.

The barcode diagram and thus the persistent homology groups are very often the end result wanted in TDA. As we will see in the next section, we can actually apply ideas of discrete Morse theory to the theory of persistent homology groups to get some interesting results which also have implications on real life applications of TDA.

6 Morse Filtration

In this section we will show how discrete Morse theory can be applied to the theory of persistent homology. Specifically we want to apply the idea of critical cells capturing the key information of abstract cell complexes, but on a filtration of complexes rather than a single one. The end result is we can construct a smaller Morse filtration in Definition 6.2.2 whose persistent homology groups are isomorphic to those of the original filtration.

6.1 Filtered Chain Map

In Section 5 we only worked with one filtration of abstract cell complexes at a time. Hence, in order to work with different filtrations, which will be crucial in the development of the Morse complex, we need additional tools. The central tool for our purpose will be filtered chain maps, a kind of structure-preserving map between filtrations of abstract cell complexes.

Definition 6.1.1. Filtered Chain Map

Let $\mathcal{F} = \{X^p, b\}$ and $\mathcal{F}' = \{X'^p, b\}$ be filtrations of abstract cell complexes. A filtered chain map $\Phi : \mathcal{F} \to \mathcal{F}'$ *is a sequence $\{\phi^p = \{\phi_d^p\} : C(X^p) \to C(X'^p)\}$ of chain maps so that for each p and d the following diagram commutes*

$$\begin{array}{ccc} C(X_d^p) & \xrightarrow{i_d^p} & C(X^{p+1}) \\ \downarrow{\scriptstyle \phi_d^p} & & \downarrow{\scriptstyle \phi_d^{p+1}} \\ C(X_d'^p) & \xrightarrow{i_d^p} & C(X_d'^{p+1}) \end{array}$$

Remark 6.1.2. *Let $\Phi : \mathcal{F} \to \mathcal{F}'$ be a filtered chain map. By considering $i^{p\to p'} = i^{p'-1} \circ ... \circ i^p$ we see that for each p, p' and d, the following diagram also commutes.*

$$\begin{array}{ccc} C(X_d^p) & \xrightarrow{i_d^{p\to p'}} & C(X_d^{p'}) \\ \downarrow{\scriptstyle \phi_d^p} & & \downarrow{\scriptstyle \phi_d^{p'}} \\ C(X_d'^p) & \xrightarrow{i_d^{p\to p'}} & C(X_d'^{p'}) \end{array}$$

Analogous to the chain maps in the theory of homology, a filtered chain map also induces homomorphisms between the corresponding persistent homology groups.

Proposition 6.1.3. *Given filtered chain map $\Phi : \mathcal{F} \to \mathcal{F}'$, Φ induces a family of homomorphisms $\Phi_* := \{(\phi_d^{p\to p'})_* : H_d^{p\to p'}(\mathcal{F}) \to H_d^{p\to p'}(\mathcal{F}')\}$ between the persistent homological groups.*

Proof. Let $\mathcal{F} = \{X^p, b\}$, $\mathcal{F}' = \{X'^p, b'\}$.

Given $\bar{x} \in H_d^{p\to p'}(\mathcal{F})$, let $x = i_d^{p\to p'}(z)$ for some $z \in Z_d(X^p)$. Consider $\phi_d^p(z)$. Since ϕ^p is a chain map, $\partial_d \phi_d^p(z) = \phi_d^p \partial_d(z) = \phi_d^p(0) = 0$. Hence $\phi_d^p(z) \in Z_d(X'^p)$ and we can define $(\phi_d^{p\to p'})_*(\bar{x}) := \overline{i_d^{p\to p'} \phi_d^p(z)}$. By Remark 6.1.2, $i_d^{p\to p'} \phi_d^p(z) = \phi_d^{p'} i_d^{p\to p'}(z) = \phi_d^{p'}(x)$. So $(\phi_d^{p\to p'})_*(\bar{x})$ is also $\overline{\phi_d^{p'}(x)}$.

It is trivial to check that addition is preserved, so it remains to prove $(\phi_d^{p\to p'})_*$ maps zero elements in $H_d^{p\to p'}(\mathcal{F})$ to zero elements in $H_d^{p\to p'}(\mathcal{F}')$ which also implies that $(\phi_d^{p\to p'})_*$ is a well-defined. Since for $\bar{x} = \bar{y}$, we will have $\overline{x - y} = 0$ and so $(\phi_d^{p\to p'})_*(\bar{x}) = (\phi_d^{p\to p'})_*(\bar{y} + \overline{x - y}) = (\phi_d^{p\to p'})_*(\bar{y}) + 0 = (\phi_d^{p\to p'})_*(\bar{y})$.

Let $\bar{x} = 0 \in H_d^{p\to p'}(\mathcal{F})$. By Definition 5.1.5, $x \in i_d^{p\to p'}(Z_d(X^p)) \cap B_d(X^{p'})$.

Since $x \in B_d(X^{p'})$, $x = \partial_{d+1}(z')$ for some $z' \in C_{d+1}(X^{p'})$. Again since $\phi^{p'}$ is a chain map, $\phi_d^{p'}(x) = \phi_d^{p'}\partial_{d+1}(z') = \partial_{d+1}\phi_d^{p'}(z') \in B_d(X'^{p'})$. So $(\phi_d^{p\to p'})_*(\bar{x}) = \overline{\phi_d^{p'}(x)} = \overline{\partial_{d+1}\phi_d^{p'}(z')} = 0 \in H_d^{p\to p'}(\mathcal{F}')$.

□

The idea of chain homotopy in the theory of chain complexes and homologies also translates well into our context and with the same important property of showing two different filtered chain maps induces the same homomorphisms.

Definition 6.1.4. Filtered Chain Homotopy

Let $\Phi, \Phi' : \mathcal{F} \to \mathcal{F}'$ where $\Phi = \{\phi^p\}$ and $\Phi' = \{\phi'^p\}$ be filtered chain maps. A filtered chain homotopy *between Φ and Φ' is a collection of chain homotopies $\Theta = \{\theta^p\}$ where each θ^p is a chain homotopy between ϕ^p and ϕ'^p. If such Φ exists, we say Φ and Φ' are filtered chain homotopic.*

Proposition 6.1.5. *Let $\Phi, \Phi' : \mathcal{F} \to \mathcal{F}'$ be filtered chain maps which induce the families of homomorphism Φ_*, Φ'_*. If Φ and Φ' are filtered chain homotopic. Then $\Phi_* = \Phi'_*$*

Proof. Let $\mathcal{F} = \{X^p, b\}$, $\mathcal{F}' = \{X'^p, b'\}$, $\Phi = \{\phi_d^p : C_d(X^p) \to C_d(X'^p)\}$, and $\Phi' = \{\phi_d'^p : C_d(X^p) \to C_d(X'^p)\}$. The families of induced homomorphisms would then be $\Phi_* = \{(\phi_d^{p\to p'})_*\}$, $\Phi'_* = \{(\phi_d'^{p\to p'})_*\}$. We want to show $(\phi_d^{p\to p'})_* = (\phi_d'^{p\to p'})_*$ for every p and d.

Given $x \in Z_d(X^{p'})$. From the proof of Proposition 6.1.3, $(\phi_d^{p\to p'})_*(\bar{x}) = \overline{i_d^{p\to p'}\phi_d^p(z)}$ where $z \in Z_d(X^p)$ and $i_d^{p\to p'}(z) = x$. By Proposition 5.1.7, we can alternatively define $H_d^{p\to p'}(\mathcal{F}')$ as $(i_d^{p\to p'})_*(H_d(X'^p))$, in which case $\overline{i_d^{p\to p'}\phi_d^p(z)}$ becomes $(i_d^{p\to p'})_*(\overline{\phi_d^p(z)})$. Note also that $\overline{\phi_d^p(z)} = \phi_{d*}^p(\bar{z})$.

Putting everything together, $(\phi_d^{p\to p'})_*(\bar{x}) = (i_d^{p\to p'})_*\phi_{d*}^p(\bar{z})$. Similarly, $(\phi_d'^{p\to p'})_*(\bar{x}) = (i_d^{p\to p'})_*\phi_{d*}'^p(\bar{z})$.

Let $\Theta = \{\theta^p\}$ be a filtered chain homotopy between Φ and Φ'. In particular θ^p is a chain homotopy between the chain maps ϕ^p and ϕ'^p. But this means ϕ^p and ϕ'^p induces the same homomorphisms, in particular $\phi_{d*}^p = \phi_{d*}'^p$. So $(\phi_d^{p\to p'})_* = (\phi_d'^{p\to p'})_*$ as desired.

□

We end this subsection by introducing a useful lemma for our next subsection.

Lemma 6.1.6. *Let $\Phi : \mathcal{F} \to \mathcal{F}'$ $\{\phi^p : C(X^p) \to C(X'^p)\}$, $\Psi : \mathcal{F}' \to \mathcal{F}$ $\{\psi^p : C(X'^p) \to C_d(X^p)\}$ be filtered chain maps.*

If each pair of chain maps ϕ^p and ψ^p has the property that $\psi^p \circ \phi^p$ and $\phi^p \circ \psi^p$ are chain homotopic to the identity chain map, then $\mathcal{F}$ and $\mathcal{F}'$ have isomorphic persistent homology groups.

Proof. Since each $\psi^p \circ \phi^p$ is chain homotopic to the identity, there is a chain homotopy θ_p between $\psi^p \circ \phi^p$ and $id_{C(X^p)}$ for every p. Take this collection of chain homotopies as our filtered chain homotopy and we conclude that $\Psi \circ \Phi$ is filtered chain homotopic to the identity filtered chain map, so by Proposition 6.1.5, $(\Psi \circ \Phi)_* = id_*$.

Similarly $(\Phi \circ \Psi)_* = id_*$. So together we conclude that $\Phi_* = (\Psi_*)^{-1}$ is invertible and thus $\mathcal{F}$ and $\mathcal{F}'$ have isomorphic persistent homology groups.

□

6.2 Filtered Morse Complex

The construction of Morse filtrations is basically the same as the construction of Morse complexes as discussed in Section 4. We construct acyclic partial matchings for each abstract cell complex with some compatibility condition to be defined.

Then each acyclic partial matching will give us a Morse complex, and it can be shown that the compatibility condition implies that these Morse complexes form a filtration of abstract cell complexes with isomorphic persistent homology groups.

We begin by defining those acyclic partial matchings which we call a filtered acyclic partial matchings.

Definition 6.2.1. Filtered Acyclic Partial Matching

Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes. A filtered acyclic partial matching *of $\mathcal{F}$ assigns to each frame X^p an acyclic partial matching $(A^p, w^p : Q^p \to K^p)$. Moreover, for each p that $A^p \subset A^{p+1}$, $K^p \subset K^{p+1}$, $Q^p \subset Q^{p+1}$, and $w^p = w^{p+1}\,|_{Q^p}$. Since the function w^p coincide we will omit the index and use $w := w^p$.*

And the definition of a Morse filtration follows very naturally.

Definition 6.2.2. Morse Filtration

Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matchings $\{(A^p, w : Q^p \to K^p)\}$. The Morse filtration *of $\mathcal{F}$ induced by $\{(A^p, w : Q^p \to K^p)\}$, denoted by $\mathcal{M} = \{A^p, b'\}$, is the sequence of complexes $... \subset A^{p-1} \subset A^p \subset A^{p+1} \subset ...$ where each (A^p, b') is the Morse complex associated to the acyclic partial matching $(A^p, w : Q^p \to K^p)$ of (X^p, b).*

We then establish that a Morse filtration is indeed a filtration of abstract cell complexes, which also justifies using b' as the sole boundary incidence function for every frame A^p of $\mathcal{M}$. We first establish the following lemma.

Lemma 6.2.3. *Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$. Let $\rho = \{\alpha = q_0, \beta_1, ..., \beta_{l+1} = \beta\}$ be a $Q^{p'}$-sequence and let $p < p'$. If $\mu_b(\rho) \neq 0$, and $\alpha \in X^p$, then ρ is also a Q^p-sequence.*

Proof. By Proposition 4.3.4, ρ is a saturated gradient path in $X^{q'}$, i.e. $\alpha \succcurlyeq q_1 \succcurlyeq ... \succcurlyeq q_l \succcurlyeq \beta$. In other words, $\alpha \succ q_1$, $w(q_1) \succ q_2$,...,$w(q_l) \succ \beta$.

Since X^p is a subcomplex of $X^{p'}$ and $\alpha \in X^p$, by Definiton 3.1.6, $q_1 \prec \alpha$ implies $q_1 \in X^p$. Meanwhile by the way we defined w in Definition 6.2.1, $q_1 \in X^p$ implies $w(q_1) \in X^p$. Repeating the above procedure we conclude that β and all the q_i's are in X^p, so ρ is also a Q^p-sequence. □

Proposition 6.2.4. *Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$. The associated Morse filtration $\mathcal{M} = \{A^p, b'\}$ is indeed a filtration of abstract cell complexes.*

Proof. We want to show that each complex A^p is indeed a subcomplex of $A^{p'}$ for all $p' > p$. Recall from Definition 3.1.6, that means $A^p \subset A^{p'}$, $b'_p = b'_{p'}$, and given any $\alpha \in A^p$ and $\beta \in A^{p'}$, $\beta \prec \alpha$ implies $\beta \in A^p$.

Now we fix p. Then Definition 6.2.1 implies that $A^p \subset A^{p'}$.

We first let $(Q^p, \lessdot_p, \lhd_p)$ and $(Q^{p'}, \lessdot_{p'}, \lhd_{p'})$ be the gradients induced by their corresponding partial matchings. By Definition 4.2.2 we see that $\lessdot_p$ and $\lessdot_{p'}$ depend exclusively on b, which is shared among all the frames. Hence when restricted to elements in Q^p the two relations have identical behavior. For convenience we denote both of them as $\lessdot$.

We then will show that the boundary incidence functions of each complex coincide. Before that is established, we temporarily denote each Morse complex in the Morse filtration by (A^p, b'_p) with the goal of showing $b'_{p'} \mid_{A^p} = b'_p$. This will then justify calling all b'_p's merely b'. We further expand the meaning of $\lessdot$ as we did in Notation 4.3.1, which also depends exclusively on b.

Let $\alpha, \beta \in A^p$. By Remark 4.3.6, we can define $b'_p(\alpha, \beta) = \sum_{\rho \in Q^p(\alpha,\beta)} \mu_b(\rho)$, where $Q^p(\alpha, \beta)$ is the set of all Q^p-sequences (Definition 4.3.2) from α to β in X^p with respect to the acyclic partial matching $(A^p, w : Q^p \to K^p)$. We want to show $b'_p = b'_{p'}$, that is, we want to show the following equations holds.

$$\sum_{\rho \in Q^p(\alpha,\beta)} \mu_b(\rho) = \sum_{\rho \in Q^{p'}(\alpha,\beta)} \mu_b(\rho) \qquad (*)$$

To see that, we first note that the multiplicity function μ_b (Definition 4.3.2) used on both sides of $(*)$ are in fact the same, since both are derived from b which is simultaneously the boundary incidence function of all X^p. Furthermore, $Q^p(\alpha, \beta) \subset Q^{p'}(\alpha, \beta)$ by Definition 4.3.2 and $Q^p \subset Q^{p'}$.

Now, given $\rho \in Q^{p'}(\alpha, \beta)$. Let $\rho = (\alpha = q_0, ..., q_{l+1} = \beta)$. If $\mu_b(\rho) = 0$, it is a zero summand and for our purpose ρ can be completely ignored. Assume then that $\mu_b(\rho) \neq 0$. Since $\alpha \in X^p$ Lemma 6.2.3 shows $\rho \in Q^p(\alpha, \beta)$. From this we conclude equation $(*)$ holds, and $b'_p = b'_{p'}$.

Lastly we check that given $\alpha \in A^p$, for every $\beta \in A^{p'}$, $\beta \prec \alpha$ in $A^{p'}$ implies $\beta \in A^p$. Now $\beta \prec \alpha$ in $A^{p'}$ implies $b'(\alpha, \beta) \neq 0$ as the sum of some which means there must be at least one $\rho \in Q^{p'}(\alpha, \beta)$ such that $\mu_b(\rho) \neq 0$, as $b'(\alpha, \beta) \neq 0$ is the sum of these $\mu_b(\rho)$'s. So we can apply Lemma 6.2.3 again to show ρ is a Q^p-sequence, which by extension means $\beta \in A^p$, as desired.

□

It is particularly worth noting that the proof of Proposition 6.2.4 does not depend on most of the finiteness conditions we used in this paper. And thus it translates very nicely to other variants of the theory such as the one suggested in Remark 5.1.4.

With the above proposition established, we now show our main theorem of this section.

Theorem 6.2.5. *Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$ which induces the Morse filtration $\mathcal{M} = \{A^p, b'\}$. Then the persistent homology groups of $\mathcal{F}$ and $\mathcal{M}$ are isomorphic, that is for every p, p', d, the following equation holds.*

$$H_d^{p \to p'}(\mathcal{F}) \cong H_d^{p \to p'}(\mathcal{M})$$

The proof of this theorem will be done inductively and will be similar to the proof of Theorem 4.3.7. Before that however we will have to identify a key property of every cell in a filtration of abstract cell complexes.

Definition 6.2.6. Birthframe

Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes. We define the birthframe *of each cell $\beta \in \bigcup_p X^p$ by $X^{B(\beta)}$ where B is defined as follows.*

$$B(\beta) := \min\{p \mid \beta \in X^p\}$$

Remark 6.2.7. *Note that $\beta \in X^p$ for all $p \geq B(\beta)$ by the subcomplex relation, see Definition 3.1.6. Moreover if $\beta \prec \alpha$ in any X^p, $B(\beta) \leq B(\alpha)$, since by the subcomplex property $\alpha \in X^{B(\alpha)}$ implies $\beta \in X^{B(\alpha)}$.*

Then we introduce a filtered version of a q-reduced complex as in Definition 4.3.8.

Definition 6.2.8. q-reduced Filtration

Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$. For $q \in \bigcup_p Q^p$, the q-reduced filtration *of $\mathcal{F}$ is $\mathcal{F}_q = \{X^p_q, b_q\}$, where (X^p_q, b_q) is the q-reduced subcomplex of (X^p, b) when $p \geq B(q)$, and $(X^p_q, b_q) = (X^p, b)$ otherwise.*

Obviously, "reduce by q" only makes sense for X^p when $q \in X^p$, thus the definition. The next proposition shows that what we get is indeed a filtration of abstract cell complexes. In fact similar to the case of just abstract cell complexes, the (persistent) homology groups are also preserved.

Proposition 6.2.9. *Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$. For any $q \in \bigcup_p Q^p$, the q-reduced filtration $\mathcal{F}_q = \{X^p_q, b_q\}$ is indeed a filtration of abstract cell complexes.*

Proof. We first define a simple filtered acyclic partial matching $\{(A^p_q, w : Q^p_q \to K^p_q)\}$. For $p < B(q)$, note $q \notin Q^p \subset X^p$, and we define $A^p_q = X^p$, $Q^p_q := \emptyset$, $K^p_q := \emptyset$. For $p \geq B(q)$, note that $q \in Q^p \subset X^p$, then we define $A^p_q := X^p \setminus \{q, w(q)\}$, $Q^p_q := \{q\}$, $K^p_q := \{w(q)\}$. The acyclic condition is trivially true since Q^p_q is at most a one-element set, while the other conditions of Definition 6.2.1 are clearly satisfied.

Now we make an observation that the Morse filtration induced by $\{(A^p_q, w : Q^p_q \to K^p_q)\}$ is exactly the q-reduced filtration. More precisely, the Morse complex induced by $(A^p_q, w : Q^p_q \to K^p_q)$ on each X^p is the q-reduced complex.

For $p < B(q)$ this is again trivial since both complexes are X^p itself. For $p \geq B(q)$, it follows from Remark 4.3.11.

Hence, Proposition 6.2.4 proves that the q-reduced filtration is indeed a filtration of abstract cell complexes.

□

Proposition 6.2.10. *Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$. For any $q \in \bigcup_p Q^p$, let the q-reduced filtration be $\mathcal{F}_q = \{X^p_q, b_q\}$. Then for every p, p', d, the following equation holds.*

$$H_d^{p \to p'}(\mathcal{F}) \cong H_d^{p \to p'}(\mathcal{F}_q)$$

Proof. Let $q \in \bigcup_p Q^p$.

For $p < B(q)$, define $\phi^p : C(X^P) \to C(X_q^p)$ and $\psi^p : C(X_q^p) \to C(X^p)$ to be the identity chain map.

For $p \geq B(q)$, we know from Proposition 4.3.13 that there are chain maps $\phi^p : C(X^p) \to C(X_q^p)$ and $\psi^p : C(X_q^p) \to C(X^p)$ such that $\phi^p \circ \psi^p$ and $\psi^p \circ \phi^p$ are chain homotopic to the identity map, since X_q^p is the q-reduced complex of X^p.

It remains to show that the collection of chain maps defined by $\Phi = \{\phi^p\}$ and $\Psi = \{\psi^p\}$ are indeed filtered chain maps, and Lemma 6.1.6 will finish the proof. Hence we need to show that the following diagrams commute.

$$\begin{array}{ccc} C_d(X^p) & \xrightarrow{i_d^{p\to p'}} & C_d(X^{p'}) \\ \downarrow{\scriptstyle \phi_d^p} & & \downarrow{\scriptstyle \phi_d^{p'}} \\ C_d(X_q^p) & \xrightarrow{i_d^{p\to p'}} & C_d(X_q^{p'}) \end{array} \qquad \begin{array}{ccc} C_d(X^p) & \xrightarrow{i_d^{p\to p'}} & C_d(X^{p'}) \\ \uparrow{\scriptstyle \psi_d^p} & & \uparrow{\scriptstyle \psi_d^{p'}} \\ C_d(X_q^p) & \xrightarrow{i_d^{p\to p'}} & C_d(X_q^{p'}) \end{array}$$

For the first diagram, note that $C_d(X^p)$ is a submodule of $C_d(X^{p'})$. When $p \geq B(q)$, by definition ϕ_d^p and $\phi_d^{p'}$ are both identities on every basis element except on q and $w(q)$ and also act the same way on q and $w(q)$. Hence $\phi_d^p = \phi_d^{p'} \mid_{C_d(X^p)}$. When $p < B(q)$, ϕ_d^p is always the identity and so $\phi_d^p = \phi_d^{p'} \mid_{C_d(X^p)}$ still holds. Hence the first diagram commutes.

For the second diagram, similarly, when $p \geq B(q)$, $\phi_d^p = \phi_d^{p'} \mid_{C_d(X_q^p)}$ by definition, which also holds when $p < p' < B(q)$ when both functions are the identity. When $p < B(q) \leq p'$, note that $\psi_d^{p'}(\alpha) = \alpha - \dfrac{b(\alpha, q)}{b(w(q), q)} w(q)$ only differs from the identity when $b(\alpha, q) \neq 0$, i.e. $q \prec \alpha$. But then as remarked in Definition 6.2.6, $B(\alpha) \geq B(q) = p'$ and so $\alpha \notin C_d(X_q^p)$. That is $\psi_d^{p'}$ is the identity on $C_d(X_q^p)$ and $\phi_d^p = \phi_d^{p'} \mid_{C_d(X_q^p)}$ still holds. Hence the second diagram commutes.

□

Proof of Theorem 6.2.5. Let $\mathcal{F} = \{X^p, b\}$ be a filtration of abstract cell complexes with filtered acyclic partial matching $\{(A^p, w : Q^p \to K^p)\}$, which induces the Morse filtration $\mathcal{M} = \{A, b'\}$.

Similar to the proof of Theorem 4.3.7, we inductively remove pairs of non-critical cells $q \in \bigcup_p Q^p$ and $w(q) \in \bigcup K^p$ from $\mathcal{F}$.

More precisely, let $\bigcup_p Q^p = \{q_1, ..., q_n\}$ and define filtrations $\mathcal{F}_0 = \{X_0^p, b_0\}, ...,$
$\mathcal{F}_n = \{X_n^p, b_n\}$ so that $\mathcal{F}_0 = \mathcal{F}$ and for every i, $\mathcal{F}_{i+1}$ is the q_{i+1}-reduced filtration of $\mathcal{F}_i$. Then by Proposition 6.2.10, for every p, p', d, $H_d^{p\to p'}(\mathcal{F}_i) \cong H_d^{p\to p'}(\mathcal{F}_{i+1})$. So by induction we can conclude that every $\mathcal{F}_i$ has isomorphic persistent homology groups.

Lastly, we can fix p and look at each i. Let $\{(A_i^p, w : Q_i^p \to K_i^p)\}$ be the acycling partial matching it adapted in the process. If $q_{i+1} \in X_i^p$, then X_{i+1}^p is the q_{i+1}-reduced complex of X_i^p. Otherwise $X_i^p = X_{i+1}^p$. Note that at X_n^p, all q_i's in $\bigcup_p Q^p = \{q_1, ..., q_n\}$ have been removed and Q^p in particular is empty, but by Lemma 4.3.15 that means (X_n^p, b_n) is just (A, b'). Since this is true for every p, $\mathcal{F}_n$ is precisely $\mathcal{M}$ and so $\mathcal{M}$ and $\mathcal{F}$ have isomorphic persistent homology groups as desired.

□

7 Conclusion and Future Direction

The work of Forman [3] laid the foundation of discrete Morse theory. In this paper we have demonstrated one of the applications of discrete Morse theory. Specific to TDA, we have shown how Mischaikow and Nanda [4] connected discrete Morse theory and persistent homology and showed how to use discrete Morse theory to accelerate the computation of persistent homology by drastically reducing the number of cells involved, as discussed in Section 6. They have also provided what has now become a commonly used algorithm, see [4] for more details.

Moreover, the framework of abstract cell complexes, introduced by Tucker [7] and refined by Mischaikow and Nanda, has the potential to be further expanded and applied to other theories of CW complexes. One also could find applications to not just persistent homology, but also many of its variants. Here we list some related questions and research subjects.

- Zigzag persistence [13] is a generalization of persistent homology with various potential applications, and the same framework to optimize computation introduced in this paper also seems applicable. Some work has already been done on this particular topic, see [14].
- In the construction of our framework we employed a few finiteness and discreteness conditions with impunity, since they align very well with the practical applications in data analysis. For example, an abstract cell complex is restricted to having finitely many cells, and a filtration is restricted to having finitely many frames and thus is indexed discretely. What results can we get if we consider some of these parameters to be infinite?
- A generalization of persistent homology would be to replace the ordered index set by some partially ordered set. The work done in this paper seems to translate very well to this particular setting. How about if we employ an even more complicated index set?

The rise of TDA provided algebraic topology many new applications of its theories and stimulated many new ideas, and discrete Morse theory is one good example of both. It is exciting how many abstract results have found its way to practical uses in real life because of it, and I look forward to more of this happening in the future.

References

[1] A. Lundell and S. Weingram, *The Topology of CW Complexes*, Van Nostrand Reinhold, New York, (1969).

[2] R. Ghrist, *Barcodes: the persistent topology of data*, Bulletin of the American Mathematical Society (New Series) 45, 1 (2008), 61–75.

[3] R. Forman, *Morse Theory for Cell Complexes*, Advances in Mathematics, Volume 134, issue 1, (1998), 90-145.

[4] K. Mischaiwkow and V. Nanda. *Morse Theory for Filtrations and Efficient Computation of Persistent Homology*, Discrete Comput Geom 50, (2013), 330–353.

[5] D.N. Kozlov, *Discrete Morse theory for free chain complexes*, C. R. Acad. Sci. Paris, Ser. I 340 (2005).

[6] R. Forman, *A User's Guide to Discrete Morse Theory*, Seminaire Lotharingien de Combinatoire 48 (2002), Article B48c.

[7] A.W. Tucker *Cell Spaces*, Annals of Mathematics, Second Series, Vol 37, No. 1 (1936), pp. 92-100.

[8] J. Milnor *Morse Theory*, Princeton University Press, Princeton (1963).

[9] A. Hatcher *Algebraic Topology*, Cambridge University Press, Cambridge (2002).

[10] J.R. Munkres *Elements of Algebraic Topology*, Westview Press (1984).

[11] G. Carlsson *Topology and Data*, Bulletin (New Series) of the American Mathematical Society Volume 46, Number 2, (2009), 255–308.

[12] J. Milnor *Lectures on the h-cobordism theorem*, "Princeton Mathematical Notes", Princeton Univ. Press, Princeton, N.J., (1965).

[13] G. Carlsson and V. De Silva, *Zigzag Persistence*, Found Comput Math 10, (2010), 367–405.

[14] C. Maria and H. Schreiber *Discrete Morse Theory for Computing Zigzag Persistence*, Algorithms and Data Structures. WADS 2019. Lecture Notes in Computer Science, vol 11646. Springer, Cham, (2019), 538-552.