

DISSERTATION

THE MÖBIUS HOMOLOGY LAPLACIAN FOR PERSISTENCE

Submitted by

Tatum D. Rask

Department of Mathematics

In partial fulfillment of the requirements

For the Degree of Doctor of Philosophy

Colorado State University

Fort Collins, Colorado

Spring 2026

Doctoral Committee:

Advisor: Amit Patel

Christopher Peterson

Clayton Shonkwiler

Dustin Tucker

Copyright by Tatum D. Rask 2026

All Rights Reserved

ABSTRACT

THE MÖBIUS HOMOLOGY LAPLACIAN FOR PERSISTENCE

Persistent homology provides a framework for studying the evolution of topological features across filtered spaces. While persistence diagrams give discrete invariants that record the birth and death of homological features, they do not directly encode analytic or geometric structure associated to persistence modules. This dissertation develops a spectral framework for persistent homology by introducing and studying a Laplace operator on the Möbius chain complex, called the *Möbius homology Laplacian*.

Möbius homology categorifies Möbius inversion for persistence modules by replacing integer-valued data with vector-space-valued chain complexes. By equipping these chain complexes with inner products, we define a Laplace operator on the Möbius chain complex whose kernel provides canonical cycle representatives for the associated Möbius homology space. We analyze the spectral properties of this operator in two primary settings: one-parameter persistence modules and birth-death modules arising from filtrations of simplicial complexes. In both cases, we show that the spectrum reflects the combinatorial structure of the associated persistence diagram.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank Amit Patel. Your wisdom and guidance has been instrumental in my success and growth while at CSU. Thank you to my committee members - Chris Peterson, Clay Shonkwiler, and Dustin Tucker - for your belief and interest in this project.

To my collaborators: your insights and guidance played an enormous role in my development as a mathematician. Though none of my collaborative projects are present in this dissertation, I am grateful for so many opportunities to do applied topology with my colleagues. Thank you Alex Elchesen, Iryna Hartsock, Jose Perea, Luis Scoccola, Hitesh Gakhar, Johnathan Bush, Niko Schon-scheck, Ling Zhou, Lisbeth Fajstrup, Brittany Fasy, Wenwen Li, Lydia Mezrag, Francesca Tombari, and Živa Urbančič. I would also like to acknowledge my mentors and coworkers from the various internships I completed. From NASA, thank you Alan Hylton, Brendan Mallery, Jihun Hwang, Mark Ronnenberg, Miguel Lopez, and Oliver Chiriac. From GTRI, thank you Branden Stone, Andrew Tawfeek, Sunia Tanweer, and Vicente Gonzalez-Bosca. I am a better mathematician and person for having worked with you.

To my Math 271 students: thank you for such a joyful teaching experience over the past two years. Your interest and insightful questions were inspiring to me both as a teacher and researcher. Our discussions, particularly on the geometry of linear transformations, certainly had an influence on this work. I am certain that your hard work and curiosity will carry you far.

Finally, I would like to thank the friends and family who supported me throughout my graduate studies. Every inner tube water polo game (go Sitting Ducks!), Growing Food Security event, and Gnar Runners trail race made the past six years some of my favorite ones. I am so grateful to have such an extensive support system, and there are far too many amazing people in my life to list here. Special thanks to Mom, Dad, Keenan, Sean, Grandma, and Grandpa for the unwavering support and love. Thank you Seth, Madelyn, Joe, Erin, Kelly, Andrew, Amaury, Cori, Bela, Grace, Katie, Anna, and Lauren: you all bring so much meaning and joy to my life.

TABLE OF CONTENTS

| | | |
|-----------|--|-----|
| | ABSTRACT | ii |
| | ACKNOWLEDGEMENTS | iii |
| Chapter 1 | Introduction | 1 |
| 1.1 | Category Theoretic Language | 4 |
| Chapter 2 | Combinatorial Hodge Laplacians | 8 |
| 2.1 | Inner Product Spaces | 9 |
| 2.1.1 | Operators on Inner Product Spaces | 12 |
| 2.2 | Simplicial Homology | 14 |
| 2.3 | Combinatorial Hodge Laplacians | 18 |
| 2.3.1 | Graph Laplacians | 19 |
| 2.3.2 | Simplicial Hodge Laplacians and Discrete Hodge Theory | 23 |
| 2.3.3 | Simplicial Cosheaf Laplacians | 26 |
| Chapter 3 | Möbius Inversion in One-Parameter Persistence | 31 |
| 3.1 | Möbius Inversion on Posets | 32 |
| 3.2 | One-Parameter Persistence | 35 |
| 3.2.1 | Interval Decomposition of One-Parameter Modules | 36 |
| 3.2.2 | Rank Function and Möbius inversion | 38 |
| 3.2.3 | Birth-Death Function and Möbius Inversion | 41 |
| 3.2.4 | Equivalence of Methods | 46 |
| Chapter 4 | Möbius Homology and Multiparameter Persistence | 49 |
| 4.1 | Non-Functoriality of Images | 50 |
| 4.2 | Categorification of the Birth-Death Function | 51 |
| 4.3 | Möbius Homology | 56 |
| 4.3.1 | Möbius Homology of One-Parameter Modules | 60 |
| 4.3.2 | Möbius Homology of the Birth-Death Module | 64 |
| 4.4 | Motivation for a Laplacian on the Möbius Chain Complex | 66 |
| Chapter 5 | The Möbius Homology Laplacian | 68 |
| 5.1 | Constructing a Laplacian on the Möbius Chain Complex | 69 |
| 5.1.1 | Möbius Homology Laplacian for One-Parameter Modules | 72 |
| 5.2 | Discussion and Directions for Future Work | 82 |
| Chapter 6 | Möbius Homology Laplacian for Simplicial Filtrations | 84 |
| 6.1 | Birth-Death Modules from Simplicial Filtrations | 85 |
| 6.2 | Hilbert Space Structure and Möbius Homology Laplacians | 88 |
| 6.3 | Matrix Representation of the Degree-Zero Laplacian | 91 |
| 6.4 | Harmonic Representatives and Newly Added Simplices | 92 |
| 6.5 | Combinatorial Spectral Properties of the Degree-Zero Laplacian | 97 |

| | | |
|------------------------|---|-----|
| 6.5.1 | Spectrum in the One-Parameter Setting | 97 |
| 6.5.2 | General Spectral Properties | 102 |
| 6.6 | Discussion and Future Directions | 105 |
| 6.6.1 | Relationship to Related Work | 106 |
| 6.6.2 | Future Work | 107 |
| Bibliography | | 109 |
| Appendix A | Examples of Filtrations | 114 |
| A.1 | One-Parameter Examples | 114 |
| A.2 | Multiparameter Examples | 117 |

Chapter 1

Introduction

Persistent homology was developed as a tool for extracting topological information from data [1–3]. In topological data analysis, one often constructs a filtered simplicial complex from a point cloud of data (for example, the Vietoris-Rips or Čech complexes [4, Section 2.2]) and tracking the evolution of homological features as the filtration value varies. A key feature of persistent homology is *stability*: persistence diagrams vary continuously under perturbations of the data, ensuring robustness to noise [5, 6].

Although motivated by data analysis, persistent homology has developed into a rich mathematical theory. Filtered simplicial complexes give rise to *persistence modules*: sequences of vector spaces with linear maps between them. In the one-parameter setting, persistence modules admit complete discrete invariants in the form of persistence diagrams [7]. Persistence diagrams record the birth and death of homological features, and they can be computed using algebraic and combinatorial methods such as rank invariants [7], birth-death functions [8], and interval decompositions [9]. However, while these constructions are equivalent in the one-parameter case, they are combinatorial in nature. They record when features exist, but they do not provide canonical cycle representatives or analytic structure. This dissertation is motivated by the following question:

Can persistence admit an intrinsic analytic structure, analogous to Hodge theory in classical topology?

In classical topology, algebraic and combinatorial invariants are enriched by analytic methods such as the Hodge Laplacian on a Riemannian manifold or its combinatorial counterpart on a simplicial complex. The kernel of the Hodge Laplacian recovers homology spaces and provides canonical harmonic representatives for each homology class, and its non-zero spectrum reflects geometric and combinatorial structure [10–13]. Persistence modules have not traditionally admit-

ted such operators. This dissertation develops a framework in which an analytic structure naturally emerges. To do so, we utilize *Möbius homology*, a categorification of Möbius inversion.

Möbius inversion is a classical tool for extracting local contribution from cumulative data on finite posets [14]. In persistence theory, rank and birth-death functions encode cumulative information about the presence of homological features across intervals. The persistence diagram can be computed by applying Möbius inversion to these functions, as described in Chapter 3.

Möbius homology categorifies Möbius inversion [15]. Instead of assigning integers to intervals, it assigns vector spaces organized into a chain complex indexed by intervals (see Chapter 4). In the one-parameter setting, Möbius homology recovers the same interval decomposition data as persistence diagrams (see Section 4.3.1). In the case of mutliparameter modules, where no complete discrete invariants exist, Möbius homology distinguishes non-isomorphic modules that are indistinguishable by Möbius inversion and persistence alone. The additional structure of chain complexes allow for the introduction of inner products, adjoint operators, and ultimately a combinatorial Laplacian acting on the Möbius chain spaces (see Chapter 5).

The central idea of this dissertation is that Möbius homology admits a natural Laplace operator, called the *Möbius homology Laplacian*, that provides analytic structure for persistence modules.

Main Contributions

This dissertation establishes foundational elements of a spectral theory for persistence modules.

First, we define the Möbius homology Laplacian for persistence modules indexed by finite posets. This operator is self-adjoint and positive-semidefinite, and its kernel is canonically isomorphic to Möbius homology (see Theorem 5.1.2).

Second, in the one-parameter setting, we provide an explicit spectral description of the degree-one Möbius homology Laplacian. The eigenvalues are determined by the combinatorial organization of the persistence diagram (see Theorem 5.1.13).

Third, when persistence modules arise from simplicial filtrations, the analytic viewpoint reflects the geometry of the underlying simplicial complex: this is discussed in detail in Chapter 6.

The degree-zero Laplacian acts on a birth-death module constructed directly from the filtration, and harmonic representatives emphasize newly added simplices (see Theorem 6.4.2). Moreover, even in the multiparameter setting, spectral invariants such as the trace admit combinatorial descriptions in terms of the birth-death function (see Theorem 6.5.5).

In all, this work introduces the operator, establishes its fundamental properties, and connects it to existing perspectives in persistence theory. Many deeper analytic and geometric consequences remain to be explored.

Related Work

Spectral methods have previously been introduced for filtered simplicial complexes through constructions such as persistent Laplacians [16–18] and harmonic persistent homology [16, 19]. These approaches act on filtered chain complexes at individual parameter values. In contrast, the Möbius homology Laplacian acts on the *Möbius* chain complex, whose construction depends on the combinatorics of the entire indexing poset. A similar perspective appears in recent work defining the Grassmannian persistence diagram [20, 21]. Nonetheless, the cycle representatives obtained via the Möbius homology Laplacian coincide with these previously studied constructions. We make the relationship between our work and existing literature more precise in Section 6.6.1.

Organization

This dissertation is organized as follows. Chapter 2 develops analytic background from combinatorial Hodge theory. Chapter 3 reviews Möbius inversion and one-parameter persistence. Chapter 4 introduces Möbius homology as a categorification of Möbius inversion. Chapter 5 constructs the Möbius homology Laplacian and analyzes its spectral properties in the one-parameter case. Chapter 6 specializes to simplicial filtrations, establishing geometric interpretations of cycle representatives and combinatorial descriptions of spectral invariants.

1.1 Category Theoretic Language

This section reviews only the categorical notions required for this dissertation: categories, functors, and natural transformations. These concepts provide a concise language for objects central to this work, such as persistence modules and cosheaves. No further categorical machinery is assumed, but the interested reader can refer to [22] for further discussion.

Definition 1.1.1: A **category** \mathcal{C} consists of **objects** A, B and **morphisms** $f : A \rightarrow B$ satisfying the following properties. Each object has an identity morphism $1_A : A \rightarrow A$. For every $f : A \rightarrow B$ and $g : B \rightarrow C$, the composition $gf : A \rightarrow C$ is also a morphism in \mathcal{C} . For $f : A \rightarrow B$, composing with the identity does not change the morphism f : that is, $1_B f = f 1_A = f$. Finally, composition is associative.

Below are examples of categories that we will see throughout this work.

Example 1.1.2: Denote by $\text{Vec}_{\mathbb{F}}$ the category of vector spaces over the field \mathbb{F} . Morphisms are given by linear maps. We will often consider vector spaces over the real numbers: $\text{Vec}_{\mathbb{R}}$.

Definition 1.1.3: A **partially ordered set**, or a **poset**, is a set P endowed with an order relation \leq such that, for all elements $a, b, c \in P$, the following hold.

1. $a \leq a$ (*reflexivity*),
2. If $a \leq b$ and $b \leq a$, then $a = b$ (*anti-symmetry*),
3. If $a \leq b$ and $b \leq c$, then $a \leq c$ (*transitivity*).

If $p \leq q$ and $p \neq q$, we write $p < q$.

Example 1.1.4: Any poset P is a category. Elements are given by poset elements $p, q \in P$ and morphisms $p \rightarrow q$ are given by relations $p \leq q$.

We will also consider various categories associated to simplicial complexes. First, we review the definition of a simplicial complex.




Figure 1.1: The standard 2-simplex, Δ^2 , is given by *all* non-empty subsets of the vertex set $\{v_1, v_2, v_3\}$.

Definition 1.1.5: A **simplicial complex** K is a collection of sets, called **simplices**, that is closed under taking subsets. That is, if $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$ and we call τ a **face** of σ .

See Figure 1.1 for a visualization of a simplicial complex.

Example 1.1.6: Any simplicial complex K is a poset (thus, a category) where elements are given by simplices σ in K and the order is given by reverse inclusion. That is, we say $\sigma \leq \tau$ if τ is a face of σ .

See Figure 1.2 for the *Hasse diagram* of a simplicial complex. A Hasse diagram is a visualization of a poset P , constructed as follows. We say q covers p , denoted $p <_1 q$, whenever $p < q$ and there is no $t \in P$ so that $p < t < q$. In the Hasse diagram, every $p <_1 q$ is shown as a line connecting p and q , with the larger element q drawn above the smaller element p . Properties such as transitivity and reflexivity are implied.

Example 1.1.7: Denote by ΔK the poset (thus, category) of *subcomplexes* of K . Elements are subsets $L \subseteq K$ such that L is a simplicial complex itself. The order is given by inclusion \subseteq . That is, there is a morphism $L \rightarrow L'$ whenever $L \subseteq L'$.

It is common in mathematics to map one mathematical object to another type of object: for example, as we will see in Section 2.2, simplicial homology maps simplicial complexes to vector spaces. It is expected that such maps preserve important mathematical properties. Namely, it is expected that such maps respect not only the objects in a category, but also the morphisms. These maps are called *functors* in the language of category theory.

Definition 1.1.8: Let \mathcal{C}, \mathcal{D} be categories. A **functor** $F : \mathcal{C} \rightarrow \mathcal{D}$ assigns to every $A \in \mathcal{C}$ an object $F(A)$ in \mathcal{D} and to every morphism $f : A \rightarrow B$ in \mathcal{C} a morphism $F(f) : F(A) \rightarrow F(B)$




Figure 1.2: The Hasse diagram of the simplicial complex Δ^2 shown in Figure 1.1.

in \mathcal{D} . This morphism must obey composition and identity morphisms. That is, $F(1_A) = 1_{F(A)}$ and, for $gf : A \rightarrow C$, $F(gf) = F(g)F(f)$.

Throughout this dissertation, we will study various functors. Below is a type of functor that is commonly studied in persistent homology.

Definition 1.1.9: Let P be a poset and K a simplicial complex. A **filtration** of K is a functor from P to the category of subcomplexes of K , $F : P \rightarrow \Delta K$.

Intuitively, throughout a filtration of a simplicial complex, we slowly build the complex K as we move through the poset P . Since poset relations map to inclusions of subcomplexes, the simplicial complex K is "growing" throughout the filtration F .

Given two categories \mathcal{C} and \mathcal{D} , the collection of functors $\mathcal{C} \rightarrow \mathcal{D}$ also forms a category. The morphisms in such a category are called *natural transformations*.

Definition 1.1.10: Given two functors $F, G : \mathcal{C} \rightarrow \mathcal{D}$, a **natural transformation** $\varphi : F \Rightarrow G$ between them assigns to every object $A \in \mathcal{C}$ a morphism $\varphi_A : F(A) \rightarrow G(A)$ so that for every morphism $f : A \rightarrow B$ in \mathcal{C} the following diagram commutes.

$$\begin{array}{ccc}
 F(A) & \xrightarrow{F(f)} & F(B) \\
 \varphi_A \downarrow & & \downarrow \varphi_B \\
 G(A) & \xrightarrow{G(f)} & G(B)
 \end{array}$$

That is, $G(f)\varphi_A = \varphi_B F(f)$. We call φ_A the **component** of the natural transformation at A . We call φ a **natural isomorphism** if every φ_A is an isomorphism.

Below are two other types of functors that we will encounter throughout this work.

Definition 1.1.11: Let P be a poset. A **persistence module** is a functor $M : P \rightarrow \text{Vec}_{\mathbb{R}}$.

That is, persistence modules map poset elements to vector spaces and poset relations to linear maps so that everything commutes. Morphisms between persistence modules $M, N : P \rightarrow \text{Vec}_{\mathbb{R}}$, then, are natural transformations $\varphi : M \Rightarrow N$. Two modules are isomorphic if there is a natural isomorphism between them.

Given two persistence modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ and $N : P \rightarrow \text{Vec}_{\mathbb{R}}$, their direct sum is another persistence module $M \oplus N : P \rightarrow \text{Vec}_{\mathbb{R}}$ defined as follows: for poset elements $a \in P$, $(M \oplus N)(a) = M(a) \oplus N(a)$. For poset relations, $(M \oplus N)(a \leq b) = M(a \leq b) \oplus N(a \leq b)$.

Definition 1.1.12: Let K be a simplicial complex. A **simplicial cosheaf** valued in $\text{Vec}_{\mathbb{R}}$ is a functor $\mathcal{F} : K \rightarrow \text{Vec}_{\mathbb{R}}$.

Intuitively, simplicial cosheaves assign vector spaces to each simplex, and the linear maps propagate information *down* in dimension. That is, if $\sigma \leq \tau$ (τ is a face of σ), then there is a linear map $\mathcal{F}(\sigma \rightarrow \tau)$. Again, morphisms of simplicial cosheaves are natural transformations.

Overall, category theory provides a concise and unified language to describe the central objects in this dissertation. Persistence modules, simplicial cosheaves, and filtrations of simplicial complexes will be the central objects of study in future chapters.

Chapter 2

Combinatorial Hodge Laplacians

The central objects in this dissertation are combinatorial Laplace operators. In this chapter, we develop the mathematical framework for studying these operators, which, in our setting, are square matrices acting on finite-dimensional inner product spaces. These operators provide an analytic method for studying discrete structures such as graphs, simplicial complexes, and simplicial cosheaves. In particular, the kernel of the combinatorial Laplacian provides canonical "harmonic" representatives for homology classes, while the spectrum of the operator reflects combinatorial and geometric structure of the underlying object.

The motivation for studying such operators originates in differential geometry. Building upon de Rham's work on cohomology theories for Riemannian manifolds and algebraic varieties, Hodge sought to relate topological and geometric properties using analytic tools. To do so, he introduced the Hodge Laplace operator, which acts on differential forms on compact manifolds. The kernel of the Hodge Laplacian consists of harmonic forms, which provide canonical energy-minimizing representatives for de Rham cohomology classes. Furthermore, the Hodge decomposition theorem states that any differential form decomposes into harmonic, exact, and coexact components. This decomposition separates topological information from non-harmonic contributions.

Scalar functions are examples of differential 0-forms. Related to the degree-zero Hodge Laplacian (differentiated by a sign and a constant related to curvature) is the classical Laplacian on Euclidean space and the Laplace-Beltrami operator on Riemannian manifolds. Intuitively, this operator measures how the value of a function at a point deviates from the average of the function at nearby points. Elements in the kernel of the Laplacian are called harmonic functions and are known to minimize Dirichlet energy, which measures local variation of a function.

In this dissertation, we are interested in combinatorial Laplace operators defined on discrete objects such as graphs, simplicial complexes, and simplicial cosheaves. Much like in the smooth setting, these operators are constructed using boundary maps and their adjoints. By equipping chain

spaces with inner product structures, we obtain combinatorial Laplace operators whose kernels recover the associated homology and whose spectra reflect combinatorial and geometric features of the underlying data.

We begin this chapter by introducing prerequisite material from inner product spaces and simplicial homology. Using these tools, we introduce Laplace operators on graphs, simplicial complexes, and simplicial cosheaves. Throughout, we discuss combinatorial analogs of harmonicity and Hodge decomposition and summarize known results on the spectra of these operators.

2.1 Inner Product Spaces

We begin by recalling structures from linear algebra that will be necessary in our discussion of combinatorial Hodge Laplacians and their spectrum. Throughout, all vector spaces are finite-dimensional over \mathbb{R} . See [23, Chapter I] and [24, Chapters 6-7] for more details.

Definition 2.1.1: An **inner product** on a real vector space V is a map $\langle \cdot, \cdot \rangle_V : V \times V \rightarrow \mathbb{R}$ such that for all $x, y, z \in V$ and $\alpha \in \mathbb{R}$:

1. $\langle x, y \rangle_V = \langle y, x \rangle_V$ (*symmetry*),
2. $\langle \alpha x + \beta y, z \rangle_V = \alpha \langle x, z \rangle_V + \beta \langle y, z \rangle_V$ (*linearity*),
3. $\langle x, x \rangle_V \geq 0$ with equality if and only if $x = 0$ (*positive definiteness*).

When the space V is clear from the context, we write $\langle \cdot, \cdot \rangle$.

Definition 2.1.2: We call $(V, \langle \cdot, \cdot \rangle_V)$ an **inner product space**. When the inner product structure is clear from the context, we will simply write V .

Definition 2.1.3: The **norm** induced by $\langle \cdot, \cdot \rangle_V$ is $\|x\|_V = \sqrt{\langle x, x \rangle_V}$. The **energy** of a vector x is given by its norm squared $\|x\|_V^2 = \langle x, x \rangle_V$.

Definition 2.1.4: A **Hilbert space** is an inner product space V that is complete with respect to the norm induced by the inner product $\langle \cdot, \cdot \rangle_V$.

For the purposes of this dissertation, we will focus on finite-dimensional vector spaces over \mathbb{R} . In finite dimensions, every inner product space over \mathbb{R} is complete, so any finite-dimensional inner product space is a Hilbert space.

We denote by $\text{Hilb}_{\mathbb{R}}$ the category of Hilbert spaces over \mathbb{R} . The morphisms are bounded linear maps. We say a linear map $f : V \rightarrow W$ is *bounded* if there exists a constant $L > 0$ such that, for all $v \in V$, $\|f(v)\|_W \leq L\|v\|_V$. Indeed, any linear map between finite-dimensional vector spaces (which is the setting we are working in) is bounded.

Definition 2.1.5: Vectors $x, y \in V$ are **orthogonal** if $\langle x, y \rangle_V = 0$. For a subspace $W \subseteq V$, the **orthogonal complement** of W in V is

$$W^\perp = \{ v \in V \mid \langle v, w \rangle_V = 0 \text{ for all } w \in W \}.$$

Given an inner product on V and a subspace $W \subseteq V$, we can define a linear map that projects vectors in V onto the subspace W .

Lemma 2.1.6: Let V be a finite-dimensional inner product space and let $W \subseteq V$ be a subspace. Then there exists a unique linear map $\text{proj}_W : V \rightarrow W$ such that $x - \text{proj}_W(x) \in W^\perp$ for all $x \in V$. If $\{w_i\}_{i=1}^k$ is an orthonormal basis of W , then for all $v \in V$,

$$\text{proj}_W(v) = \sum_{i=1}^k \langle v, w_i \rangle_V w_i.$$

Lemma 2.1.7: For any subspace $W \subseteq V$, we have a direct sum decomposition

$$V = W \oplus W^\perp.$$

In particular, every $v \in V$ decomposes uniquely as $v = w + z$ with $w \in W$ and $z \in W^\perp$.

When taking a quotient of an inner product space by a subspace, it is not immediately obvious how the quotient space inherits an inner product structure. One method for understanding the inner

product structure on the quotient space is to consider an isomorphic inner product space. We give such an isomorphism in the lemma below.

Lemma 2.1.8: Let V be a Hilbert space and let W be a subspace of V . Then, $V/W \cong W^\perp$.

Proof. By Lemma 2.1.7, each $v \in V$ decomposes as $v = w + z$ with $w \in W$ and $z \in W^\perp$. Consider the projection map from Lemma 2.1.6: proj_{W^\perp} is surjective and $\text{proj}_{W^\perp}(v) = z$. Moreover, $\text{proj}_{W^\perp}(v) = 0 \iff z = 0 \iff v \in W$. That is, $\ker \text{proj}_{W^\perp} = W$ and by the first isomorphism theorem $V/W \cong W^\perp$. \square

Given a bounded linear map $f : V \rightarrow W$ between Hilbert spaces, we are interested in a "backwards" map $W \rightarrow V$ that preserves inner products in the following sense.

Definition 2.1.9: Let V, W be Hilbert spaces and let $f : V \rightarrow W$ be a morphism between them. Then, the **adjoint** of f , denoted f^* , is the unique map $f^* : W \rightarrow V$ such that

$$\langle f(v), w \rangle_W = \langle v, f^*(w) \rangle_V$$

for all $v \in V, w \in W$.

Again, since we are only considering finite-dimensional Hilbert spaces, every linear map has a unique adjoint. This follows from the Riesz Representation Theorem: see Definition 7.1 and the following discussion in [24, Chapter 7].

Let F be the matrix representation for $f : V \rightarrow W$ in terms of the bases $\{v_i\}$ for V and $\{w_i\}$ for W . We can compute the adjoint matrix F^* as follows. There is a unique square matrix M_V so that $\langle v, x \rangle_V = v^T M_V x$ for all vectors $v, x \in V$, and similarly there is a unique matrix M_W so that $\langle w, y \rangle_W = w^T M_W y$ for all vectors $w, y \in W$. We call these the *Gram matrices* for the inner products on V and W , respectively. By the definition of adjoints, $\langle Fv, w \rangle_W = \langle v, F^*w \rangle_V$ and, thus, $v^T F^T M_W w = v^T M_V F^* w$ for all $v \in V, w \in W$. The equality $F^T M_W = M_V F^*$, then, results in the following equation for adjoint matrices: $F^* = M_V^{-1} F^T M_W$. If both $\{v_i\}$ and $\{w_i\}$ are orthonormal bases, then $F^* = F^T$.

The following proposition relates the image of a bounded linear map to the kernel of its adjoint.

Proposition 2.1.10: Let V and W be Hilbert spaces, and let $f : V \rightarrow W$ be a bounded linear map.

Then, $\ker f^* = (\operatorname{im} f)^\perp$.

Proof. Let $x \in W$. Then,

$$\begin{aligned} x \in (\operatorname{im} f)^\perp &\iff \langle x, w \rangle_W = 0 \text{ for all } w \in \operatorname{im} f \\ &\iff \langle x, f(v) \rangle_W = 0 \text{ for all } v \in V \\ &\iff \langle f^*(x), v \rangle_V = 0 \text{ for all } v \in V \\ &\iff x \in \ker f^*. \end{aligned}$$

□

In the following section, we will focus specifically on *automorphisms*, also called *operators*, from a Hilbert space V to itself.

2.1.1 Operators on Inner Product Spaces

Given a (finite-dimensional) Hilbert space V , an *operator* on V is a bounded linear map from V to itself, $f : V \rightarrow V$. In this section, we will discuss various methods for studying such operators and their matrix representations, focusing specifically on spectral analysis. For the purposes of this section, fix a basis $\{v_i\}$ for V and let F be the matrix representation of $f : V \rightarrow V$ according to this basis. Notice that F is a square matrix with dimensions given by $\dim V$.

Definition 2.1.11: The **characteristic polynomial** of F is given by $\det(\lambda I_n - F)$, where n is the dimension of V .

Recall that a *zero* λ of the characteristic polynomial is an eigenvalue of F , and that λ comes associated with eigenvectors v such that $Fv = \lambda v$. That is, the eigenvectors of F are the vectors in V that are simply rescaled by the action of F . The rescaling factor is precisely the associated eigenvalue.

Remark 2.1.12: Observe that we defined the characteristic polynomial and, thus, the eigenvalues and vectors of an operator f in terms of its matrix representation F according to some basis $\{v_i\}$. Both the characteristic polynomial and the eigenvalues of f are independent of choice of basis. Indeed, suppose U^{-1} is the change of basis matrix for V that transforms the basis $\{v_i\}$ to another basis $\{u_i\}$. Then, $U^{-1}FUv = \lambda v \implies (FU)v = \lambda(Uv)$. Thus, the eigenvalues of f are basis-independent. The eigenvectors of f , as abstract elements in the vector spaces V , are also basis independent. However, we have shown that the *coordinates* of the eigenvectors will change according to the change of basis.

Below, we give various spectral properties of certain nice operators f . Motivated by the remark above, we will state such spectral properties in terms of the linear map f rather than a specific matrix representation.

Definition 2.1.13: Let f have eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ counted with algebraic multiplicity. The **trace** of f is $\text{tr}(f) = \sum_{i=1}^n \lambda_i$. The **determinant** of f is $\det(f) = \prod_{i=1}^n \lambda_i$.

As an immediate consequence of the definition above, the determinant is the constant term of the characteristic polynomial multiplied by $(-1)^n$. The trace is the negative of the coefficient in front of λ^{n-1} .

Definition 2.1.14: Given an operator $f : V \rightarrow V$ such that $f = f^*$, we call f **self-adjoint**.

Self-adjoint operators have various nice spectral properties that we will rely on throughout this dissertation.

Proposition 2.1.15: If $f : V \rightarrow V$ is self-adjoint, then its eigenvalues are real.

Proof. See [24, Proposition 7.27] for a proof that the characteristic polynomial of f has all real roots. Thus, all eigenvalues of f are real. □

Proposition 2.1.16: The operator $f : V \rightarrow V$ is self-adjoint if and only if V admits an orthonormal basis consisting of eigenvectors of f . Equivalently, the matrix representation of f with respect to such a basis is diagonal.

Proof. This is the real spectral theorem: see [24, Proposition 7.29]. \square

In other words, if F is the matrix representation of $f : V \rightarrow V$ according to some basis $\{v_i\}$, then there exists an orthogonal matrix U so that $U^{-1}FU = D$ for a diagonal matrix D . Indeed, in this case, the columns of U are eigenvectors of f with eigenvalue given by the respective diagonal entry of D .

Definition 2.1.17: We say that the operator $f : V \rightarrow V$ is **positive-semidefinite** if it is self-adjoint and $\langle f(v), v \rangle_V \geq 0$ for all $v \in V$.

Proposition 2.1.18: Let $f : V \rightarrow V$ be self-adjoint. Then f is positive-semidefinite if and only if all eigenvalues of f are nonnegative.

Proof. By Proposition 2.1.16, there exists an orthonormal basis $\{e_i\}$ of V consisting of eigenvectors of f , with $f(e_i) = \lambda_i e_i$. For any vector $v = \sum_i c_i e_i$, the inner product $\langle f(v), v \rangle$ is given by the sum $\sum_i \lambda_i c_i^2$. Thus, $\langle f(v), v \rangle \geq 0$ for all v if and only if $\lambda_i \geq 0$ for all i . \square

The central objects of study in this dissertation, combinatorial Hodge Laplacians, are self-adjoint, positive-semidefinite operators that act on graphs, simplicial complexes, or simplicial cosheaves. The fact that the eigenvalues of such operators are always non-negative is a useful insight for extracting combinatorial and geometric information about the underlying objects. Thus, before introducing combinatorial Hodge Laplacians, we will investigate simplicial complexes and their homology spaces.

2.2 Simplicial Homology

Now, we recall standard definitions and conventions for simplicial homology. See [25] for more details. The chain spaces and boundary maps defined here will be central to our discussion on combinatorial Hodge Laplacians. Indeed, simplicial complexes are the discrete structures for which we want to extend Hodge Laplacians. We restate the definition of a simplicial complex below, adding additional language.

Definition 2.2.1: A **simplicial complex** K is a collection of sets, called **simplices**, that is closed under taking subsets. That is, if $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$. We call σ an **i -simplex** if $|\sigma| = i + 1$. Let K^i be the **i -skeleton** of K (that is, all σ such that $|\sigma| \leq i + 1$), and denote the set of all i -simplices of K by $K^{[i]} = K^i - K^{i-1}$.

For the purposes of this dissertation, we will focus on *finite* simplicial complexes: that is, simplicial complexes K where each $K^{[i]}$ is a finite set. Furthermore, we will consider simplicial complexes where the empty set is *not* a simplex.

Fix a total order on the vertices of K . Such an order induces an orientation on simplices: that is, we write an i -simplex as an *ordered* list of vertices $\sigma = [v_0 v_1 \dots v_i]$. If $\tau = [v_{\pi(0)} v_{\pi(1)} \dots v_{\pi(i)}]$ for an odd permutation π of $\{j\}_{j=0}^i$, we say that σ and τ have opposite orientation and write $\tau = -\sigma$. If π was, instead, an even permutation, we say that σ and τ have the same orientation: $\tau = \sigma$.

Definition 2.2.2: For $i \geq 0$, the **i -th chain space** $C_i(K)$ is the real vector space with basis given by the set of *ordered* i -simplices $K^{[i]}$. That is, an i -chain has the following form:

$$z = \sum_{j=0}^n a_j \sigma_j,$$

where each $\sigma_j \in K^{[i]}$ and each $a_j \in \mathbb{R}$.

Define the **boundary map** $\partial_i : C_i(K) \rightarrow C_{i-1}(K)$ on i -simplices as follows:

$$\partial_i(\sigma) = \sum_{j=0}^i (-1)^j [v_0 \dots \hat{v}_j \dots v_i],$$

and extend linearly to obtain a map on all of $C_i(K)$. Here, \hat{v}_j means that the j th vertex has been removed.

Definition 2.2.3 ([19, Definition 4.2]): The **support** of an i -chain $z = \sum_{j=0}^n a_j \sigma_j \in C_i(K)$ is given by the set $\text{supp}(z) = \{\sigma_j \in K^{[i]} \mid a_j \neq 0\}$.

It is easy to check that $\partial_i \partial_{i+1} = 0$. Thus, we can collect all of the chain spaces and boundary maps into a chain complex:

$$\dots \xrightarrow{\partial_{i+1}} C_i(K) \xrightarrow{\partial_i} C_{i-1}(K) \xrightarrow{\partial_{i-1}} \dots \xrightarrow{\partial_1} C_0(K) \xrightarrow{\partial_0} 0.$$

Definition 2.2.4: For $i \geq 0$, define the **cycle** and **boundary** subspaces as

$$Z_i(K) = \ker \partial_i \quad \text{and} \quad B_i(K) = \text{im } \partial_{i+1}$$

respectively. The i -th **homology space** is given by the quotient $H_i(K) = Z_i(K)/B_i(K)$.

It is worth noting that different orderings on vertices and, thus, orientations on simplices will result in the same homology spaces up to isomorphism.

Definition 2.2.5: The **Euler characteristic** of K is given by

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

Now, consider a *subcomplex* A of the simplicial complex K . We wish to describe homology of K *relative* to the subcomplex A .

Definition 2.2.6: For $A \subseteq K$, the **relative i -chain space** is defined as the quotient

$$C_i(K, A) = \frac{C_i(K)}{C_i(A)}.$$

A quick check shows that the boundary maps on K restrict to boundary maps on A and, thus, respect the quotient. That is, we obtain maps $\partial_i : C_i(K, A) \rightarrow C_{i-1}(K, A)$. These maps also satisfy $\partial_i \partial_{i+1} = 0$ and, thus, we write the *relative* chain complex as

$$\dots \xrightarrow{\partial_{i+1}} C_i(K, A) \xrightarrow{\partial_i} C_{i-1}(K, A) \xrightarrow{\partial_{i-1}} \dots \xrightarrow{\partial_1} C_0(K, A) \xrightarrow{\partial_0} 0.$$

Definition 2.2.7: For $i \geq 0$, define the i -th **relative cycle and boundary spaces** as

$$Z_i(K, A) = \ker \partial_i \quad \text{and} \quad B_i(K, A) = \text{im } \partial_{i+1},$$

respectively. Then, the i -th **relative homology space** is $H_i(K, A) = Z_i(K, A)/B_i(K, A)$.

Example 2.2.8: Let Δ^2 be the standard 2-simplex: that is, Δ^2 is the simplex composed of all possible (non-empty) subsets of $\{v_1, v_2, v_3\}$. See Figure 1.1. Place the lexicographic order on vertices, $v_1 < v_2 < v_3$, and order the basis of $C_0(\Delta^2)$ accordingly. Order the basis for $C_1(\Delta^2)$ similarly: $[v_1v_2] < [v_2v_3] < [v_1v_3]$. Then, the boundary maps are defined as follows on basis elements:

$$\partial_2 [v_1v_2v_3] = [v_2v_3] - [v_1v_3] + [v_1v_2], \quad \partial_1 [v_iv_j] = v_j - v_i.$$

Their matrix representations are

$$\partial_2 = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix} \quad \text{and} \quad \partial_1 = \begin{pmatrix} -1 & 0 & -1 \\ 1 & -1 & 0 \\ 0 & 1 & 1 \end{pmatrix}.$$

Indeed, $H_0(\Delta^2) \cong \mathbb{R}$ whereas $H_i(\Delta^2) = 0$ for all $i > 0$. Thus, $\chi(\Delta^2) = 1$.

On the other hand, consider the subcomplex given by the *boundary* of the 2-simplex: $\partial\Delta^2$. This is the simplicial complex consisting of the simplices $\{v_1, v_2, v_3, [v_1v_2], [v_1v_3], [v_2v_3]\}$. Then, the only non-trivial relative chain space is $C_2(\Delta^2, \partial\Delta^2) \cong \mathbb{R}$. Thus, $H_2(\Delta^2, \partial\Delta^2) \cong \mathbb{R}$ whereas all other relative homology is trivial: $H_i(\Delta^2, \partial\Delta^2) = 0$ for $i \neq 2$.

By introducing an inner product structure on the chain spaces of a finite simplicial complex, the vector spaces $C_i(K)$ are, actually, Hilbert spaces. A common inner product structure for $C_i(K)$ is given by the standard inner product in \mathbb{R}^n . In particular, define $\langle \cdot, \cdot \rangle_{C_i(K)}$ by declaring that the set

of i -simplices $K^{[i]}$ forms an orthonormal basis. That is, for all $\sigma, \sigma' \in K^{[i]}$, we define

$$\langle \sigma, \sigma' \rangle_{C_i(K)} = \begin{cases} 1 & \text{if } \sigma = \sigma' \\ 0 & \text{else} \end{cases}$$

and extended linearly. Thus, $C_i(K)$ is an inner product space and so are the cycle and boundary subspaces, $Z_i(K)$ and $B_i(K)$. Since each space is finitely generated, the boundary maps ∂_i are bounded. However, it is not immediately obvious how to induce an inner product on the homology spaces $H_i(K)$. Recall that, instead, we can identify $H_i(K)$ with an intersection of two inner product spaces.

Proposition 2.2.9: Given a finite simplicial complex K and $i \geq 0$, $H_i(K)$ is isomorphic to the subspace of $C_i(K)$ given by $Z_i(K) \cap B_i(K)^\perp$.

Proof. Use Lemma 2.1.8 to obtain an isomorphism $H_i(K) = Z_i(K)/B_i(K) \cong B_i(K)^\perp$. To highlight that this is the space of *cycles* that are orthogonal to all boundaries, rather than *chains* that are orthogonal to all boundaries, we write $H_i(K) \cong Z_i(K) \cap B_i(K)^\perp \subseteq C_i(K)$. \square

Simplicial homology is an algebraic method that describes topological properties of the associated simplicial complex K . By introducing an inner product structure on each chain space $C_i(K)$, we can study the simplicial complex K *analytically*. We do so by investigating combinatorial analogs of the Hodge Laplacian.

2.3 Combinatorial Hodge Laplacians

By placing an inner product on simplicial chain spaces, we can study analytic and geometric properties of the underlying simplicial complexes. We will do so by introducing *combinatorial* Hodge Laplacians, slowly increasing in generality. We will begin by introducing the graph Laplacian, then we will extend the Laplacian to simplicial complexes and simplicial cosheaves on them.

2.3.1 Graph Laplacians

The graph Laplacian was first introduced by Kirchoff in his study of electrical circuits, motivated by calculating quantities such as current flow and effective resistance between nodes [12]. The graph Laplacian is defined as follows. Let $G = (V, E)$ be a finite, simple, undirected graph. That is, the graph G consists of a finite vertex set $V = \{v_i\}_{i \in \mathcal{I}}$ and edge set E that consists of unordered pairs of vertices $\{v_i, v_j\}$ with $i \neq j$. We construct the following matrices to analyze G .

Define the *adjacency matrix* $A = (a_{ij})$ by

$$a_{ij} = \begin{cases} 1, & \text{if } \{v_i, v_j\} \in E, \\ 0, & \text{otherwise.} \end{cases}$$

Observe that, because we are considering simple graphs with no vertices of the form $\{v_i, v_i\}$, all diagonal entries of A are equal to 0.

The *degree matrix* is the diagonal matrix D where $D(i, i)$ is the degree of vertex i . That is, the entry $D(i, i)$ counts the number of edges that contain v_i .

Definition 2.3.1: The **graph Laplacian** is the $|V| \times |V|$ matrix $L = D - A$. When the graph G needs emphasis or clarification, we will write L^G .

Observe that, since we are considering undirected graphs, the graph Laplacian is *symmetric*.

Intuitively, think of the graph Laplacian as follows. The i th entry of a vector $x \in \mathbb{R}^{|V|}$ assigns a real number to vertex v_i . When applying the Laplacian to x , entry i of Lx describes how much the value assigned to v_i differs from the average of the values assigned to each of v_i 's neighbors. That is, much like the continuous setting, the graph Laplacian measures "local variation."

Alternatively, we may define the graph Laplacian by viewing a graph as a one-dimensional simplicial complex and using the boundary matrices defined in Section 2.2. That is, we can compute the graph Laplacian as follows.

Proposition 2.3.2 ([26, Lemma 1]): $L = \partial_1 \partial_1^T$.

Proof. First, focus on the diagonal entries. Entry $\partial_1 \partial_1^T(i, i)$ is given by multiplying row i of ∂_1 by column i of ∂_1^T . That is, we get a factor of $\pm 1 \times \pm 1$ (signs matching) for each edge that contains the vertex v_i . That is, $\partial_1 \partial_1^T(i, i) = D(i, i) = L(i, i)$.

Next, focus on the entries $\partial_1 \partial_1^T(i, j)$ for $i \neq j$. Since we have a simple graph, row i of ∂_1 and column j of ∂_1^T will both have a non-zero entry in precisely one entry if and only if $\{v_i, v_j\} \in E$, and in this case, $\partial_1 \partial_1^T(i, j) = -1 = -A(i, j) = L(i, j)$. \square

In Proposition 2.3.2, we are implicitly assuming that the chain spaces $C_1(G)$ and $C_0(G)$ are endowed with inner product structures so that the edges and vertices form orthonormal bases. By considering a different inner product structure on chain spaces, we obtain a "new" graph Laplacian defined by $\partial_1 \partial_1^*$.

Kirchhoff's original paper, along with the applications to electrical engineering, includes interesting theoretical results about the graph Laplacian and the connection between its spectrum and combinatorial properties. For example, Kirchhoff proves the matrix tree theorem, which we will state later. However, it took more than 100 years for the graph Laplacian to become a popular object of study: Fiedler is often credited for popularizing the graph Laplacian with his study on the algebraic connectivity of graphs in 1973 [27]. Indeed, since the mid-twentieth century, extensive research has focused on understanding the combinatorial and geometric properties of the graph Laplacian and its spectrum. For one, Fisher studied the *normalized* graph Laplacian (as a discrete analogue to the Laplace-Beltrami operator) and its connection to random walks on lattices in [28]. Moreover, Anderson and Morley prove important results about the spectrum of the graph Laplacian in [26]. For example, [26, Theorem 1] states that the eigenvalues of a graph with n vertices will always lie in the (closed) interval $[0, n]$: in doing so, they show that the graph Laplacian is positive-semidefinite. The proof of their result utilizes the following proposition, which will play a central role in later chapters of this dissertation.

Proposition 2.3.3: Let K_n be the complete graph on n vertices. Then, the eigenvalues of L^{K_n} are 0 with multiplicity 1 and n with multiplicity $n - 1$.




Figure 2.1: The complete graph on three vertices, K_3 .

Proof. Since each vertex in K_n has degree $n - 1$, the adjacency matrix for K_n satisfies $A = J - I_n$, where J is the all-ones matrix and I_n is the $n \times n$ identity matrix. Thus,

$$L^{K_n} = D - A = (n - 1)I_n - (J - I_n) = nI_n - J.$$

Let $\vec{1}$ denote the all-ones vector. Then, J satisfies $J\vec{1} = n\vec{1}$ and $L^{K_n}\vec{1} = \vec{0}$. Moreover, $Jx = \vec{0}$ for all x orthogonal to $\vec{1}$. Thus, $L^{K_n}x = nx$. Therefore, the eigenvalues of L^{K_n} are 0 with multiplicity 1 and n with multiplicity $n - 1$. \square

Remark 2.3.4: In the proof of Proposition 2.3.3, we showed that the eigenvectors x of L^{K_n} with non-zero eigenvalue are orthogonal to $\vec{1}$. Indeed, since the Laplacian is self-adjoint, eigenvectors with distinct eigenvalues must be orthogonal. In this case, $0 = \langle x, \vec{1} \rangle = \sum_{i=1}^n x_i$. That is, the entries of any eigenvector x for K^{K_n} with eigenvalue n must sum to 0.

Example 2.3.5: Consider the complete graph on three vertices, K_3 (see Figure 2.1). The graph Laplacian L^{K_3} is given by

$$L^{K_3} = \begin{pmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix} - \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}.$$

The eigenvalues of L^{K_3} are $\{0, 3, 3\}$ as expected, and their respective eigenvectors are

$$\begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}, \begin{pmatrix} 1 \\ -1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}.$$

Moreover, in [26, Section 5], the authors fully characterize the eigenvalues for other nice classes of graphs (i.e., complete bipartite graphs, cycles, paths, and wheels). These results demonstrate the ability of the spectrum of the graph Laplacian to reflect the combinatorial structure of these objects. We will see throughout that this trend generalizes.

While investigating the eigenvalues of the graph Laplacian L , it is natural to investigate the *characteristic polynomial* of L : $\det(\lambda I_{|V|} - L) = \sum_{i=0}^{|V|} (-1)^i q_i \lambda^{|V|-i}$. Indeed, recall that the eigenvalues of a square matrix are the zeroes of its characteristic polynomial. For the graph Laplacian, it is known that the characteristic polynomial also reflects the combinatorial structure of the underlying graph G . In particular, the following result of Biggs shows that the characteristic polynomial is a generating function for *forests* of G .

Definition 2.3.6: A **forest** of a graph G is a subgraph $F = (V(F), E(F))$ such that each pair of vertices in F is connected by at most one path. Equivalently, a forest is an acyclic subgraph of G . A **spanning tree** of G is a subgraph that contains all vertices of G and each pair of vertices is connected by exactly one path. Equivalently, a spanning tree is a connected, acyclic subgraph on all vertices of G .

The coefficients of the characteristic polynomial are a generating function for the number of vertices in forests with i many edges. This statement is made precise in the following proposition.

Proposition 2.3.7 ([29, Theorem 7.5]): The coefficient q_i of $\det(\lambda I_{|V|} - L)$ is given by

$$q_i = \sum_{F, |E(F)|=i} |V(F)|.$$

That is, the coefficient q_i is the sum of the number of vertices in each forest with i many edges.

In particular, Kirchhoff's matrix tree theorem states that, for a connected graph, the number of spanning trees of G is given by $\frac{1}{|V|}q_{|V|-1}$ (the product of the non-zero eigenvalues of L divided by the number of vertices in G). Furthermore, the trace of L is given by $q_1 = 2|E|$. Finally, observe that $q_{|V|} = \det(L) = 0$ since 0 is always an eigenvalue (the row sums are 0).

Example 2.3.8: Again, consider the complete graph on three vertices, K_3 . The characteristic polynomial of L^{K_3} is $\det(\lambda I_3 - L^{K_3}) = \lambda^3 - 6\lambda^2 + 9\lambda$. Indeed, $q_1 = 6 = 0 + 3 + 3 = 2 \cdot |E(K_3)|$. Further, q_2 is the product of the non-zero eigenvalues of L^{K_3} : $9 = 3 \cdot 3$. By Kirchhoff's matrix tree theorem, there are $9/3 = 3$ spanning trees for K_3 . This is easy to verify.

In the next section, we extend this construction of a Laplacian matrix from graphs to simplicial complexes by using boundary maps to encode higher-dimensional adjacency.

2.3.2 Simplicial Hodge Laplacians and Discrete Hodge Theory

By writing the graph Laplacian as a composition of the transpose of a boundary operator on a one-dimensional simplicial complex with the boundary operator itself, it is natural to ask whether there are generalizations. Eckmann was the first to generalize Kirchhoff's graph Laplacian to a combinatorial Laplacian on simplicial complexes [11]. We will review his constructions here.

Definition 2.3.9: Let K be a (finite) simplicial complex, and endow each chain space $C_i(K)$ with an inner product. The i -th **combinatorial Laplacian** is defined as:

$$L_i = \partial_{i+1} \partial_{i+1}^* + \partial_i^* \partial_i : C_i(K) \rightarrow C_i(K).$$

It is common to decompose L_i into an up part and a down part. The **up Laplacian** is $L_i^{\text{up}} = \partial_{i+1} \partial_{i+1}^*$ and the **down Laplacian** is $L_i^{\text{down}} = \partial_i^* \partial_i$. Observe that $L_i = L_i^{\text{up}} + L_i^{\text{down}}$.

Furthermore, if we equip each chain space $C_i(K)$ with the inner product for which the i -simplices form an orthonormal basis, then adjoints are given by transposes. That is, the combina-

torial Laplacian is written as $L_i = \partial_{i+1}\partial_{i+1}^T + \partial_i^T\partial_i$. In this case, when $i = 0$, L_0 coincides with the graph Laplacian as it sees only the 1-skeleton of K . For $i > 0$, the operator L_i detects higher-dimensional structure by incorporating how i -simplices meet along $(i - 1)$ -faces and $(i + 1)$ -faces.

Example 2.3.10: Consider the standard 2-simplex, Δ^2 (See Figure 1.1). The zero Laplacian L_0 will be the same as L^{K_3} : see Example 2.3.5. However, the two-simplex in Δ^2 provides extra structure. To compute L_1 , first compute the boundary maps:

$$\partial_1^{\Delta^2} = \begin{pmatrix} -1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & 1 & 1 \end{pmatrix}, \quad \partial_2^{\Delta^2} = \begin{pmatrix} 1 \\ -1 \\ 1 \end{pmatrix}.$$

Then, the Laplacian L_1 is:

$$L_1 = \begin{pmatrix} 3 & 0 & 0 \\ 0 & 3 & 0 \\ 0 & 0 & 3 \end{pmatrix}.$$

Because L_1 is diagonal, we can easily read off the eigenvectors and eigenvalues.

Not only did Eckmann generalize Kirchhoff's work to combinatorial Laplacians, but he also defined *harmonic* homology spaces associated to a simplicial complex.

Definition 2.3.11: For K a finite simplicial complex and $i \geq 0$, the i -th **harmonic homology space** is defined as $\mathcal{H}_i(K) = \ker L_i$. Elements of $\mathcal{H}_i(K)$ are called **harmonic i -cycles**.

Indeed, *harmonic* homology spaces are isomorphic to (the more familiar) simplicial homology spaces, and they coincide with our method for identifying homology with an inner product space.

Proposition 2.3.12: $\mathcal{H}_i(K) = Z_i(K) \cap B_i(K)^\perp$.

Proof. We begin by showing that $\mathcal{H}_i(K) = \ker L_i = \ker \partial_i \cap \ker \partial_{i+1}^*$. Let $z \in \ker \partial_i \cap \ker \partial_{i+1}^*$.

Then,

$$L_i(z) = \partial_{i+1} \circ \partial_{i+1}^*(z) + \partial_i^* \circ \partial_i(z) = 0,$$

and thus $z \in \ker L_i$. That is, $\ker \partial_i \cap \ker \partial_{i+1}^* \subseteq \ker L_i$.

Now, take $z \in \ker L_i$. Then, $\partial_{i+1} \circ \partial_{i+1}^*(z) + \partial_i^* \circ \partial_i(z) = 0$ and

$$\begin{aligned} \langle z, \partial_{i+1} \circ \partial_{i+1}^*(z) \rangle + \langle z, \partial_i^* \circ \partial_i(z) \rangle &= 0 \\ \implies \langle \partial_{i+1}^*(z), \partial_{i+1}^*(z) \rangle + \langle \partial_i(z), \partial_i(z) \rangle &= 0. \end{aligned}$$

However, this is only possible if $z \in \ker \partial_{i+1}^* \cap \ker \partial_i$. Thus, $\ker L_i \subseteq \ker \partial_{i+1}^* \cap \ker \partial_i$ and we have shown that $\mathcal{H}_i(K) = \ker L_i = \ker \partial_{i+1}^* \cap \ker \partial_i$.

Finally, we argue that $\ker \partial_{i+1}^* \cap \ker \partial_i = Z_i(K) \cap B_i(K)^\perp$. By definition, $Z_i(K) = \ker \partial_i$. Furthermore, by Proposition 2.1.10, $\ker \partial_{i+1}^* = (\text{im } \partial_{i+1})^\perp = B_i(K)^\perp$. That is, we have shown the following equalities:

$$\mathcal{H}_i(K) = \ker \partial_{i+1}^* \cap \ker \partial_i = Z_i(K) \cap B_i(K)^\perp.$$

□

Corollary 2.3.13: $\mathcal{H}_i(K) \cong H_i(K)$.

Proof. This follows immediately from Propositions 2.2.9 and 2.3.12. □

Much like the continuous case, elements of $\mathcal{H}_i(K)$ provide "nice" cycle representatives for homology classes in $H_i(K)$. A harmonic chain is, intuitively, a cycle of minimal energy: $\mathcal{H}_i(K)$ contains canonical representatives of homology classes that minimize the norm induced by the inner product on $C_i(K)$. Although we will defer the proof of an energy minimizing property for a later chapter, the *discrete Hodge decomposition* provides some insight [11]. The discrete Hodge decomposition states that, for $i \geq 0$ and K a finite simplicial complex, we have the following orthogonal decomposition of $C_i(K)$:

$$C_i(K) = \text{im } \partial_{i+1} \oplus \mathcal{H}_i(K) \oplus \text{im } \partial_i^*.$$

The elements of $\mathcal{H}_i(K)$, then, are orthogonal to *both* boundaries and coboundaries.

As with the graph Laplacian, the spectrum of L_i often reflects the combinatorial structure of the associated simplicial complex. Duval and Reiner, for one, show that a certain class of simplicial complexes called *shifted* simplicial complexes always have integer eigenvalues that arise from an associated integer partition [30]. Furthermore, Horak and Jost explicitly compute an upper bound on the spectrum of L_i and, in the case of *normalized* combinatorial Laplacians, the authors provide a full characterization of the spectrum of L_i for certain types of simplicial complexes [13]. Steenbergen dives further into the spectrum of combinatorial Laplacians in his doctoral dissertation [31], and in [32, 33] the characteristic polynomial of the combinatorial Laplacian is related to higher-dimensional forests and discrete Morse theory.

In all, there is a wealth of research demonstrating the connections between the spectrum of the combinatorial Laplacian and the underlying geometry and combinatorics of the associated simplicial complex. However, as we endow our simplicial complex with the data of vector spaces over each simplex (that is, if we consider a simplicial cosheaf), the story becomes more complicated. In the following section, we will explore existing work in the area of simplicial cosheaf Laplacians.

2.3.3 Simplicial Cosheaf Laplacians

Hansen and Ghrist generalize the simplicial Hodge Laplacian even further in their work on cellular (co)sheaf Laplacians [34]. Although the work of Hansen and Ghrist applies for more general (co)sheaves on *cell complexes*, we will focus on the smaller sub-class of (co)sheaves on *simplicial complexes*. See [35, Chapter 4] for more on cellular (co)sheaves, and see [36] for a more complete treatment of Laplacians for cellular (co)sheaves and their applications.

Recall Definition 1.1.12: a simplicial cosheaf is a functor $\mathcal{F} : K \rightarrow \text{Vec}_{\mathbb{R}}$. Intuitively, simplicial cosheaves provide a way to assign local data to the cells of a simplicial complex K . This assignment is compatible with how higher-dimensional simplices meet their faces. A cosheaf propagates data *downward* in dimension: that is, data on a simplex is mapped down to its faces.

The inherent structure of simplicial cosheaves allows for the construction of cosheaf-valued chain spaces and, thus, cosheaf homology.

Definition 2.3.14: Given a simplicial cosheaf \mathcal{F} , define the i -th chain space as the direct sum:

$$C_i(K; \mathcal{F}) = \bigoplus_{\sigma \in K^{[i]}} \mathcal{F}(\sigma).$$

Define the i -th boundary map as the direct sum:

$$\partial_i = \bigoplus_{\sigma \in K^{[i]}} \left(\sum_{\sigma < \tau} [\tau : \sigma] \mathcal{F}(\sigma \rightarrow \tau) \right),$$

where $[\tau : \sigma] = (-1)^j$ for $\sigma = [v_0 v_1 \cdots v_i]$ and $\tau = [v_0 \cdots \hat{v}_j \cdots v_i]$.

Combining the data of chain spaces and boundary maps, we obtain a chain complex $C_\bullet(K; \mathcal{F})$:

$$\cdots \xrightarrow{\partial_{i+1}} C_i(K; \mathcal{F}) \xrightarrow{\partial_i} C_{i-1}(K; \mathcal{F}) \xrightarrow{\partial_{i-1}} \cdots$$

When \mathcal{F} is the cosheaf that assigns \mathbb{R} to all simplices σ with $F(\sigma \rightarrow \tau) = 1$ for all morphisms (in this case we call \mathcal{F} the constant \mathbb{R} -cosheaf), the chain spaces $C_i(K, \mathcal{F})$ and boundary maps ∂_i match the *simplicial* chain spaces and boundary maps from Section 2.2.

Definition 2.3.15: Given a simplicial cosheaf \mathcal{F} , the **cosheaf homology** spaces are given by

$$H_i(K; \mathcal{F}) = \ker \partial_i / \text{im } \partial_{i+1}.$$

Cosheaf homology often provides useful information about the cosheaf itself: for one, zeroth cosheaf homology $H_0(K; \mathcal{F})$ gives the space of global cosections of our cosheaf. Again, when \mathcal{F} is the constant \mathbb{R} -cosheaf, simplicial cosheaf homology is the same as ordinary simplicial homology. Furthermore, just as we defined the Euler characteristic of a simplicial complex, we can also define the Euler characteristic of a simplicial cosheaf using cosheaf homology spaces.

Definition 2.3.16: Given a simplicial cosheaf \mathcal{F} , its **Euler characteristic** is given by

$$\chi(\mathcal{F}) = \sum_{i \geq 0} (-1)^i \dim H_i(K; \mathcal{F}).$$

Suppose, now, that each $\mathcal{F}(\sigma)$ carries the additional structure of a Hilbert space. Since we are restricting to finite simplicial complexes, all chain spaces are necessarily finitely generated and all maps $\mathcal{F}(\sigma \rightarrow \tau)$ are bounded linear maps. Thus, we will now consider cosheaves valued over the category of Hilbert spaces: $\mathcal{F} : K \rightarrow \text{Hilb}_{\mathbb{R}}$. In [34], the authors call these *weighted simplicial cosheaves*.

With an inner product on each costalk $\mathcal{F}(\sigma)$, we obtain an inner product on cosheaf chain spaces as follows. Each summand in $C_i(K; \mathcal{F}) = \bigoplus_{\sigma, \dim \sigma = i} \mathcal{F}(\sigma)$ retains its inner product structure, and all summands are mutually orthogonal. That is, $\langle \cdot, \cdot \rangle_{C_i(K; \mathcal{F})}$ is given by:

$$\langle x, y \rangle_{C_i(K; \mathcal{F})} = \begin{cases} \langle x, y \rangle_{\mathcal{F}(\sigma)} & \text{if } x, y \in \mathcal{F}(\sigma), \\ 0 & \text{if } x \in \mathcal{F}(\sigma), y \in \mathcal{F}(\sigma') \text{ for } \sigma \neq \sigma'. \end{cases}$$

Whereas cosheaf homology itself is an algebraic operation, the addition structure of inner products on each costalk $\mathcal{F}(\sigma)$ and, thus, each chain space $C_i(K; \mathcal{F})$ will allow us to investigate *analytic* properties of these cosheaves. We do so by extending the definition of the combinatorial Hodge Laplacian to simplicial cosheaves.

Definition 2.3.17 ([34]): Given a simplicial cosheaf $\mathcal{F} : K \rightarrow \text{Hilb}_{\mathbb{R}}$, the *i -th cosheaf Laplacian* is the operator on $C_i(K; \mathcal{F})$ defined as follows:

$$L_i^{\mathcal{F}} = \partial_{i+1} \partial_{i+1}^* + \partial_i^* \partial_i.$$

Although the definition above looks identical to that of the combinatorial Laplacian, it is important to remember that we have the extra information of a cosheaf \mathcal{F} over K . However, when the cosheaf structure is clear, we will simply write L_i .




Figure 2.2: Two cosheaves on K_3 . All unlabeled maps are the identity.

Example 2.3.18: Consider two cosheaves on K_3 : let $\mathcal{F} : K_3 \rightarrow \text{Hilb}_{\mathbb{R}}$ be given by $\mathcal{F}(\sigma) = \mathbb{R}$ for all simplices σ with all morphisms the identity (See Figure 2.2, left). Let each $\mathcal{F}(\sigma)$ have the standard inner product. Let $\mathcal{G} : K_3 \rightarrow \text{Hilb}_{\mathbb{R}}$ be given by $\mathcal{G}(\sigma) = \mathbb{R}$ for all simplices σ with linear maps given by $\mathcal{G}([v_1 v_3] \rightarrow v_3) = -1$ and all other maps the identity (See Figure 2.2, right). Again, endow each $\mathcal{G}(\sigma)$ with the standard inner product.

A quick computation confirms that the 0-th cosheaf Laplacian $L_0^{\mathcal{F}}$ is identical to the graph Laplacian L^{K_3} . For one, this means that the spectrum of $L_0^{\mathcal{F}}$ is $\{0, 3, 3\}$ and that $H_0(K_3; \mathcal{F}) \cong \mathbb{R}$.

On the other hand, $L_0^{\mathcal{G}}$ reflects the "twisting" that occurs from $[v_1 v_3]$ to v_3 . First, compute the chain spaces: $C_0(K_3; \mathcal{G}) = \mathbb{R}^3$ and $C_1(K_3; \mathcal{G}) = \mathbb{R}^3$. The boundary map $\partial_1^{\mathcal{G}}$ is

$$\partial_1^{\mathcal{G}} = \begin{pmatrix} -1 & -1 & 0 \\ 1 & 0 & -1 \\ 0 & -1 & 1 \end{pmatrix}$$

and, thus,

$$L_0^{\mathcal{G}} = \partial_1^{\mathcal{G}}(\partial_1^{\mathcal{G}})^T = \begin{pmatrix} 2 & -1 & 1 \\ -1 & 2 & -1 \\ 1 & -1 & 2 \end{pmatrix}, \quad L_1^{\mathcal{G}} = (\partial_1^{\mathcal{G}})^T \partial_1^{\mathcal{G}} = \begin{pmatrix} 2 & 1 & -1 \\ 1 & 2 & -1 \\ -1 & -1 & 2 \end{pmatrix}.$$

The eigenvalues of both $L_0^{\mathcal{G}}$ and $L_1^{\mathcal{G}}$ are $\{1, 1, 4\}$. Since 0 is not in the spectrum, $H_0(K_3; \mathcal{G}) = 0$.

Observe that cosheaf homology can differentiate between \mathcal{F} and \mathcal{G} whereas their Euler characteristics cannot. Indeed, $\chi(\mathcal{F}) = 1 - 1 = 0$ and $\chi(\mathcal{G}) = 0$.

In [34], Hansen and Ghrist provide a first look into the *spectral* properties of the cosheaf Laplacian $L_i^{\mathcal{F}}$. However, because a simplicial cosheaf provides so much additional information, the spectrum of $L_i^{\mathcal{F}}$ is, in general, quite mysterious. In Chapter 4, we will explore a *specific* cosheaf obtained via Möbius homology [15]. Inspired by Möbius inversion, a counting technique from algebraic combinatorics, the inherent structure of Möbius homology make it especially amenable to spectral analysis. This is especially true for modules arising from *persistent homology*. Later, in Chapters 5 and 6, we will utilize both the algebraic and analytic techniques developed in this chapter to investigate this new Laplace operator.

Chapter 3

Möbius Inversion in One-Parameter Persistence

The overarching goal of this dissertation is to apply the analytic machinery of combinatorial Hodge Laplacians from Chapter 2 to better understand the structure of *persistence modules*. Recall from Definition 1.1.11 that a persistence module is a functor $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ from a finite poset P to the category of \mathbb{R} -vector spaces. In persistence, we want to answer the following question: when do features appear (are "born") in M , and when do they get mapped to zero?

In this chapter, we will focus specifically on *one-parameter modules*. That is, we will focus on modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite totally ordered set T . Indeed, the theory and computation of persistence began in this one-parameter setting, and the structure of one-parameter modules is well understood. More specifically, the *persistence diagram* provides a summary of the birth and death of features in M , and it is known that the persistence diagram of a one-parameter module completely determines its isomorphism type.

We will provide three (equivalent) methods for computing the persistence diagram of a one-parameter module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$. We will first investigate the algebraic structure of one-parameter persistence modules using tools from representation theory. For T a finite totally ordered set, a module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ is known to decompose into a direct sum of simpler modules called "interval modules." The persistence diagram is given by the multiplicities of the interval modules in the decomposition of M .

On the other hand, Möbius inversion provides another approach for understanding the structure of persistence modules. The Möbius inversion approach to persistence dates back to original papers such as [3], but Patel was the first to make concrete connections to Rota's work on algebraic combinatorics [37]. To apply these tools to persistence modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, we will take the following steps. First, define an integer-valued function on half-open intervals in T obtained from the data of the module M . There are multiple functions to choose from, but for the purposes of this dissertation, we will focus on the rank and the birth-death functions. Then, we apply Möbius

inversion to track the *changes* in this function: this gives the persistence diagram. Despite being indexed over distinct posets, applying Möbius inversion to the rank and the birth-death functions yields equivalent persistence diagrams.

We begin this chapter with a general treatment of Möbius inversion as it plays a central role in the theory and computation of persistence. Then, we focus on the three existing methods for computing the persistence diagram for one-parameter modules: interval decompositions, Möbius inversion of the rank function, and Möbius inversion of the birth-death function. Indeed, all three methods will provide equivalent descriptions of a one-parameter module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$.

3.1 Möbius Inversion on Posets

We begin with a general treatment of Möbius inversion as it is the foundation for the remainder of this chapter. Let P be a poset.

Definition 3.1.1: A closed **interval** in P is $[a, b] = \{c \in P \mid a \leq c \leq b\}$. Denote by $I(P)$ the **set of closed intervals** in P .

We say P is *locally finite* if all intervals $[a, b]$ are finite. Given an integer-valued function $f : P \rightarrow \mathbb{Z}$ on a locally finite poset P , we define its *Möbius inverse* as follows.

Definition 3.1.2: Let P be a locally finite poset, and let $f : P \rightarrow \mathbb{Z}$ be an integer-valued function on P . Then, the **Möbius inverse** (or, inversion) of f is defined as the unique function $\partial f : P \rightarrow \mathbb{Z}$ satisfying

$$f(b) = \sum_{a \leq b} \partial f(a)$$

for all $b \in P$.

The Möbius inverse of f acts as a sort of "combinatorial derivative." Indeed, the formula above looks like a discrete analog of the fundamental theorem of calculus.

Möbius inversion is a common technique used in algebraic combinatorics: we may learn about a "harder" function by expressing it as the Möbius inverse of an "easier" function. It originated in

number theory, specifically considering functions over the poset of divisors of the natural numbers. Below, we provide one of the original examples of Möbius inversion.

Example 3.1.3: Consider the (locally finite) poset of natural numbers \mathbb{N} ordered by divisibility. That is, we say $d \leq n$ if $d|n$. Given an integer n , we wish to count the number of natural numbers coprime to n . This is Euler's totient function, and it is denoted $\Phi : \mathbb{N} \rightarrow \mathbb{Z}$. This is the "harder" function in this example. We express Φ as the Möbius inverse of the "easier" function $f : \mathbb{N} \rightarrow \mathbb{Z}$ given by the identity map, $f(n) = n$. That is, $\Phi = \partial f$, resulting in the equation $n = \sum_{d|n} \Phi(d)$.

Hall and Weisner independently extended Möbius inversion beyond the original setting of arithmetic functions, such as Euler's totient function, to the more general setting of functions over posets [38, 39]. Later, Gian-Carlo Rota systematized and unified the theory of Möbius inversion by relating it to the *incidence algebra* of a poset P [14].

Definition 3.1.4: Let P be a locally finite poset, and let $I(P)$ denote the set of intervals $[a, b]$ in P . The **incidence algebra** of P is the algebra of functions $\phi : I(P) \rightarrow \mathbb{Z}$. Addition is given by point-wise addition, and multiplication is given by *convolution*. Given two functions $\phi, \psi : I(P) \rightarrow \mathbb{Z}$, their convolution is given by the formula $\phi * \psi[a, b] = \sum_{c \in [a, b]} \phi[a, c] \psi[c, b]$.

In the incidence algebra, the unit element is the delta function $\delta : I(P) \rightarrow \mathbb{Z}$ defined as follows:

$$\delta[a, b] = \begin{cases} 1 & \text{if } a = b, \\ 0 & \text{else.} \end{cases}$$

The zeta function $\zeta : I(P) \rightarrow \mathbb{Z}$ is the element in the incidence algebra of P so that $\zeta[a, b] = 1$ for all $[a, b] \in I(P)$. Although not all elements in the incidence algebra are invertible, ζ is.

Definition 3.1.5: The inverse of ζ is the **Möbius function** $\mu : I(P) \rightarrow \mathbb{Z}$. The Möbius function is computed as follows:

$$\mu[a, b] = \begin{cases} 1 & \text{if } a = b, \\ -\sum_{a \leq c < b} \mu[a, c], & \text{if } a < b. \end{cases}$$

Indeed, the Möbius function is deeply connected to Möbius inversion.

Proposition 3.1.6 ([14, Corollary 1]): Let P be a locally finite poset, and let $f, g : P \rightarrow \mathbb{Z}$ be functions so that

$$f(b) = \sum_{a \leq b} g(a).$$

That is, $g = \partial f$. Then,

$$\partial f(b) = g(b) = \sum_{a \leq b} f(a) \mu[a, b].$$

Let's apply the proposition above to the example of Euler's totient function.

Example 3.1.7: Again, consider the poset of natural numbers \mathbb{N} ordered by divisibility. The Möbius function on this poset is given by:

$$\mu[d, n] = \begin{cases} 1 & \text{if } d = n, \\ (-1)^t & \text{if } \frac{n}{d} \text{ is a product of } t \text{ distinct primes,} \\ 0 & \text{if } p^2 \mid \frac{n}{d} \text{ for some prime } p. \end{cases}$$

By Proposition 3.1.6,

$$\Phi(n) = \sum_{d|n} d \cdot \mu[d, n].$$

The theory of Möbius functions, then, can provide useful insights to various counting problems. Indeed, Hall and Weisner were both interested in Möbius inversion as a technique for problems in group theory. Other combinatorial concepts such as the principle of inclusion-exclusion and the chromatic number, for example, can be expressed in the language of Möbius inversion. As we will discuss in Section 3.2, Möbius inversion is also central to the theory and computation of *persistence*.

3.2 One-Parameter Persistence

In this dissertation, we are particularly interested in one example of Möbius inversion: persistence. In particular, we are interested in *one-parameter persistence*. That is, we will investigate modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite totally ordered set T . We want to study the structure of the module M by answering the following question: when do features first appear in M , and when do they get mapped to zero?

Observe that features may persist indefinitely rather than disappearing at a finite parameter value. In order to capture such features alongside those that die at a finite index, it is standard to adjoin a formal top element to the finite totally ordered set T representing "infinity."

Definition 3.2.1: Let T be a finite totally ordered set. We define

$$\bar{T} = T \cup \{\infty\}$$

to be the poset obtained by adjoining a top element ∞ with the convention that $a < \infty$ for all elements $a \in T$.

Thus, we will use elements of T to identify the birth of a feature in $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, and we will use elements in \bar{T} to identify its death. We represent the birth and the death of a feature together using *symbolic persistence intervals*.

Definition 3.2.2: Let T be a finite totally ordered set and let $\bar{T} = T \cup \{\infty\}$. A half-open **persistence interval** is a formal symbol $[a, b)$ with $a \in T$ and $b \in \bar{T}$ such that $a \leq b$. Denote by $\text{Dgm}(T)$ the set of all persistence intervals. The subset $\{[a, a) \mid a \in T\}$ is called the **diagonal**.

We emphasize that $[a, b)$ is not interpreted as a subset of T , but rather as a symbolic object recording a birth index a and a death index b .

The persistence diagram of $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ assigns to each formal persistence interval $[a, b)$ the number of features that appear at a and become zero at b . We will compute the persistence diagram of a one-parameter module using three methods: interval decomposition, Möbius inversion of

the rank function, and Möbius inversion of the birth-death function. All three methods result in equivalent diagrams that completely characterize the isomorphism type of one-parameter modules.

3.2.1 Interval Decomposition of One-Parameter Modules

We first define the persistence diagram using the interval decomposition. Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a one-parameter module. It is known that, when each $c \in T$ is assigned a finitely generated vector space $M(c)$, the one-parameter module M decomposes into certain "nice" pieces. These pieces are called interval modules.

Definition 3.2.3: Let T be a finite totally ordered set, let $\bar{T} = T \cup \{\infty\}$, and let $[a, b) \in \text{Dgm}(T)$ with $a < b$. The interval module $I_{[a,b)} : T \rightarrow \text{Vec}_{\mathbb{R}}$ is defined as follows:

$$I_{[a,b)}(c) = \begin{cases} \mathbb{R}, & \text{if } a \leq c < b \\ 0, & \text{otherwise.} \end{cases}$$

For $c \leq d$ in T , the morphism

$$I_{[a,b)}(c \leq d) : I_{[a,b)}(c) \rightarrow I_{[a,b)}(d)$$

is the identity whenever $I_{[a,b)}(c) = I_{[a,b)}(d) = \mathbb{R}$, and is the zero map otherwise.

The following proposition utilizes results from representation theory and the structural theorem of finitely generated modules over principal ideal domains. We state it using the notation from persistent homology originating with Zomorodian and Carlsson's work.

Proposition 3.2.4 ([9, Section 3.2]): Let T be a finite totally ordered set and $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module such that each $M(c)$ is finitely generated. Then, there exists a function

$$m_B : \text{Dgm}(T) \rightarrow \mathbb{Z}_{\geq 0}$$

such that

$$M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}.$$

The function $m_{\mathcal{B}}$ (equivalently, the associated multiset of intervals) is uniquely determined by M .

We call $\bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}$ the **interval decomposition of M** .

In particular, because the function $m_{\mathcal{B}}$ is uniquely determined by M , the interval decomposition of a one-parameter module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ determines its isomorphism type.

Remark 3.2.5: Although the decomposition of a module M given in Proposition 3.2.4 is unique, the *isomorphism* $\varphi : M \Rightarrow \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}$ is not.

Definition 3.2.6: Let T be a finite totally ordered set and let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module. The **persistence diagram**¹ of M is the finite-support function

$$m_{\mathcal{B}} : \text{Dgm}(T) \rightarrow \mathbb{Z}_{\geq 0}$$

such that

$$M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}$$

as in Proposition 3.2.4. The value $m_{\mathcal{B}}([a, b])$ is called the **multiplicity** of the interval $[a, b]$ in the persistence diagram. We call an interval $[a, b]$ a **point** in the persistence diagram if $m_{\mathcal{B}}([a, b]) \neq 0$.

Remark 3.2.7: Observe that we excluded intervals of the form $[a, a)$ from the definition of an interval module (Definition 3.2.3). Thus, diagonal intervals $[a, a)$ necessarily have multiplicity zero in the persistence diagram for any module M .

¹In the literature, it is common to call $m_{\mathcal{B}}$ the "barcode" of the module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, with individual persistence intervals $[a, b)$ called "bars" if $m_{\mathcal{B}}([a, b]) \neq 0$. Motivated by the equivalence of methods presented in Section 3.2.4, we use the terminology "persistence diagram" to refer to all three construction presented in this section.




Figure 3.1: The persistence diagram $m_{\mathcal{B}}$ for a module with interval decomposition $I_{[1,3)} \oplus I_{[2,4)}$. We represent intervals $[a, b)$ such that $m_{\mathcal{B}}([a, b)) \neq 0$ as lines from a to b .

Example 3.2.8: Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be the module over the totally ordered set $1 < 2 < 3 < 4$ given by:

$$\mathbb{R} \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 0 & 1 \end{pmatrix}} \mathbb{R} \xrightarrow{\begin{pmatrix} 0 \end{pmatrix}} 0.$$

The interval decomposition of M is given by $M \cong I_{[1,3)} \oplus I_{[2,4)}$. See Figure 3.1 for a visualization of the persistence diagram.

In all, representation theory provides a method for understanding the algebraic structure of a persistence module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ for T a finite totally ordered set. Indeed, the persistence diagram $m_{\mathcal{B}}$ gives the isomorphism type of M . In the following sections, we provide two other methods for obtaining the persistence diagram of a one-parameter module M using *Möbius inversion*. Indeed, as we will see in Section 3.2.4, the three methods are equivalent.

3.2.2 Rank Function and Möbius inversion

Again, let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a module over a (finite) totally ordered set T . We will use Möbius inversion to describe the structure of M . However, recall that Möbius inversion is defined for *functions* on finite posets. Thus, in this section, we will introduce the *rank function*, an integer-valued function on the *poset of persistence intervals* of T .

Definition 3.2.9: Given a finite totally ordered set T , let $\bar{T} = T \cup \{\infty\}$, and let $\text{Dgm}(T)$ denote the set of symbolic persistence intervals $[a, b)$ where $a \in T$, $b \in \bar{T}$, and $a \leq b$. We define a partial order on $\text{Dgm}(T)$, denoted by \supseteq , by declaring

$$[a, b) \supseteq [c, d) \quad \text{if and only if} \quad a \leq c \text{ and } d \leq b.$$

We refer to \supseteq as the **reverse containment order**.

Remark 3.2.10: We emphasize that elements of $\text{Dgm}(T)$ are formal symbols and are *not* interpreted as subsets of T . Thus, the order \supseteq is defined symbolically and is *not* set-theoretic inclusion.

With this order, $\text{Dgm}(T)$ is a finite poset. Furthermore, observe that, for each $a \in T$, the diagonal interval $[a, a)$ is a maximal element of $\text{Dgm}(T)$ with respect to \supseteq .

The rank function, then, is defined as follows.

Definition 3.2.11 ([7]): Let T be a finite totally ordered set, let $\bar{T} = T \cup \{\infty\}$, and consider a one-parameter module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$. Define the **rank function**

$$\text{rk}_M : (\text{Dgm}(T), \supseteq) \rightarrow \mathbb{Z}_{\geq 0}$$

as follows. For $[a, b) \in \text{Dgm}(T)$, set

$$\text{rk}_M([a, b)) = \text{rank}(M(a \leq b^-)) = \dim(\text{im}(M(a \leq b^-))),$$

where $b^- \in T$ denotes the greatest element in T that is strictly less than b . In particular, we define

$$\text{rk}_M([a, \infty)) = \text{rank}(M(a \leq \max T)) \quad \text{and} \quad \text{rk}_M([a, a)) = \text{rank}(M(a \leq a)) = \dim M(a).$$

Intuitively, the rank function counts how many features exist (or, are "born") by the index a that are not yet zero (or, have not yet "died") before the index b . The value $\text{rk}_M([a, \infty))$, in particular, records the rank of the map from $M(a)$ to the final vector space $M(\max T)$, and thus counts features present at a that persist through the end of the module. However, we want to know precisely when these features are born and when they die. In order to do so, we apply Möbius inversion.

Definition 3.2.12: Let T be a finite totally ordered set, let $\bar{T} = T \cup \{\infty\}$, and let $(\text{Dgm}(T), \supseteq)$ be the poset of symbolic persistence intervals with the reverse containment order. Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$

be a persistence module, and let $\text{rk}_M : (\text{Dgm}, \supseteq) \rightarrow \mathbb{Z}_{\geq 0}$ denote its rank function. We define the **persistence diagram** of M as the Möbius inverse of rk_M , denoted

$$\partial \text{rk}_M : (\text{Dgm}(T), \supseteq) \rightarrow \mathbb{Z}.$$

That is, ∂rk_M is the unique function satisfying

$$\text{rk}_M([a, b]) = \sum_{[c, d] \supseteq [a, b]} \partial \text{rk}_M([c, d]).$$

Because T is a totally ordered set, we are able to provide explicit formulas for ∂rk_M . Let $a \in T$ and $b \in \bar{T}$ with $a < b$. Let a^- denote the greatest element in T that is strictly less than a , when it exists. Similarly, let b^+ denote the smallest element in \bar{T} strictly greater than b , when it exists. In particular, if $b = \max T$, then $b^+ = \infty$. Then, for an interval $[a, b] \in \text{Dgm}(T)$ with $b \in T$,

$$\partial \text{rk}_M([a, b]) = \text{rk}_M([a, b]) - \text{rk}_M([a^-, b]) - \text{rk}_M([a, b^+]) + \text{rk}_M([a^-, b^+]),$$

where any term involving a non-existent a^- is omitted. On the other hand, for intervals $[a, \infty)$,

$$\partial \text{rk}_M([a, \infty)) = \text{rk}_M([a, \infty)) - \text{rk}_M([a^-, \infty)).$$

Remark 3.2.13: Unlike the interval decomposition approach for obtaining the persistence diagram, the function ∂rk_M may take negative values on diagonal intervals $[a, a)$. Indeed, we see an example of this in Figure 3.2. Furthermore, when $\partial \text{rk}_M([a, \infty)) \neq 0$ for some $a \in T$, we know there is an *infinitely persistent* feature in the module M . That is, there is some feature born at the index $a \in T$ that is never mapped to zero.




Figure 3.2: From left to right: the poset $(\text{Dgm}(T), \supseteq)$ when T is given by $1 < 2 < 3 < 4$, the rank function rk_M of the module M from Example 3.2.14, and the resulting persistence diagram ∂rk_M .

Example 3.2.14: Consider the module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ over the totally ordered set $1 < 2 < 3 < 4$ from Example 3.2.8:

$$\mathbb{R} \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 0 & 1 \end{pmatrix}} \mathbb{R} \xrightarrow{\begin{pmatrix} 0 \end{pmatrix}} 0.$$

See Figure 3.2 for an illustration of the poset $(\text{Dgm}(T), \supseteq)$, the rank function rk_M , and the persistence diagram ∂rk_M .

3.2.3 Birth-Death Function and Möbius Inversion

In the previous section, we considered the persistence diagram as a function on the poset of symbolic persistence intervals with the reverse containment order, denoted $(\text{Dgm}(T), \supseteq)$. In this section, we present an alternate but equivalent description of the persistence diagram obtained by applying Möbius inversion with respect to a different partial order. The underlying indexing set remains $\text{Dgm}(T)$, but we now equip this set with a partial order induced by the product order on birth and death parameters. This new ordering allows us to separate the roles of birth and death, leading to the notion of the birth-death function.

Definition 3.2.15: Given a finite totally ordered set T , let $\bar{T} = T \cup \{\infty\}$, and let $\text{Dgm}(T)$ denote the set of symbolic persistence intervals $[a, b)$. We define a partial order \preceq on $\text{Dgm}(T)$ by declaring

$$[a, b) \preceq [c, d) \quad \text{if and only if} \quad a \leq c \text{ and } b \leq d.$$

We call \preceq the **product order** on $\text{Dgm}(T)$.

Remark 3.2.16: Again, we emphasize that the persistence intervals $[a, b)$ are formal symbols, not subsets of the totally ordered set T . Thus, the product order \preceq is defined using the birth and death indices alone, *not* using any set-theoretic ordering.

In order to compute the birth-death function of a module M , we must first compute a *free presentation* of M .

Definition 3.2.17 ([40, Definition 4.9]): A persistence module $G : T \rightarrow \text{Vec}_{\mathbb{R}}$ is **free** if it is a finite direct sum of persistence modules of the form $\mathbb{R}^{\uparrow a}$, defined by

$$\mathbb{R}^{\uparrow a}(b) = \begin{cases} \mathbb{R} & \text{if } a \leq b \\ 0 & \text{else.} \end{cases}$$

and morphisms $\mathbb{R}^{\uparrow a}(b \leq c)$ given by the identity map for $a \leq b \leq c$ and the zero map else.

Given a module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, a **free presentation** of M consists of a free module G and a surjective natural transformation $\varphi : G \Rightarrow M$.

Since we are working specifically over *finite* totally ordered sets T , all modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ admit a free presentation: see [40, Proposition 4.13].

Definition 3.2.18 ([40, Definition 5.1]): Let T be a finite totally ordered set, let $\bar{T} = T \cup \{\infty\}$, and let $(\text{Dgm}(T), \preceq)$ denote the poset of symbolic persistence intervals equipped with the product order. Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module with free presentation $\varphi : G \Rightarrow M$. The

birth-death function associated to φ is the function

$$\text{bd}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}_{\geq 0}$$

defined by

$$\text{bd}_\varphi([a, b]) = \begin{cases} \dim(G(a) \cap \ker \varphi_b), & \text{if } b \in T, \\ \dim(G(a)), & \text{if } b = \infty. \end{cases}$$

Observe that, with the product order on intervals in $\text{Dgm}(T)$, the birth-death function is a weakly increasing function.

The birth-death function has a slightly different interpretation than the rank function. It counts the number of features born by step a that have died by step b . However, we again want to compute precisely when these features are born and when they die. Much like the rank function, the persistence diagram provides this information and is given by Möbius inversion.

Definition 3.2.19: Let T be a finite totally ordered set with $\bar{T} = T \cup \{\infty\}$. Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module with free presentation $\varphi : G \Rightarrow M$. The **persistence diagram** associated to φ is the Möbius inversion of the birth-death function

$$\partial \text{bd}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}.$$

That is, it is the unique function satisfying

$$\text{bd}_\varphi[a, b] = \sum_{[c, d] \preceq [a, b]} \partial \text{bd}_\varphi[c, d].$$

Again, because T is a totally ordered set, we are able to provide an explicit formula for $\partial \text{bd}_\varphi$. Let $a \in T$ and $b \in \bar{T}$ with $a < b$. Let a^- denote the greatest element in T that is strictly less than a , when it exists. Then, for an interval $[a, b] \in \text{Dgm}(T)$ with $b \in T$,

$$\partial \text{bd}_\varphi([a, b]) = \text{bd}_\varphi([a, b]) - \text{bd}_\varphi([a^-, b]) - \text{bd}_\varphi([a, b^-]) + \text{bd}_\varphi([a^-, b^-]),$$

where any term involving a non-existent a^- is omitted. In particular, for intervals $[a, \infty)$,

$$\partial \text{bd}_\varphi([a, \infty)) = \text{bd}_\varphi([a, \infty)) - \text{bd}_\varphi([a^-, \infty)) - \text{bd}_\varphi([a, \max T)) + \text{bd}_\varphi([a^-, \max T)).$$

Remark 3.2.20: Although Möbius inversion is used to compute the persistence diagrams for both the rank and the birth-death functions, these are *different* Möbius inversions. Indeed, they act on distinct interval posets: one with the reverse containment order and the other with the product order. For example, observe that the explicit formulas for $\partial \text{rk}_M([a, b))$ and $\partial \text{bd}_\varphi([a, b))$ have different indices. Furthermore, see Figures 3.2 and 3.3 for visualizations of the two different posets. The arrows in $(\text{Dgm}(T), \supseteq)$ point down and to the right whereas the arrows in $(\text{Dgm}(T), \preceq)$ point *up* and to the right.

Remark 3.2.21: Unlike the persistence diagram obtained from the interval decomposition, the persistence diagram defined via the birth-death function may take nonzero values on the diagonal: that is, on intervals $[a, a)$. These diagonal terms reflect relations among generators in the chosen free presentation and do not correspond to persistent features. In later sections, we will show that the off-diagonal part of $\partial \text{bd}_\varphi$ agrees with the persistence diagrams obtained via the rank function and via the interval decomposition. The diagonal values, on the other hand, encode generators in the free presentation that are immediately killed by relations.

Example 3.2.22: Let $M : T \rightarrow \text{Vec}_\mathbb{R}$ be the one-parameter module from Example 3.2.8. Consider the free module $G : T \rightarrow \text{Vec}_\mathbb{R}$ given by

$$\mathbb{R} \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}} \mathbb{R}^2.$$




Figure 3.3: From left to right: the poset $(\text{Dgm}(T), \preceq)$ when T is given by $1 < 2 < 3 < 4$, the birth-death function for M and free presentation φ as in Example 3.2.22, and the persistence diagram $\partial \text{bd}_\varphi$.

There is a free presentation $\varphi : G \Rightarrow M$ with components given by

$$\varphi_1 = 1, \quad \varphi_2 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \varphi_3 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad \text{and } \varphi_4 = 0.$$

See Figure 3.3 for the birth-death function $\text{bd}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}_{\geq 0}$ and the resulting persistence diagram $\partial \text{bd}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}$.

Observe that the definition of the persistence diagram (as well as the birth-death function) given above relies on the choice of a free presentation φ . However, there are many different free presentations for one given module. As stated in the following proposition, different choices of free presentation will result in *equivalent* persistence diagrams. That is, the persistence diagrams will agree on non-diagonal elements.

Proposition 3.2.23 ([40, Corollary 5.8]): Let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module with two free presentations $\varphi : G_1 \Rightarrow M$ and $\psi : G_2 \Rightarrow M$. Then, $\partial \text{bd}_\varphi[a, b] = \partial \text{bd}_\psi[a, b]$ when $a \neq b$.

Proof. The proof relies on the *kernel function*, another integer-valued function ascribed to a persistence module. See [40, Proposition 5.7]: the result follows as a corollary. \square

3.2.4 Equivalence of Methods

In the preceding sections, we introduced three closely related constructions associated to a one-parameter persistence module. First, we described the persistence diagram via the interval decomposition of a persistence module, encoding the multiplicities of symbolic persistence intervals. Second, we defined the persistence diagram as the Möbius inverse of the rank function on the poset of symbolic persistence intervals $\text{Dgm}(T)$ equipped with the reverse containment order. Third, we introduced the birth-death function, and we defined a persistence diagram via Möbius inversion with respect to the product order on the poset of symbolic persistence intervals.

Although these constructions arise from different perspectives and use different partial orders on the set of persistence intervals, they encode the same information in the one-parameter setting. The goal of this section is to make this equivalence precise and to clarify how these definitions are related.

Proposition 3.2.24: Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a one-parameter persistence module, and let

$$\partial\text{rk}_M : (\text{Dgm}(T), \supseteq) \rightarrow \mathbb{Z} \quad \text{and} \quad \partial\text{bd}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}$$

denote the Möbius inverses of the rank and the birth-death functions, respectively. Then, for every interval $[a, b) \in \text{Dgm}(T)$ with $a < b$,

$$\partial\text{rk}_M([a, b)) = \partial\text{bd}_\varphi([a, b)).$$

Proof. Combine [8, Section 9.1] with [41, Proposition 6.3]. □

That is, the two Möbius-inversion constructions of the persistence diagram agree on all off-diagonal intervals.

Moreover, the rank function is a complete invariant on the collection of one-parameter modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, meaning it distinguishes isomorphism classes of such modules [7]. In the lan-

guage of interval decompositions, the rank function provides the same information as the interval decomposition. This is made precise in the proposition below.

Proposition 3.2.25: Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module with interval decomposition

$$M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}.$$

Then, for every interval $[a, b] \in \text{Dgm}(T)$ with $a < b$,

$$m_{\mathcal{B}}([a, b]) = \partial \text{rk}_M([a, b]).$$

Proof. See [7, Theorem 12]: to show that the rank invariant is a complete invariant in the one-parameter setting, Carlsson and Zomorodian prove a bijection between possible interval decompositions and possible rank functions. □

That is, the multiplicity of the indecomposable interval module $I_{[a,b]}$ in the interval decomposition of M is equal to the Möbius inversion of the rank function evaluated at $[a, b]$.

Combining Propositions 3.2.24 and 3.2.25, we immediately obtain the following.

Corollary 3.2.26: Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a one-parameter persistence module with interval decomposition

$$M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}.$$

Let $\varphi : G \Rightarrow M$ be a free presentation of M . Then, for every persistence interval $[a, b] \in \text{Dgm}(T)$ with $a < b$,

$$m_{\mathcal{B}}([a, b]) = \partial \text{rk}_M([a, b]) = \partial \text{bd}_{\varphi}([a, b]).$$

Consequently, all three constructions define the same persistence diagram off the diagonal.

Example 3.2.27: Consider Examples 3.2.22 and 3.2.14 along with their persistence diagrams as shown in Figures 3.2 and 3.3. Observe that $\partial \text{bd}_{\varphi}([a, b]) = \partial \text{rk}_M([a, b])$ whenever $a \neq b$. Further-

more, recall Example 3.2.8: $m_B([1, 3)) = m_B([2, 4)) = 1$, and all other intervals in $\text{Dgm}(T)$ have multiplicity zero in the persistence diagram.

In all, both the rank and the birth-death functions provide useful (and equivalent) information for describing the structure of one-parameter persistence modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$. Furthermore, the persistence diagrams of these functions are equivalent to the persistence diagram obtained via the interval decomposition of M .

Not only is the persistence diagram uniquely determined by the isomorphism class of the one-parameter module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, but "close" persistence modules will also result in "close" persistence diagrams. This is the *bottleneck stability theorem*, first proved for functions on topological spaces in [5] and extended to the case of persistence modules more generally in [6]. The bottleneck stability theorem states that, if two persistence modules are " ϵ -interleaved" (meaning they are a "distance" of ϵ away from being isomorphic), then the resulting persistence diagrams have bottleneck distance at most ϵ . The bottleneck distance, deeply related to the optimal transport distance in measure theory, quantifies the "cost" of moving points from one persistence diagram to another. Although we will not provide an in-depth treatment of stability results in persistence, such results justify the use of persistence diagrams for comparing topological spaces and persistence modules, particularly in applications.

However, when considering *multiparameter* modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ (that is, when P is not totally ordered), the results presented in this section do not necessarily hold. Although we may extend the definition(s) of a persistence diagram to such multiparameter modules, they exhibit pathological behavior that inhibits interpretability. In the following chapter, we introduce methods for studying multiparameter persistence modules. Namely, we will introduce *Möbius homology*, a categorification of Möbius inversion.

Chapter 4

Möbius Homology and Multiparameter Persistence

We now wish to extend the constructions from Chapter 3 to *multiparameter* modules. More specifically, we want to define persistence diagrams and other persistence invariants to understand the structure of modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite poset P . However, when extending the interval decomposition, rank function, and birth-death function approaches to define persistence diagrams for a multiparameter module M , pathological behavior makes the analysis of such invariants significantly more difficult.

Indecomposable modules in the multiparameter setting do, indeed, exist, but they no longer take solely the form of interval modules, like those in the one-parameter case. That is, there are more types of indecomposable modules, and they do not necessarily have an easy-to-define structure. Nonetheless, multiparameter modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ *do* admit a decomposition into indecomposable summands. However, many multiparameter modules are of "wild type," meaning their decomposition into indecomposables is quite complex and difficult to state. Furthermore, in [42], Bauer and Scoccola show that *any* multiparameter persistence module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ is arbitrarily close to an indecomposable module. Moreover, they show that small perturbations in an indecomposable module result in non-trivial changes to its decomposition. In all, a structural theorem for multiparameter modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ is considered infeasible.

Thus, we turn our attention to the approach of Möbius inversion of the rank and the birth-death functions. Unfortunately, Carlsson and Zomorodian showed in [7] that there are *no* complete discrete invariants for multiparameter persistence modules. Nonetheless, generalizations of the rank and the birth-death functions and their resulting persistence diagrams are still used to describe modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ for P a finite poset. However, the equivalence of methods that exists in the one-parameter setting no longer holds here. That is, the rank and the birth-death functions for a multiparameter module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ may yield *non-equivalent* persistence diagrams. See, for example, [40, Table 2]. A discussion of generalized rank invariants may be found in [43, 44].

However, in this dissertation, we choose to focus on the *birth-death* approach to multiparameter persistence specifically (the motivation for this choice is explained in Section 4.1). In Section 4.2, we will define the birth-death function (along with a categorification) for a module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite poset P .

Moreover, when defining the persistence diagram for a module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite poset, even the off-diagonal persistence intervals $[a, b)$ may be assigned negative values. This greatly reduces the interpretability of the persistence diagram: we may no longer interpret the value of an interval $[a, b)$ in the persistence diagram as counting features born at a that die at b . In order to understand the appearance of negative values in the persistence diagram, structure beyond that of the integer-valued Möbius inversion is necessary. Möbius homology, a recent categorification of Möbius inversion by Patel and Skraba [15], provides that structure. It utilizes the data of a module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ to construct a chain complex and, thus, homology spaces. Such spaces provide higher-order structure with which we can study multiparameter modules. For one, Möbius homology is able to distinguish non-isomorphic modules that are indistinguishable by Möbius inversion and other persistence invariants alone.

4.1 Non-Functoriality of Images

We wish to categorify existing persistence invariants in order to obtain higher-order structure. The rank function, however, does not have a natural categorification. Indeed, consider a one-parameter module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$, and let $\text{rk}_M : (\text{Dgm}(T), \supseteq) \rightarrow \mathbb{Z}_{\geq 0}$ denote its rank function. Recall that the rank function assigns to the symbolic persistence interval $[a, b)$ the integer $\text{rk}_M([a, b)) = \dim(\text{im } M(a \leq b^-))$ where b^- denotes the greatest element in T that is strictly less than b (see Definition 3.2.11). Thus, the natural choice for lifting the rank function to a vector space-valued module is the *image*. However, taking images is *not* functorial. To demonstrate this claim, we provide a counterexample below.

Example 4.1.1: Let T be the totally ordered set $1 < 2 < 3 < 4 < 5$, and let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be the module given by

$$\mathbb{R} \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{I_2} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}} \mathbb{R}^3 \rightarrow 0$$

where I_2 denotes the identity matrix.

The reverse containment order on $\text{Dgm}(T)$ admits the relation $[1, 5] \supseteq [2, 4]$. We compute the images $\text{im } M(1 \leq 5^-) = \text{im } M(1 \leq 4)$ and $\text{im } M(2 \leq 4^-) = \text{im } M(2 \leq 3)$ below:

$$\text{im } M(1 \leq 4) = \left\langle \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix} \right\rangle \quad \text{and} \quad \text{im } M(2 \leq 3) = \left\langle \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right\rangle.$$

However, observe the data of the module M does not provide a canonical map from the space $\text{im } M(1 \leq 4)$ to $\text{im } M(2 \leq 3)$. That is, the relation $[1, 5] \supseteq [2, 4]$ does not induce a map between the corresponding images.

Indeed, the example above demonstrates that images of linear maps, the natural choice for a categorification of the rank function, are *not* functorial. Thus, for the remainder of this chapter, we will focus on the *birth-death* approach. Indeed, the birth-death function does admit a natural categorification.

4.2 Categorification of the Birth-Death Function

In order to obtain higher-order structure with which to study multiparameter modules, we will use the *birth-death* approach. More specifically, we will study higher-order structure using *birth-death modules* that categorify birth-death functions. Furthermore, we will use the birth-death approach to define persistence diagrams for multiparameter modules. In order to do so, we will restate many definitions from Section 3.2, this time specifically for a finite poset P .

Let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a multiparameter module indexed by a finite poset P . Because features in M may persist infinitely, never getting mapped to zero, we will adjoin a formal top element to P representing "infinity."

Definition 4.2.1: Let P be a finite poset. We define

$$\bar{P} = P \cup \{\infty\}$$

to be the (finite) poset obtained by adjoining a top element ∞ with the convention that $a < \infty$ for all elements $a \in P$.

We will again use symbolic persistence intervals to track both the birth and the death of a feature in M .

Definition 4.2.2: Let P be a finite poset and let $\bar{P} = P \cup \{\infty\}$. A half-open **persistence interval** is a formal symbol $[a, b)$ with $a \in P$ and $b \in \bar{P}$ such that $a \leq b$. Denote by $\text{Dgm}(P)$ the set of all persistence intervals. The subset $\{[a, a) \mid a \in P\} \subseteq \text{Dgm}(P)$ is called the **diagonal**.

Again, we emphasize that persistence intervals $[a, b)$ are formal symbols rather than subsets of the indexing poset P . Because we are using the birth-death approach, we will place the product order on $\text{Dgm}(P)$.

Definition 4.2.3: Let P be a finite poset with $\bar{P} = P \cup \{\infty\}$. Denote by $(\text{Dgm}(P), \preceq)$ the poset of persistence intervals with the **product order** \preceq . That is, we define

$$[a, b) \preceq [c, d) \quad \text{if} \quad a \leq c \text{ and } b \leq d.$$

Now, we will again compute a free presentation of a module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ to capture the birth and death of generators.

Definition 4.2.4 ([40, Definition 4.9]): A persistence module $G : P \rightarrow \text{Vec}_{\mathbb{R}}$ is **free** if it is a finite direct sum of persistence modules of the form $\mathbb{R}^{\uparrow a}$, defined by

$$\mathbb{R}^{\uparrow a}(b) = \begin{cases} \mathbb{R} & \text{if } a \leq b \\ 0 & \text{else.} \end{cases}$$

and morphisms $\mathbb{R}^{\uparrow a}(b \leq c)$ given by the identity map for $a \leq b \leq c$ and the zero map else.

Given a module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$, a **free presentation** of M consists of a free module G and a surjective natural transformation $\varphi : G \Rightarrow M$.

Much like the one-parameter case, since we are specifically considering *finite* posets P , all modules $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ admit a free presentation (see, again, [40, Proposition 4.13]).

Instead of directly computing an *integer-valued function* on $(\text{Dgm}(P), \preceq)$, we will use a free presentation $\varphi : G \Rightarrow M$ to compute a birth-death *module* indexed by $(\text{Dgm}(P), \preceq)$.

Definition 4.2.5 ([40, Section 5.1]): Let P be a finite poset with $\bar{P} = P \cup \{\infty\}$. Let $(\text{Dgm}(P), \preceq)$ denote the poset of symbolic persistence intervals with the product order. Let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module, and let $\varphi : G \Rightarrow M$ be a free presentation of M . The **birth-death module** associated to φ is the module $\text{BD}_{\varphi} : (\text{Dgm}(P), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ defined on objects by

$$\text{BD}_{\varphi}([a, b]) = \begin{cases} G(a) \cap \ker \varphi_b, & \text{if } b \in P, \\ G(a), & \text{if } b = \infty. \end{cases}$$

We call $\text{BD}_{\varphi}([a, b])$ the **birth-death space** at $[a, b]$.

For a morphism $[a, b] \preceq [c, d]$ in $(\text{Dgm}(P), \preceq)$, the induced linear map

$$\text{BD}_{\varphi}([a, b] \preceq [c, d]) : G(a) \cap \ker \varphi_b \hookrightarrow G(c) \cap \ker \varphi_d$$

is given by the inclusion induced by the structure maps of G .

Intuitively, the birth-death space $\text{BD}_\varphi([a, b])$ records the elements of the free module G that are present at index a and that map to zero in $M(b)$. Thus, the birth-death module tracks generators that are born by a and have died by b .

Indeed, the birth-death module is the natural categorification of the birth-death function. Consider a persistence module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ with free presentation $\varphi : G \Rightarrow M$. We may define the birth-death *function* associated to the free presentation φ as the dimension function of the birth-death module.

Definition 4.2.6 ([40, Definition 5.1]): Let P be a finite poset, let $\bar{P} = P \cup \{\infty\}$, and let $(\text{Dgm}(P), \preceq)$ denote the poset of symbolic persistence intervals equipped with the product order. Let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module with free presentation $\varphi : G \Rightarrow M$. The **birth-death function** associated to φ is the function

$$\text{bd}_\varphi : (\text{Dgm}(P), \preceq) \rightarrow \mathbb{Z}_{\geq 0}$$

given by the dimension function of the birth-death module $\text{BD}_\varphi : (\text{Dgm}(P), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$. That is,

$$\text{bd}_\varphi([a, b]) = \dim \text{BD}_\varphi([a, b]) = \dim (G(a) \cap \ker \varphi_b).$$

Because the structure maps in the birth-death module are inclusions, the birth-death function is a weakly increasing function.

Definition 4.2.7: Let P be a finite poset with $\bar{P} = P \cup \{\infty\}$. Let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a persistence module with free presentation $\varphi : G \Rightarrow M$. The **persistence diagram** associated to φ is the Möbius inversion of the birth-death function

$$\partial \text{bd}_\varphi : (\text{Dgm}(P), \preceq) \rightarrow \mathbb{Z}.$$




Figure 4.1: The birth-death module $\text{BD}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ for the module M as in Example 3.2.14.

That is, it is the unique function satisfying

$$\text{bd}_\varphi([a, b]) = \sum_{[c, d] \preceq [a, b]} \partial \text{bd}_\varphi([c, d]).$$

The definition of the birth-death function associated to a free presentation for general multiparameter modules agrees with the definition given for one-parameter modules in Section 3.2.3.

Remark 4.2.8: Given a module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite *totally ordered set* T and a free presentation $\varphi : G \Rightarrow M$, we may still compute the birth-death module as in Definition 4.2.5. Taking the dimension function will yield the one-parameter birth-death function given in Definition 3.2.3. To illustrate this point, consider the module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over the totally ordered set T given by $1 < 2 < 3 < 4$ and free presentation $\varphi : G \Rightarrow M$ from Examples 3.2.14 and 3.2.22. The birth-death module $\text{BD}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ is shown in Figure 4.1. Taking the dimension function yields $\text{bd}_\varphi : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}$, which is shown in Figure 3.3

Observe that the birth-death module BD_φ contains strictly more information than bd_φ . Indeed, the birth-death *module* carries the structure of vector spaces and linear maps, where the birth-death function retains only the dimensions of the birth-death spaces. Recall that, in order to define the




Figure 4.2: A poset P (left) and its order complex ΔP (right).

persistence diagram of a module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$, we Möbius inverted the birth-death function. In the following section, we categorify Möbius inversion by introducing *Möbius homology*.

4.3 Möbius Homology

Observe that, in order to compute persistence diagrams, we begin with a persistence module and define integer-valued functions in order to investigate the structure of the persistence module. When reducing a module to its persistence diagram, information is lost, particularly in the multiparameter setting. In this section, we will obtain higher-order information via *Möbius homology* [15]. Möbius homology provides a homology theory for persistence modules that *categorifies* Möbius inversion. That is, we can reconstruct Möbius inversion using Möbius homology, which provides more information and structure.

We begin by introducing Möbius homology in general. We then explore applications of Möbius homology to persistence: in particular, we will compute Möbius homology for one-parameter persistence modules in Section 4.3.1. Then, in Section 4.3.2, we will see an example of Möbius homology applied to the birth-death module.

First, we define the order complex of a poset P . This is a method for obtaining a simplicial complex associated to a poset. With the structure of a simplicial complex, we can use tools from algebraic topology (namely, homology and simplicial cosheaves) to study the structure of P .

Definition 4.3.1: Let P be a poset. Its **order complex**, denoted ΔP , is the simplicial complex constructed by including an i -simplex for all i -chains $a_0 < a_1 < \dots < a_i$ in P . Since subchains are necessarily chains themselves, ΔP is closed downwards and is, indeed, a simplicial complex.

Intuitively, the order complex captures the combinatorial "shape" of a poset by encoding the elements of P along with comparability and sequences of comparable elements (that is, longer chains). See Figure 4.2 for an example of an order complex.

Observe that, when P is a finite poset, ΔP is a finite simplicial complex. Now, given a persistence module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ over such a finite poset P , we will construct a cosheaf on the order complex ΔP using the data of M .

Definition 4.3.2 ([15, Definition 3.6]): Let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a module over a finite poset P . The **order cosheaf** of M is the simplicial cosheaf $\underline{M} : \Delta P \rightarrow \text{Vec}_{\mathbb{R}}$ defined as follows. Given a simplex $\sigma \in \Delta P$, let $\underline{M}(\sigma) = M(\min \sigma)$. For $\tau \geq \sigma$, the morphism $\underline{M}(\tau \rightarrow \sigma)$ is given by the linear map $M(\min \tau \leq \min \sigma)$.

We will now focus on a specific element b in our poset P by considering its downset: that is, all elements a such that $a \leq b$. In the order complex, a similar construction gives way to a subcomplex of ΔP .

Definition 4.3.3 ([15, Definition 3.7]): The **lower complex** of $b \in P$ is the subcomplex of ΔP given by $\Delta P_{\leq b} = \{\sigma \in \Delta P \mid \max \sigma \leq b\}$. The **strict lower complex** of $b \in P$ is the subcomplex given by $\Delta P_{< b} = \{\sigma \in \Delta P \mid \max \sigma < b\}$.

That is, the lower complex of ΔP at b contains only the chains that end at or before b . The *strict* lower complex of ΔP at b contains only the chains that end (strictly) before b . We are specifically interested in chains that end exactly at b : algebraically, we study these chains by computing relative chain spaces.

Definition 4.3.4 ([15, Definition 3.8]): Let P be a finite poset, and let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a module. Fix $b \in P$. The **Möbius chain complex** of M at b is the relative simplicial cosheaf chain complex given by

$$C_{\bullet}^{\downarrow} M(b) = C_{\bullet}(\Delta P_{\leq b}, \Delta P_{< b}; M).$$

Denote by

$$\partial_i^\downarrow : C_i^\downarrow M(b) \rightarrow C_{i-1}^\downarrow M(b)$$

the boundary maps in $C_\bullet^\downarrow M(b)$.

Applying homology in dimension i yields the **Möbius homology space** of M at b :

$$H_i^\downarrow M(b) = \ker \partial_i^\downarrow / \text{im } \partial_{i+1}^\downarrow.$$

That is, Möbius homology assigns a *vector space* for every $i \geq 0$ to every element $b \in P$, retaining chain-level structure that is lost when considering integer-valued invariants. Nonetheless, when "deategorifying," we relate Möbius homology to Möbius inversion via the following proposition.

Proposition 4.3.5 ([15, Theorem 3.13]): Let P be a finite poset, let $M : P \rightarrow \text{Vec}_{\mathbb{R}}$ be a module, and let $m : P \rightarrow \mathbb{Z}$ be the dimension function of M . That is, $m(b) = \dim M(b)$. Then, for all $b \in P$,

$$\partial m(b) = \sum_{i \geq 0} (-1)^i \dim H_i^\downarrow M(b). \quad (4.1)$$

The proposition above states precisely how Möbius homology is a categorification of Möbius inversion. Indeed, the Möbius inversion of the dimension function is the *Euler characteristic* of the order cosheaf \underline{M} (see Definition 2.3.16). Recall Example 2.3.18: we have already seen an example of two cosheaves with identical Euler characteristic but different cosheaf homology. Similarly, the extra structure of Möbius homology may be able to differentiate two distinct modules that are indistinguishable by Möbius inverting their dimension functions.

Example 4.3.6: Let P be the finite poset as shown in Figure 4.2. Let $M, N : P \rightarrow \text{Vec}_{\mathbb{R}}$ be the two modules shown in Figure 4.3. Since $\dim M(p) = \dim N(p)$ for every $p \in P$, their dimension functions $m, n : P \rightarrow \mathbb{Z}$ are equal: see Figure 4.3c. Thus, the Möbius inversions $\partial m, \partial n$ are equal as well.




Figure 4.3: Two persistence modules over the same poset P and their dimension functions $m = n$.

Although dimension cannot distinguish M and N , their Möbius homology spaces can. Consider the element $d \in P$, and consider chains in P ending at d . Order the 2- and the 1-chains lexicographically. The Möbius chain complex $C_{\bullet}^{\downarrow}M(d)$ is given by:

$$\cdots 0 \xrightarrow{\partial_3^{\downarrow}} \mathbb{R}^2 \oplus \mathbb{R}^2 \xrightarrow{\partial_2^{\downarrow} = \begin{pmatrix} -1 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}} \mathbb{R}^2 \oplus \mathbb{R} \oplus \mathbb{R} \xrightarrow{\partial_1^{\downarrow}} 0 \rightarrow 0.$$

Thus, $H_i^{\downarrow}M(d) = 0$ for all i .

On the other hand, the Möbius chain complex $C_{\bullet}^{\downarrow}N(d)$ is given by:

$$\cdots 0 \xrightarrow{\partial_3^{\downarrow}} \mathbb{R}^2 \oplus \mathbb{R}^2 \xrightarrow{\partial_2^{\downarrow} = \begin{pmatrix} -1 & 0 & -1 & 0 \\ 0 & -1 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}} \mathbb{R}^2 \oplus \mathbb{R} \oplus \mathbb{R} \xrightarrow{\partial_1^{\downarrow}} 0 \rightarrow 0.$$

Thus, $H_2^{\downarrow}N(d) \cong \mathbb{R}^4/\mathbb{R}^3 \cong \mathbb{R}$ and $H_1^{\downarrow}N(d) \cong \mathbb{R}$. For all $i \neq 1, 2$, $H_i^{\downarrow}N(d) = 0$.

Recall that we are specifically interested in applying Möbius homology to *persistence*. Thus, in the following sections, we will use Möbius homology to study existing persistence invariants. First, we will compute Möbius homology for one-parameter persistence modules, relying heavily

on their interval decomposition. Then, we will show that applying Möbius homology to the birth-death module distinguishes non-isomorphic modules with equivalent persistence diagrams.

4.3.1 Möbius Homology of One-Parameter Modules

We now want to understand Möbius homology for the simplest possible type of module. That is, we want to understand Möbius homology for modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ indexed over a finite totally ordered set T . We saw in Section 3.2 that such modules have a unique decomposition into *interval modules*. Let $\text{Dgm}(T)$ denote the set of symbolic persistence intervals $[a, b)$ for $a \leq b$ such that $a \in T$ and $b \in \bar{T} = T \cup \{\infty\}$ (see Definition 3.2.2). Then, by Proposition 3.2.4, there exists a (finite-support) function $m_{\mathcal{B}} : \text{Dgm}(T) \rightarrow \mathbb{Z}_{\geq 0}$ so that M decomposes into the direct sum

$$M \cong \bigoplus_{[a,b) \in \text{Dgm}(T)} I_{[a,b)}^{\oplus m_{\mathcal{B}}([a,b))}.$$

Moreover, $m_{\mathcal{B}}$ uniquely determines the module M . Thus, we expect that Möbius homology will reflect the known structure of such one-parameter modules.

Indeed, in this section we will compute Möbius homology of a module $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ for every element in $c \in T$. As we will see, the degree-zero Möbius homology spaces will identify the birth index a of intervals $[a, b)$ in $\text{Dgm}(T)$ so that $m_{\mathcal{B}}([a, b)) \neq 0$, whereas degree-one Möbius homology will identify the death index b when $b \neq \infty$.

In order to make this statement precise, we start by computing Möbius homology for a single interval module.

Lemma 4.3.7: Let $I_{[a,b)} : T \rightarrow \text{Vec}_{\mathbb{R}}$ be an interval module (see Definition 3.2.3). Let $c \in T$. Then,

$$H_0^{\downarrow} I_{[a,b)}(c) \cong \begin{cases} \mathbb{R} & \text{if } c = a, \\ 0 & \text{else.} \end{cases}$$

$$H_1^\downarrow I_{[a,b]}(c) \cong \begin{cases} \mathbb{R} & \text{if } c = b, \\ 0 & \text{else.} \end{cases}$$

$$H_i^\downarrow I_{[a,b]}(c) = 0 \text{ for all } i \geq 0.$$

Proof. Denote by $[a, b^-]$ the subposet of T consisting of all poset elements c such that $a \leq c < b$. That is, $[a, b^-] = \{c \in T \mid a \leq c < b\}$. Recall that interval modules are not defined for diagonal persistence intervals $[a, a)$. We will consider three cases: $c = a$, $c = b$, and $c \neq a, b$.

First, let $c = a$. As there are no elements $z < a \in T$ where $I_{[a,b]}(z) \neq 0$, the Möbius chain complex $C_\bullet^\downarrow I_{[a,b]}(a)$ is given by

$$0 \rightarrow \mathbb{R} \rightarrow 0.$$

Thus, $H_0^\downarrow I_{[a,b]}(a) \cong \mathbb{R}$ and $H_i^\downarrow I_{[a,b]}(a) = 0$ for all $i \geq 1$.

Next, let $c = b$. Since $c \in T$, $b \neq \infty$ here. Suppose $n = \#\{x \in T \mid a \leq x < b\}$. First, recall that $I_{[a,b]}(x) = \mathbb{R}$ for all $a \leq x < b$. Then, there are $\binom{n}{k}$ k -chains in $\Delta T_{\leq}^\downarrow(b) \setminus \Delta T_{<}^\downarrow(b)$ that start at or after a . That is, the Möbius chain complex $C_\bullet^\downarrow I_{[a,b]}(b)$ is given by

$$0 \rightarrow \mathbb{R} \xrightarrow{\partial_n^\downarrow} \mathbb{R} \xrightarrow{\partial_{n-1}^\downarrow} \dots \rightarrow \mathbb{R} \xrightarrow{\partial_2^\downarrow} \mathbb{R} \xrightarrow{\partial_1^\downarrow} 0 \rightarrow 0.$$

First, it is clear that $H_0^\downarrow I_{[a,b]}(b) = 0$. For $i \geq 0$, we claim that $H_{i+1}^\downarrow I_{[a,b]}(b) \cong H_i(\Delta^{n-1})$ where Δ^{n-1} denotes the standard $(n-1)$ -simplex. Indeed, view the elements $x \in [a, b^-]$ as vertices with an ordering inherited from T , and consider all chains in $[a, b^-]$. This forms the standard $(n-1)$ -simplex, and the faces of dimension i are in one-to-one correspondence with the $i+1$ chains in $\Delta T_{\leq}^\downarrow(b) \setminus \Delta T_{<}^\downarrow(b)$. The boundary map $\partial_i^{\Delta^{n-1}}$ is identical to $\partial_{i+1}^\downarrow$ (the boundary map in the Möbius chain complex). That is, $H_{i+1}^\downarrow I_{[a,b]}(b) \cong H_i(\Delta^{n-1})$. Since Δ^{n-1} is contractible, we obtain $H_1^\downarrow I_{[a,b]}(b) \cong \mathbb{R}$ and $H_i^\downarrow I_{[a,b]}(b) = 0$ for $i \geq 2$ for $b \neq \infty$.

Now, suppose $c \neq a, b$. First, consider the case $a < c < b$ and let $m = \#\{x \in T \mid a \leq x < c\}$.

Using a similar argument to above, the chain space $C_{\bullet}^{\downarrow} I_{[a,b]}(c)$ is:

$$0 \rightarrow \mathbb{R}^{\binom{m}{m}} \xrightarrow{\partial_m^{\downarrow}} \dots \rightarrow \mathbb{R}^{\binom{m}{2}} \xrightarrow{\partial_2^{\downarrow}} \mathbb{R}^{\binom{m}{1}} \xrightarrow{\partial_1^{\downarrow}} \mathbb{R} \rightarrow 0.$$

First, observe that ∂_1^{\downarrow} is the all-ones m -vector. Indeed, each 1-chain $d < c$ for $d \in [a, c^-]$ induces the identity map $I_{[a,b]}(d) \rightarrow I_{[a,b]}(c)$. Thus, $H_0^{\downarrow} I_{[a,b]}(c) = 0$. Then, again relating chains in $[a, c^-]$ with the standard $(m-1)$ -simplex, we have that $H_i^{\downarrow} I_{[a,b]}(c) \cong \tilde{H}_{i-1}(\Delta^{m-1})$ (reduced homology of the standard $(m-1)$ -simplex) for all $i \geq 1$. That is, $H_i^{\downarrow} I_{[a,b]}(c) = 0$ for all i .

It is easy to see that, when $c < a$, $H_i^{\downarrow} I_{[a,b]}(c) \cong 0$ for all $i \geq 0$. For $c > b$, we get the same chain complex as above but the dimension is shifted up. That is, $H_i^{\downarrow} I_{[a,b]}(c) = 0$ whenever $c \neq a, b$. \square

Remark 4.3.8: When applying the lemma above to a persistence interval of the form $[a, \infty)$, Möbius homology of the interval module $I_{[a,\infty)} : T \rightarrow \text{Vec}_{\mathbb{R}}$ at $c \in T$ is given by

$$H_0^{\downarrow} I_{[a,\infty)}(c) \cong \begin{cases} \mathbb{R} & \text{if } c = a, \\ 0 & \text{else.} \end{cases}$$

Furthermore, $H_i^{\downarrow} I_{[a,\infty)}(c) = 0$ for all $i \geq 0$. Indeed, for $c \in T$, it is never true that $c = \infty$. Thus, we will never have a non-zero degree-one Möbius homology space for interval modules of the form $I_{[a,\infty)}$.

Using the lemma above, we can compute Möbius homology for *any* one-parameter module M at *any* index $c \in T$.

Lemma 4.3.9: Let T be a finite totally ordered set and let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a one-parameter persistence module with interval decomposition

$$M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}.$$

Then, for each $c \in T$ and each $i \geq 0$,

$$H_i^\downarrow M(c) \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} \left(H_i^\downarrow I_{[a,b]}(c) \right)^{\oplus m_{\mathcal{B}}([a,b])}.$$

Proof. Fix an isomorphism $\varphi : M \Rightarrow \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}$. Möbius homology respects isomorphisms and distributes over direct sums. Thus, we conclude. \square

Theorem 4.3.10: Let T be a finite totally ordered set and let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be a one-parameter persistence module with interval decomposition

$$M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{\oplus m_{\mathcal{B}}([a,b])}.$$

Then, for each $c \in T$,

$$\dim H_0^\downarrow M(c) = \sum_{\substack{[a,b] \in \text{Dgm}(T) \\ a=c}} m_{\mathcal{B}}([a,b]), \quad \dim H_1^\downarrow M(c) = \sum_{\substack{[a,b] \in \text{Dgm}(T) \\ b=c}} m_{\mathcal{B}}([a,b]),$$

and $H_i^\downarrow M(c) = 0$ for all $i > 1$.

Proof. The result follows from Lemmas 4.3.7 and 4.3.9. \square

In particular, degree-zero Möbius homology records the birth indices of symbolic persistence intervals $[a, b)$ so that $m_{\mathcal{B}}([a, b)) \neq 0$ whereas degree-one Möbius homology records the *finite* death indices.

Example 4.3.11: Let $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ be the module over the totally ordered set $1 < 2 < 3 < 4$ from Example 3.2.14:

$$\mathbb{R} \xrightarrow{\begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\begin{pmatrix} 0 & 1 \end{pmatrix}} \mathbb{R} \xrightarrow{\begin{pmatrix} 0 \end{pmatrix}} 0.$$




Figure 4.4: A free module $G : P \rightarrow \text{Vec}_{\mathbb{R}}$ for P as in Figure 4.2.

We saw in Example 3.2.8 that the interval decomposition of M is given by $M \cong I_{[1,3)} \oplus I_{[2,4)}$. The

Möbius chain complex $C_{\bullet}^{\downarrow}M(2)$ is given by

$$\dots \xrightarrow{\partial_3^{\downarrow}} 0 \xrightarrow{\partial_2^{\downarrow}} \mathbb{R} \xrightarrow{\partial_1^{\downarrow} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}} \mathbb{R}^2 \xrightarrow{\partial_0^{\downarrow}} 0.$$

A quick calculation confirms that $H_0^{\downarrow}M(2) \cong \mathbb{R}$ and $H_i^{\downarrow}M(2) = 0$ for all $i > 0$.

On the other hand, the Möbius chain complex $C_{\bullet}^{\downarrow}M(3)$ is given by

$$\dots 0 \xrightarrow{\partial_3^{\downarrow}} \mathbb{R} \xrightarrow{\partial_2^{\downarrow} = \begin{pmatrix} -1 \\ 1 \\ 0 \end{pmatrix}} \mathbb{R}^3 \xrightarrow{\partial_1^{\downarrow} = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}} \mathbb{R} \xrightarrow{\partial_0^{\downarrow}} 0.$$

A quick calculation confirms that $H_1^{\downarrow}M(3) \cong \mathbb{R}$ and $H_i^{\downarrow}M(3) = 0$ for all $i \neq 1$.

Thus, Möbius homology captures the well-known structure of one-parameter persistence modules $M : T \rightarrow \text{Vec}_{\mathbb{R}}$.

4.3.2 Möbius Homology of the Birth-Death Module

In this section, we provide an example of two non-isomorphic modules whose *persistence diagrams* are equivalent (in fact, equal). Nonetheless, Möbius homology of the associated birth-death modules are distinct.

Example 4.3.12: Consider the two modules $M, N : P \rightarrow \text{Vec}_{\mathbb{R}}$ from Figure 4.3. A free module $G : P \rightarrow \text{Vec}_{\mathbb{R}}$ is shown in Figure 4.4.

There is a free presentation $\varphi : G \Rightarrow M$ with components given by

$$\varphi_a = I_2, \quad \varphi_b = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad \varphi_c = \begin{pmatrix} 0 & 1 \end{pmatrix}, \quad \varphi_d = 0.$$

Here, I_2 denotes the identity matrix. The birth-death module $\text{BD}_{\varphi} : (\text{Dgm}(P), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ and the resulting persistence diagram $\partial \text{bd}_{\varphi} : (\text{Dgm}(P), \preceq) \rightarrow \mathbb{Z}$ are shown in Figure 4.5.

There is also a free presentation $\psi : G \Rightarrow N$ with components given by

$$\psi_a = I_2, \quad \psi_b = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad \psi_c = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad \psi_d = 0.$$

The associated birth-death module $\text{BD}_{\psi} : (\text{Dgm}(P), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ and the resulting persistence diagram $\partial \text{bd}_{\psi} : (\text{Dgm}(P), \preceq) \rightarrow \mathbb{Z}$ are shown in Figure 4.6.

Observe that, although the birth-death modules are distinct, the resulting persistence diagrams are equal (even on diagonal persistence intervals): $\partial \text{bd}_{\varphi} = \partial \text{bd}_{\psi}$. However, the two birth-death modules have different Möbius homology spaces.

We focus on the interval $[a, d] \in \text{Dgm}(P)$. For the free presentation $\varphi : G \Rightarrow M$, Möbius homology of $\text{BD}_{\varphi}[a, d]$ is trivial in *all* dimensions $i \geq 0$. In particular,

$$H_0^{\downarrow} \text{BD}_{\varphi}[a, d] = 0 \quad \text{and} \quad H_1^{\downarrow} \text{BD}_{\varphi}[a, d] = 0.$$

Performing the same computations for the free presentation $\psi : G \Rightarrow N$, however, yields non-trivial Möbius homology spaces. In particular,

$$H_0^{\downarrow} \text{BD}_{\psi}[a, d] \cong \mathbb{R} \quad \text{and} \quad H_1^{\downarrow} \text{BD}_{\psi}[a, d] \cong \mathbb{R}.$$

For $i > 1$, $H_i^{\downarrow} \text{BD}_{\psi}[a, d] = 0$.




Figure 4.5: From left to right: the Hasse diagram of the interval poset $(\text{Dgm}(P), \preceq)$ for P as in Figure 4.2, the birth-death module BD_φ for free presentation $\varphi : G \Rightarrow M$ as in Example 4.3.12, and the resulting persistence diagram $\partial \text{bd}_\varphi$. All unlabeled maps in BD_φ are either the identity or the zero map.

Furthermore, we can verify Proposition 4.3.5 by taking the Euler characteristic and comparing it to the persistence diagram. In this case, $\sum_{i \geq 0} (-1)^i \dim H_i^\downarrow \text{BD}_\varphi[a, d] = 0 = \partial \text{bd}_\varphi[a, d]$. Similarly, we have that $\sum_{i \geq 0} (-1)^i \dim H_i^\downarrow \text{BD}_\psi[a, d] = 1 - 1 = 0 = \partial \text{bd}_\psi[a, d]$.

Thus, Möbius *homology* of the birth-death *module* provides extra structure that may be able to distinguish non-isomorphic multiparameter modules $M, N : P \rightarrow \text{Vec}_\mathbb{R}$ that are indistinguishable by their persistence diagrams.

4.4 Motivation for a Laplacian on the Möbius Chain Complex

Throughout this dissertation, we have discussed existing invariants for persistence modules. In Chapter 3, we used the persistence diagram to characterize the isomorphism classes of *one-parameter* modules. Indeed, for modules $M : T \rightarrow \text{Vec}_\mathbb{R}$ indexed over a finite totally ordered set T , the persistence diagrams obtained via interval decomposition, Möbius inversion of the rank function, and Möbius inversion of the birth-death function were all equivalent. This is not the case in the multiparameter setting. Furthermore, there are *no* complete discrete invariants for multiparameter modules.




Figure 4.6: From left to right: the Hasse diagram of the interval poset $(\text{Dgm}(P), \preceq)$ for P as in Figure 4.2, the birth-death module BD_ψ for free presentation $\psi : G \Rightarrow N$ as in Example 4.3.12, and the resulting persistence diagram ∂bd_ψ . All unlabeled maps in BD_ψ are either the identity or the zero map.

Thus, extra structure is often desirable when investigating multiparameter persistence modules. In particular, we saw in the previous section that Möbius homology applied to the birth-death module may distinguish persistence modules that are indistinguishable by persistence diagrams alone. The extra structure of chain complexes and homology spaces obtained in the calculation of Möbius homology provide additional information with which to work with.

The Möbius chain complex, in particular, admits boundary operators. If we equip the Möbius chain spaces with an inner product structure, these boundary operators have unique adjoints, allowing for the construction of combinatorial Hodge Laplace operators on the Möbius chain complex.

In the following chapter, we will provide an *analytic* treatment of Möbius homology using the tools developed in Chapter 2 to construct a Laplacian on the Möbius chain complex. Indeed, we will use Definition 2.3.17 to obtain a cellular cosheaf Laplacian acting on the Möbius chain spaces. This Laplacian may be defined for any degree $i \geq 0$ and any index a in the underlying finite poset P . The resulting Laplace operator is the *Möbius homology Laplacian*. In the chapter to follow, we define the Möbius homology Laplacian in general and provide a first analysis.

Chapter 5

The Möbius Homology Laplacian

We saw in Chapter 4 that Möbius homology categorifies Möbius inversion. More specifically, the Möbius chain complex is a relative cosheaf chain complex that contains useful information about the local structure of a persistence module $M : P \rightarrow \text{Vec}_{\mathbb{R}}$. When M is, instead, valued in the category of Hilbert spaces, there is extra structure with which to study M . For this entire chapter, we now consider persistence modules $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$.

Much like Hodge's original motivation for developing Hodge Laplacians in the continuous setting, we are interested in taking an analytical approach to the algebraic construction of Möbius homology. We will do so by utilizing the tools from Chapter 2 to define a Laplace operator on the Möbius chain complex: the *Möbius homology Laplacian*. The analytic information encoded in the Möbius homology Laplacian will provide useful and interesting information about the persistence diagram and Möbius homology spaces of the underlying module $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$. For one, the kernel of the Möbius homology Laplacian provides canonical, energy minimizing cycle representatives for $H_i^{\downarrow} M(a)$.

Another benefit to studying the Möbius homology Laplacian is the extra information provided by the non-zero spectrum. Indeed, in Chapter 2, we saw that the spectrum of the graph and simplicial Hodge Laplacians reflect the combinatorial structure of the underlying objects. However, the spectrum of the simplicial cosheaf Laplacian is less understood. In this chapter, we will utilize the inherent structure present in persistence (particularly one-parameter persistence) in order to make precise statements about the associated Möbius homology Laplacian. Specifically, in Section 5.1.1, we exploit the well-known structure of one-parameter modules for an explicit spectral analysis of their Möbius homology Laplacians. Later, in Chapter 6, we study a *different* Möbius homology Laplacian obtained via *filtrations* of a simplicial complex. The geometric structure present from such filtrations will allow for deep spectral analysis.

We begin this chapter by introducing the Möbius homology Laplacian and stating properties that hold in general. We then narrow our focus to the Möbius homology Laplacian for one-parameter persistence modules. In Section 5.1.1, we provide the central result of this chapter: we give a complete characterization of the eigenvalues of the degree-one Möbius homology Laplacian for one-parameter modules.

5.1 Constructing a Laplacian on the Möbius Chain Complex

For the remainder of this chapter, we will consider persistence modules valued in the category of Hilbert spaces: $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$. That is, each $M(c)$ is endowed with an inner product structure. Recall from Section 2.3.3 that the Möbius chain space

$$C_i^{\downarrow} M(c) = \bigoplus_{i\text{-chains } \sigma, \max(\sigma)=c} M(\min \sigma)$$

inherits an inner product structure as follows. In $C_i^{\downarrow} M(c)$, each summand $M(\min \sigma)$ retains its inner product structure and is orthogonal to every $M(\min \sigma')$ for distinct i -chains $\sigma' \neq \sigma$. With an inner product structure on each Möbius chain space, we can define adjoint operators for each boundary map ∂_i^{\downarrow} . Thus, assuming that all of the Hilbert spaces $M(c)$ are finite-dimensional, we construct the cellular cosheaf Laplacian on the Möbius chain space by way of [34].

Definition 5.1.1: Let $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$ be a module over a finite poset P whose stalks are finite-dimensional. Then, the degree- i **Möbius homology Laplacian** of M at $a \in P$ is the operator

$$L_i^{\downarrow} M(a) = \partial_{i+1}^{\downarrow} (\partial_{i+1}^{\downarrow})^* + (\partial_i^{\downarrow})^* \partial_i^{\downarrow}$$

acting on the Möbius chain space $C_i^{\downarrow} M(a)$. Recall that ∂_i^{\downarrow} denotes the i -th boundary operator in the Möbius chain complex $C_{\bullet}^{\downarrow} M(a)$.

Observe that the Möbius homology Laplacian is defined point-wise at a fixed index $a \in P$. Thus, provided a module $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$, we may construct a Laplacian $L_i^{\downarrow} M(a)$ for *any* poset

element $a \in P$ and for any dimension $i \geq 0$. For every $a \in P$ and $i \geq 0$, we immediately obtain the following result.

Theorem 5.1.2: Let P be a finite poset, let $a \in P$, and let $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$ be a module. Then,

$$\ker L_i^\downarrow M(a) \cong H_i^\downarrow M(a).$$

Proof. This follows immediately using results from [34, Theorem 3.1]. Indeed, Hansen and Ghrist argue that $\ker L_i^\downarrow M(a) = \ker \partial_i^\downarrow \cap \text{im}(\partial_{i+1}^\downarrow)^\perp$. By Lemma 2.1.8, the projection map $\text{proj}_{(\text{im} \partial_{i+1}^\downarrow)^\perp}$ yields an isomorphism $H_i^\downarrow M(a) = \ker \partial_i^\downarrow / \text{im} \partial_{i+1}^\downarrow \cong \ker \partial_i^\downarrow \cap \text{im}(\partial_{i+1}^\downarrow)^\perp$ where we take the intersection in $C_i^\downarrow M(a)$. \square

Furthermore, as in simplicial harmonic homology, $\ker L_i^\downarrow M(a)$ provides *energy minimizing cycle representatives* for $H_i^\downarrow M(a)$, which on its own provides only an equivalence class of cycles. In the theorems and definitions below, we make this claim precise.

Definition 5.1.3: Let P be a finite poset, let $a \in P$, and let $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$ be a module. Denote by

$$\varphi : \ker \partial_i^\downarrow \rightarrow H_i^\downarrow M(a)$$

the natural surjection obtained by taking the quotient $\ker \partial_i^\downarrow / \text{im} \partial_{i+1}^\downarrow$. The space of **cycle representatives** for a homology class $[x] \in H_i^\downarrow M(a)$ is given by $\varphi^{-1}([x])$.

Obtaining cycle representatives for $H_i^\downarrow M(a)$ is akin to finding a suitable section of φ . In the following theorem, we show that $\ker L_i^\downarrow M(a)$ provides a method for obtaining such cycle representatives.

Theorem 5.1.4: Let P be a finite poset, let $a \in P$, and let $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$ be a module. The inclusion $s : \ker L_i^\downarrow M(c) \hookrightarrow \ker \partial_i^\downarrow$ induces a section of the quotient map $\varphi : \ker \partial_i^\downarrow \rightarrow H_i^\downarrow M(a)$.

Proof. Let $[x] = [y] \in H_i^\downarrow M(a)$, meaning $x = y + \partial_{i+1}^\downarrow z$ for some $z \in C_{i+1}^\downarrow M(a)$. Since

$$\text{proj}_{(\text{im}\partial_{i+1}^\downarrow)^\perp}(x) = \text{proj}_{(\text{im}\partial_{i+1}^\downarrow)^\perp}(y + \partial_{i+1}^\downarrow z) = \text{proj}_{(\text{im}\partial_{i+1}^\downarrow)^\perp}(y),$$

the map s is well-defined and, thus, yields a canonical representative for $[x] \in H_i^\downarrow M(a)$ contained in $\ker L_i^\downarrow M(a)$. \square

In other words, Theorem 5.1.4 states that for each homology class $[x] \in H_i^\downarrow M(a)$, the element $\text{proj}_{(\text{im}\partial_{i+1}^\downarrow)^\perp}(x)$ is the *unique representative* of $[x]$ lying in $\ker L_i^\downarrow M(a)$. Moreover, as we will show below, these representatives have nice geometric properties. Recall that the energy of a vector x is given by its norm squared $\|x\|^2$ (see Definition 2.1.3).

Theorem 5.1.5: Let P be a finite poset, let $a \in P$, and let $M : P \rightarrow \text{Hilb}_\mathbb{R}$ be a module. Let $x \in \ker L_i^\downarrow M(a)$ and let $y \in \varphi^{-1}([x])$ be another cycle representative of the same homology class. Then,

$$\|x\|^2 \leq \|y\|^2.$$

That is, $x \in \ker L_i^\downarrow M(a)$ is the unique representative with minimal energy.

Proof. Since $y \in \varphi^{-1}([x])$, $[y] = [x]$ and we can write $y = x + \partial_{i+1}^\downarrow z$ for some $z \in C_{i+1}^\downarrow M(a)$. Using the inner product structure on $C_i^\downarrow M(a)$,

$$\begin{aligned} \|y\|^2 &= \langle y, y \rangle = \langle y, x \rangle + \langle y, \partial_{i+1}^\downarrow z \rangle \\ &= \langle x, x \rangle + \langle \partial_{i+1}^\downarrow z, x \rangle + \langle x, \partial_{i+1}^\downarrow z \rangle + \langle \partial_{i+1}^\downarrow z, \partial_{i+1}^\downarrow z \rangle \\ &= \langle x, x \rangle + \langle \partial_{i+1}^\downarrow z, \partial_{i+1}^\downarrow z \rangle \\ &\geq \|x\|^2. \end{aligned}$$

Note that equality holds if and only if $\partial_{i+1}^\downarrow z = 0 \implies y = x$. \square

The two results above establish that the kernel of the Möbius homology Laplacian provides canonical and energy minimizing representatives for Möbius homology spaces. In the sections and

chapters to follow, we will utilize these results for two modules arising from persistent homology. Furthermore, for these modules specifically, the inherent structure of Möbius homology makes its associated cosheaf Laplacian amenable to spectral analysis.

To illustrate the claim above, we will investigate the Möbius homology Laplacian in-depth for one-parameter persistence modules. Not only will we provide a complete characterization of the eigenvalues of this Laplacian, but we will show that these eigenvalues reflect the combinatorial structure of the persistence diagram.

5.1.1 Möbius Homology Laplacian for One-Parameter Modules

In this section, we will focus specifically on *one-parameter* modules, $M : T \rightarrow \text{Hilb}_{\mathbb{R}}$ for T a finite totally ordered set. Recall that such modules have well-known structure: as we saw in Proposition 3.2.4, one-parameter persistence modules decompose into a direct sum of interval modules. That is, denote by $\text{Dgm}(T)$ the collection of symbolic persistence intervals $[a, b)$ where $a \in T$ and $b \in \bar{T} = T \cup \{\infty\}$ such that $a \leq b$. Then, there exists a finite-support function $m_{\mathcal{B}} : \text{Dgm}(T) \rightarrow \mathbb{Z}_{\geq 0}$ such that

$$M \cong \bigoplus_{[a,b) \in \text{Dgm}(T)} I_{[a,b)}^{m_{\mathcal{B}}([a,b))}.$$

Recall that we call $[a, b)$ a *point* in the persistence diagram of M if $m_{\mathcal{B}}([a, b)) \neq 0$.

To study the Möbius chain space $C_i^{\downarrow} M(c)$ analytically, we first endow each chain space $C_i^{\downarrow} M(c)$ with an inner product structure. Recall that $C_i^{\downarrow} M(c)$ is given by $\bigoplus_{i\text{-chains } \sigma} M(\min \sigma)$. Thus, we need only describe the inner product structure on each summand $M(\min \sigma)$ or, equivalently, the spaces $M(c)$. We will rely heavily on the interval decomposition of M to define an inner product structure on $M(c)$. For elements $c \in T$ such that $a \leq c < b$, endow $I_{[a,b)}(c) \cong \mathbb{R}$ with an orthonormal basis (that is, 1 with the standard inner product). Then, $M(c) \cong \bigoplus_{[a,b) \in \text{Dgm}(T)} I_{[a,b)}(c)^{m_{\mathcal{B}}([a,b))}$, by construction, also has an orthonormal basis. Thus each summand $I_{[a,b)}$ determines an orthonormal basis vector in $M(c)$.

With an inner product structure on each chain space $C_i^\downarrow M(c)$, we can compute its associated Möbius homology Laplacian $L_i^\downarrow M(c)$. In this section, we will focus specifically on the degree-one Möbius homology Laplacian $L_1^\downarrow M(c)$. Indeed, recall from Theorem 4.3.10 that degree-one Möbius homology $H_1^\downarrow M(c)$ has dimension given by the number of points $[a, c)$ in the persistence diagram of M (counted with multiplicity).

The spectrum of the Möbius homology Laplacian $L_1^\downarrow M(c)$ will depend heavily on the *length* of persistence intervals $[a, b)$ in $\text{Dgm}(T)$.

Definition 5.1.6: Let P be a finite poset, and let $[a, b)$ be a symbolic persistence interval. Define the **length** of $[a, b)$ as

$$\text{len}([a, b)) = \#\{c \in P \mid a \leq c < b\}.$$

We emphasize that the death index b is *not* counted in the length calculation. In particular,

$$\text{len}([a, a)) = 0 \quad \text{and} \quad \text{len}([a, \infty)) = \#\{c \in P \mid a \leq c\}.$$

The proof of Lemma 4.3.7 immediately shines a light on the degree-one Möbius homology Laplacian of an interval module $I_{[a,b)}$ at the (finite) element $b \in T$, $L_1^\downarrow I_{[a,b)}(b)$.

Corollary 5.1.7: Let $I_{[a,b)} : T \rightarrow \text{Hilb}_{\mathbb{R}}$ be an interval module with $b \neq \infty$, and let $n = \text{len}([a, b))$. The degree-one Möbius homology Laplacian $L_1^\downarrow I_{[a,b)}(b) : C_1^\downarrow I_{[a,b)}(b) \rightarrow C_1^\downarrow I_{[a,b)}(b)$ is the (unnormalized) graph Laplacian on K_n , the complete graph with n vertices.

Proof. From the proof of Lemma 4.3.7, recall that $C_\bullet^\downarrow I_{[a,b)}(b)$ is isomorphic to $C_\bullet(\Delta^{n-1})$. Recall that the degree-zero simplicial Hodge Laplacian $\partial_1^{\Delta^{n-1}} (\partial_1^{\Delta^{n-1}})^T$ on $C_0(\Delta^{n-1})$ is precisely the graph Laplacian on the 1-skeleton of Δ^{n-1} , which is the complete graph on n vertices. Furthermore, with our choice of inner product on $I_{[a,b)}(b)$, adjoint operators are given by the transpose. Since $\partial_1^{\Delta^{n-1}} = \partial_1^\downarrow$ (again, see the proof of Lemma 4.3.7), we conclude. \square

Corollary 5.1.8: Let $I_{[a,b)} : T \rightarrow \text{Hilb}_{\mathbb{R}}$ be an interval module with $b \neq \infty$, and let $n = \text{len}([a, b))$. Then, the eigenvalues of $L_1^\downarrow I_{[a,b)}(b)$ are 0 (with multiplicity 1) and n (with multiplicity $n - 1$).

Proof. This is a direct consequence of Corollary 5.1.7 and Proposition 2.3.3. \square

Indeed, when considering a single interval module $I_{[a,b]}$ for $b \neq \infty$, the fact that $L_1^\downarrow I_{[a,b]}(b)$ has eigenvalue 0 with multiplicity 1 agrees with Lemma 4.3.7: $\ker L_1^\downarrow I_{[a,b]}(b) \cong H_1^\downarrow I_{[a,b]}(b) \cong \mathbb{R}$.

Indeed, we can utilize the corollary above to describe the eigenvalues of $L_1^\downarrow M(c)$ for *any* one-parameter module $M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ and *any* index $c \in T$. For one, the spectrum of $L_1^\downarrow M(c)$ will contain information about the *length* and *location* of points $[a,b]$ in the persistence diagram in relation to the chosen element $c \in T$. We will make this statement precise in the following two lemmas. First, however, we will need the following proposition.

Proposition 5.1.9: Let A be an $n \times n$ matrix with eigenvalues $\lambda_1, \lambda_2, \dots, \lambda_n$ and corresponding eigenvectors $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n$. Then, the eigenvalues of $A + aI_n$ are $a + \lambda_1, a + \lambda_2, \dots, a + \lambda_n$ with corresponding eigenvectors $\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n$.

Proof. This follows from the calculation $(A + aI_n)\vec{x}_i = A\vec{x}_i + aI_n\vec{x}_i = \lambda_i\vec{x}_i + a\vec{x}_i = (\lambda_i + a)\vec{x}_i$. \square

Indeed, we can use Corollary 5.1.8 to show that the spectrum reflects the points $[a,b]$ in the persistence diagram "below" c : that is, $b \leq c$.

Lemma 5.1.10: Let $M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ be a module over the totally ordered set T and let $c \in T$. Then, for each $[a,b] \in \text{Dgm}(T)$ so that $m_{\mathcal{B}}([a,b]) \neq 0$ and $b \leq c$, $L_1^\downarrow M(c)$ has eigenvalues

$$\begin{aligned} &\text{len}([b,c]) \text{ with multiplicity } m_{\mathcal{B}}([a,b]), \\ &\text{len}([a,c]) \text{ with multiplicity } m_{\mathcal{B}}([a,b]) \cdot (\text{len}([a,b]) - 1). \end{aligned}$$

Proof. First, since $b \leq c$ and $c \in T$, we necessarily have that $b \neq \infty$.

Since the inner product of $M(c)$ is such that distinct interval modules are orthogonal, we can write $L_1^\downarrow M(c)$ as direct sum of $L_1^\downarrow M(c)|_{I_{[a,b]}}$ for all persistence intervals $[a,b]$ counted with multiplicity $m_{\mathcal{B}}([a,b])$. Thus, it suffices to look at one point in the persistence diagram at a time: for this lemma, we will focus on $[a,b]$ so that $b \leq c$.

We claim that $L_1^\downarrow M(c)|_{I_{[a,b]}}$ can be expressed as $L_0^{K_{\text{len}([a,b])}} + \text{len}([b,c])I_{\text{len}([a,b])}$. Indeed, view $L_1^\downarrow M(c)|_{I_{[a,b]}}$ as the (unnormalized) graph Laplacian on the complete graph on $\text{len}([a,b])$ vertices with an extra edge attached to each vertex for every x such that $b \leq x < c$. These edges connect an "empty" vertex corresponding to x to each element d such $a \leq d < b$. Formally, the "empty" vertices come from our module structure: $I_{[a,b]}(x) = 0$. This amounts to adding $\text{len}([b,c])$ to the degree of each vertex: that is,

$$L_1^\downarrow M(c)|_{I_{[a,b]}} = L_0^{K_{\text{len}([a,b])}} + \text{len}([b,c])I_{\text{len}([a,b])}.$$

By combining Proposition 5.1.9 with Corollary 5.1.8, the eigenvalues of $L_1^\downarrow M(c)|_{I_{[a,b]}}$ are given by $0 + \text{len}([b,c])$ with multiplicity 1 and $\text{len}([a,b]) + \text{len}([b,c]) = \text{len}([a,c])$ with multiplicity $\text{len}([a,b]) - 1$. Since the eigenvalues of a direct sum of matrices is the disjoint union of the eigenvalues of each summand, we conclude. \square

Observe that the above result agrees with Corollary 5.1.8. Indeed, consider a single interval module $I_{[a,b]} : T \rightarrow \text{Hilb}_{\mathbb{R}}$ such that $b \neq \infty$. In this case, $m_{\mathcal{B}}([a,b]) = 1$. Then, $L_1^\downarrow I_{[a,b]}(b)$ has an eigenvalue of $\text{len}([b,b]) = 0$ with multiplicity 1 and $\text{len}([a,b])$ with multiplicity $(\text{len}([a,b]) - 1)$.

Now, we will see that $L_1^\downarrow M(c)$ also contains information on the points $[a,b]$ in the persistence diagram that "intersect" c : that is, $a < c < b$.

Lemma 5.1.11: Let $M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ be a module over the totally ordered set T and let $c \in T$. Then, for each $[a,b] \in \text{Dgm}(T)$ so that $m_{\mathcal{B}}([a,b]) \neq 0$ and $a < c < b$, $L_1^\downarrow M(c)$ has eigenvalues

$$\text{len}([a,c]) \text{ with multiplicity } m_{\mathcal{B}}([a,b]) \cdot \text{len}([a,c]).$$

Proof. Again, we will consider one summand $L_1^\downarrow M(c)|_{I_{[a,b]}}$ of $L_1^\downarrow M(c)$ at a time. Now, we will focus on points $[a,b]$ in the persistence diagram so that $a < c < b$.

The Laplacian $L_1^\downarrow M(c) \Big|_{I_{[a,b]}}$ is given by $L_0^{K_{\text{len}([a,c])}} + (\partial_1^\downarrow)^* \partial_1^\downarrow$. Since $c < b$, the boundary map ∂_1^\downarrow is the $\text{len}([a, c])$ -vector of all 1's and, thus, $(\partial_1^\downarrow)^* \partial_1^\downarrow$ is the $\text{len}([a, c]) \times \text{len}([a, c])$ matrix of all 1's, denoted J . Thus, $L_1^\downarrow M(c) \Big|_{I_{[a,b]}} = L_0^{K_{\text{len}([a,c])}} + J = \text{len}([a, c]) I_{\text{len}([a,c])}$. Since the eigenvalues of a diagonal matrix are precisely the diagonal entries, we conclude. \square

Intuitively, the lemma above states that the spectrum of $L_1^\downarrow M(c)$ contains the length of the "truncation" of the point $[a, b]$ in the persistence diagram at c . The multiplicity of $\text{len}([a, c])$ in the spectrum is also given by this length.

Remark 5.1.12: In particular, Lemma 5.1.11 applies to persistence intervals of the form $[a, \infty)$. That is, let $M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ be a module over the totally ordered set T and fix $c \in T$. Then, for each $[a, \infty) \in \text{Dgm}(T)$ so that $m_{\mathcal{B}}([a, \infty)) \neq 0$ and $a < c$, then $L_1^\downarrow M(c)$ has eigenvalue $\text{len}([a, c])$ with multiplicity $m_{\mathcal{B}}([a, \infty)) \cdot \text{len}([a, c])$

Combine the previous two lemmas to get a complete characterization of the eigenvalues of the degree-one Laplacian $L_1^\downarrow M(c)$ for any $c \in T$.

Theorem 5.1.13: Let $M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ be a module over the totally ordered set T and let $c \in T$. Then, the eigenvalues of $L_1^\downarrow M(c)$ are given by

For $[a, b)$ such that $b \leq c$: $\text{len}([b, c])$ with multiplicity $m_{\mathcal{B}}([a, b))$,

$\text{len}([a, c])$ with multiplicity $m_{\mathcal{B}}([a, b)) \cdot (\text{len}([a, b)) - 1)$,

For $[a, b)$ such that $a < c < b$: $\text{len}([a, c])$ with multiplicity $m_{\mathcal{B}}([a, b))$.

Proof. Since the Laplacian is diagonalizable (as it is self-adjoint), the number of eigenvalues (with multiplicity) must equal the dimension of $C_1^\downarrow M(c)$. By definition, $C_1^\downarrow M(c)$ is given by the direct sum

$$\bigoplus_{d < c} M(d) \cong \bigoplus_{d < c} \left(\bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}(d)^{m_{\mathcal{B}}([a,b])} \right),$$




Figure 5.1: The persistence diagram $m_{\mathcal{B}}$ for module with interval decomposition $I_{[1,4]} \oplus I_{[2,7]} \oplus I_{[4,5]}$.

and, thus

$$\begin{aligned} \dim \bigoplus_{d < c} M(d) &= \dim \bigoplus_{d < c} \left(\bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}(d)^{m_{\mathcal{B}}([a,b])} \right) \\ &= \sum_{d < c} \dim \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}(d)^{m_{\mathcal{B}}([a,b])}. \end{aligned}$$

Observe that the chain space $C_1^\downarrow M(c)$ only includes interval modules $I_{[a,b]}$ such that $a < c$. The corresponding point $[a,b]$ in the persistence diagram either "intersects" c (meaning $a < c < b$) or does not (meaning $b \leq c$).

If $a < c < b$, then $I_{[a,b]}$ appears $m_{\mathcal{B}}([a,b]) \cdot \text{len}([a,c])$ many times in the direct sum $C_1^\downarrow M(c)$. In this case, Lemma 5.1.11 accounts for all possible eigenvalues. On the other hand, if a point $[a,b]$ is such that $b \leq c$, then there are $m_{\mathcal{B}}([a,b]) \cdot \text{len}([a,b])$ many copies of $I_{[a,b]}$ in $C_1^\downarrow M(c)$. In this case, Lemma 5.1.10 accounts for all possible eigenvalues.

In conclusion, Lemmas 5.1.10 and 5.1.11 provide the same number of eigenvalues (with multiplicity) as the dimension of $C_1^\downarrow M(c)$ and, thus, provide the entire spectrum. \square

Theorem 5.1.13 is best illustrated with an example.

Example 5.1.14: Let $M : T \rightarrow \text{Hilb}_{\mathbb{R}}$ be a module over the totally ordered set T given by all integers in the interval $[1, 7]$, and suppose M has interval decomposition $M \cong I_{[1,4]} \oplus I_{[2,7]} \oplus I_{[4,5]}$. See Figure 5.1 for a visualization of the persistence diagram $m_{\mathcal{B}}$.

We will compute $L_1^\downarrow M(4)$. There are $\binom{3}{2} = 3$ two-chains in T that end at 4 and $\binom{3}{1} = 3$ one-chains in T that end at 4. Order these chains lexicographically. The Möbius chain complex

$C_{\bullet}^{\downarrow}M(4)$ is given by

$$\dots \rightarrow \mathbb{R} \oplus \mathbb{R} \oplus \mathbb{R}^2 \xrightarrow{\partial_2^{\downarrow}} \mathbb{R} \oplus \mathbb{R}^2 \oplus \mathbb{R}^2 \xrightarrow{\partial_1^{\downarrow}} \mathbb{R}^2 \rightarrow 0$$

where the boundary maps are

$$\partial_2^{\downarrow} = \begin{pmatrix} -1 & -1 & 0 & 0 \\ 1 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \partial_1^{\downarrow} = \begin{pmatrix} 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}.$$

First, observe that $H_0^{\downarrow}M(4) = \mathbb{R}^2 / \text{im } \partial_1^{\downarrow} \cong \mathbb{R}$, which coincides with Theorem 4.3.10. Indeed, 4 is the birth index of the point $[4, 5)$ in the persistence diagram. In the eigenvalue computations that follow, we will see that $H_1^{\downarrow}M(4) \cong \mathbb{R}$, which coincides with 4 being the death index of the point $[1, 4)$. Then, compute the Laplacian $L_1^{\downarrow}M(4) = \partial_2^{\downarrow}(\partial_2^{\downarrow})^T + (\partial_1^{\downarrow})^T\partial_1^{\downarrow}$:

$$L_1^{\downarrow}M(4) = \begin{pmatrix} 2 & -1 & 0 & -1 & 0 \\ -1 & 2 & 0 & -1 & 0 \\ 0 & 0 & 2 & 0 & 0 \\ -1 & -1 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 2 \end{pmatrix}.$$

Its eigenvalues are $\{0, 2, 2, 3, 3\}$.

First, recall Theorem 4.3.10: $\dim \ker L_1^{\downarrow}M(4) = 1$ reflects the fact that 4 is the death index of a point in the persistence diagram. Next, recall Lemma 5.1.11: for this example, the only point $[a, b)$ in the persistence diagram with $b \leq 4$ is $[1, 4)$. Thus, the spectrum contains $\{0, 3, 3\}$. Indeed, the point $[1, 4)$ has length 3, and its death index is 0 steps away from 4. Finally, recall Lemma 5.1.10: $[2, 7)$ is the only point in the persistence diagram that intersects 4, meaning the

spectrum contains $\{2, 2\}$. Indeed, the "truncated" interval $[2, 4)$ has length 2. By Theorem 5.1.13, these two cases provide the entire spectrum: $\{0, 2, 2, 3, 3\}$.

In conclusion, we know precisely what the eigenvalues of $L_1^\downarrow M(c)$ are for any index $c \in T$. In particular, if $c = \max T$, then we can obtain even more information from the spectral decomposition of $L_1^\downarrow M(c)$. In this case, its *eigenvectors* reflect the points in the persistence diagram with *finite* death index. We make this statement precise in the following corollary.

Corollary 5.1.15: Let $M \cong \bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ be a module over a totally ordered set T and let $c = \max T$. Fix a persistence interval $[a, b] \in \text{Dgm}(T)$ with $b \neq \infty$, and let $n = \text{len}([a, b])$. Let $m = m_{\mathcal{B}}([a, b])$. Then, for each of the m distinct summands

$$I_{[a,b]}^{(1)}, \dots, I_{[a,b]}^{(m)}$$

in the interval decomposition of M , $L_1^\downarrow M(c)$ has an eigenvector \vec{x}_j with eigenvalue $\text{len}([b, c])$ so that:

1. \vec{x}_j has exactly n non-zero entries,
2. each non-zero entry of \vec{x}_j is equal to 1, and
3. the support of \vec{x}_j consists precisely of the basis elements $I_{[a,b]}^{(j)}(x) \in C_1^\downarrow M(c)$ for $a \leq x < b$.

In particular, if $m > 1$, the m eigenvectors $\vec{x}_1, \dots, \vec{x}_m$ have pairwise disjoint supports.

Proof. Recall that $L_1^\downarrow M(c)$ can be written as a block diagonal matrix. In particular, there is a block $L_1^\downarrow M(c) \Big|_{I_{[a,b]}^{(j)}}$ for each copy of $I_{[a,b]}$ in the interval decomposition of M . Thus, we focus on one such block for the remainder of the proof.

By Lemma 5.1.10, we have that $L_1^\downarrow M(c) \Big|_{I_{[a,b]}^{(j)}} = L_0^{K_n} + \text{len}([b, c])I$, where $L_0^{K_n}$ denotes the graph Laplacian for the complete graph on n vertices. By Proposition 5.1.9, the eigenvectors of $L_1^\downarrow M(c) \Big|_{I_{[a,b]}^{(j)}}$ are the same as the eigenvectors of $L_0^{K_n}$. In this case, the all-ones vector is an eigenvector with eigenvalue $\text{len}([b, c])$ (see the proof of Proposition 2.3.3).

Extending to the entire space $C_1^\downarrow M(c)$, we obtain an eigenvector of $L_1^\downarrow M(c)$ with exactly n non-zero entries, each equal to 1. These non-zero entries correspond precisely to the basis elements $I_{[a,b]}^{(j)}(x) \in C_1^\downarrow M(c)$ for $a \leq x < b$. Since distinct copies $I_{[a,b]}^{(i)}$ determine distinct blocks in $L_1^\downarrow M(c)$, the resulting eigenvectors have distinct support. \square

Example 5.1.16: Recall Example 5.1.14 and Figure 5.1. Since there are no points in the persistence diagram of the form $[a, \infty)$, we can compute the spectral decomposition of $L_1^\downarrow M(7)$ to obtain a complete picture of the interval decomposition of the module $M \cong I_{[1,4]} \oplus I_{[2,7]} I_{[4,5]}$. Order the summands in $C_1^\downarrow M(7)$ according to the lexicographic order. That is, write $C_1^\downarrow M(c)$ as follows:

$$M(1) \oplus M(2) \oplus M(3) \oplus M(4) \oplus M(5) \oplus M(6)$$

or, in terms of interval modules,

$$I_{[1,4]}(1) \oplus I_{[1,4]}(2) \oplus I_{[2,7]}(2) \oplus I_{[1,4]}(3) \oplus I_{[2,7]}(3) \oplus I_{[2,7]}(4) \oplus I_{[4,5]}(4) \oplus I_{[2,7]}(5) \oplus I_{[2,7]}(6).$$

Recall that each summand above is isomorphic to \mathbb{R} . Since the boundary maps ∂_1^\downarrow and ∂_0^\downarrow are much larger than in Example 5.1.14, we omit them here. Computing $\partial_1^\downarrow(\partial_1^\downarrow)^* + (\partial_0^\downarrow)^*\partial_0^\downarrow$ yields

$$L_1^\downarrow M(7) = \begin{pmatrix} 5 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 5 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 4 & 0 & -1 & -1 & 0 & -1 & -1 \\ -1 & -1 & 0 & 5 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 4 & -1 & 0 & -1 & -1 \\ 0 & 0 & -1 & 0 & -1 & 4 & 0 & -1 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & -1 & 0 & -1 & -1 & 0 & 4 & -1 \\ 0 & 0 & -1 & 0 & -1 & -1 & 0 & -1 & 4 \end{pmatrix}.$$

The eigenvalues of $L_1^\downarrow M(7)$ are given by $\{0, 2, 3, 5, 5, 5, 5, 6, 6\}$ with eigenvectors given by:

$$E = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & -1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & -1 & -1 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}.$$

Here, the eigenvector in column i of E corresponds to the i -th smallest eigenvalue.

Notice the first three eigenvectors: the first column of E with eigenvalue $7 - 7 = 0$ has five non-zero entries, identifying the summands $I_{[2,7)}(x)$ for $2 \leq x < 7$. The second column with eigenvalue $7 - 5 = 2$ has one non-zero entry, identifying the summand $I_{[4,5)}(4)$. The third column with eigenvalue $7 - 4 = 3$ has three non-zero entries, identifying the summands $I_{[1,4)}(x)$ for $1 \leq x < 4$.

In summary, an eigenvector of $L_1^\downarrow M(\max T)$ consisting of only 1's and 0's identifies a single summand $I_{[a,b)}$ (such that $b \neq \infty$) in the interval decomposition $M \cong \bigoplus_{[a,b) \in \text{Dgm}(T)} I_{[a,b)}^{m_{\mathcal{B}}([a,b))}$ of a one-parameter module $M : T \rightarrow \text{Hilb}_{\mathbb{R}}$.

It is, thus, tempting to claim that we can construct the persistence diagram of M using the spectral decomposition of $L_1^\downarrow M(\max T)$. However, there is an important subtlety hiding here. Recall that we considered an inner product on each $M(c)$ so that its interval decomposition has an orthonormal basis. Provided an isomorphism $\varphi : M \Rightarrow \bigoplus_{[a,b) \in \text{Dgm}(T)} I_{[a,b)}^{m_{\mathcal{B}}([a,b))}$, we can define such an inner product on each $M(c)$ as follows. Denote by $\langle \cdot, \cdot \rangle_{\oplus}$ the inner product on $M(c)$ so

that $\bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ has an orthonormal basis. Then, for any $x, y \in M(c)$, define

$$\langle x, y \rangle_{M(c)} = \langle \phi(x), \phi(y) \rangle_{\oplus}.$$

However, having an isomorphism means that we already know the interval decomposition of M . To be able to truly construct this interval decomposition from $L_1^\downarrow M(\max T)$, we need to answer the following inverse problem: *which inner products on M guarantee that its interval decomposition has an orthonormal basis?*

5.2 Discussion and Directions for Future Work

In this chapter, namely in Section 5.1.1, we provide a complete characterization of the spectrum of $L_1^\downarrow M(c)$ for a one-parameter module $M : T \rightarrow \text{Hilb}_{\mathbb{R}}$ at a fixed index $c \in T$, where M has interval decomposition $\bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$. By Theorem 5.1.13, the spectrum of $L_1^\downarrow M(c)$ captures the combinatorial structure of the persistence diagram $m_{\mathcal{B}} : \text{Dgm}(T) \rightarrow \mathbb{Z}_{\geq 0}$ for persistence intervals $[a, b]$ such that $b \leq c$. Moreover, by Corollary 5.1.15, certain eigenvalues of $L_1^\downarrow M(\max T)$ identify summands $I_{[a,b]}$ with finite death index, $b \neq \infty$.

Interesting future work in this direction is as follows. As mentioned at the end of the previous section, these results rely heavily on a choice of inner product on each $M(c)$ that is compatible with its interval decomposition. Determining which inner product structure(s) on M guarantee that $\bigoplus_{[a,b] \in \text{Dgm}(T)} I_{[a,b]}^{m_{\mathcal{B}}([a,b])}$ admits an orthonormal basis is an interesting and subtle, though potentially difficult, problem. A solution to this problem may lead to a new method for computing the interval decomposition of M .

Additionally, the results presented in this chapter rely on the interval decomposition of one-parameter modules. It would be interesting to study whether similar methods could be applied to the spectrum of the Möbius homology Laplacian for zigzag modules or other multiparameter modules that are known to decompose into indecomposable summands [45]. These questions for future work are a direct continuation of the one-parameter case.

Finally, recall that persistence homology is motivated by the study of filtered simplicial complexes. The evolving topology of such filtered complexes give rise to persistence modules: in particular, a *birth-death module*. Furthermore, the *geometry* of a filtered complex leads to a natural choice of an inner product structure on its associated birth-death spaces and hence Möbius chain spaces. Thus, we are interested in studying *analytical* properties of modules arising from filtrations using the geometry of the underlying simplicial complex. In Chapter 6, we will analyze a Möbius homology Laplacian arising from filtrations of a simplicial complex.

Chapter 6

Möbius Homology Laplacian for Simplicial

Filtrations

In the previous chapter, we defined a Laplace operator on the Möbius chain complex, called the *Möbius homology Laplacian*. This Laplacian depends on a module $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$, an index $a \in P$, and dimension $i \geq 0$. When working specifically with one-parameter modules, we saw in Section 5.1.1 that the spectrum of the Möbius homology Laplacian captured the structure of the persistence diagram. In this chapter, we will investigate the Möbius homology Laplacian for the birth-death module. Unlike Chapter 5.1.1, we are no longer constrained to the one-parameter setting. That is, we consider modules over finite posets P .

To study the Möbius homology Laplacian for the birth-death module, we will return to the original motivation for persistent homology: *filtrations* of a (finite) simplicial complex K . Given a filtration of K over a finite poset P , we will define a new *birth-death module* that captures the evolving topology of K throughout the filtration.

Moreover, the extra data of a filtered simplicial complex will allow for an inner product structure on birth-death spaces that reflects the geometry of the underlying simplicial complex. In particular, this inner product structure will allow us to study analytical properties of the birth-death module by computing its associated Möbius homology Laplacian. We will see that the kernel of the degree-zero Möbius homology Laplacian has nice geometric properties. That is, it provides a method for obtaining cycle representatives for Möbius homology that emphasizes "newly added" simplices. Furthermore, the spectrum of the Möbius homology Laplacian for the birth-death module will, again, reflect the combinatorial structure of the persistence diagram.

6.1 Birth-Death Modules from Simplicial Filtrations

We begin this chapter by describing how to construct the *birth-death module* directly from the data of a filtration of a simplicial complex K . This method will differ slightly from the method described in Section 4.2, but, as we will see, the two methods provide equivalent persistence diagrams in the end.

Recall from Definition 1.1.9 that a filtration of a finite simplicial complex K is a functor from a poset P to the category of subcomplexes of K , denoted $F : P \rightarrow \Delta K$. In this chapter, we will consider filtrations over a *finite* poset P . Intuitively, a filtration assigns a subcomplex $F(a)$ to each $a \in P$, and for every relation $a \leq b$ there is an *inclusion* $F(a) \subseteq F(b)$. In other words, the simplicial complex grows throughout P .

When working with a filtration $F : P \rightarrow \Delta K$, we construct persistence modules associated to F as follows.

Definition 6.1.1 ([9]): Let P be a finite poset, K a simplicial complex, and $F : P \rightarrow \Delta K$ a filtration of K . Fix a dimension $p \geq 0$. The degree- p **persistent homology module** associated to F is defined as the composition

$$H_p F : P \xrightarrow{F} \Delta K \xrightarrow{H_p(-)} \text{Vec}_{\mathbb{R}}.$$

The morphisms are given by the induced maps $H_p F(a \leq b) : H_p F(a) \rightarrow H_p F(b)$.

That is, the degree- p persistent homology module assigns to each $a \in P$ the p -th homology of $F(a)$. In order to compute the persistence diagram for the module $H_p F$, we begin by first obtaining a free presentation $\varphi : G \Rightarrow H_p F$ (see Definition 4.2.4). Then compute the birth-death module $\text{BD}\varphi$ and function $\text{bd}\varphi$, obtaining the persistence diagram by Möbius inverting $\partial \text{bd}\varphi$: see Section 4.2. We outline this process through an example, given below.




Figure 6.1: A filtration $F : T \rightarrow \Delta K$ for T given by $1 < 2 < 3 < 4$.

Example 6.1.2: Let K be the simplicial complex consisting of all subsets of the (oriented) two-simplices $[v_1v_2v_4]$ and $[v_2v_3v_4]$. See Figure 6.1 for an example of a filtration $F : T \rightarrow \Delta K$ (in particular, with $F(\max T) = K$).

Computing the degree-one persistent homology module $H_1F : T \rightarrow \text{Vec}_{\mathbb{R}}$ yields the same module as $M : T \rightarrow \text{Vec}_{\mathbb{R}}$ in Example 3.2.14. Thus, a free presentation $\varphi : G \Rightarrow H_1F$ is given in Example 3.2.22. See Figure 3.3 for the birth-death module and resulting persistence diagram.

However, when considering simplicial filtrations, we can compute the birth-death module using the data of F directly, no longer requiring the data of a free presentation. This module is constructed using the birth-death spaces as first defined by Ghrist and Henselman-Petrusek in [46].

Definition 6.1.3: Let K be a finite simplicial complex, P be a finite poset, and $\bar{P} = P \cup \{\infty\}$. Let $(\text{Dgm}(P), \preceq)$ be the poset of symbolic persistence intervals with the product order. Fix $p \geq 0$, and let $F : P \rightarrow \Delta K$ be a filtration. The **birth-death module** associated to F is the module $\text{BD}_pF : (\text{Dgm}(P), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ defined on objects by

$$\text{BD}_pF([a, b)) = \begin{cases} Z_pF(a) \cap B_pF(b), & \text{if } b \in P, \\ Z_pF(a), & \text{if } b = \infty, \end{cases}$$

where we view the cycle and boundary spaces $Z_pF(a)$ and $B_pF(b)$ as subspaces of $C_p(K)$, the p th chain space of the total simplicial complex K . Thus, we take the intersection in $C_p(K)$.

For each relation $[c, d) \preceq [a, b)$ in $(\text{Dgm}(P), \preceq)$, the birth-death module assigns an injective linear map $\text{BD}_pF([c, d) \preceq [a, b)) : \text{BD}_pF([c, d)) \hookrightarrow \text{BD}_pF([a, b))$ induced by the inclusions of cycle and boundary spaces: $Z_pF(c) \hookrightarrow Z_pF(a)$ and $B_pF(d) \hookrightarrow B_pF(b)$.

The birth-death module applied to a persistence interval $[a, b)$ consists of cycles at step a of the filtration F that are boundaries at step b . We will define the birth-death function and persistence diagram using the module $\text{BD}_p F$ directly, rather than choosing a free presentation.

Definition 6.1.4 ([8, Definition 5.6]): Let K be a finite simplicial complex, P be a finite poset, and $\bar{P} = P \cup \{\infty\}$. Let $(\text{Dgm}(P), \preceq)$ be the poset of symbolic persistence intervals with the product order. Let $F : P \rightarrow \Delta K$ be a filtration, and fix $p \geq 0$. The **birth-death function** associated to F is the integer-valued map

$$\text{bd}_p F : (\text{Dgm}(P), \preceq) \rightarrow \mathbb{Z}_{\geq 0}$$

given by the dimension function of the birth-death module $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$. That is,

$$\text{bd}_p F([a, b)) = \dim(Z_p F(a) \cap B_p F(b)) = \dim \text{BD}_p F([a, b)).$$

Because the structure maps in the birth-death module are inclusions, the birth-death function is a weakly increasing function.

Intuitively, the birth-death function counts the number of *cycles* in the subcomplex $F(a)$ that are boundaries in the larger subcomplex $F(b)$. In other words, $\text{bd}_p F$ counts the number of cycles born at or before a that become boundaries at or before b .

Definition 6.1.5: Let K be a finite simplicial complex, P be a finite poset, and $\bar{P} = P \cup \{\infty\}$. Let $F : P \rightarrow \Delta K$ be a filtration, and fix $p \geq 0$. The **persistence diagram** of F is the Möbius inversion of the birth-death function

$$\partial \text{bd}_p F : (\text{Dgm}(P), \preceq) \rightarrow \mathbb{Z}.$$

That is, it is the unique function satisfying

$$\text{bd}_p F([a, b)) = \sum_{[c, d) \preceq [a, b)} \partial \text{bd}_p F([c, d)).$$

Whether computing the persistent homology module $H_p F$ and the birth-death function of a free presentation $\varphi : G \Rightarrow H_p F$, or instead computing the birth-death module $\text{BD}_p F$ from the filtration F directly, the resulting persistence diagrams will be equivalent, as made clear in the following proposition.

Proposition 6.1.6 ([41, Proposition 6.3]²): Let $F : P \rightarrow \Delta K$ be a filtration of a finite simplicial complex K over a finite poset P . Let $\varphi : G \Rightarrow H_p F$ be a free presentation of the persistent homology module $H_p F$. Let bd_φ denote the birth-death function associated to φ , and let $\text{bd}_p F$ denote the birth-death function associated to F . Then, $\partial \text{bd}_\varphi([a, b]) = \partial \text{bd}_p F([a, b])$ when $a \neq b$.

Example 6.1.7: Let T be the totally ordered set $1 < 2 < 3 < 4$, and let $F : T \rightarrow \Delta K$ be the filtration shown in Figure 6.1. The birth-death module $\text{BD}_1 F : (\text{Dgm}(T), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ and the resulting persistence diagram $\partial \text{bd}_1 F : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}$ are shown in Figure 6.2. Observe that, in this case, we have *equality* of the two persistence diagrams: $\partial \text{bd}_1 F = \partial \text{bd}_\varphi$, where $\partial \text{bd}_\varphi$ is shown in Figure 3.3.

Thus, for the remainder of this chapter, we will focus specifically on the birth-death module obtained directly from a filtration $F : P \rightarrow \Delta K$. In the following section, we will discuss the Möbius chain complex on such modules. With an inner product structure on the birth-death spaces, we can apply the tools from Chapter 5 to study the associated Möbius homology Laplacian.

6.2 Hilbert Space Structure and Möbius Homology Laplacians

Recall that, for every interval $[a, b] \in (\text{Dgm}(P), \preceq)$, the Möbius chain complex $C_\bullet^\downarrow \text{BD}_p F([a, b])$ is constructed as follows (see Definition 4.3.4):

$$\cdots \xrightarrow{\partial_{i+1}^\downarrow} \bigoplus_{\substack{i\text{-chains } c \\ \max c = [a, b]}} \text{BD}_p F(\min c) \xrightarrow{\partial_i^\downarrow} \bigoplus_{\substack{(i-1)\text{-chains } c \\ \max c = [a, b]}} \text{BD}_p F(\min c) \xrightarrow{\partial_{i-1}^\downarrow} \cdots$$

²As noted in [41], Proposition 6.1.6 is due to Alex McCleary.




Figure 6.2: From left to right: the birth-death module $\text{BD}_1 F : (\text{Dgm}(T), \preceq) \rightarrow \text{Vec}_{\mathbb{R}}$ associated to the filtration F in Figure 6.1 and the resulting persistence diagram $\partial \text{bd}_1 F : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}$.

We equip each birth-death space with the following inner product and use it to induce an inner product on each chain space $C_i^\downarrow \text{BD}_p F([a, b])$. Let $K^{[p]}$ denote the collection of p -simplices in K . Given two p -cycles $z = \sum_{\sigma \in K^{[p]}} a_\sigma \sigma$ and $w = \sum_{\sigma \in K^{[p]}} b_\sigma \sigma$ in $\text{BD}_p F([a, b])$, define their inner product as follows:

$$\langle z, w \rangle_{\text{BD}_p F([a, b])} = \sum_{\sigma \in K^{[p]}} a_\sigma b_\sigma.$$

That is, the inner product on $\text{BD}_p F([a, b])$ is inherited from the standard inner product on $C_p(K)$ as discussed in Section 2.2. Indeed, this is the inner product structure commonly used in the literature: see [16–21].

Thus, we now view the birth-death module as a module valued in the category of Hilbert spaces: $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$. The Hilbert space structure allows us to define the Möbius homology Laplacian for birth-death modules as in Definition 5.1.1. That is, the degree- i Möbius homology Laplacian for $\text{BD}_p F([a, b])$ is the operator

$$L_i^\downarrow \text{BD}_p F([a, b]) = (\partial_i^\downarrow)^* \partial_i^\downarrow + \partial_{i+1}^\downarrow (\partial_{i+1}^\downarrow)^*$$

acting on $C_i^\downarrow \text{BD}_p F([a, b])$.

For the remainder of this chapter, we will focus specifically on $L_0^\downarrow \text{BD}_p F([a, b]) = \partial_1^\downarrow (\partial_1^\downarrow)^*$. That is, we will focus on the end of the Möbius chain complex for $\text{BD}_p F([a, b])$, which is given in terms of birth-death spaces below:

$$\dots \xrightarrow{\partial_2^\downarrow} \bigoplus_{[c,d] \prec [a,b]} \text{BD}_p F([c, d]) \xrightarrow{\partial_1^\downarrow} \text{BD}_p F([a, b]) \rightarrow 0.$$

By Theorems 5.1.2 and 5.1.4, we know that $\ker L_0^\downarrow \text{BD}_p F([a, b]) \cong H_0^\downarrow \text{BD}_p F([a, b])$ and that elements in $\ker L_0^\downarrow \text{BD}_p F([a, b])$ provide canonical cycle representatives for $H_0^\downarrow \text{BD}_p F([a, b])$. Thus, we focus on the degree-zero Möbius homology Laplacian, as it enables an analytic study of the birth-death spaces themselves.

Since $L_0^\downarrow \text{BD}_p F([a, b])$ is self-adjoint, the spectral theorem yields

$$C_0^\downarrow \text{BD}_p F([a, b]) = \text{BD}_p F([a, b]) = \text{im } L_0^\downarrow \text{BD}_p F([a, b]) \oplus \ker L_0^\downarrow \text{BD}_p F([a, b]).$$

In this chapter, we will study both the kernel and the image of $L_0^\downarrow \text{BD}_p F([a, b])$. Indeed, we will see that the entire spectral decomposition of $L_0^\downarrow \text{BD}_p F([a, b])$ reflects geometric and combinatorial properties of the persistence diagram $\partial \text{bd}_p F$.

In order to make computations and prove theorems about the non-zero spectrum of the Möbius homology Laplacian of the birth-death module, it will often be useful to write the matrix representation of $L_0^\downarrow \text{BD}_p F([a, b])$ in terms of some basis for $\text{BD}_p F[a, b]$. Recall that we also relied heavily on the matrix representation of $L_1^\downarrow M(c)$ for M a one-parameter module in Section 5.1.1. Thus, before beginning our analysis, we provide an in-depth description of the matrix representation for $L_0^\downarrow \text{BD}_p F([a, b])$.

6.3 Matrix Representation of the Degree-Zero Laplacian

We saw in Section 2.1 how to construct the adjoint of a linear map using Gram matrices. In this section, we review that construction, specifically for the Möbius homology Laplacian $L_0^\dagger \text{BD}_p F([a, b])$. Indeed, the construction given in this section will be useful for both computations and proofs later on.

Let P be a finite poset with $\bar{P} = P \cup \{\infty\}$. Denote by $(\text{Dgm}(P), \preceq)$ the poset of symbolic persistence intervals $[a, b]$ where $a \in P$ and $b \in \bar{P}$. Let $F : P \rightarrow \Delta K$ be a filtration of a finite simplicial complex K , and let $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be the associated birth-death module. Fix a persistence interval $[a, b]$.

Suppose $\dim \text{BD}_p F([a, b]) = m$, and let $\{z_1, z_2, \dots, z_m\}$ be cycles that form a basis for $\text{BD}_p F([a, b])$. Observe that the basis given for $\text{BD}_p F([a, b])$ is likely not an orthonormal basis. Define the Gram matrix $M_{a,b}$ for $\text{BD}_p F([a, b])$ by $(M_{a,b})_{i,j} = \langle z_i, z_j \rangle = \sum_{\sigma \in K^{[p]}} a_\sigma^i a_\sigma^j$ for basis elements $z_i = \sum_{\sigma \in K^{[p]}} a_\sigma^i \sigma$ and $z_j = \sum_{\sigma \in K^{[p]}} a_\sigma^j \sigma$.

Let $\{[c_i, d_i]\}_{i=1}^N$ be the persistence intervals strictly less than $[a, b]$ in $(\text{Dgm}(P), \preceq)$. The Gram matrix M for the chain space $C_1^\dagger \text{BD}_p F([a, b]) = \bigoplus_{i=1}^N \text{BD}_p F([c_i, d_i])$ is a block diagonal matrix with blocks given by M_{c_i, d_i} . Indeed, the block structure here reflects the orthogonality of the direct sum $\bigoplus_{i=1}^N \text{BD}_p F([c_i, d_i])$.

Consider the inclusion $f_i : \text{BD}_p F([c_i, d_i]) \rightarrow \text{BD}_p F([a, b])$. If $\dim \text{BD}_p F([c_i, d_i]) = n_i$ and the set of cycles $\{w_1, w_2, \dots, w_{n_i}\}$ forms a basis for $\text{BD}_p F([c_i, d_i])$, denote by A_i the $m \times n_i$ matrix representation of f_i . Suppose $f_i(w_j) = \sum_{k=1}^m a_k z_k$. Then, column j of A_i is given by the m -vector with entries $\{a_k\}_{k=1}^m$.

The map $\bigoplus_{i=1}^N \text{BD}_p F([c_i, d_i]) \xrightarrow{\partial_1^\dagger} \text{BD}_p F([a, b])$ can be written as the block matrix

$$\partial_1^\dagger = \begin{pmatrix} A_1 & A_2 & \cdots & A_N \end{pmatrix}.$$

Thus, the Laplacian $L_0^\downarrow \text{BD}_p F([a, b])$ can be written as:

$$\begin{aligned} L_0^\downarrow \text{BD}_p F([a, b]) &= \partial_1^\downarrow (\partial_1^\downarrow)^* = \partial_1^\downarrow M^{-1} (\partial_1^\downarrow)^T M_{a,b} \\ &= \begin{pmatrix} A_1 & A_2 & \cdots & A_N \end{pmatrix} \begin{pmatrix} M_{c_1, d_1}^{-1} & \cdots & \cdots \\ & M_{c_2, d_2}^{-1} & \cdots \\ \cdots & \cdots & M_{c_N, d_N}^{-1} \end{pmatrix} \begin{pmatrix} A_1^T \\ A_2^T \\ \cdots \\ A_N^T \end{pmatrix} M_{a,b}. \end{aligned}$$

We will use this matrix representation and its block structure to help us understand and compute both the kernel and the image of $L_0^\downarrow \text{BD}_p F([a, b])$.

6.4 Harmonic Representatives and Newly Added Simplices

In this section, we will explore the kernel of $L_0^\downarrow \text{BD}_p F([a, b])$. Again, by Theorem 5.1.4, we know that $\ker L_0^\downarrow \text{BD}_p F([a, b])$ provides cycle representatives for $H_0^\downarrow \text{BD}_p F([a, b])$. Moreover, the representatives in $\ker L_0^\downarrow \text{BD}_p F([a, b])$ have a nice geometric interpretation: by Theorem 5.1.5, elements in $\ker L_0^\downarrow \text{BD}_p F([a, b])$ have minimal energy (that is, they have minimal squared norm). In the case of the birth-death module for filtrations, we can say more about the geometric properties of elements in $\ker L_0^\downarrow \text{BD}_p F([a, b])$. Inspired by Basu and Cox's work on harmonic persistent homology [19], our choice of an inner product structure on each $\text{BD}_p F([a, b])$ results in harmonic representatives that maximize the contribution of "newly added" simplices. In the definitions and propositions to follow, we make this claim precise.

Recall Definition 2.2.3: the *support* of a chain $z = \sum_{\sigma \in K^{[p]}} a_\sigma \sigma \in \text{BD}_p F([a, b]) \subseteq C_p(K)$ is given by the collection $\text{supp}(z) = \{\sigma \in K^{[p]} \mid a_\sigma \neq 0\}$. That is, the support of a chain z consists of all p -simplices with non-zero coefficient in z .

Definition 6.4.1: Let $\sum_{[c, d] \prec [a, b]} \text{BD}_p F([c, d])$ denote the internal sum of the birth-death spaces $\text{BD}_p F([c, d])$, and let $A = \text{BD}_p F([a, b]) \setminus \sum_{[c, d] \prec [a, b]} \text{BD}_p F([c, d])$. Then, given a symbolic per-

sistence interval $[a, b) \in \text{Dgm}(P)$, define

$$\nu[a, b) = \bigcap_{z \in A} \text{supp}(z).$$

We call $\nu[a, b)$ the **set of newly added simplices** associated to the persistence interval $[a, b)$. Then, the **relative new content** of a chain $z = \sum_{\sigma \in K^{[p]}} a_\sigma \sigma \in \text{BD}_p F([a, b))$ is defined as

$$\text{New}(z) = \left(\frac{\sum_{\sigma \in \nu[a, b)} a_\sigma^2}{\sum_{\sigma \in K^{[p]}} a_\sigma^2} \right)^{1/2}.$$

The definitions of the set of newly added simplices associated to an interval and the relative new content of a chain are adapted from Basu and Cox's set of essential simplices and relative essential content, respectively. See [19, Definitions 4.4 and 4.5].

Intuitively, the relative new content of a chain z in $\text{BD}_p F([a, b))$ quantifies the proportion of simplices in z that are "newly added:" that is, they do not appear in previous birth-death spaces. Observe that, if $\text{BD}_p F([a, b)) \setminus \sum_{[c, d) \prec [a, b)} \text{BD}_p F([c, d)) = \emptyset$, then $\nu[a, b) = \emptyset$ as well. In this case, the relative new content of *any* chain $z \in \text{BD}_p F([a, b))$ will be zero.

Recall Definition 5.1.3: given the natural surjection $\varphi : \text{BD}_p F([a, b)) \rightarrow H_0^\downarrow \text{BD}_p F([a, b))$, the *space of cycle representatives* for an equivalence class $[x] \in H_0^\downarrow \text{BD}_p F([a, b))$ is given by $\varphi^{-1}([x])$.

Theorem 6.4.2: Let K be a finite simplicial complex and P a finite poset. Let $F : P \rightarrow \Delta K$ be a filtration, and fix $p \geq 0$. Consider a persistence interval $[a, b) \in \text{Dgm}(P)$ so that $\text{BD}_p F([a, b))$ has the property $\dim \text{BD}_p F([a, b)) - \dim \sum_{[c, d) \prec [a, b)} \text{BD}_p F([c, d)) = 1$.

Let $z_0 \in \ker L_0^\downarrow \text{BD}_p F([a, b))$, and let $z \in \varphi^{-1}([z_0])$ be any other cycle representative. Then,

$$\text{New}(z_0) \geq \text{New}(z).$$

Proof. The argument is nearly identical to that of [19, Theorem 4.8], as their theorem was the inspiration for this result. We provide the proof for our specific context here.

Recall that the boundary operator $\partial_1^\downarrow : C_1^\downarrow \text{BD}_p F([a, b]) \rightarrow C_0^\downarrow \text{BD}_p F([a, b])$ is the direct sum of inclusions $\text{BD}_p F([c, d] \preceq [a, b])$ for $[c, d] \prec [a, b]$. Thus, $\text{im } \partial_1^\downarrow = \sum_{[c, d] \prec [a, b]} \text{BD}_p F([c, d])$.

Let $z_0 \in \ker L_0^\downarrow \text{BD}_p F([a, b])$, and let $z = \sum_{\sigma \in K[p]} a_\sigma \sigma \in \varphi^{-1}([z_0])$ be a cycle representative of the same homology class $[z_0] \in H_0^\downarrow \text{BD}_p F([a, b])$. Define $z_1 = \sum_{\sigma \in \nu[a, b]} a_\sigma \sigma$ to be only the contribution of "newly added" simplices, and define $z_2 = z - z_1$.

Restated in our context, [19, Lemma 4.10] states that $\sigma \in \nu[a, b]$ if and only if $\sigma \in (\text{im } \partial_1^\downarrow)^\perp$. Thus, each $\sigma \in \text{supp}(z_1)$ is in $(\text{im } \partial_1^\downarrow)^\perp$, and therefore $z_1 \in (\text{im } \partial_1^\downarrow)^\perp$. Thus

$$\begin{aligned} z_0 &= \text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z) && \text{(by Theorem 5.1.4)} \\ &= \text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_1) + \text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2) \\ &= z_1 + \text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2) && \text{(since } z_1 \in (\text{im } \partial_1^\downarrow)^\perp\text{)}. \end{aligned}$$

Now, we will compute $\text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2)$. For $\sigma \in \nu[a, b]$, $\sigma \notin \text{supp}(z_2)$, and hence [19, Lemma 4.10] yields $\sigma \notin \text{supp}(\text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2))$. In particular,

$$\text{supp}(z_1) \cap \text{supp}(\text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2)) = \emptyset \implies \langle z_1, \text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2) \rangle = 0. \quad (6.1)$$

We will now compare relative new content.

$$\begin{aligned} \text{New}(z)^2 &= \frac{\|z_1\|^2}{\|z\|^2} \\ &= \frac{\|z_1\|^2}{\|z_1\|^2 + \|z_2\|^2} && \text{(since } \langle z_1, z_2 \rangle = 0 \text{ by construction)} \\ &\leq \frac{\|z_1\|^2}{\|z_1\|^2 + \|\text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2)\|^2} && \text{(orthogonal projection does not increase norm)} \\ &= \frac{\|z_1\|^2}{\|z_0\|^2} && \text{(by Equation 6.1)}. \end{aligned}$$

Since $\sigma \in \nu[a, b] \implies \sigma \notin \text{supp}(\text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2))$, and since $z_0 = z_1 + \text{proj}_{(\text{im } \partial_1^\downarrow)^\perp}(z_2)$, the only simplices in $\nu[a, b]$ appearing in z_0 are those in z_1 . That is, $\text{New}(z_0) = \frac{\|z_1\|^2}{\|z_0\|^2}$ and we obtain

$$\text{New}(z)^2 \leq \text{New}(z_0)^2,$$

as desired. □

Example 6.4.3: Consider the filtration $F : T \rightarrow \Delta K$ in Figure 6.1 and its birth-death module $\text{BD}_1 F$ as shown in Figure 6.2.

Let the birth-death space $\text{BD}_1 F([2, 4])$ be generated by the cycles $z_1 = [v_1 v_2] + [v_2 v_4] - [v_1 v_4]$ and $z_2 = [v_2 v_3] + [v_3 v_4] - [v_2 v_4]$. The Möbius chain space $C_1^\downarrow \text{BD}_1 F([2, 4])$, then, is given by the direct sum $\text{BD}_1 F([1, 3]) \oplus \text{BD}_1 F([1, 4]) \oplus \text{BD}_1 F([2, 3])$ as these are the only persistence intervals strictly less than $[2, 4]$ in the product order so that $\text{BD}_1 F([c, d]) \neq 0$. All three spaces are generated by z_1 . Thus, the internal sum $\text{BD}_1 F([1, 3]) + \text{BD}_1 F([1, 4]) + \text{BD}_1 F([2, 3])$ has dimension 1 and $\text{BD}_1 F([2, 4])$ satisfies the requirements of Theorem 6.4.2.

The end of the Möbius chain complex $C_\bullet^\downarrow \text{BD}_1 F([2, 4])$ is given by

$$\dots \xrightarrow{\partial_2^\downarrow} \mathbb{R} \oplus \mathbb{R} \oplus \mathbb{R} \xrightarrow{\partial_1^\downarrow = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 0 & 0 \end{pmatrix}} \mathbb{R}^2 \rightarrow 0.$$

We now compute the Laplacian $L_0^\downarrow \text{BD}_1 F([2, 4])$. The Gram matrix $M_{2,4}$ is given by $\begin{pmatrix} 3 & -1 \\ -1 & 3 \end{pmatrix}$.

The Gram matrix on $C_1^\downarrow \text{BD}_1 F([2, 4])$ is $M = 3 \cdot I_3$. Thus, the degree-zero Möbius homology Laplacian for $\text{BD}_1 F([2, 4])$ is

$$L_0^\downarrow \text{BD}_1 F([2, 4]) = \partial_1^\downarrow M^{-1} (\partial_1^\downarrow)^T M_{2,4} = \begin{pmatrix} 3 & -1 \\ 0 & 0 \end{pmatrix}.$$

It is easy to check that $\ker L_0^\downarrow \text{BD}_1 F([2, 4])$ is generated by $\begin{pmatrix} 1/3 \\ 1 \end{pmatrix}$, which represents the cycle

$$\begin{aligned} z &= \frac{1}{3} ([v_1 v_2] + [v_2 v_4] - [v_1 v_4]) + ([v_2 v_3] + [v_3 v_4] - [v_2 v_4]) \\ &= \frac{1}{3} [v_1 v_2] - \frac{1}{3} [v_1 v_4] + [v_2 v_3] - \frac{2}{3} [v_2 v_4] + [v_3 v_4]. \end{aligned}$$

The energy of z is $\|z\|^2 = 8/3$. The energy of z_2 , a different cycle representative for the bar $[2, 4]$, is given by $\|z_2\|^2 = 3$. Indeed, $\|z\|^2 \leq \|z_2\|^2$.

We also compute the relative new content of z and z_2 . First, observe that the set of newly added simplices for $[2, 4]$ is given by $\nu[2, 4] = \{[v_2 v_3], [v_3 v_4]\}$. Then,

$$\begin{aligned} \text{New}(z) &= \left(\frac{1^2 + 1^2}{8/3} \right)^{1/2} = \sqrt{3/4}, \text{ and} \\ \text{New}(z_2) &= \left(\frac{1^2 + 1^2}{3} \right)^{1/2} = \sqrt{2/3}. \end{aligned}$$

Indeed, $\text{New}(z) \geq \text{New}(z_2)$. Although the supports of both z and z_2 contain $\nu[2, 4]$, the relative new simplices make up a larger proportion of the support of $z \in \ker L_0^\downarrow \text{BD}_1 F([2, 4])$.

Remark 6.4.4: Let $F : P \rightarrow \Delta K$ be a filtration, and let $[a, b] \in (\text{Dgm}(P), \preceq)$ be such that $\dim \text{BD}_p F([a, b]) = 1$ and $\dim \text{BD}_p F([c, d]) = 0$ for all $[c, d] \prec [a, b]$. Since $\text{BD}_p F([a, b])$ satisfies all requirements of Theorem 6.4.2, elements in $\ker L_0^\downarrow \text{BD}_p F([a, b])$ will provide cycle representatives that emphasize newly added content. In this case, there are some subtleties worth pointing out. Computing $C_\bullet^\downarrow \text{BD}_p F([a, b])$ yields

$$\dots \rightarrow 0 \xrightarrow{\partial_1^\downarrow} \text{BD}_p F([a, b]) \rightarrow 0.$$

Thus, $L_0^\downarrow \text{BD}_p F([a, b]) = 0$. Choose a generator $x \in \text{BD}_p F([a, b])$. Then, $x \in \ker L_0^\downarrow \text{BD}_p F([a, b])$. It is vacuously true that x maximizes relative new content. Since there are *no* cycles present in previous birth-death spaces, *all* simplices present in x are new. Thus, $\text{New}(x) = 1$.

6.5 Combinatorial Spectral Properties of the Degree-Zero Laplacian

Let $F : P \rightarrow \Delta K$ be a filtration, and let $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be its associated birth-death module. In the previous section, we investigated the *kernel* of the degree-zero Möbius homology Laplacian $L_0^\downarrow \text{BD}_p F([a, b])$. In particular, Theorem 6.4.2 states that, when $\dim \ker L_0^\downarrow \text{BD}_p F([a, b]) = 1$, the elements in $\ker L_0^\downarrow \text{BD}_p F([a, b])$ have nice geometric properties. Namely, they have minimal energy, which in this case means they emphasize *relative new content*.

The Möbius homology Laplacian, however, provides additional structure *beyond* its kernel. Much like classical combinatorial Hodge Laplace operators, we expect that the non-zero spectrum of $L_0^\downarrow \text{BD}_p F([a, b])$ will reflect the combinatorial structure of the persistence diagram associated to F , the underlying object in question here. Indeed, the number of non-zero eigenvalues of $L_0^\downarrow \text{BD}_p F([a, b])$ is the same as the rank of ∂_1^\downarrow . Recall that these are the cycles in $\text{BD}_p F([a, b])$ that are born at c and die at d for $[c, d] \prec [a, b]$.

However, observe that the results from Section 5.1.1 do not apply here. Indeed, even when P is totally ordered, $\text{Dgm}(P)$ is often not.³ Thus, new methods are needed to investigate the spectrum of $L_0^\downarrow \text{BD}_p F([a, b])$.

6.5.1 Spectrum in the One-Parameter Setting

We again start by restricting to the simpler, one-parameter case. Let $F : T \rightarrow \Delta K$ be a filtration indexed over a totally ordered set T and let $\text{BD}_p F : (\text{Int } T, \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be the birth-death module associated to F . In this setting, under the condition that the cycles in the birth-death module do not *branch*, we can simplify the matrix representation of $L_0^\downarrow \text{BD}_p F([a, b])$ from Section 6.3 to extract its corresponding eigenvalues.

³For T a totally ordered set, the interval poset $(\text{Int } T, \preceq)$ is only totally ordered when T has two or fewer elements.

Definition 6.5.1: Let $F : T \rightarrow \Delta K$ be a filtration, let $(\text{Dgm}(T), \preceq)$ denote the poset of persistence intervals with the product order, and let $\text{BD}_p F : (\text{Dgm}(T), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be the birth-death module associated to F . Fix a persistence interval $[a, b)$. We say a cycle $z \in \text{BD}_p F([a, b))$ **branches** if it can be written as a sum $z = z_1 + z_2$ such that

- The cycle z_1 first appears in $\text{BD}_p F[c_1, d_1)$. That is,

$$[c_1, d_1) = \min_{\preceq} \{[c, d) \preceq [a, b) \mid z_1 \in \text{BD}_p F([c, d))\},$$

- The cycle z_2 first appears in $\text{BD}_p F[c_2, d_2)$. That is,

$$[c_2, d_2) = \min_{\preceq} \{[c, d) \preceq [a, b) \mid z_2 \in \text{BD}_p F([c, d))\},$$

- The persistence intervals $[c_1, d_1)$ and $[c_2, d_2)$ are incomparable in $(\text{Dgm}(T), \preceq)$.

Theorem 6.5.2: Let $F : T \rightarrow \Delta K$ be a filtration indexed over a finite totally ordered set T , let $\text{BD}_p F : (\text{Dgm}(T), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be the birth-death module associated to F . Fix $[a, b) \in \text{Dgm}(T)$. Suppose $\text{BD}_p F([a, b))$ does not contain any cycles that branch, and let $\{z_1, z_2, \dots, z_m\}$ be a basis that is compatible with each space $\text{BD}_p F([c, d))$ so that $[c, d) \preceq [a, b)$. That is, we can write $\text{BD}_p F([c, d)) = \text{span}\{z_i \mid [\alpha_i, \beta_i) \preceq [c, d)\}$. For every z_i , define

$$[\alpha_i, \beta_i) = \min_{\preceq} \{[c, d) \preceq [a, b) \mid z_i \in \text{BD}_p F([c, d))\}.$$

That is, $\text{BD}_p F([\alpha_i, \beta_i))$ is the first birth-death space that z_i appears in. This is guaranteed to be unique since we are considering a one-parameter filtration. Then,

$$\lambda = \text{len}\left([\alpha_i, \beta_i), [a, b)\right) = \#\{[c, d) \in \text{Dgm}(T) \mid [\alpha_i, \beta_i) \preceq [c, d) \prec [a, b)\}.$$

is an eigenvalue of $L_0^\downarrow \text{BD}_p F([a, b))$.

Proof. We prove the proposition by showing that we can write the Laplacian $L_0^\downarrow \text{BD}_p F([a, b])$ as an upper triangular matrix, where

$$(L_0^\downarrow \text{BD}_p F([a, b]))_{i,i} = \lambda_i = \{\# \text{of birth-death spaces } \text{BD}_p F[c_j, d_j] \text{ that contain } z_i\}.$$

First, order $\{z_1, z_2, \dots, z_m\}$ according to the order in which they appear in F . If two appear at the same time, order arbitrarily. Then, order each $\text{BD}_p F([c_j, d_j])$ according to dimension. Again, if two spaces have the same dimension, order arbitrarily.

Since no basis element of $\text{BD}_p F([a, b])$ branches, with our choice of basis for $\text{BD}_p F([a, b])$ we can write each inclusion $\text{BD}_p F([c_j, d_j] \preceq [a, b])$ as a block matrix of the form $A_j = \begin{pmatrix} I_{n_j} \\ 0_{(m-n_j) \times n_j} \end{pmatrix}$.

Thus, we can write

$$\begin{aligned} (\partial_1^\downarrow)^T M_{a,b} &= \begin{pmatrix} A_1^T \\ A_2^T \\ \dots \\ A_N^T \end{pmatrix} M_{a,b} \\ &= \begin{pmatrix} M_{a,b}|_{[c_1, d_1]} \\ M_{a,b}|_{[c_2, d_2]} \\ \vdots \\ M_{a,b}|_{[c_N, d_N]} \end{pmatrix}, \end{aligned}$$

where $M_{a,b}|_{[c_j, d_j]}$ is $M_{a,b}$ with the rows corresponding to $z_i \in \text{BD}_p F([a, b]) \setminus \text{BD}_p F([c_j, d_j])$ deleted.

Then, multiplying by M^{-1} , we obtain

$$M^{-1} (\partial_1^\downarrow)^T M_{a,b} = \begin{pmatrix} M_{c_1, d_1}^{-1} & \dots & \dots \\ & M_{c_2, d_2}^{-1} & \dots \\ \dots & \dots & M_{c_N, d_N}^{-1} \end{pmatrix} \begin{pmatrix} M_{a,b}|_{[c_1, d_1]} \\ M_{a,b}|_{[c_2, d_2]} \\ \vdots \\ M_{a,b}|_{[c_N, d_N]} \end{pmatrix}$$

$$= \begin{pmatrix} M_{c_1, d_1}^{-1} M_{a, b}|_{[c_1, d_1]} \\ M_{c_2, d_2}^{-1} M_{a, b}|_{[c_2, d_2]} \\ \vdots \\ M_{c_N, d_N}^{-1} M_{a, b}|_{[c_N, d_N]} \end{pmatrix}$$

Observe that, for every j , the product $M_{c_j, d_j}^{-1} M_{a, b}|_{[c_j, d_j]}$ decomposes into a block matrix of the form $\begin{pmatrix} I_{n_j \times n_j} & *_{n_j \times (m-n_j)} \end{pmatrix}$ where the matrix $*_{n_j \times (m-n_j)}$ is potentially non-zero. Thus, we obtain

$$\begin{aligned} \partial_1^\downarrow M^{-1} (\partial_1^\downarrow)^T M_{a, b} &= \begin{pmatrix} A_1 & A_2 & \cdots & A_N \end{pmatrix} \begin{pmatrix} M_{c_1, d_1}^{-1} M_{a, b}|_{[c_1, d_1]} \\ M_{c_2, d_2}^{-1} M_{a, b}|_{[c_2, d_2]} \\ \vdots \\ M_{c_N, d_N}^{-1} M_{a, b}|_{[c_N, d_N]} \end{pmatrix} \\ &= \begin{pmatrix} A_1 & A_2 & \cdots & A_N \end{pmatrix} \begin{pmatrix} I_{n_1} & *_{n_1 \times (m-n_1)} \\ I_{n_2} & *_{n_2 \times (m-n_2)} \\ \vdots & \vdots \\ I_{n_N} & *_{n_N \times (m-n_N)} \end{pmatrix} \\ &= \begin{pmatrix} I_{n_1} & & & \\ 0_{(m-n_1) \times n_1} & I_{n_2} & & \\ & 0_{(m-n_2) \times n_2} & \cdots & \\ & & & I_{n_N} \\ & & & 0_{(m-n_N) \times n_N} \end{pmatrix} \begin{pmatrix} I_{n_1} & *_{n_1 \times (m-n_1)} \\ I_{n_2} & *_{n_2 \times (m-n_2)} \\ \vdots & \vdots \\ I_{n_N} & *_{n_N \times (m-n_N)} \end{pmatrix} \end{aligned}$$

The block matrix structure above implies that $L_0^\downarrow \text{BD}_p F([a, b])$ is upper-triangular. Indeed, whenever $i > j$, $(L_0^\downarrow \text{BD}_p F([a, b]))_{i, j} = 0$. Thus, the eigenvalues are given by the diagonal entries. The block structure also implies that $(L_0^\downarrow \text{BD}_p F([a, b]))_{i, i}$ counts the number of symbolic persistence intervals $[c, d] \prec [a, b]$ such that $z_i \in \text{BD}_p F[c, d]$. That is, if the first birth-death space that z_i appears in is $\text{BD}_p F[\alpha, \beta]$, then $(L_0^\downarrow \text{BD}_p F([a, b]))_{i, i} = \text{len}([\alpha, \beta), [a, b))$. \square

The theorem above is consistent with our results on $\ker L_0^\downarrow \text{BD}_p F([a, b])$. Suppose $\text{BD}_p F([a, b])$ satisfies the conditions of Theorem 6.5.2, and let $z_i \in \text{BD}_p F([a, b])$ be such that $z_i \notin \text{BD}_p F[c, d]$ for all $[c, d] \prec [a, b]$. That is, $\text{BD}_p F([a, b])$ is the first birth-death space that z_i appears in. Then, $\text{len}([a, b], [a, b]) = 0$ is an eigenvalue of $L_0^\downarrow \text{BD}_p F([a, b])$.

Example 6.5.3: Consider the filtration $F : T \rightarrow \Delta K$ in Figure 6.1 and its birth-death module $\text{BD}_1 F$ as shown in Figure 6.2.

Observe that no cycles in $\text{BD}_1 F([2, 4])$ branch, meaning the conditions of Theorem 6.5.2 are satisfied. The Möbius homology Laplacian $L_0^\downarrow \text{BD}_1 F([2, 4])$ is computed in Example 6.4.3. Indeed, observe that $L_0^\downarrow \text{BD}_1 F([2, 4])$ is upper-triangular and has eigenvalues 0 and 3. The eigenvalue 3 is given by $\text{len}([1, 3], [2, 4]) = \#\{[1, 3], [1, 4], [2, 3]\}$.

It is worth noting that the conditions in Theorem 6.5.2 are sufficient but not necessary. In Appendix A, we provide additional examples of one-parameter filtrations and birth-death spaces that do not satisfy the branching condition. One example will result in a Möbius homology Laplacian with integer eigenvalues while the other will have non-integer eigenvalues. Furthermore, the definition of a branched cycle is not well-defined in general for multiparameter filtrations. Nonetheless, there are multiparameter filtrations who admit Möbius homology Laplacians that achieve the result of Theorem 6.5.2. Again, see Appendix A for an example.

Even without the branching condition in Theorem 6.5.2, when considering one-parameter filtrations, we know that the eigenvalues of $L_0^\downarrow \text{BD}_p F([a, b])$ are related to the points in the persistence diagram below $[a, b]$.

Proposition 6.5.4: The number of non-zero eigenvalues of $L_0^\downarrow \text{BD}_p F([a, b])$ is the same as the number of points (with multiplicity) $[a, b]$ in the persistence diagram of $\text{BD}_p F$ for $F : T \rightarrow \Delta K$ a one-parameter filtration.

Proof. The number of non-zero eigenvalues is given by $\dim \text{im} \left(L_0^\downarrow \text{BD}_p F([a, b]) \right)$. We know that $\text{im} L_0^\downarrow \text{BD}_p F([a, b]) = \text{im} \partial_1^\downarrow$. The map ∂_1^\downarrow is the direct sum $\bigoplus_{[c, d] \prec [a, b]} \text{BD}_p F([c, d] \preceq [a, b])$, where each $\text{BD}_p F([c, d] \preceq [a, b])$ is the inclusion $\text{BD}_p F[c, d] \hookrightarrow \text{BD}_p F[a, b]$. Thus, $\text{im} \partial_1^\downarrow$ con-

sists of all the cycles that are born at c and die at d for $[c, d) \prec [a, b)$. When F is a one-parameter filtration, these cycles are in one-to-one correspondence with the intervals $[c, d)$ such that $\partial \text{bd}_p F[c, d) \neq 0$ (counted with multiplicity). \square

6.5.2 General Spectral Properties

Although we will not provide a complete characterization for the eigenvalues of $L_0^\downarrow \text{BD}_p F([a, b))$ for filtrations $F : P \rightarrow \Delta K$ indexed over a finite poset P , we can still extract spectral information from the *characteristic polynomial* of $L_0^\downarrow \text{BD}_p F([a, b))$. In particular, we will compute the trace exactly and provide a lower bound for the determinant. These properties will, as we expect, reflect the combinatorial structure of the persistence diagram.

Recall that, to compute the eigenvalues of $L_0^\downarrow \text{BD}_p F([a, b))$, we solve for the roots of the characteristic polynomial $\det(\lambda I_m - L_0^\downarrow \text{BD}_p F([a, b))$ where $m = \dim \text{BD}_p F([a, b))$. It is known that the coefficient in front of λ^{m-1} is given by $-\text{tr } L_0^\downarrow \text{BD}_p F([a, b))$. In fact, we can compute the trace exactly.

Theorem 6.5.5: Let $F : P \rightarrow \Delta K$ be a filtration, and let $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be its associated birth-death module for a fixed $p \geq 0$. Fix $[a, b) \in \text{Dgm}(P)$, and let $L_0^\downarrow \text{BD}_p F([a, b))$ denote the associated Möbius homology Laplacian. Then,

$$\text{tr } L_0^\downarrow \text{BD}_p F([a, b)) = \sum_{[c, d) \prec [a, b)} \text{bd}_p F([c, d)).$$

Proof. We begin by changing basis so that, for every interval $[c, d) \preceq [a, b)$, $\text{BD}_p F([c, d))$ has an orthonormal basis. In particular, this means the direct sum $\bigoplus_{[c, d) \prec [a, b)} \text{BD}_p F([c, d))$ also has an orthonormal basis. Denote by U^{-1} and V^{-1} the change of basis matrices for $\text{BD}_p F([a, b))$ and $\bigoplus_{[c, d) \prec [a, b)} \text{BD}_p F([c, d))$, respectively. Changing basis changes the matrix representation of the Laplacian $L_0^\downarrow \text{BD}_p F([a, b))$ as follows. Denote by $\widetilde{L}_0^\downarrow \text{BD}_p F([a, b))$ the new representation

given by

$$\begin{aligned}\widetilde{L}_0^\downarrow \text{BD}_p F([a, b]) &= U^{-1} \partial_1^\downarrow V V^{-1} (\partial_1^\downarrow)^* U \\ &= U^{-1} \partial_1^\downarrow V (U^{-1} \partial_1^\downarrow V)^*,\end{aligned}$$

where the second equality follows from the fact that U and V are orthogonal matrices. Furthermore, since we have changed to orthonormal bases,

$$\widetilde{L}_0^\downarrow \text{BD}_p F([a, b]) = U^{-1} \partial_1^\downarrow V (U^{-1} \partial_1^\downarrow V)^T.$$

Let $\widetilde{\partial}_1^\downarrow = U^{-1} \partial_1^\downarrow V$. Observe that changing basis does not change eigenvalues (recall Remark 2.1.12).

Suppose $\widetilde{L}_0^\downarrow \text{BD}_p F([a, b])$ has eigenvector \vec{v} with eigenvalue λ . Then,

$$\begin{aligned}\widetilde{L}_0^\downarrow \text{BD}_p F([a, b]) \vec{v} = \lambda \vec{v} &\implies U^{-1} \partial_1^\downarrow V V^{-1} (\partial_1^\downarrow)^* U \vec{v} = \lambda \vec{v} \\ &\implies \partial_1^\downarrow (\partial_1^\downarrow)^* U \vec{v} = U (\lambda \vec{v}) \\ &\implies L_0^\downarrow \text{BD}_p F([a, b]) (U \vec{v}) = \lambda (U \vec{v}).\end{aligned}$$

In particular, $\text{tr } L_0^\downarrow \text{BD}_p F([a, b]) = \text{tr } \widetilde{L}_0^\downarrow \text{BD}_p F([a, b])$.

Now, $\widetilde{\partial}_1^\downarrow$ has block structure given by $\begin{pmatrix} \widetilde{A}_1 & \widetilde{A}_2 & \cdots & \widetilde{A}_N \end{pmatrix}$ (see Section 6.3). Thus, we can write $\widetilde{L}_0^\downarrow \text{BD}_p F([a, b]) = \widetilde{A}_1 \widetilde{A}_1^T + \widetilde{A}_2 \widetilde{A}_2^T + \cdots + \widetilde{A}_N \widetilde{A}_N^T$. Each summand $\widetilde{A}_i \widetilde{A}_i^T$ is a projection matrix that projects onto $\text{BD}_p F([c_i, d_i]) \subseteq \text{BD}_p F([a, b])$. Thus, $\text{tr } \widetilde{A}_i \widetilde{A}_i^T = \dim \text{BD}_p F([c_i, d_i])$ and, because the trace distributes over matrix addition,

$$\text{tr } L_0^\downarrow \text{BD}_p F([a, b]) = \sum_{[c, d] \prec [a, b]} \dim \text{BD}_p F([c, d])$$

as desired. □

Furthermore, the methods utilized in the proof of Theorem 6.5.5 will also provide insight into the determinant of $L_0^\downarrow \text{BD}_p F([a, b])$.

Theorem 6.5.6: Let $F : P \rightarrow \Delta K$ be a filtration, and let $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be its associated birth-death module for a fixed $p \geq 0$. Fix $[a, b] \in \text{Dgm}(P)$, and let $L_0^\downarrow \text{BD}_p F([a, b])$ denote the associated Möbius homology Laplacian. Let $m = \dim \text{BD}_p F([a, b])$. Then,

$$(\det L_0^\downarrow \text{BD}_p F([a, b]))^{1/m} \geq \# \{[c, d] \mid [c, d] \prec [a, b] \text{ and } \text{bd}_p F([c, d]) = \text{bd}_p F([a, b])\}.$$

Proof. First, change basis as in the proof of Theorem 6.5.5. Because changing basis does not change the determinant,

$$\det L_0^\downarrow \text{BD}_p F([a, b]) = \det \widetilde{L}_0^\downarrow \text{BD}_p F([a, b]) = \det(\widetilde{A}_1 \widetilde{A}_1^T + \widetilde{A}_2 \widetilde{A}_2^T + \cdots + \widetilde{A}_N \widetilde{A}_N^T).$$

By the Minkowski determinant inequality (see [47, Part II Section 4.1.8]),

$$\begin{aligned} & (\det(\widetilde{A}_1 \widetilde{A}_1^T + \widetilde{A}_2 \widetilde{A}_2^T + \cdots + \widetilde{A}_N \widetilde{A}_N^T))^{1/m} \\ & \geq (\det(\widetilde{A}_1 \widetilde{A}_1^T))^{1/m} + (\det(\widetilde{A}_2 \widetilde{A}_2^T))^{1/m} + \cdots + (\det(\widetilde{A}_N \widetilde{A}_N^T))^{1/m}. \end{aligned} \tag{6.2}$$

Since each $\widetilde{A}_i \widetilde{A}_i^T$ is a projection matrix, $\det(\widetilde{A}_i \widetilde{A}_i^T) = 0$ if $\dim \text{BD}_p F([c_i, d_i]) < \dim \text{BD}_p F([a, b])$ whereas $\det(\widetilde{A}_i \widetilde{A}_i^T) = 1$ if $\dim \text{BD}_p F([c_i, d_i]) = \dim \text{BD}_p F([a, b])$. Thus, we conclude. \square

Intuitively, the $1/\text{bd}_p F([a, b])$ -th power of the determinant of $L_0^\downarrow \text{BD}_p F([a, b])$ is bounded from below by the number of intervals $[c, d] \prec [a, b]$ where $\text{bd}_p F([c, d]) = \text{bd}_p F([a, b])$.

The lower bound for $(\det L_0^\downarrow \text{BD}_p F([a, b]))^{1/m}$ given by Theorem 6.5.6 is less informative as the number of intervals $[c, d] \prec [a, b]$ grows. Indeed, Equation 6.2 follows inductively:

$$\begin{aligned} & (\det(\widetilde{A}_1 \widetilde{A}_1^T + \widetilde{A}_2 \widetilde{A}_2^T + \cdots + \widetilde{A}_N \widetilde{A}_N^T))^{1/m} \\ & \geq (\det(\widetilde{A}_1 \widetilde{A}_1^T))^{1/m} + (\det(\widetilde{A}_2 \widetilde{A}_2^T + \cdots + \widetilde{A}_N \widetilde{A}_N^T))^{1/m} \\ & \geq (\det(\widetilde{A}_1 \widetilde{A}_1^T))^{1/m} + (\det(\widetilde{A}_2 \widetilde{A}_2^T))^{1/m} + (\det(\cdots + \widetilde{A}_N \widetilde{A}_N^T))^{1/m} \\ & \geq \cdots \\ & \geq (\det(\widetilde{A}_1 \widetilde{A}_1^T))^{1/m} + (\det(\widetilde{A}_2 \widetilde{A}_2^T))^{1/m} + \cdots + (\det(\widetilde{A}_N \widetilde{A}_N^T))^{1/m}. \end{aligned}$$

That is, each summand $\widetilde{A_i A_i}^T$ adds an inequality.

Example 6.5.7: Consider the filtration $F : T \rightarrow \Delta K$ shown in Figure 6.1 and its birth-death module $\text{BD}_1 F$ as shown in Figure 6.2. In Example 6.4.3, we computed the Möbius homology Laplacian $L_0^\downarrow \text{BD}_1 F([2, 4]) = \begin{pmatrix} 3 & -1 \\ 0 & 0 \end{pmatrix}$. The characteristic polynomial, then, is given by

$$\det(\lambda I_2 - L_0^\downarrow \text{BD}_1 F([2, 4])) = \lambda^2 - 3\lambda.$$

First, observe that $\text{tr } L_0^\downarrow \text{BD}_1 F([2, 4]) = 3 = \text{bd}_1 F[1, 3] + \text{bd}_1 F[1, 4] + \text{bd}_1 F[2, 3]$. Furthermore, $\det L_0^\downarrow \text{BD}_p F([a, b]) = 0$. Indeed, there are zero birth-death spaces below $\text{BD}_1 F([2, 4])$ with the same dimension.

See Appendix A for more examples illustrating Theorems 6.5.5 and 6.5.6.

In particular, the two theorems presented in this chapter relate the combinatorial structure of the birth-death function to the spectrum of $L_0^\downarrow \text{BD}_p F([a, b])$ through linear algebraic invariants (the trace and the determinant).

6.6 Discussion and Future Directions

In this chapter, we first introduced *filtrations* of a simplicial complex K , denoted $F : P \rightarrow \Delta K$. In order to describe the evolving topology of K , we constructed the birth-death *module* $\text{BD}_p F$ directly from the data of the filtration F . By endowing the birth-death spaces $\text{BD}_p F([a, b])$ with an inner product structure inherited from the geometry of the underlying simplicial complex K , we defined the Möbius homology Laplacian $L_i^\downarrow \text{BD}_p F([a, b])$ acting on the chain space $C_i^\downarrow \text{BD}_p F([a, b])$.

We were particularly interested in the degree-zero Möbius homology Laplacian $L_0^\downarrow \text{BD}_p F([a, b])$. Under certain conditions, we found that $\ker L_0^\downarrow \text{BD}_p F([a, b])$ provided cycle representatives for the Möbius homology space $H_0^\downarrow \text{BD}_p F([a, b])$. These representatives maximize relative new content. Furthermore, we provided a first spectral analysis of $L_0^\downarrow \text{BD}_p F([a, b])$. We computed the trace of $L_0^\downarrow \text{BD}_p F([a, b])$ exactly, and we obtained a lower bound on the determinant of $L_0^\downarrow \text{BD}_p F([a, b])$.

Both quantities were given in terms of $\text{bd}_p F$. Finally, under the conditions that F is a one-parameter filtration and $\text{BD}_p F([a, b])$ has no cycles that branch, we computed the eigenvalues of $L_0^\downarrow \text{BD}_p F([a, b])$ exactly.

In all, we have demonstrated that the Möbius homology Laplacian for the birth-death module of a filtration is an interesting analytical tool with which to study the filtration $F : P \rightarrow \Delta K$ and its associated birth-death module $\text{BD}_p F$. However, we are not the first to investigate Laplace operators and harmonic homology spaces associated to a filtration: we provide a brief summary of the literature in this area in Section 6.6.1.

6.6.1 Relationship to Related Work

We are not the first to define or study combinatorial Laplace operators or harmonic cycle representatives for persistent homology. Early work on persistent harmonic forms appeared in a 2014 talk by Lieutier [16]. Since then, many researchers have studied analytic and algebraic properties of simplicial filtrations using tools such as persistent Laplacians, harmonic persistent homology, and orthogonal Möbius inversion. In this section, we briefly summarize how our approach relates to existing methods. We refer interested readers to the cited literature for further details.

Many existing tools in the area of persistent Laplacians and harmonic persistent homology consider one-parameter filtrations. Let T be a finite totally ordered set and let $F : T \rightarrow \Delta K$ be a filtration. In Section 6.1, we obtained the persistence diagram of F by first constructing the birth-death function $\text{bd}_p F : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}_{\geq 0}$ then applying Möbius inversion to obtain $\partial \text{bd}_p F : (\text{Dgm}(T), \preceq) \rightarrow \mathbb{Z}$. The following constructions provide cycle representatives associated to a persistence interval $[a, b) \in \text{Dgm}(T)$:

- In [19], Basu and Cox define a *space of harmonic cycle representatives* for a persistence interval $[a, b) \in \text{Dgm}(T)$ so that $\partial \text{bd}_p F([a, b)) \neq 0$. This space is denoted $\mathcal{P}_p^{a,b}$.
- In [20, 21], Gülen, Mémoli, and Wan define *orthogonal Möbius inversion*, a method for applying Möbius inversion to the Hilbert-space-valued birth-death module $\text{BD}_p F$. The orthogonal Möbius inverse of $\text{BD}_p F([a, b))$ is denoted $\text{OI}(\text{BD}_p F([a, b)))$.

In the proposition to follow, we clarify how these existing approaches relate to our construction.

Proposition 6.6.1 ([20, 21]): Let T be a totally ordered set, let $F : T \rightarrow \Delta K$ be a filtration, and let $\text{BD}_p F : (\text{Dgm}(T), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ be the associated birth-death module. Then,

$$\ker L_0^\downarrow \text{BD}_p F([a, b]) \cong \text{OI}(\text{BD}_p F([a, b])) \quad (6.3)$$

$$\cong \mathcal{P}_p^{a,b}. \quad (6.4)$$

Proof. Equation 6.3 follows from combining Theorem 5.1.2 with [20, Theorem 9], where the authors prove that $\text{OI}(\text{BD}_p F([a, b])) \cong H_0^\downarrow \text{BD}_p F([a, b])$. Equation 6.4 follows from [21, Theorem 9]. \square

That is, in the one-parameter setting, orthogonal Möbius inversion, harmonic persistent homology, and the Möbius homology Laplacian all provide isomorphic spaces of cycle representatives.

Moreover, Gülen, Mémoli, and Wan prove [20, Theorem 9] for the multiparameter setting. Specifically, if $F : P \rightarrow \Delta K$ is a filtration indexed by a finite poset P , then, combining [20, Theorem 9] with Theorem 5.1.2 yields

$$\ker L_0^\downarrow \text{BD}_p F([a, b]) \cong H_0^\downarrow \text{BD}_p F([a, b]) \cong \text{OI}(\text{BD}_p F([a, b])).$$

Finally, we note that our construction differs from existing formulations of persistent Laplacians (see, for example, [17, 18]), which typically act directly on simplicial chain complexes associated to each filtration level. In contrast, the degree-zero Möbius homology Laplacian acts on Hilbert-space-valued birth-death modules defined on intervals. Thus, the two provide complementary analytical methods with which to study filtered simplicial complexes.

6.6.2 Future Work

Thus, the Möbius homology Laplacian $L_0^\downarrow \text{BD}_p F([a, b])$ is compatible with existing constructions of cycle representatives. Moreover, the results of this chapter demonstrate that the non-zero

spectrum of $L_0^\downarrow \text{BD}_p F([a, b])$ reflects the combinatorial structure of the associated persistence diagram. However, there are many avenues for future work that remain to be investigated.

For one, further investigation into the characteristic polynomial of $L_0^\downarrow \text{BD}_p F([a, b])$ would be interesting. We have already seen that the trace of $L_0^\downarrow \text{BD}_p F([a, b])$ reflects the combinatorial structure of the persistence diagram, and we have a lower bound for the determinant in terms of the birth-death function. However, when $\dim \text{BD}_p F([a, b]) > 2$, there are more than two coefficients in the characteristic polynomial $\det(\lambda I_m - L_0^\downarrow \text{BD}_p F([a, b]))$. It would be interesting to investigate whether the coefficients of this characteristic polynomial admit combinatorial or topological interpretations, as in the case of the graph and combinatorial Laplacians. A better understanding of the characteristic polynomial of $L_0^\downarrow \text{BD}_p F([a, b])$ will aid in better understanding of its spectrum as a whole.

Furthermore, recall that we focused particularly on the *degree-zero* Möbius homology Laplacian $L_0^\downarrow \text{BD}_p F([a, b])$ in this chapter. Recall from Section 4.3 that the appearance of non-diagonal persistence intervals $[a, b]$ such that $\partial \text{bd}_p F([a, b]) < 0$, in part, motivated the introduction of Möbius homology. Indeed, a *negative* value for $\partial \text{bd}_p F([a, b])$ suggests the presence of odd degree Möbius homology spaces $H_{2i+1}^\downarrow \text{BD}_p F([a, b])$. Investigating the degree-one Möbius homology Laplacian, for example, may provide useful insight.

Moreover, Proposition 6.1.6 begs the question: given a persistence module $M : P \rightarrow \text{Hilb}_{\mathbb{R}}$ with free presentation $\varphi : G \Rightarrow M$, it is natural to ask which spectral properties of $L_0^\downarrow \text{BD}_p F([a, b])$ extend to the birth-death module BD_φ associated to φ . When $M = H_p F$ for a filtration F , it would be interesting to investigate how the Möbius homology Laplacian associated to BD_φ relates to that associated to $\text{BD}_p F$.

Bibliography

- [1] Claudia Landi and Patrizio Frosini. New pseudodistances for the size function space. In *Vision Geometry VI*, volume 3168, pages 52 – 60. International Society for Optics and Photonics, SPIE, 1997.
- [2] Vanessa Robins. Towards computing homology from approximations. *Topology Proceedings*, 24:503–532, 1999.
- [3] Herbert Edelsbrunner, David Letscher, and Afra Zomorodian. Topological persistence and simplification. *Discrete & Computational Geometry*, 28(4):511–533, 2002.
- [4] Gunnar Carlsson. Topology and data. *Bulletin of The American Mathematical Society*, 46:255–308, 2009.
- [5] David Cohen-Steiner, Herbert Edelsbrunner, and John Harer. Stability of persistence diagrams. *Discrete & Computational Geometry*, 37:103–120, 2007.
- [6] Frédéric Chazal, David Cohen-Steiner, Marc Glisse, Leonidas J. Guibas, and Steve Y. Oudot. Proximity of persistence modules and their diagrams. In *Proceedings of the Twenty-Fifth Annual Symposium on Computational Geometry*, SCG '09, page 237–246. Association for Computing Machinery, 2009.
- [7] Gunnar Carlsson and Afra Zomorodian. The theory of multidimensional persistence. *Discrete & Computational Geometry*, 42(1):71–93, 2009.
- [8] Alexander McCleary and Amit Patel. Edit distance and persistence diagrams over lattices. *SIAM Journal on Applied Algebra and Geometry*, 6(2):134–155, 2022.
- [9] Afra Zomorodian and Gunnar Carlsson. Computing persistent homology. *Discrete & Computational Geometry*, 33(2):249–274, 2005.

- [10] William V.D. Hodge. *The Theory and Applications of Harmonic Integrals*. Cambridge University Press, 1941.
- [11] Beno Eckmann. Harmonische funktionen und randwertaufgaben in einem komplex. *Commentarii Mathematici Helvetici*, 17(1):240–255, 1944.
- [12] Gustav R. Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148:497–508, 1847.
- [13] Danijela Horak and Jürgen Jost. Spectra of combinatorial Laplace operators on simplicial complexes. *Advances in Mathematics*, 244:303–336, 2013.
- [14] Gian-Carlo Rota. On the foundations of combinatorial theory I. Theory of Möbius functions. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 2(4):340–368, 1964.
- [15] Amit Patel and Primož Skraba. Möbius homology. *Transactions of the American Mathematical Society*, 2026.
- [16] André Lieutier. Talk: Persistent harmonic forms, 2014.
- [17] Rui Wang, Duc Duy Nguyen, and Guo-Wei Wei. Persistent spectral graph. *International Journal for Numerical Methods in Biomedical Engineering*, 36(9), 2020.
- [18] Facundo Mémoli, Zhengchao Wan, and Yusu Wang. Persistent Laplacians: Properties, algorithms and implications. *SIAM Journal on Mathematics of Data Science*, 4(2):858–884, 2022.
- [19] Saugata Basu and Nathanael Cox. Harmonic persistent homology. *SIAM Journal on Applied Algebra and Geometry*, 8(1):189–224, 2024.
- [20] Aziz Burak Gülen, Facundo Mémoli, and Zhengchao Wan. Grassmannian persistence diagrams. arXiv:2311.06870v5, 2025.

- [21] Aziz Burak Gülen, Facundo Mémoli, and Zhengchao Wan. Grassmannian persistence diagrams: Special properties in the 1-parameter setting. *arXiv:2504.06077v2*, 2025.
- [22] Emily Riehl. *Category Theory in Context*. Aurora: Dover Modern Math Originals. Dover Publications, 2017.
- [23] Gilbert Helmberg. *Introduction to Spectral Theory in Hilbert Space*. North-Holland Series in Applied Mathematics and Mechanics. North-Holland Publishing Company, 1969.
- [24] Sheldon Axler. *Linear Algebra Done Right*. Undergraduate Texts in Mathematics. Springer, fourth edition, 2024.
- [25] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
- [26] William N. Anderson Jr. and Thomas D. Morley. Eigenvalues of the Laplacian of a graph. *Linear and Multilinear Algebra*, 18(2):141–145, 1985.
- [27] Miroslav Fiedler. Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal*, 23(2):298–305, 1973.
- [28] Michael E. Fisher. On hearing the shape of a drum. *Journal of Combinatorial Theory*, 1(1):105–125, 1966.
- [29] Norman Biggs. *Algebraic Graph Theory*. Cambridge University Press, second edition, 1993.
- [30] Art M. Duval and Victor Reiner. Shifted simplicial complexes are Laplacian integral. *Transactions of the American Mathematical Society*, 354(11):4351–4377, 2002.
- [31] John Steenbergen. *Towards a Spectral Theory for Simplicial Complexes*. PhD thesis, Duke University, 2013.
- [32] Olivier Bernardi and Caroline J. Klivans. Directed rooted forests in higher dimension. *Electronic Journal of Combinatorics*, 23(4), 2016.

- [33] Ivan Contreras and Andrew Tawfeek. On discrete gradient vector fields and Laplacians of simplicial complexes. *Annals of Combinatorics*, 28(1):67–91, 2024.
- [34] Jakob Hansen and Robert Ghrist. Toward a spectral theory of cellular sheaves. *Journal of Applied and Computational Topology*, 3, 2019.
- [35] Justin Michael Curry. *Sheaves, Cosheaves and Applications*. PhD thesis, University of Pennsylvania, 2014.
- [36] Jakob Hansen. *Laplacians of Cellular Sheaves: Theory and Applications*. PhD thesis, University of Pennsylvania, 2020.
- [37] Amit Patel. Generalized persistence diagrams. *Journal of Applied and Computational Topology*, 1(3):397–419, 2018.
- [38] Peter Hall. The Eulerian functions of a group. *Quarterly Journal of Mathematics*, 1:134–151, 1936.
- [39] Louis Weisner. Abstract theory of inversion of finite series. *Transactions of the American Mathematical Society*, 38(3):474–484, 1935.
- [40] Aziz Burak Gülen and Alexander McCleary. Galois connections in persistent homology. arXiv:2201.06650v5, 2022.
- [41] Amit Patel and Tatum Rask. Poincaré duality for generalized persistence diagrams of (co)filtrations. *Journal of Applied and Computational Topology*, 8(1):427–442, 2024.
- [42] Ulrich Bauer and Luis Scoccola. Multi-parameter persistence modules are generically indecomposable. *International Mathematics Research Notices*, 2025(5), 2025.
- [43] Woojin Kim and Facundo Mémoli. Generalized persistence diagrams for persistence modules over posets. *Journal of Applied and Computational Topology*, 5(4):533–581, 2021.

- [44] Nathaniel Clause, Woojin Kim, and Facundo Mémoli. The generalized rank invariant: Möbius invertibility, discriminating power, and connection to other invariants. arXiv:2207.11591v5, 2024.
- [45] Gunnar Carlsson and Vin de Silva. Zigzag persistence. *Foundations of Computational Mathematics*, 10(4):367–405, 2010.
- [46] Robert Ghrist and Gregory Henselman-Petrusek. Saecular persistence. arXiv:2112.04927v2, 2021.
- [47] Marvin Marcus and Henryk Minc. *A Survey of Matrix Theory and Matrix Inequalities*. Dover Books on Mathematics. Dover Publications, 1992.

Appendix A

Examples of Filtrations

In this appendix, we provide multiple examples of filtrations $F : P \rightarrow \Delta K$ indexed over a finite poset P . Denote by $(\text{Dgm}(P), \preceq)$ the poset of all persistence intervals $[a, b)$ with $a \in P$ and $b \in \overline{P} = P \cup \{\infty\}$ so that $a \leq b$. In all examples, we will compute the birth-death module $\text{BD}_p F : (\text{Dgm}(P), \preceq) \rightarrow \text{Hilb}_{\mathbb{R}}$ for $p = 1$.

These examples are meant to illustrate that the conditions of Theorem 6.5.2 are sufficient but not necessary. The conditions for this theorem are that (1) P must be totally ordered, and (2) cycles in $\text{BD}_p F([a, b))$ do not branch. Furthermore, we will explore the statements of Theorems 6.4.2, 6.5.5, and 6.5.6 in these specific examples.

A.1 One-Parameter Examples

We start by considering one-parameter filtrations. That is, $F : T \rightarrow \Delta K$ for T a finite totally ordered set. First, we provide an example where $L_0^\downarrow \text{BD}_1 F([a, b))$ has integer eigenvalues even though the branching condition is not satisfied.

Example A.1.1: Consider the filtration $F : T \rightarrow \Delta K$ in Figure A.1. Observe that $\text{BD}_1 F([2, 4))$ contains a cycle that branches: namely, $z = ([v_1 v_2] + [v_2 v_3] - [v_1 v_3]) + ([v_3 v_4] + [v_4 v_5] - [v_3 v_5])$. The cycle $[v_1 v_2] + [v_2 v_3] - [v_1 v_3]$ first appears in $\text{BD}_1 F([2, 3))$ whereas $[v_3 v_4] + [v_4 v_5] - [v_3 v_5]$ first appears in $\text{BD}_1 F([1, 4))$, but $[2, 3)$ and $[1, 4)$ are not comparable in $\text{Dgm}(T)$.

However, the Möbius homology Laplacian for $\text{BD}_1 F([2, 4))$ is the identity matrix:

$$L_0^\downarrow \text{BD}_1 F([2, 4)) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$





Figure A.1: Above: a filtration $F : T \rightarrow \Delta K$ where $T = 1 < 2 < 3 < 4$. Below: the birth-death module $\text{BD}_1 F$ (left), and the persistence diagram $\partial \text{bd}_1 F$ (right). In $\text{BD}_1 F$, all unlabeled maps are either the zero map or the identity

Indeed, the eigenvalue of 1 with multiplicity two is given by $\text{len}([2, 3], [2, 4]) = \#\{[2, 3]\}$ and $\text{len}([1, 4], [2, 4]) = \#\{[1, 4]\}$. Furthermore,

$$\det \left(\lambda I_2 - L_0^\downarrow \text{BD}_1 F([2, 4]) \right) = \lambda^2 - 2\lambda + 1.$$

In this case,

$$\text{tr } L_0^\downarrow \text{BD}_1 F([2, 4]) = 2 = \text{bd}_1 F([1, 4]) + \text{bd}_1 F([2, 3])$$

and

$$\det \left(L_0^\downarrow \text{BD}_1 F([2, 4]) \right)^{1/2} = 1^{1/2} \geq 1.$$




Figure A.2: Above: a filtration $F : T \rightarrow \Delta K$ where $T = 1 < 2 < 3 < 4$. Below: the birth-death module $\text{BD}_1 F$ (left), and the persistence diagram $\partial \text{bd}_1 F$ (right). In $\text{BD}_1 F$, all unlabeled maps are either the zero map or the identity.

Next, we provide an example of a filtration F and a birth-death space $\text{BD}_1 F([a, b))$ containing cycles that branch. The conditions of Theorem 6.5.2 do not hold, and we have non-integer eigenvalues.

Example A.1.2: Consider the filtration $F : T \rightarrow \Delta K$ in Figure A.2. The Möbius homology Laplacian for $\text{BD}_1 F([2, 4))$ is

$$L_0^\downarrow \text{BD}_1 F([2, 4)) = \begin{pmatrix} 1/2 & 1/2 \\ 1/6 & 3/2 \end{pmatrix}.$$

The eigenvalues of $L_0^\downarrow \text{BD}_1 F([2, 4])$ are $1 \pm \frac{1}{\sqrt{3}}$. The two non-zero eigenvalues represent the points in the persistence diagram at $[1, 4)$ and $[2, 3)$, but they do not give us some notion of "length." Furthermore,

$$\det \left(\lambda I_2 - L_0^\downarrow \text{BD}_1 F([2, 4]) \right) = \lambda^2 - 2\lambda + \frac{2}{3}.$$

In this case,

$$\text{tr } L_0^\downarrow \text{BD}_1 F([2, 4]) = 2 = \text{bd}_1 F([1, 4]) + \text{bd}_1 F([2, 3])$$

and

$$\det \left(L_0^\downarrow \text{BD}_1 F([2, 4]) \right)^{1/2} = (2/3)^{1/2} \geq 0.$$

A.2 Multiparameter Examples

Now, we provide examples of multiparameter filtrations $F : P \rightarrow \Delta K$: that is, P is a finite poset. First, let's see an example where $L_0^\downarrow \text{BD}_1 F([a, b])$ has non-integer eigenvalues.

Example A.2.1: Consider the filtration $F : P \rightarrow \Delta K$ in Figure A.3. The Möbius homology Laplacian of $\text{BD}_1 F[d, d]$ is the 2×2 matrix given by:

$$L_0^\downarrow \text{BD}_1 F[d, d] = \begin{pmatrix} 5 & -2/3 \\ -2/3 & 5 \end{pmatrix}.$$

The eigenvalues of $L_0^\downarrow \text{BD}_1 F[d, d]$ are $\frac{13}{3}$ and $\frac{17}{3}$. The characteristic polynomial is

$$\det \left(\lambda I_2 - L_0^\downarrow \text{BD}_1 F([d, d]) \right) = \lambda^2 - 10\lambda + 24\frac{5}{9}.$$

In this case,

$$\text{tr } L_0^\downarrow \text{BD}_1 F([d, d]) = 10 = \sum_{[x, y] \prec [d, d]} \text{bd}_1 F([x, y])$$

and

$$\det \left(L_0^\downarrow \text{BD}_1 F([d, d]) \right)^{1/2} = (221/9)^{1/2} \geq 3.$$




Figure A.3: Above: a filtration $F : P \rightarrow \Delta K$ where P is not totally ordered. Below: $(\text{Dgm}(P), \preceq)$ (left), the birth-death module $\text{BD}_1 F$ (center), and the persistence diagram $\partial \text{bd}_1 F$ (right). In $\text{BD}_1 F$, all unlabeled maps are either the zero map or the identity.

Now, let's see an example of a multiparameter filtration F where $L_0^\downarrow \text{BD}_1 F([a, b])$ has integer eigenvalues.

Example A.2.2: Consider the multiparameter filtration from Figure A.4. Let $\text{BD}_1 F([c, d])$ be generated by the cycles $z_1 = [v_1 v_2] + [v_2 v_4] - [v_1 v_4]$ and $z_2 = [v_2 v_3] + [v_3 v_4] - [v_2 v_4]$, and let $\text{BD}_1 F([a, d])$ be generated by $z_3 = z_1 + z_2$. The Möbius homology Laplacian for $\text{BD}_1 F([c, d])$ is

given by

$$L_0^\downarrow \text{BD}_1 F([c, d]) = \begin{pmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{pmatrix}.$$

It is easy to check that $\ker L_0^\downarrow \text{BD}_1 F([c, d])$ is generated by $\begin{pmatrix} 1 \\ -1 \end{pmatrix}$, which represents the cycle

$$\begin{aligned} z &= ([v_1 v_2] + [v_2 v_4] - [v_1 v_4]) - ([v_2 v_3] + [v_3 v_4] - [v_2 v_4]) \\ &= [v_1 v_2] + 2[v_2 v_4] - [v_1 v_4] - [v_2 v_3] - [v_3 v_4]. \end{aligned}$$

Observe that z_1 is another representative for $\text{BD}_1 F([c, d])$. We can write z_1 as the sum $\frac{1}{2}z + \frac{1}{2}(z_1 + z_2)$ where $z_1 + z_2$ is the image of z_3 under ∂_1^\downarrow . Up to rescaling, the harmonic representative z has minimal energy: $\|z_1\|^2 = 3$ whereas $\|\frac{1}{2}z\|^2 = \frac{1}{4}8 = 2$.

The eigenvalues of $L_0^\downarrow \text{BD}_1 F([c, d])$ are 0 and 1. Indeed, $1 = \text{len}([a, d], [c, d]) = \#\{[a, d]\}$.

The characteristic polynomial is

$$\det(\lambda I_2 - L_0^\downarrow \text{BD}_1 F([c, d])) = \lambda^2 - \lambda.$$

In this case,

$$\text{tr } L_0^\downarrow \text{BD}_1 F([c, d]) = 1 = \text{bd}_1 F([a, d])$$

and

$$\det(L_0^\downarrow \text{BD}_1 F([c, d]))^{1/2} = 0^{1/2} \geq 0.$$




Figure A.4: Above: A filtration $F : P \rightarrow \Delta K$ where P is not totally ordered. Below: $(\text{Dgm}(P), \preceq)$ of P (left), the birth-death module $BD_1 F$ (center), and the persistence diagram $\partial bd_1 F$ (right). In $BD_1 F$, all unlabeled maps are either the zero map or the identity.