COMPUTATIONAL CONNECTION MATRIX THEORY

By

KELLY SPENDLOVE

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in Mathematics

Written under the direction of

Konstantin Mischaikow

And approved by

____

____

____

____

New Brunswick, New Jersey

October, 2019

ABSTRACT OF THE DISSERTATION

Computational Connection Matrix Theory

By KELLY SPENDLOVE

Dissertation Director: Konstantin Mischaikow

We develop a computational and categorical framework for connection matrix theory. In terms of computations, we give an algorithm for computing the connection matrix based on algebraic-discrete Morse theory. The makes the connection matrix available, for the first time, as a computational tool within applied topology and dynamics. In addition, the algorithm provides a straightforward constructive proof of the existence of connection matrices. In terms of categories, our formulation resolves the non-uniqueness of the connection matrix, as well as relates the connection matrix to persistent homology.

We extend existing computational Conley theory to incorporate connection matrix theory. This is done by developing a setting, which we call *transversality models*, in which discrete approximations to continuous flows can be used to compute connection matrices for the underlying continuous system. We make applications to a Morse theory on spaces of braid diagrams.

Finally, we provide an implicit discrete Morse pairing for cubical complexes. This enables the computation of connection matrices in high-dimensional cubical complexes. We benchmark our algorithm on a set of such examples.

Acknowledgements

I owe a great deal to Konstantin Mischaikow, who has advised me throughout my time at Rutgers. His enthusiasm and optimism are inspirational. I am also indebted to Shaun, with whom I have had many discussions; Shaun has been a second adviser to me during this time. Rob – both our collaboration and friendship have been very valuable, especially so now that I will set off on my own. Tomas has always been supportive and helped me to start – now so long ago – down this particular path.

Thank you to my other committee members: Chris and Chuck – especially Chuck for helping to shape the thesis in these final days.

The hospitality and guidance of the people who have hosted me in their countries has been wonderful: thank you Rob, Hiroshi, Zin, and Fred. Thanks is also due to other senior people who have offered me counsel and conversation: Tom, Marian, Sarah, Yasu, Miro, Matt Thibault, Matt Kahle, Uli, Vidit, Sanjeevi, Sayan, Justin, and Bob.

My friends throughout have been very helpful. In particular, Justin for the d-dog dinners, Vince for each two+ course meal, Cole for all the werblin mornings and racquetball, and Alex for the summer conferences. Their friendships have been indispensable.

I would like to express my gratitude toward the NSF. I have benefited tremendously through the Graduate Research Fellowship, and the EAPSI and GROW programs. I owe a thank you to all my hosting locations: Kyoto, INRIA, MBI and VU Amsterdam. I also need to thank Gary Buschhorn for being exceptionally competent and having the ability to resolve all of my administrative bunglings.

Finally, thanks to the following coffee shops, where much of this thesis has been written and these ideas developed: in A'dam – Monks Coffee and Scandinavian Embassy; in New Jersey – OQ; in Columbus – Mission and One Line; in Zoo – Clyde, Dolce, Break, and Zootown.

Dedication

For Sandee, Cecily, and Brett.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

High-throughput technologies and high performance computing infrastructures are enabling the efficient and inexpensive collection of massive amounts of experimental and simulated data. In the 21st century, science is becoming inherently more data-driven, with data being harnessed to inform experiment and advance theory. Such data are often high-dimensional and generated by a complex, nonlinear system. A pressing challenge is to reconcile this fact with classical dynamical systems theory, which demonstrated that nonlinear systems can exhibit intricate behavior at all scales with respect to both system variables and parameters. Experimental data are often imprecise or noisy; simulated data arise from models where parameters are rarely known exactly and nonlinearities may not be derived from first principles. This suggests that any analysis designed to synthesize theory, experimentation, computation and data must be based on robust, multi-scale and multi-parameter features. That is, one is interested in the multi-scale topology and geometry of the data, as well as a robust description of the behavior of the unknown nonlinear system; both the *topology* and *dynamics* of the data.

Topological data analysis is a flourishing field dedicated to using algebraic topological tools to understand complex and nonlinear data. The most prominent tool in the field is *persistent homology*. Roughly speaking, persistent homology is a method for computing topological features at different resolutions of a space or dataset $\mathcal{X}$. Persistent homology rests on the idea of topologizing $\mathcal{X}$ (often with some cell complex structure), and then establishing a filtration $\varnothing = \mathcal{X}^0 \subset \mathcal{X}^1 \subset \ldots \subset \mathcal{X}^n = \mathcal{X}$. For instance, the sublevel sets $\{f^{-1}(-\infty, a]\}$ of a scalar function $f\colon \mathcal{X} \to \mathbb{R}$ form a filtration of $\mathcal{X}$. The inclusions $\mathcal{X}^i \hookrightarrow \mathcal{X}^j$ induce homomorphisms in homology and persistence tabulates the change in homology between $\mathcal{X}^i$ and $\mathcal{X}^j$ [11, 39].

A useful method for extracting coarse yet robust descriptions of dynamics is *computational Conley theory.* Classical Conley theory is a topological generalization of Morse theory to continuous semiflows and self-maps on a compact metric space X. The global dynamics are organized via a *Morse decomposition* M: a finite collection of mutually disjoint, compact, isolated invariant sets $M(p) \subset X$, called *Morse sets*, which are indexed by a partial order $(\mathsf{P}, \leq)$, such that if $x \in X \setminus \bigcup_{p\in\mathsf{P}} M(p)$, then in forward time (with respect to the semiflow or self-map) x limits to a Morse set $M(p)$ and in backward time an orbit through x limits to a Morse set $M(q)$.

In computational Conley theory, one first discretizes the phase space at a resolution fixed *a priori*. A multi-valued map is constructed which appropriately models the underlying continuous dynamics. One can partition the multi-valued map into recurrent and non-recurrent components efficiently using graph-theoretic algorithms [2]. This furnishes a combinatorial representation of the Morse decomposition. Moreover, one can infer the behavior of the recurrent dynamics produced from the decomposition using algebraic topological methods, viz., Conley index theory. There are strong links between the combinatorial dynamics and continuous dynamics that have been formalized in [25, 26, 27, 28].

The techniques of this dissertation lie at the intersection of applied topology, computational dynamics, and topological data analysis. For instance, a simple generalization of a filtration is to assume that a decomposition of the dataset $\mathcal{X}$ is given in the form of a distributive lattice. To be more precise, assume that L is a finite distributive lattice and $\{\mathcal{X}^a \subset \mathcal{X} \mid a \in \mathsf{L}\}$ is an isomorphic lattice (the indexing providing the isomorphism) with operations $\vee := \cup$ and $\wedge := \cap$. One aim of this thesis is to provide an efficient algorithm for computing a boundary operator, called the *connection matrix*,

$$\Delta\colon \bigoplus_{a\in\mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}) \to \bigoplus_{a\in\mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}) \tag{1.1}$$

that is strictly upper triangular with respect to $\leq$ where $\mathsf{J}(\mathsf{L})$ denotes the set of join-irreducible elements of L and $\overleftarrow{a}$ denotes the unique predecessor of a, again with respect to $\leq$. In the setting of topological data analysis, the connection matrix can be used to recover the persistent homology of the lattice of cell complexes $\{\mathcal{X}^a\}$ (see Theorem 3.9.6).

In the context of analyzing dynamical systems, the aim is to guarantee that the collection $\{\mathcal{X}^a\}$ is a collection of attracting blocks for an underlying continuous flow (see Section 4.3). In this case, the collection $\{\mathsf{M}(a)\}_{a\in\mathsf{J}(\mathsf{L})}$ where $\mathsf{M}(a) = \mathrm{Inv}(\mathcal{X}^a \setminus \mathcal{X}^{\overleftarrow{a}})$, i.e., the maximal invariant set contained in $\mathcal{X}^a \setminus \mathcal{X}^{\overleftarrow{a}}$, is a Morse decomposition. The (homological) Conley index for a Morse set is

$$CH_\bullet\mathsf{M}(p) = H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}})$$

Thus, the connection matrix of (1.1) can be rewritten as

$$\Delta\colon \bigoplus_{a\in\mathsf{J}(\mathsf{L})} CH_\bullet\mathsf{M}(a) \to \bigoplus_{a\in\mathsf{J}(\mathsf{L})} CH_\bullet\mathsf{M}(a). \tag{1.2}$$

The existence of a Δ expressed in the form of (1.2) is originally due to R. Franzosa [17]. Although Franzosa's proof of the existence of connection matrices is constructive it is not straightforwardly amenable to computation. The name connection matrix arose since Δ can be used to identify and give lower bounds on the structure of connecting orbits between Morse sets [34, 35, 37]. Ultimately, the connection matrix is the mathematical object that promotes the Conley theory to a homological theory for dynamical systems. Foremost, we consider the contributions within this dissertation as promoting the computational Conley theory to a computational homological theory for flows.

1.1 Outline and Contributions

In Chapter 3 we give a new formulation for connection matrix theory. In particular, we emphasize categorical language and homotopy categories. One contribution in this section is our use of chain equivalence and homotopy categories to formulate the definition of the connection matrix. The payoffs of this formulation are quite satisfying:

1. We can give an algorithm based on the technique of reductions, which employs algebraic-discrete Morse theory. The proof of correctness of our algorithm yields a algorithmic proof of the existence of connection matrices.

2. We address the non-uniqueness question of the connection matrix. Classically, this is an open problem. We show that in our formulation, a connection matrix is unique up to a similarity transformation.

3. We distill the construction of the connection matrix to a particular functor, which we call a *Conley functor*.

4. We readily relate persistent homology and connection matrix theory. This implies applications for connection matrix theory in topological data analysis.

In Chapter 4 we develop a setting for applying connection matrix theory to continuous dynamics. We show how one moves from Conley-theoretic approximations of continuous dynamics to the appropriate graded or filtered cell complexes on which to apply the computational connection matrix theory. As an application we use classical examples from connection matrix theory, as well as examples from a Morse theory on braids.

In Chapter 5 we present an an implicit scheme for discrete Morse theory on cubical complexes. This allows us to modify the connection matrix algorithm when computing connection matrices in the setting of cubical complexes. This enables connection matrix computations in high-dimensions (i.e., $d = 9, 10$).

Much of the work in this thesis is a result of collaborations with S. Harker, K. Mischaikow and R. Van der Vorst. In particular, Chapters 3 and 5 are joint work with S. Harker and K. Mischaikow. Chapter 4 is joint work with S. Harker, K. Mischaikow and R. Van der Vorst.

1.2 Examples

We front-load this thesis with a set of simple examples to get the point across. This allows us to refer back to this original selection of examples further along in the thesis when trying to illuminate various concepts. Unfortunately, many settings in which one is traditionally interested in applying connection matrix theory (viz., dynamics and the search for connecting orbits) require setting up quite a bit of mathematical machinery.

Such examples are given later in Section 4.7. For now, we content ourselves with the following selection of examples. The first three examples are selected from applied topology; the last example comes from dynamics.

- Example 1.2.1 gives an example of the connection matrix.
- Example 1.2.2 provides a computational perspective on connection matrix theory. Succinctly, the connection matrix is a form of 'homologically-lossless data compression'.
- Example 1.2.3 examines the relationship between the connection matrix and persistent homology.
- Example 1.2.4 bridges combinatorial and continuous dynamics, depicting that the output of our algorithm is a connection matrix for an underlying continuous flow.

Example 1.2.1 (Connection Matrix)**.** Let $\mathsf{P} = \{p, q, r\}$ with order $\leq$ generated by $p \leq q$ and $r \leq q$. Consider the cell complexes $\mathcal{X}, \mathcal{X}'$ and the maps ν, ν' below. These cell complexes will provide a working example throughout the thesis.

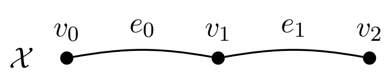

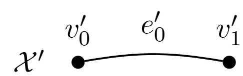

$$\nu(x) = \begin{cases} p & x = v_0 \\ r & x = v_2 \\ q & x \in \{e_0, v_1, e_1\} \end{cases} \qquad \nu'(x) = \begin{cases} p & x = v_0' \\ r & x = v_1' \\ q & x = e_0' \end{cases}$$

The maps ν and ν' are poset morphisms from the face poset of $\mathcal{X}$ and $\mathcal{X}'$ to P. The pairs $(\mathcal{X}, \nu)$ and $(\mathcal{X}', \nu')$ are called P-*graded cell complexes*, see Section 3.3.4. The associated chain complexes $C_\bullet(\mathcal{X})$ and $C_\bullet(\mathcal{X}')$ are (graded) chain equivalent (see Example 3.3.24). A P-graded cell complex may be visualized as in Figure 1.1: the Hasse diagram of P is given and each $s \in \mathsf{P}$ is annotated with itself and its fiber $\mathcal{X}^s := \nu^{-1}(s)$. In our visualization if t covers s in P then there is a directed edge $t \to s$. This orientation coincides with the action of the boundary operator and agrees with the Conley-theoretic literature. However, outside of Conley theory the Hasse diagram is often give the opposite orientation. These visualizations will be useful for the computations done in Chapter 5.

$$\begin{tikzcd}[column sep=small] & q : e_0, v_1, e_1 \arrow[dl] \arrow[dr] & \\ p : v_0 & & r : v_2 \end{tikzcd}$$

Figure 1.1: Visualization of $(\mathcal{X}, \nu)$.

$$\begin{tikzcd}[column sep=small] & q : e_0' \arrow[dl] \arrow[dr] & \\ p : v_0' & & r : v_1' \end{tikzcd}$$

Figure 1.2: Visualization of $(\mathcal{X}', \nu')$.

Let L be the lattice of down-sets of P, i.e., the lattice with operations $\cap$ and $\cup$, consisting of unions of the sets

$$\alpha := \downarrow p = \{s \in \mathsf{P} \mid s \leq p\} \qquad \beta := \downarrow q = \{s \in \mathsf{P} \mid s \leq q\} \qquad \gamma := \downarrow r = \{s \in \mathsf{P} \mid s \leq r\}$$

For $a \in \mathsf{L}$ define $\mathcal{X}^a := \nu^{-1}(a)$. In this way, ν engenders the lattice of subcomplexes $\{\mathcal{X}^a \subset \mathcal{X} \mid a \in \mathsf{L}\}$ and ν' engenders the lattice of subcomplexes $\{\mathcal{X}'^a \subset \mathcal{X}' \mid a \in \mathsf{L}\}$. These are given below, where the down-sets are annotated.

$$\begin{tikzcd}[column sep=small] & \beta : v_0, e_0, v_1, e_1, v_2 \arrow[d, dash] & \\ & v_0, v_2 \arrow[dl, dash] \arrow[dr, dash] & \\ \alpha : v_0 \arrow[dr, dash] & & \gamma : v_2 \arrow[dl, dash] \\ & \varnothing & \end{tikzcd}$$

Figure 1.3: $\{\mathcal{X}^a \mid a \in \mathsf{L}\}$.

$$\begin{tikzcd}[column sep=small] & \beta : v_0', e_0', v_1' \arrow[d, dash] & \\ & v_0', v_1' \arrow[dl, dash] \arrow[dr, dash] & \\ \alpha : v_0' \arrow[dr, dash] & & \gamma : v_1' \arrow[dl, dash] \\ & \varnothing & \end{tikzcd}$$

Figure 1.4: $\{\mathcal{X}'^a \mid a \in \mathsf{L}\}$.

The boundary operator Δ' for $C_\bullet(\mathcal{X}')$ is a connection matrix for $(\mathcal{X}, \nu)$. To unpack this a bit more, for $\Delta_1' \colon C_1(\mathcal{X}') \to C_0(\mathcal{X}')$ we have

$$\Delta_1' = \begin{array}{c c} & e_0' \\ \begin{array}{c} v_0' \\ v_1' \end{array} & \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{array}$$

A quick computation shows that

$$H_1(\mathcal{X}^\beta, \mathcal{X}^{\overleftarrow{\beta}}) \cong H_1(\mathcal{X}^\gamma) \cong \mathbb{Z}_2 = C_1(\mathcal{X}')$$

and, using $\mathcal{X}^{\overleftarrow{\alpha}} = \mathcal{X}^{\overleftarrow{\gamma}} = \varnothing$,

$$H_0(\mathcal{X}^\alpha, \mathcal{X}^{\overleftarrow{\alpha}}) \oplus H_0(\mathcal{X}^\gamma, \mathcal{X}^{\overleftarrow{\gamma}}) \cong H_0(\mathcal{X}^p) \oplus H_0(\mathcal{X}^r) \cong \mathbb{Z}_2 \oplus \mathbb{Z}_2 = C_0(\mathcal{X}')$$

Thus we may write Δ_1' as a map

$$\Delta_1' \colon H_1(\mathcal{X}^\beta, \mathcal{X}^{\overleftarrow{\beta}}) \to H_0(\mathcal{X}^\alpha, \mathcal{X}^{\overleftarrow{\alpha}}) \oplus H_0(\mathcal{X}^\gamma, \mathcal{X}^{\overleftarrow{\gamma}})$$

In context of computational dynamics – in particular Conley theory – the relative homology groups $\{H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}})\}_{a\in \mathsf{J}(\mathsf{L})}$ are the Conley indices indexed by the join-irreducibles of L. The Conley index is an algebraic invariant of an isolated invariant set, which generalizes the notion of the Morse index. The classical form of the connection matrix is a boundary operator on Conley indices

$$\Delta' \colon \bigoplus_{a\in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}) \to \bigoplus_{a\in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}})$$

This form makes it more apparent that non-zero entries in the connection matrix may force the existence of connecting orbits. See [17].

Example 1.2.2 (Data Compression)**.** In applications, the data are orders of magnitude larger than Example 1.2.1. Let $N = 9 \times 10^9$ and $\mathcal{K}$ be the cubical complex on $[0,1] \subset \mathbb{R}$ where the vertices $\mathcal{K}_0$ and edges $\mathcal{K}_1$ are given by

$$\mathcal{K}_0 = \left\{ \left[\frac{k}{N}, \frac{k}{N}\right] \right\}_{0\le k\le N} \qquad \mathcal{K}_1 = \left\{ \left[\frac{k}{N}, \frac{k+1}{N}\right] \right\}_{0\le k< N}$$

Let P be as in Example 1.2.1. Let $\mu\colon \mathcal{K} \to \mathsf{P}$ be the poset morphism

$$\mu(x) = \mu([l,r]) = \begin{cases} p & \text{if } r \le 3\times 10^9 \\ r & \text{if } l \ge 6\times 10^9 \\ q & \text{otherwise} \end{cases}$$

$\mathcal{K}$ contains a large number of cells; see Section 5.4 for examples arising from computational dynamics in which similar orders of magnitude are encountered. Due to the cardinality of $\mathcal{K}$, in Figure 1.5 we give a different visualization for the P-graded cell complex $(\mathcal{K}, \mu)$. Here $M := N/3$ and each $p \in \mathsf{P}$ is annotated with itself p and the

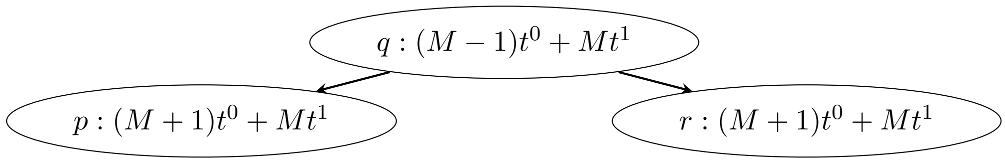

Figure 1.5: Visualization of $(\mathcal{K}, \mu)$.

polynomial $F_{\mathcal{K}^p}(t) := \sum_i \alpha_i t^i$ where α_i is the number of cells in the fiber $\mathcal{K}^p = \mu^{-1}(p)$ of dimension i, viz. the f-polynomial of $\mathcal{K}^p$.

The graded cell complex $(\mathcal{X}', \nu')$ of Example 1.2.1 may also be visualized in this fashion. A similar argument as the one in Example 1.2.1 shows that Δ' is a connection

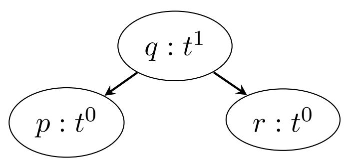

Figure 1.6: Visualization of $(\mathcal{X}', \nu')$.

matrix for $(\mathcal{K}, \mu)$. Another way to see this is as follows. Recall that the Poincare polynomial of $\mathcal{X}$ is defined as the polynomial $P_{\mathcal{X}}(t) := \sum_i \dim H_i(\mathcal{X}) t^i$. A quick computation shows that

$$P_{\mathcal{K}^q}(t) = t^1 = F_{\mathcal{X}'^q}(t)$$
$$P_{\mathcal{K}^p(t)} = P_{\mathcal{K}^r(t)} = t^0 = F_{\mathcal{X}'^p}(t) = F_{\mathcal{X}'^r}(t)$$

Therefore for each $p \in \mathsf{P}$ the f-polynomial $F_{\mathcal{X}'^p}(t)$ is precisely the Poincare polynomial $P_{\mathcal{K}^p}(t)$. This implies that the boundary operator Δ' is can be interpreted as a map on the relative homology groups of $\{\mathcal{X}^a\}$. We call the visualization in Figure 1.6 the *Conley-Morse graph*. This is a visualization of a graded complex, however the data of the connection matrix itself is not shown.

This example highlights two aspect of the connection matrix.

1. First, from a computational perspective the cell complex $(\mathcal{X}', \nu')$ and connection matrix Δ' can be thought of as a compression of $(\mathcal{K}, \mu)$. Moreover, as part of our definition of connection matrix (see Definition 3.3.25) there is a (graded)

chain equivalence ϕ between the chain complexes $C(\mathcal{K})$ and $C(\mathcal{X}')$. ϕ induces an isomorphism on any homological invariant (e.g. homology, persistent homology, graded module braids). Thus from the computational perspective, the connection matrix is a form of compression which does not lose information with respect to homological invariants. See Example 1.2.3, Section 3.9 and Theorems 3.9.3.

2. Second – and again from the computational perspective – notice that $(\mathcal{X}', \nu')$ cannot be compressed further as $P_{\mathcal{X}'p}(t) = F_{\mathcal{X}'^p}(t)$ for each p, i.e., the f-polynomial of $\mathcal{X}'^p$ is precisely the Poincare polynomial of $\mathcal{X}'^p$. In this sense, the connection matrix is maximally compressed and it is the smallest object (of the graded chain equivalence class) capable of recovering the homological invariants.

Example 1.2.3 (Persistent Homology)**.** In this example we address Theorem 3.9.6 and the interplay of persistent homology and connection matrix theory.[1] Let $(\mathcal{X}, \nu)$ and $(\mathcal{X}', \nu')$ be as in Example 1.2.1. Let Q be the poset $\mathsf{Q} = \{0, 1, 2\}$ with order $0 \leq 1 \leq 2$. Consider the poset morphism $\rho\colon \mathsf{P} \to \mathsf{Q}$ given by

$$\mu(x) = \begin{cases} 0 & x = p \\ 1 & x = r \\ 2 & x = q \end{cases}$$

Let $\mu := \rho \circ \nu\colon \mathcal{X} \to \mathsf{Q}$ and $\mu' := \rho \circ \nu'\colon \mathcal{X}' \to \mathsf{Q}$. Then $(\mathcal{X}, \mu)$ and $(\mathcal{X}', \mu')$ are Q-graded cell complexes, which may be visualized as in Figure 1.7.

Let K be the lattice of down-sets of Q; then K is the totally ordered lattice of Figure 1.8. Setting $\mathcal{X}^{\downarrow n} := \mu^{-1}(\downarrow n)$ for $0 \leq n \leq 2$ gives the filtration

$$\varnothing \subset \mathcal{X}^{\downarrow 0} \subset \mathcal{X}^{\downarrow 1} \subset \mathcal{X}^{\downarrow 2} = \mathcal{X}' \qquad \varnothing \subset v_0 \subset v_0, v_2 \subset v_0, v_2, e_0, v_1, e_1$$

and setting $\mathcal{X}'^{\downarrow n} := \mu'^{-1}(\downarrow n)$ the filtration

$$\varnothing \subset \mathcal{X}'^{\downarrow 0} \subset \mathcal{X}'^{\downarrow 1} \subset \mathcal{X}'^{\downarrow 2} = \mathcal{X}' \qquad \varnothing \subset v_0' \subset v_0', v_1' \subset v_0', v_1', e_0'$$

[1] In this example we restrict to filtrations. However, we wish to emphasize that our results hold for multi-parameter persistence; see Section 3.9, in particular Theorem 3.9.3.

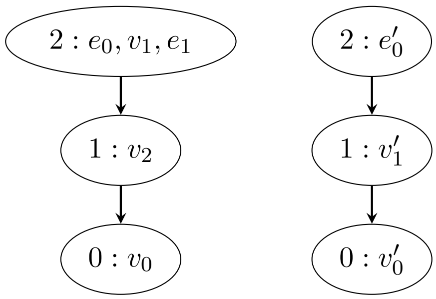

Figure 1.7: Visualization of $(\mathcal{X}, \mu), (\mathcal{X}', \mu')$.

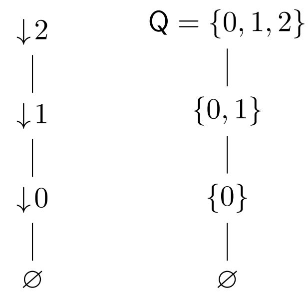

Figure 1.8: Lattice K of downsets of Q.

For each downset $\downarrow i$, the (graded) chain equivalence $\phi\colon C(\mathcal{X}) \to C(\mathcal{X}')$ induces a chain equivalence $\phi^{\downarrow i}\colon C(\mathcal{X}^{\downarrow i}) \to C(\mathcal{X}'^{\downarrow i})$ (see Section 3.4), which fit into the following commutative diagram:

$$\begin{tikzcd}
0 \arrow[r, hook] & C(\mathcal{X}^{\downarrow 0}) \arrow[r, hook] \arrow[d, "\phi^{\downarrow 0}"] & C(\mathcal{X}^{\downarrow 1}) \arrow[r, hook] \arrow[d, "\phi^{\downarrow 1}"] & C(\mathcal{X}) \arrow[d, "\phi"] \\
0 \arrow[r, hook] & C(\mathcal{X}'^{\downarrow 0}) \arrow[r, hook] & C(\mathcal{X}'^{\downarrow 1}) \arrow[r, hook] & C(\mathcal{X}')
\end{tikzcd}$$

As each $\phi^{\downarrow i}$ is a chain equivalence, and the diagram commutes, ϕ induces an isomorphism on the persistent homology. As a corollary, all computational tools that tabulate the persistent homology groups, e.g. the persistence diagrams and barcodes [11], can be computed via Δ'.

Example 1.2.4 (Computational Dynamics)**.** Consider the gradient flow $\varphi\colon \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}^2$ generated by the system in Figure 1.9. Restricting to the compact isolated invariant

$$\dot{x} = x - x^2$$
$$\dot{y} = y - y^2$$

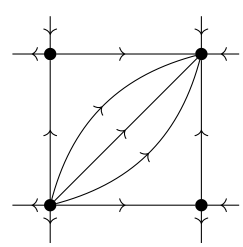

Figure 1.9: System of differential equations and associated phase portrait.

set $X = [0, 1] \times [0, 1]$, the phase space is *combinatorialized* with a cubical complex $\mathcal{X}$ (see Figure 1.10). Continuous dynamics are combinatorialized with a binary relation $\mathcal{F} \subset \mathcal{X}^+ \times \mathcal{X}^+$ where $\mathcal{X}^+ = \{\xi_i\}$ is the set of top-dimensional cubes. The relation $\mathcal{F}$ approximates the underlying flow φ. This is depicted in Figure 1.10.

Figure 1.10: $\mathcal{X}, \mathcal{F}$ and the lattice $\mathsf{Invset}^+(\mathcal{F})$.

The lattice of forward invariant sets of $\mathcal{F}$, denoted $\mathsf{Invset}^+(\mathcal{F})$, captures the long-term behavior of the relation $\mathcal{F}$. Let $\mathsf{Sub}_{Cl}(\mathcal{X})$ be the lattice of closed subcomplexes of $\mathcal{X}$. $\mathcal{F}$ is a good combinatorial approximation for φ since the map $c\colon \mathsf{Invset}^+(\mathcal{F}) \to \mathsf{Sub}_{Cl}(\mathcal{X})$ given by

$$\mathsf{Invset}^+ \ni a \mapsto \mathrm{cl}(a) \in \mathsf{Sub}_{Cl}(\mathcal{X})$$

is a lattice morphism. This implies that the collection $\{\mathcal{X}^a\}$, where $\mathcal{X}^a = c(a)$, is a lattice of attracting blocks for φ. Set $\mathsf{P} = \mathsf{J}(\mathsf{Invset}^+(\mathcal{F}))$. The cell complex $\mathcal{X}$ together with the Birkhoff dual (see Section 2.4.3) $\nu = \mathsf{J}(c)\colon \mathcal{X} \to \mathsf{P}$ is a P-graded cell complex $(\mathcal{X}, \nu)$ where $\nu^{-1}(a) = \mathcal{X}^a \setminus \mathcal{X}^{\overleftarrow{a}}$. See Figure 1.2.4. As a reminder, in our visualizations if q covers p in P then there is a directed edge $q \to p$. This orientation of directed edges respects the behavior of the dynamics and agrees with the orientation of the Conley theory literature [2, 6]. However, Hasse diagram is often given the opposite orientation outside of Conley theory. The P-graded complex $(\mathcal{X}, \nu)$ is the input into our algorithm

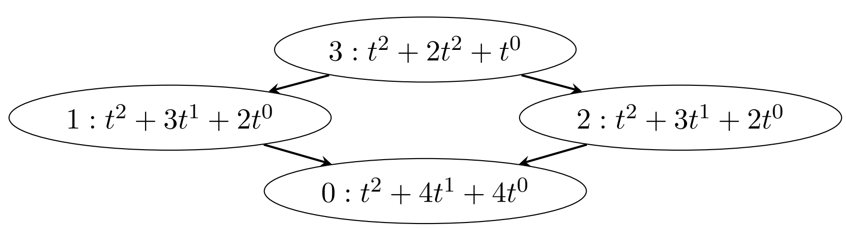

Figure 1.11: Visualization of $(\mathcal{X}, \nu)$.

ConnectionMatrix (Algorithm 3.7.8). The associated Conley complex (the output of algorithm) is in Figure 1.12.

Let $\mathsf{M} = \{\mathsf{M}(a)\}_{a \in \mathsf{J}(\mathsf{L})}$ where $\mathsf{M}(a) = \mathrm{Inv}(\mathcal{X}^a \setminus \mathcal{X}^{\overleftarrow{a}})$ and let $\mathsf{T} = \{\mathrm{cl}(\mathcal{X}^a \setminus \mathcal{X}^{\overleftarrow{a}})\}_{a \in \mathsf{J}(\mathsf{L})}$. M and T are posets by restricting the ordering of L. The individual sets $\mathsf{M}(a)$ are called

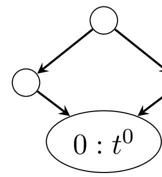

Figure 1.12: Conley-Morse Graph. Blank nodes have trivial Conley index.

Morse sets. Classically, M is called a *Morse decomposition for* φ. The pair $(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}})$ is an index pair for $\mathsf{M}(a)$ and there is an order-embedding $\mathsf{M} \hookrightarrow \mathsf{T}$. The diagram $\mathsf{M} \hookrightarrow \mathsf{T}$ is a *tessellated Morse decomposition* for φ. The connection matrix Δ for the Morse decomposition M is a boundary operator on the Conley indices of the Morse sets,

$$\Delta \colon \bigoplus_{a \in \mathsf{J}(\mathsf{L})} CH_\bullet \mathsf{M}(a) \to \bigoplus_{a \in \mathsf{J}(\mathsf{L})} CH_\bullet \mathsf{M}(a). \tag{1.3}$$

As each pair $(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}})$ is an index pair for $\mathsf{M}(a)$, (1.3) can be written as

$$\Delta \colon \bigoplus_{a \in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}) \to \bigoplus_{a \in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}).$$

This is precisely the output of the algorithm ConnectionMatrix; visualized in Figure 1.12.

Chapter 2

Preliminary material

In this chapter we recall the necessary mathematical prerequisites. For a more complete discussion the reader is referred to [9, 43] for order theory; [19, 51] for category theory and homological algebra; [30] for algebraic topology and cell complexes; [23, 38, 45] for discrete Morse theory; [26, 27, 28] for (computational) dynamics.

2.1 Notation

Boldface font is used to denote categories and Fraktur font to denote particular functors. Sans-serif font is used for order-theoretic structures, such as posets and lattices. Lower case Greek letters are used for morphisms of (graded, filtered) chain complexes. Script-like letters are used for chain complex braids and graded module braids and upper case Greek letters are used for morphisms of these objects. Calligraphic font is typically used for notation referring to combinatorial objects and cell complexes.

2.2 Category Theory

The exposition of category theory primarily follows [51], see also [19, 31].

Definition 2.2.1. A *preadditive category* or *Ab-category* is a category that is enriched over the category of abelian groups. That is, every hom-set $\mathrm{Hom}_{\mathcal{A}}(A, B)$ in $\mathcal{A}$ has the structure of an abelian group such that composition distributes over addition (the group operation). A *additive category* is a preadditive category with a zero object and a product $A \times B$ for every pair A, B of objects in $\mathcal{A}$.

Definition 2.2.2. In an additive category $\mathcal{A}$ a *kernel* of a morphism $f\colon B \to C$ is a map $i\colon A \to B$ such that $fi = 0$ and that is universal with respect to this property.

Dually, a *cokernel* of f is a map $e\colon C \to D$ which is universal with respect to having $ef = 0$. in $\mathcal{A}$, a map $i\colon A \to B$ is *monic* if $ig = 0$ implies $g = 0$ for every map $g\colon A' \to A$. A monic map is called a *monomorphism.* A map $e\colon C \to D$ is an *epi*, or *epimorphism*, if $he = 0$ implies $h = 0$ for every map $h\colon D \to D'$.

Definition 2.2.3. An *abelian category* is an additive category $\mathcal{A}$ such that

1. Every map in $\mathcal{A}$ has a kernel and a cokernel
2. Every monic in $\mathcal{A}$ is the kernel of its cokernel
3. Every epi in $\mathcal{A}$ is the cokernel of its kernel.

Let $F\colon \mathcal{A} \to \mathcal{B}$ be a functor. For a pair of objects A, B in $\mathcal{A}$, F induces a map on hom-sets

$$F_{A,B}\colon \mathrm{Hom}_{\mathcal{A}}(A, B) \to \mathrm{Hom}_{\mathcal{B}}(F(A), F(B))$$

Definition 2.2.4. An *additive functor* $F\colon \mathcal{A} \to \mathcal{B}$ between *Ab*-categories is a functor such that each $\mathrm{Hom}_{\mathcal{A}}(A, A') \to \mathrm{Hom}_{\mathcal{B}}(F(A), F(A'))$ is a group homomorphism.

Definition 2.2.5. A functor $F\colon \mathcal{A} \to \mathcal{B}$ is *full* if $F_{A,B}$ is surjective for all pairs A, B. F is *faithful* if $F_{A,B}$ is injective for all pairs A, B. A subcategory $\mathcal{A}$ of $\mathcal{B}$ is *full* if the inclusion functor $\mathcal{A} \hookrightarrow \mathcal{B}$ is full. A functor is *fully faithful* if it is both full and faithful. A functor $F\colon \mathcal{A} \to \mathcal{B}$ is *essentially surjective* if for any object B of $\mathcal{B}$ there is an object A in $\mathcal{A}$ such that B is isomorphic to $F(A)$. A functor $F\colon \mathcal{A} \to \mathcal{B}$ is an *equivalence of categories* if there is a functor $G\colon \mathcal{B} \to \mathcal{A}$ and natural isomorphisms $\epsilon\colon FG \to \mathrm{id}_{\mathcal{B}}$ and $\eta\colon \mathrm{id}_{\mathcal{A}} \to GF$. Categories $\mathcal{A}$ and $\mathcal{B}$ are *equivalent* if there a equivalence $F\colon \mathcal{A} \to \mathcal{B}$.

Proposition 2.2.6. *$F\colon \mathcal{A} \to \mathcal{B}$ is an equivalence of categories if and only if F is full, faithful and essentially surjective.*

Definition 2.2.7. A functor $F\colon \mathcal{A} \to \mathcal{B}$ is *conservative* if given a morphism $f\colon A \to B$ in $\mathcal{A}$, $F(f)\colon F(A) \to F(B)$ is an isomorphism only if f is an isomorphism.

Definition 2.2.8. Following [31, Section II.8], we say that a *congruence relation* $\sim$ on a category $\mathcal{A}$ is a collection of equivalence relations $\sim_{A,B}$ on $\mathrm{Hom}(A, B)$ for each pair of objects A, B such that the equivalence relations respect composition of morphisms. That is, if $f, f' \colon A \to B$ and $f \sim f'$ then for any $g \colon A' \to A$ and $h \colon B \to B'$ we have $hfg \sim hf'g$. If $\mathcal{A}$ is an additive category, we say a congruence relation $\sim$ is *additive* if $f_0, f_1, g_0, g_1 \colon A \to B$ with $f_i \sim g_i$ then $f_0 + f_1 \sim g_0 + g_1$.

Definition 2.2.9. Given a congruence relation $\sim$ on a category $\mathcal{A}$ the *quotient category* $\mathcal{A}/\!\sim$ is defined as the category whose objects are those of $\mathcal{A}$ and whose morphisms are equivalence classes of morphisms in $\mathcal{A}$. That is,

$$\mathrm{Hom}_{\mathcal{A}/\sim}(A, B) = \mathrm{Hom}_{\mathcal{A}}(A, B)/\!\sim_{A,B}$$

There is a quotient functor from $\mathcal{A} \to \mathcal{A}/\!\sim$ which is the identity on objects and sends each morphism to its equivalence class. If $\sim$ is additive then the quotient category $\mathcal{A}/\!\sim$ is additive, and the quotient functor $\mathcal{A} \to \mathcal{A}/\!\sim$ is an additive functor.

Proposition 2.2.10 ([31], Proposition II.8.1)**.** *Let $\sim$ be a congruence relation on the category $\mathcal{A}$. Let $F \colon \mathcal{A} \to \mathcal{B}$ be a functor such that $f \sim f'$ implies $F(f) = F(f')$ for all f and f', then there is a unique functor F' from $\mathcal{A}/\!\sim$ to $\mathcal{B}$ such that $F' \circ q = F$.*

Proposition 2.2.11. *Let $\mathcal{A}$ be a category, $\sim$ a congruence relation on $\mathcal{A}$ and $\mathcal{B} = \mathcal{A}/\!\sim$ be the quotient category. Let $\mathcal{A}'$ be a full subcategory of $\mathcal{A}$. Let $\mathcal{B}'$ be the full subcategory whose objects are the objects in $\mathcal{A}'$. Then $\mathcal{B}'$ is a quotient of $\mathcal{A}'$.*

2.3 Homological Algebra

2.3.1 In Additive Categories

Let $\mathcal{A}$ be an additive category. Unless explicitly stated, we assume that functors of additive categories are additive. While most of homological algebra takes place in the setting that $\mathcal{A}$ is an abelian category, the construction of the homotopy category may be done for an additive category $\mathcal{A}$. In this section we outline the construction.

Definition 2.3.1. A *chain complex* $(C_\bullet, \partial_\bullet)$ in $\mathcal{A}$ is a family $C_\bullet = \{C_n\}_{n \in \mathbb{Z}}$ of objects of $\mathcal{A}$ together together with morphsims $\partial_\bullet = \{\partial_n \colon C_n \to C_{n-1}\}_{n \in \mathbb{Z}}$, called *boundary operators*, or *differentials*, such that $\partial_{n-1} \circ \partial_n = 0$ for all n.

When the context is clear we abbreviate $(C_\bullet, \partial_\bullet)$ as (C, ∂) or more simply as just C. A morphism $\phi\colon C \to C'$ of chain complexes C and C' is a *chain map*, that is, a family $\phi = \{\phi_n\colon C_n \to C'_n\}_{n\in\mathbb{Z}}$ of morphisms in $\mathcal{A}$ such that $\phi_{n-1} \circ \partial_n = \partial'_n \circ \phi_n$ for all n. Chain complexes and chain maps constitute the category denoted $\mathbf{Ch} = \mathbf{Ch}(\mathcal{A})$.

Definition 2.3.2. A chain map $\phi\colon C \to C'$ is *null homotopic* if there exists a family $\gamma = \{\gamma_n\colon C_n \to C'_{n+1}\}_{n\in\mathbb{Z}}$ of morphisms in $\mathcal{A}$ such that

$$\phi = \gamma \circ \partial + \partial' \circ \gamma.$$

The morphisms $\{\gamma_n\}$ are called a *chain contraction* of ϕ.

Definition 2.3.3. Two chain maps $\phi, \psi\colon C \to C'$ are *chain homotopic* if their difference $\phi - \psi$ is null homotopic, that is, if there exists a family of morphisms $\{\gamma_n\}$ such that

$$\phi - \psi = \gamma \circ \partial + \partial' \circ \gamma.$$

The morphisms $\{\gamma_n\}$ are called a *chain homotopy* from ϕ to ψ. We write $\phi \sim \psi$ to indicate that ϕ, ψ are chain homotopic.

Definition 2.3.4. Given chain complexes C and C' in $\mathcal{A}$, a family $\gamma = \{\gamma_n\colon C_n \to C'_{n+1}\}$ of morphisms in $\mathcal{A}$ is called a *degree 1 map* from C to C'.

Definition 2.3.5. We say that $\phi\colon C \to D$ is a *chain equivalence* if there is a chain map $\psi\colon D \to C$ such that $f \circ g$ and $g \circ f$ are chain homotopic to the respective identity maps of C and D. We say that C and D are *chain equivalent* if there exists a chain equivalence $\phi\colon C \to D$.

Proposition 2.3.6. *Chain equivalence is an equivalence relation on the objects of* $\mathbf{Ch}(\mathcal{A})$.

Proposition 2.3.7. *The relation $\sim$ is an additive congruence relation on* $\mathbf{Ch}(\mathcal{A})$.

Definition 2.3.8. We define the homotopy category of $\mathbf{Ch}(\mathcal{A})$, which we denote by $\mathbf{K} = \mathbf{K}(\mathcal{A})$, to be the category whose objects are chain complexes and whose morphisms are chain homotopy equivalence classes of chain maps between chain complexes. In other words, $\mathbf{K}(\mathcal{A})$ is the quotient category $\mathbf{Ch}(\mathcal{A})/{\sim}$ formed by defining hom-sets

$$\mathrm{Hom}_{\mathbf{K}(\mathcal{A})}(A, B) = \mathrm{Hom}_{\mathbf{Ch}(\mathcal{A})}(A, B)/{\sim}$$

We define the quotient functor $q\colon \mathbf{Ch}(\mathcal{A}) \to \mathbf{K}(\mathcal{A})$ to be the functor which sends each chain complex to itself and each chain map to its chain homotopy equivalence class.

It follows from the construction of $\mathbf{K}(\mathcal{A})$ that two chain complexes are isomorphic in $\mathbf{K}(\mathcal{A})$ if and only if they are chain equivalent.

Proposition 2.3.9. *If $F\colon \mathcal{A} \to \mathcal{B}$ is a functor of additive categories then there is an associated functor $F_{\mathbf{Ch}}\colon \mathbf{Ch}(\mathcal{A}) \to \mathbf{Ch}(\mathcal{B})$ given by*

$$F_{\mathbf{Ch}}(C, \partial) = \Big(\{F(C_n)\}_{n\in\mathbb{Z}}, \{F(\partial_n\colon C_n \to C_{n-1})\}_{n\in Z}\Big)$$

Moreover, since F is additive it induces a functor $F_{\mathbf{K}}\colon \mathbf{K}(\mathcal{A}) \to \mathbf{K}(\mathcal{B})$ between the homotopy categories $\mathbf{K}(\mathcal{A})$ and $\mathbf{K}(\mathcal{B})$.

Proposition 2.3.10. *If $F\colon \mathcal{A} \to \mathcal{B}$ is a equivalence of categories then the induced functor $F_{\mathbf{Ch}}$ is an equivalence of categories.*

2.3.2 In Abelian Categories

Let $\mathcal{A}$ be an abelian category. Let C be a chain complex in $\mathcal{A}$. The elements of C_n are called the n-chains. The elements of $\ker \partial_n \subset C_n$ are called the *n-cycles*; the elements of $\operatorname{im} \partial_{n+1} \subset C_n$ are called the *n-boundaries.* Since $\mathcal{A}$ is an abelian category we may define the notion of homology of a chain complex.

Definition 2.3.11. We say that a chain complex (C, ∂) is *minimal* if $\partial_n = 0$ for all n. The minimal chain complexes form the full subcategory $\mathbf{Ch}_0 \subset \mathbf{Ch}$. A chain complex is called *acyclic* if it is exact, i.e., $\ker \partial = \operatorname{im} \partial$. A chain complex B is a *subcomplex* of C if each B_n is a subspace of C_n and $\partial^B = \partial^C|_B$. That is, the inclusions $\{i_n\colon B_n \to C_n\}$ form a chain map $i\colon B \to C$. If $\phi\colon A \to B$ is a chain map then $\ker(\phi)$ and $\operatorname{im}(\phi)$ are subcomplexes of A and B respectively. Suppose B is a subcomplex of C. The *quotient complex* C/B is the chain complex consisting of the family $\{C_n/B_n\}_{n\in\mathbb{Z}}$ together with differentials $\{x + B_n \mapsto \partial_n(x) + B_{n-1}\}_{n\in\mathbb{Z}}$. The n-th *homology* of C is the quotient $H_n(C) := \ker \partial_n / \operatorname{im} \partial_{n+1}$. We define the *homology* of C as $H_\bullet(C) := \{H_n(C)\}_{n\in\mathbb{Z}}$ equipped with boundary operators $\{0\colon H_n(C) \to H_{n-1}(C)\}_{n\in\mathbb{Z}}$ and regard it as a minimal chain complex.

A chain complex C is acyclic if and only if $H_\bullet(C) = 0$. Chain maps induce morphisms on homology: let C, C' be chain complexes and $\phi\colon C \to C'$ a chain map. There there exists a well-defined map $H_\bullet(\phi)$ called the *induced map on homology* given via

$$H_n(\phi) : z + \operatorname{im} \partial_{n+1} \mapsto \phi(z) + \operatorname{im} \partial'_{n+1}$$

Proposition 2.3.12. *Homology is a functor $H_\bullet\colon \mathbf{Ch} \to \mathbf{Ch}_0$. For each $n \in \mathbb{Z}$ the n-th homology is a functor $H_n\colon \mathbf{Ch} \to \mathcal{A}$.*

We often write $H_\bullet$ more simply as H.

Proposition 2.3.13. *Chain homotopic maps induce the same map on homology.*

Proposition 2.3.14. *A chain equivalence $\phi\colon C \to D$ induces an isomorphism on the homology $H(\phi)\colon H(C) \to H(D)$.*

The category $\mathbf{K}$ enjoys a universal property with respect to chain equivalences.

Proposition 2.3.15 ([51], Proposition 10.1.2)**.** *Let $F\colon \mathbf{Ch} \to D$ be any functor that sends chain equivalences to isomorphisms. Then F factors uniquely through $\mathbf{K}$. In particular, there exists a unqiue functor $H_{\mathbf{K}}\colon \mathbf{K} \to \mathbf{Ch}_0$ such that $H_{\mathbf{K}} \circ q = H$.*

Let $\mathbf{Vect}$ be the category of finite-dimensional vector spaces over a field $\mathbb{K}$. We have the following result for $\mathbf{Ch}(\mathbf{Vect})$.

Proposition 2.3.16 ([19], Proposition III.2.4)**.** *The pair of functors $q \circ i\colon \mathbf{Ch}_0(\mathbf{Vect}) \to \mathbf{K}(\mathbf{Vect})$ and $H_{\mathbf{K}}\colon \mathbf{K}(\mathbf{Vect}) \to \mathbf{Ch}_0(\mathbf{Vect})$ form an equivalence of categories.*

We can give an alternative to Proposition 2.3.16 using the results and perspectives of this paper. This will use Algorithm 3.7.2 and give the flavor of Theorem 3.8.1. Let $\mathbf{K}_0(\mathbf{Vect})$ denote the full subcategory of $\mathbf{K}(\mathbf{Vect})$ whose objects are the objects of are the objects of $\mathbf{Ch}_0(\mathbf{Vect})$ and whose morphisms are given by

$$\operatorname{Hom}_{\mathbf{K}_0}(C, D) = \operatorname{Hom}_{\mathbf{Ch}_0}(C, D)/\sim$$

There is a quotient functor $q\colon \mathbf{Ch}_0(\mathbf{Vect}) \to \mathbf{K}_0(\mathbf{Vect})$. The next result shows that $\mathbf{Ch}_0(\mathbf{Vect})$ may be identified with $\mathbf{K}_0(\mathbf{Vect})$.

Proposition 2.3.17. *The quotient functor* $q\colon \mathbf{Ch}_0(\mathbf{Vect}) \to \mathbf{K}_0(\mathbf{Vect})$ *is an isomorphism on hom-sets.*

Proof. Given $\psi, \phi\colon C \to D$ we have that $\psi \sim \phi$ if there exists γ such that

$$\psi - \phi = \gamma \circ \partial + \partial \circ \gamma = 0$$

Where the last equality follows since $\partial = 0$ within the subcategory of minimal objects. Thus $\psi \sim \phi$ if and only if $\psi = \phi$. □

Proposition 2.3.18. *The inclusion functor* $i\colon \mathbf{K}_0(\mathbf{Vect}) \to \mathbf{K}(\mathbf{Vect})$ *is an equivalence of categories.*

Proof. The functor i is full and faithful. Moreover, i is essentially surjective from the Theorem 3.7.3, which is the proof of correctness of Algorithm 3.7.2 (Homology). It follows from Proposition 2.2.6 that i is an equivalence of categories. □

This result implies there is an inverse functor to i, call it F, such that i and F are an equivalence of categories. In particular, $i \circ F(C)$ is minimal and $i \circ F(C)$ and C are chain equivalent.

2.4 Order Theory

2.4.1 Posets

Definition 2.4.1. A *partial order* $\leq$ is a reflexive, antisymmetric, transitive binary relation. A set $\mathsf{P} = (\mathsf{P}, \leq)$ together with a partial order is called a *partially ordered set*, or *poset*. We let $<$ be the relation on P such that $x < y$ if and only if $x \leq y$ and $x \neq y$. A function $\nu\colon \mathsf{P} \to \mathsf{Q}$ is *order-preserving* if $p \leq q$ implies that $\nu(p) \leq \nu(q)$. The *category of finite posets*, denoted **FPoset**, is the category whose objects are finite posets and whose morphisms are order-preserving maps.

Definition 2.4.2. Let P be a finite poset and $p, q \in \mathsf{P}$. We say that q and p are *comparable* if $p \leq q$ or $q \leq p$. We say that q *covers* p if $p < q$ and there does not exist an r with $p < r < q$. If q covers p then p is a *predecessor* of q. Let $Q \subset \mathsf{P}$. We say that Q is a *chain* in P if any two elements in Q are comparable. We say that p and q are *incomparable* if they are not comparable. We say that Q is an *antichain* in P if any two elements in Q are incomparable.

Definition 2.4.3. Let P be a finite poset. An *upper set* of P is a subset $U \subset \mathsf{P}$ such that if $p \in \mathsf{U}$ and $p \leq q$ then $q \in U$. For $p \in \mathsf{P}$ the *upper set at* p is $\uparrow p := \{q \in \mathsf{P} : p \leq q\}$. Following [27], we denote the collection of upper sets by $\mathsf{U}(\mathsf{P})$. A *down set* of P is a set $D \subset \mathsf{P}$ such that if $q \in D$ and $p \leq q$ then $p \in D$. The *down set at* q is $\downarrow q := \{p \in \mathsf{P} : p \leq q\}$. Following [27], we denote the collection of down sets by $\mathsf{O}(\mathsf{P})$.

Remark 2.4.4. Any down set can be obtained by a union of down sets of the form $\downarrow q$. In fact, $\mathsf{O}(\mathsf{P})$ are the closed sets of the Alexandroff topology of the poset P. Under a poset morphism, the preimage of a down set is a down set. Similarly, the preimage of an upper set is an upper set.

Definition 2.4.5. Let P be a finite poset. For $p, q \in \mathsf{P}$ the *interval* from p to q, denoted $[p, q]$, is the set $\{x \in \mathsf{P} : p \leq x \leq q\}$. A subset $I \subset \mathsf{P}$ is *convex* if whenever $p, q \in I$ then $[p, q] \subset I$. Following [17], we denote the collection of convex sets by $I(\mathsf{P})$.

Remark 2.4.6. Let P be a finite poset. Any convex set of P can be obtained by an intersection of a down and upper set. Under a poset morphism the preimage of a convex set is a convex set. See [43].

Remark 2.4.7. In [17, 15, 16] convex sets are instead called intervals. We adopt the terminology convex as this is standard in order theory literature.

2.4.2 Lattices

Some texts introduce lattices as a particular type of poset. Instead, we begin with definition of lattice as an algebraic structure. For a discussion of the relationship of these two definitions the reader may consult the section 'Lattices as algebraic structures' within Chapter 2 of [9]. In particular, see [9, Theorem 2.9].

Definition 2.4.8. A *lattice* is a set L with the binary operations $\vee, \wedge \colon \mathsf{L} \times \mathsf{L} \to \mathsf{L}$ satisfying the following four axioms:

1. (idempotent) $a \wedge a = a \vee a = a$ for all $a \in \mathsf{L}$

2. (commutative) $a \wedge b = b \wedge a$ and $a \vee b = b \vee a$ for all $a, b \in \mathsf{L}$

3. (associative) $a \wedge (b \wedge c) = (a \wedge b) \wedge c$ and $a \vee (b \vee c) = (a \vee b) \vee c$ for all $a, b, c \in \mathsf{L}$

4. (absorption) $a \wedge (a \vee b) = a \vee (a \wedge b) = a$ for all $a, b \in \mathsf{L}$

   A lattice L is *distributive* if it satisfies the additional axiom:

5. (distributive) $a \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c)$ and $a \vee (b \wedge c) = (a \vee b) \wedge (a \vee c)$ for all $a, b, c \in \mathsf{L}$

   A lattice L is *bounded* if there exist *neutral* elements 0 and 1 that satisfy the following property:

6. $0 \wedge a = 0, 0 \vee a = a, 1 \wedge a = a, 1 \vee a = 1$ for all $a \in \mathsf{L}$

A lattice homomorphism $f\colon \mathsf{L} \to \mathsf{M}$ is a map such that if $a, b \in \mathsf{L}$ then $f(a \wedge b) = f(a) \wedge f(b)$ and $f(a \vee b) = f(a) \vee f(b)$. If L and M are bounded lattices then we also require that $f(0) = 0$ and $f(1) = 1$. In particular, we are interested in finite lattices. Every finite lattice is bounded. A subset $\mathsf{K} \subset \mathsf{L}$ is a sublattice of L if $a, b \in \mathsf{K}$ implies that $a \vee b \in \mathsf{K}$ and $a \wedge b \in \mathsf{K}$. For sublattices of bounded lattices we impose the extra condition that $0, 1 \in \mathsf{K}$.

Definition 2.4.9. The *category of finite distributive lattices*, denoted **FDLat**, is the category whose objects are finite distributive lattices and whose morphisms are lattice homomorphisms.

A lattice L has an associated poset structure given by $a \leq b$ if $a = a \wedge b$ or, equivalently, if $b = a \vee b$.

Definition 2.4.10. An element $a \in \mathsf{L}$ is *join-irreducible* if it has a unique predecessor; given a join-irreducible a we denote its unique predecessor by $\overleftarrow{a}$. The set of join-irreducible elements of L is denoted by $\mathsf{J}(\mathsf{L})$. $(\mathsf{J}(\mathsf{L}), \leq)$ is a poset, where the order $\leq$ is the restriction of the partial order of L.

Definition 2.4.11. For $a \in \mathsf{L}$ the expression

$$a = b_1 \vee \cdots \vee b_n$$

where the b_i's are distinct join-rreducibles is called *irredundant* if it is not the join of any proper subset of $U = \{b_1, \ldots, b_n\}$. Clearly, if the join is irredundant, then U is an antichain.

Proposition 2.4.12 ([43],Theorem 4.29)**.** *If* L *is a finite distributive lattice then every* $a \in \mathsf{L}$ *has an irredundant join-irreducible representation*

$$a = b_1 \vee \cdots \vee b_n$$

and all such representations have the same number of terms.

Definition 2.4.13. Let L be a lattice with a minimum element 0. An element $a \in \mathsf{L}$ is an *atom* if a covers 0.

Definition 2.4.14. A *complemented lattice*, also called a *Boolean algebra*, is a bounded lattice (with least element 0 and greatest element 1), in which every element a has a complement, i.e., an element b such that $a \vee b = 1$ and $a \wedge b = 0$.

Definition 2.4.15. Let X be a finite set. The *power set* of X is the collection of all subsets of X. Following [25], we denote the power set of X as $\mathsf{Set}(X)$. The power set $\mathsf{Set}(X)$ is a Boolean algebra and the atoms of $\mathsf{Set}(X)$ are the elements of X.

Definition 2.4.16. A *relatively complemented* lattice is a lattice such that every interval $[a, b]$, viewed as a bounded lattice, is complemented.

Example 2.4.17. Let V be a vector space. The associated *lattice of subspaces*, denoted by $\mathsf{Sub}(V)$, consists of all subspaces of V with the operations $\wedge := \cap$ and $\vee := +$ (span). $\mathsf{Sub}(V)$ is a relatively complemented lattice. It is not distributive in general.

Definition 2.4.18. Let C be a chain complex. The associated *lattice of subcomplexes*, denoted by $\mathsf{Sub}(C)$, consists of all subcomplexes of C with the operations $\wedge := \cap$ and $\vee := +$ (span), i.e.,

$$(A_\bullet, \partial^A) \wedge (B_\bullet, \partial^B) := (A_\bullet \cap B_\bullet, \partial^C|_{A \cap B}),$$
$$(A_\bullet, \partial^A) \vee (B_\bullet, \partial^B) := (A_\bullet + B_\bullet, \partial^C|_{A+B}).$$

$\mathsf{Sub}(C)$ is a bounded lattice, but is not distributive in general.

2.4.3 Birkhoff's Theorem and Transforms

As indicated above, given a finite distributive lattice L, $\mathsf{J}(\mathsf{L})$ has a poset structure. In the opposite direction, given a finite poset $(\mathsf{P}, \leq)$ the collection of downsets $\mathsf{O}(\mathsf{P})$ is a bounded distributive lattice under $\wedge = \cap$ and $\vee = \cup$. The following theorem often goes under the moniker 'Birkhoff's Representation Theorem'.

Theorem 2.4.19 ([43], Theorem 10.4)**.** J *and* O *are contravariant functors from* **FDLat** *to* **FPoset** *and* **FPoset** *to* **FDLat***, respectively. Following [25], we represent this via the following diagram.*

$$\begin{tikzcd}
\mathsf{K} \arrow[d, "h"'] & & \mathsf{J}(\mathsf{K}) & \mathsf{P} \arrow[d, "\nu"'] & & \mathsf{O}(\mathsf{P}) \\
\mathsf{L} \arrow[r, phantom, "\overset{\mathsf{J}}{\Longrightarrow}"] & & \mathsf{J}(\mathsf{L}) \arrow[u, "\mathsf{J}(h)"'] & \mathsf{Q} \arrow[r, phantom, "\overset{\mathsf{O}}{\Longrightarrow}"] & & \mathsf{O}(\mathsf{Q}) \arrow[u, "\mathsf{O}(\nu)"']
\end{tikzcd}$$

The formulas for the morphisms $\mathsf{J}(h)$ *and* $\mathsf{O}(\nu)$ *are given by*

$$\mathsf{J}(h)(a) = \min h^{-1}(\uparrow a), \qquad \textit{where } a \in \mathsf{J}(\mathsf{L}), \tag{2.1}$$

$$\mathsf{O}(\nu)(a) = \nu^{-1}(a), \qquad \textit{where } a \in \mathsf{O}(\mathsf{Q}). \tag{2.2}$$

Furthermore,

$$\mathsf{L} \cong \mathsf{O}(\mathsf{J}(\mathsf{L})) \quad \textit{and} \quad \mathsf{P} \cong \mathsf{J}(\mathsf{O}(\mathsf{P})).$$

The pair of contravariant functors O and J called the *Birkhoff transforms*. Given $\nu \colon \mathsf{P} \to \mathsf{Q}$ we say that $\mathsf{O}(\nu)$ is the *Birkhoff dual* to ν. Similarly, for $h \colon \mathsf{K} \to \mathsf{L}$ we say that $\mathsf{J}(h)$ is the *Birkhoff dual* to h.

Example 2.4.20. Consider the poset P of Example 1.2.1, recalled in Figure 2.1(a). The lattice of down-sets $\mathsf{O}(\mathsf{P})$ is given in Figure 2.1(b) and the join-irreducibles $\mathsf{J}(\mathsf{O}(\mathsf{P}))$ in Figure 2.1(c).

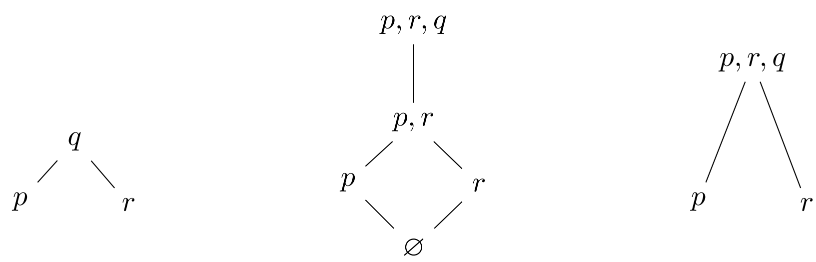

(a) (b) (c)

Figure 2.1: (a) Poset P (b) Lattice of down-sets $\mathsf{O}(\mathsf{P})$ (c) Join-irreducibles $\mathsf{J}(\mathsf{O}(\mathsf{P}))$.

2.4.4 Relations

Binary relations on finite sets arise in computational dynamics as discrete approximations to continuous dynamics [25, 26]. Relations are often viewed from three different (yet equivalent) perspectives, depending upon the intended application. First, as a true binary relation. Second, as a *multivalued map* $\mathcal{F}\colon X \rightrightarrows X$, where $(x, y) \in \mathcal{F}$ if and only if $y \in \mathcal{F}(x) := \{z : (x, z) \in \mathcal{F}\}$. Third, as a directed graph with verties X and edges given by $(x, y) \in \mathcal{F}$ if and only if there is a directed edge $x \to y$. In this section we briefly cover some tools used to manipulate relations. The exposition of relations follows [25].

Definition 2.4.21. A relation $\mathcal{F}$ on a finite set X that is both symmetric and transitive is called a *partial equivalence relation.* Given $x \in X$ the *partial equivalence class* of x is the (possibly empty) set $[x]_\mathcal{F} = \{y \in X : (x, y) \in X\}$. If, in addition, $\mathcal{F}$ is reflexive then $\mathcal{F}$ is an *equivalence relation.* In this case $[x]_\mathcal{F}$ is called an *equivalence class* of x.

Definition 2.4.22. Let $\mathcal{F}$ be a relation on a finite set X. We say that x is *reachable* from y, denoted by $x \leftarrow y$, if there exists a sequence $\{x_0, x_1, \ldots, x_k\}$ with $x_0 = y$ and $x_k = x$ such that $x_{i+1} \in \mathcal{F}(x_i)$. The reachability relation is the transitive closure $\mathcal{F}^+$ of the relation $\mathcal{F}$. If $x \leftarrow y$ and $y \leftarrow x$ then x and y are *connected*, denoted by $x \leftrightarrow y$. The relation $\leftrightarrow$ is a partial equivalence relation, which we call *connectivity.* The reflexive closure of the $\leftrightarrow$, denoted by $(\leftrightarrow)^=$, is an equivalence relation which we call *strong connectivity.* The equivalence classes are called *strongly connected components* and are denoted by $\mathsf{SC}(\mathcal{F})$.

The set $\mathsf{SC}(\mathcal{F})$ of strongly connected components of $\mathcal{F}$ is partially ordered via the reachability relation as follows. For $[x], [y] \in \mathsf{SC}(\mathcal{F})$ we say $[x] \leq [y]$ if and only if there exists $x' \in [x]$ and $y' \in [y]$ such that $x' \leftarrow y'$. It is shown in [25] that $\mathsf{SC}(-)$ is a covariant functor from the category finite binary relations to the category of finite posets.

2.5 Cell Complexes

Since our ultimate focus is on data analysis, we are interested in combinatorial topology. We make use of the following complex, whose definition is inspired by [30, Chapter III (Definition 1.1)]. Recall that $\mathbb{K}$ is a field.

Definition 2.5.1. A *cell complex* $(\mathcal{X}, \leq, \kappa, \dim)$ is an object $(\mathcal{X}, \leq)$ of **FPoset** together with two associated functions $\dim\colon \mathcal{X} \to \mathbb{N}$ and $\kappa\colon \mathcal{X} \times \mathcal{X} \to \mathbb{K}$ subject to the following conditions:

1. $\dim\colon (\mathcal{X}, \leq) \to (\mathbb{N}, \leq)$ is a poset morphism;
2. For each ξ and ξ' in $\mathcal{X}$:
$$\kappa(\xi, \xi') \neq 0 \quad \text{implies } \xi' \leq \xi \quad \text{and} \quad \dim(\xi) = \dim(\xi') + 1;$$
3. For each ξ and ξ'' in $\mathcal{X}$,
$$\sum_{\xi' \in X} \kappa(\xi, \xi') \cdot \kappa(\xi', \xi'') = 0.$$

For simplicity we typically write $\mathcal{X}$ for $(\mathcal{X}, \leq, \kappa, \dim)$. The partial order $\leq$ is the *face partial order*. $\mathcal{X}$ is a graded set with respect to dim, i.e., $\mathcal{X} = \bigsqcup_{n \in \mathbb{N}} \mathcal{X}_n$ with $\mathcal{X}_n = \dim^{-1}(n)$. An element $\xi \in \mathcal{X}$ is called a *cell* and $\dim \xi$ is the *dimension* of ξ. The function κ is the *incidence function* of the complex. The values of κ are referred to as the *incidence numbers*.

Definition 2.5.2. Given a cell complex $\mathcal{X}$ the *associated chain complex* $C(\mathcal{X})$ is the chain complex $C(\mathcal{X}) = \{C_n(\mathcal{X})\}_{n \in \mathbb{Z}}$ where $C_n(\mathcal{X})$ is the vector space over $\mathbb{K}$ with basis elements given by the cells $\xi \in \mathcal{X}_n$ and the boundary operator $\partial_n\colon C_n(\mathcal{X}) \to C_{n-1}(\mathcal{X})$

is defined by

$$\partial_n(\xi) := \sum_{\xi' \in \mathcal{X}} \kappa(\xi, \xi')\xi'.$$

Condition (3) of Definition 2.5.1 ensures $\partial_{n-1}\partial_n = 0$.

Definition 2.5.3. Given a cell complex $\mathcal{X}$ the *homology* of $\mathcal{X}$, denoted $H_\bullet(\mathcal{X})$, is defined as the homology of the associated chain complex $H_\bullet(C_\bullet(\mathcal{X}))$.

Definition 2.5.4. A cell complex $\mathcal{X}$ is *minimal* if the associated chain complex $(C(\mathcal{X}), \partial)$ is minimal.

Definition 2.5.5. Consider $\mathcal{K} \subset \mathcal{X}$ and let $(\leq', \kappa', \dim')$ be the restriction of $(\leq, \kappa, \dim)$ to $\mathcal{K}$. $(\mathcal{K}, \leq', \kappa', \dim')$ is a *subcomplex* of $\mathcal{X}$ if $(\mathcal{K}, \leq', \kappa', \dim')$ is a cell complex.

Remark 2.5.6. Given any subcomplex $\mathcal{K} \subset \mathcal{X}$ there is an associated chain complex $C(\mathcal{K})$. However the inclusion $\mathcal{K} \subset \mathcal{X}$ need not induce a chain map $C(\mathcal{K}) \to C(\mathcal{X})$. In other words, the associated chain complex $C(\mathcal{K})$ need not be a subcomplex of $C(\mathcal{X})$. For example, let $\mathcal{X}$ be as in Example 1.2.1 and set $\mathcal{K} = \{e_0, e_1\}$. $(\mathcal{K}, \leq', \kappa', \dim')$ is itself a cell complex.

Proposition 2.5.7. *Let $\mathcal{X}$ be a cell complex. If $I \subset \mathcal{X}$ is convex then $(I, \leq, \kappa, \dim)$ is a subcomplex.*

Definition 2.5.8. A subcomplex $\mathcal{K} \subset \mathcal{X}$ is *closed* if $\mathcal{K}$ is a down-set of $\mathcal{X}$. A subcomplex $\mathcal{K} \subseteq \mathcal{X}$ is *open* if it is an upper set of $\mathcal{X}$.

Proposition 2.5.9. *Let $\mathcal{K} \subset \mathcal{X}$ be a closed subcomplex. Then $C(\mathcal{K})$ is a subcomplex of $C(\mathcal{X})$.*

Remark 2.5.10. If $\mathcal{U}$ is an open subcomplex of $\mathcal{X}$ then $\mathcal{X} \setminus \mathcal{U}$ is a closed subcomplex. Therefore $C(\mathcal{U}) \cong C(\mathcal{X})/C(\mathcal{X} \setminus \mathcal{U})$. Thus open subcomplexes correspond to quotient complexes of $\mathcal{X}$. If $I \subset \mathcal{X}$ is convex then $\mathcal{K} := \downarrow I \setminus I$ is a downset; $C(I)$ is isomorphic to the subquotient $C(I) \cong C(\downarrow I)/C(\mathcal{K})$.

Definition 2.5.11. Given a cell complex $(\mathcal{X}, \leq, \kappa, \dim)$, the *lattice of closed subcomplexes* of $\mathcal{X}$ is $\mathsf{Sub}_{Cl}(\mathcal{X}) := \mathsf{O}(\mathcal{X}, \leq)$. There is a lattice monomorphism span: $\mathsf{Sub}_{Cl}(\mathcal{X}) \to$

$\mathsf{Sub}(C(\mathcal{X}))$ given by

$$a \mapsto \operatorname{span}(a) = \left\{ \sum_{i=0}^{n} \lambda_i \xi_i : n \in \mathbb{N}, \lambda_i \in \mathbb{K}, \xi_i \in a \right\}.$$

The *lattice of closed subcomplexes of* $C(\mathcal{X})$ is defined as $\mathsf{Sub}_{Cl}(C(\mathcal{X})) := \operatorname{im} \operatorname{span}$. A subcomplex of $C(\mathcal{X})$ which belongs to $\mathsf{Sub}_{Cl}(C(\mathcal{X}))$ is a *closed subcomplex* of $C(\mathcal{X})$.

Remark 2.5.12. The lattices $\mathsf{Sub}_{Cl}(C(\mathcal{X}))$ and $\mathsf{Sub}_{Cl}(\mathcal{X})$ are isomorphic. This implies that $\mathsf{Sub}_{Cl}(C(\mathcal{X}))$ is a distributive lattice, whereas in general $\mathsf{Sub}(C(\mathcal{X}))$ is not distributive. The lattice morphism span factors as

$$\mathsf{Sub}_{Cl}(\mathcal{X}) \xrightarrow{\text{span}} \mathsf{Sub}_{Cl}(C(\mathcal{X})) \hookrightarrow \mathsf{Sub}(C(\mathcal{X})).$$

We define the star and closures:

$$\operatorname{star}(\xi) := \uparrow\xi = \{\xi' : \xi \leq \xi'\} \qquad \text{and} \qquad \operatorname{cl}(\xi) := \downarrow\xi = \{\xi' : \xi' \leq \xi\}.$$

The star defines an open subcomplex while the closure defines a closed subcomplex. In order-theoretic terms these are the up and down sets of $(\mathcal{X}, \leq)$ at ξ. We use the duplicate notation star, cl to agree with the literature of cell complexes.

Definition 2.5.13. Let $\mathcal{K}$ be a subcomplex of $\mathcal{X}$. A cell $\xi \in \mathcal{K}$ is *interior to* $\mathcal{K}$ if $\operatorname{star}(\xi) \subseteq \mathcal{K}$. We denote by $\operatorname{int}(\mathcal{K})$ the set of interior cells. The *frontier* of $\mathcal{K}$ is defined as $\operatorname{fr} \mathcal{K} = \mathcal{K} \setminus \operatorname{int}\mathcal{K}$. If $\xi \in \mathcal{K}$ then $\xi \in \operatorname{fr} \mathcal{K}$ if and only if $\operatorname{star} \xi \not\subset \mathcal{K}$.

The interior of $\mathcal{K}$ satisfies the following identity:

$$\operatorname{int}(\mathcal{K}) = \bigcup_{\xi \in \operatorname{int}\mathcal{K}} \operatorname{star}(\xi).$$

It follows from this identity that $\operatorname{int}(\mathcal{K})$ is an upper set of $\mathcal{X}$.

Definition 2.5.14. Given a complex $\mathcal{X}$, a cell $\xi \in \mathcal{X}$ is a *top-cell* if it is maximal with respect to $\leq$, i.e., $\operatorname{star}(\xi) = \{\xi\}$. Following [23], we denote the set of top-cells is denoted $\mathcal{X}^+ \subset \mathcal{X}$. A cell complex $\mathcal{X}$ is called *pure* if there is an $n \in \mathbb{N}$ such that $\mathcal{X}^+ = \mathcal{X}_n$. In this case, n is called the *dimension* of $\mathcal{X}$.

Definition 2.5.15. Given a subcomplex $\mathcal{X}' \subseteq \mathcal{X}$, a pair of cells $(\xi, \xi') \in \mathcal{X}' \times \mathcal{X}'$ is a *coreduction pair* in $\mathcal{X}'$ if $\partial(\xi) = \kappa(\xi, \xi')\xi'$ with $\kappa(\xi, \xi') \neq 0$. A cell $\xi \in \mathcal{X}$ is *free in* $\mathcal{X}'$ if $\kappa(\xi, \xi') = 0$ for $\xi' \in \mathcal{X}'$.

Definition 2.5.16. Let $\mathcal{X}$ be a cell complex. The *f-vector* of $\mathcal{X}$ is the integral sequence

$$(f_0, f_1, f_2, \dots)$$

where f_i is the number of i-dimensional cells. The *f-polynomial* of $\mathcal{X}$ is the polynomial

$$\mathcal{F}_{\mathcal{X}}(t) = \sum_i f_i t^i.$$

The *Poincare polynomial* of $\mathcal{X}$ is the polynomial

$$P_{\mathcal{X}}(t) = \sum_i \dim H_i(\mathcal{X}) t^i.$$

2.6 Cubical Complexes

Cubical complexes often arise computational dynamics as grids on a phase space, e.g. [2, 6, 24], and therefore have a particular importance for Conley theory computations. In this section we give a brief review of cubical complexes which follows [24]. An *elementary interval* is a subset $I \subset \mathbb{R}$ of the form $I = [l, l+1]$ or $I = [l, l]$ for some $l \in \mathbb{Z}$. An *elementary cube* ξ in $\mathbb{R}^n$ is a finite product of elementary intervals, i.e.,

$$\xi = I_1 \times I_2 \times \cdots \times I_n \subseteq \mathbb{R}^n.$$

A *cubical set* is a union of elementary cubes. Intervals of length zero are called *degenerate* while those of length 1 are *nondegenerate*. The *dimension* of a cube ξ, denoted $\dim \xi$, is the number of nondegenerate intervals in ξ. Therefore dim is a function $\dim\colon \mathcal{X} \to \mathbb{N}$. If $\xi \subseteq \xi'$ are elementary cubes then ξ is a *face* of ξ'. If, in addition, $\dim \xi = \dim \xi' - 1$ then ξ is a *primary face* of ξ. The face relation defines a poset $\leq$ on the cubical set. Given a particular field $\mathbb{K}$, an incidence function $\kappa\colon \mathcal{X} \times \mathcal{X} \to \mathbb{K}$ can be described in detail [24]. When $\mathbb{K} = \mathbb{Z}_2$ the incidence number $\kappa(\xi, \xi')$ is nonzero if and only if ξ' is a primary face of ξ.

Definition 2.6.1. A *cubical complex* is a cubical set such that $(\mathcal{X}, \leq, \kappa, \dim)$ form a complex.

Definition 2.6.2. We call the indices $\{1, \dots, n\}$ the set of *coordinates*. An elementary cube $\xi = I_1 \times I_2 \times \cdots \times I_n \subseteq \mathbb{R}^n$ is said to have *extent* in coordinate m if I_m is nondegenerate.

2.7 Polyhedral complexes

When the cells come equipped with realization maps that form convex regions of Euclidean space we call this a polyhedral complex. The next definition follows [30, Chapter III (Definition 6.1)]. The primary difference is that we make an explicit distinction between the abstract cell complex $\mathcal{X}$ and the geometric realization of cells in Euclidean space.

Definition 2.7.1. A *polyhedral complex* in $\mathbb{R}^n$ is a cell complex $\mathcal{X}$ together with an realization map $|\cdot|\colon \mathcal{X} \to \mathbb{R}^n$ which obeys the following properties:

1. An m-cell $\xi \in \mathcal{X}_m$ has an realization $|\xi|$ which is a bounded convex region of some m-dimensional affine subspace of $\mathbb{E}^m \subset \mathbb{R}^n$;

2. The collection of realizations $\{|\xi|\}_{\xi \in \mathcal{X}}$ is disjoint in $\mathbb{R}^n$;

3. For any $\xi \in \mathcal{X}$ we have $|\operatorname{cl}(\xi)| = \operatorname{cl}|\xi|$.

The set $|\mathcal{X}| := \bigcup_{\xi \in \mathcal{X}} |\xi| \subset \mathbb{R}^n$ is called a polyhedron.

Condition (3) implies that the cellular closure relates to the topological closure, in particular $\xi' \leq \xi$ if and only if $\xi' = \xi$ or $|\xi'| \subseteq \operatorname{cl}|\xi| - |\xi|$. The polyhedral complex $\mathcal{X}$ is a discretization of the topological space $|\mathcal{X}| \subset \mathbb{R}^n$ into a cellular complex. In the sequel, we assume that any polyhedral complex $(\mathcal{X}, \leq, \kappa, \dim)$ has its incidence function $\kappa\colon \mathcal{X} \times \mathcal{X} \to \mathbb{K}$ determined as in [30, Chapter III (6.4)]. When the underlying field $\mathbb{K}$ is $\mathbb{Z}_2$ the incidence numbers may be described quite simply: $\kappa(\xi, \xi')$ is nonzero if and only if ξ covers ξ' in $(\mathcal{X}, \leq)$ (see Definition 2.4.2).

For the remainder of the section let $\mathcal{X}$ be a polyhedral complex in $\mathbb{R}^n$. The next few results relate the topology of $\mathcal{X}$ with that of $|\mathcal{X}|$.

Lemma 2.7.2. *Let $\xi \in \mathcal{X}$ and $\mathcal{K} \subset \mathcal{X}$. If $|\xi| \cap |\mathcal{K}| \neq \varnothing$ then $\xi \in \mathcal{K}$.*

Proof. $|\mathcal{K}| = \bigcup_{\xi' \in \mathcal{K}} |\xi'|$. It follows from (2) that if $|\xi| \cap \cup_{\xi' \in \mathcal{K}} |\xi| \neq \varnothing$, then $\xi \in \mathcal{K}$. □

Proposition 2.7.3. *If $\mathcal{U} \subset \mathcal{X}$ is an upper set then $|\mathcal{U}|$ is open in $|\mathcal{X}|$.*

Proof. If $\mathcal{U} \subset \mathcal{X}$ is an upper set then $\mathcal{U}^c$ is a down-set. Thus $\mathrm{cl}(\mathcal{U}^c) = \mathcal{U}^c$. (2) implies $|\mathcal{U}^c| = |\mathcal{U}|^c$. Thus

$$|\mathcal{U}|^c = |\mathcal{U}^c| = |\,\mathrm{cl}(\mathcal{U}^c)| = \mathrm{cl}\,|\mathcal{U}^c| = \mathrm{cl}(|\mathcal{U}|^c).$$

Therefore $|\mathcal{U}|^c$ is a closed set, implying $|\mathcal{U}|$ is an open set. $\square$

Proposition 2.7.4. *If* $\mathcal{K} \subset \mathcal{X}$ *then* $|\mathrm{int}\mathcal{K}| = \mathrm{int}|\mathcal{K}|$.

Proof. By definition of frontier, $\mathcal{K} = \mathrm{fr}\,\mathcal{K} \sqcup \mathrm{int}\mathcal{K}$ and (2) implies $|\mathcal{K}| = |\,\mathrm{fr}\,\mathcal{K}| \sqcup |\mathrm{int}\mathcal{K}|$. As $\mathrm{int}\mathcal{K}$ is an upper set, it follows from Proposition 2.7.3 that $|\mathrm{int}\mathcal{K}|$ is open in $|\mathcal{X}|$. Therefore $|\mathrm{int}\mathcal{K}| \subset \mathrm{int}|\mathcal{K}|$. Since $\mathrm{fr}\,\mathcal{K} \subseteq \mathcal{X} \setminus \mathcal{X}^+$ (the frontier is composed of cells that are not maximal) we have $\mathrm{int}|\,\mathrm{fr}\,\mathcal{K}| = \varnothing$. Therefore there are no open sets U such that $|\mathrm{int}\mathcal{K}| \subset U \subset \mathrm{int}|\mathcal{K}|$. As $\mathrm{int}|\mathcal{K}|$ is the largest open set contained in $|\mathcal{K}|$, it follows that $|\mathrm{int}\mathcal{K}| = \mathrm{int}|\mathcal{K}|$. $\square$

Corollary 2.7.5. *If* $\mathcal{K} \subset \mathcal{X}$ *then* $\mathrm{bd}\,|\mathcal{K}| = |\,\mathrm{fr}\,\mathcal{K}|$.

Proof. By definition $\mathcal{K} = \mathrm{fr}\,\mathcal{K} \sqcup \mathrm{int}\mathcal{K}$. From (2) and Proposition 2.7.4 we have that

$$|\mathcal{K}| = |\,\mathrm{fr}\,\mathcal{K}| \sqcup |\mathrm{int}\mathcal{K}| = |\,\mathrm{fr}\,\mathcal{K}| \sqcup \mathrm{int}|\mathcal{K}|.$$

Thus $|\,\mathrm{fr}\,\mathcal{K}| = |\mathcal{K}| \setminus \mathrm{int}|\mathcal{K}| = \mathrm{bd}\,|\mathcal{K}|$. $\square$

2.8 Discrete Morse Theory

We review the use of discrete Morse theory to compute homology of complexes. Our exposition is brief and follows [23]. See also [22, 38, 45].

Definition 2.8.1. A *partial matching* of cell complex $\mathcal{X}$ consists of a partition of the cells in $\mathcal{X}$ into three sets A, K, and Q along with a bijection $w\colon Q \to K$ such that for any $\xi \in Q$ we have that $\kappa(w(\xi), \xi) \neq 0$. A partial matching is called *acyclic* if the transitive closure of the binary relation $\ll$ on Q defined by

$$\xi' \ll \xi \text{ if and only if } \kappa(w(\xi), \xi') \neq 0$$

generates a partial order $\leq$ on Q.

Partial matchings are sometimes called *discrete vector fields*; acyclic partial matchings are sometimes called *gradient discrete vector fields*. We may lift the partial matching to a degree 1 map (see Definition 2.3.4) $V\colon C_\bullet(\mathcal{X}) \to C_{\bullet+1}(\mathcal{X})$ by defining it using the distinguished basis:

$$V(x) = \begin{cases} \kappa(\xi,\xi')w(x) & x \in Q \\ 0 & \text{otherwise} \end{cases} \tag{2.3}$$

We denote acyclic partial matchings by the tuple $(A, w\colon Q \to K)$. An acyclic partial matching $(A, w\colon Q \to K)$ of $\mathcal{X}$ can be used construct a new chain complex. This is done through the observation that acyclic partial matchings produce degree 1 maps $C_\bullet(\mathcal{X}) \to C_{\bullet+1}(\mathcal{X})$ called *splitting homotopies*. Splitting homotopies are reviewed in depth in Section 3.6. Further references to the use of splitting homotopies within discrete Morse theory can be found in [45]. The following proposition is from [23], however we make a sign change to agree with the exposition in Section 3.6.

Proposition 2.8.2 ([23], Proposition 3.9)**.** *An acyclic partial matching (A, w) induces a unique linear map $\gamma\colon C_\bullet(\mathcal{X}) \to C_{\bullet+1}(\mathcal{X})$ so that $\operatorname{im}(id_{\mathcal{X}} - \partial\gamma) \subseteq C_\bullet(A) \oplus C_\bullet(K)$, $\operatorname{im}\gamma = C_\bullet(K)$ and $\ker\gamma = C_\bullet(A) \oplus C_\bullet(K)$. It is given by the formula*

$$\gamma = \sum_{i\geq 0} V(\mathrm{id}_{C(\mathcal{X})} - \partial V)^i. \tag{2.4}$$

Let $\iota_{\mathcal{A}}\colon C_\bullet(A) \to C_\bullet(\mathcal{X})$ and $\pi_{\mathcal{A}}\colon C_\bullet(\mathcal{X}) \to C_\bullet(\mathcal{A})$ be the canonical inclusion and projection. Define $\psi\colon C_\bullet(\mathcal{X}) \to C_\bullet(A), \phi\colon C_\bullet(A) \to C_\bullet(\mathcal{X})$ and $\partial^A\colon C_\bullet(A) \to C_{\bullet-1}(A)$ by

$$\psi := \pi_A \circ (\mathrm{id}_{\mathcal{X}} - \partial\gamma) \qquad \phi := (\mathrm{id}_{\mathcal{X}} - \gamma\partial) \circ \iota_A \qquad \partial^A := \psi \circ \partial \circ \phi \tag{2.5}$$

Theorem 2.8.3 ([23], Theorem 3.10)**.** *$(C_\bullet(A), \partial^A)$ is a chain complex and ψ, ϕ are chain equivalences. In particular,*

$$\psi \circ \phi = \mathrm{id}_{C(A)} \qquad \phi \circ \psi - \mathrm{id}_{C(\mathcal{X})} = \partial\gamma + \gamma\partial$$

As a corollary $H_\bullet(C_\bullet(A)) \cong H_\bullet(C_\bullet(\mathcal{X}))$. Regarding computations, acyclic partial matchings are relatively easy to produce, see [Algorithm 3.6 (Coreduction-based

Matching)][23], which is recalled in Section 3.7. Moreover, given an acyclic partial matching there is an efficient algorithm to produce the associated splitting homotopy [23, Algorithm 3.12 (Gamma Algorithm)], also recalled in Section 3.7.

2.9 Lattice Structures in Dynamics

Traditionally, Conley theory uses the notions of attractors and repellers, but little in the way of order theory [8]. The use of order theory became more explicit in Franzosa's papers on connection matrix theory [17, 15, 16] and especially in the work of Robbin-Salamon [42]. The more modern treatment of dynamics in [26, 27, 28], relies very heavily on order theory, and our exposition of dynamics will follow this set of papers. In the sequel, X is a compact metric space.

Definition 2.9.1. A *flow* on X is a continuous map $\varphi\colon \mathbb{R}\times X\to X$ that satisfies:

1. $\varphi(0,x) = x$ for all $x\in X$, and
2. $\varphi(t,\varphi(s,x)) = \varphi(t+s,x)$ for all $s,t\in\mathbb{R}$ and $x\in X$.

Definition 2.9.2. A subset $N\subset X$ is a *regular closed* set if $N = \mathrm{cl}(\mathrm{int}N)$. The collection of regular closed sets on X forms a Boolean algebra, denoted $\mathscr{R}(X)$, with operations $\vee=\cup$ and $\wedge = \mathrm{cl}(\mathrm{int}(\cdot\cap\cdot))$.

Definition 2.9.3. Let $U\subset X$. The *omega-limit set of* U is defined as

$$\omega(U) := \bigcap_{t\geq 0} \mathrm{cl}\left(\varphi([t,\infty),U\right).$$

The *alpha-limit set of* U is defined as

$$\alpha(U) := \bigcap_{t\leq 0} \mathrm{cl}\left(\varphi(-\infty,t],U\right).$$

Definition 2.9.4. A regular closed set N is an *attracting block* if $\varphi(t,N)\subset \mathrm{int}N$ for all $t>0$. The *lattice of regular closed attracting blocks*, denoted by $\mathsf{ABlock}_{\mathscr{R}}(\varphi)$, is the collection of regular closed attracting blocks with $\vee=\cup$ and $\wedge = \mathrm{cl}(\mathrm{int}(\cdot\cap\cdot))$.

Definition 2.9.5. A set $A\subset X$ is an *attractor* if there exists an attracting block U such that $a=\omega(U)$. The *lattice of attractors*, denoted $\mathsf{Att}(\varphi)$, is the collection of attractors with operations $\vee := \cup$ and $\wedge := \omega(\cdot\cap\cdot)$.

The map $\omega(\cdot)$ induces a lattice morphism $\omega\colon \mathsf{ABlock}_{\mathscr{R}}(\varphi) \to \mathsf{Att}(\varphi)$.

2.9.1 Computational Dynamics

We review some central concepts to the computational Conley theory. See [2, 25, 26] for more material and further references. We begin with discretization of topological space using grids and the algebra of regular, closed sets [25, 26].

Grids are a widely used for discretizing the phase space of a dynamical systems, see [2, 25, 26]. We use the elegant order-theoretic definition of grid given in [26].

Definition 2.9.6. A *grid* on X is the set of atoms of a finite subalgebra $\mathscr{R}(X)$.

Example 2.9.7. If $\mathcal{X}$ is a polyhedral complex in $\mathbb{R}^n$ then $|\mathcal{X}|$ is a compact, convex subset. The topological closures of the top-cells $\{\mathrm{cl}\,|\xi| \mid \xi \in \mathcal{X}^+\}$ is a grid on $|\mathcal{X}|$.

In computational Conley theory, continuous dynamics are approximated via finite binary relations. In our case, the dynamics are approximated with a binary relation $\mathcal{F} \subset \mathcal{X}^+ \times \mathcal{X}^+$ defined on the set of top-cells of the cell complex $\mathcal{X}$. In the context of approximating flows the pair $(\mathcal{X}^+, \mathcal{F})$ is called the *discrete flow*, cf. [25]. See also Section 4.3.

There are three (equivalent) perspectives for the finite binary relation $\mathcal{F}$.

- As a binary relation in itself. This is useful for developing the appropriate category theory and extensions of Birkhoff's Theorem. See [25, Appendix A].
- As a directed graph with vertices $\mathcal{X}^+$ and edge set $\{\xi \to \xi' \mid (\xi, \xi') \in \mathcal{F}\}$. This perspective is useful for developing algorithms [6].
- As a multi-valued map. This is traditionally written as $\mathcal{F}\colon \mathcal{X}^+ \rightrightarrows \mathcal{X}^+$. In this case the notation $\mathcal{F}(\xi) = \{\xi' \mid (\xi, \xi') \in \mathcal{F}\}$. This is the perspective that $\mathcal{F}$ is a combinatorial approximation to continuous dynamics.

Definition 2.9.8. A set $\mathcal{U} \subset \mathcal{X}$ is *forward invariant* if $\mathcal{F}(\mathcal{U}) \subset \mathcal{U}$. The collection of forward invariant sets, denoted by $\mathsf{Invset}^+(\mathcal{F})$, is a finite distributive lattice with $\wedge := \cap$ and $\vee := \cup$.

Definition 2.9.9. The ω-*limit set* of a set $\mathcal{U} \subset \mathcal{X}$ is defined as

$$\omega(\mathcal{U}, \mathcal{F}) := \bigcap_{k \geq 0} \Gamma_k^+(\mathcal{U})$$

where $\Gamma_k^+(\mathcal{U}) = \cup_{n \geq k} \mathcal{F}^n(\mathcal{U})$ for $k > 0$ is the k-*forward image* of $\mathcal{U}$. When $\mathcal{F}$ is clear from context, we write $\omega(\cdot) = \omega(\cdot, \mathcal{F})$.

Definition 2.9.10. A set $\mathcal{A} \subset \mathcal{X}$ is an *attractor* for $\mathcal{F}$ if $\mathcal{F}(\mathcal{A}) = \mathcal{A}$. The collection of attractors, denoted $\mathsf{Att}(\mathcal{F})$, is a finite distributive lattice with $\wedge$ and $\vee$ given via

$$\mathcal{A} \vee \mathcal{A}' = \mathcal{A} \cup \mathcal{A}' \quad \text{and} \quad \mathcal{A} \wedge \mathcal{A}' = \omega(\mathcal{A} \cap \mathcal{A}')$$

The map $\omega \colon \mathsf{Invset}^+(\mathcal{F}) \twoheadrightarrow \mathsf{Att}(\mathcal{F})$ is a lattice epimorphism. As a relation, $\mathcal{F}$ has a poset of strongly connected components, denoted $\mathsf{SC}(\mathcal{F})$. Moreover we have $\mathsf{Invset}^+(\mathcal{F}) = \mathsf{O}(\mathsf{SC}(\mathcal{F}))$.

Chapter 3

Computational Connection Matrix Theory

We begin with expounding upon the relationship between our formulation and classical connection matrix theory. Following this comes an overview of the content of this chapter.

3.1 Historical Remarks and Context

Historically, connection matrix theory was first developed by R. Franzosa [15, 16, 17]. Franzosa's treatment uses a *chain complex braid* indexed over a poset P. The chain complex braid can be understood as a data structure that stores the singular chain data associated to a lattice A of attracting blocks. In this case, the poset P arises as the poset of join-irreducibles $\mathsf{J}(\mathsf{A})$. These objects constitute the category $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$ and are reviewed in Section 3.5. Graded module braids are data structures for storing the homological information contained in a chain complex braid. Graded module braids form a category $\mathbf{GMB}(\mathsf{P}, \mathbf{Vect}^{\mathbb{Z}})$ and there is a functor $\mathfrak{H}\colon \mathbf{ChB}(\mathsf{P}, \mathbf{Vect}) \to \mathbf{GMB}(\mathsf{P}, \mathbf{Vect}^{\mathbb{Z}})$ which is analogous to a homology functor. Connection matrix theory for continuous self-maps, as developed by D. Richeson, also employs the structures of chain complex braids and graded module braids [41].

The contribution of J. Robbin and D. Salamon to connection matrix theory both addresses maps and merges the theory with order theoretic principles [42]. They introduced the idea of a chain complex being either graded by a poset P or filtered by a lattice L. These respectively constitute the categories $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ and $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ and are described in Sections 3.3 and 3.4.

One goal of this thesis is to address how our approach, Franzosa's approach and

Robbin and Salamon's approach fit together. First, we wish to emphasize that in applications data come in the form of a P-graded cell complex, the collection of which we call $\text{Cell}(\mathsf{P})$. Let L be the lattice $\mathsf{O}(\mathsf{P})$ of downsets of P. A P-graded cell complex determines three distinct objects: a P-graded chain complex, an L-filtered chain complex, and a chain complex braid over P. This may be put into the following diagram.

$$
\begin{tikzcd}
& \text{Cell}(\mathsf{P}) \arrow[ddl, dashed, "\mathcal{L}"'] \arrow[d, dashed, "\mathcal{C}"] \arrow[ddr, dashed, "\mathcal{B}"] & \\
& \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \arrow[dl, "\mathfrak{L}"'] \arrow[d] \arrow[dr, "\mathfrak{B}"] & \\
\mathbf{Ch}(\mathbf{FVect}(\mathsf{L})) \arrow[d] & \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \arrow[dl, "\mathfrak{L}_{\mathbf{K}}"'] \arrow[dr, "\mathfrak{B}_{\mathbf{K}}"] & \mathbf{ChB}(\mathsf{P}, \mathbf{Vect}) \arrow[d] \\
\mathbf{K}(\mathbf{FVect}(\mathsf{L})) & & \mathbf{KB}(\mathsf{P}, \mathbf{Vect})
\end{tikzcd}
\tag{3.1}
$$

The dashed arrows are assignments while the solid arrows are functors. These are described in Sections 3.3–3.5. Franzosa's theory comprises the right-hand side of (3.1), while Robbin and Salamon's theory comprises the left hand side. One of our contributions to connection matrix theory is to phrase it in a homotopy-theoretic language. In particular, we introduce the appropriate homotopy categories $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$, $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ and $\mathbf{KB}(\mathsf{P}, \mathbf{Vect})$ on the bottom of (3.1). Moreover, in the context of applications we show that both of the notions of connection matrices for both $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ and $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$ (that is, the formulations of both Robbin-Salamon and Franzosa) may be computed by utilizing graded algebraic-discrete Morse theory within the category $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$.

3.2 Overview

In this chapter we develop a new formulation of connection matrix theory. This may be described roughly as follows. In Conley theory one has either some homological or chain data of interest. In Franzosa's setting the datum is a graded module braid. For Robbin and Salamon, the datum is an L-filtered chain complex. In our setting of computational dynamics and applied topology, the datum is either an L-filtered cell complex or a P-graded cell complex. The guiding principle of a connection matrix is then:

> *A connection matrix is the boundary operator of a poset-graded chain complex which recovers (up to chain equivalence) the data of interest.*

Some remarks are in order.

- First, the connection matrix is a boundary operator. Boundary operators do not exist independently of the underlying chain complex. We call this chain complex a *Conley complex*.

- Second, we are taking a cue from Franzosa here by regarding the connection matrix as residing in the graded category. Consulting (3.1), the graded category is upstream of both the filtered and braided categories. Proving the existence of a connection matrix amounts to proving that an appropriate functor is essentially surjective. For instance, Franzosa's existence proof is akin to (but not precisely) the fact that the functor

$$\mathfrak{B}_{\mathbf{K}}\colon \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{KB}(\mathsf{P}, \mathbf{Vect})$$

  is essentially surjective.

- Third, recovering up to chain equivalence is equivalent to recovering the isomorphism class in the homotopy category. Moreover, it implies that all homological invariants (graded module braids, persistent homology, etc) are recovered up to isomorphism.

Many of the ideas developed in this section may be found either in Franzosa [17] or Robbin-Salamon [42]. Our primary contribution in this section is our use of chain equivalence and homotopy categories to set up connection matrix theory. This has the following payoffs:

- We settle the non-uniqueness problems of the connection matrix. Classically, this is still an open problem. We show that in our formulation, a Conley complex is unique up to isomorphism. Thus connection matrices are unique up to a similarity transformation. See Proposition 3.3.29 and Remark 3.3.31.

- We distill the construction of the connection matrix to a particular functor, which we call a *Conley functor*. See Section 3.8.
- We readily relate persistent homology and connection matrix theory. This implies applications for connection matrix theory in topological data analysis.

3.3 Graded Complexes

In Sections 3.3 and 3.4 we introduce objects which result from a marriage of homological algebra and order theory. We introduce our notion of connection matrix, which is part of what we call a Conley complex. In particular, we employ categorical language and explicitly develop an appropriate homotopy category for connection matrix theory over fields. Along the way we provide motivation through a selection of examples, many of which build upon Example 1.2.1 in the introduction.

Important results from this section for connection matrix theory are Proposition 3.3.29, which shows that the Conley complex is an invariant of the chain equivalence class, implying the non-uniqueness of the connection matrix is captured in terms of a change of basis, cf. [18]; Theorem 3.4.21 establishes that computing a connection matrix in the sense of [42] can be done at the level of the P-graded chain complex. The relationship between posets and lattices encapsulated by Birkhoff's theorem (Section 2.4.3) is also reflected in the homological algebra. Namely, we establish a categorical equivalence – Theorem 3.4.10 – between the category of L-filtered chain complexes and P-graded chain complexes where $\mathsf{L} = \mathsf{O}(\mathsf{P})$.

For the remainder of this section, let $\mathbb{K}$ be a field and let P be a finite poset. Recall that **Vect** is the category of vector spaces over $\mathbb{K}$.

3.3.1 Graded Vector Spaces

Definition 3.3.1. A P-*graded vector space* $V = (V, \pi)$ is a vector space V equipped with a P-indexed family of idempotents (projections) $\pi = \{\pi^p \colon V \to V\}_{p\in\mathsf{P}}$ such that $\sum_{p\in\mathsf{P}} \pi^p = \mathrm{id}_V$ and if $p \neq q$ then $\pi^p \circ \pi^q = 0$. We call π a P-*grading* of V. Suppose (V, π_V) and (W, π_W) are P-graded vector spaces. A map $\phi \colon V \to W$ is P-*filtered* if for all $p, q \in \mathsf{P}$

$$\phi^{pq} := \pi_W^p \circ \phi \circ \pi_V^q \neq 0 \implies p \leq q. \tag{3.2}$$

Remark 3.3.2. Given a map $\phi\colon V \to W$, the terminology P-filtered is apt as it is readily verified that ϕ obeys Eqn. (3.2) if and only if for all $q \in \mathsf{P}$

$$\phi(V^q) \subseteq \bigoplus_{p \leq q} V^p.$$

This is in turn equivalent to

$$\phi(\bigoplus_{p \leq q} V^q) \subseteq \bigoplus_{p \leq q} V^q.$$

See Proposition 3.4.2, cf. Definition 3.4.1.

The next few results establish that working with P-graded vector spaces and P-filtered linear maps follows the rules of working with upper triangular matrices. The proofs are elementary linear algebra and matrix theory.

Proposition 3.3.3. *A P-graded vector space (V, π) admits a decomposition $V = \bigoplus_{p \in \mathsf{P}} V^p$ where $V^p = \operatorname{im} \pi^p$. Conversely, if V is a vector space and $V = \bigoplus_{p \in \mathsf{P}} V^p$ then the collection $\pi = \{\pi^p\}$ with $\pi^p(\sum_{q \in \mathsf{P}} v^q) := v^p$ where $v^q \in V^q$ is a P-grading of V.*

Proposition 3.3.4. *If $\phi\colon (U, \pi_U) \to (V, \pi_V)$ and $\psi\colon (V, \pi_V) \to (W, \pi_W)$ are P-filtered linear maps, then the composition $\psi \circ \phi$ is P-filtered and*

$$(\psi \circ \phi)^{pq} = \sum_{p \leq r \leq q} \psi^{pr} \phi^{rq} \qquad\qquad \psi \circ \phi = \sum_{p \leq q} (\psi\phi)^{pq}$$

Proof. From linearity and the P-grading properties:

$$\begin{aligned}
(\psi \circ \phi)^{pq} = \pi^p \psi \phi \pi^q = \pi^p \psi \big(\sum_{r \in \mathsf{P}} \pi^r \big) \phi \pi^q &= \sum_{r \in \mathsf{P}} \pi^p \psi \pi^r \phi \pi^q \\
= \sum_{r \in \mathsf{P}} \pi^p \psi \pi^r \pi^r \phi \pi^q = \sum_{r \in \mathsf{P}} \psi^{pr} \phi^{rq} &= \sum_{p \leq r \leq q} \psi^{pr} \phi^{rq}
\end{aligned}$$

The second equality follows from the first:

$$\begin{aligned}
\psi \circ \phi = \sum_{p \leq q'} \psi^{pq'} \sum_{p' \leq q} \phi^{p'q} = \sum_{\substack{p \leq q' \\ p' \leq q}} \psi^{pq'} \phi^{p'q} &= \sum_{p \leq q' = p' \leq q} \psi^{pq'} \phi^{p'q} \\
= \sum_{\substack{p \leq q \\ p \leq r \leq q}} \psi^{pr} \phi^{rq} = \sum_{p \leq q} \sum_{p \leq r \leq q} \psi^{pr} \phi^{rq} &= \sum_{p \leq q} (\psi \circ \phi)^{pq}
\end{aligned}$$

□

Definition 3.3.5. The *category of* P*-graded vector spaces*, denoted $\mathbf{GrVect}(\mathsf{P})$, is the category whose objects are P-graded vector spaces and whose morphisms are P-filtered linear maps. We denote by $u\colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{Vect}$ the forgetful functor taking (V, π) to V which forgets the grading.

For a P-graded vector space (V, π) any projection $\pi^p\colon V \to V$ factors as

$$V \xrightarrow{e^p} V^p \xrightarrow{\iota^p} V,$$

where $e^p\colon V \to V^p$ is the epimorphism to $\operatorname{im} \pi^p = V^p$ and $\iota^p\colon V^p \hookrightarrow V$ is the natural inclusion. We have the identities

$$\pi^p = \iota^p \circ e^p, \qquad \pi^p \circ \iota^p = \iota^p, \qquad e^p \circ \pi^p = e^p.$$

Given a linear map $\phi\colon (V, \pi_V) \to (W, \pi_W)$ we define

$$\Phi^{pq} := e^p_W \circ \phi \circ \iota^q_V \colon V^q \to W^p$$

Using the upper-case Φ^{pq} as above is our convention to refer to the restriction of ϕ to the (p, q)-matrix entry. The P-gradings imply that

$$V = \bigoplus_{p \in \mathsf{P}} V^p \qquad W = \bigoplus_{p \in \mathsf{P}} W^p$$

The linear map ϕ is equivalent to the matrix of linear maps $\{\Phi^{pq}\}_{p,q \in \mathsf{P}}$, via

$$\sum_{p,q} \Phi^{pq} e^q x = \sum_{p,q} \phi^{pq} x$$

It is straightforward that ϕ is P-filtered if and only if

$$\Phi^{pq} \neq 0 \implies p \leq q \tag{3.3}$$

Remark 3.3.6. In [17], linear maps which obey (3.3) are referred to as upper triangular with respect to P.

Given a P-graded vector space (V, π) and a subset $I \subset \mathsf{P}$ we define

$$\pi^I := \sum_{p \in I} \pi^p \qquad V^I := \operatorname{im} \pi^I = \bigoplus_{p \in I} V^p$$

The space V^I is a subspace of the underlying vector space $V = u(V,\pi)$. For a P-filtered linear map $\phi\colon (V,\pi_V) \to (W,\pi_W)$, we define $\phi^I\colon V \to V$ and $\Phi^I\colon V^I \to V^I$ via

$$\phi^I := \pi^I \circ \phi \circ \pi^I \qquad \text{and} \qquad \Phi^I := e^I_W \circ \phi \circ \iota^I_V\colon V^I \to W^I. \tag{3.4}$$

Proposition 3.3.7. *If $\phi\colon (U,\pi_U) \to (V,\pi_V)$ and $\psi\colon (V,\pi_V) \to (W,\pi_W)$ are P-filtered maps and $I \subset \mathsf{P}$ then*

$$\phi^I = \sum_{\substack{p\le q\\ p,q\in I}} \phi^{pq}.$$

Moreover, if $I \subset \mathsf{P}$ is convex then $\psi^I \circ \phi^I = (\psi \circ \phi)^I$.

Proof. The first identity is immediate. The second follows from convexity of I:

$$\begin{aligned}(\psi\circ\phi)^I = \pi^I(\psi\circ\phi)\pi^I = \sum_{\substack{p\le q\\ p,q\in I}} (\psi\circ\phi)^{pq} &= \sum_{\substack{p\le q\\ p,q\in I}} \sum_{p\le r\le q} \psi^{pr}\phi^{rq}\\ &= \sum_{\substack{p\le q\\ p,q\in I}} \psi^{pq} \circ \sum_{\substack{p\le q\\ p,q\in I}} \phi^{pq} = \psi^I \circ \phi^I.\end{aligned}$$ $\square$

The above result enables the definition of the following family of forgetful functors, parameterized by the convex sets of P.

Definition 3.3.8. Let $I \subset P$ be a convex set. The *forgetful functor* $u^I\colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{GrVect}(I)$ is defined via

$$u^I\big((V,\pi)\big) := (V^I, \{\pi^p\}_{p\in I}).$$

For $\phi\colon (V,\pi_V) \to (W,\pi_W)$, we define

$$u^I(\phi) := \Phi^I = e^I_W \circ \phi \circ \iota^I_V\colon V^I \to W^I.$$

We write a $\mathbb{Z}$-indexed family of P-graded vector spaces as $(V_\bullet,\pi_\bullet) = \{(V_n,\pi_n\})_{n\in\mathbb{Z}}$. For a fixed $p \in \mathsf{P}$ there is a family of vector spaces $V^p_\bullet = \{V^p_n\}_{n\in\mathbb{Z}}$.

3.3.2 Graded Chain Complexes

The category $\mathbf{GrVect}(\mathsf{P})$ is additive but not abelian. Following Section 2.3 we may form the category $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ of chain complexes in $\mathbf{GrVect}(\mathsf{P})$. An object C

of $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ is a chain complex of P-graded vector spaces. For short, we say that this is a P-*graded chain complex*. The data of C can be unpacked as the triple $C = (C_\bullet, \partial_\bullet, \pi_\bullet)$ where:

1. $(C_\bullet, \partial_\bullet)$ is a chain complex,
2. (C_n, π_n) is a P-graded vector space for all n, and
3. $\partial_n \colon (C_n, \pi_n) \to (C_{n-1}, \pi_{n-1})$ is a P-filtered linear map for each n.

Typically we denote C by (C, π) to distinguish it as carrying a grading. A morphism $\phi\colon (C,\pi) \to (C',\pi)$ is a chain map $\phi\colon (C,\partial) \to (C',\partial')$, such that $\phi_n\colon (C_n,\pi_n) \to (C'_n,\pi_n)$ is a P-filtered linear map for each n. We call the morphisms of $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ the P-*filtered chain maps*.

Proceeding with our convention, we define $\Delta_j^{pq} := e_j^p \circ \partial_{j-1} \circ \iota_j^q \colon C_j^q \to C_j^p$. Since (C,π) is P-graded we have

$$C_j = \bigoplus_{q\in\mathsf{P}} C_j^q \qquad C_{j-1} = \bigoplus_{p\in\mathsf{P}} C_{j-1}^p$$

The boundary operator $\partial_j \colon C_j \to C_{j-1}$ is equivalent to the matrix of maps $\{\Delta_j^{pq}\}_{p,q\in\mathsf{P}}$. The P-filtered condition of Eqn. (3.2) is equivalent to the condition that

$$\Delta_j^{pq} \neq 0 \implies p \leq q. \tag{3.5}$$

Remark 3.3.9. Viewing the boundary operator ∂ as a matrix of maps $\{\Delta^{pq}\}_{p,q\in\mathsf{P}}$ is the origin of the term 'connection matrix'.

It follows from Proposition 2.3.9 that the forgetful functor $u\colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{Vect}$ induces a forgetful functor $u_{\mathbf{Ch}}\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}(\mathbf{GrVect}(I))$, where

$$u_{\mathbf{Ch}}(C,\pi) = \big(\{u(C_n)\}_{n\in\mathbb{Z}}, \{u(\partial_n : C_n \to C_{n-1})\}_{n\in\mathbb{Z}}\big) = \big(\{C_n^I\}_{n\in\mathbb{Z}}, \{\Delta_n^I\}_{n\in\mathbb{Z}}\big).$$

Similarly, for a convex set $I \subset \mathsf{P}$ the forgetful functor $u^I\colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{GrVect}(I)$ induces a functor $u^I_{\mathbf{Ch}}\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}(\mathbf{GrVect}(I))$. When the context is clear we will write u for $u_{\mathbf{Ch}}$ and u^I for $u^I_{\mathbf{Ch}}$.

Given a convex set $I \subset \mathsf{P}$ and the forgetful functor $u\colon \mathbf{Ch}(\mathbf{GrVect}(I)) \to \mathbf{Ch}(\mathbf{Vect})$ it is often useful to consider the the composition $u \circ u^I\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}(\mathbf{Vect})$. Unpacking Definition 3.3.8 shows that $u \circ u^I$ may be written simply as

$$u \circ u^I(C, \pi) = (C^I, \Delta^I), \qquad u \circ u^I(\phi) = \Phi^I\colon C^I \to C^I.$$

Remark 3.3.10. A P-graded complex (C, π) engenders a collection of chain complexes in $\mathbf{Vect}$, $\{u \circ u^I(C, \pi)\}_{I \in I(\mathsf{P})}$, indexed by the convex sets of P. This collection is used for the functor $\mathfrak{B}$ which builds a chain complex braid out of a graded chain complex. See Section 3.5.

Proposition 3.3.11. *Let (C, π) be a P-graded chain complex. If $a \in \mathsf{O}(\mathsf{P})$, i.e., a is a down-set of P, then (C^a, Δ^a) is a subcomplex of C.*

Proof. If $a \in \mathsf{O}(\mathsf{P})$ the fact that ∂ is P-graded implies that $\partial(\bigoplus_{p \in a} C^p) \subseteq \bigoplus_{p \in a} C^p$. Moreover $\Delta^a = e^p \circ \partial \circ \iota^p = \partial|_{C^a}$. Therefore (C^a, Δ^a) is a subcomplex of C. □

3.3.3 The Subcategory of Strict Objects

Definition 3.3.12. A P-graded chain complex (C, π) is said to be *strict* if for each $j \in \mathbb{Z}$ and $p \in \mathsf{P}$

$$\partial_j^{pp} = 0. \tag{3.6}$$

The strict objects form a subcategory $\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \subset \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$, called the *subcategory of strict objects*.

Remark 3.3.13. In [17], a boundary operator ∂_j which obeys condition (3.6) is called strictly upper triangular with respect to P.

Proposition 3.3.14. *If (C, π) be a P-graded chain complex then (C, π) is strict if and only if for each $j \in \mathbb{Z}$*

$$\partial_j = \sum_{p<q} \partial_j^{pq}. \tag{3.7}$$

Corollary 3.3.15. *If (C, π) is strict, then $u^p(C, \pi) = (C_\bullet^p, \Delta_\bullet^{pp})$ is a minimal chain complex for any $p \in \mathsf{P}$. Moreover, for any $j \in \mathbb{Z}$*

$$C_j = \bigoplus_{p \in \mathsf{P}} H_j(C_\bullet^p, \Delta_\bullet^{pp}).$$

Proof. If (C, π) is strict then the boundary operators $\Delta_j^{pp} = 0$ for all $j \in \mathbb{Z}$ by definition. Therefore $H_j(C_\bullet^p, \Delta^{pp}) = C_j$. Finally, since (C, π) is P-graded we have that

$$C_j = \bigoplus_{p \in \mathsf{P}} C_j^p = \bigoplus_{p \in \mathsf{P}} H_j(C_\bullet^p, \Delta_\bullet^{pp}). \qquad \square$$

Corollary 3.3.15 implies that $\partial_j \colon C_j \to C_{j-1}$ may be regarded as a P-filtered map on homology:

$$\partial_j \colon \bigoplus_{p \in \mathsf{P}} H_j(C_\bullet^p, \Delta_\bullet^{pp}) \to \bigoplus_{p \in \mathsf{P}} H_{j-1}(C_\bullet^p, \Delta_\bullet^{pp}). \tag{3.8}$$

In the context of Conley theory, Eqn. (3.8) implies that ∂_j is a boundary operator on Conley indices.

Remark 3.3.16. The significance of Eqn. (3.8) is that in this form the nonzero entries in the boundary operator relate to connecting orbits.

Example 3.3.17. Let X be a closed manifold and $\varphi \colon \mathbb{R} \times X \to X$ a Morse-Smale gradient flow. The set P of fixed points are partially ordered by the flow and there is a poset morphism $\mu \colon \mathsf{P} \to \mathbb{N}$ which assigns each p its Morse index, i.e., the dimensionality of its unstable manifold. The associated Morse-Witten complex may be written

$$C_\bullet(X, \varphi) = \bigoplus_{p \in \mathsf{P}} C_\bullet^p$$

where $C_\bullet^p$ is the minimal chain complex in which the only nonzero chain group is in dimension $\mu(p)$, and $C_{\mu(p)}^p = \mathbb{K}$. The boundary map Δ is defined using trajectories [13, 42]. It is thus P-filtered. In particular, when $\mathbb{K} = \mathbb{Z}_2$ the entry Δ_{qp} counts the number of flow lines from q to p modulo two. It is a classical result that the homology $H_\bullet(C(X, \varphi))$ is isomorphic to the singular homology of X.

3.3.4 Graded Cell Complexes

In applications, data often come in the form of a cell complex $\mathcal{X} = (\mathcal{X}, \leq, \kappa, \dim)$ graded by a partial order P. This is codified in terms of an order preserving map $\nu \colon (\mathcal{X}, \leq) \to \mathsf{P}$. See Chapter 4 for an example of how these structures arise in the context of computational dynamics.

Definition 3.3.18. A P-*graded cell complex* is a cell complex $\mathcal{X} = (\mathcal{X}, \leq, \kappa, \dim)$ together with P and a poset morphism $\nu\colon (\mathcal{X}, \leq) \to \mathsf{P}$. The map ν is called the *grading*. We denote by $\mathrm{Cell}(\mathsf{P})$ the collection of P-graded cell complexes. For a P-graded cell complex $(\mathcal{X}, \nu)$, the underlying set $\mathcal{X}$ can be decomposed as

$$\mathcal{X} = \bigsqcup_{p \in \mathsf{P}} \mathcal{X}^p, \qquad \text{where } \mathcal{X}^p := \nu^{-1}(p)\,.$$

For each p, the fiber $\mathcal{X}^p$ together with the restriction of $(\leq, \kappa, \dim)$ to $\mathcal{X}^p$ is a subcomplex of $\mathcal{X}$. A P-graded cell complex $(\mathcal{X}, \nu)$ determines an *associated* P-*graded chain complex* $(C_\bullet(\mathcal{X}), \pi^\nu)$ (see Section 3.3.2) where for any $j \in \mathbb{Z}$

$$C_j(\mathcal{X}) = \bigoplus_{p \in \mathsf{P}} C_j(\mathcal{X}^p)\,.$$

The projection maps $\pi^\nu = \{\pi_j^p\}$ project to the fibers of ν, i.e.,

$$\pi_j^p\colon C_j(\mathcal{X}) \to C_j(\mathcal{X}^p).$$

The boundary operator

$$\partial_j\colon C_j(\mathcal{X}) \to C_{j-1}(\mathcal{X})$$

is P-filtered since ν is order-preserving; $\kappa(\xi, \xi') \neq 0$ implies that $\xi' \leq \xi$ which in turn implies $\nu(\xi') \leq \nu(\xi)$. The boundary operator ∂_j can be written as an upper triangular a matrix of maps $\{\Delta_j^{pq}\}$ where $\Delta_j^{pq}\colon C_j(\mathcal{X}^q) \to C_{j-1}(\mathcal{X}^p)$. We denote by $\mathcal{C}\colon \mathrm{Cell}(\mathsf{P}) \to \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ the assignment $(\mathcal{X}, \nu) \mapsto (C(\mathcal{X}), \pi^\nu)$.

Akin to graded chain complexes, there is a notion of being strict.

Definition 3.3.19. A P-graded cell complex $(\mathcal{X}, \nu)$ is *strict* if, for each $p \in \mathsf{P}$, the fiber $\nu^{-1}(p)$ is a minimal cell complex (see Definition 2.5.4).

Strict P-graded cell complexes engender strict P-graded chain complexes.

Proposition 3.3.20. *If $(\mathcal{X}, \nu)$ is a strict P-graded cell complex, then the associated P-graded chain complex $(C(\mathcal{X}), \pi^\nu)$ is strict.*

Example 3.3.21. Consider $(\mathcal{X}, \nu)$ and $(\mathcal{X}', \nu')$ and $\mathsf{P} = \{p, q, r\}$ of Example 1.2.1. It is a routine verification that $(\mathcal{X}, \nu)$ and $(\mathcal{X}', \nu')$ are both P-graded complexes. In particular, $(\mathcal{X}', \nu')$ is a strict P-graded complex. The underlying set $\mathcal{X}$ decomposes as $\mathcal{X} = \mathcal{X}^p \sqcup \mathcal{X}^r \sqcup \mathcal{X}^q$. More explicitly,

$$\mathcal{X} = \{v_0, e_0, v_1, e_1, v_2\}, \qquad \mathcal{X}^p = \{v_0\}, \qquad \mathcal{X}^r = \{v_2\}, \qquad \mathcal{X}^q = \{e_0, v_1, e_1\}.$$

This decomposition is reflected in the algebra as every chain group splits C_j as

$$C_j(\mathcal{X}) = C_j(\mathcal{X}^p) \oplus C_j(\mathcal{X}^r) \oplus C_j(\mathcal{X}^q).$$

As in Definition 3.3.18 the boundary operator ∂_j can be written as the P-filtered linear map (see Section 3.3.1)

$$\partial_j = \begin{array}{c} \\ C_{j-1}(\mathcal{X}^p) \\ C_{j-1}(\mathcal{X}^r) \\ C_{j-1}(\mathcal{X}^q) \end{array} \begin{array}{c} \begin{array}{ccc} C_j(\mathcal{X}^p) & C_j(\mathcal{X}^r) & C_j(\mathcal{X}^q) \end{array} \\ \begin{pmatrix} \Delta_j^{pp} & 0 & \Delta_j^{pq} \\ 0 & \Delta_j^{rr} & \Delta_j^{rq} \\ 0 & 0 & \Delta_j^{qq} \end{pmatrix} \end{array}$$

In particular, Δ_1 (the only nonzero differential) can be written as

$$\partial_1 = \begin{array}{c} \\ v_0 \\ v_1 \\ v_2 \end{array} \begin{array}{c} \begin{array}{cc} e_0 & e_1 \end{array} \\ \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \end{array} = \partial_1^{pq} + \partial_1^{rq} + \partial_1^{qq} = \begin{pmatrix} \Delta_1^{pq} \\ \Delta_1^{rq} \\ \Delta_1^{qq} \end{pmatrix}$$

Example 3.3.22. In applications the input is often a cell complex $\mathcal{X}$ and a function $\bar{\nu}\colon \mathcal{X}^+ \to \mathbb{R}$ on top cells. For instance, imaging data is often a two dimensional cubical complex with greyscale values on pixels (2-cells). Let $(\mathsf{Q}, \leq)$ be the totally ordered set where $\mathsf{Q} := \bar{\nu}(\mathcal{X}^+)$ and $\leq$ is the total order inherited from $\mathbb{R}$. We may extend $\bar{\nu}$ to a grading $\nu\colon (\mathcal{X}, \leq) \to \mathsf{Q}$ via

$$\mathcal{X} \ni \xi \mapsto \min\{\bar{\nu}(\eta) \colon \eta \in \mathrm{star}(\xi) \cap \mathcal{X}^+\} \in \mathsf{Q}$$

The map ν is a poset morphism since if $\xi \leq \eta$ then $\mathrm{star}(\eta) \subseteq \mathrm{star}(\xi)$. As $(\mathcal{X}, \nu)$ is a Q-graded cell complex we may consider the Birkhoff dual $\mathsf{O}(\nu)\colon \mathsf{O}(\mathsf{Q}) \to \mathsf{Sub}_{Cl}(\mathcal{X})$. Since Q is totally ordered, the collection $\{\mathsf{O}(\nu)(a)\}_{a \in \mathsf{O}(\mathsf{Q})}$ is a filtration of $\mathcal{X}$. This is the standard input for the topological data analysis pipeline.

Example 3.3.23. Consider $(\mathbb{R}^n, \leq)$ where $\leq$ is given by

$$(a_1, \ldots, a_n) \leq (b_1, \ldots, b_n) \iff a_i \leq b_i \quad \text{for all } i$$

Let $(\mathsf{P}, \leq)$ be a poset where $\mathsf{P} \subseteq \mathbb{R}^n$ and the partial order $\leq$ is inherited from $\mathbb{R}^n$. Let $(\mathcal{X}, \nu)$ be a P-graded cell complex and consider $\mathsf{O}(\nu)\colon \mathsf{O}(\mathsf{P}) \to \mathsf{Sub}_{Cl}(\mathcal{X}, \leq)$. In the theory of multi-parameter persistence [7], the collection $\{\mathsf{O}(\nu)(a)\}_{a \in \mathsf{O}(\mathsf{P})}$ of subcomplexes is called a *one-critical multi-filtration* of $\mathcal{X}$, since any cell enters the lattice/multi-filtration at a unique minimal element with respect to the partial order on $\mathsf{O}(\mathsf{P})$. Namely, a cell ξ enters the multi-filtration at $\downarrow\nu(\xi)$. Multi-filtrations can be converted to one-critical multi-filtrations via the mapping telescope [7].

3.3.5 Homotopy Category of Graded Complexes

It follows from the general construction of the homotopy category $\mathbf{K}(\mathcal{A})$ for an additive category $\mathcal{A}$ in Section 2.3 that there is a homotopy category $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ of the category of P-graded chain complexes $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$. To unpack this a bit, first recall the definition P-filtered chain maps in Section 3.3.2. We say that two P-filtered chain maps $\phi, \psi\colon (C, \pi) \to (C', \pi)$ are P*-filtered chain homotopic* if there is a P-filtered chain contraction $\gamma\colon C \to C'$ such that $\phi_n - \psi_n = \gamma_{n-1} \circ \partial_n + \partial'_{n+1} \circ \gamma_n$. We denote this by $\psi \sim_{\mathsf{P}} \phi$. The map γ is called a P*-filtered chain homotopy* from ϕ to ψ. A P-filtered chain map $\phi\colon (C, \pi) \to (C', \pi)$ is a P*-filtered chain equivalence* if there is a P-filtered chain map $\psi\colon (C', \pi) \to (C, \pi)$ such that $\psi\phi \sim_{\mathsf{P}} \mathrm{id}_C$ and $\phi\psi \sim_{\mathsf{P}} \mathrm{id}_{C'}$. In this case we say that (C, π) and (C', π) are P*-filtered chain equivalent.*

Following Definition 2.3.8, the *homotopy category of* P*-graded chain complexes*, denoted by $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$, is the category whose objects are P-graded chain complexes and whose morphisms are P-filtered chain homotopy equivalence classes of P-filtered chain maps. There is a quotient functor $q\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ which sends each P-graded chain complex to itself and each P-filtered chain map to its P-filtered chain homotopy equivalence class.

Example 3.3.24. The P-graded chain complexes $(C(\mathcal{X}), \pi)$ and $(C(\mathcal{X}'), \pi)$ of Example 1.2.1 are P-filtered chain equivalent via P-filtered chain maps

$$\phi\colon C(\mathcal{X}) \to C(\mathcal{X}') \quad \psi\colon C(\mathcal{X}') \to C(\mathcal{X}),$$

and P-filtered chain homotopies

$$\gamma\colon C(\mathcal{X}) \to C(\mathcal{X}) \quad \gamma'\colon C(\mathcal{X}') \to C(\mathcal{X}'),$$

which are described below. The nonzero differentials are

$$\partial_1 = \begin{array}{c c} & \begin{array}{cc} e_0 & e_1 \end{array} \\ \begin{array}{c} v_0 \\ v_1 \\ v_2 \end{array} & \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \end{array} \qquad \partial_1' = \begin{array}{c c} & \begin{array}{c} e_0' \end{array} \\ \begin{array}{c} v_0' \\ v_1' \end{array} & \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{array}.$$

The nonzero parts of the chain maps are ϕ and ψ are as follows.

$$\psi_0 = \begin{array}{c c} & \begin{array}{ccc} v_0 & v_1 & v_2 \end{array} \\ \begin{array}{c} v_0' \\ v_1' \end{array} & \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \end{pmatrix} \end{array} \qquad \phi_0 = \begin{array}{c c} & \begin{array}{cc} v_0' & v_1' \end{array} \\ \begin{array}{c} v_0 \\ v_1 \\ v_2 \end{array} & \begin{pmatrix} 1 & 0 \\ 0 & 0 \\ 0 & 1 \end{pmatrix} \end{array}$$

$$\psi_1 = \begin{array}{c c} & \begin{array}{cc} e_0 & e_1 \end{array} \\ e_0' & \begin{pmatrix} 1 & 0 \end{pmatrix} \end{array} \qquad \phi_1 = \begin{array}{c c} & \begin{array}{c} e_0' \end{array} \\ \begin{array}{c} e_0 \\ e_1 \end{array} & \begin{pmatrix} 1 \\ 1 \end{pmatrix} \end{array}$$

In this case $\gamma' = 0$. And the nonzero part of γ is

$$\gamma_0 = \begin{array}{c} \\ e_0 \\ e_1 \end{array} \!\!\begin{array}{c} \begin{array}{ccc} v_0 & v_1 & v_2 \end{array} \\ \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} \end{array}.$$

A lengthy but routine calculation shows that ψ and ϕ are P-filtered chain maps and that $\phi \circ \psi = \gamma \circ \partial + \partial \circ \gamma$ and $\psi \circ \phi = \mathrm{id}$.

We can now introduce our definition of the connection matrix. In particular, our definition of connection matrix rests on the homotopy-theoretic language.

Definition 3.3.25. Let (C, π) be a P-graded chain complex. A P-graded chain complex (C', π) is a *Conley complex for* (C, π) if

1. (C', π) is strict, i.e., $\partial_j^{pp} = 0$ for all p and j, and
2. (C', π) is isomorphic to (C, π) in $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$.

If (C', π) is a Conley complex for (C, π) then we say the associated boundary operator $\partial' = \{\Delta'^{pq}\}_{p,q \in \mathsf{P}}$ is a *connection matrix for* (C, π).

Remark 3.3.26. With the definition in place, we make some remarks about uniqueness and existence.

- Given a P-graded chain complex (C, π), a Conley complex (C', π) for (C, π) exists. This follows from the proof of correctness of Algorithm 3.7.8.
- A classical issue in Conley theory is the non-uniqueness of the connection matrix. In our treatment of connection matrix theory using chain equivalence and homotopy categories it turns out that Conley complexes are unique up to isomorphism. Thus a connection matrix is unique up to a similarity transformation in the sense that if one fixes a basis, then given two connection matrices Δ and Δ' there is a P-filtered chain map Φ such that $\Delta' = \Phi^{-1}\Delta\Phi$, cf. [18]. See Remark 3.3.31.

Example 3.3.27. Consider $(C(\mathcal{X}), \pi)$ and $(C(\mathcal{X}'), \pi)$ of Example 1.2.1. A straightforward verification shows that $(C(\mathcal{X}'), \pi)$ is strict and an object of $\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P}))$

(recall Definition 3.3.3). Moreover, from Example 3.3.24 we see that $(C(\mathcal{X}'), \pi)$ and $(C(\mathcal{X}), \pi)$ are isomorphic in $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$. Therefore $(C(\mathcal{X}'), \pi)$ is a Conley complex for $(C(\mathcal{X}), \pi)$ and $\partial' = \{\Delta^{pq}\}$ is a connection matrix for $(C(\mathcal{X}), \pi)$.

Proposition 2.2.11 allows for the following definition.

Definition 3.3.28. Let $\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$ denote the full subcategory of $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ whose objects are the objects of $\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P}))$. Then

$$\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) = \mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P}))/\sim_{\mathsf{P}}$$

and there is a quotient functor $q\colon \mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$.

Proposition 3.3.29. *Strict* P*-graded chain complexes are isomorphic in* $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ *if and only if they are* P*-filtered chain equivalent.*

Proof. The 'only if' direction is immediate; set the homotopies $\gamma = \gamma' = 0$. If (C, π) and (C', π) are P-filtered chain equivalent then there are P-filtered chain equivalences

$$\phi\colon (C, \pi) \to (C', \pi) \qquad \psi\colon (C, \pi) \to (C', \pi)$$

and P-filtered chain homotopies

$$\gamma\colon (C, \pi) \to (C, \pi) \qquad \gamma'\colon (C', \pi) \to (C', \pi).$$

It follows from Proposition 3.3.4 that

$$\psi_j^{pp}\phi_j^{pp} - \mathrm{id}_C^{pp} = (\psi_j\phi_j - \mathrm{id}_C)^{pp} = (\gamma_j\partial_j + \partial_j\gamma_j)^{pp} = \gamma_j^{pp}\partial_j^{pp} + \partial_j^{pp}\gamma_j^{pp} = 0$$

Therefore each entry ϕ_j^{pp} is an isomorphism with inverse ψ_j^{pp}. It follows from elementary matrix algebra that ϕ is an isomorphism. $\square$

Corollary 3.3.30. *The quotient functor* $q\colon \mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$ *is a conservative functor.*

Remark 3.3.31*.* Proposition 3.3.29 addresses non-uniqueness of the connection matrix in our formulation. In particular, the connection matrix is unique up to a choice of basis. Non-uniqueness manifests as a change of basis. See Section 4.4 for some applications where non-uniqueness arises. See [17, 18, 40] for more discusion of non-uniqueness in connection matrix theory.

3.3.6 Examples

Example 3.3.32 (Computing Homology)**.** Consider the situation where $\mathcal{X}$ is a cell complex and one is interested in computing the homology $H(\mathcal{K})$ of a closed subcomplex $\mathcal{K} \subset \mathcal{X}$. Connection matrix theory applies to this situation in the following fashion. Let $\mathsf{Q} = \{0, 1\}$ be the poset with $0 \leq 1$. Define the order-preserving map $\nu\colon (\mathcal{X}, \leq) \to \mathsf{Q}$ via

$$\nu(x) = \begin{cases} 0 & x \in \mathcal{K} \\ 1 & x \in \mathcal{X} \setminus \mathcal{K} \end{cases}$$

The pair $(\mathcal{X}, \nu)$ is a P-graded cell complex and $(C(\mathcal{X}), \pi^\nu)$ is the associated P-graded chain complex (see Definition 3.3.18). We have $\mathcal{K} = \nu^{-1}(0)$ and for any $j \in \mathbb{Z}$

$$C_j(\mathcal{X}) = C_j(\mathcal{X}^0) \oplus C_j(\mathcal{X}^1) = C_j(\mathcal{K}) \oplus C_j(\mathcal{X} \setminus \mathcal{K}).$$

Let (D, π) be a Conley complex for $(C(\mathcal{X}), \pi^\nu)$. Then for each $j \in \mathbb{Z}$

$$D_j = D_j^0 \oplus D_j^1.$$

Moreover, (D, π) is P-graded, the boundary operator $\partial_j\colon D_j \to D_{j-1}$ can be written as the matrix

$$\partial_j = \begin{array}{c c} & \begin{array}{cc} D_j^0 & D_j^1 \end{array} \\ \begin{array}{c} D_{j-1}^0 \\ D_{j-1}^1 \end{array} & \begin{pmatrix} \Delta_j^{00} & \Delta_j^{01} \\ 0 & \Delta_j^{11} \end{pmatrix} \end{array}.$$

The first condition in the definition of Conley complex (Definition 3.3.25) gives that (D, π) is strict. Therefore $\Delta_j^{00} = 0$ and $\Delta_j^{11} = 0$. The second condition in the definition implies that there is a P-filtered chain equivalence $\phi\colon (D, \pi) \to (C(\mathcal{X}), \pi^\nu)$. We can write $\phi_j\colon D_j \to C_j(\mathcal{X})$ as

$$\partial_j = \begin{array}{c c} & \begin{array}{cc} D_j^0 & D_j^1 \end{array} \\ \begin{array}{c} C_{j-1}(\mathcal{X}^0) \\ C_{j-1}(\mathcal{X}^1) \end{array} & \begin{pmatrix} \Phi_j^{00} & \Phi_j^{01} \\ 0 & \Phi_j^{11} \end{pmatrix} \end{array}.$$

It follows that the map

$$\Phi_\bullet^{00}\colon D_\bullet^0 \to C_\bullet(\mathcal{X}^0)$$

is a chain equivalence. Thus for all $j \in \mathbb{Z}$

$$H_j(C_\bullet(\mathcal{X}^0)) \cong H_j(D^0_\bullet, \Delta^{00}_\bullet) = D^0_j$$

where the last equality follows from Corollary 3.3.15 and the fact that (D, π) is strict.

Example 3.3.33 (Long Exact Sequence)**.** Consider $\mathsf{Q} = \{0, 1\}$ from Example 3.3.32 and $(\mathcal{X}, \nu)$, $(\mathcal{X}', \nu')$ and P from Example 1.2.1. There is an epimorphism $\rho\colon \mathsf{P} \to \mathsf{Q}$ given by

$$\rho(x) = \begin{cases} 0 & x = p \\ 0 & x = r \\ 1 & x = q \end{cases}$$

Let $\mu\colon \mathcal{X} \to \mathsf{Q}$ be the composition $\mu = \rho \circ \nu$ so that $(\mathcal{X}, \mu)$ is a Q-graded cell complex and $\mathcal{X}$ partitions as $\mathcal{X} = \mathcal{X}^0 \sqcup \mathcal{X}^1$, where $\mathcal{X}^i = \mu^{-1}(i)$. $\mathcal{X}^0$ is a closed subcomplex and $\mathcal{X}^1$ is an open subcomplex. There is a short exact sequence

$$0 \to C(\mathcal{X}^0) \to C(\mathcal{X}) \to C(\mathcal{X}^1) \to 0$$

In the associated long exact sequence on homology all homology groups are zero aside from the following:

$$\ldots \to H_1(\mathcal{X}^1) \xrightarrow{\delta} H_0(\mathcal{X}^0) \to H_0(\mathcal{X}) \to \ldots$$

A straightforward computation shows that this sequence is

$$\ldots \to \mathbb{Z}_2 \xrightarrow{\begin{pmatrix} 1 \\ 1 \end{pmatrix}} \mathbb{Z}_2 \oplus \mathbb{Z}_2 \xrightarrow{\begin{pmatrix} 1 & 1 \end{pmatrix}} \mathbb{Z}_2 \to \ldots$$

Consider the Q-graded complex $(\mathcal{X}', \mu')$ where $\mu' = \rho \circ \nu'$. A quick verification shows that the chain map $\phi\colon C(\mathcal{X}') \to C(\mathcal{X})$ of Example 3.3.24 is a Q-filtered chain equivalence. Therefore $(C(\mathcal{X}'), \pi^{\mu'})$ is a Conley complex for $(C(\mathcal{X}), \pi^\mu)$. The map ϕ induces a morphism of short exact sequences:

$$\begin{tikzcd}
0 \arrow[r] & C(\mathcal{X}^0) \arrow[r] \arrow[d, "\Phi^{00}"] & C(\mathcal{X}) \arrow[r] \arrow[d, "\phi"] & C(\mathcal{X}^1) \arrow[r] \arrow[d, "\Phi^{11}"] & 0 \\
0 \arrow[r] & C(\mathcal{X}'^0) \arrow[r] & C(\mathcal{X}') \arrow[r] & C(\mathcal{X}'^1) \arrow[r] & 0
\end{tikzcd}.$$

The morphism of short exact sequences induces a morphism of the long exact sequences. The fact that ϕ is a Q-filtered chain equivalence implies that the induced maps on homology are isomorphisms.

$$
\begin{tikzcd}
\dots \arrow[r] & H_1(\mathcal{X}^1) \arrow[r, "\delta"] \arrow[d, "\cong"] & H_0(\mathcal{X}^0) \arrow[r] \arrow[d, "\cong"] & H_0(\mathcal{X}) \arrow[r] \arrow[d, "\cong"] & \dots \\
\dots \arrow[r] & H_1(\mathcal{X}'^1) \arrow[r, "\delta'"] & H_0(\mathcal{X}'^0) \arrow[r] & H_0(\mathcal{X}') \arrow[r] & \dots \\
& C_1(\mathcal{X}') \arrow[r, "\Delta_1'^{01}"] \arrow[u, "\mathrm{id}"] & C_0(\mathcal{X}'^0) \arrow[r] \arrow[u, "\mathrm{id}"] & C_0(\mathcal{X}') \arrow[u, "\mathrm{id}"] &
\end{tikzcd}
$$

This discussion shows that in the setting of a Q-graded complex – where $\mathsf{Q} = \{0,1\}$ – the connection matrix Δ' is the connecting homomorphism of the long exact sequence.

3.4 Filtered Complexes

For the remainder of this section we fix L in **FDLat**.

3.4.1 Filtered Vector Spaces

Definition 3.4.1. An L-*filtered vector space* $V = (V, f)$ is a vector space V equipped with a lattice morphism $f\colon \mathsf{L} \to \mathsf{Sub}(V)$. We call f an L-*filtering* of V. Suppose (V, f) and (W, g) are L-filtered vector spaces. A map $\phi\colon V \to W$ is L-filtered if

$$\phi(f(a)) \subseteq g(a), \quad \text{for all } a \in \mathsf{L}.$$

The *category of* L-*filtered vector spaces*, denoted $\mathbf{FVect}(\mathsf{L})$, is the category whose objects are L-filtered vector spaces and whose morphisms are L-filtered linear maps.

Since f is a finite lattice homomorphism, we have that under f

$$0_{\mathsf{L}} \mapsto 0 \qquad \text{and} \qquad 1_{\mathsf{L}} \mapsto V.$$

We write a family of L-filtered vector spaces as $(V_\bullet, f_\bullet) = \{(V_n, f_n)\}_{n\in\mathbb{Z}}$. For a fixed $a \in \mathsf{L}$ there is a family of vector spaces $f(a) = f_\bullet(a) = \{f_n(a)\}_{n\in\mathbb{Z}}$.

Proposition 3.4.2. *Let (V, π) and (W, π) be P-graded vector spaces. A linear map $\phi\colon (V, \pi) \to (W, \pi)$ is P-filtered if and only if $\phi(V^a) \subset W^a$ for all $a \in \mathsf{J}(\mathsf{O}(\mathsf{P}))$.*

Proof. We start with showing that $\phi(V^a) \subset W^a$ if ϕ is P-filtered. Let $a \in \mathsf{J}(\mathsf{O}(\mathsf{P}))$. From Birkhoff's Theorem there exists $s \in \mathsf{P}$ such that $a = \downarrow s$. If $x \in V^a$ then $x = \pi^a(x) = \sum_{r \leq s} \pi^r(x)$. Since a is a down-set, if $p \leq s$ then $p \in a$ and $W^p \subset W^a$. Therefore

$$\phi(x) = \sum_{p \leq q} \sum_{r \leq s} \phi^{pq} \pi^r(x) = \sum_{p \leq q \leq s} \phi^{pq}(x) \in W^a.$$

Conversely, assume that $\phi(V^a) \subset W^a$ for all $a \in \mathsf{J}(\mathsf{O}(\mathsf{P}))$. Suppose $\phi^{pq}(x) \neq 0$ for $p, q \in \mathsf{P}$ and $x \in V$. Let b denote $\downarrow q$. Then $\pi^q(x) \in V^b$ and $\phi(\pi^q(x)) \subset W^b$. We have $\phi^{pq}(x) = \pi^p(\phi(\pi^q(x)) \neq 0$, which implies $p \in b$. Therefore $p \leq q$. $\square$

The previous result enables the definition of a functor which constructs a lattice-filtered chain complex from a poset-graded chain complex. Recall that u is the forgetful functor $u \colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{Vect}$ which forgets the grading.

Definition 3.4.3. Let $\mathsf{L} = \mathsf{O}(\mathsf{P})$. Define the functor $\mathfrak{L} \colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{FVect}(\mathsf{L})$ via

$$\mathfrak{L}((V, \pi)) := (u(V, \pi), f) = (V, f)$$

where the L-filtering $f \colon \mathsf{L} \to \mathsf{Sub}(V)$ sends $a \in \mathsf{O}(\mathsf{P})$ to

$$V^a = u^a((V, \pi)) \in \mathsf{Sub}(V)\,.$$

Proposition 3.4.2 states that a P-filtered map $\phi \colon (V, \pi) \to (W, \pi)$ is L-filtered. Therefore we define $\mathfrak{L}$ to be the identity on morphisms:

$$\mathfrak{L}(\phi) := \phi \in \mathrm{Hom}_{\mathbf{FVect}}((V, f), (W, g))$$

Theorem 3.4.4. *Let $\mathsf{L} := \mathsf{O}(\mathsf{P})$. The functor $\mathfrak{L} \colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{FVect}(\mathsf{L})$ is additive, full, faithful and essentially surjective.*

Proof. The functor $\mathfrak{L}$ is additive since $\mathfrak{L}$ is an identity on hom-sets. That $\mathfrak{L}$ is a bijection on hom-sets (fully faithful) follows from Proposition 3.4.2. We now show that $\mathfrak{L}$ is essentially surjective. Let (V, f) be an L-filtered vector space; first we will construct a P-graded vector space (W, π) and then we will show that it satsifies $\mathfrak{L}(W, \pi) = (V, f)$. $\mathsf{Sub}(V)$ is a relatively complemented lattice (see Definition 2.4.16 and Example 2.4.17). Therefore we may choose for each join irreducible $p \in \mathsf{J}(\mathsf{L})$ a subspace $W^p \in \mathsf{Sub}(V)$

such that $W^p + f(\overleftarrow{p}) = f(p)$ and $W^p \cap f(\overleftarrow{p}) = 0$. Thus $f(p) = W^p \oplus f(\overleftarrow{p})$. As L is in **FDLat**, Proposition 2.4.12 gives that any $a \in \mathsf{L}$ can be written as the irredundant join of join-irreducibles, i.e., we have $a = \vee_i q_i$ with $q_i \in \mathsf{J}(\mathsf{L})$. Thus

$$f(a) = f(\vee_i q_i) = \vee_i f(q_i)\,.$$

It follows from well-founded induction over the underlying poset of L, that for all $a \in \mathsf{L}$ that

$$f(a) = \bigoplus_{\substack{q \leq a,\\ q \in \mathsf{J}(\mathsf{L})}} W^q.$$

Now set $W = \bigoplus_{q \in \mathsf{J}(\mathsf{L})} W^q$ and $\pi = \{\pi^q\}_{q \in \mathsf{J}(\mathsf{L})}$ where π^q is defined to be the projection $\pi^q \colon V \to W^q$. (W, π) is a $\mathsf{J}(\mathsf{L})$-graded vector space. From Birkhoff's theorem, $\mathsf{J}(\mathsf{L})$ and P are isomorphic, which implies that (W, π) may be regarded as a P-graded vector space. Now we show that $\mathfrak{L}(W, \pi) = (V, f)$. From the definition of $\mathfrak{L}$ it suffices to choose $a \in \mathsf{O}(\mathsf{P})$ and show that $W^a = f(a)$. This follows since $f(a) = \bigoplus_{q \leq a} W^q = W^a$. $\square$

3.4.2 Filtered Chain Complexes

Similar to $\mathbf{GrVect}(\mathsf{P})$, the category $\mathbf{FVect}(\mathsf{L})$ is additive but not abelian. Following Section 2.3 once again, we may form the category $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ of chain complexes in $\mathbf{FVect}(\mathsf{L})$. An object C of $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ is a chain complex in L-filtered vector spaces. For short, we say that this is an L-*filtered chain complex.* The data of C can be unpacked as the triple $C = (C_\bullet, \partial_\bullet, f_\bullet)$ where:

1. $(C_\bullet, \partial_\bullet)$ is a chain complex,
2. (C_n, f_n) is an L-filtered vector space for each n, and
3. $\partial_n \colon (C_n, f_n) \to (C_{n-1}, f_{n-1})$ is an L-filtered linear map.

We will denote C as (C, f) to distinguish the L-filtering. A morphism $\phi \colon (C, f) \to (C', f')$ is a chain map $\phi \colon (C, \partial) \to (C', \partial')$ such that for each n, $\phi_n \colon (C_n, f_n) \to (C'_n, f'_n)$ is an L-filtered linear map. We entitle the morphisms of $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ the L-*filtered chain maps.*

If (C, f) is an L-filtered chain complex then $\partial_n(f_n(a)) \subseteq f_{n-1}(a)$. Thus $\{f_n(a)\}_{n\in\mathbb{Z}}$ together with $\{\partial_n|_{f_n(a)}\}_{n\in\mathbb{Z}}$ is a subcomplex of C. We define the map $f\colon \mathsf{L} \to \mathsf{Sub}(C)$ via

$$f(a) := \big(\{f_n(a)\}, \{\partial_n|_{f_n(a)}\}\big) \in \mathsf{Sub}(C) \tag{3.9}$$

A chain complex equipped with a lattice homomorphism $\mathsf{L} \to \mathsf{Sub}(C)$ is the object that Robbin and Salamon work with. The next two results show that these two perspectives are equivalent. The proofs are immediate, and are included for completeness.

Proposition 3.4.5. *If (C, f) is an L-filtered chain complex then $f\colon \mathsf{L} \to \mathsf{Sub}(C)$, given as in* (3.9) *is a lattice morphism.*

Proof. Let $a, b \in \mathsf{L}$. Let $A = \{f_n(a)\}_{n\in\mathbb{Z}}$ and $B = \{f_n(b)\}_{n\in\mathbb{Z}}$. Then

$$f(a) \vee f(b) = (A, \partial|_A) \vee (B, \partial|_B) = (A + B, \partial|_{A+B}) = f(a \vee b)$$
$$f(a) \wedge f(b) = (A, \partial|_A) \wedge (B, \partial|_B) = (A \cap B, \partial|_{A+B}) = f(a \vee b) \qquad \square$$

Proposition 3.4.6. *Let $C = (C_\bullet, \partial_\bullet)$ be a chain complex together with a lattice morphism $f\colon \mathsf{L} \to \mathsf{Sub}(C)$. If $\{f_n\colon \mathsf{L} \to \mathsf{Sub}(C_n)\}_{n\in\mathbb{Z}}$ is the family of maps defined as*

$$f_n(a) := A_n \subseteq C_n \qquad \textit{where } f(a) = (A_\bullet, \partial^A_\bullet),$$

then $(C_\bullet, \partial_\bullet, f_\bullet)$ is an L-filtered chain complex.

Proof. We show first show that $(C_\bullet, f_\bullet)$ is a family of L-filtered vector spaces spaces. Let $a, b \in \mathsf{L}$. Let $A_\bullet = f(a), B_\bullet = f(b)$ and $D_\bullet = f(a \vee b)$. As f is a lattice morphism, we have that

$$A_\bullet \vee B_\bullet = f(a) \vee f(b) = f(a \vee b) = D_\bullet.$$

This implies $A_n + B_n = D_n$ for all n. It follows that $f_n(a) \vee f_n(b) = A_n + B_n = D_n = f_n(a \vee b)$. Similarly, it follows that $f_n(a) \wedge f_n(b) = f_n(a \wedge b)$. Observe that ∂_n is an L-filtered linear map for each n because $f(a) \in \mathsf{Sub}(C)$ implies that $\partial_n f_n(a) \subseteq f_{n-1}(a)$. $\square$

3.4.3 The Subcategory of Strict Objects

There is a notion of strict object in the category $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$. Recall from Definition 2.4.10 the notion of a join-irreducible element $a \in \mathsf{J}(\mathsf{L})$ and its unique predecessor $\overleftarrow{a} \in \mathsf{L}$.

Definition 3.4.7. We say that an L-filtered chain complex (C, f) is *strict* if

$$\partial_n(f_n(a)) \subseteq f_{n-1}(\overleftarrow{a}), \quad \text{for all } a \in \mathsf{J}(\mathsf{L}) \text{ and } n \in \mathbb{Z}.$$

In this case we say that f is a *strict filtering*. The strict objects form a subcategory $\mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})) \subset \mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$, called the *subcategory of strict objects*.

Remark 3.4.8. Recall from Definition 2.4.3 that $\mathsf{O}(\mathsf{P})$ is the lattice of down-sets of a poset P. Recall the definition of 'connection matrix' from [42, Section 8]: a *connection matrix (in the sense of Robbin-Salamon)* is an $\mathsf{O}(\mathsf{P})$-filtered chain complex (C, f) such that for any $b \in \mathsf{O}(\mathsf{P})$ and $n \in \mathbb{Z}$

$$\partial_n(f_n(b)) \subset f_{n-1}(b \setminus \{p\})$$

whenever p is maximal in b.

For any $b \in \mathsf{O}(\mathsf{P})$, p is maximal in b if and only if b covers $b \setminus \{p\}$. Therefore, the following result shows that our notion of a strict L-filtered complex is equivalent to their definition of connection matrix.

Proposition 3.4.9. *Let (C, f) be an L-filtered chain complex. Then (C, f) is strict if and only if it obeys the following property: given $n \in \mathbb{Z}$ and $a, b \in \mathsf{L}$ such that b covers a then $\partial_n(f_n(b)) \subseteq f_{n-1}(a)$.*

Proof. The 'if' direction is immediate: a covers $\overleftarrow{a}$ for $a \in \mathsf{J}(\mathsf{L})$. Thus $\partial_n(f_n(a)) \subseteq f_{n-1}(\overleftarrow{a})$. Now suppose that (C, f) is strict and that b covers a. As L is in $\mathbf{FDLat}$, Proposition 2.4.12 states that any $b \in \mathsf{L}$ can be written as the irredundant join of join-irreducibles, i.e., we have $b = \vee_i q_i$ with $q_i \in \mathsf{J}(\mathsf{L})$. Since f is a lattice morphism,

$$f_n(b) = f_n(\vee_i q_i) = \vee_i f_n(q_i).$$

Moreover, since b covers a there is precisely one q_j such that $a \vee q_j = b$ with $q_j \not\leq a$ and $q_i \leq a$ for $i \neq j$. That b covers a implies that $\overleftarrow{q_j} \leq a$, otherwise $a < a \vee \overleftarrow{q_j} < b$. For any $x \in f_n(b)$ we have $x = \sum_i x_i$ with $x_i \in f_n(q_i)$ and $\partial_n(x) = \sum_i \partial_n(x_i)$. Since f is a strict filtering, $\partial_n(x_i) \in f_{n-1}(\overleftarrow{q_i})$ and $\partial_n(x) \in \vee_i f_{n-1}(\overleftarrow{q_i}) = f_{n-1}(\vee_i \overleftarrow{q_i}) \subseteq f_{n-1}(a)$. $\square$

3.4.4 Equivalence of Categories

We now examine the relationship between graded and filtered chain complexes. Our primary aim is to establish an equivalence between these two categories, as well as their strict subcategories. With $\mathsf{L} = \mathsf{O}(\mathsf{P})$, it follows from Proposition 2.3.9 that the functor $\mathfrak{L}\colon \mathbf{GrVect}(\mathsf{P}) \to \mathbf{FVect}(\mathsf{L})$ (see Definition 3.4.3) induces a functor

$$\mathfrak{L}_{\mathbf{Ch}}\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}(\mathbf{FVect}(\mathsf{L})). \tag{3.10}$$

Recalling the definition of $\mathfrak{L}_{\mathbf{Ch}}$ from Section 2.3, we have

$$\mathfrak{L}_{\mathbf{Ch}}(C, \pi) := (C, f)$$

where the L-filtering $f\colon \mathsf{L} \to \mathsf{Sub}(C)$ is given by

$$\mathsf{L} \ni a \mapsto (C^a_\bullet, \Delta^a_\bullet) = u^a((C, \pi)) \in \mathsf{Sub}(C)\,.$$

Here, u^a is the forgetful functor $u^a\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}$ described in Definition 3.3.8. When the context is clear we abbreviate $\mathfrak{L}_{\mathbf{Ch}}$ by $\mathfrak{L}$.

Theorem 3.4.10. *Let* $\mathsf{L} = \mathsf{O}(\mathsf{P})$. *The functor* $\mathfrak{L}\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ *is additive, fully faithful and essentially surjective (hence a categorical equivalence). Moreover,* $\mathfrak{L}$ *restricts to an equivalence of the subcategories*

$$\mathfrak{L}\colon \mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})).$$

Proof. The first part follows from from Theorem 3.4.4 and Proposition 2.3.10. For the second part, suppose (C, π) is a P-graded chain complex and (C, f) is an L-filtered chain complex such that $\mathsf{L} = \mathsf{O}(\mathsf{P})$ and $\mathfrak{L}((C, \pi)) = (C, f)$. We show that (C, π) is strict if and only if (C, f) is strict. We first show that if (C, π) is strict then (C, f) is strict. Let

$a \in \mathsf{J}(\mathsf{O}(\mathsf{P}))$. We show that $\partial(f(a)) \subseteq f(\overleftarrow{a})$. By Birkhoff's theorem, there exists $s \in \mathsf{P}$ such that $a = \downarrow s$ and $\overleftarrow{a} = \bigcup_{p<s} \downarrow p$. If $x \in C^a_\bullet$ then $x = \pi^a(x)$. Hence

$$\partial x = \partial(\pi^a(x)) = \sum_{p<q} \partial^{pq} \sum_{r \leq s} \pi^r(x) = \sum_{p<q\leq s} \partial^{pq}(x).$$

Since $\partial^{pq}(x) \in C^p$ and $\overleftarrow{a} = \bigcup_{p<s} \downarrow p$ it follows that $\partial x \in f(\overleftarrow{a})$ as desired.

We now show that if (C, f) is strict then (C, π) is strict. Let $p \in \mathsf{P}$ and a denote $\downarrow p$. Let $x \in C$. Then $\pi^p(x) \in C^p \subset C^a$. Since (C, f) is strict $\partial(\pi^p(x)) \in \partial(C^a) \subset C^{\overleftarrow{a}}$. That $p \notin \overleftarrow{a}$ implies $\partial^{pp} = \pi^p \partial(\pi^p(x))) = 0$. Therefore $\mathfrak{L}$ restricts to an equivalence of the strict subcategories. $\square$

3.4.5 Filtered Cell Complexes

We consider again the data analysis perspective, and define the appropriate concept for cell complexes. Recall from Section 2.5 that the notion of subcomplex for a cell complex is more general than for a chain complex. Given a cell complex $\mathcal{X} = (\mathcal{X}, \leq, \kappa, \dim)$ we work with $\mathsf{Sub}_{Cl}(\mathcal{X})$, the lattice of closed subcomplexes.

Definition 3.4.11. An L-*filtered cell complex* is a cell complex $\mathcal{X} = (\mathcal{X}, \leq, \kappa, \dim)$ together with a lattice morphism $f \colon \mathsf{L} \to \mathsf{Sub}_{Cl}(\mathcal{X})$. The morphism f is called an L-*filtering* of $\mathcal{X}$. We write $(\mathcal{X}, f)$ to denote an L-filtered cell complex.

Definition 3.4.12. Let $(\mathcal{X}, \nu)$ be a P-graded cell complex and $\mathsf{L} = \mathsf{O}(\mathsf{P})$. The *associated* L-*filtered chain complex* is the pair $(C(\mathcal{X}), f^\nu)$ where f^ν is the composition

$$\mathsf{L} \xrightarrow{\mathsf{O}(\nu)} \mathsf{Sub}_{Cl}(\mathcal{X}) \xrightarrow{\text{span}} \mathsf{Sub}(C(\mathcal{X}))$$

given explicitly by sending $a \in \mathsf{L}$ to

$$\text{span}(\mathsf{O}(\nu)(a)) = \left\{ \sum_{i=0}^{n} \lambda_i \xi_i : n \in \mathbb{N}, \lambda_i \in \mathbb{K}, \xi_i \in \mathsf{O}(\nu)(a) \right\} \in \mathsf{Sub}(C(\mathcal{X})).$$

We write $\mathcal{L} \colon \text{Cell}(\mathsf{P}) \to \mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ for the assignment $(\mathcal{X}, \nu) \mapsto (C(\mathcal{X}), f^\nu)$.

Example 3.4.13. Let X be a closed manifold and $\varphi \colon \mathbb{R} \times X \to X$ a Morse-Smale gradient flow. Let P be the set of fixed points and $\mu \colon \mathsf{P} \to \mathbb{N}$ the assignment of Morse indices. Each unstable manifold $W^u(p)$ is a $\mu(p)$-cell and the manifold X admits a cellular decomposition $(\mathcal{X}, \leq, \kappa, \mu)$. Furthermore, the closure $\mathrm{cl}(W^u(p))$ is a cell subcomplex of $\mathcal{X}$. The map

$$\mathsf{O}(\mathsf{P}) \ni a \mapsto \bigcup_{p \in a} \mathrm{cl}(W^u(p))$$

is an $\mathsf{O}(\mathsf{P})$-filtering of the cellular (Morse) complex $(\mathcal{X}, \leq, \kappa, \mu)$.

The next result follows from an examination of the definitions of $\mathcal{C}, \mathcal{L}$ and $\mathfrak{L}$.

Proposition 3.4.14. *Let* $\mathsf{L} = \mathsf{O}(\mathsf{P})$. *The functor* $\mathfrak{L} \colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$ *fits into the following commutative diagram with the assignments* $\mathcal{C}$ *and* $\mathcal{L}$ *(denoted by dashes arrows).*

$$\begin{tikzcd}
\mathrm{Cell}(\mathsf{P}) \arrow[r, dashed, "\mathcal{C}"] \arrow[dr, dashed, "\mathcal{L}"'] & \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \arrow[d, "\mathfrak{L}"] \\
& \mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))
\end{tikzcd}$$

3.4.6 Homotopy Category of Filtered Complexes

Once again we may follow Section 2.3 to introduce the homotopy category $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$ of the category of L-filtered chain complexes $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$. To spell this out a bit further, we say that two L-filtered chain maps $\phi, \psi \colon (C, f) \to (D, g)$ are L-*filtered chain homotopic* if there is an L-filtered chain contraction $\gamma \colon C \to D$ such that $\phi - \psi = \gamma \circ \partial + \partial \circ \gamma$. We denote this by $\psi \sim_{\mathsf{L}} \phi$.

Proceeding as in Section 2.3, the *homotopy category of* L*-filtered chain complexes*, which we denote by $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$, is the category whose objects are L-filtered chain complexes and whose morphisms are L-filtered chain homotopy equivalence classes of L-filtered chain maps. It follows from Proposition 2.3.9 that the functor $\mathfrak{L}$ induces a functor on the homotopy categories $\mathfrak{L}_{\mathbf{K}} \colon \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}(\mathbf{FVect}(\mathsf{L}))$. This functor is defined on objects as $\mathfrak{L}_{\mathbf{K}}((C, \pi)) = \mathfrak{L}_{\mathbf{Ch}}(C, \pi)$ and on morphisms as $\mathfrak{L}_{\mathbf{K}}([\phi]_{\mathsf{P}}) = [\phi]_{\mathsf{L}}$. Moreover, this functor satisfies the identity

$$\mathfrak{L}_{\mathbf{K}} \circ q = q \circ \mathfrak{L}_{\mathbf{Ch}}.$$

Definition 3.4.15. Let $\mathbf{K}_s(\mathbf{FVect}(\mathsf{L}))$ denote the full subcategory of $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$

whose objects are the objects of $\mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L}))$. Then

$$\mathbf{K}_s(\mathbf{FVect}(\mathsf{L})) = \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L}))/\sim_{\mathsf{L}}$$

and there is a quotient functor $q\colon \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})) \to \mathbf{K}_s(\mathbf{FVect}(\mathsf{L}))$.

Proposition 3.4.16. *Let* $\mathsf{L} = \mathsf{O}(\mathsf{P})$. *The functors* $\mathfrak{L}_{\mathbf{K}}\colon \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}(\mathbf{FVect}(\mathsf{L}))$ *and* $\mathfrak{L}_{\mathbf{K}}\colon \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}_s(\mathbf{FVect}(\mathsf{L}))$ *are equivalences of categories.*

In analogy to Proposition 3.3.29, strict L-filtered chain complexes which are L-filtered chain equivalent are isomorphic in $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$. The may be phrased in terms of the following result.

Proposition 3.4.17. *The functor* $q\colon \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})) \to \mathbf{K}_s(\mathbf{FVect}(\mathsf{L}))$ *is conservative.*

Proof. By Birkhoff's Theorem there is some P such that $\mathsf{L} \cong \mathsf{O}(\mathsf{P})$. Without loss of generality, we let $\mathfrak{L}$ be the composition (of equivalences of categories)

$$\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \xrightarrow{\mathfrak{L}} \mathbf{Ch}(\mathbf{FVect}(\mathsf{O}(\mathsf{P}))) \to \mathbf{Ch}(\mathbf{FVect}(\mathsf{L})).$$

We have the following diagram.

$$\begin{tikzcd}
\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \arrow[r, "\mathfrak{L}"] \arrow[d, "q"] & \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})) \arrow[d, "q"] \\
\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) \arrow[r, "\mathfrak{L}_{\mathbf{K}}"] & \mathbf{K}_s(\mathbf{FVect}(\mathsf{L}))
\end{tikzcd}$$

Theorem 3.4.10 and Proposition 3.4.16 state, respectively, that $\mathfrak{L}$ and $\mathfrak{L}_{\mathbf{K}}$ are equivalences of categories. Moreover, Proposition 3.3.30 states that $q\colon \mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$ is conservative. It follows that q is conservative. $\square$

Corollary 3.4.18. *Let* (C, f) *and* (D, g) *be strict* L*-filtered chain complexes.* (C, f) *and* (D, g) *are* L*-filtered chain isomorphic if and only if they are isomorphic in* $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$.

Remark 3.4.19. Proposition 3.4.17 implies that, up to isomorphism, the strict L-filtered chain complexes are an invariant of the L-filtered chain equivalence class.

Fix P and $\mathsf{L} = \mathsf{O}(\mathsf{P})$. Consider the following commutative diagram.

$$
\begin{tikzcd}
& \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})) \arrow[rr, dotted] \arrow[dd, dotted, "q"] & & \mathbf{Ch}(\mathbf{FVect}(\mathsf{L})) \arrow[dd, dotted] \\
\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \arrow[ur, "\mathfrak{L}"] \arrow[rr, dotted] \arrow[dd, dotted, "q"'] & & \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \arrow[ur, "\mathfrak{L}"] \arrow[dd, dotted] & \\
& \mathbf{K}_s(\mathbf{FVect}(\mathsf{L})) \arrow[rr, dashed] & & \mathbf{K}(\mathbf{FVect}(\mathsf{L})) \\
\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) \arrow[ur, "\mathfrak{L}"] \arrow[rr, dashed] & & \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \arrow[ur, "\mathfrak{L}"] &
\end{tikzcd}
\tag{3.11}
$$

Our results thus far have the following implications.

- It follows from Theorem 3.4.10 and Proposition 3.4.16 that the solid arrows are equivalences of categories.

- Propositions 3.3.29 and 3.4.17 show that the quotient functors

  $$q\colon \mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})), \quad q\colon \mathbf{Ch}_s(\mathbf{FVect}(\mathsf{L})) \to \mathbf{K}_s(\mathbf{FVect}(\mathsf{L}))$$

  are conservative.

- Finally, a bit further along in the thesis, the proof of correctness of the Algorithm establishes that the dashed arrows – inclusion functors $\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ and $\mathbf{K}_s(\mathbf{FVect}(\mathsf{L})) \to \mathbf{K}(\mathbf{FVect}(\mathsf{L}))$ – are essentially surjective. In particular Theorem 3.8.1 and Corollary 3.8.2 of Section 3.7 will show that these are equivalences of categories.

Finally, we reach our definition of Conley complex and connection matrix for an L-filtered chain complex.

Definition 3.4.20. Let (C, f) be an L-filtered chain complex and $\mathsf{L} = \mathsf{O}(\mathsf{P})$. A P-graded chain complex (C', π) is called a *Conley complex* for (C, f) if

1. (C', π) is an object of $\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P}))$;
2. $\mathfrak{L}(C', \pi)$ is isomorphic to (C, f) in $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$.

With the theory that has been built up, the following result is straightforward.

Theorem 3.4.21. *Let $(\mathcal{X}, \nu)$ be a P-graded cell complex. Let $\mathsf{L} = \mathsf{O}(\mathsf{P})$. If (C', π) is a Conley complex for $(C(\mathcal{X}), \pi^\nu)$ then (C', π) is a Conley complex for $(C(\mathcal{X}), f^\nu)$.*

Proof. Since (C', π) is a Conley complex for $(C(\mathcal{X}), \pi^\nu)$, by definition it is an object of $\mathbf{Ch}_s(\mathbf{GrVect}(\mathsf{P}))$. Moreover, by definition $q(C', \pi) \cong q(C(\mathcal{X}), \pi^\nu)$. It follows from (3.11) that

$$q \circ \mathfrak{L}(C', \pi) = \mathfrak{L}_{\mathbf{K}} \circ q(C', \pi) \cong \mathfrak{L}_{\mathbf{K}} \circ q(C(\mathcal{X}), \pi^\nu) = q \circ \mathfrak{L}(C(\mathcal{X}), \pi^\nu)$$

It follows from Proposition 3.4.14 that $q \circ \mathfrak{L}(C(\mathcal{X}), \pi^\nu) = q(C(\mathcal{X}), f^\nu)$. Therefore $q \circ \mathfrak{L}(C', \pi) \cong q(C(\mathcal{X}), f^\nu)$. $\square$

Conceptually, Theorem 3.4.21 implies that one may do homotopy-theoretic computations within the category $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ in order to compute the relevant objects of interest for $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$. At this point in the paper, we refer the reader back to the left hand side of Diagram (3.1).

3.5 Franzosa's Connection Matrix Theory

In this section we will review connection matrix theory as developed by R. Franzosa in the sequence of papers [15, 16, 17] from the late 1980's. Briefly, the connection matrix is the appropriate generalization of the Morse boundary operator for Conley theory; it is a boundary operator defined on Conley indices. The connection matrix allows one to recover the graded module braid that is obtained from the index lattice. However, unlike the Morse boundary operator, the connection matrix is not defined from trajectories, it is only related to them. The basic function of the connection matrix is to prove the existence of connecting orbits [32]. At a higher level, it serves as an algebraic representation of global dynamics and in certain cases can be used to construct (semi)-conjugacies of the global attractor [10, 34, 35]. Its preeminent function is to promote Conley index theory to a homology theory [33].

3.5.1 The Categories of Braids

It was Conley's observation [8] that focusing on the attractors of a dynamical system provides a generalization of Smale's Spectral Decomposition [47, Theorem 6.2]. There is a lattice structure to the attractors of a dynamical system [27, 28, 42] and one is often naturally led to studying a finite sublattice of attractors A and an associated sublattice of attracting blocks N with $\omega\colon \mathsf{N} \to \mathsf{A}$ (see Definition 2.9.3). This setup is expressed in the diagram below.

$$\begin{tikzcd} \mathsf{N} \arrow[r, "\subset"] \arrow[d, two heads, "\omega"] & \mathsf{ABlock}_{\mathscr{R}}(\varphi) \arrow[d, two heads, "\omega"] \\ \mathsf{A} \arrow[r, "\subset"] & \mathsf{Att}(\varphi) \end{tikzcd}$$

A sublattice of attracting blocks is what Franzosa terms an *index filtration* [17, 15, 16]. However, as these sublattices are not necessarily totally ordered, we follow [27] and call this an *index lattice*.

In his work, Franzosa introduces the notion of a *chain complex braid* as a data structure to hold the singular chain complexes that arise out of the topological data within the index lattice. The chain complex braid is organized by the poset of join-irreducibles of the index lattice. Implicit in Franzosa's work is a description of a category for chain complex braids over a fixed poset P. We now describe this category, which we label $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$. First we recall the notion of adjacent convex sets.

Definition 3.5.1. An ordered collection $(I_1, \ldots, I_N)$ of convex sets of P is *adjacent* if

1. $I_1, \ldots, I_n$ are mutually disjoint;
2. $\bigcup_{i=1}^{n} I_i$ is a convex set in P;
3. For all $p, q \in \mathsf{P}$, $p \in I_i, q \in I_j, i < j$ implies $q \not< p$.

We are primarily interested in adjacent pairs of convex sets (I, J) and for simplicity write the union $I \cup J$ as IJ. We denote the set of convex sets as $I(\mathsf{P})$ and the set of adjacent tuples and triples of convex sets as $I_2(\mathsf{P})$ and $I_3(\mathsf{P})$.

Definition 3.5.2. We say that a pair (I, J) of convex sets is *incomparable* if p and q are incomparable for any $p \in I$ and $q \in J$. This immediately implies that (I, J) and (J, I) are adjacent.

Definition 3.5.3. Following [17], we say that a sequence of chain complexes and chain maps

$$C_1 \xrightarrow{i} C_2 \xrightarrow{p} C_3$$

is *weakly exact* if i is injective, $p \circ i = 0$ and $p \colon C_2/\operatorname{im}(i) \to C_3$ induces an isomorphism on homology.

Proposition 3.5.4 ([17], Proposition 2.2)**.** *Let*

$$C_1 \xrightarrow{i} C_2 \xrightarrow{p} C_3$$

be a weakly exact sequence of chain complexes and ∂_i the boundary operator of C_i. There exists a natural degree -1 homomorphism $\partial \colon H(C_3) \to H(C_1)$ such that

1. *if $[x] \in H(C_3)$ then $\partial([x]) = [i^{-1}\partial_2 p^{-1}(x)]$,*
2. *$\ldots \to H(C_1) \xrightarrow{i} H(C_2) \xrightarrow{p} H(C_3) \xrightarrow{\partial} H(C_1) \to \ldots$ is exact.*

Proof. We sketch the proof, following [17]. Let $\partial' \colon H(C_2/\operatorname{im}(i)) \to H(C_1)$ be the connecting homomorphism for the short exact sequence of chain complexes

$$0 \to C_1 \xrightarrow{i} C_2 \xrightarrow{\rho} C_2/\operatorname{im}(i) \to 0$$

where ρ is projection onto the quotient. Since the sequence is weakly exact, ρ induces an isomorphism on homology $\rho_* \colon H(C_2)/\operatorname{im}(i)) \to H(C_3)$. Define $\partial \colon H(C_3) \to H(C_1)$ by $\partial = \partial' \rho_*^{-1}$. □

Definition 3.5.5. A *chain complex braid* $\mathscr{C}$ over P is a collection of chain complexes and chain maps in $\mathbf{Ch}(\mathbf{Vect})$ such that

1. for each $I \in I(\mathsf{P})$ there is a chain complex $(C(I), \Delta(I))$,
2. for each $(I, J) \in I_2(\mathsf{P})$ there are chain maps

$$i(I, IJ) \colon C(I) \to C(IJ) \qquad \text{and} \qquad p(IJ, J) \colon C(IJ) \to C(J),$$

which satisfy

(a) $C(I) \xrightarrow{i(I,IJ)} C(IJ) \xrightarrow{p(IJ,J)} C(J)$ is weakly exact,

(b) if I and J are incomparable then $p(JI, I)i(I, IJ) = \mathrm{id}|_{C(I)}$,

(c) if $(I, J, K) \in I_3(P)$ then the following braid diagram commutes.

$$\begin{tikzcd}[column sep=small]
 & & C(J) \arrow[dr] & & \\
 & C(IJ) \arrow[ur] \arrow[dr] & & C(JK) \arrow[dr] & \\
C(I) \arrow[ur] \arrow[rr] & & C(IJK) \arrow[ur] \arrow[rr] & & C(K)
\end{tikzcd}$$

Definition 3.5.6. The *category of chain complex braids over* P, denoted $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$, is the category whose objects are chain complex braids over P. Given two chain complex braids $\mathscr{C}$ and $\mathscr{C}'$ a morphism $\Psi\colon \mathscr{C} \to \mathscr{C}'$ is a collection of chain maps $\{\Psi(I)\colon C(I) \to C'(I)\}_{I\in I(\mathsf{P})}$ such that for $(I, J) \in I_2(\mathsf{P})$ the following diagram commutes.

$$\begin{tikzcd}
C(I) \arrow[r] \arrow[d, "\Psi(I)"] & C(IJ) \arrow[r] \arrow[d, "\Psi(IJ)"] & C(J) \arrow[d, "\Psi(J)"] \\
C'(I) \arrow[r] & C'(IJ) \arrow[r] & C'(J)
\end{tikzcd}$$

For a given sublattice of attractors A, two index lattices N, N' associated with the same sublattice of attractors A, i.e., $\omega(\mathsf{N}) = \mathsf{A}$ and $\omega(\mathsf{N}') = \mathsf{A}$, may yield different chain complex braids. However, the homology groups of the chain complexes contained in the chain complex braid are an invariant. This is the motivation for the idea of a *graded module braid*, which formalizes the notion of 'homology' for a chain complex braid. In order for our terminology to agree with Franzosa's in [17], we introduce 'graded module braids' in the generality of graded R-modules. However, for the results of this paper we specifically work in the case when R is a field, and a graded R-module is a graded vector space.

Definition 3.5.7. Let R be a ring. A *graded R-module* is a family $M_\bullet = \{M_n\}_{n\in\mathbb{Z}}$ of R-modules. A *graded R-module homomorphism* is a family $f\colon M_\bullet \to M'_\bullet$ are families of R-module homomorphisms $f = \{f_n\colon M_n \to M'_n\}_{n\in\mathbb{Z}}$. The category of *graded R-modules*, denoted $\mathbf{R\text{-}Mod}^{\mathbb{Z}}$, is the category whose objects are graded R-modules and whose morphisms are graded R-module homomorphisms.

Definition 3.5.8. Let $M_\bullet$ and $M'_\bullet$ be graded R-modules. A *degree d map* γ from $M_\bullet$ to $M'_\bullet$ is a family of R-module homomorphisms $\{\gamma_n\colon M_n \to M'_{n+d}\}_{n\in\mathbb{Z}}$.

Definition 3.5.9. A *graded R-module braid over* P $\mathscr{G}$ is a collection of graded R-modules maps satisfying:

1. for each $I \in I(\mathsf{P})$ there is a graded R-module $G(I)$;
2. for each $(I, J) \in I_2(\mathsf{P})$ there are maps:

$$\begin{aligned} i(I, IJ)\colon G(I) \to G(IJ) &\text{ of degree 0,} \\ p(IJ, J)\colon G(IJ) \to G(J) &\text{ of degree 0,} \\ \partial(J, I)\colon G(J) \to G(I) &\text{ of degree -1} \end{aligned}$$

   which satisfy

   (a) $\ldots \to G(I) \xrightarrow{i} G(IJ) \xrightarrow{p} G(J) \xrightarrow{\partial} G(I) \to \ldots$ is exact,

   (b) if I and J are incomparable then $p(JI, I)i(I, IJ) = \mathrm{id}|_{G(I)}$

   (c) if $(I, J, K) \in I_3(P)$ then the braid diagram (3.12) commutes.

$$\begin{tikzcd}
\vdots \arrow[d] & \cdots \arrow[dl] \arrow[dr] & \vdots \arrow[d] \\
G(I) \arrow[dr, "i"] \arrow[dd, bend right, "i"'] & & G(K) \arrow[dl, "\partial"'] \arrow[dd, bend left, "\partial"] \\
& G(IJ) \arrow[dl, "i"'] \arrow[dr, "p"] & \\
G(IJK) \arrow[dr, "p"] \arrow[dd, bend right, "p"'] & & G(J) \arrow[dl, "i"'] \arrow[dd, bend left, "\partial"] \\
& G(JK) \arrow[dl, "p"'] \arrow[dr, "\partial"] & \\
G(K) \arrow[dr, "\partial"] \arrow[dd, bend right, "\partial"'] & & G(I) \arrow[dl, "i"'] \arrow[dd, bend left, "i"] \\
& G(IJ) \arrow[dl, "p"'] \arrow[dr, "i"] & \\
G(J) \arrow[d] \arrow[dr] & & G(IJK) \arrow[d] \arrow[dl] \\
\vdots & \cdots & \vdots
\end{tikzcd} \tag{3.12}$$

Definition 3.5.10. A morphism $\Theta\colon \mathscr{G} \to \mathscr{G}'$ of graded R-module braids is a collection of graded R-module homomorphisms $\{\Theta(I)\colon G(I) \to G'(I)\}_{I \in I(\mathsf{P})}$ such that for each

$(I, J) \in I_2(\mathsf{P})$ the following diagram commutes:

$$
\begin{tikzcd}
\dots \arrow[r] & G(I) \arrow[r, "i"] \arrow[d, "\Theta(I)"] & G(IJ) \arrow[r, "p"] \arrow[d, "\Theta(IJ)"] & G(J) \arrow[r, "\partial"] \arrow[d, "\Theta(J)"] & G(I) \arrow[r] \arrow[d, "\Theta(I)"] & \dots \\
\dots \arrow[r] & G'(I) \arrow[r, "i'"] & G'(IJ) \arrow[r, "p'"] & G'(J) \arrow[r, "\partial'"] & G'(I) \arrow[r] & \dots
\end{tikzcd}
$$

Remark 3.5.11. Since a morphism of braids $\Theta\colon \mathscr{G} \to \mathscr{G}'$ involves a fixed map $\Theta(I)$ for each convex set I, there is a commutative diagram involving the two braid diagrams of (3.12) and Θ for any $(I, J, K) \in I_3(\mathsf{P})$. In fact, as remarked in [3, 36] one does not need to use graded R-module braids, but only a collection of long exact sequences given this definition of morphism.

Definition 3.5.12. Given a fixed ring R, the *category of graded R-module braids over* P, denoted by $\mathbf{GMB}(\mathsf{P}, \mathbf{R\text{-}Mod}^{\mathbb{Z}})$, is the category of graded R-module braids and their morphisms.

When R may be understood from the context we refer to a graded R-module braid as a graded module braid. This terminology matches Franzosa [17]. For the purposes of this paper, R is a field and we work with $\mathbf{GMB}(\mathsf{P}, \mathbf{Vect}^{\mathbb{Z}})$. Implicit in [17, Proposition 2.7] is the description of a functor from $\mathfrak{H}\colon \mathbf{ChB}(\mathsf{P}, \mathbf{Vect}) \to \mathbf{GMB}(\mathsf{P}, \mathbf{Vect}^{\mathbb{Z}})$ which is the analogy of the homology functor.

Definition 3.5.13. A pair of chain complex braid morphisms $\Psi, \Phi\colon \mathscr{C} \to \mathscr{C}'$ are P-*braided homotopic* if there is a collection $\{\Gamma(I)\colon C(I) \to C'(I)\}_{I \in I(\mathsf{P})}$ of chain contractions such that for each I

$$\Phi(I) - \Psi(I) = \Delta(I)\Gamma(I) + \Gamma(I)\Delta'(I).$$

The collection $\Gamma = \{\Gamma(I)\}_{I \in I(\mathsf{P})}$ is called a P-*braided chain homotopy*. We write $\Psi \sim_{\mathsf{P}} \Phi$ if Ψ and Φ are P-braided homotopic.

Proposition 3.5.14. *The binary relation $\sim_{\mathsf{P}}$ is a congruence relation on* $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$.

Definition 3.5.15. Let $\mathscr{C}, \mathscr{C}'$ are chain complex braids. A morphism of chain complex braids $\Phi\colon \mathscr{C} \to \mathscr{C}'$ is a P-*braided chain equivalence* if there is a P-braided chain map $\Psi\colon \mathscr{C}' \to \mathscr{C}$ such that $\Psi\Phi \sim_{\mathsf{P}} \mathrm{id}_{\mathscr{C}}$ and $\Phi\Psi \sim_{\mathsf{P}} \mathrm{id}'_{\mathscr{C}}$. The *homotopy category of chain* *complex braids over* P, which we denote $\mathbf{KB}(\mathsf{P}, \mathbf{Vect})$, is the category whose objects are chain complex braids over P and whose morphisms are P-braided chain homotopy equivalence classes of chain complex braid morphisms. In other words, $\mathbf{KB}(\mathsf{P}, \mathbf{Vect})$ is the quotient category $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})/\sim_{\mathsf{P}}$ formed by defining the hom-sets via

$$\mathrm{Hom}_{\mathbf{KB}(\mathsf{P},\mathbf{Vect})}(\mathscr{C}, \mathscr{C}') = \mathrm{Hom}_{\mathbf{ChB}(\mathsf{P},\mathbf{Vect})}(\mathscr{C}, \mathscr{C}')/\sim_{\mathsf{P}}$$

where $\sim_{\mathsf{P}}$ is the braided homotopy equivalence relation. Denote by $q\colon \mathbf{ChB}(\mathsf{P}, \mathbf{Vect}) \to \mathbf{KB}(\mathsf{P}, \mathbf{Vect})$ the quotient functor which sends each chain complex braid over P to itself and each chain complex braid morphism to its P-braided chain homotopy equivalence class. It follows from the construction that two chain complex braids $\mathscr{C}, \mathscr{C}'$ are isomorphic in $\mathbf{KB}(\mathsf{P}, \mathbf{Vect})$ if and only if $\mathscr{C}, \mathscr{C}'$ are P-braided chain equivalent.

Proposition 3.5.16. *Let $\mathscr{C}$ and $\mathscr{C}'$ be chain complex braids over P. If $\mathscr{C}, \mathscr{C}'$ are braided chain equivalent then $\mathfrak{H}(\mathscr{C}) \cong \mathfrak{H}(\mathscr{C}')$. In particular, there is a functor*

$$\mathfrak{H}_{\mathbf{K}}\colon \mathbf{KB}(\mathsf{P}, \mathbf{Vect}) \to \mathbf{GMB}(\mathsf{P}, \mathbf{Vect}^{\mathbb{Z}})$$

that sends braided chain equivalences to graded module braid isomorphisms.

3.5.2 Franzosa's Connection Matrix

We previously introduced our notion of connection matrix in Definition 3.3.25. In this section we review Franzosa's definition. Let $\mathscr{C}$ be a chain complex braid in $\mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$. Historically, the connection matrix was introduced as a P-filtered (upper triangular) boundary operator Δ on the direct sum of homological Conley indices associated to the elements of P

$$\Delta\colon \bigoplus_{p\in\mathsf{P}} H_\bullet(C(p)) \to \bigoplus_{p\in\mathsf{P}} H_\bullet(C(p))$$

which recovers the associated graded module braid $\mathfrak{H}(\mathscr{C})$. See Definition 3.5.20 for the precise notion. Δ may be thought of as a matrix of linear maps $\{\Delta^{pq}\}$ and the identification with the matrix structure is the genesis of the phrase *connection matrix.*

Recall that the functor $\mathfrak{L}$ of (3.10) is used to build an $\mathsf{O}(\mathsf{P})$-filtered chain complex from P-graded complex. The next results show that graded chain complexes can be used to build chain complex braids. First, recall the forgetful functor u as well as the family of forgetful functors $\{u^I\}$ parameterized by the convex sets $I \in I(\mathsf{P})$ defined in Definition 3.3.8 (see also Section 3.3). For a P-graded chain complex (C,π) and a convex set $I \subset \mathsf{P}$, $u \circ u^I(C,\pi) = (C^I, \Delta^I)$ is a chain complex. Given a P-filtered chain map $\phi\colon (C,\pi) \to (C',\pi)$, $u^I(\phi)$ is the chain map $\Phi^I = e^I \circ \phi \circ i^I\colon C^I \to C'^I$.

Proposition 3.5.17 ([17], Proposition 3.4)**.** *Let (C,π) be an P-graded chain complex. The collection $\mathscr{C} = \mathfrak{B}(C,\pi)$ consisting of the chain complexes $\{u^I(C,\pi)\}_{I\in(\mathsf{P})}$ and the natural chain maps $i(I,IJ)$ and $p(IJ,J)$ for each $(I,J) \in I_2(\mathsf{P})$ form a chain complex braid over P.*

Proposition 3.5.18 ([18], Proposition 3.2)**.** *Let (C,π) and (C',π) be P-graded chain complexes. If $\phi\colon (C,\pi) \to (C',\pi)$ is a P-filtered chain map then the collection $\{\Phi^I\colon C^I \to C'^I\}_{I\in I(\mathsf{P})}$ is a chain complex braid morphism from $\mathfrak{B}(C,\pi)$ to $\mathfrak{B}(C',\pi)$.*

Propositions 3.5.17 and 3.5.18 describe a functor

$$\mathfrak{B}\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{ChB}(\mathsf{P},\mathbf{Vect}) \tag{3.13}$$

That is, the functor $\mathfrak{B}$ is defined on objects as $\mathfrak{B}(C,\pi) = \{C^I, \Delta^I\}_{I\in I(\mathsf{P})}$ together with the natural inclusion and projection maps. Moreover, $\mathfrak{B}$ is defined on morphisms as

$$\mathfrak{B}(\phi) = \{\Phi^I\colon C^I \to C'^I\}_{I\in I(\mathsf{P})}$$

Proposition 3.5.19. *The functor $\mathfrak{B}\colon \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{ChB}(\mathsf{P},\mathbf{Vect})$ is additive. Moreover $\mathfrak{B}$ induces a functor on homotopy categories*

$$\mathfrak{B}_{\mathbf{K}}\colon \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{KB}(\mathsf{P},\mathbf{Vect})$$

We can now state Franzosa's definition of connection matrix. In brief, this is a P-graded chain complex capable of reconstructing the appropriate graded module braid.

Definition 3.5.20 ([17], Definition 3.6)**.** Let $\mathscr{G}$ be a graded module braid over P and (C,π) be a P-graded chain complex. The boundary operator ∂ of (C,π) is called a *C-connection matrix for $\mathscr{G}$* if

1. There is an isomorphism of graded module braids

$$\mathfrak{H} \circ \mathfrak{B}(C, \pi) \cong \mathscr{G}. \tag{3.14}$$

2. If, in addition, (C, π) is strict, i.e., $\partial^{pp} = 0$ for all $p \in \mathsf{P}$, then ∂ is a *connection matrix (in the sense of Franzosa) for* $\mathscr{G}$.

In light of Definition 3.5.20, the connection matrix is an efficient codification of data which is capable of recovering the braid $\mathscr{G}$. In Conley theory, a graded module braid $\mathscr{G}$ over $\mathsf{J}(\mathsf{L})$ is derived from a index lattice. In this way the connection matrix is a graded object (over $\mathsf{J}(\mathsf{L})$) capable of recovering (up to homology) the data of the index lattice. Moreover, both the chain complex braid $\mathscr{C} = \mathfrak{B}(C, \pi)$ and the graded module braid $\mathscr{G} = \mathfrak{H} \circ \mathfrak{B}(C, \pi)$ associated to the connection matrix are simple objects in their appropriate categories. Observe that

1. For $\mathscr{C}$ we have $C(I) = \bigoplus_{p \in I} C^p$ for all $I \in I(\mathsf{P})$.

2. For the graded module braid $\mathscr{G}$ if $[\alpha] \in G(I)$ then $\partial(J, I)([\alpha]) = [\Delta^{J,I}(\alpha)]$ from [17, Proposition 3.5].

The next result is one of Franzosa's theorems on existence of connection matrices, written in our terminology. As Franzosa works with R-modules, instead of vector spaces as we do, he assumes the chain complexes consist of free R-modules.

Theorem 3.5.21 ([17], Theorem 4.8)**.** *Let $\mathscr{C}$ be a chain complex braid over P. Let $\{B^p\}_{p \in \mathsf{P}}$ be a collection of free chain complexes such that $H(B^p) \cong H(C(p))$ and set $B = \bigoplus_{p \in \mathsf{P}} B^p$. There exists a P-filtered boundary operator Δ so that (B, π), where $\pi = \{\pi^p \colon B \to B^p\}_{p \in \mathsf{P}}$, is a P-graded chain complex. Moreover, there exists a morphism of chain complex braids $\Psi \colon \mathscr{B} \to \mathscr{C}$ where $\mathscr{B} = \mathfrak{B}(B, \pi)$ such that $\mathfrak{H}(\Psi)$ is a graded module braid isomorphism.*

Here is a simple application of Franzosa's theorem. Let $\mathscr{C}$ be a chain complex braid. Choose $B = \{C(p)\}_{p \in P}$. The theorem says that there exists a P-graded chain complex (C, π) and a morphism of chain complex braids $\Psi \colon \mathfrak{B}(C, \pi) \to \mathscr{C}$ that induces an isomorphism on graded module braids. Therefore for any chain complex braid over

P there is a simple representative (one coming from a P-graded chain complex (B, π)) that is quasi-isomorphic to $\mathscr{C}$ (in the sense that there is a morphism Ψ of chain complex braids that induces an isomorphism on graded module braids). In the case when one works with fields the homology $H(C(p))$ of each chain complex $C(p)$ is a $\mathbb{Z}$-graded vector space (see Definition 2.3.11). Therefore we may choose $B = \{H_\bullet(C(p))\}_{p\in\mathsf{P}}$. Invoking the theorem gives a P-graded chain complex (B, π) such that

$$\Delta\colon \bigoplus_{p\in\mathsf{P}} H_\bullet(C(p)) \to \bigoplus_{p\in\mathsf{P}} H_\bullet(C(p))$$

In our terminology this implies that (B, π) is a Conley complex and Δ is a connection matrix, both in the sense of our definition of connection matrix (Definition 3.3.25) and of Definition 3.5.20 of Franzosa.

Remark 3.5.22*.* The classical definition of the connection matrix (Definition 3.5.20) does not involve a chain equivalence. In particular, the connection matrix is not associated to a representative of a chain equivalence class. In fact, in Franzosa's definition the isomorphism of (3.14) is not required to be induced from a chain complex braid morphism.

Remark 3.5.23*.* We do not develop the theory of reductions (see Section 3.6) for braids, although it is straightforward to do so.

Definition 3.5.24. Let $(\mathcal{X}, \nu)$ be a P-graded cell complex. The preimage of each convex set $\mathcal{X}^I := \nu^{-1}(I)$ is a convex set in $(\mathcal{X}, \leq, \kappa, \dim)$. Therefore each $\mathcal{X}^I = (\mathcal{X}^I, \leq^I, \kappa^I, \dim^I)$ is a cell complex where $(\leq^I, \kappa^I, \dim^I)$ are the restrictions to $\mathcal{X}^I$. This implies that each $(C_\bullet(\mathcal{X}^I), \partial|_{\mathcal{X}^I})$ is a chain complex. A routine computation shows that the collection

$$\{(C_\bullet(\mathcal{X}^I), \partial|_{\mathcal{X}^I})\}_{I\in I(\mathsf{P})}$$

satisfies the axioms of a chain complex braid over P, and that this is precisely the image of $(\mathcal{X}, \nu)$ under the composition

$$\mathrm{Cell}(\mathsf{P}) \xrightarrow{\mathcal{C}} \mathbf{Ch}(\mathbf{GrVect}(\mathsf{P})) \xrightarrow{\mathfrak{B}} \mathbf{ChB}(\mathsf{P}, \mathbf{Vect}).$$

The composition defines an assignment $\mathcal{B}\colon \mathrm{Cell}(\mathsf{P}) \to \mathbf{ChB}(\mathsf{P}, \mathbf{Vect})$.

Theorem 3.5.25. *Let $(\mathcal{X}, \nu)$ be a P-graded cell complex, $(C(\mathcal{X}), \pi^\nu)$ be the associated P-graded chain complex and $\mathscr{G} = \mathfrak{H}(\mathcal{B}(\mathcal{X}, \nu))$ be the associated graded module braid. If (C', π) is a Conley complex for $(C(\mathcal{X}), \pi^\nu)$ then ∂' is a connection matrix (in the sense of Franzosa, Definition 3.5.20) for $\mathscr{G}$.*

Proof. By definition (C, π) and (C', π) are P-filtered chain equivalent. It follows from Proposition 3.5.19 that the associated chain complex braids $\mathfrak{B}(C', \pi)$ and $\mathfrak{B}(C, \pi)$ are P-braided chain equivalent. Then

$$\mathfrak{H} \circ \mathfrak{B}(C', \pi) \cong \mathfrak{H} \circ \mathfrak{B}(C, \pi) = \mathscr{G},$$

where the first isomorphism follows from Proposition 3.5.16 and the equality follows from the definition of $\mathcal{B}$. $\square$

Theorem 3.5.25 implies that one may do homotopy-theoretic computations within the category $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ in order to compute connection matrices in the classical sense of Definition 3.5.20. At this point in the paper we refer the reader back to the full Diagram (3.1), which encapsulates much of the machinery introduced thus far. Most importantly, taken together Theorems 3.4.21 and 3.5.25 imply that if one finds a strict P-graded chain complex (C', π) that is P-filtered chain equivalent to the given $(C(\mathcal{X}), \pi^\nu)$, then

1. One can construct a strict L-filtered chain complex, $\mathfrak{L}(C', \pi)$, which is chain equivalent to the associated lattice-filtered complex $(C(\mathcal{X}), f^\nu)$.

2. ∂' is a connection matrix in the classical sense of Definition 3.5.20 for the associated graded-module braid $\mathfrak{H}(\mathcal{B}(\mathcal{X}, \nu)$.

This implies that to compute connection matrices in both the sense of Franzosa [17] and Robbin-Salamon [42], it suffices to work within the category $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$.

3.6 Reductions

In this section we formalize the method of computing Conley complexes. Later, in Section 3.7.3, we detail a computational version, using discrete Morse theory, of the theory presented here. In Section 3.7 we present two algorithms based on the Morse theory: the first for computing homology using these methods, and the second for computing Conley complexes and connection matrices.

First, we review the tools for the chain complexes and the category $\mathbf{Ch}(\mathbf{Vect})$. Then we proceed to the graded and filtered versions within the categories $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$ and $\mathbf{Ch}(\mathbf{FVect}(\mathsf{L}))$, respectively. Much of the material may feel redundant, as we will port results from chain complexes to graded and filtered versions.

In computational homological algebra, one often finds a simpler representative with which to compute homology. A model for this is the notion of *reduction*, which is a particular type of chain homotopy equivalence. The notion also goes under the moniker *strong deformation retract* or sometimes chain contraction [46].[1] It appears in [12], in homological perturbation theory [4] and forms the basis for effective homology theory [44] and algebraic Morse theory [45, 46]. Our exposition of reductions primarily follows the preprint [44]. Roughly speaking, a reduction is a method of data reduction for a chain complex without losing any information with respect to homology.

Definition 3.6.1. A *reduction* is a pair of chain complexes and triple of maps

$$\begin{tikzcd} C \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & M \arrow[l, shift left, "\phi"] \end{tikzcd}$$

where ϕ, ψ are chain maps and a chain contraction γ satisfying the identities:

1. $\psi\phi = \mathrm{id}_M$
2. $\phi\psi = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$
3. $\gamma^2 = \gamma\phi = \psi\gamma = 0$.

From the definition it is clear that ϕ is a monomorphism and ψ is an epimorphism. In applications, one calls M the *reduced complex*. When reductions arise from algebraic-discrete Morse theory M is sometimes called the *Morse complex*. The point is that one

[1]We previously introduced the term *chain contraction* in Section 2.3 which agrees with [51]. This idea should not be confused with reduction.

wants $|M| \ll |C|$, then one may compute $H(M)$ (and thus $H(C)$) more efficiently. Notice that by using (3), an application γ on the left of (2) gives:

$$0 = (\gamma\phi)\psi = \gamma(\mathrm{id}_C - \gamma\partial - \partial\gamma) = \gamma - \gamma\partial\gamma \tag{3.15}$$

This equation is axiomized as the condition for a degree 1 map (see Definition 2.3.4) called a *splitting homotopy*.

Definition 3.6.2. Let C be a chain complex. A *splitting homotopy* is a degree 1 map $\gamma\colon C \to C$ such that $\gamma^2 = 0$ and $\gamma\partial\gamma = \gamma$.

The upshot is that reductions can be obtained from splitting homotopies. The conditions $\partial^2 = \gamma^2 = 0$ and $\gamma\partial\gamma = \gamma$ ensure that $\gamma\partial + \partial\gamma$ is idempotent. Therefore $\rho = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$ is a projection onto the complementary subspace to $\mathrm{im}(\gamma\partial + \partial\gamma)$. Since ρ is a projection, there is a splitting of C into subcomplexes:

$$C = \ker\rho \oplus \operatorname{im}\rho.$$

The image $(M, \partial^M) = (\operatorname{im}\rho, \partial|_{\operatorname{im}\rho})$ is a subcomplex of C. We have the following reduction:

$$\begin{tikzcd} C \arrow[loop above, "\gamma"] \arrow[r, shift left, "\rho"] & M. \arrow[l, shift left, "i"] \end{tikzcd} \tag{3.16}$$

We can calculate the differential ∂^M via

$$\partial\rho = \partial(\mathrm{id}_C - (\gamma\partial + \partial\gamma)) = \partial - \partial\gamma\partial + \partial\partial\gamma = \partial - \partial\gamma\partial.$$

Finally, it is straightforward that the remaining identities $\gamma i = \rho\gamma = 0$ are easily verified. Furthermore, $\ker\rho$ is a subcomplex of C and $\gamma|_{\ker\rho}$ is a chain contraction, since $\mathrm{id}_{\ker\rho} = \partial\gamma + \gamma\partial$. This implies that $\ker\rho$ is acyclic, i.e., $H_\bullet(\ker\rho) = 0$. As is well-known, reductions and splitting homotopies are (up to isomorphism) in bijective correspondence. We include a proof here for completeness.

Proposition 3.6.3. *Reductions and splitting homotopies are in bijective correspondence, up to isomorphism.*

Proof. From (3.15) we have that any chain homotopy in a reduction is a splitting homotopy. Moreover, any splitting homotopy can be put into the reduction in (3.16). Let γ be a splitting homotopy and consider two reductions:

$$\begin{tikzcd} M \arrow[r, shift left, "i"] & C \arrow[l, shift left, "\rho"] \arrow[loop above, "\gamma"] \arrow[r, shift left, "\rho'"] & M. \arrow[l, shift left, "i'"] \end{tikzcd}$$

A routine computation using the conditions (1)–(3) shows that the compositions $\rho \circ i'$ and $\rho' \circ i$ are inverses. Therefore M and M' are chain isomorphic. □

Example 3.6.4. Let $\mathcal{X}$ be a cell complex and let $(A, w\colon Q \to K)$ be an acyclic partial matching, see Section 2.8. By Proposition 2.8.2 there exists a unique splitting homotopy γ. From Theorem 2.8.3 defining the maps

$$\psi := \pi_A \circ (\mathrm{id}_{\mathcal{X}} - \partial\gamma) \qquad \phi := (\mathrm{id}_{\mathcal{X}} - \gamma\partial) \circ \iota_A \qquad \partial^A := \psi \circ \partial \circ \phi$$

leads to a reduction:

$$\begin{tikzcd} C_\bullet(\mathcal{X}) \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (C_\bullet(A), \partial^A) \arrow[l, shift left, "\phi"] \end{tikzcd} \tag{3.17}$$

Notice that this is a different reduction than the one defined in Diagram (3.16). However, we have $(C_\bullet(\mathcal{A}), \partial^{\mathcal{A}}) \cong (M, \partial^M)$ from Proposition 3.6.3. In contrast to Diagram (3.16), using the reduction of Diagram (3.17) has the property that the Morse complex is comprised of critical cells of the matching.

Definition 3.6.5. We say a reduction is *minimal* if the reduced complex M is minimal. We say a splitting homotopy γ is *perfect* if $\partial = \partial\gamma\partial$.

Proposition 3.6.6. *Minimal reductions and perfect splitting homotopies are in bijective correspondence.*

Proof. If the reduction is minimal, then $\partial^M = 0$. Thus $\partial i\rho = (i\partial^M)\rho = 0$. By hypothesis $i\rho = \mathrm{id}_C - \partial\gamma - \gamma\partial$. Application of ∂ to both sides yields

$$0 = \partial(i\rho) = \partial(\mathrm{id}_C - \partial\gamma - \gamma\partial) = \partial - \partial\gamma\partial.$$

Conversely, if γ is perfect then with $M = \mathrm{im}(\rho)$ the differential ∂^M is calculated as

$$\partial^M = \partial - \partial\gamma\partial = 0.$$

Therefore M is minimal and the reduction is minimal. $\square$

A perfect splitting homotopy implies $\operatorname{im}(\rho) \cong H_\bullet(C)$. This allows the homology to be read from the reduction without computation. In addition, we have $\partial i = i\partial^M = 0$. Therefore $\operatorname{im}(i) \subset \ker \partial$ and and the map $i\colon M \to \ker \partial$ gives representatives for the homology in C. In the case of fields, perfect splittings always exist.[2] This implies that a chain complex C and its homology $H_\bullet(C)$ always fit into a reduction. Moreover any reduction where C is a minimal complex is trivial in the sense that the two complexes are isomorphic.

Proposition 3.6.7. *Let C be a minimal chain complex. Any reduction*

$$\begin{tikzcd} C \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & M \arrow[l, shift left, "\phi"] \end{tikzcd}$$

is minimal. Moreover, we have $M \cong C$.

Proof. We have $\partial' = \partial'(\psi\phi) = \psi\partial\phi = 0$. We have $\psi \circ \phi = \mathrm{id}_M$. If C is minimal then $\phi \circ \psi = \mathrm{id}_C - (\gamma\partial + \partial\gamma) = \mathrm{id}_C$. $\square$

In this sense, the homology $H_\bullet(C)$ is the algebraic core of a chain complex and the minimal representative for C with respect to reductions. This result will have analogues in the graded and filtered cases. Finally, we show that reductions compose.

Proposition 3.6.8. *Given the sequence of reductions:*

$$\begin{tikzcd} C \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & M \arrow[l, shift left, "\phi"] \arrow[loop above, "\gamma'"] \arrow[r, shift left, "\psi'"] & M' \arrow[l, shift left, "\phi'"] \end{tikzcd}$$

there is a reduction

$$\begin{tikzcd} C \arrow[loop above, "\gamma''"] \arrow[r, shift left, "\psi''"] & M' \arrow[l, shift left, "\phi''"] \end{tikzcd}$$

with the maps given by the formulas

$$\phi'' = \phi \circ \phi' \qquad \psi'' = \psi' \circ \psi \qquad \gamma'' = \gamma + \phi \circ \gamma' \circ \psi.$$

[2]In fact, this is a Corollary of Algorithm 3.7.2.

Proof. Elementary computations show that

$$\rho'' \circ i'' = \mathrm{id}_{M''} \quad \text{and} \quad i'' \circ \rho'' = \mathrm{id}_C - (\partial\gamma'' + \gamma''\partial).$$

Conditions (1)–(3) follow from the same conditions for γ and γ'.

$$(\gamma'')^2 = (\gamma + i\gamma'\rho)(\gamma + i\gamma'\rho) = \gamma^2 + (\gamma i)\gamma'\rho + i\gamma'(\rho\gamma) + i\gamma'(\rho i)\gamma'\rho = i(\gamma'\gamma')\rho = 0$$

$$\gamma'' \circ i'' = (\gamma + i\gamma'\rho)(i \circ i') = (\gamma i)i' + \rho'(\rho i)\gamma'\rho = (\rho'\gamma)'\rho = 0$$

$$\rho'' \circ \gamma'' = (\rho'\rho)(\gamma + i\gamma'\rho) = \rho'(\rho\gamma) + \rho'(\rho i)\gamma'\rho = (\rho'\gamma')\rho = 0. \qquad \square$$

An inductive argument gives the following result the next result.

Proposition 3.6.9. *Given a tower of reductions*

$$\begin{tikzcd} C \arrow[loop above, "\gamma_0"] \arrow[r, shift left, "\psi_0"] & M_0 \arrow[l, shift left, "\phi_0"] \arrow[loop above, "\gamma_1"] \arrow[r, shift left, "\psi_1"] & \dots \arrow[l, shift left, "\phi_1"] \arrow[r, shift left, "\psi_{n-1}"] & M_{n-1} \arrow[l, shift left, "\phi_{n-1}"] \arrow[loop above, "\gamma_n"] \arrow[r, shift left, "\psi_n"] & M_n; \arrow[l, shift left, "\phi_n"] \end{tikzcd}$$

1. *there is a reduction*

$$\begin{tikzcd} C \arrow[loop above, "\Gamma"] \arrow[r, shift left, "\Psi_n"] & M_n \arrow[l, shift left, "\Phi_n"] \end{tikzcd} \tag{3.18}$$

   with maps given by the formulas

$$\Psi_m = \prod_{i=0}^{m} \psi_i \qquad \Phi_m = \prod_{i=0}^{m} \phi_i \qquad \Gamma = \gamma_0 + \sum_{i=0}^{n-1} \Phi_i \circ \gamma_{i+1} \circ \Psi_i.$$

2. Γ *is a splitting homotopy and* Γ *is perfect if any* γ_i *is perfect.*

Proof. Part (1) follows from Proposition 3.6.8 and an inductive argument. Given (3.18) the fact that γ is a splitting homotopy follows from the proof of Proposition 3.6.3. If γ_i is perfect, then M_i is minimal by Proposition 3.6.6. Thus M_j is minimal for $j \geq i$ by Proposition 3.6.7. In particular M_n is minimal and (3.18) is minimal. Therefore γ is a perfect splitting homotopy by Proposition 3.6.6. $\square$

3.6.1 Graded Reductions

Filtered and graded versions of the reductions are obtained by porting the definition to the appropriate category.

Definition 3.6.10. A P-*graded reduction* is a pair of P-graded chain complexes and triple of P-filtered maps

$$\begin{tikzcd} (C,\pi) \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (M,\pi) \arrow[l, shift left, "\phi"] \end{tikzcd}$$

where ϕ, ψ are chain maps and a chain contraction γ satisfying the identities:

1. $\psi\phi = \mathrm{id}_M$
2. $\phi\psi = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$
3. $\gamma^2 = \gamma\phi = \psi\gamma = 0.$

An P-graded reduction is *minimal* if (M,π) is strict.

Definition 3.6.11. A P-*filtered splitting homotopy* is a degree 1 map $\gamma\colon (C,\pi) \to (C,\pi)$ such that $\gamma^2 = 0$ and $\gamma\partial\gamma = \gamma$. A P-filtered splitting homotopy is *perfect* if $\partial^{pp} = \partial^{pp}\gamma^{pp}\partial^{pp}$ for each $p \in \mathsf{P}$.

Again, one may define $\rho = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$ and $M = \mathrm{im}(\rho)$. Then M is a P-graded subcomplex of (C,π), $p \circ i = \mathrm{id}_M$ and $i \circ p = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$.

Proposition 3.6.12. *P-filtered splitting homotopies and P-graded reductions are in bijective correspondence. Furthermore, perfect P-filtered splitting homotopy and minimal P-graded reductions are in bijective correspondence.*

Proof. The proof of the first result follows the proof of Proposition 3.6.3, except the maps are now filtered. For any reduction we have

$$i\partial'\rho = \partial(i\rho) = \partial(\mathrm{id}_C - \gamma\partial - \partial\gamma) = \partial - \partial\gamma\partial.$$

For a minimal reduction, we have

$$0 = i^{pp}\partial'^{pp}\rho^{pp} = \partial^{pp} - \partial^{pp}\gamma^{pp}\partial^{pp}.$$

Thus γ is perfect. Conversely, let γ be a perfect P-filtered splitting homotopy. The formula for the differential on $M = \mathrm{im}(\rho)$ is $\partial' = \partial - \partial\gamma\partial$. Since the maps ∂ and γ are P-filtered, we have

$$\partial'^{pp} = (\partial - \partial\gamma\partial)^{pp} = \partial^{pp} - \partial^{pp}\gamma^{pp}\partial^{pp} = 0.$$

□

Observe that in a minimal reduction $\operatorname{im}(\phi^{pp}) \subset \ker \partial^{pp}$ since $\partial^{pp}\phi^{pp} = \phi^{pp}\partial'^{pp} = 0$. Therefore the images $\phi^{pp}(M^p)$ are representatives of the homology $H_\bullet(C^p, \Delta^{pp})$. We may also show that strict P-graded chain complexes are minimal with respect to reductions. This mirrors Proposition 3.6.7. The point is that strict P-graded complexes are the graded analogue of minimal complexes.

Proposition 3.6.13. *Let (C, π) be strict. Any reduction*

$$\begin{tikzcd} (C,\pi) \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (M,\pi) \arrow[l, shift left, "\phi"] \end{tikzcd}$$

is minimal. Moreover (M, π) and (C, π) are P-filtered chain isomorphic, i.e., isomorphic in $\mathbf{Ch}(\mathbf{GrVect}(\mathsf{P}))$.

Proof. We have $\partial' = \partial'\psi\phi = \psi\partial\phi$. Thus $\partial'^{pp} = \psi^{pp}\partial^{pp}\phi^{pp} = 0$. Therefore (M, π) is strict. Since i and p are chain equivalences, invoking Proposition 3.3.29 shows that (M, π) and (C, π) are P-filtered chain isomorphic. ☐

For a tower of graded reductions, we have the following result, which is analogous to Proposition 3.6.9.

Proposition 3.6.14. *Given a tower of P-graded reductions*

$$\begin{tikzcd} (C,\pi) \arrow[loop above, "\gamma_0"] \arrow[r, shift left, "\psi_0"] & (M_0,\pi) \arrow[l, shift left, "\phi_0"] \arrow[loop above, "\gamma_1"] \arrow[r, shift left, "\psi_1"] & \ldots \arrow[l, shift left, "\phi_1"] \arrow[r, shift left, "\psi_{n-1}"] & (M_{n-1},\pi) \arrow[l, shift left, "\phi_{n-1}"] \arrow[loop above, "\gamma_n"] \arrow[r, shift left, "\psi_n"] & (M_n,\pi); \arrow[l, shift left, "\phi_n"] \end{tikzcd}$$

1. *there is a reduction*

$$\begin{tikzcd} (C,\pi) \arrow[loop above, "\Gamma"] \arrow[r, shift left, "\Psi_n"] & (M,\pi) \arrow[l, shift left, "\Phi_n"] \end{tikzcd}$$

   with maps given by the formulas

$$\Psi_m = \prod_{i=0}^{m} \psi_i \qquad \Phi_m = \prod_{i=0}^{m} \phi_i \qquad \Gamma = \gamma_0 + \sum_{i=0}^{n-1} \Phi_i \circ \gamma_{i+1} \circ \Psi_i.$$

2. *Γ is a P-filtered splitting homotopy and Γ is perfect if any γ_i is perfect.*

3.6.2 Filtered Reductions

Definition 3.6.15. An L-*filtered reduction* is a pair of L-filtered chain complexes and triple of L-filtered chain maps

$$\overset{\gamma}{\circlearrowleft}(C, f) \underset{\phi}{\overset{\psi}{\rightleftarrows}} (M, g)$$

where ϕ, ψ are chain maps and γ is a degree+1 map satisfying the identities:

1. $\psi\phi = \mathrm{id}_M$
2. $\phi\psi = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$
3. $\gamma^2 = \gamma\phi = \psi\gamma = 0.$

An L-filtered reduction is *minimal* if (M, g) is a strict L-filtered chain complex.

Definition 3.6.16. An L-*filtered splitting homotopy* is a degree 1 map $\gamma\colon C \to C$ such that $\gamma^2 = 0$ and $\gamma\partial\gamma = \gamma$. An L-filtered splitting homotopy $\gamma\colon C \to C$ for (C, f) is *perfect* if the induced map $\gamma^a\colon f(a)/f(\overleftarrow{a}) \to f(a)/f(\overleftarrow{a})$ is perfect for each $a \in \mathsf{J}(\mathsf{L})$.

Again, one may define $\rho = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$ and $M = \mathrm{im}(\rho)$. Then M is an L-filtered subcomplex of C and $\rho i = \mathrm{id}_M$ and $i\rho = \mathrm{id}_C - (\gamma\partial + \partial\gamma)$.

Corollary 3.6.17. *L-filtered splitting homotopies and L-filtered reductions are in bijective correspondence. Perfect filtered splitting homotopies are in bijective correspondence with minimal L-filtered reductions.*

Proof. The proof of the first statement follows the proof of Proposition 3.6.3, with the maps now filtered. Consider $g\colon \mathsf{L} \to \mathsf{Sub}(M, \partial^M)$ of the L-filtered reduction guaranteed by Corollary 3.6.17. If γ is perfect the differential ∂^M must obey $\partial^M(g(q)) = (\partial - \partial\gamma\partial)g(q) \subseteq g(\overleftarrow{q})$. Therefore (M, g) is strict. $\square$

We now show that strict L-filtered chain complexes are minimal with respect to reductions.

Proposition 3.6.18. *Let (C, f) be a strict L-filtered complex. Any filtered reduction*

$$\begin{tikzcd} (C,f) \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (M,g) \arrow[l, shift left, "\phi"] \end{tikzcd}$$

is minimal. Moreover (C, f) and (M, g) are filtered chain isomorphic.

Proof. We have $\partial' = \partial'\psi\phi = \psi\partial'$. For $a \in \mathsf{J}(\mathsf{L})$

$$\partial'(g(a)) = \psi\partial\phi(g(a)) \subseteq g(\overleftarrow{a}).$$

This implies that (M, g) is strict. Since (C, f) and (M, g) are both strict L-filtered chain complexes, they are L-filtered chain isomorphic by Proposition 3.4.17. $\square$

For a tower of L-filtered reductions there is an result analogous to Propositions 3.6.9 and 3.6.14.

Proposition 3.6.19. *For a tower of L-filtered reductions*

$$\begin{tikzcd} (C,f) \arrow[loop above, "\gamma_0"] \arrow[r, shift left, "\psi_0"] & (M_0,g) \arrow[l, shift left, "\phi_0"] \arrow[loop above, "\gamma_1"] \arrow[r, shift left, "\psi_1"] & \cdots \arrow[l, shift left, "\phi_1"] \arrow[r, shift left, "\psi_{n-1}"] & (M_{n-1},g) \arrow[l, shift left, "\phi_{n-1}"] \arrow[loop above, "\gamma_n"] \arrow[r, shift left, "\psi_n"] & (M_n,g); \arrow[l, shift left, "\phi_n"] \end{tikzcd}$$

1. *there is a reduction*

$$\begin{tikzcd} (C,f) \arrow[loop above, "\Gamma"] \arrow[r, shift left, "\Psi_n"] & (M_n,g_n) \arrow[l, shift left, "\Phi_n"] \end{tikzcd}$$

   with maps given by the formulas

$$\Psi_m = \prod_{i=0}^{m} \psi_i \qquad \Phi_m = \prod_{i=0}^{m} \phi_i \qquad \Gamma = \gamma_0 + \sum_{i=0}^{n-1} \Phi_i \circ \gamma_{i+1} \circ \Psi_i.$$

2. *Γ is a splitting homotopy and Γ is perfect if any γ_i is perfect.*

3.7 Connection Matrix Algorithm

In this section we introduce the algorithm for computing Conley complexes and connection matrices. The algorithm is based on (graded) Morse theory, which is described in Section 3.7.3. It is formalized via the framework of reductions developed in Section 3.6.

In Section 3.7.1 we recall the Morse theoretic algorithms of [23]. The exposition relies on the discrete Morse theory reviewed in Section 2.8. In Section 3.7.2 we demonstrate how to compute the homology of a chain complex using Morse theory and reductions. This is given via Algorithm 3.7.2. Section 3.7.4 covers how to compute a connection matrix, via Algorithm 3.7.8. The computation of a connection matrix ends up being very similar to computing homology, only adapted to the graded setting.

3.7.1 Morse Theoretic Algorithms

Our algorithm relies on [23, Algorithm 3.6] and [23, Algorithm 3.12], which are reproduced below, respectively, as the algorithms as MATCHING and GAMMA. In particular, Lemma 3.7.1, which relies on MATCHING, is used to to verify the correctness of Algorithms 3.7.2 and 3.7.8. First, recall the notion of a coreduction pair and free cell, from Definition 2.5.15, and that of acyclic partial matching, from Definition 2.8.1.

function MATCHING($\mathcal{X}$)

  $\mathcal{X}' \leftarrow \mathcal{X}$

  while $\mathcal{X}'$ is not empty **do**

    while $\mathcal{X}'$ admits a coreduction pair (ξ, ξ') **do**

      Excise (ξ, ξ') from $\mathcal{X}'$

      $K \leftarrow \xi$, $Q \leftarrow \xi'$

      $w(\xi') := \xi$

    end while

    while $\mathcal{X}'$ does not admit a coreduction pair **do**

      Excise a free cell ξ from $\mathcal{X}'$

      $A \leftarrow \xi$

    end while

  end while

  return $(A, w\colon Q \to K)$

end function

function GAMMA($\xi_{in}, w\colon Q \to K$)

  $\xi \leftarrow \xi_{in}$

  $c \leftarrow 0$

  while $\xi \notin C(A) \oplus C(K)$ **do**

Choose a $\leq$-maximal $\xi' \in Q$ with $\kappa(\xi, \xi') \neq 0$

$\xi'' \leftarrow w(\xi')$

$c \leftarrow c + \xi''$

$\xi \leftarrow \xi + \partial\xi''$

end while

return c

end function

The proof of correctness of our algorithms depend upon the following lemma.

Lemma 3.7.1. *Let $\mathcal{X}$ be a cell complex. If (A, w) is an acyclic partial matching on $\mathcal{X}$ obtained from* MATCHING$(\mathcal{X})$ *such that $A = \mathcal{X}$ then $(C_\bullet(A), \partial^A) = (C_\bullet(\mathcal{X}), \partial^{\mathcal{X}})$ is a minimal complex.*

Proof. Let $\xi \in \mathcal{X} = A$. We wish to show that $\partial(\xi) = 0$. Since $A = \mathcal{X}$ there are no coreduction pairs in the execution of the algorithm. This implies of the two secondary **while** loops in Algorithm 3.6, only the second **while** has executed. This **while** loop has iterated $n = |\mathcal{X}|$ times and each iteration adds a cell ξ to the collection of critical cells A. We may therefore regard A as a stack and label each $\xi \in A$ with the integer giving the particular iteration of the **while** loop that added ξ to A. Denote this labeling $\mu\colon A \to \mathbb{N}$. Now set $n = \mu(\xi)$ and let $U = \mu^{-1}[0, n)$. From Algorithm 3.6 we must have that ξ is a free cell in $\mathcal{X} \setminus U$. Therefore if $\kappa(\xi, \xi') \neq 0$ for some $\xi' \in \mathcal{X}$ then $\xi' \in U$. Suppose that $\kappa(\xi, \xi') \neq 0$ for some $\xi' \in U$. Let $m = \mu(\xi')$ and $U' = \mu^{-1}[0, m)$. We must have that (ξ, ξ') is a coreduction pair in $\mathcal{X} \setminus U'$. This is a contradiction of the execution of the algorithm. $\square$

3.7.2 Homology Algorithm

We first give an algorithm for computing the homology of a complex $\mathcal{X}$ based on discrete Morse theory. This will provide an intuition and the basis for the Algorithm 3.7.8, CONNECTIONMATRIX.

Algorithm 3.7.2.

function HOMOLOGY$(\mathcal{X}_{in}, \partial_{in})$

$A \leftarrow \mathcal{X}_{in}, \Delta \leftarrow \partial_{in}$

do

  $\mathcal{X} \leftarrow A, \partial \leftarrow \Delta$

  $(A, w\colon Q \to K) \leftarrow \textsc{Matching}(\mathcal{X})$

  for $\xi \in A$ **do**

    compute and store $\Delta(\xi)$ using $\textsc{Gamma}(\xi, w)$

  end for

while $|A| < |\mathcal{X}|$

return A

end function

Theorem 3.7.3. *Given a cell complex $\mathcal{X}$, Algorithm 3.7.2 (with input $\mathcal{X}$ and ∂) halts and outputs the homology of $\mathcal{X}$.*

Proof. The fact that $\mathcal{X}_{in}$ is finite, together with the fact that Matching halts [23], implies that Homology halts. Finally, if the algorithm terminates with $A_\infty = \textsc{Homology}(\mathcal{X}_{in})$ then $C(A_\infty)$ is a minimal chain complex by Lemma 3.7.1.

It remains to prove that $C(A_\infty) \cong H(\mathcal{X})$. In any iteration of the **do** loop, there are chain equivalences $\psi\colon C(\mathcal{X}) \to C(A)$ and $\phi\colon C(A) \to C(\mathcal{X})$, which are as defined in Eq. (2.5) of Section 2.8, using $\gamma(\cdot) = \textsc{Gamma}(\cdot, w)$. The pair of complexes $C(\mathcal{X})$, $C(A)$ and the triple maps ϕ, ψ, γ fit into a reduction via Example 3.6.4. Therefore an execution of the entire the **do-while** loop is associated to a tower of reductions:

$$\begin{tikzcd} C(\mathcal{X}_{in}) \arrow[loop above, "\gamma_0"] \arrow[r, shift left] & \cdots \arrow[l, shift left] \arrow[r, shift left] & C_\bullet(\mathcal{X}) \arrow[l, shift left] \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (C_\bullet(A), \partial^A) \arrow[l, shift left, "\phi"] \arrow[loop above] \arrow[r, shift left] & \cdots \arrow[l, shift left] \arrow[r, shift left] & C(A_\infty). \arrow[l, shift left] \end{tikzcd}$$

Thus the output $C(A_\infty)$ is isomorphic to the homology $H(\mathcal{X}_{in})$. ∎

Example 3.7.4. In this example we give some flavor of the concepts behind Algorithm 3.7.2 (Homology). Consider the cubical complex $\mathcal{K}$ given in Figure 3.1(a) below, which consists of four 2-cells, 14 edges and 9 vertices. We work over the field $\mathbb{Z}_2$. Therefore we have

$$C_2(\mathcal{K}) = \mathbb{Z}_2^4 \qquad C_1(\mathcal{K}) = \mathbb{Z}_2^{14} \qquad C_0(\mathcal{K}) = \mathbb{Z}_2^9.$$

The complex $\mathcal{K}$ is open on the right in order to simplify the Morse theory. For the sake of an example, we want to illustrate how multiple rounds of Morse theory work in practice. Unfortunately, the algorithm Homology as stated is too effective in this example as the Matching subroutine simplifies to a minimal complex in only one round of Morse theory. Instead, we may substitute Matching with the set of implicit cubical matchings proposed in Section 5.3; every coordinate direction provides a round of Morse theory.

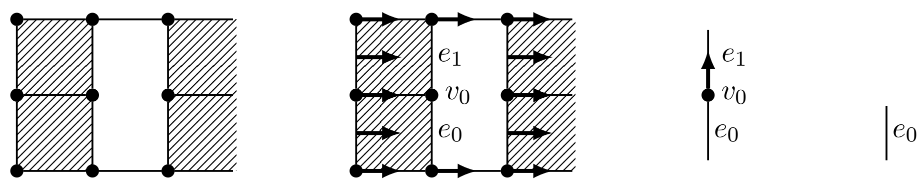

(a) (b) (c) (d)

Figure 3.1: (a) Cubical complex $\mathcal{K}$. (b) First Pairing. (c) Second Pairing. (d) Minimal complex, i.e., the homology of $\mathcal{K}$. A single 1-cell e_0 remains with $\partial(e_0) = 0$.

The algorithm begins with the executing the first iteration of the **while** loop which computes an acyclic partial matching on $\mathcal{K}$ by attempting to pair all cells right. This furnishes an acyclic partial matching (A_0, w_0) (see Section 5.3). This is visualized in Figure 3.1(b). As in much of the literature on discrete Morse theory, e.g., [14, 22], a pair $\xi \in Q$ and $\xi' \in K$ with $w(\xi) = \xi'$ is visualized with directed edge $\xi' \to \xi$; see Figure 3.1(b). The directed edges in Figure 3.1(b) may also be thought of as a graphical representation of the degree 1 map V (see Section 2.8). The cells e_0, e_1 and v_0 do not have right coboundaries (see Section 5.2.1) and are therefore critical, i.e., $A_0 = \{e_0, e_1, v_0\}$. The set of cells A_0 is the basis for a chain complex $(C(A_0), \Delta)$ where

$$C_1(A_0) = \mathbb{Z}_2\langle e_0\rangle \oplus \mathbb{Z}_2\langle e_1\rangle, \qquad C_0(A_0) = \mathbb{Z}_2\langle v_0\rangle, \qquad \Delta_1 = \begin{array}{cc} & \begin{array}{cc} e_0 & e_1 \end{array} \\ v_0 & \begin{pmatrix} 1 & 1 \end{pmatrix} \end{array}.$$

The second iteration of the **while** loop attempts to pair remaining cells, i.e., the cells in A_1, upwards. This furnishes an acyclic partial matching (A_1, w_1) on A_0, visualized in Figure 3.1(c). The cells v_0 and e_1 are paired and $A_2 = \{e_0\}$. Moreover, the set of cells A_1 is a basis for the chain complex $(C(A_1), \Delta)$ where

$$C_1(A_1) = \mathbb{Z}_2\langle e_0\rangle \qquad \text{and} \qquad \Delta = 0.$$

In the final iteration of the **while** there are no coreduction pairs and $A_2 = A_1$ is returned. These two rounds of Morse theory give a tower of reductions and the algorithm terminates with $A_\infty = \textsc{Homology}(\mathcal{K}, \partial) = A_1$. The chain complexes $C(\mathcal{X}), C(A_0)$ and $C(A_1)$, together with the maps $\{\phi_i, \psi_i, \gamma_i\}$ fit into the tower of reductions below.

$$\begin{tikzcd} C(\mathcal{K}) \arrow[loop above, "\gamma_0"] \arrow[r, shift left, "\psi_0"] & C(A_0) \arrow[l, shift left, "\phi_0"] \arrow[loop above, "\gamma_1"] \arrow[r, shift left, "\psi_1"] & C(A_1) \arrow[l, shift left, "\phi_1"] \end{tikzcd}$$

3.7.3 Graded Morse Theory

In this section, we review a graded version of discrete Morse theory. Consider a P-graded cell complex $(\mathcal{X}, \nu)$. Recall that the underlying set $\mathcal{X}$ decomposes as $\mathcal{X} = \bigsqcup_{p \in \mathsf{P}} \mathcal{X}^p$ where $\mathcal{X}^p = \nu^{-1}(p)$.

Definition 3.7.5. Let $(\mathcal{X}, \nu)$ be a P-graded cell complex and let $(A, w\colon Q \to K)$ be an acyclic partial matching on $\mathcal{X}$. We say that (A, w) is P-*graded*, or simply *graded*, if it satisfies the property that $w(\xi) = \xi'$ only if $\xi, \xi' \in \mathcal{X}^p$ for some $p \in \mathsf{P}$. That is, matchings may only occur in the same fiber of the grading.

The idea of graded matchings can be found many places in the literature, for instance, see [38] and [29, Patchwork Theorem]. Recall from Definition 3.3.18 that a P-graded cell complex $(\mathcal{X}, \nu)$ has an associated P-graded chain complex $(C(\mathcal{X}), \pi^\nu)$.

Proposition 3.7.6. *Let $(\mathcal{X}, \nu)$ be a P-graded cell complex and $(A, w : Q \to K)$ a graded acyclic partial matching. Let $A^p = A \cap \mathcal{X}^p$. Then*

1. *$(C(A), \partial^A, \pi)$ is a P-graded chain complex where the projections $\pi = \{\pi_n^p\}_{n \in \mathbb{Z}, p \in \mathsf{P}}$ are given by*
$$\pi_j^p\colon C_j(A) \to C_j(A^p). \tag{3.19}$$
2. *The maps ϕ, ψ of Eqn.* (2.5) *and γ of Eqn.* (2.4) *fit into a P-graded reduction*
$$\begin{tikzcd} (C(\mathcal{X}), \pi^\nu) \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (C(A), \pi). \arrow[l, shift left, "\phi"] \end{tikzcd} \tag{3.20}$$

Proof. It follows from Proposition 2.8.3 that $(C(A), \partial^A)$ is a chain complex. We must show that if $\partial_A^{pq} \neq 0$ then $p \leq q$. By Proposition 2.8.2 there is a unique splitting homotopy $\gamma\colon C(\mathcal{X}) \to C(\mathcal{X})$ associated to the matching (A, w). The fact that (A, w) is graded implies that V, as defined in Eqn. (2.3) of Section 2.8, is P-filtered. From the definition of γ given in (2.4) a routine verification shows that γ is P-filtered. Therefore by Proposition 3.6.12 there is an associated reduction

$$\begin{tikzcd} (C(\mathcal{X}), \pi^\nu) \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (C(A), \pi). \arrow[l, shift left, "\phi"] \end{tikzcd}$$

Let $p \in \mathsf{P}$. Consider (A^p, w^p) the matching restricted to the fiber $\mathcal{X}^p = \nu^{-1}(p)$. We have

$$A^p = A \cap \mathcal{X}^p \qquad\qquad w^p\colon Q \cap \mathcal{X}^p \to K \cap \mathcal{X}^p.$$

It follows that (A^p, w^p) is an acyclic partial matching on the fiber $\mathcal{X}^p$. Proposition 2.8.2 implies that there is a unique splitting homotopy $\gamma^p\colon C(\mathcal{X}^p) \to C(\mathcal{X}^p)$. In particular, $\gamma^{pp} = \gamma^p$. $\square$

Example 3.7.7. Consider the graded complex $(\mathcal{X}, \nu)$ of Example 1.2.1. Let

$$A := \{v_0, v_2, e_0\} \qquad Q := \{v_1\} \qquad K := \{e_1\} \qquad w(v_1) = e_1$$

This is depicted in Figure 3.2.

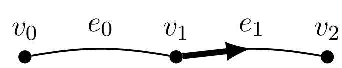

Figure 3.2: Graded matching (A, w) on $\mathcal{X}$. The pairing $w(v_1) = e_1$ is visualized with an arrow $v_1 \to e_1$. The set $A = \{v_0, e_0, v_2\}$ are the critical cells.

This is a acyclic partial matching (A, w) on $\mathcal{X}$. It is straightforward that (A, w) is graded as $v_1, e_1 \in \nu^{-1}(q)$. The maps of the associated P-graded reduction are precisely the ones described in Example 3.3.24.

3.7.4 Connection Matrix Algorithm

We can now state the algorithm for computing a connection matrix, which relies on Matching and Gamma of Section 3.7.1.

Algorithm 3.7.8.

function ConnectionMatrix($\mathcal{X}_{in}, \nu_{in}, \partial_{in}$)

  $A \leftarrow \mathcal{X}_{in}, \Delta \leftarrow \partial_{in}, \mu \leftarrow \nu_{in}$

  do

    $\mathcal{X} \leftarrow A, \partial \leftarrow \Delta, \nu \leftarrow \mu$

    for $p \in \mathsf{P}$ **do**

      $(A^p, \omega^p \colon Q^p \to K^p) \leftarrow$ Matching($\mathcal{X}^p$)

    end for

    $(A, w) \leftarrow (\bigsqcup_{p\in\mathsf{P}} A^p,\ \bigsqcup_{p\in\mathsf{P}} w^p)$

    for $\xi \in A$ **do**

      compute and store $\Delta(\xi)$ using Gamma(ξ, w)

    end for

    $\mu \leftarrow \nu|_A$

  while $|A| < |\mathcal{X}|$

  return (A, Δ, μ)

end function

Theorem 3.7.9. *Let $(\mathcal{X}, \nu)$ be a P-graded cell complex. Algorithm 3.7.8 (with input $(\mathcal{X}, \nu)$ and ∂) halts. Moreover, the returned data (A_∞, Δ, π) has the property that $(C(A_\infty), \Delta, \pi)$ is a Conley complex for $(C(\mathcal{X}), \pi^\nu)$.*

Proof. Since $\mathcal{X}_{in}$ is finite and Matching halts, it follows that ConnectionMatrix halts. Let $(A_\infty, \Delta_\infty, \mu_\infty) =$ ConnectionMatrix($\mathcal{X}_{in}, \nu, \partial$). It follows from Proposition 3.7.6 that $(C(A_\infty), \Delta, \pi)$, where π is given by Eqn. (3.19), is P-graded. It follows from Lemma 3.7.1 that for each $p \in \mathsf{P}$ the fiber A^p_∞ is minimal. This implies that the P-graded chain complex $C(A_\infty)$ is strict.

It remains to show that $(C(A_\infty), \Delta, \pi)$ is a Conley complex. In any iteration of the **do** loop, it follows from Proposition 3.7.6 that there are P-filtered chain equivalences $\psi \colon C(\mathcal{X}) \to C(A)$ and $\phi \colon C(A) \to C(\mathcal{X})$, which are as defined in (2.5), using $\gamma(\cdot) =$ Gamma($\cdot, w$). The pair of complexes $(C(\mathcal{X}), \pi^\nu)$, $(C(A), \pi^\mu)$ and the triple maps ϕ, ψ, γ

fit into the P-graded reduction of (3.20). Therefore an execution of the entire the **do-while** loop is associated to a tower of reductions:

$$
\begin{tikzcd}
(C(\mathcal{X}_{in}),\pi_{in}) \arrow[loop above, "\gamma_0"] \arrow[r, shift left] & \dots \arrow[l, shift left] \arrow[r, shift left] & (C(\mathcal{X}),\pi^{\nu}) \arrow[l, shift left] \arrow[loop above, "\gamma"] \arrow[r, shift left, "\psi"] & (C(A),\pi^{\mu}) \arrow[l, shift left, "\phi"] \arrow[loop above] \arrow[r, shift left] & \dots \arrow[l, shift left] \arrow[r, shift left] & (C(A_{\infty}),\pi). \arrow[l, shift left]
\end{tikzcd}
\tag{3.21}
$$

Thus the output $(C(A_\infty),\Delta)$ is a Conley complex. □

Example 3.7.10. We given an example of the algorithm ConnectionMatrix. Let $\mathcal{X}$ be the cubical complex in Figure 3.3(a) and let $\mathcal{K}$ be a the cubical complex from Example 3.7.4. We again work over the field $\mathbb{Z}_2$. The cubical complex $\mathcal{X}$ consists of $\mathcal{K}$ together with the 2-cells $\{\xi_0,\xi_1\}$ and the 1-cell e_2. The 2-cells in $\mathcal{X}\setminus\mathcal{K}$ are shaded, while the 2-cells in $\mathcal{K}$ are drawn with hatching. Let $\mathsf{Q}=\{0,1\}$ be the poset with order $0\leq 1$. There is a Q-graded cell complex where $(\mathcal{X},\nu)$ and $\nu\colon\mathcal{X}\to\mathsf{Q}$ is given via

$$
\nu(x)=\begin{cases}0 & x\in\mathcal{K}\\ 1 & x\in\mathcal{X}\setminus\mathcal{K}.\end{cases}
$$

Once again, using Matching would be too effective on this example to illustrate multiple rounds. We proceed as before and use the implicit cubical matchings of Section 5.3. Again, every coordinate direction gives a new round of Morse theory.

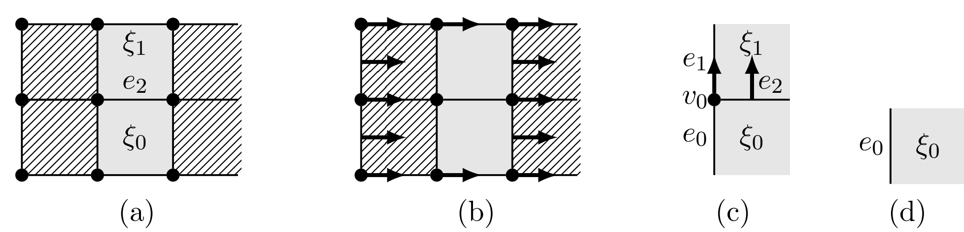

Figure 3.3: (a) Graded Cubical Complex. (b) First Graded Pairing. (c) Second Graded Pairing. (d) Conley complex. A 2-cell ξ_0 and a 1-cell e_0 remain with $\partial(\xi_0)=e_0$.

The algorithm starts by computing a graded acyclic partial matching on $\mathcal{K}$, attempting to pair all cells to the right (within their fiber). The cells e_0,e_1,v_0 have right coboundaries ξ_0,ξ_1,e_2 respectively. However, these do lie in the same fiber as $e_0,e_1,v_0\in\mathcal{X}^0$ and $\xi_0,\xi_1,e_2\in\mathcal{X}^1$. Therefore these cells cannot be paired and $A_0=\{e_0,e_1,v_0,\xi_0,\xi_1,e_2\}$. The second round of Morse theory attempts to pair the remaining cells up (within their fiber). In this case, $w(v_0) = e_1$ and $w(e_2) = \xi_1$ and $A_1 = \{e_0, \xi_0\}$. These two rounds of graded Morse theory give a tower of graded reductions. The returned data (A_∞, Δ, μ) form the strict P-graded chain complex

$$C_2(A_\infty) = \mathbb{Z}_2 \qquad C_1(A_\infty) = \mathbb{Z}_2 \qquad \Delta_2^{01} = \begin{pmatrix} 1 \end{pmatrix}.$$

Remark 3.7.11. Algorithm 3.7.8 may be refined by returning either the entire tower of reductions or the reduction defined by the compositions as in Theorem 3.6.14. Returning these data allow one to lift generators back in the original chain complex.

Remark 3.7.12. Our implementation of this algorithm is available at [1]. Also available is a Jupyter notebook for the application of the algorithm to a Morse theory on braids; see Section 4.7. The application to braids falls within the scope of a larger project, namely developing the ability to compute connection matrices for transversality models; see Section 4.3. More results of the algorithm, along with some timing data, are covered Section 5.4.

3.8 Categorical Connection Matrix Theory

We are now in a position to return to Diagram (3.11) of Section 3.4.6 and discuss the categorical setup for connection matrix theory.

Theorem 3.8.1. *The inclusion functor* $i\colon \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ *is an equivalence of categories.*

Proof. Since the subcategory $\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$ is full, the inclusion functor i is full and faithful. We show that i is essentially surjective. Let (C, π) be a P-graded chain complex and let $\mathcal{X}$ be a basis for $C_\bullet$. Define $\dim(\xi) = n$ where $\xi \in C_n$. Define $\kappa(\xi, \xi')$ as the appropriate coefficient of ∂, i.e.,

$$\partial(\xi) = \sum_{\xi' \in \mathcal{X}} \kappa(\xi, \xi')\xi',$$

and define the partial order $\leq$ via

$$\xi' \leq \xi \iff \kappa(\xi, \xi') \neq 0.$$

Then $\mathcal{X} = (\mathcal{X}, \kappa, \dim, \leq)$ is a P-graded cell complex. Consider the strict P-graded chain complex $(C(A), \Delta)$ where $(A, \Delta, \mu) = \textsc{ConnectionMatrix}(\mathcal{X}, \nu, \partial)$. It follows from Theorem 3.7.9 that this is an object of $\mathbf{K}_s(\mathbf{GrVect}(\mathsf{P})) \subset \mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$ and that it is a Conley complex for (C, π). Therefore $(C(A), \Delta)$ is isomorphic to (C, π) in $\mathbf{K}(\mathbf{GrVect}(\mathsf{P}))$. $\square$

Corollary 3.8.2. *The inclusion functor* $\mathbf{K}_s(\mathbf{FVect}(\mathsf{L})) \subset \mathbf{K}(\mathbf{FVect}(\mathsf{L}))$ *is an equivalence of categories.*

Proof. This follows from an examination of the bottom square of (3.11): three of the functors are categorical equivalences. $\square$

Corollary 3.8.3. *Let* $\mathsf{L} = \mathsf{O}(\mathsf{P})$*. There exists inverse functors, which we call Conley functors,*

1. $F\colon \mathbf{K}(\mathbf{GrVect}(\mathsf{P})) \to \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$*, and*
2. $G\colon \mathbf{K}(\mathbf{FVect}(\mathsf{L})) \to \mathbf{K}_s(\mathbf{GrVect}(\mathsf{P}))$*,*

which take a P*-graded chain complex or* L*-filtered chain complex to its Conley complex.*

Remark 3.8.4. Corollary 3.8.3 provides a functorial framework for connection matrix theory. Algorithm 3.7.8 (ConnectionMatrix) computes the functor F on objects.

3.9 Relationship to Persistent Homology

Persistent homology is a quantitative method within applied algebraic topology and the most popular tool of topological data analysis. We give a brief outline, and refer the reader to [11, 39] and their references within for further details. In this section we show that given an L-filtered chain complex, one can recover its persistent homology using a Conley complex and connection matrix. Persistent homology may be viewed as a family of functors, parameterized by pairs of elements $a, b \in \mathsf{L}$ with $a \leq b$:

$$\{PH_\bullet^{a,b} \colon \mathbf{Ch}(\mathbf{FVect}(\mathsf{L})) \to \mathbf{Ch}_0(\mathbf{Vect})\}_{a \leq b}.$$

Remark 3.9.1. To be consistent with the literature of persistent homology, our notation is $PH_\bullet^{a,b}$ where $a \leq b$. This is in contrast to our 'matrix' notation that runs through the rest of the paper.

Let $a, b \in \mathsf{L}$ with $a \leq b$. $PH_\bullet^{a,b}(-)$ is defined on objects as follows. Let (C, f) be an L-filtered chain complex (see Section 3.4.2). There is an inclusion of subcomplexes

$$\iota^{a,b} \colon f(a) \hookrightarrow f(b). \tag{3.22}$$

Recall from Section 2.3 that we view homology as a functor $H_\bullet \colon \mathbf{Ch} \to \mathbf{Ch}_0$. Applying $H_\bullet$ to Eqn. (3.22) yields a map $H_\bullet(\iota^{a,b}) \colon H_\bullet(f(a)) \to H_\bullet(f(b))$. Then

$$PH_\bullet^{a,b}\big((C, f)\big) := \operatorname{im} H_\bullet(\iota^{a,b}) \in \mathbf{Ch}_0.$$

From this setup we can recover the standard persistence: for $j \in \mathbb{Z}$ the *j-th persistent homology group of $a \leq b$* is the vector space

$$PH_j^{a,b}\big((C, f)\big) := \operatorname{im} H_j(\iota^{a,b}).$$

The *j-th persistent Betti numbers* are the integers

$$\beta_j^{a,b} = \dim \operatorname{im}(H(\iota_j^{a,b})).$$

$PH_\bullet^{a,b}(-)$ is defined on morphisms as follows. Let $\phi \colon (C, f) \to (C', f')$ be an L-filtered chain map. Since ϕ is L-filtered, ϕ restricts to chain maps $\phi^a \colon f(a) \to f'(a)$ and $\phi^b \colon f(b) \to f'(b)$, which fit into the following commutative diagram.

$$\begin{tikzcd}
H_\bullet f(a) \arrow[r, "H_\bullet(\iota^{a,b})"] \arrow[d, "H_\bullet(\phi^a)"'] & H_\bullet f(b) \arrow[d, "H_\bullet(\phi^b)"] \\
H_\bullet f'(a) \arrow[r, "H_\bullet(\iota'^{a,b})"] & H_\bullet f'(b)
\end{tikzcd}$$

As the diagram commutes, $H_\bullet(\phi^b)$ restricts to a map $H_\bullet(\phi^b) \colon \operatorname{im} H_\bullet(\iota^{a,b}) \to \operatorname{im} H_\bullet(\iota'^{a,b})$, and

$$PH_\bullet^{a,b}(\phi) := H_\bullet(\phi^b) \colon \operatorname{im} H_\bullet(\iota^{a,b}) \to \operatorname{im} H_\bullet(\iota'^{a,b}).$$

Proposition 3.9.2. *$PH_\bullet^{a,b}$ sends L-filtered chain equivalences to isomorphisms in* $\mathbf{Ch}_0$*.*

Proof. Let $a, b \in \mathsf{L}$ with $a \leq b$. Let $\phi\colon (C, f) \to (C', f')$ be an L-filtered chain equivalence. Since ϕ is L-filtered, ϕ^a and ϕ^b are chain equivalences. Proposition 2.3.14 implies $H(\phi^a)$ and $H(\phi^b)$ are isomorphisms. Thus $PH_\bullet^{a,b}(\phi)$ is an isomorphism. $\square$

Theorem 3.9.3. *Let (C, f) be a L-filtered chain complex. Let (C', π) be a Conley complex for (C, f). Then for all $j \in \mathbb{Z}$ and $a \leq b$ in L*

$$PH_\bullet^{a,b}(\mathfrak{L}(C', \pi)) \cong PH_\bullet^{a,b}((C, f)).$$

Proof. It follows from Proposition 3.9.2 and Proposition 2.2.10 that $PH_\bullet^{a,b}$ factors as $PH_{\mathbf{K}}^{a,b} \circ q$, giving the following commutative diagram.

$$\begin{tikzcd}
\mathbf{Ch}(\mathbf{FVect}(\mathsf{L})) \arrow[rr, "PH_\bullet^{a,b}"] \arrow[dr, "q"'] & & \mathbf{Ch}_0(\mathbf{Vect}) \\
& \mathbf{K}(\mathbf{FVect}(\mathsf{L})) \arrow[ur, "PH_{\mathbf{K}}^{a,b}"'] &
\end{tikzcd}$$

Since (C', π) is a Conley complex for (C, f), by definition we have that

$$q((C, f)) \cong q(\mathfrak{L}(C', \pi).$$

It follows that

$$PH_\bullet^{a,b}(C, f) = PH_{\mathbf{K}}^{a,b} \circ q(C, f) \cong PH_{\mathbf{K}}^{a,b} \circ q \circ \mathfrak{L}(C', \pi) = PH_\bullet^{a,b} \circ \mathfrak{L}(C', \pi). \quad \square$$

Remark 3.9.4. As a corollary, all computational tools that tabulate the persistent homology groups, such as the persistence diagrams and barcodes (see [11, 39]), can be computed from the Conley complex.

Let $\mathcal{X}$ be a finite cell complex and L a finite distributive lattice. Suppose that $\{\mathcal{X}^a \subset \mathcal{X} \mid a \in \mathsf{L}\}$ be an isomorphic lattice of subcomplexes. Defining f by taking $a \in \mathsf{L}$ to $\mathcal{X}^a \subset \mathcal{X}$ yields a lattice morphism $f\colon \mathsf{L} \to \mathsf{Sub}_{Cl}(\mathcal{X})$. Therefore $(\mathcal{X}, f)$ is an L filtered cell complex.

Definition 3.9.5. Let L be a finite distributive lattice and $(\mathcal{X}, f)$ be an L-filtered cell complex. The *persistent homology of* $(\mathcal{X}, f)$ is defined to be

$$PH_\bullet^{a,b}(\mathcal{X}, f) := PH_\bullet^{a,b} \circ \mathcal{L}(\mathcal{X}, f).$$

Theorem 3.9.6. *Let $\mathcal{X}$ be a finite cell complex with associated chain complex $(C(\mathcal{X}), \partial)$. Let L be a finite distributive lattice. Let $\{\mathcal{X}^a \subset \mathcal{X} \mid a \in \mathsf{L}\}$ be an isomorphic lattice of subcomplexes with operations $\cap$ and $\cup$ and minimal and maximal elements $\varnothing$ and $\mathcal{X}$, respectively. Let*

$$\Delta\colon \bigoplus_{a\in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}) \to \bigoplus_{a\in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}})$$

be an associated connection matrix. Define

$$\mathcal{M}^a := \bigoplus_{\substack{b\leq a,\\ b\in \mathsf{J}(\mathsf{L})}} H_\bullet(\mathcal{X}^b, \mathcal{X}^{\overleftarrow{b}}).$$

The persistent homology groups of $\{\mathcal{M}^a\}_{a\in\mathsf{L}}$ and $\{\mathcal{X}^a\}_{a\in\mathsf{L}}$ are isomorphic.

Proof. Define $f\colon \mathsf{L} \to \mathsf{Sub}_{Cl}(\mathcal{X})$ as the lattice morphism $f(a) = \mathcal{X}^a$ for $a \in \mathsf{L}$. Then $(\mathcal{X}, f)$ is an L-filtered subcomplex. Set

$$M_\bullet = \bigoplus_{a\in \mathsf{J}(\mathsf{L})} H_\bullet(\mathcal{X}^a, \mathcal{X}^{\overleftarrow{a}}).$$

By hypothesis $(M, \pi) = (M, \Delta, \pi)$ is a Conley complex for $\mathcal{L}(\mathcal{X}, f)$, so $\mathfrak{L}(M, \pi) \cong \mathcal{L}(\mathcal{X}, f)$ in $\mathbf{K}(\mathbf{FVect}(\mathsf{L}))$ by definition. Let $a, b \in \mathsf{L}$ with $a \leq b$. We have that

$$PH_\bullet^{a,b}(M, g) = PH_\bullet^{a,b}(\mathfrak{L}(M, \Delta, \pi)) \cong PH_\bullet^{a,b}(\mathcal{L}(\mathcal{X}, f)),$$

where the isomorphism follows from Proposition 3.4.14. $\square$

Chapter 4

Computational Conley-Morse Homology for Flows

4.1 Overview

In this chapter we develop a setting in which the computational connection matrix theory may be applied to computational dynamics. We intertwine computational Conley theory [2, 6, 25, 26, 27, 28] and the computational connection matrix theory as developed in Chapter 3. The contributions can be summarized as follows.

- We show how one moves from Conley-theoretic approximations of continuous dynamics to the appropriate graded or filtered cell complexes on which to apply the computational connection matrix theory. This is the content of Theorem 4.2.4.
- We develop a specific setting for applying the theory. We call these *transversality models*.
- Using the theory of transversality models, we make applications to some classical examples as well as to a Morse theory on braids.

4.2 Computational Conley-Morse Homology for Flows

Moving from continuous dynamics to a combinatorial dynamics is often called *combinatorialization*. There are two aspects to this:

- Combinatorialization of the phase space (i.e., the underlying topological space).
- Combinatorialization of continuous dynamics.

4.2.1 Combinatorialization of Topological Space

Recall that the starting point for the combinatorialization of a space is a choice of pure polyhedral complex $\mathcal{X}$ in $\mathbb{R}^n$ (see Section 2.7). Given a top-cell $\xi \in \mathcal{X}^+$ (see Definition 2.5.14), the realization $|\xi|$ is an open convex set of $|\mathcal{X}|$ and $\mathrm{cl}\,|\xi| = |\,\mathrm{cl}(\xi)|$ is a closed convex set. Recall from Definition 2.4.14 that a Boolean algebra is a lattice where every element has a complement. The collection of regular closed sets of $|\mathcal{X}|$ is a Boolean algebra denoted $\mathscr{R}(|\mathcal{X}|)$, see Definition 2.9.2.

Definition 4.2.1. Let $\mathcal{X}$ be a polyhedral complex in $\mathbb{R}^n$. The *associated Boolean algebra of regular closed sets*, denoted by X, is the finite sublattice of $\mathscr{R}(|\mathcal{X}|)$ which has atoms $\{\mathrm{cl}\,|\xi| \mid \xi \in \mathcal{X}^+\}$.

We are now in a situation where we have combinatorial objects, viz., $\mathsf{Set}(\mathcal{X}^+)$ (see Definition 2.4.15); topological objects, viz., X; and algebraic objects, viz., $\mathsf{Sub}_{Cl}(\mathcal{X})$ (see Definition 2.5.11). We need to introduce the proper morphisms which relate these objects, fitting into the following schematic.

$$\begin{tikzcd}
\text{Combinatorics} \arrow[rr, leftrightarrow] \arrow[dr, leftrightarrow] & & \text{Topology} \arrow[dl, leftrightarrow] \\
& \text{Algebra} &
\end{tikzcd}$$

There is a map $k\colon \mathsf{Set}(\mathcal{X}^+) \to \mathsf{X}$ which sends $\{\xi\} \in \mathsf{Set}(\mathcal{X}^+)$ to

$$k(\{\xi\}) = \mathrm{cl}\,|\{\xi\}| \in \mathsf{X}.$$

It follows from the fact that $\mathrm{cl}\,|\cdot| = |\,\mathrm{cl}(\cdot)|$ (see Definition 2.7.1) that k is a lattice isomomorphism. There is a map $c\colon \mathsf{Set}(\mathcal{X}^+) \to \mathsf{Sub}_{Cl}(\mathcal{X})$ sending $\{\xi\} \in \mathsf{Set}(\mathcal{X}^+)$ to

$$c(\{\xi\}) = \mathrm{cl}(\{\xi\}) \in \mathsf{Sub}_{Cl}(\mathcal{X}). \tag{4.1}$$

The map c may not preserve intersection because of the existence of the lower-dimensional cells of $\mathcal{X}$; c is not a lattice morphism in general. The condition that $\mathrm{cl}\,|\xi| = |\,\mathrm{cl}\,\xi|$ implies that the isomorphism k factors as

$$\mathsf{Set}(\mathcal{X}^+) \xrightarrow{c} \mathsf{Sub}_{Cl}(\mathcal{X}) \xrightarrow{|\cdot|} \mathsf{X}.$$

Let $i\colon \mathcal{X}^+ \hookrightarrow \mathcal{X}$ be the inclusion of the anti-chain of top-cells into $\mathcal{X}$. This is a poset morphism. We set $j = \mathsf{O}(i)\colon \mathsf{Sub}_{Cl}(\mathcal{X}) \to \mathsf{Set}(\mathcal{X}^+)$, the Birkhoff dual of i (see Section 2.4.3). It is readily seen that j is the map

$$\mathsf{Sub}_{Cl}(\mathcal{X}) \ni \mathcal{K} \mapsto \mathcal{K} \cap \mathcal{X}^+ \in \mathsf{Set}(\mathcal{X}^+).$$

It can also be checked that $j \circ c = \mathrm{id}$. Thus we have the following diagram; following [25] we use the notation $\leftrightarrow$ to denote an isomorphism.

$$\begin{tikzcd}
\mathsf{Set}(\mathcal{X}^+) \arrow[rr, leftrightarrow, "k"] \arrow[dr, shift right, "c"'] & & \mathsf{X} \arrow[dl, leftrightarrow, "|\cdot|"] \\
& \mathsf{Sub}_{Cl}(\mathcal{X}) \arrow[ul, shift right, "j"'] &
\end{tikzcd} \tag{4.2}$$

4.2.2 Combinatorialization of Dynamics

The starting point for the combinatorialization of dynamics is a binary relation $\mathcal{F}$ defined on the top-cells $\mathcal{X}^+$ of the complex $\mathcal{X}$. A good combinatorial approximation is one that relates the structures of the combinatorial dynamics (e.g., $\mathsf{Att}(\mathcal{F}), \mathsf{Invset}^+(\mathcal{F})$) to the structures of the underlying continuous dynamics (e.g., $\mathsf{Att}(\varphi), \mathsf{ABlock}_{\mathscr{R}}(\varphi)$); see Section 2.9 for definitions. Roughly, speaking, we add in dynamics to the following conceptual schematic.

$$\begin{tikzcd}
& \text{Dynamics} \arrow[dl, leftrightarrow] \arrow[dr, leftrightarrow] & \\
\text{Combinatorics} \arrow[rr, leftrightarrow] \arrow[dr, leftrightarrow] & & \text{Topology} \\
& \text{Algebra} \arrow[ur, leftrightarrow] &
\end{tikzcd}$$

As $\mathsf{Invset}^+(\mathcal{F}) \subseteq \mathsf{Set}(\mathcal{X}^+)$, we may restrict the maps k, c in (4.2) to $\mathsf{Invset}^+(\mathcal{F})$ to get the following diagram.

$$\begin{tikzcd}
\mathsf{Invset}^+(\mathcal{F}) \arrow[rr, "k"] \arrow[dr, "c"'] & & \mathsf{X} \\
& \mathsf{Sub}_{Cl}(\mathcal{X}) \arrow[ur, "|\cdot|"'] &
\end{tikzcd} \tag{4.3}$$

Definition 4.2.2. Let $\mathcal{X}$ be a polyhedral complex in $\mathbb{R}^n$. A binary relation $\mathcal{F}$ on the collection of top-cells $\mathcal{X}^+$ is a *discrete approximation* for the flow $\varphi\colon \mathbb{R} \times |\mathcal{X}| \to |\mathcal{X}|$ if $k(\mathsf{Invset}^+(\mathcal{F}))$ is a sublattice of $\mathsf{ABlock}_{\mathscr{R}}(\varphi)$.

This definition draws heavily on the definition of *combinatorial models*, given in [25]. For instance, given a discrete approximation $\mathcal{F}$, let $\mathsf{N} = k(\mathsf{Invset}^+(\mathcal{F}))$ and $\mathsf{A} = \omega(\mathsf{N})$ (recall ω from Definition 2.9.3). Then the following diagram commutes.

$$\begin{tikzcd}
\mathsf{Invset}^+(\mathcal{F}) \arrow[r, leftrightarrow, "k"] \arrow[d, twoheadrightarrow, "\omega"] & \mathsf{N} \arrow[r, "\subset"] \arrow[d, twoheadrightarrow, "\omega"] & \mathsf{ABlock}_{\mathscr{R}}(\varphi) \arrow[d, twoheadrightarrow, "\omega"] \\
\mathsf{Att}(\mathcal{F}) & \mathsf{A} \arrow[r, "\subset"] & \mathsf{Att}(\varphi)
\end{tikzcd}$$

The primary result of this section is Theorem 4.2.4: if $\mathcal{F}$ is a discrete approximation then c is a lattice homomorphism, i.e., the following is a diagram in **FDLat**.

$$\begin{tikzcd}
& \mathsf{N} \arrow[dr, "i"] & \\
\mathsf{Invset}^+(\mathcal{F}) \arrow[ur, leftrightarrow, "k"] \arrow[rr, "k"] \arrow[dr, "c"'] & & \mathsf{X} \\
& \mathsf{Sub}_{Cl}(\mathcal{X}) \arrow[ur, "|\cdot|"'] &
\end{tikzcd}$$

Letting $\mathsf{L} = \mathsf{Invset}^+(\mathcal{F})$, Theorem 4.2.4 implies that $(\mathcal{X}, c)$ is an L-filtered cell complex. The Conley index of an attracting block $A \in \mathsf{N}$ can be computed by computing the homology of the associated subcomplex in $\mathcal{X}$, i.e., the image of A under the composition $\mathsf{N} \to \mathsf{Invset}^+(\mathcal{F}) \xrightarrow{c} \mathsf{Sub}_{Cl}(\mathcal{X})$. For the next lemma, recall the notion of frontier, denoted fr, given in Definition 2.5.13. Note also that bd refers to the (topological) boundary of a set.

Lemma 4.2.3. *Let $\mathcal{X}$ be a polyhedral complex. Let $\mathcal{F}$ be a discrete approximation for $\varphi\colon \mathbb{R} \times |\mathcal{X}| \to |\mathcal{X}|$. Let $a \in \mathsf{Invset}^+(\mathcal{F})$ and $\sigma \in \operatorname{fr}(c(a))$. For $T > 0$:*

1. *if $\varphi(T, |\sigma|) \cap |\tau| \neq \varnothing$ then $\tau \in \operatorname{int} c(a)$;*

2. *$\varphi([0, T], |\sigma|) \cap |\operatorname{star}(\sigma) \setminus \sigma| \neq \varnothing$.*

Proof. For part (1), let $\varphi(T, |\sigma|) \cap |\tau| \neq \varnothing$. If $\mathcal{F}$ is a discrete approximation then $|c(a)|$ is an attracting block. Therefore as $|\sigma| \subset |c(a)|$ we have that $\varphi(|\sigma|, T) \subset \operatorname{int}|c(a)|$. It follows from Proposition 2.7.4 that $\operatorname{int}|c(a)| = |\operatorname{int} c(a)|$. It follows that if $|\tau| \cap \varphi(T, |\sigma|) \neq \varnothing$ then $|\tau| \cap |\operatorname{int} c(a)| \neq \varnothing$. Lemma 2.7.2 implies that as $|\tau| \cap |\operatorname{int} c(a)| \neq \varnothing$ then $\tau \in \operatorname{int} c(a)$.

For part (2), it follows from Corollary 2.7.5 that if $\sigma \in \mathrm{fr}\ c(a)$ then $|\sigma| \subset |\,\mathrm{fr}\ c(a)| = \mathrm{bd}\,|c(a)|$. Since $|c(a)|$ is an attracting block, $|\sigma|$ is not fixed by φ. This implies that

$$\varphi([0,T],|\sigma|) \cap |(\mathrm{star}\,\sigma \cup \mathrm{cl}\,\sigma) \setminus \sigma| \neq \varnothing.$$

Suppose that $\varphi([0,T],|\sigma|) \cap |\,\mathrm{star}(\sigma) \setminus \sigma| = \varnothing$. This implies $\varphi([0,T],|\sigma|) \cap |\,\mathrm{cl}(\sigma) \setminus \sigma| \neq \varnothing$. Since $\sigma \in \mathrm{fr}\ c(a)$ we have that $\mathrm{cl}(\sigma) \subset \mathrm{fr}\ c(a)$. It follows from Corollary 2.7.5 that $|\,\mathrm{cl}(\sigma)| \subset \mathrm{bd}\,|c(a)|$. However, this is a contradiction as $|c(a)|$ is an attracting block and $|\sigma|$ must be mapped into $\mathrm{int}|c(a)|$. Therefore $\varphi([0,T],|\sigma|) \cap |\,\mathrm{star}(\sigma) \setminus \sigma| \neq \varnothing$. ☐

Theorem 4.2.4. *Let $\mathcal{X}$ be a polyhedral complex and let $\mathcal{F}$ be a discrete approximation for $\varphi\colon \mathbb{R} \times |\mathcal{X}| \to |\mathcal{X}|$. The map $c\colon \mathsf{Invset}^+(\mathcal{F}) \to \mathsf{Sub}_{Cl}(\mathcal{X})$ given by $c(a) = \mathrm{cl}(a)$ is a lattice homomorphism.*

Proof. We begin with

$$c(0) = c(\varnothing) = \varnothing = 0 \quad \text{ and } \quad c(1) = c(\mathcal{X}^+) = \mathcal{X} = 1.$$

Let $a, b \in \mathsf{Invset}^+(\mathcal{F})$. A routine calculation shows that

$$\begin{aligned} c(a \vee b) = \mathrm{cl}(a \cup b) = \bigcup_{\xi \in a \cup b} \mathrm{cl}(\xi) &= \Big(\bigcup_{\xi \in a} \mathrm{cl}(\xi)\Big) \cup \Big(\bigcup_{\xi' \in b} \mathrm{cl}(\xi')\Big) \\ &= \mathrm{cl}(a) \cup \mathrm{cl}(b) \\ &= c(a) \vee c(b). \end{aligned}$$

Furthermore, we have that

$$\begin{aligned} c(a \wedge b) = \mathrm{cl}(a \cap b) = \bigcup_{\xi \in a \cap b} \mathrm{cl}(\xi) \subseteq \Big(\bigcup_{\xi \in a} \mathrm{cl}(\xi)\Big) \cap \Big(\bigcup_{\xi' \in b} \mathrm{cl}(\xi')\Big) &= \mathrm{cl}(a) \cap \mathrm{cl}(b) \\ &= c(a) \wedge c(b). \end{aligned}$$

Notice that if $a \leq b$ then $a \subseteq b$. Thus

$$c(a \wedge b) = \mathrm{cl}(a \cap b) = \mathrm{cl}(a) = \mathrm{cl}(a) \cap \mathrm{cl}(b) = c(a) \wedge c(b).$$

Assume a and b are incomparable. It remains to show that $c(a) \wedge c(b) \subseteq c(a \wedge b)$. To this end, let $\sigma \in c(a) \wedge c(b) = \mathrm{cl}(a) \cap \mathrm{cl}(b)$ and suppose that $\sigma \notin \mathrm{cl}(a \cap b)$. Since

$\sigma \in c(a) \cap c(b)$, there must exist $\xi \in a$ and $\xi' \in b$ with $\sigma \in \operatorname{cl} \xi$ and $\sigma \in \operatorname{cl} \xi'$. Moreover, since $\sigma \notin \operatorname{cl}(a \cap b)$ it follows that $\xi \in a \setminus b$ and $\xi' \in b \setminus a$. Thus $\operatorname{star}(\sigma) \not\subset c(a)$ and $\operatorname{star}(\sigma) \not\subset c(b)$. Therefore $\sigma \in \operatorname{fr}\, c(a) \cap \operatorname{fr}\, c(b))$. As $\mathcal{F}$ is a discrete approximation we have that $|c(a)|$ and $|c(b)|$ are attracting blocks for φ. For any $T > 0$ we must have $\varphi([0,T], |\sigma|) \cap |\tau| \neq \varnothing$ for some $\tau \in \operatorname{star}(\sigma) \setminus \sigma$ by Lemma 4.2.3 (2). It follows from Lemma 4.2.3 (1) that $\tau \in \operatorname{int} c(a) \cap \operatorname{int} c(b)$. Hence if $\xi'' \in \mathcal{X}^+ \cap \operatorname{star}(\tau)$ then $\xi'' \in a \cap b$. But $\sigma \leq \tau \leq \xi''$ implies that $\sigma \in \operatorname{cl}(\xi'')$. Therefore $\sigma \in \operatorname{cl}(a \cap b)$. This is a contradiction, hence $c(a) \wedge c(b) = c(a \wedge b)$. □

Theorem 4.2.4 states that for a discrete approximation $\mathcal{F}$, the map $c\colon \mathsf{Invset}^+(\mathcal{F}) \to \mathsf{Sub}_{Cl}(\mathcal{X})$ is a lattice morphism. Therefore c has a Birkhoff dual, $\mathsf{J}(c)\colon \mathsf{J}(\mathsf{Sub}_{Cl}(\mathcal{X})) \to \mathsf{J}(\mathsf{Invset}^+(\mathcal{F}))$. Since O, J are contravariant functors which form an equivalence of categories (see Theorem 2.4.19), there is a natural isomorphism $\eta\colon \mathrm{id} \to \mathsf{JO}$. This implies that there are poset isomorphisms $\eta_{\mathcal{X}}, \eta_{\mathsf{SC}(\mathcal{F})}$, and a poset morphism $\nu = \eta^{-1}_{\mathsf{SC}(\mathcal{F})} \circ \mathsf{J}(c) \circ \eta_{\mathcal{X}}$ fitting into the following diagram.

$$\begin{tikzcd}
\mathcal{X} \arrow[r, "\nu"] \arrow[d, "\eta_{\mathcal{X}}"] & \mathsf{SC}(\mathcal{F}) \arrow[d, "\eta_{\mathsf{SC}(\mathcal{F})}"] \\
\mathsf{J}(\mathsf{Sub}_{Cl}(\mathcal{X})) \arrow[r, "\mathsf{J}(c)"] & \mathsf{J}(\mathsf{Invset}^+(\mathcal{F}))
\end{tikzcd}$$

The pair $(\mathcal{X}, \nu)$ is a P-graded cell complex with associated graded chain complex $(C(\mathcal{X}), \pi^\nu)$; see Section 3.3.4. We can cook up an explicit formula for ν as follows. First, note that every top-cell ξ is a vertex in the graph $\mathcal{F}$ and belongs to the strongly connected component $[\xi]$. There is a map $\nu_0\colon \mathcal{X}^+ \to \mathsf{SC}(\mathcal{F})$ given by $\nu_0(\xi) = [\xi]$, taking a top-cell to the strongly connected component in which it belongs. Tthe collection of top-cells $\mathcal{X}^+$ form an antichain (see Definition 2.4.2) in the face poset $(\mathcal{X}, \leq)$. Therefore $\mathcal{X}^+$ may be regarded as a poset (where the order $\leq$ is trivial) and ν_0 may be regarded as a poset morphism.

Proposition 4.2.5. *The map $\nu = \eta^{-1}_{\mathsf{SC}(\mathcal{F})} \circ \mathsf{J}(c) \circ \eta_{\mathcal{X}}\colon \mathcal{X} \to \mathsf{SC}(\mathcal{F})$ is given by the formula*

$$\mathcal{X} \ni \xi \mapsto \min_{\mathsf{SC}(\mathcal{F})} \{\nu_0(\xi')\colon \xi' \in \operatorname{star}(\xi) \cap \mathcal{X}^+\} \in \mathsf{SC}(\mathcal{F}). \tag{4.4}$$

Proof. The isomorphisms $\eta_{\mathcal{X}} \colon \mathcal{X} \to \mathsf{J}(\mathsf{Sub}_{Cl}(\mathcal{X}))$ and $\eta_{\mathsf{SC}(\mathcal{F})} \colon \mathsf{SC}(\mathcal{F}) \to \mathsf{J}(\mathsf{Invset}^+(\mathcal{F}))$ are given by the formulas $\eta_{\mathcal{X}}(\xi) = \downarrow \xi$ and $\eta_{\mathsf{SC}(\mathcal{F})}([\xi]) = \downarrow [\xi]$. Let $\xi \in \mathcal{X}$. We first compute $\mathsf{J}(c) \circ \eta_{\mathcal{X}}(\xi)$. Recall the definition of $\mathsf{J}(c)$ from Eqn. (2.1). We have that

$$\uparrow(\eta_{\mathcal{X}}(\xi)) = \uparrow(\downarrow \xi) = \{\mathcal{K} \in \mathsf{Sub}_{Cl}(\mathcal{X}) \mid \downarrow \xi \subset \mathcal{K}\}.$$

By the definition of c we have $c^{-1}(\uparrow(\eta(\xi))) = \{a \in \mathsf{Invset}^+(\mathcal{F}) \mid \xi \in \mathrm{cl}(a)\}$. Since $\mathsf{J}(c)$ is well-defined, this set has a minimal element which is join-irreducible, i.e., a minimal $a_0 \in \mathsf{J}(\mathsf{Invset}^+(\mathcal{F}))$ such that $\xi \in \mathrm{cl}(a_0)$. This implies that

$$\min\{a \in \mathsf{Invset}^+(\mathcal{F}) \mid \xi \in \mathrm{cl}(a)\} = \min\{a \in \mathsf{J}(\mathsf{Invset}^+(\mathcal{F})) \mid \xi \in \mathrm{cl}(a)\}.$$

First, as $\eta_{\mathsf{SC}(\mathcal{F})}$ is a poset isomorphism we have that: $a \in \mathsf{J}(\mathsf{Invset}^+(\mathcal{F}))$ if and only if $a = \eta_{\mathsf{SC}}([\xi'']) = \downarrow [\xi'']$ for $[\xi''] \in \mathsf{SC}(\mathcal{F})$. Second, we have that: $\xi \in \mathrm{cl}(\downarrow [\xi''])$ if and only if there is $\xi' \in \downarrow [\xi'']$ such that $\xi \in \mathrm{cl}(\xi')$; $\xi \in \mathrm{cl}(\xi')$ if and only if $\xi' \in \mathrm{star}(\xi)$. Together, these facts imply that $\eta_{\mathsf{SC}(\mathcal{F})}$ furnishes an (order-preserving) bijection

$$\{[\xi'] \in \mathsf{SC}(\mathcal{F}) \mid \xi' \in \mathrm{star}(\xi) \cap \mathcal{X}^+\} \xrightarrow{\eta} \{a \in \mathsf{J}(\mathsf{Invset}^+(\mathcal{F})) \mid \xi \in \mathrm{cl}(a)\}.$$

This implies that $\{[\xi'] \in \mathsf{SC}(\mathcal{F}) \mid \xi' \in \mathrm{star}(\xi) \cap \mathcal{X}^+\}$ has a minimal value. Setting $a_0 = \min\{a \in \mathsf{J}(\mathsf{Invset}^+(\mathcal{F})) \mid \xi \in \mathrm{cl}(a)\}$ and $p_0 = \{[\xi'] \in \mathsf{SC}(\mathcal{F}) \mid \xi' \in \mathrm{star}(\xi) \cap \mathcal{X}^+\}$ we have that $a_0 = \eta_{\mathsf{SC}(\mathcal{F})}(p_0)$ and $p_0 = \eta^{-1}_{\mathsf{SC}(\mathcal{F})}(a_0)$. Therefore ν is given by (4.4). This completes the proof. □

4.3 Transversality Models

In this section we define a certain class of binary relations that yield discrete approximation for flows which are generated by a vector field $\dot{x} = f(x)$ in subsets in $\mathbb{R}^n$. For the remainder of this section, let $\mathcal{X}$ be a pure polyhedral cell complex in $\mathbb{R}^n$; see Definition 2.5.14 and Section 2.7. Notice that since $\mathcal{X}$ is pure, we have that $\mathcal{X}^+ = \mathcal{X}_n$.

Definition 4.3.1. The *adjacency relation* $\mathcal{E}$ on $\mathcal{X}^+$ is the (symmetric) relation given by

$$(\xi, \xi') \in \mathcal{E} \iff \mathcal{X}_{n-1} \cap \mathrm{cl}(\xi) \cap \mathrm{cl}(\xi') \neq \varnothing \quad \text{and} \quad \xi \neq \xi'$$

That is, $\mathcal{E}$ is the collection of pairs of top-cells that share an $(n-1)$-dimensional face. Since all cells in $\mathcal{X}$ are convex, two distinct top cells cannot share more than one $(n-1)$-dimensional face. Let $\rho\colon \mathcal{E} \to \mathcal{X}_{n-1}$ be the map that takes a pair of distinct n-cells (ξ, ξ') to the shared $(n-1)$-dimensional face; $\rho(\xi, \xi') = \rho(\xi', \xi)$. If $\sigma \in \mathcal{X}_{n-1}$ then $\operatorname{cl}|\sigma|$ is a closed convex set in $\mathbb{R}^n$. This can be regarded as an affine $(n-1)$-dimensional space. Therefore for any $x \in \operatorname{cl}\rho(\xi, \xi')$ there are two unit vectors (which are negatives of each other) normal to $\operatorname{cl}|\sigma|$ at x. For $(\xi, \xi') \in \mathcal{E}$ and $x \in \operatorname{cl}\rho(\xi, \xi')$, let $\vec{n}_{\xi\to\xi'}$ be the normal in the direction from ξ to ξ'. Let λ_x be the assignment

$$\mathcal{E} \ni (\xi, \xi') \mapsto \vec{n}_{\xi\to\xi'}.$$

Definition 4.3.2. Let $\mathcal{X}$ be a pure polyhedral complex and $\varphi : \mathbb{R} \times |\mathcal{X}| \to |\mathcal{X}|$ be a flow generated by $\dot{x} = f(x)$. We say that a pair $(\xi, \xi') \in \mathcal{E}$ is *(positively) transverse* with respect to φ if $f(x) \cdot \lambda_x(\xi, \xi') > 0$ for all $x \in \operatorname{cl}|\rho(\xi, \xi')|$. A picture of positively transversality is Figure 4.1.

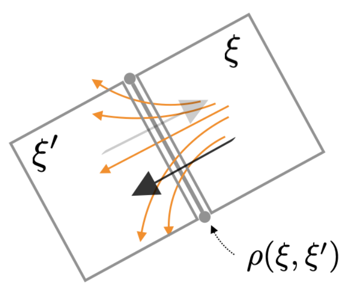

Figure 4.1: An example of a pair (ξ, ξ') that is positively transverse to φ. A qualitative picture of the orbits of the flow φ of $\dot{x} = f(x)$ is depicted.

The notion of transversality gives rise to the following restriction of $\mathcal{E}$.

Definition 4.3.3. Let $\mathcal{X}$ be a pure polyhedral complex , $\mathcal{F}$ be a binary relation on $\mathcal{X}^+$ and $\varphi : \mathbb{R} \times |\mathcal{X}| \to |\mathcal{X}|$ be a flow generated by $\dot{x} = f(x)$. The pair $(\mathcal{X}, \mathcal{F})$ is a *transversality model* for φ if

$$(\xi', \xi) \in \mathcal{E} \setminus \mathcal{F} \implies (\xi, \xi') \text{ is positively transverse with respect to } \varphi.$$

$\mathcal{F}$ is called the *discrete flow.*

The intuition is as follows. The top-cells ξ, ξ' intersect in an $(n-1)$-dimensional cell $\rho(\xi', \xi)$. If the pair is positively transverse in the direction $\lambda(\xi, \xi')$, then one is allowed remove the pair (ξ', ξ) from the relation $\mathcal{E}$.

Remark 4.3.4. Given a pure polyhedral complex $\mathcal{X}$ and flow φ generated by $\dot{x} = f(x)$, there are two extremal transversality models.

- $\mathcal{E}$ itself is the maximal transversality model.
- Define the relation $\mathcal{T} \subset \mathcal{E}$ via

$$\mathcal{T} = \{(\xi', \xi) \in \mathcal{E} \mid (\xi, \xi') \text{ is positively transverse with respect to } \varphi\}$$

The minimal transversality model $\mathcal{F} \subset \mathcal{E}$ is defined by

$$\mathcal{F} = \mathcal{E} \setminus \mathcal{T}.$$

We can now state our main result for transversality models.

Theorem 4.3.5. *If $(\mathcal{X}, \mathcal{F})$ is a transversality model for the flow $\varphi\colon \mathbb{R} \times |\mathcal{X}| \to |\mathcal{X}|$ generated by $\dot{x} = f(x)$ then $\mathcal{F}$ is a discrete approximation for φ.*

Proof. To show that $\mathcal{F}$ is a discrete approximation, we must show that there is a lattice homomorphism $\mathsf{N} = k(\mathsf{Invset}^+(\mathcal{F}))$ is a sublattice of $\mathsf{ABlock}_{\mathscr{R}}(\varphi)$; recall $\mathsf{ABlock}_{\mathscr{R}}(\varphi)$ from Definition 2.9.4. Since N is the image of a lattice homomorphism, it suffices to show that $|\operatorname{cl}(a)|$ is an attracting block for any $a \in \mathsf{Invset}^+(\mathcal{F})$. To this end, let $a \in \mathsf{Invset}^+(\mathcal{F})$ and set $A = \operatorname{cl}(a)$. We must show that $|A|$ is an attracting block. It suffices to show that if $x \in \operatorname{bd}(|A|)$ then $\varphi^t(x) \in \operatorname{int}(|A|)$ for all $t > 0$. Let $x \in \operatorname{bd}(|A|)$. Then $x \in |\sigma|$ for some $\sigma \in \operatorname{fr}(A)$ by Corollary 2.7.5. Consider $Z := \mathcal{X}_{n-1} \cap \operatorname{star}(\sigma) \cap \operatorname{fr}(A)$. Let $\tau \in Z$. Since $a \in \mathsf{Invset}^+(\mathcal{F})$ there is a pair $(\xi, \xi') \in \mathcal{E}$ such that $\tau = \rho(\xi, \xi')$ and (ξ, ξ') is positively transverse with respect to φ. Since $x \in |\sigma| \subset \operatorname{cl}|\tau|$ this implies that

$$f(x) \cdot \lambda_x(\xi, \xi') > 0. \tag{4.5}$$

Since (4.5) holds for all $\tau \in Z$, this implies that $\varphi^t(x) \in \operatorname{int}|A|$ for all t sufficiently small. Since this holds at every $x \in \operatorname{bd}(|A|)$ it follows that $\varphi^t(x) \in \operatorname{int}|A|$ for all t. Therefore $|A|$ is an attracting block. $\square$

4.4 Computational Examples

In this section we translate some well known connection matrix examples (see [17, 40]) into our framework. Importantly, these examples illustrate that transversality models are general enough to capture the interesting phenomena of the connection matrix theory. In particular, the following examples highlight that non-uniqueness of the connection matrix may occur in computations.

Example 4.4.1 (Saddle-Saddle Connection)**.** This flow appears as a qualitative example in [17, Example 6.1] and [40, Example 2.2]. If the stable and unstable manifolds of fixed points do not intersect transversely, then this may lead to non-uniqueness in the connection matrix. Figure 4.2(a) contains a transversality model $(\mathcal{X}, \mathcal{F})$ (gray) for the flow φ (orange). The cubical complex $\mathcal{X}$ has nine top-dimensional cubes (2-cells). The arrows along the boundary of $\mathcal{X}$ only serve to indicate that the entire polyhedron is an attracting block in a larger flow on $\mathbb{R}^2$; that is, they are boundary conditions. In Figures 4.2(b) and (c) are, respectively, the Conley-Morse graph (see Example 1.2.2) and the connection matrix for the Conley complex $(\mathcal{M}, \mu)$. $\mathcal{M}$ consists of five cells: three 0-cells η_0, η_1, η_2 with $\mu(\eta_i) = i$ and two 1-cells with $\mu(\sigma_i) = i$. The non-uniqueness of the connection matrix is expressed as the equation $a + b \equiv 0 \mod 2$. It is straightforward to show that different discrete Morse theoretic calculations it is possible to find both of the connection matrices. In Franzosa's model, the parameter controls the connecting

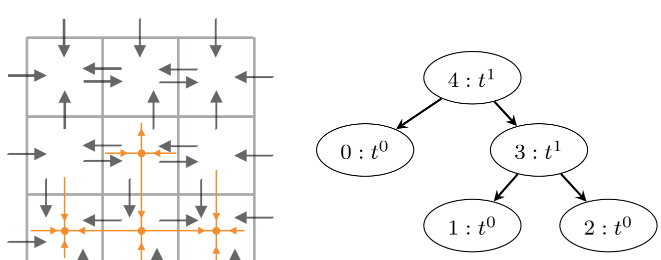

(a) (b)

$$\Delta_1 = \begin{array}{c c} & \begin{array}{cc} \sigma_3 & \sigma_4 \end{array} \\ \begin{array}{c} \eta_0 \\ \eta_1 \\ \eta_2 \end{array} & \begin{pmatrix} & 1 \\ 1 & a \\ 1 & b \end{pmatrix} \end{array}$$

(c)

Figure 4.2: (a) Transversality model for a saddle-saddle connection. (b) Conley Morse graph. (c) Connection matrix (zero entries not shown).

orbit and for a critical parameter there is a saddle-saddle connection. On either side of the critical parameter, there is no saddle-saddle connection. At this level of resolution, i.e., choice of $\mathcal{X}$, it is not possible to separate the distinct cases. If the location of the fixed point is perturbed then it becomes possible to find a transversality model which has a unique connection matrix, see Figure 4.3. The blank node has trivial index.

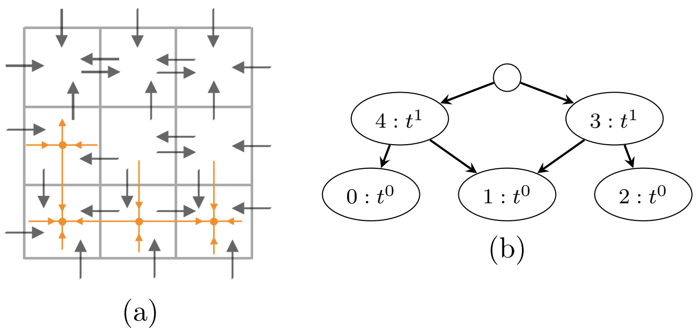

$$
\begin{array}{c} \\ \eta_0 \\ \eta_1 \\ \eta_2 \end{array}
\begin{array}{c}
\begin{array}{cc} \sigma_3 & \sigma_4 \end{array} \\
\begin{pmatrix} & 1 \\ 1 & 1 \\ 1 & \end{pmatrix}
\end{array}
$$

(c)

Figure 4.3: (a) Transversality model for perturbation and transverse intersection. (b) Conley Morse graph (c) Connection matrix.

Example 4.4.2 (Periodic Orbit)**.** In [40, Example 3.1] Reineck showed that periodic orbits may cause non-uniqueness of the connection matrix, even in the case when stable and unstable manifolds intersect transversally. A transversality model can be constructed for the example, warranting application of the computational connection matrix theory. Our system differs slightly from [40, Example 3.1] by an addition of two isolated invariant sets: a stable periodic orbit encircling the example as well as a central unstable fixed point. There are seven cells in the Conley complex: two 0-cells, η_0, η_1, four 1-cells $\sigma_0, \sigma_1, \sigma_2, \sigma_3$ and three 2-cells, ξ_4, ξ_5, ξ_6; we have $\mu(\eta_i) = i, \mu(\sigma_i) = i$ and $\mu(\eta_i) = i$. The conventional connection matrix analysis proceeds as follows: the fact that the total index $H(S) = 0$ implies that $\ker \partial_2 = 0$. This implies $\dim(\operatorname{im} \partial_2) = 3$. Thus the columns of Δ associated to η_4, η_5, η_6 must be linearly independent. This gives four possible connection matrices; all four may be obtained using different Morse-theoretic reductions.

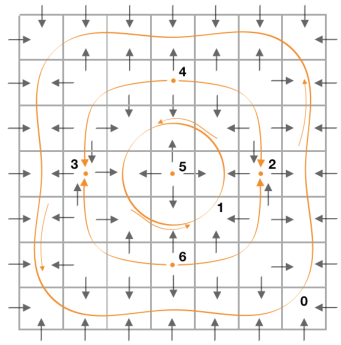

Figure 4.4: A transversality model for Reineck's example.

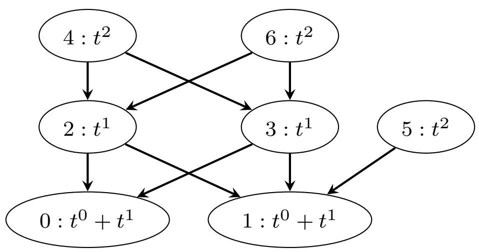

(a)

$$
\begin{array}{c}
\\ \eta_0 \\ \eta_1 \\ \sigma_0 \\ \sigma_1 \\ \sigma_2 \\ \sigma_3
\end{array}
\begin{array}{c}
\begin{array}{ccccc} \sigma_2 & \sigma_3 & \eta_4 & \eta_5 & \eta_6 \end{array} \\
\begin{pmatrix}
1 & 1 & & & \\
1 & 1 & & & \\
 & & a & & c \\
 & & b & 1 & d \\
 & & 1 & & 1 \\
 & & 1 & & 1
\end{pmatrix}
\end{array}
$$

(b)

Figure 4.5: (a) Conley-Morse graph. (b) Connection matrix.

4.5 Parabolic Recurrence Relations

In this section we provide a review of parabolic recurrence relations. This is the primary ingredient for the Morse theory on braids developed in [21]. Other references include [20, 49, 48, 50]. We also recall the space of discretized braids, which is the phase space for parabolic recurrence relations. Finally, in Section 4.6 we will build transversality models for a parabolic recurrence relation that fixes a particular braid.

Definition 4.5.1. A *parabolic recurrence relation* (of period $d > 0$) is a system of differential equations

$$\dot{x}_i = R_i(x_{i-1}, x_i, x_{i+1}) \quad \text{for } i \in \mathbb{Z}$$

where each $R_i \colon \mathbb{R}^3 \to \mathbb{R}$ is a smooth bounded function with $R_{i+d} = R_i$. Moreover each R_i is *parabolic*, i.e., $\partial_1 R_i > 0$ and $\partial_3 R_i \geq 0$. We denote the entire vector field by $\mathcal{R}$.

Examples of parabolic recurrence relations include discretizations of uniform parabolic PDE's, monotone twist maps, etc [49, 48, 20, 21, 50]. Given a sequence $\{y_i\}_{i\in\mathbb{Z}}$ with $y_i \in \mathbb{R}$ one can evolve the points according to the equation

$$\dot{y}_i = R_i(y_{i-1}, y_i, y_{i+1}) \tag{4.6}$$

The sequence $\{y_i\}$ can be regarded as a discretized function, see Figure 4.6. Multiple discretized functions can be written as $\mathbf{v} = \{y_i^\alpha\}_{i\in\mathbb{Z}}^{\alpha=1,\dots,m}$ where α indexes the function. The collection $\mathbf{v}$ may be evolved according to (4.6). A *stationary solution of* $\mathcal{R}$ is a sequence $\{y_i\}$ such that $R_i(y_i) = 0$ for all i. A collection $\mathbf{v}$ is *stationary under* $\mathcal{R}$ if each y^α is stationary. In this case we say that $\mathcal{R}$ *fixes* $\mathbf{v}$. The theory is designed to use a stationary solutions to force the existence of other solutions.

4.5.1 Discretized Braids

According to the theory of [21], the collection $\mathbf{v} = \{y_i^\alpha\}$ should be regarded as a *discretized braid* and a single sequence y^α as a *strand*. This motivations the following definition of closed positive discretized braids, which we recall from [21].

Definition 4.5.2. The space of *closed discretized period d braids* on n strands, denoted $\mathbf{D}_m^d$, is the space of all ordered sets of *strands* $\mathbf{v} = \{\mathbf{v}^\alpha\}^{\alpha=1,\dots,m}$, defined as follows:

(a) each strand $\mathbf{v}^\alpha = (v_1^\alpha, v_1^\alpha, \dots, v_{d+1}^\alpha) \in \mathbb{R}^{d+1}$ consists of $d+1$ *anchor points* v_j^α;

(b) $v_{d+1}^\alpha = v_1^{\tau(\alpha)}$ for all $\alpha = 1, \dots, m$ for some permutation $\tau \in S_m$;

(c) for any pair of distinct strands $\mathbf{v}^\alpha$ and $\mathbf{v}^{\alpha'}$ such that $v_j^\alpha = v_j^{\alpha'}$ for some j, the *transversality* condition $\left(v_{j-1}^\alpha - v_{j-1}^{\alpha'}\right)\left(v_{j+1}^\alpha - v_{j+1}^{\alpha'}\right) < 0$ holds.

The elements $\mathbf{v} \in \mathbf{D}_m^d$ will be referred to as *discretized braids*. The *dimension* of $\mathbf{v}$ is d. The number of strands is m.

The path components of $\mathbf{D}_m^d$ comprise the *discrete braid classes*. The discrete braid class of a discrete braid diagram $\mathbf{u}$ is denoted $[\mathbf{u}]$. The closure of $\mathbf{D}_m^d$, denoted $\mathrm{cl}\,\mathbf{D}_m^d$, is obtained by ignoring condition (c) of Definition 4.5.2. The *singular discrete braids* are elements of $\Sigma_n^d = \mathrm{cl}\,\mathbf{D}_m^d \setminus \mathbf{D}_m^d$.

Discretized braids can be regarded as topological braids by extending each strand to a piecewise linear function, see Figure 4.6. In this figure is the *braid diagram*: the collection (i, v_i^α), with each strand extended to a piecewise linear function. Our braids will be positive, i.e., all crossings are positive. From (b), these braids are also closed, i.e., we may identify the last coordinate of the braid with the first.

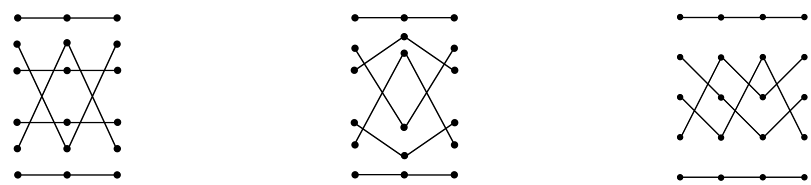

Figure 4.6: Braids $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w}$. $\mathbf{u}$ and $\mathbf{v}$ are 2-D braids. $\mathbf{w}$ is a 3-D braid.

The point is that a parabolic recurrence relation of period d relation defines a vector field on $\mathrm{cl}\,\mathbf{D}_m^d$ via Eqn. (4.6) and there is an associated flow. In order to prove forcing theorems, one needs to work with relative braids: a single strand $\mathbf{u}$ – called the *free strand* – relative to a fixed braid $\mathbf{v}$ – *the skeleton*. In this thesis we restrict to single strands relative to a fixed braid $\mathbf{v}$; relative braids can be defined in more generality, see [21]. Let $\mathcal{R}$ be a parabolic recurrence relation which fixes $\mathbf{v} \in \mathbf{D}_m^d$. Define

$$\mathbf{D}_1^d \text{ rel } \mathbf{v} := \{\mathbf{u} \in \mathbf{D}_1^d : \mathbf{u} \sqcup \mathbf{v} \in \mathbf{D}_{m+1}^d\}$$

The elements $\mathbf{D}_1^d$ are denoted $\mathbf{u}$ rel $\mathbf{v}$ and are called *discrete braids relative to* $\mathbf{v}$, or more simply *relative discrete braids*. The path components of $\mathbf{D}_1^d$ rel $\mathbf{v}$ comprise the *discrete braid classes relative to* $\mathbf{v}$, or more simply *relative discrete braid classes*, denoted $[\mathbf{u} \text{ rel } \mathbf{v}]$. A discrete braid class $[\mathbf{u} \text{ rel } \mathbf{v}]$ is *bounded* if $\mathrm{cl}[\mathbf{u} \text{ rel } \mathbf{v}]$ is a bounded set. The parabolic recurrence relation $\mathcal{R}$ defines a vector field on $\mathrm{cl}\,\mathbf{D}_1^d$ rel $\mathbf{v}$ via Eqn. (4.6) and there is an associated flow φ on $\mathrm{cl}\,\mathbf{D}_1^d$ rel $\mathbf{v}$.

Let $[\mathbf{u} \text{ rel } \mathbf{v}]$ be a bounded proper braid class. Set

$$N = \mathrm{cl}\big([\mathbf{u} \text{ rel } \mathbf{v}]\big),$$

which we may identify with a regular closed subset of $\mathbb{R}^d$. It follows from [21, Theorem 15] that N is an isolating neighborhood for φ and there is a well-defined (homological) Conley index which may be defined as follows. A boundary face of N is an occurrence where $u_i = v_i^\alpha$ for some i and some α. A boundary face of N can be given a co-orientation: such a face is regarded as positive co-oriented if $u_{i-1} - v_{i-1}^\alpha > 0$ and thus also $u_{i+1} - v_{i+1}^\alpha > 0$. Define N^- to be the closure of the union of boundary faces of N where the co-orientation is positive. Following [21], we define the *braid Conley homology* as

$$HC_k([\mathbf{u} \text{ rel } \mathbf{v}]) := H_k(N, N^-), \quad k \in \mathbb{N}. \tag{4.7}$$

It was proved in [21] that this is an invariant for $[\mathbf{u} \text{ rel } \mathbf{v}]$ and a Conley index for the flow φ.

Given $\mathbf{v} \in \mathbf{D}_m^d$, from the definition $\mathbf{v}$ has m strands and dimensionality d. These two parameters suffice to determine the cubical complex $\mathcal{X} = \mathcal{X}(d, m)$ on the d-cube $[0, m+1]^d$. Without loss of generality (up to scaling) we assume the following condition on our skeleton $\mathbf{v}$: for any fixed i the cross-section $(v_i^1, v_i^2, \dots, v_i^m)$ is a permutation of is a permutation of $(1, 2, \dots, m)$. That is, the (v_i^α) are integers and take unique values between 1 and m. This implies that the pairs (i, v_i^α) lie on the integer lattice within the box $[1, d+1] \times [1, m]$. The cubical complex $\mathcal{X}$ is comprised of all the cells ξ where

$$\xi = I_1 \times I_2 \times \ldots \times I_d \qquad \text{and } I_i = [l_i, l_i + 1] \text{ with } 1 \leq l_i \leq m.$$

Along any coordinate i there are $m+1$ top-dimensional d-cubes in $\mathcal{X}$. Given a top-cell $\xi \in \mathcal{X}^+$ we have that

$$\xi = [l_1, l_1 + 1] \times [l_2, l_2 + 1] \times \ldots \times [l_d, l_d + 1].$$

From the perspective of the cubical complex $\mathcal{X}$, a sequence $\mathbf{u} = (u_1, \cdots, u_{d+1})$, with $u_{d+1} = u_1$ and $\mathbf{u} \sqcup \mathbf{v} \in \mathbf{D}_{m+1}^d$, is a free strand. The relative discrete braid is denoted $\mathbf{u}$ rel $\mathbf{v} \in \mathbf{D}_1^d$ rel $\mathbf{v}$. The relative discrete braid class $[\mathbf{u} \text{ rel } \mathbf{v}]$ is a path component of $\mathbf{D}_1^d$. Moreover, if $[\mathbf{u} \text{ rel } \mathbf{v}]$ is bounded then it is a top-cell of the cubical complex $\mathcal{X}$. See Figure 4.7. The closure of the bounded discrete braid classes of $\mathbf{D}_1^d$ rel $\mathbf{v}$ is the cube $[0, m+1]^d \subset \mathbb{R}^d$. See Figure 4.8(b).

The free strand and braid have an associated crossing number, $\text{cross}(\mathbf{u})$ which is the number of intersections of the strand $\mathbf{u}$ with the strands in $\mathbf{v}$. For $\mathbf{u}'$ rel $\mathbf{v} \in [\mathbf{u} \text{ rel } \mathbf{v}]$ we have $\text{cross}(\mathbf{u}' \text{ rel } \mathbf{v}) = \text{cross}(\mathbf{u} \text{ rel } \mathbf{v})$, see [21]. Therefore the crossing number is

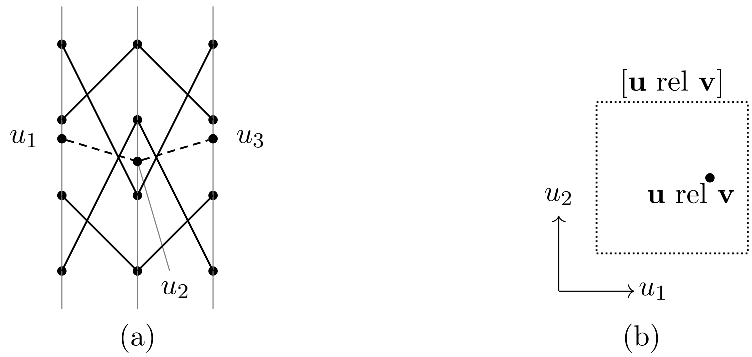

Figure 4.7: (a) Braid $\mathbf{v}$ with free strand $\mathbf{u}$. (b) A relative braid class $[\mathbf{u} \text{ rel } \mathbf{v}]$ which is a top-cell in $\mathcal{X}$.

constant on top-cells. This furnishes a function

$$\text{cross}\colon \mathcal{X}^+ \to \mathbb{N}.$$

The function cross$(\cdot)$ is a discrete Lyapunov function for the flow φ, see [21]. In Figure 4.8(a) the braid $\mathbf{v}$ is appended with two constants strands (these act as boundary conditions); (b) shows the cubical complex $\mathcal{X}$ with each top-cell (a two dimensional cube) is labeled with its crossing number. There are two kinds of relative braid classes in $\mathbf{D}_1^d$ rel $\mathbf{v}$.

Definition 4.5.3. A relative braid class $[\mathbf{u} \text{ rel } \mathbf{v}]$ is *improper* if $\mathbf{v}^\alpha \in \text{cl}[\mathbf{u} \text{ rel } \mathbf{v}]$ for some α. Otherwise a braid class is *proper*.

Here is another perspective on the sources of improperness. If a strand $\mathbf{v}^\alpha = (v_1^\alpha, \ldots, v_{d+1}^\alpha)$ in $\mathbf{v}$ obeys the condition $v_1^\alpha = v_{d+1}^\alpha$, then it corresponds to a vertex in $\mathcal{X}$ via

$$\mathbf{v}^\alpha = [v_1^\alpha, v_1^\alpha] \times \ldots \times [v_d^\alpha, v_d^\alpha]$$

We call these vertices in $\mathcal{X}$ *improper* and write $G_{\mathbf{v}}$ for the set of improper vertices. A relative braid class $[\mathbf{u} \text{ rel } \mathbf{v}]$ corresponds to a top-cell ξ in $\mathcal{X}$ and it is improper of $\mathbf{v}^\alpha \in \text{cl}\,\xi$.

4.6 Transversality Models for Parabolic Recurrence Relations

Define the relation $\mathcal{F}_0$ on set of top-cells $\mathcal{X}^+$ as follows. Recall the adjacency relation $\mathcal{E}$ on top-cells, given in Definition 4.3.1. A pair $(\xi, \xi') \in \mathcal{F}_0$ if

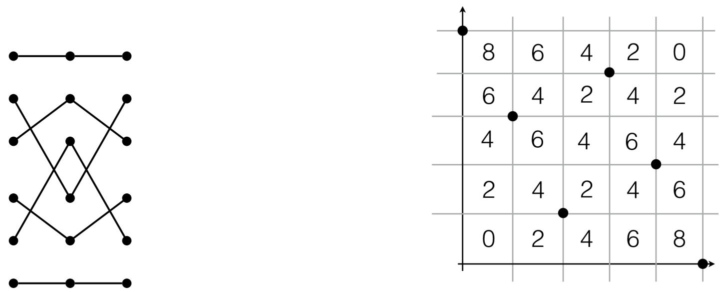

(a) (b)

Figure 4.8: (a) Braid $\mathbf{v}$ with two constant strands (top and bottom). (b). Cubical complex $\mathcal{X}$ and cross: $\mathcal{X}^+ \to \mathbb{N}$. Each dot is an improper vertex.

1. $(\xi, \xi') \in \mathcal{E}$. That is, $\xi, \xi' \in \mathcal{X}^+$ are adjacent top-cells;

2. $\operatorname{cross}(\xi) \geq \operatorname{cross}(\xi')$.

In order to construct a transversality model we must coarsen $\mathcal{F}_0$ by adding relations around improper vertices. Define the relation $\mathcal{F} \supset \mathcal{F}_0$ as follows. A pair $(\xi, \xi') \in \mathcal{F}$ if

1. $(\xi, \xi') \in \mathcal{F}_0$, or

2. $(\xi, \xi') \in \mathcal{E}$ and $\xi, \xi' \in \mathcal{X}^+ \cap \operatorname{star}(v)$ for some $v \in G_{\mathbf{v}}$.

Figure 4.9(b) shows the phase space (the two-dimensional cubical complex $\mathcal{X}$) as well as the relation $\mathcal{F}$. One-dimensional faces of two top-cells ξ and ξ' are labeled with an arrow $\xi \to \xi'$ to indicate that $(\xi, \xi') \in \mathcal{F}$. A face labeled with a bidirectional arrow $\xi \leftrightarrow \xi'$ indicates that both (ξ, ξ') and (ξ', ξ') are in $\mathcal{F}$.

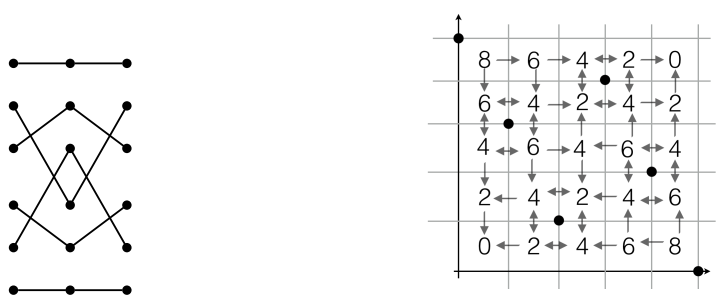

(a) (b)

Figure 4.9: (a) Braid $\mathbf{v}$ with constants strands. (b) Tranversality model $(\mathcal{X}^+, \mathcal{F})$.

Proposition 4.6.1. *Let $\mathcal{R}$ be a parabolic recurrence relation that fixes $\mathbf{v} \in \mathbf{D}^d_m$. The tuple $(\mathcal{X}, \mathcal{F})$, given as above, is a transversality model for the flow φ generated by $\mathcal{R}$ on*

$\mathrm{cl}\,\mathbf{D}_1^d$ rel $\mathbf{v}$.

Proof. $\mathcal{X}$ is a pure polyhedral complex in $\mathbb{R}^n$ by construction. Let $(\xi', \xi) \in \mathcal{E} \setminus \mathcal{F}$. We must show that (ξ, ξ') is positively transverse. Since $(\xi', \xi) \notin \mathcal{F}$ we have that $\mathrm{cross}(\xi') < \mathrm{cross}(\xi)$. Let $\mathbf{u}$ and $\mathbf{u}'$ be free strands such that $\xi = [\mathbf{u} \text{ rel } \mathbf{v}]$ and $\xi' = [\mathbf{u}' \text{ rel } \mathbf{v}]$. Since $(\xi', \xi) \in \mathcal{E} \setminus \mathcal{F}$ we must have that $[\mathbf{u} \text{ rel } \mathbf{v}]$ and $[\mathbf{u}' \text{ rel } \mathbf{v}]$ are proper classes. It follows from [21, Proposition 11] that (ξ, ξ') is positively transverse. $\square$

Since $\mathcal{F}$ is a transversality model for the flow generated by $\mathcal{R}$, it follows from Theorem 4.3.5 that $\mathcal{F}$ is a discrete approximation and it follows from Theorem 4.2.4 that $c\mathsf{Invset}^+(\mathcal{F}) \to \mathsf{Sub}(\mathcal{X})$ defined as in Eqn. (4.1) is a lattice homomorphism. Therefore there is an associated $\mathsf{SC}(\mathcal{F})$-graded cell complex $(\mathcal{X}, \nu)$ with ν by Eqn. (4.4). The graded cell complex $(\mathcal{X}, \nu)$ is input into Algorithm 3.7.8 (ConnectionMatrix).

4.7 Computational Examples for Parabolic Relations

In this section we examine the Conley-Morse graphs for the three braids of Figure 4.6 and their n-fold covers. The n-fold cover is the repetition of a braid n times. More discussion of similar braids and n-fold covers can be found in [21, Section 4.5: Example 1–3]. For all our examples, we append constant top and bottom strands to the skeleton, which act as (attracting) boundary conditions. This set of parabolic examples also serves as a collection of examples to benchmark the connection matrix algorithm. See Section 5.4 for the experimental results.

Example 4.7.1 (Two Dimensions)**.** Our first example in Figure 4.10(a) is a two-dimensional braid $\mathbf{u}$. More discussion of this particular braid can be found in [21, Section 4.5: Example 1]. In 2D one may work out the cross and $\mathcal{F}$ by hand, as in Figure 4.10(b). Since $\mathcal{F}$ is a tranversality model, there is an associated $\mathsf{SC}(\mathcal{F})$-graded cell complex $(\mathcal{X}, \nu)$; this is visualized (in the fashion of Example 1.2.2) in Figure 4.11. The output of the algorithm ConnectionMatrix is the Conley complex – an $\mathsf{SC}(\mathcal{F})$-graded complex $(\mathcal{M}, \mu)$. Moreover, the associated boundary operator ∂^M is a connection matrix for the tessellated Morse decomposition $\pi\colon \mathsf{M}(\mathsf{A}) \hookrightarrow \mathsf{T}(\mathsf{N})$. We depict this data in Figure 4.12(a).

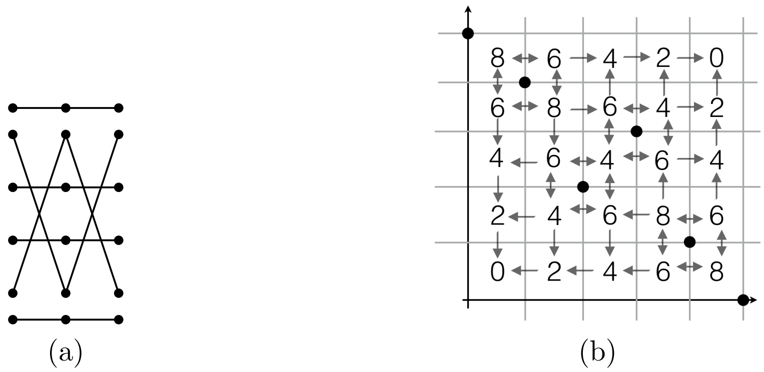

Figure 4.10: (a) The braid $\mathbf{u}$. (b) 2D cubical complex and cross: $\mathcal{X}^+ \to \mathbb{N}$. Directed graph $\mathcal{F}$ is depicted with vertex set the collection of 2-cells.

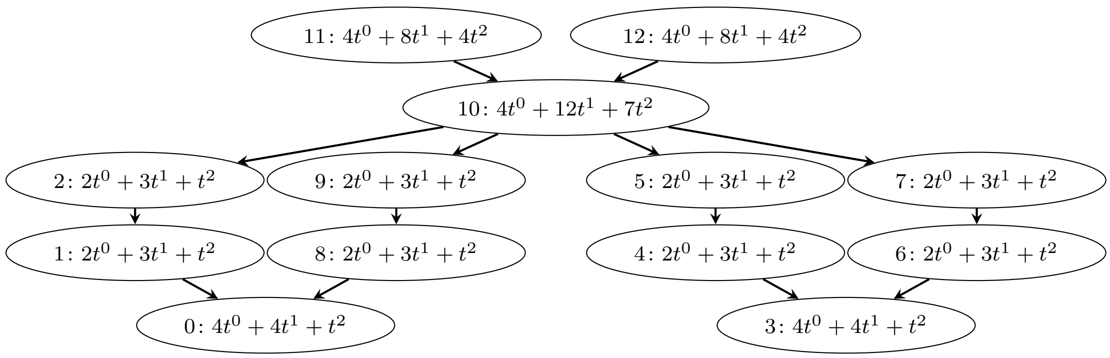

Figure 4.11: $\mathsf{SC}(\mathcal{F})$-graded complex associated to the braid $\mathbf{u}$.

The size of the poset $\mathsf{SC}(\mathcal{F})$ often increases rapidly with both dimension and number of strands. In this case it is convenient to restrict the poset $\mathsf{SC}(\mathcal{F})$ to the nodes with nontrivial Conley indices. We label this poset $\mathsf{RC}(\mathcal{F})$. $\mathsf{RC}(\mathcal{F})$ is the image of $\mathcal{M}$ under μ, and there is a factorization of μ as $\mathcal{M} \to \mathsf{RC}(\mathcal{F}) \rightarrowtail \mathsf{SC}(\mathcal{F})$. We call the restriction $(\mathcal{M}, \mu)$ the reduced Conley complex. The associated Conley-Morse graph is the *reduced Conley-Morse graph*, which is depicted in Figure 4.12 (b). Higher dimensonal examples are given in Figures 4.13 and 4.14. Figure 4.13(a) is the 2-fold cover of $\mathbf{u}$ with reduced Conley-Morse graph in (b). Figure 4.14(a) and (b) have the Conley-Morse graphs for the 4-fold cover and 5-fold cover, respectively.

Example 4.7.2 (Growth Rate)**.** A second example is similar to the braid in [21, Section 4.5: Example 3]. We examine the growth rate of $\mathsf{SC}(\mathcal{F})$ and $\mathsf{RC}(\mathcal{F})$ with respect to n-fold covers of $\mathbf{v}$. In [21] it is shown that for an n-fold cover of $\mathbf{v}$ there are at least $3^n - 2$

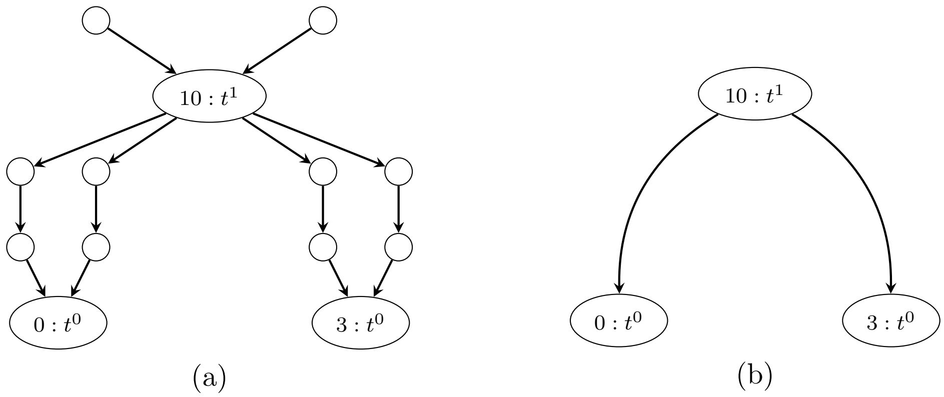

Figure 4.12: (a) Conley-Morse graph for the Conley complex $(\mathcal{M}, \mu)$. Nodes in $\mathsf{SC}(\mathcal{F})$ with trivial Conley index are blank. The tuple gives the Betti numbers (and thus the Conley index) of the associated fibers. (b) $\mathsf{RC}(\mathcal{F})$ and the reduced Conley-Morse graph. There is an order embedding $\mathsf{RC}(\mathcal{F}) \hookrightarrow \mathsf{SC}(\mathcal{F})$.

braid classes that have nontrivial Conley index. We can compare this lower bound with the true number of non-trivial classes for $2 \neq n \leq 5$. For this particular braid, the number of nontrivial indices grows so rapidly that we only depict the Conley-Morse graph for the **v** itself; see Figure 4.15. Table 4.16 gives dim (the dimensionality of the cubical complex), and $\#\mathsf{SC}$ and $\#\mathsf{RC}$, the sizes of the posets SC and RC, respectively. From Table 4.16 it is clear that the estimate misses many nontrivial classes.

	Growth Rate			
n-fold cover	dim	$\#$ SC	$\#$ RC	$3^n - 2$ Estimate
v	2	13	3	1
2-fold cover	4	114	32	7
3-fold cover	6	879	196	25
4-fold cover	8	7212	1153	79
5-fold cover	10	62 157	6724	241

Table 4.16: Growth rates for n-fold covers of **v**.

Example 4.7.3 (Pseudo-Anosov Braid)**.** Let **w** be the three-dimensional braid in Figure 4.17(a). Figures 4.18(a) gives the 2-fold cover of **w** and (b) gives the associated Conley-Morse graph. Finally, in Figure 4.19 we give the Conley-Morse graph for the three-fold cover of **w**.

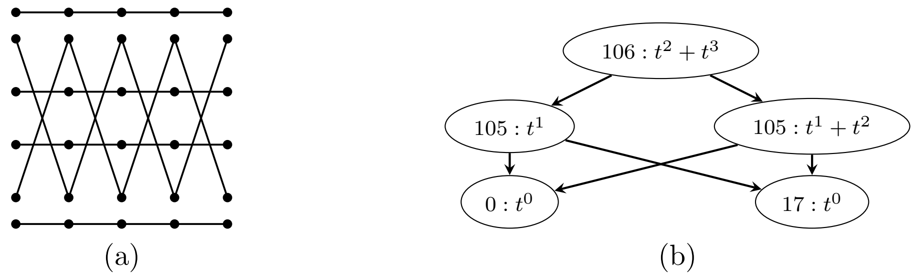

Figure 4.13: (a) 2-fold cover of **u**. (b) Reduced Conley Morse graph for 2-fold cover of **u**.

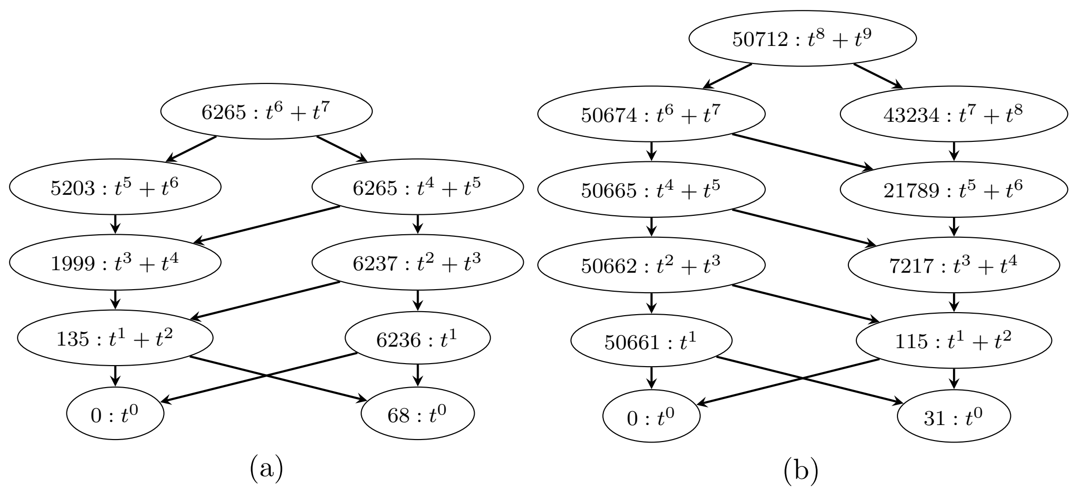

Figure 4.14: Reduced Conley Morse graph for (a) 4-fold cover and (b) 5-fold cover of **u**.

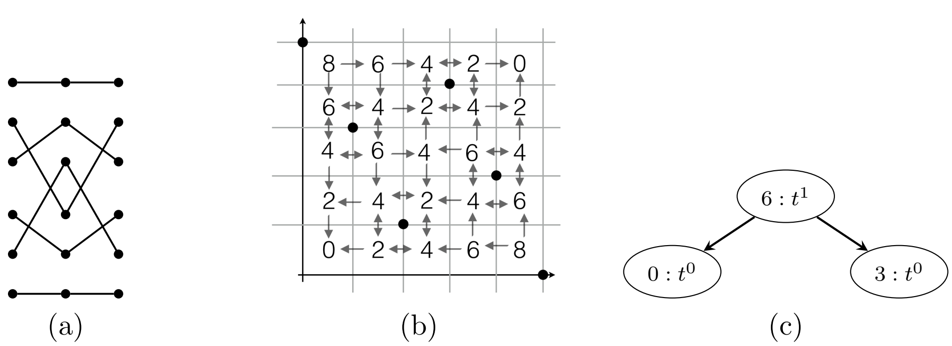

Figure 4.15: (a) Braid **v**. (b) Lap numbers. (c) Reduced Conley-Morse graph.

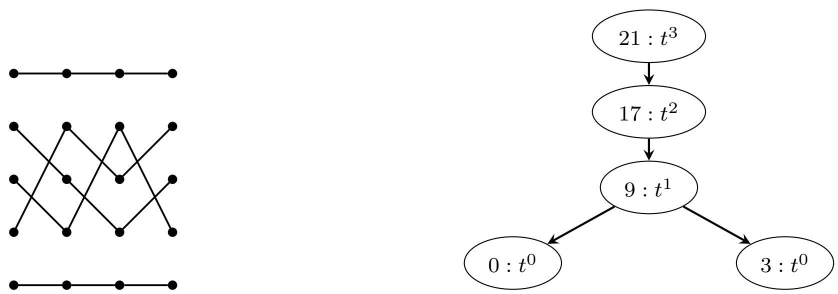

(a) (b)

Figure 4.17: (a) **w**. (b) Reduced Conley-Morse graph for **w**.

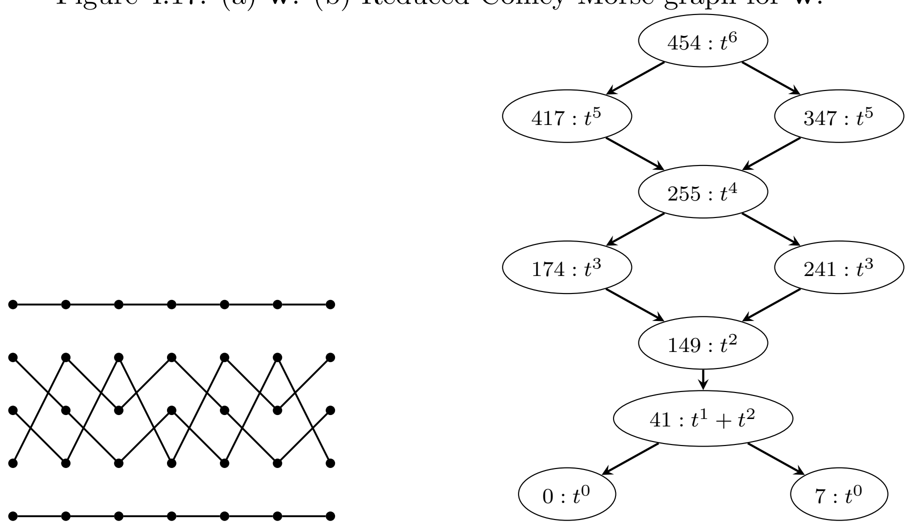

Figure 4.18: (a) 2-fold cover of **w**; Conley-Morse graph for 2-fold cover.

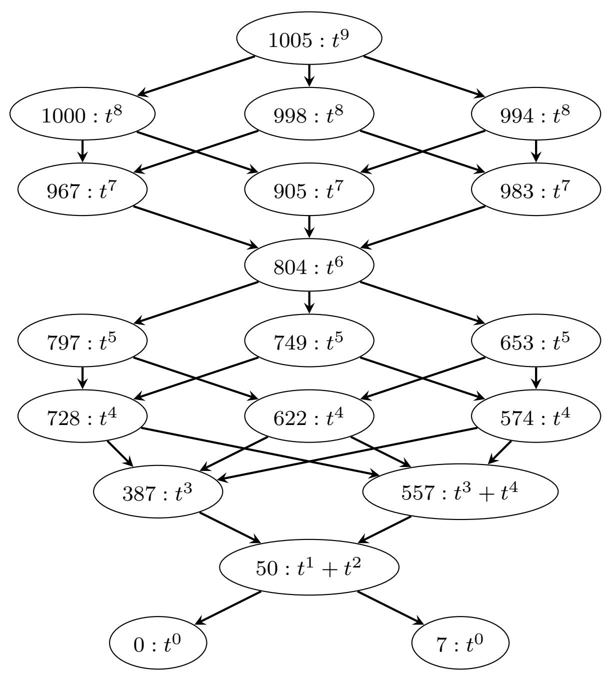

Figure 4.19: Reduced Conley-Morse Graph for 3-fold cover of $\mathbf{w}$.

Chapter 5

Elements of Connection Matrix Computation

5.1 Overview

In main goal of this chapter is to describe a modification to the algorithm for computing the connection matrix, which is applicable when computing connection matrices in the setting of cubical complexes. The modification is to use an implicit scheme for the discrete Morse theory, adapted to the cubical setting, enabling connection matrix computations in high-dimensions (e.g., $d = 9, 10$).

This is quite effective, as our computational results – in Section 5.4 – demonstrate. In fact, the cubical Morse theory we describe decouples from the connection matrix algorithm, and provides a very effective way to compute the homology of high-dimensional cubical complexes. See [5] for some examples of homology computations of cubical configuration spaces.

A few remarks are in order. First, our emphasis on cubical complexes is due to the fact that these arise naturally within computational dynamics, most often as grids on the phase space. Second, we introduce our principles in a more general setting than that of cubical complexes. We do this as we anticipate extending the results of implicit matchings to other settings, e.g., simplicial complexes and the order complexes of cubical complexes.

5.2 Principles of Partial Matchings

Our algorithm for computing a connection matrix (Algorithm 3.7.8 of Section 3.7.4) relies on discrete Morse theory (see Section 2.8). A naive implementation would require storage of both the complex $\mathcal{X}$ and the tower of partial matchings (see Section 3.7). In many situations, this would require a significant or infeasible amount of memory.

The contributions of this chapter are twofold. First, we maneuver around the memory issue with *implicit* partial matchings, that is, partial matchings that may be expressed as a formula. In this case the entire partial matching does not need to be stored. This is developed in Sections 5.2.1 and 5.3. Second, we give a general scheme which takes as input a sequence of matchings $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ and derives a new sequence of partial matchings $\boldsymbol{\omega} = (w_1, \ldots, w_n)$, which incorporates the collective matching information in $\boldsymbol{\alpha}$.

In practice, we use implicit matchings adapted to (graded) cubical complexes. The benefit is twofold: first, as the matching is implicit, the algorithm for computing the first Morse complex requires low memory usage; second, in experiments we find that the implicit matching scheme executes quickly, see Section 5.4. Within the connection matrix framework, the Morse theory described in this section is typically the very first part of the computation. In general, after one round of Morse theory, the complex is no longer cubical and the algorithm described in this section is no longer applicable. See Section 5.4 for some results on the efficacy of this approach.

5.2.1 Example of Implicit Matching

We begin with an example of an implicit matching on a cubical complex. Let $\mathcal{X}$ be a cubical complex in $\mathbb{R}^n$ with an ordering on the coordinates of the complex (see Section 2.6). Without loss of generality, we take the ordering $(1, 2, \ldots, n)$. Fix a coordinate i. We create a partial matching α_i. Any cell $\xi \in \mathcal{X}$ may be written as

$$\xi = I_1 \times \ldots \times I_i \times \ldots \times I_n.$$

Either ξ has extent in i or it does not (see Definition 2.6.2). If it does have extent i, then $I_i = [l, l+1]$. In this case ξ has a *left boundary*, ξ' (which may or may not belong to $\mathcal{X}$), given by

$$\xi' = I_1 \times \ldots \times [l, l] \times \ldots \times I_n.$$

If ξ does not have extent, then $I_i = [l, l]$. In this case, ξ has a *right coboundary*, ξ' (which may or may not be in $\mathcal{X}$), given by

$$\xi' = I_1 \times \ldots \times [l, l+1] \times \ldots \times I_n.$$

We define the following function $\mathcal{M}_i$ from cubes to cubes:

$$\mathcal{M}_i(\xi) = \begin{cases} I_1 \times \ldots \times [l, l+1] \times \ldots \times I_n & \text{if } I_i = [l, l] \\ I_1 \times \ldots \times [l, l] \times \ldots \times I_n & \text{if } I_i = [l, l+1]. \end{cases} \tag{5.1}$$

For a cube ξ we refer to $\mathcal{M}_i(\xi)$ as the *mate of ξ in coordinate i*. Let α_i be the following acyclic partial matching (see Definition 2.8.1) associated to the function $\mathcal{M}_i$:

$$\alpha_i(\xi) = \begin{cases} \mathcal{M}_i(\xi) & \text{if } \mathcal{M}_i(\xi) \in \mathcal{X} \\ \xi & \text{otherwise.} \end{cases} \tag{5.2}$$

This procedure describes a set $\{\alpha_i\}$ of acyclic partial matchings indexed by the coordinates. The ordering on coordinates imposes an ordering (*preference*) on the acyclic partial matchings. Thus we have a sequence (α_i) of acyclic partial matchings, and for a cell $\xi \in \mathcal{X}$ the set $\{\alpha_i(\xi)\}$ holds the possible mates, and the ordering of the α_i gives a preference on the possible mates. Using the sequence (α_i) we construct a final acyclic partial matching, denoted w, which uses the preference to match a cell ξ with its most preferred mate $w(\xi) \in \{\alpha_i(\xi)\}$, subject to the condition that $w(\xi)$ does not have a more preferred mate (within $\{\alpha_i(w(\xi))\}$) to ξ itself. If ξ is unmatched, then ξ is a critical cell. This is formalized in Section 5.3. For the sequence (α_1, α_2), w is shown in Figure 5.1. Recall that given a partial matching (A, w), a pair (ξ, ξ') with $\xi \in Q$, $\xi' \in K$ and $w(\xi) = \xi'$ is visualized with a directed edge $\xi \to \xi'$. See [14, 22] and also Section 3.7.

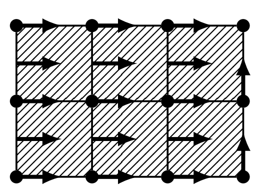

Figure 5.1: Example of the implicit matching w on 2D cubical complex.

5.2.2 Principles of Matchings

Let $\mathcal{X}$ be a cell complex. The incidence function κ engenders a binary relation which we call the *incidence relation*.

Definition 5.2.1. Let $\mathcal{X}$ be a cell complex. The *incidence relation* is the relation denoted by $\prec$ given as

$$\xi' \prec \xi \text{ if and only if } \kappa(\xi, \xi') \neq 0.$$

An equivalent formulation of the notion of partial matching (Definition 2.8.1) is an involution $w\colon \mathcal{X} \to \mathcal{X}$ which is subject to the following trichotomy. Given any $\xi \in \mathcal{X}$, either:

1. $w(\xi) = \xi$, or
2. $\xi \prec w(\xi)$, or
3. $w(\xi) \prec \xi$.

That is, a partial matching is an involution in which the image of a cell is either itself or a primary face/coface. In this case the decomposition $\mathcal{X} = A \sqcup Q \sqcup K$ is given as

$$A := \{\xi : \xi = w(\xi)\} \qquad Q := \{\xi : \xi \prec w(\xi)\} \qquad K := \{\xi : w(\xi) \prec \xi\}.$$

That is, A is recovered as the set of cells fixed under w, denoted $Fix(w)$, and w restricts to a bijection $w|_Q\colon Q \to K$.

Definition 5.2.2. Let w be a partial matching. A *discrete flowline*, or simply *flowline*, is a sequence of cells

$$\xi_0 \prec w(\xi_0) \succ \xi_1 \prec w(\xi_1) \prec \ldots \succ \xi_n \prec w(\xi_n).$$

Flowlines are sometimes called *zigzags paths*, where $\prec$ are the 'zigs' and $\succ$ are the 'zags'. Flowlines are used to derive the relation $\xi' \gg \xi$ on Q (see Definition 2.8.1). Recall from Definition 2.8.1 that a partial matching is *acyclic* if the transitive closure of $\gg$ is a partial order on Q.

For the remainder of this section let $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ be a sequence of partial matchings.

Definition 5.2.3. A partial matching $w\colon \mathcal{X} \to \mathcal{X}$ is $\boldsymbol{\alpha}$-*compatible* if for all $\xi \in \mathcal{X}$ there is an i such that $w(\xi) = \alpha_i(\xi)$.

As each α_i is an $\boldsymbol{\alpha}$-compatible matching, the set of $\boldsymbol{\alpha}$-compatible matchings is nonempty. If w is an $\boldsymbol{\alpha}$-compatible partial matching then there is a relation $\gg$ on Q and any sequence $\xi_0 \gg \xi_1$ lifts to a flowline

$$\xi_0 \overset{i}{\prec} w(\xi_0) \succ \xi_1 \overset{j}{\prec} w(\xi_1)$$

where

$$w(\xi_0) = \alpha_i(\xi_0) \text{ and } w(\xi_1) = \alpha_j(\xi_1).$$

That is, we can label the 'zigs' (denoted with $\prec$) with the appropriate index. A notion stronger than compatibility is that of *stability* – a matching which obeys the preference. Stability forces the flowlines to obey a particular forcing condition, which becomes useful when proving acyclicity, see Proposition 5.2.11.

Definition 5.2.4. Let w be an $\boldsymbol{\alpha}$-compatible matching. Consider $\xi_0 \gg \xi_1$ with associated flowline

$$\xi_0 \overset{j}{\prec} w(\xi_0) \succ \xi_1 \overset{j'}{\prec} w(\xi_1).$$

The pair $\xi_0 \gg \xi_1$ is $\boldsymbol{\alpha}$-*unstable* if there exists i such that

1. $w(\xi_0) = \alpha_i(\xi_1)$, and
2. $i < j$ and $i < j'$.

A matching is called $\boldsymbol{\alpha}$-*stable* if there are no $\boldsymbol{\alpha}$-unstable pairs.

Proposition 5.2.5. *For any sequence $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ of partial matchings there is an $\boldsymbol{\alpha}$-stable matching w.*

Proof. We inductively derive a sequence $(w_0, \ldots, w_n)$ as follows. First, let $w_0 := \mathrm{id}$. For $i \in \{1, \ldots, n\}$, let $A_{i-1} = Fix(w_{i-1})$ and

$$w_i(\xi) = \begin{cases} \alpha_i(\xi) & \xi, \alpha_i(\xi) \in A_{i-1} \\ w_{i-1}(\xi) & \text{otherwise.} \end{cases} \tag{5.3}$$

The matching of interest is

$$w := w_n. \tag{5.4}$$

We now show that w is a stable partial matching. We begin by proving that that w is a partial matching. We proceed by induction on the sequence $\{w_0, \ldots, w_n\}$. The base case is $w_0 = \mathrm{id}$, which is a partial matching. Now we assume that w_{i-1} is a partial matching and show that w_i is also a partial matching.

This amounts to showing two things: first, that w_i is an involution, and second, that it obeys the incidence trichotomy. By the inductive hypothesis we have that $w_{i-1}^2 = \mathrm{id}$. If $\xi \notin A_{i-1}$ or $\alpha_i(\xi) \notin A_{i-1}$ then $w_i^2(\xi) = w_{i-1}^2(\xi) = \mathrm{id}$. Otherwise

$$w_i(w_i(\xi)) = w_i(\alpha_i(\xi)) = \alpha_i^2(\xi) = \mathrm{id}.$$

If $\xi \notin A_{i-1}$ or $\alpha_i(\xi) \notin A_{i-1}$ then $w_i(\xi) = w_{i-1}(\xi)$ and, as w_{i-1} is a partial matching by the inductive hypothesis, either $\kappa(\xi, w_i(\xi)) \neq 0$ or $\kappa(w_i(\xi), \xi)) \neq 0$. Otherwise, $\xi, \alpha_i(\xi) \in A_{i-1}$, and $w_i(\xi) = \alpha_i(\xi)$. As α_i is a partial matching by hypothesis, we have that $\kappa(\alpha_i(\xi), \xi)) \neq 0$ or $\kappa(\xi, \alpha(\xi)) \neq 0$. Therefore the incidence trichotomy holds and w_i is a partial matching.

It remains to show that w is stable. Suppose that w had an unstable pair $\xi_0 \gg \xi_1$ with associated flowline

$$\xi_0 \overset{j}{\prec} w(\xi_0) \overset{i}{\succ} \xi_1 \overset{j'}{\prec} w(\xi_1).$$

and $\alpha_i(\xi_1) = w(\xi_0)$ and $i < j, j'$. Let $\xi' = w(\xi_0)$ and $\xi'' = w(\xi_1)$. As $\alpha_i(\xi_1) = \xi'$ but $w(\xi_1) \neq \xi'$ it follows from Eqn. (5.3) that either $\xi_1 \notin A_{i-1}$ or $\xi' \notin A_{i-1}$. That $\xi' = w(\xi_0) = \alpha_j(\xi_0)$ implies $\xi_0, \xi' \in A_{j-1}$. That $\xi'' = w(\xi_1) = \alpha_{j'}(\xi_1)$ implies $\xi_1, \xi'' \in A_{j'-1}$. By Eqn. (5.3) there is a descending sequence

$$\mathcal{X} = A_0 \supset A_1 \supset \ldots \supset A_n.$$

Note that $i < j, j'$ implies $A_j, A_{j-1} \subset A_{i-1}$. Therefore $\xi_1, \xi' \in A_i$, a contradiction. Therefore w is stable. $\square$

Remark 5.2.6. Stable partial matchings need not be unique. Here is an example.

The construction of the stable matching w as defined in Eqn. (5.4) can be expressed as the recursive algorithm, Mate.

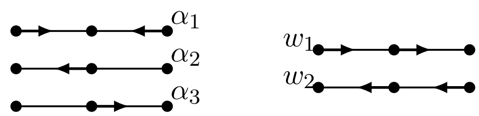

Figure 5.2: Sequence of matchings $\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \alpha_3)$. Both w_1 and w_2 are stable matchings with respect to α.

function MATE(ξ)
  return MATEHELPER(ξ, n)
end function
function MATEHELPER(ξ, i)
  if $i = 0$ **then**
    return ξ
  end if
  $\xi' \leftarrow$ MATEHELPER($\xi, i-1$)
  if $\xi' = \xi$ **then**
    if MATEHELPER($\alpha_i(\xi), i-1$) $= \alpha_i(\xi)$ **then**
      return $\alpha_i(\xi)$
    end if
  end if
  return ξ'
end function

Proposition 5.2.7. *Let $\mathcal{X}$ be a cell complex and let $\boldsymbol{\alpha} = (\alpha_i)$ be a sequence of partial matchings. Then $w(\cdot) :=$ MATE$(\cdot)$ is an $\boldsymbol{\alpha}$-stable partial matching.*

Proof. Set $w_i(\cdot) :=$ MATEHELPER$(\cdot, i)$. Notice that $w_0 = \mathrm{id}$. It follows from the proof of Proposition 5.2.5 that w is an $\boldsymbol{\alpha}$-stable partial matching. □

Proposition 5.2.8. *Let $\mathcal{X}$ be a cell complex and $\xi \in \mathcal{X}$. The algorithm* MATE(ξ) *executes in $O(2^n)$ time.*

Proof. We prove this by induction on n. In the worst case, the two **if** statements are true at each level of recursion. We define a function T for the complexity of the algorithm in the worst case. For $n = 0$, ξ is returned. Set $T(0) = c_1$. and $T(n) = c_2 + 2T(n-1)$. The closed form of this recurrence relation is $T(n) = 2^n c_1 + c_2 \sum_{i=0}^{n-1} 2^i = 2^n(c_1 + c_2) - c_2 = O(2^n)$. $\square$

5.2.3 Establishing Acyclicity

Even if each partial matching α_i in the sequence $\boldsymbol{\alpha}$ is acyclic, an $\boldsymbol{\alpha}$-stable matching need not be acyclic. In this section we review the notion of a Lyapunov function. This will be a main tool for partially establishing acyclicity of w as given in Eqn. 5.4. To establish acyclicity, it will remain to show a local acyclicity condition – namely, that the matching is acyclic on the fibers (level-sets) of the Lyapunov function.

Definition 5.2.9. A *Lyapunov function* for a partial matching $w\colon \mathcal{X} \to \mathcal{X}$ is a poset morphism $f\colon (\mathcal{X}, \leq) \to \mathbb{Z}$ such that $f(w(\xi)) = f(\xi)$ for each $\xi \in \mathcal{X}$. We say that w is *globally acyclic* if it has a Lyapunov function f.

Given any sequence $\xi_0 \gg \xi_1 \gg \ldots \gg \xi_n$ there is a corresponding a flowline

$$\xi_0 \prec \xi_0' \succ \xi_1 \prec \xi_1' \succ \ldots \succ \xi_n$$

where $w(\xi_i) = \xi_i'$. Application of the Lyapunov function f yields:

$$f(\xi_0) = f(\xi_0') \geq f(\xi_1) = f(\xi_1') \geq \ldots \geq f(\xi_n).$$

The image under f of any flowline is a (weakly) decreasing sequence, hence the terminology Lyapunov function.

Proposition 5.2.10. *Let $\boldsymbol{\alpha} = (\alpha_i)$ be a sequence of partial matchings. If f is a Lyapunov function for each α_i then f is a Lyapunov function for w, where w is given by Eqn.* (5.3).

Proof. We prove this by induction on the sequence $(w_0, w_1, \ldots, w_n = w)$. The base case is that $w_0 = \mathrm{id}$. By hypothesis f is a poset morphism, therefore it is a Lyapunov function for w_0. Now assume that f is a Lyapunov function for w_{i-1}. We show that f is a Lyapunov function for w_i. Let $\xi \in \mathcal{X}$. As f is a Lyapunov function for α_i we have that $f(\xi) = f(\alpha_i(\xi))$. If $\xi \in A_{i-1}$ and $\in A_{i-1}$ then $f(w_i(\xi)) = f(\alpha_i(\xi)) = f(\xi)$.

Otherwise, $f(w_i(\xi)) = f(w_{i-1}(\xi)) = f(\xi)$ by the inductive hypothesis. Therefore f is a Lyapunov function for w. $\square$

This technique establishes a Lyapunov function for the matching w. To establish that w is an acyclic partial matching, one needs an extra assumption that guarantees that w is acyclic on the fibers of f. If such a condition is met, then w is acyclic by the next result.

Proposition 5.2.11. *Let f be a Lyapunov function for a partial matching w. If w is acyclic on the fibers of f then w is an acyclic partial matching.*

Proof. We have $\mathcal{X} = \bigsqcup_{k\in\mathbb{Z}} f^{-1}(k)$. Let w^k be the restriction of w to $f^{-1}(k)$. By hypothesis w^k is acyclic for each $k \in \mathbb{Z}$. From the Patchwork Theorem [29, Patchwork Theorem], the disjoint union of acyclic partial matchings $\{w^k\}$ on the fibers of f is an acyclic partial matching on $\mathcal{X}$. $\square$

5.3 Implicit Matching for Cubical Complexes

Let $\mathcal{X}$ be a cubical complex in $\mathbb{R}^n$. Recall the definition of $\mathcal{M}_i$ and α_i from Eqn. (5.1) and Eqn. (5.2), respectively. Let $\boldsymbol{\alpha} = (\alpha_i)$. It follows from Proposition 5.2.7 that $w(\cdot) := \textsc{Mate}(\cdot)$ is an $\boldsymbol{\alpha}$-stable partial matching.

Proposition 5.3.1. *The partial matching $w(\cdot) = \textsc{Mate}(\cdot)$ is globally acyclic.*

Proof. Global acyclicity is established by furnishing a Lyapunov function $f\colon (\mathcal{X}, \leq) \to \mathbb{Z}$ for w. For any cube ξ we have

$$\xi = I_1 \times \ldots \times I_n$$

where $I_i = [l_i, l_i]$ or $I_i = [l_i, l_i + 1]$. Define

$$f(\xi) = -\sum_{i=1}^{n} l_i.$$

We first show that f is a poset morphism. It suffices to show that $f(\xi) \leq f(\xi')$ for any pair $\xi \prec \xi'$, i.e., where ξ is a primary face of ξ'. Let $\xi = I_1 \times \ldots \times I_n$ and $\xi' = I'_1 \times \ldots \times I'_n$. Since ξ is a primary face, ξ and ξ' agree on all intervals except one, call it I_j. In this case $I'_j = [l_j, l_j+1]$ is nondegenerate while I_j is degenerate. Either $I_j = [l_j, l_j]$ or $I_j = [l_j+1, l_j+1]$. In either case we have that $f(\xi) \leq f(\xi')$.

Now for any i we have $f(\xi) = f(\mathcal{M}_i(\xi))$. This implies that $f(\xi) = f(\alpha_i(\xi))$ and that f is a Lyapunov function for each α_i. Therefore f is a Lyapunov function for w from Proposition 5.2.10. □

It remains to establish acyclicity on the fibers of f. For the moment, consider the case that $\mathcal{X}$ is the cubical complex in $\mathbb{R}^d$ consisting of every elementary cube, and recall the notion of interval from Definition 2.4.5. In this case, a fiber of $f\colon (\mathcal{X}, \leq) \to \mathbb{Z}$ is a disjoint union $\bigsqcup H_j$ of intervals $H_j \subset \mathcal{X}$. By an appropriate translation to the origin, each interval H_j is of the form

$$H_j = [0, a_1] \times [0, a_2] \times \ldots \times [0, a_n] \qquad \text{where } a_i \in \{0, 1\}.$$

Setting $a := (a_1, \ldots, a_n) \in \{0,1\}^n$, it may be seen that each H_j is in fact a hypercube in the face poset $\mathcal{X}$.

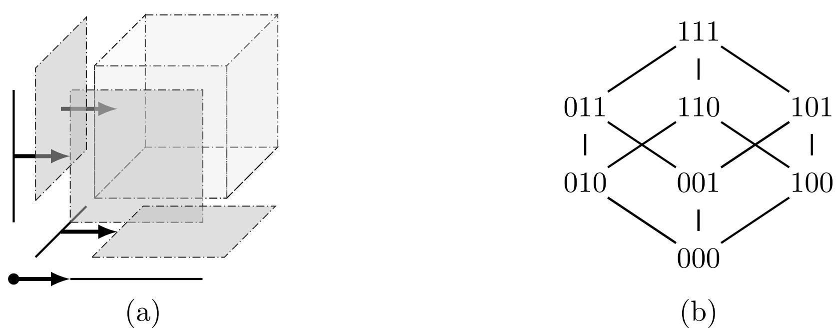

Figure 5.3: (a) Example of H and w in cubical complex. (b) Associated hypercube H in the face poset formed by the sequences (a_1, a_2, a_3) with $a_i \in \{0, 1\}$.

In the case of an arbitrary cubical complex $\mathcal{X}$ in $\mathbb{R}^n$, a fiber of the Lyapunov function $f\colon (\mathcal{X}, \leq) \to \mathbb{Z}$ will consist of a subset of this disjoint union of hypercubes.

Proposition 5.3.2. *w is acyclic.*

Proof. Propostion 5.3.1 implies that w is globally acyclic. By Proposition 5.2.11, it remains to show that w is acyclic on the fibers of f. Without loss of generality, we may restrict to the subset of the hypercube H

$$\{\xi = [0, a_1] \times \ldots [0, a_n] \in \mathcal{X} : a_i \in \{0, 1\}\} \subset \mathcal{X}.$$

Now given a sequence $\xi_0 \gg \xi_1 \gg \ldots \gg \xi_n \gg \xi_0$ there is a corresponding a zig-zag in the hypercube H

$$\xi_0 \overset{i_0}{\prec} w(\xi_0) \overset{j_0}{\succ} \xi_1 \overset{i_1}{\prec} w(\xi_1) \overset{j_1}{\succ} \ldots \overset{j_{n-1}}{\succ} \xi_n \overset{i_n}{\prec} w(\xi_n) \overset{j_n}{\succ} \xi_0.$$

Here every zig and zag is labeled with the appropriate coordinate, i.e.,

$$\mathcal{M}_{i_k}(\xi_k) = w(\xi_k) = \mathcal{M}_{j_k}(\xi_{k+1}).$$

Let $i^* = \min\{i_0, \ldots, i_n\}$, i.e., the most preferred coodinate for the zig. Since the sequence is a loop in a hypercube, there must be a corresponding zag in the sequence through the coordinate i^*, i.e., there is some k where

$$\xi_k \overset{i_k}{\prec} w(\xi_k) \overset{i^*}{\succ} \xi_{k+1}.$$

However, this would imply that $i_k < i^*$, which is a contradiction. Therefore there can be no such sequence, implying that w is acyclic in the fibers of f. Therefore w is an acyclic partial matching by Proposition 5.2.11. □

5.3.1 Graded Implicit Matching

In the graded case, the input is a P-graded complex $(\mathcal{X}, \nu)$. We may adapt the algorithm Mate of Section 5.2.2 to the graded case by using a sequence of graded partial matchings; we distinguish this by calling the new algorithm Graded-Mate. Recall from Section 3.7.3 that a partial matching is *graded* if two cells are matched only if they belong to the same fiber. We adapt our sequence α_i from Section 5.3 to the graded case as follows:

$$\beta_i(\xi) = \begin{cases} \mathcal{M}_i(\xi) & \text{if } \nu(\xi) = \nu(\mathcal{M}_i(\xi)) \\ \xi & \text{otherwise.} \end{cases}$$

The algorithm Graded-Mate is precisely the algorithm Mate with the sequence of graded partial matchings $\boldsymbol{\beta} = (\beta_i)$.

Proposition 5.3.3. *Let $(\mathcal{X}, \nu)$ be a P-graded cubical complex in $\mathbb{R}^n$. The function $w(\cdot) =$ Graded-Mate$(\cdot)$ is a P-graded acyclic partial matching.*

Proof. The fibers $\mathcal{X}^p = \nu^{-1}(p)$ partition the underlying set $\mathcal{X}$ as $\mathcal{X} = \bigsqcup \mathcal{X}^p$. Each fiber $\mathcal{X}^p$ may be regarded as a subcomplex by restricting $(\leq, \kappa, \dim)$ to $\mathcal{X}^p$. The algorithm of Graded-Mate is equivalent to applying Mate on each fiber $\mathcal{X}^p$. It follows from Propositions 5.2.7 and 5.3.2 that this furnishes an acyclic partial matching for each fiber $\mathcal{X}^p$. It follows from the Patchwork Theorem [29, Patchwork Theorem], that taken together these give an acyclic partial matching for the entire cell complex $\mathcal{X}$. $\square$

5.4 Computational Experiments

Our computational experiments are drawn from applications of the computational connection matrix theory to the Morse theory on braids (see Chapter 4). This theory provides a nice set of scalable and relevant (cubical) examples. For the purposes of this section we are interested in the performance of our algorithms. Therefore, we will briefly recall the setup.

We will also use this section to show the visuals one gets from the computation, which we call the *Conley-Morse graph*, see Example 1.2.2. In this case the f-polynomial is replaced with the f-vector, see Definition 2.5.16.

5.4.1 A Morse Theory on Braids Redux

Recall from Section 4.5, that a braid $\mathbf{v}$ has two parameters: the dimension – d – and the number of strands – m. These two parameters suffice to determine the cubical complex $\mathcal{X} = \mathcal{X}(n, m)$ on the n-cube $[0, m+1]^n$. Recall that cross: $\mathcal{X}^+ \to \mathbb{Z}$ is a map on the top-cells of $\mathcal{X}$; cross is also a Lyapunov function for a parabolic relation which fixes $\mathbf{v}$.

Recall that a directed graph $\mathcal{F}_0$ is constructed with edge set $\mathcal{X}^+$, where $\xi \to \xi'$ if $\mathrm{cross}(\xi) \geq \mathrm{cross}(\xi')$. A transversality model is made by coarsening this to a relation $\mathcal{F}$ by handling the improperness (see Definition 4.5.3). We are interested in $\mathsf{SC}(\mathcal{F})$ the strongly connected components of $\mathcal{F}$. Since every top-cell ξ is a vertex in the graph $\mathcal{F}$, it must belong to some strongly connected component. This furnishes a map $\nu_0 \colon \mathcal{X}^+ \to \mathsf{SC}(\mathcal{F})$ that takes any top-cell to the strongly connected component in which it belongs. It follows from the results in Section 4.5 that the map $\nu \colon \mathcal{X} \to \mathsf{SC}(\mathcal{F})$ given

(a) (b)

Figure 5.4: (a) 2D cubical complex and cross: $\mathcal{X}^+ \to \mathbb{N}$. Each point is an improper strand $v \in G_{\mathbf{v}}$. Directed graph $\mathcal{F}$ is depicted with vertex set the collection of 2-cells. (b) The Hasse diagram of the poset $\mathsf{SC}(\mathcal{F})$. Each node p is labeled with the f-vector of the fiber over p.

(a) (b)

Figure 5.5: (a) Conley complex for Figure 5.4. Blank nodes have no cells in the fiber. (b) $\mathsf{RC}(\mathcal{F})$-graded complex.

by

$$\nu(\xi) = \min_{\mathsf{SC}(\mathcal{F})} \{\nu_0(\eta) \colon \eta \in \operatorname{star}(\xi) \cap \mathcal{X}^+\} \in \mathsf{SC}(\mathcal{F})$$

is a well-defined poset morphism $\nu \colon \mathcal{X} \to \mathsf{SC}(\mathcal{F})$. Therefore $(\mathcal{X}, \nu)$ is a $\mathsf{SC}(\mathcal{F})$-graded cubical complex. This is shown in Figure 5.4. The graded complex $(\mathcal{X}, \nu)$ is the input into Algorithm 3.7.8 (ConnectionMatrix) of Section 3.7. The output (A, Δ, μ) is a Conley complex and Δ associated connection matrix. This is visualized in Figure 5.5(a).

Define

$$\mathsf{RC}(\mathcal{F}) := \{[\xi] \in \mathsf{SC}(\mathcal{F}) : |\mu^{-1}([\xi])| \neq 0\}.$$

In words, $\mathsf{RC}(\mathcal{F}) \subset \mathsf{SC}(\mathcal{F})$ is the subset of strongly connected components that have non-empty fiber. An $\mathsf{RC}(\mathcal{F})$-graded cell complex is formed by restriction of μ to its image, see Figure 5.5(b).

5.4.2 Experimental Results

The specific examples we give to benchmark the performance of the algorithm are described in more detail, along with their Conley-Morse graphs, in Section 4.7.

Given a $\mathsf{SC}(\mathcal{F})$-graded cell complex, we record the following parameters, which we term 'initial data': the top-dimension of the cubical complex, $\#\mathcal{X}^+$ (the number of top-dimensional cubes), and $\#\mathcal{X}$ (the total number of cells). Finally, we give $\#\mathsf{SC}(\mathcal{F})$, the size of the poset of strongly connected components.

We record the results of two algorithms: first, the cubical-Morse theoretic algorithm – Graded-Mate – which gives an initial graded Morse complex. We record its size of the Morse complex and the f-vector (that is, number of cells in each dimension). We also record the execution time of this algorithm.

Applying Algorithm 3.7.8 (ConnectionMatrix) to the graded complex $(\mathcal{X}, \nu)$ produces a Conley complex, which can be viewed as a new $\mathsf{SC}(\mathcal{F})$-graded cell complex (A, μ). The cubical Morse theory is only the initial graded reduction (see Section 3.7.3) within the tower of reductions associated to the execution of the algorithm (see Diagram 3.21). For the results of the computational experiments, we record the following parameters of the Conley complex: the number of cells in the Conley complex, the size of the tower of reductions (including the initial complex and Conley complex), and the time elapsed. We record $\#\mathsf{RC}(\mathcal{F})$, the cardinality of $\mathsf{RC}(\mathcal{F})$.

The time elapsed data come from iPython's %timeit 'magic', which displays the average of the best 3 runs out of 10 runs. All testing was done on a Intel(R) Xeon E5-2680 v3 2.50GHz CPU with 128 GB RAM. This is part of Rutgers University's Perceval cluster.

5.4.3 Braids Studied

n-fold covers

We examine the n-fold covers of $\mathbf{u}, \mathbf{v}, \mathbf{w}$ as done in Section 4.7.

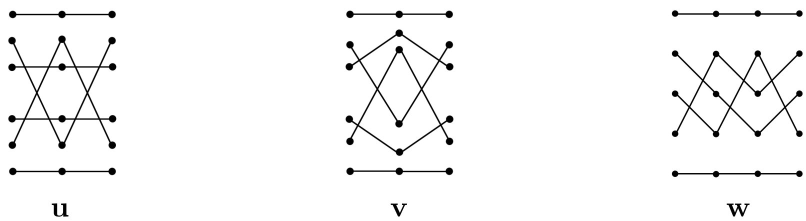

Figure 5.6: Braids **u**, **v** and **w**.

Torus Knot

Finally, we wish to examine how the algorithm scales with the number of subdivisons, of the n-cube, i.e., the number of strands in the braid. To do this, study a braid $\mathbf{t}$ which corresponds to a torus knot. In this case we fix $\mathbf{t}$ to be four-dimensional and vary the number of strands m.

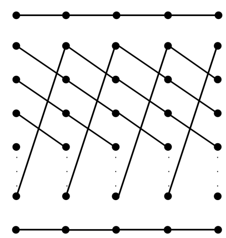

Figure 5.7: Schematic for 4-D braid $\mathbf{t}$ on m strands.

Regardless of the number of strands, the pattern of the torus remains the same: for any strand $\mathbf{t}^\alpha$, we have that the $\mathbf{t}^\alpha_{i+1}$ is obtained from $\mathbf{t}^\alpha_i$ via

$$\mathbf{t}^\alpha_{i+1} = \begin{cases} \mathbf{t}^\alpha_i - 1 & \mathbf{t}^\alpha_i > 1 \\ m & \mathbf{t}^\alpha_i = 1 \ . \end{cases}$$

We append two constant strands to $\mathbf{t}$ to act as boundary conditions.

5.4.4 Results

Initial Data for n-fold covers of $\mathbf{u}$				
	Dim	# Top Cells	# Cells	$\#\mathsf{SC}(\mathcal{F})$
$\mathbf{u}$	2	25	122	13
2-fold cover	4	625	15368	121
3-fold cover	6	15625	1992992	903
4-fold cover	8	390625	264990848	6747
5-fold cover	10	9765625	35958682112	51755

Table 5.8: Data for n-fold covers of $\mathbf{u}$.

Data for implicit Morse scheme on n-fold cover		
n	$\# M$	Time Elapsed
1	13	432 μs
2	361	60.7 ms
3	7969	8.74 s
4	166593	23min 18s
5	3417222	83 hr

Table 5.9: Experimental results for cubical Morse theory for $\mathbf{u}$.

Data for implicit Morse scheme on n-fold cover	
n	f-vector
1	$(6, 6, 1)$
2	$(46, 124, 126, 56, 9)$
3	$(321, 1332, 2361, 2280, 1257, 372, 46)$
4	$(2206, 12248, 30428, 44168, 40934, 24760, 9520, 2120, 209)$
5	(15126,105030,334260,642290,825440,741426,471290,209210,62000,11060,901)

Table 5.10: Experimental results for cubical Morse theory for $\mathbf{u}$.

Connection Matrix Data for n-fold covers of $\mathbf{u}$				
	# Cells	$\#\mathsf{RC}(\mathcal{F})$	# Tower	Time Elapsed
$\mathbf{u}$	3	3	3	475μs
2-fold	7	5	3	60.6ms
3-fold	11	7	3	8.72s
4-fold	15	9	3	23min 23s
5-fold	19	11	4	86 hr

Table 5.11: Computational results for n-fold covers of $\mathbf{u}$.

Initial Data for n-fold covers of $\mathbf{v}$				
	Dim	# Top Cells	# Cells	$\#\mathsf{SC}(\mathcal{F})$
$\mathbf{u}$	2	25	122	13
2-fold	4	625	15368	114
3-fold	6	15625	1992992	879
4-fold	8	390625	264990848	7212
5-fold	10	9765625	35958682112	62157

Table 5.12: Data for n-fold covers of $\mathbf{v}$.

Data for implicit Morse scheme on n-fold cover		
n	$\# M$	Time Elapsed
$\mathbf{v}$	7	311 μs
2	133	42.2 ms
3	1825	6.38 s
4	23281	21min 9s
5	291017	72.9 hr

Table 5.13: Experimental results for cubical Morse theory for $\mathbf{v}$.

Data for implicit Morse scheme on n-fold cover	
n	f-vector
$\mathbf{v}$	$(4, 3)$
2	$(20, 46, 43, 20, 4)$
3	$(88, 321, 525, 495, 285, 96, 15)$
4	$(380, 1884, 4322, 5988, 5483, 3408, 1412, 360, 44)$
5	(1649,10305,30255,54910,68235,60743,39320,18340,5930,1210,120)

Table 5.14: Experimental results for cubical Morse theory for $\mathbf{v}$.

Connection Matrix Data for n-fold covers of $\mathbf{v}$				
	# Cells	$\#\,\mathsf{RC}(\mathcal{F})$	# Tower	Time Elapsed
$\mathbf{u}$	3	3	3	358 μs
2-fold	33	32	3	50.6 ms
3-fold	197	196	3	7.55 s
4-fold	1155	1153	3	22min 37s
5-fold	6727	6725	3	77.2hr

Table 5.15: Computational results for n-fold covers of $\mathbf{v}$.

Initial Data for **w**.				
	Dim	# Top Cells	# Cells	#$\mathsf{SC}(\mathcal{F})$
u	3	64	756	15
2-fold cover	6	4096	631072	136
3-fold cover	9	262144	567108864	1005

Table 5.16: Data for n-fold covers of **w**.

Data for implicit Morse scheme			
n	# M	f-vector	Time Elapsed
1	27	$(7, 11, 7, 2)$	3.17 ms
2	713	$(45, 146, 215, 182, 93, 28, 4)$	2 s
3	18075	$(289, 1428, 3270, 4535, 4197, 2694, 1207, 372, 75, 8)$	46min 17s

Table 5.17: Experimental results for cubical Morse theory for **w**.

Connection Matrix Data for n-fold covers of **w**.				
	# Cells	# $\mathsf{RC}(\mathcal{F})$	# Tower	Time Elapsed
u	5	5	3	3.4ms
2-fold	11	10	3	2.24s
3-fold	21	19	3	54min 22s

Table 5.18: Computational results for n-fold covers of **w**.

Initial Data for **t**.			
#Strands	# Top Cells	# Cells	#$\mathsf{SC}(\mathcal{F})$
20	194481	3429896	38284
40	2825761	47499656	727824
60	13845841	228977416	3892564

Table 5.19: Data for n-fold covers of **t**.

Data for implicit Morse scheme			
# Strands	# M	f-vector	Time Elapsed
20	81	$(16, 32, 24, 8, 1)$	8.3 s
40	81	$(16, 32, 24, 8, 1)$	1min 54s
60	81	$(16, 32, 24, 8, 1)$	9min 7s

Table 5.20: Experimental results for cubical Morse theory for **t**.

Connection Matrix Data for n-fold covers of $\mathbf{t}$.				
# Strands	# Cells	# $\mathsf{RC}(\mathcal{F})$	# Tower	Time Elapsed
20	3	3	3	8.27 s
40	3	3	3	1min 37s
60	3	3	3	9min 3s

Table 5.21: Computational results for n-fold covers of $\mathbf{t}$.

References

[1] PyCHomP: Connection matrix software. `https://github.com/shaunharker/pyCHomP`.

[2] Z. Arai, W. Kalies, H. Kokubu, K. Mischaikow, H. Oka, and P. Pilarczyk. A database schema for the analysis of global dynamics of multiparameter systems. *SIAM Journal on Applied Dynamical Systems*, 8(3):757–789, 2009.

[3] M. Barakat and S. Maier-Paape. Computation of connection matrices using the software package conley. *International Journal of Bifurcation and Chaos*, 19(09):3033–3056, 2009.

[4] D. W. Barnes and L. A. Lambe. A fixed point approach to homological perturbation theory. *Proceedings of the American Mathematical Society*, 112(3):881–892, 1991.

[5] U. Bauer, M. Kahle, R. MacPherson, and K. Spendlove. Homological phases of matter and the configuration space of unit cubes in a rectangle. *In preparation.*

[6] J. Bush, M. Gameiro, S. Harker, H. Kokubu, K. Mischaikow, I. Obayashi, and P. Pilarczyk. Combinatorial-topological framework for the analysis of global dynamics. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(4):047508, 2012.

[7] G. Carlsson, G. Singh, and A. Zomorodian. Computing multidimensional persistence. *Journal of Computational Geometry*, 1(1), 2010.

[8] C. C. Conley. *Isolated invariant sets and the Morse index*. Number 38. American Mathematical Soc., 1978.

[9] B. A. Davey and H. A. Priestley. *Introduction to lattices and order*. Cambridge university press, 2002.

[10] S. Day, Y. Hiraoka, K. Mischaikow, and T. Ogawa. Rigorous numerics for global dynamics: A study of the Swift–Hohenberg equation. *SIAM Journal on Applied Dynamical Systems*, 4(1):1–31, 2005.

[11] H. Edelsbrunner and J. Harer. *Computational topology: an introduction*. American Mathematical Soc., 2010.

[12] S. Eilenberg and S. Mac Lane. On the groups $H(\pi, n)$, i. *Ann. of Math.(2)*, 58(1):55–106, 1953.

[13] A. Floer. Witten's complex and infinite-dimensional Morse theory. *Journal of differential geometry*, 30(1):207–221, 1989.

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

[15] R. Franzosa. Index filtrations and the homology index braid for partially ordered Morse decompositions. *Transactions of the American Mathematical Society*, 298(1):193–213, 1986.

[16] R. Franzosa. The continuation theory for Morse decompositions and connection matrices. *Transactions of the American Mathematical Society*, 310(2):781–803, 1988.

[17] R. Franzosa. The connection matrix theory for Morse decompositions. *Transactions of the American Mathematical Society*, 311(2):561–592, 1989.

[18] R. Franzosa and K. Mischaikow. Algebraic transition matrices in the Conley index theory. *Transactions of the American Mathematical Society*, 350(3):889–912, 1998.

[19] S. I. Gelfand and Y. I. Manin. *Methods of homological algebra*. Springer Science & Business Media, 2013.

[20] R. Ghrist and R. Vandervorst. Scalar parabolic pdes and braids. *Transactions of the American Mathematical Society*, 361(5):2755–2788, 2009.

[21] R. W. Ghrist, J. B. Van den Berg, and R. Vandervorst. Morse theory on spaces of braids and lagrangian dynamics. *Inventiones mathematicae*, 152(2):369–432, 2003.

[22] A. Gonzalez-Lorenzo, A. Bac, J.-L. Mari, and P. Real. Allowing cycles in discrete Morse theory. *Topology and its Applications*, 228:1–35, 2017.

[23] S. Harker, K. Mischaikow, M. Mrozek, and V. Nanda. Discrete Morse theoretic algorithms for computing homology of complexes and maps. *Foundations of Computational Mathematics*, 14(1):151–184, 2014.

[24] T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational homology*, volume 157. Springer Science & Business Media, 2006.

[25] W. D. Kalies, D. Kasti, and R. Vandervorst. An algorithmic approach to lattices and order in dynamics. *SIAM Journal on Applied Dynamical Systems*, 17(2):1617–1649, 2018.

[26] W. D. Kalies, K. Mischaikow, and R. Vandervorst. An algorithmic approach to chain recurrence. *Foundations of Computational Mathematics*, 5(4):409–449, 2005.

[27] W. D. Kalies, K. Mischaikow, and R. Vandervorst. Lattice structures for attractors i. *Journal of Computational Dynamics*, 1:307–338, 2014.

[28] W. D. Kalies, K. Mischaikow, and R. Vandervorst. Lattice structures for attractors ii. *Foundations of Computational Mathematics*, 16(5):1151–1191, 2016.

[29] D. Kozlov. *Combinatorial algebraic topology*, volume 21. Springer Science & Business Media, 2007.

[30] S. Lefschetz. *Algebraic topology*, volume 27. American Mathematical Soc., 1942.

[31] S. Mac Lane. *Categories for the working mathematician*, volume 5. Springer Science & Business Media, 2013.

[32] S. Maier-Paape, K. Mischaikow, and T. Wanner. Structure of the attractor of the cahn–hilliard equation on a square. *International Journal of Bifurcation and Chaos*, 17(04):1221–1263, 2007.

[33] C. McCord. Mappings and homological properties in the Conley index theory. *Ergodic Theory and Dynamical Systems*, 8(8*):175–198, 1988.

[34] C. McCord. Simplicial models for the global dynamics of attractors. *Journal of Differential Equations*, 167(2):316–356, 2000.

[35] C. McCord and K. Mischaikow. On the global dynamics of attractors for scalar delay equations. *Journal of the American Mathematical Society*, 9(4):1095–1133, 1996.

[36] C. McCord and J. Reineck. Connection matrices and transition matrices. *Banach Center Publications*, 47(1):41–55, 1999.

[37] K. Mischaikow. Global asymptotic dynamics of gradient-like bistable equations. *SIAM journal on mathematical analysis*, 26(5):1199–1224, 1995.

[38] K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete & Computational Geometry*, 50(2):330–353, 2013.

[39] S. Y. Oudot. *Persistence theory: from quiver representations to data analysis*, volume 209. American Mathematical Society Providence, RI, 2015.

[40] J. F. Reineck. The connection matrix in Morse-Smale flows. *Transactions of the American Mathematical Society*, 322(2):523–545, 1990.

[41] D. S. Richeson. *Connection matrix pairs for the discrete Conley index*. PhD thesis, Northwestern University, 1998.

[42] J. W. Robbin and D. A. Salamon. Lyapunov maps, simplicial complexes and the stone functor. *Ergodic Theory and Dynamical Systems*, 12(1):153–183, 1992.

[43] S. Roman. *Lattices and ordered sets*. Springer Science & Business Media, 2008.

[44] J. Rubio and F. Sergeraert. Constructive homological algebra and applications. *arXiv preprint arXiv:1208.3816*, 2012.

[45] E. Sköldberg. Morse theory from an algebraic viewpoint. *Transactions of the American Mathematical Society*, 358(1):115–129, 2006.

[46] E. Sköldberg. Algebraic Morse theory and homological perturbation theory. *arXiv preprint arXiv:1311.5803*, 2013.

[47] S. Smale. Differentiable dynamical systems. *Bulletin of the American mathematical Society*, 73(6):747–817, 1967.

[48] J. B. van den Berg, S. Day, and R. Vandervorst. Braided connecting orbits in parabolic equations via computational homology. *Journal of Differential Equations*, 255(6):1086–1108, 2013.

[49] J. B. van den Berg, R. Ghrist, R. Vandervorst, and W. Wójcik. Braid Floer homology. *Journal of Differential Equations*, 259(5):1663–1721, 2015.

[50] J. B. Van Den Berg, M. Kramár, and R. Vandervorst. The order of bifurcation points in fourth order conservative systems via braids. *SIAM Journal on Applied Dynamical Systems*, 10(2):510–550, 2011.

[51] C. A. Weibel. *An introduction to homological algebra.* Number 38. Cambridge university press, 1995.