

CAPTURING CHANGES IN COMBINATORIAL DYNAMICAL SYSTEMS VIA PERSISTENT HOMOLOGY

by


Ryan Slechta

A Dissertation

Submitted to the Faculty of Purdue University

In Partial Fulfillment of the Requirements for the degree of

Doctor of Philosophy



Department of Computer Science

West Lafayette, Indiana

May 2022

**THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF COMMITTEE APPROVAL**

Dr. Tamal K. Dey, Chair

Department of Computer Science

Dr. Saugata Basu

Department of Mathematics

Dr. Elena Grigorescu

Department of Computer Science

Dr. Paul Valiant

Department of Computer Science

Approved by:

Dr. Kihong Park

ACKNOWLEDGMENTS

In the autumn of 2012, I matriculated at the University of St. Thomas in St. Paul, Minnesota as an electrical engineering major. All engineering majors were required to take an “Introduction to Engineering” course where we did various simple, engineering-inspired projects. Our first project was to construct a water-balloon firing air cannon that could reliably hit a target. My group’s grade? A whopping 47/100 (we would have received 50/100, but my friend Zac lost us three points because he thought that he was too cool for safety glasses). As a consequence, I decided to switch to computer science. In the computer science department, one of my professors was an alumnus of The Ohio State University, which put OSU’s Ph.D. program on my radar. At OSU, I met and started working with Dr. Tamal Dey, who subsequently brought me to Purdue University. Herein lies the problem with acknowledgments: in acknowledging the chain of causality, where to terminate?

We will start with the proximate causes of this dissertation. The person who has had the most influence on the development of this dissertation is undoubtedly my advisor, Dr. Tamal Dey. Tamal’s zest for topology and algorithms has inspired me these past six years, and he has shown me more patience, instruction, and kindness than I deserve. Thank you, Tamal. It has been an honor.

I would also like to thank Dr. Marian Mrozek, who is a co-author of the three papers that make up this dissertation and who introduced me to the world of combinatorial dynamical systems. Over the past few years, he has in many ways become a second advisor. Marian’s intellect is exceeded only by his kindness, and he has been critical as a source of inspiration and knowledge when writing this dissertation.

For the past couple of years, I have had the privilege to work with Dr. Michał Lipiński on one of the papers in this dissertation. Michał has been an excellent collaborator and source of ideas. I look forward to our continued collaboration.

I would like to thank Dr. Saugata Basu and Dr. Elena Grigorescu for serving on my advisory committee, preliminary exam committee, and dissertation committee. They have both been exceptionally generous with their time, and I am grateful for the kindness they have shown me. I would also like to thank Dr. Alex Psomas and Dr. Paul Valiant for serving

on my preliminary exam committee and dissertation committee, respectively. Paul has also patiently helped me develop as an instructor, for which I am grateful.

I could not have made it through the years without the support of my fellow graduate students. Amongst the OSU clan, Dr. Alfred Rossi, Dr. Sayan Mandal, and Lucas Magee have all provided valuable academic, social, and moral support. Tao Hou, Soham Mukherjee, and Cheng Xin have been my comrades in the transition from Ohio State to Purdue. I'm grateful for the support they have given me over the years.

As an undergraduate, I had several advisors who helped me develop as a researcher. I am grateful to Dr. Nathan DeBardeleben, Dr. Laura Monroe, Dr. Peter Olver, Dr. Gerry Ruch, Dr. Jason Sawin, Dr. Cheri Shakiban, and Dr. Rob Thompson for tolerating my incompetence and showing me how it's done.

I would like to thank my friends, without whom I could not have done this. And finally, I'd like to thank my parents, who have been constant in their love and support. This dissertation is dedicated to them.

For the record, Zac left engineering, too (he got a Ph.D. in conspicuously-safety-glasses-free forest hydrology in 2020).

TABLE OF CONTENTS

LIST OF FIGURES	7
ABSTRACT	13
1 INTRODUCTION	14
2 PERSISTENT HOMOLOGY	17
3 COMBINATORIAL DYNAMICAL SYSTEMS	20
3.1 Multivector Fields and Combinatorial Dynamics	21
3.2 Conley Indices	28
3.3 Morse Decompositions and the Conley-Morse Graph	29
4 PERSISTENCE OF THE CONLEY INDEX	32
4.1 Noise-Resilient Index Pairs	40
4.1.1 Shrinking E	42
Computation	46
4.1.2 Enlarging P	50
4.2 A Simple Tracking Algorithm	53
4.2.1 Changing the Isolating Neighborhood	53
4.2.2 Finding Isolating Neighborhoods	58
5 PERSISTENCE OF CONLEY-MORSE GRAPHS	63
5.1 Conley-Morse Filtrations	63
5.2 Graph Filtrations	70
5.3 Barcodes for Conley-Morse Graphs	75
6 TRACKING ISOLATED INVARIANT SETS WITH CONTINUATION	81
6.1 Tracking Isolated Invariant Sets	81
6.2 Tracking via Continuation	85
6.2.1 Combinatorial Continuation and the Tracking Protocol	88

6.2.2	Characterizing Tracked Isolated Invariant Sets	92
6.3	Tracking via Persistence	93
6.3.1	From continuation to filtration	94
6.3.2	Tracking beyond continuation	97
6.3.3	Strategy for Step 2g of the Tracking Protocol	98
7	CONCLUSION	99
	REFERENCES	101
	VITA	104

LIST OF FIGURES

1.1	The behavior of the differential equation in Equation 1.1. From left-to-right, top-to-bottom, the depicted behavior is of the case when $\lambda = -10, 0, 1, 3, 6, 10, 15, 20, 25$	15
2.1	An example filtration and its associated barcode. Two 0-dimensional bars, in white, capture two connected components as they persist through the filtration. The 1-dimensional bar, in grey, captures the presence of the 1-dimensional cycle.	18
2.2	An example zigzag filtration and its associated barcode. The 0-dimensional bar, in white, captures the single connected component that persists through the entire filtration. Similarly, two 1-dimensional bars, in gray, capture the lifetime of the two 1-dimensional cycles, which are both present in the middle complex.	19
3.1	Three multivector fields. Multivectors are drawn by including a vector from each nonmaximal simplex in the multivector to each maximal simplex. In these three cases, each multivector has a unique maximal simplex. Multivectors that consist of a single simplex are marked with a circle.	22
3.2	Three examples of solutions, marked in yellow, over the same multivector field. The set on the left is the image of a solution under the given multivector field because one can spend infinitely long in the periodic orbit in both the positive and negative directions. For the set in the middle, one spends infinitely long in the triangle that is marked with a circle, then follows a path to the periodic orbit, and spends infinitely long in the periodic orbit in the positive direction. The set on the right is the image of a solution because one can spend infinitely long in the triangle in both the negative and positive directions.	24
3.3	Six multivectors, depicted in black, all of which contain a triangle. The three multivectors on the top row are critical, whereas the three multivectors on the bottom row are regular.	25
3.4	Three examples of an invariant set, marked in yellow. The invariant set on the left is given by the union of the images of two solutions: one which remains at the central triangle for infinitely long in the negative and positive directions, and one that stays in the periodic orbit for infinitely long in the negative and positive directions. In contrast, the invariant set in the middle is given by the union of many solutions. In particular, it is given by the union of all solutions that spend infinitely long in the central triangle in the negative direction, then follow a path to the periodic orbit, and then spend infinitely long in the periodic orbit in the positive direction. The invariant set on the right corresponds to exactly one of these solutions.	26

- 3.5 Three invariant sets on the same multivector field, marked in yellow. The invariant set on the left is isolated by the entire rectangle. The invariant set in the middle is isolated by its closure, but not by the rectangle. This is because there is a path, marked in red, that starts in the invariant set, leaves the invariant set, and then re-enters the invariant set, all while staying within the rectangle. The invariant set on the right is isolated by neither its closure nor the rectangle, because there is a path from a yellow triangle, to the red edge, to a yellow vertex. 27
- 3.6 Two index pairs for the yellow triangle, denoted σ . The left is given by $(\text{cl}(\sigma), \text{mo}(\sigma))$ where $\text{mo}(\sigma)$ is in red and $\text{cl}(\sigma) \setminus \text{mo}(\sigma)$ is exactly the gold triangle. The second index pair is $(\text{pf}(\text{cl}(\sigma)), \text{pf}(\text{mo}(\sigma)))$, where $\text{pf}(\text{mo}(\sigma))$ consists of red simplices while $\text{pf}(\text{cl}(\sigma))$ consists of all red and yellow simplices. Note that the second index pair is also an index pair in N , where N is taken to be the entire rectangle. . . 29
- 3.7 On the left, we show a multivector field and the minimal Morse decomposition for the maximal invariant set in N , where N is the entire rectangle. The maximal invariant set in N is also the entire rectangle because every colored simplex is critical and every white or black simplex is on a path from the golden critical triangle to the blue periodic attractor or from the gray critical triangle to the blue periodic attractor. Each of the six Morse sets in the minimal Morse decomposition is represented by a vertex in the Conley-Morse graph (right) with a matching color. Each vertex is annotated with a Poincaré polynomial that summarizes the Conley indices of the Morse set. 31
- 4.1 Two invariant sets, marked in gold, under two similar but not identical multivector fields. Each invariant set is exactly a gold triangle and corresponds to a repelling fixed point in the classical setting. We can obtain an index pair for each invariant set by taking the closure and the mouth of each invariant set. The simplices in the closure are those simplices that are gold or red and the simplices in the mouth are red. 33
- 4.2 All three images depict the same multivector field, which includes a yellow repelling fixed point (triangle, marked with a black circle). (left) and (right) depict two different index pairs, (P_l, E_l) and (P_r, E_r) , for the repelling fixed point: P_l and P_r consist of yellow and red simplices and E_l and E_r consist of red simplices. The intersection $(P_l \cap P_r, E_l \cap E_r)$ is depicted in the middle. Check that this pair is not an index pair because if e denotes a yellow edge, then $F_V(e) \not\subseteq P_l \cap P_r$. Beneath, we depict the barcode that is associated with the zigzag filtration $(P_l, E_l) \supseteq (P_l \cap P_r, E_l \cap E_r) \subseteq (P_r, E_r)$. Because (P_l, E_l) and (P_r, E_r) are both index pairs for the same repelling fixed point, we would expect the barcode to be full. However, as $(P_l \cap P_r, E_l \cap E_r)$ is not an index pair for the repelling fixed point, its relative homology can change drastically. 34
- 4.3 All three images depict an index pair for the yellow triangle marked with a black circle. P is given by the red and yellow simplices, while E is given by the red simplices. If N is taken to be the entire rectangle, then the index pairs on the left and the right are not index pairs in N , but the index pair on the right is. . . 39

- 4.4 Examples of index pairs computed by using the push forward on multivector fields induced by a differential equation. A sequence of multivector fields was generated from a λ -parametrized differential equation undergoing supercritical Hopf bifurcation [32, Section 11.2]. The consecutive images (from left to right) present a selection from this sequence: the case when $\lambda < 0$ and there is only an attracting fixed point inside N ; the case when $\lambda > 0$ is small and N contains a repelling fixed point, a small attracting periodic trajectory and all connecting trajectories; the case when $\lambda > 0$ is large and the periodic trajectory is no longer contained in N . In all three images, N is given by the rectangle, E is given by the simplices in red, and $P \setminus E$ is given by the simplices in yellow. Note that in the leftmost image, the only invariant set is a triangle which represents an attracting fixed point. For this invariant set, the only relative homology group which is nontrivial is $H_0(P, E)$, which has a single homology generator. In the middle image, the invariant sets represent a repelling fixed point, a periodic attractor, and heteroclinic orbits which connect the repelling fixed point with the periodic attractor. Note that the relative homology has not changed from the leftmost case, so the only nontrivial homology group is $H_0(P, E)$. In the rightmost image, the periodic attractor is no longer entirely contained within N , so the only invariant set corresponds to a repelling fixed point. Here, the only nontrivial homology group is $H_2(P, E)$, which has one generator, so the Conley index has changed. Algorithm 1 captures this change. The persistence barcode output by Algorithm 1 is below index pairs, where a H_0 generator (white bar) lasts until the periodic trajectory leaves N , at which point it is replaced by an H_2 generator (dark gray bar). 41
- 4.5 Infeasibility of the index pair $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$: The sets $E = \text{pf}(\text{mo}(S))$ are colored pink in all three images, while the invariant sets which equal $P \setminus E$ are golden in all three images. (left) $\mathcal{V}_1 : P_1 \setminus E_1$ consists of a single golden triangle; (right) $\mathcal{V}_2 : P_2 \setminus E_2$ consists of the single golden triangle; (middle) $(P_1 \cap P_2) \setminus (E_1 \cap E_2)$ consists of two golden triangles (excluding the edge between them) in the intersection field $\mathcal{V}_1 \cap \mathcal{V}_2$. The barcode for index pairs is depicted by two dark gray bars, each of which represents a 2-dimensional homology generator. Ideally, these would be a single bar. 43
- 4.6 Enlarging $P \setminus E$ which is gold in all three pictures while E is colored pink. (left) \mathcal{V}_1 ; (right) \mathcal{V}_2 ; (middle) $\mathcal{V}_1 \cap \mathcal{V}_2$. Note that there is one bar in the barcode, in contrast with Figure 4.5 45
- 4.7 Index pairs on two slightly perturbed multivector fields (left, right) and their intersection (middle). The isolating neighborhood N is the portion of the rectangle where the multivectors are drawn, E is in red, and $P \setminus E$ is in yellow. Note that we have the same difficulty as in Figure 4.5, where there are two homology generators in the intersection multivector field, so we get a broken bar code. 48
- 4.8 A zoomed in version of Figure 4.7. 49

4.9	The same index pairs as in Figure 4.7 with the same color scheme, but after applying Algorithm 2 to reduce the size of E . This forces a 2-dimensional homology generator to persist across both multivector fields (left, right) and their intersection (middle).	49
4.10	In all three images, the isolating neighborhood N is given by the blue and the yellow simplices, while an isolated invariant set in N is given by the yellow simplices. In addition, in the left and the right images, we can obtain an index pair in N by taking the yellow simplices to be P and letting $E := \emptyset$. If (P_1, E_1) denotes the index pair on the left and (P_2, E_2) denotes the index pair in N on the right, then the index pair in the center is given by $(P_1 \cap P_2, E_1 \cap E_2)$. The yellow simplices in the middle are those simplices in $P_1 \cap P_2$. Intuitively, one would expect a 0-dimensional bar and a 1-dimensional bar to persist through all three images. However, when one takes the intersection of the two index pairs, one no longer has an annulus, so the 1-dimensional bar does not persist through all three images. Instead, we obtain several short, 0-dimensional bars. We show the barcode beneath the images, while excluding several of the short 0-dimensional bars.	51
4.11	When one uses Proposition 4.13 to thicken the index pairs in Figure 4.10, a 1-dimensional homology class persists through all three images, which matches the intuition.	53
4.12	Three different index pairs generated from our scheme in Algorithm 3. The isolating neighborhood is in blue, E is in red, and $P \setminus E$ is in yellow. Note how the isolating neighborhood changes by defining a collar around the invariant sets (which are exactly equal to $P \setminus E$). Between the left and middle multivector fields, the periodic attractor partially leaves K , so the maximal invariant set in N is reduced to just a triangle. Hence, the size of N drastically shrinks between the middle and right multivector fields.	62
5.1	Three multivector fields. On the left, a yellow repelling fixed point is surrounded by a blue periodic attractor. In the middle, the repelling fixed point has split into a yellow periodic repeller and a blue attracting fixed point. On the right, the periodic repeller has collided with the periodic attractor to form a red semistable limit cycle (a semistable limit cycle is a limit cycle that is stable on one side and unstable on the other). In this example, all three multivector fields are significantly different from one another. In computing persistence, we will generally assume that there are several intermediate multivector fields, representing a gradual transition between the multivector fields shown here.	64

5.2	This figure illustrates the approach from [18], applied to capturing the changing structure of the multivector fields in Figure 5.1. The approach in [18] requires selecting a single isolated invariant set for each multivector field, with the canonical choice being the maximal isolated invariant set. In each multivector field, the maximal isolated invariant set is an attractor, highlighted in blue. Despite each multivector field giving rise to different dynamical systems, the maximal isolated invariant set is the same in each multivector field. Hence, computing the persistence using techniques from [18] gives a single, 0-dimensional bar, which we depict at the bottom in white.	64
5.3	The Conley-Morse graph (Definition 3.12) for the Morse decompositions (Definition 3.10) in Figure 5.1, where the top vertices represent the fixed points. The colors of the vertices match the colors of the Morse sets with which they are drawn in Figure 5.1. Each label captures information about the Conley indices of the Morse sets.	65
5.4	We extract a set of zigzag filtrations to capture the changing Conley indices in a sequence of Conley-Morse graphs. In this case, we extract three particular zigzag filtrations, which correspond to the sequences of Morse sets boxed in rectangles.	65
5.5	Illustrating our approach for computing the changing structure of the multivector fields in Figure 5.1. By extracting a specific set of zigzag filtrations (Section 5.1), we can compute a set of barcodes which represent the changing Conley indices of the Morse sets (Section 5.3). Each row represents an extracted filtration, and the barcode for the filtration is depicted below it. White bars are 0-dimensional, light gray bars are 1-dimensional, and dark gray bars are 2-dimensional. In addition, we extract barcodes representing the changing structure of the Conley-Morse graph (not pictured). We collate all of these bars into a single barcode by removing redundancy in Section 5.3.	66
5.6	If the multivector field on the left is \mathcal{V}_1 and the one on the right is \mathcal{V}_2 , then the multivector field in the middle is $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. The maximal invariant set is the entire rectangle in all three multivector fields, and the Morse sets in the minimal Morse decomposition for this isolated invariant set are colored. In the middle, the Morse sets in gray, green, magenta, and pink are spurious because they are not contained in a Morse set on the right. Also in the middle, there is a connection from the golden triangle to the magenta vertex. But on the left, the connection passes through the turquoise triangle. Hence, the golden triangle-to-magenta vertex connection is not relevant.	71
5.7	On the top row, we depict the Conley-Morse graphs (absent the Pointcaré polynomials) for the Morse decompositions in Figure 5.6. The bottom row depicts the Conley-Morse graphs for the left and right Morse decompositions in Figure 5.6, but it instead includes the relevant Conley-Morse graph in the center. . . .	74

5.8	All three maximal sequences extracted from the changing Conley-Morse graph for the Morse decompositions in Figure 5.1. The isolating neighborhood is given by yellow, red, and blue simplices, while if (P, E) is an index pair, the simplices in $P \setminus E$ are in yellow and the simplices in E are in red. The top three images show a periodic attractor that becomes a semistable limit cycle. The middle three images show a repelling fixed point that becomes a periodic repeller and then becomes a semistable limit cycle, and the bottom three images show a repelling fixed point that transitions into an attracting fixed point. Beneath each maximal sequence, we include the barcode from the zigzag filtration that we get by applying Theorem 4.4 to the maximal sequence to obtain a Conley-Morse filtration. White bars are 0-dimensional, light gray bars are 1-dimensional, and dark gray bars are 2-dimensional.	77
5.9	The barcode corresponding to the sequence of combinatorial dynamical systems in Figure 5.1. The white bar above the dotted line is 0-dimensional, and it represents the connected component in the Conley-Morse graph. The bars below the dotted line are obtained by extracting the barcodes from the Conley-Morse filtrations in Figure 5.8 and removing redundant bars.	78
6.1	Applying Step 1a to an invariant set (yellow, left) to get a new one (yellow, right).	85
6.2	Applying Step 1b to an invariant set (yellow, left) to get a new one (yellow, right).	85
6.3	Applying Step 1c to an invariant set (yellow, left) to get a new one (yellow, right). The merged vector is outside of the invariant set on the left, so the invariant sets are the same.	86
6.4	Applying Step 1d to an invariant set (yellow, left) to get a new one (yellow, right).	86
6.5	Applying Step 2f to an invariant set (yellow, left) to get a new one (yellow, right). The associated persistence barcode is depicted below the figures.	86
6.6	Subfigure 6.6a contains an initial multivector field and a seed isolated invariant set, which is a yellow edge. Each subsequent subfigure contains a multivector field that is an atomic refinement or atomic coarsening of the previous. The isolated invariant set that we get by iteratively applying the Tracking Protocol is depicted in yellow. Splitting and merging multivectors are in blue.	87
6.7	The barcode associated with the tracked invariant sets in Figure 6.6. Starting with subfigure 6.6h, we see the birth of a new 1-dimensional homology generator.	87
6.8	An index pair, where P is in yellow and E is empty, for the isolated invariant sets in Figure 6.4. There is a common index pair for both isolated invariant sets, so they form a continuation.	89
6.9	An index pair, where P is given by the yellow and red simplices and E is given by the red simplices, for the isolated invariant sets in Figure 6.2. Thus, they form a continuation.	89

ABSTRACT

Recent innovations in combinatorial dynamical systems permit them to be studied with algorithmic methods. One such method from topological data analysis, called persistent homology, allows one to summarize the changing homology of a sequence of simplicial complexes. This dissertation explicates three methods for capturing and summarizing changes in combinatorial dynamical systems through the lens of persistent homology. The first places the Conley index in the persistent homology setting. This permits one to capture the persistence of salient features of a combinatorial dynamical system. The second shows how to capture changes in combinatorial dynamical systems at different resolutions by computing the persistence of the Conley-Morse graph. Finally, the third places Conley's notion of continuation in the combinatorial setting and permits the tracking of isolated invariant sets across a sequence of combinatorial dynamical systems.

1. INTRODUCTION

One of the main achievements of topological data analysis is *persistent homology*, briefly *persistence*, which is a technique that permits one to summarize the changing homology of a finite sequence of simplicial complexes [1]–[4]. In particular, persistence is celebrated for assigning a homology to point clouds via the so-called Vietoris-Rips complex [5]. The output of persistence is a *persistence diagram* [1] or a *barcode* [2], both of which give one a visual description of changing homology. These objects have been vectorized [6], [7] and have found applications in a variety of areas [8]–[12].

Despite its wide applicability, there has been little work done to apply persistence to vector valued data. In this dissertation, we consider precisely this problem. As a motivating example, consider the λ -parameterized differential equation in Equation 1.1.

$$\begin{aligned}x' &= -y + x(\lambda - x^2 - y^2) \\y' &= x + y(\lambda - x^2 - y^2)\end{aligned}\tag{1.1}$$

The equation behaves very differently depending on the value of λ . Suppose that we are only interested in the behavior of the differential equation in the region $\{(x, y) \in \mathbb{R}^2 \mid -4 \leq x \leq 4 \wedge -4 \leq y \leq 4\}$. Equivalently, we are interested in the behavior of the differential equation in the 8 by 8 square in \mathbb{R}^2 that is centered on the origin and with sides parallel to the axes. As one varies the value of λ , we get the behavior depicted in Figure 1.1. When $\lambda < 0$, we have an attracting fixed point. At $\lambda = 0$, we have a bifurcation point, and the attracting fixed point splits into an attracting periodic orbit and a repelling fixed point. After $\lambda = 16$, the periodic orbit partially leaves our region of interest, and all that we are left with is a repelling fixed point. Intuitively, features of this dynamical system persist as we vary λ . Initially, we have an attracting fixed point, and when $\lambda > 0$, the “attractingness” from the fixed point is preserved in the periodic orbit. Once the periodic orbit is no longer entirely within our region of interest, a repelling fixed point “persists” as $\lambda \rightarrow \infty$. We aim to capture these and related phenomena using persistence.

To capture these phenomena, we will utilize *combinatorial dynamical systems* that are generated by *multivector fields* [13], [14]. These dynamical systems are defined on finite




Figure 1.1. The behavior of the differential equation in Equation 1.1. From left-to-right, top-to-bottom, the depicted behavior is of the case when $\lambda = -10, 0, 1, 3, 6, 10, 15, 20, 25$.

topological spaces, and they are rooted in earlier work by R. Forman [15], [16]. Developing the theory of combinatorial dynamical systems is an area of active research, and there have been several recent studies on the topic [13], [14], [17]–[22]. Some studies have placed classical notions, such as Morse decompositions and the Conley index, in the combinatorial setting, while others have proven connections between the classical and combinatorial settings.

This dissertation proposes three ways to capture changes in combinatorial dynamical systems. In Chapter 4, we show how to compute the persistence of the Conley index across a sequence of combinatorial dynamical systems. Chapter 5 shows how to compute the persistence of the Conley-Morse graph, which captures information at a finer resolution than the Conley index does alone. Finally, Chapter 6 places the classical notion of continuation in the combinatorial framework, and shows how to use continuation and persistence to track isolated invariant sets. The work in Chapters 4, 5, and 6 first appeared in [18]–[20], and passages from [18]–[20] are quoted in this dissertation verbatim.

Before we show how to use persistence to capture changes in combinatorial dynamical systems, we first briefly review persistent homology in Chapter 2 and formally introduce combinatorial dynamical systems in Chapter 3. Several passages from these chapters are also drawn verbatim from [18]–[20].

2. PERSISTENT HOMOLOGY

We begin by reviewing *persistent homology* [1]–[4] or briefly, *persistence*, which is a tool that captures the changing homology over an indexed sequence of simplicial complexes. Throughout this dissertation, we will assume that the reader is familiar with homology. The unfamiliar reader is encouraged to consult [23], [24]. Unless otherwise specified, all references to homology in this dissertation refer to simplicial homology with coefficients drawn from a finite field.

A typical setting in persistence is where one has a simplicial complex K and a function $f : K \rightarrow \mathbb{R}$ that respects the face relation. That is, if σ is a face of τ , then $f(\sigma) \leq f(\tau)$. In such a setting, the sublevel sets with respect to f are subcomplexes. Let $K_i := \{\sigma \in K \mid f(\sigma) \leq i\}$. It is easy to see that if $i \leq j$, then $K_i \subseteq K_j$. Hence, the function f admits a *filtration* $\mathcal{F} : K_{a_1} \subseteq K_{a_2} \subseteq \dots \subseteq K_{a_n}$, where a_i is the i th element of $f(K)$ in sorted order. We always assume that K is finite, so all filtrations are finite. The reader will note that if we are given a filtration $K_1 \subseteq K_2 \subseteq \dots \subseteq K_n = K$, then one can obtain a function $f : K \rightarrow \mathbb{R}$ where $f(\sigma) = \min\{i \mid \sigma \in K_i\}$. As each K_i is assumed to be a subcomplex of K , it follows immediately that f respects the face relation. Hence, a function f that respects the face relation induces a filtration, and a filtration induces a function f .

Corresponding to each filtration is a unique *barcode*, which captures how the homology changes across the complexes in \mathcal{F} . We include an example of a filtration and its associated barcodes in Figure 2.1. We refer the reader to [3], [4] for a comprehensive, modern treatment on persistence, barcodes, and algorithms for computing them.

A natural generalization of persistence is *zigzag persistence*, which arises from *zigzag filtrations*. These filtrations allow inclusions in both directions. Forward and backward inclusions in a zigzag filtration

$$K_{a_1} \supseteq K_{a_2} \subseteq K_{a_3} \supseteq K_{a_4} \supseteq \dots \subseteq K_{a_n}$$

correspond to insertion and deletions of simplices, respectively. Similar to standard filtrations, each zigzag filtration provides a unique barcode that captures the changing homology




Figure 2.1. An example filtration and its associated barcode. Two 0-dimensional bars, in white, capture two connected components as they persist through the filtration. The 1-dimensional bar, in grey, captures the presence of the 1-dimensional cycle.

of the filtration. We include an example in Figure 2.2. The interested reader is encouraged to consult [4], [25] for additional information on zigzag persistence. It is known that the barcode corresponding to a zigzag filtration can be computed in $O(M(n) + n^2 \log^2(n))$ -time, where $M(n)$ is the complexity of matrix multiplication and n is the number of simplices that are inserted and deleted in the zigzag filtration [26]. In the special case of standard (or unidirectional persistence), there exists an $O(M(n))$ algorithm to compute the barcode by using a result in [27].




Figure 2.2. An example zigzag filtration and its associated barcode. The 0-dimensional bar, in white, captures the single connected component that persists through the entire filtration. Similarly, two 1-dimensional bars, in gray, capture the lifetime of the two 1-dimensional cycles, which are both present in the middle complex.

3. COMBINATORIAL DYNAMICAL SYSTEMS

The origins of combinatorial dynamical systems lie with R. Forman, who in two seminal papers [15], [16], introduced the notion of a *combinatorial vector field* and a corresponding dynamical system. Initially, in [15], Forman introduced *discrete gradient vector fields* on cell complexes and proved various analogs of theorems from classical Morse theory. However, gradient vector fields are very restrictive and only admit simple dynamics. Roughly speaking, under a discrete gradient vector field, all cells are either *critical* (corresponding to the classical notion of a critical point) or lie on a path between two critical cells (reminiscent of the concept of an *integral line*). To address this, Forman subsequently introduced combinatorial dynamical systems and considered non-gradient dynamics in [16].

Despite permitting non-gradient behavior, non-gradient combinatorial vector fields still admit fairly restrictive dynamics. Vectors are only permitted to consist of two simplices (σ, τ) where $\dim(\tau) = \dim(\sigma) + 1$. In addition, each simplex is allowed to be part of at most one vector. Hence, one can view a combinatorial vector field V as a partition of a cell complex C into *singletons* and *doubletons*. Singletons are those cells that are not part of a vector and therefore correspond to critical points, while doubletons are those cells that are part of a vector. How flow is defined on these vector fields means that one can only move between cells of alternating dimensions. Hence, if considering a simplicial complex that consists of vertices, edges, and triangles, one can only move from a triangle to an edge and then back to a triangle - it is not permitted to move from a triangle to a vertex or vice-versa.

M. Mrozek relaxed this requirement when he introduced multivector fields in [13]. Multivector fields permit vectors that are much more complicated than those of Forman and therefore allow much more complicated dynamics. As an example, one can model a monkey saddle with a single multivector, whereas this phenomenon cannot be modeled with combinatorial vector fields. Mrozek's original formulation was subsequently generalized in [14], when the authors relaxed a requirement that multivectors contain a unique maximal element.

3.1 Multivector Fields and Combinatorial Dynamics

We now move to formally define multivector fields and combinatorial dynamical systems. In their full generality, one can define a multivector field and a corresponding combinatorial dynamical system on a finite set X with a partial order \leq . However, because we aim to use persistent homology to capture changes in combinatorial dynamical systems, we restrict our attention to finite simplicial complexes. Hence, we let K denote a finite (abstract) simplicial complex, and let \leq denote the face relation on K . Namely, for $\sigma, \tau \in K$, we let $\sigma \leq \tau$ if and only if $\sigma \subseteq \tau$. The face relation \leq lends itself to a notion of convexity: a set $A \subseteq K$ is *convex* if every trio of simplices $\sigma, \tau \in A$, $\rho \in K$ where $\sigma \leq \rho \leq \tau$ has the property that $\rho \in A$. We define multivectors and multivector fields in terms of convex sets.

Definition 3.1 (Multivector, Multivector Field). *A multivector V is a convex subset of K . A multivector field over K is a partition of K into multivectors.*

We typically denote individual multivectors with the roman V and we denote multivector fields with the script \mathcal{V} . Let V denote a multivector, and let $\sigma \in V$ denote a simplex such that for all $\tau \in V$, either $\tau \leq \sigma$ or τ and σ are incomparable under \leq . We call such a σ a *maximal* simplex in V . Typically, we depict a multivector V by drawing a vector from each nonmaximal simplex in V to each maximal simplex in V . Hence, one can determine if σ and τ are in the same multivector by checking if there is a vector from σ to τ , τ to σ , or if there are two vectors from each of σ and τ to some other simplex ρ . If a multivector consists of exactly one simplex, then we mark that simplex with a circle. We include an example of multivector fields in Figure 3.1.

The reader will note that convexity on its own is a fairly loose condition. Namely, disconnected sets may be convex, and convex sets may have multiple maximal simplices. All of the multivectors in Figure 3.1 are connected and have a unique maximal simplex. Disconnected multivectors typically do not occur in practice, while multivectors with multiple maximal simplices may occasionally occur in computer-generated examples. In the interest of simplicity, we usually only depict examples that are connected and contain a unique maximal simplex. However, none of the results in this dissertation require such an assumption.




Figure 3.1. Three multivector fields. Multivectors are drawn by including a vector from each nonmaximal simplex in the multivector to each maximal simplex. In these three cases, each multivector has a unique maximal simplex. Multivectors that consist of a single simplex are marked with a circle.

Thus far, we have defined multivector fields. However, multivector fields are not dynamical systems on their own. Hence, we move to introduce an element of dynamism that first appeared in [13]. Our dynamics generator will require the notion of the *closure* of a set of simplices. For a simplex $\sigma \in K$, we define the closure of σ , denoted $\text{cl}(\sigma)$, as $\text{cl}(\sigma) := \{\tau \in K \mid \tau \leq \sigma\}$. This idea extends to the closure of a set of simplices, where for a set $A \subseteq K$, the set $\text{cl}(A) := \bigcup_{\sigma \in A} \text{cl}(\sigma)$. We say that $A \subseteq K$ is *closed* if $\text{cl}(A) = A$. The reader familiar with the Alexandrov topology [28, Section 1.1] will immediately notice that this notation and terminology is aligned with the topology induced on K by the relation \leq . For $\sigma \in K$, we denote the multivector containing σ as $[\sigma]$. If the multivector field \mathcal{V} is not clear from context, we will use the notation $[\sigma]_{\mathcal{V}}$. Note that a multivector field is a partition of a simplicial complex, so $[\sigma]_{\mathcal{V}}$ is unique. We now use a multivector field \mathcal{V} on K to define a multivalued map $F_{\mathcal{V}} : K \multimap K$. In particular, we let $F_{\mathcal{V}}(\sigma) := \text{cl}(\sigma) \cup [\sigma]_{\mathcal{V}}$. Such a multivalued map induces a notion of flow on K . In the interest of brevity, for $a, b \in \mathbb{Z}$, we set $\mathbb{Z}_{[a,b]} := [a, b] \cap \mathbb{Z}$ and define $\mathbb{Z}_{(a,b)}$, $\mathbb{Z}_{[a,b)}$, $\mathbb{Z}_{(a,b]}$ as expected.

Definition 3.2 (Path, Solution). *A path from σ to σ' under the multivector field \mathcal{V} is a function $\rho : \mathbb{Z}_{[a,b]} \rightarrow K$, where $\rho(a) = \sigma$, $\rho(b) = \sigma'$, and for all $i \in \mathbb{Z}_{(a,b)}$, we have that $\rho(i) \in F_{\mathcal{V}}(\rho(i-1))$. Similarly, a solution to a multivector field over K is a map $\rho : \mathbb{Z} \rightarrow K$ where $\rho(i) \in F_{\mathcal{V}}(\rho(i-1))$.*

Informally, a path is a finite sequence of simplices $\sigma_0, \sigma_1, \dots, \sigma_n$ where $F_{\mathcal{V}}$ permits one to “move” from σ_i to σ_{i+1} . A solution is a bi-infinite sequence of such simplices. We include three examples of solutions in Figure 3.2. An unfortunate consequence of this definition is that every simplex σ admits a solution $\rho : \mathbb{Z} \rightarrow K$ where $\rho(i) = \sigma$ for all $i \in \mathbb{Z}$. This does not match the intuition from differential equations: solutions that stay at one point correspond exactly to critical points. Hence, so as to faithfully replicate this behavior, we need a notion of a “critical” multivector [13], [14].

To define critical multivectors, we define the *mouth* of a set as $\text{mo}(A) := \text{cl}(A) \setminus A$. The multivector V is critical if the relative homology $H_p(\text{cl}(V), \text{mo}(V)) \neq 0$ in some dimension p . Otherwise, V is *regular*. Note that $H_p(\text{cl}(V), \text{mo}(V))$ is well defined because for all $S \subseteq K$, $\text{mo}(S) \subseteq \text{cl}(S) \subseteq K$. Intuitively, a multivector V is regular if $\text{cl}(V)$ can be collapsed onto




Figure 3.2. Three examples of solutions, marked in yellow, over the same multivector field. The set on the left is the image of a solution under the given multivector field because one can spend infinitely long in the periodic orbit in both the positive and negative directions. For the set in the middle, one spends infinitely long in the triangle that is marked with a circle, then follows a path to the periodic orbit, and spends infinitely long in the periodic orbit in the positive direction. The set on the right is the image of a solution because one can spend infinitely long in the triangle in both the negative and positive directions.

$\text{mo}(V)$ via elementary collapses. We include examples of critical and regular multivectors in Figure 3.3. In Figure 3.3, the top three multivectors are critical, while the bottom three multivectors are regular. For a multivector V that is depicted in Figure 3.3, the set of simplices that are yellow or red is exactly equal to $\text{cl}(V)$, while the set of simplices that are in red is equal to $\text{mo}(V)$. Hence, if V_1 is the multivector depicted in the top-left, it is easy to verify that $H_2(\text{cl}(V_1), \text{mo}(V_1); \mathbb{Z}_2) = \mathbb{Z}_2$, and therefore, the multivector is critical. If the multivector in the top-middle is V_2 , then $H_1(\text{cl}(V_2), \text{mo}(V_2); \mathbb{Z}_2) = \mathbb{Z}_2 \neq 0$, and if V_3 is the multivector in the top-right, then $H_0(\text{cl}(V_3), \text{mo}(V_3); \mathbb{Z}_2) = \mathbb{Z}_2 \neq 0$. Hence, all of the multivectors on the top row of Figure 3.3 are critical. Similarly, if V denotes a multivector on the bottom row, then $H(\text{cl}(V), \text{mo}(V)) = 0$.

The authors in [14] used critical multivectors to define a special type of solution.




Figure 3.3. Six multivectors, depicted in black, all of which contain a triangle. The three multivectors on the top row are critical, whereas the three multivectors on the bottom row are regular.

Definition 3.3 (Essential Solution). *A solution $\rho : \mathbb{Z} \rightarrow K$ is an essential solution of the multivector field \mathcal{V} on K if for each $i \in \mathbb{Z}$ where $[\rho(i)]_{\mathcal{V}}$ is regular, there exists an $i^-, i^+ \in \mathbb{Z}$ where $i^- < i < i^+$ and $[\rho(i^-)]_{\mathcal{V}} \neq [\rho(i)]_{\mathcal{V}} \neq [\rho(i^+)]_{\mathcal{V}}$.*

Fundamentally, essential solutions are solutions where you are only permitted to stay within a single multivector V for infinitely long (in the positive or negative directions) if V is critical. In Figure 3.2, the leftmost solution is essential, because one does not stay in any given multivector for infinitely long. The solution in the middle is also essential, because

while one stays in the yellow triangle for infinitely long in the negative direction, the yellow triangle is the only simplex in a critical multivector. However, the solution on the right is not essential, because one is staying in a regular vector for infinitely long in the positive and negative directions.

For a set $A \subseteq K$, let $\mathbf{eSol}(A)$ denote the set of essential solutions ρ such that $\rho(\mathbb{Z}) \subseteq A$. We define the *invariant part* of A as $\mathbf{Inv}(A) = \{\sigma \in A \mid \exists \rho \in \mathbf{eSol}(A), \rho(0) = \sigma\}$. If the multivector field is not clear from context, we use the notation $\mathbf{Inv}_{\mathcal{V}}(A)$.

Definition 3.4 (Invariant Set). *A set A is invariant or an invariant set with respect to the multivector field \mathcal{V} if $\mathbf{Inv}_{\mathcal{V}}(A) = A$.*

Informally, an invariant set is a union of solutions to a multivector field. We include examples of invariant sets in Figure 3.4.




Figure 3.4. Three examples of an invariant set, marked in yellow. The invariant set on the left is given by the union of the images of two solutions: one which remains at the central triangle for infinitely long in the negative and positive directions, and one that stays in the periodic orbit for infinitely long in the negative and positive directions. In contrast, the invariant set in the middle is given by the union of many solutions. In particular, it is given by the union of all solutions that spend infinitely long in the central triangle in the negative direction, then follow a path to the periodic orbit, and then spend infinitely long in the periodic orbit in the positive direction. The invariant set on the right corresponds to exactly one of these solutions.

Solutions and invariant sets are defined in accordance with their corresponding concepts in the classical setting. For more information on these classical counterparts, see [29]. In general, we will focus our attention on particular types of invariant sets. The first relevant condition is \mathcal{V} -compatibility.

Definition 3.5 (\mathcal{V} -compatibility). *An invariant set S under the multivector field \mathcal{V} on K is \mathcal{V} -compatible if each multivector $V \in \mathcal{V}$ where $V \cap S \neq \emptyset$ has the property that $V \subseteq S$.*

Intuitively, an invariant set S is \mathcal{V} -compatible if S can be written as a union of multivectors. In Figure 3.4, the left two invariant sets are \mathcal{V} -compatible. The invariant set on the right is not \mathcal{V} -compatible, because the entire multivector that contains the thickened edge is not included in the invariant set.

As in the classical setting, we have a notion of isolation for combinatorial invariant sets.

Definition 3.6 (Isolated Invariant Set, Isolating Neighborhood). *A \mathcal{V} -compatible invariant set $S \subseteq N$, N closed, is isolated by N if all paths $\rho : \mathbb{Z}_{[a,b]} \rightarrow N$ for which $\rho(a), \rho(b) \in S$ satisfy $\rho(\mathbb{Z}_{[a,b]}) \subseteq S$. The closed set N is said to be an isolating neighborhood for S . If there exists an isolating neighborhood for S , then S is an isolated invariant set.*

Intuitively, S is an isolated invariant set if there exists some closed set $N \supseteq S$ such that it is not possible to leave and reenter S while staying within N . Hence, an invariant set may be isolated by some closed sets N , but not others. We elaborate on this in Figure 3.5.




Figure 3.5. Three invariant sets on the same multivector field, marked in yellow. The invariant set on the left is isolated by the entire rectangle. The invariant set in the middle is isolated by its closure, but not by the rectangle. This is because there is a path, marked in red, that starts in the invariant set, leaves the invariant set, and then re-enters the invariant set, all while staying within the rectangle. The invariant set on the right is isolated by neither its closure nor the rectangle, because there is a path from a yellow triangle, to the red edge, to a yellow vertex.

The astute reader will note that as $\text{cl}(S)$ is contained in all closed sets $N \supseteq S$, it follows trivially that a \mathcal{V} -compatible invariant set is isolated if and only if it is isolated by $\text{cl}(S)$.

3.2 Conley Indices

The Conley index of an isolated invariant set is a topological invariant used to characterize features of dynamical systems [30], [31]. In both the classical and the combinatorial settings, the Conley index is determined by index pairs.

Definition 3.7. *Let S be an isolated invariant set. The pair of closed sets (P, E) subject to $E \subseteq P \subseteq K$ is an index pair for S if all of the following hold:*

1. $F_V(E) \cap P \subseteq E$
2. $F_V(P \setminus E) \subseteq P$
3. $S = \text{Inv}(P \setminus E)$

In addition, an index pair is said to be a *saturated index pair* if $S = P \setminus E$. In Figure 3.6, on the left, the gold, critical triangle σ is an isolated invariant set. The reader can easily verify that $(\text{cl}(\sigma), \text{mo}(\sigma))$ is an index pair for σ . In fact, this technique is a canonical way of picking an index pair for an isolated invariant set. This is formalized in the following proposition.

Proposition 3.8. [14, Proposition 4.3] *Let S be an isolated invariant set. Then $(\text{cl}(S), \text{mo}(S))$ is a saturated index pair for S .*

However, there are several other natural ways to find index pairs. Figure 3.6, on the right, depicts another index pair for the same gold triangle σ . By letting $P := \text{cl}(\sigma) \cup S_P$ and $E := \text{mo}(\sigma) \cup S_E$, where S_P and S_E are the set of simplices reachable from paths originating in $\text{cl}(S)$, $\text{mo}(S)$ respectively, we obtain a much larger index pair. In Figure 3.6, P is the set of red simplices and yellow simplices, while E is the set of all red simplices.

In principle, the Conley index must be independent of the choice of index pair. Fortunately, it is also known that the relative homology given by an index pair for an isolated invariant set S is independent of the choice of index pair.

Theorem 3.9. [14, Theorem 4.15] *Let (P_1, E_1) and (P_2, E_2) be index pairs for the isolated invariant set S . Then $H_p(P_1, E_1) \cong H_p(P_2, E_2)$ for all p .*

The *Conley Index* of an isolated invariant set S in dimension p is then given by the relative homology group $H_p(P, E)$ for any index pair of S denoted (P, E) .




Figure 3.6. Two index pairs for the yellow triangle, denoted σ . The left is given by $(\text{cl}(\sigma), \text{mo}(\sigma))$ where $\text{mo}(\sigma)$ is in red and $\text{cl}(\sigma) \setminus \text{mo}(\sigma)$ is exactly the gold triangle. The second index pair is $(\text{pf}(\text{cl}(\sigma)), \text{pf}(\text{mo}(\sigma)))$, where $\text{pf}(\text{mo}(\sigma))$ consists of red simplices while $\text{pf}(\text{cl}(\sigma))$ consists of all red and yellow simplices. Note that the second index pair is also an index pair in N , where N is taken to be the entire rectangle.

3.3 Morse Decompositions and the Conley-Morse Graph

This paper heavily uses *Morse decompositions* of isolated invariant sets. For an essential solution $\rho : \mathbb{Z} \rightarrow K$, we use the notation $\alpha(\rho) := \bigcap_{i=1}^{\infty} \rho(-\infty, -i]$ and $\omega(\rho) = \bigcap_{i=1}^{\infty} \rho[i, \infty)$. Because we take K to be finite, $\alpha(\rho) \neq \emptyset$ and $\omega(\rho) \neq \emptyset$. Intuitively, $\alpha(\rho)$ and $\omega(\rho)$ capture the behavior at the beginning and end of an essential solution.

Definition 3.10 (Morse Decomposition). *Let S denote an isolated invariant set in N and (\mathbb{P}, \leq) a finite poset. The collection $\mathcal{M} = \{M_p \mid p \in \mathbb{P}\}$ is called a Morse decomposition of S if both of the following conditions are satisfied:*

1. \mathcal{M} is a family of mutually disjoint, isolated invariant subsets of S , and
2. For every essential solution $\rho : \mathbb{Z} \rightarrow S$ either $\text{im}(\rho) \subseteq M_r$ for an $r \in \mathbb{P}$ or there exist $p, q \in \mathbb{P}$ such that $q > p$, $\alpha(\rho) \subseteq M_q$, and $\omega(\rho) \subseteq M_p$.

An element of a Morse decomposition is called a Morse set. For Morse sets $M_p, M_q \in \mathcal{M}$, we often abuse notation and write $M_p \leq M_q$ if $p \leq q$. If the only Morse decomposition for an isolated invariant set S is $\mathcal{M} = \{S\}$, then S is minimal.

We frequently discuss *connections* between Morse sets. Fix a Morse decomposition \mathcal{M} of an isolated invariant set S . A connection is a path $\rho : \mathbb{Z}_{[0,n]} \rightarrow S$ where $\rho(0) \in M_1 \in \mathcal{M}$ and $\rho(n) \in M_2 \in \mathcal{M}$ and $\rho(i) \notin \bigcup_{M \in \mathcal{M}} M$ for all $i \in \mathbb{Z}_{[1,n-1]}$. In this dissertation, we will frequently use a particular type of Morse decomposition called a *minimal* Morse decomposition. A Morse decomposition \mathcal{M} is minimal if each $M \in \mathcal{M}$ is minimal. In Figure 3.7, on the left, the entire rectangle is an invariant set S . Each Morse set in the minimal Morse decomposition for S is depicted in a different color. Fortunately, there is a simple characterization of minimal isolated invariant sets.

Proposition 3.11. [14, Proposition 6.7] *Let \mathcal{V} denote a multivector field over K , and let S denote an isolated invariant set under \mathcal{V} . The set S is minimal if and only if for all $\sigma, \tau \in S$, there exists a path $\rho : [0, n] \rightarrow S$ where $\rho(0) = \sigma$ and $\rho(n) = \tau$.*

Hence, when viewing a combinatorial dynamical system as a directed graph G where each simplex σ corresponds to a vertex $v(\sigma)$, the minimal Morse decomposition of an isolated invariant set S is given by finding the strongly connected components of the subgraph of G that corresponds to S [14]. A Morse decomposition and the Conley index provide different information about an isolated invariant set. These two descriptors are often combined into the *Conley-Morse graph*, which captures the structure of a Morse decomposition and contains information about the Conley indices of the Morse sets in the decomposition.

Definition 3.12 (Conley-Morse graph). *Let \mathcal{M} denote a Morse decomposition, and let G denote the directed graph such that there is a bijection $f : \mathcal{M} \rightarrow V(G)$, and there exists a connection from $M \in \mathcal{M}$ to $M' \in \mathcal{M}$ if and only if there exists a directed edge from $f(M)$ to $f(M')$. The Conley-Morse graph for \mathcal{M} is the graph G where each vertex $f(M) = v \in V(G)$ is annotated with the Poincaré polynomial $\sum_{i=0}^m \beta_i t^i$ where m is the largest integer for which the m -dimensional Conley index is nontrivial and β_i is the rank of the i -dimensional Conley index of M .*

For convenience, we use M to refer both to a Morse set and to its corresponding vertex in the Conley-Morse graph. We include an example of a Morse decomposition and the associated Conley-Morse graph in Figure 3.7.




Figure 3.7. On the left, we show a multivector field and the minimal Morse decomposition for the maximal invariant set in N , where N is the entire rectangle. The maximal invariant set in N is also the entire rectangle because every colored simplex is critical and every white or black simplex is on a path from the golden critical triangle to the blue periodic attractor or from the gray critical triangle to the blue periodic attractor. Each of the six Morse sets in the minimal Morse decomposition is represented by a vertex in the Conley-Morse graph (right) with a matching color. Each vertex is annotated with a Poincaré polynomial that summarizes the Conley indices of the Morse set.

4. PERSISTENCE OF THE CONLEY INDEX

In the introduction, we explained that this dissertation aims to use persistent homology to capture changes in combinatorial dynamical systems. Now that we have introduced combinatorial dynamical systems and the Conley index, we move to establish the foundations of Conley index persistence. The results in this chapter are from [18], [19], and we quote passages from these papers verbatim.

Given a sequence of multivector fields $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ on a simplicial complex K , one may want to quantify the changing behavior of the vector fields. One such approach is to compute a sequence of isolated invariant sets S_1, S_2, \dots, S_n under each multivector field, and then to compute an index pair for each isolated invariant set. By Proposition 3.8, a canonical way to do this is to take the closure and mouth of each isolated invariant set to obtain a sequence of index pairs $(\text{cl}(S_1), \text{mo}(S_1)), (\text{cl}(S_2), \text{mo}(S_2)), \dots, (\text{cl}(S_n), \text{mo}(S_n))$. A first idea is to take the element-wise intersection of consecutive index pairs, which results in the zigzag filtration:

$$(\text{cl}(S_1), \text{mo}(S_1)) \supseteq (\text{cl}(S_1) \cap \text{cl}(S_2), \text{mo}(S_1) \cap \text{mo}(S_2)) \subseteq (\text{cl}(S_2), \text{mo}(S_2)) \cdots (\text{cl}(S_n), \text{mo}(S_n))$$

Taking the relative homology groups of the pairs in the zigzag sequence, we obtain a zigzag persistence module. We can extract a barcode corresponding to a decomposition of this module:

$$H_p(\text{cl}(S_1), \text{mo}(S_1)) \leftarrow H_p(\text{cl}(S_1) \cap \text{cl}(S_2), \text{mo}(S_1) \cap \text{mo}(S_2)) \rightarrow H_p(\text{cl}(S_2), \text{mo}(S_2)) \leftarrow \cdots \rightarrow H_p(\text{cl}(S_n), \text{mo}(S_n)).$$

However, the chance that this approach works in practice is low. In general, two isolated invariant sets S_1, S_2 need not overlap, and hence their corresponding index pairs need not intersect. For example, consider the invariant sets, marked with yellow triangles, in Figure 4.1. By taking the closure and the mouth of the invariant sets, one obtains the index pairs in the figure, where the closure is given by the simplices that are red and yellow and the mouth is given by the simplices that are red. The pairwise intersection of these index pairs is empty, so nothing can persist. This is problematic in computing the persistence, because intuitively there should be an H_2 generator that persists through the multivector fields. Furthermore,




Figure 4.1. Two invariant sets, marked in gold, under two similar but not identical multivector fields. Each invariant set is exactly a gold triangle and corresponds to a repelling fixed point in the classical setting. We can obtain an index pair for each invariant set by taking the closure and the mouth of each invariant set. The simplices in the closure are those simplices that are gold or red and the simplices in the mouth are red.

the intersection of two index pairs, even under the same multivector field, need not be an index pair. We include an example in Figure 4.2.




Figure 4.2. All three images depict the same multivector field, which includes a yellow repelling fixed point (triangle, marked with a black circle). (left) and (right) depict two different index pairs, (P_l, E_l) and (P_r, E_r) , for the repelling fixed point: P_l and P_r consist of yellow and red simplices and E_l and E_r consist of red simplices. The intersection $(P_l \cap P_r, E_l \cap E_r)$ is depicted in the middle. Check that this pair is not an index pair because if e denotes a yellow edge, then $F_V(e) \not\subseteq P_l \cap P_r$. Beneath, we depict the barcode that is associated with the zigzag filtration $(P_l, E_l) \supseteq (P_l \cap P_r, E_l \cap E_r) \subseteq (P_r, E_r)$. Because (P_l, E_l) and (P_r, E_r) are both index pairs for the same repelling fixed point, we would expect the barcode to be full. However, as $(P_l \cap P_r, E_l \cap E_r)$ is not an index pair for the repelling fixed point, its relative homology can change drastically.

To address these problems, we consider a special type of index pair called an *index pair* in N .

Definition 4.1. Let S be an invariant set isolated by N under \mathcal{V} . The pair of closed sets (P, E) satisfying $E \subseteq P \subseteq N$ is an index pair for S in N if all of the following conditions are met:

1. $F_V(P) \cap N \subseteq P$
2. $F_V(E) \cap N \subseteq E$
3. $F_V(P \setminus E) \subseteq N$, and
4. $S = \text{Inv}(P \setminus E)$.

As is expected, such index pairs in N are index pairs. The index pair on the right of Figure 3.6 is an index pair in N where N is the entire rectangle.

Theorem 4.2. *Let (P, E) be an index pair in N for S . The pair (P, E) is an index pair for S in the sense of Definition 3.7.*

Proof. Note that by condition three of Definition 4.1, if $\sigma \in P \setminus E$, then $F_{\mathcal{V}}(\sigma) \subseteq N$. Condition one of Definition 4.1 implies that $F_{\mathcal{V}}(\sigma) \cap N = F_{\mathcal{V}}(\sigma) \subseteq P$, which is condition two of Definition 3.7. Likewise, by condition two of Definition 4.1, if $\sigma \in E$, then $F_{\mathcal{V}}(\sigma) \cap N \subseteq E$. Note that $P \subseteq N$, so it follows that $F_{\mathcal{V}}(\sigma) \cap P \subseteq F_{\mathcal{V}}(\sigma) \cap N \subseteq E$, which is condition one of Definition 3.7. Finally, condition four of Definition 4.1 directly implies condition three of Definition 3.7. \square

An additional advantage to considering index pairs in N is that the intersection of index pairs in N is an index pair in N . We get the next two results which involve the notion of a new multivector field obtained by intersection. Given two multivector fields $\mathcal{V}_1, \mathcal{V}_2$, we define $\mathcal{V}_1 \cap \mathcal{V}_2 := \{V_1 \cap V_2 \mid V_1 \in \mathcal{V}_1, V_2 \in \mathcal{V}_2\}$.

Theorem 4.3. *Let $(P_1, E_1), (P_2, E_2)$ be index pairs in N for S_1, S_2 under $\mathcal{V}_1, \mathcal{V}_2$. The set $\text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$ is isolated by N under $\mathcal{V}_1 \cap \mathcal{V}_2$.*

Proof. To contradict, we assume that there exists a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N$ under $\mathcal{V}_1 \cap \mathcal{V}_2$ where $\rho(a), \rho(b) \in \text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$ and there exists some $i \in (a, b) \cap \mathbb{Z}$ where $\rho(i) \notin \text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$. Note that by the definition of an index pair, $F_{\mathcal{V}}(P) \cap N \subseteq P$. Hence, it follows by an easy induction argument that since $F_{\mathcal{V}_1 \cap \mathcal{V}_2}(\sigma) \subseteq F_{\mathcal{V}_1}(\sigma), F_{\mathcal{V}_2}(\sigma)$, we have that $\rho(\mathbb{Z}_{[a,b]}) \subseteq P_1, P_2$. This directly implies that $\rho(\mathbb{Z}_{[a,b]}) \subseteq P_1 \cap P_2$. In addition, it is easy to see that ρ can be extended to an essential solution in $P_1 \cap P_2$, which we denote $\rho' : \mathbb{Z} \rightarrow N$, by some simple surgery on essential solutions. This is because there must be essential solutions $\rho_1, \rho_2 : \mathbb{Z} \rightarrow (P_1 \cap P_2) \setminus (E_1 \cap E_2)$ where $\rho_1(a) = \rho(a)$ and $\rho_2(b) = \rho(b)$, as $\rho(a)$ and $\rho(b)$ are both in essential solutions. Hence, $\rho'(x) = \rho_1(x)$ if $x \leq a$, $\rho'(x) = \rho(x)$ if $a \leq x \leq b$, and $\rho'(x) = \rho_2(x)$ if $b \leq x$. Since ρ' is an essential solution, we have that $\rho(\mathbb{Z}_{[a,b]}) \subseteq \text{Inv}(P_1 \cap P_2)$, but also that $\rho(\mathbb{Z}_{[a,b]}) \not\subseteq \text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$. Therefore,

we must have that $\rho(i) \in E_1 \cap E_2$. But by the same reasoning as before, it follows that $\rho(\mathbb{Z}_{[i,b]}) \subseteq E_1 \cap E_2$. Hence, $b \notin (P_1 \cap P_2) \setminus (E_1 \cap E_2)$, a contradiction. \square

Theorem 4.4. *Let (P_1, E_1) and (P_2, E_2) be index pairs in N under $\mathcal{V}_1, \mathcal{V}_2$. The tuple $(P_1 \cap P_2, E_1 \cap E_2)$ is an index pair for $\text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$ in N under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$.*

Proof. We proceed by using the conditions in Definition 4.1 to show that $(P_1 \cap P_2, E_1 \cap E_2)$ is an index pair in N . Note that $F_{\mathcal{V}_1 \bar{\cap} \mathcal{V}_2}(P_1 \cap P_2) \cap N \subseteq F_{\mathcal{V}_1}(P_1) \cap F_{\mathcal{V}_2}(P_2) \cap N$, which is immediate by the definition of F and considering $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. Note that since (P_1, E_1) and (P_2, E_2) are index pairs in N , we know from Definition 4.1 that $F_{\mathcal{V}_1}(P_1) \cap N \subseteq P_1$ and $F_{\mathcal{V}_2}(P_2) \cap N \subseteq P_2$. Therefore $F_{\mathcal{V}_1 \bar{\cap} \mathcal{V}_2}(P_1 \cap P_2) \cap N \subseteq P_1 \cap P_2$. This implies the first condition in Definition 4.1. This argument also implies the second condition by replacing P with E .

Now, we aim to show that $(P_1 \cap P_2, E_1 \cap E_2)$ satisfies condition three in Definition 4.1. Consider $\sigma \in (P_1 \cap P_2) \setminus (E_1 \cap E_2)$. Without loss of generality, we assume $\sigma \notin E_1$. Therefore, $\sigma \in P_1 \setminus E_1$, so $F_{\mathcal{V}_1}(\sigma) \subseteq N$ by the definition of an index pair in N . Hence, since $F_{\mathcal{V}_1 \bar{\cap} \mathcal{V}_2}(\sigma) \subseteq F_{\mathcal{V}_1}(\sigma)$, condition three is satisfied.

Finally, note that $\text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$ is obviously equal to $\text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$, so condition four holds as well. \square

Hence, if (P_i, E_i) are index pairs in N , these theorems gives a meaningful notion of persistence of Conley index through the decomposition of the following zigzag persistence module:

$$H_p(P_1, E_1) \leftarrow H_p(P_1 \cap P_2, E_1 \cap E_2) \rightarrow H_p(P_2, E_2) \leftarrow \cdots \rightarrow H_p(P_n, E_n). \quad (4.1)$$

Because of the previous two theorems, when one decomposes the above zigzag module, one is actually capturing a changing Conley index. This contrasts the case where one only considers index pairs of the form $(\text{cl}(S_i), \text{mo}(S_i))$, because $(\text{cl}(S_i) \cap \text{cl}(S_{i+1}), \text{mo}(S_i) \cap \text{mo}(S_{i+1}))$ need not be an index pair for any invariant set.

As has been established, the pair $(\text{cl}(S), \text{mo}(S))$ is an index pair, but it need not be an index pair in N . We introduce a canonical approach to transform $(\text{cl}(S), \text{mo}(S))$ to an index pair in N by using the *push forward*.

Definition 4.5. The push forward $\text{pf}(S)$ of a set S in N , N closed, is the set of all simplices in S together with those $\sigma \in N$ such that there exists a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N$ where $\rho(a) \in S$ and $\rho(b) = \sigma$.

If N or \mathcal{V} are not clear from context, we use the notation $\text{pf}(S, N)$, $\text{pf}_{\mathcal{V}}(S)$, or $\text{pf}_{\mathcal{V}}(S, N)$. The next series of results imply that an index pair in N can be obtained by taking the push forward of $(\text{cl}(S), \text{mo}(S))$.

Proposition 4.6. If $S \subseteq K$ is an isolated invariant set with isolating neighborhood N under \mathcal{V} , then $\text{pf}(\text{mo}(S)) \cap \text{cl}(S) = \text{mo}(S)$.

Proof. Note that by definition, $\text{mo}(S) \subseteq \text{cl}(S)$ and $\text{mo}(S) \subseteq \text{pf}(\text{mo}(S))$, so it follows that $\text{mo}(S) \subseteq \text{cl}(S) \cap \text{pf}(\text{mo}(S))$. Hence, it is sufficient to show that $\text{cl}(S) \cap \text{pf}(\text{mo}(S)) \subseteq \text{mo}(S)$. Aiming for a contradiction, assume there exists a $\sigma \in \text{cl}(S) \cap \text{pf}(\text{mo}(S))$ where $\sigma \notin \text{mo}(S)$. This directly implies that $\sigma \in \text{cl}(S) \setminus \text{mo}(S)$. But by Proposition 3.8, $\text{cl}(S) \setminus \text{mo}(S) = S$, so $\sigma \in S$. But since $\sigma \in \text{pf}(\text{mo}(S))$, there exists a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N$ where $\rho(a) \in \text{mo}(S)$ and $\rho(b) = \sigma$. Because $\rho(a) \in \text{mo}(S)$, there exists a $\sigma' \in S$ such that $\rho(a) \leq \sigma'$. This implies that there exists a path $\rho' : \mathbb{Z}_{[a-1,b]} \rightarrow N$ where $\rho(a-1) = \sigma'$ and $\rho(b) = \sigma$, but $\rho(a) \notin S$. Hence, S is not isolated by N , a contradiction. \square

Proposition 4.7. If $S \subseteq K$ is an isolated invariant set with isolating neighborhood N under \mathcal{V} , then $\text{pf}(\text{mo}(S)) \cup \text{cl}(S) = \text{pf}(\text{cl}(S))$.

Proof. Note that since $\text{pf}(\text{mo}(S)) \subseteq \text{pf}(\text{cl}(S))$ and $\text{cl}(S) \subseteq \text{pf}(\text{cl}(S))$, it follows immediately that $\text{pf}(\text{mo}(S)) \cup \text{cl}(S) \subseteq \text{pf}(\text{cl}(S))$. Hence, it is sufficient to show that $\text{pf}(\text{cl}(S)) \subseteq \text{pf}(\text{mo}(S)) \cup \text{cl}(S)$. Aiming for a contradiction, we assume there exists a $\sigma \in \text{pf}(\text{cl}(S))$ such that $\sigma \notin \text{pf}(\text{mo}(S)) \cup \text{cl}(S)$. Note that since $\sigma \notin \text{pf}(\text{mo}(S))$, it follows that there must exist a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N$ where $\rho(i) \notin \text{mo}(S)$ for all i . Else, $\rho(b) = \sigma$ would be in $\text{pf}(\text{mo}(S))$, a contradiction. Hence, there must exist a $\rho(i) \in \text{cl}(S) \setminus \text{mo}(S) = S$ such that $\rho(i+1) \in F_{\mathcal{V}}(\rho(i))$ and $\rho(i+1) \notin \text{cl}(S)$. Note that $\rho(i+1)$ cannot be a face of $\rho(i)$, else $\rho(i+1) \in \text{cl}(S)$. Hence, $[\rho(i+1)] = [\rho(i)]$, which means one can construct a path $\rho' : \mathbb{Z}_{[0,2]} \rightarrow N$ where $\rho'(0) = \rho'(2) = \rho(i)$ and $\rho'(1) = \rho(i+1)$. But this is a path with endpoints in S which is not contained in S , which contradicts S being isolated by N . \square

Proposition 4.8. *If $S \subseteq K$ is an isolated invariant set with isolating neighborhood N , then $\text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S)) = \text{cl}(S) \setminus \text{mo}(S) = S$.*

Proof. First, we note that by Proposition 4.7, we have that $\text{pf}(\text{mo}(S)) \cup \text{cl}(S) = \text{pf}(\text{cl}(S))$. From Proposition 4.6 we have that $\text{pf}(\text{mo}(S)) \cap \text{cl}(S) = \text{mo}(S)$, which together with the fact that $\text{cl}(S) \subseteq \text{mo}(S)$ implies that $\text{pf}(\text{mo}(S)) \cup (\text{cl}(S) \setminus \text{mo}(S)) = \text{pf}(\text{cl}(S))$. Since $\text{pf}(\text{mo}(S))$ and $\text{cl}(S) \setminus \text{mo}(S)$ are disjoint, it follows that $\text{cl}(S) \setminus \text{mo}(S) = \text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S))$. By Proposition 3.8, $(\text{cl}(S), \text{mo}(S))$ is a saturated index pair for S , so it follows that $\text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S)) = \text{cl}(S) \setminus \text{mo}(S) = S$. \square

From these propositions, we get the following.

Theorem 4.9. *If S is an isolated invariant set then $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$ is an index pair in N for S .*

Proof. First, we note that since the index pair $(\text{cl}(S), \text{mo}(S))$ is saturated, it follows that $S = \text{Inv}(\text{cl}(S) \setminus \text{mo}(S)) = \text{cl}(S) \setminus \text{mo}(S)$. But since by Proposition 4.8 $\text{cl}(S) \setminus \text{mo}(S) = \text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S))$, it follows that $S = \text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S)) = \text{Inv}(\text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S)))$, which satisfies condition four of being an index pair in N .

We show that $F_V(\text{pf}(\text{cl}(S))) \cap N \subseteq \text{pf}(\text{cl}(S))$. Let $x \in \text{pf}(\text{cl}(S))$, and assume that $y \in F_V(x) \cap N$. There must be a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N$ where $\rho(a) \in \text{cl}(S)$ and $\rho(b) = x$, by the definition of the push forward. Thus, we can construct an analogous path $\rho' : \mathbb{Z}_{[a,b+1]} \rightarrow N$ where $\rho'(i) = \rho(i)$ for $i \in \mathbb{Z}_{[a,b]}$ and $\rho'(b+1) = y$. Hence, $y \in \text{pf}(\text{cl}(S))$ by definition. Identical reasoning can be used to show that $F_V(\text{pf}(\text{mo}(S))) \cap N \subseteq \text{pf}(\text{mo}(S))$, so $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$ also meets the first two conditions required to be an index pair.

Finally, we show that $F_V(\text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S))) \subseteq N$. By Proposition 4.8, this is equivalent to showing that $F_V(\text{cl}(S) \setminus \text{mo}(S)) \subseteq N$. Since $(\text{cl}(S), \text{mo}(S))$ is an index pair for S , it follows that $F_V(\text{cl}(S) \setminus \text{mo}(S)) \subseteq \text{cl}(S)$. Note that since $N \supseteq S$ is closed, it follows that $\text{cl}(S) \subseteq N$. Hence, $F_V(\text{pf}(\text{cl}(S)) \setminus \text{pf}(\text{mo}(S))) \subseteq N$, and all conditions for an index pair in N are met. \square

An example of an index pair induced by the push forward can be seen in Figure 4.3. Hence, instead of considering a zigzag filtration given by a sequence of index pairs $(\text{cl}(S_1), \text{mo}(S_1))$,




Figure 4.3. All three images depict an index pair for the yellow triangle marked with a black circle. P is given by the red and yellow simplices, while E is given by the red simplices. If N is taken to be the entire rectangle, then the index pairs on the left and the right are not index pairs in N , but the index pair on the right is.

$(\text{cl}(S_2), \text{mo}(S_2)), \dots, (\text{cl}(S_n), \text{mo}(S_n))$, a canonical choice is to instead consider the zigzag filtration given by the the sequence of index pairs

$$(\text{pf}(\text{cl}(S_1)), \text{pf}(\text{mo}(S_1))), (\text{pf}(\text{cl}(S_2)), \text{pf}(\text{mo}(S_2))), \dots, (\text{pf}(\text{cl}(S_n)), \text{pf}(\text{mo}(S_n))).$$

Choosing S_i is highly application specific, so in our implementation we choose $S_i := \text{Inv}_{\mathcal{V}_i}(N)$. This decision together with the previous theorems gives Algorithm 1 for computing the persistence of the Conley Index. Note that in the case where $\text{Inv}_{\mathcal{V}_i}(N)$ is not \mathcal{V}_i -compatible, then S_i is actually \mathcal{V} -compatible with respect to the multivector field $\{V \cap N \mid V \in \mathcal{V}_i\} \cup \{V \cap N \mid V \in \mathcal{V}_i\}$. This adjustment is handled implicitly in the algorithm. The complexity of finding S_i and (P_i, E_i) is linear in $|N|$, so the overall complexity of the scheme is $O(n|N| + M(m) + m^2 \log^2(m))$, where m is the number of simplices that are added and deleted in the filtration.

Input: Sequence of multivector fields $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$, closed set $N \subseteq K$.

Output: Barcodes corresponding to persistence of Conley Index

$i \leftarrow 1$

while $i \leq n$ **do**

$S_i \leftarrow \text{Inv}_{\mathcal{V}_i}(N)$
 $(P_i, E_i) \leftarrow (\text{pf}(\text{cl}(S_i), N), \text{pf}(\text{mo}(S_i), N))$
 $i \leftarrow i + 1$

end

return $\text{zigzagPers}((P_1, E_1) \supseteq (P_1 \cap P_2, E_1 \cap E_2) \subseteq (P_2, E_2) \supseteq \dots \subseteq (P_n, E_n))$

Algorithm 1: Scheme for computing the persistence of the Conley Index, fixed N

Index pairs and barcodes computed by Algorithm 1 can be seen in Figure 4.4.

4.1 Noise-Resilient Index Pairs

Earlier in this chapter, we showed that if one is given two isolated invariant sets S_1 and S_2 under \mathcal{V}_1 and \mathcal{V}_2 , where both are isolated by some N , one can obtain the index pairs $(\text{pf}(\text{cl}(S_1)), \text{pf}(\text{mo}(S_1)))$ and $(\text{pf}(\text{cl}(S_2)), \text{pf}(\text{mo}(S_2)))$ and intersect them using Theorem 4.4 to compute the persistence of the Conley index. However, this approach works best when the isolated invariant sets are very close to each other. In the presence of noise, these index pairs




Figure 4.4. Examples of index pairs computed by using the push forward on multivector fields induced by a differential equation. A sequence of multivector fields was generated from a λ -parametrized differential equation undergoing supercritical Hopf bifurcation [32, Section 11.2]. The consecutive images (from left to right) present a selection from this sequence: the case when $\lambda < 0$ and there is only an attracting fixed point inside N ; the case when $\lambda > 0$ is small and N contains a repelling fixed point, a small attracting periodic trajectory and all connecting trajectories; the case when $\lambda > 0$ is large and the periodic trajectory is no longer contained in N . In all three images, N is given by the rectangle, E is given by the simplices in red, and $P \setminus E$ is given by the simplices in yellow. Note that in the leftmost image, the only invariant set is a triangle which represents an attracting fixed point. For this invariant set, the only relative homology group which is nontrivial is $H_0(P, E)$, which has a single homology generator. In the middle image, the invariant sets represent a repelling fixed point, a periodic attractor, and heteroclinic orbits which connect the repelling fixed point with the periodic attractor. Note that the relative homology has not changed from the leftmost case, so the only nontrivial homology group is $H_0(P, E)$. In the rightmost image, the periodic attractor is no longer entirely contained within N , so the only invariant set corresponds to a repelling fixed point. Here, the only nontrivial homology group is $H_2(P, E)$, which has one generator, so the Conley index has changed. Algorithm 1 captures this change. The persistence barcode output by Algorithm 1 is below index pairs, where a H_0 generator (white bar) lasts until the periodic trajectory leaves N , at which point it is replaced by an H_2 generator (dark gray bar).

can perform quite poorly, for multiple reasons. We consider a couple of these reasons in turn and devise methods to combat noise.

4.1.1 Shrinking E

The previous strategy given for producing index pairs in N produces saturated index pairs. Equivalently, the cardinality of $P \setminus E$ is minimized. This is problematic in the presence of noise, where if \mathcal{V}_2 is a slight perturbation of \mathcal{V}_1 we frequently have that $\text{Inv}_{\mathcal{V}_1}(N) \neq \text{Inv}_{\mathcal{V}_2}(N)$. This gives a perturbation in our generated index pairs and in particular a perturbation in $P \setminus E$. As the Conley Index is obtained by taking relative homology, taking the intersection of index pairs (P_1, E_1) and (P_2, E_2) where $P_i \setminus E_i = \text{Inv}_{\mathcal{V}_i}(N)$ can result in a “breaking” of bars in the barcode. An example can be seen in Figure 4.5, where because of noise, the two $P \setminus E$ do not overlap, and hence a 2-dimensional homology class which intuitively should persist throughout the interval does not. In Figure 4.5, the Conley indices of the invariant sets consisting of the singleton critical triangles in \mathcal{V}_1 and \mathcal{V}_2 (the left and right multivector fields) have rank 1 in dimension 2 because the homology group H_2 of P (which is all colored simplices) relative to E (which is all red simplices) has rank 1. However, the generators for $H_2(P_1, E_1)$ and $H_2(P_2, E_2)$ are both in the intersection field $\mathcal{V}_1 \cap \mathcal{V}_2$. Hence, rather than one generator persisting through all three multivector fields, we get two bars that overlap at the intersection field. The difficulty is rooted in the fact that the sets $W_1 = P_1 \setminus E_1$, $W_2 = P_2 \setminus E_2$, and $W_{12} = (P_1 \cap P_2) \setminus (E_1 \cap E_2)$ do not have a common intersection.

To address this problem, we propose an algorithm to expand the size of $P \setminus E$. It is important to note that a balance is needed to ensure a large E as well as a large $P \setminus E$. If E_1 and E_2 are too small, then it is easy to see that E_1 and E_2 may not intersect as expected even though consecutive vector fields are very similar. The following proposition is very useful for computing a balanced index pair.

Proposition 4.10. *Let (P, E) be an index pair for S in N under \mathcal{V} . If $V \subseteq E$ is a regular multivector where $E' := E \setminus V$ is closed, then (P, E') is an index pair for S in N .*




Figure 4.5. Infeasibility of the index pair $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$: The sets $E = \text{pf}(\text{mo}(S))$ are colored pink in all three images, while the invariant sets which equal $P \setminus E$ are golden in all three images. (left) $\mathcal{V}_1 : P_1 \setminus E_1$ consists of a single golden triangle; (right) $\mathcal{V}_2 : P_2 \setminus E_2$ consists of the single golden triangle; (middle) $(P_1 \cap P_2) \setminus (E_1 \cap E_2)$ consists of two golden triangles (excluding the edge between them) in the intersection field $\mathcal{V}_1 \cap \mathcal{V}_2$. The barcode for index pairs is depicted by two dark gray bars, each of which represents a 2-dimensional homology generator. Ideally, these would be a single bar.

Proof. Note that since P does not change, it is immediate the P satisfies $F_V(P) \cap N \subseteq P$, and the first condition of an index pair in N is met.

We show that $F_V(E') \cap N \subseteq E$. For a contradiction, we assume that there exists an $x \in E'$ such that there is a $y \in F_V(x) \cap N$, $y \notin E'$. Note that if $y \leq x$, then by definition y is in the closure of x . Since E is closed and $x \in E$, it follows that $y \in E$. But since $y \notin E'$, it follows that $y \in V$. But this is a contradiction since by assumption, $E \setminus V$ is closed. Hence, by definition of F_V , since $y \notin \text{cl}(x)$, y and x must be in the same multivector. Note that in such a case $y \notin V$, as if it were, x would not be in E' since x and y are in the same multivector. This implies that $F_V(E) \cap N \not\subseteq E$, as $y \in F_V(x)$ and $y \notin E$, a contradiction. Thus, we conclude that $F_V(E') \cap N \subseteq E$.

Now, we show that $F_V(P \setminus E') \subseteq N$. Assume there exists an $x \in P$ satisfying $F_V(x) \setminus N \neq \emptyset$. Then since (P, E) is an index pair in N , it follows that $x \in E$. To contradict, we assume $x \notin E'$. Hence, $x \in V$. We let $y \in F_V(x) \setminus N$. By definition of F_V , either $y \in \text{cl}(x)$ or y and x are in the same multivector. In the later case, $y \in N$ as it was assumed that $V \subseteq E$, a contradiction. Hence, $y \leq x$. But this implies that $y \in \text{cl}(x) \subseteq E \subseteq N$, a contradiction. Ergo, $x \in E'$, and $F_V(P \setminus E') \subseteq N$.

Finally, we show that $\text{Inv}(P \setminus E') = \text{Inv}(P \setminus E)$. Trivially, $\text{Inv}(P \setminus E) \subseteq \text{Inv}(P \setminus E')$, so it is sufficient to show that $\text{Inv}(P \setminus E') \subseteq \text{Inv}(P \setminus E)$. For a contradiction, assume that there exists an $x \in \text{Inv}(P \setminus E')$, $x \notin \text{Inv}(P \setminus E)$. Thus, there exists an essential solution $\rho : \mathbb{Z} \rightarrow P \setminus E'$ where for some k , $y := \rho(k) \in V$. Since V is regular, we assume without loss of generality that $z := \rho(k+1) \notin V$. Hence, $z \in \text{cl}(y)$. In addition, since $z \in \text{cl}(V)$, we have that $z \in \text{cl}(E) = E$. Therefore, $z \in E \setminus V = E'$. But ρ is a solution in $P \setminus E'$, a contradiction. \square

Figure 4.6 illustrates how enlarging $P \setminus E$ by removing regular vectors as Proposition 4.10 suggests can help mitigate the effects of noise on computing Conley index persistence. Contrast this example with the example in Figure 4.5. Denoting $W_i = P_i \setminus E_i$ for $i = 1, 2$ and $W_{12} = (P_1 \cap P_2) \setminus (E_1 \cap E_2)$ in both figures, we see that $W_1 \cap W_{12} \cap W_2$ is empty in Figure 4.5 while in Figure 4.6 it consists of three critical simplices each marked with a circle. Hence, in Figure 4.6 a single generator persists throughout the interval, unlike in Figure 4.5.




Figure 4.6. Enlarging $P \setminus E$ which is gold in all three pictures while E is colored pink. (left) \mathcal{V}_1 ; (right) \mathcal{V}_2 ; (middle) $\mathcal{V}_1 \cap \mathcal{V}_2$. Note that there is one bar in the barcode, in contrast with Figure 4.5

Computation

We give a method for computing a noise-resilient index pair by using these techniques. Note that by Theorem 4.9, we have that $(\mathbf{pf}(\mathbf{cl}(S)), \mathbf{pf}(\mathbf{mo}(S)))$ is an index pair for invariant set S in N . Hence, we adopt the strategy of taking $P = \mathbf{pf}(\mathbf{cl}(S))$ and $E = \mathbf{pf}(\mathbf{mo}(S))$, and we aim to find some collection $R \subseteq E$ so that $(P, E \setminus R)$ remains an index pair in N . Finding an appropriate R is a difficult balancing act: one wants to find an R so that $P \setminus (E \setminus R)$ is sufficiently large, so as to capture perturbations in the isolated invariant set as described in the previous section, but not so large that E is small and perturbations in E are not captured. If R is chosen to be as large as possible, then a small shift in E may results in $(E \setminus R) \cap (E' \setminus R')$ having a different topology than E or E' leading to a “breaking” of barcodes analogous to the case described.

Before we give an algorithm for outputting such an R , we first define a δ -collar.

Definition 4.11. *We define the δ -collar of an invariant set $S \subseteq K$ recursively:*

1. *The 0-collar of S is $\mathbf{cl}(S)$.*
2. *For $\delta > 0$, the δ -collar of S is the set of simplices σ in the $(\delta - 1)$ -collar of S together with those simplices τ where τ is a face of some σ with a face τ' in the $(\delta - 1)$ -collar of S .*

For an isolated invariant set S , we will let $C_\delta(S)$ denote the δ -collar of S . Together with Proposition 4.10, δ -collars give a natural algorithm for finding an R to enlarge $P \setminus E$.

In particular, we use Algorithm 2 for this purpose.

Theorem 4.12. *Let R be the output of Algorithm 2 applied to index pair (P, E) in N for isolated invariant set S . The pair $(P, E \setminus R)$ is an index pair for S in N .*

Proof. To contradict, we assume that the R output by Algorithm 2 results in $(P, E \setminus R)$ is not an index pair. We note by inspection of the algorithm that multivectors are removed sequentially, so there exists some first V such that $(P, E \setminus R_V)$ is an index pair for S in N but $(P, E \setminus (R_V \cup V))$ is not an index pair for S in N , where R_V denotes the R variable in Algorithm 2 before V is added to it. By inspection of the algorithm, we observe that since

Input: Isolated invariant set S with respect to \mathcal{V} contained in some closed set N ,
Index pair (P, E) in N with respect to \mathcal{V} , $\delta \in \mathbb{Z}$

Output: List of simplices R such that $(P, E \setminus R)$ is an index pair for S in N .

```

 $R \leftarrow \text{new set}()$ 
 $vecSet \leftarrow \{[\sigma] \in \mathcal{V} \mid [\sigma] \subseteq E \cap C_\delta(S) \wedge [\sigma] \cap \partial(E) \neq \emptyset \wedge [\sigma] \cap \partial(P) = \emptyset\}$ 
 $vec \leftarrow \text{new queue}()$ 
appendAll( $vec, vecSet$ )
while size( $vec$ ) > 0 do
     $[\sigma] \leftarrow \text{pop}(vec)$ 
    if isClosed( $(E \setminus R) \setminus [\sigma]$ ) and  $[\sigma] \subseteq E \setminus R$  and isRegular( $[\sigma]$ ) then
         $R \leftarrow R \cup [\sigma]$ 
         $mouthVecs \leftarrow \{[\tau] \mid \tau \in \text{mo}([\sigma]) \wedge \dim(\tau) = \dim(\sigma) - 1 \wedge [\tau] \subseteq C_\delta(S)\}$ 
        appendAll( $vec, mouthVecs$ )
    end
end
return  $R$ 

```

Algorithm 2: findR($S, P, E, \mathcal{V}, \delta$)

V was added to R_V , it must be that V is a regular vector, and $E \setminus (R \cup V)$ is closed, and $V \subseteq E \setminus R_V$. Proposition 4.10 directly implies that $(P, (E \setminus R) \setminus V)$ is an index pair, which is a contradiction. Hence, there can exist no such V , so it follows that $(P, E \setminus R)$ must be an index pair for S in N . \square

Hence, Algorithm 2 provides a means by which the user may enlarge $P \setminus E$. If the dimension of K is fixed, then we can determine if $(E \setminus R) \setminus [\sigma]$ is closed in $O(|E|)$ time. Each τ can be considered at most $|E|$ times when considering adding $[\tau]$ to the queue. Hence, the total complexity is $O(|E|^3)$, assuming we have pre-computed which multivectors are in $C_\delta(S)$ (requiring $O(|E|)$ preprocessing) and which multivectors are regular (which can be done in matrix multiplication time in the size of the multivector). This bound, however, is very pessimistic, as each simplex is not the face of $|E|$ other simplices in E .

As this algorithm is parameterized, a robust choice of δ may be application-specific. We also include some demonstrations of the effectiveness of using this technique. A real instance of the difficulty can be seen in Figure 4.7, while the application of Algorithm 2 with $\delta = 5$ to solve the problem is found in Figure 4.9.




Figure 4.7. Index pairs on two slightly perturbed multivector fields (left, right) and their intersection (middle). The isolating neighborhood N is the portion of the rectangle where the multivectors are drawn, E is in red, and $P \setminus E$ is in yellow. Note that we have the same difficulty as in Figure 4.5, where there are two homology generators in the intersection multivector field, so we get a broken bar code.




Figure 4.8. A zoomed in version of Figure 4.7.




Figure 4.9. The same index pairs as in Figure 4.7 with the same color scheme, but after applying Algorithm 2 to reduce the size of E . This forces a 2-dimensional homology generator to persist across both multivector fields (left, right) and their intersection (middle).

4.1.2 Enlarging P

In the previous subsection, we addressed noise in persistence that is caused because E is too large. In this section, we consider noise that occurs because P is too small. This typically occurs when the relevant isolated invariant sets are attractors. Consider Figure 4.10. In the first and third images, we find an index pair for the isolated invariant sets M_1 , M_2 by using this technique. The simplices in yellow and blue are the simplices in N , and the simplices in blue are in P . In these two cases, $E = \emptyset$. Intuitively, we expect that a 1-dimensional homology class would persist from the left to the right. However, this is not the case. The middle image shows the index pair $(\text{cl}(M_1) \cap \text{cl}(M_2), \text{mo}(M_1) \cap \text{mo}(M_2))$ in blue, and $H_1(\text{cl}(M_1) \cap \text{cl}(M_2), \text{mo}(M_1) \cap \text{mo}(M_2)) = 0$. Thus, a one-dimensional homology class does not persist from the first to third image. We devise a new method to compute index pairs that permits us to circumvent this issue. For a set of multivectors \mathcal{A} , we let $\langle \mathcal{A} \rangle := \cup_{A \in \mathcal{A}} A$.

Proposition 4.1.1. *Let (P, E) denote an index pair in N under \mathcal{V} , and let \mathcal{A} denote a set of multivectors such that $\langle \mathcal{A} \rangle \subseteq N$, $\langle \mathcal{A} \rangle \cap E = \emptyset$ and $\text{mo}(\langle \mathcal{A} \rangle) \subseteq P$. The pair $(P \cup \langle \mathcal{A} \rangle, E)$ is an index pair in N under \mathcal{V} .*

Proof. It is immediate that $\text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E) = \text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E)$. In addition, note that E has not changed, so $F_{\mathcal{V}}(E) \cap N \subseteq E$ by hypothesis. Hence, it is sufficient to show that $F_{\mathcal{V}}(P \cup \langle \mathcal{A} \rangle) \cap N \subseteq P \cup \langle \mathcal{A} \rangle$ and $F_{\mathcal{V}}(P \cup \langle \mathcal{A} \rangle) \setminus E \subseteq N$.

First, we consider $F_{\mathcal{V}}(\langle \mathcal{A} \rangle)$. The set $\langle \mathcal{A} \rangle$ is a union of multivectors, so by definition, $F_{\mathcal{V}}(\langle \mathcal{A} \rangle) = \langle \mathcal{A} \rangle \cup \text{cl}(\langle \mathcal{A} \rangle)$. By definition, $\text{cl}(\langle \mathcal{A} \rangle) = \langle \mathcal{A} \rangle \cup \text{mo}(\langle \mathcal{A} \rangle)$. Thus, it follows that $F_{\mathcal{V}}(\langle \mathcal{A} \rangle) = \langle \mathcal{A} \rangle \cup \text{mo}(\langle \mathcal{A} \rangle)$. By assumption, $\text{mo}(\langle \mathcal{A} \rangle) \subseteq P$, so it follows that $F_{\mathcal{V}}(\langle \mathcal{A} \rangle) \subseteq \langle \mathcal{A} \rangle \cup P \subseteq N$.

Now, we show that $F_{\mathcal{V}}(P \cup \langle \mathcal{A} \rangle) \cap N \subseteq P \cup \langle \mathcal{A} \rangle$. By the definition of F , it follows that $F_{\mathcal{V}}(P \cup \langle \mathcal{A} \rangle) = F_{\mathcal{V}}(P) \cup F_{\mathcal{V}}(\langle \mathcal{A} \rangle)$. The pair (P, E) is an index pair in N , so it follows that $F_{\mathcal{V}}(P) \cap N \subseteq P$. We have already shown that $F_{\mathcal{V}}(\langle \mathcal{A} \rangle) \subseteq P \cup \langle \mathcal{A} \rangle \subseteq N$, so it follows that $F_{\mathcal{V}}(P \cup \langle \mathcal{A} \rangle) \cap N \subseteq P \cup \langle \mathcal{A} \rangle$.

Now, we move to showing that $F_{\mathcal{V}}((P \cup \langle \mathcal{A} \rangle) \setminus E) \subseteq N$. By assumption, $\langle \mathcal{A} \rangle \cap E = \emptyset$. Thus, $(P \cup \langle \mathcal{A} \rangle) \setminus E = (P \setminus E) \cup \langle \mathcal{A} \rangle$. Hence, $F_{\mathcal{V}}((P \cup \langle \mathcal{A} \rangle) \setminus E) = F_{\mathcal{V}}(P \setminus E) \cup F_{\mathcal{V}}(\langle \mathcal{A} \rangle)$.

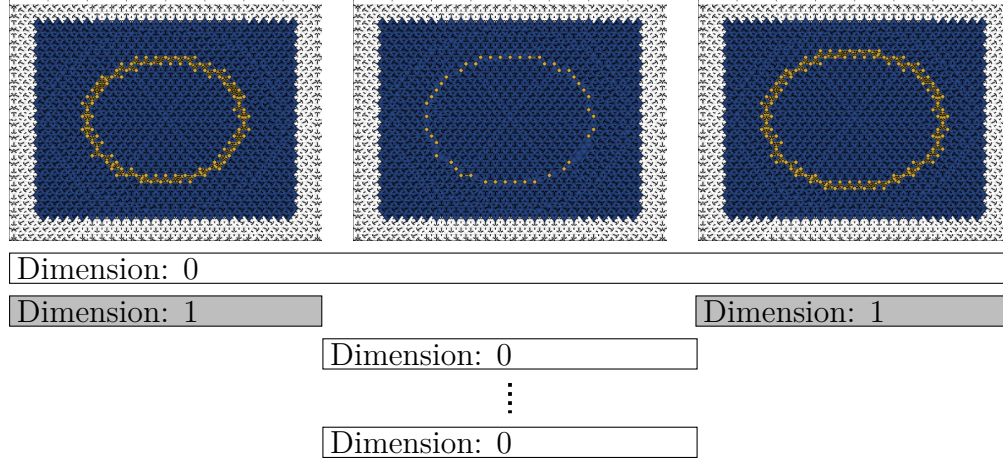


Figure 4.10. In all three images, the isolating neighborhood N is given by the blue and the yellow simplices, while an isolated invariant set in N is given by the yellow simplices. In addition, in the left and the right images, we can obtain an index pair in N by taking the yellow simplices to be P and letting $E := \emptyset$. If (P_1, E_1) denotes the index pair on the left and (P_2, E_2) denotes the index pair in N on the right, then the index pair in the center is given by $(P_1 \cap P_2, E_1 \cap E_2)$. The yellow simplices in the middle are those simplices in $P_1 \cap P_2$. Intuitively, one would expect a 0-dimensional bar and a 1-dimensional bar to persist through all three images. However, when one takes the intersection of the two index pairs, one no longer has an annulus, so the 1-dimensional bar does not persist through all three images. Instead, we obtain several short, 0-dimensional bars. We show the barcode beneath the images, while excluding several of the short 0-dimensional bars.

The pair (P, E) is an index pair in N , so $F_V(P \setminus E) \subseteq N$. We have already shown that $F_V(\langle \mathcal{A} \rangle) \subseteq N$, so it follows that $F_V((P \cup \langle \mathcal{A} \rangle) \setminus E) \subseteq N$.

Hence, we conclude that $(P \cup \langle \mathcal{A} \rangle, E)$ is an index pair for $\text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E)$ in N . \square

However, Proposition 4.1.1 does not imply that $\text{Inv}(P \setminus E) = \text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E)$. We strengthen Proposition 4.1.1 to account for this.

Proposition 4.13. *Let N denote a closed set, and let \mathcal{M} denote the minimal Morse decomposition for $\text{Inv}(N)$. Furthermore, let (P, E) denote an index pair in N for an isolated invariant set S in N . If \mathcal{A} is a set of multivectors where $\langle \mathcal{A} \rangle \subseteq N$, $\langle \mathcal{A} \rangle \cap E = \emptyset$, $\text{mo}(\langle \mathcal{A} \rangle) \subseteq P$, and $\langle \mathcal{A} \rangle \cap M = \emptyset$ for $M \in \mathcal{M}$ where $M \not\subseteq S$, then $(P \cup \langle \mathcal{A} \rangle, E)$ is an index pair in N for S .*

Proof. By Proposition 4.1.1, $(P \cup \langle \mathcal{A} \rangle, E)$ is an index pair for $\text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E)$, so it is sufficient to show that $\text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E) = S$. Trivially, $S \subseteq \text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E)$. Hence, we aim to show the reverse inclusion. To contradict, assume that $\sigma \in \text{Inv}((P \cup \langle \mathcal{A} \rangle) \setminus E) \setminus S$. Then there exists some essential solution $\rho : \mathbb{Z} \rightarrow (P \cup \langle \mathcal{A} \rangle) \setminus E$ where $\rho(0) = \sigma$. In addition, there must exist $M_1, M_2 \in \mathcal{M}$ such that $\alpha(\rho) \subseteq M_1$ and $\omega(\rho) \subseteq M_2$.

We have that $\text{im}(\rho) \subseteq (P \cup \langle \mathcal{A} \rangle) \setminus E$, so $M_1 \cap (P \cup \langle \mathcal{A} \rangle) \setminus E \neq \emptyset$, and similarly for M_2 . We claim that $M_1, M_2 \subseteq (P \cup \langle \mathcal{A} \rangle) \setminus E$. If there were a $\eta \in M_1$ but $\eta \notin (P \cup \langle \mathcal{A} \rangle) \setminus E$, then since there is a $\tau \in M_1 \cap (P \cup \langle \mathcal{A} \rangle) \setminus E$, Proposition 3.11 implies there exists a path from η to τ and there exists a path from τ to η . But by the definition of an index pair in N , it follows that $\eta \in E$. Since there is a path from η to τ , but $\tau \in (P \cup \langle \mathcal{A} \rangle) \setminus E$, this contradicts the requirement that $F_V(E) \cap N \subseteq E$. Hence, no such η can exist, and $M_1, M_2 \subseteq (P \cup \langle \mathcal{A} \rangle) \setminus E$.

By assumption, $\langle \mathcal{A} \rangle \cap M_1 = \emptyset$ if $M_1 \not\subseteq S$ (and similarly for M_2), so $M_1 \cap (\langle \mathcal{A} \rangle \setminus P) = \emptyset$ (and similarly for M_2). Thus, $M_1, M_2 \subseteq P \setminus E$. But $S = \text{Inv}(P \setminus E)$, so $M_1, M_2 \subseteq S$. But this implies there is a path from S to σ and a path from σ to S , which implies that S is not isolated by N , a contradiction. \square

Proposition 4.13 provides a natural avenue for finding index pairs that are suitable for computing the persistence of the Conley-Morse graph. If we are given an index pair (P, E) in N for a Morse set M , we can incrementally add multivectors to P that satisfy the conditions

of Proposition 4.13, up to a specified distance away from the original P . We give an example of this in Figure 4.11.




Figure 4.11. When one uses Proposition 4.13 to thicken the index pairs in Figure 4.10, a 1-dimensional homology class persists through all three images, which matches the intuition.

4.2 A Simple Tracking Algorithm

In this chapter, we established the persistence of the Conley index of invariant sets in consecutive multivector fields which are isolated by a single isolating neighborhood and we showed how to intelligently choose index pairs. In this section, we develop an algorithm to “track” an invariant set over a sequence of isolating neighborhoods. A classic example is a hurricane, where if one were to sample wind velocity at times t_0, t_1, \dots, t_n , there may be no fixed N which captures the eye of the hurricane at all t_i without also capturing additional, undesired invariant sets at some t_j .

4.2.1 Changing the Isolating Neighborhood

Thus far, we have defined a notion of persistence of the Conley Index for some fixed isolating neighborhood N and simplicial complex K . This setting is very inflexible—one may want to incorporate domain knowledge to change N to capture changing features of a sequence of sampled dynamics. We now extend our previous results to a setting where N need not be fixed. Throughout this section, we consider multivector fields $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$ with corresponding isolated invariant sets S_1, S_2, \dots, S_n . In addition, we assume that there

exist isolating neighborhoods N_1, N_2, \dots, N_{n-1} where N_i isolates both S_i and S_{i+1} . We will also require that if $1 < i < n$, the invariant set S_i is isolated by $N_i \cup N_{i+1}$. Note that for each i where $1 < i < n$, there exist two index pairs for S_i : one index pair $(P_i^{(i-1)}, E_i^{(i-1)})$ in N_{i-1} and another index pair $(P_i^{(i)}, E_i^{(i)})$ in N_i . In the case of $i = 1$, there is only one index pair $(P_1^{(1)}, E_1^{(1)})$ for S_i . Likewise, in the case of $i = n$, there is a single index pair $(P_n^{(n-1)}, E_n^{(n-1)})$.

By applying the techniques of Section 5.1, we obtain a sequence of persistence modules:

$$\begin{aligned}
H_p(P_1^{(1)}, E_1^{(1)}) &\longleftarrow H_p(P_1^{(1)} \cap P_2^{(1)}, E_1^{(1)} \cap E_2^{(1)}) \longrightarrow H_p(P_2^{(1)}, E_2^{(1)}) \\
H_p(P_2^{(2)}, E_2^{(2)}) &\longleftarrow H_p(P_2^{(2)} \cap P_3^{(2)}, E_2^{(2)} \cap E_3^{(2)}) \longrightarrow H_p(P_3^{(2)}, E_3^{(2)}) \\
H_p(P_3^{(3)}, E_3^{(3)}) &\longleftarrow H_p(P_3^{(3)} \cap P_4^{(3)}, E_3^{(3)} \cap E_4^{(3)}) \longrightarrow H_p(P_4^{(3)}, E_4^{(3)}) \\
&\vdots \\
H_p(P_{n-1}^{(n-1)}, E_{n-1}^{(n-1)}) &\longleftarrow H_p(P_{n-1}^{(n-1)} \cap P_n^{(n-1)}, E_{n-1}^{(n-1)} \cap E_n^{(n-1)}) \longrightarrow H_p(P_n^{(n-1)}, E_n^{(n-1)}).
\end{aligned} \tag{4.2}$$

In the remainder of this subsection, we develop the theory necessary to combine these modules into a single module. Without any loss of generality, we will combine the first modules into a single module, which will imply a method to combine all of the modules into one.

First, we note that by Theorem 3.9, we have that $H_p(P_2^{(1)}, E_2^{(1)}) \cong H_p(P_2^{(2)}, E_2^{(2)})$. To combine the persistence modules

$$\begin{aligned}
H_p(P_1^{(1)}, E_1^{(1)}) &\longleftarrow H_p(P_1^{(1)} \cap P_2^{(1)}, E_1^{(1)} \cap E_2^{(1)}) \longrightarrow H_p(P_2^{(1)}, E_2^{(1)}) \\
H_p(P_2^{(2)}, E_2^{(2)}) &\longleftarrow H_p(P_2^{(2)} \cap P_3^{(2)}, E_2^{(2)} \cap E_3^{(2)}) \longrightarrow H_p(P_3^{(2)}, E_3^{(2)}).
\end{aligned} \tag{4.3}$$

into a single module, it is either necessary to explicitly find a simplicial map which induces an isomorphism $\phi : H_p(P_2^{(1)}, E_2^{(1)}) \rightarrow H_p(P_2^{(2)}, E_2^{(2)})$, or to construct some other index pair for S_2 denoted (P, E) such that $P_2^{(1)}, P_2^{(2)} \subset P$ and $E_2^{(1)}, E_2^{(2)} \subset E$. This would allow substituting both occurrences of $(P_2^{(1)}, E_2^{(1)})$ or $(P_2^{(2)}, E_2^{(2)})$ for (P, E) , and allow the combining of all the modules in Equation 4.2 into a single module. Since constructing the isomorphism given by

Theorem 3.9 is fairly complicated, we opt for the second approach. First, we define a special type of index pair that is sufficient for our approach.

Definition 4.14 (Strong Index Pair). *Let (P, E) be an index pair for S under \mathcal{V} . The index pair (P, E) is a strong index pair for S if for each $\tau \in E$, there exists a $\sigma \in S$ such that there is a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow P$ where $\rho(a) = \sigma$ and $\rho(b) = \tau$.*

Intuitively, a strong index pair (P, E) for S is an index pair for S where each simplex $\tau \in E$ is reachable from a path originating in S . Strong index pairs have the following useful property.

Theorem 4.15. *Let S denote an invariant set isolated by N , N' , and $N \cup N'$ under \mathcal{V} . If (P, E) and (P', E') are strong index pairs for S in N , N' under \mathcal{V} , then the pair*

$$(\text{pf}_{N \cup N'}(P \cup P'), \text{pf}_{N \cup N'}(E \cup E'))$$

is a strong index pair for S in $N \cup N'$ under \mathcal{V} .

Proof. We proceed by showing that the pair meets the requirements to be a strong index pair in $N \cup N'$. First, note that for all $\sigma \in \text{pf}(P \cup P')$, if $\tau \in F_{\mathcal{V}}(\sigma) \cap (N \cup N')$, then it follows by the definition of the push forward that $\tau \in \text{pf}(P \cup P')$. Hence, $F_{\mathcal{V}}(\text{pf}(P \cup P')) \cap (N \cup N') \subset \text{pf}(P \cup P')$. Analogous reasoning shows that $F_{\mathcal{V}}(\text{pf}(E \cup E')) \cap (N \cup N') \subset \text{pf}(E \cup E')$.

Hence, we proceed to show that $F_{\mathcal{V}}(\text{pf}(P \cup P') \setminus \text{pf}(E \cup E')) \subset N \cup N'$. To contradict, assume that there exists a $\sigma \in \text{pf}(P \cup P') \setminus \text{pf}(E \cup E')$ so that there is a $\tau \in F_{\mathcal{V}}(\sigma)$ where $\tau \notin N \cup N'$. Since $\sigma \in \text{pf}(P \cup P')$, there must exist a $x \in P \cup P'$ such that there is a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N \cup N'$ where $\rho(a) = x$ and $\rho(b) = \sigma$. Without loss of generality, we assume that $x \in P$. Note that if $\sigma \in P$, then $\sigma \in P \setminus E$ by our assumption and hence $F_{\mathcal{V}}(\sigma) \subseteq N \subseteq N \cup N'$ contradicting our assumption that $F_{\mathcal{V}}(\sigma) \ni \tau \notin (N \cup N')$. Hence, $\sigma \notin P$. Note that since (P, E) is an index pair, we have that $F_{\mathcal{V}}(P \setminus E) \subset P$. Hence, there must be some i such that $\rho(i) \in E$. Hence, $\sigma \in \text{pf}(E \cup E')$, a contradiction. Thus, no such σ can exist.

Now, we show that $S = \text{Inv}(\text{pf}(P \cup P') \setminus \text{pf}(E \cup E'))$. Note that since S is isolated by $N \cup N'$ and every $\sigma \in E \cup E'$ is reachable by a path originating at $\tau \in S$, it follows

that $\mathbf{pf}(E \cup E') \cap S = \emptyset$, else S is not isolated by $N \cup N'$. Hence, this implies that $S \subset \mathbf{Inv}(\mathbf{pf}(P \cup P') \setminus \mathbf{pf}(E \cup E'))$.

To show that $\mathbf{Inv}(\mathbf{pf}(P \cup P') \setminus \mathbf{pf}(E \cup E')) \subset S$, we first note that if $x \in \mathbf{pf}(P \cup P')$ and $x \notin P \cup P'$, then it follows that $x \in \mathbf{pf}(E \cup E')$. This is because if there exists a $\sigma \in P$ such that there is a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N \cup N'$ which satisfies $\rho(a) = \sigma$ and $\rho(b) = x$, there must be some $\rho(i) \in E$. If there is no such $\rho(i)$, then this contradicts requirement 2 of Definition 3.7, which states that $F_V(P \setminus E) \subset P$. Hence, by the definition of the push forward, it follows that $x \in \mathbf{pf}(E \cup E')$. Thus, it follows that $\mathbf{pf}(P \cup P') \setminus \mathbf{pf}(E \cup E') \subset (P \cup P') \setminus (E \cup E')$. Thus, every essential solution in $\mathbf{pf}(P \cup P') \setminus \mathbf{pf}(E \cup E')$ is also an essential solution in $(P \cup P') \setminus (E \cup E')$. It remains to be shown that every essential solution in $(P \cup P') \setminus (E \cup E')$ is also an essential solution in $P \setminus E$.

For a contradiction, assume that there exists an essential solution $\rho : \mathbb{Z} \rightarrow (P \cup P') \setminus (E \cup E')$ such that there exists an i where $\rho(i) \notin P \setminus E$. Note that since $\rho(i) \in (P \cup P') \setminus (E \cup E')$, it follows that $\rho(i) \notin E$. Together, these two facts imply that $\rho(i) \notin P$. Hence, it follows that $\rho(i) \in P' \setminus E'$. We claim in particular that for all $j \in \mathbb{Z}$, $\rho(j) \in P' \setminus E'$. Note that if there exists a $j < i$ such that $\rho(j) \notin P' \setminus E'$, then it follows that $\rho(j) \in P \setminus E$. Since (P, E) is an index pair, we have that $F_V(P \setminus E) \subseteq P$. Thus, there must exist some k where $j < k < i$ where $\rho(k) \in E$, but this contradicts that $\rho(\mathbb{Z}) \subseteq (P \cup P') \setminus (E \cup E')$. Hence, no such j exists. Analogous reasoning shows that there is no $j > i$ where $\rho(j) \notin P' \setminus E'$. It follows that $\rho(\mathbb{Z}) \subseteq \mathbf{Inv}(P' \setminus E')$. But this implies that $\rho(\mathbb{Z}) \subseteq S$. Note that (P, E) and (P', E') are both index pairs for S , so it follows that $\rho(\mathbb{Z}) \subseteq P \setminus E, P' \setminus E'$. This contradicts our assumption that $\rho(i) \notin P \setminus E$. Thus, there is no such ρ . Hence, we have that $\mathbf{Inv}((P \cup P') \setminus (E \cup E')) \subseteq \mathbf{Inv}(P \setminus E), \mathbf{Inv}(P' \setminus E') = S$, which implies that $\mathbf{Inv}(\mathbf{pf}(P \cup P') \setminus \mathbf{pf}(E \cup E')) \subseteq S$.

To see that $(\mathbf{pf}(P \cup P'), \mathbf{pf}(E \cup E'))$ is a strong index pair, note that there must exist a path from S to every $\sigma \in E \cup E'$, and since $\mathbf{pf}(E \cup E')$ is the set of simplices σ for which there exists a path from $E \cup E'$ to σ , it follows easily from path surgery that there exists a path from S to σ . Hence, $(\mathbf{pf}(P \cup P'), \mathbf{pf}(E \cup E'))$ is a strong index pair for S in N . \square

Crucially, this theorem gives a persistence module

$$H_p(P_2^{(1)}, E_2^{(1)}) \longrightarrow H_p(\text{pf}(P_2^{(1)} \cup P_2^{(2)}), \text{pf}(E_2^{(1)} \cup E_2^{(2)})) \longleftarrow H_p(P_2^{(2)}, E_2^{(2)}) \quad (4.4)$$

where the arrows are given by the inclusion. Note that since these are all index pairs for the same S , it follows that we have $H_p(P_2^{(1)}, E_2^{(1)}) \cong H_p(\text{pf}(P_2^{(1)} \cup P_2^{(2)}), \text{pf}(E_2^{(1)} \cup E_2^{(2)})) \cong H_p(P_2^{(2)}, E_2^{(2)})$. Hence, we will substitute $H_p(\text{pf}(P_2^{(1)} \cup P_2^{(2)}), \text{pf}(E_2^{(1)} \cup E_2^{(2)}))$ into the persistence module. By using the modules in Equation 4.2, we get a new sequence of persistence modules

$$\begin{aligned} H_p(P_1^{(1)}, E_1^{(1)}) &\longleftarrow H_p(P_1^{(1)} \cap P_2^{(1)}, E_1^{(1)} \cap E_2^{(1)}) \longrightarrow H_p(\text{pf}(P_2^{(1)} \cup P_2^{(2)}), \text{pf}(E_2^{(1)} \cup E_2^{(2)})) \\ H_p(\text{pf}(P_2^{(1)} \cup P_2^{(2)}), \text{pf}(E_2^{(1)} \cup E_2^{(2)})) &\longleftarrow H_p(P_2^{(2)} \cap P_3^{(2)}, E_2^{(2)} \cap E_3^{(2)}) \longrightarrow H_p(\text{pf}(P_3^{(2)} \cup P_3^{(3)}), \text{pf}(E_3^{(2)} \cup E_3^{(3)})) \\ H_p(\text{pf}(P_3^{(2)} \cup P_3^{(3)}), \text{pf}(E_3^{(2)} \cup E_3^{(3)})) &\longleftarrow H_p(P_3^{(3)} \cap P_4^{(3)}, E_3^{(3)} \cap E_4^{(3)}) \longrightarrow H_p(\text{pf}(P_4^{(3)} \cup P_4^{(4)}), \text{pf}(E_4^{(3)} \cup E_4^{(4)})) \\ &\vdots \\ H_p(\text{pf}(P_{n-1}^{(n-2)} \cup P_{n-1}^{(n-1)}), \text{pf}(E_{n-1}^{(n-2)} \cup E_{n-1}^{(n-1)})) &\longleftarrow H_p(P_{n-1}^{(n-1)} \cap P_n^{(n-1)}, E_{n-1}^{(n-1)} \cap E_n^{(n-1)}) \longrightarrow H_p(P_n^{(n-1)}, E_n^{(n-1)}). \end{aligned} \quad (4.5)$$

which can immediately be combined into a single persistence module.

This approach is not without its disadvantages, however. Namely, if (P, E) and (P', E') are index pairs for S in N and N' , it requires that (P, E) and (P', E') are strong index pairs and that S is isolated by $N \cup N'$. Fortunately, the push forward approach to computing an index pair in N gives a strong index pair.

Theorem 4.16. *Let S be an isolated invariant set where N is an isolating neighborhood for S . The pair $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$ is a strong index pair in N for S .*

Proof. We note that by Theorem 4.15, the pair $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$ is an index pair for S in N . Hence, it is sufficient to show that the index pair is strong. Note that by definition, for all $\sigma \in \text{mo}(S)$, there exists a $\tau \in S$ such that σ is a face of τ . Hence, $\sigma \in F_V(\tau)$. Note that $\text{pf}(\text{mo}(S))$ is precisely the set of simplices σ' for which there exists a path originating in $\text{mo}(S)$ and terminating at σ' , so it is immediate that there is a path originating in S and terminating at σ . Hence, the pair $(\text{pf}(\text{cl}(S)), \text{pf}(\text{mo}(S)))$ is a strong index pair. \square

Our enlarging scheme given in Algorithm 2 does not affect the strongness of an index pair.

Theorem 4.17. *Let R be the output of applying Algorithm 2 to the strong index pair (P, E) in N for S with some parameter δ . The pair $(P, E \setminus R)$ is a strong index pair for S in N .*

Proof. Theorem 4.12 gives that $(P, E \setminus R)$ is an index pair for S in N , so it is sufficient to show that such an index pair is strong. Note that P does not change, but the strongness of index pairs only requires paths to be in P . Since all paths in (P, E) are also paths in $(P, E \setminus R)$, it follows that $(P, E \setminus R)$ is a strong index pair in N . \square

These theorems give us a canonical scheme for choosing invariant sets from a sequence of multivector fields and then computing the barcode of the persistence module given in Equation (4.5). We give our exact scheme in Algorithm 3. Note that we assume that when we let $S_i \leftarrow \text{Inv}_{\mathcal{V}_i}(N_{i-1})$, we are assuming that S_i is \mathcal{V}_i -compatible. If it is not, we can always intersect each multivector in \mathcal{V}_i with N , but this will be handled implicitly by the algorithm.

The astute reader will notice an important detail about Algorithm 3. Namely, the `find` function is parameterized by a nonnegative integer δ , and the function has not yet been defined. In particular, said function must output a closed $N_i \supseteq S_i$ such that S_i is isolated by $N_{i-1} \cup N_i$. An obvious choice is to let $N_i := N_{i-1}$, but such an approach does not allow one to capture essential solutions that “move” outside of $N_{i-1} = N_i$ as the multivector fields change. We give a nontrivial `find` function in the next subsection that can be used to capture such changes in an essential solution.

4.2.2 Finding Isolating Neighborhoods

Given an invariant set S isolated by N with respect to \mathcal{V} , we now propose a method to find a closed, nontrivial $N' \subseteq K$ such that $N \cup N'$ isolates S . Our method relies heavily on the concept of δ -collar introduced in Section 5.1. In fact, we will let $N' = C_\delta(S) \setminus R$ such that $N \cup N'$ isolates S . Hence, it is sufficient to devise an algorithm to find $C_\delta(S) \setminus R$. Before we give and prove the correctness of the algorithm, we briefly introduce the notion of the *push backward* of some set S in N , denoted $\text{pb}_N(S)$. We let $\text{pb}_N(S) = \{x \in N \mid \exists \rho : \text{flow}_\rho(x) \in S\}$.

Input: Sequence of multivector fields $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$, closed set $N_0 \subset K$, $\delta \in \mathbb{Z}$.

Output: Barcodes corresponding to persistence of Conley Index

$i \leftarrow 1$

while $i \leq n$ **do**

$S_i \leftarrow \text{Inv}_{\mathcal{V}_i}(N_{i-1})$

$(P'_{i,1}, E'_{i,1}) \leftarrow (\text{pf}_{N_{i-1}}(\text{cl}(S_i)), \text{pf}_{N_{i-1}}(\text{mo}(S_i)))$

$R_{i,1} \leftarrow \text{findR}(S_i, P'_{i,1}, E'_{i,1}, \mathcal{V}, \delta)$

$(P_i^{(1)}, E_i^{(1)}) \leftarrow (P'_{i,1}, E'_{i,1} \setminus R_{i,1})$

$N_i \leftarrow \text{find}(S_i, N_{i-1}, \mathcal{V}, \delta)$

$(P'_{i,2}, E'_{i,2}) \leftarrow (\text{pf}_{N_i}(\text{cl}(S_i)), \text{pf}_{N_i}(\text{mo}(S_i)))$

$R_{i,2} \leftarrow \text{findR}(S_i, P'_{i,2}, E'_{i,2}, \mathcal{V}, \delta)$

$(P_i^{(2)}, E_i^{(2)}) \leftarrow (P'_{i,2}, E'_{i,2} \setminus R_{i,2})$

if $i = 1$ **then**

$(P_i, E_i) \leftarrow (P_i^{(2)}, E_i^{(2)})$

else if $i = n$ **then**

$(P_i, E_i) \leftarrow (P_i^{(1)}, E_i^{(1)})$

else

$(P_i, E_i) \leftarrow (\text{pf}_{N_{i-1} \cup N_i}(P_i^{(1)} \cup P_i^{(2)}), \text{pf}_{N_{i-1} \cup N_i}(E_i^{(1)} \cup E_i^{(2)}))$

end

$i \leftarrow i + 1$

end

return

$\text{zigzagPers}((P_1, E_1) \supseteq (P_1^{(2)} \cap P_2^{(1)}, E_1^{(2)} \cap E_2^{(1)}) \subseteq (P_2, E_2) \supseteq \dots \subseteq (P_n, E_n))$

Algorithm 3: Scheme for computing the persistence of the Conley Index, variable N

$\mathbb{Z}_{[a,b]} \rightarrow N$, $\rho(a) = x, \rho(b) \in S$. Essentially, the push backward of S in N is the set of simplices $\sigma \in N$ for which there exists a path in N from σ to S .

Input: Invariant set S isolated by N under \mathcal{V} , $\delta \in \mathbb{Z}$
Output: Closed set $N' \supseteq S$ such that $N \cup N'$ isolates S under \mathcal{V}

```

 $V \leftarrow \text{new stack}()$ 
 $R \leftarrow \text{new set}()$ 
 $pb \leftarrow \text{pb}_N(S)$ 
foreach  $\sigma \in C_\delta(S) \cup N$  do
  |  $\text{setUnvisited}(\sigma)$ 
end
foreach  $\sigma \in S$  do
  |  $\text{adj} \leftarrow \text{cl}(\sigma) \cup [v]_{\mathcal{V}}$ 
  | foreach  $\tau \in \text{adj}$  do
  | | if  $\tau \notin S$  and  $\tau \in C_\delta(S) \cup N$  then
  | | |  $\text{push}(V, \tau)$ 
  | | end
  | end
end
while  $\text{size}(V) > 0$  do
  |  $v \leftarrow \text{pop}(V)$ 
  | if  $\text{!hasBeenVisited}(v)$  then
  | |  $\text{setVisited}(v)$ 
  | | if  $(\text{cl}(v) \cup [v]_{\mathcal{V}}) \cap pb \neq \emptyset$  then
  | | |  $\text{add}(R, v)$ 
  | | |  $cf \leftarrow \text{cofaces}(v)$ 
  | | |  $\text{addAll}(R, cf)$ 
  | | else
  | | | foreach  $\sigma \in (\text{cl}(v) \cup [\sigma]_{\mathcal{V}}) \cap (C_\delta(S) \cup N)$  do
  | | | |  $\text{push}(V, \sigma)$ 
  | | | end
  | | end
  | end
end
return  $C_\delta(S) \setminus R$ 

```

Algorithm 4: $\text{find}(S, N, \mathcal{V}, \delta)$

We now prove that $N \cup (C_\delta(S)) \setminus R$ isolates S . Note that since $S \subseteq C_\delta(S) \setminus R$, this also implies that $C_\delta(S) \setminus R$ isolates S .

Theorem 4.18. *Let S denote an invariant set isolated by $N \subseteq K$ under \mathcal{V} . If $C_\delta(S) \setminus R$ is the output of Algorithm 4 on inputs $S, N, \mathcal{V}, \delta$, then the closed set $N \cup (C_\delta(S) \setminus R)$ isolates S .*

Proof. For a contradiction, assume that there exists a path $\rho : \mathbb{Z}_{[a,b]} \rightarrow N \cup (C_\delta(S) \setminus R)$ so that $\rho(a), \rho(b) \in S$ where there is an i satisfying $a < i < b$ with $\rho(i) \notin S$. Note that since N isolates S , if $N \cup C_\delta(S) \setminus R$ does not isolate S , then there must exist a first $k \in \mathbb{Z}_{[a,b]}$ such that $\rho(k) \in C_\delta(S) \setminus N$ and $F_{\mathcal{V}}(\rho(k)) \cap \mathbf{pb}_N(S) \neq \emptyset$. If this were not the case, then N would not isolate S . Without loss of generality, we assume that for all $a < j < k$, we have that $\rho(j) \notin S$. Note that for all $j \in \mathbb{Z}_{[a+1,k-1]}$, when $\rho(j)$ is removed from the stack V , if $\rho(j+1)$ has not been visited, then $\rho(j+1)$ is added to the stack. Hence, this implies that if any $\rho(j)$ is visited, then $\rho(k)$ will be added to R . If this were not the case, there would exist some $\rho(j)$ such that when $\rho(j)$ was removed from the stack, $\rho(j+1)$ was not visited and was not added to the stack. This implies that $F_{\mathcal{V}}(\rho(j)) \cap \mathbf{pb}_N(S) \neq \emptyset$, which contradicts $\rho(k)$ being the first such simplex in the path.

Hence, since $\rho(a+1)$ is added to the stack, it follows that $\rho(k)$ is added to R , which implies that $\rho(\mathbb{Z}_{[a,b]}) \not\subset N \cup (C_\delta(S) \setminus R)$. Note too that $N \cup C_\delta(S) \setminus R$ must be closed, as if there is a $\sigma \in N$ such that $\rho(k) \leq \sigma$, then $\rho(k) \in N$ because N is closed, a contradiction. But when $\rho(k)$ is removed from $C_\delta(S)$, any of its cofaces which are in $C_\delta(S)$ are also removed. Hence, N is closed, $C_\delta(S) \setminus R$ is closed, so their union must be closed. \square

Hence, we use Algorithm 4 as the `find` function in our scheme given in Algorithm 3. We give an example of our implementation of Algorithm 3 using the `find` function in Figure 4.12.




Figure 4.12. Three different index pairs generated from our scheme in Algorithm 3. The isolating neighborhood is in blue, E is in red, and $P \setminus E$ is in yellow. Note how the isolating neighborhood changes by defining a collar around the invariant sets (which are exactly equal to $P \setminus E$). Between the left and middle multivector fields, the periodic attractor partially leaves K , so the maximal invariant set in N is reduced to just a triangle. Hence, the size of N drastically shrinks between the middle and right multivector fields.

5. PERSISTENCE OF CONLEY-MORSE GRAPHS

The previous chapter considered the problem of computing the persistence of the Conley index across a changing sequence of multivector fields. In [33], the authors use zigzag persistence [25] to capture the changing structure of Morse sets in a Morse decomposition. Unfortunately, the information given by only considering Morse sets or only considering the Conley index is incomplete. Given an isolated invariant set S , there may be multiple Morse sets within S , each of which is associated with a Conley index. This information is captured in a *Conley-Morse graph*, which includes a vertex for each Morse set for a given Morse decomposition and the Conley index of the given Morse set. The Conley-Morse graph is a much more precise summary of the behavior of a combinatorial dynamical system than the Morse sets or Conley index alone. In this chapter, we use persistence to capture the changing Conley-Morse graph of a sequence of multivector fields with a specified Morse decomposition. Our method tracks changes in combinatorial dynamical systems at a much finer level of detail than the methods in [33] and the previous chapter. Figure 5.1 is an example of a changing sequence of multivector fields, and Figure 5.3 shows the changing Conley-Morse graphs that correspond to these dynamical systems. Figure 5.2 depicts how the approach in the previous chapter captures the changing behavior of combinatorial dynamical systems. In contrast, Figure 5.5 depicts part of the approach in this chapter, which summarizes the changing Conley index across a select set of sequences of Morse sets, shown in Figure 5.4. By inspecting the barcodes in Figure 5.2 and 5.5, it is easy to see that this chapter presents a much more detailed representation of the changing behavior of combinatorial dynamical systems than the approach in the previous chapter.

This chapter is based on [19], and it quotes passages of [19] verbatim.

5.1 Conley-Morse Filtrations

We now move to develop a method to compute a set of filtrations for a sequence of Conley-Morse graphs. These filtrations need to capture key features of the sequence: how the structure of the Conley-Morse graphs changes throughout the sequence, and how the Conley index changes at individual vertices; see Figures 5.3, 5.4, and 5.5. Formally, we assume





Figure 5.1. Three multivector fields. On the left, a yellow repelling fixed point is surrounded by a blue periodic attractor. In the middle, the repelling fixed point has split into a yellow periodic repeller and a blue attracting fixed point. On the right, the periodic repeller has collided with the periodic attractor to form a red semistable limit cycle (a semistable limit cycle is a limit cycle that is stable on one side and unstable on the other). In this example, all three multivector fields are significantly different from one another. In computing persistence, we will generally assume that there are several intermediate multivector fields, representing a gradual transition between the multivector fields shown here.



Dimension: 0

Figure 5.2. This figure illustrates the approach from [18], applied to capturing the changing structure of the multivector fields in Figure 5.1. The approach in [18] requires selecting a single isolated invariant set for each multivector field, with the canonical choice being the maximal isolated invariant set. In each multivector field, the maximal isolated invariant set is an attractor, highlighted in blue. Despite each multivector field giving rise to different dynamical systems, the maximal isolated invariant set is the same in each multivector field. Hence, computing the persistence using techniques from [18] gives a single, 0-dimensional bar, which we depict at the bottom in white.




Figure 5.3. The Conley-Morse graph (Definition 3.12) for the Morse decompositions (Definition 3.10) in Figure 5.1, where the top vertices represent the fixed points. The colors of the vertices match the colors of the Morse sets with which they are drawn in Figure 5.1. Each label captures information about the Conley indices of the Morse sets.




Figure 5.4. We extract a set of zigzag filtrations to capture the changing Conley indices in a sequence of Conley-Morse graphs. In this case, we extract three particular zigzag filtrations, which correspond to the sequences of Morse sets boxed in rectangles.

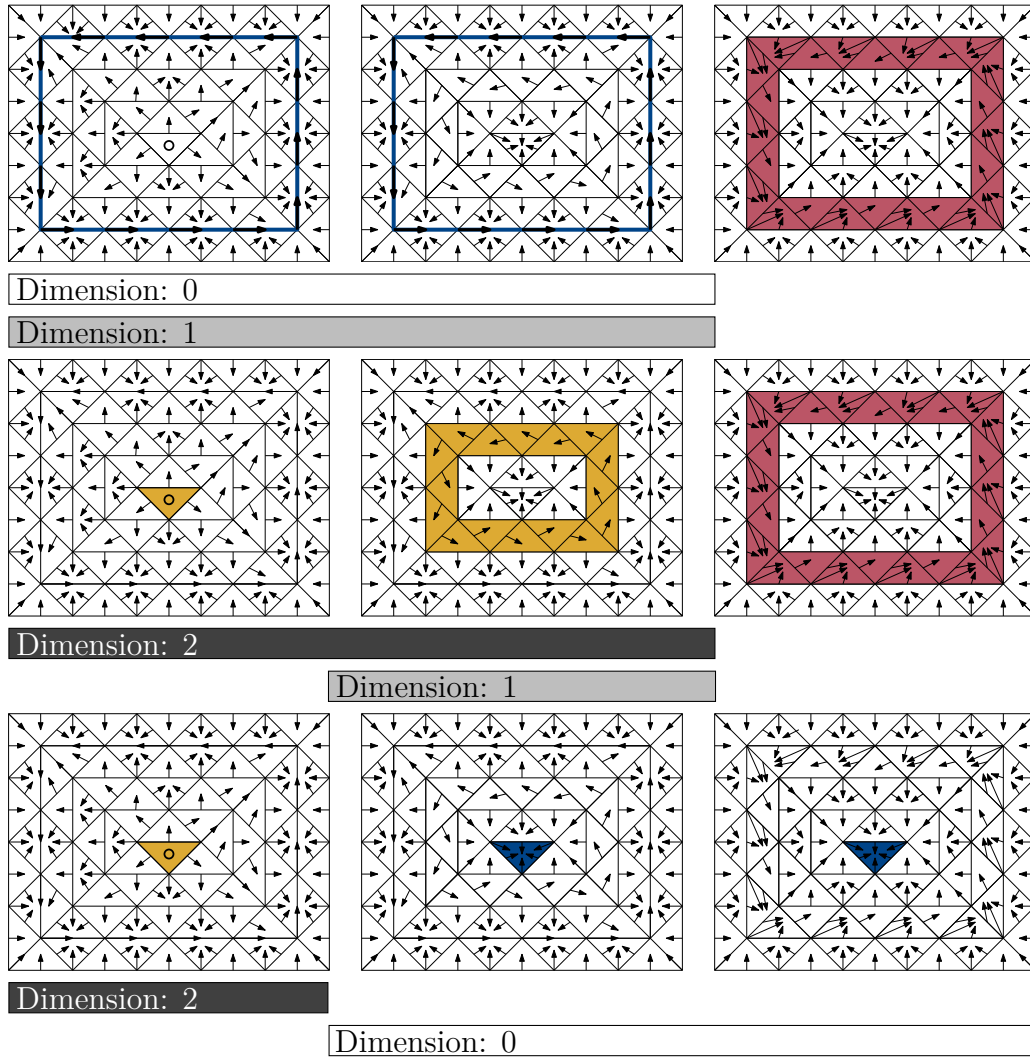


Figure 5.5. Illustrating our approach for computing the changing structure of the multivector fields in Figure 5.1. By extracting a specific set of zigzag filtrations (Section 5.1), we can compute a set of barcodes which represent the changing Conley indices of the Morse sets (Section 5.3). Each row represents an extracted filtration, and the barcode for the filtration is depicted below it. White bars are 0-dimensional, light gray bars are 1-dimensional, and dark gray bars are 2-dimensional. In addition, we extract barcodes representing the changing structure of the Conley-Morse graph (not pictured). We collate all of these bars into a single barcode by removing redundancy in Section 5.3.

that we are given a sequence of Conley-Morse graphs $\{G_i\}_{i=1}^n$. Each Conley-Morse graph corresponds to a Morse decomposition $\{\mathcal{M}_i\}_{i=1}^n$ of an isolated invariant set $\{S_i\}_{i=1}^n$ under a multivector field $\{\mathcal{V}_i\}_{i=1}^n$. We assume that each S_i is in some fixed isolating neighborhood N . We also assume that for each $M_j \in \mathcal{M}_i$, we have an associated index pair (P_j, E_j) . This may seem particularly daunting, in that our approach requires a specified isolated invariant set and Morse decomposition for every multivector field. If a practitioner does not have a particular isolated invariant set or Morse decomposition in mind, a canonical choice is always to take the maximal invariant set S_i in N under \mathcal{V}_i and to take the minimal Morse decomposition of S_i . The minimal Morse decomposition can be easily computed by converting a multivector field into a directed graph and computing maximal strongly connected components. For more details, see [14]. In addition, we require that each Morse set M is associated with a particular index pair (P, E) in N . A canonical choice of an index pair is to use Theorem 4.9 and let $P := \text{pf}(\text{cl}(M))$ and $E := \text{pf}(\text{mo}(M))$, but this is in general not resilient to noise.

In this section, we extract Conley index zigzag filtrations from a sequence of Conley-Morse graphs G_1, G_2, \dots, G_m , corresponding to Morse decompositions $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_m$. That is, we extract filtrations which represent the changing Conley index at the vertices of the Conley-Morse graph. We call these *Conley-Morse filtrations*, and we postpone discussion of *graph filtrations*, which capture the changing structure of the Conley-Morse graph itself, to the next section. A first approach to extracting Conley-Morse filtrations is to consider all possible zigzag filtrations via vertex sequences (three such sequences are shown in Figure 5.4). Each vertex corresponds to an index pair, so for each sequence of vertices taken from consecutive Conley-Morse graphs, we can use Theorem 4.4 to get a relative zigzag filtration. If $|G_i| = n$ for all i , then there are $\Theta(n^m)$ possible zigzag filtrations. This is clearly intractable for mid-to-large values of m and n . We aim to reduce the number of filtrations while still capturing the changing Conley index. Our key observation is that if (P_1, E_1) and (P_2, E_2) are index pairs, then clearly a homology class persists from (P_1, E_1) to (P_2, E_2) only if $(P_1 \setminus E_1) \cap (P_2 \setminus E_2) \neq \emptyset$, because an empty intersection leaves no room for the class to be present in both spaces under inclusions.

Consider a sequence of index pairs in N , denoted $\{(P_i, E_i)\}_{i=a}^b$, where each (P_i, E_i) is an index pair for a Morse set $M_i \in \mathcal{M}_i$. Such a sequence is *feasible* if for all $a \leq i < b$, $(P_i \setminus E_i) \cap (P_{i+1} \setminus E_{i+1}) \neq \emptyset$. A feasible sequence of index pairs is a *maximal* sequence if there does not exist a (P_{a-1}, E_{a-1}) such that $(P_{a-1} \setminus E_{a-1}) \cap (P_a \setminus E_a) \neq \emptyset$, and there does not exist a (P_{b+1}, E_{b+1}) such that $(P_b \setminus E_b) \cap (P_{b+1} \setminus E_{b+1}) \neq \emptyset$. Hence, each feasible sequence is contained in a maximal sequence. Each maximal sequence gives a relative zigzag filtration by intersecting consecutive index pairs (recall Theorem 4.4). We call a relative zigzag filtration that is given by a maximal sequence a *Conley-Morse filtration*, and we use the set of Conley-Morse filtrations to compute the changing Conley indices at the vertices of a sequence of Conley-Morse graphs.

Algorithm 5 is our formal approach for computing the set of Conley-Morse filtrations. For completeness, we prove that it successfully computes the set of Conley-Morse filtrations.

Proposition 5.1. *Given a sequence of Conley-Morse graphs $\{G_i\}_{i=1}^n$ corresponding to a sequence of Morse decompositions $\{\mathcal{M}_i\}_{i=1}^n$ of isolated invariant sets in N , where each $M_i \in \mathcal{M}_i$ is associated with a unique index pair (P, E) , Algorithm 5 outputs all Conley-Morse filtrations.*

Proof. Note that Algorithm 5 attempts to find all Conley-Morse filtrations by first finding all maximal sequences. Hence, it is sufficient to show that it correctly finds all maximal sequences, because it is trivial to convert these into Conley-Morse filtrations. Note that Algorithm 5 constructs sequences by incrementally adding index pairs to already-started sequences. By inspection, we note that the algorithm only constructs a new sequence with the index pair (P, E) from a Morse set in the Conley-Morse graph G_i if there does not exist an index pair (P', E') for a Morse set in the Conley-Morse graph G_{i-1} where $(P' \setminus E') \cap (P \setminus E) \neq \emptyset$. Similarly, the algorithm only ceases to consider a sequence ending with the index pair (P, E) for a Morse set in the Conley-Morse graph G_i if there does not exist an index pair (P', E') for a Morse set in the Conley-Morse graph G_{i+1} where $(P \setminus E) \cap (P' \setminus E') \neq \emptyset$. In addition, note that the algorithm only appends the index pair (P, E) to a sequence ending in (P', E') if $(P \setminus E) \cap (P' \setminus E') \neq \emptyset$. Hence, each sequence that is constructed by the algorithm is a maximal sequence.

Input: Sequence of Conley-Morse graphs G_i corresponding to Morse decompositions \mathcal{M}_i for isolated invariant sets in N . Each $M_i \in \mathcal{M}_i$ is associated with a unique index pair (P, E) .

Output: Set of all Conley-Morse filtrations for $\{G_i\}_{i=1}^n$

```

alive_seqs ← new set()
all_seqs ← new set()
for i ∈ {1, ..., n} do
    to_remove ← new set()
    still_alive ← new set()
    for seq ∈ alive_seqs do
        (P', E') ← LastIndexPair(seq)
        for M ∈  $\mathcal{M}_i$  do
            (P, E) ← IndexPair(M)
            if  $(P \setminus E) \cap (P' \setminus E') \neq \emptyset$  then
                new_seq ← copy(seq)
                append(new_seq, (P, E))
                add(still_alive, new_seq)
                IsInSequence(M) ← True
                add(to_remove, seq)
            end
        end
    end
    dead_seqs ← alive_seqs \ to_remove
    alive_seqs ← still_alive
    all_seqs ← all_seqs ∪ dead_seqs
    for M ∈  $\mathcal{M}_i$  do
        if IsInSequence(M) = False then
            seq ← new Sequence()
            append(seq, IndexPair(M))
            add(alive_seqs, seq)
        end
    end
end
all_seqs ← all_seqs ∪ alive_seqs
filtrations ← ConvertToFiltrations(all_seqs)
return filtrations

```

Algorithm 5: FindConleyMorseFiltrations($\{G_i\}_{i=1}^n, \{\mathcal{M}_i\}_{i=1}^n$)

It remains to be shown that the algorithm constructs all possible maximal sequences. Assume there existed some maximal sequence $\{(P_i, E_i)\}_{i=a}^b$ that was not included in *all_seqs* before they are converted to filtrations. In such a case, there must exist some (P_k, E_k) which was not appended to the sequence $\{(P_i, E_i)\}_{i=a}^{k-1}$. But by inspection, we see that (P_k, E_k) is appended to all such sequences so long as $(P_k \setminus E_k) \cap (P_{k-1} \setminus E_{k-1}) \neq \emptyset$. Hence, this cannot be the case, which implies that $\{(P_i, E_i)\}_{i=a}^b$ is included in *all_seqs*. \square

The set of Conley-Morse filtrations and the associated barcodes for the Morse decompositions in Figure 5.1 are shown in Figure 5.5.

5.2 Graph Filtrations

Now, we show how to find a zigzag filtration that corresponds to the changing structure of the Conley-Morse graph called a *graph filtration*. We begin by reviewing some properties of multivector fields. A multivector field \mathcal{V}_2 is *inscribed* in \mathcal{V}_1 , denoted $\mathcal{V}_2 \sqsubseteq \mathcal{V}_1$, if for each multivector $V_1 \in \mathcal{V}_1$, there exists a multivector $V_2 \in \mathcal{V}_2$ where $V_1 \subseteq V_2$.

Proposition 5.2 (See e.g. [13], [33]). *Let \mathcal{V}_1 and \mathcal{V}_2 denote multivector fields over K , and let $\mathcal{V}_2 \sqsubseteq \mathcal{V}_1$. Then for all $\sigma \in K$, $F_{\mathcal{V}_2}(\sigma) \subseteq F_{\mathcal{V}_1}(\sigma)$.*

Proposition 5.3 follows directly from Proposition 5.2.

Proposition 5.3. *Let \mathcal{V}_1 and \mathcal{V}_2 denote multivector fields over K , and let $\mathcal{V}_2 \sqsubseteq \mathcal{V}_1$. If there exists a path ρ in K under \mathcal{V}_2 from σ_0 to σ_n , then ρ is a path under \mathcal{V}_1 .*

From these two results, we deduce an important property of Morse sets.

Proposition 5.4. *Let M_1 denote an isolated invariant set in N under \mathcal{V}_1 and M_2 denote a minimal Morse set in N under \mathcal{V}_2 , where $\mathcal{V}_2 \sqsubseteq \mathcal{V}_1$. If $M_1 \cap M_2 \neq \emptyset$, then $M_2 \subseteq M_1$.*

Proof. Consider $\sigma \in M_1 \cap M_2$, and $\tau \in M_2$. By Proposition 3.11, there exists a path p in N under \mathcal{V}_2 from σ to τ and a path q in N under \mathcal{V}_2 from τ to σ . By concatenating the paths, we get a new path r in N under \mathcal{V}_2 that starts and ends at $\sigma \in M_1$. By Proposition 5.3, r is a path under \mathcal{V}_2 . But, M_1 is isolated by N , so $\text{im}(r) \subseteq M_1$. Hence, $\tau \in M_1$, so $M_2 \subseteq M_1$. \square

Proposition 5.4 is the starting point for constructing a zigzag filtration. For tractability, we focus on two Conley-Morse graphs G_1 and G_2 , which correspond to Morse decompositions \mathcal{M}_1 and \mathcal{M}_2 for the isolated invariant sets S_1 and S_2 in N under \mathcal{V}_1 and \mathcal{V}_2 . It is natural to consider the Conley-Morse graph corresponding to the minimal Morse decomposition of $\text{Inv}(N)$ under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. We call this graph $G_{1,2}$. There is a partial function $\iota_1 : V(G_{1,2}) \rightarrow V(G_1)$, where $\iota_1(M_{1,2}) = M_1$ if $M_{1,2} \subseteq M_1$ (here, we use M_1 and $M_{1,2}$ to denote both a Morse set and its corresponding vertex in the Conley-Morse graph). There is a similar partial function $\iota_2 : V(G_{1,2}) \rightarrow V(G_2)$. Note that ι_1 and ι_2 are not edge preserving. If there exists a directed edge from vertex $M_{1,2} \in G_{1,2}$ to vertex $M'_{1,2} \in G_{1,2}$, then there is a connection from $M_{1,2}$ to $M'_{1,2}$. By Proposition 5.3, there must exist a path from $\iota_1(M_{1,2})$ to $\iota_1(M'_{1,2})$ (likewise for $\iota_2(M_{1,2})$ to $\iota_2(M'_{1,2})$). But this path may intersect some other invariant set $M_1 \in \mathcal{M}_1$, so there could exist an edge from $M_{1,2}$ to $M'_{1,2}$, but no edge from $\iota_1(M_{1,2})$ to $\iota_1(M'_{1,2})$. This phenomenon occurs in Figure 5.6. Hence, we need a slightly modified notion of a Conley-Morse graph under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$.




Figure 5.6. If the multivector field on the left is \mathcal{V}_1 and the one on the right is \mathcal{V}_2 , then the multivector field in the middle is $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. The maximal invariant set is the entire rectangle in all three multivector fields, and the Morse sets in the minimal Morse decomposition for this isolated invariant set are colored. In the middle, the Morse sets in gray, green, magenta, and pink are spurious because they are not contained in a Morse set on the right. Also in the middle, there is a connection from the golden triangle to the magenta vertex. But on the left, the connection passes through the turquoise triangle. Hence, the golden triangle-to-magenta vertex connection is not relevant.

Definition 5.5 (Relevant and Spurious). *Let $\mathcal{M}_1, \mathcal{M}_2$ denote Morse decompositions of isolated invariant sets S_1, S_2 in N under $\mathcal{V}_1, \mathcal{V}_2$, and let $\mathcal{M}_{1,2}$ denote the minimal Morse decomposition of $S_{1,2} = \text{Inv}(N)$ under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. A Morse set $M_{1,2} \in \mathcal{M}_{1,2}$ is relevant if there exists an $M_1 \in \mathcal{M}_1$ and an $M_2 \in \mathcal{M}_2$ where $M_{1,2} \subseteq M_1 \cap M_2$. If not, then $M_{1,2}$ is spurious.*

Definition 5.6 (Relevant Connection). *Let $\mathcal{M}_1, \mathcal{M}_2$ denote Morse decompositions of isolated invariant sets S_1, S_2 in N under $\mathcal{V}_1, \mathcal{V}_2$, and let $\mathcal{M}_{1,2}$ denote the minimal Morse decomposition of $\text{Inv}(N)$ under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. Also, let $\rho : \mathbb{Z} \cap [a, b] \rightarrow N$ denote a connection from $M_{1,2} \in \mathcal{M}_{1,2}$ to $M'_{1,2} \in \mathcal{M}_{1,2}$ under $\mathcal{V}_{1,2}$ where $M_{1,2}$ and $M'_{1,2}$ are relevant Morse sets. If ρ satisfies both of the following:*

1. *If $i \in \mathbb{Z} \cap [a, b]$ satisfies $\rho(i) \in M_1 \in \mathcal{M}_1$, then $M_1 = \iota_1(M_{1,2})$ or $M_1 = \iota_1(M'_{1,2})$.*
2. *If $i \in \mathbb{Z} \cap [a, b]$ satisfies $\rho(i) \in M_2 \in \mathcal{M}_2$, then $M_2 = \iota_2(M_{1,2})$ or $M_2 = \iota_2(M'_{1,2})$.*

then ρ is a relevant connection.

We use the notions of relevant Morse sets to define the *relevant Conley-Morse graph*.

Definition 5.7 (Relevant Conley-Morse Graph). *Let $\mathcal{M}_1, \mathcal{M}_2$ denote Morse decompositions of isolated invariant sets S_1, S_2 in N under $\mathcal{V}_1, \mathcal{V}_2$, and let $\mathcal{M}_{1,2}$ denote the minimal Morse decomposition of $S_{1,2} = \text{Inv}(N)$ under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. The relevant Conley-Morse graph is the graph $G_{1,2}$ given by including a vertex in $V(G_{1,2})$ for each relevant Morse set in $\mathcal{M}_{1,2}$, and including a directed edge from the vertex corresponding to $M_{1,2}$ to the vertex corresponding to $M'_{1,2}$ if there is a relevant connection from $M_{1,2}$ to $M'_{1,2}$.*

An example of these concepts can be seen in Figure 5.7. The top three graphs in Figure 5.7 are the Conley-Morse graphs, omitting the Pointcaré polynomials, for the minimal Morse decompositions in Figure 5.6. If the top left and top right Conley-Morse graphs are Conley-Morse graphs under \mathcal{V}_1 and \mathcal{V}_2 , then the top center Conley-Morse graph is a Conley-Morse graph under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. The bottom left graph and bottom right graphs are the same as the top left and top right graphs, but the bottom center graph is the relevant Conley-Morse graph that is extracted from the top center graph. Each colored vertex represents the Morse set of the same color in Figure 5.6. Red arrows between the top center and top left or top right graphs indicate that a Morse set represented by a vertex in the top center graph is contained in a Morse set represented by a vertex in the top left or top right graphs. Hence, a vertex in the top center Conley-Morse graph is only relevant if there is a red arrow from it to a vertex in the top left and top right graphs. Relevant connections in the top center

graph are shown in blue. There is a connection from the golden triangle to the blue periodic attractor by heading directly south in all three multivector fields in Figure 5.6. Hence, this is a relevant connection, and the edge from the golden vertex to the blue vertex is included in the relevant Conley-Morse graph in the bottom center. There is a path from the golden triangle to the magenta vertex in both the left and the center vector fields in Figure 5.6, but in the left multivector field, the path first passes through the turquoise Morse set. So while the path exists in both the left and middle multivector fields, it is not direct in the left multivector field, so it is not a relevant connection. Similarly, while a connection from the gray Morse set to the green Morse set exists in both the left and center multivector fields, neither the gray nor the green Morse sets are contained in a Morse set in the right multivector field, so this is not a relevant connection. Green edges represent paths that are in both the top center and top left graphs, but there is not a corresponding path in the top right graph.

For the remainder of the paper, whenever we refer to a graph $G_{i,i+1}$ for some i , we are referring to the relevant Conley-Morse graph under $\mathcal{V}_i \bar{\cap} \mathcal{V}_{i+1}$. Note that we compute relevant Conley-Morse graphs by taking subgraphs of a Conley-Morse graph. Hence, ι_1 and ι_2 restrict to functions $\iota_1 : G_{1,2} \rightarrow G_1$ and $\iota_2 : G_{1,2} \rightarrow G_2$ (they are no longer partial).

Proposition 5.8. *Let G_1 and G_2 denote the Conley-Morse graphs for Morse decompositions of the isolated invariant sets S_1, S_2 in N under $\mathcal{V}_1, \mathcal{V}_2$, and let $G_{1,2}$ denote the relevant Conley-Morse graph for the minimal Morse decomposition of the maximal invariant set in N under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. If there is a directed edge from $M_{1,2}$ to $M'_{1,2}$ in $G_{1,2}$, then either $\iota_k(M_{1,2}) = \iota_k(M'_{1,2})$ or there exists a directed edge from $\iota_k(M_{1,2})$ to $\iota_k(M'_{1,2})$ for $k \in \{1, 2\}$.*

Proof. If there exists a directed edge from $M_{1,2}$ to $M'_{1,2}$, then there exists a relevant connection from $M_{1,2}$ to $M'_{1,2}$ under $\mathcal{V}_1 \bar{\cap} \mathcal{V}_2$. We denote this connections as $\rho : [0, n] \cap \mathbb{Z} \rightarrow N$. By Proposition 5.3, ρ must also be a path under \mathcal{V}_1 and \mathcal{V}_2 . By the definition of relevant connection, if $\rho(i) \in M \in \mathcal{M}_1$, then $\iota_1(M_{1,2}) = M$ or $\iota_1(M'_{1,2}) = M$. Hence, either $\iota_1(M_{1,2}) = \iota_1(M'_{1,2})$, or as ρ is a direct connection from $\iota_1(M_{1,2})$ to $\iota_1(M'_{1,2})$, there exists a directed edge from the former to the later. The same argument holds for ι_2 . \square




Figure 5.7. On the top row, we depict the Conley-Morse graphs (absent the Pointcaré polynomials) for the Morse decompositions in Figure 5.6. The bottom row depicts the Conley-Morse graphs for the left and right Morse decompositions in Figure 5.6, but it instead includes the relevant Conley-Morse graph in the center.

Now, we have two Conley-Morse graphs G_1, G_2 for isolated invariant sets in N under $\mathcal{V}_1, \mathcal{V}_2$ and the relevant Conley-Morse graph $G_{1,2}$ for the maximal invariant set in N under $\mathcal{V}_1 \cap \mathcal{V}_2$. We are interested in how the structure of these graphs “persist.” To do this, we will treat the graphs as simplicial complexes by ignoring the orientation on edges. The maps ι_1, ι_2 induce simplicial maps on these complexes.

Proposition 5.9. *Let $f_1 : G_{1,2} \rightarrow G_1$ and $f_2 : G_{1,2} \rightarrow G_2$ denote the maps induced by ι_1, ι_2 where $f_1(\{u, v\}) = \{\iota_1(u), \iota_1(v)\}$ and $f_2(\{u, v\}) = \{\iota_2(u), \iota_2(v)\}$. The maps f_1 and f_2 are simplicial maps.*

Proof. The maps f_1 and f_2 bring vertices to vertices, so it is sufficient to show that the image of an edge is either an edge or a vertex. If $\sigma \in G_{1,2}$ is an edge, then it corresponds to an edge $e = (u, v) \in G_{1,2}$. Such edges correspond to relevant paths from the vertex u to the vertex v . Thus, either $\iota_1(u) = \iota_1(v)$, which implies that $f_1(\{u, v\}) = \iota_1(u)$, or $\iota_1(u) \neq \iota_1(v)$, which implies that there is a connection from $\iota_1(u)$ to $\iota_1(v)$. Hence, there exists an edge $\{\iota_1(u), \iota_1(v)\} \in G_1$. Thus, f_1 is both a map and a simplicial map. The argument for f_2 follows analogously. \square

Given a sequence of n Conley-Morse graphs and $n - 1$ relevant Conley-Morse graphs, we can use Proposition 5.9 to obtain a sequence of complexes connected by simplicial maps. We show this sequence in Equation 5.1:

$$G_1 \leftarrow G_{1,2} \rightarrow G_2 \leftarrow G_{2,3} \rightarrow \cdots \leftarrow G_{n-1,n} \rightarrow G_n. \quad (5.1)$$

Hence, we have a zigzag filtration that captures the changing structure of a sequence of Conley-Morse graphs. We can compute the barcode for zigzag filtrations where the maps are simplicial by using an algorithm in [34].

5.3 Barcodes for Conley-Morse Graphs

In Section 5.1 we showed how to find Conley-Morse filtrations, which represent the changing Conley indices at the vertices of a Conley-Morse graph. Similarly, in Section 5.2, we showed how to extract a graph filtration, which represents the changing graph structure

of the Conley-Morse graph. By taking all barcodes for these filtrations together, we can straightforwardly obtain a “barcode” for a sequence of Conley-Morse graphs. However, the barcode obtained from Conley-Morse filtrations may contain a significant amount of redundancy. Recall that we compute the changing Conley index by considering all possible Conley-Morse filtrations given by maximal sequences $\{(P_i, E_i)\}_{i=j}^k$ where (P_i, E_i) is an index pair in N for a Morse set in the i th Conley-Morse graph and $(P_i \setminus E_i) \cap (P_{i+1} \setminus E_{i+1}) \neq \emptyset$. As a result of this construction, a subfiltration given by the sequence $\{(P_i, E_i)\}_{i=j}^k$ may occur in several filtrations.

Duplication of subfiltrations can lead to a duplication of bars. In Figure 5.8, consider the middle and bottom filtrations. The middle filtration shows a repelling fixed point turning into a periodic repeller which in turn merges with a periodic attractor to become a semistable limit cycle. In the bottom filtration, a repelling fixed point becomes an attracting fixed point. The 2-dimensional barcode for the middle filtration is given by a single bar, which represents the 2-dimensional homology generator for the repelling fixed point and the periodic repeller. In contrast, the 2-dimensional bar code for the bottom filtration represents only the lifetime of the repelling fixed point. This is redundant: the 2-dimensional bar for the middle filtration already captures the homology generator for the repelling fixed point in the bottom filtration. One bar is a subset of the other, and they capture the same generator.

In devising a barcode that captures a changing Conley-Morse graph, it is best that we eliminate these redundant bars. In this section, we use the notation $\mathcal{I}_{a,b}$ to denote an interval I such that $\text{cl}(I) = [a, b]$. That is, $\mathcal{I}_{a,b}$ can denote any of $[a, b]$, $[a, b)$, $(a, b]$, or (a, b) . We also consider *subfiltrations*. If \mathcal{F} denotes the zigzag filtration $(P_1, E_1) \supseteq (P_1 \cap P_2, E_1 \cap E_2) \subseteq (P_2, E_2) \supseteq \dots$, then we use the notation $\mathcal{F}_{a,b}$ to denote the subfiltration $(P_a, E_a) \supseteq (P_a \cap P_{a+1}, E_a \cap E_{a+1}) \subseteq (P_{a+1}, E_{a+1}) \supseteq \dots \subseteq (P_b, E_b)$.

Definition 5.10 (Redundant). *Let \mathcal{Z} denote the set of all maximal relative zigzag filtrations for a sequence of Conley-Morse graphs, and let $\mathcal{I}_{a,b}$ denote a k -dimensional bar extracted from $\mathcal{F} \in \mathcal{Z}$. If there exists a filtration $\mathcal{G} \in \mathcal{Z}$ where $\mathcal{G}_{a,b} = \mathcal{F}_{a,b}$, then the bar $\mathcal{I}_{a,b}$ is redundant.*

Proposition 5.11. *Let \mathcal{Z} denote the set of maximal relative zigzag filtrations for a sequence of Conley-Morse graphs. If $\mathcal{I}_{a,b}$ is a redundant k -dimensional bar in the barcode for $\mathcal{F} \in \mathcal{Z}$,*

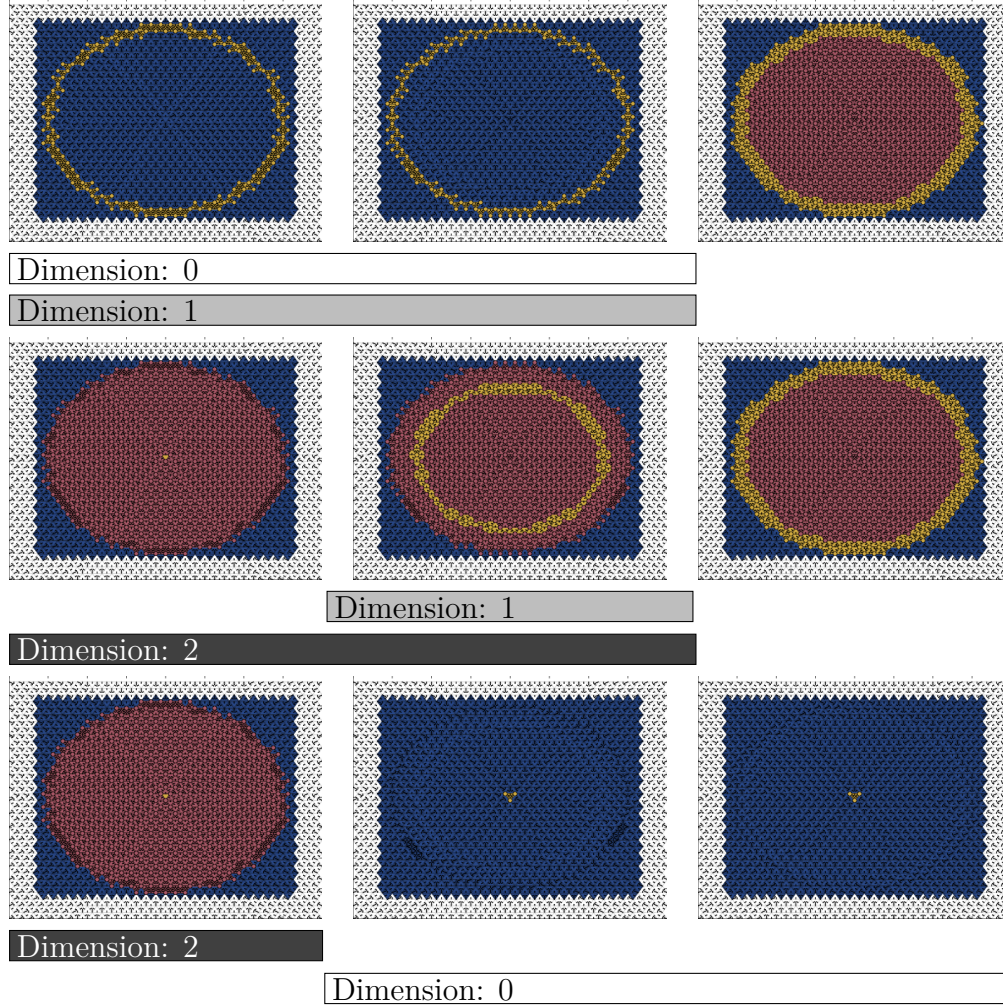


Figure 5.8. All three maximal sequences extracted from the changing Conley-Morse graph for the Morse decompositions in Figure 5.1. The isolating neighborhood is given by yellow, red, and blue simplices, while if (P, E) is an index pair, the simplices in $P \setminus E$ are in yellow and the simplices in E are in red. The top three images show a periodic attractor that becomes a semistable limit cycle. The middle three images show a repelling fixed point that becomes a periodic repeller and then becomes a semistable limit cycle, and the bottom three images show a repelling fixed point that transitions into an attracting fixed point. Beneath each maximal sequence, we include the barcode from the zigzag filtration that we get by applying Theorem 4.4 to the maximal sequence to obtain a Conley-Morse filtration. White bars are 0-dimensional, light gray bars are 1-dimensional, and dark gray bars are 2-dimensional.




Figure 5.9. The barcode corresponding to the sequence of combinatorial dynamical systems in Figure 5.1. The white bar above the dotted line is 0-dimensional, and it represents the connected component in the Conley-Morse graph. The bars below the dotted line are obtained by extracting the barcodes from the Conley-Morse filtrations in Figure 5.8 and removing redundant bars.

then there exists a filtration $\mathcal{G} \in \mathcal{Z}$ such that the k -dimensional barcode for \mathcal{G} contains a bar $\mathcal{I}_{c,d}$ where $\mathcal{I}_{a,b} \subseteq \mathcal{I}_{c,d}$.

Proof. By restricting the k -dimensional barcode for \mathcal{G} and \mathcal{F} to the interval $[a, b]$, we get two sets of bars for the zigzag filtration $\mathcal{F}_{a,b} = \mathcal{G}_{a,b}$. The barcode for $\mathcal{F}_{a,b} = \mathcal{G}_{a,b}$ is unique (see [25]). Hence, there must be a bar $\mathcal{I}_{c,d}$ in the barcode for \mathcal{G} where $\mathcal{I}_{c,d} \cap [a, b] = \mathcal{I}_{a,b}$. \square

If \mathcal{Z} is the set of zigzag filtrations for the Conley-Morse graph, and \mathcal{B} denotes the set of bars extracted from filtrations in \mathcal{Z} , then redundancy gives a partial order on \mathcal{B} . In particular, $\mathcal{I}_{a,b} \leq \mathcal{I}_{c,d}$ if $\mathcal{I}_{a,b} \subseteq \mathcal{I}_{c,d}$. Together with the bars extracted from the graph structure filtration in Section 5.2, the set of maximal bars under \leq are taken to be the barcode for the changing Conley-Morse graph.

Input: List of filtrations $\mathcal{F}_1, \dots, \mathcal{F}_n$, and the corresponding set of k -dimensional barcodes $\mathcal{B}(\mathcal{F}_1), \dots, \mathcal{B}(\mathcal{F}_n)$

Output: Set of bars \mathcal{B}

```

 $\mathcal{B} \leftarrow \text{new set}()$ 
for  $i \in 1, \dots, n$  do
     $\mathcal{F} \leftarrow \mathcal{F}_i$ 
    for  $\mathcal{I}_{a,b} \in \mathcal{B}(\mathcal{F}_i)$  do
         $\text{redundant} \leftarrow \text{False}$ 
        for  $j \in 1 \dots n$  do
            if  $j \neq i$  then
                 $\mathcal{G} \leftarrow \mathcal{G}_j$ 
                if  $\text{!forbidden}(\mathcal{G}_{a,b})$  and  $\mathcal{F}_{a,b} = \mathcal{G}_{a,b}$  then
                     $\text{redundant} \leftarrow \text{True}$ 
                     $\text{forbidden}(\mathcal{F}_{a,b}) \leftarrow \text{True}$ 
                end
            end
        end
        if  $\text{redundant} = \text{False}$  then
             $\text{add}(\mathcal{B}, \mathcal{I}_{a,b})$ 
        else
             $\mathcal{B}(\mathcal{F}_i) \leftarrow \mathcal{B}(\mathcal{F}_i) \setminus \{\mathcal{I}_{a,b}\}$ 
        end
    end
end
return  $\mathcal{B}$ 

```

Algorithm 6: $\text{EliminateRedundancies}(\{\mathcal{F}_i\}_{i=1}^n, \{\mathcal{B}(\mathcal{F}_i)\}_{i=1}^n)$

Proposition 5.12. *Algorithm 6 outputs the set of maximal bars.*

Proof. Consider the set of maximal bars, denoted \mathcal{B}_M . We consider two cases. In the first, the maximal bar $B_{a,b}$ corresponds to a zigzag subfiltration in exactly one zigzag filtration \mathcal{F} . In such a case, there is no other filtration \mathcal{G} where $\mathcal{F}_{a,b} = \mathcal{G}_{a,b}$, so when considering $B_{a,b}$ the redundant flag will not be **True** and hence $B_{a,b}$ will be added to \mathcal{B} . Now, assume that the maximal bar $B_{a,b}$ corresponds to the same subfiltration in $\mathcal{F}_{i_1}, \dots, \mathcal{F}_{i_k}$, where $i_1 \leq \dots \leq i_k$. For each filtration \mathcal{F}_{i_j} , $i_j < i_k$, we claim that $B_{a,b}$ will be considered redundant. Clearly this is the case, because $B_{a,b}$ corresponds to a subfiltration $\mathcal{F}_{i_k,a,b}$, and as \mathcal{F}_{i_k} has not been processed at the time that \mathcal{F}_{i_j} is processed for $i_j < i_k$, it follows that when each is processed and $B_{a,b}$ is considered, the redundant flag will be set to **True** and $B_{a,b}$ will not be added to \mathcal{B} . However, when \mathcal{F}_{i_k} is considered, all of the previous $\mathcal{F}_{i_j,a,b}$ will have been marked forbidden, and $B_{a,b}$ will be added to \mathcal{B} . Thus, $B_{a,b}$ will be added exactly once.

Hence, all maximal bars are included in \mathcal{B} . Note that every non-maximal bar $B_{a,b}$ will be marked redundant, because if $B_{a,b}$ is non-maximal then it corresponds to some filtration \mathcal{F} and there must exist some maximal bar $B'_{c,d}$ corresponding to the filtration \mathcal{G} where $\mathcal{G}_{a,b} = \mathcal{F}_{a,b}$. Since $B'_{c,d}$ is maximal, we have already established that at least one copy of it will not be marked redundant, so $\mathcal{F}_{a,b}$ can be compared against $\mathcal{G}_{a,b}$. Hence, $B_{a,b}$ will be marked redundant. \square

6. TRACKING ISOLATED INVARIANT SETS WITH CONTINUATION

In the previous two chapters, we have seen how one can use persistence to summarize how salient features of combinatorial dynamical systems change across multivector fields. We have also introduced an elementary tracking algorithm that automatically selects a new isolated invariant set from a seed invariant set. However, this algorithm is quite primitive. The selected isolated invariant set is only guaranteed to share an isolating neighborhood with the initial isolated invariant set. The Conley index could vary wildly. In this section, we introduce a more sophisticated tracking protocol to capture changes to a seed isolated invariant set.

To do this, we place the notion of the *continuation* of an isolated invariant set in the combinatorial setting. In the classical setting, Conley showed that the Conley index of an isolated invariant set can not change under arbitrarily small perturbations [30]. Intuitively, this means that an attractor cannot change into a repeller under an arbitrarily small perturbation, and vice-versa. This permits one to track an isolated invariant set through the space of dynamical systems. Ultimately, the Conley index can change, which is typically at a bifurcation point. Once the bifurcation has occurred, it may be possible to continue tracking an isolated invariant set, possibly with a different Conley index than before the bifurcation point. In this chapter, we introduce the notion of small changes to a multivector field, and we show how to track an isolated invariant set. When a bifurcation occurs in a combinatorial dynamical system, we show how to continue tracking with continuation. This gives us two “pieces” of continuation, and we can connect these by using persistence approaches from the prior two chapters. This enables us to view which Conley index generators persist through the bifurcation. Results in this section are from [20], and passages are quoted verbatim.

6.1 Tracking Isolated Invariant Sets

In this section, we introduce the protocol for tracking an isolated invariant set across multivector fields. Results in the continuous theory imply that under a sufficiently small

perturbation, some homological features of an isolated invariant set do not change. Hence, we require a notion of a small perturbation of a multivector field. In particular, let \mathcal{V} and \mathcal{V}' denote two multivector fields on K . If each multivector $V' \in \mathcal{V}'$ is contained in a multivector $V \in \mathcal{V}$, $|\mathcal{V} \setminus \mathcal{V}'| = 1$, and $|\mathcal{V}' \setminus \mathcal{V}| = 2$, then \mathcal{V}' is an *atomic refinement* of \mathcal{V} . It is so-called because \mathcal{V}' is obtained by “splitting” exactly one multivector in \mathcal{V} into two multivectors, while all the other multivectors remain the same. Symmetrically, we say that \mathcal{V} is an *atomic coarsening* of \mathcal{V}' . More broadly, it is said that \mathcal{V} and \mathcal{V}' are *atomic rearrangements* of each other. In Figures 6.1, 6.2, 6.3, 6.4, and 6.5, the two multivector fields are atomic rearrangements of each other. In these figures, we draw the multivectors that are splitting or merging in red.

Given an isolated invariant set S under \mathcal{V} , and an atomic rearrangement of \mathcal{V} denoted \mathcal{V}' , we aim to find an isolated invariant set S' that is a minimal perturbation of S . We accomplish this through two mechanisms: *continuation* and *persistence*. When we use continuation, or when we attempt to *continue*, we check if there exists an S' under \mathcal{V}' that is in some sense the same as S . If there is at least one such S' , then we choose a canonical one. This is explained in Section 6.2. If there is no S' to which we can continue, then we use persistence. In particular, we choose a canonical isolated invariant set S' under \mathcal{V}' , and while S does not continue to S' , we can use zigzag persistence to observe which features of S are absorbed by S' . We elaborate on this scheme in Section 6.3. To choose S' , we require the following results.

Proposition 6.1. [35, Corollary 4.1.22] *Let A be a convex and \mathcal{V} -compatible set. Then $\text{Inv}_{\mathcal{V}}(A)$ is an isolated invariant set.*

Proposition 6.2. [35, Proposition 4.1.21] *An invariant set S is isolated if and only if it is convex and \mathcal{V} -compatible.*

The set S is an isolated invariant set by assumption, so Proposition 6.2 implies that S is convex and \mathcal{V} -compatible. Thus, if S is also \mathcal{V}' -compatible, a natural choice is then to use Proposition 6.1 and take $S' := \text{Inv}_{\mathcal{V}'}(S)$. However, if S is not \mathcal{V}' -compatible, then the situation is more complicated. The set S is not \mathcal{V}' -compatible precisely when \mathcal{V}' is an atomic coarsening of \mathcal{V} , and the unique multivector $V \in \mathcal{V}' \setminus \mathcal{V}$, occasionally called the

merged multivector, has the properties that $V \cap S \neq \emptyset$ and $V \not\subseteq S$. In such a case, we use the notation $\langle S \cup V \rangle_{\mathcal{V}'}$ to denote the intersection of all \mathcal{V}' -compatible and convex sets that contain $S \cup V$. The simplicial complex K is \mathcal{V}' -compatible and convex, so $\langle S \cup V \rangle_{\mathcal{V}'}$ always exists and $S \subsetneq \langle S \cup V \rangle_{\mathcal{V}'}$. We observe that $\langle S \cup V \rangle_{\mathcal{V}'}$ is \mathcal{V}' -compatible and convex, and thus it is the minimal convex and \mathcal{V}' -compatible set that contains S . In such a case, we use Proposition 6.1 and take $S' := \text{Inv}_{\mathcal{V}'}(\langle S \cup V \rangle_{\mathcal{V}'})$. These principles are enumerated in the following Tracking Protocol.

Tracking Protocol

Given a nonempty isolated invariant set S under \mathcal{V} , and an atomic rearrangement of \mathcal{V} denoted \mathcal{V}' , use the following rules to find an isolated invariant set S' under \mathcal{V}' that corresponds to S .

1. Attempt to track via continuation:

- (a) If \mathcal{V}' is an atomic refinement of \mathcal{V} , then take $S' := \text{Inv}_{\mathcal{V}'}(S)$.
- (b) If \mathcal{V}' is an atomic coarsening of \mathcal{V} , and the unique merged multivector V has the property that $V \subseteq S$, then take $S' := \text{Inv}_{\mathcal{V}'}(S)$.
- (c) If \mathcal{V}' is an atomic coarsening of \mathcal{V} , and the unique merged multivector V has the property that $V \cap S = \emptyset$, then take $S' := \text{Inv}_{\mathcal{V}'}(S) = S$.
- (d) If \mathcal{V}' is an atomic coarsening of \mathcal{V} and the unique merged multivector V satisfies the formulae $V \cap S \neq \emptyset$ and $V \not\subseteq S$, then consider $A = \langle S \cup V \rangle_{\mathcal{V}'}$. If $\text{Inv}_{\mathcal{V}'}(A) = S$, then take $S' := \text{Inv}_{\mathcal{V}'}(A)$.
- (e) Else, it is impossible to track via continuation.

2. If it is impossible to track via continuation, then attempt to track via persistence:

- (f) If $A := \langle S \cup V \rangle_{\mathcal{V}'}$, then take $S' := \text{Inv}_{\mathcal{V}'}(A)$. If S and S' are isolated by a common isolating neighborhood, then use the technique in Equation 6.3 to find a zigzag filtration connecting them.

(g) Otherwise, there is no natural choice of S' . See Section 6.3.3 for a possible strategy.

We include an example of Step 1a in Figure 6.1, Step 1b in Figure 6.2, Step 1c in Figure 6.3, Step 1d in Figure 6.4, and Step 2f in Figure 6.5. Each figure depicts a multivector field and a seed isolated invariant set on the left, and an atomic rearrangement and the resulting isolated invariant set on the right. By iteratively applying this protocol (until $S' = \emptyset$, in which case we are done), we can track how an isolated invariant set changes across several atomic rearrangements. See Figure 6.6 and the associated barcode in Figure 6.7. The following proposition shows that any two multivector fields \mathcal{V}_1 and \mathcal{V}_2 can be related by a sequence of atomic rearrangements, and hence the Tracking Protocol can be used to track how an isolated invariant set changes across an arbitrary sequence of multivector fields.

Proposition 6.3. *For each pair of multivector fields \mathcal{V} and \mathcal{V}' over K , there exists a sequence $\mathcal{V} = \mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n = \mathcal{V}'$ where each \mathcal{V}_i is an atomic rearrangement of \mathcal{V}_{i-1} for $i > 1$.*

Proof. Let \mathcal{W} denote the multivector field over K where each simplex $\sigma \in K$ is contained in its own multivector. That is, $\mathcal{W} := \{\{\sigma\} \mid \sigma \in K\}$. We will show that there exists a sequence $\mathcal{V} = \mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n = \mathcal{W}$ and a sequence $\mathcal{V}' = \mathcal{V}'_1, \mathcal{V}'_2, \dots, \mathcal{V}'_m = \mathcal{W}$ where each \mathcal{V}_{i+1} (or \mathcal{V}'_{i+1}) is an atomic refinement of \mathcal{V}_i (or \mathcal{V}'_i). This immediately gives a sequence $\mathcal{V} = \mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n = \mathcal{W} = \mathcal{V}'_m, \dots, \mathcal{V}'_1 = \mathcal{V}'$. Note that because the reverse of an atomic refinement is an atomic coarsening, this sequence of multivectors is a sequence of $n - 1$ consecutive atomic refinements followed by $m - 1$ atomic coarsenings.

Hence, all that is necessary is to find a sequence $\mathcal{V} = \mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n = \mathcal{W}$ and $\mathcal{V}' = \mathcal{V}'_1, \mathcal{V}'_2, \dots, \mathcal{V}'_m = \mathcal{W}$. Without loss of generality, we consider the former. Consider an arbitrary multivector $V \in \mathcal{V}$ where $|V| > 1$ (if there is no such V , then $\mathcal{V} = \mathcal{W}$ and we are done). Take a maximal element $\sigma \in V$ (with respect to \leq), and let $\mathcal{V}_1 := (\mathcal{V} \setminus \{V\}) \cup \{\{\sigma\}\} \cup \{V \setminus \{\sigma\}\}$. Effectively, we have partitioned the multivector V into one multivector that only contains σ and the remainder of the multivector. If we iteratively repeat this process, then we must arrive at \mathcal{W} in a finite number of steps, because each step creates a new multivector consisting of a single simplex, and multivectors are never merged. Furthermore, note that each time we split a multivector constitutes an atomic refinement. Hence, by iterating this

process, we obtain a sequence $\mathcal{V} = \mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n = \mathcal{W}$ where each multivector field is related to its predecessor by atomic refinement. Thus, we can combine these sequences as previously described, and we are done. \square




Figure 6.1. Applying Step 1a to an invariant set (yellow, left) to get a new one (yellow, right).




Figure 6.2. Applying Step 1b to an invariant set (yellow, left) to get a new one (yellow, right).

6.2 Tracking via Continuation

Now, we introduce continuation in the combinatorial setting, and we justify the canonicity of the choices made in Step 1 of the Tracking Protocol. In addition, we show that if Step 1 is used to obtain S' from S , then S and S' are related by continuation. We will require the following two results.

Proposition 6.4. *[14, Proposition 5.6] Let (P, E) be an index pair under \mathcal{V} . Then $P \setminus E$ is convex and \mathcal{V} -compatible.*




Figure 6.3. Applying Step 1c to an invariant set (yellow, left) to get a new one (yellow, right). The merged vector is outside of the invariant set on the left, so the invariant sets are the same.




Figure 6.4. Applying Step 1d to an invariant set (yellow, left) to get a new one (yellow, right).




Figure 6.5. Applying Step 2f to an invariant set (yellow, left) to get a new one (yellow, right). The associated persistence barcode is depicted below the figures.

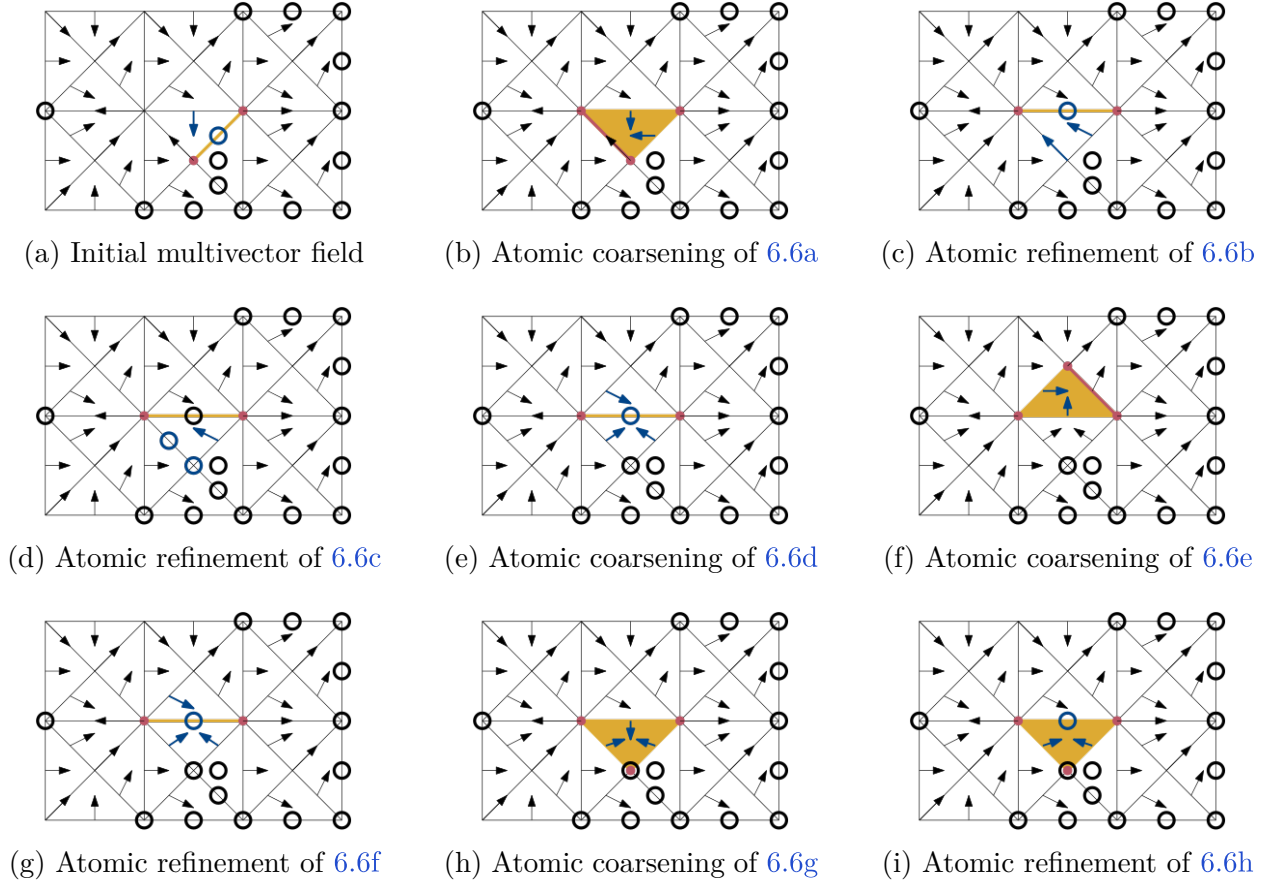


Figure 6.6. Subfigure 6.6a contains an initial multivector field and a seed isolated invariant set, which is a yellow edge. Each subsequent subfigure contains a multivector field that is an atomic refinement or atomic coarsening of the previous. The isolated invariant set that we get by iteratively applying the Tracking Protocol is depicted in yellow. Splitting and merging multivectors are in blue.



Figure 6.7. The barcode associated with the tracked invariant sets in Figure 6.6. Starting with subfigure 6.6h, we see the birth of a new 1-dimensional homology generator.

Proposition 6.5. *If A is convex and \mathcal{V} -compatible, then $(\text{cl}(A), \text{mo}(A))$ is an index pair for $\text{Inv}_{\mathcal{V}}(A)$.*

Proof. By Proposition 6.1 the set $S = \text{Inv}_{\mathcal{V}}(A)$ is an isolated invariant set. Since $\text{cl}(A) \setminus \text{mo}(A) = A$, we immediately get condition 3 from Definition 3.7.

Since A is \mathcal{V} -compatible we get $F_{\mathcal{V}}(A) = \text{cl } A$, and thus, condition 2.

To see condition 1, consider $x \in F_{\mathcal{V}}(\text{mo}(A))$. By the definition of $F_{\mathcal{V}}$, there exists a $\sigma \in \text{mo}(A)$ such that either $x \in [\sigma]_{\mathcal{V}}$ or $x \in \text{cl}(\sigma)$. In the first case $x \notin A$, because A is \mathcal{V} -compatible and $a \notin A$. Therefore $[a]_{\mathcal{V}} \cap \text{cl}(A) \subseteq \text{cl}(A) \setminus A = \text{mo}(A)$. If $x \in \text{cl}(a)$ then $x \in \text{mo}(A)$, because $\text{mo}(A)$ is closed. Hence, it follows that $F_{\mathcal{V}}(\text{mo}(A)) \cap \text{cl}(A) \subseteq \text{mo}(A)$. \square

6.2.1 Combinatorial Continuation and the Tracking Protocol

We now move to placing continuation in the combinatorial setting and explaining Step 1 of the Tracking Protocol. In essence, a continuation captures the presence of the “same” isolated invariant set across multiple multivector fields. We then show that Step 1 of the Tracking Protocol does use continuation to track an isolated invariant set.

Definition 6.6. *Let S_1, S_2, \dots, S_n denote a sequence of isolated invariant sets under the multivector fields $\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_n$, where each \mathcal{V}_i is defined on a fixed simplicial complex K . We say that isolated invariant set S_1 continues to isolated invariant set S_n whenever there exists a sequence of index pairs $(P_1, E_1), (P_2, E_2), \dots, (P_{n-1}, E_{n-1})$ where (P_i, E_i) is an index pair for both S_i and S_{i+1} . Such a sequence is a sequence of connecting index pairs.*

Each index pair (P_i, E_i) in a connecting sequence of index pairs is an index pair for a pair of consecutive isolated invariant sets S_i and S_{i+1} (see Figures 6.8 and 6.9). Hence, the isolated invariant sets in the continuation all have the same Conley index. In this sense, we are capturing the “same” isolated invariant set. In Step 1 of the Tracking Protocol, we first attempt to track the isolated invariant set S via continuation. That is, if we use Step 1, then we choose S' such that S and S' have a common index pair, say (P, E) . It so happens that (P, E) is easy to find algorithmically. We begin with the refinement case, or Step 1a.




Figure 6.8. An index pair, where P is in yellow and E is empty, for the isolated invariant sets in Figure 6.4. There is a common index pair for both isolated invariant sets, so they form a continuation.




Figure 6.9. An index pair, where P is given by the yellow and red simplices and E is given by the red simplices, for the isolated invariant sets in Figure 6.2. Thus, they form a continuation.

Theorem 6.7. *Let \mathcal{V} and \mathcal{V}' denote multivector fields where \mathcal{V}' is an atomic refinement of \mathcal{V} . Let A be a \mathcal{V} -compatible and convex set. The pair $(\text{cl}(A), \text{mo}(A))$ is an index pair for both $\text{Inv}_{\mathcal{V}}(A)$ under \mathcal{V} and $\text{Inv}_{\mathcal{V}'}(A)$ under \mathcal{V}' .*

Proof. By Proposition 6.5, if A is convex and \mathcal{V} -compatible, then $(\text{cl}(A), \text{mo}(A))$ is an index pair for $\text{Inv}_{\mathcal{V}}(A)$. By assumption, A is \mathcal{V} -compatible, so $(\text{cl}(A), \text{mo}(A))$ is an index pair for $\text{Inv}_{\mathcal{V}}(A)$. Hence, if we can show that A is \mathcal{V}' -compatible, then it will immediately follow by Proposition 6.5 that $(\text{cl}(A), \text{mo}(A))$ is an index pair for $\text{Inv}_{\mathcal{V}'}(A)$. Because A is \mathcal{V} -compatible, it follows that there exists a set of multivectors $R \subseteq \mathcal{V}$ where $A = \cup_{V \in R} V$. Recall that as \mathcal{V}' is an atomic refinement of \mathcal{V} , there is exactly one multivector $W \in \mathcal{V} \setminus \mathcal{V}'$. If $W \notin R$, then we are done, and A is necessarily \mathcal{V}' -compatible, as each multivector in R is a multivector in \mathcal{V}' . If $W \in R$, then we observe that there exist two multivectors $W_1, W_2 \in \mathcal{V}' \setminus \mathcal{V}$ where $W = W_1 \cup W_2$. In such a case, it follows easily that if $R' = (R \setminus \{W\}) \cup \{W_1, W_2\}$, then each multivector in R' is a multivector in \mathcal{V}' and $A = \cup_{V \in R'} V$. Hence, A is \mathcal{V}' -compatible, and $(\text{cl}(A), \text{mo}(A))$ is an index pair for $\text{Inv}_{\mathcal{V}'}(A)$. \square

In Step 1a of the Tracking Protocol, where \mathcal{V}' is an atomic refinement of \mathcal{V} , we choose $S' := \text{Inv}_{\mathcal{V}'}(S)$. By Proposition 6.2, it follows that S is \mathcal{V} -compatible. By identical reasoning to that presented in the proof of Theorem 6.7, it follows that S is also \mathcal{V}' -compatible. Hence, Theorem 6.7 implies that $(\text{cl}(S), \text{mo}(S))$ is an index pair for both $S = \text{Inv}_{\mathcal{V}}(S)$ and $S' = \text{Inv}_{\mathcal{V}'}(S)$. Thus, S and S' share an index pair.

The case of an atomic coarsening, corresponding to Steps 1b, 1c, and 1d of the Tracking Protocol, is more complicated. Recall that if \mathcal{V}' is an atomic coarsening of \mathcal{V} , then the unique multivector $V \in \mathcal{V}' \setminus \mathcal{V}$ is called the *merged multivector*.

Theorem 6.8. *Let \mathcal{V} and \mathcal{V}' denote multivector fields where \mathcal{V}' is an atomic coarsening of \mathcal{V} . Let A be a convex and \mathcal{V} -compatible set, and let $V \in \mathcal{V}'$ be the unique merged multivector. If $V \subseteq A$ or $V \cap A = \emptyset$, then $(\text{cl}(A), \text{mo}(A))$ is an index pair for both $\text{Inv}_{\mathcal{V}}(A)$ and $\text{Inv}_{\mathcal{V}'}(A)$.*

Proof. If $V \cap A = \emptyset$, then A is both \mathcal{V} -compatible and \mathcal{V}' -compatible. Thus, Proposition 6.5 implies that $(\text{cl}(A), \text{mo}(A))$ is an index pair for both $S = \text{Inv}_{\mathcal{V}}(A)$ and $S' = \text{Inv}_{\mathcal{V}'}(A)$.

If $V \subseteq A$, then by the same reasoning as in the proof of Theorem 6.7, it follows that A is both \mathcal{V} -compatible and \mathcal{V}' -compatible. Thus, Proposition 6.5 implies that $(\text{cl}(A), \text{mo}(A))$ is an index pair for both $\text{Inv}_{\mathcal{V}}(A)$ and $\text{Inv}_{\mathcal{V}'}(A)$. \square

By Proposition 6.2, S is convex and \mathcal{V} -compatible. Theorem 6.8 implies that if $V \subseteq S$ or $V \cap S = \emptyset$, then $(\text{cl}(S), \text{mo}(S))$ is an index pair for both $\text{Inv}_{\mathcal{V}}(S) = S$ and $\text{Inv}_{\mathcal{V}'}(S) = S'$. In Steps 1b and 1c of the Tracking Protocol, S' is chosen as $\text{Inv}_{\mathcal{V}'}(S)$. Hence, the index pair $(\text{cl}(S), \text{mo}(S))$ is an index pair for both S and S' .

A more complicated case is Step 1d, where $V \cap S \neq \emptyset$ and $V \not\subseteq S$. Recall that $A := \langle S \cup V \rangle_{\mathcal{V}'}$ denotes the intersection of all convex and \mathcal{V}' -compatible sets that contain $S \cup V$, and in particular, A is convex and \mathcal{V}' -compatible. In Step 1d of the Tracking Protocol, we first check if $S = \text{Inv}_{\mathcal{V}}(A)$. By Proposition 6.5, if $S = \text{Inv}_{\mathcal{V}}(A)$, then $(\text{cl}(A), \text{mo}(A))$ is an index pair for S . The set $\langle S \cup V \rangle_{\mathcal{V}'}$ is necessarily \mathcal{V} -compatible, because it is \mathcal{V}' -compatible by construction and it contains the unique merged multivector. Hence, Proposition 6.1 implies that $S' := \text{Inv}_{\mathcal{V}'}(A)$ is an isolated invariant set. Thus, Proposition 6.5 implies that $(\text{cl}(A), \text{mo}(A))$ is also an index pair for S' . Hence, if Step 1d gives S' , there is an index pair for S and S' .

In Step 1e of the Tracking Protocol, we claim that if $S \neq \text{Inv}_{\mathcal{V}}(A)$, then it is not possible to continue. Equivalently, there is no S' that shares an index pair with S .

Theorem 6.9. *Let S denote an isolated invariant set under \mathcal{V} and let \mathcal{V}' denote an atomic coarsening of \mathcal{V} where the unique merged multivector $V \in \mathcal{V}' \setminus \mathcal{V}$ satisfies the formulae $V \cap S \neq \emptyset$ and $V \not\subseteq S$. Furthermore, let $A := \langle S \cup V \rangle_{\mathcal{V}'}$. If $S \neq \text{Inv}_{\mathcal{V}}(A)$, then there does not exist an isolated invariant set S' under \mathcal{V}' for which there is an index pair (P, E) satisfying $\text{Inv}_{\mathcal{V}}(P \setminus E) = S$ and $\text{Inv}_{\mathcal{V}'}(P \setminus E) = S'$.*

Proof. Suppose that $S \neq \text{Inv}_{\mathcal{V}}(A)$ and there exists an index pair, (P, E) , for both S under \mathcal{V} and some S' under \mathcal{V}' . By Proposition 6.4, the set $P \setminus E$ must be convex and \mathcal{V}' -compatible. Since $S \subseteq P \setminus E$ and A is the smallest convex and \mathcal{V}' -compatible set containing S , it follows that $A \subseteq P \setminus E$. Hence, $\text{Inv}_{\mathcal{V}}(A) \subseteq \text{Inv}_{\mathcal{V}}(P \setminus E)$. By assumption, $S \subsetneq \text{Inv}_{\mathcal{V}}(A)$. Thus, $S \subsetneq \text{Inv}_{\mathcal{V}}(P \setminus E)$. This implies that (P, E) is not an index pair for S , a contradiction. \square

6.2.2 Characterizing Tracked Isolated Invariant Sets

Step 1 of the Tracking Protocol provides an avenue for tracking an isolated invariant set across a sequence of atomic rearrangements. In this subsection, we justify the canonicity of the selected isolated invariant set in Step 1 of the Tracking Protocol. First, we observe that we always have an inclusion. Theorem 6.12 follows directly from the next two results

Proposition 6.10. *Let S be an isolated invariant set under \mathcal{V} , and let S' denote an isolated invariant set under \mathcal{V}' that is obtained by applying the Tracking Protocol. If S' is obtained via Steps 1a, 1b, or 1c, then $S' \subseteq S$.*

Proof. In Steps 1a, 1b, and 1c, S' is obtained by taking $S' := \text{Inv}_{\mathcal{V}'}(S)$. By definition, $\text{Inv}_{\mathcal{V}'}(S) \subseteq S$, so $S' \subseteq S$. \square

Proposition 6.11. *Let S be an isolated invariant set under \mathcal{V} , and let S' denote an isolated invariant set under \mathcal{V}' that is obtained via applying the Tracking Protocol. If S' is obtained via Step 1d then $S \subseteq S'$ or $S' \subseteq S$.*

Proof. First, we claim that if $S \not\subseteq S'$, and if V is the unique merged multivector $V \in \mathcal{V}' \setminus \mathcal{V}$, then $V \cap S' = \emptyset$.

If $S \not\subseteq S'$, then there exists a $\sigma \in S \setminus S'$. Because $\sigma \in S$, there exists an essential solution $\rho : \mathbb{Z} \rightarrow A$ under \mathcal{V} where $\rho(0) = \sigma$. It is easy to check that because \mathcal{V}' is an atomic coarsening of \mathcal{V} , we have that for every $\tau \in K$, $F_{\mathcal{V}}(\tau) \subseteq F_{\mathcal{V}'}(\tau)$. Hence, ρ must be a solution under \mathcal{V}' . But, since $\sigma \in S \setminus S'$, it follows that ρ is not an essential solution.

Without loss of generality, we assume that there exists a $j > 0$ such that for all $i_1, i_2 \geq j$, we have that $[\rho(i_1)]_{\mathcal{V}'} = [\rho(i_2)]_{\mathcal{V}'}$. Because ρ is an essential solution under \mathcal{V} , and $|\mathcal{V}' \setminus \mathcal{V}| = 1$, it follows that $[\rho(i_1)]_{\mathcal{V}'} = [\rho(i_2)]_{\mathcal{V}'} = V$. Hence, V must not be critical, as if it were, then ρ would be an essential solution under \mathcal{V}' .

Now, aiming for a contradiction, assume there exists a $\tau \in V \cap S'$. Then there exists an essential solution $\rho' : \mathbb{Z} \rightarrow A$ under \mathcal{V}' where $\rho'(j+1) = \tau$. Thus, because $\rho(j) \in V$, we can obtain a new solution $r : \mathbb{Z} \rightarrow S'$ where $r(i) = \rho(i)$ if $i \leq j$ and $r(i) = \rho'(i)$ if $i > j$. We have shown that $\rho(j) \in V$, and by assumption, $\tau = \rho'(j+1)$ has the property that $\tau \in V$. Hence, because ρ is a solution and ρ' is an essential solution, we have the property that for

all k , there exists an $i > k$ where $[r(i)]_{\mathcal{V}'} \neq [r(k)]_{\mathcal{V}'}$. We can use the same construction to guarantee that there exists an $i < k$ where $[r(i)]_{\mathcal{V}'} \neq [r(k)]_{\mathcal{V}'}$. Hence, r is an essential solution under \mathcal{V}' , where $r(0) = \sigma$. But this implies that $\sigma \in S'$, a contradiction. Hence, there can exist no such τ , so $V \cap S' = \emptyset$.

Thus, we have the property that $V \cap S' = \emptyset$. Let $\rho : \mathbb{Z} \rightarrow S'$ denote an essential solution under \mathcal{V}' . Observe that for each i , $[\rho(i)]_{\mathcal{V}'} = [\rho(i)]_{\mathcal{V}}$. Ergo, ρ is also an essential solution under \mathcal{V} . Hence, $S' \subseteq S$. \square

Theorem 6.12. *If S' is obtained by applying Step 1 of the Tracking Protocol to S , then we have $S \subseteq S'$ or $S' \subseteq S$.*

Furthermore, isolated invariant sets chosen by Step 1 minimize the perturbation to S in terms of the number of inclusions.

Proposition 6.13. *Let S be an isolated invariant set under \mathcal{V} , and let S' be an isolated invariant set under \mathcal{V}' that is obtained by applying Step 1 of the Tracking Protocol to S . If S'' is any isolated invariant set under \mathcal{V}' that shares a common index pair with S , then $S' \subseteq S''$. Moreover, if $S'' \subseteq S$, then $S' = S''$.*

Proof. Let (P, E) be a common index pair for S under \mathcal{V} and S'' under \mathcal{V}' . Consider Steps 1a, 1b, and 1c where $S' = \text{Inv}_{\mathcal{V}'}(S)$. By definition, $S \subseteq P \setminus E$, and it follows that $S' = \text{Inv}_{\mathcal{V}'} S \subseteq \text{Inv}_{\mathcal{V}'}(P \setminus E) = S''$. Moreover, if $S'' \subseteq S$, we get that $S'' = \text{Inv}_{\mathcal{V}'} S'' \subseteq \text{Inv}_{\mathcal{V}'} S = S'$. Thus, $S' = S''$.

To prove the property for Step 1d, notice that by Proposition 6.4, and the fact that $A := \langle S \cup V \rangle_{\mathcal{V}'}$ is the minimal convex and \mathcal{V}' -compatible set containing S , we get $A \subseteq P \setminus E$. Therefore $S' = \text{Inv}_{\mathcal{V}'}(A) \subseteq \text{Inv}_{\mathcal{V}'}(P \setminus E) = S''$. Similarly, if $S'' \subseteq S \subseteq A$, then we get $S'' = \text{Inv}_{\mathcal{V}'} S'' \subseteq \text{Inv}_{\mathcal{V}'} A = S'$. Thus, $S' = S''$. \square

6.3 Tracking via Persistence

In the previous section, we explicated Step 1 of the protocol, which uses continuation to track an isolated invariant set across a changing multivector field. In this section, we

first place continuation in the persistence framework by showing how to translate the idea of combinatorial continuation into a zigzag filtration [4], [25] that does not introduce spurious information. Then, we use the persistence view of continuation to justify Step 2f of the Tracking Protocol, which permits us to capture changes in an isolated invariant set when no continuation is possible. In particular, it permits us to track an isolated invariant set even in the presence of a bifurcation that changes the Conley index. If the isolated invariant set that we are tracking collides, or *merges*, with another isolated invariant set, then we follow the newly formed isolated invariant set, and persistence captures which aspects of our original isolated invariant set persist into the new one. Conversely, if an isolated invariant set splits, we track the smallest isolated invariant set that contains all of the child invariant sets.

6.3.1 From continuation to filtration

Now, we show that a continuation of an isolated invariant set S_1 to S_{n+1} can be expressed in terms of persistence. Namely, a corresponding sequence of connecting index pairs (P_1, E_1) , (P_2, E_2) , \dots , (P_n, E_n) can be turned into a *zigzag filtration*, that is a sequence of pairs $\{(A_i, B_i)\}_{i=1}^m$ such that either $(A_i, B_i) \subseteq (A_{i+1}, B_{i+1})$ or $(A_{i+1}, B_{i+1}) \subseteq (A_i, B_i)$. Ideally, each (A_i, B_i) would be an index pair for some S_j from the initial continuation so as to not introduce spurious invariant sets or Conley indices. A connecting index pair (P_i, E_i) is an index pair for both S_i under \mathcal{V}_i and for S_{i+1} under \mathcal{V}_{i+1} . Thus, (P_i, E_i) and (P_{i+1}, E_{i+1}) are both index pairs for S_{i+1} under \mathcal{V}_{i+1} . We will construct auxiliary index pairs for S_{i+1} and then relate (P_i, E_i) and (P_{i+1}, E_{i+1}) with a zigzag filtration using these auxiliary pairs. If we can connect all adjacent pairs (P_i, E_i) and (P_{i+1}, E_{i+1}) with a zigzag filtration, then we can concatenate all of these zigzag filtrations and transform a sequence of connecting index pairs into a larger zigzag filtration. The following results are important for achieving this.

Proposition 6.14. [14, Proposition 5.2] *Let (P, E) denote an index pair for S . The set P is an isolating neighborhood for S .*

Proposition 6.15. *Let (P, E) denote an index pair for S under \mathcal{V} . The pair (P, E) is an index pair for S in P under \mathcal{V} .*

Proof. First, we observe that $S = \text{Inv}_{\mathcal{V}}(P \setminus E)$ because (P, E) is an index pair. In addition, $F_{\mathcal{V}}(P) \cap N = F_{\mathcal{V}}(P) \cap P \subseteq P$ by definition. Since (P, E) is an index pair, it has the property that $F_{\mathcal{V}}(P \setminus E) \subseteq P$. In the case of index pairs in N , we require that $F_{\mathcal{V}}(P \setminus E) \subseteq N = P$, so this case is immediately satisfied. Finally, because (P, E) is an index pair, $F_{\mathcal{V}}(E) \cap P \subseteq E$. Thus, $F_{\mathcal{V}}(E) \cap N = F_{\mathcal{V}}(E) \cap P \subseteq E$. \square

Theorem 6.16. *Let (P_1, E_1) and (P_2, E_2) denote index pairs for S in N under \mathcal{V} . The pair $(P_1 \cap P_2, E_1 \cap E_2)$ is an index pair for S in N under \mathcal{V} .*

Proof. By Theorem 4.4, the pair $(P_1 \cap P_2, E_1 \cap E_2)$ is an index pair in N under \mathcal{V} . Hence, it is sufficient to show that $\text{Inv}_{\mathcal{V}}((P_1 \cap P_2) \setminus (E_1 \cap E_2)) = S$. Furthermore, $S = \text{Inv}_{\mathcal{V}}(P_1 \setminus E_1)$ and $S = \text{Inv}_{\mathcal{V}}(P_2 \setminus E_2)$. Hence, $S \subseteq P_1 \setminus E_1$ and $S \subseteq P_2 \setminus E_2$. Ergo, $S \subseteq P_1 \cap P_2$. In addition, $S \cap E_1 = \emptyset$ and $S \cap E_2 = \emptyset$. Hence, $S \cap (E_1 \cap E_2) = \emptyset$, so it follows that $S \subseteq (P_1 \cap P_2) \setminus (E_1 \cap E_2)$. Thus, $S \subseteq \text{Inv}_{\mathcal{V}}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$. Ergo, it remains to be shown that $\text{Inv}_{\mathcal{V}}((P_1 \cap P_2) \setminus (E_1 \cap E_2)) \subseteq S$.

Aiming for a contradiction, assume that there exists an $\sigma \in \text{Inv}_{\mathcal{V}}((P_1 \cap P_2) \setminus (E_1 \cap E_2)) \setminus S$. Equivalently, there exists an essential solution $\rho : \mathbb{Z} \rightarrow (P_1 \cap P_2) \setminus (E_1 \cap E_2)$ where $\rho(0) = \sigma$. Because $\rho(\mathbb{Z}) \subseteq P_1 \cap P_2$, but $\rho(\mathbb{Z}) \not\subseteq \text{Inv}(P_1 \setminus E_1)$, there must exist an $i_1 \in \mathbb{Z}$ where $\rho(i_1) \in E_1$. Similarly, there must exist an $i_2 \in \mathbb{Z}$ where $\rho(i_2) \in E_2$.

We claim that for all $i \geq i_1$, $\rho(i) \in E_1$. To contradict, assume that this is not the case. Then there must exist some first $j > i_1$ where $\rho(j) \notin E_1$. However, $\rho(j-1) \in E_1$. By definition of an index pair in N , if $x \in E_1$ and $y \in F_{\mathcal{V}}(x) \cap N$, then $y \in E_1$. Hence, since $\rho(j) \notin E_1$, it follows that $\rho(j) \notin N$. But by assumption, $\rho(j) \in P_1 \cap P_2 \subseteq N$. Therefore, there is no such j , so for all $i > i_1$, we have that $i \in E_1$. The same argument implies that for all $i \geq i_2$, we have that $\rho(i) \in E_2$.

Thus, it follows that for all $i \geq \max\{i_1, i_2\}$, $\rho(i) \in E_1 \cap E_2$. Ergo, $\rho(\mathbb{Z}) \not\subseteq (P_1 \cap P_2) \setminus (E_1 \cap E_2)$, a contradiction. Hence, no such ρ can exist, which implies that no such σ can exist. Thus, $S = \text{Inv}_{\mathcal{V}}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$. \square

Now, we move to using these results to translate a sequence of connecting index pairs $\{(P_i, E_i)\}_{i=1}^n$ into a zigzag filtration. For $1 < i \leq n$, (P_{i-1}, E_{i-1}) and (P_i, E_i) are both index pairs for S_i . By Proposition 3.8, the pair $(\text{cl}(S_i), \text{mo}(S_i))$ is an index pair for S_i . Hence, a

natural approach is to find a zigzag filtration that connects (P_i, E_i) with $(\text{cl}(S_i), \text{mo}(S_i))$ and a zigzag filtration that connects (P_{i-1}, E_{i-1}) with $(\text{cl}(S_i), \text{mo}(S_i))$. If we can find such zigzag filtrations for all S_i , then we can concatenate all of them and obtain a zigzag filtration that connects (P_1, E_1) with (P_n, E_n) . We depict the resulting zigzag filtration in Equation 6.1.

$$(P_1, E_1) \supseteq \dots \supseteq (\text{cl}(S_2), \text{mo}(S_2)) \subseteq \dots \subseteq (P_2, E_2) \supseteq \dots \supseteq (\text{cl}(S_3), \text{mo}(S_3)) \subseteq \dots (P_n, E_n) \quad (6.1)$$

We connect $(\text{cl}(S_i), \text{mo}(S_i))$ with (P_i, E_i) , and (P_{i-1}, E_{i-1}) connects with $(\text{cl}(S_i), \text{mo}(S_i))$ symmetrically. By Proposition 6.14, P_i is an isolating neighborhood for S_i . Thus, by Theorem 4.9, $(\text{pf}_{\mathcal{V}_i}(\text{cl}(S_i), P_i), \text{pf}_{\mathcal{V}_i}(\text{mo}(S_i), P_i))$ is an index pair for S_i in P_i . Proposition 6.15 implies that (P_i, E_i) is an index pair for S_i in P_i . By Theorem 6.16, $(P_i \cap \text{pf}_{\mathcal{V}_i}(\text{cl}(S_i), P_i), E_i \cap \text{pf}_{\mathcal{V}_i}(\text{mo}(S_i), P_i))$ is an index pair for S_i in P_i . Hence, we get the following zigzag filtration:

$$\begin{aligned} (\text{cl}(S_i), \text{mo}(S_i)) \subseteq (\text{pf}_{\mathcal{V}_i}(\text{cl}(S_i), P_i), \text{pf}_{\mathcal{V}_i}(\text{mo}(S_i), P_i)) \supseteq \\ (P_i \cap \text{pf}_{\mathcal{V}_i}(\text{cl}(S_i), P_i), E_i \cap \text{pf}_{\mathcal{V}_i}(\text{mo}(S_i), P_i)) \subseteq (P_i, E_i) \end{aligned} \quad (6.2)$$

Every pair in Equation 6.2 is an index pair for S_i under \mathcal{V}_i . Thus, we do not introduce any spurious invariant sets. We can concatenate these filtrations to get Equation 6.1.

We now analyze the barcode obtained for 6.1. Our main result is Theorem 6.19, and it follows immediately from the next two results.

Lemma 6.17. [14, Lemma 5.10] *Let $(P, E) \subseteq (P', E')$ be index pairs for isolated invariant set S under \mathcal{V} such that either $P = P'$ or $E = E'$. Then the inclusion $i : (P, E) \hookrightarrow (P', E')$ induces an isomorphism in homology.*

Theorem 6.18. [20, Theorem 26] *If (P, E) and (P', E') are index pairs for S where $(P', E') \subseteq (P, E)$, then the inclusion induces an isomorphism in the Conley indices.*

Theorem 6.19. *For every $k \geq 0$, the k -dimensional barcode of a connecting sequence of index pairs $\{(P_i, E_i)\}_{i=1}^n$ has m bars $[1, n]$ if $\dim H_k(P_1, E_1) = m$.*

6.3.2 Tracking beyond continuation

In the previous subsection, we showed how to convert a connecting sequence of index pairs into a zigzag filtration. Furthermore, we observed that it produces “full” barcodes - they have one bar for each basis element of the Conley index that persists for the length of the filtration. This change of perspective allows us to generalize our protocol to handle cases when it is impossible to continue.

In particular, we consider Step 2f of the protocol. Let S denote an isolated invariant set under \mathcal{V} , and \mathcal{V}' is an atomic coarsening of \mathcal{V} where the merged multivector V has the property that $V \cap S \neq \emptyset$ and $V \not\subseteq S$. In such a case, we consider $A := \langle S \cup V \rangle_{\mathcal{V}'}$ and take $S' = \text{Inv}_{\mathcal{V}'}(A)$. Theorem 6.9 implies that if $S \neq \text{Inv}_{\mathcal{V}}(A)$, then it is impossible to continue. However, it may be possible to compute persistence in a way that resembles continuation. Let $B := \text{cl}(S) \cup \text{cl}(S')$. Trivially, B is closed. If B is an isolating neighborhood for both S and S' , then we say that S and S' are *adjacent*. By Theorem 4.9, $(\text{pf}_{\mathcal{V}}(\text{cl}(S), B), \text{pf}_{\mathcal{V}}(\text{mo}(S), B))$ is an index pair for S in B . Similarly, $(\text{pf}_{\mathcal{V}'}(\text{cl}(S'), B), \text{pf}_{\mathcal{V}'}(\text{mo}(S'), B))$ is an index pair for S' in B . Thus, we can use Theorem 4.4 to obtain the following zigzag filtration.

$$\begin{aligned} (\text{cl}(S), \text{mo}(S)) &\subseteq (\text{pf}_{\mathcal{V}}(\text{cl}(S), B), \text{pf}_{\mathcal{V}}(\text{mo}(S), B)) \\ &\supseteq (\text{pf}_{\mathcal{V}}(\text{cl}(S), B) \cap \text{pf}_{\mathcal{V}'}(\text{cl}(S'), B), \text{pf}_{\mathcal{V}}(\text{mo}(S), B) \cap \text{pf}_{\mathcal{V}'}(\text{mo}(S'), B)) \subseteq \\ &\quad (\text{pf}_{\mathcal{V}'}(\text{cl}(S'), B), \text{pf}_{\mathcal{V}'}(\text{mo}(S'), B)) \supseteq (\text{cl}(S'), \text{mo}(S')) \end{aligned} \quad (6.3)$$

Suppose that we are iteratively applying Step 1 of the Tracking Protocol, finding a sequence of isolated invariant sets where adjacent ones share an index pair, and we terminate with an isolated invariant set S and an index pair (P, E) . We can connect (P, E) with $(\text{cl}(S), \text{mo}(S))$ with techniques from the previous section. That is, if $(P, E) \neq (\text{cl}(S), \text{mo}(S))$, then we can find a filtration that connects them:

$$\begin{aligned} (P, E) &\supseteq (P \cap \text{pf}_{\mathcal{V}}(\text{cl}(S), P), E \cap \text{pf}_{\mathcal{V}}(\text{mo}(S), P)) \subseteq \\ &\quad (\text{pf}_{\mathcal{V}}(\text{cl}(S), P), \text{pf}_{\mathcal{V}}(\text{mo}(S), P)) \supseteq (\text{cl}(S), \text{mo}(S)) \end{aligned} \quad (6.4)$$

We can then concatenate this filtration with the zigzag filtration in Equation 6.3. This effectively completes the Tracking Protocol: when continuation, represented as Step 1, is impossible, we can attempt to apply Step 2f and persistence to continue to track.

In Step 2f, we choose to take $S' = \text{Inv}_{\mathcal{V}'}(A)$. In practice, there may be many isolated invariant sets under \mathcal{V}' that are adjacent to S . However, our choice of S' is canonical.

Proposition 6.20. *Let S' denote an isolated invariant set under \mathcal{V}' that is obtained from applying Step 2f of the Tracking Protocol to the isolated invariant set S under \mathcal{V} . If S'' is an isolated invariant set under \mathcal{V}' where $S \subseteq S''$, then $S' \subseteq S''$.*

Proof. By Proposition 6.1, the set S'' is convex and \mathcal{V}' -compatible. Since A is the minimal convex and \mathcal{V}' -compatible set containing S , it follows that $S \subseteq A \subseteq S''$. By definition, $S' = \text{Inv}_{\mathcal{V}'}(A)$, so $S' \subseteq A \subseteq S''$. \square

6.3.3 Strategy for Step 2g of the Tracking Protocol

We briefly consider Step 2g of Tracking Protocol, which is the case when it is impossible to continue from S to some S' , and $B = \text{cl}(S) \cup \text{cl}(S')$ does not isolate both S and S' . In such a case, we do not have two index pairs in the same isolating neighborhood, so we cannot use Theorem 4.9 to obtain index pairs in a common N . Thus, we cannot use Theorem 4.4, which permits us to intersect index pairs and guarantee that the resulting pair is an index pair under a known multivector field. Recall that $A = \langle S \cup V \rangle_{\mathcal{V}'}$. A natural choice is to take $S' = \text{Inv}_{\mathcal{V}'}(A)$, and consider the zigzag filtration

$$(\text{cl}(S), \text{mo}(S)) \supseteq (\text{cl}(S) \cap \text{cl}(S'), \text{mo}(S) \cap \text{mo}(S')) \subseteq (\text{cl}(S'), \text{mo}(S')). \quad (6.5)$$

It is easy to construct examples where $(\text{cl}(S) \cap \text{cl}(S'), \text{mo}(S) \cap \text{mo}(S'))$ is not an index pair under any natural choice of multivector field (see Figure 4.2). However, this approach may work well in practice. We leave a thorough examination to future work.

7. CONCLUSION

In this dissertation, we have introduced three different approaches for tracking changes in combinatorial dynamical systems. The first, explicated in Chapter 4, is to choose an isolating neighborhood and an isolated invariant set under each multivector field and to compute the persistence barcode by intersecting index pairs for these isolated invariant sets. Such a scheme is simple and powerful, but it can only capture the persistence of the Conley index at a single resolution. This omits information about the Morse sets within a given isolated invariant set. To address this concern, we introduced the persistence of Conley-Morse graphs in Chapter 5. This scheme allows us to obtain a barcode that represents the changing Conley index at all of the Morse sets in a Morse decomposition, while also capturing information that represents the changing structure of the Conley-Morse graph. Finally, we placed Conley's notion of continuation in the combinatorial setting, and we showed how to combine it with Conley index persistence to track changes in a combinatorial dynamical system.

There are several possible directions for future work. In Chapter 4, we showed that if (P_1, E_1) is an index pair in N under \mathcal{V}_1 and (P_2, E_2) is an index pair in N under \mathcal{V}_2 , then $(P_1 \cap P_2, E_1 \cap E_2)$ is an index pair in N under $\mathcal{V}_1 \overline{\cap} \mathcal{V}_2$. However, we know a priori that (P_1, E_1) is an index pair for some isolated invariant set S_1 and (P_2, E_2) is an index pair for S_2 . We have shown that $(P_1 \cap P_2, E_1 \cap E_2)$ is an index pair for $\text{Inv}((P_1 \cap P_2) \setminus (E_1 \cap E_2))$, but we do not know if there are any conditions under which this invariant set is somehow related to S_1 or S_2 . Further study is needed.

In Chapter 5, we encountered a scenario where some bars in our barcode were redundant. That is, they represented the same feature multiple times. We were able to remove them by checking if they corresponded to the same part of a zigzag filtration, but it would be good to find a more systematic way to do this.

Finally, in Chapter 6, we encountered Step 2g, where it is not possible to track our isolated invariant set while computing persistence in a controlled way. This is related to the problem that we first confronted in Chapter 4, when we noticed that the intersection of two arbitrary index pairs need not be an index pair under a natural choice of multivector field. We solved this problem by taking the intersection of index pairs in N . Perhaps there are

conditions where we can intersect index pairs and obtain an index pair, which would permit us to track the isolated invariant set in a controlled way. More investigation is needed.

REFERENCES

- [1] H. Edelsbrunner, D. Letscher, and A. Zomorodian, “Topological persistence and simplification,” *Discrete Comput. Geom.*, vol. 28, no. 4, pp. 511–533, Nov. 2002. DOI: [10.1007/s00454-002-2885-2](https://doi.org/10.1007/s00454-002-2885-2).
- [2] G. Carlsson, “Topology and data,” *Bull. Amer. Math. Soc. (N.S.)*, vol. 46, no. 2, pp. 255–308, Apr. 2009.
- [3] H. Edelsbrunner and J. Harer, *Computational Topology: An Introduction*. American Mathematical Society, Jan. 2010.
- [4] T. K. Dey and Y. Wang, *Computational Topology for Data Analysis*. Cambridge University Press, 2022, <https://www.cs.purdue.edu/homes/tamaldey/book/CTDAbook/CTDAbook.pdf>.
- [5] L. Vietoris, “Über den höheren zusammenhang kompakter räume und eine klasse von zusammenhangstreuen abbildungen,” *Mathematische Annalen*, vol. 97, no. 1, pp. 454–472, 1927. DOI: [10.1007/BF01447877](https://doi.org/10.1007/BF01447877).
- [6] H. Adams, T. Emerson, M. Kirby, *et al.*, “Persistence images: A stable vector representation of persistent homology,” *J. Mach. Learn. Res.*, vol. 18, no. 8, 2017.
- [7] P. Bubenik, “Statistical topological data analysis using persistence landscapes,” *J. Mach. Learn. Res.*, vol. 16, no. 1, 2015.
- [8] T. K. Dey and S. Mandal, “Protein Classification with Improved Topological Data Analysis,” in *18th International Workshop on Algorithms in Bioinformatics (WABI 2018)*, ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 113, 2018, 6:1–6:13, ISBN: 978-3-95977-082-8. DOI: [10.4230/LIPIcs.WABI.2018.6](https://doi.org/10.4230/LIPIcs.WABI.2018.6).
- [9] J. A. Perea and J. Harer, “Sliding windows and persistence: An application of topological methods to signal analysis,” *Found. Comput. Math.*, vol. 15, no. 3, pp. 799–838, 2015. DOI: [10.1007/s10208-014-9206-z](https://doi.org/10.1007/s10208-014-9206-z).
- [10] S. Mukherjee, “Denoising with discrete morse theory,” *The Visual Computer*, vol. 37, no. 9, pp. 2883–2894, 2021. DOI: [10.1007/s00371-021-02255-7](https://doi.org/10.1007/s00371-021-02255-7). [Online]. Available: <https://doi.org/10.1007/s00371-021-02255-7>.
- [11] M. E. Aktas, E. Akbas, and A. E. Fatmaoui, “Persistence homology of networks: Methods and applications,” *Applied Network Science*, vol. 4, no. 1, p. 61, 2019.

- [12] J. Townsend, C. P. Micucci, J. H. Hymel, V. Maroulas, and K. D. Vogiatzis, “Representation of molecular structures with persistent homology for machine learning applications in chemistry,” *Nature Communications*, vol. 11, no. 1, p. 3230, 2020.
- [13] M. Mrozek, “Conley–Morse–Forman theory for combinatorial multivector fields on Lefschetz complexes,” *Found. Comput. Math.*, vol. 17, no. 6, pp. 1585–1633, Dec. 2017. DOI: [10.1007/s10208-016-9330-z](https://doi.org/10.1007/s10208-016-9330-z).
- [14] M. Lipiński, J. Kubica, M. Mrozek, and T. Wanner, *Conley-Morse-Forman theory for generalized combinatorial multivector fields on finite topological spaces*, 2021. arXiv: [1911.12698](https://arxiv.org/abs/1911.12698) [[math.DS](#)].
- [15] R. Forman, “Morse theory for cell complexes,” *Adv. Math.*, vol. 134, pp. 90–145, 1998. DOI: [10.1006/aima.1997.1650](https://doi.org/10.1006/aima.1997.1650).
- [16] R. Forman, “Combinatorial vector fields and dynamical systems,” *Math. Z.*, vol. 228, pp. 629–681, 1998. DOI: [10.1007/PL00004638](https://doi.org/10.1007/PL00004638).
- [17] B. Batko, T. Kaczynski, M. Mrozek, and T. Wanner, “Linking combinatorial and classical dynamics: Conley index and Morse decompositions,” *Found. Comput. Math.*, vol. 20, no. 5, pp. 967–1012, 2020. DOI: [10.1007/s10208-020-09444-1](https://doi.org/10.1007/s10208-020-09444-1).
- [18] T. K. Dey, M. Mrozek, and R. Slechta, “Persistence of the Conley index in combinatorial dynamical systems,” in *36th International Symposium on Computational Geometry (SoCG 2020)*, Zürich, Switzerland, Jun. 2020, 37:1–37:17. DOI: [10.4230/LIPIcs.SoCG.2020.37](https://doi.org/10.4230/LIPIcs.SoCG.2020.37).
- [19] T. K. Dey, M. Mrozek, and R. Slechta, “Persistence of Conley-Morse graphs in combinatorial dynamical systems,” *SIAM J. Appl. Dyn. Syst.*, 2022, To appear.
- [20] T. K. Dey, M. Lipiński, M. Mrozek, and R. Slechta, “Tracking dynamical features via continuation and persistence,” in *38th International Symposium on Computational Geometry (SoCG 2022)*, To appear, 2022.
- [21] T. Kaczynski, M. Mrozek, and T. Wanner, “Towards a formal tie between combinatorial and classical vector field dynamics,” *J. Comput. Dyn.*, vol. 3, no. 1, pp. 17–50, 2016. DOI: [10.3934/jcd.2016002](https://doi.org/10.3934/jcd.2016002).
- [22] M. Mrozek and T. Wanner, “Creating semiflows on simplicial complexes from combinatorial vector fields,” *J. Differential Equations*, vol. 304, pp. 375–434, 2021. DOI: [10.1016/j.jde.2021.10.001](https://doi.org/10.1016/j.jde.2021.10.001).
- [23] J. Munkres, *Topology*, ser. Featured Titles for Topology Series. Prentice Hall, Incorporated, 2000, ISBN: 9780131816299.

- [24] A. Hatcher, *Algebraic Topology*. Cambridge: Cambridge University Press, 2002, pp. xii+544.
- [25] G. Carlsson and V. de Silva, “Zigzag persistence,” *Found. Comput. Math.*, vol. 10, no. 4, pp. 367–405, Aug. 2010. DOI: [10.1007/s10208-010-9066-0](https://doi.org/10.1007/s10208-010-9066-0).
- [26] N. Milosavljević, D. Morozov, and P. Skraba, “Zigzag persistent homology in matrix multiplication time,” in *Proceedings of the Twenty-Seventh Annual Symposium on Computational Geometry*, ser. SoCG ’11, Paris, France: Association for Computing Machinery, 2011, pp. 216–225. DOI: [10.1145/1998196.1998229](https://doi.org/10.1145/1998196.1998229).
- [27] J. R. Bunch and J. E. Hopcroft, “Triangular factorization and inversion by fast matrix multiplication,” *Mathematics of Computation*, vol. 28, no. 125, pp. 231–236, 1974. DOI: [10.1090/s0025-5718-1974-0331751-8](https://doi.org/10.1090/s0025-5718-1974-0331751-8).
- [28] J. A. Barmak, *Algebraic Topology of Finite Topological Spaces and Applications*, ser. Lecture Notes in Math. Springer Verlag, Berlin - Heidelberg - New York, 2011, vol. 2032.
- [29] N. P. Bhatia and G. P. Szegő, *Dynamical Systems: Stability Theory and Applications*, ser. Lecture Notes in Math. Springer Verlag, Berlin - Heidelberg - New York, 1967, vol. 35.
- [30] C. Conley, “Isolated invariant sets and the Morse index,” in *CBMS Reg. Conf. Ser. Math.*, vol. 38, 1978.
- [31] K. Mischaikow and M. Mrozek, *The Conley Index*, ser. Handbook of Dynamical Systems II: Towards Applications. (B. Fiedler, ed.) North-Holland, 2002.
- [32] J. Hale and H. Koçak, *Dynamics and Bifurcations*, ser. Texts in Applied Mathematics **3**. Springer-Verlag, 1991.
- [33] T. K. Dey, M. Juda, T. Kapela, J. Kubica, M. Lipiński, and M. Mrozek, “Persistent homology of Morse decompositions in combinatorial dynamics,” *SIAM J. Appl. Dyn. Syst.*, vol. 18, no. 1, pp. 510–530, 2019. DOI: [10.1137/18M1198946](https://doi.org/10.1137/18M1198946).
- [34] T. K. Dey, F. Fan, and Y. Wang, “Computing topological persistence for simplicial maps,” in *SOCG’14: Proceedings of the Thirtieth Annual Symposium on Computational Geometry*, Kyoto, Japan, 2014, pp. 345–354. DOI: [10.1145/2582112.2582165](https://doi.org/10.1145/2582112.2582165).
- [35] M. Lipiński, “Morse-conley-forman theory for generalized combinatorial multivector fields on finite topological spaces,” Ph.D. dissertation, Jagiellonian University, 2021.

VITA

Ryan Slechta is a PhD candidate in the Department of Computer Science at Purdue University. He received his master's degree in computer science and engineering from The Ohio State University in 2020, and he received a bachelor's degree in mathematics and computer science *summa cum laude* from the University of St. Thomas.