CATALOGING THE GLOBAL BEHAVIOR OF DYNAMICAL SYSTEMS: ADAPTIVELY SEARCHING PARAMETER SPACE USING THE CONLEY-MORSE DATABASE

by

Kelly Tulare Spendlove

A thesis submitted in partial fulfillment
of the requirements for the degree

of

Master of Science

in

Mathematics

MONTANA STATE UNIVERSITY
Bozeman, Montana

July, 2013

APPROVAL

of a thesis submitted by

Kelly Tulare Spendlove

This thesis has been read by each member of the thesis committee and has been found to be satisfactory regarding content, English usage, format, citations, bibliographic style, and consistency, and is ready for submission to The Graduate School.

Dr. Tomáš Gedeon

Approved for the Department of Mathematical Sciences

Dr. Ken Bowers

Approved for The Graduate School

Dr. Ronald W. Larsen

STATEMENT OF PERMISSION TO USE

Kelly Tulare Spendlove

July, 2013

ACKNOWLEDGEMENTS

I am most grateful to my advisor, Tomáš Gedeon, for his enthusiasm and support over the last two years. I have had many wonderful opportunities available to me due to Tomáš, for which I am very thankful. It has been a pleasure to work together.

There are many others who have been inspiring: Konstantin Mischaikow, who seems to be a machine of great ideas; Marcy Barge, for always being receptive to questions and offering a cohomology class; Sarah Day, for her excellent instruction on Conley Theory; Jesse Berwald, for our great collaborations and conversations; and Drew LaMar, for introducing me to good coffee.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Figure Page

ABSTRACT

The aim of this thesis is to build upon a combinatorial-topological framework to global dynamics of multiparameter dynamical systems. A combinatorial multivalued map of the dynamics for each subset of the parameter range is computed using rigorous numerical methods and is represented via a directed graph. The dynamics is then decomposed into the recurrent and gradient-like parts by graph theoretic algorithms using an adaptive computation. The novelty of this thesis is to introduce a similar adaptive scheme in parameter space. Furthermore, it is proven that this scheme produces an output which is naturally coarser than the output of an original computation. Incorporating previous results, we make an estimate for the savings achieved by this adaptive scheme in the setting of a saddle-node bifurcation. Furthermore, we make an empirical comparison of how well our scheme approximates an original computation.

CHAPTER 1

INTRODUCTION

1. Introduction

Bifurcation theory tells us that even seemingly simple multiparameter nonlinear dynamical systems can exhibit dramatically different dynamics that varies over multiple scales in both phase space and parameter space. However, in the context of applications the associated nonlinear systems are only meaningful down to a particular scale due to modeling assumptions and limits to accurate or precise measurements. In these settings an understanding of robust identifiable dynamic structures that exist over reasonable ranges of parameter space is of primary importance. Recently, [2, 3] introduced a computational method based on Conley's approach to dynamics [4] which provides a finite algebraic-combinatorial description of the global dynamics over large regions of parameter space. Because this description is finite and queriable it is referred to as a *database for the dynamics*.

The aim of this thesis is to introduce new techniques designed to reduce the cost of computing the database. We consider a multiparameter dynamical system given in the form of a continuous map

$$\begin{aligned} f\colon \mathbb{R}^n \times [0,1]^m &\to \mathbb{R}^n \\ (x,z) &\mapsto f_z(x) := f(x,z) \end{aligned} \tag{1.1}$$

where $\mathbb{R}^n$ is the phase space and $[0,1]^m$ is the parameter space. As indicated above our focus is on characterizing dynamics that is stable with respect to perturbations in parameter space. Thus, it is convenient to characterize (1.1) in terms of the parameterized dynamical system

$$\begin{aligned} F\colon \mathbb{R}^n \times [0,1]^m &\to \mathbb{R}^n \times [0,1]^m \\ (x,z) &\mapsto (f_z(x), z) \end{aligned} \tag{1.2}$$

Given $Z \subset [0,1]^m$, we denote the restriction of F to $\mathbb{R}^m \times Z$ by $F_Z\colon \mathbb{R}^n \times Z \to \mathbb{R}^n \times Z$.

Again for the sake of simplicity assume that $[0,1]^n \subset \mathbb{R}^n$ is the compact subset of phase space that contains the dynamics of interest. Only a finite number of computations can be performed to compute the database and thus we need combinatorial representations for $[0,1]^n$, $[0,1]^m$ and F. Since the goal is to obtain rigorous results, the approximation of F involves the use of interval arithmetic. For this reason it is natural to use cubical grids of the form

$$\mathcal{X} = \mathcal{X}^{(d_X)} := \left\{ \zeta_k = \xi_{k_1,\dots,k_n} := \prod_{i=1}^{n} \left[k_i 2^{-d_X}, (k_i+1)2^{-d_X}\right] \mid k_i = 0, \dots, 2^{d_X} - 1 \right\}$$

and

$$\mathcal{Z} = \mathcal{Z}^{(d_Z)} := \left\{ \zeta_k = \zeta_{k_1,\dots,k_m} := \prod_{i=1}^{m} \left[k_i 2^{-d_Z}, (k_i+1)2^{-d_Z}\right] \mid k_i = 0, \dots, 2^{d_Z} - 1 \right\}$$

Given a collection of grid elements $\mathcal{A} \subset \mathcal{X}$, the corresponding subset of $[0,1]^n$ is denoted by $|\mathcal{A}| \subset [0,1]^n$. The same notation is applied to subsets of $\mathcal{Z}$.

For every grid element $\zeta \in \mathcal{Z}$ let $\mathcal{F}_\zeta\colon \mathcal{X} \rightrightarrows \mathcal{X}$ denote a multivalued map that approximates $F_{|\zeta|}\colon X \times |\zeta| \to X \times |\zeta|$. In particular, given an element $\xi \in \mathcal{X}$ the multivalued map returns a set of grid elements $\mathcal{F}_\zeta(\xi) \subset \mathcal{X}$. The only requirement on $\mathcal{F}_\zeta$ is that it be an *outer approximation*, that is

$$F_{|\zeta|}(|\xi|) \subset \operatorname{int}_{[0,1]^n}\left(|\mathcal{F}_\zeta(\xi)|\right) \quad \text{for all } \xi \in \mathcal{X}. \tag{1.3}$$

The *minimal* outer approximation of $F_{|\zeta|}$ is given by

$$\widetilde{\mathcal{F}}_\zeta(\xi) := \left\{\xi' \mid \xi' \cap F_{|\zeta|}(|\xi|) \neq \emptyset\right\}$$

Notice that any other outer approximation $\mathcal{F}_\zeta$ of $F_{|\zeta|}$ satisfies $\widetilde{\mathcal{F}}_\zeta(\xi) \subset \mathcal{F}_\zeta(\xi)$ for all $\xi \in \mathcal{X}$ [2].

Observe that the size of the grid elements provides a bound on the accuracy with which the dynamics of f_z for each $z \in \zeta$ can be approximated. The strength of the database approach is that one is able to a priori choose the size of the grids, $\bar{d}_X$ and $\bar{d}_Z$, and then extract rigorous statements about the dynamics of $F_{|\zeta|}$ that are identifiable and robust with respect to those scales. Ideally, $\bar{d}_X$ and $\bar{d}_Z$ are chosen based on modeling assumptions. In general one expects that for models of well understood phenomena or physical systems for which accurate measurements can be made it is appropriate to choose finer grids. In contrast, given a crude or heuristic model or a system for which only coarse measurements can be made, one expects that the results associated with cruder grids will provide more meaningful information.

Clearly, for all but the lowest dimensional problems, the number of grid elements grows extremely rapidly as the number of divisions ($\bar{d}_X$ or $\bar{d}_Z$) is increased. Thus from a computational perspective computing directly with all of $\mathcal{X}$ or $\mathcal{Z}$ becomes prohibitively costly. As is described below, for each $\zeta \in \mathcal{Z}$ the computation of $\mathcal{F}_\zeta$ can be done in a multiscale manner and thus it is possible to carry out computations for problems with higher dimensional phase space [6]. The novelty of this thesis is that we introduce a multiscale approach for parameter space. To outline the challenges and how we address them requires a more detailed description of the database and the associated calculations. Such a description follows, but will be presented more formally in Chapter 2.

To efficiently extract information about the dynamics of $F_{|\zeta|}$ it is convenient to view $\mathcal{F}_\zeta$ as a directed graph with edge $\xi \to \xi'$ if and only if $\xi' \in \mathcal{F}_\zeta(\xi)$. Recall [5] that a *strongly connected component* of a directed graph is a maximal set of vertices $\mathcal{C}$ such that for every pair of elements $\xi_0, \xi_1 \in \mathcal{C}$, there are directed paths from ξ_0 to ξ_1 and from ξ_1 to ξ_0. In terms from graph theory, ξ_0 is *reachable* from ξ_1, and vice versa. It is important to note that by this definition, a single vertex ξ is a strongly connected component if it does not lie on a cycle. The *component graph* of $\mathcal{F}_\zeta$ is obtained by taking each strongly connected component $\{\mathcal{C}_i\}$ to be a vertex and declaring an edge from $\mathcal{C}_i \to \mathcal{C}_j$ if there exists an edge $\xi_i \to \xi_j$ in $\mathcal{F}_\zeta$ for some $\xi_i \in \mathcal{C}_i$ and $\xi_j \in \mathcal{C}_j$. Observe that the component graph of $\mathcal{F}_\zeta$ is a directed acyclic graph and hence a partially ordered set (poset).

Two elements $\xi_0, \xi_1 \in \mathcal{X}$ are defined to be in the same *strongly connected path component* of $\mathcal{F}_\zeta$ if there exist nontrivial directed paths in $\mathcal{F}_\zeta$ from ξ_0 to ξ_1 and ξ_1 to ξ_0. The set of strongly connected path component is a subset of the set of strongly connected component. Thus the set of strongly connected path components can be viewed as a poset that inherits its ordering from the component graph.

Definition 1.1. Given a directed graph $\mathcal{F}_\zeta$, the individual strongly connected path components are called *Morse sets.* Let P_ζ be an index set for the collection of all Morse sets and let $\leq_\zeta$ denote the partial order inherited from the component graph of $\mathcal{F}_\zeta$. The set

$$\{\mathcal{M}_\zeta(p) \subset \mathcal{X} \mid p \in (\mathsf{P}_\zeta, \leq_\zeta)\}$$

is called the *Morse decomposition* of $\mathcal{F}_\zeta$.

Recall that given a partially ordered set $(P, \leq)$, we say that q covers p if from the relation $q \leq r \leq p$ it follows that either $q = r$ or $r = p$.

Definition 1.2. Given a directed graph $\mathcal{F}_\zeta$, the *Morse graph* of $\mathcal{F}_\zeta$ is the directed acylic graph with nodes consisting of Morse sets, and directed edges $M_\zeta(p) \to M_\zeta(q)$ if and only if q covers p in $(P_\zeta, \leq_\zeta)$; it is denoted by $\mathsf{MG}(\mathcal{F}_\zeta)$.

The Morse graph $\mathsf{MG}(\mathcal{F}_\zeta)$ contains substantial information about the global dynamics of f_z for all $z \in |\zeta|$. In particular, as is shown in [8] the collection $M_z(p) := \mathrm{Inv}(|\mathcal{M}_\zeta(p)|, f_z)$, $p \in \mathsf{P}_\zeta$, forms a Morse decomposition, in the classical sense of Conley [4], of $\mathrm{Inv}([0,1]^n, f_z)$ under the dynamics of f_z with admissible order $\leq_\zeta$. As is indicated in [2, 3], the Conley index of the Morse sets $M_z(p)$ can be computed using $\mathcal{F}_\zeta$ under the assumption that $\mathcal{F}_\zeta$ has acyclic values; that is, if the homology of $\mathcal{F}_\zeta(\xi)$ is isomorphic to the homology of a single-point space for each ξ. This is important since the Conley index provides information about the structure of the dynamics associated with the invariant sets $M_z(p)$. For now, it is enough to say that the Conley index gives provides information about the existence, structure and stability of the Morse sets; we postpone defining the Conley index until Chapter 2.

Definition 1.3. The *Conley-Morse graph* of $\mathcal{F}_\zeta$ consists of the Morse graph $\mathsf{MG}(\mathcal{F}_\zeta)$ and the Conley index information for each Morse set. It is denoted by $\mathsf{CMG}(\mathcal{F}_\zeta)$.

In the context of the current implementation of the database $\mathsf{CMG}(\mathcal{F}_\zeta)$ provides the complete description of the dynamics for the set of parameter values $|\zeta|$.

To return to the issue of performing these computations in a multiscale manner let

$$\widetilde{\mathcal{F}}_\zeta^{(d_X,d_Z)} \colon \mathcal{X}^{(d_X)} \rightrightarrows \mathcal{X}^{(d_X)}$$

denote the minimal outer approximation working with grids $\mathcal{X}^{(d_X)}$ and $\mathcal{Z}^{(d_Z)}$, i.e. $\zeta \in \mathcal{Z}^{(d_Z)}$. Assume $d_X < d'_X$. If $\xi \in \mathcal{X}^{(d_X)}$, then there exists $\xi' \in \mathcal{X}^{(d'_X)}$ such that

$|\xi'| \subset |\xi|$. Furthermore,

$$\left|\widetilde{\mathcal{F}}_\zeta^{(d'_X,d_Z)}(\xi')\right| \subset \left|\widetilde{\mathcal{F}}_\zeta^{(d_X,d_Z)}(\xi)\right|.$$

This observation allows one to conclude [8] that if $\mathcal{M}_\zeta^{(d'_X)}(p)$ is a Morse set for $\widetilde{\mathcal{F}}_\zeta^{(d'_X,d_Z)}$, then there exists a Morse set $\mathcal{M}_\zeta^{(d_X)}(q)$ for $\widetilde{\mathcal{F}}_\zeta^{(d_X,d_Z)}$ such that

$$\left|\mathcal{M}_\zeta^{(d'_X)}(p)\right| \subset \left|\mathcal{M}_\zeta^{(d_X)}(q)\right|.$$

This implies that to identify Morse sets at a given level of resolution d_X it is sufficient to restrict ones attention to Morse sets at level $d_X - 1$. To be more precise, let $\left\{\mathcal{M}_\zeta^{(d_X-1)}(p) \mid p \in \mathsf{P}_\zeta\right\}$ be the Morse sets of $\widetilde{\mathcal{F}}_\zeta^{(d_X-1,d_Z)}$. To determine the Morse sets of $\widetilde{\mathcal{F}}_\zeta^{(d_X,d_Z)}$ it is sufficient to compute the Morse sets of

$$\widetilde{\mathcal{F}}_\zeta^{(d_X,d_Z)} \colon \mathcal{Y}(p) \rightrightarrows \mathcal{Y}(p) \tag{1.4}$$

for each $p \in \mathsf{P}_\zeta$ where $\mathcal{Y}(p) := \left\{\xi \in \mathcal{X}^{(d_X)} \mid |\xi| \subset \left|\mathcal{M}_\zeta^{(d_X-1)}(p)\right|\right\}$.

Turning to the question of computational cost, given $\mathcal{F}_\zeta$ there exist linear time algorithms for identifying the Morse sets [5]. Furthermore, at least asymptotically, the cost of computing outer approximations on the Morse sets depends to a large extent on the dimension of the Morse sets for the underlying dynamics $F_{|\zeta|}$ as opposed to n, the dimension of the phase space. Thus, in principle the computation of $\mathcal{F}_\zeta$ is feasible for models in which the recurrent dynamics takes the form of fixed points, periodic orbits, invariant circles or low dimensional chaotic sets even if the ambient dimension of the phase space is high.

It should be noted that the argument outlined above depends on using the minimal outer approximation, which is rarely obtained in practice. For the general case a more subtle approach is necessary [3].

The focus of this thesis is on developing a multiscale approach for parameter space. The key step in realizing computational savings with respect to phase space is the ability to restrict computations as indicated by (1.4). Because we are interested in the dynamics at all parameter values this type of restriction is not possible. Thus we adopt a much simpler strategy: (i) compute $\mathcal{F}_\zeta^{(\bar{d}_X,\bar{d}_Z)}$ at various values of $\zeta \in \mathcal{Z}^{(\bar{d}_Z)}$; (ii) use this information to determine larger regions $\mathcal{S} \subset \mathcal{Z}^{(\bar{d}_Z)}$ where the dynamics appears to be only slowly changing; (iii) compute the Conley-Morse graphs associated with $\mathcal{F}_\mathcal{S} \colon \mathcal{X} \rightrightarrows \mathcal{X}$; and (iv) decide if the information captured by $\mathcal{F}_\mathcal{S}$ is sufficient or if further refinement in parameter space is necessary.

There are several issues the need to be resolved to carry out this procedure. The first is how to compare Conley-Morse graphs computed at different grid elements ζ and ζ' or different levels of resolution. This is discussed in Chapter 3; in Chapter 3 we also provide a formal definition of the database, which we appropriately entitle the 'Conley-Morse database'. The second is how to choose test values of ζ and the regions $\mathcal{S}$. We refer to our approach as the 'strip method' as we consider 1-dimensional 'strips' of grid elements in $\mathcal{Z}^{(\bar{d}_Z)}$ and compute the Conley-Morse graphs at the grid elements which define the endpoints. If the Conley-Morse graphs agree at the endpoints we check to see if the Morse graph is valid over the entire strip. If not we subdivide the strip and repeat the process. This is described in Chapter 4. We prove that this approach still provides mathematically rigorous descriptions of the dynamics. However, it is possible that the dynamics described by $\mathcal{F}_\mathcal{S}$ contains less information that that available from the Conley-Morse information associated with $\mathcal{F}_\zeta^{(\bar{d}_X,\bar{d}_Z)}$. On this level of generality we know of no theorems that describe how much information is lost. Thus in Chapter 5 we present numerical results that provide some quantification of how much information is lost and how much computational gain can be achieved.

CHAPTER 2

MATHEMATICAL FRAMEWORK

1. Overview

In broad strokes, the Conley-Morse database is an attempt to extract rigorous description of global dynamics over large regions of parameter space by representing dynamics in terms of a combinatorial graph and analyzing these graphs with computational topology. In the context of what was presented in the previous Chapter, a description of global dynamics is given by Conley's classical notion of a Morse decomposition, and its combinatorial analogue is the Morse graph.

In fact, much of the foundation of this approach comes from preliminary ideas of C. Conley, and the core mathematical objects of the theory are very robust with respect to perturbations of the system. In the past decade, much work has gone into translating Conley's ideas into an algorithmic approach. In the simplest setting, the approach distills to three core components: *isolation*, *decomposition*, and *reconstruction via the Conley index*.

2. Isolation

Consider the setting of a multiparameter family of dynamical systems given by a continuous function $f : X \times \Lambda \to X$. For the purpose of this review we will work more generically than in $\mathbb{R}^n$; we require the phase space X is a locally compact metric space and the parameter space Λ to be a compact, locally contractible, connected metric space. For $\lambda \in \Lambda$ we have an associated dynamical system $f_\lambda : X \to X$. A compact set $N \subset X$ is an *isolating neighborhood* for f_λ if $\mathrm{Inv}(N, f_\lambda)$, the maximal invariant set in N under f_λ, is contained in the interior of N. Much of the focus on isolating neighborhoods comes from two reasons: first, if N is an isolating neighborhood for f_{λ_0}, then N is an isolating neighborhood for f_λ for all λ sufficiently close to λ_0; second, there are efficient algorithms for computing isolating neighborhoods.

To perform the computations we discretize the phase space X using the concept of a *grid*. This consists of a finite collection $\mathcal{X}$ of nonempty, compact subsets of X respectively, with the following properties:

I. $X = \bigcup_{\xi \in \mathcal{X}} \xi$

II. $\xi = \mathrm{cl}(\mathrm{int}(\xi))$ for all $\xi \in \mathcal{X}$

III. $\xi \cap \mathrm{int}(\xi') = \emptyset$ for all $\xi \neq \xi'$

To discretize the dynamics, we make us of a *combinatorial multivalued map* $\mathcal{F} : \mathcal{X} \rightrightarrows \mathcal{X}$. As above, $\mathcal{F}$ is an *outer approximation* of f if $f(|\xi|) \subset \mathrm{int}(|\mathcal{F}(\xi)|)$ for each grid element $\xi \in \mathcal{X}$. To compute over sets of parameter values, we impose a grid $\mathcal{Z}$ on Λ, which upholds the above properties. For $\zeta \in \mathcal{Z}$ construct $\mathcal{F}_\zeta$ such that $F_{|\zeta|}(|\xi|) \subset \mathrm{int}(|\mathcal{F}_\zeta(\xi)|)$. As introduced above, $\widetilde{\mathcal{F}}(\xi) := \{\xi' \in \mathcal{X} \mid \xi' \cap F(|\xi|)\}$ is the minimal outer approximation in the sense that given any other outer approximation $\mathcal{F}$, $\widetilde{\mathcal{F}}(\xi) \subset \mathcal{F}(\xi)$ for all $\xi \in \mathcal{X}$. In this thesis, $\widetilde{\mathcal{F}}$ will be useful for theoretical purposes, however it cannot typically be computed in practice.

3. Decomposition

Conley introduced the notion of a *Morse decomposition* of an isolated invariant set; a concept providing a method of decomposing the dynamics into gradient-like (strictly non-recurrent) and recurrent dynamics. For isolated invariant set S_z under f_z, this is a finite collection of disjoint isolated invariant sets $M_z(p) \subset S_z$ called *Morse sets* which are indexed by poset $(\mathsf{P}_z, \leq_z)$ such that for every $x \in S_z \setminus \bigcup_{p \in \mathsf{P}_z} M_z(p)$ and any complete orbit γ of f_λ through x in S_z there exists indices $p >_z q$ such that $\omega(\gamma) \subset M_\lambda(q)$ and $\alpha(\gamma) \subset M_\lambda(p)$. Thus S_z is composed of the Morse sets and the connecting orbits between them.

Since Morse sets are isolated invariant sets, they can be characterized in terms of isolating neighborhoods, and the partial order can be determined by approximating orbits between isolating neighborhoods. Further, there are linear time algorithms for determining isolating neighborhoods for Morse decompositions from a multivalued map $\mathcal{F}$. This is done by viewing $\mathcal{F}$ as a combinatorial graph, then the strongly connected components of the graph are the Morse sets. Further, one can then collapse each component to one vertex, then collapse all the edges between these new vertices to one edge to obtain the component graph, as described in Chapter 1. This gives a coarse, yet compact and efficiently computable representation of the global dynamics of the system - the transitive reduction of which is what we term a Morse Graph, $\mathsf{MG}(\mathcal{F})$.

4. Reconstruction via the Conley Index

The final ingredient in this approach is the *Conley index*, an algebraic topological invariant of isolated invariant sets.

Consider an arbitrary continuous map $g : Z \to Z$ on a locally compact metric space Z. Let (N, L) be a pair of compact subsets of Z where $L \subset N$, and let $P = (N, L)$. Consider the pointed space $(N/L, [L])$ obtained by collapsing L to a single point $[L]$, where $[L]$ denotes the equivalence class $x \simeq y$ if and only if $x, y \in L$.

Define $g_P : (N/L, [L]) \to (N/L, [L])$ by

$$g_P(x) = \begin{cases} g(x), & \text{if } x, g(x) \in N\backslash L \\ [L], & \text{otherwise} \end{cases}$$

The pair $P = (N, L)$ forms an *index pair* if the map g_P is continuous and $\mathrm{cl}(N\backslash L)$ is an isolating neighborhood under g.

Before defining the Conley index, it is important to note that for an isolated invariant set S_z, there exists at least one index pair such that $S_z = \mathrm{Inv}(\mathrm{cl}(N\backslash L), g)$. In terms of any Morse set $\mathcal{M}_z(p)$, this implies there exists an index pair $P = (N, L)$ where $F_{z,P}$ is a continuous map and $\mathcal{M}_\zeta(p) = \mathrm{Inv}(\mathrm{cl}(N\backslash L)$. This function in turn induces a homomorphism on homology. Furthermore, there may be many index pairs which isolate an isolated invariant set. Thus different choices of index pairs may lead to potentially different homomorphisms. The Conley index of an isolated invariant set must then be defined as an equivalence class of these homomorphisms. In particular, the shift equivalence class of the induced map on homology:

$$F_{z,P*} : H_*(N/L, [L]) \to H_*(N/L, [L])$$

represents the *Conley index of the isolated invariant set* S_z.

Given an outer approximation $\mathcal{F}$ there are fast algorithms for determining index pairs of Morse sets in terms of the grid $\mathcal{X}$. Furthermore, if $\mathcal{F}$ takes acylic values then there exists algorithms to compute $F_{z,P*}$.

The Conley index provides information about the existence, structure and stability of the maximal invariant set. For instance, the most basic result is that if $F_{z,P*}$ is not nilpotent, then $\mathrm{Inv}(N\backslash L, f_z) \neq \emptyset$. This description is coarse, yet it is also robust, which provides one method to overcome the problem of bifurcations occurring on all scales of parameter space. The following are three of the most important observations of the Conley index [3]:

1. One can associate a Conley index to any isolating neighborhood

2. If N and N' are isolating neighborhoods and $\mathrm{Inv}(N, f_z) = \mathrm{Inv}(N', f_z)$, then they have the same Conley index

3. If N is an isolating neighborhood for all z in a path connected subset of Z, then the Conley index associated with N is the same for all f_z.

For each grid element $\zeta \in \mathcal{Z}$, the data obtained from the computation of the Morse decomposition and the Conley indices of the Morse sets is represented via a *Conley-Morse graph*, denoted $\mathsf{CMG}(\mathcal{F}_\zeta)$. As described in Chapter 1, this is a directed acylic graph obtained from the poset $(\mathsf{P}_\zeta, \leq_\zeta)$, where the vertices represent the Morse sets, and attached to each vertex is information regarding the Conley Index of the associated Morse set. As there are only a finite number of parameter space grid elements, this provides a finite combinatorial representation of the global dynamics over some bounded parameter region.

CHAPTER 3

COMPARING DYNAMICS

1. Classifying Dynamics Over Parameter Space

As is indicated in the Chapter 1, the general strategy for computation is to choose a grid element $\zeta \in \mathcal{Z}$ or more generally a set of grid elements $\mathcal{S} \subset \mathcal{Z}$; construct a multivalued outer approximation $\mathcal{F}_{\mathcal{S}}$ of $F_{|\mathcal{S}|}$; and then determine the associated Conley-Morse graphs $\mathsf{CMG}(\mathcal{F}_{\mathcal{S}})$. Since one of the goals of the database is to give a representation of the dynamics over all parameter values, we need to be able to compare the dynamics over different sets of parameter values. To do this we make use of the following combinatorial object.

Definition 1.1. Consider $\mathcal{S}_0, \mathcal{S}_1 \subset \mathcal{Z}$ with outer approximations $\mathcal{F}_{\mathcal{S}_0}$ and $\mathcal{F}_{\mathcal{S}_1}$. The *clutching graph* $\mathcal{J}(\mathcal{S}_0, \mathcal{S}_1)$ is the bipartite graph with vertices $\mathsf{P}_{\mathcal{S}_0} \cup \mathsf{P}_{\mathcal{S}_1}$ and edges

$$\mathsf{E}_{\mathcal{S}_1,\mathcal{S}_0} := \{(p,q) \in \mathsf{P}_{\mathcal{S}_0} \times \mathsf{P}_{\mathcal{S}_1} \mid \mathcal{M}_{\mathcal{S}_0}(p) \cap \mathcal{M}_{\mathcal{S}_1}(q) \neq \emptyset\}.$$

If every vertex in $P_{\mathcal{S}_0}$ has at most one edge in $\mathsf{E}_{\mathcal{S}_0,\mathcal{S}_1}$, then the *clutching morphism* $\iota_{\mathcal{S}_1,\mathcal{S}_0} : \mathsf{P}_{\mathcal{S}_0} \to \mathsf{P}_{\mathcal{S}_1}$ is defined by

$$\iota_{\mathcal{S}_1,\mathcal{S}_0}(p) := \begin{cases} q & \text{if } (p,q) \in \mathsf{E}_{\mathcal{S}_0,\mathcal{S}_1}, \\ \emptyset & \text{otherwise.} \end{cases}$$

If every vertex in $P_{\mathcal{S}_0}$ has exactly one edge in $\mathsf{E}_{\mathcal{S}_0,\mathcal{S}_1}$, then the clutching morphism is called a *clutching function*.

The following result is fundamental and follows from the continuation of Morse decompositions and the Conley index [3].

Proposition 1.2. *Consider* $\mathcal{S}_0, \mathcal{S}_1 \in \mathcal{Z}$ *such that* $\mathcal{S}_0 \cap \mathcal{S}_1 \neq \emptyset$*. If* $\iota_{\mathcal{S}_1,\mathcal{S}_0}$ *defines a directed graph isomorphism*

$$\iota_{\mathcal{S}_1,\mathcal{S}_0} : \mathsf{MG}(\mathcal{F}_{\mathcal{S}_0}) \to \mathsf{MG}(\mathcal{F}_{\mathcal{S}_1})$$

then the Conley index of $\mathcal{M}_{\mathcal{S}_0}(p)$ *is equivalent to the Conley index of* $\mathcal{M}_{\mathcal{S}_1}(\iota_{\mathcal{S}_1,\mathcal{S}_0}(p))$

Proposition 1.2 motivates the following definition.

Definition 1.3. Two Morse graphs $\mathsf{MG}(\mathcal{F}_{\zeta_0})$ and $\mathsf{MG}(\mathcal{F}_{\zeta_1})$ are *equivalent* if

1. $\zeta_0 \cap \zeta_1 \neq \emptyset$,
2. ι_{ζ_0,ζ_1} is a clutching function and induces a directed graph isomorphism

$$\iota_{\zeta_0,\zeta_1} : \mathsf{MG}(\mathcal{F}_{\zeta_1}) \to \mathsf{MG}(\mathcal{F}_{\zeta_0}).$$

The equivalence classes of $\{\mathsf{MG}(\mathcal{F}_\zeta) \mid \zeta \in \mathcal{Z}\}$ with respect to the transitive closure of this relation are called *continuation classes.* The set of continuation classes is denoted by $\{\mathcal{CC}_j(\mathcal{X}, \mathcal{Z}, \mathcal{F}) \mid j = 1, \ldots, J\}$.

Definition 1.4. Two Conley-Morse graphs $\mathsf{CMG}(\mathcal{F}_{\zeta_0})$ and $\mathsf{CMG}(\mathcal{F}_{\zeta_1})$ are *equivalent* if

1. $\mathsf{MG}(\mathcal{F}_{\zeta_0})$ and $\mathsf{MG}(\mathcal{F}_{\zeta_1})$ are equivalent,
2. The Conley index of $\mathcal{M}_{\zeta_0}(p)$ and $\mathcal{M}_{\zeta_1}(\iota_{\zeta_1,\zeta_0}(p))$ are the same for each $p \in \mathcal{P}_{\zeta_0}$

This allows us to give a formal definition of the database.

Definition 1.5. Consider a parameterized dynamical system (1.2), grids $\mathcal{X}$ and $\mathcal{Z}$ and a choice of multivalued outer approximations $\mathcal{F}_\zeta$ for $F_{|\zeta|}$. The associated *database* $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$ is a collection of continuation classes $\{\mathcal{CC}_j \mid j = 1, \ldots, J\}$, where to each class there is associated a set of grid elements $\pi(\mathcal{CC}_j) \subset \mathcal{Z}$ and a representative Conley-Morse graph CMG_j. Furthermore, it is required that $\{\pi(\mathcal{CC}_j) \mid j = 0, \ldots, J\}$ forms a partition of $\mathcal{Z}$.

In the context of this thesis $\mathsf{DB}(\mathcal{X}^{(\bar{d}_X)}, \mathcal{Z}^{(\bar{d}_Z)}, \widetilde{\mathcal{F}})$ (where DB is constructed using the minimal multivalued map) provides the maximal amount of information about the underlying parameterized dynamical system (1.2) at the resolution $\bar{d}_X, \bar{d}_Z$. Thus any discussion concerning the loss of information due to a particular computational scheme must be measured in terms of differences from $\mathsf{DB}(\mathcal{X}^{(\bar{d}_X)}, \mathcal{Z}^{(\bar{d}_Z)}, \widetilde{\mathcal{F}})$.

To quantify these differences a few simple concepts from the theory of posets are useful.

Definition 1.6. Let $(\mathsf{P}, \leq)$ be a poset. A subset $C \subset \mathsf{P}$ is *convex* if for every $p, q \in C$ and $r \in \mathsf{P}$ such that $p \leq r \leq q$, then $r \in C$.

Given an outer approximation $\mathcal{F}_\zeta$, its Morse decomposition $\mathcal{M}_\zeta$ and associated partial order $(\mathsf{P}_\zeta, \leq_\zeta)$ and $p \in \mathsf{P}_\zeta$, define

$$\uparrow (\mathcal{M}_\zeta(p)) := \{\xi \in \mathcal{F}_\zeta \mid \mathcal{M}_\zeta(p) \text{ is reachable from } \xi\}$$
$$\downarrow (\mathcal{M}_\zeta(p)) := \{\xi \in \mathcal{F}_\zeta \mid \xi \text{ is reachable from } \mathcal{M}_\zeta(p)\}$$

Notice $\mathcal{M}_\zeta(p) = \uparrow (\mathcal{M}_\zeta(p)) \cap \downarrow (\mathcal{M}_\zeta(p))$. We may now formalize the meaning of a convexity in a Morse decomposition $\mathcal{M}_\zeta$ of outer approximation $\mathcal{F}_\zeta$.

Definition 1.7. If $C \subset \mathsf{P}_\zeta$ is convex, then

$$\mathcal{M}_\zeta(C) := \bigcup_{p,q\in C} \downarrow (\mathcal{M}_\zeta(p)) \cap \uparrow (\mathcal{M}_\zeta(q))$$

Definition 1.8. Given combinatorial Morse decompositions, $\mathcal{M}_\zeta$ and $\mathcal{M}'_\zeta$, with associated partial orders $(\mathsf{P}_\zeta, \leq_\zeta)$ and $(\mathsf{P}'_\zeta, \leq'_\zeta)$, a function $\phi : (\mathsf{P}_\zeta, \leq_\zeta) \to (\mathsf{P}'_\zeta, \leq'_\zeta)$ is a *coarsening* if

I. ϕ is order-preserving, i.e. $p \geq q$, $p, q \in \mathsf{P}_\zeta \implies \phi(p) \geq \phi(q)$

II. For any convex subset $C' \subset \mathsf{P}'_\zeta$, the Conley index of $\mathcal{M}'_\zeta(C')$ is equivalent to the Conley index of $\mathcal{M}_\zeta(\phi^{-1}(C'))$.

Definition 1.9. Consider combinatorial Morse decompositions, $\mathcal{M}_\zeta$ and $\mathcal{M}'_\zeta$, with associated partial orders $(\mathsf{P}_\zeta, \leq_\zeta), (\mathsf{P}'_\zeta, \leq'_\zeta)$. If there exists a coarsening $\phi : \mathsf{P}_\zeta \to \mathsf{Q}_\zeta$, then the set $\{\mathcal{M}_\zeta(\phi(p)) \subset \mathcal{X} \mid p \in \mathsf{P}_\zeta\}$ is a *coarsened Morse Decomposition* and the associated Morse graph is a *coarsened Morse graph.*

Theorem 1.10. *Consider a parameterized dynamical system (1.2) and choice of grids $\mathcal{X}$ and $\mathcal{Z}$ for phase and parameter space. Fix $\zeta \in \mathcal{Z}$. Let $\mathcal{F}_\zeta$ and $\mathcal{F}'_\zeta$ be two choices of multivalued outer approximations, with $\mathcal{F}_\zeta(\xi) \subset \mathcal{F}'_\zeta(\xi)$ for each $\xi \in \mathcal{X}$. Then there exists a coarsening $\phi_\zeta : \mathsf{P}_\zeta \to \mathsf{P}'_\zeta$.*

Proof. Note that as $\mathcal{F}_\zeta \subset \mathcal{F}'_\zeta$, $\mathcal{M}_\zeta(p) \subset \mathcal{F}'_\zeta$ for each $p \in \mathsf{P}_\zeta$. Let $\phi : \mathsf{P}_\zeta \to \mathsf{P}'_\zeta$ be the map induced by this inclusion of strongly connected path components in $\mathcal{F}$. We will show that ϕ_ζ has the properties of a coarsening:

I. Let $p \geq q$, $p, q \in \mathsf{P}_\zeta$. Then there is a path $\xi_0, \ldots, \xi_n$ in $\mathcal{F}_\zeta$ with $\xi_0 \in \mathcal{M}_\zeta(p)$ and $\xi_n \in \mathcal{M}_\zeta(q)$, $\xi_{i+1} \in \mathcal{F}_\zeta(\xi_i)$. As $\mathcal{F}_\zeta \subset \mathcal{F}'_\zeta$ this path is present in $\mathcal{F}'_\zeta$ and as

ϕ is constructed by inclusion, $\xi_0 \in \mathcal{M}_\zeta(\phi(p))$ and $\xi_n \in \mathcal{M}_\zeta(\phi(q))$. Therefore, $\phi(p) \geq \phi(q)$.

II. Let $C' \subset \mathsf{P}'_\zeta$ be convex and consider $z \in |\zeta|$. Define C as the convex preimage, $C := \phi^{-1}(C')$. As $|\mathcal{M}'_\zeta(C')|, |\mathcal{M}_\zeta(C')|$ are isolating neighborhoods, we will show $\mathrm{Inv}(|\mathcal{M}'_\zeta(C')|, f_z) = \mathrm{Inv}(|\mathcal{M}_\zeta(C)|, f_z)$, and invoke the second statement regarding the Conley Index from Chapter 2, giving that Conley Index of $\mathcal{M}'_\zeta(C')$ is equivalent to that of $\mathcal{M}_\zeta(C)$. Accordingly, define $S'_z := \mathrm{Inv}(|\mathcal{M}'_\zeta(C')|, f_z)$ and $S_z := \mathrm{Inv}(|\mathcal{M}_\zeta(C)|, f_z)$. We will show $S_z = S'_z$.

Notice that $\mathcal{M}_\zeta(C) \subset \mathcal{M}'_\zeta(C')$ as ϕ is defined by inclusion, thus $S_z \subset S'_z$.

Now let $x \in S'_z$. We will show $x \in S_z$. As $x \in S'_z$ there exists a full trajectory (x_k) through x such that $\{x_k \mid k \in \mathbb{Z}\} \subset |\mathcal{M}'_\zeta(C')|$, i.e. a sequence satisfying $x_0 = x$ and $x_{n+1} = f(x_n)$ for all $n \in \mathbb{Z}$, captured entirely in $|\mathcal{M}'_\zeta(C')|$. Let $\xi_k \in \mathcal{M}'_\zeta(C')$ be grid elements containing the respective x_k. Since $\mathcal{F}_\zeta$ and $\mathcal{F}'_\zeta$ are outer approximations, $\xi_{i+1} \in \mathcal{F}_\zeta(\xi_i)$ and $\xi_{i+1} \in \mathcal{F}'_\zeta(\xi_i)$ for all $i \in \mathbb{Z}$. As $(\xi_k) \subset \mathcal{M}'_\zeta(C')$, and by the pigeonhole principle, there must exist $p', q' \in C'$ $p' > q'$ such that for all $n \geq N'$ $\xi_{-n} \in \mathcal{M}'(p')$ and $\xi_n \in \mathcal{M}'(q')$. Similarly, there must exist $p, q \in \mathsf{P}_\zeta$, $p > q$, such that for all $n \geq N$ $\xi_{-n} \in \mathcal{M}_\zeta(p)$ and $\xi_n \in \mathcal{M}_\zeta(q)$. Thus $\mathcal{M}'_\zeta(p') \cap \mathcal{M}_\zeta(p) \neq \emptyset$, and as $\mathcal{F}_\zeta \subset \mathcal{F}'_\zeta$ this implies that $\mathcal{M}'_\zeta(p) \subset \mathcal{M}_\zeta(p)$ and thus $\phi(p) = p'$. By a similar argument, $\phi(q) = q'$. Hence $p, q \in C = \phi^{-1}(C')$, and thus $(\xi_k) \in \downarrow (\mathcal{M}_\zeta(p)) \cap \uparrow (\mathcal{M}_\zeta(q)) \subset \mathcal{M}_\zeta(C)$. Therefore for the full trajectory (x_k), we have $(x_k) \in |\mathcal{M}_\zeta(C)|$, so $x \in \mathrm{Inv}(|M_\zeta(C)|, f_z) = S_z$.

□

We generalize the notion of coarsening a Morse graph to that of coarsening an entire database with the following definition.

Definition 1.11. Fix a parameterized dynamical system (1.2) and choice of grids $\mathcal{X}$ and $\mathcal{Z}$ for phase and parameter space. Let $\mathcal{F}$ and $\mathcal{F}'$ be two different choices of multivalued outer approximations. Let $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$ and $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}')$ be the resulting databases. A *database coarsening*

$$\Phi\colon \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}) \to \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}').$$

consists of a collection of functions $\{\phi_\zeta\}$ such that each ϕ_ζ is a coarsening.

Corollary 1.12 then follows naturally from Theorem 1.10.

Corollary 1.12. *Consider a parameterized dynamical system (1.2) and choice of grids $\mathcal{X}$ and $\mathcal{Z}$ for phase and parameter space. Let $\mathcal{F}_\zeta$ and $\mathcal{F}'_\zeta$ be two choices of multivalued outer approximations, with $\mathcal{F}_\zeta(\xi) \subset \mathcal{F}'_\zeta(\xi)$ for each $\xi \in \mathcal{X}, \zeta \in \mathcal{Z}$. Let $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$ and $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}')$ be the resulting databases.*

Then there exists a database coarsening

$$\Phi : \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}) \to \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}')$$

Furthermore, we can apply Corollary 1.12 to the minimal-multivalued map $\widetilde{\mathcal{F}}$ to achieve the following result.

Corollary 1.13. *Consider a parameterized dynamical system (1.2) and choice of grids $\mathcal{X}$ and $\mathcal{Z}$ for phase and parameter space. Let $\mathcal{F}$ be a multivalued outer approximations. Let $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \widetilde{\mathcal{F}})$ and $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$ be the resulting databases.*

Then there exists a database coarsening

$$\Phi : \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \widetilde{\mathcal{F}}) \to \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$$

That is, Corollary 1.13 shows that every database is a coarsening of the finest database.

CHAPTER 4

COMPUTING DYNAMICS

1. Strip Method

We turn to the question of how to improve the computational cost associated with building $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$.

To simplify the presentation we restrict our attention to rectangular regions $Z = [0,1]^m \subset \mathbb{R}^m$ in parameter space. Fix $d_z \in \mathbb{N}$. We consider the cubical grid

$$\mathcal{Z} = \left\{ \zeta_k = \zeta_{k_1,\dots,k_m} := \prod_{i=1}^{m} \left[k_i 2^{-d_z}, k_i(2^{-d_z}+1) \right] \mid k_i = 0, \dots, 2^{d_z}+1 \right\}$$

An l-strip in $\mathcal{Z}$ is a collection of grid elements of the form

$$\left\{ \zeta_{k_1,\dots,k_m} \mid k_i = \bar{k}_i \text{ for } i \notin \{j_1,\dots,j_l\},\ k_{j_r} \in \{m,\dots,m+N(r)\} \right\} \subset \mathcal{Z}.$$

In words, the concept of an l-strip provides a method of selecting a contiguous set of grid elements with fixed indices in all but l dimensions. We will work with 1-strips, which, without loss of generality, are subsets of grid elements of the form

$$\mathcal{S}(K, J, \bar{k}) := \left\{ \sigma_j \in \mathcal{Z} \mid \sigma_j = \zeta_{K+j,\bar{k}_2,\dots,\bar{k}_m},\ j = 0,\dots,J \right\}$$

where $\bar{k} = (\bar{k}_2, \dots, \bar{k}_m)$. Let $\bar{\mathcal{K}} := \left\{ \bar{k} \mid \bar{k}_i = 0, \dots, 2^d - 1 \right\}$.

2. Local Bisection

Consider the set of grid elements

$$\mathcal{E} := \{\zeta_{k_0,\dots,k_m} \mid k_0 \in \{0, 2^{d_z}\}\}$$

We call $\mathcal{E}$ the set of *endpoints.* The associated 1-strips are:

$$\mathcal{S}(\mathcal{E}) = \bigcup_{\bar{k}\in\bar{\mathcal{K}}} \mathcal{S}(1, 2^d - 1, \bar{k})$$

Notice that the collection $\mathcal{S}(\mathcal{E}) \bigcup \mathcal{E}$ defines both a partition of $\mathcal{Z}$ and a grid over Z. Further, it lends itself to a method of subdiving the parameter space in regions where interesting dynamics may occur. It can be seen that for each strip $\mathcal{S} \in \mathcal{S}(\mathcal{E})$, there exist $\zeta_0, \zeta_1 \in \mathcal{E}$ such that both endpoints share a $m-1$ codimensional face with $\mathcal{S}$. Accordingly, we call these the endpoints of $\mathcal{S}$.

Instead of computing at individual grid elements, one may compute valid Morse graphs over large regions of the parameter space, i.e. computing $\mathsf{MG}(\mathcal{F}_\mathcal{S})$ for $\mathcal{S} \in \mathcal{S}(\mathcal{E})$. For each $\zeta \in \mathcal{S}$, $\mathsf{MG}(\mathcal{F}_\mathcal{S})$ may be thought of as a valid Morse graph over ζ. Intuitively, one would like to do this where the dynamics in $\mathcal{S}$ are not changing significantly. This approach of computing only one Morse graph would then provide a more efficient computation.

We take advantage of this specific grid through a local bisection algorithm, where each $\mathcal{S}$ is bisected depending upon the change in dynamics within $\mathcal{S}$. To formalize a method of capturing such a change in dynamics consider a 1-strip $\mathcal{S}$ and associated endpoints ζ_L, ζ_R. If both $\mathcal{J}(\mathcal{S}, \zeta_L)$ and $\mathcal{J}(\mathcal{S}, \zeta_R)$ define directed graph isomorphisms, we conclude that $\mathsf{MG}(\mathcal{F}_\mathcal{S}), \mathsf{MG}(\mathcal{F}_{\zeta_L}), \mathsf{MG}(\mathcal{F}_{\zeta_R})$ all belong to the same continuation class, and do not bisect $\mathcal{S}$. A specific implementation for the local bisection algorithm is given in Algorithm 1.

Algorithm 1 Local bisection algorithm for 1-strips in $\mathcal{Z}$

```
function LOCAL-BISECTION(DB(𝒳, 𝒵, ℱ), 𝒮(K, J, k̄))
    σ_K ← ζ_{K,k̄_2,...,k̄_m}
    σ_J ← ζ_{J,k̄_2,...,k̄_m}
    if J = K + 1 then
        DB(𝒳, 𝒵, ℱ) ← 𝒥(σ_K, σ_J)
        return
    end if
    if HEURISTIC-CHECK(MG(ℱ_{σ_K}),MG(ℱ_{σ_J})) then        ▷ Optional heuristic
        if 𝒥(σ_K, 𝒮) and 𝒥(σ_J, 𝒮) induce directed graph isomorphisms then
            for each ζ ∈ 𝒮 do
                MG(ℱ_ζ) ← MG(𝒮)
                DB(𝒳, 𝒵, ℱ) ← MG(ℱ_ζ)
            end for
            return
        end if
    end if
    DB(𝒳, 𝒵, ℱ) ← MG(ℱ_{σ_K}), MG(ℱ_{σ_J})
    L ← (J + K)/2
    LOCAL-BISECTION(DB(𝒳, 𝒵, ℱ), 𝒮(K, L, k̄))
    LOCAL-BISECTION(DB(𝒳, 𝒵, ℱ), 𝒮(L, J, k̄))
end function
```

For our implementation, HEURISTIC-CHECK determines whether the Morse graphs pass an isomorphism heuristic, which is a necessary condition to the clutching function defining a directed graph isomorphism. In other words, this optional heuristic prevents the computation of an unneccesary Morse graph over an entire 1-strip. Though we decided to test isomorphism of the Morse graphs, a number of other subroutines would suffice (for instance, bisecting if the clutching function itself did not define an isomorphism). Further, graph isomorphism is a notoriously difficult problem, though not known to be NP-Complete [7], which led us to utilizing a heuristic. We provide the details of the algorithm used in our implementation in Chapter 5.

3. Analysis of Strip Method

In order to relate original and strip approximations, we first assume the following condition on monotonicity in the parameter space:

A1. Consider strips $\mathcal{S}_1, \mathcal{S}_2 \subset \mathcal{Z}$, if $\mathcal{S}_1 \subset \mathcal{S}_2$, then for $\xi \in \mathcal{X}$, $\mathcal{F}_{\mathcal{S}_1}(\xi) \subset \mathcal{F}_{\mathcal{S}_2}(\xi)$.

Theorem 3.1. *Consider a parameterized dynamical system (1.2) and choice of grids $\mathcal{X}$ and $\mathcal{Z}$ for phase and parameter space and a choice of multivalued outer approximations $\mathcal{F}_\zeta$ for $F_{|\zeta|}$. Let $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}^S)$ be the database obtained using the strip method. Then there exists a database coarsening*

$$\Phi \colon \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \widetilde{\mathcal{F}}) \to \mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}^S).$$

Proof. Recall that for $\zeta \in \mathcal{Z}$, the difference between $\mathcal{F}_\zeta$ and $\widetilde{\mathcal{F}}_\zeta^S$ is that the latter may be computed on a strip $\mathcal{S}$, i.e. a set $\{\zeta \mid \zeta \in \mathcal{S}\}$. Notice $\mathcal{F}_\zeta^S$ is a valid multivalued map for all $\lambda \in |S|$, and under the monotonicity assumption **A1**, $\zeta \subset S$ implies $\widetilde{\mathcal{F}}_\zeta \subset \mathcal{F}_\zeta^S$. Therefore the result follows from Corollary 1.12.

□

Observe that this theorem implies that the strip method gives a correct, but not necessarily finest database.

4. Computational Savings

In this section we derive a lower bound on the savings provided by Algorithm 1 over the standard construction, wherein each Conley-Morse graph is computed for each $\zeta \in \mathcal{Z}$. We compute the estimate in the neighborhood of a saddle-node bifurcation, since such bifurcation is generic and we can use the estimates for the sizes of different continuation classes in the parameter space [1]. The savings are realized when the two endpoints, ζ_0 and ζ_n, of a 1-strip fall into the same continuation class, and $\mathsf{CMG}(\mathcal{F}_{\zeta_0})$ and $\mathsf{CMG}(\mathcal{F}_{\zeta_n})$ belong to the same equivalence class.

4.1. Saddle-Node Bifurcation

In this section we review work of [1] which analyzes the neighborhood of a generic saddle-node bifurcation using the Conley-Morse Database. Let $f : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, given by

$$f(x, \lambda) = x + x^2 - \lambda \tag{4.1}$$

The fixed points of this map are solutions to

$$\varphi(x, \lambda) := f(x, \lambda) - x = 0$$

and are given by $x = \pm\sqrt{\lambda}$, for $\lambda \geq 0$.

The following assumptions characterize saddle node-bifurcation at $(0, 0)$.

SN1 $f(0, 0) = 0$ and $f_x(0, 0) = 1$;

SN2 $f_\lambda(0, 0) < 0$ and $f_{xx}(0, 0) > 0$.

In addition, we have the following estimate

Lemma 4.1. *Assume* **SN1** *and* **SN2**. *Then there exist compact intervals* $X = [X_-, X_+] \subset \mathbb{R}$ *and* $Z = [\Lambda_-, \Lambda_+] \subset \mathbb{R}$ *such that there exist positive constants* $a_0 \leq |f_\lambda(x, \lambda)| \leq a_1$ *and* $b_0 \leq |f_{xx}(x, \lambda)| \leq b_1$ *for all* $(x, \lambda) \in X \times Z$.

Following [1] we assume that X is subdivided into intervals of length δ and the parameter space Z into intervals of length ν. This defines the grid $\mathcal{X} \times \mathcal{Z}$ on $X \times Z$. The main result of [1] is that in the discretized space there are three regions in the parameter space $\mathcal{A}, \mathcal{B}, \mathcal{C} \subset \mathcal{Z}$, with the corresponding Conley-Morse graphs are, respectively, an empty set, a single Morse set and two Morse sets forming an attractor-repeller decomposition Fig. 4.1. Furthermore, they show that, as $\delta, \nu \to 0$, the parameter region for which the Conley-Morse graph is a single Morse set limits to the point at which the saddle-node bifurcations occurs.

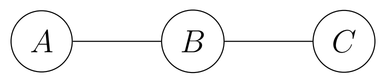

Figure 4.1: Continuation Graph for Saddle-Node Bifurcation

Finally, [1] also prove the following estimate for the region with single Morse set $|\mathcal{B}|$ in terms of discretization parameters δ and ν.

Theorem 4.2. *Assume that $\delta, \nu \to 0$ but $\frac{\delta}{\nu}$ remains bounded from above and below by positive constants. Then the set $|\mathcal{B}| = [\Lambda^a, \Lambda^c]$, where*

$$L := -\left(\frac{\delta}{a_0 - K_1\nu} + \frac{\nu}{1 - K_2\nu}\right) \leq \Lambda^a \leq 0$$

for some $K_1, K_2 > 0$, and

$$M := \frac{\delta}{a_0}\left(1 + \frac{b_1\delta}{8}\right) \geq \Lambda^c$$

4.2. Outline of the Argument

The computational savings in the bisection model will come from selection of non-neighboring parameter boxes which are both within either the $\mathcal{A}$ or $\mathcal{C}$ region, and for which their Conley-Morse graphs are deemed equivalent.

Before making an estimate for grid $\mathcal{Z}$ on Z, it is illuminating to make a minor adjustment, and give an estimate for the number of discretization points that fall into

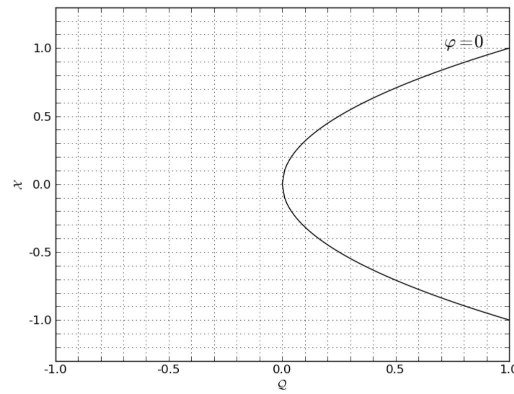

Figure 4.2: Grid defined by $\mathcal{X} \times \mathcal{Z}$, depicting solution to $\varphi(x, \lambda) = 0$.

$\mathcal{A}$ or $\mathcal{C}$ at a given level ν. From this we estimate computational savings of a database being computed just from these points, which may be regarded as the 0-dimensional faces (vertices) of 1-dimensional grid elements (edges).

Note that is a different scheme than the computation we present in this thesis, as we have discussed discretizing the entire parameter space into intersecting grid elements. However, this estimate will both prove useful to obtain an estimate for a grid $\mathcal{Z}$ on Z, as well as for ongoing work computing on only vertices and edges in $\mathcal{Z}$ (regardless of the dimension of Z).

4.3. An Estimate for Point Discretization

For sake of exposition, assume that $Z = [-1, 1]$. We mention here that our estimates still hold for the general case of $Z = [\Lambda_-, \Lambda_+]$, where $-1 < \Lambda_- < \Lambda_+ < 1$ following from Lemma 4.1, by simply rescaling Z. Furthermore, assume that there are $2^m + 1$ points $\{p_0, \ldots, p_{2^m}\}$ uniformly distributed in $[-1, 1]$ with $p_0 = -1$ and

$p_{2^m} = 1$. This implies that the discretization scale of the parameter space is

$$\nu = \frac{2}{2^m} = 2^{-m+1}.$$

Again, to aid in both the estimate for $\mathcal{Z}$ and ongoing work, we will first assume that the Morse graph is computed in the parameter space over isolated points, rather than intersecting grid elements. Later, we will use this computation to make estimates that pertain to grid elements.

We start with an example for the upper bound M. Similar behavior occurs for the lower bound L. Assume that there are nine points ($m = 3$) and assume that the upper bound on the set $\mathcal{B}$ is $M = \frac{1}{8}$, see Fig. 4.3. Let $\mathcal{F}$ be the associated outer approximation of Eqn. 4.1. Note that in our case of the saddle-node bifurcation, Algorithm 2 will give a positive answer every time two Morse Graphs belong to the same class $\mathcal{A}$, $\mathcal{B}$, or $\mathcal{C}$, and a negative answer if they do not. For point $p \in Z$, we will refer to the multi-valued map restricted to p as $\mathcal{F}_{\{p\}}$. We will abuse notation and use $\mathcal{A}, \mathcal{B}, \mathcal{C}$ as both the continuation class of Morse graphs and the corresponding region in parameter space. However, the meaning will be clear from the context.

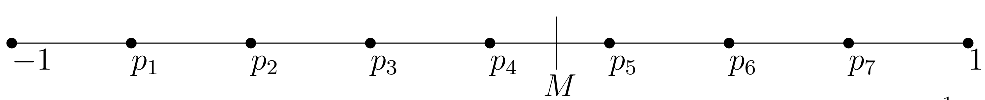

Figure 4.3: Example point discretization on Z, with bounds $M = \frac{1}{8}$

The local bisection algorithm first tests $p_0 = -1$ and $p_8 = 1$. Since $\mathsf{MG}(\mathcal{F}_{\{-1\}}) \in \mathcal{A}$ and $\mathsf{MG}(\mathcal{F}_{\{1\}}) \in \mathcal{C}$, the bisection algorithm will next tests $\mathsf{MG}(\mathcal{F}_{\{p_4\}})$. Since $\mathsf{MG}(\mathcal{F}_{\{p_4\}})$ $\in \mathcal{B}$ the interval $[p_4, 1]$ will be subdivided again. Since both $\mathsf{MG}(\mathcal{F}_{\{p_6\}})$ and $\mathsf{MG}(\mathcal{F}_{\{1\}})$, where $p_6 = \frac{1}{2}$, are in class $\mathcal{C}$, this interval will not be subdivided further. The isomorphism fails between $\mathsf{MG}(\mathcal{F}_{\{p_4\}})$ and $\mathsf{MG}(\mathcal{F}_{\{p_6\}})$, and this interval is subdivided further to $[p_4, p_5]$ and $[p_5, p_6]$. Note that even though $\mathsf{MG}(\mathcal{F}_{\{p_6\}})$ and $\mathsf{MG}(\mathcal{F}\{p_5\})$ belong to the same class, we do not save any computation over the exhaustive search, since these are neighboring points in the grid. Our savings came from not needing to compute Morse graph for p_7.

4.4. Number of Points in Regions $\mathcal{A}$ and $\mathcal{C}$

Observe that in the general case our savings come from the number of intervals of length $\frac{1}{4}, \frac{1}{8}, \ldots, \frac{1}{2^k}$ which can cover interval $[M, 1]$, starting from the right end. An interval of length $\frac{1}{2^k}$ covers $2^{m-k-1} + 1$ points, where the exponent -1 comes from the fact that the length of Z is 2. Since we always compute at both endpoints, the number of internal points, and hence points where we do not compute, is

$$2^{m-k-1} - 1$$

points for each such interval.

The number of intervals of length $\frac{1}{2^k}$ that cover $[M, 1]$ can be computed from the binary expansion up to order m of the number $1 - M$. Indeed, every 1 in this expansion represent such an interval. Let

$$1 - M = 0.\sigma_1\sigma_2\sigma_3\ldots = \sum_{k=1}^{\infty} \frac{\sigma_k}{2^k},$$

where $\sigma_i \in \{0, 1\}$.

Therefore the savings will be

$$\begin{aligned} & \sigma_1(2^{m-1-1} - 1) + \sigma_2(2^{m-2-1} - 1) + \sigma_3(2^{m-3-1} - 1) + \ldots + \sigma_{m-2}(2^1 - 1) \\ = & \sum_{k=1}^{m-2} \sigma_k(2^{m-k-1} - 1). \end{aligned} \tag{4.2}$$

Combining this with a similar estimate for the class $\mathcal{A}$ using the binary expansion of

$$L+1 = 0.\tau_1\tau_2\tau_3\ldots = \sum_{k=1}^{\infty} \frac{\tau_k}{2^k},$$

we get the following Lemma.

Lemma 4.3. *Assume that $Z = [-1,1]$ and there are 2^m+1 points P distributed uniformly in Z. Let L and M be bounds on the class $\mathcal{B}$, as above. Let $\{\sigma_i\}_{i=1}^{\infty}$ and $\{\tau_i\}_{i=1}^{\infty}$ be digits of the binary expansions of $1-M$ and $L+1$, respectively. Applying bisection algorithm to points P, let $S \subset P$ be the set of points where the Morse Graph computation will not be performed. Then*

$$|S| \geq \left(\sum_{k=1}^{m-2} \sigma_k(2^{m-(k+1)}-1) + \sum_{k=1}^{m-2} \tau_k(2^{m-(k+1)}-1)\right) \tag{4.3}$$

Notice that in the example $1-M = \frac{7}{8} = 0.\sigma_1\sigma_2\sigma_3\ldots = 0.111$ Thus our estimate is $\sum_{k=1}^{m-2} \sigma_1(2^{m-(k+1)}-1) = 1$.

The previous argument readily generalizes to a problem, where we compute Conley-Morse graphs that are valid for entire grid element ζ_i, which we discuss in the next section.

4.5. An Estimate for Grid Elements

In this section, we slightly modify the previous estimate to compute lower bound on savings when we compute Conley-Morse graphs over elementary boxes of the grid of diameter ν imposed on Z.

More specifically, assume that the parameter space $Z = [-1,1]$ is discretized into 2^m+1 intervals of size ν. This means that there is 2^m+2 points which bound these intervals and thus

$$\nu = \frac{2}{2^m + 1}.$$

Thus $\mathcal{Z} = \{\zeta_i = [\lambda_i, \lambda_{i+1}] \mid i = 0...2^m\}$, where $\lambda_{i+1} - \lambda_i = \nu$. We assume that m is sufficiently large that Conley-Morse graph at ζ_0 is in class $\mathcal{A}$ and the Conley-Morse graph at ζ_{2^m} belongs to $\mathcal{C}$. Since these disagree, the algorithm will compute the Conley-Morse graph at the middle interval, which belongs to class $\mathcal{B}$.

We observe that the situation with grid elements ζ_i is completely analogous to the computation we have performed for individual points. The only difference is that instead of finding the nearest point p_j to the right of the upper bound M of the class $\mathcal{B}$ we have to find the grid element ζ_j with minimum index that is a subset of $\mathcal{C}$, and instead of finding the nearest point to the left of the lower bound L of the class $\mathcal{B}$ we have to find the grid element with maximum index that is a subset of $\mathcal{A}$.

We have the following result.

Theorem 4.4. *Assume that $Z = [-1, 1]$ is divided into $2^m + 1$ intervals ζ_i distributed uniformly in Z. Let L and M be bounds on the class $\mathcal{B}$, as above. Let $\{\sigma_i\}_{i=1}^{\infty}$ and $\{\tau_i\}_{i=1}^{\infty}$ be digits of the binary expansions of $2(1 - \frac{k_M}{2^m})$ and $2(\frac{k_L}{2^m})$, respectively. Here the numbers*

$$k_M = \lfloor (M-1)(\frac{2^m+1}{2}) + 2^m + 1 \rfloor$$
$$k_L = \lceil \frac{1}{2}(L+1)(2^m+1) - 1 \rceil,$$

where $\lfloor \cdot \rfloor$ ($\lceil \cdot \rceil$) denotes the closest lower (upper) integer, i.e. floor (ceiling), of the argument.

Applying bisection algorithm to intervals ζ_i, let $S \subset \{0, 1, \ldots, 2^m\}$ be the set of indices of those intervals where the Morse Graph computation will not be performed.

Then

$$|S| \geq \left(\sum_{k=1}^{m-2} \sigma_k (2^{m-(k+1)} - 1) + \sum_{k=1}^{m-2} \tau_k (2^{m-(k+1)} - 1) \right) \tag{4.4}$$

Proof. We show that the number k_M is the smallest index of a grid element ζ_{k_M} which is a subset of $\mathcal{C}$, and k_L is the largest index of a grid element ζ_{k_L} which is still a subset of $\mathcal{A}$. Indeed, as our index set starts from 0, k_L is the the minimum index that satisfies

$$L < -1 + (k_L + 1)\nu = -1 + (k_L + 1)(\frac{2}{2^m + 1}).$$

To compute k_L we let $L + 1 = (k_L + 1)(\frac{2}{2^m+1})$ from which

$$k_L = \lceil \frac{1}{2}(L + 1)(2^m + 1) - 1 \rceil.$$

Similar argument for k_M finds the maximum index k such that

$$M > 1 - (2^m - (k_M - 1))\nu = 1 - (2^m - (k_M - 1))\frac{2}{2^m + 1}.$$

Setting the equality and a short computation shows

$$k_M = \lfloor (M - 1)\left(\frac{2^m + 1}{2}\right) + 2^m + 1 \rfloor.$$

□

CHAPTER 5

EXPERIMENTAL RESULTS

1. Implementation of Strip Method

As described above, for practical purposes we utilize a heuristic to determine whether the two Morse graphs are isomorphic. In other words, we view Morse graphs purely as combinatorial objects, not taking into account any phase space information or Conley index information. We then subdivide the space between two grid elements if their associated Morse graphs fail the following isomorphism heuristic. In the following section we introduce the heuristic.

1.1. Isomorphism Heuristic

The isomorphism heuristic is given by Algorithm 2, with Encode-Graph as a subroutine decomposing the graph into a list of its vertices' data. Note that Algorithm 1 may be implemented with subroutines using a different set of criteria for determining when to subdivide. For instance, one could take phase space information into account and test whether the clutching function induced a directed graph isomorphism between the two Morse graphs (at the price of being more computationally expensive). However, for the purposes of this thesis we utilize the graph isomorphism test presented in Algorithm 2.

2. Two Parameter Overcompensatory Leslie Model

In order to test the strip method, we choose to compute on the Leslie population model. Mathematically, the Leslie model is defined to be a map $g : \mathbb{R}^d \to \mathbb{R}^d$ given

Algorithm 2 Heuristic for the graph isomorphism problem

```
function ISOMORPHISM-HEURISTIC(MG(F_ζ1), MG(F_ζ2))
    Q_1 ← ENCODE-GRAPH(MG(F_ζ1))
    Q_2 ← ENCODE-GRAPH(MG(F_ζ2))
    if Q_1 = Q_2 then
        return true
    else
        return false
    end if
end function
function ENCODE-GRAPH(G := (V, E))
    Q ← ∅
    for all v ∈ V do
        a ← Number of Ancestors of v
        d ← Number of Descendents of v
        i ← In degree of v
        o ← Out degree of v
        ENQUEUE(Q, (a, d, i, o))
    end for
    Sort(Q)                          ▷ Lexicographical sorting subroutine
    return Q
end function
```

by

$$\begin{bmatrix} x_1^n \\ x_2^n \\ \vdots \\ x_d^n \end{bmatrix} \mapsto \begin{bmatrix} (\theta_1 x_1^n + \ldots + \theta_d x_d^n) e^{\alpha(x_1^n + \ldots + x_d^n)} \\ p_1 x_1^n \\ \vdots \\ p_{d-1} x_{d-1}^n \end{bmatrix}$$

The Leslie model, as well as a discussion regarding its importance in population biology, are given in detail in [11]. In the model, the population is partitioned into d generations, each with population $x_1, \ldots, x_d$, and an associated reproduction rate. The nonlinearity stems from the assumption that fertility decreases exponentially with the total size of the population.

In particular, we choose to conduct our experiments on the two-dimensional Leslie model. For visualization purposes, we also choose our parameter space to be two-dimensional, and fix $p_1 = .7$ and $\alpha = -.1$.

Figs. 5.1–5.3 are from parameter space of $Z = [8, 37] \times [3, 50]$, using six parameter space subdivisions, i.e. $d_Z = 6$, giving $(2^6)^2 = 4096$ grid elements in $\mathcal{Z}$ total. See the captions for more information on what the figures represent. Further, observe that this implies that we have computated 4096 Morse graphs for Fig. 5.1.

3. Conley-Morse Database

Figs. 5.1–5.3 demonstrate a test against an original computation of the Conley-Morse Database. In order to facilitate an accurate comparison, we use subdivide $\mathcal{Z}$ into 2^k intervals in each dimension, though one may note that 2^k leads to a slight bias in the behavior in the bisection algorithm. This is opposed to what we presented in the last section, where in the general case we use 2^{k+1} for the purposes of bisection. The computations are done with $Z = [8, 37] \times [3, 50]$ and $X = [0, 320.056] \times [0, 224.040]$, as well as 6 subdivisions in parameter space (creating $(2^6)^d$ grid elements in $\mathcal{Z}$ - in this case $d = 2$).

As can be seen from the figures, the choice of the direction chosen to form the strips results in two different databases. In particular, the results of Fig. 5.3 may be surprising, in that it shows the bisection algorithm can completely shear off regions of continuation classes. From Table 5.1, it can be seen that the orientation of the strips also affects runtime, though marginally in these cases.

Figs. 5.4–5.6 are from an extended parameter space of $Z = [8, 66] \times [3, 97]$, using eight parameter space subdivisions, or $(2^8)^2$ grid elements in $\mathcal{Z}$ total.

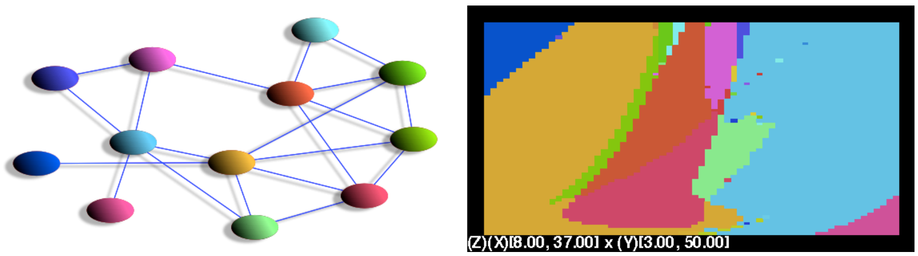

Figure 5.1: Original Conley-Morse database computation; (Left) Continuation Graph: Each vertex corresponds to a continuation class, and each edge corresponds to a clutching graph between the representative Conley-Morse graphs. (Right) Parameter space partitioned into regions according to the projections of continuation classes, $\pi(CC_j)$. The color coding of parameter space matches the color coding of the nodes in the continuation graph.

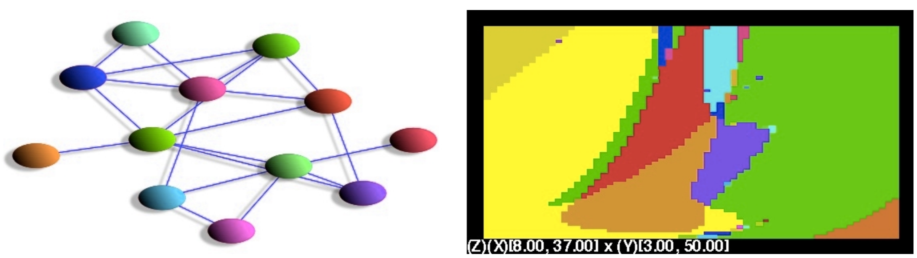

Figure 5.2: Strip Method; strips formed in θ_1 direction

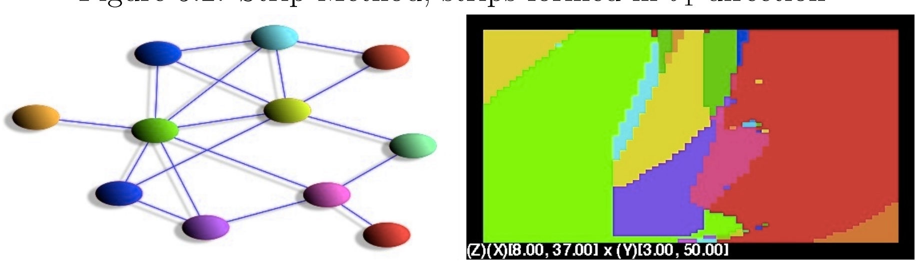

Figure 5.3: Strip Method; strips formed in θ_2 direction

Subdiv.	CMDB Time	Strips in X Time	Strips in Y Time
2^6	78	58	64.15
2^8	1934.27	690.66	702.7

Table 5.1: Timing information (minutes) for original database computation (CMDB) against the strips database with strips formed in X or Y directions.

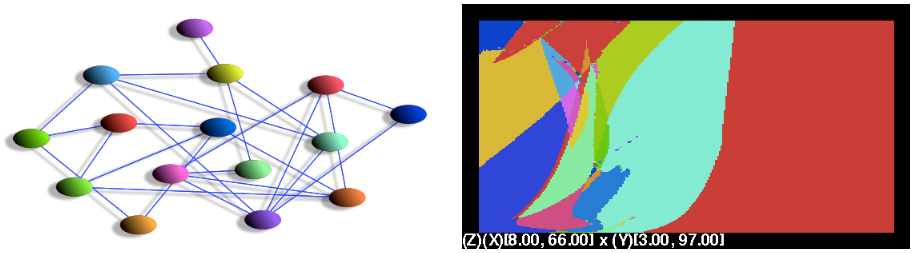

Figure 5.4: Original Conley-Morse Database Computation

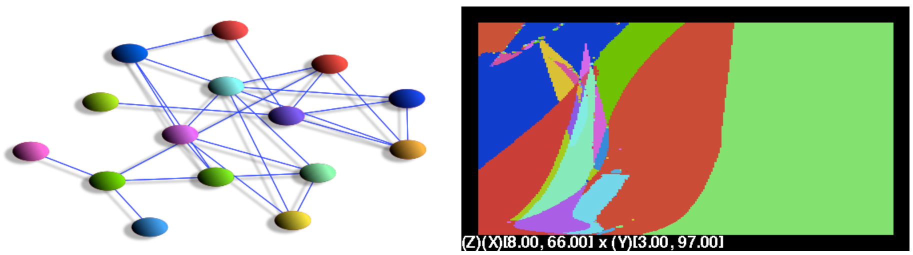

Figure 5.5: Strip Method Computation; Strips in θ_1 direction

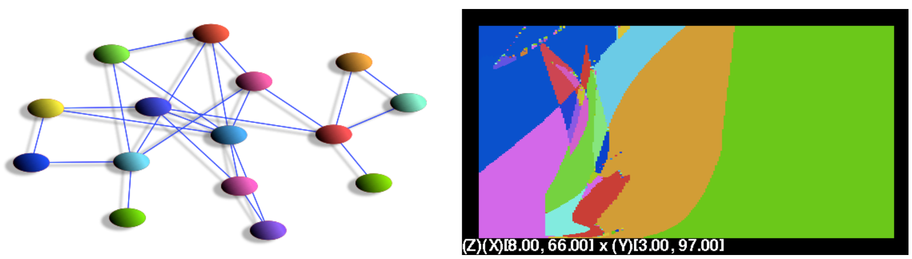

Figure 5.6: Strip Method Computation; Strips in θ_2 direction

4. Empirical Comparison

In this section, we quantify the difference between the $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}^S)$ and $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$ for our empirical tests. To do this, note that each $\zeta \in \mathcal{Z}$ from an original database calculation is contained in a unique grid element $\mathcal{S}$ in $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}^S)$. We compare $\mathsf{MG}(\mathcal{F}_\zeta)$ and $\mathsf{MG}(\mathcal{F}_\mathcal{S})$ using Algorithm 2, the graph isomorphism heuristic. To quantify this relationship we calculate the percentage of $\zeta \in \mathcal{Z}$ from $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}^S)$ whose Morse graphs, $\mathsf{MG}(\mathcal{F}_\zeta)$ pass the isomorphism heuristic with their counterpart

$\mathsf{MG}(\mathcal{F}_{\mathcal{S}})$ in $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$. While this form of comparison will not account for changes in continuation classes, it still provides a coarse estimate as to the quality of our results.

Subdiv.	# in CMDB	# Isomorphisms in X (%)	# Isomorphisms in Y (%)
2^6	4096	4089 (99%)	3952 (96.5%)
2^8	65536	65404 (99.8%)	64795 (98.9%)

Table 5.2: Comparison of grid elements which pass the isomorphism heuristic in each database. The first column presents the subdivision depth while the second column displays the number of total grid elements in an original column. The third and fourth columns depict how many grid elements in the original database pass the heuristic with their counterparts in the strips databases.

The results in Table 5.2 show that, at least on the level of isomorphism, the strip computation approximates an original computation very precisely. Note that by Corollary 1.12, for $\zeta \in \mathcal{Z}$ and $\mathcal{S}$ such that $\zeta \subset \mathcal{S}$, if $\mathsf{MG}(\mathcal{F}_\zeta)$ and $\mathsf{MG}(\mathcal{F}_{\mathcal{S}})$ do not pass the isomorphism test then there exists a coarsening from P_ζ to $\mathsf{P}_{\mathcal{S}}$. In particular, there are natural inclusion maps from the Morse graphs of $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F})$ to their counterparts in the $\mathsf{DB}(\mathcal{X}, \mathcal{Z}, \mathcal{F}^S)$.

CHAPTER 6

CONCLUSION

1. Concluding Remarks

In this thesis we introduced an approach to cataloging global changes in dynamical systems using the framework of the Conley-Morse database. This approach to the study of complex dynamical systems provides rich knowledge of the system, especially in the event that exact parameters are unknown. This is a common problem in the biological sciences, where obtaining parameters for a model often requires extensive experimentation. For instance, certain models of chemotaxis [25] involve upwards of seventeen parameters, some only approximated, ranging across eight orders of magnitude. In the database framework described in Section 3, the parameters can be taken over a range of values, with rigorous results valued over that entire region in parameter space. Much effort is spent validating and tuning model parameters [6], an aspect that the database schema helps to alleviate.

Though the database is based on fast graph theoretic algorithms, the combinatorial descriptions of both the phase space and parameter space inherently suffer from the curse of dimensionality. We have shown that adaptive construction techniques based on assumptions of the underlying dynamics can lead to accurate approximations of an original database, alleviating some of the computational burden due to high-dimensional parameter spaces.

Furthermore, the presentation of this thesis has focused on discrete time dynamical systems. From a theoretical perspective these same ideas apply to dynamics governed by differential equations. However, practically, the issue of constructing an outer approximation for flows is more difficult. Though there exist tools for doing this effectively, the computational cost is much greater than simply evaluating a map (as is done in the case of discrete time). Therefore, an adaptive scheme such as the one presented in this thesis becomes even more relevant in such a setting.

Another topic which has only been briefly discussed is the relationship between the database and classical bifurcation theory. The boundaries between continuation classes need not represent boundaries between different topological conjugacy classes of dynamics, and vice versa [2]. The beginnings of such analysis have appeared in [1], whose estimates we incorporated to provide a bound on savings for the Strip Method. However beyond this, the problem of more generally recognizing bifurcations from the database and determining how the database will represent such bifurcations in terms of grid size in both phase and parameter space remains unexplored; it is an important direction for the future of the database project.

REFERENCES CITED

[1] Z. Arai, M. Gameiro, T. Gedeon, H. Kokubu, K. Mischaikow, and H. Oka. Graph-based Topological Approximation of Saddle-Node Bifurcation in Maps. *RIMS Kokyuroku Beassatsu*, 2012.

[2] Z. Arai, W. Kalies, H. Kokubu, K. Mischaikow, H. Oka, and P. Pilarczyk. A Database Schema for the Analysis of Global Dynamics of Multiparameter Systems. *SIAM Journal on Applied Dynamical Systems*, 8(3):757, 2009.

[3] J. Bush, M. Gameiro, S. Harker, H. Kokubu, K. Mischaikow, I. Obayashi, and P. Pilzarczyk. Combinatorial-Topological Framework for the Analysis of Global Dynamics. *Chaos*, 22, 2012.

[4] C.C. Conley. *Isolated Invariant Sets and the Morse Index*. Regional conference series in mathematics. R.I.:American Mathematical Society, 1978.

[5] T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, third edition, 2009.

[6] S. Day, O. Junge, and K. Mischaikow. A rigorous numerical method for the global analysis of infinite dimensional discrete dynamical systems. *SIAM J. Appl. Dyn. Syst*, 3:117–160.

[7] S. Fortin. The graph isomorphism problem. Technical report, 1996.

[8] W. D. Kalies, K. Mischaikow, and R. C. A. M. VanderVorst. An algorithmic approach to chain recurrence. *Found. Comput. Math.*, 5(4):409–449, November 2005.

[9] K. Mischaikow. Topological Techniques for Efficient Rigorous Computation in Dynamics. *Acta Numerica*, 11:435–477, 2002.

[10] A. Szymzcak. *Index Pairs: From Dynamics to Combinatorics and Back*. PhD thesis, Georgia Institute of Technology, 1999.

[11] I. Ugarcovici and H. Weiss. Chaotic dynamics of a nonlinear density dependent population model. *Nonlinearity*, 17:1689–1711, 2004.