Acta Numerica (2002), pp. 435–477
DOI: 10.1017/S0962492902000065
© Cambridge University Press, 2002
Printed in the United Kingdom

Topological techniques for efficient rigorous computation in dynamics

Konstantin Mischaikow*
Center for Dynamical Systems and Nonlinear Studies,
Georgia Institute of Technology,
Atlanta, GA 30332, USA
E-mail: mischaik@math.gatech.edu

We describe topological methods for the efficient, rigorous computation of dynamical systems. In particular, we indicate how Conley's Fundamental Decomposition Theorem is naturally related to combinatorial approximations of dynamical systems. Furthermore, we show that computations of Morse decompositions and isolating blocks can be performed efficiently. We conclude with examples indicating how these ideas can be applied to finite- and infinite-dimensional discrete and continuous dynamical systems.

CONTENTS

1 Introduction 435
2 Chain recurrence and isolating blocks 440
3 Conley index 462
4 Examples 468
References 474

1. Introduction

This paper is an expository article on using topological methods for the *efficient*, *rigorous* computation of dynamical systems. Of course, since its inception the computer has been used for the purpose of simulating nonlinear models. However, in recent years there has been a rapid development in numerical methods specifically designed to study these models from a dynamical systems point of view, that is, with a particular emphasis on the structures which capture the long-term or asymptotic states of the system. At the risk of greatly simplifying these results, this work has followed two themes: *indirect methods* and *direct methods*.

* Research supported in part by NSF grant DMS-9805584 and DMS-0107395.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

The indirect methods are most closely associated with simulations, and as such are extremely important because they tend to be the cheapest computationally. The emphasis is on developing numerical schemes whose solutions exhibit the same dynamics as the original system: for example, if one is given a Hamiltonian system, then it is reasonable to want a numerical method that preserves the integrals of the original system. A comprehensive introduction to these questions can be found in Stuart and Humphries (1996).

The direct methods focus on the development of numerical techniques that find particular dynamical structures, for instance fixed points, periodic orbits, heteroclinic orbits, invariant manifolds, *etc.*, and are often associated with continuation methods (see Cliffe, Spence and Tavener (2000), Doedel, Champneys, Fairgrieve, Kuznetsov, Sandstede and Wang (1999) and Dieci and Eirola (2001) and references therein). To paraphrase Poincaré, these techniques provide us with a window into the rich structures that nonlinear systems exhibit.

There is no question that these methods are essential. However, they cannot capture the full dynamics. As pointed out in Stuart and Humphries (1996, p. xiii) a fundamental question for the indirect method, requiring a positive answer, is:

> Assume that the differential equation has a particular invariant set. Does the numerical method have a corresponding invariant set which converges to the true invariant set as $\Delta t \to 0$?

For non-hyperbolic systems the answer is almost surely 'no'. The set of parameters at which global bifurcations occur can be dense, which suggests that any numerical error will lead to significant distinctions in the structure of the invariant sets (Palis and Takens 1993).

Similarly, for the direct methods to compute these invariant objects effectively, they typically need to be hyperbolic, which limits the applicability for general systems. Further, one of the more interesting features of nonlinear systems is the existence of complicated or chaotic dynamics, the structure of which has not yet been captured by the above-mentioned methods.

A final observation is that, while these dynamical objects are ubiquitous, the number of specific nonlinear systems for which mathematicians have been able to prove the existence of these structures, especially as it relates to chaotic dynamics, is quite small. Therefore, it appears that there is room for a *complementary* computational method which can provide accurate information about global structures of nonlinear dynamical systems and is rigorous in the sense of a mathematical proof, albeit involving the assistance of a computer.

At the same time, such a method is of little value unless it can be shown to be computationally efficient. Hence, throughout this paper the discussion

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

of the topological techniques will be presented in the context of the computational effort required.

For the sake of simplicity we will restrict most of our discussion to the dynamics of continuous maps. However, we will begin on a very general level, specializing as we proceed. Let X be a metric space and let $f : X \to X$ be a continuous function. Since we are not assuming that f is a homeomorphism we will make extensive use of the following definition. A *full trajectory* of f through $x \in X$ is a function $\sigma_x : \mathbb{Z} \to X$ satisfying:

(1) $\sigma_x(0) = x$, and
(2) $\sigma_x(n+1) = f(\sigma_x(n))$ for all $n \in \mathbb{Z}$.

In other words, $\sigma_x(k) = f^k(x)$, where f^k denotes the kth iterate of f. Recall that $S \subset X$ is *invariant* under f if, for every $x \in S$, there exists a full trajectory $\sigma_x : \mathbb{Z} \to S$.

Ideally we want to understand the existence and topological structure of invariant sets. The same above-mentioned conundrum shows that, as stated, this goal is completely unrealistic. Therefore, we step back and ask:

Q1. *What should one compute?*

The theme of Section 2 is that ϵ-chain-recurrent sets and isolating blocks are very computable objects.

The ϵ-chain-recurrent set consists, essentially, of those points which, having allowed for a finite number of errors of size ϵ, return to themselves under the dynamics. Conley's Fundamental Decomposition Theorem, Theorem 2.1, states that, away from the chain-recurrent set, the dynamics is gradient-like. Thus, finding the chain-recurrent set provides, on a very crude level, global information about the dynamics. This is described in detail in Section 2.1.

A compact set $N \subset X$ is an *isolating neighbourhood* for f if

$$\mathrm{Inv}(N, f) := \{x \in X \mid \exists \sigma_x : \mathbb{Z} \to N\} \subset \mathrm{int}(N),$$

where $\mathrm{int}(N)$ denotes the interior of N. $\mathrm{Inv}(N, f)$ is the *maximal invariant set* in N. An invariant set S is *isolated* if there exists an isolating neighbourhood N such that $S = \mathrm{Inv}(N, f)$.

Observe that if N is an isolating neighbourhood then $\mathrm{Inv}(N, f) \cap \partial N = \emptyset$, where ∂N denotes the topological boundary of N. An *isolating block* is a special form of an isolating neighbourhood that satisfies the additional condition

$$f^{-1}(N) \cap N \cap f(N) \subset \mathrm{int}(N).$$

If N is an isolating block and $x \in \partial N$, then either $f(x) \cap N = \emptyset$ or $f^{-1}(x) \cap N = \emptyset$, that is, points on the boundary either leave N immediately or all its pre-images lie outside N. Isolating blocks and their relation to chain-recurrent sets will be discussed in Section 2.2.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Isolating neighbourhoods serve two purposes. The first is to localize the dynamics being considered. As will be made clear later, once an isolating neighbourhood N has been chosen, the focus is on the structure of the dynamics of $\mathrm{Inv}(N, f)$. The second point follows from the observation that continuity and compactness implies that, if N is an isolating neighbourhood for f, then N is an isolating neighbourhood for perturbations of f. This suggests that isolating neighbourhoods are robust objects and therefore computable.

Though the entire discussion could be carried out in greater generality, from now on $X \subset \mathbb{R}^n$ is compact and $f : \mathbb{R}^n \to \mathbb{R}^n$. Also, assume that X is an isolating neighbourhood for f. Then we can consider $f : X \to \mathbb{R}^n$.

Of course, the method of computation of isolating neighbourhoods is essential. In Section 2.3, motivated by keeping track of numerically induced errors, we consider *topological multivalued maps*. To be more precise, consider the continuous function f which generates the dynamics we are interested in analysing. Given $x \in X$ and using the computer to evaluate f, we cannot expect to obtain the value $f(x)$. Because of the associated numerical errors, at best we obtain a numerical value $f_n(x)$ and an error bound $\epsilon > 0$, from which we can conclude that $f(x) \in B_\epsilon(f_n(x))$, *i.e.*, that the true image lies in an ϵ-ball centred at the numerically generated image. In essence, to each point $x \in f^{-1}(X)$ we can associate a subset $F(x) \subset X$ with the property that $f(x) \in F(x) \subset X$. To emphasize the fact that the images of F are subsets of a topological space, we refer to F as a topological multivalued map and use the notation $F : X \rightrightarrows X$.

The next step is to obtain a combinatorial description of F. This is done by discretizing the space X via a finite grid $\mathcal{G}$, to obtain a purely combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$. Some of the material in these sections may appear redundant, as we generalize concepts from dynamical systems for continuous maps such as f to a topological multivalued map F, and then to a combinatorial multivalued map $\mathcal{F}$. However, we want to stress that the topological multivalued map provides a straightforward relationship between the dynamical structures of interest: those corresponding to f, and the computable structures of the combinatorial multivalued map $\mathcal{F}$.

Finally, in Section 2.4 we discuss the graph algorithms that are used to perform the computations. As will be made precise, the key operations are essentially linear in the number of grid elements used to do the approximations.

The discussion in all these subsections is done on the level of maps. From the perspective of modelling, the dynamics of flows generated by differential equations is at least as important. In Section 2.5 we discuss three different approaches: time τ maps, Poincaré maps, and flow-transverse polygonal approximations. The first two are fairly standard and therefore will be only briefly mentioned. The third approach is still speculative. It is presented

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

here because the same algorithms that apply to the map case work in this setting. What is lacking are computational geometry algorithms for obtaining simplicial decompositions of the phase space that lead to good approximations of the flow. Not surprisingly, this appears to be a difficult problem when the dimension of the phase space is greater than or equal to three.

Having argued that ϵ-chain-recurrent sets and isolating blocks are computable, the obvious question is:

Q2. *How can one interpret the results of the computations?*

Algebraic topology and, in particular, the Conley index are used at this stage. The Conley index is the topological generalization of Morse theory to maps and flows which do not necessarily arise from variational problems. The Conley index is very briefly described in Section 3; more complete discussions can be found elsewhere. What we wish to emphasize now is that the Conley index has three important characteristics.

(1) The Conley index is an index of isolating neighbourhoods. Furthermore, if N and N' are isolating neighbourhoods for the map f and $\mathrm{Inv}(N, f) = \mathrm{Inv}(N', f)$, then the Conley index of N is the same as the Conley index of N'. Observe that this implies that we can also consider the Conley index as an index of isolated invariant sets. Both interpretations are convenient, the first for applications and the latter for presenting definitions.

(2) **Ważewski Property:** If the Conley index of N is not trivial, then $\mathrm{Inv}(N, f) \neq \emptyset$.

(3) **Continuation:** If N is an isolating neighbourhood for a continuous family of maps f_λ, for instance,

$$\mathrm{Inv}(N, f_\lambda) \subset \mathrm{int}\, N \quad \text{for } \lambda \in [0, 1],$$

then the Conley index of N under f_0 is the same as the Conley index of N under f_1.

The first point suggests that to detect specific invariant sets does not require one to compute a unique isolating block. The second is the simplest example of how the Conley index is used to make assertions about the dynamics of f. The third point leads to the conclusion that the Conley index of the dynamical system generated by a sufficiently good numerical approximation will be the same as that of the original system. In a weak, but universal, sense this is the positive answer to the question posed for indirect methods.

As was mentioned above, little detail about the index theory is presented in this paper, though references are provided. The focus is on the computability. As will be emphasized, having computed chain-recurrent sets and isolating blocks, the input for the Conley index computations is available at essentially no cost. However, to determine the Conley index, homology

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

groups and homology maps need to be computed. The cost of the homology groups seems to be reasonable, though for higher-dimensional complexes the known complexity bounds seem rather pessimistic. Currently the cost of computing homology maps is expensive. However, algorithms in this subject are fairly new and there appears to be considerable room for improvement.

We finish the paper with four examples that demonstrate that this approach can be applied in nontrivial settings. The first example is the Hénon map. While the results are not surprising, that is, they agree with results obtained by simulations, they are mathematically rigorous and provide information about the structure of the chaotic dynamics that is observed. In particular, this approach provides the greatest rigorous lower bound on the entropy of the global attractor at the classical parameter values. The next example involves the Lorenz equations, and is included to demonstrate that, using Poincaré sections, these methods can also be applied to differential equations.

A criticism of these topological methods is that the cost grows very rapidly with dimension. This is a serious issue, which cannot be dismissed lightly. For this reason we have included two recent results that involve infinite-dimensional systems. In both examples, even though the ambient phase space is infinite-dimensional, the dynamic objects of interest are low-dimensional, for instance, fixed points, periodic orbits, horseshoes, *etc.* From the purely topological point of view one might hope that it is the dimension of the invariant set that determines the complexity of the computations. These rather special examples provide evidence that this might actually be the case.

2. Chain recurrence and isolating blocks

In this section we discuss the computation of isolating blocks and ϵ-chain-recurrent sets. The presentation begins with classical definitions and results from dynamical systems and ends with graph algorithms. The focus is on the computability of these objects and on the fact that there is a mathematically valid interpretation of the graph-theoretic structures as objects in continuous dynamical systems.

2.1. Chain recurrence

The following is a standard definition from dynamical systems (Easton 1998, Robinson 1995).

Definition 1. Let S be a subset of X. An ϵ-*chain* from x to y in S for the map f is a finite sequence of points $\{z_0, z_1, \ldots, z_n\} \subset S$, such that $x = z_0$, $y = z_n$ and $\|f(z_i) - z_{i+1}\| < \epsilon$.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Though it may seem strange at the moment, we prefer to think of an ϵ-chain of f from x to y in S as a map $\mu_x : \{0, \ldots, n\} \to S$ satisfying:

(1) $\mu_x(0) = x$ and $\mu_x(n) = y$,
(2) $\|\mu_x(i+1) - f(\mu_x(i))\| < \epsilon$ for $i = 1, \ldots, n-1$.

If $\mu_x : \{0, \ldots, n\} \to S$ is an ϵ-chain from x to itself, then $\mu_x(0) = \mu_x(n) = x$. In this case, we can extend $\mu_x : \mathbb{Z} \to S$ periodically. This leads to the following concept.

Definition 2. The *ϵ-chain-recurrent set* of S under f is defined by

$$\mathcal{R}(S, f, \epsilon) = \{x \in X \mid \text{there exists a periodic } \epsilon\text{-chain } \mu_x : \mathbb{Z} \to S\}.$$

There is a natural partition of $\mathcal{R}(S, f, \epsilon)$ into equivalence classes defined as follows.

Definition 3. Let $x, y \in S$. Set $x \sim_\epsilon y$ if and only if there exists a periodic ϵ-chain $\mu_x : \mathbb{Z} \to S$ such that $\mu_x(j) = y$ for some $j \in \mathbb{Z}$.

Let $\mathcal{P}$ be an indexing set for the resulting collection of equivalence classes, that is, let

$$\mathcal{R}(S, f, \epsilon) = \bigcup_{p \in \mathcal{P}} \mathcal{R}_p(S, f, \epsilon),$$

where $x, y \in \mathcal{R}_p(S, f, \epsilon)$ if and only if $x \sim_\epsilon y$.

Recall that a *strict partial order* $\succ$ is a nonreflexive, nonsymmetric, transitive relation. A strict partial order $\succ$ is an *admissible* order on $\mathcal{P}(S, f, \epsilon)$ if, for $p \neq q$, the existence of an ϵ-chain from an element of $\mathcal{R}_p(S, f, \epsilon)$ to an element of $\mathcal{R}_q(S, f, \epsilon)$ implies that $p \succ q$. Observe that it need not be the case that any two equivalence classes are necessarily related by $\succ$.

We are particularly interested in the ϵ-chain-recurrent sets in the case where S is a compact invariant set. In this case, simple arguments based on the continuity of f show that not only is $\mathcal{R}(S, f, \epsilon)$ open, but each equivalence class $\mathcal{R}_p(S, f, \epsilon)$ is also open. Since $\{\mathcal{R}_p(S, f, \epsilon) \mid p \in \mathcal{P}\}$ covers S, this implies that, for any given $\epsilon > 0$, there is at most a finite number of elements in $\mathcal{P}$.

Definition 4. The *chain-recurrent set* of S under f is given by

$$\mathcal{R}(S, f) := \bigcap_{\epsilon > 0} \mathcal{R}(S, f, \epsilon).$$

Again, one can define equivalence classes $\mathcal{R}_p(S, f)$, $p \in \mathcal{P}$ of $\mathcal{R}(S, f)$ by $x \sim y$ if and only if $x \sim_\epsilon y$ for all $\epsilon > 0$.

There are cases where $S = \mathcal{R}(S, f)$. For example, if S were a strange attractor, then it would be an invariant set, but every point in S is chain-recurrent to every other point.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Theorem 2.1. (Fundamental Decomposition Theorem) Let S be a compact invariant set for f. Then there exists a continuous Lyapunov function $L : S \to [0,1]$ satisfying:

(1) $L(f(x)) < L(x)$ for all $x \in S \setminus \mathcal{R}(S,f)$,

(2) for each $p \in \mathcal{P}$ there exists $c_p \in [0,1]$ such that $\mathcal{R}_p(S,f) \subset L^{-1}(c_p)$.

Furthermore, $\mathcal{R}(S,f)$ is a compact invariant set and

$$\mathcal{R}(\mathcal{R}(S,f),f) = \mathcal{R}(S,f). \tag{2.1}$$

To understand the importance of this theorem, observe that (2.1) implies that the chain-recurrent set is the minimal recurrent set. In particular, computing the chain-recurrent set while restricting ourselves to $\mathcal{R}(S,f)$ does not result in a smaller set. On the other hand, the existence of the Lyapunov function L indicates that we have captured all the recurrent dynamics. This theorem suggests that, to understand the global dynamics of f, it is sufficient to understand the dynamics in the equivalence classes $\mathcal{R}_p(S,f)$ of $\mathcal{R}(S,f)$ and the structure of the set of connecting orbits between these equivalence classes.

The Fundamental Decomposition Theorem provides a starting point for the investigation of *any* continuous dynamical system. However, within the context of this paper the question we now need to address is whether, from a practical point of view, the chain-recurrent set is computable. With this in mind, consider the following example.

Example 1. Let $f : \mathbb{R} \to \mathbb{R}$ be given by $f(x) = x + x^2 \sin(\frac{1}{x})$ Observe that $S = [-\frac{1}{\pi}, \frac{1}{\pi}]$ is a compact invariant set, and

$$\mathcal{R}(S,\varphi) = \left\{ \pm\frac{1}{n\pi} \mid n \in \mathbb{N} \right\} \cup \{0\}.$$

In particular, in this example $\mathcal{R}(S,\varphi)$ consists of an infinite number of connected components. This suggests that its explicit computation will be difficult.

To go a step further, consider a simple perturbation of this equation. In particular, let f^λ be given by

$$f^\lambda = x + x^2 \sin\left(\frac{1}{x}\right) + \lambda\left(x^2 - \frac{1}{\pi^2}\right).$$

Again, $S = [-\frac{1}{\pi}, \frac{1}{\pi}]$ is a compact invariant set for f^λ; however, for $\lambda \neq 0$ the structure of $\mathcal{R}(S,\varphi^\lambda)$ changes dramatically in the sense that the number of equivalence classes in $\mathcal{R}(S,f)$ is finite. The point to be emphasized is that small perturbations in the dynamical system can lead to significant qualitative changes in the recurrent set.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

This example is meant to suggest that while the chain-recurrent set is of fundamental importance to our understanding of the structure of invariant sets it should not be the focus of our computational efforts.

Therefore, we scale back our ambition concerning the detail of dynamical information that we seek. The simplest possibility is to preclude the study of an infinite collection of 'recurrent' sets. This leads to the following concept.

Definition 5. Let S be a compact invariant set for the map f. A *Morse decomposition* of S consists of a finite collection of mutually disjoint compact invariant sets,

$$\mathcal{M}(S) := \{M(p) \mid p \in \mathcal{P}\},$$

for which there exists a Lyapunov function $L : S \to [0, 1]$ satisfying:

(1) $L(f(x)) < L(x)$ for all $x \in S \setminus \bigcup_{p\in\mathcal{P}} M(p)$,

(2) for each $p \in \mathcal{P}$ there exists $c_p \in [0, 1]$ such that $M(p) \subset L^{-1}(c_p)$.

The individual invariant sets $M(p)$ are called *Morse sets*.

Though the motivation for this definition may appear somewhat artificial, we claim that Morse decompositions are, from the computational point of view, the *natural* global structures to study. To begin justifying this statement we will now work our way back to the ϵ-chain-recurrent set, beginning with some definitions.

Recall that the ω-*limit set* of x under f is

$$\omega(x) := \bigcap_{n>0} \mathrm{cl}\left(\bigcup_{i=n}^{\infty} f^i(x)\right),$$

where cl denotes the topological closure. We can define an α-*limit set* for x using a full trajectory, that is,

$$\alpha(\sigma_x) := \bigcap_{n<0} \mathrm{cl}\left(\bigcup_{i=n}^{-\infty} \sigma_x(i)\right).$$

Returning to the setting of a Morse decomposition, the existence of a Lyapunov function leads to the existence of partial orders on the indexing set $\mathcal{P}$. More precisely, since S is invariant, $x \in S \setminus \cup_{p\in\mathcal{P}} M(p)$ implies that there exists a full trajectory $\sigma_x : \mathbb{Z} \to S$ and Morse sets $M(p)$ and $M(q)$ such that:

$$\omega(x) \subset M(q) \qquad \text{and} \qquad \alpha(\sigma_x) \subset M(p). \tag{2.2}$$

The following definition will be used to describe the set of points with this property.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Definition 6. The set of *connecting orbits* from $M(p)$ to $M(q)$ is

$$C(p,q) := \{x \in S \mid \exists \sigma_x : \mathbb{Z} \to S, \alpha(\sigma_x) \subset M(p),\ \omega(x) \subset M(q)\}.$$

An *admissible order* is a strict partial order $\succ$ satisfying

$$C(p,q) \neq \emptyset \quad \Rightarrow \quad p \succ q.$$

This definition suggests that, given a Morse decomposition, the fundamental questions are: What is the dynamics within the Morse sets? and What is the structure of the set of connecting orbits between the Morse sets?

By choosing the same notation for the indexing sets and the partial orders we have strongly suggested that there is a relation between the equivalence classes of the ϵ-chain-recurrent set and Morse sets. The following theorem makes this precise.

Theorem 2.2. Let S be an invariant set for f. Let $\{\mathcal{R}_p(S,f,\epsilon) \mid p \in \mathcal{P}\}$ be the set of equivalence classes of the ϵ-chain-recurrent set of S. Define

$$M(p) := \mathrm{Inv}\,(\mathcal{R}_p(S,f,\epsilon), f).$$

Then $\mathcal{M}(S) := \{M(p) \mid p \in \mathcal{P}\}$ is a Morse decomposition of S. Furthermore, if $\succ$ is an admissible order for the equivalence classes of $\mathcal{R}(S,f,\epsilon)$, then $\succ$ is an admissible order for $\mathcal{M}(S)$.

Theorem 2.2 shows that an ϵ-chain-recurrent set produces a Morse decomposition. However, it is not the case that every Morse decomposition arises via an ϵ-chain-recurrent set. To explain this we shall explore a little further the structure of Morse decompositions.

Consider a Morse decomposition $\mathcal{M}(S) = \{M(p) \mid p \in (\mathcal{P}, \succ)\}$ where $\succ$ is an admissible order. A subset $I \subset \mathcal{P}$ is called an *interval* if $p, q \in I$ and $p \succ r \succ q$ implies that $r \in I$. An *attracting* interval I satisfies the additional condition that $p \in I$, and $p \succ q$ implies that $q \in I$. The set of intervals on $(\mathcal{P}, \succ)$ will be denoted by $\mathcal{I}(\mathcal{P}, \succ)$ and the set of attracting intervals by $\mathcal{A}(\mathcal{P}, \succ)$.

Proposition 2.3. Let $I \in \mathcal{I}(\mathcal{P}, \succ)$ and define

$$M(I) = \left(\bigcup_{p \in I} M(p) \right) \cup \left(\bigcup_{p,q \in I} C(M(p), M(q)) \right).$$

Then, $M(I)$ is an isolated invariant set.

The proof is fairly straightforward and follows from the compactness of S. The following proposition shows that a given Morse decomposition can give rise to a coarser Morse decomposition.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Proposition 2.4. $\mathcal{M}(S) = \{M(p) \mid p \in \mathcal{P}\backslash I\} \cup \{M(I)\}$ defines a Morse decomposition of S. Furthermore, an admissible partial order $\succ'$ is given by

$$\begin{aligned} p \succ' q \quad &\Leftrightarrow \quad p \succ q \ \text{ if } p, q \in \mathcal{P}\backslash I, \\ p \succ' I \quad &\text{if} \quad \text{there exists } q \in I \text{ such that } p \succ q, \\ I \succ' p \quad &\text{if} \quad \text{there exists } q \in I \text{ such that } q \succ p. \end{aligned}$$

Bibliographical notes

The results described in this section are by now classical. Conley's proof of the decomposition theorem in the setting of flows can be found in Conley (1978). This is also an excellent reference for Morse decompositions. For another proof and further references see Robinson (1995). The presentation in Easton (1998) (see also Norton (1989)) is closest in spirit to that of this paper.

2.2. Isolating blocks

As was mentioned in the Introduction, isolating blocks provide us with a means of localizing the dynamics of interest. The most important property of an isolating block is that it is robust with respect to perturbation.

Proposition 2.5. Let $f, g : X \to X$ be continuous functions. Assume N is an isolating block for f. Then there exists $\epsilon > 0$ such that, if $\|f - g\| < \epsilon$, then N is an isolating block for g.

The proof follows directly from the continuity of f and g and the compactness of N. From the perspective of numerics, this suggests that isolating blocks are stable with respect to numerical errors. The problem is that, to do rigorous computations, we would like to have a sufficient ϵ *a priori*, so that we know what error tolerance can be permitted. We will return to this issue in the next section; for the moment we will present a very simple example to indicate that verification of the existence of an isolating block is possible even in the presence of sizeable errors.

Example 2. Consider the linear map $f : \mathbb{R}^2 \to \mathbb{R}^2$ given by

$$f(x) = \begin{bmatrix} 2 & 0 \\ 0 & \frac{1}{3} \end{bmatrix} x, \quad x \in \mathbb{R}^2. \tag{2.3}$$

Obviously, the origin $(0, 0)$ is a fixed point and the square $N = [-1, 1]^2$ is an isolating block with $(0, 0) = \mathrm{Inv}(N, f)$. Let $g : \mathbb{R}^2 \to \mathbb{R}^2$ satisfy $||g(x) - f(x)|| \leq \frac{1}{2}$ for all $x \in N$. Then N is an isolating block for g.

While this trivial example shows that isolating blocks persist in the presence of fairly large perturbations, it does not address the question of how to find them. Again, we will address this issue from an algorithmic point of

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

view later; however, the following proposition suggests that, viewed in the context of global dynamics, they are natural objects.

Proposition 2.6. Let S be an invariant set under f. Let $\mathcal{R}_p(S, f, \epsilon)$ be an equivalence class of the ϵ-chain-recurrent set. Then $\mathrm{cl}(\mathcal{R}_p(S, f, \epsilon))$ is an isolating block.

Bibliographical notes
There is a variety of references on isolating blocks, beginning with Conley and Easton (1971) for flows, and Easton (1975). Explicit relations between ϵ-chain recurrence and isolation can be found in Norton (1989) and Easton (1989, 1998).

2.3. Multivalued maps

As was made clear in the previous section, the notion of an ϵ-chain is fundamental in dynamical systems. Unfortunately, the classical definition sheds little light on how to systematically compute chain-recurrent sets. With this in mind, let us change our perspective slightly. Consider a point $x \in X$ and a continuous map f. Let $\sigma_x : \{0, \ldots, n\} \to X$ be an ϵ-chain starting at x. By definition, this implies that $\sigma_x(1) \in B_\epsilon(f(x))$ and, more generally, that $\sigma_x(n+1) \in B_\epsilon(f(\sigma_x(n)))$. This suggests that multivalued or set-valued maps provide a convenient language in which to describe the set of possible elements in an ϵ-chain.

Definition 7. A *multivalued map* $F : X \rightrightarrows X$ is a function whose values are subsets of X, *i.e.*, for every $x \in X$, $F(x) \subset X$. Because X is a topological space, we shall on occasion refer to F as a *topological* multivalued map. A *continuous selector* of F is a continuous function $g : X \to X$ for which $g(x) \in F(x)$ for all $x \in X$.

We will treat F as a dynamical system. A *trajectory of F through x* is a function $\mu_x : I \to X$ defined on an interval $I \subset \mathbb{Z}$ and satisfying the following conditions:

(1) $\mu_x(0) = x$,
(2) $\mu_x(n+1) \in F(\mu_x(n))$.

If $I = \mathbb{Z}$, then μ_x is a *full trajectory*. $S \subset X$ is an *invariant set* of F if for every $x \in S$ there exists a full trajectory $\mu_x : \mathbb{Z} \to S$. A compact set $N \subset X$ is an *isolating block* for F if

$$F^{*-1}(N) \cap N \cap F(N) \subset \mathrm{int}(N),$$

where $F^{*-1}(N) := \{x \in X \mid F(x) \cap N \neq \emptyset\}$. At the risk of being redundant, observe that, if $x \in \partial N$, then either

$$F(x) \cap N = \emptyset \qquad \text{or} \qquad F^{*-1}(x) \cap N = \emptyset. \tag{2.4}$$

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

An invariant set S is *isolated* if there exists an isolating block N such that

$$S = \mathrm{Inv}\,(N, F) := \{x \mid \text{there exists } \mu_x : \mathbb{Z} \to N\}.$$

Example 3. In the context of this paper, the most natural way to generate a multivalued map is to begin with a continuous map $f : X \to X$ and define $F : X \rightrightarrows X$ by

$$x \mapsto B_\epsilon(f(x)).$$

The value of ϵ can be chosen to be an upper bound on the round-off errors that arise in the evaluation of the map f. Observe that $\mathcal{R}(X, f, \epsilon)$ is the same as the set of $x \in X$ for which there exists a periodic trajectory $\mu_x : \mathbb{Z} \to X$ of F. Thus, the set of periodic trajectories of F determines a Morse decomposition of f. In fact, if one assumes that ϵ is determined by the round-off error, then the resulting Morse decomposition is the finest that can be obtained given the specific level of numerical accuracy.

It is also easy to check that, if N is an isolating block for F, then by (2.4) N is an isolating block for f. However, the following observation is more important. The same argument also shows that N is an isolating block for any continuous selector g of F. In particular, if we assume that F was obtained numerically, then the values of the map f are not known precisely. Therefore, it is essential that the information that can be extracted from F be valid for all possible continuous selectors.

While this example is supposed to be suggestive of the value of using multivalued maps, it should be clear that, as presented, it is still not a computationally effective tool. In particular, the values of F differ at each point $x \in X$ and, hence, F is not a combinatorial object. To overcome this problem we discretize the phase space.

Definition 8. Let $X \subset \mathbb{R}^n$ and let $\mathcal{G} := \{G_i \subset \mathbb{R} \mid i \in \mathcal{I}\}$. $\mathcal{G}$ is a *grid covering* X if the following conditions are satisfied.

(1) $X \subset \bigcup_{i \in \mathcal{I}} G_i$.
(2) For every compact set $K \subset X$, there are only finitely many elements of $\mathcal{G}$ which intersect K nontrivially.
(3) For every $G_i \in \mathcal{G}$,
$$\emptyset \neq G_i = \mathrm{cl}\,(\mathrm{int}\,(G_i)).$$
(4) If $G_i \neq G_j$, then $\mathrm{int}\,(G_i) \cap \mathrm{int}\,(G_j) = \emptyset$.

$\mathcal{G}$ is a *convex* grid if each grid element G_i is a convex set.

Example 4. The classical example of a convex grid comes from a simplicial complex $\mathcal{K}$ in $\mathbb{R}^n$. $\mathcal{K}$ is a *full simplicial complex* if every simplex in $\mathcal{K}$ is the face of an n-dimensional simplex. Let

$$\mathcal{K}^{(l)} := \{K \in \mathcal{K} \mid \dim K = l\}.$$

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Set $\mathcal{S} := \mathcal{K}^{(n)}$. Then, $\mathcal{S}$ is a *simplicial grid*.

We can generalize this as follows. Define a *polygon* P to be a connected set obtained from the union of a set of elements of $\mathcal{K}^{(n)}$ that share $n-1$-dimensional faces. The elements of a *polygonal grid* are polygons.

Example 5. From the perspective of data structures, a particularly nice grid is the cubical grid. Consider, for example, $X = \prod_{i=1}^{n}[0,1]$. We can define a sequence of grids by

$$\mathcal{G}^{(s)} := \left\{ \prod_{i=1}^{n} \left[\frac{i_k}{2^s}, \frac{i_k+1}{2^s} \right] \;\middle|\; i_k \in \{0, \ldots, 2^s - 1\} \right\}.$$

Definition 9. A *combinatorial multivalued map* is a multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ defined on a finite grid $\mathcal{G}$.

Again, we wish to view $\mathcal{F}$ as a dynamical system, which leads to the following definitions.

Definition 10. A *full trajectory* through $G \in \mathcal{G}$ of a combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ is a function $\gamma_G : \mathbb{Z} \to \mathcal{G}$ satisfying:

(1) $\gamma_G(0) = G$,

(2) for every $n \in \mathbb{Z}$, $\gamma_G(n+1) \in \mathcal{F}(\gamma_G(n))$.

Let $\mathcal{U} \subset \mathcal{G}$. The *maximal invariant set* in $\mathcal{U}$ is given by

$$\mathrm{Inv}\,(\mathcal{U}, \mathcal{F}) := \{G \in \mathcal{U} \mid \text{there exists } \gamma_G : \mathbb{Z} \to \mathcal{U}\}.$$

There are two ways to view the grid $\mathcal{G}$. The first is as a collection of subsets of $\mathbb{R}^n$, that is, $\mathcal{G} = \{G_i \mid G_i \subset X\}$. In this case each grid element G_i carries topological information. The second is simply as a finite list of elements $\mathcal{G} = \{G_i\}$. This is the data that the computer can manipulate. We will adopt both perspectives. The latter is used to develop algorithms, while the former is used to give mathematical interpretation to the computations. In practice, we begin with the set X and define a grid $\mathcal{G} = \{G_i \mid G_i \subset X\}$. However, it is the list $\mathcal{G} = \{G_i\}$ that is entered into the computer (the computer knows no topology) and the output typically contains a sublist of $\mathcal{G}$. To use this output to draw conclusions about the dynamics will require re-introducing the topological structure of the grid elements. To emphasize this we will talk about the *support* of an element of the list, which in an abuse of notation is denoted by $|G_i| = G_i \subset X$. Similarly, given a set of elements of the list, $\mathcal{U} \subset \mathcal{G}$, the support of $\mathcal{U}$ is

$$|\mathcal{U}| := \bigcup_{G \in \mathcal{U}} |G| \subset X.$$

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

A combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ on a grid $\mathcal{G}$ (here $\mathcal{G}$ is viewed as a list) induces a topological multivalued map $F : X \rightrightarrows X$ by

$$F(x) := \bigcup_{x \in G} |\mathcal{F}(G)|. \tag{2.5}$$

Using this relationship we can impose geometric and topological constraints on $\mathcal{F}$. For example, the combinatorial multivalued map $\mathcal{F}$ is *convex* if $F(x)$ is convex for every $x \in X$.

Beginning with an arbitrary topological multivalued map, it is not evident that there is a simple procedure to produce a combinatorial multivalued map. Fortunately, we do not need to work on this level of generality. We are only interested in combinatorial representations of continuous functions. Furthermore, we are interested in combinatorial representations that will lead to isolating blocks, which is why the following condition is imposed.

Definition 11. Let $\mathcal{G}$ be a grid covering X. A combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ is an *outer approximation* of the continuous map $f : X \to X$ if, for every $G \in \mathcal{G}$,

$$\bigcup_{x \in G} f(x) \subset \operatorname{int}(|\mathcal{F}(G)|).$$

Observe that, if $\mathcal{F}$ is generated via rigorous numerical computations, then this condition is almost automatically satisfied.

We also need to be able to recognize on the combinatorial level when we have obtained an isolating neighbourhood, or better, an isolating block. To do this we need to make sense of neighbourhoods on the grid level. Let $\mathcal{U} \subset \mathcal{G}$,

$$o(\mathcal{U}) := \{G \in \mathcal{G} \mid G \cap |\mathcal{U}| \neq \emptyset\}.$$

Observe that $|o(\mathcal{U})|$ is the smallest neighbourhood of $|\mathcal{U}|$ that can be represented using elements of the grid $\mathcal{G}$. Let

$$d(\mathcal{U}) := o(\mathcal{U}) \setminus \mathcal{U}.$$

Remark 1. Recall that in the Introduction we made the assumption that $X \subset \mathbb{R}^n$ was an isolating neighbourhood under f. Since we have discretized the phase space we need to strengthen this assumption. From now on we assume the following:

(A) X is a neighbourhood of $|\operatorname{Inv}(\mathcal{G}, \mathcal{F})|$ relative to $\mathbb{R}^n$.

Theorem 2.7. Let $\mathcal{U} \subset \mathcal{G}$. If $\operatorname{Inv}(o(\mathcal{U}), \mathcal{F}) \subset \mathcal{U}$, then $|\mathcal{U}|$ is an isolating neighbourhood for f.

In Section 2.1 we asserted that Morse decompositions are the natural global objects to study. To see how they arise in the context of combinatorial

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

multivalued maps we only need to recast the definitions. Set

$$\mathcal{R}(\mathcal{G},\mathcal{F}) := \{G \in \mathcal{G} \mid \text{there exists a periodic trajectory } \gamma_G : \mathbb{Z} \to \mathcal{G}\}.$$

We can decompose $\mathcal{R}(\mathcal{G},\mathcal{F})$ into equivalence classes $\mathcal{R}_p(\mathcal{G},\mathcal{F})$, $p \in \mathcal{P}$, by setting $G \sim G'$ if there exists a periodic trajectory $\gamma_G : \mathbb{Z} \to \mathcal{G}$ such that $\gamma_G(j) = G'$ for some $j \in \mathbb{Z}$. Similarly, we can define an admissible order $\succ$ on $\mathcal{P}$.

Theorem 2.8. Let $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ be a combinatorial multivalued map which is an outer approximation of $f : X \to \mathbb{R}^n$. Let $M(p) := \mathrm{Inv}(\mathcal{R}_p(\mathcal{G},\mathcal{F}), f)$, $p \in \mathcal{P}$, and let $\succ$ be an admissible order for $\mathcal{P}$. Then, $\{M(p) \mid p \in \mathcal{P}\}$ is a Morse decomposition for f with admissible order $\succ$.

The final point that needs to be discussed is how well continuous maps can be approximated by combinatorial multivalued maps. This leads to the following definition.

Definition 12. Let $\mathcal{G}^{(n)}$ be a convex grid covering X; let $\mathcal{F}^{(n)} : \mathcal{G}^{(n)} \rightrightarrows \mathcal{G}^{(n)}$ be a combinatorial multivalued map. Set $F^{(n)} : X \rightrightarrows X$ to be the topological multivalued map satisfying (2.5). A sequence of outer approximations $\mathcal{F}^{(n)}$ *converges* to f if, for each n, $\mathcal{F}^{(n)}$ is an outer approximation of f and for every $\epsilon > 0$ there exists an integer N such that $n \geq N$ and $y \in F^{(n)}(x)$ implies that $\|f(x) - y\| < \epsilon$.

Theorem 2.9. If $f : X \to X$ is Lipschitz-continuous, then there exists a sequence of convex grids $\mathcal{G}^{(n)}$ and convex combinatorial multivalued maps $\mathcal{F}^{(n)} : \mathcal{G}^{(n)} \rightrightarrows \mathcal{G}^{(n)}$ converging to f.

Bibliographical notes

The use of multivalued maps as a means of translating between the combinatorial structures used to compute and the continuous dynamics of interest was first used by Mrozek and the author to provide a computer-assisted proof of chaotic dynamics in the Lorenz equation (Mischaikow and Mrozek 1995, 1998). For a systematic discussion of the use of combinatorial multivalued maps to understand continuous functions see Mrozek (1996). A proof of Theorem 2.7 can be found in Szymczak (1997). Assumption (A) can be removed but then additional (verifiable) conditions need to be checked to obtain the conclusions of Theorems 2.7 and 2.8. See Szymczak (1997) and Boczko, Kalies and Mischaikow (2002).

2.4. Graph algorithms

We now turn to the problem of determining the global structure of the dynamics. The general strategy is as follows. We begin with a continuous map $f : X \to X$. We choose a grid $\mathcal{G}$ that covers X and compute a combinatorial

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ that is an outer approximation of f. The goal is to determine either the global structure of the dynamics, that is, to compute a Morse decomposition, or to find an isolating block. To do this efficiently we make use of classical graph algorithms. To explain these we begin by introducing some standard definitions.

Definition 13. A *directed graph* is a pair $(\mathcal{V}, \mathcal{E})$, where the collection of *vertices* $\mathcal{V}$ consists of a finite set and the *edges* $\mathcal{E}$ are ordered pairs of vertices from $\mathcal{V}$. The number of elements in $\mathcal{V}$ and $\mathcal{E}$ are denoted by $|\mathcal{V}|$ and $|\mathcal{E}|$, respectively. Furthermore,

$$\mathcal{E}^T := \{(u, v) \mid (v, u) \in \mathcal{E}\}.$$

Observe that the combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ defines a directed graph $(\mathcal{V}, \mathcal{E})$ as follows. Set $\mathcal{V} = \mathcal{G}$ and let

$$(G_i, G_j) \in \mathcal{E} \quad \Leftrightarrow \quad G_j \in \mathcal{F}(G_i).$$

Throughout this paper we shall always assume that the directed graph $(\mathcal{V}, \mathcal{E})$ was generated by a combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ that is an outer approximation of $f : X \to X$.

Definition 14. A *path* in a directed graph $(\mathcal{V}, \mathcal{E})$ is a function $\gamma_G : \{0, \ldots, n\} \to \mathcal{V}$ with the property that $(\gamma_G(i), \gamma_G(i+1)) \in \mathcal{E}$ for all $i = 0, \ldots, n-1$; γ_{G_i} is a path from G_i to G_j if $\gamma_{G_i}(0) = G_i$ and $\gamma_{G_i}(n) = G_j$. It is a *cycle* if $\gamma_{G_i}(0) = \gamma_{G_i}(n)$.

Observe that a path in the directed graph $(\mathcal{V}, \mathcal{E})$ is equivalent to a trajectory of $\mathcal{F}$.

Because we are interested in using these directed graphs to perform the computations, it is important to consider how this information is represented. An obvious method is through an adjacency matrix $E = [e_{i,j}]$, where $e_{i,j} = 1$ if and only if $(G_j, G_i) \in \mathcal{E}$. This, of course, entails a cost of $|\mathcal{V}|^2$. However, since our directed graphs are generated by combinatorial multivalued maps which are meant to be reasonable outer approximations of a continuous function, one expects that there are only a few edges associated with any given vertex. This in turn implies that the matrix E is sparse. Therefore, for applications it is probably preferable to represent the graph as an adjacency list.

Definition 15. An *adjacency list representation* of a directed graph $(\mathcal{V}, \mathcal{E})$ is an array of $|\mathcal{V}|$ lists, one list for each vertex. For each vertex G, the list consists of all vertices such that $(G, G') \in \mathcal{E}$.

Observe that the size of an adjacency list representation is given by $|\mathcal{E}|$.

In what follows we assume that an adjacency list representation is used. This implies that we have an ordering of the vertices $\mathcal{V} = \{G_i \mid i = 1, \ldots, |\mathcal{V}|\}$

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

and then, for each vertex G_i, an ordering of the edges in the list associated with G_i.

Global dynamics

As was demonstrated in the previous subsection, a Morse decomposition for f can be obtained by determining the components of the recurrent set of $\mathcal{F}$. A standard definition from graph theory that nearly corresponds to components of recurrent sets is as follows.

Definition 16. A *strongly connected component* (SCC) of $(\mathcal{V}, \mathcal{E})$ is a maximal set of vertices $\mathcal{U} \subset \mathcal{V}$ such that, for every pair $G_i, G_j \in \mathcal{U}$, there exist paths from G_i to G_j and from G_j to G_i.

We can restate Proposition 2.6 as follows.

Theorem 2.10. Every SCC is an isolating block for f.

The problem with this definition is that, if $G \in \mathcal{V}$ and does not belong to any cycle, then vacuously $\{G\}$ defines a strongly connected component. For this reason we introduce a slightly stronger condition.

Definition 17. A *strongly connected path component* (SCPC) of $(\mathcal{V}, \mathcal{E})$ is a strongly connected component that contains at least one edge.

Example 6. Consider the directed graph $(\mathcal{V}, \mathcal{E})$ of Figure 2.1. The set of strongly connected components consists of

$$\{\{a\}, \{c, d, g\}, \{h\}, \{e\}, \{f\}, \{b\}\},$$

while the set of strongly connected path components consists of

$$\{\{c, d, g\}, \{f\}\}.$$

Observe that, if $(\mathcal{V}, \mathcal{E})$ was generated by a combinatorial multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$, then the components of the chain-recurrent set would correspond to the strongly connected path components. Furthermore, the maximal invariant set consists of $\{c, d, g, h, f\}$.

The following theorem says that Morse decompositions are the optimal global objects to compute.

Theorem 2.11. Let $(\mathcal{V}, \mathcal{G})$ be a directed graph. There is a linear-time algorithm (more precisely, the computation time is proportional to the number of vertices and edges, *i.e.*, $|\mathcal{V}| + |\mathcal{E}|$) that identifies the strongly connected path components

$$\{\mathcal{R}_p \subset \mathcal{V} \mid p = 1, \ldots n\},$$

and produces a function

$$L : \mathcal{V} \to \mathbb{Z}$$

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

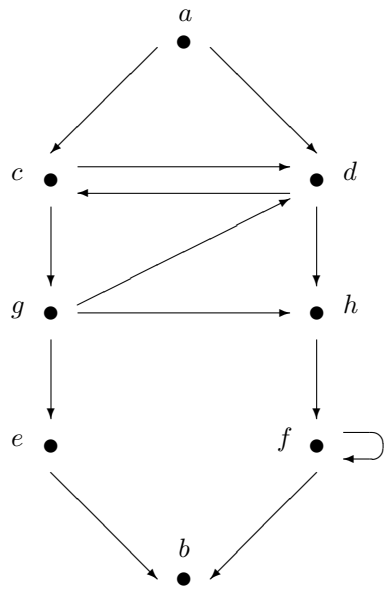


Figure 2.1. Directed graph $(\mathcal{V}, \mathcal{E})$

with the property that L is constant on each strongly connected path component and, if $G \in \mathcal{V} \setminus \cup_{p=1}^{n} \mathcal{R}_p$ and $(G, G') \in \mathcal{E}$, then

$$L(G) > L(G').$$

We shall not give a complete proof here, but, instead, briefly describe the depth-first search (DFS) algorithm which is the key step.

Depth-first search is a fairly simple recursive algorithm. To implement it we need a counter ι which takes integer values. We will also assign to each vertex G three values. The first value is $\xi(G) \in \{0, 1, 2\}$. If $\xi(G) = 0$, this indicates that the vertex G has not yet been viewed by the algorithm. The second value is an integer $\alpha(G)$ that marks the initial step at which the vertex is viewed. The third value is, also, an integer $\omega(G)$ which marks the final step at which the vertex is viewed. Finally, we can keep an ordered list of unviewed vertices, that is, the set of vertices for which $\xi = 0$. At the beginning this list is all of $\mathcal{V}$; however, as each vertex is viewed it is removed from this list.

Recall that we are assuming an adjacency list representation of $(\mathcal{V}, \mathcal{E})$ and, hence, have an ordering on the vertices and on the edges from each vertex. Initially, $\iota = 0$, and for each $G_i \in \mathcal{V}$,

$$\xi(G_i) = 0, \alpha(G_i) = 0, \omega(G_i) = 0.$$

The first rule is that ι is increased by one if and only if $\xi(G)$ is changed for some $G \in \mathcal{V}$.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

The algorithm begins on the first vertex G_1. $\xi(G_1) = 0$, which means that this is the first time that the vertex has been viewed. Reset $\xi(G_1) = 1$, augment the counter to $\iota = 1$, and set $\alpha(G_1) = \iota$. Let (G_1, G_i) be the first edge in the list associated to G_1. View G_i. (Observe that if $i \neq 1$ then $\xi(G_i) = 0$.)

We now describe a general recursive step. Assume that $\iota = k$, and we view a vertex G_i such that $\xi(G_i) = 0$. Then, reset $\xi(G_i) = 1$, augment the counter to $\iota = k+1$, and set $\alpha(G_i) = \iota$. Let (G_i, G_j) be the first edge in the list associated to G_i and view G_j.

An alternative is that $\iota = k$, and we view a vertex G_i such that $\xi(G_i) \in \{1, 2\}$. Observe that to be in this situation we must have arrived at G_i by considering an edge (G_l, G_i). In this case we return to the list of edges from G_l. We did not change the value of any $\xi(G)$, therefore $\iota = k$. As was remarked before, the edges from G_l are in an ordered list, so we consider the edge (G_l, G_j) which follows the edge (G_l, G_i) in that list. View G_j. (Notice that if we are performing this step then $\xi(G_l) = 1$.)

It is possible that the edge (G_l, G_i) is the last in the list. In this case, set $\xi(G_l) = 2$, augment the counter to $\iota = k + 1$, and set $\omega(G_l) = \iota$. There are now two cases to consider. The first is that we first viewed G_l via an edge (G_m, G_l). In this case view G_m. Otherwise view the first vertex on the list of unviewed vertices.

We will add one additional action to this algorithm. Observe that one of the ways we move through the vertices is by viewing edges (G_i, G_j) where $\xi(G_j) = 0$. If this is the case we label this edge as a *forward edge*. Let $\mathcal{E}^f \subset \mathcal{E}$ denote the set of forward edges obtained by running DFS. It is easy to check that the directed graph $(\mathcal{V}, \mathcal{E}^f)$ is a forest, that is, a disjoint collection of trees. Observe that we began constructing a tree within this forest each time we chose a vertex from the list of unviewed vertices. We declare this vertex to be the root of the tree. Thus, having run DFS, to each $G \in \mathcal{V}$ we can assign a unique vertex $\rho(G)$ that denotes the root of the tree to which it belongs. Notice that $\rho(G)$ has the lowest α value and largest ω value of any vertex G in the tree.

If we apply DFS to a directed graph $(\mathcal{V}, \mathcal{E})$ the outputs of interest are the values $\{\alpha(G_i), \omega(G_i) \mid G_i \in \mathcal{V}\}$, all of which are distinct, and the trees of $(\mathcal{V}, \mathcal{E}^f)$.

The following algorithm, which we shall refer to as SCC, satisfies Theorem 2.11.

(1) Apply DFS to the directed graph $(\mathcal{V}, \mathcal{E})$ to compute $\{\omega(G) \mid G \in \mathcal{V}\}$.
(2) Compute $(\mathcal{V}, \mathcal{E}^T)$.
(3) Re-order the vertices of $\mathcal{V}$ by decreasing values of $\omega(G)$. Using this order, apply DFS to the directed graph $(\mathcal{V}, \mathcal{E}^T)$. This computes a new set of values $\{\omega(G) \mid G \in \mathcal{V}\}$ and a new set of trees.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

(4) Each tree of $(\mathcal{V}, (\mathcal{E}^T)^f)$ is a strongly connected component of $(\mathcal{V}, \mathcal{E})$. Define

$$L(G) = |\mathcal{V}| - \omega(\rho(G)).$$

Since this algorithm produces the trees associated with the strongly connected components, we can use this information to produce a new directed graph $(\mathcal{V}^{SCC}, \mathcal{E}^{SCC})$ called the *component graph*, which is defined as follows. Let $\{\mathcal{C}_1, \ldots, \mathcal{C}_k\}$ be the set of strongly connected components of $(\mathcal{V}, \mathcal{G})$. Set $\mathcal{V}^{SCC} = \{C_1, \ldots, C_k\}$, that is, one vertex for each component. $(C_i, C_j) \in \mathcal{E}^{SCC}$ if there exists $G \in \mathcal{C}_i$, $G' \in \mathcal{C}_j$, and $(G, G') \in \mathcal{E}$.

This construction of the component graph demands a brief digression back to the realm of continuous dynamics. Conley (1978) defined a *chain-recurrent flow* (*map*) to be a flow (map) whose chain-recurrent set was the entire phase space. At the other extreme, he declared a flow (map) to be *strongly gradient-like* if the chain-recurrent set is totally disconnected and consists only of equilibria. Using this terminology he then recast his Fundamental Decomposition Theorem into the following statement:

> Every flow on a compact space is uniquely represented as the extension of a chain recurrent flow by a strongly gradient-like flow; that is the flow admits a unique subflow which is chain recurrent and such that the quotient flow is strongly gradient-like.

(The same is true for continuous maps.) The quotient flow is obtained by collapsing each connected component of the chain-recurrent set to a distinct point. Observe that the procedure for constructing the component graph is the discrete analogue of the construction of the strongly gradient-like dynamical system. The point which is meant to be emphasized is that there is a natural correspondence between Conley's approach to decomposing continuous dynamical systems and the basic decompositions of directed graphs.

It was observed earlier that paths in $(\mathcal{V}, \mathcal{E})$ correspond to trajectories of $\mathcal{F}$. Of course, from the perspective of dynamics we are interested in full trajectories of $\mathcal{F}$. The following proposition characterizes those vertices of $\mathcal{V}$ for which one can define a path $\gamma_G : \mathbb{Z} \to \mathcal{V}$ which corresponds to a full trajectory.

Proposition 2.12. Let $(\mathcal{V}, \mathcal{E})$ be a directed graph with strongly connected path components $\{\mathcal{R}_p \mid p \in \mathcal{P}\}$. Let $G \in \mathcal{V}$. There exists a path $\gamma_G : \mathbb{Z} \to \mathcal{V}$ if and only if $G \in \cup_{p\in\mathcal{P}}\mathcal{R}_p$, or there exists a path $\gamma_{G_0} : \{0, \ldots, n\} \to \mathcal{V}$ from G_0 to G_1 such that $G_i \in \cup_{p\in\mathcal{P}}\mathcal{R}_p$ and $\gamma_{G_0}(j) = G$ for some $j \in \{0, \ldots, n\}$.

Isolation within recurrent sets

Of considerable interest in nonlinear systems is the existence of complicated or chaotic dynamics. Of course, this behaviour is recurrent and will therefore be found within the strongly connected path components. As was indicated

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

in the Introduction we use the Conley index theory to extract lower bounds on the structure of the dynamics of f. The details of this will be discussed in the next section. For the moment, we remind the reader that effective computation of the index depends upon isolating blocks. In particular, the more 'structure' that the blocks possess, the more information about the dynamics can be extracted. Therefore, in this subsection we will discuss two different approaches for finding isolating blocks within strongly connected path components. The first begins with the full SCPC and reduces the number of grid elements in the complex. The second approach adopts the opposite strategy. One begins with a minimal number of grid elements which may possess a specific dynamics, and then adds grid elements until isolation is achieved.

Essential to both approaches is the ability to compute invariant sets for combinatorial multivalued maps. More precisely, given $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ and $\mathcal{U} \subset \mathcal{G}$, we need to be able to determine $\mathrm{Inv}(\mathcal{U}, \mathcal{F})$. As before, we let $(\mathcal{V}, \mathcal{E})$ be the directed graph associated to $\mathcal{F}$. Let $(\mathcal{U}, \mathcal{E}')$ be the subgraph of $(\mathcal{V}, \mathcal{E})$ obtained by restricting to those vertices associated with $\mathcal{U}$.

Proposition 2.13. $\mathrm{Inv}(\mathcal{U}, \mathcal{F})$ can be computed in linear time: more precisely, in time proportional to the number of vertices and edges in $(\mathcal{U}, \mathcal{E}')$.

As before, we will not provide a formal proof of this proposition, but, instead, briefly describe an algorithm that will perform the task. Furthermore, the algorithm we describe is presented to emphasize a graph-theoretic counterpart to the Fundamental Decomposition Theorem: a vertex is in the maximal invariant set if and only if it belongs to a SCPC or it lies on a path from one SCPC to another SCPC. The first step is to run $\mathrm{SCC}(\mathcal{U}, \mathcal{E}')$ and create $(\mathcal{U}^{SCC}, \mathcal{E}'^{SCC})$. Within $\mathcal{U}^{SCC}$, identify those which correspond to the strongly connected path components of $(\mathcal{U}, \mathcal{E}')$ (these are the vertices which have an edge to themselves). Now apply the following modified version of DFS. To each vertex $G \in \mathcal{U}^{SCC}$ we assign an additional value $\zeta(G) \in \{0, 1\}$. Initially, $\zeta(G) = 0$, unless G corresponds to a strongly connected path component, in which case $\zeta(G) = 1$. This modified DFS is only initiated on vertices that correspond to strongly connected path components. There are two occasions when $\zeta(G)$ may be changed. The first is that we are at a vertex G_i and considering an edge (G_i, G_j) where $\zeta(G_j) = 1$. Then, set $\zeta(G_i) = 1$. The second is that $\zeta(G_l) = 1$, and we first viewed G_l via an edge (G_m, G_l). Then, set $\zeta(G_m) = 1$.

At the completion of this modified DFS, $\mathrm{Inv}(\mathcal{U}, \mathcal{F})$ is given by the set of vertices for which $\zeta = 1$.

Throughout the remainder of this subsection, we assume that $(\mathcal{V}, \mathcal{E})$ is a directed graph that has a unique strongly connected component.

The first approach is based on the following theorem.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Theorem 2.14. Let $\mathcal{U} \subset \mathcal{V}$. Let $\mathcal{U}_0 := \mathcal{U}$. For $i = 1, \ldots, m$, let $\mathcal{U}_i, \mathcal{W}_i \subset \mathcal{V}$, which satisfies the following conditions:

(1) $\mathrm{Inv}\,(\mathcal{U}_{i-1}, \mathcal{F}) = \mathcal{U}_i \cup \mathcal{W}_i$,
(2) $o(\mathcal{U}_i) \cap \mathcal{W}_i = \emptyset$,
(3) $o(\mathcal{U}_m) \subset \mathcal{U}$.

Then, $\mathcal{U}_m = \mathrm{Inv}\,(o(\mathcal{U}_m) \cup \mathcal{W}_m) \cap o(\mathcal{U}_m)$.

Using Theorem 2.14 there is a variety of specific algorithms that can be created (see Szymczak (1997, Section 4)). The essential point is to generate the sets $\mathcal{U}_i$ and $\mathcal{W}_i$ inductively in such a way that conditions (1) and (2) are satisfied. The algorithm then halts when condition (3) is satisfied. Observe that in this method we begin with a 'large' set of vertices, *i.e.*, $\mathcal{U}$, and at each step we reduce the set of vertices being considered. To emphasize this point, assume that $|\mathcal{V}|$ is a connected set and that we choose $\mathcal{U} = \mathcal{V}$. Then $\mathcal{U} = \mathrm{Inv}\,(U_0, \mathcal{F}) = \mathcal{U}_1 \cup \mathcal{W}_1$. Therefore, because of condition (2), $\mathcal{W}_1 = \emptyset$ and $\mathcal{U}_1 = U_0$. In this case, no progress has been made.

In principle, these algorithms can be applied with no knowledge of paths within $(\mathcal{V}, \mathcal{E})$. The next strategy takes the opposite approach. Consider Figure 2.2. This is meant to represent that within $(\mathcal{V}, \mathcal{E})$ we have identified a cycle of length one and a cycle of length two. One might suspect that $|G_3|$ contains a fixed point of f and $|G_1| \cup |G_2|$ contains a period two orbit of f. (In the next section we will provide theorems that allow one to check such claims rigorously.) However, since $(\mathcal{V}, \mathcal{E})$ has a unique strongly connected

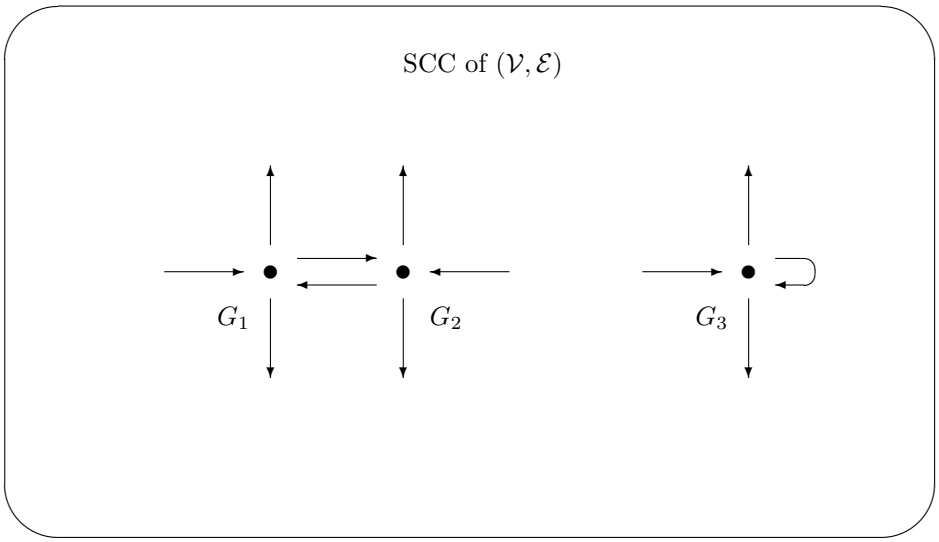


Figure 2.2. A strongly connected component which contains a period one path and a period two path

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

component, we know that there are paths from G_3 to G_1 and from G_1 to G_3. This in turn suggests the existence of horseshoe or symbolic dynamics based on trajectories of f that pass through $|G_1| \cup |G_2|$ and $|G_3|$. To prove this requires an appropriate isolating neighbourhood and is the purpose of the next algorithm.

The first step of the process is to generate an initial guess for the isolating block that will capture the desired dynamics. Let $\mathcal{B}_i \subset \mathcal{V}$, $i = 0, 1$. Let $(\mathcal{B}_i, \mathcal{E}_i)$ be the corresponding subgraphs of $(\mathcal{V}, \mathcal{E})$. Assume that $(\mathcal{B}_i, \mathcal{E}_i)$ have unique strongly connected components. Let $\gamma_0 : \{0, \ldots, n_0\} \to \mathcal{V}$ be a minimal path from any element of $\mathcal{B}_0$ to any element of $\mathcal{B}_1$ and let $\gamma_1 : \{0, \ldots, n_1\} \to \mathcal{V}$ be a minimal path from any element of $\mathcal{B}_1$ to any element of $\mathcal{B}_0$. Let $\mathcal{W}_i = \gamma_i(\{0, \ldots, n_i\})$ and let

$$\mathcal{U} := \mathcal{B}_0 \cup \mathcal{B}_1 \cup \mathcal{W}_0 \cup \mathcal{W}_1.$$

Proposition 2.15. Under the assumption that the number of elements of $\mathcal{B}_i$ is small compared to the number of elements in $\mathcal{V}$, Dijkstra's algorithm (Corman, Leiserson, Rivest and Stein 2001) will determine $\mathcal{W}_0$ and $\mathcal{W}_1$ with a running time $O(|\mathcal{V}| \ln |\mathcal{V}| + |\mathcal{E}|)$.

Observe that the resulting directed subgraph $(\mathcal{U}, \mathcal{E}')$ has a unique strongly connected component and therefore $\mathrm{Inv}(\mathcal{U}, \mathcal{F}) = \mathcal{U}$. There is no reason *a priori* to believe that $|\mathcal{U}|$ is an isolating block for f. Thus, the next step of the process is to isolate the dynamics we have captured. Let $\mathcal{U}_0 := \mathcal{U}$. Given $\mathcal{U}_i$, define $\mathcal{X}_i = o(\mathcal{U}_i)$. Set $\mathcal{U}_{i+1} = \mathrm{Inv}(X_i, \mathcal{F})$. If

$$o(\mathcal{U}_{i+1}) \subset \mathcal{X}_i$$

then $\mathcal{U}_{i+1}$ is an isolating block for f and we stop. If not we repeat the process.

Bibliographical notes

Definitions and explicit descriptions of the graph algorithms can be found in Corman *et al.* (2001). The use of these algorithms to find chain-recurrent sets, isolating neighbourhoods and index pairs (see Section 3) was first implemented by M. Eidenschink (1995) in the context of simplicial approximations of two-dimensional flows (see Section 2.5).

We did not provide any complexity bounds for either algorithm of this subsection. In part this is because there is no guarantee that these algorithms will produce an isolating block. In the first case, it is possible to reduce $\mathcal{U}$ to the empty set. In the second case, it is possible that $\mathcal{U}$ will grow to be the entire SCPC. In practice, however, these methods seem to work well (Szymczak 1997, Allili, Day, Junge and Mischaikow 2002).

The usefulness of these ideas are only as good as the software that can support them. Fortunately, M. Dellnitz, A. Hohmann and O. Junge have

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

independently developed an excellent general purpose package GAIO based on cubical grids (Dellnitz and Hohmann 1997, Dellnitz and Junge 2001, Dellnitz, Froyland and Junge 2000). Some of the above-mentioned graph algorithms are already incorporated in the code. Furthermore, the grid is constructed in an adaptive method which has significant implications for the cost of computing the combinatorial multivalued map. In particular, one can expect that for many problems the cost is determined by the dimension of the invariant set rather than the dimension of the phase space.

2.5. Flows

We now turn to a discussion of computations in the context of flows. Let $\dot{x} = g(x)$ be an ordinary differential equation defined on $\mathbb{R}^n$ that generates a flow $\varphi : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$. Currently, the only systematic approach to problems of this nature is to change it to a problem involving continuous maps. A completely general procedure is to fix a time $\tau > 0$ and define $f : \mathbb{R}^n \to \mathbb{R}^n$ by $f(x) := \varphi(\tau, x)$. The following result indicates that, on a theoretical level, the invariant sets of f which we can hope to recover using these topological techniques are equivalent to those of φ.

Proposition 2.16. Consider a flow $\varphi : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$. Let $\varphi_\tau : \mathbb{R}^n \to \mathbb{R}^n$ be defined by $\varphi_\tau(x) = \varphi(\tau, x)$. Then, the following are equivalent:

(1) S is an isolated invariant set for φ,

(2) S is an isolated invariant set for φ_t for all $t > 0$,

(3) S is an isolated invariant set for φ_t for any $t > 0$.

Unfortunately, while theoretically simple, its implementation is less obvious. In particular, some choice of τ needs to be made and the map f needs to be computed, most probably by numerically integrating the differential equation. Therefore, if τ is large, then the numerical errors will force the combinatorial multivalued map to have exponentially large values. This in turn will make it difficult to find isolating neighbourhoods. On the other hand, if τ is small then f will be a near identity map and an extremely fine grid size will be required. This suggests that one needs to look for an 'optimal' choice of τ. Unfortunately, such a value is probably dependent on the location in phase space. It should be mentioned that, in spite of these negative comments, this method can be employed successfully (Pilarczyk 1999), though at a significant computational price.

A more successful variation of this strategy involves the use of local Poincaré sections for the flow. In this case we study the dynamics of the Poincaré map $f : \Xi \to \Xi$, where Ξ consists of $(n-1)$-dimensional hypersurfaces that are transverse to the flow. Again, f is determined by integrating the differential equation and therefore one is faced with exponential growth in error.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

However, in this setting it is easier to implement multistep methods and multiple hypersurfaces to control the error. Furthermore, and this leads to significant computational savings, the dimension of the grid is reduced by one. This approach has been successfully applied to a variety of low-dimensional problems. Since, aside from the numerical issue of computing f, the technical aspects are similar to the map case, we will not discuss this approach further.

In Example 4 a polygonal grid $\mathcal{T}$ was defined in terms of a full simplicial complex $\mathcal{K}$ in $\mathbb{R}^n$. In particular, a polygon P is a connected set that can written as

$$P = \bigcup_{i=1}^{j} K_i,$$

where $K_i \in \mathcal{K}^{(n)}$. In an abuse of notation we will write $K_i \in P$. Observe that the boundary of P, denoted by ∂P, can be written as the union of elements of $\mathcal{K}^{(n-1)}$, that is,

$$\partial P = \bigcup_{i=1}^{m} L_i,$$

where $L_i \in \mathcal{K}^{(n-1)}$. We will continue the abuse of notation by writing $L_i \in \partial P$.

Given $L \in \mathcal{K}^{(n-1)}$, let

$$\nu(L)$$

denote one of the two unit normal vectors to L. To determine a unique choice of $\nu(L)$, let $K \in \mathcal{K}^{(n)}$ such that L is a face of K. Then,

$$\nu_K(L)$$

is defined to be the inward unit normal of L with respect to K.

Definition 18. $L \in \mathcal{K}^{(n-1)}$ is a *flow-transverse face* with respect to the vector field g, if

$$\nu(L) \cdot g(x) \neq 0$$

for every $x \in L$. A polygon P is *flow-transverse* if every $L \in \partial P$ is flow-transverse.

Definition 19. Let $\mathcal{T} = \{P_1, \ldots, P_N\}$ be a polygonal decomposition of X. $\mathcal{T}$ is a *flow-transverse polygonal decomposition* of X if every polygon P_i is flow-transverse.

As we shall explain, given a simplicial complex, constructing a flow-transverse polygonal decomposition can be done in linear time in the number of n-dimensional simplices. For this discussion we shall assume that X

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

is the polygonal region defined by the full simplicial complex $\mathcal{K}$, and for the sake of simplicity that X is flow-transverse (see Boczko *et al.* (2002) where this latter assumption is dropped). Let $K_i, K_j \in \mathcal{K}^{(n)}$ such that $K_i \cap K_j = L \in \mathcal{K}^{(n-1)}$. Set $K_i \sim K_j$ if $\nu(L) \cdot f(x) = 0$ for some $x \in L$. Extend this relation by transitivity. Then $\sim$ is an equivalence relation on $\mathcal{K}$. Define $\mathcal{T} := \{P_1, \ldots, P_N\}$ to be the polygons defined by the equivalence classes. This grid is a flow-transverse polygonal decomposition of X.

The next step is to define a combinatorial multivalued map $\mathcal{F} : \mathcal{T} \rightrightarrows \mathcal{T}$ that captures the dynamics of the flow. Observe that if P contains recurrent dynamics under the flow φ, then we want $P \in \mathcal{F}(P)$. The simplest way to ensure this is as follows:

$$P \in \mathcal{F}(P) \quad \Leftrightarrow \quad \|g(x)\| < A \text{ for some } x \in P, \tag{2.6}$$

where the constant A is determined by the diameter of P and $\|Dg(x)\|$. Otherwise, set

$$P_i \in \mathcal{F}(P_j) \quad \Leftrightarrow \quad \nu_{P_i}(L) \cdot g(x) > 0, \text{ for some } x \in L \in \partial P_i \cap \partial P_j \cap \mathcal{K}^{(n-1)}. \tag{2.7}$$

Observe that this implies that a polygon can be in the image of another polygon only if they share an $(n-1)$-dimensional simplex. This implies that the directed graph $(\mathcal{V}, \mathcal{E})$ generated by the combinatorial multivalued map $\mathcal{F} : \mathcal{T} \rightrightarrows \mathcal{T}$ is, in general, sparse. At this point we are, of course, free to apply the graph algorithms of the previous subsection. As the following result indicates, this allows us to compute a Morse decomposition[1] efficiently.

Theorem 2.17. Let $M(p) := \mathrm{Inv}\,(\mathcal{R}_p(\mathcal{T}, \mathcal{F}))$, $p \in \mathcal{P}$, and let $\succ$ be an admissible order for $\mathcal{P}$. Then, $\{M(p) \mid p \in \mathcal{P}\}$ is a Morse decomposition for $\mathrm{Inv}\,(X, \varphi)$ with admissible order $\succ$.

The other important structure that we need to be able to compute is an isolating block which we shall define in the context of polygonal approximations.

Definition 20. A polygon P is an *isolating block* if, for every point $x \in \partial P$ and any $\epsilon > 0$,

$$\varphi((-\epsilon, \epsilon), x) \not\subset P.$$

The fact that φ is a smooth flow implies the following result.

Proposition 2.18. If P is a convex flow-transverse polygon, then it is automatically an isolating block.

[1] The definition of a Morse decomposition for a flow is the obvious analogue of Definition 5, or see Conley (1978), Robinson (1995) and Arnold, Jones, Mischaikow and Raugel (1995).

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Theorem 2.19. Let $\mathcal{U}$ be a strongly connected component of $(\mathcal{V}, \mathcal{E})$. Then $|\mathcal{U}|$ is a isolating block for φ.

Thus, we regain two out of the three important computational results. What remains is the question of obtaining arbitrarily good approximations of the dynamics. This depends heavily on the initial simplicial complex and remains a difficult open question in computational geometry.

Bibliographical notes

References to rigorous computational results for ordinary differential equations using Poincaré maps can be found in Mischaikow and Mrozek (2001). The results on flow-transverse polygonal approximations will appear in the forthcoming paper by Boczko *et al.* (2002). Though not presented in this generality, in his thesis, Eidenschink (1995) provided code that provides good simplicial approximations for two-dimensional systems.

3. Conley index

The numerical computations and algorithms discussed in the previous section do not in themselves provide rigorous information concerning the existence and structure of invariant sets. To do this we will invoke the Conley index theory.

3.1. Topological preliminaries

Before providing a definition of the index we recall a few elementary notations from topology.

Definition 21. A *pointed space* (X, x_0) is a topological space X with a distinguished point $x_0 \in X$. A *continuous map between pointed spaces* $f : (X, x_0) \to (Y, y_0)$ is a continuous map $f : X \to Y$ with the additional condition that $f(x_0) = y_0$.

Given a pair (N_1, N_0) of spaces with $N_0 \subset N_1$, set

$$N_1/N_0 := (N_1 \setminus N_0) \cup [N_0],$$

where $[N_0]$ denotes the equivalence class of points in N_0 in the equivalence relation: $x \sim y$ if and only if $x = y$ or $x, y \in N_0$. Hereafter, we will usually use N_1/N_0 to denote the pointed space $(N_1/N_0, [N_0])$. The topology on $(N_1/N_0, [N_0])$ is defined as follows: a set $U \subset N_1/N_0$ is open if U is open in N_1 and $[N_0] \notin U$, or the set $(U \cap (N_1 \setminus N_0)) \cup N_0$ is open in N_1. If $N_0 = \emptyset$, then

$$(N_1/N_0, [N_0]) := (N_1 \cup \{*\}, \{*\}),$$

where $\{*\}$ denotes the equivalence class consisting of the empty set.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Definition 22. Let (X, x_0) and (Y, y_0) be pointed topological spaces and let $f, g : (X, x_0) \to (Y, y_0)$ be continuous functions. f is *homotopic* to g, denoted by

$$f \sim g,$$

if there exists a continuous function $F : X \times [0, 1] \to Y$ such that

$$\begin{aligned} F(x, 0) &= f(x), \\ F(x, 1) &= g(x), \\ F(x_0, s) &= y_0, \quad 0 \le s \le 1. \end{aligned}$$

Obviously $\sim$ is an equivalence relation. The equivalence class of f in this relation is called the *homotopy class* of f and denoted by $[f]$.

Definition 23. Two pointed topological spaces (X, x_0) and (Y, y_0) are *homotopic*

$$(X, x_0) \sim (Y, y_0)$$

if there exists $f : (X, x_0) \to (Y, y_0)$ and $g : (Y, y_0) \to (X, x_0)$ such that

$$f \circ g \sim \mathrm{id}_Y \quad \text{and} \quad g \circ f \sim \mathrm{id}_X.$$

Observe that homotopy defines an equivalence class on the set of topological spaces. Homotopy classes of topological spaces are extremely difficult to work with directly. One of the most useful tools in this area is algebraic topology, in particular, homology theory. As was indicated in the Introduction we use homology to provide a rigorous interpretation of the objects that have been computed numerically. For the purposes of this presentation it is sufficient to know that, given a pair of spaces $N_0 \subset N_1$, which are described in terms of grid elements or a simplicial complex, it is possible to compute the *relative homology groups*

$$H_*(N_1, N_0) = \{H_k(N_1, N_0) \mid k = 0, 1, 2, \ldots\}.$$

For each $k \in \mathbb{N}$, $H_k(N_1, N_0)$ is a finitely generated abelian group and, for the types of calculations we are interested in performing, $H_k(N_1, N_0) \cong 0$ for sufficiently large k. Furthermore, if $f : (N_1, N_0) \to (N_1', N_0')$ is a continuous map of pairs, that is, $f : N_1 \to N_1'$ is continuous and $f(N_0) \subset N_0'$, then there is a collection of induced group homomorphisms, the *homology maps*,

$$f_* : H_*(N_1, N_0) \to H_*(N_1', N_0'),$$

where $f_* = \{f_k : H_k(N_1, N_0) \to H_k(N_1', N_0') \mid k = 0, 1, 2, \ldots\}$.

As will become clear shortly, the sets N_i for which we wish to compute homologies arise as large sets of grid elements. To further complicate matters, we do not know the map f, but rather a multivalued combinatorial map $\mathcal{F}$ which is an outer approximation of f. Because of the size of these objects, the homology computations must be done by computer. It is only

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

recently that reasonably efficient algorithms for these types of computations have been developed and, in the case of computing homology maps, there is considerable room for improvement. However, the important point is that there are algorithms and code available to perform these computations (see Kaczyński, Mischaikow and Mrozek (2001) and references therein).

The following set of definitions is required to define the Conley index for maps.

Definition 24. Let $f : X \to X$ and $g : Y \to Y$ be continuous maps between topological spaces (group homomorphism between abelian groups). They are *shift-equivalent* if there exist continuous maps (group homomorphism) $r : X \to Y$, $s : Y \to X$ and a natural number m such that

$$r \circ f = g \circ r, \quad s \circ g = f \circ s, \quad r \circ s = g^m, \quad s \circ r = f^m.$$

The homotopy classes of f and g are shift-equivalent if there exist continuous maps $s : X \to Y$, $r : Y \to X$ and a natural number m such that

$$r \circ f \sim g \circ r, \quad s \circ g \sim f \circ s, \quad r \circ s \sim g^m, \quad s \circ r \sim f^m.$$

Bibliographical notes

The terminology present in this subsection can be found in any algebraic topology book (Spanier 1982). An undergraduate-level introduction to computational homology is Kaczyński, Mischaikow and Mrozek (200x).

3.2. Conley index for flows

As a means of introducing the Conley index we begin with a discussion in the context of flows, since this is the simpler case. Furthermore, since it is for motivational purposes, we shall present the definitions in a very restrictive setting. For a more general presentation the reader should consult Conley (1978), Smoller (1980), Salamon (1985), Arnold *et al.* (1995) and Mischaikow and Mrozek (2001).

As before, $\varphi : \mathbb{R} \times \mathbb{R}^n \to \mathbb{R}^n$ denotes the flow generated by the differential equation $\dot{x} = g(x)$. Let $\mathcal{T}$ be a flow-transverse polygonal decomposition of $X \subset \mathbb{R}^n$.

Let $P \in \mathcal{T}$. If $L \in \partial P$, then there exists a unique n-dimensional simplex $K \in P$ such that L is face of K. L is an *exit face* of P if

$$\nu_K(L) \cdot f(x) < 0 \quad \text{for all } x \in L.$$

Definition 25. A pair $N = (N_1, N_0)$ of compact sets with $N_0 \subset N_1$ is an *index pair* if N_1 is a flow-transverse polygonal isolating block and

$$N_0 = \{L \in \partial N_1 \mid L \text{ is an exit face of } N_1\}.$$

Observe that by Proposition 2.18 polygonal isolating blocks, and hence polygonal index pairs, can be computed efficiently.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Let $N = (N_1, N_0)$ be an index pair. Let

$$S := \mathrm{Inv}\,(\mathrm{cl}\,(\mathrm{int}\,N_1 \setminus N_0), \varphi) = \mathrm{Inv}\,(N_1, \varphi) \subset \mathrm{int}\,(N_1 \setminus N_0). \tag{3.1}$$

Definition 26. The (*homotopy*) *Conley index* of S is

$$h(S) \sim (N_1/N_0, [N_0]).$$

Remark 2. Here we have presented the Conley index as an index of the isolated invariant set S. It can be proved that $h(S)$ is independent of the index pair used to compute it. In other words, if $N' = (N_1', N_0')$ is another index pair with the property that

$$S = \mathrm{Inv}\,(\mathrm{cl}\,(\mathrm{int}\,N_1' \setminus N_0'), \varphi),$$

then

$$(N_1/N_0, [N_0]) \sim (N_1'/N_0', [N_0']).$$

Observe that, as presented in Definition 26, the index is the homotopy equivalence class of a pointed topological space. This is an extremely difficult object to compute. For this reason we work with a coarser invariant, the *homological Conley index*

$$CH_*(S) :\cong H_*(N_1/N_0, [N_0]) \cong H_*(N_1, N_0).$$

The isomorphism on the right is due to the fact that N_1 and N_0 are simplicial complexes.

3.3. Index pairs and index filtrations

Having defined the Conley index for flows, we return to the setting of continuous maps. Let $N = (N_1, N_0)$ be a pair of compact sets with $N_0 \subset N_1$. Define $f_N : (N_1/N_0, [N_0]) \to (N_1/N_0, [N_0])$ by

$$f_N(x) = \begin{cases} f(x), & \text{if } x, f(x) \in N_1 \setminus N_0, \\ [N_0], & \text{otherwise.} \end{cases}$$

Definition 27. The pair of compact sets $N = (N_1, N_0)$ is an *index pair* if $\mathrm{cl}\,(N_1 \setminus N_0)$ is an isolating neighbourhood and the *index map* f_N is continuous.

The Conley index for maps will be defined in terms of the index map f_N. Therefore, it is essential that we be able to compute index pairs. As the theorems of this section indicate, the same algorithms that determined isolating blocks and Morse decompositions provide us with index pairs.

Theorem 3.1. Let $\mathcal{F}$ be an outer approximation of f. Let $\mathcal{U} = \mathrm{Inv}\,(o(\mathcal{U}) \cup \mathcal{V}, \mathcal{F}) \cap o(\mathcal{U})$ where $\mathcal{V} \cap o(\mathcal{U}) = \emptyset$. Then the pair

$$N = (N_1, N_0) := (|(d(\mathcal{U}) \cap \mathcal{F}(\mathcal{U})) \cup \mathcal{U}|, |d(\mathcal{U}) \cap \mathcal{F}(\mathcal{U})|)$$

is an index pair for f.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

If we let $\mathcal{V} = \emptyset$, then we obtain the following corollary. Observe that the assumption is precisely the halting condition for our algorithm to find isolating blocks within recurrent sets.

Corollary 3.2. Let $\mathcal{U} = \mathrm{Inv}\,(o(\mathcal{U}), \mathcal{F})$. Then the pair

$$N = (N_1, N_0) := (|(d(\mathcal{U}) \cap \mathcal{F}(\mathcal{U})) \cup \mathcal{U}|, |d(\mathcal{U}) \cap \mathcal{F}(\mathcal{U})|)$$

is an index pair for f.

As will be indicated in the next subsection, we can use these theorems to gain additional information about the structure of the dynamics within recurrent sets.

The construction which follows will allow us to understand the structure of connecting orbits between Morse sets.

Definition 28. Let $\mathcal{M}(S) := \{M(p) \mid p \in (\mathcal{P}, \succ)\}$ be a Morse decomposition of S with admissible order $\succ$. An *index filtration* of $\mathcal{M}(S)$ is a collection of compact sets $\mathcal{N} := \{N(I) \mid I \in \mathcal{A}(P, \succ)\}$ satisfying the following properties:

(1) for each $I \in \mathcal{A}(\mathcal{P}, \succ)$, $(N(I), N(\emptyset))$ is an index pair for $M(I)$,
(2) for any $I, J \in \mathcal{A}(\mathcal{P}, \succ)$,

$$N(I \cap J) = N(I) \cap N(J) \qquad \text{and} \qquad N(I) \cup N(J) = N(I \cup J).$$

Theorem 3.3. Let $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ be an outer approximation of f. Let $\mathcal{S}$ be an isolated invariant set for $\mathcal{F}$ and let $\{\mathcal{R}_p \mid p \in (\mathcal{P}, >)\}$ be the set of SCPC for the directed graph associated to $\mathcal{F}$ restricted to the isolating neighbourhood of $\mathcal{S}$. Let $\mathcal{M}(S) := \{M(p) \mid p \in (\mathcal{P}, \succ)\}$ be the associated Morse decomposition for f, that is, $S = \mathrm{Inv}\,(|\mathcal{S}|, f)$ and $M(p) := \mathrm{Inv}\,(|\mathcal{R}_p|, f)$. Then there exists a collection of subsets of $\mathcal{G}$ such that $\{|\mathcal{N}(I)| \mid I \in \mathcal{A}(P, \succ)\}$ is an index filtration of S under f.

A modification of the depth-first search algorithm that was used to compute the set of SCPC and the discrete Lyapunov function also provides for an index filtration. In particular, let $(\mathcal{N}(\mathcal{P}), \mathcal{N}(\emptyset))$ be an index pair for $\mathcal{S}$. For each strongly connected path component $\mathcal{R}_p$ of $\mathcal{F}$ restricted to $\mathcal{N}(\mathcal{P})$, define $\mathcal{E}_p$ to be the set of grid elements G for which there exists a path $\gamma_G : \{0, 1, \ldots, n\} \to \mathcal{N}(\mathcal{P})$ such that $\gamma_G(n) \in \mathcal{R}_p$. One obtains an index filtration by defining

$$\mathcal{N}(I) := \mathcal{N}(\mathcal{P}) \setminus \bigcup_{p \notin I} \mathcal{E}_p.$$

Bibliographical notes

The proof of Theorem 3.1 is due to Szymczak (1997). Index filtrations were first developed by R. Franzosa (1986) in the context of flows. The

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

particular construction described here has been used on several occasions (Franzosa and Mischaikow 1988, Eidenschink 1995, Richeson 1997).

3.4. Conley index for maps

Let

$$S := \mathrm{Inv}\,(\mathrm{cl}\,(\mathrm{int}\,N_1 \setminus N_0), f) \subset \mathrm{int}\,(N_1 \setminus N_0). \tag{3.2}$$

(compare with (3.1)). One is tempted at this point to define the Conley index of S in terms of the homotopy type of the pointed space $(N_1/N_0, [N_0])$. Unfortunately, it is easy to create examples where (N_1, N_0) and (N_1', N_0') are index pairs for S, that is,

$$\mathrm{Inv}\,(\mathrm{cl}\,(\mathrm{int}\,N_1 \setminus N_0), f) = S = \mathrm{Inv}\,(\mathrm{cl}\,(\mathrm{int}\,N_1' \setminus N_0'), f),$$

but

$$(N_1/N_0, [N_0]) \not\sim (N_1'/N_0', [N_0']).$$

The appropriate definition is as follows.

Definition 29. Let $N = (N_1, N_0)$ be an *index pair* for S. The (*homotopy*) *Conley index* for S is the homotopy shift equivalence class of the index map $f_N : (N_1/N_0, [N_0]) \to (N_1/N_0, [N_0])$.

For computational reasons we use the (*homology*) *Conley index* which is the shift equivalence class for group homomorphisms of

$$f_{N*} : H_*(N_1/N_0, [N_0]) \to H_*(N_1/N_0, [N_0]),$$

which in our context is equivalent to the shift equivalence class of

$$f_{N*} : H_*(N_1, N_0) \to H_*(N_1, N_0),$$

since N_1 and N_0 are constructed from grid elements.

Bibliographical notes

Defining the Conley index in terms of shift equivalence is due to Franks and Richeson (2000).

It should be noted that, because of the nature of our approximations, we cannot compute f_{N*} directly. Therefore, it is important to remark that, if $\mathcal{F}$ is a convex-valued outer approximation of f, then it is possible to compute f_{N*} from the data of $\mathcal{F}$. Details can be found in Kaczyński *et al.* (2001) and Kaczyński *et al.* (200x).

3.5. The structure of invariant sets

As was indicated in the Introduction, there is a variety of other references that provide information about the Conley index. What we provide here is a very curt description of how knowledge of the index provides information about the existence and structure of dynamics.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

The first theorem is the most fundamental: it implies that there exists a nontrivial invariant set.

Theorem 3.4. (Ważewski Property) Let $N = (N_1, N_0)$ be an index pair for f. If f_{N*} is not shift-equivalent to the trivial map, then

$$S := \mathrm{Inv}\,(\mathrm{cl}\,(\mathrm{int}\,N_1 \setminus N_0), f) \neq \emptyset.$$

There are additional theorems based on finer topological invariants that allow one to deduce the existence of fixed points and periodic orbits (McCord, Mischaikow and Mrozek 1995, McCord 1988, Mrozek 1989) or even that the invariant set has positive topological entropy (Baker 1998).

The next result provides a method by which one can move from local results to global results.

Theorem 3.5. (Summation property) Let $N = (N_1, N_0)$ and $N' = (N_1', N_0')$ be index pairs where $N_1 \cap N_1' = \emptyset$. Assume $N \cup N'$ is also an index pair for f. Then $f_{N \cup N'*}$ is shift-equivalent to $f_{N*} \oplus f_{N'*}$.

The simplest application is to assume that one has a Morse decomposition that consists of two Morse sets, that is, $\mathcal{M}(S) = \{M(0), M(1) \mid 1 \succ 0\}$. If the Conley index of S is not the direct sum of the Conley index of $M(0)$ and the Conley index of $M(1)$, then there must exist a connecting orbit from $M(1)$ to $M(0)$. This type of calculation can be extended to general Morse decompositions via the *connection matrix* (Franzosa 1989, Franzosa 1988, Arnold *et al.* 1995, Richeson 1997, Robbin and Salamon 1992, Mischaikow and Mrozek 2001) which is a generalization of the Morse inequalities. These types of arguments can also be used to describe the structure of the set of connecting orbits (Arnold *et al.* 1995, Mischaikow and Mrozek 2001, McCord and Mischaikow 1996, Mischaikow 1995, Mischaikow and Morita 1994) and global bifurcations (McCord and Mischaikow 1992, Franzosa and Mischaikow 1998, Mischaikow 1989).

In an alternate direction, by comparing the structure of $f_{N \cup N'*}$ with that of $f_{N*} \oplus f_{N'*}$ one can derive conditions under which one must have chaotic symbolic dynamics where the symbols correspond to the regions $\mathrm{cl}\,(\mathrm{int}\,N_1 \setminus N_0)$ and $\mathrm{cl}\,(\mathrm{int}\,N_1' \setminus N_0')$ (Mischaikow and Mrozek 1995, Szymczak 1996, Szymczak 1995, Carbinatto, Kwapisz and Mischaikow 2000, Carbinatto and Mischaikow 1999).

4. Examples

In this section we present four examples that show how the ideas of the previous sections can be applied. Space does not permit a full description of the work; therefore we only present the most pertinent aspects of each result.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

4.1. Hénon map

Recall that the Hénon map is given by the formula

$$f(x, y) = (1 - ax^2 + y, bx).$$

Szymczak (1997) obtained a computer-assisted proof of the following theorem.

Theorem 4.1. For $a = 1.4$ and $b = 0.3$ the Hénon map admits periodic orbits of all minimal periods except for 3 and 5.

We take this opportunity to present an outline of the proof that indicates how the constructions and results of the previous sections are applied.

The first step is to choose a region X of phase space for which a multivalued map is computed. Let

$$X = \left(\left[-\frac{85}{69}, \frac{85}{69}\right] \times \left[-\frac{86}{230}, \frac{86}{230}\right]\right) \setminus \left(\left[-\frac{85}{69}, \frac{10}{69}\right) \times \left[-\frac{86}{230}, 0\right)\right).$$

The second step is to construct a grid $\mathcal{G}$ that covers X. This was done using rectangles of size $\frac{1}{69} \times \frac{1}{230}$. Simple estimates allow one to construct a multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ that is an outer approximation of f.

The third step is to find an interesting isolating block that is a subset of $\mathcal{G}$. This was done using the reduction technique described in Theorem 2.14. An isolating block with five components was chosen, from which an index pair was constructed using Theorem 3.1. From this a Conley index map f_N was computed and the Conley index was used to guarantee the existence of periodic orbits with all minimal periods except 3 and 5.

The final step is to exclude the periodic orbits of period 3 and 5. This was done by examining all cubes G with the property that $G \in \mathcal{F}^3(G)$ or $G \in \mathcal{F}^5(G)$, and then showing that, within the region of phase space containing these cubes, the only recurrent dynamics consisted of a fixed point.

4.2. Lorenz equations

As was indicated earlier, one way to study complicated dynamics for ordinary differential equations is to study the dynamics of a Poincaré map $f : \Xi \to \Xi$ by computing a multivalued outer approximation $\mathcal{F}$ and following similar steps as described for the Hénon map computation. Of course the computational expense is considerably more than in the previous example, since one needs to rigorously bound the trajectories of a differential equation. This approach was used by M. Mrozek, A. Szymczak and the author (Mischaikow, Mrozek and Szymczak 2001) to study the Lorenz equations with the following results.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

Let $\Sigma(k)$ be the set of bi-infinite sequences on k symbols. For a $k \times k$ matrix $A = (A_{ij})$ over $\mathbb{Z}_2$ let $\Sigma(A)$ be the set of bi-infinite sequences $\alpha : \mathbb{Z} \to \{1, 2, \ldots, k\}$ on k symbols that satisfy the following restriction:

$$A_{\alpha(i)\alpha(i+1)} = 1 \quad \text{for all } i \in \mathbb{Z}.$$

Clearly, $\Sigma(A) \subset \Sigma(k)$. Let $\sigma : \Sigma(A) \to \Sigma(A)$ be the shift map given by

$$\sigma(\alpha)(i) = \alpha(i+1).$$

Let $n_{(10,28,8/3)} = 6$, $n_{(10,60,8/3)} = n_{(10,54,45)} = 4$,

$$A_{(10,28,8/3)} = \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 0 \end{bmatrix},$$

$$A_{(10,60,8/3)} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \qquad A_{(10,54,45)} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}.$$

Theorem 4.2. Consider the Lorenz equations

$$\begin{aligned} \dot{x} &= s(y - x), \\ \dot{y} &= Rx - y - xz, \\ \dot{z} &= xy - qz, \end{aligned} \tag{4.1}$$

and the plane $P := \{(x, y, z) \mid z = R - 1\}$. For all parameter values in a sufficiently small neighbourhood of (σ, R, b) there exists a Poincaré section $N \subset P$ such that the associated Poincaré map f is Lipschitz and well defined. Furthermore, there is a continuous map $\rho : \mathrm{Inv}(N, f) \to \Sigma(n_{(\sigma,R,b)})$ such that

$$\rho \circ f = \sigma \circ \rho$$

and $\Sigma(A_{(\sigma,R,b)}) \subset \rho(\mathrm{Inv}(N, f))$. Moreover, for each periodic sequence $\alpha \in \Sigma(A)$ there exists an $x \in \mathrm{Inv}(N, f)$ on a periodic trajectory of the same minimal period such that $\rho(x) = \alpha$.

4.3. A partial differential equation

As was indicated in the Introduction, these ideas can be carried out for infinite-dimensional systems. A long-range goal is to be able to perform efficient rigorous computations to prove the existence of chaotic dynamics

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

in partial differential equations. There are three obvious difficulties that need to be overcome to achieve our goal.

D1 Because of the finite nature of a computer it is impossible to compute directly on an infinite-dimensional system. Therefore, it is necessary to use an appropriate finite-dimensional reduction.

D2 Given a finite-dimensional system, *i.e.*, an ordinary differential equation, we need to be able to find isolating blocks and compute Conley indices.

D3 The index results of the finite-dimensional system need to be lifted to the full infinite-dimensional system.

However, as was discussed in the subsection of flows, we still do not have efficient methods for producing flow-transverse polygonal grids that are good approximations of the dynamics. Therefore, in general, we are not yet able to overcome difficulty D2. On the other hand, recent work by P. Zgliczyński and the author (Zgliczyński and Mischaikow 2001) shows that efficient, accurate and rigorous computation of fixed points is possible.

They considered the Kuramoto–Sivashinsky equation

$$u_t = -\nu u_{xxxx} - u_{xx} + 2uu_x \qquad (t,x) \in [0,\infty) \times (-\pi,\pi)$$

subject to periodic and odd boundary conditions

$$u(t,-\pi) = u(t,\pi) \quad \text{and} \quad u(t,-x) = -u(t,x). \tag{4.2}$$

The finite-dimensional reduction is done using a standard Galerkin projection based on a Fourier series expansion. In particular, one obtains the following infinite set of ordinary differential equations:

$$\dot{a}_k = k^2(1-\nu k^2)a_k - k\sum_{n=1}^{k-1} a_n a_{k-n} + 2k\sum_{n=1}^{\infty} a_n a_{n+k} \quad k = 1,2,3,\ldots, \tag{4.3}$$

which is then truncated to the finite-dimensional system

$$\dot{a}_k = k^2(1-\nu k^2)a_k - k\sum_{n=1}^{k-1} a_n a_{k-n} + 2k\sum_{n=1}^{m-k} a_n a_{n+k} \quad k = 1,\ldots,m. \tag{4.4}$$

Let $W \subset \mathbb{R}^m$ be a compact region. Within W, fixed points for (4.4) are identified. One can linearize around these points and use the eigenvectors to identify m-dimensional cubes with the property that the vector field generated by (4.4) is normal to the boundary faces.

Of course, (4.4) is just an approximation of the true vector field given by (4.3). To control the errors, it is assumed that

$$|a_k| \leq \pm\frac{C_s}{k^s} \quad \text{for all } k > m,$$

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

where $s \geq 4$. Observe that this is a regularity condition. However, solutions to the Kuramoto–Sivashinsky equation that exist for all time are analytic, and therefore the higher Fourier coefficients of any solution that lies in a bounded invariant set will satisfy these constraints for sufficiently large C_s. It should also be remarked that this high degree of regularity is not particular to Kuramoto–Sivashinsky (see Hale and Raugel (2001) and references therein).

These constraints imply that, in the function space, we are studying the dynamics on the compact region

$$W \times \prod_{k=m+1}^{\infty} \left[-\frac{C_s}{k^s}, \frac{C_s}{k^s} \right].$$

Restricted to this region, one can bound the difference between the vector fields generated by (4.4) and (4.3), and therefore one can check if the transversality conditions on the faces of the isolating blocks in $\mathbb{R}^n$ are valid, not only for the vector field (4.4) but also for (4.3).

Finally, since Kuramoto–Sivashinsky is a strongly dissipative system it is easy to believe that, for sufficiently large k, the vector field is essentially of the form $\dot{a}_k = \lambda_k a_k$, where $\lambda < 0$. In other words, for sufficiently large k, the vector field is transverse to the faces defined by the product $\prod_{k=m+1}^{\infty} \left[-\frac{C_s}{k^s}, \frac{C_s}{k^s} \right]$ and furthermore, point inward. Since the Conley index is determined by the boundary pieces where the vector field points outwards, the index is determined by the low modes. In this way one can compute the index by a finite-dimensional approximation.

The final step is to show that this information can be lifted to the full system. The reader is referred to Zgliczyński and Mischaikow (2001) for the details.

4.4. An infinite-dimensional map

This final example is included to indicate that in the case of maps one can rigorously determine a wide range of dynamical objects even in the infinite-dimensional setting. Recently, M. Allili, S. Day, O. Junge and the author (Allili *et al.* 2002) have studied the Kot–Schaffer growth-dispersal model for plants. This is a map $\Phi : L^2([-\pi, \pi]) \to L^2([-\pi, \pi])$ of the form

$$\Phi(a)(y) := \frac{1}{2\pi} \int_{-\pi}^{\pi} b(x, y)\ \mu\ a(x) \left(1 - \frac{a(x)}{c(x)} \right) \mathrm{d}x, \tag{4.5}$$

where $\mu > 0$ and $b(x, y) = b(|x - y|)$. Observe that the regularity of this map is determined by the regularity of the dispersal kernel b and the spatial heterogeneity of c in the nonlinear term.

The same difficulties D1–D3 need to be confronted with this problem. Again, using Fourier series one can rewrite the problem as an infinite system

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

of maps:

$$a'_k = f_k(a) := \mu b_k \left[a_k - \sum_{\substack{j+l+n=k \\ j,l,k \geq 0}} c_j a_l a_n \right], \quad k = 0, 1, 2, \ldots. \tag{4.6}$$

Simple projection results in the finite-dimensional system $f^{(m)} : \mathbb{R}^m \to \mathbb{R}^m$

$$a'_k = f_k^{(m)}(a_0, \ldots, a_{m-1}) := \mu b_k \left[a_k - \sum_{\substack{j+l+n=k \\ 0 \leq j,l,n \leq m-1}} c_j a_l a_n \right], \tag{4.7}$$

where $k = 0, 1, \ldots, m-1$, upon which the numerical computations are based.

Let

$$W := \prod_{k=0}^{m-1} [a_k^-, a_k^+] \subset \mathbb{R}^m,$$

and consider $f^{(m)} : W \to \mathbb{R}^m$. Assuming that the Fourier coefficients of the dispersal kernel b decay at an exponential rate, and taking advantage of the resulting regularity constraint on solutions which lie in bounded invariant sets, one can restrict attention to a compact set

$$Z := W \times \prod_{k=m}^{\infty} \left[-\frac{C_s}{s^k}, \frac{C_s}{s^k} \right]. \tag{4.8}$$

A grid $\mathcal{G}$ covering W is constructed and a multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ that is an outer approximation of f is constructed. It should be emphasized that the error bounds used to determine $\mathcal{F}$ include the errors obtained by truncation to the finite-dimensional system for elements in Z. This implies that, with the proper choice of m and C_s, the index computations for the finite-dimensional approximation are valid in the infinite-dimensional setting.

Using the techniques described in the latter part of Section 2.4, one can find isolating blocks for fixed points, periodic orbits, connecting orbits, and horseshoes. At this point the cost of computing the Conley index becomes an issue. Computation of homology grows in expense with the dimension of the complex. On the other hand, to obtain precise estimates on the solutions requires the use of more modes.

The strategy adopted in Allili *et al.* (2002) is as follows. A coarse (but mathematically rigorous) computation is done to determine an isolating block using a minimal number of modes. At this point the index is computed. One then increases the number of modes and refines the isolating block. As was described in Section 2, these computations are essentially linear in the number of boxes, and therefore the computational complexity is determined by the dimension of the invariant set. In the case of heteroclinic

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

orbits or horseshoes, these are at most 1-dimensional objects. Since the blocks isolate the same invariant set, the Conley index remains unchanged. In this way the number of modes used in the approximation can be increased until the significant errors are those introduced by the floating point errors. The reader is referred to Allili *et al.* (2002) for details.

Acknowledgements
The author would like to thank Thomas Wanner, Anthony Baker, and the editors for carefully reading the orginal manuscript and providing a multitude of corrections and useful suggestions.

REFERENCES

M. Allili, S. Day, O. Junge and K. Mischaikow (2002), A rigorous numerical method for the global analysis of an infinite dimensional discrete dynamical system. In preparation.

L. Arnold, C. Jones, K. Mischaikow and G. Raugel (1995), *Dynamical Systems* (R. Johnson, ed.), Vol. 1609 of *Lecture Notes in Mathematics*, Springer.

A. Baker (1998), Topological entropy and the first homological Conley index map. Preprint.

E. Boczko, W. Kalies and K. Mischaikow (2002), Polygonal approximations of flows. In preparation.

M. Carbinatto and K. Mischaikow (1999), 'Horseshoes and the Conley index spectrum II: The theorem is sharp', *Dis. Cont. Dyn. Sys.* **5**, 599–616.

M. Carbinatto, J. Kwapisz and K. Mischaikow (2000), 'Horseshoes and the Conley index spectrum', *Erg. Theory Dyn. Sys.* **20**, 365–377.

K. A. Cliffe, A. Spence and S. J. Tavener (2000), 'The numerical analysis of bifurcation problems with application to fluid mechanics', in *Acta Numerica*, Vol. 9, Cambridge University Press, pp. 39–131.

C. Conley (1978), *Isolated Invariant Sets and the Morse Index.* Vol. 38 of *CBMS Lecture Notes*, AMS, Providence, RI.

C. Conley and R. Easton (1971), 'Isolated invariant sets and isolating blocks', *Trans. Amer. Math. Soc.* **158**, 35–61.

T. Corman, C. Leiserson, R. Rivest, C. Stein (2001), *Introduction to Algorithms*, MIT Press.

M. Dellnitz and A. Hohmann (1997), 'A subdivision algorithm for the computation of unstable manifolds and global attractors', *Numer. Math.* **75**, 293–317.

M. Dellnitz and O. Junge (2001), Set oriented numerical methods for dynamical systems, in *Handbook of Dynamical Systems* III: *Towards Applications* (B. Fiedler, G. Iooss and N. Kopell, eds), World Scientific. To appear.

M. Dellnitz, G. Froyland and O. Junge (2000), The algorithms behind GAIO: Set oriented numerical methods for dynamical systems. Preprint.

L. Dieci and T. Eirola (2001), *Numerical Dynamical Systems.* Notes.

E. J. Doedel, A. R. Champneys, T. F. Fairgrieve, Yu. A. Kuznetsov, B. Sandstede and X. J. Wang (1999), AUTO: *Continuation and Bifurcation Software with Ordinary Differential Equations, User's Guide*, Concordia University.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

R. Easton (1975), 'Isolating blocks and symbolic dynamics', *J. Diff. Equations* **17**, 96–118.

R. Easton (1989), 'Isolating blocks and epsilon chains for maps', *Physica D* **39**, 95–110.

R. Easton (1998), *Geometric Methods for Discrete Dynamical Systems*, Oxford University Press.

M. Eidenschink (1995), Exploring global dynamics: A numerical algorithm based on the Conley index theory, PhD Thesis, Georgia Institute of Technology.

A. Floer (1987), 'A refinement of the Conley index and an application to the stability of hyperbolic invariant sets', *Erg. Theory Dyn. Sys.* **7**, 93–103.

J. Franks and D. Richeson (2000), Shift equivalence and the Conley index. *Trans. Amer. Math. Soc.* **352**, 3305–3322.

R. Franzosa (1986), 'Index filtrations and the homology index braid for partially ordered Morse decompositions', *Trans. Amer. Math. Soc.* **298**, 193–213.

R. Franzosa (1988), 'The continuation theory for Morse decompositions and connection matrices', *Trans. Amer. Math. Soc.* **310**, 781–803.

R. Franzosa (1989), 'The connection matrix theory for Morse decompositions', *Trans. Amer. Math. Soc.* **311**, 561–592.

R. Franzosa and K. Mischaikow (1988), 'The connection matrix theory for semiflows on (not necessarily locally compact) metric spaces', *J. Diff. Equations* **71**, 270–287.

R. Franzosa and K. Mischaikow (1998), 'Algebraic transition matrices in the Conley index theory', *Trans. Amer. Math. Soc.* **350**, 889–912.

J. Hale and G. Raugel (2001), Regularity, Determining Modes and Galerkin Methods, Preprint.

T. Kaczyński and M. Mrozek (1995), 'Conley index for discrete multivalued dynamical systems', *Topology Appl.* **65**, 83–96.

T. Kaczyński and M. Mrozek (1997), 'Stable index pairs for discrete dynamical systems', *Canadian Math. Bull.* **40**, 448–455.

T. Kaczyński, K. Mischaikow and M. Mrozek (2001), Computing homology. Preprint.

T. Kaczyński, K. Mischaikow and M. Mrozek (200x), *Computational Homology*, Springer. In preparation.

C. McCord (1988), 'Mappings and homological properties in the homology Conley index', *Erg. Theory Dyn. Sys.* **8***, 175–198.

C. McCord and K. Mischaikow (1992), 'Connected simple systems, transition matrices, and heteroclinic bifurcations', *Trans. Amer. Math. Soc.* **333**, 379–422.

C. McCord and K. Mischaikow (1996), 'On the global dynamics of attractors for scalar delay equations', *J. Amer. Math. Soc.* **9**, 1095–1133.

C. McCord, K. Mischaikow and M. Mrozek (1995), Zeta functions, periodic trajectories, and the Conley index, *J. Diff. Equations* **121**, 258–292.

K. Mischaikow (1989), 'Transition systems', *Proc. Roy. Soc. Edin.* **112A**, 155–175.

K. Mischaikow (1995), 'Global asymptotic dynamics of gradient-like bistable equations', *SIAM Math. Anal.* **26**, 1199–1224.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

K. Mischaikow and Y. Morita (1994), 'Dynamics on the global attractor of a gradient flow arising in the Ginzburg–Landau equation', *Jap. J. Ind. Appl. Math.* **11**, 185–202.
K. Mischaikow and M. Mrozek (1995), 'Isolating neighborhoods and chaos', *Jap. J. Ind. Appl. Math.* **12**, 205–236.
K. Mischaikow and M. Mrozek (1995), 'Chaos in Lorenz equations: A computer assisted proof', *Bull. Amer. Math. Soc.* (*N.S.*) **33**, 66–72.
K. Mischaikow and M. Mrozek (1998), 'Chaos in Lorenz equations: A computer assisted proof, Part II: Details', *Math. Comput.* **67**, 1023–1046.
K. Mischaikow and M. Mrozek (2001), Conley index, in *Handbook of Dynamical Systems*, Vol. 2 (B. Fiedler, ed.) Elsevier.
K. Mischaikow, M. Mrozek and A. Szymczak (2001), 'Chaos in Lorenz equations: A computer assisted proof, Part III: The classic case'. *J. Diff. Equations* **169** 17–56.
M. Mrozek (1989), 'Index pairs and the fixed point index for semidynamical systems with discrete time', *Fund. Mathematicae* **133**, 177–192.
M. Mrozek (1990), 'Leray functor and the cohomological Conley index for discrete dynamical systems', *Trans. Amer. Math. Soc.* **318**, 149–178.
M. Mrozek (1996), 'Topological invariants, multivalued maps and computer assisted proofs in dynamics', *Comput. Math.* **32**, 83–104.
M. Mrozek (1999), 'An algorithmic approach to the Conley index theory', *J. Dyn. Diff. Equations* **11** 711-734.
M. Mrozek, M. Żelawski (1997), 'Heteroclinic connections in the Kuramoto–Sivashinsky equation', *Reliable Computing* **3**, 277–285.
D. Norton (1989), A metric approach to the Conley decomposition theorem, PhD thesis, University of Minnesota.
J. Palis and F. Takens (1993), *Hyperbolicity and Sensitive Chaotic Dynamics at Homoclinic Bifurcations*, Cambridge University Press.
P. Pilarczyk (1999), Computer assisted method for proving existence of periodic orbits. Preprint.
M. Poźniak (1994), 'Lusternik–Schnirelmann category of an isolated invariant set', *Univ. Jag. Acta Math.* **XXXI**, 129–139.
J. Reineck (1988), 'Connecting orbits in one-parameter families of flows', *Erg. Theory Dyn. Sys.* **8***, 359–374.
D. Richeson (1997), Connection matrix pairs for the discrete Conley index. Preprint.
J. W. Robbin and D. Salamon (1988), 'Dynamical systems, shape theory and the Conley index', *Erg. Theory Dyn. Sys.* **8***, 375–393.
J. W. Robbin and D. Salamon (1992), 'Lyapunov maps, simplicial complexes and the Stone functor', *Erg. Theory Dyn. Sys.* **12**, 153–183.
C. Robinson (1995), *Dynamical Systems: Stability, Symbolic Dynamics, and Chaos*, CRC Press.
C. Robinson (1977), 'Stability theorems and hyperbolicity in dynamical systems', *Rocky Mount. J. Math.* **7**, 425–437.
D. Salamon (1985), 'Connected simple systems and the Conley index of isolated invariant sets'. *Trans. Amer. Math. Soc.* **291**, 1–41.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press

J. Smoller (1980), *Shock Waves and Reaction Diffusion Equations*, Springer, New York.

E. H. Spanier (1982), *Algebraic Topology*, Springer. (McGraw-Hill 1966.)

A. M. Stuart and A. R. Humphries (1996), *Dynamical Systems and Numerical Analysis*, Vol. 2 of *Cambridge Monographs on Applied and Computational Mathematics*, Cambridge University Press.

A. Szymczak (1995), 'The Conley index for discrete dynamical systems', *Topology and its Applications* **66**, 215–240.

A. Szymczak (1995), 'The Conley index for decompositions of isolated invariant sets', *Fund. Math.* **148**, 71–90.

A. Szymczak (1996), 'The Conley index and symbolic dynamics', *Topology* **35**, 287–299.

A. Szymczak (1997), 'A combinatorial procedure for finding isolating neighborhoods and index pairs', *Proc. Roy. Soc. Edin.* **127A**.

P. Zgliczyński and K. Mischaikow (2001), 'Rigorous numerics for partial differential equations: The Kuramoto–Sivashinsky equation'. *J. FoCM* **1**, 255–288.

https://doi.org/10.1017/S0962492902000065 Published online by Cambridge University Press