

An Algorithmic Approach to Chain Recurrence

W. D. Kalies,¹ K. Mischaikow,² and R. C. A. M. VanderVorst³

¹Department of Mathematical Sciences
Florida Atlantic University
Boca Raton, FL 33434, USA
wkalies@fau.edu

²Center for Dynamical Systems and Nonlinear Studies
Georgia Institute of Technology
Atlanta, GA 30332, USA
mischai@math.gatech.edu

³Department of Mathematics
Vrije University Amsterdam
De Boelelaan 1081
1081 HV, Amsterdam
The Netherlands
vdvorst@cs.vu.nl

Dedicated to Steve Smale on His 75th Birthday

Abstract. In this paper we give a new definition of the chain recurrent set of a continuous map using finite spatial discretizations. This approach allows for an algorithmic construction of isolating blocks for the components of Morse decompositions which approximate the chain recurrent set arbitrarily closely as well as discrete approximations of Conley's Lyapunov function. This is a natural framework in which to develop computational techniques for the analysis of qualitative dynamics including rigorous computer-assisted proofs.

Date received: November 18, 2004. Final version received: June 16, 2005. Date accepted: July 13, 2005. Communicated by Mike Shub. Online publication December 1, 2005.

AMS classification: 37B20, 37B25, 37B30, 37B35, 37C70.

Key words and phrases: Chain recurrence, Lyapunov function, Conley's decomposition theorem, Algorithms, Computation, Combinatorial dynamics.

Conley's Fundamental Decomposition Theorem and its extension to Morse decompositions is a powerful tool in dynamical systems theory. However, the framework on which the standard theory is built does not lead naturally to an algorithmic or computational approach for the approximation of the chain recurrent set, i.e., generation of Morse decompositions or the approximation of a Lyapunov function for the gradient-like part of the system. One can approximate the chain recurrent set by the ε -chain recurrent set for finite $\varepsilon > 0$, but there are no algorithmic or computational techniques for computing this set directly.

In this paper, we present an alternative approach based on finite discretizations and combinatorial multivalued maps. This approach has several advantages. The basic elements of the theory can be proved in a straightforward manner. Moreover, the methods are inherently combinatorial and hence algorithmic. The framework leads naturally to computational techniques for analyzing qualitative dynamics including rigorous computer-assisted proofs, see, e.g., [12], [16], [4], [5].

1. Introduction

1.1. Preliminaries

We begin with a suitable definition of a dynamical system. Let (X, d) be a compact metric space X with metric d .

Definition 1.1. Let \mathbb{T} denote either \mathbb{Z} or \mathbb{R} and let \mathbb{T}^+ denote either $\mathbb{Z}^+ = \{0, 1, 2, \dots\}$ or $\mathbb{R}^+ = [0, \infty)$. A **dynamical system** on X is a continuous map $\varphi: \mathbb{T}^+ \times X \rightarrow X$ that satisfies the following two properties:

- (i) $\varphi(0, x) = x$;
- (ii) $\varphi(t, \varphi(s, x)) = \varphi(t + s, x)$, for all $s, t \in \mathbb{T}^+$.

We adopt the additional property that, for $t > 0$,

- (iii) $\varphi(t, X) = X$ and, consequently,

$$\varphi(-t, x) := \{y \in X \mid \varphi(t, y) = x\} \neq \emptyset.$$

The last condition implies that φ extends to a (possibly) multivalued map (for $t < 0$) on all of $\mathbb{T} \times X$. In this case, a complete orbit of φ through x is a function $\gamma_x: \mathbb{T} \rightarrow X$ satisfying $\gamma_x(0) = x$ and $\gamma_x(t + s) = \varphi(s, \gamma_x(t))$ for all $t, s \in \mathbb{T}$. If in the above definition φ extends to a single-valued, continuous map on $\mathbb{T} \times X$, then φ is either an iteration of a homeomorphism or is a flow. In the case of discrete-time dynamical systems it is often useful to write φ in terms of its generator and define $f = \varphi(1, \cdot)$. Then $\varphi(n, \cdot) = f^n(\cdot)$ for $n \in \mathbb{Z}^+$, where f^n is the n th composition of f .

An important question in the understanding of dynamical systems is how $\varphi(t, x)$ evolves as $t \rightarrow \infty$ for all $x \in X$. Answering this question begins with the fundamental notions of invariant and isolated invariant sets.

Definition 1.2. A subset $S \subset X$ is called an **invariant set** if $\varphi(t, S) = S$ for all $t \in \mathbb{T}^+$. A compact subset $N \subset X$ is called an **isolating neighborhood** if

$$\text{Inv}(N, \varphi) \subset \text{int}(N),$$

where $\text{Inv}(N, \varphi)$ denotes the maximal invariant set under φ contained in N . An invariant set S is called **isolated** if $S = \text{Inv}(N, \varphi)$ for some isolating neighborhood N .

Natural invariant sets to study for long-term behavior are the α -**limit** and ω -**limit** sets $\omega(U) = \bigcap_{t>0} \text{cl}(\varphi([t, \infty] \cap \mathbb{T}, U))$ and $\alpha(U) = \bigcap_{t>0} \text{cl}(\varphi((-\infty, -t] \cap \mathbb{T}, U))$, and a first step toward characterizing the dynamics of a system in terms of transient and nontransient behavior is the development of Morse decompositions by Conley [2], which in their simplest form are described as follows.

A set A is called an *attractor* if there exists an open neighborhood $U \supset A$ such that $\omega(U) = A$. An attractor is an isolated invariant set. The associated *dual repeller* A^* is defined as $A^* = \{x: \omega(x) \cap A = \emptyset\}$, which is also an isolated invariant set. One can show that if $U^* = X \setminus A$, then $A^* = \alpha(U^*)$. The pair (A, A^*) is called an *attractor–repeller pair* and yields the following decomposition for X :

$$X = A \cup A^* \cup C(A^*, A),$$

where $C(A^*, A) \subset \{x \in X: \omega(x) \subset A, \alpha(x) \subset A^*\}$. The notions of attractor and repeller can also be used for invariant subsets $S \subset X$. The restriction $\varphi|_S$ is a dynamical system on S . By letting S play the role of X in the above definitions, we obtain the notion of an attractor–repeller pair decomposition for any isolated invariant set S . Attractor–repeller pair decompositions are the coarsest decompositions of X ; note that (X, \emptyset) and (\emptyset, X) are trivial attractor–repeller pairs. Morse decompositions are refinements/generalizations of attractor–repeller pair decompositions and are defined as follows.

Definition 1.3. A **Morse decomposition** of X is a finite collection of disjoint isolated invariant sets (Morse sets) S_1, \dots, S_n with a partial ordering \succ on the index set $\{1, \dots, n\}$ such that for every $x \in X \setminus \bigcup_i S_i$ and every complete orbit γ_x through x there exist indices $i \succ j$ so that

$$\omega(x) \subset S_j \quad \text{and} \quad \alpha(\gamma_x) = \bigcap_{t \leq 0} \text{cl}(\gamma_x((-\infty, t]))$$

By decomposing Morse sets into attractor–repeller pairs, Morse decompositions can be refined. While many systems have a finest Morse decomposition, there are systems which have (countably) infinitely many attractor–repeller pairs and hence no finest Morse decomposition. The process of extracting all attractor–repeller pairs leads to the chain recurrent set introduced by Conley [2].

Definition 1.4. The **chain recurrent** set $\mathcal{R}(X)$ for X is characterized by

$$\mathcal{R}(X) = \bigcap_j (A_j \cup A_j^*),$$

where the intersection is taken over all possible attractor–repeller pairs for X .

The chain recurrent set $\mathcal{R}(X)$ is an invariant set for φ which represents the smallest set outside of which the dynamics is gradient-like as stated in Conley’s Fundamental Decomposition Theorem below. The *chain components* of $\mathcal{R}(X)$ are the equivalence classes of the relation $x \sim y$ which is defined as follows: $x, y \in \mathcal{R}(X)$ are equivalent if for all attractor–repeller pairs (A_j, A_j^*) either x and y are both in A_j or both in A_j^* . The chain recurrent set has only countably many recurrent components $\mathcal{R}_i(X)$, which are not necessarily isolated invariant sets.

Theorem 1.5 (Conley’s Fundamental Decomposition Theorem [2], [15]). *Let $\varphi: \mathbb{T}^+ \times X \rightarrow X$ be a dynamical system on a compact metric space X . Then there are at most countably many chain components $\{\mathcal{R}_i\}_{i \in J}$ of $\mathcal{R}(X)$ and there exists a continuous function $V: X \rightarrow [0, 1]$ such that:*

- (a) if $x \notin \mathcal{R}(X)$, then $V(x) > V(\varphi(t, x))$, $t > 0$, and
- (b) for each $i \in J$ there exists $\sigma_i \in [0, 1]$ such that $\mathcal{R}_i \subset V^{-1}(\sigma_i)$.

Moreover, V can be chosen such that $\sigma_i \neq \sigma_j$ for all $i, j \in J$.

The function V is called a *Lyapunov function* for φ . Note that all Morse decompositions can be built by grouping chain recurrent components and their connecting orbits and, hence, corresponding Lyapunov functions can be constructed which are constant on the Morse sets.

1.2. Outline of the Main Ideas and Results

An important goal of this paper is to present a method for constructing approximations of the chain recurrent set of a dynamical system and its associated Lyapunov function. The standard way of defining the chain recurrent set is to consider the ε -chain recurrent set which consists of all points x for which there exists an ε -chain from x back to itself. While allowing for an ε -error in the dynamical system is a natural approximation to the chain recurrent set, this approach is not computable in many circumstances, and we will proceed by a different path. For a general reference on chain recurrence see [2] or [15], and for an example where ε -chains have been used successfully in certain situations, see [11] and the references therein.

The first step in our approach is to define a spatial discretization of the metric space X . In Sections 2 and 3, we introduce the notion of a grid from [13] which divides the space into finitely many compact sets. We denote the set of these grid elements by \mathcal{G} . The diameter of the largest element of the grid controls the fineness of the approximation. After discretization, we define a multivalued map on the grid

which mimics the dynamics of φ . To be more precise, we define maps $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ which “converge” in the proper sense to φ as the grid size goes to zero, as explained in Sections 3 and 5.

A multivalued map \mathcal{F} can be thought of as a discrete dynamical system and as a directed graph. In the latter case the grid elements act as vertices of the graph and two vertices are connected via an edge if \mathcal{F} maps from one vertex to the other. In Section 3 we define for \mathcal{F} the dynamical notions of attractors, repellers, and recurrent sets and put these in the context of graph theory. The purpose of these formulations is to demonstrate that the objects we wish to study are computable by standard graph-theoretic algorithms, which are often very efficient, see [3].

As we will see in Section 5, an appropriately defined sequence of multivalued maps and their associated recurrent sets will converge to the chain recurrent set of φ , providing a new definition of the chain recurrent set—(a similar result can be found in the work of Osipenko and Campbell [14]). Furthermore, the spatial realization of a recurrent set for \mathcal{F} approximates the chain recurrent set for φ , and graph theory provides a linear-time algorithm for determining the recurrent set of a graph, see [3]. Hence these approximations are computable. In Section 7 we show that the spatial realizations of components of the recurrent set of \mathcal{F} are isolating blocks for φ . Indeed, this occurs for the spatial realizations of many invariant sets of \mathcal{F} . Moreover, in Section 7 we prove that under certain conditions invariant sets in the graph converge to isolated invariant sets for φ . In particular, Morse sets and the recurrent set have this property.

Section 6 investigates another important tool in dynamical systems theory, Lyapunov functions. For multivalued maps \mathcal{F} one can easily construct (e.g., via graph theory) a function on the vertices of the graph which is constant on the components of the recurrent set and decreases strictly off the recurrent set, i.e., for each vertex that does not lie in the recurrent set, the function attains strictly smaller values on all vertices in the image under \mathcal{F} . We would like to use such Lyapunov functions for \mathcal{F} to approximate a Lyapunov function for φ . As the grid size goes to zero, Lyapunov functions constructed in an arbitrary manner will not generally converge to a Lyapunov function for φ . However, we will describe a method for constructing Lyapunov functions for \mathcal{F} that have the property that, for a family of grids with the grid size going to zero, these discrete Lyapunov functions converge along a subsequence to a continuous Lyapunov function for φ .

In Sections 2 through 6, for the sake of simplicity, the theory is explained for discrete-time dynamical systems, i.e., $\varphi(n, \cdot) = f^n$, where $f: X \rightarrow X$ is a surjective map. We utilize a particular multivalued map \mathcal{F} in this case. In Section 7 we formulate general conditions that a multivalued map must satisfy in order for the theory to hold.

2. Grids and Multivalued Maps

The first step in developing a computational foundation for Conley’s theory is the construction of a combinatorial representation of the dynamical system of interest.

This requires a finite description of both the metric space X and the dynamical system φ . In this section φ will be a discrete dynamical system generated by a surjective map $f: X \rightarrow X$. For the definition of grid and the combinatorialization of f only local compactness of X is needed. Therefore, in this section we assume that X is locally compact and separable. We begin with the following discretization due to Mrozek [13].

Definition 2.1. A **grid** on X is a collection \mathcal{G} of nonempty compact subsets of X with the following properties:

- (i) $X = \bigcup_{G \in \mathcal{G}} G$.
- (ii) $G = \text{cl}(\text{int}(G))$ for all $G \in \mathcal{G}$.
- (iii) $G \cap \text{int}(H) = \emptyset$ for all $G \neq H \in \mathcal{G}$.
- (iv) If $K \subset X$ is compact, then $\text{cov}(K, \mathcal{G}) = \{G \in \mathcal{G} \mid G \cap K \neq \emptyset\}$ is a finite set.

Observe that if X is compact, then \mathcal{G} is finite. The *diameter* of a grid is defined by

$$\text{diam}(\mathcal{G}) := \sup\{\text{diam}(G) \mid G \in \mathcal{G}\}.$$

The *realization map* $|\cdot|$ is a mapping from subsets of \mathcal{G} to subsets of X , and is defined as $|\mathcal{A}| := \bigcup_{A \in \mathcal{A}} A \subset X$. On compact spaces the existence of grids of arbitrarily small size easily follows from compactness. From this the same can be achieved in the locally compact case.

Theorem 2.2. *For any locally compact, separable metric space X and $\varepsilon > 0$, there exists a grid \mathcal{G} with $\text{diam}(\mathcal{G}) \leq \varepsilon$.*

Proof. The construction involves a simple induction argument. For simplicity first assume that X is compact. By compactness we can choose a finite subcover $\mathbf{U} = \{U_i \mid i = 1, \dots, n\}$ from the set $\{B_\varepsilon(x) \mid x \in X\}$ with $U_i = B_\varepsilon(x_i)$ for some $x_i \in X$. Let $\mathbf{V}^0 := \{V_i^0 := \text{cl}(U_i) \mid i = 1, \dots, n\}$. Set $G_1 := V_1^0$. In general define, recursively,

$$\mathbf{V}^k := \{V_i^k := \text{cl}(V_i^{k-1} \setminus G_k) \mid i = k + 1, \dots, n\}$$

for $k = 1, \dots, n - 1$, with $G_k := V_k^{k-1}$. After $n - 1$ steps we let $\mathcal{G} := \{G_i \mid i = 1, \dots, n\}$. Clearly condition (i) of Definition 2.1 is satisfied. The fact that $\text{cl}(\text{int}(V_i^0)) = V_i^0$ can be used to justify condition (ii). Condition (iii) follows from the fact that, for $i > 1$,

$$\text{int}(G_i) = \text{int}(V_i^{i-2} \setminus G_{i-1}) = \{x \mid x \in \text{int}(V_i^{i-2}) \setminus G_{i-1}\}.$$

Clearly $\text{diam } \mathcal{G} = \varepsilon$.

As for the case that X is locally compact and separable we argue as follows. Since the space X is locally compact and separable, there exists a sequence of open relatively compact sets $U_n \subset X$ such that $\text{cl}(U_n) \subset U_{n+1}$ for all $n \geq 0$, and $X = \bigcup_n U_n$. Now define $X_n := \text{cl}(X_n \setminus X_{n-1})$ for $n \geq 1$, and $X_0 = U_0$. For each X_n the above proof provides a finite grid \mathcal{G}_n with $\text{diam } \mathcal{G}_n = \varepsilon$. Therefore, $\mathcal{G} = \bigcup_n \mathcal{G}_n$ is then a grid for X with $\text{diam } \mathcal{G} = \varepsilon$. \square

For a fixed grid we can define various combinatorial descriptions of the map f . The combinatorialization of f given below is natural and best suited for explaining the theory in this paper. We will consider other combinatorial representations of dynamical systems in Section 6.

Definition 2.3. The **minimal multivalued map** associated to f on the grid \mathcal{G} is defined by $\mathcal{F}_{\mathcal{G}}(G) := \{H \in \mathcal{G} \mid H \cap f(G) \neq \emptyset\}$.

If the underlying grid is clear from context, then we simply write \mathcal{F} . To emphasize the fact that \mathcal{F} is multivalued, that is each element of \mathcal{G} is sent to a set of elements of \mathcal{G} , we write $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$. For two multivalued maps $\mathcal{F}, \mathcal{F}': \mathcal{G} \rightrightarrows \mathcal{G}$, the map \mathcal{F}' is said to *enclose* the map \mathcal{F} if $\mathcal{F}(G) \subset \mathcal{F}'(G)$ for every $G \in \mathcal{G}$. The inverse of a given multivalued map \mathcal{F} is defined by

$$\mathcal{F}^{-1}(G) := \{H \in \mathcal{G} \mid G \in \mathcal{F}(H)\}.$$

From the definition of \mathcal{F}^{-1} it follows that enclosure carries over to inverses, i.e. if \mathcal{F}' encloses \mathcal{F} , then $(\mathcal{F}')^{-1}$ encloses \mathcal{F}^{-1} .

The following property is important in the process of going from the combinatorial information contained in \mathcal{F} to the topological properties of f , and is an essential requirement for multivalued maps in the context of combinatorializing dynamical systems.

Definition 2.4. A multivalued map $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ is an **outer approximation** of $f: X \rightarrow X$ if

$$f(G) \subset \text{int}(|\mathcal{F}(G)|)$$

for every $G \in \mathcal{G}$.

Szymczak [16] was the first to identify and make explicit use of this essential concept which is used in Sections 4 through 7 to associate \mathcal{F} with the dynamics of f . The following propositions emphasize the importance of the minimal multivalued map.

Proposition 2.5. *If $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ is the minimal multivalued map associated to f , then \mathcal{F} is an outer approximation of f .*

Proof. Let $y \in f(G)$. Then $y \in H$ for at most finitely many $H \in \mathcal{G}$ all of which are, by definition, in $\mathcal{F}(G)$. If $y \notin \text{int}(|\mathcal{F}(G)|)$, then there exists a sequence $y_n \rightarrow y$ which can be chosen to lie in a single grid element $H_* \notin \mathcal{F}(G)$. However, in this case, $y \in \text{cl}(H_*) = H_*$ so that $H_* \in \mathcal{F}(G)$, which is a contradiction. \square

Observe that determining the minimal multivalued map associated to a map f requires complete knowledge of the image of f on grid elements. The following corollary guarantees that the result is true even if one needs to enlarge the images of \mathcal{F} , for example, to take into account numerical errors.

Corollary 2.6. *If $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ encloses the minimal multivalued map associated to f , then \mathcal{F} is an outer approximation of f .*

A similar proof results in the following proposition used in the theory developed in the subsequent sections.

Proposition 2.7. *If $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ encloses the minimal multivalued map associated to f , then $f^{-1}(G) \subset \text{int}(|\mathcal{F}^{-1}(G)|)$ for any $G \in \mathcal{G}$,*

The minimal multivalued map is a natural setting in which to design algorithms to approximate the discrete dynamics of a map, and general computer software has been developed for this purpose, see GAIO [7], [6]. To obtain rigorous information from these algorithms, one must take into account numerical errors. Thus, a useful extension of Definition 2.3 is to allow for an error in the image $f(G)$. This leads to the following definition of a multivalued map.

Example 2.8. Let $\varepsilon > 0$,

$$\mathcal{F}_\varepsilon(G) := \{H \mid B_\varepsilon(f(G)) \cap H \neq \emptyset\}.$$

Note that \mathcal{F}_ε encloses \mathcal{F} , the minimal multivalued map for f . In practice there are always approximation errors in computing the image of f , since f can be evaluated at only finitely many points, as well as round-off errors. If these errors can be bounded by $\varepsilon > 0$, this example shows that they can be included in the multivalued map in order to get rigorous statements about the dynamics of the original system.

In Section 4 we will describe other extensions and more general multivalued maps. In the next section we give a detailed account of the properties of multivalued maps and their interpretation as both discrete dynamical systems and as directed graphs.

3. Directed Graphs and Discrete Dynamics

Grids and multivalued maps introduced in the previous section provide a framework in which to pass from topology to combinatorics and back. In Sections 3 through 5 we will assume that the metric space X is compact and, therefore, \mathcal{G} is a finite set. It is often useful to treat $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ as a dynamical system defined on a finite set. However, from the point of view of developing efficient algorithms, it is advantageous to simultaneously consider $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ as a directed graph whose vertices are the grid elements with an edge from G to H if $H \in \mathcal{F}(G)$. Thus, throughout this section, we will alternate between concepts from dynamics and graph theory. The notation \mathcal{F} will be used for both a multivalued map and a directed graph. We begin with some graph-theoretic definitions.

Definition 3.1. A directed graph \mathcal{F} is *closed* if $\mathcal{F}(G) \neq \emptyset$ and $\mathcal{F}^{-1}(G) \neq \emptyset$ for every vertex $G \in \mathcal{G}$.

As a multivalued map the closedness property implies that the space of vertices \mathcal{G} gets mapped onto itself by both \mathcal{F} and \mathcal{F}^{-1} . This property is closely related to the surjectivity assumption in Definition 1.1. Unless specified otherwise \mathcal{F} is always assumed to be closed.

Proposition 3.2. *Let $f: X \rightarrow X$ be surjective. If $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ encloses the minimal multivalued map associated to f , then \mathcal{F} is closed.*

Proof. Since $f: X \rightarrow X$ for every $G \in \mathcal{G}$, $f(G) \subset X$, and hence $|\mathcal{F}(G)| \cap X \neq \emptyset$. The assumption that f is surjective implies that for any $G \in \mathcal{G}$ there exists $H \in \mathcal{G}$ such that $f(H) \cap G \neq \emptyset$. Hence, $|\mathcal{F}^{-1}(G)| \cap X \neq \emptyset$. \square

A *complete orbit* of \mathcal{F} through G is the sequence $\gamma_G = \{G_k\}_{k \in \mathbb{Z}}$ satisfying $G = G_0$ and $G_{k+1} \in \mathcal{F}(G_k)$ for every $k \in \mathbb{Z}$. Observe that complete orbits of $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ are equivalent to bi-infinite paths in the digraph \mathcal{F} . We say that a path in \mathcal{F} is *nontrivial* if it has at least one edge. A *forward orbit* $\{G_k \mid k \geq 0\}$ and a *backward orbit* $\{G_k \mid k \leq 0\}$ are denoted by γ_G^+ and γ_G^- , respectively.

As was indicated in the Introduction, the concepts of invariant sets and attractor–repeller pairs are essential for understanding the structure of dynamical systems. In the context of a multivalued map $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ on a finite set, the following definitions are appropriate.

Definition 3.3. A subset $S \subset \mathcal{G}$ is an **invariant set** for \mathcal{F} if $S \subset \mathcal{F}(S)$ and if $S \subset \mathcal{F}^{-1}(S)$.

By symmetry, S is an invariant set for \mathcal{F} if and only if S is an invariant set for \mathcal{F}^{-1} . The following proposition follows immediately from this definition and gives various characterizations of invariant sets which are useful in different contexts.

Proposition 3.4. *Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed, then the following statements are equivalent:*

- (a) \mathcal{S} is an invariant set for \mathcal{F} ;
- (b) $\mathcal{F}_{\mathcal{S}}$ is closed and thus $\mathcal{S} = \bigcap_{n \in \mathbb{Z}} \mathcal{F}^n(\mathcal{S})$; and
- (c) for all $G \in \mathcal{S}$ there exists a complete orbit $\gamma_G \subset \mathcal{S}$.

By restricting \mathcal{F} to an invariant set \mathcal{S} , which we denote by $\mathcal{F}_{\mathcal{S}}$, one obtains a closed graph again on which the theory can be applied.

Proof. We start by showing that (a) implies (b). If $\mathcal{F}(G) = \emptyset$ for some $G \in \mathcal{S}$, then $G \notin \mathcal{F}^{-1}(\mathcal{S})$. Similarly, if $\mathcal{F}^{-1}(G) = \emptyset$ for some $G \in \mathcal{S}$, then $G \notin \mathcal{F}(\mathcal{S})$. Both statements contradict assumption (a).

We now show that (b) implies (c). For $G \in \mathcal{S}$ we construct a complete orbit γ_G by induction. Let $G_0 := G$. Since $\mathcal{F}_{\mathcal{S}}$ is closed, we can choose $G_1 \in \mathcal{F}(G) \cap \mathcal{S}$ and $G_{-1} \in \mathcal{F}^{-1}(G) \cap \mathcal{S}$. Having determined G_k , if $k > 0$, choose $G_{k+1} \in \mathcal{F}(G_k) \cap \mathcal{S}$ and, if $k < 0$, choose $G_{k-1} \in \mathcal{F}^{-1}(G_k) \cap \mathcal{S}$. In either case the choice can be made since $\mathcal{F}_{\mathcal{S}}$ is closed.

Finally, to show that (c) implies (a) we argue as follows. For any $G_1 \in \mathcal{S}$ there exists a $G \in \mathcal{F}^{-1}(G_1) \cap \mathcal{S}$. Moreover, $\mathcal{F}(G) \subset \mathcal{F}(\mathcal{S})$ for all $G \in \mathcal{S}$ so that $G_1 \in \mathcal{F}(\mathcal{S})$, which proves that $\mathcal{S} \subset \mathcal{F}(\mathcal{S})$. The property for \mathcal{F}^{-1} follows from symmetry. \square

From Proposition 3.4(c) it is easily seen that if \mathcal{S} and \mathcal{S}' are invariant sets under \mathcal{F} , then $\mathcal{S} \cup \mathcal{S}'$ is also an invariant set under \mathcal{F} . For $k \leq l$ let

$$\Gamma^{k,l}(\mathcal{U}) := \bigcup_{k \leq n \leq l} \mathcal{F}^n(\mathcal{U}).$$

We write $\Gamma_+^k(\mathcal{U}) = \Gamma^{k,\infty}(\mathcal{U})$ as the k -forward image of \mathcal{U} and, similarly, $\Gamma_-^k(\mathcal{U}) = \Gamma^{-\infty,k}(\mathcal{U})$ as the k -backward image of \mathcal{U} . Moreover, $\Gamma(\mathcal{U})$ denotes the complete image of \mathcal{U} .

Definition 3.5. For a given set $\mathcal{U} \subset \mathcal{G}$ define the ω -limit and α -limit sets of \mathcal{U} by

$$\omega(\mathcal{U}) = \bigcap_{k \geq 0} \Gamma_+^k(\mathcal{U}) \quad \text{and} \quad \alpha(\mathcal{U}) = \bigcap_{k \leq 0} \Gamma_-^k(\mathcal{U}).$$

Proposition 3.6. *Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed and let $\mathcal{U} \subset \mathcal{G}$ with $\mathcal{F}(\mathcal{U}) \subset \mathcal{U}$. Then $\text{Inv}(\mathcal{U}) = \omega(\mathcal{U}) \neq \emptyset$ with $\mathcal{F}(\omega(\mathcal{U})) = \omega(\mathcal{U})$. Similarly, if $\mathcal{F}^{-1}(\mathcal{U}) \subset \mathcal{U}$, then $\text{Inv}(\mathcal{U}) = \alpha(\mathcal{U}) \neq \emptyset$ with $\mathcal{F}^{-1}(\alpha(\mathcal{U})) = \alpha(\mathcal{U})$.*

Proof. We begin by showing that $\text{Inv}(\mathcal{U}) \neq \emptyset$. Since \mathcal{F} is closed, for any $G \in \mathcal{U}$ there exists a forward orbit $\gamma_G^+ \subset \mathcal{U}$. Since \mathcal{U} is finite, there exist $k_1 > k_0 \geq 0$

such that $G_{k_1} = G_{k_0}$. Let $\gamma_H = \{H_k\}_{k \in \mathbb{Z}} \subset \mathcal{U}$ be the orbit $H_k = G_{k_0+k}$ for $k = 0, \dots, k_1 - k_0$ extended periodically with period $k_1 - k_0$. By Proposition 3.4(c), γ_H is invariant, hence $\text{Inv}(\mathcal{U}) \neq \emptyset$.

Now consider any complete orbit $\gamma_G \subset \mathcal{U}$. It is easily checked that $\gamma_G \subset \omega(\mathcal{U})$. Thus, $\text{Inv}(\mathcal{U}) = \omega(\mathcal{U})$. But $\mathcal{F}(\omega(\mathcal{U})) = \omega(\mathcal{U})$, hence $\omega(\mathcal{U})$ is an attractor. \square

Definition 3.7. Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed. A subset $\mathcal{A} \subset \mathcal{G}$ is an **attractor** for \mathcal{F} if $\mathcal{F}(\mathcal{A}) = \mathcal{A}$. The **dual repeller** \mathcal{A}^* of an attractor \mathcal{A} is the maximal attractor for \mathcal{F}^{-1} contained in $\mathcal{G} \setminus \mathcal{A}$, and $(\mathcal{A}, \mathcal{A}^*)$ is called an **attractor–repeller pair**. In general, a **repeller** is an invariant set \mathcal{B} that satisfies $\mathcal{F}^{-1}(\mathcal{B}) = \mathcal{B}$ and thus is an attractor for \mathcal{F}^{-1} .

The following proposition justifies this definition of an attractor–repeller pair for a closed graph.

Proposition 3.8. Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed, and let $(\mathcal{A}, \mathcal{A}^*)$ be an attractor–repeller pair in \mathcal{G} . Then the following statements are true:

- (a) \mathcal{A} and \mathcal{A}^* are invariant sets for \mathcal{F} .
- (b) If γ_G is a complete orbit through $G \in \mathcal{A}$, then $\gamma_G^+ \subset \mathcal{A}$.
- (c) If γ_G is a complete orbit through $G \in \mathcal{A}^*$, then $\gamma_G^- \subset \mathcal{A}^*$.
- (d) $(\mathcal{A}^*, \mathcal{A})$ is an attractor–repeller pair for \mathcal{F}^{-1} .
- (e) Let γ_G be a complete orbit through $G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$. Then:
 - (i) $\gamma_G^- \cap \mathcal{A} = \emptyset$;
 - (ii) $\gamma_G^+ \cap \mathcal{A}^* = \emptyset$;
 - (iii) $\omega(G) \subset \mathcal{A}$ and $\alpha(G) \subset \mathcal{A}^*$; and
 - (iv) $G \neq G_k$ for all $k \neq 0$.

That is, γ_G is a connecting orbit from \mathcal{A}^* to \mathcal{A} .

Proof. To show that \mathcal{A} is invariant it is sufficient to show that $\mathcal{A} \subset \mathcal{F}^{-1}(\mathcal{A})$. So assume that there exists $G \in \mathcal{A}$ such that $G \notin \mathcal{F}^{-1}(\mathcal{A})$. Then $\mathcal{F}(G) \cap \mathcal{A} = \emptyset$. However, $\mathcal{F}(G) \neq \emptyset$ since \mathcal{F} is closed, and hence $\mathcal{F}(\mathcal{A}) \not\subset \mathcal{A}$, a contradiction. A similar argument applies to \mathcal{A}^* . Statement (b) follows from the fact that $\mathcal{F}(\mathcal{A}) = \mathcal{A}$, and (c) follows from $\mathcal{F}^{-1}(\mathcal{A}^*) = \mathcal{A}^*$.

Suppose \mathcal{U} is an attractor in $\mathcal{G} \setminus \mathcal{A}^*$ that contains \mathcal{A} . Let $H \in \mathcal{U} \setminus \mathcal{A}$. By Proposition 3.4, there exists a backward orbit $\gamma_H^- \subset \mathcal{U}$. Since \mathcal{G} is finite, there exists $k, l > 0$ such that $H_{-k-l} = H_{-k}$. If H_{-k} were in \mathcal{A} , then by (b) also $H \in \mathcal{A}$ contrary to our assumption, thus $H_{-k} \notin (\mathcal{A} \cup \mathcal{A}^*)$. Therefore, $\mathcal{B} = \alpha(H_{-k}) \cup \mathcal{A}^* \subset \mathcal{G} \setminus \mathcal{A}$ is an attractor for \mathcal{F}^{-1} since $\mathcal{F}^{-1}(\mathcal{B}) = \mathcal{F}^{-1}(\alpha(H_{-k})) \cup \mathcal{F}^{-1}(\mathcal{A}^*) = \mathcal{B}$. Moreover, $\mathcal{B} \neq \mathcal{A}^*$ since $H_{-k} \in \mathcal{B}$ but $H_{-k} \notin \mathcal{A}^*$, which contradicts the fact that \mathcal{A}^* is the dual repeller of \mathcal{A} . Hence \mathcal{A} is the maximal attractor in $\mathcal{G} \setminus \mathcal{A}^*$ which proves (d).

In statement (e), (i) follows from (b), and (ii) follows from (c). Also (e(iii)) follows from (d) since $\omega(G)$ is an attractor by Proposition 3.6 and $\omega(G) \cap \mathcal{A}^* = \emptyset$ since $\mathcal{F}^{-1}(\mathcal{A}^*) = \mathcal{A}^*$ and $G \notin \mathcal{A}^*$. A similar argument shows $\alpha(G) \subset \mathcal{A}^*$.

Finally, if $G = G_k$ for some $k > 0$, then $G \subset \omega(G)$ which contradicts (e(iii)), and a similar argument holds for $k < 0$. \square

Refining the notion of an attractor–repeller pair leads to the following definition.

Definition 3.9. Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed. A **Morse decomposition** of \mathcal{G} is a finite collection of invariant sets $\mathcal{S}_1, \dots, \mathcal{S}_n$ for which there exists a strict partial ordering $>$ on the index set $\{1, \dots, n\}$ that satisfies the following property. Given a complete orbit $\{G_k\}_{k \in \mathbb{Z}}$ there exist unique $1 \leq i, j \leq n$ and there exist $k^\pm \in \mathbb{Z}$ such that

$$G_k \in \mathcal{S}_i \quad \text{for all } k < k^- \quad \text{and} \quad G_k \in \mathcal{S}_j \quad \text{for all } k > k^+.$$

Moreover, if $i \neq j$, then $i > j$. The sets \mathcal{S}_i are called **Morse sets**.

Off the Morse sets \mathcal{S}_i the multivalued map \mathcal{F} has a gradient-like structure. To make this precise we introduce the notion of a Lyapunov function on a directed graph.

Definition 3.10. A **Lyapunov function** for a Morse decomposition $\{\mathcal{S}_i \mid i = 1, \dots, I\}$ of $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ is a function $V: \mathcal{G} \rightarrow [0, 1]$ satisfying:

- (a) if $G, H \in \mathcal{S}_i$, then $V(G) = V(H)$;
- (b) if $H \in \mathcal{F}(G)$, then $V(G) \geq V(H)$; and
- (c) if $H \in \mathcal{F}(G)$, and G and H do not belong to the same Morse set, then $V(G) > V(H)$.

Lyapunov functions for Morse decompositions of \mathcal{F} can easily be constructed using linear time graph algorithms; see [3]. In Section 6 we discuss the construction of Lyapunov functions which limit on a Lyapunov function for the underlying dynamical system.

It is often useful to be able to partition a digraph using Lyapunov functions. We will use the following notation:

$$V_a = \{G \in \mathcal{G} \mid a \leq V(G)\}, \quad V^b = \{G \in \mathcal{G} \mid V(G) \leq b\}, \quad V_a^b = V_a \cap V^b.$$

The following proposition explains the relationship between Morse sets and attractor–repeller pairs.

Proposition 3.11. *If \mathcal{S} is a Morse set, then there exist attractor–repeller pairs $(\mathcal{A}_i, \mathcal{A}_i^*)$, $i = 1, 2$, such that*

$$\mathcal{S} = \mathcal{A}_1 \cap \mathcal{A}_2^*.$$

Proof. Let V be a Lyapunov function for the Morse decomposition such that $V^{-1}(c) = \mathcal{S}$. Then the sublevel set V^c satisfies $\mathcal{F}(V^c) \subset V^c$ and, hence, V^c

contains an attractor. Since $\mathcal{S} = V^{-1}(c) \subset V^c$ we have $\mathcal{S} \subset \omega(\mathcal{S}) \subset \omega(V^c)$ and, consequently, $\mathcal{S} \subset \omega(V^c)$. The same argument applies to the superlevel set V_c and, thus, $\mathcal{S} \subset \alpha(V_c)$. If we set $\mathcal{A}_1 = \omega(V^c)$ and $\mathcal{A}_2^* = \alpha(V_c)$, then $\mathcal{S} \subset \omega(V^c) = \mathcal{A}_1$, and $\mathcal{S} \subset \alpha(V_c) = \mathcal{A}_2^*$. Consequently, $\mathcal{S} \subset \mathcal{A}_1 \cap \mathcal{A}_2^* \subset V^c \cap V_c = \mathcal{S}$, which proves that $\mathcal{S} = \mathcal{A}_1 \cap \mathcal{A}_2^*$. \square

Using the above proposition we establish a fundamental relationship between Morse decompositions of a graph and the set of attractor–repeller pairs.

Proposition 3.12. *Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed, and let $\{\mathcal{S}_i \mid i = 1, \dots, I\}$ be a Morse decomposition of \mathcal{F} . Then there exists a collection of attractor–repeller pairs $\{(\mathcal{A}_j, \mathcal{A}_j^*) \mid j = 1, \dots, J\}$ such that*

$$\bigcup_{i=1}^I \mathcal{S}_i = \bigcap_{j=1}^J (\mathcal{A}_j \cup \mathcal{A}_j^*).$$

Morse decompositions split the dynamics of \mathcal{F} into gradient-like dynamics off the Morse sets and “recurrent” dynamics on the Morse sets. Next we introduce the notion of recurrence for \mathcal{F} in order to describe the largest set on which \mathcal{F} is gradient-like.

Definition 3.13. The **recurrent set** of a multivalued map $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ is defined as follows:

$$\mathcal{R}(\mathcal{F}) := \{G \in \mathcal{G} \mid \text{there exists a nontrivial path from } G \text{ to } G\}.$$

The following result relates the recurrent set of a graph to the set of attractor–repeller pairs.

Proposition 3.14. *Let $\mathcal{F}: \mathcal{G} \rightrightarrows \mathcal{G}$ be closed. Let the set of all attractor–repeller pairs in \mathcal{F} be given by $\{(\mathcal{A}_j, \mathcal{A}_j^*) \mid j = 1, \dots, J\}$. Then,*

$$\mathcal{R}(\mathcal{F}) = \bigcap_{j=1}^J (\mathcal{A}_j \cup \mathcal{A}_j^*).$$

Proof. We begin by showing that $\mathcal{R}(\mathcal{F}) \subset \bigcap_{j=1}^J (\mathcal{A}_j \cup \mathcal{A}_j^*)$. Let $G \in \mathcal{R}(\mathcal{F})$ and assume that there exists an attractor–repeller pair $(\mathcal{A}_j, \mathcal{A}_j^*)$ such that $G \notin \mathcal{A}_j \cup \mathcal{A}_j^*$. By Proposition 3.8(e(iv)), there is no nontrivial path from G to G , a contradiction.

Now assume $G \in \bigcap_{j=1}^J (\mathcal{A}_j \cup \mathcal{A}_j^*)$ but is not recurrent. Thus, for every complete solution γ_G through G , we have $G \neq G_k$ for all $k > 0$. Define

$$\mathcal{P} := \{H \in \mathcal{G} \mid \text{there exists a nontrivial path from } G \text{ to } H\}.$$

Clearly, $\mathcal{F}(\mathcal{P}) \subset \mathcal{P}$ and $G \notin \mathcal{P}$. By Proposition 3.6, $\omega(\mathcal{P}) = \mathcal{A} \subset \mathcal{P}$ is an attractor which does not contain G . We will reach a contradiction if $G \notin \mathcal{A}^*$. If $G \in \mathcal{A}^*$, then there is a complete path $\gamma_G \subset \mathcal{A}^*$ and $H \in \gamma_G^+ \cap \mathcal{A}$ which would imply that $\mathcal{A} \cap \mathcal{A}^* \neq \emptyset$. Therefore, $G \notin (\mathcal{A} \cup \mathcal{A}^*)$, a contradiction. \square

Definition 3.15. The **components** of $\mathcal{R}(\mathcal{F})$ are the equivalence classes defined by the relation $G \sim H$ if there exist paths from G to H and from H to G .

Observe that components of $\mathcal{R}(\mathcal{F})$ are the Morse sets of the finest Morse decomposition. In the terminology of graph theory, the components of $\mathcal{R}(\mathcal{F})$ are precisely the (nontrivial) strongly connected components of \mathcal{F} which contain at least one edge. This characterization of $\mathcal{R}(\mathcal{F})$ implies that there exists a linear time algorithm to construct a Lyapunov function. Since \mathcal{G} is a finite set, and we are considering Lyapunov functions with images in \mathbb{R} , we can choose $V: \mathcal{G} \rightarrow [0, 1]$ with the following properties.

Proposition 3.16. *There exists a Lyapunov function $V: \mathcal{G} \rightarrow [0, 1]$ such that for any $c \in \mathbb{R}$, $V^{-1}(c)$ is either a single component of $\mathcal{R}(\mathcal{F})$, an element of $\mathcal{G} \setminus \mathcal{R}(\mathcal{F})$, or empty. Furthermore, in the latter two cases $(\omega(V^c), \alpha(V^c))$ is an attractor–repeller pair.*

4. Isolating Neighborhoods

After discretizing both f and X and studying the dynamics of the discrete system \mathcal{F} , the next logical question is whether the characteristics of the discrete dynamics, such as Morse sets, have a bearing on the dynamics of f . We will see in this section that most invariant sets for the discrete dynamics yield information about isolation for the dynamics of f . Recall that a compact set $N \subset X$ is an isolating neighborhood if $\text{Inv}(N, f) \subset \text{int}(N)$. A compact neighborhood that satisfies the property

$$f^{-1}(N) \cap N \cap f(N) \subset \text{int}(N) \quad (1)$$

is clearly an isolating neighborhood. Easton called such an isolating neighborhood an *isolating block* [8]. The main result in this section is the relationship between Morse sets for \mathcal{F} and isolating blocks for f .

Theorem 4.1. *Let \mathcal{S} be a Morse set for \mathcal{F} , an outer approximation of f , then its geometric realization $N = |\mathcal{S}|$ is an isolating block for f .*

Proof. The invariant set \mathcal{S} is a Morse set and, therefore, by Proposition 3.11, it is the intersection of a single attractor \mathcal{A} and a single repeller \mathcal{A}^* , i.e. $\mathcal{S} = \mathcal{A} \cap \mathcal{A}^*$. Now $N = |\mathcal{S}| = |\mathcal{A} \cap \mathcal{A}^*|$ and, by construction, $\mathcal{F}(\mathcal{S}) = \mathcal{F}(\mathcal{A} \cap \mathcal{A}^*) \subset \mathcal{F}(\mathcal{A}) =$

\mathcal{A} , and $\mathcal{F}^{-1}(\mathcal{S}) \subset \mathcal{A}^*$. We obtain

$$\begin{aligned} f(N) &= \bigcup_{G \in \mathcal{S}} f(G) \subset \bigcup_{G \in \mathcal{S}} \text{int}(|\mathcal{F}(G)|) \\ &= \text{int}(|\mathcal{F}(\mathcal{S})|) \subset \text{int}(|\mathcal{A}|). \end{aligned}$$

Similarly for the repeller we obtain

$$\begin{aligned} f^{-1}(N) &= \bigcup_{G \in \mathcal{S}} f^{-1}(G) \subset \bigcup_{G \in \mathcal{S}} \text{int}(|\mathcal{F}^{-1}(G)|) \\ &= \text{int}(|\mathcal{F}^{-1}(\mathcal{S})|) \subset \text{int}(|\mathcal{A}^*|). \end{aligned}$$

Combining these two inclusions yields the following inclusion:

$$\begin{aligned} f^{-1}(N) \cap N \cap f(N) &\subset \text{int}(|\mathcal{A}^*|) \cap |\mathcal{A} \cap \mathcal{A}^*| \cap \text{int}(|\mathcal{A}|) \\ &\subset \text{int}(|\mathcal{A} \cap \mathcal{A}^*|) = \text{int}(N), \end{aligned}$$

which proves that N is an isolating block. \square

An important byproduct of the above theorem is that one can also find attractor–repeller pairs for f via attractor–repeller pairs for \mathcal{F} . See Proposition 5.5. The fact that attractor–repeller pairs for \mathcal{F} can be used to find attractor–repeller pairs for f can be further generalized to Morse decompositions in the sense that a Morse decomposition for \mathcal{F} will provide isolating blocks for a Morse decomposition for f .

Corollary 4.2. *Let $\{S_i\}$, $i \in (\mathcal{I}, >)$, be a Morse decomposition for \mathcal{F} . Then the collection of sets $\{S_i\}$, with $S_i = \text{Inv}(|S_i|, f)$, form a Morse decomposition for f .*

Example 4.3. A simple counterexample shows that invariant sets \mathcal{S} that are not Morse sets do not yield isolating neighborhoods in general. Indeed, consider the following example on the 2-disk $X = D^2$. Let $f: D^2 \rightarrow D^2$ be a rotation over 90 degrees. Consider the concentric circles around the origin with radii $r = 1/N, 2/N, \dots, 1$. This gives a grid \mathcal{G} for D^2 . Let \mathcal{F} be the minimal multivalued map defined in Definition 2.3. One can now find various invariant sets which do not yield isolating neighborhoods.

Morse sets are obvious candidates for finding isolating neighborhoods. However, there may be invariant sets which are not Morse sets but still yield isolating neighborhoods. If a Morse set for \mathcal{F} is minimal (recurrent component), i.e., it does not allow further attractor–repeller pair decompositions, then it may still have invariant subsets. Sometimes these sets provide additional isolating neighborhoods and reveal additional information about the structure of the dynamics of f at the recurrent components. In order to achieve this we introduce the notion of isolated invariant sets relative to a Morse set.

An invariant set \mathcal{S} is *isolated with respect to a Morse set \mathcal{S}'* if

$$\text{Inv}(\{G \in \mathcal{S}' \mid G \cap |\mathcal{S}| \neq \emptyset\}) = \mathcal{S}.$$

Note that one choice of \mathcal{S}' is the whole grid \mathcal{G} . This notion of isolation combined with the outer approximation property also leads to isolating blocks. We can now extend Theorem 4.1 as follows:

Theorem 4.4. *Let \mathcal{S} be an isolated invariant set with respect to a Morse set \mathcal{S}' . Then $|\mathcal{S}|$ is an isolating block for f .*

Proof. By assumption, $N' = |\mathcal{S}'|$ is an isolating block. Then, for the isolated invariant set $\mathcal{S} \subset \mathcal{S}'$, set $N = |\mathcal{S}|$. The condition that N is also an isolating block is equivalent to saying that boundary points leave N immediately in either forward or backward time, i.e., for $x \in \partial N$, $f(x) \notin N$, or $f^{-1}(x) \notin N$. The boundary of N splits into two parts; $\partial N = M_{\text{lat}} \cup M_{\text{int}}$, where $M_{\text{lat}} \subset \partial N'$ and $M_{\text{int}} \subset \text{int } N'$. If $x \in M_{\text{lat}}$, then since N' is an isolating block either $f(x)$ or $f^{-1}(x)$ is not in N' , and therefore not in N . What remains to be shown is that the same holds for points in M_{int} . Suppose not, then $f(x), f^{-1}(x) \in N$. There exist grid elements $G^\pm \in \mathcal{S}$ such that $f^{-1}(x) \in G^-$ and $f(x) \in G^+$. The invariance of \mathcal{S} now implies that there exist orbits $\gamma_{G^\pm}^\pm \subset \mathcal{S}$. Since $x \in M_{\text{int}}$ and \mathcal{S} is isolated in \mathcal{S}' , we can choose $H \in \mathcal{S}' \setminus \mathcal{S}$ which contains x .

Using the fact that \mathcal{F} is an outer approximation now implies that $f(x) \in f(H) \subset \text{int } |\mathcal{F}(H)|$ and, consequently, $f(H) \cap G^+ \neq \emptyset$ so that $G^+ \in \mathcal{F}(H)$. In same way one proves that $G^- \in \mathcal{F}^{-1}(H)$. We can now construct the following complete orbit

$$\gamma_{G^-}^- \cup \{H\} \cup \gamma_{G^+}^+ \subset \mathcal{S}'$$

which does not lie in \mathcal{S} , implying that \mathcal{S} is not isolated with respect to \mathcal{S}' , a contradiction. \square

The theoretical framework outlined in Sections 2–4 extends the ideas introduced by Szymczak in [16]. Using implementations on \mathbb{R}^n with rectangular grids by software packages such as GAIO [6], [7], one can produce isolating blocks (or, more generally, index pairs), index filtrations, and rigorous computer-assisted proofs of specific dynamics, such as periodic or connecting orbits, in both finite- and infinite-dimensional maps, see [16], [4], [5], [9].

5. Chain Recurrence

In this section we continue our exposition for the case of a discrete dynamical system generated by a surjective map $f: X \rightarrow X$ on a compact metric space, and for any grid \mathcal{G} the multivalued map \mathcal{F} will be the minimal multivalued map as

in Definition 2.3. In Section 7 we will state conditions which guarantee that the results of this section hold for more general dynamical systems and multivalued maps.

Definition 5.1. Let $\{\mathcal{G}_n\}_{n \in \mathbb{N}}$ be any sequence of grids for which $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$. The **chain recurrent set** for f is defined by $\mathcal{R}(X, f) := \bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|$.

Obviously, the first order of business is to prove that $\mathcal{R}(X, f)$ is well defined. We begin with a lemma which relates the multivalued maps on grids of different diameters using only the continuity of f .

Lemma 5.2. *Let \mathcal{G} be a grid with diameter $\text{diam}(\mathcal{G})$. Then there exists a $\delta = \delta(\text{diam}(\mathcal{G})) > 0$ such that if \mathcal{H} is a grid with $\text{diam}(\mathcal{H}) < \delta$, then $\mathcal{F}_{\mathcal{H}}(H) \subset \mathcal{F}_{\mathcal{G}}(G)$ for all $G \in \mathcal{G}$ and all $H \in \mathcal{H}$ with $G \cap H \neq \emptyset$. Moreover, if $H \in \mathcal{R}(\mathcal{F}_{\mathcal{H}})$ and $G \cap H \neq \emptyset$, then $G \in \mathcal{R}(\mathcal{F}_{\mathcal{G}})$. Hence $|\mathcal{R}(\mathcal{F}_{\mathcal{H}})| \subset |\mathcal{R}(\mathcal{F}_{\mathcal{G}})|$.*

Proof. Since \mathcal{G} has only finitely many elements, Proposition 2.5 implies that there exists $\varepsilon > 0$ such that $B_\varepsilon(f(G)) \subset \text{int}(|\mathcal{F}_{\mathcal{G}}(G)|)$ for all $G \in \mathcal{G}$. Since f is uniformly continuous on X , there exists $\rho > 0$ such that $\text{diam}(G) < \rho$ implies $\text{diam}(f(G)) < \varepsilon/2$. Choose $\delta > 0$ such that $\delta < \min\{\rho, \varepsilon/2\}$. Let \mathcal{H} be any grid with $\text{diam}(\mathcal{H}) < \delta$. Let $G \in \mathcal{G}$ and $H \in \mathcal{H}$ with $G \cap H \neq \emptyset$. Then $f(H) \subset B_{\varepsilon/2}(f(G))$ since $\text{diam}(H) < \rho$. Also, $|\mathcal{F}_{\mathcal{H}}(H)| \subset B_\varepsilon(f(G))$ since $\text{diam}(\mathcal{H}) < \varepsilon/2$. Therefore $|\mathcal{F}_{\mathcal{H}}(H)| \subset \text{int}(|\mathcal{F}_{\mathcal{G}}(G)|)$.

If $H \in \mathcal{R}(\mathcal{F}_{\mathcal{H}})$, then there exists $k > 0$ such that $H \in \mathcal{F}_{\mathcal{H}}^k(H)$. Note that $|\mathcal{F}_{\mathcal{H}}^k(H)| \subset \text{int}(|\mathcal{F}_{\mathcal{G}}^k(G)|)$ for all $k > 0$. Indeed, since $|\mathcal{F}_{\mathcal{H}}(H)| \subset \text{int}(|\mathcal{F}_{\mathcal{G}}(G)|)$, for any element $K \in \mathcal{F}(H)$ we have $K \cap \text{int}(G') \neq \emptyset$ for some $G' \in \mathcal{F}(G)$ which implies that $|\mathcal{F}_{\mathcal{H}}(K)| \subset \text{int}(|\mathcal{F}_{\mathcal{G}}(G')|) \subset \text{int}(|\mathcal{F}_{\mathcal{G}}^2(G)|)$; hence $|\mathcal{F}_{\mathcal{H}}^2(H)| \subset \text{int}(|\mathcal{F}_{\mathcal{G}}^2(G)|)$. Repeating this argument gives $|\mathcal{F}_{\mathcal{H}}^k(H)| \subset \text{int}(|\mathcal{F}_{\mathcal{G}}^k(G)|)$ for all $k > 0$. Now we have $H \subset \text{int}(|\mathcal{F}_{\mathcal{G}}^k(G)|)$. Since $G \cap H \neq \emptyset$, we conclude that $G \subset \mathcal{F}_{\mathcal{G}}^k(G)$, i.e., $G \in \mathcal{R}(\mathcal{F}_{\mathcal{G}})$. \square

Lemma 5.3. *The set $\bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|$ is independent of the particular sequence $\{\mathcal{G}_n\}_{n \in \mathbb{N}}$ with $\text{diam}(\mathcal{G}_n) \rightarrow 0$. Hence, $\mathcal{R}(X, f)$ is well defined.*

Proof. Suppose \mathcal{G} is a fixed grid and $\{\mathcal{H}_n\}_{n \in \mathbb{N}}$ is a sequence of grids with $\text{diam}(\mathcal{H}_n) \rightarrow 0$ as $n \rightarrow \infty$. We first show that $\bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{H}_n})| \subset |\mathcal{R}(\mathcal{F}_{\mathcal{G}})|$. Choose $n > 0$ so that $\text{diam}(\mathcal{H}_n) < \delta$ where $\delta > 0$ is associated to \mathcal{G} by Lemma 5.2. Let $y \in \bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{H}_n})|$. Then there exists $H \in \mathcal{H}_n$ such that $y \in H \in \mathcal{R}(\mathcal{F}_{\mathcal{H}_n})$. Choose $G \in \mathcal{G}$ so that $y \in G$. By Lemma 5.2 we have $G \in \mathcal{R}(\mathcal{F}_{\mathcal{G}})$ and, hence, $y \in |\mathcal{R}(\mathcal{F}_{\mathcal{G}})|$.

Finally, suppose $\{\mathcal{G}_m\}_{m \in \mathbb{N}}$ and $\{\mathcal{H}_n\}_{n \in \mathbb{N}}$ are two sequences with $\text{diam}(\mathcal{G}_m) \rightarrow 0$ and $\text{diam}(\mathcal{H}_n) \rightarrow 0$ as $m, n \rightarrow \infty$. Above we showed that $\bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{H}_n})| \subset |\mathcal{R}(\mathcal{F}_{\mathcal{G}_m})|$ for each $m > 0$. Hence $\bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{H}_n})| \subset \bigcap_m |\mathcal{R}(\mathcal{F}_{\mathcal{G}_m})|$. Interchanging

the roles of \mathcal{G}_m and \mathcal{H}_n gives $\bigcap_m |\mathcal{R}(\mathcal{F}_{\mathcal{H}_n})| = \bigcap_m |\mathcal{R}(\mathcal{F}_{\mathcal{G}_m})|$. Hence $\mathcal{R}(X, f)$ is well defined. \square

Now we establish a basic lemma relating the orbits through points x under f to the orbits of grid elements $G \ni x$ under \mathcal{F} . Here we will denote by $\gamma_x^{k,l}$ the orbit segment $\varphi([k, l] \cap \mathbb{T}, x)$ for $0 \leq k \leq l < \infty$. While the proof of this lemma is a direct consequence of uniform continuity, we provide a proof here because it is used extensively in the sequel. We denote by $h(A, B)$ the Hausdorff distance between the sets A and B .

Proposition 5.4. *Let $\varepsilon > 0$, $0 \neq n \in \mathbb{Z}$, and $x \in X$. Then there exists a $\delta > 0$ such that, for any grid \mathcal{G} with $\text{diam}(\mathcal{G}) < \delta$, we have $|\mathcal{F}^n(G)| \subset B_\varepsilon(f^n(G))$ for all $G \in \mathcal{G}$. Moreover, if $x \in G$ and $k, l \in \mathbb{Z}$ with $k \leq l < 0$ or $0 < k \leq l$, then $h(\gamma_x^{k,l}, |\Gamma^{k,l}(G)|) \rightarrow 0$ as $\text{diam}(\mathcal{G}) \rightarrow 0$.*

Proof. Let $\delta_0 = \varepsilon/2$. First suppose $n > 0$. Since $f^{n-1}(G)$ is compact and f is uniformly continuous, there exists $\delta_1 > 0$ such that $f(B_{\delta_1}(f^{n-1}(G))) \subset B_{\delta_0}(f^n(G))$. Furthermore, for each $1 < i < n$ there exists $\delta_i > 0$ such that $f(B_{\delta_i}(f^{n-i}(G))) \subset B_{\delta_{i-1}/2}(f^{n-i+1}(G))$. Let $\delta := \min_{0 \leq i \leq n-1} \delta_i/2$. If $\text{diam}(\mathcal{G}) < \delta$, then $|\mathcal{F}(G)| \subset B_{\delta_{n-1}}(f(G))$. For any $H \in \mathcal{F}(G)$, we have $f(H) \subset B_{\delta_{n-2}/2}(f^2(G))$. Hence $|\mathcal{F}^2(G)| \subset B_{\delta_{n-2}}(f^2(G))$. Repeating this argument for each $i < n$, we obtain $|\mathcal{F}^n(G)| \subset B_{\delta_0}(G) = B_\varepsilon(f^n(G))$. When $n < 0$, there exists $\delta_1 > 0$ such that $f^{-1}(B_{\delta_1}(f^{-n+1}(G))) \subset B_{\delta_0}(f^{-n}(G))$, and the rest of the arguments are analogous.

By uniform continuity of f , choose $\delta_* > 0$ such that $f^n(B_{\delta_*}(x)) \subset B_\varepsilon(f^n(x))$ for all $k \leq n \leq l$. Then, for $\text{diam}(\mathcal{G}) < \min\{\delta, \delta_*\}$ we have $|\mathcal{F}^n(G)| \subset B_\varepsilon(f^n(G)) \subset B_{2\varepsilon}(f^n(x))$ for all $k \leq n \leq l$. The convergence of orbit segments in the Hausdorff metric now follows. \square

One of the fundamental ideas in the proof of the decomposition theorem in Section 6 is the relationship between the chain recurrent set and the attractor–repeller pairs in X , which we now prove. This also establishes that our definition of the chain recurrent set yields the same set as the standard definition. First we need a proposition that relates attractor–repeller pairs for the multivalued map \mathcal{F} to attractor–repeller pairs for f .

Proposition 5.5. (i) *For every attractor–repeller pair (A, A^*) for \mathcal{F} , there exists a unique attractor–repeller pair (A, A^*) for f such that $A \subset |\mathcal{A}|$ and $A^* \subset |\mathcal{A}^*|$.*

(ii) *Let (A, A^*) be an attractor–repeller pair for f . For every $0 < \varepsilon < \text{dist}(A, A^*)/2$, there exists $\delta > 0$ such that if \mathcal{G} is a grid with $\text{diam}(\mathcal{G}) < \delta$, then there exists a unique attractor–repeller pair (A, A^*) for \mathcal{F} with the property that $A \subset |\mathcal{A}| \subset B_\varepsilon(A)$ and $A^* \subset |\mathcal{A}^*| \subset B_\varepsilon(A^*)$. Any other attractor–repeller pair $(\mathcal{B}, \mathcal{B}^*)$ for \mathcal{F} has the property that either $|\mathcal{B}|$ or $|\mathcal{B}^*|$ has nonempty intersection with both $B_\varepsilon(A)$ and $B_\varepsilon(A^*)$.*

Proof. (i) Since $\mathcal{F}(\mathcal{A}) = \mathcal{A}$, we have $f(|\mathcal{A}|) \subset \text{int}(|\mathcal{A}|)$ so that $|\mathcal{A}|$ is an attracting neighborhood. Then $A = \omega(|\mathcal{A}|) \subset |\mathcal{A}|$ is the maximal attractor in $|\mathcal{A}|$, and $A^* = \text{Inv}(\text{cl}(X \setminus |\mathcal{A}|)) \subset |\text{Inv}(\widehat{\mathcal{G}} \setminus \widehat{\mathcal{A}})| = |\mathcal{A}^*|$. Moreover, any other attractor $\widehat{A} \subset |\mathcal{A}|$ is not maximal and, hence, $\widehat{A}^* \cap |\mathcal{A}| \neq \emptyset$ so that $\widehat{A}^* \not\subset |\mathcal{A}^*|$, which implies uniqueness.

(ii) By the definition of attractor, there exists a neighborhood $V \subset B_\varepsilon(A)$ of A with the property that $f(V) \subset V$. Similarly, there exists a neighborhood $V^* \subset B_\varepsilon(A^*)$ of A^* such that $f^{-1}(V^*) \subset V^*$. Since $\text{dist}(A, A^*) > 2\varepsilon$, we have $V \cap V^* = \emptyset$.

By definition of an attractor–repeller pair and compactness, there exists $N > 0$ such that $f^N(x) \in V$ for all $x \in X \setminus f^{-1}(V^*)$. Since V is open, Proposition 5.4 provides $\delta > 0$ such that if $\text{diam}(\mathcal{G}) < \delta$, then $|\mathcal{F}^N(G)| \subset V$ for any $G \subset X \setminus f^{-1}(V^*)$. Also, we can choose $\delta < h(V, f(V))$.

Recall that for any set $S \subset X$ we define $\text{cov}(S) := \{G \in \mathcal{G} \mid S \cap G \neq \emptyset\}$. Let $\mathcal{C} = \text{cov}(f(V))$. Since $f(V) \subset V$ and $\delta < h(V, f(V))$, we have $\mathcal{F}(\mathcal{C}) \subset \mathcal{C}$. Now let $\mathcal{A} := \omega(\mathcal{C})$ and let \mathcal{A}^* be its dual repeller. Since $\mathcal{A} \subset \mathcal{C}$, we have $|\mathcal{A}| \subset B_\varepsilon(A)$. Moreover, since $f(\mathcal{A}) = \mathcal{A}$, we have $\text{cov}(\mathcal{A}) \subset \mathcal{F}(\text{cov}(\mathcal{A}))$, and hence $\text{cov}(\mathcal{A}) \subset \mathcal{F}^k(\text{cov}(\mathcal{A}))$ for all $k > 0$. Thus $\text{cov}(\mathcal{A}) \subset \omega(\text{cov}(\mathcal{A})) \subset \omega(\mathcal{C}) = \mathcal{A}$ so that $A \subset |\mathcal{A}|$.

Suppose G is an element which does not intersect $f^{-1}(V^*)$. Then by the above construction, $\mathcal{F}^{N+1}(G) \subset \mathcal{C}$ so that \mathcal{A}^* does not contain G . Thus $|\mathcal{A}^*| \subset V^* \subset B_\varepsilon(A^*)$. Moreover, $A^* \subset f^{-1}(A^*)$ implies that $\text{cov}(A^*) \subset \mathcal{F}^{-1}(\text{cov}(A^*))$, and hence $\text{cov}(A^*) \subset \mathcal{F}^{-k}(\text{cov}(A^*))$ for all $k > 0$. Thus $\text{cov}(A^*) \subset \alpha(\text{cov}(A^*)) \subset \alpha(\mathcal{G} \setminus \mathcal{C}) = \mathcal{A}^*$ so that $A^* \subset |\mathcal{A}^*|$.

Suppose $(\mathcal{B}, \mathcal{B}^*)$ is any other attractor–repeller pair for \mathcal{F} . By Proposition 3.8 all recurrent elements of \mathcal{G} are in either \mathcal{A} or \mathcal{A}^* and, likewise, in either \mathcal{B} or \mathcal{B}^* . Therefore we must show that either \mathcal{B} or \mathcal{B}^* contains recurrent elements from both \mathcal{A} and \mathcal{A}^* which will imply that either \mathcal{B} or \mathcal{B}^* intersects the ε -neighborhoods of both A and A^* .

First suppose that $\mathcal{A} \cap \mathcal{B} \neq \emptyset$ and $\mathcal{B} \not\subset \mathcal{A}$. Then \mathcal{B} contains a recurrent element of \mathcal{A} , since otherwise it would not be invariant, and likewise \mathcal{B} contains a recurrent element from \mathcal{A}^* since $\mathcal{B} \not\subset \mathcal{A}$. Conversely, if $\mathcal{B} \subset \mathcal{A}$ or $\mathcal{A} \cap \mathcal{B} = \emptyset$, then by invariance, \mathcal{B}^* contains a recurrent element of \mathcal{A} since $\mathcal{A} \not\subset \mathcal{B}$. In the first case, \mathcal{B}^* contains all of \mathcal{A}^* and, in the latter case, $\mathcal{B} \subset \mathcal{A}^*$ but $\mathcal{B} \neq \mathcal{A}^*$ and, hence, \mathcal{B}^* also contains a recurrent element of \mathcal{A}^* . \square

Theorem 5.6. *The chain recurrent set $\mathcal{R}(X, f)$ is the intersection of all attractor–repeller pairs of (X, f) .*

Proof. Let (A, A^*) be an attractor–repeller pair for f . Suppose $\{\mathcal{G}_n\}$ is a sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$. Let $\varepsilon > 0$. By Proposition 5.5 there exists $N > 0$ and an attractor–repeller pair $(\mathcal{A}_N, \mathcal{A}_N^*)$ for $\mathcal{F}_{\mathcal{G}_N}$ such that $|\mathcal{A}_N \cup \mathcal{A}_N^*| \subset B_\varepsilon(A \cup A^*)$. By definition of an attractor–repeller pair, $|\mathcal{R}(\mathcal{F}_{\mathcal{G}_N})| \subset |\mathcal{A}_N \cup \mathcal{A}_N^*|$ and, thus, $\mathcal{R}(X, f) = \bigcap_n |\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})| \subset B_\varepsilon(A \cup A^*)$. Since $\varepsilon > 0$ was

arbitrary, $\mathcal{R}(X, f) \subset A \cup A^*$ for every attractor–repeller pair; hence $\mathcal{R}(X, f) \subset \bigcap (A \cup A^*)$.

Let $x \in A \cup A^*$ for all attractor–repeller pairs and let \mathcal{G} be any grid. Suppose $x \notin |\mathcal{R}(\mathcal{F})|$. Choose G containing x , and let $c = V(G)$ where V is a Lyapunov function on \mathcal{G} from Proposition 3.16. Then the sublevel and superlevel sets V^c and V_c contain an attractor \mathcal{A} and its dual repeller \mathcal{A}^* , respectively. Moreover, \mathcal{A} and \mathcal{A}^* do not contain G . Therefore, $|\mathcal{A}|$ and $|\mathcal{A}^*|$ contain an attractor A and its dual repeller A^* for which $x \notin A \cup A^*$, which is a contradiction. Thus $x \in |\mathcal{R}(\mathcal{F})|$. Since \mathcal{G} was an arbitrary grid, $x \in \mathcal{R}(X, f)$. \square

We now investigate how well grid recurrent sets approximate the chain recurrent set. From our definition, it is not a priori clear that grid recurrent sets $\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})$ with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$ cannot have elements which are uniformly bounded away from $\mathcal{R}(X, f)$ but nevertheless disappear in their intersection. Our next lemma will show that this cannot happen.

Lemma 5.7. *If \mathcal{G}_n is any sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$, then $h(|\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|, \mathcal{R}(X, f)) \rightarrow 0$ as $n \rightarrow \infty$.*

Proof. Note that since $\mathcal{F}_{\mathcal{G}_n}$ is closed, $\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})$ is nonempty, and since f is surjective on X (compact), $\mathcal{R}(X, f)$ is also nonempty. Hence their Hausdorff distance is well defined.

Since $\mathcal{R}(X, f) \subset |\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|$, we have

$$h(|\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|, \mathcal{R}(X, f)) = \max_{y \in |\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|} \text{dist}(y, \mathcal{R}(X, f)).$$

Suppose that there exists $\alpha > 0$ such that $h(|\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|, \mathcal{R}(X, f)) \geq \alpha$ for infinitely many $n > 0$. Then one can choose a sequence $x_n \rightarrow x$ such that $x_n \in |\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|$ and $\text{dist}(x_n, \mathcal{R}(X, f)) \geq \alpha$. Thus $\text{dist}(x, \mathcal{R}(X, f)) \geq \alpha$ and, hence, $x \notin \mathcal{R}(X, f)$.

Since $x \notin \mathcal{R}(X, f)$ there must exist $m > 0$ such that $x \notin |\mathcal{R}(\mathcal{F}_{\mathcal{G}_m})|$. In the fixed grid \mathcal{G}_m there must be an element G containing x and x_n for infinitely many $n > 0$. Since $x \notin |\mathcal{R}(\mathcal{F}_{\mathcal{G}_m})|$, also $G \notin \mathcal{R}(\mathcal{F}_{\mathcal{G}_m})$. There exists $N > 0$ large enough so that $\text{diam}(\mathcal{G}_N) < \delta$ where $\delta > 0$ is associated to \mathcal{G}_m by Lemma 5.2. Since $x_N \in H_N \in \mathcal{R}(\mathcal{F}_{\mathcal{G}_N})$ and $G \cap H_N \neq \emptyset$, Lemma 5.2 implies that $G \in \mathcal{R}(\mathcal{F}_{\mathcal{G}_m})$, which is a contradiction. Thus $h(|\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})|, \mathcal{R}(X, f)) \rightarrow 0$ as $n \rightarrow \infty$. \square

The above lemma shows that the chain recurrent set can be approximated arbitrarily closely by grid recurrent sets by taking fine enough grids. For completeness we show that this property holds for isolated invariant sets, which also yields a proof that every isolating neighborhood N contains an isolating block for $\text{Inv}(N)$.

Theorem 5.8. *Let N be an isolating neighborhood and let $S = \text{Inv}(N)$. There exists $\delta > 0$ such that if \mathcal{G} is a grid on X with $\text{diam}(\mathcal{G}) < \delta$, then $S = \text{Inv}(\text{cov}(N))$*

is an isolated invariant set in \mathcal{G} with $S = \text{Inv}(|\mathcal{S}|) \subset \text{int}(|\mathcal{S}|)$. Moreover, if $\text{diam}(\mathcal{G}_n) \rightarrow 0$, then $h(|\mathcal{S}_n|, S) \rightarrow 0$ as $n \rightarrow \infty$.

Proof. Let $\mathcal{N} = \text{cov}(N)$. Let $0 < \varepsilon < \text{dist}(S, \partial N)/2$ and let $M = N \setminus B_\varepsilon(S)$. If $x \in M$, then for every complete orbit γ_x there exists $n_x \in \mathbb{Z}$ such that $\gamma_x(n_x) \notin N$, since $x \notin S$. For $n_x > 0$, by continuity, there is $r_x > 0$ such that $f^{n_x}(B_{r_x}(x)) \subset N^c$. Since M is compact, $n \in \mathbb{Z}$ can be chosen so that $|n_x| \leq n$ for every $x \in M$. By Proposition 5.4 there exists $0 < \delta < \varepsilon/4$ such that if $\text{diam}(\mathcal{G}) < \delta$, then if $x \in M$ and $x \in G$ we have $\text{dist}(f^{n_x}(x), |\mathcal{F}^{n_x}(G)|) < \varepsilon/2$ so that $G \notin \text{Inv}(\mathcal{N})$. A similar argument holds for $n_x < 0$. Therefore $|\mathcal{S}| = |\text{Inv}(\mathcal{N})| \subset \text{cl}(B_{\varepsilon/4}(S))$. Letting $\varepsilon \rightarrow 0$ yields the Hausdorff convergence of $|\mathcal{S}_n|$ to S . \square

Corollary 5.9. *Let N be an isolating neighborhood with $S = \text{Inv}(N)$. Then there exists a compact set $B \subset \text{int}(N)$ which is an isolating block for S .*

Proof. By Theorem 4.4, $B = |\mathcal{S}|$, in Theorem 5.8, is an isolating block. \square

6. Lyapunov Functions and Conley's Decomposition Theorem

In Section 3 we introduced discrete Lyapunov functions for a multivalued map \mathcal{F} on a grid \mathcal{G} which are naturally obtained from standard graph-theoretic algorithms. In Section 5 we showed that the chain recurrent set of a map can be approximated by recurrent sets of multivalued maps on grids. However, if we consider any sequence of grids \mathcal{G}_n , with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$, then an associated sequence of any such discrete Lyapunov functions need not have a limit. It is natural to ask whether the discrete Lyapunov functions on \mathcal{F} can be chosen to converge to a Lyapunov function for the original dynamical system. In this section we construct specific discrete Lyapunov functions that have continuous limits, which provides an independent proof of Conley's Fundamental Decomposition Theorem. We implement this construction for the example of the multivalued map defined in Definition 2.3 in Section 2. Later we provide axioms under which this procedure works for more general systems.

6.1. Discrete Lyapunov Functions for Attractor–Repeller Pairs

The construction and convergence of discrete Lyapunov functions is performed in several stages. We begin by constructing discrete Lyapunov functions for attractor–repeller pairs which incorporate the topology of X (compact).

6.1.1. Distance Potentials. Let $A, B \subset X$ be compact subsets of X where the intersection $A \cap B$ need not be empty.

Definition 6.1. A function $\mathbf{v}_{(A,B)}: X \setminus (A \cap B) \rightarrow [0, 1]$ is called a **distance potential** for the pair (A, B) if:

- (a) $\mathbf{v}_{(A,B)}$ is locally Lipschitz continuous on $X \setminus (A \cap B)$;
- (b) $\mathbf{v}_{(A,B)}^{-1}(0) = A \setminus B$; and
- (c) $\mathbf{v}_{(A,B)}^{-1}(1) = B \setminus A$.

If $A \cap B = \emptyset$, then a distance potential is defined on all of X and, hence, Lipschitz continuous on X since the space X is compact. Distance potentials can easily be constructed as the following example shows. For $x \in X \setminus (A \cap B)$ define

$$\mathbf{v}_{(A,B)}(x) := \frac{d(x, A)}{d(x, A) + d(x, B)}, \quad (2)$$

where the distance function $d(\cdot, \cdot)$ is derived from the metric on X . It is obvious that if $x \in A \setminus (A \cap B)$, or $x \in B \setminus (A \cap B)$, then $\mathbf{v}_{(A,B)}$ as defined in (2) is equal to 0 or 1, respectively. As for the local Lipschitz continuity we have

Lemma 6.2. *The function $\mathbf{v}_{(A,B)}$ defined in (2) is locally Lipschitz continuous on $X \setminus (A \cap B)$. If $A \cap B = \emptyset$, then $\mathbf{v}_{(A,B)}$ is globally Lipschitz with Lipschitz constant $\text{Lip}_v \leq 1/d(A, B)$.*

Proof. To establish Lipschitz continuity, let $x, y \in X \setminus (A \cap B)$, then

$$\begin{aligned} \mathbf{v}_{(A,B)}(x) - \mathbf{v}_{(A,B)}(y) &= \frac{d(x, A)}{d(x, A) + d(x, B)} - \frac{d(y, A)}{d(y, A) + d(y, B)} \\ &= \frac{d(x, A)d(y, B) - d(y, A)d(x, B)}{(d(x, A) + d(x, B))(d(y, A) + d(y, B))}. \end{aligned}$$

Using the inequality $d(x, A) \leq d(x, y) + d(y, A)$ for A , and the same for B , we obtain

$$\begin{aligned} |\mathbf{v}_{(A,B)}(x) - \mathbf{v}_{(A,B)}(y)| &\leq \frac{d(x, y)(d(y, A) + d(y, B))}{(d(x, A) + d(x, B))(d(y, A) + d(y, B))} \\ &\leq \frac{d(x, y)}{d(x, A) + d(x, B)} \end{aligned}$$

which proves local Lipschitz continuity.

As for the second part of the lemma we argue as follows. Since A and B are compact sets in a compact metric space, when $A \cap B = \emptyset$, they are separated, and thus $d(A, B) > 0$. Using the inequality $d(x, A) + d(x, B) \geq d(A, B) > 0$ we obtain that $|\mathbf{v}_{(A,B)}(x) - \mathbf{v}_{(A,B)}(y)| \leq d(x, y)/d(A, B)$, which proves global Lipschitz continuity. \square

Other examples of distance potentials are found by slightly adjusting the example given by (2):

$$v_{(A,B)}(x) := \frac{d(x, A)^n}{d(x, A)^n + d(x, B)^n}, \quad n \in \mathbb{N}.$$

This is one among many possibilities for constructing distance potentials. The advantage of the above formula is that for $n > 1$ the associated Lyapunov functions in the continuous limit may display more smoothness.

6.1.2. *Discrete Lyapunov Functions for \mathcal{F} .* Let $\mathcal{P} = (\mathcal{A}, \mathcal{A}^*)$ be an attractor–repeller pair for \mathcal{F} , and let $v_{|\mathcal{P}|}$ be any distance potential for the pair $|\mathcal{P}| = (|\mathcal{A}|, |\mathcal{A}^*|)$. The sets $|\mathcal{A}|$ and $|\mathcal{A}^*|$ are not necessarily disjoint, but their interiors are disjoint by construction. Therefore, the fact that $v_{|\mathcal{P}|}$ may not be defined in $|\mathcal{A}| \cap |\mathcal{A}^*|$ when the intersection is nonempty does not affect the following definition.

Choose representatives $x_G \in \text{int}(|G|)$ for all $G \in \mathcal{G}$, and define the following function on the graph \mathcal{F} :

$$v_{\mathcal{P}}(G) := v_{|\mathcal{P}|}(x_G).$$

For any $k \geq 0$, let

$$v_{\mathcal{P}}^*(G, k) := \max_{H \in \Gamma_+^k(G)} v_{\mathcal{P}}(H).$$

This function is almost a Lyapunov function for the attractor–repeller pair decomposition as the following lemma shows.

Lemma 6.3. *Let $H \in \Gamma_+^n(G)$ for some $n \geq 1$, then*

$$v_{\mathcal{P}}^*(H, k) \leq v_{\mathcal{P}}^*(G, k) \leq v_{\mathcal{P}}^*(G, k + 1) \quad \text{for all } k \geq 0.$$

Proof. Clearly, $\Gamma_+^k(H) \subset \Gamma_+^k(G)$, and the first inequality follows from the definition of $v_{\mathcal{P}}^*$. Since $H \in \mathcal{F}^n(G)$ for some $n > 0$, $\Gamma_+^k(H) \subset \Gamma_+^{k+1}(G)$ from which the second inequality follows. \square

The above lemma shows that $v_{\mathcal{P}}^*$ is a weak Lyapunov function in the sense that it is not yet strictly decreasing off the attractor–repeller pair. For this reason we introduce the function

$$\mathcal{V}_{\mathcal{P}}(G) := \sum_{k=0}^{\infty} 2^{-k-1} v_{\mathcal{P}}^*(G, k). \tag{3}$$

It is immediately clear that the sum is convergent for all $G \in \mathcal{G}$ and $\mathcal{V}_{\mathcal{P}}(G) \in [0, 1]$ for all $G \in \mathcal{G}$.

Lemma 6.4. *For all $H \in \mathcal{F}(G)$ it holds that $\mathcal{V}_{\mathcal{P}}(G) \geq \mathcal{V}_{\mathcal{P}}(H)$ and $\mathcal{V}_{\mathcal{P}}(G) = \mathcal{V}_{\mathcal{P}}(H)$ if and only if $G, H \in \mathcal{A}$ or $G, H \in \mathcal{A}^*$. In addition, $\mathcal{V}_{\mathcal{P}}(\mathcal{A}) = 0$ and $\mathcal{V}_{\mathcal{P}}(\mathcal{A}^*) = 1$.*

Proof. The first statement of the lemma follows immediately from the definition of $\mathcal{V}_{\mathcal{P}}$ and Lemma 6.3. If both G and H lie in either \mathcal{A} or \mathcal{A}^* , then, by the definition of $\mathcal{V}_{\mathcal{P}}$ it follows that $\mathcal{V}_{\mathcal{P}}$ is equal to either 0 or 1, respectively.

Suppose that H and G are not both contained in either \mathcal{A} or \mathcal{A}^* . The definition of \mathcal{A} and \mathcal{A}^* then implies that $G \notin \mathcal{A}$ and $H \notin \mathcal{A}^*$. We distinguish three cases. The first case is when neither H nor G is contained in $\mathcal{A} \cup \mathcal{A}^*$. Suppose that $\mathcal{V}_{\mathcal{P}}(G) = \mathcal{V}_{\mathcal{P}}(H)$. Lemma 6.3 yields that $\mathbf{v}_{\mathcal{P}}^*(H, k) \leq \mathbf{v}_{\mathcal{P}}^*(G, k)$ for all $k \geq 0$, and thus $\mathbf{v}_{\mathcal{P}}^*(H, k) = \mathbf{v}_{\mathcal{P}}^*(G, k)$ for all $k \geq 0$. Since $G \notin \mathcal{A}^*$ it follows that $\mathbf{v}_{\mathcal{P}}^*(G, k) \rightarrow 0$ as $k \rightarrow \infty$. Indeed, $\mathbf{v}_{\mathcal{P}}^*(G, k) = 0$ for $k \geq k_*$ for some $k_* > 0$. From Lemma 6.3 we get

$$0 = \mathbf{v}_{\mathcal{P}}^*(G, k_*) \geq \mathbf{v}_{\mathcal{P}}^*(H, k_* - 1) = \mathbf{v}_{\mathcal{P}}^*(G, k_* - 1) \geq 0,$$

which shows that $\mathbf{v}_{\mathcal{P}}^*(G, k_* - 1) = 0$ and, therefore, $\mathbf{v}_{\mathcal{P}}^*(G, k) = 0$ for all $k \geq 0$. Consequently, $\mathcal{V}_{\mathcal{P}}(G) = \mathcal{V}_{\mathcal{P}}(H) = 0$. However, since $G \notin \mathcal{A}$, $\mathcal{V}_{\mathcal{P}}(G) > 0$, a contradiction.

In the second case assume that $G \in \mathcal{A}^*$. Suppose again that $\mathcal{V}_{\mathcal{P}}(G) = \mathcal{V}_{\mathcal{P}}(H)$. Then $1 = \mathbf{v}_{\mathcal{P}}^*(G, k) = \mathbf{v}_{\mathcal{P}}^*(H, k) \rightarrow 0$ as $k \rightarrow \infty$ since $H \notin \mathcal{A}^*$, a contradiction. Therefore, $1 = \mathcal{V}_{\mathcal{P}}(G) > \mathcal{V}_{\mathcal{P}}(H)$.

Finally, in the third case, assume that $H \in \mathcal{A}$. Again suppose that $\mathcal{V}_{\mathcal{P}}(G) = \mathcal{V}_{\mathcal{P}}(H)$. Then $0 = \mathbf{v}_{\mathcal{P}}^*(G, k) = \mathbf{v}_{\mathcal{P}}^*(H, k)$ for all $k \geq 0$. However, $\mathbf{v}_{\mathcal{P}}^*(G, 0) > 0$, a contradiction. Therefore, $0 = \mathcal{V}_{\mathcal{P}}(H) < \mathcal{V}_{\mathcal{P}}(G)$. This completes the proof that $\mathcal{V}_{\mathcal{P}}$ is a (strong) Lyapunov function. \square

6.1.3. Distance Potentials and Lyapunov Functions for Morse Decompositions.

It is also possible to follow the philosophy of distance potentials directly on finer Morse decompositions. This is not the approach we will take, so we only give a rough idea how distance potentials generalize and how this leads to discrete Lyapunov functions for Morse decompositions.

As an example, a distance potential on three sets $A, B, C \subset X$ can be constructed by

$$\mathbf{v}_{(A,B,C)}(x) := \alpha \frac{d(x, A \cup B)}{d(x, A \cup B) + d(x, C)} + (1 - \alpha) \frac{d(x, A)}{d(x, A) + d(x, B \cup C)}$$

with $\alpha \in (0, 1)$. Here we assume the ordering $C \rightarrow B \rightarrow A$. Of course, the properties of such a function can be put into a set of axioms as we did for distance potentials for two sets. A discrete Lyapunov function can be constructed by the previous procedure.

6.2. The Limits of Discrete Lyapunov Functions for a Single Attractor–Repeller Pair

We will now link the Lyapunov functions constructed above to attractor–repeller pairs for the map $f: X \rightarrow X$. The function $\mathcal{V}_{\mathcal{P}}$ induces a piecewise constant

function $V_{|\mathcal{P}|}$ on X by

$$V_{|\mathcal{P}|}(x) = \mathcal{V}_{\mathcal{P}}(G) \quad \text{for } x \in |G|.$$

By construction $V_{|\mathcal{P}|}(f(x)) \leq V_{|\mathcal{P}|}(x)$, which implies that $V_{|\mathcal{P}|}$ is a (discontinuous) weak Lyapunov function. Notice that $V_{|\mathcal{P}|}(x)$ is not well defined on overlapping boundaries of grid elements. For all of the constructions in this section, any reasonable choice of values on these overlaps will work, for example, we could choose $V_{|\mathcal{P}|}(x) = \max_{G \ni x} \mathcal{V}_{\mathcal{P}}(G)$, but for particular applications different choices may be more appropriate.

The next step is to consider any sequence of grids $\{\mathcal{G}_n\}$ with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$, and obtain proper Lyapunov functions for attractor–repeller pairs (A, A^*) for f . Associated with $\{\mathcal{G}_n\}$ is a sequence of multivalued maps $\{\mathcal{F}_{\mathcal{G}_n}\}$. For a given attractor–repeller pair $P = (A, A^*)$ for f , Proposition 5.5 implies that there exist attractor–repeller pairs $\mathcal{P}^n = (\mathcal{A}^n, \mathcal{A}^{*n})$ for the graphs $\mathcal{F}_{\mathcal{G}_n}$ whose realizations converge to $P = (A, A^*)$ as $n \rightarrow \infty$.

For the sequence of grids we have the distance potentials $\mathbf{v}_{|\mathcal{P}^n|}$ defined for the attractor–repeller pairs $(\mathcal{A}^n, \mathcal{A}^{*n})$. Since $(|\mathcal{A}^n|, |\mathcal{A}^{*n}|) \rightarrow (A, A^*)$ and, thus, $|\mathcal{A}^n| \cap |\mathcal{A}^{*n}| = \emptyset$ (using $A \cap A^* = \emptyset$), for n sufficiently large, we conclude $\mathbf{v}_{|\mathcal{P}^n|} \rightarrow \mathbf{v}_P$ in $C^0(X)$, where \mathbf{v}_P is defined as before with respect to $P = (A, A^*)$.

Before establishing the convergence of $V_{|\mathcal{P}^n|}$ let us first consider Lyapunov functions for attractor–repeller pairs for the map $f: X \rightarrow X$.

6.2.1. *Continuous Lyapunov Functions.* Define the functions

$$\mathbf{v}_P^*(x, k) := \max_{y \in \gamma_x^k} \mathbf{v}_P(y), \quad k \geq 0.$$

Lemma 6.5. *The functions $\mathbf{v}_P^*(\cdot, k)$ are continuous on X for all k . Moreover, for each k , $\mathbf{v}_P^*(\cdot, k)$ is nonincreasing along orbits γ_x and is nonincreasing in k , i.e., $\mathbf{v}_P^*(x, k) \geq \mathbf{v}_P^*(f(x), k)$ and, $\mathbf{v}_P^*(x, k) \geq \mathbf{v}_P^*(x, k + 1)$.*

Proof. Clearly, $\mathbf{v}_P^*(f(x), k) = \mathbf{v}_P^*(x, k + 1) \leq \mathbf{v}_P^*(x, k)$, since $\gamma_x^{k+1} \subset \gamma_x^k$.

To prove continuity we argue as follows. Assume, for the sake of simplicity, that $k = 0$ and write $\mathbf{v}_P^*(x) = \mathbf{v}_P^*(x, 0)$. Indeed, $\mathbf{v}_P^*(x, k) = \mathbf{v}_P^*(f^k(x), 0)$, so continuity only needs to be established at $k = 0$.

For $x \in A^*$ it holds that $\mathbf{v}_P(x) = 1$, and thus $\mathbf{v}_P^*(x) = 1$. If for any sequence $\{x_i\}$, with $\lim x_i = x \in A^*$, then clearly $\lim \mathbf{v}_P^*(x_i) = 1$, proving continuity at points $x \in A^*$. As for points $x \in A$ we argue as follows. For each $\varepsilon > 0$ there exists a neighborhood $U_\varepsilon \supset A$, with $f(U_\varepsilon) \subset U_\varepsilon$ and $h(U_\varepsilon, A) < \varepsilon$, and there exists an $N(\varepsilon) > 0$ such that $x_i \in U_\varepsilon$ for all $i \geq N(\varepsilon)$. By the Lipschitz continuity of \mathbf{v}_P this then immediately implies that $\mathbf{v}_P^*(x_i) < C\varepsilon$, proving that $\mathbf{v}_P^*(x_i) \rightarrow 0$ as $i \rightarrow \infty$.

For $x \in X \setminus (A \cup A^*)$ it holds that $\mathbf{v}_P(f^i(x)) \rightarrow 0$ as $i \rightarrow \infty$. Let $B_\delta(x)$ be a closed neighborhood of x , then there exists a $\delta > 0$, and an $i_0(x)$ such that

$\mathbf{v}_P(f^i(x'))$ attains its maximum for $i \in [0, i_0(x)]$ for all $x' \in B_\delta(x)$. This can be seen as follows. Set $d(x, A) = d_1 > 0$ and $d(x, A^*) = d_2 > 0$, then for all $x' \in B_\delta(x)$ it holds that

$$\mathbf{v}_P(x') \geq m > 0$$

provided that $\delta \leq \frac{1}{2} \min(d_1, d_2)$. Let $N \supset A$ be an isolating attracting set for A . By choosing $h(N, A)$ sufficiently small we obtain that

$$\mathbf{v}_P|_N < m.$$

Clearly, there exists an $i_0(x)$ such that $y = f^{i_0(x)}(x) \subset \text{int}(N)$. The map f^{i_0} is uniformly continuous, thus there exists an $\delta_\varepsilon > 0$ such that $f^{i_0}(x') \in B_\varepsilon(y)$ for all $x' \in B_{\delta_\varepsilon}(x)$. By choosing $\delta = \min(\delta_\varepsilon, d_1/2, d_2/2)$ we derive that $\mathbf{v}_P(B_\delta(x)) > \mathbf{v}_P(B_\varepsilon(y))$. Consequently, $\mathbf{v}_P(f^i(x'))$ attains its maximum for $i \in [0, i_0(x)]$ for all $x' \in B_\delta(x)$.

Define $\mathbf{z}(i, x) = \mathbf{v}_P(f^i(x))$, which is continuous in x , and for each $\varepsilon > 0$ there exists a $\delta_\varepsilon > 0$ such $d(x, x') < \delta_\varepsilon$ implies that $|\mathbf{z}(i, x) - \mathbf{z}(i, x')| < \varepsilon$ uniformly in $i \in [0, i_0(N)]$. Thus,

$$-\varepsilon + \mathbf{z}(i, x) < \mathbf{z}(i, x') < \mathbf{z}(i, x) + \varepsilon.$$

We have $\mathbf{v}_P^*(x) = \mathbf{z}(i_1, x)$ and $\mathbf{v}_P^*(x') = \mathbf{z}(i_2, x)$, $i_1, i_2 \in [0, i_0(N)]$. Using the above inequalities we derive that

$$\mathbf{v}_P^*(x) = \mathbf{z}(i_1, x) < \mathbf{z}(i_1, x') + \varepsilon \leq \mathbf{z}(i_2, x') + \varepsilon = \mathbf{v}_P^*(x') + \varepsilon.$$

Similarly $\mathbf{v}_P^*(x') < \mathbf{v}_P^*(x) + \varepsilon$ so that $|\mathbf{v}_P^*(x) - \mathbf{v}_P^*(x')| < \varepsilon$. \square

For the map f the following function

$$V_P(x) = \sum_{k=0}^{\infty} 2^{-k-1} \mathbf{v}_P^*(x, k) \quad (4)$$

is a Lyapunov function for the attractor–repeller pair $P = (A, A^*)$, as was proved by several authors in the case of homeomorphisms, see [2], [10], [15].

6.2.2. Limits of Discrete Lyapunov Functions. Let $x \in G_n$ and introduce the functions

$$\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) = \max_{y \in |\Gamma_+^k(G_n)|} \mathbf{v}_{|\mathcal{P}^n|}(y), \quad k \geq 0.$$

The next lemma links the piecewise constant functions $\mathbf{v}_{|\mathcal{P}^n|}^*(x, k)$ to continuous functions $\mathbf{v}_P^*(x, k)$.

Lemma 6.6. *For each $k \geq 0$ and $\varepsilon > 0$ there exists an $N(\varepsilon, k)$ such that*

$$|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_P^*(x, k)| < \varepsilon \quad \text{for } n \geq N(\varepsilon, k),$$

uniformly in $x \in X$.

Proof. Choose compact neighborhoods $U_{\varepsilon,k} \supset A$ and $U_{\varepsilon,k}^* \supset A^*$, with $h(A, U_{\varepsilon,k}) < \varepsilon/(6 \text{Lip}_v)$ and $h(A^*, U_{\varepsilon,k}^*) < \varepsilon/(6 \text{Lip}_v)$, such that, in addition,

$$0 \leq \mathbf{v}_P(f^i(x)) \leq \varepsilon/6 \quad \text{for all } x \in U_{\varepsilon,k}$$

and

$$1 - \varepsilon/6 \leq \mathbf{v}_P(f^i(x)) \leq 1 \quad \text{for all } x \in U_{\varepsilon,k}^*$$

for all $0 \leq i \leq k$. Define $D_{\varepsilon,k} = \text{cl}(X \setminus (U_{\varepsilon,k} \cup U_{\varepsilon,k}^*))$. From the proof of Lemma 6.5 we obtain a constant $k_* \geq k$ such that $\mathbf{v}_P(f^i(x))$ attains its maximum in i on the interval $[k, k_*]$ for all $x \in D_{\varepsilon,k}$. Consequently, for $x \in D_{\varepsilon,k}$, we have that

$$\mathbf{v}_P^*(x, k) = \max_{y \in \gamma_x^{k, k_*}} \mathbf{v}_P(y).$$

Moreover, choose $G_n \ni x$ and define

$$W(x, k) := \max_{y \in |\Gamma^{k, k_*}(G_n)|} \mathbf{v}_P(y).$$

By Proposition 5.4 we have that $|\Gamma^{k, k_*}(G_n)| \rightarrow \gamma_x^{k, k_*}$ in the Hausdorff metric, as $n \rightarrow \infty$. Since \mathbf{v}_P is Lipschitz continuous it follows that $|\mathbf{v}_P(x) - \mathbf{v}_P(x')| \leq Cd(x, x')$. By choosing $\delta_\varepsilon = \varepsilon/(3 \text{Lip}_v)$ it then holds that $|\mathbf{v}_P(x) - \mathbf{v}_P(x')| < \varepsilon/3$ for all x, x' , for which $d(x, x') < \delta_\varepsilon$. Consequently, there exists an $N'(\varepsilon, k)$ such that $h(|\Gamma^{k, k_*}(G_n)|, \gamma_x^{k, k_*}) < \delta_\varepsilon$ for all $n \geq N'(\varepsilon, k)$. As a direct consequence, it holds that $|W(x, k) - \mathbf{v}_P^*(x, k)| < \varepsilon/3$ for each $x \in D_{\varepsilon,k}$. By construction, the same estimate holds for $x \in U_{\varepsilon,k} \cup U_{\varepsilon,k}^*$ and, thus, $|W(x, k) - \mathbf{v}_P^*(x, k)| < \varepsilon/3$ for $n \geq N'(\varepsilon, k)$, and for all $x \in X$.

From Proposition 5.5 it follows that $|\mathcal{A}^n| \cap |\mathcal{A}^{*n}| = \emptyset$, provided that n is sufficiently large. Therefore, for n sufficiently large the distance potential $\mathbf{v}_{|\mathcal{P}^n|}$ is a Lipschitz continuous function on X . Let $G_n \in \mathcal{A}^{*n}$, then by definition $\Gamma_+^k(G_n) \cap \mathcal{A}^{*n} \neq \emptyset$. As a consequence $\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) = 1$ for $x \in |\mathcal{A}^{*n}|$ and, similarly, $\Gamma_+^k(G_n) \cap \mathcal{A}^n \subset \mathcal{A}^n$ implies $\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) = 0$ for $x \in |\mathcal{A}^n|$.

Consider the complement $\mathcal{Y}^n = \mathcal{G}_n \setminus (\mathcal{A}^n \cup \mathcal{A}^{*n})$ and set $Y^n = |\mathcal{Y}^n|$. From continuity it now follows that, for each n , there exists a k_n such that $f^{k_n}(x) \in |\mathcal{A}^n|$ for all $x \in Y^n$. This implies that (for $k \leq k_n$)

$$\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) = \max_{y \in |\Gamma^{k, k_n}(G_n)|} \mathbf{v}_{|\mathcal{P}^n|}(y)$$

for all $x \in X$. For $k > k_n$ it holds that $\mathbf{v}_{|\mathcal{P}^n|}^*(x, k)$ is equal to either 0 or 1. Let $x \in G_n$ with $G_n \in \mathcal{Y}^n$ (and thus $x \in Y^n$), and define the function

$$\widetilde{W}(x, k) := \max_{y \in |\Gamma^{k, k_n}(G_n)|} \mathbf{v}_P(y).$$

Since $\mathbf{v}_{|\mathcal{P}^n|} \rightarrow \mathbf{v}_P$ in $C^0(X)$ as $n \rightarrow \infty$, it follows that for each $\varepsilon > 0$ there exists an $N''(\varepsilon, k)$ such that $|\widetilde{W}(x, k) - \mathbf{v}_{|\mathcal{P}^n|}^*(x, k)| < \varepsilon/3$, uniformly in $x \in X$, provided that $n \geq N''(\varepsilon, k) > 0$.

If n is chosen large enough, then $D_{\varepsilon,k} \subset Y^n$. This we can include in the choice of $N''(\varepsilon, k)$. Consequently, $h(D_{\varepsilon,k}, Y^n) < \varepsilon/(6 \text{Lip}_v)$, which yields that

$$|\tilde{W}(x, k) - W(x, k)| < \varepsilon/3$$

for $n \geq N''(\varepsilon, k) > 0$, uniformly in $x \in X$. Indeed, since $h(A, U_{\varepsilon,k}) < \varepsilon/(6 \text{Lip}_v)$ and $h(A^*, U_{\varepsilon,k}^*) < \varepsilon/(6 \text{Lip}_v)$, we can compare the function $\mathbf{v}_{|\mathcal{P}^n|}$ outside A and A^* . Using the Lipschitz continuity it then follows that $|\mathbf{v}_{|\mathcal{P}^n|}(x) - \mathbf{v}_{|\mathcal{P}^n|}(y)| \leq 2 \cdot \text{Lip}_v d(x, y) < \varepsilon/3$ for any pair $x, y \in U_{\varepsilon,k} \setminus A$. The same holds for $x, y \in U_{\varepsilon,k}^* \setminus A^*$. The estimate on W and \tilde{W} now follows. From the inequality

$$\begin{aligned} |\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_P^*(x, k)| &\leq |\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \tilde{W}(x, k)| \\ &\quad + |\tilde{W}(x, k) - W(x, k)| + |W(x, k) - \mathbf{v}_P^*(x, k)|. \end{aligned}$$

We deduce that $|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_P^*(x, k)| < \varepsilon$ for all $n \geq \max(N'(\varepsilon, k), N''(\varepsilon, k)) = N(\varepsilon, k)$ uniformly in $x \in X$, which proves the lemma. \square

By Lemma 6.6 it follows that for each $\varepsilon > 0$ and $k \geq 0$ there exists an $N(\varepsilon, k)$ such that $|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_P^*(x, k)| < \varepsilon$ for $n \geq N(\varepsilon, k)$, uniformly in $x \in X$. Choose $M(\varepsilon)$ such that $\sum_{k=M+1}^{\infty} 2^{-k-1} < \varepsilon$ for $M \geq M(\varepsilon)$. Define $N^*(\varepsilon) = \max_{k \leq M(\varepsilon)} N(\varepsilon, k)$, then $|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_P^*(x, k)| < \varepsilon$ for all $n \geq N^*(\varepsilon)$, and for all $k \leq M(\varepsilon)$. Consequently,

$$\left| \sum_{k=0}^{M(\varepsilon)} 2^{-k-1} \mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \sum_{k=0}^{M(\varepsilon)} 2^{-k-1} \mathbf{v}_P^*(x, k) \right| \leq \sum_{k=0}^{M(\varepsilon)} 2^{-k-1} \varepsilon < (1 - \varepsilon) \varepsilon < \varepsilon$$

for all $n \geq N^*(\varepsilon)$, uniformly in $x \in X$. The same bound can be achieved for k ranging from $M(\varepsilon) + 1$ to ∞ due to the choice of $M(\varepsilon)$. Therefore, $\sum_{k=0}^{\infty} 2^{-k-1} \mathbf{v}_{|\mathcal{P}^n|}^*(x, k) \rightarrow \sum_{k=0}^{\infty} 2^{-k-1} \mathbf{v}_P^*(x, k) = V_P(x)$ uniformly in $x \in X$ as $n \rightarrow \infty$. Finally, we need to link $\mathbf{v}_{|\mathcal{P}^n|}^*(x, k)$ to $\mathbf{v}_{\mathcal{P}^n}^*(G_n, k)$. A straightforward comparison shows that

$$|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_{\mathcal{P}^n}^*(G_n, k)| \leq \varepsilon \text{Lip}_v.$$

Consequently,

$$\left| \sum_{k=0}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \sum_{k=0}^K 2^{-k-1} \mathbf{v}_{\mathcal{P}^n}^*(G_n, k) \right| \leq \sum_{k=0}^K 2^{-k-1} \varepsilon \text{Lip}_v < \varepsilon \text{Lip}_v$$

for all $K > 0$ and, thus, $|V_{|\mathcal{P}^n|}(x) - V_P(x)| < \varepsilon(2 + \text{Lip}_v)$, which concludes our construction. The result can be summarized as follows:

Theorem 6.7. *Given an attractor–repeller pair $P = (A, A^*)$ and a sequence of grids $\{\mathcal{G}_n\}$ with $\text{diam}(\mathcal{G}_n) \rightarrow 0$. Then there exists a sequence of attractor–repeller pairs $\mathcal{P}^n = (A^n, A^{*n})$, with $h(|\mathcal{P}^n|, P) \rightarrow 0$ as $n \rightarrow \infty$, and associated piecewise constant Lyapunov functions $V_{|\mathcal{P}^n|}$, coming from the (discrete) Lyapunov functions $\mathcal{V}_{\mathcal{P}^n}$, such that $V_{|\mathcal{P}^n|}$ converges uniformly on $x \in X$ to the (continuous) Lyapunov function V_P , as $n \rightarrow \infty$ defined by (4).*

Remark 6.8. In equations (3) and (4) the discrete and continuous Lyapunov functions are defined via a sequence of weights that guarantee convergence of the series. There is of course freedom in choosing these weights as long as they satisfy the appropriate convergence criteria. In some instances it may be preferable to choose different weight factors.

6.3. *Lyapunov Functions for Morse Decompositions and the Chain-Recurrent Set*

In the previous subsection it is shown that Lyapunov functions for attractor–repeller pairs for f can be obtained as limits of piecewise constant Lyapunov functions on a sequence of grids. The final step in this section is to establish this result also for Morse decompositions and the chain-recurrent set $\mathcal{R}(X, f)$.

From Theorem 5.6 we have that the chain recurrent set of f is given by $\mathcal{R}(X, f) = \bigcap_{i \in \mathbb{N}} (A_i, A_i^*)$, where $\{(A_i, A_i^*)\}_{i \in \mathbb{N}}$ is the collection of all attractor–repeller pairs for (X, f) . As before, we denote the attractor–repeller pairs by $P_i = (A_i, A_i^*)$. The function

$$V(x) = 2 \sum_{i=1}^{\infty} 3^{-i} V_{P_i}(x), \quad (5)$$

where V_{P_i} are the Lyapunov functions defined in (4) for the attractor–repeller pairs $P_i = (A_i, A_i^*)$, is a proper Lyapunov function for (X, f) with values in $[0, 1]$. The function V defined in (5) now establishes Conley’s Decomposition Theorem as explained at the end of this section. However, such a Lyapunov function can be defined in many ways. First, by choosing the ordering of the attractor repeller pairs and, second, by a variation in the weight factors. This leads to uncountably many different Lyapunov functions for the chain-recurrent set. In this section we also show that such choices can be made in a way that allows approximation by piecewise constant Lyapunov functions on a sequence of grids.

The ordering of the attractor–repeller pairs can also be arranged simply by re-ordering the weight factors. Let $\mathbf{a} = \{\alpha_i\}_{i=1}^{\infty}$ be a sequence of weight factors with $\sum_{i=1}^{\infty} \alpha_i = 1$, then define

$$V_{\mathbf{a}}(x) = \sum_{i=1}^{\infty} \alpha_i V_{P_i}(x). \quad (6)$$

It is immediately clear from this definition that $V_{\mathbf{a}}$ is a function that maps onto $[0, 1]$. If all weight factors are positive, then $V_{\mathbf{a}}$ is a strong Lyapunov function due to the fact that V_{P_i} are Lyapunov functions.

6.3.1. *Lyapunov Functions for Morse Decompositions.* A Morse decomposition M can be found as a finite intersection of attractor–repeller pairs, i.e., $M =$

$\bigcap_{k=1}^{\ell} (A_{i_k} \cup A_{i_k}^*)$. By choosing weights β_k , with $\sum_k \beta_k = 1$, the function

$$V_M(x) = \sum_{k=1}^{\ell} \beta_k V_{P_{i_k}}(x) \quad (7)$$

becomes a Lyapunov function for M . We should point out that V_M is a special case of a Lyapunov function of the type defined in equation (6) by setting

$$\alpha_i = \begin{cases} \beta_k, & i = i_k, \\ 0, & i \neq i_k. \end{cases}$$

From Theorem 6.7 we obtain attractor–repeller pairs $|\mathcal{P}_{i_k}^n| \rightarrow P_{i_k}$ and piecewise constant Lyapunov functions $V_{|\mathcal{P}_{i_k}^n|} \rightarrow V_{P_{i_k}}$. To be more precise, for each $\varepsilon > 0$, there exists an $N(\varepsilon, i_k)$ such that $|V_{|\mathcal{P}_{i_k}^n|}(x) - V_{P_{i_k}}(x)| < \varepsilon$. For the graph $\mathcal{F}_{\mathcal{G}_n}$ we define the Lyapunov functions

$$\mathcal{V}_M^n(G_n) = \sum_{k=1}^{\ell} \beta_k \mathcal{V}_{\mathcal{P}_{i_k}^n}^n(G_n),$$

and the associated piecewise constant functions $V_M^n(x) = \mathcal{V}_M^n(G_n)$ for all $x \in G_n$. By setting $L(\varepsilon) = \max_{1 \leq k \leq \ell} N(\varepsilon, i_k)$ we obtain that $|V_M^n(x) - V_M(x)| < \varepsilon$ for all $n \geq L(\varepsilon)$, which yields a convergence result for Lyapunov functions for Morse decompositions.

Theorem 6.9. *Given a Morse decomposition M for (X, f) and a sequence of grids \mathcal{G}_n , with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$. Then, for any Lyapunov function V_M for M defined in (7), there exists piecewise constant functions V_M^n , coming from Lyapunov functions \mathcal{V}_M^n for the associated Morse decompositions $\mathcal{M} = \bigcap_{k=1}^{\ell} (A_{i_k}^n \cup A_{i_k}^{*n})$ for $\mathcal{F}_{\mathcal{G}_n}$, such that $V_M^n \rightarrow V_M$, uniformly in $x \in X$, as $n \rightarrow \infty$.*

6.3.2. Convergence to Lyapunov Functions on the Chain-Recurrent Set. Our next step is to construct discrete Lyapunov functions that converge to a Lyapunov function for the chain-recurrent $\mathcal{R}(X, f)$. This can be done by constructing sequences of Lyapunov functions for the Morse decompositions $M_k = \bigcap_{i=1}^k (A_i \cup A_i^*)$. We start by approximating V_a for a given weight sequence $\mathbf{a} = \{\alpha_i\}$. For the Morse decomposition M_k we define $\beta_i^k = \alpha_i / \sigma_k$, where $\sigma_k = \sum_{i=1}^k \alpha_i < 1$. Set $V_{M_k}(x) = \sum_{i=1}^k \beta_i^k V_{P_i}(x)$, then

$$\begin{aligned} |V_{M_k}(x) - V_a(x)| &= \left| \sum_{i=1}^k (\beta_i^k - \alpha_i) V_{P_i}(x) - \sum_{i=k+1}^{\infty} \alpha_i V_{P_i}(x) \right| \\ &= \left| \sum_{i=1}^k \alpha_i \left(\frac{1}{\sigma_k} - 1 \right) V_{P_i}(x) - \sum_{i=k+1}^{\infty} \alpha_i V_{P_i}(x) \right| \end{aligned}$$

$$\begin{aligned} &\leq \left(\frac{1}{\sigma_k} - 1\right) \sum_{i=1}^k \alpha_i + \sum_{i=k+1}^{\infty} \alpha_i \\ &\leq (1 - \sigma_k) + \sum_{i=k+1}^{\infty} \alpha_i. \end{aligned}$$

Consequently, for each $\varepsilon > 0$, there exists a $K(\varepsilon) > 0$ such that $|V_{M_k}(x) - V_a(x)| < \varepsilon$ for all $k \geq K(\varepsilon)$, uniformly in $x \in X$. From Subsection 6.3.1, $|V_{M_k}^n(x) - V_{M_k}(x)| < \varepsilon$ for $n \geq L(\varepsilon, M_k)$ for Morse decompositions. Now for each $\varepsilon > 0$ there exists an $N'(\varepsilon) = L(\varepsilon/3, M_{K(\varepsilon/3)})$ such that, for all $k \geq K(\varepsilon/3)$ and $n \geq N'(\varepsilon)$, it holds that

$$\begin{aligned} |V_{M_n}^n(x) - V_a(x)| &\leq |V_{M_n}^n(x) - V_{M_k}^n(x)| \\ &\quad + |V_{M_k}^n(x) - V_{M_k}(x)| + |V_{M_k}(x) - V_a(x)| \\ &\leq |V_{M_n}^n(x) - V_{M_k}^n(x)| + 2\varepsilon/3. \end{aligned}$$

The remaining term can be estimated as follows. Recall that

$$V_{M_k}^n(x) = \sigma_k^{-1} \sum_{i=1}^k \alpha_i V_{|\mathcal{P}_i^n|}(x),$$

where the attractor–repeller pairs are chosen such that $\mathcal{P}_i^n \rightarrow \mathcal{P}$ for all i , as $n \rightarrow \infty$. Fix $k, n > N$, then

$$\begin{aligned} |V_{M_n}^n(x) - V_{M_k}^n(x)| &\leq \left| \frac{1}{\sigma_n} - \frac{1}{\sigma_k} \right| \sum_{i=1}^N \alpha_i V_{|\mathcal{P}_i^n|}(x) \\ &\quad + \frac{1}{\sigma_k} \sum_{i=N+1}^k \alpha_i + \frac{1}{\sigma_n} \sum_{i=N+1}^n \alpha_i. \end{aligned}$$

Consequently, by choosing $n \geq N'(\varepsilon)$ large enough, it holds that the latter can be estimated by $\varepsilon/3$. By setting $N(\varepsilon) = \max(N'(\varepsilon), N''(\varepsilon))$ we obtain that $|V_{M_n}^n(x) - V_a(x)| < \varepsilon$ for all $n \geq N(\varepsilon)$, uniformly in $x \in X$.

We have constructed discrete Lyapunov functions $\mathcal{V}_{M_n}^n$ for Morse decompositions M_n for which the piecewise constant realizations $V_{M_n}^n$ converge to V_a . The final step is to find a sequence of Lyapunov functions for the chain-recurrent sets $\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})$. Define the function

$$V_a^n(x) = \frac{1}{\rho_n} \sum_{i=1}^{p_n} \alpha_i V_{|\mathcal{P}_i^n|}(x) = \frac{\sigma_n}{\rho_n} V_{M_n}^n(x) + \frac{1}{\rho_n} \sum_{i>n} \alpha_i V_{|\mathcal{P}_i^n|}(x),$$

where $\rho_n = \sum_{i=1}^{p_n} \alpha_i$ and p_n is the total number of attractor–repeller pairs for each n . By definition $\sigma_n \leq \sigma_n/\rho_n \leq 1$, and $\sigma_n \rightarrow 1$ as $n \rightarrow \infty$. Therefore,

$$|V_a^n(x) - V_{M_n}^n(x)| \leq \left(1 - \frac{\sigma_n}{\rho_n}\right) \sum_{j=1}^n \alpha_j + \frac{1}{\rho_n} (1 - \sigma_n) < \varepsilon,$$

provided n is large enough. These arguments prove the following theorem.

Theorem 6.10. *Let $\{\mathcal{G}_n\}$ be any sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$, and let the functions $\{V_a^n\}$ and V_a be as introduced above. Then*

$$V_a^n \rightarrow V_a \quad \text{uniformly in } x \in X$$

as $n \rightarrow \infty$.

In the above considerations we choose \mathbf{a} as a fixed “index.” We can also allow sequences \mathbf{a}_n with the property that $\mathbf{a}_n \rightarrow \mathbf{a}$ uniformly as $n \rightarrow \infty$. The statement of Theorem 6.10 then becomes $V_{\mathbf{a}_n}^n \rightarrow V_{\mathbf{a}}$ uniformly in $x \in X$ as $n \rightarrow \infty$.

6.3.3. Conley’s Fundamental Decomposition Theorem. The above constructions provide discrete Lyapunov functions for $\mathcal{R}(\mathcal{G}_n)$ which converge to a Lyapunov function for $\mathcal{R}(X, f)$. The properties of these Lyapunov functions are summarized in Theorem 6.12 below, which establishes Conley’s Fundamental Decomposition Theorem.

The following definition is equivalent to the one given in the Introduction.

Definition 6.11. The **chain components** of $\mathcal{R}(X, f)$ are the equivalence classes of the relation $x \sim y$ if for every $\varepsilon > 0$ there exists a grid \mathcal{G} with $\text{diam}(\mathcal{G}) < \varepsilon$ containing elements G, H such that $x \in G, y \in H$, and $G \sim H$ in $\mathcal{R}(\mathcal{G})$.

Theorem 6.12 (Conley’s Fundamental Decomposition Theorem). *Let $f: X \rightarrow X$ be a continuous map on a compact metric space. Then there are at most countably many chain components $\{\mathcal{R}^i\}_{i \in J}$ of $\mathcal{R}(X, f)$ and there exists a continuous function $V: X \rightarrow [0, 1]$ such that:*

- (1) if $x \notin \mathcal{R}(X, f)$, then $V(x) > V(f(x))$; and
- (2) for each $i \in J$ there exists $\sigma_i \in [0, 1]$ such that $\mathcal{R}^i \subset V^{-1}(\sigma_i)$.

Moreover, if \mathcal{G}_n is a sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$, then there exist Lyapunov functions V^n for the multivalued map on \mathcal{G}_n which are constant on grid elements such that $V^n \rightarrow V$ as $n \rightarrow \infty$ uniformly on X .

6.4. Explicit Constructions for the Chain-Recurrent Set

In the previous subsection we established the convergence of Lyapunov functions for Morse decompositions and for the chain-recurrent set. In particular, Theorem 6.10 states that any given Lyapunov function V_a can be approximated by piecewise constant approximations V_a^n which come from discrete Lyapunov functions \mathcal{V}_a^n on the grid \mathcal{G}_n .

In this subsection we address this convergence from a constructive point of view, i.e., given the Lyapunov functions \mathcal{V}_a^n for $\mathcal{R}(\mathcal{F}_{\mathcal{G}_n})$, under what conditions do they converge to a Lyapunov function for $\mathcal{R}(X, f)$? We outline a constructive way to

answer this question and thus provide an algorithm for constructing approximations for Lyapunov functions for $\mathcal{R}(X, f)$.

Let \mathcal{G}_n be a sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$. Proposition 5.5 establishes the following one-to-one correspondence between certain attractor–repeller pairs for f and attractor–repeller pairs for \mathcal{F} . Let $\mathbf{P}_c = \{P = (A, A^*) \mid \text{dist}(A, A^*) \geq c\}$ and for \mathcal{F} define $\mathbf{\Pi}_c = \{\mathcal{P} = (\mathcal{A}, \mathcal{A}^*) \mid \omega(|\mathcal{A}|) = A \text{ and } (A, A^*) \in \mathbf{P}_c\}$. Since \mathbf{P}_c is a finite set, Proposition 5.5 implies that there exists $N \geq 0$ such that, for any $n \geq N$, there is a one-to-one correspondence between \mathbf{P}_c and $\mathbf{\Pi}_c^n$ for \mathcal{F}_n . Moreover, given any $0 < \rho < c/2$, the number N can be chosen such that $h(|\mathcal{P}|, P) < \rho$ for corresponding pairs. Applying Theorem 6.7 to the elements of \mathbf{P}_c , we obtain the convergence of corresponding Lyapunov functions.

Thus, we will impose an ordering on attractor–repeller pairs by the distance between the attractor and repeller. In particular, the ordering $\{P_i\}_{i=1}^p$ (where p could be ∞) of all attractor–repeller pairs is *admissible* if $\text{dist}(A_i, A_i^*) \geq \text{dist}(A_{i+1}, A_{i+1}^*)$ for all $i > 0$ with a similar definition for attractor–repeller pairs for \mathcal{F} . The above ideas lead to the following theorem.

Theorem 6.13. *Let \mathcal{G}_n be a sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$, and let $\{P_i^n\}_{i=1}^{p_n}$ be admissible labelings of the attractor–repeller pairs for $\mathcal{F}_{\mathcal{G}_n}$, for each n . Given a (positive) sequence $\mathbf{a} = \{\alpha_i\}$ with $\sum_i \alpha_i = 1$, there exist piecewise constant Lyapunov functions $V_{\mathbf{a}}^n$ coming from discrete Lyapunov functions for $\mathcal{F}_{\mathcal{G}_n}$ such that for every $\varepsilon > 0$ there exists $N(\varepsilon)$ and (continuous) Lyapunov functions $V_{\mathbf{a}_n}$ for f with*

$$|V_{\mathbf{a}}^n(x) - V_{\mathbf{a}_n}(x)| < \varepsilon$$

uniformly for all $x \in X$ and all $n \geq N(\varepsilon)$. Moreover, the piecewise constant functions $V_{\mathbf{a}}^n$ are weak Lyapunov functions for f , i.e., $V_{\mathbf{a}}^n(f(x)) \leq V_{\mathbf{a}}^n(x)$ for all $x \in X$ and all $n > 0$.

Proof. We will consider three separate cases. First we will assume that there are infinitely many distinct attractor–repeller pairs and $\text{dist}(A_i, A_i^*) > \text{dist}(A_{i+1}, A_{i+1}^*)$ for all $i > 0$, in which case $\mathbf{a}_n = \mathbf{a}$ is constant. The case where some pairs have the same distance between attractor and repeller requires a possible rearrangement of \mathbf{a} for each n and is a technical modification of the argument below. Finally, the third case of finitely many attractor–repeller pairs is addressed.

Case 1. Assume there are infinitely many attractor–repeller pairs and the distances between attractors and repellers are all distinct.

Let $\varepsilon > 0$. Choose $L(\varepsilon)$ such that $\sum_{j=L(\varepsilon)}^{\infty} \alpha_j < \varepsilon/4$. Let $c = \text{dist}(A_{L(\varepsilon)}, A_{L(\varepsilon)}^*)$. Then $\mathbf{P}_c = \{P_1, \dots, P_{L(\varepsilon)}\}$. Define $r_i = (\text{dist}(A_i, A_i^*) - \text{dist}(A_{i+1}, A_{i+1}^*))/2$ for $i = 1, \dots, L(\varepsilon)$ and $r = \min_{1 \leq i \leq L(\varepsilon)} r_i$. Choose $\rho < \min\{\varepsilon, c/2, r/2\}$.

Now let $\{\mathcal{P}_i\}_{i=1}^{p_n}$ (where $p_n < \infty$) be an ordering of the set of attractor–repeller pairs of \mathcal{F}_n satisfying the condition $\text{dist}(|\mathcal{A}_i|, |\mathcal{A}_i^*|) \geq \text{dist}(|\mathcal{A}_{i+1}|, |\mathcal{A}_{i+1}^*|)$ for all $i > 0$. Choose $N > 0$ such that for all $n \geq N$ each attractor–repeller pair $P_i \in \mathbf{P}_c$ has a unique corresponding attractor–repeller pair \mathcal{Q}_i^n for which $h(|\mathcal{Q}_i^n|, P_i) < \rho$ by Proposition 5.5. Observe that since $A \subset |\mathcal{A}_i^n| \subset B_\rho(A)$ and $A^* \subset |\mathcal{A}_i^{*n}| \subset B_\rho(A^*)$, we have $\text{dist}(A_i, A_i^*) - 2\rho < \text{dist}(|\mathcal{A}_i^n|, |\mathcal{A}_i^{*n}|) < \text{dist}(A_i, A_i^*)$ so that

$$\begin{aligned} \text{dist}(|\mathcal{A}_i^n|, |\mathcal{A}_i^{*n}|) &> \text{dist}(A_i, A_i^*) - 2r_i \\ &> \text{dist}(A_{i+1}, A_{i+1}^*) > \text{dist}(|\mathcal{A}_{i+1}^n|, |\mathcal{A}_{i+1}^{*n}|) \end{aligned}$$

by the choices of r_i and ρ . Therefore $\{\mathcal{Q}_i^n\}$ respects the ordering by distance between the attractor and repeller so that $\mathcal{Q}_i^n = \mathcal{P}_i^n$ for all $i = 1, \dots, L(\varepsilon)$, which also implies $p_n \geq L(\varepsilon)$.

Furthermore, if \mathcal{Q}^n is any attractor–repeller pair for which $\text{dist}(|\mathcal{A}^n|, |\mathcal{A}^{*n}|) > c - 2\rho$, then the corresponding attractor–repeller pair for f from part (i) of Proposition 5.5 satisfies $\text{dist}(A, A^*) \geq c$ and hence $\mathcal{Q}^n = \mathcal{P}_i^n$ for some $1 \leq i \leq L(\varepsilon)$. Therefore $\text{dist}(|\mathcal{A}_i^n|, |\mathcal{A}_i^{*n}|) < c$ for all attractor–repeller pairs \mathcal{P}_i^n with $i > L(\varepsilon)$.

Using the same notation as in the previous subsection, define

$$V_a(x) = \sum_{i=1}^{\infty} \alpha_i V_{P_i}(x), \quad V_a^{L(\varepsilon)}(x) = \sum_{i=1}^{L(\varepsilon)} \alpha_i V_{P_i}(x),$$

and

$$V_a^n(x) = \frac{1}{\rho_n} \sum_{i=1}^{p_n} \alpha_i V_{|\mathcal{P}_i^n|}(x), \quad V_a^{n,L(\varepsilon)}(x) = \frac{1}{\rho_n} \sum_{i=1}^{L(\varepsilon)} \alpha_i V_{|\mathcal{P}_i^n|}(x).$$

We estimate $|V_a(x) - V_a^n(x)|$ by

$$\begin{aligned} |V_a(x) - V_a^n(x)| &\leq |V_a(x) - V_a^{L(\varepsilon)}(x)| + |V_a^{L(\varepsilon)}(x) - V_a^{n,L(\varepsilon)}(x)| \\ &\quad + |V_a^{n,L(\varepsilon)}(x) - V_a^n(x)| \\ &< \varepsilon/4 + \varepsilon/4 + 2\varepsilon/4 < \varepsilon. \end{aligned}$$

The first and third terms are estimated by

$$|V_a(x) - V_a^{L(\varepsilon)}(x)| \leq \sum_{i=L(\varepsilon)+1}^{\infty} \alpha_i V_{P_i}(x) < \varepsilon/4$$

and

$$|V_a^{n,L(\varepsilon)}(x) - V_a^n(x)| \leq \frac{1}{\rho_n} \sum_{i=L(\varepsilon)+1}^{p_n} \alpha_i V_{|\mathcal{P}_i^n|} \leq 2 \sum_{i=L(\varepsilon)+1}^{\infty} \alpha_i < 2\varepsilon/4$$

since $\rho_n > 1 - \varepsilon/4 > \frac{1}{2}$ for $\varepsilon < 2$. To estimate the middle term we apply Theorem 6.7 to choose N large enough so that $|V_{P_i}(x) - V_{|\mathcal{P}_i^n|}(x)| \leq \varepsilon/8$ uniformly in x for

all $1 \leq i \leq L(\varepsilon)$ and $1 - 1/\rho_n < \varepsilon/8$. Then,

$$\begin{aligned} |V_a^{L(\varepsilon)}(x) - V_a^{n,L(\varepsilon)}(x)| &\leq \sum_{i=1}^{L(\varepsilon)} \alpha_i \left| V_{P_i} - \frac{1}{\rho_n} V_{|\mathcal{P}_i^n|} \right| \\ &\leq \sum_{i=1}^{L(\varepsilon)} \alpha_i |V_{P_i} - V_{|\mathcal{P}_i^n|}| + \left(1 - \frac{1}{\rho_n}\right) \sum_{i=1}^{L(\varepsilon)} \alpha_i < \varepsilon/4. \end{aligned}$$

This completes Case 1.

Case 2. Assume that there are infinitely many attractor–repeller pairs and $\text{dist}(A_i, A_i^*) = \text{dist}(A_{i+1}, A_{i+1}^*)$ for some values of $1 \leq i < \infty$.

Note that for any given value of $\text{dist}(A_i, A_i^*)$ there are only finitely many attractor–repeller pairs which attain this distance. The argument proceeds as in Case 1 with the following modifications.

Choose $L(\varepsilon)$ as in Case 1 with the additional constraint that $\text{dist}(A_{L(\varepsilon)}, A_{L(\varepsilon)}^*) > \text{dist}(A_{L(\varepsilon)+1}, A_{L(\varepsilon)+1}^*)$. The distances associated to the attractor–repeller pairs Q_i^n respect the gaps in distances associated to P_i but it need not happen that $Q_i^n = \mathcal{P}_i^n$ for all $1 \leq i \leq L(\varepsilon)$. However, as sets $\Pi_c^n = \{Q_i^n\}$ so that any attractor–repeller pair Q^n for which $\text{dist}(|\mathcal{A}^n|, |\mathcal{A}^{*n}|) > c - 2\rho$ satisfies $Q^n = \mathcal{P}_i^n$ for some $1 \leq i \leq L(\varepsilon)$.

In terms of Lyapunov functions, for each $n \geq N$ there is a rearrangement \mathbf{a}_n of \mathbf{a} such that the same estimates hold for $|V_{\mathbf{a}_n}(x) - V_{\mathbf{a}}^n(x)|$ as for $|V_{\mathbf{a}}(x) - V_{\mathbf{a}}^n(x)|$ in Case 1.

Case 3. Assume there are finitely many attractor–repeller pairs with admissible ordering $\{P_i\}_{i=1}^p$.

Let $\varepsilon > 0$. Let $\{\mathcal{P}_i^n\}$ be an admissible ordering of the attractor–repeller pairs for \mathcal{F}_n . Choose $\rho < \min\{\varepsilon, \text{dist}(A_p, A_p^*)/2\}$. By Proposition 5.5 there exists $N > 0$ such that, for all $n \geq N$, each P_i has a unique corresponding attractor–repeller pair Q_i^n such that $h(|Q_i^n|, P_i) < \rho$. Moreover, given any attractor–repeller pair Q^n the maximal attractor $\omega(|\mathcal{A}^n|)$ in $|\mathcal{A}^n|$ must be A_j for some $1 \leq j \leq p$. Therefore, by Proposition 5.5, $Q^n = Q_j^n$ or either $|\mathcal{A}^n|$ intersects $B_\rho(A_j^*)$ or $|\mathcal{A}^{*n}|$ intersects $B_\rho(A_j)$, which implies $\text{dist}(|\mathcal{A}^n|, |\mathcal{A}^{*n}|) < \rho$. Therefore we can conclude that

$$\text{dist}(|\mathcal{A}_i^n|, |\mathcal{A}_i^{*n}|) \rightarrow \begin{cases} \text{dist}(A_i, A_i^*) & \text{for } 1 \leq i \leq p, \\ 0 & \text{for } i > p, \end{cases} \quad (8)$$

as $n \rightarrow \infty$ uniformly in i .

In this case we need to adjust our definition of the Lyapunov functions to obtain convergence. First we consider the sequence $\widehat{\mathbf{a}}$ given by $\widehat{\alpha}_i = D_i \alpha_i > 0$, where $D_i = d_i / \sum_{j=1}^p d_j \alpha_j$ and $d_i = \text{dist}(A_i, A_i^*)$ for $1 \leq i \leq p$ and $\widehat{\alpha}_i = 0$ for $i > p$ (the number of attractor–repeller pairs for f). Likewise the (nonnegative) sequences $\widehat{\mathbf{a}}_n$ are defined by $\widehat{\alpha}_{n,i} = D_i^n \alpha_{n,i}$ where $D_i^n = \delta_i^n / \sum_{j=1}^p \delta_j^n \alpha_j$ and

$\delta_i^n = \text{dist}(|\mathcal{A}_i^n|, |\mathcal{A}_i^{*n}|)$ for $1 \leq i \leq p_n$ and $\widehat{\alpha}_{n,i} = 0$ for $i > p_n$. Note that $\sum_i \widehat{\alpha}_i = \sum_i \widehat{\alpha}_{n,i} = 1$ for each $n > 0$. The functions $V_{\hat{a}}$ and $V_{\hat{a}}^n$ are now defined by

$$V_{\hat{a}}(x) = \sum_{i=1}^{\infty} \widehat{\alpha}_i V_{P_i}(x) = V_{\hat{a}}^p(x) = \sum_{i=1}^p \widehat{\alpha}_i V_{P_i}(x)$$

and

$$V_{\hat{a}_n}^n(x) = \sum_{i=1}^{p_n} \widehat{\alpha}_{n,i} V_{|\mathcal{P}_i^n|}(x), \quad V_{\hat{a}_n}^{n,p}(x) = \sum_{i=1}^p \widehat{\alpha}_{n,i} V_{|\mathcal{P}_i^n|}(x).$$

By Theorem 6.7 and (8), we have $V_{\hat{a}_n}^{n,p}(x) \rightarrow V_{\hat{a}}(x)$ uniformly as $n \rightarrow \infty$ (up to possible rearrangement if some of the distances between attractor and repeller are equal for pairs for f). Also since the convergence in (8) is uniform in i , we have $V_{\hat{a}_n}^n(x) - V_{\hat{a}_n}^{n,p}(x) \rightarrow 0$ as $n \rightarrow \infty$ uniformly in x . Therefore, we have $V_{\hat{a}_n}^n(x) \rightarrow V_{\hat{a}}(x)$ as $n \rightarrow \infty$ uniformly in x . This completes the proof of Case 3. \square

Remark 6.14. Theorem 6.13 shows that discrete, weak Lyapunov functions can be constructed which approximate a continuous Lyapunov function arbitrarily closely. For a fixed grid the approximation is based only on information obtained from the multivalued map without any a priori information about the true dynamics of the system and, hence, this construction can be implemented computationally. However, as the grid size changes, the “limiting” Lyapunov function can also change due to possible rearrangement of the order of attractor–repeller pairs. It may appear from the proof of Theorem 6.13 that the construction requires a priori knowledge about whether the number of attractor–repeller pairs of f is finite or infinite. In fact, if the distances between attractor and repeller are included in the weights which determine the Lyapunov functions as in the finite case, convergence is also obtained in the infinite case. The development of efficient algorithms for computing approximate Lyapunov functions is the subject of ongoing research and will be addressed in a future work.

7. Generalizations

For the sake of clarity, the theory presented in the previous sections was restricted to compact metric spaces. However, for a wide variety of applications the dynamical systems of interest are generated by differential equations or maps defined on \mathbb{R}^m . In this section we extend the theory to this more general setting.

7.1. Locally Compact Spaces

Throughout this subsection X denotes a locally compact, separable metric space, and f is a continuous map on X . From Theorem 2.2 we obtain the existence of a

grid \mathcal{G} with $\text{diam}(\mathcal{G}) = \varepsilon$. For a finite subset $\mathcal{Y} \subset \mathcal{G}$, the objective is to define an appropriate multivalued map on \mathcal{Y} that has the right properties with respect to the underlying dynamics of f on the compact region $|\mathcal{Y}|$. The most significant change is that we now allow the empty set to be in the range of the multivalued map \mathcal{F} on \mathcal{Y} .

From an algorithmic point of view, empty images in the multivalued map pose no additional complications. The graph-theoretic algorithms to compute, for example, the recurrent set $\mathcal{R}(\mathcal{F})$ do not change; any invariant sets resulting from such computations will naturally lie in the maximal closed subgraph, i.e., the maximal invariant set of \mathcal{Y} . The only difficulty is a technical issue arising from the possibility that the maximal invariant set of the topological map f could intersect the boundary of $|\mathcal{Y}|$, and we now present the appropriate modifications of the theory to address this issue.

Definition 7.1. A multivalued map $\mathcal{F}: \mathcal{Y} \rightrightarrows \mathcal{Y}$ is an **outer approximation on $|\mathcal{Y}|$** of f if the following conditions are satisfied:

- (a) $f(G) \cap |\mathcal{Y}| \subset \text{int}_{|\mathcal{Y}|}(|\mathcal{F}(G)|)$ for all $G \in \mathcal{Y}$; and
- (b) $f^{-1}(G) \cap |\mathcal{Y}| \subset \text{int}_{|\mathcal{Y}|}(|\mathcal{F}^{-1}(G)|)$ for all $G \in \mathcal{Y}$.

Since we are reducing our calculations to a subset of X we need to understand when isolation in the reduced combinatorial system $\mathcal{F}|_{\mathcal{Y}}$ leads to isolation for $f: X \rightarrow X$.

Definition 7.2. An invariant set \mathcal{S} is **isolated in \mathcal{Y}** if $\mathcal{S} = \text{Inv}(\{G \in \mathcal{Y} \mid G \cap |\mathcal{S}| \neq \emptyset\})$.

Proposition 7.3. Let $\mathcal{F}: \mathcal{Y} \rightrightarrows \mathcal{Y}$ be an outer approximation on $|\mathcal{Y}|$ of f . Let \mathcal{S} be an isolated invariant set in \mathcal{Y} . If $\{G \in \mathcal{G} \mid G \cap |\mathcal{S}| \neq \emptyset\} \subset \mathcal{Y}$, then $N = |\mathcal{S}|$ is an isolating block for $\text{Inv}(N, f)$.

Proof. We start by proving that $N = |\mathcal{S}|$ is an isolating block. For this we adjust the proof of Theorem 4.4, arguing by contradiction. Assume there exists a point $x \in \partial N$ such that $f(x) \in N$ and $f^{-1}(x) \cap N \neq \emptyset$. There exist grid elements $G^\pm \in \mathcal{S}$ such that $f^{-1}(x) \cap G^- \neq \emptyset$ and $f(x) \in G^+$. The invariance of \mathcal{S} now implies that there exist complete orbits $\gamma_{G^\pm} \subset \mathcal{S}$. By assumption, $H \in \mathcal{Y}$ whenever $x \in H \in \mathcal{G}$. Moreover, since \mathcal{S} is isolated in \mathcal{Y} , there exists $H \in \mathcal{Y} \setminus \mathcal{S}$ such that $x \in H$.

Using the fact that \mathcal{F} is an outer approximation now implies that $f(x) \in f(H) \cap |\mathcal{Y}| \subset \text{int}_{|\mathcal{Y}|}(|\mathcal{F}(H)|)$ and, consequently, $\{G \mid f(H) \cap |\mathcal{Y}| \cap G \neq \emptyset\} \subset \mathcal{F}(H)$. This then implies that $G^+ \in \mathcal{F}(H)$. In same way, one proves that $G^- \in \mathcal{F}^{-1}(H)$. We can now construct the following complete orbit

$$\gamma_{G^-}^- \cup \{H\} \cup \gamma_{G^+}^+ \subset \mathcal{Y}$$

which does not lie in \mathcal{S} , implying that \mathcal{S} is not isolated in \mathcal{Y} , a contradiction. \square

Example 7.4. Let $f: \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a continuous map. The following collection \mathcal{G} is a grid for \mathbb{R}^m ,

$$\mathcal{G} := \left\{ \prod_{i=1}^m \left[\frac{j_i}{2}, \frac{j_i + 1}{2} \right] \mid j_i \in \mathbb{Z} \right\}.$$

Let Y be a compact neighborhood in \mathbb{R}^m and set $\mathcal{Y} = \{G \in \mathcal{G} \mid G \cap Y \neq \emptyset\}$. Define $\mathcal{F}: \mathcal{Y} \rightrightarrows \mathcal{Y}$ by $\mathcal{F}(G) := \{H \in \mathcal{Y} \mid H \cap f(G) \neq \emptyset\}$. The multivalued map \mathcal{F} is an outer approximation of f on \mathbb{R}^m . To see why this is the case, define $\mathcal{T}: \mathcal{Y} \rightrightarrows \mathcal{G}$ by $\mathcal{T}(G) := \{H \in \mathcal{G} \mid H \cap f(G) \neq \emptyset\}$. By Proposition 2.5 for each $G \in \mathcal{G}$, $f(G) \subset \text{int}(|\mathcal{T}(G)|)$. Since $\text{int}_{|\mathcal{Y}|}(|\mathcal{F}(G)|) = \text{int}(|\mathcal{T}(G)|) \cap |\mathcal{Y}|$, then Definition 7.1(a) is satisfied. A similar argument shows that Definition 7.1(b) is also satisfied.

7.2. Convergence

For clarity of presentation, the main results of Sections 5 and 6 on the convergence of the grid recurrent set to the chain-recurrent set and the convergence of Lyapunov functions were developed for the minimal combinatorial multivalued map associated to a topological map. Here we briefly state a convergence criterion for other outer approximations of maps which guarantees that these results still hold. Note that the results of Section 4 and Section 7.1 do not involve convergence and hence already apply to any outer approximation. For simplicity, we assume that X is a compact metric space with $f: X \rightarrow X$.

Definition 7.5. Given a sequence of outer approximations $\mathcal{F}_n: \mathcal{G}_n \rightrightarrows \mathcal{G}_n$ of f , with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$. The family \mathcal{F}_n is **convergent** if for any $\varepsilon > 0$ there exists an $n_\varepsilon > 0$, such that

$$|\mathcal{F}_n(G_n)| \subset B_\varepsilon(f(G_n))$$

for all $n \geq n_\varepsilon$, and for all $G_n \in \mathcal{G}_n$.

In Sections 5 and 6 the main property of the minimal multivalued map on which the proofs rested was the convergence of combinatorial orbits to topological orbits as stated in Proposition 5.4. It is clear that this property holds for a convergent family of multivalued maps as well, hence we state the following lemma whose proof is a straightforward modification of the proof of Proposition 5.4.

Lemma 7.6. *Given a convergent sequence of outer approximations $\mathcal{F}_n: \mathcal{G}_n \rightrightarrows \mathcal{G}_n$ for any $\varepsilon > 0$ and $k \geq 1$, there exists an $n_\varepsilon(k) > 0$ such that*

$$|\mathcal{F}_n^k(G_n)| \subset B_\varepsilon(f^k(G_n))$$

for all $n \geq n_\varepsilon(k)$, and for all $G_n \in \mathcal{G}_n$.

The convergence of recurrent sets and Lyapunov functions for these maps follows. The most important example of a convergent family of maps is the case where the minimal multivalued map is outer-approximated by a computed map which includes discretization and/or round-off errors. Such maps are convergent as long as the errors can be made arbitrarily small with the diameter of the grid.

7.3. Flows

So far we have considered primarily discrete dynamical systems arising from the iteration of a continuous map. There are two different approaches one could take to generalize the above results for continuous systems or flows. First, one could consider the time- T map of the flow for fixed $T > 0$, which can be approximated by numerical quadrature routines resulting in an outer approximation of the true time- T map. In the case that the discretization error goes to zero on a family of grids so that Definition 7.5 is satisfied, the previous results imply that the grid recurrent sets converge to the chain-recurrent set of the time- T map, which is equal to the chain recurrent set of the flow. Moreover, the combinatorial Lyapunov functions described above converge to a true Lyapunov function for the time- T map, which is not necessarily a Lyapunov function for the flow. However, Lemma 7.7 below provides a method for approximating a true Lyapunov function for the flow.

In [1] another approach is proposed in which the flow is approximated directly using the vector field. A multivalued map on a polygonal decomposition of a compact region of phase space is generated from a triangulation.

7.3.1. Lyapunov Functions. Recall our discussion in Section 6 and consider Lyapunov functions for a single attractor–repeller pair. Let \mathcal{P}^n be a sequence of attractor–repeller pairs for \mathcal{F}_n approaching an attractor–repeller pair P . From the construction in Section 6 we obtain a Lyapunov function $\mathcal{V}_{\mathcal{P}^n}$ given by equation (3). The piecewise constant extension is given by $V_{|\mathcal{P}^n|}(x)$. By Lemma 7.6 $V_{|\mathcal{P}^n|}(x)$ converges uniformly for $x \in X$ to some Lyapunov function $V_P(x)$.

In order to also get results for continuous-time dynamical systems we can adopt the philosophy of deriving Lyapunov functions from the time-1 dynamics. Let $\varphi: \mathbb{R} \times X \rightarrow X$ be a flow. Let $\mathbf{v}_P^*(\cdot, k)$ be pre-Lyapunov functions for the map $f = \varphi(1, \cdot)$ as introduced in Section 6. Then, to obtain a Lyapunov function for φ , we have the following lemma.

Lemma 7.7. *The function*

$$V_{P,\varphi}(x) = 2 \log 2 \int_0^1 2^{-t} V_P^*(\varphi(t, x)) dt \quad (9)$$

is a Lyapunov function for the continuous-time system φ , where $V_P^ = \sum_{m=0}^{\infty} 2^{-m-1} \mathbf{v}_P^{**}(x, m)$ and $\mathbf{v}_P^{**}(x, m) = \max_{\theta \in [0,1]} \mathbf{v}_P^*(\varphi(\theta, x), m)$, $m \geq 0$.*

Proof. The function $V_{P,\varphi}(x) = \log 2 \int_0^\infty 2^{-t} \sup_{\tau \geq 0} (\mathbf{v}_P(\varphi(t + \tau, x))) dt$ is a continuous Lyapunov function for φ , see [15]. This function can be rewritten as follows, using the continuity properties of \mathbf{v}_P^* as derived in the previous section,

$$\begin{aligned}
V_{P,\varphi}(x) &= \log 2 \int_0^\infty 2^{-t} \sup_{\tau \geq 0} (\mathbf{v}_P(\varphi(t + \tau, x))) dt \\
&= \log 2 \int_0^\infty 2^{-t} \max_{\theta \in [0,1]} \sup_{k \geq 0} (\mathbf{v}_P(\varphi(t + \theta + k, x))) dt \\
&= \log 2 \int_0^\infty 2^{-t} \max_{\theta \in [0,1]} \sup_{k \geq 0} (\mathbf{v}_P(f^k(\varphi(t + \theta, x)))) dt \\
&= \log 2 \int_0^\infty 2^{-t} \max_{\theta \in [0,1]} (\mathbf{v}_P^*(\varphi(t + \theta, x))) dt \\
&= \log 2 \sum_{m=0}^{\infty} \left\{ \int_0^1 2^{-t-m} \max_{\theta \in [0,1]} (\mathbf{v}_P^*(\varphi(t + m + \theta, x))) dt \right\} \\
&= \log 2 \int_0^1 2^{-t} \sum_{m=0}^{\infty} \max_{\theta \in [0,1]} (2^{-m} \mathbf{v}_P^*(\varphi(t + m + \theta, x))) dt \\
&= \log 2 \int_0^1 2^{-t} \sum_{m=0}^{\infty} \max_{\theta \in [0,1]} (2^{-m} \mathbf{v}_P^*(f^m(\varphi(t + \theta, x)))) dt \\
&= \log 2 \int_0^1 2^{-t} \sum_{m=0}^{\infty} 2^{-m} \mathbf{v}_P^{**}(\varphi(t, x), m) dt \\
&= 2 \log 2 \int_0^1 2^{-t} V_P^*(\varphi(t, x)) dt,
\end{aligned}$$

which proves the lemma. \square

Finally, to define a piecewise constant analogue that approximates $V_{P,\varphi}$ one can numerically approximate $\varphi(t_i, x_{G_n})$ at discrete times $t_0 = 0 < t_i < t_l = 1$ and use these points to approximate $V_{P,\varphi}$ as follows. Compute $\mathbf{v}_{|\mathcal{P}^n|}^*(\varphi(t_i, x_{G_n}), k)$ as defined in Subsection 6.2.2 and approximate $\mathbf{v}_P^*(x_{G_n}, k)$ by $\mathbf{v}_{|\mathcal{P}^n|}^{**}(x_{G_n}, k) = \max_{0 \leq i \leq l} \mathbf{v}_{|\mathcal{P}^n|}^*(\varphi(t_i, x_{G_n}), k)$. Then, letting $l \rightarrow \infty$ as $n \rightarrow \infty$, the analogue of Lemma 6.6 holds and the discretization error in both the approximation of $\varphi(t_i, x_{G_n})$ and the integral defining $V_{P,\varphi}$ will tend uniformly to zero as well, yielding convergence to $V_{P,\varphi}$.

Acknowledgments

We thank the referee for identifying several mistakes in the original version of this paper. W. K. was partially supported by NSF grants DMS-9973331 and

DMS-0511208. K. M. was partially supported by NSF grants DMS-0107396 and DMS-0511115 and DARPA. R. V. was partially supported by NWO VIDI grant 639.032.202.

References

- [1] E. Boczko, W. Kalies, and K. Mischaikow, Polygonal approximation of flows, Preprint, 2004.
- [2] C. Conley, *Isolated Invariant Sets and the Morse Index*, CBMS Regional Conference Series in Mathematics, Vol. 38, American Mathematical Society, Providence, RI, 1978.
- [3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, The MIT Electrical Engineering and Computer Science Series, MIT Press, Cambridge, MA, 1990.
- [4] S. Day, *A Rigorous Numerical Method in Infinite Dimensions*, PhD dissertation, 2003.
- [5] S. Day, O. Junge, and K. Mischaikow, A rigorous numerical method for the global analysis of infinite-dimensional discrete dynamical systems, *SIAM J. Appl. Dyn. Syst.* **3**(2) (2004), 117–160 (electronic).
- [6] M. Dellnitz and O. Junge, GAIO, <http://www-math.uni-paderborn.de/~agdellnitz/gaio>, 1998.
- [7] M. Dellnitz, G. Froyland, and O. Junge, The algorithms behind GAIO-set oriented numerical methods for dynamical systems, in *Ergodic Theory, Analysis, and Efficient Simulation of Dynamical Systems*, pp. 145–174, 805–807, Springer-Verlag, Berlin, 2001.
- [8] R. W. Easton, Isolating blocks and symbolic dynamics, *J. Differential Equations*, **17** (1975), 96–118.
- [9] M. Eidenschink, *Exploring Global Dynamics: A Numerical Algorithm Based on Conley Index Theory*, PhD dissertation, 1995.
- [10] J. Franks and M. Misiurewicz, Topological methods in dynamics, in *Handbook of Dynamical Systems*, Vol. 1A, pp. 547–598, North-Holland, Amsterdam, 2002.
- [11] F. Y. Hunt, Finite precision representation of the Conley decomposition, *J. Dynam. Differential Equations*, **13**(1) (2001), 87–105.
- [12] K. Mischaikow, Topological techniques for efficient rigorous computation in dynamics, *Acta Numer.*, **11** (2002), 435–477.
- [13] M. Mrozek, An algorithm approach to the Conley index theory, *J. Dynam. Differential Equations*, **11**(4) (1999), 711–734.
- [14] G. Osipenko and S. Campbell, Applied symbolic dynamics: Attractors and filtrations, *Discrete Contin. Dynam. Systems*, **5**(1) (1999), 43–60.
- [15] C. Robinson, *Dynamical Systems*, Studies in Advanced Mathematics, CRC Press, Boca Raton, FL, Stability, Symbolic Dynamics, and Chaos, 2nd ed., 1999.
- [16] A. Szyczak, A combinatorial procedure for finding isolating neighbourhoods and index pairs, *Proc. Roy. Soc. Edinburgh Sect. A* **127**(5) (1997), 1075–1088.