

A Computational Approach to Conley's Decomposition Theorem

Hyunju Ban and William D. Kalies

Department of Mathematical Sciences
Florida Atlantic University
Boca Raton, FL 33431

ABSTRACT. The discrete dynamics generated by a continuous map can be represented combinatorially by an appropriate multivalued map on a discretization of the phase space such as a cubical grid or triangulation. In this paper we provide explicit algorithms and computational complexity bounds for computing dynamical structures for the resulting combinatorial multivalued maps. Specifically we focus on the computation attractor-repeller pairs and Lyapunov functions for Morse decompositions. These discrete Lyapunov functions are weak Lyapunov functions and well-approximate a continuous Lyapunov function for the underlying map.

KEYWORDS. chain recurrence, Lyapunov function, Conley's Decomposition Theorem, algorithms, computation, combinatorial dynamics.

On the occasion of Philip Holmes' 60th birthday.

1. Introduction

Identifying recurrent behavior is of fundamental importance in dynamical systems theory, and several characterizations of recurrence have proved useful, e.g. periodicity, nonwandering points, and chain recurrence. In graph theory, the notion of recurrence corresponds to periodicity, the existence of a nontrivial path from a vertex back to itself, and the vertices of a finite directed graph can be separated into recurrent and nonrecurrent components by a linear-time computation. These observations lead naturally to the question of whether dynamical systems can be faithfully represented combinatorially by a finite directed graph, which can then be used for efficient computation.

Conley's Decomposition Theorem asserts that every dynamical system can be separated into a minimal (chain) recurrent set and its complement on which the dynamics is gradient-like, i.e. there exists a Lyapunov function which is strictly decreasing along orbits. The chain recurrent set can then be divided into components which are partially-ordered by the existence of connecting orbits between them. In general, this does not lead to a finite directed graph because the number of recurrent components can be infinite. However, by grouping together recurrent components and connections, one can produce finite partially-ordered collections of invariant sets called Morse decompositions, which do provide representations of the dynamics by finite directed graphs.

The standard theoretical framework for Conley's Decomposition Theorem does not lead naturally to a computational method for approximating the chain recurrent set, constructing Morse decompositions, or approximating Lyapunov functions. In [1], an alternative approach based on finite spatial discretizations and combinatorial multivalued maps is developed. As we outline

briefly below, the main features of this approach are (1) the basic notions of recurrent set, attractor, Morse decomposition, Lyapunov function, etc. are meaningful in combinatorial systems, (2) there is convergence of these dynamical structures on the combinatorial level to corresponding structures in the underlying dynamical system, and (3) the tools of the Conley index theory can be applied to a combinatorial approximation to provide rigorous computer-assisted proofs of qualitative dynamics, see [2, 3, 4, 5].

1.1. Combinatorial multivalued maps and dynamics. The development of a computational framework for the construction of a combinatorial representation of a dynamical system requires a finite description of both the metric space X and the dynamical system φ . To introduce these ideas of combinatorialization and dynamics, we consider a discrete dynamical system generated by a surjective map $f : X \rightarrow X$. The more general cases of continuous time and noninvariant domains require a few more technicalities and are considered via examples in later sections. We begin with a general form of discretization of the phase space.

DEFINITION 1.1 ([6]). A **grid** on X is a finite collection \mathcal{G} of nonempty compact subsets of X with the following properties:

- (i) $X = \cup_{G \in \mathcal{G}} G$,
- (ii) $G = \text{cl}(\text{int}(G))$ for all $G \in \mathcal{G}$,
- (iii) $G \cap \text{int}(H) = \emptyset$ for all $G \neq H \in \mathcal{G}$.

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 for $\mathcal{S} \subset \mathcal{G}$ by $|\mathcal{S}| := \cup_{S \in \mathcal{S}} S \subset X$. The existence of grids of arbitrarily small size easily follows from compactness. In practice, $X \subset \mathbb{R}^n$ and \mathcal{G} is a cubical grid, triangulation, or polygonal tiling, but the theory requires no extra work for arbitrary grids.

We now construct a multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ which maps grid elements to sets of grid elements and is a combinatorial representation of the action of f on X . Of primary importance is the ability to go backward from the combinatorial information contained in \mathcal{F} to information about the topological dynamics of f . As described in [1], the essential requirement is that images of the multivalued map \mathcal{F} well-cover images of f , as first recognized by Szymczak [3].

DEFINITION 1.2 ([3]). 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 $G \in \mathcal{G}$.

The smallest such outer approximation is given as follows.

DEFINITION 1.3. The **minimal multivalued map** $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ associated to f on the grid \mathcal{G} is defined by $\mathcal{F}(G) := \{H \in \mathcal{G} \mid H \cap f(G) \neq \emptyset\}$.

The minimal multivalued map is a natural setting in which to design algorithms to approximate the discrete dynamics of a continuous map. However, determining the minimal multivalued map \mathcal{F} requires complete knowledge of the image of f on grid elements, which is computationally impractical. Often one needs to enlarge the images of \mathcal{F} , for example to take into account truncation error and numerical error.

PROPOSITION 1.4 ([1]). *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 .*

General computer software has been developed to compute such multivalued maps, cf. GAIO [7, 8]. If errors are incorporated into the construction of an outer approximation, then rigorous information about the dynamics of f can be obtained using the Conley index, see [2, 3, 4, 5].

Note that the combinatorial system $(\mathcal{G}, \mathcal{F})$ is simply a different way of describing a directed graph whose vertices are elements of \mathcal{G} with an edge from G to H if $H \in \mathcal{F}(G)$. For example,

consider the logistic map $f(x) = 2.5x(1 - x)$. Figure 1 shows the directed graph corresponding to the minimal multivalued map for the grid obtained by dividing the domain $X = [0, 5/8]$ into five equal-length subintervals.




FIGURE 1. Directed graph representation of the minimal multivalued map for the logistic map $f(x) = 2.5x(1 - x)$. The pair $(\{G_5\}, \{G_1\})$ is an attractor-repeller pair for \mathcal{F} .

We now adapt some dynamical notions to this combinatorial setting, taking into account the fact that \mathcal{F} is multivalued. We present here only the material necessary for the results of this paper; for a more complete, detailed description the reader is referred to [1]. Note that since f is surjective on X and \mathcal{F} is an outer approximation, the multivalued map has the property that $\mathcal{F}(G)$ and $\mathcal{F}^{-1}(G)$ are nonempty for all $G \in \mathcal{G}$. We call such a multivalued map **closed**.

DEFINITION 1.5. $\mathcal{S} \subset \mathcal{G}$ is an **invariant set** for \mathcal{F} if $\mathcal{S} \subset \mathcal{F}(\mathcal{S})$ and $\mathcal{S} \subset \mathcal{F}^{-1}(\mathcal{S})$.

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} . The **maximal invariant set** in $\mathcal{U} \subset \mathcal{G}$ is denoted by $\text{Inv}(\mathcal{U})$. 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 $\Gamma_-^k(\mathcal{U}) = \Gamma^{-\infty,k}$ as the **k -backward image** of \mathcal{U} .

DEFINITION 1.6. 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} in the subgraph $\mathcal{G} \setminus \mathcal{A}$.

Figure 1 shows an example of a combinatorial attractor-repeller pair. Refining the notion of an attractor-repeller pair leads to the following definition.

DEFINITION 1.7. 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 \succ on the index set $\{1, \dots, n\}$ that satisfies the following property. Given a complete orbit $\{G_k\}_{k \in \mathbb{Z}}$ for which there exist unique $1 \leq i, j \leq n$ and there exist $k^\pm \in \mathbb{Z}$ such that $G_k \in \mathcal{S}_i$ for all $k < k^-$ and $G_k \in \mathcal{S}_j$ for all $k > k^+$ and if $i \neq j$, then $i \succ j$. The sets \mathcal{S}_i are called **Morse sets**.

DEFINITION 1.8. Let $\mathcal{F} : \mathcal{G} \rightarrow \mathcal{G}$ be closed. The **recurrent set** of \mathcal{F} is defined by $\mathcal{R}(\mathcal{F}) = \{G \in \mathcal{G} \mid G \in \mathcal{F}^n(G) \text{ for some } n > 0\}$. The **recurrent components** are the equivalence classes of $\mathcal{R}(\mathcal{F})$ defined by the relation $G \sim H$ if there exist $n, m > 0$ such that $H \in \mathcal{F}^m(G)$ and $G \in \mathcal{F}^n(H)$.

Observe that components of $\mathcal{R}(\mathcal{F})$ are the Morse sets of the finest Morse decomposition. In graph-theoretic terms, the recurrent components correspond precisely to the nontrivial strongly-connected components or cycles of \mathcal{F} which contain at least one edge.

PROPOSITION 1.9 ([1]). *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, $\mathcal{R}(\mathcal{F}) = \bigcap_{j=1}^J (\mathcal{A}_j \cup \mathcal{A}_j^*)$*

For any Morse decomposition, including the finest Morse decomposition given by the recurrent components, the multivalued map \mathcal{F} has a gradient-like structure off of the Morse sets \mathcal{S}_i . To make this precise we introduce the notion of a Lyapunov function on a directed graph.

DEFINITION 1.10. 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 $L : \mathcal{G} \rightarrow [0, 1]$ satisfying:

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

Since \mathcal{G} is a finite set and we are considering Lyapunov functions with images in \mathbb{R} , we can choose $L : \mathcal{G} \rightarrow [0, 1]$ with the following properties.

PROPOSITION 1.11 ([1]). *There exists a Lyapunov function $L : \mathcal{G} \rightarrow [0, 1]$ such that for any $c \in \mathbb{R}$, $L^{-1}(c)$ is either a Morse set \mathcal{S}_i , an element of $\mathcal{G} \setminus \cup_i \mathcal{S}_i$, or empty.*

Discrete Lyapunov functions for Morse decompositions of \mathcal{F} can easily be constructed using linear-time graph algorithms, see [9]. One can also think of a discrete Lyapunov function as a piecewise-constant function defined on X which is constant on grid elements. However, these functions need not be weak Lyapunov functions or approximate continuous Lyapunov functions for the underlying map f since they do not take into account any topological information. A **weak Lyapunov function** for f on X satisfies $L(x) \geq L(f(x))$ for all $x \in X$.

1.2. Outline. In the previous section, we build and identify structures in a multivalued map or directed graph which parallel the dynamical notions surrounding the idea of recurrence, namely attractor-repeller pairs, Morse decompositions, Lyapunov functions, and the recurrent set. In [1], it is shown that information from these combinatorial structures can often be pulled back to provide rigorous information about the dynamics of the underlying system using the Conley index theory, see also [2, 3, 4, 5]. Moreover, it is shown that these structures in a sequence of grids \mathcal{G}_n with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ converge to the corresponding structures in the underlying dynamical system.

In this work, we focus on issues of computation. In particular we provide algorithms for finding all attractor-repeller pairs in a multivalued map or directed graph and for constructing Lyapunov functions which approximate continuous Lyapunov functions for the underlying dynamics. In Section 2 we develop an algorithm for finding all combinatorial attractor-repeller pairs and a bound on its complexity. In Section 3 we explain how to compute discrete Lyapunov functions which are related to the underlying dynamics and give an overall complexity bound in Theorem 3.10. To provide an efficient method for computing such Lyapunov functions, we design an algorithm different from that given in [1]. Some examples from both maps and flows can be found in Section 4.

There are open problems for future work. For example, to compute a nice Lyapunov function for the recurrent set, it is not necessary to use all attractor-repeller pairs in general; one can identify certain irreducible attractors on which to compute Lyapunov functions, which may reduce the overall cost, cf. [10]. Also, in certain situations it may be possible to approximate smooth Lyapunov functions, which may be useful in some applications.

Finally, in certain places the algorithms require a search of a given list. The complexity of this search depends on the data structures used to implement the algorithm. When presenting the computational complexity of an algorithm, we will use the notation $l(n)$ to denote the complexity to search a list of n elements.

2. Computation of attractor-repeller pairs

In this section, we introduce an algorithm for finding all attractor-repeller pairs in a directed graph and provide an estimate of the computational complexity.

2.1. Preliminaries. First we establish some important properties for attractor-repeller pairs in a closed direct graph.

LEMMA 2.1. *Let $\mathcal{B} \subset \mathcal{G}$ satisfy $\mathcal{F}(\mathcal{B}) \supset \mathcal{B}$. Then $\Gamma_+(\mathcal{B})$ is an attractor containing \mathcal{B} .*

Proof: $\mathcal{F}(\Gamma_+(\mathcal{B})) = \bigcup_{n \geq 1} \mathcal{F}^n(\mathcal{B}) = \bigcup_{n \geq 0} \mathcal{F}^n(\mathcal{B}) = \Gamma_+(\mathcal{B})$ since $\mathcal{F}^0(\mathcal{B}) = \mathcal{B}$ and $\mathcal{F}^1(\mathcal{B}) = \mathcal{F}(\mathcal{B}) \supset \mathcal{B}$. So by definition $\Gamma_+(\mathcal{B})$ is an attractor and contains \mathcal{B} . ■

LEMMA 2.2. *Every attractor contains at least one recurrent component, and if $\mathcal{R} \cap \mathcal{A} \neq \emptyset$ for some recurrent component \mathcal{R} and attractor \mathcal{A} , then $\mathcal{R} \subset \mathcal{A}$.*

Proof: Let \mathcal{A} be an attractor. Then $\mathcal{F}(\mathcal{A}) = \mathcal{A}$ so that \mathcal{A} is invariant set, and hence for all $G \in \mathcal{A}$ there exists a complete orbit $\gamma_G = \{G_k\}_{k \in \mathbb{Z}}$ in \mathcal{A} . Since \mathcal{F} is closed and finite, there exists H such that $G_{k_n} = H$ for some subsequence. Then H is recurrent and $\mathcal{R}(H) \subset \mathcal{A}$ where $\mathcal{R}(H)$ is the recurrent component containing H . Indeed it follows that for every complete orbit G_k there exists $k_* \geq 0$ and a recurrent component \mathcal{R} such that $G_k \in \mathcal{R}$ for all $k \geq k_*$, and a similar result holds for backward orbits. To prove the second statement, if $H \in \mathcal{R} \cap \mathcal{A}$ and $G \in \mathcal{R}$, then $G \in \mathcal{F}^n(H)$ for some $n > 0$. Thus $G \in \mathcal{A}$ since $\mathcal{F}(\mathcal{A}) = \mathcal{A}$. ■

THEOREM 2.3. *Let \mathcal{P}_R be the set of all possible unions of recurrent components. If $\mathcal{S} \in \mathcal{P}_R$, then $\Gamma_+(\mathcal{S})$ is an attractor containing \mathcal{S} . Moreover, for every attractor \mathcal{A} we have $\mathcal{A} = \Gamma_+(\mathcal{S}_\mathcal{A})$ where $\mathcal{S}_\mathcal{A} = \mathcal{R}(\mathcal{F}) \cap \mathcal{A} \in \mathcal{P}_R$, and its dual repeller $\mathcal{A}^* = \Gamma_+(\mathcal{R}(\mathcal{F}) \cap \mathcal{A}^*)$.*

PROOF. Let \mathcal{R} be a recurrent component, then \mathcal{R} is invariant so that $\mathcal{R} \subset \mathcal{F}(\mathcal{R})$. Clearly this property holds for a union $\mathcal{S} \in \mathcal{P}_R$ as well, and thus $\Gamma_+(\mathcal{S})$ is an attractor containing \mathcal{S} by Lemma 2.1.

Now let \mathcal{A} be an attractor. Then $\mathcal{S}_\mathcal{A} \neq \emptyset$ by Lemma 2.2. Let $\mathcal{A}' = \Gamma_+(\mathcal{S}_\mathcal{A})$. We now show that $\mathcal{A} = \mathcal{A}'$. Since $\mathcal{F}(\mathcal{A}) = \mathcal{A}$, the images $\mathcal{F}^n(\mathcal{S}_\mathcal{A}) \subset \mathcal{A}$ for all $n \geq 0$ so that $\mathcal{A}' \subset \mathcal{A}$. Suppose $G \in \mathcal{A}' \setminus \mathcal{A}$. Since $\mathcal{F}(\mathcal{A}) = \mathcal{A}$ and $\mathcal{F}(\mathcal{A}') = \mathcal{F}(\mathcal{A}')$, there is a backward orbit $\{G_k\}_{k=-\infty}^0 \subset \mathcal{A}' \setminus \mathcal{A}$. Using the same argument as in the proof of Lemma 2.2, there exists $k_* \leq 0$ and a recurrent component \mathcal{R} such that $G_k \in \mathcal{R}$ for all $k \leq k_*$. Hence $\mathcal{R} \subset \mathcal{A}' \setminus \mathcal{A}$ by Lemma 2.2, which is a contradiction. The analogous result for the dual repeller follows from the fact that \mathcal{A}^* is an attractor for \mathcal{F}^{-1} . ■

COROLLARY 2.4. *Let $\mathcal{R}_1, \dots, \mathcal{R}_s$ be recurrent components. Let $\mathcal{A}_i = \Gamma_+(\mathcal{R}_i)$ be their corresponding attractors with $\mathcal{S} = \bigcup_i \mathcal{R}_i$. Then $\Gamma_+(\mathcal{S}) = \mathcal{A} = \bigcup_i \mathcal{A}_i$.*

To summarize, all attractor-repeller pairs can be found by first computing the recurrent components and then taking the forward images of all possible unions of recurrent components. In the next section, we will provide an algorithm and an estimate of computational complexity based on graph-theoretic techniques.

2.2. Algorithms. As seen in Theorem 2.3 and Corollary 2.4, the recurrent components play an important role in computing attractor-repeller pairs. In fact, by taking all forward and backward images from all possible unions of recurrent components we can construct all possible attractor-repeller pairs, but we would like a more efficient algorithm. Also, attractors generated by different unions of recurrent components might be the same, and we would like to avoid redundant calculation. In order to control this problem, we introduce the *inclusion graph* between recurrent components describing which recurrent components are reachable by which

other components. Thus the inclusion graph \mathcal{I} has one vertex for each recurrent component and there is an edge from the vertex corresponding to \mathcal{R}_1 to the vertex corresponding to \mathcal{R}_2 if $\mathcal{R}_2 \subset \Gamma_+(\mathcal{R}_1)$. We now present a complete algorithm for finding all attractor-repeller pairs. Let $\mathcal{R}_1, \dots, \mathcal{R}_s$ be the set of all recurrent components.

STEP 1. Compute inclusion graph and attractors from single recurrent components.

```

ComputeInclusionGraphAt( $G$ )
1. color( $G$ )  $\leftarrow$  BLACK
2. for each  $H \in \mathcal{F}(G)$ 
3.   if  $H = \mathcal{R}_j$  then add  $j$  to  $I(G)$ 
4.   if color( $H$ )=WHITE then ComputeInclusionGraphAt( $H$ )
5.   if color( $H$ )=BLACK then add  $I(H)$  to  $I(G)$ 
6.   if  $G \in \mathcal{R}_i$  then add  $I(G)$  to  $\mathcal{I}(i)$ 

ComputeAttractorFrom( $G, i$ )
1. color( $G$ )  $\leftarrow$  BLACK
2. add  $G$  to  $ATR$ 
3. for each  $H \in \mathcal{F}(G)$ 
4.   if  $H = \mathcal{R}_j$  then add  $ATT(j)$  to  $ATR$ 
5.   else if color( $H$ )=WHITE then ComputeAttractorFrom( $H, i$ )

main()
1. for each  $G \in \mathcal{G}$ , color( $G$ )  $\leftarrow$  WHITE
2. for each  $G \in \mathcal{G}$ 
3.   if color( $G$ )=WHITE then ComputeInclusionGraphAt( $G$ )
4.  $L = \text{TopologicalSort}(\mathcal{I})$ 
5. for each  $G \in \mathcal{G}$  color( $G$ )  $\leftarrow$  WHITE
6. for each  $1 \leq i \leq s$  in ascending values of  $L(i)$ 
7.   make  $ATR$  empty
8.   ComputeAttractorFrom( $G \in L^{-1}(i), i$ )
9.   add  $ATR$  to  $ATT(i)$ 

```

The above algorithm assumes that the recurrent components have been computed, which requires linear time, and that the recurrent components have each been compressed into a single vertex which requires $O(r l(r))$ time, where r is the number of grid elements in the forward image of the recurrent set, $\mathcal{F}(\mathcal{R})$. Also, the function `TopologicalSort()` operates on the acyclic graph \mathcal{I} , the inclusion graph for the recurrent components, and produces a Lyapunov function L for \mathcal{I} in linear time, see Sections 23.4 and 23.5 in [9]. We note that during the `ComputeInclusionGraphAt()` function, $I(G)$ is stored for each grid element G because components might not be connected directly in the original graph.

The algorithm produces a list $ATT(i)$ for $1 \leq i \leq s$ of all attractors generated by a single recurrent component. Also we note we can apply the same procedure for the reverse graph to get all repellers generated by a single component, say $REP(i)$ for $1 \leq i \leq s$. Then the vectors of lists A and R of all attractor-repeller pairs is generated as follows.

STEP 2. Compute all attractor-repeller pairs.

```

IsMaximal(list  $L$ )
1. for each  $l \in L$ , if  $\mathcal{I}(l) \not\subset L$  then return FALSE

```

2. return TRUE

FindMaximalInclusions(int n , list L)

1. add n to list L
2. if IsMaximal(L), save L in PL
3. for each $n + 1 \leq j \leq s$, FindMaximalInclusions(j, L)

ComputeAttractorRepellerPairs()

1. for each $1 \leq i \leq s$
2. FindMaximalInclusions($i, L = \text{empty list}$)
3. for each list L in PL and for each $1 \leq i \leq s$
4. if $i \in L$, $A(L) = A(L) \cup ATT(i)$
5. if $i \notin L$, $R(L) = R(L) \cup REP(i)$

2.3. Computational complexity. In the above algorithms, we employ a linear-time, depth-first search style algorithm a constant number of times to compute the recurrent components, the inclusion graph, and all attractor-repeller pairs. In addition, the algorithm to test for the maximality of unions of recurrent components requires at most $s l(s)$ comparisons in IsMaximal() which are performed at most p times, where p is the number of attractor-repeller pairs. Therefore we have the following theorem. Recall that $l(\cdot)$ is the complexity of searching a list in whatever data structures are used to implement the algorithm.

THEOREM 2.5. *Given a multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$, the complexity to find all distinct attractor-repeller pairs is at most $O(|\mathcal{F}| + |\mathcal{G}| + sp + r l(r))$, where $|\mathcal{F}|$ and $|\mathcal{G}|$ are the numbers of edges and vertices in the graph respectively, s is the number of recurrent components, p is the number of attractor-repeller pairs, and r is the number of elements in the forward image of \mathcal{R} . In the worst case $p \leq 2^s$.*

3. Computation of Lyapunov functions

We will follow the general scheme in [1] with some modifications to implement an algorithm to approximate Lyapunov functions. We assume that an outer approximation $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ of f is given which is closed. We will approximate a continuous Lyapunov function by a function which is constant on the interiors of the grid elements.

Let $\mathcal{P} = (A, A^*)$ be a given attractor repeller pair. The first step is to define such a piecewise constant function which measures the relative distances from a grid element to the attractor and the repeller.

DEFINITION 3.1. *Let A and B be compact subsets of X where the intersection $A \cap B$ need not be empty. 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*

- (i) $\mathbf{v}_{(A,B)}$ is locally Lipschitz continuous on $X \setminus (A \cap B)$,
- (ii) $\mathbf{v}_{(A,B)}^{-1}(0) = A \setminus (A \cap B)$,
- (iii) $\mathbf{v}_{(A,B)}^{-1}(1) = B \setminus (A \cap B)$.

While any distance potential can be used, we consider the function

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

where the distance function $d(\cdot, \cdot)$ is derived from the metric on X . In [1] it was shown that $\mathbf{v}_{(A,B)}(x)$ is a distance potential. For each grid element G , choose a representative $x_G \in \text{int}(G)$.

Define $\mathbf{v}_{\mathcal{P}} : \mathcal{G} \rightarrow \mathbb{R}$ by $\mathbf{v}_{\mathcal{P}}(G) = \mathbf{v}_{(|\mathcal{A}|, |\mathcal{A}^*|)}(x_G)$. Since $|\mathcal{A}|$ and $|\mathcal{A}^*|$ intersect possibly only on the boundary of grid elements, $\mathbf{v}_{\mathcal{P}}$ is well-defined and $\mathbf{v}_{\mathcal{P}}(G) = 0$ for all $G \in \mathcal{A}$ and $\mathbf{v}_{\mathcal{P}}(G) = 1$ for all $G \in \mathcal{A}^*$.

Following the construction in [1], for the attractor-repeller pair \mathcal{P} a discrete Lyapunov function on \mathcal{G} , which is a piecewise constant weak Lyapunov function on X , can be obtained by $\sum_{k=1}^{\infty} 2^{-k-1} \mathbf{v}_{\mathcal{P}}^*(G, k)$ where $\mathbf{v}_{\mathcal{P}}^*(G, k) = \max_{H \in \Gamma_+^k(G)} \mathbf{v}_{\mathcal{P}}(H)$. To compute $\mathbf{v}_{\mathcal{P}}^*(G, k)$ for fixed G and k , one can use a depth-first search, but this method is highly inefficient if a search is performed for each grid element and each relevant $k > 0$. To obtain a reasonably efficient algorithm we use the fact that the graph on $\mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$ is acyclic and every complete orbit γ_G with $G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$ is a path between the attractor and repeller, cf. [1] (Proposition 3.8), and define a new Lyapunov function by

$$(3.1) \quad \mathcal{W}_{\mathcal{P}}(G) = \begin{cases} 0 & \text{if } G \in \mathcal{A}, \\ 1 & \text{if } G \in \mathcal{A}^*, \\ \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(G, 0) + \frac{1}{2} \max_{H \in \mathcal{F}(G)} \{\mathcal{W}_{\mathcal{P}}(H)\} & \text{otherwise.} \end{cases}$$

Note that \mathcal{W} is recursively defined with $\mathcal{W}_{\mathcal{P}}(G) \in [0, 1]$ for all $G \in \mathcal{G}$.

LEMMA 3.2. *The function $\mathcal{W}_{\mathcal{P}}$ defined in equation (3.1) is a Lyapunov function in the sense of Definition 1.10 with $\mathcal{W}_{\mathcal{P}}^{-1}(0) = \mathcal{A}$ and $\mathcal{W}_{\mathcal{P}}^{-1}(1) = \mathcal{A}^*$. Moreover, $\mathcal{W}_{\mathcal{P}}$ is a piecewise-constant, weak Lyapunov function for f on X provided \mathcal{F} is an outer approximation of f .*

Proof: First recall some properties of $\mathbf{v}_{\mathcal{P}}^*$ which can be found in Lemma 6.5 of [1]: $\mathbf{v}_{\mathcal{P}}^*(G, k) = 0$ if and only if $G \in \mathcal{A}$, $\mathbf{v}_{\mathcal{P}}^*(G, k) = 1$ if and only if $G \in \mathcal{A}^*$, and $\mathbf{v}_{\mathcal{P}}^*(G, k) \geq \mathbf{v}_{\mathcal{P}}(H, k)$ for all $H \in \mathcal{F}(G)$.

By definition $\mathcal{W}_{\mathcal{P}}(G) = 0$ for all $G \in \mathcal{A}$ and $\mathcal{W}_{\mathcal{P}}(G) = 1$ for all $G \in \mathcal{A}^*$. For $G \notin \mathcal{A}$, we have $\mathcal{W}_{\mathcal{P}}(G) \geq \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(G, 0) > 0$. If $G \notin \mathcal{A}^*$, then $\mathcal{W}_{\mathcal{P}}(G) \leq \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(G, 0) + \frac{1}{2} < 1$, which implies (a).

To prove (b) and (a) we will show $\mathcal{W}_{\mathcal{P}}(G) > \mathcal{W}_{\mathcal{P}}(H)$ for $G \notin (\mathcal{A} \cup \mathcal{A}^*)$. Define

$$C_i = \{G \notin \mathcal{A} \cup \mathcal{A}^* \mid i \geq 0 \text{ is the minimum integer such that } \mathcal{F}^i(G) \subset \mathcal{A}\}.$$

Note that $C_0 = \mathcal{A}$ and $\mathcal{A}^* \cap C_i = \emptyset$ for all $i \geq 0$. Now we proceed by induction on $i > 0$. For $G \in C_1$ we have $\mathcal{W}_{\mathcal{P}}(G) > 0 = \mathcal{W}_{\mathcal{P}}(H)$ for all $H \in \mathcal{F}(G)$ since $H \in \mathcal{A}$ and $G \notin \mathcal{A}$.

Let $G \in C_{i+1}$. Then for any $H \in \mathcal{F}(G)$ and for all $I \in \mathcal{F}(H)$ we have $H \in C_i$ and

$$\mathcal{W}_{\mathcal{P}}(G) \geq \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(G, 0) + \frac{1}{2} \mathcal{W}_{\mathcal{P}}(H) > \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(G, 0) + \frac{1}{2} \mathcal{W}_{\mathcal{P}}(I) \geq \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(H, 0) + \frac{1}{2} \mathcal{W}_{\mathcal{P}}(I)$$

where the strict inequality follows by induction, and thus

$$\mathcal{W}_{\mathcal{P}}(G) > \frac{1}{2} \mathbf{v}_{\mathcal{P}}^*(H, 0) + \frac{1}{2} \max_{I \in \mathcal{F}(H)} \{\mathcal{W}_{\mathcal{P}}(I)\} = \mathcal{W}_{\mathcal{P}}(H) \text{ for all } H \in \mathcal{F}(G).$$

■

Now we need to prove that the discrete function $\mathcal{W}_{\mathcal{P}}$ converges to a continuous Lyapunov function for f on X . Here we assume that \mathcal{F} is the minimal multivalued map for f on \mathcal{G} , though one can define a notion of convergent outer approximations, see [1]. We use several results from [1] which we briefly review here.

Let $P = (A, A^*)$ be an attractor-repeller pair for f . Define the function

$$V_P(x) = \sum_{k=0}^{\infty} 2^{-k-1} \mathbf{v}_P^*(x, k) \quad \text{where} \quad \mathbf{v}_P^*(x, k) = \max_{y \in \gamma_x^k} \mathbf{v}_P(y) \quad \text{for } k \geq 0.$$

Then V_P is a Lyapunov function for the attractor-repeller pair $P = (A, A^*)$, as was proved by several authors in the case of homeomorphisms [11, 12, 13] and in general in [1].

Let \mathcal{G}_n be a sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ as $n \rightarrow \infty$. By Theorem 5.5 in [1] there exist attractor-repeller pairs $\mathcal{P}^n = (\mathcal{A}^n, \mathcal{A}^{*n})$ such that $(|\mathcal{A}^n|, |\mathcal{A}^{*n}|) \rightarrow (A, A^*)$ in the Hausdorff metric. Define $\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) = \max_{y \in |\Gamma_+^k(G)|} \mathbf{v}_{|\mathcal{P}^n|}(y)$ for $k \geq 0$. We now show that $\mathcal{W}_{\mathcal{P}^n}$ converges to V_P uniformly on X .

THEOREM 3.3. *For every $\epsilon > 0$ there exists a constant $M > 0$ and $N(\epsilon) \geq 0$ such that*

$$|V_P(x) - \mathcal{W}_{\mathcal{P}^n}(x)| < \epsilon[1 + \text{Lip}_{\mathbf{V}} + M(\text{Lip}_{\mathbf{V}} + 1)] \quad \text{for all } n \geq N(\epsilon).$$

First we need the following lemma from [1].

LEMMA 3.4 ([1] Lemma 6.6). *Let $x \in G \in \mathcal{G}$. Then for each k and $\epsilon > 0$ there exists an $N(\epsilon, k)$ such that*

$$|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_P^*(x, k)| < \epsilon \quad \text{uniformly in } x \in X,$$

and

$$|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_{\mathcal{P}^n}^*(G, k)| \leq \epsilon \text{Lip}_{\mathbf{V}} \quad \text{for } n \geq N(\epsilon, k).$$

COROLLARY 3.5. *Let $x \in G \in \mathcal{G}$ and $H \in \mathcal{F}^k(G)$ for $k \geq 0$. Then for each $k \geq 0$ and $\epsilon > 0$ there exists an $N(\epsilon, k)$ such that*

$$|\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_{\mathcal{P}^n}^*(H, 0)| \leq \epsilon \text{Lip}_{\mathbf{V}} + 2k\epsilon(\text{Lip}_{\mathbf{V}} + 1) \quad \text{for } n \geq N(\epsilon, k).$$

Proof: First we prove that

$$|\mathbf{v}_{\mathcal{P}^n}^*(G, k) - \mathbf{v}_{\mathcal{P}^n}^*(H, k-1)| < 2\epsilon(\text{Lip}_{\mathbf{V}} + 1) \quad \text{for } H \in \mathcal{F}(G) \text{ and } k \geq 1.$$

Since $H \in \mathcal{F}(G)$, there is $x' \in G$ such that $f(x') \in H$, and $\mathbf{v}_P^*(f(x'), k-1) = \mathbf{v}_P^*(x', k)$. So by the Lemma 3.4 there exists N such that

$$\begin{aligned} & |\mathbf{v}_{\mathcal{P}^n}^*(H, k-1) - \mathbf{v}_{\mathcal{P}^n}^*(G, k)| \\ & \leq |\mathbf{v}_{\mathcal{P}^n}^*(H, k-1) - \mathbf{v}_{|\mathcal{P}^n|}^*(f(x'), k-1)| \\ & \quad + |\mathbf{v}_{|\mathcal{P}^n|}^*(f(x'), k-1) - \mathbf{v}_P^*(f(x'), k-1)| + |\mathbf{v}_P^*(f(x'), k-1) - \mathbf{v}_{|\mathcal{P}^n|}^*(x', k)| \\ & \quad + |\mathbf{v}_{|\mathcal{P}^n|}^*(x', k) - \mathbf{v}_{\mathcal{P}^n}^*(G, k)| \\ & \leq |\mathbf{v}_{\mathcal{P}^n}^*(H, k-1) - \mathbf{v}_{|\mathcal{P}^n|}^*(f(x'), k-1)| + |\mathbf{v}_{|\mathcal{P}^n|}^*(f(x'), k-1) - \mathbf{v}_P^*(f(x'), k-1)| \\ & \quad + |\mathbf{v}_P^*(x', k) - \mathbf{v}_{|\mathcal{P}^n|}^*(x', k)| + |\mathbf{v}_{|\mathcal{P}^n|}^*(x', k) - \mathbf{v}_{\mathcal{P}^n}^*(G, k)| \\ & \leq \epsilon \text{Lip}_{\mathbf{V}} + \epsilon + \epsilon + \epsilon \text{Lip}_{\mathbf{V}} = 2\epsilon(\text{Lip}_{\mathbf{V}} + 1). \end{aligned}$$

Thus we have

$$\begin{aligned} & |\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_{\mathcal{P}^n}^*(H, 0)| \\ & \leq |\mathbf{v}_{|\mathcal{P}^n|}^*(x, k) - \mathbf{v}_{\mathcal{P}^n}^*(G_0, k)| + |\mathbf{v}_{\mathcal{P}^n}^*(G_0, k) - \mathbf{v}_{\mathcal{P}^n}^*(G_1, k-1)| \\ & \quad + |\mathbf{v}_{\mathcal{P}^n}^*(G_1, k-1) - \mathbf{v}_{\mathcal{P}^n}^*(G_2, k-2)| + \cdots + |\mathbf{v}_{\mathcal{P}^n}^*(G_{m-1}, 1) - \mathbf{v}_{\mathcal{P}^n}^*(H, 0)| \\ & \leq \epsilon \text{Lip}_{\mathbf{V}} + 2k\epsilon(\text{Lip}_{\mathbf{V}} + 1), \end{aligned}$$

where $G_0 = G$, $G_k = H$ and $G_i \in \mathcal{F}(G_{i-1})$ for each $1 \leq i \leq k$. ■

Proof of Theorem 3.3: Using Lemma 3.4, it is shown in [1] (Theorem 6.7) that there exists $N(\epsilon)$ such that $|V_P(x) - V_{|\mathcal{P}^n|}(x)| < \epsilon$, and so we have

$$\begin{aligned} |V_P(x) - \mathcal{W}_{\mathcal{P}^n}(x)| & \leq |V_P(x) - V_{|\mathcal{P}^n|}(x)| + |V_{|\mathcal{P}^n|}(x) - \mathcal{W}_{\mathcal{P}^n}(x)| \\ & < \epsilon + |V_{|\mathcal{P}^n|}(x) - \mathcal{W}_{\mathcal{P}^n}(x)|. \end{aligned}$$

Now also applying Corollary 3.5 we have

$$\begin{aligned}
& \left| \sum_{k=0}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \mathcal{W}_{\mathcal{P}_n}(G) \right| \\
&= \left| \frac{1}{2} \mathbf{v}_{|\mathcal{P}_n|}^*(x, 0) + \sum_{k=1}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \frac{1}{2} \mathbf{v}_{\mathcal{P}_n}^*(G, 0) - \frac{1}{2} \max_{H \in \mathcal{F}(G)} \mathcal{W}(H) \right| \\
&\leq \frac{1}{2} \epsilon \text{Lip}_{\mathbf{v}} + \left| \sum_{k=1}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \frac{1}{2} \max_{H \in \mathcal{F}(G)} \left\{ \frac{1}{2} \mathbf{v}_{\mathcal{P}_n}^*(H, 0) + \frac{1}{2} \max_{I_n \in \mathcal{F}(H)} \mathcal{W}(I) \right\} \right| \\
&\leq \frac{1}{2} \epsilon \text{Lip}_{\mathbf{v}} + \left| \sum_{k=1}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \frac{1}{4} \mathbf{v}_{\mathcal{P}_n}^*(H', 0) - \frac{1}{4} \max_{I \in \mathcal{F}(H')} \mathcal{W}(I) \right|
\end{aligned}$$

where H'_n is a grid element which attains the maximum of the set $\{\frac{1}{2} \mathbf{v}_{\mathcal{P}_n}^*(H, 0) + \frac{1}{2} \max_{I \in \mathcal{F}(H)} \mathcal{W}(I) \mid H \in \mathcal{F}(G)\}$. Now

$$\begin{aligned}
& \left| \sum_{k=0}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \mathcal{W}_{\mathcal{P}_n}(G) \right| \\
&\leq \frac{1}{2} \epsilon \text{Lip}_{\mathbf{v}} + \frac{1}{4} |\mathbf{v}_{|\mathcal{P}_n|}^*(x, 1) - \mathbf{v}_{\mathcal{P}_n}^*(H', 0)| + \left| \sum_{k=2}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \frac{1}{4} \max_{I \in \mathcal{F}(H')} \mathcal{W}(G) \right| \\
&\leq \frac{1}{2} \epsilon \text{Lip}_{\mathbf{v}} + \frac{1}{4} [\epsilon \text{Lip}_{\mathbf{v}} + 2\epsilon(\text{Lip}_{\mathbf{v}} + 1)] \\
&\quad + \left| \sum_{k=2}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \frac{1}{4} \max_{I \in \mathcal{F}(H')} \left\{ \frac{1}{2} \mathbf{v}_{\mathcal{P}_n}^*(I, 0) + \frac{1}{2} \max_{J \in \mathcal{F}(I)} \mathcal{W}(J) \right\} \right|
\end{aligned}$$

Repeating this estimate we obtain

$$\begin{aligned}
& \left| \sum_{k=0}^K 2^{-k-1} \mathbf{v}_{|\mathcal{P}_n|}^*(x, k) - \mathcal{W}_{\mathcal{P}_n}(G) \right| \\
&\leq \sum_{k=1}^K 2^{-k} \epsilon \text{Lip}_{\mathbf{v}} + \sum_{k=1}^K 2^{-k} k \epsilon (\text{Lip}_{\mathbf{v}} + 1) < \epsilon [\text{Lip}_{\mathbf{v}} + M(\text{Lip}_{\mathbf{v}} + 1)]
\end{aligned}$$

for all $K > 0$ a fixed constant M , since the last series converges. Hence

$$|\mathbb{V}_{\mathcal{P}} - \mathcal{W}_{\mathcal{P}_n}| < \epsilon [1 + \text{Lip}_{\mathbf{v}} + M(\text{Lip}_{\mathbf{v}} + 1)].$$

■

3.1. Algorithms. In the previous section we proved the existence of a Lyapunov function $\mathcal{W}_{\mathcal{P}}$ for an attractor-repeller pair for \mathcal{F} which is a weak Lyapunov function for f and when applied to a sequence of grids \mathcal{G}_n with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ converge to a continuous Lyapunov function for an attractor-repeller pair of f . In this section, we provide explicit algorithms for computing $\mathcal{W}_{\mathcal{P}}$.

First we need to compute the distance potential $\mathbf{v}_{\mathcal{P}}$. While in principle one could directly compute $\mathbf{v}_{\mathcal{P}}$, particularly on regular grids, it is computationally expensive to compute the distances between sets of grid elements. Therefore, the following algorithm approximates the function $\mathbf{v}_{\mathcal{P}}$ by

$$\widehat{\mathbf{v}}_{\mathcal{P}}(G) = \frac{\widehat{d}(G, \mathcal{A})}{\widehat{d}(G, \mathcal{A}) + \widehat{d}(G, \mathcal{A}^*)}$$

where $\widehat{d}(G, \mathcal{S}) = \min_{H \in \mathcal{S}} d(x_G, x_H)$. In other words, distances between sets of grid elements are computed by the minimum distance between the representative points in each set.

The next two lemmas show that the difference between the Lyapunov function $\mathcal{W}_{\mathcal{P}}$ computed using the distance potential $\mathbf{v}_{\mathcal{P}}$ and the function $\widehat{\mathcal{W}}_{\mathcal{P}}$ computed using $\widehat{\mathbf{v}}_{\mathcal{P}}$ decreases linearly with the diameter of the grid.

LEMMA 3.6. *If $\mathcal{S} \subset \mathcal{G}$ and $G \notin \mathcal{S}$, then $0 < \widehat{d}(G, \mathcal{S}) - d(x_G, |\mathcal{S}|) \leq \text{diam}(\mathcal{G})$.*

Proof: By definition, $\widehat{d}(G, \mathcal{S}) = \min_{H \in \mathcal{S}} d(x_G, x_H) \geq d(x_G, |\mathcal{S}|) = \min_{x \in |\mathcal{S}|} d(x_G, x)$. Let $y \in |\mathcal{S}|$ be a minimizer of $d(x_G, |\mathcal{S}|)$. Then there exists $H \in \mathcal{S}$ such that $y \in H$ so that $d(x_H, y) \leq \text{diam}(\mathcal{G})$. Therefore $\widehat{d}(G, \mathcal{S}) \leq d(x_G, x_H) \leq d(x_G, |\mathcal{S}|) + \text{diam}(\mathcal{G})$. ■

LEMMA 3.7. *If $\mathcal{P} = (\mathcal{A}, \mathcal{A}^*)$ is an attractor-repeller pair for \mathcal{F} with $|\mathcal{A}| \cap |\mathcal{A}^*| = \emptyset$, then there exists $C(\text{dist}(|\mathcal{A}|, |\mathcal{A}^*|)) > 0$ such that $|\widehat{\mathbf{v}}_{\mathcal{P}}(G) - \mathbf{v}_{\mathcal{P}}(G)| \leq C \text{diam}(\mathcal{G})$ and $|\widehat{\mathbf{v}}_{\mathcal{P}}^*(G, k) - \mathbf{v}_{\mathcal{P}}^*(G, k)| \leq C \text{diam}(\mathcal{G})$ for all $G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$ and $k \geq 0$. Moreover, if $\widehat{\mathcal{W}}_{\mathcal{P}}$ is the function defined by equation (3.1) with distance potential $\widehat{\mathbf{v}}_{\mathcal{P}}$, then $|\widehat{\mathcal{W}}_{\mathcal{P}}(G) - \mathcal{W}_{\mathcal{P}}(G)| \leq C \text{diam}(\mathcal{G})$ for all $G \in \mathcal{G}$.*

Proof: Using Lemma 3.7,

$$|\widehat{\mathbf{v}}_{\mathcal{P}}(G) - \mathbf{v}_{\mathcal{P}}(G)| \leq \frac{2 \text{diam}(\mathcal{G}) \text{diam}(X)}{[d(|\mathcal{A}|, |\mathcal{A}^*|)]^2}.$$

Suppose $\widehat{\mathbf{v}}_{\mathcal{P}}^*(G, k) - \mathbf{v}_{\mathcal{P}}^*(G, k) > C \text{diam}(\mathcal{G})$. Then there is a maximizer $H \in \Gamma_+^k(G)$ such that $\widehat{\mathbf{v}}_{\mathcal{P}}^*(G, k) = \widehat{\mathbf{v}}_{\mathcal{P}}^*(H, k)$ so that $\widehat{\mathbf{v}}_{\mathcal{P}}^*(H, k) - \mathbf{v}_{\mathcal{P}}^*(H, k) \geq \widehat{\mathbf{v}}_{\mathcal{P}}^*(G, k) - \mathbf{v}_{\mathcal{P}}^*(G, k) > C \text{diam}(\mathcal{G})$, which is a contradiction. A similar contradiction is reached if $\mathbf{v}_{\mathcal{P}}^*(G, k) - \widehat{\mathbf{v}}_{\mathcal{P}}^*(G, k) > C \text{diam}(\mathcal{G})$.

Finally,

$$|\widehat{\mathcal{W}}_{\mathcal{P}}(G) - \mathcal{W}_{\mathcal{P}}(G)| \leq \frac{1}{2} |\widehat{\mathbf{v}}_{\mathcal{P}}(G, 0) - \mathbf{v}_{\mathcal{P}}(G)| + \frac{1}{2} \left| \max_{H \in \mathcal{F}(G)} \widehat{\mathcal{W}}_{\mathcal{P}}(H) - \max_{H \in \mathcal{F}(G)} \mathcal{W}_{\mathcal{P}}(H) \right|.$$

Proceeding by induction we may assume that $|\widehat{\mathcal{W}}_{\mathcal{P}}(H) - \mathcal{W}_{\mathcal{P}}(H)| \leq C \text{diam}(\mathcal{G})$ for all $H \in \mathcal{F}(G)$. Then by the same argument as in the proceeding paragraph $|\max_{H \in \mathcal{F}(G)} \widehat{\mathcal{W}}_{\mathcal{P}}(H) - \max_{H \in \mathcal{F}(G)} \mathcal{W}_{\mathcal{P}}(H)| \leq C \text{diam}(\mathcal{G})$, and the result follows. ■

From the definition of $\widehat{\mathcal{W}}_{\mathcal{P}}$ and the acyclicity of the graph on $\mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$, we provide the following algorithm to compute $\widehat{\mathcal{W}}_{\mathcal{P}}$. Note that the algorithm assumes that \mathcal{A} has been compressed into a single vertex, which requires $O(a)$ time, where a is the number of elements in \mathcal{A} , since the multivalued map need not be adjusted because $\mathcal{F}(\mathcal{A}) = \mathcal{A}$. Also Calculate- $\widehat{\mathbf{v}}_{\mathcal{P}}$ is straightforward and not presented.

Calculate- $\widehat{\mathcal{W}}_{\mathcal{P}}$

1. for each $A \in \mathcal{A}$, $\widehat{\mathbf{v}}_{\mathcal{P}}(A, 0) \leftarrow 0$, $\widehat{\mathcal{W}}_{\mathcal{P}}(A) \leftarrow 0$, $\text{color}(A) \leftarrow \text{BLACK}$
2. for each $B \in \mathcal{A}^*$, $\widehat{\mathcal{W}}_{\mathcal{P}}(B) \leftarrow 1$, $\text{color}(B) \leftarrow \text{BLACK}$
3. for each $G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$, Calculate- $\widehat{\mathbf{v}}_{\mathcal{P}}(G)$, $\text{color}(G) \leftarrow \text{WHITE}$
4. $\text{preimage} \leftarrow \mathcal{A}$
5. do until $\text{preimage} = \emptyset$
6. $\text{preimage} \leftarrow \mathcal{F}^{-1}(\text{preimage})$
7. for each $G \in \text{preimage}$
8. if for each $H \in \mathcal{F}(G)$ $\text{color}(H) = \text{BLACK}$
9. $\widehat{\mathbf{v}}_{\mathcal{P}}^*(G, 0) \leftarrow \max(\{\widehat{\mathbf{v}}_{\mathcal{P}}^*(H, 0) \mid H \in \mathcal{F}(G)\} \cup \{\widehat{\mathbf{v}}_{\mathcal{P}}(G)\})$
10. $\widehat{\mathcal{W}}_{\mathcal{P}}(G) = \widehat{\mathbf{v}}_{\mathcal{P}}^*(G, 0) + \frac{1}{2} \max\{\widehat{\mathcal{W}}_{\mathcal{P}}(H) \mid H \in \mathcal{F}(G)\}$

11. $\text{color}(G) \leftarrow \text{BLACK}$

To summarize, we have given an algorithm for computing an approximation $\widehat{\mathcal{W}}_{\mathcal{P}}$ to a Lyapunov function $\mathcal{W}_{\mathcal{P}}$ for a single attractor-repeller pair \mathcal{P} for the multivalued map \mathcal{F} . The error in this approximation decreases linearly with the diameter of the grid. For any grid, $\mathcal{W}_{\mathcal{P}}$ is a weak Lyapunov function for the continuous map f . Moreover, if \mathcal{G}_n is a sequence of grids with $\text{diam}(\mathcal{G}_n) \rightarrow 0$ and \mathcal{P}_n are suitably chosen, then $\mathcal{W}_{\mathcal{P}_n}$, and hence also $\widehat{\mathcal{W}}_{\mathcal{P}_n}$, converges to a continuous Lyapunov function for an attractor-repeller pair for f .

We can extend these results by taking convex combinations of Lyapunov functions for different attractor-repeller pairs to obtain Lyapunov functions for Morse decompositions and for the recurrent set, as described in [1]. Indeed, if one takes a suitable convex combination of Lyapunov functions for all attractor-repeller pairs of \mathcal{F} , i.e. $\mathcal{W} = \sum_{\mathcal{P}} \alpha_{\mathcal{P}} \mathcal{W}_{\mathcal{P}}$ with $\sum_{\mathcal{P}} \alpha_{\mathcal{P}} = 1$, suitably ordered, then as $\text{diam}(\mathcal{G}_n) \rightarrow 0$ these Lyapunov functions for \mathcal{F} converge to a Lyapunov function for the chain recurrent set of f , [1] (Theorem 6.13).

3.2. Computational complexity. First we assume the distance potential function $\widehat{v}_{\mathcal{P}}$ has already been computed for a given attractor-repeller pair, so that Calculate- $\widehat{v}_{\mathcal{P}}$ is not done in line 3 of Calculate- $\mathcal{W}_{\mathcal{P}}$.

THEOREM 3.8. *Assume a distance potential function $\widehat{v}_{\mathcal{P}}$ has been computed for an attractor-repeller pair \mathcal{P} . Then the complexity to compute $\mathcal{W}_{\mathcal{P}}$ is*

$$O\left(a + \sum_{G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)} O_G(M_G - m_G + 2)\right),$$

where a is the number of elements in \mathcal{A} , O_G is the out-degree, the number of edges with G as their initial vertex, $M_G = \min\{k : \mathcal{F}^k(G) \subset \mathcal{A}\}$, and $m_G = \min\{k : \mathcal{F}^k(G) \cap \mathcal{A} \neq \emptyset\}$.

Proof: To compress the attractor to a single vertex requires operations on the order of the number of elements of the attractor, $O(a)$. Since the graph $\mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$ is acyclic with each $G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$ possessing a complete orbit from \mathcal{A}^* to \mathcal{A} in \mathcal{G} , the numbers $M_G, m_G > 0$ are well-defined. To calculate $\mathcal{W}_{\mathcal{P}}$ at G we must compute $v_{\mathcal{P}}^*(G, 0)$ and $\max_{H_j \in \mathcal{F}(G)} \{\mathcal{W}_{\mathcal{P}}(H_j)\}$. Both are computed when their values at all vertices in the forward image of G have been computed, by finding the maximums over these values on $\mathcal{F}(G)$. The number of comparisons in this step is O_G . Moreover, when vertex G is reached in a particular preimage, the forward image must be checked (line 8) to determine if the function values can be computed at G , which requires O_G checks. Since each vertex can be found many different times in a backward image of the attractor and the number of times that each vertex is found is bounded by $M_G - m_G + 1$, the total number of operations is bounded by

$$O\left(\sum_{G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)} O_G(M_G - m_G + 2)\right).$$

■

To calculate the distance potential function $\widehat{v}_{\mathcal{P}}(G)$ for a given $G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)$, we must compute the minimum distances between x_G and the representatives in the attractor and repeller which gives the following theorem.

THEOREM 3.9. *Given an attractor-repeller pair \mathcal{P} , the complexity to calculate the distance potential $\widehat{v}_{\mathcal{P}}$ is $O(|\mathcal{G}| - |\mathcal{A}| - |\mathcal{A}^*|)(|\mathcal{A}| + |\mathcal{A}^*|)$, which is bounded by $O(|\mathcal{G}|^2)$.*

Proof: The first part is trivial. The maximum of $(|\mathcal{G}| - |\mathcal{A}| - |\mathcal{A}^*|)(|\mathcal{A}| + |\mathcal{A}^*|)$ over $|\mathcal{A}| + |\mathcal{A}^*| \leq |\mathcal{G}|$ is $(|\mathcal{G}| - \frac{1}{2}|\mathcal{G}|)(\frac{1}{2}|\mathcal{G}|) = \frac{1}{4}|\mathcal{G}|^2$ for sufficiently large $|\mathcal{G}|$. ■

Now we summarize the overall computational complexity for the algorithm. Note that we assume that the multivalued map \mathcal{F} is given on a grid \mathcal{G} , and hence the computational cost to compute \mathcal{F} is not included in the following theorem, which is a consequence of Theorems 2.5, 3.8, and 3.9.

THEOREM 3.10. *The overall complexity to approximate the Lyapunov function is*

$$O \left(|\mathcal{F}| + |\mathcal{G}| + spl(s) + r l(r) + p |\mathcal{G}|^2 + \sum_{i=1}^p \left(a_i + \sum_{G \in \mathcal{G} \setminus (\mathcal{A} \cup \mathcal{A}^*)} [O_G(M_G^i - m_G^i + 2)] \right) \right)$$

where s is the number of recurrent components, p is the number of attractor-repeller pairs with $p \leq 2^s$, r is the cardinality of $\mathcal{F}(\mathcal{R})$, a is the cardinality of \mathcal{A} , O_G is the cardinality of $\mathcal{F}(G)$, $M_G = \min\{k : \mathcal{F}^k(G) \subset \mathcal{A}\}$, and $m_G = \min\{k : \mathcal{F}^k(G) \cap \mathcal{A} \neq \emptyset\}$.

Recall that $l(n)$ is the complexity of searching a list of n elements for the data structure used to implement the algorithm. Using binary-search trees, at worst $l(n) = O(\log(n))$, and using a good hash table $l(n)$ can be expected to be nearly constant.

4. Examples

In this section we present the results from three simple examples drawn from the one-dimensional logistic map and the two-dimensional van der Pol ODE. For the logistic map, $f(x) = \lambda x(1 - x)$, we consider two values of the parameter. At $\lambda = 3.3$, the chain recurrent set contains two unstable fixed points and a stable period-2 orbit, and at $\lambda = 3.835$, there is a stable period-3 orbit whose dual repeller is a chaotic set.

In the first case, we took a grid of intervals by dividing $[0, 1]$ into 2^{17} equal length subintervals. In this graph there were three recurrent components and two nontrivial attractor-repeller pairs. In the second case, we divided $[0, 1]$ equally into 2^{20} subintervals, and the graph contains four recurrent components and three nontrivial attractor-repeller pairs. Figure 2 shows the graphs of Lyapunov functions. The points at the absolute minima correspond to the periodic orbits.




FIGURE 2. Lyapunov function for logistic map for (a) $\lambda = 3.3$ and $n = 2^{17} = 131,072$ and (b) with $\lambda = 3.835$ and $n = 2^{20} = 1,048,576$.

We also consider the van der Pol ODE, $\dot{x} = y$ and $\dot{y} = -x + (1 - x^2)y$. It is well known that this system has a stable periodic orbit and one unstable equilibrium point at $(0, 0)$. Figure

3 shows a Lyapunov function on $[-4, 4] \times [-4, 4]$ which was subdivided seven times in each direction to obtain a rectangular grid of $2^{14} = 16,384$ elements. We used GAIO [8] to produce a multivalued map for the time- τ map with $\tau = 0.2$. The multivalued map has two components and one nontrivial attractor-repeller pair. Note that a continuous Lyapunov function for the time- τ map may not be a strict Lyapunov function for the flow. In [1], a method is given for adjusting the above approximations to be certain to obtain an approximate Lyapunov function for the flow, but we do not implement it here.

Also we note that the domains used in these examples, the interval $[0, 1]$ for logistic map and $[-4, 4] \times [-4, 4]$ for the van der Pol ODE, are not invariant and do not yield a closed directed graph. In this case, one first restricts the multivalued map to the maximal invariant set in \mathcal{G} .




FIGURE 3. Lyapunov function for the van der Pol ODE. In (b) the grayscale represents the height.

References

- [1] Kalies, W., Mischaikow, K., and VanderVorst, R., 2005. “An algorithmic approach to chain recurrence”. *Foundations of Computational Mathematics*, **5**, pp. 409–449.
- [2] Mischaikow, K., 2002. “Topological techniques for efficient rigorous computation in dynamics”. *Acta Numer.*, **11**, pp. 435–477.
- [3] Szymczak, A., 1997. “A combinatorial procedure for finding isolating neighbourhoods and index pairs”. *Proc. Roy. Soc. Edinburgh Sect. A*, **127**(5), pp. 1075–1088.
- [4] Day, S., 2003. *A Rigorous Numerical Method in Infinite Dimensions*. Ph.D. dissertation.
- [5] Day, S., Junge, O., and Mischaikow, K., 2004. “A rigorous numerical method for the global analysis of infinite-dimensional discrete dynamical systems”. *SIAM J. Appl. Dyn. Syst.*, **3**(2), pp. 117–160 (electronic).
- [6] Mrozek, M., 1999. “An algorithm approach to the Conley index theory”. *J. Dynam. Differential Equations*, **11**(4), pp. 711–734.
- [7] Dellnitz, M., Froyland, G., and Junge, O., 2001. “The algorithms behind GAIO-set oriented numerical methods for dynamical systems”. In *Ergodic theory, analysis, and efficient simulation of dynamical systems*. Springer, Berlin, pp. 145–174, 805–807.
- [8] Dellnitz, M., and Junge, O., 1998. GAIO. <http://www-math.uni-paderborn.de/~agdellnitz/gaio>.
- [9] Cormen, T. H., Leiserson, C. E., and Rivest, R. L., 1990. *Introduction to algorithms*. The MIT Electrical Engineering and Computer Science Series. MIT Press, Cambridge, MA.
- [10] Robbin, J. W., and Salamon, D. A., 1992. “Lyapunov maps, simplicial complexes and the Stone functor”. *Ergodic Theory Dynam. Systems*, **12**(1), pp. 153–183.

- [11] Conley, C., 1978. *Isolated invariant sets and the Morse index*, Vol. 38 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, R.I.
- [12] Franks, J., and Misiurewicz, M., 2002. "Topological methods in dynamics". In *Handbook of dynamical systems, Vol. 1A*. North-Holland, Amsterdam, pp. 547–598.
- [13] Robinson, C., 1999. *Dynamical systems: Stability, Symbolic Dynamics, and Chaos*, second ed. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL.