

LYAPUNOV FUNCTIONS AND ATTRACTORS IN ARBITRARY METRIC SPACES

MIKE HURLEY

(Communicated by Mary Rees)

ABSTRACT. We prove two theorems concerning Lyapunov functions on metric spaces. The new element in these theorems is the lack of a hypothesis of compactness or local compactness. The first theorem applies to a discrete dynamical system on any metric space; the result is that if A is an attractor for a continuous map g of a metric space X to itself, then there is a Lyapunov function for A . The second theorem applies only to separable metric spaces; the theorem is that there is a complete Lyapunov function for any continuously-generated discrete dynamical system on a separable metric space. (A complete Lyapunov function is a real-valued function that is constant on orbits in the chain recurrent set, is strictly decreasing along all other orbits, and separates different components of the chain recurrent set.)

1. INTRODUCTION

The main results of this paper are two theorems concerning the relationship between attractors, chain recurrence, and the existence of Lyapunov functions for a discrete dynamical system generated by a continuous self-map g of a metric space X . These theorems extend earlier results where the same conclusions were reached under the assumption that X was compact or locally compact [1], [2], [3], [4]. The first theorem is that there is a continuous Lyapunov function for any attractor A of a continuous map g of any metric space into itself. The Lyapunov function has the properties that its value is 0 in A and 1 on the complement of the basin of A , while it is strictly decreasing along other g orbits. This result was first established by C. Conley [1] for flows on compact metric spaces; Conley's result was adapted for maps on compact spaces by J. Franks [2], and was later established for maps on locally compact metric spaces in [4]. The precise statement of this theorem follows; detailed definitions can be found in the next section.

Theorem A. *Let X be an arbitrary metric space and g a continuous map from X to itself. Suppose that U is an attractor block for g and that A is the associated attractor. Then there is a Lyapunov function λ for the attractor A . This means that λ is a continuous function taking on values in $[0, 1]$, $\lambda(g(x)) \leq \lambda(x)$ for every x , $\lambda^{-1}(0) = A$, $\lambda^{-1}(1) = X - \mathcal{O}^-(U)$, and with the property that $\lambda(g(x)) < \lambda(x)$ for any $x \in \mathcal{O}^-(U) - A$. (Here $A = \bigcap_{n \geq 0} \overline{g^n(U)}$ and $\mathcal{O}^-(U) = \bigcup_{n \geq 0} g^{-n}(U)$.)*

Received by the editors May 12, 1994.

1991 *Mathematics Subject Classification.* Primary 58F12.

Key words and phrases. Attractor, Lyapunov function, chain recurrence.

The arguments used in [1], [2] to establish Theorem A when X is compact depend upon compactness in two important ways. First, the Lyapunov function is constructed in terms of an auxiliary function ϕ , where $\phi(x)$ is defined to be the maximum of the distances from the attractor A to any point along the forward orbit of x . When x is in the basin of A , compactness implies that the orbit of x will approach A , from which it can be shown that this auxiliary function is continuous at x and that, for large enough n , the value of the auxiliary function at $g^n(x)$ will be less than its value at x . This last observation is the key to showing that the Lyapunov function that is constructed will be strictly decreasing along the orbit of such an x .

When X is not compact, things are more complicated. In particular, the empty set can be an attractor, and even if an attractor A is nonempty, there may be points in its basin that do not approach the attractor under iteration by g ; examples can be found in [3], [4], or see example 1 in section 4 of this paper. When X is locally compact the only way in which this can happen is if the orbit of such a point goes off to infinity. This observation was used in [4] to construct a Lyapunov function in the locally compact case, using the same general approach as [1], [2]. Without local compactness the argument of [4] breaks down entirely.

The construction below replaces the auxiliary function used in the earlier constructions with one of another type. Here is a (not entirely accurate) outline of the idea, at least in the case that the attractor A is nonempty. Keeping the intuition that points in the basin of attraction are being drawn toward the attractor by the map, we define a type of ‘energy function’ that is an attempt to measure how hard it is to ‘swim against the current’ so as to move from A out to a point x in the basin of A . Since A is forward invariant, no orbit can start in A and reach a point that is not in A , so we consider ‘pseudo-orbits’ that do so. In this paper a pseudo-orbit for g is any finite sequence of points x_0, x_1, \dots, x_n of X , which we imagine as being generated in the following way: start at x_0 , make a jump to some other point \tilde{x}_0 , and then apply g to obtain the point $x_1 = g(\tilde{x}_0)$; repeat the process with x_1 taking the place of x_0 , and so on. This is different from the usual definition of a pseudo-orbit in that we first make a jump and then we apply g , rather than the other way around. The ‘energy’ assigned to such a pseudo-orbit is essentially the total length of all the jumps. It is relatively easy to show that this energy function is both continuous and nonincreasing along g orbits. To construct a function that is actually decreasing along orbits in the basin of A that are not in A takes a little more work, as does overcoming the difficulties that arise because A might be empty, but we will not describe that part of the construction here.

The second theorem requires that the metric space X be separable.

Theorem B. *If X is a separable metric space and g is a continuous function from X to itself, then there is a complete Lyapunov function Λ for g . This means that (i) $\Lambda: X \rightarrow [0, 1]$ is a continuous function that is nonincreasing along all g orbits, is constant along the orbit of a point in the chain recurrent set of g , and is strictly decreasing along the orbit of any other point; (ii) Λ maps the chain recurrent set into a subset of the Cantor middle-third set, taking each chain transitive component to a single value and taking distinct chain transitive components to distinct values; and (iii) if C, C' are chain transitive components and if for each $\epsilon \in \mathcal{P}$ there is an ϵ chain from C to C' then $\Lambda(C) > \Lambda(C')$.*

The complete Lyapunov function of Theorem B is constructed as a sum $\sum \lambda_n$ of the Lyapunov functions given by Theorem A for a family of attractors A_n of g . The general idea of the construction is much the same as the construction of complete Lyapunov functions in [1], [2], [4], although some technical difficulties arise because X is not assumed to be locally compact.

2. DEFINITIONS AND BACKGROUND

Standing notation. Throughout this paper g denotes a continuous map from the metric space (X, d) to itself, and \mathcal{P} denotes the set of all continuous, strictly positive, real-valued functions on X .

2.1. Attractor blocks and attractors.

Definition. A nonempty open subset U of X is an *attractor block* for g if $\overline{g(U)} \subset U$. An attractor block *determines the attractor* A , where A is defined as $A = \bigcap_{n \geq 0} \overline{g^n(U)}$. The *basin of A relative to U* is the open set consisting of the union of all inverse images of U , and is denoted as $\mathcal{O}^-(U)$. (To say that a set A is an attractor means that there exists an attractor block U with $A = \bigcap \overline{g^n(U)}$.) At times we will refer to an (attractor, attractor block) pair (A, U) , meaning that U is an attractor block that determines A as an attractor.

The definition of attractor that is given above is the one used in [4], [5], and is one way of generalizing the definition given in [1] to noncompact spaces. Even though it allows the pathologies described in the introduction (the empty set can be an attractor; there may be points in the basin of an attractor A that do not approach A under iteration by g), one measure of the appropriateness of the definition is that with it Conley's theorem relating the chain recurrent set of g to its collection of attractors holds without compactness; the theorem is briefly described in section 2.3 below.

2.2. Chains.

Definition. A *chain* in X is any finite nonempty sequence ψ of points of X ,

$$\psi = (x_0, x_1, \dots, x_n).$$

The integer $n \geq 0$ called the *length* of the chain. If $\epsilon \in \mathcal{P}$, then ψ is an ϵ -*chain* for g provided that its length n is at least 1 and that $d(g(x_j), x_{j+1}) < \epsilon(x_j)$ for every j with $0 \leq j \leq n - 1$.

The following result is established in [5] (lemma 5.2).

Lemma 2.1. *If f is a continuous function from X to itself and $\epsilon \in \mathcal{P}$, then there is a $\delta \in \mathcal{P}$ with the property that if y is in the ball of radius $\delta(x)$ that is centered at x , then $d(f(x), f(y)) < \epsilon(f(x))$.*

Lemma 2.2. *Suppose that f is a continuous map from X to itself and that $\epsilon \in \mathcal{P}$. There is an $\eta \in \mathcal{P}$ such that $\eta \leq \epsilon/2$ and with the property that if $x, y \in X$ and $d(x, y) < \eta(x)$, then $2\epsilon(g(x)) > \epsilon(g(y)) > \epsilon(g(x))/2$.*

Proof. The second inequality is lemma 5.3 of [5], and the proof of the first inequality is essentially the same. \square

Lemma 2.3. *If Y, Z are disjoint nonempty closed subsets of X , then there is an $\epsilon \in \mathcal{P}$ with the property that if $y \in Y$, then the $\epsilon(y)$ ball centered at y does not intersect Z . In particular, if U is an attractor block for g , then there is an $\epsilon \in \mathcal{P}$ such that the $\epsilon(y)$ ball centered at y does not meet $X - U$ for any $y \in \overline{g(U)}$.*

Proof. $\epsilon(x) = [d(x, Y) + d(x, Z)]/2$ is such a function. \square

The next lemma provides a connection between attractors and ϵ -chains.

Lemma 2.4. *If U is an attractor block, then there is an $\epsilon \in \mathcal{P}$ with the property that any ϵ -chain beginning in U must also end in U .*

Proof. If $U = X$ there is nothing to prove. Otherwise the sets $X - U$ and $\overline{g(U)}$ are closed, nonempty and disjoint. Using these closed sets in the lemma gives us the continuous function $\epsilon \in \mathcal{P}$, and a trivial induction shows that any ϵ -chain beginning in U also ends in U . \square

2.3. Chain recurrence.

Definition. A point $x \in X$ is said to be *chain recurrent* for g if for each $\epsilon \in \mathcal{P}$ there is an ϵ -chain for g that both begins and ends at x (recall that an ϵ -chain must consist of at least two points). The *chain recurrent set* of g is the set of all its chain recurrent points. The chain recurrent set is closed and forward invariant.

Conley's theorem describes a fundamental connection between the chain recurrent set of g and the collection of attractors of g . The theorem states that if x is chain recurrent and if U is an attractor block determining an attractor A , then either $x \in A$ or else $x \notin \mathcal{O}^-(U)$, and conversely, if x is not chain recurrent, then there is an (attractor, attractor block) pair (A, U) with $x \in \mathcal{O}^-(U) - A$. Conley's original result applies to flows on compact spaces [1]; it was extended to maps on any metric space in [5].

Definition. There is a natural equivalence relation on the chain recurrent set, defined as follows: if x, y are chain recurrent, say $x \sim y$ if and only if for each $\epsilon \in \mathcal{P}$ there is an ϵ -chain from x to y and one from y to x . The equivalence classes are called the *chain transitive components* of g .

2.4. Lyapunov functions. If (A, U) is an (attractor, attractor block) pair, then there is an associated three-set partition of X ,

$$X = A \cup [\mathcal{O}^-(U) - A] \cup [X - \mathcal{O}^-(U)].$$

Heuristically, one might describe these three sets as the attractor (A), its 'complementary repeller' ($X - \mathcal{O}^-(U)$), and the 'connecting orbits' ($\mathcal{O}^-(U) - A$). A Lyapunov function for A captures this intuition in a precise way.

Definition. A *Lyapunov function* for the (attractor, attractor block) pair (A, U) is a continuous map $\lambda: X \rightarrow [0, 1]$ with the properties that $\lambda^{-1}(0) = A$, $\lambda^{-1}(1) = X - \mathcal{O}^-(U)$, and $\lambda(g(x)) < \lambda(x)$ whenever $x \in \mathcal{O}^-(U) - A$.

Note that Conley's theorem can be rephrased in terms of the three-set partition by saying that the points that are not chain recurrent are those that lie on a connecting orbit from some (attractor, attractor block) pair. The connection between attractors and chain recurrence described in section 2.3 will be exploited in section 6 to construct a 'complete Lyapunov function' for the chain recurrent set of g from the Lyapunov functions for various attractors A of g .

Definition. A *complete Lyapunov function* for g is a continuous map $\Lambda: X \rightarrow [0, 1]$ with the following properties:

- (1) $\Lambda(g(x)) = \Lambda(x)$ whenever x is chain recurrent, and otherwise $\Lambda(g(x)) < \Lambda(x)$.
- (2) Λ takes distinct chain transitive components of g to distinct points of a nowhere dense subset of $[0, 1]$.
- (3) If C, C' are distinct chain transitive components of g that have the property that for each $\epsilon \in \mathcal{P}$ there is an ϵ -chain from C to C' , then $\Lambda(C) > \Lambda(C')$.

3. HOPSCOTCH

The definition of an ϵ -chain that we gave in the last section is that it is a sequence of points obtained by repeating the steps “apply g to the current point, then make a small jump to reach the next point”. It will be useful to consider instead sequences obtained by first making small jumps, and then applying g .

Definition. A *J-chain* for g is any finite nonempty sequence of ordered pairs of points of X , $\psi = ((x_0, \tilde{x}_0), (x_1, \tilde{x}_1), \dots, (x_n, \tilde{x}_n))$, with the property that $g(\tilde{x}_j) = x_{j+1}$ whenever $0 \leq j < n$. The integer $n \geq 0$ is called the *length* of the chain, which is said to *begin* at x_0 and *end* at \tilde{x}_n . If $\epsilon \in \mathcal{P}$, then ψ is said to be an ϵ -*J-chain* for g if $d(x_j, \tilde{x}_j) < \epsilon(x_j)$ for every j .

The chain recurrent set and the chain transitive components would be unaffected if J -chains were used in place of chains in the definitions of those sets:

Lemma 3.1. *Suppose that $\epsilon \in \mathcal{P}$. Then there is a $\delta \in \mathcal{P}$ with the property that if $((x_j, \tilde{x}_j))$ is a δ - J -chain for g , then (x_j) is an ϵ -chain for g . Conversely, if (x_j) is an ϵ -chain for g , then $((x_0, x_0), (g(x_0), x_1), (g(x_1), x_2), \dots, (g(x_{n-1}), x_n))$ is an ϵ - J -chain for g .*

Proof. Lemma 2.1 gives a positive function δ such that $d(x, y) < \delta(x)$ implies that $d(g(x), g(y)) < \epsilon(g(x))$, which is all that is required for the first assertion. The second assertion is evident. □

Corollary. *Given $\epsilon \in \mathcal{P}$, there is a $\delta \in \mathcal{P}$ with the property that if there is a δ - J -chain going from a to b , then there is an ϵ -chain going from a to b . Conversely, if there is an ϵ -chain going from a to b , then there is an ϵ - J -chain going from a to b .*

Proof. The second statement is an obvious consequence of the second part of the lemma, and the first statement follows if we use the lemma to select a δ corresponding to $\epsilon/2$ and note that without loss of generality we may assume that $\delta \leq \epsilon/2$. □

4. ENERGY FUNCTIONS FROM J-CHAINS

In this section we will outline the construction of ‘energy’ functions that will be used later to measure the dynamical distance from a point of X to the attractor A . Things are complicated by the fact that when X is not compact an attractor may not have all of the properties that one might hope for. As an example, A can be empty (an example is the map $x \rightarrow x^2 + 1$ on R with $U = (-\infty, \infty)$). Even if A is nonempty there may well be points in the attractor block U whose forward orbits do not approach A , as the following example on R^2 shows.

Example 1. Consider the map f on R^2 defined by $f(x, y) = (x + x^2 + y^2, y)$, which has a fixed point at the origin and maps any other point horizontally to the right. If U is the open half-plane $\{(x, y) \mid x > -1\}$ then clearly $\overline{f(U)} \subset U$, so that U is an

attractor block and the attractor determined by U is the positive x -semiaxis. All of the points in the right half-plane are in U , but none are ever mapped any closer to A than they were at the beginning.

To avoid difficulties caused by this sort of pathology we will define the energy functions in terms of the sets $\overline{g^k(U)}$, $k \geq 1$, instead of in terms of A , thinking of these sets as approximations to A . The energy function at x will measure the size of the jumps needed to construct a J -chain that begins in $\overline{g^k(U)}$ and ends at x .

Definition. Given a function $\delta \in \mathcal{P}$, define an operator Γ_δ on the set of all J -chains as follows. Let $\psi = ((x_0, \tilde{x}_0), (x_1, \tilde{x}_1), \dots, (x_n, \tilde{x}_n))$ be a J -chain, and define

$$\Gamma_\delta(\psi) = \sum_{j=0}^n \frac{d(x_j, \tilde{x}_j)}{\delta(x_j)}.$$

(Although the definition of a J -chain makes reference to a continuous map f , namely the requirement that $f(\tilde{x}_j) = x_{j+1}$, the definition of Γ_δ is independent of f .)

If Y is a nonempty closed subset of X and f is any continuous map from X to itself, let $\mathcal{Z}(x; Y, f)$ to be the set of all J -chains for f that begin in Y and end at x , and define

$$L_\delta(x; Y, f) \equiv \inf_{\psi \in \mathcal{Z}(x; Y, f)} \Gamma_\delta(\psi).$$

Remark 1. If $x \in Y$, then $((x, x)) \in \mathcal{Z}(x; Y, f)$, so that $L_\delta(x; Y, f) = 0$.

Lemma 4.1. L_δ is nonincreasing along f orbits: $L_\delta(x; Y, f) \geq L_\delta(f(x); Y, f)$ for every $x \in X$.

Proof. Given $\psi = ((x_0, \tilde{x}_0), \dots, (x_n, x)) \in \mathcal{Z}(x; Y, f)$, define $\psi' \in \mathcal{Z}(f(x); Y, f)$ by appending $(f(x), f(x))$ to the end of ψ . Clearly $\Gamma_\delta(\psi') = \Gamma_\delta(\psi)$, so that taking the infimum over all such ψ gives the desired inequality:

$$L_\delta(f(x); Y, f) \leq \inf_{\psi \in \mathcal{Z}(x; Y, f)} \Gamma_\delta(\psi') = \inf_{\psi \in \mathcal{Z}(x; Y, f)} \Gamma_\delta(\psi) = L_\delta(x; Y, f).$$

□

Lemma 4.2. The map $x \rightarrow L_\delta(x; Y, f)$ is continuous.

Proof. Choose a point $x \in X$; there are two cases depending upon whether $x \in Y$ or not. In the first case $x \in Y$, so $L_\delta(x; Y, f) = 0$. Given $y \in X$, let ψ be the J -chain $((x, y)) \in \mathcal{Z}(y; Y, f)$, so that

$$L_\delta(y; Y, f) \leq \Gamma_\delta(\psi) = d(x, y)/\delta(x),$$

which clearly goes to 0 as $y \rightarrow x$, establishing the lemma in this case.

In the remaining case $x \notin Y$. Given $\psi \in \mathcal{Z}(x; Y, f)$, define $\psi' \in \mathcal{Z}(y; Y, f)$ by changing the last pair (z, x) of ψ to (z, y) . Then

$$|\Gamma_\delta(\psi) - \Gamma_\delta(\psi')| = \left| \frac{d(z, x)}{\delta(z)} - \frac{d(z, y)}{\delta(z)} \right| \leq \frac{d(x, y)}{\delta(z)},$$

which goes to 0 as y approaches x , but at a rate depending upon ψ . Given any positive constant ϵ , there is a $\psi \in \mathcal{Z}(x; Y, f)$ satisfying $\Gamma_\delta(\psi) < L_\delta(x; Y, f) + \epsilon$. Using this ψ in the calculation above shows that $L_\delta(y; Y, f) < L_\delta(x; Y, f) + \epsilon$ for all y sufficiently close to x , which shows that $L_\delta(\cdot; Y, f)$ is upper semicontinuous at x . Now the only way that $L_\delta(\cdot; Y, f)$ could fail to be continuous at x is if there is a sequence y_n converging to x and a constant $\beta > 0$ with $L_\delta(y_n; Y, f) <$

$L_\delta(x; Y, f) - \beta$. In this situation there would be J -chains $\psi'_n \in \mathcal{Z}(y_n; Y, f)$ with $\Gamma_\delta(\psi'_n) < L_\delta(x; Y, f) - \beta$. Let $\psi_n \in \mathcal{Z}(x; Y, f)$ be associated with ψ'_n as above (change the y_n at the end of ψ'_n to an x). Using the calculation above with y_n taking the place of y gives

$$|\Gamma_\delta(\psi_n) - \Gamma_\delta(\psi'_n)| \leq d(x, y_n)/\delta(z_n)$$

where (z_n, y_n) is the last pair in ψ'_n . If we can show that $d(x, y_n)/\delta(z_n)$ gets arbitrarily close to 0, then we will be done, since there would then be an n such that $\Gamma_\delta(\psi_n) < L_\delta(x; Y, f)$, which is absurd.

We know that

$$L_\delta(x; Y, f) - \beta > \Gamma_\delta(\psi'_n) \geq d(z_n, y_n)/\delta(z_n)$$

(the fraction is the last of the nonnegative terms whose sum is $\Gamma_\delta(\psi'_n)$), so that $\delta(z_n) \geq Md(z_n, y_n)$, where M is a positive constant that is independent of n . It follows that

$$d(x, y_n)/\delta(z_n) \leq d(x, y_n)/Md(z_n, y_n),$$

which will go to 0 as we want, unless $d(z_n, y_n)$ goes to 0. If this happens, then it follows that $z_n \rightarrow x$, which allows us to conclude from the continuity of δ that $\delta(z_n) > \delta(x)/2$ for all large n . This means that $\delta(z_n)$ is bounded away from 0, so that $d(x, y_n)/\delta(z_n)$ goes to 0 as required. \square

Note: The reason for using J -chains instead of regular chains is to ensure that L_δ is continuous at points of Y . An analogous construction with chains can be made, provided that one is considering chains that go from x to Y , instead of the other way around as we have done above. If one defines another energy function in this case by setting $B_\delta(x; Y, f)$ equal to the minimum over all chains from x to Y of the sum of all the terms $d(f(x_j), x_{j+1})/\delta(f(x_j))$, then essentially the same arguments as above show that B_δ is continuous and is *nondecreasing* along f orbits.

5. LYAPUNOV FUNCTIONS FOR ATTRACTORS

As before, assume that U is an attractor block for g , and that A is the attractor determined by U . By lemma 2.4 there is a positive function $\epsilon \in \mathcal{P}$ with the property that if $x \in \overline{g(U)}$ then the ball centered at x with radius $\epsilon(x)$ is contained in U . Without loss of generality we can assume that ϵ is globally bounded above by 1. We begin by defining a sequence of energy functions for g .

Definition. For each $k \geq 1$ define $L'_k(x) = L_\epsilon(x; \overline{g^k(U)}, g^k)$.

Remark 2. By the results of the last section, each L'_k is continuous, nonnegative, and $L'_k(g^k(x)) \leq L'_k(x)$ for all x .

Remark 3. If $x \in A$ then $x \in \overline{g^k(U)}$ for every k , so that $L'_k(x) = 0$ for every k (by Remark 1 from the last section).

Remark 4. If $\psi = ((x_j, \tilde{x}_j))$ is an ϵ - J -chain for g^k with $x_0 \in \overline{g^k(U)}$, then $\tilde{x}_0 \in U$ and $x_1 = g^k(\tilde{x}_0) \in \overline{g^k(U)}$, and so by induction ψ ends in U . (This is the J -chain analogue of lemma 2.4.)

Lemma 5.1. *If $x \notin U$, then $L'_k(x) \geq 1$ for every k .*

Proof. By remark 4, if $\psi \in \mathcal{Z}(x; \overline{g^k(U)}, g^k)$ then ψ is not an ϵ - J -chain, so there is a value of j with $d(x_j, \tilde{x}_j) \geq \epsilon(x_j)$, so that $\Gamma_\epsilon(\psi) \geq d(x_j, \tilde{x}_j)/\epsilon(x_j) \geq 1$. Since this inequality holds for all $\psi \in \mathcal{Z}(x; \overline{g^k(U)}, g^k)$, the lemma follows. \square

Lemma 5.2. $(L'_k)^{-1}(0) = \overline{g^k(U)}$.

Proof. By remark 1 (following the definition of L_δ in the last section) it is enough to show that $L'_k(x) > 0$ whenever $x \notin \overline{g^k(U)}$. If there is no ϵ - J -chain for g^k from $\overline{g^k(U)}$ to x , then the proof of the last lemma shows that $L'_k(x) \geq 1$, so we can assume that there are ϵ - J -chains for g^k from $\overline{g^k(U)}$ to x , and it is enough to deal with those J -chains. By remark 4, if $\psi = ((x_j, \tilde{x}_j))$ is an ϵ - J -chain from $\overline{g^k(U)}$ to x , then $x_j \in \overline{g^k(U)}$ for every j . In particular, if (x_n, x) is the last pair in ψ , then $x_n \in \overline{g^k(U)}$, so that

$$\Gamma_\epsilon(\psi) \geq \frac{d(x_n, x)}{\epsilon(x_n)} \geq d(x_n, x) \geq \text{dist}(\overline{g^k(U)}, x) > 0.$$

\square

Definition. For each $k \geq 1$ define

$$L''_k(x) = \frac{1}{k} \sum_{j=0}^{k-1} L'_k(g^j(x)).$$

Clearly L''_k inherits the properties of being nonnegative and continuous from L'_k .

Lemma 5.3. L''_k is nonincreasing along g orbits.

Proof. $L''_k(g(x)) - L''_k(x) = [L'_k(g^k(x)) - L'_k(x)]/k \leq 0$ by remark 2. \square

Now we are in position to define our Lyapunov function λ .

Definition. Let $L_k(x) = \min\{L''_k(x), 1\}$, so that L_k is continuous, nonincreasing along g orbits, $0 \leq L_k(x) \leq 1$ for all x , and $(L_k)^{-1}(0) = \overline{g^k(U)}$. Now define

$$L(x) = \sum_{k=1}^{\infty} \frac{L_k(x)}{2^k} \quad \text{and} \quad \lambda(x) = \sum_{j=0}^{\infty} \frac{L(g^j(x))}{2^{j+1}}.$$

The infinite sums in the definitions of L and λ are uniformly convergent, so these functions are continuous. In addition these functions are nonincreasing along orbits of g , and their values are in the interval $[0, 1]$.

Remark 5. If $\lambda(x) = 0$, then $L(x) = 0$, so that $L_k(x) = 0$ for every k , and thus $x \in \bigcap \overline{g^k(U)} = A$. The opposite implication is clear, so $\lambda^{-1}(0) = A$.

Lemma 5.4. If $x \in \mathcal{O}^-(U) - A$, then $\lambda(g(x)) < \lambda(x)$.

Proof. Since each term in the series defining λ is no larger at $g(x)$ than it is at x , it is enough to show that there is one of these terms that is actually smaller at $g(x)$ than at x . We will prove the lemma in two stages. In the first stage assume that $x \in \overline{g(U)} - A$. It follows that there is a smallest integer k with the property that $x \notin \overline{g^k(U)}$, and that $k \geq 2$. Since $x \in \overline{g^j(U)}$ for $j < k$, $L_j(x) = 0$ for all such j . Since $x \notin \overline{g^k(U)}$, lemma 5.2 shows that $L_k(x) > 0$. Now $g(x) \in \overline{g^j(U)}$ for $j \leq k$, and in particular $L_k(g(x)) = 0 < L_k(x)$. Thus $L(g(x)) < L(x)$, which completes the proof in this case.

In the remaining case $x \in \mathcal{O}^-(U) - g(U)$. In this situation there is a $j \geq 1$ with $g^j(x) \in g(U) - A$, so by the previous case $L(g^j(x)) > L(g^{j+1}(x)) = L(g^j(g(x)))$, which says that the $(j + 1)^{st}$ term in the sum for $\lambda(x)$ is strictly larger than the corresponding term in the sum for $\lambda(g(x))$, which finishes the proof. \square

Lemma 5.5. $\lambda^{-1}(1) = X - \mathcal{O}^-(U)$.

Proof. If $x \notin \mathcal{O}^-(U)$, then no point on the forward orbit of x lies in U , and so by lemma 5.1 $L_k(g^j(x)) = 1$ for every k and j , implying that $\lambda(x) = 1$. Conversely, if $\lambda(x) = 1$, then $L(g^j(x)) = 1$ for every j , and so $L_k(g^j(x)) = 1$ for every j and k . If $x \in \mathcal{O}^-(U)$, then there are choices of $j \geq 0$ and $k \geq 1$ with $g^j(x) \in \overline{g^k(U)}$, implying that $L_k(g^j(x)) = 0$, so that $\lambda(x) < 1$. \square

We have completed the proof of Theorem A.

Theorem 5.6 (Theorem A). *Let X be an arbitrary metric space and g a continuous map from X to itself. Suppose that U is an attractor block for g and that A is the associated attractor. The function λ defined above is a Lyapunov function for the attractor A .*

6. COMPLETE LYAPUNOV FUNCTIONS FOR SEPARABLE METRIC SPACES

This section is devoted to the proof of Theorem B. We begin with some technical lemmas that are necessary because we will be dealing with chains defined in terms of positive, continuous functions rather than positive constants.

Lemma 6.1. *If $\epsilon \in \mathcal{P}$, then there is a function $\delta_1 \in \mathcal{P}$ with the property that if (x_0, x_1, \dots, x_n) is an ϵ -chain for g and if p is some point of X satisfying $d(x_0, p) < \delta_1(x_0)$, then (p, x_1, \dots, x_n) is a 3ϵ -chain for g .*

Proof. First choose $\eta \in \mathcal{P}$, $\eta \leq \epsilon/2$, corresponding to ϵ as in lemma 2.2. Then apply lemma 2.1 to η to obtain $\delta \equiv \delta_1 \in \mathcal{P}$ with the property that $d(g(x), g(y)) < \eta(g(x))$ whenever $d(x, y) < \delta(x)$. Now if $d(x_0, p) < \delta_1(x_0)$, then $d(g(x_0), g(p)) < \eta(g(x_0))$, so by lemma 2.2 we conclude that $\epsilon(g(p)) > \epsilon(g(x_0))/2$. Finally,

$$\begin{aligned} d(x_1, g(p)) &\leq d(x_1, g(x_0)) + d(g(x_0), g(p)) \\ &< \epsilon(g(x_0)) + \eta(g(x_0)) \leq 3 \cdot \epsilon(g(x_0))/2, \end{aligned}$$

and this last value is less than $3\epsilon(g(p))$, which is all that is required. \square

Lemma 6.2. *Given $\epsilon \in \mathcal{P}$, there is $\delta_2 \in \mathcal{P}$ with $\delta_2 \leq \epsilon$ and with the property that if $(x_0, x_1, \dots, x_n, x_{n+1})$ is a δ_2 -chain for g and if $q \in X$ satisfies $d(x_{n+1}, q) < \epsilon(x_{n+1})$, then $(x_0, x_1, \dots, x_n, q)$ is a 3ϵ -chain for g .*

Proof. Taking f as the identity function, let $\eta \in \mathcal{P}$ correspond to ϵ as in lemma 2.2, and set $\delta_2 = \eta$. Since $d(g(x_n), x_{n+1}) < \delta_2(g(x_n))$, lemma 2.2 allows us to conclude that $\epsilon(x_{n+1}) < 2\epsilon(g(x_n))$. Now

$$\begin{aligned} d(g(x_n), q) &\leq d(g(x_n), x_{n+1}) + d(x_{n+1}, q) < \delta_2(g(x_n)) + \epsilon(x_{n+1}) \\ &\leq \epsilon(g(x_n)) + 2\epsilon(g(x_n)) = 3\epsilon(g(x_n)), \end{aligned}$$

which completes the proof. \square

Corollary. *For each $\epsilon \in \mathcal{P}$ there is a $\gamma \in \mathcal{P}$ with $\gamma \leq \epsilon$ and with the property that if $(x_0, x_1, \dots, x_n, x_{n+1})$ is a γ -chain for g of length at least 2 (so $n \geq 1$), then (p, x_1, \dots, x_n, q) is a 3ϵ -chain for g whenever both $d(p, x_0) < \gamma(x_0)$ and $d(q, x_{n+1}) < \gamma(x_{n+1})$.*

Proof. This follows directly from the two lemmas when γ is defined to be the minimum of δ_1 and δ_2 . \square

Proposition 6.3. *Suppose that p, q are points of X and $\epsilon \in \mathcal{P}$ is such that there is no 3ϵ -chain for g of length at least 2 from p to q . There are neighborhoods N_1 of p and N_2 of q and an $\eta \in \mathcal{P}$ with the property that no η -chain for g can begin in N_1 and end in N_2 .*

Proof. Let γ be as in the corollary and define $\eta = \gamma/2$. Let N_1 be a neighborhood of p that is contained in the $\eta(p)$ -ball about p and small enough that $\gamma(x) > \gamma(p)/2 = \eta(p)$ for any $x \in N_1$. Let N_2 be a neighborhood of q with the analogous properties. If there were an η -chain (v, x_1, \dots, x_n, w) with $v \in N_1$ and $w \in N_2$, then, because $v \in N_1$, $d(p, v) < \gamma(p)/2 < \gamma(v)$. Similarly we conclude that $d(q, w) < \gamma(w)$. But now the corollary would give the impossible conclusion that (p, x_1, \dots, x_n, q) is a 3ϵ -chain from p to q , and so the proposition is established. \square

Proposition 6.4. *Suppose that X is a separable metric space and that $g: X \rightarrow X$ is continuous. There is a countable subset $\mathcal{E} \subset \mathcal{P}$ with the property that if v, w are points of X and if for some $\delta \in \mathcal{P}$ there does not exist a δ -chain of length at least 2 from v to w , then there is an $\alpha \in \mathcal{E}$ such that there is no α -chain of length at least 2 from v to w .*

Proof. Define $G = \{(x, y) \mid \text{there is no } \epsilon\text{-chain from } x \text{ to } y \text{ for some } \epsilon \in \mathcal{P}\}$. By proposition 6.3, if $(x, y) \in G$ then there are open subsets of X , $N(x)$ and $N(y)$, containing x and y respectively, and $\eta \in \mathcal{P}$ such that no η -chain starts in $N(x)$ and ends in $N(y)$. In particular, the open subset $N_{xy} \equiv N(x) \times N(y)$ of $X \times X$ is a neighborhood of (x, y) in G . The collection of all such N_{xy} is an open cover of G . Since X is separable and G is open, G is also separable, and so there is a countable subcover $N_j \equiv N(x_j) \times N(y_j)$ of G . Let η_j be the element of \mathcal{P} corresponding to N_j as in the last proposition, and let \mathcal{E} be the set of all the η_j 's. \square

Theorem 6.5 (Theorem B). *If X is a separable metric space and $g: X \rightarrow X$ is continuous, then there is a complete Lyapunov function Λ for g .*

To prove the theorem, first we construct the function Λ and then we prove several propositions establishing that Λ has the properties required of a complete Lyapunov function as given in the definition in section 2.4. Let Z be a countable dense subset of X . For each $z \in Z$ and $\epsilon \in \mathcal{E}$ let

$$U(z, \epsilon) = \{x \in X \mid \text{there is an } \epsilon\text{-chain of length at least 2 from } z \text{ to } x \}.$$

It is not hard to check that each of these sets $U(z, \epsilon)$ is open, nonempty, and an attractor block [3], [4], [5]. If \mathcal{U} denotes the collection of all these attractor blocks, then \mathcal{U} is countable, and so can be enumerated as $\mathcal{U} = \{U_n \mid n \geq 1\}$. For each n let A_n be the attractor that is determined by U_n , and let $\lambda_n: X \rightarrow [0, 1]$ be the associated Lyapunov function given by Theorem A. Define the complete Lyapunov function by

$$\Lambda(x) = \sum_{n=1}^{\infty} \frac{2\lambda_n(x)}{3^n}.$$

It is clear that Λ is continuous, nonincreasing along orbits, and that $0 \leq \Lambda(x) \leq 1$. It remains to establish that Λ has the other properties required of a complete Lyapunov function.

Lemma 6.6. *If C is a chain transitive component for g and U is an attractor block for g that determines A as an attractor, then either $C \subset A$ or else $C \subset X - \mathcal{O}^-(U)$.*

Proof. This is an immediate consequence of the existence of a function $\epsilon \in \mathcal{P}$ with the property that any ϵ -chain that begins in U must also end in U . (See lemma 2.4.) □

Corollary. *For each n , $\lambda_n(C)$ is either 0 or 1. In particular, $\Lambda(C)$ is a point of the usual Cantor middle-third set.*

Proposition 6.7. *Let Z, \mathcal{E} be as in the definition of $\mathcal{U} = \{U_k\}$, and let G be as in the proof of proposition 6.4. Suppose that $(x, y) \in G$.*

- (1) *There are an $\eta \in \mathcal{E}$ and $z \in Z$ such that $y \notin U(z, \eta)$ and $x \in \mathcal{O}^-(U(z, \eta))$.*

For the rest of the propositions let $U_n \equiv U(z, \eta)$ be in the enumeration of \mathcal{U} and let λ_n be the associated Lyapunov function given by Theorem A.

- (2) *If y is a chain recurrent point of g , then $\lambda_n(x) < \lambda_n(y)$.*
- (3) *If $y = x$, then $\lambda_n(g(x)) < \lambda_n(x)$.*
- (4) *If x is not chain recurrent, then $\Lambda(g(x)) < \Lambda(x)$.*
- (5) *If x, y are both chain recurrent but are not in the same chain transitive component, then $\Lambda(x) \neq \Lambda(y)$.*
- (6) *If C, C' are distinct chain transitive components, and if for each $\epsilon \in \mathcal{P}$ there is an ϵ -chain beginning in C and ending in C' , then $\Lambda(C) > \Lambda(C')$.*

Proof. As in the proof of proposition 6.4, the assumption that $(x, y) \in G$ means that $(x, y) \in N_j \equiv N(a) \times N(b)$, where N_j is one of the elements of the countable subcover of G constructed in the proof of that proposition. $N(a)$ is an open subset of X , so $Z \cap N(a)$ is dense in $N(a)$. For any $z \in Z \cap N(a)$ proposition 6.4 shows that there is no η_j -chain from z to y , so $y \notin U(z, \eta_j) \equiv U_n$. On the other hand, if the point z is chosen sufficiently close to x , then there will certainly be an η_j -chain of length 2 from z to $g^2(x)$, showing that $g^2(x) \in U_n$, and so $x \in \mathcal{O}^-(U_n)$. Setting $\eta = \eta_j$, this finishes the proof of part 1. If y is chain recurrent, then combining part 1 with the preceding lemma shows that $y \notin \mathcal{O}^-(U)$, so that $\lambda_n(y) = 1$, while $\lambda_n(x) < 1$ by Theorem A; this establishes part 2. To prove part 3, note that if $y = x$ then $x \in \mathcal{O}^-(U_n) - U_n$, so that $\lambda_n(g(x)) < \lambda_n(x)$ by Theorem A. Part 4 follows because the conditions that $(x, x) \in G$ and that x is not chain recurrent are equivalent, and Λ will decrease along an orbit anytime that one of the λ_k 's does. To establish part 5, note that if x, y lie in distinct chain transitive components then one of (x, y) or (y, x) is in G ; for definiteness assume that $(x, y) \in G$. By part 2, $\lambda_n(x) < \lambda_n(y)$, so by the corollary to the previous lemma, $\lambda_n(x) = 0$ and $\lambda_n(y) = 1$, which is sufficient to show that $\Lambda(x)$ and $\Lambda(y)$ are distinct points of the Cantor middle-third set. The assumption in part 6 shows that if C is contained in some attractor A , then so is C' (if $C' \not\subset A$ then $C' \subset X - \mathcal{O}^-(U)$, where U is an attractor block determining A , and for some $\epsilon \in \mathcal{P}$ there is no ϵ -chain from A to $X - \mathcal{O}^-(U)$). In other words, if $\lambda_n(C) = 0$ then $\lambda_n(C') = 0$, implying that $\lambda_n(C) \geq \lambda_n(C')$ for all n , so that $\Lambda(C) \geq \Lambda(C')$, which in view of part 5 is enough to establish part 6. □

This finishes the proof of Theorem B.

REFERENCES

- [1] C. Conley, *Isolated Invariant Sets and the Morse Index*, C.B.M.S. publication #38, A.M.S., Providence, 1978. MR **80c**:58009
- [2] J. Franks, A variation on the Poincaré-Birkhoff theorem, in *Hamiltonian Dynamical Systems*, volume 81 of *Contemporary Mathematics*. (A.M.S., Providence, 1988), 111-117. MR **90e**:58095
- [3] M. Hurley, Chain recurrence and attraction in non-compact spaces, *Ergodic Theory & Dynamical Systems*, **11** (1991), 709-729. MR **83b**:58096
- [4] _____, Noncompact chain recurrence and attraction, *Proc. of the A.M.S.* **115** (1992), 1139-1148. MR **92j**:58063
- [5] _____, Chain recurrence, semiflows, and gradients, *J. Dynamics & Differential Equations* **7** (1995), 437-456. MR **96k**:58190

DEPARTMENT OF MATHEMATICS, CASE WESTERN RESERVE UNIVERSITY, CLEVELAND, OHIO
44106-7058

E-mail address: mgh3@po.cwru.edu