CHAIN RECURRENCE RATES AND TOPOLOGICAL ENTROPY

DAVID RICHESON AND JIM WISEMAN

ABSTRACT. We investigate the properties of chain recurrent, chain transitive, and chain mixing maps (generalizations of the well-known notions of non-wandering, topologically transitive, and topologically mixing maps). We describe the structure of chain transitive maps. These notions of recurrence are defined using ε-chains, and the minimal lengths of these ε-chains give a way to measure recurrence time (chain recurrence and chain mixing times). We give upper and lower bounds for these recurrence times and relate the chain mixing time to topological entropy.

1. INTRODUCTION

Pseudo-orbits, or ε-chains, are important tools for investigating properties of discrete dynamical systems. These ε-chains detect recurrent and mixing behaviors that may not be evident by studying actual orbits. However, what is missing from the definitions of terms such as chain recurrence, chain transitivity, and chain mixing is any information about the lengths of the recurrence or mixing times. They also do not give fine detail about the "dynamics" of the ε-chains.

Let X be a compact metric space and $f : X \to X$ be a continuous map. The map f is chain recurrent if for every $\varepsilon > 0$ and every point x, there is an ε-chain from x to itself (see section 2 for definitions). The chain recurrence time (that is, the length of the shortest such chain) depends on ε. Similarly, if f is chain mixing, then the chain mixing time also depends on ε. We discuss the properties and relationship of these two quantities, and in particular relate the chain mixing time to topological entropy.

We also study the structure of chain transitive maps, which resembles that of the symbolic dynamics determined by irreducible graphs and is rather different from that of topologically transitive maps. For example, an irrational rotation of a circle is topologically transitive (actually minimal), but certainly not topologically mixing. However, as we will see (Cor. 14), a map on a connected space is chain mixing if

1

and only if it is chain transitive (and, in fact, if and only if it is chain recurrent).

The paper is organized as follows. In section 2 we give definitions. In section 3 we discuss the structure of chain transitive maps. In section 4 we explore the properties of the chain recurrence and mixing times, and in section 5 we relate the chain mixing time to topological entropy. In section 6 we study examples.

2. Definitions

Let X be a compact metric space. (To avoid technical difficulties later, we assume that X is not a single point.) Throughout the paper, let $f : X \to X$ be a continuous map. An *(ε, f)-chain (or (ε, f)-pseudo-orbit)* from x to y is a sequence $(x = x_0, x_1, \ldots, x_n = y)$ such that $d(f(x_{i-1}), x_i) \leq \varepsilon$ for $i = 1, \ldots, n$. To simplify notation, in most cases we write simply "ε-chain" without referring explicitly to the map f. We define the *length* of the ε-chain $(x_0, x_1, \ldots, x_n)$ to be n.

There are a variety of ways to use ε-chains to describe recurrence. In this paper we investigate three such notions—chain recurrence, chain transitivity, and chain mixing. They generalize the well-known topological notions of recurrence of non-wandering, topological transitivity, and topological mixing.

For the sake of comparison, let us recall the definitions of these topological notions of recurrence. A point $x \in X$ is *non-wandering* if for any open neighborhood $U \subset X$ of x, there exists $n > 0$ such that $f^n(U) \cap U \neq \emptyset$. The map f is non-wandering if every point of X is non-wandering. A map f is *topologically transitive* if for any nonempty open sets $U, V \subset X$, there is an $n \geq 0$ such that $f^n(U) \cap V \neq \emptyset$. A map f is *topologically mixing* if for any nonempty open sets $U, V \subset X$ there is an $N > 0$ such that $f^n(U) \cap V \neq \emptyset$ for all $n \geq N$.

Chain recurrence. A point x is *chain recurrent* if for every $\varepsilon > 0$, there is an ε-chain from x to itself. The map f is *chain recurrent* if every point of X is chain recurrent.

If x is a chain recurrent point, define the *ε-chain recurrence time* $r_\varepsilon(x, f)$ to be the smallest n such that there is an ε-chain of length n from x to itself. If f is chain recurrent, define $r_\varepsilon(f)$ to be the maximum over all x of $r_\varepsilon(x, f)$. (To see that this maximum exists, observe that if there is an $\frac{\varepsilon}{2}$-chain of length n from x to itself, then there is a neighborhood U of x such that for all $y \in U$, there is an ε-chain of length n from y to itself. Then the compactness of X gives an upper bound on r_ε.) For the sake of brevity, we usually write "$r_\varepsilon(x)$" and "r_ε."

Chain transitivity. A map f is *chain transitive* if for every $x, y \in X$ and every $\varepsilon > 0$, there is an ε-chain from x to y. Following [10] we say that a map f is *totally chain transitive* if f^n is chain transitive for all $n \geq 0$.

Chain mixing. By compactness, a map f is topologically mixing if and only if for any $\varepsilon, \delta > 0$, there is an N such that for any $x, y \in X$ and any $n \geq N$, there are points $x' \in B_\delta(x)$ and $y' \in B_\varepsilon(y)$ with $f^n(x') = y'$. To mimic this, we say that f is *ε-chain mixing* if there is an $N > 0$ such that for any $x, y \in X$ and any $n \geq N$, there is an ε-chain from x to y of length exactly n. The map f is *chain mixing* if it is ε-chain mixing for every $\varepsilon > 0$. If $0 < \varepsilon < \delta$ and $x \in X$, define the *chain mixing time* $m_\varepsilon(x, \delta, f)$ to be the smallest N such that for any $n \geq N$ and any $y \in X$, there is an ε-chain of length exactly n from some point in $B_\delta(x)$ to y. (In other words, $m_\varepsilon(x, \delta, f)$ is the smallest N such that $(B_\varepsilon \circ f)^N(B_\delta(x)) = X$.) We define $m_\varepsilon(\delta, f)$ to be the maximum over all x of $m_\varepsilon(x, \delta, f)$. (This maximum exists by compactness.) Again, to simplify notation, in most cases we write "$m_\varepsilon(x, \delta)$" and "$m_\varepsilon(\delta)$."

If f is a physical system, then δ represents the uncertainty in the measurement of the initial state of the system, and ε the possible error in evaluating the map at each step. The chain mixing time $m_\varepsilon(\delta)$ is the time after which all information about the initial state of the system is lost, since any point could be anywhere by then.

We are indebted to the referee for some extremely helpful comments and ideas.

3. Structure of chain transitive maps

In this section we investigate the structure of chain transitive dynamical systems. To the best of our knowledge, the main result about this structure, Theorem 6, appears only in the exercises of [1]. Thus a self-contained discussion of it and its consequences may be useful.

We show that for a fixed $\varepsilon > 0$, the space X decomposes into finitely many subsets; f permutes these n subsets cyclically and f^n is ε-chain mixing on each of them. At one extreme, if ε is greater than the diameter of X, f is ε-chain mixing on X itself. As ε goes to zero, the subsets divide repeatedly, either finitely or infinitely many times.

For a fixed $\varepsilon > 0$, the dynamical structure of ε-chains for a chain transitive map is very similar to that of the symbolic dynamics determined by a finite irreducible graph. (See [19, §4.5], from which we borrow some terminology.) As ε goes to zero, the size of the graph may grow, possibly without bound.

Let us consider two elementary examples.

Example 1. Let $X = S^1$ (considered as $\mathbb{R}/\mathbb{Z}$). The doubling map $x \mapsto 2x$ is chain mixing.

Example 2. Let X be the disjoint union of two circles and f be the map sending a point x to the point $2x$ in the other circle. The map f is chain transitive but not chain mixing, because it is not ε-chain mixing for any ε smaller than the distance between the two circles. However, f^2 restricted to one circle *is* chain mixing.

We should not be misled by these examples; they are not prototypes for all chain transitive maps. As ε decreases the number of permuted subsets may go to infinity. As we see in the following example, it is possible that no power of f has associated invariant subsets on which it is chain mixing.

Example 3. Let X be the Cantor set of sequences $\mathbf{a} = (a_1, a_2, \ldots)$, where each $a_i \in \{0, 1\}$. Give X the metric $d(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^{\infty} \delta(a_i, b_i)/2^i$, where $\delta(a_i, b_i) = 0$ if $a_i = b_i$ and 1 otherwise. Define $f : X \to X$ by adding $(1, 0, 0, 0, \ldots)$, where addition is defined componentwise base 2, with carrying to the right. (For example, $f(1, 0, 0, 0, \ldots) = (1, 0, 0, 0, \ldots) + (1, 0, 0, 0, \ldots) = (0, 1, 0, 0, \ldots)$, and $f(1, 1, 1, 0, 0, \ldots) = (0, 0, 0, 1, 0, \ldots)$.) The map f, which is minimal, is called an *adding machine* or *odometer map*.

If we look only at the first digit, then f has period two. On the first two digits, f has period four $((0, 0) \mapsto (1, 0) \mapsto (0, 1) \mapsto (1, 1) \mapsto (0, 0))$, and so on. Thus f is ε-chain mixing for $\varepsilon \geq \frac{1}{2}$, but not for $\varepsilon < \frac{1}{2}$. However, when $\frac{1}{4} \leq \varepsilon < \frac{1}{2}$, f^2 restricted to $\{(0, a_2, a_3, \ldots)\}$ or $\{(1, a_2, a_3, \ldots)\}$ is ε-chain mixing, and in general, for $\frac{1}{2^k} \leq \varepsilon < \frac{1}{2^{k-1}}$, X decomposes into 2^k disjoint sets on which f^{2^k} is ε-chain mixing, but $f^{2^{k-1}}$ is not.

In fact, this is essentially the only example of a chain transitive map f such that no f^k is a disjoint union of finitely many chain mixing maps, in the sense that any such map factors onto an adding machine map, which we now define more generally. (Adding machines arise in many dynamical contexts; see, for example, [3, 7, 8, 9, 18, 20, 23]. See [6] for more detailed definitions and background.)

Definition 4 ([6])**.** Let $J = (j_1, j_2, \ldots)$ be a sequence of integers greater than or equal to 2. Let X_J be the Cantor set of all sequences $(a_1, a_2, \ldots)$ where $a_i \in \{0, 1, \ldots, j_i - 1\}$ for all i, with $d(\mathbf{a}, \mathbf{b}) = \sum_{i=1}^{\infty} \delta(a_i, b_i)/2^i$. Define the *adding machine map* $f_J : X_J \to X_J$ by $f_J(a_1, a_2, \ldots) = (a_1, a_2, \ldots) + (1, 0, 0, \ldots)$, where addition is defined componentwise mod j_i, with carrying to the right.

It is easy to see that every adding machine f_J is minimal. The following alternative definition of an adding machine as an inverse limit of cyclic groups will be useful.

Lemma 5 ([11, p. 21]). *Let $m_1, m_2, \ldots$ be an increasing sequence of positive integers such that m_i divides m_{i+1} for all i. Let $p_{m_i} : \mathbb{Z}/m_i \to \mathbb{Z}/m_i$ be the cyclic permutation given by $p_{m_i}(n) = n + 1$ mod m_i. Then the inverse limit map $p(\{n_i\}) = \{n_i + 1 \mod m_i\}$ is topologically conjugate to the adding machine map f_J, where $J = (m_1, m_2/m_1, m_3/m_2, \ldots)$.*

The following result is essentially contained in exercises 8.22 and 9.18 of [1]. We include it here for completeness and to fill in the gaps in the proofs sketched there.

Theorem 6 ([1]). *Let $f : X \to X$ be a chain transitive map. Then either*

(1) *There is a period $k \geq 1$ such that f cyclically permutes k closed and open equivalence classes of X, and f^k restricted to each equivalence class is chain mixing; or*

(2) *f factors onto an adding machine map.*

The general idea for the proof is to define an equivalence relation $\sim_\varepsilon$ on X. The induced map on the quotient space $X/\sim_\varepsilon$ is a periodic orbit. Sending ε to zero yields an equivalence relation $\sim$; thus f factors onto the inverse limit of the periodic orbits, which will be either a periodic orbit or an adding machine.

Lemma 7. *Let $f : X \to X$ be a chain transitive map and let $\varepsilon > 0$. There exists $k_\varepsilon \geq 1$ such that for any $x \in X$, k_ε is the greatest common divisor of the lengths of all ε-chains from x to itself.*

Proof. For each $x \in X$ define $k_\varepsilon(x)$ to be the greatest common divisor of the lengths of ε-chains from x to itself. Consider two points $x, y \in X$. Let $(y_0 = y, y_1, \ldots, y_n = y)$ be any ε-chain from y to itself. We show that $k_\varepsilon(x)$ divides n. Let $(x_0 = x, x_1, \ldots, y, \ldots, x_{mk_\varepsilon(x)} = x)$ be an ε-chain from x to y and back to x. The ε-chain $(x_0 = x, x_1, \ldots, y, y_1, \ldots, y_n = y, \ldots, x_{mk_\varepsilon(x)} = x)$ from x to itself has length $mk_\varepsilon(x) + n$. Since $mk_\varepsilon(x) + n$ is necessarily a multiple of $k_\varepsilon(x)$, n is as well. $\square$

Define a relation on X by setting $x \sim_\varepsilon y$ if there is an ε-chain from x to y of length a multiple of k_ε. By concatenating chains, it is easy to see that this is an equivalence relation. In fact, by the definition of k_ε, if $x \sim_\varepsilon y$, then *any* ε-chain from x to y must have length a multiple of

k_ε. We may now define another equivalence relation on X by saying $x \sim y$ if $x \sim_\varepsilon y$ for all $\varepsilon > 0$.

Remark 8. Since f is chain transitive, it is easy to see that $x \sim_\varepsilon y$ if and only if there is a point z such that there are ε-chains from x to z and from y to z of the same length. Thus another way to define the equivalence relation $\sim$ is as the set of pairs (x, y) such that the set of points reachable from (x, y) with $(f \times f)$-chains $(\mathcal{C}(f \times f)(x, y)$, in the notation of [1]) meets the diagonal.

Lemma 9. *Let $f : X \to X$ be a chain transitive map. For any $\varepsilon > 0$ the equivalence relation $\sim_\varepsilon$ is both open and closed. The equivalence relation $\sim$ is closed.*

Proof. It is clear that the relations $\sim_\varepsilon$ and $\sim$ are closed. To see that $\sim_\varepsilon$ is open, observe that an $\frac{\varepsilon}{2}$-chain is a fortiori an ε-chain. Since f is chain transitive, there is an $\frac{\varepsilon}{2}$-chain from x to y for any points x and y, and, if $x \sim_\varepsilon y$, the length of the chain must be a multiple of k_ε. Thus there is an ε-chain of the same length from x to any point within $\frac{\varepsilon}{2}$ of y (use the same $\frac{\varepsilon}{2}$ chain and change only the last element). $\square$

Now, let $\tilde{K} = X/\sim$ be the quotient space and $\tilde{f} : \tilde{K} \to \tilde{K}$ be the map induced by f.

Proof of Theorem 6. To begin, let $\varepsilon > 0$ be fixed. There are k_ε equivalence classes for $\sim_\varepsilon$, f cycles among the classes periodically, and each class is invariant under f^k. The quantity k_ε is nondecreasing as $\varepsilon \to 0$, and in fact k_{ε_2} divides k_{ε_1} if $\varepsilon_1 \leq \varepsilon_2$. Either k_ε stabilizes at some k, or it grows without bound. We consider the two cases separately.

If k_ε stabilizes at $k = 1$, then there is only one $\sim$ equivalence class, and f is chain mixing. To see this, observe that for any ε and any x, since 1 is the GCD of the lengths of ε-chains from x to itself, there exist ε-chains from x to itself of lengths m and n with m and n relatively prime. Sylvester's formula ([26]) tell us that by concatenating copies of these loops, we can get an ε-chain from x to itself of length N for any $N > mn - m - n$. By compactness, there is an $M > 0$ such that between any two points in X there is an ε-chain of length less than or equal to M. By adding a loop at x to a chain from x to y, we can get a chain from x to y of any length greater than $M + N$, so f is chain mixing.

More generally, if $k = k_{\varepsilon_0}$, then the equivalence relation $\sim$ is the same as $\sim_{\varepsilon_0}$. Thus there are k equivalence classes, f cycles among the classes periodically, and each class is invariant under f^k. An argument similar to that for $k = 1$ shows that f^k is chain mixing on each equivalence class.

If, on the other hand, k_ε grows without bound, then the period of f's cycling goes to infinity as ε shrinks. In this case, f factors onto an adding machine map. More precisely, let $\tilde{K}_\varepsilon$ be the quotient space X modulo $\sim_\varepsilon$, with the quotient topology, and let $\tilde{f}_\varepsilon : \tilde{K}_\varepsilon \to \tilde{K}_\varepsilon$ be the induced map. Let $\tilde{f} : \tilde{K} \to \tilde{K}$ be the induced map on the quotient space $\tilde{K} = X/\sim$. The map $\tilde{f}_\varepsilon : \tilde{K}_\varepsilon \to \tilde{K}_\varepsilon$ is conjugate to the cyclic permutation $p_{k_\varepsilon} : \mathbb{Z}/k_\varepsilon\mathbb{Z} \to \mathbb{Z}/k_\varepsilon\mathbb{Z}$ given by $p_{k_\varepsilon}(n) = n + 1$ mod k_ε. Thus, if we take a sequence $\{\varepsilon_i\}$ decreasing to 0 such that the sequence $\{k_{\varepsilon_i}\}$ is strictly increasing, then $\tilde{f}$ is conjugate to the map $g : Z \to Z$, where Z is the inverse limit of the sequence $\{p_{k_{\varepsilon_i}} : \mathbb{Z}/k_{\varepsilon_i}\mathbb{Z} \to \mathbb{Z}/k_{\varepsilon_i}\mathbb{Z}\}$ and g is given by $g(\{x_i\}) = \{x_i + 1 \mod k_{\varepsilon_i}\}$. Then, by Lemma 5, g is topologically conjugate to the adding machine map f_J, where $J = (k_{\varepsilon_1}, k_{\varepsilon_2}/k_{\varepsilon_1}, k_{\varepsilon_3}/k_{\varepsilon_2}, \dots)$. Thus the conjugacy class of f_J is independent of the choice of the sequence $\{\varepsilon_i\}$. In fact, it is shown in [8] that the conjugacy class of an adding machine $f_{(j_1, j_2, \dots)}$ depends only on the sums of the powers to which prime numbers occur as factors of the j_i's. $\square$

We have a number of nice corollaries that follow from Theorem 6.

Corollary 10. *Let $f : X \to X$ be a chain transitive map. If f has a point of period n, then there is a k dividing n such that X decomposes into k closed and open equivalence classes, f permutes the equivalence classes cyclically, and f^k restricted to each equivalence class is chain mixing.*

Proof. If $x \in X$ has period n under f, then the image of x in $\tilde{K}$ has period k under $\tilde{f}$, for some k dividing n. Thus $\tilde{f}$ cannot be an adding machine map. $\square$

Corollary 11. *Let $f : X \to X$ be a chain transitive map. If one of the $\sim$-equivalence classes is open, then there is a k such that X decomposes into k closed and open equivalence classes, f permutes the equivalence classes cyclically, and f^k restricted to each equivalence class is chain mixing.*

Proof. If one of the $\sim$-equivalence classes is open, then its projection in $\tilde{K}$ is a single point that is both open and closed. Thus $\tilde{K}$ cannot be a Cantor set. $\square$

By definition, if f is chain mixing, then it is totally chain transitive. By Theorem 6, the converse is also true.

Corollary 12. *$f : X \to X$ is totally chain transitive if and only if it is chain mixing.*

In fact, as is shown in [28], these two conditions are also equivalent to chain transitivity for $f \times f$. To put this result in perspective, we contrast it with topologically transitive maps. (See also [5].)

Example 13. An irrational rotation $f : S^1 \to S^1$ is totally topologically transitive, but not topologically mixing.

Conley's fundamental theorem of dynamical systems ([27, §9.1]), says that the interesting dynamics of f takes place on the chain recurrent set $\mathcal{R}$, with all other orbits simply going from one piece of $\mathcal{R}$ to another. In some special cases (if $\mathcal{R}$ has a hyperbolic structure, for example ([27, §9.5])), $\mathcal{R}$ decomposes into finitely many chain transitive pieces (basic sets), so there is not much difference between chain recurrence and chain transitivity. In general, the situation is more complicated, but we do have the following result.

Corollary 14. *Let X be connected and $f : X \to X$ be continuous. Then the following are equivalent:*

(1) *f is chain recurrent.*
(2) *f is chain transitive.*
(3) *f is totally chain transitive.*
(4) *f is chain mixing.*

Proof. Clearly $(4) \Rightarrow (3) \Rightarrow (2) \Rightarrow (1)$. By Theorem 6, connectivity and chain transitivity imply that f is chain mixing. So it is enough to show that chain recurrence implies chain transitivity. This follows from the proof of Proposition 1.1 in [2], but for completeness we now give a more direct proof. Assume that f is chain recurrent and pick $\varepsilon > 0$. We say that x and y are ε-chain equivalent if there are ε-chains from x to y and from y to x. Since f is chain recurrent, this is an equivalence relation. So to finish the proof it suffices to show that this is an open condition. Let x and y be ε-chain equivalent. Choose $\delta \leq \varepsilon/2$ such that if $d(y, y') < \delta$, then $d(f(y), f(y')) < \varepsilon/2$. It suffices to show that x is ε-chain equivalent to an arbitrary $z \in B_\delta(y)$. Let $(x_0 = x, x_1, \ldots, x_n = y)$ be an ε-chain and let $(z_0 = y, z_1, \ldots, z_m = y)$ be an $\varepsilon/2$-chain from y to itself. Then $(x_0 = x, x_1, \ldots, x_n = y = z_0, \ldots, z_{m-1}, z)$ is an ε-chain from x to z. Similarly, let $(y_0 = y, y_1, \ldots, y_r = x)$ be an ε-chain from y to x. Then $(z, z_1, \ldots, z_m = y = y_0, \ldots, y_n = x)$ is an ε-chain from z to x. Thus x is ε-chain equivalent to z. $\square$

Recall (from [4], for example) the definition $R_\pi = \{(x, y) : \pi(x) = \pi(y)\}$ for a factor map $\pi : (X, f) \to (Y, g)$. The system (X, f) is *weak mixing* if $f \times f$ is topologically transitive; the factor map π is *weak mixing* if the restriction of $f \times f$ to R_π is topologically transitive.

Analogously, we will define the factor map π to be *chain mixing* if the restriction of $f \times f$ to R_π is chain transitive. The following proposition shows that when π is the quotient map associated to $\sim$, then π is chain mixing.

Proposition 15. *The relation $R_\pi = \sim$ is a maximal closed, chain transitive, invariant subset of $X \times X$.*

Proof. It is clearly closed and invariant. To see that it is chain transitive, fix $\varepsilon > 0$ and points $x_0, y_0, x_1, y_1 \in X$ with $x_0 \sim y_0$ and $x_1 \sim y_1$. We must show that there is an ε-chain for $f \times f$ from (x_0, y_0) to (x_1, y_1). Let $\tilde{K}_0$ be the $\sim_\varepsilon$-equivalence class containing x_0 and y_0, and $\tilde{K}_1$ be the $\sim_\varepsilon$-equivalence class containing x_1 and y_1. There is an n such that f^n maps K_0 onto K_1. Since f^{k_ε} restricted to $\tilde{K}_0$ is chain mixing, there is an ε-chain for $f \times f$ from (x_0, y_0) to $(f^{-n}(x_1), f^{-n}(y_1))$, which we can extend by an actual orbit to (x_1, y_1).

To see that $\sim$ is a maximal chain transitive subset, observe that if $x_0 \sim_\varepsilon y_0$ and $(x_0, y_0), (x_1, y_1), \ldots, (x_m, y_m)$ is an ε-chain, then $x_i \sim_\varepsilon y_i$ for all i. Thus there cannot be an ε-chain from (x_0, y_0) to (x, y) unless $x \sim_\varepsilon y$. $\qquad \square$

We conclude this section with a few results from the literature that are similar to Theorem 6. However, instead of looking at the induced map on the space of equivalence classes $\tilde{K}$, these relate to the induced map on K, the set of connected components of X. See [8, 12, 21] for details and more general results. The first is a "folklore" result similar to Theorem 6 for topologically transitive maps.

Theorem 16 ([8]). *Let X be a compact metric space, $f : X \to X$ be topologically transitive map, and $\bar{f} : K \to K$ be the induced map on the connected components of X. Then either*
   (1) *K is finite and $\bar{f}$ is a cyclic permutation, or*
   (2) *K is a Cantor set and $\bar{f}$ is topologically transitive.*

Notice that under these hypotheses if x and y are in the same connected component, then necessarily $x \sim y$. Thus $\bar{f} : K \to K$ factors onto $\tilde{f} : \tilde{K} \to \tilde{K}$. If X has only finitely many connected components, then the two decompositions and the two maps are the same.

Definition 17. Let $A \subset X$ be a compact invariant set.
   (1) A is topologically transitive (resp. chain transitive) if $f|_A : A \to A$ is topologically transitive (resp. chain transitive).
   (2) A is *Liapunov stable* if every neighborhood U of A contains a neighborhood V of A such that $f^n(V) \subset U$ for all $n \geq 0$.

(3) A is an *attractor* if there exists a neighborhood U of A such that $\mathrm{cl}(f(U)) \subset \mathrm{Int}\, U$ and $A = \bigcap_{n\geq 0} f^n(U)$.

Theorem 18 ([8, 12]). *Let X be a locally compact, locally connected metric space and $f : X \to X$ be a continuous map. Let $A \subset X$ be a compact, Liapunov stable transitive set, and $\bar{f} : K \to K$ be the induced map on the connected components of A. Then either*

(1) *K is finite and $\bar{f}$ is a cyclic permutation, or*
(2) *K is a Cantor set and $\bar{f}$ is topologically conjugate to an adding machine map.*

Theorem 19 ([12]). *Let X be a locally compact, locally connected metric space and $f : X \to X$ be a continuous map. Let $A \subset X$ be a compact, chain transitive attractor. Then $\bar{f} : K \to K$, the induced map on the connected components of A, is a cyclic permutation of a finite set.*

Theorems 6 and 19 imply the following result.

Corollary 20. *Let X be a locally compact, locally connected metric space and $f : X \to X$ be a continuous map. Let $A \subset X$ be a compact, chain transitive attractor. Then there is a k such that A decomposes into k closed and open equivalence classes, f permutes the equivalence classes cyclically, and f^k restricted to each equivalence class is chain mixing.*

It is worth remarking here that on many compact metric spaces, including manifolds, homeomorphisms generically have no chain transitive attractors. See [3, 15].

We conclude this section with an example.

Example 21. Construct a Denjoy map $f : S^1 \to S^1$ by starting with an irrational rotation R_α and replacing $\{R_\alpha^n(0)\}_{n\in\mathbb{Z}}$, the orbit of 0, with a sequence of intervals $\{I_n\}$ such that $\sum_{n\in\mathbb{Z}} \mathrm{length}(I_n) = 1$ and I_n is mapped into I_{n+1} by an orientation-preserving homeomorphism (see [27, §2.8] for details).

So defined, f is chain mixing, but not topologically transitive. In fact, the omega limit set of any $x \in S^1$ is the invariant Cantor set $C = S^1 \backslash \bigcup_{n\in\mathbb{Z}} \mathrm{Int}(I_n)$.

If we restrict the map to C, then $f|_C$ is still chain mixing and $f|_C$ is topologically transitive. Thus the space K of connected components of C is C itself, while $\tilde{K}$ is a single point.

4. Estimating recurrence rates

Let $f : X \to X$ be a chain transitive map. We wish to estimate the chain recurrence time r_ε and, in the case that f is chain mixing, the

chain mixing time $m_\varepsilon(\delta)$ for a given ε. Our technique is to model ε-chains with a subshift of finite type, an idea which has been used before in some other contexts ([1, 14, 16, 24], for example). The approach we use here is adapted from that in chapter 5 of [1]. For calculations of r_ε and $m_\varepsilon(\delta)$ for a variety of dynamical systems, see section 6.

Proposition 22. *Let d' be the upper box dimension of X (see [25, §6] for a definition of upper box dimension). There exists a constant $C > 0$ such that for small enough ε:*

(1) *if $f : X \to X$ is chain transitive, then $r_\varepsilon \leq C/\varepsilon^{d'}$.*
(2) *if $f : X \to X$ is chain mixing, then $m_\varepsilon(\delta) \leq C/\varepsilon^{2d'}$.*

Proof. Let $\{X_i\}_{i=1}^{N_{\mathrm{span}}(\frac{\varepsilon}{2})}$ be a minimal cardinality collection of balls of radius $\frac{\varepsilon}{2}$ whose interiors cover X. Observe that for small enough ε, $N_{\mathrm{span}}(\frac{\varepsilon}{2}) + 1 \leq C/\varepsilon^{d'}$ for some constant C, where d' is the upper box dimension of X ([25, §6]). Construct the $N_{\mathrm{span}}(\frac{\varepsilon}{2}) \times N_{\mathrm{span}}(\frac{\varepsilon}{2})$ transition matrix A_ε by setting $a_{ij} = 1$ if $f(X_j) \cap X_i \neq \emptyset$ and 0 otherwise. If $(i_0, \ldots, i_k) \in \{1, \ldots, N_{\mathrm{span}}(\frac{\varepsilon}{2})\}^{k+1}$ is an allowable word in the subshift of finite type determined by A_ε, then for any $x \in f^{-1}(X_{i_0})$ and any $x_k \in X_{i_k}$, there is an ε-chain $(x, x_0, \ldots, x_k)$ with $x_j \in X_{i_j}$ for each j. Conversely (see Lemma 5.1 of [1]), if $(x_0, \ldots, x_k)$ is an α-chain, where α is less than the Lebesgue number of the cover $\{X_i\}$ and $x_j \in X_{i_j}$, then the word $(i_0, \ldots, i_k)$ is an allowable word. Thus, because f is chain transitive, for any i and j, there is an allowable word starting at i and ending at j. By removing any subwords that begin and end at the same symbol, we can assume that the word has length less than or equal to $N_{\mathrm{span}}(\frac{\varepsilon}{2})$, and so for any two points x and y there is an ε-chain from x to y of length less than or equal to $N_{\mathrm{span}}(\frac{\varepsilon}{2}) + 1 \leq C/\varepsilon^{d'}$ for small enough ε. This proves the first part of the result.

To prove the second part, observe that since f is chain mixing, the matrix A_ε is eventually positive. So if p is the smallest integer such that each entry of A_ε^p is strictly positive, then $m_\varepsilon(\delta) \leq p + 1$. Since $p \leq (N_{\mathrm{span}}(\frac{\varepsilon}{2}) - 1)^2 + 1$ ([13]), we have that $m_\varepsilon(\delta) \leq C/\varepsilon^{2d'}$ for small enough ε. $\qquad\square$

We also have a lower bound for $m_\varepsilon(\delta)$ in the case that f is Lipschitz. We know that the Lipschitz constant c is at least 1, since otherwise the Banach fixed point theorem says that f has an attracting fixed point, contradicting the fact that f is chain mixing.

Proposition 23. *Let $f : X \to X$ be chain mixing and have Lipschitz constant c. Let D be the diameter of X. Then for δ sufficiently small,*

$$m_\varepsilon(\delta) \geq \log_c \left(\frac{D(c-1) + 2\varepsilon}{2\delta(c-1) + 2\varepsilon} \right) \text{ if } c > 1, \text{ and } m_\varepsilon(\delta) \geq \frac{D - 2\delta}{2\varepsilon} \text{ if } c = 1.$$

Proof. The diameter of the nth iterate of a ball of radius δ under $B_\varepsilon \circ f$ is

$$\mathrm{diam}((B_\varepsilon \circ f)^n(B_\delta(x))) \leq c^n(2\delta) + c^{n-1}(2\varepsilon) + c^{n-2}(2\varepsilon) + \cdots + 2\varepsilon.$$

So, if $c > 1$, then

$$\mathrm{diam}((B_\varepsilon \circ f)^n(B_\delta(x))) \leq c^n(2\delta) + \frac{1 - c^n}{1 - c}(2\varepsilon).$$

Thus if $\mathrm{diam}((B_\varepsilon \circ f)^n(B_\delta(x)))$ is D, then n is at least $\log_c \left(\frac{D(c-1)+2\varepsilon}{2\delta(c-1)+2\varepsilon} \right)$.

When $c = 1$, $\mathrm{diam}((B_\varepsilon \circ f)^n(B_\delta(x))) \leq 2\delta + 2n\varepsilon$.

$\square$

Give the product space $X \times Y$ the sup metric, i.e., $d((x,y),(x',y')) = \max(d_X(x,x'), d_Y(y,y'))$.

Proposition 24. *Let X and Y be compact, $f : X \to X$ and $g : Y \to Y$ be chain recurrent, and $k \in \mathbb{Z}^+$. Then $f^k : X \to X$ and $f \times g : X \times Y \to X \times Y$ are also chain recurrent, and for all $\varepsilon > 0$, $x \in X$, and $y \in Y$,*

(1) $r_\varepsilon((x,y), f \times g) \geq \max(r_\varepsilon(x,f), r_\varepsilon(y,g))$,
(2) $r_\varepsilon((x,y), f \times g) \leq \mathrm{lcm}(r_\varepsilon(x,f), r_\varepsilon(y,g))$,
(3) $r_\varepsilon(x, f^k) \geq \dfrac{1}{k} r_\varepsilon(x,f)$, and
(4) *there exists an $\varepsilon' \leq \varepsilon$ such that $r_\varepsilon(x, f^k) \leq r_{\varepsilon'}(x,f)$.*

Proof. Observe that $((x_0, y_0), \ldots, (x_n, y_n))$ is an ε-chain for $f \times g$ if and only if $(x_0, \ldots, x_n)$ and $(y_0, \ldots, y_n)$ are ε-chains for f and g respectively. Statement (1) and the fact that $f \times g$ is chain recurrent if and only if f and g both are follow immediately.

To prove (2), let $(x_0 = x, x_1, \ldots, x_n = x)$ and $(y_0 = y, y_1, \ldots, y_m = y)$ be ε-chains. Then the ε-chain

$$(x_0, \ldots, x_n = x_0, x_1, \ldots, x_n = x_0, \ldots \ldots, x_n = x_0)$$

formed by concatenating $(x_0 = x, x_1, \ldots, x_n = x)$ with itself $\dfrac{m}{\gcd(n,m)}$ times has length $\mathrm{lcm}(m,n)$, as does the ε-chain

$$(y_0, \ldots, y_m = y_0, y_1, \ldots, y_m = y_0, \ldots \ldots, y_m = y_0)$$

formed by concatenating $(y_0 = y, y_1, \ldots, y_m = y)$ with itself $\dfrac{n}{\gcd(n,m)}$ times. Combining the two gives an ε-chain from (x,y) to itself.

To show that f is chain recurrent if f^k is, and to prove (3), it suffices to observe that if $(x_0, \ldots, x_n)$ is an (ε, f^k)-chain, then

$$(x_0, f(x_0), \ldots, f^{k-1}(x_0), x_1, f(x_1), \ldots \ldots, f^{k-1}(x_{n-1}), x_n)$$

is an (ε, f)-chain of length nk.

To see that f^k is chain recurrent if f is, pick ε' small enough that for any (ε', f)-chain $(x_0, \ldots, x_k)$, x_k is within ε of $f^k(x_0)$. (Such an ε' exists because of the uniform continuity of f.) Thus, if $(x_0, \ldots, x_{nk})$ is an (ε', f)-chain of length nk, then $(x_0, x_k, x_{2k}, \ldots, x_{nk})$ is an (ε, f^k)-chain of length n. To prove (4), observe that by taking an (ε', f)-chain of length n from x to x and concatenating it with itself k times, we get an (ε', f)-chain of length nk, and thus an (ε, f^k)-chain of length n. $\square$

Remark 25. If we give $X \times Y$ the Euclidean metric, then (2) becomes "$r_\varepsilon((x, y), f \times g) \le \mathrm{lcm}(r_{\varepsilon/2}(x, f), r_{\varepsilon/2}(y, g))$."

Proposition 26. *Let X and Y be compact, $f : X \to X$ and $g : Y \to Y$ be chain mixing, and $k \in \mathbb{Z}^+$. Then $f^k : X \to X$ and $f \times g : X \times Y \to X \times Y$ are also chain mixing, and for all $\varepsilon > 0$,*

(1) $m_\varepsilon(\delta, f \times g) = \max(m_\varepsilon(\delta, f), m_\varepsilon(\delta, g))$

(2) $m_\varepsilon(\delta, f^k) \ge \dfrac{1}{k} m_\varepsilon(\delta, f)$

(3) *there exists an $\varepsilon' \le \varepsilon$ such that $m_\varepsilon(\delta, f^k) \le m_{\varepsilon'}(\delta, f)$.*

Proof. Statement (1) and the fact that $f \times g$ is chain mixing if and only if both f and g are follow from the definition of chain mixing. The proofs of (2) and (3) and of the fact that f^k is chain mixing are essentially identical to those of the corresponding statements for chain recurrent maps in Proposition 24. $\square$

Notice that if $f : X \to X$ is chain recurrent and $S \subset X$ is an invariant subset, then $r_\varepsilon(f|_S)$ may be greater than or less than $r_\varepsilon(f)$; while clearly $r_\varepsilon(x, f|_S) \ge r_\varepsilon(x, f)$ for $x \in S$, we take the maximum over a smaller set (S instead of X) to get $r_\varepsilon(f|_S)$. Consider the following example.

Example 27. Let $X = S_1 \cup S_2 \cup S_3$ and $X' = S_1 \cup S_2$ where S_1, S_2, and S_3 are the circles shown in Figure 1. Let $f : X \to X$ be a homeomorphism with fixed points at the south poles and motion indicated by the arrows. Assume that away from the fixed points the motion on S_1 is much faster than the motion on S_2 which is much faster than the motion on S_3. If S_1 stays close enough to S_2 near their common fixed point, then a point x near the fixed point can return to itself more quickly via an ε-chain passing through S_1 than via one that
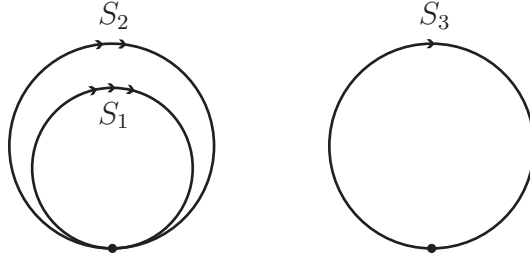
FIGURE 1.

stays in S_2, so $r_\varepsilon(x, f|_{S_2}) \geq r_\varepsilon(x, f|_{X'})$, and in fact $r_\varepsilon(f|_{S_2}) > r_\varepsilon(f|_{X'})$. On the other hand, points move very slowly on S_3, so $r_\varepsilon(f|_{S_2}) < r_\varepsilon(f)$.

5. TOPOLOGICAL ENTROPY

The following theorem shows that the growth of the chain mixing times can give information about topological entropy.

Theorem 28. *Let $f : X \to X$ be chain mixing. Then the topological entropy, $h(f)$, satisfies*

$$h(f) \geq d \cdot \limsup_{\delta \to 0} \frac{\log(1/\delta)}{\lim_{\varepsilon \to 0} m_\varepsilon(\delta)},$$

where d is the lower box dimension of X.

Proof. In [22], Misiurewicz shows that we can use the growth rate of ε-chains to calculate topological entropy. Specifically,

$$h(f) = \lim_{\delta \to 0} \lim_{\varepsilon \to 0} \limsup_{n \to \infty} \frac{1}{n} \log S(n, \delta, \varepsilon),$$

where $S(n, \delta, \varepsilon)$ is the maximum cardinality of a set of δ-separated ε-chains of length n. (Two ε-chains $(x_0, \ldots, x_{n-1})$ and $(y_0, \ldots, y_{n-1})$ are δ-separated if $d(x_i, y_i) \geq \delta$ for some i.)

For $\alpha > 0$, let $N(\alpha) = S(0, \alpha, 0)$ (i.e., the maximum cardinality of an α-separated collection of points). We claim that for $k \geq 0$, $S(km_\varepsilon(\delta), \delta, \varepsilon) \geq (N(3\delta))^{k+1}$. To see this, let $\{x_1, \ldots, x_{N(3\delta)}\}$ be a 3δ-separated set of points. By the definition of $m_\varepsilon(\delta)$, for each of the $(N(3\delta))^{k+1}$ sequences $(i_0, \ldots, i_k)$ $(1 \leq i_j \leq N(3\delta))$ there is a sequence of points $(x_{i_1,1}, \ldots, x_{i_k,k})$ such that for each j, $x_{i_j,j} \in B_\delta(x_i)$ and there is an ε-chain of length $m_\varepsilon(\delta)$ from $x_{i_j,j}$ to $x_{i_{j+1},j+1}$. Since the points x_i are 3δ-separated, the sequences $(x_{i_1,1}, \ldots, x_{i_k,k})$ are δ-separated. Thus $S(km_\varepsilon(\delta), \delta, \varepsilon) \geq (N(3\delta))^{k+1}$.

Next, let $N_{\text{span}}(\alpha)$ be the minimum number of balls of radius α necessary to cover X. By the definition of lower box dimension ([25, §6]),

for small enough α, $N_{\text{span}}(\alpha) \geq C(1/\alpha)^d$ for some positive constant C. Clearly $N(\alpha) \geq N_{\text{span}}(\alpha)$ ([27, Lemma 8.1.10]), so $N(\alpha) \geq C(1/\alpha)^d$ as well.

Finally, we have that

$$
\begin{aligned}
h(f) &= \lim_{\delta \to 0} \lim_{\varepsilon \to 0} \limsup_{n \to \infty} \frac{1}{n} \log S(n, \delta, \varepsilon) \\
&\geq \limsup_{\delta \to 0} \lim_{\varepsilon \to 0} \limsup_{k \to \infty} \frac{1}{km_\varepsilon(\delta)} \log(N(3\delta))^{k+1} \\
&= \limsup_{\delta \to 0} \lim_{\varepsilon \to 0} \frac{\log N(3\delta)}{m_\varepsilon(\delta)} \\
&\geq \limsup_{\delta \to 0} \lim_{\varepsilon \to 0} \frac{\log C/(3\delta)^d}{m_\varepsilon(\delta)} \\
&= d \limsup_{\delta \to 0} \frac{\log 1/\delta}{\lim_{\varepsilon \to 0} m_\varepsilon(\delta)},
\end{aligned}
$$

since $\lim_{\delta \to 0} m_\varepsilon(\delta) = \infty$ (recall that X is not a single point, and it cannot be a finite collection of points since f is chain mixing). $\qquad \square$

In the following section we see that the formula in Theorem 28 can give the precise value of the topological entropy.

6. Examples

We conclude with three examples—rigid rotations of the circle, subshifts of finite type, and the doubling map on the circle—in which we illustrate the ideas in the previous sections.

Example 29 (Rigid rotations). Consider the dynamical system R_α : $S^1 = \mathbb{R}/\mathbb{Z} \to S^1$ which is the rigid rotation by $\alpha \in \mathbb{R}$ given by $R_\alpha(x) = \alpha + x \pmod 1$. We know that when α is rational every point is periodic and when α is irrational every point has a dense orbit. Thus, in either case S^1 is chain mixing. The chain mixing time is just the least integer greater than or equal to $\frac{1}{2\varepsilon}$.

The chain recurrence time for R_α depends on the Diophantine properties of α. Let $\varepsilon > 0$ be given. Clearly, $r_\varepsilon(x, R_\alpha) = r_\varepsilon(R_\alpha) = r_\varepsilon(R_{\alpha+n})$ for all $x \in S^1$ and $n \in \mathbb{Z}$. So we may assume that $0 < \varepsilon < \alpha < 1$ and $x = 0$. Notice that for a rotation, the set of points that can be reached with an ε-chain of length n beginning at 0 is precisely those points within $n\varepsilon$ of $R_\alpha^n(0)$. Therefore r_ε is the smallest integer $n > 0$ such that $d(R_\alpha^n(0), 0) < n\varepsilon$. Equivalently, we want to find the smallest $n > 0$ such that

$$
|n\alpha - m| < n\varepsilon, \text{ or equivalently, } \left| \alpha - \frac{m}{n} \right| < \varepsilon
$$

for some $m \in \mathbb{Z}$.

If α is rational then 0 is periodic with some period, q. In this case, for ε sufficiently small, the shortest ε chain from 0 to 0 is the periodic orbit. Thus, $r_\varepsilon = q$.

Now assume α is irrational. To analyze this case we need to use continued fractions (for details, see [17]). Let

$$[a_0, a_1, a_2, ...] = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cdots}}$$

be the continued fraction expansion for α, and $p_k/q_k = [a_0, a_1, ..., a_k]$ the kth convergent (assume that p_k/q_k is in lowest terms). We say that the rational number p/q ($q > 0$) is a *best approximation* of α if the inequalities $p/q \neq r/s$ and $0 < s \leq q$ imply

$$|q\alpha - p| < |s\alpha - r|.$$

In other words, the only way to get a better rational approximation of α than a best approximation is by increasing the denominator.

Theorem 30 ([17]). *Suppose α is irrational. The rational number p/q ($q > 0$) is a best approximation of α if and only if it is a convergent of α.*

From Theorem 30 we can conclude that for a given ε, the values of n and m that we wish to find are precisely q_k and p_k for some convergent p_k/q_k of α. That is, the function r_ε is a nonincreasing piecewise constant function taking on the values $\{q_k\}$. Thus, to understand the behavior of r_ε as $\varepsilon \to 0$ we must investigate the rate of convergence of p_k/q_k to α. The following theorem gives a bound on this rate of convergence.

Theorem 31 ([17]). *Let $\{p_k/q_k\}$ be the sequence of convergents for $\alpha \in \mathbb{R}$. Then for all $k > 1$*

$$\left| \alpha - \frac{p_k}{q_k} \right| < \frac{1}{q_k^2}.$$

Thus, for a given ε, if p_k/q_k is the first convergent of α with $q_k \geq 1/\sqrt{\varepsilon}$, then

$$\left| \alpha - \frac{p_k}{q_k} \right| < \frac{1}{q_k^2} \leq \varepsilon.$$

In other words, the fastest that r_ε can grow is at a rate of $O(1/\sqrt{\varepsilon})$.

In fact, this worst case can be realized. For example, when $\alpha = (\sqrt{5} - 1)/2 = [1, 1, 1, 1, \ldots]$ we have ([17])

$$\left| \alpha - \frac{p_k}{q_k} \right| = \frac{1}{(\sqrt{5} + \beta_k)q_k^2}$$

where $\beta_k \to 0$ as $k \to \infty$. Thus, when $\varepsilon = 1/q_k^2$ we have

$$\left| \alpha - \frac{p_k}{q_k} \right| = \frac{\varepsilon}{\sqrt{5} + \beta_k} \sim \frac{\varepsilon}{\sqrt{5}}$$

for large k, so $r_\varepsilon(R_\alpha) \approx \dfrac{1}{\sqrt{\sqrt{5}\varepsilon}}$.

On the other hand, there are α for which $r_\varepsilon(R_\alpha)$ grows slowly—as slowly as desired—as $\varepsilon \to 0$. In particular, if $\psi : \mathbb{R}^+ \to \mathbb{R}^+$ is any decreasing function with $\psi(x) \to \infty$ as $x \to 0$, then there is an α such that $r_{\varepsilon_i}(R_\alpha) \leq \psi(\varepsilon_i)$ for some sequence $\varepsilon_i \to 0$. To see this we need the following theorem.

Theorem 32 ([17])**.** *For any function $\varphi : \mathbb{Z}^+ \to \mathbb{R}^+$, there is an irrational number α such that the inequality $|\alpha - p/q| < \varphi(q)$ has an infinite number of solutions in p and q $(q > 0)$.*

Define $\varphi = \psi^{-1}|_{\mathbb{Z}^+}$. Let α be the irrational number guaranteed by Theorem 32 for this φ. Then there exist infinitely many p and q with $|\alpha - p/q| < \varphi(q)$. Let $\{q_i\}$ denote this collection of denominators, listed in increasing order, and let $\varepsilon_i = \varphi(q_i)$. Then $r_{\varepsilon_i}(R_\alpha) \leq q_i = \varphi^{-1}(\varepsilon_i) = \psi(\varepsilon_i)$.

Example 33 (Subshifts of finite type)**.** Let $\sigma : \Sigma \to \Sigma$ be a subshift of finite type given by an $n \times n$ transition matrix A, with $k \geq 2$, and metric $d(\mathbf{a}, \mathbf{b}) = \sum \dfrac{\delta(a_i, b_i)}{2^{|i|}}$. Since an ε-jump can take any element of a cylinder set of diameter ε to any other element in the same cylinder set, it is easy to see that σ is chain transitive if and only if it is topologically transitive. (This is because, for a one-sided shift, we can consider only ε-chains that have a jump only at the first step, and afterwards are actual orbits; for a two-sided shift, they can have a jump at the beginning and at the end.) The shift space Σ decomposes into finitely many topologically transitive pieces ([19, §§6.3, 4.4]), so we assume that σ is topologically transitive. For notational convenience, we consider one-sided shifts (the proofs for two-sided shifts are essentially the same).

Given $\varepsilon > 0$ and $\mathbf{a} = (a_0, a_1, \ldots) \in \Sigma$, we want to find $r_\varepsilon(\mathbf{a})$. Let k be the smallest nonnegative integer satisfying $2^{-k-1} < \varepsilon$. Since σ is topologically transitive, the transition matrix A is irreducible ([19,

§6.3]), so there is an allowable word $(a_k, a'_1, a'_2, \ldots, a'_{m-1}, a_0)$ of length $m \le n$ from a_k to a_0. Thus the sequence $(\mathbf{b}_0 = \mathbf{a}, \mathbf{b}_1, \ldots, \mathbf{b}_{k+m} = \mathbf{a})$ defined by $\mathbf{b}_1 = (a_1, a_2, \ldots, a_k, a'_1, a'_2, \ldots, a'_{m-1}, a_0, a_1, \ldots)$ and $\mathbf{b}_i = \sigma(\mathbf{b}_{i-1})$ for $i > 1$ is an ε-chain from $\mathbf{a}$ to itself of length $k + m$. Notice that for some $\mathbf{a} \in \Sigma$ we can find an ε-chain from $\mathbf{a}$ to itself that is shorter than this, but for most points in Σ this is the shortest possible ε-chain. Thus, the worst case is that $r_\varepsilon(\mathbf{a}) = O(\log_2(1/\varepsilon))$.

Since we assume that $\varepsilon < \delta$, a similar argument shows that if σ is chain mixing, then $m_\varepsilon(\delta, \sigma)$ is $O(\log_2(1/\delta))$ if σ is one-sided, and $O(\log_2(1/\varepsilon))$ if σ is two-sided.

Example 34 (The doubling map). Let $f(x) = 2x \mod 1$ be the doubling map on the circle $S^1 = \mathbb{R}/\mathbb{Z}$. If we consider binary expansions, then $S^1 = \{(0.a_1a_2 \ldots) : a_i \in \{0,1\}\}$, with the identifications $(0.11 \ldots) = (0.00 \ldots)$, $(0.011 \ldots) = (0.100 \ldots)$, etc. Thus the argument in Example 33 shows that $r_\varepsilon(f) = \lceil \log_2(1/\varepsilon) \rceil$, where $\lceil \cdot \rceil$ is the least integer function. Arguing as in the proof of Proposition 23, we see that $m_\varepsilon(\delta, f) = \lceil \log_2\left(\frac{1+2\varepsilon}{2\delta+2\varepsilon}\right) \rceil$. Then Theorem 28 says that

$$h(f) \ge \limsup_{\delta \to 0} \frac{\log(1/\delta)}{\lim_{\varepsilon \to 0} \lceil \log_2\left(\frac{1+2\varepsilon}{2\delta+2\varepsilon}\right) \rceil} = \limsup_{\delta \to 0} \frac{\log(1/\delta)}{\log_2(1/2\delta)} = \log 2,$$

so in this case the estimate is exact.

References

[1] Ethan Akin, *The general topology of dynamical systems*, Graduate Studies in Mathematics, vol. 1, American Mathematical Society, Providence, RI, 1993. MR MR1219737 (94f:58041)

[2] _____, *Stretching the Oxtoby-Ulam theorem*, Colloq. Math. **84/85** (2000), part 1, 83–94. MR MR1778842 (2001e:37005)

[3] Ethan Akin, Mike Hurley, and Judy A. Kennedy, *Dynamics of topologically generic homeomorphisms*, Mem. Amer. Math. Soc. **164** (2003), no. 783. MR MR1980335 (2004j:37024)

[4] Joseph Auslander, *Minimal flows and their extensions*, North-Holland Mathematics Studies, vol. 153, North-Holland Publishing Co., Amsterdam, 1988, Notas de Matemática [Mathematical Notes], 122. MR MR956049 (89m:54050)

[5] John Banks, *Regular periodic decompositions for topologically transitive maps*, Ergodic Theory Dynam. Systems **17** (1997), no. 3, 505–529. MR MR1452178 (98d:54074)

[6] Louis Block and James Keesling, *A characterization of adding machine maps*, Topology Appl. **140** (2004), no. 2-3, 151–161. MR MR2074913 (2005f:37032)

[7] Alexander M. Blokh, *The "spectral" decomposition for one-dimensional maps*, Dynamics reported, Dynam. Report. Expositions Dynam. Systems (N.S.), vol. 4, Springer, Berlin, 1995, pp. 1–59. MR MR1346496 (96e:58087)

[8] Jorge Buescu and Ian Stewart, *Liapunov stability and adding machines*, Ergodic Theory Dynam. Systems **15** (1995), no. 2, 271–290. MR MR1332404 (96i:58083)

[9] Ethan M. Coven, Marcus Pivato, and Reem Yassawi, *Prevalence of odometers in cellular automata*, Proc. Amer. Math. Soc. **135** (2007), no. 3, 815–821 (electronic). MR MR2262877

[10] Fatemeh Helen Ghane and Abbas Fakhari, *Totally chain-transitive attractors of generic homeomorphisms are persistent*, Bull. Korean Math. Soc. **42** (2005), no. 3, 631–638. MR MR2162219 (2007f:37021)

[11] Martin Golubitsky and Ian Stewart, *The symmetry perspective*, Progress in Mathematics, vol. 200, Birkhäuser Verlag, Basel, 2002. MR MR1891106 (2003e:37068)

[12] Morris W. Hirsch and Mike Hurley, *Connected components of attractors and other stable sets*, Aequationes Math. **53** (1997), no. 3, 308–323. MR MR1444181 (98f:58131)

[13] John C. Holladay and Richard S. Varga, *On powers of non-negative matrices*, Proc. Amer. Math. Soc. **9** (1958), 631–634. MR MR0097416 (20 #3885)

[14] Fern Y. Hunt, *Finite precision representation of the Conley decomposition*, J. Dynam. Differential Equations **13** (2001), no. 1, 87–105. MR MR1822213 (2002c:37037)

[15] Mike Hurley, *Generic homeomorphisms have no smallest attractors*, Proc. Amer. Math. Soc. **123** (1995), no. 4, 1277–1280. MR MR1273500 (95e:58109)

[16] W. D. Kalies, K. Mischaikow, and R. C. A. M. VanderVorst, *An algorithmic approach to chain recurrence*, Found. Comput. Math. **5** (2005), no. 4, 409–449. MR MR2189545 (2006i:37035)

[17] A. Ya. Khinchin, *Continued fractions*, The University of Chicago Press, Chicago, Ill.-London, 1964. MR MR0161833 (28 #5037)

[18] Marcin Kulczycki, *Adding machines as invariant sets for homeomorphisms of a disk into the plane*, Ergodic Theory Dynam. Systems **26** (2006), no. 3, 783–786. MR MR2237478

[19] Douglas Lind and Brian Marcus, *An introduction to symbolic dynamics and coding*, Cambridge University Press, Cambridge, 1995. MR MR1369092 (97a:58050)

[20] L. Markus and Kenneth R. Meyer, *Periodic orbits and solenoids in generic Hamiltonian dynamical systems*, Amer. J. Math. **102** (1980), no. 1, 25–92. MR MR556887 (81g:58013)

[21] Ian Melbourne, Michael Dellnitz, and Martin Golubitsky, *The structure of symmetric attractors*, Arch. Rational Mech. Anal. **123** (1993), no. 1, 75–98. MR MR1218687 (94m:58141)

[22] Michael Misiurewicz, *Remark on the definition of topological entropy*, Dynamical systems and partial differential equations (Caracas, 1984), Univ. Simon Bolivar, Caracas, 1986, pp. 65–67. MR MR882013 (88g:58163)

[23] Zbigniew Nitecki, *Topological dynamics on the interval*, Ergodic theory and dynamical systems, II (College Park, Md., 1979/1980), Progr. Math., vol. 21, Birkhäuser Boston, Mass., 1982, pp. 1–73. MR MR670074 (84g:54051)

[24] George Osipenko and Stephen Campbell, *Applied symbolic dynamics: attractors and filtrations*, Discrete Contin. Dynam. Systems **5** (1999), no. 1, 43–60. MR MR1664473 (99j:58135)

[25] Yakov B. Pesin, *Dimension theory in dynamical systems*, Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1997. MR MR1489237 (99b:58003)

[26] J. L. Ramírez Alfonsín, *The Diophantine Frobenius problem*, Oxford Lecture Series in Mathematics and its Applications, vol. 30, Oxford University Press, Oxford, 2005. MR MR2260521

[27] Clark Robinson, *Dynamical systems*, Studies in Advanced Mathematics, CRC Press, Boca Raton, FL, 1995. MR MR1396532 (97e:58064)

[28] Run-sheng Yang, *Topological Anosov maps of non-compact metric spaces*, Northeast. Math. J. **17** (2001), no. 1, 120–126. MR MR1848409 (2002e:54024)

DICKINSON COLLEGE, CARLISLE, PA 17013
E-mail address: richesod@dickinson.edu

AGNES SCOTT COLLEGE, DECATUR, GA 30030
E-mail address: jwiseman@agnesscott.edu