

Topological and symbolic dynamics for hyperbolic systems with holes

STEFAN BUNDFUSS†, TYLL KRÜGER‡ and SERGE TROUBETZKOY§

† Department of Mathematics, Technische Universität Darmstadt, Germany
(e-mail: bundfuss@mathematik.tu-darmstadt.de)

‡ FB Mathematik, Technische Universität Berlin and Fakultät für Physik, Universität
Bielefeld, Germany
(e-mail: tkrueger@math.tu-berlin.de)

§ Centre de Physique Théorique, Fédération de Recherche des Unités Mathématiques de
Marseille, Institut de Mathématiques de Luminy, and Université de la Méditerranée,
Marseille, France
(e-mail: troubetz@iml.univ-mrs.fr)

(Received 9 July 2009 and accepted in revised form 16 June 2010)

Abstract. We consider an axiom A diffeomorphism or a Markov map of an interval and the invariant set Ω^* of orbits which never falls into a fixed hole. We study various aspects of the symbolic representation of Ω^* and of its non-wandering set Ω^{nw} . Our results are on the cardinality of the set of topologically transitive components of Ω^{nw} and their structure. We also prove that Ω^* is generically a subshift of finite type in several senses.

1. Introduction

Let f be an axiom A diffeomorphism of a compact s -dimensional Riemannian manifold M . We cut an open hole H (with a finite number of simply connected components) out of M and consider the invariant set of non-wandering points whose orbits (forward and backward) never fall in the hole. In part of this article we will consider non-invertible interval maps, in this case we consider the set of points whose forward orbit never falls in the hole. We fix a Markov partition, which yields a natural representation of the non-wandering set as a subshift of finite type. We will consider the subshift corresponding to the set of points which do not fall in the hole; we call this subshift an *exclusion subshift*. We are interested in several questions about the topological structure of exclusion subshifts and their non-wandering sets. More precisely, for which holes are exclusion shifts subshifts of finite type, sofic shifts, or coded systems? As we vary the hole, what is the typical type of an exclusion subshift? How many transitive components can the non-wandering set of an exclusion shift have?

After giving background definitions in §2, we show in §3 that every subshift of finite type (SFT) and every β -shift is an exclusion shift, but certain sofic systems are not. In particular, there are exclusion shifts that are sofic and ones that are not sofic. In this section

we also show that the baker's map and multiplication by n are universal in the sense that any rectangular (respectively interval) exclusion subshift can be constructed there. In §4 we give a criterion for when the hole leads to an SFT. We apply this criterion to several different classes of holes (rectangles, polyhedra, holes with continuous boundary) with corresponding topologies/measures to show that the 'typical' hole leads to an SFT. In §5, we compare the non-wandering set for the orbits which do not fall in the hole and the non-wandering set of the exclusion shift. In §6 we prove the finiteness of the number of transitive components in the one-dimensional case and give an explicit upper bound in terms of the number of holes. In §7 we show that an exclusion subshift has at most countably many topologically transitive components with dense periodic points and that exclusion subshifts are always coded systems on such topologically transitive components.

Dynamical systems with holes have been studied extensively by physicists. Cvitanovich and his coworkers have investigated how to characterize the exclusion subshift in various settings (see the survey [13]). They have introduced the notion of a pruning front which corresponds to ∂H in our setting. Some aspects of Cvitanovich's work have been carried out by Carvalho in a mathematical framework for Smale's horseshoe [7]. Another kind of question about dynamical systems with holes has also been extensively studied by physicists, namely construction of a physical semi-invariant measure and the understanding of the speed of mass disappearance into the holes (the escape rate formula) [19]. Starting with the work of Pianigiani and Yorke [24], a series of mathematical works have confirmed the expectations of the physicists in many settings (see the survey [15] for details). Bunimovich and Yurchenko have studied an aspect of this question close to the spirit of our article [6].

2. Definitions

In this article, we consider axiom A diffeomorphisms $f : M \rightarrow M$ of compact s -dimensional ($s \geq 2$) Riemannian manifolds M , i.e., the non-wandering set Λ of f is hyperbolic and periodic points are dense in Λ . Sometimes we will, for the purposes of illustration, also consider the n -fold baker's map which is not formally axiom A since it is not a diffeomorphism. The n -fold baker's map $B : [0, 1]^2 \rightarrow [0, 1]^2$ is defined by $B(x, y) = (nx \bmod 1, (y + i)/n)$ if $x \in [i/n, (i + 1)/n]$. We will call the 'partition' $\{[i/n, (i + 1)/n] : i = 0, 1, \dots, n - 1\}$ the standard Markov partition.

Consider the alphabet $\{1, 2, \dots, n\}$ and a finite collection of forbidden words of length m (m fixed). The subset of all sequences in $\{1, \dots, n\}^{\mathbb{Z}}$ where the forbidden words never appear is called a *subshift of finite type*. In the one-sided case we call this a one-sided subshift of finite type.

In this article, we will consider a Markov partition \mathcal{P} and the resulting coding by an SFT $(\Sigma^{\mathcal{P}}, \sigma)$; see [21] for a definition. A Markov partition is called proper if each element of the partition is the closure of its interior. Let $\pi : \Sigma^{\mathcal{P}} \rightarrow M$ be the projection map.

Throughout this article the word *hole* refers to an open set $H \subset M$ whose boundary $\partial H = \bar{H} \setminus H$ consists of a finite union of compact topological codimension-one manifolds and $H = \text{int}(\bar{H})$. In particular, this implies that H has a finite number of simply connected components. Fix a hole H and consider the invariant set $\Omega^* = \Omega_H^*$ of points whose orbit (forward and backward) never falls in the hole.

For a continuous map g of a compact topological space M , a point $x \in M$ is called *non-wandering* if for any neighborhood \mathcal{U} of x there exists a positive integer n such that $g^n \mathcal{U} \cap \mathcal{U} \neq \emptyset$. The set of non-wandering points of g is always closed, invariant and non-empty. The non-wandering set of $f|_{\Omega^*}$ is denoted by Ω^{nw} . Let $\Sigma^* = \Sigma_H^* = \pi^{-1}\Omega^*$ with the following convention[†]: if $x \in \Omega^*$ is on the boundary of the hole and $\pi^{-1}(x)$ is not unique, then we only consider those preimages $\pi^{-1}(x)$ to be in Σ^* which can be approximated from outside the closure of H , i.e., those $s \in \pi^{-1}(x)$ such that there exists $\{x_j\} \subset M \setminus \overline{H}$ such that $x_j \rightarrow x$ and there exists $s_j \in \pi^{-1}(x_j)$ with $s_j \rightarrow s$. We call the subshift Σ^* an *exclusion subshift*. Furthermore, let Σ^{nw} be the set of non-wandering points of $\sigma|_{\Sigma^*}$. As we will see in §5, $\pi(\Sigma^{\text{nw}}) = \Omega^{\text{nw}}$ and thus Σ^{nw} is the non-wandering set of $\pi^{-1}(\Omega^{\text{nw}})$.

If M is two-dimensional, we call a hole H a *rectangle-like hole* if it is the finite union of holes with local product structure, i.e., $H = \bigcup_{i=1}^n H_i$ such that for any $x, y \in H_i \cap \Lambda$ the intersection of the local stable manifold $W_{\text{loc}}^s(x)$ with the local unstable manifold $W_{\text{loc}}^u(y)$ is a unique point belonging to H_i . Two holes H, H' are called *equivalent* if $\Omega_H^{\text{nw}} = \Omega_{H'}^{\text{nw}}$. It is easy to show that any rectangle-like hole is equivalent to one whose boundary consists of a finite union of local stable and unstable manifolds. We call the corresponding exclusion shift a *rectangle exclusion shift* (RES).

We will also consider the one-dimensional situation, where M is an interval or a circle. We will consider maps which we will call Markov maps (the assumptions may be somewhat stronger than the usual usage of this term in other articles). A *Markov map* is a topologically transitive, uniformly expanding continuous map of the circle or piecewise continuous map of the interval which, when considered as a map of the circle, is continuous and which has a finite generating Markov partition \mathcal{P} whose elements consist of intervals. Since the map is not invertible we only require that the forward orbit never falls into a hole. If the hole consists of a finite union of intervals, we call the corresponding exclusion shift an *interval exclusion shift* (IES). Our standard example in this framework is the multiplication by a positive integer $f(x) = nx \bmod 1$.

Since we study only topological properties of exclusion shifts, the uniformly expanding assumption is convenient, but not necessary. It can be replaced by a topological conjugacy to a uniformly hyperbolic map. This topological conjugacy takes a hole which is an interval to an interval and thus preserves the notion of interval exclusion shift. Similarly, the assumption of axiom A in higher dimensions can be replaced by a topological conjugacy to an axiom A map; this preserves the notion of open hole.

3. SFTs and sofic systems

PROPOSITION 3.1. *Every SFT is an exclusion subshift.*

Proof. Consider an SFT Σ . Let n be the cardinality of the alphabet of Σ and m the length of the forbidden blocks. Consider a horseshoe map f with n -branches. We consider the standard Markov partition for f . Define H' to be the (finite) union of Markov rectangles (for f^m) which correspond to the forbidden m -blocks which define Σ . Markov rectangles are by definition closed. Since a horseshoe is a Cantor set it is totally disconnected, thus we

[†]This convention is not needed when the coding is always unique, for example for horseshoes.

can find an open hole H containing H' such that the intersection of H with the horseshoe is exactly H' . This yields Σ as an exclusion subshift. \square

Remark. The same construction works for the n -fold baker's map, except that we define the hole H to be the interior of H' and use the coding convention from §2.

The following theorem shows that the baker's map is 'universal' for rectangle-like holes and multiplication by n is 'universal' for interval holes.

THEOREM 3.2.

- (1) Any subshift which is representable as an RES for some axiom A diffeomorphism is representable as an RES for the baker's map.
- (2) Any subshift which is representable as an interval exclusion subshift for some Markov map is representable as an interval exclusion subshift for multiplication by n .

Proof. The relationship between the various objects constructed in this proof are summarized as follows:

$$\begin{array}{ccc} \Omega_{H,g}^* & \xleftarrow{\pi} & \Sigma^* \xrightarrow{\pi'} \Omega_{H'',B}^* \\ \cap & \cap & \cap \\ \Lambda_g & \xleftarrow{\pi} & \Sigma \xrightarrow{\pi'} \Omega_{H',B}^* \end{array} \quad (3.1)$$

Fix the axiom A system $g : M \rightarrow M$, the Markov partition \mathcal{Q} yielding an SFT Σ , and the rectangle-like hole $H \subset M$ yielding the RES Σ^* . Let $B : [0, 1]^2 \rightarrow [0, 1]^2$ be the baker's map with the standard Markov partition \mathcal{P} .

The idea of the proof is quite simple. By Proposition 3.1 we can find a rectangle-like hole for the baker's map which yields Σ as an RES for B . Then lift the hole H to the Σ and then project to a hole for the baker's map.

Denote the projection $\pi : \Sigma \rightarrow \Lambda_g$ where Λ_g is the non-wandering set of g . The set $\Omega_{H,g}^*$ of points whose g orbit never falls into H by definition satisfies $\Sigma^* = \pi^{-1}(\Omega_{H,g}^*)$.

Remember that $H = \cup H_i$ with each hole H_i having product structure. Let $H_{i,j}$ be the interior of the intersection of H_i with the j th element of the Markov partition; note that $H \subset \cup H_{i,j}$, and if some H_i intersects several Markov partition elements this inclusion is strict. Nevertheless, the coding convention implies that $\Omega_{\cup H_{i,j},g}^* = \Omega_{H,g}^*$.

By Proposition 3.1 we can find a rectangle-like hole $H' \subset [0, 1]^2$ which yields Σ as an RES for (B, \mathcal{P}) . Let $\Omega_{H',B}^*$ be the set of points whose B orbit never falls into H' and let $\pi' : \Sigma \rightarrow \Omega_{H',B}^*$ be the projection.

Consider

$$H_{i,j}^{\text{symb}} := \{s \in \Sigma : \pi(s) \in H_{i,j}\} = \pi^{-1}(H_{i,j}).$$

Note that $H_{i,j}^{\text{symb}}$ is open since π is continuous and $H_{i,j}$ is open. Let

$$H''_{i,j} := \pi'(H_{i,j}^{\text{symb}}) \quad \text{and} \quad H'' = \bigcup_{i,j} H''_{i,j}.$$

Clearly the set of points $\Omega_{H'',B}^*$ whose B orbit never falls into H'' satisfies $\Omega_{H'',B}^* \subset \pi'(\Sigma^*)$. With our coding convention we have $\pi'^{-1}(\Omega_{H'',B}^*) = \Sigma^*$, i.e. Σ^* is an RES for B with respect to the hole H'' .

The proof in the one-dimensional case is similar. \square

Let Σ_{even}^+ be the one-sided even shift, i.e., the set of all 0–1 half-infinite sequences with the constraint that the number of consecutive ones that occur in between two zeros is always even.

PROPOSITION 3.3. *The one-sided even shift Σ_{even}^+ is not an interval exclusion system.*

Proof. By Theorem 3.2 it is enough to consider the doubling map, i.e. $M := \Sigma^1$, $f(x) := 2x \bmod 1$, $\mathcal{P} = \{[0, \frac{1}{2}], [\frac{1}{2}, 1]\}$. Suppose we have a hole which is a union of intervals and yields Σ_{even}^+ as an exclusion system. Consider the point $x_n = 01^{2n+1}0^\infty$. The whole orbit of x_n , except x_n itself, is in Σ_{even}^+ . Thus x_n must lie in the hole.

Let $y_n = 01^{2n}0^\infty$. Clearly $y_n \in \Sigma_{\text{even}}^+$ and $x_{n-1} < y_n < x_n$. Thus the hole must consist of an infinite number of intervals, i.e., Σ_{even}^+ cannot be an IES. \square

Let Σ_{even} be the two-sided even shift, i.e., the set of all 0–1 bi-infinite sequences with the constraint that the number of consecutive ones which occur in between two zeros is always an even number.

PROPOSITION 3.4. *The even shift Σ_{even} is not an RES.*

Proof. By Theorem 3.2 it is enough to consider the twofold baker's map B . Let $x_n = 0^\infty.1^{2n+1}0^\infty$ (here the decimal point marks the position between the -1 th and 0 th elements of the sequence). The point x_n is not in the even shift; it must fall into the hole under some iteration of B .

We treat several cases. First of all, suppose that x_n falls into the hole at the boundary of M . Then at the instant that x_n falls into the hole all the 1s are to the right (or all are to the left) of the decimal point (i.e., $B^j x_n \in \partial M$ with $j \leq 0$ or $j \geq 2n + 1$). Note that the intersection of the rectangular holes with the boundary consists of a finite union of intervals. Thus we can apply the argument of the previous example to conclude that it is impossible to have an infinite number of the x_n fall into the hole when they are on the boundary of M .

In other words, all but finitely many x_n fall into the hole away from the boundary of M . For such an x_n consider the code a at the instance of falling into the hole. It has the form $0^\infty 1^p.1^q 1^{2n+1}0^\infty$ where $p + q = 2n$.

Consider the sequence $(u_i)_{i \in \mathbb{Z}}$ with $u_i = 0^\infty 1^{2i} 01^{p-1} 1^q 1^{2n+1}0^\infty$ and the sequence $(v_i)_{i \in \mathbb{Z}}$ with $v_i = 0^\infty 1^p.1^q 01^{2i} 0^\infty$. These two sequences get arbitrary near from the left and the bottom to a and all their elements are in the even shift. Thus a must be a corner of the hole. This contradicts the assumption that the number of corners of the hole is finite. \square

A closed shift invariant subset of $\{1, 2, \dots, n\}^{\mathbb{Z}}$ (respectively $\{1, 2, \dots, n\}^{\mathbb{N}}$) is called *sofic* if it is a factor of an SFT. The even shift is an example of a sofic system. A subshift Σ is called a *coded* system if it can be represented by an irreducible countable labeled graph [3]. We do not know if every sofic system—or even every coded system—is an exclusion subshift. However, if we drop the requirement that the boundary consists of codimension-one manifolds then we get that *all* subshifts are exclusion subshifts in a weaker sense than considered in this article: they are the complement of an open hole without any further assumptions on the boundary. To see this, simply define the hole to be the complement of the invariant set which gives rise to the subshift.

Consider the dynamical system $h_\beta : [0, 1] \rightarrow [0, 1]$ given by $h_\beta x = \beta x \bmod 1$. Suppose that $\beta \in (1, 2]$ and consider the ‘partition’ $I_0 := [0, \beta^{-1}]$ and $I_1 := [\beta^{-1}, 1]$. Given $\omega, \alpha \in \{0, 1\}^{\mathbb{N}}$, we say that $\omega < \alpha$ if there exists a positive integer N such that $\omega_i = \alpha_i$ for $i = 0, 1, \dots, N - 1$ and $\omega_N = 0 < \alpha_N = 1$. It is well known that the code $\hat{\beta}$ of the orbit of β satisfies $\sigma^n \hat{\beta} < \hat{\beta}$ for all $n > 0$. The set of all sequences which are the code of some orbit is called the *one-sided β -shift*, and is characterized by [2] as

$$\{x \in \{0, 1\}^{\mathbb{N}} : \sigma^n x < \hat{\beta} \forall n > 0\}.$$

The natural extension of the one-sided β -shift is called the *two-sided β -shift*.

THEOREM 3.5. *Every one-sided β -shift is an IES and every two-sided β -shift is an RES.*

Proof. Let $g : [0, 1] \rightarrow [0, 1]$ be defined by $g(x) = 2x \bmod 1$. Inspired by the above, a number β is called a β -number if $g^n \beta < \beta$ for all $n > 0$. If β is a β -number then the set $X_\beta := \{x \in [0, 1] : g^n x < \beta \forall n > 0\}$ is conjugate (up to a countable number of points) to the one-sided β -shift. If $H \subset [0, 1]$ is the interval $(\beta, 1)$ then the set of points whose forward g orbit never falls in H is exactly X_β .

Now let f be the baker’s map. Fix a β -number β and let $H \subset M$ be the rectangular hole $\{(x, y) : x > \beta, 0 \leq y \leq 1\}$. Let $(x_{-i}, y_{-i}) = f^{-i}(x, y)$. Since the hole stretches from the bottom to the top of M it is easy to see that $(x, y) \in \Omega^*$ if and only if $g^n(x_{-i}) < \beta$ for all $n \geq 0$ and all $i \geq 0$. Thus the exclusion shift in this example is precisely the two-sided β -shift. We remark that the two-sided β -shift is of finite type, sofic, or not sofic if and only if the one-sided one has the same property. It is well known that the β -shift is of finite type if the binary expansion of β is finite, it is sofic if it eventually periodic, and otherwise it is not sofic [2]. \square

Remark. The notion of β -shifts can be extended to $\beta > 2$. The theorem extends easily to this case; for $\beta \in (n - 1, n]$ we must use the map $g = nx \bmod 1$ and the n -fold baker’s map.

4. Genericity results

4.1. A criterion for SFTs. In this subsection we give the criterion that will be applied throughout the rest of the section to prove genericity of SFTs in various settings.

PROPOSITION 4.1. *If for each $x \in \partial H$ there is an i such that $f^i x \in H$ then Σ^* is an SFT.*

Proof. Fix a generating Markov partition \mathcal{P} and let $\mathcal{P}^{(n)} := \bigvee_{i=-n}^n f^i \mathcal{P}$. Let

$$P_x^{(n)} := \bigcup_{\{P \in \mathcal{P}^{(n)} : x \in P\}} P.$$

If $x \in \partial H$ and $f^i x \in H$ then, since \mathcal{P} is generating by continuity, there is an $n(x)$ such that $f^i P_x^{(n(x))} \subset H$. Since ∂H is compact we can cover ∂H by a finite collection of the sets $P_x^{(n(x))}$ to obtain a neighborhood N of ∂H such that $N \cap \Omega = \emptyset$. Since we used a finite collection of P_x , the hole $H' := N \cup H$ consists of a finite union of elements of $\mathcal{P}^{(N)}$ for some sufficiently large integer N and thus Σ^* is an SFT. \square

4.2. Results in the Hausdorff metric. Let s denote the Hausdorff dimension of M . Consider the set C of all holes such that ∂H is a continuous, i.e., there is a continuous map $h : \mathbf{S}^{s-1} \rightarrow M$ whose image is ∂H . For $H \in C$, let $H^\epsilon := \bigcup_{x \in \partial H} B(x, \epsilon)$. For $H_1, H_2 \in C$, we define

$$d(H_1, H_2) := \inf\{\epsilon > 0 : H_1 \subset H_2^\epsilon, H_2 \subset H_1^\epsilon\}.$$

A set is called *totally disconnected* if no two points are in the same connected component.

LEMMA 4.2. *If the set Ω_H is totally disconnected then, for every $\epsilon > 0$, there is a hole $H' \in C$ with $d(H', H) < \epsilon$ and an open neighborhood $\mathcal{U} \subset C$ of H' such that for all holes $H'' \in \mathcal{U}$ the subshift Σ^* is an SFT.*

Proof. Consider the set $\mathcal{V} := \{H' \in C : H \subset H', d(H, H') < \epsilon\}$. Since Ω_H is totally disconnected there are holes $H' \in \mathcal{V}$ such that $\partial H' \cap \Omega_H = \emptyset$. Since $\Omega_H \subset \Omega_{H'}$ such an H' satisfies the requirements of Proposition 4.1 and so defines an SFT.

Next we will show that an open set \mathcal{U} of holes satisfies the requirements of Proposition 4.1. We claim that we can find an open set $O \subset H'$ with $d(O, H') > 0$ and a positive integer N such that for all $x \in \partial H$ there exists an integer i , satisfying $|i| \leq N$ such that $f^i x \in O$. To see this, note that for each $x \in \partial H'$ there is an $n(x) > 0$ such that $f^{n(x)} \in H'$. By continuity of f we can choose $\epsilon(x) > 0$ so small that, for all $y \in B(x, \epsilon(x))$, we have $f^{n(x)} y \in H'$ and $d(f^{n(x)} y, \partial H') d(f^{n(x)} x, \partial H') / 2$. By compactness of $\partial H'$ we can cover it by a finite number of such balls $\{B(x_i, \epsilon(x_i))\}$. Then O is $\cup f^{n(x_i)} B(x_i, \epsilon(x_i))$.

Then, just as in the proof of Proposition 4.1, there is a neighborhood U of $\partial H'$ such that, for each $x \in U$, for some i satisfying $|i| \leq N$ we have $f^i x \in O$. This immediately implies that we can choose a small neighborhood $\mathcal{U} \subset O$ of H' such that, for any hole $H'' \in \mathcal{U}$, the boundary of H'' satisfies the requirements of Proposition 4.1 and thus the corresponding shift Σ'' is an SFT. \square

THEOREM 4.3. *In the two-dimensional case the set of $H \in C$ for which Σ^* is an SFT is open and dense.*

Proof. Consider an arbitrary hole $H \in C$ and the associated invariant set Ω_H . If x is a generic point in the sense that it visits (in both forward and backward time) any cylinder set (defined by the Markov partition) with the correct frequency, then $W^s(x)$ and $W^u(x)$ completely fall into H and thus are both disjoint from Ω_H . Both $W^s(x)$ and $W^u(x)$ are curves that are dense in M , thus since M is two-dimensional the complement of their union is totally disconnected. Since Ω_H is a subset of this set it is also totally disconnected. Applying Lemma 4.2 finishes the proof. \square

4.3. Rectangle-like holes. In this section we assume that M is two-dimensional, that f is a C^2 axiom A attractor and that μ is its SBR measure, i.e., μ has absolutely continuous conditional measures on unstable manifolds. Recall that a hole is a rectangle-like hole if ∂H consists of a finite number of curves, each of which is a piece of a stable or unstable manifold of f . We assume that at each point where two curves meet, one is stable and the other is unstable. Thus the number of corners is always even, and if we fix an orientation of the boundary we only need to give the coordinates of every other corner point to describe

a rectangle-like hole. If we fix the number $2n$ of corners of a rectangle-like hole ($n \geq 2$) then we can parametrize the set of all such rectangles by an open subset $R^{(n)}$ of M^{2n} . We will consider the Lebesgue measure on $R^{(n)}$.

THEOREM 4.4. *For every $n \geq 2$, the set of rectangle-like holes with $2n$ corners for which Σ^* is an SFT is of full Lebesgue measure and contains an open dense subset of $R^{(n)}$.*

Similarly, for the one-dimensional case one can show that for any k the set of holes which are a union of k -intervals for which Σ^+ is an SFT is of full Lebesgue measure and contains an open dense subset of the set of all such holes.

Proof. Consider the set G of generic points in the same sense as in the proof of Theorem 4.3. The set G is of full Lebesgue measure. Suppose that H is a rectangle-like hole with the property that every other corner point (those which are noted in the description of $H \in R^{(n)}$) is in G . Such a hole satisfies the requirements of Proposition 4.1 and thus Ω is an SFT. Clearly the set of such holes is dense and of full measure. The proof of openness is the same as in the proof of the previous theorem. \square

4.4. Polyhedral holes. In this section we suppose that M is s -dimensional and has a flat structure (i.e., $M \subset \mathbf{R}^s$ or $M \subset \mathbf{T}^s$). We consider holes which are the interior of arbitrary polyhedra, i.e., whose boundary consists of straight line segments. Fix the number of corners n . The set of n -gons is an open subset of M^n which we denote by $P^{(n)}$. We consider the Lebesgue measure on M^n . We call a polyhedral hole $H \in P^{(n)}$ *large* if the Hausdorff dimension of the associated invariant set Ω is strictly less than one. Let $\hat{P}^{(n)} \subset P^{(n)}$ be the set of all large polyhedral holes.

THEOREM 4.5. *For every n the set of large polyhedral holes for which Σ^* is an SFT is of full Lebesgue measure in $\hat{P}^{(n)}$ and contains an open dense subset of $\hat{P}^{(n)}$.*

For natural families of maps (hyperbolic toral automorphisms, baker's maps, horseshoes) it is not hard to see that the set of large polygonal holes for which Ω is uncountable has positive measure for all even n . For horseshoes we can show this for all n .

Proof. We note that if $H \subset H'$ then $\Sigma_{H'}^* \subset \Sigma_H^*$. Thus if H is a large hole then so is H' . Our proof is local. Fix an open set B . Consider the set Σ_B^* of points that never fall into B . Suppose that B is large enough that $a := \dim(\Omega_B) < 1$. This implies that $\dim(\text{proj}_\theta(\Omega_B)) = a < 1$ for Lebesgue almost every $\theta \in \mathbf{S}^{s-1}$ [18, Theorem 6.9]. We will call such θ generic. Here proj_θ denotes the orthogonal projection from M onto L_θ , the line through the origin in the direction $\theta \in \mathbf{S}^{s-1}$.

Consider the set of codimension-one hyperplanes with normal direction θ and parametrize these hyperplanes by a parameter $t \in R \equiv l_\theta$. For any generic θ , since $a < 1$ the set $\text{proj}_\theta(\Omega_B)$ is totally disconnected and thus for a.e. t , the hyperplane with parameter t will be disjoint from Ω_B . We will also refer to this hyperplane as generic.

Now consider any polyhedron H that contains B in its interior and for which all the faces are contained in generic hyperplanes. The faces of this polyhedron do not intersect the set Ω_B . Thus we can apply Proposition 4.1 to conclude that Σ^* is an SFT. The set

of such polyhedra is locally of full $d\theta \times dt$ measure. It is easy to conclude that it is also locally of full Lebesgue measure. Sets of full measure are also locally dense. Lemma 4.2 implies that it is also open. \square

Remark. We believe that the set of polyhedral holes defining an SFT is of full measure and contains an open and dense set. Our strategy of proof cannot be used since the projection onto a line L_θ of a set of dimension greater than one has positive one-dimensional measure for a.e. θ .

5. Transitive components of Σ^{nw}

A homeomorphism f of a compact metric space M is called *topologically transitive* if there exists a dense orbit, or equivalently if every proper closed f -invariant subset is nowhere dense. A closed invariant set X (i.e., $f^{-1}X = fX = X$) is called a *topologically transitive component* if X contains a dense orbit and there is no closed invariant set $X' \supsetneq X$ containing a dense orbit. Note that for non-invertible maps studied in the last section we will modify this definition.

In [21] a different definition of ‘topologically transitive’ is used—there f is topologically transitive if there exists a point whose forward and backward orbits are both dense. We call this notion *strong topological transitivity*. In particular, they prove that a subshift is strongly topologically transitive if and only if it is irreducible (i.e., for any pair of cylinders U, V there exists n_0 such that $\sigma^n U \cap V \neq \emptyset$).

We call a subshift a *non-wandering subshift of finite type* if it is the non-wandering set of a subshift of finite type. In this section we are interested in the non-wandering set Σ^{nw} of Σ^* and will be considering the approximation of Σ^{nw} by non-wandering subshifts of finite type. The non-wandering set of a subshift of finite type is the (finite) disjoint union of irreducible subshifts of finite type. Each of these components is strongly topologically transitive, and thus topologically transitive. Thus the decomposition into topologically transitive components for a non-wandering subshift of finite type is equivalent to the usual decomposition into irreducible components. In particular, in this case, strong and normal topological transitivity are equivalent. We illustrate this with an example. Consider the SFT Σ^* defined in Figure 1(a). The SFT Σ^* has infinitely many topologically transitive components: for each $j \geq 1$ the orbit of the point $x_j := 1^\infty 2^j \cdot 3^\infty$ is a dense orbit in the transitive component $O(x_j) \cup \{1^\infty\} \cup \{3^\infty\}$. There are two other transitive components, $O(1^\infty \cdot 2^\infty) \cup \{1^\infty\} \cup \{2^\infty\}$ and $O(2^\infty \cdot 3^\infty) \cup \{2^\infty\} \cup \{3^\infty\}$. There are only three strong topological transitive components of Σ^* —the three fixed points which form the non-wandering set Σ^{nw} shown in Figure 1(b). These three fixed points are the usual topologically transitive components of Σ^{nw} ; they are disjoint.

Furthermore, periodic points are dense in a topologically transitive component of a non-wandering SFT or in a topologically transitive component of an axiom A diffeomorphism.

Returning to the general case, let F^+ (respectively F^-) be the set of points whose forward (respectively backward) images do not fall into a hole. Remember that $\Omega^* = F^+ \cap F^-$.

PROPOSITION 5.1. $NW(f|_{F^+}) \subset \Omega^*$ and $NW(f^{-1}|_{F^-}) \subset \Omega^*$.




FIGURE 1. (a) An SFT. (b) Its non-wandering set.

Proof. Fix $x \in NW(f|_{F^+})$. Suppose that $x \notin F^-$, thus we can choose $i > 0$ to be the first time such that $f^{-i}x$ is in the hole. Choose a neighborhood U of x so small that $f^{-i}U$ is contained in the hole. Choose $n = n(U) > 0$ such that $f^n(U \cap F^+) \cap (U \cap F^+) \neq \emptyset$. Choose a point y in this intersection; i.e., y and $f^{-n}y$ are in $U \cap F^+$. Let $z = f^{-n}y$; then z and $f^n z$ are in $U \cap F^+$. By choosing U even smaller we may assume that $n > i$. Since $f^n z \in U$ and $f^{-i}U \subset H$ we get $f^{-i}(f^n z) = f^{n-i}z \in H$. This is a contradiction since $z \in F^+$ and $n - i > 0$. Thus $x \in F^-$. The other containment is proven similarly. \square

Throughout the rest of the article the projection of a symbolic point will be denoted by the same letter without a hat, for example $x = \pi(\hat{x})$.

PROPOSITION 5.2. $\Sigma^{\text{nw}} = \pi^{-1}(\Omega^{\text{nw}})$.

Proof. Consider $\hat{x} \in \Sigma^{\text{nw}}$. Since \hat{x} is non-wandering, for any $n < 0 < m$ there exists $j > 0$ such that the cylinder $\mathcal{U} := [\hat{x}]_n^m$ satisfies $\sigma^j \mathcal{U} \cap \mathcal{U} \neq \emptyset$. Let \hat{y} be a point of this intersection. For any neighborhood \mathcal{V} of $x := \pi(\hat{x})$, we can find $n < 0 < m$ such that the cylinder $\pi(\mathcal{U}) \subset \mathcal{V}$. Thus $y := \pi(\hat{y}) \in f^j \mathcal{V} \cap \mathcal{V}$. It follows that $\pi(\Sigma^{\text{nw}}) \subset \Omega^{\text{nw}}$.

To finish we need to show that π is onto. Choose any $x \in \Omega^{\text{nw}}$. First suppose that there is a unique $\hat{x} \in \Sigma^*$ such that $\pi(\hat{x}) = x$. By the non-wandering of (Ω^{nw}, f) for any neighborhood \mathcal{V} of x there exists $j > 0$ such that $f^j \mathcal{V} \cap \mathcal{V} \neq \emptyset$. Given a cylinder $\mathcal{U} := [\hat{x}]_n^m$, we can choose a neighborhood \mathcal{V} of x such that $\pi^{-1}(\mathcal{V}) \subset \mathcal{U}$. Thus

$$\sigma^j \mathcal{U} \cap \mathcal{U} \supset \pi^{-1}(f^j \mathcal{V} \cap \mathcal{V}) \neq \emptyset$$

and therefore \hat{x} is non-wandering and the claim follows.

Note that for any neighborhood \mathcal{V} of x there are times $0 < j_1 < j_2 < \dots$ such that $f^{j_i} \mathcal{V} \cap \mathcal{V} \neq \emptyset$. Now suppose that the coding of x is not unique, i.e., there are several points $\hat{x}_i \in \Sigma^*$ such that $\pi(\hat{x}_i) = x$. It is well known that the number of such points is uniformly bounded by a constant L , which depends only on the dimension of M . In this case, if we try to argue as in the previous paragraph, we cannot choose \mathcal{V} such that $\pi^{-1}(\mathcal{V}) \subset \mathcal{U}$. For $i \in \{1, \dots, L\}$, let \mathcal{U}_i be a cylinder set containing \hat{x}_i such that $\pi^{-1}\mathcal{V} \subset U := \bigcup_{i=1}^L \mathcal{U}_i$. Thus

$$\sigma^{j_i} \mathcal{U}_i \cap U \supset \pi^{-1}(f^{j_i} \mathcal{V} \cap \mathcal{V}) \neq \emptyset.$$

Note that the point x is on the boundary of the sets $\pi(\mathcal{U}_i)$ for all i . We can assume that all the \mathcal{U}_i are cylinders of the same length; thus we can refine our original Markov partition \mathcal{P} by intersecting with a finite number of Markov partitions $f^i \mathcal{P}$, yielding a new Markov partition (which we will call \mathcal{P}') such that each of the sets $\mathcal{V}_i := \pi(\mathcal{U}_i)$ is an element of \mathcal{P}' . Suppose that x is in the stable boundary of the \mathcal{V}_i s. Then by the Markov property for

each j_i and each l , $f^{j_i}\mathcal{V}_l$ has a Markov intersection with some \mathcal{V}_k , and for fixed j_i the map $l \rightarrow k$ is invertible. Fixing \mathcal{V}_k , we iterate and produce a sequence of Markov intersections $f^{j_i}\mathcal{V}_{l_i}$ with \mathcal{V}_k . By the pigeonhole principle some l must appear as an l_i twice. This implies that $f^{j_{i_1}}\mathcal{V}_l$ has a Markov intersection with $f^{j_{i_2}}\mathcal{V}_l$, but since both of these sets have a Markov intersection with \mathcal{V}_k we can conclude that $\sigma^{j_{i_1}-j_{i_2}}\mathcal{V}_k$ has a Markov intersection with \mathcal{V}_k . Since this is true for any \mathcal{V}_k , we conclude in this case that each of the points \hat{x}_i is non-wandering.

If x is in the unstable bound of the \mathcal{V}_i 's then, since σ is invertible and $\sigma^{j_i}\mathcal{U} \cap \mathcal{U} \neq \emptyset$ we have $\sigma^{-j_i}\mathcal{U} \cap \mathcal{U} \neq \emptyset$. Now we can repeat the above argument. \square

6. Finiteness of topologically transitive components in dimension one

Suppose that M is one-dimensional, i.e., an interval or circle, (f, \mathcal{P}) is as described in §2 and $H = \bigcup_{i=1}^p H_i$ where the H_i are disjoint open intervals. By the continuity assumption on f we can always think of M as the circle. The definitions of Ω^* , Ω^{nw} , Σ^* and Σ^{nw} are similar to the corresponding definitions in the invertible case, with the difference that since f is not invertible, we only require that the forward orbit does not fall into H . The set $(\Omega^{\text{nw}})^c := I \setminus \Omega^{\text{nw}}$ is open. Let H_i^{ext} be the maximal interval containing H_i which is a subset of $(\Omega^{\text{nw}})^c$ and $H^{\text{ext}} = \bigcup_{i=1}^r H_i^{\text{ext}}$. It is possible for several H_i to amalgamate into one H_j^{ext} , thus $r \leq p$. Call H^{ext} the extended hole. We also assume that our Markov partition is fine enough that the left and right end points of each of the H_i^{ext} are in different elements of the Markov partition. This assumption can be made without loss of generality by replacing \mathcal{P} by $\mathcal{P} \vee f\mathcal{P} \vee \dots \vee f^n\mathcal{P}$ for a large enough n .

Since our maps are not invertible, we must modify the definition of topologically transitive component: a closed *forward invariant* (i.e., $f(X) = X$) set is called *topologically transitive* if it has a *forward orbit which is dense* and there is no closed *forward invariant* set $X' \supsetneq X$ containing a dense orbit. We gather together here some simple facts about topologically transitive components of continuous maps which we will use.

PROPOSITION 6.1. *Suppose f is a continuous map of the interval or circle.*

- (1) *If a topologically transitive component contains an isolated point, then this point is periodic and the component coincides with this periodic orbit.*
- (2) *Each topologically transitive component is finite or uncountable.*
- (3) *Any dense orbit in a forward invariant set is recurrent.*
- (4) *Topologically transitive components are non-wandering.*

Proof. (1) Let X denote the topologically transitive component. If z is isolated and the orbit of x is dense, then $f^i x = z$ for some $i \geq 0$. The point x is also isolated; if not, then by continuity z is not isolated. If the orbit of x is not periodic then since x is isolated we have $fX = X \setminus \{x\}$ and X is not forward invariant. This contradiction implies that x is periodic. Since x 's orbit is dense, it must coincide with the topologically transitive component.

(2) We claim that if the set of non-wandering points of a topologically transitive component is not a finite set, then it must be a perfect set (i.e., it equals the set of its limit points). It is well known that the non-wandering set is closed. Since there is a dense forward orbit $\{f^i x : i \geq 0\}$, any non-wandering point is a limit point of this orbit. It is well known that perfect sets are uncountable.

(3) If X is finite then this is clear. Suppose that X is uncountable and the forward orbit of x is dense in X . By part (1) the point x cannot be isolated. Since x is not isolated and the forward orbit of x is dense this orbit must come arbitrarily close to x to be dense, i.e., the orbit is recurrent.

(4) By definition every topologically transitive component X is forward invariant and has a forward dense orbit $\{f^i x : i \geq 0\}$. For any $y \in X$ and any neighborhood \mathcal{U} of y there is a positive i such that $f^i y \in \mathcal{U}$. Since x is recurrent, there is an n such that $f^{i+n} x \in \mathcal{U}$ as well, which shows that $f^n \mathcal{U} \cap \mathcal{U} \neq \emptyset$. \square

LEMMA 6.2. *The set of points whose forward orbit never hits the extended hole H^{ext} is precisely the non-wandering set of the set of points whose forward orbit never hits H , i.e., $\Omega_{H^{\text{ext}}}^* = \Omega_H^{\text{nw}}$.*

Proof. Consider $G := \bigcup_{n \geq 0} f^{-n} H^{\text{ext}}$. The set G consists of a countable union of intervals and is dense in M . The complement of G is a Cantor set. Clearly $\Omega_H^{\text{nw}} \subset G^c \subset \Omega_{H^{\text{ext}}}^*$. For each $x \in G$ let $n(x) \geq 0$ be the first time $f^n x \in H$ and let $k(x)$ be the corresponding component $H_{k(x)}^{\text{ext}}$. By continuity n and k are locally constant. Consider a point $x \in G$ and let $G(x)$ be the maximal interval containing x for which n and k are constant. We call $G(x)$ the gap of x . The image $f^{n(x)} G(x) \subset H_{k(x)}^{\text{ext}}$ and by maximality it is onto. Thus the boundary points of $G(x)$ get mapped onto the boundary points of H_k^{ext} . In particular, they are in the set Ω_H^{nw} by definition of the set H_k^{ext} .

We have shown that G^c is a Cantor set and all of its boundary points are in Ω_H^{nw} , i.e., non-wandering. Since the boundary points of a Cantor set are dense in itself, every point in G^c is non-wandering, i.e., $G^c \subset \Omega_H^{\text{nw}}$. Thus the claim follows. \square

The main theorem of this section is as follows.

THEOREM 6.3. *Every interval exclusion system has at most $4r$ topologically transitive components.*

To prove this theorem we will first prove that the number of topologically transitive components is at most $2r$ in the case of an SFT (Theorem 6.4). In the general case we will deduce that the number of topologically transitive components which stay a bounded distance away from the boundary of the extended hole is also bounded by $2r$. In Proposition 6.5 we develop a local concatenation technique (similar to bracketing for hyperbolic maps) for orbits that accumulate at a boundary point of the extended hole. This yields a correspondence between topologically transitive components which accumulate on the boundary of the extended hole and the boundary points of the extended hole (Proposition 6.7) and gives the upper bound of $2r$ on such components.

Let $\{X_i : X_i \subset M\}$ be the topologically transitive components of $f|_{\Omega^{\text{nw}}}$ and $\{Y_i : Y_i \subset \Sigma\}$ the topologically transitive components of $\sigma|_{\Sigma^{\text{nw}}}$. Let \mathcal{I} be the index set of the topologically transitive components of $f|_{\Omega^{\text{nw}}}$.

THEOREM 6.4. *Suppose that the interval exclusion system is an SFT. Then the number of topologically transitive components of Ω^{nw} is at most $2r$.*

Proof. If M is the circle we puncture it at a point of $\partial \mathcal{P}$ to view it as an interval. This allows us to define an order of M . Each topologically transitive component X_i is closed, thus we

can define $a_i := \min(X_i)$ and $b_i := \max(X_i)$. Since Σ^{nw} is a non-wandering SFT, it has only a finite number of topologically transitive components Y_i and these components are disjoint. For the moment, we suppose that all the points a_i and b_i have unique coding (i.e. the orbits of a_i and b_i do not intersect $\partial\mathcal{P}$). With this additional assumption the disjointness of the Y_i implies that a_i and b_i do not belong to any transitive component other than X_i . This implies that there are gaps G_{a_i} (G_{b_i}) on the left (right) side of a_i (b_i). Fix $x \in G_{a_i}$. As we saw above, all z between x and a_i fall into the same hole at the same time $n(z)$. By the continuity of f this implies that $f^{n(z)}a_i \in \partial H^{\text{ext}}$. A similar statement holds for G_{b_i} .

Consider the order on Σ^{nw} which is compatible with the order on Ω^{nw} . This order can always be defined in an inductive manner by simply considering the relative order of the elements of $\mathcal{P}^{(n)} := \bigvee_{i=0}^n f^i \mathcal{P}$. If f is locally order preserving then this is simply the lexicographical order on Σ^{nw} . Fix an $i \in \mathcal{I}$. Consider the symbolic coding of a_i and b_i . Call these codings $s = (s_j)_{j \in \mathbb{N}}$ and $t = (t_j)_{j \in \mathbb{N}}$ (here the dependence on i is suppressed since i is fixed). We claim that if X_i is uncountable then s is not a preimage of t and t is not a preimage of s . To see this, fix a higher block coding which defines the Markov transition graph of Σ^{nw} . The fact that a_i is defined via a minimum implies that if there are several followers of the symbol s_j in the Markov graph, then s_{j+1} is the minimal follower in the sense that in the order it is smaller than all other followers. In a similar fashion the sequence t is maximal.

If s is a preimage of t or vice versa, then s and t are eventually maximal and minimal at the same time. This means that starting from the point where they agree, the maximal follower of s_j is also the minimal follower of s_j , so there is only one follower of s_j . Thus s and t are eventually periodic and Y_i is simply this periodic orbit. In particular, Y_i is finite, finishing the proof of the claim.

If X_i is uncountable, then we have just shown the disjointness of the codes of a_i and b_i . Since we are still assuming that the points a_i and b_i have unique coding, this implies the disjointness of their f -orbits. Thus at least two points of X_i lie on ∂H^{ext} .

On the other hand if X_i is at most countable, then by Proposition 6.1 it is finite and consists of a single periodic orbit. The points a_i and b_i lie on this orbit.

Under the above assumptions we have shown that for every uncountable X_i there are at least two points of X_i on ∂H^{ext} , and for every at most countable X_i there is at least one such point. Thus, since the endpoints can only belong to a single X_i we have shown that

$$\#\{i \in \mathcal{I} : X_i \text{ at most countable}\} + 2\#\{i \in \mathcal{I} : X_i \text{ uncountable}\} \leq 2r.$$

Finally we turn to the case where some of the a_i and b_i have non-unique coding. Note that each point can have at most two codes. Repeat the above proof with (at most) $4r$ codes corresponding to these points. This proof shows that if two topologically transitive components intersect at an a_i or b_i then the coding of this intersection point is different in each of the components.

Now suppose that the coding of a_i is not unique and that two topologically transitive components intersect at a_i . Remember that $f^{n(z)}a_i \in \partial H^{\text{ext}}$. Consider the two Markov partition elements P_1, P_2 which share the point a_i . One of these elements, say P_1 , satisfies $f^{n(z)}P_1 \subset \overline{H^{\text{ext}}}$. This implies that the point a_i is isolated in the corresponding topologically transitive component. By Proposition 6.1 this topologically transitive component consists




FIGURE 2. The flag at \hat{x}_j is covered by the flag at \hat{x}_i , the flag at \hat{x}_k is uncovered by the flags at \hat{x}_i or at \hat{x}_j .

of the orbit of a_i and this orbit is periodic. Thus it is part of the other topologically transitive component. \square

Suppose that X is a topologically transitive component of Ω^{nw} . Fix $x \in X$ such that the orbit of x is dense in X . Consider any $\hat{x} = \hat{x}_1\hat{x}_2 \dots \in \Sigma^{\text{nw}}$ such that $\pi(\hat{x}) = x$. Let $\{c^{(j)}\}$ be the boundary points of the extended hole H^{ext} . If the point $c^{(j)}$ is not on the boundary of a Markov partition element then let $\hat{c}^{(j)} = \pi^{-1}c^{(j)} \in \Sigma^{\text{nw}}$; in the case where $c^{(j)}$ is on the boundary we define $\hat{c}^{(j)}$ to be the preimage of $c^{(j)}$ which is the coding of the orbit defined from the exterior of the extended hole. For each $k \geq 0$ we define $n_k(\hat{x})$ in the following way: $n_k(\hat{x})$ is the length of the longest initial block of one of the $\hat{c}^{(j)}$ which agrees with the initial block of the same length of $\sigma^k \hat{x}$, i.e., $\hat{x}_{i+k} = \hat{c}_i^{(j)}$ for $i = 0, 1, \dots, n_k(\hat{x}) - 1$. We call $n_k(\hat{x})$ the *flag* of x at time k (see Figure 2). We remark that $n_k(\hat{x}) < \infty$ if $f^k x \notin \partial H$.

By Proposition 6.1 either x is an isolated periodic point which hits ∂H or it visits ∂H only a finite number of times. In the second case, since, by Proposition 6.1, x is recurrent, we can assume by replacing x by $f^i x$ for sufficiently large i that the orbit of x does not visit ∂H at all.

Since x is fixed we will drop the x -dependence of the notation. We define a partial order on flags $n_k \prec n_{k'}$ if $k' < k$ and $n_{k'} + k' \geq n_k + k$, and in this case we say that the flag n_k is *covered* by the flag $n_{k'}$. If a flag is not covered by any other flag we say it is *uncovered*. The motivation for this terminology is the following lemma.

PROPOSITION 6.5. *If the flag of \hat{x} is uncovered at time k then we can concatenate \hat{x} 's initial segment of length k with any orbit in Σ^{nw} whose initial segment agrees with the next n_k entries of \hat{x} and the concatenated orbit belongs to Σ^* (it may be wandering and thus not belong to Σ^{nw}).*

To prove this fact we need the following lemma.

LEMMA 6.6. *Let $c \in \partial H^{\text{ext}}$ and $\hat{c}_0\hat{c}_1 \dots$ be the code of c . Suppose that $a \notin H^{\text{ext}}$ has coding $\hat{a} = \hat{c}_0\hat{c}_1 \dots \hat{c}_{l-1}\hat{a}_l\hat{a}_{l+1} \dots$ with $\hat{a}_l \neq \hat{c}_l$. Any point b with coding $\hat{b} = \hat{c}_0\hat{c}_1 \dots \hat{c}_{l-1}\hat{a}_l\hat{b}_{l+1}\hat{b}_{l+2} \dots$ satisfies $b \notin H^{\text{ext}}$.*

Proof. The hole H^{ext} consists of a finite union of intervals. We have assumed that the Markov partition \mathcal{P} is fine enough that the left and right end points of each of these intervals is in a different partition element. The point c is in the Markov partition element $[\hat{c}_0\hat{c}_1 \dots \hat{c}_{l-1}]$. Since $a \notin H^{\text{ext}}$ and $\hat{a}_l \neq \hat{c}_l$, the point \hat{c} is not in the Markov partition element $[\hat{c}_0\hat{c}_1 \dots \hat{c}_{l-1}\hat{a}_l]$. Thus $b \notin H^{\text{ext}}$. \square

Proof of Proposition 6.5. For concreteness suppose that we concatenate \hat{x} with \hat{y} to produce a point \hat{z} . First of all, if $j \geq k$ then $f^j z = f^j y$, thus since y never falls in a hole, z cannot fall in a hole after time k .

Suppose that $f^j z \in H$ for some $j < k$. Since the $n_k(\hat{x})$ flag of x is uncovered, we have $\hat{z}_j \cdots \hat{z}_{j+n_j(\hat{x})} = \hat{x}_j \cdots \hat{x}_{j+n_j(\hat{x})}$ and $j + n_j(\hat{x}) < k + n_k(\hat{x})$. Thus applying the previous lemma to $a = f^j y$ and $b = f^j z$ yields a contradiction and concludes the proof. \square

Assume that we start with a fine enough Markov partition that each $c^{(j)}$ belongs to a different element of the time zero partition. We define the color of the uncovered flag of x at time k to be the unique $c^{(j)}$ such that the initial blocks of length $n_k(\hat{x})$ of $\hat{c}^{(j)}$ and $\sigma^k \hat{x}$ coincide.

Let X be a (non-wandering) topologically transitive component and $x \in X$ a dense orbit. Define

$$I(X) = I(X, x) := \{c^{(j)} : \text{the uncovered flags of } x \text{ of color } j \text{ are of unbounded length}\}.$$

PROPOSITION 6.7.

- (1) *If $X \cap \partial H^{\text{ext}} \neq \emptyset$ then $I(X)$ is non-empty.*
- (2) *If $I(X) \cap I(Y) \neq \emptyset$ then $X = Y$.*

Proof of Proposition 6.7. (1) First note that

$$\limsup\{n_k(\hat{x}) : \text{the color of the flag at time } k \text{ is } j\} = \infty$$

for any $c^{(j)} \in X \cap \partial H^{\text{ext}}$. Any chain of covered flags $n_{k_1} \prec n_{k_2} \prec n_{k_3} \prec \dots$ contains a maximal element since the k_i 's are decreasing and bounded from below (positive). The maximal element is an uncovered flag, and it is longer than any of the flags it covers. Thus if the covered flags have unbounded length, then the uncovered flags also have unbounded length. Thus $I(X) \neq \emptyset$.

(2) Fix two topologically transitive components, X and Y , with dense orbits x and y . Consider $c := c^{(j)} \in I(X, x) \cap I(Y, y)$.

First consider the case where the forward orbit of x contains c . Proposition 6.1(3) implies that the forward orbit of c is dense, so we may assume that $x = c$. Let $\omega(y)$ denote the forward limit set of the orbit of y . Since $c \in \omega(y)$ then by continuity $w(c) \subset \omega(y)$, thus $X \subset Y$. Since topologically transitive components must be maximal it follows that $X = Y$. Similarly, we can conclude that $X = Y$ if the forward orbit of y contains c .

Now suppose that the forward orbits of x and y do not hit the point c . We will recursively construct a point $\hat{z} \in \Sigma^{\text{nw}}$ such that the orbit of $z = \pi(\hat{z})$ is dense in $X \cup Y$. The idea of the construction is the following. Take longer and longer segments of \hat{x} and \hat{y} 's orbit, so that the union of these segments is dense in X and Y . Proposition 6.5 gives a condition when one can concatenate segments of the orbit of \hat{x} and \hat{y} together. Since $I(X) \cap I(Y) \neq \emptyset$ we can apply Proposition 6.5 to arbitrarily long (i.e., arbitrarily dense) orbit segments.

More formally, fix a positive sequence $\varepsilon_m \rightarrow 0$. Let $\hat{x}(m, n)$ be the finite sequence $\hat{x}_m, \hat{x}_{m+1}, \dots, \hat{x}_n$. Define $x(m, n)$ similarly to be x_m, x_{m+1}, \dots, x_n . Let $n_k := n_k(\hat{x})$ and $m_l := n_l(\hat{y})$ be a sequence of the uncovered flags of color j with unbounded length and with $c^{(j)} \in I(X) \cap I(Y)$. Consider the first time t_1 such that the orbit segment $x(0, t_1)$

is ϵ_1 dense in X . Consider the smallest integer k_1 such that $k_1 \geq t_1$ such that the x -flag n_{k_1} is uncovered and of color j . Consider now the first time l_1 such that the y -flag m_{l_1} satisfies $m_{l_1} \geq n_{k_1}$. Let s_1 be the smallest integer such the orbit segment $y(m_{l_1}, s_1)$ is at least ϵ_1 dense in Y . Next fix $i_1 \geq s_1$ so that the y -flag m_{i_1} is uncovered and of color j . Finally, consider the first time p_1 such that the x -flag n_{p_1} satisfies $n_{p_1} \geq m_{i_1}$. For two finite strings a, b we denote by $a \star b$ the string a concatenated with b . Thus Proposition 6.5 guarantees that

$$\hat{z}_1 := \hat{x}(0, k_1) \star \hat{y}(l_1, i_1) \star \hat{x}(p_1, \infty) \in \Sigma^*.$$

Let $z_1 = \pi(\hat{z}_1)$. For $i \in \{0, 1, \dots, k_1\}$ the symbolic codes of $f^i x$ and $f^i z_1$ agree for at least n_{k_1} symbols and for $i \in \{l_1, \dots, i_1\}$ the symbolic codes of $f^i y$ and $f^i z_1$ agree for at least n_{i_1} symbols. Thus, since f is uniformly hyperbolic, we can find constants $C > 0$ and $\lambda \in (0, 1)$ which depend only on f and \mathcal{P} , such that $|f^i x - f^i z_1| < C\lambda^{k_1}$ for $i \in \{0, 1, \dots, k_1\}$ and $|f^i y - f^i z_1| < C\lambda^{i_1}$ for $i \in \{l_1, l_1 + 1, \dots, i_1\}$. Thus the orbit segment $z_1(0, i_1)$ is at least $\epsilon_1 + C\lambda^{\min(k_1, i_1)}$ dense in $X \cup Y$.

We now repeat this procedure. Consider the first time t_2 such that the orbit segment $x(p_1, t_2)$ is ϵ_2 dense in X . Consider the smallest integer k_2 such that $k_2 \geq t_2$ such that the x -flag n_{k_2} is uncovered and of color j . Consider now the first time l_2 such that the y -flag m_{l_2} satisfies $m_{l_2} \geq n_{k_2}$. Let s_2 the smallest integer such the orbit segment $y(m_{l_2}, s_2)$ is at least ϵ_2 dense in Y . Next fix $i_2 \geq s_2$ so that the y -flag m_{i_2} is uncovered of color j . Finally, consider the first time p_2 such that the x -flag n_{p_2} satisfies $n_{p_2} \geq m_{i_2}$. Thus Proposition 6.5 guarantees that the orbit

$$\hat{z}_2 := \hat{x}(0, k_1) \star \hat{y}(l_1, i_1) \star \hat{x}(p_1, k_2) \star \hat{y}(l_2, i_2) \star \hat{x}(p_2, \infty) \in \Sigma^*.$$

As above, the orbit segment $z_2(p_1, i_2)$ is at least $\epsilon_2 + C\lambda^{\min(k_2, i_2)}$ dense in $X \cup Y$.

Define $\hat{z}_n \in \Sigma^*$ similarly. Clearly the limit $\hat{z} := \lim_{n \rightarrow \infty} \hat{z}_n$ exists. Since Σ^* is closed, we have $\hat{z} \in \Sigma^*$. We have $k_n \rightarrow \infty$ since $\epsilon_n \rightarrow 0$. Thus orbit of $z = \pi(\hat{z})$ is dense in $X \cup Y$. In particular, z is non-wandering and thus $z \in \Sigma^{\text{nw}}$. \square

Proof of Theorem 6.3. If Σ^{nw} is not an SFT we approximate H^{ext} from the outside by Markov holes $H^{(n)}$ as follows. Let \mathcal{P} be the generating Markov partition; then $\mathcal{P}^{(n)} := \bigvee_{i=-n}^n f^i \mathcal{P}$ is also a generating Markov partition. Let

$$H^{(n)} := \text{int} \left(\bigcup_{\{P \in \mathcal{P}^{(n)} : P \cap \overline{H} \neq \emptyset\}} P \right).$$

Clearly $H^{(1)} \supset H^{(2)} \supset \dots$ and $\overline{H} \subset \bigcap_n H^{(n)}$. Furthermore, we have $\text{diam}(\mathcal{P}^{(n)}) \rightarrow 0$ as $n \rightarrow \infty$ and thus $\bigcap_n H^{(n)} = \overline{H}$. The exclusion system $\Sigma^{\text{nw}(n)} := \Sigma_{H^{(n)}}^{\text{nw}}$ is an SFT. Clearly $\Sigma^{\text{nw}(1)} \subset \Sigma^{\text{nw}(2)} \subset \dots$. Let

$$\tilde{\Sigma}^{\text{nw}} := \overline{\bigcup_n \Sigma^{\text{nw}(n)}}.$$

Clearly $\tilde{\Sigma}^{\text{nw}} \subset \Sigma^{\text{nw}}$.

The number of transitive components of $\overline{\bigcup_n \Sigma^{\text{nw}(n)}}$ is less than or equal to the limsup of the number of components of the approximating sequence. Applying Theorem 6.4, this limsup is at most $2r$.

The orbit of any point in $\Sigma^{\text{nw}} \setminus \overline{\bigcup_n \Sigma^{\text{nw}(n)}}$ accumulates on ∂H^{ext} . Applying Proposition 6.7 we see that there is at most one such topologically transitive component for each color $c^{(j)}$, i.e., $2r$. \square

7. Coded systems

A subshift Σ is called a *coded* system if it is topologically conjugate to the shift on an irreducible countable labeled graph [3]. Equivalently, Σ is called coded if Σ contains an increasing sequence of irreducible SFTs whose union is dense in Σ [4]. This alternative definition indicates that topologically transitive components of hole systems should be coded since we can approximate holes from the outside by Markov holes. It turns out that this is true under mild assumptions. We begin with a preparatory proposition. Remember that a non-wandering subshift of finite type consists of a finite union of irreducible SFTs.

PROPOSITION 7.1. *A subshift Σ which contains an increasing sequence of non-wandering subshifts of finite type whose union is dense in Σ has at most countably many topologically transitive components and each topologically transitive component is coded.*

Proof. Suppose $\Sigma = \overline{\bigcup_n \Sigma^{(n)}}$ where $\Sigma^{(n)}$ is a non-wandering subshift of finite type. If each $\Sigma^{(n)}$ were topologically transitive then Σ would be coded. Let $A_i^{(n)}$ be the topologically transitive components of $\Sigma^{(n)}$. The $A_i^{(n)}$ form a filtration in the sense that for each $A_i^{(n)}$ there exists $A_j^{(n+1)}$ such that $A_i^{(n)} \subset A_j^{(n+1)}$. In other words, the transitive components $A_i^{(n)}$ form an at most countable union of directed trees where each node's out degree is exactly one. Each topologically transitive component of $\sigma|_{\Sigma}$ contains a set $\cup A_i^{(n)}$ where the union is taken over a directed path in one of the trees (we call this a path limit). Noticing that such a path is uniquely defined by the root of the tree since the out degree is always one, there are at most countably many such paths and thus $\sigma|_{\Sigma}$ has at most countably many topologically transitive components. This finishes the proof of the countability of the claim of the theorem.

Suppose that $C \subset \Sigma$ is a topologically transitive component of $\sigma|_{\Sigma}$. To see that C is coded we will define a new filtration. Since $\Sigma^{(n)}$ is a non-wandering subshift of finite type, it has finitely many topologically transitive components $A_i^{(n)}$ which are pairwise disjoint. Consider those components $A_i^{(n)} : i = 1, \dots, k_n$ which are strictly contained in C . We can assume that the $A_i^{(n)}$ are so ordered that $A_1^{(n)} \subset A_1^{(n+1)}$ for all n . We only need to show that

$$\overline{\bigcup_n A_1^{(n)}} = C. \quad (7.2)$$

The rest of the proof is devoted to establishing this equality.

If equation (7.2) is not true, there is an n_0 such that for all $n \geq n_0$ we can find another $A_i^{(n)} \subset C$, which we denote without loss of generality by $A_2^{(n)}$, such that $A_2^{(n)} \subset A_2^{(n+1)}$ for all $n \geq n_0$ but $\bigcup_n A_2^{(n)} \cap \bigcup_n A_1^{(n)} = \emptyset$. In the terminology introduced above, this means we can find two disjoint paths in the trees whose path limits are both contained in C .

Fix $n \geq n_0$. For $i = 1, 2$ consider a finite word $w_i \in A_i^{(n)}$ where each symbol and each transition which characterize $A_i^{(n)}$ appear in w_i . Since C is topologically transitive, there

is a point $\underline{x} \in C$ where the words w_1 and w_2 both appear in \underline{x} . Thus, we can find l (which we assume positive without loss of generality) so that $\underline{x} \in \sigma^l w_1 \cap w_2$.

We form a new SFT $\hat{\Sigma}$ as follows. We start by defining an SFT with the permitted transitions being exactly the transitions $\underline{x}_i \underline{x}_{i+1}$ for $i = 0, \dots, N$ for sufficiently large N . This defines an SFT $\hat{\Sigma}$, but it may be reducible. If N is sufficiently large, then all the transitions of $A_1^{(n)}$ and of $A_2^{(n)}$ are allowed in $\hat{\Sigma}$. Furthermore, transitions from each of these sets to the other are allowed in $\hat{\Sigma}$. Take the irreducible component of $\hat{\Sigma}$ which includes $A_1^{(n)} \cup A_2^{(n)}$ and call it Σ' . But, by the construction of $\Sigma^{(m)}$, we have $\Sigma' \subset \Sigma^{(m)}$ for sufficiently large m and thus $\Sigma' \subset A_1^{(m)}$ and $\Sigma' \subset A_2^{(m)}$, a contradiction. \square

We call a topologically transitive component A *approximable* (aTTC) if periodic points whose orbit avoids \overline{H} are dense in A ; otherwise we call A *non-approximable* (nTTC).

THEOREM 7.2. *Every aTTC is coded. There are at most countably many such topologically transitive components.*

Proof. Let Σ^{nw} be an exclusion subshift, i.e., f is an axiom A map, \mathcal{P} a fixed proper, generating Markov partition, and H an open hole with boundary consisting of a finite union of compact codimension-one manifolds. We approximate by Markov holes as in the proof of Theorem 6.3, the only difference being that for the ‘exclusion’ system $\Sigma^{\text{nw}(n)} := \Sigma_{H^{(n)}}^{\text{nw}}$ the boundary of $H^{(n)}$ may not be the union of compact codimension-one manifolds, thus the quotes on the word exclusion. By Proposition 7.1, $\tilde{\Sigma}^{\text{nw}}$ has at most countably many topologically transitive components and each component is coded.

Since $\Sigma^{\text{nw}(n)}$ is a non-wandering SFT disjoint from \overline{H} , periodic points disjoint from \overline{H} are dense in $\Sigma^{\text{nw}(n)}$, and thus they are also dense in $\tilde{\Sigma}^{\text{nw}}$. Thus by definition any topologically transitive component which is a subset of $\tilde{\Sigma}^{\text{nw}}$ is an aTTC. On the other hand if $x \in \Sigma^{\text{nw}} \setminus \tilde{\Sigma}$ then x is not an accumulation point of periodic points x_n which avoid \overline{H} , since any such periodic point necessarily belongs to $\Sigma^{\text{nw}(n)}$ for some n . Thus we have shown that every aTTC is contained in $\tilde{\Sigma}^{\text{nw}}$. \square

We can construct a hole such that any periodic orbit is an nTTC. However, we do not know if less trivial nTTCs can exist, and, in particular, if non-coded nTTCs can exist and if uncountably many nTTCs can exist.

We actually have a stronger property than coded since our subshifts are well approximable from outside as well as inside.

8. Open problems

Many interesting open problems remain. The results on the characterization of topologically transitive components in the multi-dimensional case are much less satisfactory than in the one-dimensional case. It is hard to see how to obtain more precise results without stronger assumptions.

In the one-dimensional case, is it possible to classify those interval exclusion subshifts which are SFTs via the arithmetical properties of the boundary points of the intervals.

One would like to develop a relationship between the escape rate properties and topological and/or metric invariants of the invariant set.

Acknowledgements. We would like to thank the referee for useful remarks. This paper was written during reciprocal visits which were supported by SFB 288, PRODYN and Universität Bielefeld.

REFERENCES

- [1] H. van den Bedem and N. Chernov. Expanding maps of an interval with holes. *Ergod. Th. & Dynam. Sys.* **22** (2002), 637–654.
- [2] F. Blanchard. β -shifts and symbolic dynamics. *Theoret. Comput. Sci.* **65** (1989), 131–141.
- [3] F. Blanchard and G. Hansel. Systèmes codés. *Theoret. Comput. Sci.* **44** (1986), 17–49.
- [4] F. Blanchard and G. Hansel. Sofic constant-to-one extensions of subshifts of finite type. *Proc. Amer. Math. Soc.* **112** (1991), 259–265.
- [5] H. Bruin, M. Demers and I. Melbourne. Existence and convergence properties of physical measures for certain dynamical systems with holes. *Ergod. Th. & Dynam. Sys.* **30**(3) (2010), 687–728.
- [6] L. Bunimovich and A. Yurchenko. Where to place a hole to achieve a maximal escape rate, 2008, arXiv:0811.4438.
- [7] A. de Carvalho. Pruning fronts and the formation of horseshoes. *Ergod. Th. & Dynam. Sys.* **19** (1999), 851–894.
- [8] N. Chernov and R. Markarian. Ergodic properties of Anosov maps with rectangular holes. *Bol. Soc. Bras. Mat.* **28** (1997), 271–314.
- [9] N. Chernov and R. Markarian. Anosov maps with rectangular holes. Nonergodic cases. *Bol. Soc. Bras. Mat.* **28** (1997), 315–342.
- [10] N. Chernov, R. Markarian and S. Troubetzkoy. Conditionally invariant measures for Anosov maps with small holes. *Ergod. Th. & Dynam. Sys.* **18** (1998), 1049–1073.
- [11] N. Chernov, R. Markarian and S. Troubetzkoy. Invariant measures for Anosov maps with small holes. *Ergod. Th. & Dynam. Sys.* **20** (2000), 1007–1044.
- [12] P. Collet, S. Martínez and B. Schmitt. The Yorke–Pianigiani measure and the asymptotic law on the limit Cantor set of expanding systems. *Nonlinearity* **7** (1994), 1437–1443.
- [13] P. Cvitanović. Dynamical averaging in terms of periodic orbits. *Phys. D* **83** (1995), 109–123.
- [14] M. Demers, P. Wright and L.-S. Young. Escape rates and physically relevant measures for billiards with small holes. *Comm. Math. Phys.* **294**(2) (2010), 353–388.
- [15] M. Demers and L.-S. Young. Escape rates and conditionally invariant measures. *Nonlinearity* **19** (2006), 377–397.
- [16] M. Demers. Markov extensions and conditionally invariant measures for certain logistic maps with small holes. *Ergod. Th. & Dynam. Sys.* **25** (2005), 1139–1171.
- [17] M. Demers. Markov extensions for dynamical systems with holes: an application to expanding maps of the interval. *Israel J. Math.* **146** (2005), 189–221.
- [18] K. Falconer. *The Geometry of Fractal Sets (Cambridge Tracts in Mathematics, 85)*. Cambridge University Press, Cambridge, 1985.
- [19] P. Gaspard and J. Dorfman. Chaotic scattering theory, thermodynamic formalism, and transport coefficients. *Phys. Rev. E* **52** (1995), 3525–3552.
- [20] G. Keller and C. Liverani. Rare events, escape rates and quasistationarity: some exact formulae. *J. Stat. Phys.* **135** (2009), 519–534.
- [21] D. Lind and B. Marcus. *Symbolic Dynamics and Coding*. Cambridge University Press, Cambridge, 1995.
- [22] C. Liverani and V. Maume-Deschamps. Lasota–Yorke maps with holes: conditionally invariant probability measures and invariant probability measures on the survivor set. *Ann. Inst. H. Poincaré Probab. Statist.* **39** (2003), 385–412.
- [23] A. Lopes and R. Markarian. Open billiards: Cantor sets, invariant and conditionally invariant probabilities. *SIAM J. Appl. Math.* **56** (1996), 651–680.
- [24] G. Pianigiani and J. Yorke. Expanding maps on sets which are almost invariant: decay and chaos. *Trans. Amer. Math. Soc.* **252** (1979), 351–366.
- [25] D. Ruelle. Entropy production in nonequilibrium statistical mechanics. *Comm. Math. Phys.* **189** (1997), 365–371.