

SUBSYSTEM ENTROPIES OF SHIFTS OF FINITE TYPE AND SOFIC SHIFTS ON COUNTABLE AMENABLE GROUPS

ROBERT BLAND, KEVIN MCGOFF, AND RONNIE PAVLOV

ABSTRACT. In this work we study the entropies of subsystems of shifts of finite type (SFTs) and sofic shifts on countable amenable groups. We prove that for any countable amenable group G , if X is a G -SFT with positive topological entropy $h(X) > 0$, then the entropies of the SFT subsystems of X are dense in the interval $[0, h(X)]$. In fact, we prove a “relative” version of the same result: if X is a G -SFT and $Y \subset X$ is a subshift such that $h(Y) < h(X)$, then the entropies of the SFTs Z for which $Y \subset Z \subset X$ are dense in $[h(Y), h(X)]$. We also establish analogous results for sofic G -shifts.

1. INTRODUCTION

Let G be a countable group and let \mathcal{A} be a finite alphabet of symbols. In symbolic dynamics, the central objects of study are the subsystems of the so-called *full shift*, the dynamical system (\mathcal{A}^G, σ) , where σ denotes the action of G on \mathcal{A}^G by translations (Definition 2.7). *Shifts of finite type* (Definition 2.16) and *sofic shifts* (Definition 2.17) are the most widely studied and well understood examples of symbolic dynamical systems. In each of these cases, the system of interest is completely specified by a finite amount of information. This allows for combinatorial, finitary arguments to be applied to the analysis of the dynamics of such systems.

Entropy is one of the most fundamental invariants of a topological dynamical system. Many fundamental results from classical entropy theory (i.e., in the case where $G = \mathbb{Z}$) only generalize if G is an *amenable* group (Definition 2.2). Amenability allows one to “approximate” the group by a sequence of finite subsets in a way that is useful for studying dynamics. See Definition 2.22 for the definition of the entropy of a symbolic dynamical system on an amenable group.

In general, one would like to understand the structure of the collection of subsystems of a given subshift. In this paper we study the entropies of the SFT subsystems of a given SFT, as well as the entropies of the sofic subsystems of sofic shifts. There are many existing results in the literature in the case where $G = \mathbb{Z}$. For example, the Krieger Embedding Theorem [12] characterizes the irreducible SFT subsystems of a given irreducible \mathbb{Z} -SFT. Additionally, Lind [13] has provided an algebraic characterization of the real numbers that are realized as the entropy of a \mathbb{Z} -SFT.

However, the situation is very different in cases where $G \neq \mathbb{Z}$. Even in the case where $G = \mathbb{Z}^d$ for $d > 1$, the classes of SFTs and sofic shifts behave quite differently. For example, Boyle, Pavlov, and Schraudner [5] have shown by example that the subsystems of \mathbb{Z}^d sofic shifts can be badly behaved for $d > 1$ (in contrast with the case where $d = 1$). Moreover, Hochman and Meyerovitch [9] have characterized

1991 *Mathematics Subject Classification.* Primary 37B10, Secondary 37B40, 37B51.

Key words and phrases. Amenable groups, Shifts of finite type, Entropy.

the real numbers that are realized as entropy of a \mathbb{Z}^d -SFT (with $d > 1$), but in contrast to the result of Lind mentioned above, the characterization is in algorithmic terms and unavoidably involves concepts from computability and recursion theory. Nonetheless, Desai [6] has shown that a \mathbb{Z}^d -SFT with positive entropy has a wealth of SFT subsystems (sharpening an earlier result of Quas and Trow [16]).

Theorem 1.1 ([6]). *Let $G = \mathbb{Z}^d$ for some $d \in \mathbb{N}$ and let X be a G -SFT such that $h(X) > 0$. Then*

$$\{h(Y) : Y \subset X \text{ and } Y \text{ is an SFT}\}$$

is dense in $[0, h(X)]$.

In recent years, several results of the $G = \mathbb{Z}$ and $G = \mathbb{Z}^d$ cases have seen extensions to larger classes of groups, especially amenable groups. To name a few: Barbieri [2] has classified the real numbers that are realized as the entropy of a G -SFT for many types of amenable G (extending the result of Hochman and Meyerovitch mentioned above); Frisch and Tamuz [8] have investigated the (topologically) generic properties of G -subshifts for arbitrary amenable G ; Barbieri and Sablik [3] have shown how an arbitrary effective G -subshift, where G is finitely generated, may be simulated by a G' -SFT, where G' is the semidirect product $G' = \mathbb{Z}^2 \rtimes G$; and Huczek and Kopacz [10] have (very recently) obtained a partial generalization of Boyle's lower entropy factor Theorem [4] to countable amenable groups with the comparison property. In this vein, we prove the following generalization of Theorem 1.1 to arbitrary countable amenable groups.

Theorem 4.2. *Let G be a countable amenable group, let X be a G -SFT, and let $Y \subset X$ be any subsystem such that $h(Y) < h(X)$. Then*

$$\{h(Z) : Y \subset Z \subset X \text{ and } Z \text{ is an SFT}\}$$

is dense in $[h(Y), h(X)]$.

Choosing $G = \mathbb{Z}^d$ and $Y = \emptyset$ in the above theorem recovers the result of Desai (Theorem 1.1 above). Note that a shift space $X \subset \mathcal{A}^G$ has at most countably many SFT subsystems, and therefore the set of entropies of SFT subsystems is at most countable. In this sense, Theorem 4.2 is "the most one could hope for."

Remark 1.2. After a preprint of this work was made public, the authors of [8] made us aware that a short alternate proof of Theorem 4.2 can be derived from their main results. Specifically, they prove there that for any countable amenable group G and any real $c \geq 0$, the set of G -subshifts with entropy c is dense (in fact residual) within the space of G -subshifts with entropy at least c with respect to the Hausdorff topology. This result immediately implies that for any G -SFT X , there exist G -subshifts contained in X that achieve all possible entropies in $[0, h(X)]$; then, some simple approximations with G -SFTs (in the sense of our Theorem 2.27) can be used to obtain a proof of Theorem 4.2.

For sofic shifts, we obtain the following result.

Theorem 5.2. *Let G be a countable amenable group, let W be a sofic G -shift, and let $V \subset W$ be any subsystem such that $h(V) < h(W)$. Then*

$$\{h(U) : V \subset U \subset W \text{ and } U \text{ is sofic}\}$$

is dense in $[h(V), h(W)]$.

From this result, we can quickly derive the fact (Corollary 5.3) that if X is a sofic G -shift, then each real number in $[0, h(X)]$ can be realized as the entropy of some (not necessarily sofic) subsystem of X . (Recall that the alternate proof of Theorem 4.2 described in Remark 1.2 above relies on a version of this result requiring X to be an SFT.) The tool for proving Theorem 5.2 (from Theorem 4.2) is provided by the following theorem, which may be of independent interest. We note that this result generalizes another theorem of Desai [6, Proposition 4.3], which addressed the case $G = \mathbb{Z}^d$.

Theorem 5.1. *Let G be a countable amenable group and let W be a sofic G -shift. For every $\varepsilon > 0$, there exists an SFT \tilde{X} and a one-block code $\tilde{\phi} : \tilde{X} \rightarrow W$ such that the maximal entropy gap of $\tilde{\phi}$ satisfies $\mathcal{H}(\tilde{\phi}) < \varepsilon$.*

The maximal entropy gap $\mathcal{H}(\tilde{\phi})$ is defined in §2 (Definition 2.28). In particular, this result implies that if Y is sofic and $\varepsilon > 0$, then there is an SFT X that factors onto Y and satisfies $h(X) < h(Y) + \varepsilon$.

Our proofs of Theorems 4.2, 5.1, and 5.2 take the same general approach as the arguments given by Desai for the $G = \mathbb{Z}^d$ case. However, the extension to the general amenable setting requires substantial new techniques. Indeed, our proofs are made possible by the existence of *exact tilings* (Definition 3.1) of the group G that possess nice dynamical properties. Such exact tilings are trivial to find for \mathbb{Z}^d (by tiling the group using large hypercubes), but for arbitrary amenable groups were only recently constructed by Downarowicz, Huczek, and Zhang [7]; their construction is the main technical tool employed in this paper.

As mentioned in Remark 1.2 above, Theorem 4.2 can be alternately derived from results in [8]. We present a self-contained proof here for two reasons. Firstly, we would like to present a direct adaptation of the techniques from [6], since it demonstrates the power of the improved tiling results of [7]. Secondly, this presentation provides a unified approach to all of our proofs, since our proofs in the sofic setting (where we are not aware of alternative proofs) also rely on tiling-based constructions that are similar to those in our proof of Theorem 4.2.

The paper is organized as follows. In §2 we discuss basic notions and elementary theorems of symbolic dynamics, set in terms appropriate for countable amenable groups. In §3 we define and explore the concept of tilings and exact tilings of amenable groups, appealing to Downarowicz, Huczek, and Zhang for the existence of certain desirable tilings. In §4 we prove our main results for G -SFTs, and in §5 we prove our main results for sofic G -shifts. Finally, in §6 we provide an example of a \mathbb{Z}^2 sofic shift whose only SFT subsystem is a fixed point.

2. BASICS OF SYMBOLIC DYNAMICS

2.1. Amenable groups. We begin with a brief overview of amenable groups.

Definition 2.1 (Group theory notations). Let G be a group and let $K, F \subset G$ be subsets. We employ the following notations.

- i. The *group identity* is denoted by the symbol $e \in G$,
- ii. $KF = \{kf : k \in K \text{ and } f \in F\}$,
- iii. $K^{-1} = \{k^{-1} : k \in K\}$,
- iv. $Kg = \{kg : k \in K\}$ for each $g \in G$,
- v. $K \sqcup F$ expresses that K and F are disjoint, and is their (*disjoint*) *union*,
- vi. $K \triangle F = (K \setminus F) \sqcup (F \setminus K)$ is the *symmetric difference* of K and F , and

vii. $|K|$ is the *cardinality* of the (finite) set K .

Definition 2.2 (Følner condition for amenability). Let G be a countable group. A *Følner sequence* is a sequence $(F_n)_n$ of finite subsets $F_n \subset G$ which *exhausts* G (in the sense that for each $g \in G$, we have $g \in F_n$ for all sufficiently large n) and for which it holds that

$$\lim_{n \rightarrow \infty} \frac{|KF_n \Delta F_n|}{|F_n|} = 0$$

for every finite subset $K \subset G$. If such a sequence exists, then G is said to be an *amenable* group.

Throughout this paper, G denotes a fixed countably infinite amenable group and $(F_n)_n$ is a fixed Følner sequence for G .

Definition 2.3 (Invariance). Let $K, F \subset G$ be finite subsets, and let $\varepsilon > 0$. We say F is (K, ε) -*invariant* if

$$\frac{|KF \Delta F|}{|F|} < \varepsilon.$$

If $e \in K$ and F is (K, ε) -invariant, then F is also (K', ε') -invariant for any $\varepsilon' > \varepsilon$ and any $K' \subset K$ such that $e \in K'$. If F is (K, ε) -invariant, then so is the translate Fg for each fixed $g \in G$. Invariance is the primary way by which we say a large finite subset $F \subset G$ is a “good finite approximation” of G , according to the finitary quantifiers K and ε . The amenability of G provides a wealth of nearly invariant sets, which enables such approximation for the purpose of studying the dynamics of G -actions.

Next we develop concepts related to the geometry of finite subsets of G .

Definition 2.4 (Boundary and interior). Let $K, F \subset G$ be finite subsets. The K -*boundary* of F is the set

$$\partial_K F = \{f \in F : Kf \not\subset F\},$$

and the K -*interior* of F is the set

$$\text{int}_K F = \{f \in F : Kf \subset F\}.$$

Observe that $F = (\partial_K F) \sqcup (\text{int}_K F)$.

If F is sufficiently invariant with respect to K , then the K -boundary of F is a small subset of F (proportionally), by the following lemma.

Lemma 2.5. *Suppose $K, F \subset G$ are nonempty finite subsets and $e \in K$. Then*

$$\frac{1}{|K|} |KF \Delta F| \leq |\partial_K F| \leq |K| |KF \Delta F|.$$

In particular, if F is (K, ε) -invariant then $|\partial_K F| < \varepsilon |K| |F|$.

Proof. If $e \in K$, then $KF \Delta F = KF \setminus F$. If $g \in KF \setminus F$, then $g = kf$ for some $k \in K$ and $f \in \partial_K F$, by Definition 2.4. Therefore $KF \setminus F \subset K\partial_K F$, in which case $|KF \setminus F| \leq |K| |\partial_K F|$.

For the second inequality, note that $f \in \partial_K F$ implies $\exists k \in K$ such that $kf \notin F$, therefore $g = kf \in KF \setminus F$ is a point such that $f \in K^{-1}g \subset K^{-1}(KF \setminus F)$. Consequently $\partial_K F \subset K^{-1}(KF \setminus F)$, in which case $|\partial_K F| \leq |K| |KF \setminus F|$.

Finally if F is (K, ε) -invariant, then $|\partial_K F| \leq |K| |KF \setminus F| < \varepsilon |K| |F|$. \square

Given finite subsets $K, F \subset G$, in this paper we focus on the KK^{-1} -boundary and KK^{-1} -interior of F (rather than the K -boundary and K -interior), and we make use of the following lemma.

Lemma 2.6. *Let $K, F \subset G$. For any translate Kg of K (for any $g \in G$), either $Kg \subset F$ or $Kg \subset (\text{int}_{KK^{-1}} F)^c$ (or both are true).*

Proof. Suppose $Kg \not\subset (\text{int}_{KK^{-1}} F)^c$. Then $\exists f \in \text{int}_{KK^{-1}} F$ such that $f \in Kg$, which implies $g \in K^{-1}f$ and hence $Kg \subset KK^{-1}f \subset F$. \square

2.2. Shift spaces. Here we present necessary definitions from symbolic dynamics. See Lind and Marcus [14] for an introductory treatment of these concepts.

Definition 2.7 (Shifts and subshifts). Let \mathcal{A} be a finite set of symbols equipped with the discrete topology. A function $x : G \rightarrow \mathcal{A}$ is called an \mathcal{A} -labelling of G . By convention, we write x_g for the symbol $x(g) \in \mathcal{A}$ which is placed by x at $g \in G$. The set of all \mathcal{A} -labellings of G is denoted \mathcal{A}^G , which we equip with the product topology. For each $g \in G$, let $\sigma^g : \mathcal{A}^G \rightarrow \mathcal{A}^G$ denote the map given by

$$(\sigma^g x)_h = x_{hg} \quad \forall h \in G$$

for each $x \in \mathcal{A}^G$. The collection $\sigma = (\sigma^g)_{g \in G}$ is an action of G on \mathcal{A}^G by homeomorphisms. The pair (\mathcal{A}^G, σ) is a dynamical system called the *full shift* over the alphabet \mathcal{A} . A subset $X \subset \mathcal{A}^G$ is called *shift-invariant* if $\sigma^g x \in X$ for each $x \in X$ and $g \in G$. A closed, shift-invariant subset $X \subset \mathcal{A}^G$ is called a *subshift* or a *shift space*. For a given $x \in \mathcal{A}^G$, the *orbit* of x is the subset $\mathcal{O}(x) = \{\sigma^g x : g \in G\} \subset \mathcal{A}^G$. The subshift *generated by x* is the topological closure of $\mathcal{O}(x)$ as a subset of \mathcal{A}^G , and is denoted $\overline{\mathcal{O}}(x) \subset \mathcal{A}^G$.

Definition 2.8 (Codes and factors). Let $\mathcal{A}_X, \mathcal{A}_W$ be finite alphabets and let $X \subset \mathcal{A}_X^G$ and $W \subset \mathcal{A}_W^G$ be subshifts. A map $\phi : X \rightarrow W$ is *shift-commuting* if $\phi \circ \sigma^g = \sigma^g \circ \phi$ for each $g \in G$; the map ϕ is said to be a *sliding block code* if it is continuous and shift-commuting; and ϕ is said to be a *factor map* if it is a surjective sliding block code. If a factor map exists from X to W , then W is said to be a *factor* of X and X is said to *factor onto* W . If a sliding block code $\phi : X \rightarrow W$ is invertible and bi-continuous, then ϕ is said to be a *topological conjugacy*, in which case X and W are said to be *topologically conjugate*.

Definition 2.9 (Products of shifts). If \mathcal{A} and Σ are finite alphabets, then $\mathcal{A} \times \Sigma$ is also a finite alphabet (of ordered pairs). If $X \subset \mathcal{A}^G$ and $T \subset \Sigma^G$ are subshifts, then we view the *dynamical direct product* $X \times T$ as a subshift of $(\mathcal{A} \times \Sigma)^G$, defined by $(x, t) \in X \times T$ if and only if $x \in X$ and $t \in T$. The shift space $X \times T$ factors onto both X and T via the projection maps π_X and π_T , given by $\pi_X(x, t) = x$ and $\pi_T(x, t) = t$ for each $(x, t) \in X \times T$.

Remark 2.10. Definition 2.9 above introduces an abuse of notation, as technically we have $(x, t) \in \mathcal{A}^G \times \Sigma^G \neq (\mathcal{A} \times \Sigma)^G$. However, if equipped with the G -action ς given by $\varsigma^g(x, t) = (\sigma^g x, \sigma^g t)$, then $\mathcal{A}^G \times \Sigma^G$ becomes a dynamical system that is topologically conjugate to $(\mathcal{A} \times \Sigma)^G$.

2.3. Patterns. In this section we describe *patterns* and their related combinatorics.

Definition 2.11 (Patterns). Let \mathcal{A} be a finite alphabet and let $F \subset G$ be a finite set. A function $p : F \rightarrow \mathcal{A}$ is called a *pattern*, said to be of *shape F* . The set of all

patterns of shape F is denoted \mathcal{A}^F . The set of all patterns of any finite shape is denoted $\mathcal{A}^* = \bigcup_F \mathcal{A}^F$, where the union is taken over all finite subsets $F \subset G$.

Remark 2.12. Given a point $x \in \mathcal{A}^G$ and a finite subset $F \subset G$, we take $x(F)$ to mean the *restriction* of x to F , which is itself a pattern of shape F . Usually this is denoted $x|_F \in \mathcal{A}^F$, but we raise F from the subscript for readability.

Definition 2.13 (One-block code). Let \mathcal{A}_X and \mathcal{A}_W be finite alphabets and let $X \subset \mathcal{A}_X^G$ and $W \subset \mathcal{A}_W^G$ be subshifts. A factor map $\phi : X \rightarrow W$ is said to be a *one-block code* if there exists a function $\Phi : \mathcal{A}_X \rightarrow \mathcal{A}_W$ with the property that

$$\phi(x)_g = \Phi(x_g), \quad \forall g \in G$$

for each $x \in X$.

Definition 2.14 (Occurrence). Let \mathcal{A} be a finite alphabet and let $F \subset G$ be a finite set. A pattern $p \in \mathcal{A}^F$ is said to *occur* in a point $x \in \mathcal{A}^G$ if there exists an element $g \in G$ such that $(\sigma^g x)(F) = p$. If $X \subset \mathcal{A}^G$ is a subshift, then the collection of all patterns of shape F occurring in any point of X is denoted by

$$\mathcal{P}(F, X) = \{(\sigma^g x)(F) \in \mathcal{A}^F : x \in X \text{ and } g \in G\}.$$

If $X \subset \mathcal{A}^G$ is a subshift and $F \subset G$ is a finite subset, then $|\mathcal{P}(F, X)| \leq |\mathcal{A}|^{|F|}$. If $F' \subset G$ is another finite subset, then $|\mathcal{P}(F \cup F', X)| \leq |\mathcal{P}(F, X)| \cdot |\mathcal{P}(F', X)|$. If $F' \subset F$ and $X' \subset X$, then $|\mathcal{P}(F', X')| \leq |\mathcal{P}(F, X)|$.

Definition 2.15 (Forbidden patterns). Let \mathcal{A} be a finite alphabet, let $F \subset G$ be a finite set and let $X \subset \mathcal{A}^G$ be a subshift. A pattern $p \in \mathcal{A}^F$ is said to be *allowed* in X if $p \in \mathcal{P}(F, X)$ (if p occurs in at least one point of X).

Given a (finite or infinite) collection of patterns $\mathcal{F} \subset \mathcal{A}^*$, a new subshift $X' \subset X$ may be constructed by expressly *forbidding* the patterns in \mathcal{F} from occurring in points of X . We denote this by

$$X' = \mathcal{R}(X, \mathcal{F}) = \{x \in X : \forall p \in \mathcal{F}, p \text{ does not occur in } x\}.$$

For a single pattern p , we abbreviate $\mathcal{R}(X, \{p\})$ as $X \setminus p$. The shift X is said to be *specified* by the collection \mathcal{F} if $X = \mathcal{R}(\mathcal{A}^G, \mathcal{F})$.

2.4. Shifts of finite type. In this section, we define shifts of finite type and sofic shifts over G . We also discuss many related elementary facts.

Definition 2.16 (SFTs). A subshift $X \subset \mathcal{A}^G$ is a *shift of finite type (SFT)* if there is a finite collection $\mathcal{F} \subset \mathcal{A}^*$ such that $X = \mathcal{R}(\mathcal{A}^G, \mathcal{F})$. For an SFT, it is always possible to take \mathcal{F} in the form $\mathcal{F} = \mathcal{A}^K \setminus \mathcal{P}(K, X)$ for some large finite subset $K \subset G$. In this case, we say X is *specified* by (patterns of shape) K .

If $X \subset \mathcal{A}^G$ is an SFT specified by a finite subset $K \subset G$, then it holds that

$$x \in X \iff \forall g \in G ((\sigma^g x)(K) \in \mathcal{P}(K, X))$$

for each $x \in \mathcal{A}^G$. If K specifies X , then so does K' for any (finite) subset $K' \supset K$. If X and T are SFTs, then so is the dynamical direct product $X \times T$.

Definition 2.17 (Sofic shifts). A subshift W is *sofic* if there exists an SFT X which factors onto W .

The following elementary facts are needed; we abbreviate the proofs as they are similar to the well-known the proofs in the case where $G = \mathbb{Z}$ (see [14]).

Proposition 2.18. *Let X be an SFT, let W be a sofic shift, and let $\phi : X \rightarrow W$ be a factor map. Then there exists an SFT \tilde{X} and a topological conjugacy $\tilde{\phi} : \tilde{X} \rightarrow X$ such that the composition $\phi \circ \tilde{\phi} : \tilde{X} \rightarrow W$ is a one-block code.*

Proof. Because ϕ is continuous and shift-commuting, there exists a large finite subset $K \subset G$ such that for each $x, x' \in X$ and each $g \in G$, it holds that

$$(\sigma^g x)(K) = (\sigma^g x')(K) \implies \phi(x)_g = \phi(x')_g.$$

Suppose that $e \in K$ and that $\mathcal{P}(K, X)$ specifies X as an SFT. Let $\tilde{\mathcal{A}} = \mathcal{P}(K, X)$ be a new finite alphabet, and let $\tilde{X} \subset \tilde{\mathcal{A}}^G$ be the set of all points $\tilde{x} \in \tilde{\mathcal{A}}^G$ such that

$$\exists x \in X, \forall g \in G, \tilde{x}_g = (\sigma^g x)(K).$$

Then \tilde{X} is an SFT specified by patterns of shape $K^{-1}K$. The map $\tilde{\phi} : \tilde{X} \rightarrow X$ desired for the theorem is given by

$$\tilde{\phi}(\tilde{x})_g = (\tilde{x}_g)_e \in \mathcal{A}, \quad \forall g \in G, \forall \tilde{x} \in \tilde{X}.$$

□

Proposition 2.19. *For any subshift $X \subset \mathcal{A}^G$, there is a descending family of SFTs $(X_n)_n$ such that $X = \bigcap_n X_n$.*

Proof. Let $(p_n)_n$ enumerate $\{p \in \mathcal{A}^* : p \text{ does not occur in } X\}$, and for each n let

$$X_n = \mathcal{R}(\mathcal{A}^G, \{p_1, p_2, \dots, p_n\}).$$

Then $(X_n)_n$ witnesses the result. □

Proposition 2.20. *Let $X \subset \mathcal{A}^G$ be a subshift and let $X_0 \subset \mathcal{A}^G$ be an SFT such that $X \subset X_0$. If $(X_n)_n$ is any descending family of subshifts such that $X = \bigcap_n X_n$, then $X_n \subset X_0$ for all sufficiently large n .*

Proof. Take $K \subset G$ to specify X_0 as an SFT. Note $(\mathcal{P}(K, X_n))_n$ is a descending family of finite sets, and it is therefore eventually constant. In particular, we have

$$\mathcal{P}(K, X_n) = \mathcal{P}(K, X) \subset \mathcal{P}(K, X_0)$$

for all sufficiently large n . □

When $G = \mathbb{Z}^d$, SFTs are often reduced via conjugacy to so-called *1-step* SFTs, in which the allowed patterns are specified by a d -hypercube of side-length 1. Such SFTs are often desired because they allow for a kind of “surgery” of patterns. If two patterns occur in two different labellings from a 1-step SFT, and yet they agree on their 1-boundaries, then the first may be *excised* and *replaced* by the second. This yields a new labelling which also belongs to the 1-step SFT. Although there is no obvious notion of 1-step SFTs when $G \neq \mathbb{Z}^d$, we do have the following result which allows for this sort of excision and replacement of patterns.

Lemma 2.21. *Let $X \subset \mathcal{A}^G$ be an SFT specified by $K \subset G$, let $F \subset G$ be a finite subset, and let $x, y \in X$ be two points such that x and y agree on $\partial_{KK^{-1}}F$. Then the point z , defined by $z_g = y_g$ if $g \in F$ and $z_g = x_g$ if $g \notin F$, also belongs to X .*

Proof. Let $g \in G$. By Lemma 2.6, either $Kg \subset F$ or $Kg \subset (\text{int}_{KK^{-1}}F)^c$. In the first case, we have $(\sigma^g z)(K) = (\sigma^g y)(K)$ which is an allowed pattern in X . In the second case, we have $Kg \subset (F^c) \sqcup (\partial_{KK^{-1}}F)$. Since x and y agree on $\partial_{KK^{-1}}F$, we have $(\sigma^g z)(K) = (\sigma^g x)(K)$ which is again an allowed pattern in X . In either case, $(\sigma^g z)(K)$ is allowed in X for every g , hence $z \in X$. □

2.5. Entropy. Let $X \subset \mathcal{A}^G$ be a nonempty subshift. Recall that for a given large finite set $F \subset G$, the number of patterns of shape F that occur in any point of X is $|\mathcal{P}(F, X)|$, which is at most $|\mathcal{A}|^{|F|}$. As this grows exponentially (with respect to $|F|$), we are interested in the *exponential growth rate* of $|\mathcal{P}(F, X)|$ as F becomes very large and approaches the whole group G . For nonempty finite sets $F \subset G$, we let

$$h(F, X) = \frac{1}{|F|} \log |\mathcal{P}(F, X)|.$$

If $F, F' \subset G$ are disjoint finite subsets, then $h(F \sqcup F', X) \leq h(F, X) + h(F', X)$. This is because $|\mathcal{P}(F \sqcup F', X)| \leq |\mathcal{P}(F, X)| \cdot |\mathcal{P}(F', X)|$ and

$$\frac{1}{|F \sqcup F'|} = \frac{1}{|F| + |F'|} \leq \min\left(\frac{1}{|F|}, \frac{1}{|F'|}\right).$$

Definition 2.22 (Entropy). Let X be a nonempty subshift. The (*topological*) entropy of X is the nonnegative real number $h(X)$ given by the limit

$$h(X) = \lim_{n \rightarrow \infty} h(F_n, X),$$

where $(F_n)_n$ is again the Følner sequence of G . For the empty subshift, we adopt the convention that $h(\emptyset) = 0$.

It is well-known that the limit above exists, does not depend on the choice of Følner sequence for G , and is an invariant of topological conjugacy (see [11]).

For any subshift $X \subset \mathcal{A}^G$ and any finite subset $F \subset G$ it holds that $h(F, X) \leq \log |\mathcal{A}|$, and consequently $h(X) \leq \log |\mathcal{A}|$. More generally, if X and X' are subshifts such that $X \subset X'$, then $h(F, X) \leq h(F, X')$ for every finite subset $F \subset G$ and consequently $h(X) \leq h(X')$. If X and X' are subshifts over \mathcal{A} , then so is $X \cup X'$ and $h(X \cup X') = \max(h(X), h(X'))$.

The following proposition is a classical fact; a proof is given in [11].

Proposition 2.23. *Let G be a countable amenable group. If a G -shift W is a factor of a G -shift X , then $h(W) \leq h(X)$.*

Frequently in this paper we refer to “measuring” or *approximating* the entropy of a subshift via a large set F . We give a precise definition as follows.

Definition 2.24. Let $X \subset \mathcal{A}^G$ be a subshift, and let $\delta > 0$. A finite subset $F \subset G$ is said to δ -*approximate* the entropy of X if

$$h(X) - \delta < h(F, X) < h(X) + \delta.$$

We shall more commonly write $h(X) < h(F, X) + \delta < h(X) + 2\delta$.

Infinitely many such sets exist for any δ , as provided by the Følner sequence and the definition of $h(X)$. We introduce this notion so that we may layer invariance conditions and entropy-approximating conditions as needed.

Proposition 2.25. *For finitely many choices of i , let $K_i \subset G$ be any finite subsets, and let $\varepsilon_i > 0$ be any positive constants. For finitely many choices of j , let $X_j \subset \mathcal{A}_j^G$ be any subshifts over any finite alphabets, and let $\delta_j > 0$ be any positive constants. There exists a finite subset $F \subset G$ which is (K_i, ε_i) -invariant for every i , and which δ_j -approximates the entropy of X_j for every j .*

Proof. Choose $F = F_n$ for sufficiently large n . □

The following theorem is an elementary generalization of a classical statement (see [14] for a proof in the case where $G = \mathbb{Z}$). We omit the proof here for brevity.

Proposition 2.26. *Let $(X_n)_n$ be a descending family of subshifts, and let $X = \bigcap_n X_n$. Then*

$$h(X) = \lim_{n \rightarrow \infty} h(X_n).$$

It is desirable to work with SFTs as much as possible while preserving (or, in our case, approximating) relevant dynamical quantities. We shall make frequent use of the next theorem, which we justify with several of the above results.

Theorem 2.27. *Let $X \subset \mathcal{A}^G$ be a subshift and suppose that $X_0 \subset \mathcal{A}^G$ is an SFT such that $X \subset X_0$. For any $\varepsilon > 0$, there exists an SFT $Z \subset \mathcal{A}^G$ such that $X \subset Z \subset X_0$ and $h(X) \leq h(Z) < h(X) + \varepsilon$.*

Proof. By Proposition 2.19, there is a descending family of SFTs $(X_n)_n$ such that $X = \bigcap_n X_n$. By Proposition 2.20, we have $X_n \subset X_0$ for all sufficiently large n . By Proposition 2.26, we have $h(X) \leq h(X_n) < h(X) + \varepsilon$ for all sufficiently large n . Choose $Z = X_n$ for n large enough to meet both conditions. \square

If $\phi : X \rightarrow W$ is a factor map of subshifts, then we have already seen that $h(W) \leq h(X)$. The “entropy drop” or *entropy gap* between X and W is the quantity $h(X) - h(W)$. A subsystem $X' \subset X$ induces a corresponding subsystem $\phi(X') = W' \subset W$, and later in this paper we will want a uniform bound for the entropy gap between every X' and W' pair. We make this idea precise in the following definition.

Definition 2.28. Suppose $\phi : X \rightarrow W$ is a factor map. The *maximal entropy gap* of ϕ is the quantity

$$\mathcal{H}(\phi) = \sup_{X'} (h(X') - h(\phi(X'))),$$

where the supremum is taken over all subshifts $X' \subset X$. In particular, it holds that

$$h(W) \leq h(X) \leq h(W) + \mathcal{H}(\phi).$$

Recall that if X and T are subshifts, then the dynamical direct product $X \times T$ factors onto both X and T via the projection map(s) $\pi_X(x, t) = x$ and $\pi_T(x, t) = t$.

Proposition 2.29. *Let X and T be shift spaces. The maximal entropy gap of the projection map $\pi_X : X \times T \rightarrow X$ is*

$$\mathcal{H}(\pi_X) = h(T).$$

Proof. It is classically known that $h(X \times T) = h(X) + h(T)$, in which case $h(T) = h(X \times T) - h(X) \leq \mathcal{H}(\pi_X)$. For the converse inequality, suppose $Z \subset X \times T$ is any subshift. Note by Definition 2.9 that $z \in Z$ implies $z = (z^X, z^T)$, where $z^X = \pi_X(z) \in \pi_X(Z) \subset X$ and $z^T \in T$. Therefore $Z \subset \pi_X(Z) \times T$, in which case it follows that $h(Z) \leq h(\pi_X(Z)) + h(T)$. Since Z was arbitrary, we have

$$h(T) \leq \mathcal{H}(\pi_X) = \sup_Z (h(Z) - h(\pi_X(Z))) \leq h(T),$$

where the supremum is taken over all subshifts $Z \subset X \times T$. \square

A quick corollary is that when $h(T) = 0$, we have $h(Z) = h(\pi_X(Z))$ for any subsystem $Z \subset X \times T$.

3. TILINGS OF AMENABLE GROUPS

3.1. Definition and encoding. In this section we consider the notion of *tilings* of G . The existence of tilings of G with certain properties is essential in our constructions in subsequent sections.

Definition 3.1 (Quasi-tilings and exact tilings). A *quasi-tiling* of G is a pair (\mathcal{S}, C) , where \mathcal{S} is a finite collection of finite subsets of G (called the *shapes* of the tiling) and C is a function that assigns each shape $S \in \mathcal{S}$ to a subset $C(S) \subset G$, called the set of *centers* or *center-set* attributed to S . We require that e is in S for each $S \in \mathcal{S}$. The following properties are also required.

- i. For distinct shapes $S, S' \in \mathcal{S}$, the subsets $C(S)$ and $C(S')$ are disjoint.
- ii. The shapes in \mathcal{S} are “translate-unique”, in the sense that

$$S \neq S' \implies Sg \neq S', \quad \forall g \in G,$$

for each $S, S' \in \mathcal{S}$.

- iii. The map $(S, c) \mapsto Sc \subset G$ defined on the domain $\{(S, c) : S \in \mathcal{S} \text{ and } c \in C(S)\}$ is injective.

We may refer to both the pair (\mathcal{S}, C) and the collection

$$\mathcal{T} = \mathcal{T}(\mathcal{S}, C) = \{Sc \subset G : S \in \mathcal{S} \text{ and } c \in C(S)\}$$

as “*the quasi-tiling*.” Each subset $\tau = Sc \in \mathcal{T}$ is called a *tile*. For a quasi-tiling \mathcal{T} , we denote the union of all the tiles by $\bigcup \mathcal{T}$. A quasi-tiling \mathcal{T} may not necessarily cover G in the sense that $\bigcup \mathcal{T} = G$; nor is it necessary for any two distinct tiles $\tau, \tau' \in \mathcal{T}$ to be disjoint. However, if both of these conditions are met (that is, if \mathcal{T} is a *partition* of G), then \mathcal{T} is called an *exact tiling* of G .

Ornstein and Weiss [15] previously constructed quasi-tilings of G with good dynamical properties, and this construction has become a fundamental tool for analyzing the dynamics of G -actions. Downarowicz, Huczczek, and Zhang [7] sharpened this construction, showing that a countable amenable group exhibits many *exact* tilings with good dynamical properties, as we describe below (see Theorem 3.6).

A quasi-tiling \mathcal{T} of G may be encoded in symbolic form, allowing for dynamical properties to be attributed to and studied for quasi-tilings. The encoding method presented here differs from the one presented in [7], as we will only require exact tilings in this paper. See Remark 3.4 below for further discussion of the relation between our encoding and the encoding given in [7].

Definition 3.2 (Encoding). Let \mathcal{S} be a finite collection of finite shapes, and let

$$\Sigma(\mathcal{S}) = \{(S, s) : s \in S \in \mathcal{S}\},$$

which we view as a finite alphabet. If \mathcal{T} is an exact tiling of G over \mathcal{S} , then it corresponds to a unique point $t \in \Sigma^G$ as follows. For each $g \in G$, there is a unique tile $Sc \in \mathcal{T}$ containing g ; let $s = gc^{-1} \in S$ and set $t_g = (S, s)$.

In the above definition, note that s is the “relative position” of g in the translate Sc of S . In other words, t labels each element g of G with both the type of shape of the tile containing g and the relative position of g within that tile. In particular, $g \in C(S) \iff t_g = (S, e)$.

Note that the correspondence $\mathcal{T} \mapsto t \in \Sigma^G$, when regarded as a map on the set of all exact tilings of G over \mathcal{S} , is injective. However, the correspondence is not surjective in general. Let $\Sigma_E \subset \Sigma^G$ be the set of all encodings of exact tilings of G

over \mathcal{S} . It may be the case that *no* exact tiling of G over \mathcal{S} exists, in which case $\Sigma_E = \emptyset$. In general, we have the following useful theorem.

Proposition 3.3. *Let \mathcal{S} be a finite collection of finite shapes drawn from G . Then $\Sigma_E(\mathcal{S}) \subset \Sigma(\mathcal{S})^G$ is an SFT.*

Proof. Let Σ_1 be the set of all points $t \in \Sigma^G$ that satisfy the following local rule: for each $g \in G$, if $t_g = (S_0, s_0) \in \Sigma$ then

$$(R1) \quad t_{sc} = (S_0, s), \quad \forall s \in S_0,$$

where $c = s_0^{-1}g$. It is easy to see that Σ_1 is an SFT, and from Definition 3.2 it is immediate that $\Sigma_E \subset \Sigma_1$.

For the reverse inclusion, let $t \in \Sigma_1$ be an arbitrary point satisfying the local rule (R1) everywhere. For each $S \in \mathcal{S}$, let $C(S) = \{g \in G : t_g = (S, e)\}$. Then $\mathcal{T} = \mathcal{T}(\mathcal{S}, C)$ is a quasi-tiling. To complete the proof, it suffices to show that \mathcal{T} is exact and encoded by t , since that would give $t \in \Sigma_E$ and then $\Sigma_E(\mathcal{S}) = \Sigma_1$.

Let $g \in G$, suppose $t_g = (S, s)$, and let $c = s^{-1}g$. By rule (R1) and the fact that $e \in S$, we have $t_c = t_{ec} = (S, e)$ and therefore $c \in C(S)$. Hence, $g = sc \in Sc \in \mathcal{T}$. This demonstrates that $\bigcup \mathcal{T} = G$. Next, suppose $Sc, S'c' \in \mathcal{T}$ are not disjoint and let $g \in Sc \cap S'c'$. Then $g = sc = s'c'$ for some $s \in S$ and $s' \in S'$. From $c \in C(S)$ we have $t_c = (S, e)$, and by the rule (R1) we have

$$t_g = t_{sc} = (S, s).$$

By identical proof we have $t_g = (S', s')$, from which it follows that $S = S'$ and $s = s'$. The latter implies that

$$c = s^{-1}g = s'^{-1}g = c',$$

and hence Sc and $S'c'$ are the same tile. This demonstrates that \mathcal{T} is a partition of G , and therefore \mathcal{T} is an exact tiling of G over \mathcal{S} . Finally, we note that it is straightforward to check that \mathcal{T} is encoded by t , which completes the proof. \square

Remark 3.4. Before we move on, we note here that the encoding method presented above (Definition 3.2) differs from the one presented in [7]. The encoding method in that work gives symbolic encodings for all quasi-tilings, which is not necessary for our present purposes. Indeed, the encoding in [7] uses the alphabet $\Lambda = \mathcal{S} \cup \{0\}$, and a point $\lambda \in \Lambda^G$ encodes a quasi-tiling (\mathcal{S}, C) when $\lambda_g = S \iff g \in C(S)$ and $\lambda_g = 0$ otherwise. This is a prudent encoding method for the study of general quasi-tilings, as any quasi-tiling may be encoded in this manner. Our encoding method works only for exact tilings, but is well-suited to our purposes. In fact, if one is only interested in exact tilings, then the two encodings are equivalent. Indeed, if $\Lambda_E \subset \Lambda^G$ is the collection of all encodings of exact tilings of G over \mathcal{S} , then there is a topological conjugacy $\phi : \Sigma_E \rightarrow \Lambda_E$ given by $\phi(t)_g = S \iff t_g = (S, e)$ and $\phi(t)_g = 0$ otherwise.

Next we turn our attention to the dynamical properties of tilings, as derived from their encodings.

Definition 3.5 (Dynamical tiling system). Let \mathcal{S} be a finite collection of finite shapes, let \mathcal{T} be an exact tiling of G over \mathcal{S} , and let \mathcal{T} be encoded by the point $t \in \Sigma_E(\mathcal{S})$. The *dynamical tiling system* generated by \mathcal{T} is the subshift generated by t in Σ^G , denoted $\Sigma_{\mathcal{T}} = \overline{\mathcal{O}}(t) \subset \Sigma_E$.

This allows for the dynamical properties (e.g., entropy) of $\Sigma_{\mathcal{T}}$ as a subshift of Σ^G to be ascribed to \mathcal{T} . The *tiling entropy* of \mathcal{T} is $h(\mathcal{T}) = h(\Sigma_{\mathcal{T}})$, the entropy of $\Sigma_{\mathcal{T}}$ as a subshift of Σ^G .

The tiling entropy of \mathcal{T} is a measure of the “complexity” of tile patterns that occur in large regions of G . In particular, when \mathcal{T} has entropy zero, the number of ways to cover a large region $F \subset G$ by tiles in \mathcal{T} grows subexponentially (with respect to $|F|$).

The following theorem is quickly deduced from the main result of Downarowicz, Huczek, and Zhang [7], which we state in this form for convenience. It is this result that allows us to utilize exact tilings of G in this paper.

Theorem 3.6 ([7]). *Let $K \subset G$ be a finite subset, and let $\varepsilon > 0$. Then there exists a finite collection of finite shapes \mathcal{S} with the following properties.*

- i. Each shape $S \in \mathcal{S}$ is (K, ε) -invariant.*
- ii. $K \subset S$ and $|S| > \varepsilon^{-1}$ for each shape $S \in \mathcal{S}$.*
- iii. There exists a point $t_0 \in \Sigma_E(\mathcal{S})$ such that $h(\overline{\mathcal{O}}(t_0)) = 0$.*

The point t_0 encodes an exact tiling \mathcal{T}_0 of G over \mathcal{S} with tiling entropy $h(\mathcal{T}_0) = 0$.

3.2. Approximating sets with tiles. Entropy and other dynamical properties of G -shifts are well measured by sets with strong invariance properties (the Følner sequence F_n provides a wealth of such sets). However, we would instead like to utilize an (appropriately selected) exact tiling \mathcal{T} for this purpose. In this section, we build good *tile approximations* of sets: finite collections of tiles $\mathcal{T}^* \subset \mathcal{T}$ attributed to large, suitably invariant subsets $F \subset G$ that are *good* in the sense that the symmetric difference $F \Delta \bigcup \mathcal{T}^*$ is small (as a proportion of $|F|$).

Definition 3.7 (Tile approximation). Let $F \subset G$ be a finite subset. An exact tiling \mathcal{T} of G induces two finite collections of tiles: the *outer approximation* of F by \mathcal{T} , denoted

$$\mathcal{T}^\times(F) = \{\tau \in \mathcal{T} : \tau \cap F \neq \emptyset\},$$

and the *inner approximation* of F by \mathcal{T} , denoted

$$\mathcal{T}^\circ(F) = \{\tau \in \mathcal{T} : \tau \subset F\}.$$

Denote $F^\times(\mathcal{T}) = \bigcup \mathcal{T}^\times(F)$ and $F^\circ(\mathcal{T}) = \bigcup \mathcal{T}^\circ(F)$. Observe that $F^\circ \subset F \subset F^\times$.

Lemma 3.8. *Let \mathcal{S} be a finite collection of shapes from G , and let $U = \bigcup \mathcal{S}$. Let $\varepsilon > 0$, and choose $\delta > 0$ such that $\delta|U||UU^{-1}| < \varepsilon$. Let $F \subset G$ be a finite subset that is (UU^{-1}, δ) -invariant. For any exact tiling \mathcal{T} of G over \mathcal{S} , the following statements hold:*

- i. $|F^\times(\mathcal{T}) \setminus F^\circ(\mathcal{T})| < \varepsilon|F|$,*
- ii. $(1 - \varepsilon)|F| < |F^\circ(\mathcal{T})| \leq |F|$, and*
- iii. $|F| \leq |F^\times(\mathcal{T})| < (1 + \varepsilon)|F|$.*

Proof. First, we observe that each tile $\tau \in \mathcal{T}$ is contained in a translate Ug for some $g \in G$; indeed, we have $\tau = Sc$ for some $S \in \mathcal{S}$ and $c \in C(S) \subset G$, then $S \subset U$ implies $\tau \subset Uc$. This fact also gives that $|\tau| \leq |U|$ for every tile $\tau \in \mathcal{T}$.

We claim that every tile $\tau \in \mathcal{T}^\times(F) \setminus \mathcal{T}^\circ(F)$ intersects $\partial_{UU^{-1}}F$. To establish the claim, we first note that for each such tile τ it holds that $\tau \cap F \neq \emptyset$ and $\tau \not\subset F$. So let $f \in \tau \cap F$, and note that $f \in \tau \subset Ug$ for some $g \in G$. From $\tau \not\subset F$ we

also have $Ug \not\subset F$. By Lemma 2.6 we have $f \in Ug \subset (\text{int}_{UU^{-1}} F)^c$, and hence $f \in F \setminus (\text{int}_{UU^{-1}} F) = \partial_{UU^{-1}} F$, which establishes our claim.

By the claim in the previous paragraph, there is a map $\gamma : \mathcal{T}^\times \setminus \mathcal{T}^\circ \rightarrow \partial_{UU^{-1}} F$ with the property that $\gamma(\tau) \in \tau$ for each τ . Observe that γ is injective, as distinct tiles are disjoint, and therefore $|\mathcal{T}^\times \setminus \mathcal{T}^\circ| \leq |\partial_{UU^{-1}} F|$. We also have that

$$|\partial_{UU^{-1}} F| < \delta |UU^{-1}| |F|,$$

by the invariance hypothesis on F and Lemma 2.5. Then

$$\begin{aligned} |F^\times \setminus F^\circ| &= \sum_{\tau \in \mathcal{T}^\times \setminus \mathcal{T}^\circ} |\tau| \\ &\leq |\mathcal{T}^\times \setminus \mathcal{T}^\circ| |U| \\ &\leq |\partial_{UU^{-1}} F| |U| \\ &< \delta |U| |UU^{-1}| |F| \\ &< \varepsilon |F|. \end{aligned}$$

This establishes statement (i). The remaining two statements are easy to check using $F^\times = F^\circ \sqcup (F^\times \setminus F^\circ)$, $F^\circ \subset F \subset F^\times$, and statement (i). \square

One more notion is necessary to develop before moving on from tilings: the *frame* of a given subset with respect to a given tiling.

Definition 3.9 (Frame of a tiling). Let $F, K \subset G$, and let \mathcal{T} be an exact tiling of G . The inner (\mathcal{T}, K) -*frame* of F is the subset

$$\text{fr}_{\mathcal{T}, K}(F) = \bigcup_{\tau} \partial_K(\tau),$$

where the union ranges over all $\tau \in \mathcal{T}^\circ(F)$. See Figure 1 for an illustration.

4. RESULTS FOR SFTS

Having discussed everything about tilings relevant for our purposes, we are now ready to begin discussing our main results. In this section we present our results for SFTs, and in the following section we turn our attention to sofic shifts.

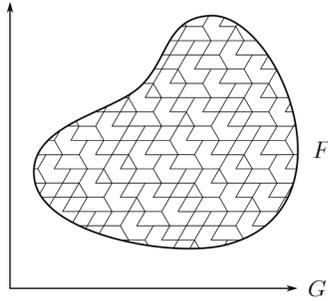
Theorem 4.1. *Let G be a countable amenable group, and let X be a G -SFT such that $h(X) > 0$. Then*

$$\{h(Y) : Y \subset X \text{ and } Y \text{ is an SFT}\}$$

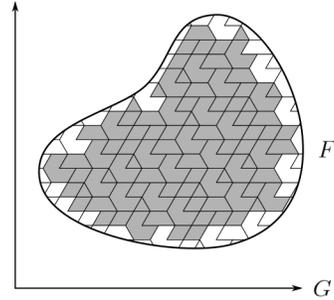
is dense in $[0, h(X)]$.

Before we begin the proof, let us give a short outline of the main ideas. The broad strokes of this proof come from Desai [6], whose argument in the case where $G = \mathbb{Z}^d$ we are able to extend to the case where G is an arbitrary countable amenable group. This is possible by utilizing the exact tilings of G constructed by Downarowicz, Huczek and Zhang [7].

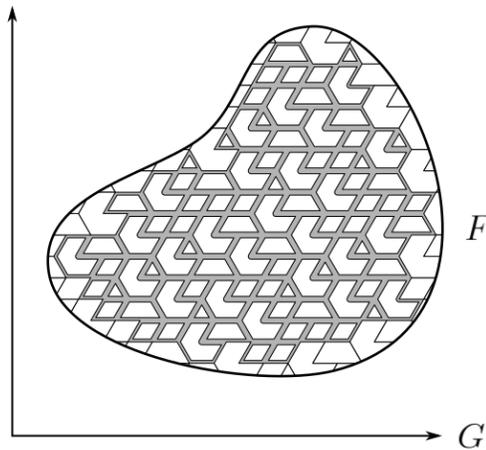
Given an arbitrary $\varepsilon > 0$, we produce a family of SFT subshifts of X whose entropies are 2ε -dense in $[0, h(X)]$. We accomplish this by first selecting an exact, zero entropy tiling \mathcal{T}_0 of G with suitably large, invariant tiles. Then we build subshifts with strongly controlled entropies inside the product system $Z_0 = X \times \Sigma_0$, where Σ_0 is the dynamical tiling system generated by \mathcal{T}_0 .



(A) A hypothetical region $F \subset G$ with illustrated tiling \mathcal{T} .



(B) The \mathcal{T} -interior of F is shaded.



(C) The inner frame of F (with respect to \mathcal{T} , K) is shaded. The K -boundary of each tile inside F is taken.

FIGURE 1. A sketch of the construction of $\text{fr}_{\mathcal{T},K}(F)$

To construct these subshifts from Z_0 , we control which patterns in the X layer can appear in the “interior” of the tiles in the Σ_0 layer. We are able to finely comb away entropy from Z_0 by forbidding these patterns one at a time. This process generates a descending family of subsystems for which the entropy drop between consecutive subshifts is less than ε . After enough such patterns have been forbidden, the overall entropy is less than ε . This collection of subshifts therefore has entropies that are ε -dense in $[0, h(Z_0)]$. Then we project the subshifts into X and utilize Theorem 2.27 to produce SFTs subsystems of X with entropies that are 2ε -dense in $[0, h(X)]$.

Proof. Let $X \subset \mathcal{A}^G$ be an SFT such that $h(X) > 0$, let $K \subset G$ be a large finite subset such that $\mathcal{P}(K, X)$ specifies X as an SFT, and let ε be any constant such that $0 < \varepsilon < h(X)$. Choose $\delta > 0$ such that

$$2\delta + \delta \log 2 + 2\delta \log |\mathcal{A}| < \varepsilon.$$

By Theorem 3.6, there exists a finite collection \mathcal{S} of finite subsets of G with the following properties.

- i. Each shape $S \in \mathcal{S}$ is (KK^{-1}, η) -invariant, where $\eta > 0$ is a constant such that $\eta|KK^{-1}| < \delta$. By Lemma 2.5, this implies that $|\partial_{KK^{-1}}S| < \delta|S|$ for each shape $S \in \mathcal{S}$.
- ii. $KK^{-1} \subset S$ and $|S| > \delta^{-1}$ for each $S \in \mathcal{S}$.
- iii. There is a point $t_0 \in \Sigma_E(\mathcal{S})$ such that $h(\overline{\mathcal{O}}(t_0)) = 0$. Consequently, t_0 encodes an exact tiling \mathcal{T}_0 of G over \mathcal{S} with tiling entropy zero.

For the remainder of this proof, these are all fixed. We shall abbreviate $\partial F = \partial_{KK^{-1}}F$ for any finite subset $F \subset G$. For a pattern p on F , we take ∂p to mean $p(\partial F)$ and call this the *border* of p (with respect to KK^{-1}).

Let $\Sigma_0 = \overline{\mathcal{O}}(t_0) \subset \Sigma_E$ be the dynamical tiling system generated by the tiling \mathcal{T}_0 , which has entropy zero. Of central importance to this proof is the product system $X \times \Sigma_0$, which factors onto X via the projection map $\pi : X \times \Sigma_0 \rightarrow X$ given by $\pi(x, t) = x$ for each $(x, t) \in X \times \Sigma_0$. Let us establish some terminology for certain patterns of interest which occur in this system.

Given a shape $S \in \mathcal{S}$, we shall refer to a pattern $b = (b^X, b^{\mathcal{T}}) \in \mathcal{P}(S, X \times \Sigma_0)$ as a *block* (to distinguish from patterns of any general shape). If a block $b \in (\mathcal{A} \times \Sigma)^S$ satisfies $b_s^{\mathcal{T}} = (S, s)$ for every $s \in S$, then we shall say b is *aligned*. See Figure 2 for an illustration of the aligned property.

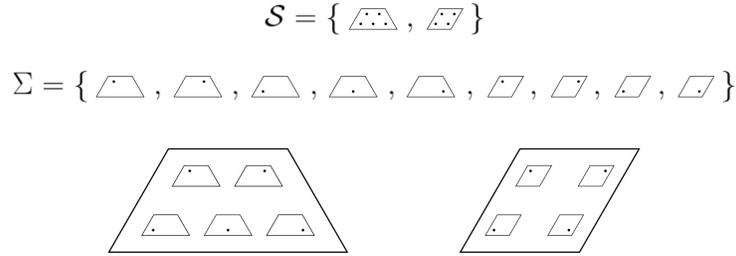


FIGURE 2. A hypothetical collection of shapes \mathcal{S} , the appropriate alphabet $\Sigma = \Sigma(\mathcal{S})$, and (the \mathcal{T} -layer of) two aligned blocks are pictured. Each point of each block is labelled with the correct shape type and relative displacement within that shape.

For a subshift $Z \subset X \times \Sigma_0$, we denote the subcollection of aligned blocks of shape S that occur in Z by

$$\mathcal{P}^a(S, Z) \subset \mathcal{P}(S, Z) \subset (\mathcal{A} \times \Sigma)^S,$$

where the superscript a identifies the subcollection. Given a shape $S \in \mathcal{S}$ and an aligned block b of shape S , consider the border $\partial b \in (\mathcal{A} \times \Sigma)^{\partial S}$. We are interested in the number of ways that the border ∂b may be extended to all of S - that is, the

number of allowed (and in particular, aligned) *interiors* for S which agree with ∂b on the boundary ∂S . For a subshift $Z \subset X \times \Sigma_0$, we denote this collection by

$$\text{ints}^a(\partial b, Z) = \{b' \in \mathcal{P}^a(S, Z) : \partial b' = \partial b\}.$$

We shall extend all the same terminology described above (blocks, aligned blocks, borders, interiors) to tiles $\tau = Sc \in \mathcal{T}_0$, as there is a bijection between $\mathcal{P}(S, Z)$ and $\mathcal{P}(\tau, Z) = \mathcal{P}(Sc, Z)$. For a given tile $\tau \in \mathcal{T}_0$, a block $b \in (\mathcal{A} \times \Sigma)^\tau$ is aligned if $b_{sc}^T = (S, s)$ for each $sc \in Sc = \tau$. The subcollection of aligned blocks of shape τ occurring in a shift $Z \subset X \times \Sigma_0$ is denoted $\mathcal{P}^a(\tau, Z)$. Given a border $\partial b \in (\mathcal{A} \times \Sigma)^{\partial\tau}$, the collection of aligned blocks of shape τ occurring in Z agreeing with ∂b on $\partial\tau$ is also denoted $\text{ints}^a(\partial b, Z) \subset \mathcal{P}^a(\tau, Z)$.

For the theorem, we shall inductively construct a descending family of subshifts $(Z_n)_n$ of $X \times \Sigma_0$ as follows. Begin with $Z_0 = X \times \Sigma_0$, then assume Z_n has been constructed for $n \geq 0$. If there exists a shape $S_n \in \mathcal{S}$ and an aligned block $\beta_n \in \mathcal{P}^a(S_n, Z_n)$ such that

$$|\text{ints}^a(\partial\beta_n, Z_n)| > 1,$$

then let $Z_{n+1} = Z_n \setminus \beta_n$. If no such block exists on any shape $S \in \mathcal{S}$, then Z_n is the final subshift in the chain and the chain is finite in length.

Let us first argue that in fact, the chain *must* be finite in length. For each $n \geq 0$ we have $Z_{n+1} \subset Z_n$, in which case $\mathcal{P}^a(S, Z_{n+1}) \subset \mathcal{P}^a(S, Z_n)$ for every shape $S \in \mathcal{S}$. Moreover, for the distinguished shape S_n (the shape of the forbidden block β_n), it holds that $\mathcal{P}^a(S_n, Z_{n+1}) \sqcup \{\beta_n\} \subset \mathcal{P}^a(S_n, Z_n)$. This implies that

$$\sum_{S \in \mathcal{S}} |\mathcal{P}^a(S, Z_n)|$$

strictly decreases with n . There is no infinite strictly decreasing sequence of positive integers, hence the descending chain must be finite in length. Let $N \geq 0$ be the index of the terminal subshift, and note by construction that the shift Z_N satisfies

$$|\text{ints}^a(\partial b, Z_N)| = 1$$

for every aligned block $b \in \mathcal{P}^a(S, Z_N)$ on any shape $S \in \mathcal{S}$.

Most of the rest of the proof aims to establish the following two statements:

$$(U1) \quad h(Z_{n+1}) \leq h(Z_n) < h(Z_{n+1}) + \varepsilon \quad \text{for each } n < N, \text{ and}$$

$$(U2) \quad h(Z_N) < \varepsilon.$$

To begin, let $F \subset G$ be a finite subset satisfying the following two conditions:

(F1) F is (UU^{-1}, ϑ) -invariant, where $U = \bigcup \mathcal{S}$ and ϑ is a positive constant such that $\vartheta|U||UU^{-1}| < \delta$. Note this implies that F may be well approximated by tiles from any exact tiling of G over \mathcal{S} , in the sense of Lemma 3.8.

(F2) F is large enough to δ -approximate (Definition 2.24) the entropy of Σ_0 and Z_n for every $n \leq N$. This implies in particular that $h(F, \Sigma_0) < \delta$.

Such a set exists by Proposition 2.25. We fix F for the remainder of this proof.

Now for each $n \leq N$, we claim that

$$(E1) \quad |\mathcal{P}(F, Z_n)| \geq \sum_t \sum_f \prod_\tau |\text{ints}^a(f(\partial\tau), Z_n)|, \quad \text{and}$$

$$(E2) \quad |\mathcal{P}(F, Z_n)| \leq |\mathcal{A}|^{\delta|F|} \cdot \sum_t \sum_f \prod_\tau |\text{ints}^a(f(\partial\tau), Z_n)|,$$

where the indices t , f , and τ are as follows. The variable t ranges over $\mathcal{P}(F, \Sigma_0)$, and therefore t is the restriction to F of an encoding of an exact, zero entropy tiling \mathcal{T}_t of G over \mathcal{S} . The variable f ranges over all $(\mathcal{A} \times \Sigma)$ -labellings of the (\mathcal{T}_t, KK^{-1}) -frame of F (Definition 3.9) that are allowed in Z_0 and for which $f^\mathcal{T}$ agrees with t . Lastly, the variable τ ranges over the tiles in $\mathcal{T}_t^\circ(F)$.

To begin the argument towards the claims (E1) and (E2), let $n \leq N$ be arbitrary. To count patterns $p \in \mathcal{P}(F, Z_n)$, write $p = (p^X, p^\mathcal{T})$ and sum over all possible labellings in the tiling component. We have

$$(1) \quad |\mathcal{P}(F, Z_n)| = \sum_t |\{p \in \mathcal{P}(F, Z_n) : p^\mathcal{T} = t\}|,$$

where the sum ranges over all $t \in \mathcal{P}(F, \Sigma_0)$. This is valid because $Z_n \subset Z_0 = X \times \Sigma_0$, hence any $z = (z^X, z^\mathcal{T}) \in Z_n$ must have $z^\mathcal{T} \in \Sigma_0$.

Next, let $t \in \mathcal{P}(F, \Sigma_0)$ be fixed. The pattern t extends to/encodes an exact, zero entropy tiling \mathcal{T}_t of G over \mathcal{S} , possibly distinct¹ from the original selected tiling \mathcal{T}_0 .

Recall that $F^\circ(\mathcal{T}_t) = \bigcup \mathcal{T}_t^\circ(F) \subset F$ is the inner tile approximation of F by the tiling \mathcal{T}_t (Definition 3.7). Recall also that the (\mathcal{T}_t, KK^{-1}) -frame of F is the subset $\bigcup_\tau \partial\tau$ where the union is taken over all $\tau \in \mathcal{T}_t^\circ(F)$ (Definition 3.9). Since K is fixed for this proof, we shall abbreviate the frame as $\text{fr}_t(F)$. From Equation (1), we now split over all allowed labellings of $\text{fr}_t(F)$. We have

$$(2) \quad |\mathcal{P}(F, Z_n)| = \sum_t \sum_f |\{p \in \mathcal{P}(F, Z_n) : p^\mathcal{T} = t \text{ and } p(\text{fr}_t(F)) = f\}|,$$

where the first sum is taken over all $t \in \mathcal{P}(F, \Sigma_0)$, and the second sum is taken over all $f \in \mathcal{P}(\text{fr}_t(F), Z_0)$ for which $f^\mathcal{T}$ agrees with t .

We have the pattern $t \in \mathcal{P}(F, \Sigma_0)$ fixed from before; next we fix a frame pattern $f \in \mathcal{P}(\text{fr}_t(F), Z_0)$ such that $f^\mathcal{T}$ agrees with t . We wish to count the number of patterns $p \in \mathcal{P}(F, Z_n)$ such that $p^\mathcal{T} = t$ and $p(\text{fr}_t(F)) = f$. Let this collection be denoted by $D_n = D(t, f; Z_n) \subset \mathcal{P}(F, Z_n)$. Observe that each D_n is finite and $D_{n+1} \subset D_n$ for each $n < N$. Consider the map $\gamma : D_0 \rightarrow \prod_\tau \mathcal{P}(\tau, Z_0)$ given by $\gamma(p) = (p(\tau))_\tau$, which sends a pattern $p \in D_0$ to a vector of blocks indexed by $\mathcal{T}_t^\circ(F)$. We claim the map γ is at most $|\mathcal{A}|^{\delta|F|}$ -to-1, and for each $n \leq N$ we have

$$(3) \quad \gamma(D_n) = \prod_\tau \text{ints}^a(f(\partial\tau), Z_n).$$

Together, these claims will provide a bound for $|D_n|$ from above and below, which combine with Equation (2) to yield the claims (E1) and (E2).

First we argue that γ is at most $|\mathcal{A}|^{\delta|F|}$ -to-1. This is where we first invoke the invariance of F . Suppose $(b_\tau)_\tau \in \prod_\tau \mathcal{P}(\tau, Z_0)$ is a fixed vector of blocks. If $p \in D_0$ is a pattern such that $\gamma(p) = (b_\tau)_\tau$, then $p^\mathcal{T}$ is determined by t and $p(\tau) = b_\tau$ for each tile $\tau \in \mathcal{T}_t^\circ(F)$. Therefore, p is uniquely determined by $p^X(F \setminus F_t^\circ)$, hence $|\gamma^{-1}(b_\tau)_\tau| \leq |\mathcal{A}|^{|F \setminus F_t^\circ|}$. By property (F1) of the set F and by Lemma 3.8 we have $|F \setminus F_t^\circ| < \delta|F|$, and thus the map γ is at most $|\mathcal{A}|^{\delta|F|}$ -to-1.

Next, we shall prove the set equality (3). Let $n \leq N$, and let $p \in D_n$. For each tile $\tau \in \mathcal{T}_t^\circ(F)$, the block $p(\tau) \in \mathcal{P}(\tau, Z_n) \subset (\mathcal{A} \times \Sigma)^\tau$ is aligned; this is because $p^\mathcal{T} = t$ and t encodes the tiling \mathcal{T}_t itself. Moreover, p agrees with f on $\text{fr}_t(F)$ by

¹As Σ_0 is generated by \mathcal{T}_0 , one may take \mathcal{T}_t to be a translation of \mathcal{T}_0 that agrees with t on F .

assumption that $p \in D_n = D(t, f; Z_n)$, in which case $p(\partial\tau) = f(\partial\tau)$ for each tile $\tau \in \mathcal{T}_t^\circ(F)$. This demonstrates that

$$\gamma(D_n) \subset \prod_{\tau} \text{ints}^a(f(\partial\tau), Z_n).$$

We shall prove the reverse inclusion by induction on n . For the $n = 0$ case, let $(b_\tau)_\tau$ be a vector of blocks such that $b_\tau \in \mathcal{P}^a(\tau, Z_0)$ and $\partial b_\tau = f(\partial\tau)$ for each $\tau \in \mathcal{T}_t^\circ(F)$. To construct a γ -preimage of $(b_\tau)_\tau$ in D_0 , begin with a point $x \in X$ such that $x(\text{fr}_t(F)) = f^X$. Such a point exists because f occurs in some point of $Z_0 = X \times \Sigma_0$. Note that

$$x(\partial\tau) = f^X(\partial\tau) = \partial b_\tau^X$$

for each $\tau \in \mathcal{T}_t^\circ(F)$, because $\partial\tau \subset \text{fr}_t(F)$ for each τ . Moreover, for each τ it holds that the block b_τ^X occurs in a point of X , as each block $b_\tau = (b_\tau^X, b_\tau^T)$ occurs in a point of $Z_0 = X \times \Sigma_0$. Because X is an SFT specified by patterns of shape K , we may repeatedly apply Lemma 2.21 to excise the block $x(\tau)$ and replace it with b_τ^X for every $\tau \in \mathcal{T}_t^\circ$. Every tile is disjoint, so the order in which the blocks are replaced does not matter. After at most finitely many steps, we obtain a new point $x' \in X$ such that $x'(\text{fr}_t(F)) = f^X$ and $x'(\tau) = b_\tau^X$ for each $\tau \in \mathcal{T}_t^\circ$.

Recall that the point $t \in \Sigma_0$ is fixed from before. The point $(x', t) \in X \times \Sigma_0$ is therefore allowed in $Z_0 = X \times \Sigma_0$. Let $p = (x', t)(F) \in \mathcal{P}(F, Z_0)$. We have that $p^T = t$ and $p(\text{fr}_t(F)) = f$ by the selection of x' . This implies that $p \in D_0$. It also holds that $p(\tau) = b_\tau$ for each $\tau \in \mathcal{T}_t^\circ(F)$, as t itself encodes the tiling \mathcal{T}_t from which the tiles $\tau \in \mathcal{T}_t^\circ(F)$ are drawn (and each block b_τ is aligned, by assumption). We then finally have $\gamma(p) = (b_\tau)_\tau$, which settles the case $n = 0$.

Now suppose the set equality (3) holds for some fixed $n < N$, and let $(b_\tau)_\tau \in \prod_{\tau} \text{ints}^a(f(\partial\tau), Z_{n+1})$. From the inclusion $Z_{n+1} \subset Z_n$ and the inductive hypothesis, it follows there is a pattern $p \in D_n$ such that $\gamma(p) = (b_\tau)_\tau$. Suppose $p = (x, t)(F)$ for some $(x, t) \in Z_n$ (by induction, t is the point fixed from before). We need to modify p only slightly to find a γ -preimage of $(b_\tau)_\tau$ which occurs in Z_{n+1} (and hence belongs to D_{n+1}).

Consider the block β_n determined at the beginning of this proof, which is forbidden in the subshift Z_{n+1} . If β_n occurs anywhere in the point (x, t) , then (by the assumption that β_n is aligned) it must occur on a tile² $\tau \in \mathcal{T}_t$. It does *not* occur on any of the tiles from $\mathcal{T}_t^\circ(F)$, because for each tile $\tau \in \mathcal{T}_t^\circ(F)$ we have $(x, t)(\tau) = b_\tau$ which is allowed in Z_{n+1} by assumption.

Yet, β_n may occur in (x, t) outside of F_t° . By the construction of Z_{n+1} , we have

$$|\text{ints}^a(\partial\beta_n, Z_n)| > 1,$$

and therefore there is an aligned block \tilde{b} which occurs in Z_n such that $\tilde{b} \neq \beta_n$ and $\partial\tilde{b} = \partial\beta_n$. Apply Lemma 2.21 at most countably many times³ to excise β_n^X wherever it may occur in x , replacing it with \tilde{b}^X . This yields a new point $x' \in X$.

Then $(x', t) \in Z_0$ also belongs to Z_{n+1} . It was already the case that none of the blocks $\beta_0, \dots, \beta_{n-1}$ could occur anywhere in (x, t) by the assumption $(x, t) \in Z_n$,

²This sentence is the reason why we consider the product system $X \times \Sigma_0$ rather than working in X directly; the extra information on the tiling layer of a labelling (x, t) allows us to control where an aligned block may occur within (x, t) .

³If x_k is the point constructed after k excisions, then $(x_k)_k$ is a Cauchy sequence in X and therefore has a limit $x_\infty \in X$, the desired point with infinitely many excisions applied. The order in which the blocks are replaced does not matter because the tiles are all disjoint.

and now neither does β_n occur anywhere in (x', t) . The pattern $p' = (x', t)(F)$ may be distinct from $p = (x, t)(F)$ (the labelling may change on $F \setminus F_t^\circ$), but we did not replace any of the blocks within F_t° . We still have $p'(\tau) = b_\tau$ for each tile $\tau \in \mathcal{T}_t^\circ(F)$, and hence $p' \in D_{n+1}$ and $\gamma(p') = (b_\tau)_\tau$.

This completes the induction, and we conclude that the set equality (3) holds for each $n \leq N$. From this equality and the fact that γ is at most $|\mathcal{A}|^{\delta|F|-1}$, we obtain

$$\prod_{\tau} |\text{ints}^a(f(\partial\tau), Z_n)| \leq |D(t, f; Z_n)| \leq |\mathcal{A}|^{\delta|F|} \cdot \prod_{\tau} |\text{ints}^a(f(\partial\tau), Z_n)|.$$

Notice that the above inequalities hold for each fixed $t \in \mathcal{P}(F, \Sigma_0)$, each fixed $f \in \mathcal{P}(\text{fr}_t(F), X \times \Sigma_0)$ such that $f^T = t(\text{fr}_t(F))$, and each $n \leq N$. From these inequalities and Equation (2), we conclude that (E1) and (E2) hold, i.e.,

$$(E1) \quad |\mathcal{P}(F, Z_n)| \geq \sum_t \sum_f \prod_{\tau} |\text{ints}^a(f(\partial\tau), Z_n)|, \quad \text{and}$$

$$(E2) \quad |\mathcal{P}(F, Z_n)| \leq |\mathcal{A}|^{\delta|F|} \cdot \sum_t \sum_f \prod_{\tau} |\text{ints}^a(f(\partial\tau), Z_n)|,$$

where the first sum is taken over all $t \in \mathcal{P}(F, \Sigma_0)$, the second sum over all $f \in \mathcal{P}(\text{fr}_t(F), Z_0)$ for which f^T agrees with t , and the product over all $\tau \in \mathcal{T}_t^\circ(F)$.

Property (F2) of the set F implies that $h(F, \Sigma_0) < \delta$, in which case $|\mathcal{P}(F, Z_0)| < e^{\delta|F|}$. Consequently, the variable t in (E1) and (E2) ranges over fewer than $e^{\delta|F|}$ terms. Moreover, by the selection of \mathcal{S} , we have $|\partial\tau| < \delta|\tau|$ for each $\tau \in \mathcal{T}_t^\circ(F)$, in which case it follows that

$$|\text{fr}_t(F)| = \left| \bigcup_{\tau} \partial\tau \right| = \sum_{\tau} |\partial\tau| < \sum_{\tau} \delta|\tau| = \delta \left| \bigcup_{\tau} \tau \right| = \delta|F_t^\circ| \leq \delta|F|.$$

Here we have used that distinct tiles from \mathcal{T}_t are disjoint. From this estimate, we deduce that there are fewer than $|\mathcal{A}|^{\delta|F|}$ labellings of the frame of F that agree with a fixed t on the \mathcal{T} -layer⁴. Consequently, the variable f in (E1) and (E2) ranges over fewer than $|\mathcal{A}|^{\delta|F|}$ terms. Observe also that the size of $\mathcal{T}_t^\circ(F)$ as a collection is small compared to F . Indeed, $|S| > \delta^{-1}$ for each shape $S \in \mathcal{S}$, in which case

$$|F| \geq |F_t^\circ| = \sum_{\tau} |\tau| \geq |\mathcal{T}_t^\circ(F)| \cdot \left(\min_{S \in \mathcal{S}} |S| \right) > |\mathcal{T}_t^\circ(F)| \delta^{-1},$$

and therefore $|\mathcal{T}_t^\circ(F)| < \delta|F|$. Consequently, the variable τ in (E1) and (E2) ranges over fewer than $\delta|F|$ terms.

Before returning to (U1) and (U2), one more estimate is necessary. For each $n < N$, each shape $S \in \mathcal{S}$, and each aligned block $b \in \mathcal{P}^a(S, Z_n)$, we claim that

$$(4) \quad |\text{ints}^a(\partial b, Z_n)| \leq 2 |\text{ints}^a(\partial b, Z_{n+1})|.$$

Let $S \in \mathcal{S}$, and let $b \in \mathcal{P}^a(S, Z_n)$ be an aligned block occurring in Z_n , distinct from the forbidden block β_n . Say $b = (x, t)(S)$ for some $(x, t) \in Z_n$. We claim that b occurs in a point of Z_{n+1} . Note that t extends to/encodes an exact, zero entropy tiling \mathcal{T}_t of G over \mathcal{S} , and note that S is a tile of \mathcal{T}_t . This is because b is aligned by assumption, in which case $t_e = b_e^T = (S, e)$, and therefore $e \in C_t(S)$.

Suppose the forbidden block β_n occurs anywhere in (x, t) . Because β_n is aligned, it must occur on a tile $\tau \in \mathcal{T}_t$. It does *not* occur on S , because $b \neq \beta_n$. By the

⁴It is important to specify that the \mathcal{T} -layer is fixed, else the number is at most $|\mathcal{A} \times \Sigma|^{\delta|F|}$, which cannot be bounded because δ must be selected before Σ is constructed.

assumption that $|\text{ints}^a(\partial\beta_n, Z_n)| > 1$, we know there is an aligned block $\tilde{b}_n \neq \beta_n$ that occurs in Z_n such that $\partial\tilde{b}_n = \partial\beta_n$.

Recall X is an SFT specified by patterns of shape K , and \tilde{b}_n^X is allowed in X . Again we may apply Lemma 2.21 at most countably many times, excising β_n^X wherever it may occur in x and replacing it with \tilde{b}_n^X . At the end we receive a new point $x' \in X$, within which β_n^X does not occur. Then (x', t) is allowed in Z_{n+1} and $(x', t)(S) = b$, hence $b \in \mathcal{P}^a(S, Z_{n+1})$.

The conclusion is that β_n is the only aligned block lost from Z_n to Z_{n+1} . For each $b \in \mathcal{P}^a(S, Z_n)$, we have either $\text{ints}^a(\partial b, Z_n) = \text{ints}^a(\partial b, Z_{n+1})$ or $\text{ints}^a(\partial b, Z_n) = \text{ints}^a(\partial b, Z_{n+1}) \sqcup \{\beta_n\}$. If two positive integers differ by at most 1 then their ratio is at most 2, hence the inequality (4) follows.

Finally, we shall use the estimates (E1) and (E2) to argue for the ultimate claims (U1) and (U2) made before. For the first, consider a fixed $n < N$. It is clear that $h(Z_{n+1}) \leq h(Z_n)$ by inclusion. For the second inequality in (U1), we have

$$\begin{aligned} |\mathcal{P}(F, Z_n)| &\leq |\mathcal{A}|^{\delta|F|} \cdot \sum_t \sum_f \prod_\tau |\text{ints}^a(f(\partial\tau), Z_n)| \\ &\leq |\mathcal{A}|^{\delta|F|} \cdot \sum_t \sum_f \prod_\tau 2 |\text{ints}^a(f(\partial\tau), Z_{n+1})| \\ &< |\mathcal{A}|^{\delta|F|} \cdot 2^{\delta|F|} \cdot \sum_t \sum_f \prod_\tau |\text{ints}^a(f(\partial\tau), Z_{n+1})| \\ &\leq |\mathcal{A}|^{\delta|F|} \cdot 2^{\delta|F|} \cdot |\mathcal{P}(F, Z_{n+1})|, \end{aligned}$$

where the inequalities are justified by (E2), (4), the fact that $|\mathcal{T}_t^\circ(F)| < \delta|F|$, and (E1), respectively. Taking logs and dividing through by $|F|$, we obtain

$$\begin{aligned} h(Z_n) &< h(F, Z_n) + \delta \\ &< (\delta \log |\mathcal{A}| + \delta \log 2 + h(F, Z_{n+1})) + \delta \\ &< \delta \log |\mathcal{A}| + \delta \log 2 + (h(Z_{n+1}) + \delta) + \delta \\ &= h(Z_{n+1}) + 2\delta + \delta \log 2 + \delta \log |\mathcal{A}| \\ &< h(Z_{n+1}) + \varepsilon, \end{aligned}$$

where we have used the property (F2) of F , the previous display, the property (F2) again, and our choice of δ . This inequality establishes (U1). For (U2), recall that the terminal shift Z_N has the property that any aligned border $\partial b \in \mathcal{P}^a(\partial S, Z_N)$ on any shape $S \in \mathcal{S}$ has exactly 1 allowed aligned interior. Hence, we see that

$$\begin{aligned} |\mathcal{P}(F, Z_N)| &\leq |\mathcal{A}|^{\delta|F|} \cdot \sum_t \sum_f \prod_\tau |\text{ints}^a(f(\partial\tau), Z_N)| \\ &= |\mathcal{A}|^{\delta|F|} \cdot \sum_t \sum_f \prod_\tau 1 \\ &< |\mathcal{A}|^{\delta|F|} \cdot e^{\delta|F|} \cdot |\mathcal{A}|^{\delta|F|} \cdot 1, \end{aligned}$$

where the first inequality is justified by (E2) and the last inequality is justified by our bounds on the number of terms in the sums (established previously). Taking logs and dividing through by $|F|$, we finally have

$$\begin{aligned} h(Z_N) &< h(F, Z_N) + \delta \\ &< (\delta + 2\delta \log |\mathcal{A}|) + \delta \end{aligned}$$

$$\begin{aligned}
&= 2\delta + 2\delta \log |\mathcal{A}| \\
&< \varepsilon,
\end{aligned}$$

where we have used the property (F2) of the set F , the previous display, and our choice of δ . We have now established (U2).

With (U1) and (U2) in hand, the rest of the proof is easy. By (U1) and (U2), we have that $(Z_n)_{n \leq N}$ is a family of subshifts of $X \times \Sigma_0$ such that $(h(Z_n))_{n \leq N}$ is ε -dense in $[0, h(X \times \Sigma_0)]$.

For each $n \leq N$, let $X_n = \pi(Z_n) \subset X$, where π is the projection map $\pi(x, t) = x$. From Lemma 2.29, $\mathcal{H}(\pi) = h(\Sigma_0) = 0$, hence $h(X_n) = h(Z_n)$ for every $n \leq N$.

Then $(X_n)_{n \leq N}$ is a descending family of subshifts of X such that $(h(X_n))_{n \leq N}$ is ε -dense in $[0, h(X)]$. Though each X_n may not be an SFT, we do know that X is an SFT. One may therefore apply Theorem 2.27 to construct a family of SFTs $(Y_n)_{n \leq N}$ such that for each $n \leq N$, we have $X_n \subset Y_n \subset X$ and $h(X_n) \leq h(Y_n) < h(X_n) + \varepsilon$. Hence $(h(Y_n))_{n \leq N}$ is 2ε -dense in $[0, h(X)]$. As ε was arbitrary, we conclude that the entropies of the SFT subsystems of X are dense in $[0, h(X)]$. \square

The following ‘‘relative’’ version of Theorem 4.1 is stronger and easily obtained as a consequence of Theorem 4.1.

Corollary 4.2. *Let G be a countable amenable group, let X be a G -SFT, and let $Y \subset X$ be any subsystem such that $h(Y) < h(X)$. Then*

$$\{h(Z) : Y \subset Z \subset X \text{ and } Z \text{ is an SFT}\}$$

is dense in $[h(Y), h(X)]$.

Proof. We prove the density directly. Suppose $(a, b) \subset [h(Y), h(X)]$ for positive reals $a < b$, and let $\varepsilon < (b - a)/2$. By Theorem 4.1, there exists an SFT $Z_0 \subset X$ such that $a < h(Z_0) < a + \varepsilon$. Note that these inequalities give $h(Y) < h(Z_0)$. Consider the subshift $Y \cup Z_0 \subset X$, which has entropy

$$h(Y \cup Z_0) = \max(h(Y), h(Z_0)) = h(Z_0) \in (a, a + \varepsilon).$$

Because X is an SFT and $Y \cup Z_0 \subset X$, by Theorem 2.27 there is an SFT Z such that $Y \cup Z_0 \subset Z \subset X$ and $h(Y \cup Z_0) \leq h(Z) < h(Y \cup Z_0) + \varepsilon$. Thus we have

$$a < h(Y \cup Z_0) \leq h(Z) < h(Y \cup Z_0) + \varepsilon < a + 2\varepsilon < b.$$

Since (a, b) was arbitrary, the proof is complete. \square

5. SOFIC SHIFTS

5.1. An extension result for sofic shifts. In order to address the case of sofic shifts, we seek to leverage our results on SFTs. In particular, given a sofic shift W , we would like an SFT X such that W is a factor of X and such that the maximal entropy drop across the factor map is very small. The following theorem guarantees the existence of such SFTs.

Theorem 5.1. *Let $W \subset \mathcal{A}_W^G$ be a sofic shift. For every $\varepsilon > 0$, there exists an SFT \tilde{X} and a one-block code $\tilde{\phi} : \tilde{X} \rightarrow W$ such that the maximal entropy gap of $\tilde{\phi}$ satisfies $\mathcal{H}(\tilde{\phi}) < \varepsilon$.*

Proof. Since W is sofic, there exists an SFT $X \subset \mathcal{A}_X^G$ and a factor map $\phi : X \rightarrow W$. Without loss of generality, we assume that

- i. ϕ is a one-block code, witnessed by the function $\Phi : \mathcal{A}_X \rightarrow \mathcal{A}_W$, and

ii. \mathcal{A}_X and \mathcal{A}_W are disjoint.

We abbreviate $\mathcal{A}_{XW} = \mathcal{A}_X \sqcup \mathcal{A}_W$. Let $\varepsilon > 0$, and select $\delta > 0$ such that

$$4\delta + \delta(1 + \delta) \log |\mathcal{A}_X| < \varepsilon/2.$$

Let $K \subset G$ be a large finite subset that specifies X as an SFT. The set K is fixed for the remainder of this proof, and thus we shall denote $\partial_{KK^{-1}}F$ by ∂F and $\text{int}_{KK^{-1}}F$ by $\text{int}F$ for any finite set $F \subset G$. By Theorem 3.6, there exists a finite set of finite shapes \mathcal{S} such that the following conditions are met.

- i. Each shape $S \in \mathcal{S}$ is (KK^{-1}, η) -invariant, where $\eta > 0$ is a constant such that $\eta|KK^{-1}| < \delta$. By Lemma 2.5, this implies $|\partial S| < \delta|S|$ for each $S \in \mathcal{S}$.
- ii. $KK^{-1} \subset S$ and $|S| > \delta^{-1}$ for each $S \in \mathcal{S}$.
- iii. There is a point $t_0 \in \Sigma_E(\mathcal{S})$ such that $h(\overline{\mathcal{O}}(t_0)) = 0$.

Recall by Proposition 3.3 that Σ_E is an SFT. By Theorem 2.27, there is an SFT T such that $\overline{\mathcal{O}}(t_0) \subset T \subset \Sigma_E$ and $h(T) < h(\overline{\mathcal{O}}(t_0)) + \delta$. Consequently, each point $t \in T$ is an encoding of an exact tiling \mathcal{T}_t of G over \mathcal{S} (possibly distinct from the original tiling \mathcal{T}_0), with tiling system entropy

$$h(\mathcal{T}_t) = h(\overline{\mathcal{O}}(t)) \leq h(T) < h(\overline{\mathcal{O}}(t_0)) + \delta = 0 + \delta.$$

Because X and T are SFTs, we have that $X \times T \subset (\mathcal{A}_X \times \Sigma)^G$ is also an SFT.

Let $t \in T$ be arbitrary, and recall that \mathcal{T}_t is a partition of G . Thus, for each $g \in G$, there is a unique tile $\tau \in \mathcal{T}_t$ such that $g \in \tau$. We define the notation $\mathcal{T}_t(g)$ by setting $\mathcal{T}_t(g) = \tau$. Next we define a map $\phi_t : X \rightarrow \mathcal{A}_{XW}^G$ by the following rule: for each $g \in G$ and $x \in X$,

$$\phi_t(x)_g = \begin{cases} x_g & \text{if } g \in \partial\mathcal{T}_t(g), \\ \Phi(x_g) & \text{if } g \in \text{int}\mathcal{T}_t(g). \end{cases}$$

This map is well-defined, as $\tau = \partial\tau \sqcup \text{int}\tau$. The map ϕ_t applies the one-block code ϕ to “most” of a point x , by relabelling the *interiors* of each tile $\tau \in \mathcal{T}_t$.

We now define a sliding block code $\varphi : X \times T \rightarrow (\mathcal{A}_{XW} \times \Sigma)^G$ by applying the map(s) ϕ_t fiber-wise: for each point $(x, t) \in X \times T$, let

$$\varphi(x, t) = (\phi_t(x), t).$$

It is straightforward to check that φ is indeed a sliding block code (Definition 2.8). For the theorem, the desired shift \tilde{X} is identified with the range of this map. Let

$$\tilde{X} = \varphi(X \times T) \subset (\mathcal{A}_{XW} \times \Sigma)^G.$$

See Figure 3 for an illustration of the construction. It remains to show that there is a one-block code $\tilde{\phi} : \tilde{X} \rightarrow W$, that the shift \tilde{X} is an SFT, and that $\mathcal{H}(\tilde{\phi}) < \varepsilon$.

First, let us show that \tilde{X} factors onto W . The factor map is induced by the function $\tilde{\Phi} : \mathcal{A}_{XW} \rightarrow \mathcal{A}_W$, which is an extension of Φ , defined by the following rule: $\tilde{\Phi}(\alpha) = \alpha$ if $\alpha \in \mathcal{A}_W$, and $\tilde{\Phi}(\alpha) = \Phi(\alpha)$ if $\alpha \in \mathcal{A}_X$. Let $\tilde{\phi} : \tilde{X} \rightarrow \mathcal{A}_W^G$ be given by

$$\tilde{\phi}(\tilde{x}, t)_g = \tilde{\Phi}(\tilde{x}_g), \quad \forall g \in G \text{ and } \forall(\tilde{x}, t) \in \tilde{X}.$$

It is clear that $\tilde{\phi}$ is a one-block code. Let us now show that $\tilde{\phi}(\tilde{X}) = W$. Let $x \in X$ and $t \in T$, in which case $(\phi_t(x), t) \in \tilde{X}$ is an arbitrary point. The effect of applying the map ϕ_t to x is to apply the one-block code ϕ to “most” of x . The map $\tilde{\phi}$ then “completes” the relabelling, via the extended function $\tilde{\Phi}$. In fact, we have that $\tilde{\phi}(\phi_t(x), t) = \phi(x) \in W$, hence $\tilde{\phi}(\tilde{X}) \subset W$. For the reverse inclusion, let $w \in W$. Since $\phi : X \rightarrow W$ is onto, there exists a point $x \in X$ such that $\phi(x) = w$.

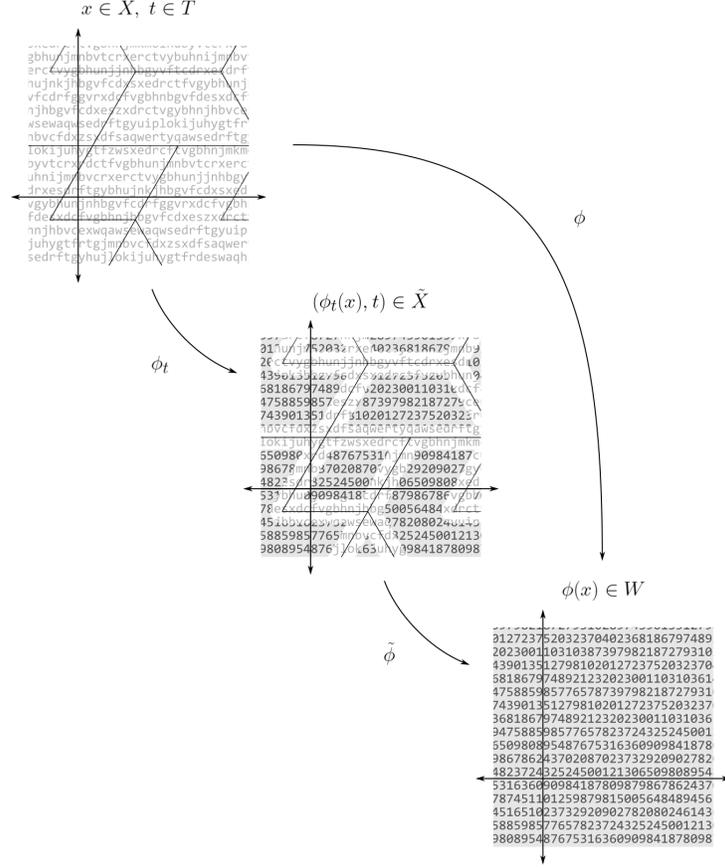


FIGURE 3. A hypothetical point $x \in X$ with a tiling $t \in T$ overlaid; the partially-transformed point $\phi_t(x)$ is pictured, which is labelled with symbols from both X and W ; finally, the wholly-transformed image point $\phi(x) \in W$ is reached.

Choose $t \in T$ arbitrarily; then $(\phi_t(x), t) \in \tilde{X}$ and $\tilde{\phi}(\phi_t(x), t) = \phi(x) = w$. We conclude that $\tilde{\phi} : \tilde{X} \rightarrow W$ is a genuine factor map (and a one-block code).

Let us now show that \tilde{X} is an SFT. We repeat that the shift \tilde{X} can be written in the following instructive form:

$$\tilde{X} = \{(\phi_t(x), t) : x \in X \text{ and } t \in T\} \subset (\mathcal{A}_{XW} \times \Sigma)^G.$$

In order to show that \tilde{X} is an SFT, we will construct an SFT $\tilde{X}_1 \subset (\mathcal{A}_{XW} \times \Sigma)^G$ and then prove that $\tilde{X} = \tilde{X}_1$. Recall that $K \subset G$ specifies X as an SFT. Let $K_T \subset G$ be a finite subset such that $\mathcal{P}(K_T, T)$ specifies T . We define \tilde{X}_1 to be the set of points $(\tilde{x}, t) \in (\mathcal{A}_{XW} \times \Sigma)^G$ that satisfy the following local rules.

- (R1) Any pattern of shape K_T that occurs in t must belong to $\mathcal{P}(K_T, T)$, and any pattern of shape K that occurs in \tilde{x} and belongs to \mathcal{A}_X^K must also belong

to $\mathcal{P}(K, X)$ (recall $\mathcal{P}(K, \tilde{X}) \subset \mathcal{A}_{\tilde{X}W}^K = (\mathcal{A}_X \sqcup \mathcal{A}_W)^K$ in general). Note by Definition 2.14 that this condition is shift-invariant.

- (R2) For any shape $S \in \mathcal{S}$ and any $c \in G$, if t satisfies $(\sigma^c t)_s = (S, s)$ for each $s \in S$, then $\exists b \in \mathcal{P}(S, X)$ such that $(\sigma^c \tilde{x})_s = b_s \in \mathcal{A}_X$ for all $s \in \partial S$ and $(\sigma^c \tilde{x})_s = \Phi(b_s) \in \mathcal{A}_W$ for all $s \in \text{int } S$.

As these are local rules, they define an SFT; call it $\tilde{X}_1 \subset (\mathcal{A}_{XW} \times \Sigma)^G$. Moreover, it is easily checked that any point $(\phi_t(x), t) \in \tilde{X}$ satisfies these rules everywhere (by construction of \tilde{X}), and so we have $\tilde{X} \subset \tilde{X}_1$.

For the reverse inclusion, consider a point $(\tilde{x}, t) \in \tilde{X}_1$. From (R1) it follows that $t \in T$, as T is an SFT specified by K_T . Therefore, t encodes an exact tiling \mathcal{T}_t of G over \mathcal{S} with $h(\mathcal{T}_t) < \delta$. Let $(\tau_n)_n$ enumerate the tiles of \mathcal{T}_t , and for each n let $\tau_n = S_n c_n$ for some $S_n \in \mathcal{S}$ and $c_n \in G$. Recall $\{\tau_n : n \in \mathbb{N}\}$ is a partition of G .

Let $n \in \mathbb{N}$, and consider $c = c_n$ and $S = S_n$. Observe that, because t encodes the tiling \mathcal{T}_t , we have $(\sigma^c t)_s = (S, s)$ for each $s \in S$. Then by (R2), there exists a block $b = b_n \in \mathcal{P}(S, X)$ such that $(\sigma^c \tilde{x})_s = b_s$ for all $s \in \partial S$ and $(\sigma^c \tilde{x})_s = \Phi(b_s)$ for all $s \in \text{int } S$.

Define a point $x \in \mathcal{A}_X^G$ by setting $x(\tau_n) = b_n$ for each $n \in \mathbb{N}$. We claim that x is an allowed point of X and that $\phi_t(x) = \tilde{x}$. Toward this, let $g \in G$ be arbitrary, and consider the translate Kg (recall that K specifies X as an SFT).

If Kg intersects the *interior* of any tile $\tau_n = S_n c_n$, then $Kg \subset \tau_n$ by Lemma 2.6. In this case, the pattern $(\sigma^g x)(K)$ is a subpattern of b_n , and must therefore be allowed in X as $b_n \in \mathcal{P}(S_n, X)$. The alternative is that Kg is disjoint from the interior of every tile, in which case $Kg \subset \bigcup_n \partial \tau_n$. By (R2), we also have $\tilde{x}_g \in \mathcal{A}_X$ for every $g \in \bigcup_n \text{int } \tau_n$. In this case we have $(\sigma^g x)(K) = (\sigma^g \tilde{x})(K)$, which is again allowed in X by (R1).

In either case we have that $(\sigma^g x)(K)$ is allowed in X for any $g \in G$, and hence $x \in X$. Then by the definition of ϕ_t , we see that $\phi_t(x) = \tilde{x}$. Thus, we have found a point $(x, t) \in X \times T$ such that $\varphi(x, t) = (\phi_t(x), t) = (\tilde{x}, t)$, and hence $(\tilde{x}, t) \in \tilde{X}$. We conclude that $\tilde{X} = \tilde{X}_1$, and therefore \tilde{X} is an SFT.

Finally, let us show that $\mathcal{H}(\tilde{\phi}) < \varepsilon$. Towards this end, let $\tilde{X}' \subset \tilde{X}$ be any subsystem of \tilde{X} , and let $W' = \tilde{\phi}(\tilde{X}') \subset W$. We will show that $h(\tilde{X}') - h(W') < \varepsilon/2$.

Let $F \subset G$ be a finite subset such that the following conditions are met.

- (F1) F is (UU^{-1}, ϑ) -invariant, where $U = \bigcup \mathcal{S}$ and $\vartheta > 0$ is a constant such that $\vartheta|U||UU^{-1}| < \delta$ (recall δ was selected at the beginning of this proof). Note this implies that F may be well approximated by tiles from any exact tiling of G over \mathcal{S} , in the sense of Lemma 3.8.
- (F2) F is large enough to δ -approximate (Definition 2.24) the entropy of the shifts X', W' and T (recall that $h(T) < \delta$, in which case $h(F, T) < 2\delta$).

Such a set exists by Proposition 2.25. This set is fixed for the remainder of this proof. Recall that $\tilde{\phi}$ is a one-block code, and therefore there is a well defined map $\tilde{\Phi}_F : \mathcal{P}(F, \tilde{X}') \rightarrow \mathcal{P}(F, W')$ which takes a pattern $p \in \mathcal{P}(F, \tilde{X}')$ and applies the one-block code to p (at each element of F).

Recall also that a pattern $p \in \mathcal{P}(F, \tilde{X}')$ is of the form $p = (\phi_t(x), t)(F)$ for some points $x \in X$ and $t \in T$. The point t encodes an exact tiling \mathcal{T}_t of G over \mathcal{S} . For each tile $\tau \in \mathcal{T}_t$, the definition of ϕ_t implies that

$$(5) \quad \phi_t(x)(\text{int } \tau) \in \mathcal{A}_W^*, \quad \text{and} \quad \phi_t(x)(\partial \tau) \in \mathcal{A}_X^*.$$

Let $q = \tilde{\Phi}_F(p) \in \mathcal{P}(F, W')$. Recall that every element $g \in F$ belongs to a unique tile $\tau = \mathcal{T}_t(g) \in \mathcal{T}_t^\times(F)$, where $\mathcal{T}_t^\times(F) \subset \mathcal{T}_t$ is the outer approximation of F by the tiling \mathcal{T}_t (Definition 3.7). By (5), we have that

$$q_g = \tilde{\Phi}_F(p)_g = \tilde{\Phi}(p_g^{\tilde{X}}) = \begin{cases} \Phi(p_g^{\tilde{X}}) & \text{if } g \in \partial\mathcal{T}_t(g) \\ p_g^{\tilde{X}} & \text{if } g \in \text{int } \mathcal{T}_t(g) \end{cases}$$

In particular, we have $q_g = p_g^{\tilde{X}}$ whenever g belongs to the set

$$F \cap \left(\bigcup_{\tau} \text{int } \tau \right),$$

where the union is taken over all $\tau \in \mathcal{T}_t^\times(F)$.

In light of these observations, we are ready to estimate $|P(F, \tilde{X}')|$ in terms of $|P(F, W')|$. We first use $\tilde{\Phi}_F$ to split over $\mathcal{P}(F, W')$, and then we split again over all possible T -layers. Indeed, we have

$$(6) \quad |\mathcal{P}(F, \tilde{X}')| = \sum_q \sum_t |\{p \in \tilde{\Phi}_F^{-1}(q) : p^T = t\}|$$

where the sums are taken over all patterns $q \in \mathcal{P}(F, W')$ and $t \in \mathcal{P}(F, T)$. Choose and fix patterns q and t . If $p \in \mathcal{P}(F, \tilde{X}')$ is a pattern such that $\tilde{\Phi}_F(p) = q$ and $p^T = t$, then the observations above imply that p is uniquely determined by

$$p^{\tilde{X}} \left(F \cap \left(\bigcup_{\tau} \partial\tau \right) \right) \in \mathcal{A}_X^*$$

where the union is taken over all tiles $\tau \in \mathcal{T}_t^\times(F)$. Moreover, our choice of \mathcal{S} and the property (F1) of F together yield that

$$\left| \bigcup_{\tau} \partial\tau \right| < \delta \left| \bigcup_{\tau} \tau \right| = \delta |F_t^\times| < \delta(1 + \delta)|F|.$$

Therefore, there are at most $|\mathcal{A}_X|^{\delta(1+\delta)|F|}$ patterns p such that $\tilde{\Phi}_F(p) = q$ and $p^T = t$. From this and Equation (6), we have

$$|\mathcal{P}(F, \tilde{X}')| \leq |\mathcal{P}(F, W')| \cdot |\mathcal{P}(F, T)| \cdot |\mathcal{A}_X|^{\delta(1+\delta)|F|}.$$

By taking logs and dividing through by $|F|$, we obtain the following:

$$\begin{aligned} h(\tilde{X}') &< h(F, \tilde{X}') + \delta \\ &\leq (h(F, W') + h(F, T) + \delta(1 + \delta) \log |\mathcal{A}_X|) + \delta \\ &< (h(W') + \delta) + (2\delta) + \delta(1 + \delta) \log |\mathcal{A}_X| + \delta \\ &= h(W') + 4\delta + \delta(1 + \delta) \log |\mathcal{A}_X| \\ &< h(W') + \varepsilon/2, \end{aligned}$$

where we have used property (F2) of the set F , the above inequality, property (F2) again, and our choice of δ respectively. Since $\tilde{X}' \subset \tilde{X}$ was arbitrary, we have that

$$\mathcal{H}(\tilde{\phi}) = \sup_{\tilde{X}' \subset \tilde{X}} (h(\tilde{X}') - h(\tilde{\phi}(\tilde{X}'))) \leq \varepsilon/2 < \varepsilon,$$

which completes the proof. \square

5.2. Subsystem entropies for sofic shifts. Here we present our main result concerning subsystem entropies for sofic shifts. The proof follows easily by combining our extension result (Theorem 5.1) with our result for SFTs (Theorem 4.2).

Theorem 5.2. *Let G be a countable amenable group, let W be a sofic G -shift and let $V \subset W$ be any subsystem such that $h(V) < h(W)$. Then*

$$\{h(U) : V \subset U \subset W \text{ and } U \text{ is sofic}\}$$

is dense in $[h(V), h(W)]$.

Proof. We prove the density directly. Let $(a, b) \subset [h(V), h(W)]$ for some real numbers $a < b$. Let $\varepsilon < (b - a)/2 < h(W) - h(V)$. By Theorem 5.1, there exists an SFT X and a factor map $\phi : X \rightarrow W$ such that $\mathcal{H}(\phi) < \varepsilon$.

Consider the preimage $Y = \phi^{-1}(V) \subset X$, which is a subshift. Note that $\phi(Y) = V$ because ϕ is surjective, hence $\phi|_Y : Y \rightarrow V$ is a factor map. We then have that

$$h(Y) \leq h(V) + \mathcal{H}(\phi) < h(V) + \varepsilon < h(W) \leq h(X).$$

Note also that $b \leq h(W) \leq h(X)$ and that $a \geq h(V) > h(Y) - \varepsilon$, which together yield that $(a + \varepsilon, b) \subset [h(Y), h(X)]$. By Theorem 4.2, there exists an SFT Z such that $Y \subset Z \subset X$ and $h(Z) \in (a + \varepsilon, a + 2\varepsilon) \subset (a, b)$. It follows that $U = \phi(Z)$ is a sofic shift for which $V \subset U \subset W$ and $h(U) \leq h(Z) < h(U) + \varepsilon$. Then we have

$$a < h(Z) - \varepsilon < h(U) \leq h(Z) < a + 2\varepsilon < b.$$

Thus $h(U) \in (a, b)$, which completes the proof. \square

If one selects $V = \emptyset$ for the above theorem, then one recovers the statement that the entropies of the sofic subsystems of W are dense in $[0, h(W)]$. Next, we present our result concerning the entropies of arbitrary subsystems of sofic shifts.

Corollary 5.3. *Let W be a sofic shift. For every nonnegative real $r \leq h(W)$, there exists a subsystem $R \subset W$ for which $h(R) = r$.*

Proof. If $h(W) = 0$, then $r = 0$, in which case one may simply select $R = W$. If $h(W) > 0$, then let $W_0 = W$ and let $(\varepsilon_n)_n$ be a sequence of positive real numbers converging to zero. We have that W_0 is sofic and $r \leq h(W_0)$, and without loss of generality we assume that $h(W_0) < r + \varepsilon_0$.

Inductively construct a descending sequence of sofic shifts as follows. If $W_n \subset W$ is a sofic shift such that $r \leq h(W_n) < r + \varepsilon_n$, then by Theorem 5.2 there exists a sofic shift $W_{n+1} \subset W_n$ for which $r \leq h(W_{n+1}) < r + \varepsilon_{n+1}$.

Then $R = \bigcap_n W_n \subset W$ is a subshift such that $h(R) = \lim_n h(W_n) = r$ by Proposition 2.26. \square

6. A COUNTER-EXAMPLE

Theorem 5.2 implies that the entropies of the *sofic* subsystems of a sofic shift space W are dense in $[0, h(W)]$. One may wonder if this can be somehow “sharpened”; that is, one may wonder whether

$$\{h(W') : W' \subset W \text{ and } W' \text{ is an SFT}\}$$

is dense in $[0, h(W)]$. However, this statement is nowhere close to true in general, as we illustrate in this section by counterexample. This example is an adaptation of a construction of Boyle, Pavlov, and Schraudner [5].

Proposition 6.1. *There exists a sofic \mathbb{Z}^2 -shift with positive entropy whose only SFT subsystem is a singleton.*

Proof. We first construct a certain point in $\{0, 1\}^{\mathbb{Z}}$ as the limit of a sequence of finite words, then consider the subshift it generates. Let $\delta = 0.1$ and let $(T_n)_n$ be the sequence of natural numbers given by

$$T_n = 2n \cdot 2^n \cdot \delta^{-1} + 1$$

for each n . Let $w^1 = 010 \in \{0, 1\}^3$, and for each n define the word

$$(7) \quad w^{n+1} = w^n w^n \cdots w^n w^n 0^n 10^n,$$

where the w^n term is repeated exactly T_n times. The limit word $w^\infty \in \{0, 1\}^{\mathbb{N}_0}$ is an infinite one-sided sequence. Define a two-sided sequence $x^* \in \{0, 1\}^{\mathbb{Z}}$ by $x_i^* = w_{|i|}^\infty$ for each $i \in \mathbb{Z}$. Let $X = \overline{\mathcal{O}}(x^*) \subset \{0, 1\}^{\mathbb{Z}}$ be the subshift generated by x^* . We claim that X exhibits the following three properties.

- (P1) X is *effective*, meaning that there exists a finite algorithm which enumerates a set of words $\mathcal{F} \subset \{0, 1\}^*$ such that $X = \mathcal{R}(\{0, 1\}^{\mathbb{Z}}, \mathcal{F})$.
(P2) There exists a point $x \in X$ such that

$$\limsup_{n \rightarrow \infty} \frac{|\{k \in [-n, n] : x_k = 1\}|}{2n + 1} > 0.1.$$

- (P3) For each $x \in X$, either $x = 0^{\mathbb{Z}}$ or x contains the word $0^n 10^n$ for every n .

For (P1), let N be arbitrary. Note that because $X = \overline{\mathcal{O}}(x^*)$, any word of length N occurring in any point $x \in X$ is also a word occurring in x^* . By the recursive definition (7) and the fact that the sequence $\{T_n\}_{n=1}^\infty$ is recursive, there is an algorithm which, upon input N , prints all the words of length N that do *not* appear as subwords of x^* . The shift X is therefore effective.

For (P2), we argue that x^* satisfies the condition. For each n , let L_n be the length of the word w^n . Note that by the recurrence (7), we have

$$L_{n+1} = T_n L_n + 2n + 1 \quad \forall n.$$

For each n , let f_n be the *frequency* of 1s in w^n , given by

$$f_n = \frac{|\{i : w_i^n = 1\}|}{L_n}.$$

Observe that $f_n \leq 1$ for each n and $f_1 = \frac{1}{3}$. It follows from the recurrence (7) that

$$f_{n+1} = \frac{f_n T_n L_n + 1}{T_n L_n + 2n + 1}$$

for each n . This implies that

$$\begin{aligned} f_n - f_{n+1} &= \frac{f_n T_n L_n + f_n(2n + 1) - f_n T_n L_n - 1}{T_n L_n + 2n + 1} \\ &\leq \frac{1(2n + 1) - 1}{T_n} \end{aligned}$$

in which case $f_n - f_{n+1} \leq \frac{2n}{T_n} < \frac{\delta}{2^n}$ for each n . Hence, we have that

$$\begin{aligned} f_1 - f_n &= (f_1 - f_2) + (f_2 - f_3) + \cdots + (f_{n-1} - f_n) \\ &< \frac{\delta}{2} + \frac{\delta}{4} + \cdots + \frac{\delta}{2^{n-1}} \end{aligned}$$

$$< \delta$$

in which case $\frac{1}{3} - \delta < f_n$ for each n . By the recurrence (7), we therefore have that

$$\limsup_{n \rightarrow \infty} \frac{|\{k \in [-n, n] : x_k^* = 1\}|}{2n+1} \geq \frac{1}{3} - \delta > 0.1.$$

and the subsequence along $(L_n)_n$ is a witness.

For (P3), let n be given. First, observe that the infinite sequence w^∞ is the concatenation of a sequence of blocks, where each block is either the word w^n or $0^m 10^m$ for some $m \geq n$. Moreover, each w^n begins with 0 and ends with 0^n . This implies that 1s in distinct blocks are separated by at least $n+1$ appearances of the symbol 0. Therefore, if for any $k \leq n$ we have that $10^k 1$ appears anywhere in w^∞ , then it must appear as a subword of a single block (rather than overlapping two distinct blocks), and that block must be w^n .

Next, let $x \in X$ be arbitrary. If the symbol 1 appears in x at most one time, then (P3) trivially holds. Otherwise, assume that $10^k 1$ appears somewhere in x for some $k \geq 1$. Without loss of generality, suppose $x_0 = x_{k+1} = 1$ and $x_i = 0$ for $i \in [1, k]$. Now consider the subword $\omega = x([-L_n, L_n])$ for any n such that $k < L_n$. Because $X = \overline{\mathcal{O}(x^*)}$, the word ω must be a subword of x^* . Then, either ω is a subword of $x^*([-2L_n, 2L_n])$, or ω is a subword of w^∞ or a mirror reflection of one. In the first case, the definitions of x^* and w^∞ imply that ω contains the word w^n or its mirror. In the latter two cases, the observation of the previous paragraph implies that ω must contain w^n or its mirror. In any case, $0^n 10^n$ is a subword of x . As n can be made arbitrarily large, this proves (P3).

We now use the shift X to construct the \mathbb{Z}^2 -shift which is desired for the theorem. For each point $x \in X$, let $x^{\mathbb{Z}} \in \{0, 1\}^{\mathbb{Z}^2}$ denote the \mathbb{Z}^2 -labelling given by

$$(x^{\mathbb{Z}})_{(i,j)} = x_i$$

for each $(i, j) \in \mathbb{Z}^2$. That is, $x^{\mathbb{Z}}$ is a \mathbb{Z}^2 -labelling such that the symbols along each column are constant, and each row is equal to x itself. We shall also denote

$$X^{\mathbb{Z}} = \{x^{\mathbb{Z}} : x \in X\} \subset \{0, 1\}^{\mathbb{Z}^2}.$$

It is a theorem of Aubrun and Sablik [1] that if X is effective, then $X^{\mathbb{Z}}$ is sofic.

Next, consider the alphabet $\{0, 1, 1'\}$, where we have artificially created two independent 1 symbols. Let $\pi : \{0, 1, 1'\}^{\mathbb{Z}^2} \rightarrow \{0, 1\}^{\mathbb{Z}^2}$ be the one-block code which collapses 1 and $1'$. Let $Y = \pi^{-1}(X^{\mathbb{Z}}) \subset \{0, 1, 1'\}^{\mathbb{Z}^2}$. The shift Y is a copy of the shift $X^{\mathbb{Z}}$, in which the 1 symbols of every point have been replaced either by 1 or $1'$ in every possible combination.

We claim that the shift Y is the desired subshift for the theorem. Specifically, we claim that Y is sofic, that Y has positive entropy, and that the only nonempty SFT subsystem of Y is the singleton $\{0^{\mathbb{Z}^2}\}$.

To prove that Y is sofic, we construct an SFT S' and a factor map $\phi' : S' \rightarrow Y$ to witness the soficity of Y . Since $X^{\mathbb{Z}}$ is sofic, there is an SFT $S \subset \mathcal{A}^{\mathbb{Z}^2}$ and a factor map $\phi : S \rightarrow X^{\mathbb{Z}}$. Without loss of generality, assume that ϕ is a one-block code induced by the function $\Phi : \mathcal{A} \rightarrow \{0, 1\}$.

Define a new finite alphabet $\mathcal{A} \times \{1, 1'\}$ and a one-block code ϕ' induced by the function $\Phi' : \mathcal{A} \times \{1, 1'\} \rightarrow \{0, 1, 1'\}$ which is given by

$$\Phi'(a, b) = \begin{cases} 0 & \text{if } \Phi(a) = 0 \\ b & \text{if } \Phi(a) = 1 \end{cases}$$

for each $(a, b) \in \mathcal{A} \times \{1, 1'\}$.

Let $S' = S \times \{1, 1'\}^{\mathbb{Z}^2}$, which we regard as a subshift of $(\mathcal{A} \times \{1, 1'\})^{\mathbb{Z}^2}$. Note that S' is an SFT, because both S and $\{1, 1'\}^{\mathbb{Z}^2}$ (the full \mathbb{Z}^2 -shift on two symbols) are SFTs. A point $s' \in S'$ is of the form $s' = (s, \iota)$, where s is a point of S and $\iota \in \{1, 1'\}^{\mathbb{Z}^2}$ is an arbitrary 2-coloring of \mathbb{Z}^2 . The reader may easily check that $(\pi \circ \phi')(s, \iota) = \phi(s) \in X^{\mathbb{Z}}$, from which it follows that $\phi'(S') = \pi^{-1}(X^{\mathbb{Z}}) = Y$. Then $\phi' : S' \rightarrow Y$ is a factor map. Since S' is an SFT, we conclude that Y is sofic.

Next, we will show that $h(Y) > 0$. From property (P2), the point $x^* \in X$ exhibits 1s in more than 10% of the positions in each of infinitely many symmetric intervals, say of the form $[-\ell_n, \ell_n]$ for an increasing sequence of natural numbers $(\ell_n)_n$. Therefore, the point $(x^*)^{\mathbb{Z}}$ exhibits 1s in more than 10% of the positions in each square $F_n = [-\ell_n, \ell_n]^2$. Each 1 in the pattern $(x^*)^{\mathbb{Z}}(F_n)$ may be replaced by 1 or 1' independently to yield an allowed pattern of Y , which implies that

$$|\mathcal{P}(F_n, Y)| \geq 2^{0.1|F_n|} \quad \forall n.$$

As $(F_n)_n$ is a Følner sequence for \mathbb{Z}^2 , we then have $h(Y) \geq 0.1 \log 2 > 0$.

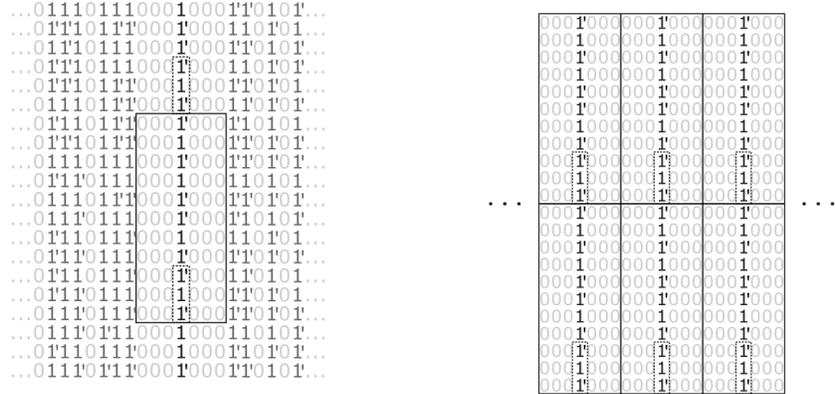
It remains to show that the only nonempty SFT subsystem of Y is the singleton $\{0^{\mathbb{Z}^2}\}$. Suppose to the contrary that $Z \subset Y$ is an SFT subsystem of Y which contains a nonzero point. Since Z is an SFT, we may find a constant $k \in \mathbb{N}$ such that the allowed patterns of Z are specified by the shape $K = [0, k]^2 \subset \mathbb{Z}^2$.

Let $z \in Z$ be a point different from $0^{\mathbb{Z}^2}$ and note $\pi(z) = x^{\mathbb{Z}} \in X^{\mathbb{Z}}$ for some $x \in X$ with $x \neq 0^{\mathbb{Z}}$. By property (P1), the string $0^n 10^n$ appears in x for every n . Let $n > k$ be fixed. Suppose without loss of generality that $0^n 10^n$ appears centered at the origin of x (with $x_0 = 1$ and $x_i = 0$ for $0 < |i| \leq n$). Thus we have $z_{(0,0)} = 1$ or 1'. In fact, by the definition of Y , we have $z_{(0,j)} \in \{1, 1'\}$ for every $j \in \mathbb{Z}$.

Consider the $i = 0$ column of the point z . Starting with each index $\ell \in \mathbb{Z}$ and looking up, there is a corresponding vertically oriented word $\omega^\ell \in \{1, 1'\}^n$ given by $\omega_j^\ell = z_{(0, \ell+j)}$ for each $j \in [0, n)$. By the pigeonhole principle, there must exist a word $\varsigma \in \{1, 1'\}^n$ such that $\varsigma = \omega^\ell$ for infinitely many choices of ℓ . That is, for infinitely many choices of ℓ , we have $z_{(0, \ell+j)} = \varsigma_j$ for each $j \in [0, n)$.

Let $\ell_1 < \ell_2$ be two such indices where a repetition occurs, with $\ell_2 - \ell_1 > n$. That is, we have $z_{(0, \ell_1+j)} = z_{(0, \ell_2+j)} = \varsigma_j$ for every $j \in [0, n)$. Now consider the rectangle $r = z([-n, n] \times [\ell_1, \ell_2])$. Tile \mathbb{Z}^2 with infinitely many translated copies of r to obtain a new point $z' \in \{0, 1, 1'\}^{\mathbb{Z}^2}$. Figure 4 illustrates the construction.

Every pattern of shape $K = [0, k]^2$ which occurs in z' is a pattern which occurs in z (including the pattern of all zeroes), hence they are all allowed in Z . Because Z is an SFT specified by K , it then follows that $z' \in Z$. Because $Z \subset Y = \pi^{-1}(X^{\mathbb{Z}})$, there must exist a point $x' \in X$ such that $\pi(z') = (x')^{\mathbb{Z}}$. We obtain a contradiction, as the point x' cannot satisfy the property (P3) of X . For instance, the word $0^{3n} 10^{3n}$ cannot appear in x' (as each row of z' is periodic in the horizontal direction with period $2n + 1 < 3n$). This demonstrates that if Z is an SFT, then it contains no nonzero point. Therefore, the only nonempty SFT subsystem of Y is $\{0^{\mathbb{Z}^2}\}$. \square



(A) A hypothetical point z is illustrated around $[-n, n] \times [\ell_1, \ell_2]$. The repeated vertical word ζ is indicated by the dotted box, and the rectangle r by the solid box.

(B) The rectangle r is used to tile \mathbb{Z}^2 and thereby construct z' . Every $k \times k$ block which occurs in z' also occurs in z .

FIGURE 4. An illustration of the construction of the contradictory point z' , in a hypothetical case where $n = 3$ and $k = 2$.

REFERENCES

- [1] Nathalie Aubrun and Mathieu Sablik. “Simulation of effective subshifts by two-dimensional subshifts of finite type”. In: *Acta Appl. Math.* 126 (2013), pp. 35–63. ISSN: 0167-8019. DOI: 10.1007/s10440-013-9808-5. URL: <https://doi-org.du.idm.oclc.org/10.1007/s10440-013-9808-5>.
- [2] Sebastián Barbieri. “On the entropies of subshifts of finite type on countable amenable groups”. In: *Groups Geom. Dyn.* 15.2 (2021), pp. 607–638. ISSN: 1661-7207. DOI: 10.4171/GGD/608. URL: <https://doi-org.du.idm.oclc.org/10.4171/GGD/608>.
- [3] Sebastián Barbieri and Mathieu Sablik. “A generalization of the simulation theorem for semidirect products”. In: *Ergodic Theory Dynam. Systems* 39.12 (2019), pp. 3185–3206. ISSN: 0143-3857. DOI: 10.1017/etds.2018.21. URL: <https://doi-org.du.idm.oclc.org/10.1017/etds.2018.21>.
- [4] Mike Boyle. “Lower entropy factors of sofic systems”. In: *Ergodic Theory Dynam. Systems* 3.4 (1983), pp. 541–557. ISSN: 0143-3857. DOI: 10.1017/S0143385700002133. URL: <https://doi.org/10.1017/S0143385700002133>.
- [5] Mike Boyle, Ronnie Pavlov, and Michael Schraudner. “Multidimensional sofic shifts without separation and their factors”. In: *Trans. Amer. Math. Soc.* 362.9 (2010), pp. 4617–4653. ISSN: 0002-9947. DOI: 10.1090/s0002-9947-10-05003-8. URL: <https://doi-org.du.idm.oclc.org/10.1090/s0002-9947-10-05003-8>.
- [6] Angela Desai. “Subsystem entropy for \mathbb{Z}^d sofic shifts”. In: *Indag. Math. (N.S.)* 17.3 (2006), pp. 353–359. ISSN: 0019-3577. DOI: 10.1016/S0019-3577(06)

- 80037-6. URL: [https://doi-org.du.idm.oclc.org/10.1016/S0019-3577\(06\)80037-6](https://doi-org.du.idm.oclc.org/10.1016/S0019-3577(06)80037-6).
- [7] Tomasz Downarowicz, Dawid Huczek, and Guohua Zhang. “Tilings of amenable groups”. In: *J. Reine Angew. Math.* 747 (2019), pp. 277–298. ISSN: 0075-4102. DOI: 10.1515/crelle-2016-0025. URL: <https://doi-org.du.idm.oclc.org/10.1515/crelle-2016-0025>.
 - [8] Joshua Frisch and Omer Tamuz. “Symbolic dynamics on amenable groups: the entropy of generic shifts”. In: *Ergodic Theory Dynam. Systems* 37.4 (2017), pp. 1187–1210. ISSN: 0143-3857. DOI: 10.1017/etds.2015.84. URL: <https://doi-org.du.idm.oclc.org/10.1017/etds.2015.84>.
 - [9] Michael Hochman and Tom Meyerovitch. “A characterization of the entropies of multidimensional shifts of finite type”. In: *Ann. of Math. (2)* 171.3 (2010), pp. 2011–2038. ISSN: 0003-486X. DOI: 10.4007/annals.2010.171.2011. URL: <https://doi-org.du.idm.oclc.org/10.4007/annals.2010.171.2011>.
 - [10] Dawid Huczek and Sebastian Kopacz. *Factoring strongly irreducible group shift actions onto full shifts of lower entropy*. 2021. arXiv: 2106.10687.
 - [11] David Kerr and Hanfeng Li. *Ergodic theory*. Springer Monographs in Mathematics. Independence and dichotomies. Springer, Cham, 2016, pp. xxxiv+431. ISBN: 978-3-319-49845-4. DOI: 10.1007/978-3-319-49847-8. URL: <https://doi-org.du.idm.oclc.org/10.1007/978-3-319-49847-8>.
 - [12] Wolfgang Krieger. “On the subsystems of topological Markov chains”. In: *Ergodic Theory Dynam. Systems* 2.2 (1982), 195–202 (1983). ISSN: 0143-3857. DOI: 10.1017/S0143385700001516. URL: <https://doi-org.du.idm.oclc.org/10.1017/S0143385700001516>.
 - [13] Douglas Lind. “The entropies of topological Markov shifts and a related class of algebraic integers”. In: *Ergodic Theory Dynam. Systems* 4.2 (1984), pp. 283–300. ISSN: 0143-3857. DOI: 10.1017/S0143385700002443. URL: <https://doi-org.du.idm.oclc.org/10.1017/S0143385700002443>.
 - [14] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge, 1995, pp. xvi+495. ISBN: 0-521-55124-2. DOI: 10.1017/CB09780511626302. URL: <https://doi-org.du.idm.oclc.org/10.1017/CB09780511626302>.
 - [15] Donald Ornstein and Benjamin Weiss. “Entropy and isomorphism theorems for actions of amenable groups”. In: *J. Analyse Math.* 48 (1987), pp. 1–141. ISSN: 0021-7670. DOI: 10.1007/BF02790325. URL: <https://doi-org.du.idm.oclc.org/10.1007/BF02790325>.
 - [16] Anthony Quas and Paul Trow. “Subshifts of multi-dimensional shifts of finite type”. In: *Ergodic Theory Dynam. Systems* 20.3 (2000), pp. 859–874. ISSN: 0143-3857. DOI: 10.1017/S0143385700000468. URL: <https://doi-org.du.idm.oclc.org/10.1017/S0143385700000468>.

ROBERT BLAND, DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF NORTH CAROLINA AT CHARLOTTE, 9201 UNIVERSITY CITY BLVD., CHARLOTTE, NC 28223

Email address: rbland5@uncc.edu

KEVIN MCGOFF, DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF NORTH CAROLINA AT CHARLOTTE, 9201 UNIVERSITY CITY BLVD., CHARLOTTE, NC 28223

Email address: kmcgoff1@uncc.edu

URL: <https://clas-math.uncc.edu/kevin-mcgoff/>

RONNIE PAVLOV, DEPARTMENT OF MATHEMATICS, UNIVERSITY OF DENVER, 2390 S. YORK ST., DENVER, CO 80210

Email address: rpavlov@du.edu

URL: <https://cs.du.edu/~rpavlov/>