

SYMBOLIC DYNAMICS ON AMENABLE GROUPS: THE ENTROPY OF GENERIC SHIFTS

JOSHUA FRISCH AND OMER TAMUZ

ABSTRACT. Let G be a finitely generated amenable group. We study the space of shifts on G over a given finite alphabet A . We show that the zero entropy shifts are generic in this space, and that more generally the shifts of entropy c are generic in the space of shifts with entropy at least c . The same is shown to hold for the space of transitive shifts and for the space of weakly mixing shifts.

As applications of this result, we show that for every entropy value $c \in [0, \log |A|]$ there is a weakly mixing subshift of A^G with entropy c . We also show that the set of strongly irreducible shifts does not form a G_δ in the space of shifts, and that all non-trivial, strongly irreducible shifts are non-isolated points in this space.

CONTENTS

1. Introduction	2
1.1. Definitions and results	3
1.2. Organization and notation	8
2. Quasi-tiling shifts	9
3. Low entropy approximation	11
3.1. Overview	11
3.2. Deletion and local completion maps	12
3.3. Global completion maps	14
4. Fixed entropy approximation	17
4.1. Overview	17
4.2. Q and Q_p	18
4.3. Enumerating the completion maps	19
4.4. The shifts X^j	19
5. The topology of $\mathfrak{S}_{\geq c}$	23
5.1. Compatible quasi-tilings	24
5.2. Constructing the shifts X^j	24
5.3. Constructing the shift Y	25

Date: August 20, 2015.

J. Frisch was supported by MIT's Undergraduate Research Opportunities Program. This research was partially conducted at Microsoft Research, New England.

5.4. Realizing entropies	25
Appendix A. Combinatorial lemmata for bounding entropy	26
Appendix B. Semi-continuity of the entropy function	27
B.1. Bounding entropy using quasi-tilings	27
B.2. Proof of semi-continuity	29
References	30

1. INTRODUCTION

Let G be a countable, finitely generated amenable group, and let A be a finite set of symbols. G acts by shifts on A^G , which is endowed with the product topology. A *symbolic dynamical system* or a *shift* is a closed, shift-invariant subset of A^G . The space of all shifts $\mathfrak{S} = \mathfrak{S}(G, A)$ admits a natural topology induced from the Hausdorff topology; in this topology two shifts are close if they coincide on large finite subsets of G .

An important dynamical property of a shift is its *entropy*, which roughly measures the exponential growth rate of its projections on finite sets. We show that for every $c, \varepsilon \geq 0$ the set of shifts with entropy between c and $c + \varepsilon$ is dense in $\mathfrak{S}_{\geq c}$, the set of shifts with entropy at least c . We furthermore show that this still holds for some interesting subsets of the space of shifts, such as the weakly mixing shifts and the transitive shifts.

This result is novel even for the case of $G = \mathbb{Z}$ and $c = 0$, although in this case it admits a simpler proof. The case of $G = \mathbb{Z}^2$ already seems to require all of our machinery.

We show that entropy is upper semi-continuous, from which it follows that the shifts of entropy c are dense in $\mathfrak{S}_{\geq c}$, and in fact are *generic*. The study of genericity in dynamical systems has a long and fruitful history. For example, Halmos [5] showed that a generic measure preserving transformation is weakly mixing, and Rohlin [9] showed that it is not strongly mixing; this constituted the first proof that there exist weakly mixing transformations that are not strongly mixing.

An immediate consequence of the fact that the shifts of entropy c are generic in $\mathfrak{S}_{\geq c}$ is that they *exist*. Hence for every $c \in [0, \log |A|]$ there exists a shift in A^G with entropy c (and in fact, this can be achieved using weakly mixing shifts). These results use the Baire Category Theorem and thus are not explicitly constructive; for $G = \mathbb{Z}$, an explicit construction of shifts of every entropy is given by Weiss [10] (and see also the forthcoming book by Coornaert [3]).

Hochman [6] studies the space of transitive shifts over \mathbb{Z} , but where the symbols are in the Hilbert cube. Among many results, he shows that the zero entropy shifts are generic. Note that the space of shifts over the Hilbert cube is different than the one we study: for example, it is connected, while ours is zero dimensional and in fact has isolated points. These differences induce different generic properties of the two spaces. For example, Hochman shows that the weakly mixing shifts are dense in the space of transitive shifts; this is not true over finite alphabets, since there are isolated transitive shifts that are not weakly mixing. Accordingly, our techniques are different than Hochman's, and specifically are more combinatorial in nature.

Our main tool are *quasi-tiling shifts*. A *tiling* of a group is a decomposition of its elements into disjoint finite sets, where each set is a translate of some finite number of "tiles". In a *quasi-tiling* there are still only a finite number of tiles, but not all the group elements need be covered, and some may be covered by more than one tile. A "good" quasi-tiling will have few such "errors". Another desirable property of a quasi-tiling is that the tiles have small boundaries. Ornstein and Weiss [8] show that finite subsets of amenable groups can be arbitrarily well quasi-tiled, using tiles that have arbitrarily small boundaries.

Quasi-tiling shifts are closed, shift-invariant sets of quasi-tilings of the entire group. We show that "good" ones exist, in the sense that they have few errors, have tiles with small boundaries, and have disjoint tiles. Furthermore, they are strongly irreducible and have low entropy.

In a very recent paper by Downarowicz, Huczek and Zhang [4], it is shown that amenable groups can in fact be *tiled* by tiles with small boundaries, and that there exist tiling shifts with low entropy. We use some of their intermediate results in this paper, but not the actual tiling shifts, since these are not guaranteed to be strongly irreducible.

1.1. Definitions and results.

1.1.1. *The space of shifts.* Let G be a countable, finitely generated amenable group. We fix $d(\cdot, \cdot)$, a left-invariant word length metric on G . Let A be a finite set of symbols. A^G is endowed with the product topology and with the G -action by shifts, given by $[gx](h) = x(g^{-1}h)$.

A closed, G -invariant subset X of A^G is called a *shift*. The space of all shifts is denoted by $\mathfrak{S} = \mathfrak{S}(G, A)$ and is equipped with the Hausdorff topology, or, more precisely, with its restriction to the shifts. As such it is a compact Polish space.

This topology admits a simple geometrical definition. For a finite set $K \subset G$, a shift $X \in \mathfrak{S}$, and an $x \in X$, let $x_K: K \rightarrow A$ be the

restriction of x to K , and let X_K be the projection of X on K :

$$X_K = \{x_K : x \in X\}.$$

Then a sequence of shifts $\{X^n\}_{n=1}^\infty$ converges to X if and only if for all finite $K \subset G$ it holds that $\lim_n X_K^n = X_K$, or equivalently that $X_K^n = X_K$ for n large enough.

A subset of \mathfrak{S} is a G_δ if it is a countable intersection of open sets. A set that contains a dense G_δ subset is called generic (equivalently: residual or comeagre) and its complement is called meager. By the Baire Category Theorem, a countable intersection of dense open sets is comeagre.

1.1.2. *Følner sets and entropy.* Let $B_r \subset G$ be the ball of radius r around the origin in G :

$$B_r = \{g \in G : d(e, g) \leq r\}.$$

Definition 1.1. *Let $F \subset G$ be finite. For $r \in \mathbb{N}$, the r -boundary of F is*

$$\partial_r F = \{g \in G : gB_r \cap F \neq \emptyset \text{ and } gB_r \cap (G \setminus F) \neq \emptyset\},$$

and the r -boundary ratio of F is

$$\rho_r(F) = \frac{|\partial_r F|}{|F|},$$

where $|\cdot|$ is the counting measure on G . We say that F is (r, ε) -invariant if $\rho_r(F) \leq \varepsilon$.

A characterization of amenable groups is the existence of *Følner sequences*: a sequence $\{F_n\}$ of finite subset of G is Følner if $\lim_n \rho_r(F_n) = 0$ for some (equivalently, for every) $r > 0$.

Definition 1.2. *The entropy of $X \in \mathfrak{S}$ is given by*

$$h(X) = \lim_{n \rightarrow \infty} \frac{1}{|F_n|} \cdot \log |X_{F_n}|.$$

That this limit exists and is independent of the choice of Følner sequence was shown by Ornstein and Weiss [8].

Proposition 1.3. *The entropy map $h: \mathfrak{S} \rightarrow \mathbb{R}^+$ is upper semi-continuous.*

The (relatively straightforward) proof of this proposition appears in Appendix B; see Lindenstrauss and Weiss [7, Appendix 6] for a related proof that uses similar ideas. It follows that the set of shifts with entropy at least c is a closed subset of \mathfrak{S} , for any $c \geq 0$. We denote it by $\mathfrak{S}_{\geq c}$, and denote by \mathfrak{S}_c the set of shifts with entropy exactly c .

1.1.3. *Main results.* Our main result is the following.

Theorem 1.4. *For every $c \geq 0$, \mathfrak{S}_c is comeagre in $\mathfrak{S}_{\geq c}$.*

We in fact prove a more general statement. Strongly irreducible shifts (see definition below) have good mixing properties, and in particular are strongly mixing. We call a class \mathcal{F} of shifts *strongly-stable* if it is closed under the operations of taking factors and of taking products with strongly irreducible shifts. For example, the transitive shifts, the recurrent shifts, the weakly mixing shifts, the strongly mixing shifts and the strongly irreducible shifts are all strongly-stable classes. The first three are also G_δ subsets of \mathfrak{S} [6]. Let $\mathfrak{S}^{\mathcal{F}}$ be the subset of \mathfrak{S} that is in \mathcal{F} , and define $\mathfrak{S}_c^{\mathcal{F}}$ and $\mathfrak{S}_{\geq c}^{\mathcal{F}}$ similarly.

Theorem 1.5. *Let \mathcal{F} be a strongly-stable class such that $\mathfrak{S}^{\mathcal{F}}$ is a G_δ subset of \mathfrak{S} . Then for every $c \geq 0$, $\mathfrak{S}_c^{\mathcal{F}}$ is comeagre in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$.*

More generally, when $\mathfrak{S}^{\mathcal{F}}$ is not necessarily a G_δ , for every $c \geq 0$ and $\varepsilon > 0$, the subset of shifts in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$ with entropy in $[c, c + \varepsilon)$ is dense in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$.

Since the class of all shifts is also strongly-stable, Theorem 1.5 implies Theorem 1.4. Note again that the weakly mixing shifts and the transitive shifts both form a G_δ [6].

Theorem 1.5 admits a number of interesting immediate corollaries.

Corollary 1.6. *For every $c \in [0, \log |A|]$ there exists a weakly mixing shift $X \subseteq A^G$ with $h(X) = c$.*

It is not known whether the set of strongly mixing shifts is a G_δ (see Hochman [6]). However, the next result follows directly from Theorem 1.5.

Corollary 1.7. *The set of strongly irreducible shifts is not a G_δ .*

This is a consequence of the fact that non-trivial (i.e., of cardinality greater than one) strongly irreducible shifts have positive entropy. Another consequence on strongly irreducible shifts is the following corollary, which follows immediately from the second part of Theorem 1.5 and the fact that non-trivial strongly irreducible shifts have positive entropy.

Corollary 1.8. *All non-trivial strongly irreducible shifts are non-isolated points in \mathfrak{S} .*

Another direct consequence is the following.

Corollary 1.9. *The isolated points in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$ all have entropy c .*

Since any isolated point in $\mathfrak{S}_{\geq c}$ has to be a shift of finite type, and since there are only countably many such shifts, it also follows that

Corollary 1.10. *For all but a countable set of values of c , the space $\mathfrak{S}_{\geq c}$ has no isolated points, and is therefore homeomorphic to the Cantor space.*

We furthermore show that for a dense set of entropy values $\mathfrak{S}_{\geq c}$ indeed has isolated points.

Theorem 1.11. *There is a dense subset $C \subset [0, \log |A|]$ such that for each $c \in C$ there exists a strongly irreducible shift $X \subseteq A^G$ with $h(X) = c$, and where X is an isolated point in $\mathfrak{S}_{\geq c}$.*

1.1.4. *Quasi-tiling shifts.* To prove Theorem 1.5 we show the existence of good *strongly irreducible quasi-tiling shifts*. A quasi-tiling of a countable group G is a partial covering of G by translates of a finite number of “tiles” or finite subsets of G . A quasi-tiling is “good” if there are few overlaps and uncovered regions. See Ornstein and Weiss [8], and also Ceccherini-Silberstein and Coornaert [1, Section 5.6], who use a somewhat different formulation. Note that usually tile-translates in quasi-tilings are allowed to overlap. We, however, will only use *disjoint* quasi-tilings, where tile-translates do not overlap.

Strong irreducibility is a strong form of mixing; in particular stronger than strong mixing. We now define strong irreducibility.

Definition 1.12. *A shift X is strongly irreducible if there exists an $r \in \mathbb{N}$ such that, for any two finite subsets $K, H \in G$ which satisfy $d(k, h) > r$ for all $k \in K$ and $h \in H$, and any $x, y \in X$, there exists a $z \in X$ with $z_K = x_K$ and $z_H = y_H$.*

Definition 1.13. *A tile set $\mathcal{T} = (T_1, \dots, T_n)$ is a finite collection of finite subsets of G , each of which includes the identity.*

Given an $r > 0$ and $\varepsilon > 0$, a tile set \mathcal{T} is said to be (r, ε) -invariant if each T_i is (r, ε) -invariant.

We denote by $r(\mathcal{T}) = \max\{d(e, g) : g \in T_i, i = 1, \dots, n\}$ the radius of \mathcal{T} . This is the radius of the smallest ball that contains all the tiles in \mathcal{T} .

A tile-translate of T_i is a set of the form hT_i , for some $h \in G$.

We next define a \mathcal{T} -tiling, which is a covering of G by disjoint translates of the tiles in \mathcal{T} .

Definition 1.14. *Let \mathcal{T} be a tile set. $T: G \rightarrow \mathcal{T} \cup \emptyset$ is called a \mathcal{T} -tiling if $gT(g) \cap hT(h) = \emptyset$ for all $g \neq h \in G$, and if, for each $g \in G$, there exists a unique $h \in G$ with $g \in hT(h)$.*

Here h can be thought of as the “corner” of the tile-translate $hT(h)$. As mentioned above, recently Downarowicz, Huczek and Zhang [4] showed that all amenable groups admit good tilings. We will not use this result.

A \mathcal{T} -quasi-tiling is much like a tiling, except that not all of G need be covered. We will still require the tile-translates to be disjoint (thus parting from the usual terminology), and will sometimes refer to these quasi-tilings explicitly as disjoint quasi-tilings.

Definition 1.15. *Let \mathcal{T} be a tile set. $T: G \rightarrow \mathcal{T} \cup \emptyset$ is called a (disjoint) \mathcal{T} -quasi-tiling if $gT(g) \cap hT(h) = \emptyset$ for all $g, h \in G$.*

A $g \in G$ is said to be a T -error if it is not an element of any tile-translate $gT(g)$.

In a slight abuse of notation, we denote the set of (disjoint) \mathcal{T} -quasi-tilings by \mathcal{T}^G .

A “good” quasi-tiling $T \in \mathcal{T}^G$ will have few errors. The following definition is a useful way to quantify that.

Definition 1.16. *Let $T \in \mathcal{T}^G$ be a quasi-tiling. Denote by $e(T, F)$ be the number of T -errors in $F \subset G$. The error density of a quasi-tiling T is said to be at most ε if there exists some $\delta > 0$ such that*

$$\frac{e(T, F)}{|F|} \leq \varepsilon$$

whenever F is $(1, \delta)$ -invariant.

1.1.5. *Shifts of quasi-tilings.* Given a tile set \mathcal{T} , a *quasi-tiling shift* is simply a subshift of \mathcal{T}^G .

Definition 1.17. *The error density of a \mathcal{T} -quasi-tiling shift Q is said to be at most ε if there exists some $\delta > 0$ such that*

$$\frac{e(T, F)}{|F|} \leq \varepsilon$$

for all $T \in Q$ and whenever F is $(1, \delta)$ -invariant.

The following theorem shows that “good” disjoint quasi-tiling shifts exist. Its proof is straightforward, given the results of Ornstein and Weiss [8] and a recent paper of Downarowicz, Huczek and Zhang [4]; the only missing ingredient is strong irreducibility.

Theorem 1.18 (Ornstein and Weiss [8], Downarowicz, Huczek and Zhang [4]). *For every r, ε there exists an (r, ε) -invariant tile set \mathcal{T} and a shift $Q \subset \mathcal{T}^G$ that has the following properties:*

- (1) *The error density of Q is at most ε .*

- (2) $h(Q) < \varepsilon$.
 (3) Q is strongly irreducible.

The following last definition of this section will be useful.

Definition 1.19. *A tile set \mathcal{T} is ε -good if there exists a quasi-tiling shift $Q \subseteq \mathcal{T}^G$ that satisfies the conditions of Theorem 1.18 with $r = 1$.*

1.2. Organization and notation. The remainder of the paper proceeds as follows. In Section 2 we show the existence of good quasi-tiling shifts. In Section 3 we show that every shift can be approximated by a shift with entropy close to zero. We use this construction to prove in Section 4 that every shift can be approximated by shifts of any lower entropy, and then prove our main theorem. Finally, in Section 5 we show that for a dense set of entropy values c the space $\mathfrak{S}_{\geq c}$ has isolated points.

We provide below an overview of the notation and nomenclature used in this paper.

G	Finitely generated amenable group
A	Finite alphabet
B_r	The ball of radius r around the identity in G
$\partial_r K$	The boundary of radius r of a set $K \subseteq G$
$\rho_r K$	The ratio between $\partial_r K$ and $ K $
(r, ε) -invariant	A set F satisfying $\rho_r F \leq \varepsilon$
$h(X)$	The entropy of a shift X
\mathfrak{S}	The space of subshifts of A^G
\mathfrak{S}_c	The space of subshifts of A^G with entropy c
$\mathfrak{S}_{\geq c}$	The space of subshifts of A^G with entropy c or higher
\mathcal{F}	A strongly stable class of shifts
$\mathfrak{S}^{\mathcal{F}}$	The space of subshifts of A^G belonging to \mathcal{F}
tile	A finite subset of G containing the identity
tile-translate	A set of the form hT_i where T_i is a tile and $h \in G$
X, Y	Shifts, or closed, shift-invariant subsets of A^G
X_K	The projection of X to $K \subseteq G$
x	Element of a shift X
x_K	The projection of x to $K \subseteq G$
\mathcal{T}	A tile set, or a finite set of tiles
ε -good tile set	A tile set allowing good quasi-tilings (as per Theorem 1.18)
$r(\mathcal{T})$	The radius of the tile set \mathcal{T}
T, S	A quasi-tiling
\mathcal{T}^G	The set of disjoint \mathcal{T} -quasi-tilings
$M^{\mathcal{T}}$	The set of maximal disjoint \mathcal{T} -quasi-tilings
$e(T, K)$	The number of errors of T on a finite $K \subset G$

2. QUASI-TILING SHIFTS

Let \mathcal{T} and \mathcal{S} be tile sets, and let T and S be \mathcal{T} - and \mathcal{S} -quasi-tilings, respectively. A natural way of combining T and S to a new (disjoint) quasi-tiling is to add to T any tile-translate of S that is disjoint from all the tile-translates of T . Formally, we define the map

$$\psi: \mathcal{T}^G \times \mathcal{S}^G \rightarrow (\mathcal{T} \cup \mathcal{S})^G$$

by

$$[\psi(T, S)](g) = \begin{cases} T(g) & \text{if } T(g) \neq \emptyset \\ S(g) & \text{if } S(g) \neq \emptyset \text{ and } gS(g) \cap \cup_h hT(h) = \emptyset \\ \emptyset & \text{otherwise.} \end{cases}$$

It is easy to see that this map is continuous and commutes with the shift; continuity follows from the fact that whether or not $gS(g) \cap \cup_h hT(h) = \emptyset$ depends only on the values of S and T within distance $r(\mathcal{T}) + r(\mathcal{S})$ of g .

We extend ψ to a function on a product of more than two quasi-tilings, as follows. Let

$$\psi: \mathcal{T}_1^G \times \cdots \times \mathcal{T}_n^G \rightarrow (\mathcal{T}_1 \cup \cdots \cup \mathcal{T}_n)^G$$

be defined recursively by

$$\psi(T_1, \dots, T_n) = \psi(\psi(T_1, \dots, T_{n-1}), T_n).$$

That is, we start with T_1 , add to it any tile in T_2 that is disjoint, add to that any tile in T_3 that is disjoint, etc. By the same reasoning used above, ψ here is still continuous and commutes with the shift.

Given a tile set \mathcal{T} , the \mathcal{T} -quasi-tilings can be ordered by inclusion; namely, $T' \geq T$ if $T(g) \neq \emptyset$ implies $T'(g) = T(g)$. By a straightforward application of Zorn's lemma, maximal disjoint \mathcal{T} -quasi-tilings exist. It is clear that these maximal quasi-tilings form a closed, shift-invariant set, and are hence a subshift of \mathcal{T}^G , which we denote by $M^{\mathcal{T}}$.

Given tile sets $\mathcal{T}_1, \dots, \mathcal{T}_n$, we can apply ψ to $M^{\mathcal{T}_1} \times \cdots \times M^{\mathcal{T}_n}$. The result will be a subshift of $(\mathcal{T}_1 \cup \cdots \cup \mathcal{T}_n)^G$, since ψ is continuous and commutes with the shift.

The following theorem is a more detailed restatement of Theorem 1.18.

Theorem 2.1. *For every r, ε there exists an (r, ε) -invariant tile set $\mathcal{T} = \mathcal{T}_1 \cup \cdots \cup \mathcal{T}_n$ such that $Q = \psi(M^{\mathcal{T}_1} \times \cdots \times M^{\mathcal{T}_n})$ has the following properties:*

- (1) *The error density of Q is at most ε .*
- (2) *$h(Q) < \varepsilon$.*

(3) Q is strongly irreducible.

This construction is based on (and is nearly identical to) the original construction of Ornstein and Weiss; they however use ε -disjoint quasi-tilings rather than disjoint ones, and prove a version of (1) for finite subsets of G . The complete proof of (1) and (2) appears in [4, Lemmata 4.1 and 4.2]. It therefore remains to be shown that Q is strongly irreducible. To show this, it is enough to show that each $M^{\mathcal{T}^i}$ is strongly irreducible, since Q is formed from these shifts by taking products and factors, and strong irreducibility is closed under these operations. This is done in the following lemma, which thus concludes the proof of Theorem 2.1.

Lemma 2.2. *Let \mathcal{S} be a tile set. Then $M^{\mathcal{S}}$ is strongly irreducible.*

Proof. Let $r = 4r(\mathcal{S})$. Fix two finite subsets $K, H \in G$ which satisfy $d(k, h) > r$ for all $k \in K$ and $h \in H$. Choose any $T, S \in M^{\mathcal{S}}$. We will prove the claim by showing that there exists an $R \in M^{\mathcal{S}}$ with $R_K = T_K$ and $R_H = S_H$.

Let $K' = K \cup \partial_{r(\mathcal{S})}K$ and likewise let $H' = H \cup \partial_{r(\mathcal{S})}H$. Then K' and H' are disjoint, since $d(K, H) > 4r(\mathcal{S})$, and furthermore the tile-translates whose corners are in $K' \cup H'$ are disjoint. Consider the set of \mathcal{S} -quasi-tilings that include all the tile-translates that T has on K' and that S has on H' . By Zorn's lemma there exists a maximal (with respect to inclusion) element in this set, which we will call R . By definition, R includes all the tile-translates that T has on K and that S has on H . It thus remains to be shown that (1) R does not include any additional tile-translates on $K \cup H$, and (2) R is an element of $M^{\mathcal{S}}$.

To see (1), note that any tile that can be added to R in K (while keeping it a disjoint quasi-tiling) can also be added to T , since R includes all the tile-translates of T that are close enough to K to intersect tile-translates in K ; those are the precisely the tile-translates in K' . But it is not possible to add more tile-translates to T , since it is maximal. The same applies to adding to H .

Finally, to see (2), note that R is maximal in the set of quasi-tilings that are greater (again with respect to inclusion) than the quasi-tiling which only includes the tiles of T on K' and the tiles of S on H' . Hence it is also maximal in the set of all quasi-tilings. \square

Rather than using Theorem 2.1 as stated for our work, we will need a corollary which we are now almost in a position to prove. We will need the following additional definition.

Definition 2.3. Let T be a \mathcal{T} -quasi-tiling, and fix $r \in \mathbb{N}$. The r -interior of T is the set

$$\bigcup_{g \in G} gT(g) \setminus \bigcup_{g \in G} \partial_r gT(g),$$

which is the union of all the tile-translates in T , from which is removed the union of the boundaries (of radius r) of all the tile-translates. The r -exterior of T is the complement of the interior; alternatively, it is the set of elements of G that are not covered in T , or are within distance r of the boundary of a tile-translate in T .

A \mathcal{T} -quasi-tiling shift Q has r -exterior density at most ε if there exists a $\delta > 0$ such that for any finite, $(1, \delta)$ -invariant $F \subset G$, and any $T \in Q$ it holds that the intersection of the r -exterior of T with F is of size at most $\varepsilon|F|$.

With these definitions in place we are ready to state our corollary.

Theorem 2.4. For all ε, r there exist a tile set \mathcal{T} and a strongly irreducible \mathcal{T} -quasi-tiling shift Q with r -exterior density at most ε , and such that $h(Q) < \varepsilon$.

Proof. Fix r, ε . Let \mathcal{T} be a tile set and Q a \mathcal{T} -quasi-tiling shift with the following properties.

- (1) \mathcal{T} is $(r, \varepsilon/2)$ -invariant.
- (2) The error density of Q is at most $\varepsilon/2$.
- (3) $h(Q) < \varepsilon$.

The existence of such \mathcal{T} and Q is guaranteed by Theorem 2.1. It is straightforward to show that the r -exterior density of Q is at most ε . \square

3. LOW ENTROPY APPROXIMATION

In this section we prove the following proposition.

Proposition 3.1. Let \mathcal{F} be a strongly-stable class of shifts. For every $\varepsilon > 0$, the set of shifts in \mathcal{F} with entropy in $[0, \varepsilon)$ is dense in $\mathfrak{S}^{\mathcal{F}}$.

3.1. Overview. To prove this proposition we construct shifts that are in \mathcal{F} , have entropy in $[0, \varepsilon)$, and arbitrarily well approximate a given shift X .

To construct an approximating shift Y we use strongly irreducible quasi-tiling shifts. Fix $r > 0$, let $T = (T_1, \dots, T_n)$ be a tile set, and let Q be a \mathcal{T} -quasi-tiling shift. Let $T \in Q$, and recall that the r -exterior of T is the subset of G that is within radius r of the boundary of a tile-translate, or is not covered by a tile-translate.

The idea behind the construction of Y is the following. Given a configuration $x \in X$, we first choose a quasi-tiling $T \in Q$, and map the pair (x, T) to a $y' \in Y' \subset A^G$, as follows. On the r -exterior of T we let $y'(g) = x(g)$, and elsewhere we let $y'(g) = \emptyset$, for some fixed symbol $\emptyset \in A$. We then map the pair (y', T) to a $y \in Y$, by leaving $y(g) = y'(g) = x(g)$ on the r -exterior of T , and choosing $y(g)$ elsewhere using a “completion map” that is locally compatible with X , but that is completely determined by the values of y on the exterior. Both of the maps $X \times Q \mapsto Y'$ and $Y' \times Q \mapsto Y$ are continuous and commute with the shift, and therefore both Y' and Y are shifts. The entropy of Y' is controlled by the fact that the number of exterior points is low. The entropy of Y is bounded by $h(Y') + h(Q)$, and hence is also low. Y is a good approximation of X since we can take r to be large, and x and y agree on the r -exterior of T .

3.2. Deletion and local completion maps. Let K be a subset of G , and let $e: K \rightarrow \{\mathbf{ext}, \mathbf{int}\}$ be a labeling of K into exterior and interior points. Let $\emptyset \in A$ be an arbitrary “distinguished” symbol. Given $y \in A^K$, we define $y \cdot e \in A^G$ by

$$(1) \quad [y \cdot e](g) = \begin{cases} y(g) & \text{if } e(g) = \mathbf{ext} \\ \emptyset & \text{otherwise.} \end{cases}$$

Intuitively, multiplying y by e leaves the configuration unchanged on the exterior of K , and substitutes the \emptyset symbol on the interior.

Fix $r \in \mathbb{N}$ for the remainder of this section; this will be a parameter in the various constructions that follow. Given a \mathcal{T} -quasi-tiling $T \in \mathcal{T}^G$, we let $e^T \in \{\mathbf{int}, \mathbf{ext}\}^G$ be given by $e^T(g) = \mathbf{ext}$ for any g on the r -exterior of T , and $e^T(g) = \mathbf{int}$ elsewhere. As can be easily verified, the map $T \mapsto e^T$ is a continuous and shift-equivariant map.

We next define the “deletion” map $\delta: A^G \times \mathcal{T}^G \rightarrow A^G$. Given a configuration $y \in A^G$ and a quasi-tiling T , δ outputs a configuration in which the r -exterior is left unchanged, and the interior of the tiles is “deleted”, leaving there the symbol $\emptyset \in A$. Formally,

$$\delta(y, T) = y \cdot e^T = \begin{cases} y(g) & \text{when } e^T(g) = \mathbf{ext} \\ \emptyset & \text{otherwise.} \end{cases}$$

The product $y \cdot e^T$ is given in the sense of (1).

Note that δ is continuous and shift-equivariant, since the map $T \mapsto e^T$ is. Hence the image $\delta(X \times \mathcal{T}^G)$ is a shift, as is $\delta(X \times Q)$, for any quasi-tiling shift Q . We show that the entropy of this shift can be controlled, using an appropriate choice of Q .

Proposition 3.2. *Let the r -exterior density of Q be at most ε . Then*

$$h(\delta(X \times Q)) \leq \varepsilon \log \frac{3|A|}{\varepsilon},$$

where $X \subseteq A^G$.

Proof. By the definition of δ , on every sufficiently invariant finite $F \subset G$, and every $y \in \delta(X \times Q)$ it holds that y_F is equal to \emptyset on all but at most $\varepsilon|F|$ of the elements of F . Hence, by Lemma A.2, the entropy of $h(\delta(X \times Q))$ is at most $\varepsilon \log \frac{3|A|}{\varepsilon}$. \square

Given a shift X and a finite $K \subset G$, a *local (X, K) -completion map* is a function $m: A^K \times \{\mathbf{int}, \mathbf{ext}\}^K \rightarrow A^K$ with the following properties. Let $(y, e) \in A^K \times \{\mathbf{int}, \mathbf{ext}\}^K$, and let there exist some $x \in X_K$ such that $y \cdot e = x \cdot e$ (i.e. x and y agree on the exterior). Then

- (1) $m(y, e) \in X_K$.
- (2) $m(y, e) \cdot e = y \cdot e$.
- (3) $m(y \cdot e, e) = m(y, e)$.

Intuitively, the map m completes any configuration that is compatible with X_K on the exterior of K to a complete configuration in X_K ; that is the first property. The second property ensures that this is indeed a completion of the configuration on the exterior; $m(y, e)$ and y agree on the exterior. The third property ensures that the completion is independent of the configuration in the interior; $m(y, e)$ only depends on $y(g)$ if $f(g) = \mathbf{ext}$.

The following claim formalizes the idea that for every $x \in X$, there exists a local (K, X) -completion map that will output x_K whenever its input matches x on the exterior. Its proof is straightforward, if tedious, and we omit it.

Claim 3.3. *For every finite $K \subset G$ and $x \in X_K$ there exists a local (X, K) -completion map m such that $m(y, e) = x$ for all $e \in \{\mathbf{int}, \mathbf{ext}\}^K$ and $y \in A^K$ such that $y \cdot e = x \cdot e$.*

Any such local (K, X) -completion function m is said to be (x, K) -compatible.

Note that a particular implication of this claim is that for every K there exists a local (X, K) -completion map.

Let T_i be a tile in \mathcal{T} . We will consider local (X, T_i) -completion maps; these will complete configurations on the exterior of the tile to the rest of the tile. They exist by Claim 3.3. In a slight abuse of notation, we will apply a local (X, T_i) -completion map $m: A^{T_i} \times \{\mathbf{int}, \mathbf{ext}\}^{T_i} \rightarrow A^{T_i}$ to any tile-translate hT_i in the natural way. That is, the function

$m': A^{hT_i} \times \{\text{int}, \text{ext}\}^{hT_i} \rightarrow A^{hT_i}$ given by

$$m'(y, e) = h[m(h^{-1}y, h^{-1}e)]$$

will also be denoted by m .

3.3. Global completion maps. Let (m_1, \dots, m_n) be local (X, T_i) -completion maps of each of the tiles in the tile set $\mathcal{T} = (T_1, \dots, T_n)$. We would like to define a global completion map $M: A^G \times \mathcal{T}^G \rightarrow A^G$ that applies m_i to each T_i -tile-translate separately, after calculating its exterior.

Definition 3.4. *Given a configuration $y \in A^G$ and a quasi-tiling $T \in \mathcal{T}^G$, define the global completion map $M(y, T)$ as follows. For g that is not in any tile-translate, set*

$$(2) \quad [M(y, T)](g) = y(g).$$

For any T tile-translate hT_i set

$$(3) \quad M(y, T)_{hT_i} = m_i(y_{hT_i}, e_{hT_i}^T).$$

It follows that

$$(4) \quad [M(y, T)](g) = y(g) \quad \text{whenever } e^T(g) = \text{ext}.$$

The next claim follows immediately from the local nature of this definition.

Claim 3.5. *M is continuous and shift-equivariant.*

Hence $Y = M(X \times Q)$ is a subshift of A^G . Furthermore, since Q is strongly irreducible, then $Y \in \mathfrak{S}^{\mathcal{F}}$ whenever $X \in \mathfrak{S}^{\mathcal{F}}$.

We next show that M inherits the three properties of the local maps m_1, \dots, m_n .

Claim 3.6. *Fix $(y, T) \in A^G \times \mathcal{T}^G$ such that $y \cdot e^T = x \cdot e^T$ for some $x \in X$. Then*

- (1) $M(y, T)_{hT_i} \in X_{hT_i}$ for any T -tile-translate hT_i .
- (2) $M(y, T) \cdot e^T = y \cdot e^T$.
- (3) $M(y \cdot e^T, T) = M(y, T)$.

The first property means that the projection of $M(y, T)$ on a tile-translate hT_i coincides with the projection to hT_i of some $x \in X$. The second property means that M leaves y unchanged on the exterior. The third means that M depends only on the values of y on the exterior.

Proof. Let (y, T) satisfy the claim hypothesis. The first property follows immediately from (3) and the first property of local completion maps. The second property is a restatement of (4).

To see the third property, consider two cases: that g is in the r -exterior of T , and that it is in the interior. In the first case,

$$[M(y \cdot e^T, T)](g) = [y \cdot e^T](g) = y(g) = [M(y, T)](g),$$

where the first equality follows from (4) (since g is in the exterior), the second from the definition of $y \cdot e^T$ and the third again from (4).

In the second case, g is in the r -interior of T , and therefore an element of a tile-translate hT_i . Then by (3)

$$[M(y, T)](g) = [m_i(y_{hT_i}, e_{hT_i}^T)](g)$$

where m_i is the local (X, T_i) -completion map used to construct M . By the third property of local completion maps this can be written as

$$\begin{aligned} [M(y, T)](g) &= [m_i(y_{hT_i} \cdot e_{hT_i}^T, e_{hT_i}^T)](g) \\ &= [m_i((y \cdot e^T)_{hT_i}, e_{hT_i}^T)](g) \\ &= [M(y \cdot e^T, T)](g). \end{aligned}$$

□

Our approximating shift is simply going to be $Y = M(X \times Q)$, the image of $X \times Q$ under M , for an appropriate choice of Q .

To control the entropy of $Y = M(X \times Q)$, we first prove the following claim. Define the map δ_* by

$$\begin{aligned} \delta_*: A^G \times \mathcal{T}^G &\longrightarrow A^G \times \mathcal{T}^G \\ (y, T) &\longmapsto (\delta(y, T), T). \end{aligned}$$

That is, δ_* performs the same operation as δ , but also returns T , in a new, second coordinate. The next claim shows that first deleting and then completing is the same as just completing.

Claim 3.7. *For all $(x, T) \in X \times Q$ it holds that $M(x, T) = [M \circ \delta_*](x, T)$. It follows that $M(X \times Q) = [M \circ \delta_*)(X \times Q)$.*

Proof. Fix $(x, T) \in X \times Q$, and recall that $\delta(x, T) = x \cdot e^T$. It then follows from the third part of Claim 3.6 that $M(x, T) = M(x \cdot e^T, T) = M(\delta(x, T), T)$, and so

$$M(x, T) = [M \circ \delta_*](x, T).$$

□

Since

$$\delta_*(X \times Q) \subseteq \delta(X \times Q) \times Q,$$

it follows that

$$(5) \quad h(\delta_*(X \times Q)) \leq h(\delta(X \times Q)) + h(Q).$$

We can therefore now show that the entropy of Y is small.

Claim 3.8. $h(Y) \leq h(\delta(X \times Q)) + h(Q)$.

Proof.

$$\begin{aligned} h(Y) &= h(M(X \times Q)) \\ &= h([M \circ \delta_*)(X \times Q)) \\ &\leq h(\delta_*(X \times Q)) \\ &\leq h(\delta(X \times Q)) + h(Q), \end{aligned}$$

where the second equality is Claim 3.7, the next inequality is a consequence of the fact that factors decrease entropy, and the last is (5). \square

We next show that $M(X \times Q)$ is a good approximation of X . Recall that $r \in \mathbb{N}$ is a parameter in the construction of M .

Proposition 3.9. *Let Q be any quasi-tiling shift, and let $Y = M(X \times Q)$. Then $Y_{B_r} = X_{B_r}$.*

Proof. We will show containment in both directions. First, let $x \in X$. There exists a $T \in Q$ such that $\partial_r hT(h)$ contains B_r , by the shift-invariance of Q . Hence B_r is contained in the r -exterior of T . By the definition of M , $M(x, T)_{B_r} = x_{B_r}$, and so there exists a $y = M(x, T) \in Y$ such that $y_{B_r} = x_{B_r}$. Hence $X_{B_r} \subseteq Y_{B_r}$.

Now, let $y = M(x, T) \in Y$. B_r intersects the r -interior of at most one T -tile-translate, by the definition of the r -interior. If B_r is contained in the r -exterior of T then y agrees with x on B_r by the second part of Claim 3.6, and then $x_{B_r} = y_{B_r}$. Otherwise, B_r intersects the interior of some tile-translate hT_i . In this case it must be contained in hT_i , and so y_{B_r} is the projection to B_r of $M(x, T)_{hT_i}$. But the latter is in X_{hT_i} by the first part of Claim 3.6, and so $y_{B_r} \in X_{B_r}$. Hence $Y_{B_r} \subseteq X_{B_r}$. \square

We are now ready to prove Proposition 3.1, the main result of this section.

Proof of Proposition 3.1. Let $X \in \mathfrak{S}^{\mathcal{F}}$, $\varepsilon > 0$ and $r \in \mathbb{N}$. Define the deletion map δ and the global completion map M as above, using the parameter r . Let \mathcal{T} be a tile set, and let Q be a strongly irreducible \mathcal{T} -quasi-tiling shift with entropy at most $\varepsilon/2$, and r -exterior density at most ε' , where

$$3\varepsilon' \log \frac{|A|}{\varepsilon'} \leq \varepsilon/2.$$

The existence of such a Q is guaranteed by Theorem 2.4. Let $Y = M(X \times Q)$. By Claim 3.8 we have that

$$\begin{aligned} h(M(X \times Q)) &\leq h(\delta(X \times Q)) + h(Q) \\ &\leq h(\delta(X \times Q)) + \varepsilon/2. \end{aligned}$$

By Proposition 3.2 the first addend is at most $3\varepsilon' \log \frac{|A|}{\varepsilon'} \leq \varepsilon/2$, and we have shown that $h(Y) \leq \varepsilon$. Furthermore, by Proposition 3.9, $Y_{B_r} = X_{B_r}$.

Since $X \in \mathcal{F}$ and Q is strongly irreducible, $X \times Q \in \mathcal{F}$. By Claim 3.5 M is continuous and shift-equivariant. Hence $Y = M(X \times Q) \subset A^G$ is also in $\mathfrak{S}^{\mathcal{F}}$. It follows that for every ε, r there exists a $Y \in \mathfrak{S}^{\mathcal{F}}$ with $h(Y) < \varepsilon$ and $Y_{B_r} = X_{B_r}$, and so the set of shifts in \mathcal{F} with entropy in $[0, \varepsilon)$ is dense in $\mathfrak{S}^{\mathcal{F}}$. \square

4. FIXED ENTROPY APPROXIMATION

In this section we prove the following proposition. We then deduce from it the main result of this paper, Theorem 1.5.

Proposition 4.1. *Let \mathcal{F} be a strongly-stable class of shifts, let $c \geq 0$ and let $\varepsilon > 0$. The set of shifts in \mathcal{F} with entropy in $[c, c + \varepsilon)$ is dense in $\mathfrak{S}_c^{\mathcal{F}}$, the set of shifts in \mathcal{F} with entropy at least c .*

4.1. Overview. Our strategy is the following. Given a shift X , we construct for every $r \in \mathbb{N}$ and $\varepsilon > 0$ a sequence of shifts X^0, X^1, \dots, X^ℓ with the following properties:

- (1) $X^j \in \mathfrak{S}^{\mathcal{F}}$.
- (2) $X_{B_r}^j = X_{B_r}$; these are good approximations of X .
- (3) $h(X^0) \leq \varepsilon$.
- (4) $h(X^j) - h(X^{j-1}) \leq \varepsilon$.
- (5) $X \subseteq X^\ell$ and so $h(X^\ell) \geq h(X)$.

It follows that for all $0 \leq c \leq h(X)$ and some j , $h(X^j) \in [c, c + \varepsilon)$. Since this X^j is a good approximation of X , it follows that X can be arbitrarily well approximated by shifts with entropy in $[c, c + \varepsilon)$, which implies Proposition 4.1.

To construct X^0 we simply apply Proposition 3.1, using an appropriate quasi-tiling shift Q with low entropy. To construct the rest of these shifts, we first define a quasi-tiling shift Q_p , related to the shift Q , and which also has low entropy. We then construct a sequence of shifts $X_Q^1, X_Q^2, \dots, X_Q^\ell$ such that each X_Q^{j+1} is a factor of $X_Q^j \times Q_p$. It follows that $h(X_Q^{j+1}) \leq h(X_Q^j) + h(Q_p)$. The shifts X^j are each a factor

of X_Q^j , and we show that $h(X_Q^j) - h(Q) \leq h(X^j) \leq h(X_Q^j)$. It thus follows that $h(X^j) - h(X^{j-1}) \leq h(Q) + h(Q_p)$.

4.2. Q and Q_p . Given a \mathcal{T} -quasi-tiling shift Q , let Q_p be the \mathcal{T} -quasi-tiling shift which includes the quasi-tilings in Q , with some (or no) tile-translates removed:

$$Q_p = \{T' : \exists T \in Q \text{ s.t. } \forall g \in G, T'(g) = T(g) \text{ or } T'(g) = \emptyset\}.$$

Alternatively, Q_p is the set of quasi-tilings bounded from above (according to the inclusion relation) by some quasi-tiling in Q . The next proposition implies that the entropy of Q_p can be controlled by an appropriate choice of Q .

Proposition 4.2. *Let $\mathcal{T} = (T_1, \dots, T_n)$ be an (r, ε) -invariant tile set, and let Q be a \mathcal{T} -quasi-tiling. Then $h(Q_p) \leq h(Q) + \frac{2}{r} \log(3r)$.*

Proof. To prove this proposition, we construct Q_p somewhat differently. Given Q , let $Z \subseteq \{0, 1\}^G$ be the shift given by

$$Z = \{z \in \{0, 1\}^G : \exists T \in Q \text{ s.t. } z(g) = 1 \text{ implies } T(g) \neq \emptyset\}$$

and define

$$\iota: \mathcal{T}^G \times \{0, 1\}^G \rightarrow \mathcal{T}^G$$

by

$$[\iota(T, z)](g) = \begin{cases} T(g) & \text{if } z(g) = 1 \\ \emptyset & \text{otherwise.} \end{cases}$$

It is straightforward to verify that $Q_p = \iota(Q \times X)$. It follows that

$$(6) \quad h(Q_p) \leq h(Q) + h(X).$$

Since \mathcal{T} is (r, ε) -invariant, each of the tiles in \mathcal{T} is of size at least r . Hence, for any sufficiently invariant $F \subset G$, and for any $T \in Q$, the support of any T_F is of size at most $2/r|F|$. It follows that the same applies to Z . by Lemma A.2, $h(X) \leq \frac{2}{r} \log(3r)$. Hence, by (6), $h(Q_p) \leq h(Q) + \frac{2}{r} \log(3r)$. \square

The next claim follows immediately from the definitions of Q_p and strong irreducibility.

Proposition 4.3. *If Q is strongly irreducible then so is Q_p .*

4.3. Enumerating the completion maps. As in the previous section, fix $r \in \mathbb{N}$. For each tile T_i there only exist a finite number of local (X, T_i) -completion maps (with parameter r), since they map a finite set to a finite set. Hence there only exists a finite number ℓ of global completion maps, since each corresponds to a choice of local completions maps m_1, \dots, m_n . Enumerate them by M^1, M^2, \dots, M^ℓ .

For each global completion map $M^j: A^G \times \mathcal{T}^G \rightarrow A^G$ we define the corresponding map

$$N^j: A^G \times \mathcal{T}^G \times \mathcal{T}^G \longrightarrow A^G$$

which, given (y, T, S) , applies M^j to the configuration on a tile-translate hT_i only if $T(h) = S(h)$, and leaves the configuration unchanged elsewhere. Formally, for each tile-translate hT_i that is both in T and in S we set

$$N^j(y, T, S)_{hT_i} = M^j(y, T)_{hT_i},$$

and for g outside such tile-translates we set

$$[N^j(y, T, S)](g) = y(g),$$

N^j applies the same local completion maps that M^j does, but it only does so for tile-translates on which T and S agree. And it leaves the configuration elsewhere unchanged. It is immediate from this definition that N^j is continuous and shift-equivariant. It also follows that $N^j(y, T, S) \cdot e_T = y \cdot e_T$, in analogy to a property of M^j .

4.4. The shifts X^j . For $(x, T) \in X \times Q$, let

$$x^0 = M^1(x, T).$$

Given a $S^1 \in Q_p$, let

$$x^1 = N^1(x^0, T, S^1).$$

Likewise, given an additional $S^2 \in Q_p$, let

$$x^2 = N^2(x^1, T, S^2).$$

repeating the same logic, given x^{j-1} , and S^j , $1 \leq j \leq \ell$, let

$$x^j = N^j(x^{j-1}, T, S^j).$$

The next claim follows immediately from this definition, and from the fact that $N^j(\cdot, T, \cdot)$ leaves the configuration unchanged on the r -exterior of T .

Claim 4.4. *For all $0 \leq j \leq \ell$ it holds that $x^j \cdot e_T = x \cdot e_T$.*

Denote by X^j the set of all x^j that can be thus constructed, by applying the preceding maps to some $x \in X$, $T \in Q$, and $S^1, \dots, S^j \in Q_p$.

Claim 4.5. *X^j is a shift. Furthermore, if X is in some strongly-stable class \mathcal{F} and Q is strongly irreducible, then $X^j \in \mathcal{F}$.*

Proof. Define

$$M_*^j: A^G \times \mathcal{T}^G \longrightarrow A^G \times \mathcal{T}^G \\ (y, T) \longmapsto (M^j(y, T), T).$$

and likewise

$$N_*^j: A^G \times \mathcal{T}^G \times \mathcal{T}^G \longrightarrow A^G \times \mathcal{T}^G \\ (y, T, S) \longmapsto (N^j(y, T, S), T).$$

Let

$$X_Q^1 = M_*^1(X \times Q) \subset A^G \times Q,$$

and for $1 < j \leq \ell$ let

$$X_Q^j = N_*^j(X_Q^{j-1} \times Q_p).$$

Then $X^j = N^j(X_Q^{j-1} \times Q_p)$; equivalently, it is the projection to the first coordinate of X_Q^j , and is therefore a shift. Since Q is strongly irreducible then so is Q_p , by Proposition 4.3. Finally, since X^j is constructed by taking a series of factors and products involving X , Q and Q_p , it follows from the definition of strongly-stable classes that X^j is also in \mathcal{F} . \square

Proposition 4.6. *For $0 < j \leq \ell$*

- (1) $h(X^j) \leq h(X^{j-1}) + H(Q) + H(Q_p)$.
- (2) $X_{B_r} = X_{B_r}^j$.

Proof. (1) Recall that $X_Q^j = N^j(X_Q^{j-1} \times Q_p)$. Recall also that the projections of X_Q^{j-1} to the first and second coordinates are X^{j-1} and Q , respectively. Hence

$$h(X^j) \leq h(X_Q^{j-1}) + h(Q_p) \\ \leq h(X^{j-1}) + h(Q) + h(Q_p).$$

- (2) $X_{B_r}^0 = X_{B_r}$ by Proposition 3.9. Fix $(x, T) \in X \times Q$, and let $x^0 = M^1(x, T)$, $x^1 = N^1(x^0, T, S^1)$, etc. Recall that $x^j \cdot e_T = x \cdot e_T$ (Claim 4.4); the maps N^j do not alter the configuration on the r -exterior of T . Now, inside each tile-translate applying N^j either does nothing, or else is the same as applying M^j . Hence

the interiors of the tile-translates are also compatible with X , and the same argument of Proposition 3.9 applies in this case too.

□

We would next like to show the following proposition.

Proposition 4.7. $X \subseteq X^\ell$.

Proof. To prove this proposition, we construct for each $x \in X$ a sequence of quasi-tilings $\{S^j\}_{j=1}^\ell$ such that the associated x^ℓ is equal to x .

To this end, fix some $x \in X$ and $T \in Q$. Let hT_i be a T -tile-translate. Then by Claim 3.3 there exists a local (X, hT_i) -completion map m_i that is (x, hT_i) -compatible; that is, it completes to x_{hT_i} any configuration on hT_i that is compatible with x on the exterior of hT_i .

Now, there will an N^j that will use m_i to complete translates of T_i . Hence for such an N^j it will hold that

$$(7) \quad M^j(y, T, T')_{hT_i} = x_{hT_i}$$

for all T' such that $T'(h) = T(h) = i$ and for all $y \in A^G$ such that $(y \cdot e_T) = (x \cdot e_T)$. In this case we say that N^j is (x, hT_i) -compatible.

Let $C_{x,T}^j \subset G$ be the set of T -tile-translate locations h for which N^j is $(x, hT(h))$ -compatible:

$$C_{x,T}^j = \{h \in G : N^j \text{ is } (x, hT(h))\text{-compatible}\}.$$

Then for each T -tile-translate hT_i there is some $1 \leq j \leq \ell$ such that $h \in C_{x,T}^j$. We furthermore choose the sets $C_{x,T}^j$ so that each such h appears in exactly one set.

Let S^j be given by

$$S^j(h) = \begin{cases} T(h) & \text{if } h \in C_{x,T}^j \\ 0 & \text{otherwise} \end{cases}.$$

Clearly $S^j \in Q_p$. It follows immediately from the definition of $C_{x,T}^j$ that if $h \in C_{x,T}^j$ then

$$N^j(y, T, S^j)_{hT(h)} = x_{hT(h)}.$$

for all $y \in A^G$ such that $y \cdot e_T = x \cdot e_T$.

Following our construction of the shifts X^j , let

$$\begin{aligned} x^0 &= M^1(x, T) \\ x^1 &= N^1(x^0, T, S^1) \\ x^2 &= N^2(x^1, T, S^2) \\ &\vdots \\ x^\ell &= N^\ell(x^{\ell-1}, T, S^\ell). \end{aligned}$$

By the definition of the configurations S^j , for each T -tile-translate hT_i there will be a j such that $S^j(h) = T(h)$. Hence for that j it will hold that

$$x_{hT_i}^j = N^j(x^{j-1}, T, S^j)_{hT_i} = x_{hT_i}.$$

For all $j' > j$ the configuration on this tile will remain unchanged, and so $x_{hT_i}^{j'} = x_{hT_i}^j = x_{hT_i}$. Since this holds for all tile-translates, and since x^ℓ agrees with x on the r -exterior of T (Claim 4.4), it follows that $x^\ell = x$. Hence $x \in X^\ell$. \square

We are now ready to prove Proposition 4.1.

Proof of Proposition 4.1. Let $X \in \mathfrak{S}^{\mathcal{F}}$, $\varepsilon > 0$ and $r \in \mathbb{N}$. Define maps M^j as above, using the parameter r .

Let \mathcal{T} be a tile set, and let Q be a strongly irreducible \mathcal{T} -quasi-tiling shift such that $h(Q) + h(Q_p) < \varepsilon$ and such that if $X^0 = M^1(X \times Q)$ then $h(X^0) < \varepsilon$. The existence of such a quasi-tiling shift is guaranteed by Theorem 2.4 and Propositions 3.1 and 4.2, for r large enough.

Let X^1, X^2, \dots, X^ℓ be the shifts defined above. They are in $\mathfrak{S}^{\mathcal{F}}$, since they are constructed from X and Q by taking products and factors only.

By Proposition 4.7 $X \subseteq X^\ell$, and so $h(X^\ell) \geq h(X)$. On the other hand, by the first part of Proposition 4.6, $h(X^j) - h(X^{j-1}) \leq \varepsilon$. It follows that $h(X^j) \in [c, c + \varepsilon)$ for some $1 \leq j \leq \ell$. Finally, by the second part of Proposition 4.6, $X_{B_r}^j = X_{B_r}$. \square

The proof of our main theorem is now straightforward.

Proof of Theorem 1.5. We first prove that for all ε and all $c > 0$ the set of shifts with entropy in the range $[c, c + \varepsilon)$ is dense in the set of shifts with entropy greater than or equal to c .

Let X be a shift with entropy $h(X) \geq c$ and fix ε . Let Y_n be a subshift that agrees with X on B_r and has entropy in $[c, c + \varepsilon)$; the existence of these subshifts is guaranteed by Proposition 4.1. Since the sets B_r exhaust G , it follows that $\lim_n Y_n = X$. Thus the set of shifts with entropy in $[c, c + \varepsilon)$ is dense in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$.

By Proposition 1.3 the entropy function is upper semi-continuous. Hence the set of shifts with entropy in the range $[0, c + \varepsilon)$ is open, for all $c + \varepsilon > 0$. Thus it follows that the set of shifts with entropy in the range $[c, c + \varepsilon)$ is open in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$.

Since $\mathfrak{S}_c^{\mathcal{F}}$ is the intersection of the sets of shifts with entropy in the range $[c, c + 1/n)$ (for n in \mathbb{N}), and since, by the above, each of these is open and dense in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$, then by the Baire Category Theorem $\mathfrak{S}_c^{\mathcal{F}}$ is comeagre in $\mathfrak{S}_{\geq c}^{\mathcal{F}}$. \square

5. THE TOPOLOGY OF $\mathfrak{S}_{\geq c}$

In this section we prove Theorem 1.11.

Theorem 1.11. *There is a dense subset $C \subset [0, \log |A|]$ such that for each $c \in C$ there exists a strongly irreducible shift $X \subseteq A^G$ with $h(X) = c$, and where X is an isolated point in $\mathfrak{S}_{\geq c}$.*

We first note the following fact.

Proposition 5.1. *If X is a strongly irreducible shift of finite type with $h(X) = c$ then X is an isolated point in $\mathfrak{S}_{\geq c}$.*

Proof. Let $\lim_n X^n = X$ with $X^n \neq X$ for all n . Since X is of finite type, all but finitely many of the X^n have to be proper subshifts of X . Since X is strongly irreducible, each of its proper subshifts have entropy strictly less than c [2, Proposition 4.2]. Hence all but finitely many of the X^n are outside $\mathfrak{S}_{\geq c}$, and X is an isolated point in $\mathfrak{S}_{\geq c}$. \square

In light of this proposition, we prove Theorem 1.11 by finding a countable family of strongly irreducible shifts of finite type, whose entropies form a dense set in $[0, \log |A|]$. To this end, we employ a strategy similar to the one used in Section 4: for every $\varepsilon > 0$ we construct a sequence of shift X^1, X^2, \dots, X^ℓ , where

- (1) Each X^j is a strongly irreducible shift of finite type.
- (2) $X^{j-1} \subset X^j$.
- (3) $h(X^0) = 0$.
- (4) $h(X^j) - h(X^{j-1}) \leq \delta(\varepsilon)$, where $\lim_{\varepsilon \rightarrow 0} \delta(\varepsilon) = 0$.
- (5) $h(X^\ell) = \log |A|$.

It follows that the set of entropies $\{h(X^j)\}$ is dense in $[0, \log |A|]$.

Our proof proceeds as follows: for every ε we choose a sufficiently good tile set \mathcal{R} . The shift X^j will be the shift which is supported on at most an $j/|R_1|$ proportion of any \mathcal{R} -tile-translate, where R_1 is the largest tile in \mathcal{R} .

This is clearly a strongly irreducible shift of finite type. It is immediate that $h(X^{|R_1|}) = \log |A|$, and we show that $h(X^j)$ is low when j is low.

To show that $h(X^j) - h(X^{j-1}) \leq \varepsilon$ for large j we again use a strategy similar to that of Section 4: We show that X^j is a subshift of a factor of $X^{j-1} \times Y$, where Y is some shift of entropy at most $\delta(\varepsilon)$. This implies that $h(X^j) \leq h(X^{j-1} \times Q) = h(X^{j-1}) + \delta(\varepsilon)$.

5.1. Compatible quasi-tilings. Before defining the shifts X^j , we take a short intermission to define a technical tool which will be useful to that end, and state a simple claim regarding it.

Definition 5.2. Let $\mathcal{T} = (T_1, \dots, T_n)$ and $\mathcal{R} = (R_1, \dots, R_m)$ be tile sets, and fix $\varepsilon > 0$. A \mathcal{T} -quasi-tiling T is said to be $(\mathcal{R}, \varepsilon)$ -compatible if, for any tile-translate $K = hR_i$ it holds that $e(T, K) \leq \varepsilon|K|$.

Given tile sets \mathcal{T} and \mathcal{R} , and given an $\varepsilon > 0$, the set of $(\mathcal{R}, \varepsilon)$ -compatible \mathcal{T} -quasi-tilings is a shift. In fact, it is a shift of finite type. The next claim shows that it is non-empty, if \mathcal{R} is sufficiently good. It follows immediately from the definitions.

Claim 5.3. Let \mathcal{T} be an $\frac{1}{2}\varepsilon$ -good tile set, and let \mathcal{R} be a tile set with $\rho_1 R_i$ sufficiently small. Then there exists a \mathcal{T} -quasi-tiling that is $(\mathcal{R}, \varepsilon)$ -compatible.

5.2. Constructing the shifts X^j . Fix $\varepsilon > 0$. Let \mathcal{T} be an $\frac{1}{2}\varepsilon$ -good tile set, and let \mathcal{R} be an ε -good tile set such that $\rho_1 R_i$ is small enough so that, by Claim 5.3, there exist $(\mathcal{R}, \varepsilon)$ -compatible \mathcal{T} -quasi-tilings. Furthermore, let each tile in \mathcal{R} be of size at least $1/\varepsilon$, and be $(r(\mathcal{T}), \varepsilon)$ -invariant. Let Q be the \mathcal{T} -quasi-tiling shift of tilings with these properties.

Let R_1 be a largest tile in \mathcal{R} . For every $0 \leq j \leq |R_1|$ let $X^j \subseteq A^G$ be the shift of all $x \in A^G$ whose projections to any \mathcal{R} -tile-translate hR_i , x_{hR_i} , have support of size at most $j|R_i|/|R_1|$.

We next show that X^j has low entropy for low values of j .

Claim 5.4. $h(X^j) \leq \frac{j}{|R_1|} \log \frac{3|A||R_1|}{j} + \varepsilon \log |A|$.

Proof. Denote $\alpha = j/|R_1|$, so that the support of $X_{R_i}^j$ is of size at most $\alpha|R_i|$. It follows from Lemma A.1 that $X_{R_i}^j$ is of size at most $(3|A|/\alpha)^{\alpha|R_i|}$. We can now apply Proposition B.4 to X^j and \mathcal{R} , with $p = (3|A|/\alpha)^\alpha$. This yields

$$h(X^j) \leq \alpha \log \frac{3|A|}{\alpha} + \log(|A|)\varepsilon.$$

□

5.3. Constructing the shift Y . Fix $\varepsilon > 0$. Recall that Q is the \mathcal{T} -quasi-shift of (R, ε) -compatible \mathcal{T} -quasi-tilings.

Let $Y^0 \subset Q \times A^G$ be a shift defined as follows: if $(T, y) \in Y$ then for every tile-translate hT_i in T it holds that $y(g) \neq 0$ for at most one $g \in hT_i$. For g outside the tile-translates $y(g)$ can take any value.

Let $Y \subset A^G$ be the projection of Y^0 on its second coordinate. We can bound the size of the support of Y_F by the number of tile-translates fully contained in a finite set F , plus the size of F 's boundary $\partial_{r(\mathcal{R})}F$. Since each tile is of size at least $1/\varepsilon$, the total number of tiles is at most $\varepsilon|F|$. If we choose F so that $\rho_{r(\mathcal{R})}F$ is small enough, then the size of $\partial_{r(\mathcal{R})}F$ will be at most $\varepsilon|F|$. It follows that the support of Y_F is of size at most $2\varepsilon|F|$. Hence, by Lemma A.2 we have that

$$(8) \quad h(Y) \leq 2\varepsilon \log \frac{|A|}{\varepsilon}.$$

5.4. Realizing entropies. Let $\varphi: A^G \times A^G \rightarrow A^G$ be given by

$$[\varphi(x, y)](g) = \begin{cases} x(g) & \text{if } x(g) \neq 0 \\ y(g) & \text{otherwise.} \end{cases}$$

Claim 5.5. *If $j/|R_1| > 3\varepsilon$ then X^j is a subshift of $\varphi(X^{j-1} \times Y)$.*

Proof. Let $x^j \in X^j$. Choose a $T \in Q$. Let $x^{j-1} \in A^G$ vanish outside the T tile-translates. Inside each T -tile-translate hT_i let x^j coincide with x^{j-1} , except for at a single point $g \in hT_i$ in which $x^j(g) \neq 0$ (assuming one exists), where we set $x^{j-1}(g) = 0$. Let L be the set of these locations g .

Let $y \in A^G$ be given by $y(g) = x^j(g)$ outside the T -tile-translates, and let $y(g) = 0$ inside the tile-translates, except for $g \in L$, where we set $y(g) = x^j(g)$. Hence $y \in Y$, since L intersects each T -tile-translate in at most one element. It is immediate that $\varphi(x^{j-1}, y) = x^j$. The claim will thus be proved if we show that $x^{j-1} \in X^{j-1}$.

To see that $x^{j-1} \in X^{j-1}$, we fix a an \mathcal{R} -tile-translate hR_i , and consider two cases. First, if the support of $x^j_{hR_i}$ is of size less than $j|R_i|/|R_1|$ (i.e., the maximum it can be in X^j), then $x^j_{hR_i} \in X^{j-1}_{hR_i}$, and hence $x^{j-1}_{hR_i} \in X^{j-1}_{hR_i}$, since its support is at most that of $x^j_{hR_i}$.

Second, consider the case that the support of $x^j_{hR_i}$ is of size $j|R_i|/|R_1|$, and so it is larger than $3\varepsilon|R_i|$. Each \mathcal{R} -tile-translate hR_i is well covered by T , since T is $(\mathcal{R}, \varepsilon)$ -compatible; in particular, $e(T, hR_i) \leq \varepsilon|R_i|$. Since $\rho_{r(\mathcal{T})}(R_i) \leq \varepsilon$, it follows that the union of the T -tile-translates that are *fully contained* in hR_i is at least of size $(1 - 2\varepsilon)|R_i|$. Since the support of $x^j_{hR_i}$ is of size $3\varepsilon|R_i|$ it follows that at least one

T -tile-translate has an element g such that $x^j(g) \neq 0$. This will be removed in x^{j-1} , and so the size of the support of $x_{hR_i}^{j-1}$ is at most $j|R_i|/|R_1| - 1 \leq (j-1)|R_i|/|R_1|$. Hence $x^{j-1} \in X^{j-1}$. \square

It follows that if $j/|R_1| > 3\varepsilon$ then

$$h(X^j) - h(X^{j-1}) \leq h(Y) \leq 2\varepsilon \log \frac{|A|}{\varepsilon},$$

where the second inequality is 8.

It follows from Claim 5.4 that if $j/|R_1| \leq 3\varepsilon$ then

$$h(X^j) - h(X^{j-1}) \leq h(X^j) \leq 3\varepsilon \log \frac{|A|}{\varepsilon} + \log(|A|)\varepsilon.$$

Hence, if we set

$$\delta(\varepsilon) = 3\varepsilon \log \frac{|A|}{\varepsilon} + \log(|A|)\varepsilon$$

then we have shown that for all $0 \leq j \leq \ell$

$$h(X^j) - h(X^{j-1}) \leq \delta(\varepsilon),$$

thus proving Theorem 1.11.

APPENDIX A. COMBINATORIAL LEMMATA FOR BOUNDING ENTROPY

In this appendix we prove two combinatorial lemmata that are useful for bounding entropy.

Lemma A.1. *Let $X \in A^G$ be a shift, and let 0 be a distinguished element of A . Suppose that for some finite $F \subset G$ it holds that for every $x \in X$ the projection x_F vanishes (i.e., equals 0) on all but at most $\varepsilon|F|$ of the elements of F . Then the size of the projection of X on F can be bounded by*

$$|X_F| \leq \left(\frac{3|A|}{\varepsilon} \right)^{\varepsilon|F|}.$$

Proof. Without loss of generality, we may assume that $\varepsilon|F|$ is an integer. Let $f: F \rightarrow A$ be the projection of $x \in X$ to F . To choose f , we can first choose an $H \subset F$ of size $\varepsilon|F|$ which contains the support of f . Given such an H there are $|A|^{\varepsilon|F|}$ functions f supported on H . Since there are $\binom{|F|}{\varepsilon|F|}$ different choices of H , it follows that there are at most $\binom{|F|}{\varepsilon|F|} |A|^{\varepsilon|F|}$ distinct functions.

A standard bound on binomial coefficients which follows from Stirling's approximation is

$$\binom{n}{k} \leq \left(\frac{ne}{k}\right)^k < \left(\frac{3n}{k}\right)^k.$$

It follows that there are at most

$$\binom{|F|}{\varepsilon|F|} A^{\varepsilon|F|} \leq \left(\frac{3|F|}{\varepsilon|F|}\right)^{\varepsilon|F|} A^{\varepsilon|F|} = \left(\frac{3|A|}{\varepsilon}\right)^{\varepsilon|F|}$$

possible projections of X on F . \square

Lemma A.2. *Let $X \in A^G$ be a shift, and let 0 be a distinguished element of A . Suppose that for some $\delta > 0$ and every finite, $(1, \delta)$ -invariant $F \subset G$, it holds that for every $x \in X$ the projection x_F vanishes (i.e., equals 0) on all but at most $\varepsilon|F|$ of the elements of F . Then $h(X) \leq \varepsilon \log \frac{3|A|}{\varepsilon}$.*

Proof. Let $F \subset G$ be finite and $(1, \delta)$ -invariant. By Lemma A.1

$$|X_F| \leq \left(\frac{3|A|}{\varepsilon}\right)^{\varepsilon|F|}.$$

Since this holds for all $(1, \delta)$ -invariant F , it follows that

$$h(X) \leq \varepsilon \log \frac{3|A|}{\varepsilon}.$$

\square

APPENDIX B. SEMI-CONTINUITY OF THE ENTROPY FUNCTION

In this appendix we prove Proposition 1.3.

Proposition 1.3. *The entropy map $h: \mathfrak{S} \rightarrow \mathbb{R}^+$ is upper semi-continuous.*

Our proof uses ideas which are similar to those used by Lindenstrauss and Weiss [7, Appendix 6]

B.1. Bounding entropy using quasi-tilings. Recall that given a shift X and a finite subset $K \in G$, X_K is the projection of X onto K . We first state an easy lemma about the size of this set.

Lemma B.1. *Let X be a shift. If K_1, \dots, K_n are finite sets and $K = \bigcup_{i \in \{1, \dots, n\}} K_i$ then*

$$|X_K| \leq \prod_{i=1}^n |X_{K_i}|$$

Proof. By induction it suffices to prove the lemma when $n = 2$ which we will do. Let $K = K_1 \cup K_2$. Then any element of X_K can be projected to an element of X_{K_1} and also into an element of X_{K_2} , by restriction. We can thus map an element of X_K to the Cartesian product of its mappings on X_{K_1} and X_{K_2} . Since this mapping is injective, the claim follows. \square

We will need two simple corollaries of this lemma before proving the main proposition.

Corollary B.2. *Let $X \subseteq A^G$ be a shift, let K be finite subset of G and H a subset of K with $p|K|$ elements, for some $0 \leq p \leq 1$. Then*

$$|X_K| \leq |X_H| \cdot |A|^{(1-p)|K|}$$

Proof. There are $(1-p)|K|$ elements of $K \setminus H$. Since there are $|A|^{(1-p)|K|}$ colorings of $K \setminus H$ in the full shift, there are at most that many colorings of $K \setminus H$ in X_K . The corollary then follows by applying lemma B.1 to the sets $K_1 = H$ and $K_2 = K \setminus H$. \square

Our second corollary is an immediate consequence Lemma B.1 to sets satisfying certain conditions; these conditions will later be satisfied by tile-translates.

Corollary B.3. *Let $X \subseteq A^G$ be a shift and suppose that for all elements of some collection K_1, K_2, \dots, K_n of disjoint, finite subsets of G , and for some constant $p > 0$, we have $|X_{K_i}| \leq p^{|K_i|}$ for $1 \leq i \leq n$. Let $K = \bigcup_i K_i$. Then*

$$|X_K| \leq p^{|K|}.$$

We are now ready to prove that the entropy of a shift can be controlled, given the size of its projections to tiles of a good tile set. We will use this proposition in this section, as well as in Section 5.

Proposition B.4. *Let $X \subseteq A^G$ be a shift. Let \mathcal{T} be a $\frac{1}{2}\varepsilon$ -good tile set. Suppose that $|X_{T_i}| \leq p^{|T_i|}$ for some constant $p > 0$ and every tile $T_i \in \mathcal{T}$. Then*

$$h(X) \leq \log p + \log(|A|)\varepsilon.$$

Proof. Let F be any finite, $(r(\mathcal{T}), \frac{1}{2}\varepsilon)$ -invariant subset of G . It follows from Theorem 1.18 that there exists a \mathcal{T} -quasi-tiling T with the following properties:

- (1) Each tile-translate in T is fully contained in F .
- (2) $e(T, F) \leq \varepsilon|F|$.

By (2), if we denote by $E \subseteq F$ the set of elements covered the tile-translates in T , then $|E| \geq (1 - \varepsilon)|F|$.

Fix a tiling T of F with the properties described above. Since X is shift-invariant,

$$|X_{hT(h)}| \leq p^{|hT(h)|},$$

for all tile-translates $hT(h)$ in T . Since E is the union of all the tile-translates in T , and since the sum of their sizes is $|E|$, by applying Corollary B.3 we see that

$$|X_E| \leq \exp(\log p|E|) \leq \exp(\log p|F|).$$

applying Corollary B.2 to $F = E \cup (F \setminus E)$ yields

$$\begin{aligned} |X_F| &\leq \exp(\log p|F|) \cdot |A|^{\varepsilon|F|} \\ &= \exp(\log p|F| + \log(|A|)\varepsilon|F|). \end{aligned}$$

Hence

$$\frac{1}{|F|} \log |X_F| \leq \log p + \log(|A|)\varepsilon.$$

Since this holds for all F with a small enough boundary, it follows that

$$h(X) \leq \log p + \log(|A|)\varepsilon.$$

□

B.2. Proof of semi-continuity. To prove the semi-continuity of the entropy function it suffices to show that for every sequence X^1, X^2, \dots of shifts with limit $\lim_n X^n = X$ it holds that $\limsup_n h(X^n) \leq h(X)$.

We will prove Proposition 1.3 by showing that for every $\varepsilon > 0$ there is an N large enough so that for all $n \geq N$ it holds that

$$h(X^n) \leq h(X) + \varepsilon + \log(|A|)\varepsilon.$$

Fix $\varepsilon > 0$. Let $\mathcal{T} = (T_1, \dots, T_n)$ be a δ -good tile set, with each T_i being $(1, \delta)$ -invariant, for some $\delta < \frac{1}{2}\varepsilon$ small enough so that

$$(9) \quad \frac{1}{|T_i|} \log |X_{T_i}| \leq h(X) + \varepsilon$$

for all $T_i \in \mathcal{T}$.

Let N be large enough so that, for all $n \geq N$, $X_{T_i}^n = X_{T_i}$ for all $T_i \in \mathcal{T}$. Hence

$$\frac{1}{|T_i|} \log |X_{T_i}^n| = \frac{1}{|T_i|} \log |X_{T_i}|,$$

and by (9)

$$\frac{1}{|T_i|} \log |X_{T_i}^n| \leq h(X) + \varepsilon.$$

Rearranging, we get that for every tile T_i

$$|X_{T_i}^n| \leq e^{(h(X)+\varepsilon)|T_i|},$$

and by shift-invariance, the same holds for every tile-translate hT_i :

$$|X_{hT_i}^n| \leq e^{(h(X)+\varepsilon)|hT_i|},$$

We can now apply Proposition B.4 to X^n and \mathcal{T} , by setting $p = e^{h(X)+\varepsilon}$. This yields

$$h(X^n) \leq h(X) + \varepsilon + \log(|A|)\varepsilon,$$

proving Proposition 1.3.

REFERENCES

- [1] Tullio Ceccherini-Silberstein and Michel Coornaert, *Cellular automata and groups*, Springer, 2010.
- [2] ———, *The myhill property for strongly irreducible subshifts over amenable groups*, Monatshefte für Mathematik **165** (2012), no. 2, 155–172.
- [3] Michel Coornaert, *Topological dimension and dynamical systems*, 2015. Springer, forthcoming.
- [4] Tomasz Downarowicz, Dawid Huczek, and Guohua Zhang, *Tilings of amenable groups*, Arxiv preprint arXiv:1502.02413 (2015).
- [5] Paul Halmos, *In general a measure preserving transformation is mixing*, Annals of Mathematics (1944), 786–792.
- [6] Michael Hochman, *Genericity in topological dynamics*, Ergodic Theory Dynam. Systems **28** (2008), no. 1, 125–165.
- [7] Elon Lindenstrauss and Benjamin Weiss, *Mean topological dimension*, Israel Journal of Mathematics **115** (2000), no. 1, 1–24.
- [8] Donald Ornstein and Benjamin Weiss, *Entropy and isomorphism theorems for actions of amenable groups*, Journal d’Analyse Mathématique **48** (1987), no. 1, 1–141.
- [9] Vladimir Rohlin, *A “general” measure-preserving transformation is not mixing*, Doklady Akad. Nauk SSSR (NS), 1948, pp. 349–351.
- [10] Benjamin Weiss, *Intrinsically ergodic systems*, Bulletin of the American Mathematical Society **76** (1970), no. 6, 1266–1269.

DEPARTMENT OF MATHEMATICS, MASSACHUSETTS INSTITUTE OF TECHNOLOGY, CAMBRIDGE MA 02139, USA.