Contractible subshifts

Leo Poirier Ville Salo

February 28, 2024

Abstract

We introduce the notion of a contractible subshift. This is a strengthening of the notion of strong irreducibility, where we require that the gluings are given by a block map. We show that a subshift is a retract of a full shift if and only if it is a contractible SFT with a fixed point. For virtually polycyclic groups, contractibility implies dense periodic points. We introduce a "homotopy theory" framework for working with this notion, and "contractibility" is in fact simply an analog of the usual contractibility in algebraic topology. We also explore the symbolic dynamical analogs of homotopy equivalence and equiconnectedness of subshifts. Contractibility is implied by the map extension property of Meyerovitch, and among SFTs, it implies the finite extension property of Briceño, McGoff and Pavlov. We include thorough comparisons with these classes. We also encounter some new group-geometric notions, in particular a periodic variant of Gromov's asymptotic dimension of a group.

1 Introduction

If G is a group, Σ a finite set (alphabet) and $X \subset \Sigma^G$ is topologically closed, and closed under the shift maps induced by the regular action (the formula $(g \cdot x)_h = x_{g^{-1}h}$ where $g, h \in G, x \in X$) then X is a *subshift*. Subshifts are the main object of study in the field of symbolic dynamics.

An important notion in symbolic dynamics is *gluing*. Suppose X is a subshift, $x, y \in X$, and $A, B \subset G$. Then a *gluing* of x, y on the respective areas A, B is $z \in X$ such that $z|A = x|A$ and $z|B = y|B$. By a *gluing notion* we understand a property of a subshift, which states the existence of gluings under some conditions.

The classical notion of topological transitivity of a $\mathbb{Z}$-system can be seen as a gluing notion. Recall that the definition of this is that if U, V are nonempty open sets, then we can always find $x \in U$ and $n \geq 1$ such that $\sigma^n(x) \in V$ (where the action of n is written as σ^n). When we consider a $\mathbb{Z}$-subshift X, topological transitivity can be seen as a gluing property, as it is equivalent to the following statement: for any finite $A, B \Subset G$ and $x, y \in X$, there exists $n \geq 1$ such that x and $\sigma^n(y)$ can be glued on the areas A and $B + n$. (Note that with our definition, σ shifts symbols in the positive direction.)

Topological transitivity is too weak of an assumption for many symbolic dynamical purposes, especially on groups other than $\mathbb{Z}$. A much stronger gluing property is *strong irreducibility* (see e.g. [10]). It states that a gluing can be performed for any two points x, y whenever the minimal distance between an element of A and an element of B is bounded from below by some constant (depending on the subshift).

Strong irreducibility is perhaps the strongest possible "pure" gluing notion. Yet, it is not strong enough for many purposes. In particular, it is an open problem whether $\mathbb{Z}^d$-subshifts of finite type (see Section 2 for the definition) have dense periodic points for $d \geq 3$, and it was recently shown that general strongly irreducible $\mathbb{Z}^d$-subshifts do *not* always have dense periodic points for $d \geq 2$ [11].

Thus, it makes sense to try to strengthen gluing notions by adding other ingredients. A natural possibility is to require that instead of gluings simply existing, we have some control on the gluing. For example, Bowen's classical notion of specification for one-dimensional systems [4] includes (slightly rephrasing) the requirement that if the areas A, B are periodic, then the gluing can also be made periodic.

One situation where we want very fine control on the gluing is when constructing block maps to a subshift (for example, automorphisms). In this case, it is typically not useful to know that a gluing exists, but it should have some predictable form – in fact, it should be produced by a block map itself. This is the approach taken in the present paper.

1.1 Contractibility

In the present paper, we consider the notion of strong irreducibility with the requirement that the gluing is produced by a block map. This means that we should have a shift-commuting continuous map h which is given two points x, y from the subshift X, and another "time" configuration t that designates areas where we glue from the first point and areas where we glue from the second point. Given this data, it produces a gluing that is valid sufficiently deep inside the given areas.

More precisely, letting I denote the full shift $\{0,1\}^G$, we should have a *block map* (meaning shift-commuting continuous map)

$$h : I \times X \times X \to X$$

such that $h(t, x, y)$ copies the symbol from x (resp. y) at $a \in G$ if t contains a large patch of 0s (resp. 1s) near position a. Of course, by compactness of the Cantor space $I \times X \times X$, this is precisely equivalent to the requirements $h(\bar{0}, x, y) = x$, $h(\bar{1}, x, y) = y$, where $\bar{a}$ is the constant-a configuration.

Note that $h(\bar{0}, x, y) = x$, $h(\bar{1}, x, y) = y$ is analogous to the usual definition of a homotopy in algebraic topology between two projection maps, except the interval is replaced by a full shift, and in addition to continuity we require shift-commutation.

Definition 1.1. *If $f, g : X \to Y$ are morphisms (= block maps), a* homotopy *between f and g, denoted $h : f \cong g$, is a block map $h : I \times X \to Y$ satisfying $h(\bar{0}, x) = f(x)$ and $h(\bar{1}, x) = g(x)$. We then say f and g are* homotopic*.*

We can then mimic (one of the standard definitions of) the notion of contractibility from algebraic topology.

Definition 1.2. *A subshift $X \subset \Sigma^G$ is* contractible *if there is a homotopy between the two projections $\pi_1, \pi_2 : X \times X \to X$.*

We show in Section 2.1 that this corresponds to the standard notion of contractibility in algebraic topology.

Unwrapping the definition, we see that this means precisely what we wrote above: a subshift X is contractible if there is a continuous, shift-commuting map $h : I \times X \times X \to X$ such that $h(\bar{0}, x, y) = x$, and $h(\bar{1}, x, y) = y$. We call such h a *contraction homotopy*, though it is also natural to see it as a *gluing morphism.*

1.2 Results

In this section, we highlight some of the results of the paper. First, in the most classical case of one-dimensional SFTs, we have the expected result:

Theorem 1.3. *A one-dimensional subshift of finite type is topologically mixing if and only if it is contractible.*

This is proved in Theorem 5.18.

The following shows that contractible subshifts (even without the SFT assumption) avoid some pathologies of strongly irreducible subshifts.

Theorem 1.4. *Let G be an infinite finitely-generated residually finite group with finite periodic asymptotic dimension. Then every contractible G-subshift has dense periodic points.*

This is proved in Theorem 6.10.

Periodic asymptotic dimension is a new notion that we introduce. We show in Section 9.1 that strongly polycyclic groups (such as the group $\mathbb{Z}^d$), the lamplighter groups $\mathbb{Z}_p \wr \mathbb{Z}$, and the metabelian Baumslag-Solitar groups $BS(1, n)$, have finite periodic asymptotic dimension (and they are of course infinite finitely-generated residually finite groups).

In Section 4 we characterize retracts of full shifts in terms of contractibility, showing that contractible subshifts (of finite type, with a fixed point) arise naturally in the category of subshifts of finite type on any group (subshifts form a category in a natural way by taking morphisms to be block maps [16]).

Theorem 1.5. *On every group, a subshift X is a retract of a full shift if and only if it has a fixed point, is of finite type, and is contractible.*

This is proved in Theorem 4.5. An analogous result is more generally true when the full shift is replaced by an arbitrary contractible subshift Y, though the existence of a fixed point is replaced by the assumption that a map exists from Y to X.

While our notion of contractibility is stated as a strengthening of strong irreducibility, its connection to the finite extension property (FEP) of [6] seems much tighter. This is not particularly surprising, as the FEP was also introduced as a requirement on the target subshift of a block map.

Theorem 1.6. *On every group, every contractible SFT has the FEP.*

This is proved in Theorem 8.5. We show by examples in Section 8 that on many groups, contractible SFTs are strictly contained in the class of FEP subshifts, which in turn are a strict subclass of factors of contractible SFTs.

The FEP class was introduced in [6] as the codomain condition in a general result about factoring SFTs onto smaller entropy SFTs. In Theorem 8.8, we prove the following weaker result, which applies on any group, and has a very simple proof.

Theorem 1.7. *Suppose $Y \subset A^G$ is a contractible SFT. Then the following are equivalent:*

- *there is a block map from $f : X \to Y$ and a block map $g : X \to A^G$ such that $g(X) \supset Y$,*
- *there is a factor map $f : X \to Y$.*

Here, the first item lists the obvious necessary conditions for the existence of a factor map. It is not difficult to recover the special case of the result of [6] for contractible SFTs from this theorem.

Contractibility and FEP are incomparable properties, and we do not know a natural common generalization, but there are many connections. In particular in Theorem 6.13, we prove the analog of Theorem 1.4 for FEP subshifts, with essentially the same proof.

Another class of subshifts, partially motivated by similar concerns as our notion, called subshifts with the map extension property, was recently introduced by Meyerovitch in [14]. A variant of Theorem 1.7 is proved there as well.

Our property is weaker than that of Meyerovitch. Specifically, the definition of the map extension property in [14], which we recall in Section 8.4, consists of two somewhat orthogonal items. The first item is related to retracts/gluing, and the second to periodic points. We have nothing to say about periodic points, but the first item indeed directly corresponds to our class:

Theorem 1.8. *On $\mathbb{Z}^d$, contractible SFTs are precisely the subshifts satisfying the first item in the definition of the map extension property stated in [14].*

This is proved in Theorem 8.22. We also give a characterization of this property on a general group in Theorem 8.21, using an auxiliary notion we call the map existence property.

The map extension property introduced in [14] arises precisely from the study of retracts, so in light of Theorem 1.5, it is not surprising that the notions are strongly connected.

There are many other things one can do with the symbolic dynamical homotopy notion in Definition 1.1, in particular we study (the analog of) equiconnectedness, meaning that there is a homotopy between the two projections of $X \times X$ which fixes the diagonal.

Theorem 1.9. *Suppose the group G has the patching property. Then all equiconnected subshifts are contractible SFTs. If the group is $\mathbb{Z}^d$, or the subshift has a fixed point, the converse holds.*

The first direction is proved in Theorem 7.4. The second is proved in Theorem 8.22 ($\mathbb{Z}^d$) and Theorem 7.6 (subshifts with fixed points).

The patching property is a new technical notion that we introduce. We show that it is implied by finite asymptotic dimension, and by subexponential growth, in Section 9.2. Thus it includes the groups listed above, but also for example nonabelian free groups, and the Grigorchuk group. We do not have any example of a group that does not have the patching property.

We also take a look at homotopy equivalence, and show among other things the following:

Theorem 1.10. *Two one-dimensional transitive SFTs are homotopy equivalent if and only if they map into each other.*

Theorem 1.11. *A subshift of finite type is homotopy-equivalent to a point if and only if it is contractible and has a fixed-point.*

The first is proved in Theorem 5.19, and the latter in Theorem 5.12. (Note that this is the same condition that characterizes the retracts of full shifts.)

In the important special case $G = \mathbb{Z}^d$, we summarize the equivalence results obtained above, first in the case of SFTs with a fixed point, and then in general:

Theorem 1.12. *Let $G = \mathbb{Z}^d$, and let X be a subshift of finite type with a fixed point. Then TFAE:*

- *X is contractible;*
- *X is a retract of a full shift;*
- *X is a retract of a contractible subshift of finite type;*
- *X is equiconnected;*
- *X satisfies the first item of the definition of the map extension property in [14];*
- *X is homotopy-equivalent to a point.*

Theorem 1.13. *Let $G = \mathbb{Z}^d$, and let X be a subshift. Then TFAE:*

- *X is a contractible SFT;*
- *X is a retract of a contractible subshift of finite type;*
- *X is equiconnected;*
- *X satisfies the first item of the definition of the map extension property in [14].*

Besides specific results, we give many examples of contractible subshifts and general constructions throughout the paper.

2 Preliminaries

2.1 Contractibility in algebraic topology

Our definition of contractibility indeed corresponds to the usual notion in the setting of topological spaces. The most common definition of contractibility in algebraic topology is likely that a topological space is contractible if there is a homotopy between the identity map on X, and the map from X to one of its points. We show below that in the topological setting, this is equivalent to the definition we use in this paper.

Lemma 2.1. *Let X be a topogical space. Then the following are equivalent:*

- *there is a homotopy between the identity map on X, and the constant map from X to one of its points;*
- *the two projections from $X \times X$ to X are homotopic.*

Proof. Suppose there is a homotopy between the identity map on X, and the map from X to one of its points, say $h : \mathrm{id}_X \cong (x \mapsto x_0)$ for some $x_0 \in X$. Then

$$h'(t, x, y) = \begin{cases} h(2t, x) & \text{if } t \leq 1/2 \\ h(2 - 2t, y) & \text{if } t \geq 1/2 \end{cases}$$

is a homotopy between the two projections. Conversely if h' is a homotopy between the two projections, then

$$h(t, x) = h'(t, x, x_0)$$

gives a homotopy between the identity map on X, and the map from X to $\{x_0\}$. $\square$

In our symbolic dynamical category, these notions are not equivalent, the basic issue being that a morphism cannot talk about any individual point unless there is a fixed point. We will later see, however, that these definitions are equivalent for subshifts of finite type containing at least one fixed point. We find the second definition more useful in our symbolic dynamical category, so we use that as the main definition.

We sometimes say that a subshift is *fixed-point contractible* if the analog of more common definition of contractibility holds, i.e. the identity map is homotopic to the map to some one-point subsystem $\{\bar{a}\}$. We will see later that this is equivalent to being contractible and having a fixed point, for subshifts of finite type (Corollary 5.10).

2.2 General and symbolic dynamical definitions

If $f : A \to B$ is a function and $C \subset A$, then $f|C$ is the domain restriction of f to C. By $A \subset B$ we mean A is a not necessarily proper subset of B. By $\mathbb{Z}_k$ we denote the finite ring or additive group $\mathbb{Z}/k\mathbb{Z}$. We write $A \subset B \iff A \subseteq B$.

Throughout, G denotes a group. Elements of a group are usually called a, b, c (because g and h will be morphisms). We mostly consider infinite finitely-generated groups with a preferred generating set S. We usually (but not always) let S be a finite symmetric ($S = S^{-1}$) generating set. We write $d(a, b)$ for the left-invariant word metric of G with respect to this generating set.

The *full shift* is the set A^G (functions from G to A), for a finite set A called the *alphabet*, whose elements are called *symbols* or sometimes *colors*. Its elements $x \in A^G$ are called *configurations* and we write x_a instead of $x(a)$, but we use the same restriction notation as for functions. A full shift A^G is *non-trivial* if $|A| \geq 2$.

The full shift is a topological dynamical system when A^G is given the product topology (making it a Cantor set if $|A| \geq 2$). G acts by $ax_b = x_{a^{-1}b}$ for $x \in A^G$ and $a, b \in G$. Note that this shift-convention means that one should picture the Cayley graph of G as having edges (a, ab) for $b \in S$ (i.e. the right Cayley graph).

The shift map by a translates the identity 1_G to a, and translates everything else by the unique graph automorphism of the S-labeled Cayley graph that this determines.

If $D \subset G$, and $p \in A^D$, then p is a *pattern*. A pattern is *finite* if its domain D is finite, which we denote by $D \Subset G$. We translate patterns as follows: if $p \in A^D$ and $a \in G$, then $ap \in A^{aD}$ and $ap_b = p_{a^{-1}b}$. Note that configurations are patterns with a full domain, and their translation agrees with the translation of patterns. If $x \in A^G$ then our syntactic convention is that translation happens *before* restriction, so $ax|N = (ax)|N$. We note the useful formula $a^{-1}(x|aN) = a^{-1}x|N$.

If p, q are patterns, then $q \sqsubset p$ where $q \in A^D$, $p \in A^E$, means $p|aD = aq$ for some $a \in G$ such that $aD \subset E$. We say q *appears* in p. We extend this naturally to the case where $p \in A^E$ for possibly infinite $E \subset G$ (e.g. a p a configuration in some subshift).

We extend this to the case where q is a single symbol s, which we can identify with the pattern $1_G \mapsto s$. We write $q \subset p$ if q is a subpattern of p as a positioned pattern, meaning $p|D = q$ where q has domain D. If $X \subset A^D$ is a set of patterns (possibly with infinite D), we say p *appears* in X, or X *contains* p, if p appears in x for some $x \in X$.

The *cylinder* corresponding to a finite pattern p is $[p] = \{x \in A^G \mid \forall a \in D : x_a = p_a\}$. The topology of A^G has a basis of clopen sets. The clopen sets are precisely the finite unions of *cylinders*.

A *subshift* is a topologically closed and G-closed subset X of a full shift. Equivalently, it is a subset of A^G defined as the set of points $x \in A^G$ whose *orbit* Gx does not intersect any cylinder in a countable (possibly finite) set $\{[p_1], [p_2], \ldots\}$, where the p_i are finite patterns called *forbidden patterns* of X. This is equivalent to saying that X is the set of points in A^G which do not contain any of the forbidden patterns.

If we can pick a finite set of forbidden patterns, then X is a *subshift of finite type* or *SFT* for shift. We may always suppose the forbidden patterns of an SFT have the same domain D, which we call a *window* for the SFT. Note that by our shift convention, forbidden patterns are "checked on the right".

A subshift $X \subset \Sigma^G$ is *strongly irreducible* if there exists R such that for all $A, B \subset G$ such that $d(A, B) = \min_{a \in A, b \in B} d(a, b) \geq R$, and for all $x, y \in X$, there exists $z \in X$ such that $z|A = x|A, z|B = y|B$. We say X is *(topologically) mixing* if for all $A, B \Subset G$ (finite subsets!) there exists R such that whenever $x, y \in X$, and $a \in G$ is such that $d(A, aB) \geq R$, there exists $z \in X$ such that $z|A = x|A, z|aB = y|aB$.

Especially when we have homotopy in mind, we denote by I_n the full shift on symbols $\{0, 1, \ldots, n-1\}$. The binary full shift $I = I_2$ plays the role of the interval, and we often refer to its points as *time parameters*. For a symbol i, write $\bar{i}$ or sometimes i^G for the unique point in $\{i\}^G$. For a function $p : D \to A$ we also use the usual notation $p \equiv i$ for $\forall a \in D : p_a = i$, where $i \in A$. So $\bar{i}$ is the unique configuration x with $x \equiv i$. Such configurations are often called *unary*. We also use such terminology and notation for patterns, writing i^D for the unique element for $\{i\}^D$.

We work in the category of subshifts on G, and a *morphism* or *block map* is a shift-commuting continuous function between two subshifts. Isomorphisms, embeddings and factor relations are naturally interpreted with these morphisms, as bijective, injective and surjective morphisms, respectively. The image of a surjective morphism is also called a *factor*, and a preimage in a surjective morphism is a *cover*. Isomorphisms are often called *conjugacies*. Almost all of our definitions and notions are preserved under conjugacy (with some obvious exceptions like safe symbols and the Z0 property).

A morphism $f : X \to Y$ between two subshifts $X \subset A^G, Y \subset B^G$ has a *neighborhood* $N \Subset G$ and a *local rule* $f_{\text{loc}} : A^N \to B$, such that $f(x)_a = f_{\text{loc}}(a^{-1}(x|aN))$ (or equivalently $f(x)_a = f_{\text{loc}}(a^{-1}x|N)))$. We usually deal with only finitely-generated groups with a preferred generating set S, and then we also say r is a *radius* for f it the radius-r ball $B_r = S^{\leq r} = \{g \mid d(1_G, g) \leq r\}$ is a neighborhood for it.

If $X \subset A^H$ is a subshift and $\pi : G \to H$ is a group epimorphism, then the *pullback* (of X with respect to π) is the subshift $Y \subset A^G$ with

$$y \in Y \iff \exists x \in X : \forall a \in G : y_a = x_{\pi(a)}.$$

If $H \leq G$, then corresponding the *free extension* (a.k.a. *induction*) of a subshift $X \subset A^H$ is the subshift $Y \subset A^G$ with the same forbidden patterns. We write it as $Y = X^{G/H}$ Equivalently, configurations of Y consist of independently chosen configurations of X on different cosets of H.

If $X \subset A^G$ is a subshift, the *SFT approximation* of X with *window* $W \Subset G$ is the SFT defined by forbidden patterns $\{p \in A^W \mid [p] \cap X = \emptyset\}$.

If $A \times B$ is a (Cartesian) product alphabet, we think of subshifts $X \subset (A \times B)^G$ also as subshifts of $A^G \times B^G$ with the diagonal action. We often refer to the different projections of subshifts (and configurations) as *tracks*.

A *vertex shift* is a subshift of $A^{\mathbb{Z}}$ defined by forbidden patterns with domain $\{0, 1\}$. In the one-dimensional case (i.e. the group $\mathbb{Z}$) we also call patterns whose domain is an interval $\{0, 1, \dots, k-1\}$ *words*, and the *length* of a word $w \in A^k = A^{\{0,1,\dots,k-1\}}$ is $|w| = k$. We identify them with the free monoid on the alphabet. We can *concatenate* two words $u \in A^k, v \in A^m$ in an obvious way to a word $uv \in A^{k+m}$. For A an alphabet, write A^* for the set of all words over this alphabet, including the unique word of length 0.

2.3 A few words on groups

We mostly assume the reader is familiar with group theory, but we recall some definitions and a basic lemma.

A *polycyclic group* is G such that there is a sequence of subgroups $1 = G_0 < G_1 < \dots < G_n = G$ such that each G_i is normal in G_{i+1} and each quotient G_{i+1}/G_i is cyclic. If the quotients are isomorphic to $\mathbb{Z}$, the group is *strongly polycyclic*.

A group is *just-infinite* if it is infinite and has no non-trivial normal subgroup of infinite index. Just-infinite groups are an important and interesting class of groups, see [2]. Nevertheless, arguably most groups are not just-infinite. We mention only the trivial fact that infinite strongly polycyclic groups are never just-infinite (except for $\mathbb{Z}$).

We write $1 \to K \to G \to H \to 1$ (an *exact sequence*) and say G is a *group extension* is there is a surjective group homomorphism $\pi : G \to H$, such that K is the kernel of this homomorphism. If π has a section $\gamma : H \to G$ meaning γ is a homomorphism and $\pi \circ \gamma = \mathrm{id}_H$, then we say π (or the group extension) *splits*, and then G is a semidirect product $K \rtimes H$.

We sometimes need long or infinite geodesics. We recall the simple proof that these exist universally.

Lemma 2.2. *Let G be infinite and finitely-generated by a symmetric set $S \Subset G$. Then there exists $p : \mathbb{Z} \to G$ such that $d(p(i), p(j)) = |j - i|$ for all $i < j$.*

Proof. By Kőnig's lemma we have $p : \mathbb{N} \to G$ with this property, by considering optimal paths to the sphere S^r for increasing r. Then for each n, reparametrize the path to $p_n(j) : [-n, \infty) \to G$ by $p_n(j) = p(n)^{-1}p(j + n)$, extend each p_n arbitrarily to domain $\mathbb{Z}$, and take any limit point of the p_n in the product topology. □

2.4 Two lemmas about SFTs

Lemma 2.3. *The following are equivalent for a subshift $Y \subset B^G$:*

1. *Y is SFT;*
2. *for any local rule $f_{\text{loc}} : A^N \to B$ defining a valid map $f : X \to Y$ for a subshift $X \subset A^G$, for all sufficiently good SFT approximations X' of X, f_{loc} also gives a well-defined map $\tilde{f} : X' \to Y$;*
3. *for any local rule $f_{\text{loc}} : B^N \to B$, for all sufficiently good SFT approximations Y' of Y, f_{loc} also gives a well-defined map $\tilde{f} : Y' \to Y$.*

Proof. For (1) to (2), let W be a window size for Y, and let X' be the SFT approximation of X with window WN or greater (with respect to inclusion).

Now when applied to configurations of X', $\tilde{f}$ sees the same WN-patterns as when applied to configurations of X. Thus, it produces the same W-patterns in the image, in particular if $f(X) \subset Y$ then it cannot produce forbidden patterns of Y.

The implication (2) $\implies$ (3) is trivial.

For (3) to (1), improve the SFT approximation Y' to Y'' so that f_{loc} does not see any forbidden patterns of Y on points in Y''. Then f must be the identity on Y'', in particular $Y = Y''$ is SFT. □

The following is a simple compactness argument.

Lemma 2.4. *Let $X \subset A^G$ be a subshift of finite type with a fixed set of forbidden patterns $P \subset A^K$. Then for all $T_1 \Subset G$, there exists $T_2 \Subset G$ such that when $p \in A^{T_2}$ contains no pattern from P, then p appears in a configuration of X.*

Remark 2.5. *We note that the previous lemma is only an existential result, in that for the groups $G = \mathbb{Z}^d$, $d \geq 2$ (and more generally a large class of one-ended groups), there is no computable function that produces T_2 from T_1 and the forbidden patterns defining X. This is because we can encode seeded Turing machine computation into T_1-patterns, and the size of the corresponding T_2 roughly corresponds to the busy beaver function from computability theory.*

2.5 The natural extension of a contraction homotopy

A very common modification we need to perform on a contraction homotopy is to extend its local rule, so that it can also be applied in invalid contexts. We do this in a specific way.

Definition 2.6. *Let $h : I \times X \times X \to X$ be a contraction homotopy, where $X \subset A^G$. Then a* natural extension *is any $\tilde{h} : I \times A^G \times A^G \to A^G$ with the same neighborhood as h such that $\tilde{h}(\bar{0}, x, y) = x$, $\tilde{h}(\bar{1}, x, y) = y$ for all $x, y \in A^G$ and $\tilde{h}(t, x, y) = h(t, x, y)$ whenever $x, y \in X$.*

A natural extension always exists: simply use the same local rule when you can, extend it to copy from the second input when the first input is all 0, from the third when it is all 1, and finally extend arbitrarily to other patterns.

3 First examples of contractible SFTs

An important definition, and the simplest possible example of a homotopy, is the following.

Definition 3.1. *Let G be any group, $X = \Sigma^G$, and $f, g : X \to X$ arbitrary. For $a \in G, t \in I, x \in X$, define*

$$h(t, x)_a = \begin{cases} f(x)_a & \text{if } t_a = 0, \\ g(x)_a & \text{if } t_a = 1. \end{cases}$$

Then h is the naive homotopy *between f and g.*

More generally, we refer to any homotopy defined by this formula a naive homotopy.

An important source of strongly irreducible subshifts are safe symbols. Recall that a *safe symbol* for a subshift $X \subset A^G$ is a symbol $0 \in A$ such that for all $a \in G$ and $x \in X$ we have $y \in X$, where $y_b = x_b$ for $b \neq a$, and $y_a = 0$.

Proposition 3.2. *Every subshift of finite type which has a safe symbol is contractible.*

Proof. Let $F \Subset G$ be a window for the SFT $X \subset A^G$ and suppose $0 \in A$ is a safe symbol. Pick any $M \Subset G$ such that $M \supset F^{-1}$ and define $h : I \times X \times X \to X$ by

$$h(t, x, y)_a = \begin{cases} 0 & \text{if } t_{aM} \notin \{0^{aM}, 1^{aM}\} \\ x_a & \text{if } t_{aM} = 0^{aM} \\ y_a & \text{if } t_{aM} = 1^{aM}. \end{cases}$$

If the image of h contains a forbidden pattern for some input, by shifting we find $t \in I, x, y \in X$ such that $h(t, x, y)|F$ is forbidden in X. Then the two latter cases must be used in the definition of h when computing $h(t, x, y)|F$, since otherwise $h(t, x, y)|F$ is obtained from $x|F$ or $y|F$ by inserting safe symbols and thus cannot be forbidden. This means that $t|aM = 0^{aM}$ for some $a \in F$, and $t|bM = 1^{bM}$ for some $b \in F$. This is impossible because $1_G \in aM \cap bM$ so we would have $0 = t_{1_G} = 1$. $\square$

Without the assumption of finite type, safe symbols do not imply contractibility. We later discuss the subshift P_2, which is a union of two full shifts sharing a unary point $\bar{1}$. It has safe symbol 1, but is not contractible. See also the discussion after Corollary 5.10.

In particular, every *hereditary subshift*, meaning one where the alphabet A is totally ordered and any symbol can be replaced by a smaller one, in any valid configuration, without introducing a forbidden pattern, has safe symbol $\min A$. A concrete example (which is also SFT so the previous proposition applies) is the golden mean shift in one dimension, i.e. the set $X \subset \{0,1\}^{\mathbb{Z}}$ of configurations $x \in \{0,1\}^{\mathbb{Z}}$ such that $11 \not\sqsubset x$.

One can produce many examples without safe symbols by conjugating a subshift with a safe symbol to an isomorphic one without a safe symbol.

The following proposition gives an example of a contractible subshift of finite type that is not produced this way (though we do not prove that it cannot be conjugated to have a safe symbol). This example is from [8], and it was the first example of a strongly irreducible subshift of finite type with multiple measures of maximal entropy. We thus state it as a proposition (although it gives examples of contractible SFTs on any group).

Proposition 3.3. *On $\mathbb{Z}^d$, $d \geq 2$, there exist contractible subshifts of finite type with no safe symbol and with multiple measures of maximal entropy.*

Thus, contractibility (even with the additional assumption of SFTness) does not remove the a priori "pathological" possibility of multiple measures of maximal entropy. On the other hand, the example from [8] is extremely natural, so including it is difficult to see it as a shortcoming of our definition.

Proof. Let G be a group with symmetric generating set $S \Subset G \setminus \{1_G\}$. Consider the G-SFT X over alphabet $\Sigma = \{-M, \dots, -1, 1, \dots, M\}$ where for each $s \in S$, we forbid $p \in \Sigma^{\{1_G, s\}}$ when $\prod_i p_i \leq -2$. In other words, adjacent symbols cannot have different signs, unless both are ± 1.

For $G = \mathbb{Z}^d$, this is Example 1.5 of [8]. They show that if M is large enough and $d \geq 2$, there are multiple measures of maximal entropy for X.

By definition, X is a subshift of finite type. It has no safe symbol (for $M \geq 2$), as the symbol $s \in \Sigma$ cannot be positioned next to $\mathrm{sign}(s)2$.

The subshift X is contractible. Namely, given $x, y \in X$ and $t \in I$, we define

$$h(t,x,y)_a = \begin{cases} x_a & \text{if } t|aS = 0^{aS} \\ y_a & \text{if } t|aS = 1^{aS} \\ \mathrm{sign}(x_{ab}) & \text{if } t|aS \notin \{0^{aS}, 1^{aS}\} \wedge \exists b \in S : t|abS = 0^{abS} \\ \mathrm{sign}(y_{ab}) & \text{if } t|aS \notin \{0^{aS}, 1^{aS}\} \wedge \exists b \in S : t|abS = 1^{abS} \\ 1 & \text{otherwise.} \end{cases}$$

This is well-defined, since if $\exists b \in S : t|abS = 0^{abS}$ then $t_a = 0$, so the third and fourth cases do not collide. We claim that $h(t,x,y) = z$ can not contain a forbidden pattern. Otherwise, by shifting we may suppose $z_{1_G} z_a \leq -2$ for $a \in S$. Then certainly for $T = \{1_G, a\}$, $z|T \neq x|T$ and $z|T \neq y|T$, so $t|TS$ is non-unary.

If neither $t|S$ and $t|aS$ is unary, then both entries of $z|T$ have absolute value 1. We conclude that exactly one of $t|S, t|aS$ is unary. By symmetry we may assume $t|aS$ is unary, and by symmetry we may assume $t|aS = 0^{aS}$. Then $z_a = x_a$ and $z_{1_G} = \mathrm{sign}(x_a)$, so in fact $z_{1_G} z_a = |z_a| \geq 1$. This contradicts the assumption that T contains a forbidden pattern.

This concludes the proof that h is a contraction homotopy for X. $\square$

The examples above all have fixed points. There are also contractible subshifts without fixed points. The following example also appears in [14]. It is often referred to as the subshift of graph colorings, but in this paper colors refer to symbols in a general configuration, so we avoid this name.

Proposition 3.4. *Let G be a group and let $S \not\ni 1_G$ be a finite symmetric subset of G. Let $|A| \geq |S|+1$, and let $X \subset A^G$ be the subshift containing configurations $x \in A^G$ such that $x_a \neq x_{ab}$ for all $b \in S$. Then X is a contractible SFT.*

Proof. We define a contraction homotopy $h : I \times X \times X \to X$. Let $f : (A \cup \{\#\})^S \to A$ be any function such that $f(p) \not\sqsubset p$ (in the sense that the symbol $f(p)$ does not appear in p) for all $p \in (A \cup \{\#\})^S$. Let $A = \{1, 2, \dots, k\}$.

Define $h_0 : I \times X \times X \to (A \cup \{\#\})^G$ by

$$h_0(t,x,y)_a = \begin{cases} x_a & \text{if } t_a = 0, \\ y_a & \text{if } t|aS = 1^{aS}, \\ \# & \text{otherwise.} \end{cases}$$

Then for each $i \in A$ define $h_i : I \times X \times X \to (A \cup \{\#\})^G$ by

$$h_i(t,x,y)_a = \begin{cases} h_{i-1}(t,x,y)_a & \text{if } h_{i-1}(t,x,y)_a \neq \#, \\ f(a^{-1}(h_{i-1}(t,x,y)|aS)) & \text{if } h_{i-1}(t,x,y)_a = \# \wedge x_a = i, \\ \# & \text{otherwise.} \end{cases}$$

Note that each h_i will preserve all non-# symbols picked in previous stages. The image of $h = h_k$ cannot contain symbol #: if $x_a = i$, then the definition of $h_i(t,x,y)_a$ will set it to a non-# symbol unless it was already set in $h_{i-1}(t,x,y)$. The image of h is in fact contained in X, as clearly no map h_i will produce a forbidden pattern: h_0 copies all symbols from x or y, and never picks the value of two adjacent cells from different points; other h_i pick values by applying f to the pattern in the S-adjacent cells (and f picks some value that does not appear in the neighborhood), and f is never applied in consecutive cells in the same step.

Also, $h(\bar{0},x,y) = x$ and $h(\bar{1},x,y) = y$ are already clear from the definition of h_0. □

More examples of contractible subshifts are given throughout the paper. In particular in Proposition 5.14, in Lemma 6.6, and several examples in Section 8.

4 Characterization of retracts of full shifts

One natural situation where contractible subshifts appear is the characterization of retracts of full shifts.

Definition 4.1. *Let X, Y be two subshifts. We say X is a* retract *of Y if there are morphisms $f : X \to Y$ and $r : Y \to X$ such that $r \circ f = \mathrm{id}_X$.*

Note that while retracts (with this same definition) appear prominently in algebraic topology, this is also a standard existing notion in (symbolic) dynamics.

In this situation we say f is a *section* of r, and that r is the *retraction* of f. In this case one also says that f is *split monic* [16], though we do not use this terminology here. Note that any f admitting a retraction needs to be injective, thus a conjugacy to its image, and it is easy to show that once f is indeed injective, being split monic only depends on the image $f(X)$. In particular, when classifying retracts of a subshift Y, it suffices to classify retracts of Y which are literally subshifts of Y.

Lemma 4.2. *For every subgroup $H \leq G$, if Y has an H-periodic point so do its retracts.*

Proof. The retraction map $r : Y \to X$ commutes with G so $hy = y \implies r(hy) = hr(y)$. Thus if H fixes y, it also fixes $r(y) \in X$. □

Lemma 4.3. *A retract of a subshift of finite type is of finite type.*

Proof. Suppose Y is a subshift of finite type and suppose $f : X \to Y$ and $r : Y \to X$ with $r \circ f = \mathrm{id}_X$. Pick a local rule g_{loc} for $r \circ f$ (for example by composing the local rules of f and r in the obvious way). Since Y is of finite type, by Lemma 2.3 for any good enough SFT approximation X' of X, there is an extension $\tilde{f} : X' \to Y$ of f with the same local rule.

Then $r \circ \tilde{f} : X' \to X$ is an extension of id_X, and for good enough approximation uses the same local rule g_{loc}. By Lemma 2.3, X is of finite type. □

Lemma 4.4. *A retract of a contractible subshift is contractible.*

Proof. Suppose Y is contractible, and consider directly a subshift $X \subset Y$, and a retraction $r : Y \to X$ with $r|X = \mathrm{id}_X$. Let $h : I \times Y \times Y \to Y$ be the contraction homotopy. Define

$$h'(t, x, x') = r(h(t, x, x')).$$

for $(t, x, x') \in I \times X \times X$. Observe that the image of $h(t, x, x')$ is in Y, so the composition makes sense, and we can take the codomain of g' to be X. Now $h'(\bar{0}, x, x') = r(x) = x$ since h is a homotopy and r a retraction, and similarly $h'(\bar{1}, x, x') = x'$. □

Theorem 4.5. *Let $X \subset \Sigma^G$ be a subshift. Then X is a retract of Σ^G (through its natural embedding, equivalently abstractly) if and only if all of the following hold:*

- *X has a fixed-point,*
- *X is of finite type, and*
- *X is contractible.*

Proof. The previous three lemmas show that the three properties hold for retracts of Σ^G.

Now suppose the properties hold. Let $\bar{0} \in X$, and let $h : I \times X \times X \to X$ be the contraction homotopy. Consider the natural extension $\tilde{h} : I \times \Sigma^G \times \Sigma^G \to \Sigma^G$.

Fix a window $W \Subset G$ for X. Let $M \subset G$ be finite, and define $t_M : \Sigma^G \to I$ by

$$t_M(x)_a = 0 \iff x|aM \text{ is globally valid in } X.$$

Define $r_M : \Sigma^G \to \Sigma^G$ by

$$r_M(x) = \tilde{h}(t_M(x), x, \bar{0}).$$

For $x \in X$, we have $r_M(x) = \tilde{h}(\bar{0}, x, \bar{0}) = x$, so it suffices to show is that for a suitable choice of M, the codomain of r_S can be restricted to X.

Fix some M, and suppose that the codomain restriction cannot be done, i.e. the image of r_M is not always in X. Then, using the fact W is a window size for X and possibly shifting, we find $x_M \in \Sigma^G$ such that

$$\tilde{h}(t_M(x_M), x_M, \bar{0})|W \not\sqsubset X.$$

Note that since $0^W \sqsubset X$, we must have $\tilde{h}(t_M(x_M), x_M, \bar{a})_b \neq 0$ for some $b \in W$. Thus by the construction of $\tilde{h}$, $0 \sqsubset t_M(x_M)|WN$ where N is the neighborhood of $\tilde{h}$. By the definition of t_M, $x_M|bcM$ is globally valid in X for some $b \in W, c \in N$. Note that W and N are fixed sets, while M can be picked arbitrarily large, so for large M, x_M will have arbitrarily large globally valid central patterns.

From this we conclude that we can have $\tilde{h}(t, x_M, \bar{0})|W \not\sqsubset X$ with x_M with $d(x_M, X)$ arbitrarily small in the Hausdorff metric. This means $\tilde{h}$ is not uniformly continuous in $I_2 \times X \times X$. This contradicts the compactness of $I_2 \times X \times X$, continuity of $\tilde{h}$, and the fact $\tilde{h}$ is an extension of h. □

4.1 Adaptation to general contractible subshifts

It seems difficult to give a full characterization of retracts of subshifts. However, given that the characterization of retracts of full shifts is in terms of contractible subshifts, it is natural to characterize the retracts of contractible subshifts. As we see, the proof of this is essentially the same.

Definition 4.6 (Subshift-relative contractibility)**.** *Let $X \subset Y \subset A^G$ be subshifts. We say that X is* contractible relatively to Y *if there exists a block-map $h : I \times X \times X \to Y$ such that for all $x, y \in X$, $h(\bar{0}, x, y) = x$, $h(\bar{1}, x, y) = y$.*

We call h a *relative contraction homotopy.* The definition says that, given configurations in X, h only introduces patterns that are forbidden in X, but not ones forbidden in Y, and at endpoints it behaves like a standard contraction homotopy.

Theorem 4.7. *Then X is a retract of Y if the following hold:*

- *X is a relative SFT of Y,*
- *there exists a map from Y to X, and*
- *X is contractible.*

Conversely, the first two items hold whenever X is retract of Y, and the third also holds if X is additionally relatively contractible in Y.

If Y is itself contractible, then any subshift of it is relatively contractible, as the restriction of the contraction homotopy has the required properties. Thus in this case the theorem provides a full characterization of retracts, generalizing Theorem 4.5 (noting that the existence of a morphism from a full shift to X is equivalent to the existence of a fixed point).

Corollary 4.8. *If Y is contractible, then X is a retract of Y if and only if X is a relative SFT of Y, there exists a map from Y to X, and X is itself contractible.*

Proof of Theorem 4.7. Suppose first that the three properties hold. In other words, let X be a relative SFT of Y, let $f : Y \to X$ be a morphism, and let $h : I \times X \times X \to X$ be the contraction homotopy. We then construct a retract as in Theorem 4.5: we take $\tilde{h}$ the natural extension of h, and the retraction is given by $r(x) = \tilde{h}(t(x), x, f(x))$ where $t(x)$ has large areas of 1s where there are relative forbidden patterns of X nearby, and 0s elsewhere. We leave the details to the interested reader.

Conversely, if $r : Y \to X$ is a retraction, then

- X is a relative SFT of Y because it is the set of fixed points of r, and
- $r : Y \to X$ is a morphism from Y to X.

For the third item, if $h' : I \times X \times X \to Y$ is a relative contraction homotopy for X relative to Y, then $h = r \circ h'$ is a contraction homotopy for X. $\square$

5 Contractibility and homotopy equivalence

In this section, we attempt to work out some basic "algebraic topology" for subshifts, based on the idea of a homotopy. We obtain the useful fact that in the case of contractible SFTs, homotopies can be composed. We also characterize contractibility of one-dimensional SFTs, and homotopy equivalence for transitive SFTs.

Definition 5.1. *Two subshifts X, Y are* homotopy equivalent*, and we write $X \cong Y$, if there are morphisms $f : X \to Y$, $g : Y \to X$ such that $g \circ f \cong \mathrm{id}_X$ and $f \circ g \cong \mathrm{id}_Y$.*

It is useful to allow auxiliary symbols in the parameter space of a homotopy:

Definition 5.2. *A k-*symbol homotopy *is $h : I_k \times X \to Y$ such that $h(\bar{0}, x) = f(x), h(\overline{k-1}, x) = g(x)$ for all $x \in X$.*

It is very easy to see that this does not change the meaning of homotopy. We give a somewhat roundabout complicated proof, as it contains some ideas that are useful later.

For the proof, and also later in the text, we think of $I = I_2$ as a *two-pointed subshift*, meaning it has a *left endpoint* $\bar{0}$ and a *right endpoint* $\bar{1}$, and these are taken to be part of the structure. We say a morphism between two two-pointed subshifts *respects endpoints*, or is *mod endpoints* if it maps the left (resp. right) endpoint to the left (resp. right) endpoint.

In general, for any two-pointed subshift J we have a notion of J-homotopy between f and g, meaning $h : J \times X \to Y$ such that $h(y, x) = f(x), h(z, x) = g(x)$ for all $x \in X$, where y, z are the two endpoints of J.

Let's say two two-pointed subshifts are *homotopy equivalent mod endpoints* if we can pick the f, g in the definition of homotopy equivalence so that they respect endpoints, and the homotopies between $g \circ f \cong \mathrm{id}_X$ and $f \circ g \cong \mathrm{id}_Y$ are mod endpoints meaning $h(t, x) = x$ for both endpoints x and all t.

Lemma 5.3. *I_m and I_n are homotopy equivalent mod endpoints for $m, n \geq 2$.*

Proof. Define $f : I_m \to I_n$ by

$$f(x)_a = \begin{cases} 0 & \text{if } x_a \neq m-1, \\ n-1 & \text{otherwise.} \end{cases}$$

and $g : I_m \to I_n$ by

$$g(x)_a = \begin{cases} 0 & \text{if } x_a \neq n-1, \\ m-1 & \text{otherwise.} \end{cases} .$$

Clearly these maps respect endpoints. The map $g \circ f$ preserves $(m-1)$-symbols, and maps everything else to 0. For the homotopy $h : g \circ f \to \mathrm{id}_{I_m}$ we can use the naive homotopy:

$$h(t,x)_a = \begin{cases} g(f(x))_a & \text{if } t_a = 0, \\ x_a & \text{if } t_a = 1. \end{cases}$$

It is easy to check that $h(t,x) = x$ for all $t \in I_2$ when x is an endpoint. The case of $f \circ g$ is symmetric. □

Lemma 5.4. *Two morphisms are m-symbol homotopic if and only if they are n-symbol homotopic.*

Proof. Suppose we have an m-symbol homotopy $h : I_m \times X \to Y$ between morphisms f, g. Let $\pi : I_m \to I_n$ be any morphism respecting endpoints. Precompose h with $\pi \times \mathrm{id}_X$ to get $\hat{h} : I_n \times X \to Y$, i.e. $\hat{h}(t,x) = h(\pi(t), x)$. Then $\hat{h}(\bar{0}, x) = h(\bar{0}, x) = f(x)$ and $\hat{h}(\overline{m-1}, x) = h(\overline{n-1}, x) = g(x)$. □

Two two-pointed subshifts can be joined in an obvious way by identifying the right endpoint of the first with the left endpoint of the second, analogously to the usual definition of path composition when defining the fundamental group in algebraic topology. In our symbolic category, this turns out a little more subtle, since already the composition of two copies of I_2, is not isomorphic to I_2. Indeed, it is not even a subshift of finite type. We will see that, nevertheless, in the category of subshifts of finite type, the resulting generalized notion of homotopy does not change.

Denote by $I_{a,b}$ the subshift isomorphic to I_2 on symbols $\{a, b\}$, with left and right endpoints $\bar{a}, \bar{b}$.

Definition 5.5. *Let $k \geq 1$. The* k-path shift *is $P_k = I_{0,1} \cup I_{1,2} \cup \cdots \cup I_{k-1,k}$. A* k-step homotopy *is $h : P_k \times X \to Y$ such that $h(\bar{0}, x) = f(x), h(\bar{k}, x) = g(x)$ for all $x \in X$. A* multistep homotopy *is a k-step homotopy for some k.*

Note that P_k is naturally a two-pointed subshift with endpoints $\bar{0}, \bar{k}$. Thus a k-step homotopy is simply a P_k-homotopy in the terminology introduced below Definition 5.2. We have $P_1 = I_2$, and P_k is of finite type if and only if $k = 1$.

Lemma 5.6. *For subshifts X, Y, multistep homotopy is an equivalence relation on morphisms $f : X \to Y$.*

Proof. One can invert and compose homotopies in an obvious way. □

Lemma 5.7. *Every SFT approximation of P_m is homotopy equivalent to I_2 mod endpoints.*

Proof. An SFT approximation of P_m simply checks that the symbols at a, b differ by at most 1 if a and b belong to a translate of a particular finite subset of the group. The idea is then that for the map from I_2 to P_m, if we are deep inside an area of 0s, then we write 0, and otherwise we write a naive transient towards the maximal symbol, depending on the distance from an area of 0s.

In formulas, define $\delta : G \times I_2 \to \mathbb{N} \cup \{\infty\}$ by $\delta(a, x) = \inf\{d_G(a, b) \mid b \in G, x_b = 0\}$. Then for some large r define $f : I_2 \to P_m$ by

$$f(x)_a = \min(\lfloor \delta(a, x)/r \rfloor, m).$$

We claim that f is well-defined, i.e. if r is sufficiently large, then the image is contained in any given SFT approximation of P_m. For this, simply note that $\delta(a, x)/r$ is $\frac{1}{r}$-Lipschitz in a for a fixed x, and combine this with the first sentence of the proof.

For $g : P_m \to I_2$ we use simply the map

$$g(x)_a = \min(x_a, 1).$$

Note that $g \circ f$ and $f \circ g$ indeed respect endpoints. The map $g \circ f$ is homotopy-equivalent to id_{I_2} by the naive homotopy, which is mod endpoints. Now consider the map $f \circ g$. What this map does is it forgets the positions of all non-zero symbols, and replaces them by a (discrete) speed-$\frac{1}{r}$ transient depending on the distance from the nearest 0-symbol, terminating at the m-symbol once far enough.

We can construct a homotopy $h : \mathrm{id}_{P_m} \to f \circ g$ analogously to the above definition of f: define $L : [0, 1] \times \{0, \dots, m+1\}^2 \to \{0, \dots, m+1\}$ by $L(t, i, j) = \lfloor i + (j - i)r \rfloor$ (one can think of this as a discrete analog of a linear interpolation map – or as a discrete homotopy). Then define

$$h(t, x)_a = L(\min(\delta(a, t)/r, 1), x_a, f(g(x))_a).$$

The calculation $h(\bar{0}, x)_a = L(0, x_a, f(g(x))_a) = x_a$ for all $x \in P_m, a \in G$ and $h(\bar{1}, x)_a = L(1, x_a, f \circ g) = f(g(x))_a$ shows that this is a homotopy, and it respects endpoints because f, g respect endpoints and $L(t, i, i) = i$ for all t, i. □

The reader may notice that in the previous proof, no properties of the specific morphism $f \circ g$ were used to define the homotopy (other than checking the preservation of endpoints). Indeed, any two endomorphisms are homotopic on this subshift, by the same formula. This reflects the fact that SFT approximations of P_m are contractible subshifts, as we will see.

Lemma 5.8. *If Y is a subshift of finite type, then k-step homotopy is equivalent to 1-step homotopy for morphisms $f : X \to Y$.*

Proof. Consider a k-symbol homotopy $h : I_k \times X \to Y$ from f to g (recall that k-symbol homotopy is equivalent to 2-symbol homotopy). Let $\pi : P_k \to I_k$ be the natural inclusion. Then $h \circ (\pi \times \mathrm{id}_X) : P_k \times X \to Y$ is a homotopy from f to g, thus is a k-step homotopy.

Consider then a k-step homotopy $h : P_k \times X \to Y$ from f to g. Since Y is a subshift of finite type, there is an SFT approximation P of P_k such that a domain extension $\tilde{h} : P \times X \to Y$ of h exists. Let π be a morphism from I_2 to P which respects endpoints (endpoints of P being those of P_k) from the previous lemma. Then $\tilde{h} \circ (\pi \times \mathrm{id}_X) : I_2 \times X \to Y$ is a homotopy from f to g. □

Lemma 5.9. *Let X be a subshift containing a fixed point $\bar{a}$, let $\pi_1, \pi_2 : X \times X \to X$ be the projections. Then*

$$\mathrm{id}_X \cong_k (x \mapsto \bar{a}) \implies \pi_1 \cong_{2k} \pi_2$$

and

$$\pi_1 \cong_k \pi_2 \implies \mathrm{id}_X \cong_k (x \mapsto \bar{a}).$$

Proof. We simply mimic the proof of Lemma 2.1, though in the first case we have to replace the explicit formula by an abstract composition of two homotopies. More precisely, suppose $h \cong_k \mathrm{id}_X \to (x \mapsto \bar{a})$. Then

$$h'(t, x, y) = h(t, x)$$

is a homotopy from π_1 to $((x, y) \mapsto \bar{a})$, and

$$h'(t, x, y) = h(\bar{1} - t, y)$$

is a homotopy from $((x, y) \mapsto \bar{a})$ to π_2. The proof of Lemma 5.6 shows that $\pi_1 \cong_{2k} \pi_2$.

In the second case we can directly copy the formula from Lemma 2.1: if h' is a homotopy between the two projections, then

$$h(t, x) = h'(t, x, \bar{0})$$

gives a homotopy between the identity map on X, and the map from X to $\{\bar{0}\}$. □

Corollary 5.10. *If X is a subshift of finite type with a fixed point, then it is fixed-point contractible if and only if it is contractible. In particular, in this case fixed-point contractibility does not depend on the fixed-point.*

This indeed needs the finite type assumption, as $P_2 = I_{0,1} \cup I_{1,2}$ is fixed-point contractible to $\bar{1}$, but not $\bar{0}$, and is not contractible. Of course, by the above lemma, we do have a 2-step homotopy $\pi_1 \cong_2 \pi_2$ between the two projections $\pi_1, \pi_2 : P_2 \times P_2 \to P_2$.

Lemma 5.11. *Homotopy equivalence is an equivalence relation for subshifts of finite type.*

Proof. Suppose X, Y, Z are subshifts of finite type, and suppose $X \cong Y \cong Z$. There exist $f_1 : X \to Y, f_2 : Y \to Z, g_1 : Y \to X, g_2 : Z \to Y$ such that $g_1 \circ f_1 \cong \mathrm{id}_X, f_1 \circ g_1 \cong \mathrm{id}_Y, g_2 \circ f_2 \cong \mathrm{id}_Y, f_2 \circ g_2 \cong \mathrm{id}_Z$.

Define $f = f_2 \circ f_1 : X \to Z, g : g_1 \circ g_2 : Z \to X$. We claim that these give a homotopy equivalence. Consider $g \circ f : X \to X$. This is just the map $g_1 \circ g_2 \circ f_2 \circ f_1$. We first form a homotopy from $g_1 \circ g_2 \circ f_2 \circ f_1$ to $g_1 \circ f_1$. For this, let $h_2 : g_2 \circ f_2 \cong \mathrm{id}_Y$. We define the block map

$$h(t, x) = g_1(h_2(t, f_1(x)))$$

and observe that

$$h(\bar{0}, x) = g_1(h_2(\bar{0}, f_1(x))) = g_1(g_2 \circ f_2(f_1(x))))$$

and

$$h(\bar{1}, x) = g_1(h_2(\bar{1}, f_1(x))) = g_1(\mathrm{id}_Y(f_1(x)))) = g_1(f_1(x))$$

as desired. From $g_1 \circ f_1 \cong \mathrm{id}_X$ it follows from Lemma 5.6 that $g \circ f$ and id_X are 2-step homotopic, and since X is of finite type, they are homotopic. Similarly, one shows that $f \circ g \cong \mathrm{id}_Y$. □

Theorem 5.12. *Let X be a subshift of finite type. Then X is homotopy equivalent to a one-point subshift if and only if it is contractible and has a fixed-point.*

Proof. If X does not have a fixed-point, then it of course cannot be homotopy equivalent to a one-point subshift (as there is no map from the one-point subshift into X). So we simply need to show that for subshifts of finite type with a fixed-point, homotopy equivalence with a one-point subshift is equivalent to contractibility.

Suppose $\bar{0} \in X$. Suppose first that X is contractible. By the previous corollary, it is fixed-point contractible to $\bar{0}$, meaning $h : \mathrm{id}_X \cong \bar{0}$. Let $f : X \to \{\bar{0}\}$ be the constant map, and let $g : \{\bar{0}\} \to X$ be the embedding. Then $g \circ f$ is the endomorphism of X that maps all points to $\bar{0}$. By assumption, we have $h^R : g \circ f \cong \mathrm{id}_X$, where h^R is the inverse homotopy obtained by changing the rules of 0 and 1. On the other hand, $f \circ g = \mathrm{id}_{\{\bar{0}\}}$ so certainly these maps are homotopic. Thus, f and g form a homotopy equivalence.

Next suppose that X is homotopy equivalent to a one-point subshift. We may assume this subshift is $\bar{a}$. It follows that there exists maps $f : X \to \{\bar{a}\}$ and $g : \{\bar{a}\} \to X$ such that $g \circ f \cong \mathrm{id}_X$ and $f \circ g \cong \mathrm{id}_{\{\bar{a}\}}$. Now $g \circ f$ is an endomorphism fo X that maps all points to some $\bar{b} = g(\bar{a})$. So already $g \circ f \cong \mathrm{id}_X$ implies that X is fixed-point contractible, thus contractible. □

Lemma 5.13. *If X is contractible and X is homotopy equivalent to Y, then Y is 3-step contractible. In particular, if Y is of finite type, it is contractible.*

Proof. This is clear from the previous lemma if X is a subshift of finite type with a fixed point. We give a direct proof for the general case. The assumptions are that there is a homotopy $h : \pi_1 \cong \pi_2$ where $\pi_1, \pi_2 : X \times X \to X$ are the two projections, and that there exist $f : X \to Y$ and $g : Y \to X$ such that $g \circ f \cong \mathrm{id}_X$ and $h' : f \circ g \cong \mathrm{id}_Y$.

Concretely $h' : f \circ g \cong \mathrm{id}_Y$ means

$$h' : I_2 \times Y \to Y$$

such that $h'(\bar{0}, y) = f \circ g(y)$ and $h'(\bar{1}, y) = y$.

We define

$$h''(t, y, y') = f(h(t, g(y), g(y'))) : I_2 \times Y \times Y \to Y.$$

Clearly this is indeed a morphism (and the codomain choice is valid). We have

$$h''(\bar{0}, y, y') = f(h(\bar{0}, g(y), g(y'))) = f(g(y))$$

and

$$h''(\bar{1}, y, y') = f(h(\bar{1}, g(y), g(y'))) = f(g(y'))$$

so this is a homotopy from $f \circ g \circ \pi_1$ to $f \circ g \circ \pi_2$.

Next we observe that $\pi_1 \cong f \circ g \circ \pi_1$: simply use $h'''(t, y, y') = h'(\bar{1} - t, y)$ (where subtraction is performed cellwise, $(z - z')_a = z_a - z'_a$). Similarly, $f \circ g \circ \pi_2 \cong \pi_2$. We conclude $\pi_1 \cong_2 \pi_2$, i.e. 3-step contractibility. Finally, when Y is SFT, we obtain contractibility by Lemma 5.8. $\square$

We now conclude that P_m has contractible SFT approximations, as was claimed previously.

Proposition 5.14. *Every SFT approximation of P_m is contractible.*

Proof. We showed above that every SFT approximation of P_m is homotopy equivalent to I_2 (even mod endpoints). Since I_2 is contractible and SFT approximations are SFT, we conclude from the previous lemma that indeed SFT approximations of P_m are contractible. $\square$

Lemma 5.15. *If Y is contractible, then any morphisms $f, g : X \to Y$ are homotopic.*

Proof. Let $h : I_2 \times Y \times Y \to Y$ be a homotopy from π_1 to π_2. Map $h'(t, y) = h(t, f(y), g(y))$. $\square$

The following very simple "time dilation lemma" shows that in the time parameter t, we may always assume all 0-areas are contained in large balls. It would be possible to simultaneously say something nice about 1-areas, at least on abelian groups, but we do not know in which generality one can make both types of areas consist of literal balls.

Lemma 5.16. *Let $J \subset I_2$ be the subshift where for each $x \in J$, we require that $x_a = 0 \implies x|bB_n \equiv 0$ for some $b \in aB_n$ (where B_n is the radius n ball for some symmetric generating set S). Then the notion of J-homotopy (with the same endpoints $\bar{0}, \bar{1}$) is equivalent to the notion of I-homotopy.*

Proof. An I_2-homotopy immediately gives a J-homotopy by restriction. From a J-homotopy, we obtain an I_2-homotopy by constructing an endpoint-respecting morphism from I_2 to J.

Define $f : I_2 \to I_2$ by $f(t)_a = \min(t|aB_n)$. This clearly fixes the endpoints. For $t \in I$, whenever $f(t)_a = 0$, we have $0 \in t|aB_n$, say $t_{ab} = 0$ for $b \in B_n$. Then of course $f(t)|abB_n \equiv 0$, and we conclude that $f(t) \in J$. $\square$

For the group $\mathbb{Z}$, we can easily turn both 0- and 1-areas into long intervals.

Lemma 5.17. *For $G = \mathbb{Z}$, let $J \subset I_2$ be the subshift where for each $x \in J$, we require that $x_a = 0 \implies x|[i, i+n] \equiv 0$ for some $i \in \mathbb{Z}$ such that $a \in [i, i+n]$. Then the notion of J-homotopy (with the same endpoints $\bar{0}, \bar{1}$) is equivalent to the notion of I_2-homotopy.*

Proof. Apply first $f(t)_a = \min(t|[a, a+10n])$ and then $g(t)_a = \max(t|[a, a+n])$. $\square$

Theorem 5.18. *In one dimension, an SFT is contractible if and only if it is mixing.*

Proof. If a subshift is contractible, it is strongly irreducible, thus mixing.

To show that a mixing $\mathbb{Z}$-SFT X is contractible, observe first that contractibility is conjugacy-invariant, so we may assume X has window $\{0,1\}$. Recall [13] that mixing is equivalent to the existence of n such that for any two symbols a, b, there exists $v = v(a,b)$ such that $|v| = n$ and $avb \in L(X)$. Let us index the word v as $v_0 v_1 \ldots v_{n-1}$.

By the previous lemma, it suffices to construct a J-homotopy between the two projections, where in J every 0- and 1-symbol belongs to an interval of length at least $n+1$. Define $h : J \times X \times X \to X$ by

$$h(t,x,y)_i = \begin{cases} x_i & \text{if } t|[i,i+n] \equiv 0 \\ v(x_{i+j-n-1}, y_{i+j})_{n-j} & \text{if } \exists j \in [1,n] : t|[i,i+j-1] \equiv 0, t_{i+j} = 1 \\ y_i & \text{if } t|[i,i+n] \equiv 1 \\ v(y_{i+j-n-1}, x_{i+j})_{n-j} & \text{if } \exists j \in [1,n] : t|[i,i+j-1] \equiv 1, t_{i+j} = 0 \end{cases}$$

Since we restrict to configurations in J, it is straightforward to verify that this is a homotopy between the two projections, as desired, namely between coordinates $i < j$ where we copy from different points, there is at some point a change in t, and at the change we introduce a valid transient using v. $\square$

Theorem 5.19. *In one dimension, two transitive SFTs are homotopy equivalent if and only if they map into each other.*

Proof. If there is no map from X to Y (or vice versa), then X and Y clearly cannot be homotopy equivalent. We show that they are homotopy equivalent in all other cases.

For mixing SFTs, the claim is immediate from the previous theorem, namely if we have maps $f : X \to Y$ and $g : Y \to X$, by contractibility of X and Y, and by the previous lemma, we have $f \circ g \cong \mathrm{id}_Y$ and $g \circ f \cong \mathrm{id}_X$.

To deal with general transitive SFTs, we recall basic structure theory of SFTs (see [13]). Namely, we consider the *primitive period* of a transitive SFT, which is the least $p \in \mathbb{N}_+$ such that we can find a clopen cross-section C, meaning that every $x \in C$ returns to C in exactly p steps, and the first-return map to C is mixing. Each $x \in X$ then has a well-defined *phase* $i \in \mathbb{Z}_p$.

The primitive period cannot increase under a morphism (by mixing of the first-return map), so two transitive SFTs that map into each other must have the same primitive period. Consider now two transitive SFTs X, Y which map into each other. These maps $f : X \to Y, g : Y \to X$ must act consistently on the phase, in the sense that if some $x \in X$ (resp. $y \in Y$) has phase i and $f(x)$ (resp. $g(y)$) has phase j, then every point with phase i' is mapped to a point with phase $i' + (j - i)$. Thus by composing f and g with shifts we may assume they map the chosen cross-sections of X and Y into each other.

The first-return map is well-known to be of finite type, so as a mixing SFT it is contractible by the above proof. It follows that the maps f, g give a homotopy equivalence between the first-return maps of X, Y to their chosen cross-sections, by the first part of the proof. The homotopy must automatically respect phases of points, so in fact f, g are a homotopy equivalence between X and Y. $\square$

Two transitive $\mathbb{Z}$-SFTs X and Y map into each other if and only if $P(X) = P(Y)$ where $P(Z)$ is the set of (not necessarily least) periods of points in a subshift Z [13].

6 Finite periodic asymptotic dimension and periodic points

In this section, we show that contractibility implies dense periodic points, on groups which are finite-dimensional in an appropriate sense. We also show that the FEP (which is incomparable with contractibility) implies dense periodic points with a very similar proof. We suspect that there is a natural way to unify these results, but we do not have one at present.

The idea behind both proofs is the following variant of Gromov's asymptotic dimension. Let G be a group.

Definition 6.1. *Let $X_{G,d,r,m}$ (or just $X_{d,r,m}$ when dealing with a single group G) be the subshift of the full shift $\{0,\dots,d-1\}^G$, containing those configurations $x \in \{0,\dots,d-1\}^G$ such that for each $0 \leq c < d$, if we define a graph with nodes $\{a \in G \mid x_a = c\}$ and edges $\{(a,b) \in G^2 \mid x_a = x_b = c \wedge |a^{-1}b| \leq r\}$, then the connected components of the graph are of cardinality at most m.*

Note that $X_{G,d,r,m}$ is a subshift of finite type for any G, d, r, m. In words, in $X_{G,d,r,m}$ we require that for each color $0 \leq c < d$, the graph with c-colored nodes and edges between r-separated nodes has components of cardinality at most m. If $x \in X_{d,r,m}$, we call the sets $x^{-1}(i) \subset G$ *color classes*, and refer to the (bounded) components of the aforementioned graph *(monochromatic) r-components*. We usually think in terms of r-paths, meaning sequences $a_1, a_2, \dots$ in G, and the assumption that $x \in X_{d,r,m}$ for some m is equivalent to showing that there is a uniform bound on the lengths of such injective paths.

Recall that one of the equivalent definitions of Gromov's asymptotic dimension is that it is the infimum of all d such that $\forall r : \exists m : X_{d+1,r,m}$ is nonempty.

We define a periodic variant of this:

Definition 6.2. *The* periodic asymptotic dimension $\mathrm{pdim}(G)$ *of a group G is the infimum of all d such that*

$$\forall r : \exists m : X_{d+1,r,m} \text{ has a finite subsystem.}$$

A finite subsystem is the same as having a point with finite orbit, equivalently a finite-index stabilizer.

We can of course generalize asymptotic dimension and periodic asymptotic dimension to a general countable group G by instead of word norm picking $S_r \Subset G$, where $G = \bigcup_r S_r$ as an increasing union, and by r-paths $p : I \to G$ referring to paths with offsets $p(i)^{-1}p(i+1) \in S_r$. It is easy to see that then the dimension of a group (in any sense) is just the supremum of dimensions of its finitely-generated subgroups, which is why the definition above is given in terms of word norm. Only in Theorem 9.6 do we need this generalized notion.

Recall that the *residual finiteness core* [7] of a group G is the intersection of all of its finite-index subgroups.

Lemma 6.3. *If an infinite f.g. group G has finite periodic asymptotic dimension, then its residual finiteness core is locally finite. In particular, if G is torsion-free, then it must be residually finite.*

Proof. Suppose G has dimension d. Suppose that the residual finiteness core contains an infinite finitely-generated subgroup H, generated by say $T \Subset H$.

Pick r such that $T \subset B_r$. Consider now a periodic point $x \in X_{d,r,m}$ of G. Since the stabilizer K of this configuration contains H, we have $x|H = i^H$ where $i = x_{1_G}$. Since H is generated by T, this is a single monochromatic component, contradicting $x \in X_{d,r,m}$ (for all m). □

We show in Section 9.1 for example that all virtually polycyclic groups have finite periodic asymptotic dimension, and so do metabelian Baumslag-Solitar groups and the lamplighter group $\mathbb{Z}_2 \wr \mathbb{Z}$. Furthermore, the dimension of $\mathbb{Z}^d$ is the expected d for this notion.

There are many equivalent definitions of the asymptotic dimension. We will need the analog of one of them.

Consider a subshift $Y_{d+1,r,m}$ which, instead of recording actual color classes, only records their shapes. This can be coded into a finite alphabet by explicitly recording the relative shape of an r-component containing a cell into its color. Specifically, the constraint is that a configuration of $Y_{d+1,r,m}$ codes a partition where each component is r-connected and of cardinality at most m, and each B_r-ball touches at most $d+1$ distinct components. We indeed think of the points of this subshift as partitions, mostly ignoring the (trivial) coding issues.

Lemma 6.4. *There is a factor map $\phi : X_{d+1,2r,m} \to Y_{d+1,r,m}$. Conversely, if r' is sufficiently large compared to r, and m' is sufficiently large compared to m, then there is a factor map $\phi : Y_{d+1,r',m} \to X_{d+1,r,m'}$.*

Proof. The first claim is clear, we can simply write down the arrows to other cells in the color component of a cell (w.r.t. r-paths). Since even $2r$-components have cardinality at most m, the r-components do as well. A ball of radius r cannot touch more than two different monochromatic $2r$-components in a point of $X_{d,2r,m}$, by the definition of a monochromatic component.

If $y \in Y_{d+1,r',m}$ then for cell a we pick its color $\phi(y)_a$ to be the maximal $1 \leq k \leq d+1$ such that the rk-ball around it touches k or more components of $y \in Y_{d+1,r',m}$ (so it actually has to touch exactly k, or we would not reach the maximum, unless $k = d+1$).

It is easy to see that in fact r-components are now bounded, assuming that indeed when the maximal color k is used, the $r(k+1)$-ball does not touch $k+1$ or more color components (which is automatic if $r' > r(d+2)$). Namely, if a path with r-jumps stays in a single color class, then it has to actually stay near the same set of components of y. The reason is simply that if the color stays k and some component is exchanged for another, the one dropped is now at distance at most $r(k+1)$, so in fact we should in fact have used the color $k+1$ (or larger) in the previous step. □

We use the same idea in Theorem 9.1.

Note that by the lemma, we can also use $Y_{d+1,r,m}$ to characterize periodic asymptotic dimension, since factor maps preserve periods.

Definition 6.5. *Say that a subset A of a metric space B is an R-*packing *if $\min\{d(a,b) \mid a,b \in A, a \neq b\} \geq R$, an R-*covering *if $\max_{b \in B} d(b,A) \leq R$, and an R-*net *if it is both.*

Lemma 6.6. *On the group $\mathbb{Z}^d$, there exists D such that for all r, the subshift $Y_{D,r,m}$ contains a contractible SFT for large enough m.*

We can surely pick $D = d + 1$, but the proof is geometrically simpler with a larger D.

Proof. We use the ℓ_∞ metric on $\mathbb{Z}^d$. To an R-net of vectors with integer coordinates, we associate a partition of the plane, and vice versa. There are many ways to do this, the most obvious possibility being Voronoi cells. However, the non-uniqueness of centers poses some technical difficulties, so we use a different method.

For each point $\vec{v} \in \mathbb{Z}^d$ consider the cone of all points whose coordinates are at least as large as those in the *corner point* $\vec{v}$. For $\vec{u}$ in this cone, associate $\vec{u}$ to the component $\vec{v}$ if $\vec{v}$ is the closest corner point in that it could be associated to. We use the lexicographic order to break distance ties. It is then easy to see that components are 1-connected, since if $\vec{u}$ is associated to the cone of corner $\vec{v}$, we stay in the same component when moving toward $\vec{v}$ along standard generators.

Clearly the sets of components are in shift-commuting continuous bijection with the corner points, so we see that this set of partitions is an SFT Y (we simply check the R-net condition for the corner points).

It is easy to see that there is a bound on the cardinality of the components. What we need to show now is that there is a bound (independent of R) on how large D needs to be so that at most D distinct components can meet an r-ball. For this a simple volume argument suffices. We observe that a component contains at least $\Omega(R^d)$ elements because of the growth rate of the group and that we have an R-packing, and on the other hand it is contained in a $O(R)$-radius since we have covering radius R (the invisible constants can be taken to be, for example, 10).

We have shown that the SFT Y is contained in $Y_{D,r,m}$, where D is uniformly bounded, and m depends on R which in turn depends on r.

We claim that the SFT Y is contractible. For this, it suffices to pick where the contraction homotopy writes corner points, given a time parameter t, and the sets of corner points of $x, y \in Y$. Then we simply write back the corresponding partition to get a point of Y.

In time parameter t, when t has a large area of 0s, we of course want to copy from x, Specifically, in such areas, we pick corners from the first input point. We do the same for areas of 1s.

After this, we add all corner points from y which are not at distance less than R from point already included. After this step, we have a set of corners points which are a R-packing and a $2R$-covering. For the latter claim, any $\vec{u}$ will be at distance R from a corner point $\vec{v}$ from y, and because this point was not included, there is an already-included point at distance R from $\vec{v}$.

On an abelian group,[1] there is a general method for turning an nR-covering R-packing into an R-net, in a shift-invariant continuous way. This is just a greedy version of the usual existential proof that such nets exist. Namely, list all vectors of length at most nR, and when considering the vector $\vec{v}_i \in B_{nR}$, add $\vec{v} + \vec{v}_i$ to our set of points if $\vec{v}$ is an already included point and $\vec{v} + \vec{v}_i$ does not break the R-packing constraint. If an element of the group is not included in the end, then this can only be because there is is an element at distance R, so we have an R-covering, yet at no step do we break the R-packing constraint. $\square$

[1]The argument specifically needs a bi-invariant metric, which implies that our group is virtually abelian.

Remark 6.7. *It is possible to extend the definition of periodic asymptotic dimension so that the dimension is allowed to grow with scale, in the sense that if $d : \mathbb{N} \to \mathbb{N}$ is a function, we can say G* admits periodic asymptotic dimension *d if*

$$\forall r : \exists m : X_{d(r)+1,r,m} \text{ has a finite subsystem.}$$

Our proofs should go through with d growing sufficiently slowly rather than being constant.

6.1 Stitching morphisms together

We show that a contraction homotopy can be used to stitch morphisms together. While the statements look quite technical, the idea is simple: if we have multiple morphisms from a subshift X to a full shift which sometimes output $\#$, and sometimes valid content of a subshift Y, then if the non-$\#$ parts of the images together cover Y, and Y is a contractible SFT, then we can use a contraction homotopy to glue a morphism out of them, and we obtain a valid morphism from X to Y.

Lemma 6.8. *Suppose $Y \subset B^G$ is a contractible SFT defined by forbidden patterns $P \subset B^K$, $K \Subset G$. Then there exists a symmetric finite set $1_G \in N \Subset G$ such that the following holds: Suppose we have morphisms $f, g : X \to (B \sqcup \{\#\})^G$ such that no patterns in P appear in the images. Then there is a morphism $k : X \to (B \cup \{\#\})^G$ such that*

- *the image of k contains no patterns from P,*
- *whenever $k(x)_a = \#$, both $f(x)|aN$ and $g(x)|aN$ must contain $\#$, and*
- *whenever $f(x) \in Y$, we have $k(x) = f(x)$.*

Furthermore, the construction is uniform in the sense that for any $T_1 \Subset G$, by picking N large enough we can make sure that there exists T_2 such that $k(x)|T_1 = f(x)|T_1$ whenever $\# \not\sqsubset f(x)|T_2$ (where N, T_2 do not depend on f, g).

Proof. We simply feed $f(x)$ and $g(x)$ into the contraction homotopy, and pick the time parameter so that we copy from $f(x)$ whenever possible. When neither map is producing a large area of non-$\#$-symbols, we simply output $\#$.

Let us say the same in formulas. Let $h : I \times Y \times Y \to Y$ be the contraction homotopy with neighborhood W, and let $\tilde{h} : I \times B^G \times B^G \to B^G$ be the natural extension. We may suppose $1_G \in K$.

Pick $M \Subset G$ such that $W^{-1}K^{-1}KW \subset M$. Define $t : \Sigma^G \to I$ by

$$t(x)_a = 0 \iff f(x)|aM \text{ appears in } X.$$

Now pick $M' \Subset G$ such that $K^{-1}KWM \subset M'$ and define

$$k(x)_a = \begin{cases} \# & \text{if } f(x)|aM' \text{ and } g(x)|aM' \text{ do not appear in } Y \\ \tilde{h}(t(x), f(x), g(y))_a & \text{otherwise.} \end{cases}$$

We claim that the image of k contains no pattern from P. Suppose on the contrary that for some $x \in X$, $k(x)|K \in P$. Then in particular $k(x)_a \neq \#$ for

$a \in K$ and thus we always use the second rule of k. Thus, for all $a \in K$, either $f(x)|aM'$ or $g(x)|aM'$ appears in X.

Suppose first that $f(x)|aM'$ appears in X, for some $a \in K$. Then because $M' \supset K^{-1}KWM$, in fact $f(x)|KWM$ appears in Y. This means $t(x)|KW = 0^{KW}$, so $P \ni \tilde{h}(t(x), f(x), g(x))|K = f(x)|K$ by the definition of the natural extension, and this is a contradiction by the assumption that f does not write P-patterns in its image.

Suppose then that $g(x)|aM'$ is always valid. Certainly then $g(x)|KW \sqsubset X$. If $f(x)|KW$ is not valid, then since $M \supset W^{-1}K^{-1}KW$, all calculations of the t-value in KW see this invalid pattern, and we have $t(x)|KW = 1^{KW}$. This means that $P \ni \tilde{h}(t(x), f(x), g(x))|K = g(x)|K$, again a contradiction.

This proves the first item of the claim.

For the second item of the statement, suppose $k(x)_a = \#$. Then the first item of the definition of k is used at a. This means that neither $f(x)|aM'$ nor $g(x)|aM'$ appears in Y. By Lemma 2.4, there exists N such that if $z|aN$ has no forbidden pattern from P, then $z|aM'$ is globally valid in Y. Since $f(x)$ and $g(x)$ do not contain forbidden patterns of Y (over alphabet B), both $f(x)|aN$ and $g(x)|aN$ must contain $\#$ whenever neither $f(x)|aM'$ nor $g(x)|aM'$ appears in Y. We can of course then add 1_G to N and symmetrize it, as required in the statement.

For the third item of the statement, if $f(x) \in Y$, then k will always use the second rule of the definition, and $t(x) = \bar{0}$, so $k(x) = \tilde{h}(\bar{0}, f(x), g(x)) = f(x)$.

The last claim is clear because the event $t(x)|KW = 0^{KW}$ is only a function of which parts of the configuration $f(x)$ are valid. $\square$

Lemma 6.9. *Suppose $Y \subset B^G$ is a contractible SFT, and let N be as in the previous lemma. Suppose that $f_1, f_2, \ldots, f_\ell : X \to (B \cup \{\#\})^G$ are morphisms such that no forbidden patterns of Y over B appear in their images, and for all $x \in X$ and $a \in G$, for some i we have $\# \not\sqsubset f_i(x)|aN^{\ell-1}$. Then there is a morphism $k : X \to Y$ such that if $f_1(x) \in Y$, then $k(x) = f_1(x)$. Furthermore, the construction is uniform in the sense that for any $T_1 \Subset G$, by picking large enough N, we may ensure that there exists T_2 such that $k(x)|T_1 = f_1(x)|T_1$ whenever $\# \not\sqsubset f_1(x)|T_2$.*

Intuitively k locally copies the image from $f_i(x)$ for minimal possible i, and performs the necessary transitions. Often, we just use this lemma to conclude the existence of some morphism $k : X \to Y$.

Proof. The idea is to simply iterate the previous lemma.

More precisely, let k_1 be given by the previous lemma for the pair f_1, f_2, and define k_{i+1} by applying the previous lemma to k_i and f_{i+1}. Note that this inductive definition makes sense, since the previous lemma can be applied under just the assumption that no forbidden patterns of Y over alphabet B appear, and the morphism k it produces has this property as well.

As already stated, inductively we see that k_i does not have any forbidden patterns of Y over B. Furthermore, a straightforward induction shows that whenever $k_i(x)_a = \#$, for all $j \leq i+1$ the pattern $f_j(x)|aN^i$ contains $\#$. We also inductively see that whenever $f_1(x) \in Y$, we have $k_i(x) = f_1(x)$.

The last uniformity claim is immediate from the uniformity claim in the previous lemma.

We pick $k = k_{\ell-1}$. If $k(x)_a = \#$ for some $x \in X$, $a \in G$, then for all i we have $\# \sqsubset f_i(x)|aN^{\ell-1}$ by the previous paragraph. But we assumed $\exists i : \# \not\sqsubset f_i(x)|aN^{\ell-1})$, so this cannot happen, and we conclude that $k(x) \in B^G$ for all x. Then $k(x) \in Y$, since there are no forbidden patterns in the image. This concludes the proof. $\square$

6.2 Contractible implies DPP

Theorem 6.10. *Let $Y \subset B^G$ be a contractible subshift on an infinite finitely-generated residually finite group G. If G has finite periodic asymptotic dimension, then Y has dense periodic points.*

Proof. Since G has finite periodic asymptotic dimension, the subshift $X_{d+1,r,m}$ has a periodic point for arbitrarily large r and some m. Let y be any periodic point in it. We may assume the color 0 is used, by picking the minimal possible d.

The most lengthy part of the proof is showing that we can pick "centers" for monochromatic r-components in y – not necessarily in a shift-invariant way, but at least in a periodic way. We begin with this.

By the assumption, the group is residually finite. Thus we can pass to an arbitrarily sparse finite-index subgroup K that fixes y (without actually changing y), where sparse means that the smallest word norm of a non-identity element is large. We may also suppose K is normal.

We show that if K is sufficiently sparse, we can pick an element from each color class K-periodically. Specifically, we claim that there is a K-periodic configuration $x \in \{0,1\}^G$ such that every r-component of a color class $y^{-1}(i)$ of y contains exactly one element $a \in G$ with $x_a = 1$.

We can simply pick centers greedily: as long as there is a monochromatic r-component $C \Subset G$ without a center in y, we pick a center $a \in C$ for it arbitrarily, and then make each element of Ka the center of the component it touches. In other words, we set $x|Ka \equiv 1$. Eventually, every component has a center.

We must check that this process is not contradictory. Specifically, we must check that

- we never pick a second center for C with this procedure (by shift-symmetry this means we also do not pick two centers for any other component that Ka touches); and
- no component D touching Ka that gets a center this way had a center previously.

For the first item, note that $Ka = aK$ by normality, and suppose we get a second center for C. Then $|C \cap aK| \geq 2$. Then $ab, ac \in C$ for some distinct $b, c \in K$ so $1_G \neq b^{-1}c \in C^{-1}C \cap K^{-1}K$. Since $K = K^{-1}K$, we have that K contains a non-1_G element from $C^{-1}C$. Since $C \subset eB_{rm}$ for some $e \in G$, $C^{-1}C \subset B_{2rm}$, and thus we can prevent this by simply picking K sparse enough.

For the second item, if we pick center a for C and $Ka \cap D \neq \emptyset$ (meaning this gives a center for some other existing component D), then suppose for a contradiction that D already has a center b. We have $KD \cap C \neq \emptyset$ so in fact by K-periodicity of y we have $D = cC$ for some $c \in K$. Then $b \in D \implies c^{-1}b \in C$ so C already has a center.

We conclude that we have a K-periodic configuration $(x, y) \in (\{0,1\} \times \{0,1,\dots,d\})^G$ where x marks centers for all r-components of $y \in X_{d,r,m}$. Let X be the orbit closure of (x, y); note that $|X| \leq [G : K]$.

The previous construction can be performed for any r. Let N be as in Lemma 6.9, and apply the previous construction to r such that that $N^{3d} \subset B_r$.

Next, we will produce for each $i = 0, 1, \dots, d$ a partial morphism $f_i : X \to (B \cup \{\#\})^G$, so that their images cover G in the sense of Lemma 6.9.

Let $z \in Y$ be any legal configuration. The way f_i works is very simple: For each i-colored component aC in the input point $q \in X$ (in the projection to the orbit of y), with a the marked center point of C, it positions the origin of z on $a \in C$ and copies content from z, in the sense of setting $f_i(q)|aCN^d = a(z|CN^d)$ in this case.

Since $N^{3d} \subset B_r$, these output patterns do not overlap (by the bound on the size of color components in y). We fill the rest of the configuration $f_i(q)$ with #. No forbidden patterns of Y appear in the image configuration, if N was originally taken large enough (the separate patches being taken from a legal configuration and at least N^d-separated).

We have $d + 1$ morphisms f_i. If $q \in X$ has q_a has color i in the component coming from the orbit of y, then $f_i|aN^d$ is copied from z, so has no #-symbols. Thus the assumptions of Lemma 6.9, and we conclude there is a morphism $k : X \to Y$ combining the morphisms.

The last sentence of Lemma 6.9 shows that if N is chosen large enough, then $k(q)|T_1 = f_1(q)|T_1$ if $f_1(q)$ has a large enough non-#-pattern. For density of periodic points, if U is open and nonempty, we can pick $z \in Y$ and $T_1 \Subset G$ so that $[z|T_1] \in U$, and then $k(q)|T_1 \in U$ if $q \in X$ is such that the central cell is marked and of color 0. $\square$

We note that the proof is somewhat easier in the case that G is torsion-free, since then no finite component can have a self-symmetry, and we can immediately pick centers by a shift-invariant rule without needing to pick a new period N and add artificial center markers. (And in this case we can also remove the explicit assumption of residual finiteness by Lemma 6.3.)

Question 6.11. *On a free group with $k \geq 2$ generators, does there exists a contractible SFT without periodic points?*

We suspect that even strong irreducible SFTs have periodic points (and thus dense periodic points [9]) on free groups, but we do not have a proof.

Remark 6.12. *Both strong irreducibility and having periodic points are decidable properties for SFTs on free groups (the latter observation is due to Piantadosi, the former we are simply claiming here), so one direction of this problem may be attacked by brute force. We implemented the algorithm for strong irreducibility partially, and Piantadosi's algorithm completely, and generated random SFTs for a few hours without finding a counterexample.*

6.3 FEP implies DPP

We have shown that contractibility and SFT imply FEP, and that contractibility implies dense periodic points. In this section we show that FEP also implies dense periodic points.

Theorem 6.13. *If G is a finitely-generated infinite residually finite group with finite periodic asymptotic dimension, then FEP subshifts on it have dense periodic points.*

Proof. Consider an FEP subshift X. Let R be the FEP gap, i.e. if a pattern on domain $W + B_R$ is locally valid (with respect to some fixed finite set of forbidden patterns), then its restriction to W is globally valid.

Pick a generating set S for G. Pick some periodic point $y \in X_{d,r,m}$ where $r \geq 10Rd$. For each color $0 \leq c \leq d$, let C_c be the set of cells $a \in G$ such that $y_a = c$. As in the previous proof, we may assume that we also have a periodic binary configuration x that marks a center for each r-component of each C_c.

On each R-component of $C_0 B_{Rd}$, we position some pattern P, where P only depends on the shape of this component – as in the previous proof, we may assume the color components have marked centers, so this is possible. Note also that $3Rd < r$, so R-components of $C_0 B_{Rd}$ correspond to R-components of C_0, and thus are bounded.

The pattern obtained is locally legal. By the FEP assumption, after we shrink it by R (in the sense of dropping from the support all $a \in G$ such that the ball aB_R is not contained in the support), it is a globally valid pattern with domain containing $C_0 B_{R(d-1)}$. Next, consider $C_1 B_{Rd} \setminus C_0 B_{R(d-1)}$. Again the R-components are bounded (they are bounded even before removing the already used cells $C_0 B_{R(d-1)}$).

For each R-component, we again pick a pattern, only depending on the shape of the component, and the nearby contents of the partial configuration we have from the 0th step. This gives a periodic configuration, which is locally valid. As in the previous step, we shrink it by R, to get something globally valid, and continue by induction.

In the next step, we will have a locally valid pattern on $C_2 B_{Rd} \cup C_1 B_{R(d-1)} \cup C_0 B_{R(d-2)}$ followed by a globally valid pattern on $C_2 B_{R(d-1)} \cup C_1 B_{R(d-2)} \cup C_0 B_{R(d-3)}$, and so on.

Finally after the dth step (or $2d + 1$ if we count the steps where the configuration is only locally valid) we have a periodic configuration defined on all of $\bigcup_c C_c$ and which is periodic, because the configuration stays periodic at all times.

This gives one periodic point, and picking patterns P suitably we get a dense set of them. Alternatively, we can simply observe that an FEP subshift is strongly irreducible, and a strongly irreducible SFT with a periodic point has dense periodic points [9] on any residually finite group. $\square$

7 Equiconnectedness

Definition 7.1. *We say a subshift X is* equiconnected*, if there is a homotopy $h : I_2 \times X \times X \to X$ between the left and right projections $\pi_1, \pi_2 : X \times X \to X$, with the additional property that $h(t, x, x) = x$ for all $x \in X$.*

One may also call an equiconnected subshift one that is contractible mod(ulo) diagonal, to emphasize that this is a natural strengthening of contractibility.

In the classical topological setting, equiconnectedness is a true strengthening of contractibility. For example, equiconnectedness implies that for any base point x_0, there is a base point preserving homotopy between the space and the constant map at x_0. This is called *strong deformation retract contractibility.* It is well-known that there are spaces that are contractible, but one cannot preserve a particular base point. The usual example is the comb space

$$((\{1/n \mid n \in \mathbb{N}\} \cup \{0\}) \times [0,1]) \cup ([0,1] \times \{0\}) \subset \mathbb{R}^2$$

with the induced topology. There is clearly no base point preserving homotopy between the identity map and the constant map at $(0,1)$.

It can be shown that in the topological setting, equiconnectedness is equivalent to $X \times X$ strong deformation retracting to the diagonal Δ_X. The same is true in our case, by the same proof, at least when X is SFT. For this, analogously to the topological setting we say that X is a *strong deformation retract* of Y if $X \subset Y$ and there is a homotopy $h : \mathrm{id}_Y \to g$ where $g(Y) = X$ and $h(t,x) = x$ for all $t \in I$ and $x \in X$.

Lemma 7.2. *An SFT $X \subset A^G$ is equiconnected if and only if $X \times X$ strong deformation retracts to Δ_X.*

Proof. Suppose first that $h : I \times X \times X \to X \times X$ is a strong deformation retraction to the diagonal. Then define $g(x,y) = \pi_1(h(\bar{1},x,y))$. Then $\pi_1 : X \times X \to X$ is homotopic to g by $h'(t,x,y) = \pi_1(h(t,x,y))$. Similarly, $\pi_2 : X \times X \to X$ is homotopic to $(x,y) \mapsto \pi_2(h(t,x,y))$.

Of course we have $\pi_1(h(1,x,y)) = \pi_2(h(1,x,y))$ so π_1 and π_2 are two-step homotopic. Furthermore, the homotopies respect the diagonal because h is a strong deformation retraction. Since X is SFT, π_1 and π_2 are homotopic by Lemma 5.8. Furthermore it is easy to see that this construction yields a morphisms that respects the diagonal.

On the other hand suppose π_1, π_2 are homotopic mod diagonal by $h : I \times X \times X \to X$. A strong deformation retraction to the diagonal is given by the formula $h'(t,x,y) = (h(t,x,y), y)$. □

As in the topological setting, equiconnectedness is not equivalent to contractibility in our symbolic dynamical category. Indeed, plenty of contractible subshifts are not of finite type, but we show below that on a large family of groups, equiconnectedness implies finite type. (On the groups $\mathbb{Z}^d$, we will later show the converse as well.)

If $A \subset G$, write $A^{\circ r} = \{a \in G \mid aB_r \subset A\}$, where B_r is the ball of radius r (around the identity element, with respect to the fixed word metric).

Definition 7.3. *Consider a group G with generating set S. Let $r, R \in \mathbb{N}$, and consider the smallest family of subsets containing all balls of radius R, which is closed under the* almost-union *operation defined by*

$$A \overset{r}{\cup} B = A^{\circ r} \cup B^{\circ r} \cup ((A \cup B) \setminus (AB_r \cap BB_r)),$$

and under taking subsets. If for all r, there exists R such that left translates aB_R, $a \in G$, generate all larger balls in this sense, then we say G has the patching property.

The operation is monotone decreasing in r, so the choice of metric used to define the operation does not matter. The operation is not monotone in A, B, but since we explicitly close the family under subsets, we could replace the balls

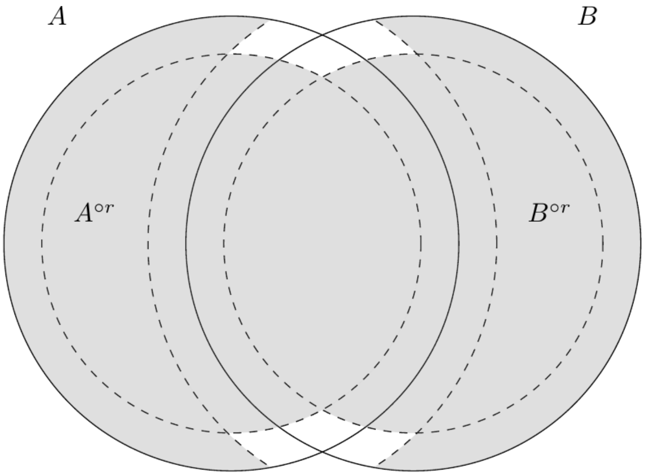

Figure 1: Almost-union of two Euclidean balls A, B, shown in gray.

aB_R with left translates of any family of sets $S_i \Subset G$ such that every $b \in G$ is eventually contained in every S_i. In particular, the choice of metric does not matter in the definition.

Theorem 7.4. *If G has the patching property, then every G-subshift X which is equiconnected is a contractible SFT.*

Proof. Contractibility is obvious (even without the patching property).

Suppose X is equiconnected. Let $h : I_2 \times X \times X \to X$ be the homotopy in the definition of equiconnectedness, say it has radius $r/3$. Let R be as in the definition of the patching property for G, with parameter r, and let X' be the SFT approximation of X where balls B_R of radius R are required to be globally valid for X.

Consider $x \in X'$. Let $\mathcal{C}$ be the family of sets A such that $x|A$ is globally valid in X. By definition, the balls of radius R belong to $\mathcal{C}$, and of course this set is closed under taking subsets. It now suffices to show closure under almost-unions, as then every finite pattern in x is globally valid for X, meaning $x \in X$.

So consider two A, B such that $x|A$ and $x|B$ are globally valid, say $y, z \in X$ satisfy $y|A = x|A$, $z|B = x|B$. Define $t \in \{0,1\}^G$ by $t_a = 0 \iff d'_G(a, A) \leq d'_G(a, B)$. Here $d'_G(a, C)$ is defined as the distance to C when this distance is positive, and otherwise to be $-n$ where n is maximal such that $aB_n \subset C$.

Then $w = h(t, y, z) \in X$, and it suffices to show that w agrees with x in the set

$$A^{\circ r} \cup B^{\circ r} \cup (A \cup B) \setminus (B_r(A) \cap B_r(B)),$$

as then this set must be globally valid, thus in $\mathcal{C}$.

Suppose first that $a \in A^{\circ r}$. Then $d'_G(a, A) \leq -r$. If $d'_G(a, B) > -r/3$, then by the 1-Lipschitz property of d'_G, $t|aB_{r/3} \equiv 0$. Thus $w_a = y_a = x_a$ since h is a homotopy. If $d'_G(a, B) \leq -r/3$, then $aB_{r/3} \subset A \cap B$, thus $y|aB_{r/3} = z|aB_{r/3}$, and thus $w_a = x_a$ since h is respects the diagonal. The case $a \in B^{\circ r}$ is symmetric.

Suppose then that $a \in (A \cup B) \setminus (B_r(A) \cap B_r(B))$. Suppose $a \in A$. Then $d'_G(a, A) \leq 0$ and $d'_G(a, B) > r$, so the first case of the previous argument applies, i.e. $w_a = y_a = x_a$. The case $a \in B$ is symmetric. $\square$

We list some examples of groups that this covers. These are proved in Section 9.2.

Corollary 7.5. *Every equiconnected subshift is SFT on every group with finite asymptotic dimension or subexponential growth. Thus the implication holds for example for virtually polycyclic groups, metabelian Baumslag-Solitar groups, the lamplighter group $\mathbb{Z}_2 \wr \mathbb{Z}$, free groups, and the Grigorchuk group.*

On the other hand, on every group, a contractible SFT is equiconnected as soon as it has a fixed point.

Theorem 7.6. *Every contractible SFT with a fixed point is equiconnected.*

Proof. A contractible SFT X with a fixed point is a retract of a full shift A^G, suppose $X \subset A^G$ and let $r : A^G \to X$ be the retraction. Let

$$c(t, x, y)_a = \begin{cases} x_a & \text{if } t_a = 0 \\ y_a & \text{otherwise.} \end{cases}$$

be the naive homotopy. Define

$$h(t, x, y) = r(c(t, x, y)).$$

We have $r(A^G) \subset X$,

$$h(\bar{0}, x, y) = r(c(\bar{0}, x, y)) = r(x) = x,$$

$$h(\bar{1}, x, y) = r(c(\bar{1}, x, y)) = r(y) = y,$$

and

$$h(t, x, x) = r(c(t, x, x)) = r(x) = x$$

for all $t \in I_2$. $\square$

Question 7.7. *Is there a contractible SFT that is not equiconnected?*

We see in Section 8.4 that for abelian groups, the two properties coincide, so for these groups the answer to the question is "no".

8 Comparison with other dynamical properties

8.1 Contractibility and finite type

As we have seen, homotopy theory works somewhat more nicely for subshifts of finite type than general ones, in that homotopies are composable in the former case. However, contractible subshifts are not necessarily finite type, and we have shown that density of periodic points holds without this additional assumption (on a large class of groups).

We give here a special property that guarantees contractibility of a subshifts, which allows constructing many non-SFT examples.

Definition 8.1. *Let $A \ni 0$ and let $X \subset A^G$ be any subshift defined by forbidden patterns that do not contain* 0*, and whose domains are connected with respect to some generating set S of G. Then we say X has property Z0.*

This implies the existence of a safe symbol, and also the property of zero gluing studied in [17].

Proposition 8.2. *Every subshift with property Z0 is contractible.*

Proof. We may use the contraction homotopy

$$h(t, x, y)_a = \begin{cases} x_a & \text{if } t|aS = 0^{aS}, \\ y_a & \text{if } t|aS = 1^{aS}, \\ 0 & \text{otherwise.} \end{cases}$$

Clearly $h : I \times X \times X \to A^G$ and $h(\bar{0}, x, y) = x, h(\bar{1}, x, y) = y$. On the other hand if $h(t, x, y)|N$ is one of the forbidden patterns defining X, then it must not involve coordinates where h writes 0, nor can it only involve coordinates coming from only one of the first two cases. Say $h(t, x, y)_a$ comes from the first item (is copied from x) meaning $t|aS = 0^{aS}$ and $h(t, x, y)_b$ comes from the second meaning $t|bS = 1^{bS}$.

Since N is S-connected, there is a path $p : [0, k] \to G$ with $p(i) \in N$ for all $i \in [0, k]$ and $p(i)^{-1}p(i + 1) \in S$. Then there must be a first i such that $t|p(i + 1)S \not\equiv 0$. Since $p(i + 1) \in p(i)S$, $p(i + 1) = 0$ and thus $t|p(i + 1)S$ is not of the form $c^{p(i+1)S}$, meaning $h(t, x, y)_{p(i+1)} = 0$, contradicting the assumption that $h(t, x, y)|N$ is one of the forbidden patterns defining X. □

Proposition 8.3. *On every infinite f.g. group G, there exist uncountably many contractible subshifts over the three-symbol alphabet.*

Proof. Let S be the generating set, and require that every S-path contains a configuration from a given $\mathbb{Z}$-subshift over symbols $1, 2$. This always gives a G-subshift with property Z0, and using Lemma 2.2 it is easy to see that they are distinct for any two $\mathbb{Z}$-subshifts that are disting up to reversal. There are uncountably many $\mathbb{Z}$-subshifts, so this concludes the proof. □

The two-symbol alphabet suffices, even using the Z0 property, but the proof is slightly trickier, and we leave this to the interested reader.

Proposition 8.4. *On every infinite f.g. group G, there exists a proper sofic contractible full shift factor.*

Proof. For the full shift, we use the alphabet $(S' \times \{0, 1\}) \cup \{0, 1, \#\}$ where $S' = \{(s, s') \in S^2 \mid s \neq s'\}$ meaning each cell contains a bit or # (which is in fact going to be the zero symbol), and when it contains a bit it can also contain two distinct arrows from the symmetric generating set $S \not\ni 1_G$. The factor map is into the same alphabet. We refer to the $(s, s') \in S^2$ as the *forward* and *backward* arrows, respectively.

The factor map will only modify the bits. Its behavior is as follows: first, we identify *active edges*, namely pairs $(a, as) \in G$ such that the forward edge of x_a is s, and the backward edge of x_{as} is s^{-1}. The active edges form paths and cycles in the configuration.

Next, we identify *first cells.* These are those cells where a path of positive length through active edges begins; and *last cells* where such a path ends.

In the first cell of a path, we simply write down the current bit. In the last cell, if the forward arrow points to a cell with a bit b, and the backward arrow points to a cell with bit b', we write $b \oplus b'$, where $\oplus$ is addition modulo 2. In other active cells – and also in the last cell if it point to # – we add the bit in the present cell to the bit in the cell that points to the present cell.

It is easy to check that in the image, on every path of active cells, if the last cell points to a bit then the sum of the bits on the path, and the bit the last cell points to, is 0. (If the last cell points to #, there is no restriction.)

The sum constraint on paths where the last cell points to a bit (which does not continue the path) is in fact the exact subshift constraint defining the image. Thus, the subshift is of type Z0 (with # the 0 symbol), so it is contractible. By considering long geodesic segments (obtained from Lemma 2.2) connected to paths using arrows (and writing # elsewhere), we can easily find pseudo-orbits that are not actual orbits, so the image is proper sofic. $\square$

8.2 Contractible SFTs have FEP

In this section, we show that contractibility and finite type together imply the finite extension property. We will see later in Theorem 8.15 that FEP does not imply contractibility.

Recall that once we have fixed a window W for an SFT, *locally valid* for a pattern means there are no forbidden patterns in it (of shape W), and *globally valid* means it extends to a full configuration in the SFT.

Recall that the *finite extension property* or *FEP* [6] for an SFT X means that for some window size W, there exists $1_G \in N \Subset G$ such that if a pattern $p \in \Sigma^D$ admits a locally valid extension to DN, then p is globally valid.

Theorem 8.5. *Every contractible SFT has the finite extension property.*

Note that there is no standard definition of finite extension property for subshifts that are not of finite type, so in that sense the SFT assumption cannot be dropped.

Proof. Let X be of finite type, and fix a symmetric window W for it, with forbidden patterns $\mathcal{F} \subset \Sigma^W$. Let X be contractible with homotopy $h : I \times X \times X \to X$ with symmetric neighborhood R. Let $\tilde{h} : I \times \Sigma^G \times \Sigma^G \to \Sigma^G$ be the natural extension.

Claim 8.6. *Under the above assumptions, for any $x, y \in A^G$, if $\tilde{h}(t,x,y)|aW \in \mathcal{F}$, then one of the following holds, where $K = aWR$:*

1. *$t|K \equiv 0$ and $x|aW$ is a forbidden pattern of X;*
2. *$t|K \equiv 1$ and $y|aW$ is a forbidden pattern of X; or*
3. *$t|K \notin \{0^K, 1^K\}$, and either $x|K$ or $y|K$ is not globally valid in X.*

Proof of claim. Suppose $\tilde{h}(t,x,y)|aW$ contains a forbidden pattern. If $t|K \equiv 0$, then $\tilde{h}$ just copies the left argument, so this forbidden pattern is $x|aW$ and we are in the first item. If $t|K \equiv 1$, we are similarly in the second item. If t is not constant, we must show $x|K$ or $y|K$ is not globally valid in X.

Suppose the contrary, and let $x', y' \in X$ satisfy $x'|K = x|K, y'|K = y|K$. Then $h(t, x', y')|aW = \hat{h}(t, x', y')|aW = \hat{h}(t, x, y)|aW$ is a forbidden pattern, in particular $h(t, x', y') \notin X$, contradicting the type of h. ■

We now show that under the above assumptions, X is indeed FEP. For the forbidden patterns of the FEP we will use the window WR and as forbidden patterns $\mathcal{F}'$ we pick patterns $p \in A^{WR}$ which are not globally valid in X. As the FEP gap we will use $N = WR \cup R^2W^2R$.

If X is empty, the FEP trivially holds, so suppose X is nonempty, and let $y \in X$. Let $M \subset G$ be arbitrary, and let $Q \in A^{MN}$ be any pattern which has no subpattern from $\mathcal{F}'$. Note in particular that there are then no patterns from $\mathcal{F}$ inside the subpattern $Q|MR$: every W-shaped subset $aW \subset MW$ is contained in $aWR \subset MN$, and $Q|aWR \notin \mathcal{F}'$ implies that the subpattern $Q|aW \notin \mathcal{F}$.

Now, to prove FEP, it now suffices to show that $P = Q|M$ is globally valid in X. For this, let $x \in A^G$ be any configuration with $x|MN = Q$. Define t by $t|MR \equiv 0$, $t|(G \setminus MR) \equiv 1$, and consider $\tilde{h}(t, x, y) = z \in A^G$. We have $z|M = P$ because $t|MR \equiv 0$ and by our choices of x and the extension $\hat{h}$. Thus it suffices to show that z is in X.

Suppose on the contrary that $z \notin X$. Then $z|aW \in \mathcal{F}$ for some $a \in G$. Let $K = aWR$ and consider the cases of the previous lemma. The second case (which says $t|K \equiv 1$ and $y|aW \in \mathcal{F}$) is impossible, since $y \in X$. The first case (which says $t|K \equiv 0$ and $x|aW \in \mathcal{F}$) is also impossible: $t|K \equiv 0$ means that $K = aWR \subset MR$ so in particular $aW \subset MR$, but then $x|aW = Q|aW$ would contain a pattern from $\mathcal{F}$ in the subset MN, which we argued above is impossible.

Now we look at the third case, i.e. $t|K \not\equiv 0$, $t|K \not\equiv 1$, and either $x|K$ or $y|K$ is not globally valid in X. Of course it must be $x|K$ that is not globally valid. Since $K = aWR$ contains a 0, $aWR \cap MR \neq \emptyset$ so $a \in MR^2W$. But then Q has a globally forbidden pattern in a translate of WR contained in $K = aWR \subset MR^2W^2R \subset MN$. But this contradicts the assumption on Q. □

The reason FEP was originally introduced was the following theorem of Briceño, McGoff and Pavlov [6]:

Theorem 8.7. *Let $X \subset \mathbb{Z}^d$ be a block gluing subshift, and Y an FEP subshift with a fixed point, such that Y contains a fixed point. Then X factors onto Y.*

Here, recall that a $\mathbb{Z}^d$-subshift X is *block gluing* [5] if there exists $r \geq 0$ for any $x, y \in X$, if A and B are half-spaces separated by at least r, then there is a point $z \in X$ with $z|A = x|A, z|B = y|B$.

The fixed point is only used to guarantee the existence of a block map $f : X \to Y$. The use of FEP is more subtle, and seems to use the fact $\mathbb{Z}^d$ has finite asymptotic dimension. The point of block gluing and $h(X) > h(Y)$ is to guarantee that some patterns of suitable form in X are able to encode patterns of Y. This may be considered as simply a natural way to ensure that there exists a map $g : X \to A^{\mathbb{Z}^d}$ to a full shift, whose image contains Y.

With this understanding, the following theorem can be seen as a variant of the above theorem. In this generality, the existence of a block map from $f : X \to Y$ and a block map $g : X \to A^G$ such that $g(X) \supset Y$ do not have obvious general conditions (note that we are not even assuming soficity of G, so $h(X) > h(Y)$ is meaningless).

Theorem 8.8. *Suppose $Y \subset A^G$ is a contractible SFT. Then the following are equivalent:*

- *there is a block map from $f : X \to Y$ and a block map $g : X \to A^G$ such that $g(X) \supset Y$,*
- *there is a factor map $f : X \to Y$.*

Proof. Let $h : \{0,1\}^G \times Y \times Y \to Y$ be the homotopy, with symmetric neighborhood M. Let $\tilde{h} : \{0,1\}^G \times A^G \times A^G \to A^G$ be any extension that still satisfies $\tilde{h}(0^G, x, y) = x$ and $\tilde{h}(1^G, x, y) = y$ and has the same neighborhood. Let $F \subset A^N$ be the forbidden patterns of Y. Let N be a symmetric window for Y. Given $x \in X$, define $t(x) \in \{0,1\}^G$ by

$$\forall a \in G : (t(x)_a = 0 \iff g(x)|aMNM \text{ has no forbidden patterns from } F).$$

We now claim that

$$x \mapsto \tilde{h}(t(x), f(x), g(x))$$

is a factor map onto Y. First, it is surjective onto Y because if $y \in Y$, then $g(x) = y$ for some x. Then $g(x)$ does not contain forbidden patterns of Y so $t(x) = 1^G$, meaning $\tilde{h}(t(x), f(x), g(x)) = \tilde{h}(1^G, f(x), g(x)) = g(x)$.

We now explain why the image is contained in Y. Suppose it is not, say $\tilde{h}(t(x), f(x), g(x))_{aN} \in aF$. Then some symbol in aN must come from $g(x)$ (since $f : X \to Y$). Thus $t(x)_b = 1$ for some $b \in aNM$ (as $\tilde{h}$ would copy from $f(x)$ otherwise). In particular, by the definition of t, $x|bMNM \supset g(x)|aM$ does not contain any forbidden patterns of Y. In this case, $\tilde{h}(t(x), f(x), g(x)) = h(t(x), f(x), g(x)) \in Y$ since $\tilde{h}$ extends h. □

8.3 Separations

We have seen that retractions of contractible subshifts are contractible. It is tempting to think that a general factor of a contractible SFT already has nice dynamical properties. A contractible SFT is in particular strongly irreducible, so these factors are strongly irreducible.

In this section we give some indication that one cannot say much more than this, at least using the classes discussed in the present paper. Namely we show:

Theorem 8.9. *On the groups $\mathbb{Z}^d$, $d \geq 2$, we have the chain of inclusions*

$$\textit{contractible SFTs} \subsetneq \textit{FEP} \subsetneq \textit{SFT factors of contractible SFTs} \ (\subset \textit{SI SFTs}).$$

Note that the first inclusion was already shown in Section 8.2 (and holds for all groups). The second inclusion is immediate from [6] for FEP subshifts with a fixed point. We give a general proof in Lemma 8.13 (for groups $\mathbb{Z}^d$ only). The last inclusion is obviously true for all groups, and is included for context only. We suspect this last inclusion is strict as well. At least it cannot be proved non-strict with current knowledge, since if it is not strict, then SI $\mathbb{Z}^3$-SFTs have dense periodic points (which is a well-known open problem).

The strictness of the second inclusion is more generally shown for all groups which have a subgroup that is infinite, finitely-presented and one-ended (and the contractible SFT is a full shift). We also show that on any infinite group that is not just-infinite, there is a separation of the first and third class, i.e. a contractible SFT with a non-contractible SFT factor.

We need some basic closure properties of contractible subshifts.

Lemma 8.10. *If X is a contractible H-shift then the free extension to a G-shift is contractible.*

Proof. We simply apply the homotopy of X on each coset separately.

More precisely, let $h : I \times X \times X \to X$ be the contraction homotopy. Then we define $h' : I' \times X^{G/H} \times X^{G/H} \to X^{G/H}$ (where $I' = \{0,1\}^G$) by

$$h'(t,x,y)|rH = rh(r^{-1}(t|rH), r^{-1}(x|rH), r^{-1}(y|rH))$$

for all $r \in G$. Here, we note interpret $z|rH$ as a pattern with support rH, and the shift $r^{-1}(z|rH)$ has support H. We can see such a pattern again as a configuration, and for $z \in \{x,y\}$ it is clearly in X by the definition of the free extension. Now h produces an H configuration, so again seeing it as a pattern and shifting it by r produces the contents of rH as required in the formula.

It is clear that this is well-defined, in the sense that the definition of the contents of rH does not depend on the choice of left coset representative r. We give a calculation for completeness: if $rH = r'H$, then if $a \in rH$, we can write $a = rb = r'b'$ and then

$$\begin{aligned}
(h'(t,x,y)|rH)_a &= rh(r^{-1}(t|rH), r^{-1}(x|rH), r^{-1}(y|rH))_{rb} \\
&= h(r^{-1}(t|rH), r^{-1}(x|rH), r^{-1}(y|rH))_b \\
&= b'b^{-1}h(r^{-1}(t|rH), r^{-1}(x|rH), r^{-1}(y|rH))_{b'} \\
&= h(b'b^{-1}r^{-1}(t|rH), b'b^{-1}r^{-1}(x|rH), b'b^{-1}r^{-1}(y|rH))_{b'} \\
&= h(r'^{-1}(t|rH), r'^{-1}(x|rH), r'^{-1}(y|rH))_{b'} \\
&= h(r'^{-1}(t|r'H), r'^{-1}(x|r'H), r'^{-1}(y|r'H))_{b'} \\
&= r'h(r'^{-1}(t|r'H), r'^{-1}(x|r'H), r'^{-1}(y|r'H))_{r'b'} \\
&= (h'(t,x,y)|r'H)_a
\end{aligned}$$

We have

$$h'(\bar{0},x,y)|rH = rh(r^{-1}(\bar{0}|rH), r^{-1}(x|rH), r^{-1}(y|rH)) = rr^{-1}(x|rH) = x|rH$$

and symmetrically for $\bar{1}$, so $h'(\bar{0},x,y) = x$ and $h'(\bar{1},x,y) = y$. Continuity of h' is also clear. For shift-commutation,

$$\begin{aligned}
(bh'(t,x,y))|rH &= b(h'(t,x,y)|b^{-1}rH) \\
&= b(b^{-1}rh(r^{-1}b(t|b^{-1}rH), r^{-1}b(x|b^{-1}rH), r^{-1}b(y|b^{-1}rH))) \\
&= rh(r^{-1}b(t|b^{-1}rH), r^{-1}b(x|b^{-1}rH), r^{-1}b(y|b^{-1}rH)) \\
&= rh(r^{-1}((bt)|rH), r^{-1}((bx)|rH), r^{-1}((by)|rH)) \\
&= h'(bt,bx,by).
\end{aligned}$$

□

Lemma 8.11. *If X, Y are contractible then $X \times Y$ is contractible.*

Proof. If $h_X : I \times X \times X \to X$ and $h_Y : I \times Y \times Y \to Y$ are contraction homotopies, then $h : I \times (X \times Y) \times (X \times Y) \to (X \times Y)$ defined by

$$h(t, (x, y), (x', y')) = (h_X(t, x, x'), h_Y(t, y, y'))$$

is easily verified to be a contraction homotopy. $\square$

We also need the following. The SFT case is in [15]. We leave the FEP case to the reader.

Lemma 8.12. *If X is an H-subshift and $H < G$, then X is SFT (resp. FEP) if and only if $X^{G/H}$ is SFT (resp. FEP).*

Lemma 8.13. *Let $G = \mathbb{Z}^d$. Then every FEP subshift Y on G is a factor of a contractible SFT.*

Proof. Let R be the FEP gap, i.e. if a pattern on domain $W + B_R$ is locally valid (with respect to some fixed finite set of forbidden patterns), then its restriction to W is globally valid.

Let r' be large. By Lemma 6.6, $Y_{D,r',m}$ contains a contractible SFT Z for large m. The product of Z with a large full shift is still a contractible SFT.

As a factor of $Y_{D,r',m}$, we have $X_{D,r,m'}$ for sufficiently large m', by Lemma 6.4, and r is still large if r' is. We show that $X_{D,r,m'} \times A^G$ for a large alphabet A factors onto Y. This is analogous to the proof in [6], so we give only the outline. (The same idea is also used in Theorem 6.13.)

We proceed in steps $i = 1, \ldots, D+1$. On step i, we have a partially specified configuration, where for all $j < i$, we have already specified valid contents for the $(r - iR)$-neighborhood of each j-component. Furthermore, at all times the partially specified configuration must have a globally valid extension.

In the inductive step, we pick new contents for each i-colored component, as follows. We note that the $(d+1)r$-neighborhoods around the i-components are disjoint (and separated by distance at least R). In the area not yet determined in previous steps, we now simply pick any locally valid contents. This is possible because the configuration is globally valid by the inductive assumption, and we can perform the choice in a shift-invariant continuous way since the components are bounded and disjoint.

Now we shrink the determined area by R steps, to ensure that the remaining configuration is globally valid (using the FEP assumption). Note that in the induction step, we allow the determined area to shrink by R, so indeed the induction stays valid. $\square$

Theorem 8.14. *Suppose G infinite f.g. and not just-infinite. Then G admits a contractible SFT which admits a non-contractible SFT factor.*

Proof. Let $a \in G \setminus \{1_G\}$. Take finite generating set for G containing a. Use colors $\{0, 1, 2\} \times \{0, 1, 2\}$. Require that on the first track 0 and 2 cannot be next to each other, and that on second track you cannot have same color between a-neighbors. Call the resulting SFT X. It is a Cartesian product of two SFTs. The first one has safe symbol 1 so it is contractible. The second is a free extension of a mixing $\mathbb{Z}$-SFT, so it is contractible.

Now project the second track away when the first track contains 0 or 2. The image subshift Y over alphabet $\{0, 2\} \cup \{(1, 0), (1, 1), (1, 2)\}$ contains precisely those configurations where on the first track 0 and 2 cannot be next to each other, and on the second, on top of two a-neighboring 1s we require values to differ (here one should note that the mixing distance of the one-dimensional SFT we are freely extending is 1). This is clearly an SFT.

Since G is not just-infinite, a can be chosen so that the normal closure H of a is not of finite index in G. We claim that in this case, Y is not contractible. Suppose the contrary, and let $h : I \times Y \times Y \to Y$ be the contraction homotopy. Suppose $t \in \{0,1\}^G$ is constant on cosets of H, and consider $x = h(t, \bar{0}, \bar{2})$. Since H is normal, every shift of t is a-periodic, therefore also every shift of x is a-periodic.

Now suppose there are infinitely many cosets for H. Let the neighborhood of h be N. Let $Ha_1, ..., Ha_k$ be the cosets of H intersecting N. Let $t_b = 0$ whenever $b \in \bigcup_i Ha_i$, and $t_b = 1$ otherwise. Now $x = h(t, \bar{0}, \bar{2})$ of course contains 0 at the origin on the first track. We claim $\pi_1(x)_b = 2$ for some $b \in G$. If this is not the case, then $1 \sqsubset t|bN$ for all $b \in G$. Thus, bN intersects one of the finitely many cosets Ha_i for every $b \in G$. In particular, for every $b \in G$, $b \in Ha_i c$ for some $c \in N^{-1}, i = 1, \dots, k$, meaning $G \subset \bigcup_{i \leq k, c \in N^{-1}} Ha_i c$ and in fact H has finite index.

Since x contains both 0 and 2 on the first track, it must contain 1 on the first track as well, by the first rule of the SFT Y. Let b be such that $bx_{e_G} = (1, \alpha)$. Then since bx is H-periodic, $abx_{e_G} = bx_{a^{-1}} = (1, \alpha)$, contradicting the SFT rule. $\square$

Theorem 8.15. *If $|A|$ is large enough, then $A^{\mathbb{Z}^d}$ admits an SFT factor which is FEP but not contractible.*

Proof. By the result of [6], every block-gluing subshift with strictly greater entropy than X factors on to X, if X has the FEP and a fixed point. The subshift X from the previous theorem has the single-site filling property if we replace the second $\{0, 1, 2\}$ track with a large alphabet. Then, it has the FEP as well. This proves the result when $|A|$ is sufficiently large. $\square$

In fact, every non-trivial full shift $A^{\mathbb{Z}^d}$ admits an SFT factor which is FEP but not contractible – for this it suffices to find non-contractible FEP subshifts with arbitrarily small entropy and with a fixed point. For this, we can reduce the entropy in the previous construction for example by requiring that symbols 0 and 2 all belong to large monochromatic areas, and replacing the binary full shift on 1-symbols by a low-entropy mixing one-dimensional SFT. We leave the details to the interested reader.

We next give an example that shows that SFT factors of full shifts need not be FEP. This gives another construction of SFT factors of full shifts which are not contractible, in light of Theorem 8.5.

Theorem 8.16. *Suppose H contains a group G which is infinite, finitely presented and one-ended. Then any non-trivial full shift A^H admits an SFT factor which is not FEP.*

Proof. It suffices to show that if G is infinite, finitely presented and one-ended, then any non-trivial full shift A^G admits an SFT factor which is not FEP. Namely, if H contains such a group, then we can apply said factor cosetwise to get that $(A^G)^{H/G} = A^H$ has factor $Y^{H/G}$ where Y is SFT and not FEP. Then $Y^{H/G}$ is SFT but not FEP by Lemma 8.12.

Furthermore, we may assume the alphabet is $A = \{0,1\}$ by projecting away unnecessary symbols. We now construct a factor map from A^G. One should think of the image configurations as having binary values on edges of the Cayley graph. To code this into a literal subshift, we can take alphabet $\{0,1\}^S$ for the image subshift, where $S \not\ni 1_G$ is a symmetric generating set for G, and we will respect the SFT constraint $(y_a)_s = (y_{as})_{s^{-1}}$ for all $s \in S$ and $a \in G$.

We now simply write the differences modulo 2 on the edges, i.e. if $x \in A^G$ then $(f(x)_a)_s = 1 \iff x_a \neq x_{as}$ (note that then $(f(x)_{as})_{s^{-1}} = 1 \iff x_{as} \neq x_a$ so indeed the SFT constraint is satisfied). For algebraic purposes it is more useful to write $(f(x)_a)_s = x_{as} - x_a$.

We claim that the image is SFT, namely we check that the SFT rule defining the image is just that the bits on every cycle sum to 0, and finite presentation then means it suffices to constrain parity of finitely many sums.

We check this in more detail. Let $p : \mathbb{Z}_k \to G$ be any cycle such that $p(i)^{-1}p(i+1) \in S$ for all $i \in \mathbb{Z}_k$. Then any $f(x)$ satisfies

$$\sum_i f(x)_{p(i)} = \sum_i (x_{p(i+1)} - x_{p(i)}) = 0$$

by telescoping.

Let $R \subset S^*$ be a finite set of relations defining G. The geometric meaning is that every $w \in R$ represents a cycle in the Cayley graph of G (when we follow successive prefixes of w), and every cycle in the Cayley graph can be written as a concatenation of such cycles, started at different positions, and taking into account cancellations when a path literally moves backwards.

We can then define an SFT where cycles corresponding to $w \in R$ must sum to 0, and it is then clear that in this SFT, over any cycle, we will see sum 0, by presenting the cycle as a succession of cycles of shape R. It follows that f has SFT image.

We claim that the image subshift is not FEP. Suppose on the contrary that $N \Subset G$ is an FEP window.

Take a geodesic path $p : \mathbb{Z} \to G$ with $\forall i : d(p(i), p(i+1)) = 1$ (so in particular $p(i)^{-1}p(i+1) \in S$). Such a path exists by Lemma 2.2, and we may assume $p(0) = 1_G$. Now take a large n and observe that $p(i) \in B_n$ (the ball of radius n around the origin) if and only if $i \in [-n, n]$. Both $p(-n-1)$ and $p(n+1)$ have infinite connected components in $G \setminus B_n$, thus they correspond to the same end, and we can join them by a finite path in the subgraph of the Cayley graph of G induced by $G \setminus B_n$. This gives us a path $q : \mathbb{Z}_k \to G$ with $q(i)^{-1}q(i+1) \in S$ for all i, which agrees with p when $i \in [-n, n]$, and stays outside of B_n on other steps $i \in \mathbb{Z}_k$.

Now consider the configuration $x \in A^G$ defined by $x_a = 1$ if for all large enough i we have $d(p(i), a) \leq i$ (this is the indicator function of the Busemann horoball corresponding to p). Clearly $x_{p(i)} = 1 \iff i \geq 0$. Thus on the edges corresponding to the path $p(i)$, we have a single 1-symbol in the image $f(x)$.

Consider the restriction of $f(x)$ to the tube around the path p, i.e. $y = f(x)|p(\mathbb{Z})N$, and observe that since word norm is 1-Lipschitz, y has finite support. Now consider the configuration $z \in (A^S)^G$ where $z|q([-n,n])N = y|q([-n,n])N$ and $(z_a)_s = 0$ for $a \in G$, $s \in S$ not specified by this rule. Then z is a pseudo-orbit, since it corresponds to $y \in f(A^G)$ in the inner tube $z|q([-n,n])N$, and in the connecting path it corresponds to the (f-image of the) constant zero configuration.

It does not correspond to an actual orbit even along the path $q(\mathbb{Z}_k)$, since the sum of values on the edges of this cycle is 1. This contradicts the FEP. $\square$

The proof of the previous theorem is clearly cohomological, and one can replace finite presentability with the finiteness property that after adding finitely many G-orbits of 2-cells corresponding to relations, the $\mathbb{Z}_2$-valued 2-cohomology becomes trivial, i.e. all cocycles are coboundaries (coboundaries corresponding exactly to configurations with an f-preimage).

In the one-dimensional case, all the classes mentioned in Theorem 8.9 coincide.

Proposition 8.17. *Let $G = \mathbb{Z}$. Then an SFT is a strongly irreducible if and only if it is a factor of a contractible SFT if and only if it is FEP if and only if it is contractible.*

Proof. It suffices to observe that a one-dimensional SFT is mixing if and only if it is strongly irreducible, as we showed in Theorem 5.18 that mixing is also equivalent to contractibility. $\square$

Example 1: There is a sofic $\mathbb{Z}$-shift which is strongly irreducible but not contractible. Namely, take alphabet $\{0, 1, 2\}$ and forbid words $2bw2$ where $w \in \{0, 1\}^*$ and $b \equiv \sum_i w \bmod 2$. $\bigcirc$

Conjecture 8.18. *On a large class of groups, not every SI SFT is an SFT factor of a contractible SFT.*

As suggested in the beginning of this section, it would be difficult to refute this conjecture even for $\mathbb{Z}^3$. Namely, factors of contractible SFTs (even general contractible subshifts) have dense periodic points by Theorem 6.10, so showing that their SFT factors are SI would imply that SI SFTs have dense periodic points, which is a well-known open problem. We suspect, however, that there is a simple way to separate these classes.

8.4 Map extension property

Write $X \rightsquigarrow Y$ if for all $x \in X$ there exists $y \in Y$ with $\mathrm{stab}(x) \leq \mathrm{stab}(y)$. This is just the obvious necessary periodic point condition for the existence of a morphism from X to Y.

Recently [14], Meyerovitch introduced the following notion.

Definition 8.19. *A G-subshift Y has the* map extension property *if:*

1. *Let X be any G-subshift such that $X \rightsquigarrow Y$. Then for any morphism $g : X' \to Y$ with $X' \subset X$ a subshift, there exists a morphism $f : X \to Y$ that extends g.*

2. *There exists a finite set $\mathcal{G} \subset \mathrm{Sub}(G)$ of non-trivial subgroups of G, such that for every $G_0 \leq G$ either there exists $y \in Y$ with $G_0 \leq \mathrm{stab}(y)$ or there exists $G_1 \in \mathcal{G}$ such that $G_1 \leq G_0$.*

(Meyerovitch writes that X' is a closed shift-invariant set, instead of a subshift, since his subshifts are required to be non-empty.) The motivation for introducing this class is quite similar to ours, and this class also has many similar properties as the contractible subshifts.

In this section, we give some explicit connections between the first item of this definition, and contractibility. First, we isolate a piece of the above definition's first item.

Definition 8.20. *We say Y has the* map existence property *if for any subshift X with $X \rightsquigarrow Y$, there is a morphism $f : X \to Y$.*

We state the main results of this section, and then prove the lemmas needed in the proofs.

Theorem 8.21. *On any finitely generated infinite group G, the following are equivalent for a subshift X:*

1. *X satisfies the first item of Definition 8.19;*
2. *X is a contractible SFT with the map existence property;*
3. *X is an equiconnected SFT with the map existence property.*

Proof. ($1 \implies 3$): Suppose X satisfies the first item of Definition 8.19. Then by Lemma 8.23 below it is an equiconnected SFT. A subshift with the map extension property has the map existence property because we may take $X' = \emptyset$ in the definition.

($3 \implies 2$): An equiconnected subshift is clearly contractible.

($2 \implies 1$): A contractible SFT with the map existence property satisfies the first item of Definition 8.19 by Lemma 8.24. ☐

Theorem 8.22. *On the group $\mathbb{Z}^d$ with $d \geq 2$, the following are equivalent for a subshift X:*

1. *X satisfies the first item of Definition 8.19;*
2. *X is a contractible SFT;*
3. *X is equiconnected.*

Proof. The group $\mathbb{Z}^d$ has the patching property, so being equiconnected is equivalent to being equiconnected and SFT. By Lemma 8.27, every contractible SFT (thus also every equiconnected SFT) has the map existence property. Thus, the items are equivalent to the items of the previous theorem for the groups $\mathbb{Z}^d$. ☐

Lemma 8.23. *Every subshift satisfying the first item in the definition of the map extension property is an equiconnected SFT.*

Proof. Assume Y has the map extension property. Proposition 4.22 in [14] shows Y is of finite type. We show equiconnectedness. Take $X = I \times Y \times Y$ and observe that $\pi_2 : X \to Y$, so certainly $X \rightsquigarrow Y$.

Let

$$X' = (\{\bar{0}, \bar{1}\} \times Y \times Y) \cup (I \times \Delta_Y)$$

and define $g : X' \to Y$ by the formulas $g(\bar{0}, x, y) = x, g(\bar{1}, x, y) = y, g(t, x, x) = x$. Then the extension $h = f : X \to Y$ provided by the definition of the map extension property proves equiconnectedness. ☐

Lemma 8.24. *Every contractible SFT with the map existence property satisfies the first item in the definition of the map extension property.*

Proof. Assume Y is a contractible SFT with symmetric window N, let $h : I \times Y \times Y \to Y$ be a contraction homotopy with symmetric neighborhood F, and take the natural extension $\tilde{h} : I \times A^G \times A^G \to A^G$ of h, so $\tilde{h}$ has the same neighborhood and still satisfies $\tilde{h}(\bar{0}, x, y) = x, \tilde{h}(\bar{1}, x, y) = y$.

Suppose $X \rightsquigarrow Y$. Let $c : X \to Y$ be some morphism, guaranteed by $X \rightsquigarrow Y$ and the map existence property. Let $X' \subset X$ be any subshift, and let $g : X' \to Y$ be any morphism with neighborhood M. We must show that g extends to a morphism $f : X \to Y$.

Let $W \Subset G$ and let X_W be the SFT approximation of X' with window W, and $\tilde{g} : B^G \to A^G$ the morphism using the same local rule as g (extended arbitrarily). By Lemma 2.3, if W is large enough then $\tilde{g}$ maps X_W into Y. Using Lemma 2.4, let D be such that if $x \in A^G$ does not contain a forbidden pattern of X_W (of shape W) in $x|D$, then $x|NFM$ is globally valid in X_W.

Define $f : X \to A^G$ by

$$x \mapsto \tilde{h}(t(x), c(x), \tilde{g}(x))$$

where $t(x)_a = 1$ if $x|aFND$ does not contain forbidden patterns of X_W. If $x \in X'$, then $x \in X_W$ so $t(x) = \bar{1}$, and thus $f(x) = \tilde{h}(\bar{1}, c(x), g(x)) = g(x)$.

Now we are left with showing that the f-image of X is contained in Y. Suppose not, and say $\tilde{h}(t(x), c(x), \tilde{g}(x))|aN$ is forbidden in Y. Then $t(x)|aNF$ must contain the symbol 1, as otherwise $\tilde{h}$ would have copied the contents from $c(x) \in Y$. This means that for some $h \in NF$, $x|ahFND \supset x|aD$ does not contain any forbidden pattern of X_W. This means that $x|aNFM$ is globally valid in X'. Thus, $\tilde{g}(x)|aNF$ cannot be a forbidden pattern of Y since $\tilde{g}$ has neighborhood M and maps X_W into Y.

This contradicts the assumption that the image of f can have forbidden patterns of Y, and we conclude that the codomain restriction $f : X \to Y$ is an extension of g. $\square$

We now show that contractible SFTs on abelian groups have the map existence property.

We need the usual marker lemma [12], which states that if a configuration locally does not have a period, then there is a shift-invariant way of dropping "markers" on it so that the configuration of markers does not have a period. The following version is proved in [14]. We only need the abelian case, and Meyerovitch also proves such version, but we state the general version here and then deduce our own abelian statement.

Lemma 8.25. *Let $X \subset A^G$ be a subshift. Let $P \Subset G$ be a finite symmetric set with $1_G \notin P$, and let $V \subset X$ be a clopen set of all $x \in X$ such that $ax \neq x$ for every $a \in P$. Then there exists a clopen set $C \subset V$ so that*

$$C \cap aC = \emptyset \text{ for every } a \in P, \text{ and } V \subset C \cup \bigcup_{a \in P} aC$$

The clopen set C is referred to as a *marker*. Our viewpoint is that the characteristic morphism $g = \chi_C : X \to \{0,1\}^G$ defined by $g(x)_a = 1 \iff a^{-1}x \in C$ identifies a set of preferred positions (also called *markers*) in G, and these markers are picked in a shift-invariant continuous way.

We state an abelian version in this terminology. If $x \in A^G$ with G abelian, we say a is a *period-breaker* for $b \in G$ if $x_a \neq x_{ab}$.

Lemma 8.26. *Let G be an abelian group, let $X \subset A^G$ be a subshift, and let $r, R \in \mathbb{N}$ be arbitrary. Then there exists a morphism $g : X \to Y$ such that for all $x \in X$, $(g(x))^{-1}(1)$ is an r-packing (in G); and if x has the property that for every $\vec{v} \in B_r \setminus \{\vec{0}\}$, there is a period-breaker in the set B_R, then $(g(x))^{-1}(1)$ is an r-covering of B_R.*

Proof. Pick $P = B_r \setminus \{\vec{0}\}$ and let V be the set of configurations that contain period breakers for all $\vec{v} \in P$ in the ball B_{2R}. Apply the previous lemma and set $g = \chi_C$ where C is the clopen set in the conclusion of the lemma.

The property

$$C \cap aC = \emptyset \text{ for every } a \in P$$

means that $(g(x))^{-1}(1)$ is always an r-packing. Now consider the property

$$V \subset C \cup \bigcup_{a \in P} aC.$$

If in x, we can break all P-periods in the R-ball, then we can break all of them in the $2R$-ball around any $a \in B_R$. This means $ax \in V$ for all $a \in B_R$, and then the condition means that for any $a \in B_R$, $g(x)$ has a 1 somewhere in aB_r. □

Note also that one can alternatively take V to be the set of configurations that see period-breakers in B_R, and then use the trick from the proof of Lemma 6.6 to spread the markers to the surrounding R-ball while preserving the r-packing property.

Lemma 8.27. *Every contractible G-SFT with G finitely-generated abelian has the map existence property.*

Proof. Suppose $Y \subset B^G$ is a contractible G-SFT. We may inductively suppose that for all proper quotients of G, the claim is true, since f.g. abelian groups have the ascending chain condition for subgroups, and if $K \leq G$ is normal (which for abelian G is automatic) then the K-periodic points of a contractible SFT are themselves a contractible SFT.

Suppose $X \rightsquigarrow Y$. For G the trivial group, the claim is trivial: every $x \in X$ is simply mapped to any configuration in Y.

For the general case, pick a symmetric generating set S. Let r' be large. By default, we will apply the previous marker lemma (with $r = r'$) to identify an r'-net of markers (in the sense that the set is r'-separated and is r'-dense other than near the boundary).

In regions containing such markers, we look at the Voronoi cells of the markers. They give us patterns satisfying the SFT constraint of the subshift $Y_{D,r',m'}$ from Lemma 6.4 for sufficiently large D and m'. Note that D can be taken independent from r', by a volume argument and the doubling property of G.

We can now construct points of $X_{D',r,m}$ for some D' (again independent of r') using the proof of Lemma 6.4. As $r' \to \infty$, r can also be taken arbitrarily large. We then follow the proof of Lemma 8.13 to construct valid contents in Y, by proceeding one color at a time. This is possible since D' stays bounded while r can be taken arbitrarily large.

All in all, this discussion gives us a morphism $f_{\vec{0}} : X \to (B \cup \{\#\})^G$ such that, whenever it outputs a patch over alphabet B, this patch contains no forbidden patterns, and furthermore, when the marker lemma applies sufficiently nearby (say, at distance at most $R/2$), we actually produce a large patch over alphabet B.

In areas where the marker lemma does not apply, i.e. large areas with at least one small period, we apply induction. More precisely, for each $\vec{v}$, we have a factor map from $X/\langle\vec{v}\rangle$ to $Y/\langle\vec{v}\rangle$ where $Z/\langle\vec{v}\rangle$ denotes the $G/\langle\vec{v}\rangle$-subshift of $\vec{v}$-periodic points of a subshift Z.

By applying the local rule of such a morphism whenever we can and outputting # elsewhere, we have for each $\vec{v} \in P$ a morphism to the full shift with alphabet $B \cup \{\#\}$ which, when not outputting #, outputs locally valid content.

All in all, we have morphisms $f_{\vec{0}}$, and $f_{\vec{v}}$ for each $\vec{v} \in P$. We observe that without changing r' (thus with a fixed number $|P| + 1$ of morphisms), we can ensure that for each $\vec{u} \in \mathbb{Z}^d$, an arbitrarily large ball is covered in the non-# part of the image of at least one of the morphisms $f_{\vec{v}}$, $v \in B_{r'}$. For this, it suffices to increase the number R to make the morphism f look further for period-breakers before outputting #.

Thus the images of the morphisms cover G in the precise sense of Lemma 6.9, and the lemma guarantees the existence of a morphism from X to Y (with some additional properties that do not interest us here). □

9 The new group properties

In the main text, we abstracted away some group properties that we needed in proofs, namely the finite periodic asymptotic dimension, and the patching property. To our knowledge, these notions have not previously appeared in the literature. Here, we give examples of groups with these properties.

9.1 Groups with finite periodic asymptotic dimension

Theorem 9.1. *The periodic asymptotic dimension of $\mathbb{Z}^d$ is d.*

Proof. The usual proofs that asymptotic dimension of $\mathbb{Z}^d$ is at most d already use a periodic grid.

Let us recall this argument. We should prove that for all r, there exists m such that the d-dimensional subshift $X_{d+1,r,m}$ has a periodic point.

Let N be large, and for $\vec{v} \in \mathbb{Z}^d$, calculate maximal k such that there are at least k coordinates i such that $d(\vec{v}_i, N\mathbb{Z}) \leq kr$. Let C_k be the set of such vectors, and observe that $\{C_k \mid 0 \leq k \leq d\}$ is a partition of $\mathbb{Z}^d$.

Suppose $\vec{u}$ and $\vec{v}$ are in T_k and the ℓ_∞ distance between $\vec{u}$ and $\vec{v}$ is at most r. Let I be the k many coordinates such that $d(\vec{u}_i, N\mathbb{Z}) \leq kr$ for $i \in I$, and let J be the corresponding set for $\vec{v}$. Observe that for $i \notin J$, we must have $d(\vec{v}_i, N\mathbb{Z}) > (k+1)r$, or k is not maximal for $\vec{v}$.

Then we see that $I = J$: Namely if $i \in I \setminus J$, then we would have

$$kr \geq d(\vec{u}_i, N\mathbb{Z}) \geq d(\vec{v}_i, N\mathbb{Z}) - r > kr.$$

We conclude that any path $p : \mathbb{N} \to T_k$ that satisfies $d(p(i), p(i+1)) \leq r$ must satisfy that there is a set I of k coordinates such that $\vec{v} = p(i)$ has $d(\vec{v}_i, N\mathbb{Z}) \leq kr$ for $i \in I$ and $d(\vec{v}_i, N\mathbb{Z}) > (k+1)r$ for $i \notin I$.

Now if $N > 100(d+1)r$, it is easy to see that in each individual coordinate, the set of possible values it can reach is finite. Namely if $i \in I$, then we stay in a single hyperplane $U_i = \mathbb{Z}^{i-1} \times [cN - kr, cN + kr] \times \mathbb{Z}^{d-i-2}$, since such hyperplanes are at distance more than r for distinct c; and if $i \notin I$, then we stay in a single hyperplane $U_i' = \mathbb{Z}^{i-1} \times [cN + (k+1)r, (c+1)N - (k+1)r] \times \mathbb{Z}^{d-i-2}$, since such hyperplanes are at distance more than r for distinct c

This gives $\mathrm{pdim}(\mathbb{Z}^d) \leq d$. For the lower bound $\mathrm{pdim}(\mathbb{Z}^d) \geq d$, we simply observe $\mathrm{pdim}(G) \geq \dim(G) = d$ where the latter is well-known. $\square$

Lemma 9.2. *Let $H \leq G$ be f.g. groups. Then* $\mathrm{pdim}(H) \leq \mathrm{pdim}(G)$.

Proof. If $x \in X_{G,d,r,m}$, then clearly $x|H \in X_{H,d,r,m}$ (taking the generating set of G to contain the one used for H). $\square$

Lemma 9.3. *If $H < G$ is of finite index and both groups are finitely-generated, then* $\mathrm{pdim}(H) = \mathrm{pdim}(G)$.

Proof. The group H cannot have larger periodic asymptotic dimension than G by the previous lemma. We now show that G does not have larger periodic asymptotic dimension than H, which concludes the proof.

Let S be the finite generating set used for H, $S' \supset S$ the one for G, pick coset representatives $G = \bigsqcup_{t \in T} Ht$. For each $t \in T, s' \in S'$ write $ts' = s''t'$ with $s'' \in H, t' \in T$. Let ℓ be the maximal word norm of s'' that appear this way.

Now if $x \in X_{H,d,r,m}$ is periodic with stabilizer $K \leq H$, then we construct a point $y \in \{0, 1, \ldots, d\}^G$ by the formula $y_{hs} = x_h$ for $h \in H, s \in T$. Then that stabilizer of y still contains K which is of finite index in G.

We claim that $y \in X_{G,d,r',m'}$ for some r', m', where $r' \to \infty$ as $r \to \infty$. Namely, consider a monochromatic r'-path in y. Then we can construct a path in H by projecting ht to h for $t \in T, h \in H$. Consider an S'-move from ht to hts' with $s' \in S'$. We have $hts' = hs''t'$ with $s'' \in H$, $t' \in T$ with s'' of length at most ℓ. Thus, a monochromatic r'-path with S'-moves in G projects to an $r'm$-path in x with S-moves. So we can pick $r' = \lfloor r/\ell \rfloor$ and $m' = m|T|$. $\square$

Proposition 9.4. *Let $1 \to K \to G \to H \to 1$ be an exact sequence, and suppose K is locally finite. Then* $\mathrm{pdim}(G) = \mathrm{pdim}(H)$.

Proof. The claim $\mathrm{pdim}(G) \geq \mathrm{pdim}(H)$ is clear. Let's show the other direction.

For G use a generating set S which projects to the generating set used for H. For any r, let m be such that there is a periodic $x \in X_{H,d,r,m}$ where $d = \mathrm{pdim}(H) + 1$. Pull this back to a G-configuration through π by $y_a = x_{\pi(a)}$. Then y obviously has finite orbit. We claim that $y \in X_{G,d,r,m'}$ for some m'. For this, observe that while staying within a monochromatic component, the projection will stay in the same component of x. It suffices to give a bound on the size of the components in the fiber $F' = \pi^{-1}(F)$ where F is the monochromatic component of x that contains the identity.

Consider a path by right S-translations which stays stays in the fiber F'. Picking a lift of F and cancelling and reintroducing the right coset representative of elements encountered during the path, we get a new path which stays inside K, and moves by finitely many generators $T \Subset K$. Since K is locally finite, there is a bound m' on the components. This bound can be taken to only depend on the diameter of F, so we are done. $\square$

Corollary 9.5. *The lamplighter group $G = \mathbb{Z}_2 \wr \mathbb{Z}$ has* $\mathrm{pdim}(G) = 1$.

A similar idea works in general, but leads to some technical issues. We give a presumably highly non-optimal formula for the split exact case.

Theorem 9.6. *Let G be a countable group. If $1 \to K \to G \to H \to 1$ is split exact, then* $(\mathrm{pdim}(K)+1)(\mathrm{pdim}(H)+1)-1$.

Proof. Recall that in the case that a group K is not finitely generated, by r-paths we simply mean paths by jumps with sets S_r, where $K = \bigcup_r S_r$.

We need $\mathrm{pdim}(K)+1$ colors D for the colorings of K, and $\mathrm{pdim}(H)+1$ colors D' for the colorings of H. We will use colors $D \times D'$ for G to get the stated dimension formula.

Let r be given. We construct a periodic point y in $X_{G,D\times D',r,m}$ for some m. First, take a periodic point $x \in X_{H,D',r,m'}$ for sufficiently large m'. If $\pi : G \to H$ is the projection, set the D'-component of y_a to be $x_{\pi(a)}$.

Next, consider a fiber corresponding to a single finite r-component $F \Subset H$ of a color class of x. If we take $m \geq m'$, then all monochromatic r-components of elements in the fiber $\pi^{-1}(F)$ stay in this fiber. Thus we are effectively left with picking the D-colors in individual fibers, so that r-components within fibers are bounded, and the global configuration is periodic.

Let us consider the fiber F containing the identity element. We observe that K acts on $\pi^{-1}(1_H)$-configurations (i.e. K-configurations) by shifting. Picking any representatives R for fibers of $\pi^{-1}(f)$, $f \in F$ (i.e. coset representatives for cosets of K that fit inside $\pi^{-1}(F)$), we observe that for any configuration $z \in X_{K,D,r',m''}$ with period $N \leq K$, the configuration where we copy the color at k to all of kR is still N-periodic as a pattern with domain $\pi^{-1}(F)$ (in our usual sense of translation of patterns).

Furthermore, suppose that r' is larger than $m+2r$, where m is the maximal word norm of an element in B_r^g where $g \in R$ and B_r is the ball of radius r. Then any r-path in a monochromatic component of this new graph corresponds to an r'-path in a monochromatic component of the K-configuration z, since we can always conjugate the path back to K. So with this choice of r', the monochromatic components in $\pi^{-1}(F)$ are bounded by some m, which is a function of m''.

We can use this scheme to pick the D-symbols in any individual fiber $\pi^{-1}(F)$ of a monochromatic r-component $F \Subset H$, and it is then automatic that the r-components are bounded if we use only finitely many types of coloring m. We are left with doing this so that the global configuration is periodic.

Since the extension is split, there is a canonical way to perform a coloring, which will lead to periodicity. First, we may assume that the period lattice L used for x is normal and sparse enough, so that, as in the proof of Lemma 6.6, we can pick "centers" for the sets F so that the set of centers is L-periodic.

Now the scheme we use to color $\pi^{-1}(F)$ is that we take the center $h \in F$, see h as an element of G, translate it to the origin, color the fiber, and move it back.

The canonicality of this procedure implies that the resulting configuration is automatically L-periodic. We claim that its K-orbit is also finite. To see this, note that we use only finitely many different periodic points z in the construction, and thus we have only finitely many period lattices N. Let L' be any finite index characteristic subgroup of K which is contained in all of them.

Observe that a translation by k is effectively translation by $g^{-1}kg \in K$ in the fiber $\pi^{-1}(gF)$, in the sense that if the contents of $\pi^{-1}(gF)$ are obtained by taking $z \in D^{\pi^{-1}(F)}$ and translating it to $\pi^{-1}(gF)$ by shifting $z \mapsto gz$, then $kgz = g(g^{-1}kg)z$ meaning the new contents of $\pi^{-1}(gF)$ corresponds to the g-shift of $(g^{-1}kg)z$. For $k \in L'$, $(g^{-1}kg)z = z$ for the L'-periodic configurations used in the fibers.

We conclude that the stabilizer of the configuration we have built projects onto a finite index subgroup (namely L) of H, and contains a finite index subgroup (namely L') of K. It is a general group-theoretic fact that such a subgroup has finite index in G, thus we have a periodic point. $\square$

Corollary 9.7. *Virtually polycyclic groups have finite periodic asymptotic dimension.*

Proof. A strongly polycyclic group is inductively defined as a group extension G with exact sequence $1 \to P \to G \to \mathbb{Z} \to 1$ where P is itself strongly polycyclic or trivial. Since $\mathbb{Z}$ is a free group, any such extension splits, so the previous lemma and induction shows that they have finite periodic asymptotic dimension.

Now the result follows from Lemma 9.3, since a virtually polycyclic group by definition has a finite-index subgroup which is polycyclic, and it is well-known that a polycyclic group has a strongly polycyclic finite-index subgroup. $\square$

Corollary 9.8. *The Baumslag-Solitar group* $\mathrm{BS}(1,n)$ *has periodic asymptotic dimension* 2 *or* 3*.*

Proof. The group $\mathrm{BS}(1,n)$ is a semidirect product $\mathbb{Z}[\frac{1}{n}] \rtimes \mathbb{Z}$ where $\mathbb{Z}[\frac{1}{n}] = \bigcup_k \frac{1}{n^k}\mathbb{Z}$ (as a subgroup of $(\mathbb{Q}, +)$) is locally cyclic and thus has dimension 1, and $\mathbb{Z}$ of course also has dimension 1. The previous theorem gives the bound 3 for the dimension, and the well-known asymptotic dimension 2 gives a lower bound. $\square$

By looking more carefully at this group, it is not difficult to show that the dimension is indeed 2. Generally, one expects dimension to behave according to $\dim(A \times B) \leq \dim(A) + \dim(B)$, where $A \times B$ is some type of product of A and B.

It is known that in the case of asymptotic dimension of groups, such a formula holds for group extensions [3], and in the case of Lebesgue covering dimension of metric spaces, it holds for the direct product.

Our naive formula suffices for us because in our applications we are only interested in finiteness of the dimension. However, it seems likely that one can use standard dimension-theoretic tricks to improve on it.

Question 9.9. *Does periodic asymptotic dimension stay finite in general group extensions? Does the formula* $\mathrm{pdim}(G) \leq \mathrm{pdim}(K) + \mathrm{pdim}(H)$ *hold for periodic asymptotic dimension for (split or non-split) group extensions?*

Conjecture 9.10. *Nonabelian free groups have infinite periodic asymptotic dimension.*

Note that the usual asymptotic dimension of free groups is 1. The groups $\mathrm{BS}(m,n)$ with $1 < m \leq n$ contain nonabelian free subgroups, so if the conjecture is true, then they are also infinite-dimensional.

9.2 Groups with the patching property

We have no examples of groups that do not have the patching property from Definition 7.3. However, we give two (incomparable) classes of groups that do have it.

Lemma 9.11. *Let G be a group with finite asymptotic dimension. Then G has the patching property.*

Note that here we talk about asymptotic dimension rather than periodic asymptotic dimension, so this covers the free group.

Proof. Take $C_1 \cup \cdots \cup C_d$ a cover of the group so that each C_i has connected $2r$-components bounded, and where each element $a \in G$ has its $2rd$-ball contained in some C_i (for this start from slightly more separated sets C_i, and then artificially increase their sizes).

Then balls of radius B_R, where R comes from the bound on size of connected components, gives the patching property for r. First, any finite subset of the union of connected components in a given C_i can be constructed using almost-unions, since $B_r(A) \cap B_r(B) = \emptyset$ for distinct connected components A, B, so for distinct components

$$A \overset{\circ r}{\cup} ((A \cup B) \supset (A \cup B) \setminus (AB_r \cap BB_r)) = A \cup B$$

After growing any finite subsets of components from each C_i, in d steps of almost-unioning we get an arbitrarily large finite pattern in G from the fact that the almost-union of A and B contains $A^{\circ r} \cup B^{\circ r}$. □

Lemma 9.12. *Let G be a finitely-generated group with subexponential growth. Then G has the patching property.*

Proof. Let $r \in \mathbb{N}_+$ be arbitrary. Suppose that there exists k such that B_{R+1} can be covered by at most 2^k many left translates of B_{R-rk}, for all large enough R. Then we claim that R proves the patching property for R.

Namely, let us cover B_{R+1} with 2^k many translates $a_i B_{R-rk}$, $a_i \in G$. Now combining the corresponding translates $a_i B_R$ pairwise in a tournament-tree fashion, we get a common almost-union for them with at most $2^k - 1$ almost-unions, so that each individual $a_i B_R$ only partakes in an almost-union on k steps.

The almost-union of A and B contains $A^{\circ r} \cup B^{\circ r}$ and the operation $A \mapsto A^{\circ r}$ is superadditive in the sense that $A^{\circ r} \cup B^{\circ r} \subset (A \cup B)^{\circ r}$. From this, we deduce that the almost-union of the $a_i B_R$ constructed in the previous paragraph contains each $a_i B_R^{\circ rk} = a_i B_{R-rk}$. By assumption their union covers B_{R+1}, and by induction we can show that every B_{r+j} can be constructed.

Now observe that in a general finitely-generated group, we have $B_{R+1} = \bigcup_{h \in B_{rk+1}} hB_{R-rk}$ for all large enough R. So if the property fails, we must have $|B_{rk+1}| \geq 2^k$ for infinitely many k. In other words, the group has exponential growth.[2] □

Note that the free group has finite asymptotic dimension and exponential growth, while the Grigorchuk group has infinite asymptotic dimension and subexponential growth.

[2] Recall that the size of balls is submultiplicative, so $|B_{rk+1}| \geq 2^k$ for infinitely many k implies $B_k \geq \alpha^k$ for some $\alpha > 1$ and all $k \in \mathbb{N}$.

Acknowledgements

The notion of homotopy arose as a tangent of the second-named author's work with Thomas Worsch in 2016.

The observation that contractibility modulo diagonal corresponds to equiconnectedness was given to us by user Tyrone on MathOverflow [1].

Theorem 8.16 is roughly based on a geometric idea of Ilkka Törmä on the group $\mathbb{Z}^2$.

References

[1] Tyrone (https://mathoverflow.net/users/54788/tyrone). *Spaces that are contractible mod diagonal*. MathOverflow. URL:https://mathoverflow.net/q/457117 (version: 2023-10-25). eprint: https://mathoverflow.net/q/457117. URL: https://mathoverflow.net/q/457117.

[2] Laurent Bartholdi, Rostislav I. Grigorchuk, and Zoran Šuni. "Branch groups". In: ed. by M. Hazewinkel. Vol. 3. Handbook of Algebra. North-Holland, 2003, pp. 989–1112. DOI: https://doi.org/10.1016/S1570-7954(03)80078-5. URL: https://www.sciencedirect.com/science/article/pii/S1570795403800785.

[3] G. C. Bell and A. N. Dranishnikov. "A Hurewicz-Type Theorem for Asymptotic Dimension and Applications to Geometric Group Theory". In: *Transactions of the American Mathematical Society* 358.11 (2006), pp. 4749–4764. ISSN: 00029947. URL: http://www.jstor.org/stable/3845399 (visited on 12/18/2023).

[4] Rufus Bowen. "Periodic Points and Measures for Axiom a Diffeomorphisms". In: *Transactions of the American Mathematical Society* 154 (1971), pp. 377–397. ISSN: 00029947. URL: http://www.jstor.org/stable/1995452 (visited on 12/21/2023).

[5] Mike Boyle, Ronnie Pavlov, and Michael Schraudner. "Multidimensional sofic shifts without separation and their factors". In: *Transactions of the American Mathematical Society* 362.9 (2010), pp. 4617–4653.

[6] Raimundo Briceño, Kevin McGoff, and Ronnie Pavlov. "Factoring onto $\mathbb{Z}^d$ subshifts with the finite extension property". In: *Proceedings of the American Mathematical Society* 146.12 (2018), pp. 5129–5140.

[7] Nic Brody and Kasia Jankiewicz. "Generalized Residual Finiteness of Groups". In: *arXiv preprint arXiv:2312.15120* (2023).

[8] Robert Burton and Jeffrey E Steif. "Non-uniqueness of measures of maximal entropy for subshifts of finite type". In: *Ergodic Theory and Dynamical Systems* 14.2 (1994), pp. 213–235.

[9] Tullio Ceccherini-Silberstein and Michel Coornaert. "On the density of periodic configurations in strongly irreducible subshifts". In: *Nonlinearity* 25.7 (2012), p. 2119.

[10] Tullio Ceccherini-Silberstein and Michel Coornaert. "The Myhill property for strongly irreducible subshifts over amenable groups". In: *Monatshefte für Mathematik* 165 (2012), pp. 155–172.

[11] Michael Hochman. *Irreducibility and periodicity in $\mathbb{Z}^2$ symbolic systems.* 2024. arXiv: 2401.02273 [math.DS].

[12] Wolfgang Krieger. "On the subsystems of topological Markov chains". In: *Ergodic Theory and Dynamical Systems* 2.2 (1982), pp. 195–202.

[13] Douglas Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge university press, 2021.

[14] Tom Meyerovitch. *An embedding theorem for multidimensional subshifts.* 2023. arXiv: 2312.05650 [math.DS].

[15] Ville Salo. *When are group shifts of finite type?* 2018. arXiv: 1807.01951 [math.GR].

[16] Ville Salo and Ilkka Törmä. "Category theory of symbolic dynamics". In: *Theoretical Computer Science* 567 (2015), pp. 21–45. ISSN: 0304-3975. DOI: https://doi.org/10.1016/j.tcs.2014.10.023. URL: https://www.sciencedirect.com/science/article/pii/S0304397514008020.

[17] Ville Salo and Ilkka Törmä. "Nilpotent endomorphisms of expansive group actions". In: *International Journal of Algebra and Computation* 31.03 (2021), pp. 393–452.