

STRONGLY APERIODIC SFTS ON HYPERBOLIC GROUPS: WHERE TO FIND THEM AND WHY WE LOVE THEM

YO'AV RIECK

ABSTRACT. D. B. Cohen, C. Goodman-Strauss, and the author [*Ergodic Theory Dynam. Systems* **42** (2022), no. 9, 2740–2783] proved that a hyperbolic group admits an SA SFT if and only if it has at most one end. This paper has two distinct parts: the first is a conversation explaining what an SA SFT is and how it may be of use, while in the second part I attempt to explain both old and new ideas that go into that proof. References to specific claims in the work cited above are given, with the hope that any interested reader may be able to find the details there more accessible after reading this exposition.

This paper is about a structure called *strongly aperiodic subshift of finite type*, abbreviated as SA SFT, that exists on some groups and not on others (all groups considered are infinite and finitely generated). The paper has two parts: in Part 1 we discuss the reasons for constructing SA SFTs and explain what they are. In Part 2 we explain the basic ideas behind the construction which led to the theorem of D. B. Cohen, C. Goodman-Strauss, and the author [10]:

Theorem 0.1. *A hyperbolic group admits an SA SFT if and only if it has at most one end.*

This paper is written for topologists and geometric group theorists, and we are not assuming any familiarity with SA SFTs. We do assume familiarity with hyperbolic groups, which were introduced by Gromov in an extremely influential paper [18]. Gromov outlined there the study of SFTs on hyperbolic groups, and a detailed investigation was carried out by Coornaert and Papadopoulos [11, 12]. Our interest is promoting the SFT to an SA SFT.

Remark 0.2. Much of this paper reviews old and well-known results. In order to facilitate the reading of [10], we included many references to that work. This is *not* an indication that all the claims are original.

2020 *Mathematics Subject Classification.* 20F67, 20F65, 05B45.

The author was partially supported by the Simons Foundation (award number 637880).

Part 1. Why do we do what we do?

1. THE GEOGRAPHY OF A GROUP

Fix, for the entirety of Part 1, an infinite, finitely generated group G , together with a finite generating set $S = S^{-1}$.

As usual, we view the group as its Cayley graph Γ , which we define here for completeness. The vertices of Γ are the elements of G , and the edges out of a vertex g have the form (g, ga) for all $a \in S$. The edges of Γ are labeled and directed — the edge (g, ga) is labeled a and points from g to ga . Since $S = S^{-1}$, there is an edge pointing in the opposite direction, namely, (ga, gaa^{-1}) . Connectivity of Γ follows from the fact that S generates G . Assigning length 1 to each edge induces a metric on Γ . Left multiplication defines an action of G on Γ by a graph isomorphism that preserves the labels, directions, and distances. This action is transitive on the vertices.

Since G acts on the vertices of Γ transitively, we cannot use Γ to distinguish elements, that is, for any $g_1, g_2 \in G$ and any $R > 0$, the balls $B(R, g_1)$ and $B(R, g_2)$ are identical.

How are we to know where we are?

One way is to name each and every group element. For example, if the group is the integers \mathbb{Z} , each integer has a name and we may use it. This will not do. When labeling group elements we may only use *finitely many* labels (imagine a small device attached to each group element that records the label) and the group G is, by assumption, infinite.

Here's another suggestion: label the identity 0 and all other elements 1. This poses a different problem: once the labeling of the group is completed we want to be able to *guarantee* that it is “correct”, and this must be done locally. If we allow arbitrarily large balls to be labeled 1, what is to stop us from labeling the entire group 1?

This pinpoints our goal: to construct a finite set of labels and a collection of local rules that will allow us to distinguish elements of G . Specifically, we would like to (and here we give an intuitive description of an SA SFT):

- (1) Choose a finite set A .
- (2) Construct *local rules* for labeling.
- (3) Guarantee that in *any* labeling that obeys the local rules, and for any $g_1 \neq g_2 \in G$, there is $R > 0$, so that the labelings on $B(R, g_1)$ and $B(R, g_2)$ are distinct (R depends on g_1 and g_2 and this is unavoidable).
- (4) Prove that there exists a labeling obeying the rules, so condition (3) is not vacuous.

Remark 1.1. A few remarks are in order:

- (i) A finite set carries no information other than its cardinality. (In reality, in order to work with A , each label will carry useful information; a similar comment can be made about the states of a Turing machine.)
- (ii) The local rules are the key here. *Local rules* form an atlas, or a finite collection of *charts*, where each chart is a labeling of a ball of radius R , for some fixed R .¹ We say that the labeling satisfies the local rules if the labeling about each point g agrees with one of the charts (translated to be centered at g).
- (iii) From this point on we focus on labeling the vertices, but the fixed labeling of the edges plays a key role: since the edges are labeled and directed, once we center a chart at g the exact position of each vertex of the ball is determined (we cannot “rotate” the ball).
- (iv) An SFT is only required to satisfy the first two conditions above: a finite set, and local rules for labeling. The other two conditions may not be satisfied: on the one hand, the rules may allow for a labeling in which some (even all!) vertices look the same; on the other hand, the rules may not allow for *any* labeling at all (this is an empty SFT). Adding conditions (3) and (4) is the content of “SA”.

2. RELATION TO TILINGS

An SFT on a group G is sometimes called a “tiling” of G . This is very reasonable, as we now explain. Starting from the second point of Remark 1.1, we see that an SFT is given by an atlas of charts. Say there are n charts and number them 1 through n .

Now the labeling about each g , with labels from A , must coincide with one of the charts. We record this, using the number of the chart at g (thus obtaining a function $G \rightarrow \{1, \dots, n\}$). Conversely, any function $G \rightarrow \{1, \dots, n\}$ gives a labeling $G \rightarrow A$ provided that overlapping charts are compatible. An easy argument shows that if the labelings about neighboring vertices agree, then the labeling is consistent.

And so here is the tiling: each tile covers exactly one group element g and carries a label from $\{1, \dots, n\}$ that corresponds to the chart used for the labeling about g . Let T_g denote the tile that covers g . We can visualize T_g as a “polygon” that covers only g , where the boundary of T_g consists of finitely many “edges”, one for each neighbor of g , where “edges” can be glued together if the corresponding charts are compatible. Here *edge* means an “edge” of the “polygon”, not an edge of the Cayley graph (in fact every “edge” of the glued “polygons” crosses exactly one edge of the Cayley graph and *vice-versa*). By construction, T_g depends only on the chart at g ; thus there are exactly n possible tile shapes. For obvious reasons this is called a *nearest neighbor SFT*.

¹The reader may have seen local rules defined on finite sets that are not necessarily balls, but there is no loss of generality in assuming that they are all balls and of the same radius.

We conclude that, for a group G , the concepts of *SFT* and *tiling* are the same. The notions of “aperiodicity” (to be defined in the next section) carry over. We will focus on SFTs.

One final remark: SFTs on \mathbb{Z}^2 correspond to tilings of \mathbb{E}^2 (where by \mathbb{E}^2 we mean \mathbb{R}^2 with the Euclidean metric), but the converse does not hold. There are tilings of \mathbb{E}^2 (for example, see [22]) that are not aligned along a lattice and hence do not induce an SFT on \mathbb{Z}^2 . Similarly, strongly aperiodic tiles for \mathbb{H}^2 were first constructed by Goodman-Strauss in [17], but do not provide an SFT on any group of isometries of \mathbb{H}^2 ; the first SFT on a hyperbolic was constructed by Cohen and Goodman-Strauss twelve years later [9].

3. FORMAL DEFINITIONS

After reviewing the concepts the formal definitions should be easy to read. For a finitely generated group G we provide the following definitions:

- (1) The *full shift* on G with labels in a finite set A is

$$A^G := \{\omega : G \rightarrow A\}.$$

In other words, these are all possible labelings of the elements of G by labels from A .

- (2) We endow A with the discrete topology and A^G with the product topology. This turns A^G into a compact space that carries a right G -action; the action of $g \in G$ on $\omega \in A^G$ is defined as follows (here, we evaluate $\omega \cdot g \in A^G$ on $h \in G$):

$$(\omega \cdot g)(h) = \omega(gh).$$

This action is by homeomorphisms.

- (3) A *subshift* is a closed, invariant subset $\Omega \subset A^G$. The action of G on A^G induces an action of G by homeomorphisms on the compact space Ω . Note that $\Omega = \emptyset$ is a subshift. More interesting examples are given by the closure of the orbit of any $\omega \in A^G$.
- (4) Given a subshift Ω we call $\omega \in \Omega$ a *configuration*. Thus any configuration is a function

$$\omega : G \rightarrow A.$$

Consistent with the terminology above, we call the value $\omega(g)$ the *label* at g .

- (5) A subshift Ω is called *strongly aperiodic (SA)* if it is not empty and, for every $\omega \in \Omega$, we have that $\text{Stab}(\omega)$ is **trivial** (here and below $\text{Stab}(\cdot)$ stands for *stabilizer*).
- (6) A non-empty subshift is called *weakly aperiodic* if $\text{Stab}(\omega)$ has infinite index in G for every $\omega \in \Omega$. Since we restricted ourselves to infinite groups, strong aperiodicity implies weak aperiodicity (but at this point of our discussion, it is not clear if the converse holds). We will focus on strong aperiodicity.
- (7) A *subshift of finite type (SFT)* is a subset of the full shift A^G that is defined by local rules, or an atlas of charts, as described above. It is not defined

as a subshift satisfying an extra condition, but in fact it is: it is not hard to see that an SFT is closed and G -invariant, and hence a subshift.

- (8) Combining (5) and (7), we define a *strongly aperiodic subshift of finite type (SA SFT)* to be an SFT in which the stabilizer is trivial for every configuration.

4. ORIGINS

The origins of the theory, in the very early 1960s, are closely linked to axiomatization and computational theory (as we shall see shortly). H. Wang [26] asked if the *Domino Problem* is decidable for tilings of \mathbb{E}^2 :

The Domino Problem. Is there an algorithm that decides if a given finite set of tiles in \mathbb{E}^2 can be used to tile the plane?

Obviously, one can replace \mathbb{E}^2 with other spaces of interest, for example, \mathbb{H}^2 , \mathbb{E}^n , and \mathbb{H}^n . (We will not discuss higher dimensions.)

Wang proposed an “algorithm”: try! If the tiles do not tile \mathbb{E}^2 , an attempt to tile arbitrarily large balls would eventually fail (and one will certainly detect this). On the other hand, if they do tile, one “should” find a tiled domain. Of course, in the latter the tiling would be a lift of a tiling of the torus. But does this actually work? Perhaps there is a set of tiles that tiles the plane but not the torus? This seemed rather unlikely.

Wang himself found such tiles, but they are *seeded*, which means that they contain a special tile (the seed) that must be used. With this, Wang constructed tiles that emulate an arbitrary Turing machine and tile the plane if and only if the given machine never halts; undecidability of the Domino Problem for seeded tiles follows from undecidability of the Halting Problem. Figure 1 indicates a Turing machine and seeded tiles that emulate it. Note that without the seed (in the second row, second column) copies of the blank tiles (in the first row) can always be used to tile \mathbb{E}^2 (even the torus), and thus without the seed these tiles give nothing of interest.

A few years later, R. Berger showed how to emulate an arbitrary Turing machine with an *unseeded* set of tiles (see Figure 2, which shows a much-simplified set of tiles due to R. Robinson). These tiles can be used to show two things:

- (1) The Domino Problem in \mathbb{E}^2 is undecidable.
- (2) There exists a set of strongly aperiodic tiles for \mathbb{E}^2 . This is indicated in Figure 2. As suggested by that figure, the tiles assemble to form disjoint squares of arbitrarily large scale. Hence any element of the stabilizer will need to move these squares off themselves or not move them at all, showing that the stabilizer is trivial.

Later, J. Kari [20] gave a different and very simple construction that also implies undecidability of the Domino Problem in \mathbb{E}^2 .

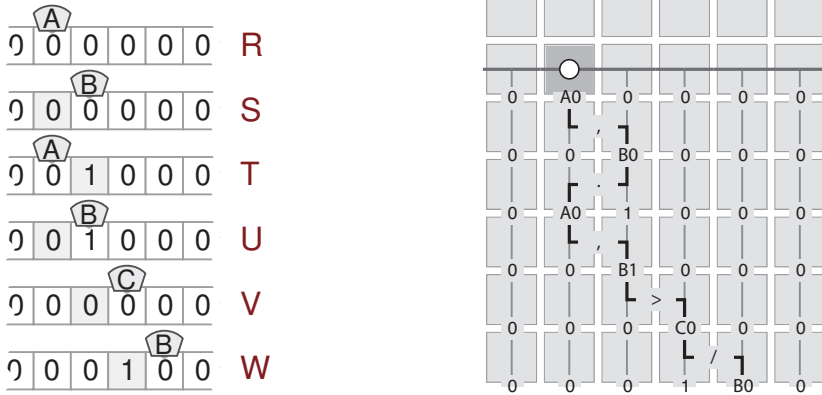


FIGURE 1. Wang emulates the run of a Turing machine as a tiling problem.




FIGURE 2. Unseeded tiles.

5. OTHER GEOMETRIES

Existence of SA SFTs has subtle connections to the geometry of the group. For example, Cohen [8] showed that for finitely presented groups, existence is a quasi-isometry invariant.² I find this quite surprising because a quasi-isometry is very coarse and can destroy any local information, and SFTs are defined using local rules only.

Here is a sample of results. Some show existence of an SA SFT, and some show that the Domino Problem is undecidable (implying the existence of a weakly aperiodic set of tiles, but not implying the existence of an SA SFT). More results can be found in [10]. The 2018 survey [1] contains many results about the Domino Problem. All groups considered in this section are assumed to have decidable word problem and to be finitely presented (although some of the results require only finite generation).

In 1997 Mozes [21] constructed strongly aperiodic tilings on any simple Lie group Γ of rank greater than one. The tiles themselves are Voronoi cells of a uniform lattice $G \leq \Gamma$.³ The Voronoi cells are then decorated with *combinatorial* information, given by the pattern of intersection they form with the Voronoi cells of a second, incompatible, uniform lattice. Mozes showed that incompatibility of the lattices breaks all symmetry and yields a strongly aperiodic tiling of Γ ; for our purposes, it also gives an SA SFT on G . (This situation is similar to the tiling shown in Figure 3, although these tiles come from a different construction in \mathbb{H}^2 .)

We remark that uniform lattices in rank-one simple Lie groups are hyperbolic and therefore admit an SA SFT by Theorem 0.1. Thus a uniform lattice in any simple Lie group admits an SA SFT.

In [17] Goodman-Strauss constructed the first set of strongly aperiodic tiles in \mathbb{H}^2 . A simpler construction is shown in Figure 3 where the grid “rectangles” are all isometric pentagons, each with five edges of length 1, four vertices of angle $\pi/2$, and one vertex of angle π , in the middle of the bottom edge. These overlay an incompatible shaded grid that has two types of tiles, hexagons and pentagons. We note, however, that these tiles do not correspond to a lattice and therefore do not provide an example of an SA SFT.

In a recent preprint [6], Barbieri, Sablik, and Salo showed that the product of two non-amenable groups with decidable word problem admits an SA SFT.

In 2019 Barbieri [5] showed that the product of three infinite groups, each with decidable word problem, admits an SA SFT.

In [3] Aubrun and Kari constructed SFTs on $BS(1, n)$ groups, which were shown to be SA by Esnay and Moutot in [16]. Furthermore, in a recent paper Aubrun and Kari [4] showed that the Domino Problem is undecidable for all $BS(m, n)$.

The reader may have noticed that not even one of the groups discussed above is hyperbolic. This changed in 2017 when Cohen and Goodman-Strauss showed that

²For the definition of *quasi-isometry*, see, for example, [7].

³The Voronoi cell corresponding to $g \in G$ consists of all the $x \in \Gamma$ satisfying $d(x, g) \leq d(x, g')$ for all $g' \in G$.




FIGURE 3. Two patterns of “rectangles” are shown, each rectangle having some predecessor above and some successors below. In the pattern drawn with dark lines, the number of rectangles doubles from row to row. In the gray pattern, light rectangles (which are all congruent) have one light and one dark rectangle as successors, and dark rectangles (which are all congruent) have one light and two dark successors.

surface groups admit an SA SFT [9]. Since then Aubrun, Barbieri, and Moutot [2] showed that the Domino Problem for surface groups is undecidable.

6. TWO NECESSARY CONDITIONS

From this point on we are only considering finitely presented groups. There are two known necessary conditions for existence of an SA SFT on a group G :

- (1) G must have a decidable word problem.
- (2) G must have at most one end.

This leads us to ask:

Question 6.1. *Does every finitely presented, one-ended group with decidable word problem admit an SA SFT?*

The first condition, proved by E. Jeandel [19], fits well with the philosophy that an SA SFT gives us a way to “address” group elements. After all, being able to construct arbitrarily large balls of the Cayley graph is equivalent to having a decidable word problem.

Here are a few details regarding the second condition.

It was shown by Cohen [8] that no group with more than one end admits an SA SFT. We demonstrate this for \mathbb{Z} , which is a well-known and easy case. Suppose we are given a non-empty SFT on \mathbb{Z} with label set A and local rules that are defined on sets of (say) n adjacent integers. Note that an interval of length n can be labeled in finitely many ways. It follows that there are two disjoint intervals of length n with identical labeling; say that one starts at a and the other starts at $b > a$. It is now an easy exercise to show that the labeling of the integers from a to $b - 1$ can be repeated to produce a periodic configuration.

Let us return for a moment to Wang’s tiles. Recall that these tiles had a *seed* that we were required to use. The following example demonstrates the dramatic effect of seeding:

- $A := \{\mathbf{C}, \mathbf{L}, \mathbf{R}\}$. The meaning of the labels: \mathbf{C} is the “center” of \mathbb{Z} ; \mathbf{L} and \mathbf{R} mean “left/right of \mathbf{C} ”, respectively.
- Local rules (allowable configurations, defined on adjacent pairs):
 - $\mathbf{C-R}$
 - $\mathbf{L-C}$
 - $\mathbf{R-R}$
 - $\mathbf{L-L}$.
- The seed is \mathbf{C} .

A moment’s reflection shows that any allowable configuration has the form

$$\cdots\mathbf{-L}\cdots\mathbf{-L-C-R}\cdots\mathbf{-R}\cdots.$$

Of course, the stabilizer of such a configuration is trivial, since any non-trivial translation will move \mathbf{C} to a point labeled \mathbf{L} or \mathbf{R} . So the seeded SFT is SA. This is in contrast to *unseeded* SFTs, where we saw that \mathbf{Z} admits no SA SFT. Indeed, if we remove the seeding requirement we have two periodic configurations:

$$\cdots\mathbf{-R}\cdots\mathbf{-R}\cdots \quad \text{and} \quad \cdots\mathbf{-L}\cdots\mathbf{-L}\cdots.$$

This concludes the first part, in which we attempted to do three things: explain what an SA SFT is, motivate our interest in them, and survey some of the many results. The reader should be aware that in recent years there has been a flurry of activity and this discussion is far from complete.

Part 2. How do we do what we do?

In the second part of this paper we discuss the construction of an SA SFT on a 1-ended hyperbolic group G which we fix once and for all. We also fix a finite generating set $S = S^{-1}$. We start with the construction of an SFT that is not required to be strongly aperiodic, and has been known for a long time (Ω_S below). We then enhance it to attain the desired SA SFT.

A finitely presented group G (with a fixed finite generating set $S = S^{-1}$) is called *hyperbolic* if there is some $\delta > 0$ so that every triangle in the Cayley graph of G is δ -slim, that is, any point on one edge is within δ from the union of the other two edges. We fix $\delta > 0$ for the remainder of this paper.

Hyperbolicity, due to Gromov [18], turns out to be both very natural (satisfied by many groups that appear in application) and very useful. As a simple example, hyperbolicity implies a very efficient solution to the word problem known as *Dehn’s algorithm*; surprisingly, this turns out to be equivalent to hyperbolicity. As a far more sophisticated example we mention that the isomorphism problem for hyperbolic groups is decidable (Sela [23] for torsion free hyperbolic groups, later extended to all hyperbolic groups by Dahmani and Guirardel [14]). These are, of course, just examples of what one can do with hyperbolic groups; very many other

beautiful results are known and at this point the theory of hyperbolic groups seems very well understood.

An excellent reference is [7]. Another very useful reference is [15], in particular for the *shortlex FSA*, which will be used extensively below.

7. SHORTLEX SHELLINGS

In this section we explain how to construct a certain SFT, called the *shortlex SFT* and denoted by Ω_S , on a hyperbolic group G ; these SFTs will serve as the backbone for our main construction. The group G need not be one-ended, and the resulting SFT is not necessarily strongly aperiodic (in fact, if G is infinite, this SFT is necessarily not strongly aperiodic). The ideas presented here are similar to ones that date back to Gromov's paper [18]; see Coornaert and Papadopoulos [11, 12] for a detailed treatment.

A key to our construction is *Cannon's Shortlex FSA* (finite state automaton)⁴. Fixing an arbitrary order on S induces a lexicographic order on the finite words $(S \cup S^{-1})^*$. A path in the Cayley graph is called *shortlex* if and only if it is a geodesic, and is first in the lexicographic order among all geodesics with the given endpoints. It is clear that for each element $g \in G$ there is a unique shortlex geodesic from the identity e to g (and this is true for any finitely generated group). It is far less obvious how to calculate shortlex representatives. In fact, shortlex representatives in a group G can be calculated if and only if G has a decidable word problem, see [15].

A hyperbolic group G admits an FSA, described in [15], that accepts a word if and only if it is a shortlex representative. Denoting the states for the FSA as

$$\{s_1, \dots, s_n\},$$

we can now label each group element with the three labels

$$(P(g), \text{dist}(e, g), s_{i(g)}),$$

defined as follows:

- (1) $P(g)$ is the generator that points towards e on the shortlex geodesic to g ; in other words, the product $gP(g)$ is the last group element on the shortlex geodesic before arriving at g . $P(e)$ is not defined. P is called the *parent function*.
- (2) $\text{dist}(e, g)$ is the distance from e to g .
- (3) $s_{i(g)}$ is the state of the FSA at g (so $i(g) \in \{1, \dots, n\}$).

Of course, $\text{dist}(e, g)$ cannot be used as a label for an SFT since it takes on infinitely many values. We replace it with the function that describes the *difference* of the distance to the origin between g and its neighbors, that is, for each $a \in S$, we define

$$\bar{\text{d}}(g)(a) := \text{dist}(e, g) - \text{dist}(e, ga).$$

⁴For those unfamiliar with it, an FSA is a Turing machine without a tape; the "memory" is contained in finitely many states.

Since g and ga are neighbors, this function can only take the values ± 1 or 0 . Thus $\bar{\delta}$ is a function

$$\bar{\delta} : G \rightarrow \{-1, 0, 1\}^S.$$

This completes the labeling on g for each group element. This labeling is *not* a configuration in any SFT but rather a blueprint for constructing Ω_S , as we now explain.

We now define a full shift of G with labels

$$(\bar{\delta}(g), s_{i(g)}, P(g)) \in \{-1, 0, 1\}^S \times \{s_1, \dots, s_n\} \times S. \tag{7.1}$$

Let Ω_S be the subshift consisting of all configurations satisfying the following condition:

The labeling of any ball, of any radius, coincides with labeling above on some ball not containing e .

It is not clear that:

- (1) This is an SFT. The requirement above says “balls of any radius”, and we must ensure that this can be enforced by considering balls of a fixed radius. This is one of the many places in which hyperbolicity is used.
- (2) This SFT is not empty.

The second point is actually not too hard. Since we can label balls of arbitrary radius, a diagonalization argument shows that there is a labeling of all of G , that is, the SFT is not empty.

Remark 7.1. It is known that Ω_S is necessarily *not* SA. More needs to be done.

Remark 7.2. We give an example of a parent function P defined on F_2 , the free group on generators a and b . The Cayley graph of F_2 , a regular 4-tree, is shown in Figure 4. At each vertex there are two gray edges (corresponding to a and a^{-1}) and two black edges (corresponding to b and b^{-1}). Each edge is marked with an arrow that points from g to $P(g)$. As the figure suggests, $P(g)$ can be *any* neighbor of g ; the only rule is that at any g we see exactly one triangle pointing “out” (towards $P(g)$) and three pointing “in” (from $P^{-1}(g)$). Picking any $g \in F_2$ and following the arrows we arrive at a point on $\partial_\infty F_2$, and this point is independent of the choice of g .

8. RELATION TO THE BOUNDARY AND WEAK APERIODICITY

We are assuming that the reader is familiar with $\partial_\infty G$, the *boundary at infinity* of G .

Given any configuration $\omega \in \Omega_S$ and any $g \in G$, we construct the path

$$g, Pg, P^2g, \dots$$

It follows from the definition of P that this path is a geodesic ray and hence defines a point at infinity. Hyperbolicity implies that this point is independent of g , and we denote it by ξ (of course, $\xi(\omega)$ depends on ω). The association $\omega \mapsto \xi$ defines a function

$$\Omega_S \rightarrow \partial_\infty G.$$




FIGURE 4. The function P .

It is clear that this function is compatible with the actions of G on Ω_S and on $\partial_\infty G$. This means that the action of G on $\partial_\infty G$ is a *factor* of an SFT, which motivated Gromov to study them.

We would like to exploit the function $\Omega_S \rightarrow \partial_\infty G$ differently: we will use it to show that Ω_S is weakly aperiodic. The map $\Omega_S \rightarrow \partial_\infty G$ shows that, for any $\omega \in \Omega_S$, we have

$$\text{Stab}(\omega) < \text{Stab}(\xi(\omega)).$$

It is known that $\text{Stab}(\xi)$ is virtually cyclic for any $\xi \in \partial_\infty G$.⁵ We conclude that $\text{Stab}(\omega)$ is virtually cyclic (for any $\omega \in \Omega_S$). By assumption, G is one-ended, and

⁵To see that $\text{Stab}(\xi)$ is virtually cyclic use the following: (1) every non-torsion element of G fixes exactly two points on the boundary; (2) if two elements share one fixed point, they share both fixed points; and (3) the elements that fix the same two points form a virtually cyclic group (this allows for a finite group, as finite groups are virtually trivial). See [18, Sections 8.1 and 8.2] and [7, proof of Proposition III.Γ.3.20 (p. 467)].

hence G is not virtually cyclic. This shows that $\text{Stab}(\omega)$ has infinite index; in other words, Ω_S is weakly aperiodic.

9. HOROSPHERES

Fix a configuration $\omega \in \Omega_S$. We may “integrate” $\bar{\partial}$ (notation as in (7.1) above) to get a function $h : G \rightarrow \mathbb{Z}$ (defined up to an additive constant). The level sets of h are called *horospheres*.

Any $g \in \text{Stab}(\omega)$ preserves $\bar{\partial}$ and hence g preserves h up to an additive constant, that is, there exists an integer C so that, for any $x \in G$, we have

$$h(g \cdot x) = h(x) + C.$$

If g is a torsion element then $C = 0$. Conversely, if g has infinite order then $C \neq 0$ ([10, Lemma 9.1]). This simple observation will prove quite useful.

Since a hyperbolic group has finitely many conjugacy classes of torsion, getting rid of the torsion in the stabilizers is very easy ([10, Proposition 3.3]). Thus we may assume that g has infinite order and $C \neq 0$.

10. THE DIVERGENCE GRAPH

Fix $\omega \in \Omega_S$. As we discussed in Section 8, for every $g \in G$, the sequence

$$g, Pg, P^2g, \dots$$

converges to $\xi(\omega) \in \partial_\infty G$. For convenience we describe P as moving “down” towards ξ .

We now move “up” from $g \in G$:

$$g, P^{-1}g, P^{-2}g, \dots$$

This defines sets that move away from g ; unlike the downwards path that limits on ξ , the limit of the sets $P^{-n}g$ is more complex. It *does* depend on g , and will often be an uncountable subset of $\partial_\infty G$. The union of these sets is called the *future cone* of g , namely,

$$P^{-*}g = \bigcup_{n=0}^{\infty} P^{-n}g.$$

Remark 10.1 (Growth rate). The hyperbolic group G has a well-defined growth rate, which we will denote by $\lambda > 0$. Given a configuration $\omega \in \Omega_S$, the future cone of each $g \in G$ is completely determined by the state of the FSA at g , denoted by $s_{i(g)}$ in (7.1) above (this was analyzed in detail in [13]). In particular, $s_{i(g)}$ determines the growth rate of the future cone, and we call it the *growth rate of the state*. In what follows we only consider states whose future cone has growth rate λ (for a precise discussion see [10, Definition 6.3]). We denote the set of all elements of G satisfying this condition by G^+ . Note that each $g \in G^+$ accumulates to an uncountable subset of $\partial_\infty G$. By [10, Proposition 6.5] any 2δ -ball, in any configuration $\omega \in \Omega_S$, contains a point of G^+ .

Back to our discussion, fix a configuration $\omega \in \Omega_S$.

Let h be as in Section 9. For $i \in \mathbb{Z}$, set $H_i := h^{-1}(i)$ and $H_i^+ := G^+ \cap H_i$. On H_i^+ we define the *divergence graph* as follows:

- v The vertices of the divergence graph are the vertices of H_i^+ .
- e Two vertices $g_1, g_2 \in H_i^+$ are connected by an edge if and only if their futures remain a bounded distance apart. In other words, for some $C > 0$, and for each integer $n \geq 1$, there are $v_1 \in P^{-n}g_1$ and $v_2 \in P^{-n}g_2$ with $\text{dist}(v_1, v_2) < C$.

The following holds (see see [10, Lemma 7.4] and its proof):

- The union of the limit sets of the future cones of all of H_i^+ is $\partial_\infty G - \{\xi\}$.
- $g_1, g_2 \in H_i^+$ are connected by a divergence graph edge if and only if the limit sets of their future cones intersect.
- We may therefore describe the divergence graph as a discrete approximation of $\partial_\infty G$.
- Most importantly, *the divergence graph is connected*. This reflects the fact that $\partial_\infty G - \xi$ is connected (the Cut Point Conjecture, proved by Swarup [25]).

The plan is now as follows. Any infinite order element $\phi \in \text{Stab}(\omega)$ will translate the levels of h by $C \neq 0$, as explained in Section 9. It is our goal to enhance Ω_S by associating an integer $\Delta(g)$ to each $g \in G^+$ in a way that cannot be periodic; the new SFT will no longer have ϕ in its stabilizer.

We enhance the labels of Ω_S (compare this with the labels presented in (7.1))

$$(\partial(g), s_{i(g)}, P(g), \wp(g), \Delta(g), m(g)).$$

The enhanced SFT is called *populated shelling*, denoted by Ω_P . The name comes from the fact that \wp defines a “population” of “villagers” on each “village” $g \in G$ (or, if the reader prefers, any village $v \in G^+$, since the population of any $g \in G - G^+$ is zero anyway).

11. \wp AND m

The number of “villagers” defines the function \wp , with $\wp(g)$ being the population at g :

$$\wp : G^+ \rightarrow \mathbb{Z}_{\geq 1}.$$

This function is required to be bounded, with the population bound N fixed in advance.

Clearly, we need something that will help us relate \wp to the geometry of the group, for otherwise the population values will be arbitrary numbers. This is m , which stands for *matching*. More precisely, it is *parent-child matching*.

Each “villager” in $v \in H_i^+$ has q^{Δ_i} children (q^{Δ_i} will be described in the next section; for now just take it to be “some number”). We list the children as follows:

$$(v, j, k),$$

where v is the village, $1 \leq j \leq \wp(v)$ is the villager, and $1 \leq k \leq q^{\Delta_i}$ is the child. The function m “places” this child as a villager in H_{i+1}^+ :

$$m(v, j, k) = (u, l);$$

here, $u \in H_{i+1}^+$ is a village and $1 \leq l \leq \wp(u)$ is a villager.

The geometry of the group comes to play when we force the child to be placed not too far from the parent. The precise condition is the following: the child may take up to 3 steps on the divergence graph of H_i^+ , and then move one step up. A succinct description is this: the child of a villager in $v \in H_i^+$ is a villager in $u \in H_{i+1}^+$ with

$$\text{DivDist}(v, Pu) \leq 3.$$

Naturally, DivDist denotes that divergence graph distance.

Remark 11.1. The reader probably finds the constant 3 rather arbitrary (not to say mysterious). It comes from an application of a theorem in graph theory which states that if a graph is connected, then its cube admits a Hamiltonian path between any two vertices.⁶ This is then used to construct a “translation-like \mathbb{Z} action” (in the sense of Seward [24]) whose *defect* is 3, and this is the origin of the constant.

12. STRAYING AWAY AND COMING BACK HOME

It would be quite natural to worry that we are too loose with the geometry of the group here. We populate the group with the goal of considering the population after arbitrarily many generations, and descendants may stray 3 divergence-graph steps each generation.

Indeed, after many generations, a descendant of the villager (v, j) may be in a village which is very far from the future cone $P^{-*}v$. It may be worth emphasizing that we are using two distinct “futures” here: on the one hand, we have the future cone $P^{-*}v$, which is the collection of *villages* u for which

$$v = P^n u$$

for some n ; on the other hand, there are *villagers* that are descendants of (v, j) , and may stray away from $P^{-*}v$.

Hyperbolic geometry to the rescue. As described above, let (u, l) be a descendant of (v, j) , after, say, n generations. We use the notation $v = v_0, v_1, \dots, v_n = u$ for the villages so that (u, l) is a descendant of (v_{n-1}, l_{n-1}) (for some l_{n-1}), (v_{n-1}, l_{n-1}) is a descendant of (v_{n-2}, l_{n-2}) (for some l_{n-2}), and so on. It is not hard to show that an edge of a divergence graph connects vertices of Cayley distance at most 2δ , and so the apple doesn’t fall too far from the tree:

$$\text{CayDist}(v_i, v_{i+1}) \leq 6\delta + 1.$$

On the other hand, because $v_i \in H_i^+$ and $v_j \in H_j^+$, we get that

$$\text{CayDist}(v_i, v_j) \geq |i - j|.$$

⁶The *cube* of a graph is obtained by adding an edge between any two vertices of distance at most 3.

This produces a *quasi-geodesic* that can be compared with the geodesic

$$v_n, P v_n, \dots, P^n v_n.$$

It is a feature of hyperbolic geometry that quasi-geodesics remain a bounded distance away from geodesics, which means that, for some fixed $R > 0$, we have

$$\text{CayDist}(v_0, P^n v_n) \leq R.$$

A precise statement is given in [10, Lemma 9.2]. This allows for sufficient control over the population growth, since it shows that all the descendants of villagers in $S \subset H_i^+$, where S is any finite set, live in the future cone of $\mathcal{N}(S)$, where here $\mathcal{N}(S)$ is the 3-neighborhood of S in the divergence graph on H_i^+ .

13. Δ

We finally describe Δ , focusing on G^+ ; Δ is extended to G by setting it to be zero on $G - G^+$. It is Δ that will ultimately be responsible for aperiodicity.

For each $g \in G^+$, the growth rate of the population at g is controlled by $\Delta(g)$. To be precise, every villager at g has exactly $q^{\Delta(g)}$ children (for q to be decided momentarily). We use, intentionally, a number q which is *not* compatible with λ : q is an integer, which we may take to be either 2 or 3, so that

$$\frac{\log(q)}{\log(\lambda)} \notin \mathbb{Q}.$$

The function

$$\Delta : G^+ \rightarrow \{ \lfloor \log_q(\lambda) \rfloor, \lceil \log_q(\lambda) \rceil \}$$

is required to satisfy the following conditions:

- (1) (And this is key) Δ is constant along levels of h , that is, for each integer i , $\Delta|_{H_i^+}$ is constant. We denote this value by Δ_i . This defines the sequence

$$(\Delta_i)_{i \in \mathbb{Z}}.$$

- (2) The sequence Δ_i approximates λ (in a sense made precise in [10, Corollary 9.3]).
- (3) The condition above and the incompatibility of λ and q guarantee that Δ_i is not periodic ([10, Corollary 9.4]).

We saw in Section 9 that any infinite-order elements in the stabilizer of a configuration $\omega \in \Omega_S$ must translate the levels $\{h = i\}$ by a non-zero amount. Once we enhance Ω_S by populating G^+ , the function Δ will not be invariant under such an element. This shows that there is no infinite-order element in the stabilizer of any configuration in Ω_P , as desired.

Remark 13.1 (The role of the (necessary!) assumption of one-endedness of G). The question of existence of an SA SFT is irrelevant for zero-ended groups (that is, finite groups), where the answer is always “yes”, as well as two-ended groups (that is, virtually- \mathbb{Z} 's), where the answer is always “no”. So we ignore these groups in this remark and consider only one-ended and infinitely-ended groups. Our main result is that an SA SFT exists *only* for the former. This begs the question: where

exactly was the “one-endedness” assumption used? The only place is *imposing that Δ be constant along each H_i^+* . This must be enforced via local rules, as part of the SFT. What allows us to do this is *connectivity of the divergence graph*. This is the one and only place where the assumption is used, as connectivity of the divergence graph is equivalent to G being one-ended. Swarup’s resolution of the cut-point conjecture [25] plays a key role here; it states that $\partial_\infty G \setminus \{\xi\}$ is connected (for any $\xi \in \partial_\infty G$).

14. ONE LAST ISSUE

It is not the goal of this paper to give a complete proof of Theorem 0.1. Our goal is to explain some of the elements that go into the proof in a way that would facilitate its reading. However, it is hard to ignore the fact that we have not addressed the following question:

Does a populated shelling even exist?

This should not be taken lightly as it is quite possible that we defined the empty subshift here. In fact, if the population bound is too small, this is probably the case. In [10, Proposition 8.5] it is shown that Ω_P described above is indeed an SFT, and in [10, Proposition 9.5] it is shown that no configuration has an infinite order element in its stabilizer, but neither result addresses existence of a configuration.

Much of the work in [10] is devoted to Proposition 8.12, showing that (for an appropriately chosen population bound) Ω_P is indeed not empty. This is the most technical and longest part of the proof, and here is my attempt at explaining the idea. We proceed in three steps:

Level: The first step is populating each level H_i^+ . To discuss that we need to dig a little deeper into the maximal growth states of the FSA.⁷ In Remark 10.1 we explained that the future of some states must have growth rate λ (the growth rate of the group itself) and denoted the set of vertices that have this growth rate by G^+ . From that point we concentrated on G^+ and on $H_i^+ := G^+ \cap H_i$. In fact, more is true; although the transitions of the FSA need not satisfy the assumptions of the Perron–Frobenius Theorem, it is possible to associate to them a measure that behaves just like a Perron–Frobenius eigenvector; see [13] or [10, Section 6]. Now in the first step we populate each level H_i^+ so that the population of *any* finite subset approximates its total measure (it is not possible to get this to be exact; the ratio of the measures of distinct states is usually irrational). This is [10, Lemma 8.7] (for one level). The average ratio of population to total measure (called the population density) can be chosen freely, within a reasonable range ([10, Definition 8.6 and Lemma 8.7]).

⁷There are two distinct notions of growth rate at play: growth rate of the population, controlled by Δ_i , and growth rate of the group elements, which is λ ; here we discuss the latter. Having two incompatible growth rates is a little confusing but it is the very thing that leads to strong aperiodicity.

Levels: We apply this to all of G (one level at a time, and still without matching levels). This is given in [10, Corollary 8.9]. In order to be able to match parents and children, we ensure that the sequence of densities behaves well (grows when small, shrinks when big). The precise description is given in [10, Definition 8.10].

Matching: Having populated the group as described above, we apply the Hall Matching Theorem to prove existence of a matching function m as required. This is [10, Proposition 8.11].

Ta da! That's all, folks.

ACKNOWLEDGMENTS

I would like to thank the organizers of the special session, José Ayala Hoffman, Mario Eudave Muñoz, and Jennifer Schultens, for giving me this opportunity. I thank Chaim Goodman-Strauss for helping in the preparation of this paper and the beautiful illustrations. I am very grateful to the anonymous referee for very helpful suggestions.

REFERENCES

- [1] N. Aubrun, S. Barbieri, and E. Jeandel, About the domino problem for subshifts on groups, in *Sequences, Groups, and Number Theory*, 331–389, Trends Math, Birkhäuser/Springer, Cham, 2018. MR 3799931.
- [2] N. Aubrun, S. Barbieri, and E. Moutot, The domino problem is undecidable on surface groups, in *44th International Symposium on Mathematical Foundations of Computer Science*, Art. 46, 14 pp., LIPIcs. Leibniz Int. Proc. Inform., 138, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019. MR 4008435.
- [3] N. Aubrun and J. Kari, Tiling problems on Baumslag-Solitar groups, in *Proceedings: Machines, Computations and Universality 2013*, 35–46, Electron. Proc. Theor. Comput. Sci. (EPTCS), 128, EPTCS, 2013. MR 3594268.
- [4] N. Aubrun and J. Kari, On the domino problem of the Baumslag-Solitar groups, *Theoret. Comput. Sci.* **894** (2021), 12–22. MR 4334698.
- [5] S. Barbieri, A geometric simulation theorem on direct products of finitely generated groups, *Discrete Anal.* **2019**, Paper No. 9, 25 pp. MR 3964145.
- [6] S. Barbieri, M. Sablik, and V. Salo, *Groups with Self-Simulable Zero-Dimensional Dynamics*, arXiv:2104.05141 [math.GR], 2021.
- [7] M. R. Bridson and A. Haefliger, *Metric Spaces of Non-positive Curvature*, Grundlehren der mathematischen Wissenschaften, 319, Springer-Verlag, Berlin, 1999. MR 1744486.
- [8] D. B. Cohen, The large scale geometry of strongly aperiodic subshifts of finite type, *Adv. Math.* **308** (2017), 599–626. MR 3600067.
- [9] D. B. Cohen and C. Goodman-Strauss, Strongly aperiodic subshifts on surface groups, *Groups Geom. Dyn.* **11** (2017), no. 3, 1041–1059. MR 3692905.
- [10] D. B. Cohen, C. Goodman-Strauss, and Y. Rieck, Strongly aperiodic subshifts of finite type on hyperbolic groups, *Ergodic Theory Dynam. Systems* **42** (2022), no. 9, 2740–2783. MR 4461690.
- [11] M. Coornaert and A. Papadopoulos, *Symbolic Dynamics and Hyperbolic Groups*, Lecture Notes in Mathematics, 1539, Springer-Verlag, Berlin, 1993. MR 1222644.

- [12] M. Coornaert and A. Papadopoulos, Horofunctions and symbolic dynamics on Gromov hyperbolic groups, *Glasg. Math. J.* **43** (2001), no. 3, 425–456. MR 1878587.
- [13] F. Dahmani, D. Futer, and D. T. Wise, Growth of quasiconvex subgroups, *Math. Proc. Cambridge Philos. Soc.* **167** (2019), no. 3, 505–530. MR 4015648.
- [14] F. Dahmani and V. Guirardel, The isomorphism problem for all hyperbolic groups, *Geom. Funct. Anal.* **21** (2011), no. 2, 223–300. MR 2795509.
- [15] D. B. A. Epstein, J. W. Cannon, D. F. Holt, S. V. F. Levy, M. S. Paterson, and W. P. Thurston, *Word Processing in Groups*, Jones and Bartlett, Boston, MA, 1992. MR 1161694.
- [16] J. Esnay and E. Moutot, *Weakly and Strongly Aperiodic Subshifts of Finite Type on Baumslag-Solitar Groups* arXiv:2004.02534 [math.DS], 2021.
- [17] C. Goodman-Strauss, A strongly aperiodic set of tiles in the hyperbolic plane, *Invent. Math.* **159** (2005), no. 1, 119–132. MR 2142334.
- [18] M. Gromov, Hyperbolic groups, in *Essays in Group Theory*, 75–263, Math. Sci. Res. Inst. Publ., 8, Springer, New York, 1987. MR 0919829.
- [19] E. Jeandel, *Aperiodic Subshifts of Finite Type on Groups*, arXiv:1501.06831v2 [math.GR], 2015.
- [20] J. Kari, A small aperiodic set of Wang tiles, *Discrete Math.* **160** (1996), no. 1-3, 259–264. MR 1417578.
- [21] S. Mozes, Aperiodic tilings, *Invent. Math.* **128** (1997), no. 3, 603–611. MR 1452434.
- [22] C. Radin, The pinwheel tilings of the plane, *Ann. of Math. (2)* **139** (1994), no. 3, 661–702. MR 1283873.
- [23] Z. Sela, The isomorphism problem for hyperbolic groups. I, *Ann. of Math. (2)* **141** (1995), no. 2, 217–283. MR 1324134.
- [24] B. Seward, Burnside’s Problem, spanning trees and tilings, *Geom. Topol.* **18** (2014), no. 1, 179–210. MR 3158775.
- [25] G. A. Swarup, On the cut point conjecture, *Electron. Res. Announc. Amer. Math. Soc.* **2** (1996), no. 2, 98–100. MR 1412948.
- [26] H. Wang, *Computation, Logic, Philosophy, Mathematics and its Applications* (Chinese Series), 2, Science Press Beijing, Beijing, 1990. MR 1112395.

Yo’av Rieck

Department of Mathematics, University of Arkansas, Fayetteville, AR 72701, USA
yoav@uark.edu

Received: October 14, 2021

Accepted: February 2, 2022