

Realization of aperiodic subshifts and uniform densities in groups

Nathalie Aubrun*, Sebastián Barbieri*
and Stéphan Thomassé*
`nathalie.aubrun@ens-lyon.fr`
`sebastian.barbieri@ens-lyon.fr`
`stephan.thomasse@ens-lyon.fr`

Abstract

A theorem of Gao, Jackson and Seward, originally conjectured to be false by Glasner and Uspenskij, asserts that every countable group admits a 2-coloring. A direct consequence of this result is that every countable group has a strongly aperiodic subshift on the alphabet $\{0, 1\}$. In this article, we use Lovász local lemma to first give a new simple proof of said theorem, and second to prove the existence of a G -effectively closed strongly aperiodic subshift for any finitely generated group G . We also study the problem of constructing subshifts which generalize a property of Sturmian sequences to finitely generated groups. More precisely, a subshift over the alphabet $\{0, 1\}$ has uniform density $\alpha \in [0, 1]$ if for every configuration the density of 1's in any increasing sequence of balls converges to α . We show a slightly more general result which implies that these subshifts always exist in the case of groups of subexponential growth.

Keywords: Symbolic dynamics, Countable groups, Amenable groups, Sturmian sequences, Aperiodic subshift.

Introduction

Symbolic dynamics is concerned with the study of subshifts on groups. Subshifts are sets of colorings of a group G by some finite alphabet \mathcal{A} that respect local constraints given by forbidden patterns, or equivalently, subsets of \mathcal{A}^G that are both closed for the product topology and shift-invariant. They can be used to model dynamical systems [12, 22], but can also be seen as computational models [13]. Subshifts of finite type (SFT) – those which can be defined by forbidding a finite set of patterns – constitute an interesting class of subshifts since they are defined by local conditions, and can model real-world phenomena.

*LIP, ENS de Lyon – CNRS – INRIA – UCBL – Université de Lyon

Classical symbolic dynamics has focused on the one-dimensional case $G = \mathbb{Z}$ [12, 15] and more recently $G = \mathbb{Z}^d$ with $d \geq 2$ [13, 21]. Subshifts on free groups have also been studied [19, 20]. Very recent results tackle computational aspects of subshifts on finitely generated groups [1, 4, 6, 7, 14].

In this article we consider two general aspects of realizability which concern subshifts in groups: The first one asks if it is possible to construct a strongly aperiodic subshift, that is, one such that the stabilizer of every element is the trivial subgroup. The second aspect is inspired by Sturmian words, by the fact that the factors of length n carry a density of 1's which converges to the slope of the irrational rotation which generates the word. Here, given a finitely generated group and its word metric given by a set of generators, we ask if it is possible to construct a subshift $X \subset \{0, 1\}^G$ such that for every configuration $x \in X$ and every sequence $(g_n)_{n \in \mathbb{N}}$ of group elements, the density of 1's over the balls $B(g_n, n)$ always converges to a fixed density $\alpha \in [0, 1]$. We call this property *uniform density*.

The existence of a countable group which does not admit a non-empty strongly aperiodic subshift over the alphabet $\{0, 1\}$ was asked in [11] and subsequently answered negatively in [10]. Nevertheless, their proof is very technical. In this article we combine the asymmetric version of Lovász local lemma [3] and compactness of the set of configurations to get a nice tool to prove non-emptiness of subshifts defined by forbidden patterns. This technique, which in some sense is the analogue of the probabilistic method in graph theory, provides very short proofs of the existence of configurations in subshifts. We use it to prove again in a succinct way the existence of a strongly aperiodic subshift on any countable group. We also extend the previous result by showing that in finitely generated groups it is also possible to construct non-empty strongly aperiodic subshifts which satisfy the condition of being G -effectively closed, that is, that they can be defined as the complement of a recursively enumerable union of cylinders by a Turing machine which has access to the word problem of G . That is, we show:

Theorem 2.4. *Every countable group G has a non-empty, strongly aperiodic subshift on the alphabet $\{0, 1\}$.*

Theorem 2.6. *Every finitely generated group G has a non-empty G -effectively closed strongly aperiodic subshift.*

A bi-infinite sequence of 0's and 1's is balanced if for every $n \in \mathbb{N}$, every factor of size n can have only two possible quantities of symbols 1. Most famous examples of balanced sequences are Sturmian sequences, which are both balanced and aperiodic, or equivalently, codifications of irrational rotations in the circle [16, 18]. Some attempts have been made to generalize Sturmian sequences to dimension 2 [5, 9]. In general groups a configuration such that the amount of 1's over any finite connected support of size n has at most two values is not possible as the group's geometry is too wild. Instead, we propose this uniform density property of a subshift mentioned above as a generalization of the aforementioned property of Sturmian subshifts. In this direction, we are able to prove

a more technical result which implies the existence of non-empty subshifts with uniform density in every group of subexponential growth. Formally, we show the following result.

Theorem 3.1. *Let G be an infinite and finitely generated amenable group and $\alpha \in [0, 1]$. Then there exists a non-empty subshift $X_\alpha \subset \{0, 1\}^G$ with $\lim_{n \rightarrow \infty} \text{dens}(1, F_n, x) = \alpha$ for any $x \in X_\alpha$ and any Følner sequence $(F_n)_{n \in \mathbb{N}}$.*

That is, by replacing the sequence of intervals of \mathbb{Z} for a general Følner sequence we produce a subshift where the densities of 1's converge to a fixed value α . We show that the subshifts given by our construction are weakly aperiodic if $\alpha \notin \mathbb{Q}$. Our example loses some of the properties of those constructions, like balancedness, but instead uses a two-symbol alphabet and is available on any infinite, finitely generated and amenable group.

1 Preliminaries

Throughout this article the groups G considered will always be countable; we denote their identity element 1_G . When G is finitely generated we associate a finite set $S \subset G$ of generators and the undirected right Cayley graph $\Gamma(G, S) = (G, \{\{g, gs\} \mid g \in G, s \in S\})$ so that (G, d) is a metric space where d is the distance induced on G by $\Gamma(G, S)$. If two words w_1, w_2 in S^* represent the same element in G , we write $w_1 =_G w_2$. We shall denote by $B(g, n) = \{h \in G \mid d(g, h) \leq n\}$ the ball of size n centered in $g \in G$. In general we denote $B_\Gamma(v, n)$ the ball of size n centered in v of an arbitrary graph Γ . For $g \in G$, we denote by $|g|$ the length of a shortest path from 1_G to g in $\Gamma(G, S)$, that is to say $|g| = d(1_G, g)$. We also denote by $\text{WP}(G) := \{w \in (S \cup S^{-1})^* \mid w =_G 1_G\}$ the set of words which can be written using elements from S and their inverses which are equal to 1_G in the group G . If $\text{WP}(G)$ is a decidable language we say that G has decidable word problem. For more references see [17].

We now give some basic definitions of symbolic dynamics. For a more complete introduction the reader may refer to [15, 8]. Let \mathcal{A} be a finite alphabet and G a countable group. The set $\mathcal{A}^G = \{x : G \rightarrow \mathcal{A}\}$ equipped with the left group action $\sigma : G \times \mathcal{A}^G \rightarrow \mathcal{A}^G$ defined by $(\sigma_g(x))_h = x_{g^{-1}h}$ is the G -fullshift. The elements $a \in \mathcal{A}$ and $x \in \mathcal{A}^G$ are called *symbols* and *configurations* respectively. By taking the discrete topology on \mathcal{A} we obtain that the set of configurations \mathcal{A}^G is compact and metrizable. In the case of a countable group, given an enumeration $1_G = g_0, g_1, \dots$ of G , the topology is generated by the metric $d(x, y) = 2^{-\inf\{n \in \mathbb{N} \mid x_{g_n} \neq y_{g_n}\}}$. If E is a subset of \mathcal{A}^G , we denote by \bar{E} its topological closure. In the case of a finitely generated group another possibility which is more practical is $d(x, y) = 2^{-\inf\{|g| \mid g \in G: x_g \neq y_g\}}$. This topology is generated by a clopen basis given by the *cylinders* $[a]_g = \{x \in \mathcal{A}^G \mid x_g = a \in \mathcal{A}\}$. A *support* is a finite subset $F \subset G$. Given a support F , a *pattern with support F* is an element p of \mathcal{A}^F , i.e. a finite configuration and we write $\text{supp}(p) = F$. We also denote the cylinder generated by p centered in g as $[p]_g = \bigcap_{h \in F} [p_h]_{gh}$. One says that a pattern $p \in \mathcal{A}^F$ *appears* in a configuration $x \in \mathcal{A}^G$ if there

exists $g \in G$ such that for any $h \in F$, $x_{gh} = p_h$, said otherwise, if there exists g such that $x \in [p]_g$. In this case we write $p \sqsubset x$. We denote the set of finite patterns over G as $\mathcal{A}_G^* := \bigcup_{F \subset G, |F| < \infty} \mathcal{A}^F$.

Definition 1.1. *A subset X of \mathcal{A}^G is a G -subshift if it is σ -invariant – $\sigma(X) \subset X$ – and closed for the cylinder topology. Equivalently, X is a G -subshift if and only if there exists a set of forbidden patterns $\mathcal{F} \subset \mathcal{A}_G^*$ that defines it.*

$$X = X_{\mathcal{F}} := \{ x \in \mathcal{A}^G \mid \forall p \in \mathcal{F}, p \not\sqsubset x \} = \bigcap_{p \in \mathcal{F}, g \in G} \mathcal{A}^G \setminus [p]_g.$$

That is, a G -subshift is a shift-invariant subset of \mathcal{A}^G which can be written as the complement of a union of cylinders. If it is clear from the context, we will drop the G and simply refer to a subshift. A subshift $X \subseteq \mathcal{A}^G$ is of *finite type* – G -SFT for short – if there exists a finite set of forbidden patterns \mathcal{F} such that $X = X_{\mathcal{F}}$.

Consider a group which is generated by a finite set S . A *pattern coding* c is a finite set of tuples $c = (w_i, a_i)_{1 \leq i \leq n}$ where $w_i \in (S \cup S^{-1})^*$ and $a_i \in \mathcal{A}$. We say that a pattern coding is *consistent* if for every pair of tuples such that $w_i =_G w_j$ (w_i and w_j represent the same element under G) then $a_i = a_j$. We say a consistent pattern coding c *codifies* a pattern P if every w_i represents an element of $\text{supp}(P)$ and for every $g \in \text{supp}(P)$ there exists a tuple $(w_i, a_i) \in c$ such that $g =_G w_i$ and $P_g = a_i$.

Definition 1.2. *For a finitely generated group G we say a subshift $X \subseteq \mathcal{A}^G$ is G -effectively closed if there exists a Turing machine with oracle $\text{WP}(G)$ which recognizes a set of pattern codings such that the consistent ones codify a set of patterns \mathcal{F} such that $X = X_{\mathcal{F}}$. If the same property is valid without the oracle we say X is effectively closed.*

Being G -effectively closed is a generalization of the same concept for \mathbb{Z} -subshifts where the set of forbidden patterns is a recursively enumerable set of words.

Let $x \in \mathcal{A}^G$ be a configuration. The *orbit* of x is the set of configurations $\text{orb}_{\sigma}(x) = \{\sigma_g(x) \mid g \in G\}$, and the *stabilizer* of x is the set of group elements $\text{stab}_{\sigma}(x) = \{g \in G \mid \sigma_g(x) = x\}$.

Definition 1.3. *A G -subshift $X \subseteq \mathcal{A}^G$ is weakly aperiodic if for every configuration $x \in X$, $|\text{orb}_{\sigma}(x)| = \infty$. A G -subshift $X \subseteq \mathcal{A}^G$ is strongly aperiodic if for every configuration $x \in X$, $\text{stab}_{\sigma}(x) = \{1_G\}$.*

For infinite groups the weak concept of aperiodicity is relevant and implied by strong aperiodicity.

2 Non-empty strongly aperiodic subshifts

In this section we construct non-empty strongly aperiodic subshifts on any countable group. As mentioned in the introduction, the question of the existence of an infinite countable group G that does not admit a non-empty strongly

aperiodic subshift over the alphabet $\{0, 1\}$ was asked in [11] and subsequently answered in the negative in [10]. In this section we present a short proof based on Lovász local lemma [3]. For finitely generated groups, we also give a second proof – inspired by the use of local lemma in [2] – which is quite easy to visualize and gives a G -effectively closed subshift, but this proof uses a large alphabet and only works in finitely generated groups.

We begin by introducing the asymmetric version of the local lemma. We then extract a corollary to show how it can be used in order to produce configurations in subshifts by using the compactness of the set of configurations and then we proceed to the construction of the strongly aperiodic subshifts.

2.1 Lovász local lemma

Lemma 2.1. *[Asymmetric Lovász local lemma] Let $\mathcal{A} := \{A_1, A_2, \dots, A_n\}$ be a finite collection of measurable sets in a probability space (X, μ, \mathcal{B}) . For $A \in \mathcal{A}$, let $\Gamma(A)$ be the smallest subset of \mathcal{A} such that A is independent of the collection $\mathcal{A} \setminus (\{A\} \cup \Gamma(A))$. Suppose there exists a function $x : \mathcal{A} \rightarrow (0, 1)$ such that:*

$$\forall A \in \mathcal{A} : \mu(A) \leq x(A) \prod_{B \in \Gamma(A)} (1 - x(B))$$

then the probability of avoiding all events in \mathcal{A} is positive. Specifically

$$\mu \left(X \setminus \bigcup_{i=1}^n A_i \right) \geq \prod_{A \in \mathcal{A}} (1 - x(A)) > 0.$$

The sets A_1, A_2, \dots, A_n can be seen as bad events that we want to avoid. In the context of the present article where \mathcal{A} is a finite alphabet and G a countable group, we will choose the probability space to be $X = \mathcal{A}^G$ with μ any Bernoulli probability measure on \mathcal{A}^G . Suppose X is a subshift defined by a set of forbidden patterns $\mathcal{F} = \bigcup_{n \geq 1} \mathcal{F}_n$ where $\mathcal{F}_n \subset \mathcal{A}^{S_n}$ is a finite set of patterns with finite support S_n . We will consider the bad events $A_{n,g} = \bigcup_{p \in \mathcal{F}_n} [p]_g = \{x \in \mathcal{A}^G : x|_{gS_n} \in \mathcal{F}_n\}$, that is to say one of the forbidden patterns $p \in \mathcal{F}_n$ appears in position g . Subshifts might be defined by an infinite amount of forbidden patterns while the lemma only holds for a finite collection of bad events. Nevertheless the compactness of \mathcal{A}^G allows us to use the lemma anyway, as explained in what follows.

Lemma 2.2. *Let G a countable group and $X \subset \mathcal{A}^G$ a subshift defined by the set of forbidden patterns $\mathcal{F} = \bigcup_{n \geq 1} \mathcal{F}_n$, where $\mathcal{F}_n \subset \mathcal{A}^{S_n}$. Let μ be a Bernoulli probability measure on \mathcal{A}^G . Suppose that there exists a function $x : \mathbb{N} \times G \rightarrow (0, 1)$ such that:*

$$\forall n \in \mathbb{N}, g \in G, \mu(A_{n,g}) \leq x(n, g) \prod_{\substack{gS_n \cap hS_k \neq \emptyfty \\ (k,h) \neq (n,g)}} (1 - x(k, h)), \quad (*)$$

where $A_{n,g} = \{x \in \mathcal{A}^G : x|_{gS_n} \in \mathcal{F}_n\}$. Then the subshift X is non-empty.

Proof. Consider an enumeration $\{g_k\}_{k \in \mathbb{N}}$ of G . For every $n \in \mathbb{N}$, we apply Lemma 2.1 to construct a configuration $x_n \in \mathcal{A}^G$ that satisfies the following property: for every forbidden pattern $p \in \mathcal{F}_k$ such that $k \leq n$ and every $g \in (g_k)_{k \leq n}$ such that $gS_k \subseteq \{g_k\}_{k \leq n}$, we have $(x_n) \notin [p]_g$ – in other terms, the configuration x_n avoids all the forbidden patterns in $\bigcup_{k \leq n} \mathcal{F}_k$ on the finite set $\{g_1, \dots, g_n\} \subset G$. Indeed, in order to show the existence of x_n we only need that for every $k \leq n$ and $g \in G$ such that $gS_k \subseteq \{g_k\}_{k \leq n}$,

$$\mu(A_{k,g}) \leq x(k,g) \prod_{\substack{gS_k \cap hS_{k'} \neq \emptyset \\ hS_{k'} \subseteq \{g_0, \dots, g_n\} \\ (k',h) \neq (k,g), k \leq n}} (1 - x(k',h))$$

which is a relaxation of condition (*) by the fact that $0 \leq x(k',h) \leq 1$. The local lemma holds since the set $\{g_0, \dots, g_n\}$ is finite and we only consider a finite number of forbidden patterns, consequently we only consider a finite number of bad events $A_{k,g}$.

By compactness, we can extract from this sequence of configurations $(x_n)_{n \in \mathbb{N}}$ a subsequence $(x_{\phi(n)})_{n \in \mathbb{N}}$ converging to some $x \in \mathcal{A}^G$. Then x does not contain any forbidden pattern $p \in \mathcal{F}$. Suppose it were the case, that is to say, there exists some $g \in G$ and $p \in \mathcal{F}_m$ such that $x \in [p]_g$. Since there exists also n, l such that $g = g_l$ and $gS_m \subset \{g_0, \dots, g_n\}$, by definition of the metric there exists some $N \geq \max\{m, n, l\}$ sufficiently big such that $\phi(N)$ appears in the subsequence $(\phi(n))_{n \in \mathbb{N}}$. Then x_N contains the forbidden pattern p somewhere in $(x_N)|_{\{g_1, \dots, g_N\}}$. This contradicts the construction of the sequence $(x_n)_{n \in \mathbb{N}}$, thus $x \in X_{\mathcal{F}}$. \square

2.2 A non-empty strongly aperiodic subshift over $\{0, 1\}$ in any countable group.

Consider a configuration $x \in \{0, 1\}^G$. We say that x has *the distinct neighborhood property* – in [10] they call x a 2-coloring – if for every $h \in G \setminus \{1_G\}$ there exists a finite subset $T \subset G$ such that:

$$\forall g \in G : x|_{ghT} \neq x|_{gT}.$$

Proposition 2.3. *If a configuration $x \in \{0, 1\}^G$ has the distinct neighborhood property, then the G -subshift $X := \overline{\text{orb}_\sigma(x)}$ is strongly aperiodic.*

Proof. Let $y \in X$. By definition there exists a sequence $(g_i)_{i \in \mathbb{N}}$ such that $\sigma_{g_i}(x)$ converges to y in the product topology. Suppose there is $h \neq 1_G$ such that $\sigma_h(y) = y$, then by continuity of the shift action under the product topology we have that $\sigma_{hg_i}(x) \rightarrow \sigma_h(y) = y$. Since x has the distinct neighborhood property, there exists a finite subset T of G – associated to h^{-1} – such that $\forall g \in G : x|_{gh^{-1}T} \neq x|_{gT}$. By definition of convergence in the metric, there exists $n \in \mathbb{N}$ such that $T \subset \{g_0, g_1, \dots, g_n\}$ and $m \in \mathbb{N}$ satisfying:

$$\sigma_{hg_m}(x)|_{\{g_0, g_1, \dots, g_n\}} = y|_{\{g_0, g_1, \dots, g_n\}} = \sigma_{g_m}(x)|_{\{g_0, g_1, \dots, g_n\}}$$

Therefore $\sigma_{hg_m}(x)|_T = \sigma_{g_m}(x)|_T$ which implies that $x|_{g_m^{-1}h^{-1}T} = x|_{g_m^{-1}T}$, a contradiction. \square

Theorem 2.4. *Every countable group G has a non-empty, strongly aperiodic subshift on the alphabet $\{0, 1\}$.*

Proof. The case where G is finite is trivial, as the G -SFT given by

$$X := \{x \in \{0, 1\}^G \mid |x^{-1}(1)| = 1\}$$

is strongly aperiodic. Indeed, let $x \in X$ and $g \in G$ be the only element such that $x_g = 1$. Let $h \in \text{stab}_\sigma(x)$ then $\sigma_h(x) = x$ which implies that $x_{h^{-1}g} = x_g = 1$ and thus $h = 1_G$. For the rest of the proof we suppose that G is infinite.

Let $(s_i)_{i \in \mathbb{N}}$ be an enumeration of G such that $s_0 = 1_G$. Choose $(T_i)_{i \in \mathbb{N}}$ a sequence of finite subsets of G such that for every $i \in \mathbb{N}$, $T_i \cap s_i T_i = \emptyset$ and $|T_i| = C \cdot i$, where C is a constant to be defined later. These sets always exist as G is infinite.

Consider the uniform Bernoulli probability μ in $\{0, 1\}^G$ and the collection of sets $\mathcal{A} := \{A_{n,g}\}_{n \geq 1, g \in G}$ given by $A_{n,g} = \{x \in \{0, 1\}^G \mid x|_{gT_n} = x|_{gs_n T_n}\}$. Note that each set is a union of cylinders and that the existence of a configuration \tilde{x} in the intersection of the complement of these sets allows us to conclude the theorem by producing a configuration with the distinct neighborhood property. Our strategy is to apply Lemma 2.2 to ensure its existence.

As the intersection $s_n T_n \cap T_n$ is empty we have that $\mu(A_{n,g}) = 2^{-|T_n|} = 2^{-Cn}$. Consider one set $A_{n,g}$. The number of sets $A_{m,g'}$ for a fixed $m \in \mathbb{N}$ which are not independent from $A_{n,g}$ is at most $4C^2 nm$ – observe that $A_{n,g}$ is independent from $A_{m,g'}$ if and only if $(gT_n \cup gs_n T_n)$ does not intersect $(g'T_m \cup g's_m T_m)$. We also define $x(A_{n,g}) := 2^{-\frac{Cn}{2}}$. Therefore, in order to conclude we must prove that:

$$2^{-Cn} \leq 2^{-\frac{Cn}{2}} \prod_{m=1}^{\infty} (1 - 2^{-\frac{Cm}{2}})^{4C^2 nm}.$$

Using the fact that $1 - x \geq 2^{-2x}$ if $x \leq 1/2$ we obtain the following bound:

$$\begin{aligned} 2^{-\frac{Cn}{2}} \prod_{m=1}^{\infty} (1 - 2^{-\frac{Cm}{2}})^{4C^2 nm} &\geq 2^{-\frac{Cn}{2}} \prod_{m=1}^{\infty} 2^{-8C^2 nm 2^{-\frac{Cm}{2}}} \\ &= 2^{-\frac{Cn}{2}} 2^{-8C^2 \sum_{m=1}^{\infty} nm 2^{-\frac{Cm}{2}}} \end{aligned}$$

Therefore, it suffices to prove that:

$$\begin{aligned}
2^{-\frac{Cn}{2}} &\leq 2^{-8C^2 \sum_{m=1}^{\infty} nm2^{-\frac{Cm}{2}}} \\
\iff 1 &\geq 16C \sum_{m=1}^{\infty} m2^{-\frac{Cm}{2}} \\
\iff 1 &\geq 16C \frac{2^{\frac{C}{2}}}{(2^{\frac{C}{2}} - 1)^2}
\end{aligned}$$

The previous inequality holds true for $C \geq 17$. Therefore choosing $C = 17$ completes the proof by application of Lemma 2.2. \square

2.3 A graph-oriented construction and some computational properties

In this subsection we present another construction of a non-empty strongly aperiodic subshift. This construction is not as general as the previous one, as it only holds for finitely generated groups, and the size of the alphabet is rather large. Nevertheless, it can be defined by a natural property which allows us to use it in computability constructions with ease.

Let $\Gamma = (V, E)$ be a simple graph, consider a finite alphabet \mathcal{A} and a coloring $x \in \mathcal{A}^V$ of the vertices of Γ . We say x contains a *vertex-square path* if there exists an odd length path $p = v_1 \dots v_{2n}$ such that $x_{v_i} = x_{v_{i+n}}$ for every $1 \leq i \leq n$. If the coloring x does not contain any vertex-square path then we say it is a *square-free vertex coloring*. Next we show a proposition which is a slight modification of a proof which can be found in [2, Theorem 1].

Proposition 2.5. *Let G be a group which is generated by the finite set S . Then there exists a square-free vertex coloring of the undirected right Cayley graph $\Gamma(G, S)$ with $2^{19}|S|^2$ colors.*

Proof. Consider a finite alphabet \mathcal{A} and let $X = \mathcal{A}^{\Gamma(G, S)}$ be the set of all vertex colorings of the Cayley graph $\Gamma(G, S)$. We define μ as the uniform Bernoulli probability, that is, for $a \in \mathcal{A}$ and $g \in G$ then

$$\mu(\{x \in X \mid x_g = a\}) = \frac{1}{|\mathcal{A}|}.$$

Consider \mathcal{P} as the set of all odd length paths in $\Gamma(G, S)$. For $p \in \mathcal{P}$ let A_p be the set of colorings of $\Gamma(G, S)$ such that p is a square under that coloring and note that if p is of length $2n-1$ then $\mu(A_p) = |\mathcal{A}|^{-n}$ if there exists a path of said length. Consider $\mathcal{A}_n = \{A_p \mid p \text{ has length } 2n-1\}$ and $\mathcal{A} = \{A_p \mid p \in \mathcal{P}\} = \bigcup_{n \geq 1} \mathcal{A}_n$. In order to apply Lemma 2.2, we define $x(A_p) := (8|S|^2)^{-n}$ for $A_p \in \mathcal{A}_n$. The lemma holds if for every $A \in \mathcal{A}$ then $\mu(A) \leq x(A) \prod_{B \in \Gamma(A)} (1 - x(B))$. Replacing both sides yields the necessary condition:

$$\forall n \geq 1, |\mathcal{A}|^{-n} \leq (8|S|^2)^{-n} \prod_{j \geq 1} (1 - (8|S|^2)^{-j})^{|\Gamma(A_p) \cap \mathcal{A}_j|}.$$

$|\Gamma(A_p) \cap \mathcal{A}_j|$ corresponds to the amount of paths of length $2j-1$ which share a vertex with p . If p has length $2n-1$ this can be bounded by $4nj(2|S|)^{2j}$. Indeed, there are at most $(2|S|)^{2j}$ paths of length $2j-1$ starting from a vertex v . Each of these paths can intersect a given vertex of p in $2j$ positions and there are $2n$ vertices in p . Hence, we need to show:

$$\forall n \geq 1, |\mathcal{A}|^{-n} \leq (8|S|^2)^{-n} \prod_{j \geq 1} (1 - (8|S|^2)^{-j})^{4nj(2|S|)^{2j}}.$$

Using the inequality $1 - x \geq 2^{-2x}$ if $x \leq 1/2$, the requirement to apply the lemma can be restrained further so that the following is required to conclude:

$$\begin{aligned} \forall n \geq 1, |\mathcal{A}|^{-n} &\leq (8|S|^2)^{-n} \prod_{j \geq 1} 2^{-8nj(8|S|^2)^{-j}(4|S|^2)^j} \\ &= (8|S|^2)^{-n} \prod_{j \geq 1} 2^{-8nj2^{-j}} \end{aligned}$$

or equivalently:

$$\begin{aligned} |\mathcal{A}| &\geq (8|S|^2) 2^{8 \sum_{j \geq 1} j 2^{-j}} \\ &\geq 2^{19} |S|^2. \end{aligned}$$

The latter inequality is satisfied by hypothesis, therefore, there exists a coloring of the graph such that no path of odd length is a square under that coloring. \square

Theorem 2.6. *Every finitely generated group G has a non-empty G -effectively closed strongly aperiodic subshift.*

Proof. Let S be a set of generators of G and \mathcal{A} an alphabet such that $|\mathcal{A}| \geq 2^{19}|S|^2$. Consider the set of forbidden patterns \mathcal{F} defined as follows: Take \mathcal{P} the set of all finite paths of odd length of $\Gamma(G, S)$. For every $p \in \mathcal{P}$ we define the set of patterns Π_p as those with support p and such that they are vertex-square paths. Let $\mathcal{F} = \bigcup_{p \in \mathcal{P}} \Pi_p$ and let $X = X_{\mathcal{F}}$ be the G -effectively closed subshift – vertex-square paths can be recognized with a Turing machine with access to $\text{WP}(G)$ – defined by this set of forbidden patterns. By Proposition 2.5 this subshift is non-empty. We claim it is strongly aperiodic.

Let $x \in X$ and $g \in \text{stab}_{\sigma}(x)$. We are going to show that if $g \neq 1_G$ then x contains a vertex-square path. Consider an expression of g as an element of $(S \cup S^{-1})^*$ such it can be factorized as $g =_G u w v$ with $u =_G v^{-1}$. This can always be done by choosing $u = v = \varepsilon$ and w a product of generators producing g . Amongst all those possible representations choose one such that

$|w|$ is minimal. Clearly $|w| = 0$ implies that $g = 1_G$, so we suppose $|w| > 0$. Let $w = w_1 \dots w_n$ and consider the odd length walk $\pi = v_0 v_1 \dots v_{2n-1}$ defined by:

$$v_i = \begin{cases} 1_G & \text{if } i = 0 \\ w_1 \dots w_i & \text{if } i \in \{1, \dots, n\} \\ w w_1 \dots w_{i-n} & \text{if } i \in \{n+1, \dots, 2n-1\} \end{cases}$$

We claim that π is a path. Indeed, by definition, w can not be reduced and thus there are no repeated vertices in $v_0 v_1 \dots v_n$ nor in $v_{n+1} \dots v_{2n-1}$. Therefore if there is a repeated vertex then it appears once in both parts. Suppose that it happens, thus we can consider two factorizations $w = ab$ and $w = cd$ such that $a =_G abc$. We obtain that $b = c^{-1}$. Obviously $|c| = |b|$, if not, w can be written as $abcc^{-1} =_G ac^{-1}$ or $b^{-1}bcd =_G b^{-1}d$ which contradicts the minimality of $|w|$. Without loss of generality, we can replace c by the word obtained by reversing the order and inverting the letters of b . Moreover, $|c| > 0$ and thus $|b| > 0$ which means that w is written as follows:

$$w = a_1 \dots a_k b_1 \dots b_l = b_l^{-1} \dots b_1^{-1} d_1 \dots d_k$$

Therefore we can factorize b_l^{-1} and b_l from both sides obtaining a smaller word w' in the representation of g . This contradiction shows that indeed π is a path. We now show that it is a square-vertex path. As $g \in \text{stab}_\sigma(x)$, we also have $g^{-1} \in \text{stab}_\sigma(x)$ and thus for each $h \in G$ we have $x_{gh} = x_h$. Fix $h = u w_1 \dots w_j$ for some $j \in \{1, \dots, n\}$, we obtain that

$$x_{u w_1 \dots w_j} = x_{u w u^{-1} u w_1 \dots w_j} = x_{u w w_1 \dots w_j}.$$

Applying σ^u to both sides we obtain that $x_{w_1 \dots w_j} = x_{w w_1 \dots w_j}$ and therefore $x_{v_j} = x_{v_{j+n}}$, yielding a square-vertex path. Therefore $|w| = 0$ and thus $g = u w = 1_G$. \square

Theorem 2.6 provides a non-empty strongly aperiodic G -effectively closed subshift. Recently Jeandel [14] has shown that for recursively presented groups, if the group admits an effectively closed strongly aperiodic subshift then its word problem is decidable. Moreover, he has shown that the same conclusion stands when we allow every configuration to have a finite –instead of trivial– stabilizer. Our result actually shows the converse, that is, that every recursively presented group with decidable word problem admits a non-empty strongly aperiodic effectively closed subshift. In the remainder of this section we prove this equivalence.

Lemma 2.7. *Let G be a finitely generated group and $X \subset \mathcal{A}^G$ a non-empty strongly aperiodic subshift. There exists a function $f : \mathbb{N} \rightarrow \mathbb{N}$ such that for every $x \in X$, if $g \in B(1_G, n) \setminus \{1_G\}$ then $x|_{B(1_G, f(n))} \neq x|_{B(g, f(n))}$.*

Proof. Suppose f does not exist, thus there exists $n \in \mathbb{N}$ and a sequence $(x_j, g_j)_{j \in \mathbb{N}} \subset X \times B(1_G, n) \setminus \{1_G\}$ such that $x_j|_{B(1_G, j)} = x_j|_{B(g_j, j)}$. As $B(1_G, n)$ is finite there exists $\bar{g} \neq 1_G$ which appears infinitely often in $(g_j)_{j \in \mathbb{N}}$. Consider

the subsequence $(x_k)_{k \in \mathbb{N}, g_k = \bar{g}}$ and from there extract a converging subsequence $(x_{k_\alpha}) \rightarrow \bar{x} \in X$. We claim that $\bar{g}^{-1} \in \text{stab}_\sigma(\bar{x})$. By definition of convergence, for every $m \in \mathbb{N}$ there exists $N_\alpha \geq m$ such that $\bar{x}|_{B(1_G, m+n)} = (x_{N_\alpha})|_{B(1_G, m+n)}$ and thus

$$\bar{x}|_{B(1_G, m)} = (x_{N_\alpha})|_{B(1_G, m)} = (x_{N_\alpha})|_{B(\bar{g}, m)} = \bar{x}|_{B(\bar{g}, m)}$$

So for every $m \in \mathbb{N}$ we have $\bar{x}|_{B(1_G, m)} = \bar{x}|_{B(\bar{g}, m)}$ and therefore $\forall g \in G : \bar{x}_g = \bar{x}_{\bar{g}g}$. Which yields a contradiction as X is strongly aperiodic. \square

Theorem 2.8. *Let G be a recursively presented and finitely generated group. There exists a non-empty strongly aperiodic effectively closed subshift if and only if $\text{WP}(G)$ is decidable.*

Proof. Every G -effectively closed subshift is effectively closed when $\text{WP}(G)$ is decidable. Therefore Theorem 2.6 yields the desired construction. Conversely, suppose there is a non-empty effectively closed subshift X which is strongly aperiodic. As G is recursively presented then $\text{WP}(G)$ is recognizable. Let T be a Turing machine which accepts every inconsistent pattern coding and a maximal set of consistent pattern codings which generates \mathcal{F} such that $X = X_{\mathcal{F}}$. The existence of such a machine in the case of a recursively presented group is given in [1].

Let $w \in (S \cup S^{-1})^*$. We present an algorithm which accepts if and only if $w \neq_G 1_G$. Consider the ball of size n in the free monoid over the alphabet $(S \cup S^{-1})^*$, that is $\Lambda_n = \{u \in (S \cup S^{-1})^* \mid |u| \leq n\}$ and consider the set $\Lambda_n \cup w\Lambda_n$. For each one of these sets we construct the set Π_n of all pattern codings c such that for $u \in \Lambda_n$ then $(u, a) \in c$ if and only if $(wu, a) \in c$. That is, we force the ball of size n around the empty word ϵ and w to be the same. Consider the algorithm which iteratively runs T on every pattern coding of $\Pi_1, \Pi_2, \dots, \Pi_j$ up to j steps and then does $j \leftarrow j + 1$ and which accepts w if and only if every pattern coding in a particular Π_i is accepted by T . If $w =_G 1_G$ the algorithm can never accept as it would mean no patterns are constructible around 1_G and thus $X = \emptyset$. Conversely, if $w \neq_G 1_G$ then using the function f given by Lemma 2.7 we get that for every $x \in X$ if $w \neq_G 1_G$ then $x|_{B(1_G, f(|w|))} \neq x|_{B(w, f(|w|))}$ thus every pattern in $\Pi_{f(|w|)}$ is either inconsistent or represents a forbidden pattern, and therefore T must accept every pattern of $\Pi_{f(|w|)}$. \square

One may ask if it is possible to construct non-empty strongly aperiodic subshifts which satisfy stronger constraints, such as being of finite type, sofic or effectively closed. The previous result shows that our construction is in this sense optimal for recursively presented groups with undecidable word problem.

3 Realization of densities

In this section we construct over any infinite and finitely generated group a non-empty subshift over $\{0, 1\}$ with the property that the density of 1's over any Følner sequence converges to a fixed $\alpha \in [0, 1]$. From this result we derive

the existence of uniform density subshifts for infinite groups of subexponential growth for any finite set of generators. Furthermore, we show that said subshifts are always weakly aperiodic.

Definition 3.1. Let $F \subset G$ be a finite subset of a group and $x \in \{0, 1\}^G$ be a configuration. We define the density of 1's in F and x as:

$$\text{dens}(1, x|_F) = \text{dens}(1, F, x) = \frac{|\{g \in F \mid x_g = 1\}|}{|F|}.$$

Similarly if $P \in \mathcal{A}^F$ is a pattern with support F , we denote by $\text{dens}(1, P)$ the ratio $\frac{|\{g \in F \mid P_g = 1\}|}{|F|}$.

Definition 3.2. Let G be a finitely generated group and S a finite set of generators. We say a G -subshift over $\{0, 1\}$ has uniform density $\alpha \in [0, 1]$ if for every configuration $x \in X$ and for every sequence $(g_n)_{n \in \mathbb{N}}$ of elements in G , one has $\text{dens}(1, B(g_n, n), x) \rightarrow \alpha$.

In a way similar to the previous definition, we could say a configuration $x \in \{0, 1\}^G$ has density $\alpha \in [0, 1]$ for some sequence of subsets $(T_n)_{n \in \mathbb{N}}$ if for each sequence of elements $(g_n)_{n \in \mathbb{N}}$ we have that $\text{dens}(1, g_n T_n, x) \rightarrow \alpha$. Nevertheless, contrary to the preceding section, Lovász local lemma cannot directly be applied to prove the existence of such configurations. If we define the forbidden sets to be $A_{n,g} = \{x \in \{0, 1\}^G \mid |\text{dens}(1, g T_n, x) - \alpha| > \delta_n \alpha\}$ for some sequence of error terms $\delta_n \rightarrow 0$ we obtain that the measure of this set can be bounded above using the Chernoff bounds by $2 \exp(-\delta_n^2 \alpha |T_n|/3)$. For any function which bounds these values above, and after some elimination of exponents, we obtain that the left hand side of the inequality required by the local lemma depends on δ_n while the right hand side is constant. Therefore we tackle this problem with a different approach.

Nevertheless, if the condition that the group is amenable is added, not only it is possible to obtain a result like the one defined in the previous paragraph, moreover, the density over every Følner sequence can be asked to converge to the same fixed α .

A group G is called *amenable* if there exists a left-invariant finitely additive probability measure $\mu : \mathcal{P}(G) \rightarrow [0, 1]$ on G . The amenability of a group has many equivalent definitions – many of which can be found in [8]. From a combinatorial point of view the Følner condition states that a group is amenable if and only if it admits a Følner net, that is, a net F_α of non-empty finite subsets of G such that $\forall g \in G$:

$$\lim_{\alpha} \frac{|g F_\alpha \triangle F_\alpha|}{|F_\alpha|} = 0.$$

Let $\text{Int}(F, K) := \{g \in F \mid \forall k \in K, gk \in F\}$ be the interior of F with respect to K and $\partial_K F := F \setminus \text{Int}(F, K)$ the boundary of F with respect to K . In the case of countable groups the net can be just taken to be a sequence and thus amenability can be shown to be equivalent to the fact for every finite $K \subset G$ we have $\lim_{n \rightarrow \infty} \frac{|\partial_K F_n|}{|F_n|} = 0$. That is to say, for any finite set K the boundaries of the sets F_n with respect to K grow slower than themselves.

Theorem 3.1. *Let G be an infinite and finitely generated amenable group and $\alpha \in [0, 1]$. Then there exists a non-empty subshift $X_\alpha \subset \{0, 1\}^G$ with $\lim_{n \rightarrow \infty} \text{dens}(1, F_n, x) = \alpha$ for any $x \in X_\alpha$ and any Følner sequence $(F_n)_{n \in \mathbb{N}}$.*

Before proving this theorem we prove a useful property of metric spaces.

Definition 3.3. *Let (X, d) be a metric space. We say $F \subset X$ is r -covering if for each $x \in X$ there is $y \in F$ such that $d(x, y) \leq r$. We say F is s -separating if for each $x \neq y \in F$ then $d(x, y) > s$.*




Figure 1: In green, an example of 2-covering and 2-separating set in $PSL(2, \mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/3\mathbb{Z}$. Green vertices are at distance at least 3 from each other, and every vertex is at distance at most 2 from a green vertex.

Lemma 3.2. *Let (X, d) be a metric space and $r \in \mathbb{N}$. There exists a set $F_r \subset X$ which is r -separating and r -covering.*

Proof. Suppose we have an r -separating set F which is not r -covering. Then the set $K := \{x \in X \mid d(F, x) > r\}$ is not empty and F can be extended by an element of K . Thus any maximal r -separating set is also r -covering.

We only need to show that maximal r -separating sets exist. Let (\mathcal{S}, \subset) be the set of r -separating subsets of X ordered by inclusion. Clearly $\emptyset \in \mathcal{S}$ and given a chain $\{A_i\}_{i \in I} \subset \mathcal{S}$ we have that $A = \bigcup_{i \in I} A_i$ is an upper-bound. Indeed, if $x, y \in A$ then $x \in A_i$ and $y \in A_j$ for some $i, j \in I$. As $\{A_i\}_{i \in I}$ is a chain, then either $A_i \subset A_j$ or $A_j \subset A_i$. As any of these two sets is r -separating we get that $d(x, y) > r$ and hence $A \in \mathcal{S}$. By Zorn's lemma there exists a maximal r -separating set. \square

Now we are ready to begin the proof of Theorem 3.1.

Proof. If $\alpha \in \{0, 1\}$ the result is trivial. Let $\alpha \in (0, 1)$, and define $K_n := B(1_G, 5^n)$ and consider the subshift X_α given by the set of forbidden patterns \mathcal{F} such that for $P \in \{0, 1\}^F$ (where $F \subset G$, $|F| < \infty$) belongs to \mathcal{F} if and only if the following condition is not satisfied:

$$2n|\partial_{K_n} F| < |F| \implies |\text{dens}(1, P) - \alpha| \leq \frac{1}{n}$$

In other words, we forbid a pattern P with support F if the ratio $\frac{|\partial_{K_n} F|}{|F|}$ is sufficiently small and the density of ones in P is further than $\frac{1}{n}$ from α .

Consider a Følner sequence $(F_n)_{n \in \mathbb{N}}$ and let $m \in \mathbb{N}$ and $x \in X_\alpha$. As $\lim_{n \rightarrow \infty} \frac{|\partial_{K_m} F_n|}{|F_n|} = 0$ there exists $N \in \mathbb{N}$ such that

$$\forall n \geq N \quad \frac{|\partial_{K_m} F_n|}{|F_n|} < \frac{1}{2m}$$

Therefore, for every $n \geq N$ we get that $|\text{dens}(1, x|_{F_n}) - \alpha| \leq \frac{1}{m}$. As m can be made arbitrarily big we obtain that $\lim_{n \rightarrow \infty} \text{dens}(1, F_n, x) = \alpha$.

We only need to show that $X_\alpha \neq \emptyset$. Our strategy will be to inductively construct an infinite covering forest of G , and then put a Sturmian word along an enumeration of the leaves of each of its trees. The configuration $x \in \{0, 1\}^G$ obtained by this process will belong to X_α . The following objects – that are formally described below – will be used to formalize this idea: a sequence $(A_n)_{n \in \mathbb{N}} \subset 2^G$ of subsets of G , a sequence $(p_n)_{n \in \mathbb{N}} : G \rightarrow A_n$ of functions and a sequence $(\Gamma_n)_{n \in \mathbb{N}}$ of graphs on vertex sets $(A_n)_{n \in \mathbb{N}}$ respectively. They are defined by the following recurrences, with base cases $A_0 = G$, $p_0 = id$ and $\Gamma_0 = \Gamma(G, S)$ where S is a finite set of generators of G :

1. The set A_{n+1} is chosen as a 2-separating and 2-covering subset of A_n for the distance induced by Γ_n . In particular, the sets $(A_n)_{n \in \mathbb{N}}$ are nested.
2. Suppose $p_n : G \rightarrow A_n$ is already defined, we first define p_{n+1} on A_n and then extend it to G . Consider an element $g \in A_n$. Since A_{n+1} 2-covers A_n in Γ_n , there are only three possible cases.
 - $g \in A_{n+1}$: in this case we set $p_{n+1}(g) = g$.
 - $d_{\Gamma_n}(g, A_{n+1}) = 1$: there exists a unique $h \in A_{n+1}$ that satisfies $d_{\Gamma_n}(g, h) = 1$ – uniqueness comes from the fact that A_{n+1} is 2-separating – and we set $p_{n+1}(g) = h$.
 - $d_{\Gamma_n}(g, A_{n+1}) = 2$: we arbitrarily choose one $h \in A_{n+1}$ that realizes $d_{\Gamma_n}(g, h) = 2$ and set $p_{n+1}(g) = h$.

For $g' \in G \setminus A_n$ we finally extend this function by $p_{n+1} := p_{n+1} \circ p_n$.

3. For $g \in A_n$ define the n -cluster of g by $\mathcal{C}_n(g) := \{h \in G \mid p_n(h) = g\}$. The element $g \in A_n$ is called the *center* of the cluster $\mathcal{C}_n(g)$. The graph Γ_{n+1}

has vertex set A_{n+1} , and there is an edge in Γ_{n+1} between two elements $g, h \in A_{n+1}$ if and only if there exist $g' \in \mathcal{C}_n(g)$ and $h' \in \mathcal{C}_n(h)$ that are neighbors in $\Gamma(G, S)$.

The *covering forest* defined by the sequence $(A_n, p_n, \Gamma_n)_{n \in \mathbb{N}}$ is (V, E) , where the set of vertices V is the multiset $\bigsqcup_{n \in \mathbb{N}} A_n$, and the edges are given by the parent functions: $(g, h) \in E$ if and only if $g \in A_n$, $h \in A_{n+1}$ and $p_n(g) = h$. In particular the successive applications of p_1, \dots, p_n to an element $g \in G = A_0$ gives the path from the leaf labeled by g to its height n parent. The cluster $\mathcal{C}_n(g)$ corresponds to the set of labels of descendants of the node labeled by g that appears at height n in the covering forest. The cluster $\mathcal{C}_{n+1}(g)$ is obtained as the union of the cluster $\mathcal{C}_n(g)$, all clusters $\mathcal{C}_n(h)$ for $h \in A_n$ such that $d_{\Gamma_n}(g, h) = 1$ and clusters $\mathcal{C}_n(h')$ for $h' \in A_n$ such that $d_{\Gamma_n}(g, h') = 2$ for which the parent function $p_{n+1}(h')$ has been chosen to be g (see Figure 2). Remark that every cluster $\mathcal{C}_n(g)$ is connected in Γ as it is the finite union of adjacent connected sets in Γ .




Figure 2: A covering forest of G . In the left section of the image the edge structure is emphasized by writing explicitly the parent functions. In the right section we remark the cluster structure for $g \in A_2$.

Note that definition 3 above is equivalent to what follows: for $g, h \in A_{n+1}$ then the edge (g, h) is in $E(\Gamma_{n+1})$ if and only if there exists a path $g_1 = g, g_2, \dots, g_m = h$ from g to h in $\Gamma(G, S)$ such that for every $i \in \{1, \dots, m\}$ we have $p_{n+1}(g_i) \in \{g, h\}$.

Claim 3.1. *Let $g \in A_n$, then $B(g, n) \subset \mathcal{C}_n(g) \subset B(g, \frac{1}{2}(5^n - 1))$.*

Proof. We prove the claim by induction. It stands true for $n = 0$.

- Consider $\mathcal{C}_n(g)$. By induction hypothesis, $B(g, n-1) \subset \mathcal{C}_{n-1}(g) \subset \mathcal{C}_n(g)$. Let $h \in B(g, n) \setminus B(g, n-1)$. Either $h \in \mathcal{C}_{n-1}(g)$ and we are done, or $h \in \mathcal{C}_{n-1}(g')$ for some $g' \in A_{n-1}$. Then necessarily $d_{\Gamma_{n-1}}(g, g') = 1$, since

$hs \in B(g, n-1) \subset \mathcal{C}_{n-1}(g)$ for some $s \in S \cup S^{-1}$. Finally as A_n is a 2-separating subset of the vertices of Γ_{n-1} we get that $\mathcal{C}_{n-1}(g') \subset \mathcal{C}_n(g)$ thus $h \in \mathcal{C}_n(g)$. We conclude that $B(g, n) \subset \mathcal{C}_n(g)$. Note that the same argument proves that $\mathcal{C}_n(g') \cdot (S \cup S^{-1}) \subset \mathcal{C}_{n+1}(g')$ for every $n \in \mathbb{N}$ and $g' \in A_{n+1}$.

- Suppose inductively that for every $g \in A_{n-1}$ the inclusion $\mathcal{C}_{n-1}(g) \subset B(g, \frac{1}{2}(5^{n-1} - 1))$ holds. Fix one $g \in A_n$ and consider an element h in the cluster $\mathcal{C}_n(g)$. We show that $d_G(h, g) \leq \frac{1}{2}(5^n - 1)$ by constructing a path of length at most $\frac{1}{2}(5^n - 1)$ from h to g . By definition of the cluster $\mathcal{C}_n(g)$, we know that the element $h' \in A_{n-1}$ such that $h \in \mathcal{C}_{n-1}(h')$ satisfies $d_{\Gamma_{n-1}}(g, h') \leq 2$. In the sequel we will only consider the case where this distance is exactly 2 as it is the worst case. Thus we assume that there exists a path $h' \rightarrow h'' \rightarrow g$ of length 2 between this h' and g in Γ_{n-1} . By definition of the graph Γ_{n-1} , this implies the existence of $k' \in \mathcal{C}_{n-1}(h')$ and $k'' \in \mathcal{C}_{n-1}(h'')$ that are neighbors in $\Gamma(G, S)$ and $\ell'' \in \mathcal{C}_{n-1}(h'')$ and $\ell \in \mathcal{C}_{n-1}(g)$ that are neighbors in $\Gamma(G, S)$. Putting everything together, we can build the following path in $\Gamma(G, S)$ (see Figure 3):

$$h \rightarrow \dots \rightarrow h' \rightarrow \dots \rightarrow k' \rightarrow k'' \rightarrow \dots \rightarrow h'' \rightarrow \dots \rightarrow \ell'' \rightarrow \ell \rightarrow \dots \rightarrow g.$$




Figure 3: A path from an element h of $\mathcal{C}_n(g)$ to g which inductively proves the inclusion $\mathcal{C}_n(g) \subset B(g, \frac{1}{2}(5^n - 1))$.

Since they all link an element of a cluster of level $n-1$ to the center of this cluster, the induction hypothesis implies that we can choose the five subpaths $h \rightarrow \dots \rightarrow h'$, $h' \rightarrow \dots \rightarrow k'$, $k'' \rightarrow \dots \rightarrow h''$, $h'' \rightarrow \dots \rightarrow \ell''$ and $\ell \rightarrow \dots \rightarrow g$ of length at most $\frac{1}{2}(5^{n-1} - 1)$. Thus the total length of the path is at most $5 \cdot \frac{1}{2}(5^{n-1} - 1) + 2 \leq \frac{1}{2}(5^n - 1)$. Therefore $\mathcal{C}_n(g) \subset B(g, \frac{1}{2}(5^n - 1))$.

□

Let $x \in \{0, 1\}^G$ be a configuration such that for every $n \in \mathbb{N}$ and $g \in A_n$

$$|\alpha|\mathcal{C}_n(g)|| \leq |\{h \in \mathcal{C}_n(g) \mid x_h = 1\}| \leq \lfloor \alpha|\mathcal{C}_n(g)|| \rfloor + 1. \quad (1)$$

Claim 3.2. *There exists a configuration x that satisfies condition (1).*

Proof. Consider the covering forest given by some sequence $(A_n, p_n, \Gamma_n)_{n \in \mathbb{N}}$ as specified above. For every component C of this forest, take ϕ_C a convex enumeration of its leaves: if g and g' are two leaves of C with the same parent of height n for some $n \in \mathbb{N}$ – i.e. $p_n(g) = p_n(g')$ – then the preimage h of every integer between $\phi_C(g)$ and $\phi_C(g')$ satisfies that $p_n(h) = p_n(g)$. Such an enumeration always exists.

Let $(w_k)_{k \in \mathbb{N}}$ be a Sturmian word of slope α . We can build a configuration x by putting the Sturmian sequence $(w_k)_{k \in \mathbb{N}}$ along the convex enumeration chosen for every component of the forest. Since Sturmian words are balanced, we deduce that the configuration x satisfies condition (1). \square

Claim 3.3. *If a configuration $x \in \{0, 1\}^G$ satisfies condition (1), then x belongs to X_α .*

Proof. Let x be such a configuration and take some $n \in \mathbb{N}$. Let F be a set such that $2n|\partial_{K_n} F| < |F|$ – remember that K_n is $B(1_G, 5^n)$ – and consider the pattern $P := x|_F$. Let $V := \text{Int}(F, B(1_G, \frac{1}{2}(5^n - 1))) \cap A_n$ and $R = F \setminus \bigcup_{v \in V} \mathcal{C}_n(v)$. As $\bigcup_{v \in V} \mathcal{C}_n(v) \subset \bigcup_{v \in V} vB(1_G, \frac{1}{2}(5^n - 1)) \subset F$ we get that:

$$\frac{1}{|F|} \sum_{v \in V} (\lfloor \alpha|\mathcal{C}_n(v)|| \rfloor) \leq \text{dens}(1, P) \leq \frac{1}{|F|} \sum_{v \in V} (\lfloor \alpha|\mathcal{C}_n(v)|| \rfloor + 1) + \frac{|R|}{|F|}.$$

Before working on those inequalities we remark two facts:

1. $R \subset \partial_{K_n} F$. Therefore $\frac{|R|}{|F|} < \frac{1}{2n}$.
2. $|V| \leq \frac{|F|}{|B(1_G, n)|}$.

Indeed, let $r \in \text{Int}(F, K_n)$. That is, for all $g \in K_n$ then $rg \in F$. As $d(r, p_n(r)) \leq \frac{1}{2}(5^n - 1)$ then $p_n(r) \in \text{Int}(F, B(1_G, 5^n - \frac{1}{2}(5^n - 1))) \subset \text{Int}(F, B(1_G, \frac{1}{2}(5^n - 1)))$ therefore $p_n(r) \in V$. That means that $r \notin R$, therefore $R \subset F \setminus \text{Int}(F, K_n) = \partial_{K_n} F$. The second remark is a consequence of Claim 3.1

From the left side we get:

$$\begin{aligned}
\text{dens}(1, P) &\geq \frac{1}{|F|} \sum_{v \in V} (\lfloor \alpha |\mathcal{C}_n(v)| \rfloor) \\
&\geq \frac{\alpha}{|F|} \sum_{v \in V} |\mathcal{C}_n(v)| - \frac{|V|}{|F|} \\
&\geq \alpha \frac{|\bigcup_{v \in V} \mathcal{C}_n(v)|}{|F|} - \frac{|F|}{|F| |B(1_G, n)|} \\
&\geq \alpha \frac{|F \setminus R|}{|F|} - \frac{1}{|B(1_G, n)|} \\
&\geq \alpha \left(1 - \frac{1}{2n}\right) - \frac{1}{2n} \\
&\geq \alpha - \frac{1}{n}
\end{aligned}$$

While from the right side:

$$\begin{aligned}
\text{dens}(1, P) &\leq \frac{1}{|F|} \sum_{v \in V} (\lfloor \alpha |\mathcal{C}_n(v)| \rfloor + 1) + \frac{|R|}{|F|} \\
&\leq \frac{\alpha}{|F|} \sum_{v \in V} (|\mathcal{C}_n(v)|) + \frac{|V|}{|F|} + \frac{1}{2n} \\
&\leq \alpha \frac{|\bigcup_{v \in V} \mathcal{C}_n(v)|}{|F|} + \frac{|F|}{|F| |B(1_G, n)|} + \frac{1}{2n} \\
&\leq \alpha + \frac{1}{n}
\end{aligned}$$

□

Putting together Claims 3.2, 3.1 and 3.3, we obtain that $X_\alpha \neq \emptyset$ which completes the proof of Theorem 3.1. □

Remark. *In the case where α is a computable number, the subshift X_α given in the previous proof is G -effectively closed.*

By noting that in the case of a group of subexponential growth, the sequence of balls always forms a Følner sequence [8, Theorem 6.11.2], we obtain the following result.

Corollary 3.3. *Let G be a group of subexponential growth, for every set of generators S and $\alpha \in [0, 1]$ there exists a non-empty G -subshift with uniform density.*

Definition 3.4. Let G be a group and S a finite set of generators. The rate of convergence of a subshift X with uniform density α is the function

$$\theta_X(n) := \inf\{k \in \mathbb{N} \mid \sup_{g \in G, x \in X} |\text{dens}(1, B(g, k), x) - \alpha| \leq \frac{1}{n}\}.$$

As the construction given in Theorem 3.1 is explicit, we can give bounds for the rate of convergence of X_α . Indeed, let γ denote the growth of a group G , that is, $\gamma(k) = |B(1_G, k)|$ for a fixed set of generators S . Let $x \in X$, $g \in G$ and $B_k := B(g, k)$. By definition of X_α , if $2n|\partial_{B(1_G, 5^n)} B_k| < |B_k|$ then $|\text{dens}(1, B_k, x) - \alpha| < \frac{1}{n}$ for each $x \in X$ and $g \in G$.

As $\partial_{B(1_G, 5^n)} B_k = B_k \setminus B_{k-2 \cdot 5^n}$ if $k \geq 2 \cdot 5^n$, we obtain that:

$$\theta_{X_\alpha}(n) = \inf\{k \geq 2 \cdot 5^n \mid 2n(\gamma(k) - \gamma(k - 2 \cdot 5^n)) < \gamma(k)\}.$$

Therefore a lower bound is always $\theta_{X_\alpha}(n) \geq 2 \cdot 5^n$. The upper bound depends on the growth rate of the group. For instance, if G has polynomial growth then $\theta_{X_\alpha}(n) = O(n \cdot 5^n)$. Indeed, if $\gamma(k) = k^d$ for some $d \geq 1$ we can write:

$$\begin{aligned} 2n(k^d - (k - 2 \cdot 5^n)^d) < k^d &\iff 1 - \left(1 - \frac{2 \cdot 5^n}{k}\right)^d < \frac{1}{2n} \\ &\iff \left(1 - \frac{2 \cdot 5^n}{k}\right)^d > 1 - \frac{1}{2n} \end{aligned}$$

By Bernoulli's inequality, $\left(1 - \frac{2 \cdot 5^n}{k}\right)^d \geq 1 - \frac{2d \cdot 5^n}{k}$. Hence it suffices to choose $k > 4nd \cdot 5^n$. This shows that $\theta_{X_\alpha}(n) = O(n \cdot 5^n)$. In the case of a group of subexponential growth, an upper bound can be computed with an analogous reasoning given the exact rate of growth 2^{k^β} for some $\beta \in (0, 1)$.

Sturmian sequences are classical examples of aperiodic sequences [18]. As X_α shares this uniform density property which makes it similar to Sturmian sequences, it is a natural question to ask if it satisfies a form of aperiodicity.

Proposition 3.4. Let $\alpha \in [0, 1] \setminus \mathbb{Q}$. Then X_α is weakly aperiodic.

Proof. Suppose there exist a configuration $x \in X_\alpha$ and an integer $n \in \mathbb{N}$ such that $|\text{Orb}_\sigma(x)| = n$. Let $D := \{g_i\}_{1 \leq i \leq n} \subset G$ such that each $\sigma_{g_i}(x)$ represents a different element of $\text{Orb}_\sigma(x)$, with $g_1 = 1_G$. Consider also $\alpha' = \text{dens}(1, x|_D) \in \mathbb{Q}$ and $N = \max_{1 \leq i \leq n} d(1_G, g_i)$.

Let $m \in \mathbb{N}$ such that $\frac{2}{m} < |\alpha - \alpha'|$ and $5^m > \frac{N}{2}$. Recall that $K_m := B(1_G, 5^m)$ and consider a finite subset $F \subset G$ such that $2m|\partial_{K_m} F| < |F|$ – by amenability of G such a subset always exists. As $x \in X_\alpha$ we get that $|\text{dens}(1, x|_F) - \alpha| < \frac{1}{m}$. Let $V = \text{Int}(F, B(1_G, N)) \cap \text{stab}_\sigma(x)$ and $R = F \setminus VD$. Note that by definition of N , $VD = \bigcup_{v \in V} vD \subset F$ and that as each $v \in \text{stab}_\sigma(x)$ then $\text{dens}(1, x|_{VD}) = \text{dens}(1, x|_D) = \alpha'$. We obtain:

$$\text{dens}(1, x|_D) \frac{|VD|}{|F|} \leq \text{dens}(1, x|_F) \leq \text{dens}(1, x|_D) \frac{|VD|}{|F|} + \frac{|R|}{|F|}$$

Let $g \in \text{Int}(F, K_m)$. Since the configuration x is supposed to have finite orbit $\{x, \sigma_{g_2}(x), \dots, \sigma_{g_n}(x)\}$, there exists $l \in \{1, \dots, n\}$ such that $\sigma_{g^{-1}}(x) = \sigma_{g_l}(x)$. Therefore $g_l^{-1}g^{-1} \in \text{stab}_\sigma(x)$ which is a subgroup, thus $gg_l \in \text{stab}_\sigma(x)$. As $d(g, gg_l) \leq N$ and $gg_l \in V$ then we conclude that $\text{Int}(F, K_m) \subset VD$ (because we have chosen $g_1 = 1_G$) and therefore $R \subset \partial_{K_m}F$.

Similarly to the previous proof, we bound each side using this relation, obtaining:

$$\begin{aligned} \alpha' \frac{|VD|}{|F|} &\leq \text{dens}(1, x|_F) \leq \alpha' \frac{|VD|}{|F|} + \frac{|R|}{|F|} \\ \alpha' \frac{|F \setminus R|}{|F|} &\leq \text{dens}(1, x|_F) \leq \alpha' + \frac{|\partial_{K_m}F|}{|F|} \\ \alpha' \left(1 - \frac{1}{2m}\right) &\leq \text{dens}(1, x|_F) \leq \alpha' + \frac{1}{2m} \\ \alpha' - \frac{1}{2m} &\leq \text{dens}(1, x|_F) \leq \alpha' + \frac{1}{2m} \end{aligned}$$

Therefore $|\text{dens}(1, x|_F) - \alpha'| < \frac{1}{m}$ and $|\text{dens}(1, x|_F) - \alpha| < \frac{1}{m}$ which implies that $|\alpha - \alpha'| < \frac{2}{m}$ contradicting the definition of m . \square

In the case of \mathbb{Z}^2 , the subshift X_α defined in the proof of Theorem 3.1 is not strongly aperiodic, since it contains the following configurations with non-trivial stabilizer: take a bi-infinite Sturmian word and repeat it vertically so that a configuration $x \in \{0, 1\}^{\mathbb{Z}^2}$ is defined. Then $x \in X_\alpha$ since no forbidden pattern defining X_α appears in x . Thus Proposition 3.4 is in some sense the best we can do for this particular construction.

The statement of Theorem 3.1 itself requires amenability for the group G in order to be meaningful, since we want the density to converge to α for every Følner sequence. Therefore, it doesn't say anything about non-amenable groups. For free groups, we can build configurations (and therefore, subshifts) with uniform density by constructing a regular covering tree and putting a Sturmian sequence on every level of this tree. Nevertheless, we still don't know if this kind of construction is always possible. To our knowledge the following question remains open:

Question. *Let G be an infinite group generated by a finite set S and $\alpha \in [0, 1]$. Does there exist a subshift $Y_\alpha \subset \{0, 1\}^G$ with uniform density α ?*

Acknowledgements: This work was partially supported by the ANR project CoCoGro (ANR-16-CE40-0005).

References

- [1] Nathalie Aubrun, Sebastián Barbieri, and Mathieu Sablik. A notion of effectiveness for subshifts on finitely generated groups. *Theoretical Computer Science*, 661:35–55, 2017.

- [2] Noga Alon, Jaroslaw Grytczuk, Mariusz Haluszczak, and Oliver Riordan. Nonrepetitive colorings of graphs. *Random Structures & Algorithms*, 21(3-4):336–346, 2002.
- [3] Noga Alon and Joel H. Spencer. *The probabilistic method*. Wiley, 2008.
- [4] Alexis Ballier and Maya Stein. The domino problem on groups of polynomial growth. *arXiv preprint arXiv:1311.4222*, 2013.
- [5] Valérie Berthé and Laurent Vuillon. Tilings and rotations on the torus: a two-dimensional generalization of sturmian sequences. *Discrete Mathematics*, 223(1-3):27–53, 2000.
- [6] David B. Cohen. The large scale geometry of strongly aperiodic subshifts of finite type. *Advances in Mathematics*, 308:599–626, 2017.
- [7] David Carroll and Andrew Penland. Periodic points on shifts of finite type and commensurability invariants of groups. *New York Journal of Mathematics*, 21:811–822, 2015.
- [8] Tullio Ceccherini-Silberstein and Michel Coornaert. *Cellular Automata and Groups*. Springer, 2009.
- [9] Thomas Fernique. Multidimensional Sturmian Sequences and Generalized Substitutions. *International Journal of Foundations of Computer Science*, 17:575–600, 2006.
- [10] Su Gao, Steve Jackson, and Brandon Seward. A coloring property for countable groups. *Mathematical Proceedings of the Cambridge Philosophical Society*, 147:579–592, 11 2009.
- [11] Eli Glasner and Vladimir V. Uspenskij. Effective minimal subflows of bernoulli flows. *Proc. Am. Math. Soc.*, 137(9):3147–3154, 2009.
- [12] Gustav Hedlund and Marston Morse. Symbolic dynamics. *American Journal of Mathematics*, 60(4):815–866, 1938.
- [13] Mike Hochman. On the dynamics and recursive properties of multidimensional symbolic systems. *Inventiones Mathematicae*, 176(1):131–167, 2009.
- [14] Emmanuel Jeandel. Some notes about subshifts on groups. *arXiv preprint arXiv:1501.06831*, 2015.
- [15] Douglas A. Lind and Brian Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, 1995.
- [16] M. Lothaire. *Algebraic combinatorics on words*. Cambridge: Cambridge University Press, 2002.
- [17] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Dover Books on Mathematics Series. Dover Publications, 2004.

- [18] N. Pytheas Fogg. *Substitutions in Dynamics, Arithmetics and Combinatorics*, volume 1794 of *Lecture Notes in Mathematics*. Springer, 2002.
- [19] Steven T. Piantadosi. Symbolic dynamics on free groups. Master's thesis, University of North Carolina, Chapel Hill, 2006.
- [20] Steven T. Piantadosi. Symbolic dynamics on free groups. *Discrete and Continuous Dynamical Systems*, 20(3):725–738, 2008.
- [21] Ronnie Pavlov and Michael Schraudner. Classification of sofic projective subdynamics of multidimensional shifts of finite type. *Transactions of the American Mathematical Society*, 367(5), 2014.
- [22] Mark Pollicott and Michiko Yuri. *Dynamical Systems and Ergodic Theory*. London Mathematical Society Student Texts. Cambridge University Press, 1998.