



HAL
open science

Graph Subshifts

Pablo Arrighi, Amélia Durbec, Pierre Guillon

► **To cite this version:**

Pablo Arrighi, Amélia Durbec, Pierre Guillon. Graph Subshifts. 27th International Workshop on Cellular Automata and Discrete Complex Systems, Jul 2021, Marseille, France. 10.48550/arXiv.2302.07249 . hal-03238087v3

HAL Id: hal-03238087

<https://hal.science/hal-03238087v3>

Submitted on 17 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Graph subshifts

Pablo Arrighi, Amélia Durbec, and Pierre Guillon

¹ Université Paris-Saclay, Inria, CNRS, LMF, 91190 Gif-sur-Yvette, France and IXXI, Lyon, France

² Université Paris-Saclay, CNRS, LISN, 91190 Gif-sur-Yvette, France

³ Aix-Marseille Université, CNRS, I2M, Marseille, France

Abstract. We propose a definition of graph subshifts of finite type that can be seen as extending both the notions of subshifts of finite type from classical symbolic dynamics and finitely presented groups from combinatorial group theory. These are sets of graphs that are defined by forbidding finitely many local patterns. In this paper, we focus on the question whether such local conditions can enforce a specific support graph, and thus relate the model to classical symbolic dynamics. We prove that the subshifts that contain only infinite graphs are either aperiodic, or feature no residual finiteness of their period group, yielding non-trivial examples as well as two natural undecidability theorems.

1 Introduction

Subshifts of finite type are well studied objects in symbolic dynamics [17] and ergodic theory. Given an alphabet Σ , a subshift of finite type (SFT) on \mathbb{Z}^d is a set of configurations $\mathcal{Y} \subseteq \Sigma^{\mathbb{Z}^d}$ that do not contain a given finite set of forbidden patterns $\mathcal{F} \subseteq \Sigma^{\{-n, n\}^d}$. In spite of their relatively benign and local definition, SFT have proven to have complex global behaviours. One example of such results is the existence of aperiodic SFTs in dimension 2, but not in dimension 1. Even some natural problems such as determining whether an SFT contains a periodic configuration at all, are proven to be undecidable [5]. This local-to-global complexity is in fact shared by multiple dynamical systems and most notably with Cellular Automata (CA) [16]. However, the two models are deeply connected as CA of dimension d can be seen as subshift of dimension $d + 1$. This connection has been used to prove multiple theorems on different subclasses of CA [14]. This result is also one of the original motivations of this paper: Graph subshifts ought to be a mean to study a generalization of CA called Causal Graph Dynamics (CGD) [1]. CGD extend CA in two complementary ways: first because they are defined on arbitrary graphs of bounded degree m , and second because they allow for the graph itself to evolve, according to a local, shift-invariant rule. Graph subshifts aim to encompass CGD spacetime diagrams as a subclass.

Subshifts have already been generalized to configurations over Cayley graphs of finitely presented groups. However this is not enough to simulate CGD as all configurations share the same support graph. In this paper, we provide a formalism which relaxes this constraint, allowing the support of the graphs itself to also be prescribed by set of forbidden patterns. We introduce and formalize the notion of graph subshifts, i.e. sets of graphs that are defined by forbidding finitely many local patterns.

The natural question to ask, then, is whether we may enforce, by means of such local constraints only, that the set of graphs be of a particular shape. In particular, notice that a graph subshift may contain a finite or an infinite number of graphs (up to shift), and that the graphs themselves may be of finite or of infinite size. In this first paper we focus on the question whether there exist subshifts whose graphs are all of infinite size. The question is non-trivial; for instance we prove that the problem whether a given set of forbidden patterns induces a graph subshift without finite configuration is undecidable. The problem whether it uniquely fixes the support graph is also shown undecidable. Still, we prove that the graph subshifts that contain only infinite graphs are either aperiodic, or feature no residual finiteness of their period group, yielding us with non-trivial examples and establishing connections with different areas of mathematics : periodicity from symbolic dynamics; residually finite groups from combinatorial group theory [4][18]; and graph covers as used in graph theory and distributed computing.

2 Graphs

Graphs, pointed graphs modulo and their operations are intuitive enough as summarized here, but they have been introduced several times before (see for instance [2]).

Let π be a finite set, $\Pi = \pi^2$, and V some universe of names.

Graphs. Our graphs are the usual, connected, undirected, possibly infinite, bounded-degree labeled graphs, but with a few added twists:

- Vertices are connected through their ports. An edge is an unordered pair $\{x : a, y : b\}$ of $V \times \pi$, or a singleton $\{x : a\}$ standing for a self-loop, where x, y are vertices and $a, b \in \pi$ are ports. Each port is used at most once per node: if $e, e' \in E(G)$ intersect, they must be equal. As a consequence the degree of the graph is bounded by $|\pi|$.
- Each vertex x is given a label $\sigma(x)$, taken within a finite set Σ , also referred to as an internal state.

The set of all graphs having ports π , vertex labels Σ is denoted $\mathcal{G}_{\Sigma, \pi}$, or simply \mathcal{G} .




Fig. 1. *The different types of graphs.* (a) A graph G . (b) A pointed graph $(G, 1)$. (c) A pointed graph modulo X .

We single out a vertex as the origin:

Pointed graph non-modulo. A *pointed graph* is a pair (G, p) with $p \in V(G)$. The set of pointed graphs with states in Σ and ports π is written $\mathcal{P}_{\Sigma, \pi}$, or simply \mathcal{P} .

Here is when graph differ only modulo names of vertices:

Graph Isomorphism. An *isomorphism* R is a function from \mathcal{G} to \mathcal{G} which is specified by a bijection $R(\cdot)$ from V to V . The image of a graph G under the isomorphism R is a graph RG having :

- A set of vertices $R(V(G))$.
- A set of edges $\{\{R(u) : a, R(v) : b\} \mid \{u : a, v : b\} \in E(G)\}$.
- Internal states given by $\sigma \circ R^{-1}$.

Similarly, the image of a pointed graph $P = (G, p)$ is the pointed graph $RP = (RG, R(p))$. When P and Q are isomorphic we write $P \approx Q$, defining an equivalence relation on the set of pointed graphs.

Notice that pointed graph isomorphism rename the pointer in the same way as it renames the vertex upon which it points; which effectively means that the relative position of vertices w.r.t to the pointer is unchanged.

Pointed graphs modulo. The goal of this twists is to work with the set of graphs that is a compact space. Indeed, consider a graph G , a vertex $x \in V(G)$, and $D = G_x^r$ the disk of radius r around x , *i.e.* the subgraph of G induced by the r -neighbors of x in G . Now forgetting about G , this D is one amongst the many possible disks of radius r that can appear in \mathcal{G} . In order for the set of graphs to form a compact space, one requirement is that the set of disks of radius r be finite. Morally this is the case, since the degree is bounded by $|\pi|$. Vertex names x, y, z jeopardize this however; we must get rid of them. Thus,

- Graphs are endowed with a privileged pointed vertex p playing the role of an origin.
- These pointed graphs are considered modulo isomorphism, so that only the relative positions of vertices matter.

The set of all pointed graphs modulo is denoted $\mathcal{X}_{\Sigma, \pi}$, or simply \mathcal{X} . Given a graph non-modulo G , and a vertex $p \in V(G)$, we write $X((G, p))$ for the corresponding pointed graph modulo. Additionally the set of disks of radius r with states Σ and ports π is written \mathcal{X}^r .

Paths and vertices. Over graphs modulo isomorphism without pointer, vertices no longer have a unique identifier, which makes designating a vertex a daunting task. Fortunately, our graphs are pointed, so that any vertex \hat{u} can always be designated by some sequence of ports u that leads to it, starting from the origin. Given a pointed graph modulo X , its set of paths, starting from the origin, forms a subset $L(X)$ of Π^* with $\Pi = \pi^2$. Whenever two paths u and v lead to the same vertex, we write $u \equiv_X v$. Thus vertex \hat{u} is really the \equiv_X -equivalence class of u , hence the hat notation. The set of these equivalence classes, a.k.a vertices, is denoted $V(X)$. Notice that, starting from origin, following path $v = (ab) \cdots (cd)$ and then path $\bar{v} = (dc) \cdots (ba)$, leads back to the vertex at the origin, namely $\hat{\varepsilon}$.

Remark 1 Notice that two paths u, u' designate the same vertex $\hat{u} = \hat{u}'$ if and only if for every $v \in L(X)$, $\widehat{uv} = \widehat{u'v}$.

Remark 2 A reduced path is a path that does not go through the same edge twice consecutively in opposite directions, i.e. $u.abba.v$ is not simple as it can be reduced to $u.v$. The language up to this reduction equivalence is a groupoid for concatenation (and inversion), i.e. alike a group but its law is partially defined.

The set of (centered) cycles, i.e. reduced paths from the pointer back to it, is a subgroup of it (concatenation and inversion are always defined among them).

The set of vertices is the quotient of the set of reduced paths by the following equivalence relation: $u \equiv_X v$ if and only if there exists a cycle c such that $u = c.v$. Indeed, whenever they designate the same vertex $\hat{u} = \hat{v}$, we have $u = (u\bar{v})v$, where $u\bar{v}$ is a cycle.

Nevertheless, concatenation, does not induce a semigroup at the level of vertices, since we may have $\hat{u} = \hat{v}$ and yet $\widehat{w\hat{u}} \neq \widehat{w\hat{v}}$.

2.1 Operations over pointed Graphs modulo

Shift. Let $X \in \mathcal{X}$ be a pointed graph modulo and G its associated graph. Consider $u \in L(X)$. Then X_u is $X((G(X), \hat{u}))$. The pointed graph modulo X_u is referred to as X shifted by u .

Back to graphs non-modulo. Sometimes we still need to manipulate usual graphs, where vertices do have names. For this purpose we use a canonical naming function $X \in \mathcal{X} \mapsto G(X) \in \mathcal{G}$, which names each vertex of X by the set of paths that lead to it, starting from ε . This $G(X)$ is referred to as ‘the associated graph’. One has $X(G(X), \hat{\varepsilon}) = X$.

2.2 Topology on pointed graphs modulo

Having a well-defined notion of disks allows us to define a topology upon \mathcal{X} , which is the natural generalization of the well-studied Cantor metric upon CA configurations [12].

Definition 1 (Gromov-Hausdorff-Cantor metrics) Consider the function

$$\begin{aligned} d : \mathcal{X} \times \mathcal{X} &\longrightarrow \mathbb{R}^+ \\ (X, Y) &\mapsto d(X, Y) = 0 \quad \text{if } X = Y \\ (X, Y) &\mapsto d(X, Y) = 1/2^r \quad \text{otherwise} \end{aligned}$$

where r is the minimal radius such that $X^r \neq Y^r$.

The function $d(\cdot, \cdot)$ is such that for $\epsilon > 0$ we have (with $r = \lfloor -\log_2(\epsilon) \rfloor$):

$$d(X, Y) < \epsilon \iff X^r = Y^r.$$

Remark 3 This defines an ultrametric distance, and the topology can be seen as a product of discrete topologies over the (unions of finitely many) discs of each radius.

3 Subshifts

Given a pointed graph modulo X , a word in Π^* designates a vertex, and a language $\mathcal{L} \subseteq \Pi^*$ designates a finite set of vertices. If, moreover, \mathcal{L} is stable under taking the prefix, then the designated vertices induce a connected subgraph of X which we refer to as ‘cut’. For instance, a cut could represent a spacelike surface within a spacetime diagram. Given a set \mathcal{Z} of graphs, one may be interested in selecting those whose cut according to \mathcal{L} equals some graph Z , which we refer to as ‘cylinder’. This could be, for instance, because \mathcal{Z} represents the set of every possible spacetime diagram, and Z a particular initial condition.

Definition 2 (Cuts and cylinders) Consider $\mathcal{L} \subseteq \Pi^*$ a prefix-stable language, i.e. such that $wv \in \mathcal{L}$ implies $u \in \mathcal{L}$. The \mathcal{L} -cut of a graph $X \in \mathcal{X}$ is the subgraph induced by the vertices $\{\hat{u} \mid u \in \mathcal{L} \cap L(X)\}$. It is denoted $X|_{\mathcal{L}}$. Consider $\mathcal{Z} \subseteq \mathcal{X}$, we write $\mathcal{Z}|_{\mathcal{L}}$ for $\{X|_{\mathcal{L}} \mid X \in \mathcal{Z}\}$. Consider $Z \in \mathcal{Z}|_{\mathcal{L}}$, the cylinder of Z within \mathcal{Z} is $\{X \in \mathcal{Z} \mid X|_{\mathcal{L}} = Z\}$. It is denoted $[Z]_{\mathcal{Z}}$.

Definition 3 (Subshift) Let \mathcal{F} be a set of tuples (F, L) , where each F is a finite graph and each L a finite prefix-stable language. The subshift forbidding \mathcal{F} is

$$\mathcal{Z} = \{X \in \mathcal{G}_{\Sigma, \pi} \mid \forall v \in L(X), \forall (F, L) \in \mathcal{F}, (X_v)|_L \neq F\}.$$

It is of finite type if \mathcal{F} can be chosen finite.

Remark 4 The following are equivalent.

1. \mathcal{Z} is a subshift of finite type.
2. \mathcal{Z} is the subshift forbidding some $\mathcal{F}' \times \{M\}$ where \mathcal{F}' is a finite set of finite graphs and M a single finite prefix-stable language.
3. \mathcal{Z} is the set of graphs allowing some finite set \mathcal{A} of tuples (A, L) , where A is a finite graph and L a finite prefix-stable language—in the sense that:

$$\mathcal{Z} = \{X \in \mathcal{G}_{\Sigma, \pi} \mid \forall v \in L(X), \exists (A, L) \in \mathcal{A}, (X_v)|_L = A\}.$$

4. \mathcal{Z} is the set of graphs allowing $\mathcal{A}' \times \{M\}$ where \mathcal{A}' is a finite set of finite graphs and M a single finite prefix-stable language.

Proof.

- 1 \implies 2 For every \mathcal{F} , define $M = \bigcup_{(F, L) \in \mathcal{F}} L$ and $\mathcal{F}' = (\mathcal{X} \setminus \mathcal{Z})|_M$. $\mathcal{F}' \times \{M\}$ defines the same subshift as \mathcal{F} .
- 3 \iff 4 Similarly, for every \mathcal{A} , define $M = \bigcup_{(F, L) \in \mathcal{F}} L$ and $\mathcal{A}' = \mathcal{Z}|_M$. $\mathcal{A}' \times \{M\}$ defines the same subshift as \mathcal{A} .
- 2 \iff 4 The subshift \mathcal{Z} obtained by forbidding $\mathcal{F}' \times M$ is the same as that obtained from allowing $\mathcal{A}' \times M = (\mathcal{X}|_M \setminus \mathcal{F}') \times M$. Note that $\mathcal{X}|_M$ is finite.
- 2 \implies 1; 4 \implies 3 Obvious.

Definition 4 The prefix-stable language M is the same in Conditions 2 and 4 of Remark 4. It is called a defining window for \mathcal{Z} . We can assume that it is the minimal possible one.

Remark 5 The duality between forbidding and allowing no longer holds in general when the subshifts are not of finite type. For example we can think about the set of graphs authorizing exactly those finite graphs which have exactly one red vertex. This set however is not closed topologically (cf. [2] as summarized by the Appendix) and hence it is not a subshift. Its closure is the so-called one-dimensional sunny-side-up subshift, as defined by forbidding those finite graphs having two red vertices.

Remark 6 In the pairs (F, L) , L cannot always be assumed to be the language labelling the paths of F . For example, requiring every vertex to have an edge with port $a \in \Sigma$ is a very natural finite-type condition, but this SFT is not of the form $\{X \in \mathcal{G}_{\Sigma, \pi} \mid \forall v \in L(X), \forall F \in \mathcal{F}, (X_v)_{|\mathcal{L}_F} \neq F\}$ as the \mathcal{L}_F do not contain a and hence cannot tell that it is absent.

The same remark holds for the dual definitions, i.e. through allowed patterns: think of the SFT consisting of those graphs having no port $a \in \Sigma$.

The equivalent definitions all capture the concept of local constraint, and in the following sections we often define our SFTs by describing some conditions that graphs must locally satisfy. Among these conditions, *nearest-neighbor constraints* are those defined with a single edge between two vertices, that is $L \in \Pi$.

4 Examples

Example 1. A *directed 2-regular SFT* is an SFT with ports $\pi = \{a, a', b, b'\}$ obtained by imposing constraints: 1/ enforcing that edges be of the form pp' ; 2/ enforcing that all ports be occupied. It contains, amongst many others, the Cayley graph of the free group over two generators (a tree), as well as those of all of its quotients.




Fig. 2. *Locally grid-like SFT corresponding to the Cayley graph of $\langle a, b \mid aba^{-1}b^{-1} \rangle$.* The SFT is defined by forbidding all but these patterns, on their corresponding languages.

Example 2 (Locally grid-like SFT). We can define a graph SFT with ports $\pi = \{a, a', b, b'\}$ corresponding to the monochromatic fullshift on $\langle a, b \mid aba^{-1}b^{-1} = \varepsilon \rangle$, by imposing three families of constraints:

1. enforcing that edges be of the form pp' ;
2. enforcing that all ports be occupied;
3. enforcing that the subgraph induced by $\{\varepsilon, aa', bb', aa'bb', bb'aa'\}$ obeys the relation i.e. form a square as in Fig. 2).

This graph SFT obviously contains the Cayley graph of $\langle a, b \mid aba^{-1}b^{-1} = \varepsilon \rangle$, but also the Cayley graphs of most quotient subgroups (excluding really small groups such as the trivial group).

A similar construction can be performed for any finitely presented group, and via a more precise third constraint, get a cleaner statement (with all possible quotients).

Example 3. Consider a finitely presented group $\Gamma = \langle \mathcal{J} \mid \mathcal{R} \rangle$, that is, a group generated by the abstract finite set \mathcal{J} , and \mathcal{R} represents a minimal set of relations, that is words over $J \sqcup J^{-1}$ that represent the identity element (and from which all such words can be derived). For example $\mathbb{Z}^2 = \langle a, b \mid aba^{-1}b^{-1} \rangle$ (it will always be implicitly given by this canonical presentation) and $\mathbb{F}_2 = \langle a, b \mid \rangle$.

For every alphabet Σ , there is a canonical way to associate a *Cayley SFT* $\mathcal{Z}_{\Sigma}^{\Gamma}$ to this group. Define $\pi = J \sqcup J^{-1}$ and the following constraints:




Fig. 3. Configuration of the hard-square model that is present in \mathbb{Z}^2 . The colours represent the 0/1 internal states of the vertices.

1. each edge involves two inverse ports;
2. every vertex must have all possible ports;
3. for every word $u \in \mathcal{R}$ and L its language of prefixes, one forbids the pair (Y, L) for every graph Y supported by L in which path u is not a cycle back to the origin.

This SFT obviously contains the Cayley graph of Γ , but also all of the Cayley graphs of its quotient subgroups, as will be seen in the next section. For example, $\mathcal{Z}_{\Sigma}^{\mathbb{Z}^2}$ contains the infinite grid as well as tori and cylinders, and $\mathcal{Z}_{\Sigma}^{\mathbb{Z}^2}$ actually contains all Cayley graphs of all groups with two generators.




Fig. 4. Configuration of the hard-square model whose support graph is not \mathbb{Z}^2 . In order to preserve clarity, only half of the ports are written and some edges have been dashed.

Let Γ be a finitely presented group. A SFT is defined by Γ -NN constraints if they include precisely the forbidden patterns defining \mathcal{Z}^{Γ} and some additional nearest-neighbor forbidden patterns. Note that if $\Gamma = \mathbb{F}_k$, then \mathbb{F}_k -NN constraints are simply constraints on the edges, provided the correct regularity of the graph.

Example 4 (Hard-square model). The hard-square model SFT is defined by \mathbb{Z}^2 -NN constraints: $\mathcal{Z} \subseteq \mathcal{X}_{\{0,1\},a,a',b,b'}$ such that for every configuration X and edge $\{u : a, v : b\} \in E(X)$, one has that $\sigma_x(u) = 0$

or $\sigma_x(v) = 0$. This can be expressed as an SFT by forbidding the black/white coloured versions of Fig. 2 that have two adjacent black vertices. This is a direct generalization of a well-known SFT on \mathbb{Z}^2 and other subgroup that go by the same name. As a graph SFT, it not only contains all the configuration that are already present on \mathbb{Z}^2 such as in Fig. 3, but also new infinite configurations such as the cylinder of Fig. 4, and even finite configuration such as tori only containing 0.

Besides finitely presented groups, the notion of quotient of a group can be extended to periodic graphs, which is the subject of the following section.

Example 5. Given any subshift \mathcal{Z} , one can define the subshift of its infinite graphs, by additionally forbidding $(F, L\pi^2)$ for every prefix-stable language $L \in \Pi^*$ and every graph F such that $F|_L = F$. Indeed this condition exactly means that L was already describing all possible vertices in F , and longer paths do not add any. Nevertheless, this subshift is very rarely of finite type, as we will see in the following section.

On the contrary, all SFTs presented so far contain finite graphs. It is not obvious a priori whether some SFTs contain only finite graphs, or even whether they are *homogeneous*, i.e. all the graphs share a common support. We will discuss these properties in the next sections.

5 Periodicity and finiteness

5.1 Quotient and Periods

Definition 5 (Period) *Any word $u \in L(X)$ such that $X = X_u$ is called a path-period. The corresponding vertex \hat{u} is called a period. The stabilizer of X is its set of periods; it is denoted by $\mathcal{S}(X)$. X is called weakly periodic if $\mathcal{S}(X) \supsetneq \{\hat{\varepsilon}\}$, strongly aperiodic otherwise.*

Notice that if X is weakly periodic, then either $\mathcal{S}(X)$ is infinite or every period \hat{u} is *torsion*, i.e. $\widehat{u^n} = \hat{\varepsilon}$ for some $n \geq 2$.

Remark 7 *The set of reduced path-periods is a subgroup of the groupoid of reduced paths defined in Remark 2. The set of cycles is a normal subgroup of the group of path-periods (since if u is a period and v is a cycle, then $\bar{u}v\bar{u}$ is still a cycle). Like in Remark 2, the stabilizer is the quotient of the group of reduced path-periods by the normal subgroup of cycles (since if u and v are paths to the same vertex, then $u = (u\bar{v})v$, where $u\bar{v}$ is a cycle).*

All the statements in Remark 7 are easy to prove, but we concentrate on proving the next proposition, which will be useful for the rest of the paper.

Proposition 1 *Let X be a graph. For \hat{u}, \hat{v} in $\mathcal{S}(X)$, the concatenation operation $\hat{u}\hat{v} \mapsto \widehat{u\bar{v}}$, and the reverse operation $\widehat{\bar{u}} \mapsto \widehat{\bar{u}}$ are well-defined, and endow $\mathcal{S}(X)$ with a group structure.*

Proof.

- Let $u \equiv_X u'$ and $v \equiv_X v'$. Remark 1 already gives that $uv \equiv_X u'v'$. Now from $X_{u'} = X$, one can deduce that $v \equiv_{X_{u'}} v'$, which means that $u'v \equiv_X u'v'$. We have proven that $uv \equiv_X u'v'$, so that the concatenation operation is well defined over vertices.
- Remark 1 gives that $u'\bar{u}' = \varepsilon = u\bar{u} \equiv_X u'\bar{u}$. This means that $\bar{u}' \equiv_{X_{u'}} \bar{u}$, and since $X = X_{u'}$, $\bar{u}' \equiv_X \bar{u}$. We have proven that the inversion is well defined over vertices.
- *Neutral element:* the origin $\hat{\varepsilon}$ is a period, because $X_\varepsilon = X$.
- *Concatenation:* If $\hat{u}, \hat{v} \in \mathcal{S}(X)$, then $(X_{u\bar{v}}) = (X_u)_{\bar{v}} = X_v = X$. It follows that $\widehat{u\bar{v}}$ is a period.
- *Inversion:* If $\hat{u} \in \mathcal{S}(X)$, then since $X = X_{u\bar{u}} = (X_u)_{\bar{u}} = X_{\bar{u}}$. It follows that $\widehat{\bar{u}}$ is a period. This period provides a right and left inverse for \hat{u} , since $\widehat{u\bar{u}} = \hat{\varepsilon} = \widehat{\bar{u}u} = \widehat{\bar{u}u}$.

Remark 8 If H is a stabilizer subgroup of X , then for every vertex $\hat{u} \in V(X)$, $\widehat{u}H\hat{u}$ is a stabilizer subgroup of X_u . $(X_u)/\widehat{u}H_u$ is well-defined and turns out to be equal to $(X/H)_u$.

Notice that monochromatic Cayley graphs are exactly the graphs whose all vertices are periods. As we have seen in our Cayley SFT examples, periodic graphs can typically be “folded into smaller graphs” that also belong in the graph SFT. The following formalizes the idea of a graph folding into another one.

Definition 6 (Homomorphism) A homomorphism from graph $X \in \mathcal{X}$ to graph $X' \in \mathcal{X}$ is a function φ from $V(X)$ to $V(X')$ such that for all $u \in L(X)$,

$$\varphi(\hat{u}) = \hat{u} \text{ (in particular } L(X) = L(X')) \quad \text{and} \quad \sigma_X(u) = \sigma_{X'}(u).$$

When it is surjective, we say that X covers X' and that φ is a covering.

Remark 9 A homomorphism sends cycles to cycles and preserves exactly the set of path-periods. Hence if \hat{u} is a period of X , $\varphi(\hat{u})$ is a period of $\varphi(X)$.

Consider $X \in \mathcal{X}$ and a stabilizer subgroup H . We can define a quotient graph X/H to be $X((G, p))$ where:

- $V(G) = \{H\hat{u} \mid \hat{u} \in V(X)\}$, i.e. each vertex is labelled by a set $H\hat{u}$.
- $p = H$, i.e. the pointer is the vertex labeled by H .
- $E(G) = \left\{ (H\hat{u} : a, H\widehat{uab} : b) \mid uab \in L(X) \right\}$.
- $\sigma_G(H\hat{u}) = \sigma_X(u)$.

Proposition 2 (Quotient graph) X/H is a well-defined graph, and $\hat{u} \mapsto H\hat{u}$ is a homomorphism.

Proof. First of all, notice that $H\hat{u} = H\hat{u}'$ implies that there exists $\hat{h} \in H$ such that $\hat{u}' = \hat{h}\hat{u}$, so that $X_{u'} = (X_h)_u = X_u$.

Let us prove that $\sigma_{X/H}$ is well-defined. Let $\hat{u}, \hat{u}' \in V(X)$ such that $H\hat{u} = H\hat{u}'$. From the remark above, $X_u = X_{u'}$, and therefore $\sigma_X(u) = \sigma_X(u')$.

Now let us prove that if there is an edge $\{H\hat{u} : a, H\hat{v} : b\}$ in $G(X/H)$, then there is an edge $\{\hat{u} : a, \widehat{h''v} : b\}$ in $G(X)$. Indeed, by definition and the remark above, there is an edge $\{\hat{h}\hat{u} : a, \hat{h}'\hat{v} : b\}$ in $G(X)$ for some $\hat{h}, \hat{h}' \in H$. Since $X_u = X_{hu}$, there is also an edge $\{\hat{u} : a, \widehat{h\hat{h}'v} : b\}$. Thus $H\hat{u}$ inherits from \hat{u} the absence of port conflict. The function φ satisfies trivially the two homomorphism conditions.

We say that a set of vertices H is M -separated in graph X if for every $\hat{u} \in V(X)$, $H \cap (X_u)_{|M} = \{\varepsilon\}$. If $U, V \subset \Sigma^*$, then we note $UV = \{uv \mid u \in U, v \in V\}$, and $\overline{U} = \{\bar{u} \mid u \in U\}$.

Proposition 3 Let X be a graph and $H \leq \mathcal{S}(X)$ a stabilizer subgroup. If H is $M\overline{M}$ -separated in X , then the covering preserves the language of support M in the sense that for all $u \in L(X)$, $(X_u)_{|M} = (X/H_u)_{|M}$ (modulo isomorphism).

In particular, if X belongs in some SFT \mathcal{X} with defining window M , then X/H also belongs in \mathcal{X} .

Proof. By definition, if u is a valid path in X , then it is also a valid path in X/H , and the internal state $\sigma_X(u) = \sigma_{X/H}(u)$ is the same. It remains to prove that if v, w designate distinct vertices in $(X_u)_{|M}$, then they still designate distinct vertices in $(X/H_u)_{|M}$. Assume on the contrary that $v \equiv_{(X/H_u)_{|M}} w$. By the fact that concatenation is well defined (Proposition 2), there exists h such that $\hat{h} \in H$ and $h\bar{v} \equiv_{X_u} w$, i.e. $h \equiv_{X_u} w\bar{v}$. Since $w\bar{v} \in (X_u)_{|M\overline{M}}$ and $\hat{h} \in H$, separation gives that $w\bar{v} \equiv_{X_u} h \equiv_{X_u} \varepsilon$. This means that $w \equiv_{X_u} v$. This proves that the homomorphism induces an isomorphism over $(X_u)_{|M}$. By definition, this homomorphism does not change the labeling alphabet, so that (modulo isomorphism) $(X_u)_{|M} = (X/H_u)_{|M}$. If \mathcal{X} has defining window M , then by definition $X \in \mathcal{X} \iff X/H \in \mathcal{X}$.

5.2 Strong periodicity

Definition 7 (Strong periodicity) A language $L \in \Pi^*$ is said to be dense in X if there exists a finite prefix-stable language $M \in \Sigma^*$ (we say that L is M -dense) such that for all $\hat{u} \in V(X)$, $\hat{u}M \cap L \neq \emptyset$. A graph X is said to be strongly periodic if its stabilizer $\mathcal{S}(X)$ is dense in X . Equivalently, there are only finitely many distinct (modulo isomorphisms) graphs X_u , for $u \in L(X)$.

Note that finite graphs are strongly periodic, but may not always be weakly periodic. Apart from this degenerate case, strongly periodic graphs are weakly periodic.

Notice also that strong periodicity is almost equivalent to being a Cayley graph, up to replacement of each vertex by a fixed finite graph (with possible edges between them). This replacement (factor map) corresponds somehow to the application of a causal graph dynamics [2] to the Cayley graph.

Corollary 1 A graph X is strongly periodic if and only if X/H is finite for some stabilizer subgroup H .

Proof. First, assume that $H = \mathcal{S}(X)$ is M -dense for some $M \in \Sigma^*$: For every $u \in L(X)$, there exist $\hat{h} \in \mathcal{S}(X)$ and $v \in \overline{M}$ such that \hat{u} can be written $\hat{h}\hat{v}$. Now for every $\hat{h}' \in H$, $\hat{h}'\hat{u} = \hat{h}'\hat{h}\hat{v} \in H\hat{v}$, so $H\hat{u} \subseteq H\hat{v}$ and $\hat{h}'\hat{v} = \hat{h}'\hat{h}\hat{h}\hat{v} = \hat{h}'\hat{h}\hat{u} \in H\hat{u}$, so $H\hat{v} \subseteq H\hat{u}$. Thus $H\hat{u} = H\hat{v}$. Hence, the set $\{H\hat{u} \mid \hat{u} \in V(X)\} = \{H\hat{v} \mid v \in M \cap L(X)\}$ of vertices of X/H is finite. We conclude by Proposition 2. Conversely, if X admits a finite quotient X/H , there exists M such that $\{H\hat{u} \mid u \in L(X)\} = \{H\hat{v} \mid v \in M \cap L(X)\}$. For all $\hat{u} \in V(X)$, since $\hat{u} \in H\hat{u}$, we get that $\hat{u} \in H\hat{v}$ for some $v \in M$. Therefore, $\hat{u} = \hat{h}\hat{v}$ with $\hat{h} \in \mathcal{S}(X)$ and $v \in M$, so that H is \overline{M} -dense.

The following result establishes a surprising connection between the SFT that admit finite graphs, and the well-established notion (see for instance [7, Chap. 2]) of residual finiteness from group theory:

Definition 8 (Residually finite group) A group \mathcal{S} is residually finite if and only if for any non-identity element $\hat{u} \in \mathcal{S}$, there exists a normal subgroup $H \trianglelefteq \mathcal{S}$ of finite index such that $\hat{u} \notin H$.

This notion is very robust. In particular, the property extends to finite sets of non-identity elements I , i.e. there exists a normal subgroup $H \trianglelefteq \mathcal{S}$ of finite index such that $H \cap I \subseteq \{\varepsilon\}$, simply because any finite intersection of finite-index subgroups is a finite-index subgroup.

Let us relate residual finiteness to quotient graphs, at the intuitive level first. Say X belongs in some SFT \mathcal{Y} and let $\hat{u}_1, \dots, \hat{u}_n$ be distinct elements of $V(X)$. We have that $H(\hat{u}_1, \dots, \hat{u}_n) \subseteq \mathcal{S}(X)$ is a subgroup of its stabilizer group, and thus that $X/H(\hat{u}_1, \dots, \hat{u}_n)$ is a possible quotient graph; one which maintains that $\hat{u}_1, \dots, \hat{u}_n$ are distinct. The correct choice of $\hat{u}_1, \dots, \hat{u}_n$ will help us make sure that $X/H(\hat{u}_1, \dots, \hat{u}_n)$ still belongs in SFT \mathcal{Y} .

Theorem 1 Let \mathcal{Y} be a SFT. \mathcal{Y} contains a finite graph if and only if \mathcal{Y} contains a strongly periodic graph X such that $\mathcal{S}(X)$ is residually finite.

Proof. The right-to-left implication directly comes from the fact that finite graphs are strongly periodic and finite groups are residually finite.

Now, suppose that \mathcal{Y} has defining window M and admits a strongly periodic graph X such that $\mathcal{S}(X)$ is residually finite. Since X is strongly periodic, $\bigcup_{\hat{u} \in V(X)} V((X_u)_{|M\overline{M}})$ is a finite union of finite subgraphs. Since $\mathcal{S}(X)$ is residually finite, there exists a subgroup $H \leq \mathcal{S}(X)$ of finite index which does not intersect $\bigcup_{\hat{u} \in V(X)} V((X_u)_{|M\overline{M}}) \setminus \{\varepsilon\} \subseteq \Sigma^{M\overline{M}}$. Hence, H is $M\overline{M}$ -separated, so that we can use Proposition 3: for all $\hat{u} \in V(X)$ we have that $(X/H_u)_{|M} = (X_u)_{|M}$. Given that \mathcal{Y} has defining window M , we get that $X/H \in \mathcal{Y}$.

Thus there are two complementary ways in which a nonempty SFT may have only infinite configurations: either by preventing strongly periodic configurations, or by forcing them to have stabilizers that are not residually finite.

In the non-residually-finite case (like the Tarski monster groups, see for instance [20]), one can build SFTs with strongly periodic configurations that do not fold uniformly onto finite graphs.

Remark 10 Γ is not residually finite if and only if $\mathcal{Z}_{\{0\}}^\Gamma$ includes an SFT containing the Cayley graph of Γ but no finite quotient of it. Indeed, one can consider the SFT \mathcal{Z} included in $\mathcal{Z}_{\{0\}}^\Gamma$ defined by the additional constraint that u is never a cycle, where u is a noncyclic path that witnesses that Γ is not residually finite, in the sense that all finite quotients of Γ map u to ε .

5.3 Quotients in Cayley SFTs

We have seen that the notions of quotients and SFTs are linked, because finite-type constraints cannot easily tell a graph from some of its quotients. On the opposite point of view, it is noticeable that some SFTs admit a maximal support with respect to coverings.

Proposition 4 Let Γ be a finitely presented group, X a graph, and \mathcal{Z} some SFT defined by Γ -NN constraints. The following are equivalent:

1. $X \in \mathcal{Z}$.
2. X is covered by some colored Cayley graph Y of Γ which is in \mathcal{Z} .

Proof.

- 1 \implies 2 Suppose that the support of X is in \mathcal{Z} . We note that paths designating vertex \hat{u} in Y all designate the same vertex in X : indeed, if \hat{v} also designates \hat{u} in Y , then it can be obtained from \hat{u} by iteratively applying the reduction rules from the presentation. Then thanks to the third condition in the definition of \mathcal{Z}_Y^Γ , it also designates the same vertex as \hat{u} in X . Hence one can simply define $\phi(\hat{u}) = \hat{u} \in V(X)$ for every $\hat{u} \in V(Y)$, and $\sigma_Y(\hat{u}) = \sigma_X(\hat{u})$. This clearly satisfies the conditions of Definition 6.
- 2 \implies 1 Conversely, consider a graph X covered by some Cayley graph Y of Γ which is in \mathcal{Z} . Consider a cycle constraint of \mathcal{Z} , that is: every path labeled by some u (in the group presentation of Γ) must be a cycle. This constraint is satisfied in X because loops are preserved by homomorphism. Now consider an edge constraint of \mathcal{Z} . These are also preserved by homomorphism, by definition. We have proven that X satisfies all constraints defining \mathcal{Z} .

In graphs of $\mathcal{Z}_{\Sigma}^{\mathbb{Z}^2}$, all stabilizers are subgroups of the grid generated by $\{hh', vv'\}$, which makes many properties easier.

Proposition 5 Let \mathcal{Z} be an SFT defined by \mathbb{Z}^2 -NN constraints. Then exactly one of the following occurs:

1. Every graph in \mathcal{Z} is supported by the infinite grid, and is strongly aperiodic.
2. \mathcal{Z} contains a finite graph and a strongly periodic infinite grid.

Proof. Assume that Condition 1 is not satisfied. A classical pumping argument from multidimensional symbolic dynamics (see for instance [3, Lemma 1.16]) shows that \mathcal{Z} then contains a strongly periodic configuration. In that case, its stabilizer, as a subgroup of \mathbb{Z}^2 , is residually finite (see for example [7, Prop.2.2.1]). Theorem 1 then implies that \mathcal{Z} contains a finite graph.

On the other hand, it is clear that both conditions cannot be satisfied simultaneously, so that they indeed form a dichotomy.

It is already interesting to note that the dichotomy from Proposition 5 is nontrivial, directly from the existence of strongly aperiodic SFTs over the grid (and other ones). But the next subsection will show a stronger result.

In contrast with $\mathcal{Z}_{\Sigma}^{\mathbb{Z}^2}$, $\mathcal{Z}_{\Sigma}^{\mathbb{F}_2^2}$ is much less rigid in the sense of the following statement, because it can be related to tilings over small-square surfaces, as studied in [8,19], and formalized in [13, Subsection 1.3.3].

Proposition 6 *Let \mathcal{Z} be an SFT defined by \mathbb{F}_2 -NN constraints. Then the following are equivalent:*

1. \mathcal{Z} is nonempty.
2. \mathcal{Z} contains a weakly periodic Cayley graph of \mathbb{F}_2 .
3. \mathcal{Z} contains a Cayley graph of a group different from \mathbb{F}_2 .

Besides, the following are equivalent:

4. \mathcal{Z} contains a finite graph.
5. \mathcal{Z} contains a finite Cayley graph.
6. \mathcal{Z} contains a strongly periodic Cayley graph of \mathbb{F}_2 .
7. \mathcal{Z} contains an amenable Cayley graph.

Nevertheless, there exist nonempty SFTs containing only finite graphs and defined by \mathbb{F}_2 -NN constraints.

The second equivalence is very close to the main result of [9], about how tilesets tile Cayley graphs. They also show the equivalence with containing a *cylindric* graph, *i.e.* with having an infinite cyclic dense stabilizer.

Proof.

- 1 \implies 2 If \mathcal{Z} is nonempty, then by Proposition 4, it contains a Cayley graph of \mathbb{F}_2 . It is known that every SFT over \mathbb{F}_2 contains a periodic configuration (see for instance [21, Theorem 2.2]).
- 2 \implies 3 Let X be a weakly periodic Cayley graph of \mathbb{F}_2 and \hat{u} be a nontrivial period. Then the infinite cyclic group $\langle \hat{u} \rangle$ is a residually finite stabilizer subgroup: if M is a defining window for \mathcal{Z} , then there exists an infinite (cyclic) subgroup $H < \langle \hat{u} \rangle$ which is $M\bar{M}$ -separated. By Proposition 3, X/H belongs in \mathcal{Z} .
- 3 \implies 1 This is trivial.
- 5 \implies 4 This is trivial.
- 6 \implies 5 This is due to Theorem 1.
- 4 \implies 6 If \mathcal{Z} contains a finite graph X , then from Proposition 4, it also contains a colored Cayley cover Y of X . All preimages of ε by the homomorphism are period in Y (because they are periods in X), and their set is M -dense, where M is a finite language representing all possible paths in X .
- 5 \iff 7 This is trivial because finite groups are amenable.
- 7 \implies 4 This is remarked in [19], based on [8].
- The last statement is also inherited from symbolic dynamics over groups [21, Example 4.1].

5.4 Undecidability

Theorem 2 *The dichotomy from Proposition 5 is undecidable, given a list of \mathbb{Z}^2 -NN constraints.*

Proof. Every set \mathcal{F} of \mathbb{Z}^2 -NN constraints defines both a (graph) SFT \mathcal{Z} , and a classical SFT \mathcal{Z}' over the infinite grid (which is included in \mathcal{Z}). If \mathcal{Z}' is strongly aperiodic, then \mathcal{Z} is in Class 1, whereas if it contains a weakly periodic configuration, then \mathcal{Z} is in Class 2. The dichotomy is then directly equivalent to strong aperiodicity of \mathcal{Z}' , which is known to be undecidable [11].

Consequently, the class of homogeneous graph SFTs that contains only finite graphs is computably unseparable from the class of graph SFTs containing both finite and infinite graphs. This means the undecidability of any property that is implied by one of the two classes, and which implies the negation of the other one.

Corollary 2 *The problem whether a graph SFT contains only finite graphs, and the problem whether a graph SFT is homogeneous, are both undecidable.*

More precisely, the reduction directly gives that the dichotomy is Σ_1^0 -complete, so that all general properties separating it are Σ_1^0 -hard. The encoding of Cayley graphs into graph SFTs also, more directly, implies the following undecidability result.

Theorem 3 *There exists a Cayley SFT $\mathcal{Z}_{\{0\}}^\Gamma$ such that the problem whether, given a word $u \in \Pi^*$, u represents a cycle from the origin back to itself in every graph of $\mathcal{Z}_{\{0\}}^\Gamma$, is undecidable.*

It is actually Σ_1^0 -complete.

Proof. From Proposition 4, u represents a cycle from the origin back to itself in every graph of \mathcal{Z}_Γ if and only if it represents the identity element in Γ , which means that it is in the language of the *word problem*. It is known that there exists a finitely presented group Γ whose word problem is undecidable.

6 Conclusion

To sum up, we provided a formalism generalizing SFTs on Cayley graphs by allowing the support graph to be determined by forbidden subgraphs. The generalization also allows the SFT to have finite configurations and we provided examples of such SFT. By defining quotient graphs, we have shown that the only way for an SFT to not have finite configurations is to only have configurations that are either aperiodic, or with a stabilizer that is not residually finite. Similarly, we have shown that having only one support graph requires configurations to have finite stabilizers. We used both results to prove that the finite configuration problem and the support graph unicity problem are undecidable.

In the future, we wish to study determinism and study directly the relation between causal graph dynamics and graph subshifts, and use this connection to prove structure and complexity results on causal graph dynamics. Moreover, graph subshifts could provide a new point of view for studying group subshifts. On the other hand, many notions and results known upon group subshifts such as sofic subshifts, determinism or entropy, would be interesting to lift to graph subshift.

We have seen that graph subshifts could generalize both classical symbolic dynamics and finite-type group geometry. They could be used to simulate a larger class of objects; let us discuss one example. Let \mathcal{T} be a finite set of *tiles*, *i.e.* polygons of the plane \mathbb{R}^2 . A *tiling* is then a covering of \mathbb{R}^2 by translated copies of tiles from \mathcal{T} , the intersection of any two of them being either empty, or a vertex, or a whole edge, or the whole tile. Now let us define subshift $\mathcal{Z}_\mathcal{T}$ with $\Sigma = \mathcal{T}$, π as the set of tile edges up to translation, and the following constraints: each vertex labeled $t \in \mathcal{T}$ must have exactly those ports which correspond to the edges of t , and every simple cycle must correspond to a valid vertex in a tiling of the plane by \mathcal{T} (*i.e.* it involves tiles that can be put next to each other around this vertex). The fact that the latter condition can be forced by finitely many constraints (so that $\mathcal{Z}_\mathcal{T}$ be an SFT) is equivalent to the classical *finite-local-complexity* property of \mathcal{T} . With this construction, there is a one-to-one correspondence between tilings by \mathcal{T} of the plane and of all surfaces and graphs in $\mathcal{Z}_\mathcal{T}$. Now if one is interested in the shift-orbit closure of a single tiling of the plane (*i.e.* the Penrose tiling), it is an interesting question whether this corresponds to an SFT in $\mathcal{Z}_\mathcal{T}$. Some work has been devoted to understanding for which tilings one can force locally tilings of the plane to be in the shift-orbit closure [10,6]; one says that the tilings admits *local rules*. Additionnally forbidding (via finite constraints) all possible graphs which represent tilings on surfaces but not on the plane is a new question. Our work shows that if the tiles are rectangles and if we allow for some decorations (that could also be encoded locally as small hanging subgraphs), then this is possible. It is tempting to think that this is still true in the case of general strongly aperiodic tilings, but the proofs would require a new formalism, because the support is no longer a Cayley graph. The same could be done in the tridimensional space, or in the hyperbolic plane.

Let us make a final remark. Graph theorists are used to considering families of graphs defined by forbidden finite graph minors. Of course, seen as including possibly infinite graphs, these families are subshifts, but they are a priori not SFTs (not even sofic). However, even if the *minor* operation involves vertices at possibly unbounded distance, it is known to have some kind of representation as local constraints [15] (involving a nondeterministic coloring). The relation with our setting could be investigated.

Acknowledgements

We are indebted to Nicolas Schabanel who encouraged us to study graph subshifts and contributed to the first four definitions. This publication was made possible through the support of the ID# 62312 grant from the John Templeton Foundation, as part of the ‘The Quantum Information Structure of Spacetime’ Project (QISS) . The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation.

References

1. Pablo Arrighi, Amélia Durbec, and Aurélien Emmanuel. Reversibility vs local creation/destruction. In Michael Kirkedal Thomsen and Mathias Soeken, editors, *Reversible Computation - 11th International Conference, RC 2019, Lausanne, Switzerland, June 24-25, 2019, Proceedings*, volume 11497 of *Lecture Notes in Computer Science*, pages 51–66. Springer, 2019. doi:10.1007/978-3-030-21500-2_4.
2. Pablo Arrighi, Simon Martiel, and Vincent Nesme. Cellular automata over generalized cayley graphs. *Mathematical Structures in Computer Science*, 28(3):340–383, 2018. doi:10.1017/S0960129517000044.
3. Alexis Ballier. *Propriétés structurelles, combinatoires et logiques des pavages*. Ph.d. thesis, Université de Provence, Nov 2009.
4. Gilbert Baumslag and Donald Solitar. Some two-generator one-relator non-hopfian groups. *Bulletin of the American Mathematical Society*, 68(3):199–201, 1962. 00471. URL: <https://www.ams.org/bull/1962-68-03/S0002-9904-1962-10745-9/>, doi:10.1090/S0002-9904-1962-10745-9.
5. Robert Berger. *The undecidability of the domino problem*. Phd, Harvard University, jul 1964.
6. Nicolas Bédaride and Thomas Fernique. When periodicities enforce aperiodicity. *Communications in Mathematical Physics*, 335(3):1099–1120, May 2015. arXiv: 1309.3686. URL: <http://arxiv.org/abs/1309.3686>, doi:10.1007/s00220-015-2334-8.
7. Tullio Ceccherini-Silberstein and Michel Coornaert. *Cellular Automata and Groups*. Springer, 2009.
8. J.-R. Chazottes, J.-M. Gambaudo, and F. Gautero. Tilings of the plane and Thurston semi-norm. *arXiv.org*, May 2012. URL: <https://arxiv.org/abs/1205.5118v2>.
9. Benjamin Hellouin de Menibus. Tiling an arbitrary amenable group is decidable. private communication, 2021.
10. Thomas Fernique and Mathieu Sablik. Weak colored local rules for planar tilings. *arXiv:1603.09485 [cs, math]*, March 2016. URL: <http://arxiv.org/abs/1603.09485>.
11. Yuri Gurevich and I. Koryakov. A remark on berger’s paper on the domino problem. *Siberian Journal of Mathematics*, 13:459–463, 1972.
12. G. A. Hedlund. Endomorphisms and automorphisms of the shift dynamical system. *Math. Systems Theory*, 3:320–375, 1969.
13. Emmanuel Jeandel and Pascal Vanier. The Undecidability of the Domino Problem. In Shigeki Akiyama and Pierre Arnoux, editors, *Substitution and Tiling Dynamics: Introduction to Self-inducing Structures*, Lecture Notes in Mathematics, page 66. Springer, 2017.
14. Jarkko Kari. The nilpotency problem of one-dimensional cellular automata. *siam Journal on Computing*, 21(3):571–586, 1992.
15. Dietrich Kuske and Markus Lohrey. Logical aspects of Cayley-graphs: the group case. *Annals of Pure and Applied Logic*, 131(1-3):263–286, 2005. URL: <http://linkinghub.elsevier.com/retrieve/pii/S0168007204000831>, doi:10.1016/j.apal.2004.06.002.
16. C. G. Langton. Self-reproduction in cellular automata. *Physica D*, 10:135–144, 1984.
17. Douglas A. Lind and Brian Marcus. *An introduction to symbolic dynamics and coding*. Cambridge University Press, Cambridge ; New York, 1995. 03085.
18. Roger C. Lyndon and Paul E. Schupp. *Combinatorial Group Theory*. Classics in Mathematics. Springer-Verlag, Berlin Heidelberg, 2001. 04184. URL: <https://www.springer.com/gp/book/9783540411581>.
19. Thierry Monteil. Down with hierarchy !, 2013.
20. A Yu Ol’shanskiĭ. *Geometry of defining relations in groups*, volume 70. Springer Science & Business Media, 2012.
21. Steven Piantadosi. Symbolic dynamics on free groups. *Discrete & Continuous Dynamical Systems - A*, 20(3):725–738, 2008. URL: <http://aimsciences.org//article/doi/10.3934/dcds.2008.20.725>, doi:10.3934/dcds.2008.20.725.