THÈSE

Pour obtenir le diplôme de doctorat

Spécialité **INFORMATIQUE**

Préparée au sein de l'**Université de Caen Normandie**

Indécidabilité des invariants géométriques dans les pavages

Présentée et soutenue par

LEO PAVIET SALOMON

Thèse soutenue le 17/12/2024

devant le jury composé de :

M. PASCAL VANIER	Professeur des universités - Université de Caen Normandie (UCN)	Directeur de thèse
M. NICOLAS OLLINGER	Professeur des universités - Université d'Orléans	Président du jury
M. JULIEN CERVELLE	Professeur des universités - UNIVERSITE PARIS 12 VAL DE MARNE	Membre du jury
M. MATHIEU HOYRUP	Chargé de recherche - Université de Lorraine	Membre du jury
M. SAMUEL PETITE	Professeur des universités - UNIVERSITE AMIENS PICARDIE JULES VERNE	Membre du jury
MME MARIE-PIERRE BEAL	Professeur des universités - Université Gustave Eiffel	Rapporteur du jury

Thèse dirigée par **PASCAL VANIER** (Groupe de recherche en informatique, image et instrumentation de Caen)

Contents

Acknowledgements

Notations and conventions

$D \subset_f \mathbb{Z}^d$	D is a *finite* subset of $\mathbb{Z}^d$
$f\colon A \hookrightarrow B, g\colon A \twoheadrightarrow B$	f, g are injective, surjective from A to B
π_A	The natural projection on A
$\mathfrak{S}_X$	The symmetric group (set of bijections, with composition) on X
$\mathbf{u}$	Point (or vector) of $\mathbb{Z}^d$
$\mathbf{e_i}$	$\mathbf{e_i} = (\underbrace{0,\ldots,0}_{i}, 1, \underbrace{0,\ldots,0}_{d-i-1})$, the i-th basis vector of $\mathbb{Z}^d$.
$[\![-n, n]\!]$	The set $\{i \in \mathbb{Z} \vert -n \leq i \leq n\}$
$\mathcal{B}_n$	The $\mathbb{Z}^d$-ball $[\![-n, n]\!]^d$
$\mathcal{Q}_n$	The set $[\![0, n-1]\!]^d$
$\mathcal{A}$	A finite set of symbols, called an alphabet
$\mathcal{A}^*$	The set of all the d-dimensional square patterns over $\mathcal{A}$
$\vert w\vert_a$	Number of occurrences of $a \in \mathcal{A}$ in the pattern w
$\mathcal{X}_\mathcal{F}$	The subshift defined by the forbidden patterns $\mathcal{F}$
$w \sqsubseteq x$	The configuration (or pattern) x contains the pattern w
$\sigma_\mathbf{u}$	Shift operator by $\mathbf{u} \in \mathbb{Z}^d$
$\mathcal{L}(X), \mathcal{L}_n(X)$	The language of X, the $\mathcal{Q}_n$-supported language of X
$[w]$	Cylinder $[w] = \left\{x \in X \middle\vert x\vert_{\mathrm{supp}(w)} = w\right\}$
$\mathfrak{s}$	A substitution
$\bar{S}$	The set $S \cup \{s^{-1} \mid s \in S\}$
$\pi_1(X)$	Fundamental group of a topological space X
$\pi_1^{proj}(X)$	Projective fundamental group of X
$\langle S \mid R\rangle$	Presentation of a group G by generators S and relations R
TOTAL	A decision problem
$M(x)\Downarrow, M(x)\Downarrow^n, M(x)\Uparrow, M(x)\Uparrow^n$	On entry x, the machine M halts, halts in at most n steps, runs forever, runs for at least n steps

Introduction

Tilings, colourings, decorations

This thesis studies problems which are all related to *tilings*. In everyday life, tilings are usually decorative coverings of walls or floors, which contrary to mosaics use a limited amount of different *tiles* to produce geometric and aesthetically pleasing *patterns*. To the computer scientist's eye, a tiling's artistic value is second to other questions:

1. Given a box with sufficiently many copies of a few different tiles, will a tiler be able to properly tile an entire wall ? More importantly, will they be able to tile an *infinite* wall (unlikely as it is for a regular tiler to encounter such a wall, computer scientists have plenty of time to devote to tiling infinite surfaces) ?

2. Assuming you find a clever way to ensure that a given surface can in fact be tiled, can you say anything about the kind of patterns you can produce ? After all, even with finitely many different tiles, they might allow for surprisingly complex arrangements: can you count them, or describe them in any way ?

3. Suppose that you are not particularly happy with the appearance of your walls' tiling. Can you locally modify it so as to replace the parts you are not satisfied with, without having to change the rest of the tiling ? More generally, does the structure of the valid tilings allow for purely local changes ? If not, can you describe the kind of modifications that must be done in order to change a finite part of tiling (for example, do you need to modify a finite region but larger than the one you initially wanted to modify ? Do you need to modify tiles in arbitrary locations, or only in a specific region of the plane, *e.g.* to the right of your initial changes ?)

4. Given that some of your walls are not exactly planar, but have some depth, relief, and non-trivial curvature, does the set of possible tilings change when taking this additional geometrical constraints into account ? More generally, can you imagine different surfaces or geometries such that tilings would be completely different from what you would obtain by tilings plain bi-dimensional walls ?

Fortunately, our computer scientist can entirely formalize these problems. When done in the simplest possible way, they become questions about Wang Tiles: these are square tiles, which all have the same size, and which have some decorations (or colours) on each one of their four sides. In order to model the fact that tiles must fit together, two Wang tiles can be placed next to one another only if they share the same colour on the sides that are in contact. If we moreover enforce that tiles must all aligned, each tiling of some surface D using tiles from a set $\mathcal{A}$ can be viewed as an element of $\mathcal{A}^D$ where $D \subseteq \mathbb{Z}^2$. Using this formalism, it becomes easy to ask the same questions for surfaces $D \subseteq \mathbb{Z}^d, d \geq 1$ instead, giving the higher-dimensional computer scientist some work to do. We can even further generalize the kind of tilings being considered: instead of only enforcing matching rules between neighbouring tiles, we can enforce additional constraints between tiles which might further apart. To be able to give precise instructions to the tiler, the computer scientist chooses to specify each such constraint as a *forbidden pattern*: that is, a pattern, or specific arrangement of finitely many tiles, that must not appear in any tiling. Letting $\mathcal{F}$ be a finite set of such rules, we will call $\mathcal{X}_{\mathcal{F}} \subseteq \mathcal{A}^{\mathbb{Z}^2}$ the set of tilings of the infinite plane by the tiles $\mathcal{A}$ that respect all the rules, that is, which avoid any pattern of $\mathcal{F}$, everywhere. This is called a **subshift of finite type**, and will be an important object studied throughout the thesis. If the computer scientist wants to be especially annoying to their tiler, they can instead specify infinitely many such rules: the set of valid tilings will then simply be called a **subshift**, although there is no telling what its patterns might look like. Surprisingly, this is all we need to answer by the negative most of the questions asked so far:

1. The first question is known as the **Domino problem**. We describe it in Section 1.2.3, and show that there is no general method which solves it in Theorem 1.77 ([Ber66]).

2. As a consequence, the second question does not admit a satisfying answer either in general. The specific question of counting the number of valid patterns of a given size will be studied in Definition 1.23, and the broader question of describing tilings and their patterns, up to some loss of specificity, is the idea behind the definition of conjugacy invariants Definition 1.21: this as a way to formalize what it means for tilings to be equivalent, or to say that the patterns of one can be "derived" from the patterns of the other.

3. The third question will be studied in depth in Chapter 2. Our main question will be to quantify in some way the amount of patterns of various sizes which can freely be replaced by one another, or swapped out, in larger tilings, and to find some precise quantitative restrictions on the number of such patterns.

4. The fourth question is excessively general, and will be studied from two different points of view in the thesis: in Chapter 3, we will try to understand what happens if you try to tile two dimensional surfaces, where you possibly leave some large holes, or gaps, in the tiling. It turns out that this can significantly alter the set of valid tilings, and we will try to understand the structure of those new tilings from an algebraic perspective. In Chapter 4, we instead completely relax the condition that tiles must assembled in any kind of geometric way, and will study abstract arrangements, known as tilings on **graphs**. In fact, tilings of $\mathbb{Z}^2$ and $\mathbb{Z}^d$ for $d \geq 2$ often share similar properties, but it turns out that we can relax generalize subshifts to be defined on much more general spaces, such as groups or even arbitrary graphs. We will focus on a specific class of tilings, namely **substitutive tilings**, and try to relate their structure when defined as subshifts on graphs compared to the well-known case of substitutive subshifts of $\mathbb{Z}^d$.

An important *leitmotiv* underpinning the present document is the fact that despite their elementary definition, tilings, or more precisely subshifts, exhibit an immense variety of complex properties, ranging from algorithmic undecidability to algebraic characterizations. A recurring, informal but important question that we will try to study from several angles is therefore the following: how *complex* can tilings be (algorithmically, combinatorially, algebraically), even when we restrict ourselves to the simplest possible classes of subshifts ?

Some historical background

The study of subshifts is an important part of a mathematical domain called *symbolic dynamics*. The "dynamical" aspect of the informal definitions given above is not apparent.

Indeed, although the first part of this introduction introduces subshifts as a way to model arrangements of tiles which respect some local constraints, one of the reasons behind for the variety and large panorama of interesting questions about subshifts can be understood by looking at subshifts from another point of view, namely, the one of dynamical systems. Before that, let us take a step back and give some additional details on the first approach.

We generally attribute to Hao Wang the definition of Wang tiles, introduced in [Wan61, Section 4.1], while studying whether there existed procedure deciding the satisfiability of logical formulaes of the $\forall\exists\forall$. His main conjecture was that it was decidable whether or not a given set of tiles $\mathcal{A}$ tiled the plane, that is, if one could place a tile of $\mathcal{A}$ on each position of the infinite grid $\mathbb{Z}^2$ while respecting the adjacency constraints between neighbouring tiles. Although false, this conjecture highlights the importance of decision problems, and more generally computability theory, in the study of multi-dimensional subshifts. Indeed, the proofs of the undecidability of the domino problem, first by Berger [Ber66] and Robinson [Rob71], and later by Kari [Kar96] using a completely different method, have successfully been adapted to show that those links between symbolic dynamics and computability theory exist in very different settings (see for example and among many others [BS16], [ABM19], [AK21], [Bar22], [EGL23]). More recently, computability theory has found other uses to precisely characterize some behaviours of subshifts (formally, the values taken by conjugacy invariants, which we define in Section 1.1.2), in *e.g.* [HM10] or [JV13], and we will ourselves continue this line of research in Chapter 2 – we say a few more words about this at the end of this introduction.

On the other hand, the earliest results stated with a formalism which is similar to the modern are proven as early as [MH38], and several decades ago if we consider some specific problems (for example, [Thu12] studies the so-called Prouhet-Thue-Morse sequence). This historical approach studies problems which are mainly related to infinite *sequences* of symbols, which can be viewed as discrete versions of trajectories in a dynamical system: a dynamical system is simply a space X with some self-map T, that you iterate on X, possibly infinitely many times, to obtain trajectories $(T^{n(x)})_{n\in\mathbb{Z}}$ when T is furthermore invertible. If instead of recording precisely $T^n(x)$ at each timestep, you only keep track of a coarser information, *e.g.* some element $X_i \ni T^n(x)$ in a partition $X = \bigsqcup_{i=0}^{k} X_i$ of the space, you obtain an infinite sequence which approximates the original trajectory. This point of view motivates the study of those sequences from another perspective, using tools from the theory of dynamical systems, and some objects already mentioned in this introduction (conjugacy invariants, entropy ...) are in fact typical objects from this theory, and their interplay with the more combinatorial definition of a subshift makes symbolic dynamics a new, interesting field on its own.

Structure of the document

Each chapter contains a more detailed introduction, and motivating examples and questions. We briefly summarize here the contents of each chapter, as well as state our main theorems.

The Chapter 1 is a general introductory chapter, and has to be read first. We recall the main definitions that will be used throughout the entire document, and the necessary background and preliminary results in order to familiarize with the various objects. In particular, as we study subshifts and tilings using various tools, we recall the important notions of computability theory and group theory that we need in other chapters.

The three chapters Chapter 2, Chapter 3 and Chapter 4 are independent, and can be read in any order:

- Chapter 2 is devoted to the study of a conjugacy invariant named **extender entropy**, from the point of view of computability theory. Introduced in [FP19], it is based on classical notions studied in formal languages ([Myh57], [Ner58]) and in subshifts themselves (see for example [KM13], [FOP16]). The idea behind the definition is that instead of trying to count the number of valid patterns in a subshift, we are interested in how many of those patterns are "equivalent", in the sense that they can be exchanged, or replaced by one another, in any valid configuration. This quantitatively captures the informal idea that there is some amount of information "flowing" from the inside of a pattern to the outside regarding how it can legally be extended while respecting the constraints of the subshift. Our main results completely characterize the possible values of extender entropy on various classical classes of subshifts, both one-dimensional and multi-dimensional, using classifications of real numbers in terms of their computability properties. Our main results are the following:

Theorem (Theorem 2.43)**.** *The extender entropies of effective $\mathbb{Z}$-subshifts are exactly the non-negative Π_3 real numbers.*

Theorem (Theorem 2.48)**.** *For any $d \geq 2$, the extender entropies of sofic $\mathbb{Z}^d$-subshifts are exactly the non-negative Π_3 real numbers.*

- We study in Chapter 3 another conjugacy invariant introduced in [GP95], called the **projective fundamental group** $\pi_1^{proj}(X)$ of a subshift X, an object algebraic nature which captures some geometrical properties of the subshift. Contrary to the extender entropy studied in Chapter 2, which is interested in how finite patterns can be extended, we can view the projective fundamental group as a way to investigate how "holes" in an infinite tiling can be filled to produce a valid complete configuration of the subshift. Like the usual fundamental group of a topological space, this group moreover gives precise information about the obstructions that some "tilings with holes" have with regards to being completed into entire configurations. We study various properties of this group, and relate some of its properties with well-known notions of symbolic dynamics. Our main result is about the possibility of realizing a large class of groups as projective fundamental groups of "simple" subshifts:

  Theorem (Theorem 3.81)**.** *Let $G = \langle S \mid R \rangle$ be a finitely presented group. Then, there exists an SFT X such that:*

  - *X is projectively connected.*
  - *$\pi_1^{proj}(X) \simeq G$*

- Chapter 4 is more exploratory, and tries to understand how the expressive properties of subshifts of finite type change with the geometry of the ambient space (until know, $\mathbb{Z}^d$). More precisely, we try to adapt a classical result, the Mozes theorem [Moz89], in a much more elementary setting: this theorem states that a specific class of subshifts, namely substitutive subshifts, can be "enforced" by subshifts of finite type (formally, we say that they are **sofic** sunshifts, a definition we give in Section 1.1.4 and revisit in Section 4.3.5). This is a somewhat surprising result, as substitutive subshifts exhibit a hierarchical, global structure, while subshifts of finite type are only able to enforce finitely many purely local constraints. Nevertheless, this result has been adapted in a variety of settings, ([Goo98], [FO10], [BS16]), and the goal of this chapter is to understand whether this result still holds when considering subshifts on *graphs* rather than $\mathbb{Z}^d$, using only combinatorial arguments, without using any underlying geometrical properties of the ambient space being tiled. Another motivation for the introduction of this more general setting is the fact that a similar direction has already been explored in symbolic dynamics, with subshifts on more general groups than $\mathbb{Z}^d$, or graphs of algebraic origin, with similar questions to the one we ask in this chapter ([Sil20], [Bar22]). Our main result is a weaker version of Mozes theorem:

Theorem (Theorem 4.68)**.** *Let* $\mathfrak{s}$ *be a graph substitution, and* $\mathfrak{s}_c$ *a coloured* $\mathfrak{s}$-*substitution. Suppose that* $\mathfrak{s}$ *is quasi-connected. Then, there exists a sofic graph subshift* $Y_{\mathfrak{s}_c}$ *which is* $X_{\mathfrak{s}_c}^{\infty}$*-sheeted and contains* $X_{\mathfrak{s}_c}^{\infty}$*.*

Chapter 1

First definitions

All the problems studied throughout this thesis are related to **tilings**. We will not try to give any definition general enough to encompass all the different kinds of tilings one might encounter, neither in the literature nor even in this thesis. However, we will see how two different formalisms help us characterize what we mean, in context, by “tiling”, and how they influence the kind of questions one might then ask about tilings.

Far from being independent, these two formalisms merely are two distinct point of views on the same underlying mathematical objects: although their history differs, one must play with these multiple facets in order to solve most problems related to tilings. The first formalism that we describe in Section 1.1.1 is one dating back to Wang [Wan61], and is more combinatorial in nature. In this setting, tilings are defined as some generalized colourings, which verify some local, puzzle-like constraints. This formalism will in particular influence the definitions of Chapter 4. The second point of view, presented in Section 1.1.3, is by contrast closely linked with the field of *dynamical systems*, and carries with it the tools and the questions from this domain. It defines tilings as subsets of an abstract space satisfying some dynamical conditions. In fact, most definitions of Section 1.1.3 will be reformulations in our specific setting of the more general ones for dynamical systems; the combinatorial nature of tilings then provides a more concrete representation of these properties. In this chapter, Section 1.1 presents both points of view and gives other general results about subshifts, as well as a few classical classes of subshifts that will be further studied in this thesis.

We also introduce in Section 1.2 the theory of *computability*. It is mainly used in Chapter 2, but we will nevertheless use some results in other sections of the thesis.

Finally, in Section 1.3, we give a quick overview of some algebraic notions that we need in this thesis: the notion of group presentation which will be essential for Chapter 3 and in particular Section 3.5, and the notion of Cayley graphs, which motivates some constructions in Chapter 4.

This document assumes that the reader already has a basic knowledge of some mathematical objects, most notably metric spaces and topological spaces, and basic group theory. We do not assume any advanced knowledge of these objects, and will recall in this first chapter the definitions that the reader might not be familiar with. We do not assume any familiarity with dynamical systems in general, and tiling spaces and symbolic dynamics in particular. In any case, we try to give some useful intuitions and motivations behind our definitions and the questions we investigate, and we do not try to give the most general possible statements or smartest proofs if we believe that doing so would be detrimental to clarity.

1.1 Subshifts and tilings

1.1.1 Subshifts: patterns and local rules

Wang Tiles: an historical point of view

In 1961, while studying problems related to fragments of first order logic, Hao Wang introduced an object now called **Wang tiles** [Wan61], and the associated tilesets. In the following definition, $\mathcal{C}$ is a finite set of *colours*.

Definition 1.1: Wang tiles

A **Wang tile** is a quadruplet $c = (c_W, c_S, c_E, c_N) \in \mathcal{C}^4$.
A **tileset** is a finite set of Wang tiles.

We typically represent Wang tiles as unit squares, with coloured sides:

Using these tiles, the goal is to find a valid **tiling** of the plane: the rules are simple, one wants to place one tile at each position of $\mathbb{Z}^2$, and neighbouring tiles must have the same colour on their common side.

Definition 1.2: Tiling

Given a tileset T, we define a T-**tiling** as a map $\tau\colon \mathbb{Z}^2 \to T$ which verifies that for all $(i, j) \in \mathbb{Z}^2$:

- $\tau(i, j)_E = \tau(i + 1, j)_W$
- $\tau(i, j)_N = \tau(i, j + 1)_S$

The set of all the T-tilings is called a **tiling space**, denoted by $\mathcal{X}_T$.

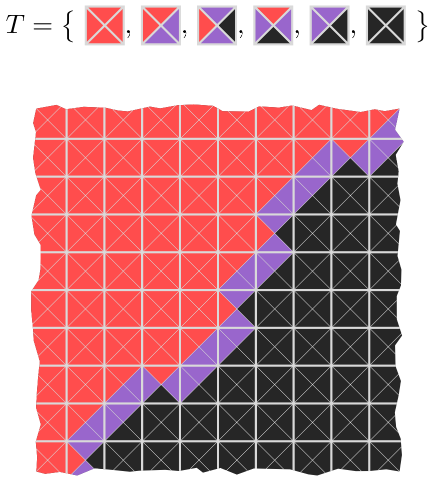

Figure 1.1: A Wang tileset T and a configuration of $\mathcal{X}_T$.

One can then ask all sorts of questions about the tiling space of some tileset T, and we will see in particular in Section 1.2.3 how seemingly simple and natural problems about Wang tiles and tiling spaces happen to be computationally involved. For now, let us take a step back and introduce a slightly different formalism, which will both be easier to generalize and to work with. Most results and constructions presented in this section can be found in [LM21].

From tilings to subshifts and forbidden patterns

We call **alphabet** a finite set $\mathcal{A}$ of **symbols**. We call a mapping from $\mathbb{Z}^d$ to $\mathcal{A}$ a **configuration**: thinking of $\mathbb{Z}^d$ as a graph and of $\mathcal{A}$ as a set of colours, a configuration is then a colouring of the grid $\mathbb{Z}^d$ by $\mathcal{A}$. The set of all configurations is $\mathcal{A}^{\mathbb{Z}^d}$. For a configuration x and some position $\mathbf{i} \in \mathbb{Z}^d$, we usually use the subscript notation $x_{\mathbf{i}}$ rather than the functional notation $x(\mathbf{i})$.

In Chapter 2 and Chapter 3, we will only consider the case of $\mathbb{Z}^d$ configurations, as explained above. Most of the time, we will even restrict ourselves to the particular cases $d = 1$ and $d = 2$. This is the most common setting in the literature, for reasons that will be detailed below. However, Chapter 4 tries to extend results to a case of more general structures, that we call *self-similar graphs*. In that case, configurations will then be elements of $\mathcal{A}^G$ for some infinite graph G, but we defer the precise definitions to this chapter to simplify this first exposition.

We will mainly be interested in *sets* of configurations. However, we will not talk about generic subsets of $\mathcal{A}^{\mathbb{Z}^d}$, and we will only consider sets called **subshifts**, corresponding to the following intuition:

- A subshift is the set of all the configurations verifying some common condition, or restriction.
- This restriction has to be:
  - Local: in order to determine if a configuration satisfies the condition, it should be enough to verify it locally, without looking at the entire configuration at once.
  - Homogeneous: the restriction, which is local as required above, should be verified everywhere in the configuration.

These conditions try to capture the fact that we want to model and study *tilings*, in which the tiles have to "fit" together – this is the locality condition – and we tile the surface using the same tiles everywhere – this corresponds to the required homogeneity.

More formally, we will define a subshift as the set of configurations avoiding some **patterns**.

Definition 1.3: Pattern

A **pattern** is a function $u\colon D \to \mathcal{A}$, where D is a finite subset of $\mathbb{Z}^d$ called the **domain** (or **support**) of u, denoted by $\mathrm{dom}(u)$.

It will sometimes be convenient to consider patterns only up to translation. In that case, two patterns u, v would be considered the same if there exists some translation $\mathbf{i} \in \mathbb{Z}^d$ such that $u = \mathbf{i} + v$. This is in fact what we usually mean by a pattern, but to avoid having to explicitly work with equivalence classes, we often talk about "the" support of a pattern – we will try to be explicit if and where the distinction matters.

For example, we say that a configuration x **contains** a pattern u at position $\mathbf{i} \in \mathbb{Z}^d$ and we write $u \sqsubseteq x$ if for all $\mathbf{j} \in \mathrm{dom}(u)$, $x_{\mathbf{i}+\mathbf{j}} = u_{\mathbf{j}}$. A configuration is said to **avoid** a pattern u if it does not contain it, in which case we naturally write that $u \not\sqsubseteq x$.

Definition 1.4: Subshift

Given an alphabet $\mathcal{A}$, a dimension d and a family $\mathcal{F}$ of finite patterns, we define the d-dimensional subshift $\mathcal{X}_{\mathcal{F}}$ as the set of all the configurations that do not contain any pattern of $\mathcal{F}$:

$$\mathcal{X}_{\mathcal{F}} = \left\{ x \in \mathcal{A}^{\mathbb{Z}^d} \mid \forall u \in \mathcal{F}, u \not\sqsubseteq x \right\}$$

We call **points** of the subshift the configurations $x \in \mathcal{X}_{\mathcal{F}}$. The family $\mathcal{F}$ is *a* family of **forbidden patterns** defining $\mathcal{X}_{\mathcal{F}}$.

Given $\mathcal{F}$, we will also say that the configurations $x \in \mathcal{X}_{\mathcal{F}}$ are **valid**.

The definition of a subshift uses a negative condition – subshifts are sets of configuration which *avoid* a specific family of forbidden patterns – which tells us little about what those configurations actually look like. In particular, given a family $\mathcal{F}$, it is possible for a pattern u not to contain any pattern of $\mathcal{F}$, while still being forbidden in $\mathcal{X}_{\mathcal{F}}$ as a consequence of the other patterns being explicitly rejected: in this case, we say that u is **locally admissible**, but not **globally admissible** or **extensible**. For example, on the binary alphabet $\mathcal{A} = \{0, 1\}$ and considering subshifts over $\mathbb{Z}$, forbidding the patterns $\mathcal{F} = \{01, 10, 111\}$ necessarily prevents 11 from appearing in any valid configuration. In particular, this means that several families of forbidden patterns could be used to define the same subshift.

It is possible to state the definition using a positive property, and specifying which patterns are *allowed* instead:

Proposition 1.5: Allowed patterns

If $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ is a subshift, there exists a family of allowed finite patterns $\mathcal{G}$ such that

$$X = \left\{ x \in \mathcal{A}^{\mathbb{Z}^d} \mid \forall D \subset_f \mathbb{Z}^d, x\big|_D \in \mathcal{G} \right\}$$

Proof. Define

$$\mathcal{G} = \bigcup_{x \in X} \bigcup_{D \subset_f \mathbb{Z}^d} x\big|_D$$

□

However, we will almost always define subshifts using families of forbidden patterns. Indeed, for an allowed family $\mathcal{G}$, it is not clear that any given $g \in \mathcal{G}$ actually appears in some configuration of X – said differently, specifying the locally admissible patterns is not sufficient to determine those which are extensible. When defining a subshift more informally, we will however switch between the two points of view.

Some basic definitions

The definition of a subshift suggests that we can quantify how complicated a subshift is by considering the "simplest" family $\mathcal{F}$ defining it. The most natural such class of subshifts is the class of **subshifts of finite type**, or SFTs for short:

Definition 1.6: SFT

An **subshift of finite type** (SFT) is a subshift $X = \mathcal{X}_{\mathcal{F}}$ for some finite family of forbidden patterns $\mathcal{F}$.

The simplest example of an SFT is the **full-shift**:

Definition 1.7: Full-shift

The **full-shift** over an alphabet $\mathcal{A}$ is $\mathcal{A}^{\mathbb{Z}^d} = \mathcal{X}_{\emptyset}$.

We can now explain where the names sub*shift* and full-*shift* come from: given the previous definition, we can see that a subshift is a subset of the full-shift (obtained by forbidding some patterns in the full-shift). Now, the "shift" comes from the **shift functions**:

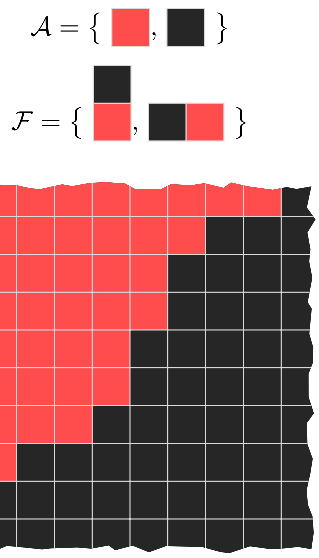

Figure 1.2: An example of a configuration from an SFT defined by two forbidden patterns on a binary alphabet

Definition 1.8: Shift

For any $\mathbf{i} \in \mathbb{Z}^d$, we denote $\sigma_{\mathbf{i}}$ and call **i**-shift the map

$$\begin{aligned} \sigma_{\mathbf{i}} \colon \mathcal{A}^{\mathbb{Z}^d} &\to \mathcal{A}^{\mathbb{Z}^d} \\ x &\mapsto (\mathbf{j} \in \mathbb{Z}^d \mapsto x_{\mathbf{i}+\mathbf{j}}) \end{aligned}$$

In the case of $\mathbb{Z}$-subshifts, we simply write $\sigma = \sigma_1$.

The shift functions are translations of configurations: $\sigma_{\mathbf{i}}$ simply translates entire configurations by **i**. The aforementioned "homogeneity property" of subshifts can be reformulated more formally as the fact that a subshift is σ-invariant:

Proposition 1.9

Let $X = \mathcal{X}_{\mathcal{F}} \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a subshift. Then for any $\mathbf{u} \in \mathbb{Z}^d$, $\sigma_{\mathbf{u}}(X) = X$.

Proof. If x does not contain any pattern from $\mathcal{F}$, neither does $\sigma_{\mathbf{u}}(x)$. Moreover, $\sigma_{\mathbf{u}} \circ \sigma_{-\mathbf{u}} = \sigma_{-\mathbf{u}} \circ \sigma_{\mathbf{u}} = \mathrm{id}_X$, hence $\sigma_{\mathbf{u}}$ is bijective and so $\sigma_{\mathbf{u}}(X) = X$. □

For an arbitrary subset of the full-shift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$, requiring that for all **i**, $\sigma_{\mathbf{i}}(X) = X$ is not sufficient to ensure that X is a subshift, and we will detail in Section 1.1.3 the missing condition. We call **orbit** of a configuration the set of all its shifts:

Definition 1.10: Orbit

Let $x \in \mathcal{A}^{\mathbb{Z}^d}$ be a configuration. The **orbit** of x is the set

$$\mathrm{Orb}(x) = \bigcup_{\mathbf{u} \in \mathbb{Z}^d} \sigma_{\mathbf{u}}(x)$$

We give one more definition, to be able to refer to the extensible patterns of a subshift in a more precise and succinct way.

Notation. For any $n \geq 1$, we write $\mathcal{Q}_n = [\![0, n-1]\!]^d \subset \mathbb{Z}^d$ ($\mathcal{Q}$ standing for "quadrant"). For $n \geq 0$, write $\mathcal{B}_n = [\![-n, n]\!]^d$.

In order to lighten the notation, we do not specify the dimension d and use $\mathcal{Q}_n, \mathcal{B}_n$ regardless of d, context making it clear.

Notation. For any alphabet $\mathcal{A}$, we write $\mathcal{A}^* = \bigsqcup_{n\geq 1} \mathcal{A}^{\mathcal{Q}_n}$. Here again, the dimension is implicit and should be clear from the context. In dimension 1, this is the standard notation for the set of all finite words (or finite sequences) on $\mathcal{A}$.

Definition 1.11: Language

Let X be a subshift. For a finite domain $D \subset_f \mathbb{Z}^d$, we note

$$\mathcal{L}_D(X) = \left\{ x\big|_D \mid x \in X \right\}$$

the D-language of X. In the specific case $D = \mathcal{Q}_n$, we write $\mathcal{L}_n(X) = \mathcal{L}_{\mathcal{Q}_n}(X)$ the n-language of X.
The **language** of X is then the set of all its extensible patterns, *i.e.*

$$\mathcal{L}(X) = \bigcup_{D \subset_f \mathbb{Z}^d} \mathcal{L}_D(X)$$

It will sometimes be useful to also talk about the language of a single configuration $x \in X$. In that case, we naturally write $\mathcal{L}_D(x), \mathcal{L}_n(x)$ and $\mathcal{L}(x)$ without ambiguity. Note that in those expressions, we consider patterns supported by D or $\mathcal{Q}_n$ up to translation:

$$\begin{aligned} \mathcal{L}_n(x) &= \bigcup_{\mathbf{u} \in \mathbb{Z}^d} x\big|_{\mathbf{u}+\mathcal{Q}_n} \\ &= \bigcup_{y \in \mathrm{Orb}(x)} y\big|_{\mathcal{Q}_n} \end{aligned}$$

Back to Wang tiles

It is now pretty clear that the definition of subshifts is more general than what we could express with Wang tiles and tiling spaces; in particular, given a tileset T, the set of all the T-tilings is a $\mathbb{Z}^2$ subshift on the alphabet T: moreover, it is an SFT, as we can define it by forbidding all the pairs of adjacent tiles whose shared sides do not have the same colour. In fact, this is an equivalence: all SFTs on $\mathbb{Z}^2$ can be re-encoded into a Wang tileset. The precise definition of what we mean by "re-encoding" will be given in Section 1.1.2, but it suffices for now to detail this particular example.

Consider an alphabet $\mathcal{A}$, and X a $\mathbb{Z}^2$ subshift of finite type over $\mathcal{A}$. By definition, there exists a finite family of forbidden two-dimensional patterns $\mathcal{F}$ such that $X = \mathcal{X}_\mathcal{F}$. Let r be the maximal size of a pattern of $\mathcal{F}$: as forbidden patterns are only considered up to translation when defining $\mathcal{X}_\mathcal{F}$, this means that we can consider that $\mathrm{dom}(w) \subseteq \mathcal{Q}_r$ for all $w \in \mathcal{F}$. By completing those patterns in all the possible ways, we can even assume that $\mathcal{F} \subset \mathcal{A}^{\mathcal{Q}_r}$.

Now, let $\mathcal{G} = \mathcal{A}^{\mathcal{Q}_r} \setminus \mathcal{F}$, and let $\mathcal{C}_H = \mathcal{A}^{[\![0,r]\!]\times[\![0,r-1]\!]}$ and $\mathcal{C}_V = \mathcal{A}^{[\![0,r-1]\!]\times[\![0,r]\!]}$. For each pattern $w \in \mathcal{G}$, we construct a Wang tile t_w with side colours in $\mathcal{C}_H \cup \mathcal{C}_V$: the left (respectively bottom, right, top) colour is given by the $r-1$ leftmost columns (respectively bottommost rows, rightmost columns, topmost rows) of w. The resulting Wang tileset is $T = \{t_w, w \in \mathcal{L}_r(X)\}$.

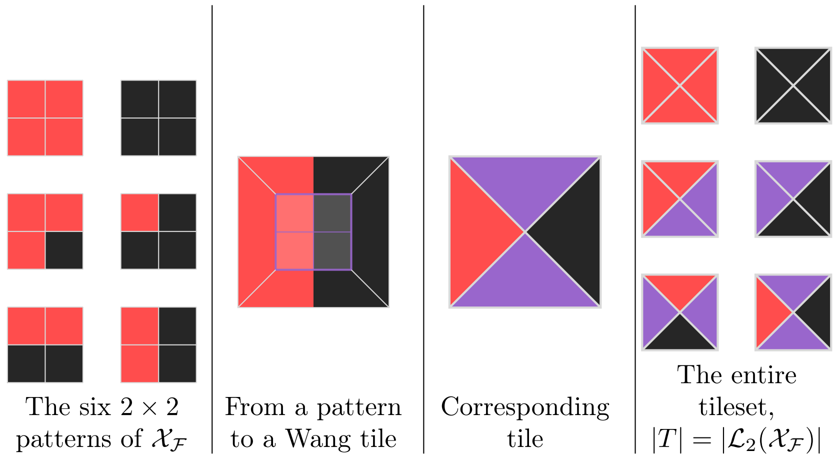

Figure 1.3: Equivalence between the SFT of Figure 1.2 and the Wang tileset of Figure 1.1

As $\mathcal{A}$ is finite, $\mathcal{L}_r(X)$ and therefore T are finite; however, X being an SFT defined by $\mathcal{F}$ implies that $\mathcal{L}_r(X)$ is enough to fully describe the subshift: more precisely, any configuration $x \in \mathcal{A}^{\mathbb{Z}^2}$ such that $\mathcal{L}_r(x) \subseteq \mathcal{L}_r(X)$ does belong to X, as it clearly avoids all the forbidden patterns of $\mathcal{F}$. In particular, any valid configuration $x \in X$ can easily be converted to a T-tiling, and the converse is also true. Intuitively, the tileset T defines an "isomorphic" subshift to X: this is the idea that we try to make precise below. This construction is called the r-**higher-block code** of X, of which we give a precise definition in Section 1.1.2.

As an (a priori) intermediate class of subshifts, we also define the class of **nearest-neighbour** SFTs, or NN-SFT. For any dimension $d > 0$ and $0 \leq i < d$, write $\mathbf{e_i}$ for the i-th basis vector, that is, $\mathbf{e_i} = (\underbrace{0, \dots, 0}_{i}, 1, \underbrace{0, \dots, 0}_{d-i-1})$:

Definition 1.12: Nearest-Neighbour

A $\mathbb{Z}^d$-subshift X is a **nearest-neighbour SFT** if all its forbidden patterns have a support of the form $\{\mathbf{0}, \mathbf{e_i}\}$ for some $0 \leq i < d$.

In particular, the subshift defined by any Wang tileset is a nearest-neighbour SFT. We will see in Section 1.1.4 yet another representation of one-dimensional subshifts of finite type, using graphs.

1.1.2 Factor maps, conjugacy and invariants

In a lot of problems, we do not really care about the concrete set of tilings used to define a subshift X, and as in Figure 1.3 we want to be able to define an equivalence notion on subshifts. In particular, subshifts X and Y should be considered equivalents if one can "recover" any configuration of X starting from configurations of Y, and vice-versa, *via* some invertible operation. For now, we will use definitions motivated by such high-level considerations, and we will see in Section 1.1.3 that they in fact satisfy some general and more abstract properties, and in particular, are far from being *ad hoc*.

Block maps

As explained in Section 1.1.1, subshifts are defined using local and homogeneous conditions and restrictions. It seems natural then to study functions and maps that act on subshifts in a similar way, in order for the image of a subshift to remain a subshift, on a possibly different alphabet. This high-level motivation is enough to propose the following definition, and we will see in Section 1.1.3 why it corresponds to a natural construction coming from more mathematical considerations.

Definition 1.13: Block map

Let $r \geq 0$, $d \geq 1$ and let $\mathcal{A}, \mathcal{B}$ be some alphabets. We call **local map** a function $f\colon \mathcal{A}^{\mathcal{B}_r} \to \mathcal{B}$. The **block map** Φ whose local map is f is defined as

$$\Phi\colon \mathcal{A}^{\mathbb{Z}^d} \to \mathcal{B}^{\mathbb{Z}^d}$$
$$x \mapsto \begin{array}{l} \mathbb{Z}^d \to \mathcal{B} \\ \mathbf{u} \mapsto f(x|_{\mathbf{u}+\mathcal{B}_r}) \end{array}$$

We call r the **radius** of Φ, written radius(Φ).

Block maps are sometimes called **sliding block codes** in the literature. Most of the time, what we really care about is the image of some subshift X under the block map, and in particular we do not need to define the local map $f\colon \mathcal{A}^{\mathcal{B}_r} \to \mathcal{B}$ but can only restrict the definition to be $f\colon \mathcal{L}_{\mathcal{B}_r}(X) \to \mathcal{B}$. We still call block map the associated block map defined only on X rather than on the full shift $\mathcal{A}^{\mathbb{Z}^d}$. Moreover, we will also sometimes abuse the notation and write $\Phi(u)$ for a pattern u. In this case, the image has domain $\{\mathbf{v} \mid \mathbf{v} + \mathcal{B}_r \subset \mathrm{dom}(u)\}$, as those are the only point at which we can apply the local map, and in particular $\mathrm{dom}(\Phi(u)) \subsetneq \mathrm{dom}(u)$ whenever $u \neq \mathbb{Z}^d$. An important case is the case $\Phi\colon \mathcal{A}^{\mathcal{B}_n} \to \mathcal{A}^{\mathcal{B}_{n-r}}$ for $n \geq r$.

Note that block maps defined in this way *a priori* satisfy our "local and homogeneous" conditions: indeed, for any configuration $x \in \mathcal{A}^{\mathbb{Z}^d}$, $\Phi(x)$ is a configuration of $\mathcal{B}^{\mathbb{Z}^d}$, and:

- the value of a cell $\Phi(x)_{\mathbf{u}}$ depends only on its $\mathcal{B}_r$-neighbourhood in x where $r = \mathrm{radius}(\Phi)$.
- two cells with the same $\mathcal{B}_r$-neighbourhood in x will be mapped to the same symbol in $\Phi(x)$.

In fact, this is enough to guarantee the following proposition:

Proposition 1.14

Let $X \subset \mathcal{A}^{\mathbb{Z}^d}$ be a subshift and $\Phi\colon X \to \mathcal{B}^{\mathbb{Z}^d}$ be a block map. Then $\Phi(X)$ is a subshift.

Proof. Let $r = \mathrm{radius}(\Phi)$, and define $L = \bigcup_{n \geq r} \Phi(\mathcal{L}_{\mathcal{B}_n}(X))$. Let $\mathcal{F} = \mathcal{B}^* \setminus L$. We claim that $\Phi(X) = \mathcal{X}_{\mathcal{F}}$, the subshift on $\mathcal{B}$ whose forbidden patterns are the complement of L. To show this, we need to show that no pattern of $\mathcal{F}$ appears in configurations of $\Phi(X)$, and that moreover if a configuration $y \in \mathcal{B}^{\mathbb{Z}^d}$ is such that for all $u \in \mathcal{F}$, $u \not\sqsubseteq y$, then $y \in \Phi(X)$. The former point is obvious by definition of $\mathcal{F}$. Let then $y \in \mathcal{B}^{\mathbb{Z}^d}$ be such that it avoids all patterns of $\mathcal{F}$. Write $v_n = y|_{\mathcal{B}_n}$. By definition of $\mathcal{F}$, for any n we have $v_n \in L$, and in particular there exists $u_n \in \mathcal{L}_{\mathcal{B}_{n+r}}(X)$ such that $\Phi(u_n) = v_n$. We define a configuration x and sequences of patterns $(u_n^{(i)})$ for all n by induction as follows:

- As (u_n) is infinite and $\mathcal{A}$ is finite, there exists $a \in \mathcal{A}$ and subsequence $(u_n^{(0)})$ of (u_n) such that $u_n^{(0)}(\mathbf{0}) = a$ for all n. Let $x_{\mathbf{0}} = a$.
- Assume that $x|_{\mathcal{B}_i}$ and $(u_n^{(i)})_{n\in\mathbb{N}}$ have already been defined, satisfying for all $n \in \mathbb{N}$ $u_n^{(i)}\Big|_{\mathcal{B}_i} = x|_{\mathcal{B}_i}$. Then, there exists a subsequence $(u_n^{(i+1)})_{n\in\mathbb{N}}$ of elements which coincide on $\mathcal{B}_{i+1}$. We can then define $x|_{\mathcal{B}_{i+1}\setminus\mathcal{B}_i} = u_0^{(i+1)}\Big|_{\mathcal{B}_{i+1}\setminus\mathcal{B}_i}$, and so we get a definition of $x|_{\mathcal{B}_{i+1}}$.

As $\bigcup_{\mathcal{B}_i} = \mathbb{Z}^2$, we eventually define x on the entire plane. Moreover, for $i \geq 0$ we have by construction $x|_{\mathcal{B}_{i+r}} = u_0^{(i+r)}\Big|_{\mathcal{B}_{i+r}}$. In particular, $\Phi(x)|_{\mathcal{B}_i} = v_i$, and so $\Phi(x) = y$. □

Note that we cannot deduce from the previous proof that if X is an SFT, then its image $\Phi(X)$ is an SFT itself. It The consequences of this remark will explored be in Section 1.1.4.

Definition 1.15: Embedding, factor map

Let $\Phi\colon X \to Y$ be a block map. If Φ is injective, we say that it is an **embedding** of X in Y. If Φ is surjective, we say that Y is a **factor** of X, and that X factors onto Y – in that case, Φ will also be called a **factor map**.

The previous remark can then be reformulated as the fact that factors of SFTs need not be SFTs.

We give an additional easy property on block maps, directly following from the definitions, which we informally refer to when saying that “block maps commute with the shifts”.

Proposition 1.16

Let $\Phi\colon X \to Y$ be a block map. Then, $\Phi \circ \sigma_X = \sigma_Y \circ \Phi$, where σ_X, σ_Y are the respective shift functions on X and Y.

Conjugacy and conjugacy invariants

A somewhat surprising fact is that inverses of block maps are also block maps:

Proposition 1.17

Let $\Phi\colon X \to Y$ be a bijective block map. Then its inverse $\Phi^{-1}\colon Y \to X$ is a block map.

We could prove this result by hand, but we will show it in a more elegant manner in Section 1.1.3 using a topological characterization of subshifts and block maps. For now, we simply observe that block maps behave particularly well with subshifts, in the sense that they map subshifts to subshifts, and are stable under a lot of operations, most notably composition and inverse when it makes sense.

Definition 1.18: Conjugacy

Let $\Phi\colon X \to Y$ be a block map. If it is bijective, we say that it is a **conjugacy**, and X and Y are **conjugate** subshifts.

Conjugacy is the "correct" notion of isomorphism for subshift, and captures the ideas already hinted at in Section 1.1.1 of being able to "recode" or recover configurations of one subshift from the other. In particular, the construction of higher-block codes of Section 1.1.1 can now be formalized:

Definition 1.19: Higher-block code

Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be any subshift, and $r > 0$. We call r-**higher-block map** of X the map $\Phi_r\colon X \to \mathcal{L}_{\mathcal{B}_r}(X)^{\mathbb{Z}^d}$ whose local map of radius r is

$$\begin{aligned} f\colon \mathcal{A}^{\mathcal{B}_r} &\to \mathcal{A}^{\mathcal{B}_r} \\ w &\mapsto w \end{aligned}$$

the domain being viewed as a set of patterns, the co-domain being seen as an alphabet, *i.e.* a set of symbols.
We call r-**higher-block code** of X the subshift $\Phi_r(X) \subseteq \mathcal{L}_{\mathcal{B}_r}(X)^{\mathbb{Z}^d}$.

It is pretty clear that Φ_r is always a conjugacy, whose inverse has radius 0 and a local map given by $g(w) = w_{\mathbf{0}}$.

The proof that was sketched in Section 1.1.1 can be formalized to show the following result, stated in our new terminology:

Theorem 1.20

Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be any subshift. The following are equivalent:

- X is an SFT.
- X is conjugate to a nearest-neighbour SFT.
- X is conjugate to a Wang tiling space.

Proof. If $X = \mathcal{X}_{\mathcal{F}}$, simply let $r = \max_{u\in\mathcal{F}} \min_k \operatorname{dom}(u) \subseteq \mathcal{Q}_k$. Then X is conjugate to its r-higher-block code, which is easily seen to be a nearest neighbour SFT. □

Now, having a notion of isomorphism, we might wonder which kind of properties are stable under isomorphism.

Definition 1.21: Conjugacy invariant

A **conjugacy invariant** is a mathematical object $f(X)$ associated to a subshift X, such that $f(X) = f(Y)$ when X and Y are conjugate.

This definition is necessarily quite vague, as conjugacy invariant can by nature be pretty much any mathematical object. For example, in this thesis, we will explore an invariant called the extender entropy in Chapter 2, which is a real number, and another invariant called the projective fundamental group in Chapter 3, which is a group. Generally, we do not have the reciprocal, that is, if $f(X) = f(Y)$ we don't necessarily have that X and Y are conjugate, even when one restricts to small classes of subshifts.

Let us now give a few examples of examples of properties and objects that are invariant by conjugacy:

Proposition 1.22

Being a subshift of finite type is invariant by conjugacy.

Proof. Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}, Y \subseteq \mathcal{B}^{\mathbb{Z}^d}$ be two subshifts, and $\Phi\colon X \to Y$ be a conjugacy with local map ϕ. A minor difficulty is that ϕ is not defined outside of the language of X, and so we cannot really consider patterns u of X such that $\Phi(u) \notin \mathcal{L}(Y)$, as Φ might simply not be defined on u. Suppose that Y is of finite type. Define, for $n \geq 0$ and any subshift Z, the set $\mathcal{F}_n(Z) = \mathcal{B}^{\mathcal{Q}_n} \setminus \mathcal{L}_n(Z)$. As Y is an SFT, there exists N such that $Y = \mathcal{X}_{\mathcal{F}_N(Y)}$. Let $r = \text{radius}(\Phi)$, and let then $x \in \mathcal{X}_{\mathcal{F}_{N+2r}(X)} \subseteq \mathcal{A}^{\mathbb{Z}^d}$. As any pattern u of support $\mathcal{B}_r$ in x contains only subpatterns of size $N + 2r$ that are valid in X, we can compute $\Phi(x)$, by applying ϕ everywhere. We obtain a configuration $y \in \mathcal{X}_{\mathcal{F}_N(Y)} = Y$, and so $x \in X$. As Φ is bijective, we get that $X = \mathcal{X}_{\mathcal{F}_{N+2r}(X)}$ and in particular it is an SFT. □

In dimension 1, Proposition 1.22 is much simpler to prove, as any sufficiently large locally admissible pattern is globally admissible in an SFT. We will see in Section 1.2.3 that the situation is very different for higher-dimensional SFTs.

We now present another invariant, which will inspire our definitions in Chapter 2. This is a special case of an invariant defined for any dynamical system, but which is especially easy to define in the case of subshifts:

Definition 1.23: Entropy

Let X be any subshift on $\mathbb{Z}^d$. The **entropy** of X is the real

$$h(X) = \inf_{n\to+\infty} \frac{\log|\mathcal{L}_n(X)|}{|\mathcal{Q}_n|} = \lim_{n\to+\infty} \frac{\log|\mathcal{L}_n(X)|}{n^d}$$

The fact that the limit exists is not obvious, and is a consequence of a classical lemma, known as the subadditivity (or Fekete's) lemma. We will prove in Proposition 2.20 that a variant of entropy, namely, the extender entropy, is well-defined, using this technique. Entropy measures how many globally admissible patterns exist in X, and can be interpreted as follows: $h(X)$ is the average amount of information (in bits) that you learn about a *configuration* x whenever a random cell of x is revealed to you – this is from the point of view of information theory: of course, the two-points subshift $\{0^{\mathbb{Z}^d}, 1^{\mathbb{Z}^d}\}$ has zero entropy, but knowing the value of a single cell reveals the entire configuration: this is consistent with our informal description, as every cell after the first one is completely determined, and far from "unexpected".

Proposition 1.24

If X and Y are conjugate subshifts then $h(X) = h(Y)$.

Proof. Let $\Phi\colon X \to Y$ be a conjugacy, and assume without loss of generality that $\text{radius}(\Phi) = \text{radius}(\Phi^{-1}) = r$. Then for $n \geq 2r$,

$$|\mathcal{L}_{n-2r}(X) \leq |\mathcal{L}_n(Y)| \leq |\mathcal{L}_{n+2r}(X)||$$

Dividing everything by n^d, we can rewrite those inequalities as

$$\underbrace{\frac{|\mathcal{L}_{n-2r}(X)|}{(n-2r)^d}}_{\to h(X)} \underbrace{\frac{(n-2r)^d}{n^d}}_{\to 1} \leq \underbrace{\frac{|\mathcal{L}_n(Y)|}{n^d}}_{\to h(Y)} \leq \underbrace{\frac{|\mathcal{L}_{n+2r}(X)|}{(n+2r)^d}}_{\to h(X)} \underbrace{\frac{(n+2r)^d}{n^d}}_{\to 1}$$

and we get $h(X) \leq h(Y) \leq h(X)$ and so $h(X) = h(Y)$. □

Finally, and in order to show that conjugacy invariants can be very diverse mathematical objects, we introduce an algebraic invariant:

Definition 1.25: Automorphism group

Let X be a subshift, and G be the set of bijective block maps $X \to X$. Then $\mathfrak{S}_X = (G, \circ)$ is a group, called the **automorphism group** of X.

Proposition 1.26

Let $\Phi\colon X \to Y$ be a conjugacy. Then $\mathfrak{S}_X \simeq \mathfrak{S}_Y$.

Proof.

$$\begin{aligned} \Psi\colon \mathfrak{S}_X &\to \mathfrak{S}_Y \\ f &\mapsto \Phi \circ f \circ \Phi^{-1} \end{aligned}$$

is an isomorphism between $\mathfrak{S}_X$ and $\mathfrak{S}_Y$. □

This group has been extensively studied, but many important questions remain open. For example, writing X_2 and X_3 the full-shifts on respectively 2 and 3 symbols, we do not know if $\mathfrak{S}_{X_2}$ and $\mathfrak{S}_{X_3}$ are isomorphic, although they embed into each other. There are also deep links between this automorphism group and dynamical properties (presented in Section 1.1.3), see for example [Hoc09a; Sal15; Sal16], and with the complexity of the subshift (that is, the sequence $(|\mathcal{L}_n(X)|)_{n\in\mathbb{N}}$), see *e.g.* [Don+15; CK15a; CK15b; Don+17]. We will not study this group in the thesis, but another algebraic invariant in Chapter 3.

1.1.3 Topological aspects and some operations

We now turn our attention to a second point of view about subshifts, which was hinted at in Section 1.1.2. Although most definitions given until now are combinatorial in nature, and explain why subshifts lend themselves especially well to analysis using tools from computability theory, they are only a specific subclass of **dynamical systems**. This second point of view allows us to use powerful results and techniques of the theory of dynamical systems, or point set topology, to study general properties of subshifts.

Dynamical systems

In general, a dynamical system can be seen as a space on which we have an action that we wish to apply repeatedly: that is, given X a space and $f\colon X \to X$, we try to understand the **orbits** of the points $(x, f(x), f^2(x)\dots)$, or $(\dots, f^{-1}(x), x, f(x), f^2(x)\dots)$ in the case of a reversible map f. In general, studying orbits is a difficult problem, and many other – hopefully easier – questions naturally arise:

- Are there stable subspaces, that is, $Y \subset X$ such that $f(Y) \subseteq Y$?
- Does some orbit, or even every orbit, pass arbitrarily close to every point ?
- Do close points have "close" orbits ? Otherwise, how unpredictable or chaotic are the orbits ?
- What about the same questions, asked for the systems (X, f^n) for $n > 1$ instead ?

We will see how some of these questions can be translated in combinatorial terms when studying subshifts in particular. The specific study of subshifts and related objects constitutes the field of symbolic dynamics, and we refer to [Kur03] for a complete introduction.

Definition 1.27: Dynamical system

A (discrete-time) dynamical system is a pair (X, f) where X is a space, and $f\colon X \to X$ is a function.

Generally, and in order to diminish the complexity of the system, we consider a compact (topological) space X and a continuous action f. As we will now see, those conditions are satisfied with subshifts, and a subshift X with the shift action σ can be viewed as a dynamical system. One encounters a minor subtlety when considering higher-dimensional subshifts in this general framework of dynamical systems: indeed, over $\mathbb{Z}$-subshifts, we have $\sigma_{\mathbf{k}} = \sigma_{\mathbf{1}}^k$ for any $k \in \mathbb{Z}$, and so (X, σ_1) satisfies the definition given in Definition 1.27 (of course, we still need to define a topology on X). On the other hand, we do not have a *single* action for $\mathbb{Z}^d$ subshifts with $d > 1$. We will briefly mention in Section 1.1.3 how some classical notions of the study of dynamical systems can be adapted in this case. In general though, this is only a minor difficulty, as most definitions can straightforwardly be adapted to accommodate the existence of multiple actions $(\sigma_{\mathbf{u}})_{\mathbf{u}\in\mathbb{Z}^d}$.

Cylinders, compactness, continuity

We first give a definition of a distance on $\mathcal{A}^{\mathbb{Z}^d}$, which in turn induces a topology that we describe below. For a point $\mathbf{u} = (u_1, \dots u_d) \in \mathbb{Z}^d$, we write $\|\mathbf{u}\|_\infty = \max_{1\leq i\leq d} |u_i|$. In particular, $\mathcal{B}_r$ is the ball of radius r for this norm.

Definition 1.28: Distance - subshifts

Let $\mathcal{A}$ be a finite alphabet, and define:

$$d\colon \mathcal{A}^{\mathbb{Z}^d} \times \mathcal{A}^{\mathbb{Z}^d}$$
$$(x, y) \mapsto 2^{-\inf_{\mathbf{u}\in\mathbb{Z}^d}\{\|\mathbf{u}\|_\infty, x_\mathbf{u}\neq y_\mathbf{u}\}}$$

Then d is a distance on $\mathcal{A}^{\mathbb{Z}^2}$.

An intuitive way to understand this distance is the following: configurations are close if they agree on a large central ball. Other distances, or pseudo-distances, have been considered in the literature, see the recent thesis [Ben23] for a survey. A basis of open sets is the set of **cylinders**:

Definition 1.29: Cylinders

For any pattern $u \in \mathcal{A}^{\mathbb{Z}^d}$, we call **cylinder** and write $[u]$ the set

$$[u] = \{x \in \mathcal{A}^{\mathbb{Z}^d}, x|_{\text{dom}(u)} = u\}$$

We sometimes abuse the notation and also write $[u]$ for $[u] \cap X$ for some subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ when X is clear from the context.

Depending on the familiarity of the reader with metric spaces or the more abstract topological spaces, one can use either point of view to think about the topology of subshifts thanks to the following proposition:

Proposition 1.30

The topology induced by the distance d has the family $([u])_{u\in\mathcal{A}^*}$ as a basis of open sets. It is the product topology of the discrete topology on each of the spaces $\mathcal{A}$.

As any subshift is a subset of some $\mathcal{A}^{\mathbb{Z}^d}$, both the distance and the topology on X itself are obtained as the induced distance and topology. From now on, we will always consider the space $\mathcal{A}^{\mathbb{Z}^d}$ with this topology and distance for any finite alphabet $\mathcal{A}$.

We can now give a completely topological characterization of subshifts, equivalent to Definition 1.4. This in turn will imply that block maps (and therefore, conjugacy) have nice topological descriptions, making some results easier to prove.

Proposition 1.31

For any finite alphabet $\mathcal{A}$, the space $\mathcal{A}^{\mathbb{Z}^d}$ is compact.

Proof. Let $X = \mathcal{A}^{\mathbb{Z}^d}$. As each $\mathcal{A}$ is finite and so obviously compact, an abstract proof is given by Tychonoff's theorem, stating that a product of compact spaces is compact, and so X is compact. A more concrete proof can be derived from the proof of Proposition 1.14. Once again, we can simply use the "metric" point of view, and show that any sequence of configurations $(x_n) \in X^{\mathbb{N}}$ contains a converging subsequence. Using the same diagonal argument than in Proposition 1.14, we can construct by successive extractions sequences (x^i_n) and a configuration $x \in \mathcal{A}^{\mathbb{Z}^2}$, with $x|_{\mathcal{B}_n} = x^n_k|_{\mathcal{B}_n}$ for all $k \geq 0$. □

Proposition 1.32

For a finite alphabet $\mathcal{A}$, the subshifts of $\mathcal{A}^{\mathbb{Z}^d}$ are exactly the closed, shift-invariant subsets of $\mathcal{A}^{\mathbb{Z}^d}$.

Proof. Let $\mathcal{F}$ be a family of finite forbidden d-dimensional patterns on some alphabet $\mathcal{A}$, and let $X = \mathcal{X}_{\mathcal{F}}$. Then X is shift-invariant by Proposition 1.9. Then, let $(x_n) \in X^{\mathbb{N}}$ be a converging sequence of configurations of X, $x_n \to x$. We show that $x \in X$. Up to some extraction, we can assume that for any $n \in \mathbb{N}$ we have $x|_{\mathcal{B}_n} = x_n|_{\mathcal{B}_n}$. In particular, $x|_{\mathcal{B}_n}$ is locally admissible, that is, it does not contain patterns from $\mathcal{F}$. Therefore, x itself does not contain any forbidden pattern, and so $x \in X$, hence X is closed.

Let X be closed and shift-invariant in $\mathcal{A}^{\mathbb{Z}^d}$. We claim that $X = \mathcal{X}_{\mathcal{A}^{\mathbb{Z}^d} \setminus \mathcal{L}(X)}$. The inclusion $X \subseteq \mathcal{X}_{\mathcal{A}^* \setminus \mathcal{L}(X)}$ is clear. For the other inclusion, let $x \in \mathcal{X}_{\mathcal{A}^* \setminus \mathcal{L}(X)}$. By definition, for all $n \in \mathbb{N}$, $x|_{\mathcal{B}_n} \in \mathcal{L}(X)$, and so there exists $x_n \in X$ with $x|_{\mathcal{B}_n} = x_n|_{\mathcal{B}_n}$, and moreover $x_n \to x$ (in $\mathcal{A}^{\mathbb{Z}^d}$). As X is closed, we in fact get $x \in X$. □

Corollary 1.33: Compactness

Subshifts are compact spaces.

Proof. $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ is closed in a compact space, so it is compact. □

One of the consequences is that a sequence of larger and larger *locally admissible* patterns $(u_n)_{n\in\mathbb{N}}$ for a subshift X, for example with $\mathrm{dom}(u_n) = \mathcal{B}_n$, admits a converging subsequence to some $x \in X$. In particular, if $u_n \sqsubseteq u_{n+1}$ for all n, then $u_n \to x \in X$.

The following theorem, attributed to Curtis, Hedlund and Lyndon by Hedlund himself in [Hed69], explains why we argued that the definition of block maps presented in Section 1.1.2 was indeed the correct one:

Theorem 1.34: Curtis-Hedlund-Lyndon [Hed69, Thm. 3.4]

Let $(X, \sigma_X), (Y, \sigma_Y)$ be subshifts on $\mathbb{Z}^d$. Then, the block maps $\Phi\colon X \to Y$ are exactly the continuous maps Φ which satisfy for all $\mathbf{u} \in \mathbb{Z}^d$ that $\Phi \circ \sigma_{X,\mathbf{u}} = \sigma_{Y,\mathbf{u}} \circ \Phi$.

Proof. Let $\mathcal{A}$ and $\mathcal{B}$ be the respective alphabets of X, Y. The fact that block maps commute with the shift functions is Proposition 1.16. As block maps are defined using local functions, they are also clearly continuous: this is easily seen using the "metric" definition of continuity, as for a block map Φ and any $n > 0$, $x, x' \in X$, $d(x, x') \geq 2^{-n}$ implies that $d(\Phi(x), \Phi(x')) \leq 2^{-n+\mathrm{radius}(\Phi)}$ – in fact, we obtain that block maps are Lipschitz-continuous, with a Lipschitz constant depending on their radius, which is a much stronger property.

For the other direction, let Φ be a shift-commuting continuous function. As X is compact, any continuous function is uniformly continuous by Heine theorem. As X is furthermore a metric space, there exists $r \geq 0$ such that for any $x, x' \in X$, $d(x, x') \leq$

$2^{-r} \implies d(\Phi(x), \Phi(x')) \leq \frac{1}{2}$, that is, $\Phi(x)_{\mathbf{0}} = \Phi(y)_{\mathbf{0}}$. We can therefore define the map:

$$\begin{aligned} \phi\colon \mathcal{L}_r(X) &\to \mathcal{B} \\ u &\mapsto \Phi(u)_{\mathbf{0}} \end{aligned}$$

where $\Phi(u)_{\mathbf{0}}$ is the value of $\Phi(x)_{\mathbf{0}}$ for x in $[u]$, which is well-defined by the above remark. The fact that Φ is shift-commuting means that Φ is the block map whose local function is ϕ. □

This theorem has been generalized to spaces other than subshifts on $\mathbb{Z}^d$, but the proofs are always very similar, see for example [CC10].

We are now ready to prove Proposition 1.17 using tools from topology rather than constructing inverse block maps "by hand":

Proof of Proposition 1.17. $\Phi\colon X \to Y$ is continuous and defined on a compact space, so Φ^{-1} is also continuous. It is easy to check that Φ^{-1} also commutes with all the shift functions on Y, and so by Theorem 1.34, it is a block map. □

Dynamical and mixing properties

As dynamical systems, subshifts can be studied from this point of view, and we can try to understand how classical definitions and properties from the theory of dynamical systems can be reformulated, or better understood, in the case of subshifts. The difficulty of (X, σ) not being an actual dynamical system mentioned in Section 1.1.3, at least when using Definition 1.27, will have minor consequences on some definitions, most importantly what it means for a subshift to be **mixing**. In particular, this explains why some results of Section 2.3.2 will differentiate between the one-dimensional and the higher-dimensional case. We give a few examples of properties that will be used later in the thesis, mainly in Section 2.4 and Section 3.3. Some other but lesser-known and sometimes more technical properties will also be introduced directly in those chapters whenever needed.

Periodicity A first example of dynamical property is **periodicity**. We give the higher-dimensional definition directly, and mention an interesting detail about this definition.

Definition 1.35: Periodic subshift

Let $x \in \mathcal{A}^{\mathbb{Z}^d}$ be a configuration. We say that x is **periodic** if there exists $n_0, \dots, n_{d-1} > 0$ such that for all $0 \leq i < d$, $\sigma_{n_i \mathbf{e_i}}(x) = x$.
A subshift X is periodic if all its configurations are periodic. If X has no periodic point, then it is **aperiodic**.

Note that aperiodicity is not the negation of periodicity: a subshift containing periodic and non-periodic configurations is neither periodic not aperiodic. An interesting result, which is a consequence of compactness, is the following:

Theorem 1.36 [BDJ08, Thm. 3.8]

If X is periodic, then it is finite.

Note that this theorem is interesting only when we use the definition of "periodic subshift" given in Definition 1.35. The traditional definition is obtained here as a corollary, showing that both are in fact equivalent:

Corollary 1.37

If $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ is a periodic subshift, there exists $(n_i)_{0 \leq i < d}$ such that for any $x \in X$ and $0 \leq i < d$, $\sigma_{n_i \mathbf{e_i}}(x) = x$.

In dimensions $d \geq 2$, other weaker notions of periodicity exist, and we might simply require that for *some* $\mathbf{e_i}$ there exists n_i such that $\sigma_{n_i \mathbf{e_i}}(x) = x$. In this thesis, whenever we say that a subshift is periodic, we mean that it satisfies Definition 1.35, which is classically known as **strong periodicity** in the literature.

Minimality Another strong condition, which admits a nice characterization in the special case of subshifts, is minimality. In full generality, the usual definition – which also explains the name of *minimal* subshifts – is the following:

Definition 1.38: Minimal subshift

A subshift X is **minimal** if it contains no proper non-empty subshift.

This is a general definition from the theory of dynamical systems, which even in the general case admits several other equivalent definitions. In the case of subshifts, this can be reformulated in a very simple condition:

Proposition 1.39

Let X be any subshift, defined by some family of forbidden patterns $\mathcal{F}$. Then X is minimal if and only for all $u \in \mathcal{L}(X)$, $\mathcal{X}_{\mathcal{F} \cup \{u\}} = \emptyset$. Equivalently, for all $x \in X$, $u \sqsubseteq x$.

Proof. If $u \in \mathcal{L}(X)$ then $\mathcal{X}_{\mathcal{F} \cup \{u\}}$ is a proper subshift of X. □

We will see in Section 2.4.1 that minimality can prevent some otherwise complicated combinatorial properties from appearing in multidimensional subshifts.

Transitivity and mixing properties We give here some of the most important "mixing notions" studied in dynamical systems in general, and tiling spaces in particular. By "mixing notion", we mean any property which quantifies the way arbitrarily large patterns can be "glued" with each other, that is, to what extent can we find configurations containing both $u, v \in \mathcal{L}_n(X)$ for any pair of patterns u, v as n grows ? A rather weak notion is **transitivity**:

Definition 1.40: Transitivity

A subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ is **transitive** if for any pair of patterns $v, w \in \mathcal{L}(X)$, there exists $x \in X$ such that $v \sqsubseteq x$ and $w \sqsubseteq x$.

A formulation that makes the link clearer with some other, more quantitative mixing notions can easily be obtained:

Proposition 1.41

A subshift X is transitive if and only if

$$\forall n > 0, \exists \mathbf{u} \in \mathbb{Z}^d, \forall v, w \in \mathcal{L}_n(X), \exists x \in X, x\big|_{\mathcal{B}_N} = v, x\big|_{\mathcal{B}_N+\mathbf{u}} = w$$

A stronger property is the one of weakly mixing dynamical systems.

Definition 1.42: Weak mixing

A dynamical system (X, T) is **weakly mixing** if $(X \times X, T \times T)$ is transitive.

Finally, the last definition which we will discuss in Section 2.4.2 and Section 3.3.2 is the **mixing** property. We give a definition which is only valid for one-dimensional subshifts:

Definition 1.43: Mixing

A subshift $X \subseteq \mathcal{A}^{\mathbb{Z}}$ is **mixing** if for $n > 0$, there exists $N > 0$ such that for any $z \in \mathbb{Z}$ with $|z| \geq N$ and any pair of patterns $v, w \in \mathcal{L}_n(X)$, there exists $x \in X$ such that $x|_{\mathcal{B}_N} = v, x|_{\mathcal{B}_N+z} = w$.

In the case of subshifts, we try to give a high-level overview of the difference between being transitive, weakly mixing, and mixing:

- Transitive means that any two patterns can be glued, without any control over "how" (their relative position, their distance ...)
- Weakly mixing means that we still do not control anything on how we can glue an arbitrary pair of patterns, but for any two pairs of patterns (P_1, P_2) and (Q_1, Q_2), we can glue P_1 and P_2 on the one hand, Q_1 and Q_2 on the other hand, in the same relative positions.
- Mixing means that any two patterns can be glued in any relative position, provided that they are placed sufficiently far from one another.

In the case of higher-dimensional subshifts, being mixing is often too strong of a condition. Indeed, one has to deal with patterns of complex shapes, such as rings, and more generally patterns with non-rectangular support, which can then be at large distance but in complex geometric configurations (for example, sequences of concentric rings that "belong" alternatively to v or w). We often choose to study weaker conditions, specifying only how rectangular blocks can be glued together, and avoid imposing anything on general patterns – see Section 3.4 for such a notion.

In particular, we have the easy chain of implications:

Proposition 1.44

Mixing $\implies$ Weakly mixing $\implies$ Transitive.
Those implications are strict.

1.1.4 Walks on graphs, regular languages, sofic subshifts

Multidimensional sofic subshifts

The main class of subshifts that we presented up to this point was the class of SFTs. However, as already mentioned about the proof of Proposition 1.14, the class of SFTs is not stable by factor map:

Definition 1.45: Sofic subshift

Let $Y \subseteq \mathcal{B}^{\mathbb{Z}^d}$ be a subshift. We say that X is **sofic** if there exists an SFT $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ and a factor map $\Phi\colon X \to Y$. In that case, X is an **SFT extension** of Y.

More succinctly, sofic shifts are factors of SFTs. We will often use the next proposition to make some simplifying assumptions on sofic subshifts:

Proposition 1.46

Let $Y \subseteq \mathbb{Z}^d$ be a sofic subshift. There exists a nearest-neighbour SFT X and a factor map Φ of radius 0 such that $Y = \Phi(X)$.

Proof. Let Z be the $\mathbb{Z}^d$ SFT and $\Psi\colon Z \to Y$ be the factor map given by Definition 1.45. By Theorem 1.20, Z is conjugate *via* $\Theta\colon Z \to Z'$ to a nearest-neighbour SFT Z'. Let $r = \text{radius}(\Psi \circ \Theta^{-1})$, and let $X = \Phi_r(Z')$ be the r-higher-block code of Z' (see Definition 1.19). Then X is a nearest-neighbour SFT, and $\Phi = \Psi \circ \Theta^{-1} \circ \Phi_r^{-1}$ is a factor map $X \to Y$ of radius 0.

$$X \xleftarrow{\ \Phi_r\ } Z' \xleftarrow{\ \Theta\ } Z \xrightarrow{\ \Psi\ } Y$$

□

In general, in later constructions, we will often view sofic subshifts as being “SFTs up to construction lines”. The typical example will be a sofic subshift $Y \subseteq \mathcal{B}^{\mathbb{Z}^d}$ obtained as the natural projection of an SFT $X \subseteq (\mathcal{A} \times \mathcal{B})^{\mathbb{Z}^d} \simeq \mathcal{A}^{\mathbb{Z}^d} \times \mathcal{B}^{\mathbb{Z}^d}$. We will refer to this situation as X being a subshift with two (or more) **layers**, the sofic subshift Y being then the subshift consisting only of one (or more) of X's layers.

Example 1 (Sunny-side-up)**.** *We define the **sunny-side-up** subshift in dimension d as the subshift of $\{0,1\}^{\mathbb{Z}^d}$ defined by*

$$Y_d = \{y \in \{0,1\}^{\mathbb{Z}^d}, \sum_{\mathbf{u} \in \mathbb{Z}^d} y_i \leq 1\}$$

We show that $Y = Y_1$, the one-dimensional sunny-side-up, is a sofic subshift. For a visual illustration of what we mean by layers, let us represent Y as a subshift on the alphabet $\mathcal{B} = \{□, □\}$, Y being the set of configurations containing at most 1 yellow square.

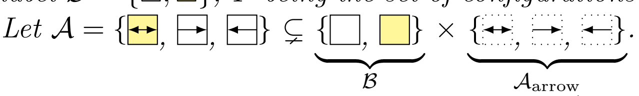

We define a nearest-neighbour SFT X on $\mathcal{A}$ by defining the allowed *patterns* $\mathcal{F} =$ 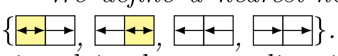*. More informally, the arrows on the $\mathcal{A}_{\text{arrow}}$ layer must be continued in the same direction. Define then the block-map $\Phi\colon X \to Y$ of radius 0 which* *simply forgets the $\mathcal{A}_{\text{arrow}}$ layer, projecting only on the $\mathcal{B}$ layer of $\mathcal{A}$. Then Φ is a factor map, and so Y is sofic as the image of the nearest neighbour SFT X by a block map.*

The proof that Y_d is sofic for $d > 1$ is presented in Section 3.2.3, but follows the same ideas of using additional colours and layers to identify the only point that can be mapped to the unique 1 in configurations of Y_d.

Notation. For any product $A = \prod_{i=1}^n A_i$, we will denote $\pi_{A_j} \colon A \to A_j$ the natural projection. We will sometimes use this notation in the following ways:

- $\pi_{\prod_{j\in J\subset [\![1,n]\!]} A_j} = \prod_{j\in J} \pi_{A_j}$ is a map $\prod_{i=1}^n A_i \to \prod_{j\in J} A_j$.
- If $X \subseteq (\prod_{i=1}^n A_i)^{\mathbb{Z}^d}$ is a subshift with several layers, we write $\pi_{A_j}(X) \subseteq A_j^{\mathbb{Z}^d}$ the subshift obtained by applying $\pi_{A_i} \colon \prod_{i=1}^n \to A_j$ point-wise on each configuration.

Proposition 1.47

The class of sofic shifts strictly contains the class of SFTs (in any dimension).

Proof. Every SFT is clearly sofic, as for example they are the images of themselves by the identity map. We show that the sunny-side-up in dimension d is sofic but not an SFT. The fact that it is sofic is proven in Example 1 for the case $d = 1$, and Section 3.2.3 for the case $d = 2$ which easily generalizes to higher dimensions. Suppose that $Y_d = \mathcal{X}_\mathcal{F}$ is an SFT, defined by some finite family of forbidden patterns $\mathcal{F}$. Let then k be such that $\mathcal{F} \subseteq \{0,1\}^{\mathcal{Q}_k}$. Let $x \in \{0,1\}^{\mathbb{Z}^d}$ be the configuration defined by $x_\mathbf{u} = 1 \iff \mathbf{u} \in \{\mathbf{0}, (k+1)\mathbf{e_0}\}$. As x contains two symbols "1", we have $x \notin Y_d$. However, as those "1" are at distance $k+1$, we also have that for any $\mathbf{u} \in \mathbb{Z}^d$, $x|_{\mathbf{u}+\mathcal{Q}_k} \notin \mathcal{F}$. Indeed, every subpattern of x of size k is globally admissible in Y_d, hence locally admissible. By definition, this means that $x \in \mathcal{X}_\mathcal{F}$, that is $x \in Y_d$. This is a contradiction. □

In Chapter 2, we will present some other tools and techniques to distinguish SFTs from sofic subshifts, and even sofic subshifts from non-sofic subshifts.

Interlude: some graph theory

A convenient way to represent SFTs and sofic subshifts in dimension one is to use graphs. We briefly present here the kind of graphs we consider, bearing in mind that graphs are used at several places in the thesis (most importantly in Section 3.4 and Chapter 4) with slightly different assumptions, that will therefore be highlighted when needed. In particular, some additional definitions will be given in Section 4.3.1, and that we choose not to mention for now as they are only needed in this chapter.

Definition 1.48: Graphs

A **graph** G is a pair $(V(G), E(G))$ of **vertices** $V(G)$ and **edges** $E(G) \subseteq V(G) \times V(G)$. If for all $(u,v) \in E(G)$, we have $(v,u) \in E(G)$, then G is said to be **undirected**, and **directed** otherwise.

In the case of undirected graphs, we will sometimes abuse notation and terminology, and talk about "the" edge $e = (u,v)$ between u and v for both (u,v) and $(v,u) \in E(G)$.

Notation. For a graph $G = (V, E)$, and an edge $e \in E$, we write $s(e), t(e) \in V$ the vertices such that $e = (s(e), t(e))$. Those vertices are respectively called the starting and ending vertices of e. If it exists, the edge $(t(e), s(e))$ is written e^{-1}, and the map $e \in E \to e^{-1}$ is then an involution.

It is sometimes convenient to have some additional information carried by the graph:

Definition 1.49: Labeled graph

Let $G = (V, E)$ be a graph. A **vertex-labeling** function is a map $\lambda_V : V \to \mathcal{C}$ into some set $\mathcal{C}$; An **edge-labeling** function is a map $\lambda_E : E \to \mathcal{D}$ into some set $\mathcal{D}$. In this case, $(V, E, \lambda_V), (V, E, \lambda_E)$ and $(V, E, \lambda_V, \lambda_E)$ are respectively called vertex-labeled, edge-labeled and labeled graphs.

Finally, we need a last set of definitions:

Definition 1.50: Paths, walks, cycles

Let $G = (V, E)$ be a graph. A **path**, or **walk**, in G, is a sequence $v_1, \ldots, v_n$ of vertices such that for $1 \leq i \leq n-1$, $(v_i, v_{i+1}) \in E$. It is a **cycle** if $v_1 = v_n$, is simple if $v_i \neq v_j$ for $i \neq j$, and is non-backtracking if $v_i \neq v_{i+2}$ for $1 \leq i \leq n-2$.

Unless specified otherwise, all the graphs considered in Section 1.1.4 are finite, *i.e.* have finite sets of vertices and edges.

The one-dimensional case

In dimension 1, there exist links between graphs and subshifts. A first easy observation, following from Theorem 1.20, shows that we can use directed graphs to represent $\mathbb{Z}$ subshifts of finite type:

Proposition 1.51

Let $G = (V, E)$ be a directed graph. Then, the set of bi-infinite walks on G is a nearest-neighbour SFT on the alphabet V.

Proof. The bi-infinite walks on G are exactly the $\mathbb{Z}$-indexed sequences x of elements of V such that for all $i \in \mathbb{Z}$, $(x_i, x_{i+1}) \in E$. In other words, this set of walks is exactly $\mathcal{X}_{V^2 \setminus E}$. □

However, the interesting direction is the other one, which shows that many examples of algorithmic problems on SFTs can be solved using general techniques of graph theory. Let us write Walks(G) for the set of bi-infinite walks on G, so that Proposition 1.51 states that Walks(G) is a NN-SFT.:

Proposition 1.52

Let X be an SFT. Then, there exists a graph $G = (V, E)$ such that X is conjugate to $\mathrm{Walks}(G)$.

Proof. By Theorem 1.20, X is conjugate to a nearest-neighbour SFT $Y = \mathcal{X}_{\mathcal{F}} \subseteq \mathcal{B}^{\mathbb{Z}}$, with $\mathcal{F} \in \mathcal{B}^2$. Define $G = (\mathcal{B}, \mathcal{B}^2 \setminus \mathcal{F})$. Then $X \simeq Y = \mathrm{Walks}(G)$. □

This point of view is very useful, as we can use the structure of directed graphs to prove general claims about $\mathbb{Z}$ SFTs. An important example concerns periodic points, which as we will see in Section 1.2.3 does not hold in higher dimensions:

Proposition 1.53

Let X be a non-empty $\mathbb{Z}$-subshift. Then X contains a periodic point.

Proof. As periodic points are preserved by conjugacy, we can assume by Proposition 1.52 that $X = \mathrm{Walks}(G)$ for some finite graph G. As X is non-empty, there exists a bi-infinite walk x in G. In particular, G must contain a cycle, which can then be used to produce a periodic point: more precisely, as G is finite, there exists $i, j \in \mathbb{Z}$ such that $x_i = x_j$. Up to shifting x, we can even assume $i = 0$. The configuration $k \in \mathbb{Z} \mapsto x_{(k \mod j)}$ is then a point in X and is periodic of period j. □

For sofic subshifts, a very similar characterization holds, which in turns allows us to use the theory of **formal languages** to study problems about them, rather than graph theory in general:

Definition 1.54: Labeled walks

Let $G = (V, E, \lambda_E)$ be an *edge*-labeled graph. An **edge-labeled walk** is a sequence $y = (y_i)_{i \in [\![a,b]\!]}$ for $a \leq b \in \mathbb{N} \cup \{\pm\infty\}$, such that there exists a walk $x = (x_i)_{i \in [\![a,b+1]\!]}$ with $\lambda_E(x_i, x_{i+1}) = y_i$ for all i.

Said differently, an edge-labeled walk is the sequence of labels read along the edges taken by a walk, in the sense of Definition 1.50. One of the main tools to study sofic subshifts over $\mathbb{Z}$ is then the following result:

Theorem 1.55 [Wei73, Thm. 3]

$\mathbb{Z}$-sofic subshifts are exactly the sets of bi-infinite edge-labeled walks on graphs.

One important nuance with SFTs is that an edge-labeling map need not be injective: in particular, the same edge-labeled walk might correspond to different walks in the graph, and therefore, be extended in *infinite* edge-labeled walks in different ways.

Proof. Suppose that $Y \subseteq \mathcal{B}^{\mathbb{Z}}$ is sofic, so that by Proposition 1.46 it is a factor of a nearest-neighbour SFT $X \subseteq \mathcal{A}^{\mathbb{Z}^2}$ by some 0-block-map Φ with local map f. Let G be the graph given by Proposition 1.52 such that X is the set of walks on G. Define then the edge-labeled graph H as $H = (V(G), E(G), \lambda_E)$ where:

$$\begin{aligned} \lambda_E \colon E(H) &\to \mathcal{B} \\ (a,b) &\mapsto f(a) \end{aligned}$$

One can then check that edge-labeled walks on H are exactly the elements of $\Phi(X) = Y$.

Suppose now that Y is the set of bi-infinite walks on some graph $G = (V, E, \lambda_E)$. Simply define X as $X = \mathcal{X}_{V^2 \setminus E}$, and Φ the 1-block-map induced by by the local map $f \colon (\cdot, a, b) \in V^3 \mapsto (\lambda_E(a,b)) \in \mathcal{B}$. □

This result makes sofic subshifts amenable to analysis using graph-theoretic tools: for example, the same proof as Proposition 1.53 gives that $\mathbb{Z}$-sofic subshifts are either empty or contain a periodic point. However, as *labeled* walks on graphs, we can study sofic subshifts with another tool, the theory of formal languages, and in particular of regular languages. In order to define regular languages, we will introduce a combinatorial device, and which will be generalized in some sense in Section 1.2. This is a rather short introduction to the topic, and we only give the necessary definitions; a complete introduction to formal languages and automata theory can be found in [And06], although more details will be given in Section 2.1.1 as needed:

Definition 1.56: Finite Automaton

A **finite automaton** A is a quintuple $A = (\mathcal{A}, Q, \delta, I, F)$, satisfying:

- $\mathcal{A}$ is a finite alphabet, and Q a finite set of states.
- $\delta \colon Q \times \mathcal{A} \to 2^Q$ is a partial transition function.
- $I, F \subseteq Q$ are respectively an initial and final set of states.

This is the typical definition of automata as found in the literature. However, we can easily see the link between the formalism of Definition 1.56 and the one of edge-labeled walks: consider the graph G whose vertex set is Q, with edges $E = \{(q,q') \in Q^2 \mid q' \in \delta(q, \mathcal{A})\}$, edge labels $2^{\mathcal{A}}$ and an edge-labeling map defined by $\lambda_E(q,q') = \{a \in \mathcal{A} \mid q' \in \delta(q,a)\} \subseteq \mathcal{A}$.

This automaton is to be thought of as a machine, reading a word letter by letter and transitioning from to state according to the letters it reads:

Definition 1.57: Automaton run

Let $A = (\mathcal{A}, Q, \delta, I, F)$ be a finite automaton. A **valid run** of A on a word $w \in \mathcal{A}^n$ is a sequence of states $(q_0, q_1, q_2, \ldots, q_n) \in Q^{n+1}$ such that:

- $q_0 \in I$
- For $0 \leq i < n$, $q_{i+1} \in \delta(q_i, w_i)$

It is an **accepting run** if moreover $q_n \in F$.

In other words, a valid run on w is an edge-labeled walk on the graph defined after Definition 1.56, where the sequence of labels is the word w. We will therefore usually extend the map $\delta \colon Q \times \mathcal{A} \to 2^Q$ to a map $\delta \colon Q \times \mathcal{A}^* \to 2^Q$, setting $\delta(q,w)$ as the set of states that can be reached by a run on w starting from the state q. This allows us to define the language recognized by the automaton:

Definition 1.58: Regular language

The **language** $\mathcal{L}(A)$ of an automaton A is the set of words that admit some accepting run of A.
A language L is regular if there exists an automaton A such that $L = \mathcal{L}(A)$.

Using Theorem 1.55, the next proposition becomes an easy reformulation in the setting of formal languages:

Proposition 1.59

Let X be a sofic subshift. Then, $\bigcup_{n\in\mathbb{N}} \mathcal{L}_n(X)$ is a regular language.
Reciprocally, if L is regular, the subshift whose allowed words are the elements of L is sofic.

The second point of Proposition 1.59 is somewhat misleading, although correct: for it to be a non-empty statement, the language L must be infinite, and be – at least for some infinite family – downwards closed, in the sense that $(u \in L \wedge w \sqsubseteq u) \implies w \in L$.

Proposition 1.59 is a powerful way to show that some $\mathbb{Z}$-subshifts are not sofic, by showing that their languages are not regular. This is often an easy solution, as we will see in Section 2.1, as regular languages have been studied extensively and admit a large variety of characterizations.

A final theorem about finite automata which will make proofs of Chapter 2 easier relates the general finite automata introduced in Definition 1.56, and deterministic automata: an automaton $A = (\mathcal{A}, Q, \delta, I, F)$ is **deterministic** if $I = \{q_0\}$ for some $q_0 \in Q$, and $|\delta(q, a)| \leq 1$ for all $q \in Q, a \in \mathcal{A}$ – in that case, we write $\delta(q, a) = q'$ instead of $\delta(q, a) = \{q'\}$.

Theorem 1.60

Let $L \subseteq \mathcal{A}^*$ be a language. Then L is regular if and only there exists a **deterministic** automaton $A = (\mathcal{A}, Q, \delta, q_0, F)$ such that $L = \mathcal{L}(A)$.

1.2 Computability

1.2.1 Turing Machines and decision problems

We quickly introduce the main computability notions that we will use in this thesis, either directly or as a motivation for the kind of problems that we choose to look at. Although we recall the main definitions, we expect the reader to be somewhat familiar with the basic objects. We refer to [Soa16, Introduction, Turing Machines] for a more detailed exposition.

Historically, the "building block" of computability theory was the notion of **Turing Machine**. Although computability can be defined in a way that avoids mentioning Turing Machines altogether, this point of view makes the link with tilings quite clear, as we will see in Section 1.2.3. A Turing Machine can be viewed a generalization of finite automata (see Definition 1.56), with an additional *memory*.

Definition 1.61: Turing Machine

A (deterministic) Turing Machine M is a 6-tuple $M = (\mathcal{A}, Q, \square, \delta, q_0, F)$, where:

- $\mathcal{A}$ is a finite alphabet, and $\square \in \mathcal{A}$ is the **blank** symbol.
- Q is a finite set of states, and $q_0 \in Q$ is the initial state, $F \subseteq Q$ is a (possibly empty) set of final states.
- $\delta \colon Q \times \mathcal{A} \to Q \times Q \times \mathcal{A} \times \{-1, 0, 1\}$ is a *partial* transition function.

Just as we defined runs for finite automata Definition 1.57, we can define a run of a Turing Machine.

Definition 1.62: Turing Machine global state

A **global state** of some Turing Machine M is a triplet (Γ, h, q) where $\Gamma \in \mathcal{A}^{\mathbb{Z}}$ is a **tape**, $h \in \mathbb{Z}$ is the **head position**, and $q \in Q$ is the internal state of M.

We can now define the action of M on a global state. The idea is that a machine M acts on bi-infinite tapes, in which each cell contains a symbol from $\mathcal{A}$, using a read-write *head* located somewhere on the tape. At each step, the machine updates its internal state $q \in Q$ depending on the content of the tape at its head position, possibly rewrites the symbol in this position, and moves its head to an adjacent cell (or stays in place) – the new internal state, and the actions on the tape, are all given by δ. More formally:

Definition 1.63: Turing Machine run

Let $M = (\mathcal{A}, Q, \square, \delta, q_0, F)$ be a Turing Machine and $S = (\Gamma, h, q)$ be some global state. Let $(q', b, k) = \delta(q, \Gamma_h) \in Q \times \mathcal{A} \times \{-1, 0, 1\}$. A **step** of M on S is the global state $S' = (\Gamma', h', q')$, where:

- $h' = h + k \in \mathbb{Z}$.
- $\Gamma'_h = b \in \mathcal{A}$, and for $i \neq h$, $\Gamma'_i = \Gamma_i$.

If $\delta(q, \Gamma_h)$ is not defined, then there are no valid steps from this global state.

A **run** of M on some input $x = a_0 \dots a_{n-1} \in (\mathcal{A} \setminus \{\square\})^n$ is a (possibly infinite) sequence of steps $S^0, S^1, \dots$, where S^0 is the global state $(\Gamma^0, 0, q_0)$, with initial global state $\Gamma_0(i) = \begin{cases} a_i & \text{if } 0 \leq i < n \\ \square & \text{otherwise} \end{cases}$.

Just as with automata, we can associate to a Turing Machine a language, and more precisely, as we have a tape on which we can write, we can compute *functions*:

Definition 1.64: Turing Machine function

Let $M = (\mathcal{A}, Q, \Box, \delta, q_0, F)$ be a Turing Machine, with a *total* transition function δ, and $q_0 \notin F$. Let $(S^0, \dots, S^n)$ be a finite run on $x = a_0 \dots a_{n-1}$, where $S^i = (\Gamma^i, h_i, q_i)$ for $0 \leq i \leq n$ satisfying that for $i < n$, we have $q_i \notin F$ and $q_n \in F$. Let $y = \Gamma^n|_{[\![0, \min_{i>0}\{\Gamma^n_i = \Box\}-1]\!]}$ be the word on $\mathcal{A} \setminus \Box$ read at the origin of the last tape Γ^n. We write:

- $M(x)\Downarrow^n$ as M stops in n steps, or more generally $M(x)\Downarrow$, and say that M **halts** on input x.
- $M(x) = y$ is the value of the function computed by M on x.
- If there exists no such finite run, we write $M(x)\Uparrow$, and say that M does not halt.

In particular, $M\colon \mathcal{A}^* \to \mathcal{A}^*$ is not necessarily total, even if δ is.

We can now define decision problems, and what it means for a problem to be solvable:

Definition 1.65: Decision problem

A **decision problem** Problem is a subset of $\{0,1\}^*$. An element $x \in$ Problem is a solution, or positive instance, of the problem.

Rather than subsets of $\{0,1\}^*$, we will often define decision problems as *closed* questions about some objects:

Decision Problem Pseudo-problem

Input: Some object X.
Output: Whether X satisfies some property $\mathcal{P}$.

This is an equivalent point of view on the same object: indeed, this is a way to define the set of (encodings of) objects X in $\{0,1\}^*$ such that $\mathcal{P}(X)$ holds. This is also why it is important to restrict ourselves to closed questions. We often choose this formulation of decision problems, as we find that it makes more apparent the fact that one has to carefully define what the "valid inputs" are, and in particular, to ensure that they have a representation that could be given to a Turing Machine to operate on.

Definition 1.66: Decidable, enumerable

A decision problem Problem is **decidable** (or **solvable**) if there exists a total Turing Machine M such that for $x \in \{0,1\}^*$, $M(x) = 1$ if and only if $x \in$ Problem.
It is (recursively) **enumerable** if there exists M such that for all $x \in \{0,1\}^*$, $M(x)\Downarrow$ and $M(x) = 1$.

The first and most well-known example of a problem which is undecidable is the Halting

Problem:

Decision Problem HALT

Input: A Turing Machine M.
Output: Whether M halts on the empty input.

Theorem 1.67 [Tur+36]

HALT is undecidable.

For any property P, we say that a subset X of $\{0,1\}^*$ is **co-**P if and only if $\{0,1\}^* \setminus X$ is P.

We will need another last notion to classify the difficulty of certain decision problems. The notion of **reduction** is a way to formalize that one problem is harder than another:

Definition 1.68: Many-one reduction

Let $A, B \subset \{0,1\}^*$ be two decision problems. We say that B is harder than A, and write $A \leq_m B$, if there exists a total Turing Machine M such that $x \in A \iff M(x) \in B$. In that case, we say that M is a **(many-one) reduction** from A to B, or that A is reducible to B.

The definition of Turing Machines given in this section is very robust, in the sense that many natural variants have the same expressive power (that is, compute the same functions). For example, we can have more than a single tape, we can use mono-infinite rather than bi-infinite tapes, we can generalize the possible moves from $\{-1,0,1\}$ to any finite set, and so on. This will be useful when considering tilings that "represent" computations of Turing Machines, in Section 1.2.3.

1.2.2 Arithmetical hierarchy

In a lot of applications, knowing whether a problem is decidable or not can in itself be difficult; depending on the domain, one hopes that the problem being considered is solvable, and can then try to find the most "efficient" algorithm deciding it. On the other hand, we will see in Section 1.2.3 that most problems about higher-dimensional subshifts are undecidable, and that it is generally easy to show that a given problem is also undecidable. We therefore want to obtain more fine-grained information about its "degree of undecidability": we can in fact construct a hierarchy of undecidable problems, and one of our goals will then be to find exactly where our questions about subshifts lie in this hierarchy. This "hierarchy of undecidability" is called the **arithmetical hierarchy**. There are many approaches to define this hierarchy, and we choose one based on logical formulae. For definiteness, we consider *first-order* formulae in the language of Peano arithmetic.

Definition 1.69: Arithmetical hierarchy - Formulae

We recursively define for $n \geq 0$ the sets Σ_n^0 and Π_n^0 of logical formulae as follows:

- If ψ contains only bounded quantifiers, it belongs both to Σ_0^0 and Π_0^0.
- If ψ is equivalent to a formula $\exists x, \phi(x)$ with $\phi(x) \in \Pi_n^0$, then $\psi \in \Sigma_{n+1}^0$.
- If ψ is equivalent to a formula $\forall x, \phi(x)$ with $\phi(x) \in \Sigma_n^0$, then $\psi \in \Pi_{n+1}^0$.

This defines a hierarchy on the set of first-order formulae:

Proposition 1.70

For any $n \geq 0$:

- $\Sigma_n^0 \subsetneq \Sigma_{n+1}^0$ and $\Pi_n^0 \subsetneq \Pi_{n+1}^0$.
- $\Sigma_n^0 \subsetneq \Pi_{n+1}^0$ and $\Pi_n^0 \subsetneq \Sigma_{n+1}^0$.

Proof. The inclusions are immediate consequences of the definition. The fact that they are strict is a classical diagonal argument (see [Soa16, Corollary 4.2.3]). □

This is sufficient to assign a "level" to any formula: as there exists an encoding of pairs (and so of any finite sequence) of integers which can be written in first-order Peano arithmetic, the usual example being the Cantor pairing function $(x, y) \in \mathbb{N}^2 \mapsto \frac{(x+y+1)(x+y)}{2} + y \in \mathbb{N}$, a formula $\psi = \exists x_1, \exists x_2 \ldots \exists x_n \phi$ can always be written as an equivalent formula $\psi' = \exists x, \phi'(x)$ where ϕ and ϕ' are in the same level of the hierarchy (and the same holds for universal quantifiers).

We can now define an equivalent hierarchy of *sets* of natural numbers:

Definition 1.71: Arithmetical hierarchy - Sets

A set of integers $X \subset \mathbb{N}$ is a Π_n^0 (resp. Σ_n^0) set if and only if there exists a Π_n^0 (resp. Σ_n^0) formula ϕ such that for all $x \in \mathbb{N}$,

$$x \in X \iff \phi(x) \text{ is true}$$

As a special case, a set $X \in \Sigma_0^0 = \Pi_0^0$ is said to be recursive, or **computable**. We refer to [Soa16] for more details, but try to give some intuitions. There are some profound links between the arithmetical hierarchy thus defined, and computability theory. The first result is the following:

Proposition 1.72

A set X is Σ_1^0 if and only if there exists a Turing Machine M such that for $x \in \mathbb{N}$, $x \in X \iff M(x)\downarrow$, that is if X is recursively enumerable. It is Π_1^0 if and only there exists M such that for $x \in \mathbb{N}$, $x \in X \iff M(x)\uparrow$.

Proof sketch. We prove it for recursively enumerable sets. Let X be recursively enumerable, and M such that $x \in X \iff M(x)\Downarrow$ for all $x \in \mathbb{N}$. We can then write a formula $\psi(n)$ corresponding to the statement "M stops in at most n steps on input x" in first-order Peano arithmetic with only bounded quantifiers (we do not detail the exact encoding scheme, but it is enough to notice that in n steps, the machine only ever visits finitely many cells of its tape, and so we can encode its entire run so far with a finite integer). In particular, the formula $\exists n \psi(n)$ is equivalent to the fact that M halts. Conversely, as any formula with only bounded-quantifier in first-order Peano can be decided by a Turing Machine, it is sufficient to decide a Σ_1^0 formula $\psi(x)$ of the form $\exists n, \phi(x, n)$ to test for each possible value of n the (thus decidable) formula $\phi(x, n)$ and halt when it is true. □

We can define within each of those classes some "reference" problems, which are at least as hard as any other problem of the same class:

Definition 1.73: Complete problems

A problem P is said to be Σ_n^0**-hard** (resp. Π_n^0-hard) if for any Σ_n^0 problem A (resp. Π_n^0 problem A), we have $A \leq_m P$. If furthermore P is itself a Σ_n^0 problem, it is Σ_n^0**-complete** (resp. Π_n^0-complete).

By Proposition 1.72, Halt is Σ_1^0-complete. We will see in Section 1.2.3 that a natural problem related to tilings is also Σ_1^0-complete, and give in Section 2.3 other examples of complete problems, both related and unrelated to tilings, for higher levels in the hierarchy.

In order to characterize some conjugacy invariants of tilings, we need yet another hierarchy. As presented in Section 1.1.2, the entropy Definition 1.23 is a *real number* which is invariant by conjugacy. However, the hierarchy defined so far only makes sense for integers – in the language of Peano arithmetic. More generally, Turing Machines are discrete objects, and we have no way to make computations with real numbers. A general introduction to the field of computable analysis is [Wei12], and we only state the basic ideas. Our goal here is to characterize how "complicated" any real number is, from the point of view of computability theory. As real numbers can be defined using only rational numbers (which can themselves easily be encoded, so that Turing Machines can be assumed to work with them), we will characterize reals using the rationals which define it: more precisely, a real $x \in \mathbb{R}$ can be viewed as the set $\{q \in \mathbb{Q} \mid q \leq x\}$ (this is akin to "Dedekind cuts" construction of the real numbers). Up to any reasonable encoding, this is a set of natural numbers, which falls in some class of the arithmetic hierarchy:

Definition 1.74: Arithmetical hierarchy - Reals

We define for any $n \geq 0$ the following sets of real numbers:

$$\Sigma_n = \{x \in \mathbb{R} \mid \{q \in \mathbb{Q} \mid q \leq x \text{ is a } \Sigma_n^0 \text{ set}\}\}$$

and

$$\Pi_n = \{x \in \mathbb{R} \mid \{q \in \mathbb{Q} \mid q \leq x \text{ is a } \Pi_n^0 \text{ set}\}\}$$

There exists another definition, more analytical in nature, of those same classes. We will provide an even further simplified definition, which we will use in Section 2.3.

Proposition 1.75

For any $n \geq 0$,

$$\Sigma_n = \{\sup_{i_1} \inf_{i_2} \sup_{i_3} \dots f(i_1, \dots, i_n) \text{ for a computable } f \colon \mathbb{N}^n \to \mathbb{Q}\}$$

and

$$\Pi_n = \{\inf_{i_1} \sup_{i_2} \inf_{i_3} \dots f(i_1, \dots, i_n) \text{ for a computable } f \colon \mathbb{N}^n \to \mathbb{Q}\}$$

Proof. Let $x \in \Sigma_1$ (as defined in Definition 1.74). Therefore, $\{q \in \mathbb{Q} \mid q \leq x\}$ is a Σ_1^0 set, *i.e.* it is recursively enumerable, so there exists M a Turing Machine enumerating it. Then $x = \sup_i M(i)$. On the other hand, suppose that $x = \sup_i f(i)$ for a computable map $f \colon \mathbb{N} \to \mathbb{Q}$. Then $\{q \in \mathbb{Q} \mid q \leq x\} = \{q \in \mathbb{Q} \mid \exists i, q \leq f(i)\}$ which is clearly a Σ_1^0 set.

The same proof shows the equivalence for the class Π_1, and an induction on n prove in the same way that this holds for any n. □

In fact, Proposition 1.75 is the way the arithmetical hierarchy of real numbers if defined in [ZW01]. We can even make some additional assumptions on this expression. For notational convenience, we state and prove it only in the case $n = 3$, which is the one we need in Chapter 2, the general case being proven in exactly the same way:

Proposition 1.76

Let $x = \sup_i \inf_j \sup_k f(i, j, k)$ for some computable function $f \colon \mathbb{N}^3 \to \mathbb{Q}$. Then there exists $g \colon \mathbb{N}^3 \to \mathbb{Q}$ computable such that:

- $x = \sup_i \inf_j \sup_k \dots g(i, j, k)$
- For any $i, j \in \mathbb{N}$, $g_{ij} \colon k \in \mathbb{N} \mapsto g(i, j, k) \in \mathbb{Q}$ is non-decreasing.
- For any $i \in \mathbb{N}$, $g_i \colon j \in \mathbb{N} \mapsto \lim_{+\infty} g_{ij}$ is non-increasing.
- $i \in \mathbb{N} \mapsto \lim_{+\infty} g_i$ is non-decreasing.

Proof. This is already proven in [ZW01, Lemma 3.1], for the case $n = 2$ and with a slightly different presentation. For $i, j, k \in \mathbb{N}$, define $g(i, j, k) = \max_{i' \leq i} \min_{j' \leq j} \max_{k' \leq k} f(i, j, k)$. We claim that g satisfies the required conditions:

- g is computable: indeed, it suffices to compute all the images $f(i', j', k')$ for $i' \leq i, j' \leq j, k' \leq k$ to determine $g(i, j, k)$, which is a finite number of computable values.

- $g_{ij} \colon k \mapsto g(i, j, k)$ is clearly non-decreasing, the maps $g_i \colon j \mapsto \lim_{k \to +\infty} g_{ij}(k)$ is non-increasing, and $i \to \lim_{j \to +\infty} g_i(j)$ is also clearly non-decreasing.

- In the next computations, the sup and inf are replaced by limits by the previous remark about monotonicity:

$$\begin{aligned}
\sup_i \inf_j \sup_k g(i,j,k) &= \sup_i \inf_j \sup_k (\max_{i'\leq i} \min_{j'\leq j} \max_{k'\leq k} f(i',j',k')) \\
&= \sup_i \inf_j \lim_k (\max_{i'\leq i} \min_{j'\leq j} \max_{k'\leq k} f(i',j',k')) \\
&= \sup_i \inf_j (\max_{i'\leq i} \min_{j'\leq j} \lim_k \max_{k'\leq k} f(i',j',k')) \\
&= \sup_i \inf_j (\max_{i'\leq i} \min_{j'\leq j} \sup_k f(i',j',k)) \\
&= \sup_i \inf_j \sup_k f(i,j,k)
\end{aligned}$$

where the last line is either obtained by repeating the process of replacing and commuting the various $\sup, \inf, \lim$ operators, or more directly by applying [ZW01, Lemma 3.1, (7)].

□

1.2.3 Some natural links with subshifts

The domino problem

One of the first links between computability theory and subshifts is the fact that even the apparently simple problems on higher-dimensional SFTs are in fact undecidable. Hence, when asking any non-trivial question about subshifts, we cannot avoid the fact that it is most the time undecidable, and either try to find partial algorithms, or additional conditions on the subshifts that we consider to make the problem solvable, or use more advanced concepts in computability theory to fully characterize and understand the extent to which these problems are "hard".

The simplest problem one can ask about an SFT, or as historically introduced, about a set of Wang tiles, is whether one can find an infinite tiling of the plane using these tiles. This problem is commonly known as the DOMINO problem:

Decision Problem — DOMINO

Input: A finite family of $\mathbb{Z}^2$ forbidden patterns $\mathcal{F}$.
Output Whether $\mathcal{X}_\mathcal{F}$ is empty.

In dimension 1, this problem is easy: indeed, as shown in Proposition 1.52, it is equivalent to finding an infinite walk in a graph, that is, a cycle. A more abstract but equivalent (in this setting) decision procedure consists in using Proposition 1.53, and enumerate periodic points: this is known as Wang's Algorithm, which is in fact a semi-algorithm as it does not necessarily terminate. Given $\mathcal{F}$:

- Enumerate all the locally admissible rectangular patterns.
- If one has the same right and left, and top and bottom, sides, it can be "glued" to itself in all four directions to form a periodic tiling, and so $\mathcal{X}_\mathcal{F}$ is not empty.
- On the other hand, if there are no locally admissible patterns of support $\mathcal{Q}_n$, then there is *a fortiori* no globally admissible pattern of this size, and so $\mathcal{X}_\mathcal{F}$ is empty.

In dimension 2 and higher, this procedure does not terminate, due to the existence of aperiodic SFTs, which were an important tool for Berger to prove the following theorem:

Theorem 1.77 [Ber66]

DOMINO is Σ_1^0-complete.

Sketch of the proof. The upper-bound is a consequence of compactness: using the same procedure than in dimension 1, we can simply enumerate all the larger and larger locally admissible patterns. As local admissibility is decidable, and as the existence of arbitrarily large locally admissible patterns is equivalent to the subshift being non-empty, we get that DOMINO $\in \Sigma_1^0$.

The proof of the lower-bound is based on the following ideas:

- One can represent a run of a Turing Machine M using Wang Tiles, and more generally SFTs: each row in a configuration x represents a global state, with each cell of $x|_{\mathbb{Z}\times\{i\}}$ being a symbol of the alphabet of M, and some other layers indicate the position of the head, and of the current state. As each step (see Definition 1.63) is purely local (the tape, head and internal state only changes around the head), we can enforce the consistency of consecutive steps using local rules. See Figure 1.4 for how one can define such a tileset, and obtain a corresponding subshift X_M.

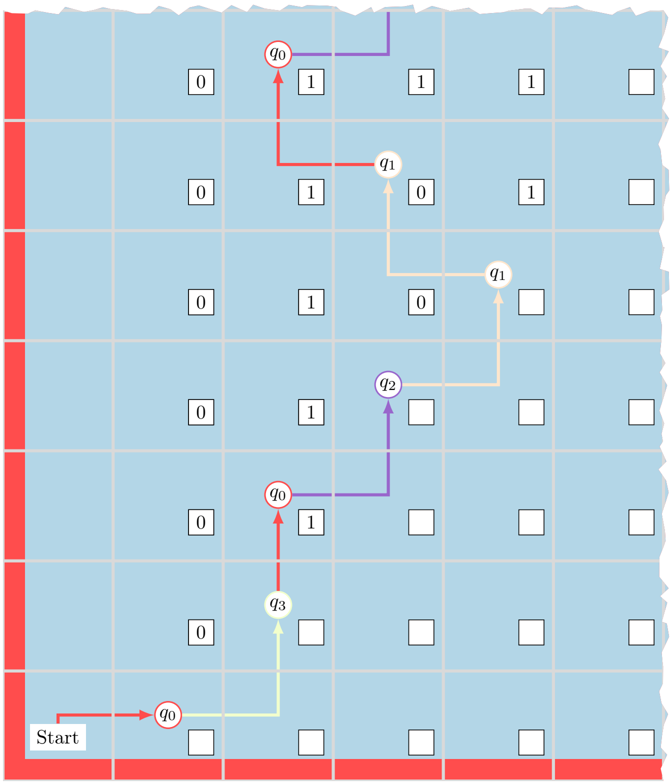

Figure 1.4: Part of a configuration encoding a run of a Turing Machine M on the empty input. We consider machines with one-way infinite tapes. Each row represents a global state, and time goes upwards. The initial state is q_0. The letter from the alphabet $\mathcal{A} = \{0, 1, \square\}$ is in the square node drawn on each cell, and the current state is in the coloured circle of a cell. There is only one state per row, corresponding to the position of the head on the tape. The precise encoding of M and the matching rules are left to the reader, but should easily be inferred from the figure: for example, we see here that $\delta(q_0, 1) = (q_2, 1, +1)$, and $\delta(q_3, \square) = (q_0, 1, 0)$

.

- The goal is then to reduce DOMINO to HALT, ensuring that the subshift is empty exactly when M does not halt on the empty input. To do this, we simply add to the list of forbidden patterns all the tiles containing any final state of M. This means that no configuration can encode a finite run ending in a final state of M.

- This is however not sufficient: for example, configurations corresponding to empty tapes $\square^{\mathbb{Z}}$ with no head and no internal state must exist by compactness, and so the reduction fails. In order to ensure that every infinite configuration actually corresponds to a run, it suffices to ensure that the *beginning* of a run appears somewhere; we do this by forcing it to appear "everywhere", using an elaborate extra layer. A visual representation is given in Figure 1.6, using the tileset of Figure 1.5 initially presented in [Rob71].

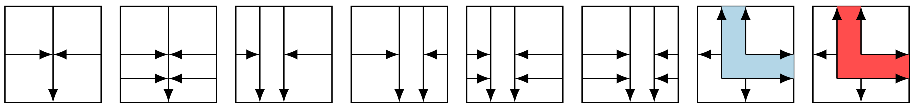

Figure 1.5: The Robinson tileset consists of these tiles, as well as their rotations by $\frac{\pi}{2}$, $-\frac{\pi}{2}$ and π. The arrows must be continued by an arrow going in the same direction. An additional constraint, not depicted in the tileset, is that there must a "blue corner" in each 2×2 square, and these blue tiles must be placed in a $2\mathbb{Z} \times 2\mathbb{Z}$-sublattice.

Using these tiles, the only valid configurations consist of "nested squares" of larger and larger size. A precise description of the valid configurations can be found *e.g.* in [GS21], in which the authors furthermore modify the tileset to ensure additional mixing properties; a more visual understanding of the valid configurations is given in Figure 1.6:

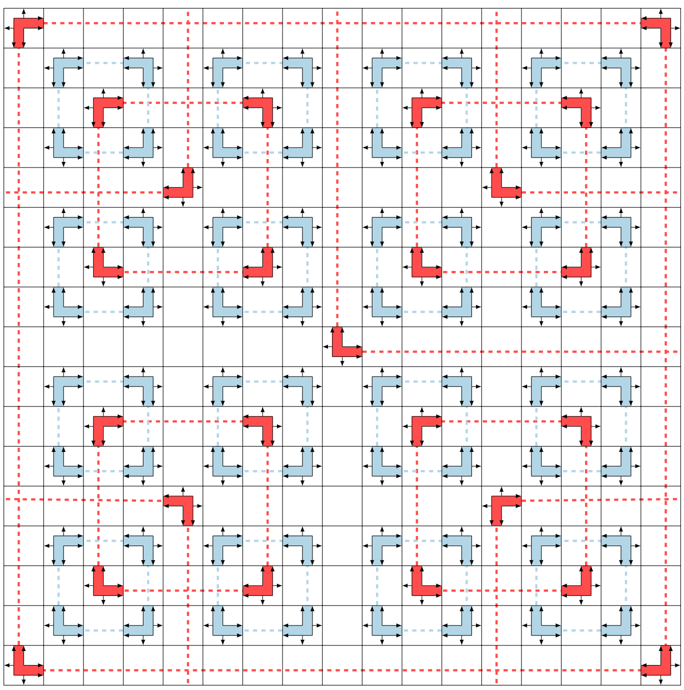

Figure 1.6: Part of a configuration of the Robinson SFT, where only the coloured tiles have been represented. Dashed lines correspond to tiles with a double arrow in this direction, the colour on the dashed lines being there for visual purposes only.

The idea is then to superimpose a layer from Robinson's subshift with the original subshift X_M simulating M, but not as a cartesian product. We instead use the "empty space" *between* the squares to embed longer and longer runs of M. Using this hierarchical structure, we can therefore correctly initialize the computation, simulating the same computations infinitely many times in parallel, so that the reduction finally works: the final subshift is empty if and only if M halts on the empty input.

□

Effective subshifts

We now present one last class of subshifts that we will be interested in during this thesis, namely effective subshifts. Although this class naturally appears when trying to consider larger classes than sofic subshifts defined by "algorithmic" conditions (rather than dynamical ones, as in Section 1.1.3), we will conclude this section by a theorem showing that they are in fact a very natural class, which arises when studying some problems which would *a priori* be unrelated to computability theory.

Definition 1.78: Effective subshifts

Let X be a $\mathbb{Z}^d$ subshift. If there exists a recursively enumerable family of finite forbidden patterns $\mathcal{F}$ such that $X = \mathcal{X}_{\mathcal{F}}$, then X is an **effective** subshift. If furthermore $\mathcal{L}(X)$ itself is recursively enumerable (and therefore computable), we say that X is a **computable** subshift.

Whenever we will consider decision problems for which the input can be an effective subshifts, we assume that it is given as a machine M which enumerates a family of forbidden patterns. In general, we will not consider algorithmic problems on subshifts which are not effective, as there is no clear way to specify what the input should even be.

We will see in Chapter 2 some ways to prove that effective subshifts are not sofic. Let us give a family of subshifts which are not even necessarily effective, but we will nevertheless consider in Chapter 3. We first give some background about balanced words. A general introduction to the objects presented here can be found in [Pyt+02, Sturmian Sequences] or [Lot02, Chapter 2].

Definition 1.79: Balanced word

A (possibly infinite) word x on the alphabet $\{0, 1\}$ is **balanced** if for all sub-words $u, v \sqsubseteq x$ of the same length $|u| = |v|$, we have $||u|_1 - |v|_1| \leq 1$.

Definition 1.80: Mechanical word

Let $\alpha \in [0, 1], \beta \in \mathbb{R}$. We define the **mechanical word** $x_{\alpha,\beta}$ of slope α and offset β as the bi-infinite sequence

$$\begin{aligned} x_{\alpha,\beta} \colon \mathbb{Z} &\mapsto \{0, 1\} \\ n &\to \lfloor \alpha(n + 1) + \beta \rfloor - \lfloor \alpha n + \beta \rfloor \end{aligned}$$

A visual illustration of mechanical words is given in Figure 1.7:

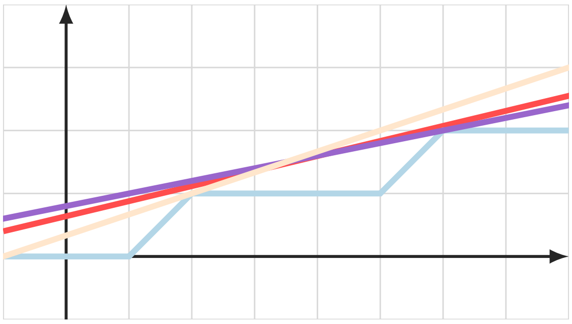

Figure 1.7: Examples of parameters (α, β) producing the same mechanical word 001000100 on the interval $[\![-1, 7]\!]$. The purple and orange lines have rational slope, and intersect the grid; any arbitrarily small variation of their slope avoids this technicality, while producing the same word.

Lemma 1.81 [Lot02, Lem. 2.1.14]

Mechanical infinite words are balanced.

Proposition 1.82 [Lot02, Prop. 2.1.17]

Let x be a finite balanced word. Then there exists an infinite mechanical word containing it as a subword.

It is false that infinite balanced words are mechanical – the simplest counter-example being ${}^{\infty}010^{\infty}$. Some partial converses do hold (see [Lot02, Lemma 2.1.15]), but they are not as easy to formulate for balanced words on $\mathbb{Z}$ as they are on $\mathbb{N}$.

We can then define the subshift associated to some parameters α, β. A particularly important example is the class of Sturmian subshifts:

Definition 1.83: Sturmian subshift

Let $\alpha \in [0, 1]$ be an irrational real. We define the **Sturmian subshift** X_α as

$$X_\alpha = \overline{\mathrm{Orb}(x_{\alpha,0})},$$

the closure of the orbit of the mechanical word $x_{\alpha,0}$.

As there are uncountably many non-computable irrational reals α, this class contains uncountably many non-effective subshifts, as $X_\alpha, X_{\alpha'}$ are not conjugate for $\alpha \neq \alpha'$. There are many equivalent definitions of Sturmian subshifts, we give another classical one:

Proposition 1.84

Let X be a $\mathbb{Z}$ subshift which is not eventually periodic. Then X is Sturmian if and only if for all $n \geq 1$, $|\mathcal{L}_n(X)| = n + 1$.

See [Lot02, Chapter 2.1] for a proof. From this, we can deduce the next characterization on mechanical words:

Proposition 1.85

If α is irrational, X_α contains no periodic point and is therefore not sofic. It is effective if and only if $\alpha \in \Sigma_1 \cap \Pi_1$.

Proof. We combine Proposition 1.84 and the fact that all the points of X_α have the same language. As periodic configurations have complexity $(\mathcal{L}_n(X))_{n\in\mathbb{N}}$ bounded by a constant, this proves that X_α is aperiodic. See [Lot02, Theorem 2.1.5] for another proof.

The fact that effectiveness of X_α is equivalent to the computability of α is a consequence of the fact that for $x \in X_\alpha$, we have $\lim_{n\to+\infty} \frac{|x|_{[0,n-1]}|_1}{n} \to \alpha$. In particular, if X_α was effective, we could approximate α, and the reciprocal is clear using the definition of mechanical words. □

As a convention, we will try whenever possible to use Z to denote effective subshifts, Y for sofic subshifts, and X either for an SFT or to designate any arbitrary subshift.

Lifting constructions

We define two variants of an operation defined on subshifts of dimension d which produce a new subshift of dimension $d+1$:

Definition 1.86: Lifts

Let $z \in \mathcal{A}^{\mathbb{Z}^d}$ be any configuration. We define its **periodic lift** as the $d+1$-dimensional configuration

$$\begin{aligned} z^{\uparrow} : \mathbb{Z}^{d+1} &\to \mathcal{A} \\ (i_1, \ldots, i_d, i_{d+1}) &\mapsto z_{(i_1,\ldots,i_d)} \end{aligned}$$

The (periodic) lift of a $\mathbb{Z}^d$-subshift X is $X^{\uparrow} = \{x^{\uparrow} \mid x \in X\}$.
We define the **free lift** of a subshift X as:

$$X^{\Uparrow} = \{x \in \mathcal{A}^{\mathbb{Z}^{d+1}} \mid \forall i \in \mathbb{Z}, x|_{\mathbb{Z}^d \times \{i\}} \in X\}$$

In other words, $X^{\uparrow}$ contains configurations which are obtained by stacking infinitely many copies of the same configuration $x \in X$ on top of one another, in the extra dimension; and $X^{\Uparrow}$ is obtained by stacking configurations of X, but not necessarily the same one.

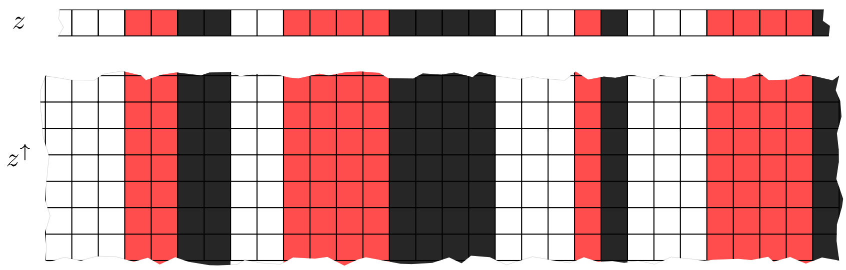

Figure 1.8: An example of a periodically-lifted configuration, from some non-sofic $\mathbb{Z}$-subshift.

Theorem 1.87: Lift [Hoc09b],[AS13, Thm 3.1],[DRS12, Thm. 10]

X is an effective $\mathbb{Z}^d$-subshift if and only if $X^{\uparrow}$ is a sofic $\mathbb{Z}^{d+1}$-subshift.

Brief sketch of the proof. There are two main ways by which this theorem has been proven: the first one, used in [Hoc09b] and [AS13], is a refinement on the ideas of Theorem 1.77. Instead of simply simulating a machine in the configurations, the simulated machine now also needs to check whether any row of the $\mathbb{Z}^2$ configurations contain a forbidden pattern of the original $\mathbb{Z}$-effective subshift. In order to do this, a complex system of information routing, synchronization mechanism ... are developed and "embedded" into valid configurations, so that the computations of the simulated machine eventually detect any forbidden pattern in themselves, and "forbid" them by entering some special, explicitly forbidden state.

The proof of [DRS12] is quite different, and relies on a analogous to Kleene's fixed point theorem. The idea is to build a self-simulating tileset, the configurations of which can be decomposed into larger and larger $N \times N$ "grids", each simulating a valid configuration of X itself. Here again, one embeds computations of universal Turing Machines *within* these self-simulating configurations, which are then able to perform longer and longer computations at higher "levels" of the simulation, and eventually detect forbidden patterns too.

Both proofs are quite technical, and we simply refer to the original articles for more details. The very high-level overview in both case is the same: we embed, in additional layers, computations of Turing Machines that enumerate the forbidden patterns of the effective subshift, and check that they do not appear in a configuration. Using a block-map, we then "remove" these computations, keeping only the underlying layer of the original lifted effective $\mathbb{Z}$-subshift, which is therefore sofic. □

This theorem will mainly be used as a black-box: in order to construct a sofic subshift (or SFT) in dimension $d \geq 2$ satisfying some properties, it will be convenient to first define an effective subshift in dimension $d-1$, with the required properties (which is usually easier, as effective subshifts are much less constrained than sofic subshifts), and then lift it to obtain a d-dimensional sofic shift (and its SFT extension), which hopefully still satisfies the required conditions. This will be especially true when we want to impose computability-type restrictions, as they will naturally tend to define effective subshifts.

1.3 Some notions of group theory

In this thesis, we will explore some of the links between tilings and groups in two different ways: in Chapter 3, we will study a conjugacy invariant which happens to be a group; in order to obtain characterizations of groups that can be obtained in this way, we need to introduce a way to abstractly describe general groups, and to (coarsely) classify them according to how "complicated" they are from an algorithmic point of view. In Chapter 4, we will introduce tilings of *graphs* (rather than tilings of $\mathbb{Z}^d$), using some motivations coming from group theory presented in Section 1.3.2 and Section 1.3.3. The necessary tools are presented in Section 1.3.1. For yet another approach to the study of links between group theory and symbolic dynamics, see [Van19].

1.3.1 Group presentations

We define in this section the notion of **presentation** of a group, using a combinatorial point of view rather than a purely algebraic one, making the results of Section 3.5 more intuitive. We assume that the reader is already familiar with the notion of (possibly infinite, non-necessarily abelian) group. The point of view adopted here can be found with more details in [MKS04].

Given a group $(G, *)$, we say that $S \subseteq G$ is a **generating set** if any element g of G can be written as a product $s_1 * s_2 * \cdots * s_n = g$. It is **symmetric** if for any $s \in S$, we also have $s^{-1} \in S$. From now on, we will always use the multiplicative notation for groups, and omit the operator $*$, writing *e.g.* gh for the product $g * h$.

We are now going to view group elements as being *words*, on an alphabet of generators, where words are considered the same if they can be rewritten into one another by adding, or removing, subwords which "represent" identity words in G. Let us make this idea more precise. Let S be some (non-necessarily finite) alphabet, and define $S^{-1} = \{s^{-1} \mid s \in S\}$ a set of formal inverses. A **word** in S is then an element of $(S \cup S^{-1})^*$, where the empty word is denoted ε. The fact that w, w' are equal (as words in S) is denoted by $w \equiv w'$. We use the standard notation $s^n = \underbrace{(s, s, \ldots, s)}_{n \text{ times}}$ for any $s \in S$, and similarly $s^{-n} = (s^{-1}, \ldots, s^{-1})$ for $s^{-1} \in S^{-1}$. We write as usual $s_1 s_2 \ldots s_n$ the word $(s_1, \ldots, s_n)$. For a word $w = w_1 \ldots w_n \in (S \cup S^{-1})^n$, we define the word w^{-1} as $w^{-1} = w'_n \ldots w'_1$, where $w'_i = \begin{cases} s^{-1} & \text{if } w_i = s \in S \\ s & \text{if } w_i = s^{-1} \in S^{-1} \end{cases}$. We define $\varepsilon^{-1} = \varepsilon$. Let now R be a (non-necessarily finite) set of words, called **relators**. We define a binary relation $\rightarrow_{S,R}$ on $((S \cup S^{-1})^*)^2$, that we write $w \rightarrow_{S,R} w'$ or $w \rightarrow w'$ when the context is clear, by $w \rightarrow w' \iff \exists r \in R, 1 \leq k \leq |w|, w' = w_1 \ldots w_k r w_{k+1} \ldots w_n$. In other words, $w \rightarrow w'$ if w' can be obtained by inserting a relator $r \in R$ somewhere in w. We denote by $\leftrightarrow^*_{S,R}$ the reflexive and transitive closure of $\rightarrow$, and $[w]$ the equivalence class of w for this relation.

Definition 1.88: Presentation of a group

Let S be an alphabet, S^{-1} a set of formal inverses of S, and R a set of finite words on $S \cup S^{-1}$. Let $R' = R \cup \{ss^{-1} \mid s \in S\} \cup \{s^{-1}s \mid s \in S\}$. The group defined by the **presentation** $\langle S \mid R \rangle$, written $G = \langle S \mid R \rangle$, is the group whose elements are the classes $\{[w] \mid w \in (S \cup S^{-1})^*\}$ for $\leftrightarrow^*_{S,R'}$, with inverses w^{-1}, and the group operation being word concatenation.

See [MKS04, Theorem 1.1] and more generally the book's first chapter for more details on this construction, in particular the fact that it is indeed a group, and that the unit of this group is $[\varepsilon]$.

Example 2. *We can define $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/3\mathbb{Z}$ with the presentation $\left\langle a, b \mid a^2, b^3, aba^{-1}b^{-1} \right\rangle$. We give an example of successive rewriting showing that $abbab \leftrightarrow^* \varepsilon$:*

$$\begin{aligned} abbab &\rightarrow abb\color{red}{bb^{-1}}\color{black}ab \\ &\leftarrow ab^{-1}ab \\ &\rightarrow ab^{-1}a\color{red}{aba^{-1}b^{-1}}\color{black}b \\ &\leftarrow ab^{-1}ba^{-1} \\ &\leftarrow aa^{-1} \\ &\leftarrow \varepsilon \end{aligned}$$

Relators are sometimes in the form of a **relation**, that is, an equality $w = w'$, which defines another relation like $\rightarrow$ where we can *replace* w by w' (instead of inserting): this is an equivalent point of view, as we can consider the relator ww'^{-1} instead, defining the same group in the end.

Algebraic properties are not clearly related to group presentations. In particular, there is *a priori* no clear relation between a presentation of a group, and whether it is *e.g.* empty, cyclic, simple... We will see below that those are indeed hard decision problems. For now, let us simply define some classes of groups which are amenable to algorithmic analysis, and on which we will mainly focus in the remainder of this thesis:

Definition 1.89: Finitely, recursively presented group

Let G be a group. It is **finitely generated** if there exists a finite S and a set R such that $G \simeq \langle S \mid R \rangle$. It is **finitely presented** if R can moreover be chosen finite. It is **recursively presented** if there exists a (non-necessarily finite) set S, and a recursively enumerable family R, such that $G \simeq \langle S \mid R \rangle$.

As with subshifts, where decision problems on SFTs were undecidable, most non-trivial decision problems on groups are also undecidable:

Decision Problem — WORD

Input: A finitely presented group $\langle S \mid R \rangle$ and a word w on $S \cup S^{-1}$.
Output Whether $[w] = [\varepsilon]$ in the group $\langle S \mid R \rangle$.

Theorem 1.90 [Nov55; Nov58], [Boo58]

The WORD problem is Σ_1^0-complete.

Note that the word problem is equivalent to asking whether $\langle S \mid R \rangle$ itself is trivial, in the case of finitely presented groups: indeed, it is trivial if and only each generator s is, of which there are finitely many.

1.3.2 Cayley graphs

Finally, we show how each group can be viewed as a graph, and how one can generalize tilings to arbitrary finitely generated groups G rather than simply $\mathbb{Z}^d$.

Definition 1.91: Cayley graph

Let G be a group, and S be a set of generators of G. The **Cayley graph** is an edge-labeled directed graph $\Gamma_{G,S} = (V, E, \lambda_E)$, defined by:

- $V = G$, that is each group element is a vertex.
- $(g, gs) \in E \iff s \in S$.
- $\lambda_E(g, gs) = s$.

With this definition, each path $(s_1, \ldots, s_n)$ from the identity vertex 1_G to any element $g \in G$ "represents" the same element of the group, that is, $s_1 s_2 \ldots s_n =_G g$. Indeed, whenever we follow an edge, we multiply on the right the starting vertex to obtain the ending vertex of this edge. In particular, any cycle in the graph $\Gamma_{G,S}$ is a relation, that is, a word on the alphabet of generators which is equal (in G) to the identity of G. Note though that the definition is not symmetric with respect to left or right multiplication: the vertices g and sg for any $s \in S, g \in G$ might be at a large distance in $\Gamma_{G,S}$, while g, gs are neighbours by definition. Also note that the graph depends on the chosen generating set: for example, the groups $\mathbb{Z}/n\mathbb{Z}$ have cyclic Cayley graphs if one considers the presentation $\mathbb{Z}/n\mathbb{Z} = \langle a \mid a^n = 1 \rangle$, but the Cayley graph is a clique (up to the edge-labeling) for the presentation $\langle a_0, \dots, a_{n-1} \mid a_i a_j = a_{i+j \mod n}, 0 \leq i, j < n \rangle$, see Figure 1.9.

Notation. The cyclic graph of order n, denoted C_n, is the undirected graph $C_n = (\llbracket 0, n-1 \rrbracket, E)$ where $\{i, j\} \in E \iff j = i + 1 \mod n$.

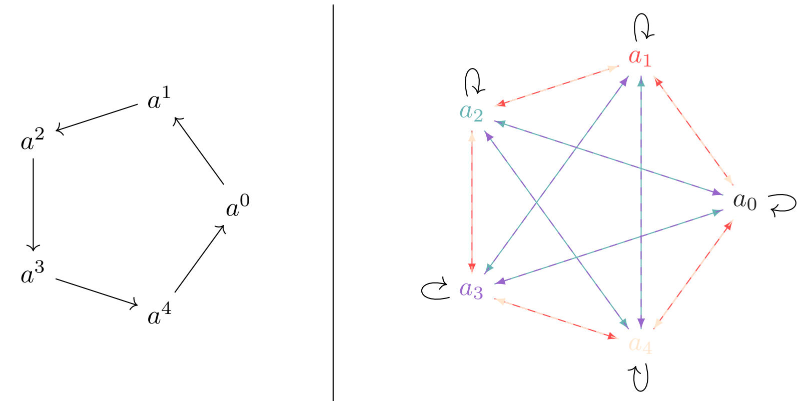

Cayley graph of $\mathbb{Z}/5\mathbb{Z}$ for the standard presentation $\langle a \mid a^5 \rangle$

Cayley graph of $\mathbb{Z}/5\mathbb{Z}$ for the presentation $\langle a_0, \dots, a_4 \mid a_{(i+j \mod 5)} a_i^{-1} a_j^{-1}, 0 \leq i, j \leq 4 \rangle$

Figure 1.9: Two presentations of the same group, and the corresponding Cayley graphs

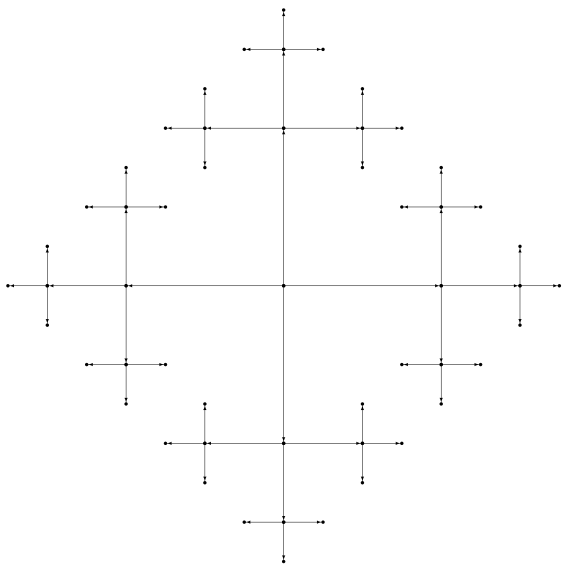

Figure 1.10: The ball of radius 3 of the Cayley graph of the free group F_2 with the standard presentation.

We will consider in Chapter 4 graphs that are not necessarily Cayley graphs, although they will present some similarities. There are theorems characterizing when a graph is indeed a Cayley graph of some group (and, in this case, of which group), but we will not really be interested in this question within this thesis.

1.3.3 Tilings on groups

One can define a notion of subshift on groups, just as we defined subshifts on $\mathbb{Z}^d$ in Definition 1.4. The simplest definition uses the characterization given in Proposition 1.32: one defines a topology on G using cylinders, and subshifts on some alphabet $\mathcal{A}$ are then the shift-invariants subsets of $\mathcal{A}^G$, where the shift σ_g is defined for $g \in G$ by $\sigma_g : x \in \mathcal{A}^g \mapsto (h \mapsto x_{hg})$ – other conventions exist in the literature, but they are all obviously equivalent. As was the case over $\mathbb{Z}^d$, there is also a combinatorial characterization of subshifts, using forbidden patterns. We will explore a very similar formalism in Section 4.3. Although we do not yet give precise definitions, we can also consider an equivalent to Wang Tiles, which would be $2n$-uplets where $n = |S|$. A valid tiling x is then a way to associate a tile to each element $g \in G$, and the tiles x_g, x_{gs} for $s \in S$ must coincide on their "s and and s^{-1}-sides". As an informal example, let us consider the following example over the free group F_2, whose Cayley graph is depicted in Figure 1.10:

Example 3 (Wang tiling on F_2)**.** *Figure 1.11 shows a locally (and provably globally) admissible pattern of a Wang Tiling on F_2, using a Wang tileset for which it is easy to easy to see that no tiling of $\mathbb{Z}^2$ exists (using the obvious morphism $F_2 \to \mathbb{Z}^2$, we can also view the "F_2-Wang tiles" as the usual Wang tiles).*

$T = \{$ 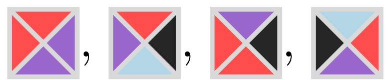 $\}$

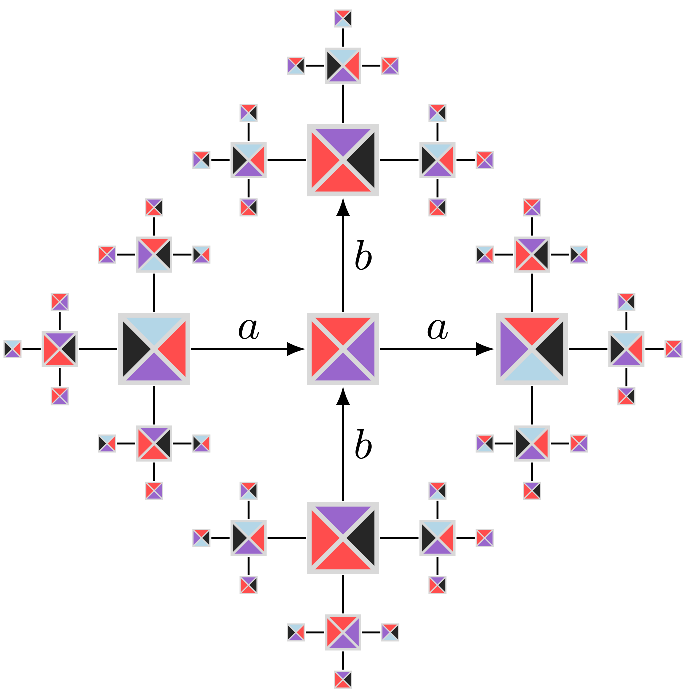

Figure 1.11: Tiling of the ball of radius 3 of F_2. Wang tiles are represented as squares, where the right, top, left and bottom sides correspond respectively to the generators a, b, a^{-1}, b^{-1} in F_2.

There are some important results concerning tilings on Cayley graphs of finitely presented groups, but the most important ones state that the "natural" problems do not depend on the chosen set of generators for the group. More precisely, the domino problem on the Cayley graph $\Gamma_{G,S}$ is decidable if and only if it is decidable on $\Gamma_{G,S'}$ for any other generating set S' of G. This is a well-known fact in this area, for which we did not find the first reference: we refer to the recent thesis [Bit24] for the usual constructions and references about decision problems on subshifts on groups, and the aforementioned result is in particular proven as Lemma 2.0.7.

Chapter 2

Extender entropies

In this chapter, we try to look at a conjugacy invariant defined by Thomas French and Ronnie Pavlov in [FP19], called the **extender entropy** of a subshift. This is a real number, which tries to quantify the number of patterns of a given size that can be freely exchanged within any configuration. One possible interest of looking at this quantity is to study the following problem: given a family of forbidden patterns $\mathcal{F}$, is the subshift $\mathcal{X}_\mathcal{F}$ sofic ? It turns out that this in Section 1.1.4: indeed, if one-dimensional question is much harder in dimension $d \geq 2$, for reasons already presented sofic subshifts can be seen as automata and walks on labeled graphs, multi-dimensional subshifts are on the contrary complicated objects, from an algorithmic perspective. In particular, deciding whether an SFT (and *a fortiori* a sofic subshift) is empty is already undecidable (see Section 1.2.3). In higher dimensions, a less ambitious approach than solving this problem then consists in studying criteria, typically in the form of conjugacy invariants h, which take strictly more values on the class $\mathcal{E}_d$ of d-dimensional effective subshifts than on the class $\mathcal{S}$ of d-dimensional sofic subshifts. For invariants h for which $h(\mathcal{E}_d) = h(\mathcal{S}_d)$, we still obtain the interesting fact, which is possibly a surprise, that sofic subshifts and behind them SFTs can "express" as much complexity as the much more general class $\mathcal{E}_d$. In this case, it is then natural to study more precisely what are the values taken by h, and to find some characterizations or subclasses $\mathcal{T}$ for which $h(\mathcal{T}) \subsetneq h(\mathcal{S})$.

The general organisation of this chapter is as follows: in Section 2.1, we present the motivations behind the introduction of this specific invariant, the extender entropy h_E, and we explain in particular how it relates to well-known problems on regular languages. In Section 2.2, we give the actual definition of h_E, and prove some basic facts and properties. Section 2.3 then presents the main tool used to study h_E in the rest of the chapter, and shows that computability theory is a natural object to characterize extender entropies. Finally, Section 2.4 gives some characterizations of the values taken by h_E on some usual classes of subshifts, such as minimal or mixing subshifts, and proves the main results of the chapter, which is a complete characterization of $h_E(\mathcal{E}_d)$ and $h_E(\mathcal{S}_d)$ for any $d \geq 1$, and we prove that for $d \geq 2$, they are identical.

Some of the results presented in this chapter are also available in [CPV24], unpublished at the time of writing.

2.1 Extender sets

As briefy mentioned above, the main notions considered in this chapter are in fact inspired by classical problems on regular languages, and their natural links to sofic $\mathbb{Z}$ subshifts. This translation from results on automata and languages to subshifts has already been exploited, and most of the content of this section can already be found in [LM21, Chapter 3]. A systematic study of extender sets for one-dimensional subshifts as defined in this section can be found in [Fre16b], in particular in the case of sofic $\mathbb{Z}$-subshifts.

2.1.1 Regular languages

We give in this section additional results and ideas to study regular languages that were not already mentioned in Section 1.1.4. We refer to this section for the most basic definitions.

Follower and predecessor sets

Definition 2.1: Follower, predecessor sets - languages

Let $\mathcal{A}$ be a finite alphabet and $\mathcal{L} \subset \mathcal{A}^*$ be a language. For any $x \in \mathcal{A}^*$, we call **follower set** of x the set

$$F_{\mathcal{L}}(x) = \{y \in \mathcal{A}^*, x \cdot y \in \mathcal{L}\}$$

We call **predecessor set** of x the set

$$P_{\mathcal{L}}(x) = \{y \in \mathcal{A}^*, y \cdot x \in \mathcal{L}\}$$

In other words, $F_{\mathcal{L}}(x)$ is the set of valid (right-)completions of x in $\mathcal{L}$. Note that this set can be empty, even if $x \in \mathcal{L}$, and that it might be non-empty for $x \notin \mathcal{L}$. A classical result from the theory of formal languages provides a complete characterization for regular languages:

Theorem 2.2: Myhill-Nerode Theorem [Myh57; Ner58]

Let $\mathcal{L} \subset \mathcal{A}^*$ be an arbitrary language. Then $\mathcal{L}$ is regular if and only $F = \{F_{\mathcal{L}}(x), x \in \mathcal{A}^*\}$ is finite.

Note that each set $F_{\mathcal{L}}(x)$ might be an infinite set; the theorem simply states that there are finitely many classes of words admitting different valid completions in $\mathcal{L}$ exactly when $\mathcal{L}$ is regular.

Proof. The easy direction is to show that $\mathcal{L}$ being regular implies that F is finite. In this case, by Theorem 1.60, there exists an automaton $A = (\mathcal{A}, Q, \delta, q_0, T)$ recognizing $\mathcal{L}$. We can assume that $\mathcal{A}$ is complete, that is, $\delta\colon Q \times \mathcal{A}$ is total (and so the extension $\delta\colon Q \times \mathcal{A}^*$ is also total), and that A is deterministic. Then, for $x, y \in \mathcal{A}^*$, if we have $\delta(q_0, x) = \delta(q_0, y)$ then $F_{\mathcal{L}}(x) = F_{\mathcal{L}}(y)$. Indeed, for any $w \in \mathcal{A}^*$,

$$\begin{aligned} w \in F_{\mathcal{L}}(x) &\iff x \cdot w \in \mathcal{L} \\ &\iff \delta(q_0, x \cdot w) \in T \\ &\iff \delta(\delta(q_0, x), w) \in T \\ &\iff \delta(\delta(q_0, y), w) \in T \\ &\iff w \in F_{\mathcal{L}}(y) \end{aligned}$$

In particular, $|F| \leq |Q|$.

The other direction is only slightly more involved: if F is finite, define the automaton $A = (\mathcal{A}, F, \delta_F, F_{\mathcal{L}}(\emptyset), \{F_{\mathcal{L}}(x), x \in \mathcal{L}\})$, where $\delta_F(F_{\mathcal{L}}(y), a) = F_{\mathcal{L}}(y \cdot a) \in F$. It is routine to show that this is a well-defined transition function, and that the automaton A accepts exactly $\mathcal{L}$. □

This theorem is mainly useful to prove that some language $\mathcal{L}$ is not regular. Indeed, it suffices to exhibit an infinite family $(x_n)_{n \in \mathbb{N}}$ of words having pairwise distinct follower sets to ensure that $\mathcal{L}$ is not regular by Theorem 2.2, and to show that $F_{\mathcal{L}}(x_i) \neq F_{\mathcal{L}}(x_j)$, it is enough to exhibit some y_i such that $x_i \cdot y_i \in \mathcal{L}$ and $x_j \cdot y_i \notin \mathcal{L}$.

Example 4. *The language $\mathcal{L} = \{a^n b^n, n \in \mathbb{N}\}$ is not regular. Indeed, for $i \neq j$, $b^i \in F_{\mathcal{L}}(a^i) \setminus F_{\mathcal{L}}(a^j)$, so $(F_{\mathcal{L}}(a^i))_{i \in \mathbb{N}}$ is an infinite family of follower sets.*

As regular languages are a stable class, we also explain why we usually consider only follower sets rather than predecessor sets:

Lemma 2.3

Let $\mathcal{L} \subseteq \mathcal{A}^*$ be an arbitrary language. Then $\{F_{\mathcal{L}}(x), x \in \mathcal{A}^*\}$ is finite if and only if $\{P_{\mathcal{L}}(x), x \in \mathcal{A}^*\}$ is finite.

Proof. This is simply because $\overline{\mathcal{L}} = \{w_n w_{n-1} \dots w_0 \mid w_0 \dots w_n \in \mathcal{L}\}$ is regular if and only if $\mathcal{L}$ is regular, and $F_{\mathcal{L}}(x) = \overline{P_{\overline{\mathcal{L}}}(\overline{x})}$. □

Extender sets and syntactic monoids

Adapting the Myhill-Nerode theorem to the case of one-dimensional sofic subshifts is quite straightforward. We introduce an intermediate step, which will make it easier to generalize the definitions to higher-dimensional subshifts, while not changing anything for $\mathbb{Z}$-subshifts.

Definition 2.4: Extender set - language

Let $\mathcal{L} \subset \mathcal{A}^*$ be a language, and $x \in \mathcal{A}^*$. We call **extender set** of x the set

$$E_{\mathcal{L}}(x) = \{(u, v) \in (\mathcal{A}^*)^2, u \cdot x \cdot v \in \mathcal{L}\}$$

For languages, this does not change the previous results:

Proposition 2.5

A language $\mathcal{L}$ is regular if and only if $E = \{E_{\mathcal{L}}(x)\}$ is finite.

Proof. Denote F, P, E respectively the set of follower sets, predecessor sets, and extender sets of $\mathcal{L}$. Clearly, $|E| \geq |F|$ so E being finite implies that $\mathcal{L}$ is regular by Theorem 2.2. Indeed, if $E_{\mathcal{L}}(x) = E_{\mathcal{L}}(y)$ for any words $x, y \in \mathcal{A}^*$, then $F_{\mathcal{L}}(\varepsilon \cdot x) = F_{\mathcal{L}}(\varepsilon \cdot y)$ as for any v, $\varepsilon x v \in \mathcal{L} \iff \varepsilon y v \in \mathcal{L}$.

Reciprocally, if $\mathcal{L}$ is regular, then we can consider a minimal automaton $A = (\mathcal{A}, Q, \delta, q_0, T)$ recognizing $\mathcal{L}$. It is easy to adapt the proof of Theorem 2.2 to show that for any $x, y \in \mathcal{A}^*$, we have $E_{\mathcal{L}}(x) = E_{\mathcal{L}}(y)$ if and only if $\delta(q_0, x) = \delta(q_0, y)$. □

This construction is also standard in the literature, and is one of the main ideas in the algebraic study of formal languages. Denoting $\sim_{\mathcal{L}}$ the equivalence relation on $\mathcal{A}^*$ defined by $x \sim_{\mathcal{L}} y$ if and only if $E_{\mathcal{L}}(x) = E_{\mathcal{L}}(y)$ – it is easy to see that this indeed defines an equivalence relation – we can define a central object in the theory of formal languages, the **syntactic monoid**[1] of a language:

Definition 2.6: Syntactic monoid [And06, Def 3.6]

Let $\mathcal{L}$ be a language. Then $M(L) = L/\sim_L$ with the concatenation operation is a monoid, called the **syntactic monoid** of L.

We will state, at the end of this section, a minor result about syntactic monoids, and more precisely their growth rate, in a sense that we will precise in due time. A general introduction to this domain can be found in [And06].

Extender sets and subshifts

In the case of subshifts, we are no longer dealing with finite words, but infinite configurations on $\mathbb{Z}^d$. However, as subshifts are compact, the definitions are sufficiently robust to accomodate for this difference:

[1] Formally, a monoid is a set with an associative binary operation and an identity element. Informally, a monoid is a group in which some (or all) elements might not have inverses.

Notation. For possibly infinite patterns u, v over a common alphabet $\mathcal{A}$, if $\mathrm{dom}(u) \cap \mathrm{dom}(v) = \emptyset$, we define $u \sqcup v$ as the pattern

$$u \sqcup v \colon (\mathrm{dom}(u) \cup \mathrm{dom}(v)) \to \mathcal{A}$$
$$z \mapsto \begin{cases} u_z & \text{if } z \in \mathrm{dom}(u) \\ v_z & \text{otherwise} \end{cases}$$

Definition 2.7: Extender set - subshift

Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a subshift, and let $D \subset_f \mathbb{Z}^d$, $u \in \mathcal{A}^D$. We call **extender set** of u the set

$$E_X(u) = \{x \in \mathcal{A}^{\mathbb{Z}^d \setminus D}, x \sqcup u \in X\}$$

Lemma 2.8

Let $X \subseteq \mathcal{A}^{\mathbb{Z}}$ be a subshift, $n \geq 0$ and $u, v \in \mathcal{A}^n$. Then $E_X(u) = E_x(v) \iff E_{\mathcal{L}(X)}(u) = E_{\mathcal{L}(X)}(v)$.

Proof. If $E_X(u) = E_X(v)$ then for any $(w, w') \in E_{\mathcal{L}(X)}(x)$, we have by definition $wuw' \in \mathcal{L}(X)$ and so there exists a configuration $x \in X$ with $x_{[0,n-1]} = wuw'$. As $E_X(u) = E_X(v)$, we can replace u by v within x to obtain a configuration containing wvw' and so $wvw' \in \mathcal{L}(X)$, and therefore $(w, w') \in E_{\mathcal{L}(X)}(y)$.

The other direction is a simple application of compactness: if $E_{\mathcal{L}(X)}(u) = E_{\mathcal{L}(X)}(v)$, then for all x such that $x \sqcup u \in X$, we can consider for $n \geq 0$ the words $w_n = x_{[-n,-1]}$ and $w'_n = x_{[|u|,|u|+n-1]}$. Then, $w_n u w'_n \in \mathcal{L}(X)$, so $w_n v w'_n \in \mathcal{L}(X)$. Then $(w_n v w'_n)$ converges to the configuration $x \sqcup v$, which by compactness is in X. □

Remark. *In the case of $\mathbb{Z}$ subshifts, we can also immediately extend the notions of predecessor (resp. follower) set $P_X(u)$ (resp. $F_X(u)$) of a word u, as being the set of left-infinite (resp. right-infinite) sequences w such that $wu \in \mathcal{L}(X)$ (resp. $uw \in \mathcal{L}(X)$). As for $E_X(n)$, we write $P_X(n)$ (resp. $F_X(n)$) for $|\{P_X(u), u \in \mathcal{A}^n\}|$ (resp. $|\{F_X(u), u \in \mathcal{A}^n\}|$).*

To some extent, we can view extender sets (and the number of distinct extender sets) as a way to quantify how much information is carried by each pattern, regarding the way it can be extended in a valid configuration. Intuitively, having a small number of extender sets can be seen as the fact that the extensions of patterns are determined only by a small part of the pattern. Although not rigorous, this idea can be formalized (see for example [DR22] for a slightly different point of view), and should in any case be the general idea to keep in mind when trying to construct subshifts with prescribed $(E_X(n))_{n \in N}$.

Unsurprisingly given Proposition 2.5, we have the following proposition for $\mathbb{Z}$-subshifts:

Proposition 2.9

Let $X \subseteq \mathcal{A}^{\mathbb{Z}}$ be a subshift. Then X is sofic if and only if $\{E_{\mathcal{L}}(x), x \in \mathcal{A}^*\}$ is finite.

Proof. X is sofic if and only $\mathcal{L}(X)$ is a regular language by Proposition 1.59, and so Proposition 2.5 gives the result. □

This proposition, as Theorem 2.2, is mainly used to prove that some subshifts are not sofic. However, the theorem relies on properties of regular languages, and sofic subshifts in higher dimensions are not easily definable by their *language*, but as factors of SFTs. It is then quite natural to determine if similar characterizations hold for sofic $\mathbb{Z}^d$-subshifts. In fact, some results are already known. In order to study quantitatively the number of extender sets, it will be convenient to use the following notation:

Notation. Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a subshift, and let $D \subset_f \mathbb{Z}^d$. We denote $E_X(D) = \{E_X(u), u \in \mathcal{A}^D\}$.

For $n \geq 0$, we also denote $E_X(n) = |E_X(\mathcal{Q}_n)| = |E_X([\![0, n-1]\!]^d)|$.

Note that $E_X(n)$ is the cardinal of the set $E_X(\mathcal{Q}_n)$: we will indeed very rarely be interested in what this set is, but rather, the behaviour of the sequence $(E_X(n))_{n\in\mathbb{N}}$.

Theorem 2.10 [OP16, Thm 1.1]

Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a subshift. If there exists n such that $E_X(n) \leq n$, then X is sofic.

An important thing to note here is the fact that the bound $E_X(n) \leq n$ is a very restrictive bound: indeed, the volume of $\mathcal{Q}_n$ in dimension d is n^d, and there are therefore up to 2^{n^d} patterns of this size.

In dimension 1, we also have the following easier result, which improves on Proposition 2.9:

Proposition 2.11 [OP16, Thm 3.4]

Let $X \subseteq \mathcal{A}^{\mathbb{Z}}$ be a subshift. X is sofic if and only if $(E_X(n))_{n\in\mathbb{N}}$ is bounded.

2.1.2 First examples and constructions

We are now ready to compute $E_X(n)$ for some classes of subshifts, in any dimension.

Example 5 (Full-shifts). *Let $X = \mathcal{A}^{\mathbb{Z}^d}$. Then $E_X(n) = 1$ for all n. In particular, despite being of maximal pattern complexity, this subshift has the minimal number of extender sets.*

The situation is also simple in the case of subshifts with very low pattern complexity, in particular the periodic subshifts:

Example 6 (Periodic subshifts). *Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be a periodic subshift, so that there exists $p_1, \ldots, p_d \in \mathbb{Z}$ such that for any $x \in X$ and $1 \leq i \leq d$, $x = \sigma_{p_i \mathbf{e_i}}(x)$. For $n \geq \max_i p_i$, it is therefore clear that for $u \in \mathcal{A}^{\mathcal{Q}_n}$ the set $E_X(u)$ depends only on $u_{[\prod_{i=1}^d [\![0, p_i-1]\!]]}$. Therefore, noting $p = \prod_{i=1}^d p_i$, we have $E_X(n) \leq p^{|A|}$, which in particular does not depend on n.*

Recall that for a subshift X in dimension d, we defined $X^{\Uparrow}$ as the $d+1$-dimensional subshift whose "hyperplanes" in direction e_0 were independent but configurations of X, and $X^{\uparrow}$ the $d+1$-dimensional subshift of $X^{\Uparrow}$ when all those hyperplanes are equal to the same configuration (see Section 1.2.3).

Proposition 2.12

Let X be a $\mathbb{Z}^d$-subshift. Then for any $n > 0$,

$$E_{X^{\Uparrow}}(n) = E_X(n)^n \text{ and } E_{X^{\uparrow}}(n) = E_X(n).$$

Proof. Let $n > 0$ and $u, v \in \mathcal{L}_n(X^{\Uparrow})$. Then

$$E_{X^{\Uparrow}}(u) = E_{X^{\Uparrow}}(v) \iff \forall 0 \leq j < n, E_X\left(u|_{[\![0,n-1]\!]\times\{j\}}\right) = E_X\left(v|_{[\![0,n-1]\!]\times\{j\}}\right)$$

Similarly, for $u, v \in \mathcal{L}_n(X^{\uparrow})$, we have

$$E_{X^{\uparrow}}(u) = E_{X^{\uparrow}}(v) \iff E_X\left(u|_{[\![0,n-1]\!]\times\{0\}}\right) = E_X\left(v|_{[\![0,n-1]\!]\times\{0\}}\right)$$

□

Finally, we show how this notion can be used to distinguish SFT from more complicated subshifts:

Definition 2.13: Border

Let $D \subset_f \mathbb{Z}^d$. Denote d_1 the distance induced by the norm $\|\|_1$, and define the r-border of D as

$$\partial_r D = \{z \in D, d_1(z, \mathbb{Z}^d \setminus D) \leq r\}$$

In particular,

$$\partial_r \mathcal{Q}_n = \{(i_1, \ldots, i_d) \in \mathcal{Q}_n, \exists k, i_k \in [\![0, r-1]\!] \cup [\![n-r, n-1]\!]\}$$

For a pattern u, we might write $\partial_r(u)$ for $u|_{\partial_r \mathrm{dom}(u)}$.

Example 7 (SFT, Folklore, see for example [KM13, Section 2])**.** *Let $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ be an SFT, defined by a family $\mathcal{F}$. Then $E_X(n) = 2^{O(n^{d-1})}$. Without loss of generality, we can assume that $\mathcal{F} \subseteq \mathcal{A}^{\mathcal{Q}_r}$ for some $r \geq 0$. Then, for $n \geq 2r$, consider $u, v \in \mathcal{A}^{\mathcal{Q}_n}$ equal on their r-border. Now, if x is such that $u \sqcup x \in X$, then clearly $v \sqcup x \in X$. Contraposing, we get that $E_X(u) \neq E_X(v) \implies \exists (i,j) \in \partial_r(\mathcal{Q}_n), u_{i,j} \neq v_{i,j}$, and so $E_X(n) \leq \mathcal{A}^{|\partial_r(\mathcal{Q}_n)|^d} \leq |A|^{2rdn^{d-1}}$*

A reader already familiar with usual arguments about sofic shifts might think that this immediately implies that the same inequality holds for sofic subshifts, using the fact that for a factor map $\phi\colon X \to Y$ of radius r we have $E_Y(n) \leq E_X(n+r)$. Unfortunately, such an inequality does not hold in general. We will see in Section 2.2.1 some results of this kind, which relate $E_X(n)$ and $E_{\Phi(X)}(n)$ for some factor map Φ, but the usual arguments applied when dealing with the classical entropy (see Definition 1.23) cannot be applied here.

Of course, the bound of Example 7 is not a characterization: using Proposition 2.12 and the lifting construction, one can construct subshifts which are not SFT, starting from a $\mathbb{Z}$ subshift (possibly not even effective), repeatedly lifting its configurations to obtain d-dimensional subshifts with few extender sets. In particular, we have the following result, which shows that Theorem 2.10 is in a sense optimal:

Proposition 2.14 [OP16, Thm 1.4]

For any $d \geq 1$, there exists a non-sofic subshift X satisfying for all $n \in \mathbb{N}$, $E_X(n) = n + 1$.

Proof. Consider a non-effective Sturmian subshift $X \subseteq \{0,1\}^{\mathbb{Z}}$. Then, as Sturmian subshifts have pattern complexity $n+1$ (see 1.84), we have $E_X(n) \leq n+1$, and as X is not sofic, Theorem 2.10 implies that $E_X(n) = n+1$. We can then consider the subshifts $X_1 = X^{\uparrow}, X_2 = X^{\uparrow^{\uparrow}} \dots$ to obtain subshifts of arbitrary dimensions, and $E_{X_d}(n) = E_X(n) = n+1$ by Proposition 2.12. They are not sofic, for otherwise, by Theorem 1.87, X would be effective. □

We finally give here another construction, from which we draw inspiration in Section 2.4.3 and Section 2.4.5. The idea itself comes from [DR22], where the authors studied a similar problem, not in terms of extender sets but in terms of what they call *epitome*. This is a similar object, although quite technical, and this example is particularly interesting in both settings.

For any $\mathbb{Z}^2$ subshift X over some alphabet $\mathcal{A}$, we define the subshifts $X_{\text{mirror}}, X_{\text{semi-mirror}}$, depicted in Figure 2.1 as follows:

- X_{mirror} is a subshift on alphabet $\mathcal{A} \sqcup \{■\}$, and $X_{\text{semi-mirror}}$ is defined on alphabet $\mathcal{A} \sqcup \{■, □\}$.
- Each configuration x of either subshifts $X_{\text{mirror}}, X_{\text{semi-mirror}}$ might contain an horizontal infinite line of ■, splitting x in an upper and a lower half-plane. Those are the only ■ symbols in x.
- The upper half-plane must be a globally admissible pattern of X.
- In X_{mirror}, the lower half-plane must be the reflection of the upper half-plane by the horizontal ■-row.
- In $X_{\text{semi-mirror}}$, the lower-half plane consists only of □-coloured cells, except for at most one cell coloured with $\mathcal{A}$. This non-□ cell must be the mirror of the corresponding cell in the upper half-plane. Formally, if for some $i, j \in \mathbb{Z}, k > 0$, we have $x_{i,j} = ■$ and $x_{i,j-k} \in \mathcal{A}$, then $x_{i,j-k} = x_{i,j+k}$.

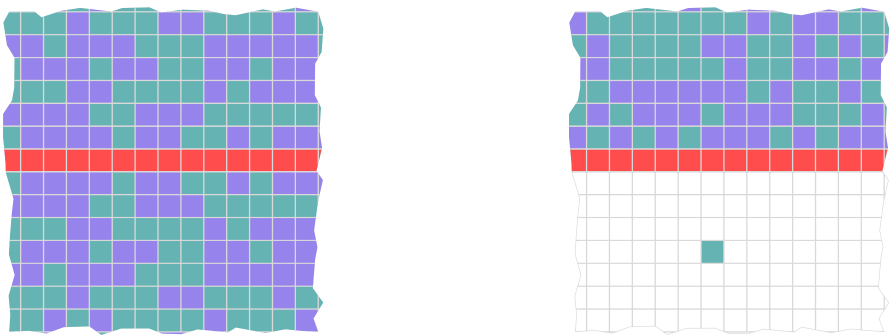

Figure 2.1: Configurations of the mirror and semi-mirror subshifts associated to the binary full-shift. This is [DR22, Example 5″].

The next proposition is somewhat folklore, and its proof illustrates one of the classical techniques used to show that some effective subshifts are not sofic.

Proposition 2.15 [ABJ18, Prop. 57]

Let $Y \subseteq \mathcal{A}^{\mathbb{Z}^2}$ be a subshift satisfying $h(Y) > 0$. Then Y_{mirror} is not sofic.

Proof. We prove this by contradiction. Suppose that Y_{mirror} is sofic, and let then $X \subseteq \mathcal{B}^{\mathbb{Z}^2}$ be a nearest-neighbour SFT and $\Phi\colon X \to Y_{\text{mirror}}$ be a 1-block factor map. Let $D_n = [\![-n, n]\!]^2$. Then, by definition of the entropy, for large enough n and any $0 < h < h(Y)$:

- $|\mathcal{L}_{D_n}(Y)| \geq 2^{h(2n+1)^2}$
- $|\partial_1 D_n| = 8n$ and so $|X| \leq |\mathcal{B}|^{8n}$

As $Y \subset Y_{\text{mirror}}$, there are therefore at least $2^{h(2n+1)^2}$ patterns of support D_n in Y_{mirror} that do not contain a symbol ■. Let $S \subset \mathcal{L}_{D_n}(Y_{\text{mirror}})$ be this set of patterns. Note that if $u, v \in \mathcal{L}_{D_n}(X)$ are such that $\partial_1(u) = \partial_1(v)$, as X is a nearest-neighbour SFT, we have $E_X(u) = E_X(v)$. For n large enough, $|S| \geq 2^{h(2n+1)^2} \geq |\mathcal{B}|^{8n}$ and so there exists $u, v \in X$ such that $\Phi(u), \Phi(v) \in S$, but $\Phi(u) \neq \Phi(v)$, and $\partial_1(u) = \partial_1(v)$. Choosing any configuration $x \in X$ containing u and such that $\Phi(x)$ contains a ■-row, as $E_X(u) = E_X(v)$, we can replace the occurrence of u in x by v to obtain another configuration $x' \in X$. Now, as $\Phi(u) \neq \Phi(v)$, and Φ being a 1-block map, $\Phi(x') \notin Y_{\text{mirror}}$ as the mirror is broken where we replaced u by v. □

We see that this proof relies on similar ideas than those formalized by extender sets: the obstruction to some Y being sofic is the fact that there are too few extender sets in any SFT cover X, which means that the "rest" of a configuration $x \in X$ cannot distinguish between several patterns $u, v \in E_X(w)$ and must be valid for all (or none) of them. In the case of positive entropy, this means in particular that there exists multiple patterns with the same extender set but different images under the hypothetical factor map.

However, this arguments breaks if we slightly relax the mirroring condition:

Proposition 2.16

Let Y be a sofic $\mathbb{Z}^2$ subshift. Then $Y_{\text{semi-mirror}}$ is sofic. Furthermore, if $u \neq v \in \mathcal{L}(X)$ then $E_{Y_{\text{semi-mirror}}}(u) \neq E_{Y_{\text{semi-mirror}}}(v)$.

Proof. The first part of the proposition is easy, using standard constructions from symbolic dynamics. For example, using variants of the sunny side up subshift (see Example 1), one can mark specific positions in a configuration, and so ensure that at most one cell is mirrored, or only one "mirror" row appears in a configuration.

Now, let $n \geq 0$ and $u \neq v \in \mathcal{L}_n(Y)$. One can use the reflected cell to pinpoint any difference between u and v: considering a configuration $y \in Y_{\text{semi-mirror}}$ containing u, such that the mirrored cell is the cell $u_{i,j}$ where $u_{i,j} \neq v_{i,j}$, the mirroring condition then ensures that we cannot replace u by v in y, and so $E_{Y_{\text{semi-mirror}}}(u) \neq E_{Y_{\text{semi-mirror}}}(v)$. □

2.2 Another kind of entropy

The sequence $(E_X(n))_{n \in \mathbb{N}}$ can be a complicated and counter-intuitive object, even in the case of sofic $\mathbb{Z}$-subshifts. In particular, let us mention some properties that might be unexpected:

Proposition 2.17 [OP16, Ex. 3.5]

There exists a sofic $\mathbb{Z}$ subshift X such that for all n, $E_X(2n) = 46$ and $E_X(2n+1) = 44$. In particular, $E_X(n)$ is periodic and is not non-decreasing.

This example is attributed to Martin Delacourt. More generally, we have the stronger theorem:

Theorem 2.18 [Fre16a, Thm. 1.3]

Let $n \geq 0$. For any partition $\bigsqcup_{i=1}^{k} A_i = [\![0, n-1]\!]$ and any sequence $0 = r_1 < \cdots < r_k \in \mathbb{N}$, there exists $m \geq 0$ and X a sofic $\mathbb{Z}$-subshift such that for all sufficiently large ℓ, when $\ell \mod n \in A_j$, we have $E_X(\ell) = m + r_j$.

More precise quantitative bounds are given in the article [Fre16a, Thm. 1.3] and are not reproduced here: the important remark is that pretty much any periodic sequence can be realized as the sequence $(E_X(n))$ for some sofic subshift X. This motivates the introduction of a simpler, although less precise, quantity to study the behaviour of extender sets in any subshift. We show in this section how to define a conjugacy invariant based on the sequence $E_X(n)$, called the **extender entropy**.

2.2.1 Extender entropy: a conjugacy invariant

Extender entropy has already been defined and studied for $\mathbb{Z}$ subshifts in [FP19, Def 2.17]. We show here how to adapt the definition from one-dimensional to multi-dimensional subshifts.

Definition 2.19: Extender entropy [FP19, Def 2.17]

Let X be a $\mathbb{Z}^d$-subshift. We call **extender entropy** of X the real

$$h_E(x) = \lim_{n \to +\infty} \frac{\log E_X(n)}{n^d}$$

Before stating properties about the extender entropy, we need to show that it actually exists. It was already proven for $\mathbb{Z}$-subshift by [FP19, Theorem 3.1], and we extend the argument to arbitrary subshifts:

Proposition 2.20

For any non-empty subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$,

$$\lim_{n \to +\infty} \frac{\log E_X(n)}{n^d} = \inf_{n \in \mathbb{N}} \frac{\log E_X(n)}{n^d}$$

In particular, $h_E(X)$ exists.

Proof. The proof relies on the subadditivity (or Fekete's) lemma. The usual strategy is to prove that the sequence $f = (\frac{\log E_X(n)}{n^d})_{n\in\mathbb{N}}$ satisfies $f_{m+n} \leq f_m + f_n$, or equivalently, $E_X(m+n)^{(m+n)^d} \leq E_X(m)^{m^d} E_X(n)^{n^d}$. However, we do not know if this holds in general. We can nevertheless follow the usual proof of the subadditivity lemma, as $E_X(n)$ is still regular enough to obtain a similar inequality.

Fix some $k, n \in \mathbb{N}$ and write $n = kq + r$ the Euclidean division of n by k. Define the map

$$\begin{aligned}\phi\colon \mathcal{L}_n(X) &\to E_X(k)^{q^d} \times \mathcal{A}^{\mathcal{Q}_n \setminus \mathcal{Q}_{kq}} \\ w &\mapsto \left(\prod_{z\in\mathcal{Q}_{kq}} E_X(w|_{\mathcal{Q}_k + kz}\right) \times w|_{\mathcal{Q}_n\setminus\mathcal{Q}_{kq}}\end{aligned}$$

Ideally, we would like to partition a square of side n into squares of size k, and relate the extender set of the larger pattern and the extender sets of all the smaller pattern. What ϕ does is precisely this "partition", in the case where k does not divide n, so there remains some small rectangular strips not coverered by the k-squares within $\mathcal{Q}_n$, as illustrated in Figure 2.2.

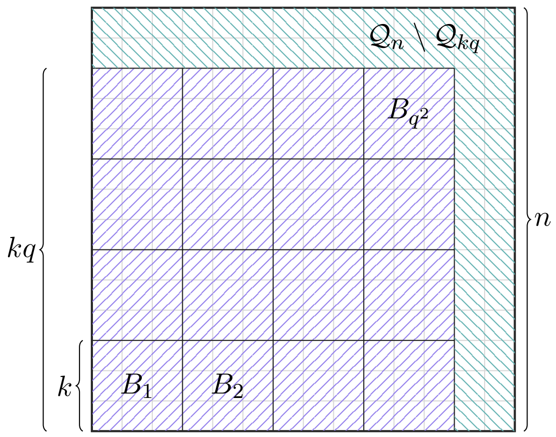

Figure 2.2: Example in dimension $d = 2$ for $n = 14$ and $k = 3$, so $n = 3 \times 4 + 2$: an $n \times n$ square contains $q^2 = 16$ squares of size 3×3.

We now claim that for any patterns $u, v \in \mathcal{A}^{\mathcal{Q}_n}$, we have $\phi(u) = \phi(v) \implies E_X(u) = E_X(v)$. This is clear if we interpret the equality of extender sets as meaning that two patterns can be freely exchanged in any configuration. Indeed, let now x be such that $x \sqcup u \in X$. Call (i,j)-block the domain $\mathcal{Q}_k + k(i,j)$, and fix any ordering of the blocks $B_1, \dots, B_{q^d}$. Denote $w^\ell \in \mathcal{A}^{\mathcal{Q}_n}$ the pattern equal to u on the blocks $B_{\ell'}$ for $\ell' \leq_{\text{lex}} \ell$ and equal to v elsewhere. In particular, as $\phi(u) = \phi(v)$, we have $w_0 = v$ and $w_{q^d} = u$. Consider x such that $x \sqcup v \in X$. Then, by $\phi(u) = \phi(v)$, we have $E_X(v|_{B_1} = u|_{B_1})$. Therefore, $x \sqcup w^1 \in X$, and $\phi(v) = \phi(w^1)$. We can therefore repeat the argument q^d times, and we obtain $x \sqcup w^n = x \sqcup u \in X$. This gives $E_X(v) \subseteq E_X(u)$, and so by symmetry $E_X(v) = E_X(u)$.

Now, as $n = kq + r$ is a Euclidean division, we have $0 \leq r < k$ and so $|\mathcal{Q}_n \setminus \mathcal{Q}_{kq}| \leq drn^{d-1} \leq dkn^{d-1}$. As we have shown that ϕ factored through E_X, we get

$$E_X(n) \leq E_X(k)^{q^d} \times \mathcal{A}^{dkn^{d-1}}$$

Hence,

$$\begin{aligned}\frac{\log E_X(n)}{n^d} &\le \frac{q^d E_X(k)}{n^d} + \frac{O(n^{d-1})}{n^d}\\ &\le \frac{\log E_X(k)}{k^d} + O(\frac{1}{n})\end{aligned}$$

As n goes to $+\infty$ we obtain $\limsup_n \frac{\log E_X(n)}{n^d} \le \frac{\log E_X(k)}{k^d}$, and taking the infimum over k gives $\limsup_n \frac{\log E_X(n)}{n^d} \le \inf_k \frac{\log E_X(k)}{k^d}$. But clearly the inequality holds in the other direction, so $\frac{\log E_X(n)}{n^d}$ converges to $\inf \frac{\log E_X(n)}{n^d}$. □

In fact, the previous proof can be adapted to prove a slightly stronger result. We state it and give a sketch of the proof, the details being identical to the case of Proposition 2.20.

Lemma 2.21

Let $r = \frac{a}{b} \in \mathbb{Q}$ and X any $\mathbb{Z}^2$-subshift. Then, denoting $E_X(am, bm) = |E_X([\![0, am-1]\!] \times [\![0, bm-1]\!])|$, we have:

$$lim_{m\to+\infty} \frac{\log E_X(am, bm)}{abm^2} = h_E(X)$$

Sketch of the proof of Lemma 2.21. We follow the strategy of the proof of Proposition 2.20, using rectangles instead of squares.

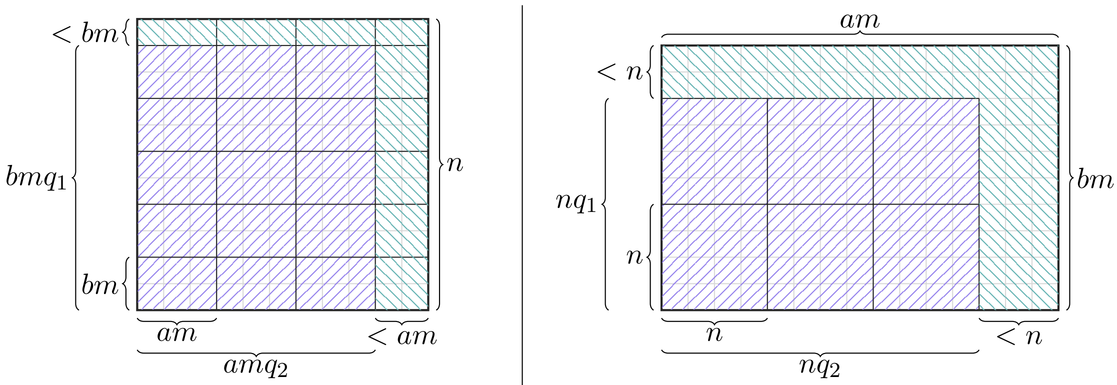

(a) Tiling a square of side $n = 11$ by a rectangle $am \times bm$, $m = 1$

(b) Tiling a rectangle $am \times bm$ with $m = 5$ by a square of side $n = 4$

Figure 2.3: Example for a rectangle of ratio $r = \frac{a}{b}$, with $a = 3, b = 2$

Figure 2.3 shows how we adapt the previous proof to deal with rectangles. What is fixed in this construction is the ratio $r = \frac{a}{b}$. The same proof gives the following inequalities for all $m, n \ge 1$:

$$\frac{\log E_X(n)}{n^2} \le \frac{\log E_X(am, bm)}{abm^2} + O(\frac{1}{n})$$

and

$$\frac{\log E_X(am, bm)}{abm^2} \le \frac{\log E_X(n)}{n^2} + O(\frac{1}{n})$$

As we already know that $\frac{\log E_X(n)}{n^2} \to h_E(X)$ as $n \to +\infty$, we obtain

$$\lim_{m\to+\infty} \frac{\log E_X(am, bm)}{abm^2} = h_E(X)$$

We give another useful way to write this equality. Denoting $A_m = abm^2$ the area of the rectangle, we can write $E_X(am, bm) = 2^{h_E(X)A_m + o_m(m^2)}$. $\square$

Now that h_E is proven to be well-defined for any subshift in any dimension, we furthermore claim that it is a conjugacy invariant. As before, it is already proven in the case of $\mathbb{Z}$ subshift in [FP19, Theorem 3]. The geometry of $\mathbb{Z}^d$ makes the proof slightly more heavy in terms of notations, but the general strategy is exactly the same.

Proposition 2.22

Let X be a $\mathbb{Z}^d$ subshift, and $\Phi_r \colon X \to X_r$ the natural conjugacy from X to its r-higher-block code X_r. Then $h_E(X) = h_E(X_r)$.

This is a consequence of the following lemma:

Lemma 2.23

For $n \geq r$ and $u, v \in \mathcal{L}_{\mathcal{B}_n}(X)$, $E_{X_r}(\Phi_r(u)) = E_{X_r}(\Phi_r(v))$ if and only if $E_X(u) = E_X(v)$ and furthermore $\partial_{2r}(u) = \partial_{2r}(v)$.

Proof of Lemma 2.23. We prove both directions:

$\Rightarrow$: Let $u, v \in \mathcal{L}_{\mathcal{B}_n}(X)$, and suppose that $E_{X_r}(\Phi_r(u)) = E_{X_r}(\Phi_r(v))$. $\Phi_r(u), \Phi_r(v)$ are patterns of support $\mathcal{B}_{n-r}$. Let then $w \in E_{X_r}(\Phi_r(u))$, and let $\mathbf{z} \in \partial_1(\mathcal{B}_{n-r+1})$, that is, a point adjacent to $\mathrm{dom}(\Phi_r(u))$, neighbour to some $\mathbf{z}' \in \mathcal{B}_{n-r}$ for the $\|\|_1$ norm. By definition of Φ_r, $w_{\mathbf{z}}$ and $\Phi_r(u)_{\mathbf{z}'}$ must coincide (when viewed as patterns of support $\mathcal{B}_r$) on a set of the form $[\![-r, r-1]\!] \times [\![-r, r]\!]^{d-1}$. The same holds for $\Phi_r(v)_{\mathbf{z}'}$, as $w \in E_{X_r}(\Phi_r(v))$. This in turn implies that u and v coincide on $\partial_{2r}(\mathcal{B}_n)$, as $\Phi_r(v)_{\mathbf{z}'} = v|_{\mathbf{z}'+\mathcal{B}_r}$.

Now, for any $w \in E_X(u)$, we show that $w \in E_X(v)$, which by symmetry gives the equality $E_X(u) = E_X(v)$. As $w \sqcup u \in X$, we have $\Phi_r(w \sqcup u) \in X_r$, and we can decompose $\Phi_r(w \sqcup u) = \Phi_r(w) \sqcup s\Phi_r(u)$, with up to translation $\mathrm{dom}(s) = \mathcal{B}_{n+r} \setminus \mathcal{B}_{n-r}$, and s depends only on $u|_{\partial_r \mathcal{Q}_n}$ and $w|_{\partial_r \mathcal{Q}_{n+r}}$. By the previous point, $u|_{\partial_r \mathcal{Q}_n} = v|_{\partial_r \mathcal{Q}_n}$ and so $\Phi_r(w \sqcup v) = \Phi_r(w) \sqcup s \sqcup \Phi_r(v)$. As $E_{X_r}(\Phi_r(u)) = E_{X_r}(\Phi_r(v))$, this means that $\Phi_r(w \sqcup v) \in X_r$, which by definition of Φ_r implies $w \sqcup v \in X$.

$\Leftarrow$: Suppose that $E_X(u) = E_X(v)$ and furthermore $\partial_{2r}(u) = \partial_{2r}(v)$. Consider any $w \in E_{X_r}(\Phi_r(u))$. By definition, $w \sqcup \Phi_r(u) \in X_r$ and so $\Phi_r^{-1}(w \sqcup \Phi_r(u)) \in X$. We can write $\Phi_r^{-1}(w \sqcup \Phi_r(u)) = \overline{w} \sqcup u$. By definition of Φ_r^{-1}, $\overline{w}|_{\partial_r \mathcal{B}_{n+r}}$ depends only on w and $\partial_{2r}(u)$. Therefore, as $\partial_{2r}(u) = \partial_{2r}(v)$, we have $\Phi_r^{-1}(w \sqcup \Phi_r(v)) = \overline{w} \sqcup v$, and as $E_X(u) = E_X(v)$, this means that $\overline{w} \in E_X(v)$ and so $w \sqcup \Phi_r(v) \in X_r$, which means that $E_{X_r}(\Phi_r(u)) \subseteq E_{X_r}(\Phi_r(v))$, and by symmetry we have equality. $\square$

Proof of Proposition 2.22. By Lemma 2.23, we have that for $n > 4r$, $E_X(n) \leq E_{X_r}(n - 2r) \leq E_X(n)|A|^{2drn^{d-1}}$. As r is fixed, we obtain $h_E(X) = h_E(X_r)$. $\square$

Lemma 2.24

Let X, Y be subshifts conjugate via a block map $\Phi \colon X \to Y$ of radius 0. Let s be the radius of Φ^{-1}. Then, for all $u, v \in \mathcal{L}_{\mathcal{B}_n}(X)$ such that $E_X(u) = E_X(v)$, and any $w \in E_{\mathcal{L}(X)}(u)$ of support $\partial_s(\mathcal{B}_{n+s})$, we have $E_Y(\Phi(w \sqcup u)) = E_Y(\Phi(w \sqcup v))$.

Proof. Observe that

$$E_Y(\Phi(w \sqcup u)) = \bigcup_{x \in \Phi^{-1}(\Phi(w \sqcup u))} \Phi(E_X(x)) = \bigcup_{\overline{w} \in \Phi^{-1}(w), \overline{w} \sqcup u \in \mathcal{L}(X)} \Phi(E_X(\overline{w} \sqcup u))$$

The first equality always holds and is a set-theoretic equality; the second equality comes from the fact that if $x \in \Phi^{-1}(\Phi(w \sqcup u))$, then as Φ^{-1} has radius s and w has "width" s, one must have $x|_{\text{dom}(u)} = u$. We assumed that $E_X(u) = E_X(v)$, therefore $\{\overline{w} \in \Phi^{-1}(w), \overline{w} \sqcup u \in \mathcal{L}(X)\} = \{\overline{w} \in \Phi^{-1}(w), \overline{w} \sqcup v \in \mathcal{L}(X)\}$, hence $E_X(w \sqcup u) = E_X(w \sqcup v)$, which finally gives $E_Y(\Phi(u)) = E_Y(\Phi(v))$. □

We can now generalize [FP19, Thm 3] to arbitrary subshifts:

Theorem 2.25 [FP19, Thm 3]

Extender entropy is a conjugacy invariant.

Proof. Suppose that X and Y are $\mathbb{Z}^d$-subshifts, conjugate *via* $\phi \colon X \to Y$. By Proposition 2.22, we can assume that ϕ has radius 0, up to taking a higher block code of X. Let s be the radius of ϕ^{-1}. By Lemma 2.24, we have that for $n > 2s$, $E_Y(n) \leq E_X(n-2s)|A|^{2dsn^{d-1}}$. As s is fixed, we get $h_E(Y) \leq h_E(X)$. Applying this result to the conjugacy $\phi^{-1} \colon Y \to X$ instead, we obtain the reverse inequality and finally $h_E(X) = h_E(Y)$. □

2.2.2 Preliminary results on extender entropies

We state here without several other results of [FP19] and easy corollaries of our previous remarks of Section 2.1.2.

Theorem 2.26 [FP19, Thm 4.3]

For any $x \leq y$, there exists a $\mathbb{Z}$-subshift X satisfying $h_E(X) = x, h(X) = y$.

Note that this theorem says nothing about the dynamical or computational properties of X. In particular, X is never assumed to be effective. In fact, the proof shows that if x, y are both Π_1 real numbers, then X is effective. Using Proposition 2.12, we can then immediately extend the result to any dimension $d \geq 1$.

Theorem 2.27 [FP19, Thm 3.2]

For any $\mathbb{Z}^d$-subshifts X, Y, we have $h_E(X \times Y) = h_E(X) + h_E(Y)$.

Finally, we have the immediate corollaries of the previous sections:

Corollary 2.28

Let X be a $\mathbb{Z}^d$ SFT. Then $h_E(X) = 0$.

Proof. This is Example 7. □

Corollary 2.29

Let X be a $\mathbb{Z}$-sofic subshift. Then $h_E(X) = 0$.

Proof. If X is sofic, then $(E_X(n))_{n\in\mathbb{N}}$ is bounded by Proposition 2.11, and so $h_E(X) = 0$. □

However, we also have some positive results, showing how to construct a large class of real numbers as extender entropies:

Proposition 2.30

Let X be a sofic $\mathbb{Z}^2$ subshift. Then there exists Y a sofic $\mathbb{Z}^2$ subshift such that $h_E(Y) = h(X)$.

Proof. Let $\mathcal{A}$ be the alphabet of X. We can always assume that $h(X) > 0$, as otherwise taking any SFT Y works by Corollary 2.28. By Proposition 2.16, we know that $Y = X_{\text{semi-mirror}}$ is sofic and satisfies $h_E(Y) \geq h(X)$. For the other inequality, we simply show that $h(Y) = h(X)$. Then, as $h_E(Y) \leq h(Y)$, we get equality.

Any pattern of Y might contain:

- A row of ■, at some height $0 \leq m < n$.
- A pattern of $\mathcal{L}_{n\times(n-m-1)}(X)$, which we simply approximate by a pattern of $\mathcal{L}_n(X)$.
- A possible reflected cell, with colour in $\mathcal{A}$ and position in $\mathcal{Q}_n$.

In total, we get $|\mathcal{L}_n(Y)| \leq |\mathcal{A}|n^3|\mathcal{L}_n(X)|$, and so as $\frac{log(|\mathcal{A}|n^3)}{n^2} \to 0$, we get $h(Y) = h(X)$. □

In particular, we obtain the following corollary:

Corollary 2.31

Every Π_1 real number is the extender entropy of some sofic $\mathbb{Z}^2$-subshift.

Proof. Every Π_1 real number is the entropy of a sofic $\mathbb{Z}^2$-subshift (in fact, even of an SFT) by [HM10], and we conclude by Proposition 2.30. □

We will also be able to transfer results from construction of subshifts on $\mathbb{Z}$ or $\mathbb{Z}^2$ to higher dimensions:

Proposition 2.32

Let X be a $\mathbb{Z}^d$ effective (resp. sofic, computable) subshift. Then for all $d' > d$ there exists an effective (resp. sofic, computable) $\mathbb{Z}^{d'}$-subshift Y such that $h_E(Y) = h_E(X)$

Proof. This is a simple application of Proposition 2.12, noticing that if X is effective (resp. sofic, computable) then so is $X^{\Uparrow}$. □

From now on, our constructions will therefore be done for $\mathbb{Z}$ and $\mathbb{Z}^2$ subshifts.

2.3 Computability considerations

In this section, we now show that the class of real numbers that can be realized as extender entropies of sofic or effective subshifts is much larger than the Π_1 class that Corollary 2.31 shows how to construct. We will in fact provide a complete characterization.

2.3.1 Inclusion of extender sets

In a first step, we prove an upper bound on the set of realizable extender entropies, of effective, computable, and sofic subshifts. In order to do this, we will define some decision problems, the difficulty of which will correspond to classes in the arithmetic hierarchy of real numbers.

Decision Problem — EXTENDER-INCLUSION

Input: An effective subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$, and $n \geq 0$, $u, v \in \mathcal{L}_n(X)$.
Output: Whether $E_X(u) \subseteq E_X(v)$.

Proposition 2.33

EXTENDER-INCLUSION is Π_2^0-complete.

In order to prove Proposition 2.33, we will reduce EXTENDER-INCLUSION to another classical problem.

Decision Problem — DET-REC-STATE

Input: A deterministic Turing Machine M and a state q.
Output: Does M visit q infinitely often during its run starting from the empty input ?

This problem is known to be Π_2^0-complete. It is a simple reformulation of INF (the problem of determining if a machine halts on infinitely many inputs), detailed for example in [Soa16, Theorem 4.3.2].

Proof of Proposition 2.33. We prove both directions:

Proof that EXTENDER-INCLUSION is a Π_2^0 problem: By definition, if $\mathrm{dom}(u) = \mathrm{dom}(v)$ then:

$$E_X(u) \subseteq E_X(v) \iff \forall D \subset_f \mathbb{Z}^d \setminus \mathrm{dom}(u), \forall B \in \mathcal{A}^D, B \sqcup u \notin \mathcal{L}(X) \vee B \sqcup v \in \mathcal{L}(X)$$

If X is effective, then testing if a pattern belongs to $\mathcal{L}(X)$ is a Π_1^0 problem, so testing $B \sqcup u \notin \mathcal{L}(X)$ is a Σ_1^0 problem. Therefore, testing it for all $B \in \mathcal{A}^*$ is a Π_2^0 problem.

Proof that EXTENDER-INCLUSION is Π_2^0-hard: It suffices to show it for $\mathbb{Z}$-subshifts. We reduce the problem to DET-REC-STATE. Let (M, q) some instance of DET-REC-STATE, and define the following subshift X_M:

- X_M is a $\mathbb{Z}$-subshift over the alphabet $\{0, 1, \square\}$.
- 0 and 1 cannot appear in the same configuration.
- 1 appears at most once in a configuration.
- If 0 appears twice in a configuration x, so that there exists i, n with $x_i = x_{i+n} = 0$, then x is n-periodic. Moreover, if M visits the state q at least m times, then we enforce that $n \geq m$.

This is an effective subshift: the constraints on 0, 1 and on the periodicity are clearly effective, and we can enumerate the forbidden small periods by simply executing the machine M, and forbidding additional patterns each time it enters the state q.

The instance of EXTENDER-INCLUSION that we consider is then $(X, 0, 1)$. Indeed, we claim that $E_X(0) \subseteq E_X(1)$ if and only if M enters q infinitely often starting from the empty input: the symbol 0 either appears in a configuration ${}^{\infty}\square 0 \square^{\infty}$, or in a periodic configuration ${}^{\infty}(0\square^n)^{\infty}$. The first kind of configurations also belong to $E_X(1)$, but not the second, as 0s and 1s cannot occur in the same configuration. However, those n-periodic configurations exist exactly when the machine visits q less than n times. ☐

Note that, although not crucial in the previous proof, we rely on the observations made in the case of periodic subshifts in Example 6: ensuring that configurations are periodic is a straightforward way to enforce that different patterns u, v have different extender sets, as replacing a single occurrence of u by v makes the configuration no longer periodic.

Proposition 2.34

If we restrict the possible input to computable subshifts, then EXTENDER-INCLUSION is Π_1^0-complete.

Proof. Both directions are proven as in the case of effective subshift in Proposition 2.33:

Proof that EXTENDER-INCLUSION is a Π_1^0 problem for computable subshifts: We still write

$$E_X(u) \subseteq E_X(v) \iff \forall D \subset_f \mathbb{Z}^d \setminus \mathrm{dom}(u), \forall B \in \mathcal{A}^D, B \sqcup u \notin \mathcal{L}(X) \vee B \sqcup v \in \mathcal{L}(X)$$

Now, as X is computable, testing for any particular B whether $B \sqcup u \notin \mathcal{L}(X)$ is decidable, and so testing it for all $B \in \mathcal{A}^*$ is a Π_1^0 problem.

Proof that EXTENDER-INCLUSION is Π_1^0-hard for computable subshifts: We define another problem to which we reduce EXTENDER-INCLUSION:

Decision Problem CO-HALT

Input: A Turing Machine M.
Output: Whether M runs forever on the empty input.

As the complement of Halt, this problem is Π_1^0.

The reduction and the proof are then exactly the same as in Proposition 2.33, where we instead define the subshift X_M by forbidding periods less than n if M runs for at least n steps. This is a computable subshift, as it suffices to run M for n steps to decide whether or not any given pattern is allowed. □

2.3.2 Number of extender sets

The Section 1.2.2 shows how to relate decision problems and computability properties of real numbers. We can therefore obtain restrictions on the reals that might be realized as extender entropies:

Decision Problem — Extender-count

Input: An effective $\mathbb{Z}^d$-subshift X and $k, n \geq 0$.
Output: Whether $k \leq E_X(n)$.

Lemma 2.35

Extender-count is a Σ_2^0 problem.

Proof. We simply write a first-order formula $\phi(X, k, n)$ which is true if and only if $k \leq E_X(n)$:

$$\phi(X, k, n) = \exists v_1, \ldots, v_k \in \mathcal{L}_n(X), \bigwedge_{1 \leq i < j \leq k} E_X(v_i) \neq E_X(v_j)$$

For any v_i, v_j, the proposition $E_X(v_i) \neq E_X(v_j)$ can simply be rewritten as $\neg(E_X(v_i) \subseteq E_X(v_j) \wedge E_X(v_j) \subseteq E_X(v_i))$, which by Proposition 2.33 is a Σ_2^0 problem. We can then rewrite

$$\phi(X, k, n) = \exists v_1, \ldots, v_k \in \mathcal{A}^*, \underbrace{\bigwedge_{1 \leq i \leq k} v_i \in \mathcal{L}_n(X)}_{\Pi_1^0 \subset \Sigma_2^0} \wedge \underbrace{\bigwedge_{1 \leq i < j \leq k} E_X(v_i) \neq E_X(v_j)}_{\Sigma_2^0}$$

Moreover, for any Σ_2^0 formula ψ, the formula $\exists x \psi$ is still Σ_2^0. Hence, $\phi(X, k, n) \in \Sigma_2^0$. □

Proposition 2.36

Let X be an effective $\mathbb{Z}^d$-subshift. Then $h_E(X) \in \Pi_3$.

Proof. Given $X, n \geq 0$, the *set* $\{k \in \mathbb{N}, k \leq E_X(n)\}$ is a Σ_2^0 set by Lemma 2.35. This immediately implies that $\{x \in \mathbb{R}, x \leq \frac{\log E_X(n)}{n^d}\}$ is a Σ_2^0 set as logarithms and quotients are computable functions. By definition, this means that for all n, $\frac{\log E_X(n)}{n^d} \in \Sigma_2$. As $h_E(X) = \inf_n \frac{\log E_X(n)}{n^d}$ by Proposition 2.20, we get $h_E(X) \in \Pi_3$. □

Proposition 2.37

Let X be a computable $\mathbb{Z}^d$-subshift. Then $h_E(X) \in \Pi_2$.

Proof. The proof is the same as for Proposition 2.36. We simply show that if X is computable, then for any n the set $\{k \in \mathbb{N}, k \leq E_X(n)\}$ is a Σ_2^0 set. Indeed, as in the proof of Lemma 2.35, we can write a formula $\phi(X,k,n)$ which holds if and only if $k \leq E_X(n)$ as:

$$\phi(X,k,n) = \exists v_1,\dots,v_k \in \mathcal{A}^*, \underbrace{\bigwedge_{1\leq i\leq k} v_i \in \mathcal{L}_n(X)}_{\text{decidable}} \wedge \underbrace{\bigwedge_{1\leq i<j\leq k} E_X(v_i) \neq E_X(v_j)}_{\Sigma_1^0}$$

and therefore $\phi(X,k,n)$ is a Σ_1^0 formula. □

2.4 Characterizations of extender entropies

Before moving on to the proof of our main results about the extender entropy of sofic and effective subshifts in general, we prove some easier results, about the possible behaviours of the extender entropy when the subshifts satisfy additional dynamical properties.

2.4.1 Minimal subshifts

Recall from Proposition 1.39 that a minimal subshift is a subshift in which no additional pattern can be forbidden, for otherwise it becomes empty. From this characterization, we deduce a nice result on the extender sets of minimal subshifts:

Proposition 2.38

Let X be a minimal subshift. For any $n \geq 0$ and $u, v \in \mathcal{L}_n(X)$, $E_X(u) \subseteq E_X(v) \implies u = v$.

Proof. We show by contradiction that $E_X(u) \subseteq E_X(v)$ and $u \neq v$ imply that X is not minimal. More precisely, if $\mathcal{F}$ is a family of forbidden patterns defining X, we prove that $\mathcal{X}_{\mathcal{F}\cup\{u\}} \neq \emptyset$. Indeed, let $x \in X$ be some configuration. As $E_X(u) \subseteq E_X(v)$, we can replace any occurrence of u by v in x, and construct for any $n \geq 0$ configurations x_n such that $u \not\sqsubseteq x_n|_{\mathcal{B}_n^d}$. By compactness, we obtain at the limit a configuration $x_\infty \in X$ containing no u, so $x_\infty \in \mathcal{X}_{\mathcal{F}\cup\{u\}}$, which is a contradiction. □

Corollary 2.39

Let X be a $\mathbb{Z}^d$-minimal subshift. Then $h(X) = h_E(X)$.

This means that any result on the entropy of minimal subshifts immediately gives the same result about its extender entropy. In particular, we have the following results as immediate consequences from previous theorems of the literature:

Corollary 2.40

If X is a minimal sofic subshift, then $h_E(X) = 0$.

Proof. This is because for a $\mathbb{Z}^d$-sofic subshift X, minimality implies that X has zero entropy (folklore, see for example [Gan18, Proposition 6.1], in which one can find other results about how dynamical properties constrain the possible values of the entropy). □

Proposition 2.41

The extender entropies of minimal $\mathbb{Z}^d$ effective subshifts are exactly the Π_1 real numbers.

Proof. This is a consequence of the fact that the *entropies* of minimal effective subshifts are the Π_1 real numbers, see for example [Kur03, Theorem 4.77] for a proof. □

2.4.2 Mixing properties

Recall from Definition 1.43 that a $\mathbb{Z}$-subshift is **mixing** if for any two patterns $u, v \in \mathcal{L}_n(X)$, there exists some $N \geq 0$ such that for any $z \in \mathbb{Z}, |z| \geq N$, there exists some configuration x such that $x|_{\mathcal{Q}_n} = u$ and $x|_{z+n+\mathcal{Q}_n} = v$ (this is a slight but equivalent reformulation of the definition given in Definition 1.43). If N can be chosen independently from n, we say that X is N-mixing. Less formally, mixing means that provided that we place them sufficiently far, any patterns u and v can appear in the same configuration. Such a property could *a priori* prevent some behaviours in $(E_X(n))_{n\in\mathbb{N}}$: intuitively, being mixing means that little information "escapes" from the pattern to determine the rest of the configuration, and the only influence exerted by the pattern is in a bounded neighbourhood around it. This is formally incorrect, and this intuition is in fact false:

Proposition 2.42

Let X be a $\mathbb{Z}$-subshift. Then there exists a $\mathbb{Z}$-subshift Y which is 1-mixing, and such that $h_E(X) = h_E(Y)$.
If X is effective, then Y can be chosen effective too.

Proof. Let $X = \mathcal{X}_\mathcal{F} \subseteq \mathcal{A}^\mathbb{Z}$ be defined by a family of forbidden patterns $\mathcal{F} \subset \mathcal{A}^*$. We can assume that $\mathcal{F} = \mathcal{A}^* \setminus \mathcal{L}(X)$. In particular, we assume that the domain of each $f \in \mathcal{F}$ is an interval, and for any $u \in \mathcal{A}^*$, if u contains no element of $\mathcal{F}$ then $u \in \mathcal{L}(X)$. Consider the alphabet $\mathcal{B} = \mathcal{A} \sqcup \{\#\}$, and let Y be the $\mathbb{Z}$-subshift with alphabet $\mathcal{B}$ and the same family of forbidden patterns $\mathcal{F}$. Then, $y \in Y$ contains arbitrary (possibly empty, possibly infinite) words of $\mathcal{L}(X)$, separated by $\#$. Now, clearly Y is 1-mixing as for any words $u, v \in \mathcal{L}(Y)$, we have $u\#v \in \mathcal{L}(Y)$. Moreover, Y is clearly effective if X is effective.

It remains to show that $h_E(X) = h_E(Y)$.

Claim 1. For any $n \geq 0$, the following inequality holds:

$$E_X(n) \leq E_Y(n) \leq E_X(n) + \sum_{0\leq i,j\leq n, i+j<n} P_X(i)F_X(j)$$

.

Proof. We prove the two inequalities separately:

- For the leftmost inequality, notice that $X \subseteq Y$, and in particular for any $u, v \in \mathcal{L}_n(X)$ and x such that $x \sqcup u \in X, x \sqcup v \notin X$, we have $x \sqcup u \in Y$ but still $x \sqcup v \notin Y$. More informally, a configuration "witnessing" that $E_X(u) \neq E_X(v)$ also proves that $E_Y(u) \neq E_Y(v)$.

- For the rightmost inequality, let $u, v \in \mathcal{L}_n(Y)$:

  - If $\# \notin u$ and $\# \notin v$, then we have $E_Y(u) = E_Y(v) \iff E_X(u) = E_X(v)$. Indeed, by assumption on $\mathcal{F}$ this means that $u, v \in \mathcal{L}(X)$, and if $E_Y(u) = E_Y(v)$ then in particular any $x \in E_Y(u)$ not containing $\#$ is also an extender for u in X. Therefore, $E_X(u) = E_X(v)$. On the other hand, if $E_X(u) = E_X(v)$, then let $x \in E_Y(u) \setminus E_X(u)$. Then x must contain some $\#$. Let $w \sqsubseteq x$ be the maximal subpattern so that $\# \notin w \sqcup u$ and $\mathrm{dom}(w \sqcup u)$ is an interval. Then $w \in E_{\mathcal{L}(Y)}(u)$ (indeed, we assumed that $\mathcal{F}$ was such that locally admissible patterns on intervals are globally admissible), and as $E_X(u) = E_X(v)$ we also have $w \in E_{\mathcal{L}(Y)}(v)$, and so $x \in E_Y(v)$ as w is bounded by $\#$'s in x (or is infinite).
  - If u contain a $\#$, then let $0 \leq i < n$ be the leftmost position such that $u_i = \#$, and let $0 \leq j < n$ be the rightmost such position. It is clear that $E_Y(u)$ depends only on $P_Y(u|_{[\![0,i-1]\!]})$ and $F_Y(u|_{[\![j+1,n-1]\!]})$. Moreover, for the same reasons as above, this depends only on $P_X(u|_{[\![0,i-1]\!]})$ and $F_X(u|_{[\![j+1,n-1]\!]})$.

  Summing the two cases separately, we get

$$E_Y(n) \leq E_X(n) + \sum_{0 \leq i \leq j < n} P_X(i) F_X(n-1-j) \leq E_X(n) + \sum_{i+j<n} P_X(i) F_X(j)$$

■

Now, we also have for all n that $P_X(n) \leq E_X(n)$ and $F_X(n) \leq E_X(n)$. Moreover, by definition of $h_E(X)$, we have $E_X(n) = 2^{h_E(X)n + o(n)}$. Denoting $\alpha = h_E(X)$, we have:

$$\begin{aligned} E_Y(n) &\leq E_X(n) + \sum_{i+j<n} P_X(i) F_X(j) \\ &\leq E_X(n) + \sum_{i+j<n} E_X(i) E_X(j) \\ &\leq 2^{\alpha n + o(n)} + \sum_{i+j<n} 2^{\alpha i + o(i)} 2^{\alpha j + o(j)} \\ &\leq \mathrm{poly}(n) 2^{\alpha n + o(n)} \end{aligned}$$

Taking the log and dividing by n, and with n going to $+\infty$ we finally get $h_E(X) \leq h_E(Y) \leq h_E(X)$. □

2.4.3 One-dimensional effective subshifts

We will now turn our attention to the general case of effective subshifts. By Proposition 2.36, we cannot realize more than then Π_3 real numbers. We will show that in fact, effective-$\mathbb{Z}$ subshifts are already enough to realize this class.

Theorem 2.43

The extender entropies of effective $\mathbb{Z}$-subshifts are exactly the non-negative Π_3 real numbers.

The proof relies on an explicit construction of a subshift Z_α satisfying $h_E(Z_\alpha) = \alpha$ for any positive real $\alpha \in \Pi_3$. The high-level strategy is rather simple: the main idea is to first try to enforce $h(Z_\alpha) \approx 2^{\alpha n}$, and then ensure that for $u \neq v \in \mathcal{L}(Z_\alpha)$ we have $E_{Z_\alpha}(u) \neq E_{Z_\alpha}(v)$. However, even the first requirement is already impossible to satisfy, as if it were true, we would have $h(Z_\alpha) = \alpha$ an arbitrary Π_3 real while entropies of effective subshifts are always Π_1 reals. Nevertheless, this is still the strategy we try to follow. The rest of the Section 2.4.3 is the proof of this theorem. Note that thanks to Theorem 2.27, it is enough to prove it for any $\alpha \in [0,1] \cap \Pi_3$.

Encoding integers into configurations

Part of the construction will use ideas from the proof of Proposition 2.33, where configurations are periodic, with the allowed periods depending on the behaviour of some machine. We formally define this construction and define a few helpful notations.

Let $\mathcal{A}_* = \{*, \square\}$.

Notation. For $i > 0, 0 \leq k_1 < i$, note $\langle i \rangle_{k_1} = \sigma_{k_1}(\dots \square \cdot * \square^{i-1} * \square^{i-1} * \dots)$. More formally:

$$\langle i \rangle_{k_1} : \mathbb{Z} \to \mathcal{A}_*$$
$$p \mapsto \begin{cases} * & \text{if } p = k_1 \mod i \\ \square \text{ otherwise} & \end{cases}$$

Denote $\langle \infty \rangle = \{x \in \mathcal{A}_*^{\mathbb{Z}}, |x|_* \leq 1\}$.

Taking the orbit of all those configurations, we obtain the subshift

$$X_* = \overline{\bigcup \left\{ \langle i \rangle_{k_1}, i > 0, 0 \leq k_1 < i \right\}} = \bigcup \left\{ \langle i \rangle_{k_1}, i > 0, 0 \leq k_1 < i \right\} \sqcup \langle \infty \rangle$$

We call the configurations of $\langle \infty \rangle$ **degenerate** configurations.

Claim 2. $|\mathcal{L}_n(X_*)| = O(n^2)$.

Proof. By definition:

$$\mathcal{L}_n(X_*) = \{0^n\} \cup \bigcup_{i \leq n} \{0^i * 0^{n-i-1}\} \bigcup_{i \leq n} \bigcup_{0 \leq k_1 < i} 0^{k_1} * 0^{i-1} \dots * 0^{n-k_1 \mod i}$$

□

Configurations with controlled density

We now explain how to construct configurations with a controlled density, that is, configurations on $\{0,1\}$ where the number of 1 in large patterns converges to some value. More precisely, for some $\alpha \in [0,1]$, we want to construct a subshift $T_\alpha \subseteq \{0,1\}^{\mathbb{Z}}$ which verifies for any $x \in T_\alpha$:

$$\lim_{n \to +\infty} \frac{|x|_{[\![0,n-1]\!]}|_1}{n} = \alpha$$

There are several explicit construction of such subshifts T. The one we choose here is based on **Toeplitz** subshifts. A general introduction to many equivalent definitions and properties of Toeplitz subshifts can be found in [Sel20], but we only restate what is necessary for our proofs, without giving general results about Toeplitz subshifts.

Define the **2-adic valuation** map $v_2 \colon \mathbb{N} \to \mathbb{N}$ by $v_2(n) = \max_m 2^m | n$, and consider the sequence $T = 12131214\cdots = (v_2(2n))_{n>0}$ (also known as the ruler function, see OEIS A001511). Using T, we can associate to any sequence $(u_n)_{n>0}$ another sequence Tu, called its **Toeplitzification**, defined by $(Tu)_n = u_{T_n}$ for all $n > 0$.

Let $\beta \in [0,1]$ be a real number, and let $(\beta_n)_{n>0}$ be its binary expansion, so that $\beta = \sum_{n=1}^{+\infty} \beta_n 2^{-n}$. Define $T\beta$ the sequence defined for all n by $(T\beta)_n = \beta_{T_n} = \beta_1\beta_2\beta_1\beta_3\beta_1\dots$.

Claim 3. For any $\beta \in [0,1]$, $n \geq 1$ and for $w \sqsubseteq T\beta$ any subsequence of length n, we have $|w|_1 = \beta n + O(1)$.

Proof. It is enough to prove the claim for $\operatorname{dom}(w) = [\![1, 2^k]\!]$, as $v_2(2^k + \ell) = v_2(\ell)$ for $0 < \ell < 2^k$. Then:

$$\begin{aligned}
|w|_1 &= |\{1 \leq i \leq 2^k \mid (T\beta)_i = 1\}| \\
&= |\{1 \leq i \leq 2^k \mid \beta_{T_i} = 1\}| \\
&= \sum_{j=1}^{k} |\{1 \leq i \leq 2^k \mid T_i = j\}| \mathbf{1}_{\beta_j = 1} \\
&= \sum_{j=1}^{k} \beta_j |\{1 \leq i \leq 2^k \mid T_i = j\}| \\
&= \sum_{j=0}^{k-1} \beta_j |\{1 \leq i \leq 2^k \mid v_2(i) = j\}| \\
&= \sum_{j=0}^{k-1} \beta_j |\{p \in \mathbb{N} \mid 1 \leq 2^j(2p+1) \leq 2^k\}| \\
&= \sum_{j=0}^{k-1} \beta_j |\{0 \leq p < 2^{k-j-1}\}| \\
&= \sum_{j=0}^{k-1} \beta_j 2^{k-j-1} \\
&= 2^k \sum_{j=1}^{k} \beta_j 2^{-j} \\
&= n(\beta - \sum_{j=k+1}^{\infty} (\beta_j 2^{-j})) \\
&= n(\beta + O(\frac{1}{n}))
\end{aligned}$$

□

We can then define a subshift associated to a real number β, by considering the orbit of the $T\beta$. However, our goal is not to control a *density* of 1 in configurations, but to control the number of extender sets. We still use the previous idea of constructing periodic configurations to separate extender sets.

Notation. Let $\alpha \in [0,1]$, $i > 0$, and $0 \leq k_1 < i$. Define

$$\begin{aligned}
T(\beta, i)_{k_1} \colon \mathbb{Z} &\to \{0,1\} \\
p &\mapsto (T\beta)_{p+k_1 \mod i}
\end{aligned}$$

In other words, $T(\beta, i)_0$ is the sequence obtained by repeating periodically the word $((T\beta)_1 \dots (T\beta)_i)$, and $T(\beta, i)_{k_1} = \sigma_{k_1}(T(\beta, i)_0)$.

Notation. We denote by $\beta(x)$ the inverse operation: more precisely, for $x = (T(\beta, i)_{k_1})$, we define

$$\beta(x) = \sum_{k=1}^{\log i} \beta_k 2^{-k} = \sum_{k=0}^{\log i - 1} x_{k_1+2^k} 2^{-k-1}$$

We now define, for every $\alpha \in [0,1]$ and $i > 0$, the $\mathbb{Z}$-subshift $T_{\alpha,i}$ on alphabet $\{0,1\}$ as

$$T_{\leq \alpha, i} = \left\{ T(\beta, i)_{k_1} \in \{0,1\}^{\mathbb{Z}}, \beta \leq \alpha, 0 \leq k_1 < i \right\}$$

Claim 4. $T_{\leq \alpha, i}$ is a subshift.

Claim 5. For $\alpha \in [0,1] \cap \Pi_1$ and any $i > 0$, $T_{\leq \alpha, i}$ is an SFT. Moreover, there exists an enumerating Turing Machine M such that $T_{\leq \alpha, i} = M(\alpha, i)$ eventually; in other words, one can enumerate a finite family $\mathcal{F}$ such that $T_{\leq \alpha, i} = \mathcal{X}_{\mathcal{F}}$ uniformly in α.

Proof. By definition of the class Π_1, the set $\{r \in \mathbb{Q}, r > \alpha\}$ is recursively enumerable. One can then forbid the non-i-periodic configurations, or those which do not correspond to some prefix of the ruler sequence. Moreover, we forbid the finitely many patterns corresponding to Toeplitzification of rationals $r = \sum_{k=1}^{\log i} r_k 2^{-k}$ with $r > \alpha$. More precisely, we forbid all the patterns of length $2i$ such that for all $k_1 < i$, $w_{k_1} w_{k_1+1} \dots w_{k_1+i-1}$ is the prefix of $T(r, i)_0$ for some $r > \alpha$. □

An auxiliary subshift

Let $\alpha \in [0,1] \cap \Pi_3$ be a positive real number. By definition, there exists $(\alpha_{i,j})_{i,j \in \mathbb{N}}$ such that $\alpha = \inf_i \sup_j \alpha_{i,j}$. We can make the following assumptions:

- For all $i \in \mathbb{N}$, we have that $(\alpha_{i,j})_{j \in \mathbb{N}}$ is a non-decreasing sequence. Denote $\alpha_i = \sup_j \alpha_{i,j}$.
- The sequence $(\alpha_i)_{i \in \mathbb{N}}$ is non-increasing.

We do not define yet the subshift Z_α satisfying $h_E(Z_\alpha) = \alpha$. For the moment, we construct an intermediary subshift, from which we will easily obtain Z_α. Define therefore W_α the following $\mathbb{Z}$-subshift on various layers:

1. **First layer L_1:** we set the first layer to be $L_1 = X_*$ as defined in Section 2.4.3. Informally, in a configuration x, this layer will be some $\langle i \rangle$ and will act as a witness for some i such that x "approximates" $\alpha_i = \sup_j \alpha_{i,j} \in \Sigma_2$.

2. **Second layer L_2:** we also set the second layer to be $L_2 = X_*$. Informally, in a configuration x with first layer $\langle i \rangle$, this layer will be some $\langle j \rangle$ and will act as a witness for some j such that x "approximates" $\alpha_{i,j} = \Pi_1$.

3. **Density layer L_d:** the density layer is $L_d = \mathcal{A}_d^{\mathbb{Z}} = \{0,1\}^{\mathbb{Z}}$. We impose an additional condition: in a configuration x having its first two layers non-degenerate (in the sense of Section 2.4.3), respectively $\langle i \rangle$ and $\langle j \rangle$, the density layer will be a configuration of $T(\alpha_{i,j}, i)$. As $\alpha_{i,j} \in \Pi_1$, we have by Claim 5 that $T(\alpha_{i,j}, i)$ is an SFT.

Finally, define W_α as:

$$W_\alpha = \left\{(z^{(1)}, z^{(2)}, z^{(d)}) \in L_1 \times L_2 \times L_d \;\middle|\; z^{(2)} \in \langle\infty\rangle \right\}$$
$$\cup \bigcup_{i>0} \bigcup_{j\geq i} \Big\{ \left(\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, T(\beta, i)_{k_1}\right) \in L_1 \times L_2 \times L_d \;\Big|$$
$$i \leq j, 0 \leq k_1 < i, 0 \leq k_2 < j, 0 \leq \beta \leq \alpha_{i,j} \Big\}$$

We show in Figure 2.4 an example of a configuration of W_α. In accordance with Section 2.4.3, we say that a configuration $z \in W_\alpha$ is **degenerate** if its second layer is degenerate in X_*, and **proper** otherwise. We extend this distinction to patterns of $\mathcal{L}(W_\alpha)$: a pattern w is degenerate if it only appears in degenerate configurations, and proper if there exists a proper configuration z such that $w \sqsubseteq z$.

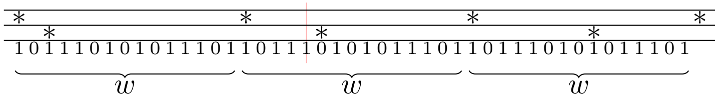

Figure 2.4: A proper pattern: L_d contains a Toeplitz encoding of $\overline{.1010}^2 = \frac{5}{8}$. $z = (\langle 15\rangle_{11}, \langle 18\rangle_1, T(\frac{5}{8}, 15)_{10})$. The vertical red line indicates the origin.

We now prove that W_α is effective:

Claim 6. If $\alpha \in [0,1] \cap \Pi_3$, then W_α is an effective $\mathbb{Z}$-subshift.

Proof. The first two layers are clearly effective, as they are simply periodic configurations. One can enforce that in a configuration $z = (\langle i\rangle, \langle j\rangle, z^{(d)})$, we have $i \leq j$ by forbidding all the patterns $(*0^{n-1}, P, z^{(d)}) \in \mathcal{L}_n(L_1 \times L_2 \times L_d)$ where P contains at least 2 $*$. The only remaining condition is then to ensure that in a proper configuration $z = (\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, z^{(d)})$, we have $z^d \in T_{\leq \alpha_{i,j}, i}$. But this is also straightforward, as we can simply enumerate the patterns of $T_{\leq\alpha_{i,j},i}$ only when the first two layers each contain two $*$, so that we know the respective i, j encoded in those first and second layers, and therefore the Π_1 real $\alpha_{i,j}$. By Claim 5, this is an effective procedure, and so W_α is effective. □

Let $\mathfrak{D}_E(n) = |\{E_{W_\alpha}(w), w \in \mathcal{L}_n(W_\alpha) \text{ and } w \text{ is degenerate}\}|$, and $\mathfrak{P}_E(n) = |\{E_{W_\alpha}(w), w \in \mathcal{L}_n(W_\alpha) \text{ and } w \text{ is proper}\}|$. The idea is that $\mathfrak{D}_E(n)$ should be negligible, while $\mathfrak{P}_E(n)$ relates to α in a precise way.

Claim 7. If u is a degenerate pattern and v is a proper pattern, then $E_{W_\alpha}(u) \neq E_{W_\alpha}(v)$.

Proof. By definition of proper patterns, there exists a proper configuration extending v, which therefore cannot extend u. □

Claim 8. $\mathfrak{D}_E(n) = O(n^3)$.

Proof. Let $n \geq 1$, $(u = (u^{(1)}, u^{(2)}, u^{(d)}), v = (v^{(1)}, v^{(2)}, v^{(d)})) \in \mathcal{L}_n(W_\alpha)$ be two degenerate patterns. By definition of W_α, for any $z^{(d)} \in \mathcal{L}_n(L_d) = \{0,1\}^n$, we have that $(u^{(1)}, u^{(2)}, z^{(d)}) \in \mathcal{L}_n(W_\alpha)$ is a degenerate pattern. In particular, if $u^{(1)} = v^{(1)}$ and $u^{(2)} = v^{(2)}$, then $E_{W_\alpha}(u) = E_{W_\alpha}(v)$. As u, v are degenerate, for any $z = (z^{(1)}, z^{(2)}, z^{(d)}) \in W_\alpha$ such that $u \sqsubseteq z$, we have $|z^{(2)}|_* \leq 1$, so there are $O(n)$ such patterns. Moreover, $\mathcal{L}_n(X_*) = O(n^2)$ by Claim 2, so in total, $\mathfrak{D}_E(n) = O(n^3)$. □

Claim 9. If $u \neq v$ are proper patterns, then $E_{W_\alpha}(u) \neq E_{W_\alpha}(v)$.

Proof. There exists $z = (\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, z^{(d)}) \in W_\alpha$ such that $z|_{[\![0,n-1]\!]} = u$. Then, z is $i \times j$ periodic as $\langle i\rangle_{k_1}, z^{(d)}$ are i-periodic and $\langle j\rangle_{k_2}$ is j-periodic, so $z|_{\mathbb{Z}\setminus[\![0,n-1]\!]}$ completely determines $z|_{[\![0,n-1]\!]}$ and in particular $z|_{\mathbb{Z}\setminus[\![0,n-1]\!]} \notin E_{W_\alpha}(v)$. □

Those claims imply that $E_{W_\alpha}(n) = O(n^3) + \mathfrak{P}_E(n)$. However, $(\mathfrak{P}_E(n))_{n\in\mathbb{N}}$ grows polynomially, and not exponentially. It will nonetheless be possible to slightly alter the construction to obtain a precisely controlled growth rate of $(\mathfrak{P}_E(n))$ thanks to the next remark:

Lemma 2.44

Let $z = (\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, z^{(d)}) \in W_\alpha$ be a proper configuration. Then $|z^{(d)}|_{[\![0,i-1]\!]}|_1 \leq \alpha_i i + O_i(1)$, and for $i \geq n$, $|z^{(d)}|_{[\![0,n-1]\!]}|_1 \leq \alpha_n n + O_n(1)$.

Proof. The first inequality holds because $z^{(d)}$ is $T(\beta,i)_{k_1}$ for some $\beta \leq \alpha_{i,j}$, which is i-periodic, and so $z^{(d)}|_{[\![0,i-1]\!]}$ is (up to a cyclic permutation) $T(\beta,i)|_{[\![0,i-1]\!]}$. By Claim 3 it contains $\beta i + O(i)$ symbols 1, and $\beta \leq \alpha_{i,j} \leq \alpha_i$.

To get the second claim, we assumed that $(\alpha_k)_{k\in\mathbb{N}}$ was non-increasing so $\alpha_n \leq \alpha_i$. □

Multiplying the number of patterns

A common technique from the literature to transform a density-like property of a subshift into an entropy-like property is to add another layer acting almost as a full-shift (see for example [HM10, Section 8] or [CV21, Section 4.6]). More precisely, consider the following construction. Let $\mathcal{A}_f = \{\square, 0, 1\}$ and $L_f = \mathcal{A}_f^{\mathbb{Z}}$, and define $\pi_{\text{sync}} \colon \mathcal{A}_f \to \{0,1\}$ be the map defined by $\pi_{\text{sync}}(\square) = 0$ and $\pi_{\text{sync}}(0) = \pi_{\text{sync}}(1) = 1$. We naturally extend π_{sync} to a map $L_f \to \{0,1\}^{\mathbb{Z}}$. We can then add a fourth layer to W_α:

4. **Free layer L_f:** We add a layer L_f which is synchronized with the density layer in the following sense: if $z = (z^{(1)}, z^{(2)}, z^{(d)}, z^{(f)})$ is such that $(z^{(1)}, z^{(2)}, z^{(d)})$ is a proper W_α configuration then we require $\pi_{\text{sync}}(z^{(f)}) = z^{(d)}$, and that $z^{(f)}$ is periodic (with the same period than $z^{(1)}$ and $z^{(d)}$), see Figure 2.5.

More precisely, we define the subshift Z_α from W_α as follows:

$$
\begin{aligned}
Z_\alpha = &\left\{ (z^{(1)}, z^{(2)}, z^{(d)}, z^{(f)}) \in L_1 \times L_2 \times L_d \times L_f \;\middle|\; z^{(2)} \in \langle\infty\rangle \right\} \\
\cup \bigcup_{i>0}\bigcup_{j\geq i} \Big\{ &\left(\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, T(\beta,i)_{k_1}, z^{(f)} \right) \in L_1 \times L_2 \times L_d \times L_f \;\Big| \\
&0 \leq k_1 < i, 0 \leq k_2 < j, 0 \leq \beta \leq \alpha_{i,j}, \\
&\pi_{\text{sync}}(z^{(f)}) = T(\beta,i)_{k_1}, z^{(f)} \text{ is } i\text{-periodic} \Big\}
\end{aligned}
$$

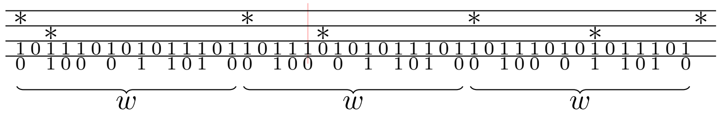

Figure 2.5: A proper pattern: L_d contains a Toeplitz encoding of $\overline{.1010}^2 = \frac{5}{8}$. $z = (\langle 15\rangle_{11}, \langle 18\rangle_1, T(\frac{5}{8}, 15)_{10})$. The vertical red line indicates the origin.

The Figure 2.5 shows a proper pattern of Z_α, corresponding to the addition of the free layer to the pattern of W_α already depicted in Figure 2.4.

Claim 10. Z_α is an effective $\mathbb{Z}$-subshift.

Proof. This is a consequence of W_α being effective by Claim 6. Indeed, it suffices to forbid on top of the patterns defining W_α all the patterns $(u^{(1)}, u^{(2)}, u^{(d)}, u^{(f)})$ where $u^{(2)}$ contains two $*$, in which $(u^{(d)}, u^{(f)})$ contains either of $(0,0), (0,1), (1,\square)$, and to enforce the periodicity of $u^{(f)}$ according to the one of $u^{(1)}$. $\square$

In Z_α, we say that a pattern u (resp. a configuration z) is proper (resp. degenerate) if $\pi_{L_1\times L_2\times L_d}(u) \in \mathcal{L}(W_\alpha)$ (resp. $\pi_{L_1\times L_2\times L_d}(z) \in W_\alpha$) is proper (resp. degenerate).

Counting patterns and extender sets

Note that the proofs of Claim 8 and Claim 9 also give the equivalent claims for Z_α:

Claim 11. $|E_X(u), u \in \mathcal{L}_n(Z_\alpha)$ is degenerate$| = O(n^3)$ and if $u \neq v \in \mathcal{L}_n(Z_\alpha)$ are proper then $E_{Z_\alpha}(u) \neq E_{Z_\alpha}(v)$.

In fact, we have the following quantitative claim, relating the patterns of W_α to those of Z_α:

Claim 12. Let $(\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, z^{(d)})$ be proper in W_α, and let d be the number of 1 in any i-period of $z^{(d)}$. Then:

$$|\{(\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, z^{(d)}, z^{(f)}) \in Z_\alpha\}| = 2^d$$

Proof. $(\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, z^{(d)}, z^{(f)}) \in \mathcal{L}(Z_\alpha)$ if and only if $\pi_{\text{sync}}(z^{(f)}) = z^{(d)}$, if and only for all k such that $z^{(d)}_k = 1$, we have $z^{(d)} \in \{0,1\}$. As $z^{(f)}$ is also i-periodic, and as the symbols of $z^{(d)}$ within a period can be chosen independently, we have the required equality. $\square$

The key lemma is then the following, mirroring Lemma 2.44.

Lemma 2.45

Let $P(n)$ be the number of proper patterns in $\mathcal{L}_n(Z_\alpha)$. Then

$$2^{\alpha_n n + O_n(1)} \leq P(n) \leq \text{poly}(n) \sum_{i=1}^{n} 2^{\alpha_i i + O_i(1)}$$

Proof. We first prove the leftmost inequality. For any $n \geq 0$, it suffices to exhibit a family of patterns of the required cardinality, all having a different extender set. Consider for $j \geq n$ the pattern $w_j \in \mathcal{L}_n(W_\alpha)$, defined by $w_j = (\langle n\rangle_0, \langle j\rangle_0, T(\alpha_{n,j}, 0))|_{[\![0,n-1]\!]}$. By Claim 3, w_j contains $d_j = \alpha_{n,j} n + O_n(1)$ symbols 1 in its density layer. Taking $j \to +\infty$, we get $\alpha_{n,j} \to \alpha_n$ and therefore at the limit, we obtain a pattern $w \in \mathcal{L}_n(W_\alpha)$ which is proper and with $d = \alpha_n n$ symbols 1 in its density layer. By Claim 12, we obtain the lower bound.

For the rightmost inequality, we will do a rough overestimation of the number of proper patterns in Z_α. To do this, we first count the number of proper patterns in W_α. For a proper pattern $u = (\langle i\rangle_{k_1}, \langle i\rangle_{k_2}, T(\beta, i)_{k_1}) \in \mathcal{L}_n(W_\alpha)$:

- If $i < n$, then the pattern u is i-periodic and so it suffices to bound the number of i-periods. In particular, by Lemma 2.44, there are at most $\alpha_i i + O(1)$ symbols 1 in $T(\beta, i)_{k_1}$.
- Otherwise, Lemma 2.44 also gives that there are at most $\alpha_n n + O(1)$ symbols 1 in $T(\beta, i)_{k_1}$.

Summing everything and given Claim 12, we get that

$$\begin{aligned}P(n) &\leq \sum_{i=1}^{n}\sum_{k_1=0}^{i-1}\sum_{j=i}^{n}\sum_{k_2=0}^{j-1} 2^{\alpha_i i+O_i(1)} + \sum_{k_1=0}^{n}\sum_{k_2=0}^{n} 2^{\alpha_n n+O(1)}\\ &\leq \text{poly}(n)\sum_{i=1}^{n} 2^{\alpha_i i+O_i(1)}\end{aligned}$$

□

We can now prove the claimed theorem, using a small computational lemma:

Lemma 2.46

For any non-zero polynomial P with $P(n) > 0$ for $n > 0$, $d > 0$, and any converging positive sequence $\alpha_n \to \alpha$, we have

$$\frac{\log\left(\sum_{i=0}^{n} P(i)2^{\alpha_i i^d}\right)}{n^d} \to \alpha$$

Proof. Fix $\varepsilon > 0$, and let $I \in \mathbb{N}$ be such that for all $i \geq I$, $|\alpha_i - \alpha| \leq \varepsilon$ and moreover $P(i+1) \geq P(i)$. Let us note $S_n = \sum_{i=0}^{n} P(i)2^{\alpha_i i^d}$, and fix some $i \geq I$. Now, $S_n \geq P(n)2^{\alpha_n n^d} \geq P(n)2^{(\alpha-\epsilon)n^d}$ and so $\frac{\log S_n}{n^d} \geq \alpha - \epsilon$.

On the other hand,

$$\begin{aligned}S_n &= \underbrace{\sum_{i=0}^{I} P(i)2^{\alpha_i i^d}}_{K} + \sum_{k=I+1}^{n} P(i)2^{\alpha_i i^d}\\ &\leq K + \sum_{i=I+1}^{n} P(i)2^{(\alpha+\varepsilon)i^d}\\ &\leq K + nP(n)2^{(\alpha+\varepsilon)n^d}\\ &\leq (n+1)P(n)2^{(\alpha+\varepsilon)n^d} \text{ for large enough } n\end{aligned}$$

and so $\frac{\log S_n}{n^d} \leq \alpha + \varepsilon + \frac{\log(n+1)P(n)}{n^d}$. By taking n large enough so that $\frac{\log(n+1)P(n)}{n^d} \leq \epsilon$, we finally have

$$\alpha - \varepsilon \leq \frac{\log S_n}{n^d} \leq \alpha + 2\varepsilon$$

for all sufficiently large n. As this holds for any $\varepsilon > 0$, we get $\frac{\log S_n}{n^d} \to \alpha$. □

Proof of Theorem 2.43. Let $\alpha \in \Pi_3$ and consider the subshift Z_α defined in Section 2.4.3. It is effective by Claim 10. Denoting $P(n)$ the number of proper patterns of size n of Z_α, we have by Claim 11 that $E_{Z_\alpha}(n) = O(n^3) + P(n)$, and so by Lemma 2.45 we have

$$2^{\alpha_n n+O_n(1)} \leq E_{Z_\alpha}(n) \leq \text{poly}(n)\sum_{i=1}^{n} 2^{\alpha_i i+O_i(1)}$$

Taking $n \to +\infty$ we get $\alpha_n \to \alpha$ and therefore $h_E(Z_\alpha) = \alpha$ by Lemma 2.46. □

2.4.4 Computable subshifts

In the case of computable subshift, the construction used to prove Theorem 2.43 also proves the next theorem:

Theorem 2.47

The extender entropies of computable $\mathbb{Z}^d$ subshifts are exactly the non-negative Π_2 real numbers.

Proof. The upper bound is given by Proposition 2.37. It suffices to prove the theorem for $\mathbb{Z}$ subshifts, thanks to Proposition 2.32. Now, notice that if α a Π_2 real number, the subshift Z_α constructed in the proof of Theorem 2.43 is in fact computable. To see this, note that we can now write $\alpha = \inf_i \sup_j \alpha_{i,j}$ where the $\alpha_{i,j}$ are computable (rather than Π^1). This means that given a pattern w of $L_1 \times L_2 \times L_d \times L_f$, assuming that the bits of L_f are consistent with the density bits of L_d:

- If $w^{(2)}$ contains at most one $*$ then it is allowed (this is already the case in the proof of Theorem 2.43).

- Otherwise, $w \sqsubseteq (\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, T(\beta, i)_{k_1}, w^{(f)})$. As $j \geq i$ and $w^{(2)}$ contains two $*$, so does $w^{(1)}$, and so we know the period of w. It then suffices to check if $\beta \leq \alpha_{i,j}$, which is decidable as $\alpha_{i,j}$ is computable.

□

2.4.5 Multi-dimensional sofic subshifts

We would now like to show that this characterization of extender entropies also holds for sofic multi-dimensional subshifts. In order to transfer the result of Theorem 2.43 to sofic $\mathbb{Z}^d$-subshifts, one could try to follow one of those ideas:

- Use Proposition 2.32 and free lifts. This does not work, as the free lift of an effective subshift is effective but not sofic in general.

- Adapt the ideas of Section 2.4.3 and Section 2.4.3 to an higher-dimensional construction, by using periodic square blocks. However, arguments similar to those used in Proposition 2.15 show that we have no hope of such a subshift, containing configurations of arbitrarily large periods, being sofic.

- Use ideas from Proposition 2.16 to avoid the problem of having arbitrarily large periods. Now, the problem is the number of extender sets will be very hard to bound: indeed, a key element in the proof of Theorem 2.43, which is explicit in Lemma 2.46, is the fact that configurations "implementing" approximations of $\alpha_{i,j}$ have a small period, and therefore, become negligible compared to the number of patterns of larger periods.

Combining those ideas, we will nevertheless manage to show that the following result holds:

Theorem 2.48

For any $d \geq 2$, the extender entropies of sofic $\mathbb{Z}^d$-subshifts are exactly the non-negative Π_3 real numbers.

The proof relies on an observation already present in Proposition 2.16: it is not necessary to ensure that configuration are *periodic* to distinguish the extender sets of any two "good" patterns, it is enough that there exists a configuration which witnesses any difference between those patterns. One can understand this difference by considering the following logical formulas. The first one corresponds to the case of periodic configurations:

$$\forall u, \exists z \in E_X(u), \forall v, z \in E_X(v) \implies u = v.$$

This is of course under-specified, notably, u and v must be chosen in some specific subset of $\mathcal{L}(X)$ for the formula to hold. However, such a formula would indeed imply that all such patterns u, v have different extender sets. The semi-mirror construction instead makes another kind of property hold:

$$\forall u, \forall v, \exists z \in E_X(u), u \neq v \implies z \notin E_X(v).$$

Using this observation, we will prove Theorem 2.48 with the following strategy:

- For each $\alpha \in \Pi_3 \cap [0, 1]$, consider the subshift W_α constructed in Section 2.4.3.
- Using Theorem 1.87, we obtain a sofic $\mathbb{Z}^2$ subshift ${W_\alpha}^\uparrow$, where $i \times i$ squares have a controlled density of symbols 1 in their density layer.
- In each $i \times i$ square, we identify a single bit. This bit is made $i \times i$ periodic using some additional background layers.
- Arguments similar to those used in Proposition 2.16 then ensure that all the "proper" patterns have different extender sets, and the computations of Section 2.4.3 can be adapted to show that the number of such patterns is what we need to obtain an extender entropy α.

Note that:

- We can restrict ourselves to the case $\alpha \in [0, 1]$ thanks to Theorem 2.27.
- We can restrict ourselves to the case of $\mathbb{Z}^2$-subshifts thanks to Proposition 2.32, as the free lift of a *sofic* subshift is sofic.

Marked offsets instead of periods

As in Section 2.4.3, we will introduce some specific intermediate subshifts to build the actual subshift Y_α with $h_E(Y_\alpha) = \alpha$. In order to implement this "marker" layer, which identifies a single bit per square as explained above, let $\mathcal{A}_m = \{\square, \blacksquare\}$, and define for any $i > 0, m_1, m_2 \in [\![0, i-1]\!]^2$ the configuration:

$$[2i]_{m_1,m_2} : \mathbb{Z}^2 \to \mathcal{A}_m$$
$$p \mapsto \begin{cases} \blacksquare & \text{if } p = m_1 \mod 2i \wedge p = m_2 \mod 2i \\ \square & \text{otherwise} \end{cases}$$

As a shortcut, we will use $p = (m_1, m_2) \mod (2i, 2i)$ for $p = m_1 \mod 2i \wedge p = m_2 \mod 2i$. Note that $[2i]_{m_1,m_2}$ is $(2i, 2i)$-periodic. For $x = [2i]_{m_1,m_2}$, say that a position

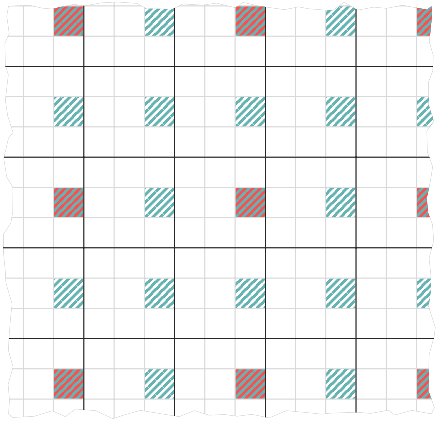

Figure 2.6: The configuration $[2 \cdot 3]_{2,1} \in X_m$, where the dashed cells represent the marked positions.

$p \in \mathbb{Z}^2$ is **marked** if $p \in (m_1 + i\mathbb{Z}, m_2 + i\mathbb{Z})$. Note that some marked positions p satisfy $x_p = \square$. Note Marked(x) this set of position.

We now define the subshift X_m of those configurations: let $[\infty] = \{x \in \mathcal{A}_m, |x|_{\blacksquare} \leq 1\}$. Then, define

$$X_m = \{[2i]_{m_1,m_2}, i > 0, 0 \leq m_1, m_2 < i\} \cup [\infty]$$

Claim 13. X_m is a sofic subshift.

A sofic marking subshift

Keeping the notations used in Section 2.4.3, we will define for $\alpha \in [0,1] \cap \Pi_3$ a $\mathbb{Z}^2$-subshift Y_α using several layers:

1. **Lifted layers:** The first three of Y_α are $L_1{}^{\uparrow}, L_2{}^{\uparrow}$ and $L_d{}^{\uparrow}$ where L_1, L_2, L_d are the layers of W_α.
2. **Marker layer L_m:** We define the marker layer $L_m = X_m$.
3. **Free layer L_f:** As in Section 2.4.3, we define $L_f = \{\square, 0, 1\}$.

and we can now define the subshift Y_α formally as follows:

$$\begin{aligned} Y_\alpha = & \left\{ (y^{(1)\uparrow}, y^{(2)\uparrow}, y^{(d)\uparrow}, y^{(m)}, y^{(f)}) \in \langle\infty\rangle^{\uparrow} \times \langle\infty\rangle^{\uparrow} \times L_d{}^{\uparrow} \times [\infty] \times L_f \right\} \\ & \cup \bigcup_{i>0} \left\{ (\langle i \rangle_{k_1}{}^{\uparrow}, y^{(2)\uparrow}, y^{(d)\uparrow}, [2i]_{m_1,m_2}, y^{(f)}) \in L_1{}^{\uparrow} \times \langle\infty\rangle^{\uparrow} \times L_d{}^{\uparrow} \times L_m \times L_f \mid 0 \leq m_1, m_2 < i \right\} \\ & \cup \bigcup_{i>0} \bigcup_{j \geq i} \left\{ (\langle i \rangle_{k_1}{}^{\uparrow}, \langle j \rangle_{k_2}{}^{\uparrow}, T(\beta, i)_{k_1}{}^{\uparrow}, [2i]_{m_1,m_2}, y^{(f)}) \in L_1{}^{\uparrow} \times L_2{}^{\uparrow} \times L_d{}^{\uparrow} \times L_m \times L_f \mid \right. \\ & \qquad 0 \leq k_1 < i, 0 \leq k_2 < j, 0 \leq m_1, m_2 < i, 0 \leq \beta \leq \alpha_{i,j}, \\ & \qquad \left. \pi_{\text{sync}}(y^{(f)}) = T(\beta, i)_{k_1}{}^{\uparrow}, y^{(f)}\Big|_{\text{Marked}([2i]_{m_1,m_2})} \text{ is constant} \right\} \end{aligned}$$

Informally, the configurations of Y_α are obtained as follows:

- They contain a lift $z^{\uparrow}$ of a configuration $z \in W_\alpha$.
- If z's first layer is non-degenerate equal to some $\langle i \rangle$, then the marker layer $y^{(m)}$ "marks" one position per (i,i)-square, and is $(2i, 2i)$-periodic with a single $\blacksquare$-cell per period.

- If furthermore z's second layer is non-degenerate, so that z is in fact a proper configuration in W_α, then the value of the free bits in position $\text{Marked}(y^{(m)})$ are all the same, *i.e.* one bit per (i,i)-square is periodic. In that case, the other free bits can only be superimposed onto non-zero bits of the underlying Toeplitz configuration of $z^{\uparrow}$, just as in Section 2.4.3.

We naturally extend the terminology of Section 2.4.3 to Y_α, and say that $y \in Y_\alpha$ is proper (resp. degenerate) if $\pi_{L_1\times L_2\times L_d^{\uparrow}}$ is proper (resp. degenerate) in W_α, and the same for patterns rather than for entire configurations.

We claim that Y_α is sofic and satisfies $h_E(Y_\alpha) = \alpha$. The fact that is has the required extender entropy is proven in a similar fashion than in Section 2.4.3, so we first prove the main claim that it is sofic, showing that those few modifications to Z_α are sufficient.

Lemma 2.49

Y_α is sofic.

There are three main steps in this proof. We will build an SFT that factors down to Y_α, using several layers:

- The first layers will be used to implement the marked positions, and in particular to enforce that if the first layer of Y_α is some $\langle i\rangle_{-}^{\uparrow}$ then the marker layer is some $[2i]_{-}$.

- The next layer is used to synchronize the periodic bit in all the marked position of Y_α's Free Layer.

- Finally, we need to slightly alter the construction so that there is a forced periodic free bit only if the first layers of Y_α are a (lift of a) proper configuration of W_α.

Using grids to mark positions We define a sofic subshift Y_{grid}, whose configurations are grids of a specific mesh. Let $\mathcal{A}_b = \{□, ◫\}$, $\mathcal{A}_r = \{□, ◫, ◫, ⊞\}$ and $\mathcal{A}_v = \{□, ◫, ◫, ⊞\}$. Then define:

- **The column layer:** The column layer $L_c \subset \mathcal{A}_b^{\mathbb{Z}^2}$ contains a $(i,1)$-periodic configuration for some $i > 0$, with one ◫ per period. Configurations of L_b are then blue columns, at distance i from one another.

- **The red and violet layers L_r, L_v:** They are respectively subshifts on $\mathcal{A}_r, \mathcal{A}_v$. Their configurations are grids (as illustrated in Figure 2.7) of some mesh $2i$.

Define $Y_{grid} \subset L_c \times L_r \times L_v$ as the subshift where the blue of columns of L_c are separated by any $i > 0$, the red and violet grids have mesh $2i$, and the red and violet grids are offset by (i,i), as illustrated in Figure 2.7.

Claim 14. L_b, L_r and L_v are sofic.

Proof. It is easy to construct L_r or L_v, that is, a subshift of grids, given L_b. Indeed, one can construct such a grid using a layer L_b, a layer which is L_b rotated so that the lines are horizontal, and using a third layer to synchronize the periods in the horizontal and vertical directions by using diagonal signals. □

Claim 15. Y_{grid} is sofic.

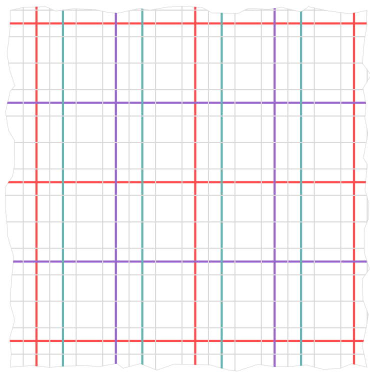

Figure 2.7: A configuration of Y_{grid}. The vertical blue lines are 3-periodic, and the two grids each have mesh 6 and are offset by $(3, 3)$

Proof. Each of L_b, L_v, L_r are sofic, and it remains to show how to synchronize the periods of the three layers and the offset of the two coloured grids. The offset can be solved as in the previous claim, by sending signals in diagonal which must only intersect grids on $\boxplus$ and $\boxplus$, and the periods can also be made the same by noticing that there must be exactly one $\boxbar$ between consecutive horizontal $\boxplus$ and $\boxplus$. Such a constraint can be checked by an SFT, so finally Y_{grid} is sofic. □

Making a free bit periodic Define the subshift $Y_{\text{sync}} \subseteq Y_{\text{grid}} \times L_f$ as

$$\begin{aligned} Y_{\text{sync}} = \{y^{(c)}, y^{(r)}, y^{(v)}, y^{(f)} \in L_c \times L_r \times L_v \times L_f \mid & (y^{(c)}, y^{(r)}, y^{(v)}) \in Y_{\text{grid}} \\ & \exists b \in \{0,1\}, \\ & \forall p \in \mathbb{Z}^2, (y_p^{(r)} = \boxplus \vee y_p^{(v)} = \boxplus) \implies y_p^{(f)} = b\} \end{aligned}$$

Claim 16. Y_{sync} is sofic

Proof. Define the **global subshift** $L_g = \{0^{\mathbb{Z}^2}, 1^{\mathbb{Z}^2}\}$, which is clearly an SFT. Then, $Y_{\text{grid}} \times L_f \times L_g$ is sofic by Claim 15, and therefore Y_{sync} is sofic as it is the image of $Y_{\text{grid}} \times L_f \times L_g$ by the following factor map whose block-map is $\pi_{\text{grid}\to\text{sync}}$ defined by:

$$\begin{aligned} \pi_{\text{grid}\to\text{sync}} \colon \mathcal{A}_c \times \mathcal{A}_r \times \mathcal{A}_v \times \mathcal{A}_f \times \{0,1\} &\to \mathcal{A}_c \times \mathcal{A}_r \times \mathcal{A}_v \times \mathcal{A}_f \\ (a_c, a_r, a_v, a_f, a_g) &\mapsto \left(a_c, a_r, a_v, \begin{cases} a_g & \text{if } (a_r = \boxplus \vee a_v = \boxplus) \\ a_f & \text{otherwise} \end{cases} \right) \end{aligned}$$

□

Synchronization only in proper configurations As explained above, in Y_α, we impose that there is a periodic free bit in each $i \times i$-square only when the underlying lifted W_α configuration is proper. The reason for that shall become clear when counting extender sets, and can be understood as the fact that we want degenerate configurations to amount to a very low number of extender sets, and we therefore do not want to have *any* configuration that could "differentiate" degenerate patterns using a marked periodic bit.

To do this, we need to "know" whether the configuration is degenerate. This can only be done in an effective subshift, which means that the lifted configurations should already "carry" this information and transmit it to Y_α.

Define $L_p = \{\mathbf{p}^{\mathbb{Z}}, \mathbf{d}^{\mathbb{Z}}\}$, the **proper layer**, which is clearly an SFT. Define then a subshift $W'_\alpha \subseteq W_\alpha \times L_p$ as

$$\begin{aligned} W'_\alpha = \{ & (z^{(1)}, z^{(2)}, z^{(d)}, z^{(p)}) \in L_1 \times L_2 \times L_d \times L_p \mid \\ & (z^{(1)}, z^{(2)}, z^{(d)}) \in W_\alpha, \\ & z^{(2)} \notin \langle \infty \rangle \implies z^{(p)} = \mathbf{p}^{\mathbb{Z}} \} \end{aligned}$$

In other words, the layer L_p can be any of $\mathbf{p}^{\mathbb{Z}}$ and $\mathbf{d}^{\mathbb{Z}}$ in degenerate configurations of W'_α, but must be $\mathbf{p}^{\mathbb{Z}}$ in proper configurations.

Claim 17. W'_α is effective.

Proof. If suffices to forbid for all $n > 0$ all the patterns $(w^{(1)}, w^{(2)}, w^{(d)}, \mathbf{d}^n) \in \mathcal{L}_n(W_\alpha) \times \{\mathbf{d}^n\}$ in which $w^{(1)}$ and $w^{(2)}$ each contain at least two $*$. This is clearly an effective procedure, and as W_α is effective by Claim 6, we have W'_α effective. $\square$

We are now ready to prove that Y_α itself is a sofic subshift.

Proof of Lemma 2.49. To show that Y_α is sofic, we will define yet another subshift Y'_α using Y_{sync} and W'_α. A simple 1-block map will then be enough to obtain Y_α as a factor of Y'_α, which will be enough to conclude. Hence, we define:

$$\begin{aligned} Y'_\alpha = \Big\{ & (y^{(1)\uparrow}, y^{(2)\uparrow}, y^{(d)\uparrow}, y^{(p)\uparrow}, y^{(c)}, y^{(r)}, y^{(v)}, y^{(f)} \in W'^{\uparrow}_\alpha \times Y_{\text{grid}} \times L_f) \mid \\ & \forall p \in \mathbb{Z}^2, y^{(1)}_p = * \iff y^c_p = ◫, \\ & \pi_{\text{sync}}(y^{(f)}) = y^{(d)\uparrow}, \\ & y^{(p)} = \mathbf{p}^{\mathbb{Z}} \implies (y^{(c)}, y^{(r)}, y^{(v)}, y^{(f)}) \in Y_{\text{sync}} \Big\} \end{aligned}$$

In other words, the configurations of Y'_α consist of a lifted configuration of $y_\alpha \in W_\alpha$, a configuration of $y_{\text{grid}} \in Y_{\text{grid}}$, a layer of free bits $y^{(f)}$ and a constant layer $y^{(p)}$ equal to $\mathbf{p}$ or $\mathbf{d}$. The blue columns of y_{grid} are aligned with the $*$ of the first layer of y_α. The free bits of $y^{(f)}$ can only appear on the "density bits" of $y^{(d)}_\alpha$, and one $i \times i$-periodic position has a synchronized bit *via* y_{grid} when the global constant layer is $\mathbf{p}^{\mathbb{Z}^2}$. By Claim 17, Theorem 1.87 and Claim 16, Y'_α is sofic.

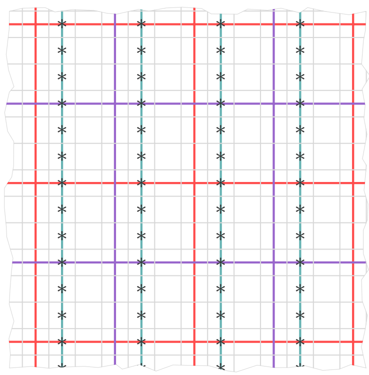

Figure 2.8: The layers L_1, L_c, L_r, L_v and parts of L_f of some configuration of Y'_α. Only the synchronized buit of L_f is shown, all the other ones are independent.

Define now the 1-block map ϕ:

$$\phi \colon \mathcal{A}_* \times \mathcal{A}_* \times \mathcal{A}_d \times \{\mathbf{p}, \mathbf{d}\} \times \mathcal{A}_c \times \mathcal{A}_r \times \mathcal{A}_v \times \mathcal{A}_f \to \mathcal{A}_* \times \mathcal{A}_* \times \mathcal{A}_d \times \{\square, \blacksquare\} \times \mathcal{A}_f\}$$

$$(a_1, a_2, a_d, a_g, a_c, a_r, a_v, a_f) \mapsto \left(a_1, a_2, a_d, \begin{cases} \blacksquare & \text{if } a_r = \boxplus \\ \square & \text{otherwise} \end{cases}, a_f \right)$$

$\square$

Counting extender sets

It remains to show that Y_α indeed has the claimed number of extender sets. We do not need to prove precise bounds, and can once again overestimate the number of extender sets.

Claim 18. $|\{E_{Y_\alpha}(w), w \in \mathcal{L}_n(Y_\alpha) \text{ degenerate}\}| = poly(n)$.

Proof. By Claim 11 it suffices to show that if u, v are degenerate and coincide on L_1, L_2, L_d, L_m then they have the same extender sets. But this is the case as we do not impose any kind of periodicity on the free bits in degenerate configurations. □

Say that two patterns $u, v \in \mathcal{L}_n(Y_\alpha)$ are **similar** if they coincide on their layers $L_1{}^\uparrow, L_2{}^\uparrow, L_d{}^\uparrow$ and L_m.

Claim 19. For u, v two similar patterns, $E_{Y_\alpha}(u) \neq E_{Y_\alpha}(v)$ if and only if there exists $y \in E_{Y_\alpha}(u)$ and $p \in \text{Marked}(\pi_{L_m}(y \sqcup u))$ such that $u_p^{(f)} \neq v_p^{(f)}$.

Proof. If there exists such a configuration y and marked position p, then clearly y does not extend v, as the marked positions need all have the same value on the free layer in proper configurations. For the other direction, notice that $\pi_{L_1 \times L_2 \times L_d}(E_{Y_\alpha}(u))$ depends only on $\pi_{L_1 \times L_2 \times L_d}(u)$ by definition of Y_α. Now, consider $y \in E_{Z_\alpha}(u)$ such that for any $p \in \text{dom}(u) \cap \text{Marked}(\pi_{L_m}(y \sqcup u))$, we have $u^{(f)}(p) = v_p^{(f)}$. As all the bits of L_f outside of $\text{Marked}(\pi_{L_m}(y \sqcup u))$ are independent, this means in particular that they do not depend on $\pi_{L_m}(u)$ outside of marked positions, and so this implies that y also extends v. Thus, if $E_{Z_\alpha}(u) \neq E_{Z_\alpha}(v)$, no such y, p exist. □

Figure 2.9 shows an example of similar patterns, a possible configuration extending one but not the other, and tries to picture why we need to consider similar patterns with an additional condition on the position of their differences in Claim 19.

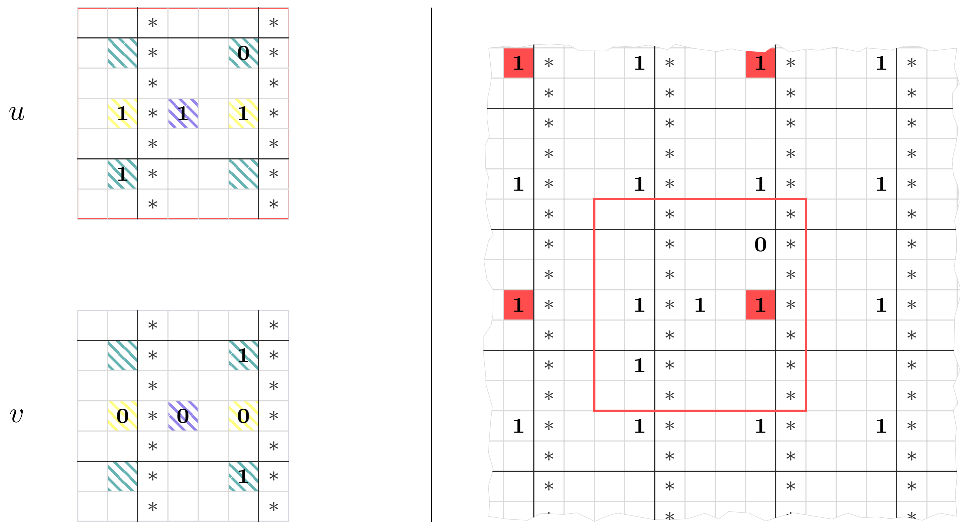

Figure 2.9: Example of similar patterns. Only some of the free bits are depicted. The dashed patterns are simply decorations: although u, v differ on some blue position, those positions do not satisfy the hypothesis of Claim 19 as they cannot be marked in neither u nor v. Yellow positions can be marked in both, and u, v differ on them, so they satisfy the hypothesis of Claim 19. The violet position obviously satisfies the hypothesis. On the right, an example of configuration extending u but not v.

This is enough to allow to prove an analogous of Lemma 2.45 without having the direct analogous of Claim 11:

Lemma 2.50

Let $P(n)$ be the number of proper patterns in $\mathcal{L}_n(Y_\alpha)$. Then

$$2^{\alpha_n n^2+O_n(n)} \leq P(n) \leq \text{poly}(n)\sum_{i=1}^{n} 2^{\alpha_i i^2+O_i(i)}$$

Proof. We first prove the leftmost inequality, by exhibiting a specific family of patterns with distinct extender sets. Fix $n \geq 0$, and define for any $j \geq n$ the set $P_j = \{(\langle n\rangle_0{}^{\uparrow}, \langle j\rangle_0{}^{\uparrow}, T_{\alpha_{n,j}}(n)^{\uparrow}, y^{(m)}, y^{(f)}) \in Y_\alpha\}$. By Claim 3, for $y_j \in P_j$ there are $\alpha_{n,j}n^2{+}O(n)$ symbols 1 in $\pi_{L_d}(y_j)|_{\mathcal{Q}_n}$. As $\alpha_{n,j} \to \alpha$, we can take $j \to \infty$ to obtain a set P_j of configurations which all have $\alpha_n n^2{+}O(n)$ symbols 1 in their $L_d{}^{\uparrow}$ layer. Therefore, for those j, we have that for any pattern $w \in P_j|_{\mathcal{Q}_n}$, there are $2^{\alpha_n n^2+O(n)}$ different free layers possible in P_j. Consider any two distinct such patterns u, v, differing in some position p in their layer L_f and containing no ■ on their layer L_m – such patterns exist, as the L_m layers extending u are of the form $[2n]_{-}$. Then, the configuration $[2n]_{p+(n,n) \mod (2n,2n)}$ extends $\pi_{L_m}(u)$ and marks in particular the position p. With this L_m layer, the patterns u, v are similar, and differ in a marked position, so by Claim 19 they have different extender sets. This gives $P(n) \geq 2^{\alpha_n n^2+O(n)}$

For the upper bound, we bound the number of extender sets by bounding the number of proper patterns $\{y|_{\mathcal{Q}_n}, y = (\langle i\rangle_{k_1}, \langle j\rangle_{k_2}, T(\beta, i)_{k_1}, [2i]_{m_1,m_2})\}$ – this is sufficient, as degenerate patterns are negligible by Claim 18. We make a rough overestimation, as we will pretend that for every such pattern w, and every position $p \in \mathcal{Q}_n$ such that $w_p^{(d)} = 1$, w can be extended in a configuration y marking p.

- If $i \leq n$, there are at most $\alpha_i i^2 + O_i(i)$ symbols 1 in the density layer of any $i \times i$ square, which are then (i, i)-periodic.
- If $i > n$, as α_i is decreasing we can simply bound the number of 1 by $\alpha_n n^2 + O(n^2)$.

Finally, we get:

$$\begin{aligned} P(n) &\leq \sum_{i=1}^{n}\sum_{k_1=0}^{i-1}\sum_{j=i}^{n}\sum_{k_2=0}^{j-1} 2^{\alpha_i i^2+O_i(i)} + \sum_{k_1=0}^{n}\sum_{k_2=0}^{n} 2^{\alpha_n n^2+O(n)} \\ &\leq \text{poly}(n)\sum_{i=1}^{n} 2^{\alpha_i i^2+O_i(n)} \end{aligned}$$

□

Proof of Theorem 2.48. As observed in Section 2.4.5, it suffices to prove the case $\alpha \in [0, 1]$. Then, by Lemma 2.49, the subshift Y_α defined in Section 2.4.5 is sofic. By Lemma 2.50 and Lemma 2.46, we also get $h_E(Y_\alpha) = \alpha$. □

2.4.6 A short note about syntactic monoids

As explained in Section 2.1, extender sets are usually important in the theory of formal languages as they are used to define the syntactic monoid of a language. As for any monoid or group, we can try to determine what its growth rate is, in the following sense:

Definition 2.51: Reduced length

Let $\mathcal{L} \subset \mathcal{A}^*$ be a language and $u \in \mathcal{L}$. The **reduced length** of u is $\|u\|_{\mathcal{L}} = \min_{v \sim_L u} |v|$.

Definition 2.52: Growth rate

Let $\mathcal{L} \subseteq \mathcal{A}^*$ be a language. The **exponential growth rate** of $M(\mathcal{L})$ is

$$h(M(\mathcal{L})) = \lim_{n \to +\infty} \frac{\log|\{[u] \in M(\mathcal{L}) \mid u \in \mathcal{L}, \|u\|_L \leq n\}|}{n}$$

In the case of groups rather than monoids, this object – more precisely, the study of the sequence of $|\{g \in G, \|g\|_G \leq n\}|$ where $\|g\|_G$ is defined relative to a presentation of the group G rather than *via* syntactic constructs – is of prime importance. For monoids, it is generally seen as a more anecdotal question, although some work has been done to understand the growth rates of some classes of monoids [Kam+24; IYN12]. In any case, we get the following result almost for free with our construction:

Theorem 2.53

The exponential growth rates of syntactic monoids of languages of effective $\mathbb{Z}$-subshifts are exactly the non-negative Π^3 real numbers:

$$\{h(M(\mathcal{L}(X))), X \text{ effective } \mathbb{Z}\text{-subshift}\} = \Pi^3 \cap \mathbb{R}^+$$

Proof. The key remark is that for the specific case of Z_α constructed in the proof of Theorem 2.43, for almost all patterns $w \in \mathcal{L}(Z_\alpha)$ we have $\|w\| = |w|$, and so $h(M(\mathcal{L}(Z_\alpha))) = h_E(Z_\alpha)$. More precisely, proper patterns $w \in \mathcal{L}(Z_\alpha)$ whose first and second layers contain at most one $*$ are already reduced, so the lower bound of Lemma 2.45 holds in the context of the syntactic monoid by the same proof, and the overestimation made for the upper bound also holds, as we prove it by counting patterns rather than extender sets anyway. Hence, the proof of Lemma 2.45, and therefore of Theorem 2.43, also prove that $h(M(\mathcal{L}(Z_\alpha))) = h_E(Z_\alpha) = \alpha$ for any $\alpha \in \Pi^3 \cap \mathbb{R}^+$. □

2.5 Summary

We give in Figure 2.10 a summary of the computability characterizations presented in this chapter.

<table>
<tr><th></th><th>$\mathbb{Z}$</th><th>$\mathbb{Z}^d, d \geq 2$</th></tr>
<tr><td>SFT</td><td colspan="2">$\{0\}$ (Corollary 2.28)</td></tr>
<tr><td>Sofic</td><td>$\{0\}$ (Corollary 2.29)</td><td>Π_3 (Theorem 2.48)</td></tr>
<tr><td>Effective</td><td colspan="2">Π_3 (Theorem 2.43)</td></tr>
<tr><td>Computable</td><td colspan="2">Π_2 (Theorem 2.47)</td></tr>
<tr><td>Sofic and minimal</td><td colspan="2">$\{0\}$ (Corollary 2.40)</td></tr>
<tr><td>Effective and 1-Mixing</td><td>Π_3 (Proposition 2.42)</td><td></td></tr>
<tr><td>Effective and minimal</td><td colspan="2">Π_1 (Proposition 2.41)</td></tr>
</table>

Figure 2.10: Characterizations of the extender entropies realized by some classes of subshifts.

Chapter 3

The projective fundamental group of subshifts

This chapter will focus on a conjugacy invariant called the **Projective Fundamental Group** of subshifts, introduced by Geller and Propp in [GP95]. The main motivations behind the introduction of this invariant come from two different directions, highlighted in Section 3.1:

- Given a subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ (not necessarily of finite type) and a partial configuration $x \in \mathcal{A}^{\mathbb{Z}^d \setminus D}$ for some bounded D, can we fill x in a valid configuration of X ? Variants of this question are also of interest, for example one might allow a finite number of changes in x before filling it.

- Given a subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$ and a configuration $x \in \mathcal{A}^{\mathbb{Z}^d}$, does there exist a block map Φ such that $\Phi(x) \in X$? This is a very general question, and needs to be refined to be answered meaningfully. Usually, one does not consider arbitrary configurations $x \in \mathcal{A}^{\mathbb{Z}^d}$, but instead focuses on those that are "defective" relative to X (in a sense explained below).

The main result of this chapter is that any finitely presented group is the projective fundamental group of some $\mathbb{Z}^2$-SFT, and will be proven in Section 3.5. The rest of the chapter will be devoted to studying some other properties of this group, or closely-related questions:

- Section 3.2 will formally define the projective fundamental group, and give a few examples and computations in easy cases.

- Section 3.3 will focus on a property, called **projective connectedness**, which is analogous to path-connectedness in the usual topological setting, and state some sufficient conditions for a subshift to be projectively connected.

- Section 3.4 will study a specific class of subshifts, the Hom shifts, which are particularly well-suited for the algebraic approach to tilings taken in this chapter.

- Section 3.5 will show that a large class of groups, namely finitely presented groups (see Definition 1.89) can be realized as the projective fundamental groups of $\mathbb{Z}^2$ SFTs.

Some results, mainly of Section 3.5, have been published in [PV23].

3.1 Filling holes and patching defects

We start this section with a very famous and classical combinatorial problem, attributed to Max Black: using standard 2×1 and 1×2 dominoes, can you tile the "mutilated 8×8-chessboard" depicted in Figure 3.1 – in the sense that dominoes do not overlap, or fall out of the bounds of the mutilated chessboard ?

Presented as above, the problem is made easier by the presence of black and white cells on the chessboard: of course, this does not change the solution, as the only constraints are of geometric nature, but it makes it easier to notice that the "mutilation" consisted in removing two black squares. As a domino must cover exactly one black and one white square, this counting (or colouring) argument shows that it is impossible to tile the mutilated chessboard with dominoes. This simple problem can be refined much further, and an advanced version of this colouring argument will be our first starting point to introduce our main object, the projective fundamental group.

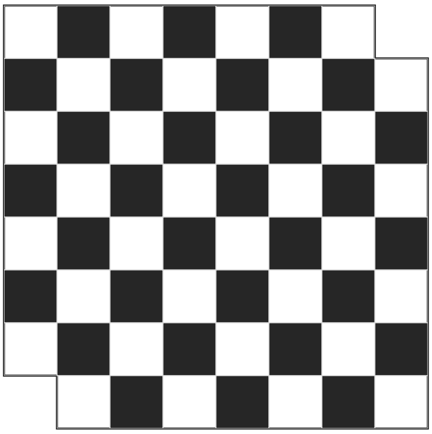

Figure 3.1: The 8×8 mutilated chessboard

3.1.1 Conway's tiling group

Conway and Lagarias introduced in [CL90] algebraic objects associated with a tileset, called the tile boundary group and tile homology group, which generalizes all the counting or colouring arguments of this kind. In particular, [CL90, Section 5] explicitly considers the example of the mutilated chessboard. In, [Thu90], Thurston then considers a number of other problems, working with shapes more complex than dominoes to try to tile some specific surfaces. We do not reproduce the full constructions and arguments, but rather try to give an overview of the main idea. As we will mainly be working with Wang Tiles, the point of view that we adopt is different from the original presentation that can be found in [CL90]. A lighter introduction to the initial constructions can be found in [Pro97].

Recall that a set of Wang Tiles is a set $T = \{(t_W, t_S, t_E, t_N) \mid t \in T\}$ where each t_D for $t \in T, D \in \{W, S, E, N\}$ is a colour belonging to some set C. For simplicity, we can always assume that the vertical colours $\{(t_W, t_E) \mid t \in T\}$ and the horizontal colours $\{(t_S, t_N) \mid t \in T\}$ are disjoints. We can now define the **tiling group** of T:

Definition 3.1: Tiling Group [Thu90]

For a Wang tileset T, we define the **tiling group** of T as

$$\Gamma(T) = \left\langle C \mid t_W^{-1} t_N^{-1} t_E t_S, t = (t_W, t_S, t_E, t_N) \in T \right\rangle$$

For a Wang tileset T, the tiling group $\Gamma(T)$ is a finitely presented group, where elements of the groups are sequences of colours of C – or their formal inverses. The relation $t_W^{-1} t_N^{-1} t_E t_S = 1_{\Gamma(T)}$ can then be seen as follows: starting from the bottom-left corner of the tile $t = (t_W, t_S, t_E, t_N)$, we write in order the colours encountered when going counterclockwise around t, counting a colour $c \in C$ as itself if we encounter it when moving from left to right or from bottom to top, and as c^{-1} otherwise.

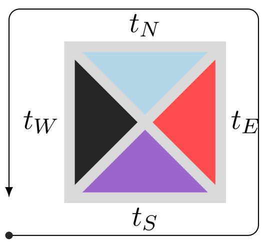

This group can be used to define the **contour words** of tilings of simply connected finite subset $D \subset_f \mathbb{Z}^2$ (see [MRR01] for a general introduction and several results): starting from one of the boundary vertices of D, the contour word associated with a tiling $x : D \to T$ is the element of $\Gamma(T)$ obtained when following the edges of D clockwise:

Example 8 (Dominoes – contour word)**.** *We give an example of a contour word for a region tiled by dominoes. As we work in the framework of Wang tiles, we need to explicitly define what we mean by a domino. Consider the following set T of Wang tiles:*

In a tiling, those tiles must be paired as follows to produce what we call dominoes:

We already distinguished the vertical and horizontal colours, in order to properly define the tiling group. More precisely, in order to make the computations readable, we use the following symbols instead of colours: H for ■, h for □, V for 🟪 and v for 🟦. The tiling group can then be written as follows, where the relations are given in the same order as the tiles above:

$$\Gamma(T) = \left\langle H, V, h, v \mid V^{-1}h^{-1}VH, V^{-1}H^{-1}vH, V^{-1}H^{-1}Vh, v^{-1}H^{-1}VH \right\rangle$$

Consider now the tiled region shown in Figure 3.2:

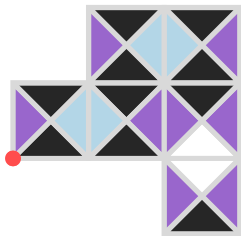

Figure 3.2: A region tiled by dominoes represented as Wang tiles.

The contour word associated to this tiled region, starting from the marked vertex, is $V^{-1}H^{-1}V^{-1}H^{-2}V^3HV^{-1}H^2$. Some easy rewriting using the relations of $\Gamma(T)$ show that this is the identity element of the group:

$$\begin{aligned}
V^{-1}H^{-1}V^{-1}H^{-2}V^3HV^{-1}H^2 &= V^{-1}H^{-1}V^{-1}H^{-2}V^3HV^{-1}(h^{-1}h)H^2 \\
&= V^{-1}H^{-1}V^{-1}H^{-2}V^2hH^2 \\
&= V^{-1}H^{-1}V^{-1}H^{-2}V^2h(V^{-1}H^{-1}HV)H^2 \\
&= V^{-1}H^{-1}V^{-1}H^{-2}VHVH^2 \\
&= \dots \\
&= 1_{\Gamma(T)}
\end{aligned}$$

As this word does not depend on tiles other than those at the boundary of D, we can also associate a contour word to any finite polygonal domain D with colours on its boundary, rather than an actual tiling of D (see Example 9 and in particular Figure 3.3

for an example). This word obviously depends on the starting point of the contour of this polygonal domain, but different starting points generate words that are conjugate in $\Gamma(T)$. This is generally not important, as one of the most important results about those contour words is the following:

Proposition 3.2: Conway's criterion

Let T be a Wang tileset, $D \subset_f \mathbb{Z}^2$ a finite simply connected domain with coloured boundary. Then, D can be tiled only if its contour word is the identity of $\Gamma(T)$.

In general, this is not an equivalence – that is, there exists Wang tilesets and finite regions with coloured boundaries that cannot be tiled, but whose contour word is trivial. Nevertheless, some refinements of this idea can give information on the tilings of the region, for example, the number of valid tilings, or their structure (see for example [MRR01, Propositions 3, 4, 5]).

Example 9 (Mutilated chessboard – tiling group)**.** *We show using contour words that the mutilated chessboard cannot be tiled by dominoes. We follow the computations of [Sch98, Proposition 6.1], and refer to it for the proofs of the claims. The Wang tileset used to define dominoes is the one of Example 8.*

Claim 20. $H^2 = h^2, V^2 = v^2$, and H^2, V^2 are in the center of Γ (that is to say, for any $g \in \Gamma(T)$, we have $gH^2 = H^2 g$).

Claim 21. Let $\Gamma' = \Gamma(T)/\langle H^2, V^2 \rangle$, let h', v', H', V' be the images of respectively h, v, H, V in the quotient and let $\phi\colon \Gamma' \to GL_3(\mathbb{Z})$ be the following map:

$$\phi(h') = \begin{pmatrix} -1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad \phi(H') = \begin{pmatrix} -1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

$$\phi(v') = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \end{pmatrix}, \qquad \phi(V') = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & -1 \end{pmatrix}$$

Then ϕ is an injective group morphism.

Now, in order to show that the mutilated chessboard Figure 3.1 cannot be tiled by dominoes, we want to apply Proposition 3.2:

In Figure 3.3, we can see a mutilated chessboard with the only possible colours on the border that would result from a tiling by dominoes, when represented as Wang tiles from T. The contour word g obtained from this, starting from the marked point on the figure, is $g = V^{-1}H^{-1}V^{-7}H^{-7}VHV^7H^7$. By Claim 20, we have $g = (VH)^4(V^2H^2)^{-2}$. By Claim 21, it suffices to show that $\phi((V'H')^4) \neq I_3 \in GL_3(\mathbb{Z})$ to show that the mutilated chessboard is not tileable by dominoes. As $\phi((V'H')^4) = \begin{pmatrix} 1 & 0 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ we have the desired claim.

Although the Conway tiling group is a powerful algebraic tool to study whether or not some concrete regions can be tiled by a given Wang tileset, the group $\Gamma(T)$ is not

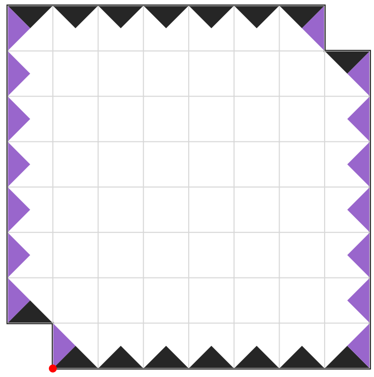

Figure 3.3: A mutilated chessboard with colours on its border, corresponding to a possible tiling by complete dominoes.

a conjugacy invariant of the subshift X_T: for example, adding a tile with a fresh colour which does not appear in any valid infinite tiling adds a generator to $\Gamma(T)$, and there exist completely different tilesets T' for which we neither have $\Gamma(T) \triangleleft \Gamma(T')$ nor $\Gamma(T') \triangleleft \Gamma(T)$. This is one of the shortcomings that the projective fundamental group tries to solve.

3.1.2 Defects in tilings

Another motivation to study the projective fundamental group of subshifts comes from a different perspectives, namely, studying "defects" in tilings. This approach has mostly been used in [Piv07]. We will mainly adopt the perspective of this article, as there is a very rich literature studying similar problems, with no unanimous definitions or motivating examples. Broadly speaking, given a subshift X, the idea is to study configurations which are "almost" in X, up to some local "defects", and in particular, whether some specific classes of defects can be corrected by a simple process (mainly block maps, or any other combinatorial process, with the goal of modifying x in as few positions as possible so that it now belongs to X). Let us be a bit more precise:

Definition 3.3: Defects [Piv07, Sec. 1]

Let $X \subset \mathcal{A}^{\mathbb{Z}^2}$ be a subshift and $x \in \mathcal{A}^{\mathbb{Z}^2}$. The **defect field** of x is the map

$$\begin{aligned} \mathcal{F}_x \colon \mathbb{Z}^2 &\to \mathbb{N} \cup \{\infty\} \\ \mathbf{u} &\mapsto \max\{r \mid x|_{\mathcal{B}_r+\mathbf{u}} \in \mathcal{L}_r(X)\} \end{aligned}$$

The **defect set** of x is the set of local minima of $\mathcal{F}_x$.
For $R \geq 0$, the R-defect points is $\mathcal{D}_R(x) = \{\mathbf{u} \in \mathbb{Z}^2, \mathcal{F}_x(\mathbf{u}) \geq R\}$.
A configuration x is **defective** if $x \notin X$ but $\sup_{\mathbf{u}\in\mathbb{Z}^2} \mathcal{F}_x(\mathbf{u}) = \infty$.

Defective configurations are configurations which contain arbitrary large balls which are globally admissible patterns of X. In [Piv07], the author tries to understand the subshifts X, and the kind of defective configurations x, which are such that there exists block maps ϕ such that $\phi(x) \in X$. More precisely, we have the following important proposition:

Proposition 3.4 [Piv07, Prop. 1.2]

Let $X \subset \mathcal{A}^{\mathbb{Z}^2}$ be a subshift, and let $\phi\colon \mathcal{A}^{\mathbb{Z}^2} \to \mathcal{A}^{\mathbb{Z}^2}$ be a block map such that $\phi(X) \subseteq X$. Let $x \in \mathcal{A}^{\mathbb{Z}^2}$ be a defective configuration. Then $\phi(x)$ is either defective or in X.

In particular, this motivates the following definition, further characterizing defective configurations:

Definition 3.5

Let $X \subset \mathcal{A}^{\mathbb{Z}^2}$ be a subshift, and let $\phi\colon \mathcal{A}^{\mathbb{Z}^2} \to \mathcal{A}^{\mathbb{Z}^2}$ be a block map such that $\phi(X) = X$. Let $x \in \mathcal{A}^{\mathbb{Z}^2}$ be a defective configuration.

- If $\phi^n(x)$ is defective for all $n > 0$, then x has a **ϕ-persistent** defect.
- If there exists $R \geq 0$ and $y \in X$ such that $x|_{\mathcal{D}_R(x)} = y|_{\mathcal{D}_R(x)}$, then x has a removable defect; otherwise, it has an essential defect.

Persistent defects are defects that cannot be removed by applying block maps; essential defects are defects that cannot be removed by modifying a finite region around some specific points of the configuration. One can then wonder whether a defective configuration a has a ϕ-persistent defect for any block map ϕ, and if a has an essential defect. It turns out that in a number of (technical) cases, the projective fundamental group of X can help answer those questions. We refer to [Piv07] for the precise statements, as we would need to introduce significantly more objects in order to simply state the main results – see in particular the propositions and theorems 2.11, 2.15, 3.8 and 4.10 of this article.

3.2 Projective fundamental group: adaptation to subshifts

The ideas of Section 3.1 are quite similar to another classical construction in topology, known as the fundamental group. This is a homeomorphism invariant, associated with any topological space X and denoted $\pi(X)$ – with a small caveat detailed in Section 3.2.1 – which tries to capture the structure of "holes" in the space, by considering paths and loops winding around those holes. A complete introduction to the general field of algebraic topology can be found in [Hat00], and we simply give the basic definition to highlight the similarities and links with the projective fundamental group.

3.2.1 The classical fundamental group

The idea of the fundamental group, and algebraic topology in general, is to associate algebraic structures with topological spaces, intuitively capturing the shape, dimension and relative disposition of "holes" in the space. Here, "holes" are to be understood at an intuitive level, as a formal definition would otherwise end up being circular, defining holes as what is being measured by the invariant of algebraic topology such as homotopy or homology groups. In our particular setting, we will only consider the first homotopy group, called the **fundamental group**, that we now define. This section is a rather short introduction, and adopts a classical point-of-view - that is, we do not use any category theory, or more abstract axiomatization of algebraic topology invariants. An introduction to those topics with a more categorical and "modern" point-of-view than [Hat00] can be found in [May99] or [Die08].

Paths and loops

The first notion that we need is the one of **paths**:

Definition 3.6: Path - topology

Let X be a topological space. A **path** p in X is a continuous map $p\colon [0,1] \to X$.
$p(0), p(1)$ are the endpoints of the path p. If $p(0) = p(1)$, p is called a **loop**, based at $p(0)$.

Paths need not be injective, non-constant, or any other restriction. Paths can nonetheless be quite complicated objects: paradigmatic examples are the numerous space-filling curves, such as the Hilbert Curve [Hil35], which are surjective paths $p\colon [0,1] \to [0,1]^2$. There is however a nice structure on the set of paths:

Definition 3.7: Path concatenation

Let $p, q\colon [0,1] \to X$ be two paths. Suppose that $p(1) = q(0)$. We then define the **concatenation** $p * q$ as the path:

$$\begin{aligned} p * q\colon [0,1] &\to X \\ t &\mapsto \begin{cases} p(2t) & \text{if } t \leq \frac{1}{2} \\ q(2t-1) & \text{otherwise} \end{cases} \end{aligned}$$

The concatenation of p, q can be seen as the path which follows p and then q, with a renormalization so that it is still a proper path $[0,1] \to X$. It is easy to see that it is indeed a path, as $p * q$ is continuous.

We can define a last operation on paths:

Definition 3.8: Path inverse

Let $p\colon [0,1] \to X$ be a path. The **inverse** of p is the path

$$\begin{aligned} p^{-1}\colon [0,1] &\to X \\ t &\mapsto p(1-t) \end{aligned}$$

With those two operations, one can define a group, in fact a group per point $x_0 \in X$:

Proposition 3.9

Let X be a topological space and $x_0 \in X$ be a **basepoint**. Let $L = \{\text{loops of } X \text{ based at } x_0\}$. Then, $(L, *)$ is a group.

This group is however rather unwieldy. Indeed, there are in general uncountably many loops based at x_0 even in simple cases. Moreover, most of those loops ought to be intuitvely considered the same, and do not carry different "topological" information on the space –

for example, taking $X = [0,1], x_0 = 0$, all the loops $t \mapsto \sin(2\pi x^\alpha), \alpha > 0$, which simply go from 0 to 1 and back at different "speeds".

Homotopy and the fundamental group

For this reason, we will consider a quotient of this space, by an operation called **homotopy**, or more informally, deformation:

Definition 3.10: Path homotopy

Let $p, q\colon [0,1] \to X$ be two paths. We say that they are **homotopic**, and write $p \sim q$, if $p(0) = q(0)$ and $p(1) = q(1)$ and if there exists a continuous map $H\colon [0,1]^2 \to X$, called a **path homotopy**, such that:

- $\forall t \in [0,1], H(t,0) = p(t)$, that is, $H(\cdot, 0) = p$
- $\forall t \in [0,1], H(t,1) = q(t)$, that is, $H(\cdot, 1) = q$
- $\forall x \in [0,1], H(0,x) = p(0) = q(0)$ and $H(1,x) = p(1) = q(1)$, that is, $H(\cdot, x)$ is a also a path between $p(0)$ and $q(0)$

We denote $[p]$ the homotopy class of a path p.

This is a special case of the following more general definition:

Definition 3.11: Homotopy

An homotopy is a continuous map $H\colon X \times [0,1]\colon Y$. The map $H(t,\cdot)\colon X \to Y$ is sometimes denoted H_t.

The key observation is the following:

Lemma 3.12

Let p_0, p_1, q_0, q_1 be paths in X, such that:

- $p_0(0) = p_1(0), p_0(1) = p_1(1)$, that is p_0, p_1 have the same endpoints.
- $q_0(0) = q_1(0), q_0(1) = q_1(1)$, that is q_0, q_1 have the same endpoints.
- $p_0(1) = q_0(0)$, so that $p_i * q_j$ is well-defined for any $i, j \in \{0,1\}$.
- $p_0 \sim p_1, q_0 \sim q_1$

Then $(p_0 * q_0) \sim (p_1 * q_1)$.

Corollary 3.13

The concatenation $*$ extends to loop homotopy classes: for any paths p, q with $p(1) = q(0)$, we have $[p * q] = [p] * [q]$.

We are ready to define the main object of this section:

Definition 3.14: Fundamental group

Let X be a topological space, and $x_0 \in X$. The fundamental group based at x_0 is the group $\pi_1(X, x_0)$ of loops based at x_0, with the concatenation operation: denoting $L = \{\text{loops based at } x_0\}$, we define

$$\pi_1(X, x_0) = (\{[p], p \in L\}, *)$$

Proposition 3.15

The fundamental group is a homeomorphism invariant.

Proof. It is easy to check that homeomorphisms send paths and path homotopies to paths and path homotopies respectively. Hence, loop homotopy classes are sent to loop homotopy classes. □

In some cases, we can avoid the reference to the basepoint, if we are only interested to $\pi_1(X, x_0)$ up to group isomorphism:

Definition 3.16: Path-connectedness

A topological space X is **path-connected** if for any $x, y \in X$, there exists a path $p\colon [0,1] \to X$ such that $p(0) = x$, $p(1) = y$.

Proposition 3.17

If X is path-connected, then for any $x_0, y_0 \in X$, $\pi_1(X, x_0) \simeq \pi_1(X, y_0)$, and we write $\pi_1(X)$ without reference to a basepoint.

Proof. Let $x_0, y_0 \in X$. As X is path connected, consider any path γ from x_0 to y_0, and define

$$\begin{aligned} \phi\colon \pi_1(X, x_0) &\to \pi_1(X, y_0) \\ [p] &\mapsto [\gamma * p * \gamma^{-1}] \end{aligned}$$

By Corollary 3.13, ϕ is well-defined, and as it is a conjugacy, it is in particular an isomorphism. □

Example 10. $\pi_1(\mathbb{R}) = \{e\}$.

Proof. As X is clearly path-connected, it suffices to show that $\pi_1(\mathbb{R}, 0) = \{e\}$. We will show this by explicitly defining a path homotopy between any loop p based at 0 and the trivial loop $[0,1] \to 0$. Let then p be a loop based at 0 and define

$$\begin{aligned} H\colon [0,1]^2 &\to X \\ (t, x) &\to xp(t) \end{aligned}$$

Then: $H(\cdot,0) = 0 \times p = 0$, $H(\cdot,1) = 1 \times p = p$, and for any x we have:

$$H(0,x) = x \times p(0) = x \times 0 = 0 = H(1,x).$$

Hence H is a path homotopy between the constant loop and p, so there is a single loop class and $\pi_1(X)$ is trivial. □

In general, computing this fundamental group, even in simple cases, can be quite hard, even when the group itself is *e.g.* finite or free. A typical example, motivating the introduction of a key tool that we will adapt and use in the setting of subshifts in Section 3.4, is the circle. To compute $\pi_1(S^1) = \pi_1([0,1]/(0 \sim 1))$, we need the notion of **covering**.

Covering spaces

Definition 3.18: Covering space

Let X be a topological space. A **covering** of X is a space Y and a map $\rho\colon Y \to X$ such that for any point $x \in X$, there exists a neighbourhood $x \in U \subset X$, and a discrete space D_x, such that:

- $\rho^{-1}(U) = \bigsqcup_{i \in D_x} V_i$, that is, $\rho^{-1}(U)$ is a disjoint union of open sets.
- For any $i \in D_X$, $\rho|_{V_i} \to U$ is an homeomorphism.

ρ is a **covering map**. We sometimes use "covering" for ρ or Y alone, when the other one is clear from context.

One can think of a covering (Y,ρ) as an "unrolled" version of the base space X: indeed, around each y point of Y one can find a small neighbourhood V which is isomorphic to $\rho(V) \subseteq X$. We will often call **lift** of (a point, a path ...) x the object $\tilde{x}$ such that $\rho(\tilde{x}) = x$, where ρ might have been extended to *e.g.* $[0,1] \times Y \to [0,1] \times X$, or later $Y^{\mathbb{Z}^2} \to X^{\mathbb{Z}^2}$ in the obvious way.

In particular, we have the following important theorem:

Proposition 3.19 [Hat00, Prop. 1.30]

If $\rho\colon \tilde{X} \to X$ is a covering, then for any homotopy $H\colon Y \times [0,1] \to X$ and map $\tilde{H}_0\colon Y \to \tilde{X}$ lifting $H_0 = H(\cdot,0)$, then there exists a unique homotopy $\tilde{H}\colon Y \times [0,1] \to \tilde{X}$ of $\tilde{H}_0$ lifting H.

As this holds for any space Y, we obtain the following interesting special cases:

Proposition 3.20

Let $\rho\colon \tilde{X} \to X$ be a covering. Let $p, p'\colon [0,1] \to X$ be two homotopic paths starting at $x_0 \in X$. Then, for any lift $\tilde{x}_0 \in \tilde{X}$ of x_0:

- There exists a unique path $\tilde{p}\colon [0,1] \to \tilde{X}$ lifting p starting at $\tilde{x}_0$
- $\tilde{p}$ and $\tilde{p}'$ are homotopic paths.

Proof. For the first point, we apply Proposition 3.19 with Y a one-point space $\{y\}$, so that $p\colon Y \to X$ can be viewed as any arbitrary point $p(y) = x_0 \in X$.

For the second point, we apply Proposition 3.19 with $Y = [0,1]$, so that $p\colon Y \to X$ is a path. The fact that $\tilde{p} \sim \tilde{p}'$ is a consequence of the fact that the homotopy preserves the paths' endpoints at all time, that is to say, for all $x \in [0,1]$ we have $H(0,x) = p(0) = p'(0)$, and so $\tilde{H}(0,x) = \tilde{p}(0) = \tilde{p}'(0)$. □

More details can be found in [Hat00, Lifting Properties]. The last proposition we state explains why we needed to introduce covering spaces in order to study fundamental groups. A few key properties are hidden here (for example, why we need to consider covering spaces with trivial fundamental groups, whether those exist at all, or whether it is unique up to isomorphism), but we simply want to highlight part of the link between coverings and the fundamental group itself:

Proposition 3.21 [Hat00, Prop. 1.39]

Let $\rho\colon \tilde{X} \to X$ be a covering, and suppose that $X, \tilde{X}$ are path-connected, X is locally path-connected, and that $\pi_1(\tilde{X})$ is the trivial group. Then, $\pi_1(X)$ is isomorphic to the set of homeomorphisms d of $\tilde{X}$ satisfying $\rho \circ d = \rho$.

This gives us the following non-trivial example:

Example 11 ([Hat00, Thm 1.7])**.** $\pi_1(S^1) = \mathbb{Z}$.

Proof. As a covering space with trivial fundamental group for S^1, we can consider $\mathbb{R}$ with the map $\rho\colon x \in \mathbb{R} \mapsto (\cos 2\pi x, \sin 2\pi x) \in S^1$ – in fact, it is a theorem that this is the only such covering up to isomorphism.

Now, let d be a homeomorphism d of $\mathbb{R}$ satisfying $\rho \circ d = \rho$. In particular, for any $x \in \mathbb{R}$, there exists $n_x \in \mathbb{Z}$ such that $d(x) = x + n_x$. But as d is continuous, we must have n_x independent from x, and so d is simply a translation by some $n \in \mathbb{Z}$, written d_n. As any translation by $m \in \mathbb{Z}$ induces such a homeomorphism $d_m\colon \mathbb{R} \to \mathbb{R}$, and as $d_m \circ d_{m'} = d_{m+m'}$, the group of those homeomorphisms is $\mathbb{Z}$. By Proposition 3.21, we get $\pi_1(S^1) \simeq \mathbb{Z}$. Intuitively, each loop in S^1 is homotopic to one which winds exactly n times around the circle. Said differently, $\pi_1(X)$ is generated by the loop $p\colon x \in [0,1] \mapsto (\cos 2\pi x, \sin 2\pi x)$. □

This strategy – using covering spaces in order to compute fundamental groups – will also be used in the case of the projective fundamental group of subshifts, in particular in Section 3.4.

3.2.2 Definition of the group and links with other notions

In the case of subshifts, we have a first technical difficulty: although a subshift X is a compact topological space, it is totally disconnected, meaning in particular that the path-connected components are trivial, and therefore paths in X are constant. To go around this difficulty, the idea introduced in [GP95] is to consider a (projective, or inverse) limit of fundamental groups $\pi_1(X_n)$, defined for spaces $(X_n)_{n\in\mathbb{N}}$ which "converge" to X itself in some sense. We will present two equivalent point of views on the projective fundamental group, the first one (Section 3.2.2) closely resembling the classical objects introduced in Section 3.2.1, and the second one (Section 3.2.2) more combinatorial in nature and similar to other classical constructions in symbolic dynamics.

In what follows, we keep some of the terminology of the original article of Geller and Propp [GP95].

An actual fundamental group of scene spaces

For now, let us stick to the formal, topological definitions rather than the combinatorial ones. We will call **aperture window** a bounded subset $B \subset \mathbb{R}^d$. Instead of considering configurations $x \in \mathcal{A}^{\mathbb{Z}^d}$ as colourings of the discrete grid, we view them as colourings $\tilde{x}$ of the euclidean plane $\mathbb{R}^d$ as follows: for any $\mathbf{z} =\in \mathbb{R}^d, \tilde{x}_{\mathbf{z}} = x_{\lfloor \mathbf{z} \rfloor}$, where we compute $\lfloor \mathbf{z} \rfloor$ coordinate-wise. We drop the notation $\tilde{x}$ and use x even for $\mathbb{R}^d$ configurations when context makes it clear. We naturally extend the shift $\sigma_{\mathbf{v}}$ to those configurations for any $\mathbf{v} \in \mathbb{Z}^d$. We then define $\tilde{X} = \{\tilde{x}, x \in X\}$. Note in particular that $\tilde{X}$ is not an $\mathbb{R}^d$ subshift in the general sense, as we only consider the $\mathbb{Z}^d$ action on configurations.

Notation. For any two configurations $x, y \in \mathcal{A}^{\mathbb{R}^d}$, point $\mathbf{z} \in \mathbb{R}^d$, and aperture window $B \subset \mathbb{R}^d$, define the equivalence relation $\sim_B$ as

$$(x, \mathbf{z}) \sim_B (y, \mathbf{z}) \iff x\big|_{\mathbf{z}+B} = y\big|_{\mathbf{z}+B}$$

In particular, $(x, \mathbf{z}) \not\sim_B (x, \mathbf{z}')$ for $\mathbf{z} \neq \mathbf{z}'$, and $(x, \mathbf{z}) \not\sim_B (\sigma_{\mathbf{u}}(x), \mathbf{z} + \mathbf{u})$ for $\mathbf{u} \neq \mathbf{0}$.

Definition 3.22: Scene space

For a subshift $X \subseteq \mathbb{Z}^d$ and aperture window $B \subset \mathbb{R}^d$, define the B-**scene space** as $S_B(X) = (\tilde{X} \times \mathbb{R}^d)/\sim_B$.

Whenever X is clear from the context, we will simply write S_B. Each point of S_B can then be seen as a pattern (*not* up to translation), whose support is some translation of B. As for any topological space, we can now try to compute the fundamental group of S_B. In general, we will not however directly compute $\pi_1(S_B, (x|_{B+\mathbf{z}}, \mathbf{z}))$:

- The space S_B is a “complicated” space, and we do not have any straightforward way to compute its fundamental group in general, even for simple subshifts X such as SFTs.
- It is *a priori* unclear how $\pi_1(S_B, (x|_{B+\mathbf{z}}, \mathbf{z}))$ depends on B.
- It is also unclear how $\pi_1(S_B, (x|_{B+\mathbf{z}}, \mathbf{z}))$ depends on the concrete space X being considered, rather than the conjugacy class of X in the space of subshifts.

For all those reasons, we will need to use another technical construction that we introduce now. The definitions are given in a general setting, and will be reframed for the specific setting of computing projective fundamental groups of $\mathbb{Z}^d$ subshifts afterwards.

Definition 3.23: Restriction maps

Let $B' \subseteq B \subset \mathbb{R}^d$. Define the **canonical restriction map**

$$\begin{aligned} \mathrm{restr}_{B',B} \colon \mathcal{A}^B &\to \mathcal{A}^{B'} \\ P &\mapsto (\mathbf{z} \in B' \mapsto P(\mathbf{z})) \end{aligned}$$

We naturally extend those maps to $\bigcup_{\mathbf{z} \in \mathbb{R}^d} \mathcal{A}^{\mathbf{z}+B}$, so that if $\mathrm{dom}(P) = B + \mathbf{z}$ then $\mathrm{dom}(\mathrm{restr}_{B',B}(P)) = B' + \mathbf{z}$, and to S_B so that $\mathrm{restr}_{B',B} \colon S_B \to S_{B'}$.

Those maps forget the part $B \setminus B'$ of a pattern of support B. In particular, they are all surjective, but not necessarily injective, if a B'-supported pattern can be extended in more than one way.

Using this and a general algebraic construction, we can define a single group associated to X, called the Projective Fundamental Group. We give here a first definition, and another point of view, which will be shown to be equivalent, will be given in Section 3.2.2:

Definition 3.24: Projective limit

Let $(\mathcal{I}, \leq)$ be a partially ordered set, and let $(G_i)_{i\in\mathcal{I}}$ be a family of sets. Suppose that we have a family of functions $f_{ij} \colon G_j \to G_i$ for all $i \leq j \in \mathcal{I}$, satisfying:

- For all $i \in \mathcal{I}$, $f_{i,i} = \mathrm{id}_{G_i}$
- For all $i \leq j \leq k \in \mathcal{I}$, $f_{i,k} = f_{i,j} \circ f_{j,k}$

Then, the **projective limit** (or **inverse limit**) of the projective (or inverse) system $((G_i)_{i\in\mathcal{I}}, (f_{i,j})_{i\leq j\in\mathcal{I}})$ is a subset of the potentially infinite direct product of all the G_i's, denoted here by G_∞

$$G_\infty = \varprojlim_{i\in\mathcal{I}} G_i = \left\{ \vec{g} \in \prod_{i\in\mathcal{I}} G_i \;\middle|\; g_i = f_{i,j}(g_j) \text{ for all } i \leq j \text{ in } \mathcal{I} \right\}$$

If the G_i's are groups, and the $f_{i,j}$'s are homomorphisms, then G_∞ is a group with the group operation defined pointwise.

In order to lighten the notations, we might write $\pi_1(S_B, x_0)$ for $\pi_1(S_B, (x_0|_B, \mathbf{0}))$ where $\mathbf{0} = (0, \ldots, 0) \in \mathbb{R}^d$.

Lemma 3.25

If $B' \subseteq B \subset \mathbb{R}^d$ and $p \sim_B p'$ are homotopic paths in S_B, then $\mathrm{restr}_{B',B}(p) \sim_{B'} \mathrm{restr}_{B',B}(p')$. In other words, for any $x_0 \in X$, $\mathrm{restr}_{B',B} \colon \pi_1(S_B, x_0) \to \pi_1(S_{B'}, x_0)$ is well-defined.

Proof. It H is a homotopy between p, p' then $\mathrm{restr}_{B',B} \circ H$ is a homotopy between $\mathrm{restr}_{B',B}(p)$ and $\mathrm{restr}_{B',B}(p')$. □

Definition 3.26: Projective Fundamental Group

Let X be a $\mathbb{Z}^d$ subshift, and $(B_n)_{n\in\mathbb{N}}$ an increasing sequence of subsets of $\mathbb{R}^d$, so that $\bigcup_{n\in\mathbb{N}} B_n = \mathbb{R}^d$.
We define the **projective fundamental group** of X based at $x_0 \in X$ as the inverse limit of the system $(\pi_1(S_{B_n}), (\mathrm{restr}_{B_n,B_m})_{0<m<n})$:

$$\pi_1^{proj}(X, x_0) = \varprojlim_{i\in\mathbb{N}} \pi_1(S_{B_n}, (x_0, \mathbf{0}))$$

The next proposition shows that the actual sequence of supports $(B_n)_{n\in\mathbb{N}}$ used in Definition 3.26 does not matter, as long as it is increasing and eventually covers the entire space $\mathbb{R}^d$. This is a special case of a very general statement about inverse limits:

Proposition 3.27

Let $(B_n)_{n\in\mathbb{N}}$ and $(B'_n)_{n\in\mathbb{N}}$ be two increasing sequences of $\mathbb{R}^d$ subsets, satisfying $\bigcup_n B_n = \bigcup_n B'_n = \mathbb{R}^d$. Then

$$\varprojlim_{n\in\mathbb{N}} S_{B_n} = \varprojlim_{n\in\mathbb{N}} S_{B'_n}$$

Remark. *Proposition 3.27 also implies that it is only necessary to specify the maps $f_{n,n+1}$ in order to entirely define the projective limit. Indeed, we can obtain map $f_{m,n}$ for $m < n$ by composing these ones. In particular, in order to show that an element $(x_1, x_2, \dots)$ belongs to the projective limit, it suffices to check that $f_{n,n+1}(x_n) = x_{n+1}$ for all n.*

Just as the fundamental group is a homeomorphism invariant (Proposition 3.15), we can show that the *projective* fundamental group is a conjugacy invariant for subshifts:

Theorem 3.28

Let X, Y be two subshifts and $\phi\colon X \to Y$ be a conjugacy map. Then, for any $x_0 \in X$, $\pi_1^{proj}(X, x_0) = \pi_1^{proj}(Y, \phi(x_0))$.

We defer the proof of this theorem to Section 3.3, in which we introduce and study in more details an equivalent to the path-connectedness property. In the meantime, we simply give an intuition as to why the projective fundamental group is a conjugacy invariant, while the Conway's tiling group (Definition 3.1) is not. The idea is that because we are looking at a limit of groups, which are defined using paths which "see" larger and larger patterns, the small-scale irregularities eventually disappear.

In this chapter, we will often use the "projective P" terminology for some object P (a path, a homotopy class ...). In our setting, this has to be understood as a sequence $(P_n)_{n>0}$, where each P_n or is an element of $S_{\mathcal{B}_n}$, and $\mathrm{restr}_{n,n+1}(P_{n+1}) = P_n$. In particular, a projective path class (or projective loop class) is a sequence of paths (or loops) that are all homotopy-equivalent after applying the suitable restriction map.

A combinatorial point of view

We keep the terminology introduced in Section 3.2.2, but use different definitions adapted to the case of $\mathbb{Z}^d$ subshifts. We will show that this is legitimate, as the notions coincide.

Definition 3.29: Path - subshifts

Let $B \subset_f \mathbb{Z}^d$ a finite subset called **aperture window**. A B-path is a finite sequence $(P_t, \mathbf{v_t})_{0\leq t\leq N}$ such that for any t with $0 \leq t \leq N$:

- P_t is an extensible pattern of X of support $B + \mathbf{v_t}$,
- $\mathbf{v_t}$ is adjacent to $\mathbf{v_{t+1}}$, *i.e.*, $\mathbf{d_t} = \mathbf{v_{t+1}} - \mathbf{v_t}$ has euclidean norm exactly

1,

- $P_t(u) = P_{t+1}(u)$ for any $u \in B \cap \sigma_{\mathbf{d_t}}(B)$, *i.e.*, consecutive patterns overlap,
- the pattern $P_t \cup P_{t+1}$ obtained by merging P_t and P_{t+1} is extensible in X.

The first and last element of the sequence are respectively called the **starting point** and the **ending point** of the path. If they are equal, the path is called a **loop**. The path $(P_{N-t}, \mathbf{v_{N-t}})_{0\leq t\leq N}$ is called its **inverse path**, denoted by p^{-1}.
We call **trajectory** of the path the sequence $(\mathbf{v_t})_{0\leq t\leq N}$.

This construction is similar to another, classical definition in symbolic dynamics, of the Rauzy graphs (see for example [Ber+15, Section 4] or [Pel+09]): keeping the notations of Definition 3.29, it is the graph whose vertices are patterns P_t, and there is an edge between two "coherent" patterns P_t and P_{t+1} labelled by $P_{t+1} \setminus P_t$. It is mainly used when studying one-dimensional subshifts, and in our notion of path, we also keep track of the position of the patterns within $\mathbb{Z}^d$.

In the remainder of this chapter, we let $B_n = [\![-n, n-1]\!]^2$. Note that $B_n \subsetneq \mathcal{B}_n$. The reason we choose this support if for symmetry reasons when dealing specifically with the fundamental group: the point $(0,0) \in \mathbb{R}^2$ is at the center of the square $(-n, n)^2$, which is the area covered by (open) unit tiles whose bottom-left corners are placed on the positions of B_n. This makes computations easier in the remainder of the chapter. Unless stated otherwise, all the aperture windows considered will be of this form.

The canonical restriction maps have the same definition as in the case of tilings of the euclidean plane:

Notation. For $0 < n \leq m$, we write $\mathrm{restr}_{n,m}$ the map

$$\begin{aligned}\mathrm{restr}_{n,m} \colon \mathcal{A}^{\mathcal{B}_m} &\to \mathcal{A}^{\mathcal{B}_n}\\ P &\mapsto P\big|_{\mathcal{B}_n}\end{aligned}$$

Two paths may be composed when the first one ends where the second one starts:

Definition 3.30: Path composition

Given $p = (P_t, \mathbf{v_t})_{0\leq t\leq N}$ and $p' = (P'_t, \mathbf{v'_t})_{0\leq t\leq N'}$ two paths such that $(P_N, \mathbf{v_N}) = (P'_0, \mathbf{v'_0})$ we denote by $p * p'$ the path

$$p * p' = (P_0, \mathbf{v_0}) \ldots (P_N, \mathbf{v_N})(P'_1, \mathbf{v'_1}) \ldots (P'_{N'}, \mathbf{v'_{N'}}).$$

This gives a quantitative way to look at paths, which will be useful in later proofs:

Definition 3.31: Coherent path

A path $p = (P_i, \mathbf{v_i})_{i\leq N}$ is coherent if all its patterns are equal on the points where their supports overlap, and furthermore, the pattern obtained by merging all the P_i is globally admissible in X. In that case, for any $x \in X$

containing $\bigcup_{i\leq N} P_i$, we say that p can be **traced** in x.

Definition 3.32: Coherent path decomposition

A **coherent decomposition** of a path p is a sequence $p_1, \ldots, p_L$ of coherent paths such that $p = p_1 * p_2 \ldots * p_L$, and L is called the **length** of the decomposition.

We can now define a corresponding homotopy notion, using this notion of coherent path:

Definition 3.33: Elementary deformation

Let $p = p_1 * p_2 * p_3$ be a path and suppose that p_2 is coherent and can be traced in some configuration $x \in X$. Then, for any p_2' traced in x with the same starting and ending point as p_2, the path $p_1 * p_2' * p_3$ is called an **elementary deformation** of p.

Note that as paths might be empty or consist of a single point, they can be deformed by inserting or removing loops traced in a single configuration at any step.

Using those elementary deformations, we define a general notion of homotopy:

Definition 3.34: Homotopy - subshifts

Two paths p, p' are said to be **homotopic** if there exists a finite sequence of elementary deformations from p to p'. This defines an equivalence relation between paths, and we denote by $[p]$ the equivalence class of p. If p and p' are paths with an aperture window $B \subset \mathbb{Z}^d$, we denote by $p \sim_B p'$ the fact that they are homotopic.

Remark. *By definition, when two paths are homotopic, they necessarily have the same starting and ending points. This coincides with Definition 3.10. When B is clear from the context, we will simply write $p \sim p'$.*

We give in Figure 3.4 an illustration of a projective path between $x, y \in X$.

With this definition of a path and of homotopy, we can define a fundamental group for each possible aperture window $B \subset \mathbb{Z}^d$.

Definition 3.35: Fundamental Group - subshifts

Let X be an SFT, $B \subset \mathbb{Z}^d$ an aperture window, $x_0 \in X$ and $\mathbf{v} \in \mathbb{Z}^d$. The **fundamental group** of X based at $(x_0, \mathbf{v})$ for the aperture window B, denoted by $\pi_1^B(X, (x_0, \mathbf{v}))$, is the group of all the equivalence classes of loops starting and ending at $(x_{0|B}, \mathbf{v})$ for the homotopy equivalence relation, along with the $*$ operation.

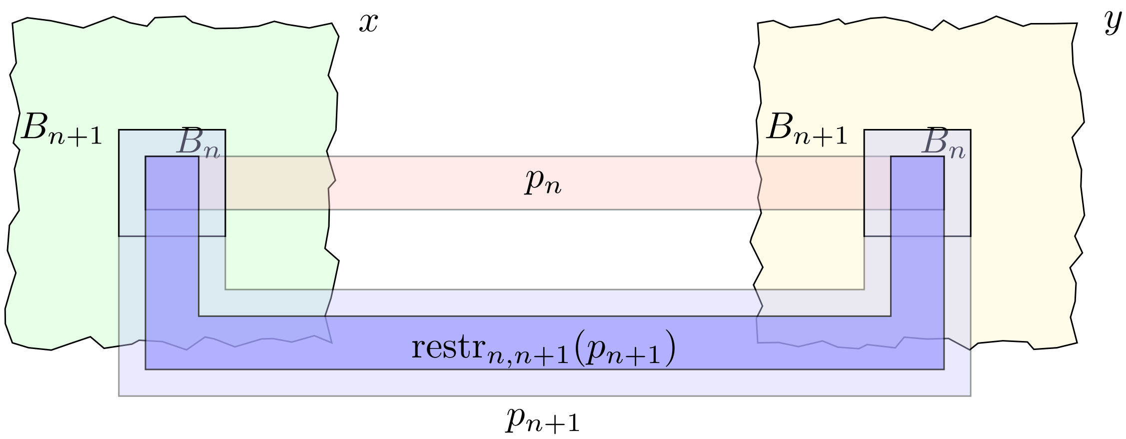

Figure 3.4: An illustration of two paths between the central patterns of two configurations x and y, along with a restriction map.

In other words, an element of $\pi_1^{proj}(X)$ is a sequence of paths (p_n) (in the sense of Definition 3.29) such that for any $n > 0$, $\text{restr}_{n+1,n}(p_{n+1}) \sim_{\mathcal{B}_n} p_n$.

It is important to note that in the inverse system that we consider here to define $\pi_1^{proj}(X, x_0)$, as in any inverse system, the maps play a major role, and we cannot know what the inverse limit is simply by independently computing the groups $\pi_1^{\mathcal{B}_n}(x_0, \mathbf{0})$ – an exception being the case where they are all trivial groups, in which case the inverse limit is also trivial.

We now show that the topological definitions of Section 3.2.2 and the more combinatorial point of view presented in Section 3.2.2 coincide. In particular, calling homotopy and deformations the operations defined in Definition 3.33 or Definition 3.34 is not an abuse of terminology:

Proposition 3.36

For a subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$, $n \geq 1$, and $x_0 \in X, \mathbf{v} \in \mathbb{Z}^d$,

$$\pi_1^{\mathcal{B}_n(\mathbb{Z}^d)}(X, (x_0|_{\mathcal{B}_n(\mathbb{Z}^d)}, \mathbf{v})) \simeq \pi_1(S_{\mathcal{B}_{\mathbb{R}^d}}, (\tilde{x_0}|_{\mathcal{B}_n(\mathbb{R}^d)}, \mathbf{v}))$$

Proof. Fix $n > 0$. For convenience, we prove the case $\mathbb{Z}^2$, although the proof is identical for other dimensions. The idea is to show that the loop classes are the same for both definitions Definition 3.10 and Definition 3.34 – in fact, this holds whether or not the path is a loop. Notice that given a path p in the scene space $S_B(X) = (\tilde{X} \times \mathbb{R}^2)/\sim_{\mathcal{B}_n(\mathbb{R}^2)}$, we can always homotopically deform it into p' so that the trajectory of p' only uses points of $\mathbb{R}\times\mathbb{Z}\cup\mathbb{Z}\times\mathbb{R}$, that is, it consists of straight segments of length 1 between the points of $\mathbb{Z}^2$. We can furthermore ensure that the trajectory is never constant on any non-trivial interval. We can then associate with this path p' an equivalent path (defined as in Definition 3.29) easily, by simply considering the sequence of timesteps $0 \leq t_1 < t_2 < \cdots < t_N \leq 1$ such that $\{t_i, 0 \leq i \leq N\} = \{t \in [0,1], p'(t) \in (\cdot, \mathbb{Z})^2\}$, and the path $(p'(t_i))_{0\leq i\leq N}$, with each $p(t_i)$ satisfying $p(t_i) \in \mathcal{L}_n(X) \times \mathbb{Z}^2$. The only subtlety is that N is indeed a finite number, and that we cannot have infinitely many timesteps t at which $p'(t) \in (\cdot, \mathbb{Z}^2)$. This is simply because if p' is continuous, then the trajectory must also be continuous. □

3.2.3 First examples and properties

In order to compute the projective fundamental groups of some subshifts, one more lemma from [GP95] is useful. Of course, Proposition 3.27 tells us that we can choose any suitable sequence of supports making computations easier, as long as they eventually cover $\mathbb{Z}^d$, but we can make some additional assumptions. The definitions are given for the case $d = 2$ but easily generalize to higher dimensions:

Definition 3.37: Straight path

Let $n > 0$. We say that a path $(P_t, \mathbf{v}_t)_{0\leq t\leq N}$ is n-**straight** if for $0 \leq t \leq N$, $\mathbf{v_t} \in (\mathbb{Z} \times n\mathbb{Z}) \cup (n\mathbb{Z} \times \mathbb{Z})$, and $\mathbf{v_t} \neq \mathbf{v_{t+2}}$ for $t \leq N - 2$, and moreover $(\mathbf{v_0}), \mathbf{v_t} \in (n\mathbb{Z})^2$.

In other words, an n-straight path moves between points of the $n\mathbb{Z} \times n\mathbb{Z}$ sublattice of $\mathbb{Z}^2$, the trajectory being a straight non-backtracking path between two closest such points.

Lemma 3.38: Straightening lemma [GP95]

Let X be a $\mathbb{Z}^2$ subshift and $([p_n])_{n>0}$ a projective path class between $(x, \mathbf{v_x})$ and $(y, \mathbf{v_y})$, and let $m > 0$ be such that $\mathbf{v_x}, \mathbf{v_y} \in (m\mathbb{Z}^2)$. Let $n > 0$. Then $[p_n]$ contains an m-straight path.

Proof. The idea is to use a large aperture window M compared to m. Then, each step of the corresponding path will p_M contains several points of the $(n\mathbb{Z})^2$ sublattice. We can then use those points to determine another homotopic path, and use the fact that $n|_M(p_M) \sim p_n$.

Consider the path $p_{n+m+1} = (P_t, \mathbf{v_t})_{0\leq t\leq N}$. For each step $0 \leq t \leq N$, let $\mathbf{u}_t$ be a point in $(n\mathbb{Z})^2$ which minimizes $\{\|\mathbf{v_t} - \mathbf{w}\|_\infty \mid \mathbf{w} \in (n\mathbb{Z})^2\}$. As P_t is a pattern of support $\mathcal{B}_{n+m+1}$, we have that $\mathbf{u_t} + \mathcal{B}_n \subseteq \mathrm{dom}(P_t) \cap \mathrm{dom}(P_{t+1})$. We can then completely define a path p'_n for the aperture window B_n using those new points $(\mathbf{u}_t)$: let p'_n be the path whose trajectory consists of straight paths between $\mathbf{u_t}$ and $\mathbf{u_{t+1}}$ for any $t < N$, its patterns being the corresponding subpatterns of P_t. Then by definition of p'_n we have $p'_n \sim n|_{n+m+1}(p_{n+m+1})$, and as $([p_k]_{n\in\mathbb{N}})$ is a projective path-class, $n|_{n+m+1}(p_{n+m+1}) \sim_{\mathcal{B}_n} p_n$. We can always assume that p'_n is not backtracking: indeed, if for some t we have $\mathbf{u_t} = \mathbf{u_{t+2}}$ then the corresponding portion of the path can be entirely traced within P_{t+1}, and in particular, can be homotopically contracted to the trivial path. Therefore, p'_n is m-straight, and is homotopic to p_n. □

Note that this lemma only applies when we consider (classes of) paths which we know appear in *projective* path classes: in particular, we do not say anything about an arbitrary loop-class in $\pi_1^{\mathcal{B}_n(\mathbb{Z}^2)}(X, (x_0|_{\mathcal{B}_n}, \mathbf{v}))$, which might not admit any preimage by $\mathrm{restr}_{n,n+1}$ in $\pi_1^{\mathcal{B}_{n+1}(\mathbb{Z}^2)}(X, (x_0|_{\mathcal{B}_{n+1}}, \mathbf{v}))$.

Example 12. *We will compute the projective fundamental group of an SFT extension of the $\mathbb{Z}^2$ sunny-side up subshift,*

$$X = \{x \in \{0,1\}^{\mathbb{Z}^2}, |x|_1 \leq 1\}$$

The example of the sunny-side up X itself is implicitly considered in [GP95, Section 2]. However, an SFT extension of X is more complicated, as it is not a covering (in the *sense of Definition 3.18) or even a constant-to-one extension. The specific extension we consider here is the one illustrated in Figure 3.5.*

$T = \{$  $\}$

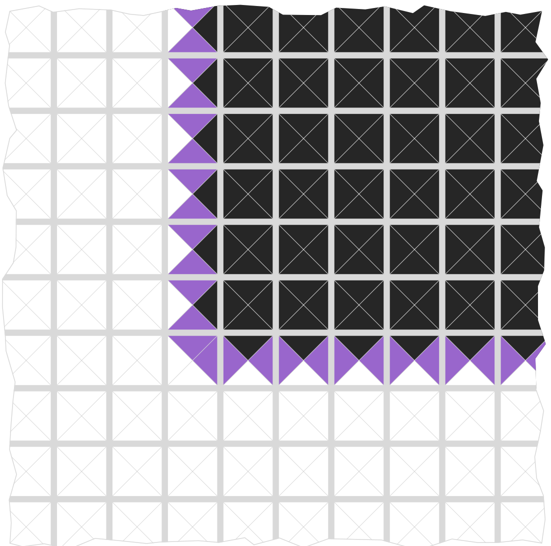

Figure 3.5: An SFT extension of the sunny-side up subshift, defined by a Wang tileset. The projection map sends  to 1 and all the other tiles to 0.

We do not show projective connectedness for now, and defer to Section 3.3 for a general argument showing that this kind of subshift is necessarily projectively connected. In any case, we prove now that the projective path component of $(\square^{\mathbb{Z}^2}, (0,0))$ *has a trivial projective fundamental group, and proving projective connectedness will then ensure that this component is in fact the only one.*

We show that for any $n > 1$*, and any projective loop class* $([p_n])_{n>0}$ *based at* $(\square^{\mathcal{B}_n}, (0,0))$*,* $[p_n]$ *is trivial.*

Let us call the tile  *the corner tile. We write* $x^{(i,j)}$ *for the unique configuration of* X *such that* $x^{(i,j)}_{(i,j)} =$ . *We first claim that any loop* $p = (P_t, \mathbf{v}_t)_{0 \le t \le N}$ *based at* $(\square^{\mathcal{B}_n}, (0,0))$ *can be homotopically deformed so as not to contain any pattern containing the corner tile. Indeed, let* $t_1 = \min_t\{$ $\sqsubseteq P_t\}$ *and* $t_2 = \min_{t>t_1}\{$ $\not\sqsubseteq P_t\}$*. Then, there must exist some* $(i,j) \in \mathbb{Z}^2$ *such that for all* $t_1 \le t \le t_2$*, we have* $P_{t,(i,j)} =$ , *and therefore the path* $(P_t, \mathbf{v}_t)_{t_1 - 1 \le t \le t_2}$ *must be traced in a single configuration* $x^{(i,j)}$*. In particular, we can perform an elementary deformation of* p *into any other path* $p \sim_{\mathcal{B}_n} (P_t, \mathbf{v}_t)_{t<t_1} * p' * (P_t, \mathbf{v}_t)_{t \ge t_2}$*, where* p' *is traced in* $x^{(i,j)}$ *and avoids* . *Repeating this process, we can remove all occurrences of*  *from* p*.*

Let us assume then that p *does not contain* . *Note that for geometric reasons, as our aperture window is a square* $\mathcal{B}_n$*, a pattern which does not contain*  *cannot contain both tiles*  *and* . *We can therefore split the path* p *into* $p = p_1 * p_2 * \ldots * p_k$ *where each* p_k *is of the following form:*

- *Its endpoints are* $\blacksquare^{\mathcal{B}_n}$ *or* $\square^{\mathcal{B}_n}$*. They are equal if and only if* p_k *can be traced in* $\square^{\mathcal{B}_n}$ *or* $\blacksquare^{\mathcal{B}_n}$*.*

- *It does not contain any*  *or it does not contain any* .

- *For each i, writing $p_i = (P_t, \mathbf{v_t})_{s_i \le t_i \le N_i}$, the set of timesteps $\{t \in [\![s_i, N_i]\!], P_t$ contains or $\}$ is a single interval. In particular, all the tiles in such a path are on the same row $\mathbb{Z} \times \{j\}$, and all the tiles are on the same column $\{i\} \times \mathbb{Z}$.*

Now, the idea is to show how to deform each such path p in order to obtain a path that can be traced in a single configuration (in fact, in $\square^{\mathbb{Z}^2}$). If p itself is coherent then it is homotopic to the trivial path, so suppose that it is not, so that $k \ge 2$. Consider the path $q = p_1 * p_2$. We show how to deform it to a coherent path; we can then repeat the process and eventually obtain a contractible path. Without loss of generality, we can assume that p_1 starts in $\square^{\mathcal{B}_n}$ and ends in $\blacksquare^{\mathcal{B}_n}$ after having seen tiles , at some height $j \in \mathbb{Z}$:

- *If p_2 is such that appears in some of its patterns, then let i be the column at which they appear. We can then deform $p_1 * p_2$ homotopically as follows: we concatenate a path r from the endpoint of p_1 to $(\blacksquare^{\mathcal{B}_n}, (i+n, j+n))$ traced entirely in $\blacksquare^{\mathbb{Z}^2}$, and concatenate r^{-1} before p_2. Let $t_1 = \min\{t \le N_1, \sqsubseteq P_t\}$ and $t_2 = \max\{N_1 + 1 \le t \le N_2, \sqsubseteq P_t\}$. Then:*
  - *The paths $(P_t, \mathbf{v_t})_{0 \le t < t_1}$ and $(P_t, \mathbf{v_t})_{t > t_2}$ can be traced in $\square^{\mathbb{Z}^2}$.*
  - *The path $(P_t, \mathbf{v_t})_{t_1 - 1 \le t \le t_2 + 1}$ can be traced in $x^{(i,j)}$, but its starting and ending patterns are both $\square^{\mathcal{B}_n}$. We can therefore deform it so that it doesn't see any tile besides $\square$, and so we it can be traced in $\square^{\mathbb{Z}^2}$ after this deformation.*

  Finally, we get a path that can be traced in $\square^{\mathbb{Z}^2}$.

- *Otherwise, p_2 also contains a tile , let j' be the height at which they all appear in p_2. Let $(o_x, o_y) = \mathbf{v_{N_1}}$ be the final position of p_1. We can homotopically deform $p_1 * p_2$ as above, by inserting a path $r * r^{-1}$ between p_1 and p_2. Let r be the straight path traced in $x^{(i,j)}$ for some sufficiently small i between $(\blacksquare^{\mathcal{B}_n}, \mathbf{N_1})$ and $(\square^{\mathcal{B}_n}, (i-n, o_y))$. In particular, r and r_1 do not contain any or . Hence, the paths $p_1 * r$ on the one hand, $r^{-1} * p_2$ on the other hand, are as in the first case, and so can independently be deformed to be traced in $\square^{\mathbb{Z}^2}$. Their concatenation can therefore also be deformed to be traced in $\square^{\mathbb{Z}^2}$.*

Repeating this argument $k-1$ times, we finally obtain that p is contractible, and so $\pi_1^{proj}(X) = \{e\}$.

3.3 Projective connectedness

Just as in the case of the usual fundamental group for a topological space X, where path-connectedness implies that the group $\pi_1(X, x_0)$ does not depend on the basepoint $x_0 \in X$, we can define an analogous notion for the projective fundamental group. Naturally called *projective* path connectedness, we can think of it as a kind of mixing property, as already observed in [GP95]. In this section, we try to give a few conditions ensuring that a subshift is projectively connected, to highlight the differences between $\mathbb{Z}$ and $\mathbb{Z}^d$ subshifts, and relate projective connectedness with other more classical mixing notions.

We leave open an important question, about which we will say a few words later in this section. First, define a natural decision problem:

Decision Problem — PROJECTIVE-CONNECTEDNESS

Input: An effective subshift $X \subseteq \mathcal{A}^{\mathbb{Z}^d}$.
Output: Whether X is projectively connected.

The main question not answered here is the following:

Question 1. What is the computational complexity of PROJECTIVE-CONNECTEDNESS, and is it in the arithmetical hierarchy ?

The question is still open even if we restrict it to the case of subshifts of finite type. We will briefly investigate this question in Section 3.3.4, and highlight a few reasons why the answer is still unknown to us.

3.3.1 Definition and basepoint independence

The definition is pretty straightfoward:

Definition 3.39: Projective connectedness

Let X be a $\mathbb{Z}^d$ subshift. We say that X is **projectively connected** if for any $x, y \in X$, there exists a projective path class between $(x, \mathbf{0})$ and $(y, \mathbf{0})$, *i.e.*

$$\exists (p_n)_{n>0}, \forall n > 0, \begin{cases} p_n \text{ is a path in } S_{\mathcal{B}_n} \text{ between } (x|_{\mathcal{B}_n}, \mathbf{0}) \text{ and } (y|_{\mathcal{B}_n}, \mathbf{0}) \\ \mathrm{restr}_{n,n+1}(p_{n+1}) \sim_{\mathcal{B}_n} p_n \end{cases}$$

It is then easy to see that the following proposition holds, using Proposition 3.17 for each space $\mathcal{B}_n$:

Proposition 3.40

Let X be a projectively connected $\mathbb{Z}^d$ subshift. Then, for any $x, y \in X$, we have $\pi_1^{proj}(X, x) \simeq \pi_1^{proj}(X, y)$. In this case, we simply write $\pi_1^{proj}(X)$ for the projective fundamental group of X (up to isomorphism).

Furthermore, this notion behaves as expected with respect to factor maps:

Lemma 3.41: Factor Lemma [GP95]

If $\phi: X \to Y$ is a factor map between $\mathbb{Z}^d$ subshifts, and if X is projectively connected then so is Y.

Proof. Let $r = \mathrm{radius}(\phi)$. Then, for any $n \geq 0$, we can naturally define a map $\phi_n: S_{\mathcal{B}_{n+r}}(X) \to S_{\mathcal{B}_n}(Y)$, by sending $(x|_{\mathcal{B}_{n+r}+\mathbf{v}}, \mathbf{v})$ to $(\phi(x)|_{\mathcal{B}_n+\mathbf{v}}, \mathbf{v})$. As $\phi: X \to Y$ is a factor map, it is in particular surjective, so ϕ_n is surjective. Moreover, it is easy to check that for any n, ϕ_n commutes with the canonical restriction maps, in the sense that $\mathrm{restr}_{n,n+1} \circ \phi_{n+1} = \phi_n \circ \mathrm{restr}_{n+r,n+r+1}$, where the restriction maps in the previous equality are respectively maps $S_{\mathcal{B}_{n+1}}(Y) \to S_{\mathcal{B}_n}(Y)$ and $S_{\mathcal{B}_{n+r+1}}(X) \to S_{\mathcal{B}_{n+r}}$. Therefore, ϕ induces a surjective map $\Phi: \varprojlim_{n\in\mathbb{N}} S_{\mathcal{B}_n}(X) \to \varprojlim_{n\in\mathbb{N}} S_{\mathcal{B}_n}(Y)$ between the respective projective limits, and in particular sends projective path classes to projective path classes, so Y is projectively connected. $\square$

Using similar ideas, we can now prove the important Theorem 3.28:

Proof of Theorem 3.28. We keep the notations of the proof of Lemma 3.41. Let $r = \text{radius}(\phi)$ and $r' = \text{radius}(\phi^{-1})$. We still have for all $n \geq 0$ that $\text{restr}_{n,n+1} \circ \phi_{n+1} = \phi_n \circ \text{restr}_{n+r,n+r+1}$. Moreover, it is easy to check that the induced map $\pi_1(S_{\mathcal{B}_{n+r}}(X), x_0)$ and $\pi_1(S_{\mathcal{B}_n}(Y), y_0)$ where $y_0 = \phi(x_0)$. In fact, for all $n \geq m \geq k \geq r'$, we have morphisms $\pi_1(S_{\mathcal{B}_{n+r}}(X), x_0) \to \pi_1(S_{\mathcal{B}_m}(Y), y_0) \to \pi_1(S_{\mathcal{B}_{k-r'}}(X), x_0)$. The idea is now to consider a new, single inverse system, indexed by $\mathcal{I} = \mathbb{N} \times \{X, Y\}$, with a partial order

$$(n, b) \geq (m, b') \iff \begin{cases} b = b', n \geq m \text{ or} \\ b = X, b' = Y, n \geq m + r \text{ or} \\ b = Y, b' = X, m \geq n + r' \end{cases}$$

with the restriction maps being either the canonical restriction maps $\text{restr}_{m,n}$, or the maps induced by ϕ, ϕ' described above. As in Proposition 3.27, the inverse limit of this system is isomorphic to the inverse limit of either of $\pi_1(S_{\mathcal{B}_n}(X), x_0))_{n\in\mathbb{N}}, \pi_1(S_{\mathcal{B}_m}(X), y_0))_{m\in\mathbb{N}}$, as any element of either subsystem is eventually dominated by an element of the order for the partial order described above. ☐

We immediately obtain the following result:

Proposition 3.42

If X is a projectively connected subshift and Y is conjugate to X, then $\pi_1^{proj}(X) \simeq \pi_1^{proj}(Y)$.

Proof. This is simply Theorem 3.28 and Lemma 3.41. ☐

An important remark, which partially explains the difficulty of deciding whether a subshift is projectively connected or not, is the fact that the definition of projective connectedness (Definition 3.39) requires to construct a *sequence* of "consistent" paths (in the sense that they are homotopic after applying the canonical restriction maps). It is *a priori* not enough to ensure that for all $n > 0$, there exists a path p_n between $(x|_{\mathcal{B}_n}, \mathbf{0})$ and $(y|_{\mathcal{B}_n}, \mathbf{0})$. We will say a few words about this problem in Section 3.3.4.

3.3.2 Projective connectedness as a mixing property

There are some easy examples of subshifts which are not projectively connected. The simplest such example is the following:

Proposition 3.43

If $X \subset \mathcal{A}^{\mathbb{Z}^2}$ is a periodic subshift, then it is not projectively connected, unless $X = \{a^{\mathbb{Z}^2}\}$ for some $a \in \mathcal{A}$.

Proof. Suppose that X is not reduced to a single configuration, let n be larger than the period of X, and let $P \in \mathcal{L}_n(X)$ a pattern that is not $(1,0)$-periodic – this is without loss of generality, up to rotating X by a quarter-turn. Then, there is a single configuration $x \in X$ such that $x|_{\mathcal{B}_n} = P$. In particular, any path in $S_{\mathcal{B}_n}$ starting from P must be traced in x, and so $(P, (0,0))$ cannot be linked by any path to $(P, (1,0))$. ☐

We get the immediate corollary:

Corollary 3.44

Let X be a $\mathbb{Z}^2$ subshift. If X admits a non-trivial periodic factor, then it is not projectively connected.

Proof. This is Proposition 3.43 and Lemma 3.41. □

Projective connectedness is a strong mixing notion, but quite hard to fully characterize. We give in this section an overview of how it might relate to other mixing notions in higher-dimensional subshifts. We first give a condition which ensures that a subshift is projectively connected.

Cones and cone-connected subshifts

We will see in later sections that most of the usual mixing properties do not imply projective connectedness. We give here a new kind of mixing property, inspired by [Sch95], based on **cones**. This property will be used in particular to prove the projective connectedness of a large class of subshifts in Section 3.4.

Let us start with a weaker result, which will nonetheless be useful to understand the later ideas of what we call the cone-connected property:

Definition 3.45: Strong irreducibility

Let X be a $\mathbb{Z}^2$ subshift, and d the distance induced by the infinite norm $\|\|_\infty$ on $\mathbb{Z}^2$. We say that X is **strongly irreducible** if there exists $N \in \mathbb{N}$ such that for any $x, y \in X$ and $A, B \subset \mathbb{Z}^2$ (not necessarily finite) with $d(A, B) \geq N$, there exists $z \in X$ such that $z|_A = x|_A, z|_B = y|_B$.

Proposition 3.46

If X is strongly irreducible then it is projectively connected.

Proof. Let $x, y \in X$. We construct explicitly a projective path-class between $(x, (0,0))$ and $(y, (0,0))$. By strong irreducibility of X, there exists $N > 0$ and $z \in X$ such that:

- $z|_{-\mathbb{N}\times\mathbb{Z}} = x|_{-\mathbb{N}\times\mathbb{Z}}$.
- $z|_{(N+\mathbb{N})\times\mathbb{Z}} = y|_{(N+\mathbb{N})\times\mathbb{Z}}$.

For $n > 0$, we define a path p_n between in $(x|_{\mathcal{B}_n}, (0,0))$ and $(y|_{\mathcal{B}_n}, (0,0))$ as follows:

- Starting from $(x|_{\mathcal{B}_n}, (0,0))$, move to the left in the configuration x to $(-n-1, 0)$, and let p_n^1 be this path.
- As $x|_{\mathcal{B}_n-(n+1,0)} = z|_{\mathcal{B}_n-(n+1,0)}$, we can continue this path in z: move straight to the right in the configuration z, until the point $(N+n, 0)$. We call this path p_n^2?
- As $y|_{\mathcal{B}_n+(N+n,0)} = z|_{\mathcal{B}_n+(N+n,0)}$, we can continue this path in y: move straight to the left in the configuration z, until the point $(0,0)$. We call this path p_n^3

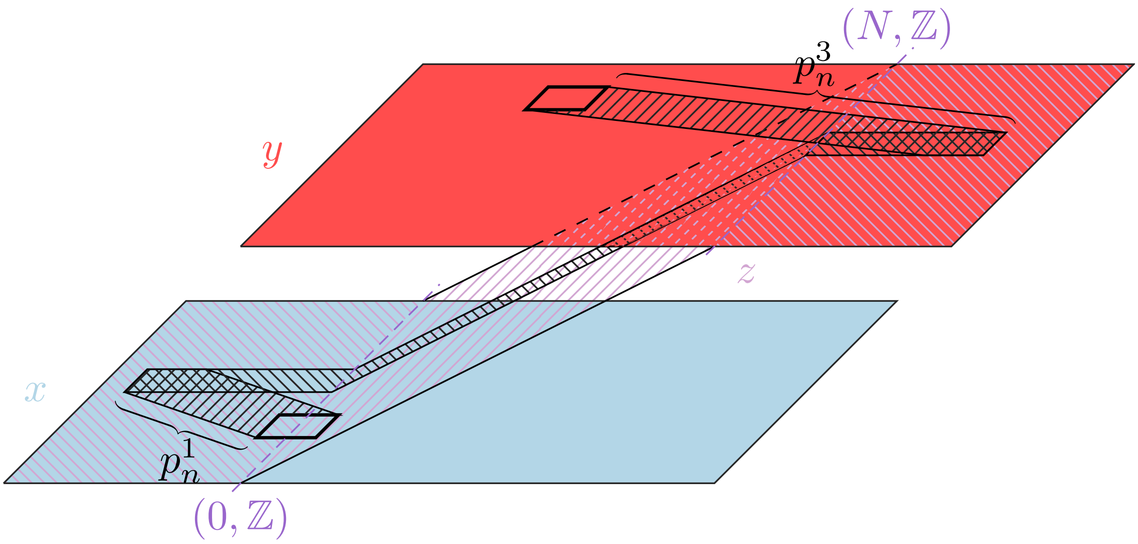

Figure 3.6: Path p_n. For visualization purposes, the trajectory does not remain on the $\{\mathbb{Z}\} \times \{0\}$ row. The configurations x and y are respectively represented in blue and red, one above the other; the configuration z has overlaps with both, and is distinct on the $[\![0, N]\!] \times \mathbb{Z}$ stride.

Finally, define $p_n = p_n^1 * p_n^2 * p_n^3$.

This path is illustrated in Figure 3.6.

It is now easy to see that for all $n > 0$, $\text{restr}_{n,n+1}(p_{n+1}) \sim_{\mathcal{B}_n} p_n$. Indeed, notice that $\text{restr}_{n,n+1}(p_{n+1}^1)$ can be obtained by adding a single step at the end to p_n, moving to $(-n-2, 0)$ in x, and $\text{restr}_{n,n+1}(p_{n+1}^1)$ can be obtained by prepending an extra step to p_n^2, moving from $(-n-2, 0)$ to $(-n-1, 0)$ in x. As those two additions are inverse from one another, this is a loop, and so can be contracted homotopically. Similarly, one can modify the end of p_n^2 and the start of p_n^3 to achieve the same deformation on the second half of the path, and so $\text{restr}_{n,n+1}(p_{n+1}) \sim_{\mathcal{B}_n} p_n$.

Hence, $(p_n)_{n>0}$ is a projective path-class between x and y. □

This argument does not actually require us to be able to "glue" arbitrary parts of configurations of X, as in the strong irreducibility property. In fact, it suffices to be able to glue **cones**:

Definition 3.47: Cone

Let $\alpha \in \mathbb{R}, \mathbf{u} \in \mathbb{Z}^2$. We define the **cone** $C(\alpha, \mathbf{u})$ as:

$$C(\alpha, \mathbf{u}) = \{\mathbf{v} \in \mathbb{Z}^2, \langle \mathbf{u}, \mathbf{v} \rangle \geq \alpha \|\mathbf{u}\|_2 \|\mathbf{v}\|_2\}$$

where $\langle \cdot, \cdot \rangle$ is the standard euclidean scalar product.

Definition 3.48: Cone-connected

Let X be a $\mathbb{Z}^2$ subshift. We say that X is **cone-connected** if there exists $x \in X$ such that for any $y \in X$, there exists $\alpha, \beta \in \mathbb{R}, \mathbf{u}, \mathbf{v}, \mathbf{o} \in \mathbb{Z}^2$ and $z \in X$ such that $z|_{C(\alpha,\mathbf{u})} = x|_{C(\alpha,\mathbf{u})}$ and $z|_{C(\beta,\mathbf{v})+\mathbf{o}} = y|_{C(\beta,\mathbf{v})+\mathbf{o}}$.

This is a rather strong mixing condition, but we still have some margin: we do not ask for *any* cones $C(\alpha, \mathbf{u})$ and $C(\beta, \mathbf{v})$ to be able to be glued in a common configuration, and furthermore, we only require the cones of a single specific configuration x to be glued to cones of other configurations of X. Nevertheless, this is a sufficient condition to ensure projective connectedness:

Proposition 3.49

Let X be a cone-connected $\mathbb{Z}^2$ subshift. Then X is projectively connected.

Proof. The proof is the same as in Proposition 3.46. We show how to define a projective path class between x and any configuration y. Let z be as in Definition 3.48, and define for $n > 0$ a path p_n between $(x|_{\mathcal{B}_n}, (0,0))$ and $(y|_{\mathcal{B}_n}, (0,0))$ as follows:

- Start by moving in x from $(0,0)$ to a point $\mathbf{w}_n \in C(\alpha, \mathbf{u})$ such that $\mathbf{w}_n + \mathcal{B}_n \subset C(\alpha, \mathbf{u})$ – such a point $\mathbf{w}_n$ exists, as $C(\alpha, \mathbf{u})$ is a cone.
- Now, move in z to a point $\mathbf{w'}_n$ so that $\mathbf{w'}_n + \mathcal{B}_n \subset C(\beta, \mathbf{v})$.
- Come back in y.

It is then easy to see that $(p_n)_{n>0}$ is a projective path class. Indeed, $\mathrm{restr}_{n,n+1}(p_{n+1})$ is homotopic to p_n, as it suffices to remove the paths in z from w_n to w_{n+1} and back, and from w'_n to w'_{n+1} and back. □

In particular, for subshifts having any kind of "safe symbol" – *i.e.* a symbol $*$ such that for any $x \in X$, changing $x_{\mathbf{u}} = *$ at any $\mathbf{u} \in \mathbb{Z}^2$ produces a configuration that is still valid in X – even with restriction such as the ones considered in [Jen01], we can show that they are projectively connected using cone-connectedness, even if they might not be strongly irreducible due to some non-isotropic properties of the subshift.

Chain-mixing properties

We first recall a few classical definitions from the theory of dynamical systems. The definitions are not specific to subshifts, but are given in the case of a topological space (X, T) with $T\colon X \to X$. In particular, for higher-dimensional subshifts, we need to consider slightly different versions of the usual definitions. A longer introduction to these properties and their relative implications can be found in [Kur03, Section 2.1].

Definition 3.50: Chain

Let (X, T) be a dynamical system, and let d be a distance on X. For $x, y \in X$ and $\varepsilon > 0$, an ε-**chain** between x and y is a sequence $(x_0 = x, x_1, \ldots, x_N = y)$ of elements of X such that for $0 \leq i < N$, $d(T(x_i), x_{i+1}) < \varepsilon$.

In other words, an ε-chain can be viewed as almost being an orbit, where at each new application of the map T, we only know the image up to a precision ε. This easily generalizes to the case of $\mathbb{Z}^d$-actions.

Definition 3.51: Chain transitive

A dynamical system (X, T) is **chain-transitive** if for all $x, y \in X$ and $\varepsilon > 0$, there exists an ε-chain between x and y.

In the case of subshifts, the definition of chains is very similar to the notion of path defined in Definition 3.29, and the next result is not surprising:

Proposition 3.52

Let X be a projectively connected SFT. Then X is chain transitive.

Proof. Let $\varepsilon > 0$, and n such that $2^{-n} < \frac{\varepsilon}{2}$. Let $x, y \in X$. We show that there exists an ε-chain between x and y.

X is projectively connected, so there is a path $p_n = (P_k, \mathbf{v}_k)_{0 \le k \le r}$ between $x_{|\mathcal{B}_n}$ and $y_{|\mathcal{B}_n}$, and each P_k is a pattern of support $\mathcal{B}_n$ – see Figure 3.7 for an illustration. By definition of a path, each P_k is globally admissible, so there exists $z^k \in X$ so that $z^k_{|\mathcal{B}_n} = P_k$. For $k < r$, let $e_k = \mathbf{v}_{k+1} - \mathbf{v}_k$. By definition,

$$\sigma_{e_k}(z^k_{|\mathcal{B}_{n-1}}) = z^{k+1}_{|\mathcal{B}_{n-1}}$$

i.e. $d(\sigma_{e_k}(z^k), z^{k+1}) < 2^{-(n-1)} \le \varepsilon$ and so (z^k) is an ε-chain between x and y in X.

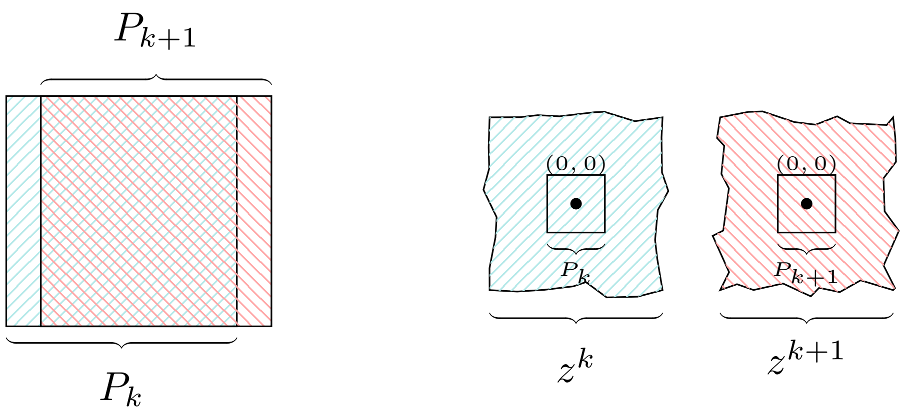

Figure 3.7: Illustration of the ε-chain between x and y

□

As already observed by [GP95] (which prove it for mixing subshifts, but the exact same proof gives the result for a weakly mixing subshift (see Definition 1.42), already present in [Piv07, Prop. 3.3]), we have the following proposition:

Proposition 3.53

Let X be a weakly mixing subshift. Then for any $n \in \mathbb{N}$, $S_{\mathcal{B}_n}$ is path-connected.

Proof. Let $x, y \in X$, and let $P \in \mathcal{L}_n(X)$ be any globally admissible pattern of X. Let $P \in \mathcal{L}_n(X)$ be an arbitrary pattern of X. As X is weakly mixing, there exists $\mathbf{u} \in \mathbb{Z}^2$ and $z_x, z_y \in X$ such that:

- $z_x|_{\mathcal{B}_n} = x|_{\mathcal{B}_n}$, $z_y|_{\mathcal{B}_n} = y|_{\mathcal{B}_n}$
- $z_x|_{\mathbf{u}+\mathcal{B}_n} = z_y|_{\mathbf{u}+\mathcal{B}_n} = P$.

We can then find a path in $S_{\mathcal{B}_n}$ between $(x|_{\mathcal{B}_n}, (0,0))$ and $(y|_{\mathcal{B}_n}, (0,0))$ by considering any path traced in z_x with a trajectory from $(0,0)$ to $\mathbf{u}$, concatenated with a path from $(P, \mathbf{u})$ to $(y|_{\mathcal{B}_n}, (0,0))$ traced in z_y. □

Transitivity

On the other hand, there are examples of SFT that are projectively connected but not mixing – and in fact, not even transitive.

Proposition 3.54

For any $d \geq 1$, there exists a non-transitive projectively connected $\mathbb{Z}^d$ subshift

Proof. We construct an example for $d = 1$, and the same example works for arbitrary dimension by considering the lifts $X^{\uparrow}, X^{\uparrow^{\uparrow}} \dots$ and so on.

Consider the one-dimensional SFT X on the alphabet $\{0, 1, 1', 2'\}$ defined by the following adjacency graph:

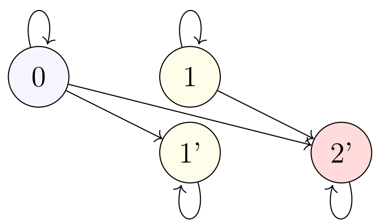

Figure 3.8: Graph giving the horizontal adjacency rules for X. An edge (u, v) indicates that the colour v can be placed to the right of u.

This SFT is obviously non-transitive, as the patterns $01'$ and $02'$ can never belong to the same configuration, but it is projectively connected. Indeed, we show how to make a projective path from any configuration to $x_0 = 0^{Z^2}$. This implies projective connectedness.

Let $x \in X$.

- If x contains a 0, then we define a path for any aperture window by simply moving to the left far enough in x so as to only see 0, and coming back in the x_0 configuration. This obviously defines an inverse system of paths.
- Otherwise, let n be such that $[\![-n, n]\!]^2$ is an aperture window large enough to see all the colours of x. For $m > n$, define a path p_n for this window as follows. If x contains at least two colours, or is constantly some i':
  - Start by moving $3n$ steps to the right in x
  - Move left for $6n$ steps in the configuration containing 0 on the negative columns, and the rightmost colour of x on the positive columns.
  - Move right for $3n$ steps in x_0.

  If x is constantly 1, then:

– Start by moving $3n$ steps to the left.
– Move right for $6n$ steps in the configuration containing 1 on the negative columns, $2'$ on the positive columns.
– Move left for $6n$ steps in the configuration containing $2'$ on the positive columns, 0 on the negative ones.
– Move right for $3n$ steps in x_0.

This defines an inverse system of paths, hence X is projectively connected but not even transitive.

Instead of considering $X^{\uparrow}$ as a $\mathbb{Z}^2$ example, we can also perform some minor modifications: the same argument showing projective connectedness works even we allow columns to be non-constant, but *e.g.* if we can shift each row one cell to the left or to the right compared to the one below it. This becomes an uncountable subshift with no isolated points, which is still not transitive but projectively connected. ☐

Contractibility

Denote $I = \{0,1\}^{\mathbb{Z}^2}$ the binary full shift, and $\bar{0} = 0^{\mathbb{Z}^2}, \bar{1} = 1^{\mathbb{Z}^2}$ its two fixed points. Following [PS24], we define a contractibility notion for subshifts:

Definition 3.55: Contractible subshift

A $\mathbb{Z}^2$ subshift X is **contractible** if there exists some block map $h\colon I\times X\times X$ such that for any $x, y \in X$:

- $h(\bar{0}, x, y) = x$
- $h(\bar{1}, x, y) = y$

This definition closely mirrors Definition 3.10: this is in fact an analogous, for subshifts, of the general property of a *space* (rather than a path), namely, having the homotopy type of a point.

In the case where X is an SFT with a fixed point, contractibility implies that we can even take h to be such that $h(\cdot, x, x) = x$ for all x, see the original article [PS24] for more details.

We now try to see the link of this property with the projective fundamental group of the subshift X. Links between contractibility and other classical notions from the literature are explored in more details in the original paper [PS24]:

Lemma 3.56

Any contractible subshift is projectively connected.

Proof. Contractibility implies strong irreducibility. It is even one the motivations behind the introduction of the notion of a contractible subshift, which can be seen a strengthening of strong irreducibility where the configuration containing the required patterns can be obtained by a block map: let r be the radius of the map h given by Definition 3.55, and let P, Q be two patterns of X. Let $x, y \in X$ and $\mathbf{v} \in \mathbb{Z}^2$ be such that:

- $x|_{\mathrm{dom}(P)} = P$.

- $y|_{\mathbf{u}+\mathrm{dom}(Q)} = Q$.
- $d(\mathrm{dom}(P), \mathbf{u} + \mathrm{dom}(Q)) > 2r$.

Then, we claim that there exists $z \in X$ containing P at $(0,0)$ and Q at $\mathbf{u}$, and as r does not depend on P, Q, this implies strong irreducibility. Let $b \in \{0,1\}^{\mathbb{Z}^2}$ be such that $b_{\mathbf{v}} = 0 \iff d(v, \mathrm{dom}(P)) \leq r$. Then, for any $\mathbf{v} \in \mathrm{dom}(P)$, $h(b, x, y)_{\mathbf{v}} = h(\bar{0}, x, y)_{\mathbf{v}} = x_{\mathbf{v}}$, and similarly for $\mathbf{v} \in \mathrm{dom}(Q)$ we have $h(b, x, y)_{\mathbf{v}} = h(\bar{1}, x, y)_{\mathbf{v}} = y_{\mathbf{v}}$. □

However, it is unclear what this implies for the value of the fundamental group. Classically, and by definition, contractible spaces have trivial fundamental group. There exists, however, spaces with trivial π_1 which are not contractible: indeed, contractibility implies that **all** the homotopy groups are trivial – we did not formally define these groups in this chapter, but they are defined using "higher-dimensional equivalents" to paths, *i.e.* continuous maps from higher-dimensional spheres rather than S^1 to the space X. In particular, the 2-sphere is an example of such a space, which has trivial fundamental group but *e.g.* $\pi_2(S^2) = \mathbb{Z}$.

Proposition 3.57

Let X be a $\mathbb{Z}^2$ SFT. If X is contractible and has a fixed point, it has trivial projective fundamental group.

Proof. To show this, we prove that X is in fact cohomologically trivial (in the sense of [Sch98]), which is *a priori* stronger. We do not need to introduce the general definitions here, and will simply state the relevant results that we need in this proof.

We use a condition implying triviality of all the cocycles from [Sch95], namely, the $\mathbf{u}$-specification property. More precisely:

- We prove that X satisfies a *specification property*, introduced in [Sch95, Definition 2.2] that we recall below.
- This property implies that X is cohomologically trivial by [Sch95, Corollary 3.3].
- We do not give precise definition of cohomological triviality, but use another result from Klaus Schmidt, [Sch98, Corollary 5.8] which implies that $\pi_1^{proj}(X) = \{e\}$.

It then suffices to prove that X indeed satisfies the required specification property. Let $\bar{x} = 0^{\mathbb{Z}^2}$ be the fixed point of X, and let $\Delta = \{x \in X \mid |\mathbf{u} \in \mathbb{Z}^2, \bar{x}_{\mathbf{u}} \neq x_{\mathbf{u}}| < +\infty\}$ be the set of configurations differing from $\bar{x}$ in finitely many places. We say that X has the **specification property** if:

- Δ is dense in X: for all $n > 0$ and $P \in \mathcal{L}_n(X)$, there exists $x \in X$ and $m > n$ such that $x|_{\mathcal{B}_n} = P$ and for $\mathbf{u} \in \mathbb{Z}^2 \setminus \mathcal{B}_m$ we have $x_{\mathbf{u}} = 0$.
- There exists a direction $\mathbf{u} \in \mathbb{Z}^2$ and an angle $\alpha > 0$, and parameters $s \geq 1, t \geq 0$ such that for all $x, y \in \Delta$ and $n \geq 0$, if $x|_{sn+t} = y|_{sn+t}$, there exists $z \in \Delta$ such that:
  - $z|_{C(\alpha,\mathbf{u})+\mathcal{B}_n} = x|_{C(\alpha,\mathbf{u})+\mathcal{B}_n}$
  - $z|_{C(\alpha,-\mathbf{u})+\mathcal{B}_n} = y|_{C(\alpha,-\mathbf{u})+\mathcal{B}_n}$

  Said differently, if x, y agree on a sufficiently large ball (of size $sn + t > r$), we can glue a cone of x and a cone of y with no gap between them, see the illustration in Section 3.3.2.

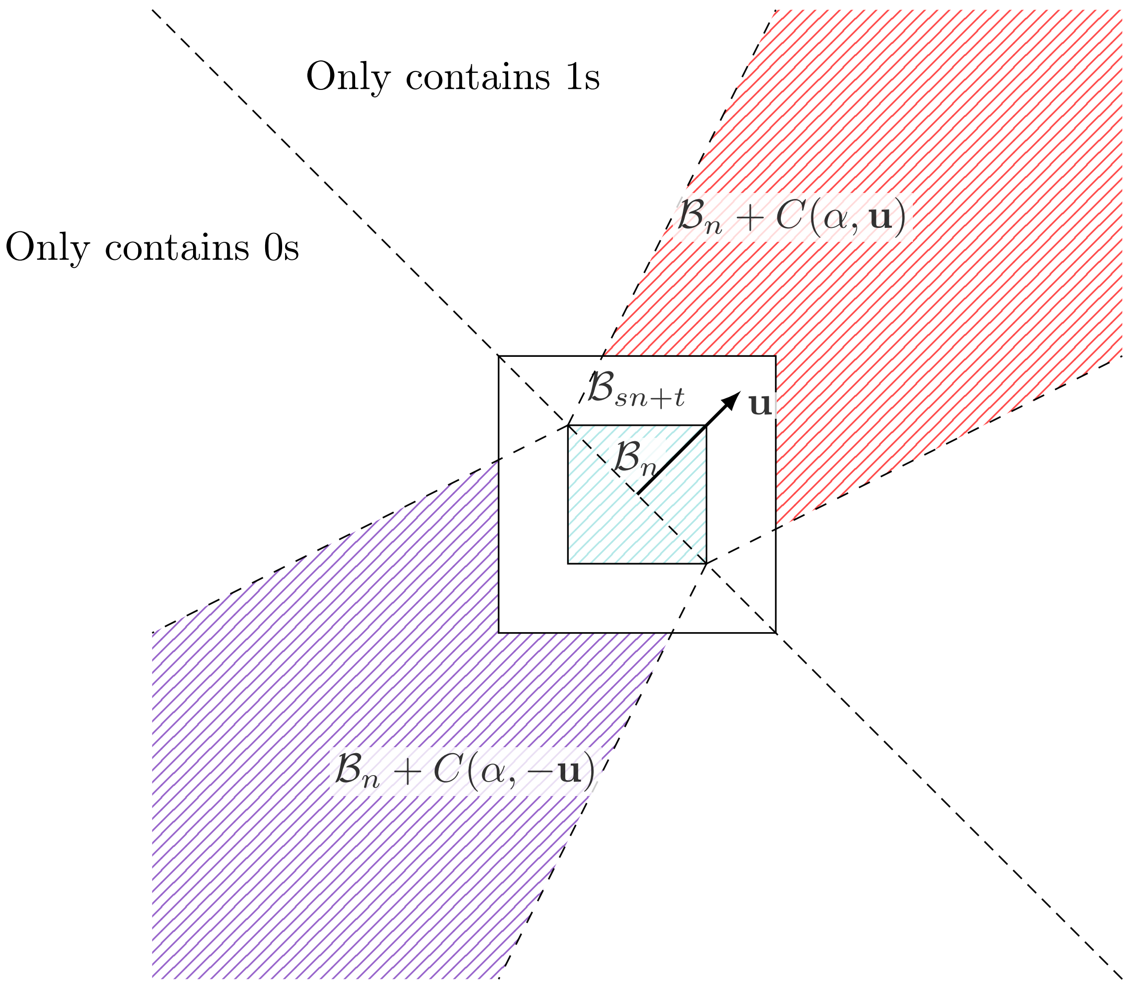

Let h the map given by Definition 3.55 for X, and let $r = \text{radius}(h)$. We prove that X has the specification property with $\mathbf{u} = (1,1), \alpha = \frac{1}{2}, s = 1, t = 2r+1$

By strong irreducibility, Δ is dense. Now, pick $n \geq 0$, and any two points x, y differing on finitely many positions from $\bar{x}$, and equal on positions of $\mathcal{B}_{n+2r+1}$.

We define a configuration b of $I = \{0,1\}^{\mathbb{Z}^2}$ as follows: for $\mathbf{v} \in \mathbb{Z}^2$, $b_{\mathbf{v}} = 0$ if and only if $\langle \mathbf{u}, \mathbf{v} \rangle \geq 0$.

Define $z = h(b, x, y) \in X$. By definition of h:

- For $\mathbf{v} \in \mathcal{B}_{n+r+1}$, $z_{\mathbf{v}} = h(\cdot, x, y)$ depends only on b, x, y restricted to $\mathcal{B}_{n+2r+1}$, and as x, y coincide there, and so $z_{\mathbf{v}} = x_{\mathbf{v}} = y_{\mathbf{v}}$.

- For $\mathbf{v} \in (\mathcal{B}_n + C(\alpha, \mathbf{u})) \setminus \mathcal{B}_{n+r+1}$, $z_{\mathbf{v}} = h(\bar{1}, x, y)_{\mathbf{v}}$ so by definition of h $z_{\mathbf{v}} = x_{\mathbf{v}}$.

- The same arguments prove that $z|_{\mathcal{B}_n + C(\alpha, -\mathbf{u})} = y|_{\mathcal{B}_n + C(\alpha, -\mathbf{u})}$.

□

The reciprocal is unclear:

Question 2. Does having a trivial fundamental group and a fixed point imply contractibility ?

We will see in Section 3.4 what could be a likely counterexample, using Hom-shifts.

3.3.3 One-dimensional SFT

The projective fundamental group of one-dimensional subshifts is not necessarily a particularly interesting object, as the main idea of homotopically deforming paths is rendered trivial here by the fact that a contractible path is essentially a sequence of (possibly nested) backtracking paths. Nevertheless, we say a few words about the projective connectedness of $\mathbb{Z}$-SFTs.

Consider a $\mathbb{Z}$-SFT X over some alphabet $\mathcal{A}$. Without loss of generality, we see X as bi-infinite walks on a graph G with vertices $V(G) = \mathcal{A}$. In this section, we always assume that graphs are connected whenever we consider them as non-directed graphs, that is, when each edge can be traversed in both directions – if G is disconnected even for this notion, X is clearly not projectively connected.

Proposition 3.58

Let X be a $\mathbb{Z}$-SFT defined by some directed graph $G = (V, E)$. Suppose that G has at least two (not necessarily simple or disjoint) cycles of relatively prime lengths, which are in the same strongly connected component. Then, X is projectively connected.

Proof. The proof is a simple generalization of the procedure already highlighted in the proof of Proposition 3.54. Most ideas are already present in [Kur03, Chapter 3.6.1] in which the author studies *attractors* of one-dimensional SFTs.

For a vertex $v \in G$, we denote $\mathrm{scc}(v)$ its strongly connected component in G. Let $C_p = (u_0, \ldots, u_{p-1}), C_q = (v_0, \ldots, v_{q-1})$ be the two cycles of relatively prime length p, q of G. As we assumed that those cycles were in the same strongly connected component of G, there exists a path between u_0 and v_0 of length k_{pq}, and a path between v_i and u_j of length k_{qp}. Let C be the strongly connected component containing C_p, C_q, that is $C = \mathrm{scc}(u_0) = \mathrm{scc}(v_0)$.

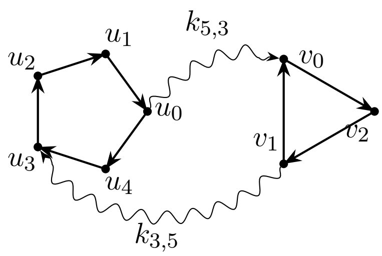

Figure 3.9: An example with $p = 5, q = 3$

Let k be the length of the path $u_0 \to v_0 \to v_i \to u_j \to u_0$. Then, starting from u_0 and taking a times the C_p cycle, b times the C_q cycle, gives a loop of length $k + ap + bq$. Finding a, b so that the result is any sufficiently large integer is always possible, as p, q are coprime: this problem is classically known as the Frobenius problem (and in that particular case, it simply derives from Bézout's identity). In the example of Figure 3.9, we have $p = 5, q = 3, k = k_{5,3} + k_{3,5} + 5$; any number greater or equal $k + 8$ can be obtained with suitable non-negative a, b.

Consider the directed acyclic graph $G' = (V', E')$ whose vertices V' are the non-trivial strongly connected components of G, (*i.e.*, components not reduced to a single vertex which has no self-loops), with an edge $(v, v') \in E' \subseteq V' \times V'$ if and only if there is a directed path between some vertices of the components in the original graph G. We can then define a partial order $\leq_{\mathrm{scc}}$ using G', by defining $v' \leq_{\mathrm{scc}} v$ if there is a (directed) path from v to v' in G'.

Now, define:

$$l\colon X \to V' \qquad\qquad r\colon X \to V'$$
$$x \mapsto \sup_{i\leq 0} \mathrm{scc}(x_i) \qquad\qquad x \mapsto \inf_{i\geq 0} \mathrm{scc}(x_i)$$

By definition of X and of G', and because G is finite, it is clear that for each $x \in X$, the function $i \mapsto \mathrm{scc}(x_i)$ is non-increasing and so *e.g.* $r(x) = \lim_{i\to+\infty} \mathrm{scc}(x_i)$.

Fix $x \in X$, with $l(x) = r(x) = C$, and let $y \in X$ be arbitrary. This is always possible, as C is a non-trivial strongly connected component and so it contains a cycle.

Let then γ be a path in G' between $l(x) = C$ and $l(y)$, when G' is viewed as a non-directed graph. We show by induction on p how to find a projective path between x and y. Without loss of generality, we can assume that $\mathrm{scc}(y_0) = l(y)$, and therefore $\mathrm{scc}(y_i) = l(y)$ for any $i \leq 0$, up to shifting both x and y by the same and sufficiently large amount.

If $l(x) = l(y) = C$ In that case, we use the remarks made at the beginning of the proof. We define the path p_n for the window $\mathcal{B}_n$. Start by moving $2n$ steps to the right in x. It is then sufficient to show that there exists $i < 0$ such that there exists a path γ' of length i in G between x_0 and y_i. But such an i always exist, by Frobenius theorem. We can then trace the second part of the path in the configuration z equal to x on $\mathbb{N}$, to y on $i-\mathbb{N}$, and to the above path between i and 0. Note that i does not depend on n, the width of the current aperture window, nor does it depend on y, but only on G' and C more specifically, and so z does not depend on n. In particular, $\mathrm{restr}_{n,n+1}(p_{n+1}) \sim p_n$.

If γ's last edge in G' is taken with the correct orientation, *i.e.* $\gamma = \gamma' \cdot (w, l(y))$ with $(w, l(y)) \in E'$. Then we show that we can find a projective path between y and some configuration z with $l(z) = w$. By induction hypothesis, we know how to construct a projective path between x and z, and so this would be enough to conclude. By definition of E', there exists a path $\eta = (\eta_{-k}, \eta_{-k+1}, \ldots, \eta_0 = y_0)$ in G between some vertex η_{-k} with $\mathrm{scc}(\eta_{-k}) = w$, and y_0. Let $(\lambda_0 = \eta_{-k}, \lambda_1, \lambda_{m-1})$ be a cycle in $w = \mathrm{scc}(\eta_{-k})$. Define z as:

$$z\colon \mathbb{Z} \to V$$
$$i \mapsto \begin{cases} y_i & \text{if } i > 0 \\ \eta_i & \text{if } i \in [\![-k, 0]\!] \\ \lambda_{(i+k) \mod m} & \text{otherwise} \end{cases}$$

Then we can define the path p_n as follows: starting from $(y|_{\mathcal{B}_n}, 0)$, move to the right in y for $2n$ steps, and come back to 0 by moving to the left for $2n$ steps in z. This defines a projective path-class between y and z.

If γ's last edge in G' is taken with the opposite orientation, *i.e.* $\gamma = \gamma' \cdot (w, l(y))$ with $(l(y), w) \in E'$. This case is completely symmetrical to the previous one. We construct $z \in X$ with $l(z) = w$ in the same way, and there is a projective path between y and z – taking its inverse and using the induction hypothesis, we get a path between x and y.

This shows that any $y \in X$ is connected to some fixed $x \in X$, and so X is projectively connected. $\square$

This shows that bearing some trivial "parity-like" restrictions, of the kind already mentioned in Corollary 3.44, any irreducible SFT is in fact projectively connected. Moreover, we have quantitative bounds on the number of different configurations needed to trace a path – said differently, we can (uniformly) bound for any n the decomposition length of p_n for p_n a path of minimal length between any two points $x, y \in X$ in the scene-space S_{B_n} (more precisely, between any two pairs $(x|_{\mathcal{B}_n}, 0)$ and $(y|_{\mathcal{B}_n}, 0)$).

3.3.4 Deciding projective connectedness

As mentioned in Section 3.3.1, in order to ensure that a subshift is projectively connected, it is *a priori* not enough to ensure the following:

$$\forall n > 0, \exists p_n \text{ a path between } (x\big|_{\mathcal{B}_n}, \mathbf{0}) \text{ and } (y\big|_{\mathcal{B}_n}, \mathbf{0}).$$

In fact, we do not know whether this implies projective connectedness, and we do not have any counter-examples. We first prove the easy fact that the problem of whether a subshift a projectively connected is already undecidable; the above discussion is simply a way to show that we have no clear idea of the actual difficulty of this problem, that is, where it falls in the arithmetical hierarchy (see Section 1.2.2).

Proposition 3.59

The problem PROJECTIVE-CONNECTEDNESS is Σ_1^0-hard.

Proof. The proof is a simple reduction to the domino problem Theorem 1.77. Let Y be the one-point subshift (which is obviously projectively connected), and let Z be any non-projectively connected subshift, such as non-trivial periodic subshift (see Proposition 3.43). Then, for any SFT X, the subshift $X \times Z \sqcup Y$ is projectively connected if and only if X is empty. ☐

This proof is not particularly insightful, and similar easily-shown-to-be-undecidable results can be proven using the general machinery developed in [Car24].

Let us now highlight a few reasons, or counter-intuitive examples, for why we do not know whether it is enough to ensure path connectedness of all the individual scene spaces $S_{\mathcal{B}_n}(X)$ in order to have projective connectedness.

Example 13 (Solenoid)**.** *Consider the* ***solenoid*** X *obtained as the inverse limit of the following system:*

- *For all* $n > 0$*, let* $S_n = S^1 = [0, 1]/(0 \sim 1)$ *be the circle.*
- *For* $m \geq n > 0$*, let* $f_{m,n} \colon x \in S^1 \mapsto 2^{m-n} x \in S^1$*, and so in particular* $f_{n+1,n} \colon x \in S^1 \mapsto 2x \in S^1$*.*

Then the solenoid X *is defined as* $X = \varprojlim_{n>0} (S_n, f_{m,n})_{0<n\leq m}$*. This is a well-known object, initially introduced in [Vie27]. It is known to be a connected, compact, metrizable topological space, but not path-connected. However, each one of the "intermediate" space* $X_n = \{(x_i)_{1\leq i\leq n} \in \prod_{i=1}^n S^1, \forall i < n, 2x_{i+1} = x_i\}$ *is homeomorphic to a circle* S^1*, via the map* $(x_n, 2x_n, \dots) \in X_n \mapsto x_n$*. Hence, any space* X_n *is clearly path-connected, as it is a circle. In fact,* X_n *can be visualized as a closed coil, with* 2^n *"spirals", the first and last spirals being connected to one another. However, the solenoid itself is not projectively connected. To show this, consider the point* $x = (0, 0, \dots) \in X$*, and the point* $y = (\frac{1}{2}, \frac{1}{4}, \frac{5}{8}, \dots)$*, where for all* $i > 0$*,* y_{i+1} *is the only preimage of* y_i *by* $x \mapsto 2x$ *in* $[\frac{1}{4}, \frac{3}{4})$*. In* *particular, the distance (in S^1) between $x_i = 0$ and any y_i is at least $\frac{1}{4}$ for any i. Suppose that there exists an inverse system of paths $(p_n)_{n>0}$ between x and y, that is, p_n is a path in $X_n \simeq S^1$ between 0 and y_i, and $2p_{n+1} = p_n$. But this is impossible: for any n, the path p_n must be of length at least $\frac{1}{4}$ by the previous remark, and we can assume that it is non-backtracking. As for any $m \geq 0$, $4^m p_{n+2m} \sim p_n$, we get that p_n is homotopic to a non-backtracking path in S^1 of length 4^{m-1}. By Example 11, all those paths cannot be homotopic, so we have a contradiction.*

This example shows that even for compact metrizable connected spaces, the fact that the intermediate spaces in the inverse limit are all path-connected is not enough to guarantee projective connectedness. On the other hand, this is obviously not enough to conclude for the specific construction considered in the fundamental group of subshifts.

Let us give an additional example, this time using subshifts:

Example 14 (Balanced subshift)**.** *Let Z be the one-dimensional subshift consisting of all the infinite balanced words (considering this set as a single subshift, rather than studying individual Sturmian subshifts, is not a new idea, and was already done in [BK98] under the name Grand Sturmian subshift). Let $z_0 = 0^{\mathbb{Z}}$. As z_0 is balanced, it is a point of Z. We show that the scene-space $S_{\mathcal{B}_n}$ of Z is path-connected, by showing that for any balanced word $u \in \mathcal{L}(Z)$, there exists a path between $(z_0|_{\mathcal{B}_n}, \mathbf{0})$ and $(u, \mathbf{0})$. We do not provide the complete computations, and only give the general idea. There might exist simpler proofs, using more advanced results on Sturmian subshifts (see [Pyt+02, Chapter 6] for example), but we simply use elementary ideas.*

Clearly, Z is far from being transitive: for any n, if $u, v \in \mathcal{L}_n(Z)$ are such that $|v|_1 \geq |u|_1 + 2$, by definition of balanced words there exist no configuration $z \in Z$ such that $u, v \in \mathcal{L}(Z)$.

However, any finite balanced word is a subword of a "mechanical word", so there exists parameters α_u, β_u such that the mechanical word x_u given by the line $x \mapsto \alpha_u x + \beta_u$ satisfies $x_u|_{\mathcal{B}_n} = u$, and we can even take α irrational. The key remark is that because u is finite, there exists several such pairs (α, β) (more precisely, by [PR12, Proposition 1], u is a factor of an infinite balanced word of slope α if and only if $|w|_1 - 1 < \alpha|w| < |w|_1 + 1$ for all $w \sqsubseteq u$), see Figure 1.7.

The idea is now the following:

- *Within an infinite aperiodic balanced word, there exists some k such that factors of length $|u|$ contain k or $k+1$ symbols 1.*
- *Using the previous remark, for any given finite balanced word u, we can find a configuration z containing u such that factors of z of length $|u|$ contain $|u|_1$ or $|u|_1 - 1$ symbols 1.*
- *Therefore, "moving" within the configuration z, we can find another factor z of weight $|u|_1 - 1$.*
- *We can then repeat to find patterns with less and less weight, until we reach the pattern $0^{\mathcal{B}_n} \sqsubseteq z_0$.*

This process is illustrated in the Figure 3.10.

However, by definition of balanced words, we have the following property: for $n > 1$, for any path (in the sense of Definition 3.29) p between $(1^n, \mathbf{0})$ and $(0^n, \mathbf{0})$, any coherent decomposition of p must be of length at least $\frac{n}{2}$. In particular, for any projective path class $([p_n])_{n>0}$ between $0^{\mathbb{Z}}$ and $1^{\mathbb{Z}}$ – the existence of which has not been proven ! – there does not exist any finite set of configuration $z_0, z_1, \ldots, z_N$ such that p_n can be decomposed in paths $p_n = q_0 * q_1 * \ldots * q_N$, with q_i traced in z_i for all i.

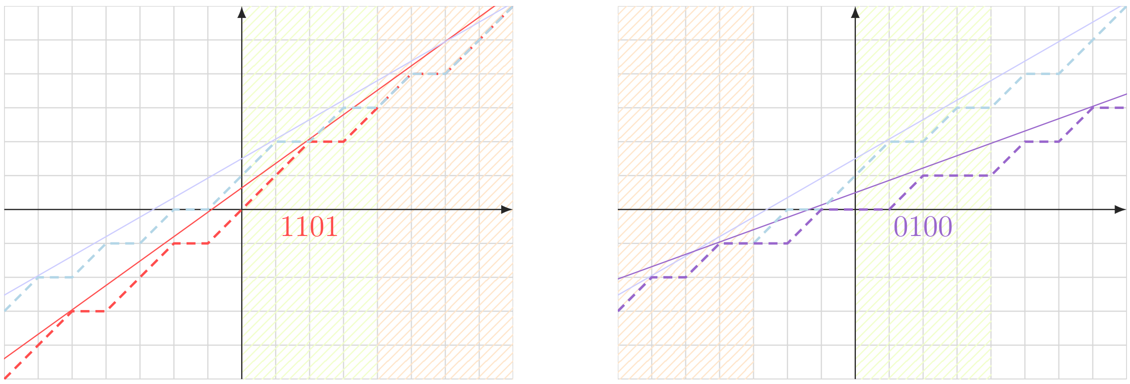

Figure 3.10: A path between $(1101, \mathbf{0})$ and $(0100, \mathbf{0})$, using an extra configuration. The red and cyan lines induce mechanical words which are equal on the interval $[\![4, 7]\!]$, and the blue and purple lines, words which are equal on $[\![-7, -4]\!]$. Note that there exists no single configuration containing both 1101 and 0100.

To sum up: if Z is projectively connected, then any projective path class between $0^{\mathbb{Z}}$ and $1^{\mathbb{Z}}$ must be non-trivial, in the sense that it cannot be obtained using a finite number of configurations $x_1, \ldots, x_N$, in which every path p_n could be decomposed as coherent paths $q_1, \ldots, q_N$, which could respectively be traced in $x_1, \ldots, x_N$. On the other hand, if Z is not projectively connected, it is a counter-example to the fact that having each scene-space path-connected implies projective connectedness.

Considering the subshift $X = Z^{\uparrow}$ shows that this behaviour also exists in $\mathbb{Z}^2$ subshifts.

3.4 Hom-shifts

We now turn our attention to a specific class of subshifts, the Hom-shifts. Hom-shifts are a subclass of higher-dimensional subshifts of finite type, introduced in [Cha16], that are defined using graphs.

3.4.1 Definition and first results

In this section and unless specified otherwise, all graphs are assumed to be finite, connected, undirected, simple, but can have self-loops (that is, an edge $\{v, v\}$). For a general overview of the terminology used in this section, see Section 4.3.1.

Definition 3.60: Hom-shift

Let $G = (V, E)$ be an undirected graph. The $\mathbb{Z}^d$ Hom-shift $X_G \subset V^{\mathbb{Z}^d}$ is defined by

$$X_G = \{x \in V^{\mathbb{Z}^d} \mid \forall \mathbf{u}, \mathbf{v} \in \mathbb{Z}^d, \|\mathbf{u} - \mathbf{v}\|_\infty = 1 \implies (x_{\mathbf{u}}, x_{\mathbf{v}}) \in E\}$$

In other words, X_G is the set of V-colourings of $\mathbb{Z}^d$ where adjacent cells of $\mathbb{Z}^d$ must respect the constraints given by the edges of G: this can be seen as a generalization of proper colourings (colourings of $\mathbb{Z}^d$ where adjacent cells must not coloured with the same colour), as the set of proper n-colourings is X_{K_n}, the Hom-shift associated with the n-clique K_n. We will only consider the case $d = 2$ in this section, and we write X_G for the $\mathbb{Z}^2$ Hom-shift associated with G.

Hom-shifts also correspond exactly to the SFTs where the constraints (or forbidden patterns) are the same in any direction. In particular, for any graph G, if $x \in X_G$ then its rotations are also in X_G, for example, $((i,j) \mapsto x_{(j,-i)}) \in X_G$. The name "Hom-shift" comes from the fact that the corresponding configurations of X_G are graph homomorphisms, from $\mathbb{Z}^d$ to G, where $\mathbb{Z}^d$ is viewed as a graph with vertex set $\mathbb{Z}^d$ and edges between adjacent vertices.

Most of the ideas presented in this section can be found, although sometimes in a slightly different form, in the original article [Cha16], or in continuations of this work on Hom-shifts [HGO22]. Some ideas are already implicit in [GP95, Theorem 3] and [Sch98, Section 7], in the special case of proper 3-colourings of the plane, which is the Hom-shift associated with the cyclic graph C_3. We prove a few additional results, and apply those ideas to compute the projective fundamental group of some Hom-shifts.

We start with a proposition about the projective connectedness of Hom-shifts:

Proposition 3.61

Let G be any connected graph. Then, X_G is projectively connected if and only if it is not bipartite. If G is bipartite, it admits exactly two projective path components, $X = Y \sqcup \sigma_{(1,0)}(Y)$.

Proof. If G is bipartite, it is easy to see that it is not projectively connected, for example using Corollary 3.44 and the block map $\phi\colon X_G \to \{0,1\}^{\mathbb{Z}^2}$ sending each vertex to its bipartite component in G. Then, $\phi(X_G)$ is the 2-points subshift containing the two proper 2-colourings of $\mathbb{Z}^2$, which is periodic.

Suppose then that G is not bipartite. By Proposition 3.49, it suffices to show that X_G is cone-connected. Let u, v be adjacent vertices of G, and define a configuration $x\colon (i,j) \in \mathbb{Z}^2 \mapsto \begin{cases} u & \text{if } i+j = 0 \mod 2 \\ v & \text{otherwise} \end{cases}$. It is clear that $x \in X_G$. Let now $y \in X_G$ be arbitrary. Let $u_0 = u = x_{(0,0)}, u_1, \ldots, u_N = y_{(0,0)}$ be a path in G of even length – such a path exists as G is not bipartite. Then, consider the configuration z presented in Figure 3.11:

As N is even, it contains the cone $(-N,0)+C(\frac{1}{2},(-1,0))$ of x, and the cone $C(\frac{1}{2},(1,0))$ of y. All the other cells are obtained by repeating in the suitable diagonal the value of the border of the cone $C(\frac{1}{2},(1,0))$ in y. Hence, X_G is cone-connected, and so projectively connected. Note that this construction can also be performed when G is bipartite, provided that the chosen vertex u and $y_{(0,0)}$ are in the same bipartite component, so that there exists a path of even length between them. □

The case of trees

We start with the simplest possible graphs, namely, trees, which are graphs without cycles. In fact, those results will be useful when studying more general Hom-shifts, as we will show in Section 3.4.1 how we can relate the projective fundamental group of graphs to those computations on trees. In order to study paths and their deformation in Hom-shifts, we introduce an operation on configurations, that we call the **pivot** operation following [Cha18]. This is a well-known operation of the literature, also mentioned for example in [Rém05, Section 4.1] in the case of tilings by rectangles. We do not study this operation in depth, and simply state what is necessary in our context.

Notation. For G a graph and $u \in V(G)$, we note $N(u) = \{v \in V(G), (u,v) \in E(G)\}$ the neighbourhood of u in G.

Figure 3.11: A configuration containing a cone of x (the u-Teal/v-Red checkerboard pattern) and a cone of y (the purple dashed cells). The path $u_0, u_1, \ldots, u_4$ between u and $y_{(0,0)}$ in G is represented with other colours.

Definition 3.62: Pivot

Let X be a subshift. A **pivot** from $x \in X$ is a configuration y differing from x in at most a single point.

Lemma 3.63 [Cha18]

Let G be a graph, and let $u, v \in V(G)$ be vertices such that $N(u) \subseteq N(v)$. Then, for $x \in X_G$, replacing any occurrence of u in x by v produces a pivot y in X_G.

This property will be important to understand how we can deform paths in Hom-shifts:

Proposition 3.64: Projective Fundamental Group - trees

Let T be a finite tree. Then X_T has exactly two projective path components, each of those having a trivial projective fundamental group.

Proof. The strategy used here is the same as the one of [GP95, Sections 5 and 6], in the specific case of 3-colourings and their universal covering, the infinite line graph. Indeed, those ideas are easily generalized to the case of any Hom-shift X_T when T is a tree.

Let r be an arbitrary vertex of T, which we call the *root* of T. Write $d_T : T^2 \to \mathbb{N}$ the distance in T, and $\text{prev}_T : T \to T$ the map associating with each vertex its father in the rooted tree (T, r), with $\text{prev}_T(r)$ being an arbitrary neighbour of r, say $r' \in T$. Let $X_{(T,r)}$ the set of configurations of the Hom-shift X_T whose value at the origin is a vertex at an even distance of r in T, *i.e.* $X_{(T,r)} = \{x \in X_T \mid d_T(r, x_{(0,0)}) \text{ is even}\}$. By Proposition 3.61, $X_{(T,r)}$ is one of the two projective path components of X, the other one being $X_{(T,r')}$. Let $\bar{x} \in X_{(T,r)}$ be the configuration containing only r and r' with r at the origin. We show that $X_{(T,r)}$ has a trivial fundamental group. To do this, it suffices to show that for any projective loop class $(p_n)_{n>0}$ where $\mathcal{B}_n$ is a loop in $S_{\mathcal{B}_n}$ based at $(\bar{x}|_{\mathcal{B}_n}, \mathbf{0})$, the path p_n is homotopic to the trivial path. Fix then some projective loop-class $([p_n])_{n>0}$, and $n > 0$. By Lemma 3.38, we can assume that p_n is N-straight, for some $N >> n$ and $N >> \text{diam}(T)$ where $\text{diam}(T) = \max_{u,v \in V(T)}(d_T(u, v))$. For a rectangle $R \subset \mathbb{Z}^2$, we define a map $\phi_R : \mathcal{L}_R(X_{(T,r)}) \to X_{(T,r)}$ as follows:

$$\phi_R(x)_{i,j} = \begin{cases} x_{(i,j)} & \text{if } (i,j) \in R \\ \text{prev}^k(x_{v^*}) & \text{otherwise, with } k = d((i,j), R), v^* = \operatorname*{argmin}_{v \in R} d((i,j), v) \end{cases}$$

where d is the distance induced by the norm $\|\|_\infty$ in $\mathbb{Z}^2$. We write ϕ_n for $\phi_{\mathcal{B}_n}$, and we extend ϕ_R to $X_{(T,r)}$ by defining $\phi_R(x) = \phi_R(x|_R)$ for $x \in X_{(T,r)}$.

Claim 22. ϕ_R is well-defined.

Proof. As R is a rectangle, for any $\mathbf{u} \in \mathbb{Z}^2$, there exists a single $\mathbf{v} \in R$ minimizing $d(\mathbf{u}, v)$. In particular, $\phi_R(x)$ is indeed a map. We now need to check that if $P \in \mathcal{L}_R(X_{(T,r)})$, then $\phi_R(P) \in X_{(T,r)}$. This holds if and only if for any $\mathbf{u}, \mathbf{v} \in \mathbb{Z}^2$ such that $\|\mathbf{u} - \mathbf{v}\|_\infty = 1$ then $(x_\mathbf{u}, x_\mathbf{v}) \in E(G)$. Let then $\mathbf{u}, \mathbf{v}$ be neighbouring cells in $\mathbb{Z}^2$, $\mathbf{u}^*, \mathbf{v}^*$ be the respective projections in R, and $d_\mathbf{u} = d(\mathbf{u}, \mathbf{u}^*), d_\mathbf{v} = d(\mathbf{v}, \mathbf{v}^*)$ be the respective distances. We are then in one of the following cases:

- If $\mathbf{u}^* = \mathbf{v}^*$, then $d_\mathbf{u} = d_\mathbf{v} \pm 1$. Without loss of generality, we then have $\phi_R(\mathbf{u}) = \text{prev}^{d_\mathbf{u}}(P_{\mathbf{u}^*}) = \text{prev}^{d_\mathbf{v}+1}(P_{\mathbf{u}^*}) = \text{prev}(\phi_R(\mathbf{v}))$.
- Otherwise, then $\mathbf{u}^* - \mathbf{v}^* = \mathbf{u} - \mathbf{v}$ and $d_\mathbf{u} = d_\mathbf{v}$, so $\text{prev}^{d_\mathbf{u}}(P_{\mathbf{u}^*})$ and $\text{prev}^{d_\mathbf{v}}(P_{\mathbf{v}^*})$ are also neighbours in G.

■

The idea of the proof is illustrated in Figure 3.12. Let us call "segment" of the path p_n a part of the path where the trajectory is comprised between two consecutive points of the $(N\mathbb{Z})^2$ sublattice. As X_G is a Hom-shift, the concatenation of the patterns seen along a horizontal (resp. vertical) segment is a globally admissible rectangular pattern of support a rectangle $(N+2n+1) \times (2n+1)$ (resp. $(2n+1) \times (N+2n+1)$). We now show how to use ϕ_R to deform each segment independently to a path which can almost be traced in $\bar{x}$. Fix a segment s of p_n, with starting point $s_s = (P_s, (0,0))$ and ending point $s_e = (P_e, (N, 0))$

without loss of generality, and let P be the corresponding pattern of support R as defined above. Let $y_0 = \phi_R(P)$, and let $S = \mathrm{dom}(P_s) \cup \mathrm{dom}(P_e)$ be the support of the two ending points of this segment.

Claim 23. s is homotopic to a segment s' covering a pattern P' of support R, such that for all $\mathbf{u} \in R$, if $d(\mathbf{u}, S) > \mathrm{diam}(T)$, then $P'(\mathbf{u}) \in \{r, r'\}$.

The proof of this claim is illustrated in Figure 3.12.

Proof. For any leaf u in T, we have $N(u) = \{\mathrm{prev}(u)\} \subseteq N(\mathrm{prev}^2(u))$. Denote $\mathring{T}$ the tree obtained by removing the leaves of T, except r and r'. We can then apply Lemma 3.63, and replace any occurrence of any leaf of T in $\phi_R(P)$ outside of S to obtain a configuration $y_0 \in X_T$. Now, $y_0|_{\mathbb{Z}^2\setminus S} \sqsubseteq X_{T_0}$. Repeating this operation, we can define for any $1 \leq k < \max_{u\in V(T)}(d_T(r,u))$ a tree T_k and a configuration y_k such that:

- $T_k = \mathring{T}_{k-1}$, in particular $X_{T_k} \subset X_T$.
- $y_k \in X_T$.
- $y_k|_S = P|_S$
- $y_k|_{\mathbb{Z}^2\setminus(\mathcal{B}_{n+k}\cup((N,0)+\mathcal{B}_{n+k}))} \sqsubseteq X_{T_k}$, that is, outside of a "thickened ball" around S, the support of the endpoints of s, y_k contains only vertices of T that are "close" to the root r.

Moreover, by definition of ϕ_R, for any $\mathbf{u}$ in $([\![-n,n]\!] \cup [\![N-n,N+n]\!]) \times \mathbb{Z}^2$, $y_k(\mathbf{u})$ depends only on $y_k|_S = P|_S$. In particular, the path traced in y_k with the following trajectory is homotopic to s_1, as it can also be traced in $\phi_R(P)$:

- Starting from $(0,0)$, move down to $(0,-N)$.
- Then move right to $(N,-N)$.
- Then move up to $(N,0)$.

Therefore, the path traced in $y_{\max_{u\in T}}(d_T(r,u))$ with the same trajectory as s_1 from $\mathcal{B}_n$ to $(N,0) + \mathcal{B}_n$ satisfies the claim.

■

Consider now the two first segments s_1 and s_2. The goal will be to deform $s_1 * s_2$ into a path which also verifies Claim 23, in the sense that it is a chessboard on r, r' sufficiently far from the ending point of s_2:

Claim 24. The path $s_1 * s_2$ is homotopic to a path $s_1' * s_2'$, where s_1 and s_1' (resp. s_2 and s_2') have the same trajectory, such that s_1' can b traced in $\bar{x}$.

This is enough to prove that p_n is contractible: we can indeed inductively repeat the argument on s_2' and s_3 and so on, and finally get a path traced entirely in $\bar{x}$.

Sketch of the proof of Claim 24. The figure to keep in mind is Figure 3.13. There are in fact three cases to consider, or two up to symmetry, depending on whether s_2 ends at the point (N,N) or $(2N,0)$, the former being the harder case depicted in Figure 3.13: otherwise, we can just apply Claim 23 to the path $s_1 * s_2$.

The key observation is that as T is a tree, the sequence of vertices encountered between the "chessboard part" of s_n and the "chessboard part" of s_{n+1} is itself a contractible path *in the tree* T, that is, it is essentially a backtracking walk in T – in Figure 3.13, this is the sequence $(r, v_0, v_1, v_2, v_1, v_2, v_1, v_0, r)$. This implies that $s_1 * s_2$ can in fact be traced in a single configuration, which contains only r and r' except at a bounded distance from the

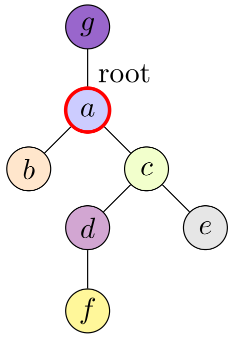

(a) An example of a tree T, rooted in a. We define $\mathrm{prev}(a) = b$.

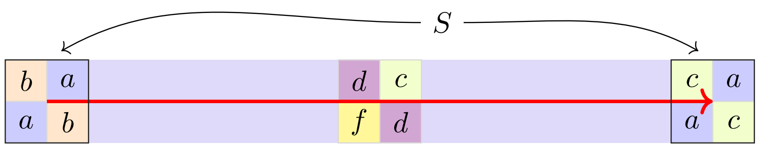

(b) An original segment in some path of $X_{(T,a)}$, covering a pattern P of support some rectangle R.

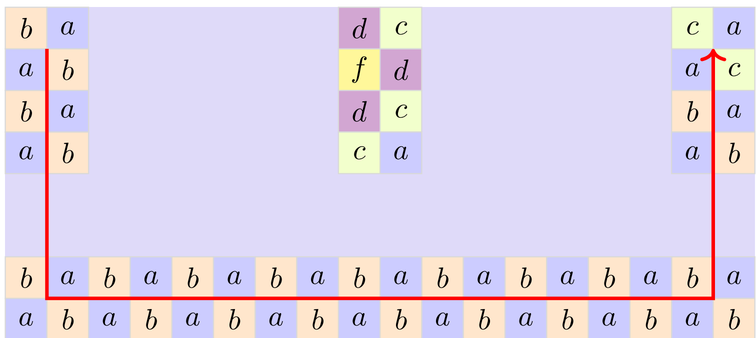

(c) Part of $\phi_R(P)$, and a path taking another trajectory in this configuration.

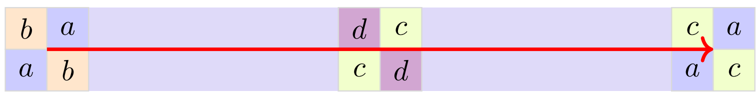

Figure 3.12: Example of a pivot used to remove occurrences of a leaf of some tree T in a path of X_T.

ending points of s_1 and s_2. In particular, in this configuration, the path whose trajectory starts from $(0, 0)$ and goes straight up to $(0, N)$ can be traced in $\bar{x}$. Repeating this process, we can deform this path back so that it has the same trajectory as the original $s_1 * s_2$. ■

□

We can even extend this to infinite trees. Note that the definitions of the fundamental group (Definition 3.26), projective connectedness (Definition 3.39) or of Hom-shifts (Defi-

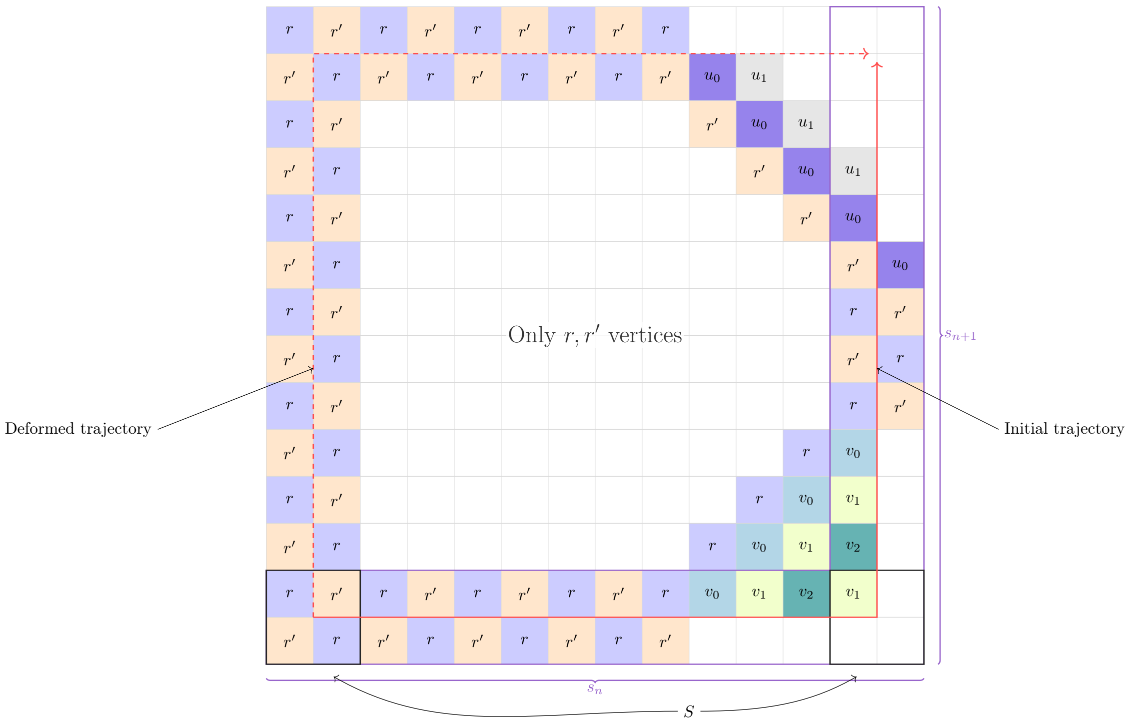

Figure 3.13: Deformation of a path $s_n * s_{n+1}$ into one which uses only two symbols, except in a bounded neighbourhood around its endpoints.

nition 3.60) also make sense for infinite graphs G, and for the associated Hom-shifts X_G. We can then prove the following:

Proposition 3.65

Let T be an infinite tree. Then X_T has exactly two projective path components, each of those having a trivial projective fundamental group.

Proof. The proof of the fact that X_T has two projective path components is the same as in Proposition 3.64. For triviality of the projective fundamental group, it suffices to notice the following: for a projective loop class $([p_n])_{n>0}$, any specific path p_n is finite. In particular, it is a valid path in $X_{T'}$ for some finite subtree $T' \sqsubseteq T$. We can now apply Proposition 3.64 and deduce that $[p_n]$ is trivial in $X_{T'}$, and therefore in X_T itself. □

Universal graph coverings

The main tool to study Hom-shifts for graphs other than trees is the **universal covering**: this is a general tool in the study of fundamental groups, as already explained in Section 3.2.1, but in the specific case of Hom-shifts, those coverings take a specific form. For the general, abstract point-of-view linking the more combinatorial notions explored here to the classical definitions of algebraic topology, we refer to [Sta83]. The point-of-view adopted here is the same as the one of [HGO22].

For simplicity, we will restrict ourselves to a specific subclass of graphs, and therefore of Hom-shifts. This restriction is the same as in [Cha16] and [CM18], and the technical tools required to generalize the main constructions to *any* graph are developed in [HGO22]. As we do not know how to adapt those additional techniques to our specific problem of computing projective fundamental groups, we stick to the easier case of **four-cycle free graphs**:

Definition 3.66: Four-cycle free

Let G be a graph. We say that it is four-cycle free if for any cycle $(v_1, v_2, v_3, v_4 = v_1)$, we have $v_1 = v_3$ or $v_2 = v_4$.

This restriction is not standard from a graph-theoretic point-of-view, and is mainly useful when looking at homomorphisms from $\mathbb{Z}^2$ to a graph; the concrete reason for why this restriction is needed will be given below, but a more abstract reason for it is still unclear, in the sense that we do not know what the corresponding condition would be for Hom-shifts on groups other than $\mathbb{Z}^d$. Most of the definitions and lemmas already appear in [Cha16, 6.Universal Covers].

Definition 3.67: Universal covering [HGO22, Def. 4.1]

Let $G = (V, E)$ be a four-cycle free graph, and fix $v \in V$. The **universal covering** of G is the undirected graph $\mathcal{U}_G(v)$ whose vertices are the non-backtracking walks in G starting from v, and whose edges are the pairs (p, q) such that $p = qu$ or $q = pu$ for some $u \in V(G)$.
We denote by $\Phi\colon \mathcal{U}_G(v) \to G$ the map sending each non-backtracking walk to its last vertex.

We will abuse notation and also write $\Phi\colon X_{\mathcal{U}_G} \to X_G$ which simply applies $\Phi\colon \mathcal{U}_G \to G$ pointwise.

We state without proofs a number of easy lemmas:

Lemma 3.68 [HGO22, Lem. 4.4]

For G a connected graph, the graphs $\mathcal{U}_G(v), v \in V(G)$ are all isomorphic.

We then simply denote by $\mathcal{U}_G$ the universal covering of a graph G. The next lemmas Lemma 3.69 and Lemma 3.71 are well-known:

Lemma 3.69

If G is a tree, then $\mathcal{U}_G \simeq T$.

Corollary 3.70

If $G = (V, E)$ is a tree with a single self-loop $(v, v) \in E$, then let G_1, G_2 be disjoint copies of G, with v_1, v_2 their respective vertices with a self-loop. Then $\mathcal{U}_G = (V_1 \sqcup V_2, E_1 \sqcup E_2 \sqcup \{(v_1, v_2)\})$

Lemma 3.71

If G admits a non-trivial cycle, then $\mathcal{U}_G$ is infinite.

Proof. Let $(v_1, v_2, \ldots, v_n)$ be a simple cycle, *i.e.* that for all $1 \leq i, j < n$, $v_i \neq v_j$ and $v_1 = v_n$. Then the walks $\underbrace{(v_1, \ldots, v_n, v_1, \ldots, v_n, \ldots)}_{k \text{ times}}$ are all non-backtracking for $k \in \mathbb{N}$, so they are different vertices in $\mathcal{U}_G$. □

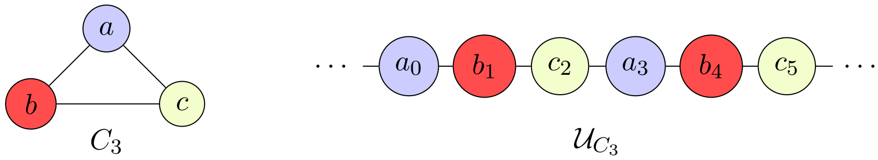

Figure 3.14: The universal covering of the cyclic graph with 3 vertices C_3

The construction of this universal covering is important due to the following property:

Proposition 3.72 [Cha16, Prop. 6.2]

Let G be a four-cycle free graph. For all $x \in X_G$, and $\tilde{v} \in \Phi^{-1}(x_{(0,0)})$, there exists a unique configuration $\tilde{x} \in X_{\mathcal{U}_G}$ such that $\tilde{x}_{(0,0)} = \tilde{v}$ and $\Phi(\tilde{x}) = x$.

Figure 3.15: Lifting a pattern from X_{C_3} to a pattern $X_{\mathcal{U}_{C_3}}$. Note that cells that were coloured with the same vertex of C_3 are not necessarily coloured with the same colour of $\mathcal{U}_{C_3}$.

Proposition 3.72 is false if we consider graphs that have four-cycles, and we need to consider a slightly different construction than universal coverings to obtain an equivalent result.

Fundamental group of graphs

The results and definitions of Section 3.4.1 are motivated by a simple observation: given a graph $G = (V, E)$, there is a natural way to view it as a topological space. Indeed, one can consider the usual topology on $[0, 1] \subset \mathbb{R}$, and then consider the following space:

- For each edge $e = (u, v)$, consider a distinct space $X_e \simeq [0, 1]$, with endpoints $x_{e,u}, x_{e,v} \in X_e$.
- Define an equivalence relation $\sim_G$ on $\bigsqcup_{e \in E} X_e$ by $x_{e,u} \sim_G x_{e',u}$ for all $e, e' \in E$ such that $u \in e \cap e'$.
- Then, $G = \bigsqcup_{e \in E} X_e / \sim_G$, with the quotient topology.

For more details on this construction, see [Hat00, 1.A Graphs and Free Groups]. In this setting, the universal covering of G defined in Definition 3.67 is a covering (in the sense of Definition 3.18) of G considered as a topological space. In particular, when viewed as a topological space, one can wonder what the fundamental group of a given graph is. In fact, this group is very easy to compute, and at least in the case of four-cycle free graphs, we will see that it is the same as the projective fundamental group of the corresponding Hom-shift.

Definition 3.73: Spanning Tree

Let G be a graph. A **spanning tree** of G is a tree included in G which is maximal for inclusion.

Lemma 3.74

Let $G = (V, E)$ be a graph. Then, any spanning tree of G has exactly $|V|$ vertices and $|V| - 1$ edges.

Proposition 3.75 [Hat00, Prop. 1.A.2]

Let $G = (V, E)$ be a graph. Then, $\pi_1(G)$ is a free group on $|E| - |V| + 1$ generators.

We do not give a proof of Proposition 3.75, which can be formally proven using the traditional tools from algebraic topology (see [Hat00, Prop. 1.A.2]), but try to give an interpretation of this result. Fix any spanning tree T of G, and consider the $m = |E| - |V| + 1$ edges of $E \setminus T$. Now, consider a cycle p in G, based at some arbitrary vertex v. As T is a tree, this cycle either uses only edges from T, in which case it is homotopic to the trivial path as it contains only backtracking subpaths, or uses edges from $E \setminus T$. As for any edge $e = (s, t) \in E \setminus T$, there exists a unique path γ_e in T (up to homotopy) from v to s, and another unique path γ'_e from t to v. One can then prove that those are non-homotopic paths, and that any path can be obtained by concatenating paths of the form $\gamma'_e e \gamma_e$ or their inverse. We get that $\pi_1(G, v) = \langle [\gamma'_e e \gamma_e], e \in E \rangle$, and in particular $\pi_1(G, v) \simeq F_m$.

Theorem 3.76

Let $G = (V, E)$ be a non-bipartite four-cycle free graph. Then, X_G is projectively connected and $\pi_1^{proj}(X_G) = \pi_1(G) = F_{|E|-|V|+1}$.

Proof. The high-level idea is to use Lemma 3.71 and Proposition 3.72 to lift loops in X_G to paths in $X_{\mathcal{U}_G}$, and then Proposition 3.65 and Proposition 3.75 to understand the structure of those paths, and what they imply on the original loops. Projective connectedness comes from Proposition 3.61, and we only need to compute the projective fundamental group.

Fix some vertex $v \in V$, and let $T = \mathcal{U}_G$ be the universal cover of G. By Lemma 3.71, T is infinite. Let $\Phi\colon X_T \to X_G$ be a covering map. From this point on, T is seen as a tree rooted in $\tilde{v}$, with $\Phi(\tilde{v}) = v$.

By Proposition 3.72, each configuration of X_G lifts to a configuration of X_T. By Propo-

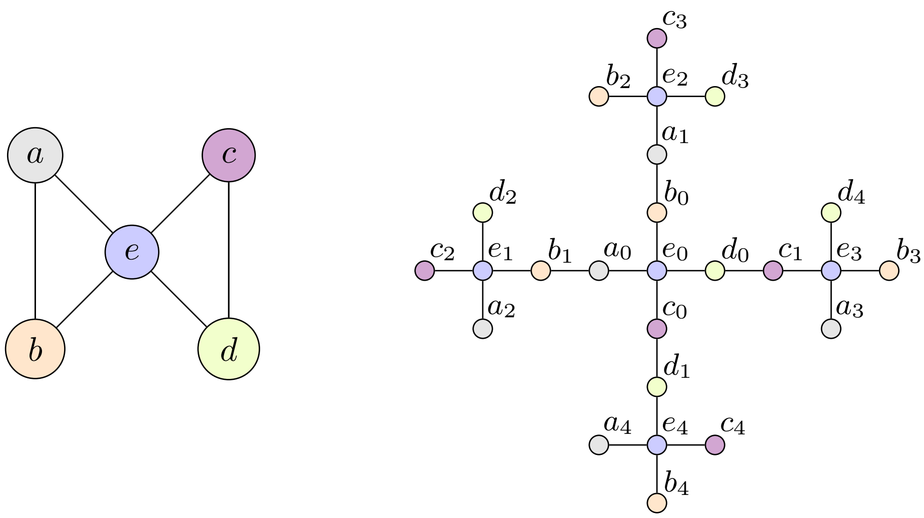

Figure 3.16: Example of a graph G and a part its universal covering, which is infinite as G has cycles. The fundamental group of G is F_2, obtained from Proposition 3.75, or from a direct computation using $\pi_1(S^1) = \mathbb{Z}$.

sition 3.65, we know that each projective path-class of X_T has a trivial fundamental group. Consider an arbitrary configuration $x \in X_G$ with $x_{(0,0)} = v$, and p a projective loop-class based on x, the components of which are the $[p_n]$ for $n > 0$, with p_n a loop in $S_{\mathcal{B}_n}(X_G)$. Fix $n > 0$. Then, p_n lifts to paths in $S_{\mathcal{B}_n}(X_T)$, and let $\tilde{p}_n$ be such a path, *i.e.* $\Phi(\tilde{p}_n) = p_n$. Note that $\tilde{p}_n$ is not necessarily a loop. Let $y^0, y^1 \in X_T$ be respectively configurations containing the starting point and the ending point of $\tilde{p}_n$, *i.e.* $\tilde{p}_n(0) = (y^0|_{\mathcal{B}_n}, (0,0))$. Let $\tilde{c} = \mathrm{restr}_{0,n}(p_n)$ be the path from $y^0_{(0,0)}$ to $y^1_{(0,0)}$ in T obtained by only looking at the central cell of the pattern seen by p_n at each timestep. Now, by definition of Φ, we have that $\Phi(y^0)|_{\mathcal{B}_n} = \Phi(y^1)|_{\mathcal{B}_n} = x|_{\mathcal{B}_n}$ and so in particular $c = \Phi(\tilde{c}_n)$ is a cycle in G.

We need to prove that this cycle is independent from the choice of y^0 and y^1 when considered as an element of the free group $\pi_1(G)$, to obtain a result which is similar to Proposition 3.21. Consider another lift $\tilde{p_n}'$ of p_n, with endpoints $y^{0'}$ and $y^{1'}$ defined as above, and let $\tilde{c}'_n$ be the corresponding path in T. We claim that $\Phi(\tilde{c}_n) = \Phi(\tilde{c}'_n)$. Indeed, by definition of $\mathcal{U}_G$ and Φ, for any $u \in V(G)$ and $w \in N(u)$, and any $\tilde{u} \in \Phi^{-1}(u)$, there exists a unique vertex $\tilde{w} \in N(\tilde{u}) \cap \Phi^{-1}(w)$. In particular, $\tilde{c}_n$ is entirely determined by its starting point, after which there is a unique lift at each timestep.

Moreover, as $\mathrm{restr}_{0,n+1}(p_{n+1}) = \mathrm{restr}_{0,n}(\mathrm{restr}_{n,n+1}(p_{n+1})) \sim \mathrm{restr}_{0,n}(p_n)$, we have that $\tilde{c}_{n+1} \sim_T \tilde{c}_n$, and therefore $\Phi(\tilde{c}_n)$ does not depend on n up to homotopy in G. This induces an injective morphism between projective loop-classes in X_G and $\pi_1(G, x_{(0,0)})$, sending each projective path-class to the associated $\Phi(\tilde{c}_n)$ for any n described above; this morphism is also easily seen to be surjective, by constructing an explicit set of configurations and a projective path in X_T between any two lifts of x in X_T, and so $\pi_1^{proj}(X_G) \simeq \pi_1(G) \simeq F_m$. □

This result was already obtained in the specific case of 3-colourings, that is, to the case of X_{C_3}, in [GP95], but combining the original ideas with the observation of Section 3.4 in general allow us to extend this to a much larger class of subshifts.

We give another informal way to understand why the projective fundamental group $\pi_1^{proj}(X_G)$ is isomorphic to $\pi_1(G)$, using the example G of Figure 3.16, with an example configuration in Figure 3.17. A point $x \in X_G$ can be seen as a "sea of vertices e in positions $S \subseteq \mathbb{Z}^2$", where connected components $\mathbb{Z}^2 \setminus S$ are "patches" of X_{C_3}, or more concretely, of X_{G_1} or X_{G_2} where G_1 is the three-cycle on vertices $a, b, e \in V(G)$, while G_2 is the three-cycle on vertices $c, d, e \in V(G)$. A path inside this configuration can therefore be viewed as a list of paths, each staying inside its "X_{C_3} component" of the configuration X. This representation is illustrated in Figure 3.17.

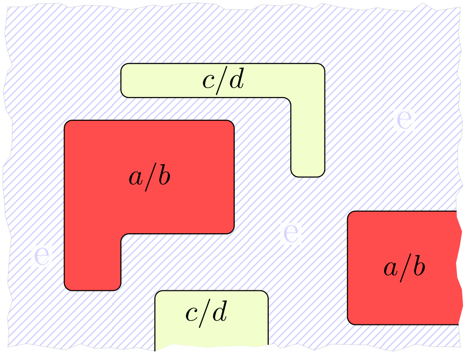

Figure 3.17: Representation of a configuration of X_G. Dashed lines represent e symbols

A similar construction has already been used in Section 2.4.2. More precisely, we can define an analogous of free products for subshifts:

Notation. For any pair of subshifts X, Y with disjoint alphabets $\mathcal{A}_X, \mathcal{A}_Y$, define the subshift $X * Y$ as the subshift on $\mathcal{A} = \mathcal{A}_X \sqcup \mathcal{A}_Y$ with forbidden patterns $\overline{\mathcal{L}(X)} \sqcup \overline{\mathcal{L}(Y)}$.

This construction is part of a more general group-subshift analogy, developed in particular in [Van19, Chapter 4, Section 3.3]

Question 3. Under what conditions do we have $\pi_1^{proj}(X * Y) = \pi_1^{proj}(X) * \pi_1^{proj}(Y)$?

3.4.2 Non-contractible Hom-shifts

We say a few words about a possible counter-example to Question 2, in the form of any $\Theta(\log)$-block-gluing Hom-shift (such examples exist, see [HGO22, Theorem 7.1]). By recent but still unpublished work of Nishant Chandgotia, it has trivial cohomology and therefore trivial projective fundamental group, using the arguments highlighted in Proposition 3.57, but it is not contractible as it is not strongly irreducible (as there exists some non-constant block-gluing examples). As the main argument is not ours and still unpublished, we formulate this as a conjecture rather than a proposition. We also state and prove lemmas only on the specific graph depicted in Figure 3.18, but a more general characterizations of graphs with the same properties exist and can be found in [HGO22, Section 4], with the important results about those graphs in [HGO22, Section 6].

Conjecture 1. *There exists subshifts with a fixed point with trivial projective fundamental group which are not contractible.*

An example of a probable example of such a subshift is the Hom-shift associated to the graph depicted in

We would need a few additional definitions and lemmas to answer Question 2 by the negative:

Definition 3.77: Block-gluing [HGO22, Def. 2.17]

Let X be a $\mathbb{Z}^2$-subshift, $f\colon \mathbb{N}^* \to \mathbb{N}$ and $k \in \mathbb{N}$. We say that X is (f, k)-phased block gluing if for $n > 0$ and any $P, P' \in \mathcal{L}_n(X)$, for any $\mathbf{u} \in \mathbb{Z}^2$ with $\|\mathbf{u}\|_\infty \geq f(n) + n$, there exists $\mathbf{v} \in \mathbb{Z}^2$ and $x \in X$ such that:

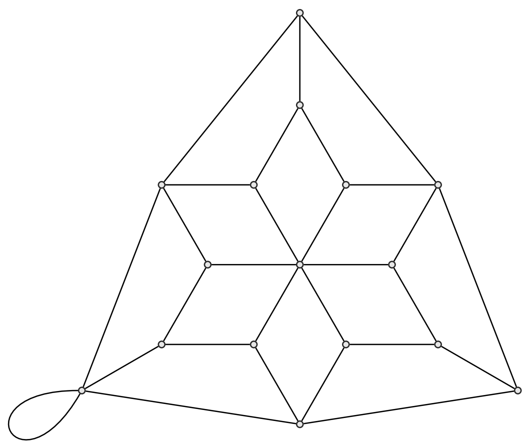

Figure 3.18: Ken-Katabami graph from [HGO22, Section 7], with an additional self-loop. We use the name "Ken-Katabami" for the graph without the self-loop.

- $\|\mathbf{v}\|_\infty < k$
- $x|_{\mathcal{B}_n} = P$
- $x|_{\mathcal{B}_n+\mathbf{u}+\mathbf{v}} = P'$.

We say that X is f-block-gluing if it is $(f, 1)$-phased block gluing.

A (f, k)-block gluing subshift is a subshift in which any $n \times n$ square patterns can be glued together, provided they are at distance at least $f(n)$, with possibly a small "phase correction" which is an offset of norm at most k.

Theorem 3.78 [HGO22, Thm. 7.1]

The Hom-shift associated with the Ken-Katabami graph is $(\Theta(\log n), 2)$-phased-block gluing.

Using the non-trivial result of Theorem 3.78, we can obtain an interesting result about projective fundamental groups of Hom-shifts:

Lemma 3.79

Let G be any graph obtained by adding a self-loop to the Ken-Katabami graph, as depicted in Figure 3.18. Then X_G is $(\Theta(\log n), 1)$-block gluing.

Proof. We give a sketch of the proof. Write K for the original Ken-Katabami graph without the self-loop. Clearly, X_G is $O(\log n)$-block gluing, as K is already $O(\log n)$-(phased). block gluing. We show that if X_G is f-block gluing then X_K is $(O(f), 2)$-phased block gluing, which then proves that X_G is $\Theta(f)$-block gluing. It suffices to consider patterns of rectangular $n \times 1$ rather than full $n \times n$ squares to compute the gluing functions of X_G, X_K.

Suppose X_G is f-block gluing for some $f\colon \mathbb{N}^* \to \mathbb{N}$. Without loss of generality, assume that f only takes even values. For a sufficiently large n, consider two $n \times 1$ "rows" r_n, r'_n which can only be glued at distance $\Theta(\log n)$ in X_K – such patterns exist by Theorem 3.78. We also only consider rows with $r_n(0,0), r'_n(0,0)$ in the same bipartite component of K. As $K \sqsubseteq G$, those patterns are also globally admissible in X_G. As we assumed that X_G is f-block gluing, we can then consider a rectangular pattern P of size $f(n) \times n$, such that $P|_{[\![0,n-1\times\{0\}]\!]} = r_n$ and $P|_{[\![0,n-1]\!]\times\{f(n)-1\}} = r'_n$. Let v be the unique vertex of G with a self-loop, and define $\hat{\mathcal{U}}_G = (V_1 \sqcup V_2, (E_1 \sqcup E_2 \sqcup \{(v_1, v_2)\}))$ where $(V_1, v_1) = (V_2, v_2) = (V, v)$, and $E_1 = E_2 = E$ (this is two copies of K with an extra, single edge between the two). Let $\Phi\colon \hat{\mathcal{U}}_G \to G$ be the obvious projection map. The notation $\hat{\mathcal{U}}_G$ is chosen for the following reason:

Claim 25. For any $P \in \mathcal{L}(X_G)$, and $\tilde{v} \in \hat{\mathcal{U}}_G$ with $\Phi(\tilde{v}) = P_{(0,0)}$, there exists a unique pattern $\tilde{P} \in X_{\hat{\mathcal{U}}_G}$ such that $\Phi(\tilde{P}) = P$ and $\tilde{P}_{(0,0)} = \tilde{v}$.

Proof. Write $\hat{\mathcal{U}}_G = G_1 \sqcup G_2$ with $G_1 \simeq G_2 \simeq G$, and $\Psi_1\colon G \to G_1 \sqsubseteq \hat{\mathcal{U}}_G$, $\Psi_2\colon G \to G_2 \sqsubseteq \hat{\mathcal{U}}_G$ the two sections of Φ. Suppose without loss of generality that $\tilde{v} \in G_1$. Consider the pattern $\tilde{P}$ defined by:

$$\begin{aligned}\tilde{P}\colon \mathrm{dom}(P) &\to V(\hat{\mathcal{U}}_G)\\ \mathbf{u} &\mapsto \begin{cases}\Psi_1(P_{\mathbf{u}}) & \text{if } \|\mathbf{u}\|_\infty = d_G(P_{(0,0)}, P_{\mathbf{u}}) \mod 2\\ \Psi_2(P_{\mathbf{u}}) & \text{otherwise}\end{cases}\end{aligned}$$

One can check that it is a valid pattern of $X_{\hat{\mathcal{U}}_G}$. ■

As K contains no self-loops, $\tilde{r_n} = \tilde{P}\Big|_{[\![0,n-1]\!]\times\{0\}}$ on the one hand, and $\tilde{r'_n} = \tilde{P}\Big|_{[\![0,n-1]\!]\times\{f(n)-1\}}$ on the other hand, must be coloured entirely with either G_1 or G_2. Suppose without loss of generality that $\tilde{r_n} \in G_1^{[\![0,n-1]\!]\times\{f(n)-1\}}$. For

parity reasons, we must also have $\tilde{r'_n} \in G_1^{[\![0,n-1]\!]\times\{f(n)-1\}}$. Fix $i \in [\![0, n-1]\!]$. Let $j_{min} = \min\{j \in [\![0, f(n)-1]\!] \tilde{P}_{(i,j)} \in G_2\}$ and $j_{max} = \max\{j \in [\![0, f(n)-1]\!] \tilde{P}_{(i,j)} \in G_2\}$. Let $u \in N(v_1) \cap V(G_1)$. Replace all the vertices of $\{i\} \times [\![j_{min}, j_{max}]\!]$ in $\tilde{P}$ alternatively by $v_1 \in V(G_1)$ and u. Because there is a single edge (v_1, v_2) between the $G_1, G_2 \sqsubseteq \hat{\mathcal{U}}_G$, one can check that the resulting pattern is valid in $X_{\hat{\mathcal{U}}_G}$. This new pattern is mapped by Φ to a pattern of X_K, whose first and last rows are r_n, r'_n. Hence, X_K is $(O(f), 2)$-phased block gluing. □

Corollary 3.80

Let G be the graph of Figure 3.18. Then X_G is projectively connected, has a fixed point, but is not contractible.

Proof. The addition of the self-loop means that G is no longer bipartite, and therefore Proposition 3.61 shows that it is projectively connected. If the self-loop is the edge $(v, v) \in E(G)$, then $v^{\mathbb{Z}^2} \in X_G$ so X_G has a fixed point, and the fact that X_G is not constant-block gluing by Lemma 3.79 shows that it is not strongly irreducible, hence not contractible. □

3.5 Finitely presented groups and SFTs

We are now going to prove our main result: any finitely presented group is the fundamental projective group of some SFT. This section closely follows the article [PV23].

Theorem 3.81

Let $G = \langle S \mid R \rangle$ be a finitely presented group. Then, there exists an SFT X such that:

- X is projectively connected.
- $\pi_1^{proj}(X) \simeq G$

3.5.1 The construction

The subshift X that we construct will informally consist of oriented wires, drawn on an empty background, each wire corresponding to a generator $\mathsf{s} \in S$ of the group $G = \langle S, R \rangle$. This idea had also been used in [Ein01], in a slightly different context – the author studied so-called fundamental cocycles, see [Sch98] – which are closely related to projective fundamental groups, to realize free groups as "groups in which the fundamental cocycles take their values". Unfortunately, we need to modify its construction to realize more general groups, and we prove our results using a more combinatorial and less algebraic approach.

We only authorize the wires to go up, perhaps in some kind of "zigzag" manner, but never down or horizontally. More precisely, we define the following tiles: first of all, a tile that we call **empty**, visually represented by $\square$, and we denote by $\mathcal{T}_{\text{empty}}$ the singleton containing this tile. We denote by $x_\square \in X$ the configuration which only contains empty tiles, and its patterns are called **empty patterns**. Then, for each element $\mathsf{s} \in \bar{S} = S \cup \{\mathsf{s}^{-1} | \mathsf{s} \in S\}$, we also consider the set $\mathcal{T}_\mathsf{s}$ of the 5 following tiles:

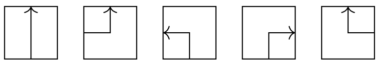

If $\mathsf{s} \neq \mathsf{s}'$, then $\mathcal{T}_\mathsf{s} \cap \mathcal{T}_{\mathsf{s}'} = \emptyset$. Distinct $\mathcal{T}_\mathsf{s}$ will be represented by wires of different colours in the figures. These tiles will, intuitively, be used to represent generators of the group in valid configurations of X. Finally, we use some other tiles that will play the role of representing the group relations. We can always assume that R contains the trivial relators ss^{-1} and $\mathsf{s}^{-1}\mathsf{s}$ for all $\mathsf{s} \in S$. Now, for each relator $\mathsf{r} = \mathsf{r}_1\mathsf{r}_2 \dots \mathsf{r}_\mathsf{n} \in R$, we let $\mathcal{T}_r$ be the tiles described by Figure 3.19.

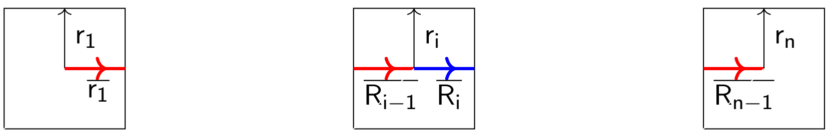

(a) Start (b) For $2 \leq i < n$ (c) End

Figure 3.19: The relation tiles.

The wire exiting from the right side of the tile Figure 3.19a does *not* have the same colour as the one exiting from the top. The former colour is denoted by $\overline{\mathsf{r}_1}$, to differentiate it from the actual r_1 wires. In the other tiles, $\overline{\mathsf{R}_\mathsf{i}} = \overline{\mathsf{r}_1\mathsf{r}_2 \dots \mathsf{r}_\mathsf{i}}$. Hence, for each relator $\mathsf{r}_1 \dots \mathsf{r}_\mathsf{n}$, we have one tile of type Figure 3.19a and one of type Figure 3.19c, and $n-2$ tiles of type Figure 3.19b. Tiles belonging to some $\mathcal{T}_r$ are called relation tiles. Note that if $\mathsf{u} \in R$ is such that it is the prefix of two different relators, *i.e.*, there exists $\mathsf{v}, \mathsf{v}' \in \bar{S}^*$ such that $\mathsf{uv} \in R, \mathsf{uv}' \in R$ then the colours $\bar{\mathsf{u}}$ are shared by the tiles used to represent those relators and so $\mathcal{T}_{\mathsf{uv}} \cap \mathcal{T}_{\mathsf{uv}'} \neq \emptyset$. X is the subshift generated by the tileset $\mathcal{T} = \mathcal{T}_{\text{empty}} \cup \bigcup_{\mathsf{s} \in \bar{S}} \mathcal{T}_\mathsf{s} \cup \bigcup_{\mathsf{r} \in R} \mathcal{T}_\mathsf{r}$ along with the obvious adjacency rules: any wire must be extended, by a wire with the same orientation given by the arrows – *e.g.*, 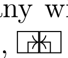 and 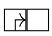 are forbidden patterns, but 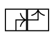 is allowed (assuming the two tiles contain a wire of the same colour). Note that the tiles defined here are not Wang tiles, although they define a nearest-neighbour SFT.

We now formalize what we really mean by a wire.

Definition 3.82: Wire

A **wire** is a sequence $\mathcal{U} = (T_t, \mathbf{v_t})_{t\in I}$, $I \subseteq \mathbb{Z}$ a non-necessarily finite interval, of pairs of non-empty tiles and $\mathbb{Z}^2$ points, such that

- $\|\mathbf{v_{t+1}} - \mathbf{v_t}\|_\infty = 1$,
- The tile T_{t+1} in position $\mathbf{v_{t+1}}$ extends the wire of tile T_t in position $\mathbf{v_t}$: placing a tile ⍐ above or below another tile ⍐ does extend it, while placing it on its right or left side does not, although they are valid patterns of X.
- $\mathcal{U}$ does not contain two consecutive relation tiles.

Remark. *We do not prevent a wire from moving back and forth: it is possible to have* $(T_t, \mathbf{v_t}) = (T_{t+2}, \mathbf{v_{t+2}})$.

Definition 3.83: Coherent wire

We say that a wire is **coherent** if there exists a configuration $x \in X$ such that for any tile $(T_i, \mathbf{v_i})$ of the wire, $x_{\mathbf{v_i}} = T_i$.

Remark. *Valid configurations of X can contain non-intersecting infinite wires, and possibly some relation tiles with wires originating from them. Any relation tile belongs to one horizontal line of k relation tiles, corresponding to a valid relator* $\mathsf{r_1} \dots \mathsf{r_k}$.

One important concept associated to paths on this subshift is the idea that paths can cross wires. Informally, this is what happens when the window, and in particular, its center, moves from one side to the other of a given wire in a path.

Definition 3.84: Crossing a wire tile

Let $n > 0$, and let $\mathbf{v}, \mathbf{v}' \in \mathbb{Z}^2$ be two adjacent points, and P, P' two patterns of respective support $\mathbf{v} + B_n, \mathbf{v}' + B_n$ such that $(P, \mathbf{v}), (P', \mathbf{v}')$ is a valid path. For $(i,j) \in B_n$, let $T_{(i,j)}$ be the tile whose bottom-left corner is on (i,j) in P. We say that this path crosses a wire tile if

- $\mathbf{v}' - \mathbf{v} = e_0 = (1,0)$ (resp. $-e_0$) and the tile $T_\mathbf{v}$ (resp. $T_{\mathbf{v}-e_0}$) was of one of the following forms:

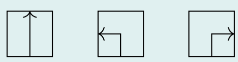

- $\mathbf{v}' - \mathbf{v} = e_1 = (0,1)$ (resp. $-e_1$) at the next step $t+1$ and the tile $T_\mathbf{v}$ (resp. $T_{\mathbf{v}-e_1}$) was of one of the following form:

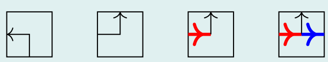

Definition 3.85: Seeing a wire

A path $p = (P_i, \mathbf{v_i})_{i \leq N}$ **sees** a wire $\mathcal{U}$ if there exists a timestep $i \leq N$, and $(T_j, \mathbf{v_j}) \in \mathcal{U}$ such that the tile in position $\mathbf{v_j}$ in P_i is T_j.

Definition 3.86: Crossing a wire

A path **crosses** a wire if it crosses one of its tiles.

3.5.2 Only Crossed Wires Matter

Our final goal is to prove that the projective fundamental group of this subshift X is the group $G = \langle S|R \rangle$. To do so, the idea will be to associate an element of the group to each path, according to the wires that it crosses. The following lemmas can be seen as a procedure to put paths in some kind of normal form via homotopies, depending only the sequence of crossed wires, regardless of the underlying geometry of the path. All the lemmas consider paths that both start and end in empty patterns, but this is not really a restriction as we will later prove that the subshift X is projectively connected, and so we will only consider loops based at $x_\square$. Unless stated otherwise, all the considered paths are using some B_n as aperture window. We start with some easy statements about patterns of support B_n, and the wires they may contain.

Lemma 3.87: Wire Order Lemma

Let $x \in X$, and let $\mathcal{U}, \mathcal{V}$ be two infinite wires in x. Suppose that $\mathcal{U}, \mathcal{V}$ do not contain relation tiles.

- For all $z \in \mathbb{Z}$, there exists between one and two $z^0_{\mathcal{U}} \in \mathbb{Z}$ such that $\mathcal{U}$ passes through the position $(z^0_{\mathcal{U}}, z)$. If there are two such $z^0_{\mathcal{U}}$, then they are necessarily adjacent, *e.g.*, 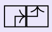 side-by-side.
- Let $z \in \mathbb{Z}$, and $z^0_{\mathcal{U}}, z^0_{\mathcal{V}} \in \mathbb{Z}$ as in the previous point respectively for $\mathcal{U}$ and $\mathcal{V}$. If $z^0_{\mathcal{U}} < z^0_{\mathcal{V}}$, then for all $z_{\mathcal{U}}, z_{\mathcal{V}}, z \in \mathbb{Z}$ such that $(z_{\mathcal{U}}, z) \in \mathcal{U}, (z_{\mathcal{V}}, z) \in \mathcal{V}$, we have $z_{\mathcal{U}} < z_{\mathcal{V}}$. Intuitively, this means that wires can globally be ordered from left to right.

If $\mathcal{U}$ or $\mathcal{V}$ contains a relation tile, then the previous claims are true only for z large enough.

Remark. *Note that the previous lemma is true because we consider wires $\mathcal{U}, \mathcal{V}$ belonging to some configuration. It is clearly false for arbitrary wires.*

Lemma 3.88

Let P be a globally admissible pattern of support B_n for some $n > 0$. Let $\mathcal{U}$ be a wire in P without relation tiles. Suppose that $\mathcal{U}$ passes to the right (resp. left) of $(0, 0)$ in P. Then, $\mathcal{U}$ neither enters nor exits P on its left (resp. right) edge.

Proof. This directly follows from the fact that no tile contains a horizontal wire, and that B_n is a square. □

Corollary 3.89

If P is a globally admissible pattern that sees a wire $\mathcal{U}$ with no relation tiles, and $x \in X$ is such that $x_{|B_n} = P$, then $\sigma^{4n}_{(0,1)}(x)_{|B_n}$ and $\sigma^{-4n}_{(0,1)}(x)_{|B_n}$ do not see $\mathcal{U}$.

In order to show that the homotopy class of a path p is indeed only determined by the wires it crosses, we will need several lemmas in which the proof will always be similar: an induction on the length L of a coherent decomposition (Definition 3.32) of p:

- for $L = 1$ (*i.e.* p is coherent), we explicitly show how to deform p to obtain the required property.
- for $L = 2$ we use the Corollary 3.94 to "normalize" both coherent subpaths of p using the base case $L = 1$.
- In general, if $p = p_1 * \ldots * p_N$, we can deform both p_1 and p_2 so that $p \sim p'_1 * p'_2 * \ldots * p_N$, in such a way that we can apply the base case to p'_1, and the induction case to $p'_2 * \ldots * p_N$.

The key step is therefore to properly show how to deal with the case $L = 2$; this is the purpose of the Corollary 3.94 that we now show, after some preliminary results.

Lemma 3.90: Finite Extension Lemma

Let P be an extensible finite pattern of X, there exists $x \in X$ containing P, such that x contains a finite number of wires.

The next lemma is especially important: it will be useful to prove projective connectedness of the subshift X, and to understand how we can deform paths that are not coherent, which is necessary to prove most of the lemmas in Section 3.5.3.

Lemma 3.91: Extensibility Lemma

Let $n > 0$, $\mathbf{u} = (0, \pm 1)$ and $\mathbf{v} = -\mathbf{u}$. There exists $\mathbf{o} \in \mathbb{Z}^2$ such that for any $y \in X$, there exists $z \in X$ with:

- $z|_{C(\frac{1}{2},\mathbf{u})+\mathbf{o}} = y|_{C(\frac{1}{2},\mathbf{u})+\mathbf{o}}$
- $z|_{C(\frac{1}{2},\mathbf{v})} = x_{\square}|_{C(\frac{1}{2},\mathbf{v})}$

Proof. We prove the case $\mathbf{u} = (0, 1)$, the case $\mathbf{u} = (0, -1)$ being similar. The picture to keep in mind in this proof is Figure 3.20.

Let r be the length of the longest relator in the finite presentation of $G = \langle S, R \rangle$, and let $\mathbf{o} = (0, r)$. Let $W \subset \mathbb{Z}^2$ be the set of positions of tiles that are part of a wire of y that:

- either passes by $C(\frac{1}{2}, \mathbf{v})$

- or originates from a relation tile which is itself part of a relator intersecting $C(\frac{1}{2}, \mathbf{v})$.

Let $\bar{C} = C(\frac{1}{2}, \mathbf{v}) + (\llbracket -r, r \rrbracket \times \{0\})$. This is a "thickened" version of the cone $C(\frac{1}{2}, \mathbf{v})$. As in Definition 2.13, we denote $\partial\bar{C}$ the points of $\bar{C}$ adjacent to $\mathbb{Z}^2 \setminus \bar{C}$.

Now, construct z as follows:

- for $(i, j) \in \bar{C} \cap W$, set $z_{(i,j)} = y_{(i,j)}$. The other tiles of $\bar{C}$ are empty.
- for $(i, j) \in (\partial\bar{C}) \cap W$ if $i < 0$, or $\square$ and $\square$ if $i \geq 0$.
- all the other tiles are empty.

Then, z is a valid configuration of X and:

- By definition of W, z, y coincide on $C(\frac{1}{2}, \mathbf{v})$.
- $\partial\bar{C}$ contains no relation tile, by definition of W and r.
- $C(\frac{1}{2}, \mathbf{u}) + \mathbf{o}$ is empty, see for example Figure 3.20 or Figure 3.22 for an illustration.

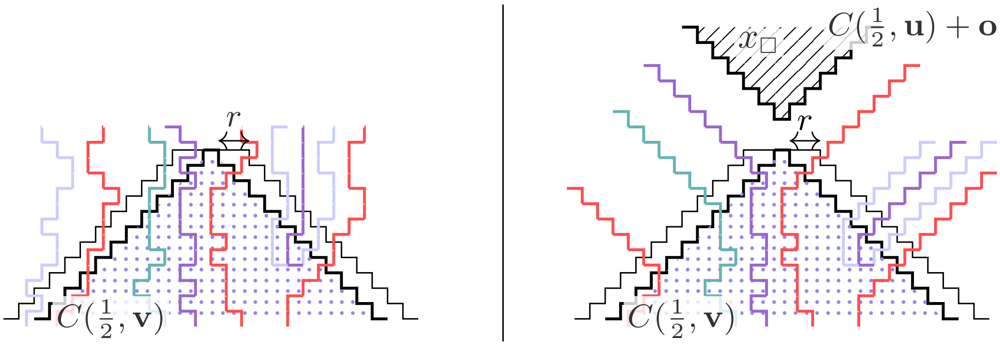

Figure 3.20: Construction of z (on the right) from y (on the left).

$\square$

As an immediate corollary, we get:

Corollary 3.92

X is cone-connected.

More importantly, we obtain the following:

Lemma 3.93: Projective connectedness

X is projectively connected.

Proof. This is an immediate consequence of Proposition 3.49 and Corollary 3.92. $\square$

We can say something a bit more precise on paths and their homotopy classes:

Corollary 3.94: Path Co-extensibility Lemma

Let $p = ((P_t, \mathbf{u}_t))_{t \leq N_p}$ and $q = ((Q_t, \mathbf{v}_t))_{t \leq N_q}$ be two paths with the same aperture window B_n, satisfying:

- Both p and q are coherent paths
- $(P_{N_p}, \mathbf{u}_{N_p}) = (Q_0, \mathbf{v}_0)$ (equivalently, $p * q$ is well-defined)
- $u_0^1 = v_{N_q}^1$ (*i.e.* q ends at the same height as p starts)

Then, there exists p', q', r paths such that:

- r ends on an empty pattern
- $p' * r$ and $r^{-1} * q'$ are well-defined and are both coherent paths.
- $p \sim p'$ and $q \sim q'$

Proof. We may assume that $u_0^1 \leq u_{N_p}^1$, *i.e.* the ending point of p is higher than its starting point, the other case being similar. We can also assume that $u_{N_p}^1$ is the highest point in the entire trajectory of both p and q (we can always homotopically deform p and q so that this is true), and up to some shift, we can assume that $\mathbf{u}_{N_p} = (0,0)$. Consider now $P \subset \mathbb{Z}^2$ so that P contains all the P_t and Q_t. Let x_p, x_q be configurations in which p, q can respectively be traced. Take N large enough so that $P \subset C(\frac{1}{2}, (0,-1)) + (0, N)$. Then, applying the Lemma 3.91 to x_p on one hand, x_q on the other hand, we obtain two configurations $z_p, z_q \in X$. With $\mathbf{o}$ as in Lemma 3.91, let r be the path obtained by moving up to $(0, 2N+1) + \mathbf{o}$ in either z_p or z_q, starting from the origin, which is the same path in both cases. Then r satisfies the conditions of Corollary 3.94. □

3.5.3 A normal form for paths

We are now ready to prove the main lemmas needed to show Theorem 3.81.

Lemma 3.95: No Relation Tile

Let p be a path starting and ending on an empty pattern. Then there exists $p' \sim p$ that does not contain any relation tile.

Proof. As explained above, the proof is by induction on the length of a coherent path decomposition of p. Figure 3.21 is an illustration of the base case, when p is a coherent path.

Let L be the minimal length of a path decomposition of p.

Base case: $\mathbf{L = 1}$ p can be traced entirely in a configuration $x \in X$.

We can assume that x contains a finite number of wires. p being finite, such a configuration exists by the Lemma 3.90. Let $(P_N, \mathbf{v_N})$ be the final point of p. Up to a translation of both p and x we can always assume that p starts at $(0,0)$, and without loss of generality, suppose that $\mathbf{v_N}$ is on the right, *i.e.,* it has a non-negative x-coordinate. This is a legitimate assumption, up to considering the path p^{-1} instead of p, which also starts and ends with empty patterns. Deform p into a path p' in x, whose trajectory only consists of moving right, and then up or down, depending on whether $\mathbf{v_N}$ is above or below $(0,0)$.

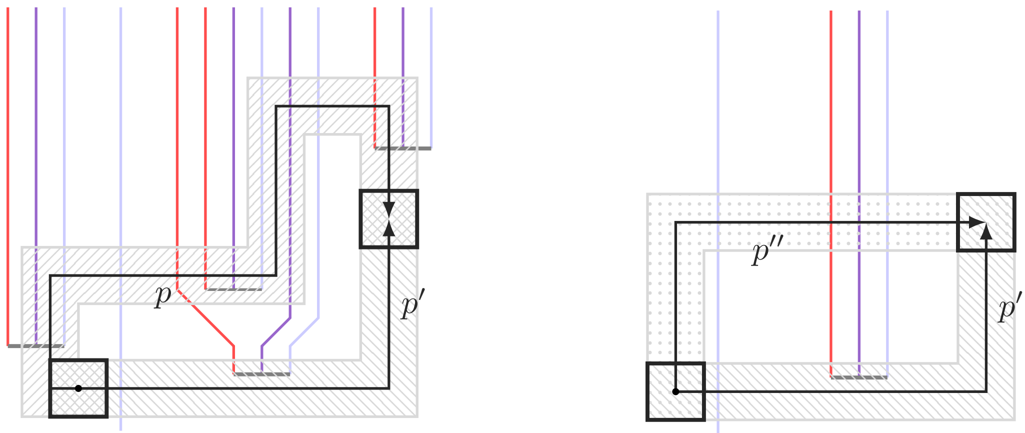

Deformation of p into an L-shaped path p'.

Deformation of p' into p'' to pass above relation tiles.

Figure 3.21: A coherent path deformed so as not to see relation tiles

Let i_{min} (resp. i_{max}) be the leftmost (resp. rightmost) position of a relation tile in x, the topmost one. We can deform p' as follows:

- Move left until the position $i_{\text{min}} - 2n$ (or don't move if $i_{\text{min}} - n \geq 0$).
- Move up until the position $j + 2n$
- Move right until $i_{\text{max}} + 2n$
- Finally, move to $\mathbf{v_N}$, by moving vertically first and then horizontally.

Let p'' be the resulting path. Then, p'' does not see any relation tile. Figure 3.21 shows this process in a simple case, with the first and third steps being trivial, and how deforming p into p' simplifies the analysis by bounding the positions of the possible relation tiles seen by p', that p'' can then avoid.

Base case: $\mathbf{L = 2}$ $\quad p = p_1 * p_2$

Let $(P_t, \mathbf{v_t})$ be the endpoint of p_1 and the starting point of p_2, with $\mathbf{v_t} = (v_t^0, v_t^1)$. Suppose that $v_t^0 \geq 0, v_t^1 \geq 0$. Let $\mathbf{v}_N$ be the $\mathbb{Z}^2$ point at which p ends – by assumption, the associated pattern P_N is only made of empty tiles. Let $x_1, x_2 \in X$ be two configurations such that p_1, p_2 can respectively be traced entirely within them, and containing a finite number of wires using the Lemma 3.90.

In order to be able to use the previous case $L = 1$, we modify the path as follows: consider the path q, traced in x_2, that:

- starts from $(P_t, \mathbf{v_t})$
- follows the inverse trajectory to p_1
- upon reaching $(0,0)$, continues horizontally until it sees an empty pattern (which always eventually happens, as x_2 contains a finite number of wires)

Let $p'_1 = p_1 * q$ be and let $p'_2 = q^{-1} * p_2$, so that $p = p'_1 * p'_2$. By construction, p'_2 can be traced entirely within x_2, and so can be appropriately deformed according to the case $L = 1$.

Like p, p'_1 has a decomposition of length 2, but we can further simplify it. Indeed, using the Corollary 3.94, we obtain a loop $r = r_1 * r_1^{-1}$, based at $(P_t, \mathbf{v_t})$, such that r_1

ends in an empty pattern and each of $p_1 * r_1$ and $r_1^{-1} * q^{-1}$ can be traced within a single configuration. This is enough to prove the case $L = 2$, using three times the case $L = 1$.

The construction is shown in Figure 3.22.

Finally, we have that

$$p \sim_{B_n} \underbrace{p_1 * r_1}_{\text{coherent}} * \underbrace{r_1^{-1} * q}_{\text{coherent}} * \underbrace{q^{-1} * p_2}_{\text{traced in } x_2}$$

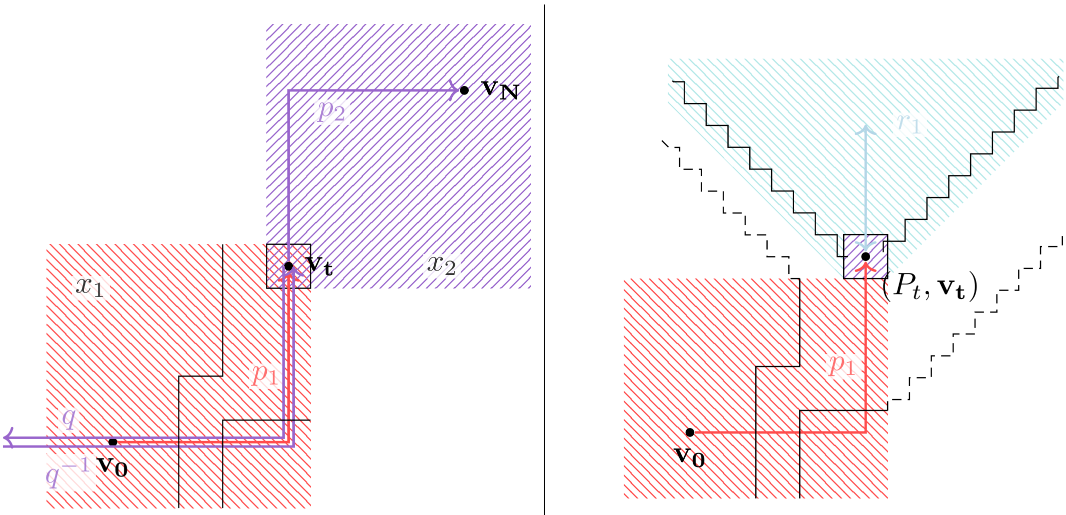

After inserting q and q^{-1} into $p = p_1 * p_2$ | The x_1' configuration.

Figure 3.22: Red paths are traced in x_1, purple ones in x_2. Wires are drawn in black.

General case: $\mathbf{L > 2}$ $p = p_1 * \ldots * p_L$.

Consider the timestep t at which p_1 ends and p_2 starts. By definition of a coherent decomposition, there exists $x_2 \in X$ such that p_2 can be entirely traced within x_2. Using the Lemma 3.90, we can suppose that x_2 contains finitely many wires. Consider a loop $r = r_1 * r_1^{-1}$ that moves to an empty pattern in x_2 by moving left (this is always possible according to Lemma 3.88) and then comes back. We have

$$p = p_1 * p_2 \ldots * p_L = \underbrace{p_1 * r_1}_{p_1'} * \underbrace{r_1^{-1} * p_2 \ldots * p_L}_{p'}$$

p_1' and p' are then respectively paths of length 2 and $L - 1$, and so using the induction hypothesis, they can be deformed so at to avoid any relation tile. $\square$

Lemma 3.96: Single Wire

Let $p = (P_i, \mathbf{v_i})_{0 \le i \le N}$ be a path starting and ending with empty patterns. There exists a path p', homotopic to p, such that the union of any two consecutive patterns in p' contains at most a single wire.

Proof. The result is also proved by induction on the length L of a path decomposition of p. As for the Lemma 3.95, we illustrate in Figure 3.23 the case where p is itself coherent.

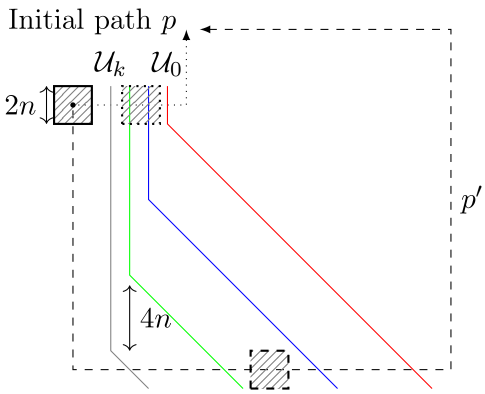

Figure 3.23: Deformation of p into p' in a single configuration to see only one wire per pattern.

Base case: $\mathbf{L} = \mathbf{1}$ p can be traced entirely in a configuration $x \in X$. Using the Lemma 3.95, we may assume that p does not see any relation tile. Without loss of generality, we may assume that x does not contain any wire that is not seen by p and that p starts at $(0, 0)$ and ends at $\mathbf{v_N} = (v_N^0, v_N^1)$, with $v_N^0 \geq 0, v_N^1 \geq 0$. For simplicity, we assume that the trajectory is made out of two straight segments, so that p first moves horizontally from $(0, 0)$ to $(v_N^0, 0)$ and then vertically to $\mathbf{v_N}$. Let $\mathcal{U}_0, \ldots, \mathcal{U}_k$ be the wires seen from right to left by p (so p sees $\mathcal{U}_k$ first, then $\mathcal{U}_{k-1}$ and so on until $\mathcal{U}_0$).

Now consider a configuration x' satisfying (see Figure 3.23):

- x' does not contain any other wire than the $\mathcal{U}_i$'s
- for $0 \leq i \leq k$, let $(z_i, -n)$ be the position of the only tile of $\mathcal{U}_i$ whose second coordinate is $-n$, and whose wire enters it from its bottom edge. Then, for $-n - 4ik \leq z \leq -n$, we define $x'(z_i, z)$ to be a tile of the form □, and all the tiles of $\mathcal{U}_i$ below that are of the form □ and □. This uniquely determines all the $\mathcal{U}_i$'s below p.

For $z \in \mathbb{Z}$, no pattern of support B_n centered at $(z, -4n(k+1))$ can see tiles belonging to two different wires at the same time in x'. Therefore, we can deform p in x' into p', where p' starts by moving down for $4n(k+1)$ steps, then right until crossing $\mathcal{U}_0$, and finally up and either right or left as needed to reach $\mathbf{v_N}$. Any two consecutive patterns on this path see at most one wire.

Base case: $\mathbf{L} = \mathbf{2}$ The proof works in exactly the same way as in the proof of the Lemma 3.95.

General case: $\mathbf{L} > \mathbf{2}$ $p = p_1 * \ldots * p_L$.

As before, consider the timestep t at which p_1 ends and p_2 starts. As p_2 is coherent, there exists $x_2 \in X$ such that p_2 can be entirely traced within x_2, and we can assume that x_2 contains finitely many wires. Let r_1 be any path that reaches to an empty pattern in x_2 by moving horizontally left (this always eventually happens, according to Lemma 3.88). We have

$$p = p_1 * p_2 \ldots * p_L = \underbrace{p_1 * r_1}_{p_1'} * \underbrace{r_1^{-1} * p_2 \ldots * p_L}_{p'}$$

p_1' and p' are respectively paths of length 2 and $L - 1$, and the induction hypothesis ensures that they can be homotopically deformed so at not to see $\mathcal{U}$. The resulting path then only sees one wire at a time. □

Lemma 3.97: No Uncrossed Wire

Let p be a path starting and ending with empty patterns, and $\mathcal{U}$ some wire seen but not crossed by p. There exists a path p', homotopic to p, which does not see $\mathcal{U}$.

Proof. The idea is that using the previous Lemma 3.96, we can deal with each wire independently. In particular, the uncrossed wire $\mathcal{U}$ is the only wire seen by some subpath p' of p, and is not seen by p neither before nor after p'. Hence, it suffices to show the result for paths seeing a single wire overall. In that case, one observes that $\mathcal{U}$ has to stay in the same "side" of the aperture window along p', that can therefore be deformed without crossing $\mathcal{U}$ by moving sufficiently far in the opposite direction.

We proceed by induction on the length L of a coherent decomposition of the path, and we assume that $\mathcal{U}$ is on the right side of the patterns. Using the Lemma 3.96, we can assume that all the patterns of p contain at most a single wire.

Base case: $\mathbf{L = 1}$ p can be traced entirely in a configuration $x \in X$.

In that case, we can simply deform p in x by changing its trajectory so that it always stays more than n units left from $\mathcal{U}$. This path can then be traced in the configuration x', equal to x except for the tiles of $\mathcal{U}$ in x that are empty tiles in x'.

Base case: $\mathbf{L = 2}$ $p = p_1 * p_2$

Let $(P_t, \mathbf{v_t})$ be the final point of p_1 and the first one of p_2. We also assume that the second coordinate of $\mathbf{v_t} = (v_t^0, v_t^1)$ is non-negative. Let $\mathbf{v_N} = (v_N^0, v_N^1)$ be the final point of the path.

Let $x_1 \in X$ (resp. x_2) be a configuration, containing a minimal number of wires (which exists according to the Lemma 3.90), such that p_1 (resp. p_2) can entirely be traced within it. Let $\mathcal{U}$ be the uncrossed wire. We can always assume that $\mathcal{U}$ appears in P_t, otherwise, we could consider p_1 and p_2 separately and apply twice the case $L = 1$.

We deform p_1 into p_1' inside x_1:

- Starting from $(0, 0)$, it first moves to the right, until $\mathcal{U}$ appears on the central tile of the pattern seen by p_1.
- It then moves up, left or right, following $\mathcal{U}$: up if the central tile is $\boxed{\uparrow}$, left then up if it is $\boxed{\Lsh}$, and so on.
- Finally, once it attains the height v_t^1, it moves left until $\mathbf{v_t}$ if needed, which takes at most n steps.

We can also deform p_2 into another path p_2' as follows:

- Starting from $\mathbf{v_t}$, move left for $\max(2n, (v_t^0 - v_N^0))$ steps. This ensures that we are far enough so as to not see $\mathcal{U}$ anymore.
- Then, move vertically to height v_N^1.
- Finally, move right until $\mathbf{v_N}$.

Let $\mathbf{w_1}$ be the last point of p_1' before seeing $\mathcal{U}$, and $\mathbf{w_2}$ the first point of p_2' after having seen $\mathcal{U}$ for the last time. The Lemma 3.96 ensures that the patterns seen at both $\mathbf{w}_1$ and $\mathbf{w}_2$ are empty. This gives a decomposition

$$p \sim p_1' * p_2' \sim p_{\text{start}} * p_{\mathcal{U}} * p_{\text{end}}$$

where p_{start} ends at $\mathbf{w_1}$, $p_{\mathcal{U}}$ is the part of the path between $\mathbf{w_1}$ and $\mathbf{w_2}$, and p_{end} starts at $\mathbf{w_2}$.

$p_{\mathcal{U}}$ can be traced entirely in a configuration x_3 whose only wire is $\mathcal{U}$. In this configuration, it can be homotopically deformed to $p'_{\mathcal{U}}$ which never sees $\mathcal{U}$ according to the case $n = 1$.

The final path $p' = p_{\text{start}} * p'_{\mathcal{U}} * p_{\text{end}}$ does not see $\mathcal{U}$.

General case: $\mathbf{L > 2}$ $p = p_1 * \ldots * p_L$

In that case, the proof is exactly the same as in the Lemma 3.95 and the Lemma 3.96: we insert a loop before p_2 starts that extends it, and from a decomposition of length L we obtain two decompositions of length respectively 2 and $L-1$, which are solved inductively. □

Lemma 3.98: Cross Anywhere

Let p be a path starting and ending with empty patterns. If p sees no relation tiles, but sees and crosses a single wire $\mathcal{U}$ exactly once, then for all $\mathbf{v} = (v^0, v^1) \in \mathbb{Z}^2$, p is homotopic to a path p' which crosses $\mathcal{U}$ exactly on $\mathbf{v}$.

Proof. The idea is that if $\mathcal{U}$ exits the aperture window B_n of p in position $(i,j) \in \mathbb{Z}^2$, it can be extended using tiles and , or and , to pass anywhere inside $(i,j) + C(\frac{1}{2}, (0,1))$ or $(i,j) + C(\frac{1}{2}, (0,-1))$ (see Lemma 3.91 for an explanation of where those cones come from). The path p can then be deformed to cross it anywhere in those two cones. Using several such deformations, we can deform p so that it crosses $\mathcal{U}$ anywhere in the plane. Note that even if p is initially coherent, it might happen that p' is not, depending on $\mathbf{v}$ and where p initially crossed $\mathcal{U}$.

Let $p = (P_i, \mathbf{v_i})_{0 \le i \le N}$ be such a path, and let t be the timestep at which p crosses $\mathcal{U}$. Without loss of generality, we can then assume that the wire is crossed from left to right, *i.e.* $\mathcal{U}$ is on the right side of P_{t-1} and on the left side of P_t.

Let x be any configuration containing $P_{t-1} \cup P_t$. We can suppose that $\mathbf{v_t} = e_0 + \mathbf{v_{t-1}}$, by deforming p in x if needed, and that $\mathbf{v_{t-1}} = (0,0)$. Let r_1 be the path starting from $(P_{t-1}, (0,0))$ which moves left for $4n + 2|v^0|$ steps in x, and let $r = r_1 * r_1^{-1}$. Let q_1 be the path starting from $(P_t, (1,0))$ which moves right for $4n + 2|v^0|$ steps in x, and let $q = q_1 * q_1^{-1}$.

We can deform p in x by inserting the loops r and q respectively at the timesteps $t-1$ and t. Using the the Lemma 3.97 twice, this path can itself be deformed into $p_{start} * p' * p_{end}$ with $p' = r_1^{-1} * (P_t, (0,0)) * q_1$, and p_{start}, p_{end} paths that only see empty patterns. The trajectory of p' is a straight horizontal line on the $x-$axis of length $8n + 2|v^0| + 1$.

Let x' be the configuration obtained by extending $\mathcal{U}$ as seen by p' using only tiles of the form . Without loss of generality, suppose that $v^1 \le 0$. We can deform p' in x' so that it moves up for $8n + 2|v^0|$ steps, then right for $8n + 2|v^0| + 1$ as before and finally down to the endpoint of p'. Call p'' the horizontal part of this path. There exists a configuration x'' in which $\mathcal{U}$ passes by $\mathbf{v}$ and in which p' can be traced. Then, p'' can be deformed in x'' to cross $\mathcal{U}$ on $\mathbf{v}$. This finally gives the result. □

3.5.4 Computing the projective fundamental group

We can now compute $\pi_1^{proj}(X)$, which is independent of the basepoint since X is projectively connected. Hence, unless stated otherwise, all the loops in this proof are based at $(x_\square, (0,0))$. With any such loop p, we associate a word $[\![p]\!]$ on the alphabet $\bar{S}$ in the following way, illustrated in Figure 3.24:

- If p does not cross any wire, we associate the empty word with it, $[\![p]\!] = \varepsilon$.
- If p crosses a single wire $\mathcal{U}$, then:
  - If $\mathcal{U}$ is not a horizontal wire found on a relation tile, and $\mathsf{s} \in \bar{S}$ is the generator corresponding to $\mathcal{U}$ (see Section 3.5.1)
    * if p crosses it from left to right, or from top to bottom on a tile shaped as ◳, or from bottom to top on a tile ◱, then $[\![p]\!] = \mathsf{s} \in \bar{S}$.
    * if p crosses it in any other direction, we set $[\![p]\!] = \mathsf{s}^{-1} \in \bar{S}$
  - Otherwise, $\mathcal{U}$ is a horizontal wire on a relation tile. Let $\overline{\mathsf{R}_\mathsf{i}} = \overline{\mathsf{r}_0 \ldots \mathsf{r}_\mathsf{i}}$ be its colour.
    * If it is crossed from top to bottom, then $[\![p]\!] = \mathsf{r}_\mathsf{i}^{-1} \ldots \mathsf{r}_0^{-1} \in \bar{S}^*$
    * Otherwise, $[\![p]\!] = \mathsf{R}_\mathsf{i} = \mathsf{r}_0 \ldots \mathsf{r}_\mathsf{i}$
- If $p = p_1 * p_2$, then $[\![p]\!] = [\![p_1]\!] \cdot [\![p_2]\!] \in \bar{S}^*$ where $\cdot$ represents the concatenation in $\bar{S}^*$.

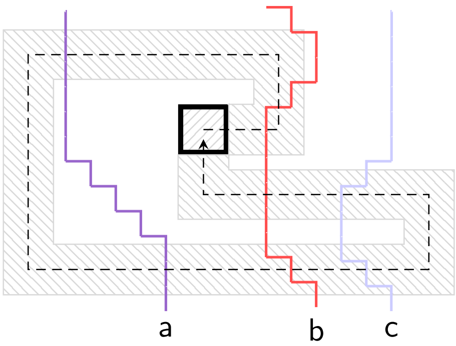

The word associated with this loop is $\mathsf{bb^{-1}a^{-1}abcc^{-1}b^{-1}} =_G 1_G$.

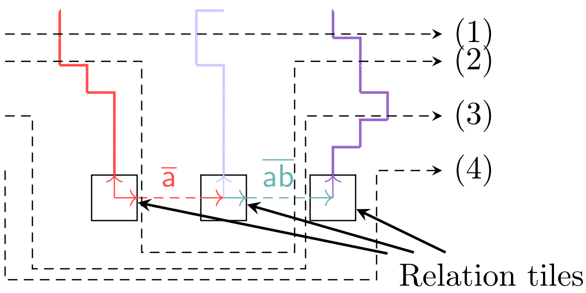

Widget for the relator $abc = 1_G$. From top to bottom, the words associated with the paths (1) to (4) are respectively $abc = 1_G, aa^{-1}(ab)c = 1_G, (ab)c = 1_G$ and 1_G. For clarity, the relation tiles are not adjacent on the figure.

Figure 3.24: Examples of words associated to coherent paths.

As in Section 1.3.1, for any two words w, w' on $\bar{S}$, we write $\mathsf{w} \equiv \mathsf{w}'$ if they are equal as words on this alphabet, and $\mathsf{w} =_G \mathsf{w}'$ if they represent the same element of the group G.

In order to prove that the projective fundamental group of this subshift is G, we will prove that the operation $[\![p]\!]$ entirely characterizes a loop up to homotopy, in the sense that loops associated with the same element of G are exactly a projective loop-class:

Lemma 3.99: Homotopic Implies Equal

For $n > 0$ and any two loops p_n, p'_n starting at $\left(x_{\square|B_n}, (0,0)\right)$,

$$p_n \sim_{B_n} p'_n \implies [\![p_n]\!] =_G [\![p'_n]\!]$$

Proof. As any two homotopic loops can be obtained from one another by a sequence of elementary deformations, we can restrict ourselves to the special case of a single deformation that is a loop based at $(P_t, \mathbf{v_t})$. By definition, this deformation is made in a single configuration $x \in X$. We consider two disjoint cases, according to the presence of relation tiles in x.

- Suppose that x does not contain any relation tile. Any bi-infinite wire splits the space in two disjoint regions (a "left" one and a "right" one). Each time a loop crosses such a wire, it has to cross it in the other direction to come back to its initial region. Because wires do not intersect, the associated word will be some kind of Dyck word, where each $\mathsf{s} \in \bar{S}$ can act as an opening or a closing bracket (in which case, the associated closing (resp. opening) bracket is s^{-1}), so it is clearly equal to 1_G in G. This is the simple case depicted in Figure 3.24.

- Now, suppose that x does contain some relation tiles. In this case, notice that any two relation tiles are either part of the same relator and are therefore linked by a finite sequence of horizontal relation tiles, or they are independent (not linked by any wire).

  Hence, we can consider each one of those patterns separately. Consider such a pattern, with relation tiles that implement a relator $\mathsf{r} = \mathsf{r_0} \dots \mathsf{r_k} \in R$, and a configuration x' that only contains this pattern. Figure 3.24 represents this in a configuration corresponding to relation $\mathsf{abc} = 1$.

  We show that, due to how $[\![\cdot]\!]$ has been defined, all the homotopy-equivalent paths in x' are associated with the same element of G. Let $\mathcal{U}_0, \dots, \mathcal{U}_k$ be the wires corresponding respectively to $\mathsf{r_0}, \dots, \mathsf{r_k}$, and suppose that the relation tiles in x' are placed on $(0,0), \dots, (k,0)$. We will show that for any p joining $(0,0)$ to $(k+1,0)$ in x', $[\![p]\!] =_G 1_G$. Let $\mathcal{R} \subset \mathbb{Z}^2$ be the set of points above the $(\mathbb{Z}, 1)$ line and between $\mathcal{U}_0$ and $\mathcal{U}_k$. We can always suppose that no wire is crossed consecutively in opposite directions, as the word associated to a path that crosses a wire in a direction and immediately crosses it in the other direction is $\mathsf{ss}^{-1} =_G 1_G$ for some $\mathsf{s} \in \bar{S}^*$. We can also suppose that p only enters and then leaves $\mathcal{R}$ once. Otherwise, we can simply split it into several such paths and prove the claim for each of them independently.

  - If p crosses $\mathcal{U}_0, \dots, \mathcal{U}_k$, then $[\![p]\!] \equiv \mathsf{r_0} \dots \mathsf{r_k} =_G 1_G$ by definition.
  - If p crosses $\mathcal{U}_0, \dots, \mathcal{U}_i, \mathcal{U}_{\overline{\mathsf{r_0} \dots \mathsf{r_i}}}$, where $\mathcal{U}_{\overline{\mathsf{r_0} \dots \mathsf{r_i}}}$ is a wire of a relation tile which is necessarily crossed from top to bottom, by definition, $[\![p]\!] \equiv \mathsf{r_0} \dots \mathsf{r_i}(\mathsf{r_i}^{-1} \dots \mathsf{r_0}^{-1}) =_G 1_G$
  - Otherwise, p crosses $\mathcal{U}_{\overline{\mathsf{r_0} \dots \mathsf{r_i}}}, \mathcal{U}_{i+1}, \dots, \mathcal{U}_j, \mathcal{U}_{\overline{\mathsf{r_0} \dots \mathsf{r_j}}}$, the first relation tile being crossed from bottom to top to enter $\mathcal{R}$ and the last one being crossed from top to bottom to exit it. By definition, $[\![p]\!] \equiv (\mathsf{r_0} \dots \mathsf{r_i})\mathsf{r_{i+1}} \dots \mathsf{r_j}(\mathsf{r_j}^{-1} \dots \mathsf{r_0}^{-1}) =_G 1_G$

This shows that all the paths traced in a single configuration are associated with the same element of the group G. As all homotopies are deformations in a given configuration, this implies that for any homotopically equivalent paths p, p', we have $[\![p]\!] =_G [\![p']\!]$. ☐

Lemma 3.100: Equal Implies Homotopic

For any window B_n, and for any pair of loops p_n, p'_n starting at $(x_{\square|B_n}, (0,0))$,

$$[\![p_n]\!] =_G [\![p'_n]\!] \implies p_n \sim_{B_n} p'_n.$$

Proof. Using the Lemma 3.95, we can always start by deforming p_n and p'_n so that they do not see any relation tile. As each elementary deformation is by definition occurring in some given configuration, Lemma 3.99 ensures that we still have $[\![p_n]\!] =_G [\![p'_n]\!]$. We will first prove that $[\![p_n]\!] \equiv [\![p'_n]\!] \implies p_n \sim_{B_n} p'_n$, which is a stronger assumption. Next, we prove that given p_n and p'_n with $[\![p_n]\!] =_G [\![p'_n]\!]$, there exists a loop p''_n such that $p_n \sim_{B_n} p''_n$ and $[\![p''_n]\!] \equiv [\![p'_n]\!]$. We then have that $p''_n \sim_{B_n} p'_n$ according to the first part of the proof, and so $p_n \sim_{B_n} p'_n$.

- We show that $[\![p_n]\!] \equiv [\![p'_n]\!] \implies p_n \sim_{B_n} p'_n$. The paths p_n and p'_n can be deformed using the Lemma 3.97 so that they cross all the wires that they see. The Lemma 3.96 can then be used to deform them so that there is at most one of those wires per pattern. Let $\hat{p_n}$ and $\hat{p'_n}$ be the resulting paths, which by assumption cross the same wires. Using the Lemma 3.98 for each of those crossed wires, we can finally deform $\hat{p_n}$ into $\hat{p'_n}$, and so $p_n \sim_{B_n} p'_n$.

- Now, we show the existence of a loop p''_n satisfying $p_n \sim_{B_n} p''_n$ and $[\![p''_n]\!] \equiv [\![p'_n]\!]$. By definition of $=_G$, there exists a finite sequence $(\mathsf{u_i})_{0 \leq i \leq N}$ of words on the alphabet $\bar{S}$ such that $[\![p_n]\!] \equiv \mathsf{u_0}$, $[\![p'_n]\!] \equiv \mathsf{u_N}$, and for all $i < N$, $\mathsf{u_i} \leftrightarrow \mathsf{u_{i+1}}$. To prove the result, it is therefore enough to show that for any word v such that $[\![p_n]\!] \leftrightarrow \mathsf{v}$, we can deform p_n in another loop p^{v}_n such that $[\![p^{\mathsf{v}}_n]\!] \equiv \mathsf{v}$.

  Suppose that v is obtained from $[\![p_n]\!]$ by deleting a relator. More formally, there exists words $\mathsf{u_1}, \mathsf{u_2}$ and a relator $\mathsf{r} \in R$ such that $\mathsf{v} \equiv \mathsf{u_1 u_2}$ and $[\![p_n]\!] \equiv \mathsf{u_1 r u_2}$. Using the Lemma 3.96 followed by the Lemma 3.97, we obtain a loop $q \sim p_n$, such that q crosses exactly wires of the same type as p_n, but it only ever sees one wire at a time, and crosses all the wires that it sees. The Lemma 3.98 then ensures that we can deform q into a loop that crosses wires corresponding to the letters of $\mathsf{u_1 r u_2}$, in order, on a horizontal line. Let $p_{\mathsf{u_1}}$ (resp. $p_{\mathsf{r}}, p_{\mathsf{u_2}}$) be the part of this path which crosses the wires corresponding to $\mathsf{u_1}$ (resp. $\mathsf{r}, \mathsf{u_2}$), starting and ending with empty patterns. Let $x_{\mathsf{r}} \in X$ be such that p_{r} can be traced in x_{r}, and in which all those wires originate from the same set of relation tiles (see Figure 3.24). We can then deform p_{r} in x_{r} into a path p'_{r} that passes below the relation tiles. The resulting path $p^{\mathsf{v}}_n = p_{\mathsf{u_1}} * p'_{\mathsf{r}} * p_{\mathsf{u_2}}$ is then a solution.

$\square$

We can finally prove the main result of this chapter, Theorem 3.81:

Proof of Theorem 3.81. Let $n > 0$ and let $\Phi_n \colon p \in \pi_1^{B_n}(X, (x_\square, (0,0))) \mapsto [\![p]\!] \in G$ be the function which associates with a loop-class with aperture window B_n the corresponding element of G. The Lemma 3.99 and Lemma 3.100 show that it is well-defined and injective. Let $[p], [p']$ be two loop-classes based at $(x_{\square|B_n}, (0,0))$. We have shown that $[p] \sim_{B_n} [p'] \iff \Phi_n([p]) =_G \Phi_n([p'])$. Now notice that $\Phi_n([p * p']) =_G \Phi_n(p) \cdot_G \Phi_n(p')$, *i.e.*, Φ_n is a group morphism. To show that it is surjective, let $g \in G$ any element, and $u_1 \ldots u_n \in \bar{S}^*$ such that $u_1 \ldots u_\ell =_G g$. Let x^g the following configuration:

- For $1 \leq i \leq \ell$ and $j \in \mathbb{Z}$, $x^g(i,j)$ is a tile of type $\boxed{\uparrow}$ and of colour u_i
- Otherwise, $x^g_{(i,j)} = \square$

Now, consider the following loop: define p_n as the loop based at $(x_{\square|B_n}, (0,0))$, which:

- moves left for n steps in $x_\square$
- moves right for $2n + \ell$ steps in x^g – at this point, it sees an empty pattern, after having crossed all the wires of x^g

- comes back to $(0,0)$ in $x_\square$.

By definition, $[\![p_n]\!] \equiv u_1 \dots u_n =_G g$.

Furthermore, notice that for any loop-class $[p_{n+1}]$ based at $(x_{\square|B_{n+1}}, (0,0))$, if p_{n+1} projects down to p then $\Phi_{n+1}([p_{n+1}]) =_G \Phi_n([p])$. This shows that $\pi_1^{proj}(X, (x_\square, (0,0)))$ is isomorphic to G, and the final result follows from the fact that X is projectively connected. □

3.5.5 Open questions: beyond finitely presented groups

Besides the open questions about projective connectedness presented in Section 3.3, we also give a few problems that we have not successfully solved, regarding the possible fundamental groups of SFTs and sofic subshifts.

Infinitely generated groups

In this chapter, we only presented groups whose projective fundamental group was finitely generated. This is not always the case, and we do not give any characterization of when this holds. For an example of such a subshift, consider the one-dimensional subshift X on $\{\square, \blacksquare\}$, whose only forbidden pattern is $\{\blacksquare\square\}$. Configurations of X are then of the form ${}^\infty\square\blacksquare^\infty$, and so the configurations of $X^\uparrow$ are either monochromatic or contain two monochromatic half-planes. As $X^\uparrow$ is easily seen to be cone-connected, it is also projectively connected by Proposition 3.49. To see that it is infinitely generated, consider any path p_i crossing the "frontier" between the two half-planes at some point $(i,j) \in \mathbb{Z}^2$. Such a path can only ever be deformed to cross this frontier $\square\blacksquare^\uparrow$ in the same horizontal column $\{i\} \times \mathbb{Z}$. We can also explicitly construct projective loop class containing all those paths, which must therefore all be distinct. In particular, $\pi_1^{proj}(X^\uparrow)$ is infinitely generated.

Recursively presented groups

Using Theorem 1.87, one might expect to realize *recursively presented* groups as fundamental groups of sofic subshifts. Given a recursively but not finitely presented group $G = \langle S \mid R \rangle$, one can indeed try to define a subshift X_G as in Section 3.5.1, where the relations are not directly imposed using the matching rules of an SFT but are enforced by a Turing Machine enumerating R. We could even hope to simplify all the proofs of Section 3.5.2 and Section 3.5.3: as we are defining a sofic subshift, we can easily enforce that at most one sequence of relation tiles appears in any given configuration, using *e.g.* a sunny-side-up layer. However, as the relations of R can be arbitrarily long (and must in fact be, for otherwise G would be finitely presented), we need to make sure that we cannot have "limit configurations", corresponding to meaningless relations, appearing by compactness of the subshift. Although similar difficulties are routinely encountered in the literature when considering subshifts with specific recursive properties, all the techniques known by the author to solve this issue fail here, mainly because they all drastically change the value of the fundamental group, an object which is more geometric than recursive in nature.

Chapter 4

Substitutive subshifts on graphs

In this chapter, we take a step back from the tilings of the grid that we have studied so far and from purely computational aspects of subshifts. The goal of this chapter is to propose a formalism which encompasses several similar results of the literature, that we call the **Mozes theorem**:

Meta-Theorem : Meta Mozes Theorem

Substitutive subshifts are sofic.

Thus stated, the theorem is only very loosely specified and informal. Even within the specific setting of tiling of $\mathbb{Z}^2$, several definitions of "substitutive subshifts" exist, and on more exotic spaces, the definition of a sofic subshift might not be as clear as what was done in Section 1.1.4. Nonetheless, we will see that this informal statement can be made precise and holds in a number of cases: Section 4.2.1 presents the original and simplest case of $\mathbb{Z}^d$ tilings as proven in [Moz89], which in particular explains why we refer to Chapter 4 as "Mozes" theorem, Section 4.2.2 focuses on the $\mathbb{R}^d$ case studied by [Goo98] or [FO10], and more combinatorial structures will be considered in Section 4.2.3, for which results have already been obtained for example in [BS16] and [Sil20]. This section will try to show that despite this diversity, the arguments used to prove the "soficness" of the relevant subshifts are often very similar, and rely on a kind of fixpoint argument. More precisely, the general idea, initiated by Mozes in the case of $\mathbb{Z}^2$ tilings [Moz89], is based on the following ideas:

- One can associate to a substitutive tiling a hierarchical, inductive structure, which partitions the tiles into larger and larger "meta-tiles", corresponding to successive applications of the substitution rule.
- In order to ensure that tilings by the considered tileset do indeed respect this substitutive structure, it is therefore enough to ensure that each tile belongs to a meta-tile, itself belonging to a larger meta-tile and so on.
- The key observation is that it is enough for each tile to remember a finite amount of information, locating it in this hierarchical structure. Part of this information is used to ensure the consistency of a given meta-tile, and part of it is used to guarantee it at the "next step" of the hierarchy. With some additional bookkeeping, this is enough to guarantee, in an inductive-like fashion, the consistency of the entire tiling.

The major technical difficulty therefore lies in ensuring that all the necessary information is communicated and synchronized, both between meta-tiles of a same level of the hierarchical structure, and to the meta-tiles of higher level. This restriction can be formalized as a combinatorial condition, and can be studied without actually looking at the geometry of the tiling, when embedded in $\mathbb{Z}^d$ or $\mathbb{R}^d$. For this reason, we introduce in Section 4.2.3 a notion of **substitutive graph**, and prove that under some conditions on the connectivity of the graph, Mozes theorem holds in this combinatorial setting: the key idea is to show that "nice" graphs are connected-enough so that the problem highlighted in the above paragraph, synchronizing the information of all the tiles, can be solved. We will say that a substitution has the **Mozes property** if the meta-Mozes theorem holds for this substitution:

Meta-Definition : Mozes property

A substitution $\mathfrak{s}$ satisfies the **Mozes property** if the meta Mozes theorem holds for $\mathfrak{s}$-substitutive subshifts.

The goal of this chapter is then to find some sufficient combinatorial conditions for substitutions to satisfy the Mozes property. One of the main motivations is to provide a framework that recovers at least some of the results of the literature, but which could also be generalized to tilings for which no general satisfactory notion of "substitution" exists, the main example being graphs of algebraic origin such as Cayley or Schreier graphs. In particular, we would like to find a general way to represent some classes of finitely presented groups in our setting, and although we do not give any general class of examples, this still motivates our definitions and explains some of the choices made when defining our notion of combinatorial substitution. This is still an exploratory idea, and the results and constructions of this chapter should be understood as a step in this general direction of research.

This chapter tries to be consistent with the terminology used in [BS16]. It is organised as follows:

- Section 4.1 tries to present the main ideas behind the numerous notions of substitutions that exist in the literature.
- Section 4.2 is a review of those various definitions, with an emphasis put on the ones for which an equivalent to Chapter 4 is known.
- Section 4.3 defines graph subshifts following [ADG23], proves some properties of those subshifts, and proposes a definition of *substitutive* graph subshifts.
- Finally, Section 4.4 proves the main result of this section, which is a partial Mozes theorem for the case of substitutive graph subshifts:

Theorem: Mozes theorem - graphs

Let $\mathfrak{s}$ be a graph substitution, and $\mathfrak{s}_c$ a coloured $\mathfrak{s}$-substitution. Suppose that $\mathfrak{s}$ is quasi-connected. Then, there exists a sofic graph subshift $Y_{\mathfrak{s}_c}$ which is $X^{\infty}_{\mathfrak{s}_c}$-sheeted and contains $X^{\infty}_{\mathfrak{s}_c}$.

The notion of of "sheeted graph" is presented in Section 4.3.4, and is necessary in our case, as we consider tilings on graphs which do not belong to a larger ambient space such as $\mathbb{Z}^d$ or $\mathbb{R}^d$, and so this assumption can be dropped when considering the usual tilings on a given, fixed space. The quasi-connected condition mirrors the usual required conditions on substitutions to satisfy the Mozes property, as highlighted in Section 4.2.

4.1 Substitutions

The modern study of substitutive, or self-similar structures, originates in the work of Berger [Ber66] who proved that the Domino Problem was undecidable: the proof strategy was to show that Wang's Algorithm had no chance to terminate, as seen in Section 1.2.3, because there exists aperiodic tilings. The construction of this aperiodic tiling relies on enforcing *via* local rules a hierarchical structure, similar to what is now called a substitutive tiling. A general survey of properties of self-similar tilings in various settings can be found in [AA20], and in particular the chapters 2.1 to 2.4 of this book are especially relevant to our approach here. We give the main ideas underlying how substitutions are to be thought of in all those settings (discrete grids, euclidean planes, Cayley graphs ...), and will state more formal definitions in the next sections.

The main idea behind a substitution is to replace each tile $t \in \tau$ by a larger **patch** $\mathfrak{s}(t)$, that is, a set of tiles. Starting from one tile, we can then repeat this replacement process by successively replacing each tile t' from each path $\mathfrak{s}(t)$ by its corresponding path $\mathfrak{s}(t')$. The difficulty in this approach lies in deciding how patches $\mathfrak{s}(t_1), \mathfrak{s}(t_2)$ coming from adjacent tiles t_1, t_2 should be "merged", or "glued" together. This is sometimes referred to as the **compatibility problem**, and is studied for example from a categorical standpoint in [Fer22] and [MS15], or as a decision problem in [JK12].

Another point-of-view, preferred for example in [Bar22], is to define the image $\mathfrak{s}(t)$ of a tile t by partitioning t using smaller versions of the tiles of τ. This process can be iterated by inflating the patch obtained at each step, and subdividing each tile in the way specified by $\mathfrak{s}$.

Both points of view can informally be summarized as follows:

Meta-Definition : Meta-substitution

A **substitution** on a set of tiles τ and a structure set G, consists of:

- A map $\mathfrak{s}\colon \tau \to \tau^{\mathcal{P}(G)}$, sending a tile $t \in \tau$ to a **patch**, that is, a set of tiles along with some structure given by a part of G (*e.g.* a discrete rectangle, a finite graph, a topological disk)
- An iteration rule (or gluing rule, or induction rule ...), which extends $\mathfrak{s}$ to a map $\mathfrak{s}'\colon \tau^G \to \tau^G$. We usually denote it by $\mathfrak{s}$ too.

Barring some technical difficulties, this allows one to define, given a suitable substitution $\mathfrak{s}$, a tiling space, consisting in all the tilings obtained by iterating this process starting from some tile $t \in \tau$. Once again, one needs to be careful in order to properly define the resulting subshift:

- We can consider the set of all the tilings $x \in \tau^G$ that admit arbitrarily many preimages by $\mathfrak{s}$, *i.e.*, there exists for any $n > 0$ a configuration y_n such that $\mathfrak{s}^n(y_n) = x$. This is a dynamical point-of-view, where substitutions act on entire configurations and the final subshift is viewed as a limit space.
- We can also consider the set of tilings x whose patterns are all contained in some finite patch obtained by iterating $\mathfrak{s}$ on some tile t. This is a more combinatorial approach, and is in particular reminiscent of the definition of subshifts by *allowed*, or locally admissible, patterns.

In general, those spaces are different, and might not share all their dynamical properties, in particular mixing properties. However, in simple cases, the distinction usually does not matter when trying to prove that either one is sofic. The formal definitions will be given in Section 4.2.1 for completeness.

In Section 4.2.3, we will define substitutions on graphs, and we will not talk at all about the compatibility problem. Our purpose is indeed to show that a sufficiently connected structure is enough to ensure global self-similarity properties of the limit space using only local constraints. In particular, some difficulties encountered in similar works can be avoided, as we explicitly avoid talking about *e.g.* an embedding in the euclidean plane. In this sense, our point-of-view is closer to the one [BS16] or even [Bar22] than the one of [BH13] or [Pri03] – they will all be presented in Section 4.2.3.

4.2 Substitutive subshifts are sofic

The definition of sofic subshifts given in Section 1.1.4 is not sufficient for our purposes. Nevertheless, we can adapt it quite easily to more general settings, combining the following ideas:

- A sofic subshift is the image under a local map of an SFT.

- An SFT is a set of tilings given by local, matching rules.

The idea of matching rules can quite readily be generalized to tilings on arbitrary structures on which we can define a notion of "neighbouring tiles". In all the cases studied below, we believe that one gets a good understanding of what we mean by a sofic subshift by considering that each tile comes with an additional set of extra colours, thought of as decorations, that are used to build configurations of an SFT, and which are then all forgotten, simultaneously and independently, to give a configuration of the sofic subshift.

We fix some general terminology, which will need to be adapted to the concrete cases studied in the later sections, but which should act as a general framework: for a given set of tiles T and a substitution $\mathfrak{s}$, we say that $\mathfrak{s}(t)$ is a **patch**. Starting from any tile t and iterating the substitution, we obtain **meta-tiles**: the n-meta-tiles, or meta-tiles of level n, are the patterns $\{\mathfrak{s}^n(t), t \in T\}$. In general, we will use a somewhat standard genealogical terminology: if $G \sqsubseteq H$ are respectively n and $n+1$ meta-tiles, we will call H the parent of G, and G will then be a child of H; if $G' \sqsubseteq H$ is another child of H, G and G' will be siblings.

4.2.1 The discrete grid

We first present the simplest case, in a slightly simplified version compared to what was originally proven in [Moz89]. The differences will be highlighted below, and we only give the simpler definition for the time being.

Definition 4.1: Rectangular substitution

Let $n > 0$, and $\mathcal{A}$ some finite alphabet. We define an $m \times n$ **rectangular substitution** as a map $\mathfrak{s}\colon \mathcal{A} \to \mathcal{A}^{m\times n}$.
We extend $\mathfrak{s}$ to entire $\mathbb{Z}^2$ configurations as follows: for any $(i,j) \in \mathbb{Z}^2, 0 \leq k < m, 0 \leq l < n$, we set $\mathfrak{s}(x)_{(ni+k),(mj+l)} = \mathfrak{s}(x_{i,j})_{(k,l)}$.

It is easy to see that this definition gives a colour to every point of the grid, that one can determine using the euclidean division of its coordinates by the size of the initial rectangle. With this definition, there is no compatibility problem and we can easily iterate the substitution. We can therefore look at the limit space obtained when iterating the substitution infinitely often, and define two subshifts:

Definition 4.2: Substitutive $\mathbb{Z}^2$ subshift

Let $\mathfrak{s}$ be a rectangular substitution on alphabet $\mathcal{A}$. We define:

$$X_{\mathfrak{s}} = \left\{x \in \mathcal{A}^{\mathbb{Z}^2} \mid \forall w \sqsubseteq x, \exists a \in \mathcal{A}, n \in \mathbb{N}, w \sqsubseteq \mathfrak{s}^n(a)\right\}$$

and

$$X_{\mathfrak{s}}^{\infty} = \bigcup_{\mathbf{i}\in\mathbb{Z}^2} \sigma^{\mathbf{i}}\left(\left\{x \in \mathcal{A}^{\mathbb{Z}^2} \mid \forall n \in \mathbb{N}, \exists y_n \in \mathcal{A}^{\mathbb{Z}^2}, \mathfrak{s}^n(y_n) = x\right\}\right)$$

For a given substitution $\mathfrak{s}$, $X_{\mathfrak{s}}$ is defined locally, *via* admissible patterns, while $X_{\mathfrak{s}}^{\infty}$ is defined implicitly as the of configurations that can be "desubstituted" infinitely many times. It is easy to show that $X_{\mathfrak{s}} \subseteq X_{\mathfrak{s}}^{\infty}$, and we will prove it in the more general setting of Section 4.3.4, and we show that the converse might not hold with Figure 4.1:

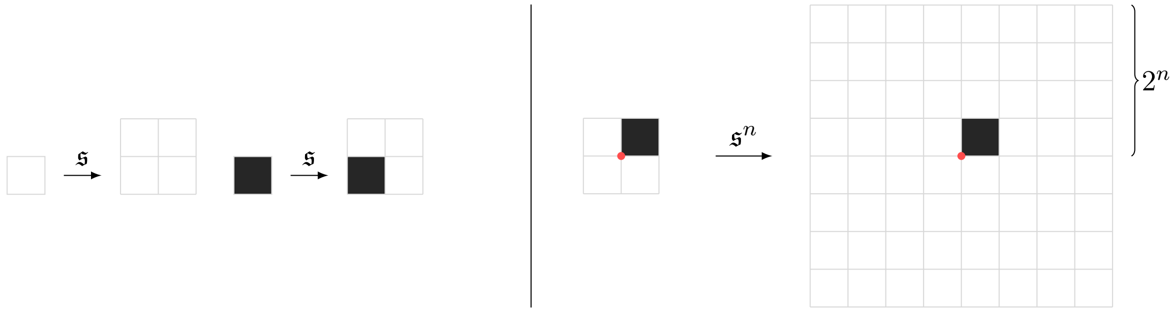

Figure 4.1: An example of substitution for which $X_\mathfrak{s} \neq X_\mathfrak{s}^\infty$ [AS11, Example 1]. $X_\mathfrak{s}$ contains a single point, $\square^{\mathbb{Z}^2}$, but $X_\mathfrak{s}^\infty$ also contains all the configurations containing a single $\blacksquare$.

Remark. *The definitions given here can be generalized in several ways, a few of which we briefly mention:*

- *All the definitions can easily be generalized to higher-dimensional subshifts, although we choose to focus on the two-dimensional case for simplicity.*
- *Instead of considering substitutions with a single rectangle, one can study substitutions in which the domain of the patches $\mathfrak{s}(a)$ might depend on the letter $a \in \mathcal{A}$. In that case, one needs to take some care when iterating the substitution, as it is no longer straightforward to extend $\mathfrak{s}$ to entire $\mathbb{Z}^2$ configurations.*
- *Even more generally, one can consider a sequence of substitutions $(\mathfrak{s}_n)_{n\in\mathbb{N}}$. Once again, the definitions get significantly messier, but under some suitable compatibility conditions, most results that we will state in this chapter can be shown to hold in that case, see for example [AS11].*

In this setting, as we are working with $\mathbb{Z}^2$ subshifts, we can still use the same definition of sofic subshifts. In particular, we have the following version of the meta-Mozes theorem:

Theorem 4.3: Mozes theorem [Moz89]

Let $\mathcal{A}$ be a finite alphabet, $m, n \geq 2$ and $\mathfrak{s}\colon \mathcal{A} \to \mathcal{A}^{m\times n}$ be a rectangular substitution. Then, $X_\mathfrak{s}$ and $X_\mathfrak{s}^\infty$ are sofic subshifts.

This is the first version of the meta-Mozes theorem, actually proven in a more general way by Mozes in [Moz89]. We give a very short outline of the proof:

Proof. Let $R = [\![0, m-1]\!] \times [\![0, n-1]\!]$. In order to create an SFT which is an extension of $X_\mathfrak{s}^\infty$, the idea is to define a tileset, in which every tile will remember some extra information, on top of the symbol it should be mapped to in $X_\mathfrak{s}^\infty$:

- Each tile will remember the rectangle $\mathfrak{s}(a)$ it belongs to, as well as its position (i, j) in this rectangle.
- If (i, j) is a corner of this rectangle $\mathfrak{s}(a)$, then the tile also remembers another rectangle $\mathfrak{s}(a')$ and a position (i', j'), such that $\mathfrak{s}(a')_{(i',j')} = a$.
- Some other tiles, which are not in the corner of their own rectangle $\mathfrak{s}(a)$, will also carry and propagate some "signals" which are used to ensure that neighbouring meta-tiles are consistent (that is, to the right of some meta-tile which thinks it is $\mathfrak{s}(a)_{(i,j)}$,

we make sure that we have a meta-tile which thinks it is $\mathfrak{s}(a)_{(i+1,j)}$, assuming (i,j) itself is not on the right border of R).

There are various signals of this form, depending on some minor differences in the specific information that needs to be propagated. The main point is that they are sufficient to identify a meta-tile T of some arbitrary level k, and propagate this information to an arbitrary distance in $\mathbb{Z}^2$, required to check the consistency of the level k. Some signals ensure that T itself is well-formed (that is, it can be decomposed in $k-1$ meta-tiles), and some other ones are used to ensure that its neighbourhood is well-formed (that is, it is correctly aligned with n meta-tiles, which carry consistent information regarding the $n+1$ meta-tile they might collectively belong to).

The key observation is that a meta-tile needs only finitely many "active sites", collecting the information of this tile (which substitution $\mathfrak{s}(a)$ does it originate from, what is its position in the image ...) and checking consistency with the neighbouring meta-tiles, and with the parent meta-tile. The difficulty lies in transmitting this information from the children tiles to those specific active sites. □

4.2.2 The euclidean plane

We now say a few words about another well-studied case in the literature, which is the case of euclidean substitutions: tiles are no longer finite subsets of $\mathbb{Z}^d$, but subsets of $\mathbb{R}^d$, usually homeomorphic to a closed ball – sometimes, only compactness of the tiles is required. In this setting, a substitution over some tileset $\mathcal{A}$ is usually defined as an expanding linear map $\mathfrak{s}$, such that for any $a \in \mathcal{A}$, $\mathfrak{s}(a)$ is a union of isometric copies of tiles of $\mathcal{A}$ with disjoint interiors. The specific set of isometries being considered varies, but is usually assumed to consist of translations and rotations. A simple example of this kind of substitutions over $\mathbb{R}^2$ is presented in Figure 4.2.

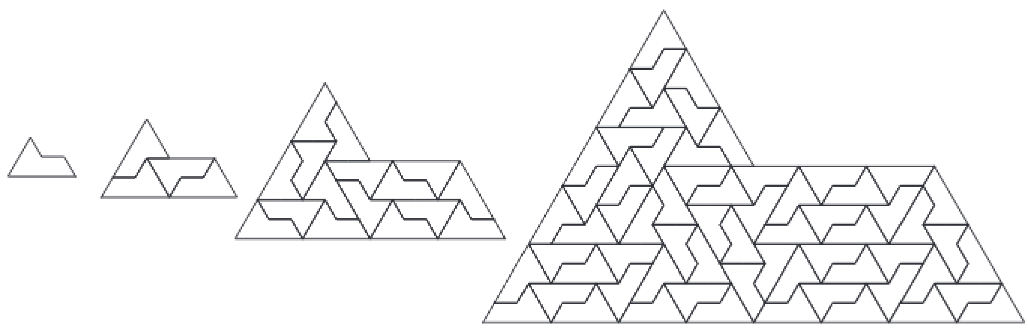

Figure 4.2: The Sphinx substitution. Figure taken from [Goo16], in which this specific substitution is studied in detail.

In the following, we call "sofic subshift" a subshift (or set of tilings) on the set of tiles $\mathcal{A}$ that can be enforced by adding some decorations on top of $\mathcal{A}$: each tile $a \in \mathcal{A}$ is endowed with (possibly several) decorations on its border ∂a, and the tiling is valid if adjacent tiles carry the same information on the common part of their borders (which is non-empty, as tiles are closed subsets of $\mathbb{R}^d$).

Matching rules for euclidean substitutive tilings

For our purposes of determining conditions under which the Mozes theorem holds, the main result about this kind of euclidean tilings is due to Goodman-Strauss in [Goo98]. We reproduce this article's main theorem here, and comment it briefly afterwards:

Theorem 4.4: Mozes theorem - Euclidean [Goo98]

If $\mathfrak{s}$ is a substitution in $\mathbb{R}^n$ which is "hereditary" and "sibling edge-to-edge", then $X_{\mathfrak{s}}$ is sofic.

The hereditary and sibling edge-to-edge conditions are technical restrictions, which more or less require that the edges of the meta-tiles are concatenations of edges of the initial tiles, and that adjacent meta-tiles must touch along entire edges. We refer to [Goo98] for a precise description of those conditions, which are considered to be very mild.

This is a very general result, which nevertheless relies on an embedding of the tiles in the euclidean space.

Combinatorial point-of-view on euclidean tilings

Following [Goo98], the authors of [FO10] prove a similar result. They consider tiles that are polytopes in $\mathbb{R}^d$, with finitely many facets, and define a **combinatorial substitution** as a map $\mathfrak{s}$ sending a tile $a \in \mathcal{A}$ to a finite tiling, *i.e.* a tiling of a bounded region of $\mathbb{R}^d$ by $\mathcal{A}$. In particular, a tile a need not be similar to $\mathfrak{s}(a)$ (said differently, $\mathfrak{s}$ need not be an expanding linear map). A large number of known tilings from the literature fall under this definition, for example the Rauzy tilings (see [AI01] or [Fer07] for a discussion), or even more famously the Penrose tilings (see [GS87]). We restate the main result of [FO10] with our notations:

Theorem 4.5: Mozes theorem - combinatorial euclidea [FO10, Thm. 1.1]

If $\mathfrak{s}$ is a *good* combinatorial substitution, then $X_{\mathfrak{s}}^{\infty}$ is sofic.

As is Theorem 4.4, we need an additional assumption on the substitution for the theorem to hold: a combinatorial substitution $\mathfrak{s}$ is called *good* if it satisfies some conditions, which ensure that each image $\mathfrak{s}(a)$ is "connected enough", and that a valid tiling by the 1-meta-tiles $\{\mathfrak{s}(a) \mid a \in \mathcal{A}\}$ is the image of an actual tiling by $\mathcal{A}$. More details on these conditions can be found in the original article, but deciding whether a given substitution is "good" or not is still an open problem.

Both the definition of a "combinatorial substitution", and the proof of Theorem 4.5, are indeed much more combinatorial in nature. However, the tilings are ultimately still tilings of $\mathbb{R}^d$, and it is unclear how this approach could be generalized to settings with a completely different (or absent) geometry, for example in the case of tilings on groups (see Section 1.3.3).

4.2.3 Beyond the geometry

Soficity relative to substitutive discrete subshifts

In order to study the boundary between sofic and non-sofic subshifts, at least from the point-of-view of substitutions, has been explored in [BS16]. The original question asked by the authors was about the difficulty of solving the Domino problem over self-similar structures, and they tried to find criteria to distinguish whether the domino problem was solvable or not. However, they also prove that an equivalent to Mozes theorem holds over various such self-similar structures, independently from the solvability of the Domino problem. The self-similar structures they consider are still defined relative to the discrete grid $\mathbb{Z}^2$, and this is one of the generalizations we try to study later on. The results and definitions also hold for higher dimensions, we restrict ourselves to the $\mathbb{Z}^2$ case for simplicity.

Definition 4.6: Self-similar colouring [BS16]

Let $\mathfrak{s}$ be a rectangular $m \times n$ substitution, with binary alphabet $\{0, 1\}$, and such that $\mathfrak{s}(0) = 0^{m\times n}$. A substitution $\mathfrak{s}_c$ is **compatible** with $\mathfrak{s}$ if it has the same shape, and is defined over $\mathcal{A} \supseteq \{0\}$, such that $\mathfrak{s}_c(0) = 0^{m\times n}$ and for all $0 \le i < m, 0 \le j < n$, and $a \in \mathcal{A} \setminus \{0\}$, $\mathfrak{s}_c(a)_{i,j} = 0 \iff \mathfrak{s}(1)_{i,j} = 0$. The set of $\mathfrak{s}_c$-self-similar colourings is

$$X_{\mathfrak{s}_c} = \left\{x \in \mathcal{A}^{\mathbb{Z}^2} \mid \forall w \sqsubseteq x, \exists a \in \mathcal{A}, n \in \mathbb{N}, w \sqsubseteq \mathfrak{s}_c^n(a)\right\}$$

The definition can be considered to be in two steps: in a first step, we define a structure using a binary substitution $\mathfrak{s}$, which has a blank symbol; over this structure, we can define other substitutions, provided they produce blank symbols exactly at the same positions than $\mathfrak{s}$.

We can then define SFTs relative to this structure, and state a general property:

Definition 4.7: Relative Mozes property [BS16, Def. 1]

A binary rectangular substitution $\mathfrak{s}$ has the **Mozes property** if for every compatible substitution $\mathfrak{s}_c$ over $\mathcal{A}$, there exists an alphabet $\mathcal{B} \supset \{0\}$, and SFT $X_{\mathcal{F}} \subseteq \mathcal{B}^{\mathbb{Z}^2}$ and a factor map $\phi \colon \mathcal{B} \to \mathcal{A}$ such that $\Phi^{-1}(0) = \{0\}$ and $\Phi \colon X_{\mathcal{F}} \twoheadrightarrow \mathcal{A}^{\mathbb{Z}^2}$ defined by $\Phi(x)_{(i,j)} = \phi(x_{(i,j)})$ is surjective onto $X_{\mathfrak{s}_c}$.

In other words, a substitution has the Mozes property if every compatible substitutive colouring of the "binary self-similar structure" it generates is sofic in a strong sense: there is an SFT factoring onto it, which furthermore places 0-tiles on top of the zeros of the structure.

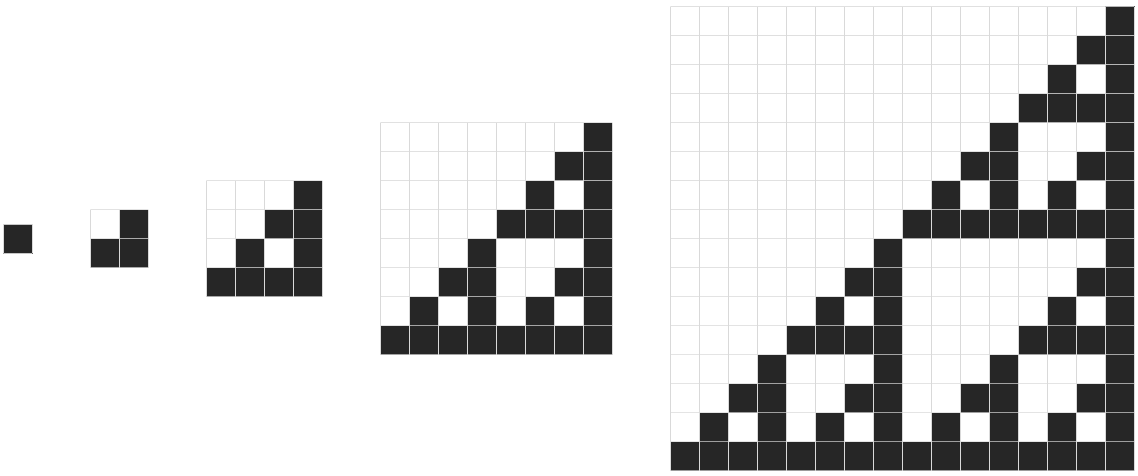

Figure 4.3: The first five iterations of the grid-Sierpinski substitution.

The authors then demonstrate that a few substitutions have the Mozes property, in particular:

Proposition 4.8 [BS16, Th.4]

The grid-Sierpinski substitution (see Figure 4.3) has the relative Mozes property.

We do not detail the proof here, as it is quite similar, although in a simpler setting, to the proof we do in Section 4.4. In particular, we will use the terminology of this article to present the general ideas in our proof.

Substitutions on groups

One of the motivations behind our definitions of Section 4.3 is to understand how one could define substitutions on groups, and use classical results about $\mathbb{Z}^2$ substitutions in this new setting, although a general, satisfying definition of substitution on groups remains open. We give an example of a case where Mozes theorem has successfully been adapted to a class of finitely-presented, non-abelian groups, namely the Baumslag-Solitar groups:

Definition 4.9: Baumslag-Solitar groups [BS62]

Let $m, n \geq 1$. The **Baumslag-solitar group** $BS(m, n)$ is defined by

$$BS(m, n) = \langle a, b \mid a^n b = ba^m \rangle$$

For $N > 1$, the group $BS(1, N)$ is non-abelian but amenable.

We do not give a definition of amenability here, as our focus is not solely algebraic, and it suffices to say that this class of group is particularly interesting for symbolic dynamics (see for example [AK13] or [EM22]), as it exhibits many “nice” properties while being significantly more general than $\mathbb{Z}^2 = BS(1, 1)$, being for example non-abelian. We show in Figure 4.4 part of the Cayley graph of $BS(1, 2)$ for the presentation given in Definition 4.9 – this is also [Sil20, Figure1.1], from which we present an important result.

It so happens that there exists a recursive decomposition of the Cayley graph of $BS(1, N)$ into what the author of [Sil20] calls **rectangles** ([Sil20, Definition 3.1]), which are finite subgraphs, defined by $R_m = \{a^j b^k \mid 0 \leq j < N^m, 0 \leq k < m\}$ for $m \geq 1$. A visual representation of the corresponding graph, and the decomposition, is shown in Figure 4.5.

Using this decomposition, one can define a notion of substitution, the limit graph of which is the Cayley graph of $BS(1, N)$ for any $N \geq 2$. To the best of our understanding, this substitution is defined in a rather *ad hoc* way, and in particular one needs to differentiate multiple cases when iterating a substitution, depending on which “vertex” is being substituted. Nevertheless, we have the following result:

Theorem 4.10: Mozes theorem - $BS(1, N)$ [Sil20, Thm. 3.20]

For $m \geq 2$, $\mathfrak{s}\colon \mathcal{A} \to \mathcal{A}^{R_m}$ a substitution on $BS(1, N)$, the subshift $X_{\mathfrak{s}} \subset \mathcal{A}^{BS(1,N)}$ is sofic.

The proof follows the same ideas as the other versions of the Mozes theorem. As mentioned above, the definition of the substitution uses a rather fortunate decomposition of the group into specific sets, and some care still needs to be taken when specifying how one

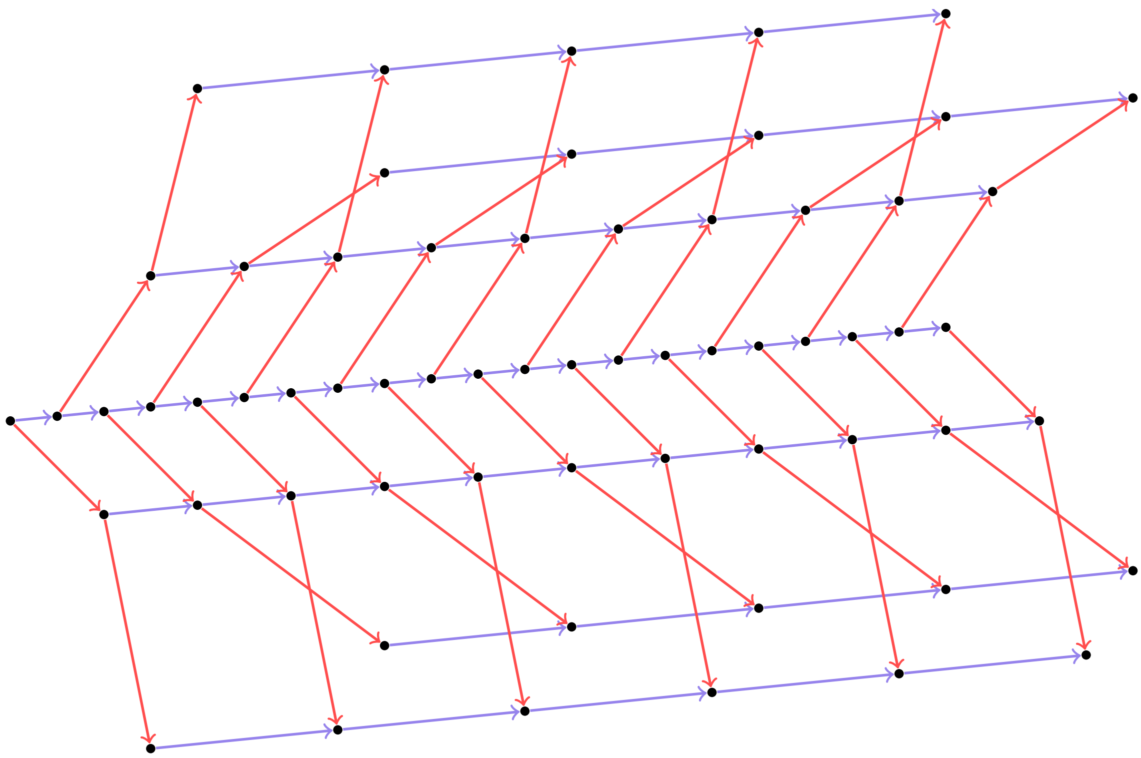

Figure 4.4: Part of the Cayley graph of $BS(1,2) = \langle a, b \mid a^2 = bab^{-1} \rangle$. Red arrows correspond to the generator b, purple arrows to a. The convention is the one of Definition 1.91.

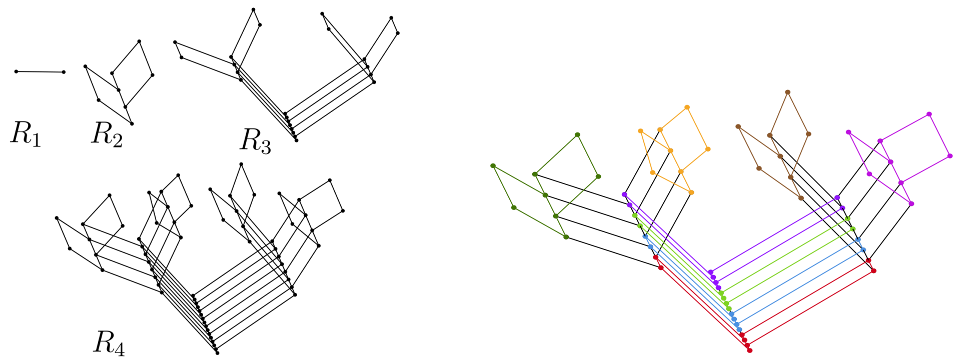

(a) The first four rectangles of the group $BS(1,2)$ ([Sil20, Figure 3.1])

(b) Decomposition of R_4 into 8 copies of R_1 ([Sil20, Figure 3.2])

Figure 4.5: The recursive decomposition of the right Cayley graph of $BS(1,2)$ in rectangles R_m.

should iterate the substitution. See [Sil20, Section 3.2, Section 3.4] for more details. The approach we take in Section 4.3 is different, as it defines substitutions on spaces which are more general than Cayley graphs. On the other hand, it is unclear which kinds of groups fit into our definition, and for those who do, whether there is an effective procedure to derive appropriate substitutions from *e.g.* a finite presentation.

Other combinatorial notions of substitutions

We give a few additional examples of various notions of substitutions which try to give a purely combinatorial description of the objects they consider, and try to highlight the main differences with our approach. A survey of similar constructions can be found in [Pri08], as well as some general results about substitutive subshifts which are out of the scope of this chapter.

Planar graphs Another definition of combinatorial substitutions using graphs, which inspires the one used in [FO10] described in Section 4.2.2, was proposed in [Pri03]. We do not describe it fully there, as a lot of definitions are similar to the ones we use in Section 4.3.4. The general idea is to consider *planar* graphs (informally, graphs that can be drawn in a plane without any two edges intersecting). Now, the substitution $\mathfrak{s}$ simply replaces each vertex v, each edge e, and each facet f of the graph by another graph, using additional maps which ensure that if $v \in e \in f$, we have $\mathfrak{s}(v) \sqsubseteq \mathfrak{s}(e) \sqsubseteq \mathfrak{s}(f)$. We highlight one important difference: in Section 4.2.2, the author's goal is to *model* euclidean substitutions using more combinatorial objects, and as such, uses *facets* of the graph to iterate the substitution. A definition that could be suitable for groups must not reference any kind of planar embedding, as Cayley graphs need not be planar. Moreover, even if the substitution $\mathfrak{s}$ is defined on planar graphs, it is unclear whether the image of any planar graph Γ by $\mathfrak{s}$ remains planar.

Topological substitution Finally, we briefly describe the approach taken in [BHJ18], in which the authors define substitution from a topological point-of-view. The question they try to answer is the following, using a reformulation in terms of the objects introduced so far: given a substitutive subshift $X_{\mathfrak{s}} \subset \mathcal{A}^{\mathbb{R}^2}$ defined by an expanding linear map $\mathfrak{s}$, is there a *topological substitution* $\mathfrak{s}'$ such that $X_{\mathfrak{s}}, X_{\mathfrak{s}'}$ are "equivalent" ? The notion of topological substitution was introduced in [BH13]: the actual definition uses the notion of CW-complex, which has not been introduced in this thesis, and we rely on a more informal description.

We call n-polygon the cyclic graph C_n. A tileset is then a set $\mathcal{A} = \{t_1, \ldots, t_n\}$ where each t_i is an n_i-polygon for some $n_i \geq 3$. An $\mathcal{A}$-patch is then a graph G such that there exists a decomposition of $G = (V, E)$ as tiles, that is, $V = \bigcup_i V_i, E = \bigcup_i E_i$ where for each i, we have $(V_i, E_i) \in T$ and moreover each edge $e \in E$ belongs to at most two distinct $(V_i, E_i) \sqsubseteq G$ – there are some additional restrictions, as the authors consider complexes which can be embedded in the plane, but we skip over those for conciseness purposes. A substitution is then defined as a map $\mathfrak{s}$ sending each tile t_i to a patch. Some care needs to be taken to define how a patch itself should be substituted, that is, to ensure that one can "glue" the images of each tile computed separately in order to obtain another patch – an important part of [BHJ18] is devoted to the description of combinatorial conditions for a substitution to be iterated infinitely many times, which can effectively be checked from the images of the t_i's – see Section 3 of the article for more details.

Although very general, this definition of substitution suffers from the same drawbacks for our purposes than the ones already highlighted in Section 4.2.3: as the author's main goal is to study and model euclidean substitution, the setting emphasizes the fact that each patch must be planar. This is once again not suitable for modeling substitutions on groups, or graphs of algebraic origin, and we relax this condition later on.

4.3 Graph subshifts

In this section, we will try to propose yet another definition of substitution, which avoids any explicit reference to the geometry of an ambient space, solely relying on combinatorial objects, namely graphs. In particular, and contrary to most of the examples described in Section 4.2, substitutive subshifts will not simply be a specific class of colourings of some space (whether $\mathbb{Z}^d$ or $\mathbb{R}^d$) obtained by iterating a "substitutive" map $\mathfrak{s}$, but the space itself will be obtained in the same way. We therefore need to define what a subshift is in this context:

- Section 4.3 will define **graph subshifts**, which are families of graphs respecting some local constraints.
- Section 4.3.4 will introduce or notion of substitution on graphs, and explain how they can be iterated.
- Section 4.3.5 will define what it means for a subshift to be sofic in this context, and more generally, adapt some of the definitions of Section 1.1.2 and Section 1.1.4 in particular.

4.3.1 Basic graph theory

Among the previous examples, we especially try to generalize the substitutive subshifts defined in [BS16], as it tries to only look at the combinatorial aspects of the graph underlying a discrete limit space generated by substitutions. In particular, a discrete structure, akin to a graph, is initially defined, and despite being subshifts on $\mathbb{Z}^2$, the subshifts studied by the authors on this structure behave as if they were defined on the underlying graph only. Nevertheless, it is possible to eliminate the grid from the intermediary step, by only considering easy-to-iterate substitutions directly defined on graphs.

We give here the main additions to the definitions that were already presented in Section 1.1.4:

- G' is a **subgraph** of G if $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$
- If furthermore $E(G') = E(G) \cap V(G')^2$, we say that G' is an **induced subgraph** of G, that we denote $G|_{V(G')}$.
- For a graph G that are oriented and edge-labeled with directions $\mathcal{D}$, we say that some vertex u is adjacent to a d-edge if there exists $e = (u, v) \in E(G)$ with $\lambda_E(e) = d$. It is adjacent to a d^{-1}-edge if there exists an edge $e = (v, u)$ with direction d.
- In the case of labeled graphs $G = (V, E, \lambda_V, \lambda_E)$, the notation $G|_{V'}$ also implies that the labeling function is now defined on the new vertices and edges sets, *i.e.* $G|_{V'} = (V', E' = E \cap V'^2, \lambda_V|_{V'}, \lambda_E|_{E'})$.
- For labeled graphs $G_1 = (V_1, E_1, \lambda_{V_1}, \lambda_{E_1})$, $G_2 = (V_2, E_2, \lambda_{V_2}, \lambda_{E_2})$, with $\lambda_{,1}, \lambda_{,2}$ taking values in some common sets $\mathcal{C}$ and $\mathcal{D}$, we call **disjoint graph union** of G_1 and G_2 the graph $G = (V_1 \sqcup V_2, E_1 \sqcup E_2, \lambda_V : V_1 \sqcup V_2 \to \mathcal{C}), \lambda_E : E_1 \sqcup E_2 \to \mathcal{D})$, the labeling functions having the obvious definition. We might also use the notation $G \sqcup E'$ for some set of edges $E' \subseteq V(G)^2$, in which case we mean that the edges of E' are not already in $E(G)$.
- A graph morphism $\phi : G \to \Gamma$ is a map satisfying that for all $e = (u, v) \in E(G)$, $(\phi(u), \phi(v)) \in E(\Gamma')$. We will usually consider labeled graph morphisms, in which case G, Γ' have the same types and directions sets, and that ϕ preserves types and directions.
- Unless explicitly mentioned, we use **labeled graph** for vertices and edge-labeled graphs, that is, a graph $G = (V, E)$ along with some maps $\lambda_V : V \to \mathcal{C}$ and $\lambda_E : E \to \mathcal{D}$. The sets of labels $\mathcal{C}$ and $\mathcal{D}$ are always assumed to be disjoints. We respectively refer to them as the set of vertices **types** and edges **directions**. In what follows, it is useful to think of those as structural properties of the vertices or edges, rather than

as colourings. This distinction is purely conceptual, but later on, when talking about subshifts on graphs, we will add another layer of colouring to the vertices. In this setting, the vertices types will be used to determine the "shape" of the substitution, regardless of their colour.

- In some problems, (possibly labeled) graphs $G = (V, E)$ that we define might naturally be directed, but it will sometimes be convenient to consider them as undirected. What we mean by this is that we instead consider the graph the undirected graph $G' = (V, E')$, where $E' = \{(u, v) \in V^2, (u, v) \in E \text{ or } (v, u) \in E\}$. If G was labeled, we give both edges in E' the label of the original edge in E. In order for this to be well-defined, we will ensure that this operation is only ever performed on graphs $G = (V, E, \lambda_E)$ in which $(u, v), (v, u) \in E \implies \lambda_E(u, v) = \lambda_E(v, u)$.

4.3.2 A specific class of graphs

We will consider graphs up to isomorphism. This will induce some difficulty later on, when trying to refer to some specific vertices in the graph, or when trying to define a topology on the space of graphs. We will therefore fix some restriction on the set of graphs considered in the rest of the chapter. The next definitions are simply a way to explicit some of the conditions imposed in [ADG23]. We adopt a slightly different presentation: rather than using a set of "ports" on each vertex, with edges potentially joining two vertices using arbitrary ports, we enforce *a priori* the fact that an edge uses the same port on both vertices, and edges are now oriented. This is an equivalent point of view, but it will be more suited to our definition of substitution later on. A graph is always assumed to be defined as in Section 4.3.1:

Definition 4.11: Pointed graph

A pointed graph is a pair (G, v) where $v \in V(G)$. We denote $\mathrm{basepoint}((G, v))$ the vertex v, called the base-point of (G, v).

From now on, all the graphs are considered up to isomorphism. We say that a (Γ, u) and (Γ', u') are isomorphic as pointed graphs if Γ and Γ' are isomorphic *via* some isomorphism ϕ, and furthermore $\phi(u) = u'$. When talking about pointed graphs Γ, Γ', we will write $\Gamma = \Gamma'$ when they are isomorphic as pointed graphs, and $\Gamma \simeq \Gamma'$ if they are isomorphic as non-pointed graphs. Equivalently, $(\Gamma, u) \simeq (\Gamma', v) \iff \exists w \in V(\Gamma), (\Gamma, w) = (\Gamma', v)$.

Definition 4.12: Graph class

Let $\mathcal{C}$ and $\mathcal{D}$ be finite sets of directions and types respectively. The **graph class** defined by $\mathcal{C}, \mathcal{D}$ is the set $\mathcal{G}_{\mathcal{C},\mathcal{D}}$ of pointed, connected, possibly infinite graphs, with vertices labeled by $\mathcal{C}$ and edges labeled by $\mathcal{D}$, where each vertex is adjacent to at most a single edge of each given direction and orientation.

In particular, the degree of any vertex of a graph $\mathcal{G}_{\mathcal{C},\mathcal{D}}$ is bounded by $2|\mathcal{D}|$, both the in-degree and out-degree being independently bounded by $|\mathcal{D}|$.

As in Section 1.1.3, we can define a few topological notions on those graphs:

Definition 4.13: Shift

Let (Γ, u) be a pointed graph, and $d \in \mathcal{D}$ a direction. The **shift** of (Γ, u) in direction d is the pointed graph $\sigma_d(\Gamma, u) = (\Gamma, v)$ where v is the unique vertex such that $(u, v) \in E(\Gamma)$ and (u, v) has direction d if it exists, or $v = u$ otherwise.

Definition 4.14: Ball

Let (Γ, u) be a pointed graph and $n \geq 0$. The **ball** of radius n is the subgraph $\mathcal{B}_r(\Gamma, u)$ of Γ induced by the vertices $\{v \in \Gamma, d_\Gamma(u, v) \leq n\}$, pointed in u.

Given this notion of ball, we can derive the usual notion of distance for configurations, just as in the $\mathbb{Z}^d$ case: configurations will be close to each other if they agree on a large central ball.

Definition 4.15: Distance

Let (Γ, u) be a pointed graph, $\mathcal{A}$ a finite alphabet and $x, y \in \mathcal{A}^{V(\Gamma)}$ two colourings of $V(\Gamma)$. The **distance** between x and y is

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ 2^{-\min_n \{x|_{B_n} \neq y|_{B_n}\}} & \text{otherwise} \end{cases}$$

We use the notation $\mathcal{B}_r(\mathcal{G}_{\mathcal{C},\mathcal{D}})$ for the set of all the finite pointed graphs of $\mathcal{G}_{\mathcal{C},\mathcal{D}}$ containing only vertices at distance at most r from u.

4.3.3 Graph subshifts and SFTs

As mentioned above, in order to talk about individual vertices in a pointed graph up to isomorphism, we need some kind of addressing scheme. Whenever we restrict ourselves to the case of connected pointed graphs, any vertex can be designed by a path from the base-point of the graph to it. Note however that this path is not necessarily unique.

Notation. We will $\bar{\mathcal{D}} = \bar{\mathcal{D}}$.

Definition 4.16: Path address

Let $(\Gamma, u) \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. Then, for any finite sequence $D = (D_i)_{1 \leq i \leq n}, D_i \in \bar{\mathcal{D}}$, called **path address**, we write $\text{path}_{(\Gamma,u)}(D)$ the vertex obtained by following the edges labeled by $D_1, \ldots, D_n$ in this order starting from u in Γ, if this vertex exists.
We denote $\text{Paths}(\Gamma, u)$ the set of valid path addresses of Γ starting from u.

Note that in any given pointed graph (Γ, u), the map $\text{path}_{(\Gamma,u)} \colon \bar{\mathcal{D}}^* \to V(\Gamma)$ is not a total map, is generally not injective, but is surjective as $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ is connected.

Some more discussion about the property of this addressing scheme can be found in [ADG23].

In order to define subshifts, we need to define what a pattern is. A natural idea would be to define patterns as finite (labeled) graphs, and say that a larger graph Γ avoids a pattern F if it contains no subgraph H isomorphic to F – one would also need to properly study the difference between enforcing H to be *induced* or not. However, such a definition lacks expressive power: as mentioned in [ADG23, Remark 6], this makes it impossible to specify that a subgraph *must* appear: one could reasonably argue that, for example, it is natural to be able to enforce conditions of the kind "every vertex of type $c \in \mathcal{C}$ has to be the starting point of an edge of direction $d \in \mathcal{D}$. We therefore need to have a slightly more general definition of what a pattern is, in order to be able to express this kind of condition. We will first give the original definitions of the authors of [ADG23], and then give some equivalent characterizations, some of which can already be found in this article.

Definition 4.17: Prefix-stable language

Let $\mathcal{D}$ be a finite set of directions. A **language** is a set $\mathcal{L} \subset \bar{\mathcal{D}}^*$ of finite words. It is **prefix-stable** if $uv \in \mathcal{L} \implies u \in \mathcal{L}$ for any $u, v \in \bar{\mathcal{D}}^*$.

Using Definition 4.16, we can therefore use languages to designate sets of vertices in a pointed graph. The condition of being prefix-stable then implies that the set of vertices thus considered is a connected subgraph of Γ.

Definition 4.18

Let $\mathcal{L}$ be a prefix-stable language, and $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. We denote $\Gamma|_{\mathcal{L}}$ the induced subgraph of Γ given by the path addresses of $\mathcal{L}$:

$$\Gamma\big|_{\mathcal{L}} = \Gamma\big|_{\{\text{path}_{(\Gamma,u)}(D) | D \in \text{Paths}(\Gamma,u) \cap \mathcal{L}\}}$$

We can now give the definition of a pattern and of a subshift in terms of forbidden patterns. This detour by languages is required in order to express some natural conditions, as explained above. The difference with *e.g.* subshifts on $\mathbb{Z}^d$ can be understood as follows: in $\mathbb{Z}^d$ subshifts, the underlying "geometry" is fixed, and does not depend on the subshift: we colour each cell of the d-dimensional grid, but each cell will always have the same $2d$ adjacent cells. In our case, graphs are not *a priori* regular, that is, vertices might neighbourhoods of different sizes; this relaxation is useful in several examples that we deal with – for example, our substitutive graphs (see Section 4.3.4 will not be regular in general) – but we therefore need a way to express some new kind of conditions, such as explicitly requiring the presence of some edge in some specific subgraphs.

Definition 4.19: Graph pattern - Language

A (language-)**pattern** is a pair $((\Gamma, u), \mathcal{L})$ where $(\Gamma, u) \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ and $\mathcal{L}$ is a prefix-stable language. It is finite if Γ and $\mathcal{L}$ are finite.

We can now define what containing or avoiding a pattern means:

Definition 4.20: Pattern avoidance

Let $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$, and let $(F, \mathcal{L})$ be a pattern. We say that Γ contains $(F, \mathcal{L})$ in position $v \in V(\Gamma)$ and write $(F, \mathcal{L}) \sqsubseteq \Gamma$ if $(\Gamma, v)|_{\mathcal{L}} = F$. If Γ does not contain $(F, \mathcal{L})$ in any position, we say that it **avoids** $(F, \mathcal{L})$.

Although the role of $\mathcal{L}$ in the previous definition might still be unclear, one can think of it as a way of imposing, or preventing, the presence of specific edges, on some vertices of F. In particular, $\mathcal{L}$ cannot be completely arbitrary: the following proposition shows that we can consider only languages that designate actual vertices of the underlying graph F, or possibly neighbouring ones:

Proposition 4.21

Let $((F, u), \mathcal{L})$ be a finite pattern, and define the language $\mathcal{M} = \left\{D = (D_1, \dots, D_n) \in \mathcal{D} \mid \mathrm{path}_{(F,u)}(D_1, \dots, D_{n-1}) \in V(F)\right\}$.
Then, for any $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$, $(F, \mathcal{L}) \sqsubseteq \Gamma \iff (F, \mathcal{M}) \sqsubseteq \Gamma$.

Proof. Let $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. For the direct implication, suppose that Γ contains $(F, \mathcal{L})$ in position v, *i.e.* $(\Gamma, v)|_{\mathcal{L}} = F$. In what follows, we consider that Γ is pointed at v. It is clear that $\Gamma|_{\mathcal{M}} \sqsubseteq \Gamma|_{\mathcal{L}}$, as $\mathcal{M} \subset \mathcal{L}$. We show that the other inclusion also holds, showing that $\Gamma|_{\mathcal{M}} = F$. By contradiction, suppose that there is path $D \in \mathcal{L}$ of minimal length such that $\mathrm{path}_{(\Gamma,v)}(D) \in \Gamma|_{\mathcal{L}} \setminus \Gamma|_{\mathcal{M}}$. As D is of minimal length and $\mathcal{L}$ is prefix-stable, denoting $D' = D_1, \dots, D_{|D|-1}$, we have $D' \in \mathcal{L}$. As $\Gamma|_{\mathcal{L}} = F$, we obtain $\mathrm{path}_{(\Gamma,v)}(D') = \mathrm{path}_{(F,u)}(D')$, and so $D \in \mathcal{M}$, which is a contradiction.

For the other inclusion, suppose that $\Gamma|_{\mathcal{M}} = F$. We show that $\Gamma|_{\mathcal{L}} = F$. As above, the inclusion $\Gamma|_{\mathcal{M}} \subset \Gamma|_{\mathcal{L}}$ is clear. Let $D \in \mathcal{L}$ such $w = \mathrm{path}_{(\Gamma,v)}(D) \in \Gamma|_{\mathcal{L}}$. By assumption, we get that $\mathrm{path}_{(F,u)}(D) \in V(F)$. By definition of $\mathcal{M}$ and $\mathcal{L}$ being prefix-stable, we obtain $D \in \mathcal{M}$, and so $w \in \Gamma|_{\mathcal{M}}$. Hence, $\Gamma|_{\mathcal{L}} = \Gamma|_{\mathcal{M}} = F$. □

If $((F, u), \mathcal{L})$ is such that for any $D = (D_1, \dots, D_n) \in \mathcal{L}$, we have $\mathrm{path}_{(F,u)}(D_1, \dots, D_{n-1}) \in V(F)$, we say that the pattern $(F, \mathcal{L})$ is **reduced**.

Definition 4.22: Graph subshift [ADG23, Def. 3]

Let $\mathcal{F}$ be a set of finite patterns. The **graph subshift** forbidding $\mathcal{F}$ is

$$X = \{\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}} \mid \forall (F, \mathcal{L}) \in \mathcal{F}, (F, \mathcal{L}) \not\sqsubseteq \Gamma\}$$

It is a **graph subshift of finite type** if $\mathcal{F}$ can be chosen finite.

We will give below a few examples of graph subshifts, and try to illustrate why we needed to use prefix-stable languages in the definition of patterns and pattern avoidance. Before that, we give a few equivalent definitions of a graph SFT. The first reformulation is easy, and says that SFTs can be defined using allowed patterns instead of forbidden patterns.

Proposition 4.23 [ADG23, Remark 4]

The following are equivalent:

1. X is a subshift of finite type on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$.
2. There exists a finite family of allowed patterns $\mathcal{A}$ such that

$$X = \left\{\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}} \mid \forall v \in V(\Gamma), \exists (A, \mathcal{L}) \in \mathcal{A}, (\Gamma, v)\big|_{\mathcal{L}} = A\right\}$$

As the link between the language $\mathcal{L}$ and the graph F used to define a language-pattern $(F, \mathcal{L})$ might be hard to grasp, we give an equivalent and hopefully clearer definition of subshifts, and a way to pass from one presentation to the other:

Definition 4.24: Annotated graph

An **annotated graph** is a graph $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ with a (partial) map $A\colon V(\Gamma) \times \bar{\mathcal{D}} \to \{\texttt{required}, \texttt{forbidden}\}$, where $A(v, d)$ might be defined only if v is not already adjacent to a d-edge in Γ.

Definition 4.25

Let $((G, v), A)$ be an annotated graph, and $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. We say that G appears in Γ in position u as an annotated graph if there exists an injective graph morphism $\phi\colon (G, v) \hookrightarrow (\Gamma, u)$ and furthermore, for all $(w, d) \in V(G) \times \bar{\mathcal{D}}$:

- If $A(w, d) = \texttt{required}$ then $\phi(w)$ must be adjacent to some d-edge in Γ.
- If $A(w, d) = \texttt{forbidden}$ then $\phi(w)$ must not be adjacent to a d-edge in Γ.

Proposition 4.26

For any finite pattern $(F, \mathcal{L})$, there exists a finite annotated graph (F', A) such that for any $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$, $(F, \mathcal{L}) \sqsubseteq \Gamma$ if and only if (F', A) appears in Γ as an annotated graph.

Proof. Let $(F, \mathcal{L})$ be a finite pattern. Suppose that it is reduced, that is, $\mathcal{L}$ is such that for $D = (D_1, \ldots, D_n) \in \mathcal{L}$, we have $\mathrm{path}_{(F,u)}(D_1, \ldots, D_{n-1}) \in V(F)$. We define an annotation function A on F. Let $D \in \mathcal{L}$ be such that $D = (D_1, \ldots, D_n)$ is not a valid path address in (F, u). Noting $D' = (D_1, \ldots, D_{n-1})$, this means that $w = \mathrm{path}_{(F,u)}(D')$ is not adjacent to an edge with direction $D_n \in \bar{\mathcal{D}}$. For all such pairs (w, D_n), set $A(w, D_n) = \texttt{forbidden}$. We claim that $(F, \mathcal{L}) \sqsubseteq \Gamma$ if and only if Γ contains (F, A) as an annotated graph.

Suppose that $(\Gamma, v)|_{\mathcal{L}} = F$. Then, for any $D = D' \cdot D_n \in \mathcal{L}$, if D is not a valid path address for (Γ, v), then D' is a valid path address as we assumed that $(F, \mathcal{L})$ is reduced, and by construction $A(\mathrm{path}_{(F,u)}(D'), D_n) = \texttt{forbidden}$. Moreover, $(\Gamma, u)|_{\mathcal{L}} = F$ implies that $\mathrm{path}_{(\Gamma,v)}(D')$ is not adjacent to a D_n-edge. Therefore, (F, A) appears in Γ as an annotated graph.

On the other hand, suppose that $((F,u),A)$ appears in Γ in position v. In particular, there is a injective graph morphism $\phi\colon F \hookrightarrow \Gamma$ sending u to v. Let $D = D' \cdot D_n \in \mathcal{L}$. As $\mathcal{L}$ is reduced, $\mathrm{path}_{(F,u)}(D')$ is a valid path address, and so $w = \mathrm{path}_{(\Gamma,v)}$ is also a valid path address. Now, for any direction d, if $A(\mathrm{path}_{(F,u)}(D'), d) = \texttt{forbidden}$, we know that w is not adjacent to any d-edge in Γ. In particular, $\mathrm{path}_{(\Gamma,v)}(D)$ is not a valid path address, and therefore we have $(\Gamma,v)|_{\mathcal{L}} = F$. □

As an immediate corollary, we get:

Corollary 4.27

The following are equivalent:

- X is a graph subshift of finite type.
- There exists a finite family of annotated patterns $\mathcal{F}$ such that

$$X = \left\{\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}} \mid \forall v \in V(\Gamma), \exists (F,A) \in \mathcal{F}, (\Gamma,v)\big|_{\mathcal{L}} = (F,A)\right\}$$

In particular, we can define graph subshifts in a more informal way, talking about edges or vertices forcing or preventing some other edges of vertices from appearing.

We also define an operation on graphs and their subgraphs, which will be useful to define patterns in a clearer way later on:

Definition 4.28

Let G be an induced subgraph of some $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. We define $\mathrm{Annot}_\Gamma(G)$ the annotated graph (G,A), where A is the following annotation function. For any $v \in G$ and $d \in \bar{\mathcal{D}}$ such that v is not adjacent to a d-edge in G,

- If v is adjacent to a d-edge in Γ, set $A(v,d) = \texttt{required}$.
- Otherwise, set $A(v,d) = \texttt{forbidden}$.

The next lemma shows that it is not abusive those spaces *subshifts*, as they satisfy the same topological conditions than what we usually call subshifts on $\mathbb{Z}^d$ or on arbitrary groups:

Lemma 4.29

For any finite alphabet $\mathcal{A}$ and family $\mathcal{F}$, the subshift $X_{\mathcal{F}}$ is closed for the topology induced by d, and invariant by any $\sigma_d, d \in \mathcal{D}$.

Proof. The shift invariance is clear from the definition: the fact that no graph $H \in \mathcal{F}$ appears as a subgraph of (Γ,u) does not depend on the base point u. It is also closed for the same reason that $\mathbb{Z}^d$ subshifts are closed: let $(\Gamma_i,u_i)_{i\in\mathbb{N}}$ be a converging sequence of pointed graphs of $X_{\mathcal{F}}$, with limit $(\Gamma,u) \in \mathcal{G}_{\mathcal{A},\mathcal{D}}$. Suppose that there is some $H \in \mathcal{F}$ such that $H \sqsubseteq \Gamma$, and let then r be large enough so that $\mathcal{B}_r((\Gamma,u))$ contains H. Then, by definition of (Γ,u), there exists n so that $\mathcal{B}_r(\Gamma_n,u_n) = \mathcal{B}_r((\Gamma,u))$. In particular, $H \sqsubseteq \Gamma_n$, which is a contradiction. □

We stop here for now our discussion about graph subshifts, and turn to the second major component object of this chapter, namely substitutions.

4.3.4 Substitutive graphs and Lindenmayer systems

The presentation adopted here for substitutions is the one of [Kna24], and is based on Lindenmayer systems. The goal of this section is not to make an overview of the rich literature about Lindenmayer systems, but to present this specific new and general framework, encompassing other graph-rewriting systems, and show how the subshifts defined on this kind of graph exhibit strong similarities with the more classical definition of substitutive subshifts exposed in Section 4.1 and Section 4.2.3. In particular, we will show that in some cases, the graphs generated by this kind of Lindenmayer system can be viewed as particular cases of graph subshifts as defined in Section 4.3.2.

Substitutions will be denoted $\mathfrak{s}$. As before, we consider graphs with labels on their vertices and edges, respectively in $\mathcal{C}_\mathfrak{s}$ and $\mathcal{D}_\mathfrak{s}$, and that edges are oriented. We distinguish a special vertex type, called the **axiom** of $\mathfrak{s}$ and denoted by $\bullet_\mathfrak{s} \in \mathcal{C}_\mathfrak{s}$. We also identify this type with the graph containing a single vertex of type $\bullet_\mathfrak{s}$ and no edges.

We now explain how one defines a substitution operation on graphs. The general idea is that vertices are to be replaced by disjoint graphs, depending only on the type of the vertex; and edges between two vertices u, v are to be replaced by a set of edges between the two graphs obtained by substituting u and v independently, where this set depends only on the direction of the edge and on the types of its endpoints u and v.

Definition 4.30: Vertex substitution

A **vertex substitution** is a set of pairs, or **rules**, (v, G) where $v \in \mathcal{C}$ and $G \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. We also denote $v \rightsquigarrow G$ the pair (v, G). If for all $v \in \mathcal{C}$ there is a single $G \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ such that $v \rightsquigarrow G$, we say that this vertex substitution is **vertex deterministic**.

Definition 4.31: Edge substitution

An **edge substitution** is a set of **rules** of the form $((u \rightsquigarrow G_u, d, v \rightsquigarrow G_v), \mathrm{E})$, where:

- $u \rightsquigarrow G_u$ and $v \rightsquigarrow G_v$ are vertex substitution rules
- $d \in \mathcal{D}$ is a direction
- $\mathrm{E} \subseteq V(G_u) \times V(G_v) \times \bar{\mathcal{D}}$, is to be viewed as a set of labeled edges between G_u and G_v, both orientations being possible.

We denote this rule by $(u \rightsquigarrow G_u) \xrightarrow{d} (v \rightsquigarrow G_v) \rightsquigarrow E$. If for all pairs of vertex rules P_1, P_2 and direction d, there is a single edge substitution rule $P_1 \xrightarrow{d} P_2 \rightsquigarrow E$, we say that this substitution is **edge deterministic**.

Notation. We use a few notations to lighten the presentation:

- We note $\mathfrak{s}_V$ the multi-valued map $\mathcal{C} \rightrightarrows \mathcal{G}_{\mathcal{C},\mathcal{D}}$ such that $v \rightsquigarrow \mathfrak{s}_V(v)$ is a vertex substitution rule. We call left-hand side and right-hand side of this rule respectively v and $\mathfrak{s}_V(v)$.

- We note $\mathfrak{s}_E$ the multi-valued map such that $P_1 \xrightarrow{d} P_2 \rightsquigarrow \mathfrak{s}_E(P_1, P_2, d)$ is an edge substitution rule. We call left-hand side and right-hand side of this rule respectively $P_1 \xrightarrow{d} P_2$ and $\mathfrak{s}_E(P_1, P_2, d)$.

- If the substitutions is vertex deterministic, we write $\mathfrak{s}_E(u, v, d)$ for $\mathfrak{s}_E((u \rightsquigarrow \mathfrak{s}_V(u)), (v \rightsquigarrow \mathfrak{s}_V(v)), d)$. If furthermore there is a single vertex type, $|\mathcal{C}| = 1$, then we write $\mathfrak{s}_E(d)$.

In order to simplify the exposition, we will assume from now on that the right-hand side of the vertex rules are always graphs with no non-trivial automorphism. This will be useful to explain how the edges E, E' arising from substituting two intersecting edges, $(u \rightsquigarrow G_u) \xrightarrow{d} (v \rightsquigarrow G_v) \rightsquigarrow E$ and $(u \rightsquigarrow G_u) \xrightarrow{d'} (v' \rightsquigarrow G_{v'}) \rightsquigarrow E'$, are to be chosen consistently and without ambiguity in the graph G_u originating from the substitution of u. Another solution, chosen in [Kna24], is to work with *concrete* graphs, with vertices being subsets of a fixed universe. One can then without ambiguity refer to particular vertices, and "forget" the concrete labeling later on. In practical cases, this will not make a difference for us. In particular, this is not in contradiction with *e.g.* the case of substitutions in $\mathbb{Z}^2$. Indeed, although substitutions defined in Section 4.2.1 use squares as supports, which clearly have automorphisms such as rotations and reflections, the substitutions themselves are not defined up to rotation: in a sense, every point in the square can tell apart its up neighbour from its right one: in order to model $\mathbb{Z}^2$ substitutions in this framework, we would need to use edge directions (typically `up` and `right`), see for example Figure 4.10 for a variant on the usual $\mathbb{Z}^2$ substitution which uses more than those two directions.

We are ready to give the definition of a graph substitution. This is a restriction of a slightly more general class, known as "0L graph grammars" – the "L" stands for Lindenmayer, and the 0 indicates that the substitution rule is context-free, in other words, vertices and edges are all substituted independently. In all the statement and unless specified otherwise, we do not assume that substitutions are vertex or edge deterministic. However, in order to lighten the notations, we might sometimes talk about "the" image of some vertex or graph by $\mathfrak{s}$ and state properties about it. In that case, the statement must be understood as a universal quantification, and what we really mean is that any such graph satisfies the property:

Definition 4.32: Graph substitution [Kna24]

A **graph substitution** $\mathfrak{s}$ on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$ is a triplet $(\mathfrak{s}_V, \mathfrak{s}_E, \bullet_{\mathfrak{s}})$, such that:

- For all $G \in \mathfrak{s}_V(\mathcal{C})$, G admits no-non trivial automorphism and is connected.

- For every edge substitution $E = \mathfrak{s}_E(c \rightsquigarrow G_c, c' \rightsquigarrow G_{c'}, d)$, the vertices $\pi_{V(G_c)\times\mathcal{D}}(E)$ only depend on d and $c \rightsquigarrow G_c$, and $\pi_{V(G_{c'})\times\mathcal{D}}(E)$ only depend on d and $c' \rightsquigarrow G_{c'}$.

- For any $c, c_1 \neq c_2 \in \mathcal{C}$, $d, d_1 \neq d_2 \in \mathcal{D}$, any vertex rule $c \rightsquigarrow G, c_1 \rightsquigarrow G_1, c_2 \rightsquigarrow G_2$, any $u \in V(G)$, and any edge rule $(c \rightsquigarrow G) \xrightarrow{d_1} (c_1 \rightsquigarrow G_1) \rightsquigarrow E_1$, $(c \rightsquigarrow G) \xrightarrow{d_2} (c_2 \rightsquigarrow G_2) \rightsquigarrow E_2$, there is at most one d-edge adjacent to u in all of G, E_1, and E_2.

- For any edge $c \rightsquigarrow G_c, c' \rightsquigarrow G_{c'}, d \in \mathcal{D}$, there is an edge substitution rule $(c \rightsquigarrow G_c) \xrightarrow{d} (c' \rightsquigarrow G_{c'}) \rightsquigarrow E$.

The second item means that in some sense, each edge could be cut in two, and substituted separately on both ends. In the case of a single image in each $\mathfrak{s}_E$, that is, $\mathfrak{s}_E$ is a deterministic map, this would correspond in the language of [Kna24] to a slight restriction of complete eD0L grammars, but we do not need to restrict ourselves to edge-deterministic graph substitutions. The consequence is that in the proof of Mozes theorem, instead of being able to hardcode some value, we will need to perform extra checks which roughly correspond to the finite state automaton constructed in *e.g.* [Goo98, Section 1.2]. The lengthy third condition ensures that we can always compute the image of a graph, as a consequence of Lemma 4.36, as will be shown below. For now, we still need to explain how to actually use those graph substitutions to substitute a graph.

In fact, in most of our examples, the substitution will indeed be edge-deterministic.

We first give an example of a classical graph substitution in Figure 4.6, called the Sierpinski triangle graph substitution. This example is an easy one, as it has a single vertex type (and therefore $\mathcal{C}_\mathfrak{s} = \{\bullet_\mathfrak{s}\}$) and is completely deterministic with a single edge per edge substitution. In our figures, we will usually represent the type of a vertex by a geometrical shape, and the type of an edge by its colour.

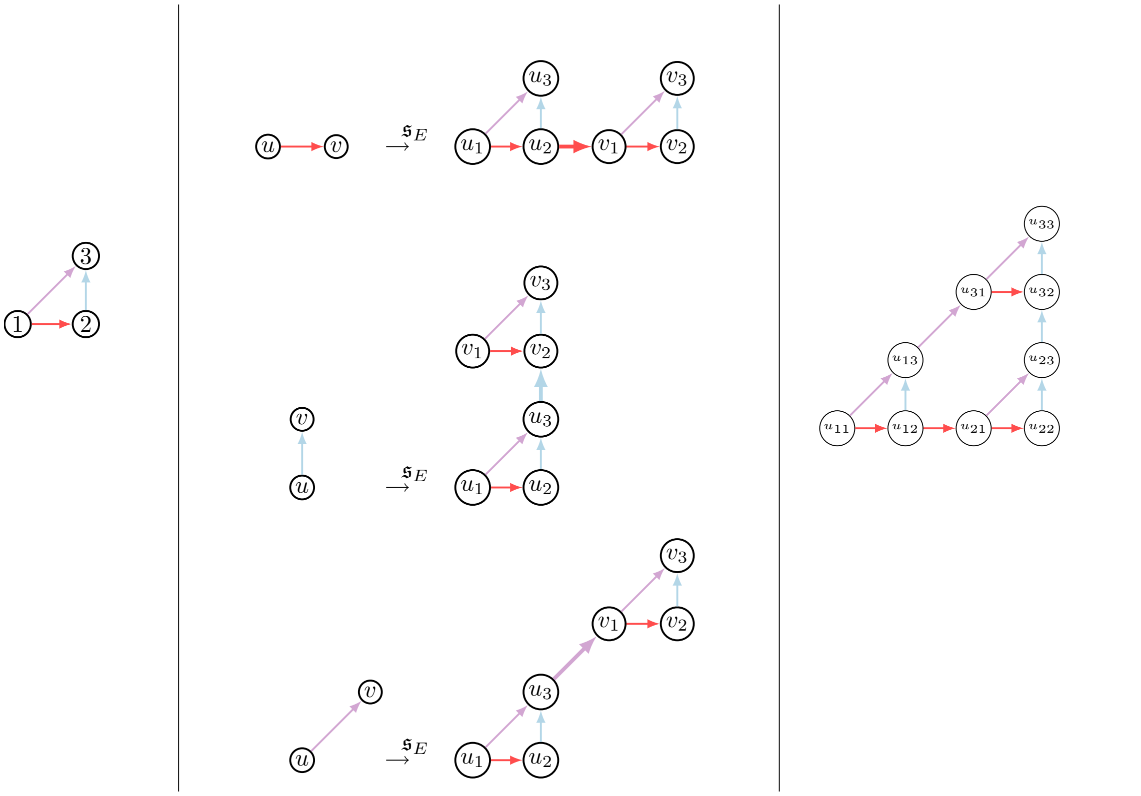

Figure 4.6: Example of the Sierpinski triangle graph substitution

The third sub-figure of Figure 4.6 shows an additional iteration of the substitution, which we have not yet defined. Hopefully, one gets from this picture an intuitive idea of the "semantics" of graph substitutions. We now properly define what it means to *apply* a substitution $\mathfrak{s}$ to a graph, and how we can then iterate to define $\mathfrak{s}^n$ for $n > 0$.

Definition 4.33: Substitution application

Let $\Gamma = (V, E, \lambda_V, \lambda_E) \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ and $\mathfrak{s} = (\mathfrak{s}_V, \mathfrak{s}_E, \bullet_\mathfrak{s})$ be a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. We define the graph $\Gamma' = (V', E', \lambda_{V'}, \lambda_{E'}) = \mathfrak{s}(\Gamma)$ as follows:

- $V' = \bigsqcup_{v\in V} V(\mathfrak{s}_V(\lambda_V(v)))$
- $E' = \bigsqcup_{v\in V} E(\mathfrak{s}_V(\lambda_V(v)))$
  $\sqcup \bigsqcup_{((u,v),d)\in E} \left\{x \xrightarrow{d'} y \mid (x,y,d') \in \mathfrak{s}_E(u,v,d)\right\}$
- $\lambda_{V'}, \lambda_{E'}$ are inherited from the labeling obtained from the choices of rules made for $\mathfrak{s}_V$ and $\mathfrak{s}_E$.

We note $\Gamma \rightsquigarrow \Gamma'$ to denote the fact that Γ' is obtained from Γ by applying the substitution rule. Moreover, as the substitution is computed by replacing independently each vertex and adding edges between the resulting graphs using $\mathfrak{s}_E$, for $u \in V(\Gamma)$, assuming Γ' has no automorphisms, we can refer to *the* subgraph $\mathfrak{s}_V(u) \sqsubseteq \Gamma'$.

In the definition of the edge set of $\mathfrak{s}(\Gamma)$, we can distinguish two parts:

- The first set corresponds to the edges of the graphs obtained by substituting the vertices $V(\Gamma)$. Following the terminology of [Kna24], we call those edges *innate*.
- The second set of edges corresponds to the ones obtained from substituting an edge $u \xrightarrow{d} v$ already present in Γ. Those edges are called *inherited*.

As an easy lemma, we have the following result:

Lemma 4.34: Increasing substitutions

Let $G \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ be a finite graph and $\mathfrak{s}$ a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. Then $|V(G)| < |V(\mathfrak{s}(G))|$.

Proof. This is simply a consequence of Definition 4.32: as the substitution is applied separately on each vertex, we have $|V(\mathfrak{s}(G))| = \sum_{v\in V(G)} |\mathfrak{s}_V(\lambda_V(v))|$. But by definition, the image of a vertex cannot be reduced to a single vertex. □

Corollary 4.35

Let $G \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ be a finite graph, $\mathfrak{s}$ a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, and $u, v \in V(G)$. Denote $G' = \mathfrak{s}(G)$. Then, for any $u' \in \mathfrak{s}_V(u) \sqsubseteq G', v' \in \mathfrak{s}_V(v) \sqsubseteq G'$, we have $d_{G'}(u', v') \geq d_G(u, v)$.

An important consequence of Definition 4.32 is that those substitutions can be iterated. This is an immediate consequence of the following lemma:

Lemma 4.36: Local finiteness

For any $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$, $d \in \mathcal{D}$ and any vertex $u \in V(\mathfrak{s}(\Gamma))$, u is adjacent to at most one incoming edge with direction d, and to at most one outgoing edge with direction d. In particular, $\mathfrak{s}(\Gamma) \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$.

Proof. Let $v \in \Gamma$ and consider a vertex $u \in \mathfrak{s}_V(v) \sqsubseteq V(\mathfrak{s}(\Gamma))$. Then u is adjacent to two types of edges: inherited edges, coming from edge substitution of an edge $(v, w) \in E(\Gamma)$, and innate edges in $E(\mathfrak{s}(v))$. We explicitly required in Definition 4.32 that edge substitution rules could never add an inherited edge with the same direction as an already-existing innate edge, and moreover, no two inherited edges adjacent to a given vertex u could arise from the substitution of edges with distinct directions d, d'. In particular, as v is adjacent to at most one edge of each direction and orientation, this property is preserved for $u \in \mathfrak{s}_V(v)$. □

Corollary 4.37: Substitution iteration

For any graph substitution $\mathfrak{s}$ over $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, and any graph $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$, the graphs $(\mathfrak{s}^n(\Gamma))_{n\geq 0}$ are well-defined and all in $\mathcal{G}_{\mathcal{C},\mathcal{D}}$.

We also provide another point of view on graph substitutions, using annotated graphs (see Definition 4.24):

Definition 4.38

Let $\mathfrak{s} = (\mathfrak{s}_V, \mathfrak{s}_E, \bullet_{\mathfrak{s}})$ be a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. We define the associated **annotated graph substitution** $\mathfrak{s}_A$, the map sending an annotated graph $(\Gamma = (V, E, \lambda_V, \lambda_E), A)$ to the graph $(G', E', \lambda'_V, \lambda'_E, A')$ defined by:

- $(G', E', \lambda'_V, \lambda'_E) = \mathfrak{s}(\Gamma)$, that is, we apply $\mathfrak{s}$ as usual on the non-annotated graph.
- For $v \in V(\Gamma), d \in \bar{\mathcal{D}}$ with $\mathfrak{s}_V(v) = G \sqsubseteq \Gamma'$, if $A(v, d) = \texttt{required}$ then for any edge $(e = (v', \_, d')) \in \mathfrak{s}_E(v \rightsquigarrow G, \_, d)$ we set $A'(v', d') = \texttt{required}$.

The map $\mathfrak{s}_A$ is well-defined: indeed, by Definition 4.32, if e is an edge adjacent to v, the endpoints and direction of the edges obtained by substituting e adjacent to $\mathfrak{s}(v)$ do not depend on the other endpoint of e. In particular, we can substitute an "annotated edge".

Skeletons, borders, meta-tiles

In order to prove our variant of Mozes theorem, we will follow the general strategy of *e.g.* [Moz89], [Goo98] or [FO10]: we partition the edges of the meta-tiles in two sets, that we call, in accordance with [BS16], **borders** and **skeletons**. Skeletons are used to synchronize the information between the siblings in a common meta-tile, while borders will be part of higher-level skeletons. If the graphs are sufficiently connected, and if the skeletons and borders are well-defined, each edge will be used to synchronize a finite number of meta-tiles and a fixed-point argument then shows that the entire tilings must obey the recursive structure given by the substitution.

Definition 4.39: Skeletons and borders

For each derivation rule of the form $c \rightsquigarrow \Gamma' \rightsquigarrow \Gamma$:

- For each $v \in \Gamma'$, fix an arbitrary vertex $\text{meta}(v) =$

$\text{meta}(v, \Gamma', \mathfrak{s}_V(v)) \in \mathfrak{s}_V(v) \sqsubseteq \Gamma$, depending only on $\Gamma', \mathfrak{s}_V(v)$ but not on the entire Γ.

- For each edge $e = (v, v') \in \Gamma'$, fix an arbitrary simple path $\text{meta}(e) = \text{meta}(e, \Gamma', \mathfrak{s}_E(e)) \subset E(\Gamma)$ from $\text{meta}(v)$ to $\text{meta}(v')$, such that $\text{meta}(e) \cap \mathfrak{s}_V(v)$ (resp. $\text{meta}(e) \cap \mathfrak{s}_V(v')$) depends only on v and Γ' (resp. v' and Γ').

We then call **border** of this 2-derivation the set of paths $B_{c,\Gamma',\Gamma} = \bigcup_{e \in E(\Gamma')} \text{meta}(e)$, and **skeleton** the graph obtained by deleting the edges of $B_{c,\Gamma',\Gamma}$ from Γ,

$$S_{c,\Gamma',\Gamma} = \Gamma \setminus B_{c,\Gamma',\Gamma}$$

We will simply write B, S in the case of a single vertex type, the axiom $\bullet_{\mathfrak{s}}$, and deterministic substitution. This is the case for the Sierpinski triangle of Figure 4.6 and the weak-grid of Figure 4.7.

Figure 4.7 shows the skeleton of the weak-grid substitution. The border is obtained by considering only the edges that do not belong to S.

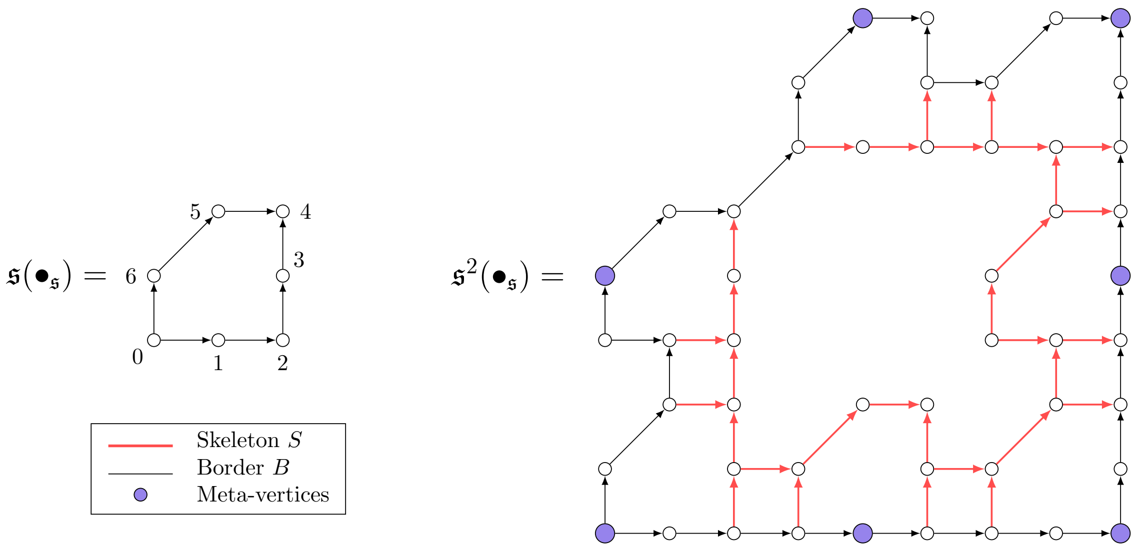

Figure 4.7: The 2-meta-tile, its skeleton S_2 and its meta-vertices. Note that the dashed edge $(46, 35)$ does not belong to the skeleton, but to the border. Indeed, if we were to consider this edge as part of the skeleton (this would require another definition than the one given in Definition 4.39), this edge could then belong to infinitely many skeletons of "higher-order meta-tiles", which means that we would need graph subshifts with infinite alphabets in our proof of Mozes theorem.

We also need the notion of **facet**. A facet will be the analogous of edges, for meta-tiles, or more precisely the edges endpoints.

Definition 4.40: Facet

Let $\mathfrak{s}$ be a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, and $c \rightsquigarrow G$ a vertex substitution rule of $\mathfrak{s}$. Let $(c \rightsquigarrow G) \xrightarrow{d} (c' \rightsquigarrow G') \rightsquigarrow E$ be an edge substitution rule of $\mathfrak{s}$.

We call d-facet of G the set

$$F_d(G) = \pi_{V(G)}(E)$$

the endpoints in G of inherited edges obtained from substituting d.

The facets are well-defined: indeed, by Definition 4.32, the set $F_d(G)$ only depends on c and d in the previous definition, and not on c' or on the specific edge substitution rule.

Let $G_1, G_2 \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ be two finite graphs, such that $(c \rightsquigarrow G) \xrightarrow{d} (c' \rightsquigarrow G') \rightsquigarrow E$ is an edge substitution rule, and $\Gamma \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ containing both G_1 and G_2. We say that G_1 and G_2 **match along a d-facet** (in Γ) if the subgraph induced by $\Gamma|_{G_1 \cup G_2}$ is isomorphic to $\mathfrak{s}(c \xrightarrow{d} c')$. For annotated graphs $(G_1, A_1), (G_2, A_2)$, they match along their d-facet if any edge e of $\mathfrak{s}(c \xrightarrow{d} c')$ is either induced in $\Gamma|_{G_1 \cup G_2}$, or there exists $v_1 \in G_1, v_2 \in G_2$ such that $A_1(v_1, \lambda_E(e)) = A_2(v_2, \lambda_E(e)^{-1}) = \texttt{required}$ (that is, there are annotations on both "sides" of the missing edge).

In order to see how facets relate to the previously defined skeletons and borders, we need to impose some conditions on how we choose to determine a substitution's skeleton. Indeed, facets are well-defined for any substitution, but there are multiple options for borders, and therefore skeletons.

Definition 4.41: Quasi-connected substitution

Let $\mathfrak{s}$ be a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. We say that it is **quasi-connected** if we can define all the borders $B_{c,\Gamma',\Gamma}$ in such a way that:

- the skeleton $S_{c,\Gamma,\Gamma'}$ is a connected graph, up to the inherited edges. That is, the union of $S_{c,\Gamma,\Gamma'}$ and all the inherited edges is a connected subgraph of Γ.
- For any $v \in V(\Gamma')$, the skeleton intersects each $G = \mathfrak{s}_V(v) \sqsubseteq \Gamma$ along at least an edge.
- For any $u, v \in E(\Gamma')$, the border $\mathrm{meta}(e) \subset E(\Gamma)$ intersects $\mathfrak{s}_V(v) \sqsubseteq \Gamma$ along at least an edge.
- For any edge substitution $(c \rightsquigarrow \Gamma') \xrightarrow{d} \_ \rightsquigarrow E$, for any $v \in F_d(\Gamma')$ and $(v \xrightarrow{d'} \_) \in E$, for any $G = \mathfrak{s}_V(v)$, we have $\mathrm{meta}(v, \Gamma', G) \cap F_{d'}(G) \neq \emptyset$.

The last condition informally ensures that in any meta-tile $\Gamma_n = \mathfrak{s}^n(v)$ with $\Gamma_1 = \mathfrak{s}(v)$, for each $v_i \in \Gamma_1$, we can uniquely identify a vertex whose neighbourhood in Γ_n "looks like" the one of v_i in Γ_1. We give in Figure 4.8 an example of a choice of meta-vertices for which the graph substitution does not respect the last condition of Definition 4.41.

The quasi-connectivity condition implies in particular that for any $v' \in V(\Gamma')$ and any direction d of an edge adjacent to v' in Γ', $F_d(\mathfrak{s}(v')) \cap S_{c,\Gamma',\Gamma} \neq \emptyset$. In Figure 4.7, this quasi-connectivity property can be seen with the dashed edge, which form a connected graph together with the plain red edges of the skeleton.

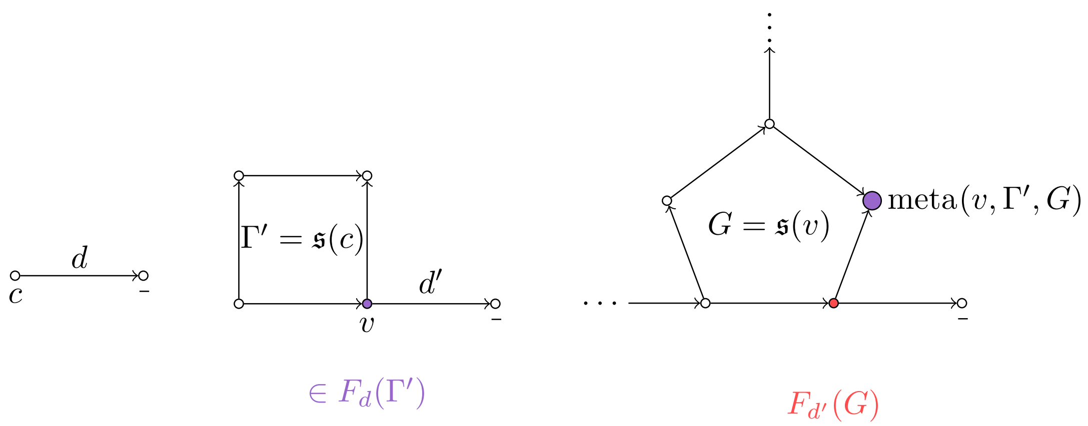

Figure 4.8: An example of substitution, where the images of individual vertices are not deterministic, which is not quasi-connected for this choice of meta-vertices.

Lemma 4.42

Let $\mathfrak{s}$ be a quasi-connected substitution. Then for any derivation $c \rightsquigarrow \Gamma' \rightsquigarrow \Gamma$, and for any $v \in V(\Gamma')$, there exists $u \in \mathfrak{s}_V(v) \cap S_{c,\Gamma',\Gamma} \cap B_{c,\Gamma',\Gamma}$.

Proof. This is immediate by the definition of the skeleton: by definition, every edge of $\mathfrak{s}_V(v)$ is either in the skeleton or in the border. There exists both an edge $e_B \in \mathfrak{s}_V(v) \cap B_{c,\Gamma',\Gamma}$ and $e_S \in \mathfrak{s}_V(v) \cap S_{c,\Gamma',\Gamma}$ by quasi-connectivity, and so there must an edge with an endpoint in both subgraphs, for otherwise they would be disconnected but $\mathfrak{s}_V(v)$ is connected. □

The last point of Definition 4.41 will be useful for the following reason:

Lemma 4.43

Let $\mathfrak{s}$ be a quasi-connected substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, and $c \xrightarrow{d} c'$ some edge for $c, c' \in \mathcal{C}, d \in \mathcal{D}$. Then, for any $G = \mathfrak{s}^2(c \xrightarrow{d} c')$, there exists meta-vertices $v \in \mathfrak{s}^2(c) \sqsubseteq G$, $v' \in \mathfrak{s}^{2(c')} \sqsubseteq G$, and an edge $e' = (v, v')$.

Proof. Consider any edge e_1 of direction d' between $v_1 = \mathfrak{s}(c)$ and $v_2 = \mathfrak{s}(c')$. It suffices to consider the edge obtained in the last point of Definition 4.41 applied to $(c \rightsquigarrow \mathfrak{s}(c)) \xrightarrow{d} (c' \rightsquigarrow \mathfrak{s}(c'))$ and $v_1 \xrightarrow{d'} v_2$. □

This might seem like a restrictive condition, but it is in fact satisfied in many cases. We say that a graph G is k-vertex-connected (resp. k-edge connected) if for any vertices $v_1, \ldots, v_k \in V(G)$ (resp. edges $e_1, \ldots, e_k \in E(G)$), the graph $G \setminus \{v_1, \ldots, v_k\}$ remains connected (resp. $G \setminus \{e_1, \ldots, e_k\}$ remains connected). It is easy to see that being k-vertex connected implies being k-edge connected: intuitively, this is because it is "worse", as far as paths are concerned, to remove an entire vertex and therefore all its adjacent edges, than simply removing a single edge.

While apparently not very restrictive for small values of k, this property implies strong conditions on the structure of the graph:

Theorem 4.44: Menger Theorem

Let G be a k-edge-connected graph, and $u, v \in V(G)$. Then there exists k edge-disjoint paths between u and v.
If G is k-vertex-connected, and u, v are not adjacent, then there exist k vertex-disjoint paths between u and v.

This property is in fact sufficient in numerous cases for a substitution to be quasi-connected. More precisely:

Proposition 4.45

For any 2-vertex connected graph G, there exists an edge substitution $\mathfrak{s}_E$ such that $\mathfrak{s} = (\mathfrak{s}_V : \bullet_{\mathfrak{s}} \mapsto G, \mathfrak{s}_E, \{\bullet_{\mathfrak{s}}\})$ is quasi-connected.

Proof. For any edge $(u, v) \in E(G)$, set $\mathfrak{s}_E(u \xrightarrow{d} v) = (G_u \sqcup G_v \sqcup (v \in G_u, u \in G_v))$, *i.e.* replace the d-edge (u, v) by a d-edge from the copy of v in $\mathfrak{s}_V(u)$ to the copy of u in $\mathfrak{s}_E(v)$. Defining the meta-vertices of Definition 4.39 as the vertices $u \in \mathfrak{s}_V(u)$ for all $u \in V(G)$, and the border by the set of edges $(u, v) \in E(\mathfrak{s}_V(u))$ and the inherited edges, we now have the conditions of Definition 4.41. Indeed, removing the single vertex meta(u) in $\mathfrak{s}_V(u)$ does not disconnect the graph, which is 2-vertex-connected. □

This is in general not the only way to define the edge substitution in order to obtain a quasi-connected substitution. For example, Figure 4.9 shows a deterministic variant of the grid substitution, which produces non-isomorphic graphs (for example, there exist adjacent vertices of degree 3, which is not the case for the grid substitutions considered until here).

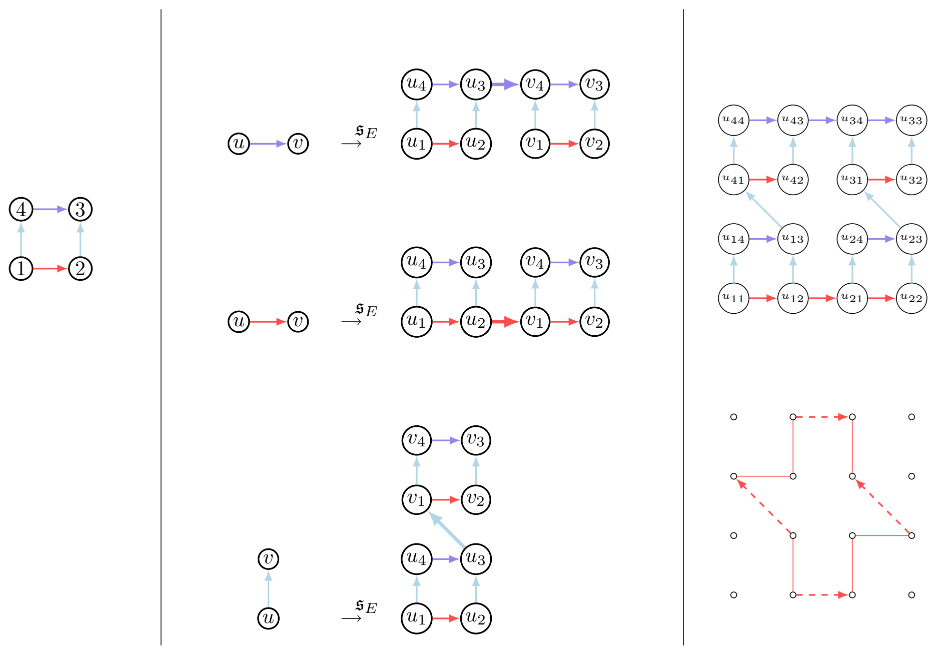

Figure 4.9: A variant of the grid substitution: the last subfigure shows the skeleton, which is connected up to the inherited dashed edges.

It is in fact a subgraph of the dual graph of the (informally defined) substitution pictured in Figure 4.10, where cells are considered adjacent if they share a vertex.

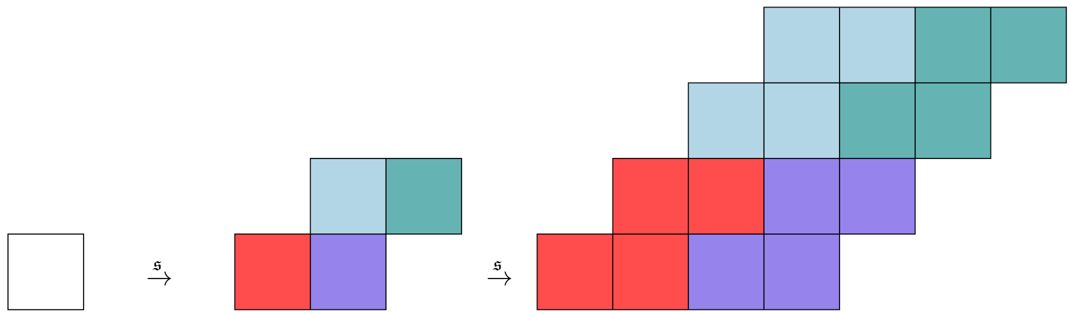

Figure 4.10: A substitution on $\mathbb{Z}^2$, a combinatorial version of which is presented in Figure 4.9

A more systematic study of this kind of constant-shape substitution, mainly from the point of view of dynamical systems, can be found in [Cab23], where the author characterizes many properties of such substitutive systems (directions of determinism, their periodicity properties, algebraic representations ...).

Sheets and subgraphs

An unavoidable difficulty that we will encounter is that, using graph subshifts of finite type, we are unable to distinguish between a graph and some of its quotients, in a sense that can be made precise. Intuitively, as the graphs are only defined by forbidding (or allowing) graphs smaller than a certain uniform size, whenever we follow a path p starting from some vertex $u \in V(\Gamma)$ which is longer than this bound, we cannot detect if p is a cycle (and therefore $\mathrm{path}_{(\Gamma,u)}(p) = u$), or if it simply ends on a vertex whose neighbourhood is the same as u's. As we will not be studying Cayley graphs in the following sections, we define a weaker notion than the quotient defined in [ADG23, Section 5], and work in a less algebraic but more combinatorial setting. We define a notion of **sheet**, and of **sheeted graph**: this is the kind of graph on which we will be able to ensure that the meta-Mozes theorem (Chapter 4) holds, rather than precisely the substitutive graph subshifts defined in Definition 4.50.

Definition 4.46: Sheet

Let $\Gamma, G \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$. We say that G contains a Γ-sheet at $u \in V(G)$ if there exists a connected subgraph $H \sqsubseteq G$ containing u, and $\psi\colon H \to \Gamma$ an injective graph morphism, with $\psi(V(H)) = V(\Gamma)$ and such that for any $u \xrightarrow{d} v \in E(\Gamma)$, there exists $\psi^{-1}(u) \xrightarrow{d} \_ \in E(G)$ and $\_ \xrightarrow{d} \psi^{-1}(v) \in E(G)$ (but not necessarily in $E(H)$).
For a class $\mathcal{G} \subset \mathcal{G}_{\mathcal{C},\mathcal{D}}$, we say that G is $\mathcal{G}$-sheeted if for any $u \in V(G)$, there exists $\Gamma \in \mathcal{G}$ such that G contains a Γ-sheet at u.

Another equivalent definition is given by the following proposition:

Proposition 4.47

G contains a Γ-sheet at $u \in V(G)$ if and only if there exists a spanning tree $T \sqsubseteq \Gamma$ and an injective morphism $\phi\colon T \to G$ such that for $u \xrightarrow{d} v \in E(\Gamma) \setminus E(T)$, $\phi(u)$ is adjacent to some edge with direction d in G.

Proof. Suppose that G contains a Γ sheet at some vertex u. Let $H, \psi\colon H \hookrightarrow \Gamma$ be as in Definition 4.46. Then, let $T \sqsubseteq \Gamma$ be any spanning tree of $\psi(H)$. It is clear that T satisfies the required conditions.

For the other direction, suppose that we have $T \sqsubseteq \Gamma, \phi\colon T \hookrightarrow G$ as in Proposition 4.47. Then $H = \phi^{-1}(T), \psi\colon u = \phi^{-1}(v) \mapsto v$ satisfy Definition 4.46. □

This notion is weaker than the one of covering (see Definition 3.18 for a general definition): indeed, if G is $\mathcal{G}$-sheeted, although each vertex $u \in G$ has to belong to some sheet Γ_u, there is no reason for $\Gamma_u, \Gamma_v \in \mathcal{G}$ containing distinct u, v to be isomorphic. In particular, we cannot always define a single morphism $\psi^*\colon G \to \Gamma$, and there need not be a global section. The condition that if $u \xrightarrow{d} v \in E(\Gamma)$, there is an edge $\psi^{-1}(u) \xrightarrow{d} \_ \in E(G)$, ensures that all the vertices of H have the same neighbourhoods in G as their image in Γ. In particular, a graph G minus some of its edges cannot have a G-sheet. An example of graph admitting (finite) sheets at every point which are all isomorphic is depicted in Figure 4.11.

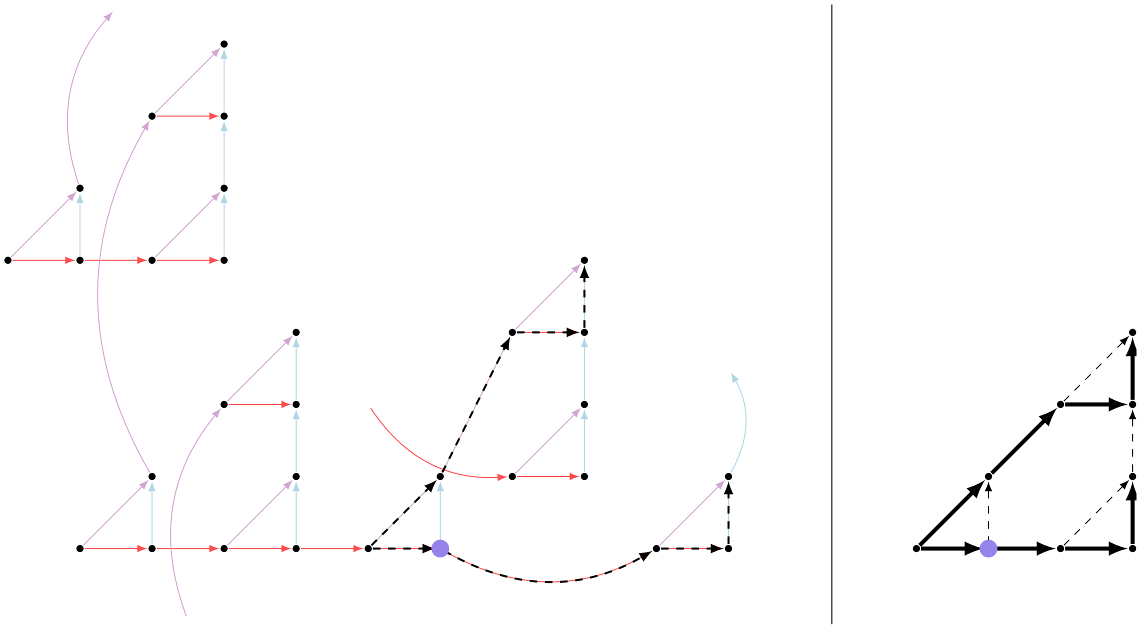

Figure 4.11: Part of a graph admitting a $\mathfrak{s}^2(\bullet_{\mathfrak{s}})$-sheet at every point, where $\mathfrak{s}$ is the Sierpinski substitution of Figure 4.6. An example of a sheet, represented as a spanning tree of $\mathfrak{s}^2(\bullet_{\mathfrak{s}})$, is represented on the right for the large purple vertex. The morphism is shown as dashed black edges on the left picture.

We give another definition of sheets, using languages as in Definition 4.17. A Γ-sheet in G can then be viewed as a way to consistently pick a path to each vertex of Γ from its base-point, the sheet in G then being the subset of $V(G)$ reached by following these paths.

Lemma 4.48: Sheets and languages

Let $\Gamma, G \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ be two pointed graphs, $u \in V(G)$. Then, G contains a Γ-sheet at u if and only if there exists a prefix-stable language $\mathcal{L} \subseteq \text{Paths}(\Gamma)$, such that:

- For any $v \in \Gamma$, there exists $p \in \mathcal{L}$ such that $\text{path}_\Gamma(p) = v$.
- For each path $p \in \mathcal{L}$ and $\text{path}_\Gamma(p) \xrightarrow{d} \_ \in E(\Gamma)$, there exists an edge $\text{path}_G(p) \xrightarrow{d} \_ \in E(G)$.
- For $p_1, p_2 \in \mathcal{L}$, p_1 is a valid path in (G, u) and $\text{path}_\Gamma(p_1) = \text{path}_\Gamma(p_2) \iff \text{path}_G(p_1) = \text{path}_G(p_2)$
- For any $p = p' \cdot d \in \mathcal{L}$ with $d \in \bar{\mathcal{D}}$, $\lambda_V(\text{path}_G(p)) = \lambda_V(\text{path}_\Gamma(p))$ and $\lambda_E(\text{path}_G(p'), \text{path}_G(p)) = \lambda_E(\text{path}_\Gamma(p'), \text{path}_\Gamma(p))$.

Proof. Suppose that G contains a Γ-sheet H at u, with $\psi \colon H \hookrightarrow \Gamma$ the corresponding injective morphism. Let $\mathcal{L} = \text{Paths}(H, u)$, and let $p = p' \cdot d \in \mathcal{L}$. Let $e \in E(H) = (\text{path}_H(p'), \text{path}_H(p))$. Then, as ψ is a morphism, we immediately have $\lambda_V(\text{path}_G(p)) = \lambda_V(\text{path}_\Gamma(p))$ and $\lambda_E(e) = \lambda_E(\text{path}_\Gamma(p'), \text{path}_\Gamma(p))$. For $p_1, p_2 \in \mathcal{L}$ such that $\text{path}_\Gamma(p_1) = \text{path}_\Gamma(p_2)$, as ψ is injective we also have $\text{path}_H(p_1) = \text{path}_H(p_2)$. Let now $v \in V(\Gamma)$. As $\psi(V(H)) = V(G)$ by definition of a sheet, there exists $v' \in H$ with $\psi(v') = v$. Pick any path p such that $\text{path}_H(p) = v'$. Then by the previous point we have $\text{path}_\Gamma(p) = v$.

For the other direction, suppose that we have a prefix-stable language $\mathcal{L}$ satisfying the properties of Lemma 4.48. Let $H \sqsubseteq G$ be the graph:

- $V(H) = \{\text{path}_G(p), p \in \mathcal{L}\}$
- $E(H) = \{(\text{path}_G(p'), \text{path}_G(p)) \in V(H)^2, p = p' \cdot d \in \mathcal{L}, d \in \bar{\mathcal{D}}\}$

Then we claim that H is a Γ-sheet containing u. We define the morpsism ψ as $\psi(\text{path}_H(p)) = \text{path}_\Gamma(p)$. We first need to show that ψ is well-defined, that is, it does not depend on the path p chosen to describe a vertex $v \in H$. Let then $p, p' \in \mathcal{L}$ be such that $\text{path}_H(p) = \text{path}_H(p')$. By definition of H, if p is a valid path in H, it is also valid in Γ. By assumption on $\mathcal{L}$, we have $\text{path}_\Gamma(p) = \text{path}_\Gamma(p')$. Injectivity and vertex-surjectivity follow. □

This is where the geometry, that we conveniently removed from the picture when defining substitutions in Definition 4.32, can help us in some cases. This is for example the case if we are studying the substitutions as defined by [BS16] and presented in Section 4.2.3: indeed, the fact that they are defined on the $\mathbb{Z}^2$ grid effectively means that for many properties P, in a $\mathbb{Z}^2$ graph G where all sheets satisfy P, there will in fact be a single sheet and so the graph G itself satisfies P due to the geometry of $\mathbb{Z}^2$. In the formalism of [BS16], this can be seen as the fact that the *structure* of the substitution is given by a binary, rectangular substitution over $\mathbb{Z}^2$ (which happens to be sofic by the classical Theorem 4.3), and then proving that substitutive subshifts defined using this structure are "sofic" in a stronger sense, relative to this structure, without referring to the "soficity" of the structure itself. This is akin to studying the properties of a single sheet, itself belonging to a graph on which we cannot guarantee *global* properties.

4.3.5 Sofic graphs and coloured substitutions

Coloured substitutions on graphs

The previous results showed how the substitutive *graphs* could be recursively decomposed in meta-tiles, in which one can define several objects, such as the skeletons of these meta-tiles, that admit nice recursive properties. In this sections, we explain how one can adapt the usual notions of symbolic dynamics to the case of graph colourings, and how we can use the notion of substitutive graphs to define substitutive subshifts, as well as subshifts of finite type. For now, let us focus on the definition of a **coloured substitution**:

Definition 4.49: Coloured graph substitution

Let $\mathcal{A}$ be a finite alphabet of colours, and $\mathfrak{s} = (\mathfrak{s}_V, \mathfrak{s}_E, \bullet_{\mathfrak{s}})$ a graph substitution on some class of graphs $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. A **coloured graph substitution** based on $\mathfrak{s}$ is a graph substitution $\hat{\mathfrak{s}}$ on $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ such that $\pi_{\mathcal{C}} \circ \hat{\mathfrak{s}} = \mathfrak{s} \circ \pi_{\mathcal{C}}$. In other words, for any $u \in \mathcal{C} \times \mathcal{A}$, $\pi_{\mathcal{C}}(\hat{\mathfrak{s}}_V(G_u)) = \mathfrak{s}_V(\pi_{\mathcal{C}}(G_u))$

This definition can be understood as follows: the vertices' labels of a coloured graph substitution can be decomposed in two parts, a "structural" part and a "colour" part. The structural part of the image of any vertex by the substitution only depends on the structural part of the vertex, while the colours in the image also depend on the original colour.

As before, this definition relies on the absence of non-trivial automorphisms of G: the colouring of the vertices is not ambiguous, and so we can iterate the coloured substitutions just as we did for the purely combinatorial ones: a coloured substitution is a graph substitution, and so all the definitions and propositions of Section 4.3.4 still apply.

We can now define the *set* of all the substitutive colourings associated to a coloured substitution $\mathfrak{s}_c$ using annotated graphs (see Definition 4.28):

Definition 4.50: Substitutive subshift

Let $\mathcal{A}$ be a finite alphabet of colours, $\mathfrak{s} = (\mathfrak{s}_V, \mathfrak{s}_E, \bullet_{\mathfrak{s}})$ a graph substitution on some class of graphs $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, and $\mathfrak{s}_c$ a coloured graph substitution on $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ based on $\mathfrak{s}$. We define:

$$X_{\mathfrak{s}_c} = \big\{x \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}} \,\big|\forall G \sqsubseteq x, \exists c_a \in \mathcal{C}\times\mathcal{A}, n \geq 0, \Gamma \sqsubseteq \mathfrak{s}_c^n(c_a), \\ G = \mathrm{Annot}_{\mathfrak{s}_c^{n+1}(c_a)}(\Gamma)\big\}$$

$$X_{\mathfrak{s}_c}^{\infty} = \{x \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}} \mid \forall n \geq 0, \exists y_n \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}, x \simeq \mathfrak{s}_c^n(y_n)\}$$

Note that in the definition of $X_{\mathfrak{s}_c}^{\infty}$, we do not care about the base point of x or its preimages y_n by $\mathfrak{s}_c^n$: this is the same condition as the one given in Definition 4.2, in which we consider the orbits of the points admitting arbitrarily many preimages to define $X_{\mathfrak{s}}^{\infty}$.

There is a subtlety in our definition of $X_{\mathfrak{s}_c}$: for an arbitrary subgraph $G \sqsubseteq x$, the natural restriction would be to require that there exists c_a, n, Γ with $G \sqsubseteq \Gamma$, or even $\Gamma|_{\mathrm{Paths}(G)} = G$. As discussed in Section 4.3.3, such a definition is not always satisfying. In that case, this would in fact mean that all the *finite* subgraphs of all the graphs of $\mathfrak{s}_c^n$

would be in the subshift $X_{\mathfrak{s}_c}$. This is not a suitable property, and is more an artifact of the structure at hand. In particular, this behaviour does not occur when considering Cayley graphs of groups. The problem is not solved by requiring x to be infinite: for example, when considering the $\mathbb{Z}^2$ grid, the infinite line isomorphic to $\mathbb{Z}$ would also have all its finite subgraphs contained in some iteration of $\mathfrak{s}_C$.

Note that for a given substitution $\mathfrak{s}$, there might exist several non-isomorphic infinite substitutive graphs in the subshift, even when considering isomorphism of non-pointed graphs. Moreover, even if all the meta-tiles $\mathfrak{G}^n$ admitted no non-trivial automorphisms, it is possible that the points of a substitutive graph subshift do admit some.

Proposition 4.51

For any graph substitution $\mathfrak{s}$, we have $X_{\mathfrak{s}} \subseteq X_{\mathfrak{s}}^{\infty}$.

Proof. Let $(x, u) \in X_{\mathfrak{s}}$ be a pointed substitutive graph for $\mathfrak{s}$ and $N \in \mathbb{N}$. We construct a graph y such that $\mathfrak{s}^N(y) \simeq x$.

For $n \geq 0$, let $f(n)$ be such that $\mathcal{B}_n(x) \sqsubseteq \mathfrak{s}^{f(n)}(c_{a,n})$ for some $c_a \in \mathcal{C} \times \mathcal{A}$. More precisely, let $(\Gamma_n, u_n) \sqsubseteq \mathfrak{s}^{f(n)}(c_{a,n})$ be such that $\mathcal{B}_n(x, u) = (\Gamma_n, u_n)$. Up to extracting, as $\mathcal{C} \times \mathcal{A}$ is finite, we can suppose that $(c_{a,n})_{n\in\mathbb{N}}$ is constant equal to some c_a. For n large enough, let $v_n \in V(\mathfrak{s}^{f(n)-N}(c_a))$ be such that $u_n \in V(\mathfrak{s}_V^N(v_n)) \sqsubseteq \mathfrak{s}^{f(n)}(c_a)$. Said differently, we set v_n the specific vertex of the Nth preimage of $\mathfrak{s}^{f(n)}(c_a)$ whose substitution generates the graph containing u_n, the base point of Γ_n in $\mathfrak{s}^{f(n)}(c_a)$.

Let $y_n = \mathfrak{s}_V^{f(n)-N}(c_a)$. We have by construction $\mathcal{B}_M(\mathfrak{s}^N(y_n, v_n)) \simeq \mathcal{B}_M(x, u)$, and by compactness, (y_n, v_n) converges to some graph $y \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$. Up to another extraction, we can assume that $\mathcal{B}_n(y) = \mathcal{B}_n(y_n)$ for all n. We claim that y satisfies the required properties.

To show this, fix $M \in \mathbb{N}$, and show that $\mathcal{B}_M(\mathfrak{s}^N(y)) \simeq \mathcal{B}_M(x)$. By definition, $\mathcal{B}_M(x, u) \simeq \mathcal{B}_M(\mathfrak{s}^N(y_M), u_M)$, and as $\mathcal{B}_M(y_M) = \mathcal{B}_M(y)$, we have $\mathcal{B}_M(\mathfrak{s}(y_M)) = \mathcal{B}_M(\mathfrak{s}(y))$ and therefore $\mathcal{B}_M(x) \sqsubseteq \mathcal{B}_M(\mathfrak{s}^N(y_m)) = \mathcal{B}_M(\mathfrak{s}^N(y))$. □

The following corollary derives immediately from Proposition 4.51 and Lemma 4.34:

Corollary 4.52

For any graph substitution $\mathfrak{s}$, $X_{\mathfrak{s}}$ contains no finite graph.

Sofic graph subshifts

Just as we were able to generalize the class of SFTs over $\mathbb{Z}^d$ by considering their images by factor maps, we can do the same thing for graph subshifts. Our definition will conceptually only be meaningful for *coloured* graph substitutions: the factor maps will recolour the vertices, but will not modify the underlying graph structure.

Definition 4.53: Graph factor map

Fix some class of graphs $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, $\mathcal{A}, \mathcal{B}$ finite alphabets and $r \geq 0$. A (graph) **block map** of radius r is a map $\phi\colon \mathcal{B}_r(\mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}}) \to \mathcal{A}$.
A **factor map** of radius r is a map $\Phi\colon \mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}} \to \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$, such that there exists a block map $\phi\colon \mathcal{B}_r(\mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}}) \to \mathcal{A}$, such that for any $x \in \mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}}$:

- $\pi_{\mathcal{C}}(x) = \pi_{\mathcal{C}}(\Phi(x))$ as pointed graphs in $\mathcal{G}_{\mathcal{C},\mathcal{D}}$.
- $\forall D, \pi_{\mathcal{A}}(\text{path}_{\Phi(x)}(D)) = \phi(\mathcal{B}_r(\sigma_D(x)))$

In other words, the colour of the point at position D in the image of x depends only on the r-ball around the position D in x, and the graph structure is unchanged. We can think of Φ as recolouring the "colour" part of the vertices' type, according to some sliding window of finite radius around the vertex being considered.

This corresponds to the usual definition of factor map. Indeed, we have an equivalent of the Curtis-Hedlund-Lyndon theorem. We say that a map $\Phi\colon \mathcal{G}_{\mathcal{C},\mathcal{D}} \to \mathcal{G}_{\mathcal{C},\mathcal{D}}$ is shift-invariant if for all $(\Gamma, \text{basepoint}(\Gamma)) \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ and $d \in \bar{\mathcal{D}}$, $d \in \text{Paths}(\Gamma, \text{basepoint}(\Gamma)) \iff d \in \text{Paths}(\Phi(\Gamma), \text{basepoint}(\Phi(\Gamma)))$, and $\Phi(\sigma_d(\Gamma, \text{basepoint}(\Gamma))) = \sigma_d(\Phi(\Gamma), \text{basepoint}(\Phi(\Gamma)))$.

Proposition 4.54

Factor maps are exactly the shift-invariant continuous maps on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$.

Proof. The proof is the same as in the case of tilings on groups, see for example [CC10], or more simply Section 1.1.2. □

Definition 4.55: Sofic graph subshift

Let $X \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ be a graph subshift. We say that X is a **sofic** graph subshift if there exists $Y \subset \mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}}$ a graph SFT and $\Phi\colon \mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}} \to \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ a factor map such that $X = \Phi(Y)$.

This is the usual definition of sofic subshifts, given our definition of factor map.

Example 15. *Define the sunny-side up subshift on some class $\mathcal{G} \subseteq \mathcal{G}_{\mathcal{C},\mathcal{D}}$ the subshift $X \subset \mathcal{G}_{\mathcal{C}\times\{0,1\},\mathcal{D}}$ such that $\pi_{\mathcal{C}}(X) = \mathcal{G}$ and for $x \in X$, there is at most one vertex of x with type in $C \times \{1\}$.*

Let $\mathcal{D} = \{a\}$, $c = \{\bullet\}$, and $\mathcal{A} = \{0,1\}, \mathcal{B} = \{1, 0_L, 0_R\}$. We define an SFT Y in $\mathcal{G}_{\mathcal{C}\times\mathcal{B},\mathcal{D}}$ by requiring that:

- *every vertex is adjacent to two non-loop edges a and a^{-1}.*
- *every 0 or 1_R vertex is connected via its a-edge to a 0_R vertex*
- *every 0 or 1_L vertex is connected via its a^{-1}-edge to a 0_L vertex.*

The first condition ensures that the valid graphs are either isomorphic to $\mathbb{Z}$ or to some cycle, when only looking at their projection onto $\mathcal{C}$. Then, for $x \in X$:

- *If it contains a vertex of type $(\bullet, 1)$, then it must only see vertices of type $(\bullet, 0_R)$ when following the a-edges, and $(\bullet, 0_L)$ when following the a^{-1}-edges; in particular, the graph cannot be a cycle.*
- *Otherwise, it is a monochromatic configuration, and the graph can be a cycle.*

The block map $\phi\colon 1 \to 1, 0_L \to 0, 0_R \to 0$ then sends Y to a sunny side-up subshift.

4.4 An equivalent to Mozes theorem

4.4.1 Self-simulation in graphs

The goal of this section is to prove the main theorem of this chapter:

Theorem: Mozes theorem - graphs

Let $\mathfrak{s}$ be a graph substitution, and $\mathfrak{s}_c$ a coloured $\mathfrak{s}$-substitution. Suppose that $\mathfrak{s}$ is quasi-connected. Then, there exists a sofic graph subshift $Y_{\mathfrak{s}_c}$ which is $X^{\infty}_{\mathfrak{s}_c}$-sheeted and contains $X^{\infty}_{\mathfrak{s}_c}$.

In order to prove this result, we roughly follow the proof of [FO10], although we need to make some modifications as our notions of skeleton and border are slightly different and satisfy other conditions than the networks and facets used in this article. The general idea is as follows:

- We define a notion of **$\mathfrak{s}$-simulation**, and an associated property of being a self-simulating graph subshift. This property consists of a few conditions which ensure that any valid configuration of a subshift can be "desubstituted" infinitely many times by $\mathfrak{s}$.
- For a sufficiently nice graph substitution $\mathfrak{s}$, we construct a set of coloured graphs which generate a self-simulating graph subshift.
- A few lemmas then show that up to decorations, the graph subshifts associated to them are in fact exactly the $\mathfrak{s}$-substitutive graph subshifts.

Definition 4.56: Annotated partition

Let $\Gamma = (V, E, A) \in \mathcal{G}_{\mathcal{C},\mathcal{D}}$ be an annotated graph. An **annotated partition** of Γ is a family of annotated graphs $(G_i = (V_i, E_i, A_i))_{i\in I}$ such that:

- $\bigsqcup_{i\in I} V_i = V$
- For any edge $e \in E$ of direction $d \in \bar{\mathcal{D}}$, either there exists a unique $i \in I$ such that $e \in E_i$, or there exists distinct $i, j \in I$ and $v_i \in V_i, v_j \in V_j$ such that $A_i(v_i, d) = A_j(v_j, d^{-1}) = \texttt{required}$ and $e = v_i \xrightarrow{d} v_j$.
- For any direction $d \in \bar{\mathcal{D}}$ and $v \in V$, if $A(v, d)$ is defined then there exists a unique $i \in I$ and $v_i \in V_i$ such that $A_i(v_i, d)$ is defined and $A_i(v_i, d) = A(v, d)$.

We say that a (possibly annotated) pointed graph (G, u) is a **nearest-neighbour** annotated graph if contains a single vertex u, and possibly some annotations. In the following lemma, for a family of graphs $\mathcal{M}$, we denote $X_{\mathcal{M}}$ the subshifts whose *allowed* patterns are $\mathcal{M}$.

Definition 4.57: Self-simulation [FO10, Def. 4.1]

Let $\mathfrak{s}$ be a graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, let $\mathcal{A}$ be a finite alphabet, and let X be a subshift of $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$. Let λ_V be the usual vertex-labeling map.
We say that X is **$\mathfrak{s}$-self-simulating** if there exists $\mathcal{M}$, $\mathcal{M}'$ finite sets of finite annotated graphs of $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$, where $\mathcal{M}$ contains only nearest-neighbour annotated graphs, and a $\phi\colon \mathcal{M}' \to \mathcal{M}$ such that:

1. For any $G' \in \mathcal{M}'$, there is a graph $G \in \mathcal{M}$ such that $\mathfrak{s}(\pi_{\mathcal{C}}(G)) = \pi_{\mathcal{C}}(\phi(G'))$.
2. $X = X_{\mathcal{M}} \subseteq X_{\mathcal{M}'}$.
3. For $x \in X_{\mathcal{M}'}$, there exists an annotated partition $(\Gamma_i = (V_i, E_i, A_i))_{i\in I}$ of x such that for all $i, j \in I$:
   - $\pi_{\mathcal{C}}(\Gamma_i)$ is isomorphic to some $G'_k \in \mathcal{M}'$.
   - Γ_i, Γ_j match along their facet d if and only if $\phi(\Gamma_i), \phi(\Gamma_j)$ annotations respectively require an edge in direction d, d^{-1}.

We briefly explain the idea behind the definition:

- $\mathcal{M}, \mathcal{M}'$ are to be thought of respectively as vertices and their image by $\mathfrak{s}$, with some of their neighbours.
- The first point ensures that the elements of $\mathcal{M}'$ do in fact contain meta-tiles of $\mathfrak{s}$, up to the decorations.
- The second point ensures that a tiling which is "locally valid for $\mathcal{M}$" enforces the structure one step further for $\mathfrak{s}$, that is, it is also a valid tiling by the meta-tiles of $\mathcal{M}'$.
- The third point ensures that this tiling can be desubstituted.

This property of self-simulation is somewhat harder to define in the case of graph subshifts, compared to the original case of planar geometric tilings studied in [FO10]. Indeed, the difficulty here is that we do not have a fixed "geometry" at all: the underlying structure is $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$, and so it is less obvious what we mean by "valid tiling" for some set of tiles, and what we really need to specify as matching conditions. In particular, the third condition ensures that the family $\mathcal{M}'$ is not arbitrary, and actually induces a proper decomposition of the tilings of $X_{\mathcal{M}'}$. We can adapt this key proposition of [FO10]:

Proposition 4.58: Self-simulation [FO10, Prop 4.2]

If X is $\mathfrak{s}$-self-simulating, then $\pi_{\mathcal{C}}(X) \subseteq X_{\mathfrak{s}}^{\infty}$.

Proof. Let $x \in X$. By Definition 4.57, Item 2, we know that $x \in X_{\mathcal{M}'}$. Let $(\Gamma_i)_{i\in I}$ be the graphs obtained by Definition 4.57, Item 3. Consider the following graph y:

- Its vertices are $\{v_i, i \in I\}$, with v_i being the of the type of the only vertex of $\pi_{\mathcal{C}}(\phi(\Gamma_i))$ according to Item 1 of Definition 4.57.
- There is an edge $v_i \xrightarrow{d} v_j$ if and only Γ_i and Γ_j match along their facets d and d^{-1} respectively.

By construction, $y \in X_{\mathcal{M}} = X$. Indeed, all the vertices $u \in y$ are adjacent exactly to their required neighbours, thanks to Definition 4.57, Item 3, which enforces it at the level of the Γ_is. Repeating the argument, we obtain by induction an infinite sequence $(y_n) \in X^{-\mathbb{N}}$ with $y_0 = x$ such that $\mathfrak{s}(y_n) \simeq y_{n+1}$, and so $x \in X_{\mathfrak{s}}^{\infty}$. □

Note that in the original setting of [FO10], there is an additional condition on $\mathfrak{s}$, which ensures that we can always construct a pre-image of a "valid tiling by meta-tiles". Said differently, there are additional conditions to ensure that a tiling that can be locally de-substituted produces a valid pre-image when de-substituting independently all the meta-tiles. This condition is not required here: the original need for it is to prevent some geometric obstructions to taking a pre-image of the entire tiling by meta-tiles; in our case, our only "geometric" condition is for graphs to satisfy a local finiteness condition expressed in Lemma 4.36: it is here a consequence of meta-tiles matching along facets. However, in all generality, we cannot construct $\mathcal{M}, \mathcal{M}'$ satisfying Definition 4.57 for a graph subshift $X_{\mathfrak{s}_c}$, for reasons highlighted in Section 4.3.4. We therefore need a weaker version of the definition, and consequently, of Proposition 4.58.

Definition 4.59: Sheeted self-simulation

Let $\mathfrak{s}$, X, $\mathcal{M}, \mathcal{M}'$ and ϕ be as in Definition 4.57. We say that X is $\mathfrak{s}$-**sheeted self-simulating** if:

1. For any $G' \in \mathcal{M}'$, there is a graph $G \in \mathcal{M}$ such that $\mathfrak{s}(\pi_{\mathcal{C}}(\mathcal{G})) = \pi_{\mathcal{C}}(\phi(G'))$.
2. $X_{\mathcal{M}}$ is $X_{\mathcal{M}'}$-sheeted.
3. For $x \in X_{\mathcal{M}'}$, there exists an annotated partition $(\Gamma_i = (V_i, E_i, A_i))_{i \in I}$ of x such that for all $i, j \in I$:
   - $\pi_{\mathcal{C}}(\Gamma_i)$ is G'_k-sheeted for some $G'_k \in \mathcal{M}'$.
   - Γ_i, Γ_j match along their facet d if and only if $\phi(\Gamma_i), \phi(\Gamma_j)$ annotations respectively require an edge in direction d, d^{-1}.

Proposition 4.60

If X is $\mathfrak{s}$-sheeted self-simulating, then $\pi_C(X)$ is $X_{\mathfrak{s}_c}^{\infty}$-sheeted.

Proof. Let X be a $\mathfrak{s}$-sheeted self-simulating graph subshift, and let then $\mathcal{M}, \mathcal{M}', \phi$ be as in Definition 4.59. Let $x \in X = X_{\mathcal{M}}$, and let $u \in V(x)$. Our goal is to construct a y-sheet at u for some $y \in X_{\mathfrak{s}_c}^{\infty}$, that is, a subgraph $z \sqsubseteq x$ and a morphism $\phi\colon z \to y$ satisfying a few conditions.

We know that X is $X_{\mathcal{M}'}$-sheeted, so there exists $u \in z_0 \sqsubseteq x$ and $\psi_0\colon z_0 \hookrightarrow y_0$ an injective vertex-surjective morphism for some $y_0 \in X_{\mathcal{M}'}$ by Definition 4.46. As in the proof of Proposition 4.58, we can define another graph $x_1 \in X_{\mathcal{M}}$ using ϕ, such that $\mathfrak{s}(x_1) \simeq y_1$. Repeating the argument with y_1, we obtain inductively an infinite sequence

$$x = x_0 \sqsupseteq z_0 \xrightarrow{\psi_0} y_0 \xrightarrow{\phi} x_1 \sqsupseteq z_1 \xrightarrow{\psi_1} y_1 \xrightarrow{\phi} x_2 \dots$$

where the z_is are defined as the sheets at u and its images, or more precisely, z_i is a sheet at $\phi \circ \psi_{i-1} \circ \phi \dots \psi_0(u) \in V(x_i)$. For all $i > 0$, by definition of ϕ, we have $\mathfrak{s}(x_i) \simeq y_{i-1}$

and so $\mathfrak{s}(z_i) \sqsubseteq y_{i-1}$. Now, as ψ_i is injective and vertex-surjective, we can slightly abuse notation and define $\psi_i^{-1}\colon y_i \to z_i$, the map sending each vertex to its pre-image (this is well-defined), and which is a partial but well-defined and injective map on the edges by Lemma 4.48. We then have $\psi_{i-1}^{-1}(\mathfrak{s}(z_i)) \sqsubseteq z_{i-1}$. Define $w_i = \psi_0^{-1} \circ \mathfrak{s} \dots \psi_{i-1}^{-1} \circ \mathfrak{s}(z_i) \sqsubseteq x$. It is clear that for any i, w_i is an infinite connected graph containing u, and $w_{i+1} \sqsubseteq w_i$ by the previous remark, so we can define $w = \cap_{n\in\mathbb{N}} w_i \sqsubseteq x$. By compactness, w is an infinite connected graph containing u. We claim that w is a $X_\mathfrak{s}$-sheet at u in x. It suffices to prove that $\psi_0\colon w \to y_0$ satisfies the properties of the sheet morphism, that is, it is injective, vertex-surjective, and every edge of y_0 has corresponding edges in $x_0 \sqsupseteq w$. Injectivity is clear as $w \sqsubseteq z_0$, but we still need to prove vertex-surjectivity; in other words, we need to show that $V(w) = V(z_0)$. But that is also clear, as each of the ψ_i is vertex-surjective, and so for all i, we have $V(w_i) = V(z_0)$. □

4.4.2 Construction of a self-simulating graph SFT

We are now ready to prove Theorem 4.68. The main tool to prove the theorem will be Proposition 4.60. In order to apply this proposition, we need to define a set of decorations $\mathcal{A}_\mathfrak{s}$ for a substitution $\mathfrak{s}$ so that it satisfies Definition 4.59. We will start by defining precisely this set of decorations, and Section 4.4.3 will prove the fact that using these decorations, we define an $\mathfrak{s}$-sheeted self-simulating subshift.

Let $\mathfrak{s}$ be a substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. We describe a set of coloured vertices and coloured meta-tiles which self-simulate for $\mathfrak{s}$. To do this, we will need to consider skeletons and borders of the meta-tiles (see Definition 4.39): as these are defined on the meta-tiles of order 2, most decorations will depend on graphs $\mathfrak{s}^2(v)$ for $v \in \mathcal{C}$.

Let us first recall that we always assumed that the meta-tiles of $\mathfrak{s}$ had no non-trivial automorphism. We can therefore fix an ordering on the graphs of $\cup_{c\in\mathcal{C}}\mathfrak{s}^2(c) = \{G_1, \dots, G_N\}$ and in each G_i, an ordering of its vertices $\{v_{i,1}, \dots, v_{i,N_i}\} = V(G_i)$ and of its edges $\{e_{i,1}, \dots, e_{i,M_j}\} = E(G_i)$. This is in fact a multiset, as in the case of non-unique derivations, there might exist several isomorphic graphs $G_i \simeq G_j$, and each graph G_i comes with an implicit derivation of length 2 from some vertex type $c \in \mathcal{C}$.

We do not attempt in any way to give a minimal set of decorations: in particular, there is some clear redundancy in the vertices' decorations, but we believe that it helps to see the general idea, and how various decorations serve different purposes. The general idea behind the decorations, and the matching rules between them that we discuss below, are the following:

- In a decorated graph, each vertex will be part of a well-formed G_i-sheet for some i. This is done by "hardcoding" in the decorations the position of each individual vertex and edge in all these graphs.
- The meta-vertices of each G_i will carry additional meaningful information, namely, a position v_{jk} in another G_j. This corresponds to the fact that this meta-vertex is part of some higher-level meta-tile obtained by substituting G_j several times.
- According to this "meta" decoration, meta-vertices will perform some other checks, to ensure that they are actually part of the meta-tile G_j the believe to be a meta-vertex of: if $v_{jk} \in G_j$ is adjacent to some vertex $v_{j\ell}$, then the meta-vertex will send "signals" *via* its decorations to its edges to ensure that there is indeed some vertex which is also decorated by $v_{j\ell}$. As the skeleton is sufficiently connected thanks to the quasi-connectivity condition, we can make sure that the entire meta-tile is well-formed, by checking each neighbourhood individually.
- We need quite a few extra decorations, to ensure that every single edge that ought to appear in a graph of $X_\mathfrak{s}^\infty$ is present in the graph: indeed, the decorations are not

only enforcing a "substitutive colouring" on an already-defined space, but they must ensure that the graph itself is substitutive.

Section 4.4.2 will precisely define the decorations, and Section 4.4.2 the matching rules we impose on them. We prove in Section 4.4.3 that this defines a self-simulating graph subshift, and we deduce Theorem 4.68. We give in Section 4.4.4 some additional results which follow from the main result.

The decorations

Let us define precisely the set of decorations that we use in our construction. We decorate each vertex with the following extra colours:

- Each vertex carries two *derivations* $(c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2)$ such that for $i \in \{1, 2\}$, $c_i \rightsquigarrow \Gamma_i$ is a valid derivation for $\mathfrak{s}$, $v_i \in V(\Gamma_i)$, and the vertex $v_2 \in V(\Gamma_2)$ has type c_1. This decoration represents the pair of derivations $c_2 \rightsquigarrow \Gamma_2 \ni v_2$ of type c_1, and $c_1 \rightsquigarrow \Gamma_1 \ni v_1$.

- Each vertex carries a *position* in one of the 2-meta-tiles G_i, that is, a vertex $v_{i,j}$ for some $1 \leq j \leq N_i$.

- Each vertex carries an *annotation function* $A_v \colon \Gamma_1 \times \bar{\mathcal{D}} \to \{\texttt{required}, \texttt{forbidden}\}$, with $|A^{-1}(\texttt{required})| \geq 1$, *i.e.* there it at least annotated required edge.

If Γ is a graph thus decorated, a vertex $v \in V(\Gamma)$ decorated with $((c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2), v_{i,j}, A)$ will in fact be the vertex v_1 of some meta-tile $\Gamma_1 \sqsubseteq \Gamma$, itself in position $v_2 \in V(\Gamma_2)$ in some meta-tile of order 2. The decoration $v_{i,j}$ is used in order to ensure that the graph Γ is well-formed for arbitrary levels of meta-tiles, and not only the first two. We discuss the annotation function A_v below.

We also decorate edges with the following colours (remember that this is legitimate in our framework: this is a way to lighten the notation, and could be implemented directly in the vertices' decorations with a larger alphabet):

- Each edge carries a *derivation* (c_e, v_e, Γ_e), with $v_e \in V(\Gamma_e)$ and $c_e \rightsquigarrow \Gamma_e$ a valid substitution rule.

- Each edge also carries a *position*, that is, an edge $e_{m,n} \in E(G_m)$ for some $1 \leq n \leq M_m$.

- Each edge will also carry an *annotation function* $A_e \colon V(\Gamma) \times \bar{\mathcal{D}} \to \{\texttt{required}, \texttt{forbidden}\}$, where Γ is such that $v_e \rightsquigarrow \Gamma$ is a valid derivation rule, and as for vertices, we force $|A_e^{-1}(\texttt{required})| \geq 1$.

- Each edge will also carry a map of *meta-positions* $\texttt{Meta}_e \colon \Gamma \to \bigsqcup_i V(G_i)$, with $\mathrm{dom}(\texttt{Meta}_e) = \mathrm{dom}(A_e)$. This map represents the *position* carried by each meta-vertex of some meta-tile.

- Edges might carry no decoration at all.

There are finitely many such decorations: in particular, there are finitely many annotation functions, as the sets of vertices types and edge directions are finite.

In a decorated graph Γ, an edge $e \in E(\Gamma)$ carrying such a set of decorations has to be viewed as follows: e is part of a skeleton of some meta-tile; in this skeleton, it plays the role of the edge $e_{m,n}$, and the meta-vertices of this meta-tile are decorated according to $\texttt{Meta}_e$. This entire meta-tile itself comes from the substitution of the vertex $v_e \in \Gamma_e$, where $\Gamma_e \in \mathfrak{s}_C(c_e)$.

As for why we need an extra annotation function, Figure 4.12 shows an example of why the necessary information required to enforce the presence or absence or some edges cannot be available by looking only at the skeleton or border of a graph. In some sense, the information flows in two different directions: the other decorations are used to ensure that siblings meta-tiles can correctly and consistently be glued to form the parent meta-tile, while the annotation function contains information transmitted from the parent to their children. A very informal high-level description of its purpose is as follows: in a given 2-meta-tile G_i given by $c \rightsquigarrow \Gamma \rightsquigarrow G_i$, this function will be shared between all the vertices of G_i, and be of the form $A_e \colon \Gamma_2 \times \mathcal{D}$, where each direction d_k such that $A_e(c, d_k) = \texttt{required}$ was already adjacent to the vertex of type $c \in \Gamma_2$ from which G_i originates by $\Gamma_2 \rightsquigarrow G_i$. This is a slight difference with the other constructions of Section 4.2: in all these examples, the geometry of the ambient space meant that we did not need to think about this problem at all. For example, if two rectangles of $\mathbb{R}^2$ are placed next to one another with two adjacent corners, they will necessarily share an entire, common side. On the other hand, we need to compute by ourselves the entire set of edges linking two adjacent meta-tiles of any order, and this can no longer be done simply by ensuring that they share a predetermined set of common "meta-vertices".

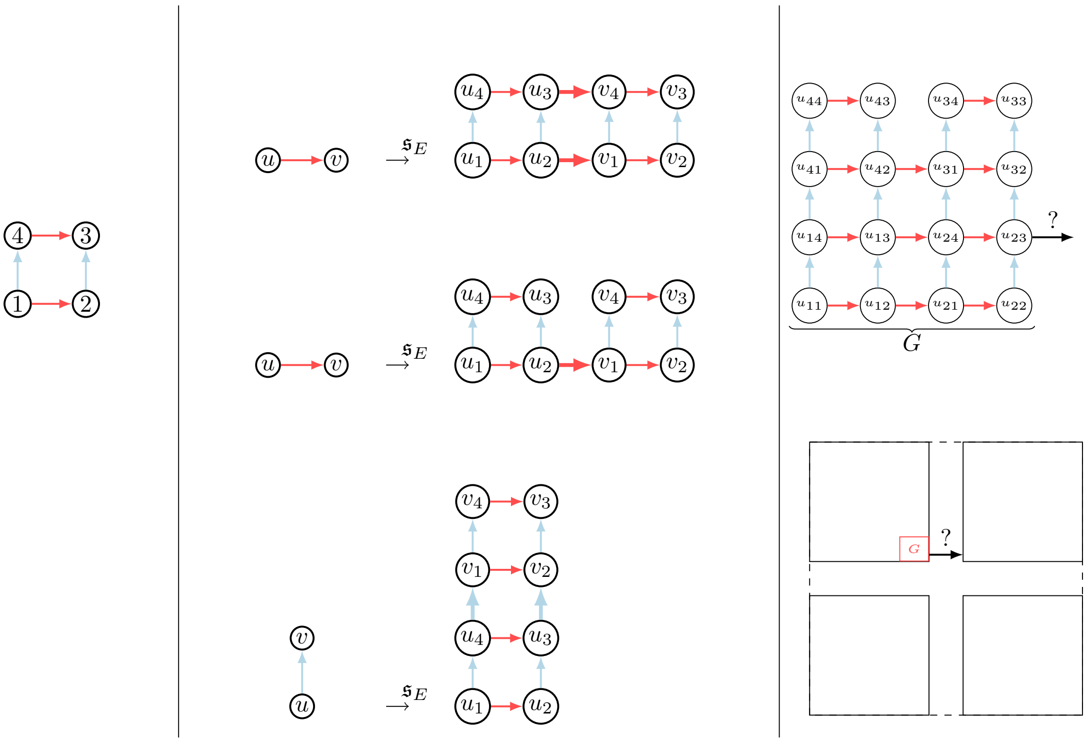

Figure 4.12: An example of non-deterministic substitution, whose iterates are subgraphs of $\mathbb{Z}^2$. The presence of the edge marked by a ? depends on information not directly available in G, and might depend on choices made arbitrarily far in the derivation sequence.

Formally, let us define the vertices decorations as

$$\begin{aligned}\texttt{Dec}_V = \Big\{ ((c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2), v_{i,j}, A) \,\Big|\, & \Gamma_b = \mathfrak{s}_V(c_b), v_b \in V(\Gamma_b), \text{ for } b = 1, 2, \\ & \lambda_V(v_2) = c_1, v_{i,j} \in V(G_i), \\ & A \colon \Gamma_1 \times \bar{\mathcal{D}} \to \{\texttt{required}, \texttt{forbidden}\} \Big\}\end{aligned}$$

and the edge decorations as

$$\begin{aligned}\mathtt{Dec}_E = \Big\{((c_e, v_e, \Gamma_e), e_{m,n}, A_e) \,\Big|\, & \Gamma_e \in \mathfrak{s}_V(c_e), v_e \in V(\Gamma_e), \\ & e_{m,n} \in E(S_{G_m}), \\ & A_e \colon \mathfrak{s}(v_e) \times \bar{\mathcal{D}} \to \{\mathtt{required}, \mathtt{forbidden}\} \\ & \mathtt{Meta}_e \colon \mathfrak{s}(v_e) \to \bigsqcup_i V(G_i)\Big\} \\ \cup\{\emptyset\} & \end{aligned}$$

These are the decorations more informally described above. The element $\emptyset \in \mathtt{Dec}_E$ corresponds to undecorated edges. Now, in order to get a nearest neighbour SFT where vertices carry the decorations, we need to decorate the vertices with $\mathtt{Dec}$, defined by:

$$\mathtt{Dec} = \mathtt{Dec}_V \times \mathtt{Dec}_E^{(\bar{\mathcal{D}})}$$

where $\mathtt{Dec}_E^{(\bar{\mathcal{D}})}$ is the set of *partial* functions from $\bar{\mathcal{D}}$ to $\mathtt{Dec}_E$. The first component corresponds to the decorations of the vertex itself, and the second component contains the decorations of each possible edge adjacent to the vertex itself. Let $\pi_{\text{vertex}}, \pi_{\text{edges}}$ be the respective projections.

Matching rules

We impose a set of rules on these decorations, split in several categories depending on which part of the hierarchical structure they enforce. We use the terminology of [BS16, Theorem 4]:

- The **structure rule** will ensure that the each vertex belongs to some $\mathfrak{s}_V(c)$, obtained from a derivation $c \rightsquigarrow \Gamma$.
- The **base rule** implements a way to synchronize the information of the 1-meta-tiles $\mathfrak{s}_V(c)$ with the meta-tiles of the next order, each isomorphic to some G_i. The skeleton edges contain all the information of this meta-tile.
- The **pasting rule** is used to ensure that siblings meta-tiles are carrying consistent information. Said differently, the analogous of the skeleton in the higher-order meta-tiles is well-formed, and all its edges have the same decorations.
- Finally, the **extension rule** ensures that the information of a child meta-tile is passed to its parent.

Example. *The Figure 4.13 shows part of a valid tiling for the weak-grid substitution, see also Figure 4.7 for how the skeleton is defined*

These rules are not optimal, in the sense that some of the matching rules might not be needed as they are implied by others. We choose to present an admittedly complex decoration scheme and set of rules, in hope that the proof of self-simulation becomes easier: any constraint on the decorations that is explicitly required need not be derived as a consequence of the other rules, using *e.g.* structural and geometrical properties of a well-decorated graph.

More formally, for a vertex v carrying the derivations $(c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2)$, a position $v_{i,j}$, and an annotation function A:

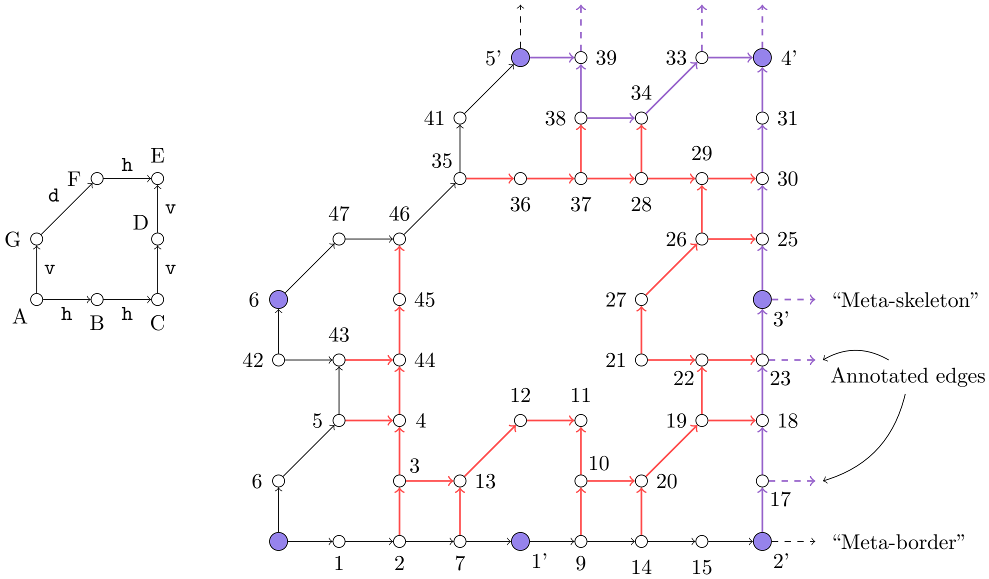

Figure 4.13: Part of a valid tiling for the weak-grid (see also Figure 4.7), assuming it is in the bottom-left corner of the higher-order substitution. Numbers correspond to the *position* $v_{i,j}$ in the decoration. The red edges are checked by the base rule, the plain purple edges are checked by the pasting rule, and the dashed purple edges are enforced by the annotation functions.

Structure rule: We always require that:

1. v has the same type as v_1 in Γ_1.

2. For each edge $(v_1, v_k) \in E(\Gamma_1)$ of some direction d, v is adjacent *via* a d-edge to some vertex u which carries the derivations $(c_1, v_k, \Gamma_1), (c_2, v_2, \Gamma_2)$ and the annotation function A.

3. For $d \in \bar{\mathcal{D}}$, v is adjacent to an edge in direction d if and only if $A(v_1, d) = \texttt{required}$.

Example. *On the Figure 4.13, this means that e.g. the vertices numbered from 7 to 13, and the purple meta-vertex between 7 and 9, are all carrying the same derivations* $(\bullet_{\mathfrak{s}}, v, \Gamma), (\bullet_{\mathfrak{s}}, B, \Gamma)$ *except for* v*, and also the same annotation function* A.

Lemma 4.61

If Γ satisfies the structure rule, then for any $u \in \Gamma$ decorated by $((c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2), \_, A) \in \texttt{Dec}_V$, Γ has a (Γ_1, A)-annotated sheet at u. These sheets can moreover be chosen to form an annotated partition of Γ.

We do not have uniqueness of the decomposition. However, starting from any vertex u, decorated with (c_1, v_1, Γ_1), and as the $\mathfrak{s}_c(c_i)$ are connected, we can pick any graph traversal of Γ_1 starting from $v_1 \in V(\Gamma_1)$ (or equivalently a spanning tree of Γ_1) and we obtain a prefix-stable language corresponding to the traversal, which in turns guarantees the existence of the Γ_1-sheet at u in Γ by Lemma 4.48.

Corollary 4.62

If Γ satisfies the structure rule, then for any $u \in \Gamma$ decorated by $((c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2), \_, A) \in \mathtt{Dec}_V$, the vertices of the Γ_i-sheet given by Lemma 4.61 are all decorated with $((c_1, \_, \Gamma_1), \_, \_, A)$.

This is an immediate consequence of Item 2 and the fact that Γ_1 is connected.

Base rule: Suppose that, according to the derivations appearing in its decorations, v is not a meta-vertex for the derivation $c_2 \rightsquigarrow \Gamma_2 \rightsquigarrow G_i$ as defined in Definition 4.39, and the vertex $v_1 \in \Gamma_1 \sqsubseteq G_i$ belongs to S_{c_2,Γ_2,G_i}. Then:

4. $v_{i,j}$ is the position of the vertex $v_1 \in \Gamma_1 \sqsubseteq G_i$ in the graph S_{c_2,Γ_2,G_i}.

5. Moreover, for each edge $e = (v_{i,j}, v_{i,k}) \in S_{c_2,\Gamma_2,G_i}$ with direction d, we require:

   (A) v is adjacent to some vertex u *via* a d-edge, and u carries derivations $\_, (c_2, v_2, \Gamma_2)$.

   (B) This edge has to carry the position given by the edge $(v_{i,j}, v_{i,k}) \in S_{c_2,\Gamma_2,G_i}$.

   (C) All such innate (not inherited) edges adjacent to v have to carry the same annotation function A_e, the same $\mathtt{Meta}_e$ function, and the same derivation (c_e, v_e, Γ_e), where the vertex $v_e \in V(\Gamma_e)$ is of type c_2.

   (D) If $v_{i,j}$ is moreover adjacent to some $v_{i,k} \in S_{c_2,\Gamma_2,G_i}$ *via* an inherited d-edge in G_i, then we require that there is some d-edge (v, u) adjacent to v, and moreover, the skeleton edges adjacent respectively to $v_{i,j}$ and $v_{i,k}$ in G_i are all present and carry the same derivations and annotation functions.

$$\text{In } G_i \qquad ? \underset{\in S}{\overset{d_1}{\longrightarrow}} v_{i,j} \underset{\text{inherited}}{\overset{d}{\longrightarrow}} v_{i,k} \underset{\in S}{\overset{d_2}{\longrightarrow}} ?$$

$$\text{Restriction on } v \qquad ? \overset{d_1}{\longrightarrow} v \overset{d}{\longrightarrow} u \overset{d_2}{\longrightarrow} ?$$

   (E) Moreover, v also enforces some conditions on the annotation function of the edges. Recall that A is the annotation function carried by the vertex v itself. We require that $\mathrm{dom}(A_e) = \mathrm{dom}(\mathtt{Meta}) = \Gamma_2 = \mathfrak{s}_V(c_e)$, and that the annotated graph (Γ_1, A) (which is well-defined by Lemma 4.61) is the annotated substitution $\mathfrak{s}(\{v_2\}, A_e|_{v_2}\})$. In other words, for d such that $A_e(v_2, d) = \mathtt{required}$, we require that $\{(v' \xrightarrow{d_e} \_) \mid A(v', d_e) = \mathtt{required}\} = \cup_{d \in \bar{\mathcal{D}}} \mathfrak{s}_E(A(v_2, d))$.

6. If v does not belong to the skeleton $S_{c_2,\Gamma_2,\Gamma_i}$ but is not a meta-vertex either, we instead simply require that all the edges adjacent to it carry exactly the same information.

On the other hand, if v is a meta-vertex according to the two derivations it is decorated with:

7. Let e be an edge of $\Gamma_1 \cap S_{c_2,\Gamma_2,\mathfrak{s}(\Gamma_2)}$, carrying a meta-position $\mathtt{Meta}_e \colon \Gamma_2 \to \bigsqcup_k V(G_k)$ (by Item 5(C) and Item 5(E), this is well-defined and does not depend on e). Then $v_{i,j} = \mathtt{Meta}_e(v_1)$.

Example. *On the Figure 4.13, this means that e.g. the vertex 10, which belongs to the skeleton, checks that the edges* $(9, 10)$ *and* $(10, 11)$ *which belong to the skeleton carry the same annotation function* A *and meta-positions* $\mathtt{Meta}$*, and the same derivation* $(\bullet_{\mathfrak{s}}, A, \Gamma)$.

The base rule enforces the presence of the inherited edge $(10, 20)$*, and checks that the derivation carried by the edges* $(14, 20)$ *and* $(20, 19)$ *is also* $(\bullet_{\mathfrak{s}}, A, \Gamma)$*.*

Vertices with no labels simply propagate the information on all their adjacent edges.

Lemma 4.63

If Γ satisfies the structure and base rules, $V(\Gamma)$ can be partitioned in graphs G_i, each one being a $\mathfrak{s}_c^2(c_i)$-sheet for some $c_i \in \mathcal{C}$. Each vertex v of G_i is decorated by $(c_1, \_, \Gamma_1), (c_2, v_2, \Gamma_2), \_, \_)$ such that $c_2 \rightsquigarrow \Gamma_2 \rightsquigarrow G_i$ is a valid derivation.

This holds because in each derivation $c_2 \rightsquigarrow \Gamma_2 \ni v_2 \rightsquigarrow \Gamma_1$, there is at least one non-meta vertex in Γ_1; by Lemma 4.61, it enforces the consistency of the graph $\Gamma_1 = \mathfrak{s}_c(c_1)$ given by the structure rule. As the skeleton is quasi-connected, this vertex also enforces by Item 4 the consistency of some of the adjacent subgraphs $\Gamma_1' \sqsubseteq \Gamma_2$. As in Lemma 4.61, a sheet at a vertex u carrying decorations $(c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2)$ can then be obtained by fixing a graph traversal of G_i obtained by the two derivations $c_2 \rightsquigarrow \Gamma_2 \rightsquigarrow G_i$ starting from $v_1 \in \Gamma_1 \sqsubseteq \Gamma_2$.

Corollary 4.64

If Γ satisfies the structure and base rule, then for any u and G_i-sheet given by Lemma 4.63, the second derivation carried by all the vertices of G_i is constant, and the edges of S_{G_i} all carry the same derivation and annotation function.

This is imposed by Item 5(C) on the 1-meta-tiles, and Item 5(D), which synchronizes it in adjacent 1-meta-tiles. By quasi-connectivity of $\mathfrak{s}$, all the 1-meta-tiles in S_{G_i} are therefore synchronized.

Pasting rule: If v carries derivations telling it that it is a meta-vertex (that is, the vertex $v_1 \in \Gamma_1 \sqsubseteq \mathfrak{s}(\Gamma_2)$ is the vertex $\text{meta}(v_2)$), then:

8. For each edge $e_{i,n} = (v_{i,j}, v_{i,k})$ adjacent to $v_{i,j}$ in $S_{c_{2,i},\Gamma_{2,i},G_i}$ of direction d, mirroring the base rule, we require:

   (A) v is adjacent to an edge in the direction given by $\text{meta}((v_{i,j}, v_{i,k}), \Gamma_{2,i}, \Gamma_1)$.

   (B) This edge carries the position $e_{i,n}$.

   (C) All these edges have to carry the same derivations (c_e, v_e, Γ_e), annotation function A_e and meta-positions $\texttt{Meta}_e$.

   (D) If $v_{i,j}$ is adjacent to some $v_{i,k} \in G_i$ via an edge $e \in E(G_i)$ of direction d_1 which is not in the skeleton, $v_{i,k}$ itself being adjacent to some edge e' of direction d_2 in G_i's skeleton, then as in Item 5(D) we moreover ensure the following: Γ_1 (containing v, given by Corollary 4.62) is adjacent *via* an edge of direction given by $\text{meta}(e)$ to a vertex u, itself adjacent to an edge of direction given by $\text{meta}(e')$, carrying the same derivation (c_e, v_e, Γ_e) and annotation function.

$$\text{In } G_i \qquad ?\ \frac{\lambda_E(e) = d_1}{\in S}\ v_{i,j}\ \frac{d}{\text{inherited}}\ v_{i,k}\ \frac{\lambda_E(e') = d_2}{\in S}\ ?$$

$$\text{Restriction on } v \qquad ?\ \frac{\text{meta}(e)}{}\ v' \in \Gamma_1\ \frac{\text{meta}(v_{i,j}, v_{i,k})}{}\ u\ \frac{\text{meta}(e')}{}\ ?$$

9. Regardless of whether v is a meta-vertex or not, if $v_{i,j} \in B_{c_{2,i},\Gamma_{2,i},G_i} \sqsubseteq G_i$, then we require the following: for any edge $e \in E(G_i)$ adjacent to $v_{i,j}$, we require that all the decorated edges of $\text{meta}(e, \Gamma_{2,i}, \Gamma_1)$ carry exactly the same information. In particular, there cannot be a single decorated edge adjacent to v.

Considering the edges of $\text{meta}(e, \Gamma_{2,i}, \Gamma_1)$ is simply a way to locally decide, in the concrete graph Γ_1 containing v given by the structure rule, which actual edges of the border of Γ_2 "represent" the meta-edges $(v_{i,j}, v_{i,k})$ that v has to enforce.

Example. *On the Figure 4.13, this rule enforces that:*

- *The large purple vertices carrying the positions 1', 2', ... to 6' have to check that the purple edges – plain and dashed – are carrying the correct (edge) position; for example, 2' checks that the edge from $2'$ going in direction* $\mathtt{v}$ *(up on the figure) carries the position $(2, 3)$, and some derivation $(\bullet_5, v, \Gamma)$. We cannot tell from the picture alone what v should be here.*
- *Moreover, the other vertices found along these purple edges have to transmit the information: for example, the vertex 18 ensures that all these edges carry the same derivation and position $(2, 3)$.*
- *The Item 8(D) ensures that the decorations carried by the plain purple edge* $\mathtt{v}$ *starting from 2' are transmitted to the neighbouring macro-tile, at the ending point of the dashed purple edge* $\mathtt{h}$ *starting from $2'$.*

Lemma 4.65

If Γ satisfies the structure, base and pasting rules, for any vertex $v \in \Gamma$ decorated by $\mathtt{dec} \in \mathtt{Dec}$ such that $\pi_{\text{vertex}}(\mathtt{dec}) = ((c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2), v_{i,j}, A)$, and for any edge $e_{i,n} = (v_{i,j}, v_{i,k}) \in S_{G_i}$, there exists a path in Γ starting from v containing only edges carrying the same decoration $((c_e, v_e, \Gamma_e), e_{i,n}, A_e, \mathtt{Meta}_e)$, which is either infinite or ends in another meta-vertex $v' \in V(\Gamma)$.

Proof. Following the edges carrying the position $e_{i,n}$ starting from v (at least one exists, by Item 8), we can either continue the path infinitely or the path ends in some vertex $v' \in \Gamma$. This path only uses edges from the border of the 2-meta-tiles given by Lemma 4.63: indeed, by Item 8, the first edge of this path is a border edge, and so by Item 9, it must be continued by another border edge carrying the same decorations. In particular, v' is necessarily decorated by some $v_{j,\ell}$ which belongs to a skeleton, but not to a border. Hence, v' must be a meta-vertex in its own 2-meta-tile, for otherwise Item 4 imposes that $v_{j,\ell}$ is its actual position in this tile, which is a contradiction with the fact that it is not on a border. □

Extension rule: Suppose that $v_{i,j}$ is both in the border B and in the skeleton S of some G_i, where G_i is the graph obtained as a derivation $c_{2,i} \rightsquigarrow \Gamma_{2,i} \rightsquigarrow G_i$, with furthermore $c_{2,i} \rightsquigarrow \Gamma_{2,i} \ni v_{2,i} \rightsquigarrow \Gamma_{1,i} \ni v_{i,j}$.

By the pasting rule, v is adjacent to the same (meta-)edges in Γ as $v_{i,j}$ is in G_i:

- By Item 8(C), all those belonging to S carry the same derivations (c_s, v_s, Γ_s), annotation function $A_S \colon \Gamma_S \times \bar{D} \to \{\texttt{required}, \texttt{forbidden}\}$ and meta-positions $\texttt{Meta}_S \colon \Gamma_S \to \bigsqcup_i V(G_i)$.
- By Item 9, all those belonging to the border B carry the same information, *i.e.* a derivation (c_b, v_b, Γ_b), a position e_{i_b,j_b} in some $G_{i_b} \in \mathfrak{s}^2(\mathcal{C})$, annotation function $A_B \colon \Gamma_B \times \bar{D} \to \{\texttt{required}, \texttt{forbidden}\}$ and meta-positions $\texttt{Meta}_B \colon \Gamma_B \to \bigsqcup_i V(G_i)$.

10. v checks that the vertex $v_s \in \Gamma_s$ is of type $c_{2,i}$.

11. As $v_{i,j}$ is in the border of $G_i = \mathfrak{s}(\Gamma_{2,i})$, there exists an edge $e \in \Gamma_{2,i}$ such that $v_{i,j} \in \mathrm{meta}(e, \Gamma_{2,i}, \Gamma_{1,i})$. Consider e as being part of the graph $e \in \Gamma_{2,i} = \mathfrak{s}(c_{2,i}) \sqsubseteq \mathfrak{s}(\Gamma_s)$, which is legitimate by Item 10. Suppose then that $e \in S_{c_s,\Gamma_s,\mathfrak{s}(\Gamma_s)}$. In that case:

   (A) v checks that the vertex v_b of Γ_b is of type c_s.

   (B) v checks that $\Gamma_s \rightsquigarrow G_{i_b}$ is a valid derivation rule.

   (C) If the edge $e \in S_{c_s,\Gamma_s,G_{i_b}}$ is numbered $e_{m,n}$, then we enforce $(m, n) = (i_b, j_b)$.

   (D) Similar to Item 5(E), we require that $\Gamma_B = \Gamma_s$, that is, $\mathrm{dom}(A_B) = \mathfrak{s}_V(c_s)$, and that the annotated graph (Γ_S, A_S) is the annotated substitution $\mathfrak{s}(\{v_s\}, A_B|_{v_s}\})$. In other words, for any direction d such that $A_B(v_s, d) = \texttt{required}$, we require that $\{(v' \xrightarrow{d_e} \_) \mid A_S(v', d_e) = \texttt{required}\} = \cup_{d \in \bar{\mathcal{D}}} \mathfrak{s}_E(A_S(v_s, d))$.

   (E) Let $u = \mathrm{meta}(v_s, \Gamma_s, G_{2,i}) \in V(G_{2,i})$. We force $\texttt{Meta}_S(u) = \texttt{Meta}_B(v_s)$, and for $u' \neq u \in V(G_{2,i})$, we define $\texttt{Meta}_S(u')$ as the position of $u' \in G_{2,i} \sqsubseteq G_{i_b}$.

The conditions on $v_{i,j}$ under which we impose the restrictions Item 11 on the edges adjacent to v are necessary: a given meta-tile will have (meta-)borders which play different roles in different level of other meta-tiles. In particular, not all such borders belong to the skeleton of the *parent* meta-tile (they could belong to the skeleton of a meta-tile of an arbitrarily large order): the conditions ensure that only the vertices $v_{i,j}$ of the "relevant" border propagate the information from the "child"'s skeleton to the parent meta-tile.

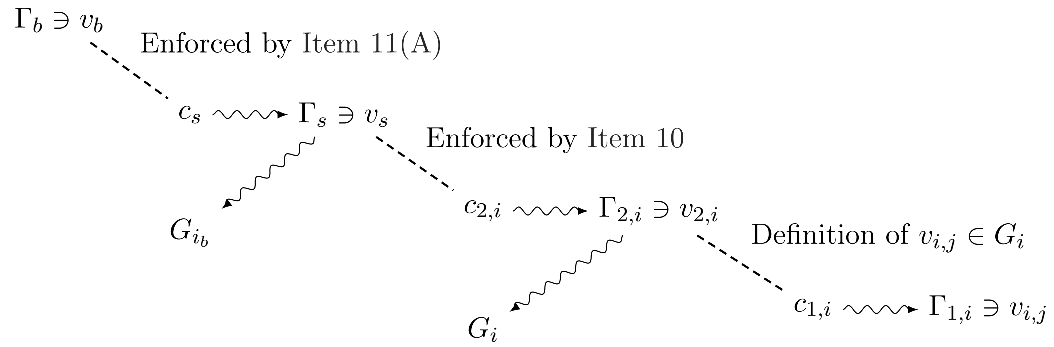

Figure 4.14: Summary of some of the relations enforced by the extension rule. The dashed lines represent a "type" relation, *e.g.* $\lambda_V(v_S) = c_{2,i} \in \mathcal{C}$, and squiggly lines are substitution rules.

We show on Figure 4.15 an example of non-deterministic substitution, with a description of which graphs are described by the various notations used in the extension rule.

Example. *On the Figure 4.13, the extension rule ensures that:*

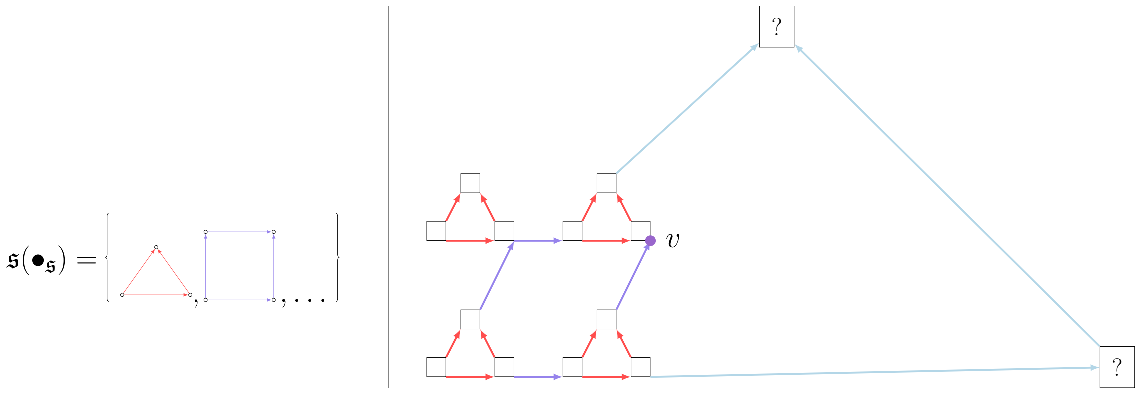

Figure 4.15: Schematic and partial representation of a non-deterministic substitution. We do not give a complete description of the skeletons and borders on the figure. For simplicity, all the meta-tiles of a same level are obtained using the same graph in $\mathfrak{s}(\bullet_{\mathfrak{s}})$. Assuming a valid decoration, the extension rule applied to the purple vertex denoted by v gives the following (recall that $\Gamma_1, \Gamma_2, \Gamma_s$ are elements of $\mathfrak{s}(axioms)$ here): Γ_1 is the "small square" of which v is the bottom-left corner. As this square is contained in a "red triangle", Γ_2 is a triangle. Moreover, v belongs to the border of the "square" containing the purple edges, and to the skeleton of the large cyan triangle. The purple square is therefore the graph $\Gamma_{2,i}$, while G_i is the entire meta-tile containing the four red triangles. Γ_s is the blue triangle, and we cannot determine Γ_b from the figure alone. The graph G_{i_b} would be obtained by taking twice the preimage of the figure, as the cyan triangle whose bottom-left vertex as been substituted by the purple square.

- *The vertices belonging to B and S, for example 23, make sure that e.g. the edges $(18, 23)$ and $(23, 3')$ carry the same information.*
- *Moreover, they also check that the information carried by these edges is consistent with the decorations of the edge $(22, 23)$, in the case of the vertex 17, in particular regarding the annotation function carried by these red edges and used in the base rule.*
- *These checks are only done by all the vertices: for example, 5 and 43, although they belong to the skeleton and to the border of the 2-meta-tile, do not enforce Item 11. On the other hand, the vertex 18 enforces it, as it belongs to the image of the edge (C, D) by $\mathfrak{s}$, and (C, D) is the edge $(2, 3)$ in the inclusion $(C, D) \in \mathfrak{s}(A) \sqsubseteq \mathfrak{s}^2(\bullet_{\mathfrak{s}})$.*
- *Item 11(E) is used to force the decoration of some meta-vertices: for example, the vertex 18 ensures that the meta-positions $\texttt{Meta}_S$ carried by all the red edges is such that $\texttt{Meta}_S(B) = 1', \cdots, \texttt{Meta}_S(G) = 6'$, and ensures that $\texttt{Meta}_S(A)$ is the same value than the one given by the border edge (17, 18).*

4.4.3 Self-simulation

Recall that we defined the decorations $\texttt{Dec}$ as

$$\texttt{Dec} = \texttt{Dec}_V \times \texttt{Dec}_E^{(\bar{\mathcal{D}})}$$

so that each vertex carries "its own" decorations, and the ones of the edges it is adjacent to. It is straightforward to convert the structure, base, pasting and extension rules into adjacency constraint on vertices decorated by $\texttt{Dec}$. Define then $\mathcal{M}_{\mathfrak{s}}$ as the set of nearest-neighbour graphs appearing in some configuration of $X_{\mathfrak{s}}^{\infty}$, with vertices carrying decorations of $\texttt{Dec}$ that satisfy the previous rules (we show in Lemma 4.66 that this is not an empty set).

Substitutive graphs can be decorated

To prove Theorem 4.68, we first explain how to decorate each element $x \in X_{\mathfrak{s}}^{\infty}$.

Lemma 4.66

For $x \in X_{\mathfrak{s}}^{\infty}$, there exists a decorated configuration $y \in X_{\mathcal{M}_{\mathfrak{s}}}$ such that $\pi_{\mathcal{C}}(y) = x$.

Proof. As $x \in X_{\mathfrak{s}}^{\infty}$, there exist arbitrary pre-images $x_1, x_2 \dots$ such that for any $n \geq 1$, $\mathfrak{s}(x_{n+1}) \simeq x_n$.

Using $\mathfrak{s}(x_1) \simeq x$, we can partition x in 1-meta-tiles, and we obtain (part of) a decoration $\mathtt{Dec}_V$ by decorating a vertex $v \in \mathfrak{s}_V(u)$ with (c_1, v_1, Γ_1), where $c_1 = \lambda_V(u)$, $\Gamma_1 = \mathfrak{s}_V(u)$ and v_1 the position of v within Γ_1. In the same way, we obtain (c_2, v_2, Γ_2) by applying this decoration scheme to $x_1 \simeq \mathfrak{s}(x_2)$. We obtain the annotation function by considering $A = \mathfrak{s}(\mathrm{Annot}_{x_1}\{u\})$. In this partition of x into 2-meta-tiles, we can decorate all the vertices which are not meta-vertices in the sense of Definition 4.39 by the correct position $v_{i,j}$ within this tile.

We decorate some of the edges of any $G_i = \mathfrak{s}^2(u)$ for $u \in x_2$ by the appropriate decorations: if $e \in S_{G_i}$, we decorate it by $((c_e, v_e, \Gamma_e), e_{i,j})$ where $e_{i,j}$ is its position in $G_i = \mathfrak{s}^2(u)$, and c_e, v_e, Γ_e are obtained by considering $u' \in x_3$ such that $u' \in \mathfrak{s}(u')$, with $\Gamma_e = \mathfrak{s}(u')$, $c_e = \lambda_V(u')$ and v_e the position of u in $\Gamma_e \sqsubseteq x_2$. For an edge e

Now, consider the position of the vertices of $\mathfrak{s}(u) \sqsubseteq x_1$ belonging to the skeleton of $\mathfrak{s}(\Gamma_e \ni u) \sqsubseteq x_1$. For each such vertex w, of some position $v_{i,j}$, decorate the meta-vertex $v \in x$ of $\mathfrak{s}(w) \sqsubseteq x$ by the position $v_{i,j}$, and define for the edges of the skeleton S_{G_i} part of their meta-positions, $\mathtt{Meta}_e(w) = v_{i,j}$. This completely decorates some of the meta-vertices of x, and partially defines $\mathtt{Meta}$ for some skeleton edges.

In order to obtain the decorations for larger parts of x, we proceed as follows. Say that some graph x is *2-decorated* if we decorated it with the previous process, so that the skeleton edges of any 2-meta-tile are decorated. We define an n-decoration for $n > 2$ inductively, by decorating x using an $n-1$-decoration of x_1. Suppose that x_1 is $n-1$-decorated. For any fully decorated meta-vertex $u \in x_1$ (that is, which carries a position $v_{i,j}$ and annotation function A, as we already determined the rest of the decorations given by π_{vertex}), we decorate the meta-vertex of $\mathfrak{s}(u) \sqsubseteq x$ with the same position $v_{i,j}$ and annotation function A. For each decorated edge $e = (u, u') \in x_1$ which is the edge $e_{m,n}$ in the skeleton of some 2-meta-tile $G_m \sqsubseteq x_i$, we decorate the meta-edge $\mathrm{meta}(e) \sqsubseteq x$ by the same decoration as e itself.

By this process, we fully decorate an increasing subset of $V(x)$ and $E(x)$. This scheme is increasing, in the sense that if $v \in V(x)$ or $e \in E(x)$ is n-decorated in some way, then it will be similarly $n+1$-decorated. In particular, we can define our final decoration pointwise, as the limit decoration given to the vertex or the edge by successive n-decorations.

If an edge remains undecorated by this process, we leave it as is, as we allow for undecorated edges. If an edge e is only decorated by a partial meta-position function $\mathtt{Meta}$, then it must contain a (meta-)vertex v such that for any $u_n \in x_n$ such that $v \in \mathfrak{s}^n(u_n) \sqsubseteq x$, u_n is a meta-vertex itself, which is not 2-decorated in x_n. We can then simply consider the meta-tile G_i of order 2 containing v in x, and set the position $v_{i,j}$ carried by x as the position of v itself in x. We can then back-propagate this decoration to decorate the pre-images $u_n \in x_n$ by the position $v_{i,j}$, obtaining n-decorations of x_1 with no undecorated vertex. Using the same construction, we can then define the *total* meta-position maps $\mathtt{Meta}_e$ for the edges, as no vertex remains partially decorated.

It is routine to check that this decoration scheme decorates x so that it belongs to

$X_{\mathcal{M}_\mathfrak{s}}$, by checking that it satisfies the structure, base, pasting and extension rules which have been defined exactly so that this is a valid decoration (the difficult part being to check that they do not allow for non $X^\infty_\mathfrak{s}$-sheeted graphs). □

Decorated graphs are sheeted-substitutive

The final step is to show that $X_{\mathcal{M}_\mathfrak{s}}$ is indeed self-simulating, using Proposition 4.60 to conclude:

Lemma 4.67

Let $\mathfrak{s}$ be a quasi-connected graph substitution. Then $X_{\mathcal{M}_\mathfrak{s}}$ is $\mathfrak{s}$-sheeted self-simulating.

Proof. We need to construct a set of graphs $\mathcal{M}'_\mathfrak{s}$ and a map $\phi\colon \mathcal{M}_{\mathfrak{s}'} \to \mathcal{M}_\mathfrak{s}$ satisfying Definition 4.59.

Definition of $\mathcal{M}'_\mathfrak{s}, \phi$: Define $\mathcal{M}'_\mathfrak{s}$ as the set of annotated pointed graphs (G', v) satisfying $\pi_\mathcal{C}(G') = \mathfrak{s}(M)$ for some $(M, u) \in \mathcal{M}_\mathfrak{s}$, with G' being furthermore locally valid for $\mathcal{M}_\mathfrak{s}$. Recall that by definition M is an annotated nearest-neighbour graphs so u is the only vertex of M, so in particular we have $v \in \mathfrak{s}(u) \sqsubseteq G'$. Consider $(G', v) \simeq \mathfrak{s}(M, u)$ a graph of $\mathcal{M}'_\mathfrak{s}$. Our first step will be to define ϕ so that $\mathfrak{s}(\pi_\mathcal{C}(M)) \simeq \pi_\mathcal{C}(G')$ as annotated graphs – in particular, ϕ must be a well-defined function, and even if $G' = \mathfrak{s}(M)$, we can obviously not use properties of M itself when defining $\phi(G')$.

Let `dec` be the decoration carried by v, $\pi_{\text{vertex}}(\texttt{dec}) = ((c_1, v_1, \Gamma_1), (c_2, v_2, \Gamma_2), v_{i,j}, A)$. Define then:

$$\pi_\mathcal{C}(\phi(G', v)) = c_1 \tag{4.1}$$

We now need to define $\pi_{\texttt{Dec}}(\phi(G', v))$. As G' satisfies the structure rule, by Corollary 4.62 v is contained in a subgraph of G' which is a Γ_1-sheet, and all the vertices of $\Gamma_1 \sqsubseteq G'$ carry the same $((c_1, \_, \Gamma_1), \_, \_)$ on the π_{vertex} part of their decoration. As $\mathfrak{s}$ is quasi-connected, the skeleton of the higher-order meta-tile containing Γ_1 intersects it along at least an edge $e \in \Gamma_1 \sqsubseteq G'$, but we do not know yet which edge it is. However, by Definition 4.39, it only depends on Γ_1. By Corollary 4.64, if there are several such edges, they all carry the same derivation (c_e, v_e, Γ_e), annotation function A_e and meta-positions $\texttt{Meta}_e$. For the same reason, we can also find the meta-vertex v' of $\Gamma_1 \sqsubseteq G'$ (which might be v itself), given by Definition 4.39. Let $v_{i',j'}$ be the position carried by this vertex (that is, the third field of π_{vertex}, which might be empty). We can now define some part of the decorations of $\phi(G', v)$:

$$\pi_{\text{vertex}}(\pi_{\texttt{Dec}}(\phi(G', v))) = ((c_2, v_2, \Gamma_2), (c_e, v_e, \Gamma_e), v_{i',j'}, A_e, \texttt{Meta}_e) \tag{4.2}$$

We now try to define the edges decorations $\phi(G', v)$, that is, $\pi_{\text{edges}}(\pi_{\texttt{Dec}}(\phi(G', v)))$. Recall that $(G', v) \simeq \mathfrak{s}(M, u)$ for some $M \in \mathcal{M}_\mathfrak{s}$. For each edge $e_B = (v_2, \_) \in \Gamma_2$ with $\lambda_E(e) = d_{e_B}$, we know by Lemma 4.65 that there exists a path of meta-edges $\text{meta}(e_B) \in G'$ starting from $v' \in \Gamma_1 \sqsubseteq G'$ containing edges that are labeled by the same $\texttt{dec}_{e_B} = ((c_{e_B}, v_{e_B}, \Gamma_{e_B}), e_{m,n}, A_{e_B}, \texttt{Meta}_{e_B})$. We can therefore define for all these edges e_B:

$$\pi_{\text{edges}}(\pi_{\texttt{Dec}}(\phi(G', v)))(d_{e_B}) = \texttt{dec}_{e_B} \tag{4.3}$$

For any other decorated edge e of direction d'_e adjacent to $w \in \Gamma_1 \sqsubseteq G'$ and decorated by some $\texttt{dec}_e$, we know by definition of a graph substitution that there it a unique direction d_e such that $(w \xrightarrow{d'_e} \_)$ is adjacent to $F_{d_e}(\Gamma_1)$. We therefore define:

$$\pi_{\text{edges}}(\pi_{\texttt{Dec}}(\phi(G', v)))(d_e) = \texttt{dec}_e \tag{4.4}$$

ϕ **is well-defined:** It is clear from the definition that ϕ sends any graph $G' \in \mathcal{M}'_{\mathfrak{s}}$ to an annotated graph but we still need to check is that the decorations carried by $\phi(G')$ are valid, that is, belong to Dec:

- $\pi_{\text{vertex}}(\pi_{\texttt{Dec}}(\phi(G', v))) = ((c_2, v_2, \Gamma_2), (c_e, v_e, \Gamma_e), v_{i',j'}, A_e) \in \texttt{Dec}_V$. Indeed, we have $v_2 \in \Gamma_2 = \mathfrak{s}_V(c_2)$ as it was already part of the decoration of v, and $v_e \in \Gamma_e = \mathfrak{s}_V(c_e)$ as it was part of a valid edge decoration. Moreover, $\lambda_V(v_e) = c_2$ by Item 5(C). Finally, $\text{dom}(A_e) = \Gamma_2$ by Item 5(E), and there is at least a required annotated edge, because this is also required in the definition of $\texttt{Dec}_E$.
- For any d, $\pi_{\text{edges}}(\pi_{\texttt{Dec}}(\phi(G', v)))(d)$ belongs to $\texttt{Dec}_E$. This is obvious by Equation (4.3).

We need to ensure that $\phi(G', v)$ thus defined still satisfies the structure, base rule, pasting and extension rules:

- $\phi(G', v)$ satisfies the structure rule because G' satisfies the base rule: Item 1 is ensured by the fact that $\lambda_V(v_2) = c_2$, Item 2 is a consequence of the fact G' satisfies Item 5 and Equation (4.3).We also obtain by Item 5(E) that for any $d \in \bar{\mathcal{D}}$, $A_e(\phi(G', v), d) = \texttt{required}$ if and only $\pi_{\text{edges}}(\pi_{\texttt{Dec}}(\phi(G', v)))(d)$ is defined, and so $\phi(G', v)$ also respects Item 3.
- For the base rule, suppose that for the derivation $c_e \rightsquigarrow \Gamma_e \rightsquigarrow G_{i'}$, the vertex $v_2 \in \mathfrak{s}(v_e) \sqsubseteq G'_i$ is not a meta-vertex and belongs to $S_{G'_i}$. Let $v_{\texttt{ext}} \in G'$ satisfying the conditions of Item 11 (in Figure 4.13, this would be the case if G was the meta-tile $\mathfrak{s}(\bullet_{\mathfrak{s}})$ containing the vertices 14 to 20, and $v_{\texttt{ext}}$ would be the vertex 18). By Item 11(E), we know the value of the meta-position $\texttt{Meta}_S$ carried by the skeleton edges in G', and Item 7 ensures that the position carried by v', the meta-vertex in G', is correct. In particular, as by Equation (4.2) we have that the position $\phi(G', v)$ is the position carried by v', we deduce that $\phi(G', v)$ satisfies Item 4 and Item 7.

  The conditions from Item 5(A) to Item 5(D) on $\phi(G', v)$ are consequences of the corresponding item from Item 8(A) to Item 8(D). Finally, the condition Item 6 on $\phi(G', v)$ is a consequence of the fact that G' satisfies Item 9.
- For the pasting rule, Item 8 and Item 9 hold in $\phi(G', v)$ because they hold in particular at the vertex $v' \in G'$, and from the definition of $v_{i',j'}$ as being a decoration by v' itself, and by the construction of the edges made in Equation (4.3), as we obtain them from the meta-edges adjacent to v'. In particular, Item 8(D) holds because the substitution is quasi-connected, and so v' being a meta-vertex means that it is itself adjacent to the inherited edges that we need to check in this point (see Lemma 4.43).
- Item 10 and Item 11 are also derived from themselves applied to G'. Indeed, they only depend on $v_{i',j'}$ and the decorations of its adjacent edges, which by Equation (4.2) and Equation (4.3) are the same in $\phi(G', v)$.

$\mathcal{M}'_{\mathfrak{s}}, \phi$ **define a sheeted-self-simulation:** We now need to check that this map satisfies Definition 4.59.

1. By definition, for any $G' \in \mathcal{M}'_{\mathfrak{s}}$, there exists $M \in \mathcal{M}_{\mathfrak{s}}$ such that $\mathfrak{s}(M) = \pi_{\mathcal{C}}(G')$. Our definition of ϕ, in particular Equation (4.1), implies that $M \simeq \phi(G', v)$.
2. $X_{\mathcal{M}}$ is $X'_{\mathcal{M}}$-sheeted. We show it using Proposition 4.47. Let $x \in X_{\mathcal{M}_{\mathfrak{s}}}$ and $u \in x$. Fix any spanning tree T of x. We claim that T satisfies Proposition 4.47, with the obvious injective morphism $\psi \colon T \hookrightarrow x$. This is an immediate consequence of Corollary 4.64 and the definition of $\mathcal{M}'_{\mathfrak{s}}$.

3. Fix $x \in X_{\mathcal{M}_{\mathfrak{s}}}$. By Corollary 4.62, we can find an annotated partition of x into $\mathfrak{s}_V(\mathcal{C})$-sheets, and we let $(\Gamma_i, A_i)_{i\in I}$ be this partition.

   - Clearly, for each Γ_i we have $\pi_{\mathcal{C}}(\Gamma_i)$ is $\mathfrak{s}(c)$-sheeted for some $s \in \mathcal{C}$: it suffices to take $c = c_1$, following the notation of Corollary 4.62, with the additional remark that the annotation function A_i is equal to the correct substituted annotation function by Item 5(E). It is easy to decorate $\mathfrak{s}(c)$ to ensure it is in $\mathcal{M}'_{\mathfrak{s}}$ (which by definition is the set of all decorated graphs $\mathfrak{s}(M)$ for $M \in \mathcal{M}_{\mathfrak{s}}$, using for example the strategy of Lemma 4.66.
   - The fact that G_i and Γ_j match along their facets is given by the fact that the only edges of x are those required by annotation functions in the decorations of the vertices $V(x)$ by Item 3, so A_i and A_j "match", and Item 5(E) and Equation (4.2) ensure that this holds for $\phi(\Gamma_i)$ and $\phi(\Gamma_j)$ computed separately.

□

This is exactly what we needed to prove the theorem:

Theorem 4.68: Mozes theorem - graphs

Let $\mathfrak{s}$ be a graph substitution, and $\mathfrak{s}_c$ a coloured $\mathfrak{s}$-substitution. Suppose that $\mathfrak{s}$ is quasi-connected. Then, there exists a sofic graph subshift $Y_{\mathfrak{s}_c}$ which is $X^{\infty}_{\mathfrak{s}_c}$-sheeted and contains $X^{\infty}_{\mathfrak{s}_c}$.

Proof. This is exactly Lemma 4.66 and Lemma 4.67. □

4.4.4 Some consequences of the construction

Link between the two substitutive subshifts

We briefly explain how we can relate properties of $X^{\infty}_{\mathfrak{s}}$ proven in Section 4.4.2 to the smaller subshift $X_{\mathfrak{s}}$.

Definition 4.69: Locally invalid graph

Let $\mathfrak{s}_c$ be a coloured graph substitution on $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$. We say that a finite graph $G \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ is **locally invalid** if there exists $H \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ and $n \geq 0$ such that:

- $G \sqsubseteq \mathfrak{s}_c^n(H)$
- For all n, $G \not\sqsubseteq \mathfrak{s}_c^n(\bullet_{\mathfrak{s}})$
- There exists $N \geq 0$ and an embedding $H \sqsubseteq \mathfrak{s}_c^N(H)$.

In general, one cannot assume that G itself can be found as a subgraph of one of its iterate by $\mathfrak{s}_c$. The kind of "periodicity" property of H ensures that one will be able to find arbitrary many preimages to H, and therefore to G, but in the case of non-deterministic substitutions, it is possible that different choices need to be made to go from H to G than to H back to itself.

Proposition 4.70

Let $\mathfrak{s}_c$ be a coloured graph substitution on $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ and $x \in X^\infty_{\mathfrak{s}_c}$. Then $x \in X_{\mathfrak{s}_c}$ if and only if there exists a finite locally invalid graph G and an embedding $G \sqsubseteq x$.

Proof. If there exists such a graph G then by definition x is not in $X_{\mathfrak{s}_c}$. For the other direction, suppose that $x \in X^\infty_{\mathfrak{s}_c} \setminus X_{\mathfrak{s}_c}$, that is, x has arbitrary preimages by $\mathfrak{s}_c$ but contain some connected $G \sqsubseteq x$ such that $G \not\sqsubseteq \mathfrak{s}_c^n(\bullet_{\mathfrak{s}})$ for any n. By definition, for all $n \geq 0$, there exists $y_n \in \mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$ such that $x \simeq \mathfrak{s}_c^n(y)$, so let H_n be the smallest (for inclusion) connected subgraph of y such that $G \sqsubseteq \mathfrak{s}_c^n(H_n)$. By definition, we have $H_n \sqsubseteq \mathfrak{s}_c(H_{n+1})$, and so $(H_n)_{n\in\mathbb{N}}$ is a decreasing sub-sequence (by number of vertices, and then by number of edges). In particular, it eventually contains only graphs with the same number of vertices and edges. We claim that we can in fact take $(H_n)_{n\in\mathbb{N}}$ eventually periodic (the choice is because $\mathfrak{s}_c$ might not be vertex deterministic). This is immediate, as there are finitely many such graphs, so we eventually find a pair $H_n = H_{n+N} = H$, so $H = H_{n+N}$ is therefore embedded in $\mathfrak{s}_c^N(H)$. □

Using this proposition, we can show that in some cases, there exists an additional set of decorations that one could add to obtain $X_{\mathfrak{s}_c}$ from $X^\infty_{\mathfrak{s}_c}$. The idea is that it suffices to ensure that no locally invalid graph G is contained in a configuration $x \in X^\infty_{\mathfrak{s}_c}$, as Proposition 4.70 then ensures that we in fact have $x \in X_{\mathfrak{s}_c}$. For that, we need to check that these graphs can be forbidden using decorations. A simple sufficient condition is to bound their size:

Definition 4.71: Bounded local invalidity

A coloured graph substitution $\mathfrak{s}_c$ is said to have N**-bounded local invalidity** if any locally invalid graph $G \sqsubseteq X^\infty_{\mathfrak{s}_c} \setminus X_{\mathfrak{s}_c}$ contains a locally invalid subgraph G' with at most N vertices. We simply say that it has **bounded local invalidity** if there exists such an $N > 0$.

We immediately obtain the next theorem:

Theorem 4.72

Let $\mathfrak{s}_c$ be a graph substitution with bounded local invalidity. Then, there exists a sofic graph subshift $Y_{\mathfrak{s}_c}$ which is $X_{\mathfrak{s}_c}$-sheeted.

Proof. Let N be such that $\mathfrak{s}_c$ has N-bounded local invalidity. It then suffices to forbid, on top of the SFT $Y_{\mathfrak{s}_c}$ of Theorem 4.68 the finitely many locally invalid subgraphs with N vertices or less. By definition, this means that the resulting configurations do not contain *any* locally invalid graph, and by Proposition 4.70 the valid configurations are then exactly $X_{\mathfrak{s}_c}$. □

Removing sheets

In a number of settings, the discussion of Section 4.4 about sheets is not required. Indeed, we have to work with this definition of $X_{\mathfrak{s}}$-sheeted subshifts for the sole reason that we cannot in general impose that arbitrarily large paths are cycles, using graph subshifts of finite type. Local consistency does not even always grant a covering, and we have to use to this weaker notion of sheets. However, in several settings, we are not constructing the "structure" – that is, the graphs $\pi_C(\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}})$ – while we define its "colouring" – that is, $\pi_A(graphclassC \times AD)$, which is what was explained in Section 4.3.4.

Following immediately from Theorem 4.68, we in fact have the following, possibly stronger result:

Theorem 4.73: Mozes theorem - coloured graphs

Let $\mathfrak{s}$ be a quasi-connected graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$, and and $\mathfrak{s}_c$ a coloured $\mathfrak{s}$-substitution on $\mathcal{G}_{\mathcal{C}\times\mathcal{A},\mathcal{D}}$. Then, there exists a graph SFT Y and a factor map Φ such that $\phi(Y) \cap X_{\mathfrak{s}}^{\infty} = X_{\mathfrak{s}_c}^{\infty}$.

Proof. It suffices to take $Y = Y_{\mathfrak{s}_c}$ from Theorem 4.68 with the associated factor map, and the intersection with $X_{\mathfrak{s}}^{\infty}$ then simply ensures that there is a single sheet containing the entire configurations. □

This version of the theorem is probably the most useful one: for example, it recovers the usual Theorem 4.3, as well as the results of [BS16].

A monotonicity result

We conclude this section with a proposition, which is an easy consequence of our definitions:

Proposition 4.74

Let $\mathfrak{s}$ be a quasi-connected graph substitution on $\mathcal{G}_{\mathcal{C},\mathcal{D}}$. For any other graph substitution $\mathfrak{s}'$ such that for all $u, v \in \mathcal{C}, d \in \mathcal{D}$:

- $V(\mathfrak{s}_V(c)) = V(\mathfrak{s}'_V(c))$
- $E(\mathfrak{s}_E(u \xrightarrow{d} v)) \subseteq E(\mathfrak{s}'_E(u \xrightarrow{d} v))$

then $\mathfrak{s}'$ also satisfies Theorem 4.68.

Proof. $\mathfrak{s}'$ is also quasi-connected, as we obtain all its images by adding edges to the images of $\mathfrak{s}$ without adding any new vertex, and so we can define the borders in the same way. □

This is not completely obvious *a priori*, as contrary to the other kinds of substitutions described in Section 4.2, any "new" edge e in the image $\mathfrak{s}_V(v)$ must be explicitly enforced by adding decorations, and its images by $\mathfrak{s}_E$ must also be hierarchically enforced. Still, our construction shows that adding a finite number of decorations suffices to ensure that this entire set of new edges is locally enforced.

References

[AA20] Shigeki Akiyama and Pierre Arnoux, eds. *Substitution and Tiling Dynamics: Introduction to Self-Inducing Structures: Cirm Jean-Morlet Chair, Fall 2017.* Springer International Publishing, 2020. ISBN: 978-3-030-57666-0. DOI: 10.1007/978-3-030-57666-0.

[ABJ18] Nathalie Aubrun, Sebastián Barbieri, and Emmanuel Jeandel. "About the Domino Problem for Subshifts on Groups". In: *Sequences, Groups, and Number Theory.* Cham: Springer International Publishing, 2018. Chap. 9, pp. 331–389. ISBN: 978-3-319-69152-7. DOI: 10.1007/978-3-319-69152-7_9.

[ABM19] Nathalie Aubrun, Sebastián Barbieri, and Etienne Moutot. "The Domino Problem is Undecidable on Surface Groups". In: *44th International Symposium on Mathematical Foundations of Computer Science (MFCS 2019).* Ed. by Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen. Vol. 138. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019, 46:1–46:14. ISBN: 978-3-95977-117-7. DOI: 10.4230/LIPIcs.MFCS.2019.46.

[ADG23] Pablo Arrighi, Amélia Durbec, and Pierre Guillon. "Graph Subshifts". In: *CoRR* (2023). arXiv: 2302.07249 [cs.DM]. URL: http://arxiv.org/abs/2302.07249v1.

[AI01] Pierre Arnoux and Shunji Ito. "Pisot Substitutions and Rauzy Fractals". In: *Bulletin of the Belgian Mathematical Society-Simon Stevin* 8.2 (2001), pp. 181–207.

[AK13] Nathalie Aubrun and Jarkko Kari. "Tiling Problems on Baumslag-Solitar Groups". In: *Electronic Proceedings in Theoretical Computer Science* 128 (Sept. 2013). DOI: 10.4204/EPTCS.128.12.

[AK21] Nathalie Aubrun and Jarkko Kari. "On the Domino Problem of the Baumslag-Solitar Groups". In: *Theoretical Computer Science* 894 (2021), pp. 12–22.

[And06] James A. Anderson. *Automata Theory With Modern Applications.* Cambridge University Press, 2006. DOI: 10.1017/CBO9780511607202.

[AS11] Nathalie Aubrun and Mathieu Sablik. *Multidimensional Effective S-Adic Systems Are sofic.* 2011. arXiv: 1103.0895v1 [cs.DM].

[AS13] Nathalie Aubrun and Mathieu Sablik. "Simulation of Effective Subshifts by Two-Dimensional Subshifts of Finite Type". In: *Acta applicandae mathematicae* 126 (2013), pp. 35–63.

[Bar22] Laurent Bartholdi. "Monadic Second-Order Logic and the Domino Problem on Self-Similar Graphs". In: *Groups, Geometry, and Dynamics* 16.4 (2022), pp. 1423–1459. DOI: 10.4171/ggd/689.

[BDJ08] Alexis Ballier, Bruno Durand, and Emmanuel Jeandel. "Structural Aspects of Tilings". In: *STACS 2008, 25th Annual Symposium on Theoretical Aspects of Computer Science, Bordeaux, France, February 21-23, 2008, Proceedings.* Ed. by Susanne Albers and Pascal Weil. Vol. 1. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2008, pp. 61–72. DOI: 10.4230/LIPICS.STACS.2008.1334.

[Ben23] Firas Ben Ramdhane. "Symbolic Dynamical Systems in Topological Spaces Defined Via Edit Distances". PhD thesis. Aix Marseille Université (AMU), Marseille, FRA.; University of Sfax, Tunisia, 2023.

[Ber+15] Valérie Berthé, Clelia De Felice, Francesco Dolce, Julien Leroy, Dominique Perrin, Christophe Reutenauer, and Giuseppina Rindone. "Acyclic, Connected and Tree Sets". In: *Monatshefte für Mathematik* 176.4 (2015), pp. 521–550.

[Ber66] Robert L. Berger. "The Undecidability of the Domino Problem". In: *Memoirs of the American Mathematical Society* (1966). URL: https://api.semanticscholar.org/CorpusID:122079099.

[BH13] Nicolas Bédaride and Arnaud Hilion. "Geometric Realizations of Two-Dimensional Substitutive Tilings". In: *The Quarterly Journal of Mathematics* 64.4 (2013), pp. 955–979. DOI: 10.1093/qmath/has025.

[BHJ18] Nicolas Bédaride, Arnaud Hilion, and Timo Jolivet. "Topological Substitution for the Aperiodic Rauzy Fractal Tiling". In: *Bulletin de la Société mathématique de France* 146.3 (2018), pp. 575–612. DOI: 10.24033/bsmf.2762.

[Bit24] Nicolás Bitar. "Subshifts of Finite Type on Groups: Emptiness and Aperiodicity". PhD thesis. Université Paris-Saclay, 2024.

[BK98] François Blanchard and Petr Kůrka. "Language Complexity of Rotations and Sturmian Sequences". In: *Theoretical computer science* 209.1-2 (1998), pp. 179–193.

[Boo58] William W Boone. "The Word Problem". In: *Proceedings of the National Academy of Sciences* 44.10 (1958), pp. 1061–1065.

[BS16] Sebastián Barbieri and Mathieu Sablik. "The Domino Problem for Self-similar Structures". In: *Pursuit of the Universal.* Pursuit of the Universal. Springer International Publishing, 2016, pp. 205–214. DOI: 10.1007/978-3-319-40189-8_21.

[BS62] Gilbert Baumslag and Donald Solitar. "Some Two-Generator One-Relator Non-Hopfian groups". In: *Bulletin of the American Mathematical Society* 68.3 (1962), pp. 199–201.

[Cab23] Christopher Cabezas. "Homomorphisms Between Multidimensional Constant-Shape Substitutions." In: *Groups, Geometry & Dynamics* 17.4 (2023).

[Car24] Nicanor Carrasco-Vargas. *Undecidability of Dynamical Properties of Sfts and Sofic Subshifts on $\mathbb{Z}^2$ and Other Groups.* 2024. arXiv: 2401.10347 [math.DS]. URL: http://arxiv.org/abs/2401.10347v1.

[CC10] Tullio Ceccherini-Silberstein and Michel Coornaert. *Cellular Automata and Groups.* Elsevier, Jan. 2010. ISBN: 978-1-4614-1799-6. DOI: 10.1007/978-3-642-14034-1.

[Cha16] Nishant Chandgotia. "Four-Cycle Free Graphs, Height Functions, the Pivot Property and Entropy Minimality". In: *Ergodic Theory and Dynamical Systems* 37.4 (2016), pp. 1102–1132. DOI: 10.1017/etds.2015.88.

[Cha18] Nishant Chandgotia. *A Short Note on the Pivot Property.* 2018. URL: https://nishantchandgotia.github.io/Research_files/Notes_files/Pivot.pdf.

[CK15a] Van Cyr and Bryna Kra. "The Automorphism Group of a Shift of Linear Growth: Beyond Transitivity". In: *Forum of Mathematics, Sigma* 3 (2015), e5. DOI: 10.1017/fms.2015.3.

[CK15b] Van Cyr and Bryna Kra. "The Automorphism Group of a Shift of Subquadratic Growth". In: *Proceedings of the American Mathematical Society* 144.2 (2015), pp. 613–621. DOI: 10.1090/proc12719.

[CL90] John H. Conway and Jeffrey C. Lagarias. "Tiling With Polyominoes and Combinatorial Group Theory". In: *J. Comb. Theory, Ser. A* 53.2 (1990), pp. 183–208. DOI: 10.1016/0097-3165(90)90057-4.

[CM18] Nishant Chandgotia and Brian Marcus. "Mixing Properties for Hom-Shifts and the Distance Between Walks on Associated Graphs". In: *Pacific Journal of Mathematics* 294.1 (2018), pp. 41–69. DOI: 10.2140/pjm.2018.294.41.

[CPV24] Antonin Callard, Léo Paviet Salomon, and Pascal Vanier. *Computability of Extender Sets in Multidimensional Subshifts.* 2024. arXiv: 2401.07549 [cs.DM].

[CV21] Antonin Callard and Pascal Vanier. "Computational Characterization of Surface Entropies for Z^2 Subshifts of Finite Type". In: (2021).

[Die08] Tammo tom Dieck. *Algebraic Topology.* Vol. 8. European Mathematical Society, 2008.

[Don+15] Sebastián Donoso, Fabien Durand, Alejandro Maass, and Samuel Petite. "On Automorphism Groups of Low Complexity Subshifts". In: *Ergodic Theory and Dynamical Systems* 36.1 (2015), pp. 64–95. DOI: 10.1017/etds.2015.70.

[Don+17] Sebastian Donoso, Fabien Durand, Alejandro Maass, and Samuel Petite. "On Automorphism Groups of Toeplitz Subshifts". In: *Discrete Analysis* (2017). DOI: 10.19086/da.1832.

[DR22] Julien Destombes and Andrei Romashchenko. "Resource-Bounded Kolmogorov Complexity Provides an Obstacle to Soficness of Multidimensional Shifts". In: *Journal of Computer and System Sciences* 128 (2022), pp. 107–134.

[DRS12] Bruno Durand, Andrei Romashchenko, and Alexander Shen. "Fixed-Point Tile Sets and Their Applications". In: *Journal of Computer and System Sciences* 78.3 (2012), pp. 731–764. DOI: 10.1016/j.jcss.2011.11.001.

[EGL23] Louis Esperet, Ugo Giocanti, and Clément Legrand-Duchesne. *The Structure of Quasi-Transitive Graphs Avoiding a Minor With Applications to the Domino Problem.* 2023. arXiv: 2304.01823 [math.CO]. URL: http://arxiv.org/abs/2304.01823v2.

[Ein01] Manfred Einsiedler. "Fundamental Cocycles of Tiling Spaces". In: *Ergodic Theory and Dynamical Systems* 21.03 (2001). DOI: 10.1017/s0143385701001389.

[EM22] Solène J Esnay and Etienne Moutot. "Aperiodic SFTs on Baumslag-Solitar Groups". In: *Theoretical Computer Science* 917 (2022), pp. 31–50.

[Fer07] Thomas Fernique. "Local Rule Substitutions and Stepped Surfaces". In: *Theoretical computer science* 380.3 (2007), pp. 317–329.

[Fer22] Alexandre Fernandez. "Théorie et Pratique des Transformations Globales". s207008. PhD thesis. 2022.

[FO10] Thomas Fernique and Nicolas Ollinger. "Combinatorial Substitutions and Sofic Tilings". In: *Computing Research Repository - CORR* 13 (Sept. 2010). arXiv: 1009.5167v2.

[FOP16] Thomas French, Nic Ormes, and Ronnie Pavlov. “Subshifts With Slowly Growing Numbers of Follower Sets”. In: *Ergodic theory, dynamical systems, and the continuing influence of John C. 506 Oxtoby* 678 (2016), pp. 175–186.

[FP19] Thomas French and Ronnie Pavlov. “Follower, Predecessor, and Extender Entropies”. In: *Monatshefte für Mathematik* 188 (2019), pp. 495–510.

[Fre16a] Thomas French. “Characterizing Follower and Extender Set Sequences”. In: *Dynamical Systems* 31.3 (2016), pp. 293–310.

[Fre16b] Thomas French. “Follower and Extender Sets in Symbolic Dynamics”. PhD thesis. University of Denver, 2016.

[Gan18] Silvère Gangloff. “Algorithmic Complexity of Growth-Type Invariants of Sft Under Dynamical Constraints”. PhD thesis. Aix-Marseille Université, 2018. URL: `http://www.theses.fr/2018AIXM0231/document`.

[Goo16] Chaim Goodman-Strauss. *Matching Rules for the Sphinx Tiling Substitution.* 2016. arXiv: `1608.07168 [math.CO]`. URL: `https://arxiv.org/abs/1608.07168`.

[Goo98] Chaim Goodman-Strauss. “Matching Rules and Substitution Tilings”. In: *The Annals of Mathematics* 147.1 (1998), p. 181. DOI: `10.2307/120988`.

[GP95] William Geller and James Propp. “The Projective Fundamental Group of a $\mathbb{Z}^2$-shift”. In: *Ergodic Theory and Dynamical Systems* 15.6 (1995), pp. 1091–1118. DOI: `10.1017/s0143385700009810`.

[GS21] Silvère Gangloff and Mathieu Sablik. “Quantified Block Gluing for Multidimensional Subshifts of Finite Type: Aperiodicity and Entropy”. In: *Journal d’Analyse Mathématique* 144.1 (2021), pp. 21–118. DOI: `10.1007/s11854-021-0172-5`.

[GS87] Branko Grünbaum and Geoffrey Colin Shephard. *Tilings and Patterns*. Courier Dover Publications, 1987.

[Hat00] Allen Hatcher. *Algebraic Topology*. Cambridge University Press, 2000. URL: `https://cds.cern.ch/record/478079`.

[Hed69] G. A. Hedlund. “Endomorphisms and Automorphisms of the Shift Dynamical System”. In: *Mathematical Systems Theory* 3.4 (1969), pp. 320–375. DOI: `10.1007/bf01691062`.

[HGO22] Benjamin Hellouin de Menibus, Silvère Gangloff, and Piotr Oprocha. “Short-Range and Long-Range Order: a Transition in Block-Gluing Behavior in Hom Shifts”. In: (2022). DOI: `10.48550/ARXIV.2211.04075`.

[Hil35] David Hilbert. “Über Die Stetige Abbildung Einer Linie Auf Ein Flächenstück”. In: *Dritter Band: Analysis· Grundlagen der Mathematik· Physik Verschiedenes: Nebst Einer Lebensgeschichte* (1935), pp. 1–2.

[HM10] Michael Hochman and Tom Meyerovitch. “A Characterization of the Entropies of Multidimensional Shifts of Finite Type”. In: *Annals of Mathematics* 171.3 (2010), pp. 2011–2038. DOI: `10.4007/annals.2010.171.2011`.

[Hoc09a] Michael Hochman. “On the Automorphism Groups of Multidimensional Shifts of Finite Type”. In: *Ergodic Theory and Dynamical Systems* 30.3 (2009), pp. 809–840. DOI: `10.1017/s0143385709000248`.

[Hoc09b] Michael Hochman. “On the Dynamics and Recursive Properties of Multidimensional Symbolic Systems”. In: *Inventiones Mathematicae* 176.1 (Apr. 2009), p. 2009.

[IYN12] Zaffar Iqbal, Shamaila Yousaf, and Sadia Noreen. "Growth Rate of Braid Monoids $MB_{n+1}, N < 5$". In: *AIP Conference Proceedings*. Vol. 1450. 1. American Institute of Physics. 2012, pp. 346–350.

[Jen01] Oliver Jenkinson. "Strong Cocycle Triviality for Z2 Subshifts". In: *Theoretical Computer Science* 262.1-2 (2001), pp. 191–213. DOI: `10.1016/s0304-3975(00)00189-4`.

[JK12] Timo Jolivet and Jarkko Kari. "Consistency of Multidimensional Combinatorial Substitutions". In: *Theoretical Computer Science* 454 (2012), pp. 178–188. DOI: `10.1016/j.tcs.2012.03.050`.

[JV13] Emmanuel Jeandel and Pascal Vanier. "Characterizations of Periods of Multi-Dimensional Shifts". In: *Ergodic Theory and Dynamical Systems* 35.2 (2013), pp. 431–460. DOI: `10.1017/etds.2013.60`.

[Kam+24] Mark Kambites, Carl-Fredrik Nyberg-Brodda, Nóra Szakács, and Richard Webb. "The Growth of Free Inverse Monoids". In: *arXiv preprint arXiv:2407.10489* (2024).

[Kar96] Jarkko Kari. "A Small Aperiodic Set of Wang Tiles". In: *Discrete Mathematics* 160.1 (1996), pp. 259–264. ISSN: 0012-365X. DOI: `10.1016/0012-365X(95)00120-L`.

[KM13] Steve Kass and Kathleen Madden. "A Sufficient Condition for Non-Soficness of Higher-Dimensional Subshifts". In: *Proceedings of the American Mathematical Society* 141.11 (2013), pp. 3803–3816.

[Kna24] Teodor Knapik. *Lindenmayer Graph Languages, First-Order Theories and Expanders*. 2024. arXiv: `2405.17629 [cs.FL]`.

[Kur03] Petr Kurka. *Topological and Symbolic Dynamics*. Société mathématique de France, 2003.

[LM21] Douglas Lind and Brian Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge university press, 2021.

[Lot02] M Lothaire. *Algebraic Combinatorics on Words*. Cambridge University Press, 2002.

[May99] J Peter May. *A Concise Course in Algebraic Topology*. University of Chicago press, 1999.

[MH38] Marston Morse and Gustav A Hedlund. "Symbolic Dynamics". In: *American Journal of Mathematics* 60.4 (1938), pp. 815–866.

[MKS04] Wilhelm Magnus, Abraham Karrass, and Donald Solitar. *Combinatorial Group Theory: Presentations of Groups in Terms of Generators and Relations*. Courier Corporation, 2004. ISBN: 9780486438306.

[Moz89] Shahar Mozes. "Tilings, Substitution Systems and Dynamical Systems Generated by Them". In: *Journal d'Analyse Mathématique* 53.1 (1989), pp. 139–186. DOI: `10.1007/bf02793412`.

[MRR01] Cristopher Moore, Ivan Rapaport, and Eric Rémila. *Tiling Groups for Wang Tiles*. Research Report LIP RR-2001-32. Laboratoire de l'informatique du parallélisme, Sept. 2001, 2+14p. URL: `https://hal-lara.archives-ouvertes.fr/hal-02101843`.

[MS15] Luidnel Maignan and Antoine Spicher. "Global Graph Transformations." In: *GCM@ ICGT*. 2015, pp. 34–49.

[Myh57] John Myhill. "Finite Automata and the Representation of Events". In: *WADD Technical Report* 57 (1957), pp. 112–137.

[Ner58] Anil Nerode. “Linear Automaton Transformations”. In: *Proceedings of the American Mathematical Society* 9.4 (1958), pp. 541–544.

[Nov55] P. S. Novikov. *Über Die Algorithmische Unentscheidbarkeit des Wortproblems in Der Gruppentheorie*. Russian. Tr. Mat. Inst. Steklova 44, 140 S. (1955). 1955.

[Nov58] P. S. Novikov. “Algorithmic Unsolvability of the Word Problem in Group Theory”. In: *Journal of Symbolic Logic* 23.1 (1958), pp. 50–52. DOI: 10.2307/2964487.

[OP16] Nic Ormes and Ronnie Pavlov. “Extender Sets and Multidimensional Subshifts”. In: *Ergodic Theory and Dynamical Systems* 36.3 (2016), pp. 908–923.

[Pel+09] Edita Pelantová et al. “A Note on Symmetries in the Rauzy Graph and Factor Frequencies”. In: *Theoretical computer science* 410.27-29 (2009), pp. 2779–2783.

[Piv07] Marcus Pivato. *Algebraic Invariants for Crystallographic Defects in Cellular Automata*. 2007. arXiv: math/0507167 [math.DS].

[PR12] Dominique Perrin and Antonio Restivo. “A Note on Sturmian Words”. In: *Theoretical Computer Science* 429 (2012), pp. 265–272.

[Pri03] Natalie Priebe Frank. “Detecting Combinatorial Hierarchy in Tilings Using Derived Voronoi Tessellations”. In: *Discrete and Computational Geometry* 29.3 (2003), pp. 459–467. DOI: 10.1007/s00454-002-0758-3.

[Pri08] Natalie Priebe Frank. “A Primer of Substitution Tilings of the Euclidean Plane”. In: *Expositiones Mathematicae* 26.4 (2008), pp. 295–326.

[Pro97] James Propp. “A Pedestrian Approach to a Method of Conway, or a Tale of Two Cities”. In: *Mathematics Magazine* 70.5 (1997), pp. 327–340. DOI: 10.1080/0025570x.1997.11996571.

[PS24] Léo Poirier and Ville Salo. *Contractible Subshifts*. 2024. arXiv: 2401.16774 [math.DS]. URL: http://arxiv.org/abs/2401.16774v1.

[PV23] Léo Paviet Salomon and Pascal Vanier. “Realizing Finitely Presented Groups as Projective Fundamental Groups of SFTs”. In: *48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*. Ed. by Jérôme Leroux, Sylvain Lombardy, and David Peleg. Vol. 272. Leibniz International Proceedings in Informatics (LIPIcs). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 75:1–75:15. ISBN: 978-3-95977-292-1. DOI: 10.4230/LIPIcs.MFCS.2023.75.

[Pyt+02] N. Pytheas Fogg, Valérie Berthé, Sébastien Ferenczi, Christian Mauduit, and Anne Siegel. *Substitutions in Dynamics, Arithmetics and Combinatorics*. Springer Berlin Heidelberg, 2002. ISBN: 978-3-540-44141-0. DOI: 10.1007/b13861.

[Rém05] Eric Rémila. “Tiling a Polygon With Two Kinds of Rectangles”. In: *Discrete & Computational Geometry* 34 (2005), pp. 313–330.

[Rob71] Raphael M Robinson. “Undecidability and Nonperiodicity for Tilings of the Plane”. In: *Inventiones mathematicae* 12 (1971), pp. 177–209.

[Sal15] Ville Salo. “Groups and Monoids of Cellular Automata”. In: *Cellular Automata and Discrete Complex Systems*. Cellular Automata and Discrete Complex Systems. Springer Berlin Heidelberg, 2015. Chap. 3, pp. 17–45. DOI: 10.1007/978-3-662-47221-7_3.

[Sal16] Ville Salo. "A Note on Subgroups of Automorphism Groups of Full Shifts". In: *Ergodic Theory and Dynamical Systems* 38.4 (2016), pp. 1588–1600. DOI: 10.1017/etds.2016.95.

[Sch95] Klaus Schmidt. "The Cohomology of Higher-Dimensional Shifts of Finite Type". In: *Pacific Journal of Mathematics* 170.1 (1995), pp. 237–269. DOI: 10.2140/pjm.1995.170.237.

[Sch98] Klaus Schmidt. "Tilings, Fundamental Cocycles and Fundamental Groups of Symbolic $\mathbb{Z}^d$-actions". In: *Ergodic Theory and Dynamical Systems* 18.6 (1998), pp. 1473–1525. DOI: 10.1017/s0143385798118060.

[Sel20] Daniel Sell. "Simple Toeplitz Subshifts: Combinatorial Properties and Uniformity of Cocycles". PhD thesis. Friedrich Schiller University Jena, 2020. URL: https://arxiv.org/pdf/2006.15348.

[Sil20] Eduardo Alejandro Silva Müller. "Subshifts En Los Grupos de Baumslag-Solitar Solubles No-Abelianos". Master Thesis. Universidad de Chile, 2020.

[Soa16] Robert I. Soare. *Turing Computability: Theory and Applications*. Vol. 300. Springer, 2016. DOI: 10.1007/978-3-642-31933-4.

[Sta83] John R. Stallings. "Topology of Finite Graphs". In: *Inventiones Mathematicae* 71.3 (Mar. 1983), pp. 551–565. ISSN: 1432-1297. DOI: 10.1007/bf02095993.

[Thu12] Axel Thue. *Über Die Gegenseitige Lage Gleicher Teile Gewisser Zeichenreihen*. 1. Jacob Dybwad, 1912.

[Thu90] William P. Thurston. "Conway's Tiling Groups". In: *The American Mathematical Monthly* 97.8 (1990), p. 757. DOI: 10.2307/2324578.

[Tur+36] Alan Mathison Turing et al. "On Computable Numbers, With an Application to the Entscheidungsproblem". In: *J. of Math* 58.345-363 (1936), p. 5.

[Van19] Pascal Vanier. "Subshifts: Aperiodicity, Complexity and Groups". Habilitation à Diriger des Recherches. Université Paris-Est Créteil, 2019.

[Vie27] Leopold Vietoris. "Über Den höheren Zusammenhang Kompakter Räume Und Eine Klasse Von Zusammenhangstreuen Abbildungen". In: *Mathematische Annalen* 97.1 (1927), pp. 454–472.

[Wan61] Hao Wang. "Proving Theorems by Pattern Recognition-II". In: *Bell system technical journal* 40.1 (1961), pp. 1–41.

[Wei12] Klaus Weihrauch. *Computable Analysis: an Introduction*. Springer Science & Business Media, 2012.

[Wei73] Benjamin Weiss. "Subshifts of Finite Type and Sofic Systems." In: *Monatshefte für Mathematik* 77 (1973), pp. 462–474. URL: http://eudml.org/doc/177666.

[ZW01] Xizhong Zheng and Klaus Weihrauch. "The Arithmetical Hierarchy of Real Numbers". In: *Mathematical Logic Quarterly: Mathematical Logic Quarterly* 47.1 (2001), pp. 51–65.

Indécidabilité des invariants géométriques dans les pavages

Mots-clés: Invariants, Sous-décalages, Dynamique symbolique, Calculabilité

Résumé: Cette thèse est consacrée à l'étude des sous-décalages, et en particulier leurs propriétés calculatoires. De façon générale, un sous-décalage est défini par un ensemble fini de symboles, un ensemble de règles spécifiant les agencements valides et invalides de ces symboles, et un espace ambiant que l'on cherche à paver: une configuration valide consiste alors en un agencement de ces symboles couvrant l'espace entier et respectant toutes les contraintes. Le sous-décalage est alors défini comme l'ensemble de toutes les configurations valides. Dans le cas le plus simple, ces règles interdisent simplement à certains symboles d'être placés côte-à-côte, et sont donc en nombre fini. Cependant, même dans ce cas restreint, les pavages de $\mathbb{Z}^d$ pour $d > 1$ sont étonnament complexes, cette complexité se manifestant sous plusieurs aspects étudiés dans cette thèse.

Cette thèse est divisée en trois chapitres essentiellement indépendants, précédés d'une introduction générale aux différents objets étudiés. Dans un premier temps, nous étudierons l'entropie d'extension des pavages de $\mathbb{Z}^d$, un nombre réel associé à un sous-décalage qui quantifie le nombre de motifs qui peuvent être librement interchangés dans n'importe quelle configuration valide. Nous montrerons que les entropies d'extension possibles sont caractérisées par des restrictions calculatoires, et correspondent exactement à des niveaux de la hiérarchie arithmétique, le niveau exact dépendant de la classe de sous-décalages considérée. Dans un second chapitre, nous nous intéresserons au Groupe Fondamental Projectif des pavages du plan $\mathbb{Z}^2$. Il s'agit d'un groupe associé à certains sous-décalages, qui permet de classifier les obstructions possibles qu'ont certaines configurations partielles ne pouvant être étendues en configurations valides sur tout l'espace. Nous montrerons là aussi que des classes simples de pavages, notamment les sous-décalages de type fini, peuvent exhiber un comportement complexe, et en particulier peuvent avoir comme groupe fondamental n'importe quel groupe finiment présenté. Enfin, nous étudierons dans un troisième chapitre les sous-décalages substitutifs, dans le contexte particulier des graphes. Nous proposerons une définition de graphe substitutif, et de sous-décalage substitutif défini sur ces graphes, et montrerons qu'une large classe de ces sous-décalages peuvent être obtenus à l'aide d'un nombre fini de règles locales. Ce résultat généralise partiellement un résultat classique de Mozes, dans un cadre plus combinatoire et moins géométrique.

Abstract: This thesis is devoted to the study of subshifts, and in particular their computational properties. A subshift is defined by a finite set of symbols, a set of rules specifying authorized and forbidden arrangements of these symbols, and an ambient space that we try to tile: a valid configuration is then an arrangement of these symbols, covering the entire space and respecting all the rules. A subshift is then defined as the set of all the valid configurations. In the simplest case, the rules are adjacency rules, which prevent some symbols from being placed next to one another. However, even in this restricted setting, tilings of $\mathbb{Z}^d$ for $d > 1$ can be surprinsingly complicated, in several ways studied in this thesis.

The thesis is divided in three independent chapters, with a preliminary chapter introducing all the relevant background knowledge for the various objects being considered. In a first chapter, we study the extender entropy of $\mathbb{Z}^d$ subshifts, a real number which quantifies for any subshift the number of patterns that can freely be exchanged in all the valid configurations. We show that the possible values of extender entropies are fully characterized by computability restrictions, more precisely, they correspond exactly to levels in the arithmetical hierarchy of real numbers, the precise level depending on the specific class of subshifts being considered. In a second chapter, we study the Projective Fundamental Group of $\mathbb{Z}^2$-subshifts, a group which aims at classifying the various kinds of obstructions encountered when trying to extend a partial configuration to a complete, valid configuration of the subshift. We show that even subshifts of finite type can have as fundamental group any finitely presented group. Finally, we study in a third chapter a kind of substitutive subshift defined on graphs. We propose a definition of substitutive graph, as well as substitutive graph subshift, and show that an important class of these subshifts can be obtained using only finitely many local rules. This partially generalizes a classical result from Mozes, in a more combinatorial but less geometrical setting.