



Arithmetical Hierarchy of the Besicovitch-Stability of Noisy Tilings

Léo Gayral¹ · Mathieu Sablik¹

Accepted: 10 August 2023 / Published online: 29 August 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

Abstract

The purpose of this article is to study the algorithmic complexity of the Besicovitch stability of noisy subshifts of finite type, a notion studied in a previous article [1]. First, we exhibit an unstable aperiodic tiling, and then see how it can serve as a building block to implement several reductions from classical undecidable problems on Turing machines. It will follow that the question of stability of subshifts of finite type is undecidable, and the strongest lower bound we obtain in the arithmetical hierarchy is Π_2 -hardness. Lastly, we prove that this decision problem, which requires to quantify over an uncountable set of probability measures, has a Π_4 upper bound.

Keywords Subshift of finite type · Stability · Besicovitch distance · Robinson tiling · Undecidability · Arithmetical hierarchy

1 Introduction

Let \mathcal{A} a finite alphabet. A subshift of finite type (SFT), denoted $\Omega_{\mathcal{F}}$, is a set of \mathcal{A} -colourings of \mathbb{Z}^d induced by a finite set of forbidden patterns \mathcal{F} which cannot appear in any configuration. One of the main topics of interest in the study of multidimensional SFTs is how a global structure can emerge from local rules. In particular, *aperiodic* SFTs have been studied by Berger [3], Robinson [4] and Kari [5] among others. One

This work is an extension upon preliminary results first introduced as an exploratory paper by the corresponding author at Automata 2021 [2].

✉ Léo Gayral
leo.gayral@math.cnrs.fr

Mathieu Sablik
msablik@math.univ-toulouse.fr

¹ IMT, Université Toulouse III - Paul Sabatier, 118 route de Narbonne, Toulouse Cedex 9 F-31062, France

of the most useful properties of the Robinson tiling is that its hierarchical structure leaves room for a *relatively* easy embedding of Turing machines into it [4, 6].

In the last decade, a lot of studies focused on the links between dynamical properties of SFTs and their algorithmic complexity. The values taken by some dynamical invariants can be characterised as some classes of (non-)computable values: possible entropies [7], or dimension entropies [8], subactions [9–11], possible periods [12], or some classes of SFTs [13]... These works help to understand the limits of what global behaviours can be enforced by local rules.

These classes of *numbers* relate to the arithmetical hierarchy of computable *sets* through the identification between $x \in \mathbb{R}$ and the interval $\{q \in \mathbb{Q}, q < x\}$. Another way to highlight the complexity of tilings is then to understand the complexity of a decision problem about a dynamical property of the SFTs. These problems are usually undecidable, but may fit into the arithmetical (or analytical) hierarchy. Regarding the arithmetical hierarchy, the Domino problem is Π_1 -complete [4], the conjugacy problem is Σ_1 -complete and the factorisation problem is Σ_3 -complete [14]... Regarding the analytical hierarchy, deciding whether a tiling has a completely positive topological entropy or not is Π_1^1 -complete [15], in dimension $d \geq 4$ the aperiodic Domino problem is Π_1^1 -complete [16]... To obtain these results, the proofs always involve the embedding of Turing machines into *complex* (and aperiodic) tilings. This is interesting since few natural problems (not directly related to a computation model) are known to be complete in these hierarchies.

In this article, we study the algorithmic complexity of the Besicovitch-stability of noisy SFTs. In a previous article [1], we introduced this notion of stability using the Besicovitch distance, which quantifies the closeness between measures through the average frequency of differences between their configurations. This framework is a natural bridge from the notion of stability described by Durand, Romashchenko and Shen [10] to ergodic theory, with a viewpoint focusing more on measure theory. The purpose is to understand if SFTs are stable in the presence of noise, if computations can survive if a small proportion of forbidden patterns is permitted. Such studies already exist for cellular automata [17] or Turing machines [18]. A digest of this framework will be introduced in Section 2, followed by a few notions about undecidability and the arithmetical hierarchy.

In the aforementioned article [1], we proved a simple computable criterion (using a word automaton) to decide stability for one-dimensional SFTs. Then, we proved the existence of both stable *and* unstable SFTs in *any* dimension, and a specific variant of the Robinson tiling was proven to be stable; before this, the only known stable aperiodic tilings were complex constructions that can be repaired locally, which is not the case for this variant of the Robinson tiling [10, 19]. However, the interface between stable and unstable examples in general was yet to be seen.

In this article, we will prove that a known two-coloured Robinson tiling is unstable in Section 3, and describe a general framework to obtain stability for some quasi-periodic SFTs in Section 4. By iterating upon both the stable and unstable constructions, we will step-by-step craft simulating tilings to show that deciding if an SFT is stable is Π_1 -hard, Σ_1 -hard and finally Π_2 -hard in Section 5.

After this, we will obtain a Π_4 upper bound for stability in Section 6. This bound may be surprising *a priori* since the definition of stability requires to quantify over uncountable sets (of translational-invariant probability measures). To obtain such a bound we will dig deeper into the technicalities of computable analysis on measures, to rewrite the stability property using only elements from a countable basis. This section is independent of the previous constructs for the lower bounds, and relies only on the definitions of Section 2.

To conclude the paper, in Section 7, we summarise some related perspectives and open questions.

2 General Framework

In this section, we define the general framework for the rest of the paper.

First, we introduce noisy SFTs and stability, which were defined more in-depth in a previous paper [1, Sections 2 and 3.1]. This subsection explains most of the notations used later on, and provides a baseline of ergodic theory for readers with a computer science background in particular.

Second, we define what decidability and the arithmetical hierarchy mean in our context, so that readers without a background in computability or mathematical logic in particular can still follow the rest.

2.1 Noisy SFTs and Besicovitch Stability

Definition 2.1 (Subshift of Finite Type) Let \mathcal{A} be a finite alphabet, and denote $\Omega_{\mathcal{A}} := \mathcal{A}^{\mathbb{Z}^d}$, endowed with the product topology and corresponding Borel algebra. Let \mathcal{F} be a finite set of *forbidden patterns* $w \in \mathcal{A}^{I(w)}$, defined on finite windows $I(w) \subseteq \mathbb{Z}^d$. An SFT is the set $\Omega_{\mathcal{F}}$ induced by \mathcal{F} as follows:

$$\Omega_{\mathcal{F}} := \left\{ \omega \in \Omega_{\mathcal{A}}, \forall w \in \mathcal{F}, \forall k \in \mathbb{Z}^d, \sigma_k(\omega)|_{I(w)} \neq w \right\},$$

with $\sigma_k : (\omega_l)_{l \in \mathbb{Z}^d} \mapsto (\omega_{k+l})_{l \in \mathbb{Z}^d}$ a translation operator for $k \in \mathbb{Z}^d$.

In other words, $\Omega_{\mathcal{F}}$ is the set of configurations such that no forbidden pattern occurs. This set is σ -invariant, invariant for any translation σ_k . Thus, if we denote $(e_i)_{1 \leq i \leq d}$ the canonical basis of \mathbb{Z}^d , $(\Omega_{\mathcal{F}}, \sigma_{e_1}, \dots, \sigma_{e_d})$ is a commutative dynamical system.

Now, we twist this notion to include noise through *obscured cells*:

Definition 2.2 (Noisy SFT) Consider the alphabet $\tilde{\mathcal{A}} = \mathcal{A} \times \{0, 1\}$, with the identification $\mathcal{A} \approx \mathcal{A} \times \{0\}$. Formally, we denote $\pi_1 : \tilde{\mathcal{A}} \rightarrow \mathcal{A}$ and $\pi_2 : \tilde{\mathcal{A}} \rightarrow \{0, 1\}$ the canonical projections. In a configuration $(\omega, b) \in \Omega_{\tilde{\mathcal{A}}}$, cells $c \in \mathbb{Z}^d$ such that $b_c = 0$ are called *clear*, and those such that $b_c = 1$ are called *obscured*. We can likewise define the set of forbidden patterns $\tilde{\mathcal{F}} := \{(w, 0^{I(w)}) \in \tilde{\mathcal{A}}^{I(w)}, w \in \mathcal{F}\}$ and the corresponding SFT $\Omega_{\tilde{\mathcal{F}}}$ on $\tilde{\mathcal{A}}$.

In general, if μ is a measure on Ω and $\varphi : \Omega \rightarrow \Omega'$ is a measurable mapping, we can define the pushforward measure $\varphi^*(\mu)$ on Ω' , such that for any measurable set $A \subset \Omega'$, we have $[\varphi^*(\mu)](A) = \mu(\varphi^{-1}(A))$.

Definition 2.3 (Noisy Probability Measures) A measure μ is σ -invariant if for any $k \in \mathbb{Z}^d$, the pushforward measure $\sigma_k^*(\mu)$ is equal to μ . Denote $\mathcal{M}_{\mathcal{F}}$ the set of σ -invariant probability measures supported by $\Omega_{\mathcal{F}}$.

Let $\mathcal{B} := \left\{ \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}, 0 \leq \varepsilon \leq 1 \right\}$ be the class of Bernoulli noises. Define:

$$\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon) := \left\{ \lambda \in \mathcal{M}_{\mathcal{F}}, \pi_2^*(\lambda) \in \mathcal{B} \text{ and } \pi_2^*(\lambda)(\{1\}) \leq \varepsilon \right\}.$$

Likewise, $\mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon) := \pi_1^* \left(\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon) \right)$ consists of probability measures on $\Omega_{\mathcal{A}}$.

The measures of $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ have a low probability of containing obscured cells in a given finite window. However, we still need a way to globally quantify the structural effect of these few local errors:

Definition 2.4 (Besicovitch Distance) We define d_H the Hamming-Besicovitch pseudo-distance on $\Omega_{\mathcal{A}}$ as $d_H(x, y) = \lim_{n \rightarrow \infty} d_n(x|_{B_n}, y|_{B_n})$, with the Hamming distances $d_n(u, v) = \frac{1}{(2n+1)^d} \# \{k \in B_n, u_k \neq v_k\}$ and $B_n := \llbracket -n, n \rrbracket^d$.

A coupling (or joining) between two measures μ on $\Omega_{\mathcal{A}_1}$ and ν on $\Omega_{\mathcal{A}_2}$ is a measure λ on $\Omega_{\mathcal{A}_1 \times \mathcal{A}_2}$ such that $\pi_1^*(\lambda) = \mu$ and $\pi_2^*(\lambda) = \nu$. Denote $J(\mu, \nu)$ the set of such couplings, and more generally $J(U, V) = \bigcup_{\mu \in U, \nu \in V} J(\mu, \nu)$. The Besicovitch distance between two σ -invariant measures is then:

$$d_B(\mu, \nu) := \inf_{\lambda \in J(\mu, \nu)} \int d_H(x, y) d\lambda(x, y).$$

By σ -invariance of the measure λ , an ergodic theorem [20, Chapter 6] gives us a link between global and local scales through the equality $\int d_H(x, y) d\lambda(x, y) = \lambda([x_0 \neq y_0])$ with the cylinder set being defined as $[x_0 \neq y_0] := \{(x, y) \in \Omega_{\mathcal{A}}^2, x_0 \neq y_0\}$. This equivalent definition of the distance can be in particular found as the distance \bar{d} in *Ergodic Theory via Joinings* [21, Chapter 15].

For two ergodic measures, d_B quantifies how well we can align their generic configurations so that they coincide on a high density subset of \mathbb{Z}^d . Using this distance, we can intuitively define stability as follows:

Definition 2.5 (Stability) The SFT $\Omega_{\mathcal{F}}$ is stable (for d_B on \mathcal{B}) if it is non-empty and there is a non-decreasing $f : [0, 1] \rightarrow \mathbb{R}^+$, continuous at 0 with $f(0) = 0$, such that:

$$\forall \varepsilon \in [0, 1], \sup_{\mu \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)} d_B(\mu, \mathcal{M}_{\mathcal{F}}) \leq f(\varepsilon).$$

In such a case, f is the speed of convergence, of stability of the system.

The general idea to keep in mind afterwards is that this framework allows us to compare the average distance between configurations, hence we will always go back to *generic* configurations in some sense, and compare these with d_H to obtain a bound for d_B . More generally, when the tilings considered are not uniquely ergodic (*i.e.* when $\mathcal{M}_{\mathcal{F}}$ is not a singleton), the system will exhibit several non-generic behaviours, and all of them must be taken into account (what matters is that we can safely ignore “strongly atypical” behaviours, of null probability regardless of the chosen measure $\mu \in \mathcal{M}_{\mathcal{F}}$).

Now that stability has been defined, we want to study its computational complexity. As we will see later on, this problem is actually undecidable, so we will want to see *how much* undecidability it contains. This is why we now need to introduce the notion of arithmetical hierarchy, which allows for a classification of the complexity of undecidable problems.

2.2 Decidability and the Arithmetical Hierarchy

The goal of this subsection is to introduce the general vocabulary and key ideas, so we will not plunge deep into the formalism, but we refer the interested reader to the classical books by Rogers [22] or Soare [23]. A less formal introduction on the topic can also be found on the mathematical blog Rising Entropy [24].

A *problem* is formally defined as a subset of integers $P \subset \mathbb{N}$, usually described implicitly as the set of integers satisfying some mathematical property. Such a problem is said to be *decidable* if there exists an algorithm (or more formally a Turing machine) that answers in finite time when asked whether $x \in \mathbb{N}$ belongs to P or not. If P cannot be decided, it is called undecidable.

This notion (and the following ones) naturally extends to any countable set that can be explicitly encoded into \mathbb{N} , such as \mathbb{Z}^d for $d \geq 2$, or the set of *finite collections of (forbidden) patterns* \mathcal{F} . Hence, we define P_{stab} as the set of families of forbidden patterns \mathcal{F} that induce a stable SFT. The goal of the arithmetical hierarchy is to further classify these undecidable problems.

Definition 2.6 [Π_k and Σ_k Problems] We say that $P \in \Pi_k$ (with $k \in \mathbb{N}$) if we have $\varphi(x, n_1, \dots, n_k)$ an algorithm (that always answers with *yes* or *no* in finite time) on \mathbb{N}^{k+1} such that $x \in P$ iff the following formula holds true:

$$\underbrace{\forall n_1 \in \mathbb{N}, \exists n_2 \in \mathbb{N}, \forall n_3 \in \mathbb{N}, \dots, \varphi(x, n_1, \dots, n_k)}_{k \text{ alternating quantifiers starting with } \forall}.$$

Likewise, we say that P is Σ_k if we have the analogous property but starting with an \exists quantifier. Note in particular how $\Pi_0 = \Sigma_0$ simply describes decidable questions.

It follows directly from the definition that $\Pi_k \cup \Sigma_k \subset \Pi_{k+1} \cap \Sigma_{k+1}$. What’s more, one can prove that this inclusion is actually strict (but the result is non-trivial).

Definition 2.7 [Π_k -hardness] We say that a problem P is Π_k -hard if, for any problem $Q \in \Pi_k$, there exists a computable reduction function $\varphi : \mathbb{N} \rightarrow \mathbb{N}$ such that $x \in Q$ iff $\varphi(x) \in P$. A problem P is then Π_k -complete if $P \in \Pi_k$ and it is Π_k -hard.




Fig. 1 The 6 basic Robinson tiles, with the leftmost one being a bumpy-corner

Notoriously, the *halting* problem P_{halt} (Does a Turing machine M halt on the empty input?) is Σ_1 -complete, and the *totality* problem P_{total} (Does M halt on all of its inputs?) is Π_2 -complete [23, Part A, Chapter IV, Theorem 3.2]. In Section 5, we will establish a computable reduction from these problems to P_{stab} to obtain a lower bound on its computational hardness. In Section 6, we will prove a Π_4 upper bound on the computational complexity of P_{stab} .

3 The Red-Black Robinson Tiling is Unstable


Consider the Robinson tiling [4] in Fig. 1, using the bumpy-corners variant (with diagonal interactions) instead of Wang tiles. The tileset uses these 6 tiles and their rotations and symmetries, for a total of 32 tiles in the alphabet. The corresponding set of forbidden patterns is self-evident, such that two laterally neighbouring tiles must have matching edges, and each square of four tiles must use exactly one bumpy-corner to fill the hole in the middle.

This tileset induces a self-similar hierarchical structure: we first define the 1-macro-tiles as the four rotated bumpy-corners tiles, and a $(n + 1)$ -macro-tile is then obtained by sticking four n -macro-tiles in a square-like pattern, around a central cross with two *arms* (with four possible orientations), as in Fig. 3.

In a previous paper [1, Theorem 7.10], we proved that an extension of this tileset, enhanced to locally enforce the alignment of macro-tiles, was stable with a polynomial speed $O(\sqrt[3]{\varepsilon})$. Durand, Romashchenko and Shen [10] had introduced before the notion of (c_1, c_2) -robustness, and constructed a fixed-point tileset satisfying this property. For such robust tilesets, they proved that in the presence of independent random errors, it is possible to correct a sequence of isolated “islands of errors” and obtain an admissible tiling (and the density of corrections goes to 0 as the density of errors ε does). This method implies the current notion of Besicovitch stability, but the speed obtained is a very slow $O\left(\frac{1}{\sqrt{\ln(\varepsilon^{-1})}}\right)$ [1, Section 6]. Note that the Robinson tiling is not (c_1, c_2) -robust.

Here, we will use the two-coloured extension of this Robinson tileset in Fig. 2, which naturally projects onto the previous tiling, so all the structural properties of the Robinson tiling still hold, and most notably aperiodicity. We will denote \mathcal{A} the tileset, RB the corresponding set of forbidden patterns, and Ω_{RB} the resulting SFT.

Fig. 2 The 11 basic Red-Black Robinson tiles



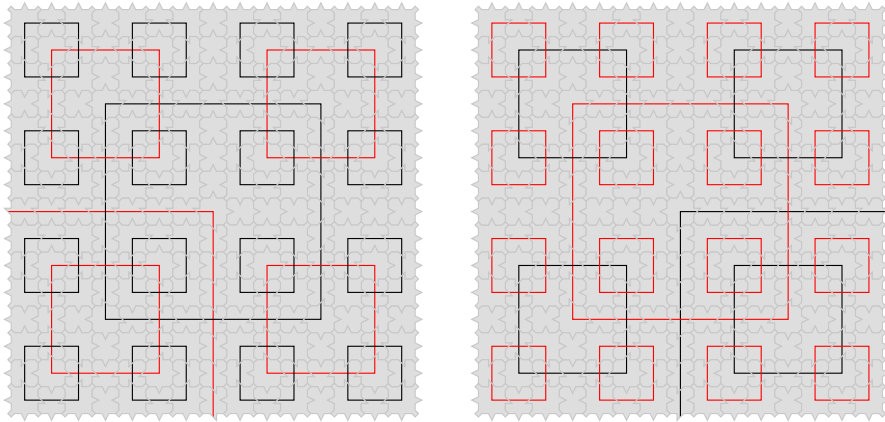


Fig. 3 Alternating colours in Red-Black Robinson macro-tiles

Because \mathcal{A} contains no tile with a monochromatic cross, only small crosses made of a straight Red line crossing with a Black one, any two squares of the same colour in the hierarchical structure of a tiling do not intersect, as we can see on the 5-macro-tiles in Fig. 3. In Subsection 5.1.1, these non-intersecting Red squares will be used to encode arbitrarily large space-time diagrams of Turing machines.

For the rest of this paper, a *generic* Robinson tiling will refer to a configuration without an infinite *cut*, such that any two tiles of \mathbb{Z}^d end up being in the same n -macro-tile for big-enough values of n . In such a *generic* configuration $\omega \in \Omega_{RB}$, by induction, the n -macro-tiles all have central arms with the same colour. In particular, a generic configuration will only contain Red *or* Black bumpy-corners, never both.


Proposition 3.1 *Let Ω_{RB} be the Red-Black Robinson tiling. For any $\varepsilon > 0$, there is $\mu \in \mathcal{M}_{RB}^B(\varepsilon)$ such that $d_B(\mu, \mathcal{M}_{RB}) \geq \frac{1}{8}$. Thus, the SFT is unstable.*

Proof The goal of this proof is to convert a generic tiling $\omega \in \Omega_{RB}$ into a random noisy tiling $\lambda_{\omega,b}$ on $\widetilde{\Omega}_{RB}$, with b a random variable on $\Omega_{\{0,1\}}$. Using a generic Bernoulli noise b in the input, we will obtain a noisy tiling for which its bumpy-corners are now half Red and half Black, which will yield the announced result since bumpy-corners have frequency $\frac{1}{4}$ in the Robinson tiling.

We will build this measure λ iteratively, as a limit of a locally-defined (thus trivially measurable) transformations. At each step of the construction, the actual monochromatic structure of the Robinson tiling will be preserved, and only the colours will be mismatched, so we may still consider n -macro-tiles in this structural sense, even though they are not actually locally admissible. We initialise $\lambda_1 = \delta_{(\omega,b)}$ as a constant Dirac measure.

Let us now explain how we obtain λ_2 out of λ_1 . This transformation will be done independently on *each* of the 2-macro-tiles of ω . We distinguish two cases, both illustrated in Fig. 4, where the black cells c represent *obscured* tiles with a noise $b_c = 1$. A macro-tile is said to be flippable if both of its bi-coloured crosses, highlighted with green borders in the figure, are obscured tiles. In such a situation, we will flip its

Fig. 4 A locally admissible 3-macro-tile with obscured cells



colours (Black lines become Red and conversely) with probability $\frac{1}{2}$, independently of the rest, which still preserves the local rules inside the macro-tile. In the figure, the top-left macro-tile is flipped, the top-right macro-tile is flippable but not flipped, and the two bottom macro-tiles are not flippable.

Likewise, we go from λ_{n-1} to λ_n by flipping independently at random any flippable n -macro-tile (except the two ends of the central arms that are “after” the bi-coloured crossed tiles, which must match the colour of the yet-unflipped $(n + 1)$ -macro-tile that contains them). This process guarantees that, if we denote $\omega' \sim \lambda_n$ the new colouring, then $(\omega', b) \in \Omega_{RB}$ necessarily.

Notice how the highlighted cells that decide whether a given macro-tile is flippable are disjoint for each macro-tile. Hence, assuming that $B \sim \mathcal{B}(\varepsilon)^{\mathbb{Z}^2}$ is a random Bernoulli noise, each macro-tile at each scale is flippable with probability ε^2 , independently of the rest. For any initial choice of $\omega \in X = \Omega_{RB}$, the weak- $*$ limit measure $\lambda_{\omega,B} = \lim_n \lambda_n$ is well-defined, supported by Ω_{RB} .

Notice how $\lambda_{\omega,B}$ is not σ -invariant (as the monochromatic Robinson structure of ω is preserved), but it “commutes” with σ in that $\sigma_k^*(\lambda_{\sigma_k(\omega),B}) = \lambda_{\omega,B}$. It follows that, if we average $\lambda_{\omega,B}$ over $\omega \sim \mu_0 \in \mathcal{M}_{RB}$ (chosen independently from the noise B), we obtain a σ -invariant measure $\mu \in \mathcal{M}_{RB}^B(\varepsilon)$.

Now, consider $G(\omega) \subset \mathbb{Z}^2$ the set of cells containing a bumpy corner in a generic configuration $\omega \in \Omega_{RB}$. For a given cell $c \in G(\omega)$, we denote by $\text{flip}_{c,n}$ the random variable equal to 1 when the n -macro-tile containing c is flippable for the pair (ω, B) . Hence the variables $\text{flip}_{c,n} \sim \mathcal{B}(\varepsilon^2)$ are *iid*. Conditionally to the event $\text{flip}_{c,n} = 1$, the colour of the cell c is uniformly distributed in λ_n after rank n . Thence, by Borel-Cantelli lemma, the colour of c is uniformly distributed in $\lambda_{\omega,B}$ (and thus μ).

Likewise, consider two distant cells $c, d \in G(\omega)$. As $d_\infty(c, d) \rightarrow \infty$, the smallest rank $n_0(c, d)$ such that c and d belong to the same n -macro-tile of ω goes to infinity. The families $(\text{flip}_{c,n})_{n < n_0}$ and $(\text{flip}_{d,n})_{n < n_0}$ are independent, and conditionally to the fact that both of these sequences contain at least a 1, the colours of cells c and d are independently uniform (in the measures λ_n after rank n_0 , hence for $\lambda_{\omega,B}$ and μ).

Now, let $\omega \sim \mu$. Without loss of generality, assume $0 \in G(\omega)$, so that $G = (2\mathbb{Z})^2$. Then the family (colour of the cell $2c$) $_{c \in \mathbb{Z}^2}$ describes a σ -invariant ergodic dynamical system, so that we may apply a pointwise ergodic theorem. This implies that the frequency of both Black and Red bumpy-corners is generically equal to $\frac{1}{2}$ in μ . As G has density $\frac{1}{4}$ in \mathbb{Z}^2 , we conclude that for almost-any $\omega' \sim \mu$ and any generic $\omega_0 \in \Omega_{RB}$ (with monochromatic bumpy-corners), we have the bound $d_H(\omega_0, \omega) \geq \frac{1}{2} \times \frac{1}{4} = \frac{1}{8}$ assuming bumpy-corners overlap between the two configurations, and even a $\frac{1}{2}$ bound if they are misaligned. In conclusion, μ satisfies $d_B(\mu, \mathcal{M}_{RB}) \geq \frac{1}{8}$. \square

The result still holds with the very same proof if we replace the bi-coloured Robinson tiling by a bi-coloured variant of the structurally enhanced Robinson tiling from our previous paper [1].

However, as the proof relies heavily on flipping the colours of bumpy-corners, by keeping only *one* of the two colours specifically for this tile, we obtain a stable tiling again. This will be useful later on, when we want to encode Turing machines into Robinson (which requires this bi-coloured setting) in a stable way. In such situations, stability will follow from the result of the next section.

4 Generalising Aperiodic Stability

In order to prove the stable cases later on, we will state a direct generalisation of one of the main results in our previous article [1, Proposition 7.9]. This proposition was proven in the specific context of the enhanced Robinson tiling, but we will here reformulate the result in a general framework for quasi-periodic tilings with a well-behaved reconstruction function, so that it applies as a black box to the tilesets described in the next section. This section is here mostly for the sake of technical completeness, and can be skipped to focus on the core of the paper to which we go back right after.

Definition 4.1 (Almost Periodic SFT) Let $\Omega_{\mathcal{F}}$ be an SFT on the alphabet \mathcal{A} , and consider $p \in \mathbb{N}^*$ a positive integer and $\rho > 0$.

We say that $\Omega_{\mathcal{F}}$ is ρ -almost p -periodic if there is a p -periodic set $G \subset \mathbb{Z}^d$ (invariant under translations in $(p\mathbb{Z})^d$) acting as a “grid” of density at most ρ , such that any configuration restricted to a translation of G^c is made periodic. By this, we mean that for any $\omega \in \Omega_{\mathcal{F}}$, there is a unique translation of G (given by a non-necessarily unique $k \in \llbracket 0, p - 1 \rrbracket^d$) such that $\omega|_{G^c+k}$ is p -periodic.

In this case, assuming $\square \notin \mathcal{A}$, we can define ω^\square by overwriting $\omega|_{G+k}$ by the blank symbol \square . Thence, $\Omega_{\mathcal{F}}^\square = \left\{ \omega^\square, \omega \in \Omega_{\mathcal{F}} \right\}$ is a finite p -periodic SFT.

The non-uniqueness of k comes from the fact that, for example, we may want to consider G a $\frac{p}{2}$ -periodic grid instead, with some more redundancy in its structure.

Definition 4.2 (C -Reconstruction Function) Consider $\Omega_{\mathcal{F}}$ a ρ -almost p -periodic SFT and G the associated grid.

The SFT has the C -reconstruction property if, for any locally admissible tiling ω of $B_{\lceil \frac{p}{2} \rceil + C}$ there is a unique translation of G such that $\left[\omega|_{B_{\lceil \frac{p}{2} \rceil} \cap (G^c+k)} \right]^\square$ (obtained

by filling $B_{\lceil \frac{p}{2} \rceil} \cap (G + k)$ with \square symbols) is globally admissible in $\Omega_{\mathcal{F}}^{\square}$ (thence $\omega|_{B_{\lceil \frac{p}{2} \rceil} \cap (G^c + k)}$ is globally admissible in $\Omega_{\mathcal{F}}$). What’s more, the translation of G depends only on what appears in any p -square included in the central window $B_{\lceil \frac{p}{2} \rceil}$ (which is either a $(p + 1)$ or $(p + 2)$ -square depending on the parity).

As $\Omega_{\mathcal{F}}^{\square}$ is p -periodic, there is a unique choice $\omega^{\square} \in \Omega_{\mathcal{F}}^{\square}$ of configuration that will match the pattern $\left[\omega|_{B_{\lceil \frac{p}{2} \rceil} \cap (G^c + k)} \right]^{\square}$.

Proposition 4.1 [*Besicovitch Bound*] Consider a ρ -almost p -periodic SFT with C -reconstruction. Then, for any $\varepsilon > 0$ and $\mu \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$, we have the bound $d_B(\mu, \mathcal{M}_{\mathcal{F}}) \leq 48 \left(2 \left(C + \lceil \frac{p}{2} \rceil \right) + 1 \right)^d \varepsilon + \rho$.

Proof The proof is really similar to the source result [1, Proposition 7.9], so we will just give the general idea.

Consider $\lambda \in \widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ and $(\omega, b) \in \Omega_{\mathcal{F}}$ a λ -generic noisy configuration. A percolation argument [1, Proposition 5.6] tells us that, almost-surely, we can forget about the $(C + \lceil \frac{p}{2} \rceil)$ -neighbourhood of *obscured* cells (cells $c \in \mathbb{Z}^d$ with $b_c = 1$) and still have a unique connected component of *clear* cells ($b_c = 0$) with density of at least $1 - 48 \left(2 \left(C + \lceil \frac{p}{2} \rceil \right) + 1 \right)^d \varepsilon$.

Each clear cell c of this connected component is the center of a clear window I_c of diameter $2 \lceil \frac{p}{2} \rceil + 1$, the C -neighbourhood of which is clear and locally admissible, so by the C -reconstruction property, there is a unique translation of G and a unique periodic configuration $\omega_c^{\square} \in \Omega_{\mathcal{F}}^{\square}$ that matches ω on $I_c \cap (G + k)^c$ (for the right translation). We can do likewise for any other cell.

Now, two neighbouring cells $c, c' \in \mathbb{Z}^d$ share a common p -square window which fixes the same choice of translation for G . Hence, ω_c^{\square} and $\omega_{c'}^{\square}$ overlap on this p -square, and $\Omega_{\mathcal{F}}^{\square}$ is p -periodic, so they are equal. Thus, all the cells of the infinite connected component $I(b)$ must agree on the same ω^{\square} . The map $\varphi : (\omega, b) \mapsto \omega^{\square}$ is measurable (for ε small-enough, so that I has density greater than $\frac{1}{2}$).

In particular, ω and ω^{\square} can only differ outside of $I(b)$, or on the translation of G , so $d_H(\omega, \omega^{\square}) \leq \text{density}(I^c) + \text{density}(G) \leq 48 \left(2 \left(C + \lceil \frac{p}{2} \rceil \right) + 1 \right)^d \varepsilon + \rho$, and the same bound holds for $d_B(\pi_1^*(\lambda), \varphi^*(\lambda))$. At last, we can fill-in the \square symbols of G in an appropriate random way, in order to send $\varphi^*(\lambda)$ into $\mathcal{M}_{\mathcal{F}}$, without changing the bound on d_B . □

In particular, this proposition gives us a linear $O(\varepsilon)$ bound for f the speed of stability of any actually periodic tiling (which will be 0-almost periodic, with $G = \emptyset$ and C -reconstruction for some C).

However, it doesn’t apply to the Red-Black tiling from the previous section, for which we can juxtapose side by side a Red and a Black n -macro-tile at any scale in a locally admissible way, which breaks the desired quasi-periodicity.

Corollary 4.2 (Stability) Assume there is a *sequence* of triplets (p_n, ρ_n, C_n) for which Proposition 4.1 applies to $\Omega_{\mathcal{F}}$. Then, as soon as $\rho_n \rightarrow 0$, we conclude that $\Omega_{\mathcal{F}}$ is a stable SFT.

Lemma 4.3 (Meta Multi-Scale-to-Polynomial Bound) Consider $D_k = \varepsilon\alpha^k + \beta^k$ with $k \in \mathbb{Z}$ and $0 < \beta < 1 < \alpha$. Denote $\theta = \log_\alpha \left(\frac{1}{\beta} \right) = \frac{-\ln(\beta)}{\ln(\alpha)} > 0$. Then, for any choice $K \in \mathbb{Z}$, the following bound holds as long as $\varepsilon \leq \frac{\theta}{\alpha^{K(1+\theta)}}$:

$$\min_{k \geq K} D_k \leq \max \left(\sqrt{\alpha}, \sqrt{\frac{1}{\beta}} \right) \times \left(\theta^{\frac{1}{1+\theta}} + (1/\theta)^{\frac{1}{1+\theta}} \right) \times \varepsilon^{\frac{\theta}{1+\theta}}.$$

Proof We will later on find the optimal bound on the right assuming $k \in \mathbb{R}$. Then, by replacing k with the nearest integer we will either increase the power of α by $\frac{1}{2}$ or decrease the one of β by $\frac{1}{2}$. Note that this bound works best under the assumption that $\alpha \approx \frac{1}{\beta}$. If one is much bigger than the other, we may simply decide on which side we always round k , with an added factor α or $\frac{1}{\beta}$ instead.

Now, consider the parameter $x := \alpha^k \in \mathbb{R}_{>0}$. Thus, $k = \log_\alpha(x)$ so:

$$\beta^k = \exp \left(\frac{\ln(x)}{\ln(\alpha)} \times \ln(\beta) \right) = \exp(-\theta \ln(x)) = x^{-\theta}.$$

With this rewriting, $D(x) := \varepsilon x + x^{-\theta}$ is much easier to minimise. Indeed, D can be seen as a positive convex function that goes to $+\infty$ on 0^+ and $+\infty$, hence is minimised when $D'(x) := \varepsilon - \theta x^{-\theta-1} = 0$, thus at $x = \left(\frac{\theta}{\varepsilon} \right)^{\frac{1}{\theta+1}}$. Using this value of x in D directly gives us the rest of the expected bound.

Now, for the domain of validity, for us to be able to round k properly, we simply require $k = \log_\alpha(x) \geq K$, which translates as $\varepsilon \leq \frac{\theta}{\alpha^{K(1+\theta)}}$. When the bound doesn't hold, when K is greater than the optimal value, the optimal choice is simply D_K .

Corollary 4.4 (Polynomial Stability) Assume there is a sequence of triplets (p_n, ρ_n, C_n) for which Proposition 4.1 applies to $\Omega_{\mathcal{F}}$. If $C_n + p_n = O\left(\alpha^{\frac{n}{d}}\right)$ and $\rho_n = O(\beta^n)$, then using the previous lemma gives us a polynomial bound $O\left(\varepsilon^{\frac{\theta}{1+\theta}}\right)$ on the speed of convergence, with $\theta = \frac{-\ln(\beta)}{\ln(\alpha)}$.

Remark 4.1 To illustrate how this framework applies, let us use it to obtain the polynomial stability for the enhanced Robinson tiling.

Unlike the usual Robinson tiling, the enhanced variant enforces alignment of neighbouring macro-tiles in a local way. At the scale of N -macro-tiles, if we forget about the grid around these tiles, of density $\rho_n = 1 - \frac{(2^n - 1)^2}{4^n} = O\left(\frac{1}{2^n}\right)$, we obtain a p_n -periodic tiling with $p_n = 2 \times 2^n$. What's more, we can prove the tiling has C_n -reconstruction [1, Proposition 7.8], with a radius $C_n = 2^n - 1$. As we have $C_n + p_n = O(2^n) = O\left(4^{\frac{n}{2}}\right)$, we can apply the previous corollary with parameters $(\alpha, \beta) = \left(4, \frac{1}{2}\right)$, so $\theta = \frac{1}{2}$ and $\frac{\theta}{1+\theta} = \frac{1}{3}$. Hence, we fall back on the $O\left(\sqrt[3]{\varepsilon}\right)$ bound of the previous article [1, Theorem 7.10] (with a comparable multiplicative constant) which is to be expected as we basically generalised the scheme of the proof used in that paper.

More generally, in a hierarchical tiling, at the scale of “macro-tiles” of diameter x , the typical reconstruction radius we may hope for is of order x at least (*i.e.* the size of a macro-tile), and likewise for the quasi-periodicity. Conversely, among the x^d cells in a macro-tile, we may have to ignore at least a one-dimensional “wire” that crosses the whole macro-tile, hence ρ of order $\frac{1}{x^{d-1}}$ at least. Following the same general computations as in the previous lemma, we conclude that in d dimensions, the best speed of convergence we may obtain is $O\left(\varepsilon^{\frac{d-1}{2d-1}}\right)$. With $d = 2$, we have $\frac{d-1}{2d-1} = \frac{1}{3}$, the order of convergence obtained for the enhanced Robinson tiling. The question of whether we can obtain a faster bound for the convergence speed of aperiodic tilings, whether by improving upon the minimal values of (C, p, ρ) conjectured here (and in particular on the C -reconstruction), or by using another method altogether, is still open.

5 Undecidability of the Stability

In the previous sections, we showed how a simple bi-coloured tiling can be unstable, and how a class of well-behaved quasi-periodic SFTs can be stable. We will now make full use of these ideas in order to equate the notion of stability with some undecidable problems in the arithmetical hierarchy through the emergence of said unstable structure.

Matter-of-factly, proving Π_2 -hardness would directly imply the weaker bounds we introduce first. However, the Π_2 -hard construction relies on the Σ_1 -hard one, and we believe the Π_1 -hard one uses a complementary and more intuitive idea that will help get the point across.

5.1 Π_1 -hard Construction

First, we will make use of the halting problem P_{halt} . What we want to do here is to encode computations into the Robinson tiling in a *stable* way, and make an *unstable* phase emerge *iff* the machine terminates. This will equate the Σ_1 -complete halting problem with instability among a class of SFTs, hence Π_1 -hardness of P_{stab} in general.

In the previous Red-Black example of Section 3, the main ingredient allowing instability was the existence of two kinds of n -macro-tiles at any scale (widely different for the finite Hamming distance) instead of just one (four similar tiles, up to the orientation of their low-density central cross) in the monochromatic case. The two kinds of macro-tiles cannot coexist in the same generic Robinson configuration, but we can replace one with the other for a small price in the presence of noise.

5.1.1 Description of the Tiletset

Let us first describe the tileset used in this section. We won't explain in details how Turing machines can be implemented inside the Robinson tiling, but the interested reader may look at the original article by Robinson [4] or lecture notes by Jeandel and Vanier [6] for a formal study of this simulation result.




Fig. 5 Main tiles of the alphabet \mathcal{A}_R

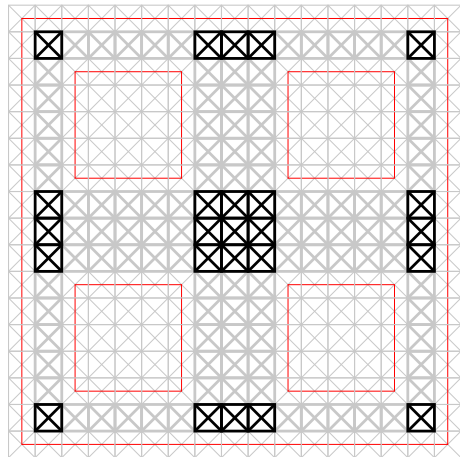
We will use a variant of this construction more suited to our needs, with two layers, defined on the alphabet $\mathcal{A} \subset \mathcal{A}_R \times \mathcal{A}_M$, where \mathcal{A}_R stands for the common Robinson layer, and \mathcal{A}_M for the layer specific to a given Turing machine M . Consequently, we will denote $\Omega_{P_1(M)}$ the corresponding SFT.

Let's first describe the common layer \mathcal{A}_R . As we can see in Fig. 5, the tileset uses four main colours, as well as grey dotted and dashed lines. These grey lines must match with one of the same type (either dotted *or* dashed on both sides of an edge), and serve to enforce alignment of the Robinson macro-tiles locally, to guarantee stability of the structure itself, just like for the enhanced Robinson tiling [1, Proposition 7.8]. Notice how bumpy-corners must be Black, after which we alternate between Black and Red. At some point, to-be-decided by the layer \mathcal{A}_M , we may transition from the Red-Black (stable) regime to the Blue-Green (unstable) regime using one of the two transition tiles on the bottom-right of Fig. 5. The whole set \mathcal{A}_R is given by all the rotations of the first three columns containing the tiles with a monochromatic V-shaped arm (but no symmetry, so that we may preserve the chirality of macro-tiles, so that each arm of the central cross may indicate the overall orientation of the macro-tile) and rotations and symmetries of the rest, which brings us to a total of $|\mathcal{A}_R| = 172$ tiles.

Now, without detailing the intricacies of \mathcal{A}_M and how it is coupled with \mathcal{A}_R in \mathcal{A} , let us give the general idea and specificities of our construction. Here, each Red square (of length $4^n + 1$, in the center of a $(2n + 1)$ -macro-tile) will contain a limited space-time diagram of the Turing machine M with a semi-infinite tape, while avoiding smaller red squares which contain their own space-time diagram. This is illustrated in Fig. 6, where the black crossed cells represent the patches of space-time diagram, and the grey cells are communication channels that synchronise the otherwise disconnected patches of the diagram. The n -th scale of simulation, occurring in a $(2n + 1)$ -macro-tile, thus has a space-time horizon of $2^n + 1$ tiles, initiated on the empty input on the bottom row.

The main difference with the canonical construction is how it behaves when M stops. In Robinson's article, the tiling doesn't allow for M to stop, in order to prove that the tileability problem is undecidable. Here, when M halts in the n -th scale of simulation, it idles until the border of the square can "notice" the halting, and decide freely whether it will force a transition from Black to Blue or Green on its border. After which, at higher scales, no more computations occur. Still, whether or not this

Fig. 6 Space-time diagram of a Turing machine in a 5-macro-tile



transition occurs, we have arbitrarily big macro-tiles, and thus $\Omega_{P_1(M)} \neq \emptyset$. Note that a description of M can be algorithmically converted into the set of forbidden patterns $P_1(M)$ (and its corresponding alphabet) in finite time.

Theorem 5.1 [P_{stab} is Π_1 -hard] Consider a Turing machine M . Then the SFT defined by $P_1(M)$ is stable ($P_1(M) \in P_{stab}$) iff M does not halt on the empty input ($M \notin P_{halt}$). As P_{halt} is Σ_1 -complete, we deduce that P_{stab} is Π_1 -hard.

The following subsections will each focus on one of the implications, which put together directly give the previous result.

5.1.2 The Stable Case

For the stable case, assume that $M \notin P_{halt}$. Because of this, at any scale of admissible macro-tiles, the previously described transition from the Red-Black to the Blue-Green regime cannot occur, and the two last lines of tiles in Fig. 5 may as well not exist in \mathcal{A}_R . Our goal is to prove that the framework of Section 4 applies here.

Notice how we can project the alphabet \mathcal{A} onto its first coordinate \mathcal{A}_R and then erase the information on which of the four colours is used for the lines atop of a tile. This way, we fall back on the enhanced Robinson tiling studied in our previous paper. In particular, the following structural result applies:

Proposition 5.2 ([1, Proposition 7.8]) Consider the enhanced Robinson SFT. Let us denote $B_k = \llbracket -k, k \rrbracket^2$. For any scale of macro-tiles $n \geq 2$, the constant $R_n = 2^n - 1$ is such that, for any $k \geq 0$ and any clear locally admissible pattern ω on B_{k+R_n} , its restriction $\omega|_{B_k}$ is made of well-aligned and orientated n -macro-tiles, plus the grid around them which we do not control.

Thence, at the n -th scale of simulation (i.e. in $(2n + 1)$ -macro-tiles), the tiling $\Omega_{P_1(M)}$ is $\frac{1}{4^n}$ -almost p_n -periodic (with $p_n = 4^{n+1}$) with C_n -reconstruction (defined as $C_n = R_{2n+1}$) if we specifically look at the layer \mathcal{A}_R . However, we need to be a bit

more careful to obtain the desired periodic behaviour on the other coordinate of the alphabet \mathcal{A} , and we will actually specifically extend the grid around $(2n + 1)$ -macro-tiles into a larger set G_n to do so.

Lemma 5.3 *Using the constant choices from the previous paragraph, the SFT $\Omega_{P_1(M)}$ is ρ_n -almost p_n -periodic with C_n -reconstruction, with G_n the area outside of Red squares up to the n -th scale and ρ_n its density (that will be computed afterwards).*

Proof Because $\Omega_{P_1(M)}$ enforces alignment in a local way, for any tiling ω we obtain the same set G_n (up to translation) by looking at all the tiles outside of Red squares up to the n -th scale of simulations. This G_n is p_n -periodic, and in particular includes the grid surrounding $(2n + 1)$ -macro-tiles so that we have ρ_n -quasi p_n -periodicity with C_n -reconstruction on the layer \mathcal{A}_R .

Regarding alignment, notice that G_n has the same periodicity as the grid around $(2n + 1)$ -macro-tiles, whose alignment is fixed by the C_n -reconstruction on the layer \mathcal{A}_R , hence its alignment is fixed in the same way.

Remark that, on the layer \mathcal{A}_M , because the Turing machine is deterministic, everything that happens on the *inside* of a given admissible Red square is fixed, insulated from outside interference. Hence, on this layer (and using of course the alignment of G_n given by the layer \mathcal{A}_R) we obtain a $\frac{p_n}{2}$ -periodic behaviour outside of G_n , as it does not depend on the orientation of the $(2n + 1)$ -macro-tiles. \square

In order to conclude, we need to compute ρ_n the density of G_n .

Lemma 5.4 *In a $(2n + 1)$ -macro-tile, we have $O(12^n)$ tiles outside of the Red squares.*

Proof The general idea of the proof is that Red squares form a kind of Sierpiński carpet inside macro-tiles.

Denote r_n the number of tiles *inside* the Red squares in a $(2n + 1)$ -macro-tile. As we can see on Fig. 3, in the process of forming a $(2n + 3)$ -macro-tile, we will create a big central square around four $(2n + 1)$ -macro-tiles, surrounded by twelve $(2n + 1)$ -macro-tiles. As we already know the size of this big square, we obtain the following recurrence:

$$r_{n+1} = 12r_n + (4^{n+1} + 1)^2 \geq 12r_n + 16^{n+1}.$$

As $r_1 = 25 \geq 16$, we obtain by induction $r_n \geq 4^{n+1} (4^n - 3^n)$. At the same time, a $(2n + 1)$ -macro-tile has $(2^{2n+1} - 1)^2 \leq 4^{2n+1}$ tiles, so *at most* $4^{n+1} 3^n = 4 \times 12^n$ tiles in total outside the Red squares. \square

Hence, as $(2n + 1)$ -macro-tiles use $\Theta(16^n)$ tiles in total, we conclude that G_n has density $\rho_n = O\left(\left(\frac{3}{4}\right)^n\right)$.

Proposition 5.5 *Consider $M \notin P_{halt}$. Then $\Omega_{P_1(M)}$ is polynomially stable, with convergence speed $O(\varepsilon^r)$ at rate $r = \frac{2 - \log_2(3)}{6 - \log_2(3)} \approx 0.094$.*

Proof We apply Corollary 4.4, with constants $\alpha = 4$ and $\beta = \frac{3}{4}$, so $r = \frac{\theta}{1+\theta}$ gives the announced rate. \square

5.1.3 The Unstable Case

Proposition 5.6 *Assume $M \in P_{halt}$. Then for any $\varepsilon > 0$ we have a measure $\mu \in \mathcal{M}_{P_1(M)}^B(\varepsilon)$ such that $d_B(\mu, \mathcal{M}_{P_1(M)}) \geq \frac{1}{4^{n+1}}$, where n denotes the last scale of simulation, at which M halts.*

Proof Consider $N = 2(n + 1)$ the first scale at which N -macro-tiles have a big Blue or Green square in the middle. Assuming two aligned N -macro-tiles don't use the same colour for the square (of diameter $d = 2^{N-1} + 1$ tiles), then we obtain at least $p = 4 \times (d - 1) = 2^{N+1}$ differences.

By following the very same colour-flipping process as in Proposition 3.1, but on the Blue-Green bit starting at the scale of N -macro-tiles, we obtain a generic colour-flipped configuration ω (with monochromatic Blue or Green squares in the N -macro-tiles).

Thus, for any generic $\omega \in \Omega_{P_1(M)}$ that aligns with ω' up to the scale of N -macro-tiles, we obtain a lower bound $d_H(\omega, \omega') \geq \frac{1}{2} \times \frac{p}{4^N} = \frac{1}{2^N} = \frac{1}{4^{n+1}}$, with the factor $\frac{1}{2}$ coming from the frequency of Blue and Green big squares in ω' , whereas *all* such squares of ω must be of the same colour.

Now, assume that N -macro-tiles in ω and ω' don't align well. By choosing the best pairing of N -macro-tiles between ω and ω' , we still have a rectangle with both sides of length at least $2^{N-1} - 1$ (the size of a $(N - 1)$ -macro-tile) where the N -macro-tiles of both tilings overlap. In this area, both macro-tiles have a Blue or Green corner of their big square, made of at least $2 \times (2^{N-1} + 1)$ tiles. As these two corners intersect in at most 2 tiles, and the rest of the area is guaranteed to use only Black or Red communication channels, we have at least 2^N differences between ω and ω' in this window. As this process repeats 2^N periodically in both directions, without even having to take the colour-flipping into account, we still obtain a lower bound $d_H(\omega, \omega') \geq \frac{1}{2^N} = \frac{1}{4^{n+1}}$. \square

Remark 5.1 More generally, as long as we can guarantee *one* difference between the two kinds of macro-tiles which we colour-flip, we obtain a lower bound on d_B of order $\frac{1}{\text{tile area}}$. We will directly invoke this “obvious” lower bound for further unstable cases.

Still, the order of magnitude $\frac{1}{\text{tile diameter}}$ obtained in the previous proposition is the best one can reasonably hope for in general, as a signal that transits through a macro-tile will typically only cross a number of tiles proportional to the diameter, normalised by the tile area.

5.2 Σ_1 -hard Construction

We can “flip around” the previous construction, by adding an *unstable* information atop of the structure simulating the Turing machine, in such a way that the information gets frozen and becomes *stable* if the machine halts. We will first describe the construction of $S_1(M)$ out of a machine M , and then state the corresponding undecidability result.

In the previous tileset $P_1(M)$, the Robinson layer \mathcal{A}_R used one communication channel with four different colours. Here, for $S_1(M)$, we use *two* communication channels in the lines of the Robinson structure, each one having two possible values. First, the Red-Black channel must be initialised as Black in bumpy corners, and then

alternate, in order to have the right structure to simulate the machine M . Second, the Blue-Green channel can be freely initialised. However, if M halts at a given scale of simulation, then the border of the Red square *must* be Blue on the other channel, which we call a *freeze*. Note that here, we can keep simulating M at higher scales after it halts for the first time, as subsequent freezes will just occur at scales of macro-tiles where the Blue-Green channel would be frozen into Blue anyway.

Proposition 5.7 *We have $S_1(M) \in P_{stab}$ iff $M \in P_{halt}$. Thus, P_{stab} is Σ_1 -hard.*

Proof First, assume that $M \notin P_{halt}$. Then we can freely do a colour-flipping process starting from any $\mu \in \mathcal{M}_{S_1(M)}$, just like in Proposition 3.1. We can start flipping the Blue-Green channel at the scale of bumpy corners, hence instability with a $\frac{1}{8}$ lower bound on d_B .

Now, assume $M \in P_{halt}$. Then, in any tiling $\omega \in \Omega_{S_1(M)}$, the Blue-Green channel is retroactively frozen all the way down to the Green bumpy-corners. By using the same grid G_n as in Lemma 5.3, we can likewise ignore everything that happens outside of Red squares, and control everything inside, hence a ρ_n -almost p_n -periodic tiling with $p_n = O(4^n)$ and $\rho_n = O\left(\left(\frac{3}{4}\right)^n\right)$.

Finally, denote n_{halt} the first scale of simulation at which M halts in $S_1(M)$. If we try to reconstruct things locally at steps lower than n_{halt} , then we will reach a family of well-aligned and well-oriented $(2n+1)$ -macro-tiles, but without any freezing happening in the tiles, hence this Blue-Green channel that may not behave in a globally admissible way, all the way down to the high-density set of bumpy-corners. Still, as long as $n \geq n_{halt}$, the freezing prevents this from happening, and using the same $C_n = O(4^n)$ as in Lemma 5.3, we conclude that this scale of the tiling has indeed C_n -reconstruction.

Still, starting at high-enough scales, for low-enough values of ε , the proof of Proposition 5.5 applies *verbatim*, so we have stability with a polynomial $O(\varepsilon^r)$ convergence rate. \square

5.3 Π_2 -hard Construction

In the construction for $S_1(M)$, we obtained stability *iff* there exists a time step such that M halts *on the empty input*. Consequently, if we manage to twist the construction to include *any possible input*, then we may equate stability with the Π_2 -complete totality problem P_{total} .

There are several ways to proceed, but we choose here to use the method of Toeplitz encoding of the input, because it is quite versatile, and may more generally be able to convert a (structurally close to) uniquely ergodic SFT encoding a Σ_k -hard problem into a (definitely not uniquely ergodic anymore) SFT encoding a Π_{k+1} -hard problem.

5.3.1 Toeplitz Input

The Toeplitz encoding of an infinite sequence $u \in \Gamma^{\mathbb{N}^*}$ on an alphabet Γ (with \mathbb{N}^* the set of positive integers) consists of inductively filling with u_n half of the holes still

free after the previous iterations, which gives a sequence $u_1 * u_1 * u_1 * u_1 * \dots$, then $u_1 u_2 u_1 * u_1 u_2 u_1 * \dots$, and so on. Toeplitz sequences have been studied as dynamical systems for a long while [25], and have since been encoded in higher-dimensional SFTs [26].

The idea of the method is to sequentially write the wanted input u into the consecutive scales hierarchical structure, which will appear as a Toeplitz encoding $u_1 u_2 u_1 \dots$ from the point of view of the simulated Turing machine, and then adapt the machine to decode it back into its original form u at first. This method was already used by Barbieri and Sablik [27] in particular.

More precisely, we build the tileset $P_2(M)$ as follows. For the Robinson structure, we use the same parallel Red-Black and Blue-Green bits as for $S_1(M)$. We add another channel that can take values in $\Sigma \sqcup \{\#, \$^\Sigma, \$^\#\}$ where Σ is the input alphabet of the machine M , $\#$ the blank tape symbol, and the $\* symbols two supplementary letters. On Black channels, we can freely use any symbol $\* following a letter from Σ on the previous Red scale, but we *must* use $\$^\#$ following $\#$. On Red channels, we *must* use a letter from Σ following a $\$^\Sigma$ symbol on the previous Black scale, and use $\#$ following $\$^\#$. If we look only at the Red channels, this gives an infinite word $u \in \Sigma^* \#^\mathbb{N} \sqcup \Sigma^\mathbb{N}$. When $u \in \Sigma^* \#^\mathbb{N}$, we will identify it with its prefix in Σ^* , followed by $\#^\mathbb{N}$.

Quite importantly, the choice of a letter is not only communicated along the regular Red-Black channels in two directions from the center of a macro-tile arm, but also along the alignment channels of the enhanced Robinson self-aligning structure, the dotted and dashed lines in the other two directions. Thus, any two (well-aligned) neighbouring N -macro-tiles must encode the same sequence.

On the simulation layer, the Turing machine is able to read which symbol is written down in the column on the right of its current position. Hence, from the point of view of the Turing machine simulated in a Red square, this represents a *read-only* second tape. In order to adequately use u as an input, we first need to explain what the machine sees.

Lemma 5.8 (*Toeplitz Encoding of the Input*) *Let $u \in \Sigma^* \#^\mathbb{N} \sqcup \Sigma^\mathbb{N}$. Define the words $w_n = w_{n-1} u_n w_{n-1}$ by induction, initialised with the empty word w_0 . The word w_n is a prefix of the Toeplitz encoding of the whole sequence u .*

At the n -th scale of simulation, from the point of view of the Turing machine, the read-only tape reads as $w_{n-1} \$^ \$^* w_{n-1} u_n$.*

Proof The last letter of the read-only tape obviously corresponds to the right border of the n -th Red square, hence reads as u_n . The central $\* symbols come from the fact, as highlighted by the blue columns in Fig. 7, they correspond to the $(n + 1)$ -th scale for Black squares followed by the first scale of bumpy corners.

The rest of the word, that reads as w_{n-1} on both ends of this central line, can be explained by the inductive construction of macro-tiles. Indeed, each quarter of the n -th Red square is actually a whole $(2n - 1)$ -macro-tile with a central Red square, and the Red squares are themselves stacked in a Toeplitz way within the macro-tile, with a gap in-between each that allows to read the letter on them. \square




Fig. 7 Structure of the read-only input, with read-only values stored in the highlighted columns, whereas the machine operates within the black patches

5.3.2 From Decoding the Input to Computations

Let us explain what Turing machine is encoded into $P_2(M)$, and how it affects the Blue-Green channel.

First, the machine will have to *decode* the Toeplitz input, while keeping the Blue-Green channel *stable* (by using a third non-alternating colour). More precisely, the machine will step by step read the letters on the read-only tape, with the k -th letter being at position 2^k , and write them one after another at the beginning of its working tape. This process will decode the Toeplitz encoding w_n back into the sequence $u_1 \dots u_n$. Using a unary counter, which we multiply by two after reading each letter, reaching the k -th letter will require about $\Theta(4^k)$ steps of computation.

Now, this process can halt in two ways. First, we read a $\* symbol, meaning that we reached halfway through the read-only tape. In this case, the machine simply idles for the rest of its finite runtime, without unfreezing the Blue-Green bit when it reaches the top border of the Red square. This won't happen at big-enough scales of simulation, considering it would take about $\Theta(4^n)$ steps but the n -th machine only has a finite horizon of 2^n steps, but it can occur at the initial scales of simulation and in particular at the very first one where the first symbol is $\* . Second, we read a $\#$ symbol before reaching the central $\* , in which case the decoding of the word $u \in \Sigma^*$ is complete. Without waiting, the machine starts then simulating M on u (this will occur roughly at the $2|u|$ -th scale of simulation). This will signal the Red square to *ignite* the unstable

Blue-Green bit (if it was not already done at a lower scale), as was the case for $P_1(M)$ in Subsection 5.1. Then, if M halts on u , this will signal the Red square to freeze the Blue-Green bit, as was the case for $S_1(M)$ in Subsection 5.2.

5.3.3 Undecidability of the Stability

Lemma 5.9 *Assume $M \notin P_{total}$ does not halt on the input $u \in \Sigma^*$. Consider $\mu_u \in \mathcal{M}_{\mathcal{F}}$ an invariant measure with $u^{\#\mathbb{N}}$ written in the Red scales of any generic configuration. Then, by colour-flipping the Blue-Green channel after the initial decoding scales, we obtain the measure $\mu_u^\varepsilon \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$, such that $\inf_{\varepsilon > 0} d_B(\mu_u^\varepsilon, \mathcal{M}_{\mathcal{F}}) > 0$.*

Proof As in Proposition 5.6, if we compare two macro-tiles with a Blue or Green square, corresponding to the same input $u^{\#\mathbb{N}}$, we can obtain a lower bound on their density of mismatching purely through the Blue-Green square.

Likewise, if we compare such a macro-tile with a macro-tile corresponding to another input, then they must differ in one of the first $|u| + 1$ letters in a Red channel. In this case, we can also obtain a lower bound independent of ε , even if they are perfectly aligned, using this mismatching letter in the input. \square

Proposition 5.10 *Denote $\varphi(n)$ the first scale of simulation at which, for any input $u \in \Sigma^{\leq n}$, both decoding and computation are over. Assume $M \in P_{total}$, so that $\varphi(n) < \infty$. Then, using the notations of Lemma 5.3, the SFT $\Omega_{P_2(M)}$ is ρ_n -almost p_n -periodic with $C_{\varphi(n)}$ -reconstruction.*

Proof If we follow the same scheme of proof as in the lemma, then we almost obtain a ρ_n -almost p_n -periodic SFT with C_n -reconstruction up to one detail.

Here, the Red-Black channel and all the computations in Red squares behave deterministically, so they are fixed for a given input (which synchronises between neighbouring tiles), but the Blue-Green channel is not. However, if we exploit the $C_{\varphi(n)}$ -reconstruction of the Robinson structure, then either:

- a given Red square isn't done decoding its input, so the Blue-Green bit is still frozen, uniquely determined,
- a given Red square has decoded its input u , at a scale of simulation lower than n , which implies $|u| < n$, but the Red square actually fits into a bigger $(2\varphi(n) + 1)$ -macro-tile, which will terminate its simulation of M on u , thus freeze the Blue-Green bit of this Red square.

In both cases, we indeed guaranteed that the area inside Red squares is globally admissible, hence the n -th scale of simulation admits $C_{\varphi(n)}$ -reconstruction. \square

In particular, $\rho_n \rightarrow 0$ so $\Omega_{P_2(M)}$ is stable according to Corollary 4.2. However, because φ can be roughly as big as any computable function, we can't possibly exhibit a good bound to apply Corollary 4.4, and will not obtain polynomial stability this time. The next theorem directly follows:

Theorem 5.11 *Let M a Turing machine. We have $M \in P_{total}$ iff $P_2(M) \in P_{stab}$. As P_{total} is Π_2 -complete, we deduce P_{stab} is Π_2 -hard.*

Note how this process doesn't adapt to translate the Π_1 -hard construction into a Σ_2 -hard one. In order to do this, we would need an added universal quantifier, which cannot work if we encode only *one* input at a time in a ground configuration. Hence, in this case we would need to enumerate the inputs inside the tiling in any case.

Remark 5.2 (Alternate Construction for Π_2 -Hardness) Let us conclude this section by briefly describing another construction relating P_{stab} to P_{total} , this time without having to encode any input.

The main idea is here to stack ignition-freezing blocks onto each other. In the tiling $P'_2(M)$, we enumerate the words of Σ^* , *e.g.* following a lexicographical order biased by increasing lengths. After enumerating a new word u we simulate M on it. Once this simulation ends, we both freeze the lower scales of the Blue-Green bit and ignite an independent Blue-Green bit for higher scales. Then, we enumerate the next word, rinse and repeat.

If $M \notin P_{total}$, then M will never end computing on u , never freeze this Blue-Green block which we will be able to colour-flip. If $M \in P_{total}$, then at any given scale of simulation, there exists a higher scale of simulation at which M terminates on *some* word, which will guarantee the current scale is frozen.

Remark 5.3 (Extension to Higher Dimensions) In our previous article [1, Section 4], we established that in the one-dimensional case, stability is decidable. In this section, we just established that in the two-dimensional case, stability is undecidable, and Π_2 -hard. In higher dimensions $d > 2$, we can simply use an independent copy of $\Omega_{P_2(M)}$ for each coordinate of \mathbb{Z}^{d-2} , which gives a d -dimensional SFT that is stable *iff* $\Omega_{P_2(M)}$ is [1, Corollary 4.13], thus we have a computable reduction from totality to stability for *any* fixed dimension $d \geq 2$.

6 Π_4 Upper Bound on the Stability

As announced, we will now need to dig deeper into the framework of computable analysis on measures to describe how much computation power is actually needed to decide our notion of stability.

The general idea of computable analysis is to study problems relating to real numbers, involving continuous functions or differential equations for example, from the point of view of effective computations [28]. Here, we are interested in doing computable analysis specifically on *probability measures*. The topic has already been studied [29–31] but, given the lack of a widespread theory, we will introduce all the needed notations and keep things self-contained.

For the rest of this section, we will consider an alphabet \mathcal{A} and a set of forbidden patterns \mathcal{F} , without any more assumptions such as $\Omega_{\mathcal{F}} \neq \emptyset$. From there, our goal in this section is to explain a process to conclude on whether $\mathcal{F} \in P_{stab}$ or not.

Even though we are interested in convergence for the Besicovitch distance d_B , we will actually need to use the weak- $*$ topology. Indeed, endowed with this topology, the set of σ -invariant measures on \mathcal{A} (*i.e.* the full-shift $\mathcal{M}_{\mathcal{A}}$) admits an explicit countable dense basis, which is the bedrock upon which most of computable analysis relies. Hence, before doing anything meaningful with this topology, we will first introduce

our notations to work with it in Subsection 6.1, and in particular the (family of) computable distances we will use later on.

Once this preliminary work is done, we will see how the measure sets $\mathcal{M}_{\mathcal{F}}$, $\mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ and $J(\mu, \mathcal{M}_{\mathcal{F}})$ can be described in this framework. At last, we will use these descriptions to prove a Π_4 upper bound on the problem P_{stab} .

6.1 A Crash Course in Computable Analysis

Definition 6.1 (Weak-* Topology) The weak-* topology on a set of probability measures $\mathbb{P}(\mathcal{A}^X)$ with \mathcal{A} finite and X countable is defined as follows. We have the convergence $\mu_n \xrightarrow{*} \mu$ when, for any finite subset $U \subseteq X$, and any pattern $w \in \mathcal{A}^U$, we have the convergence $\mu_n([w]) \rightarrow \mu([w])$.

Definition 6.2 The weak-* topology is metrisable, induced notably by:

$$d_r^+(\mu, \nu) := \sum_{n \in \mathbb{N}} \frac{1}{2^n} \times \frac{1}{r^{|U_n|}} \times \frac{1}{|\mathcal{A}^{U_n}|} \sum_{w \in \mathcal{A}^{U_n}} |\mu([w]) - \nu([w])|,$$

with $(U_n)_{n \in \mathbb{N}}$ an increasing sequence of sets that covers the space X , and $r \geq 1$ a normalisation factor (with the convention $d^+ = d_1^+$). In particular, when $X = \mathbb{Z}^d$, we can take $U_n = B_n$. $\mathbb{P}(\mathcal{A}^X)$ is a compact space for this topology. When $X = \mathbb{Z}^d$, the space of σ -invariant measures $\mathcal{M}_{\mathcal{A}}$ is a closed subspace. In this subspace, we can and will instead use $U_n := \llbracket 0, n \rrbracket^d$ to define d^+ for the rest of this article.

Note how, if $V_1 \subset \mathcal{M}_{\mathcal{A}_1}$ and $V_2 \subset \mathcal{M}_{\mathcal{A}_2}$ are both closed, then so is the set of their joinings $J(V_1, V_2) \subset \mathcal{M}_{\mathcal{A}_1 \times \mathcal{A}_2}$ (in their respective weak-* topologies).

Definition 6.3 (Closed Ball) We denote the closed ball around $\mu \in \mathcal{M}_{\mathcal{A}}$ of radius $\varepsilon > 0$ as $\overline{\mathbb{B}}(\mu, \varepsilon) = \{ \nu \in \mathcal{M}_{\mathcal{A}}, d^+(\mu, \nu) \leq \varepsilon \}$, and this definition extends to the ε -neighbourhood of any set of measures.

Definition 6.4 (Periodic Measure) For $w \in \mathcal{A}^{U_n}$, we denote $w^{\mathbb{Z}^d} \in \Omega_{\mathcal{A}}$ the configuration obtained through periodic repetition of w in each direction, and then $\widehat{\delta}_w = \frac{1}{|U_n|} \sum_{k \in U_n} \delta_{\sigma_k(w^{\mathbb{Z}^d})}$ the corresponding σ -invariant measure. We call such measures *periodic*.

Lemma 6.1 (Covering Lemma for $\mathcal{M}_{\mathcal{A}}$) There is $\psi : \mathbb{Q}_{>0} \times \mathbb{N}^2 \rightarrow \mathbb{N}$ a computable map such that, for any finite alphabet \mathcal{A} , any dimension d and any rational $\delta > 0$:

$$\mathcal{M}_{\mathcal{A}} = \bigcup_{w \in \mathcal{A}^{U_{\psi(\delta, |\mathcal{A}|, d)}}} \overline{\mathbb{B}}(\widehat{\delta}_w, \delta).$$

Proof Denote s_n the partial sum up to rank n associated to the distance d^+ . We can bound $r_n = d^+ - s_n \leq \frac{1}{2^{n-1}}$ independently of the dimension d , of \mathcal{A} and of any pair of measures. Hence, for a given value of δ , we first compute $n(\delta) = 2 + \lceil \log_2(\frac{1}{\delta}) \rceil$, such

that $r_n \leq \frac{\delta}{2}$. We now want to cover $\mathcal{M}_{\mathcal{A}}$ with balls of radius $\frac{\delta}{2}$ for the pseudo-distance s_n .

Notice that for any $k \leq n$ and any pattern $w \in \mathcal{A}^{U_k}$, we have the decomposition $\mu([w]) = \sum_{v \in \mathcal{A}^{U_n}, v|_{U_k} = w} \mu([v])$. It follows that:

$$\begin{aligned} s_n(\mu, \nu) &= \sum_{k \leq n} \frac{1}{2^k |\mathcal{A}^{U_k}|} \sum_{w \in \mathcal{A}^{U_k}} |\mu([w]) - \nu([w])| \\ &\leq \sum_{k \leq n} \frac{1}{2^k} \sum_{w \in \mathcal{A}^{U_k}} |\mu([w]) - \nu([w])| \\ &\leq \sum_{k \leq n} \frac{1}{2^k} \sum_{w \in \mathcal{A}^{U_n}} |\mu([w]) - \nu([w])| \\ &\leq 2 \sum_{w \in \mathcal{A}^{U_n}} |\mu([w]) - \nu([w])| \\ &\leq 2 |\mathcal{A}^{U_n}| \sup_{w \in \mathcal{A}^{U_n}} |\mu([w]) - \nu([w])|. \end{aligned}$$

Hence, we now need to uniformly approximate any $\mu \in \mathcal{M}_{\mathcal{A}}$ on the window U_n by a periodic measure to conclude.

To do so, consider μ_m the restriction of μ to U_m . We identify μ_m , a measure on \mathcal{A}^{U_m} , with the measure on $\Omega_{\mathcal{A}}$ that charges a periodic word $w^{\mathbb{Z}^d}$ (with $w \in \mathcal{A}^{U_m}$) with probability $\mu([w]) = \mu_m(\{w\})$. Remark that μ_m is not σ -invariant, but is $m\mathbb{Z}^d$ -periodic under translation, so we can define the corresponding averaged measure $\widehat{\mu}_m := \sum_{w \in \mathcal{A}^{U_m}} \mu_m(\{w\}) \times \widehat{\delta}_w$ which is σ -invariant.

In particular for any $w \in \mathcal{A}^{U_n}$, as long as $k + U_n \subset U_m$ (i.e. $k \in U_{m-n}$), then $[w]_k := \sigma_k([w])$ is still a cylinder defined inside U_m . Hence, for any such translation we have $\mu_m([w]_k) = \mu([w]_k) = \mu([w])$. Now, we have:

$$\widehat{\mu}_m([w]) := \frac{1}{|U_m|} \sum_{k \in U_m} \mu_m([w]_k) = \frac{|U_{m-n}|}{|U_m|} \mu([k]) + \frac{1}{|U_m|} \sum_{k \in U_m \setminus U_{m-n}} \mu_m([w]_k),$$

hence $\widehat{\mu}_m([w]) = \mu([w]) + O\left(\frac{n}{m}\right)$, where the computable domination constant depends on d . Now, if we use instead μ_m^k a dyadic approximation of μ_m on \mathcal{A}^{U_m} , with precision $\frac{1}{2^k}$, we obtain a measure μ_m^k for which:

$$\widehat{\mu}_m^k([w]) = \mu([w]) + O\left(\frac{n}{m}\right) + O\left(\frac{|\mathcal{A}^{U_m}|}{2^k}\right).$$

This new term simply uses the domination constant 1. Remark in particular that there is only a finite amount of such dyadic measures on the window U_m with precision $\frac{1}{2^k}$. We just need to be able to approximate these by periodic measures to conclude.

We can decompose any such dyadic measure as $\widehat{\mu}_m^k = \frac{1}{2^k} \sum_{w \in \mathcal{A}^{U_m}} p(w) \widehat{\delta}_w$ with weights $p(w) \in \mathbb{N}$ that sum to 2^k . Consider now $M = (m + 1) \times 2^k - 1$. On the corresponding window U_M , we can fit a total of 2^k slices, each made of windows U_m stacked in all directions but one. In $p(w)$ such consecutive slices, we write w in each box U_m . This gives us a configuration $\bar{w} \in \mathcal{A}^{U_M}$ such that, for any $w \in \mathcal{A}^{U_n}$, we have

$\widehat{\delta}_w([w]) = \widehat{\mu}_m^k([w]) + O\left(\frac{n}{m}\right)$, once again with a computable domination constant that depends on d , so that:

$$|\widehat{\delta}_w([w]) - \mu([w])| = O\left(\frac{n}{m}\right) + O\left(\frac{|\mathcal{A}^{U_m}|}{2^k}\right).$$

Thus, we can actually compute integers $m(\delta, |\mathcal{A}|, d)$ and $k(\delta, |\mathcal{A}|, d)$ such that $|\widehat{\delta}_w([w]) - \mu([w])| \leq \frac{\delta}{2} \times \frac{1}{2^{|\mathcal{A}^{U_n}|}}$, which we can replace in the supremum bound for s_n . At last, we proved that there exists a pattern $w \in \mathcal{A}^{U_{\psi(\delta, |\mathcal{A}|, d)}}$ such that $\mu \in \overline{B}(\widehat{\delta}_w, \delta)$, with $\psi(\delta, |\mathcal{A}|, d) := (m(\delta, |\mathcal{A}|, d) + 1)2^{k(\delta, |\mathcal{A}|, d)} - 1$ a map that can be computed by a Turing machine. \square

In particular, the density of the family $(\widehat{\delta}_w, w \in \mathcal{A}^{U_n}, n \in \mathbb{N})$ of all periodic measures directly follows from the Covering lemma. Note how we always have $d_r^+ \leq d_1^+$ when $r \geq 1$, so the previous Covering lemma more generally applies for all these distances.

Let us conclude this subsection with a technical lemma that relates weak-* distances when projecting.

Lemma 6.2 (Projection Lemma) *Consider two measures $\lambda, \lambda' \in \mathcal{M}_{\mathcal{A}_1 \times \mathcal{A}_2}$. Then $d_{|\mathcal{A}_2|}^+(\pi_1^*(\lambda), \pi_1^*(\lambda')) \leq d^+(\lambda, \lambda')$.*

Proof We have:

$$\begin{aligned} & d_{|\mathcal{A}_2|}^+(\pi_1^*(\lambda), \pi_1^*(\lambda')) \\ &= \sum_{n \in \mathbb{N}} \frac{1}{2^n} \times \frac{1}{|\mathcal{A}_2^{U_n}|} \times \frac{1}{|\mathcal{A}_1^{U_n}|} \sum_{w_1 \in \mathcal{A}_1^{U_n}} |\pi_1^*(\lambda)([w_1]) - \pi_1^*(\lambda')([w_1])| \\ &= \sum_{n \in \mathbb{N}} \frac{1}{2^n} \times \frac{1}{|\mathcal{A}_2^{U_n}|} \times \frac{1}{|\mathcal{A}_1^{U_n}|} \sum_{w_1 \in \mathcal{A}_1^{U_n}} \left| \sum_{w_2 \in \mathcal{A}_2^{U_n}} \lambda([w_1, w_2]) - \lambda'([w_1, w_2]) \right| \\ &\leq \sum_{n \in \mathbb{N}} \frac{1}{2^n} \times \frac{1}{|(\mathcal{A}_1 \times \mathcal{A}_2)^{U_n}|} \sum_{(w_1, w_2) \in (\mathcal{A}_1 \times \mathcal{A}_2)^{U_n}} |\lambda([w_1, w_2]) - \lambda'([w_1, w_2])| \\ &= d^+(\lambda, \lambda'), \end{aligned}$$

i.e. the announced bound.

6.2 Computable Descriptions of Measure Sets

Now, from the point of view of Turing machines, the main obstruction to discuss the notion of stability is that it is not obvious how we should proceed to compute $d_B(\mu, \mathcal{M}_{\mathcal{F}})$. To do so, we will step-by-step reach a characterisation of the sets $\mathcal{M}_{\mathcal{F}}$, $\mathcal{M}_{\mathcal{F}}^B(\varepsilon)$ and at last $J(\mu, \mathcal{M}_{\mathcal{F}})$.

Lemma 6.3 (Covering Lemma for $\mathcal{M}_{\mathcal{F}}$) Consider ψ the radius-to-scale function obtained in the Covering Lemma 6.1 for $\mathcal{M}_{\mathcal{A}}$. Assume $\mathcal{F} \subset \mathcal{A}^{U_k}$ for some $k \in \mathbb{N}$. Then:

$$\mathcal{M}_{\mathcal{F}} \subset \bigcup_{w \in \mathcal{W}_{\mathcal{F}}(\rho)} \overline{\mathbb{B}}(\widehat{\delta}_w, \rho)$$

where $\mathcal{W}_{\mathcal{F}}(\rho) \subset \mathcal{A}^{U_{\psi(\rho, |\mathcal{A}|, d)}}$ is the set of patterns on the window $U_{\psi(\rho, |\mathcal{A}|, d)}$ that contain at most $\varphi(\rho, k, |\mathcal{A}|, d)$ forbidden patterns from \mathcal{F} , with the computable map $\varphi(\rho, k, |\mathcal{A}|, d) := \lfloor 2^k \times |\mathcal{A}^{U_k}| \times \rho \times |U_{\psi(\rho, \dots)}| \rfloor$.

Note that the set $\mathcal{W}_{\mathcal{F}}(\rho)$ depends on $k, |\mathcal{A}|$ and d , but we hide them from the notation as they are either directly “written” in the computer representation of \mathcal{F} or can be deduced from it.

Proof The first covering lemma gives us $\mathcal{M}_{\mathcal{A}} = \bigcup_{w \in \mathcal{A}^{U_{\psi(\rho, |\mathcal{A}|, d)}}} \overline{\mathbb{B}}(\widehat{\delta}_w, \rho)$. If we prove that $\mathcal{M}_{\mathcal{F}} \cap \overline{\mathbb{B}}(\widehat{\delta}_w, \rho) = \emptyset$ whenever w has more than $\varphi(\rho, \dots)$ forbidden patterns, this will conclude the proof.

Consider such a pattern $w \notin \mathcal{W}_{\mathcal{F}}(\rho)$ and $\mu \in \mathcal{M}_{\mathcal{F}}$. For any forbidden pattern $u \in \mathcal{F} \subset \mathcal{A}^{U_k}$, we have $\mu([u]) = 0$, hence:

$$d^+(\mu, \widehat{\delta}_w) \geq \frac{1}{2^k} \times \frac{1}{|\mathcal{A}^{U_k}|} \sum_{u \in \mathcal{F}} \widehat{\delta}_w([u]).$$

Now, by summing the number of occurrences of each forbidden pattern in w , we obtain at least $\varphi(\rho, \dots) + 1 > 2^k |\mathcal{A}^{U_k}| \times \rho \times |U_{\psi(\rho, \dots)}|$. The last factor is precisely the normalisation constant used to define $\widehat{\delta}_w$, so that $d^+(\mu, \widehat{\delta}_w) > \rho$, which concludes the proof. \square

Corollary 6.4 Using the previous notations, we have:

$$\mathcal{M}_{\mathcal{F}} = \bigcap_{\rho > 0} \bigcup_{w \in \mathcal{W}_{\mathcal{F}}(\rho)} \overline{\mathbb{B}}(\widehat{\delta}_w, \rho).$$

Proof The covering lemma for $\mathcal{M}_{\mathcal{F}}$ holds for any $\rho > 0$, hence by taking the intersection we directly obtain the inclusion (\subset) .

Conversely, consider $\mu \in \bigcap \bigcup(\dots)$. There exists a sequence of radii $\rho_n \rightarrow 0$ and corresponding patterns w_n such that $d^+(\mu, \widehat{\delta}_{w_n}) \leq \rho_n$. For any given forbidden pattern $u \in \mathcal{F}$, we have:

$$\widehat{\delta}_{w_n}([u]) \leq \frac{\varphi(\rho_n, \dots) + O(\psi(\rho_n, \dots)^{d-1})}{|U_{\psi(\rho_n, \dots)}|} = O(\rho_n) + O\left(\frac{1}{\psi(\rho_n, \dots)}\right) \rightarrow 0,$$

with the first bound relating to the occurrences of u within w_n and the second to the occurrences on the interface between two square blocks of $w_n^{\mathbb{Z}^d}$. Thus, for the weak-* limit μ we have $\mu([u]) = 0$, so $\mu \in \mathcal{M}_{\mathcal{F}}$ [1, Lemma 2.11]. \square

On one hand, the covering lemma tells us that any measure $\mu \in \mathcal{M}_{\mathcal{F}}$ can be approximated by some periodic measures with an explicit bound on the speed of convergence. More precisely, all measures $\mu \in \mathcal{M}_{\mathcal{F}}$ are ρ -close to some $\widehat{\delta}_w$ with a pattern $w \in \mathcal{W}_{\mathcal{F}}(\rho)$, but not *all* such $\widehat{\delta}_w$ are necessarily ρ -close to $\mathcal{M}_{\mathcal{F}}$. This is roughly correlated to the nuance between locally admissible and globally admissible tilings. On the other hand, the corollary tells us that, while we have no computable bound on the speed of convergence, we still necessarily converge to measures $\mathcal{M}_{\mathcal{F}}$, *i.e.* the ρ -coverings converge to the set $\mathcal{M}_{\mathcal{F}}$ in the corresponding Hausdorff topology as $\rho \rightarrow 0$. This will allow us to computationally describe $\mathcal{M}_{\mathcal{F}}$ as the set of all adherence values of these measures.

We now want to move onto the noisy framework, to obtain a similar result for $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$, as a subset of $\mathcal{M}_{\widetilde{\mathcal{F}}}$. In the following proposition, we denote s_n the partial sum for $d_{|\mathcal{A}|}^+$ on the space of noises $\Omega_{\{0,1\}}$. In particular, if we use the computable rank $n(\rho)$ introduced in the proof of the covering lemma for $\mathcal{M}_{\mathcal{A}}$, then we can guarantee $s_{n(\rho)} \leq d_{|\mathcal{A}|}^+ \leq s_{n(\rho)} + \rho$.

Lemma 6.5 (Covering Lemma for $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$) *As in the Covering Lemma 6.3 for $\mathcal{M}_{\mathcal{F}}$, we assume here that $\mathcal{F} \subset \mathcal{A}^{U_k}$. Note that $\mathcal{M}_{\widetilde{\mathcal{F}}}$ uses the alphabet $\mathcal{A} \times \{0, 1\}$, of cardinality $2|\mathcal{A}|$. For any $\varepsilon \in [0, 1]$, we can refine the covering of $\mathcal{M}_{\widetilde{\mathcal{F}}}$ into a covering of $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$:*

$$\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon) = \bigcap_{\rho > 0} \bigcup_{(w,b) \in \widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\rho)} \overline{\mathcal{B}}(\widehat{\delta}_{(w,b)}, \rho),$$

where $\widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\rho) \subset \mathcal{W}_{\widetilde{\mathcal{F}}}(\rho)$ is the subset for which $s_{n(\rho)}(\widehat{\delta}_b, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) \leq \rho$ holds.

Proof As before, we need to prove both inclusions.

Consider first (C). As $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon) \subset \mathcal{M}_{\widetilde{\mathcal{F}}}$, we already have the inclusion for any $\rho > 0$ if we forget about the new condition on $s_{n(\rho)}$, using the Covering Lemma 6.3 for $\mathcal{M}_{\widetilde{\mathcal{F}}}$. Hence, it suffices to prove that if $(w, b) \in \mathcal{W}_{\widetilde{\mathcal{F}}}(\rho)$ does not satisfy the new condition, then $d^+(\widehat{\delta}_{(w,b)}, \widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)) > \rho$. This directly follows from the fact that for any $\lambda \in \widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$, using the Projection Lemma 6.2 we have:

$$d^+(\widehat{\delta}_{(w,b)}, \lambda) \geq d_{|\mathcal{A}|}^+(\widehat{\delta}_b, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) \geq s_{n(\rho)}(\widehat{\delta}_b, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) > \rho.$$

Conversely, consider (\supset) . For any $\lambda \in \bigcap \bigcup (\dots) \subset \mathcal{M}_{\mathcal{F}}$, we have a sequence $\rho_n \rightarrow 0$ and patterns $(w_n, b_n) \in \widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\rho_n)$ such that $d^+(\widehat{\delta_{(w_n, b_n)}}, \lambda) \leq \rho_n$. We just need to prove that $\pi_2^*(\lambda) = \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}$. This comes from the fact that:

$$\begin{aligned} d_{|\mathcal{A}|}^+(\pi_2^*(\lambda), \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) &\leq d_{|\mathcal{A}|}^+(\pi_2^*(\lambda), \widehat{\delta_{b_k}}) + d_{|\mathcal{A}|}^+(\widehat{\delta_{b_k}}, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) \\ &\leq d^+(\lambda, \widehat{\delta_{(w_k, b_k)}}) + d_{|\mathcal{A}|}^+(\widehat{\delta_{b_k}}, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) \\ &\leq d^+(\lambda, \widehat{\delta_{(w_k, b_k)}}) + s_{n(\rho_k)}(\widehat{\delta_{b_k}}, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) + \rho_k \\ &\leq 3\rho_k. \end{aligned}$$

As $k \rightarrow \infty$, we have $d_{|\mathcal{A}|}^+(\pi_2^*(\lambda), \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d}) = 0$, which concludes the proof. \square

As long as $\varepsilon, \rho \in \mathbb{Q}_{>0}$, then $(b, \varepsilon, \rho) \mapsto s_{n(\rho)}(\widehat{\delta_b}, \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^d})$ is computable. Hence, as for the case of $\mathcal{M}_{\mathcal{F}}$, this covering lemma tells us *both* that we *can* mathematically approximate some $\lambda \in \widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ with an explicit bound ρ , and that we have a way of describing $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ as a set of adherence values of a computable sequence but without an explicit control on the speed of convergence.

Remark 6.1 (Approximating $J(\mu, \mathcal{M}_{\mathcal{F}})$) Later on, in order to approximate a joining $\lambda \in J(\mu, \mathcal{M}_{\mathcal{F}})$, we will consider a periodic measure $\widehat{\delta_{(w_1, w_2)}}$ on $\Omega_{\mathcal{A} \times \mathcal{A}}$ such that $\widehat{\delta_{w_1}}$ is close to $\widehat{\delta_w}$, with $\widehat{\delta_{(w, b)}}$ on $\Omega_{\mathcal{A} \times \{0,1\}}$ an approximation of a measure that projects to $\mu \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ obtained through the Covering Lemma 6.5, and that $\widehat{\delta_{w_2}}$ is close to $\widehat{\delta_{w'}}$ obtained through the Covering Lemma 6.3 for $\mathcal{M}_{\mathcal{F}}$.

6.3 Equivalent Characterisations of Stability

In the current context, if $\Omega_{\mathcal{F}} = \emptyset$, then there is some window U_k that doesn't admit locally admissible tilings. Consequently, as long as $\varepsilon < 1$, by Borel-Cantelli lemma, with probability 1, a measure $\lambda \in \widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ should have infinitely many clear windows the size of U_k , which cannot be tiled, thus $\widetilde{\mathcal{M}}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon) = \emptyset$. Likewise, it follows that under some rank ρ_0 , the computable sets $\widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\rho)$ must also be empty. The idea to take out of this paragraph is that the non-emptiness of $\Omega_{\mathcal{F}}$ is directly implied by all incoming characterisations of stability.

Let the measurable event $\Delta := \bigcup_{a \neq b \in \mathcal{A}} [a, b] \subset \Omega_{\mathcal{A}^2}$. Because we consider σ -invariant measures, by an ergodic theorem, $d_B(\mu, \nu) = \inf_{\lambda \in J(\mu, \nu)} \lambda(\Delta)$. In particular, as $\lambda \mapsto \lambda(\Delta)$ depends only on a finite window in $\Omega_{\mathcal{A}^2}$, it is a continuous map for the weak-* topology.

Let us remind what it means for \mathcal{F} to induce a stable SFT:

$$\forall \delta > 0, \exists \varepsilon > 0, \forall \mu \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon), d_B(\mu, \mathcal{M}_{\mathcal{F}}) \leq \delta.$$

Notice how, by monotonicity of the definition, we can restrict this formula by quantifying ε and δ over the countable set $\mathbb{Q}_{>0}$ instead. What’s more, using the previous rewriting of d_B through joinings, we obtain the following characterisation:

$$\forall \delta \in \mathbb{Q}_{>0}, \exists \varepsilon \in \mathbb{Q}_{>0}, \forall \mu \in \mathcal{M}_{\mathcal{F}}^B(\varepsilon), \exists \lambda \in J(\mu, \mathcal{M}_{\mathcal{F}}), \lambda(\Delta) \leq \delta. \tag{1}$$

Note that, by embedding all the patterns of \mathcal{F} in a big-enough square box U_k and enumerating all the $w \in \mathcal{A}^{U_k}$ that contain at least one forbidden pattern, we can trivially compute \mathcal{F}' such that $\Omega_{\mathcal{F}} = \Omega_{\mathcal{F}'}$. As stability with Bernoulli noise is a conjugacy invariant [1, Corollary 3.16], we can equivalently decide whether \mathcal{F}' is stable instead. Hence, we will without loss of generality assume that $\mathcal{F} \subset \mathcal{A}^{U_k}$ in the following theorem, so that the covering lemmas may apply.

Proposition 6.6 *The SFT $\Omega_{\mathcal{F}}$ is stable iff it satisfies the following formula:*

$$\begin{aligned} &\forall \delta \in \mathbb{Q}_{>0}, \exists \varepsilon \in \mathbb{Q}_{>0}, \forall \rho \in \mathbb{Q}_{>0}, \\ &\forall \mu \in \mathcal{M}_{\mathcal{F}}^B(\varepsilon), \exists (w_1, w_2) \in (\mathcal{A}^2)^{U_{\psi(\rho, |\mathcal{A}^2|, d)}}, \tag{2} \\ &\left[d_{|\mathcal{A}|}^+(\widehat{\delta_{w_1}}, \mu) \leq \rho \right] \wedge \left[d_{|\mathcal{A}|}^+(\widehat{\delta_{w_2}}, \mathcal{M}_{\mathcal{F}}) \leq \rho \right] \wedge \left[\widehat{\delta_{(w_1, w_2)}}(\Delta) \leq \delta + |\mathcal{A}|^2 \rho \right]. \end{aligned}$$

Proof Consider first the implication (1)⇒(2). Assuming (1) is satisfied, let us fix δ , ε and μ such that there exists a joining $\lambda \in J(\mu, \mathcal{M}_{\mathcal{F}})$ for which $\lambda(\Delta) \leq \delta$. Using the Covering Lemma 6.1 for $\mathcal{M}_{\mathcal{A}^2}$, we know that for any rational $\rho \in \mathbb{Q}_{>0}$, we have some couple $(w_1, w_2) \in (\mathcal{A}^2)^{U_{\psi(\rho, |\mathcal{A}^2|, d)}}$ such that $d^+(\widehat{\delta_{(w_1, w_2)}}, \lambda) \leq \rho$. The first two inequalities in (2) follow directly from the Projection Lemma 6.2. For the third one:

$$\begin{aligned} \widehat{\delta_{(w_1, w_2)}}(\Delta) &\leq \lambda(\Delta) + \sum_{a, b \in \mathcal{A}} \left| \lambda([(a, b)]) - \widehat{\delta_{(w_1, w_2)}}([(a, b)]) \right|, \\ &\leq \lambda(\Delta) + |\mathcal{A}|^2 d^+ \left(\lambda, \widehat{\delta_{(w_1, w_2)}} \right), \\ &\leq \delta + |\mathcal{A}|^2 \rho. \end{aligned}$$

Remark that the consecutive universal blocks $\forall \mu \in \mathcal{M}_{\mathcal{F}}^B(\varepsilon)$ and $\forall \rho \in \mathbb{Q}_{>0}$ do not depend on each other, so we can freely reorder them as in (2), which concludes this implication.

Conversely, suppose (2) holds and let us prove (2)⇒(1). Likewise, fix δ , ε and μ for which the rest of the formula (i.e. $\forall \rho \dots$) is satisfied. Consider a sequence $\rho_n \rightarrow 0$ and the consequent patterns $\widehat{\delta_{w_1^n, w_2^n}} \in (\mathcal{A}^2)^{U_{\psi(\rho_n, |\mathcal{A}^2|, d)}}$. Up to extraction of a subsequence, we can assume that the sequence weakly converges to a measure $\lambda \in \mathcal{M}_{\mathcal{A}^2}$. The first inequality of (2) gives us $\pi_1^*(\lambda) = \mu$ at the limit. The second inequality gives us $\pi_2^*(\lambda) \in \mathcal{M}_{\mathcal{F}}$ (as $\mathcal{M}_{\mathcal{F}}$ is closed), so $\lambda \in J(\mu, \mathcal{M}_{\mathcal{F}})$. At last, the third inequality becomes $\lambda(\Delta) \leq \delta$ at the limit by continuity of $\lambda \mapsto \lambda(\Delta)$, which concludes the proof. □

We now want to use the Covering Lemma 6.5 for $\widetilde{\mathcal{M}}_{\mathcal{F}}^B(\varepsilon)$ to replace the $\forall \mu$ block by a universal block that quantifies over rational numbers instead.

Proposition 6.7 *The SFT $\Omega_{\mathcal{F}}$ is stable iff it satisfies the following formula:*

$$\begin{aligned} \forall \delta \in \mathbb{Q}_{>0}, \exists \varepsilon \in \mathbb{Q}_{>0}, \forall \rho \in \mathbb{Q}_{>0}, \exists \gamma \in \mathbb{Q}_{>0}, \gamma \leq \rho, \\ \forall (w, b) \in \widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\gamma), \exists w_0 \in \mathcal{W}_{\mathcal{F}}(\rho), \exists (w_1, w_2) \in (\mathcal{A}^2)^{U_{\psi}(\rho, |\mathcal{A}^2|, d)}, \\ \left[d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_1}, \widehat{\delta}_w) \leq 2\rho \right] \wedge \left[d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_2}, \widehat{\delta}_{w_0}) \leq 2\rho \right] \\ \wedge \left[\widehat{\delta}_{(w_1, w_2)}(\Delta) \leq \delta + |\mathcal{A}|^2 \rho \right]. \end{aligned} \tag{3}$$

Proof Let us prove that (2) \Rightarrow (3). Assume (2) holds true and fix δ, ε and ρ so that the rest of the formula holds true. The Covering Lemma 6.5 for $\mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ implies that $\max_{(w,b) \in \widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\gamma)} d^+(\widehat{\delta}_{(w,b)}, \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)) \xrightarrow{\gamma \rightarrow 0} 0$. In particular, there is a rational $\gamma \leq \rho$ such that $\max d^+(\dots) \leq \rho$. This will allow us to merge the universal block that should replace $\forall \mu$ directly into the already existing $\forall \rho \in \mathbb{Q}_{>0}$.

Now, for such a choice of γ , and any $(w, b) \in \widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\gamma)$, there always exists some $\mu \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$ such that $d^+(\widehat{\delta}_{(w,b)}, \mu) \leq \rho$. As $\pi_1^*(\mu) \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$, (2) applies to it, and we can chose a corresponding pair $(w_1, w_2) \in (\mathcal{A}^2)^{U_{\psi}(\rho, |\mathcal{A}^2|, d)}$. Hence:

$$\begin{aligned} d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_1}, \widehat{\delta}_w) &\leq d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_1}, \pi_1^*(\mu)) + d_{|\mathcal{A}|}^+(\pi_1^*(\mu), \widehat{\delta}_w) \\ &\leq d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_1}, \pi_1^*(\mu)) + d^+(\mu, \widehat{\delta}_{(w,b)}) \\ &\leq 2\rho. \end{aligned}$$

Likewise, we have $v \in \mathcal{M}_{\mathcal{F}}$ such that $d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_2}, v) \leq \rho$ in (2), thus by the Covering Lemma 6.3 for $\mathcal{M}_{\mathcal{F}}$ we have $w_0 \in \mathcal{W}_{\mathcal{F}}(\rho)$ such that $d_{|\mathcal{A}|}^+(v, \widehat{\delta}_{w_0}) \leq \rho$, hence $d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_2}, \widehat{\delta}_{w_0}) \leq 2\rho$. The third inequality does not change, which concludes the implication (2) \Rightarrow (3).

Now, suppose (3) is true and let us prove (3) \Rightarrow (1). Fix δ, ε in the formula. Consider any sequence $\rho_n \rightarrow 0$, and the corresponding γ_n in the formula.

Let $\mu \in \mathcal{M}_{\mathcal{F}}^{\mathcal{B}}(\varepsilon)$. Using the Covering Lemma 6.5, we know there exists a sequence $(w^n, b^n) \in \widetilde{\mathcal{W}}_{\mathcal{F}}^{\varepsilon}(\gamma_n)$ such that $d_{|\mathcal{A}|}^+(\mu, \widehat{\delta}_{w^n}) \leq \gamma_n \leq \rho_n \rightarrow 0$. At any rank, we may chose $w_0^n \in \mathcal{W}_{\mathcal{F}}(\rho_n)$ and $(w_1^n, w_2^n) \in (\mathcal{A}^2)^{U_{\psi}(\rho_n, |\mathcal{A}^2|, d)}$ accordingly in (3). Up to extraction, $\widehat{\delta}_{w_0^n}$ converges to $v \in \mathcal{M}_{\mathcal{F}}$ and $\widehat{\delta}_{(w_1^n, w_2^n)}$ to $\lambda \in \mathcal{M}_{\mathcal{A}^2}$.

At the limit $\rho_n \rightarrow 0$, the first inequality of (3) tells us that $\pi_1^*(\lambda) = \mu$. Likewise, the second one tells us that $\pi_2^*(\lambda) = v \in \mathcal{M}_{\mathcal{F}}$, thence $\lambda \in J(\mu, \mathcal{M}_{\mathcal{F}})$. The third inequality naturally becomes $\lambda(\Delta) \leq \delta$, hence (1) holds true, the SFT is stable. \square

Theorem 6.8 (Upper Bound for Stability) The problem P_{stab} is in Π_4 .

Proof We just proved that $\mathcal{F} \in P_{stab}$ iff it satisfies (3). For the two blocks $d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_i}, \widehat{\delta}_w) \leq 2\rho$ in (3), we can replace 2ρ by 3ρ to have a strict inequality instead. In particular, the proof of (3) \Rightarrow (1) applies to this variant, so it is indeed an equivalent characterisation of stability. The interest of this variant point of view is

that, as $d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_i}, \widehat{\delta}_w)$ is a computable real number, $d_{|\mathcal{A}|}^+(\widehat{\delta}_{w_i}, \widehat{\delta}_w) < 3\rho$ becomes a *semi-decidable* problem, adding a countable existential quantifier that can be merged into the $\exists\gamma$ block.

This variant version starts with $[\forall\delta \in \mathbb{Q}_{>0}, \exists\varepsilon \in \mathbb{Q}_{>0}, \forall\rho \in \mathbb{Q}_{>0}, \exists\gamma \in \mathbb{Q}_{>0}]$, *i.e.* four alternating layers of countable quantifiers. The following quantifiers are over finite computable sets, and then the three inequalities are decidable. Hence, this whole block can be decided in finite time. \square

7 Perspectives

In this article, we studied the decision problem of whether an SFT is stable, given a tiling set (*i.e.* the set \mathcal{F} of forbidden patterns on a finite alphabet) as the input. First, we described how to obtain stable and unstable structures related to the Robinson tiling. Then, by embedding Turing machines in a simulating structure based on the Robinson tiling, we obtained a reduction from the problem of totality of the support of a Turing machine (which is a known Π_2 -complete problem) to the question of stability. We couldn't improve the lower bound by simulating more complex questions in the same structure, as we then reach obstructions regarding how much communication happens in the tiling between macro-tiles (we would need a lot of communication for the stable case, and simultaneously much less of it for the color flip argument to still hold and give the unstable case). Secondly, by using computable analysis arguments, we proved that this property is at most Π_4 -computable in the arithmetic hierarchy. In this case, our best guess for a tighter bound may be to add restrictions on the structure of $\mathcal{M}_{\mathcal{F}}$, such as unique ergodicity (which implies $\mathcal{M}_{\mathcal{F}}$ is computable), but our attempts so far don't reduce the complexity of the set of ε -noisy measures, which still imposes an overall Π_4 upper bound. Thus, the precise position of this undecidable problem in the arithmetical hierarchy, between Π_2 and Π_4 , is still unknown, and may require a new approach to be found.

Some other questions related to stability remain open, notably whether the informal heuristic for the convergence speed of aperiodic tilings at the end of Section 4 actually holds. A related approach would perhaps be to study stability for other classical SFTs, such as the aperiodic Kari-Culik tiling [5]. This aperiodic tiling is notable for not having the kind of hierarchical self-similar structures we used in this article, and its aperiodicity comes instead from the simulation of a transformation of the interval. Thus, Besicovitch-stability for this tiling may be interpreted as a kind of robustness to perturbations of the corresponding transformation.

References

1. Gayral, L., Sablik, M.: On the Besicovitch-stability of noisy random tilings. *Electron. J. Probab.* **28**, 1–38 (2023). <https://doi.org/10.1214/23-EJP917>
2. Gayral, L.: The Besicovitch-stability of noisy tilings is undecidable, Marseille (2021). *Automata*. <http://www.hal.science/hal-03233596>
3. Berger, R.: The undecidability of the domino problem (66) (1966). <https://doi.org/10.1090/memo/0066>

4. Robinson, R.: Undecidability and nonperiodicity for tilings of the plane. *Invent. Math.* **12**, 177–209 (1971). <https://doi.org/10.1007/BF01418780>
5. Kari, J.: A small aperiodic set of Wang tiles. *Discrete. Math.* **160**, 259–264 (1996). [https://doi.org/10.1016/0012-365X\(95\)00120-L](https://doi.org/10.1016/0012-365X(95)00120-L)
6. Jeandel, E., Vanier, P.: The undecidability of the domino problem. In: *Substitution and Tiling Dynamics: Introduction to Self-inducing Structures. Lecture Notes in Mathematics*, vol. 2273, pp. 293–357. Springer, Cham (2020). <https://doi.org/10.1007/978-3-030-57666-06>
7. Hochman, M., Meyerovitch, T.: A characterization of the entropies of multidimensional shifts of finite type. *Ann. Math.* **171**(3), 2011–2038 (2010). <https://doi.org/10.4007/annals.2010.171.2011>
8. Meyerovitch, T.: Growth-type invariants for \mathbb{Z}^d subshifts of finite type and arithmetical classes of real numbers. *Invent. Math.* **184**, 567–589 (2011). <https://doi.org/10.1007/s00222-010-0296-1>
9. Aubrun, N., Sablik, M.: Simulation of effective subshifts by two-dimensional subshifts of finite type. *Acta. Appl. Math.* **126**, 35–63 (2013). <https://doi.org/10.1007/s10440-013-9808-5>
10. Durand, B., Romashchenko, A., Shen, A.: Fixed-point tile sets and their applications. *J. Comput. Syst. Sci.* **78**(3), 731–764 (2012). <https://doi.org/10.1016/j.jcss.2011.11.001>
11. Hochman, M.: On the dynamics and recursive properties of multidimensional symbolic systems. *Invent. math.* **176**(1), 131–167 (2009). <https://doi.org/10.1007/s00222-008-0161-7>
12. Jeandel, E., Vanier, P.: Characterizations of periods of multi-dimensional shifts. *Ergod. Theory. Dyn. Syst.* **35**(2), 431–460 (2015). <https://doi.org/10.1017/etds.2013.60>
13. Westrick, L.: Seas of squares with sizes from a Π_1^0 set. *Isr. J. Math.* **222**(1), 431–462 (2017). <https://doi.org/10.1007/s11856-017-1596-6>
14. Jeandel, E., Vanier, P.: Hardness of conjugacy, embedding and factorization of multidimensional subshifts. *J. Comput. Syst. Sci.* **81**(8), 1648–1664 (2015). <https://doi.org/10.1016/j.jcss.2015.05.003>
15. Westrick, L.: Topological completely positive entropy is no simpler in \mathbb{Z}^2 – *SFTs*(2022). <https://doi.org/10.48550/arXiv.1904.11444>
16. Callard, A., Menibus, B.: The aperiodic domino problem in higher dimension. In: *STACS*, vol. 219. Marseille, pp. 1–15 (2022). <https://doi.org/10.4230/LIPIcs.STACS.2022.19>
17. Gács, P.: Reliable cellular automata with self-organization. *J. Stat. Phys.* **103**(1–2), 45–267 (2001). <https://doi.org/10.1023/A:1004823720305>
18. Asarin, E., Collins, P.: Noisy Turing machines. In: *ICALP. Lecture Notes in Computer Science*, vol. 3580, pp. 1031–1042. Springer, Lisbon (2005). https://doi.org/10.1007/11523468_83
19. Ballier, A., Durand, B., Jeandel, E.: Structural aspects of tilings. In: *STACS*, vol. 1. Bordeaux, pp. 61–72 (2008). <https://doi.org/10.4230/LIPIcs.STACS.2008.1334>
20. Krengel, U.: *Ergodic Theorems*. De Gruyter, Berlin (1985). <https://doi.org/10.1515/9783110844641>
21. Glasner, E.: *Ergodic Theory Via Joinings* vol. **101**. AMS, Providence (2003). <https://doi.org/10.1090/surv/101>
22. Rogers, H.: *Theory of Recursive Functions and Effective Computability*. MIT press, Cambridge (1987)
23. Soare, R.: *Recursively Enumerable Sets and Degrees*. Springer, Berlin (1987)
24. Rising Entropy: The Arithmetic Hierarchy and Computability. Accessed 20 July 2022 (2020). risingentropy.com/the-arithmetic-hierarchy-and-computability
25. Jacobs, K., Keane, M.: 0–1-sequences of Toeplitz type. *Zeitschrift für Wahrscheinlichkeitstheorie und Verw. Gebiete* **13**(2), 123–131 (1969). <https://doi.org/10.1007/BF00537017>
26. Callard, A., Vanier, P.: Computational Characterization of Surface Entropies for \mathbb{Z}^2 Subshifts of Finite Type. In: *ICALP. LIPIcs*, vol. 198, pp. 122–112220. Online (2021). <https://doi.org/10.4230/LIPIcs.ICALP.2021.122>
27. Barbieri, S., Sablik, M.: A generalization of the simulation theorem for semidirect products. *Ergod. Theory. Dyn. Syst.* **39**(12), 3185–3206 (2019). <https://doi.org/10.1017/etds.2018.21>
28. Weihrauch, K.: *Computable Analysis. Texts in Theoretical Computer Science. An EATCS Series*. Springer, Berlin (2000). <https://doi.org/10.1007/978-3-642-56999-9>
29. Galatolo, S., Hoyrup, M., Rojas, C.: Dynamics and abstract computability: Computing invariant measures. *Discrete. Contin. Dyn. Syst.* **29**(1), 193–212 (2011). <https://doi.org/10.3934/dcds.2011.29.193>
30. Menibus, B., Sablik, M.: Characterization of sets of limit measures of a cellular automaton iterated on a random configuration. *Ergod. Theory. Dyn. Syst.* **38**(2), 601–650 (2018). <https://doi.org/10.1017/etds.2016.46>
31. Weihrauch, K.: Computability on measurable functions. *Computability* **6**(1), 79–104 (2017). <https://doi.org/10.3233/COM-160058>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.