

The Besicovitch-Stability of Noisy Tilings is Undecidable

Léo Gayral

► To cite this version:

Léo Gayral. The Besicovitch-Stability of Noisy Tilings is Undecidable. Automata 2021, Jul 2021, Marseille, France. hal-03233596

HAL Id: hal-03233596

<https://hal.archives-ouvertes.fr/hal-03233596>

Submitted on 25 May 2021

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

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

The Besicovitch-Stability of Noisy Tilings is Undecidable*

Gayral Léo[†]

Abstract

In this exploratory paper, I will first introduce a notion of stability, more lengthily described in a previous article [6]. In this framework, I will then exhibit an unstable aperiodic tiling. Finally, by building upon the folkloric embedding of Turing machines into Robinson tilings, we will see that the question of stability is undecidable.

1 Introduction

A subshift of finite type (usually called a *SFT*) is the dynamical version of a tiling defined by local rules. A subshift Ω is a set of \mathcal{A} -colourings of \mathbb{Z}^d invariant under any translation σ_k ($k \in \mathbb{Z}^d$). This subshift is induced by a set of patterns \mathcal{F} when it is the set of configurations that contain no translation of these patterns. When \mathcal{F} is finite, $\Omega_{\mathcal{F}}$ is a SFT.

One of the main topics of interests in the study of SFTs is how a global structure can emerge from local rules; in particular, aperiodicity of SFTs have been widely studied. First proven by Berger [4] using 20426 Wang tiles, the existence of SFTs on \mathbb{Z}^2 that force any configuration to be aperiodic has been obtained through simpler examples since, notably by Robinson [11] with 56 Wang tiles and Kari [9] with 14 Wang tiles.

One of the most interesting properties of the Robinson tiling is that its hierarchical structure leaves room for a *relatively* easy embedding of Turing machines into it [11, 10]. Berger already used tilings as static geometrical models of computation in his proof of the undecidability of the *domino problem* [4]. In the last decade, the embedding of Turing machines has been used to construct *complex* (and always aperiodic) tilings [8, 2, 5].

From this, it is natural to wonder whether tilings are stable in the presence of noise, to see if computations can survive in an imperfect setting. Such studies already exist for cellular automata [7] or Turing machines [1].

*This work was supported by the CIMI LabEx “Computability of asymptotic properties of dynamical systems” project (ANR-11-LABX-0040). The author wishes to thank their advisor M. Sablik for his careful guidance and proofreading.

[†]Université Toulouse III - Paul Sabatier, France. lgayral@math.univ-toulouse.fr
<https://perso.ens-lyon.fr/leo.gayral/en/>

A previous article [6] introduced a notion of stability of noisy SFTs using the Besicovitch distance, which quantifies the closeness between two measures with the average frequency of the differences between their generic configurations. This framework is a natural bridge from the notion of stability described by Durand, Romashchenko and Shen [5] to ergodic theory, with a more measure-theoretical viewpoint on the same general notion. The article stated in particular that stability is conjugacy-invariant, thus an intrinsic property of a global tiling, regardless of how the rules are chosen locally. A digest of this framework will be introduced in Section 2.

In this context, the article provided a fully computable criterion to decide stability in the one-dimensional case, using the fact a lot of properties of one-dimensional SFTs can be extracted from a word automaton. Then, the article proved the existence of both stable *and* unstable SFTs in *any* dimension, using independent ε -Bernoulli noises on each cell. In particular, a variant of the Robinson tiling was proven to be stable. This example is particularly interesting because the only known stable aperiodic tilings were complex constructions that can be repaired locally, which is not the case for the Robinson tiling [12, 3, 5]. However, the interface between stable and unstable examples was yet to be seen. I will here study a two-coloured variant to exhibit an unstable aperiodic example in Section 3. As of now, it remains unknown whether the canonical Robinson tiling is stable or not.

Ideally, we would seek a computable criterion on a set \mathcal{F} to decide whether $\Omega_{\mathcal{F}}$ is stable or not in any dimension, as in the one-dimensional case. As we know both stable *and* unstable variants of the Robinson tiling exist, it was natural to focus on a criterion using this structure. By combining both of these variants with the folkloric embedding of Turing machines into Robinson tilings [11], I will show in Section 4 that stability of SFTs is undecidable. This result implies that, while we may find criteria that imply (un)stability in the future, the interface between stable and unstable examples will be out of any computable reach.

2 Noisy SFTs and Besicovitch Stability

Let us concisely define noisy SFTs and stability, which were introduced more in-depth in a previous paper [6, Sections 2 and 3.1], so that we may understand the scope of the results of the further sections.

Definition 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 $w \in \mathcal{A}^{I(w)}$. A SFT is the set $\Omega_{\mathcal{F}}$ induced by \mathcal{F} as follows:

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

i.e. configurations of the SFT are such that no forbidden pattern occurs.

This set is σ -invariant, invariant under the action of any translation $\sigma_k : (\omega_l)_{l \in \mathbb{Z}^d} \mapsto (\omega_{k+l})_{l \in \mathbb{Z}^d}$. 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 (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. We can likewise define the set of forbidden patterns $\tilde{\mathcal{F}} := \left\{ (w, 0^{I(w)}) \in \tilde{\mathcal{A}}^{I(w)}, w \in \mathcal{F} \right\}$ and the corresponding SFT $\Omega_{\tilde{\mathcal{F}}}$ on the alphabet $\tilde{\mathcal{A}}$.

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 $[\varphi^*(\mu)](A) = \mu(\varphi^{-1}(A))$ for any measurable set $A \subset \Omega'$.

Definition 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:

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

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

Definition 4 (Besicovitch Distance). We define the Hamming-Besicovitch pseudo-distance d_H on $\Omega_{\mathcal{A}}$ as $d_H(x, y) = \overline{\lim}_{n \rightarrow \infty} d_n(x_n, y_n)$, with the finite Hamming distances $d_n(u, v) = \frac{1}{(2n+1)^d} \# \{k \in \llbracket -n, n \rrbracket^d, u_k \neq v_k\}$.

A coupling between two measures μ and ν on $\Omega_{\mathcal{A}}$ is a measure λ on the product space $\Omega_{\mathcal{A} \times \mathcal{A}}$ such that $\pi_1^*(\lambda) = \mu$ and $\pi_2^*(\lambda) = \nu$. The Besicovitch distance between two σ -invariant measures is then:

$$d_B(\mu, \nu) = \inf_{\lambda \text{ a coupling}} \int d_H(x, y) d\lambda(x, y).$$

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

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

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 .

3 The Red-Black Robinson Tiling is Unstable

Consider the Robinson tiles [11] in Figure 1, using the bumpy-corners variant (with diagonal interactions) instead of Wang tiles. The tileset not only uses these 6 tiles but also their rotations and symmetries, for a total of 32 tiles in the alphabet. The corresponding set of forbidden patterns is self-evident, such that we force the borders of laterally adjacent to correspond (both geometrically and with the coloured lines atop of them) and that the corner in-between any square of four tiles is filled with exactly one bumpy corner.




Figure 1: The 6 basic Robinson tiles.

This tileset induces a self-similar hierarchical structure: we define the 1-macro-tiles as the four rotations of the leftmost tile in Figure 1, with bumpy corners, 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, as in Figure 3.

A previous paper [6, Theorem 7.9] proved that an extension of this tileset, enhanced to locally enforce alignment of macro-tiles, was stable with a polynomial speed f . Note that the Robinson tiling is not *robust* in the sense of Durand, Romashchenko and Shen [5], so their stability result would not apply there. What is more, the scheme of proof used to obtain the stability of this enhanced Robinson tileset does not apply to the canonical Robinson tileset, the stability of which is still unknown.




Figure 2: The 11 basic Red-Black Robinson tiles.

Here, we will use the two-coloured variant of this Robinson tileset in Figure 2, which projects onto the previous tiling if we forget about the colour of the lines, 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. Because \mathcal{A} contains no tile with a monochromatic cross, any two squares of the same colour in the hierarchical structure do not intersect, as we can see on the 5-macro-tiles in Figure 3. As we will see in Subsection 4.1, these non-intersecting squares can be used to encode arbitrarily large space-time diagrams of Turing machines with a fixed initial empty ribbon.




Figure 3: The alternating structure of a Red-Black Robinson macro-tile.

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 a *generic* tiling $\omega \in \Omega_{RB}$, by induction, the n -macro-tiles all have a central arm with the same colour. In particular, such a generic configuration will only contain Red *or* Black 1-macro-tiles with bumpy corners, never both.

Proposition 1. *The SFT Ω_{RB} is unstable. More precisely, for any $\varepsilon > 0$, we have $\mu \in \mathcal{M}_{RB}^B(\varepsilon)$ such that $d_B(\mu, \mathcal{M}_{RB}) \geq \frac{1}{8}$.*

Proof. See Appendix A. In a few words, the idea is that we may obtain generic noisy tilings such that a half of their 1-macro-tiles is Red and the other half is Black. As these tiles with bumpy corners have a density $\frac{1}{4}$ in \mathbb{Z}^2 , these configurations are at distance d_H at least $\frac{1}{8}$ of any generic tiling of Ω_{RB} (with monochromatic bumpy-corners), thus the announced result. \square

We obtain the same instability bound by incorporating this two-coloured alternating structure into the enhanced Robinson tiles [6]. However, by keeping only *one* kind of 1-macro-tiles with bumpy-corners, whether the Red or Black ones, the previous scheme of proof cannot apply, and we could even prove stability in this specific Red-Black enhanced case just as we did for the enhanced case. This phenomenon will be useful in the next section

to encode Turing machines in these non-intersecting squares in a *stable* way (with two colours Red and Black but only one kind of 1-macro-tiles) but such that an *unstable* structure may appear if the machine terminates (by using a second set of colour Blue and Green, without the same constraint on 1-macro-tiles).

4 Undecidability of the Stability

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

By carefully encoding Turing machines into the tiling, we will equate the emergence of these two widely different macro-tiles to the appearance of a final state in the space-time diagram of the machine, thus equating stability to the halting problem:

Theorem 1. *Let us denote by Ω_T the SFT that embeds the Turing machine T into a variant of the Robinson tiling, as explained in Subsection 4.1.*

Then Ω_T is stable (for d_B on the class \mathcal{B}) if and only if the Turing machine T does not end on the empty input. In the stable case, Ω_T is polynomially stable.

Proof. See Subsection 4.2 for the proof of the stable case and Subsection 4.3 for the proof of the unstable case. \square

Corollary 1. *The problem of deciding whether the SFT $\Omega_{\mathcal{F}}$ is stable or not given the set of forbidden patterns \mathcal{F} is undecidable.*

4.1 Folkloric Embedding of Turing Machines in Tilings

We will use a product alphabet $\mathcal{A}_T \subset \mathcal{A}_R \times \mathcal{A}_T \times \mathcal{A}_F$. The layer \mathcal{A}_R corresponds to the Robinson structure, and will be given explicitly. The layer \mathcal{A}_T will correspond to the space-time diagrams of the Turing machine T , and its construction will not be detailed. The *flow* layer \mathcal{A}_F will permit the two other layers to communicate, so that we may force a transition from the stable phase (Red-Black) to the unstable phase (Blue-Green) in \mathcal{A}_R when the Turing machine terminates in \mathcal{A}_T .

The Robinson alphabet \mathcal{A}_R is explicitly described by Figure 4. The dotted and dashed grey lines on some tiles come from the enhanced Robinson tiling [6], and serve to enforce local alignment of macro-tiles; we only showed about half of the cases for brevity's sake, inverting dashed and dotted lines also gives a valid tile, *except* on the three first columns where they must stay

as-is to represent the orientation of the central crosses of macro-tiles. The actual alphabet \mathcal{A}_R is then made of all the rotations of the three leftmost columns, as well as the rotations and symmetries of the other ones, which brings us to a total of 172 tiles.




Figure 4: Main tiles of the alphabet \mathcal{A}_R .

Here, the 1-macro-tiles *must* be Black and the 2-macro-tiles Red, so that instability does not emerge at start. Then, \mathcal{A}_R produces increasingly big Red and Black squares in an alternating fashion. For a reason explained later on, a “phase transition” from Red-Black to Blue-Green squares may occur right after a certain scale of Red-armed $2n$ -macro-tiles, using the two bottom-right tiles on Figure 4. Because $(2n + 1)$ -macro-tiles may then have a Blue or Green arm without restrictions, the proof of instability from Proposition 1 may adapt here.

The Turing alphabet \mathcal{A}_T can be algorithmically derived from a Turing machine T . The intricacies of how it is built and coupled with \mathcal{A}_R were already introduced in Robinson’s article [11], and are well explained in Kari’s lecture notes [10, Subsection 5.5]. The main idea of the construction is that each $(2n + 1)$ -macro-tile contains a big Red square (of length $4^n + 1$ tiles) which will encode a space-time diagram initialised on the semi-infinite empty ribbon. Because Turing computations already occur in the smaller Red squares, they can only occur on “free” rows and columns of a Red square, that do not intersect any *smaller* Red square. Thus, a $(2n + 1)$ -macro-tile encodes a space-time diagram of length $2^n + 1$ tiles (both its space horizon on the right and its time horizon on upwards). In the folkloric construction, the tiling “freezes” when the diagram reaches a final state (because the goal was to study decidability of whether $\Omega_{\mathcal{F}}$ is empty); here, we twist this construction so that instead of freezing the tiling when the machine halts, we just erase the state of the machine and then keep copying the ribbon until we reach the upper end of the Red square.

The flow layer \mathcal{A}_F shown in Figure 5 permits to force (and restrict) the transition from the Red-Black phase to the Blue-Green phase in \mathcal{A}_R . The rules in this layer are self-evident: non-empty neighbouring tiles must have the same colour, and arrows must correspond on the borders of the tiles. The *source* tiles can only occur on the cells that contain a final state in

the Turing layer \mathcal{A}_T (thus *at most* once in any space-time diagram, and in a *free* row and column). The *flow* tiles propagate this information to the edges of the Red square, and then the *border* tiles propagate it along the edges. When it reaches what should be a tile with a Red-Black cross (for both Black arms of the central cross), it forces the transition to a Blue or Green central arm outside the Red square with the *escape* tiles.




Figure 5: Source, flow, border, escape and empty tiles (left to right).

4.2 Proof of the Stable Case

Assume here that T is a Turing machine that does not halt on the empty input. In this case, a source will never appear on the flow layer \mathcal{A}_F , which will be overall empty in any tiling, so the macro-tiles will never transition from the Red-Black phase to the Blue-Green phase. To prove stability here, we follow the tracks of the proof of stability of the enhanced Robinson tiling [6, Subsections 7.3 and 7.4]. Without colours, \mathcal{A}_R projects to the enhanced Robinson tileset, so the following result still holds here:

Proposition 2 ([6, Proposition 7.7]). *Let us denote $B_k = \llbracket -k, k \rrbracket^2$. For any scale of macro-tiles $n \geq 2$, the constant $C_n = 2^n - 1$ is such that, for any $k \geq 0$ and any clear locally admissible pattern ω on B_{k+C_n} , its restriction $\omega|_{B_k}$ is almost globally admissible, in the sense that up to a low-density set, $\omega|_{B_k}$ is the restriction of a tiling from Ω_T , with well-aligned and well-oriented n -macro-tiles.*

In the original result, the low-density set was simply the grid surrounding the n -macro-tiles, of density bounded by $\frac{1}{2^n-1}$ in \mathbb{Z}^2 . That much is still true *on the layer \mathcal{A}_R* , but not necessarily on the alphabet \mathcal{A}_T as a whole, and we need to proceed a bit more carefully.

The important thing to keep in mind is that, for the layers \mathcal{A}_T and \mathcal{A}_F , a Red square on \mathcal{A}_R acts as an insulator and entirely determines the tiles inside it on these layers, by encoding a given space-time diagram. Hence, we just need to prove that inside a macro-tile, (filled) red squares have a high-density:

Proposition 3. *If we look solely at n -macro-tiles, the number of tiles outside a Red square is a $O\left((2\sqrt{3})^n\right)$.*

Proof. See Appendix B. The general idea of the proof is to remark that Red squares form a kind of Sierpiński carpet inside macro-tiles. \square

As the overall surface area of a n -macro-tile is a $\Theta(4^n)$, whereas the tiles outside a red square occupy $O(\sqrt{12}^n)$, we finally conclude that the *so-called* low-density set has a density $O((\sqrt{3}/2)^n)$.

From there, by using this new bound, we can adapt the rest of the proof [6, Proposition 7.8 and Theorem 7.9] to obtain once again polynomial stability, though at a slower speed than with the $O(1/2^n)$ bound on the low-density set.

4.3 Proof of the Unstable Case

Assume now that the machine T halts on the empty input, and consider the smaller integer n such that the machine halts before time $2^n + 1$. Then, by construction, a final state of T first appears in $(2n + 1)$ -macro-tiles. When it appears, the *flow* layer \mathcal{A}_F sends a signal to the Red square, as shown on Figure 6. This signal then forces a Black-to-Green transition (or a Black-to-Blue one) on the central arm.




Figure 6: Structure of a 5-macro-tile with a non-empty flow layer.

These Blue and Green $(2n + 1)$ -macro-tiles will play the very same role Red and Black 1-macro-tiles did in the proof of Proposition 1.

Without going into much more details, let us underline that instead of looking at the colour of a Red or Black 1-macro-tile inside a 2×2 -square, which has density $\frac{1}{4}$, we look here at the colour of a Blue or Green tile (at least one) among 4^{2n+1} , hence a lower bound $d_B(\mu, \mathcal{M}_T) \geq \frac{1}{8 \times 16^n}$ for any amount of noise $\varepsilon > 0$.

References

- [1] Eugene Asarin and Pieter Collins. Noisy Turing machines. In *ICALP 2005: Automata, Languages and Programming, 32nd International Colloquium*, volume 3580 of *Lecture Notes in Computer Science*, pages 1031–1042. Springer, 2005.
- [2] Nathalie Aubrun and Mathieu Sablik. Simulation of effective subshifts by two-dimensional subshifts of finite type. *Acta Applicandae Mathematicae. An International Survey Journal on Applying Mathematics and Mathematical Applications*, 126:35–63, 2013.
- [3] Alexis Ballier, Bruno Durand, and Emmanuel Jeandel. Tilings robust to errors. In *LATIN 2010: Theoretical Informatics*, pages 480–491. Springer Berlin Heidelberg, 2010.
- [4] Robert Berger. *The Undecidability of the Domino Problem*. Number 66 in Memoirs of the American Mathematical Society. AMS, 1966.
- [5] Bruno Durand, Andrei Romashchenko, and Alexander Shen. Fixed-point tile sets and their applications. *Journal of Computer and System Sciences*, 78(3):731–764, 2012.
- [6] Léo Gayral and Mathieu Sablik. On the Besicovitch-stability of noisy random tilings. <https://arxiv.org/abs/2104.09885>, 2021.
- [7] Peter Gács. Reliable cellular automata with self-organization. *Journal of Statistical Physics*, 103(1-2):45–267, 2001.
- [8] Michael Hochman. On the dynamics and recursive properties of multidimensional symbolic systems. *Inventiones Mathematicae*, 176(1):131–167, 2009.
- [9] Jarkko Kari. A small aperiodic set of Wang tiles. *Discrete Mathematics*, 160:259–264, 1996.
- [10] Jarkko Kari. Tilings and patterns. <https://users.utu.fi/jkari/tilings2019/part2.pdf>, 2019.
- [11] Raphael Robinson. Undecidability and nonperiodicity for tilings of the plane. *Inventiones mathematicae*, 12:177–209, 1971.
- [12] Siamak Taati. A finite-range lattice gas model with quasicrystal phases at positive temperatures. Work in progress.

A Proof of Proposition 1

In this appendix I will only explain how, starting from any generic Robinson configuration $\omega_0 \in \Omega_{RB}$ (without a cut), we can build a random ε -Bernoulli perturbation λ (a probability measure on $\Omega_{\mathcal{A}} \times \Omega_{\{0,1\}}$ with $\pi_2^*(\lambda) = \mathcal{B}(\varepsilon)^{\mathbb{Z}^2}$) such that $d_H(\omega_0, \omega) \geq \frac{1}{8}$ for λ -almost-any noisy configuration (ω, b) .

To conclude the proof from here would not be conceptually hard, but it would only distract us from the main purpose of this exploratory paper: undecidability.

We will build λ iteratively, as a limit of a sequence (λ_n) . Let us start with the trivial coupling $\lambda_1 = \delta_{\omega_0} \otimes \mathcal{B}(\varepsilon)^{\otimes \mathbb{Z}^2}$, constant on the first coordinate.

Let us properly explain how we obtain λ_2 from λ_1 . This transformation will be done independently on *each* of the 2-macro-tiles of ω_0 . We distinguish two cases, both illustrated in Figure 7. In this figure, we *obscured* the cells c with a noise $b_c = 1$, and we highlighted with a green border the crossed tiles on the central arm of a 2-macro-tile, the cells that guarantee 1-macro-tiles and 2-macro-tiles do not have the same colour on their central arm.




Figure 7: A locally admissible 3-macro-tile with obscured cells.

Either both highlighted cells are obscured, like in the top row, or at least one of them is clear, as in the bottom row, in which case we leave the 2-macro-tile unchanged. In the first case, when both highlighted cells are obscured, we *flip* the colours in the whole 2-macro-tile, meaning that any Red line turns Black and conversely, with probability $\frac{1}{2}$ using added independent randomness. By doing this, we trivially preserve local rules *inside* the 2-macro-tile, and we know no forbidden pattern is created at the

interface between the 2-macro-tile and the central arm of the 3-macro-tile containing it, precisely because that interface is partially hidden under an obscured cell.

Likewise, we go from λ_{n-1} to λ_n by flipping at random all the colours inside any n -macro-tile, *except* the two parts of the central arm that are “after” the crossed tiles, whenever those two crossed cells are obscured.

At last, let us consider λ a weak-* limit of (λ_n) . For a given cell $c \in \mathbb{Z}^2$ containing a 1-macro-tile with bumpy corners, we denote by $\text{flip}(c, n)$ the random variable equal to 1 when we flipped the colours of the n -macro-tile containing c in the process of building λ_n from λ_{n-1} .

As a flip requires *two* independent ε -Bernoulli cells to be obscured, and then occurs with probability $\frac{1}{2}$, we actually have $\text{flip}(c, n) \sim \mathcal{B}\left(\frac{\varepsilon^2}{2}\right)$. Remark how the sets of cells used at each step are disjoint from each other, thus the noises used in each flip are independent. Consequently, $(\text{flip}(c, n))_n$ is an *iid* sequence.

Conditionally to the event $\text{flip}(c, n) = 1$, the colour of the tile ω_c is uniformly distributed, and this fact still holds even if *other* flips happened on c . By the Borel-Cantelli lemma, λ -almost-surely, $\text{flip}(c, n) = 1$ for some rank n . Thence, the colour of ω_c is uniformly distributed under the probability measure λ .

Likewise, consider two distant 1-macro-tiles c and d . 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_0 goes to infinity. The variables $(\text{flip}(c, n))_{n < n_0}$ and $(\text{flip}(d, n))_{n < n_0}$ are independent, and conditionally to the fact that both of these sequences contain a 1, the colours of ω_c and ω_d are *iid*. This fact still holds if flips occur at scales greater than n_0 , so that it holds true for λ .

Without loss of generality, assume that 0 contains a 1-macro-tile, so that 1-macro-tiles are aligned on the network $(2\mathbb{Z})^2$. Consequently, the random configuration $[\text{colour of } \omega_{2c}]_{c \in \mathbb{Z}^2}$ induces an ergodic dynamical system, so that we may apply a pointwise ergodic theorem. To put it simply, this implies that the *frequency* of 1-macro-tiles with a Black arm in (ω, b) is λ -almost-surely equal to the probability of any 1-macro-tile having a Black arm, which we proved to be equal to $\frac{1}{2}$. Finally, the density of 1-macro-tiles among all the cells of \mathbb{Z}^2 is equal to $\frac{1}{4}$, so we can conclude that, for a λ -generic configuration (ω, b) , we have $d_H(\omega_0, \omega) \geq \frac{1}{2} \times \frac{1}{4} = \frac{1}{8}$.

B Proof of Proposition 3

Without loss of generality, by monotonicity, we just need to prove the bound for $(2n+1)$ -macro-tiles, with a big Red square in the middle. From now on, we seek a $O(12^n)$ bound.

Let us now denote r_n the number of tiles inside a Red square in the $(2n+1)$ -macro-tile. As we can see on Figure 3, in the process of forming a $(2n+3)$ -macro-tile, we will create a big central square, which is surrounded by 12 $(2n+1)$ -macro-tiles (and not $4 \times 4 = 16$ because 4 of these end up stuck *inside* the big square). 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 $r_1 = 25 \geq 16$ (there is only one Red square of length 5 tiles), we obtain $r_n \geq 4^{n+1}(4^n - 3^n)$. At the same time, a $(2n+1)$ -macro-tile has $(2^{2n+1} - 1)^2 \leq 4^{2n+1}$ tiles in total, so *at most* $4^{n+1}3^n = 4 \times 12^n$ tiles outside a Red square, which concludes the proof.