

Stochastic Flips on Dimer Tilings

Thomas Fernique and Damien Regnault

LIF, CNRS & Univ. de Provence
39 rue Joliot-Curie 13453 Marseille – France
{thomas.fernique,damien.regnault}@ens-lyon.org

Abstract. This paper introduces a Markov process inspired by the problem of quasicrystal growth. It acts over dimer tilings of the triangular grid by randomly performing local transformations, called *flips*, which do not increase the number of identical adjacent tiles (this number can be thought as the tiling energy). Fixed-points of such a process play the role of quasicrystals. We are here interested in the worst-case expected number of flips to converge towards a fixed-point. Numerical experiments suggest a $\Theta(n^2)$ bound, where n is the number of tiles of the tiling. We prove a $O(n^{2.5})$ upper bound and discuss the gap between this bound and the previous one. We also briefly discuss the average-case.

Introduction

Tilings are often used as a toy model for quasicrystals, with minimal energy tilings being characterized by local properties called *matching rules*. In this context, a challenging problem is to provide a theory for quasicrystals growth. One of the proposed theories relies on a relaxation process ([5] p. 356): a tiling with many mismatches is progressively corrected by local transformations called *flips*. Ideally, the tiling eventually satisfies all the matching rules and thus shows a quasicrystalline structure. It is compatible with experiments, where quasicrystals are obtained from a hot melt by a slow cooling during which flips really occur (Bridgman-Stockbarger method). It is however unclear whether only flips can explain successful coolings or if other mechanisms should be taken into account. This is deeply related with the convergence rate of such a flip-correcting process.

A cooling process aiming to be physically realist is described in [1]. It considers so-called *cut and project tilings* of any dimension and codimension, and performs flips which modify by ΔE the energy of the tiling with a probability proportional to $\exp(-\Delta E/T)$, so that the stationary distribution at fixed temperature T is the Boltzmann one.

A simplified cooling process is obtained by performing equiprobably at random only flips whose corresponding ΔE is above a fixed threshold. It has been studied on tilings of dimension one and codimension one (two-letter words) in [2]. We here focus on tilings of dimension two and codimension one (dimer tilings).

The paper is organized as follows. Sec. 1 introduces notations and basic definitions. We then describe, in Sec. 2, the cooling process we consider. We also formally state the convergence time we want to bound, and make conjectures based on numerical experiments. Sec. 3 is then devoted to the proof of an upper bound on the convergence time. This proof relies on a concentration result for some well-chosen function. We conclude the paper by a short section discussing prospects of this work, namely the non-tightness of the obtained theoretical bound, and the average convergence time, *i.e.*, when the initial tiling is chosen at random (instead of considering the one with the greatest convergence time).

1 Settings

Dimer tiling

Let $(\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ be the unit vectors of the Euclidean plane of direction $\frac{\pi}{6} + k\frac{2\pi}{3}$, $k = 1, 2, 3$. They generate the so-called triangular grid. A *dimer* is a lozenge tile made of two adjacent triangles of the grid. A *domain* is a connected subset of the grid. Then, a *dimer tiling* is a tiling of a domain by dimers (Fig. 1, left).

Lift

Let $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ be the canonical basis of the Euclidean space. The *lift* of a dimer tiling is its image by a map ϕ which is linear over tiles and satisfies $\phi(\mathbf{x} + \mathbf{v}_k) = \phi(\mathbf{x}) + \mathbf{e}_k$ for any two vertices \mathbf{x} and $\mathbf{x} + \mathbf{v}_k$ connected by an edge of the tiling. This map is uniquely defined up to a translation; we assume that $(0, 0, 0)$ belongs to the image of the domain boundary. Tiles are thus mapped onto facets of three-dimensional unit cubes, and dimer tilings are mapped onto stepped surfaces of \mathbb{R}^3 . This can be easily seen by shading tiles (Fig. 1, center).

Height

Following [6], we define the *height* of a point in a dimer tiling by the distance of its image under ϕ to the plane $x + y + z = 0$. We also define the *height* of a tile as the height of its center: this yields a third representation of dimer tilings, where tile colors depend on the height (Fig. 1, right).




Fig. 1. Three representations of the same dimer tiling.

Errors and energy

An *error* in a dimer tiling is an edge shared by two identical tiles (up to a translation). Error-free regions thus perfectly alternate tiles. The corresponding subsets of the triangular grid can be tiled by six-triangle hexagons, and the lift of an error-free region approximate the plane $x + y + z = 0$. In particular, all the tiles have the same height. The *energy* $E(\omega)$ of a dimer tiling ω is its total number of errors.

Islands and holes

Suppose that a dimer tiling ω contains a finite, connected and simply connected set of tiles σ whose boundary edges are errors and whose domain can be tiled by six-triangle hexagons. The number $A(\sigma)$ of such hexagons is called the *area* of σ . One checks that the height of tiles with an edge on the boundary of σ take only two values, say $h(\sigma)$ and $h(\bar{\sigma})$, respectively depending whether the tile is in σ or not. One says that σ is an *island* if $0 < |h(\bar{\sigma})| < |h(\sigma)|$, a *hole* otherwise. In both cases, $h(\sigma)$ is called the *height* of σ . Fig. 2 illustrates this.

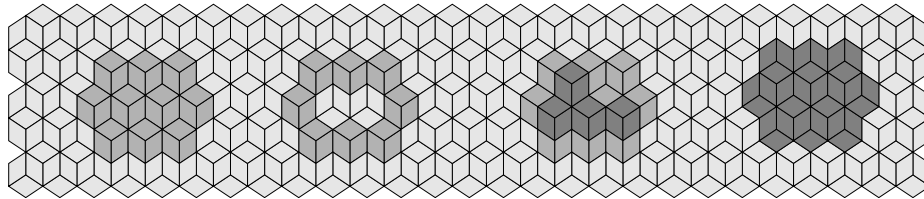


Fig. 2. Boundary tiles of this dimer tiling have the same height, say 0. From left to right: an island of height 1 and area 10; the same island with a hole of height 0 and area 2; once again the same island with a superposed island of height 2 and area 4; an island of height -1 and area 10 (which looks like a hole because its height is negative).

Volume

The *volume* $V(\omega)$ of a dimer tiling ω is defined as the sum of areas of its islands minus the sum of areas of its holes. Error-free tilings thus have volume zero. One checks that if the domain is finite, simply connected and admits an error-free tiling, then the errors of any dimer tiling of this domain are the boundaries of its islands and holes. Hence, tilings of volume zero are error-free. Otherwise, non-closed paths of errors can run across the domain, with endpoints (if any) on the boundary of the domain.

Partial order

A dimer tiling ω is said to be smaller than or equal to a dimer tiling ω' of the same domain, written $\omega \leq \omega'$, if each point has, in modulus, a height in ω smaller than or equal to its height in ω' . The set of tilings of a domain becomes a distributive lattice, whose extremal elements are the tilings of extremal volume.

Flip

Whenever a vertex x of a dimer tiling belongs to exactly three tiles, one gets a new dimer tiling by translating each of these tiles along the edges shared by the two other ones (or, equivalently, by rotating these tiles by $\frac{\pi}{3}$ around x). This local rearrangement of tiles is called a *flip* in x (Fig. 3). In the lift, it corresponds to add or remove one cube. One checks that a flip modifies the volume of a tiling by $\Delta V = \pm 1$ and its energy by $\Delta E \in \{0, \pm 2, \pm 4, \pm 6\}$. The total number of flips which can be performed on a tiling ω is denoted by $F(\omega)$.

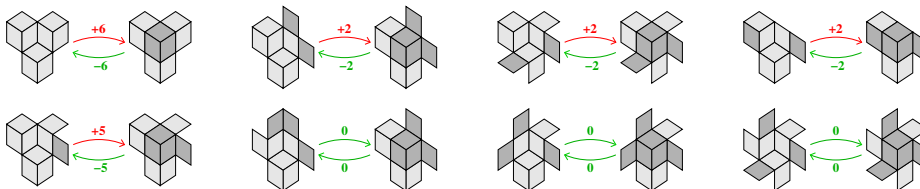


Fig. 3. A flip and all its possible immediate neighborhoods (up to a rotation), with the according variation of energy of the tiling being reported. The four leftmost cases correspond to merging of island or creation of holes. In the cooling process considered in Sec. 2, only flips which do not increase the energy will be allowed.

Note that the volume of a tiling turns out to be the minimal number of flips required to transform this tiling into an error-free tiling (if the domain admits such a tiling).

Flip-accessibility

In [6], Thurston has shown that any two dimer tilings can be connected by performing a sequence of flips, *i.e.* are mutually *flip-accessible*, if they tile the same finite and simply connected subset of the triangular grid. The case of dimer tilings of the whole triangular grid has been considered in [3]. Here, the physical motivations exposed in the introduction lead us to focus on a *constrained* flip-accessibility, where only flips which do not increase the number of errors are allowed. Moreover, we are not interested in a flip-accessibility where flips can be carefully chosen, but in the case where flips are *randomly* performed. The next section makes this more precise.

2 The cooling process

Let us fix a finite, connected and simply connected subset of the triangular grid which does admit an error-free dimer tiling made of n tiles. Let Ω_n denotes the set of all the possible dimer tilings of this subset (flip-accessibility yields that all these tilings have n tiles).

The so-called *cooling process* we are interested in is the Markov chain defined over Ω_n as follows. It starts from $\omega_0 \in \Omega_n$ and produces a sequence $(\omega_t)_{t=1,2,\dots}$ where ω_{t+1} is obtained by performing on ω_t a flip, uniformly chosen at random among the flips which do not increase the number of errors¹. In other words, the ΔE threshold discussed in the introduction is equal to zero. If there is no such flip, then the process stops (the tiling is said to be *frozen*).

Fig. 4 shows the cooling of a tiling of a hexagonal-shaped domain of 3779 tiles, chosen uniformly at random among the tilings of this domain. In this particular case, the process stops in 4290 steps, with the obtained frozen tiling turning out to be error-free. How general is such an evolution?




Fig. 4. Snapshots each 1000 flips of the cooling of a random tiling (from left to right). The evolution of island boundaries seems to be curvature driven: islands becomes smoother and smaller until they eventually disappear.

The following proposition will give us a first insight into this evolution:

Proposition 1. *Whenever a tiling in Ω_n has an error, a flip which do not increase the number of errors can be performed onto. Moreover, this flip can be chosen so that it decreases the volume of the tiling.*

Proof. The existence of a (finite) error-free tiling ensures that errors form the boundary of islands. Consider an island whose height is, in modulus, maximal. Such an island is made of hexagons tiled by three-tiles. Performing a flip on such a hexagon decreases the area of the island (hence the volume of the tiling) and decreases the number of errors by $2(k-3)$, where k is the number of error edges of this hexagon. It thus suffices to find a hexagon with at least three error edges. Let us follow clockwise the boundary of the island: the direction between two consecutive edges changes by $\pm\frac{\pi}{6}$, and there is more negative variations because the boundary is closed. We can thus find two consecutive negative variations: the three corresponding edges are errors on the boundary of the wanted hexagon. \square

¹ We discuss other ways to define such a process at the end of the last section.

This proposition ensures that a frozen tiling is necessarily error-free. Moreover, it shows that the process almost surely stops in a finite time, since the volume of tilings in Ω_n is uniformly bounded and decreases with probability at least $\frac{1}{n}$ at each step.

In order to describe more precisely this cooling process, we are interested in the probability distribution of the random variable T , called the *convergence time*, which counts the number of steps before it stops. Here, we focus only on the *worst expected convergence time* \hat{T} , defined by

$$\hat{T}(n) := \max_{\omega \in \Omega_n} \mathbb{E}(T \mid \omega_0 = \omega),$$

The fact that the cooling process almost surely stops in a finite time now reads

$$\hat{T}(n) < \infty.$$

Conversely, since the volume of a tiling decreases at most by one at each step, the worst expected convergence time is bounded below by the maximal volume of tilings in Ω_n . Since the maximal volume of a tiling of the hexagonal-shaped domain with n tiles and sides $\Theta(\sqrt{n})$ is $\Theta(n\sqrt{n})$, this yields the lower bound²

$$\hat{T}(n) = \Omega(n\sqrt{n}).$$

Fig. 5 shows the cooling of the maximal tiling of the same domain as Fig. 4.




Fig. 5. Snapshots each 18000 flips of the cooling of a tiling of maximal volume (from left to right). Islands shrink and disappear from top to bottom.

Our goal is to obtain a theoretical tight bound for the worst expected convergence time. Numerical experiments suggest that it is quadratic and correspond to tilings of maximal volume (Fig. 6). We thus conjecture:

$$\hat{T}(n) = \Theta(n^2).$$

² O , Ω and Θ are the usual Bachmann-Landau notations.




Fig. 6. Left, markers show averages of the convergence time on numerical experiments, starting from the maximal volume tiling (thought to have the worst convergence time). We also drew the curve cn^2 , $c = 6.689 \cdot 10^{-3}$, which matches rather well the markers. Right, dividing these experimental values by n^2 does not show a subpolynomial factor. We thus conjecture $\hat{T}(n) = \Theta(n^2)$.

3 A theoretical upper bound

Given a Markov chain (x_t) on Ω and $f : \Omega \rightarrow \mathbb{R}$, let $\Delta f(x_t)$ denote the variation

$$\Delta f(x_t) := f(x_{t+1}) - f(x_t).$$

Our upper theoretical bound on the expected convergence time will be obtained via a *drift function* ϕ (also called *Lyapunov-Foster function*) whose variation is uniformly bounded by some negative constant. More precisely, we will rely on a variation of the Foster's theorem (see, *e.g.*, [4]):

Theorem 1. *Let $(x_t)_{t \geq 0}$ be a Markov chain over a space Ω . Assume that there is $\varepsilon > 0$ and a map $\phi : \Omega \rightarrow [a, b] \subset (0, \infty)$ such that, whenever $\phi(x_t) > a$:*

$$\mathbb{E}[\Delta \phi(x_t) | x_t] \leq -\varepsilon.$$

Then, the expected value of the random variable $T := \min\{t \mid \phi(x_t) = a\}$ satisfies

$$\mathbb{E}(T) \leq \frac{b - a}{\varepsilon}.$$

We first need to introduce the notion of *triconvexity* (see Fig. 7):

Definition 1. *A dimer tiling is said to be triconvex if any segment of slope $0 \bmod \frac{\pi}{3}$ which connects two vertices x and y contains only vertices whose height is between the heights of x and y . The triconvex hull $\bar{\omega}$ of a dimer tiling ω is the smallest³ triconvex dimer tiling greater than or equal to ω .*

Let us then associate with any map X defined over dimer tilings (for example, the volume V or the energy E) a map \bar{X} defined as follows:

$$\bar{X}(\omega) := X(\bar{\omega}).$$

³ W.r.t. the partial order defined on page 3.




Fig. 7. A non-triconvex tiling (left, with fault segments) and its triconvex hull (right).

This will help us to avoid ambiguities, especially when considering variations. Indeed, let us stress that $\Delta\bar{X}(\omega)$ and $\Delta X(\bar{\omega})$ are generally different, since

- $\Delta\bar{X}(\omega) = X(\bar{\omega}_a) - X(\bar{\omega})$, where ω_a is obtained by performing a flip on ω ;
- $\Delta X(\bar{\omega}) = X(\omega_b) - X(\bar{\omega})$, where ω_b is obtained by performing a flip on $\bar{\omega}$.

Now, in order to use Th. 1, let us define a map ϕ over dimer tilings by:

$$\phi := 4V + E$$

We first consider the most simple case:

Lemma 1. *If ω is a triconvex dimer tiling with only one island, then*

$$\mathbb{E}[\Delta\phi(\omega)|\omega] \leq -\frac{12}{F(\omega)}.$$

Proof. Follow clockwise the boundary of the island: two consecutive edges make either a *salient* or a *reflex* angle, depending on whether we turn leftwards or rightwards (see Fig. 8). If there is a flip around some vertex x , then the six-triangle hexagon of center x has at least 3 error edges on its boundary. If these error edges are consecutive, then they form $i \geq 2$ consecutive similar angles, and by performing the flip:

- if the i angles are salient, then $\Delta V = -1$ and $\Delta E = 4 - 2i$, that is, $\Delta\phi = -2i$. This yields $\Delta\phi = -2$ per salient angle.
- if the i angles are reflex, then $\Delta V = 1$ and $\Delta E = 4 - 2i$, that is, $\Delta\phi = 8 - 2i$. This yields $\Delta\phi \leq +2$ per reflex angle.

If there are only such flips, then the claimed bound follows by checking by induction on its length that the boundary of an island (triconvex or not) has always six more salient than reflex angles. Otherwise, consider a flip in x such that the error edges on the boundary of the six-triangle hexagon of center x are non consecutive. At least two of these edges must be parallel. Thus, by triconvexity, the vertex x is in ω . Hence, such a flip decreases by 1 the volume (and disconnects the island). In particular, it decreases ϕ by at least 4, so that it can only improve the bound. \square

Since islands are not always triconvex, we need the following technical lemma, which links the variations $\Delta\bar{\phi}(\omega)$ and $\Delta\phi(\bar{\omega})$ (see Fig. 9):




Fig. 8. A single triconvex island with 24 salient angles (downwards triangles) and 18 reflex angles (upwards triangles). A number in a circle shows the variation of ϕ if a flip is performed around this circle. There are $F(\omega) = 12$ performable flips, and one computes $\mathbb{E}(\Delta\phi) = -\frac{3}{2} \leq -\frac{12}{F(\omega)}$. The bound is not tight because there is one flip (in the thick circle) which disconnects the island.

Lemma 2. *If ω is a dimer tiling with only one island, then*

$$F(\omega)\mathbb{E}[\Delta\bar{\phi}(\omega)|\omega] \leq F(\bar{\omega})\mathbb{E}[\Delta\phi(\bar{\omega})|\omega]$$

Proof. Let us fix ω and let ΔX_y denote the X -variation by performing a flip around y . Proving the lemma is then equivalent to prove

$$\sum_{x \in F(\omega)} \Delta\bar{\phi}_x(\omega) \leq \sum_{x \in F(\bar{\omega})} \Delta\phi_x(\bar{\omega}).$$

This easily follows from these three sublemmas (see Fig. 9):

- If $x \in F(\omega) \setminus F(\bar{\omega})$, then $\Delta\bar{\phi}_x(\omega) = 0$ since the flip does not modify $\bar{\omega}$.
- If $x \in F(\bar{\omega}) \setminus F(\omega)$, then $\Delta\phi_x(\bar{\omega}) > 0$.

Let us first prove by contradiction that x is not in the island of $\bar{\omega}$. Assume the contrary: x thus belongs to a segment with endpoints, say y and z , in the island of ω . Consider the six edges of the six-triangle hexagon modified by the flip around x . Since $x \in F(\bar{\omega})$, at least three of these edges are error edges of $\bar{\omega}$. If x is in the island of ω , these (at least) three error edges are also error edges of ω , and thus $x \in F(\omega)$, what is ruled out by the hypothesis. Otherwise, then consider a path in the island of ω which links y to z (such a path exists because ω has only one island). This path and the segment $[y, z]$ form a circuit in $\bar{\omega}$, and the triconvexity of $\bar{\omega}$ ensures that all the vertices enclosed by this circuit are in the island of $\bar{\omega}$, that is, there is no error edges enclosed by this circuit. Now, consider the six-triangle hexagon of center x . It has two opposite edges on the segment $[y, z]$, which thus cannot be error edges of $\bar{\omega}$. The remaining four edges are splitted in two pairs of adjacent edges. One of these pair is enclosed in the above circuit, hence cannot contain error edges of $\bar{\omega}$. The other pair can contain error edges of $\bar{\omega}$, but obviously no more than two, what contradicts $x \in F(\bar{\omega})$. We conclude that x is not in the island of $\bar{\omega}$.

Let us then show that $\Delta\phi_x(\bar{\omega}) > 0$. On the one hand, $\Delta V_x(\bar{\omega}) = 1$ since the height of a vertex not in the island of $\bar{\omega}$ can only increase. On the other hand, $\Delta E_x(\bar{\omega}) = 0$. Indeed, if there were four or more error edges around x , then two of them would be opposite and the triconvexity of $\bar{\omega}$ would yield that x is in the island of $\bar{\omega}$. Finally, $\Delta\phi_x(\bar{\omega}) = 4 > 0$.

• If $x \in F(\omega) \cap F(\bar{\omega})$, then $\Delta\bar{\phi}_x(\omega) \leq \Delta\phi_x(\bar{\omega})$.

If x is not in the island of $\bar{\omega}$ (hence not in the one of ω), then one checks that $\Delta E_x(\bar{\omega}) = \Delta\bar{E}_x(\omega) = 0$ and $\Delta V_x(\bar{\omega}) = \Delta\bar{V}_x(\omega) = 1$, hence $\Delta\bar{\phi}_x(\omega) = \Delta\phi_x(\bar{\omega})$.

If x is in the island of $\bar{\omega}$ but not in the one of ω , then arguments already used above to prove the second sublemma lead to $x \notin F(\bar{\omega})$: this case is ruled out.

The last case is when x is in the island of ω (hence in the one of $\bar{\omega}$). If the error edges around x are the same in ω and $\bar{\omega}$, then one easily checks that, once again, $\Delta\bar{\phi}_x(\omega) = \Delta\phi_x(\bar{\omega})$. Otherwise, one checks that x is in a position similar to the bold square on Fig. 9 (top-left): one has $\Delta E_x(\bar{\omega}) = \Delta\bar{E}_x(\omega) = 0$, but $\Delta\bar{\phi}_x(\omega)$ can be much smaller than $\Delta V_x(\bar{\omega}) = -1$, since by taking the triconvex hull after performing a flip one can recover much lesser than by simply performing a flip on the triconvex hull. Hence, $\Delta\bar{\phi}_x(\omega) \leq \Delta\phi_x(\bar{\omega})$. \square




Fig. 9. An island of a tiling ω . Its triconvex hull (thick boundary) corresponds to Fig. 8. Flips are indicated by circles, triangles or squares, respectively depending whether they are in $F(\omega) \setminus F(\bar{\omega})$, $F(\bar{\omega}) \setminus F(\omega)$ or $F(\omega) \cap F(\bar{\omega})$. Lemma 2 links $\Delta\bar{\phi}(\omega)$ and $\Delta\phi(\bar{\omega})$. Note that performing the flip indicated by the bold square (top-left) yields $\Delta V(\bar{\omega}) = -1$ but $\Delta\bar{V}(\omega) = -4$ (see proof of Lem. 2).

Now, Lem. 1 applied to $\bar{\omega}$ ensures

$$F(\bar{\omega})\mathbb{E}[\Delta\phi(\bar{\omega})|\omega] \leq -12.$$

Hence, Lem. 2 yields

$$\mathbb{E}[\Delta\bar{\phi}(\omega)|\omega] \leq -\frac{12}{F(\omega)}.$$

Let us stress that this bound involves $F(\omega)$ and not $F(\bar{\omega})$, although the flips performable on ω but not on $\bar{\omega}$ seem to have no effect on the triconvex hull we

look at, because they are “hidden” inside. But they do have an effect: they slow down the dynamics we are interested in! This is why the above bound cannot be improved by replacing $F(\omega)$ by $F(\bar{\omega})$, which is possibly much smaller⁴.

Let us now extend this result to the case of several islands. The only problem which can occur is that flips can merge islands. However, one big advantage of considering $\bar{\phi}$ instead of ϕ is that mergings become rather trivial, as we will see.

Lemma 3. *If ω is a dimer tiling whose islands have all the same height, then*

$$\mathbb{E}[\Delta\bar{\phi}(\omega)|\omega] \leq -\frac{12}{F(\omega)}.$$

Proof. If there is no flip which can merge two or more islands of $\bar{\omega}$, then the result just follows from Lem. 1 and 2. Otherwise, consider such a flip. It is thus performed around a vertex, say x , which is not in an island of $\bar{\omega}$. The triconvexity of $\bar{\omega}$ ensures that this is possible only if error and non-error edges alternate on the boundary of the six-triangle hexagon of center x . This flip thus merges exactly three islands of $\bar{\omega}$. Then, the triconvexity of $\bar{\omega}$ ensures that these islands are necessarily stick-shaped, that is, formed of aligned six-triangle hexagons (Fig. 10, left). At least the endpoints of these stick-shaped islands are in ω , and one easily checks that the claimed bound holds (Fig. 10, right). \square




Fig. 10. A flip which merges islands of $\bar{\omega}$ necessarily merges exactly three islands, which are moreover stick-shaped (left: $\bar{\omega}$). The flips in ω which modify $\bar{\omega}$ are such that $\bar{\phi}$ decreases on expectation (right: ω).

We are now in a position to prove:

Theorem 2. *The worst expected convergence time satisfies:*

$$\hat{T}(n) = O(n^2\sqrt{n}).$$

⁴ Imagine, for example, a tiling having a round island of size k with $\Theta(k)$ small holes, each of them allowing a flip. By taking the triconvex hull, only the flips on the boundary of this island are still performable (there can be as few as 6 such flips).

Proof. Let us consider the highest level of the tiling (that is, the highest islands). The only problem to apply Lem. 3 is that some flips can be unperformable, due to lower islands. However, only flips which would increase the volume of $\bar{\omega}$ become unperformable. Moreover, these flips could not decrease the energy of $\bar{\omega}$ (we have already seen in the proof of Lem. 2 that a flip which increases the volume of a triconvex island cannot modify its energy). The bound of Lem. 3 thus holds *a fortiori* for this highest level. Since both the volume and the energy of the highest level of the tiling are at most n , and $F(\omega) \leq n$, one can apply Th. 1 to $\bar{\phi}$ with $\varepsilon = 1/n$: this yields that the highest level of the tiling disappears in time $O(n^2)$. The claimed bound then follows since the tiling has $O(\sqrt{n})$ levels. \square

4 Prospects

Towards a tight bound

Numerical experiments led us to conjecture that the expected convergence time is in $O(n^2)$. However, Th. 2 proves only a $O(n^2\sqrt{n})$ upper bound.

A possible explanation for this gap is that we have considered triconvex hulls of islands instead of real islands. The main advantage was that we thus not had to consider holes or complicated merging/splitting of islands. But the flips inside convex hulls can also speed up the dynamics, effect that is not taken into account with triconvex hulls. However, considering tilings whose islands have all the same height (*i.e.*, two-level tilings) suggests that this effect is neglectable, at least in the worst case. Indeed, it follows from the proof of Th. 2 that the worst expected convergence time for such tilings is $\hat{T}(n) = O(n^2)$ (since the volume is in $O(n)$), and this bound seems tight for tilings made of concentric ring-shaped islands (Fig. 11). Another explanation of this gap is thus necessary.

Recall that, in the proof of Th. 2, we applied Lem. 3 only on the highest level of the tiling. This way, we make as if the $F(\omega)$ factor in the bound of Lem. 3 would count only flips on this level. But it also counts the flips on all the lower levels, which likely tend to shrink islands as well. One thus could expect that the bound of Lem. 3 still holds if it is multiplied by the number of levels. Since they are around \sqrt{n} levels, this would precisely fill the gap between our theoretical and experimental upper bounds. But the problem is that stacked islands do not evolve independently. In particular, a flip which would shrink an isolated island can be blocked by stacking another island onto it. Hence, an island can force a lower island to grow, and thus slows down the dynamics (consider, for example, the lower islands on Fig. 5: they cannot freely shrink while the higher islands do not let them enough space). In practice, this effect is limited for (at least) two reasons. First, blocking also have the contrary effect: an island can forbid the growth of a higher island. Second, the higher an island is, the faster it shrinks, and thus the lesser it blocks the islands it is stacked onto. It seems however not so easy to turn this into a formal proof. . .




Fig. 11. Snapshots each 3000 flips of the cooling of a tiling with concentric ring-shaped islands all of height 1 (from left to right). Ring radii follow a power law with scaling exponent greater than one in order to avoid ring merging when islands shrink, but close to one in order to have as many as possible rings which slow down the dynamics. The cooling of such a small volume tiling looks like the one of a maximal volume tiling (compare to Fig. 5).

Last, let us point out that the simple theoretical lower bound we provided is $\Omega(n\sqrt{n})$, that is, we also need to gain a factor \sqrt{n} .

Average expected convergence time

In this paper, we only consider the worst expected convergence time, $\widehat{T}(n)$. However, it is also natural (even more, according to the physical motivations exposed in the introduction), to consider the average expected convergence time:

$$\bar{T}(n) := \frac{1}{\#\Omega_n} \sum_{\omega \in \Omega_n} \mathbb{E}(T \mid \omega_0 = \omega).$$

In this case, numerical experiments (Fig. 12) led us to conjecture:

$$\bar{T}(n) = \Theta(n\sqrt{n}).$$

In order to bound the average expected convergence time by similar technics as for the worst expected convergence time, we first need to compute the average values of some functions, *e.g.*, the volume V , the energy E or the number of levels H . This is what we did in the one-dimensional case in [2]. However, this two-dimensional case is much harder. Let us just mention that numerical experiments indicate that both the average volume and the average energy of a tiling are linear (in the number of tiles), while the number of levels is logarithmic.

Another cooling process

The cooling process we considered performs, at each step, a flip randomly chosen among the possible flips. This yields a non-reversible Markov chain since tilings do not all have the same number of possible flips (and this point is one




Fig. 12. Analog of Fig. 6 for the average expected convergence time. Here, the drawn curve (left) has equation $cn\sqrt{n}$, $c = 1.669 \cdot 10^{-2}$. We thus conjecture $\bar{T}(n) = \Theta(n\sqrt{n})$.

of the difficulties we had to handle to bound convergence times). Another way to define the cooling, maybe more usual, is to pick up at random a vertex of the tiling and to try to perform a flip around it, if possible⁵. Clearly, this slows down the dynamics. However, numerical experiments suggest that the worst convergence time is still in $\Theta(n^2)$, probably because most of the cooling time of a maximal volume tiling is spent when the tiling has errors almost everywhere, so that both processes run similarly. Note that our (not tight) theoretical upper bound still holds by simply replacing the number $F(\omega)$ of performable flips by the number of vertices in the analysis. On the contrary, according to numerical experiments, the average convergence time is in $\Theta(n^2)$ (against $\Theta(n\sqrt{n})$ for the former cooling process): many randomly chosen vertices do not lead to a flip.

Acknowledgments. We would like to thank Olivier Bodini (LIP6, Paris), Éric Rémila (LIP, Lyon) and Mathieu Sablik (LATP, Marseille) for useful discussions. We also thank the anonymous referees, who suggested valuable corrections.

References

1. O. Bodini, Th. Fernique, D. Regnault, *Crystallization by stochastic flips*, to appear in Journal of Physics: Conference Series.
2. O. Bodini, Th. Fernique, D. Regnault, *Stochastic flips on two-letter words*, Analytic Algorithmics & Combinatorics (2010), p. 48–55.
3. O. Bodini, Th. Fernique, É. Rémila, *A Characterization of flip-accessibility for rhombus tilings of the whole plane*, Inf. Comput. **206** (2008), p. 1065–1073.
4. G. Fayolle, V. A. Malyshev, M. V. Menshikov, *Topics in the constructive theory of countable Markov chains*, Cambridge Univ. Press, 1995.
5. C. Janot, *Quasicrystals: The State of the art*, Oxford University Press, 1992.
6. W. P. Thurston, *Conway tilings group*, Amer. Math. Month. **97** (1990), p. 757–773.

⁵ This would define a reversible Markov chain if all the flips were allowed, but not in our case, where only the ones which do not increase the number of errors are allowed.