

Stochastic Flips on Two-letter Words

O. Bodini¹, Th. Fernique², and D. Regnault²

¹ LIP6, CNRS & Univ. Paris 6,
4 place Jussieu 75005 Paris - France,
olivier.bodini@lip6.fr

² LIF, CNRS & Univ. de Provence
39 rue Joliot-Curie 13453 Marseille – France
{thomas.fernique,damien.regnault}@lif.univ-mrs.fr

Abstract. This paper introduces a simple Markov process inspired by the problem of quasicrystal growth. It acts over two-letter words by randomly performing *flips*, a local transformation which exchanges two consecutive different letters. More precisely, only the flips which do not increase the number of pairs of consecutive identical letters are allowed. Fixed-points of such a process thus perfectly alternate different letters. We show that the expected number of flips to converge towards a fixed-point is bounded by $O(n^3)$ in the worst-case and by $O(n^{5/2} \ln n)$ in the average-case, where n denotes the length of the initial word.

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 ([4] 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. It is however unclear whether only flips can explain successful coolings or if other mechanisms should also be taken into account. This question is deeply related with the convergence rate of such a flip-correcting process.

A relaxation process which aims to be physically realist is described in [1]. It considers general cut and project tilings of any dimension and codimension, and performs each flip which modifies by ΔE the energy of the tiling with a probability depending not only on ΔE but also on a temperature parameter T such that the stationary distribution (at fixed T) is the Boltzmann distribution. In this paper, we focus on a very rough version of this general process. First, we consider only tilings of dimension and codimension one, which correspond to two-letter words. Second, the flips with a corresponding ΔE being above a fixed threshold are equiprobably performed, while the other flips are simply forbidden.

The paper is organized as follows. In Sec. 1, we clearly and formally state the definition of the stochastic process we consider, as well as the main question we are interested in, that is its *convergence time*. In Sec. 2, we state and prove the main result of this paper, which shows that the expected number of flips performed by the process to converge is cubic in the length of the initial word (Theorem 1). The proof mainly relies on a well-chosen function, called *variant*, whose expected value strictly decreases. This result is then extended in Sec. 3 to the case where the initial words are randomly chosen according to some particular distribution which aims to be physically realist (Theorem 2). We conclude the paper by a short section discussing perspectives.

1 Settings

A *configuration of length n* is a word $w = w_1 \dots w_n$ over $\{1, 2\}$ with $|w|_1 = |w|_2$, *i.e.*, with as many occurrences of the letter 1 as of the letter 2. Such an object is also sometimes called a *grand Dyck path*, or a *bridge*. We denote by \mathcal{W}_n the set of configurations of length n .

It is convenient to represent a configuration w as the broken line of the Euclidean plane linking the points $(k, |w_1 \dots w_k|_1 - |w_1 \dots w_k|_2)_{k=0, \dots, |w|}$ (see Fig. 1).



Fig. 1. The configuration 112111212221122122222112221121112111212.

A *flip* at position $1 \leq i < n$ on a configuration $w \in \mathcal{W}_n$ is the local transformation which exchanges w_i and w_{i+1} , provided that these letters are different. The *height* of such a flip is the integer $|w_1 \dots w_{i-1}|_1 - |w_1 \dots w_{i-1}|_2$. Geometrically, performing a flip on a configuration corresponds to move upwards or downwards by 2 a point of the broken line which represents the configuration.

Flip thus acts over configurations, and it is not hard to check that any two configurations with the same length are connected by a sequence of flips.

A configuration w is said to have a *mismatch* at position i if $w_i = w_{i+1}$. For example, the configuration on Fig. 1 has 18 mismatches, while configurations on Fig. 2 both have 4 mismatches. The total number of mismatches of w is denoted




Fig. 2. A flip of height 2 performed at position 5 (leftwards). On the representation, the height and the position of a flip are the abscissa and the ordinate of the center between the two positions of the moved vertex (that is, the middle of the arrow).

by $E(w)$. One easily checks that it ranges from 0 for configurations $(12)^n$ and $(21)^n$ to $2(n-1)$ for configurations $1^n 2^n$ and $2^n 1^n$. Physically, $E(w)$ can be thought as the *energy* of the configuration w , with the configurations $(12)^n$ and $(21)^n$ thus being ground states.

We now define a Markov chain on \mathcal{W}_n that we call *cooling process*. It starts from $w_0 \in \mathcal{W}_n$ and produces a sequence $(w_t)_{t=1,2,\dots}$ defined by:

$$w_{t+1} = \begin{cases} w_t & \text{if } E(w_t) = 0, \\ w' & \text{otherwise,} \end{cases}$$

where w' is obtained by performing on w_t a flip uniformly chosen among the flips which do not increase the number of mismatches³. Note that a flip modifies the number of mismatches by at most 2. Flips which decrease the number of mismatches are said to be *irreversible* (they cannot be performed back) while the ones which do not modify the number of mismatches are said to be *reversible*.



Fig. 3. Flips can increase the number of mismatches (red arrows), let it unchanged (blue arrows) or decrease it (green arrows). The two last are respectively called reversible and irreversible flips: they are the only ones performed during the cooling process.

Since the number $E(w_t)$ of mismatches is non-increasing, it is natural to ask whether it reaches zero or not, and at which rate. We therefore introduce the *convergence time* of this process, which is the random variable T counting the

³ In other words, the ΔE threshold discussed into the introduction is equal to zero.

number of steps required to transform the initial configuration w_0 into one of the two stable configurations without any mismatch, that is

$$T := \min\{t \geq 0 \mid E(w_t) = 0\}.$$

Then, the *worst expected convergence time* is defined by

$$\widehat{T}(n) := \max_{w \in \mathcal{W}_n} \mathbb{E}(T \mid w_0 = w).$$

We are not only interested in the worst case, but also in the average case. Given a distribution μ over \mathcal{W}_n , we define the μ -*averaged expected convergence time* by

$$T_\mu(n) := \sum_{w \in \mathcal{W}_n} \mu(w) \mathbb{E}(T \mid w_0 = w).$$

2 Bounding the worst expected convergence time

In this section, we prove the main result of this paper:

Theorem 1. *The worst expected convergence time is cubic:*

$$\widehat{T}(n) = O(n^3).$$

We will rely on the following proposition, proven in [2], which bounds from above the expected convergence time of so-called *variants*:

Proposition 1 ([2]). *Consider a markov chain $(x_t)_{t \geq 0}$ over a space Ω and a positive real map $\phi : \Omega \rightarrow [a, b]$, called variant. Assume that there is $\varepsilon > 0$ such that whenever $\phi(x_t) > a$, $\mathbb{E}[\phi(x_{t+1}) - \phi(x_t) \mid x_t] \leq -\varepsilon$. Then, the expected value of the random variable $T := \min\{t \mid \phi(x_t) = a\}$ satisfies*

$$\mathbb{E}(T) \leq \frac{\mathbb{E}(\phi(x_0))}{\varepsilon}.$$

Variants are thus sort of “strict” supermartingales, and we need to define a suitable one in order to bound the expected convergence time towards mismatch-free configurations. Since the number E of mismatches is non-increasing by definition of the cooling process, we are naturally tempted to use it as a variant. However, there are configurations having only reversible flips, whose expected variation of mismatches is thus equal to zero (as, for example, $w = 1211212212$). In order to refine, we introduce the notion of *Dyck factors* (see Fig. 4):

Definition 1. *Consider a configuration $w = p \cdot v \cdot s$. The factor $v = v_1 \cdots v_k$ is said to be a positive Dyck factor of length k and height $|p|_1 - |p|_2$ if:*

$$|v|_1 = |v|_2, \quad \forall i \in \{1, \dots, k\}, \quad |v_1 \cdots v_i|_1 \geq |v_1 \cdots v_i|_2, \quad |p|_1 - |p|_2 \geq 0.$$

Exchanging letters 1 and 2 in the above conditions defines a negative Dyck factor. A Dyck factor is maximal if no Dyck factor with the same height contains it.



Fig. 4. A positive maximal Dyck factor of height 2 and length 6 (in blue), and a negative maximal Dyck factor of height -1 and length 12 (in red).

Let us stress that Dyck factors are either positive or negative, *i.e.*, their representations do not cross the $y = 0$ axis. In particular, any configuration can be decomposed in its maximal Dyck factors of height zero, with signs alternating.

We use Dyck factors to define our variant, which depends on a parameter α :

Definition 2. For $\alpha \in (0, 1)$, let ϕ_α be defined on a configuration w by:

$$\phi_\alpha(w) = \sum_{v \in DF(w)} (1 + |v|_1)^\alpha,$$

where $DF(w)$ denotes the set of maximal Dyck factors of w .

Contrary to the number E of mismatches, ϕ_α can be increased by some flips. However, the concavity of $x \rightarrow x^\alpha$, which gives to small Dyck factors a weight proportionally bigger than the weight given to long Dyck factors, will ensure that the average variation over all the flips of a configuration is always negative. More precisely:

Lemma 1. Let $w \in \mathcal{W}_n$. If $E(w_t) > 0$, then the variant ϕ_α satisfies

$$\mathbb{E}(\phi_\alpha(w_{t+1}) - \phi_\alpha(w_t) | w_t) \leq -\frac{\alpha(1-\alpha)}{2} n^{\alpha-2}.$$

Proof. Consider a configuration $w = w_1 \cdots w_n$, with $E(w) > 0$. This ensures that at least one flip (reversible or irreversible) can be performed on w .

Let us first assume that only reversible flips can be performed on w , that is, neither 1122 nor 2211 are factors of w . This is equivalent to say that maximal Dyck factors contain at least four letters. We will group these reversible flips by pairs and prove that the average variation of ϕ_α over each pair is bounded from above by $\frac{\alpha(1-\alpha)}{2} n^{\alpha-2}$.

Each flip of positive height which transforms 21 into 12 (it increases ϕ_α) is between two positive flips which transform 12 into 21 (they decrease ϕ_α). Reversibility ensures that one of these two flips has the same high as the central flip, while the other one is higher; we group the central flip and its higher neighbor. We proceed symmetrically for flips of negative height. Fig. 5 illustrates this.



Fig. 5. Each flip which increases ϕ_α (red arrows) can be paired (light green boxes) with a flip which decreases ϕ_α (blue arrows), such that the average variation of ϕ_α over the two flips is negative. Isolated flips can only decrease ϕ_α .

Consider two paired reversible flips, say of positive height (the negative height case is symmetric). Performing the lowest flip increases by 1 the number p of 1's in the maximal positive Dyck factor of w starting at the first letter between the pair of flips, while performing the highest flip decreases by 1 the number q of 1's in some other maximal positive Dyck factor of w ; one moreover has $p \geq q$ since the latter Dyck factor is higher than the former one (see Fig. 6).




Fig. 6. Consider two paired flips (central configuration): the maximal Dyck factor extended by the flip increasing ϕ_α (red arrow, performed leftwards) is greater than the one shortened by the flip decreasing ϕ_α (blue arrow, performed rightwards), *i.e.*, $p \geq q$.

Thus, the average variation of ϕ_α is

$$\Delta := \frac{1}{2} ((p+2)^\alpha - (p+1)^\alpha + q^\alpha - (q+1)^\alpha).$$

Since $x \rightarrow x^\alpha$ is concave for $\alpha \in (0, 1)$, $p \geq q$ yields

$$\Delta \leq \Delta' := \frac{1}{2} ((p+2)^\alpha - (p+1)^\alpha + p^\alpha - (p+1)^\alpha).$$

For $k \geq 1$, let $u_k = \frac{\alpha(\alpha-1)\dots(\alpha-k+1)}{k!}$. Note that $u_{2k} < 0$. One computes:

$$\begin{aligned} \frac{2\Delta'}{(p+1)^\alpha} &= \left(1 + \frac{1}{p+1}\right)^\alpha + \left(1 - \frac{1}{p+1}\right)^\alpha - 2 \\ &= 1 + \sum_{k \geq 1} \frac{u_k}{(p+1)^k} + 1 + \sum_{k \geq 1} (-1)^k \frac{u_k}{(p+1)^k} - 2 \\ &= \sum_{k \geq 1} 2 \frac{u_{2k}}{(p+1)^{2k}} \leq 2 \frac{u_2}{(p+1)^2} = \frac{\alpha(\alpha-1)}{(p+1)^2}. \end{aligned}$$

With $p < n$, this yields the claimed bound for the pair:

$$\Delta \leq \Delta' \leq \frac{(p+1)^\alpha}{2} \times \frac{\alpha(\alpha-1)}{(p+1)^2} = -\frac{\alpha(1-\alpha)}{2}(p+1)^{\alpha-1} \leq -\frac{\alpha(1-\alpha)}{2}n^{\alpha-2}.$$

Last, note that there are some flips which have not been paired, namely the ones which are immediately after or before a crossing of the $y = 0$ axis. They however do not cause trouble, because each of them decreases a maximal Dyck factor of length at most n , hence decreases ϕ_α by at least

$$(n-1)^\alpha - n^\alpha \leq -\alpha n^{\alpha-1} \leq -\frac{\alpha(1-\alpha)}{2}n^{\alpha-2}.$$

Let us now consider the general case, *i.e.*, when there are also irreversible flips. We will show that an irreversible flip increases ϕ_α lesser than (resp. decreases ϕ_α more than) two reversible flips do. Hence, the average variation of ϕ_α can only be smaller than in the previous case (which is already negative). Intuitively, it is as if each irreversible flip is splitted into two reversible flips before performing the above pairing process.

Let us be more precise. Consider, first, the case of a positive flip which transforms 21 into 12. It replaces two maximal Dyck factors, say with respectively p and q letters 1, by one with $p+q+1$ letters 1 (see Fig. 7). Thus, on the one hand:

$$\Delta = (p+q+2)^\alpha - (p+1)^\alpha - (q+1)^\alpha.$$

On the other hand, the variation that would result of two reversible flips is:

$$\Delta' = (p+2)^\alpha - (p+1)^\alpha + (q+2)^\alpha - (q+1)^\alpha.$$

One thus has $\Delta \leq \Delta'$, since for $0 < \alpha < 1$:

$$(p+2)^\alpha + (q+2)^\alpha \geq ((p+2) + (q+2))^\alpha = (p+q+4)^\alpha \geq (p+q+2)^\alpha.$$




Fig. 7. An irreversible flip which replaces two maximal Dyck factors with respectively p and q letters 1 by a single maximal Dyck factor with $p+q+1$ letters 1.

We then consider the case of a positive irreversible flip which transforms 12 into 21. This just removes a maximal Dyck factor with one letter 1 (see Fig. 8). Thus, on the one hand:

$$\Delta = -2^\alpha.$$




Fig. 8. An irreversible flip which removes a maximal Dyck factor with one letter 1.

On the other hand, each of the two neighbor flips of this flip (one on the left, the other on the right, both transforming 21 into 12) either decreases the variant if it is negative, or increases it by:

$$\Delta' = (p+2)^\alpha - (p+1)^\alpha \leq 3^\alpha - 2^\alpha.$$

The average variation of ϕ_α is thus bounded from above by:

$$\frac{2\Delta' + \Delta}{3} \leq \frac{2(3^\alpha - 2^\alpha) - 2^\alpha}{3}.$$

A computation shows that, for $n \geq 2$, this quantity is lesser than $-\frac{\alpha(1-\alpha)}{2}n^{\alpha-2}$ for any $\alpha \in (0, 1)$, with equality for $\alpha = 1$. The case of a negative irreversible flip is symmetric. This concludes the proof. \square

This lemma thus provides a suitable ε (depending on α and n) for Prop. 1. Let us now bound ϕ_α and show that it is minimal for mismatch-free configurations:

Lemma 2. *For any configuration $w \in \mathcal{W}_n$, one has*

$$(1+n)^\alpha \leq \phi_\alpha(w) \leq (2n)^{\alpha+1},$$

with the lower bound being reached if and only if $E(w) = 0$.

Proof. Let us first focus on the upper bound. We prove by induction on $|v|$ that, for any positive configuration v ,

$$\phi_\alpha(v) \leq (1 + |v|_1)^{\alpha+1}.$$

The inequality is satisfied for $n = 2$ since one has $\phi_\alpha(12) = 2^\alpha$. Let us now consider a configuration v of length n , and assume that the result holds for any shorter positive configuration. Let u_1, \dots, u_p be the maximal (positive) Dyck factors of v of height 1. One thus has

$$\phi_\alpha(v) = (1 + |v|_1)^\alpha + \sum_{i=1}^p \phi_\alpha(u_i) \leq (1 + |v|_1)^\alpha + \sum_{i=1}^p (1 + |u_i|_1)^{\alpha+1}.$$

We then use the convexity of $x \rightarrow x^{\alpha+1}$ and compute

$$\phi_\alpha(v) \leq (1 + |v|_1)^\alpha + \left(\sum_{i=1}^p 1 + |u_i|_1 \right)^{\alpha+1} = (1 + |v|_1)^\alpha + |v|_1^{\alpha+1} \leq (1 + |v|_1)^{\alpha+1}.$$

The claimed result follows by induction. It also holds for negative configurations. Let us now consider the case of a general configuration w . Let v_1, \dots, v_q be its maximal Dyck factors of height 0. Since they are alternatively positive and negative configurations, one relies on the previous result to compute

$$\phi_\alpha(w) = \sum_{i=1}^q \phi_\alpha(v_i) \leq \sum_{i=1}^q (1 + |v_i|_1)^{\alpha+1} \leq \left(\sum_{i=1}^q 1 + |v_i|_1 \right)^{\alpha+1} = (q + |w|_1)^{\alpha+1}.$$

Since $q \leq |w|_2$, this yields $\phi_\alpha(w) \leq |w|^{\alpha+1}$. This proves the claimed upper bound.

Let us now focus on the lower bound. Let w be a configuration and v_1, \dots, v_q be its maximal Dyck factors of height 0. By relying on the concavity of $x \rightarrow x^\alpha$, one computes

$$\phi_\alpha(w) = \sum_{i=1}^q \phi_\alpha(v_i) = \sum_{i=1}^q (1 + |v_i|_1)^\alpha \geq \left(\sum_{i=1}^q 1 + |v_i|_1 \right)^\alpha = (q + |w|_1)^\alpha.$$

Since $|w|_1 = n$ and $q \geq 1$, this proves the claimed lower bound. Moreover, if this bound is reached, then $q = 1$, that is, w is either positive or negative. In both cases, w has a maximal Dyck factor with n letters 1, which has weight $(1+n)^\alpha$. This ensures that there is no other maximal Dyck factors. Hence, w is one of the configuration $(12)^n$ or $(21)^n$, which are exactly the mismatch-free ones and are both mapped onto $(1+n)^\alpha$ by ϕ_α . This proves the claimed result. \square

This bound is optimal up to a factor $2^{\alpha+1}$ since one computes $\phi_\alpha(1^n 2^n) = n^{\alpha+1}$.

Th. 1 then follows from Prop. 1 with $\varepsilon = \frac{\alpha(1-\alpha)}{2} n^{\alpha-2}$ and $f_0 = \phi_\alpha(w_0) \leq n^{\alpha+1}$:

$$\mathbb{E}(\widehat{T}(n)) \leq n^{\alpha+1} \times \frac{2n^{2-\alpha}}{\alpha(1-\alpha)} = \frac{2}{\alpha(1-\alpha)} n^3.$$

Here, the best value for α is $\frac{1}{2}$. Other values are useful only in the next section.

3 Bounding the weighted expected convergence time

Given a distribution μ over \mathcal{W}_n , one deduces from Prop. 1 a bound on the μ -weighted expected convergence time, with $\varepsilon = \frac{\alpha(1-\alpha)}{2} n^{\alpha-2}$ (Lem. 1):

$$\mathbb{E}(T_\mu(n)) \leq \frac{1}{\varepsilon} \sum_{w \in \mathcal{W}_n} \mu(w) \phi_\alpha(w).$$

Thus, to bound $\mathbb{E}(T_\mu(n))$, we just need to compute the μ -weighted value of ϕ_α . It is however generally easier to compute the μ -weighted value of the volume. The volume $V(w)$ of a configuration w is the area between its representation and the horizontal axis (for example, Fig. 2, both configurations have volume

15). One can then deduce bound on the μ -weighted value of ϕ_α from the simple following inequality, which holds for any $\alpha \in (0, 1)$ and for any configuration w :

$$\phi_\alpha(w) \leq V(w).$$

For the uniform distribution ν over \mathcal{W}_n , the ν -weighted value of the volume is known to be equivalent to $\frac{\sqrt{2\pi}}{8}n\sqrt{n}$ (see, *e.g.*, [3], p. 533). This allows to prove:

Theorem 2. *The uniformly-weighted expected convergence time T_ν satisfies*

$$\mathbb{E}(T_\nu(n)) = O(n^{5/2} \ln n).$$

Proof. For any $\alpha \in (0, 1)$, one has

$$\mathbb{E}(T_\nu(n)) \leq \frac{2n^{2-\alpha}}{\alpha(1-\alpha)} \times \sum_{w \in \mathcal{W}_n} \nu(w)V(w) = \frac{\sqrt{2\pi}}{4\alpha(1-\alpha)} n^{7/2-\alpha}.$$

One checks that this upper bound, seen as a function of α , is minimal in $1 - \frac{1}{\ln n}$. For such an α , one has $n^{7/2-\alpha} = n^{1/\ln n} n^{5/2} = e n^{5/2}$. This yields

$$\mathbb{E}(T_\nu(n)) \leq \frac{\sqrt{2\pi}}{4(1 - \frac{1}{\ln n}) \frac{1}{\ln n}} e n^{5/2} = \frac{\sqrt{2\pi}e}{4(1 - \frac{1}{\ln n})} n^{5/2} \ln n = O(n^{5/2} \ln n).$$

This shows the claimed result. \square

However, recall from the introduction that the context of this problem is the cooling of a quasicrystal from melt at very high temperature. In particular, the natural distributions over initial configurations should be the one of the melt. In the melt, the very high temperature imposes a very low threshold for ΔE , so that *all* the flips are equiprobable, with no flip being any more forbidden.

Hence, we define the *natural distribution* ν over \mathcal{W}_n as the stationary distribution of the process which, at each step, performs uniformly at random one of the possible flips, without any restriction (contrarily to the cooling process described in Sec. 1). It is a classical result of random walks on graph that this stationary distribution gives to each configuration a weight proportional to the number of flips which can be performed onto. In particular, the configurations of maximal weight are $(12)^n$ and $(21)^n$, *i.e.*, the stable configurations of the cooling process, while the configurations of minimal weight are $1^n 2^n$ and $2^n 1^n$, *i.e.*, the worst configurations of the cooling process. One can thus hope that the naturally-weighted expected convergence time T_ν is lower than the uniformly-weighted one. But we get only the same bound, because the naturally-weighted value of the volume turns out to be equivalent to its uniformly-weighted value:

Proposition 2. *The naturally-weighted value of the volume satisfies*

$$\mathbb{E}_\nu(V) := \sum_{w \in \mathcal{W}_n} \nu(w)V(w) \sim \frac{\sqrt{2\pi}}{8} n\sqrt{n}.$$

Proof. We first introduce a combinatorial enumerative series f describing only the positive configurations. Let $a_{n,p,q}$ denotes the number of positive configurations of size n , volume q and onto which p flips can be performed. We set:

$$f(z, u, v) = \sum_{(n,p,q) \in \mathbb{N}^3} a_{n,p,q} z^n u^p v^q.$$

Let us find a functional equation for f . Each positive configuration w can be written $w = 1u2v$, where both u and v are positive configurations, possibly empty. The volume satisfies $V(w) = V(u) + |u| + 1 + V(v)$. The number F of flips satisfies either $F(w) = F(u)$ if $|v| = 0$, or $F(w) = F(u) + 1 + F(v)$ otherwise. This leads, using symbolic methods described in [3], to

$$f(z, u, v) = (z^2uv + z^2vf(zv, u, v))(1 + uf(z, u, v)).$$

We introduce a similar combinatorial enumerative series g describing all the configurations. Each configuration w can be written $w = v_1 \cdots v_p$, where the v_i 's are alternating positive or negative configurations, hence described by f . Once again, symbolic methods described in [3] thus lead to the functional equation

$$g(z, u, v) = \frac{2z^2(vf(zv, u, v) + uv)}{1 - (z^2uv + z^2u^2v + z^2vf(zv, u, v) + z^2uv)}.$$

The naturally-weighted value of the volume can then be derived from g as follows:

$$\mathbb{E}_\nu(V) = \frac{[z^n] \frac{\partial^2}{\partial v \partial u} g(z, u, v) |_{u=1, v=1}}{[z^n] \frac{\partial}{\partial u} g(z, u, v) |_{u=1, v=1}},$$

where $[z^n]h$ denotes the coefficient of z^n in the series h . The claim then follows from classic computations in analytic combinatorics. \square

4 Perspectives

Th. 1 and 2 provide *upper* bounds on the expected convergence time of the cooling process, respectively in the worst case and in the average case (uniform and natural distributions). Experiments suggest, especially in the worst case, that these bounds are tight (see Fig. 9 and 10). This remains to be rigorously proven.

Moreover, this paper focuses only on the expectation of the convergence time, as the physically most significant moment. It would be worth to consider also higher moments, in order to obtain the limit distribution of the convergence time.

Last but not least, two-letter words correspond to tilings of dimension one and codimension one. But the physically interesting cases have dimension two or three and codimension at least two. For example, the celebrated *Penrose tiling* has dimension two and codimension three. The analysis of these cases seems harder but rewarding: experiments show much faster convergence rates (see [1]).




Fig. 9. The blue points represent the average over 10 simulations of the worst convergence time (on the x -axis, the length n of configurations). The pink curve $0.17n^3$ fits remarkably well with experimental values. The standard deviation of the worst convergence time is experimentally very small, so that only 10 simulations already provide a good approximation of its expectation.




Fig. 10. The blue points represent the average over 1000 simulations of the uniformly-weighted convergence time (on the x -axis, the length n of configurations). The pink curve $0.24n^{5/2} \ln n$ fits rather well with experimental values. The standard deviation of the uniformly-weighted convergence time is experimentally much greater than for the worst convergence time (indeed, each simulation starts from a configuration chosen uniformly at random), so that the number of simulations needed to get a good approximation is much greater too.

References

1. O. Bodini, Th. Fernique, D. Regnault, *Quasicrystallization by stochastic flips*, preprint, see <http://www.lif.univ-mrs.fr/~fernique/info/aperiodic.pdf>
2. N. Fatès, M. Morvan, N. Schabanel, É. Thierry, *Fully asynchronous behavior of double-quiescent elementary cellular automata*, *Theor. Comput. Sci.* **362** (2006), pp. 1–16.
3. Ph. Flajolet, R. Sedgewick, *Analytic combinatorics*, Cambridge Univ. Press, 2009.
4. C. Janot, *Quasicrystals: The State of the art*, Oxford University Press, 1992.