



Active Phase for Activated Random Walks on the Lattice in all Dimensions

Nicolas Forien, Alexandre Gaudillière

► To cite this version:

Nicolas Forien, Alexandre Gaudillière. Active Phase for Activated Random Walks on the Lattice in all Dimensions. Annales de l'Institut Henri Poincaré (B) Probabilités et Statistiques, inPress, 10.48550/arXiv.2203.02476 . hal-03611854

HAL Id: hal-03611854

<https://hal.science/hal-03611854>

Submitted on 27 Nov 2022

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.

Active Phase for Activated Random Walks on the Lattice in all Dimensions

Nicolas Forien^{1,a} and Alexandre Gaudillière^{2,b}

¹Aix Marseille Univ, CNRS, I2M, Marseille, France, Sapienza Università di Roma, Dipartimento di Matematica, Roma, Italy,

^anicolas.forien@uniroma1.it

²Aix Marseille Univ, CNRS, I2M, Marseille, France, ^balexandre.gaudilliere@math.cnrs.fr

Abstract. We show that the critical density of the Activated Random Walk model on \mathbb{Z}^d is strictly less than one when the sleep rate λ is small enough, and tends to 0 when $\lambda \rightarrow 0$, in any dimension $d \geq 1$. As far as we know, the result is new for $d = 2$.

We prove this by showing that, for high enough density and small enough sleep rate, the stabilization time of the model on the d -dimensional torus is exponentially large. To do so, we fix the set of sites where the particles eventually fall asleep, which reduces the problem to a simpler model with density one. Taking advantage of the Abelian property of the model, we show that the stabilization time stochastically dominates the escape time of a one-dimensional random walk with a negative drift. We then check that this slow phase for the finite volume dynamics implies the existence of an active phase on the infinite lattice.

Résumé. Nous démontrons que la densité critique du modèle des Marches Aléatoires Activées sur \mathbb{Z}^d est strictement inférieure à 1 quand le taux d'endormissement λ est suffisamment petit, et tend vers 0 quand $\lambda \rightarrow 0$, en toute dimension $d \geq 1$. À notre connaissance, le résultat est nouveau pour $d = 2$.

Nous obtenons ce résultat en prouvant que, pour une densité suffisamment élevée et un taux d'endormissement suffisamment petit, le temps de stabilisation du modèle sur le tore en dimension d est exponentiellement grand. Pour cela, nous fixons l'ensemble des sites sur lesquels les particules s'endorment, ce qui réduit le problème à un modèle plus simple avec densité 1. En utilisant la propriété d'Abélianité du modèle, nous montrons que le temps de stabilisation domine stochastiquement le temps d'atteinte de 0 pour une marche aléatoire en dimension 1 avec une dérive négative. Nous vérifions ensuite que cette phase de stabilisation lente pour la dynamique en volume fini implique l'existence d'une phase active sur le réseau infini.

MSC2020 subject classifications: Primary 60K35; secondary 82B26

Keywords: Activated random walks, phase transition, self-organized criticality

1. Introduction

The Activated Random Walk model is a model of interacting particles which attracted much interest since the seminal works of Rolla, Sidoravicius and Dickman [11, 25, 27]. This model emerged as a modification of another process suggested by R. Durrett in 1996 and known as the frog model [1, 2, 14, 16, 20], and it can also be seen as a particular case of some models for the spread of a rumor or an infection [21].

1.1. Informal definition of the model

Informally, the Activated Random Walk model is defined as follows. Let us consider a vertex-transitive and locally finite graph $G = (V, E)$, with, on each site of the graph, a certain number of particles, which may be either active or sleeping.

Active particles perform independent continuous-time random walks on the graph, with jump rate 1, meaning that each active particle is equipped with a Poisson clock of intensity 1 and, each time this clock rings, this particle jumps to a uniformly chosen neighbouring site of the graph.

In addition to this, when an active particle is alone on a site, it falls asleep with rate $\lambda \in [0, \infty]$. Thus, each particle is equipped with a second Poisson clock, with intensity λ : when this clock rings, if there is no other particle on the same site, the active particle falls asleep. Sleeping particles stop moving, until another active particle arrives on the same site. When this happens, the sleeping particle is waken up and turns back into the active state.

The initial configuration of particles is assumed to be distributed according to a translation-ergodic distribution, with no sleeping particles and an average density of active particles $\mu \in (0, \infty)$. Such a model is well defined if the graph is finite, and the above informal definition can also be made rigorous on \mathbb{Z}^d [26].

One motivation to study this model is its connection with the concept of self-organized criticality [5, 9, 13, 15, 18], which aims to describe some physical systems which present a critical-like behaviour without the need to finely tune their parameters to precise values (unlike in an ordinary phase transition, where the critical behaviour is only observed at the critical point).

The connection with self-organized criticality lies in the following variant of the model. Consider the model on a finite box where particles are killed when they jump out of the box. At each step, we add one active particle to a uniformly chosen site, and we let the box evolve until it reaches a stable configuration. It is conjectured that, when the size of this box tends to infinity, the average density of particles converges to the critical density of the model on \mathbb{Z}^d , as defined below (see for example [6] for a presentation of this conjecture). This property attracted some attention on the Activated Random Walk model, with the hope that this model would be more analytically tractable than other related models showing self-organized criticality, like the Stochastic Sandpile model [31] or the Manna model [24].

1.2. Some previous results on the phase transition for the model on the lattice

The Activated Random Walk model on \mathbb{Z}^d presents what is called an absorbing state phase transition. Namely, there is a critical curve in the plane (μ, λ) separating two different regimes: a fixating phase where, almost surely, the model eventually reaches an absorbing state, and an active phase where the model never stabilizes. Here, the absorbing states of the system, also called stable configurations, are the states where all the particles are sleeping.

We say that the system fixates if each site is visited only finitely many times. In this case, the configuration on any finite set eventually becomes stable, with only sleeping particles and empty sites. Otherwise, we say that the system stays active. The following key result of Rolla, Sidoravicius and Zindy establishes the existence of a universal phase transition for the model on \mathbb{Z}^d :

Theorem 1 ([28]). *For every dimension $d \geq 1$, and for every sleep rate $\lambda \in (0, \infty]$, there exists $\mu_c(\lambda)$ such that, for every translation-ergodic initial distribution with no sleeping particles and an average density of active particles μ , the Activated Random Walk model on \mathbb{Z}^d with sleep rate λ almost surely fixates if $\mu < \mu_c(\lambda)$, whereas it almost surely stays active if $\mu > \mu_c(\lambda)$.*

This threshold $\mu_c(\lambda)$ is called the critical density of the model. These last years, much effort has been directed towards finding upper and lower bounds on this critical density. As a first step, it was proved in [3, 30] that $\mu_c(\lambda) \leq 1$ for any value of the sleep rate, by showing that the density of particles is conserved over time.

Then, a natural question is: do we have $\mu_c(\lambda) \in (0, 1)$ for any finite value of the sleep rate λ in any dimension? It is known that $\mu_c(\lambda) \geq \lambda/(1 + \lambda)$ in wide generality [27, 31, 32], but this bound is sharp only when the jump distribution is totally asymmetric [8].

When the jumps are biased, it is known that $\mu_c(\lambda) < 1$ for any finite value of the sleep rate λ and that $\mu_c(\lambda) \rightarrow 0$ when $\lambda \rightarrow 0$, in any dimension $d \geq 1$ [29, 33].

For unbiased jumps, it was proved that $\mu_c(\lambda) \rightarrow 0$ when $\lambda \rightarrow 0$ in dimension 1 by [6] and in dimensions $d \geq 3$ in [32], and it was also shown that $\mu_c(\lambda) < 1$ for any finite value of the sleep rate λ , in dimension 1 [17] and in dimensions $d \geq 3$ [34]. In dimension 1, a more precise bound on the critical density showing that $\mu_c(\lambda) = O(\sqrt{\lambda})$ when $\lambda \rightarrow 0$ was obtained in [4].

1.3. Existence of an active phase in all dimensions

For the unbiased case on \mathbb{Z}^2 , the best bound available up to now was $\mu_c(\lambda) \leq 1$, and we did not know any choice of $\mu < 1$ and $\lambda > 0$ such that there is no fixation at (μ, λ) , which would imply that $\mu_c(\lambda) < 1$ for sufficiently low sleep rate.

Our main result stated below shows that $\mu_c(\lambda) \rightarrow 0$ when $\lambda \rightarrow 0$ in any dimension $d \geq 1$ which, in particular, answers the open problem of the existence of a non-trivial active phase on \mathbb{Z}^2 .

Theorem 2. *In any dimension $d \geq 1$, we have*

$$\lim_{\lambda \rightarrow 0} \mu_c(\lambda) = 0.$$

More precisely, there exists a constant $\kappa_d > 0$ such that, for every $\mu \in (0, 1)$ and $\lambda > 0$ which satisfy

$$(1) \quad \kappa_d (2d\lambda)^\mu < \mu^\mu (1 - \mu)^{1-\mu},$$

for every translation-ergodic initial distribution with no sleeping particles and an average density of active particles μ , the Activated Random Walk model on \mathbb{Z}^d with sleep rate λ almost surely stays active.

Thus, for any initial active particle density $\mu \in (0, 1)$, there exists $\lambda > 0$ such that, for every translation-ergodic initial distribution with no sleeping particles and an average density of active particles μ , the Activated Random Walk model on \mathbb{Z}^d with sleep rate λ almost surely stays active, and for all $\lambda > 0$ small enough, there exists $\mu < 1$ such that, for every translation-ergodic initial distribution with no sleeping particles and an average density of active particles μ , the Activated Random Walk model on \mathbb{Z}^d with sleep rate λ almost surely stays active.

Our result that $\mu_c(\lambda) \rightarrow 0$ when $\lambda \rightarrow 0$ was already known for the transient case $d \geq 3$ and the special case $d = 1$, but our proof works in any dimension $d \geq 1$.

Note that we only prove that $\mu_c(\lambda) < 1$ for $\lambda < 1/(2d\kappa_d)$, leaving open the conjecture that $\mu_c(\lambda) < 1$ for any finite value λ of the sleep rate. It might be possible to improve our technique to obtain this, but our efforts in this direction have not succeeded so far (see paragraph 1.7 for more comments on this).

1.4. Slow phase for the finite volume dynamics

In order to prove Theorem 2, we first study the model on the d -dimensional torus $\mathbb{Z}_n^d = (\mathbb{Z}/n\mathbb{Z})^d$. Recall that a configuration is said to be stable if all the particles are sleeping. Stable configurations are absorbing states of the model, meaning that if the dynamic reaches such a configuration, then it cannot change anymore. Hence, on a finite graph, if the initial number of particles is not greater than the number of sites, the system eventually fixates almost surely. Therefore, instead of the probability for the system to fixate, we are interested in the time it takes for the system to fixate, i.e., to reach a stable configuration.

Denoting this random time \mathcal{T}_n , and writing \mathbb{P}_μ^λ for the probability measure relative the Activated Random Walk model on the torus with sleep rate λ starting from an initial configuration with independent Poisson random numbers of active particles with mean μ on each site, we prove the following result:

Theorem 3. *In any dimension $d \geq 1$, for every $\mu \in (0, 1)$ and $\lambda > 0$ which satisfy the condition (1), the fixation time \mathcal{T}_n of the Activated Random Walk model on the torus $\mathbb{Z}_n^d = (\mathbb{Z}/n\mathbb{Z})^d$ satisfies*

$$(2) \quad \exists c > 0, \quad \forall n \geq 1, \quad \mathbb{P}_\mu^\lambda(\mathcal{T}_n < e^{cn^d}) < e^{-cn^d}.$$

Note that the result remains true if \mathcal{T}_n is replaced with the number of times a particle jumps or falls asleep (later on we will call this the number of topplings) until the system fixates (see Lemma 12).

To the best of our knowledge, this slow phase for finite systems was only studied up to now in the one-dimensional setting. In [7], it is shown that (2) holds on $\mathbb{Z}/n\mathbb{Z}$ as soon as $\mu > \mu_c(\lambda)$.

1.5. Slow phase on the torus implies active phase in infinite volume

Once Theorem 3 is proved, Theorem 2 is obtained through the following result, which relates the existence of a slow phase for the dynamics on the torus to the existence of an active phase for the model on the infinite lattice \mathbb{Z}^d .

Theorem 4. *In any dimension $d \geq 1$, for every sleep rate $\lambda \in (0, \infty)$ and every $\mu < \mu_c(\lambda)$, the property (2) about the fixation time \mathcal{T}_n of the model on the torus \mathbb{Z}_n^d does not hold.*

This Theorem is a very weak version of a conjecture mentioned in section 6 of [26], which predicts that for every density $\mu < \mu_c(\lambda)$, there should exist $\kappa > 0$ and $\delta > 0$ such that, for n large enough,

$$\mathbb{P}_\mu^\lambda(\mathcal{T}_n \leq \kappa n^d) \geq 1 - n^{-\delta}.$$

Combined with Theorem 3, the above Theorem entails that $\mu \geq \mu_c(\lambda)$ for every $\mu \in (0, 1)$ and $\lambda > 0$ satisfying the condition (1) and, this condition being a strict inequality, it implies that $\mu > \mu_c(\lambda)$ for every $\mu \in (0, 1)$ and $\lambda > 0$ satisfying this condition which, given Theorem 1, yields Theorem 2.

1.6. Key ingredients and organization of the article

In section 2, we recall the site-wise representation of the Activated Random Walk model and the properties which follow from it, including the monotonicity property and the Abelian property, which are key ingredients in our method. In section 3, we prove Theorem 3 about the fixation time on the torus, following the general strategy presented below. We conclude in section 4 with the proof of Theorem 4 which, together with Theorems 3 and 1, implies our main result, the existence of an active phase on \mathbb{Z}^d for all $d \geq 1$ (Theorem 2).

Let us sketch below the main elements of the proofs of our two main Theorems.

1.6.1. Sketch of the proof of Theorem 3

We first sketch the proof of Theorem 3, that is to say, for λ and μ satisfying the condition (1), we want to show that the stabilization time on the torus is exponentially large with overwhelming probability. All the intuitive facts claimed below are formally proved in section 3.

Fixing the initial number of particles: Taking μ' slightly lower than μ , but such that the condition (1) still holds for λ and μ' , we may reason conditionally on the fact that there are at least $\mu'n^d$ particles in the initial configuration (which happens with overwhelming probability). Then, using the monotonicity result of Lemma 1, it is enough to show that the stabilization time is exponential starting from any deterministic initial configuration with exactly $k = \lceil \mu'n^d \rceil$ active particles.

Fixing the set of sleeping sites: This is the first key ingredient of our proof. Once the number k of particles is fixed, the probability that these k particles fixate within a given time can be written as a sum over all the subsets $A \subset \mathbb{Z}_n^d$ with cardinality k of the probability that they fixate on A within a given time. Thus, we fix an arbitrary set A with cardinality k , and we try to evaluate the probability that the particles all fall asleep on A within a given time, that is to say, that we eventually reach the final stable configuration with one sleeping particle on each site of A . We will then have a combinatorial factor $\binom{n^d}{k}$ relative to the choice of A , but we will be able to compensate it by choosing λ small enough.

Forbidding particles to sleep elsewhere: Then, a second key idea is to modify the model by forbidding particles to fall asleep outside of A . Indeed, if the particles eventually fixate on A , it means that the sleep instructions which happened outside of A have not been useful, because all the particles which fell asleep outside of A have eventually been waken up by another particle. Therefore, this modification of the model only increases the probability that the particles fixate on A within a given time (see section 3.1 for a formal proof of this intuitive fact).

Taking advantage of the Abelian property of the site-wise representation: Instead of reasoning with the continuous time with Poisson clocks, we bound the number of topplings, i.e., the number of times a particle jumps or tries to fall asleep until the system stabilizes. If we show that this number of topplings is exponentially large, then it easily follows that the fixation time of the continuous-time model is also exponentially large.

To evaluate this number of topplings, we use the site-wise representation of the model (see section 2.1) and we take advantage of its Abelian property (Lemma 3), which says that the number of topplings and the final configuration do not depend on the order with which the sites are toppled. Thus, our goal is to find a toppling strategy which enables us to show that, with overwhelming probability, we can perform an exponentially large number of topplings before the system stabilizes on A (remember that, with the above modification of the model, the only stable configuration which can be reached is the configuration with one sleeping particle on each site of A).

Having one particle on each site of the settling set: In order to stabilize all the particles on the sites of A , as a first step we may let all the particles move until they arrive on an unoccupied site of A . Hence, the number of topplings necessary to fixate on A starting from any initial configuration with exactly $k = |A|$ particles stochastically dominates the number of topplings necessary to fixate on A starting with one active particle on each site of A (see section 3.2 for a formal proof of this). Thus, in what follows we may as well assume that we start with the initial configuration with one active particle on each site of A .

Toppling the sites in the right order: Then, our third key idea is to order the sites of A by writing $A = \{x_1, \dots, x_k\}$ in a way such that the distances $d(x_j, x_{j+1})$ are not too big (see section 3.3). Given this order, our toppling strategy is composed of a certain number of steps ending with exactly one particle (either active or sleeping) on each site of A .

At each step, we start by toppling a site $x_j \in A$ which contains an active particle and such that for all $i > j$, there is also an active particle on x_i , as drawn on figure 1 (at the first step, we topple x_1). With probability $\lambda/(1+\lambda)$, the particle falls asleep on x_j before moving. In this case, we proceed to the next step, where we topple the site x_{j+1} . Otherwise we let the particle walk with successive topplings until it comes back to its starting point x_j (remember that it is not allowed to fall asleep anywhere else, all the other sites of A being occupied). While walking from x_j back to x_j , this active particle may wake up some of the particles which were sleeping on previous sites of A . In particular, if it visits x_{j-1} before returning to x_j , we may choose, in the next step (i.e., once the particle has returned on x_j), to topple the site x_{j-1} , where the particle is necessarily awake. Otherwise, we topple again x_j in the next step. This procedure ends when, toppling the last particle x_k , it falls asleep before moving. At this point, some of the previous particles may still be active, but we obtain a lower bound for the number of topplings necessary to stabilize the configuration (see section 3.4 for the precise description of this toppling strategy).

During each step, the particle at site x_j can:




FIG 1. At each step of the toppling procedure used in the proof of Theorem 3, the particle at x_j either falls asleep (left), or visits the site x_{j-1} , waking the particle there if it was sleeping (middle), or goes back to x_j without visiting x_{j-1} (right). In this last case, the particle may wake up other particles at sites x_i with $i < j - 1$, but we do not care about this. Sleeping particles are denoted by 's', while active particles are denoted by '1'.

Reduction to a one-dimensional random walk with a negative drift: The toppling strategy described above is equivalent to a simplification of the model where, for each j , the particle at x_j is only allowed to wake up the particle at x_{j-1} , and we ignore collisions with other sleeping particles. Intuitively, this modification can only increase the probability that the particles fixate on A within a given time. In this modified model, at each step, either we add a sleeping particle, or we wake up the preceding particle, or nothing changes. Thus, the dynamics is reduced to a random walk on the set of integers $\{1, \dots, k+1\}$ because if j particles are sleeping, it can only be on the first j sites of A .

At this point, it only remains to check that this random walk (represented on figure 2) has a negative drift. To compute this drift, we combine two geometric estimates. The first estimate, proved in section 3.5, shows that, for any two distinct points $x, y \in \mathbb{Z}_n^d$, the probability that a symmetric random walk started at x reaches y before returning to x satisfies

$$P_x(T_y < T_x^+) \geq \frac{1}{2dd(x, y)},$$

where $d(x, y)$ is the graph distance between x and y . The second estimate, obtained in section 3.6, translates the constraint of the finiteness of the available space in the torus, saying that, if the order $A = \{x_1, \dots, x_k\}$ is well chosen, then

$$\sum_{j=1}^{k-1} \ln d(x_j, x_{j+1}) = O(n^d).$$

Combining these two estimates, we show in section 3.7 that the random walk has a negative drift, which makes the reaching time of $k+1$ exponentially small with high probability.

We then conclude in three last steps. In section 3.8, we prove that the number of topplings is exponentially large but starting from a deterministic initial configuration, checking that the combinatorial factor coming from the choice of the set A can be outweighed by the drift if λ is small enough. In section 3.9, we extend this result to an i.i.d. Poisson initial distribution, obtaining a version of Theorem 3 where the stabilization time \mathcal{T}_n is replaced with the number of topplings, before eventually proving this very Theorem in section 3.10.

Summary: To put it in a nutshell, by fixing a sleeping set $A \subset \mathbb{Z}_n^d$, we are able to reduce the dynamics to the case of density one (there is exactly one particle on each site of A , either active or sleeping), but on a different graph which is the trace graph on A . We then take advantage of the abelianity and monotonicity properties of the model to choose the order with which the sites are toppled, which allows us to control the fixation time of the model with a one-dimensional random walk. The geometric constraints on the set A then allow us to control the drift of this random walk, and therefore to show that it takes an exponential time to fixate the particles on A . Because this drift depends on λ , while the combinatorial factor $\binom{n^d}{k}$ does not, we are able to compensate it by choosing λ sufficiently small.




FIG 2. The one-dimensional random walk which governs the toppling procedure in the proof of Theorem 3.

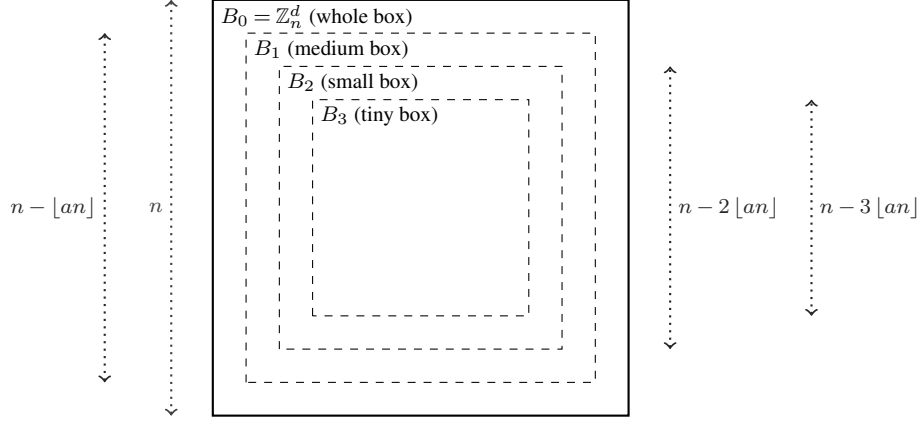


FIG 3. The four sizes of boxes used for the proof of Theorem 4.

1.6.2. Sketch of the proof of Theorem 4

We now present our strategy to prove Theorem 4, which says that if $\mu < \mu_c(\lambda)$, then the stabilization time on the torus \mathbb{Z}_n^d cannot be exponentially large in the sense of property (2). Thus, we fix $c > 0$ and we show that, for n large enough,

$$(3) \quad \mathbb{P}_\mu^\lambda(\mathcal{T}_n < e^{cn^d}) \geq e^{-cn^d}.$$

Again, we use the site-wise representation of the model, and we study the number of topplings rather than the continuous time. To stabilize the torus, we use the following toppling procedure (which is explained in more details in section 4.3). Denoting by $p: \mathbb{Z}^d \rightarrow \mathbb{Z}_n^d$ the canonical projection application on the torus and writing, for every $k \in \mathbb{N}$,

$$(4) \quad \Lambda_k = \left(-\frac{k}{2}, \frac{k}{2}\right] \cap \mathbb{Z}^d = \left\{ -\left\lfloor \frac{k-1}{2} \right\rfloor, \dots, \left\lfloor \frac{k}{2} \right\rfloor \right\}^d,$$

we define four sizes of sub-boxes of the torus, as represented on figure 3, by writing $B_j = p(\Lambda_{n-j[an]})$ for every $j \in \{0, 1, 2, 3\}$, with a certain parameter $a > 0$. We call $B_0 = \mathbb{Z}_n^d$ the whole box, B_1 the medium box, B_2 the small box, and B_3 the tiny box.

If the parameter a is chosen small enough, then the volume of the remaining space between the small box and the boundary of the whole box is not too big, and thus we can assume with a “reasonable” probabilistic cost that we start with no particles outside this small box. What we call a “reasonable” probabilistic cost is that of an event with possibly exponentially small probability, but much larger than e^{-cn^d} when c is fixed and a is chosen small enough (and n large enough). For example the probability $e^{-\mu|A|}$ of starting with an empty annulus $A = B_0 \setminus B_2$ is exponentially small in n^d , but for $c > 0$ fixed it is larger than $e^{-cn^d/24}$ for a small enough and n large enough. Thus, to obtain the bound (3) we will write the event $\{\mathcal{T}_n < e^{cn^d}\}$ as an intersection of several events which have a “reasonable” probabilistic cost. Starting with no particles outside the small box is the first of these events, and we call it step 0 of the procedure.

Then, during step 1, we stabilize the small box, ignoring particles once they jump out of the box. Because we assumed that $\mu < \mu_c(\lambda)$, we have an upper bound on the number of particles which jump out of this box during this step, which is due to Rolla and Tournier (see Lemma 4): for every $\varepsilon > 0$, with positive probability, uniformly in n , a proportion less than ε of the particles jump out of the small box.

After this, during step 2 we force these particles which jumped out of the small box to walk until they reach the boundary of the medium box, forbidding them to fall asleep before. Intuitively, this can only increase the stabilization

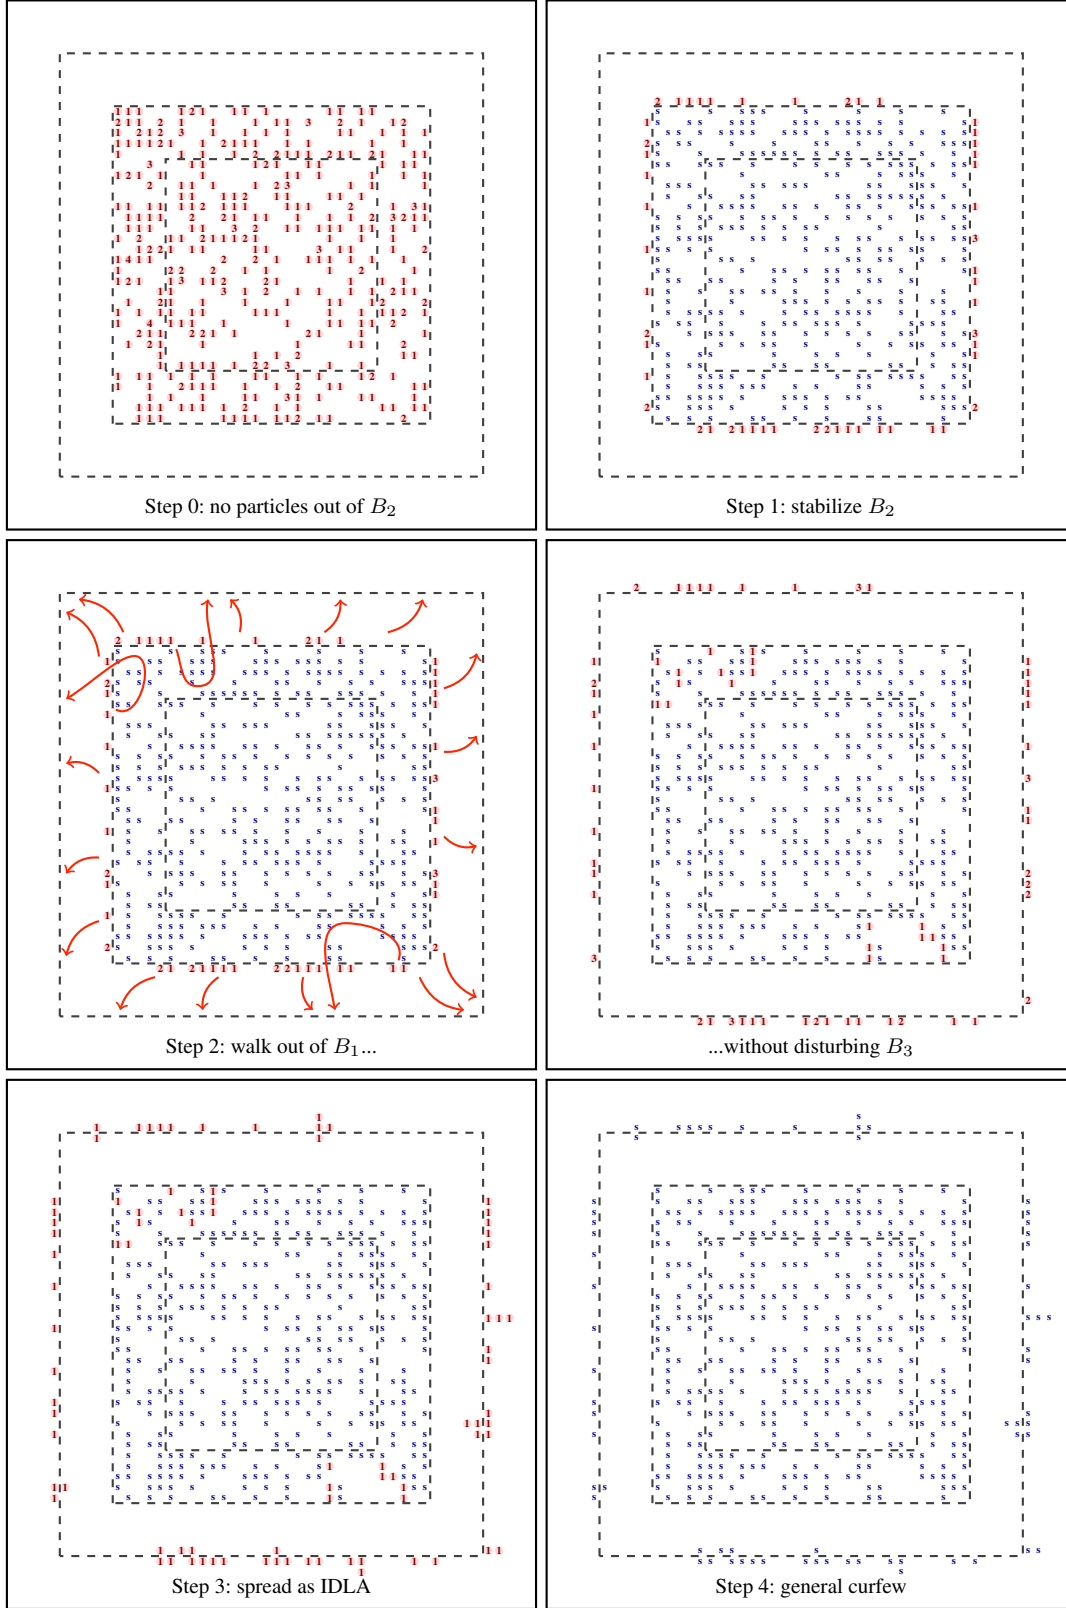


FIG 4. The toppling procedure used in the proof of Theorem 4.

time, and this argument can be made rigorous using the notion of acceptable topplings (see section 2.1). By symmetry, the probability for a random walk starting from the border of the small box to reach the border of the medium box before reaching the tiny box is at least one half. Hence, we can show that, with a “reasonable” probabilistic cost, we can move all these particles to the boundary of the medium box, without any particle visiting the tiny box, where all the particles are sleeping.

Then, during step 3 we consider these particles which reached the boundary of the medium box during step 2 and we let each of them walk until it reaches an empty site. Thus, these particles spread as an Internal Diffusion Limited Aggregation, which is a model corresponding to Activated Random Walks with an infinite sleep rate. By using a much weaker estimate than Jerison, Levine and Sheffield’s sharp bound to control its outer fluctuations (see section 4.2), we can show that, with high probability, these particles can spread without hitting neither the small box nor the boundary of the whole box.

The fourth and last step consists in forcing all the active particles to fall asleep, which again, represents a “reasonable” probabilistic cost.

If all the steps of this procedure succeed, we end up with a stable configuration, and the fact that no particle ever reaches the boundary of the whole box during this procedure allows us to upper bound the number of topplings performed (see section 4.1), eventually yielding the property (3) about the stabilization time on the torus.

1.7. Could our technique be extended for arbitrary sleep rates?

As explained above, we only prove that $\mu_c(\lambda) < 1$ for λ sufficiently small, and it is natural to wonder whether our technique could be improved to show that $\mu_c(\lambda) < 1$ for any finite value of the sleep rate λ .

Looking into the various steps of our proof detailed in section 1.6.1, we can see that the most brutal approximation we make is when we only allow a particle to wake up the particle immediately preceding it. Indeed, if μ is very close to 1, then we can expect that most of the distances $d(x_j, x_{j+1})$ will be of order 1, and thus when λ is large, the one-dimensional random walk will have a strong positive drift, making our technique fail.

If we allow every particle to wake up any other particle, and then topple the first awoken particle for the chosen order on the sites of A , then the dynamics becomes much more complicated and our reduction to a simple one-dimensional process does not work anymore.

A softer approach could be to allow the particle on x_j to wake up either no particle, or the particle on x_{j-1} , or the two particles on x_{j-1} and x_{j-2} , etc. Doing so, it would remain true that, after each step, the set of sleeping particles would remain of the form $\{1, \dots, j\}$, and thus we could control the procedure with a one-dimensional but long-range random walk. However, our preliminary computations to control the drift of this random walk seem to indicate that this simple extension of the present strategy is not sufficient to prove $\mu_c(\lambda) < 1$ for all $\lambda < \infty$.

2. Diaconis-Fulton representation and properties

We now describe a site-wise construction known as the Diaconis-Fulton representation, which was introduced in [10] and adapted to the framework of Activated Random Walks in [27]. This representation presents crucial properties of Abelianity and monotonicity, on which we rely heavily. We essentially reproduce here notations from [26].

2.1. The site-wise construction

Let $G = (V, E)$ be a vertex-transitive and locally finite graph (in what follows, we will only be interested in $G = \mathbb{Z}_n^d$ or $G = \mathbb{Z}^d$). A configuration of the Activated Random Walk model on G can be represented by a vector $\eta : V \rightarrow \mathbb{N}_s$, with $\mathbb{N}_s = \mathbb{N} \cup \{\mathfrak{s}\}$, where $\eta(x) = k \in \mathbb{N}$ means that there are k active particles on the site x , while $\eta(x) = \mathfrak{s}$ means that there is one sleeping particle on the site x . We equip the set \mathbb{N}_s with the total order $0 < \mathfrak{s} < 1 < 2 < \dots$ and, for $\eta, \eta' \in \mathbb{N}_s^V$, we write $\eta \leq \eta'$ whenever $\eta(x) \leq \eta'(x)$ for every $x \in V$. If $\eta \in \mathbb{N}_s^V$ and $A \subset V$, we introduce the following notation for the total number of particles on the sites of A in the configuration η , regardless of their state, active or sleeping:

$$|\eta|_A = \sum_{x \in A} |\eta(x)|,$$

with the convention that $|\mathfrak{s}| = 1$. If $A = V$, we simply write $|\eta| = |\eta|_V$.

We also define two operations on the configurations called toppling operations, which correspond to a particle moving or falling asleep. A site $x \in V$ is said to be stable in the configuration η if $\eta(x) \in \{0, \mathfrak{s}\}$, otherwise it is called unstable.

We say that a configuration η is stable in a subset $U \subset V$ is $\eta(x) \in \{0, \mathfrak{s}\}$ for all $x \in U$. If $\eta \in \mathbb{N}_{\mathfrak{s}}^V$ and $x \in V$ is unstable in η , we define

$$(5) \quad \tau_{x\mathfrak{s}}\eta : z \in V \mapsto \begin{cases} \mathfrak{s} & \text{if } z = x \text{ and } \eta(x) = 1, \\ \eta(z) & \text{otherwise,} \end{cases}$$

and for every $y \in V$ neighbour of x in the graph G , we define

$$(6) \quad \tau_{xy}\eta : z \in V \mapsto \begin{cases} \eta(x) - 1 & \text{if } z = x, \\ \eta(y) + 1 & \text{if } z = y, \\ \eta(z) & \text{otherwise,} \end{cases}$$

with the convention that $\mathfrak{s} + 1 = 2$. The key idea of the site-wise representation is that the toppling instructions are attached to the sites of the graph rather than to the particles. Let us explain this.

We consider a field of instructions $\tau = (\tau^{x,j})_{x \in V, j \in \mathbb{N}}$ where, for every $x \in V$ and $j \in \mathbb{N}$, the instruction $\tau^{x,j}$ is either a sleep instruction $\tau_{x\mathfrak{s}}$ or a jump instruction τ_{xy} to one of the neighbours y of x in the graph. This field τ will be taken random later on, but for the moment we consider a deterministic field of instructions. We also consider a function $h : V \rightarrow \mathbb{N}$ which we call odometer, and which will allow us to remember how many instructions have already been used at each site. Given a field of instructions $\tau = (\tau^{x,j})_{x \in V, j \in \mathbb{N}}$, a configuration $\eta \in \mathbb{N}_{\mathfrak{s}}^V$ and an odometer $h \in \mathbb{N}^V$, if $x \in V$ is unstable in η (i.e., if $\eta(x) \geq 1$), we define the toppling operation at x as

$$(7) \quad \Phi_x^\tau(\eta, h) = (\tau^{x, h(x)}\eta, h + \delta_x),$$

where $\delta_x(y) = 1$ if $x = y$ and 0 otherwise. We say that Φ_x^τ is legal for (η, h) if x is unstable in η .

One can think of this representation as taking, above each site of the graph, an infinite pile of instructions which are either sleep instructions or jump instructions to a neighbouring site. If, for every $x \in V$, $h(x)$ is the number of instructions already used at x , then the toppling operation Φ_x^τ looks at the first instruction in the pile above x which has not already been used, and applies this instruction to the configuration η .

In section 4 (where we prove Theorem 4), it will be convenient to be allowed to topple also sites where a particle is sleeping, which is not a legal toppling, but which we will call an acceptable toppling. The above definitions (6), (5) and (7) extend to the case when $\eta(x) = \mathfrak{s}$ by letting $\mathfrak{s} - 1 = 0$, and we say that Φ_x^τ is acceptable for (η, h) if $\eta(x) \geq \mathfrak{s}$. One can think of an illegal but acceptable toppling as a forced awakening of the sleeping particle at x , followed by a toppling of x .

Given a sequence of sites $\alpha = (x_1, \dots, x_l)$, we say that α is τ -legal (respectively, τ -acceptable) for (η, h) if, for every $j \in \{1, \dots, l\}$, the operation $\Phi_{x_j}^\tau$ is legal (resp., acceptable) for $\Phi_{x_{j-1}}^\tau \Phi_{x_{j-2}}^\tau \dots \Phi_{x_1}^\tau(\eta, h)$ and, in this case, we write

$$\Phi_\alpha^\tau(\eta, h) = \Phi_{x_l}^\tau \Phi_{x_{l-1}}^\tau \dots \Phi_{x_2}^\tau \Phi_{x_1}^\tau(\eta, h).$$

If $\Phi_\alpha^\tau(\eta, 0) = (\eta', h)$ where the resulting configuration η' is stable in $U \subset V$, we say that the toppling sequence α stabilizes η in U . For every sequence of sites $\alpha = (x_1, \dots, x_l)$, we define its odometer $m_\alpha \in \mathbb{N}^V$ as

$$m_\alpha = \sum_{j=1}^l \delta_{x_j},$$

which counts the number of times each site is toppled while applying Φ_α^τ . For every sequence of sites $\alpha = (x_1, \dots, x_l)$, and every subset $U \subset V$, we write $\alpha \subset U$ if $x_j \in U$ for every $j \in \{1, \dots, l\}$. If $\eta \in \mathbb{N}_{\mathfrak{s}}^V$ is a configuration, we define the odometer of η in U as

$$(8) \quad m_{U, \eta}^\tau = \sup_{\alpha \subset U, \alpha \text{ legal}} m_\alpha.$$

In the particular case $U = V$, we omit U and simply write m_η^τ . Let us also introduce the general notation

$$\|m\|_A = \sum_{x \in A} m(x)$$

for every $m \in (\mathbb{N} \cup \{+\infty\})^V$.

2.2. Monotonicity properties

A first notable property of the model is that adding particles to a configuration can only increase the number of topplings necessary to stabilize the configuration in a given set. More precisely, we have (see Lemma 2.5 of [26]):

Lemma 1. *If $\eta, \eta' \in \mathbb{N}_s^V$ are such that $\eta \leq \eta'$ and $U \subset V$, then*

$$m_{U,\eta}^\tau \leq m_{U,\eta'}^\tau.$$

This follows from the fact that for every toppling sequence $\alpha \subset U$, if α is τ -legal for η then it is also τ -legal for η' .

We also have the following comparison property between acceptable and legal topplings (we reproduce here Lemma 2.1 of [26]):

Lemma 2. *If α is an acceptable sequence of topplings that stabilizes η in a finite subset $U \subset V$, and $\beta \subset U$ is a legal sequence of topplings for η , then $m_\beta \leq m_\alpha$.*

This Lemma allows us to consider acceptable sequences of topplings when we are looking for upper bounds on $m_{U,\eta}^\tau$. We will use it in section 4.

2.3. The Abelian property

We now describe the Abelian property, which is a key feature of the Activated Random Walk model, which we use extensively. Let $U \subset V$ be a finite subset of the sites of the graph. The following Lemma corresponds to Lemma 2.4 of [26]:

Lemma 3 (Abelian property). *If α and β are both τ -legal toppling sequences for η that are contained in U and stabilize η in U , then $m_\alpha = m_\beta = m_{U,\eta}^\tau$ and $\Phi_\alpha^\tau(\eta, 0) = \Phi_\beta^\tau(\eta, 0)$.*

If $A \subset U$, the Abelian property tells us that

$$\|m_{U,\eta}^\tau\|_A = \sum_{x \in A} m_{U,\eta}^\tau(x)$$

is the number of topplings on the sites of A necessary to stabilize the configuration η in U using the instructions in τ , regardless of the order with which the instructions are used (note that this number may be infinite). In the particular case where $A = U$, we simply write $\|m_{U,\eta}^\tau\| = \|m_{U,\eta}^\tau\|_U$. If $\|m_{U,\eta}^\tau\| < \infty$, which means that η can be stabilized in U with a finite number of topplings, then the Abelian property ensures that the configuration reached when stabilizing η in U does not depend on the toppling sequence used, which allows us to write this final configuration $\sigma_{U,\eta}^\tau$ (which is defined only if $\|m_{U,\eta}^\tau\| < \infty$). In the case where V is finite and $U = V$, we simply write $\sigma_\eta^\tau = \sigma_{U,\eta}^\tau$. With these notations, if V is finite, then for any τ -legal sequence α which stabilizes η , we have $\Phi_\alpha^\tau(\eta, 0) = (\sigma_\eta^\tau, m_\eta^\tau)$.

2.4. Random instructions and dynamics

To recover the Activated Random Walk model described in paragraph 1.1, one needs to draw a random field of instructions with a certain probability distribution. Let \mathcal{P}^λ be a probability measure on the set of all possible fields of instructions, which is such that the instructions $(\tau^{x,j})_{x \in V, j \in \mathbb{N}}$ are independent, with, for every $x \in V$ and $j \in \mathbb{N}$,

$$(9) \quad \mathcal{P}^\lambda(\tau^{x,j} = \tau_{xs}) = \frac{\lambda}{1 + \lambda}$$

and, for every y neighbour of x in the graph,

$$(10) \quad \mathcal{P}^\lambda(\tau^{x,j} = \tau_{xy}) = \frac{1}{(1 + \lambda)D},$$

where D is the common degree of all the vertices in the graph G (which is assumed to be vertex-transitive). Given this random field $\tau = (\tau^{x,j})_{x \in V, j \in \mathbb{N}}$ distributed according to \mathcal{P}^λ , the Activated Random Walk model is a process $(\eta_t, h_t)_{t \geq 0}$, where $h_0 \equiv 0$ and the initial distribution η_0 is independent of τ and follows a translation-ergodic distribution on \mathbb{N}_s^V with finite density, and with, for each site $x \in V$, a transition rate of $(1 + \lambda)|\eta_t(x)|\mathbf{1}_{\eta_t(x) \geq 1}$ from (η_t, h_t) to $\Phi_x^\tau(\eta_t, h_t)$. It is clear that such a process exists on any finite graph (like the torus \mathbb{Z}_n^d), and we refer to [26] for a proof that such a process also exists on \mathbb{Z}^d .

Thanks to Theorem 1, to study the critical curve $\mu_c(\lambda)$, we may restrict ourselves to the case where the initial distribution η_0 is i.i.d. with a Poisson number of active particles with parameter μ on each site (and no sleeping particles). Let us denote by \mathcal{P}_μ the corresponding probability measure on the set \mathbb{N}_s^V . We write $\mathcal{P}_\mu^\lambda = \mathcal{P}_\mu \otimes \mathcal{P}^\lambda$ for the probability measure where the initial configuration η_0 and the field of instructions τ are independent and respectively distributed according to \mathcal{P}_μ and \mathcal{P}^λ , and we write \mathcal{E}_μ , \mathcal{E}^λ and \mathcal{E}_μ^λ for the corresponding expectations.

In what follows, we will only work with the model on finite graphs, using Lemma 4 below to relate the behaviour of the model on \mathbb{Z}^d to its behaviour in finite boxes. If the graph V is finite, the Activated Random Walk model with Poisson i.i.d. initial distribution can be constructed as follows. We draw the initial configuration η_0 and the field of toppling instructions τ distributed according to \mathcal{P}_μ^λ , and we draw an infinite collection of Poisson point processes $\theta_i \subset \mathbb{R}_+$ for $i \in \mathbb{N}$, independent of (η_0, τ) , these processes being independent and having intensity $1 + \lambda$. We draw an infinite number of clocks so that the collection $(\theta_i)_{i \in \mathbb{N}}$ is independent of the number of particles in η_0 but, in fact, we will only use the clocks θ_i for $i < |\eta_0|$. The dynamics is then defined as follows. At each time when one of these clocks rings, we choose a particle at random and, if this particle is active, we topple the site where it stands (in other words, we chose a site at random with the probability of choosing each site proportional to the number of particles on it, and we topple this site if it contains active particles).

2.5. A property of the fixating phase

To show that the slow phase on the torus \mathbb{Z}_n^d implies an active phase on the infinite lattice \mathbb{Z}^d , we rely on the following result, which is a rephrasing of Theorem 2.11 of [26] which gives a sufficient condition for activity on \mathbb{Z}^d and which was originally proved by [29]. Remember that Λ_n is the square box of \mathbb{Z}^d containing n^d sites and centered at 0, as defined by (4).

Lemma 4. *Let us write*

$$M_n = |\eta_0|_{\Lambda_n} - |\sigma_{\Lambda_n, \eta_0}^\tau|_{\Lambda_n},$$

which is the number of particles which jump out of the box $\Lambda_n \subset \mathbb{Z}^d$ when stabilizing the configuration η_0 in Λ_n via legal topplings, with particles being killed upon leaving Λ_n . For every $\lambda \in (0, \infty]$ and $\mu < \mu_c(\lambda)$, we have

$$\lim_{n \rightarrow \infty} \frac{\mathcal{E}_\mu^\lambda[M_n]}{|\Lambda_n|} = 0.$$

In other words, this Lemma tells us that if, when stabilizing an arbitrary large box, a uniformly positive fraction of particles jump out of the box with a uniformly positive probability, then $\mu \geq \mu_c(\lambda)$.

3. Proof of Theorem 3

We now turn to the study of the fixation time on the torus \mathbb{Z}_n^d , following the strategy presented in section 1.6.1. We start with a series of preliminary estimates in sections 3.1 to 3.7, before turning to the proof that the fixation time is exponentially large, first in terms of number of topplings from a deterministic initial configuration in section 3.8, then with a Poisson initial distribution in section 3.9, before the final proof of Theorem 3 in section 3.10.

3.1. Preventing sleep instructions outside a given set

As explained in paragraph 1.6.1, we first fix a set of sites $A \subset \mathbb{Z}_n^d$, and we aim to find an upper bound on the probability that the system eventually fixates with one sleeping particle on each site of A (i.e., that it reaches the final stable configuration $\sigma_\eta^\tau = \mathfrak{s} \mathbf{1}_A$) within a given time or, rather, within a given number of topplings. Since we are looking for a lower bound on this number of topplings, we may study the number of topplings on the sites of A , which gives a lower bound on the total number of topplings. Thus, our aim is to bound from above the probability

$$\mathcal{P}^\lambda \left(\{ \|m_\eta^\tau\|_A < M \} \cap \{ \sigma_\eta^\tau = \mathfrak{s} \mathbf{1}_A \} \right),$$

for a certain number M which will be chosen later, where we recall that $\|m_\eta^\tau\|_A$ denotes the number of toppling instructions performed on the sites of A during the complete stabilization of the configuration η in the torus using the instructions in τ (do not get mixed with $\|m_{A, \eta}^\tau\|$, which counts the number of topplings necessary to stabilize the configuration on A , i.e., until there are no more active particles on the sites of A).

The idea is now to modify the model by not allowing particles to sleep outside of A which, as we will show in Lemma 5 below, can only increase the above probability. Thus, we define a new distribution $\mathcal{P}^{\lambda,A}$ on the set of the fields of toppling instructions, which is such that we do not draw anymore sleep instructions on the sites outside of A . Namely, for every $x \in A$ and $j \in \mathbb{N}$, we have the same distribution for the instruction $\tau^{x,j}$ given by (9) and (10), but whenever $x \notin A$, we set

$$\mathcal{P}^{\lambda,A}(\tau^{x,j} = \tau_{xy}) = \frac{1}{2d}$$

for every y neighbour of x in \mathbb{Z}_n^d , and $\mathcal{P}^{\lambda,A}(\tau^{x,j} = \tau_{xs}) = 0$, and different instructions remain of course independent. We now prove the following:

Lemma 5. *For each initial configuration $\eta : \mathbb{Z}_n^d \rightarrow \mathbb{N}_s$, every $A \subset \mathbb{Z}_n^d$ and every $M \in \mathbb{N}$, we have the comparison*

$$\mathcal{P}^\lambda\left(\left\{\|m_\eta^\tau\|_A < M\right\} \cap \left\{\sigma_\eta^\tau = \mathbf{s} \mathbf{1}_A\right\}\right) \leq \mathcal{P}^{\lambda,A}\left(\|m_\eta^\tau\|_A < M\right).$$

Proof. Let $\eta : \mathbb{Z}_n^d \rightarrow \mathbb{N}_s$, let $A \subset \mathbb{Z}_n^d$ and $M \in \mathbb{N}$. Let us consider the natural coupling between \mathcal{P}^λ and $\mathcal{P}^{\lambda,A}$ which consists in drawing τ' according to $\mathcal{P}^{\lambda,A}$ and constructing τ by inserting, between each jump instruction of τ' on the sites $x \notin A$, a number of sleep instructions which follows a geometric distribution with parameter $\lambda/(1+\lambda)$, re-indexing the obtained stacks of instructions.

We now assume that $\|m_\eta^\tau\|_A < M$ and $\sigma_\eta^\tau = \mathbf{s} \mathbf{1}_A$, and we show that it implies that $\|m_{\eta'}^{\tau'}\|_A < M$. Since $\sigma_\eta^\tau = \mathbf{s} \mathbf{1}_A$, we know that there exists a τ -legal toppling sequence $\alpha = (x_1, \dots, x_l)$ such that $\Phi_\alpha^\tau(\eta, 0) = (\mathbf{s} \mathbf{1}_A, m_\eta^\tau)$. We now consider the toppling sequence α' which is obtained from α by removing the topplings which correspond to sleep instructions out of A .

For each $i \in \{0, \dots, l\}$, let us write $\alpha_i = (x_1, \dots, x_i)$ (with the convention that α_0 is an empty toppling sequence), and let us write α'_i for the sequence obtained from α_i by removing the sleep instructions out of A (so that α'_i may be shorter than α_i), and let us consider the configurations η_i and η'_i defined by $\Phi_{\alpha_i}^\tau(\eta, 0) = (\eta_i, m_{\alpha_i})$ and $\Phi_{\alpha'_i}^{\tau'}(\eta, 0) = (\eta'_i, m_{\alpha'_i})$, if $\Phi_{\alpha'_i}^{\tau'}$ is legal for $(\eta, 0)$.

Let us show by induction that, for every $i \in \{0, \dots, l\}$,

$$(11) \quad \Phi_{\alpha'_i}^{\tau'} \text{ is legal for } (\eta, 0), \quad \eta_i \leq \eta'_i \quad \text{and} \quad \forall x \in A \quad \eta_i(x) = \eta'_i(x).$$

This is true for $i = 0$ since $\eta_0 = \eta'_0 = \eta$. Assume that $i \in \{0, \dots, l-1\}$ is such that the induction hypothesis (11) holds. Since x_{i+1} is by definition unstable in η_i (because α is a τ -legal toppling sequence) and $\eta'_i \geq \eta_i$, we know that x_{i+1} is also unstable in η'_i , therefore $\Phi_{\alpha'_{i+1}}^{\tau'}$ is legal for $(\eta, 0)$.

Let $j = m_{\alpha_i}(x_{i+1})$. If $x_{i+1} \in A$ or $\tau^{x_{i+1},j}$ is a jump instruction, then both η_{i+1} and η'_{i+1} are obtained from η_i and η'_i by performing the toppling operation $\tau^{x_{i+1},j}$, thus (11) remains satisfied at rank $i+1$. Otherwise, if $x_{i+1} \notin A$ and $\tau^{x_{i+1},j}$ is a sleep instruction, then we have $\eta'_{i+1} = \eta'_i$ and $\eta_{i+1} = \tau_{x_{i+1}\mathbf{s}}\eta_i$, whence $\eta_{i+1} \leq \eta'_i$ and $\eta_{i+1}(x) = \eta_i(x)$ for all $x \in A$. Thus, we also still have (11) at rank $i+1$.

Therefore, it follows by induction that for all $x \in A$, $\eta'_l(x) = \eta_l(x) = \mathbf{s}$. The number of particles being conserved, we deduce that $\eta'_l = \mathbf{s} \mathbf{1}_A$, which means in particular that α' is a τ' -legal toppling sequence which stabilizes η . By the Abelian property, this implies that $m_\eta^{\tau'} = m_{\alpha'}^{\tau'}$. Besides, by definition of α' , we have

$$\|m_{\alpha'}^{\tau'}\|_A = \|m_\alpha\|_A = \|m_\eta^\tau\|_A < M,$$

whence $\|m_\eta^{\tau'}\|_A < M$, which concludes the proof, since τ' is distributed according to $\mathcal{P}^{\lambda,A}$ and τ is distributed according to \mathcal{P}^λ . \square

3.2. Starting with one active particle on each sleeping site

We now show that, if we look for a lower bound on $\|m_\eta^\tau\|_A$, the worst initial configuration is the configuration $\eta = \mathbf{1}_A$, where there is exactly one active particle on each site of A . Indeed, we prove:

Lemma 6. *For any subset $A \subset \mathbb{Z}_n^d$ and for any configuration $\eta : \mathbb{Z}_n^d \rightarrow \mathbb{N}$ with $|\eta| = |A|$ active particles, under the distribution $\mathcal{P}^{\lambda,A}$ defined in paragraph 3.1, the number of topplings performed on A to stabilize the configuration η in \mathbb{Z}_n^d , namely $\|m_\eta^\tau\|_A$, stochastically dominates $\|m_{\mathbf{1}_A}^\tau\|_A$, the number of topplings performed on A to stabilize the configuration $\mathbf{1}_A$ in \mathbb{Z}_n^d .*

Proof. Let $A \subset \mathbb{Z}_n^d$ and $\eta : \mathbb{Z}_n^d \rightarrow \mathbb{N}$ such that $|\eta| = |A|$. We have to prove that, for all $M \in \mathbb{N}$,

$$\mathcal{P}^{\lambda, A} \left(\|m_\eta^\tau\|_A < M \right) \leq \mathcal{P}^{\lambda, A} \left(\|m_{\mathbf{1}_A}^\tau\|_A < M \right).$$

Thus, we fix $M \in \mathbb{N}$. By the definition (8) of the odometer m_η^τ , we have (since we now deal with only one field of toppling instructions τ , we simply write “legal” instead of “ τ -legal”)

$$\|m_\eta^\tau\|_A = \sup_{\alpha \text{ legal for } (\eta, 0)} \|m_\alpha\|_A,$$

whence

$$(12) \quad \mathcal{P}^{\lambda, A} \left(\|m_\eta^\tau\|_A < M \right) = \mathcal{P}^{\lambda, A} \left(\forall \alpha \text{ legal for } (\eta, 0), \|m_\alpha\|_A < M \right).$$

Let τ be distributed according to $\mathcal{P}^{\lambda, A}$, and let us consider the following toppling procedure: as long as there is at least an unstable site which is not in A , or a site of A containing more than one particle, we topple one of these sites. Note that, doing so, we never topple a site of A with a single particle, so that no particle can fall asleep during this procedure (remember that there are no sleep instructions out of A because τ is drawn according to $\mathcal{P}^{\lambda, A}$). We continue this procedure until there are no more sites of A with at least two particles and no more occupied site out of A , and we denote by \mathcal{F} the event that this procedure terminates in a finite number of topplings. Since we took $|\eta| = |A|$, we have $\mathcal{P}^{\lambda, A}(\mathcal{F}) = 1$. If this event \mathcal{F} is realized, the procedure yields a finite legal sequence of topplings β such that $\Phi_\beta^\tau(\eta, 0) = (\mathbf{1}_A, m_\beta)$, where $\mathbf{1}_A$ is the configuration where there is exactly one particle on each site of A . In this case, for any sequence of topplings γ which is τ -legal for $(\mathbf{1}_A, m_\beta)$, the concatenation of β and γ yields a τ -legal sequence of topplings for $(\eta, 0)$, whence

$$\begin{aligned} \mathcal{P}^{\lambda, A} \left(\forall \alpha \text{ legal for } (\eta, 0), \|m_\alpha\|_A < M \right) &\leq \mathcal{P}^{\lambda, A} \left(\forall \gamma \text{ legal for } (\mathbf{1}_A, m_\beta), \|m_\beta + m_\gamma\|_A < M \mid \mathcal{F} \right) \\ &\leq \mathcal{P}^{\lambda, A} \left(\forall \gamma \text{ legal for } (\mathbf{1}_A, m_\beta), \|m_\gamma\|_A < M \mid \mathcal{F} \right) \\ &= \mathcal{P}^{\lambda, A} \left(\forall \gamma \text{ legal for } (\mathbf{1}_A, 0), \|m_\gamma\|_A < M \right). \end{aligned}$$

because, conditioned on \mathcal{F} , the remaining toppling instructions in τ are still distributed according to $\mathcal{P}^{\lambda, A}$ and are independent of the instructions revealed to construct β (see Proposition 4 in [22] for a formal statement of this strong Markov property). Combining this with equation (12), we obtain the claimed result. \square

Therefore, we proved that the configuration $\mathbf{1}_A$ is the worst initial configuration for the lower bound we are looking for. Thus, in the sequel we forget about this initial configuration η , and we reason as if we started with this new initial configuration $\mathbf{1}_A$.

3.3. Ordering the sleeping sites

We now start from this configuration $\mathbf{1}_A$ with exactly one active particle on each site of A , and we try to construct a random legal sequence of topplings α such that $\|m_\alpha\|_A$ is exponentially large with overwhelming probability (under $\mathcal{P}^{\lambda, A}$).

Because particles are not allowed to fall asleep neither out of A nor on sites already occupied by another particle, we may consider a dynamics where, at each step, we choose a particle of A , we topple it and, in case it chooses to move, we let it walk on \mathbb{Z}_n^d with successive topplings (in fact, we topple sites rather than particles) until it goes back to its starting point. The sleeping particles met along this path are awoken by this particle. Thus, the state space simplifies to $\mathcal{P}(A) = \{S \subset A\}$, because, after each step (if we call step a loop from one site $x \in A$ back to x), each site of A is either occupied by a sleeping particle or by an active one.

To choose which particle we topple at each step, we define an ordering of the sites of A . The aim is that each site should not be too far apart from the preceding site. We define recursively an ordering $A = \{x_1, \dots, x_k\}$.

First, we choose $x_1 \in A$ arbitrarily. Then, if $j \in \{1, \dots, k-1\}$ is such that x_1, \dots, x_j are constructed, we define the set $A_j = A \setminus \{x_1, \dots, x_j\}$ and we choose $x_{j+1} \in A_j$ such that

$$d(x_j, x_{j+1}) = \min_{x \in A_j} d(x_j, x),$$

where d denotes the graph distance on \mathbb{Z}_n^d . In other words, at each step we choose among the remaining points the closest (or one of the closest) to the last point added. For each $j \in \{1, \dots, k-1\}$, we write $l_j = d(x_j, x_{j+1})$. Doing so, we obtain an ordering $A = \{x_1, \dots, x_k\}$ which is such that

$$(13) \quad \forall j \in \{1, \dots, k-1\} \quad l_j = d(x_j, x_{j+1}) = \min_{i>j} d(x_j, x_i).$$

3.4. The toppling dynamic

Using a toppling procedure based on the order on the sites of A , we now prove a comparison between the Activated Random Walk model and a one-dimensional process.

Lemma 7. *Consider $A \subset \mathbb{Z}_n^d$, equipped with an order $A = \{x_1, \dots, x_k\}$ where $k = |A|$. Let us consider a Markov chain with state space $\{1, \dots, k+1\}$ and transition probabilities given by*

$$\begin{cases} \forall j \in \{1, \dots, k\} & p(j, j+1) = \frac{\lambda}{1+\lambda}, \\ p(1, 1) = \frac{1}{1+\lambda}, & p(k+1, k+1) = 1, \\ \forall j \in \{2, \dots, k\} & p(j, j-1) = \frac{1}{1+\lambda} P_{x_j}(T_{x_{j-1}} < T_{x_j}^+), \\ \forall j \in \{2, \dots, k\} & p(j, j) = \frac{1}{1+\lambda} P_{x_j}(T_{x_j}^+ < T_{x_{j-1}}), \end{cases}$$

where P_{x_j} denotes the probability with respect to the symmetric random walk on the torus \mathbb{Z}_n^d started from x_j , $T_{x_{j-1}}$ is the first hitting time of x_{j-1} and $T_{x_j}^+$ is the first time the random walk comes back to its starting point x_j . We write \mathbf{P}_j^A for the probability measure of this Markov chain started from j , and we denote by T_{k+1} its first hitting time of $k+1$. Then, the law of $\|m_{\mathcal{P}^{\lambda,A}}^\tau\|_A$ under $\mathcal{P}^{\lambda,A}$ stochastically dominates the law of T_{k+1} under \mathbf{P}_1^A .

Proof. Let $A = \{x_1, \dots, x_k\} \subset \mathbb{Z}_n^d$, with $k = |A|$. Let us consider a random field of toppling instructions τ distributed according to $\mathcal{P}^{\lambda,A}$. We now specify a procedure which will almost surely yield a sequence of topplings α legal for $(\mathbf{1}_A, 0)$, with a dynamics coupled with the aforementioned Markov chain in a way such that $\|m_\alpha\|_A \geq T_{k+1}$.

This procedure is defined as follows. Remember that we start with exactly one active particle on each site of A . Our toppling procedure consists of a certain number of steps, indexed by $t \in \mathbb{N}$, with a counter $J(t)$ which indicates which site of A we have to topple at each step.

We start with $J(0) = 1$, and the procedure will terminate whenever we reach $J(t) = k+1$. We will ensure that, before each step t , there is exactly one particle on each site of A and all the particles located on the sites x_i with $i \geq J(t)$ are active (while the particles on the sites x_i with $i < J(t)$ may be active or sleeping).

At each time step t , if $J(t) \leq k$, we topple the particle at site $x_{J(t)}$ (see figure 1). With probability $\lambda/(1+\lambda)$, this particle falls asleep before walking. In this case, we set $J(t+1) = J(t) + 1$, and we proceed with the next step. It remains true that, at time $t+1$, all the particles on the sites x_i with $i \geq J(t+1)$ are active.

Otherwise, with probability $1/(1+\lambda)$, the particle at site $x_{J(t)}$ chooses to move to one of the neighbouring sites on the torus. We then let it walk with successive topplings until it comes back to its starting point $x_{J(t)}$ (this is legal because it is not allowed to fall asleep outside of A and all the other sites of A are already occupied by other particles). This loop almost surely terminates in a finite number of topplings, and then we leave the particle awake at site $x_{J(t)}$. If $J(t) \geq 2$, along this path, with probability $P_{x_{J(t)}}(T_{x_{J(t)-1}} < T_{x_{J(t)}}^+)$, this active particle meets the particle at site $x_{J(t)-1}$, which becomes active if it was sleeping (it may be sleeping or active before, but at least after this visit we know that it is active). Thus, we set $J(t+1) = J(t) - 1$, meaning that, at the next step, we will topple the particle at site $x_{J(t)-1}$. This operation also ensures that at time $t+1$, there are active particles on the sites x_i for $i \geq J(t+1)$.

Note that the path may visit other sites x_i with $i \notin \{J(t), J(t)-1\}$, but this does not matter, since the particles on the sites $i > J(t)$ are already awake and we do not care about the state of the particles on the sites $i < J(t)-1$. As represented on figure 1, some of the latter particles may be active and some may be sleeping.

During this procedure, the process $J(t)$ evolves exactly like the Markov chain described above (at least almost surely, because the procedure can also fail with probability 0 if at some step the random walk on the torus never comes back to its starting point). Thus, the number T of steps before the procedure stops is distributed as the reaching time T_{k+1} under \mathbf{P}_1^A . What's more, at each step of the procedure, we topple at least one site of A . Therefore, the procedure almost surely yields a finite sequence of topplings α which is legal for $(\mathbf{1}_A, 0)$ and which is such that $\|m_\alpha\|_A \geq T$, which implies that $\|m_{\mathbf{1}_A}^\tau\|_A \geq T$. Therefore, we constructed a coupling between $\mathcal{P}^{\lambda,A}$ and \mathbf{P}_1^A which is such that $\|m_{\mathbf{1}_A}^\tau\|_A \geq T_{k+1}$ almost surely, which proves the claimed stochastic domination. \square

3.5. An estimate about random walks on the torus

We now prove the following geometric estimate about random walks on the torus:

Lemma 8. Denoting by P_x the probability measure relative to the symmetric random walk on the torus $\mathbb{Z}_n^d = (\mathbb{Z}/n\mathbb{Z})^d$ started at x , writing T_y for its first hitting time of y and T_x^+ its first return time to x and d for the graph distance on \mathbb{Z}_n^d , we have

$$\forall n \geq 1 \quad \forall x, y \in \mathbb{Z}_n^d \quad x \neq y \quad \Rightarrow \quad P_x(T_y < T_x^+) \geq \frac{1}{(2d)d(x, y)}.$$

Proof. The result is a straightforward consequence of the theory connecting random walks and electric networks (see for example [12] or [23]). Indeed, the above probability is related to the effective resistance $R(x, y)$ between x and y on the torus where each edge has a resistance equal to $2d$, through the relation

$$R(x, y) = \frac{1}{P_x(T_y < T_x^+)}.$$

According to Rayleigh's monotonicity law, this resistance can only increase if we remove edges from the graph. Thus, removing all the edges except the ones belonging to one of the shortest paths from x to y , we obtain that

$$R(x, y) \leq 2d d(x, y),$$

whence the claimed inequality. \square

3.6. A geometric estimate

We now prove the following property of the ordering that we defined in paragraph 3.3:

Lemma 9. There exists a universal constant $\kappa_d > 0$, which depends only on the dimension d , such that, for all $n \geq 1$ and for every subset $A \subset \mathbb{Z}_n^d$, the distances $l_1, \dots, l_{|A|-1}$ defined in section 3.3 satisfy

$$\prod_{j=1}^{|A|-1} l_j \leq (\kappa_d)^{n^d}.$$

Proof. Let $A \subset \mathbb{Z}_n^d$. We consider the order $A = \{x_1, \dots, x_k\}$ with $k = |A|$ and the distances l_1, \dots, l_{k-1} , as defined in paragraph 3.3. For every $L \in \mathbb{N} \setminus \{0\}$, let us consider the set

$$\mathcal{A}_L = \left\{ j \in \{1, \dots, k-1\} : l_j \geq L \right\}.$$

If $L > dn/2$, then we have $\mathcal{A}_L = \emptyset$, because there is no pair of points in \mathbb{Z}_n^d separated by a distance L . Hence, we assume that $L \leq dn/2$. This set \mathcal{A}_L has the property that any two different points with indices in \mathcal{A}_L are always separated by a distance at least L , because if $i, j \in \mathcal{A}_L$ with for example $i > j$, then the property (13) of the ordering ensures that $d(x_i, x_j) \geq l(x_j) \geq L$. Defining the open balls for the graph distance d on the torus \mathbb{Z}_n^d by

$$\forall x \in \mathbb{Z}_n^d \quad \forall r \in \mathbb{N} \quad B_{\mathbb{Z}_n^d}(x, r) = \{y \in \mathbb{Z}_n^d : d(x, y) < r\},$$

we know that the open balls $B_{\mathbb{Z}_n^d}(x, \lceil L/2 \rceil)$ for $x \in \mathcal{A}_L$ are disjoint, whence

$$(14) \quad |\mathcal{A}_L| \leq \frac{|\mathbb{Z}_n^d|}{|B_{\mathbb{Z}_n^d}(0, \lceil L/2 \rceil)|} \leq \frac{|\mathbb{Z}_n^d|}{|B_{\mathbb{Z}_n^d}(0, \lceil L/(2d) \rceil)|}.$$

Now note that, since $L/(2d) \leq n/4$, we have

$$\left| B_{\mathbb{Z}_n^d} \left(0, \left\lceil \frac{L}{2d} \right\rceil \right) \right| = \left| B_1 \left(0, \left\lceil \frac{L}{2d} \right\rceil \right) \right| \geq \left| B_\infty \left(0, \left\lceil \frac{\lceil L/(2d) \rceil}{d} \right\rceil \right) \right|,$$

where B_1 denotes the open ball on \mathbb{Z}^d for the graph distance on \mathbb{Z}^d , which corresponds to the usual L_1 -norm $\|\cdot\|_1$, while B_∞ denotes the open ball on \mathbb{Z}^d for the supremum norm $\|\cdot\|_\infty$. Thus, we can write

$$\left| B_{\mathbb{Z}_n^d} \left(0, \left\lceil \frac{L}{2d} \right\rceil \right) \right| \geq \left| B_\infty \left(0, \left\lceil \frac{L}{2d^2} \right\rceil \right) \right| = \left(2 \left\lceil \frac{L}{2d^2} \right\rceil - 1 \right)^d \geq \left(\frac{L}{2d^2} \right)^d.$$

Plugging this into (14), we obtain, for every $L \leq dn/2$,

$$|\mathcal{A}_L| \leq \left(\frac{2d^2 n}{L} \right)^d.$$

The above is also true when $L > dn/2$, since in this case we noted that $\mathcal{A}_L = \emptyset$. Therefore, we may write

$$\sum_{j=1}^{k-1} \ln(l_j) = \sum_{m=0}^{+\infty} (|\mathcal{A}_{2^m}| - |\mathcal{A}_{2^{m+1}}|) \ln(2^{m+1}) = \sum_{m=0}^{+\infty} |\mathcal{A}_{2^m}| \ln 2 \leq n^d \ln \kappa_d,$$

where the constant κ_d is given by

$$\kappa_d = 2 \exp \left((2d^2)^d \sum_{m=0}^{+\infty} \frac{m+1}{2^{md}} \right) = 2 \exp \left(\frac{(2d^2)^d}{1-2^{-d}} \right),$$

which is indeed finite for any $d \geq 1$. □

3.7. Study of the absorption time of the one-dimensional Markov chain

We now study the law of the reaching time T_{k+1} of the absorbing state $k+1$ for the Markov chain that we defined in paragraph 3.4, and we prove:

Lemma 10. *Let $\lambda > 0$, $\mu \in (0, 1)$ and $A \subset \mathbb{Z}_n^d$ with $|A| = \lceil \mu n^d \rceil$. We consider the order $A = \{x_1, \dots, x_k\}$ defined in paragraph 3.3, with $k = |A|$, and we consider the Markov chain described in Lemma 7, with law \mathbf{P}_1^A . Then, for any $M \in \mathbb{N}$, the absorption time T_{k+1} of this Markov chain satisfies*

$$\mathbf{P}_1^A(T_{k+1} \leq M) \leq \frac{M}{1 \wedge (2d\lambda)} (\kappa_d (2d\lambda)^\mu)^{n^d}.$$

Proof. Let $\lambda > 0$, $\mu \in (0, 1)$, $A = \{x_1, \dots, x_k\} \subset \mathbb{Z}_n^d$ with $|A| = k = \lceil \mu n^d \rceil$, as constructed in paragraph 3.3, and let $M \in \mathbb{N}$. We can modify the Markov chain defined in Lemma 7 by letting

$$p(k+1, k) = p(k, k+1) = \frac{\lambda}{1+\lambda} \quad \text{and} \quad p(k+1, k+1) = \frac{1}{1+\lambda},$$

since this does not change the reaching time of $k+1$, which we are interested in. With this modification, the Markov chain becomes reversible with respect to the measure ν defined on $\{1, \dots, k+1\}$ by

$$\forall j \in \{1, \dots, k+1\} \quad \nu(j) = \prod_{1 \leq i < j} \frac{p(i, i+1)}{p(i+1, i)},$$

with the convention that $\nu(1) = 1$. This reversibility allows us to write

$$\begin{aligned} \mathbf{P}_1^A(T_{k+1} \leq M) &= \sum_{t=1}^M \mathbf{P}_1^A(T_{k+1} = t) = \sum_{t=1}^M \sum_{1=j_0, j_1, \dots, j_t=k+1} \prod_{i=1}^t p(j_{i-1}, j_i) \\ &= \sum_{t=1}^M \sum_{1=j_0, j_1, \dots, j_t=k+1} \prod_{i=1}^t \frac{\nu(j_i)}{\nu(j_{i-1})} p(j_i, j_{i-1}) = \sum_{t=1}^M \nu(k+1) \mathbf{P}_{k+1}^A(J(t) = 1) \\ (15) \quad &\leq \sum_{t=1}^M \nu(k+1) = M \nu(k+1). \end{aligned}$$

We now study $\nu(k+1)$ to show that it is exponentially small. We have

$$\nu(k+1) = \prod_{j=1}^k \frac{p(j, j+1)}{p(j+1, j)} = \prod_{j=1}^{k-1} \frac{p(j, j+1)}{p(j+1, j)} = \lambda^{k-1} \prod_{j=1}^{k-1} \frac{1}{P_{x_{j+1}}(T_{x_j} < T_{x_{j+1}}^+)}.$$

Using the results of Lemmas 8 and 9, this becomes

$$\nu(k+1) \leq (2d\lambda)^{k-1} \prod_{j=1}^{k-1} l_j \leq (2d\lambda)^{k-1} (\kappa_d)^{n^d}.$$

Plugging this into (15), we obtain

$$\mathbf{P}_1^A(T_{k+1} \leq M) \leq M(2d\lambda)^{k-1} (\kappa_d)^{n^d}.$$

Recalling that $k = \lceil \mu n^d \rceil$, we have

$$(2d\lambda)^{k-1} \leq \frac{(2d\lambda)^{\mu n^d}}{1 \wedge (2d\lambda)},$$

whence the announced result. \square

3.8. Exponential number of topplings for a deterministic initial configuration

We now show how to combine the above results to prove that, still on the torus \mathbb{Z}_n^d , starting with a deterministic initial configuration with a prescribed density of active particles, the number of topplings necessary to reach a stable configuration is exponentially large with big probability, as stated in the following Lemma:

Lemma 11. *In any dimension $d \geq 1$, for every $\mu \in (0, 1)$ and $\lambda > 0$ satisfying the condition (1), there exists $c > 0$ such that, for n large enough, we have*

$$\sup_{\eta \in \mathbb{N}^{\mathbb{Z}_n^d} : |\eta| \geq \mu n^d} \mathcal{P}^\lambda \left(\|m_\eta^\tau\| < e^{cn^d} \right) < e^{-cn^d}.$$

Proof. Let $\lambda \in (0, \infty)$ and $\mu \in (0, 1)$ which satisfy the condition (1). This condition being a strict inequality, we may take $c > 0$ such that

$$(16) \quad \frac{\kappa_d(2d\lambda)^\mu}{\mu^\mu(1-\mu)^{1-\mu}} < e^{-2c}.$$

Let us start with a configuration $\eta \in \mathbb{N}^{\mathbb{Z}_n^d}$ such that $|\eta| = k = \lceil \mu n^d \rceil$. Let us recall that $\|m_\eta^\tau\|$ denotes the total number of topplings necessary to stabilize the configuration η using the toppling instructions in τ . For any $M \in \mathbb{N}$, using Lemmas 5, 6, 7 and 10, we can write

$$\begin{aligned} \mathcal{P}^\lambda \left(\|m_\eta^\tau\| \leq M \right) &= \sum_{|A|=k} \mathcal{P}^\lambda \left(\{ \|m_\eta^\tau\| \leq M \} \cap \{ \sigma_\eta^\tau = \mathbf{s} \mathbf{1}_A \} \right) \\ &\leq \sum_{|A|=k} \mathcal{P}^\lambda \left(\{ \|m_\eta^\tau\|_A \leq M \} \cap \{ \sigma_\eta^\tau = \mathbf{s} \mathbf{1}_A \} \right) \\ &\stackrel{(\text{Lemma 5})}{\leq} \sum_{|A|=k} \mathcal{P}^{\lambda, A} \left(\|m_\eta^\tau\|_A \leq M \right) \\ &\stackrel{(\text{Lemma 6})}{\leq} \sum_{|A|=k} \mathcal{P}^{\lambda, A} \left(\|m_{\mathbf{1}_A}^\tau\|_A \leq M \right) \\ &\stackrel{(\text{Lemma 7})}{\leq} \sum_{|A|=k} \mathbf{P}_1^A(T_{k+1} \leq M) \\ &\stackrel{(\text{Lemma 10})}{\leq} \left(\frac{n^d}{\lceil \mu n^d \rceil} \right) \frac{M}{1 \wedge (2d\lambda)} (\kappa_d(2d\lambda)^\mu)^{n^d}. \end{aligned} \quad (17)$$

Now note that Stirling's formula ensures that, when $m \rightarrow +\infty$,

$$\begin{aligned} \binom{m}{\lceil \mu m \rceil} &= \frac{m!}{\lceil \mu m \rceil! (m - \lceil \mu m \rceil)!} \stackrel{m \rightarrow \infty}{\sim} \frac{m^m \sqrt{2\pi m}}{\lceil \mu m \rceil^{\lceil \mu m \rceil} \sqrt{2\pi \mu m} (m - \lceil \mu m \rceil)^{m - \lceil \mu m \rceil} \sqrt{2\pi(1-\mu)m}} \\ &= O\left(\frac{m^m}{(\mu m)^{\mu m} (m - \mu m)^{m - \mu m} \sqrt{m}}\right) = O\left(\frac{1}{\mu^{\mu m} (1-\mu)^{(1-\mu)m} \sqrt{m}}\right). \end{aligned}$$

Thus, there exists $C = C(\mu) < \infty$ such that, for all $m \geq 1$,

$$\binom{m}{\lceil \mu m \rceil} \leq \frac{C}{\mu^{\mu m} (1-\mu)^{(1-\mu)m} \sqrt{m}}.$$

Plugging this into (17) and taking $M = \lfloor e^{cn^d} \rfloor$, we obtain

$$\mathcal{P}^\lambda \left(\|m_\eta^\tau\| < e^{cn^d} \right) \leq \frac{C}{(1 \wedge (2d\lambda)) n^{d/2}} \left[\frac{\kappa_d (2d\lambda)^\mu e^c}{\mu^\mu (1-\mu)^{1-\mu}} \right]^{n^d}.$$

With our assumption (16) on the constant c , this becomes

$$\mathcal{P}^\lambda \left(\|m_\eta^\tau\| < e^{cn^d} \right) \leq \frac{C}{(1 \wedge (2d\lambda)) n^{d/2}} e^{-cn^d}.$$

Since this constant c does not depend on the initial configuration η , we deduce that, for n large enough,

$$\sup_{\eta \in \mathbb{N}^{\mathbb{Z}_n^d} : |\eta| = \lceil \mu n^d \rceil} \mathcal{P}^\lambda \left(\|m_\eta^\tau\| < e^{cn^d} \right) < e^{-cn^d}.$$

By virtue of the monotonicity property stated in Lemma 1, this estimate remains true if the supremum is taken over all the configurations η with $|\eta| \geq \mu n^d$. \square

3.9. Exponential number of topplings with a Poisson initial distribution

The result of Lemma 11 easily extends to an initial configuration with Poisson i.i.d. numbers of active particles on each site of the torus, that is to say, distributed according to \mathcal{P}_μ .

Lemma 12. *In any dimension $d \geq 1$, for every $\mu \in (0, 1)$ and $\lambda > 0$ satisfying the condition (1), there exists $c > 0$ such that, for n large enough, we have*

$$\mathcal{P}_\mu^\lambda \left(\|m_{\eta_0}^\tau\| < e^{cn^d} \right) < e^{-cn^d}.$$

Proof. Let $\lambda \in (0, \infty)$ and $\mu \in (0, 1)$ satisfying the condition (1). Because this condition is a strict inequality, we may fix $0 < \mu' < \mu$ such that the condition (1) still holds for λ and μ' . Let us take $c > 0$, which will be chosen later.

Remember that \mathcal{P}_μ^λ is a joint probability measure where the initial configuration η_0 is distributed according to \mathcal{P}_μ , i.e., with i.i.d. Poisson numbers of active particles on each site, and the field of toppling instructions τ is distributed according to \mathcal{P}^λ , η_0 and τ being independent. By conditioning on η_0 , we may write

$$\begin{aligned} \mathcal{P}_\mu^\lambda \left(\|m_{\eta_0}^\tau\| < e^{cn^d} \right) &= \sum_{\eta \in \mathbb{N}^{\mathbb{Z}_n^d}} \mathcal{P}_\mu(\eta_0 = \eta) \mathcal{P}^\lambda \left(\|m_\eta^\tau\| < e^{cn^d} \right) \\ (18) \quad &\leq \mathcal{P}_\mu \left(|\eta_0| < \mu' n^d \right) + \sup_{\eta \in \mathbb{N}^{\mathbb{Z}_n^d} : |\eta| \geq \mu' n^d} \mathcal{P}^\lambda \left(\|m_\eta^\tau\| < e^{cn^d} \right). \end{aligned}$$

Remember that, under \mathcal{P}_μ , the initial configuration η_0 contains a total number of active particles $|\eta_0|$ which follows a Poisson distribution with parameter μn^d , whence, using again Stirling's formula,

$$\mathcal{P}_\mu \left(|\eta_0| < \mu' n^d \right) = \sum_{j < \mu' n^d} \frac{(\mu n^d)^j}{j!} e^{-\mu n^d} \leq e^{-\mu n^d} \lceil \mu' n^d \rceil \frac{(\mu n^d)^{\mu' n^d}}{(\mu' n^d)!}$$

$$\stackrel{n \rightarrow \infty}{\sim} e^{-\mu n^d} (\mu' n^d) \frac{(\mu n^d)^{\mu' n^d}}{(\mu' n^d)^{\mu' n^d} e^{-\mu' n^d} \sqrt{2\pi \mu' n^d}} = \sqrt{\frac{\mu' n^d}{2\pi}} e^{-c_1 n^d} \quad \text{with} \quad c_1 = (\mu - \mu') - \mu' \ln \left(\frac{\mu}{\mu'} \right).$$

Note that since we took $\mu' < \mu$, the convexity of the logarithm function ensures that $c_1 > 0$. Thus, for n large enough, we have

$$(19) \quad \mathcal{P}_\mu(|\eta_0| < \mu' n^d) \leq e^{-c_1 n^d/2}.$$

Besides, let us take $c_2 > 0$ given by Lemma 11, applied with λ and μ' , so that

$$(20) \quad \sup_{\eta \in \mathbb{N}^{\mathbb{Z}^d} : |\eta| \geq \mu' n^d} \mathcal{P}^\lambda(\|m_\eta^\tau\| < e^{c_2 n^d}) \leq e^{-c_2 n^d}.$$

If we choose $c > 0$ such that $c < \min(c_1/2, c_2)$, then plugging (19) and (20) into (18), we get

$$\mathcal{P}_\mu^\lambda(\|m_{\eta_0}^\tau\| < e^{cn^d}) \leq e^{-c_1 n^d} + e^{-c_2 n^d} < e^{-cn^d},$$

for n large enough. □

3.10. Concluding proof of Theorem 3

We are now in a position to conclude the proof of Theorem 3, which easily follows from Lemma 12.

Proof of Theorem 3. There only remains to deal with the Poisson clocks, to translate the result of Lemma 12 about the number of topplings into a result on the fixation time of the continuous-time process.

As explained in paragraph 2.4, the Activated Random Walk model on the torus can be constructed by drawing the initial configuration η_0 and the field of toppling instructions τ distributed according to \mathcal{P}_μ^λ , and then a collection of independent Poisson process $(\theta_i)_{i \in \mathbb{N}}$ of intensity $1 + \lambda$, independent of η_0 and τ . Each time one of the clocks $(\theta_i)_{i < |\eta_0|}$ rings, we choose a particle (not a site) uniformly at random, and topple its site if this particle is active at that moment (and we do nothing if the chosen particle is inactive). With this construction, we can write, for $c > 0$ which will be chosen later,

$$(21) \quad \mathbb{P}_\mu^\lambda(\mathcal{T}_n < e^{cn^d}) \leq \mathbb{P}_\mu^\lambda(\{\mathcal{T}_n < e^{cn^d}\} \cap \{|\eta_0| \leq n^d\} \cap \{\|m_{\eta_0}^\tau\| \geq e^{3cn^d}\}) \\ + \mathcal{P}_\mu(|\eta_0| > n^d) + \mathcal{P}_\mu^\lambda(\|m_{\eta_0}^\tau\| < e^{3cn^d}).$$

We know that the clocks $(\theta_i)_{i < |\eta_0|}$ ring a total number of times at least $\|m_{\eta_0}^\tau\|$ during the time interval $[0, \mathcal{T}_n]$, whence

$$(22) \quad \mathbb{P}_\mu^\lambda(\{\mathcal{T}_n < e^{cn^d}\} \cap \{|\eta_0| \leq n^d\} \cap \{\|m_{\eta_0}^\tau\| \geq e^{3cn^d}\}) \leq \mathbb{P}_\mu^\lambda\left(\sum_{i < n^d} |\theta_i \cap [0, e^{cn^d}]| \geq e^{3cn^d}\right) \\ \leq n^d \mathbb{E}_\mu^\lambda[|\theta_0 \cap [0, e^{cn^d}]|] e^{-3cn^d} = n^d(1 + \lambda) e^{-2cn^d},$$

where we used Chebychev's inequality. Besides, using a Chernoff bound, we can write, for $a > 0$,

$$(23) \quad \mathcal{P}_\mu(|\eta_0| > n^d) \leq \mathcal{E}_\mu(e^{a|\eta_0|}) e^{-an^d} = \exp((e^a - 1)\mu n^d - an^d).$$

Since $\mu < 1$ and $(e^a - 1)\mu - a \sim (\mu - 1)a$ when $a \rightarrow 0$, we may choose $a > 0$ such that $(e^a - 1)\mu - a < 0$. Also, we take $c_1 > 0$ given by Lemma 12, so that, for n large enough,

$$(24) \quad \mathcal{P}_\mu^\lambda(\|m_{\eta_0}^\tau\| < e^{c_1 n^d}) < e^{-c_1 n^d}.$$

Plugging (22), (23) and (24) into (21), and choosing $c > 0$ such that

$$c < \min(a - (e^a - 1)\mu, c_1),$$

we get, for n large enough,

$$\mathbb{P}_\mu^\lambda(\mathcal{T}_n < e^{cn^d}) \leq n^d(1+\lambda)e^{-2cn^d} + e^{-(a-(e^a-1)\mu)n^d} + e^{-c_1n^d} < e^{-cn^d}.$$

Decreasing the constant c if necessary so that the result holds for all $n \geq 1$, we obtain Theorem 3. \square

4. From the torus to the infinite lattice

We now turn to the proof of Theorem 4, which shows that the fixation time on the torus in the sub-critical regime cannot satisfy (2). Before this, we start with two elementary results about the stabilization time of a box with open boundary condition and about Internal Diffusion Limited Aggregation.

4.1. Upper bound on the stabilization time

We start by proving the following rough estimate:

Lemma 13. *We have the following upper bound on the number $\|m_{\Lambda_n, \eta_0}^\tau\|$ of toppling instructions necessary to stabilize the box $\Lambda_n \subset \mathbb{Z}^d$ with particles being killed when they jump out of Λ_n :*

$$\forall b > 0 \quad \forall \mu \in (0, 1) \quad \mathcal{P}_\mu^\lambda(\|m_{\Lambda_n, \eta_0}^\tau\| > e^{bn^d}) = O\left(\frac{n^{1+3d/2}}{e^{bn^d/2}}\right).$$

Proof. Let $b > 0$ and $\mu \in (0, 1)$. Let us write $c_b = 2 + b/2$. Conditioning on the initial configuration η_0 , we can write

$$\begin{aligned} \mathcal{P}_\mu^\lambda(\|m_{\Lambda_n, \eta_0}^\tau\| > e^{bn^d}) &= \sum_{\eta \in \mathbb{N}^{\Lambda_n}} \mathcal{P}_\mu(\eta_0 = \eta) \mathcal{P}^\lambda(\|m_{\Lambda_n, \eta}^\tau\| > e^{bn^d}) \\ &\leq \mathcal{P}_\mu(|\eta_0| > c_b n^d) + \sup_{\eta \in \mathbb{N}^{\Lambda_n} : |\eta| \leq c_b n^d} \mathcal{P}^\lambda(\|m_{\Lambda_n, \eta}^\tau\| > e^{bn^d}). \end{aligned}$$

To handle the first term, we use a Chernoff bound, recalling that $|\eta_0|$ follows a Poisson distribution with mean μn^d . Recalling that \mathcal{E}_μ denotes the expectation relative to \mathcal{P}_μ , we have

$$\mathcal{P}_\mu(|\eta_0| > c_b n^d) \leq \mathcal{E}_\mu(e^{|\eta_0|})e^{-c_b n^d} = e^{(e\mu - \mu - c_b)n^d} \leq e^{-bn^d/2}.$$

Thus, it remains to show that

$$(25) \quad \sup_{\eta \in \mathbb{N}^{\Lambda_n} : |\eta| \leq c_b n^d} \mathcal{P}^\lambda(\|m_{\Lambda_n, \eta}^\tau\| > e^{bn^d}) = O\left(\frac{n^{1+3d/2}}{e^{bn^d/2}}\right).$$

To do so, we fix $\eta \in \mathbb{N}^{\Lambda_n}$ such that $|\eta| \leq c_b n^d$, and we consider the number J_n of jump instructions (i.e., topplings which are not sleep instructions) which are necessary to stabilize the box Λ_n with particles ignored once they jump out of this box. Writing $M = \lceil e^{bn^d} \rceil$ and $M' = \lfloor e^{bn^d}/(2+2\lambda) \rfloor$, we have

$$(26) \quad \mathcal{P}^\lambda(\|m_{\Lambda_n, \eta}^\tau\| \geq M) \leq \mathcal{P}^\lambda(J_n > M') + \mathcal{P}^\lambda(\{\|m_{\Lambda_n, \eta}^\tau\| \geq M\} \cap \{J_n \leq M'\}).$$

If $\|m_{\Lambda_n, \eta}^\tau\| \geq M$ but $J_n \leq M'$, then among the first M topplings revealed, at most M' are jump instructions. Thus, if X_1, \dots, X_M are i.i.d. Bernoulli random variables with parameter $p = 1/(1+\lambda)$, we have

$$\mathcal{P}^\lambda(\{\|m_{\Lambda_n, \eta}^\tau\| \geq M\} \cap \{J_n \leq M'\}) \leq \mathbb{P}(X_1 + \dots + X_M \leq M'),$$

whence, by Chebychev's inequality,

$$(27) \quad \mathcal{P}^\lambda(\{\|m_{\Lambda_n, \eta}^\tau\| \geq M\} \cap \{J_n \leq M'\}) \leq \frac{\text{Var}(X_1 + \dots + X_M)}{(Mp - M')^2} \leq \frac{Mp(1-p)}{(Mp/2)^2} = \frac{4\lambda}{M} \leq 4\lambda e^{-bn^d}.$$

We now deal with the other term $\mathcal{P}^\lambda(J_n > M')$ appearing in (26). In the initial configuration η , let us label the particles present at each site $x \in \Lambda_n$ with labels $i \in \{1, \dots, \eta(x)\}$. When we topple a site containing more than one particle, let

us say that we topple the particle carrying the minimal label, with respect to an arbitrary order on $\Lambda_n \times \mathbb{N}$. This allows us to define, for each particle labelled (x, i) , the number $j^{x,i}$ of jumps (i.e., topplings which are not sleep instructions) it makes during the stabilization, and its trajectory $X^{x,i}(0) = x, \dots, X^{x,i}(j^{x,i})$, where $X^{x,i}(m)$ is the position of the particle labelled (x, i) after m jump instructions. With these notations, we have

$$J_n = \sum_{x \in \Lambda_n} \sum_{i=1}^{\eta(x)} j^{x,i}.$$

Given that $|\eta| \leq c_b n^d$, this implies that

$$\mathcal{P}^\lambda(J_n > M') \leq \sum_{x \in \Lambda_n} \sum_{i=1}^{\eta(x)} \mathcal{P}^\lambda\left(j^{x,i} > \frac{M'}{c_b n^d}\right).$$

Let us complete the trajectory of each particle with independent jumps to obtain, for each particle labelled (x, i) , a trajectory $(X^{x,i}(m))_{m \in \mathbb{N}} \in (\mathbb{Z}^d)^\mathbb{N}$. Thus, for each particle labelled (x, i) , the random walk $X^{x,i}$ is distributed as a symmetric random walk on \mathbb{Z}^d . We then note that if we have $j^{x,i} > M'/(c_b n^d)$, it implies that $X^{x,i}(\lfloor M'/(c_b n^d) \rfloor) \in \Lambda_n$, whence

$$(28) \quad \mathcal{P}^\lambda(J_n > M') \leq \sum_{x \in \Lambda_n} \sum_{i=1}^{\eta(x)} \mathcal{P}^\lambda\left(X^{x,i}\left(\left\lfloor \frac{M'}{c_b n^d} \right\rfloor\right) \in \Lambda_n\right) \leq c_b n^d P_0^{\mathbb{Z}^d}\left(\left\|X\left(\left\lfloor \frac{M'}{c_b n^d} \right\rfloor\right)\right\|_\infty \leq n\right),$$

where, under $P_0^{\mathbb{Z}^d}$, X is a simple random walk on \mathbb{Z}^d started at the origin. By the pigeonhole principle, among the first $\lfloor M'/(c_b n^d) \rfloor$ steps of a random walk on \mathbb{Z}^d , at least $\lfloor M'/(d c_b n^d) \rfloor$ of them are in the same direction. Therefore, we have

$$P_0^{\mathbb{Z}^d}\left(\left\|X\left(\left\lfloor \frac{M'}{c_b n^d} \right\rfloor\right)\right\|_\infty \leq n\right) \leq d P_0^{\mathbb{Z}}\left(\left|X\left(\left\lfloor \frac{M'}{d c_b n^d} \right\rfloor\right)\right| \leq n\right).$$

A classical computation shows that, when $k \rightarrow \infty$,

$$P_0^{\mathbb{Z}}(|X(k)| \leq n) \leq \frac{n+1}{2^k} \binom{k}{\lfloor k/2 \rfloor} \stackrel{k \rightarrow \infty}{\sim} \sqrt{\frac{2}{\pi}} \frac{n+1}{\sqrt{k}},$$

whence, recalling that $M' = \lfloor e^{b n^d} / (2 + 2\lambda) \rfloor$,

$$P_0^{\mathbb{Z}^d}\left(\left\|X\left(\left\lfloor \frac{M'}{c_b n^d} \right\rfloor\right)\right\|_\infty \leq n\right) = O(n^{1+d/2} e^{-b n^d/2}).$$

Plugging this into (28), we obtain

$$(29) \quad \mathcal{P}^\lambda(J_n > M') = O(n^{1+3d/2} e^{-b n^d/2}),$$

the bound being uniform with respect to the initial configuration η . Replacing (27) and (29) into (26), we obtain the desired bound (25), concluding the proof. \square

4.2. A bound on Internal Diffusion Limited Aggregation

In the proof of Theorem 4, we will use a rough estimate about a related model called Internal Diffusion Limited Aggregation, which corresponds to the Activated Random Walk model with an infinite sleep rate. In this model, active particles fall asleep instantaneously as soon as they are alone on a site.

Starting from an initial configuration of active particles $\eta \in \mathbb{N}^{\mathbb{Z}^d}$ with a finite number of particles, the model almost surely reaches a stable configuration which we write σ_η^τ , as defined in section 2.3. The result we need is the following Lemma, which is a weak version of Lemma A of [19] (note that \mathcal{P}^∞ is simply \mathcal{P}^λ with $\lambda = \infty$):

Lemma 14. *For $\beta, R \in (0, \infty)$, let us denote by $B(0, R)$ the Euclidian ball of radius R centered on the origin, and let $\mathcal{C}_{\beta, R}$ be the set of all possible initial configurations with not more than βR^d particles and no particle inside $B(0, R)$, that is to say:*

$$\mathcal{C}_{\beta, R} = \left\{ \eta \in \mathbb{N}^{\mathbb{Z}^d} : |\eta| \leq \beta R^d \quad \text{and} \quad \forall x \in \mathbb{Z}^d \cap B(0, R) \quad \eta(x) = 0 \right\}.$$

Then there exists $K = K(d) > 0$ such that, for β small enough, when $d \geq 3$, we have

$$\sup_{\eta \in \mathcal{C}_{\beta, R}} \mathcal{P}^\infty(\sigma_\eta^\tau(0) > 0) \leq e^{-KR^2},$$

whereas for $d = 2$ we have

$$\sup_{\eta \in \mathcal{C}_{\beta, R}} \mathcal{P}^\infty(\sigma_\eta^\tau(0) > 0) \leq \exp\left(-K \frac{R^2}{\ln R}\right).$$

Note that in dimension $d = 1$, we simply have $\mathcal{P}^\infty(\sigma_\eta^\tau(0) > 0) = 0$ for every initial configuration $\eta \in \mathcal{C}_{1, R}$.

In fact, we do not need a bound as sharp as above, and we only use that, whatever the dimension $d \geq 1$, we have, for β small enough,

$$(30) \quad \forall R \geq 1 \quad \sup_{\eta \in \mathcal{C}_{\beta, R}} \mathcal{P}^\infty(\sigma_\eta^\tau(0) > 0) \leq e^{-K(d)R}.$$

4.3. Proof of Theorem 4

Proof of Theorem 4. Let $\lambda \in (0, \infty)$ and let $\mu < \mu_c(\lambda)$. We want to prove that the property (2) does not hold on the torus. Hence, we have to prove that

$$(31) \quad \forall c > 0 \quad \exists n \geq 1 \quad \mathbb{P}_\mu^\lambda(\mathcal{T}_n < e^{cn^d}) \geq e^{-cn^d},$$

where \mathcal{T}_n is the stabilization time of the model on the torus. Thus, we start by fixing an arbitrary real number $c > 0$. Let $\beta > 0$ be given by Lemma 14 about Internal Diffusion Limited Aggregation, and let us choose two real parameters a and ε such that

$$(32) \quad 0 < a < \frac{c}{96d\mu} \wedge \frac{c}{72d \ln(1 + 1/\lambda)} \wedge \frac{1}{3}$$

and

$$(33) \quad 0 < \varepsilon < \frac{\beta a^d}{4^d} \wedge \frac{c}{24}.$$

As explained in section 1.6.2, we consider four sizes of sub-boxes of the torus, by writing $p: \mathbb{Z}^d \rightarrow \mathbb{Z}_n^d$ for the canonical projection application and by considering the projection of the boxes Λ_k defined by (4). We call $B_0 = \mathbb{Z}_n^d = p(\Lambda_n)$ the whole box, $B_1 = p(\Lambda_{n-\lfloor an \rfloor})$ the medium box, $B_2 = p(\Lambda_{n-2\lfloor an \rfloor})$ the small box and $B_3 = p(\Lambda_{n-3\lfloor an \rfloor})$ the tiny box. For every box Λ_k , we denote its internal boundary by

$$\partial^i \Lambda_k = \left\{ x \in \Lambda_k : \exists y \in \mathbb{Z}^d \setminus \Lambda_k \quad x \sim y \right\} = \Lambda_k \setminus \Lambda_{k-2},$$

and its external boundary by

$$\partial^e \Lambda_k = \left\{ x \in \mathbb{Z}^d \setminus \Lambda_k : \exists y \in \Lambda_k \quad x \sim y \right\},$$

and, for every $j \in \{0, 1, 2, 3\}$, we write $\partial^i B_j = p(\partial^i \Lambda_{n-j\lfloor an \rfloor})$ and $\partial^e B_j = p(\partial^e \Lambda_{n-j\lfloor an \rfloor})$ for the projections on the torus of these boundaries.

As explained in paragraph 1.6.2, the idea is to assume that we start with all the particles inside the small box B_2 , then to stabilize this small box, and then to try to stabilize the particles which escaped from this box, without waking up the particles which are sleeping inside of the tiny box B_3 .

We now describe a toppling procedure, which, when successful, yields a sequence of topplings which stabilizes the configuration in the torus B_0 , with no particle ever visiting the boundary $\partial^i B_0$. During this procedure, we allow ourselves to use acceptable topplings, as defined in section 2.1. Thanks to the monotonicity property of Lemma 1, this yields an upper bound on the number of topplings necessary to stabilize the torus with only legal topplings and, thereby, on the stabilization time of the torus in the continuous-time process.

Let us now describe our toppling procedure, which is composed of several steps, represented on figure 4. At each step, in certain cases, the procedure can fail, in which case we stop the procedure.

- Step 0 is to draw the initial configuration η_0 according to \mathcal{P}_μ , that is to say, with i.i.d. Poisson numbers of active particles with mean μ on each site. We say that step 0 is successful if there are no particles outside of the small box B_2 , i.e., if $\eta_0(x) = 0$ for all $x \in B_0 \setminus B_2$.
- In the first step, we stabilize the small box B_2 , with particles being stopped in place upon jumping out of B_2 . We say that this step is successful if it terminates with a finite number of topplings and at most εn^d particles jump out of the small box. Then, we obtain a configuration with some sleeping particles in B_2 , and at most εn^d active particles all located on the external boundary $\partial^e B_2$.
- Then, during the second step, the idea is to force all the particles left on $\partial^e B_2$ to walk until they jump out of the medium box B_1 . It is during this step that we use acceptable topplings. First, we choose a non-empty site of $\partial^e B_2$ and we topple it. If it was a sleep instruction, we topple it again (this might not be a legal toppling but it is acceptable), until we get a jump instruction to one of the neighbouring sites. We then topple this neighbouring site, until we get a jump instruction, and so on, until we reach a site of $\partial^i B_1$. We then repeat this procedure until all the sites of $\partial^e B_2$ are empty.

Although formally we are not allowed to say that we topple a specific particle instead of a site, one can think of this step as forcing a particle to walk, ignoring the sleep instructions and ignoring all the other particles. Thus, during this step, some particles which were asleep in the small box B_2 after step 1 might wake up, but the number of particles on each site of this small box is left unchanged after step 2.

We say that this step 2 is successful if it terminates in a finite number of topplings and if no active particle ever visits the tiny box B_3 during this step. In this case, we obtain a configuration with some sleeping particles in the tiny box B_3 , some sleeping and some active particles in $B_2 \setminus B_3$ (because some particles left asleep there after step 1 may have been awoken during step 2), and at most εn^d active particles on the boundary $\partial^e B_1$, with no particles elsewhere. Note in particular that, after step 2, no site of the small box B_2 can contain more than one particle.

- During step 3, as long as there is a site containing at least two particles, we topple it. We say that this step is successful if it terminates and if no particle ever visits neither $\partial^i B_0$ nor $\partial^e B_2$ during this step. In this case, the configuration inside the small box B_2 does not change during this step, and we end up with at most εn^d sites of $B_0 \setminus B_2$ with exactly one active particle, the other sites of $B_0 \setminus B_2$ being empty.
- During the fourth and last step, we topple exactly once each of the unstable sites. We say that step 4 is successful if all the topplings revealed during this step are sleep instructions. Since we start step 4 with at most one particle on each site, if this step is successful it leads to a stable configuration, with all the particles asleep.

We now estimate the probability that each of these steps is successful. For each $i \in \{0, \dots, 4\}$, we denote by \mathcal{S}_i the event that the steps 0 to i are all successful.

First, we estimate $\mathcal{P}_\mu(\mathcal{S}_0)$. At each site $x \in B_0$, the initial number of particles $\eta_0(x)$ is drawn according to a Poisson distribution with parameter μ , the number of particles at distinct sites being independent, whence

$$(34) \quad \mathcal{P}_\mu(\mathcal{S}_0) = (e^{-\mu})^{|B_0 \setminus B_2|} \geq e^{-4d\mu \lfloor an \rfloor n^{d-1}} \geq e^{-4d\mu an^d} \geq e^{-cn^d/24},$$

where we have used our assumption (32) on the parameter a in the last inequality.

We now turn to the event \mathcal{S}_1 . Writing $M_{n-2\lfloor an \rfloor}$ for the number of particles which jump out of the small box B_2 during step 1, it follows from Lemma 4 that, when $n \rightarrow \infty$,

$$\mathcal{E}_\mu^\lambda(M_{n-2\lfloor an \rfloor}) = o(|B_2|).$$

In fact, Lemma 4 deals with the model on \mathbb{Z}^d instead of the torus but, if one is only interested in the number of particles which jump out of a square box of side $n - 2\lfloor an \rfloor$ when stabilizing only inside this box with particles ignored once they jump out of it, the fact that this box is included in a torus or in \mathbb{Z}^d has no importance. Thus, for n large enough, we have

$$\mathcal{E}_\mu^\lambda(M_{n-2\lfloor an \rfloor}) \leq \frac{\varepsilon}{4(1-2a)^d} |B_2| \stackrel{n \rightarrow \infty}{\sim} \frac{\varepsilon n^d}{4},$$

which implies by Markov's inequality that, for n large enough,

$$\mathcal{P}_\mu^\lambda(M_{n-2\lfloor an \rfloor} \leq \varepsilon n^d) \geq \frac{1}{2}.$$

This event being independent of \mathcal{S}_0 , combining this with (34) we get

$$(35) \quad \mathcal{P}_\mu^\lambda(\mathcal{S}_1) = \mathcal{P}_\mu^\lambda(\mathcal{S}_0 \cap \{M_{n-2\lfloor an \rfloor} \leq \varepsilon n^d\}) \geq \frac{1}{2} e^{-cn^d/24}.$$

Next, we come to the event \mathcal{S}_2 that the procedure is successful until step 2 included. Let us denote by $\eta_1 : B_0 \rightarrow \mathbb{N}_s$ the configuration reached after step 1, with sleeping particles inside B_2 and $M_{n-2\lfloor an \rfloor}$ active particles on the boundary $\partial^e B_2$. For every $x \in \partial^e B_2$, the probability that a particle starting at x reaches the boundary $\partial^e B_1$ before reaching the tiny box B_3 is the probability that a simple random walk on \mathbb{Z}^d hits $\partial^e B_1$ before B_3 when started at $x \in \partial^e B_2$, which by symmetry is at least $1/2$. Hence, for every $\eta : \partial^e B_2 \rightarrow \mathbb{N}$ such that $|\eta| \leq \varepsilon n^d$ (so that we condition on an event which has positive probability), we have

$$\mathcal{P}_\mu^\lambda(\mathcal{S}_2 \mid \mathcal{S}_1 \cap \{\eta_1|_{\partial^e B_2} = \eta\}) \geq \frac{1}{2^{|\eta|}} \geq \frac{1}{2^{\varepsilon n^d}} \geq e^{-cn^d/24},$$

where in the last inequality we used our assumption (33) on ε which ensures that $\varepsilon \ln 2 < \varepsilon < c/24$. This bound being uniform with respect to η , we deduce that

$$\mathcal{P}_\mu^\lambda(\mathcal{S}_2 \mid \mathcal{S}_1) \geq e^{-cn^d/24}.$$

Combining this with our previous estimate (35), we obtain

$$(36) \quad \mathcal{P}_\mu^\lambda(\mathcal{S}_2) \geq \frac{1}{2} e^{-cn^d/12}.$$

We now study the probability of success of the third step. Recall that this step starts with a certain number of active particles on $\partial^e B_1$, the other particles all being alone on their sites inside the small box B_2 (see figure 4).

Recall also that, during this step, we only topple the sites with at least two particles. Thus, the dynamic on the particles, regardless of their state, active or sleeping, exactly corresponds to Internal Diffusion Limited Aggregation, as described in paragraph 4.2. Furthermore, we can note that, as soon as a particle reaches a site of $\partial^i B_0 \cup \partial^e B_2$, step 3 can be stopped and declared unsuccessful. Doing so, no site of the small box B_2 is ever toppled during step 3, so that, during this step, we may forget the particles which lie in B_2 .

Conditioned on the realization of \mathcal{S}_2 and on the configuration η_2 reached after step 2, the probability of \mathcal{S}_3 is given by the probability that the Internal Diffusion Limited Aggregation (IDLA) starting with η_2 never reaches neither the internal boundary of the whole box, $\partial^i B_0$ nor the external boundary of the small box, $\partial^e B_2$. With the above remark, we can see that this probability does not change if, instead of η_2 , we start the IDLA with the configuration η_2 restricted to $\partial^e B_1$ (because the other particles of η_2 lie inside the small box B_2 and thus have no influence on the probability of success of step 3). Therefore, fixing a configuration $\eta : \partial^e B_1 \rightarrow \mathbb{N}$ such that $|\eta| \leq \varepsilon n^d$, we may write

$$(37) \quad \begin{aligned} \mathcal{P}_\mu^\lambda(\mathcal{S}_3 \mid \mathcal{S}_2 \cap \{\eta_2|_{\partial^e B_1} = \eta\}) &= \mathcal{P}^\infty(\forall y \in \partial^i B_0 \cup \partial^e B_2 \quad \sigma_\eta^\tau(y) = 0) \\ &\geq 1 - \sum_{y \in \partial^i B_0 \cup \partial^e B_2} \mathcal{P}^\infty(\sigma_\eta^\tau(y) > 0). \end{aligned}$$

Yet, defining $R_n = an/2 - 2$, it follows from our assumption (33) on ε that, for n large enough,

$$\varepsilon n^d \leq \frac{\beta a^d n^d}{4^d} \leq \beta \left(\frac{an}{2} - 2 \right)^d = \beta R_n^d,$$

implying that $|\eta| \leq \beta R_n^d$. Thus, we may apply Lemma 14, or rather its weakened version (30), to obtain that for every point $y \in \partial^i B_0 \cup \partial^e B_2$, we have

$$\mathcal{P}^\infty(\sigma_\eta^\tau(y) > 0) \leq e^{-KR_n}.$$

Plugging this into (37), we obtain that

$$\mathcal{P}_\mu^\lambda(\mathcal{S}_3 \mid \mathcal{S}_2 \cap \{\eta_2|_{\partial^e B_1} = \eta\}) \geq 1 - |\partial^i B_0 \cup \partial^e B_2| e^{-KR_n} \geq 1 - 4dn^{d-1} e^{2K-Kan} \geq \frac{1}{2},$$

for n large enough. This bound being uniform with respect to η , we deduce that, for n large enough,

$$\mathcal{P}_\mu^\lambda(\mathcal{S}_3 \mid \mathcal{S}_2) \geq \frac{1}{2}.$$

Combined with our previous estimate (36) on the probability of \mathcal{S}_2 , this yields

$$(38) \quad \mathcal{P}_\mu^\lambda(\mathcal{S}_3) \geq \frac{1}{4} e^{-cn^d/12}.$$

We now deal with the event \mathcal{S}_4 . At the beginning of step 4, the unstable sites are the sites where the particles spread during step 3, which are all in $B_0 \setminus B_2$, and the sites of the particles in $B_2 \setminus B_3$ which have been awakened during step 2. In any case, there is no unstable site in the tiny box B_3 . Thus, the number of topplings to perform during step 4 does not exceed

$$|B_0 \setminus B_3| \leq 3d \lfloor an \rfloor n^{d-1} \leq 3dan^d,$$

whence, recalling our assumption (32) on the parameter a ,

$$\mathcal{P}_\mu^\lambda(\mathcal{S}_4 | \mathcal{S}_3) \geq \left(\frac{\lambda}{1+\lambda} \right)^{3dan^d} \geq e^{-cn^d/24}.$$

Combining this with (38), we get

$$(39) \quad \mathcal{P}_\mu^\lambda(\mathcal{S}_4) \geq \frac{1}{4} e^{-cn^d/8}.$$

As explained above, if the event \mathcal{S}_4 is realized, then the whole procedure is successful and it yields an acceptable sequence of topplings α which stabilizes the initial configuration η_0 , and which is such that the sites on the boundary $\partial^i B_0$ are never toppled. Thanks to the monotonicity property given by Lemma 2, this implies that $m_{\eta_0}^\tau(x) = 0$ for all $x \in \partial^i B_0$, where we recall that $m_{\eta_0}^\tau$ is the odometer of a legal sequence which stabilizes η in the torus using the toppling instructions in τ , as defined in section 2.1. In this case, $m_{\eta_0}^\tau$ is a legal sequence of topplings which only topples the sites of $B'_0 = B_0 \setminus \partial^i B_0$ and which stabilizes the configuration η_0 on the torus, and, especially, it also stabilizes η_0 on B'_0 . Hence, by the Abelian property (Lemma 3), the occurrence of \mathcal{S}_4 implies that $m_{\eta_0}^\tau = m_{B'_0, \eta_0}^\tau$. Thus, our result (39) becomes

$$(40) \quad \mathcal{P}_\mu^\lambda(m_{\eta_0}^\tau = m_{B'_0, \eta_0}^\tau) \geq \frac{1}{4} e^{-cn^d/8}.$$

Besides, Lemma 13 tells us that

$$\mathcal{P}_\mu^\lambda(\|m_{B'_0, \eta_0}^\tau\| > e^{cn^d/2}) = O\left(\frac{n^{1+3d/2}}{e^{c(n-2)^d/4}}\right),$$

whence, for n large enough,

$$\mathcal{P}_\mu^\lambda(\|m_{B'_0, \eta_0}^\tau\| > e^{cn^d/2}) \leq \frac{1}{8} e^{-cn^d/8}.$$

Combining this with (40), we have

$$(41) \quad \begin{aligned} \mathcal{P}_\mu^\lambda(\|m_{\eta_0}^\tau\| \leq e^{cn^d/2}) &\geq \mathcal{P}_\mu^\lambda(\{m_{\eta_0}^\tau = m_{B'_0, \eta_0}^\tau\} \cap \{\|m_{B'_0, \eta_0}^\tau\| \leq e^{cn^d/2}\}) \\ &\geq \mathcal{P}_\mu^\lambda(m_{\eta_0}^\tau = m_{B'_0, \eta_0}^\tau) - \mathcal{P}_\mu^\lambda(\|m_{B'_0, \eta_0}^\tau\| > e^{cn^d/2}) \geq \frac{1}{8} e^{-cn^d/8}. \end{aligned}$$

We now turn to the stabilization time \mathcal{T}_n of the continuous-time process on the torus, and we write

$$(42) \quad \begin{aligned} \mathbb{P}_\mu^\lambda(\mathcal{T}_n < e^{cn^d}) &\geq \mathbb{P}_\mu^\lambda(\{\mathcal{T}_n < e^{cn^d}\} \cap \{\|m_{\eta_0}^\tau\| \leq e^{cn^d/2}\}) \\ &= \mathcal{P}_\mu^\lambda(\|m_{\eta_0}^\tau\| \leq e^{cn^d/2}) - \mathbb{P}_\mu^\lambda(\{\mathcal{T}_n \geq e^{cn^d}\} \cap \{\|m_{\eta_0}^\tau\| \leq e^{cn^d/2}\}). \end{aligned}$$

If we have $\mathcal{T}_n \geq e^{cn^d}$ while $\|m_{\eta_0}^\tau\| \leq e^{cn^d/2}$, it means that, during the time interval $[0, e^{cn^d}]$ the Poisson clocks triggering a toppling event rang at most $e^{cn^d/2}$ times, although there were always at least one active particle during this time frame, which entails that a toppling event happens with rate at least $1 + \lambda$. Hence, taking $\theta \subset \mathbb{R}_+$ a Poisson process of intensity $1 + \lambda$, we have, using Chebychev's inequality,

$$\begin{aligned} \mathbb{P}_\mu^\lambda(\{\mathcal{T}_n \geq e^{cn^d}\} \cap \{\|m_{\eta_0}^\tau\| \leq e^{cn^d/2}\}) &\leq \mathbb{P}(|\theta \cap [0, e^{cn^d}]| \leq e^{cn^d/2}) \\ &\leq \frac{\text{Var}(|\theta \cap [0, e^{cn^d}]|)}{(e^{cn^d} - e^{cn^d/2})^2} = \frac{e^{cn^d}}{(e^{cn^d} - e^{cn^d/2})^2} = O(e^{-cn^d}). \end{aligned}$$

whence, for n large enough,

$$(43) \quad \mathbb{P}_\mu^\lambda \left(\{ \mathcal{T}_n \geq e^{cn^d} \} \cap \{ \|m_{\eta_0}^\tau\| \leq e^{cn^d/2} \} \right) \leq \frac{1}{16} e^{-cn^d/8}.$$

Plugging (41) and (43) into (42), we obtain that, for n large enough,

$$\mathbb{P}_\mu^\lambda \left(\mathcal{T}_n < e^{cn^d} \right) \geq \frac{1}{16} e^{-cn^d/8},$$

which proves (31), concluding the proof of Theorem 4. \square

Acknowledgments

A. G. thanks Leonardo Rolla for his introduction to the model and his kind hospitality at NYU Shanghai where this work could start. A. G. and N. F. thank the referee for his comments which helped to clarify several parts of the manuscript.

References

- [1] ALVES, O., MACHADO, F. and POPOV, S. (2002). Phase transition for the frog model. *Electronic Journal of Probability* **7** 1–21.
- [2] ALVES, O. S., MACHADO, F. P. and POPOV, S. Y. (2002). The shape theorem for the frog model. *The Annals of Applied Probability* **12** 533–546.
- [3] AMIR, G. and GUREL-GUREVICH, O. (2010). On fixation of activated random walks. *Electron. Commun. Probab.* **15** 119–123. <https://doi.org/10.1214/ECP.v15-1536> MR2643591
- [4] ASSELAH, A., ROLLA, L. T. and SCHAPIRA, B. (2019). Diffusive bounds for the critical density of activated random walks. *arXiv preprint arXiv:1907.12694*.
- [5] BAK, P., TANG, C. and WIESENFIELD, K. (1987). Self-organized criticality: An explanation of the $1/f$ noise. *Physical review letters* **59** 381.
- [6] BASU, R., GANGULY, S. and HOFFMAN, C. (2018). Non-fixation for conservative stochastic dynamics on the line. *Comm. Math. Phys.* **358** 1151–1185. <https://doi.org/10.1007/s00220-017-3059-7> MR3778354
- [7] BASU, R., GANGULY, S., HOFFMAN, C. and RICHEY, J. (2019). Activated random walk on a cycle. *Ann. Inst. Henri Poincaré Probab. Stat.* **55** 1258–1277. <https://doi.org/10.1214/18-aihp918> MR4010935
- [8] CABEZAS, M., ROLLA, L. T. and SIDORAVICIUS, V. (2014). Non-equilibrium phase transitions: activated random walks at criticality. *J. Stat. Phys.* **155** 1112–1125. <https://doi.org/10.1007/s10955-013-0909-3> MR3207731
- [9] DHAR, D. (2006). Theoretical studies of self-organized criticality. *Phys. A* **369** 29–70. <https://doi.org/10.1016/j.physa.2006.04.004> MR2246566
- [10] DIACONIS, P. and FULTON, W. (1991). A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Sem. Mat. Univ. Pol. Torino* **49** 95–119.
- [11] DICKMAN, R., ROLLA, L. T. and SIDORAVICIUS, V. (2010). Activated random walkers: Facts, conjectures and challenges. *Journal of Statistical Physics* **138** 126–142.
- [12] DOYLE, P. G. and SNELL, J. L. (1984). *Random walks and electric networks*. Carus Mathematical Monographs **22**. Mathematical Association of America, Washington, DC. MR920811
- [13] FORIEN, N. (2020). About self-organized criticality, PhD thesis, Université Paris-Saclay.
- [14] GANTERT, N. and SCHMIDT, P. (2009). Recurrence for the frog model with drift on \mathbb{Z} . *Markov Process. Related Fields* **15** 51–58.
- [15] HELMRICH, S., ARIAS, A., LOCHEAD, G., WINTERMANTEL, T., BUCHHOLD, M., DIEHL, S. and WHITLOCK, S. (2020). Signatures of self-organized criticality in an ultracold atomic gas. *Nature* **577** 481–486.
- [16] HOFFMAN, C., JOHNSON, T. and JUNGE, M. (2017). Recurrence and transience for the frog model on trees. *The Annals of Probability* **45** 2826–2854.
- [17] HOFFMAN, C., RICHEY, J. and ROLLA, L. T. (2020). Active phase for activated random walk on \mathbb{Z} . *arXiv preprint arXiv:2009.09491*.
- [18] JENSEN, H. J. (1998). *Self-organized criticality: emergent complex behavior in physical and biological systems* **10**. Cambridge university press.
- [19] JERISON, D., LEVINE, L. and SHEFFIELD, S. (2012). Logarithmic fluctuations for internal DLA. *J. Amer. Math. Soc.* **25** 271–301. <https://doi.org/10.1090/S0894-0347-2011-00716-9> MR2833484
- [20] JOHNSON, T. and JUNGE, M. (2018). Stochastic orders and the frog model. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* **54** 1013–1030.
- [21] KESTEN, H. and SIDORAVICIUS, V. (2005). The spread of a rumor or infection in a moving population. *The annals of probability* **33** 2402–2462.
- [22] LEVINE, L. and SILVESTRI, V. (2021). How far do activated random walkers spread from a single source? *J. Stat. Phys.* **185** Paper No. 18, 27. <https://doi.org/10.1007/s10955-021-02836-9> MR4334780
- [23] LYONS, R. and PERES, Y. (2016). *Probability on trees and networks*. Cambridge Series in Statistical and Probabilistic Mathematics **42**. Cambridge University Press, New York. <https://doi.org/10.1017/9781316672815> MR3616205
- [24] MANNA, S. S. (1991). Two-state model of self-organized criticality. *Journal of Physics A: Mathematical and General* **24** L363.
- [25] ROLLA, L. T. (2008). Generalized hammersley process and phase transition for activated random walk models. *PhD thesis, arXiv:0812.2473*.
- [26] ROLLA, L. T. (2020). Activated random walks on \mathbb{Z}^d . *Probab. Surv.* **17** 478–544. <https://doi.org/10.1214/19-PS339> MR4152668
- [27] ROLLA, L. T. and SIDORAVICIUS, V. (2012). Absorbing-state phase transition for driven-dissipative stochastic dynamics on \mathbb{Z} . *Invent. Math.* **188** 127–150. <https://doi.org/10.1007/s00222-011-0344-5> MR2897694
- [28] ROLLA, L. T., SIDORAVICIUS, V. and ZINDY, O. (2019). Universality and sharpness in activated random walks. *Ann. Henri Poincaré* **20** 1823–1835. <https://doi.org/10.1007/s00023-019-00797-0> MR3956161
- [29] ROLLA, L. T. and TOURNIER, L. (2018). Non-fixation for biased activated random walks. *Ann. Inst. Henri Poincaré Probab. Stat.* **54** 938–951. <https://doi.org/10.1214/17-AIHP827> MR3795072

- [30] SHELLEF, E. (2010). Nonfixation for activated random walks. *ALEA Lat. Am. J. Probab. Math. Stat.* **7** 137–149. [MR2651824](#)
- [31] SIDORAVICIUS, V. and TEIXEIRA, A. (2017). Absorbing-state transition for stochastic sandpiles and activated random walks. *Electron. J. Probab.* **22** Paper No. 33, 35. <https://doi.org/10.1214/17-EJP50> [MR3646059](#)
- [32] STAUFFER, A. and TAGGI, L. (2018). Critical density of activated random walks on transitive graphs. *Ann. Probab.* **46** 2190–2220. <https://doi.org/10.1214/17-AOP1224> [MR3813989](#)
- [33] TAGGI, L. (2016). Absorbing-state phase transition in biased activated random walk. *Electron. J. Probab.* **21** Paper No. 13, 15. <https://doi.org/10.1214/16-EJP4275> [MR3485355](#)
- [34] TAGGI, L. (2019). Active phase for activated random walks on \mathbb{Z}^d , $d \geq 3$, with density less than one and arbitrary sleeping rate. *Ann. Inst. Henri Poincaré Probab. Stat.* **55** 1751–1764. <https://doi.org/10.1214/18-aihp933> [MR4010950](#)