Macroscopic flow out of a segment for Activated Random Walks in dimension 1

Nicolas Forien

Macroscopic flow out of a segment for Activated Random Walks in dimension 1

Nicolas Forien*

May 29, 2024

Abstract

Activated Random Walk is a system of interacting particles which presents a phase transition and a conjectured phenomenon of self-organized criticality. In this note, we prove that, in dimension 1, in the supercritical case, when a segment is stabilized with particles being killed when they jump out of the segment, a positive fraction of the particles leaves the segment with positive probability.

This was already known to be a sufficient condition for being in the active phase of the model, and the result of this paper shows that this condition is also necessary, except maybe precisely at the critical point. This result can also be seen as a partial answer to some of the many conjectures which connect the different points of view on the phase transition of the model.

Keywords and phrases. Activated random walks, phase transition, self-organized criticality.

MSC 2020 *subject classifications.* 60K35, 82B26.

1 Introduction

We begin with a brief informal presentation of some aspects of the model. The reader who is familiar with Activated Random Walks may skip the two following subsections, passing directly to Subsection 1.3 where our results are presented. Some illustrated sketches of proofs are given in Subsection 1.5.

1.1 Presentation of the model

The model of Activated Random Walks consists of particles performing independent random walks on a graph, which fall asleep with a certain rate and get reactivated in the presence of other particles on the same site. The model was popularized by Rolla, Sidoravicius and Dickman [Rol08, RS12, DRS10], and can be seen as a variant of the frog model [AMP02a, AMP02b]. Its study is motivated by its connection with the concept of self-organized criticality, which was introduced by the physicists Bak, Tang and Wiesenfeld [BTW87] to describe physical systems which present a critical-like behaviour but without the need to tune the parameters of the system to particular values (like is the case for an ordinary phase transition). To illustrate this concept, Bak, Tang and Wiesenfeld introduced an interacting particle system called the Abelian sandpile model, which shares some features with Activated Random Walks. We refer to [LL21, BS22] for a comparison of the two models in particular for what concerns their mixing properties, indicating that the model of Activated Random Walks mixes faster. This can explain why Activated Random Walks are expected to have a behaviour which is more universal, in that it is less sensitive to microscopic details of the system.

Let us now define informally the Activated Random Walk model on $\mathbb{Z}^d$. A configuration of the model consists of a certain number of particles on each site of $\mathbb{Z}^d$, each of these particles being in one of two possible states: active or sleeping.

The model evolves as follows. Each active particle performs a continuous-time random walk on $\mathbb{Z}^d$ with jump rate 1, with a certain translation-invariant jump distribution $p : \mathbb{Z}^d \to [0, 1]$. This means that, after a

*CEREMADE, CNRS, Université Paris-Dauphine, Université PSL, 75016 Paris, France

random time distributed as an exponential with parameter 1, the active particle at x jumps to some other site, the probability of jumping from x to y being $p(y-x)$.

In parallel, each active particle also carries another exponential clock with a certain parameter $\lambda > 0$ and, when this clock rings, if there are no other particles on the same site, the particle falls asleep (otherwise, if the particle is not alone, nothing happens). A sleeping particle stops moving (its continuous-time random walk is somewhat paused), until it wakes up, which happens when another particle arrives on the same site. Then, the reactivated particle resumes its continuous-time random walk with jump rate 1. Equivalently, one may also consider that a particle can fall asleep even when it is not alone on a site but, whenever this happens, the particle is instantaneously waken up by the presence of the other particles.

Thus, there can never be two sleeping particles at a same site. Hence, at every time $t \geqslant 0$, the configuration of the model at time t can by encoded into a function $\eta_t : \mathbb{Z}^d \to \mathbb{N} \cup \{\mathfrak{s}\}$, where $\eta_t(x) = k \in \mathbb{N}$ means that there are k active particles at the site x, while $\eta_t(x) = \mathfrak{s}$ means that there is one sleeping particle at x. Note that, with this notation, particles are indistinguishable: we only keep track of the number and states of particles on each site, but not of the individual trajectory of each particle.

Regarding the initial configuration η_0, various setups are interesting to consider. One possibility is to take η_0 which follows a translation-invariant and ergodic probability distribution on the set of all possible configurations, with a finite mean number of particles per site. This initial configuration may have only active particles, or both active and sleeping particles. Another case of interest is that of η_0 with only finitely many particles, for example n particles on the origin. The model can also be defined on different graphs, or with slight modifications of the dynamics, like for example adding a sink vertex where particles get trapped forever.

For a rigorous construction of the process $(\eta_t)_{t\geqslant 0}$, we refer the reader to [RS12], or to the review [Rol20]. See also [LS24] for a presentation of various different settings of interest and fascinating conjectures connecting these different points of view on the model.

1.2 Phase transition

Let us consider the model on $\mathbb{Z}^d$ starting with η_0 following a translation-ergodic distribution with mean particle density ζ. Depending on the sleep rate λ, on the jump distribution p and on the particle density ζ, the model can exhibit very different behaviours. A natural question is: if we start with only active particles, do they eventually all fall asleep, or is activity maintained forever? Note that, on the infinite lattice $\mathbb{Z}^d$, almost surely there exists no finite time when all the particles are sleeping. However, we have the following notion of fixation: we say that the system fixates if the origin is visited finitely many times by an active particle during the evolution of the process. Then, up to events of 0 probability, fixation is equivalent to the configuration on every finite set of $\mathbb{Z}^d$ eventually being constant and with only sleeping particles or, if we follow the trajectory of each individual particle, fixation also turns out to be equivalent to each particle walking only a finite number of steps, or to one given particle walking only a finite number of steps [AGG10]. If the system does not fixate, we say that the system stays active.

Due to the ergodicity assumption on η_0, the probability of fixation can only be 0 or 1 (see [RS12]). Thus, we can have two different regimes, depending on the sleep rate λ, on the jump distribution p and on the law of the initial configuration: either the system almost surely fixates (this regime is called the fixating phase, or stable phase), or the system almost surely stays active (this is called the active phase, or exploding phase). Moreover, we have the following key result about this phase transition:

Theorem 1 ([RSZ19])**.** *In any dimension $d \geqslant 1$, for every sleep rate $\lambda \in (0,\infty]$ and every jump distribution p which generates all $\mathbb{Z}^d$, there exists ζ_c such that, for every translation-ergodic initial distribution with no sleeping particles and an average density of active particles ζ, the Activated Random Walk model on $\mathbb{Z}^d$ with sleep rate λ almost surely fixates if $\zeta < \zeta_c$, whereas it almost surely stays active if $\zeta > \zeta_c$.*

This result shows in particular that the critical density is in some sense universal, in that it depends on the initial configuration only through the mean density of particles ζ. Thus, to study this critical density, it is enough to consider the particular case where the configuration is i.i.d., with a given probability distribution on $\mathbb{N}$ with finite mean.

An important challenge in the study of this phase transition is to relate the property of fixation, which concerns the model on the infinite lattice with infinitely many particles, to some finite counterparts of the model. A key example is the sufficient condition for activity given by Theorem 2 below.

On the finite box $V_n = (-n/2,\, n/2\,]^d \cap \mathbb{Z}^d$, let us consider a variant of the model where particles are killed and removed from the system when they jump out of V_n (or equivalently, we can consider the model on $\mathbb{Z}^d$ where particles are frozen outside V_n, so that particles which start out of V_n or which jump out of V_n are frozen forever and cannot move any more). Let M_n count the number of particles that jump out of V_n when we let this system evolve until all the sites of V_n become stable (a site x is called stable if it is either empty or it contains a sleeping particle).

Theorem 2 ([RT18])**.** *With the notation defined above, for every sleep rate λ and every jump distribution p which generates all $\mathbb{Z}^d$, if the initial configuration η_0 is i.i.d. and if*

$$\limsup_{n\to\infty} \frac{\mathbb{E}\big[M_n\big]}{|V_n|} \;>\; 0\,, \tag{1}$$

then the model on $\mathbb{Z}^d$ with sleep rate λ, jump distribution p and initial configuration η_0 almost surely stays active.

This result relies on the following intuitive idea: if with positive probability a large box looses a positive fraction of its particles during stabilization, then a particle starting at the origin in the model on $\mathbb{Z}^d$ has a positive probability of walking arbitrarily far away, which shows that the system stays active with positive probability, and thus with probability 1 (because we have a 0-1 law).

1.3 Main results

The main result of this paper consists in the addition of a reciprocal to the implication of Theorem 2, in the particular case of dimension 1. Recall that M_n denotes the number of particles that jump out of $V_n = (-n/2,\, n/2\,] \cap \mathbb{Z}$ during the stabilization of V_n with particles being killed upon leaving V_n.

Theorem 3. *In dimension $d = 1$, for every sleep rate $\lambda > 0$ and every nearest-neighbour jump distribution p, if the initial configuration η_0 is i.i.d. with mean ζ and all particles are initially active, then we have the equivalence:*

$$\zeta \;>\; \zeta_c \quad\iff\quad \liminf_{n\to\infty} \frac{\mathbb{E}[M_n]}{n} \;>\; 0\,.$$

This shows that the sufficient condition for activity given by Theorem 2 is also necessary, except maybe exactly at the critical point. Indeed, very few things are known rigorously about the critical regime $\zeta = \zeta_c$, with the exception of the particular case of directed walks in dimension 1 starting with η_0 i.i.d., for which a proof of non-fixation at criticality, due to Hoffman and Sidoravicius, appears in [CRS14].

Theorem 3 answers a conjecture of Levine and Silvestri [LS24] (in the particular case of dimension 1), showing that the density ζ_w that they define in Section 5.3, and which corresponds to the infimum of the ζ for which condition (1) holds when η_0 is i.i.d. Poisson, is in fact equal to the critical density ζ_c.

Our result is made more precise by the following theorem, which indicates an explicit positive fraction which exits with positive probability, as a function of the sleep rate λ and the density ζ. For every deterministic initial configuration $\eta : V_n \to \mathbb{N}$ (with only active particles), let us denote by $\|\eta\| = \sum_{x\in V_n} \eta(x)$ the total number of particles in the configuration η, and let us write $\mathbb{P}_\eta$ for the probability relative to the system started with deterministic initial configuration equal to η.

Theorem 4. *In dimension $d = 1$, for every sleep rate $\lambda > 0$ and every nearest-neighbour jump distribution p, for every $\zeta > \zeta_c$ we have*

$$\forall\, \varepsilon \in \left[0,\, \frac{\lambda(\zeta-\zeta_c)}{4(1+\lambda)\zeta_c}\right) \qquad \liminf_{n\to\infty} \inf_{\substack{\eta:V_n\to\mathbb{N}:\\ \|\eta\|\geqslant \zeta n}} \mathbb{P}_\eta\big(M_n > \varepsilon n\big) \;\geqslant\; 1 - \frac{\zeta_c}{\zeta}\left(1 + \frac{4(1+\lambda)\varepsilon}{\lambda}\right) \;>\; 0\,.$$

To show this, as a first step we prove the following result, which gives an explicit upper bound on the probability that no particle exits during stabilization: This bound is not optimal but, as explained later, it allows us to obtain the bound of Theorem 4.

Theorem 5. *In dimension $d = 1$, for every sleep rate $\lambda > 0$ and every nearest-neighbour jump distribution p, for every $\zeta > \zeta_c$, for every $n \geqslant 1$ and every initial configuration $\eta : V_n \to \mathbb{N}$ with $\|\eta\| \geqslant \zeta n$ particles, initially all active, we have*

$$\mathbb{P}_\eta\big(M_n = 0\big) \ \leqslant\ \frac{\zeta_c}{\zeta}\,.$$

Let us stress that the bounds of Theorem 4 and 5 hold for any deterministic initial configuration with at least ζn particles, which includes in particular the case of interest where all the particles start from the origin (see comments about this in Section 1.4).

Lastly, in the course of the proof of Theorem 3, to deal with the case $\zeta = \zeta_c$, we establish the following fact, which might be of independent interest:

Proposition 1. *In any dimension $d \geqslant 1$, for every sleep rate $\lambda > 0$ and every jump distribution p on $\mathbb{Z}^d$ whose support generates all the group $\mathbb{Z}^d$, if η_0 is i.i.d. with mean ζ_c then we have*

$$\lim_{n\to\infty} \frac{\mathbb{E}[M_n]}{|V_n|} \ =\ 0\,.$$

Notably, this last result shows that, at least in the case of directed walks in dimension 1, for which it is known that there is no fixation at criticality (see the remark following the statement of Theorem 3), the sufficient condition for activity given by Theorem 2 is not necessary.

1.4 Some perspectives

Since the seminal works which established general properties of the phase transition, various techniques have been developed to study Activated Random Walks. In particular, a series of works [RS12, ST18, ARS22, Tag19, BGH18, HRR23, FG22, Hu22, AFG22] established that the critical density is always strictly between 0 and 1 and obtained bounds on ζ_c as a function of λ. But many of the techniques used only work far from criticality, when the density is either much larger or much smaller than ζ_c, and few results have been proved to hold up to the critical density. For example, [BGHR19] shows that the model on the torus stabilizes fast when ζ is very small, and slowly when ζ is close to 1, but we lack sharper results about a transition exactly at ζ_c from fast to slow stabilization.

Some exceptions giving insight about the behaviour at or close to ζ_c are the study of the critical regime in the case of directed walks in one dimension (see [CRS14], with an argument due to Hoffman and Sidoravicius, and [CR21]), the continuity of ζ_c as a function of λ [Tag23], and the recent work [JMT23] which considers the model on the complete graph and computes the exact value of the critical density.

In this regard, the results of the present article have the merit to hold up to the critical point. However, the bounds presented here are far from being optimal, and there remains a lot of space for improvement. For example, Theorem 3 can be seen as a partial answer to the so-called hockey stick conjecture (conjecture 17 in [LS24]), which predicts that $M_n/|V_n|$ should converge in probability to $\max(0, \zeta - \zeta_c)$, at least in the particular case when the initial distribution is i.i.d. Poisson.

Similarly, the bound given in Theorem 5 is not optimal, and it is expected that when $\zeta > \zeta_c$, the probability that $M_n = 0$ in fact decays exponentially fast with n (see conjecture 20 of [LS24]).

Note that Theorem 5 can also be seen as a partial answer to the so-called ball conjecture (see conjectures 1 and 12 in [LS24]). This conjecture predicts that, when starting with n particles at the origin, if we let these particles stabilize in $\mathbb{Z}^d$, the random set of visited sites A_n is such that, for every $\varepsilon > 0$, with probability tending to 1 as $n \to \infty$, the set A_n contains all the sites of $\mathbb{Z}^d$ that belong to the origin-centred Euclidean ball of volume $(1-\varepsilon)n/\zeta_c$ and is contained in the origin-centred Euclidean ball of volume $(1+\varepsilon)n/\zeta_c$. Theorem 5 implies that the probability that A_n is included into the ball of volume $(1-\varepsilon)n/\zeta_c$ is less than $1-\varepsilon$. Note that another partial result was obtained in this direction in [LS21], also in dimension 1, showing an inner and an outer bound on A_n.

Last but not least, the proofs of the present paper are very specific to the one-dimensional case, and it would be interesting to obtain at least similar results in higher dimension.

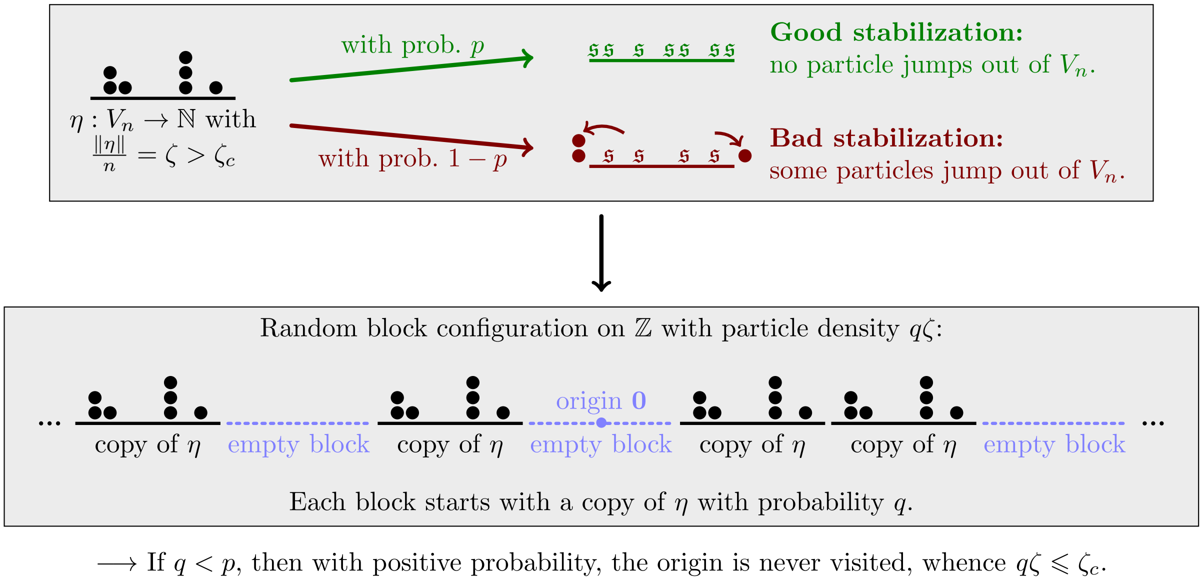

Figure 1: Outline of the proof of Theorem 5: assuming that a given deterministic initial configuration $\eta : V_n \to \mathbb{N}$ produces a good stabilization with probability p, we build a block configuration on $\mathbb{Z}$ with density $q\zeta$ which fixates if $q < p$, showing that $q\zeta \leqslant \zeta_c$. Since this holds for every $q < p$, we get $p\zeta \leqslant \zeta_c$.

1.5 Sketches of the proofs

Let us now summarize the strategy of the proofs. We start with Theorem 5, before explaining how the other results follow.

1.5.1 Sketch of the proof of Theorem 5:

Let $\eta : V_n \to \mathbb{N}$ with $\|\eta\|/n = \zeta > \zeta_c$, and let $p = \mathbb{P}_\eta(M_n = 0)$. The main idea is that, if p is too high, then we can construct a configuration on $\mathbb{Z}$ which fixates but which has a supercritical density, which is a contradiction.

More precisely, we consider a configuration on $\mathbb{Z}$ which is obtained by placing a copy of the configuration η inside each block of length n. When we stabilize each block separately in this periodic configuration, typically a fraction $1-p$ of the blocks have particles jumping out. We call these bad stabilizations. This is a problem because these bad stabilizations can disturb the neighbouring blocks. To counter this, we kill some blocks by deciding that a fraction q of the blocks start empty instead of starting with configuration η inside the block (for each block, we draw an independent Bernoulli variable with parameter q to decide whether the block starts empty or with a copy of η). Thus, we consider the initial configuration represented on Figure 1.

The idea is to take $q < p$, so that the fraction of blocks which start empty is strictly larger than the fraction of blocks whose stabilization is bad. Then we can show that, with positive probability, the blocks can be stabilized sequentially with no particle ever visiting the origin.

This implies that the particle density of the initial configuration, namely $q\zeta$, can be at most ζ_c (sweeping under the carpet the small issue that the configuration that we constructed on $\mathbb{Z}$ is not exactly translation-invariant, but this problem is easy to overcome, see Section 3.5). This being true for every $q < p$, we obtain $p\zeta \leqslant \zeta_c$.

To show that with positive probability the blocks can be stabilized sequentially without visiting the origin, we adapt the strategy that was used in [RS12] to prove that, in dimension 1, if $\zeta < \lambda/(1+\lambda)$, with positive probability the origin is never visited. This strategy consists in stabilizing each particle as close as possible to the origin, by moving it with acceptable topplings (see Section 2.1 for the definition of acceptable topplings) and exploring instructions in advance to discover where is the site closest to the origin where the particle can manage to fall asleep without waking up the particles already sleeping. In our context, we

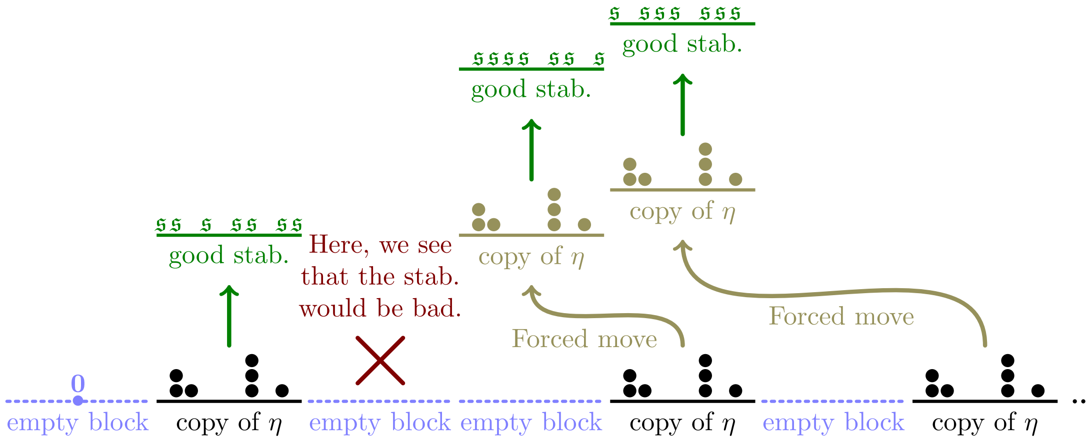

Figure 2: Strategy used in the proof of Theorem 5 to stabilize the blocks on one side of the origin without ever visiting the origin: each block is translated by means of acceptable topplings, to a place where it performs a good stabilization, as close as possible to the place where the previous block was stabilized. To decide where each block should be translated, we explore the instructions in advance to see where the stabilization will be good.

stabilize each block as close as possible to the origin, translating it and looking for the place closest to the other stabilized blocks where the block can manage to do a good stabilization, i.e., with no particle jumping out of the block (see Figure 2).

In some sense, our proof strategy amounts to building a coupling between the stabilization of the blocks on one side of the origin and an Activated Random Walk model on $\mathbb{N}$ whose sites are the blocks of the configuration. This coarse-grained model starts with a configuration which is i.i.d. Bernoulli with parameter q, each particle corresponding to a non-empty block. When we stabilize a block and no particles exit (a good stabilization), this corresponds to a sleep instruction for the block, thus the coarse-grained model has a sleep rate λ given by $p = \lambda/(1+\lambda)$. If the stabilization of the block is bad (i.e., particles jump out of the block) and if the neighbouring block in the direction of the origin is empty, then we can force the particles to move to this neighbouring block, using acceptable topplings, without disturbing the other blocks already stabilized around the origin. This means that the jump distribution of our coarse-grained model on $\mathbb{N}$ only has jumps to the left.

Then, the fact that our configuration with blocks fixates with positive probability follows from the fact that the Activated Random Walk model with a density strictly smaller than $\lambda/(1+\lambda)$ fixates whatever the jump distribution.

The proof of Theorem 5 is detailed in Section 3, and further illustrated in Figure 4.

Proving that Theorem 5 implies Theorem 4: To deduce Theorem 4 from Theorem 5, the idea is that, if less than εn particles jump out of the segment, then, taking a slightly larger segment, with enough empty space around to easily accommodate these particles which jumped out, we can stabilize this larger segment with no particles jumping out of it.

We proceed in two steps, as represented in Figure 3. We start with a configuration $\eta : V_n \to \mathbb{N}$ with density $\|\eta\|/n = \zeta > \zeta_c$ and empty strips of length $2\alpha n$ on each side of V_n.

During the first step, we stabilize $V_{n+2\alpha n}$, ignoring the sleeps out of V_n. In other words, we perform legal toppings in V_n and acceptable topplings in $V_{n+2\alpha n} \setminus V_n$, until the configuration is stable in V_n and empty in $V_{n+2\alpha n} \setminus V_n$. The number of particles which jump out of $V_{n+2\alpha n}$ during this step, that we denote by M'_n, may differ from M_n, which is the number of particles jumping out of V_n when we just stabilize V_n. Yet, it turns out that M'_n is stochastically dominated by M_n: this is the content of Lemma 3.

Then, during the second step, we adapt the trapping procedure of [RS12] to try to stabilize these M'_n particles in $V_{n+4\alpha n} \setminus V_n$. This procedure is said to be successful if, doing so, no particles jump out of $V_{n+4\alpha n}$ or

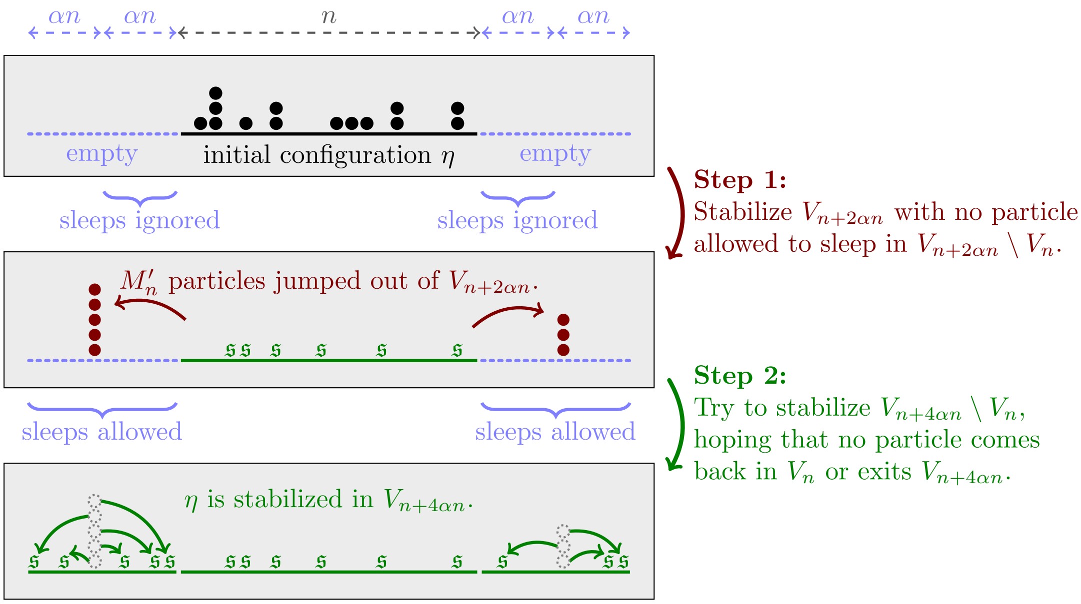

Figure 3: The two steps of the proof of Theorem 4.

come back in V_n. In this case, we obtain an acceptable toppling sequence which stabilizes the configuration η in $V_{n+4\alpha n}$, with no particle jumping out of $V_{n+4\alpha n}$, which entails that $M_{n+4\alpha n} = 0$.

Note that the overall density in the enlarged segment with the empty spaces around is $\zeta/(1+4\alpha)$. Thus, if α is chosen small enough so that $\zeta/(1+4\alpha) > \zeta_c$, then Theorem 5 gives an upper bound on the probability that $M_{n+4\alpha n} = 0$.

Then, if we consider $\varepsilon > 0$ small enough so that $\varepsilon/\alpha < \lambda/(1+\lambda)$, then we can show that the second stage succeeds with high probability, conditioned on the event that $M'_n \leqslant \varepsilon n$. Thus, we can translate the upper bound on $\mathbb{P}_\eta\big(M_{n+4\alpha n} = 0\big)$ coming from Theorem 5 into an upper bound on $\mathbb{P}_\eta(M'_n \leqslant \varepsilon n)$. The stochastic domination given by Lemma 3 then allows us to translate this into the claimed upper bound on $\mathbb{P}_\eta(M_n \leqslant \varepsilon n)$.

The proof of Theorem 4 is the object of Section 4.

Obtaining Theorem 3 and Proposition 1: Given Theorem 4, our Theorem 3 easily follows. The only detail is that, to show that we have indeed an equivalence, there remains to show Proposition 1, which states that, at criticality, there is not a positive fraction which jumps out of the box. This is presented in Section 5.

2 Some useful tools

We now describe the site-wise representation of the model, with an array of sleep and jump instructions above the sites. We refer to the survey [Rol20] for a more detailed presentation.

2.1 The site-wise construction of the model

A key ingredient in the study of Activated Random Walks is the site-wise representation, also known as Diaconis-Fulton representation [DF91, RS12]. Let $\lambda > 0$, let $p : \mathbb{Z}^d \to [0,1]$ be a jump distribution, and let $\eta_0 : \mathbb{Z}^d \to \mathbb{N} \cup \{\mathfrak{s}\}$ be the (possibly random) initial configuration.

Let us consider an array of i.i.d. variables $\tau = (\tau_{x,j})_{x\in\mathbb{Z}^d,\, j\geqslant 1}$ where, for every $x \in \mathbb{Z}^d$ and every $j \geqslant 1$, the variable $\tau_{x,j}$ is an instruction which can be either a sleep instruction, with probability $\lambda/(1+\lambda)$, or a jump instruction to some site $y \in \mathbb{Z}^d$, with probability $p(y-x)/(1+\lambda)$.

The idea is that we can construct the evolution of the system by looking at these instructions each time that something happens at some site. As we use instructions of the array, we keep track of which instructions have already been used, with the help of a function called the odometer $h : \mathbb{Z}^d \to \mathbb{N}$, which counts, at each site, how many instructions have already been used.

When we use an instruction at a site x, we say that we topple x. For a given fixed configuration η, we say that it is legal (respectively, acceptable) for η to topple a site x if x contains at least one active particle (respectively, at least one particle) in η.

If a toppling is legal or acceptable, then this toppling consists in using the next instruction $\tau_{x,\,h(x)+1}$ to update the configuration η: if this instruction is a sleep instruction, then the particle at x falls asleep if it is alone (whereas nothing happens if there are at least two particles at x), and if it is a jump instruction to another site y, one particle at x jumps to site y, waking up the sleeping particle there if there is one. If the toppling was only acceptable but not legal, we first wake up the particle at x before applying the toppling. The resulting configuration is denoted by $\tau_{x,\,h(x)+1}\eta$. Thus, for a fixed realization of the array τ, the toppling at a site x consists of an operator

$$\Phi_x \ : \ (\eta,\, h) \ \longmapsto \ (\tau_{x,\,h(x)+1}\eta,\, h+\delta_x)\,,$$

which is only defined if the toppling is acceptable.

If $\alpha = (x_1,\,\ldots,\,x_k)$ is a certain sequence of sites of $\mathbb{Z}^d$, we say that the toppling sequence α is legal (resp., acceptable) for $(\eta,\, h)$ if for every $i \in \{1,\,\ldots,\,k\}$, it is legal (resp., acceptable) for $\Phi_{x_{i-1}} \circ \cdots \circ \Phi_{x_2} \circ \Phi_{x_1}(\eta,\, h)$ to topple x_i, that is to say, if the configuration resulting from the first $i-1$ topplings has at least one active particle (resp., at least one particle) on the site x. If α is acceptable, applying the toppling sequence α means applying $\Phi_\alpha = \Phi_{x_k} \circ \cdots \circ \Phi_{x_1}$. We define the odometer of a toppling sequence α as $m_\alpha = \delta_{x_1} + \cdots + \delta_{x_k}$, which simply counts how many times each site appears in the sequence α. We also define, for every $V \subset \mathbb{Z}^d$,

$$m_{V,\,\eta} \ = \ \sup_{\alpha \subset V,\, \alpha \text{ legal}} m_\alpha\,, \tag{2}$$

where the notation $\alpha \subset V$ means that all the sites appearing in α must belong to V. The total stabilization odometer associated with the configuration η is defined as:

$$m_\eta \ = \ \sup_{V \subset \mathbb{Z}^d,\, V \text{ finite}} m_{V,\,\eta}\,. \tag{3}$$

This stabilization odometer may be infinite and, in fact, when the initial configuration η_0 follows a translation-ergodic distribution, we have (see [RS12])

$$\mathbb{P}\big(\text{the system fixates}\big) \ = \ \mathbb{P}\big(m_{\eta_0}(\mathbf{0}) < \infty\big)\,. \tag{4}$$

Thus, to know whether the system fixates or not, it is enough to look at this array of instructions and to determine whether $m_{\eta_0}(\mathbf{0})$ is finite or not.

2.2 Abelian property and the use of acceptable topplings

A key advantage of the site-wise construction is the following property, which states that the order with which we perform the topplings is irrelevant, allowing us to choose whatever convenient strategy to choose which sites to topple. We say that a sequence of topplings α stabilizes η in V if the configuration resulting from the application of the toppling sequence is stable in V, meaning that there are no active particles in V.

Lemma 1 (Abelian property, Lemma 2 in [RS12])**.** *If α and β are both legal toppling sequences for η that are contained in V and stabilize η in V, then $m_\alpha = m_\beta = m_{V,\,\eta}$ and the resulting configurations are equal.*

We also have the following monotonicity property, which shows that acceptable topplings may be used whenever one is looking for upper bounds on the legal odometer:

Lemma 2 (Lemma 2.1 in [Rol20])**.** *If α is an acceptable sequence of topplings that stabilizes η in V, and $\beta \subset V$ is a legal sequence of topplings for η, then $m_\alpha \geqslant m_\beta$. Thus, if α is an acceptable sequence of topplings that stabilizes η in V, then $m_\alpha \geqslant m_{V,\,\eta}$.*

3 Proof of Theorem 5: the probability that no particle exits

This section is devoted to the proof of Theorem 5. Let $\lambda > 0$, let p be a nearest-neighbour jump distribution, let $n \geqslant 1$ and consider a fixed deterministic initial configuration $\eta : V_n \to \mathbb{N}$ with $\|\eta\| = \zeta n$, where $\zeta > \zeta_c$. Let us write $p = \mathbb{P}_\eta\big(M_n = 0\big)$. Our aim is to show that $p \leqslant \zeta_c/\zeta$.

3.1 The block configuration

As explained in the sketch of proof in Section 1.5, the idea is to build a random initial distribution on $\mathbb{Z}$ which will fixate with positive probability. As a first step, we define a configuration which is not translation-invariant, and we will deal with this detail afterwards. Let $q \in (0,\, p)$ (we assume that $p \neq 0$, otherwise there is nothing to prove), and let $(X_i)_{i\in\mathbb{Z}}$ be i.i.d. Bernoulli variables with parameter q. We define a random initial configuration η_0 by letting, for every $x \in \mathbb{Z}$,

$$\eta_0(x) \;=\; X_{\lfloor x/n \rfloor} \times \eta\big(x - \lfloor x/n \rfloor\big)\,.$$

That is to say, the configuration η_0 consists in a bi-infinite sequence of blocks of length n, with each block being, with probability q, a translated copy of the deterministic configuration η and, with probability $1-q$, an empty block. For every $i \in \mathbb{Z}$, let us define the block number i as

$$B_i \;=\; V_n + in\,,$$

so that $\mathbb{Z} = \cup_{i\in\mathbb{Z}} B_i$, this union being disjoint. Also, it will be convenient to enumerate the occupied sites of η, counted with multiplicity: let us take $x_1,\,\ldots,\,x_{\|\eta\|} \in V_n$ such that

$$\eta \;=\; \sum_{\ell=1}^{\|\eta\|} \mathbf{1}_{\{x_\ell\}}\,. \tag{5}$$

3.2 Stabilization procedure

We now describe our toppling strategy. We first construct a procedure to stabilize the blocks $B_1 \cup \cdots \cup B_k$, for $k \geqslant 1$, using acceptable topplings. Recall that an acceptable toppling sequence is allowed to ignore some sleep instructions (or, equivalently, to move sleeping particles), and thus gives an upper bound on the odometer for legal stabilization (see Lemma 2). Recall that our stabilization procedure is inspired by the method used in [RS12], as explained in the sketch of the proof (see Section 1.5).

Aim of the procedure. We proceed step by step, stabilizing the blocks one after another, with at each step a certain probability that the procedure fails. For every $k \geqslant 1$, if the procedure is successful until step k, we will obtain an acceptable toppling sequence α_k which stabilizes the configuration $\eta_0 \mathbf{1}_{B_1\cup\cdots\cup B_k}$ in $\mathbb{Z}$, only performing topplings on blocks B_i with $i \geqslant 1$ (so that the origin is never toppled) and which is such that the resulting configuration has no particles inside the blocks B_i for $i > S_k$, for a certain integer $S_k \in \{0,\,\ldots,\,k\}$. The idea is that we try to stabilize the successive blocks as close as possible to the block B_0 (but without visiting this block B_0), and this index S_k indicates which is the rightmost box where we left some sleeping particles. Then, in the subsequent steps of the procedure we take care not to visit any more the blocks $B_1,\,\ldots,\,B_{S_k}$, so that we do not wake up the particles stabilized before. Note that the toppling sequence α_k not only stabilizes η_0 inside $B_1 \cup \cdots \cup B_k$, but stabilizes all the particles which start in these blocks.

Corrupted sites. The construction of this toppling sequence α_k may look at future instructions at the sites of the blocks labelled with $i < S_k$ and eventually choose not to use these instructions: we say that such sites are corrupted. But the sites of the blocks B_i with $i \geqslant S_k$ are not corrupted at step k, meaning that the procedure until step k is independent of the remaining instructions at these sites. These corrupted sites cause no problem because, after step k, we no longer perform topplings on the sites of the blocks $B_1,\,\ldots,\,B_{S_k}$.

Step 1 of the procedure. Step 1 of the procedure deals with the block B_1. If this block is empty, we do nothing and move on to the next block. In this case, we have $S_1 = 0$ and α_1 is an empty sequence. Otherwise, if this block B_1 is not empty (i.e., if $X_1 = 1$), then we stabilize the configuration inside B_1 with legal topplings inside B_1 until all sites of this block become stable. We say that this stabilization is a good stabilization if no particle jumps out of the block. If the stabilization of B_1 is good, then we let $S_1 = 1$ and we take α_1 to be a legal toppling sequence in B_1 which stabilizes the configuration in B_1, and we can proceed, moving on to step 2. Otherwise, if this first stabilization is not good, then we declare the procedure to fail and we stop everything. This was step 1 of the procedure, let us now describe step k for $k \geqslant 1$.

Step k of the procedure. For convenience, we let $S_0 = 0$ and we let α_0 be the empty sequence. Let $k \geqslant 1$ and assume that the procedure was successful until step $k-1$, yielding an acceptable toppling sequence α_{k-1} and an integer S_{k-1} as described above. We now describe step k. First case: if $X_k = 0$, meaning that the block B_k is empty, we simply let $\alpha_k = \alpha_{k-1}$ and $S_k = S_{k-1}$.

Second case: assume now that $X_k = 1$. We start with the odometer $m_{\alpha_{k-1}}$ of the instructions used in the previous steps. To stabilize the particles of the block B_k, we can use any of the blocks B_i for $S_{k-1} < i \leqslant k$, on which no sleeping particles are left after step $k-1$. Then, for every $i \in \{S_{k-1}+1, \ldots, k\}$, we do the following trial, which consists in translating the block B_k in block B_i and trying to stabilize B_i:

Recall the enumeration of η defined in (5). With acceptable topplings, we move one particle from the site $x_1 + kn$ until it reaches the site $x_1 + in$ (if $i = k$ there is nothing to do). We then repeat this operation for each particle: for every $\ell \in \{1, \ldots, \|\eta\|\}$, we move one particle from the site $x_\ell + kn$ with acceptable topplings, until it reaches the site $x_\ell + in$. Almost surely, after a finite number of acceptable topplings, we obtain a copy of the configuration η in the block B_i. Let us denote by $\beta_{k,i}$ the acceptable toppling sequence which corresponds to this move of the copy of η from block B_k to block B_i (for $i = k$, $\beta_{k,i}$ is empty). Note that this operation does not topple any site of the blocks B_j for $j < i$.

Then, we perform a stabilization of the block B_i with legal topplings inside B_i, which can be good or not. Let us denote by $\gamma_{k,i}$ a toppling sequence which corresponds to this legal stabilization.

We repeat this identical trial for every $i = S_{k-1}+1, \ldots, k$, in this increasing order. For each i, we start fresh again with the odometer $m_{\alpha_{k-1}}$ resulting from step $k-1$, so that instructions used for example in $\beta_{k,i-1}$ may be used again in $\beta_{k,i}$ (the different trials need not to be coherent). As soon as we find i such that the stabilization in block B_i is good, we stop the trials and we define:

$$S_k \;=\; \min\Big\{ i \in \{S_{k-1}+1, \ldots, k\} \;:\; \text{the stabilization of the block } B_i \text{ is good} \Big\}.$$

Recall that we say that the stabilization of a block is good if no particle jumps out of this block during its stabilization with legal topplings. If none of these $k - S_{k-1}$ trials leads to a good stabilization, we declare the procedure to fail and we stop there. Note that, for every i, the event $\{S_k = i\}$ is independent of the instructions in the blocks B_j for $j > i$, since we can interrupt the procedure as soon as we see a good stabilization happening, without needing to try the subsequent blocks. Here there is a small subtlety which lies in the fact that during the positioning phase to bring the particles from block B_k to block B_i, we perform topplings on the blocks B_j with $j > i$, but the event that the stabilization of the block B_i is good only depends on the odometer of $\beta_{k,i}$ inside B_i and not on the instructions used on the right of this block, since these instructions anyway almost surely have the effect of bringing each particle to the rightmost site of B_i, and making particles which eventually jump out of B_i by the right side come back inside B_i, so that we do not even need to reveal these instructions to see whether the stabilization in the block B_i will be good or not (these instructions that we do not need to reveal are those in the hatched region in Figure 4).

Then, if the procedure does not fail, we define α_k to be the concatenation of α_{k-1} followed by the two sequences β_{k,S_k} and γ_{k,S_k} defined above. By construction, this acceptable toppling sequence α_k stabilizes the configuration $\eta_0 \mathbf{1}_{B_1 \cup \cdots \cup B_k}$ in $\mathbb{Z}$, only performing topplings on blocks B_i with $i \geqslant 1$, and the resulting configuration has no particle in the blocks B_i for $i > S_k$, as required. See Figure 4 for an illustration of one step of the procedure.

3.3 Stabilizing both sides without visiting the origin

Imagine now that we want to stabilize the configuration η_0 inside the set $V_{(2k+1)n} = B_{-k} \cup \cdots \cup B_k$, for a certain $k \geqslant 1$. To do this, we may simply apply the above procedure until step k, then repeat the same

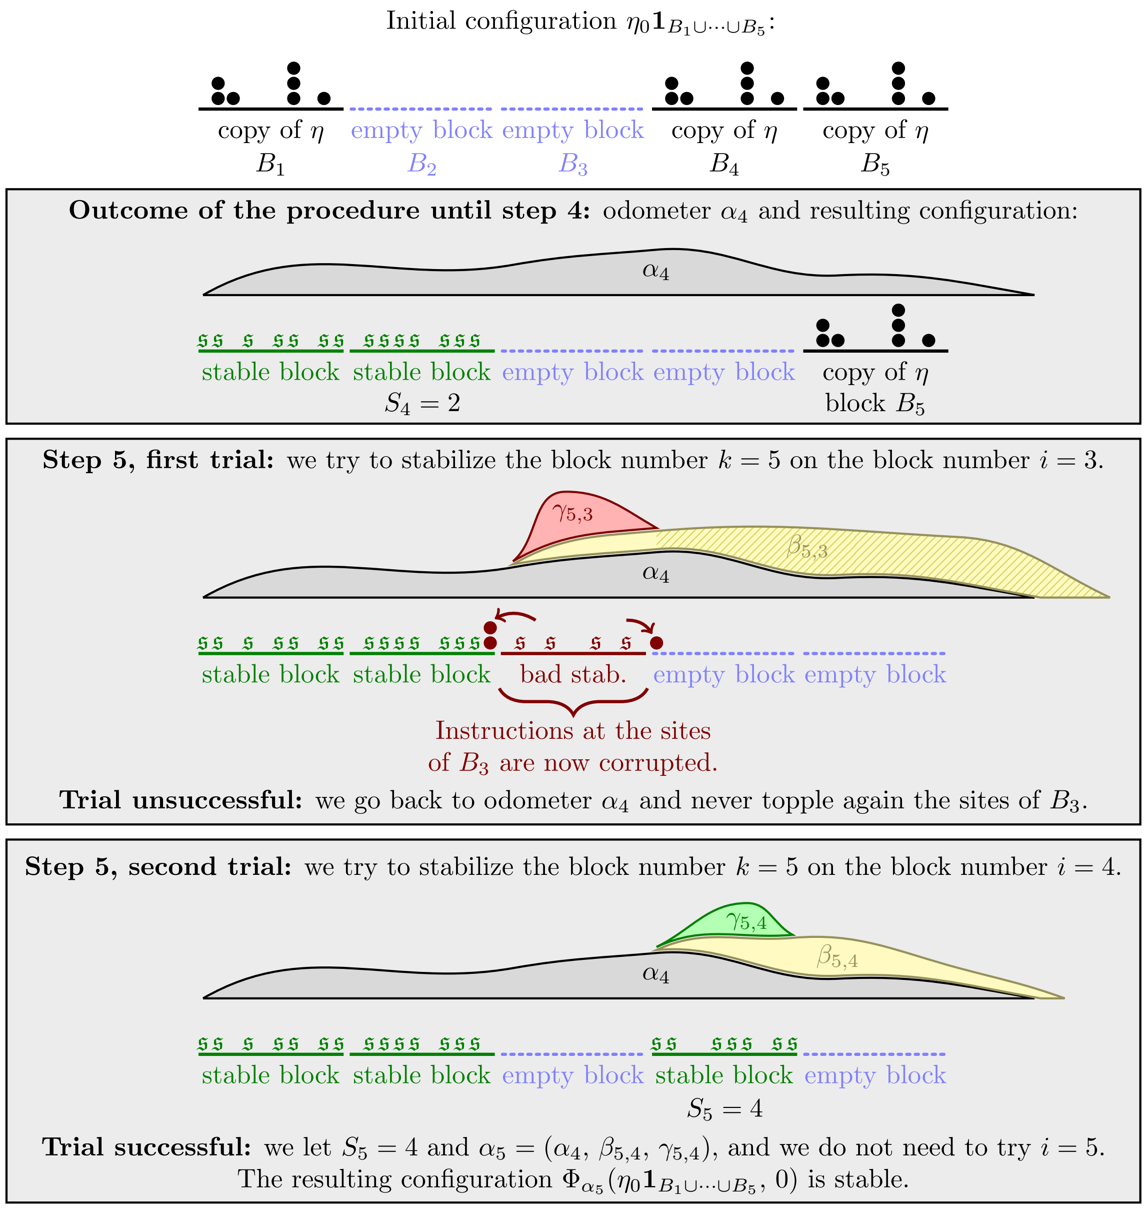

Figure 4: Step $k = 5$ of the procedure used to stabilize the blocks on the right side of the origin. This step consists in trying to stabilize the block B_5 on one of the blocks B_i for $S_4 < i \leqslant 5$. During the trial aimed at B_i, we move the particles from B_k to B_i, using an acceptable toppling sequence $\beta_{k,i}$ and we stabilize the block B_i with a legal toppling sequence $\gamma_{k,i}$. The trial is successful if this stabilization is a good stabilization. Note that, to know whether the first trial is successful, we need to reveal only the instructions on B_3, and not all the instructions of $\beta_{5,3}$. Therefore, the event that the first trail fails is independent of the instructions in the hatched part of $\beta_{5,3}$.

procedure on the negative blocks $B_{-1}, \ldots, B_{-k}$. Let us denote by $\mathcal{A}_k$ the event that the procedure on the positive blocks succeeds until step k, and let us denote by $\mathcal{A}_{-k}$ its counterpart for the negative blocks. If both procedures succeed and if moreover the block B_0 turns out to be initially empty, then we obtain an acceptable toppling sequence α which stabilizes the configuration $\eta_0 \mathbf{1}_{V_{(2k+1)n}}$ in $\mathbb{Z}$, and which is such that $m_\alpha(\mathbf{0}) = 0$. By the monotonicity property given by Lemma 2, this implies that $m_{\eta_0 \mathbf{1}_{V_{(2k+1)n}}}(\mathbf{0}) = 0$. Thus, we have

$$\mathbb{P}\big(m_{\eta_0 \mathbf{1}_{V_{(2k+1)n}}}(\mathbf{0}) = 0\big) \;\geqslant\; \mathbb{P}\Big(\big\{X_0 = 0\big\} \cap \mathcal{A}_k \cap \mathcal{A}_{-k}\Big) \;=\; (1-q)\,\mathbb{P}\big(\mathcal{A}_k\big)^2 .$$

Furthermore, recalling the definition (2) of $m_{V,\eta}$, note that we always have $m_{V,\eta} \leqslant m_{\eta \mathbf{1}_V}$, because any sequence $\alpha \subset V$ which is legal for η is also legal for $\eta \mathbf{1}_V$. Thus, we obtain

$$\mathbb{P}\big(m_{V_{(2k+1)n},\,\eta_0}(\mathbf{0}) = 0\big) \;\geqslant\; \mathbb{P}\big(m_{\eta_0 \mathbf{1}_{V_{(2k+1)n}}}(\mathbf{0}) = 0\big) \;\geqslant\; (1-q)\,\mathbb{P}\big(\mathcal{A}_k\big)^2 . \tag{6}$$

3.4 The success probability

Now recall that, at each step of the procedure, with probability $1-q$ we find an empty block, and otherwise we do trials of stabilization for each $i \in \{S_{k-1}+1, \ldots, k\}$, each trial being good with probability p, independently of everything else. Thus, for $k \geqslant 1$, for every $s \in \{0, \ldots, k-1\}$ and $s' \in \{s, \ldots, k\}$, we have

$$\mathbb{P}\big(\mathcal{A}_k \cap \{S_k = s'\} \;\big|\; \mathcal{A}_{k-1} \cap \{S_{k-1} = s\}\big) \;=\; \begin{cases} 1-q & \text{if } s' = s\,, \\ qp(1-p)^{s'-s-1} & \text{if } s' > s\,. \end{cases}$$

Therefore, for every sequence $0 = s_0 \leqslant s_1 \leqslant \cdots \leqslant s_k$ such that $s_i \leqslant i$ for every $i \leqslant k$, we have

$$\begin{aligned} \mathbb{P}\Big(\mathcal{A}_k \cap \big\{(S_1, \ldots, S_k) = (s_1, \ldots, s_k)\big\}\Big) \;&=\; \prod_{i=1}^{k} \big[(1-q)\mathbf{1}_{\{s_i = s_{i-1}\}} + qp(1-p)^{s_i - s_{i-1} - 1}\mathbf{1}_{\{s_i > s_{i-1}\}}\big] \\ &=\; \mathbb{P}\Big(\forall i \in \{1, \ldots, k\}\,,\; X_i G_i = s_i - s_{i-1}\Big)\,, \end{aligned}$$

where $(G_i)_{i \geqslant 1}$ are i.i.d. Geometric variables with parameter p, independent of $(X_i)_{i \geqslant 1}$. Thus, writing

$$S'_i \;=\; \sum_{j=1}^{i} X_j\, G_j$$

for every $i \geqslant 1$, we get

$$\mathbb{P}\big(\mathcal{A}_k\big) \;=\; \sum_{\substack{0 \leqslant s_1 \leqslant \cdots \leqslant s_k \\ \forall i \leqslant k\,,\; s_i \leqslant i}} \mathbb{P}\Big(\forall i \in \{1, \ldots, k\}\,,\; S'_i = s_i\Big) \;=\; \mathbb{P}\Big(\forall i \in \{1, \ldots, k\}\,,\; S'_i \leqslant i\Big)\,.$$

Plugging this into (6) and recalling the definition (3) of the odometer m_{η_0}, we get

$$\begin{aligned} \mathbb{P}\big(m_{\eta_0}(\mathbf{0}) = 0\big) \;&=\; \mathbb{P}\left(\bigcap_{k \geqslant 1} \big\{m_{V_{(2k+1)n}, \eta_0}(\mathbf{0}) = 0\big\}\right) \;=\; \lim_{k \to \infty} \mathbb{P}\big(m_{V_{(2k+1)n},\,\eta_0}(\mathbf{0}) = 0\big) \\ &=\; \lim_{k \to \infty} (1-q)\,\mathbb{P}\big(\mathcal{A}_k\big)^2 \;=\; (1-q)\,\mathbb{P}\Big(\forall i \geqslant 1\,,\; S'_i \leqslant i\Big)^2 . \end{aligned}$$

Yet, the law of large numbers ensures that, almost surely,

$$\frac{S'_i}{i} \;\overset{i \to \infty}{\longrightarrow}\; \mathbb{E}\big[X_1 G_1\big] \;=\; \frac{q}{p} \;<\; 1\,,$$

which implies that

$$\mathbb{P}\Big(\forall i \geqslant 1\,,\; S'_i \leqslant i\Big) \;>\; 0\,,$$

whence

$$\mathbb{P}\big(m_{\eta_0}(\mathbf{0}) = 0\big) \;>\; 0\,.$$

3.5 Making the initial configuration translation-invariant

Now, we would like to deduce that the density of particles in this configuration η_0 is at most ζ_c. But there remains to deal with a small issue: this configuration η_0 does not follow a translation-invariant distribution. To obtain a random initial configuration that is invariant by translation, we simply apply a translation by a random offset. Thus, we take Y a uniform variable in $V_n = (-n/2,\, n/2\,] \cap \mathbb{Z}$, independent of everything else, and we define the random initial configuration $\tilde{\eta}_0$ by writing, for every $x \in \mathbb{Z}$,

$$\tilde{\eta}_0(x) \;=\; \eta_0(x+Y)\,.$$

Then, this random initial configuration is translation-ergodic and has a density of particles

$$\mathbb{E}\big[\tilde{\eta}_0(\mathbf{0})\big] \;=\; \mathbb{E}\big[X_0\,\eta(Y)\big] \;=\; q \times \frac{\|\eta\|}{n} \;=\; q\zeta\,. \tag{7}$$

3.6 Conclusion

To conclude, we simply write

$$\mathbb{P}\big(m_{\tilde{\eta}_0}(\mathbf{0})=0\big) \;\geqslant\; \mathbb{P}\Big(\big\{Y=0\big\}\cap\big\{m_{\eta_0}(\mathbf{0})=0\big\}\Big) \;=\; \frac{1}{n}\,\mathbb{P}\big(m_{\eta_0}(\mathbf{0})=0\big) \;>\; 0\,.$$

Recalling the relation (4) between fixation and the odometer $m_{\tilde{\eta}_0}$, we deduce that the system with initial distribution $\tilde{\eta}_0$ fixates with positive probability. By virtue of Theorem 1, this implies that the configuration $\tilde{\eta}_0$ is not supercritical, that is to say, $\mathbb{E}[\tilde{\eta}_0(\mathbf{0})] \leqslant \zeta_c$. Given (7), we obtain that $q\zeta \leqslant \zeta_c$. This being true for every $q \in (0,\, p)$, we eventually deduce that $p\zeta \leqslant \zeta_c$, which concludes the proof of Theorem 5. □

4 Proof of Theorem 4: a fraction jumps out of the segment

The aim of this section is to deduce Theorem 4 from Theorem 5.

4.1 Preliminary: a no man's land around a segment

The following Lemma tells us that adding empty intervals around the segment V_n where particles are not allowed to sleep and stabilizing the configuration in V_n and in these intervals does not increase the number of particles which exit during stabilization, at least in distribution.

Lemma 3. *Let $\lambda > 0$ and let p be a nearest-neighbour jump distribution on $\mathbb{Z}$. Let $n \geqslant 1$, let $\eta : V_n \to \mathbb{N}$ be a fixed deterministic initial configuration on V_n with only active particles, and let $a,\, b \in \mathbb{Z}$ be such that $W = \{a,\, \ldots,\, b\} \supset V_n$. Starting from the initial configuration η and performing legal topplings in V_n and acceptable topplings in $W \setminus V_n$ until the resulting configuration is stable in V_n and empty in $W \setminus V_n$, we denote by M_n^W the number of particles which jump out of W. Then M_n stochastically dominates M_n^W.*

Note that the Lemma only gives a stochastic domination, and it is not always true that $M_n^W \leqslant M_n$ for a given array of instructions, since a particle which leaves V_n and comes back in V_n before jumping out of W may wake up many sleeping particles and cause more of them to leave W than if we just stabilize V_n.

Proof of Lemma 3. Let $\lambda,\, p,\, n,\, \eta$ be as in the statement. Since $M_n = M_n^W$ with $W = V_n$, it is enough to show that for any $a,\, b,\, a',\, b' \in \mathbb{Z}$ such that $V_n \subset W \subset W'$ where $W = \{a,\, \ldots,\, b\}$ and $W' = \{a',\, \ldots,\, b'\}$, we have that M_n^W stochastically dominates $M_n^{W'}$. Moreover, it is in fact enough to treat the case $b = b'$, because the general case then follows by applying two times the result. Thus, we consider $a' < a \leqslant 0$ and $b \geqslant n-1$, and we let $W = \{a,\, \ldots,\, b\}$ and $W' = \{a',\, \ldots,\, b\}$.

Now, we consider the following toppling strategy. First, we perform legal topplings in V_n and acceptable topplings in $W \setminus V_n$ until we reach a configuration which is stable in V_n and has no particles at all in $W \setminus V_n$. But the important point is that, doing so, we decide to always topple the leftmost particle that we can topple, that is to say, we topple either the leftmost non-empty site of W or the leftmost site containing at

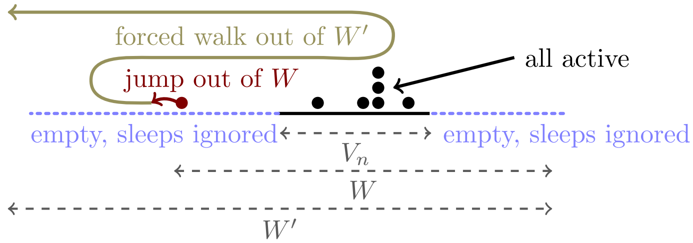

Figure 5: Strategy to prove Lemma 3, which shows that increasing the no man's lands regions around V_n leads to a decrease, in distribution, of the number of particles which jump out of the enlarged segment. The key point is that we always topple the leftmost possible particle, so that whenever a particle jumps out of W from the left, all the other particles are active, allowing us to force this particle to walk out of W' with no effect on the other particles.

least an active particle in V_n, always choosing the leftmost among these two sites. That way, all the sleeping particles always remain located on the left of the leftmost active particle. That is to say, at any time of the procedure, for any two sites $x,\, y \in W$ with $x < y$, it cannot be that x contains an active particle while y contains a sleeping one.

In the course of this stabilization, some particles may leave W by the left exit. As soon as one particle does so, we pause the procedure. Note that, at such a moment, we know that all the particles in W are active, because we just toppled the leftmost site of W and sleeping particles cannot be on the right site of this site (see Figure 5).

During this pause of the procedure, we move this particle which jumped out of W, using acceptable topplings, until it leaves W'. Doing so, this particle may visit some sites of W (and it may even visit all W, for example if it leaves W' via the right exit), but this is harmless because, as said above, all the particles are active at that moment. Once this particle jumped out of W', we resume the procedure, toppling again the leftmost possible site of W which contains a (necessarily active) particle, and we go on until we obtain a configuration which is stable inside V_n and empty in $W' \setminus V_n$. Let us call N the number of particles which jump out of W' during the whole process, which is also the number of particles which jumped out of W.

First, we can note that N is equal in distribution to M_n^W. Indeed, compared to the procedure defining M_n^W, we just added pauses during which we take a particle which had jumped out of W and we move it with acceptable topplings until it leaves W', on top of the configuration which only has active particles. Hence, inside W, the situation is the same just before and just after a pause, the only change being that some instructions have been used, resulting in a change of the odometer.

Since we force every particle which jumps out of W to jump out of W', what we obtain is therefore an acceptable toppling procedure in W', whose outcome is a stable configuration in V_n and empty sites in $W' \setminus V_n$, and which is such that N, the number of particles which jump out of W' is equal in distribution to M_n^W.

Now note that, for a fixed array of instructions, the number of particles which jump out of W' when applying a toppling procedure is an increasing function of the odometer of this procedure. Thus, by the monotonicity property given by Lemma 2, we deduce that $N \geqslant M_n^{W'}$. Thus, we eventually obtain the claimed stochastic domination. $\square$

4.2 If few particles jump out, then no one leaves a slightly larger segment

We now prove the following lower bound on the cost to stabilize in $V_m \setminus V_n$ all the particles which jump out of V_n:

Lemma 4. *In dimension $d = 1$, for every $\lambda > 0$ and every nearest-neighbour jump distribution p, for every $n \geqslant 1$ and every deterministic initial configuration $\eta : V_n \to \mathbb{N}$, for any $k,\, \ell \in \mathbb{N}$, we have*

$$\mathbb{P}_\eta\big(M_{n+4\ell} = 0\big) \ \geqslant\ \mathbb{P}_\eta\big(M_n \leqslant k\big) \times \mathbb{P}\big(G_1 + \cdots + G_k \leqslant \ell\big)\,,$$

where $(G_j)_{j\geqslant 1}$ are i.i.d. Geometric variables with parameter $\lambda/(1+\lambda)$.

Note that, when we write $\mathbb{P}_\eta\big(M_{n+4\ell} = 0\big)$, we implicitly extend the configuration $\eta : V_n \to \mathbb{N}$ to the configuration on $V_{n+4\ell}$ which coincides with η on V_n and has only zeros on $V_{n+4\ell} \setminus V_n$.

Proof. Let λ, p, n, η, k, ℓ and $(G_j)_{j\geqslant 1}$ be as in the statement. To stabilize η in $V_{n+4\ell}$, we proceed in two steps, as explained in the sketch of the proof in Section 1.5.

First, we perform legal topplings in V_n and acceptable topplings in $V_{n+2\ell} \setminus V_n$ until all the sites of V_n are stable and all the sites of $V_{n+2\ell} \setminus V_n$ are empty. Let us denote by M'_n the number of particles which jump out of $V_{n+2\ell}$ during this step. It follows from Lemma 3 that M_n stochastically dominates M'_n, which, with the notation of Lemma 3, corresponds to $M_n^{V_{m+2\ell}}$. Thus, we have

$$\mathbb{P}_\eta\big(M'_n \leqslant k\big) \;\geqslant\; \mathbb{P}_\eta\big(M_n \leqslant k\big)\,. \tag{8}$$

Then, in the second stage, we try to stabilize these M'_n particles inside $V_{n+4\ell} \setminus V_n$, adapting again the trapping procedure presented in [RS12]. This procedure shows that, if $M'_n \leqslant k$, then with probability at least $\mathbb{P}\big(G_1 + \cdots + G_k \leqslant \ell\big)$, this second stage succeeds, yielding an acceptable toppling sequence which stabilizes these M'_n particles inside $V_{n+4\ell} \setminus V_n$ with none of these particles jumping out of $V_{n+4\ell} \setminus V_n$.

Thus, we deduce that

$$\mathbb{P}_\eta\big(M_{n+4\ell} = 0\big) \;\geqslant\; \mathbb{P}_\eta\big(M'_n \leqslant k\big) \times \mathbb{P}\big(G_1 + \cdots + G_k \leqslant \ell\big)$$

which, combined with (8), concludes the proof of the Lemma. □

4.3 Concluding proof of Theorem 4

We now put the pieces together to obtain the claimed bound.

Proof of Theorem 4. Let $\lambda > 0$, let p be a nearest-neighbour jump distribution on $\mathbb{Z}$, let $\zeta > \zeta_c$ and consider ε, α and β such that

$$0 \;\leqslant\; \frac{(1+\lambda)\varepsilon}{\lambda} \;<\; \alpha \;<\; \beta \;<\; \frac{\zeta - \zeta_c}{4\zeta_c}\,. \tag{9}$$

Let $(G_j)_{j\geqslant 1}$ be i.i.d. Geometric variables with parameter $\lambda/(1+\lambda)$. Since $(1+\lambda)\varepsilon/\lambda < \alpha$, the weak law of large numbers ensures that

$$\lim_{n\to\infty} \mathbb{P}\big(G_1 + \cdots + G_{\lfloor \varepsilon n\rfloor} \leqslant \alpha n\big) \;=\; 1\,.$$

Thus, we can take $n_0 \geqslant 1$ such that, for every $n \geqslant n_0$,

$$\mathbb{P}\big(G_1 + \cdots + G_{\lfloor \varepsilon n\rfloor} \leqslant \alpha n\big) \;\geqslant\; \frac{1+4\alpha}{1+4\beta}\,.$$

Now, let $n \geqslant n_0$ and let $\eta : V_n \to \mathbb{N}$ be a fixed deterministic initial configuration such that $\|\eta\| \geqslant \zeta n$. Applying Lemma 4 with $k = \lfloor \varepsilon n\rfloor$ and $\ell = \lfloor \alpha n\rfloor$, we get

$$\mathbb{P}_\eta\big(M_{n+4\ell} = 0\big) \;\geqslant\; \mathbb{P}_\eta\big(M_n \leqslant \varepsilon n\big) \times \mathbb{P}\big(G_1 + \cdots + G_{\lfloor \varepsilon n\rfloor} \leqslant \alpha n\big) \;\geqslant\; \mathbb{P}_\eta\big(M_n \leqslant \varepsilon n\big) \times \frac{1+4\alpha}{1+4\beta}\,. \tag{10}$$

Then, note that

$$\frac{\|\eta\|}{n+4\ell} \;\geqslant\; \frac{\zeta}{1+4\alpha} \;>\; \zeta_c\,,$$

by virtue of (9). Thus, applying Theorem 5 to η, seen as a configuration on $V_{n+4\ell}$, we have

$$\mathbb{P}_\eta\big(M_{n+4\ell} = 0\big) \;\leqslant\; \frac{(1+4\alpha)\zeta_c}{\zeta}\,. \tag{11}$$

Combining (10) and (11), we get

$$\mathbb{P}_\eta\big(M_n \leqslant \varepsilon n\big) \;\leqslant\; \frac{1+4\beta}{1+4\alpha} \times \frac{(1+4\alpha)\zeta_c}{\zeta} \;=\; \frac{(1+4\beta)\zeta_c}{\zeta}\,.$$

Taking the supremum over all configurations $\eta : V_n \to \mathbb{N}$ with $\|\eta\| \geqslant \zeta n$, we obtain that

$$\forall \beta \in \left(\frac{(1+\lambda)\varepsilon}{\lambda}, \frac{\zeta-\zeta_c}{4\zeta_c}\right) \quad \exists\, n_0 \geqslant 1 \quad \forall n \geqslant n_0 \qquad \sup_{\substack{\eta: V_n \to \mathbb{N}:\\ \|\eta\| \geqslant \zeta n}} \mathbb{P}_\eta\big(M_n \leqslant \varepsilon n\big) \;\leqslant\; \frac{(1+4\beta)\zeta_c}{\zeta}\,,$$

which is precisely the claim of Theorem 4. □

5 Proof of Theorem 3

We now prove the equivalence claimed in Theorem 3. Let $\lambda > 0$, let p be a nearest-neighbour jump distribution on $\mathbb{Z}$ and let η_0 be an i.i.d. initial distribution with mean ζ and all particles initially active.

5.1 Direct implication

The direct implication is an easy consequence of Theorem 4. It follows from the Central Limit Theorem that $\mathbb{P}\big(\|\eta_0\|_{V_n} \geqslant \zeta n\big) \to 1/2$ when $n \to \infty$. Thus, choosing whatever $\varepsilon > 0$ in the range indicated by Theorem 4, we have

$$\liminf_{n\to\infty} \frac{\mathbb{E}\big[M_n\big]}{n} \;\geqslant\; \liminf_{n\to\infty} \mathbb{P}\big(\|\eta_0\|_{V_n} \geqslant \zeta\big) \inf_{\substack{\eta: V_n \to \mathbb{N}:\\ \|\eta\| \geqslant \zeta n}} \mathbb{P}_\eta\big(M_n > \varepsilon n\big)\,\varepsilon \;\geqslant\; \frac{1}{2} \times \frac{\zeta_c}{\zeta}\left(1 + \frac{2(1+\lambda)\varepsilon}{\lambda}\right) \times \varepsilon \;>\; 0\,.$$

5.2 Reciprocal: proof of Proposition 1

The reciprocal implication of Theorem 3 follows from Theorem 2 and Proposition 1, which deals with the particular case of $\zeta = \zeta_c$, and which we now prove.

Proof of Proposition 1. Let d, λ, p and η_0 be as in the statement. By monotonicity (see for example Lemma 2.5 of [Rol20]), we can assume without loss of generality that all the particles are active in the configuration η_0. Let $\varepsilon \in (0, \zeta_c)$. Then, let us consider another i.i.d. initial distribution η_0' with mean $\zeta_c - \varepsilon$, which is coupled with η_0 in such a way that $\eta_0'(x) \leqslant \eta_0(x)$ for every $x \in \mathbb{Z}^d$. This can be done for example by taking $\eta_0'(x) = Y_x \eta_0(x)$, where $(Y_x)_{x\in\mathbb{Z}^d}$ are i.i.d. Bernoulli variables with parameter $(\zeta_c - \varepsilon)/\zeta_c$, independent of everything else.

Then, for every $n \geqslant 1$, denoting by M_n and M_n' the numbers of particles which jump out of the box V_n starting respectively with η_0 and with η_0', we claim that

$$\mathbb{E}\big[M_n\big] \;\leqslant\; \varepsilon\,|V_n| + \mathbb{E}\big[M_n'\big]\,. \tag{12}$$

Indeed, starting from the configuration η_0, we may first apply acceptable topplings to the configuration $\eta_0 - \eta_0'$, until all particles exit, leaving us with only the configuration η_0' remaining inside V_n. During this first stage, the average number of particles which jump out of the box is equal to $\mathbb{E}\big[\|\eta_0\| - \|\eta_0'\|\big] = \varepsilon\,|V_n|$. Then, we stabilize in V_n with legal topplings, which gives a number of particles jumping out of the box which is distributed as M_n'. Since we performed acceptable topplings during the first stage, we obtain an upper bound on M_n, whence (12).

Then, by the contrapositive of Proposition 2, we know that

$$\lim_{n\to\infty} \frac{\mathbb{E}\big[M_n'\big]}{|V_n|} \;=\; 0\,.$$

Combining this with (12), we deduce that

$$\limsup_{n\to\infty} \frac{\mathbb{E}\big[M_n\big]}{|V_n|} \;\leqslant\; \varepsilon \,+\, \lim_{n\to\infty} \frac{\mathbb{E}\big[M_n'\big]}{|V_n|} \;=\; \varepsilon\,.$$

This being true for every $\varepsilon \in (0, \zeta_c)$, the proof of Proposition 1 is complete. □

References

[AFG22] Amine Asselah, Nicolas Forien, and Alexandre Gaudillière. The Critical Density for Activated Random Walks is always less than 1. *arXiv preprint arXiv:2210.04779*, 2022.

[AGG10] Gideon Amir and Ori Gurel-Gurevich. On fixation of activated random walks. *Electron. Commun. Probab.*, 15:119–123, 2010.

[AMP02a] O. S. M. Alves, F. P. Machado, and S. Yu. Popov. Phase transition for the frog model. *Electron. J. Probab.*, 7:no. 16, 21, 2002.

[AMP02b] O. S. M. Alves, F. P. Machado, and S. Yu. Popov. The shape theorem for the frog model. *Ann. Appl. Probab.*, 12(2):533–546, 2002.

[ARS22] Amine Asselah, Leonardo T. Rolla, and Bruno Schapira. Diffusive bounds for the critical density of activated random walks. *ALEA Lat. Am. J. Probab. Math. Stat.*, 19(1):457–465, 2022.

[BGH18] Riddhipratim Basu, Shirshendu Ganguly, and Christopher Hoffman. Non-fixation for conservative stochastic dynamics on the line. *Comm. Math. Phys.*, 358(3):1151–1185, 2018.

[BGHR19] Riddhipratim Basu, Shirshendu Ganguly, Christopher Hoffman, and Jacob Richey. Activated random walk on a cycle. *Ann. Inst. Henri Poincaré Probab. Stat.*, 55(3):1258–1277, 2019.

[BS22] Alexandre Bristiel and Justin Salez. Separation cutoff for activated random walks. *arXiv*, 2022.

[BTW87] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Physical review letters*, 59(4):381, 1987.

[CR21] Manuel Cabezas and Leonardo T. Rolla. Avalanches in critical activated random walks. In *In and out of equilibrium 3. Celebrating Vladas Sidoravicius*, volume 77 of *Progr. Probab.*, pages 187–205. Birkhäuser/Springer, Cham, [2021] ©2021.

[CRS14] M. Cabezas, L. T. Rolla, and V. Sidoravicius. Non-equilibrium phase transitions: activated random walks at criticality. *J. Stat. Phys.*, 155(6):1112–1125, 2014.

[DF91] P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Sem. Mat. Univ. Politec. Torino*, 49(1):95–119 (1993), 1991. Commutative algebra and algebraic geometry, II (Italian) (Turin, 1990).

[DRS10] Ronald Dickman, Leonardo T. Rolla, and Vladas Sidoravicius. Activated random walkers: facts, conjectures and challenges. *J. Stat. Phys.*, 138(1-3):126–142, 2010.

[FG22] Nicolas Forien and Alexandre Gaudillière. Active Phase for Activated Random Walks on the Lattice in all Dimensions. *arXiv preprint arXiv:2203.02476*, 2022.

[HRR23] Christopher Hoffman, Jacob Richey, and Leonardo T. Rolla. Active phase for activated random walk on $\mathbb{Z}$. *Comm. Math. Phys.*, 399(2):717–735, 2023.

[Hu22] Yiping Hu. Active Phase for Activated Random Walk on $\mathbb{Z}^2$. *arXiv preprint arXiv:2203.14406*, 2022.

[JMT23] Antal A Járai, Christian Mönch, and Lorenzo Taggi. Law of large numbers for an elementary model of self-organised criticality. *arXiv preprint arXiv:2304.10169*, 2023.

[LL21] Lionel Levine and Feng Liang. Exact sampling and fast mixing of Activated Random Walk. *arXiv preprint arXiv:2110.14008*, 2021.

[LS21] Lionel Levine and Vittoria Silvestri. How far do activated random walkers spread from a single source? *J. Stat. Phys.*, 185(3):Paper No. 18, 27, 2021.

[LS24] Lionel Levine and Vittoria Silvestri. Universality conjectures for activated random walk. *Probab. Surv.*, 21:1–27, 2024.

[Rol08] Leonardo T Rolla. Generalized Hammersley process and phase transition for activated random walk models. *PhD thesis, arXiv:0812.2473*, 2008.

[Rol20] Leonardo T. Rolla. Activated random walks on $\mathbb{Z}^d$. *Probab. Surv.*, 17:478–544, 2020.

[RS12] Leonardo T. Rolla and Vladas Sidoravicius. Absorbing-state phase transition for driven-dissipative stochastic dynamics on $\mathbb{Z}$. *Invent. Math.*, 188(1):127–150, 2012.

[RSZ19] Leonardo T. Rolla, Vladas Sidoravicius, and Olivier Zindy. Universality and sharpness in activated random walks. *Ann. Henri Poincaré*, 20(6):1823–1835, 2019.

[RT18] L. T. Rolla and L. Tournier. Non-fixation for biased activated random walks. *Ann. Inst. Henri Poincaré Probab. Stat.*, 54(2):938–951, 2018.

[ST18] Alexandre Stauffer and Lorenzo Taggi. Critical density of activated random walks on transitive graphs. *Ann. Probab.*, 46(4):2190–2220, 2018.

[Tag19] Lorenzo Taggi. 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(3):1751–1764, 2019.

[Tag23] Lorenzo Taggi. Essential enhancements in Abelian networks: Continuity and uniform strict monotonicity. *Ann. Probab.*, 51(6):2243–2264, 2023.