



HAL
open science

The critical density of the Stochastic Sandpile Model

Concetta Campailla, Nicolas Forien, Lorenzo Taggi

► **To cite this version:**

Concetta Campailla, Nicolas Forien, Lorenzo Taggi. The critical density of the Stochastic Sandpile Model. 2024. hal-04780325

HAL Id: hal-04780325

<https://hal.science/hal-04780325v1>

Preprint submitted on 13 Nov 2024

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

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

The critical density of the Stochastic Sandpile Model

Concetta Campailla*

Nicolas Forien†

Lorenzo Taggi*

October 25, 2024

Abstract

We study the stochastic sandpile model on \mathbb{Z}^d and demonstrate that the critical density is strictly less than one in all dimensions. This generalizes a previous result by Hoffman, Hu, Richey, and Rizzolo (2022), which was limited to the one-dimensional case. In addition, we show that the critical density is strictly positive on any vertex-transitive graph, extending the earlier result of Sidoravicius and Teixeira (2018) and providing a simpler proof.

1 Introduction

Many natural and complex phenomena, such as earthquakes, snow avalanches, forest fires, and others, exhibit sudden large-scale events triggered by small disturbances. Bak, Tang, and Wiesenfeld proposed a general explanation for these phenomena, which they referred to as *self-organized criticality* [4]. Their research brought about a major shift in understanding and became one of the most frequently referenced papers in physics throughout the following decade.

Various processes have been suggested as universal models of self-organized criticality [3]. Bak, Tang, and Wiesenfeld introduced the abelian sandpile model, which operates through a deterministic process of redistribution [4]. The model's intricate mathematical framework has been studied extensively by various branches of mathematics. However, physicists later observed that the model lacks the universality typically expected from systems displaying self-organized criticality. For example, it is highly dependent on the initial conditions and other details [10, 8, 17].

Manna introduced a probabilistic variant of the abelian sandpile model [21], later known as the *stochastic sandpile model* (SSM). This version demonstrates more robust universality properties compared to the abelian sandpile. It is widely believed that this model is the appropriate framework for investigating the phenomenon of self-organized criticality and its universal characteristics [9, 20, 16]. However, its rigorous mathematical analysis presents various challenges. For this reason, despite some progress, in recent years, the focus of the mathematics community has shifted to a close variant of SSM, the activated random walk model (ARW), which is more tractable mathematically and is believed to belong to the same universality class as SSM [20]. In particular, progress has been made toward understanding the properties of the critical curve for ARW. These advancements include the demonstration that it is strictly between zero and one [2, 11, 15, 5, 14, 25, 28], continuity as a function of the deactivation parameter [29] and universality [26]. Furthermore, the equivalence of critical thresholds has been established in one dimension [13] and explored in wider generality [19]. Additionally, some bounds on the mixing time of the so-called 'driven-dissipative' dynamics have been provided [18, 6].

In contrast, our understanding of the stochastic sandpile model remains more limited. In particular, one of the key open questions for this class of models — whether the critical density is below one — remains unresolved for any graph beyond \mathbb{Z} [12]. In this paper, we resolve this conjecture for \mathbb{Z}^d in any dimension $d \geq 1$. Additionally, we provide a new, simplified proof that the critical density is strictly positive on any vertex-transitive graph, thereby extending previous results [22, 25, 27].

*Sapienza Università di Roma, Dipartimento di Matematica Guido Castelnuovo, Roma, Italy.

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

The central tool in deriving the upper bound for the critical threshold is a novel stochastic comparison with a variant of the activated random walk model. This comparison enables us to apply recently developed techniques from the analysis of ARW to the stochastic sandpile model. We hope that this new stochastic comparison may help to make further progress in the rigorous understanding of the stochastic sandpile model in the future.

1.1 Definition and main results

Let us consider a finite graph $G = (V, E)$. At time zero, each vertex has a number of particles, which is an independent Poisson random variable with mean $\mu > 0$. At any time, each vertex is declared *stable* if it hosts at most one particle and *unstable* otherwise. The dynamics evolve in continuous time. Each vertex is associated with an independent Poisson clock with rate one. When the clock rings, if the vertex is unstable, two particles from that vertex perform an independent simple random walk step, while if the vertex is stable, nothing happens.

On infinite vertex-transitive graphs (like, for example, \mathbb{Z}^d), the dynamics can be defined by taking the limit of large finite balls with absorbing boundary conditions [1, 25]. Criticality of the SSM on infinite graphs is defined with respect to whether the system remains active or not. We say that the system *locally fixates* if every vertex is visited only finitely many times by the particles and eventually becomes stable; otherwise, we say that the system stays *active*.

Rolla and Sidoravicius [25] showed in wide generality the existence of a critical value $\mu_c \geq 0$ such that if $\mu > \mu_c$, the system is almost surely active, and if $\mu < \mu_c$, the system fixates almost surely. They proved that $\mu_c \in [\frac{1}{4}, 1]$ in \mathbb{Z} . The lower bound in \mathbb{Z} was then improved to $\mu_c \geq \frac{1}{2}$ in [22]. Sidoravicius and Teixeira [27] employed renormalisation techniques to show that $\mu_c > 0$ in \mathbb{Z}^d for each $d \in \mathbb{N}$.

Our first theorem shows that the critical density of the stochastic sandpile is strictly positive on vertex-transitive graphs.

Theorem 1.1. *For any vertex-transitive graph G of degree at least two, the critical density of the stochastic sandpile model satisfies*

$$\mu_c > 0.$$

The novelty of the result is not only in extending the class of graphs for which it holds but also in presenting a simpler proof.

Our second theorem shows that the critical density of SSM is strictly below one in \mathbb{Z}^d for each dimension $d \geq 1$, thus extending previous results from [12], which are restricted to one dimension.

Theorem 1.2. *In any dimension $d \geq 1$, the critical density of the stochastic sandpile model on \mathbb{Z}^d satisfies*

$$\mu_c < 1.$$

Theorem 1.2 is one of the main results of our paper. Together with Proposition 4.1 below, which presents the new coupling argument, these form the core contributions of our work.

1.2 Proof outline

The proof of Theorem 1.1 requires the adaptation to the stochastic sandpile model of the “weak stabilisation” method, which was introduced in [28] for the activated random walk model. The proof is short and is presented in Section 3, together with a short description of the method.

For the proof of Theorem 1.2, the new central idea is a stochastic comparison between the SSM and a version of the activated random walk model, which we refer to as the activated random walk with instantaneous deactivation (in short: ARWD), see Proposition 4.1 below.

Similarly to the activated random walk model, in ARWD, particles can be either active or sleeping. Active particles get activated whenever they share a vertex with other active particles, and they fall asleep with probability $\frac{\lambda}{1+\lambda}$ independently at each discrete step if they are alone, where $\lambda > 0$. Informally, the main difference with the activated random walk model is that, in addition, the active particles fall asleep *instantaneously* whenever they jump to a vertex.

The coupling between the two models requires moving particles ‘in parallel’ by using the same jump instructions. Even though in SSM particles can only be of one type, we can view each particle as ‘sleeping’ if it has just jumped to a vertex that has been toppled an even number of times. We think of a particle as ‘activated’ whenever such a vertex is visited by some other particle afterward. The coupling is then constructed in such a way that, whenever a particle becomes stable in SSM, the corresponding particle also falls asleep in ARWD. In some cases, it may happen that we are allowed to move a particle in SSM, but we are not allowed to move the corresponding particle in ARWD since it is sleeping. The opposite, however, never happens. This allows us to *dominate from below* the activity in SSM by the activity in ARWD.

For the coupling to work, we need to restrict the set of vertices where the reactivation of sleeping particles is possible in both models, ignoring the fact that particles outside this set may still be active. Such a restriction is necessary to dominate from below using independent random variables the reactivation of sleeping particles in SSM within this set. The proof method also requires disallowing the deactivation of particles outside a certain arbitrary set of vertices A in both models (for SSM, this is formalized through the introduction of the so-called ‘ A -stabilisation’, see Section 2.3 below) and then considering all possible choices of A . We refer to Section 5.1 for a more formal description of the coupling argument.

Once the stochastic comparison is established, it remains to show that ARWD is active, thus deducing by the stochastic comparison that SSM is also active. For this, we adapt to ARWD the strategy which was developed in [2]. We first show that for μ close enough to 1 and less than one, the stabilisation time of the SSM on the torus is exponentially large:

Theorem 1.3. *For every $d \geq 1$, for $\mu \in (0, 1)$ close enough to 1, the stabilisation time \mathcal{T}_n of the stochastic sandpile model on the torus $\mathbb{Z}_n^d = (\mathbb{Z}/n\mathbb{Z})^d$ started with i.i.d. Poisson numbers of particles on each site, with parameter μ , satisfies*

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

Theorem 1.3 implies Theorem 1.2 through the following result, which is an analog of Theorem 4 in [11]:

Theorem 1.4. *For every $d \geq 1$, for every $\mu < \mu_c$, where μ_c is the critical density of the stochastic sandpile model on \mathbb{Z}^d , the property (1) does not hold.*

This theorem can be proved by adapting the proof of Theorem 4 in [11]. Since it is already known that $\mu_c < 1$ for the stochastic sandpile model in dimension $d = 1$ [12], we provide the detailed proof only for $d \geq 2$.

While the overall strategy to prove that ARWD is active is similar to that in [2], the adaptation requires addressing some technical challenges. These include the fact that, unlike in [2], not all particles are active at time zero, the reactivation of sleeping particles is restricted to a certain set of sites, and there is an additional rule in ARWD governing when particles fall asleep. These features make it harder for the model to be active in comparison to the activated random walk. We believe these generalizations may have some relevance on their own. We also remark that ARWD is an example of a non-abelian particle system, and exploring this model may offer further insights.

Notation. We let $d_G(\cdot, \cdot)$ be the graph distance in $G = (V, E)$, $B(x, r)$ be the set of vertices which have graph distance at most r from x . We use the notation $B(r)$ for $B(o, r)$, where $o \in V$ is the root of G . We also use the notation $\mathbb{N}_0 = \{0, 1, \dots\}$ and $\mathbb{N} = \{1, 2, \dots\}$.

2 Diaconis-Fulton representation

In this section we describe the Diaconis-Fulton graphical representation for the dynamics of SSM and introduce several stabilisation procedures.

2.1 Stabilisation

Let $G = (V, E)$ be an undirected, locally finite, finite or infinite graph. A configuration of the stochastic sandpile model on G can be represented by a vector $\eta : V \rightarrow \mathbb{N}_0$, with $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$, where $\eta(x) = k \in \mathbb{N}_0$

means that there are k particles on the site x . A site $x \in V$ is *stable* in the configuration η if $\eta(x) = 0$ or 1 , and it is *unstable* if $\eta(x) \geq 2$. We say that a configuration η is stable in a subset $V' \subset V$ if $\eta(x) \leq 1$ for all $x \in V'$. We now define, for every $x \in V$ and every y neighbour of x in the graph G , an operator τ_{xy} which corresponds to one particle jumping from x to y : namely, if $\eta \in \mathbb{N}_0^V$ is such that $\eta(x) \geq 1$, we define

$$\tau_{xy}\eta(z) = \begin{cases} \eta(x) - 1 & \text{if } z = x \\ \eta(y) + 1 & \text{if } z = y \\ \eta(z) & \text{otherwise.} \end{cases}$$

Let us fix for now an array of instructions $\tau = (\tau^{x,j} : x \in V, j \in \mathbb{N})$, where for every $x \in V$ and $j \in \mathbb{N}$, $\tau^{x,j} \in \{\tau_{xy} : y \sim x\}$. Once this array is fixed, the configuration of the model at a certain time may be encoded in a couple (η, h) , where $\eta : V \rightarrow \mathbb{N}_0$ is the current configuration and $h : V \rightarrow \mathbb{N}_0$ is a function which counts the number of pairs of instructions already used at each site. We say that we topple x when we act on this couple (η, h) through the operator Φ_x , which is defined as

$$\Phi_x(\eta, h) = (\tau^{x,2h(x)+2}\tau^{x,2h(x)+1}\eta, h + \mathbb{1}_x),$$

provided that x is unstable in η . The toppling operation Φ_x at x is termed *legal* for a couple (η, h) when x is unstable in η , and *illegal* otherwise. We let $h = 0$ indicate $h(y) = 0$ for all $y \in V$, and we abbreviate $\Phi_x(\eta, 0)$ as $\Phi_x\eta$.

Given a sequence $\alpha = (x_1, x_2, \dots, x_k)$ of sites, we denote by $\Phi_\alpha = \Phi_{x_k}\Phi_{x_{k-1}}\dots\Phi_{x_2}\Phi_{x_1}$ the composition of the topplings at the sites $x_1, x_2, \dots, x_{k-1}, x_k$, in that order. We call Φ_α legal for a configuration (η, h) if Φ_{x_1} is legal for (η, h) and Φ_{x_ℓ} is legal for $\Phi_{(x_1, \dots, x_{\ell-1})}(\eta, h)$ for each $\ell = 2, \dots, k$. In this case we call α a legal sequence of topplings for (η, h) .

For every $x \in V$, we define

$$m_\alpha(x) = \sum_{\ell=1}^k \mathbb{1}_{\{x_\ell=x\}}.$$

In other words, $m_\alpha(x)$ indicates the number of times that the site x appears in the sequence α . Given a toppling sequence $\alpha = (x_1, x_2, \dots, x_k)$ and a subset V' of V , we say that α is contained in V' , and we write $\alpha \subset V'$, if all its elements are in V' , and we say that α stabilises a given configuration η in V' if every site $x \in V'$ is stable in the particle configuration of $\Phi_\alpha\eta$.

For any subset $V' \subset V$, any particle configuration η and any array of instructions τ , we define the *odometer*

$$m_{V',\eta} = 2 \sup_{\alpha \subset V', \alpha \text{ legal for } (\eta, 0)} m_\alpha.$$

In the particular case $V' = V$, we simply write $m_\eta = m_{V,\eta}$. It follows from the Abelian property (see [25, Lemma 2]) that, if there exists a legal sequence $\alpha \subset V'$ which stabilises η in V' , then $m_{V',\eta} = 2m_\alpha$, that is to say, $m_{V',\eta}$ counts the number of times that a particle jumps from x in the stabilisation of V' starting from configuration η using the instructions in τ , using whatever legal stabilising sequence. For any $V'' \subset V' \subset V$, we write

$$\|m_{V',\eta}\|_{V''} = \sum_{x \in V''} m_{V',\eta}(x)$$

for the number of topplings on the sites of V'' during the stabilisation of η in V' . In the particular case where $V'' = V'$, we simply write $\|m_{V',\eta}\| = \|m_{V',\eta}\|_{V'}$ and, if $V'' = V' = V$, we write $\|m_\eta\|$. If $\|m_{V',\eta}\| < \infty$, then the Abelian property ensures that the configuration reached when stabilising η in V' does not depend on the toppling sequence used and thus we can denote by $\eta_{V',\infty}$ the final configuration. In the case where $V' = V$, we simply write η_∞ .

We now introduce a probability measure on the space of instructions and of particle configurations. Let \mathcal{P} 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 any $j \in \mathbb{N}$, $\mathcal{P}(\tau^{x,j} = \tau_{xy}) = \frac{1}{d_x}$ for any $y \in V$ neighboring x , where d_x is the degree of the vertex x in G .

The initial distribution η is i.i.d. with a Poisson number of particles with parameter μ on each site. Let us denote by P_μ the corresponding probability measure on the set \mathbb{N}_0^V . We write $\mathcal{P}^\mu = \mathcal{P} \otimes P_\mu$ for the

joint law of τ and η , and we let \mathbb{P}^μ denote the probability measure induced by the SSM process. That is to say, compared to \mathcal{P}^μ , the measure \mathbb{P}^μ also contains the randomness relative to the times at which each site topples. Write E , E^μ , \mathbb{E}^μ for the expectations with respect to the measures \mathcal{P} , \mathcal{P}^μ , \mathbb{P}^μ .

The following lemma gives a 0-1 law for local fixation and relates it to a property of the stabilisation odometer m_η :

Lemma 2.1 (Lemma 4 in [25]). *Let $G = (V, E)$ be an infinite vertex-transitive graph and let $x \in V$ be arbitrary. Then,*

$$\mathbb{P}^\mu(\text{the system locally fixates}) = \mathcal{P}^\mu(m_\eta(x) < \infty) \in \{0, 1\}.$$

This lemma tells us that, in order to establish local fixation almost surely, it is enough to show that a site is toppled a finite number of times with positive probability.

2.2 Stabilisation with half-topplings

Given an array of instructions $\tau = (\tau^{x,j} : x \in V, j \in \mathbb{N})$ and a configuration η , we introduce a stabilisation strategy with the half-toppling operation, which moves the particles one by one.

A *half-toppling* consists in sending only one particle to a neighboring site chosen at random (i.e., following the appropriate instruction) and is possible when $\eta(x) \geq 1$. Fix an array of instructions τ . We say we half-topple x when we act on a particle configuration η satisfying $\eta(x) \geq 1$ and on the vector h through the operator ϕ_x , which is defined as

$$\phi_x(\eta, h) = (\tau^{x,2h(x)+1}\eta, h + \frac{1}{2}\mathbb{1}_x).$$

As before, we abbreviate $\phi_x(\eta, 0)$ by $\phi_x\eta$. For a particle configuration $\eta : V \rightarrow \mathbb{N}_0$ and an odometer $h : V \rightarrow (1/2)\mathbb{N}_0$, we say that a site x is half-unstable in (η, h) if either $\eta(x) \geq 2$ or $\eta(x) = 1$ and $h(x) \notin \mathbb{N}_0$ (and otherwise the site is said half-stable). For a sequence $\alpha = (x_1, \dots, x_k)$ we denote, for any $\ell \in \{1, \dots, k\}$, $\alpha^{(\ell)} = (x_1, \dots, x_\ell)$, and we define the operator ϕ_α as the composition $\phi_{x_k}\phi_{x_{k-1}}\cdots\phi_{x_1}$. We say that the sequence $\alpha = (x_1, \dots, x_k)$ is *admissible* for (η, h) if for any $\ell \in \{1, \dots, k\}$, the particle configuration of $\phi_{\alpha^{(\ell)}}(\eta, h)$ contains at least one particle at x_ℓ . Given α admissible for η , we say that $\alpha = (x_1, \dots, x_k)$ is *legal* for the stabilisation with half-topplings of (η, h) (in short: half-legal) if for every $\ell \in \{1, \dots, k\}$, the site x_ℓ is half-unstable in $\phi_{\alpha^{(\ell)}}(\eta, h)$. Moreover, we say that the sequence of topplings α *stabilises* (η, h) in V' via half-topplings (in short: it is half-stabilising) if it is half-legal for η , it is contained in V' and every site $x \in V'$ is half-stable in $\phi_\alpha(\eta, h)$. We have

$$m_{V', \eta} = \sup_{\alpha \subset V', \alpha \text{ half-legal for } \eta} m_\alpha,$$

for any $V' \subset V$ and the Abelian property also holds for the stabilisation with half-topplings (see [25, Section 6] for details). Thus, $m_{V', \eta}$ counts the number of times that a particle jumps from x both in the stabilisation and in the half-stabilisation of η in V' using the instructions in τ , using whatever stabilising or half-stabilising sequence α .

2.3 Modified stability rules with respect to a fixed settling set

Definition 2.1. Given a subset $A \subset V$, a site $x \in A$ is said *A-stable* for a couple (η, h) if $\eta(x) = 0$ or if $\eta(x) = 1$ and $h(x) \in \mathbb{N}_0$, whereas a site $x \in V \setminus A$ is said *A-stable* for (η, h) if $\eta(x) = 0$. Otherwise, the site is said *A-unstable*. We call an admissible sequence of half-topplings $\alpha = (x_1, \dots, x_k)$ *A-legal* for η if for any $\ell \in \{1, \dots, k\}$, the site x_ℓ is *A-unstable* in $\phi_{\alpha^{(\ell)}}\eta$. We say that a couple (η, h) is *A-stable* in a set $V' \subset V$ if all the sites of V' are *A-stable* for (η, h) .

Definition 2.2. We say that a sequence of half-topplings α *A-stabilises* η in $V' \subset V$ if α is admissible and $\phi_\alpha\eta$ is *A-stable* in V' . We also define the odometer of the *A-stabilisation* of V' as

$$m_{V', \eta}^A = \sup_{\alpha \subset V', \alpha \text{ is } A\text{-legal}} m_\alpha.$$

In the particular case $V' = V$, we omit V' and simply write m_η^A .

Similarly to stabilisation and half-stabilisation, the Abelian property also holds for the A -stabilisation with minor changes in the proofs (see [25, Section 3]). For any $V'' \subset V'$, denote by

$$\|m_{V'',\eta}^A\|_{V''} = \sum_{x \in V''} m_{V'',\eta}^A(x)$$

the number of half-topplings on the sites of V'' during the A -stabilisation of η in V' . In the particular case where $V'' = V'$, we simply write $\|m_{V',\eta}^A\| = \|m_{V',\eta}^A\|_{V'}$ and, if $V'' = V' = V$, we write $\|m_\eta^A\|$. If $\|m_{V',\eta}^A\| < \infty$, which means that η can be A -stabilised in V' with a finite number of half-topplings, then the Abelian property ensures that the configuration reached when stabilising η in V' does not depend on the toppling sequence used, which allows us to write this final configuration $\eta_{V',\infty}^A$. In the case where $V' = V$, we simply write η_∞^A . Note that if $V' = V = A$ then $m_\eta^A = m_\eta$ and $\eta_\infty^A = \eta_\infty$.

Lemma 2.2. *If $\eta_\infty = \mathbb{1}_A$ then $m_\eta = m_\eta^A$.*

Proof. If $\eta_\infty = \mathbb{1}_A$ then there exists a half-legal sequence α such that $\phi_\alpha \eta = (\mathbb{1}_A, h)$ with $h(x)$ integer for every $x \in A$. Then we have $m_\eta = m_\alpha$.

Yet, this final configuration $(\mathbb{1}_A, h)$ is not only stable but also A -stable. Besides, as any half-legal sequence, this sequence α is also A -legal. Thus, α is a A -legal sequence which A -stabilises η , whence by the Abelian property $m_\eta^A = m_\alpha = m_\eta$. \square

3 Proof of Theorem 1.1

The proof of Theorem 1.1 requires the introduction of the weak stabilisation technique for the stochastic sandpile model. This method was introduced for the activated random walk model in [28].

3.1 Weak stabilisation

We introduce the weak stabilisation strategy for the stochastic sandpile model. Given an array of instructions $\tau = (\tau^{x,j} : x \in V, j \in \mathbb{N})$, a configuration η and α a legal sequence for $(\eta, 0)$, we use the notation η_α for the particle configuration which is obtained by toppling the vertices of the sequence α , i.e., the first component of the pair $\Phi_\alpha(\eta)$, using the convention that $\eta_\emptyset = \eta$. For any $x \in V$, let $I_x = B(x, 1)$ be the set of vertices whose graph distance from x is at most one.

Definition 3.1 (Weakly stable configurations). Given a set $V' \subset V$, we say that the configuration η is *weakly stable* in V' with respect to $x \in V'$ if one of the following two conditions is fulfilled:

- (i) η is stable in V' ,
- (ii) η is stable in $V' \setminus I_x$, there exists a site $z \in I_x$ such that $\eta(z) = 2$ and the other sites of I_x are empty in η .

For conciseness, we simply say that η is weakly stable in (x, V') .

Definition 3.2 (Weak stabilisation). Given a set $V' \subset V$, a vertex $x \in V'$, and an array τ , we say that the sequence of topplings $\alpha = (x_1, \dots, x_k)$ is *legal for the weak stabilisation* of η in V' with respect to x (in short: x -weakly legal) if it is legal for $(\eta, 0)$ and if

$$x_\ell \in I_x \quad \text{and} \quad \eta_{\alpha^{(\ell-1)}}(x_\ell) = 2 \implies \exists z \in I_x : z \neq x_\ell \text{ and } \eta_{\alpha^{(\ell-1)}}(z) \geq 1,$$

for each $\ell \in \{1, 2, \dots, k\}$. We say that the sequence of topplings $\alpha = (x_1, \dots, x_k)$ *weakly stabilises* η in (x, V') if it is contained in V' , it is x -weakly legal for $(\eta, 0)$, and η_α is weakly stable in (x, V') .

Similarly to stabilisation, half-stabilisation and A -stabilisation, the Abelian property holds also for the weak stabilisation with minor changes in the proofs (see [21, Section 3]). An x -weakly legal sequence ensures that, whenever a vertex in I_x is toppled, there will still be a particle in I_x in the subsequent step. This observation leads to the following lemma.

Lemma 3.1. *For arbitrary $V' \subset V$ and $x \in V'$, let $\alpha = (x_1, \dots, x_k)$ be a sequence which weakly stabilises η in (x, V') . Suppose that there exists $\ell \in \{0, 1, \dots, k\}$ such that $\eta_{\alpha^{(\ell)}}(z) \geq 1$ for some $z \in I_x$. Then there exists $z \in I_x$ such that $\eta_\alpha(z) \geq 1$.*

Proof. The lemma follows from the definition of weakly legal sequence. Indeed, according to the definition of a (x, V') -weakly legal sequence, we can only topple a vertex in I_x if it contains at least three particles, or if it contains exactly two particles and another vertex in I_x hosts at least one particle. These conditions ensure that, if a particle jumps to I_x at any point, I_x will continue to host at least one particle at the end of the weak stabilisation, when we obtain a weakly stable configuration. \square

Weak stabilisation is used for the proof of the next proposition, providing a lower bound for the probability that stabilisation ends with at least one particle in I_x .

Proposition 3.2. *Let $G = (V, E)$ be a regular graph with degree d . Let $A_{x,L}$ be the event that there is at least one particle in I_x at the end of the stabilisation of $B(L)$. Then, for any $L \in \mathbb{N}$ and $x \in B(L)$,*

$$\mathcal{P}^\mu(A_{x,L}) \geq \frac{d-1}{d^3} - \mathcal{P}^\mu(m_{B(L),\eta}(x) = 0).$$

The strategy for the proof of Proposition 3.2 is to stabilise $B(L)$ by first weakly stabilising $(x, B(L))$, which gives a weakly stable configuration η^w for $(x, B(L))$. Then either I_x was never visited during the weak stabilisation, or I_x was visited and η^w is stable in $B(L)$, in which case Lemma 3.1 ensures that we finish the stabilisation procedure with one particle in I_x , or there exists $z \in I_x$ such that $\eta^w(z) = 2$, all the other vertices of I_x are empty and η^w is stable in $V' \setminus I_x$. In the latter case, with some positive probability, we stabilise $B(L)$ in at most two steps in such a way that some particle remains in I_x .

Proof of Proposition 3.2. Let η be the initial particle configuration. We stabilise η in $B(L)$ by first performing a weak stabilisation of $(x, B(L))$, thus obtaining a weakly stable configuration η^w , and then, if η^w is not stable in $B(L)$, by proceeding with the stabilisation.

Let $Q_{x,L}$ be the event that at least one particle has been in I_x during the weak stabilisation of $(x, B(L))$, i.e., there exists $\ell \in \{0, 1, \dots, k\}$ such that $\eta_{\alpha^{(\ell)}}(z) \geq 1$ for some $z \in I_x$, where $\alpha = (x_1, \dots, x_k)$ is a sequence which stabilises the initial configuration weakly in $(x, B(L))$. We observe that

$$\mathcal{P}^\mu(A_{x,L}) \geq \mathcal{P}^\mu(A_{x,L} \cap Q_{x,L} \cap \{\eta^w \text{ is stable in } B(L)\}) + \mathcal{P}^\mu(A_{x,L} \cap \{\eta^w \text{ is not stable in } B(L)\}). \quad (2)$$

Consider the second term on the right-hand side of (2). If η^w is not stable in $B(L)$, by definition of a weakly stable configuration this implies that there exists a site $z \in I_x$ such that $\eta^w(z) = 2$ and $\eta^w(y) = 0$ for all $y \in I_x \setminus \{z\}$. If $z \in I_x \setminus \{x\}$, it then happens with probability $\frac{d-1}{d^3}$ that such two particles both jump to x in one step and, at the next step, they jump to two distinct neighbours of x , thus giving a stable configuration in $B(L)$ with two particles in I_x . Instead, if $z = x$, it happens with probability $\frac{d-1}{d^3}$ that such two particles jump to two distinct neighbours of x , again giving a stable configuration in $B(L)$ with two particles in I_x . Hence,

$$\mathcal{P}^\mu(A_{x,L} \cap \{\eta^w \text{ is not stable in } B(L)\}) \geq \frac{d-1}{d^3} \mathcal{P}^\mu(\eta^w \text{ is not stable in } B(L)).$$

We consider now the first term on the right-hand side of (2). By Lemma 3.1 we have that

$$Q_{x,L} \cap \{\eta^w \text{ is stable in } B(L)\} \subset A_{x,L}.$$

Hence,

$$\begin{aligned} \mathcal{P}^\mu(A_{x,L} \cap Q_{x,L} \cap \{\eta^w \text{ is stable in } B(L)\}) &= \mathcal{P}^\mu(Q_{x,L} \cap \{\eta^w \text{ is stable in } B(L)\}) \\ &\geq 1 - \mathcal{P}^\mu(Q_{x,L}^c) - \mathcal{P}^\mu(\eta^w \text{ is not stable in } B(L)). \end{aligned}$$

By replacing the previous expression in (2), we then obtain that

$$\begin{aligned} \mathcal{P}^\mu(A_{x,L}) &\geq 1 - \mathcal{P}^\mu(Q_{x,L}^c) - \mathcal{P}^\mu(\eta^w \text{ is not stable in } B(L)) + \frac{d-1}{d^3} \mathcal{P}^\mu(\eta^w \text{ is not stable in } B(L)) \\ &\geq 1 - \mathcal{P}^\mu(Q_{x,L}^c) - \left(1 - \frac{d-1}{d^3}\right) \mathcal{P}^\mu(\eta^w \text{ is not stable in } B(L)) \\ &\geq 1 - \mathcal{P}^\mu(Q_{x,L}^c) - \left(1 - \frac{d-1}{d^3}\right) = \frac{d-1}{d^3} - \mathcal{P}^\mu(Q_{x,L}^c). \end{aligned}$$

The proof is then concluded by observing that the event $Q_{x,L}^c$ implies that $m_{B(L),\eta}(x) = 0$. \square

The next corollary is central in the proof of Theorem 1.1. The corollary implies that, if we assume that the stochastic sandpile model is active, then the probability that any ball I_x which is far enough from the boundary of $B(L)$ hosts at least one particle after stabilisation is bounded from below.

Corollary 3.3. *Let $G = (V, E)$ be an infinite vertex-transitive graph with degree d , and suppose that the SSM in G is active. Then, for any $\varepsilon > 0$, for r large enough, for any $L \in \mathbb{N}$ and $x \in B(L)$ with $B(x, r) \subset B(L)$,*

$$\mathcal{P}^\mu(A_{x,L}) \geq \frac{d-1}{d^3} - \varepsilon. \quad (3)$$

Proof. Since the system is active and the graph is vertex-transitive, Lemma 2.1 implies that we can find $r = r(\varepsilon) > 0$ such that for any $x \in V$,

$$\mathcal{P}^\mu(m_{B(x,r),\eta}(x) = 0) \leq \varepsilon.$$

By monotonicity [25, Lemma 3], we have that for every $x \in V$ such that $B(x, r) \subset B(L)$,

$$\mathcal{P}^\mu(m_{B(L),\eta}(x) = 0) \leq \mathcal{P}^\mu(m_{B(x,r),\eta}(x) = 0) \leq \varepsilon.$$

This concludes the proof, using Proposition 3.2. \square

3.2 Concluding the proof of Theorem 1.1

Let us briefly present the idea of the proof of Theorem 1.1 for SSM in \mathbb{Z}^d . Corollary 3.3 implies that, if the system is active, then the density of particles on large balls after stabilisation is at least a positive constant c , which depends only on d and does not depend on μ (this argument fails on non-amenable graphs, since (3) does not hold at vertices which are “too close” to the boundary of the ball, and such graphs have too many vertices on the boundary). Since the density of particles cannot increase during stabilization, this implies that $\mu \geq c$, and this being true for every μ such that SSM is active, we deduce that $\mu_c \geq c$.

The proof on general vertex-transitive graphs (which are not necessarily amenable) requires an additional tool, namely the method of ghost particles. We will still look for a lower bound on the density of particles after stabilization, using Corollary 3.3, to obtain *in fine* a lower bound on μ such that SSM stays active.

We let $X(t)$ denote a simple random walk in G , and P_x denote its law when $X(0) = x \in V$. We let E_x denote the corresponding expectation. Given a set $Z \subset V$, we define $\tau_Z := \inf\{t \geq 0 : X(t) \in Z\}$. For any $x, y \in V$, we define the Green’s function,

$$G_Z(x, y) := E_x \left[\sum_{t=0}^{\tau_Z c - 1} \mathbb{1}_{\{X(t)=y\}} \right].$$

We state an auxiliary result from [7] which will be used in the proof.

Lemma 3.4 (Lemma 4.1 in [7]). *Given an infinite vertex-transitive graph $G = (V, E)$, for any $r \in \mathbb{N}$ there exists $L_0 = L_0(r)$ such that, for any $L \geq L_0$,*

$$\sum_{x \in B(L)} G_{B(L)}(x, o) < 10 \sum_{\substack{x \in B(L): \\ B(x,r) \subset B(L)}} G_{B(L)}(x, o).$$

Proof of Theorem 1.1. Let $\mu > 0$ such that the SSM is active. We will show that this implies a lower bound on μ .

Let $L \in \mathbb{N}$. We perform a stabilisation of the set $B(L)$. Once the stabilisation of $B(L)$ is complete, we let each stable vertex in $B(L)$ with precisely one particle perform a *ghost* simple independent random walk starting from that vertex and being killed upon leaving $B(L)$. Let $\bar{m}_L(o)$ be the number of particles or ghosts jumping from the root vertex $o \in V$, let $w_L(o)$ be the number of ghosts jumping from o and recall that

$m_{B(L),\eta}(o)$ denotes the odometer, namely the number of particles jumping from o during the stabilisation of $B(L)$. We then have that

$$\bar{m}_L(o) = m_{B(L),\eta}(o) + w_L(o) \geq w_L(o). \quad (4)$$

Denote by $\bar{\mathcal{P}}^\mu$ the probability measure in the enlarged probability space of particles and ghosts, by \bar{E}^μ the expectation with respect to $\bar{\mathcal{P}}^\mu$, and recall that $\eta_{B(L),\infty}$ denotes the particle configuration after the stabilisation of $B(L)$.

It follows from the definition of ghosts that

$$\bar{E}^\mu[\bar{m}_L(o)] = \mu \sum_{x \in B(L)} G_{B(L)}(x, o), \quad (5)$$

$$\bar{E}^\mu[w_L(o)] = \sum_{x \in B(L)} \bar{\mathcal{P}}^\mu(\eta_{B(L),\infty}(x) = 1) G_{B(L)}(x, o). \quad (6)$$

The first identity holds since, in order to sample $\bar{m}_L(o)$, we need to start an independent simple random walk trajectory killed upon leaving $B(L)$ for each initial particle in $B(L)$ and count the number of visits at o . The second identity holds since a ghost starts from a vertex x if and only if a particle is at x after the stabilisation of $B(L)$.

Let $\varepsilon > 0$. We define the set $K_L := \{x \in B(L) : B(x, r) \subset B(L)\}$, choosing the constant r large enough so that, by Corollary 3.3, $\mathcal{P}^\mu(A_{x,L}) \geq \frac{d-1}{d^3} - \varepsilon$ for every $L \in \mathbb{N}$ and $x \in K_L$. Note that we can identify a set of vertices $U \subset K_L$ such that for each $x, y \in U$, if $x \neq y$ then $I_x \cap I_y = \emptyset$, and

$$\forall x \in K_L \quad \exists y \in U \quad d_G(x, y) \leq 2 \quad (7)$$

(indeed, if there was $x \in K_L$ such that $d_G(x, y) > 2$ for each $y \in U$, we could simply take x in U). Moreover, we observe that from the Markov property, for $x, y \in V$ such that $d_G(x, y) = k$ we have that

$$\frac{1}{d^k} G_{B(L)}(y, o) \leq G_{B(L)}(x, o) \leq d^k G_{B(L)}(y, o). \quad (8)$$

From (7) and (8) we deduce the crude bound

$$\begin{aligned} \sum_{x \in K_L} G_{B(L)}(x, o) &\leq \sum_{y \in U} \sum_{\substack{x \in K_L: \\ d(x, y) \leq 2}} G_{B(L)}(x, o) \\ &\leq \sum_{y \in U} \sum_{\substack{x \in K_L: \\ d(x, y) \leq 2}} d^2 G_{B(L)}(y, o) \leq d^2(d^2 + 1) \sum_{y \in U} G_{B(L)}(y, o). \end{aligned} \quad (9)$$

Finally, note that, by the Markov property, for any $y \in U$ we have that

$$\sum_{x \in I_y} \bar{\mathcal{P}}^\mu(\eta_\infty(x) = 1) G_{B(L)}(x, o) \geq \frac{1}{d} \bar{\mathcal{P}}^\mu(A_{y,L}) G_{B(L)}(y, o), \quad (10)$$

where we used that the expected number of visits at o from ghosts starting in I_y is at least the expected number of visits from one of the ghosts starting from I_y assuming that, if such ghost starts from $x \in I_y$ with $x \neq y$, then in the first step it jumps to y .

Thus, we have that, for any large enough L ,

$$\begin{aligned}
\bar{E}^\mu[\bar{m}_L(o)] &\geq \bar{E}^\mu[w_L(o)] \geq \sum_{y \in U} \sum_{x \in I_y} \bar{\mathcal{P}}^\mu(\eta_{B(L), \infty}(x) = 1) G_{B(L)}(x, o) \\
&\geq \frac{1}{d} \sum_{y \in U} \bar{\mathcal{P}}^\mu(A_{y, L}) G_{B(L)}(y, o) \\
&\geq \frac{1}{d} \left(\frac{d-1}{d^3} - \varepsilon \right) \sum_{y \in U} G_{B(L)}(y, o) \\
&\geq \frac{1}{d^3(d^2+1)} \left(\frac{d-1}{d^3} - \varepsilon \right) \sum_{x \in K_L} G_{B(L)}(x, o) \\
&\geq \frac{1}{10d^3(d^2+1)} \left(\frac{d-1}{d^3} - \varepsilon \right) \sum_{x \in B(L)} G_{B(L)}(x, o)
\end{aligned}$$

where the first inequality follows from (4), for the second inequality we use (6), for the third inequality we use (10), for the fourth inequality we use that r was given by Corollary 3.3, for the fifth inequality we use (9) and for the last inequality we use Lemma 3.4. Then, using (5) we deduce that

$$\bar{E}^\mu[\bar{m}_L(o)] \geq \frac{1}{10d^3(d^2+1)} \left(\frac{d-1}{d^3} - \varepsilon \right) \times \frac{1}{\mu} \bar{E}^\mu[\bar{m}_L(o)],$$

which, given that $\bar{E}^\mu[\bar{m}_L(o)] \neq 0$, implies that

$$\mu \geq \frac{1}{10d^3(d^2+1)} \left(\frac{d-1}{d^3} - \varepsilon \right).$$

This being true for every $\varepsilon > 0$ and for every $\mu > 0$ such that the SSM is active, we obtain that

$$\mu_c \geq \frac{d-1}{10d^6(d^2+1)}, \quad (11)$$

which concludes the proof that $\mu_c > 0$. \square

4 Stochastic comparison

In this section we introduce the activated random walks with instantaneous deactivation (ARWD) and we present one of our main results, namely the stochastic comparison between SSM and ARWD.

4.1 Definition of the ARWD process

Our strategy to obtain the upper bound is based on a comparison with another model, which we call activated random walks with instantaneous deactivation (ARWD). As for activated random walks, a configuration of the ARWD model on G can be represented by a vector $\eta : V \rightarrow \mathbb{N}_{\mathfrak{s}}$, with $\mathbb{N}_{\mathfrak{s}} = \mathbb{N}_0 \cup \{\mathfrak{s}\}$, where $\eta(x) = k \in \mathbb{N}_0$ 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 . 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 $U \subset V$ if $\eta(x) \in \{0, \mathfrak{s}\}$ for all $x \in U$.

For $\eta \in \mathbb{N}_{\mathfrak{s}}^V$, we denote by $|\eta| = \sum_{x \in V} |\eta(x)|$ the total number of particles in η (using the convention that $|\mathfrak{s}| = 1$). We equip the set $\mathbb{N}_{\mathfrak{s}}$ with the order $0 < \mathfrak{s} < 1 < 2 < \dots$ and, for $\eta, \eta' \in \mathbb{N}_{\mathfrak{s}}^V$, we write $\eta \leq \eta'$ if $\eta(x) \leq \eta'(x)$ for every $x \in V$.

Given $\eta \in \mathbb{N}_{\mathfrak{s}}^V$, $U \subset V$, we denote by η^U and η_U the configurations respectively obtained by waking up the particles in U or making them fall asleep, that is to say, for every $x \in V$,

$$\eta^U(x) = \begin{cases} 1 & \text{if } x \in U \text{ and } \eta(x) = \mathfrak{s}, \\ \eta(x) & \text{otherwise} \end{cases}, \quad \text{and} \quad \eta_U(x) = \begin{cases} \mathfrak{s} & \text{if } x \in U \text{ and } \eta(x) = 1, \\ \eta(x) & \text{otherwise.} \end{cases} \quad (12)$$

We now present the ARWD process, first informally, before giving the precise definition below. This process is associated with a certain fixed subset $A \subset V$, and is a variant of the loop representation of ARW introduced in [2]. At each time step of this process, an unstable site $x \in V$ is chosen, with a certain deterministic rule which depends on the past history of the process. This rule is given by what we call a toppling strategy:

Definition 4.1. A toppling strategy is an application $f : \cup_{t \geq 0} (\mathbb{N}_s^V)^{t+1} \rightarrow V$ such that for every $t \geq 0$ and every configurations η_0, \dots, η_t with η_t unstable, the site $f(\eta_0, \dots, \eta_t)$ is unstable in η_t .

Note that we will indeed use a strategy that depends not only on the current configuration but on the whole history of the process (see Section 6.1 for the precise definition of the toppling strategy used during the so-called ping-pong rally).

Then, if the chosen unstable site x belongs to A and contains one single particle, this particle falls asleep with probability $\lambda/(1 + \lambda)$. Otherwise, if it does not fall asleep, one particle at x makes a random walk on V until it finds an empty site of A and is left there, sleeping. This is a key difference with respect to the process considered in [2], where the particle was left active after its walk. Besides, along this walk made by the particle, the sleeping particles met are waken up, but only if they belong to a certain set which is given by the sleep mask, defined as follows:

Definition 4.2. A sleep mask is an application $g : \mathbb{N}_0 \times \cup_{t \geq 0} (\mathbb{N}_s^V)^{t+1} \rightarrow \mathcal{P}(V) = \{W \subset V\}$.

The sleep mask is used to ignore some reactivation events during the process. It depends on an integer parameter which is the color of the step: at each step we draw a random color in \mathbb{N}_0 and the set of sites that the walk can reactivate depends on this color (see Section 6.1 for more explanations). The idea is that this sleep mask enables us to gain some independence at certain key steps of the process. It is used in particular in the proofs of Proposition 4.1 and Lemma 6.5.

We now turn to the formal definition of the ARWD process.

Definition 4.3. Let $G = (V, E)$ be a finite connected graph, and let $A \subset V$ be fixed. We consider the state space $\mathcal{S} = \{\eta \in \mathbb{N}_s^V : |\eta| = |A|\}$. For a given parameter $\lambda \geq 0$, a toppling strategy f , a sleep mask g and an initial configuration $\eta_0 \in \mathcal{S}$, we call $\text{ARWD}(\lambda, A, f, g, \eta_0)$ the process $(\eta_t)_{t \geq 0}$ with state space \mathcal{S} and deterministic initial configuration η_0 , which is such that for every $t \geq 0$ and $\eta' \in \mathcal{S}$, we have

$$\mathbb{P}(\eta_{t+1} = \eta' \mid \eta_0, \dots, \eta_t) = \sum_{j \geq 0} \frac{1}{2^{j+1}} Q(\eta_t, f(\eta_0, \dots, \eta_t), g(j, \eta_0, \dots, \eta_t), \eta'),$$

where Q is defined as follows:

- The stable configurations are absorbing, i.e., for every $\eta \in \mathcal{S} \cap \{0, \mathfrak{s}\}^V$ and every $x \in V$ and $W \subset V$ we have $Q(\eta, x, W, \eta) = 1$.
- Let $\eta \in \mathcal{S}$ be an unstable configuration, let $x \in V$ and $W \subset V$. We write $p = \mathbb{1}_{\{x \in A\} \cap \{\eta(x)=1\}} \lambda / (1 + \lambda)$. Then, we have $Q(\eta, x, W, \eta') = p Q_{\text{sleep}}(\eta, x, \eta') + (1 - p) Q_{\text{walk}}(\eta, x, W, \eta')$, with Q_{sleep} and Q_{walk} defined as follows.

On the one hand, recalling the notation $\eta_{\{x\}}$ defined in (12), we have

$$Q_{\text{sleep}}(\eta, x, \eta_{\{x\}}) = 1,$$

that is to say, if the site x belongs to A and contains one single particle, this particle falls asleep with probability $p = \lambda / (1 + \lambda)$.

On the other hand, with probability $1 - p$ the particle at x walks until it finds an empty site of A and is left there sleeping, having woken up the sleeping particles met along its path, but only those which are in W . Thus, for every $y \in A$ and $R \subset V$, we have

$$Q_{\text{walk}}(\eta, x, W, \eta^{R \cap W} - \mathbb{1}_x + \mathfrak{s} \mathbb{1}_y) = P_x(X(\tau_B^+) = y \text{ and } \{X(t), 0 \leq t < \tau_B^+\} = R),$$

where, under P_x , $(X(t))_{t \geq 0}$ is a simple random walk on G started at x and τ_B^+ is its first hitting time of the set B of the empty sites of A , that is to say,

$$B = \{z \in A : (\eta - \mathbb{1}_x)(z) = 0\}$$

and $\tau_B^+ = \inf\{t \geq 1 : X(t) \in B\}$, which is almost surely finite because we work with configurations with exactly $|A|$ particles (and the graph is connected).

Note that this model is not abelian: to our knowledge, there is no monotonicity property of the number of steps before stabilisation, with respect to the initial configuration or with respect to the sleep mask.

Definition 4.4. For given parameters A, λ, f, g, η_0 , we denote by $T(A, \eta_0)$ the stabilisation time of the $\text{ARWD}(\lambda, A, f, g, \eta_0)$ process, that is to say, $T(A, \eta_0) = \inf\{t \geq 0 : \eta_t \text{ stable}\}$ where $(\eta_t)_{t \geq 0}$ is an instance of the $\text{ARWD}(\lambda, A, f, g, \eta_0)$ process. The dependence on λ, f and g is omitted to shorten the notation.

4.2 Connection between SSM and ARWD

Our interest for ARWD lies in the comparison established by Proposition 4.1 below, between A -stabilisation for the SSM and the stabilisation time for ARWD. For a technical reason we need to restrict ourselves to a subset $B \subset A$, and only consider the time for ARWD to stabilise this subset.

Proposition 4.1. *Let $G = (V, E)$ be a finite connected graph with maximal degree $\Delta \geq 2$ and let $B \subset A \subset V$ such that each site of B has degree at least 2 (in G) and B is totally disconnected, i.e., it does not contain any pair of neighbouring sites. Let $\eta \in \mathbb{N}_0^V$ be a SSM configuration such that $|\eta| = |A|$ and let $\tilde{\eta} \in \mathbb{N}_s^V$ be the ARW configuration with one particle on each site of A defined by $\tilde{\eta} = \mathbb{1}_U + s\mathbb{1}_{A \setminus U}$, where $U = \{x \in A : \eta(x) \geq 2\}$.*

Let $\lambda = \Delta^3$, let f be whatever toppling strategy and let g be whatever sleep mask. Then we have that $\|m_\eta^A\|$ (the number of half-topplings to A -stabilise the SSM configuration η , see Definition 2.2) stochastically dominates $T(B, \tilde{\eta}\mathbb{1}_B)$, the stabilisation time of the ARWD process on B starting from $\tilde{\eta}\mathbb{1}_B$ (see Definition 4.4).

The proof of this proposition is the object of Section 5. It relies on a coupling between the ARWD process and the A -stabilisation for the SSM such that the topplings made in ARWD are A -legal for the SSM. To construct this coupling we show that, conditioned on the evolution of the SSM process, each time that we start a walk at a site $x \in B$ which contains one single active particle, then with probability at least $1/(1+\lambda)$ the odometer at this site is odd, which implies that this walk is A -legal for the SSM.

4.3 Bound on the stabilisation time of ARWD

Once the problem is reduced to the study of the stabilisation time of ARWD, we show the following result:

Proposition 4.2. *Let $d \geq 2$. For every $\lambda \geq 0, \mu \in (0, 1)$ and $\rho \in (0, 1)$, there exists $\kappa > 0$ such that, for every $n \geq 1$, for every $A \subset \mathbb{Z}_n^d$ such that $|A| \geq \mu n^d$, there exists an toppling strategy f and a sleep mask g such that, starting with a random initial configuration η_0 with one particle on each site of A , each particle being active with probability ρ and sleeping with probability $1 - \rho$, independently of other particles, we have that $1 + T(A, \eta_0)$ dominates a geometric variable with parameter $\exp(-\kappa n^d)$.*

We prove this result in Section 6, adapting the strategy developed in the context of ARW by [2]. This strategy consists in a multiscale argument, constructing a hierarchical structure on the set A with “well connected” components at the first level of the hierarchy and pairs of clusters which merge at each level of the hierarchy. We first show that the clusters at the first level have a stabilisation time which is exponentially large in their size, before showing that this property is inherited at further levels of the hierarchy, loosing only a little in the constant in the exponential.

Then, in Section 7, we explain how Theorem 1.3 follows from Propositions 4.1 and 4.2.

5 Connection SSM-ARWD: proof of Proposition 4.1

This section is devoted to the proof of Proposition 4.1 which relates the number of toppling for the A -stabilisation of the SSM to the stabilisation time of the ARWD process.

5.1 Strategy: sleep if the odometer is even

To prove this proposition, we construct a so-called coupling between the ARWD process and the SSM, at least up to a certain stopping time T . More precisely, we construct the ARWD process using a field of

instructions τ from which the random walks used at each step are read, and uniform random variables which are used to decide whether the site falls asleep or performs a random walk.

We want to show that, following these dynamics, we obtain at each time step a toppling sequence which is A -legal for the SSM. The key point there is that, to show that the sequence of topplings performed is A -legal, we need to show that each time that we choose a site $x \in A$ with only one particle and that the particle at x makes a random walk, the odometer at this site is odd (which is a condition for the toppling at x to be legal).

The idea is to show that, in this case, knowing the configurations up to this time (but not the field τ nor the odometers), the odometer at x is odd with probability at least $1/(1+\lambda)$, allowing us to couple the Bernoulli variable which decides whether the particle sleeps so that when the odometer is even this variable tells to sleep and when the odometer is odd then we have a sleep with a certain probability, so that the overall probability to sleep, whatever the odometer, is $\lambda/(1+\lambda)$. This bound on the conditional probability that the odometer is odd is the content of Lemma 5.2 below.

The proof uses the fact that if $x \in A$ is active and $\eta_t(x) = 1 \neq \eta_0(x)$, then it entails that the site x was visited since the last time that we chose x (because upon its last arrival on x the particle fell asleep, so it must have been reactivated by the visit of another particle), or since the beginning if x was not chosen yet. Then, once we know that the site was visited, we can bound from below the probability that the number of visits was odd, by showing that one visit can be added with a reasonable probabilistic cost.

5.2 Coupled ARWD dynamics

Definition 5.1. Let $G = (V, E)$ be a finite connected graph with maximal degree Δ , let $B \subset A \subset V$, let f be a toppling strategy, g a sleep mask, let $\eta \in \mathbb{N}_0^V$ be a fixed SSM configuration with $|\eta| = |A|$ and let $U = \{x \in A : \eta(x) \geq 2\}$.

Then the coupled ARWD process $(\eta_t, \alpha_t)_{t \geq 0}$ associated with these parameters is defined as follows.

Using an arbitrary order on V , we define another toppling strategy f' by writing, for every $t \geq 0$ and every $\eta_0, \dots, \eta_t \in \mathbb{N}_s^V$,

$$f'(\eta_0, \dots, \eta_t) = \begin{cases} \min \{x \in V \setminus A : \eta_t(x) \geq 1\} & \text{if not empty;} \\ \min \{x \in A : \eta_t(x) \geq 2\} & \text{else if not empty;} \\ f(\eta_S \mathbb{1}_B, \dots, \eta_t \mathbb{1}_B) \text{ with } S = \min \{s \leq t : \eta_s \leq \mathbb{1}_A\} & \text{else if } \eta_t \text{ is unstable in } B; \\ f(\eta_0, \dots, \eta_t) & \text{else if } \eta_t \text{ is unstable;} \\ \min V & \text{if } \eta_t \text{ is stable.} \end{cases}$$

Note that the value $\min V$ in the last case has no importance, and nor does the value in the penultimate case, apart from the fact that it is an unstable site in η_t , to match the definition of a toppling strategy.

We also define a modified sleep mask g' by writing, for any $t, j \geq 0$ and $\eta_0, \dots, \eta_t \in \mathbb{N}_s^V$,

$$g'(j, \eta_0, \dots, \eta_t) = \begin{cases} \emptyset & \text{if } \eta_t \not\leq \mathbb{1}_A, \\ g(j, \eta_S \mathbb{1}_B, \dots, \eta_t \mathbb{1}_B) \text{ with } S = \min \{s \leq t : \eta_s \leq \mathbb{1}_A\} & \text{otherwise.} \end{cases}$$

In other words, as long as there is not one particle on each site of A , no sleeping particle can be awakened, while once this is the case we follow the sleep mask g .

Let $\lambda = \Delta^3$. Let $\tau = (\tau^{x,j})_{x \in V, j \geq 0}$ be a random field of toppling instructions distributed according to \mathcal{P} , that is to say, each instruction at x is a jump to a uniformly chosen neighbour of x (as defined in Section 2). We also consider two independent sequences of i.i.d. variables $(Z_t)_{t \geq 0}$ and $(J_t)_{t \geq 0}$, independent of τ , where Z_0 has uniform distribution on $[0, 1]$ and $1 + J_0$ is geometric with parameter $1/2$.

Now, we define recursively a sequence of configurations $(\eta_t)_{t \geq 0}$ and a sequence of half-toppling sequences $(\alpha_t)_{t \geq 0}$. For every $t \geq 0$ we will denote by m_t the odometer of α_t , i.e., $m_t = m_{\alpha_t}$.

We start by taking $\eta_0 = \eta_A$, that is to say, the configuration obtained from η by declaring sleeping the particles which stand alone on a site of A , and we let α_0 be the empty toppling sequence. Then, let $t \geq 0$ be such that η_0, \dots, η_t and α_t are constructed.

If $\eta_t \leq \mathbb{1}_{A \setminus B} + \mathfrak{s}\mathbb{1}_B$, that is to say, if each site of A already contains one particle with all the particles on B being asleep, then we do nothing and simply let $\eta_{t+1} = \eta_t$ and $\alpha_{t+1} = \alpha_t$.

Assume now that $\eta_t \not\leq \mathbb{1}_{A \setminus B} + \mathfrak{s}\mathbb{1}_B$, and let $X_t = f'(\eta_0, \dots, \eta_t)$ and $W_t = g'(J_t, \eta_0, \dots, \eta_t)$. Writing

$$\pi_t = \mathbb{P}(m_t(X_t) \text{ odd} \mid \eta_1, \dots, \eta_t),$$

we consider the variable I_t which can take two values, sleep and walk, defined as

$$I_t = \begin{cases} \text{sleep} & \text{if } X_t \in B, \eta_t(X_t) = 1 \text{ and } m_t(X_t) \text{ is even,} \\ \text{sleep} & \text{if } X_t \in B, \eta_t(X_t) = 1, m_t(X_t) \text{ is odd and } Z_t > \frac{1}{(1+\lambda)\pi_t}, \\ \text{walk} & \text{otherwise.} \end{cases}$$

On the one hand, if $I_t = \text{sleep}$, then the particle at X_t falls asleep: we let $\eta_{t+1} = (\eta_t)_{\{X_t\}}$ and $\alpha_{t+1} = \alpha_t$.

On the other hand, if $I_t = \text{walk}$, then we let the particle at X_t walk with successive topplings drawn from τ and applied to (η_t, m_t) , until it reaches a site of A which is empty (in the configuration $\eta_t - \mathbb{1}_{X_t}$). Let Y_t be the first such site reached by the particle, let β_t be the sequence of sites toppled along this walk, and let R_t be the set of sites visited by this walk, namely $R_t = \{x \in V : m_{\beta_t}(x) \geq 1\}$. We then let $\eta_{t+1} = (\eta_t)^{R_t \cap W_t} - \mathbb{1}_{X_t} + \mathfrak{s}\mathbb{1}_{Y_t}$, that is to say, the particle falls asleep at Y_t at the end of its walk, and we define α_{t+1} as the concatenation of α_t and β_t .

The process $(\eta_t, \alpha_t)_{t \geq 0}$ defined above is called the coupled ARWD process associated with G, A, B, f, g and η_0 .

Note that, with this toppling strategy f' and this sleep mask g' , the evolution of the process consists in the following three successive stages, which correspond to the three first cases in the definition of f' :

- During a first stage, we move the particles from $V \setminus A$ to vacant sites of A until all particles are in A , and particles cannot wake up during this stage.
- Then comes a second stage during which we move particles from sites of A with two or more particles to empty sites of A , until each site of A contains exactly one particle. Reactivation events are also ignored during this stage.

At the end of this second stage, the particles in U are all active, because a site in U has no chance to fall asleep during these first two phases, while the particles in $A \setminus U$ are all sleeping, that is to say, we have reached the configuration $\tilde{\eta} = \mathbb{1}_U + \mathfrak{s}\mathbb{1}_{A \setminus U}$.

- Then, during the third and last stage, active particles on sites of B perform loops and come back to their position, sleeping, until all particles on B are sleeping. Then, once all particles on B are sleeping, the process does not evolve anymore.

5.3 The coupling lemma

We now show that the coupled process defined above enjoys the following properties:

Lemma 5.1. *Let $G = (V, E)$ be a finite connected graph with maximal degree $\Delta \geq 2$, let $B \subset A \subset V$ with B totally disconnected and $d_x \geq 2$ for every $x \in B$. Let $\eta \in \mathbb{N}_0^V$ be a SSM configuration with $|\eta| = |A|$. We also define the ARW configuration $\tilde{\eta} = \mathbb{1}_U + \mathfrak{s}\mathbb{1}_{A \setminus U}$, where $U = \{x \in A : \eta(x) \geq 2\}$. Let $\lambda = \Delta^3$, let f be a toppling strategy and let g be a sleep mask. Let $(\eta_t, \alpha_t)_{t \geq 0}$ be the coupled ARWD process defined above, with f' the modified toppling strategy, g' the modified sleep mask and τ the associated array of instructions. Then, defining the two stopping times $S = \inf\{t \geq 0 : \eta_t \leq \mathbb{1}_A\}$ and $T = \inf\{t \geq S : \eta_t \text{ stable in } B\}$, the following hold:*

- (i) *the stopping times S and T are almost surely finite;*
- (ii) *for every $t \geq 0$, the half-toppling sequence α_t is A -legal for η (with respect to the field of instructions τ);*
- (iii) *the stopped process $(\eta_{t \wedge T})_{t \geq 0}$ is distributed as the ARWD($\lambda, A, f', g', \eta_A$) process stopped at time T ;*

(iv) we have $\eta_S = \tilde{\eta}$ (if $S < \infty$) and $(\eta_{S+t}\mathbb{1}_B)_{t \geq 0}$ is distributed as the ARWD($\lambda, B, f, g, \tilde{\eta}\mathbb{1}_B$) process;

(v) for every $t \geq 0$ we have $|m_t| \geq \mathbb{1}_{\{t \geq S\}}((t \wedge T) - S + N_t)$, where $N_t = |\{x \in A : \eta_t(x) = 1\}|$.

Before turning to the proof of Lemma 5.1, we first explain how Proposition 4.1 follows from it.

Proof of Proposition 4.1. Recall that the notation and hypotheses of Proposition 4.1 are the same as that of Lemma 5.1. First, note that point (i) ensures that the half-toppling sequence α_T is almost surely well defined. Point (ii) of Lemma 5.1 tells us that this half-toppling sequence is A -legal for η . By definition of m_η^A , this implies that $m_\eta^A \geq m_{\alpha_T} = m_T$, whence $\|m_\eta^A\| \geq |m_T|$. Yet, by virtue of (v), we have $|m_T| \geq T - S$. Then, point (iv) entails that $T - S$ is distributed as $T(B, \tilde{\eta}\mathbb{1}_B)$. Combining all this, we obtain that $\|m_\eta^A\|$ stochastically dominates $T(B, \tilde{\eta}\mathbb{1}_B)$, as claimed. \square

5.4 Proof of the coupling lemma

As explained in Section 5.1, one key point to obtain Lemma 5.1, and in particular (iii) and (iv), is to establish the following:

Lemma 5.2. *With the notation and hypotheses of Lemma 5.1, for every $t \geq 0$ and $x \in B$ we have the implication*

$$\{\eta_t(x) = 1\} \subset \left\{ \mathbb{P}(m_t(x) \text{ odd} \mid \eta_1, \dots, \eta_t) \geq \frac{1}{1 + \lambda} \right\}.$$

Let us postpone the proof of Lemma 5.2 to Section 5.5 and first prove Lemma 5.1 admitting this Lemma.

Proof of Lemma 5.1. First, point (i) follows from our assumptions that $|\eta| = |A|$ and that G is connected, along with the fact that $\lambda > 0$.

Point (ii) follows from the fact that, when we topple a site of B which contains only one particle, we never topple it if the odometer at this point is even (because in this case $I_t = \text{sleep}$ and we do not topple, marking the particle as sleeping) and by definition of f' we never topple a site of $A \setminus B$ which contains only one particle (because we stop at T , i.e., before being forced to do so). Also, recall that topplings at sites in $V \setminus A$ with at least one particle or at sites of A with at least two particles are always A -legal. Yet, along a random walk, apart from the very first toppling, we never topple a site of A which contains only one particle, because we stop as soon as the particle reaches a site of A where it is alone. Hence, point (ii) holds.

We now prove point (iii). To this end, we first show that, for every $t \geq 0$, we have the implication

$$\{t < T\} \subset \left\{ \mathbb{P}(I_t = \text{sleep} \mid \eta_1, \dots, \eta_t) = \frac{\lambda}{1 + \lambda} \mathbb{1}_{\{X_t \in B, \eta_t(X_t) = 1\}} \right\}. \quad (13)$$

Let $t \geq 0$ and let $\bar{\eta}_0, \dots, \bar{\eta}_t \in \mathcal{S}$ with $\bar{\eta}_0 = \eta_0$. Let us write $\mathcal{A} = \{\eta_1 = \bar{\eta}_1, \dots, \eta_t = \bar{\eta}_t\}$ and assume that $\mathbb{P}(\mathcal{A}) > 0$. Let $x = f'(\bar{\eta}_0, \dots, \bar{\eta}_t)$. We also assume that $\bar{\eta}_t \not\leq \mathbb{1}_{A \setminus B} + \mathfrak{s}\mathbb{1}_B$, so that $\mathcal{A} \subset \{t < T\}$. By definition of our strategy f' , this implies that $x \in B \cup (V \setminus A)$.

First, if $x \notin A$ or $\bar{\eta}_t(x) \geq 2$, then $\mathbb{P}(I_t = \text{sleep} \mid \mathcal{A}) = 0$ by definition of I_t , which shows (13) in this case.

Assume now that $x \in B$ and that $\bar{\eta}_t(x) = 1$. Let $\pi = \mathbb{P}(m_t(x) \text{ odd} \mid \mathcal{A})$. It follows from Lemma 5.2 that $\pi \geq 1/(1 + \lambda)$. Hence, we have

$$\mathbb{P}(I_t = \text{walk} \mid \mathcal{A}) = \mathbb{P}\left(m_t(x) \text{ odd}, Z_t \leq \frac{1}{(1 + \lambda)\pi} \mid \mathcal{A}\right) = \pi \times \mathbb{P}\left(Z_t \leq \frac{1}{(1 + \lambda)\pi}\right) = \pi \times \frac{1}{(1 + \lambda)\pi} = \frac{1}{1 + \lambda}.$$

Thus, we proved (13).

Let now $t \geq 0$, let $\bar{\eta}_0 = \eta_0$, let $\bar{\eta}_1, \dots, \bar{\eta}_t \in \mathcal{S}$ and take an integer $j \geq 0$. Let us assume that the event $\mathcal{A} = \{\eta_1 = \bar{\eta}_1, \dots, \eta_t = \bar{\eta}_t, J_t = j\}$ has positive probability. We also assume that $\bar{\eta}_t \not\leq \mathbb{1}_{A \setminus B} + \mathfrak{s}\mathbb{1}_B$, so that $\mathcal{A} \subset \{t < T\}$. Let us write $x = f'(\bar{\eta}_0, \dots, \bar{\eta}_t)$ and $W = g'(j, \bar{\eta}_0, \dots, \bar{\eta}_t)$. To obtain (iii), we have to show that for every $\eta' \in \mathcal{S}$,

$$\mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}) = Q(\bar{\eta}_t, x, W, \eta'), \quad (14)$$

with Q given by Definition 4.3. Let $\eta' \in \mathcal{S}$. Writing $p = \lambda/(1 + \lambda)$ if $x \in A$ and $\bar{\eta}_t(x) = 1$ (in which case we have in fact $x \in B$, by definition of our strategy f'), and $p = 0$ otherwise, it follows from (13) that $\mathbb{P}(I_t = \text{sleep} \mid \mathcal{A}) = p$. Thus, we have

$$\mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}) = p \mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}, I_t = \text{sleep}) + (1 - p) \mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}, I_t = \text{walk}).$$

Therefore, (14) follows if we show that

$$\mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}, I_t = \text{sleep}) = Q_{\text{sleep}}(\bar{\eta}_t, x, \eta') \quad \text{and} \quad \mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}, I_t = \text{walk}) = Q_{\text{walk}}(\bar{\eta}_t, x, W, \eta').$$

First, if $\eta' = (\bar{\eta}_t)_{\{x\}}$ then we have

$$\mathbb{P}(\eta_{t+1} = \eta' \mid \mathcal{A}, I_t = \text{sleep}) = 1 = Q_{\text{sleep}}(\bar{\eta}_t, x, \eta'),$$

with both members being equal to 0 for all other values of η' . The case of Q_{walk} works similarly, using that the instructions in τ used to construct the random walk at step t are independent of \mathcal{A} . Thus, we proved (14), which concludes the proof of (iii).

Let us now check (iv). First note that, since $|\eta| = |A|$, the definition of S implies that at time S there is exactly one particle on each site of A . Then, note that by definition of U , for each $x \in A \setminus U$ we have $\eta_0(x) \in \{0, \mathfrak{s}\}$. For each of these sites which starts empty, the first particle which visits the site settles there, sleeping. Moreover, recall that by definition of the sleep mask g' , sleeping particles cannot be awakened before time S . Therefore, we must have $\eta_S(x) = \mathfrak{s}$ for every $x \in A \setminus U$. Besides, note that the sites of U start with at least two particles in η_0 and before time S they are never toppled when they contain only one particle (by definition of our strategy f'). Thus, these sites cannot become stable before time S , so that we have $\eta_S(x) = 1$ for every $x \in U$. In the end, provided that $S < \infty$, the only possibility is that $\eta_S = \tilde{\eta}$.

Then, recalling the definition of f' and g' (see Definition 5.1), note that for every $t \geq 0$, we have

$$\{S \leq t < T\} \subset \{f'(\eta_0, \dots, \eta_t) = f(\eta_S \mathbb{1}_B, \dots, \eta_t \mathbb{1}_B) \quad \text{and} \quad g'(J_t, \eta_0, \dots, \eta_t) = f(J_t, \eta_S \mathbb{1}_B, \dots, \eta_t \mathbb{1}_B)\}.$$

Combined with point (iii), this implies (iv).

We now turn to point (v), that we prove by induction on t . Recall that by construction no site of A starts with one single active particle, whence $N_0 = 0$. Thus, for $t = 0$ there is nothing to prove, both sides being equal to 0. Let now $t \geq 0$ be such that (v) holds at time t and let us check that (v) also holds at time $t + 1$.

If $t \geq T$ then (v) at time $t + 1$ is identical to (v) at time t .

If $t + 1 < S$ then there is nothing to prove, the right-hand side being equal to zero.

If $t + 1 = S$ then we have to show that $|m_S| \geq N_S$. Recall that (iv) ensures that $\eta_S = \tilde{\eta}$, which entails that $N_S = |\{x \in A : \tilde{\eta}(x) = 1\}| = |U|$. Yet, for every $x \in U$ we have $\eta_S(x) = 1 < \eta_0(x)$, whence $m_S(x) \geq 1$: each site of U must have been toppled at least once before time S because it has lost at least one particle. Hence, we have $|m_S| \geq |U| = N_S$.

Assume now that $S \leq t < T$. Thus, we have $X_t \in B$ and $\eta_t(X_t) = 1$. If moreover $I_t = \text{sleep}$, then we have $m_{t+1} = m_t$ and $N_{t+1} = N_t - 1$, so that both members of the inequality do not change between rank t and rank $t + 1$. Otherwise if $I_t = \text{walk}$, we have

$$\begin{aligned} N_{t+1} &= N_t - 1 + |R_t \cap W_t \cap \{y \in A : \eta_t(y) = \mathfrak{s}\}| \\ &\leq N_t - 1 + |R_t \setminus \{X_t\}| \leq N_t - 1 + |\beta_t| - 1 = N_t + |m_{t+1}| - |m_t| - 2, \end{aligned}$$

which, combined with (v) at rank t , yields (v) at rank $t + 1$. This concludes the proof of Lemma 5.1, given Lemma 5.2 that we prove in the next section. \square

5.5 Proof of Lemma 5.2: adding jump instructions to make the odometer odd

We now prove Lemma 5.2. Note that the proof crucially relies on the facts that the coupling is restricted to the totally disconnected set B and that in the dynamics considered, the particles first move to occupy all the sites of A .

Proof of Lemma 5.2. Let $t \geq 0$, let $\bar{\eta}_0 = \eta_0$, let $\bar{\eta}_1, \dots, \bar{\eta}_t \in \mathcal{S}$ and let also $j_0, \dots, j_t \in \mathbb{N}_0$. Let us write $\mathcal{A} = \{\eta_1 = \bar{\eta}_1, \dots, \eta_t = \bar{\eta}_t\}$ and $\mathcal{B} = \{J_0 = j_0, \dots, J_t = j_t\}$ and assume that $\mathbb{P}(\mathcal{A} \cap \mathcal{B}) > 0$. Let $x \in B$ and assume that $\bar{\eta}_t(x) = 1$. We want to show that

$$\mathbb{P}(m_t(x) \text{ odd} \mid \mathcal{A} \cap \mathcal{B}) \geq \frac{1}{\Delta^3 + 1} = \frac{1}{\lambda + 1}. \quad (15)$$

Note that, in the process, there are three random ingredients: the field of instructions τ , the uniform variables (Z_t) and the colors (J_t) , recalling that the initial configuration is fixed and equal to η_0 . Thus, on the event \mathcal{B} we may express the process until time t with deterministic functions of τ and $\mathbf{Z} = (Z_0, \dots, Z_{t-1})$, writing $\eta_1^{\tau, \mathbf{Z}}, \dots, \eta_t^{\tau, \mathbf{Z}}, m_0^{\tau, \mathbf{Z}}, \dots, m_t^{\tau, \mathbf{Z}}$ and $R_0^{\tau, \mathbf{Z}}, \dots, R_{t-1}^{\tau, \mathbf{Z}}$, with the convention $R_s = \emptyset$ when $I_s = \text{sleep}$.

Let us denote the set of all possible fields of instructions by

$$\mathcal{I} = \{\bar{\tau} = (\bar{\tau}^{z,j})_{z \in V, j \geq 1} : \forall z \in V, \forall j \geq 1, \bar{\tau}^{z,j} \in \{\tau_{zz'}, z' \sim z\}\}$$

and, for every $\mathbf{z} = (z_0, \dots, z_{t-1}) \in [0, 1]^t$, let us consider

$$\mathcal{I}^{\mathbf{z}} = \{\bar{\tau} \in \mathcal{I} : \eta_1^{\bar{\tau}, \mathbf{z}} = \bar{\eta}_1, \dots, \eta_t^{\bar{\tau}, \mathbf{z}} = \bar{\eta}_t\},$$

so that

$$\mathbb{P}(\mathcal{A} \mid \mathcal{B}) = \int_{[0,1]^t} \mathcal{P}(\mathcal{I}^{\mathbf{z}}) d^t \mathbf{z}.$$

Similarly, defining

$$\mathcal{I}_{\text{odd}}^{\mathbf{z}} = \{\bar{\tau} \in \mathcal{I} : \eta_1^{\bar{\tau}, \mathbf{z}} = \bar{\eta}_1, \dots, \eta_t^{\bar{\tau}, \mathbf{z}} = \bar{\eta}_t \text{ and } m_t^{\bar{\tau}, \mathbf{z}}(x) \text{ odd}\},$$

we have

$$\mathbb{P}(\mathcal{A}, m_t(x) \text{ odd} \mid \mathcal{B}) = \int_{[0,1]^t} \mathcal{P}(\mathcal{I}_{\text{odd}}^{\mathbf{z}}) d^t \mathbf{z}.$$

Hence, to obtain (15) it is enough to show that for every $\mathbf{z} = (z_0, \dots, z_{t-1}) \in [0, 1]^t$, we have

$$\mathcal{P}(\mathcal{I}_{\text{odd}}^{\mathbf{z}}) \geq \frac{1}{\Delta^3 + 1} \mathcal{P}(\mathcal{I}^{\mathbf{z}}). \quad (16)$$

Henceforth, we fix $\mathbf{z} = (z_0, \dots, z_{t-1}) \in [0, 1]^t$. We want to use a counting argument, but a small difficulty lies in the fact that the probability \mathcal{P} has no atoms: the probability of every single field $\bar{\tau}$ is zero. Thus, we reason instead with truncated fields: for every $\bar{\tau} \in \mathcal{I}$ and $m \in \mathbb{N}_0^V$, let us denote the m -truncation of the array $\bar{\tau}$ by

$$\bar{\tau}^{\leq m} = (\bar{\tau}^{z,j})_{z \in V, j \leq m(z)}.$$

For every $m \in \mathbb{N}_0^V$, we call $\mathcal{I}_m^{\mathbf{z}}$ the subset of truncated arrays of $\mathcal{I}^{\mathbf{z}}$ with odometer m , namely,

$$\mathcal{I}_m^{\mathbf{z}} = \{\bar{\tau} = (\bar{\tau}^{z,j})_{z \in V, j \leq m(x)} : \exists \bar{\tau}_0 \in \mathcal{I}^{\mathbf{z}}, \bar{\tau}_0^{\leq m} = \bar{\tau} \text{ and } m_t^{\bar{\tau}_0, \mathbf{z}} = m\},$$

so that, writing $w(m) = \mathcal{P}(\bar{\tau}^{\leq m} \in \mathcal{I}_m^{\mathbf{z}}) = \mathcal{P}(\bar{\tau} \in \mathcal{I}^{\mathbf{z}}, m_t^{\bar{\tau}, \mathbf{z}} = m)$, we have

$$\mathcal{P}(\mathcal{I}^{\mathbf{z}}) = \sum_{m \in \mathbb{N}_0^V} w(m) \quad \text{and} \quad \mathcal{P}(\mathcal{I}_{\text{odd}}^{\mathbf{z}}) = \sum_{\substack{m \in \mathbb{N}_0^V : \\ m(x) \text{ odd}}} w(m). \quad (17)$$

We now claim that, for every $m \in \mathbb{N}_0^V$ such that $m(x)$ is even, we have

$$w(m) \leq \Delta^2 \sum_{y \sim x} w(m + \mathbb{1}_x + \mathbb{1}_y). \quad (18)$$

Admitting this claim for now and plugging it into (17), we can write

$$\begin{aligned} \mathcal{P}(\mathcal{I}^{\mathbf{z}} \setminus \mathcal{I}_{\text{odd}}^{\mathbf{z}}) &= \sum_{\substack{m \in \mathbb{N}_0^V : \\ m(x) \text{ even}}} w(m) \leq \sum_{\substack{m \in \mathbb{N}_0^V : \\ m(x) \text{ even}}} \Delta^2 \sum_{y \sim x} w(m + \mathbb{1}_x + \mathbb{1}_y) \\ &= \Delta^2 \sum_{y \sim x} \sum_{\substack{m' \in \mathbb{N}_0^V : \\ m'(x) \text{ odd}, m'(y) \geq 1}} w(m') \leq \Delta^3 \mathcal{P}(\mathcal{I}_{\text{odd}}^{\mathbf{z}}), \end{aligned}$$

which implies the desired inequality (16).

Thus, to conclude the proof of the lemma, there only remains to prove the claim (18). Note that, for every $m \in \mathbb{N}_0^V$, we have

$$w(m) = \sum_{\bar{\tau} \in \mathcal{I}_m^{\mathbf{z}}} \mathcal{P}(\tau^{\leq m} = \bar{\tau}) = |\mathcal{I}_m^{\mathbf{z}}| \prod_{z \in V} \left(\frac{1}{d_z} \right)^{m(z)},$$

where we recall that d_z is the degree of z in the graph G , so that the claim (18) will follow if we show that for every $m \in \mathbb{N}_0^V$ such that $m(x)$ is even,

$$|\mathcal{I}_m^{\mathbf{z}}| \leq \sum_{y \sim x} |\mathcal{I}_{m+\mathbb{1}_x+\mathbb{1}_y}^{\mathbf{z}}|. \quad (19)$$

We now fix $m \in \mathbb{N}_0^V$ with $m(x)$ even and we construct an injective mapping from $\mathcal{I}_m^{\mathbf{z}}$ to $\cup_{y \sim x} \mathcal{I}_{m+\mathbb{1}_x+\mathbb{1}_y}^{\mathbf{z}}$. The idea is to modify the array of instructions by adding two jump instructions at x and at a well chosen neighbouring site y , at a well chosen position in the array.

Note that our assumption that $\bar{\eta}_t(x) = 1$ implies that x was visited at least once before time t , because by definition $\bar{\eta}_0(x) \neq 1$. Thus, for every $\bar{\tau} \in \mathcal{I}_m^{\mathbf{z}}$ we can define $\sigma(\bar{\tau}) = \max\{s < t : R_s^{\bar{\tau}, \mathbf{z}} \ni x\}$.

Now, for every $\bar{\tau} \in \mathcal{I}_m^{\mathbf{z}}$, writing $\sigma = \sigma(\bar{\tau})$ and $u = u(\bar{\tau}) = f'(\bar{\eta}_0, \dots, \bar{\eta}_\sigma)$ we define, using an arbitrary order on the sites of V ,

$$y(\bar{\tau}) = \min \left(\{y \sim x : y \notin A\} \cup \{y \sim x : \bar{\eta}_\sigma(y) - \mathbb{1}_u(y) \neq 0\} \right).$$

Let us explain why the above set cannot be empty, distinguishing between two cases.

One the one hand, if the above set is empty and $u \sim x$, then we have $u \in A$ and $\bar{\eta}_\sigma(u) = 1$, while the other neighbours of x are empty in $\bar{\eta}_\sigma$, but also belong to A . This situation is impossible due to the priority rules of our toppling strategy f' and to the fact that $|\eta_0| = |A|$, because a site of A with one single particle can be chosen by f' only once that each site of A is occupied by one particle (in other words, only after time S). Note that we used here that x has at least two neighbours, following our assumption that each site of B has degree at least 2 in the graph G .

One the other hand, let us now assume that u is not a neighbour of x and that the set in the definition of $y(\bar{\tau})$ is empty. Then for every y neighbour of x we have $y \in A$ and $\bar{\eta}_\sigma(y) = 0$. Note that it also follows from the priority rules of f' (and from the fact that $|\eta_0| = |A|$) that an occupied site of A cannot become empty in a later configuration of the ARWD process. Therefore, our assumptions imply that no particle can have visited x before time σ , i.e., $m_{\sigma}^{\bar{\tau}, \mathbf{z}}(x) = 0$ (otherwise this particle would have visited one of the neighbours of x , and would have settled there). Moreover, since $\bar{\eta}_t(x) = 1$ and σ is the last visit of x before t , we have $\bar{\eta}_{\sigma+1}(x) = \bar{\eta}_t(x) = 1$. Yet, by definition of σ the site x must be visited at step σ , so that we must have $u = x$ and $\bar{\eta}_\sigma(x) = 2$, so that in the next step one particle at x jumps to one of the neighbouring sites and then terminates its walk there (because we assumed that all the neighbouring sites are in A and empty in $\bar{\eta}_\sigma$). This entails that $m_{\sigma+1}^{\bar{\tau}, \mathbf{z}}(x) = 1$. Yet, by definition of σ we have $m(x) = m_t^{\bar{\tau}, \mathbf{z}}(x) = m_{\sigma+1}^{\bar{\tau}, \mathbf{z}}(x)$, whence $m(x) = 1$, which contradicts our assumption that $m(x)$ is even. Therefore, we checked that the set in the minimum defining $y(\bar{\tau})$ is never empty.

Now, once $y(\bar{\tau})$ is well defined, we are ready to construct our application $\varphi : \mathcal{I}_m^{\mathbf{z}} \rightarrow \cup_{y \sim x} \mathcal{I}_{m+\mathbb{1}_x+\mathbb{1}_y}^{\mathbf{z}}$. For every $\bar{\tau} \in \mathcal{I}_m^{\mathbf{z}}$, every $z \in V$ and every $1 \leq j \leq (m + \mathbb{1}_x + \mathbb{1}_y)(z)$, writing $\sigma = \sigma(\bar{\tau})$ and $y = y(\bar{\tau})$, we set

$$\varphi(\bar{\tau})^{z,j} = \begin{cases} \bar{\tau}^{z,j} & \text{if } z \notin \{x, y\}, \\ \bar{\tau}^{z,j} & \text{if } z \in \{x, y\} \text{ and } j \leq m_{\sigma}^{\bar{\tau}, \mathbf{z}}(z), \\ \tau_{xy} & \text{if } z = x \text{ and } j = m_{\sigma}^{\bar{\tau}, \mathbf{z}}(x) + 1, \\ \tau_{yx} & \text{if } z = y \text{ and } j = m_{\sigma}^{\bar{\tau}, \mathbf{z}}(y) + 1, \\ \bar{\tau}^{z,j-1} & \text{if } z \in \{x, y\} \text{ and } j \geq m_{\sigma}^{\bar{\tau}, \mathbf{z}}(z) + 2. \end{cases}$$

The application φ consists in adding two jump instructions from x to $y = y(\bar{\tau})$ and from y to x so that these instructions add one round trip between these two vertices in the path of the random walk at step σ , thus changing the parity of the odometer m_t at x and y , but leaving the resulting configuration unchanged.

Note that, using the field of instructions $\bar{\tau}$ (i.e., before applying φ), the walk at step σ did not terminate at x because, if it was the case, then the particle would have settled at x , sleeping, i.e., we would have $\bar{\eta}_{\sigma+1}(x) = \mathfrak{s}$, which would imply that $\bar{\eta}_t(x) = \mathfrak{s}$, since x is no longer visited in the steps between $\sigma + 1$ and t (by definition of σ). This would contradict our assumption that $\bar{\eta}_t(x) = 1$. Hence, since the walk at step σ visits x but does not terminate there, although $x \in A$, we must have $\bar{\eta}_\sigma(x) - \mathbb{1}_u(x) \neq 0$. As a consequence, also after adding the two jump instructions, the walk at step σ still cannot stop at x .

Moreover, by definition of y , we have either $y \notin A$ or $\bar{\eta}_\sigma(y) - \mathbb{1}_u(y) \neq 0$. Thus, the walk cannot terminate neither at y . Hence, the two added instructions are indeed used during the step σ : if before applying φ , the walk visited x before y , then φ adds a round trip $x \rightarrow y \rightarrow x$, while if it visited y first, then it adds a round trip $y \rightarrow x \rightarrow y$.

Moreover, this added round trip does not change anything else than this increase in the odometer at x and y , so that for every $s \leq t$ we have $\eta_s^{\varphi(\bar{\tau}), \mathbf{z}} = \eta_s^{\bar{\tau}, \mathbf{z}} = \bar{\eta}_s$. Indeed, by definition of σ , the site x is never visited again. The site y may be visited again between time σ and time t but, since $y \notin B$ (recall that we assumed that B does not contain neighbouring points), the change of odometer at y has no impact on the evolution of the coupled ARWD process (which looks at the parity of the odometer only when we topple a site of B with one particle). All this shows that $\varphi(\bar{\tau}) \in \mathcal{I}_{m+\mathbb{1}_x+\mathbb{1}_y}^{\mathbf{z}}$.

Lastly, note that applying φ does not change the value of σ , i.e., $\sigma(\varphi(\bar{\tau})) = \sigma(\bar{\tau})$, and neither does it change the value of y or m_σ . Hence, the initial array $\bar{\tau}$ may be recovered from the knowledge of $\varphi(\bar{\tau})$ by simply removing the two instructions on x at index $m_\sigma(x) + 1$ and on y at index $m_\sigma(y) + 1$, and renumbering the other instructions. This shows that φ is injective, which concludes the proof of the claim (19), and therefore of the lemma. \square

6 The stabilisation time of ARWD: proof of Proposition 4.2

We now prove Proposition 4.2 which states that the stabilisation time of the ARWD model is exponentially large. The proof is an adaptation of the technique developed in [2] in the context of the activated random walk model. Yet, since there are some important differences and the proof is technical, we present it in details, stressing the differences with respect to [2].

6.1 Some notation and preliminaries

For convenience, in this section we use the following notion of distance on the torus, which does not correspond to the graph distance. Let $d, n \geq 1$ and $G = \mathbb{Z}_n^d$. Denoting by $\pi_n : \mathbb{Z}^d \rightarrow \mathbb{Z}_n^d$ the standard projection from \mathbb{Z}^d onto the torus, we define the distance between two points $x, y \in \mathbb{Z}_n^d$ as

$$d(x, y) = \inf \{ \|a - b\|_\infty : a, b \in \mathbb{Z}^d, \pi_n(a) = x, \pi_n(b) = y \}. \quad (20)$$

We consider this distance because we rely on some results of [2] which use this distance (the reason behind being that then the volume of balls has a simple expression), but things would work equally with the graph distance. For every non-empty set $C \subset \mathbb{Z}_n^d$, we define its diameter

$$\text{diam } C = \max_{x, y \in C} d(x, y).$$

If $r \in \mathbb{N}$, a set $C \subset \mathbb{Z}_n^d$ is said to be r -connected if, for any two points $x, y \in C$, there exists $k \in \mathbb{N}$ and a sequence $x_0, \dots, x_k \in C$ such that $x_0 = x$, $x_k = y$ and $d(x_j, x_{j+1}) \leq r$ every $j < k$.

For $\rho \in (0, 1)$ and $C \subset \mathbb{Z}_n^d$, we denote by $\mathcal{P}_{C, \rho}$ the Bernoulli site percolation measure on C with parameter ρ , that is to say, the probability measure on $\mathcal{P}(C)$ defined by $\mathcal{P}_{C, \rho}(U) = \rho^{|U|} (1 - \rho)^{|C \setminus U|}$ for each $U \subset C$.

Since in this section we are only concerned with the setting in which there is always one particle on each site of A , we choose to represent configurations with the subset $U \subset A$ of the unstable sites. We change accordingly the notation for the toppling strategy f and for the sleep mask g , writing $f : \cup_{t \geq 0} \mathcal{P}(A)^{t+1} \rightarrow V$ (in fact taking values in A) and $g : \mathbb{N}_0 \times \cup_{t \geq 0} \mathcal{P}(A)^{t+1} \rightarrow \mathcal{P}(V)$. Then, for $U_0 \subset A$, instead of $(\eta_t)_{t \geq 0}$ we denote by $(U_t)_{t \geq 0}$ the ARWD(λ, A, f, g, U_0) process, i.e., for every $t \geq 0$ we write $U_t = \{x \in A : \eta_t(x) = 1\}$.

Hitting probabilities on the torus

We need to introduce the following function, which measures the chance to wake up a distant site in a loop on the torus \mathbb{Z}_n^d , for $d \geq 1$:

$$\Upsilon_d : r \in \mathbb{N} \rightarrow \inf \{P_x(\tau_y < \tau_x^+), n \in \mathbb{N}, x, y \in \mathbb{Z}_n^d : d(x, y) \leq r\}.$$

We use the following estimates on this function (this lemma is identical to Lemma 2.5 of [2]):

Lemma 6.1. *We have the following estimates:*

- In dimension $d = 1$, we have $\Upsilon_1(r) = 1/(2r)$ for every $r \geq 1$;
- In dimension $d = 2$, there exists $K > 0$ such that $\Upsilon_2(r) \geq K/\ln r$ for every $r \geq 2$;
- In dimension $d \geq 3$, there exists $K = K(d) > 0$ such that $\Upsilon_d(r) \geq K$ for every $r \geq 1$.

6.2 The dormitory hierarchy

As explained in [2], given a settling set $A \subset \mathbb{Z}_n^d$, we introduce a hierarchical structure called the dormitory hierarchy which is deterministically associated to the set A and also depends on some parameters v and $(D_j)_{j \in \mathbb{N}_0}$.

Definition

A dormitory hierarchy is defined as follows, exactly as in [2]:

Definition 6.1. Let $d, v, n \geq 1$ and let $D = (D_j)_{j \in \mathbb{N}_0} \in \mathbb{N}^{\mathbb{N}_0}$. For every $A \subset \mathbb{Z}_n^d$ we call we call a (v, D) -dormitory hierarchy on A a non-increasing sequence of subsets $A \supset A_0 \supset \dots \supset A_J$, with $J \in \mathbb{N}_0$ and, for every $j \leq J$, a partition \mathcal{C}_j of A_j such that:

- (i) For $0 \leq j \leq J$, for every $C \in \mathcal{C}_j$, we have $|C| \geq 2^{\lfloor j/2 \rfloor} v$;
- (ii) For $0 \leq j \leq J - 1$, for every $C \in \mathcal{C}_{j+1} \setminus \mathcal{C}_j$, we have $\text{diam } C \leq D_j$ and there exists two sets $C_0, C_1 \in \mathcal{C}_j$ such that $C = C_0 \cup C_1$;
- (iii) The last partition \mathcal{C}_J contains one single set.

Given a dormitory hierarchy $(A_j, \mathcal{C}_j)_{j \leq J}$, for every $j \leq J$ and every $x \in A_j$, we define $\mathcal{C}_j(x)$ to be the set $C \in \mathcal{C}_j$ such that $x \in C$. When $x \in \mathbb{Z}_n^d \setminus A_j$ or $j > J$, we set $\mathcal{C}_j(x) = \emptyset$. The sets $C \in \mathcal{C}_j$ are called clusters at the level j . The parameter $v \in \mathbb{N}$ controls the volume of the clusters at each level of the hierarchy, while the sequence of diameters D_j ensures that we only merge clusters which are not too far apart.

Construction of the hierarchy for the two-dimensional case

We now build the dormitory hierarchy used for our proof in dimension 2. Compared with [2], we can afford to use a simpler hierarchy because we only aim to show that the critical density is less than one, whereas in [2] the parameters were optimized to obtain good bounds on the critical density. Note that in the simpler case of dimension at least 3 we use the trivial hierarchy with only one level.

Lemma 6.2. *Let $d = 2$, let $\mu \in (0, 1)$ and $v \in \mathbb{N}$. Let us write $r = 2 \lfloor \sqrt{2v/\mu} \rfloor$ and $D_j = 6^j \times 12vr$ for every $j \in \mathbb{N}_0$. For every $n \geq r + 1$ and every $A \subset \mathbb{Z}_n^2$ with $|A| \geq \mu n^2$, there exists $J \in \mathbb{N}_0$ and a (v, D) -dormitory hierarchy $(A_j, \mathcal{C}_j)_{j \leq J}$ on A , with $|A_J| \geq |A|/4$ and where every set $C \in \mathcal{C}_0$ is r -connected.*

Proof. Let d, μ, v, r, D_j, n and A be as in the statement. First, we simply let \mathcal{C}_0 be the set of the r -connected components of A of cardinality at least v , and $A_0 = \cup_{C \in \mathcal{C}_0} C$. Recall that we consider the distance coming from the infinite norm, as defined in (20). Thus, each square of side $r + 1$ can intersect at most

one r -connected component of A . Since each of the r -connected components of $A \setminus A_0$ contains at most $v - 1$ vertices, we have

$$|A \setminus A_0| \leq \left\lceil \frac{n}{r+1} \right\rceil^2 (v-1) \leq \left(\frac{n}{r+1} + 1 \right)^2 v \leq \frac{4n^2 v}{(r+1)^2} \leq \frac{\mu n^2}{2} \leq \frac{|A|}{2},$$

whence $|A_0| \geq |A|/2$. Then, we apply Lemma 5.1 of [2] to obtain the remainder of the hierarchy. The conditions required for this lemma are satisfied because $D_0/(12v) = r$ and

$$|A_0| \geq \frac{|A|}{2} \geq \frac{\mu n^2}{2} \geq \frac{2n^2}{r^2 v} = 8v \left(\frac{6n}{D_0} \right)^2.$$

This lemma completes (A_0, \mathcal{C}_0) into a (v, D) -dormitory hierarchy $(A_j, \mathcal{C}_j)_{j \leq J}$ on A , with the final set satisfying

$$|A_J| \geq |A_0| - 4v \left(\frac{6n}{D_0} \right)^2 \geq |A_0| - \frac{\mu n^2}{4} \geq \frac{|A|}{4},$$

as announced. \square

Distinguished vertices

Let $A \subset \mathbb{Z}_n^d$, and let $(A_j, \mathcal{C}_j)_{j \leq J}$ be a dormitory hierarchy on A . We define recursively a distinguished vertex in each set of the partitions. The distinguished point of a set C is written x_C^* and the particle sitting in x_C^* is called the distinguished particle of the cluster C . For every $C \in \mathcal{C}_0$, we simply set $x_C^* = \min C$, for an arbitrary order on the vertices of the torus. Then, for $1 \leq j \leq J$, if $C \in \mathcal{C}_j \setminus \mathcal{C}_{j-1}$, the property (ii) of the hierarchy tells us that C is the union of two clusters of \mathcal{C}_{j-1} . In this case, we let x_C^* be the distinguished vertex of the biggest of these two clusters (in terms of number of vertices, and with an arbitrary rule to break ties). We say that a vertex x is distinguished at the level $j \leq J$ if there exists $C \in \mathcal{C}_j$ such that $x = x_C^*$, that is to say, if $x = x_{\mathcal{C}_j(x)}^*$. If $j > J$, we say that no vertex is distinguished at level j . Note that if x is distinguished at a certain level j , then it is also distinguished at all levels j' for $j' < j$.

6.3 Construction of the toppling strategy and sleep mask

We now explain which toppling strategy and sleep mask we choose to consider. Let $A \subset \mathbb{Z}_n^d$, and let $(A_j, \mathcal{C}_j)_{j \leq J}$ be a fixed dormitory hierarchy on A .

At the first level of the hierarchy, we use a Markovian strategy, that is to say, a strategy which only looks at the current configuration. More precisely, we use the notion of C -toppling procedure, which was already used in [2]:

Definition 6.2. For every level $j \leq J$ and every cluster $C \in \mathcal{C}_j$, a C -toppling procedure is an application $p : \mathcal{P}(C) \setminus \{\emptyset\} \rightarrow C$ which is such that for every $U \in \mathcal{P}(C) \setminus \{\emptyset\}$, we have $p(U) \in U$ and, if $x_C^* \in U$ then $p(U) = x_C^*$.

In words, a C -toppling procedure is a rule to choose an active site of C in a configuration which is unstable on C , with the rule that the distinguished particle is chosen in priority, as soon as it is unstable.

We construct in Lemma 6.3 below a collection of C -toppling procedures p_C for each $C \in \mathcal{C}_0$. Let us assume for the moment that this collection of C -toppling procedures is constructed, and let us define recursively for every $j \leq J$ and every $C \in \mathcal{C}_j$ a toppling strategy f_C to stabilise the set C . Recall that we change the notation, denoting the configurations by subsets of A instead of elements of \mathbb{N}_s^V , which is made possible by the fact that we only consider configurations with one particle on each site of A .

First, if $C \in \mathcal{C}_0$, then for every $t \geq 0$ and $U_0, \dots, U_t \in \mathcal{P}(A)$ we simply define $f_C(U_0, \dots, U_t) = p_C(U_t \cap C)$ if $U_t \cap C \neq \emptyset$, and whatever site of U_t else if $U_t \neq \emptyset$, and whatever site of V otherwise if $U_t = \emptyset$.

Now, let $0 \leq j < J$, assume that we defined a toppling strategy f_C for each $C \in \mathcal{C}_j$, and let $C \in \mathcal{C}_{j+1}$. If C also belongs to \mathcal{C}_j then f_C is already constructed. Otherwise, if $C = C_0 \cup C_1$ with $C_0, C_1 \in \mathcal{C}_j$, then

for every $t \geq 0$ and every $U_0, \dots, U_t \in \mathcal{P}(A)$, for $k \in \{0, 1\}$ we denote by T_k the last instant before t when the cluster C_k was stable, that is to say,

$$T_k = \sup \{s \in \{0, \dots, t\} : U_s \cap C_k = \emptyset\},$$

with the convention that $T_k = -\infty$ if the set is empty, and we let

$$f_C(U_0, \dots, U_t) = \begin{cases} f_{C_0}(U_0, \dots, U_t) & \text{if } T_0 \leq T_1, \\ f_{C_1}(U_0, \dots, U_t) & \text{otherwise.} \end{cases}$$

The above definition encodes the so-called ping-pong rally used in [2]: we first stabilise C_0 using the strategy f_{C_0} , until C_0 is stable, and then stabilise C_1 using the strategy f_{C_1} . Then, if some sites of C_0 have been reactivated during the stabilisation of C_1 , we stabilise again C_0 , and then again C_1 if necessary, and so on until both sets are stabilised.

This recursive procedure yields a toppling strategy f_C for every cluster $C \in \mathcal{C}_j$ at every level $0 \leq j \leq J$. Then, the global toppling strategy that we consider is f_{A_J} , the strategy associated to the top-level cluster A_J with the above recursive construction. We simply denote it by f .

We now define the sleep mask associated with this strategy. For every $j \in \mathbb{N}_0$, every $t \geq 0$ and every $U_0, \dots, U_t \in \mathcal{P}(A)$, writing $x = f(U_0, \dots, U_t)$, we define

$$g(j, U_0, \dots, U_t) = \begin{cases} \emptyset & \text{if } x \text{ is distinguished at level } j+1 \text{ and } U_s \cap C_j(x) \neq \emptyset \text{ for all } s \leq t, \\ C_{j+1}(x) \setminus C_j(x) & \text{otherwise if } x \text{ is distinguished at level } j, \\ \emptyset & \text{if } x \text{ is distinguished at level } 0 \text{ but not at level } j, \\ C_0(x) & \text{if } x \text{ is not distinguished at any level.} \end{cases}$$

Compared to the definition of the sleep mask $w(x, j)$ in [2], we added the exception of the first case, which is meant to prevent the distinguished vertex of C_0 to wake up sites in C_1 during the first stabilisation of C_0 in the ping-pong rally, if C_0 and C_1 are two clusters which merge in the hierarchy. This will be useful for the proof of Lemma 6.5, to avoid a dependence issue which comes from the fact that we do not start from a fully active configuration.

Given this toppling strategy f and this sleep mask g , for every $j \leq J$ and every $C \in \mathcal{C}_j$, we denote the total number of steps (sleeps or loops) performed by the distinguished vertex of C during the stabilisation of C , starting with the configuration U_0 , by

$$H^*(C, U_0) = |\{t < T(C, U_0) : f(U_0, \dots, U_t) = x_C^*\}|,$$

where $(U_t)_{t \geq 0}$ is the ARWD(λ, A, f, g, U_0) process and $T(C, U_0)$ is the stabilisation time of C , as given by Definition 4.4.

6.4 Initialization: stabilisation of an elementary cluster

We now show that, for every cluster of the first level of the hierarchy constructed in Section 6.2, the number of steps made at the distinguished vertex is exponentially large in the size of this cluster.

Lemma 6.3. *Assume that $d = 2$ and let $\lambda > 0$, $\mu \in (0, 1)$ and $\rho_0 \in (0, 1]$. There exists $v_0 \in \mathbb{N}$ such that for every integer $v \geq v_0$ and every sequence $D \in \mathbb{N}^{\mathbb{N}_0}$, writing $r = 2\lfloor\sqrt{2v/\mu}\rfloor$, for every $n \geq 1$, if $A \subset \mathbb{Z}_n^2$ and $(A_j, C_j)_{j \leq J}$ is a (v, D) -dormitory hierarchy on A such that every set $C \in \mathcal{C}_0$ is r -connected, then there exists a collection of C -toppling procedures p_C for each $C \in \mathcal{C}_0$ such that, with the toppling strategy f and the ARWD sleep mask g defined in Section 6.3, for every $C \in \mathcal{C}_0$ and $\rho \in [\rho_0, 1]$, starting from a random initial configuration U_0 with distribution $\mathcal{P}_{C, \rho}$, the variable $1 + H^*(C, U_0)$ dominates a geometric variable with parameter $\exp(-|C|/\sqrt{v})$.*

This lemma is an analog of Lemma 3.4 of [2], and our proof follows the same strategy, but with some differences. In particular, we start with a random initial configuration instead of starting fully active, and after each loop, the particle falls asleep. Also, the stochastic domination concerns the total number of steps performed by the distinguished vertex, whereas in [2] we needed an estimate on the number of loops of color 0.

The metastability phenomenon

We first show the following result, which is the analog of Lemma 6.2 in [2]. The difference with respect to [2] lies in the fact that, after each loop, the particle which performed the loop falls asleep. This results in a slight change in the assumption (21) below, but is in fact harmless. For completeness, we include nevertheless the proof. As in [2], we use the convention that $\Upsilon_d(\infty) = \lim_{r \rightarrow \infty} \Upsilon_d(r)$ (we also use this lemma with $r = \infty$ to deal with the transient case).

Lemma 6.4. *Let $d \geq 1$, $\lambda > 0$, $\alpha > 0$, $\beta \in (0, 1)$, $v \in \mathbb{N}$, $r \in \mathbb{N} \cup \{\infty\}$, $D \in \mathbb{N}^{\mathbb{N}_0}$ and assume that*

$$(1 + \lambda)(1 - e^{-\alpha}) \leq \Upsilon_d(16vr)(1 - e^{-\alpha(1-\beta)v}). \quad (21)$$

Let $n \in \mathbb{N}$ and $A \subset \mathbb{Z}_n^d$, let $(A_j, \mathcal{C}_j)_{j \leq J}$ be a (v, D) -dormitory hierarchy on A such that every set $C \in \mathcal{C}_0$ is r -connected, then there exists a collection of C -toppling procedures p_C for each $C \in \mathcal{C}_0$ such that, with the toppling strategy f and the sleep mask g defined in Section 6.3, for every $C \in \mathcal{C}_0$ and every fixed subset $U_0 \subset C$ with $|U_0| = \lfloor \beta|C| \rfloor$, denoting by $(U_t)_{t \geq 0}$ the ARWD(λ, C, f, g, U_0) process and defining the stopping time $\tau = \inf\{t \geq 0 : |U_t| > \beta|C|\}$, it holds that

$$\mathbb{P}(\tau = \infty) \leq e^{-\alpha(\beta|C|-2)},$$

that is to say, with probability at least $1 - e^{-\alpha(\beta|C|-2)}$, the proportion of active sites in C exceeds β before hitting 0, i.e., before the stabilisation of C .

Proof. Let $d, \lambda, \alpha, \beta, v, r, D, n, A$ and $(A_j, \mathcal{C}_j)_{j \leq J}$ be as in the statement. For every $C \in \mathcal{C}_0$, Lemma 6.1 of [2] (applied with vr instead of r) tells us that there exists a function $\pi_C : \mathcal{P}(C) \setminus \{\emptyset\} \rightarrow C$ such that, for every $U \subset C$ with $0 < |U| \leq \beta|C|$, we have $\pi_C(U) \in U$ and

$$|(C \setminus U) \cap B(\pi_C(U), 16vr)| \geq (1 - \beta)v. \quad (22)$$

We turn these functions π_C into C -toppling procedures p_C by writing, for every $U \in \mathcal{P}(C) \setminus \{\emptyset\}$,

$$p_C(U) = \begin{cases} x_C^* & \text{if } x_C^* \in U, \\ \pi_C(U) & \text{otherwise.} \end{cases}$$

Then, we let f and g be the toppling strategy and sleep mask derived from these toppling procedures, as explained in Section 6.3.

We now fix a particular cluster $C \in \mathcal{C}_0$, we let $U_0 \subset C$ be such that $|U_0| = \lfloor \beta|C| \rfloor$, we denote by $(U_t)_{t \geq 0}$ the ARWD(λ, C, f, g, U_0) process and we let τ be the stopping time defined in the statement. Then, for every $t \geq 0$, we consider

$$M_t = e^{-\alpha N_t} \quad \text{where} \quad N_t = |U_t \setminus \{x_C^*\}|$$

and we show that the stopped process $(M_{t \wedge \tau})_{t \geq 0}$ is a supermartingale with respect to the natural filtration $(\mathcal{F}_t)_{t \geq 0}$ of $(U_t)_{t \geq 0}$. On the one hand, on the event $\{t \geq \tau\} \cup \{x_C^* \in U_t\}$, we have $M_{(t+1) \wedge \tau} = M_{t \wedge \tau}$, because our toppling procedure p_C chooses the distinguished particle in priority, and this particle is not allowed to wake up anyone in C , by definition of the sleep mask g .

On the other hand, on the event $\{t < \tau\} \cap \{x_C^* \notin U_t\}$, recalling that at each step the chosen particle falls asleep with probability $\lambda/(1 + \lambda)$ and makes a loop otherwise, falling asleep at the end of its loop, and noting that, on this event, our assumption (22) ensures that there are at least $k = \lceil (1 - \beta)v \rceil$ sleeping particles in the ball of radius $16vr$ around the chosen particle, we may write, on this event,

$$\begin{aligned} \mathbb{E}[M_{(t+1) \wedge \tau} | \mathcal{F}_t] &= \mathbb{E}[M_{t+1} | \mathcal{F}_t] = M_t \mathbb{E}\left[e^{-\alpha(N_{t+1} - N_t)} \mid \mathcal{F}_t\right] \\ &\leq M_t e^\alpha \left(\frac{\lambda}{1 + \lambda} + \frac{1}{1 + \lambda} \mathbb{E}\left[e^{-\alpha(Y_1 + \dots + Y_k)}\right] \right), \end{aligned} \quad (23)$$

where the variables Y_1, \dots, Y_k are Bernoulli with parameter $\Upsilon_d(16vr)$, but non necessarily independent. Then, using Lemma 6.3 of [2], we have

$$\mathbb{E}\left[e^{-\alpha(Y_1 + \dots + Y_k)}\right] \leq 1 - \Upsilon_d(16vr)(1 - e^{-\alpha k}) \leq 1 - \Upsilon_d(16vr)(1 - e^{-\alpha(1-\beta)v}).$$

Using our assumption (21), this becomes

$$\mathbb{E}[e^{-\alpha(Y_1+\dots+Y_k)}] \leq 1 - (1 + \lambda)(1 - e^{-\alpha}).$$

Plugging this into our computation (23), we deduce that, still on the event $\{t < \tau\} \cap \{x_C^* \notin U_t\}$,

$$\mathbb{E}[M_{(t+1)\wedge\tau} | \mathcal{F}_t] \leq M_t e^\alpha \left(\frac{\lambda}{1+\lambda} + \frac{1}{1+\lambda} - (1 - e^{-\alpha}) \right) = M_t = M_{t\wedge\tau},$$

which concludes the proof that $(M_{t\wedge\tau})_{t \geq 0}$ is a supermartingale.

Hence, using Doob's optional stopping theorem, we deduce that

$$\mathbb{P}(\tau = \infty) = \mathbb{P}(T(C, U_0) < \tau) \leq \mathbb{E}[M_{\tau \wedge T(C, U_0)}] \leq M_0 \leq e^{-\alpha(\beta|C|^{-2})},$$

which concludes the proof of the lemma. \square

Consequence on the number of steps performed by the distinguished particle

We are now ready to show Lemma 6.3.

Proof of Lemma 6.3. Let $d = 2$, $\lambda > 0$, $\mu \in (0, 1)$ and $\rho_0 \in (0, 1]$. Let $\beta = \rho_0/2$ and let K be the constant given by Lemma 6.1. We claim that, for v large enough, we have

$$(1 + \lambda) \left[1 - \exp\left(-\frac{4}{\beta\sqrt{v}}\right) \right] \leq \frac{K}{\ln(64v^{3/2}/\sqrt{\mu})} \left[1 - \exp\left(-\frac{4(1-\beta)\sqrt{v}}{\beta}\right) \right]. \quad (24)$$

Indeed, when $v \rightarrow \infty$ the left-hand side of (24) is equivalent to $4(1 + \lambda)/(\beta\sqrt{v})$, while the right-hand side is equivalent to $2K/(3 \ln v)$. Hence, we can fix $v_1 \in \mathbb{N}$ such that (24) holds for all $v \geq v_1$.

Let also $v_2 \in \mathbb{N}$ be such that, for every $v \geq v_2$,

$$8v^2 e^{-\sqrt{v}} \leq 1 \quad \text{and} \quad \left(1 - \frac{K}{(1 + \lambda) \ln(2v^{3/2}\sqrt{2/\mu})} \right)^{v^2} \leq \frac{e^{-v}}{4}. \quad (25)$$

We then let $v_0 = v_1 \vee v_2 \vee \lceil 1/\beta^4 \rceil$, we take $v \geq v_0$, $D \in \mathbb{N}^{\mathbb{N}_0}$ and $r = 2 \lfloor \sqrt{2v/\mu} \rfloor$ and we let $n \geq 1$ and we consider a set $A \subset \mathbb{Z}_n^2$ equipped with a (v, D) -dormitory hierarchy $(A_j, \mathcal{C}_j)_{j \leq J}$ such that every set $C \in \mathcal{C}_0$ is r -connected.

Then, the inequality (24) allows us to use Lemma 6.4 with $\alpha = 4/(\beta\sqrt{v})$. Let $(p_C)_{C \in \mathcal{C}_0}$ be the collection of toppling procedures given by this lemma, and let f and g be the toppling strategy and sleep mask derived from this collection, as detailed in Section 6.3.

Let $C \in \mathcal{C}_0$. As in [2] we consider the set

$$\mathcal{B} = \{U \subset C : |U| > \beta|C|\}$$

and its so-called boundary, defined as

$$\partial\mathcal{B} = \left\{ U \subset C \setminus \{x_C^*\} : |U| = \lfloor \beta|C| \rfloor \right\}.$$

Let us fix a deterministic initial configuration $U_0 \in \partial\mathcal{B}$, and denote by $(U_t)_{t \geq 0}$ the ARWD(λ, C, f, g, U_0) process. By Lemma 6.4, starting from this configuration U_0 , the probability that the process stabilises without ever visiting \mathcal{B} is $\mathbb{P}(\tau = \infty) \leq e^{-\alpha(\beta|C|^{-2})}$. Using that $\beta|C| \geq \beta v \geq 1/\beta^3 \geq 8 \geq 4$, we may deduce that

$$\mathbb{P}(\tau = \infty) \leq e^{-\alpha\beta|C|/2} \leq e^{-2|C|/\sqrt{v}}, \quad (26)$$

Furthermore, note that, since a C -toppling procedure only depends on the current configuration, the process $(U_t)_{t \geq 0}$ is in fact a Markov chain. Thus, since (26) holds for all $U_0 \in \partial\mathcal{B}$, denoting by $\mathcal{N}_{\mathcal{B}}$ the number of times that the process enters \mathcal{B} (from a configuration outside this set) before the complete stabilisation of C , the variable $1 + \mathcal{N}_{\mathcal{B}}$ dominates a geometric variable with parameter $e^{-2|C|/\sqrt{v}}$.

Still starting from a fixed $U_0 \in \partial\mathcal{B}$, let us introduce two random times: first,

$$T_\star = \inf \{t \in \mathbb{N}_0 : x_C^\star \in U_t\}$$

recalling that $x_C^\star \notin U_0$ by definition of $\partial\mathcal{B}$ and, for $t \in \mathbb{N}_0$,

$$T_\mathcal{B}^t = \inf \{t' \geq t : U_{t'} \in \mathcal{B}\}.$$

We also define a deterministic time $t_0 = |C|^2$, and we consider the event

$$\mathcal{E} = \{T_\star \leq t_0\} \cap \{T_\mathcal{B}^{T_\star} < \infty\}.$$

We then control the probability of this event \mathcal{E} as in Section 6.3 of [2], obtaining

$$\begin{aligned} \mathbb{P}(\mathcal{E}^c) &\leq 2t_0 e^{-2|C|/\sqrt{v}} + \left(1 - \frac{\Upsilon_2(r|C|)}{1+\lambda}\right)^{t_0} \\ &\leq 2|C|^2 e^{-2|C|/\sqrt{v}} + \left(1 - \frac{K}{(1+\lambda) \ln(2|C|\sqrt{2v/\mu})}\right)^{|C|^2} \\ &\leq 2|C|^2 e^{-\sqrt{|C|}} e^{-|C|/\sqrt{v}} + \left(1 - \frac{K}{(1+\lambda) \ln(2|C|^{3/2}\sqrt{2/\mu})}\right)^{|C|^2}, \end{aligned}$$

where in the last inequality we used that $|C| \geq v$. Then, since $|C| \geq v \geq v_2$ we may use the two inequalities in (25) to get

$$\mathbb{P}(\mathcal{E}^c) \leq \frac{e^{-|C|/\sqrt{v}}}{4} + \frac{e^{-|C|}}{4} \leq \frac{e^{-|C|/\sqrt{v}}}{2}.$$

If this event \mathcal{E} is realized, it means that, starting from $U_0 \in \partial\mathcal{B}$, the distinguished particle is awakened within time t_0 , after which we reach the set \mathcal{B} . As explained above, when we exit \mathcal{B} , we necessarily end up on a configuration in $\partial\mathcal{B}$.

Thus, since the bound on the probability of \mathcal{E}^c is uniform over all $U_0 \in \partial\mathcal{B}$, we deduce that, starting from any configuration in $\partial\mathcal{B}$, one plus the number of times that the distinguished vertex is waken up dominates a geometric variable with parameter $(e^{-|C|/\sqrt{v}})/2$.

Since each time that the distinguished vertex is waken up, it is toppled at the next step, we deduce that, still starting from any fixed configuration $U_0 \in \partial\mathcal{B}$, the variable $1 + H^\star(C, U_0)$ dominates a geometric variable with parameter $(e^{-|C|/\sqrt{v}})/2$.

There now remains to go back to the case of a random initial configuration U_0 with distribution $\mathcal{P}_{C,\rho}$, that is to say, with i.i.d. particles on C , each unstable with probability ρ , where $\rho \geq \rho_0$. In this setting, Hoeffding's inequality entails that

$$\mathcal{P}_{C,\rho}(U_0 \notin \mathcal{B}) = \mathcal{P}_{C,\rho}\left(|U_0| \leq \frac{\rho}{2}|C|\right) \leq e^{-\rho^2|C|/2} \leq e^{-2|C|/\sqrt{v}} \leq \frac{e^{-|C|/\sqrt{v}}}{2},$$

where in the second inequality we used our assumption that $v \geq 1/\beta^4 \geq 16/\rho^4$ and in the last inequality we used that $\sqrt{v} \geq \ln 2$.

Now recall that if $U_0 \in \mathcal{B}$, then the process will visit at some time the boundary set $\partial\mathcal{B}$. Thus, what we showed above implies that, conditioned on the event that $U_0 \in \mathcal{B}$, the variable $1 + H^\star(C, U_0)$ dominates a geometric variable with parameter $(e^{-|C|/\sqrt{v}})/2$.

Therefore, without this conditioning, writing

$$H^\star(C, U_0) \geq \mathbb{1}_{\{U_0 \in \mathcal{B}\}} H^\star(C, U_0)$$

and using that if X is a Bernoulli variable with parameter $1-a$ and G is a geometric variable with parameter b , independent of X , then $1 + X(G-1)$ dominates a geometric variable with parameter $a + b - ab \leq a + b$, we conclude that the variable $1 + H^\star(C, U_0)$ dominates a geometric with parameter

$$\mathcal{P}_{C,\rho}(U_0 \notin \mathcal{B}) + \frac{e^{-|C|/\sqrt{v}}}{2} \leq e^{-|C|/\sqrt{v}},$$

which concludes the proof. \square

6.5 The induction step: two clusters playing ping-pong

Recall that for every cluster C of our hierarchy we denote by $H^*(C, U_0)$ the number of steps performed by the distinguished vertex x_C^* during the stabilisation of C starting from an initial configuration $U_0 \subset C$.

The induction hypothesis $\mathcal{P}(j)$: Let $\rho_0 \in (0, 1]$ and let $(\alpha_j)_{j \in \mathbb{N}_0}$ be a sequence of positive real numbers, to be chosen later. Our induction hypothesis, written $\mathcal{P}(j)$, is the following: for every $j \in \{0, \dots, J\}$, we define

$$\mathcal{P}(j) : \quad \forall C \in \mathcal{C}_j \quad \forall \rho \in [\rho_0, 1] \quad U_0 \sim \mathcal{P}_{C, \rho} \implies 1 + H^*(C, U_0) \succeq \text{Geom}(e^{-\alpha_j |C|}). \quad (27)$$

Note that, the ARWD model being a priori non-abelian, we have no monotonicity property with respect to the initial distribution at our disposal. This is why the induction hypothesis is stated for every $\rho \geq \rho_0$ instead of just taking U_0 with distribution \mathcal{P}_{C, ρ_0} . In practice, at each level of the hierarchy, we will only use the cases $\rho = \rho_0$ (during the first ping-pong round) and $\rho = 1$ (for the successive ping-pong rallies).

Note that Lemma 6.3 showed $\mathcal{P}(0)$, under certain hypotheses.

Here is now the induction step of our proof, which is adapted from Lemma 3.1 of [2]. The constant \bar{p} is a universal constant which comes from an elementary property of geometric sums of Bernoulli variables, given by Lemma 6.6 below.

Lemma 6.5. *Let $\bar{p} \in (0, 1)$ be given by Lemma 6.6. Let $d \geq 1$, $\rho_0 \in (0, 1]$, $v \geq 1$, $(D_j)_{j \geq 0} \in \mathbb{N}^{\mathbb{N}_0}$ and let $(\alpha_j)_{j \geq 0}$ be a non-increasing sequence of positive real numbers and $(p_j)_{j \geq 0} \in (0, 1)^{\mathbb{N}_0}$ be such that for every $j \geq 0$,*

$$p_j \leq \bar{p} \wedge \frac{1}{2^{j+1}(1+\lambda)} \wedge (32\alpha_j(D_j + 1)^d)^{-1/2} \quad (28)$$

and

$$2^{j/2+2}v \leq p_j^6 \Upsilon_d(D_j) \exp((\alpha_j - \alpha_{j+1})2^{j/2}v). \quad (29)$$

For every $n \geq 1$ and every $A \subset \mathbb{Z}_n^d$ equipped with a (v, D) -dormitory hierarchy $(A_j, \mathcal{C}_j)_{j \leq J}$ and with a collection of toppling procedures $(p_C)_{C \in \mathcal{C}_0}$, with the toppling strategy f and sleep mask g defined in Section 6.3, if the property $\mathcal{P}(0)$ holds, then $\mathcal{P}(j)$ also holds for every $j \leq J$.

The proof being similar to the proof of Lemma 3.1 in [2], but with some important changes, we concentrate on the novelties in our setting and refer to [2] for more detailed explanations.

General setting for the proof of Lemma 6.5

Let $d, \rho_0, v, (D_j)_{j \geq 0}, (\alpha_j)_{j \geq 0}, (p_j)_{j \geq 0}, n, A, (A_j, \mathcal{C}_j)_{j \leq J}, (p_C)_{C \in \mathcal{C}_0}, f$ and g be as in the statement. We assume that $0 \leq j < J$ is such that the property $\mathcal{P}(j)$ holds and we wish to establish the property $\mathcal{P}(j+1)$, i.e., that for every $C \in \mathcal{C}_{j+1}$ and $\rho \in [\rho_0, 1]$, if $U_0 \sim \mathcal{P}_{C, \rho}$ then

$$1 + H^*(C, U_0) \succeq \text{Geom}(e^{-\alpha_{j+1}|C|}). \quad (30)$$

We start by distinguishing between two cases. First, if $C \in \mathcal{C}_{j+1} \cap \mathcal{C}_j$, then (30) directly follows from induction hypothesis $\mathcal{P}(j)$, using our assumption that $\alpha_{j+1} \leq \alpha_j$.

Suppose now that $C \in \mathcal{C}_{j+1} \setminus \mathcal{C}_j$. Following property (ii) of the dormitory hierarchy, we have $\text{diam } C \leq D_j$ and we can write $C = C_0 \cup C_1$ with $C_0, C_1 \in \mathcal{C}_j$. To simplify the notation, we write $x_0^* = x_{C_0}^*$ and $x_1^* = x_{C_1}^*$. Exchanging the names C_0 and C_1 if necessary, we may assume that $|C_0| \geq |C_1|$ and $x_C^* = x_0^*$.

To stabilise C we perform a ping-pong rally, namely, we first stabilise C_0 then stabilise C_1 . Then, some sites of C_0 may have been reactivated during the stabilisation of C_1 , so we stabilise again C_0 , and then again C_1 if necessary, going on until both sets are stabilised.

As explained before, the dynamics contains the following exception, which is encoded into the sleep mask g defined in Section 6.3: the loops emitted during the first stabilisation of C_0 are not allowed to wake up the sites of C_1 , so that the first stabilisation of C_1 starts from the initial configuration inside C_1 .

As in [2], denoting by U_i the set of unstable sites of C before the stabilisation number i (starting the numbering at 0, with U_0 the random configuration defined above), we define

$$\mathcal{N} = \inf \{ i \geq 1 : C_0 \not\subset U_{2i} \text{ or } C_1 \not\subset U_{2i+1} \},$$

which indicates the first stabilisation which does not wake up all the sites of the other cluster (with the exception that we do not require anything on the first stabilisation of C_0 , which in any case is not allowed to wake up anyone in C_1). Our notation U_i corresponds to the notation R_i in [2].

Key differences with respect to [2]

We now explain the main differences with respect to the proof of Lemma 3.1 in [2].

The first difference is that, while in [2] we were interested in the stabilisation time starting with all the particles active (i.e., $U_0 = C$), here we start instead with an i.i.d. initial configuration, where each site of A is unstable with probability ρ . This difference comes from the fact that, in ARW, we may let the particles walk until each site of A contains one particle, and then obtain the configuration with one active particle sitting on each site of A , whereas in ARWD, due to the fact that a particle falls asleep at the end of its loop, the configuration reached once all the sites of A are occupied is not necessarily fully active.

Another difference is that, during the first stabilisation of C_0 , the configuration in C_1 is frozen with no reactivation possible, due to the exception inserted in the definition of the sleep mask g (see Section 6.3). Thus, in our setting the first stabilisation of C_0 and the first stabilisation of C_1 both start with an i.i.d. configuration with parameter ρ . Then, as in [2], each of the subsequent stabilisation rounds of C_0 or C_1 , until the \mathcal{N} -th round, starts with the corresponding set fully active.

As a consequence, the variable \mathcal{N} is no longer exactly geometric. Nevertheless, this difference is harmless. Indeed, we will show that \mathcal{N} dominates a geometric variable, using that at each stabilisation round we can use the induction hypothesis (with initial fraction of active particles ρ for the first round and 1 for the successive rounds).

Another crucial difference is that, in our setting, we do not have an equivalent of Lemma 2.3 of [2]. This lemma states that, in the ARW case, knowing the total number \mathcal{T} of sleeps and loops of colour $0, \dots, j-1$ produced by the distinguished site of a cluster $C \in \mathcal{C}_j$ during the stabilisation of this cluster (starting from all the sites active), then the number of loops of colour j produced by x_C^* is distributed as a sum of \mathcal{T} geometric variables (minus one) with an explicit parameter, independent of \mathcal{T} . This comes from the key observation that, during the stabilisation of a cluster $C \in \mathcal{C}_j$, the loops of colour at least j have no effect on \mathcal{T} , because these loops are not allowed to wake up anyone inside C . Thus, in ARW, to determine \mathcal{T} one does not care about how many loops of colour j are inserted between any two loops of colour less than j .

In [2], the induction hypothesis was formulated in terms of the number of loops of colour j produced by the distinguished site of a cluster $C \in \mathcal{C}_j$. Thus, we used this lemma to go from one colour to the next one, translating an information about the loops of colour j into an information about the loops of colour $j+1$.

Now recall that ARWD differs from ARW in that, after each loop, the particle is forced to fall asleep. Thus, all the loops have an effect on the stabilisation of $C \in \mathcal{C}_j$ because, even if a loop is not allowed to wake up anyone, it has the effect of making the distinguished particle fall asleep. This dependency makes it more difficult to translate a statement on the number of loops of colour j into a statement of colour $j+1$.

To circumvent this, our induction hypothesis is simply formulated with the total number of steps and, at each step, we translate the lower bound on the total number of steps into a lower bound on the loops of colour j . Knowing that $H^*(C, U_0) = h$, the number of loops of colour j is a sum of h independent Bernoulli variables with parameter $1/((1+\lambda)2^{j+1})$, but these Bernoulli variables are not independent of $H^*(C, U_0)$. Yet, loosing on the parameter of the geometric variables, we are able to obtain a lower bound, using Lemma 6.6 below, whose elementary proof is deferred to Section 6.6.

Lemma 6.6. *There exists $\bar{p} \in (0, 1)$ such that, for every $p \in (0, \bar{p})$, for every ε such that*

$$\exp\left(-\frac{1}{32p^2}\right) \leq \varepsilon < p^3, \tag{31}$$

if $T+1$ is a geometric variable with parameter ε and $(X_n)_{n \geq 1}$ are i.i.d. Bernoulli variables with parameter p , not necessarily independent of T , then the variable $1 + X_1 + \dots + X_T$ stochastically dominates a geometric variable with parameter ε/p^3 .

Then, to obtain the induction hypothesis at rank $j + 1$ we simply bound from below the total number of steps by the number of loops of colour j .

The loop representation

As explained above, we want to study the number of loops of colour j produced by the distinguished vertex during the successive rounds of the ping-pong rally. Yet, as ARWD is defined, the number of loops of colour j cannot be expressed as a function of the ARWD process. Thus, we use the loop representation presented in [2], adapting it to our setting, which enables us to study these coloured loops.

This representation is based on independent random variables

$$(I(x, h))_{x \in A, h \in \mathbb{N}_0} \in \{0, 1\}^{A \times \mathbb{N}_0}, \quad (J(x, \ell))_{x \in A, \ell \in \mathbb{N}} \in \mathbb{N}^{A \times \mathbb{N}_0} \quad \text{and} \quad (R(x, \ell, j))_{x \in A, \ell \in \mathbb{N}_0, j \in \mathbb{N}_0} \in \mathcal{P}(A)^{A \times \mathbb{N}_0^2},$$

where the variables $I(x, h)$ are Bernoulli with parameter $\lambda/(1 + \lambda)$, the variables $1 + J(x, \ell)$ are geometric with parameter $1/2$ and $R(x, \ell, j)$ is distributed as the support of a symmetric random walk on the torus started and killed at x . Probabilities and expectations are simply denoted by \mathbb{P} and \mathbb{E} , which depend implicitly on the set A .

We now describe the update rules of the model. Recall that a configuration is a subset $U \subset A$ indicating which sites are active. To update the configuration we need to recall the numbers of visits and loops already used at each vertex. This is the role of what we call the odometer function $m : A \rightarrow \mathbb{N}_0$ and the loop odometer function $\ell : A \rightarrow \mathbb{N}_0$.

Let $k \in \{0, 1\}$ and let us describe the procedure to stabilise the cluster C_k using these variables. Let f_{C_k} be the toppling strategy associated with this cluster C_k and let g be the sleep mask, both defined in Section 6.3.

Given a fixed configuration $U_0 \subset A$, an odometer $m_0 : A \rightarrow \mathbb{N}_0$ and a loop odometer $\ell_0 : A \rightarrow \mathbb{N}_0$, we define a process $(U_t, m_t, \ell_t)_{t \geq 0}$ as follows. Let $t \geq 0$ and assume that (U_s, m_s, ℓ_s) is constructed for every $s \leq t$. If $U_t \cap C_k = \emptyset$ then we define $(U_{t+1}, m_{t+1}, \ell_{t+1}) = (U_t, m_t, \ell_t)$. Otherwise, if $U_t \cap C_k \neq \emptyset$ then, writing $x = f_{C_k}(U_0, \dots, U_t)$, $i = I(x, m_t(x))$, $j = J(x, \ell_t(x))$, $R = R(x, \ell_t(x), j)$ and $W = g(j, U_0, \dots, U_t)$, we define the new odometer $m_{t+1} = m_t + \mathbb{1}_x$, the new loop odometer $\ell_{t+1} = \ell_t + \mathbb{1}_{\{i=0\}} \mathbb{1}_x$, and the next configuration

$$U_{t+1} = \begin{cases} U_t \setminus \{x\} & \text{if } i = 1, \\ (U_t \cup (R \cap W)) \setminus \{x\} & \text{otherwise.} \end{cases}$$

Then, we define the stabilisation operator

$$\text{Stab}_k : \begin{cases} \mathcal{P}(A) \times (\mathbb{N}_0^A)^2 & \longrightarrow \mathcal{P}(A) \times (\mathbb{N}_0^A)^2 \\ (U_0, m_0, \ell_0) & \longmapsto (U_\tau, m_\tau, \ell_\tau), \end{cases}$$

where

$$\tau = \inf\{t \in \mathbb{N}_0 : U_t \cap C = \emptyset\},$$

which is almost surely finite (we restrict ourselves to the event on which it is always finite). This notation allows us to define, for every $k \in \mathbb{N}_0$, the number of loops of colour k produced by the distinguished vertex of k during the stabilisation of C , which is

$$\mathcal{L}(C, U_0, k) = |\{\ell < \ell_\tau(x_C^*) : J(x_C^*, \ell) = k\}|.$$

Then, with this notation, it turns out that $m_\tau(x_C^*)$ is distributed as $H^*(C, U_0)$.

Translating the induction hypothesis

Going back to our cluster $C = C_0 \cup C_1$ with $C_0, C_1 \in \mathcal{C}_j$, our induction hypothesis $\mathcal{P}(j)$ tells us that for $k \in \{0, 1\}$, for any random subset $U_0 \sim \mathcal{P}_{C_k, \rho}$ with $\rho \geq \rho_0$, the variable $1 + H^*(C_k, U_0)$ dominates a geometric random variable with parameter $\exp(-\alpha_j |C_k|)$.

We wish to translate this information in terms of the number of loops of colour j . Recall that each step made by the distinguished particle is a loop of colour j with probability $1/((1 + \lambda)2^{j+1})$. Hence, we may write

$$\mathcal{L}(C_k, U_0, j) = \sum_{i=1}^{H^*(C_k, U_0)} X_i \succeq \sum_{i=1}^{H^*(C_k, U_0)} X'_i,$$

where $(X_i)_{i \geq 1}$ are i.i.d. Bernoulli random variables with parameter $1/((1 + \lambda)2^{j+1})$ and $(X'_i)_{i \geq 1}$ are i.i.d. Bernoulli random variables with parameter p_j (recall that $p_j \leq 1/((1 + \lambda)2^{j+1})$), but not necessarily independent of $H^*(C_k, U_0)$.

We now check that the conditions required to apply Lemma 6.6 are satisfied. On the one hand, since C_k is included into C which has diameter at most D_j , we have $|C_k| \leq (D_j + 1)^d$. Thus, using our assumption (28), we have

$$\alpha_j |C_k| \leq \alpha_j (D_j + 1)^d \leq \frac{1}{32p_j^2}. \quad (32)$$

On the other hand, it follows from our assumption (29) that

$$\exp(-\alpha_j |C_k|) \leq \exp(-\alpha_j 2^{j/2} v) \leq \frac{p_j^6}{4} \leq \frac{p_j^3}{2}. \quad (33)$$

These two bounds (32) and (33) allow us to apply Lemma 6.6 with p_j and $\varepsilon = e^{-\alpha_j |C_k|}$ to deduce that, for an initial configuration $U_0 \sim \mathcal{P}_{C_k, \rho}$, the variable $1 + \mathcal{L}(C_k, U_0, j)$ dominates a geometric variable with parameter

$$q_k = \frac{1}{p_j^3} \exp(-\alpha_j |C_k|). \quad (34)$$

The ping-pong rally and the infinite Sisyphus sequence

For every $i \in \mathbb{N}_0$, we write $\varepsilon(i) = 0$ if i is even and $\varepsilon(i) = 1$ if i is odd. Then, starting with U_0 distributed as $\mathcal{P}_{C, \rho}$ and $m_0 = \ell_0 = 0$, we define, for every $i \geq 0$,

$$(U_{i+1}, m_{i+1}, \ell_{i+1}) = \text{Stab}_{\varepsilon(i)}(U_i, m_i, \ell_i).$$

For every $i \geq 0$, we denote by \mathcal{L}_i the number of loops of colour j performed by the distinguished site of $C_{\varepsilon(i)}$ during the i -th stabilisation (the first stabilisation of C_0 having number 0).

Besides, we also consider another sequence which is built by instead assuming that each stabilisation (except the first stabilisation of C_0) wakes up all the sites of the other cluster. Namely, for $i \leq 2$ we let $(U'_i, h'_i, \ell'_i) = (U_i, m_i, \ell_i)$, and for $i \geq 2$ we let

$$(U'_{i+1}, h'_{i+1}, \ell'_{i+1}) = \text{Stab}_{\varepsilon(i)}(C_{\varepsilon(i)}, m_i, \ell_i).$$

We define \mathcal{L}'_i similarly with this new sequence and, as in [2] we note that $\mathcal{L}'_i = \mathcal{L}_i$ for every $i < 2\mathcal{N}$, allowing us to write

$$H^*(C, U_0) \geq \mathcal{L}(C, U_0, j) \geq \sum_{i=0}^{\mathcal{N}-1} \mathcal{L}_{2i} = \sum_{i=0}^{\mathcal{N}-1} \mathcal{L}'_{2i}.$$

Note that the sequence $(\mathcal{L}'_i)_{i \geq 0}$ is independent, with \mathcal{L}'_0 distributed as $\mathcal{L}(C_0, U_0, j)$, \mathcal{L}'_1 distributed as $\mathcal{L}(C_1, U_0, j)$ and, for every $i \geq 2$, \mathcal{L}'_i distributed as $\mathcal{L}(C_{\varepsilon(i)}, C, j)$. Thus, as explained above, for every $i \geq 0$, the variable $1 + \mathcal{L}'_i$ dominates a geometric variable with parameter $q_{\varepsilon(i)}$ given by (34). And, as in [2], we have the key property that this sequence is independent of $(R(x_k^*, \ell, j))_{k \in \{0,1\}, \ell \geq 0}$, i.e., the sets visited by the loops of colour j emitted by the two distinguished vertices. Note that this independence property would not hold if we had not introduced the exception that during the first stabilisation of C_0 , the distinguished particle of C_0 is not allowed to wake up anyone in C_1 : because then the first stabilisation of C_1 would start with a configuration which depends on the first stabilisation of C_0 . Note also that we use here the assumption that the initial configuration U_0 is i.i.d., so that $U_0 \cap C_0$ and $U_0 \cap C_1$ are independent.

Resulting stochastic domination

Using this independence property and proceeding as in [2], we have a positive correlation inequality which allows us to write, for every $m \geq 1$,

$$\mathbb{P}(H^*(C, U_0) \geq m) \geq \sum_{k=1}^{+\infty} \mathbb{P}\left(\left\{\sum_{i=0}^{k-1} \mathcal{L}'_{2i} \geq m\right\} \cap \{\mathcal{N} = k\}\right) \geq \sum_{k=1}^{+\infty} \mathbb{P}\left(\sum_{i=0}^{k-1} \mathcal{L}'_{2i} \geq m\right) \mathbb{P}(\mathcal{N} = k).$$

Thus, taking \mathcal{N}' a copy of \mathcal{N} which is independent of \mathcal{L}' , we get the stochastic domination

$$H^*(C, U_0) \succeq \sum_{i=0}^{\mathcal{N}'-1} \mathcal{L}'_{2i}.$$

Then, as explained above, this variable \mathcal{N}' is not exactly a geometric variable, but it dominates a geometric variable with parameter $1 - \pi_0\pi_1$, with

$$\pi_0 = \mathbb{E}[\psi_0(\mathcal{L}'_0)] \wedge \mathbb{E}[\psi_0(\mathcal{L}'_2)] \quad \text{and} \quad \pi_1 = \mathbb{E}[\psi_1(\mathcal{L}'_1)] \wedge \mathbb{E}[\psi_1(\mathcal{L}'_3)],$$

where the function ψ_0 is defined by

$$\psi_0 : x \in \mathbb{N} \mapsto \mathbb{E}\left(C_1 \subset \bigcup_{0 \leq \ell < x} R(x_0^*, \ell, j)\right)$$

and ψ_1 is defined similarly, replacing C_1 with C_0 and x_0^* with x_1^* . Now, we use the following elementary result (Lemma 2.4 in [2]):

Lemma 6.7. *Let N be a geometric random variable with parameter $a \in (0, 1)$, and let $(X_n)_{n \in \mathbb{N}}$ be i.i.d. geometric variables with parameter $b \in (0, 1)$, independent of N . Then, the variable*

$$S = 1 + \sum_{n=1}^N (X_n - 1) \quad \text{is geometric with parameter} \quad \frac{ab}{1 - b + ab}.$$

Using this, we deduce that $1 + H^*(C, U_0)$ dominates a geometric variable with parameter

$$q' = \frac{(1 - \pi_0\pi_1)q_0}{1 - \pi_0\pi_1q_0}.$$

Then, the same computations as in the end of Section 7 of [2] yield

$$q' \leq \frac{|C|}{(1 - q_0)\Upsilon_d(D_j)} q_0 q_1 = \frac{|C| \exp(-\alpha_j |C|)}{(1 - q_0)\Upsilon_d(D_j) p_j^6} \leq \frac{2|C| \exp(-\alpha_j |C|)}{\Upsilon_d(D_j) p_j^6}, \quad (35)$$

using the expression (34) of q_0 and q_1 and then the crude bound $q_0 \leq 1/2$ which follows from (33). Then, using our assumption (29), we have

$$1 \leq (\alpha_j - \alpha_{j+1})2^{j/2}v \leq (\alpha_j - \alpha_{j+1})|C|,$$

which, using that $x \mapsto xe^{-x}$ decreases on $[1, \infty)$, implies that

$$|C| \exp(-(\alpha_j - \alpha_{j+1})|C|) \leq 2^{j/2}v \exp(-(\alpha_j - \alpha_{j+1})2^{j/2}v).$$

Plugging this into (35), we obtain

$$q' \leq \frac{2^{j/2+1}v \exp(-(\alpha_j - \alpha_{j+1})2^{j/2}v)}{\Upsilon_d(D_j) p_j^6} \exp(-\alpha_{j+1}|C|) \leq \exp(-\alpha_{j+1}|C|)$$

where in the last inequality we used again our assumption (29). This concludes the proof of Lemma 6.5. \square

6.6 Proof of Lemma 6.6

Proof of Lemma 6.6. Let us consider the function

$$f : p \in (0, 1) \mapsto \ln p - \frac{\ln(1 - p + p^2)}{8p^3}.$$

When $p \rightarrow 0$, we have $f(p) \sim 1/(8p^2)$. Hence, we can find $\bar{p} \in (0, 1)$ such that $f(p) \geq 1/(16p^2)$ for every $p \in (0, \bar{p})$. Upon decreasing \bar{p} if necessary, we can assume that $\bar{p} \leq 1/8$.

Let now $p \in (0, \bar{p})$ and ε satisfying (31), and let T and $(X_n)_{n \geq 1}$ be as in the statement. For every $n \geq 1$, let us write $S_n = X_1 + \dots + X_n$. To prove the result, we have to show that

$$\forall k \geq 1, \quad \mathbb{P}(S_T \geq k) \geq \left(1 - \frac{\varepsilon}{p^3}\right)^k. \quad (36)$$

Let $k \geq 1$. Defining $n = \lceil k/(8p^3) \rceil$, we simply write

$$\mathbb{P}(S_T \geq k) \geq \mathbb{P}(\{T \geq n\} \cap \{S_n \geq k\}) \geq \mathbb{P}(T \geq n) - \mathbb{P}(S_n < k). \quad (37)$$

We then bound separately these two terms.

On the one hand, we have

$$\mathbb{P}(T \geq n) = (1 - \varepsilon)^n = \exp\left(\left\lceil \frac{k}{8p^3} \right\rceil \ln(1 - \varepsilon)\right).$$

First, note that

$$\left\lceil \frac{k}{8p^3} \right\rceil \leq \frac{k}{8p^3} + 1 \leq \frac{k}{4p^3},$$

using that $p \leq 1/2$. Besides, since $\varepsilon \leq p^3 \leq 1/2$, we have, by concavity of the logarithm function,

$$\ln(1 - \varepsilon) \geq 2\varepsilon \ln\left(\frac{1}{2}\right) \geq -2\varepsilon.$$

Hence, we obtain

$$\begin{aligned} \mathbb{P}(T \geq n) &\geq \exp\left(-\frac{k\varepsilon}{2p^3}\right) = \exp\left(\frac{k\varepsilon}{2p^3}\right) \times \exp\left(-\frac{k\varepsilon}{p^3}\right) \\ &\geq \left(1 + \frac{k\varepsilon}{2p^3}\right) \exp\left(-\frac{k\varepsilon}{p^3}\right) \geq (1 + \varepsilon) \exp\left(-\frac{k\varepsilon}{p^3}\right). \end{aligned} \quad (38)$$

On the other hand, it follows from Chebychev's inequality that

$$\begin{aligned} \mathbb{P}(S_n < k) &\leq \mathbb{E}[p^{S_n}]p^{-k} = (1 - p + p^2)^n p^{-k} \leq \exp(-kf(p)) \\ &\leq \exp\left(-\frac{k}{16p^2}\right) \leq \exp\left(-\frac{1}{32p^2}\right) \exp\left(-\frac{k}{32p^2}\right) \leq \varepsilon \exp\left(-\frac{k\varepsilon}{p^3}\right) \end{aligned} \quad (39)$$

where, in the last inequality, we used our hypotheses $\varepsilon \geq \exp(-1/(32p^2))$ and $\varepsilon \leq p^3 \leq p/32$.

Now, plugging (38) and (39) into (37), we get

$$\mathbb{P}(S_T \geq k) \geq \exp\left(-\frac{k\varepsilon}{p^3}\right) \geq \left(1 - \frac{\varepsilon}{p^3}\right)^k,$$

using the inequality $\ln(1 - \varepsilon/p^3) \leq -\varepsilon/p^3$. This concludes the proof of (36), which is the claim of the Lemma. \square

6.7 Concluding proof of Proposition 4.2 in two dimensions

We now put the pieces together to obtain the lower bound on the stabilisation time of ARWD on the two-dimensional torus.

Proof of Proposition 4.2, case $d = 2$. Let $d = 2$, $\lambda > 0$, $\mu \in (0, 1)$ and $\rho \in (0, 1)$. We start by tuning the various parameters. Let K be the constant given by Lemma 6.1, let \bar{p} be the constant given by Lemma 6.6 and let

$$p = \bar{p} \wedge \frac{\sqrt{\mu}}{384} \wedge \frac{1}{2(1+\lambda)}.$$

Let v_0 be given by Lemma 6.3 (applied with $\rho_0 = \rho$), and let $v_1 \in \mathbb{N}$ be such that for all $v \geq v_1$, we have

$$2v^{10} \ln \left(\frac{1152v^3}{\mu} \right) \leq p^6 K \exp \left(\frac{1 - 2^{-1/4}}{2} v^{1/8} \right). \quad (40)$$

Let now $v = v_0 \vee v_1$ and $r = 2 \lfloor \sqrt{2v/\mu} \rfloor$. Then, for $j \in \mathbb{N}_0$ we define

$$\alpha_j = \frac{1 + 2^{-j/4}}{2\sqrt{v}}, \quad D_j = 6^j \times 12vr \quad \text{and} \quad p_j = \frac{p}{6^j v^{3/2}}.$$

Lastly, we take $\kappa = \alpha_0 \mu / 8$.

Let now $n \geq r+1$ and $A \subset \mathbb{Z}_n^2$ such that $|A| \geq \mu n^2$. Using Lemma 6.2, we can consider a (v, D) -dormitory hierarchy $(A_j, \mathcal{C}_j)_{j \leq J}$ on A such that $|A_j| \geq |A|/4$ and every set $C \in \mathcal{C}_0$ is r -connected.

Then, for every $C \in \mathcal{C}_0$ Lemma 6.3 provides us with a C -toppling procedure p_C such that the induction hypothesis $\mathcal{P}(0)$ defined in (27) holds, with $\rho_0 = \rho$. We then consider the toppling strategy f and the sleep mask g derived from this hierarchy and these toppling procedures, as explained in Section 6.3.

Let us now check that the assumptions of Lemma 6.5 are satisfied. For every $j \geq 0$ have $p_j \leq \bar{p}$ and $p_j \leq 1/(2^{j+1}(1+\lambda))$ and, moreover, using the definitions of p , of r and of D_j , along with the fact that $\alpha_j \leq 1$, we can write

$$p_j \leq \frac{\sqrt{\mu}}{384 \times 6^j v^{3/2}} \leq \frac{1}{96\sqrt{2} \times 6^j vr} = \frac{1}{8\sqrt{2} D_j} \leq \frac{1}{4\sqrt{2} (D_j + 1)} \leq \frac{1}{\sqrt{32\alpha_j (D_j + 1)}},$$

which shows that the condition (28) is satisfied. To check the condition (29), we write

$$\begin{aligned} \frac{p_j^6 \Upsilon_2(D_j) \exp((\alpha_j - \alpha_{j+1})2^{j/2}v)}{2^{j/2+2}v} &\geq \frac{p^6 K}{6^{6j+j/2+2}v^{10} \ln(6^j \times 12vr)} \exp\left(\frac{1 - 2^{-1/4}}{2} 2^{j/4} \sqrt{v}\right) \\ &\geq \frac{p^6 K}{4(4^j v)^{10} [\ln(24\sqrt{2/\mu}) + (3/2) \ln(4^j v)]} \exp\left(\frac{1 - 2^{-1/4}}{2} (4^j v)^{1/8}\right) \geq 1 \end{aligned}$$

following (40) applied with $4^j v$ instead of v . This shows that the condition (29) is also satisfied, allowing us to apply Lemma 6.5, to deduce that $\mathcal{P}(j)$ holds for all $j \leq J$. In particular we deduce that, starting with an initial configuration $U_0 \sim \mathcal{P}(A, \rho)$, the variable $1 + H^*(A_J, U_0 \cap A_J)$ dominates a geometric variable with parameter

$$\exp(-\alpha_J |A_J|) \leq \exp\left(-\frac{\alpha_0}{2} \times \frac{|A|}{4}\right) \leq \exp(-\kappa n^2).$$

Then, writing $\eta_0 = \mathbb{1}_{U_0} + \mathfrak{s} \mathbb{1}_{A \setminus U_0}$ and recalling that our toppling strategy f chooses sites in A_J in priority and that both the toppling strategy and the sleep mask do not depend on the configuration outside A_J , we have $T(A, \eta_0) \succeq T(A_J, \eta_0 \mathbb{1}_{A_J}) \geq H^*(A_J, U_0 \cap A_J)$, so that the stochastic domination also holds for $T(A, \eta_0)$. Decreasing the constant κ if necessary so that the result also holds when $n \leq r$, the proof of Proposition 4.2 in the two-dimensional case is complete. \square

6.8 Concluding proof of Proposition 4.2 in the transient case

We now study the case of dimension at least 3.

Proof of Proposition 4.2, case $d \geq 3$. Let $d \geq 3$, $\lambda > 0$, $\mu \in (0, 1)$ and $\rho \in (0, 1)$. Let $\beta = \rho/2$, let K be the constant given by Lemma 6.1 applied in dimension d and let

$$v = \left\lfloor \frac{(1+\lambda)}{(1-\beta)K} \right\rfloor + 1.$$

By definition of v , for α small enough we have

$$(1+\lambda)(1-e^{-\alpha}) \leq K(1-e^{-\alpha(1-\beta)v}).$$

Let us fix such an α , and let $\kappa < (\alpha\beta\mu) \wedge (\rho^2\mu/2)$.

Then, given $n \geq \lceil (v/\mu)^{1/d} \rceil$ and $A \subset \mathbb{Z}_n^d$ such that $|A| \geq \mu n^d$, we consider the trivial hierarchy with only one level, defined by $J = 0$, $A_0 = A$ and $\mathcal{C}_0 = \{A\}$ and we apply Lemma 6.4 with $r = \infty$ (note that $|A| \geq \mu n^d \geq v$). Thus, we obtain an A -toppling procedure p_A such that, for every fixed $U_0 \subset A$ with $|U_0| = \lfloor \beta|A| \rfloor$, the number $\mathcal{N}_{\mathcal{B}}$ of times that the number of unstable sites passes above the threshold $\beta|A|$ is such that $1 + \mathcal{N}_{\mathcal{B}}$ dominates a geometric random variable with parameter $e^{-\alpha(\beta|A|-2)}$ (applying a reasoning similar to that used in the proof of Lemma 6.3). Since $T(A, U_0) \geq \mathcal{N}_{\mathcal{B}}$, we deduce that this stochastic domination also holds for $T(A, U_0)$.

Then, to obtain the claimed result starting from a random initial configuration U_0 distributed as $\mathcal{P}_{A, \rho}$, we proceed as in the end of the proof of Lemma 6.3 using Hoeffding's inequality to estimate $\mathcal{P}_{A, \rho}(|U_0| < \beta|A|)$. We then obtain that, starting from $U_0 \sim \mathcal{P}_{A, \rho}$, the variable $T(A, U_0)$ dominates a geometric with parameter

$$e^{-\alpha(\beta|A|-2)} + e^{\rho^2\mu n^d/2} \leq e^{-\kappa n^d},$$

provided that n is taken large enough. Decreasing κ so that the estimate holds for any $n \in \mathbb{N}$, we obtain Proposition 4.2 in the transient case. \square

7 Exponential stabilisation time on the torus: proof of Theorem 1.3 given Propositions 4.1 and 4.2

Proof of Theorem 1.3. Let $d \geq 2$, let $\lambda = (2d)^3$ and let κ be given by Proposition 4.2 applied with $\mu_0 = 1/5$ and ρ_0 being the probability that a Poisson variable with parameter $4/5$ is at least 2, i.e.,

$$\rho_0 = 1 - \frac{1 + 4/5}{e^{4/5}}. \quad (41)$$

Let us consider the function

$$\psi : \mu \in (0, 1) \mapsto -\mu \ln \mu - (1 - \mu) \ln(1 - \mu).$$

Since $\psi(\mu) \rightarrow 0$ when $\mu \rightarrow 1$, we can find $\mu_1 \in (4/5, 1)$ such that $\psi(\mu_1) < \kappa/8$.

Let now $\mu \in (\mu_1, 1)$. Let $u > 0$ be such that $J(u) > 0$, where the function J is given by

$$J : u \mapsto \mu(1 - e^{-u}) - \mu_1 u.$$

The existence of such a u follows from the fact that $J(0) = 0$ and $J'(0) = \mu - \mu_1 > 0$.

Let $n \in \mathbb{N}$ be large enough so that

$$\frac{(n-2)^d}{2} \geq \frac{2n^d}{5} \quad (42)$$

and let η be a random SSM configuration on the torus \mathbb{Z}_n^d with i.i.d. Poisson numbers of particles on each site, with parameter μ . Let $M \in \mathbb{N}$. Let us start by writing

$$\mathcal{P}^\mu(\|m_\eta\| < M) \leq P_\mu(|\eta| < \mu_1 n^d) + \sum_{k=\lceil \mu_1 n^d \rceil}^{n^d} \mathcal{P}^\mu(\|m_\eta\| < M, |\eta| = k). \quad (43)$$

Note that if $|\eta| > n^d$ then $\|m_\eta\| = \infty$: this is why we consider the sum only until $k = n^d$.

To handle the first term in (43), recalling that $|\eta|$ is a Poisson variable with parameter μn^d , we use the following Chernoff bound:

$$P_\mu(|\eta| < \mu_1 n^d) \leq E_\mu(e^{-u|\eta|})e^{u\mu_1 n^d} = e^{-J(u)n^d}. \quad (44)$$

We now deal with the second term of (43). Let us take $\lceil \mu_1 n^d \rceil \leq k \leq n^d$. We first write

$$\mathcal{P}^\mu(\|m_\eta\| < M, |\eta| = k) = \sum_{A \subset \mathbb{Z}_n^d : |A|=k} \mathcal{P}^\mu(\|m_\eta\| < M, |\eta| = k, \eta_\infty = \mathbb{1}_A). \quad (45)$$

Let us fix a set $A \subset \mathbb{Z}_n^d$ with $|A| = k$. Then, recall that Lemma 2.2 tells us that if $\eta_\infty = \mathbb{1}_A$ then $m_\eta = m_\eta^A$, whence

$$\mathcal{P}^\mu(\|m_\eta\| < M, |\eta| = k, \eta_\infty = \mathbb{1}_A) \leq \mathcal{P}^\mu(\|m_\eta^A\| < M, |\eta| = k). \quad (46)$$

Then, conditioning on the initial configuration, we get

$$\mathcal{P}^\mu(\|m_\eta^A\| < M, |\eta| = k) = \sum_{\eta_0 \in \mathbb{N}_0^V : |\eta_0|=k} P_\mu(\eta = \eta_0) \mathcal{P}(\|m_{\eta_0}^A\| < M) \quad (47)$$

We wish to apply Proposition 4.1. To this end, we need a totally disconnected subset of A . Considering the standard projection application $\pi_n : \mathbb{Z}^d \rightarrow \mathbb{Z}_n^d$, we consider the subset B_0 of the ‘‘even sites’’ of the torus, removing a line to avoid problems if n is odd, namely

$$B_0 = \pi_n\left(\{(x_1, \dots, x_d) \in \{0, \dots, n-2\}^d : x_1 + \dots + x_d \text{ even}\}\right).$$

This set is totally disconnected and we have $|B_0| \geq (n-2)^d/2 \geq 2n^d/5$, following our assumption (42) on n . Then, we simply consider the set $B = A \cap B_0$, so that B is also totally disconnected and satisfies $|B| \geq |A| - 3n^d/5 \geq n^d/5$. Let f be a toppling strategy and g be a sleep mask provided by Proposition 4.2 applied with λ, μ_0 and ρ_0 defined above, with κ already fixed above and with this set B , so that if U is distributed according to \mathcal{P}_{B, ρ_0} then $1 + T(B, \mathbb{1}_U + \mathfrak{s}\mathbb{1}_{B \setminus U})$ dominates a geometric with parameter $\exp(-\kappa n^d)$.

At this point, Proposition 4.1 entails that, for every fixed $\eta_0 \in \mathbb{N}_0^V$ with $|\eta_0| = k$, we have

$$\mathcal{P}(\|m_{\eta_0}^A\| < M) \leq \mathbb{P}(T(B, \eta'_0) < M),$$

where $\eta'_0 = \mathbb{1}_U + \mathfrak{s}\mathbb{1}_{B \setminus U}$, with $U = \{x \in B : \eta_0(x) \geq 2\}$ and $T(B, \eta'_0)$ is the stabilisation time of the ARWD($\lambda, B, f, g, \eta'_0$) process. Plugging this into (47) leads to

$$\begin{aligned} \mathcal{P}^\mu(\|m_\eta^A\| < M, |\eta| = k) &\leq \sum_{\eta_0 \in \mathbb{N}_0^V : |\eta_0|=k} P_\mu(\eta = \eta_0) \mathbb{P}(T(B, \eta'_0) < M) \\ &\leq \sum_{\eta_0 \in \mathbb{N}_0^V} P_\mu(\eta = \eta_0) \mathbb{P}(T(B, \eta'_0) < M) \\ &= \sum_{U \subset B} \rho^{|U|} (1 - \rho)^{|B \setminus U|} \mathbb{P}(T(B, \mathbb{1}_U + \mathfrak{s}\mathbb{1}_{B \setminus U}) < M), \end{aligned}$$

where ρ is the probability that a Poisson variable with parameter μ is at least 2. Note that since $\mu \geq 4/5$ we have $\rho \geq \rho_0$, with ρ_0 defined by (41). Now, since the toppling strategy f and the sleep mask g were provided by Proposition 4.2, we have

$$\sum_{U \subset B} \rho^{|U|} (1 - \rho)^{|B \setminus U|} \mathbb{P}(T(B, \mathbb{1}_U + \mathfrak{s}\mathbb{1}_{B \setminus U}) < M) \leq \mathbb{P}(G_n \leq M),$$

where G_n is a geometric variable with parameter $\exp(-\kappa n^d)$. Now, taking $M = e^{an^d}$ with $a = \kappa/2$, we have

$$\mathbb{P}(G_n \leq e^{an^d}) = 1 - \exp\left(\lfloor e^{\kappa n^d/2} \rfloor \ln(1 - e^{-\kappa n^d})\right) \leq e^{-\kappa n^d/4},$$

provided that n is large enough. Combining all this and going back to (45) and (46), we obtain

$$\mathcal{P}^\mu(\|m_\eta\| < e^{an^d}, |\eta| = k) \leq \binom{n^d}{k} e^{-\kappa n^d/4}.$$

Plugging this and (44) into (43) then yields

$$\mathcal{P}^\mu(\|m_\eta\| < e^{an^d}) \leq e^{-J(u)n^d} + \sum_{k=\lceil \mu_1 n^d \rceil}^{n^d} \binom{n^d}{k} e^{-\kappa n^d/4} \leq e^{-J(u)n^d} + n^d \binom{n^d}{\lceil \mu_1 n^d \rceil} e^{-\kappa n^d/4}.$$

Using now that

$$\binom{n^d}{\lceil \mu_1 n^d \rceil} = O(e^{\psi(\mu_1)n^d}),$$

along with our assumption that $\psi(\mu_1) < \kappa/8$, we deduce that, taking $b = \min(J(u)/2, \kappa/16)$, provided that n is large enough,

$$\mathcal{P}^\mu(\|m_\eta\| < e^{an^d}) \leq e^{-bn^d}.$$

Then, adapting Section 3.10 of [11] to deal with the continuous-time model with the exponential clocks, we obtain the claimed property (1). \square

8 Proof of Theorem 1.4

Theorem 1.4 is an analog of Theorem 4 of [11], which concerns activated random walks. The proof extends to the SSM with only one minor change that we explain below.

The proof of Theorem 4 of [11] relies on a criterion for non-fixation for ARW established by Rolla and Tournier (Proposition 3 in [23]). As explained by Rolla in [24], this result extends to the SSM with no substantial change in the proof. The proof of Theorem 4 of [11] also uses a rough upper bound for the stabilisation time of a box in ARW (Lemma 13 of [11]), which also holds for the SSM, with the same proof (apart from the point dealing with sleep instructions, but this is harmless).

Then, the proof of Theorem 4 of [11] extends to the SSM without problem, apart from the fourth step of the procedure, which uses the property that, in activated random walks, a configuration with k sites each containing one active particle, all other sites being stable, can be stabilised with no particle moving, with probability at least $(\lambda/(1+\lambda))^k$: this simply happens if at each of these k sites, the particle falls asleep before moving.

Thus, to adapt the proof to the SSM, it is enough to show the following Lemma, which is an analog of this property for the SSM:

Lemma 8.1. *For every graph $G = (V, E)$ with maximal degree Δ , for every $A \subset V$ finite and connected, for every SSM configuration $\eta : V \rightarrow \mathbb{N}_0$ with at most one particle on each site of A and at least one empty site in A , for every odometer $h : V \rightarrow (1/2)\mathbb{N}_0$, the probability that no particle exits A during the stabilisation of (η, h) in A with half-legal topplings is at least $\Delta^{-2|A|}$.*

Proof. Let $G = (V, E)$ be a graph with maximal degree Δ . We prove the result by induction on $|A|$. If $|A| = 1$ then there is nothing to show, since we assume that at least one site of A is empty.

Assume that $k \in \mathbb{N}$ is such that the result holds for every $A \subset V$ with cardinality $|A| = k$.

Then, let $A \subset V$ connected with $|A| = k + 1$, let $\eta : V \rightarrow \mathbb{N}_0$ with at most one particle on each site of A and at least one empty site in A , and let $h : V \rightarrow (1/2)\mathbb{N}_0$. Let x and y be two distinct sites of A such that both $A \setminus \{x\}$ and $A \setminus \{y\}$ remain connected (take for example two leaves of a rooted spanning tree of A , or the root and the leaf if the tree has only one leaf).

If $\eta(x) = \eta(y) = 0$, then the result follows from the induction hypothesis applied to $A \setminus \{x\}$.

Assume now that at least one of these two points contains a particle: assume for example that $\eta(x) = 1$. If $h(x)$ is integer, then the result also follows from the induction hypothesis applied to $A \setminus \{x\}$, because x is already half-stable in (η, h) . Assume now that $h(x)$ is not integer. Let z be a neighbour of x in A . Then, we

perform one half-legal half-toppling at x . With probability at least $1/\Delta$, a particle jumps from x to z . Then, we perform an admissible half-toppling at z . With probability at least $1/\Delta$, a particle jumps back from z to x , and we obtain the configuration $(\eta, h + \frac{1}{2}\mathbb{1}_x + \frac{1}{2}\mathbb{1}_z)$, in which the site x is now half-stable. Applying the induction hypothesis to this configuration in $A \setminus \{x\}$, we know that with probability at least Δ^{-2k} , the configuration stabilises in $A \setminus \{x\}$ with no particle jumping out of $A \setminus \{x\}$. Thus, with probability at least Δ^{-2k-2} , there exists an admissible half-toppling sequence which half-stabilises (η, h) in A with no particle jumping out of A . This implies the result by monotonicity (noting that the number of particles which jump out of A is a non-decreasing function of the odometer). \square

With the notation of the proof of Theorem 4 in [11], this lemma enables to lower bound the success probability of the fourth step of the procedure described in [11], applying the lemma to $A = B_0 \setminus B_3$ (or to its two connected components if $d = 1$), and adjusting the constants accordingly.

Acknowledgements

The authors thank the German Research Foundation (project number 444084038, priority program SPP2265) for financial support.

References

- [1] Enrique Daniel Andjel. Invariant measures for the zero range processes. *Ann. Probab.*, 10(3):525–547, 1982.
- [2] Amine Asselah, Nicolas Forien, and Alexandre Gaudillière. The critical density for activated random walks is always less than 1. *Ann. Probab.*, 52(5):1607–1649, 2024.
- [3] Per Bak. *How Nature Works: The Science of Self-Organized Criticality*. Springer Science & Business Media, 2013.
- [4] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Phys. Rev. Lett.*, 59:381, 1987.
- [5] Riddhipratim Basu, Shirshendu Ganguly, and Christopher Hoffman. Non-fixation for conservative stochastic dynamics on the line. *Comm. Math. Phys.*, 358(3):1151–1185, 2018.
- [6] Alexandre Bristiel and Justin Salez. Separation cutoff for Activated Random Walks. *arXiv preprint arXiv:2209.03274*, 2022.
- [7] Elisabetta Candellero, Alexandre Stauffer, and Lorenzo Taggi. Abelian oil and water dynamics does not have an absorbing-state phase transition. *Trans. Amer. Math. Soc.*, 374(4):2733–2752, 2021.
- [8] Deepak Dhar. The abelian sandpile and related models. *Physica A*, 263:4–25, 1999.
- [9] Ronald Dickman. Nonequilibrium phase transitions in epidemics and sandpiles. *Physica A: Statistical Mechanics and its Applications*, 306:90–97, 2002. Invited Papers from the 21th IUPAP International Conference on Statistical Physics.
- [10] Anne Fey, Lionel Levine, and David B. Wilson. Approach to criticality in sandpiles. *Physical Review E*, 82(3):031121, 2010.
- [11] Nicolas Forien and Alexandre Gaudillière. Active phase for activated random walks on the lattice in all dimensions. *Ann. Inst. Henri Poincaré Probab. Stat.*, 60(2):1188–1214, 2024.
- [12] Christopher Hoffman, Yiping Hu, Jacob Richey, and Douglas Rizzolo. Active Phase for the Stochastic Sandpile on \mathbb{Z} . *preprint arXiv:2212.08293*, 2022.
- [13] Christopher Hoffman, Tobias Johnson, and Matthew Junge. The density conjecture for activated random walk. *arXiv*, 2024. 69 pages.

- [14] 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.
- [15] Yiping Hu. Active Phase for Activated Random Walk on \mathbb{Z}^2 . *arXiv preprint arXiv:2203.14406*, 2022.
- [16] Hoai Nguyen Huynh, Gunnar Pruessner, and Lock Yue Chew. The Abelian Manna model on various lattices in one and two dimensions. *Journal of Statistical Mechanics: Theory and Experiment*, 2011(09):P09024, 2011.
- [17] Hang-Hyun Jo and Hyeong-Chai Jeong. Comment on “driving sandpiles to criticality and beyond”. *Physical Review Letters*, 105(1):019601, 2010.
- [18] Lionel Levine and Feng Liang. Exact sampling and fast mixing of Activated Random Walk. *arXiv preprint arXiv:2110.14008*, 2021.
- [19] Lionel Levine and Vittoria Silvestri. Universality conjectures for activated random walk. *Probab. Surv.*, 21:1–27, 2024.
- [20] Sven Lübeck. Universal scaling behavior of non-equilibrium phase transitions. *International Journal of Modern Physics B*, 18(31n32):3977–4118, 2004.
- [21] Subhrangshu Sekhar Manna. Two-state model of self-organized criticality. *Journal of Physics A: Mathematical and General*, 24(7):L363–L369, 1991.
- [22] Moumanti Podder and Leonardo T. Rolla. Uniform threshold for fixation of the stochastic sandpile model on the line. *J. Stat. Phys.*, 182(3):Paper No. 52, 16, 2021.
- [23] L. T. Rolla and L. Tournier. Non-fixation for biased activated random walks. *Ann. Inst. Henri Poincaré Probab. Stat.*, 54(2):938–951, 2018.
- [24] Leonardo T. Rolla. Activated random walks on \mathbb{Z}^d . *Probab. Surv.*, 17:478–544, 2020.
- [25] 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.
- [26] Leonardo T. Rolla, Vladas Sidoravicius, and Olivier Zindy. Universality and sharpness in activated random walks. *Ann. Henri Poincaré*, 20(6):1823–1835, 2019.
- [27] Vladas Sidoravicius and Augusto Teixeira. Absorbing-state transition for stochastic sandpiles and activated random walks. *Electron. J. Probab.*, 22:Paper No. 33, 35, 2017.
- [28] Alexandre Stauffer and Lorenzo Taggi. Critical density of activated random walks on transitive graphs. *Ann. Probab.*, 46(4):2190–2220, 2018.
- [29] Lorenzo Taggi. Essential enhancements in Abelian networks: Continuity and uniform strict monotonicity. *Ann. Probab.*, 51(6):2243–2264, 2023.