Probability Surveys
Vol. 17 (2020) 478–544
ISSN: 1549-5787
https://doi.org/10.1214/19-PS339

Activated Random Walks on $\mathbb{Z}^d$*

Leonardo T. Rolla

e-mail: leorolla@dm.uba.ar

Abstract: Some stochastic systems are particularly interesting as they exhibit critical behavior without fine-tuning of a parameter, a phenomenon called self-organized criticality. In the context of driven-dissipative steady states, one of the main models is that of Activated Random Walks. Long-range effects intrinsic to the conservative dynamics and lack of a simple algebraic structure cause standard tools and techniques to break down. This makes the mathematical study of this model remarkably challenging. Yet, some exciting progress has been made in the last ten years, with the development of a framework of tools and methods which is finally becoming more structured. In these lecture notes we present the existing results and reproduce the techniques developed so far.

MSC2020 subject classifications: Primary 60K35, 82C22, 82C26.
Keywords and phrases: Absorbing-state phase transition.

Received December 2019.

Contents

*This is an original survey paper

1. Overview

In the study of critical phenomena, there is a large class of non-equilibrium lattice systems that naturally evolve to a critical state, characterized by power-law distributions for the sizes of relaxation events. A typical example is the occurrence of huge avalanches caused by small perturbations. In many cases, such systems are attracted to a stationary critical state without being specifically tuned to a critical point.

This seems to be the explanation for the emergence of random fluctuations at a macroscopic or mesoscopic scale, and creation of self-similar shapes in a variety of growth systems. Among attempts to explain long-ranged spatial-temporal correlations, the physical paradigm called *self-organized criticality* is a widely accepted theory, although it is still very poorly understood from a mathematical point of view.

For non-equilibrium steady states, it eventually became evident that self-organized criticality is related to conventional critical behavior of a system undergoing a phase transition. In the case of driven-dissipative systems, it is related to an *absorbing-state phase transition*. These are systems whose dynamics drives them towards, and then maintains them at the edge of stability. The phase transition arises from the conflict between spread of activity and a tendency for this activity to die out, and the critical point separates a phase with sustained activity and an absorbing phase in which the dynamics is eventually extinct in any finite region.

The main stochastic models in this class are the Manna Sandpile Model, its Abelian variant which we call the Stochastic Sandpile Model, and the Activated Random Walks (ARW). Due to long-range effects intrinsic to their conservative dynamics, classical analytic and probabilistic techniques fail in most cases of interest, making the rigorous analysis of such systems a major mathematical challenge. Yet some exciting progress has been made in the last ten years. In particular, a more structured framework of tools and methods started to develop and emerge.

In these notes we recall the existing results, and describe in a unified framework the tools and techniques currently available. Most of the material is devoted to the ARW. In §12.3 we mention the Stochastic Sandpile Model, for which much less is known. The Manna model so far seems intractable.

In §1.1 we describe the local rules for the ARW evolution, in §1.2 we discuss the relation between self-organized criticality and absorbing-state phase transitions, and in §1.3 we cast some of the physical predictions. Most of them are far outside the reach of currently known mathematical techniques, as discussed in §1.4. Then in §1.5 we quote all the known results (which, as the reader will see, leave most of the predictions as open conjectures), and in §1.6 we describe the main methods developed in the past ten years. Finally, in §1.7 we discuss the structure and interdependence of the remaining sections.

1.1. Activated Random Walks

The ARW evolution is defined as follows. Particles sitting on the graph $\mathbb{Z}^d$ can be in state A for *active* or S for *sleeping*. Each active particle, that is, each particle in the A state, performs a continuous-time random walk with *jump rate* $D_A = 1$. The walks follow a translation-invariant *jump distribution*, that is, they jump from x to $x+z$ with probability $p(z)$ for some fixed distribution $p(\cdot)$ on $\mathbb{Z}^d$. We assume that the support of p is restricted to the nearest-neighbors $\pm e_1, \dots, \pm e_d$ and spans all of $\mathbb{Z}^d$.

Several active particles can be at the same site. When a particle is alone, it may fall asleep, a *reaction* denoted by $A \to S$, which occurs at a *sleep rate* $0 < \lambda \leqslant \infty$. So each particle carries two clocks, one for jumping and one for sleeping. Once a particle is sleeping, it stops moving, *i.e.* it has jump rate $D_S = 0$, and it remains sleeping until the instant when another particle is present at the same site. At such an instant the particle which is in the S state flips to the A state, giving the *reaction* $A + S \to 2A$.

If the clock rings for a particle to sleep while it shares a site with other particles, the tentative transition $A \to S$ is overridden by the instantaneous reaction $A + S \to 2A$, so this attempt to sleep has no effect on the system configuration.

A particle in the S state stands still forever if no other particle ever visits the site where it is located. When a site has no particles or one sleeping particle, it is called *stable*, otherwise it has one or more particles, all active, and is called *unstable*. A stable site stays stable indefinitely, and can only become unstable if later on it is visited by an active particle. An *absorbing configuration* is one for which every site is stable.

We note that, at the extreme case $\lambda = \infty$, when a particle visits an unoccupied site, it falls asleep instantaneously. This case is equivalent to *internal diffusion-limited aggregation* with multiple sources.

We have described *local rules* for the system to evolve. In order to fully describe the system, we need to specify on which subset of $\mathbb{Z}^d$ this dynamics will occur, what are the boundary conditions, and the initial state at $t = 0$. By "state" we mean a probability distribution on the space of configurations, unless it refers to the state A or S of a particle.

1.2. Phase transition and self-organized criticality

We consider two different dynamics which follow the above local rules.

Infinite-volume conservative system. On the infinite lattice $\mathbb{Z}^d$, at $t = 0$ we start from a translation-ergodic state with average density of particles ζ. It turns out that, because the dynamics does not create or destroy particles, this average density ζ is conserved during the evolution. We say that this system fixates, or is absorbed, if each site is visited only finitely many times and eventually becomes stable. Otherwise, if each site is visited infinitely many times, we say that the system stays active. This model shows an *ordinary phase transition* in the sense that, for some ζ_c, the dynamics a.s. fixates when $\zeta < \zeta_c$ and a.s. stays active when $\zeta > \zeta_c$. This is called an *absorbing-state phase transition*.

Driven-dissipative system. On a finite box $V_L = \{-L, \dots, L\}^d$, we define a system with three components. At rate one, a new active particle is added to a site $x \in V$ chosen uniformly at random. The ARW dynamics is run with time being accelerated by a factor of $\kappa > 1$. The box has open boundary, *i.e.* particles are killed when they exit V.

We first let $\kappa \to \infty$, so that the whole box is stabilized right after a particle is added, so we obtain a Markov chain on the space of absorbing configurations called *driven-dissipative dynamics*. We then let $t \to \infty$ to reach a stationary state ν_s^L supported on absorbing configurations. We finally let $L \to \infty$ to have a state ν_s on $\mathbb{Z}^d$ with mean density ζ_s, see §1.3.

Self-organized criticality for the driven-dissipative system loosely means the following. When the average density ζ inside the box is too small, mass tends to accumulate. When it is too large, there is intense activity and a substantial number of particles are killed at the boundary. With this mechanism, the model

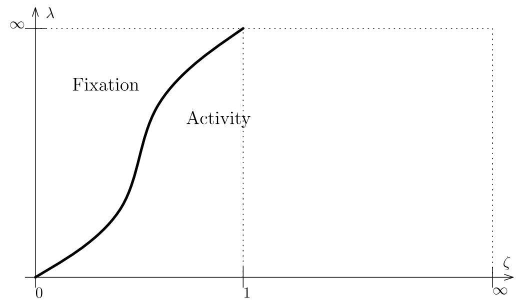

FIG 1.1. *Prediction for the phase space.*

is attracted to a steady state with an average density given by $0 < \zeta_s < \infty$, and this state has several features associated to criticality. The *density conjecture* says that ζ_c and ζ_s should coincide. Moreover, the critical exponents of the driven-dissipative system should be related to those of the infinite-volume one.

1.3. Predictions

Consider a system running on the whole graph $\mathbb{Z}^d$. At $t = 0$, sites have an i.i.d. Poisson number of active particles with parameter ζ.

We say that the system *fixates* if, for each site, there is a random time after which the site is either vacant or contains one sleeping particle. We say that the system *stays active* if, for each site, there are arbitrarily large times at which the site has at least one active particle. There is a *critical density* ζ_c, which is non-decreasing in λ, such that the system will a.s. fixate for $\zeta < \zeta_c$ and the system will a.s. stay active for $\zeta > \zeta_c$.

In this subsection we describe some aspects of the ARW behavior. Some of these claims have been proved, most of them remain widely open.

Phase space. The critical density satisfies $\zeta_c < 1$ for every $\lambda < \infty$ and $\zeta_c \to 0$ as $\lambda \to 0$. That $\zeta_c \leqslant 1$ should be obvious from the fact that each site can accommodate at most one sleeping particle. But the particles should be able to sustain activity even at densities lower than unit (except for the case $\lambda = \infty$). Moreover, $\zeta_c > 0$ for every $\lambda > 0$ and $\zeta_c \to 1$ as $\lambda \to \infty$. In particular, $\zeta_c = 1$ when $\lambda = \infty$. More generally, ζ_c is continuous and strictly increasing in λ. See Figure 1.1.

Uniqueness of the critical density. The value of ζ_c depends on the dimension d, on the jump distribution $p(\cdot)$ and on the sleep rate λ. But it does not depend on the choice of i.i.d. Poisson for the initial state. More precisely, for every translation-ergodic active initial state (see next paragraph) with density ζ, the system a.s. fixates if $\zeta < \zeta_c$ and a.s. stays active if $\zeta > \zeta_c$. At $\zeta = \zeta_c$ the system should also stay active (except for the trivial case of $\lambda = \infty$ and non-random initial condition with density $\zeta = 1$).

Invariant measures and convergence. For each $\zeta < \zeta_c$, all the translation-ergodic stationary distributions with average density ζ are *absorbing states*, that is, they are measures supported on absorbing configurations. Starting from any state with such density, the evolution a.s. converges to an absorbing configuration having the same density. In general, different initial states are attracted to different absorbing states.

For each $\zeta > \zeta_c$, there is a unique translation-ergodic stationary *active state* (*i.e.* a state which is not absorbing, which by ergodicity means a positive fraction of the particles are active) with density ζ. For every translation-ergodic active initial state with density ζ, the system converges in law to this unique stationary active state as $t \to \infty$. So in terms of basin of attraction, the active state is stable and absorbing states are unstable, in conflict with our terminology for stable and unstable sites at the microscopic level.

At $\zeta = \zeta_c$ and active initial states, the system converges in law to a unique absorbing state, but the system a.s. stays active and convergence is in law only. So the critical case mixes features from both phases.

Power laws at and near criticality. At $\zeta = \zeta_c$, the average density of activity (number of active particles per site) at time t decays as a power of t as $t \to \infty$. On the other hand, if $\zeta > \zeta_c$, the density of activity seen in the stationary regime $t = \infty$ decays as a power of $\zeta - \zeta_c$ as $\zeta \downarrow \zeta_c$. Two-point correlations in space also decay as a power of the distance Δx, and same-site time correlations decay as a power of Δt. Outside criticality, correlation decays exponentially with a typical correlation length for space and another one for time, and the correlation lengths themselves diverge as powers of $|\zeta - \zeta_c|$ as $\zeta \to \zeta_c$. More details on critical exponents can be found in [DRS10].

Driven-dissipative dynamics. Consider the driven-dissipative dynamics described in §1.2. Let ζ_s^L be the average density of particles in ν_s^L, and let $\zeta_s = \lim_L \zeta_s^L$. The *density conjecture* says that $\zeta_s = \zeta_c$. Moreover, the state ν_s should have the same two-point correlation decay exponents as conservative infinite-volume system at criticality.

A stronger version of the density conjecture is the following. Let $\zeta \in [0, \infty)$. Consider an i.i.d. Poisson configuration with density ζ on the box V_L, and run the ARW dynamics with open boundary until it reaches an absorbing configuration. Then the final state has density concentrated around some value which depends on ζ and L, and this value tends to $\min\{\zeta, \zeta_c\}$ as $L \to \infty$.

Fixed-energy dynamics. Consider the ARW dynamics on a large torus $\mathbb{Z}_n^d = (\mathbb{Z}/n\mathbb{Z})^d$ instead of $\mathbb{Z}^d$, starting with approximately ζn^d particles. Almost surely, this system will eventually fixate if and only if it has fewer than n^d particles, so the question is not *whether* it fixates. The relevant quantity is *how long* it will take the system to fixate. For parameters λ and ζ inside the active phase for $\mathbb{Z}^d$ as shown in Figure 1.1, it should take a long time to fixate (exponential in n^d), with high probability. For parameters inside the fixating phase, the corresponding system on a large torus should fixate in a short time (faster than any positive power of n). For parameters on the critical curve, the time to fixate should be in between. For the critical and near-critical regimes, several quantities should decay or blow up as power laws, such as space and time correlation lengths, activity decay, etc.

1.4. Open problems and challenges

In principle, all the statements in §1.3 that do not appear in §1.5 are open problems. But most of them are far beyond the reach of current techniques. Throughout these lecture notes we will explicitly mention some more realistic open problems, after the background needed to properly state each question has been introduced.

The first difficulty in studying the ARW lies in the fact that this system is not attractive. This is overcome by considering a site-wise kind of construction, or an explicit construction in terms of a collection of random walks, rather than the Harris graphical construction. These frameworks allow for different kinds of arguments which have proven to be very useful.

Another feature of this model –particle conservation– still causes tremendous difficulties. This has so far restrained most attempts to apply arguments of the type "energy vs. entropy." These arguments typically go as follows. One first identifies some structure that is intrinsic to the occurrence of events that conjecturally should not occur. All possible structures are then enumerated, and their number is overwhelmed by the high probabilistic cost needed for them to occur. One then concludes that such events have vanishing probability. This approach has been very successful in many branches of statistical mechanics. However, for reaction-diffusion dynamics, the conservation of particles in the system gives rise to intricate long-range effects, which makes it difficult to find suitable structures within the occurrence of events of interest. In §5 we present the only case where an approach involving enumeration of events and compensation by an extreme choice of parameters has been implemented with some success.

1.5. Results

We now briefly summarize known results towards the above predictions. The results will be stated under the common assumption that the jumps are to nearest neighbors only. In the next sections we will go through the proof of all the results mentioned here. For bibliographical references, see §12.

Phase space. Results regarding the phase space for the infinite-volume system on $\mathbb{Z}^d$ are summarized in Figure 1.2. It is known that $\zeta_c \geqslant \frac{\lambda}{1+\lambda}$ in general (§4 and §7). In particular, there is a fixation phase for every fixed $\lambda > 0$ by taking ζ small, and there is a fixation phase for every fixed $\zeta < 1$ by taking λ large. It is also known that, for every $\lambda \leqslant \infty$, there is no fixation at $\zeta = 1$ (§10).

In the case of unbiased walks (*i.e.* $\sum_y yp(y) = \mathbf{0}$) on $d = 2$, this is all we know, and proving existence of a non-trivial active phase (*i.e.* for some $\lambda > 0$

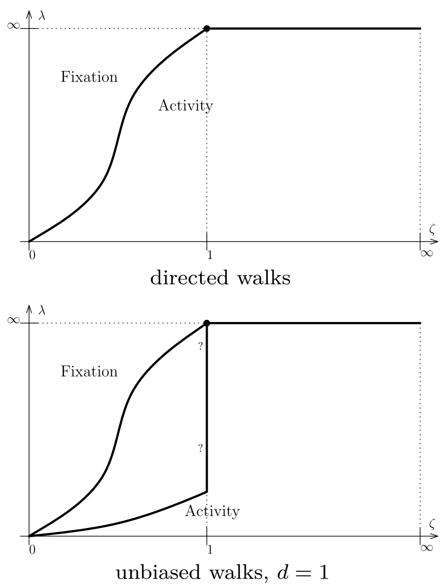

Fig 1.2. *Results about fixation vs. activity on the phase space.*

and $\zeta < 1$) remains an open problem. On $d = 1$, for every $\zeta > 0$ fixed there is an active phase by taking λ small (§5), but it remains open to show existence of a non-trivial active phase for every fixed $\lambda < \infty$ and ζ close enough to 1.

For transient walks, the picture is fairly complete: for every fixed $\zeta > 0$ the system stays active if λ is small enough, and for every fixed $\lambda < \infty$ the system stays active if ζ is close enough to 1. There are different proofs for unbiased (§7) and biased (§3) walks. If the walks are not only biased but directed (*i.e.* $p(-e_j) = 0$ for every j), the critical curve can be described explicitly: $\zeta_c = \frac{\lambda}{1+\lambda}$, and for $d = 1$, there is no fixation at $\zeta = \zeta_c$ (§3).

Fixed-energy dynamics. For the one-dimensional torus (*i.e.* the ring $\mathbb{Z}_n$), it has been shown that there is a slow stabilization phase and a fast stabilization phase. Consider the average activity time $\mathcal{T}$ given by the sum of the total time each particle is active, divided by n. For every $0 < \zeta < 1$ fixed, if λ is large enough then $\mathcal{T} \leqslant C \log^2 n$, and if λ is small enough then $\mathcal{T} \geqslant e^{cn}$, with high probability as $n \to \infty$ (§6).

Uniqueness of the critical density. The prediction given in §1.3 says that translation-ergodic distributions with average particle density $\zeta > \zeta_c$ stay active as long as they are not supported on absorbing configurations. The partial result presented in §8 holds for distributions supported on completely active configurations, that is, without any sleeping particle. Closing this gap is a major question and would probably be an important step towards the density conjecture.

1.6. Methods

Most proofs rely on the properties of the site-wise representation described in §2.2. To study the phase space and establish regions of fixation and activity, we normally check one of the conditions stated in §2.3. Uniqueness of the critical density also allows us to make convenient assumptions about the initial distribution.

In practice, to check one of these conditions we need to describe a toppling procedure for which probabilistic estimates can be obtained. *Toppling* is a one-step update of the current configuration according to the dynamics described in §1.1, and the *Abelian property* allows us to choose which site should be toppled ignoring the actual order in the continuous-time dynamics. A *toppling procedure* is a recipe that specifies the next site to be toppled, usually (but not necessarily) in terms of the outcome of previous topplings. The simpler examples of this general strategy are gathered in §3 and illustrate this principle well. From §4 to §9 all proofs follow this common setup, each one with its own specific elements.

Another method of analysis is the use of a particle-wise construction. This construction allows different uses of the mass transport principle, coupling, resampling, and ergodicity. These arguments are shown in §10.

1.7. Structure of these lecture notes

The main results and some open problems are stated at the beginning of each section. The reader may want to first have a quick glance at each section to have a sense of what is going on, then read §2 skipping the proofs, and again skim through the other sections. After that, the advice is to read the rest of the text linearly, maybe skipping computations.

The ordering of sections was decided taking into account relevance, difficulty and interdependence. Proofs are intended to be self-contained and have the level of detail of a research article. Although each section is written assuming that the reader is familiar with the material presented before it, reading the text linearly is not a strict requirement. To follow a section in full detail, going through §2 is mandatory, and going through §3 is highly recommended for most parts. The exception is §6 which uses a result from §5 and arguments from §4.

These notes are organized as follows. In §2 we describe the site-wise representation and state the main criteria to study the absorbing-state phase transition, establishing the main tools used in subsequent sections. In §3 we provide the simplest examples of a toppling procedure being used to prove fixation and activity by verifying the criteria provided in §2. In §4 we give a more sophisticated toppling procedure used to prove lower bounds for ζ_c on $d = 1$. In §5 we describe a two-scale toppling procedure and an enumerative argument to prove upper bounds for ζ_c (or lower bounds for λ_c) for unbiased walks on $d = 1$. In §6 we combine arguments from §4 and §5 plus a new argument based on a certain urn process to study fast and slow fixation on a large ring. In §7 we describe a toppling procedure based on the notion of weak and strong stabilizations to obtain a general lower bound for ζ_c valid in any dimension as well as upper bounds on ζ_c valid in the transient case $d > 2$. In §8 we use a locally-finite infinite-step parallel-update toppling procedure to show that the value of ζ_c is independent of the choice of Poisson as the initial state. In §9 we briefly sketch a recursive multi-scale estimate based on a toppling procedure that uses ideas of decoupling to prove lower bounds for ζ_c unbiased walks on $d \geqslant 2$. In §10 we depart from the framework of site-wise representation and use a different type of construction where particles are labeled. We then devise other properties of the ARW such as mass conservation, and prove an averaged condition for activity. In §11 we prove that the continuous-time evolution is well-defined and can be constructed explicitly using both the site-wise as well as particle-wise constructions. We also prove the equivalence between fixation and a condition on the site-wise representation. We then use a hybrid construction to prove a comparison lemma used in the proof of the averaged condition. Finally, in §12, we describe how and when these results were first proved, then comment on some of the arguments that extend to other graphs, initial conditions and jump distributions, and mention some of the arguments which have meaningful counter-parts for the Stochastic Sandpile Model.

2. Definitions and main tools

In this section we define precisely the stochastic process to be studied, describe the site-wise representation, and give conditions for fixation and activity. We then state mass conservation and ergodicity properties used later on, and conclude collecting frequently used notation.

2.1. The stochastic process and notation

We will denote by $\eta_t(x)$ the number and type of particles at site x at time t, as follows. Let $\mathbb{N}_0 = \{0, 1, 2, \dots\}$ and $\mathbb{N}_\mathfrak{s} = \mathbb{N}_0 \cup \{\mathfrak{s}\}$. The configuration of the ARW at time $t \geqslant 0$ is given by $\eta_t \in (\mathbb{N}_\mathfrak{s})^{\mathbb{Z}^d}$. The interpretation is that, at time t, site x contains $\eta_t(x)$ active particles if $\eta_t(x) \in \mathbb{N}_0$, or one sleeping particle if $\eta_t(x) = \mathfrak{s}$.

We turn $\mathbb{N}_\mathfrak{s}$ into an ordered set by letting $0 < \mathfrak{s} < 1 < 2 < \cdots$. We also let $|\mathfrak{s}| = 1$, so $|\eta_t(x)|$ counts the number of particles regardless of their state. To add a particle to a site, we define $\mathfrak{s} + 1 = 2$, which represents the $A + S \to 2A$ transition. We also define $1 \cdot \mathfrak{s} = \mathfrak{s}$ and $n \cdot \mathfrak{s} = n$ for $n \geqslant 2$, which represent the transitions $A \to S$ and $2A \to A + S \to 2A$, respectively.

The process has a parameter $0 < \lambda < \infty$ and evolves as follows. For each site x, a clock rings at rate $(1 + \lambda)\eta_t(x)\mathbb{1}_{\eta_t(x) \neq \mathfrak{s}}$. When this clock rings, the system goes through the transition $\eta \to \mathfrak{t}_{x\mathfrak{s}}\eta$ with probability $\frac{\lambda}{1+\lambda}$, otherwise $\eta \to \mathfrak{t}_{xy}\eta$ with probability $p(y - x)\frac{1}{1+\lambda}$. The transitions are given by

$$\mathfrak{t}_{x\mathfrak{s}}\eta(z) = \begin{cases} \eta(x) \cdot \mathfrak{s}, & z = x, \\ \eta(z), & z \neq x, \end{cases} \qquad \mathfrak{t}_{xy}\eta(z) = \begin{cases} \eta(x) - 1, & z = x, \\ \eta(y) + 1, & z = y, \\ \eta(z), & \text{otherwise}, \end{cases}$$

and only occur if $\eta(x) \geqslant 1$. The operator $\mathfrak{t}_{x\mathfrak{s}}$ represents a particle at x trying to fall asleep, which will effectively happen if there are no other particles present at x. Otherwise, by definition of $n \cdot \mathfrak{s}$ the configuration will not change. The operator $\mathfrak{t}_{xy}$ represents a particle jumping from x to y, where possible activation of a sleeping particle previously found at y is represented by the convention that $\mathfrak{s}+1=2$. The case $\lambda=\infty$ is left aside until §10.3.

Given a translation-ergodic distribution ν on $(\mathbb{N}_\mathfrak{s})^{\mathbb{Z}^d}$ with finite mean

$$\int |\eta(\mathbf{0})| \nu(\mathrm{d}\eta) < \infty,$$

there exists a process $(\eta_t)_{t\geqslant 0}$ with the above transition rates and such that η_0 has law ν, see §11 for a proof. We will use $\mathbf{P}^\nu$ to denote the underlying probability measure in a space where this process is defined.

2.2. Site-wise representation

The site-wise representation enables us to exploit an algorithmic approach to fixation. Due to particle exchangeability, this representation extracts precisely the part of the randomness that is relevant for the absorbing-state phase transition, focusing on the total number of jumps and leaving aside the order in which they take place. It is suitable for studying path traces, total occupation times, and final particle positions. But it precludes the analysis of quantities for which the order and instant of the jumps do matter, such as correlation functions, time needed for fixation, invariant measures, etc.

In this subsection we do not deal with a time evolution, and η denotes simply an element of $(\mathbb{N}_\mathfrak{s})^{\mathbb{Z}^d}$ on which one can perform certain operations. We say that site x is *unstable for the configuration* η if $\eta(x) \geqslant 1$. An unstable site x can *topple*, by applying $\mathfrak{t}_{xy}$ or $\mathfrak{t}_{x\mathfrak{s}}$ to η.

We consider a *field of instructions* $\mathcal{I} = (\mathfrak{t}^{x,j})_{x\in\mathbb{Z}^d, j\in\mathbb{N}}$. More precisely, at each site x there is a sequence or *stack* of instructions $\mathfrak{t}^{x,1}, \mathfrak{t}^{x,2}, \dots$, such that each one of the $\mathfrak{t}^{x,j}$ equals either $\mathfrak{t}_{x\mathfrak{s}}$ or $\mathfrak{t}_{xy}$ for some y. Later on we will choose $\mathcal{I}$ random, but for now $\mathcal{I}$ denotes a field that is fixed, and η denotes an arbitrary configuration.

We also introduce the *odometer field* $h = \big(h(x); x \in \mathbb{Z}^d\big)$ which counts the number of topplings already performed at each site, usually started from $h \equiv 0$. The toppling operation at x is defined by

$$\Phi_x(\eta, h) = \big(\mathfrak{t}^{x,h(x)+1}\eta, h + \delta_x\big).$$

We say that Φ_x is *legal for* (η, h) or simply *legal for* η if $\eta(x) \geqslant 1$. We may write $\Phi_x\eta$ as a short for $\Phi_x(\eta, 0)$.

Sometimes it will be convenient to topple a site x when it contains any particle at all, even if $\eta(x)=\mathfrak{s}$. To achieve that, we define on $\mathbb{N}_\mathfrak{s}$ the operations $\mathfrak{s}-1=0$ and $\mathfrak{s}\cdot\mathfrak{s}=\mathfrak{s}$. We say that toppling x is *acceptable* if $\eta(x) \geqslant \mathfrak{s}$, and we say that Φ_x is *acceptable for* (η, h) if $\eta(x) \geqslant \mathfrak{s}$. One should think of this operation as first forcing the particle to activate and then toppling the site. A legal toppling is also acceptable. We remark that it is not possible to define the operations $0 \cdot \mathfrak{s}$ and $0-1$ while preserving the local Abelian property stated below, so these two operations will not be considered as acceptable.

Sequences of topplings and local properties

Let $\alpha = (x_1, \ldots, x_k)$ denote a finite sequence of sites, and define the operator $\Phi_\alpha = \Phi_{x_k}\Phi_{x_{k-1}} \cdots \Phi_{x_1}$. We say that Φ_α is *legal* for (η, h) if Φ_{x_j} is legal for $\Phi_{(x_1,\ldots,x_{j-1})}(\eta, h)$ for each $j = 1, \ldots, k$. In this case we say that α is a *legal sequence* of topplings for (η, h). We define an *acceptable sequence* of topplings analogously. When $h \equiv 0$ we may write η instead of (η, h). Let $m_\alpha = \big(m_\alpha(x); x \in \mathbb{Z}^d\big)$ denote *the odometer of* α, given by $m_\alpha(x) = \sum_j \mathbb{1}_{x_j = x}$, so m_α is the field which counts how many times each site x appears in α. We write $m_\alpha \geqslant m_\beta$ if $m_\alpha(x) \geqslant m_\beta(x) \ \forall\ x$, and $\tilde\eta \geqslant \eta$ if $\tilde\eta(x) \geqslant \eta(x) \ \forall\ x$. We also write $(\tilde\eta, \tilde h) \geqslant (\eta, h)$ if $\tilde\eta \geqslant \eta$ and $\tilde h = h$.

We now state the four properties that make the ARW an Abelian model. The next three lemmas are based on these properties alone rather than on specific details of the ARW. Let x be a site in $\mathbb{Z}^d$ and $\eta, \tilde\eta$ be configurations.

Property 1 (Local Abelian property)**.** *If α and β are acceptable sequences of topplings for the configuration η, such that $m_\alpha = m_\beta$, then $\Phi_\alpha\eta = \Phi_\beta\eta$.*

Property 2 (Mass comes from outside)**.** *If α and β are acceptable sequences of topplings for η such that $m_\alpha(x) \leqslant m_\beta(x)$ and $m_\alpha(z) \geqslant m_\beta(z)$ for all $z \neq x$, then $\Phi_\alpha\eta(x) \geqslant \Phi_\beta\eta(x)$.*

Property 3 (Monotonicity of stability)**.** *If site x is unstable for the configuration η, and if $\tilde\eta(x) \geqslant \eta(x)$, then x is unstable for the configuration $\tilde\eta$.*

Property 4 (Monotonicity of topplings)**.** *If $\tilde\eta \geqslant \eta$ and Φ_x is legal for η, then Φ_x is legal for $\tilde\eta$ and $\Phi_x\tilde\eta \geqslant \Phi_x\eta$.*

Proof. The last two properties are immediate from the previous definitions. For convenience, define the operators $n\oplus = n + 1$ on $\mathbb{N}_\mathfrak{s}$, as well as $n\ominus = n - 1$ and $n\odot = n \cdot \mathfrak{s}$ on $\mathbb{N}_\mathfrak{s} \setminus \{0\}$. With this notation, whenever $n\ominus$ is acceptable (*i.e.* $n \neq 0$) we have $n \ominus \oplus = n \oplus \ominus$. Analogously, whenever $n\odot$ is acceptable (*i.e.* $n \neq 0$) we have $n \odot \oplus = n \oplus \odot$. Therefore, within any acceptable sequence of operations, replacing $\ominus\oplus$ by $\oplus\ominus$ and $\odot\oplus$ by $\oplus\odot$ yields an acceptable sequence with the same final outcome.

We first prove Property 1 as a warm up. Suppose that $m_\alpha = m_\beta$. Notice that $\Phi_\alpha\eta(x)$ is given by $\eta(x)$ followed by a sequence of $\oplus$, $\ominus$ and $\odot$'s. The number of times each operator appears is determined by $\mathcal{I}$ and m_α only, hence it is the same number when we write $\Phi_\beta\eta(x)$ as $\eta(x)$ followed by a sequence of $\oplus$, $\ominus$ and $\odot$'s. Their actual order depends on the sequence, but the internal order of the $\ominus$'s and $\odot$'s is determined by $(\mathfrak{t}^{x,j})_j$ and is thus the same for both $\Phi_\alpha\eta(x)$ and $\Phi_\beta\eta(x)$. As a consequence, we can apply the above identities to move the $\oplus$'s to the left, yielding then identical sequences for $\Phi_\alpha\eta(x)$ and $\Phi_\beta\eta(x)$, proving the claimed property.

For Property 2 we make a similar observation. Suppose $m_\alpha(x) \leqslant m_\beta(x)$ and $m_\alpha(z) \geqslant m_\beta(z)$ for $z \neq x$. Again $\Phi_\alpha\eta(x)$ is given by $\eta(x)$ followed by a number of $\oplus$, $\ominus$ and $\odot$'s. The number of times that operator $\oplus$ appears depends on $m_\alpha(z), z \neq x$, and is thus bigger in $\Phi_\alpha\eta(x)$ than in $\Phi_\beta\eta(x)$. The number of times that operators $\ominus$ and $\odot$ appear depend on $m_\alpha(x)$, and is thus smaller than in $\Phi_\beta\eta(x)$. Pushing the $\oplus$'s to the left as before, we get that $\Phi_\alpha\eta(x)$ is written in the same way as $\Phi_\beta\eta(x)$, perhaps with a few extra $\oplus$'s in the beginning, and a few missing $\ominus$ and $\odot$'s in the end, so $\Phi_\alpha\eta(x) \geqslant \Phi_\beta\eta(x)$. ☐

Stabilization via sequential topplings

Let V denote a finite subset of $\mathbb{Z}^d$. We say that a configuration η is *stable in V* if every $x \in V$ is stable for η. We say that α is *contained in V*, and write $\alpha \subseteq V$, if every x appearing in α is an element of V. We say that α *stabilizes η in V* if α is acceptable for η and $\Phi_\alpha\eta$ is stable in V.

Lemma 2.1. *If α is an acceptable sequence of topplings that stabilizes η in V, and $\beta \subseteq V$ is a legal sequence of topplings for η, then $m_\beta \leqslant m_\alpha$.*

Proof. Let $\beta \subseteq V$ be legal and $m_\alpha \ngeqslant m_\beta$. Write $\beta = (x_1, \dots, x_k)$ and $\beta^{(j)} = (x_1, \dots, x_j)$ for $j \leqslant k$. Let $\ell = \max\{j : m_{\beta^{(j)}} \leqslant m_\alpha\} < k$ and $y = x_{\ell+1} \in V$. Since β is legal, y is unstable in $\Phi_{\beta^{(\ell)}}\eta$. But $m_{\beta^{(\ell)}} \leqslant m_\alpha$ and $m_{\beta^{(\ell)}}(y) = m_\alpha(y)$. By the Properties 2 and 3, y is unstable for $\Phi_\alpha\eta$ and therefore α does not stabilize η in V. ☐

Let $V \subseteq \mathbb{Z}^d$ be a finite set. We define *the odometer of η in V* by

$$m_{V,\eta} = \sup_{\beta \subseteq V \text{ legal}} m_\beta \tag{2.2}$$

the supremum being taken over sequences β which are legal for η and contained in V. Lemma 2.1 says that

$$m_{V,\eta} \leqslant m_\alpha \tag{2.3}$$

for every acceptable sequence α stabilizing η in V. These two together provide very good sources of lower and upper bounds for $m_{V,\eta}$. Note that the sequence α need not be legal, nor contained in V. This allows us to choose a convenient sequence of topplings, even wake up some particles if we wish, if we are looking for upper bounds.

Lemma 2.4 (Abelian property)**.** *If α and β are both legal toppling sequences for η that are contained in V and stabilize η in V, then $m_\alpha = m_\beta = m_{V,\eta}$. In particular, $\Phi_\alpha\eta = \Phi_\beta\eta$.*

Proof. Applying (2.2), (2.3) and Lemma 2.1: $m_\beta \leqslant m_{V,\eta} \leqslant m_\alpha \leqslant m_\beta$. ☐

If there is an acceptable sequence α that stabilizes η in V, then there is a legal sequence β contained in V that also stabilizes η in V. Indeed, if one tries to perform legal topplings in V indefinitely, on the one hand one has to eventually stop since V is finite and there is a finite upper bound $\sum_x m_\alpha(x)$ for the total number of topplings, and on the other hand this procedure only stops if there are no more unstable sites in V.

Lemma 2.5 (Monotonicity)**.** *If $V \subseteq \tilde{V}$ and $\eta \leqslant \tilde{\eta}$, then $m_{V,\eta} \leqslant m_{\tilde{V},\tilde{\eta}}$.*

Proof. Let $\beta \subseteq V$ be legal for η. By successively applying Properties 3 and 4, β is also legal for $\tilde{\eta}$. Since $\beta \subseteq \tilde{V}$, the inequality follows directly from definition (2.2). $\square$

By monotonicity, the limit

$$m_\eta = \lim_{V \uparrow \mathbb{Z}^d} m_{V,\eta} \tag{2.6}$$

exists and does not depend on the particular sequence $V \uparrow \mathbb{Z}^d$. The limit is also given by the supremum of $m_{V,\eta}$ over finite V. A configuration η on $\mathbb{Z}^d$ is said to be *stabilizable* if $m_\eta(x) < \infty$ for every $x \in \mathbb{Z}^d$. In the next subsection we relate the above concepts to the question of fixation vs. activity for the continuous-time process.

2.3. Criteria for fixation and activity

Assume the initial configuration $\eta_0 \in (\mathbb{N}_\mathfrak{s})^{\mathbb{Z}^d}$ has a translation-ergodic distribution denoted by ν. Assume also that the support of the jump distribution $p(\cdot)$ generates the group $\mathbb{Z}^d$ and not a subgroup.

Consider now the field of instructions $\mathcal{I}$ being random, and distributed as follows. For each $x \in \mathbb{Z}^d$ and $j \in \mathbb{N}$, choose $\mathfrak{t}^{x,j}$ as $\mathfrak{t}_{xy}$ with probability $\frac{p(y-x)}{1+\lambda}$ or $\mathfrak{t}_{x\mathfrak{s}}$ with probability $\frac{\lambda}{1+\lambda}$, independently over x and j. Now sample η_0 according to ν and independently of $\mathcal{I}$, and let $\mathbb{P}^\nu$ denote the probability measure in a space where $\mathcal{I}$ and η_0 are defined.

Recall that $\mathbf{P}^\nu$ denotes the probability measure in whatever probability space where the process $(\eta_t)_{t\geqslant 0}$ described in §2.1 is defined. The following result is proved in §11.1.

Theorem 2.7. $\mathbf{P}^\nu\big($*fixation of* $(\eta_t)_{t\geqslant 0}\big) = \mathbb{P}^\nu(m_{\eta_0}(\mathbf{0}) < \infty) = 0$ *or* 1.

Combining Theorem 2.7 with (2.6) and Lemma 2.5 we get the following.

Corollary 2.8. *If the condition*

$$\sup_k \inf_{V \text{ finite}} \mathbb{P}^\nu(m_{V,\eta_0}(\mathbf{0}) \leqslant k) > 0 \tag{2.9}$$

is satisfied, the system a.s. fixates. If the condition

$$\inf_k \sup_{V \text{ finite}} \mathbb{P}^\nu(m_{V,\eta_0}(\mathbf{0}) \geqslant k) > 0 \tag{2.10}$$

is satisfied, the system a.s. stays active.

A typical usage of Condition (2.9) is to rely on (2.3) and use the odometer of an acceptable stabilizing sequence of topplings as a stochastic upper bound for m_{V,η_0}. Likewise, a typical usage of Condition (2.10) is to rely on (2.2) and use the odometer of a legal sequence of topplings as a stochastic lower bound for m_{V,η_0}.

In §10.2 we will prove the following sufficient condition for activity.

Theorem 2.11. *Let M_n count the number of particles that jump out of $V_n = \{-n,\dots,n\}^d$ when V_n is stabilized via legal topplings, so particles are ignored after leaving V_n. If η_0 is i.i.d. and the condition*

$$\limsup_n \frac{\mathbb{E}M_n}{|V_n|} > 0 \tag{2.12}$$

is satisfied, then the system a.s. stays active.

Similar to Condition (2.10), a typical usage of Condition (2.12) is to rely on the Abelian property and use the expected number of particles exiting V_n during some specific legal sequence of topplings $\beta \subseteq V_n$ as a stochastic lower bound for $\mathbb{E}M_n$.

An important property of the ARW is that ζ_c has a sharp definition.

Theorem 2.13 (Uniqueness of the critical density)**.** *Given the dimension d, jump distribution $p(\cdot)$, and sleep rate λ, there is a number ζ_c such that, for every translation-ergodic distribution ν supported on $(\mathbb{N}_0)^{\mathbb{Z}^d}$ with average density ζ, the ARW dynamics satisfies*

$$\mathbf{P}^{\nu}(\textit{system stays active}) = \begin{cases} 0, & \zeta < \zeta_c, \\ 1, & \zeta > \zeta_c. \end{cases}$$

This property will be proved in §8. Mathematically, it is useful because every statement about bounds for ζ_c can be proved assuming an i.i.d. initial state with whatever marginal distribution is more convenient. In fact, we can even take non-i.i.d. distributions if that helps.

We conclude with a monotonicity property.

Theorem 2.14. *The critical density ζ_c is non-decreasing in λ.*

The proof is given in §11.4.

Open Problem. Show that ζ_c is continuous and strictly increasing in λ.

2.4. Factors, ergodicity and the mass transport principle

We will not use these tools directly until §8. Let $f : \mathbb{Z}^d \times \mathbb{Z}^d \to \mathbb{R}_+$ be a translation-invariant random function, that is, $f(x, y; \omega) = f(\theta x, \theta y; \theta\omega)$ for every translation θ of $\mathbb{Z}^d$. Here ω can be any random process with translation-invariant distribution.

The *mass transport principle* is given by the identity

$$\mathbb{E}\big[\sum_y f(x,y)\big] = \mathbb{E}\big[\sum_y f(y,x)\big].$$

Informally, the mass transport principle says that on average the amount of mass transmitted from a site x is equal to the amount of mass entering x. It may seem like nothing but an obvious identity, but its strength lies in its versatility, since it holds for every translation-invariant function. The proof consists in re-indexing the sum and using translation invariance. See [LP16, Chapter 8] for applications and generalizations to other settings.

Suppose a field $\tilde{\omega} = g(\omega)$ is a translation-covariant function of ω. That is, for all ω for which $g(\omega)$ is defined, $g(\theta\omega)$ is also defined and $g(\theta\omega) = \theta g(\omega)$. Then we say that $\tilde{\omega}$ is a *factor* of ω. Note that factors inherit properties such as translation invariance and translation ergodicity.

Finally, we mention how pairs of ergodic and mixing fields behave together. Suppose ω^1 and ω^2 are independent. Suppose ω^1 is mixing (with respect to translations). If ω^2 is mixing, then the pair (ω^1, ω^2) is mixing. If ω^2 is ergodic, then (ω^1, ω^2) is ergodic. In particular, the pair $(\eta_0, \mathcal{I})$ introduced in the previous subsection is translation-ergodic. See [KL16, Thm. 2.25].

2.5. Frequently used notation

Let $\|\eta\|_V$ denote the total number of particles in the box V in the configuration η, given by $\|\eta\|_V = \sum_{x\in V} |\eta(x)|$.

The discrete ball of radius r centered at y is denoted B_r^y with $B_r = B_r^{\mathbf{0}}$. Any choice of norm to define the balls would work fine, but for the sake of aesthetics let us fix $B_r = \{-r, \dots, r\}^d$.

Normally, the letters η and ξ will denote configurations, ζ denotes density, while α and β denote sequences of topplings. The letter ν usually denotes a probability measure on $(\mathbb{N}_{\mathfrak{s}})^{\mathbb{Z}^d}$ or most often on $(\mathbb{N}_0)^{\mathbb{Z}^d}$. The letters L and r are usually used to denote sizes of some domains, whereas i, j, k are generally used to index sites, events, particles, etc. The letter n may have either of these uses.

The boxes $V_n \uparrow \mathbb{Z}^d$ are not necessarily the same in different parts of the text, and may in fact be random. Most of the time they equal B_n but we still use V_n to indicate that they are being used in the context of §2.3.

The letters M and N usually denote random variables defined by counting how many events of a family occur, such as how many particles do this and that, how many steps go wrong in a certain procedure, etc. We normally use $|\cdot|$ to denote the discrete volume of a finite subset of $\mathbb{Z}^d$. The symbol # rarely appears, technically it also denotes cardinality but we use it when such cardinality comes from counting random objects. When we have to name events, we will mostly use the letter $\mathcal{A}$ and maybe add indices.

Most of the time we deal with constructions based on independent random variables defined explicitly, and we use the letter $\mathbb{P}$ for probability (hence $\mathbb{E}$ for expectation). We use $\mathbf{P}$ to emphasize that certain statements refer to whatever probability space where the particle system $(\eta_t)_{t\geqslant 0}$ is defined. We often say "a.s." and "with high probability" without bothering about probability space formalities that are hardly relevant.

The letters κ, δ and ε usually denote a large fixed number, a small fixed number, and an arbitrarily small number. The letter K is an integer parameter in some constructions or toppling procedures, sometimes it is not fixed beforehand but instead made larger and larger throughout each proof.

3. Counting arguments

From the previous section, the Abelian property says that the odometer of a given configuration in a given region can be obtained through any legal sequence α of topplings. Note that such a sequence need not be given beforehand: it can in fact be constructed algorithmically as the configuration evolves.

This opens the door to proofs consisting of (i) the prescription of a *toppling procedure* followed by (ii) a probabilistic analysis to check Conditions (2.9), (2.10) or (2.12) in terms of such procedure. In this section we will see some simple instances of this approach being applied.

As a warm-up example for the use of Condition (2.10) we prove the following.

Theorem 3.1. *For $d = 1$ and initial state i.i.d. with mean $\zeta = 1$ and positive variance, the system a.s. stays active.*

We then show the simplest argument that uses a toppling procedure based on the Abelian property to check Conditions (2.9) or (2.10), and prove the following. A *directed walk* on $d = 1$ is the one with $p(+1) = 1$.

Theorem 3.2. *For $d = 1$ and directed walks, $\zeta_c = \frac{\lambda}{1+\lambda}$. If the initial state is i.i.d. with critical density $\zeta = \zeta_c$ and positive variance, then the system a.s. stays active.*

The proof of the previous theorem breaks down in case the walks are not totally directed, because even a small probability of jumping left would require some control on the interaction among particles which wander in the wrong direction. The following is proved via an argument that uses Condition (2.12) and Theorem 2.13.

Theorem 3.3. *For $d \geqslant 1$ and biased walks, $\zeta_c < 1$ for every $\lambda < \infty$ and $\zeta_c \to 0$ as $\lambda \to 0$.*

In proving the above we get a quantitative upper bound for ζ_c which gives the following corollary.

Corollary 3.4. *For directed walks in any dimension d, we have $\zeta_c \leqslant \frac{\lambda}{1+\lambda}$.*

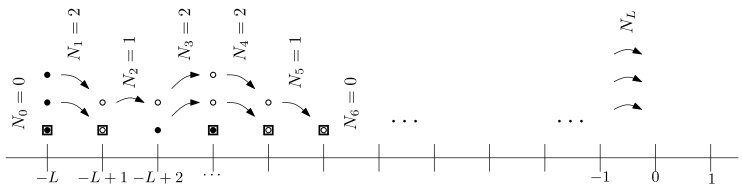

FIG 3.1. *Stabilizing $[-L, 0]$ from left to right. Disks represent particles initially present at each site, and circles represent particles that arrive from the left. Boxes indicate particles which fell asleep after they were left alone.*

Together with Theorem 7.1, this implies that $\zeta_c = \frac{\lambda}{1+\lambda}$ for directed walks.

We now proceed to the proofs.

Proof of Theorem 3.1. Let $\zeta = 1$. By the CLT, the probability that η_0 contains at least $L + 2\sqrt{L}$ particles in $V_L = [0, L]$ is at least $2\delta > 0$, for all large L. On this event, at least $2\sqrt{L}$ particles will visit either $x = \mathbf{0}$ or $x = L$ when we stabilize $[0, L]$.

Therefore, using a union bound, translation invariance and monotonicity,

$$2\delta \leqslant \mathbb{P}\big(m_{V_L,\eta_0}(\mathbf{0}) \geqslant \sqrt{L}\big) + \mathbb{P}\big(m_{V_L,\eta_0}(L) \geqslant \sqrt{L}\big) \leqslant 2\mathbb{P}\big(m_{\eta_0}(\mathbf{0}) \geqslant \sqrt{L}\big).$$

Since this is true for all large L, Condition (2.10) is satisfied and therefore the system a.s. stays active. $\square$

Proof of Theorem 3.2. Write $V_L = \{-L, \dots, L\}$. We will consider $m_{V_L,\eta_0}(\mathbf{0})$ and see in which cases it satisfies Condition (2.9) or (2.10) as $L \to \infty$.

We will describe a legal sequence of topplings that stabilizes V_L by exhausting one site after the other, from left to right. The sequence of topplings is itself random, since it is given by an algorithm which decides the next site to topple in terms of outcome of the previous topplings.

We start by toppling site $x = -L$ until each of the $\eta_0(-L)$ particles either moves to $x = -L + 1$ or falls asleep. All particles but the last one have to jump (possibly after a few frustrated attempts to sleep). The last one may sleep or jump. Let Y_0^L denote the indicator of the event that the last particle remains sleeping at $x = -L$. Conditioned on $\eta_0(-L)$, the distribution of Y_0^L is Bernoulli with parameter $\frac{\lambda}{1+\lambda}$ (in case $\eta_0(-L) = 0$, sample Y_0^L independently of everything else). The number of particles which jump from $x = -L$ to $x = -L+1$ is given by $N_1^L := [\eta_0(-L) - Y_0^L]^+$. See Figure 3.1.

Note that, after stabilizing $x = -L$, there are $N_1^L + \eta_0(-L+1)$ particles at $x = -L+1$. We now topple site $x = -L+1$ until it is stable, and denote by Y_1^L the indicator of the event that the last particle remains sleeping at $x = -L+1$. The number of particles which jump from $x = -L+1$ to $x = -L+2$ is given by $N_2^L := [N_1^L + \eta_0(-L+1) - Y_1^L]^+$. By iterating this procedure, the number N_{i+1}^L of particles which jump from $x = -L+i$ into $x = -L+i+1$ after stabilizing

$x = -L, -L+1, \ldots, -L+i$ is given by

$$N_{i+1}^L = [N_i^L + \eta_0(-L+i) - Y_i^L]^+,$$

where $N_0^L = 0$. The number of particles which jump into $\mathbf{0}$ while stabilizing V_L equals N_L^L. After that, we stabilize $x = \mathbf{0}$ and then $\{1, \ldots, L\}$, but the latter no longer affects $m_{V_L,\eta_0}(\mathbf{0})$.

Note that the sequence $(N_i^L)_{i=0,1,\ldots,L}$ is distributed as a random walk on $\mathbb{N}_0$, with independent increments distributed as $\eta_0(x) - Y$, reflected at 0. So the relevant quantity is

$$\mathbb{E}[\eta_0(-L+k) - Y_k^L] = \zeta - \frac{\lambda}{1+\lambda}.$$

If $\zeta < \frac{\lambda}{1+\lambda}$, the walk is positive recurrent. This implies that the family $\{N_L^L\}_{L\in\mathbb{N}}$ is stochastically bounded. Since $m_{V_L,\eta_0}(\mathbf{0})$ is conditionally distributed as a sum of $N_L^L + \eta_0(\mathbf{0})$ independent geometric variables with parameter $\frac{1}{1+\lambda}$, Condition (2.9) holds, and thus the system a.s. fixates. If $\zeta > \frac{\lambda}{1+\lambda}$, the walk is transient, and $\mathbb{P}\big(N_L^L \geqslant \frac{1}{2}(\zeta - \frac{\lambda}{1+\lambda})L\big)$ is large for large L, so Condition (2.10) holds and the system a.s. stays active. If $\zeta = \frac{\lambda}{1+\lambda}$, the walk is null-recurrent, so $N_L^L \to \infty$ in probability as $L \to \infty$, which again implies Condition (2.10) and therefore the system a.s. stays active. □

The proof of Theorem 3.3 uses a toppling procedure in which particles help each other progressing in the right direction. The idea is to keep the particles spread so that they form a safe zone where the last particle can transit without sleeping. This way the particle has a reasonable chance of catching up with the others without falling asleep outside the safe zone.

For $\boldsymbol{v} \in \mathbb{R}^d$, let $\mathcal{H}_{\boldsymbol{v}} = \{z \in \mathbb{Z}^d : z \cdot \boldsymbol{v} \leqslant 0\}$. Consider a continuous-time random walk which starts at $\mathbf{0}$, jumps at rate 1 according to $p(\cdot)$, and is killed at rate λ when it is at $\mathcal{H}_{\boldsymbol{v}}$. Denote by $F_{\boldsymbol{v}}(\lambda)$ the probability that this walk is ever killed. If $\boldsymbol{v}$ is such that $\boldsymbol{v} \cdot \sum_y y p(y) > 0$, we have that $F_{\boldsymbol{v}}(\lambda) < 1$ for every $\lambda < \infty$, and moreover $F_{\boldsymbol{v}}(\lambda) \to 0$ as $\lambda \to 0$. Theorem 3.3 thus follows from the following proposition.

Proposition 3.5. *For any dimension d and any $\boldsymbol{v} \in \mathbb{R}^d$, $\zeta_c \leqslant F_{\boldsymbol{v}}(\lambda)$.*

Proof. By Theorem 2.13 we can assume that the initial state is i.i.d. Bernoulli with parameter ζ. We will show that Condition (2.12) holds if $\zeta > F_{\boldsymbol{v}}(\lambda)$.

Let $V_n = B_n$ and consider a labeling $V_n = \{x_1, \ldots, x_r\}$, where $r = |V_n|$ and $x_1 \cdot \boldsymbol{v} \leqslant \cdots \leqslant x_r \cdot \boldsymbol{v}$. For $i = 0, \ldots, r$, let $A_i = \{x_{i+1}, \ldots, x_r\} \subseteq V_n$.

The toppling procedure consists of r steps. For $i = 1, \ldots, r$, Step i goes as follows. If there is a particle at x_i then we do a sequence of legal topplings, starting by a toppling at x_i, then at the location where this toppling pushed the particle to, and so on until this particle either (i) exits V_n, (ii) reaches an unoccupied site in A_i, or (iii) falls asleep in $V_n \setminus A_i$. If case (iii) occurs, we say that the particle is *left behind*.

By induction we can see that, for $i = 1, \dots, r$, after Step $i-1$, each site $x \in A_{i-1}$ (including x_i) is either unoccupied or contains exactly one active particle (recall that we are starting with Bernoulli). Hence, at most one particle is left behind at each step.

Since $V_n \setminus A_i \subseteq x_i + \mathcal{H}_{\boldsymbol{v}}$, at each step the probability of leaving a particle behind is bounded from above by $F_{\boldsymbol{v}}(\lambda)$. Moreover, during this procedure each particle either ends up exiting V_n or being left behind. Indeed, if a given step ends due to condition (ii), the corresponding particle will be moved again in a later step. Let N_n denote the total number of particles left behind, and notice that $M_n \geqslant \|\eta_0\|_{V_n} - N_n$. Therefore, $\mathbb{E}M_n \geqslant \zeta\,|V_n| - F_{\boldsymbol{v}}(\lambda)\,|V_n|$. This completes the proof. $\square$

Proof of Corollary 3.4. A jump distribution $p(\cdot)$ being directed means that, for some $\boldsymbol{v} \in \mathbb{R}^d$, we have $p(y) = 0$ for every y such that $\boldsymbol{v} \cdot y \leqslant 0$. In this case, a random walk that starts at $\mathbf{0}$ and jumps according to $p(\cdot)$ only visits $\mathcal{H}_{\boldsymbol{v}}$ once (at time zero), and so it follows from definition of $F_{\boldsymbol{v}}$ that $F_{\boldsymbol{v}}(\lambda) = \frac{\lambda}{1+\lambda}$. The result then follows from Proposition 3.5. $\square$

4. Exploring the instructions in advance

In this section we introduce a toppling procedure to prove phase transition in dimension $d = 1$ by using (2.3) to check Condition (2.9). Compared to the previous sections, it introduces a novel element which will be used in §6–§9: the heavy usage of acceptable topplings as a way to enforce activity and conveniently displace some particles away from their current position.

An element which is specific to the argument presented in this section and which will be absent in §7–§9 is that the toppling procedure used here is not "Markovian." By this we mean that in order to decide on the next topplings we use more information than just the outcome of the previous ones.

Theorem 4.1. *For $d = 1$, for every $\lambda > 0$, we have $\zeta_c \geqslant \frac{\lambda}{1+\lambda}$.*

The above lower bound is a particular case of Theorem 7.1. We find it instructive to understand the proof given in this section, in particular because it will be used in §6, it provides Corollary 4.4, and it can be adapted to study the ARW on other graphs such as regular trees. We now move to proving the theorem.

General strategy

The idea is to try and stabilize all the particles from η_0, following the instructions in $\mathcal{I}$, with the help of acceptable topplings.

After describing the procedure, we will show that, whenever it is successful, it implies that $m_{\eta_0}(\mathbf{0}) = 0$. We finally show that the procedure is successful with positive probability if $\zeta < \frac{\lambda}{1+\lambda}$, implying Condition (2.9).

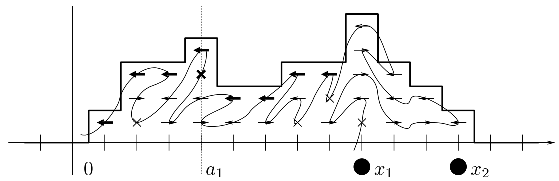

Fig 4.1. *First exploration path. It starts at position x_1 of the first particle and stops when it reaches the origin. The horizontal axis represents the lattice, and above each site x there is a stack of instructions $(\mathsf{t}^{x,j})_j$. The bold arrows indicate the last jump found at each site $x \in [1, x_1 - 1]$, and the bold cross indicates a sleep instruction found just before the last jump, this being the leftmost such cross, whose location defines a_1.*

Description of the toppling procedure

We will try to find a trap for one particle at a time. To find the trap, we launch an exploration that reveals some instructions in $\mathcal{I}$ until it identifies a suitable trap. To do that, the exploration follows the path that the particle would perform if we always toppled the site it occupies, and stop when the trap has been chosen. In the absence of a suitable trap, we declare the procedure to have failed.

An important issue is that some of the explored instructions are actually not going to be used by the corresponding particle which will instead remain sleeping at the trap. In particular, the instructions revealed in one step could interfere with the conditional distribution of subsequent steps. For this reason, we will call *corrupted* the sites where instructions have been revealed but not used.

If there is a particle at $\mathbf{0}$, we declare the procedure unsuccessful and stop. Otherwise, label the initial positions of the particles on $\mathbb{Z}$ by $\cdots \leqslant x_{-3} \leqslant x_{-2} \leqslant x_{-1} < 0 < x_1 \leqslant x_2 \leqslant x_3 \leqslant \cdots$.

Let $a_0 = 0$. We now describe the trapping of each particle. Suppose that the first $k-1$ traps have been successfully set up at positions $0 < a_1 < a_2 < \cdots < a_{k-2} < a_{k-1} \leqslant x_{k-1}$, and suppose also that the interval $[0, a_{k-1}]$ contains all traps and corrupted sites found in the previous steps.

We now launch an *exploration*. Starting at x_k, the explorer examines the next unexplored instruction at its current position, and moves to the site indicated by such instruction. If it is a sleep instruction, the explorer does not move. Repeat this indefinitely, and stop upon reaching a_{k-1}.

Next we *set up the trap*. Notice that, during the k-th exploration, a.s. each site is visited a finite number of times. Moreover, the explorer either stops at a_{k-1} or drifts to $+\infty$. In the latter case we simply take $a_k = a_{k-1}$. Suppose the former occurs. In this case, the explorer visits every site in $D_k = [a_{k-1} + 1, x_k - 1]$. Moreover, the last instruction explored at each site is a jump to the left, see Figures 4.1 and 4.2. Note that, for each $x \in D_k$, the second last instruction may or may not be a sleep instruction. We take a_k as the leftmost site at which the second last instruction explored was a sleep instruction, and call this second last instruction the k-th trap. If there is no such site in D_k, we declare the entire procedure unsuccessful and stop.

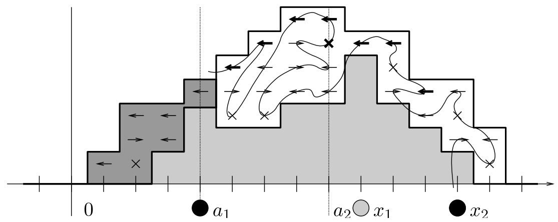

FIG 4.2. *Second exploration path. It starts at position x_2 of the second particle and stops when it reaches a_1. The regions in gray indicate the instructions already examined by the first explorer. The dark gray contains instructions examined but not used, whose locations determine the set of corrupted sites.*

So the trap is a sleep instruction found immediately before the last instruction, which is a jump to the left. Hence, we know that the exploration path has not been to the right of a_k after it revealed the instruction that we now declare as being the trap. Hence, *all the corrupted sites will be in* $[a_{k-1}+1, a_k]$, see Figure 4.2. Therefore, this process can be carried on indefinitely, as long as all the steps are successful.

If all the previous steps are successful, we repeat a symmetric construction on the negative half-line.

If the procedure is successful, then $m_{\eta_0}(\mathbf{0}) = 0$

We will show that, following the instructions of $\mathcal{I}$, η_0 is stabilized in $V_n = [x_{-n}, x_n]$ with finitely many acceptable topplings, without toppling $\mathbf{0}$.

Let us first stabilize the particle that starts at x_1. To this end, we successively topple the sites found by the first explorer, until it reaches the trap at a_1. At this moment the particle will fall asleep, and site a_1 will be stable. Note that these are acceptable topplings and the particle is following the same path that the explorer did, even if sometimes it will be sleeping.

We also know that, after the last visit to a_1, the explorer did not go further to the right, so when settling the first particle we use all the instructions examined so far, except some lying in $[a_0+1, a_1]$. Therefore, the second particle can be stabilized in the same way, as it will find the same instructions that determined the second exploration path.

Notice also that the first particle does not visit $\mathbf{0}$, and the second particle neither visits $\mathbf{0}$ nor a_1, thus it is settled without activating the first particle. Likewise, the k-th particle is settled at a_k, without ever visiting $\{\mathbf{0}, a_1, a_2, \dots, a_{k-1}\}$, for all $k = 1, \dots, n$. After settling the n first particles in $\mathbb{Z}_+$, we perform the analogous procedure for the first n particles in $\mathbb{Z}_-$.

This means that η_0 can be stabilized in V_n with finitely many acceptable topplings, not necessarily in V_n, and never toppling the origin. By (2.3), $m_{V_n,\eta_0}(\mathbf{0}) = 0$. Since it holds for all $n \in \mathbb{N}$ and $V_n \uparrow \mathbb{Z}$ as $n \to \infty$, this gives $m_{\eta_0}(\mathbf{0}) = 0$.

The procedure is successful with positive probability

For each site $x \in D_1$, the probability of finding a sleep instruction just before its last jump equals $\frac{\lambda}{1+\lambda}$, and this happens independently of the path and independently for each site. Thus, $a_1 - a_0$ is a geometric random variable with parameter $\frac{\lambda}{1+\lambda}$ truncated at $x_1 - a_0$.

Since no corrupted sites were left outside $[a_0+1, a_1]$ in the first exploration, the interdistance $a_2 - a_1$ is independent of a_1. Moreover, its distribution is also geometric with parameter $\frac{\lambda}{1+\lambda}$. The same holds for $a_3 - a_2$, $a_4 - a_3$, etc. By the law of large numbers, $a_n \sim \frac{1+\lambda}{\lambda} n$. On the other hand, $n \sim \zeta x_n$ (again by law of large numbers). Therefore, if $\zeta < \frac{\lambda}{1+\lambda}$, the event that $a_k < x_k$ for all k has positive probability. Finally, occurrence of this event implies that the procedure is successful, and the proof of Theorem 4.1 is finished.

Some immediate extensions and corollaries

In the above estimates, we obtained the following.

Remark 4.2. The distances $\{a_k - a_{k-1}\}_{k\geqslant 1}$ are i.i.d. geometric variables with parameter $\frac{\lambda}{1+\lambda}$.

Simple modifications in the choice of the trap would give the following.

Remark 4.3. Let $X = (X_n)_{n\geqslant 0}$ denote the Doob's h-transform of a walk starting at 0 and jumping according to $p(\cdot)$. That is, X is the walk conditioned to be positive for all $n > 0$. Let N be a geometric variable with parameter $\frac{\lambda}{1+\lambda}$ independent of X, and let $Z = \max\{X_1, \dots, X_N\}$. The distances $\{a_k - a_{k-1}\}_{k\geqslant 1}$ can be made i.i.d. and distributed as Z. With this modification, the above proof gives $\zeta_c \geqslant \frac{1}{\mathbb{E}Z}$ rather than just $\zeta_c \geqslant \frac{1}{\mathbb{E}N}$.

Corollary 4.4. *For $d = 1$, if the walks are not directed, we have $\zeta_c > \frac{\lambda}{1+\lambda}$.*

5. Particle flow between sparse sources

In this section we prove the following.

Theorem 5.1. *For $d = 1$ and symmetric walks, we have $\zeta_c \to 0$ as $\lambda \to 0$.*

The theorem is proved by checking Condition (2.12). Define the region $D_r = \{1, \dots, r-1\} \subseteq \mathbb{Z}^d$ with $d = 1$, and denote by η' the configuration obtained after legally stabilizing η_0 in D_r. The following proposition is more than enough to prove the theorem, and will also be used in §6.

Proposition 5.2. *For $\rho > 0$, there are $\lambda > 0$, $c > 0$ and $C < \infty$ such that*

$$\mathbb{P}\big(\|\eta'\|_{D_r} \geqslant \rho r \,\big|\, \eta_0\big) \leqslant C e^{-cr}$$

for all $r \in \mathbb{N}$ and $\eta_0 \in (\mathbb{N}_0)^{D_r}$.

Proof of Theorem 5.1. Let $\zeta > 0$. By Theorem 2.13, we can assume the initial state is i.i.d. with mean ζ. To check Condition (2.12), we can work with D_r instead of V_n. Taking $\rho < \zeta$ and λ as in Proposition 5.2, we get that $\limsup_r \frac{\mathbb{E}M_r}{r} = \limsup_r \frac{\mathbb{E}\|\eta_0\|_{D_r} - \mathbb{E}\|\eta'\|_{D_r}}{r} \geqslant \zeta - \rho > 0$. By Theorem 2.11, this implies a.s. activity, which means $\zeta_c \leqslant \zeta$, concluding the proof. $\square$

The remainder of this section is devoted to proving Proposition 5.2.

5.1. General framework

Fix some natural $K \geqslant 2\rho^{-1}$. Each site of the form iK for $i \in \mathbb{Z}$ is called a *source*. We can suppose $r = (n+1)K$ for some $n \in \mathbb{N}$. We can also suppose that $\eta_0 \in \{0,1\}^{V_r}$, otherwise we simply topple every site containing two or more particles until there is no longer such a site, and start from the resulting configuration.

We start presenting two auxiliary dynamics, the r-ARW and t-ARW. We then analyze how each block of the t-ARW behaves individually for all possible inputs. Later we consider global constraints given by *mass balance equations* for the flow of particles between sources, and see how the proposition follows from these constraints and an estimate involving a single source. We finally prove the single-source estimate.

Restricted ARW

We introduce a toppling procedure that gives a lower bound for the activity in the ARW. The *restricted ARW procedure* (r-ARW for short) goes as follows. We assign a different color to each source. Particles get the color of the last source they visited (those initially located between two sources can be assigned any of the two colors). When a particle finds a sleep instruction, we declare it to be *frozen*. When an unfrozen particle is found at the same site, it will unfreeze the frozen particle, but *only if they have the same color*. So a site might contain two frozen particles of different colors, or even a frozen particle of one color plus several unfrozen particles of another color. At each step, we topple a site in D_r containing unfrozen particles, and we do this until all particles in D_r are frozen.

We now discuss what this procedure says about the ARW on D_r.

A frozen particle may be active or sleeping in the ARW. But every unfrozen particle is also active, thus all topplings performed during this procedure are legal for the ARW. Hence, the configuration obtained at the end of this procedure gives an upper bound for $\|\eta'\|_{D_r}$.

Two-layer ARW

We now introduce the *two-layer ARW dynamics* (t-ARW for short), which is given by the ARW dynamics on the *two-layer graph* shown in Figure 5.1. Sites

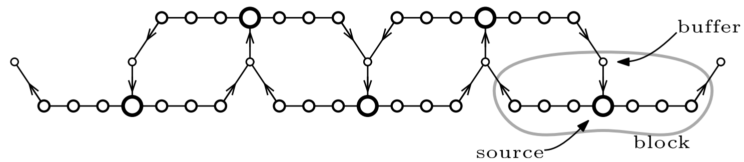

Fig 5.1. *The two-layer graph for the t-ARW with $K = 4$ and $n = 5$. It has 5 sources, at (horizontal) distance 4 from each other, represented by big sites. Between the two layers there are $5+2$ buffers, represented by tiny sites. The first and last buffers correspond to sites $x = \mathbf{0}$ and $x = (n+1)K = 24$. Directed edges indicate that buffers receive particles from neighboring blocks and release them to the corresponding sources.*

are grouped into *blocks* numbered $i = 1, \dots, n$. Each block contains $2K - 1$ regular sites, including a *source*, plus one *buffer* site. Particles never sleep at the buffers. For a configuration to be considered stable, all the regular sites must be either vacant or occupied by a sleeping particle, and each buffer that is linked to a source must be empty.

Remark that the t-ARW is equivalent to the r-ARW if we identify sites with the same horizontal coordinate and add up their particles. Given the horizontal position of a particle, it has two possible colors in the r-ARW, or equivalently it is on the upper or lower layer in the t-ARW. A sleeping particle in the t-ARW corresponds to a frozen in the r-ARW. The first and last buffers will never be toppled, they correspond to sites $\mathbf{0}$ and r.

The t-ARW is more convenient to work with because it is Abelian.

To define the initial configuration ξ on the two-layer graph, we distribute each particle in configuration $\eta_0 \in \{0,1\}^{D_r}$ to one of the two layers, according to its color. In the remainder of this section, we assume that the initial configuration ξ on the two-layer graph is fixed, and omit it in the notation. The estimates will hold uniformly with respect to ξ.

Single-block dynamics

Consider any sequence of legal topplings performed on the two-layer graph until the configuration is stable. By Abelianness of the t-ARW, the final configuration does not depend on the order of topplings. And by the above considerations, it provides a stochastic upper bound for $\|\eta'\|_{D_r}$.

Now notice that the interaction between a given block and the other ones is only through the input of particles from the buffer into its source and the output of particles from its leftmost and rightmost sites into a neighboring buffer. In order to analyze the relation between input and output, we fix a block i and study all possible values of inflow $m = 0, 1, 2, 3, \dots$.

For $m \in \mathbb{N}_0$, consider the stabilization of the configuration $\xi + m\delta_{iK}$ inside the i-th block. That is, m particles are added to the source iK and the configuration is toppled until it is stable in the block. By the Abelian property, it does not

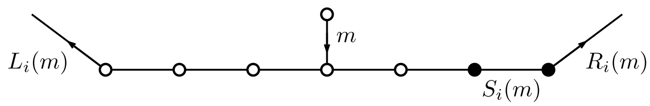

FIG 5.2. *Illustration of the dynamics inside a single block.*

matter whether the m particles are all added at the beginning or added one by one with some topplings being performed in between.

We now define random functions denoted by $L_i(\cdot)$, $R_i(\cdot)$ and $S_i(\cdot)$, illustrated in Figure 5.2. Let $L_i(m)$ count the number of particles that exit the block from the left when the configuration $\xi + m\delta_{iK}$ is stabilized in the i-th block, let $R_i(m)$ count the number of particles that exit the block from the right, $S_i(m)$ the number of particles sleeping in the block. Let $T_i(m) = L_i(m)+R_i(m)+S_i(m)$ be the total number of particles in $\xi_i + m\delta_{iK}$, where ξ_i is the restriction of ξ to the i-th block.

Remark the following about the change of these functions as m increases to $m' > m$. First, $T_i(m)$ equals m plus the number of particles initially found in the block, so it always increases by $m' - m$. The function S_i assumes values on $\{0, \dots, 2K-1\}$, so it can change by at most $2K-1$. The functions L_i and R_i are non-decreasing. It follows from these observations that

$$L_i(m') \leqslant L_i(m) + m' - m + 2K \tag{5.3}$$

for all $0 \leqslant m < m'$.

Note that these functions are random because they depend on the field of instructions $\mathcal{I}$, but they are independent across different blocks i.

Mass balance equations and proof of active phase

After globally stabilizing the two-layer graph, the odometer at the internal buffer sites will be given by $\boldsymbol{m}^* = (m_1^*, \dots, m_n^*)$. Writing $R_0 \equiv 0$ and $L_{n+1} \equiv 0$, the vector $\boldsymbol{m}^*$ satisfies the *mass balance equations*

$$m_i = R_{i-1}(m_{i-1}) + L_{i+1}(m_{i+1}) \qquad \text{for } i = 1, \dots, n. \tag{5.4}$$

We say that a non-negative integer vector $\boldsymbol{m}$ is *realizable* if it satisfies the above system of equations. A fixed deterministic $\boldsymbol{m}$ being realizable is a random event, because the functions $R_i(\cdot)$ and $L_i(\cdot)$ are random. Note that the random odometer $\boldsymbol{m}^*$ defined above is always realizable.

We now rewrite the above system as

$$L_i(m_i) = m_{i-1} - R_{i-2}(m_{i-2}) \tag{5.5}$$

for $i = 1, \dots, n+1$, where $R_{-1} \equiv 0$ and m_0 can be taken as $L_1(m_1)$.

For a non-negative vector $\boldsymbol{m}$, define

$$S(\boldsymbol{m}) = \sum_{i=1}^{n} S_i(m_i).$$

The total number of particles present in the blocks after global stabilization of the two-layer graph is given by

$$S^* = S(\boldsymbol{m}^*).$$

Lemma 5.6 (Single-block estimate)**.** *If λ is small enough depending on K, then for every $n \in \mathbb{N}$, $i = 1, \dots, n$, and initial configuration ξ, we have*

$$\sup_{\ell} \mathbb{E}\left[\sum_{m} e^{S_i(m)} \cdot \mathbb{1}_{\{L_i(m)=\ell\}}\right] \leqslant 3.$$

Proof of Proposition 5.2. Recall that $K \geqslant 2\rho^{-1}$ is fixed, and choose λ according to the previous lemma. Also recall that $r = (n+1)K$. Finally, recall that (5.5) is satisfied for $i = 1, \dots, n$ when $\boldsymbol{m} = \boldsymbol{m}^*$. Therefore,

$$\begin{aligned}
\mathbb{P}(S^* \geqslant \rho r) &= \sum_{\boldsymbol{m}} \mathbb{P}(S(\boldsymbol{m}) \geqslant \rho r,\ \boldsymbol{m}^* = \boldsymbol{m}) \\
&\leqslant \sum_{\boldsymbol{m}} \mathbb{P}\big(e^{S(\boldsymbol{m})} \geqslant e^{\rho r},\ \boldsymbol{m} \text{ is realizable}\big) \\
&\leqslant e^{-\rho r} \sum_{\boldsymbol{m}} \mathbb{E}\Big[e^{S(\boldsymbol{m})} \mathbb{1}_{\{\boldsymbol{m} \text{ is realizable}\}}\Big] \\
&= e^{-\rho r} \sum_{m_0} \mathbb{E}\Big[\sum_{m_1} \cdots \sum_{m_n} \prod_{i=1}^{n} e^{S_i(m_i)} \mathbb{1}_{\{L_i(m_i) = m_{i-1} - R_{i-2}(m_{i-2})\}}\Big].
\end{aligned}$$

Now notice that the random functions $S_n(\cdot)$ and $L_n(\cdot)$ are independent of the family of random functions $\mathcal{X}_n = (L_j(\cdot), R_j(\cdot), S_j(\cdot))_{j=1,\dots n-1}$. Hence,

$$\begin{aligned}
&\mathbb{E}\Big[\sum_{m_1} \cdots \sum_{m_n} \prod_{i=1}^{n} e^{S_i(m_i)} \mathbb{1}_{\{L_i(m_i) = m_{i-1} - R_{i-2}(m_{i-2})\}} \Big| \mathcal{X}_n\Big] = \\
&\quad = \sum_{m_1} \cdots \sum_{m_{n-1}} \mathbb{E}\Big[\sum_{m_n} e^{S_n(m_n)} \mathbb{1}_{\{L_n(m_n) = m_{n-1} - R_{n-2}(m_{n-2})\}} \Big| \mathcal{X}_n\Big] \times \\
&\qquad\qquad \times \prod_{i=1}^{n-1} \Big[e^{S_i(m_i)} \mathbb{1}_{\{L_i(m_i) = m_{i-1} - R_{i-2}(m_{i-2})\}}\Big].
\end{aligned}$$

The last conditional expectation is bounded from above by

$$\sup_{\ell} \mathbb{E}\Big[\sum_{m_n} e^{S_n(m_n)} \mathbb{1}_{\{L_n(m_n)=\ell\}}\Big].$$

Hence, regarding the previous chain of inequalities, we can pull the last term in the product out of the expectation. The same reasoning works for $i = n-1, n-2, \dots$, giving

$$\begin{aligned}\mathbb{P}(S^* \geqslant \rho r) &\leqslant e^{-\rho r} \sum_{m_0} \mathbb{E}\Big[\sum_{m_1} \cdots \sum_{m_n} \prod_{i=1}^{n} e^{S_i(m_i)} \mathbb{1}_{\{L_i(m_i)=m_{i-1}-R_{i-2}(m_{i-2})\}}\Big] \\ &\leqslant e^{-\rho r} \sum_{m_0} \prod_{i=1}^{n} \sup_{\ell} \mathbb{E}\Big[\sum_{m_i} e^{S_i(m_i)} \mathbb{1}_{\{L_i(m_i)=\ell\}}\Big] \\ &\leqslant r e^{-\rho r} 3^n \leqslant r e^{-cr}.\end{aligned}$$

Since $\|\eta'\|_{D_r}$ is stochastically dominated by S^*, this concludes the proof. $\square$

5.2. Single-block estimate

We now prove Lemma 5.6. Take $M_0 \in \mathbb{N}$ so that

$$e^{2K}(\tfrac{5}{9})^j \leqslant (\tfrac{3}{5})^j \quad \text{for all } j \geqslant M_0 \tag{5.7}$$

and take $\varepsilon > 0$ so that

$$6\varepsilon e^{2K} + (\tfrac{5}{9})^j \leqslant \tfrac{6}{5}(\tfrac{3}{5})^j \quad \text{for all } j \leqslant M_0. \tag{5.8}$$

Now take $M_1 \in \mathbb{N}$ such that the probability of getting at least one tail out of M_1 fair coin tosses is at least $1-\varepsilon$, and take $M_2 > 2M_1$ such that the probability of getting at least $M_1 + 2K + 2$ tails out of M_2 fair coin tosses is at least $1-\varepsilon$. Finally, take λ small enough so that, with probability at least $1-\varepsilon$, M_2 independent random walks all reach distance $2K$ before sleeping.

We now show how to put all these pieces together.

Let $\ell \in \mathbb{N}_0$. For the process

$$\big(L_i(m), R_i(m), S_i(m)\big)_{m=0,1,2,\dots},$$

define the following stopping times:

$$\begin{aligned}\mathcal{T}_1 &= \min\{m : L_i(m) \geqslant \ell - M_1 - 2K - 2\}, \\ \mathcal{T}_2 &= \mathcal{T}_1 + M_1, \\ \mathcal{T}_3 &= \min\{m : L_i(m) \geqslant \ell\}, \\ \mathcal{T}_4 &= \min\{m : L_i(m) \geqslant \ell + 1\}.\end{aligned}$$

From the above definitions and (5.3), we have

$$\mathcal{T}_1 < \mathcal{T}_2 < \mathcal{T}_3 \leqslant \mathcal{T}_4.$$

Consider the event

$$\mathcal{A} = \big\{S_i(m) > 0 \text{ for some } m \in [\mathcal{T}_3, \mathcal{T}_4)\big\}.$$

Since $S_i(m) < 2K$ for every m, by definition of $\mathcal{A}$ we have,

$$\sum_{j\geqslant 0} e^{S_i(\mathcal{T}_3+j)} \cdot \mathbb{1}_{\{\mathcal{T}_4-\mathcal{T}_3>j\}} \leqslant \sum_{j\geqslant 0} e^{2K\mathbb{1}_{\mathcal{A}}} \cdot \mathbb{1}_{\{\mathcal{T}_4-\mathcal{T}_3>j\}}. \tag{5.9}$$

We will show that

$$\mathbb{P}(\mathcal{T}_4 - \mathcal{T}_3 > j) \leqslant (\tfrac{5}{9})^j \quad \text{for all } j \geqslant 0 \tag{5.10}$$

and

$$\mathbb{P}(\mathcal{A}) \leqslant 6\varepsilon. \tag{5.11}$$

Let us first see how these imply the lemma. Using (5.9), (5.10), (5.11), (5.8), and (5.7),

$$\begin{aligned}
\mathbb{E}\left[\sum_m e^{S_i(m)} \cdot \mathbb{1}_{\{L_i(m)=\ell\}}\right] &\leqslant \sum_{j\geqslant 0} \mathbb{E}\left[e^{2K\mathbb{1}_{\mathcal{A}}} \cdot \mathbb{1}_{\{\mathcal{T}_4-\mathcal{T}_3>j\}}\right] \\
&\leqslant \sum_{j=0}^{M_0} \left[\mathbb{P}(\mathcal{T}_4 - \mathcal{T}_3 > j) + e^{2K}\mathbb{P}(\mathcal{A})\right] + \\
&\qquad\qquad + \sum_{j>M_0} \left[e^{2K}\mathbb{P}(\mathcal{T}_4 - \mathcal{T}_3 > j)\right] \\
&\leqslant \sum_{j=0}^{M_0} \left[6\varepsilon e^{2K} + (\tfrac{5}{9})^j\right] + \sum_{j>M_0} \left[e^{2K}(\tfrac{5}{9})^j\right] \\
&\leqslant \sum_{j\geqslant 0} \tfrac{6}{5}(\tfrac{3}{5})^j = 3.
\end{aligned}$$

So let us prove (5.10). Given that $L_i(m) = \ell$, and given all the information revealed when stabilizing $\xi + m\delta_{iK}$ in the i-th block, we claim that the conditional probability of $L_i(m+1) > \ell$ is at least $\frac{1-\varepsilon}{2}$, which is greater than $\frac{4}{9}$. Estimate (5.10) then follows by successive conditioning. Now to see why the claim holds true, consider the following toppling procedure for $\xi + (m+1)\delta_{iK}$. We keep the $(m+1)$-st particle in the buffer and let $\xi + m\delta_{iK}$ stabilize. We then add said particle to the source, and move it until it either finds a sleep instruction or exits the block. By the choice of λ, the probability of exiting before finding a sleep instruction has probability at least $1 - \varepsilon$, and by symmetry the probability of leaving the block from the left will be half of it. After that, we stabilize the remaining active particles, if any.

To finish the proof of Lemma 5.6, it remains to show (5.11).

We consider only $\ell \geqslant M_1 + 2K + 2$. The case of smaller ℓ uses similar but simpler arguments, and will be omitted. We will indicate with a "(*)" some events which occur with conditional probability at least $1-\varepsilon$, as a consequence of our choices of M_1, M_2 and λ. This chain of events all together imply $\mathcal{A}^c$.

Denote by $(\xi + \mathcal{T}_1\delta_{iK})'$ the configuration obtained after stabilizing the configuration $\xi + \mathcal{T}_1\delta_{iK}$ in the i-th block. We now stabilize the configuration $(\xi + \mathcal{T}_1\delta_{iK})' + M_1\delta_{iK}$ obtained by adding M_1 new active particles at the source iK. Let each of these M_1 new active particles move until it exits the block or finds a sleep instruction. Suppose none of them finds a sleep instruction before exiting (*). Suppose at least one of them exits the block from the left (*) and at least one from the right (*).

In this case, all the sleeping particles present in $(\xi + \mathcal{T}_1\delta_{iK})'$ have been activated. One by one, let each particle move until it exits the block or finds a sleep instruction. Suppose none of them finds a sleep instruction (*). When all the above events occur, $S_i(\mathcal{T}_2) = 0$.

So we resume from $\mathcal{T}_2$ and assuming $S_i(\mathcal{T}_2) = 0$. Suppose the next M_2 particles to be added to the source iK exit the i-th block before finding a sleep instruction (*). Given this event, the conditional probability that at least $M_1 + 2K + 2$ of them exit from the left is also at least $1 - \varepsilon$.

Suppose the latter event also occurs. Then $S_i(m) = 0$ for $m = \mathcal{T}_2, \mathcal{T}_2 + 1, \dots, \mathcal{T}_2 + M_2$, and moreover $L_i(\mathcal{T}_2 + M_2) \geqslant L_i(\mathcal{T}_2) + M_1 + 2K + 2$. It remains to check that these two events imply $\mathcal{A}^c$. The last inequality implies that $L_i(\mathcal{T}_2 + M_2) \geqslant \ell + 1$, so $\mathcal{T}_4 \leqslant \mathcal{T}_2 + M_2$. On the other hand, as noted after these stopping times were defined, it follows from (5.3) that $\mathcal{T}_2 < \mathcal{T}_3$. Hence, when the above events occur we have $S_i(m) = 0$ for $\mathcal{T}_3 \leqslant m \leqslant \mathcal{T}_4$ and event $\mathcal{A}$ cannot occur, so its probability is at most 6ε.

This concludes the proof of (5.11) and hence that of Proposition 5.2.

6. Fast and slow phases for finite systems

Consider the ARW model on the ring $\mathbb{Z}_n = \mathbb{Z}/n\mathbb{Z}$ instead of $\mathbb{Z}^d$. When a particle jumps, it chooses one of the two nearest neighbors according to a fair coin. The initial configuration η_0 is taken as i.i.d. Poisson with parameter ζ and all particles starting active. Let $\mathcal{T} = \sum_x m_{\mathbb{Z}_n,\eta_0}(x)$ denote the total number of topplings performed during stabilization of η_0.

Theorem 6.1. *Let $0 < \zeta < 1$. If λ is small enough, there exist $\delta > 0$ and $\delta' > 0$ such that, for all large n,*

$$\mathbb{P}(\mathcal{T} \geqslant e^{\delta' n}) \geqslant 1 - e^{-\delta n}. \tag{6.2}$$

If λ is large enough, there exist $\delta > 0$ and $\kappa < \infty$ such that, for all large n,

$$\mathbb{P}(\mathcal{T} \leqslant \kappa n \log^2 n) \geqslant 1 - n^{-\delta}. \tag{6.3}$$

Open Problem. Show that (6.2) or (6.3) must hold for every (ζ, λ) outside the critical curve of Figure 1.1.

Open Problem. Improve the estimate (6.3) to one without the $\log^2 n$ term.

Open Problem. Show that the family $m_{\mathbb{Z}_n,\eta_0}(\mathbf{0})$ is tight on the fast phase.

Open Problem. Show similar behavior in case of biased jumps.

Open Problem. Study fast to slow transition in higher dimensions.

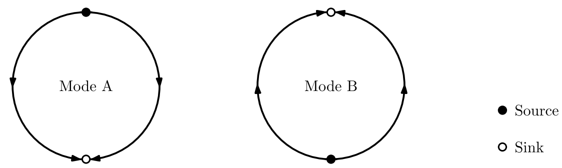

Fig 6.1. *Two alternating modes of the toppling procedure.*

6.1. Slow phase

We first prove (6.2). For simplicity we consider the model on $\mathbb{Z}_{2n}$ (*i.e.* we work with an even number of sites). The proof is based on a cyclic toppling procedure described as follows. At all times, particles will be declared to be *alive* or *steady*. Initially, declare all particles to be *alive*.

The toppling procedure consists in alternating between two modes, as illustrated in Figure 6.1. In Mode A, site $x = 0$ is called *source* and $x = n$ is called *sink*. In Mode B the roles are reversed. During each mode, we topple all sites containing active alive particles, except the sink. We keep toppling until there are no more such sites. At the beginning of each mode, we declare all active particles to be alive, and all sleeping particles to be steady. Note that particles that are declared steady at the beginning of a mode will not move during that mode, but they may be activated in the meantime, in which case they will be alive in the next mode. Note also that in this process we only use legal topplings.

Fix some $\rho < \frac{\zeta}{5}$, so the expected number of particles in $\mathbb{Z}_{2n}$ equals $10\rho n$. If the initial configuration has fewer than $9\rho n$ particles, we stop the procedure (before it even starts). Otherwise we run Mode A.

At the end of Mode A, we stop the procedure if (i) more than $2\rho n$ alive particles are sleeping, or (ii) some site of $\mathbb{Z}_{2n}$ is not visited by any of the alive particles during this mode. Otherwise we switch to Mode B. At the end of Mode B, we stop the procedure if at least one of the above conditions is met, otherwise we switch to Mode A. We continue this indefinitely until a condition to stop is met. The first run of Mode A is rather special: it does not start with many particles at the source, but it starts with all particles alive. For this first run we do not check condition (ii).

We now argue that, at the beginning of each run of Mode A or Mode B, at least $7\rho n$ particles will be active and alive, of which at least $5\rho n$ will be at the source. Indeed, at the end of the previous mode, no more than $2\rho n$ of the alive particles were found sleeping, hence at least $5\rho n$ of them finished at the sink. Moreover, all sites of $\mathbb{Z}_{2n}$ have been visited, hence all the previously steady particles were activated, implying that at least $7\rho n$ particles were active when the previous mode ended.

We are ready to make the probabilistic estimates leading to (6.2).

A crude bound on the number of topplings is the following: at each mode we perform at least one toppling. Hence, to get (6.2) we only need to prove that the probability (conditioned on the past) of stopping the procedure at the end of each mode is less than e^{-cn} for some c.

Moreover, it is enough to prove this estimate for Mode B. Indeed, Mode A is analogous (and its first run is simpler), and the probability of starting with fewer than $9\rho n$ particles decays exponentially fast by Chernoff bound for sums of i.i.d. Poisson variables.

Condition (i) is the most laborious part. But all the work has been done in §5, and at this point it suffices to choose λ according to Proposition 5.2. Indeed, since the sink is never toppled, this process is the same as the process on $D_{2n} = \{1, \dots, 2n-1\}$ and Proposition 5.2 says that condition (i) is met with exponentially small probability.

For condition (ii), label $5\rho n$ alive particles initially found at $x = 0$ and see whether each one of them ends up sleeping, reaches $x = n$ in the clockwise sense, or reaches $x = n$ in the counter-clockwise sense. For the particles which end up sleeping, use extra randomness to complete their tentative path which is stopped only upon reaching $x = n$. By symmetry, each labeled particle tentatively reaches $x = n$ from either direction with probability $\frac{1}{2}$, independently of the other labeled particles. Now if condition (i) is not met, it implies that at least $3\rho n$ of these $5\rho n$ particles perform a random walk which does make it to $x = n$, so they will perform all their tentative paths. On the other hand, the probability that $3\rho n$ or more out of the $5\rho n$ independent tentative paths reach n through the same direction is exponentially small by Chernoff bound for sums of i.i.d. Bernoulli variables. This gives the estimate for the probability of stopping the procedure at each mode, which concludes the proof of (6.2).

6.2. Fast phase

We now prove (6.3). Let

$$\frac{\lambda}{1+\lambda} > \zeta' > \zeta. \tag{6.4}$$

Letter K denotes a constant which will be chosen large during the proof. Letters C and c denote constants whose precise values are not crucial for the arguments and may change from line to line. They depend on ζ, ζ' and λ but neither on K nor on n. Constants δ, δ' and κ also depend on the choice of K.

Let $r = \lfloor K \log n \rfloor$ and split the ring $\mathbb{Z}_n$ into *arcs* containing $2r + 1$ sites (we can leave a spare site between some pairs of consecutive arcs to keep n sites in total). The middle point of each arc will be called the *source*.

We consider a toppling procedure split into two stages. In Stage 1, we force each particle to move by means of acceptable topplings, until it reaches one of the sources. If a source gets more than $2\zeta' r$ particles, we declare the procedure to have failed and stop. In Stage 2, we let the configuration stabilize normally, by using only legal topplings. If a particle leaves the corresponding arc during this stage, we declare the procedure to have failed. Otherwise it is successful.

We will show that the probability of failure is bounded by $n^{-\delta}$ for large n. Moreover, with exponentially high probability, on the event that the procedure succeeds, the total odometer will be bounded by Cnr^2, which proves (6.3).

Estimate for the total odometer

We start by stating a crude bound on exit times for random walks. Consider a collection of n or fewer independent lazy walks, each one stopped upon reaching distance $3r$. The number of steps each walk makes is stochastically dominated by Cr^2X, where X is a geometric random variable with mean 2. Indeed, by taking C large, the probability of exiting $[-3r, 3r]$ within Cr^2 steps is larger than $\frac{1}{2}$ regardless of the starting point, so if it fails to exit in Cr^2 steps we can start over. Hence, by enlarging C and using Chernoff bound for sums of i.i.d. geometric variables, the probability that the overall total number of jumps exceeds Cnr^2 is less than e^{-cn}.

We can assume that $\|\eta_0\| < n$, otherwise the procedure will necessarily fail at Stage 1. Now during this stage, each particle performs a lazy random walk stopped upon reaching a source. The laziness comes from the fact that sleep instructions keep the particle at the same site. So by the previous paragraph, the total odometer produced during this stage is bounded by Cnr^2 with probability at least $1 - e^{-cn}$.

During Stage 2, each particle performs a random path which is stopped earlier than reaching distance r, unless the procedure fails. In order to compare the path performed by the particles with a collection of independent lazy walks, we extend the paths by using extra randomness, so that the resulting collection of paths is distributed as a collection of independent lazy symmetric walks. So the same argument applies to the total odometer obtained at this stage, concluding the estimates on the total odometer.

Estimates for Stage 1

Let x be a given source. Particles which reach x during Stage 1 must have started at sites between the sources immediately to the left and right of x. Conditioning on the initial configuration, for each site y, a given particle starting at y will, independently of other particles, reach x before another source with a probability p_y which is proportional to the distance between its initial location and the other source.

Now the i.i.d. Poisson assumption about the initial configuration η_0 simplifies the argument, in that the number of particles reaching x will be a Poisson variable with parameter $\sum_y \zeta p_y$, independently over y. Since the probabilities p_y increase linearly from 0 to 1 and the nearest sources are at distance less than $2r + 2$, this parameter is less than $2\zeta r + 2$. So by Chernoff bound for a Poisson variable with large parameter, the probability that x gets more than $2\zeta' r$ particles during Stage 1 is bounded by e^{-cr} for large r. Summing over all sources, the probability that some source gets more than $2\zeta' r$ particles is bounded by $\frac{n}{2r+1} e^{-cr}$ for large n. By choosing K large enough, this is less than $n^{-\delta}$ for all large n.

Estimates for Stage 2

The central statement about this stage is the following.

Proposition 6.5. *For $\zeta' < \frac{\lambda}{1+\lambda}$, consider the symmetric ARW on $\mathbb{Z}$ starting with $m \leqslant 2\zeta' r$ particles at $\mathbf{0}$ and no other particles elsewhere. Then the probability that some particle ever leaves the interval $[-r, r]$ is less than e^{-cr} for all large r.*

Note that, unless the procedure failed during Stage 1, there are fewer than $2\zeta' r$ particles at each source. Therefore, applying the proposition and summing over all arcs we can bound the probability of failure during Stage 2 by $\frac{n}{2r+1} e^{-cr}$ for large n, with a possibly different constant c. By further enlarging K, again we can find δ such that this bound stays below $n^{-\delta}$ for all large n.

It remains to prove the proposition.

First, recall the toppling procedure presented in §4. Translating that procedure by $-r$, we initially have a barrier at $a_0 = -r$, and a new barrier $a_j > a_{j-1}$ is added when an exploration reaches the previous barrier a_{j-1}. A modification that we make in order to prove Proposition 6.5 is to confine the explorations from both directions. Namely, there is a second barrier initially at $b_0 = +r$, and a new barrier $b_i < b_{i-1}$ is added each time the exploration reaches the previous barrier b_{i-1}. We can carry this exploration m times, as long as the condition $a_j < 0 < b_i$ is preserved. If some of the m explorations fails to find a suitable trap, we declare the procedure to have failed and stop, otherwise it is successful.

We now show that this procedure is successful with high probability.

Recalling Remark 4.2, the distances $a_j - a_{j-1}$ are i.i.d. until the moment j_* when the procedure fails due to the explorer hitting a_{j_*-1} but being unable to find a suitable trap. Moreover, even the failure to find the trap can be coupled to the event that $a_{j_*-1} + Y \geqslant x_0$ for a geometric random variable Y, where $x_0 = 0$ is the starting position of the explorer. For convenience, after the procedure is finished we continue sampling independent geometric variables just so that we get two independent i.i.d. sequences $(a_j - a_{j-1})_{j\in\mathbb{N}}$ and $(b_{i-1} - b_i)_{i\in\mathbb{N}}$.

Define $J(0) = I(0) = 0$. When the $(k+1)$-th explorer starts, the barriers are at $a_{J(k)}$ and $b_{I(k)}$. If the explorer hits $a_{J(k)}$ before $b_{I(k)}$, we set $J(k+1) = J(k)+1$ and $I(k+1) = I(k)$, otherwise we set $I(i+1) = I(k) + 1$ and $J(k+1) = J(k)$. This way $J(k) + I(k) = k$ throughout the whole procedure. The goal then is to show that

$$\mathbb{P}(k_* \leqslant m) \leqslant e^{-cr}, \tag{6.6}$$

where $k_* = \min\{k : a_{J(k)} \geqslant 0 \text{ or } b_{I(k)} \leqslant 0\}$.

If we were assuming $\zeta < \frac{1}{2}$, we could choose $\frac{\lambda}{1+\lambda} > 2\zeta'$ in (6.4). In this case, the conclusion of Proposition 6.5 would follow immediately from the analysis made in §4, with the use of Chernoff bound for sums of i.i.d. geometric variables.

However, we want to show a stronger result that extends to arbitrary $\zeta < 1$, and this requires a more delicate argument. We would like to argue that about half of the particles will go to each direction. But there is an inconvenient reinforcement here: if many explorations have chosen left, it increases the chances that the next explorations will make the same choice. Fortunately, this effect is not strong enough to produce a macroscopic unbalance between I and J, as shown below.

The law of $(-a_{J(k)}, b_{I(k)})_{k=0,\dots,k_*}$ can be described as follows. Consider an urn containing $X_0 = r$ purple and $Z_0 = r$ yellow balls. At each turn k, a ball is sampled uniformly from the urn. The sampled ball is returned and a random number Y_k of balls with opposite color are destroyed, where Y_k has geometric distribution with parameter $\zeta'' = \frac{\lambda}{1+\lambda} > \zeta'$. That is,

$$(X_k, Z_k) - (X_{k-1}, Z_{k-1}) = \begin{cases} (0, -Y_k), & \text{with probability } \frac{X_k}{X_k+Z_k}, \\ (-Y_k, 0), & \text{with probability } \frac{Z_k}{X_k+Z_k}. \end{cases}$$

Finally, the urn process is stopped when one of the colors disappears, that is at step $k_* = \min\{k : X_k \leqslant 0 \text{ or } Z_k \leqslant 0\}$.

Lemma 6.7. *For $0 < \zeta' < \zeta'' < 1$, we have $\mathbb{P}(k_* \leqslant 2\zeta' r) \leqslant e^{-cr}$.*

The proof is given in Section 3.4 of [BGHR19] using a decoupling of the urn process as in [KV03] and applying estimates for sums of non-i.i.d. geometric and exponential random variables from [Jan18].

Lemma 6.7 implies (6.6), which in turn implies Proposition 6.5.

7. Weak and strong stabilization

In this section we prove the following for the ARW on $\mathbb{Z}^d$.

Theorem 7.1. *For any jump distribution in any dimension, $\zeta_c \geqslant \frac{\lambda}{1+\lambda}$.*

Theorem 7.2. *If $d > 2$, then $\zeta_c < 1$ for every $\lambda < \infty$ and $\zeta_c \to 0$ as $\lambda \to 0$.*

Open Problem. Prove a similar statement for unbiased walks on $\mathbb{Z}^2$.

Open Problem. At least prove that $\zeta_c < 1$ for some λ.

We will consider toppling sequences for any finite $V \subseteq \mathbb{Z}^d$, but changing the stability condition at site $\mathbf{0}$. In what follows, we will define weak and strong stabilizations of a configuration η in the box V.

7.1. Definitions and first corollaries

We say that $\mathbf{0}$ is *w-stable* if $\eta(\mathbf{0}) \leqslant 1$, and we say that $\mathbf{0}$ is *s-stable* if $\eta(\mathbf{0}) = 0$. Otherwise we say that $\mathbf{0}$ is *w-unstable* or *s-unstable*. For $y \neq \mathbf{0}$ we say that y is stable, w-stable, and s-stable if $\eta(y) \leqslant \mathfrak{s}$.

Toppling a site z will be called *w-legal* if z is w-unstable, and *s-legal* if z is s-unstable (recall from §2.2 that toppling a site containing a sleeping particle is a well-defined operation). We say that a sequence α of acceptable topplings *weakly stabilizes* η in V if $\Phi_\alpha \eta$ is w-stable in V. We say that it *strongly stabilizes* η in V if $\Phi_\alpha \eta$ is s-stable in V.

Let us restate Lemma 2.1 in this context.

Lemma 7.3. *If α is an acceptable sequence of topplings that w-stabilizes η in V, and $\beta \subseteq V$ is a w-legal sequence of topplings for η, then $m_\beta \leqslant m_\alpha$. The same holds replacing 'w' by 's'.*

Proof. It is the very same proof as that of Lemma 2.1. □

Define

$$m^w_{V,\eta} = \sup_{\beta \subseteq V \text{ w-legal}} m_\beta, \qquad m^s_{V,\eta} = \sup_{\beta \subseteq V \text{ s-legal}} m_\beta.$$

Notice that w-legal topplings are always legal, which are always s-legal, which in turn are always acceptable. In particular, we have by inclusion

$$m^w_{V,\eta} \leqslant m_{V,\eta} \leqslant m^s_{V,\eta}.$$

Now let η_0 be random and independent of $\mathcal{I}$, as described in §2.3. Let η'_V be the resulting configuration obtained by stabilizing η_0 in V with legal topplings, and η^w_V be the result of weakly stabilizing η_0 in V with w-legal topplings. For finite V, these are a.s. well-defined since a stable or w-stable configuration is a.s. achieved after finitely many topplings. The Abelian property implies that neither η'_V nor η^w_V depends on the order at which these legal or w-legal topplings are performed.

Proof of Theorem 7.1. By Abelianness, one way to stabilize η_0 in V is to first weakly stabilize η_0 in V, and then stabilize η^w_V in V.

With the procedure in mind, we claim that

$$m_{V,\eta_0}(\mathbf{0}) \geqslant 1 \quad \Longrightarrow \quad \eta^w_V(\mathbf{0}) = 1.$$

If there is ever a particle at $\mathbf{0}$, weak stabilization does not let this particle leave, and $\eta^w_V(\mathbf{0}) = 1$. Otherwise, $\mathbf{0}$ is never visited during weak stabilization and $\eta^w_V(\mathbf{0}) = 0$, in which case η^w_V is not only w-stable but also stable, so $\eta'_V = \eta^w_V$ and $m_{V,\eta_0}(\mathbf{0}) = 0$, which proves the claim. On the other hand,

$$\mathbb{P}\left(\eta'_V(\mathbf{0}) = \mathfrak{s}\right) \geqslant \tfrac{\lambda}{1+\lambda}\, \mathbb{P}\left(\eta^w_V(\mathbf{0}) = 1\right).$$

Indeed, if $\eta^w_V(\mathbf{0}) = 1$ then $\mathbf{0}$ is the only site in V where η^w_V is unstable, so stabilization of η^w_V starts with a toppling at $\mathbf{0}$. In this case, with probability $\frac{\lambda}{\lambda+1}$, stabilization is achieved immediately, with a particle sleeping at $\mathbf{0}$.

From these two observations we get

$$\mathbb{P}\left(\eta'_V(\mathbf{0}) = \mathfrak{s}\right) \geqslant \tfrac{\lambda}{1+\lambda}\, \mathbb{P}\left(m_{V,\eta_0}(\mathbf{0}) \geqslant 1\right).$$

Now suppose $m_{\eta_0}(\mathbf{0}) = \infty$ a.s. and let $\varepsilon > 0$. Take r such that $\mathbb{P}(m_{V,\eta_0}(\mathbf{0}) \geqslant 1) \geqslant 1 - \varepsilon$ for every $V \supseteq B_r$. Take R such that $|B_{R-r}| \geqslant (1-\varepsilon)|B_R|$. From the previous inequality, $\mathbb{E}\left[\#\{x \in B_{R-r} : \eta'_{B_R}(x) = \mathfrak{s}\}\right] \geqslant \frac{\lambda}{1+\lambda}(1-\varepsilon)^2 |B_R|$. On the other hand, all particles in η'_{B_R} were initially present in B_R, therefore $\frac{\lambda}{1+\lambda}(1-\varepsilon)^2|B_R| \leqslant \mathbb{E}\|\eta_0\|_V = \zeta |B_R|$. Since ε was arbitrary, we must have $\zeta \geqslant \frac{\lambda}{1+\lambda}$. □

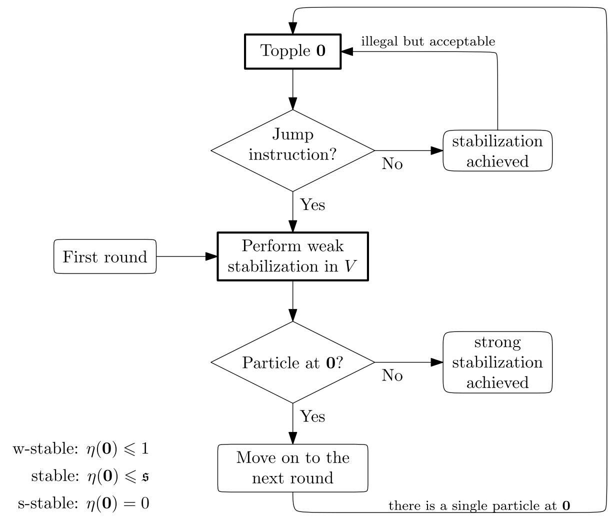

FIG 7.1. *Flow diagram showing a way to obtain stabilization by alternating between weak stabilizations and legal topplings at* $\mathbf{0}$*. If stabilization is achieved with a sleeping particle at* $\mathbf{0}$*, the process can be continued using acceptable topplings until strong stabilization is also achieved.*

7.2. Stabilization via successive weak stabilizations

Let $d > 2$, and define

$$1 < G = \mathbb{E}\left[\text{visits of a random walk to its starting point}\right] < \infty.$$

In this subsection we let $V \ni \mathbf{0}$ and $\eta \in (\mathbb{N}_0)^V$ be finite and fixed.

In the following arguments we consider a toppling procedure for obtaining *stabilization and strong stabilization via successive weak stabilizations*, shown in Figure 7.1. Let T_V and T_V^s count the number of rounds needed for stabilization and strong stabilization to be achieved, respectively (weak stabilization is always achieved in the first round). From this definition we have

$$T_V = 1 \quad \iff \quad \eta_V^w(\mathbf{0}) = 0 \quad \iff \quad T_V^s = 1. \tag{7.4}$$

Before we start using this procedure and studying T_V, we need a couple of lemmas. We define the "jump odometer" $\bar{m}_{V,\eta}$ by counting only the number of jump instructions performed at each site when η is stabilized in V. Define $\bar{m}^s_{V,\eta}$ and $\bar{m}^w_{V,\eta}$ similarly. Let $\eta^+ = \eta + \delta_{\mathbf{0}}$ denote the result of adding an active particle at $\mathbf{0}$ to a configuration η.

Lemma 7.5 (Strong−weak=extra particle)**.** *We have* $\bar{m}^s_{V,\eta} = \bar{m}^w_{V,\eta^+}$.

Proof. A sequence of topplings β is w-legal for η^+ if and only if it is s-legal for η. This proves the lemma. $\square$

Lemma 7.6 (Getting rid of the extra particle)**.** *We have*

$$\mathbb{E}\big[\bar{m}^w_{V,\eta^+}(\mathbf{0})\big] \leqslant G + \mathbb{E}\big[\bar{m}^w_{V,\eta}(\mathbf{0})\big].$$

Proof. Consider the following toppling procedure. First, make the extra particle at $\mathbf{0}$ jump until it leaves V. Then, weakly stabilize the resulting configuration (which is η). Since the resulting configuration is weakly stable, by Lemma 7.3 the jump odometer of this procedure gives an upper bound for $\bar{m}^w_{V,\eta^+}(\mathbf{0})$. Now the expected number of visits to $\mathbf{0}$ in the first stage is bounded by G (in fact it tends to G as V increases), proving the inequality. □

From the two previous lemmas, we get the following.

Corollary 7.7. *We have* $\mathbb{E}\left[\bar{m}^s_{V,\eta}(\mathbf{0}) - \bar{m}^w_{V,\eta}(\mathbf{0})\right] \leqslant G$.

We finally derive estimates for T_V and T^s_V.

Lemma 7.8. *We have* $\bar{m}^s_{V,\eta}(\mathbf{0}) \geqslant \bar{m}^w_{V,\eta}(\mathbf{0}) + T^s_V - 1$.

Proof. Consider the strong stabilization of η on V via successive weak stabilizations as shown in Figure 7.1, run until strong stabilization is achieved. The first round starts in the middle of the diagram and consists of a weak stabilization. For each of the other $T^s_V - 1$ rounds, a jump instruction is eventually performed at $\mathbf{0}$ before strong stabilization is achieved. □

From the two previous statements, we get the following.

Corollary 7.9. *We have* $\mathbb{E}T_V \leqslant \mathbb{E}T^s_V \leqslant 1 + G \leqslant 2G$.

7.3. Main estimates

Theorem 7.2 will follow from two propositions, based on the previous properties of stabilization and strong stabilization via successive weak stabilizations. We now take the initial configuration η_0 random and i.i.d.

We use the following decomposition:

$$\mathbb{P}\left(\eta'_V(\mathbf{0}) = \mathfrak{s}\right) = \sum_{n=2}^{\infty} \mathbb{P}\left(\eta'(\mathbf{0}) = \mathfrak{s}, T_V = n\right), \tag{7.10}$$

where $n = 1$ is excluded by (7.4).

Proposition 7.11. $\mathbb{P}(\eta'_V(\mathbf{0}) = \mathfrak{s}) \leqslant 4G\sqrt{\lambda}$ *for every V finite.*

Proof. We want to control the summand in (7.10). Given that after $n-1$ rounds there is an active particle at $\mathbf{0}$, the conditional probability that the next instruction at $\mathbf{0}$ is a sleep or jump instruction remains unaffected. In case it is a sleep instruction, the procedure stops at round n and $\eta'_V(\mathbf{0}) = \mathfrak{s}$. In case it is a jump instruction, the procedure continues and might reach round $n+1$. So by induction on n we get, for $n \geqslant 2$,

$$\mathbb{P}(\eta'_V(\mathbf{0}) = \mathfrak{s}, T_V = n) \leqslant \tfrac{\lambda}{1+\lambda}\left(\tfrac{1}{1+\lambda}\right)^{n-2}.$$

We now split the sum in (7.10) at an arbitrary point n_0 to get

$$\mathbb{P}(\eta'_V(\mathbf{0})=\mathfrak{s}) \leqslant \sum_{n=2}^{n_0} \tfrac{\lambda}{1+\lambda}\left(\tfrac{1}{1+\lambda}\right)^{n-2} + \sum_{n=n_0+1}^{\infty} \mathbb{P}(T_V=n).$$

By simple Markov inequality, minimization over n_0 and Corollary 7.9,

$$\mathbb{P}(\eta'_V(\mathbf{0})=\mathfrak{s}) \leqslant \lambda\cdot n_0 + \frac{\mathbb{E}T_V}{n_0+1} \leqslant 2\sqrt{\lambda \mathbb{E}T_V} \leqslant 2\sqrt{2G\,\lambda},$$

which completes the proof. □

To boost the previous estimate we need a more careful analysis.

Proposition 7.12. $\mathbb{P}(\eta'_V(\mathbf{0})=\mathfrak{s}) \leqslant 1-(2+2\lambda)^{-4G}$ *for every* V *finite.*

Proof. Let $n \geqslant 2$. We start observing that

$$\mathbb{P}\left(\eta'_V(\mathbf{0})=\mathfrak{s}, T_V=n \,\middle|\, T_V \geqslant n\right) = \tfrac{\lambda}{1+\lambda}. \tag{7.13}$$

We want an upper bound for $\mathbb{P}(\eta'_V(\mathbf{0})=\mathfrak{s}, T_V=n)$ which remains smaller than 1 when summed over n, even for large λ. For that we will relate stabilization to strong stabilization in a way that retains independence.

Let $k \geqslant 2$ be fixed. The event $T^s_V \geqslant k$ is equal to the event that, on rounds $1,\dots,k-1$ of strong stabilization via successive weak stabilizations shown in Figure 7.1, the answer to "Particle at $\mathbf{0}$?" is "Yes."

Now the main observation is that this event is *independent* of the number of times the upper cycle (the one where $\mathbf{0}$ is toppled persistently until a jump instruction is found) is performed at each of the rounds $1,\dots,k-1$. Indeed, inserting or removing sleep instructions on the stack $(\mathfrak{t}^{\mathbf{0},j})_j$ at the origin may only add or remove cycles in the upper part of the flow diagram, but has no effect on the outcome of the lower part.

Hence, for $n<k$ we have $\mathbb{P}\left(T_V>n \,\middle|\, T_V\geqslant n, T^s_V \geqslant k\right) = \tfrac{1}{1+\lambda}$, which corresponds to the answer to "Jump instruction?" being "Yes" right at the beginning of round n. Indeed, this ensures that $T_V>n$ because $T^s_V\geqslant k$ implies that the next question is also being answered affirmatively. From this identity, we get

$$\mathbb{P}\left(T_V \geqslant n+1 \,\middle|\, T^s_V \geqslant k\right) = \tfrac{1}{1+\lambda}\,\mathbb{P}\left(T_V\geqslant n \,\middle|\, T^s_V\geqslant k\right).$$

Iterating the above equality for $n-1, n-2, \dots$ and using (7.4) yields

$$\mathbb{P}\left(T_V \geqslant n+1 \,\middle|\, T^s_V \geqslant k\right) = \left(\tfrac{1}{1+\lambda}\right)^{n-1}\mathbb{P}\left(T_V\geqslant 2 \,\middle|\, T^s_V\geqslant k\right) = \left(\tfrac{1}{1+\lambda}\right)^{n-1}.$$

Finally, taking $k=n+1$, and since $T_V \geqslant n+1$ implies $T^s_V \geqslant k$, the equality becomes

$$\mathbb{P}\left(T_V \geqslant n+1\right) = \left(\tfrac{1}{1+\lambda}\right)^{n-1}\mathbb{P}\left(T^s_V \geqslant n+1\right).$$

Substituting this and (7.13) into (7.10) gives

$$\mathbb{P}(\eta'_V(\mathbf{0})=\mathfrak{s}) = \sum_{n=2}^{\infty} \tfrac{\lambda}{1+\lambda}\,\mathbb{P}\left(T_V\geqslant n\right)$$

$$= \tfrac{\lambda}{1+\lambda} \sum_{n=0}^{\infty} \left(\tfrac{1}{1+\lambda}\right)^n \mathbb{P}\left(T_V^s \geqslant n+2\right).$$

Splitting the sum at $n_0 = \lfloor 4G \rfloor \geqslant 2\mathbb{E}[T_V^s - 2]$ by Corollary 7.9, we get

$$\begin{aligned}\mathbb{P}(\eta'_V(\mathbf{0}) = \mathfrak{s}) &\leqslant \tfrac{\lambda}{1+\lambda} \sum_{n=0}^{n_0-1} \left(\tfrac{1}{1+\lambda}\right)^n + \tfrac{\lambda}{1+\lambda} \sum_{n=n_0}^{\infty} \tfrac{1}{2} \left(\tfrac{1}{1+\lambda}\right)^n \\ &= 1 - \tfrac{1}{2}(1+\lambda)^{-n_0} \leqslant 1 - (2+2\lambda)^{-4G},\end{aligned}$$

which concludes the proof. □

Proof of Theorem 7.2. Assume that $\zeta > 4G\sqrt{\lambda}$. The number N of particles that exit V during stabilization of V equals $\|\eta_0\|_V$ minus the number of sites z such that $\eta'_V(z) = \mathfrak{s}$. Using Proposition 7.11, $EN \geqslant (\zeta - 4G\sqrt{\lambda}) \times |V|$, so Condition (2.12) is satisfied and therefore the system a.s. stays active. The same argument works assuming $\zeta > 1 - (2+2\lambda)^{-4G}$. □

8. Uniqueness of the critical density

In this section we prove Theorem 2.13, or the following equivalent formulation.

Theorem 8.1. *Let d, $p(\cdot)$ and λ be given. Let ν_1 and ν_2 be two spatially ergodic distributions on $(\mathbb{N}_0)^{\mathbb{Z}^d}$, with respective densities $\zeta_1 < \zeta_2$. If the ARW system is a.s. fixating with initial state ν_2, then it is also a.s. fixating with initial state ν_1.*

Below we briefly sketch the proof of Theorem 8.1, and then give the complete proof in three parts: embedding the initial configuration into another one with higher density, stabilization of the embedded configuration, and finally stabilization of the original configuration. An interesting feature that distinguishes the toppling procedure used here is that it is not a sequential procedure for ever-growing finite domains as in the previous sections, but rather a sequence of parallel-update type of operation that topples infinitely many sites at once, in order to use ergodicity and mass conservation.

Open Problem. Suppose ν is a translation-ergodic active state (active means ν is supported on $(\mathbb{N}_{\mathfrak{s}})^{\mathbb{Z}^d} \setminus \{0,\mathfrak{s}\}^{\mathbb{Z}^d}$) with density $\zeta > \zeta_c$. Show that the ARW with initial state ν a.s. stays active.

Open Problem. Suppose ν is a translation-ergodic absorbing state (absorbing means supported on $\{0,\mathfrak{s}\}^{\mathbb{Z}^d}$) with density $\zeta > \zeta_c$. Show that the ARW with initial state ν plus one active particle at the origin stays active with positive probability. Note that this property is false if the graph is a tree [JR19].

Open Problem. Under which conditions besides $\zeta < \zeta_c$ does $\mathbb{E}[m(\mathbf{0})] < \infty$?

Open Problem. When does $\mathbb{P}(m(\mathbf{0}) \geqslant 1) = 1$ imply $\mathbb{P}(m(\mathbf{0}) = \infty) = 1$?

Let us give a brief sketch before moving to the proof.

The proof is algorithmic and has two stages, both stages being infinite. The idea is very simple and is related to what is sometimes called decoupling. Let η_0 and ξ_0 be independent and distributed as ν_1 and ν_2. In the first stage, we evolve η starting from η_0 until it gives a configuration $\eta'_0 \leqslant \xi_0$. In the second stage, we use the same set of instructions to evolve both systems. Since the evolution of ξ starting from ξ_0 a.s. fixates, so does the evolution of η starting from η'_0, concluding the proof. More precisely, in the first stage we force each particle in the system η to move (by waking it up if needed) until it meets a particle of ξ_0; once they meet, they are paired and will not be moved until the second stage. Even if it takes infinitely many steps to finish pairing globally, a.s. every particle in the system η will eventually be paired, and the resulting odometer will be a.s. finite at every site (if the odometer were infinite somewhere, by ergodicity it would be infinite everywhere, so every particle in ξ_0 would be paired, implying $\zeta_2 \leqslant \zeta_1$ by mass conservation). This yields a configuration $\eta'_0 \leqslant \xi_0$. In the second stage, we simply evolve the system using the remaining instructions. Since they are independent of ξ_0, η_0, and of the instructions used in the first stage, by assumption the remaining instructions a.s. stabilize ξ_0 leaving a locally-finite odometer. By monotonicity of the final odometer with respect to the configuration, the same set of remaining instructions also stabilizes η'_0, again with a locally-finite odometer. Adding the odometer of both stages would give the final locally-finite odometer given by stabilization of η_0, except that in the first stage we have not followed the toppling rules correctly. But it still gives an upper bound due to Lemma 2.1.

We now turn to the proof.

To make the argument precise we will not exactly move particles as in the previous sketch, since the embedding requires an infinite number of topplings. We instead explore the instructions and define a sequence of configurations in terms of η_0, ξ_0 and $\mathcal{I}$. We end up concluding that a.s. the result of this exploration implies that η_0 is stabilizable, which in turn implies the statement of the theorem.

Embedding of the smaller configuration

Without loss of generality, we assume that ν_1 or ν_2 is not only ergodic but also mixing (otherwise consider ν_3 as i.i.d. Poisson with mean $\frac{\zeta_1+\zeta_2}{2}$ which is mixing, and apply the result from ν_2 to ν_3 and from ν_3 to ν_1). So suppose ν_2 is mixing. Recall that $\mathcal{I}$ is an i.i.d. field, thus the pair $(\xi_0, \mathcal{I})$ is mixing and hence the triple $\omega = (\eta_0, \xi_0, \mathcal{I})$ is ergodic, see §2.4.

Let η_0, ξ_0 and $\mathcal{I}$ be given, and take $h_0 \equiv 0$. Fix some $k = 0, 1, 2, 3, 4, \dots$ and suppose η_k and h_k have been defined as a factor of ω (see §2.4). Denote by A_k the set given by

$$A_k = \{x : \eta_k(x) > \xi_0(x)\},$$

and consider an arbitrary enumeration

$$A_k = \{x_1^k, x_2^k, x_3^k, \dots\}.$$

Let

$$(\eta_k^j, h_k^j) = \Phi_{(x_1^k, x_2^k, \dots, x_j^k)}(\eta_k, h_k) \quad \text{and} \quad (\eta_{k+1}, h_{k+1}) = \lim_j (\eta_k^j, h_k^j) \tag{8.2}$$

in case A_k is infinite – in case it is finite, by ergodicity it is a.s. empty in which case we let $(\eta_{k+1}, h_{k+1}) = (\eta_k, h_k)$. Note that the condition $\eta_k(x) > \xi_0(x)$ is also satisfied when $\eta_k(x) = \mathfrak{s}$ and $\xi_0(x) = 0$, so this operation may require waking up particles. That is, in going from (η_k, h_k) to (η_{k+1}, h_{k+1}), every site in A_k is toppled once. These topplings are legal when $\eta_k(x) \geqslant 1$, and they are acceptable but illegal in case $\eta_k(x) = \mathfrak{s} > 0 = \xi_0(x)$.

As we go through $j = 1, 2, 3, \dots$ in (8.2), for each j the field h_k^j is increased by one unit at x_j^k, so h_{k+1} is well-defined and satisfies

$$h_{k+1}(x) = h_k(x) + \mathbb{1}_{A_k}(x).$$

To see that η_k is also well-defined in the limit (8.2), observe that, for each site x, the sequence $(\eta_k^j(x))_j$ decreases for at most one value of j. In case it decreases, it may send one particle to another site $z \neq x$. Thus, by a standard use of the mass transport principle, the configuration at each site x increases a finite number of times (the expected number of times is less than one, as can be seen by taking $f(x, y)$ as the indicator of the event that $x \in A_k$ and toppling x sends a particle to y), so the limit η_k is a.s. finite. Since there are a.s. finitely many sites $z \in A_k$ that send a particle to x when toppled, it follows from the local Abelian property that the limit (η_k, h_k) does not depend on the enumeration of A_k, so it is a factor of ω.

Let $k \to \infty$ and define

$$h_0'(x) = \lim_k h_k(x).$$

We now prove that, if $\mathbb{P}(h_0'(\mathbf{0}) = \infty) > 0$ then we must have $\zeta_1 \geqslant \zeta_2$.

First, we claim that $\mathbb{P}(h_0'(\mathbf{0}) = \infty) = 0$ or 1. Since (η_k, h_k) is a factor of ω, it is translation-ergodic. Moreover, a.s. the event that $h_0'(\mathbf{0}) = \infty$ implies the event that $h_0'(z) = \infty$ for every z such that $p(z) > 0$. These two facts together imply that, either $h_0'(x) = \infty$ a.s. for every x, or $h_0'(x) < \infty$ a.s. for every x, see §11.1. This proves the zero-one law.

Suppose $h_0'(\mathbf{0}) = \infty$ with positive probability. By the zero-one law we have $h_0'(\mathbf{0}) = \infty$ a.s., which means that $\mathbb{P}(\limsup_k \{\mathbf{0} \in A_k\}) = 1$. But if $\mathbf{0} \in A_{k_0}$ for some k_0, then necessarily $\eta_{k_0-1}(\mathbf{0}) > \xi_0(\mathbf{0})$, thus $\eta_k(\mathbf{0}) \geqslant \xi_0(\mathbf{0})$ for all $k \geqslant k_0$ by definition of A_k, and therefore $\liminf_k |\eta_k(\mathbf{0})| \geqslant |\xi_0(\mathbf{0})|$. On the other hand, from the mass transport principle we have $\mathbb{E}|\eta_k(\mathbf{0})| = \mathbb{E}|\eta_{k-1}(\mathbf{0})| = \cdots = \mathbb{E}|\eta_0(\mathbf{0})| = \zeta_1$ (to show the first identity, we let $f_k(x, y)$ be the indicator that, on step k, x sends a particle to y, and let $f_k(x, x)$ be the number of particles that were present at x at the beginning of stage k and stayed at x). By Fatou's Lemma, $\zeta_1 \geqslant \zeta_2$.

Since we are assuming $\zeta_1 < \zeta_2$, we must have $h'_0(\mathbf{0}) < \infty$ a.s. Now, as we go through $k = 1, 2, 3, \dots$, the value of $\eta_k(\mathbf{0})$ can decrease only when $\mathbf{0} \in A_k$, *i.e.* only when $h_k(\mathbf{0})$ increases. Hence, $(\eta_k(\mathbf{0}))_k$ is a.s. eventually non-decreasing, so it converges. Its limit $\eta'_0(\mathbf{0})$ satisfies $\eta'_0(\mathbf{0}) \leqslant \xi_0(\mathbf{0})$, otherwise $\mathbf{0}$ would be in A_k for all large enough k and $h'_0(\mathbf{0})$ would be infinite. By translation invariance, a.s. $h'_0(x) < \infty$ and $\eta'_0(x) = \lim_k \eta_k(x) \leqslant \xi_0(x)$ for every x.

Stabilization of the original configuration

In the previous stage we obtained a pair (η'_0, h'_0) a.s. satisfying $h'_0(x) < \infty$ and $\eta'_0(x) \leqslant \xi_0(x)$ for every $x \in \mathbb{Z}^d$. Let $\tilde{\mathcal{I}}$ be the set of instructions given by

$$\tilde{\mathfrak{t}}^{x,j} = \mathfrak{t}^{x,h'_0(x)+j}, \qquad x \in \mathbb{Z}^d, j \in \mathbb{N}.$$

that is, the field obtained by deleting the instructions used in the embedding stage described above. Since the first $h'_0(x)$ instructions have been deleted at each site x, stabilizing a system with the instructions in $\tilde{\mathcal{I}}$ instead of $\mathcal{I}$ is equivalent to starting with odometer at h'_0 instead of $h_0 \equiv 0$.

Now note that the collection of instructions $\big(\mathfrak{t}^{x,j} : x \in \mathbb{Z}^d, j > h'_0(x)\big)$ played no role in the construction of η'_0 and h'_0, so they are independent of ξ_0 and h'_0. Hence, $\tilde{\mathcal{I}}$ is an i.i.d. field just like $\mathcal{I}$, and it is also independent of ξ_0.

Therefore, $\mathbb{P}\big(\xi_0 \text{ is } \tilde{\mathcal{I}}\text{-stabilizable}\big) = \mathbb{P}\big(\xi_0 \text{ is } \mathcal{I}\text{-stabilizable}\big)$, and the latter equals 1 by assumption. Since $\eta'_0 \leqslant \xi_0$, we have $\mathbb{P}\big(\eta'_0 \text{ is } \tilde{\mathcal{I}}\text{-stabilizable}\big) \geqslant \mathbb{P}\big(\xi_0 \text{ is } \tilde{\mathcal{I}}\text{-stabilizable}\big) = 1$. This means that a.s. there exists h'_1 such that, for all finite $V \subseteq \mathbb{Z}^d$ and all $x \in \mathbb{Z}^d$, $m_{V,\eta'_0;\tilde{\mathcal{I}}}(x) \leqslant h'_1(x) < \infty$.

Stabilization of the original configuration

Given the properties of the two previous stages, we now give the (perhaps tedious) proof that η_0 is a.s. $\mathcal{I}$-stabilizable. More precisely, we will show that

$$m_{\eta_0;\mathcal{I}}(x) \leqslant h'_0(x) + h'_1(x) < \infty, \quad \forall\, x \in \mathbb{Z}^d.$$

In the first stage, the limits η'_0 and h'_0, which are determined by η_0, ξ_0 and $\mathcal{I}$, almost surely exist and satisfy $h'_0 < \infty$ and $\eta'_0 \leqslant \xi_0$. Suppose this event occurs, and let V be a fixed finite set.

If we start from (η_0, h_0) and perform all topplings in V as well as particle additions to V (coming from V^c), following the same order as in the first stage, only a finite number of operations will be performed, and we end up with a configuration that equals (η'_0, h'_0) on V.

By the local Abelian property, we can add the particles first and then topple the sites in V as in the first stage, obtaining the same result. This means that there is some $\bar{\eta}_V \geqslant \eta_0$ and an acceptable sequence $\alpha_V = (x_1, \dots, x_n)$ for $(\bar{\eta}_V, h_0)$ such that $m_{\alpha_V} = h'_0$ on V and

$$\Phi_{\alpha_V}(\bar{\eta}_V, h_0) = (\eta'_0, h'_0) \text{ on } V.$$

Now, in the second stage, we showed that a.s. there exists $h_1'(x) < \infty$ such that $m_{V',\eta_0';\tilde{\mathcal{I}}}(x) \leqslant h_1'(x)$ for every finite V'. Suppose this event occurs.

Notice that $m_{V,\eta_0',h_0';\mathcal{I}}(x) = m_{V,\eta_0';\tilde{\mathcal{I}}}(x)$, that is, to stabilize η_0' in V using the shifted field of instructions is the same as stabilize η_0' in V using the original field of instructions and shifted odometer. Therefore, there exists $\beta_V = (x_{n+1}, \dots, x_m)$ contained in V such that $m_{\beta_V} \leqslant h_1'$ on V and $\Phi_{\beta_V}(\eta_0', h_0')$ is stable in V.

By the above identity, $\Phi_{\beta_V} \circ \Phi_{\alpha_V}(\bar{\eta}_V, h_0) = \Phi_{\beta_V}(\eta_0', h_0')$, on V. Since the latter is stable in V, by Lemma 2.1 we have

$$m_{V,\bar{\eta}_V;\mathcal{I}}(x) \leqslant m_{\alpha_V}(x) + m_{\beta_V}(x), \qquad \text{for all } x \in \mathbb{Z}^d.$$

Thus, by monotonicity,

$$m_{V,\eta_0;\mathcal{I}}(x) \leqslant m_{V,\bar{\eta}_V;\mathcal{I}}(x) \leqslant m_{\alpha_V}(x) + m_{\beta_V}(x) \leqslant h_0'(x) + h_1'(x) < \infty.$$

We now note that the above bound does not depend on V, so

$$m_{\eta_0;\mathcal{I}}(x) = \sup_{V \text{ finite}} m_{V,\eta_0;\mathcal{I}}(x) \leqslant h_0'(x) + h_1'(x) < \infty, \quad \forall\, x \in \mathbb{Z}^d,$$

which means that η_0 is stabilizable, concluding the proof of Theorem 8.1.

9. A recursive multi-scale argument

In this section we comment on a multi-scale argument used to prove the following.

Theorem 9.1. *If the jumps are unbiased, $\zeta_c > 0$ for every $\lambda > 0$.*

We give an overview of the general strategy, referring the reader to [ST17] for the complete argument. Note that the above theorem is a particular case of Theorem 7.1.

The main step is to show that an initial configuration restricted to a very large box stabilizes within a slightly larger box, with high probability. This is then used to show Condition (2.10). This is proved by recursion on the scale of the box, and in fact the proof does not rely much on specific details of the actual ARW dynamics. In a sense, this kind of approach fits to our intuition that no matter how big a defect is, it will only affect a neighborhood of comparable size.

The box at scale k is a cube V_k of side length L_k, defined as follows. Let $\delta = \frac{1}{10}$, $L_0 = 10^4$ and

$$L_{k+1} = \lfloor L_k^\delta \rfloor^2 L_k.$$

Notice that L_k increases as a doubly exponential of k. We also define $R_{k+1} = \lfloor L_k^\delta \rfloor L_k$ as an intermediate scale between L_k and L_{k+1}. In Figure 9.1, we see an *inner box* V_{k+1}', an intermediate box, and a *full box* V_{k+1} of level $k+1$.

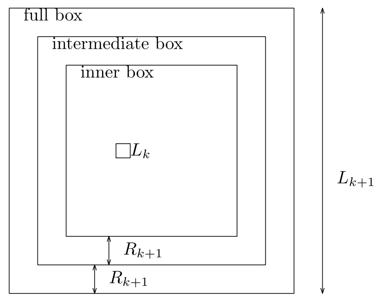

FIG 9.1. *Boxes and scales;* $L_k \ll R_{k+1} \ll L_{k+1}$.

Let p_k denote the probability that, starting from a Poisson configuration in V_k', some particle exits V_k. That $p_k \to 0$ fast as $k \to \infty$ follows from the recursion relation

$$p_{k+1} \leqslant \frac{L_{k+1}^{2d}}{L_k^{2d}} \, {p_k}^2 + e_{k+1},$$

consisting of a combinatorial term, the probability ${p_k}^2$ that stabilization fails twice at scale k, and the probability e_{k+1} that something goes wrong at scale $k+1$. Indeed, if p_{k_0} is small enough and $e_k \to 0$ fast enough, then the square power above beats the $1+2\delta$ power in the definition of L_k, and p_k vanishes doubly-exponentially fast in k.

Let us describe some aspects of this recursion step, depicted in Figure 9.2.

Configurations in light gray have Poisson product distribution with the right density. They are restricted to the inner box of level $k+1$ for the initial configuration, and the inner boxes of level k for the "sieved configurations". Configurations in dark gray are absorbing configurations, typically attained by the dynamics. Configurations in gray with a grid are "balanced configurations". Thick arrows represent typical events, while thin arrows represent events of low probability, either e_{k+1} or p_k.

Starting fresh. To let the dynamics run on boxes of the previous scale and use recursion, it is important to start with a Poisson product distribution within their inner boxes. This is achieved by a sieving procedure described below.

Worst case scenario. If the dynamics does not to stabilize all the $\frac{L_{k+1}^d}{L_k^d}$ boxes of level k, the configuration inside these boxes is no longer i.i.d. Poisson. In the absence of any useful knowledge about the resulting distribution of particles in this case, we use only the fact that the total number of particles within each box is still a Poisson random variable, and thus cannot be much larger than its mean. A *balanced configuration* is such that the number of particles within each box of level k is appropriately bounded.

Sieving procedure. Starting from a balanced configuration, we let each particle move for a certain time, so as to uniformize its relative position within whichever level-k box contains it. After performing all these jumps, if the particle happens not to be in the inner box of the level-k box containing it, we repeat the procedure again, as many times as needed. This *reshuffling* with *sieving* results

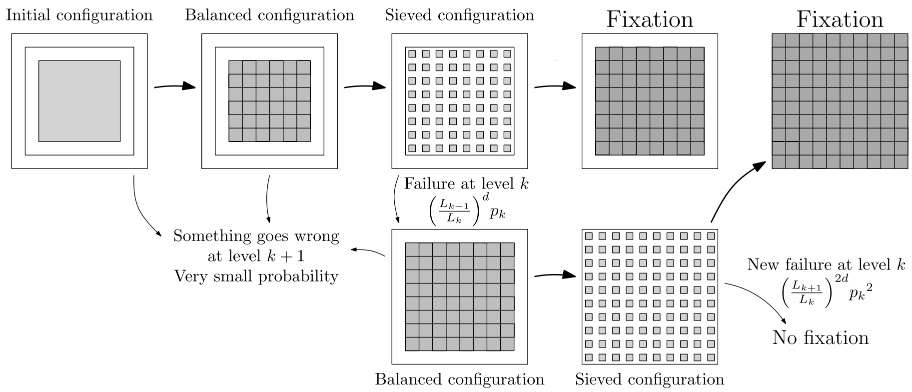

Fig 9.2. *Illustrative diagram of events for the recursion relation.*

in a state that with high probability, can be coupled with an i.i.d. Poisson configuration. This is one of the heaviest parts in [ST17]. In order for this coupling to be possible, a slight increase in the density is necessary, analogous to the *sprinkling* technique in percolation (this increase should decay just fast enough to be summable over k).

The chain of events. By hypothesis, we start with a Poisson product measure inside the inner box V'_{k+1}. Such a configuration is typically balanced, that is, each box of level k has an appropriately bounded number of particles. We then let these particles move around using acceptable topplings so that their distribution is now close to i.i.d. Poisson inside the inner boxes of level k (the sieving procedure). During this procedure they cannot exit the intermediate box. We now let the evolution run normally within each box of level k, and typically each box stabilizes nicely without letting particles leave. It may happen however that some of these boxes of level k is not stabilized as intended (which is atypical). In spite of failing to stabilize as intended, the resulting configuration is still balanced. So we let the particles move around again, now obtaining a sieved configuration in the full box V_{k+1}. Typically, the resulting configuration is properly sieved. The system is then given a second chance to stabilize. This second attempt will typically be successful, but may fail again if some of the level-k boxes does not stabilize as expected.

In the proof there are many aspects to keep under control, and many delicate statements that we omit here. The above description is not intended to serve as a sketch of proof, but hopefully gives a general flavor of the main argument. For all the details, the reader is referred to [ST17].

10. Arguments using particle-wise constructions

The techniques presented so far used the site-wise representation and its properties, as described in §2. In the particle-wise construction, the randomness of the jumps is not attached to the sites, but to the particles.

In this section we use a particle-wise construction to prove Theorem 2.11 and the following ones. We always assume that $\mathbb{E}|\eta_0(\mathbf{0})| < \infty$.

Theorem 10.1 (Mass conservation)**.** *Consider the ARW on $\mathbb{Z}^d$ with i.i.d. initial configuration η_0. If the system a.s. fixates, then $\mathbb{E}|\eta_\infty(\mathbf{0})| = \mathbb{E}|\eta_0(\mathbf{0})|$, where $\eta_\infty(x) = \lim_{t\to\infty} \eta_t(x)$.*

Now if the system fixates, then $\eta_\infty(\mathbf{0}) \in \{0, \mathfrak{s}\}^{\mathbb{Z}^d}$, hence by mass conservation $\zeta = \mathbb{E}|\eta_0(\mathbf{0})| = \mathbb{E}|\eta_\infty(\mathbf{0})| \leqslant 1$. This gives the following corollary.

Corollary 10.2. *$\zeta_c \leqslant 1$ for every λ.*

Combining the above with Theorems 2.14 and 7.1 we get the following.

Corollary 10.3. *For $\lambda = \infty$, $\zeta_c = 1$.*

The process with $\lambda = \infty$ is discussed in §10.3, where we also consider an equivalent model to prove the following.

Theorem 10.4. *For i.i.d. η_0 with density $\zeta = 1$ and positive variance, for all $\lambda \in (0, \infty]$, the ARW a.s. stays active.*

The requirement of positive variance cannot be waived. Indeed, if $\lambda = \infty$ and $\eta_0 \equiv 1$ then the configuration is already stable at time zero.

Open Problem. Prove Theorem 10.1 without using the particle-wise construction.

Open Problem. Prove Theorems 10.1 and 10.4 replacing the the i.i.d. assumption by translation ergodicity.

Remark. We highlight once more that, for unbiased walks on $\mathbb{Z}^2$, Corollary 10.2 is the best bound we have, even for small λ.

We now give a formal description of the particle-wise construction, and then move on to proving the above results.

10.1. Description and basic properties

We want to revisit the description of §1.1 and now interpret that *particles are labeled*. Each existing particle at time $t = 0$ is assigned a label (x, j), where $x \in \mathbb{Z}^d$ denotes its starting position and $j = 1, \dots, |\eta_0(x)|$ distinguishes particles starting at the same site x. Let $Y^{x,j} = (Y^{x,j}_t)_{t\geqslant 0}$ be given by the position of particle (x, j) at each time t. Let $\gamma^{x,j} = (\gamma^{x,j}(t))_{t\geqslant 0}$ be given by $\gamma^{x,j}(t) = 1$ if particle (x, j) is active at time t or $\gamma^{x,j}(t) = \mathfrak{s}$ if it is sleeping. Write $\mathbf{Y} = (Y^{x,j})_{x,j}$ and $\boldsymbol{\gamma} = (\gamma^{x,j})_{x,j}$. Then the triple $\boldsymbol{\eta} = (\eta_0, \mathbf{Y}, \boldsymbol{\gamma})$ describes the whole evolution of the system.

Whereas the process $(\eta_t)_{t\geqslant 0}$, given by

$$\eta_t(z) = \sum_x \sum_{j\leqslant|\eta_0(x)|} \delta_{Y^{x,j}(t)}(z) \cdot \gamma^{x,j}(t),$$

only counts the number of particles at a given site at a given time, having each particle labeled gives a lot more information and allows different techniques to be employed.

For a system whose initial configuration contains finitely many particles, the evolution described above is always well-defined. It is a simple continuous-time Markov chain on a countable space, and many different explicit constructions will produce $\boldsymbol{\eta}$ with the correct distribution. We now describe one which can be extended to infinite initial configurations and has proved particularly useful.

The particle-wise construction

Assign to each particle (x, j) a continuous-time walk $X^{x,j} = (X_t^{x,j})_{t \geqslant 0}$, independently of anything else, as well as a Poisson clock $\mathcal{P}^{x,j} \subseteq \mathbb{R}_+$ according to which the particle will try to sleep. $X^{x,j}$ is the path of the particle parameterized by its *inner time*, which may be slowed down with respect to the system time, depending on the interaction with other particles (denoting the inner time of particle (x, j) at time t by $\sigma^{x,j}(t)$ we have $Y_t^{x,j} = X_{\sigma^{x,j}(t)}^{x,j}$). So $X^{x,j}$ will be called the *putative trajectory* of particle (x, j). Write $\mathbf{X} = (X^{x,j})_{x,j}$ and $\boldsymbol{\mathcal{P}} = (\mathcal{P}^{x,j})_{x,j}$.

For a deterministic initial configuration ξ containing finitely many particles, $\boldsymbol{\eta}$ is a.s. determined by $(\xi, \mathbf{X}, \boldsymbol{\mathcal{P}})$ in the obvious way. The construction for infinite η_0 is done via limits over the sequences of balls $(B_n^y)_{n\in\mathbb{N}}$, centered at each site $y \in \mathbb{Z}^d$. This family of sequences is countable and translation-invariant. For $\eta \in (\mathbb{N}_\mathfrak{s})^{\mathbb{Z}^d}$ and finite $V \subseteq \mathbb{Z}^d$, let $\eta^V = \eta \cdot \mathbb{1}_V$ denote the restriction of η to V.

Definition 10.5 (Well-definedness)**.** We say that the above construction is *well-defined* if: (i) for each $x, y \in \mathbb{Z}^d$, $j \in \mathbb{N}$ and $t > 0$, both $(Y_s^{x,j})_{s\in[0,t]}$ and $(\gamma_s^{x,j})_{s\in[0,t]}$ are the same in the systems $(\eta_0^{B_n^y}, \mathbf{X}, \boldsymbol{\mathcal{P}})$ for all but finitely many n; (ii) the limiting process $\boldsymbol{\eta} = (\eta_0, \mathbf{Y}, \boldsymbol{\gamma})$ does not depend on y.

The benefit of requiring the limit not to depend on y is that $\boldsymbol{\eta}$ is a factor of ω (see §2.4). In particular, the system with labeled particles is translation-ergodic, satisfies the mass transport principle, and probabilities of local events can be approximated by finite systems regardless of which construction is used. This last property implies that different explicit constructions all yield a process $(\eta_t)_{t\geqslant 0}$ with the same law.

Theorem 10.6. *If* $\sup_x \mathbb{E}|\eta_0(x)| < \infty$*, then the above particle-wise construction is a.s. well-defined.*

The proof is deferred to §11.3.

Fixation equivalence and mass conservation

Fixation as defined in §1.2 concerns the state of sites. Now that the particles are being labeled, it makes sense to consider fixation of particles. We say that *particle* (x, j) *stays active* if $|\eta_0(x)| \geqslant j$ and $\gamma^{x,j}(t)$ is not eventually $\mathfrak{s}$. Likewise, we say that *site* x *stays active* if $\eta_t(x)$ is not eventually constant.

Theorem 10.7. *Suppose η_0 is i.i.d. The following are equivalent:*

(*i*) $\mathbf{P}$(*some site stays active*) > 0*;*
(*ii*) $\mathbf{P}$(*all sites stay active*) $= 1$*;*
(*iii*) $\mathbb{P}$(*all particles stay active*) $= 1$*;*
(*iv*) $\mathbb{P}$(*some particle stays active*) > 0*.*

The proof is given in §10.4. We are ready to show mass conservation.

Proof of Theorem 10.1. We use the construction provided by Theorem 10.6. If particle (x, j) fixates, then $\sigma^{x,j}(t)$ and $Y_t^{x,j}$ are eventually constant, and we say that particle (x, j) fixates at site $Y_\infty^{x,j}$. Let $A(x, j, y)$ denote the event that particle (x, j) fixates at site y, and let $f(x, y) = \sum_j \mathbb{1}_{A(x,j,y)}$.

Assume that a.s. all sites fixate. Note that $\eta_\infty(\mathbf{0}) = \mathfrak{s}$ if and only if we have $\sum_y f(y, \mathbf{0}) = 1$, otherwise $\eta_\infty(\mathbf{0}) = 0$ and $\sum_y f(y, \mathbf{0}) = 0$. On the other hand, by Theorem 10.7 a.s. no particles stay active, whence $\sum_y f(\mathbf{0}, y)$ equals $|\eta_0(\mathbf{0})|$. Applying the mass transport principle concludes the proof. □

10.2. Averaged condition for activity

In this subsection we prove Theorem 2.11. We are assuming that η_0 is i.i.d. We can moreover assume that $\mathbb{E}|\eta_0(\mathbf{0})| < \infty$ (otherwise truncate η_0 and use Corollary 10.2 combined with Lemma 2.5).

The variable M_n in Condition (2.12) refers to the site-wise representation of a finite system restricted to V_n. This is equivalent to a particle-wise construction with particles being killed when they exit V_n. To show that Condition (2.12) implies non-fixation, we consider a different variable M_n^* which counts how many labeled particles start in V_n and ever visit V_n^c during the evolution of the infinite system without killing. More precisely, we consider the system with labeled particles as provided by Theorem 10.6.

Lemma 10.8. $\mathbb{E}M_n \leqslant \mathbb{E}M_n^*$.

The proof of Lemma 10.8 is deferred to §11.5. Define

$$\widetilde{V}_n = V_{n-L_n},$$

where L_n is an integer sequence (e.g. $\lfloor \log n \rfloor$) such that

$$L_n \to \infty \qquad \text{but} \qquad \frac{|V_n \setminus \widetilde{V}_n|}{|V_n|} \to 0.$$

For $n \in \mathbb{N}$, introduce the event

$$\mathcal{A}_n = \text{“}\sup_t |Y_t^{\mathbf{0},1}| \geqslant L_n\text{”} = \text{“particle } (\mathbf{0}, 1) \text{ reaches distance } L_n\text{”},$$

where the requirement that $|\eta_0(\mathbf{0})| \geqslant 1$ is implicit.

Let $\widetilde{M}_n^*$ be the number of labeled particles starting in $\widetilde{V}_n$ which ever exit V_n. By translation invariance and particle exchangeability, for every K,

$$\begin{aligned}
\mathbb{E}\widetilde{M}_n^* &= \sum_{x\in\widetilde{V}_n}\sum_{i\in\mathbb{N}}\mathbb{P}(\text{particle } Y^{x,i} \text{ exits } V_n) \\
&\leqslant \sum_{x\in\widetilde{V}_n}\sum_{i\in\mathbb{N}}\mathbb{P}(\eta_0(x)\geqslant i \text{ and particle } Y^{x,i} \text{ reaches distance } L_n \text{ from } x) \\
&= |\widetilde{V}_n|\sum_{i\in\mathbb{N}}\mathbb{P}(\eta_0(\mathbf{0})\geqslant i \text{ and particle } Y^{\mathbf{0},i} \text{ reaches distance } L_n \text{ from } \mathbf{0}) \\
&= |\widetilde{V}_n|\sum_{i\in\mathbb{N}}\mathbb{P}(\eta_0(\mathbf{0})\geqslant i \text{ and particle } Y^{\mathbf{0},1} \text{ reaches distance } L_n \text{ from } \mathbf{0}) \\
&\leqslant |\widetilde{V}_n|\sum_{1\leqslant i\leqslant K}\mathbb{P}(\mathcal{A}_n) + |\widetilde{V}_n|\sum_{i>K}\mathbb{P}(\eta_0(\mathbf{0})\geqslant i) \\
&= |\widetilde{V}_n|K\,\mathbb{P}(\mathcal{A}_n) + |\widetilde{V}_n|\,\mathbb{E}[(|\eta_0(\mathbf{0})|-K)^+]
\end{aligned}$$

Hence,

$$\begin{aligned}
\mathbb{E}M_n^* &\leqslant \mathbb{E}\widetilde{M}_n^* + \zeta\,|V_n\setminus\widetilde{V}_n| \leqslant \\
&\leqslant K\,|\widetilde{V}_n|\,\mathbb{P}(\mathcal{A}_n) + |\widetilde{V}_n|\,\mathbb{E}[(|\eta_0(\mathbf{0})|-K)^+] + \zeta\,|V_n\setminus\widetilde{V}_n|,
\end{aligned}$$

and using Lemma 10.8,

$$\begin{aligned}
\mathbb{P}(\text{particle } (\mathbf{0},1) \text{ stays active}) &= \lim_n \mathbb{P}(\mathcal{A}_n) \geqslant \\
&\geqslant \frac{1}{K}\Big(\limsup_n \frac{\mathbb{E}M_n}{|\widetilde{V}_n|} - \mathbb{E}[(\eta_0(\mathbf{0})-K)_+]\Big)
\end{aligned}$$

which is positive provided K is chosen large enough.

From Theorem 10.7, we conclude that a.s. all sites stay active, which finishes the proof of the theorem.

10.3. Resampling

In this subsection we prove Theorem 10.4.

By Theorem 2.14, we can assume $\lambda = \infty$. In this case, the sleep Poisson clocks $\boldsymbol{\mathcal{P}}$ play no role in the previous construction: a particle is sleeping if and only if there are no other particles at the same site.

We consider the following dynamics instead of the ARW.

The particle-hole model. Particles perform continuous-time random walks independently of each other. Sites not containing any particle are called *holes*. When a particle is alone at some site, it *settles* there forever, filling the corresponding hole. After the hole has been filled, the site becomes available for other particles to go through. If a site is occupied by several particles at $t = 0^-$, we choose one of them uniformly to fill the hole at $t = 0$, and the other particles remain free to move. This is well-defined as in Theorem 10.6, with the same proof.

This model is very similar to the ARW with $\lambda = \infty$. In both models, once a site has at least one particle, it will always retain one particle. The differences are (i) sites with $n \geqslant 2$ particles are toppled at rate $n - 1$ instead of n and (ii) each site retains forever the first particle to arrive there, whereas in the ARW the particles can take turns replacing each other. Nevertheless, both models have the same site-wise representation (described in §11.4) and same fixation properties in the sense of Theorem 2.7.

We will show that, under the assumption of site fixation,

$$\mathbb{P}(\mathbf{0} \text{ is never visited}) > 0.$$

This implies that $\zeta < 1$ by Theorem 10.1, therefore proving Theorem 10.4.

Now assuming site fixation, necessarily there exists $k \in \mathbb{N}$ such that

$$\mathbb{P}\big(\text{the number of particles which ever visit } \mathbf{0} \text{ equals } k\big) > 0.$$

Hence, there exist $x_1, \dots, x_k \in \mathbb{Z}^d$ such that $\mathbb{P}(\mathcal{A}) > 0$, where

$$\mathcal{A} = \text{“the particles which ever visit } \mathbf{0} \text{ are initially at the sites } x_1, \dots, x_k.\text{”}$$

Consider two systems ω and $\tilde{\omega}$, coupled as follows. We take $\tilde{\mathbf{X}} = \mathbf{X}$, and $\tilde{\eta}_0(x) = \eta_0(x)$ for $x \notin \{x_1, \dots, x_k\}$. For $x \in \{x_1, \dots, x_k\}$, we sample $\tilde{\eta}_0$ and η_0 independently. Now notice that

$$\begin{aligned}
&\mathbb{P}(\mathcal{A} \text{ occurs for } \tilde{\omega}, \text{ and } \eta_0(x_1) = \cdots = \eta_0(x_k) = 0) = \\
&\qquad = \mathbb{P}(\mathcal{A} \text{ occurs for } \tilde{\omega}) \times \mathbb{P}(\eta_0(x_1) = \cdots = \eta_0(x_k) = 0 \text{ for } \omega) > 0.
\end{aligned}$$

To conclude we claim that, on the above event, no particle ever visits $\mathbf{0}$ in the system ω. Indeed, on the above event, the initial configuration of ω is the same as that of $\tilde{\omega}$ except for the deletion of the particles present in $\{x_1, \dots, x_k\}$. In particular, all the particles which visit the origin in $\tilde{\omega}$ are deleted in ω. Recalling that ω and $\tilde{\omega}$ share the same putative trajectories, by following how the effect of deleting such particles propagates in the system evolution, one can see that in the system ω no particle can possibly visit $\mathbf{0}$, finishing the proof.

10.4. Fixation equivalence

In this subsection we prove Theorem 10.7. We use the construction provided by Theorem 10.6. Three implications are immediate: $(i) \Rightarrow (ii)$ by the 0-1 law in Theorem 2.7, $(ii) \Rightarrow (iii)$ because if some particle fixates then it has to fixate at some site, $(iii) \Rightarrow (iv)$ is trivial, so we only have to show $(iv) \Rightarrow (i)$.

Let $\mathcal{A}^{x,j}$ denote the event that particle (x, j) stays active, and write $\mathcal{A}^x = \mathcal{A}^{x,1}$. Assuming (iv) holds, $a := \mathbb{P}(\mathcal{A}^{\mathbf{0}}) > 0$. Indeed, (iv) implies that

$\mathbb{P}(\mathcal{A}^{\mathbf{0},j}) > 0$ for some j and, by interchangeability of particles, we have $\mathbb{P}(\mathcal{A}^{\mathbf{0},j}) = \mathbb{P}(\mathcal{A}^{\mathbf{0}}, |\eta_0(\mathbf{0})| \geqslant j) \leqslant \mathbb{P}(\mathcal{A}^{\mathbf{0}})$.

We make a side remark before giving more details. By the mass transport principle, the number N_t of particles which stay active and are present at site $\mathbf{0}$ at time t satisfies $\mathbb{E}N_t \geqslant a$, hence $\liminf_t \mathbb{E}N_t > 0$. But to show site activity we need $\limsup_t \mathbb{P}(N_t \geqslant 1) > 0$ instead. The idea is to introduce extra randomness so as to spread out the effect of these particles.

Since the system $\boldsymbol{\eta}$ is a measurable function of the randomness ω, for each $\varepsilon > 0$ there is $k \in \mathbb{N}$ such that the event $\mathcal{A}^{\mathbf{0}}$ can be ε-approximated by some event $\mathcal{A}^{\mathbf{0}}_\varepsilon$ that depends only on $(\eta_0(x), X^x, \mathcal{P}^x)_{\|x\| \leqslant k}$. Let $\mathcal{A}^x_\varepsilon$ denote the corresponding translation of the event $\mathcal{A}^{\mathbf{0}}_\varepsilon$. When $\mathcal{A}^x_\varepsilon$ occurs, we say that particle $(x,1)$ is a *candidate*. It is a *good candidate* if $\mathcal{A}^x$ also occurs, otherwise it is a *bad candidate*.

Fix $t > 0$. Let $n \in \mathbb{N}$ be a large number. The trick is to add more randomness to the system by choosing Z^y uniformly among the first n different sites in the putative trajectory $X^{y,1}$ after time t, independently over y. Define $\mathcal{C}(y,x)$ as the event that $\mathcal{A}^y_\varepsilon$ occurs and $Z^y = x$. Let

$$q(y,x) = \mathbb{P}\left(\mathcal{C}(y,x) \,\middle|\, \omega\right) \qquad \text{and} \qquad Q(x) = \sum_y q(y,x).$$

By the mass transport principle,

$$\mathbb{E}[Q(\mathbf{0})] = \sum_y \mathbb{P}(\mathcal{C}(y,\mathbf{0})) = \sum_y \mathbb{P}(\mathcal{C}(\mathbf{0},y)) = \mathbb{P}(\mathcal{A}^{\mathbf{0}}_\varepsilon) =: b > a - \varepsilon.$$

Notice that $q(y,x) \leqslant \frac{1}{n}$. Notice also that $q(y,x)$ and $q(z,x)$ are independent if $\|y - z\| > 2k$. Using these two facts, it follows that $\mathbb{V}[Q(\mathbf{0})]$ becomes small when n is large, and thus $Q(\mathbf{0})$ converges to b in distribution.

Let $N(x) = \sum_y \mathbb{1}_{\mathcal{C}(y,x)}$ count the number of candidates for which $Z^y = x$. Then

$$\mathbb{P}\left(N(\mathbf{0}) = 0 \,\middle|\, \omega\right) = \prod_y (1 - q(y,\mathbf{0})) \leqslant e^{-Q(\mathbf{0})} \to e^{-b}$$

in probability as $n \to \infty$. Also, let $\tilde{N}(x) = \sum_y \mathbb{1}_{\mathcal{C}(y,x) \setminus \mathcal{A}^y}$ count the number of bad candidates for which $Z^y = x$. Then, using the mass transport principle,

$$\mathbb{E}[\tilde{N}(\mathbf{0})] = \sum_y \mathbb{P}(\mathcal{C}(y,\mathbf{0}) \setminus \mathcal{A}^y) = \sum_y \mathbb{P}(\mathcal{C}(\mathbf{0},y) \setminus \mathcal{A}^{\mathbf{0}}) = \mathbb{P}(\mathcal{A}^{\mathbf{0}}_\varepsilon \setminus \mathcal{A}^{\mathbf{0}}) \leqslant \varepsilon.$$

Let $\mathcal{D}^x$ denote the event that there exists a good candidate $(y,1)$ such that $Z^y = x$. Using the two last estimates we get

$$\mathbb{P}(\mathcal{D}^{\mathbf{0}}) \geqslant \mathbb{P}(N(\mathbf{0}) \geqslant 1) - \mathbb{P}(\tilde{N}(\mathbf{0}) \geqslant 1) \geqslant 1 - e^{-a+\varepsilon} - \delta_n - \varepsilon,$$

where $\delta_n \to 0$ as $n \to \infty$. Choosing ε small and n large, we have $\mathbb{P}(\mathcal{D}^{\mathbf{0}}) > \frac{a}{2}$.

To conclude, notice that, on the event $\mathcal{D}^{\mathbf{0}}$, there is a particle $(y,0)$ which stays active, and some inner time $s > t$ such that $X^{y,1}_s = \mathbf{0}$, implying that site $\mathbf{0}$ is visited by an active particle after time t. Letting $t \to \infty$, we get $\mathbb{P}(\text{site } \mathbf{0} \text{ stays active}) \geqslant \frac{a}{2} > 0$, concluding the proof that $(iv) \Rightarrow (i)$.

11. Analysis of explicit constructions

An evolution $(\eta_t)_{t\geqslant 0}$ starting with only finitely many particles can be constructed explicitly in innumerous ways, using Poisson processes, exponential variables, random walks, tossing some coins, etc. Then we want to say that a system starting from an infinite random configuration η_0 exists and can be approximated in distribution by finite ones. Namely, denoting by $\mathbf{P}^\nu_V$ the law of the process starting from the finite truncation $\eta_0 \cdot \mathbb{1}_V$ for finite $V \subseteq \mathbb{Z}^d$, we wonder whether

$$\mathbf{P}^\nu\Big((\eta_t)_{t\geqslant 0} \in \mathcal{A}\Big) = \lim_{V\uparrow\mathbb{Z}^d} \mathbf{P}^\nu_V\Big((\eta_t)_{t\geqslant 0} \in \mathcal{A}\Big) \tag{11.1}$$

for every *local event* $\mathcal{A}$, *i.e.* every event $\mathcal{A}$ whose occurrence is determined by $(\eta_s(x))_{x\in B_k, s\in[0,t]}$ for some finite t and k.

Assume the limit on the right-hand side exists for some construction which is consistent with the rates specified in §2.1. Then the limit is obviously the same for any other construction consistent with §2.1. So if there exists a process $(\eta_t)_{t\geqslant 0}$ on a space $\mathbf{P}^\nu$ whose distribution satisfies (11.1), then its distribution is unique.

There are at least three ways to show existence of a $\mathbf{P}^\nu$ satisfying (11.1). One is to consider a certain norm on a subset of $(\mathbb{N}_\mathfrak{s})^{\mathbb{Z}^d}$ and use abstract theory of generators and semigroups adapted to non-compact spaces. Such a norm has to be more restrictive than product topology (indeed, one can always make many particles visit $\mathbf{0}$ in short time by placing enough particles far away), but it still gives (11.1) with a good level of generality on ν. Another way is to consider the particle-wise construction described in §10.1, which is well-defined as we prove in §11.3. The third way is to add Poisson clocks to the site-wise representation of §2.3, which is done in §11.2. These constructions work under the assumption that $\int |\eta(x)|\nu(\mathrm{d}\eta) \leqslant \zeta_{\max}$ for some finite $\zeta_{\max}$ uniformly over x.

11.1. Conditions for fixation and activity

In this subsection we prove Theorem 2.7 assuming that an explicit site-wise construction of the process satisfies (11.2), and that (11.3) holds.

We start with the 0-1 law. For almost every $\mathcal{I}$, if $m_\eta(\mathbf{0}) = \infty$ for a given configuration η, then $m_\eta(y) = \infty$ for all y with $p(y) > 0$. Write $W_p = \{z : p(z) > 0\} \subseteq \mathbb{Z}^d$. Let us omit the tedious proof of the following fact: if the elements of a set W generate the group $(\mathbb{Z}^d, +)$, then as a semigroup they generate a set that contains some $w + U \cap \mathbb{Z}^d$, where $U \subseteq \mathbb{R}^d$ is a cone with non-empty interior. By the previous remark, if $z \in -(w+U)$ and $m_{\eta_0}(z) = \infty$, then $m_{\eta_0}(\mathbf{0}) = \infty$. Assume that $\mathbb{P}^\nu\big(m_{\eta_0}(\mathbf{0}) = \infty\big) > 0$. As $-w-U$ contains balls of arbitrarily large radius, since $(\eta_0, \mathcal{I})$ is translation-ergodic, the $\mathbb{P}^\nu$-probability of finding a site $z \in -(w+U)$ with $m_{\eta_0}(z) = \infty$ is equal to 1, and therefore $\mathbb{P}^\nu\big(m_{\eta_0}(\mathbf{0}) = \infty\big) = 1$.

We now prove that $\mathbf{P}^\nu\big(\text{fixation of } (\eta_t)_{t\geqslant 0}\big) = \mathbb{P}^\nu(m_{\eta_0}(\mathbf{0}) < \infty)$. Let $h_t(x)$ denote the number of topplings at site x during the time interval $[0,t]$, meaning any action performed at x, including unsuccessful attempts to sleep. Write $h_\infty(x) = \lim_{t\to\infty} h_t(x)$. This limit exists as $h_t(x)$ is non-decreasing in t.

The core of the proof is to add some Poisson clocks to $\mathbb{P}^\nu$ and use it to construct $\mathbf{P}^\nu$ explicitly, so that

$$\mathbf{P}^\nu\big(h_\infty(x) \geqslant k\big) = \mathbb{P}^\nu\big(m_{\eta_0}(x) \geqslant k\big) \qquad \text{for each} \quad k > 0 \tag{11.2}$$

and use it to show that

$$\mathbf{P}^\nu\big(h_t(x) \geqslant k\big) \to 0 \quad \text{ as } \quad k\to\infty \qquad \text{for each fixed } t, \tag{11.3}$$

which is done in the next subsection.

Let us show that these imply the theorem. Assume $\mathbb{P}^\nu\big(m_{\eta_0}(x) < \infty\big) = 1$. It follows from (11.2) that $\mathbf{P}^\nu\big(h_t(x) \text{ eventually constant}\big) = 1$, thus x is eventually stable in η_t and in particular $\eta_t(x)$ remains bounded for large t. But $\eta_t(x)$ can only decrease when x is unstable, so $\mathbf{P}^\nu\big(\eta_t(x) \text{ converges}\big) = 1$.

Otherwise, $\mathbb{P}^\nu\big(m_{\eta_0}(x) = \infty\big) = 1$ by the 0-1 law, then (11.2) gives $\mathbf{P}^\nu\big(h_t(x) \to \infty \text{ as } t\to\infty\big) = 1$. Now by (11.3) we know that $\big(h_t(x)\big)_{t\geqslant 0}$ cannot blow up in finite time, whence for each x, the value of $\eta_t(x)$ changes for arbitrarily large times, and the system stays active.

11.2. The site-wise construction

In this subsection we provide a coupling $\mathbb{P}^\nu$ that produces $\mathbf{P}^\nu$ and all $\mathbf{P}^\nu_V$ on the same probability space, and show that (11.1) holds. We also show that this coupling satisfies (11.2) and (11.3). The translation-invariant distribution ν is fixed and will be omitted in the notation.

Start by adding Poisson clocks to the site-wise representation. More precisely, sample $\mathcal{I}$ following the distribution described in §2.3, sample η_0 according to the distribution ν, and sample an i.i.d. collection of Poisson point processes with intensity $(1+\lambda)\mathrm{d}t$, all independently. Let $\mathbb{P}$ denote the underlying probability.

For a finite deterministic initial configuration ξ, the evolution is constructed as follows. At $t=0$, let $L_0(x) = 0$ for all x. Fix $\eta'_t(x) = \xi(x)$ for all small t, and let $L_t(x)$ increase by $\frac{\mathrm{d}}{\mathrm{d}t}L_t(x) = \big(\eta_t(x)\big)\mathbb{1}_{\eta_t(t)\neq\mathfrak{s}}$. Denote the Poisson point process at each site x by $(T_n(x))_n$ where $T_0 = 0$ and $T_{n+1} - T_n$ are i.i.d. exponentials with parameter $1+\lambda$. Writing $h'_t(x) = \max\{n \in \mathbb{N}_0 : L_t(x) \geqslant T_n(x)\}$ for all t, let $\eta'_t(x)$ remain constant until the moment t_1 of the first jump of h'_t, which happens a.s. at a unique site y_1 that must be unstable for ξ. At this point, take $\alpha_1 = (y_1)$ and $\eta'_{t_1} = \Phi_{y_1}\xi$. Notice that $h'_{t_1} = m_{\alpha_1}$. Continue evolving L_t with the same rule, keeping $\eta'_t = \eta'_{t_1}$, until the moment t_2 of the next jump of h'_t, which happens a.s. at a unique site y_2, that again must be unstable for η'_{t_1}. As before, take $\alpha_2 = (y_1, y_2)$ and $\eta'_{t_2} = \Phi_{y_2}\eta'_{t_1} = \Phi_{\alpha_2}\xi$. Again $h'_{t_2} = m_{\alpha_2}$. Carry this procedure until $\eta'_t(x) = 0$ or $\mathfrak{s}$ for all x. After this time, L_t will be constant and the configuration will no longer change.

In this construction, at each time $t \geqslant 0$, η'_t is given by Φ_{α_j} for some j, α_j is a legal sequence of topplings for ξ, and $m_{\alpha_j} = h_t$. Hence, by the local Abelian property, η'_t can be read from ξ, $\mathcal{I}$ and h'_t. Moreover, the occupation times $L_t(x)$, and thus the toppling counter $h'_t(x)$, are increasing in the initial configuration ξ (proof below). We use η^V_t and h^V_t to denote the processes obtained by taking $\xi = \eta^V_0 = \eta_0 \cdot \mathbb{1}_V$. Then for each fixed x and t the counter $h^V_t(x)$ will be increasing in V, so it has a limit $h_t(x)$ that does not depend on the particular increasing sequence $V \uparrow \mathbb{Z}^d$.

Below we will show that, denoting by $N^V_t(x)$ the number of times that a particle jumps from some $z \neq x$ into x in the process $(\eta^V_s)_{s\in[0,t]}$, we have $\mathbb{E}[N^V_t(x)] \leqslant \zeta_{\max} \times t$.

Hence, the set of sites z such that $\mathfrak{t}^{z,k} = x$ for some $k \leqslant h_t(z)$ is finite and $h_t(x)$ is also finite. Now for each z in this set, $h^V_t(z)$ eventually equals $h_t(z)$ as $V \uparrow \mathbb{Z}^d$, and the same holds for $h^V_t(x)$. It thus follows from the local Abelian property that $\eta^V_t(x)$ will also be eventually constant, and we take $\eta_t(x)$ as the limit. Hence, almost surely, convergence holds simultaneously for all $x \in \mathbb{Z}^d$ and $t \in \mathbb{Q}_+$, and taking $\eta_s(x) = \lim_{t\downarrow s, t\in\mathbb{Q}} \eta_t(x)$ gives (11.1).

We now move on to the proof of (11.2). First, for every k,

$$\mathbb{P}\big(h^V_t(x) \geqslant k\big) \underset{V\uparrow\mathbb{Z}^d}{\longrightarrow} \mathbb{P}\big(h_t(x) \geqslant k\big) \underset{t\to\infty}{\longrightarrow} \mathbb{P}\big(h_\infty(x) \geqslant k\big).$$

We now show, on the other hand, that

$$\mathbb{P}\big(h^V_t(x) \geqslant k\big) \underset{t\to\infty}{\longrightarrow} \mathbb{P}\big(m_{\eta^V_0}(x) \geqslant k\big) \underset{V\uparrow\mathbb{Z}^d}{\longrightarrow} \mathbb{P}\big(m_{\eta_0}(x) \geqslant k\big).$$

The second limit follows from $m_{\eta_0} \geqslant m_{\eta^V_0} \geqslant m_{\eta_0,V} \to m_{\eta_0}$. Let us prove the first limit. Since finite configurations are a.s. stabilized after a finite number of topplings, there is some $t_* \geqslant 0$ such that $h^V_t = h^V_{t_*}$ for all $t \geqslant t_*$, and moreover $\eta^V_{t_*}$ is stable. Since h^V_t counts the number of topplings performed at each site up to time t, all of which are legal for η^V_0, we have $h^V_{t_*} = m_{\eta^V_0}$ by the Abelian property. By monotonicity the above limits commute, proving (11.2).

Now the tedious proof that $L_t(x)$ is non-decreasing in ξ. Let $\xi_1 \leqslant \xi_2$ be finite configurations. In order to show that these yield $L^1_t(x)$ and $L^2_t(x)$ satisfying $L^1 \leqslant L^2$, we will show that the set $G = \big\{t \geqslant 0 : L^1_s \leqslant L^2_s \text{ for all } s \leqslant t\big\} \ni 0$ is both open and closed on $[0,\infty)$. Since $t \mapsto L_t$ is continuous, G is closed. Let $t \in G$ and fix $x \in \mathbb{Z}^d$. If $L^1_t(x) < L^2_t(x)$, by continuity there is some $\varepsilon_x > 0$ such that $L^1_s(x) < L^2_s(x)$ for all $s \leqslant t + \varepsilon_x$. Otherwise $L^1_t(x) = L^2_t(x)$, which means $h^1_t(x) = h^2_t(x)$. At the same time, $h^1_t(z) \leqslant h^2_t(z)$ for $z \neq x$. Hence, by Property 2 in §2.2, $\eta^1_t(x) \leqslant \eta^2_t(x)$. Since $t \mapsto \eta^1_t(x)$ and $t \mapsto \eta^2_t(x)$ are piecewise constant and right-continuous, there is some $\varepsilon_x > 0$ such that, for all $s \in (t, t+\varepsilon_x]$, $\eta^1_s(x) \leqslant \eta^2_s(x)$, hence $\frac{\mathrm{d}}{\mathrm{d}s} L^1_s(x) \leqslant \frac{\mathrm{d}}{\mathrm{d}s} L^2_s(x)$, and thus $L^1_s(x) \leqslant L^2_s(x)$. Finally, since there are a.s. finitely many sites x such that $L_s^1(x) > 0$ for some s, we can take the $\varepsilon > 0$ as the smallest ε_x over all such x, so that $t + \varepsilon \in G$.

The last missing statement is that $\mathbb{E}[N_t^V(x)] \leqslant \zeta_{\max} \times t$. This was used in the proof of (11.1) and it also implies (11.3). Instead of making a rather convoluted argument, we resort to the finite particle-wise construction of §10.1 in terms of walks X and Y, which produces a process $(\eta_t^V)_{t\geqslant 0}$ with the same distribution. Using translation invariance and re-indexing sums,

$$\begin{aligned}
\mathbb{E}\big[N_t^V(x)\big] &= \mathbb{E}\Big[\sum_y \sum_{j=1}^{\eta_0(y)} \text{number of jumps of } Y^{y,j} \text{ into } x \text{ during } [0,t]\Big] \\
&\leqslant \mathbb{E}\Big[\sum_y \sum_{j=1}^{\eta_0(y)} \text{number of jumps of } X^{y,j} \text{ into } x \text{ during } [0,t]\Big] \\
&= \textstyle\sum_y \mathbb{E}|\eta_0(y)| \times \mathbb{E}\Big[\text{jumps of } X^{y,1} \text{ into } x \text{ during } [0,t]\Big] \\
&\leqslant \zeta_{\max} \textstyle\sum_y \mathbb{E}\Big[\text{number of jumps of } X^{y,1} \text{ into } x \text{ during } [0,t]\Big] \\
&= \zeta_{\max}\, \mathbb{E}\Big[\textstyle\sum_y \text{number of jumps of } X^{x,1} \text{ into } y \text{ during } [0,t]\Big] \\
&= \zeta_{\max} \times t,
\end{aligned}$$

which concludes the proof.

11.3. Well-definedness of the particle-wise construction

In this subsection we prove Theorem 10.6. The triple $(\eta_0, \mathbf{Y}, \boldsymbol{\gamma})$ describes the system by specifying the location and state of each labeled particle. An equivalent description is to specify the labels and states of all particles present at each site.

Consider a realization of the walks $\mathbf{X}$ and clocks $\boldsymbol{\mathcal{P}}$. For a finite initial configuration ξ, we consider the evolution obtained from $(\xi, \mathbf{X}, \boldsymbol{\mathcal{P}})$, and define

$$\bar{\eta}_t(z;\xi) = \Big\{(x,j,i) \in \mathbb{Z}^d \times \mathbb{N} \times \{1,\mathfrak{s}\} \,:\, j = 1,\dots,|\xi(x)|, Y_t^{x,j} = z, \gamma_t^{x,j} = i\Big\}.$$

Dependence of Y's and γ's on ξ is omitted in the notation. Our goal is to show that, almost surely, for each $z, y \in \mathbb{Z}^d$ and $t > 0$, $(\bar{\eta}_s(z;\eta_0^{B_n^y}))_{s\in[0,t]}$ is the same for all but finitely many n, and that the limiting process $(\bar{\eta}_t)_{t\geqslant 0}$ does not depend on y. Since each walk $X^{x,j}$ only visits finitely many sites by time t, this implies well-definedness in the sense of Definition 10.5.

A convenient observation is the following. It suffices to prove that, for every sequence of finite sets $W_n \uparrow \mathbb{Z}^d$, a.s. the process $(\bar{\eta}_s(z;\eta_0^{W_n}))_{s\in[0,t]}$ is the same for all but finitely many n. Indeed, assuming this holds true, a.s. it will be the case simultaneously for all the sequences $(B_n^y)_n$ as well as some deterministic increasing sequence $(W_n)_n$ that has infinitely many terms in common with each one of them. This in turn implies that the limit $(\bar\eta_s)_{s\in[0,t]}$ is a.s. the same for each of them, for every $t>0$.

To prove this a.s. eventually constant limit, we study how the "influence" of a particle addition propagates through the system by time t. We show that the set of sites z for which the configuration $\bar\eta_t(z)$ is affected for some $s\in[0,t]$ by a particle addition at x is stochastically dominated by a branching random walk started with a single particle at x.

We make this precise now. If $\xi(x)>0$, define the event $x \overset{\xi,t}{\rightsquigarrow} z$ that

$$\bar\eta_s(z;\xi-\delta_x,\mathbf{X},\boldsymbol{\mathcal{P}}) \neq \bar\eta_s(z;\xi,\mathbf{X},\boldsymbol{\mathcal{P}}) \text{ for some } s\in[0,t].$$

Define the random set

$$Z_t^x(\xi)=\big\{z\in\mathbb{Z}^d \,:\, x \overset{\xi,t}{\rightsquigarrow} z\big\},$$

which is the set of sites influenced during $[0,t]$ by the removal of the last particle at x from configuration ξ. We will later on prove that

$$\mathbb{P}\big(z\in Z_t^x(\xi)\big)\leqslant \mathbb{P}\big(U_t(z-x)\geqslant 1\big) \tag{11.4}$$

for all z,x,t,ξ, where $(U_t)_{t\geqslant 0}$ denotes the following branching process.

At $t=0$, let $U_0=\delta_{\mathbf{0}}$ be the configuration with a single particle at $\mathbf{0}$. For each $t>0$, a transition $U_t\to U_t+\delta_x$ occurs at rate $\lambda U_t(x)$, and a transition $U_t\to U_t+2\delta_y$ occurs at rate $\sum_x U_t(x)p(y-x)$. In words, each particle produces at rate λ a new copy at the same site, and at rate 1 two new copies at a site chosen at random. Particles never disappear in $(U_t)_t$.

Before showing (11.4), let us derive Theorem 10.6 by proving the a.s. eventual constant limit as $W_n\uparrow\mathbb{Z}^d$. We can assume that $|W_n|=n$ and write $W_n=\{x_1,\dots,x_n\}$. We want to rule out that, for infinitely many n,

$$\bar\eta_s(z;\eta_0^{W_n})\neq\bar\eta_s(z;\eta_0^{W_{n-1}}) \text{ for some } s\in[0,t]. \tag{11.5}$$

For each n, occurrence of the above event implies the occurrence of

$$z\in Z_t^{x_n}(\eta_0^{W_{n-1}}+k\delta_{x_n}) \text{ for some } k=1,\dots,\eta_0(x_n)$$

in case $\eta_0(x_n)\in\mathbb{N}_0$, or $z\in Z_t^{x_n}(\eta_0^{W_{n-1}}+\delta_{x_n}\cdot\mathfrak{s})$ if $\eta_0(x_n)=\mathfrak{s}$. For simplicity we ignore the $\mathfrak{s}$ case. The implication is true because $\eta_0^{W_{n-1}}$ can be obtained from $\eta_0^{W_n}$ by removing all the $|\eta_0(x_n)|$ particles at x, one by one, and in case none of them affects the configuration at site z by time t, the event in (11.5) cannot occur. Now the bound (11.4) holds for each fixed ξ, whence

$$\mathbb{P}\Big(z\in Z_t^{x_n}(\eta_0^{W_{n-1}}+k\delta_{x_n})\,\Big|\,|\eta_0(x_n)|\geqslant k\Big)\leqslant\mathbb{P}\big(U_t(z-x_n)\geqslant 1\big).$$

Thus,

$$\begin{aligned}
\mathbb{E}\Big[\#\big\{n \in \mathbb{N} : z \in Z_t^{x_n}(\eta_0^{W_{n-1}} + k\delta_{x_n}) \text{ for some } k \leqslant |\eta_0(x_n)|\big\}\Big] \\
\leqslant \sum_n \sum_k \mathbb{P}\big(|\eta_0(x_n)| \geqslant k, z \in Z_t^{x_n}(\eta_0^{W_{n-1}} + k\delta_{x_n})\big) \\
\leqslant \zeta_{\max} \sum_n \mathbb{E}\big[U_t(z - x_n)\big] \\
= \zeta_{\max}\, \mathbb{E}\big[\textstyle\sum_y U_t(y)\big] = \zeta_{\max} \times e^{(2+\lambda)t} < \infty.
\end{aligned}$$

Hence, by Borel-Cantelli, a.s. the event in (11.5) occurs for finitely many n, concluding the proof of Theorem 10.6.

In the rest of this subsection we prove (11.4). We aim at bounding the set of particles influenced before time t by the presence or absence of an extra particle.

Recall that particles behave independently except when one of them gets reactivated by another, including the case where it is prevented from deactivating. Hence, the influence only propagates in these cases, and only if the reactivation may not be attributed to the presence of a third, uninfluenced particle. Thus, a particle is influenced by the presence of an extra particle during $[0, t]$ if either it is the extra particle itself, or if it was reactivated at some time $s \leqslant t$ by sharing the same site with particles that had been themselves influenced before time s.

Once a particle is influenced, its inner time becomes uncertain. In order to keep track of every possibility and bound the presence of influenced particles, the idea is to consider all potential paths simultaneously. In some sense, we let each such particle both stay put at its current site as well as jump to the next one, by means of replicating such particles.

Suppose ξ^+ and ξ^- are finite and differ by the presence of a single particle at a given site. By translation invariance we can assume that this particle is $(\mathbf{0}, j^*)$, so $\xi^-(\mathbf{0}) = j^* - 1$ and $\xi^+(\mathbf{0}) = j^*$ (for simplicity, we omit the case where $j^* = 1$ and $\xi^+(\mathbf{0}) = \mathfrak{s}$). We write $Y_t^{x,j,\pm}$, $\gamma_t^{x,j,\pm}$ and $\bar{\eta}_t^{\pm}$ to denote the processes obtained from $(\xi^{\pm}, \mathbf{X}, \boldsymbol{\mathcal{P}})$.

Initially, the set of potentially affected particles is $R_0^{\times} = \{(\mathbf{0}, j^*)\}$ and the set of unaffected particles is $R_0^{\circ} = \{(x,j) : j = 1, \dots, |\xi^-(x)|\}$. Let $\mathcal{T}^{x,j}$ denote the time when particle (x,j) is removed from R_t° and added to $R_t^{\times}$. From this time on, clock rings from $\mathcal{P}^{x,j}$ will be ignored, this particle will jump normally according to $X^{x,j}$ and also leave a copy behind each time it jumps, so we can safely ignore its state $\gamma^{x,j}$. More precisely, we define $\sigma^{x,j,\times}(s) = s - \mathcal{T}^{x,j} + \sigma^{x,j,\pm}(\mathcal{T}^{x,j})$ and the increasing sets

$$D_t^{x,j} = \Big\{z \in \mathbb{Z}^d : X^{x,j}_{\sigma^{x,j,\times}(s)} = z \text{ for some } s \in [\mathcal{T}^{x,j}, t]\Big\}.$$

Finally, let D_t denote the union of $D_t^{x,j}$ over all $(x,j) \in R_t^{\times}$.

On the other hand, each unaffected particle $(x,j) \in R_t^{\circ}$ evolves normally and interacts with other particles $(x',j') \in R_t^{\circ}$ normally, until the first time $\mathcal{T}^{x,j}$ when it becomes affected. This will occur when (i) particle (x,j) is sleeping at some site z and site z is added to D_t, or (ii) the clock rings for particle (x,j) to sleep at some site $z \in D_t$ and there are no other particles $(x',j') \in R_t^\circ$ with $Y_t^{x',j'} = z$. So case (ii) is triggered by a sleep clock ring of particle (x,j) itself, whereas case (i) is triggered by the jump of an affected particle.

We then define

$$\bar{\eta}_t^\times(z) = \big\{(x,j,i) : (x,j) \in R_t^\times, i \in \{1,\sigma\}, z \in D_t^{x,j}\big\}$$

as well as

$$\bar{\eta}_t^\circ(z) = \{(x,j,i) \in \eta_t^+(z) : (x,j) \in R_t^\circ\} = \{(x,j,i) \in \eta_t^-(z) : (x,j) \in R_t^\circ\},$$

and note that

$$\bar{\eta}_t^\circ(z) \subseteq \bar{\eta}_t^\pm(z) \subseteq \bar{\eta}_t^\circ(z) \cup \bar{\eta}_t^\times(z).$$

The above inclusions imply that $Z_t^{\mathbf{0}}(\xi) \subseteq D_t$.

Hence, defining

$$\tilde{U}_t(z) = \#\big\{(x,j) \in R_t \,:\, z \in D_t^{x,j}\big\},$$

to get (11.4) it suffices to show that

$$\big(\tilde{U}_t\big)_{t\geqslant 0} \leqslant \big(U_t\big)_{t\geqslant 0} \text{ in law.} \tag{11.6}$$

To prove (11.6), it is enough to show that the transition rates of $(\tilde{U}_t)_t$ are always dominated by those of $(U_t)_t$. The process $(\tilde{U}_t)_t$ can increase in two situations. First, when only one particle $(x,j) \in R_t^\circ$ is present at some site $z \in D_t$ and its sleep clock rings. This transition causes $\tilde{U}_t(z)$ to increase by 1. The rate of this transition is λ at such sites and 0 elsewhere, so it is bounded by $\lambda\tilde{U}_t(z)$. Second, when a particle $(x,j) \in R_t^\times$ jumps, in which case it may add both a new site to $D_t^{x,j}$ and add a new particle (x',j') to $R_t^\times$. This transition causes $\tilde{U}_t(z)$ to increase by at most 2, and it occurs at rate bounded by $\sum_x \tilde{U}_t(x)p(z-x)$. This concludes the proof of (11.6), thus of (11.4) and hence of Theorem 10.6.

11.4. Monotonicity and the case of infinite sleep rate

In this subsection we discuss how to adapt the site-wise construction of §11.2 to produce processes with different values of λ. We use this to prove Theorem 2.14. We also describe the site-wise representation for the case $\lambda = \infty$ as mentioned in §10.3.

Recall the construction of §11.2. At each site x, we have a sequence $(\mathfrak{t}^{x,j})_{j\in\mathbb{N}}$ and a Poisson point process which we now denote $\mathcal{P}_x \subseteq [0,\infty)$. The pair $(\mathcal{P}_x, \mathfrak{t}^{x,\cdot})$ can be seen as a marked Poisson point process where $\mathfrak{t}^{x,j}$ is the mark of the j-th point in $\mathcal{P}_x$. Alternatively, we can obtain this marked process by merging two Poisson processes, $\mathcal{P}_x^*$ for jump and $\mathcal{P}_x^\lambda$ for sleep. In this construction, different values of λ can be coupled in a standard way, so that $\mathcal{P}_x^\lambda \subseteq \mathcal{P}_x^{\lambda'}$ for $\lambda' \geqslant \lambda$.

Similarly to the proof that $L_t(x)$ is non-decreasing in ξ, one can show that $L_t(x)$ is also non-increasing in λ (details omitted). Since fixation is a.s. equivalent to $L_t(\mathbf{0})$ remaining bounded as $t \to \infty$, ζ_c is non-decreasing in λ.

We conclude by discussing the case of $\lambda = \infty$.

As we increase λ, the sleep clock rings more and more often. In the limiting case $\lambda = \infty$, it is ringing permanently and $\mathcal{P}_x^\infty$ is a.s. dense on $[0,\infty)$. So the spontaneous transition $A \to S$ happens immediately, and only the process $\mathcal{P}_x^*$ is used in the construction. If other particles are present, the reaction $A + S \to 2A$ has priority and it overrides this urgency of particles to fall asleep. That is, $A \to S$ only occurs when an A-particle is alone. Sites with $n \geqslant 2$ particles send a particle away at rate n, and when the second last particle jumps out, the remaining particle falls asleep immediately. In order to define a site-wise representation with stacks $(\mathfrak{t}^{x,j})_{j\in\mathbb{N}}$ as in §2.2, we can consider the stacks with only the jump instructions, and declare particles to be sleeping whenever they are alone. The properties mentioned above still hold in this case, and in particular the proofs of Theorems 2.7 and 2.14 work for $\lambda \in [0,\infty]$.

11.5. A hybrid construction

In this subsection we prove Lemma 10.8. The inequality may seem obvious, and our intuition says it should be a trivial consequence of monotonicity properties in the spirit of the site-wise construction. However, in the unlabeled system it is not possible to distinguish the particles that have exited and re-entered V_n from the particles which have met them after their re-entrance. To solve this, in we will introduce a two-color site-wise construction.

Recall that n is fixed. For finite $V \supseteq V_n$, define

$$\eta_0^{\square}(x) = \begin{cases} \eta_0(x), & x \in V_n, \\ -\infty, & x \notin V_n, \end{cases} \qquad \eta_0^{V}(x) = \begin{cases} \eta_0(x), & x \in V, \\ 0, & x \notin V. \end{cases}$$

Let $M_t^{\square}$ and M_t^V count the number of labeled particles that exit V_n by time t in the system with initial configuration $\eta_0^{\square}$ and η_0^V. Lemma 10.8 then becomes

$$\mathbb{E} \lim_{t\to\infty} M_t^{\square} \leqslant \mathbb{E} \lim_{t\to\infty} \lim_{V\uparrow\mathbb{Z}^d} M_t^V$$

for some fixed sequence of finite boxes $V \uparrow \mathbb{Z}^d$, where the limit in V exists by Theorem 10.6. Since M_t^V is bounded by $\|\eta_0\|_V$, using dominated convergence theorem it is enough to show that $M_t^{\square}$ is stochastically dominated by M_t^V for each $t > 0$ and $V \supseteq V_n$ fixed. This will be shown with an explicit coupling.

In this construction, particles which started in V_n and have not yet exited V_n are colored purple, and all other particles are colored yellow. There are two stacks of instructions at each site, one for each color. This way one can distinguish the particles which have not yet exited V_n from those who have, and use this distinction to define M_t^V without the need to look at individual labels.

More precisely, we consider a two-colored particle system, where we are only interested in particle counts color by color: the configuration at time t is $\eta_t = (\eta_t^{\mathrm{p}}, \eta_t^{\mathrm{y}}) \in (\mathbb{N}_{\mathfrak{s}} \times \mathbb{N}_{\mathfrak{s}})^{\mathbb{Z}^d}$. Initially, purple particles are the particles that start on V_n, and yellow particles are the particles that start outside V_n. Purple particles become yellow when they exit V_n. Sample two independent families $\mathcal{I}^{\mathrm{p}}$ and $\mathcal{I}^{\mathrm{y}}$ of instructions, to be respectively used by purple particles and yellow particles (so we never use $\mathcal{I}^{\mathrm{p}}$ outside V_n). Sample two collections $\mathcal{P}^{\mathrm{p}}$ and $\mathcal{P}^{\mathrm{y}}$ of independent Poisson point processes attached to each site, and use them to trigger purple and yellow topplings, respectively.

Similarly to the construction of §11.2, this enables to construct the process $(\eta^{\mathrm{p}}, \eta^{\mathrm{y}})$ from a finite initial configuration, with the difference that two clocks now run at each site, at speeds given by the numbers of purple and yellow particles, and each clock triggers a toppling of the respective color. Furthermore, purple topplings have the additional effect of changing the color of the jumping particle if it jumps out of V_n. Note that topplings affect the state of particles of both colors because yellow and purple at same site share activity: a yellow particle can prevent a purple particle from falling asleep, and vice-versa, etc. This construction yields a natural coupling between systems with any finite initial configuration, each system alone having the same distribution as the one obtained from the particle-wise construction.

We remark that this bi-color construction is not Abelian and the local times $L_t^{\mathrm{p}}(x)$ are not generally monotone with respect to the initial configuration ξ. Nevertheless, since each color uses a different stack of instructions, adding yellow particles has the only effect, regarding purple particles, of enforcing activation of some of them at some times. The proof that $L_t(x)$ is non-decreasing in ξ given in §11.2 can be reproduced with nearly no modifications to show that $L_t^{\mathrm{p}}(x)$ is larger with initial configuration η_0^V than with $\eta_0^{\square}$. Since yellow topplings do not change the total number of purple particles present in the system, and purple topplings can only make that number decrease, the total number of purple particles present in the system after stabilization will be lower with initial configuration η_0^V than with $\eta_0^{\square}$. Hence, the number of particles that become yellow (*i.e.* the particles that ever exit V_n) is higher for η_0^V. So in this coupling $M_t^V \geqslant M_t^{\square}$ for every $V \supseteq V_n$, concluding the proof of Lemma 10.8.

12. Historical remarks and extensions

We conclude these lecture notes by making a historical account of how the main results in this field appeared in the literature, giving appropriate credit for the arguments presented in previous sections. We then comment on the extent at which these results can be generalized to other graphs, initial distributions, jump distributions etc. We finally mention some of these arguments that have meaningful counter-parts for the Stochastic Sandpile Model.

12.1. Historical remarks

Conditions (2.9) and (2.10) for the ARW were first used in [RS12]. They rely on a construction of the evolution which is based on the particle-wise discrete representation introduced in [DF91, Eri96], formalized in [RS12] and completed in [RT18]. Uniqueness of the critical density was proved in [RSZ19], on which §8 is based.

Results in dimension $d = 1$ appeared earlier. The case of directed (*i.e.* totally asymmetric) walks marks a special case. It was shown by [HS04] and published in [CRS14] that $\zeta_c = \frac{\lambda}{1+\lambda}$, and the proof of Theorem 3.2 is based on [CRS14]. The first result on fixation for general jump distributions appeared in [RS12], showing that $\zeta_c \geqslant \frac{\lambda}{1+\lambda}$. In terms of the phase space, $\zeta_c > 0$ for all λ, and $\zeta_c \to 1$ as $\lambda \to \infty$. Although it can be obtained via a much simpler argument as in §7, we give the original proof in §4 because: the proof can be adapted to more general graphs such as regular trees, it is used to study the fixed-energy dynamics in §6, and with simple modifications it gives $\zeta_c > \frac{\lambda}{1+\lambda}$ unless the walks are directed (Corollary 4.4).

Still in $d = 1$, using Condition (2.10) it is easy to show that there is no fixation at $\zeta = 1$, so in particular $\zeta_c \leqslant 1$ for every $\lambda \leqslant \infty$, see Theorem 3.1. It was shown in [Tag16] that, when the jump distribution is biased, $\zeta_c < 1$ for every $\lambda < \infty$, and $\zeta_c \to 0$ as $\lambda \to 0$. The toppling procedure was significantly simplified and extended to higher dimensions by the introduction of Condition (2.12) in [RT18], and further simplified using Theorem 2.13. The proof of Theorem 3.3 is adapted from [RT18], which was in turn adapted from [Tag16].

For the unbiased case, it was shown in [BGH18] that $\zeta_c \to 0$ as $\lambda \to 0$, and the proof given in §5 is adapted from [BGH18]. Existence of slow and fast regimes for the fixed-energy model on $\mathbb{Z}_n$ was proved in [BGHR19]. The proof given in §6 is adapted from [BGHR19], whose analysis is substantially based on results and arguments from [BGH18, Jan18, KV03, RS12].

In dimensions $d \geqslant 2$, it was shown that $\zeta_c \leqslant 1$ by [She10], using the technique of ghost walks previously introduced in the context of Internal DLA by [LBG92]. This technique was later used in [Tag16] who considers biased jump distributions and shows that $\zeta_c < 1$ for small λ. Ghost walks are usually useful for introducing independence in order to control variance and hence bootstrap from a certain counter having high expectation to being large in probability. The introduction of Condition (2.12) dispenses the use of ghost walks on $\mathbb{Z}^d$ or other amenable graphs, but this technique is still useful in the non-amenable setting, as used in [She10, ST18]. The use of ghost walks in [She10, Tag16] is surveyed in arXiv:1507.04341.

An alternative proof that $\zeta_c \leqslant 1$ was given in [AGG10], where the inequality follows from the mass transport principle after establishing the property that site fixation is equivalent particle fixation, under the assumption that the particle-wise construction of the process is well-defined. This assumption was proved to hold true in [RT18], where moreover the equivalence property was used to obtain Condition (2.12). As an application, it is shown that $\zeta_c < 1$ for all $\lambda < \infty$ and $\zeta_c \to 0$ as $\lambda \to 0$ for biased walks in any dimension. The construction in §10.1 as well as the arguments of §10.2 and §11.3 are adapted from [RT18]. The proofs in §10.3 and §10.4 are adapted from [CRS14, CRS18] and [AGG10], respectively. The presentation and proofs in §11.1 and §11.2 are adapted from [RS12] following an important observation from [RT18].

The technique of weak stabilization was introduced in [ST18], where it was shown that, in $d \geqslant 3$, $\zeta_c < 1$ for λ small enough, and $\zeta_c \to 0$ as $\lambda \to 0$. This result was strengthened in [Tag19] where it was shown that $\zeta_c < 1$ for all $\lambda < \infty$. The idea of strong stabilization was already implicit in [ST18] but was formalized and better exploited in [Tag19] under a different name. The proofs in §7 are adapted from these two articles.

The first proof of fixation for $d \geqslant 2$ appeared in [She10], who showed that $\zeta_c > 0$ when $\lambda = \infty$ using results from [Mar02] about greedy lattice animals. This was extended in [CRS14, CRS18] who also consider the extreme case $\lambda = \infty$ and show that $\zeta_c = 1$. In [ST17], it was shown that when the jump distribution is unbiased, $\zeta_c > 0$ for every $\lambda > 0$. This was significantly extended in [ST18], where the idea of weak stabilization was used to show that $\zeta_c \geqslant \frac{\lambda}{1+\lambda}$ in general.

Many of the problems listed in [DRS10] are still open. From the enormous gap between the description given in §1.3 and the results compiled in §1.5, the reader can find countless challenging open problems in this topic. Besides the problems already listed throughout these notes, we remark that more direct or otherwise insightful alternatives to some of the arguments presented here are certainly welcome.

12.2. Extension to other graphs and settings

Properties of the site-wise representation stated in §2.2 are deterministic and hold for any graph.

Many proofs about fixation and activity involved re-indexing a sum and thus rely on translation invariance. A more general context is that of jump distributions invariant under a transitive unimodular graph automorphism subgroup with infinite orbits, or *unimodular walks* for short. Translation-invariant walks on $\mathbb{Z}^d$ and uniform nearest-neighbor walks on regular trees or other Cayley graphs are good examples. Ergodicity in this case should be understood with respect to this same subgroup.

Well-definedness in the sense of (11.1) as provided through §11.2 and §11.3 still works for unimodular walks. Theorems 2.7, 2.13 and 10.7 hold with the same or essentially the same proofs. Theorem 2.11 remains true [RT18] but requires the graph to be amenable, and it is false on trees.

Theorem 3.1 remains true for unimodular walks through the proof of Theorem 10.4. The proof of Theorem 3.3 can also be extended directly to unimodular walks on amenable graphs [RT18] and partially to non-amenable graphs using ghost walks [ST18].

The argument shown in §4 can be adapted to trees [ST18]. The argument for the universal bound $\zeta_c \geqslant \frac{\lambda}{1+\lambda}$ shown in §7.1 works on any amenable graph, and it is possible to extended it to non-amenable graphs. Proofs of Theorems 7.1

and 7.2 are given in [ST18, Tag19] assuming uniform nearest-neighbor walks, but they also work for unimodular walks.

In §6 we use Poisson thinning which shortens the argument a bit, but other non-Poisson initial distributions with good concentration inequalities should suffice with little extra work.

The proof that $\zeta_c \geqslant 1$ given in [She10] in fact shows activity for deterministic initial configurations having empirical average larger than unit, and does not even assume any symmetries on the graph besides having bounded degree.

The assumption of nearest-neighbor jump simplified the exposition, but it is really needed only in §4, §5 and §6. For these sections, dropping such assumption would decrease the range of parameters for which the proofs work, or make the constructions more complicated, or both. The argument briefly outlined in §9 extends directly to a bounded-range or very light-tail jump distribution, as long as it remains unbiased. The arguments shown in the other sections do not require any assumption on the jump distribution besides translation invariance (or unimodularity). In this case, the notion of a walk being directed is the one stated in the proof of Corollary 3.4.

12.3. Stochastic Sandpile Model

For the Stochastic Sandpile Model (SSM), there is a continuous-time process with properties analogous to those of §2.2, see [RS12]. Theorem 2.7 holds with a similar proof. The proof of Theorems 10.1, 10.6 and 8.1, and hence that of Theorem 2.11, can be adapted to the SSM without substantial changes.

In one dimension, it has been shown that $\zeta_c \geqslant \frac{1}{4}$ using a toppling procedure similar to the one shown in §4, see [RS12]. Also, it should not be too complicated to combine the toppling procedure from Section 6 of [RS12] with a urn process similar to Lemma 6.7 to show fast fixation for $\zeta < \frac{1}{4}$.

For higher dimensions, the multi-scale argument described in §9 works without any modification to show that $\zeta_c > 0$ for the SSM with unbiased walks, see [ST17]. The proof of Theorem 7.1 can be partially adapted to show that $\zeta_c > 0$ for more general jump distributions.

To the best of our knowledge, this is all that is known rigorously about the SSM, as far as predictions in the spirit of §1.3 are concerned.

Acknowledgments

These lecture notes developed throughout the years in parallel with the literature, and were used in courses given at Institute for Advanced Study of Aix-Marseille University, University of Buenos Aires, Chinese Academy of Sciences, Henan University, National University of Singapore, and Peking-Tsinghua-Beijing Normal Universities joint seminar. Much of the current narrative originated from discussions with mathematicians and physicists, especially Deepak Dhar and P K Mouhanti during the workshop Universality in Random Structures: Interfaces, Matrices, Sandpiles, held on the International Centre for Theoretical Sciences, Bangalore, 2019 (Code: ICTS/urs2019/01). These events were supported by A*MIDEX grant ANR-11-IDEX-0001-02 and the National Natural Science Foundation of China grant 11688101.

The author's current understanding of Activated Random Walks and related models is the result of discussions and exchanges with many colleagues throughout more than a decade, it would be hard to name them all. Let us just mention Vladas Sidoravicius, Ronald Dickman, Manuel Cabezas, Lionel Levine and Laurent Tournier among those who had the deepest impact.

While these notes were being finalized, Shirshendu Ganguly and Moumanti Podder contributed many remarks that improved the clarity and readability. Laurent Tournier and Vittoria Silvestri went through most of these notes and made a large number of keen remarks that substantially improved the presentation. Special thanks to Alexandre Stauffer and Jacob Richey who read the entire material thoroughly and made thoughtful comments and suggestions, without which the quality of these notes would not have been the same. All the mistakes, omissions or obscure explanations that remain are the author's responsibility.

Update

Some of the open problems mentioned in these notes or similar ones have been solved in [ASR19, CR20, HRR20, PR20, Tag20].

References

[AGG10] G. Amir, O. Gurel-Gurevich. *On fixation of activated random walks.* Electron. Commun. Probab. **15**:119–123, 2010. MR2643591

[ASR19] A. Asselah, B. Schapira, L. T. Rolla. *Diffusive bounds for the critical density of activated random walks*, 2019. Preprint. arXiv:1907.12694.

[BGH18] R. Basu, S. Ganguly, C. Hoffman. *Non-fixation for conservative stochastic dynamics on the line*. Comm. Math. Phys. **358**:1151–1185, 2018. MR3778354

[BGHR19] R. Basu, S. Ganguly, C. Hoffman, J. Richey. *Activated random walk on a cycle*. Ann. Inst. Henri Poincaré Probab. Stat. **55**:1258–1277, 2019. MR4010935

[CR20] M. Cabezas, L. T. Rolla. *Avalanches in critical activated random walks*, 2020. Preprint. arXiv:2008.05783.

[CRS14] M. Cabezas, L. T. Rolla, V. Sidoravicius. *Non-equilibrium phase transitions: Activated random walks at criticality*. J. Stat. Phys. **155**:1112–1125, 2014. MR3207731

[CRS18] M. Cabezas, L. T. Rolla, V. Sidoravicius. *Recurrence and density decay for diffusion-limited annihilating systems*. Probab. Theory Relat. Fields **170**:587–615, 2018. MR3773795

[DF91] P. Diaconis, W. Fulton. *A growth model, a game, an algebra, Lagrange inversion, and characteristic classes*. Rend. Sem. Mat. Univ. Politec. Torino **49**:95–119, 1991. pdf. MR1218674

[DRS10] R. Dickman, L. T. Rolla, V. Sidoravicius. *Activated random walkers: Facts, conjectures and challenges*. J. Stat. Phys. **138**:126–142, 2010. MR2594894

[Eri96] K. Eriksson. *Chip firing games on mutating graphs*. SIAM J. Discrete Math. **9**:118–128, 1996. MR1375419

[HRR20] C. Hoffman, J. Richey, L. T. Rolla. *Active phase for activated random walk on Z*, 2020. Preprint. arXiv:2009.09491.

[HS04] C. E. Hoffman, V. Sidoravicius, 2004. Unpublished.

[Jan18] S. Janson. *Tail bounds for sums of geometric and exponential variables*. Statist. Probab. Lett. **135**:1–6, 2018. MR3758253

[JR19] T. Johnson, L. T. Rolla. *Sensitivity of the frog model to initial conditions*. Electron. Commun. Probab. **24**:29, 2019. MR3962479

[KL16] D. Kerr, H. Li. *Ergodic Theory: Independence and Dichotomies*. Springer Monographs in Mathematics. Springer, Cham, 2016. MR3616077

[KV03] J. F. C. Kingman, S. E. Volkov. *Solution to the OK Corral model via decoupling of Friedman's urn*. J. Theoret. Probab. **16**:267–276, 2003. MR1956831

[LBG92] G. F. Lawler, M. Bramson, D. Griffeath. *Internal diffusion limited aggregation*. Ann. Probab. **20**:2117–2140, 1992. MR1188055

[LP16] R. Lyons, Y. Peres. *Probability on Trees and Networks*, vol. 42 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, New York, 2016. MR3616205

[Mar02] J. B. Martin. *Linear growth for greedy lattice animals*. Stochastic Process. Appl. **98**:43–66, 2002. MR1884923

[PR20] M. Podder, L. T. Rolla. *Uniform threshold for fixation of the stochastic sandpile model on the line*, 2020. Preprint. arXiv:2001.04268.

[RS12] L. T. Rolla, V. Sidoravicius. *Absorbing-state phase transition for driven-dissipative stochastic dynamics on Z*. Invent. Math. **188**:127–150, 2012. MR2897694

[RSZ19] L. T. Rolla, V. Sidoravicius, O. Zindy. *Universality and sharpness in activated random walks*. Ann. Henri Poincaré **20**:1823–1835, 2019. MR3956161

[RT18] L. T. Rolla, L. Tournier. *Non-fixation for biased activated random walks*. Ann. Inst. H. Poincaré Probab. Statist. **54**:938–951, 2018. MR3795072

[She10] E. Shellef. *Nonfixation for activated random walks*. ALEA Lat. Am. J. Probab. Math. Stat. **7**:137–149, 2010. pdf. MR2651824

[ST17] V. Sidoravicius, A. Teixeira. *Absorbing-state transition for stochastic sandpiles and activated random walks*. Electron. J. Probab. **22**:33, 2017. MR3646059

[ST18] A. Stauffer, L. Taggi. *Critical density of activated random walks on transitive graphs*. Ann. Probab. **46**:2190–2220, 2018.MR3813989

[Tag16] L. Taggi. *Absorbing-state phase transition in biased activated random walk*. Electron. J. Probab. **21**:13, 2016.MR3485355

[Tag19] L. Taggi. *Active phase for activated random walks on Z^d, $d \geqslant 3$, with density less than one and arbitrary sleeping rate*. Ann. Inst. Henri Poincaré Probab. Stat. **55**:1751–1764, 2019. MR4010950

[Tag20] L. Taggi. *Essential enhancements in abelian networks: continuity and uniform strict monotonicity*, 2020. Preprint. arXiv:2003.00932.