

Absorbing-State Phase Transition for Driven-Dissipative Stochastic Dynamics on \mathbb{Z}

Leonardo T. Rolla and Vidas Sidoravicius

ENS-Paris and IMPA

August 30, 2011

Abstract We study the long-time behavior of conservative interacting particle systems in \mathbb{Z} : the activated random walk model for reaction-diffusion systems and the stochastic sandpile. We prove that both systems undergo an absorbing-state phase transition.

This preprint has the same numbering for sections, theorems, equations and figures as the published article “*Invent. Math.* 188 (2012): 127–150.”

Keywords absorbing-state phase transition · Diaconis-Fulton representation · reaction-diffusion · interacting particle systems · self-organized criticality

Mathematics Subject Classification (2000) 60K35 · 82C20 · 82C22 · 82C26

1 Introduction

Modern Statistical Mechanics offers a large and important class of driven-dissipative lattice systems that naturally evolve to a critical state, which is characterized by power-law distributions of the sizes of relaxation events (a paradigm example is the emergence of avalanches caused by small perturbations). In many mathematically interesting and physically relevant cases such systems are attracted to a stationary critical state without being specifically tuned to a critical point. In particular, it is believed that this phenomenon lies behind random fluctuations at the macroscopic scale, and creation of self-similar shapes in a variety of growth systems.

L. T. Rolla

École Normale Supérieure, 45, rue d’Ulm - Paris 75005 - France

Instituto de Matemática Pura e Aplicada, Estrada Dona Castorina, 110 - Rio de Janeiro 22.460-320 - Brasil

V. Sidoravicius

Instituto de Matemática Pura e Aplicada, Estrada Dona Castorina, 110 - Rio de Janeiro 22.460-320 - Brasil

Due to strong non-locality of correlations and dynamic long-range effects, classical analytic and probabilistic techniques fail in most cases of interest, making the rigorous analysis of such systems a major mathematical challenge.

Studies of the above phenomenon are confined to very few models, and its conceptual understanding is extremely fragmented. Among theories which attempt to explain long-ranged spatio-temporal correlations, the physical paradigm called ‘self-organized criticality’ takes its particular place [6,20]. These are systems whose natural dynamics drives them towards, and then maintains them at the *edge* of stability [6]. However, for non-equilibrium steady states it is becoming increasingly evident that ‘self-organized criticality’ is related to conventional critical behavior, namely that of an *absorbing-state phase transition*. The known examples are variations of underlying non-equilibrium systems which actually do have a parameter and exhibit critical phenomena. The phase transition in these systems arises from a conflict between a spread of activity and a tendency for this activity to die out [8,28], and the transition point separates an active and an absorbing phase in which the dynamics gets eventually extinct in any finite region.

In this paper we focus on two chief examples of conservative, infinite-volume systems which belong to the above mentioned family: the activated random walk model for reaction-diffusion and the stochastic sandpile model. Let us briefly describe the models (see Section 2 for the precise definitions).

The reaction-diffusion model is given by the following conservative particle dynamics in \mathbb{Z}^d . Each particle can be in one of two states: an active A -state, and a passive S -state. Particles in the A -state perform a continuous-time random walk with jump rate $D_A = 1$ without interacting. Particles in the S -state do not move, that is, $D_S = 0$. Each particle changes its state $A \rightarrow S$ at some halting rate $\lambda > 0$ and the reaction $A + S \rightarrow 2A$ happens immediately. The catalyzed transition $A + S \rightarrow 2A$ and the spontaneous transition $A \rightarrow S$ represent the spread of activity versus a tendency for this activity to die out. This system will be referred to as the model of Activated Random Walks (ARW).¹

In sandpile models the state of the system is represented by the number of particles $\eta(x) = 0, 1, 2, \dots$ at each vertex $x \in \mathbb{Z}^d$. The vertex x is *stable* when $\eta(x) < N_c$, for some threshold value N_c , and *unstable* when $\eta(x) \geq N_c$. Relaxation (update of the state) happens by toppling each unstable vertex, i.e., sending particles to its neighbors following a certain (deterministic or stochastic) rule. In this paper we study the following sandpile dynamics, which is a variation of Manna’s model frequently considered in the physics literature [5,29]. The threshold for stability of vertices is $N_c = 2$, and each unstable vertex topples after an exponentially-distributed time, sending 2 particles to neighbors chosen independently at random. We will refer to this model as the Stochastic Sandpile Model (SSM).²

¹ The Activated Random Walks may be viewed as a special case of a diffusive epidemic process. The model was introduced in the late 1970’s by F. Spitzer, but due to its tremendous technical difficulties and complexity, remained unsolved until recently, when it was studied in detail in [23–26]. This process has also been studied via renormalization group techniques and numerical simulation [16–18,22,27,31,37].

² In the so-called *Abelian Sandpile*, $N_c = 2d$, and an unstable vertex *deterministically* sends one particle to each of its $2d$ neighbors when toppling [3,4].

For these systems, the relation between self-organized and ordinary criticality is understood as follows. On the one hand, ‘self-organized criticality’ appears in their parameter-free,³ finite-volume variation: particles are added to the bulk of a finite box, and absorbed at its boundary during relaxation. The particle addition happens at a *slow* rate, or *only after* the system globally stabilizes. In this dynamics, when the average density μ 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 is absorbed at the boundary. With this *carefully designed* mechanism, the model is attracted to a *critical state* with an average density given by $0 < \mu_c < \infty$, though it was not explicitly tuned to this critical value. On the other hand, the corresponding conservative systems in infinite volume exhibit ordinary criticality in the sense that their dynamics fixate for $\mu < \mu_c$ and do not fixate for $\mu > \mu_c$, and moreover the *critical exponents* of the finite-volume addition-relaxation dynamics are related to those of the conservative dynamics in infinite volume.

The *critical behavior* of stochastic sandpiles seems to belong to the same *universality class* as the depinning of a linear elastic interface subject to random pinning potentials,⁴ roughly depicted by a nailed carpet being detached from the floor by an external force of critical intensity, where the rupture of each nail induces other ruptures nearby, giving rise to “avalanches”.

The deterministic sandpile defines a universality class *sui generis*, and is marked by strong non-ergodic effects [8].⁵ This seems to be due to the failure of the toppling procedure in eliminating certain microscopic symmetries in the configuration by the time the system becomes unstable, as a consequence of the existence of many toppling invariants. At the same time, the stationary state of the deterministic sandpile is the uniform distribution over a special subset of configurations exhibiting several combinatorial properties which, together with its strong algebraic structure, is one of the reasons why most of the mathematical literature has focused on this model (see [32] and references therein).

In the stochastic sandpile models, by contrast, the addition-relaxation operators are themselves random, leading to a set of coupled polynomial equations [6,35]. Their explicit solutions are not known in a general form, and very little can be said rigorously about the phase transition of such systems. The reaction-diffusion dynamics, in turn, might be even more apt to dispel microscopic symmetries, for not only the moves are random, but also the particles jump individually rather than in pairs.

The main pursuit in this framework is to describe the critical behavior, the scaling relations and critical exponents, and whether the critical density is the same as the long-time limit attained in the driven-dissipative version. These questions are however far beyond the reach of current techniques.

The contribution of this paper is a step forward in the mathematical understanding of these systems: we develop a general method for the analysis of diffusive-dissipative

³ Consider the *density of particles* as a parameter, and let λ be just a fixed number.

⁴ This prediction is supported by simulations, at least for dimensions $d \geq 2$, see [8].

⁵ For instance, it is known [15,30,32] that there is local fixation for $\mu < d$ and there is no fixation for $\mu > 2d - 1$, but one can construct somewhat artificial ergodic states with $\mu \in (\mu_c, 2d - 1)$ that fixate and there are as well states with $\mu \in (d, \mu_c)$ that do not fixate [15]. Furthermore, as observed in [13,14], the density μ_s attained in the large-time limit for the driven-dissipative system does not even satisfy $\mu_s = \mu_c$.

dynamics, and establish existence of the absorbing-state phase transition for both models in the one-dimensional case. Namely, we show that the system fixates if it starts with small enough density. Remark that this question remains open for $d \geq 2$ in spite of considerable efforts in the probability and mathematical physics communities.

A word about the proof. We build upon the Diaconis-Fulton representation to exploit the combinatorial nature of the problem. Due to particle exchangeability, this representation extracts precisely the part of the randomness that is relevant for the 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 precludes the analysis of quantities for which the order of the jumps does matter, such as correlation functions or local shape properties.

We introduce a tool to modify the Diaconis-Fulton instructions, which consists in forcing stable particles to move whereas in the original dynamics they would stand still, and we show that the total amount of activity is non-decreasing with respect to this change. For the ARW model this means suppressing some $A \rightarrow S$ transitions and for the SSM it is done via extra “semi-legal” single-particle topplings.

The instruction manipulations give broad possibilities of handling the configuration as it evolves, for one can keep the particles moving until they arrive at more convenient locations. We achieve fixation by successively settling the particles at appropriate places. In this procedure we set up a trap for each particle, yielding a globally stable configuration. We finally show that this scheme has positive probability of working indefinitely.

An important step is to look ahead in the future before deciding which semi-legal operations should be performed. It is crucial though that the Diaconis-Fulton instructions are explored in a careful way, in order to keep independence. This approach allows a fine control over the motion of each individual particle.

This paper is organized as follows. In Section 2 we give precise definitions of the models and statements of the results. In Section 3 we study the Diaconis-Fulton representation and relate stability of the initial configuration to local fixation of the system. In Section 4 we explain the general strategy used in the proofs and give an overview of the main ideas. In Sections 5 and 6 we prove the phase transition respectively for the ARW and SSM. In Section 7 we discuss extensions of our results and conclude with open problems.

2 The models and results

Stochastic Sandpile Model. When a site x has at least 2 particles, it is called *unstable* and topples at rate 1. When it topples it sends 2 particles to neighboring sites chosen independently at random, that is, according to the distribution $p(y-x)$, where $p(z) = \frac{1}{2d}$ if $\|z\| = 1$ and 0 otherwise.

The state of the SSM at each time $t \geq 0$ is given by $\eta_t \in (\mathbb{N}_0)^{\mathbb{Z}^d}$, where $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ and $\eta_t(x)$ denotes the number of particles found at site x at time t . For each

site $x \in \mathbb{Z}^d$, the transitions $\eta \rightarrow \tau_{xy}\tau_{xw}\eta$ happen at rate $A(\eta_t(x))p(y-x)p(w-x)$, where

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

and $A(k) = \mathbb{1}_{k \geq 2}$ indicates whether or not site x is unstable. Let \mathbb{P}^ν denote the law of $(\eta_t)_{t \geq 0}$ starting from η_0 distributed as ν . This evolution is well defined because the jump rates are bounded.

We say that the system *locally fixates* if $\eta_t(x)$ is eventually constant for each x , otherwise we say that the system *stays active*.

Theorem 1. *Consider the Stochastic Sandpile Model in the one-dimensional lattice \mathbb{Z} , with initial distribution ν given by i.i.d. Poisson random variables with parameter μ . There exists $\mu_c \in [\frac{1}{4}, 1]$ such that the system locally fixates a.s. if $\mu < \mu_c$, and stays active a.s. if $\mu > \mu_c$.*

Activated Random Walk model. Each particle in the A -state performs a continuous-time random walk with jump rate $D_A = 1$. The jumps have a probability density $p(\cdot)$ on \mathbb{Z}^d such that the set $\{z \in \mathbb{Z}^d : p(z) > 0\}$ generates the whole group $(\mathbb{Z}^d, +)$. Independently of anything else, each particle in the A -state turns to the S -state at a halting rate $\lambda > 0$. Once a particle is in the S -state, it stops moving, i.e., its jump rate is $D_S = 0$, and it remains in the S -state until the instant when another particle is present at the same vertex. At such an instant the particle which is in S -state flips to the A -state, giving the transition $A + S \rightarrow 2A$. A particle in the S -state stands still forever if no other particle ever visits the vertex where it is located. According to these rules, the transition $A \rightarrow S$ *effectively* occurs if and only if, at the instant of such a transition, the particle *does not share* its vertex with another particle (the innocuous instantaneous transition $2A \rightarrow A + S \rightarrow 2A$ is not observed). Particles in the A -state do not interact among themselves.

The state of the ARW at time $t \geq 0$ is given by $\eta_t \in \Sigma = (\mathbb{N}_{0\rho})^{\mathbb{Z}^d}$, where $\mathbb{N}_{0\rho} = \mathbb{N}_0 \cup \{\rho\}$. In this setting $\eta_t(x)$ denotes the *number of particles* at site x at time t , i.e., the *number of active particles* if $\eta_t(x) \in \mathbb{N}$, and the *one passive particle* if $\eta_t(x) = \rho$. We turn $\mathbb{N}_{0\rho}$ into an ordered set by setting $0 < \rho < 1 < 2 < \dots$. We also let $|\rho| = 1$, so $|\eta_t(x)|$ counts the number of particles regardless of their state. The addition is defined by $\rho + 0 = 0 + \rho = \rho$ and $\rho + n = n + \rho = 1 + n$ for $n > 0$, that is, the addition operation already provides the $A + S \rightarrow 2A$ transition. The $A \rightarrow S$ transition is represented by $\rho \cdot n$ given by $\rho \cdot 1 = \rho$ and $\rho \cdot n = n$ for $n > 1$. Subtractions involving ρ are not defined.

The process evolves as follows. For each site x , we have the transitions $\eta \rightarrow \tau_{xy}\eta$ at rate $A(\eta_t(x))p(y-x)$ and $\eta \rightarrow \tau_{x\rho}\eta$ at rate $\lambda A(\eta_t(x))$, where

$$\tau_{x\rho}\eta(z) = \begin{cases} \rho \cdot \eta(x), & z = x \\ \eta(z), & \text{otherwise} \end{cases}$$

and $A(k) = k\mathbb{1}_{k \geq 1}$, so $A(\eta_t(x))$ is the number of active particles at site x at time t . We denote by ν the distribution of η_0 and assume that $\eta_0(x) \in \mathbb{N}_0$ for all x a.s.

We further write ν_M for the distribution of the truncated configuration η^M given by $\eta^M(x) = \eta_0(x)$ for $|x| \leq M$ and $\eta^M(x) = 0$ otherwise, and $\mathbb{P}_M^\nu = \mathbb{P}^{\nu_M}$. \mathbb{P}_M^ν is well defined and corresponds to the evolution of a countable-state Markov chain whose configurations contain only finitely many particles.

It follows from a construction due to Andjel⁶ that, if ν is a product measure with density $\nu(|\eta(0)|) < \infty$ then \mathbb{P}^ν is well defined and, moreover,

$$\mathbb{P}^\nu(E) = \lim_{M \rightarrow \infty} \mathbb{P}_M^\nu(E) \quad (1)$$

for every event E that depends on a finite space-time window, i.e., that is measurable with respect to $(\eta_s(x) : |x| < t, s \in [0, t])$ for some $t < \infty$.

Theorem 2. *Consider the Activated Random Walk Model with nearest-neighbor jumps in the one-dimensional lattice \mathbb{Z} with halting rate λ and with initial distribution ν given by i.i.d. Poisson random variables in \mathbb{N}_0 with parameter μ . There exists $\mu_c \in [\frac{\lambda}{1+\lambda}, 1]$ such that the system locally fixates a.s. if $\mu < \mu_c$ and stays active a.s. if $\mu > \mu_c$.*

Theorem 2 should be contrasted with the particular case of totally asymmetric jumps, for which it is known [21] that $\mu_c = \frac{\lambda}{1+\lambda}$, so the lower bound is sharp. Note also that $\mu_c = 1$ when $\lambda = +\infty$. Theorems 1 and 2 remain true under more general hypotheses. See Section 7 for details.

3 The Diaconis-Fulton representation

In this section we describe the Diaconis-Fulton graphical representation for the dynamics of both SSM and ARW. The main advantage of this representation is that it leaves aside the chronological order of topplings. This will be justified *a posteriori* by showing that local fixation for the stochastic evolution is equivalent to stability properties of the discrete operations described hereafter.

Graphical constructions have become standard tools in Probability. They are usually attributed to Harris, see [19] and references therein to backtrack its history. In the construction described below, a sequence of random instructions is assigned to each site, but unlike Harris' construction, the instant when each step should be performed is not part of the representation and is determined by some external factor. As far as we know, the earliest references describing this kind of construction are the card game of Diaconis and Fulton [7] and the general framework considered by Eriksson [11], which includes the card game, the Abelian Sandpile Model, the Stochastic Sandpile Model, and the Activated Random Walks.

We first describe the *representation*, defining \mathcal{I} , \mathcal{P} , h , Φ , instructions, stable sites and legal operations. We then stress some *local* properties in an *abstract*, model-independent framework, defining m_α , Φ_α , and legal sequences. We finally consider

⁶ Let $\Sigma'' = \{\eta \in \Sigma : \sum_x |\eta(x)| < \infty\}$, define $\alpha(x) = \sum_{n=0}^{\infty} 2^{-n} p^{(n)}(x, 0)$ and $\|\eta\| = \sum_{x \in \mathbb{Z}^d} |\eta(x)| \alpha(x)$. Let $\Sigma' = \{\eta \in \Sigma : \|\eta\| < \infty\}$. Then Σ'' is dense in Σ' , $\nu(\Sigma') = 1$, and a standard adaptation of E. Andjel's construction (see [2]) implies the existence of a measure \mathbb{P}^ν on $\Omega = D([0, \infty), \Sigma')$ corresponding to a Markov process with the prescribed transition rates, having the Feller property in this topology, and thus satisfying (1).

the *global* consequences of these properties that will be central in the upcoming sections, and define $m_{\nu, \eta}$, stable configurations and stabilizing sequences.

Representation. We start with the representation for the SSM. A site $x \in \mathbb{Z}^d$ is *stable* in the configuration η if $\eta(x) = 0$ or 1 , and it is *unstable* if $\eta(x) \geq 2$. When a site is unstable it can *topple*, sending away 2 particles, each one to a neighbor of x chosen independently at random. Toppling a stable site is an *illegal* operation and it will not be allowed.

To define the random topplings, start with an independent set of *instructions* $\mathcal{I} = (\tau^{x,j} : x \in \mathbb{Z}^d, j \in \mathbb{N})$, where $\tau^{x,j} = \tau_{xy}$ with probability $p(y-x)$. \mathcal{P}^ν will denote the joint law of η and \mathcal{I} , where η has distribution ν and is independent of \mathcal{I} .

Let $h = (h(x); x \in \mathbb{Z}^d)$ count the number of topplings at each site. The toppling operation at x is defined by $\Phi_x(\eta, h) = (\tau^{x, 2h(x)+2} \cdot \tau^{x, 2h(x)+1} \cdot \eta, h + \delta_x)$ and $\Phi_x \eta$ is a short for $\Phi_x(\eta, 0)$.

The analogous representation for the ARW is described using the same notations. For this model, a site x is *stable* in the configuration η when $\eta(x) = 0$ or $\eta(x) = \rho$ and it is *unstable* when $\eta(x) \geq 1$. The instructions $(\tau^{x,j} : x \in \mathbb{Z}^d, j \in \mathbb{N})$ are independent and are equal to τ_{xy} with probability $\frac{p(y-x)}{1+\lambda}$ or $\tau_{x\rho}$ with probability $\frac{\lambda}{1+\lambda}$. The “toppling” $\Phi_x(\eta, h) = (\tau^{x, h(x)+1} \eta, h + \delta_x)$ is *legal* when x is *unstable* in η , i.e., when $\eta(x) \geq 1$.

Properties. For $\alpha = (x_1, \dots, x_k)$, we write $\Phi_\alpha = \Phi_{x_k} \Phi_{x_{k-1}} \cdots \Phi_{x_1}$ and say that Φ_α is *legal* for η if Φ_{x_l} is legal for $\Phi_{(x_{l-1}, \dots, x_1)} \eta$ for each $1 \leq l \leq k$. In this case we say that α is a *legal sequence* of topplings for η . Let $m_\alpha = (m_\alpha(x); x \in \mathbb{Z}^d)$ be given by $m_\alpha(x) = \sum_l \mathbb{1}_{x_l=x}$, the number of times the site x appears in α . We write $m_\alpha \geq m_\beta$ if $m_\alpha(x) \geq m_\beta(x) \forall x$, and $\tilde{\eta} \geq \eta$ if $\tilde{\eta}(x) \geq \eta(x) \forall x$. We also write $(\tilde{\eta}, \tilde{h}) \geq (\eta, h)$ if $\tilde{\eta} \geq \eta$ and $\tilde{h} = h$.

Let x be a site in \mathbb{Z}^d and η, η' be configurations.

Property 1. If α is a legal sequence for η , then $\Phi_\alpha \eta$ depends on α only through m_α .

Property 2. $\Phi_\alpha \eta(x)$ is non-increasing in $m_\alpha(x)$ and non-decreasing in $m_\alpha(z), z \neq x$.

Property 3. If x is unstable in η and $\eta'(x) \geq \eta(x)$, then x is unstable in η' .

Property 4. If moreover $\eta' \geq \eta$ then $\Phi_x \eta' \geq \Phi_x \eta$.

Proof. Property 1. For sandpile-like models this follows simply from the abelianess of arithmetics, but the ARW has a slight subtlety. For each site y consider the action $\eta(y) \mapsto \tau \eta(y)$ of an operator τ on $\eta(y)$. It is of the form $\tau_k : n \mapsto n + k$ for some $k \geq 0$ when $\tau = \tau_{xz}, x \neq y$, or of the form $\tau_{-1} : n \mapsto n - 1$ when $\tau = \tau_{yx}$, or $\tau_\rho : n \mapsto \rho \cdot n$ when $\tau = \tau_{y\rho}$, and the last two are legal operations only when $n \geq 1$. Now if τ_ρ is legal then $\tau_\rho \tau_k = \tau_k \tau_\rho$ are legal, idem for τ_{-1} , and moreover $\tau_k \tau_{k'} = \tau_{k'} \tau_k$. The τ_{-1} and τ_ρ do not commute, but this does not pose any problem since the order of their appearance in Φ_α is determined by the stack of instructions $(\tau^{y,j})_{j \in \mathbb{N}}$ at y .

Property 2. For each site x , toppling sites other than x cannot decrease $\eta(x)$, and toppling site x cannot increase $\eta(x)$.

Properties 3 and 4 are obvious. □

Consequences. Let V be a finite subset of \mathbb{Z}^d . A configuration η is said to be *stable* in V if all the sites $x \in V$ are stable. We say that α is *contained* in V if all its elements are in V , and we say that α *stabilizes* η in V if every $x \in V$ is stable in $\Phi_\alpha \eta$.

Lemma 1 (Least Action Principle). *If α and β are legal sequences of topplings for η such that β is contained in V and α stabilizes η in V , then $m_\beta \leq m_\alpha$.*

Remark. The sequence α need not be legal, it suffices that the topplings can be defined in some sense, and satisfy Properties 1–4. For an application that makes use of the Least Action Principle in such context, see [12].

Proof. We follow Eriksson [11]. The proof consists in observing that all topplings performed by β are necessary for stabilization. Let β be a legal sequence contained in V and $m_\alpha \not\leq m_\beta$. Write $\beta = (x_1, \dots, x_k)$ and $\beta^{(j)} = (x_1, \dots, x_j)$ for $j \leq k$. Let $\ell = \max\{j : m_{\beta^{(j)}} \leq m_\alpha\} < k$ and $y = x_{\ell+1} \in V$. Now y is unstable in $\Phi_{\beta^{(\ell)}} \eta$ since β is legal, moreover $m_{\beta^{(\ell)}} \leq m_\alpha$ and $m_{\beta^{(\ell)}}(y) = m_\alpha(y)$ by definition of ℓ . By Properties 2 and 3, y is unstable for $\Phi_\alpha \eta$ and therefore α does not stabilize η in V . \square

Lemma 2 (Abelian Property). *If α and β are both legal toppling sequences for η that are contained in V and stabilize η in V , then $m_\alpha = m_\beta$. In particular, $\Phi_\alpha \eta = \Phi_\beta \eta$.*

Proof. Apply Lemma 1 in two directions: $m_\alpha \leq m_\beta \leq m_\alpha$. \square

By Lemma 2, $m_{V,\eta} = m_\alpha$ and $\xi_{V,\eta} = \Phi_\alpha \eta$ are well defined.

Lemma 3 (Monotonicity). *If $V \subseteq V'$ and $\eta \leq \eta'$, then $m_{V,\eta} \leq m_{V',\eta'}$.*

Proof. Let α and β be legal and stabilizing sequences respectively for η' in V' and for η in V . By successively applying Properties 3 and 4, we see that β is also a legal sequence for η' in V . Thus $m_\alpha \geq m_\beta$ by Lemma 1 and the result follows. \square

By monotonicity, the limit $m_\eta = \lim_n m_{V_n,\eta}$ exists and does not depend on the particular sequence $V_n \uparrow \mathbb{Z}^d$. A configuration η is said to be *stabilizable* if $m_\eta(x) < \infty$ for every $x \in \mathbb{Z}^d$.⁷

Lemma 4. *Let ν be a translation-invariant, ergodic distribution with finite density $\nu(\eta(0))$. Then $\mathbb{P}^\nu(\text{the system locally fixates}) = \mathcal{P}^\nu(m_\eta(0) < \infty) \in \{0, 1\}$.*

Sketch of the proof. The same proof works for both the ARW and the SSM, taking $\lambda = 0$ for the latter. We present the main steps and make references to [34] when omitting details. 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 $h_\infty(x) = \lim_{t \rightarrow \infty} h_t(x)$, the limit exists as $h_t(x)$ is non-decreasing in t .

The proof is split in three parts. First, the 0-1 law for the \mathcal{P}^ν -probability that η is stabilizable at some fixed site x . Second, that \mathcal{P}^ν provides a coupling for \mathbb{P}^ν so that

$$\mathbb{P}^\nu[h_\infty(x) \geq r] = \mathcal{P}^\nu[m_\eta(x) \geq r] \quad \text{for each } r > 0. \quad (2)$$

⁷ If η is stabilizable, $m_{V_n,\eta}(x) = m_\eta(x)$ for large n , thus by Property 2 we have $\xi_{V_{n'+1}}(x) \geq \xi_{V_n}(x)$ for $n' > n$, and since $\xi_V(x) \leq 2$ for any V , the final state $\xi(x) = \lim_{n \rightarrow \infty} \xi_{V_n}(x)$ is well defined and does not depend on the particular sequence $V_n \uparrow \mathbb{Z}^d$.

Third, that blowups do not happen in finite time:

$$\mathbb{P}^{\mathcal{V}}[h_t(x) \geq r] \rightarrow 0 \quad \text{as } r \rightarrow \infty \quad \text{for each fixed } t. \quad (3)$$

Let us show that these imply the lemma. If $\mathcal{P}^{\mathcal{V}}[m_{\eta}(x) < \infty] = 1$, then it follows from (2) that $\mathbb{P}^{\mathcal{V}}[h_t(x) \text{ eventually constant}] = 1$, 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 $\mathbb{P}^{\mathcal{V}}[\eta_t(x) \text{ converges}] = 1$. Otherwise $\mathcal{P}^{\mathcal{V}}[m_{\eta}(x) = \infty] = 1$ by the 0-1 law, then (2) gives $\mathbb{P}^{\mathcal{V}}[h_t(x) \rightarrow \infty \text{ as } t \rightarrow \infty] = 1$ and by (3) we know that $(h_t(x))_{t \geq 0}$ cannot blow up in finite time, whence for each x , the value of $\eta_t(x)$ jumps for arbitrarily large times, and the system stays active.

The 0-1 law comes from the fact that $m_{\eta}(x) = \infty$ implies $m_{\eta}(x+z) = \infty$ for all z such that $p(z) > 0$, for a.e. \mathcal{S} . See [34] for the details, as well as the proof of (3).

The proof of (2) is again divided in parts. First we use $\mathcal{P}^{\mathcal{V}}$ to simultaneously produce processes $(\eta_t^M, h_t^M)_{t \geq 0}$ distributed as $\mathbb{P}_M^{\mathcal{V}}$ for all $M \in \mathbb{N}$, starting from $\eta^M(z) = \eta(z) \mathbb{1}_{\|z\| \leq M}$ at $t = 0$. It follows from (1) that

$$\mathcal{P}^{\mathcal{V}}(h_t^M(x) \geq r) = \mathbb{P}_M^{\mathcal{V}}(h_t(x) \geq r) \xrightarrow{M \rightarrow \infty} \mathbb{P}^{\mathcal{V}}(h_t(x) \geq r) \xrightarrow{t \rightarrow \infty} \mathbb{P}^{\mathcal{V}}(h_{\infty}(x) \geq r). \quad (4)$$

We check that

$$\mathcal{P}^{\mathcal{V}}(h_t^M(x) \geq r) \xrightarrow{t \rightarrow \infty} \mathcal{P}^{\mathcal{V}}(m_{\eta^M}(x) \geq r) \xrightarrow{M \rightarrow \infty} \mathcal{P}^{\mathcal{V}}(m_{\eta}(x) \geq r) \quad (5)$$

and finally show that the limits in (4) and (5) commute.

The coupling $\{(\eta_t^M, h_t^M)_{t \geq 0}\}_{M \in \mathbb{N}}$ satisfying (5) is described in details in [34]. We show that the limits in (4) and (5) commute via monotonicity. This follows from the coupling as well, for which $h_t^M(x)$ is a.s. non-decreasing in M and t . \square

4 Strategy for stabilization

In this section we give an overview of the general strategy in the proof of fixation.

The goal is to construct an algorithm that, given any pair η, \mathcal{S} , tries to stabilize all the particles initially present in η following the jump instructions in \mathcal{S} , but with the aid of some extra topplings. The algorithm must be such that $m_{\eta}(0) = 0$ whenever it is successful. It may well fail, but all we need is *a strategy that succeeds with positive probability*, which in turn implies a.s. fixation by the 0-1 law (see Lemma 4).

The algorithm consists of applying a *settling procedure* to each particle. This procedure *explores* \mathcal{S} until it identifies a suitable *trap* for the particle. The exploration follows the path that the particle would perform if we always toppled the site it occupies, and stops when the trap has been chosen. In the absence of a suitable trap, we declare the algorithm to have failed. The trap always lies on the exploration path, however not necessarily on its tip because we need to explore further away before taking the decision. Once the trap has been chosen, the particle is moved along the exploration path until it reaches the trap, where it is settled.

The major issue is of course the spread of activity. Thus the first requirement for the settling procedure is *not to disturb particles that have already been settled*, meaning that the exploration does not examine those sites.

We need to conciliate this requirement with the goal that the algorithm have positive probability of success. This imposes that the settling procedures not only succeed with a high probability, but these probabilities should actually converge to 1. We hence consider a procedure that tries to find the traps close together as much as possible, in order to *leave more space for the maneuver of subsequent steps*.

Moreover, certain control is needed on the joint distribution of the outcome of different explorations. Some of the explored instructions are actually not going to be used by the corresponding particle by the time it settles at the trap, leaving some *corrupted sites* that may interfere with the next steps. In our proof we go for *independence*, which means that *corrupted sites left by previous steps must be avoided*. Hence the settling procedure should try not only to find its trap close to those already found, but also to *keep the corrupted region as compact as possible*.

We seek a procedure that sets the trap close to the positions where the previous traps have been placed, therefore implying typically long excursions away from the particle starting position. But *we are bound to follow the instructions in \mathcal{I}* , and cannot prevent the particles from stopping before.

To circumvent this restriction, we *enforce activation and push the particles further*, in a way that \tilde{m} , the total number of topplings after enforcement, provides an *upper bound* for the true m_η . For the ARW this is done by ignoring some instructions of the form τ_{xp} , and we show in the next session that it increases m_η . The SSM case is considered in the sequel.

Let us describe the stabilizing strategy used in this paper. Label the initial positions of the particles on \mathbb{Z}_+ by $0 \leq x_1 \leq x_2 \leq x_3 \leq \dots$. Take $a_0 = 0$ and suppose the first $k - 1$ traps have been successfully set up at positions $0 < a_1 < a_2 < \dots < a_{k-2} < a_{k-1} < x_{k-1}$. The set $\llbracket 0, a_{k-1} \rrbracket$ contains all traps and corrupted sites found so far, where $\llbracket z, w \rrbracket$ denotes $[z, w] \cap \mathbb{Z}$.

The settling procedure starts with an *exploration*. Starting at x_k , examine and follow the instructions in \mathcal{I} one by one (whenever a sleep instruction is found, the next instruction at the same site is to be examined). Follow this exploration until reaching a_{k-1} .

Next we *set up the trap*. We consider first the ARW and postpone the SSM subtleties. During the k -th exploration, we are sure to visit every $x \in B_k = \llbracket a_{k-1} + 1, x_k - 1 \rrbracket$. Moreover, the last instruction explored at each $x \in B_k$ is a jump to the left, see Figures 1 and 2. For some $x \in B_k$, the second last instruction may be a sleep instruction. Let a_k be the leftmost such site. If there is none, we declare the procedure *unsuccessful* and stop.

At this point the instruction manipulation enters the game: we suppress all the sleep instructions explored meanwhile, except for the last sleep instruction at a_k , which becomes the trap.

Since the trap is a sleep instruction found immediately before the last instruction, which is a jump to the left, we are sure that the exploration path has not gone to the

right of a_k after exploring the trap, so *all the corrupted sites will be in* $\llbracket a_{k-1} + 1, a_k \rrbracket$, see Figure 2. We repeat this procedure indefinitely, unless it fails at some step.

We finally show that this strategy is *successful with positive probability*. Each site $x \in B_1$ has a sleep instruction just before its last jump independently and with probability $\bar{\mu} = \frac{\lambda}{1+\lambda}$. Thus, $a_1 - a_0$ is distributed as a geometric random variable of mean $\bar{\mu}^{-1}$, truncated at $x_1 - a_0$. Since no corrupted sites were left outside $\llbracket a_0 + 1, a_1 \rrbracket$, the interdistance $a_2 - a_1$ is independent of a_1 and has the same distribution. By the law of large numbers $a_n \sim n/\bar{\mu}$. On the other hand, $x_n \sim n/\mu$. Therefore, if $\mu < \bar{\mu}$ there is positive probability that $a_k < x_k$ for all k , which implies success.

In the SSM case, an isolated particle stands still and the site where it is located is stable. A natural way to enforce activity is to allow a single particle to move. To achieve this, we regard the toppling operation as having two distinct steps: the particles are sent to neighboring sites, but one at a time, leading to what we call *half-topplings*.

However, this has to be consistent with the specification of the model. More precisely, to determine whether a half-toppling is legal we need to know whether or not the site has been half-toppled. If $\eta(x) \geq 2$ the half-toppling is always legal. It is also legal if $\eta(x) = 1$, but only when x has been half-toppled. These rules yield the same m_η , and moreover, violating the latter condition, i.e., half-toppling stable particles, results in an upper bound. See Section 6 for the details of how this tool is used in the proof of fixation.

5 Phase transition for activated random walks

In this section we prove Theorem 2. The core of the proof of fixation is an algorithm that, given \mathcal{S} and η , selects a subset of the sleep instructions in \mathcal{S} , providing a new $\tilde{\mathcal{S}}$ under which $\tilde{m}_\eta(0) = 0$. This algorithm works with positive probability, and the theorem follows from Lemma 4 and Lemma 5 below.

Besides the τ_{xy} and $\tau_{x\rho}$ considered in Section 3, consider in addition the neutral instruction ι , given by $\iota\eta = \eta$. Replacing a sleep instruction by a neutral instruction may be seen as enforcing the active state on a passive particle, or replacing the transition $A \rightarrow S$ by $A \rightarrow S \rightarrow A$. Given two sets of instructions $\mathcal{S} = (\tau^{x,j})_{x,j}$ and $\tilde{\mathcal{S}} = (\tilde{\tau}^{x,j})_{x,j}$, we write $\tilde{\mathcal{S}} \geq \mathcal{S}$ if for every $x \in \mathbb{Z}^d$ and $j \in \mathbb{N}$, either $\tau^{x,j} = \tilde{\tau}^{x,j}$, or $\tau^{x,j} = \tau_{x\rho}$ and $\tilde{\tau}^{x,j} = \iota$. As we deal with more than one set of instructions, we need to enlarge the notations of Section 3 and write $m_{V,\eta;\mathcal{S}}$, $\Phi_{\alpha;\mathcal{S}}$, etc., to specify which set of instructions \mathcal{S} is being considered.

Lemma 5 (Monotonicity with enforced activation). *Let $\tilde{\mathcal{S}} \geq \mathcal{S}$ be sets of instructions, η be a initial state, and $V \subseteq \mathbb{Z}^d$ be a fixed finite set. Then $m_{V,\eta;\mathcal{S}} \leq m_{V,\eta;\tilde{\mathcal{S}}}$. In particular, $m_{\eta;\mathcal{S}} \leq m_{\eta;\tilde{\mathcal{S}}}$.*

Proof. Let $\tilde{\mathcal{S}} \geq \mathcal{S}$. By Lemma 1, it suffices to show that if $(\tilde{\eta}, \tilde{h}) \geq (\eta, h)$, and α is a legal sequence for (η, h) under \mathcal{S} , then α is also legal for $(\tilde{\eta}, \tilde{h})$ under $\tilde{\mathcal{S}}$. We show the stronger fact that $\Phi_{\alpha;\tilde{\mathcal{S}}}(\tilde{\eta}, \tilde{h}) \geq \Phi_{\alpha;\mathcal{S}}(\eta, h)$. The proof is by induction on the number of steps $k = \sum_z m_\alpha(z)$.

For $k = 0$ the claim is trivially true. For $k \geq 2$ we write $\alpha = \alpha' \alpha''$ and apply the induction hypothesis for both α' and α'' . It remains to consider $k = 1$. Write $\alpha = (y)$. Since $\tilde{\eta} \geq \eta$, y is also unstable in $\tilde{\eta}$. Write $\eta' = \tau^{y, h(y)+1} \eta$ and $\tilde{\eta}' = \tilde{\tau}^{y, h(y)+1} \tilde{\eta}$. Now either $\tau^{y, h(y)+1} = \tilde{\tau}^{y, h(y)+1}$, and thus $\tilde{\eta}' = \tau^{y, h(y)+1} \tilde{\eta} \geq \tau^{y, h(y)+1} \eta = \eta'$, or $\tau^{y, h(y)+1} = \tau_{x\rho}$ and $\tilde{\tau}^{y, h(y)+1} = \iota$, and thus $\eta' = \tau^{y, h(y)+1} \eta = \tau_{y\rho} \eta \leq \eta \leq \tilde{\eta} = \iota \tilde{\eta} = \tilde{\eta}'$. In any case, $\eta' \leq \tilde{\eta}'$ and $\Phi_{y, \mathcal{J}}(\tilde{\eta}, \tilde{h}) = (\tilde{\eta}', h + \delta_y) \geq (\eta', h + \delta_y) = \Phi_{y, \mathcal{J}}(\eta, h)$. \square

Proof of Theorem 2. Let \mathcal{P}^μ denote \mathcal{P}^V for V the product probability distribution having Poisson(μ) as marginals. By Lemmas 3 and 4 it suffices to show that $\mathcal{P}^\mu(m_\eta(0) = 0) > 0$ when $\mu < \frac{\lambda}{1+\lambda}$ and $\mathcal{P}^\mu(m_\eta(0) = \infty) > 0$ when $\mu > 1$.

Let $\mu > 1$. For some $\delta > 0$, the \mathcal{P}^μ -probability that there are at least μM particles in $V = [-M, 0]$ is at least 2δ , independently of M . After stabilizing η in V , at least $(\mu - 1)M$ particles will have to exit V , and because of the topological constraints of the one-dimensional lattice, $m_{V, \eta}(-M) \geq (\mu - 1)M/2$ or $m_{V, \eta}(0) \geq (\mu - 1)M/2$ must happen. Thus, at least one of these two events has probability bigger than δ . The latter case gives directly $\mathcal{P}^\mu(m_\eta(0) \geq (\mu - 1)M/2) > \delta$, and in the former case we can consider $V = [0, M]$, which gives the same inequality by translation invariance. Letting $M \rightarrow \infty$ we get $\mathcal{P}^\mu(m_\eta(0) = \infty) \geq \delta > 0$.

In the remainder of this section we prove that η is stabilizable when $\mu < \frac{\lambda}{1+\lambda}$. Given \mathcal{J} and η , the goal is to construct a stabilizing strategy in an algorithmic way, finding $\tilde{\mathcal{J}} \geq \mathcal{J}$ such that $m_{\eta; \tilde{\mathcal{J}}}(0) = 0$. We consider first the recurrent case $p(-1) = p(+1) = \frac{1}{2}$.

Start placing an imaginary *barrier* at site $a_0 = 0$ and pick the first particle to the right of this barrier. *Explore* the instructions in \mathcal{J} one by one, following a path that starts at the position x_1 of this first particle. This means following the path that would have been performed by this particle if we were to topple the site that contains it over and over. Follow this exploration until it reaches the origin. When sleep instructions are found, this explorer looks for the next instruction at the same site, until it finds a jump instruction, as in Figure 1.

We then *set up the trap* for the first particle. For each site $y \in \llbracket a_0 + 1, x_1 - 1 \rrbracket$, the last instruction examined by this explorer at y must be a jump to the left, $\tau^{y, j} = \tau_{y, z}$, where $z = y - 1$ and $j = j_1(y)$ is the total number of instructions examined at site y . It may be the case that before this jump to the left, the previous instruction $\tau^{y, j-1}$ found at y is $\tau_{y\rho}$. Let $a_1 = \min \{y \in \llbracket a_0 + 1, x_1 - 1 \rrbracket : \tau^{y, j_1(y)-1} = \tau_{y\rho}\}$ be the leftmost such site. This is the site where the *trap is set up*. Place another imaginary *barrier* at a_1 , as in Figure 1. If no such site is found, we declare the procedure to have failed and we stop it. Otherwise, we declare the first step to be *successful*.

We now *explore the path* for the *second particle* and *set up the next trap*. Start another explorer at x_2 , the position of the second particle to the right of the origin, and follow this exploration until hitting the barrier a_1 . This explorer will examine the instructions of \mathcal{J} in the order as they appear, but skipping those that have been used in the previous step, as in Figure 2. As before, for each $y \in \llbracket a_1 + 1, x_2 - 1 \rrbracket$, the last instruction at y must be a jump to the left, $\tau^{y, j} = \tau_{y, z}$, where $z = y - 1$ and $j = j_2(y)$

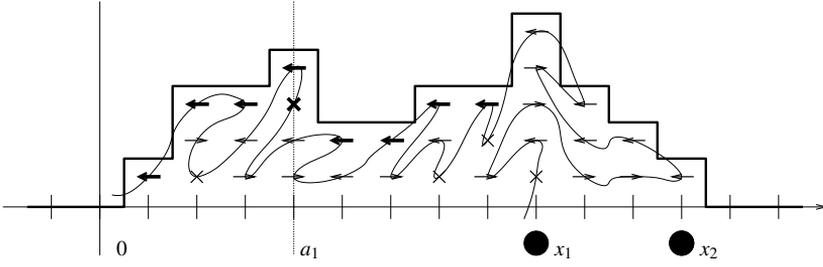


Fig. 1 First exploration path for the ARW. It starts at position x_1 of the first particle and stops when it reaches the origin. The horizontal axis represents the lattice, and above each site x there is a sequence of instructions $(\tau^{x,j})_j$. The bold arrows indicate the last jump found at each site $x \in \llbracket 1, x_1 - 1 \rrbracket$, 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 .

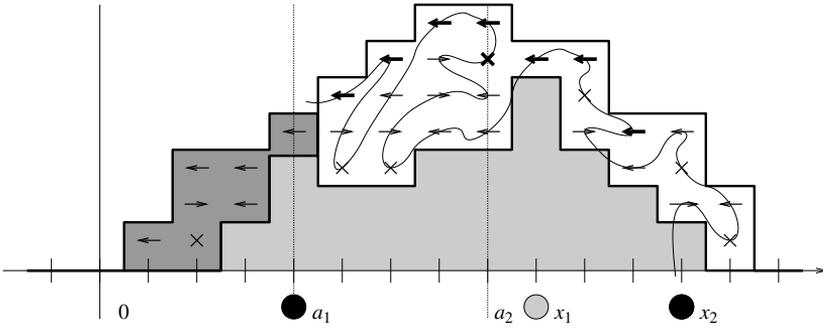


Fig. 2 Second exploration path for the ARW. It starts at position x_2 of the second particle and stops when it reaches a_1 . The regions in gray indicate the instructions already examined by the first explorer. The dark gray contains instructions examined but not used, whose locations determine the set of corrupted sites.

is the total number of instructions examined at site y by both explorers. The *trap* is set up at $a_2 = \min \{y \in \llbracket a_1 + 1, x_2 - 1 \rrbracket : \tau^{y, j_2^{(y)-1}} = \tau_{y\rho}\}$, the leftmost site where the instruction examined just before the last jump to the left is a sleep instruction, as in Figure 2. Again we place a new imaginary *barrier* at a_2 . If no such site is found, we declare the procedure to have failed and we stop it. Otherwise, we declare the second step to be *successful*.

This procedure can be carried on indefinitely, as long as *all the steps are successful*. We then perform a similar construction for the negative half line, finding the site x_{-1} with the first particle to the left of the origin, setting up a trap at $a_{-1} \in \llbracket x_{-1} + 1, -1 \rrbracket$, then at $a_{-2} \in \llbracket x_{-2} + 1, a_{-1} - 1 \rrbracket$, and so on.

Let us show that *success implies stability*. Suppose all the steps in the algorithm are successful. Replace each sleep instruction in \mathcal{I} by a neutral instruction, *except for those providing the traps*. Namely, take

$$\tilde{\tau}^{y,j} = \begin{cases} \mathbf{1}, & \tau^{y,j} = \tau_{y\rho}, (y, j) \neq (a_n, j_n(a_n)) \forall n \in \mathbb{Z}, \\ \tau^{y,j}, & \text{otherwise.} \end{cases}$$

Fix $n \in \mathbb{N}$. We need to show that, following the instructions of $\tilde{\mathcal{J}}$, η is stabilized in $V_n = [x_{-n}, x_n]$ with finitely many topplings, and moreover $m_{\eta; \tilde{\mathcal{J}}}(0) = 0$.

We first stabilize the particle that starts at x_1 . To that end, we successively topple the sites found by the first explorer. Since all the sleep instructions have been replaced, the trajectory of the particle started at x_1 will coincide with the path of the corresponding explorer, until it finds the trap $\tau^{a_1, j_1(a_1)-1} = \tau_{a_1 \rho}$. At this moment the particle will become passive, and the site a_1 will be stable.

Notice that, after the last visit to a_1 , the explorer does not go further to the right, so when settling the first particle we use all the instructions examined so far, except some lying in $\llbracket a_0 + 1, a_1 \rrbracket$. Therefore, the same procedure can be applied to the second particle, as it will find the same instructions that determined the second exploration path.

Notice also that the first particle does not visit 0, and the second particle neither visits 0 nor a_1 , thus it is settled without activating the first particle. Following the same procedure, the k -th particle is settled at a_k , without ever visiting $\{0, a_1, a_2, \dots, a_{k-1}\}$, 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 can be stabilized in V_n with finitely many topplings, not necessarily in V_n , and never toppling the origin. By Lemma 1, $m_{V_n, \eta; \tilde{\mathcal{J}}}(0) = 0$. On the other hand, it follows from Lemma 5 that $0 \leq m_{V_n, \eta; \mathcal{J}}(0) \leq m_{V_n, \eta; \tilde{\mathcal{J}}}(0)$, whence $m_{V_n, \eta; \mathcal{J}}(0) = 0$. Since it holds for all $n \in \mathbb{N}$ and $V_n \uparrow \mathbb{Z}$ as $n \rightarrow \infty$, this gives $m_{\eta; \mathcal{J}}(0) = 0$.

We conclude with the proof that there is *positive probability of success* for the algorithm described above. Namely, we show that the set of (η, \mathcal{J}) for which the above construction is successful has positive \mathcal{P}^μ -probability whenever $\mu < \frac{\lambda}{1+\lambda}$.

We claim that the position a_1 of the first barrier has a geometric distribution with parameter $\frac{\lambda}{1+\lambda}$. To be more accurate, the claim is that the probability space can be enlarged so that we can define a random variable Y_1 , independent of η , satisfying $\mathcal{P}^\mu(Y_1 > k) = \left(\frac{1}{1+\lambda}\right)^k$, with the property that the first step of the construction is successful if and only if $Y_1 < x_1$, in which case the position a_1 of the first barrier is given by $a_1 = Y_1$.

Let us prove the above claim. For each site $y \in \mathbb{Z}$, the instruction $\tau^{y, j}$ can be a sleep instruction $\tau_{y\rho}$ with probability $\frac{\lambda}{1+\lambda}$ or a jump instruction of the form τ_{yz} with probability $\frac{1}{1+\lambda}$, independently for each $j \in \mathbb{N}$. Conditioning on $J^y = \{j^1 < j^2 < \dots\}$, the subset of \mathbb{N} for which $\tau^{y, j}$ is a jump instruction, $\tau^{y, j^k} = \tau_{yz}$ with probability $p(z-y)$, independently for each k . Now the trajectory performed by the first explorer depends only on the jump instructions, and, conditioned on the trajectory, we have $\tau^{y, j_1(y)-1} = \tau_{y\rho}$ with probability $\frac{\lambda}{1+\lambda}$ and independently for each $y \in \llbracket a_0 + 1, x_1 - 1 \rrbracket$. If it fails for all such y , which happens with probability $\left(\frac{1}{1+\lambda}\right)^{x_1-1}$, sample Y_1 as $\mathcal{P}^\mu(Y_1 > k | Y_1 > x_1 - 1) = \left(\frac{1}{1+\lambda}\right)^{k-x_1+1}$. Otherwise take Y_1 as the smallest such y . This implies that Y_1 has the prescribed distribution, proving the claim.

Since each explorer skips the instructions already examined at previous steps, the sequence $(a_1, a_2 - a_1, a_3 - a_2, \dots)$ is i.i.d. More precisely, by the same argument as above, there is a sequence of i.i.d. variables Y_1, Y_2, Y_3, \dots with the property that the n -

th step is successful if and only if the previous steps are successful and $a_{k-1} + Y_k < x_k$, in which case $a_k = a_{k-1} + Y_k$.

Now the expectation of Y_1 is $\frac{\lambda+1}{\lambda}$. By the law of large numbers, the putative position $X_k = Y_1 + \dots + Y_k$ of the k -th barrier grows like $\frac{\lambda+1}{\lambda}k$. On the other hand, by the law of large numbers, the initial position x_k of the k -th particle to the right of the origin grows like k/μ . Therefore, as $\mu < \frac{\lambda}{1+\lambda}$, $X_k < x_k$ for all k with positive probability.

When we apply the same construction switching left and right, we define $X_k = Y_k + \dots + Y_{-1}$ which satisfy $x_k < X_k$ for all $k < 0$ with the same probability and independently of the first part. Therefore the construction is successful with positive probability and the proof of the recurrent case is finished.

We now consider the transient case $p(+1) = 1 - p(-1) > \frac{1}{2}$. If the n -th explorer never hits a_{n-1} , it is a.s. the case that it examines only finitely many instructions at each site. In this case we take $a_n = a_{n-1}$, no sleep instruction is needed for the n -th particle as its trap will be at ∞ , and it is still possible to define the subsequent explorers. When stabilizing the configuration η in $V_n = \llbracket x_{-n}, x_n \rrbracket$, we let the particles follow the path of the corresponding explorer as before. The total number of topplings may be infinite (and in the limit the transient particles disappear), but each site is toppled finitely many times and, as in the symmetric case, the origin never topples. \square

6 Phase transition for the stochastic sandpile model

In this section we prove Theorem 1. The proof of fixation consists in an algorithm that, given \mathcal{I} and η , stabilizes η with a sequence of topplings, *some of which are not legal*. More precisely, we introduce the half-toppling operation, which moves the particles one by one. The configuration is then stabilized with the aid of some illegal half-topplings, and without toppling the origin. This algorithm works with positive probability, and the theorem follows from Lemma 4 and Lemma 6 below.

A *half-toppling* consists of sending only one particle to a neighboring site chosen at random (i.e., following the appropriate instruction) and is possible when $\eta(x) > 0$. We say that site x is *stable* if $\eta(x) = 0$ and *unstable* if $\eta(x) \geq 2$. When $\eta(x) = 1$, the site x can be stable or unstable, depending on whether it has been half-toppled an even or an odd number of times. A half-toppling at x is considered *legal* when x is unstable, and is considered *semi-legal* when $\eta(x) \geq 1$ regardless of the site being stable or not. We denote a half-toppling at x by ϕ_x , that is, $\phi_x(\eta, h) = (\tau^{x, 2h(x)+1} \eta, h + \frac{1}{2} \delta_x)$ and $\phi_x \eta = \phi_x(\eta, 0)$. Notice that $\Phi_x = \phi_x \circ \phi_x$.

For a sequence $\beta = (y_1, \dots, y_n)$, define $\tilde{m}_\beta(y) = \sum_j \mathbb{1}_{y_j=y}$ as the number of times that the site y appears in β . Properties 1–4 also hold with half-topplings. For $\alpha = (x_1, \dots, x_k)$, denoting $\alpha^2 = (x_1, x_1, x_2, x_2, \dots, x_k, x_k)$, we have that $\tilde{m}_{\alpha^2} = 2m_\alpha$ and, given a configuration η , the sequence of half-topplings ϕ_{α^2} is legal if and only if the sequence of topplings Φ_α is itself legal, and in this case $\phi_{\alpha^2} \eta = \Phi_\alpha \eta$.

Lemma 6 (Least Action Principle, with half-topplings). *Let β be a sequence of half-topplings in the volume V that is legal for the configuration η . Let α be a sequence of half-topplings in \mathbb{Z}^d that is semi-legal for η and stabilizes η in V . Then $\tilde{m}_\beta \leq \tilde{m}_\alpha$.*

Proof. The proof is the same as for Lemma 1. \square

Corollary 7 (Abelian Property, with half-topplings). *If α is a sequence of half-topplings in V that is legal for η and stabilizes η in V , then $\tilde{m}_\alpha = 2m_{V,\eta}$.*

Proof. Take a sequence of topplings β in V that is legal for η and stabilizes η in V . Then β^2 is a legal sequence of half-topplings in V that is legal for η and stabilizes η in V , and $\tilde{m}_{\beta^2} = 2m_\beta$. Now $m_\beta = m_{V,\eta}$ by Lemma 2 and $\tilde{m}_\alpha \leq \tilde{m}_{\beta^2} \leq \tilde{m}_\alpha$ by Lemma 6. \square

Proof of Theorem 1. The outline of the proof is the following. First we show that there is no fixation for $\mu > 1$. Next we describe how to stabilize the particles initially present in \mathbb{Z}_+ , and show that this scheme has positive probability of success if μ is small enough. We then argue by symmetry that an analogous procedure works for the particles starting in \mathbb{Z}_- . The first and last parts are identical to those presented in Section 5, and we restrict this section to the specificity of the settling procedure.

Given \mathcal{S} and η , we will construct a stabilizing strategy in an algorithmic way, giving $(\tilde{m}_{V_n}(x) : x \in \mathbb{Z})_{n=1,2,\dots}$, where \tilde{m}_{V_n} counts the number of semi-legal half-topplings (not necessarily legal nor contained in V_n) that stabilize η in V_n and satisfy $\tilde{m}_{V_n}(0) = 0$, for some sequence $V_n \uparrow \mathbb{Z}_+$.

Start placing an imaginary *barrier* at site $a_0 = 0$ and pick the first particle to the right of this barrier. *Explore* the instructions in \mathcal{S} one by one, starting at the position x_1 of this first particle and following its path that would have been performed by this particle if we were to half-topple the site that contains it over and over until it reached a_0 , as in Figure 3.

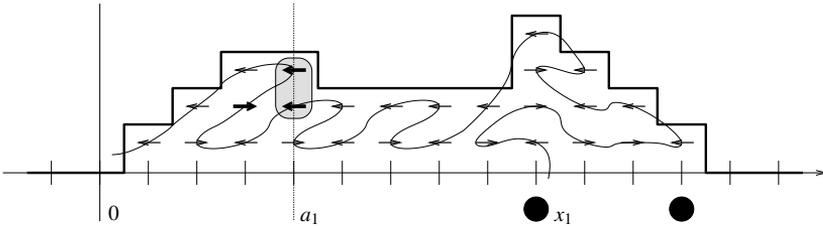


Fig. 3 First exploration for the SSM. It starts at position x_1 of the first particle and stops when it reaches the origin. The horizontal axis represents the lattice, and above each site x there is a sequence of instructions $(\tau^{x,j})_j$. The bold arrows represent the last jump to the right of this explorer and the two last jumps (necessarily to the left) on its neighboring site a_1 .

We then *set up the trap* for the first particle. Fix the site $a_1 \in \llbracket a_0 + 1, x_1 \rrbracket$ to the right of where the explorer does its *last right jump*. The *trap* will be set at a_1 , as we explain later on. Place a new barrier at this site, as in Figure 3. If no such site is found, we declare the procedure to have failed and stop it. Otherwise, we declare the first step to be *successful*.

Each step of the explorer corresponds either to a semi-legal half-toppling of a stable site, or a legal half-toppling of an unstable site. If a site is visited more than

once, one of the last two moves necessarily corresponds to the former case. Now notice that a_1 is visited at least twice by the respective explorer (which implies that at least one of the last two visits corresponds to an illegal toppling) and moreover, no site to the right of a_1 is visited after the two last jumps.

We now *explore the path* for the *second particle* and *set up the next trap*. Start another exploration at x_2 , the position of the second particle to the right of the origin, and follow this exploration until it hits the barrier a_1 . This explorer will examine the instructions of \mathcal{I} in the order as they appear, but skipping those that have been used in the previous step, as in Figure 4. As before, fix the site $a_2 \in \llbracket a_1 + 1, x_2 \rrbracket$ to the

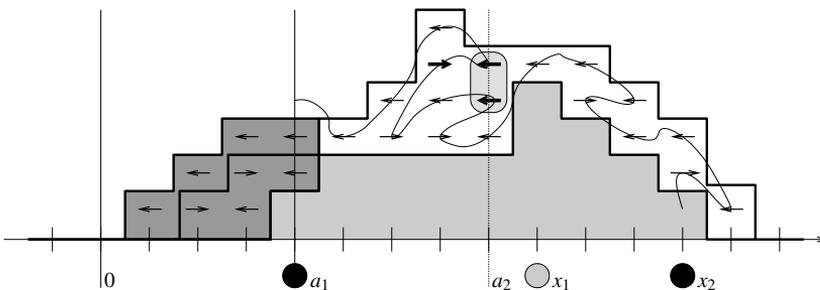


Fig. 4 Second exploration for the SSM. It starts at position x_2 of the second particle and stops when it reaches a_1 . The regions in gray indicate the instructions already examined by the first explorer. The dark gray contains instructions examined but not used, whose locations determine the set of corrupted sites. A subset of the dark gray instructions may be used, depending on how the first particle is settled at a_1 .

right of where this explorer does its last right jump. The *trap* will be set at a_2 , and we place a new barrier at this site, as in Figure 4. If no such site is found, we declare the procedure to have failed and stop it. Otherwise, we declare the second step to be *successful*. This procedure can be carried on indefinitely, as long as *all the steps are successful*.

We now show that *success implies stability*. Suppose all the steps in the algorithm are successful. Let us show how to stabilize η in $\llbracket 0, x_n \rrbracket$. We first stabilize the particle that starts at x_1 . To that end, we successively half-topple the sites found by the first explorer, except for the second last half-toppling at a_1 . At this point, if $\tilde{m}(a_1) \in 2\mathbb{N}_0$, the site a_1 is stable and we *abstain from further illegal half-topplings*, leaving the particle at a_1 . Otherwise, if $\tilde{m}(a_1) \in 2\mathbb{N}_0 + 1$, the site is unstable, in which case we keep the half-topplings until the last half-toppling at a_1 , which must be stable at this time, again leaving the particle at a_1 .

Notice that, after each of the two last visits to a_1 , the explorer does not go further to the right. Thus, when settling the first particle we use all the instructions examined by the first explorer, except some lying in $\llbracket a_0 + 1, a_1 \rrbracket$. Therefore, this procedure can be applied to the second particle as it will find the same instructions that determined the exploration path.

Notice also that the first particle does not visit 0, and the second particle neither visits 0 nor a_1 , so it is settled without causing further instabilities. Following the same

procedure, the k -th particle is settled at a_k , without ever visiting $\{0, a_1, a_2, \dots, a_{k-1}\}$, for all $k = 1, \dots, n$.

To conclude the proof of the theorem, we need to show that the construction is *successful with positive probability* when $\mu < \frac{1}{4}$.

We claim that the position a_1 of the first barrier is distributed as the time of the first jump to the left of a discrete-time simple symmetric random walk $(S_n)_{n=0,1,2,\dots}$ started at $S_0 = 0$ and conditioned to stay positive forever after. To be more accurate, the claim is that the probability space can be enlarged so that we can define a random variable Y_1 , independent of η , distributed as $\inf\{n : S_{n+1} = S_n - 1\}$, with the property that the first step of the construction is successful if and only if $Y_1 < x_1$, in which case the position a_1 of the first barrier is given by $a_1 = Y_1$.

Let us prove the above claim. A visual proof is contained in Figure 3. Let $(\tilde{S}_n)_{n=0,1,\dots,k}$ denote the path of the first explorer. If $(S_n)_{n=0,1,\dots,k}$ denotes the reversed path $S_n = \tilde{S}_{k-n}$, then S_n satisfies $S_0 = 0$, $S_k = x_1$ and $S_n > 0$ for all $n > 0$. Notice that each $(S_n)_n$ with the above properties is possible, and it happens with probability proportional to 2^{-k} , that is, $2^{-k}/Z$, where Z is the sum over all k' of $2^{-k'}$ times the number of paths $(S_n)_{n=0,1,\dots,k'}$ with the above properties. But this is the probability distribution of the path of a simple symmetric random walk $(S_n)_{n=0,1,2,\dots}$ conditioned to stay positive, up to the last visit of x_1 . In the latter each such path is also possible, and they also happen with probability proportional to 2^{-k} . Now taking $Y_1 = \inf\{n : S_{n+1} = S_n - 1\} = \tilde{S}_{n'}$ with $n' = 1 + \sup\{n : 0 \leq n < k, \tilde{S}_{n+1} = \tilde{S}_n + 1\}$, it follows that $a_1 = Y_1$ whenever $\{n : S_{n+1} = S_n - 1\} \cap \llbracket 1, x_1 - 1 \rrbracket \neq \emptyset$. The claim is proved.

As in the previous section, there is an i.i.d. sequence (Y_1, Y_2, Y_3, \dots) with the property that the n -th step is successful if and only if the previous steps are successful and $a_{k-1} + Y_k \leq x_k$, in which case $a_k = a_{k-1} + Y_k$.

Now the expectation of Y_1 can be computed explicitly, and it is equal to 4. By the law of large numbers, the putative position $X_k = Y_1 + \dots + Y_k$ of the k -th barrier grows like $4k$. On the other hand, by the law of large numbers, the initial position x_k of the k -th particle to the right of the origin grows like k/μ , and since $\mu < \frac{1}{4}$, with positive probability $X_k < x_k$ for all k .

Therefore, the construction is successful with positive probability and the proof is finished. \square

7 Concluding remarks

In our stabilizing strategy, the region to be avoided is the set of sites where particles have been settled, together with the corrupted sites, or rather the convex hull of this set. In one dimension, the volume of this region is proportional to the number of such particles. This does not seem to have a direct analogue in higher dimensions: straightforward application of a similar settling procedure would result in a Diffusion-Limited Aggregation type of growth. Thus, even though our approach seems promising for other settings, the proof of local fixation is so far restricted to one dimension.

Yet, the construction presented in Section 3, as well as Lemmas 5, 6, and Corollary 7, hold in any dimension. In particular, μ_c is always well defined. For the ARW, μ_c is non-decreasing in λ by Lemma 5.

The main results have been stated and proved under simple conditions in order to keep the presentation as transparent as possible. Below we discuss some straightforward *Generalizations*.

For the ARW, the proof that $\mu_c > 0$ when the jumps are to nearest neighbors can be adapted to bounded jumps without any difficulty. In this case we cannot get the sharp estimate $\mu_c \geq \frac{\lambda}{1+\lambda}$, though we conjecture that it should still hold. For the SSM with asymmetric jumps, the same proof still gives $\mu_c > 0$, but the lower bound for μ_c degenerates as $q \rightarrow 0$, where $q := p(+1) = 1 - p(-1)$.

The Poissonian distribution ν of the initial amount of particles at a given site plays no special role in the proof of local fixation. One can replace ν by any translation-invariant ergodic distribution with finite first moment μ and the proofs presented in the previous sections still give fixation as long as $\mu < \frac{\lambda}{1+\lambda}$ for the ARW or $\mu < \frac{1}{4}$ for the SSM. For the existence of a $\mu_c \in [0, \infty]$ separating the absorbing and active phases we have used the stochastic ordering of the family $\{\nu_\mu\}_{0 \leq \mu < \infty}$, parametrized by the density $\mu = \nu(\eta(0))$.

The proof that $\mu_c \leq 1$ in one dimension generalizes with slight modifications to any finite-range random walk and any ergodic ν . Indeed, the proof shows that the system stays active for any ergodic ν with density $\mu > 1$, or even $\mu = 1$, as long as the fluctuations diverge, that is, $\limsup_N \sum_{x=0}^N |\eta(x)| - N = \infty$. A proof that $\mu_c \leq 1$ in dimensions $d \geq 2$ was recently given in [36], using the framework described here and in [33], and a more abstract proof was later given in [1].

We finish with some *Open Problems*, see also [10].

It is remarkable that a proof of $\mu_c > 0$ is lacking for any $d \geq 2$. Whereas the trick of letting the particles evolve until hitting barriers and then setting traps worked well for $d = 1$, for higher dimensions more involved steps will be necessary.

A proof that $\mu_c < 1$ for both the SSM and ARW remains as an open problem in any dimension. For the ARW, it should hold for all λ , and moreover $\mu_c \rightarrow 0$ as $\lambda \rightarrow 0$. Yet, even a proof that $\mu_c < 1$ for *some* $\lambda > 0$ is missing. For the SSM, mean-field theory and extensive simulations confirm the existence of sustained activity for densities greater than some critical value μ_c , with $\mu_c < 1$ [8]. Simulations show that $\mu_c \approx 0.9489$ [9].

The only true parameter in the ARW model should be the density μ , in the sense that there is a value μ_c (depending only on the dimension, halting rate, and jump probabilities) such that, for any ergodic initial distribution ν , the system locally fixates if $\nu(\eta(0)) < \mu_c$ and stays active if $\nu(\eta(0)) > \mu_c$.

Besides universality of μ_c with respect to the distribution, a very interesting open problem, but also very hard to approach, is whether there is local fixation at $\mu = \mu_c$. We believe that at the critical density $\mu = \mu_c$ the two models stay active (at least when the initial condition is i.i.d. with non-degenerate marginals), in marked contrast with several lattice models that exhibit phase transition, such as percolation or the Ising model, for which there is no percolation at criticality. This has been proved

only for the very particular case of totally asymmetric, nearest-neighbor walks on the one-dimensional lattice [21].

Acknowledgments

We thank R. Dickman, who introduced to us these models, spent many hours of fruitful discussions and provided valuable numerical data. We thank C. Hoffman and A. Teixeira for deep and inspiring discussions, and D. Valesin for careful reading and valuable comments. We also thank the anonymous referees whose comments considerably improved the presentation. L. T. Rolla thanks the hospitality of CWI, where part of this research was done. This work had financial support from CNPq grant 141114/2004-5, FAPERJ, FAPESP grant 07/58470-1, and FSM-Paris.

References

1. G. AMIR AND O. GUREL-GUREVICH, *On fixation of activated random walks*, Elect. Comm. in Probab., 15 (2010), pp. 119–123.
2. E. D. ANDJEL, *Invariant measures for the zero range processes*, Ann. Probab., 10 (1982), pp. 525–547.
3. P. BAK, C. TANG, AND K. WIESENFELD, *Self-organized criticality: An explanation of the $1/f$ noise*, Phys. Rev. Lett., 59 (1987), p. 381.
4. ———, *Self-organized criticality*, Phys. Rev. A, 38 (1988), pp. 364–374.
5. D. DHAR, *The abelian sandpile and related models*, Physica A, 263 (1999), pp. 4–25. arXiv:cond-mat/9808047.
6. ———, *Theoretical studies of self-organized criticality*, Physica A, 369 (2006), pp. 29–70.
7. P. DIACONIS AND W. FULTON, *A growth model, a game, an algebra, Lagrange inversion, and characteristic classes*, Rend. Sem. Mat. Univ. Politec. Torino, 49 (1991), pp. 95–119.
8. R. DICKMAN, *Nonequilibrium phase transitions in epidemics and sandpiles*, Physica A, 306 (2002), pp. 90–97. arXiv:cond-mat/0110043.
9. R. DICKMAN, M. A. MUÑOZ, A. VESPIGNANI, AND S. ZAPPERI, *Paths to self-organized criticality*, Braz. J. Phys., 30 (2000), p. 27.
10. R. DICKMAN, L. T. ROLLA, AND V. SIDORAVICIUS, *Activated random walkers: Facts, conjectures and challenges*, J. Stat. Phys., 138 (2010), pp. 126–142.
11. K. ERIKSSON, *Chip firing games on mutating graphs*, SIAM J. Discrete Math., 9 (1996), pp. 118–128.
12. A. FEY, L. LEVINE, AND Y. PERES, *Growth rates and explosions in sandpiles*, J. Stat. Phys., 138 (2010), pp. 143–159.
13. A. FEY, L. LEVINE, AND D. B. WILSON, *Approach to criticality in sandpiles*, Phys. Rev. E, 82 (2010), p. 031121.
14. ———, *Driving sandpiles to criticality and beyond*, Phys. Rev. Lett., 104 (2010), p. 145703.
15. A. FEY-DEN BOER AND F. REDIG, *Organized versus self-organized criticality in the Abelian sandpile model*, Markov Process. Related Fields, 11 (2005), pp. 425–442. arXiv:math-ph/0510060.
16. J. E. DE FREITAS, L. S. LUCENA, L. R. DA SILVA, AND H. J. HILHORST, *Critical behavior of a two-species reaction-diffusion problem*, Phys. Rev. E, 61 (2000), pp. 6330–6336.
17. U. L. FULCO, D. N. MESSIAS, AND M. L. LYRA, *Critical behavior of a one-dimensional diffusive epidemic process*, Phys. Rev. E, 63 (2001), p. 066118.
18. ———, *Monte carlo study of the critical behavior of a diffusive epidemic process*, Physica A, 295 (2001), pp. 49–52.
19. T. E. HARRIS, *Additive set-valued Markov processes and graphical methods*, Ann. Probability, 6 (1978), pp. 355–378.
20. H. HINRICHSSEN, *Non-equilibrium critical phenomena and phase transitions into absorbing states*, Adv. Phys., 49 (2000), pp. 815–958. arXiv:cond-mat/0001070.

-
21. C. E. HOFFMAN AND V. SIDORAVICIUS, 2004. Unpublished.
 22. H.-K. JANSSEN, *Comment on "critical behavior of a two-species reaction-diffusion problem"*, Phys. Rev. E, 64 (2001), p. 058101.
 23. H. KESTEN AND V. SIDORAVICIUS, *Branching random walk with catalysts*, Electron. J. Probab., 8 (2003), pp. 1–51.
 24. ———, *The spread of a rumor or infection in a moving population*, Ann. Probab., 33 (2005), pp. 2402–2462.
 25. ———, *A phase transition in a model for the spread of an infection.*, Illinois J. Math., 50 (2006), pp. 547–634.
 26. ———, *A shape theorem for the spread of an infection*, Ann. of Math., 167 (2008), pp. 701–766.
 27. R. KREE, B. SCHAUB, AND B. SCHMITTMANN, *Effects of pollution on critical population dynamics*, Phys. Rev. A, 39 (1989), pp. 2214–2221.
 28. S. LÜBECK, *Universal scaling behavior of non-equilibrium phase transitions*, Internat. J. Modern Phys. B, 18 (2004), pp. 3977–4118. arXiv:cond-mat/0501259.
 29. S. S. MANNA, *Two-state model of self-organized criticality*, J. Phys. A: Math. Gen., 24 (1991), pp. L363–L369.
 30. R. MEESTER AND C. QUANT, *Connections between 'self-organised' and 'classical' criticality*, Markov Process. Related Fields, 11 (2005), pp. 355–370.
 31. K. OERDING, F. VAN WIJLAND, J.-P. LEROY, AND H. J. HILHORST, *Fluctuation-induced first-order transition in a nonequilibrium steady state*, J. Statist. Phys., 99 (2000), pp. 1365–1395.
 32. F. REDIG, *Mathematical aspects of the abelian sandpile model*, in Mathematical Statistical Physics – Session LXXXIII, A. Bovier, F. Dunlop, A. van Enter, F. den Hollander, and J. Dalibard, eds., Lecture Notes of the Les Houches Summer School, Elsevier, 2006, pp. 657–730. <http://www.math.leidenuniv.nl/~redig/sandpilelectures.pdf> accessed April 2011.
 33. L. T. ROLLA, *Generalized Hammersley Process and Phase Transition for Activated Random Walk Models*, PhD thesis, IMPA, Rio de Janeiro – Brazil, 2008. arXiv:0812.2473.
 34. L. T. ROLLA AND V. SIDORAVICIUS, *Absorbing-state phase transition for driven-dissipative stochastic dynamics on \mathbb{Z}* . arXiv:0908.1152v2, 2011.
 35. T. SADHU AND D. DHAR, *Steady state of stochastic sandpile models*, J. Stat. Phys., 134 (2009), pp. 427–441.
 36. E. SHELLEF, *Nonfixation for activated random walks*, Alea, 7 (2010), pp. 137–149.
 37. F. VAN WIJLAND, K. OERDING, AND H. J. HILHORST, *Wilson renormalization of a reaction-diffusion process*, Physica A, 251 (1998), pp. 179–201. arXiv:cond-mat/9706197.