Universality of Activated Random Walk

Hyojeong Son

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2026

Reading Committee:

Christopher Hoffman, Chair

Krzysztof Burdzy

Stefan Steinerberger

Shayan Oveis Gharan

Program Authorized to Offer Degree:

Mathematics

University of Washington

Abstract

Universality of Activated Random Walk

Hyojeong Son

Chair of the Supervisory Committee:
Christopher Hoffman
Department of Mathematics

Activated random walk is an interacting particle system proposed as a tractable model of self-organized criticality. Particles are active or sleeping: active particles move by random walk and may fall asleep when isolated, while sleeping particles are reactivated by incoming active particles. A central theme is universality: the critical density and critical local state should be robust under changes in initial condition, driving mechanism, and boundary condition.

This dissertation develops two one-dimensional results under the theme of universality. First, for a supercritical Bernoulli field of sleeping particles on $\mathbb{Z}$ with only one initially active particle, the probability of fixation is strictly between zero and one. Thus the supercritical active phase is not merely an artifact of starting all particles active. Second, for the point-source model started from n active particles at the origin, the probability that a bulk site contains a sleeping particle after stabilization converges to the common critical density. Consequently, any subsequential microscopic local limit of the point-source laws is shift invariant and has the correct one-site density.

TABLE OF CONTENTS

Page

ACKNOWLEDGMENTS

I would first like to thank my advisor, Christopher Hoffman. He made my Ph.D. years genuinely happy ones. I am deeply grateful for the environment he created: one in which I never felt judged, but instead welcomed, trusted, and free to wander mathematically with his full support. From him I learned many aspects of discrete mathematics, but also something equally important: how to treat people with respect and generosity, and how to speak up for what is right.

I am also very grateful to Renato Feres, my master's advisor, for his kindness and for always being willing to spend time talking with me and listening to me, even years after I left WashU. I thank Lionel Levine for his insight and vision. I feel fortunate to have seen how he thinks about problems, not only locally, but in relation to the field as a whole. I also thank Krzysztof Burdzy for teaching me how to be a good teacher of mathematics: not by trying to earn students' affection, but by thinking carefully about how to structure a course so that it truly benefits them. I am grateful to him as well for his many pieces of advice about my career and about being a good mathematician.

I am thankful to the many Ph.D. students and postdocs who made graduate school much brighter. Thank you to Feng, for being a friend with whom I could share worries and build camaraderie; to Yue, for the fun excursions around Seattle; to Jiyoung, for the wonderful chewy Dubai chocolate cookies and for her kindness; to Freda, for teaching me how to run on the Burke-Gilman Trail; to Maddy, for being my buddy when I started my first project, which was also our first project; to Miji, for inspiring me to work hard; to Xiangqian, for being a wonderfully kind sisterly presence with whom I could talk about anything; and to Jacob, for being a friend in the same research area, with whom I could speak honestly about both research and the worries.

I am especially grateful to Minjin, my dear friend from undergraduate school, for her steady support and honest advice through many obstacles from undergrad to Ph.D. As I navigated academic spaces that sometimes felt unfamiliar to me, she was the person with whom I could talk openly, think through difficulties, and find my footing. I feel very lucky to have her as my friend.

I gratefully acknowledge the scholarships and financial support that made it possible for me to pursue my studies, from undergraduate school through my Ph.D. This support changed the course of my life, and I am sincerely thankful for it.

I would like to thank the Probability Seminar for generously providing graduate students with sushi on Mondays after the weekly seminar, and especially for allowing us to choose two items. Mondays were my favorite days during my Ph.D., not only because of the sushi, although the sushi was important, but also because of the warm and caring probability community that made me feel at home.

Finally, I thank my family. I thank my parents for giving me the best childhood memories and for supporting me with so much love. I thank my brother for being such a kind person, including, very importantly, buying me IKEA furniture. Some forms of love come with an Allen wrench, and I appreciate them all.

Chapter 1
INTRODUCTION

Self-organized criticality describes systems that appear to organize themselves near a critical state without external fine tuning of a parameter. Activated random walk is a particularly useful mathematical model in this setting because it combines a simple local mechanism with a rich phase transition. Particles are either active or sleeping. Active particles move as random walks and may become sleeping when isolated; sleeping particles remain still until awakened by an active particle. This mechanism creates avalanches of activity whose large-scale behavior is expected to be insensitive to many microscopic details.

The purpose of this dissertation is to study two one-dimensional manifestations of this robustness under the unifying theme of *universality of activated random walk*. The first manifestation concerns the sharpness of the supercritical phase: above the critical density, the system can remain active even if all particles except one are initially asleep. The second manifestation concerns the point-source aggregate: in the bulk of a large stabilized aggregate, the one-site sleeping probability converges to the same critical density that appears in other formulations of the model.

1.1 Activated random walk

Fix a sleep rate $\lambda > 0$. In the continuous-time activated random walk model on $\mathbb{Z}$, active particles perform independent symmetric random walks at jump rate 1 and attempt to fall asleep at rate λ. A sleep attempt succeeds only when the particle is alone at its site. A sleeping particle is immobile, but it is immediately reactivated when an active particle arrives at its site. Thus the two local reactions may be summarized as

$$a \to s \quad \text{when an active particle is isolated,} \qquad a + s \to 2a.$$

The state of the system is stable on a set if no active particles remain there. For finite systems, one may stabilize by toppling active sites until all particles are sleeping or have exited through the boundary. The Diaconis–Fulton, or site-wise, construction attaches independent stacks of instructions to each site; the abelian property then implies that the final stable configuration is independent of the legal toppling order.

The infinite-volume question asks whether the system fixates. A configuration fixates if every site is visited by active particles only finitely often. The fundamental phase transition is governed by particle density: below a critical density, fixation occurs, while above it activity persists. Rolla, Sidoravicius, and Zindy proved a sharp universality statement for stationary ergodic initial laws with all particles initially active, giving a critical density that depends only on the model parameters and not on the particular ergodic law [20].

1.2 Universality questions

There are several natural ways to make activated random walk critical. In the fixed-energy model, one starts from an infinite or finite-volume configuration of a prescribed density. In the driven-dissipative model, particles are repeatedly added to a finite volume and particles that leave through the boundary are removed. In the point-source model, many active particles are placed at a single site and then stabilized. Universality predicts that these routes to criticality should produce the same density and compatible local critical states.

Recent work of Hoffman, Johnson, and Junge proves the equality of several one-dimensional critical densities, including the fixed-energy, driven-dissipative, cycle, and point-source formulations [9]. This result provides the density-level foundation for the present dissertation. The two chapters here address what this common critical value controls in two different regimes: the persistence of activity above the threshold and the local particle density in the point-source aggregate.

1.3 Overview of the main results

Chapter 3 studies a supercritical Bernoulli field of sleeping particles on $\mathbb{Z}$ with a single active particle inserted at the origin. If the density ρ is larger than the common critical density ρ_*, then for almost every initial sleeping configuration the probability of fixation is strictly between zero and one. The positive lower bound is immediate because the unique active particle may fall asleep before moving. The upper bound is the central result: with positive probability, the single active particle triggers unbounded activity. This shows that the supercritical phase is robust even under an initial condition that strongly favors fixation.

Chapter 4 studies point-source activated random walk on $\mathbb{Z}$. Starting with n active particles at the origin, let η be the final sleeping configuration. For sites $i(n)$ in a fixed macroscopic bulk window around the source, the theorem proves

$$\mathbb{P}\Big(\eta(i(n)) = \mathfrak{s}\Big) = \rho_c + o(1),$$

where ρ_c is the common critical density. The proof combines an IDLA flattening step with driven-dissipative estimates, and it also shows that any subsequential microscopic local limit of the point-source laws is shift invariant.

Chapter 2
PRELIMINARIES

In this chapter, we emphasize the role of the activated random walk as a candidate model for self-organized criticality, as well as the importance of the density and local-limit conjectures.

2.1 From sandpiles to activated random walk

Bak, Tang, and Wiesenfeld introduced self-organized criticality as a mechanism for scale-free behavior in driven dissipative systems [2]. The abelian sandpile model became a central mathematical example. It has a deterministic toppling rule and a rich algebraic structure, but it also exhibits features that are sensitive to the underlying lattice and boundary mechanism [6, 12, 21, 22]. Stochastic sandpiles and activated random walk were developed as more flexible stochastic alternatives [3, 19].

Activated random walk has two advantages that are central in this dissertation. First, individual particle trajectories are random walks, which allows the use of probabilistic tools such as martingales, concentration inequalities, and comparison with internal diffusion limited aggregation. Second, the sleeping mechanism gives a natural absorbing-state phase transition and a natural driven-dissipative dynamics, allowing the same critical density to be approached from several directions.

2.2 Critical density and fixation

For stationary ergodic initial laws on $\mathbb{Z}^d$ with all particles initially active, Rolla, Sidoravicius, and Zindy proved that fixation is determined by the mean density: there is a critical value such that subcritical systems fixate almost surely and supercritical systems remain active almost surely [20]. This result gives a rigorous threshold for the fixed-energy formulation.

The density conjecture asserts that the stationary density of driven-dissipative activated random walk in large finite boxes should coincide with the fixed-energy critical density. Hoffman, Johnson, and Junge proved this in one dimension and also identified the same value in the point-source and cycle formulations [9]. Thus one-dimensional activated random walk has a single critical density that appears across several natural experiments. The first contribution of this dissertation strengthens the interpretation of the supercritical side of this threshold, while the second uses the same density to identify a local statistic of the point-source aggregate.

2.3 Point-source aggregates and microscopic limits

In the point-source model, one starts with n active particles at the origin and stabilizes the system. The final sleeping configuration is an aggregate of n particles. Levine and Silvestri formulate a collection of conjectures describing the large-scale shape, microscopic local limits, translation invariance, density, and hyperuniformity of such aggregates and related stationary states [15]. In this framework, universality means not only that different constructions have the same critical density, but also that they produce the same limiting local random configuration.

The local-density result in Chapter 4 addresses the one-site marginal part of this picture in one dimension. It does not prove full convergence of the microscopic point-source law, but it shows that any subsequential limit has the correct density and is shift invariant. This is evidence that the critical state seen from the bulk of a large point-source aggregate is governed by the same critical density as the driven-dissipative and fixed-energy models.

2.4 The abelian property and legal topplings

Each site has an infinite stack of independent instructions, each of which is a left jump, right jump, or sleep instruction. A legal toppling applies the next unused instruction at an unstable site. The abelian property states that, whenever stabilization is possible, the final stable configuration and odometer do not depend on the legal order of topplings [19]. This permits the use of carefully chosen toppling procedures that are tailored to the proof.

In Chapter 3, stabilization is performed on increasing intervals, and the proof controls the probability that the next scale fails to remain good. In Chapter 4, the toppling order begins with an IDLA-type flattening phase, then relates the remaining stabilization inside a block to a driven-dissipative process. These procedures are different, but both rely on the same abelian mechanism.

2.5 Internal DLA and local flatness

Internal diffusion limited aggregation appears in Chapter 4 as a comparison process. In one dimension, the IDLA cluster generated by particles released from adjacent sources can be coupled using the Eulerian distribution of descents. This coupling shows that neighboring one-site occupation probabilities in the point-source ARW aggregate differ by a small amount. Combined with block-average estimates at the critical density, this local flatness yields convergence of single-site probabilities in the bulk.

Chapter 3

SUPERCRITICAL ACTIVATED RANDOM WALK ON $\mathbb{Z}$ IS EXPLOSIVE

This chapter is based on joint work with Madeline Brown and Christopher Hoffman. It studies a supercritical one-dimensional activated random walk started from a Bernoulli field of sleeping particles with a single initially active particle.

3.1 Introduction

In 1987, Bak, Tang, and Wiesenfeld introduced the concept of *self-organized criticality* (SOC) [2], a property observed in complex natural systems where energy accumulates slowly and is released intermittently. SOC describes systems that naturally evolve to a critical state without the need for fine-tuning parameters or external influences. Examples include financial markets, where variations in stock and commodity prices follow a power-law distribution, and forest fires, where accumulated flammable material can lead to large fires ignited by small sparks. These examples highlight the ubiquity and significance of SOC in various artificial and natural systems.

Since the concept of SOC was proposed, the search for a universal SOC model has led to extensive research on the *deterministic sandpile model*. In this model, each vertex in a graph has a nonnegative number of chips, and a vertex can topple when the number of chips at that vertex equals or exceeds its degree. A toppling vertex distributes one chip to each neighboring vertex, which can trigger a chain of topplings. This deterministic model exhibits intricate fractals and patterns [14]. However, its deterministic nature limits its ability to exhibit certain critical behaviors [5, 6, 11]. Ideally, a universal model should exhibit universality in the sense that macroscopic properties are independent of microscopic details, making the system robust to perturbations.

The probabilistic variant of the deterministic sandpile model is the *stochastic sandpile model*. In this model, instead of sending one chip to each neighboring vertex when a site topples, a fixed number of particles are sent to neighboring vertices chosen independently according to some probability distribution. There is evidence that this stochastic model features universality [3]. However, this model involves pairwise correlations in particle movements, which complicates the analysis of its dynamics.

More recently, the *Activated Random Walk* (ARW) model has emerged as a promising candidate for a universal SOC model. At each site, the initial number of particles is i.i.d. and sampled from an ergodic distribution with mean ρ. In this continuous-time interacting particle system on $\mathbb{Z}$, each particle is either active or sleeping. Active particles perform symmetric random walks at rate 1 and fall asleep at rate $\lambda \in (0, \infty)$. Sleeping particle remains stationary until awakened by an active particle and is instantly reactivated if other particles are present at the site. Since particles in the model perform random walks and fall asleep independently, the ARW model is more tractable than the stochastic sandpile model and has yielded significant findings.

One of primary questions in ARW dynamics is whether the system fixates or remains active, and under what conditions. Fixation occurs when every site is visited by active particles only finitely often, while the system stays active if there is ongoing particle movement and interaction. One might conjecture that the critical density, where a phase transition occurs, depends on the specific version of the ARW model being considered.

Recently, Hoffman, Junge, and Johnson [9] proved that for each $\lambda > 0$, the critical densities in the driven-dissipative model, the point-source model, the fixed-energy model on $\mathbb{Z}$, and the fixed-energy model on the cycle all exist and are equal, denoted by ρ_*. This result confirms the *density conjecture*, reinforcing the ARW model as a promising candidate for universality due to its well-defined critical density ρ_*.

Furthermore, the long-term behavior of a system depends on the particle density ρ and the sleep rate λ. Specifically, Rolla, Sidoravicius, and Zindy demonstrated the following result, which highlights a phase transition based on these parameters:

Theorem 3.1 (Theorem 1 of [20])**.** *Consider the ARW model on $\mathbb{Z}^d$ for a fixed $d \geq 1$ with a given sleep rate λ. There exists a critical density ρ_* such that for any spatially ergodic distribution ν supported on configurations where all particles are initially active, the following holds: a configuration sampled from ν almost surely fixates if the average density ρ is less than ρ_*, and almost surely remains active if ρ is greater than ρ_*.*

This theorem allows us to choose any shift-invariant ergodic measure for the initial configuration. However, there is an assumption that all particles should initially be active. In this chapter, we weaken this assumption to show that even with only one active particle, there is a positive probability that the system remains active.

A very similar result was recently proved by Forien [7]. He studied activated random walk on an interval where particles can leave through the endpoints. He showed that in the supercritical regime, a positive density of particles leaves the interval with positive probability. This chapter significantly strengthens his result.

We consider a starting configuration $\omega \in \{0, \mathfrak{s}\}^{\mathbb{Z}}$. In this configuration, $\omega(i) = 0$ denotes that site i is empty, and $\omega(i) = \mathfrak{s}$ indicates that site i contains a sleeping particle. The number of particles at each site is an independent Bernoulli random variable with mean $\rho \in (0, 1)$, referred to as the particle density.

Given a configuration $\omega \in \{0, \mathfrak{s}\}^{\mathbb{Z}}$ we define a new configuration ω^* by setting $\omega^*(i) = \omega(i)$ for all $i \neq 0$ and placing one active particle at the origin, that is, $\omega^*(0) = 1$. For any configuration $\nu \in \{0, \mathfrak{s}, 1, 2, \dots\}^{\mathbb{Z}}$, let $\mathbb{P}_\nu$ represent the probability measure of the system's evolution starting from ν.

We are now in position to describe our main theorem. Let `Fixation` be the event that the system eventually fixates. Our goal is to prove the following theorem:

Theorem 3.2. *Let $\omega \in \{0, \mathfrak{s}\}^{\mathbb{Z}}$ be a configuration with particle density $\rho > \rho_*$, where all particles are initially asleep and the $\omega(i)$ are i.i.d. Bernoulli random variables with expectation ρ. For almost every ω^*,*

$$0 < \mathbb{P}_{\omega^*}(\texttt{\textit{\textbf{Fixation}}}) < 1.$$

The property that $\mathbb{P}_{\omega^*}(\texttt{\textbf{Fixation}}) < 1$ is often referred to as the system being explosive.

The lower bound is easy while the upper bound is significantly more involved.

Proof of Lower Bound. At the beginning of our process, there is only one active particle at the origin. If this particle falls asleep before moving, then there are no active particles available to awaken any remaining sleeping particles. As a result, the system achieves **`Fixation`** immediately.

An active particle either jumps to a neighboring site at rate 1 or falls asleep at rate λ. The total rate of these actions is $1 + \lambda$. Therefore, the probability that the system fixates is at least $\frac{\lambda}{1+\lambda}$. □

Theorem 3.2 shows that the critical value is quite sharp. Below the critical value a configuration least likely to fixate (one with all particles initially active) still fixates. Above the critical value a configuration most likely to fixate (one with all but one particle initially asleep) still has a positive probability of not fixating.

3.2 Site-wise Construction of ARW

In our proof of the upper bound, we utilize a site-wise representation for the ARW, where infinite stacks of instructions are attached to each site. This construction, underpinned by the abelian Property (Lemma 3.3), allows us to focus on the moves taken by particles at specific sites without concerning ourselves with the order in which they occur.

Consider the ordered set $\mathbb{N} \cup \{\mathfrak{s}\}$ with $0 < \mathfrak{s} < 1 < 2 < \cdots$. A *configuration* is an element $\omega \in \{0, \mathfrak{s}, 1, 2, 3, \dots\}^{\mathbb{Z}}$, where $\omega(x)$ gives the state of site $x \in \mathbb{Z}$. If $\omega(x) = 0$, there is no particle at site x; if $\omega(x) = \mathfrak{s}$, there is one sleeping particle; and if $\omega(x) = n \geq 1$, there are n active particles.

When an active particle visits a site with a sleeping particle, it wakes up the sleeping particle, resulting in $\mathfrak{s} + 1 = 2$. If a site is given a sleep instruction, the effect depends on the number of active particles at the site. Precisely, if there is exactly one active particle at the site, the particle becomes sleeping, so the state changes from 1 to $\mathfrak{s}$. If there are $n \geq 2$ active particles at the site, the sleep instruction has no effect, and the state remains at n. We define a site x to be *stable* if $\omega(x) < 1$ (i.e., if there are no active particles at x), and *unstable* otherwise. By defining $|\mathfrak{s}| = 1$, we denote the number of particles at site x by $|\omega(x)|$.

Each site $x \in \mathbb{Z}$ has an infinite *stack* of instructions $\mathsf{Instr}_x(k)$, where k denotes the k-th instruction at site x. The distribution of $\mathsf{Instr}_x(k)$ is:

$$\mathsf{Instr}_x(k) = \begin{cases} \texttt{Left} & \text{with probability } \frac{1}{2(1+\lambda)}, \\ \texttt{Right} & \text{with probability } \frac{1}{2(1+\lambda)}, \\ \texttt{Sleep} & \text{with probability } \frac{\lambda}{1+\lambda}. \end{cases}$$

Each $\mathsf{Instr}_x(k)$ is independent for all $x \in \mathbb{Z}$ and $k \in \mathbb{N}^+$. And the rate at which we execute unused instructions at each site x is given by $\mathbb{1}_{\omega(x) \neq \mathfrak{s}} |\omega(x)|(1+\lambda)$. The `Left` (respectively, `Right`) instruction subtracts 1 from $\omega(x)$ and adds 1 to $\omega(x-1)$ (respectively, $\omega(x+1)$). Executing a `Sleep` instruction at a site with exactly one active particle changes $\omega(x) = 1$ to $\mathfrak{s}$; if there are multiple active particles at the site, the sleep instruction has no effect. We can *topple* a single site x by executing its first unused instruction. Alternatively, we can *topple* a sequence of sites $\alpha = (x_1, x_2, \dots, x_l)$ by starting with x_1 and proceeding to x_l.

To track the number of topplings at each site, we introduce an *odometer* function

U:

$$U : \mathbb{Z} \to \mathbb{N},$$

where $U(x)$ denotes the number of times site x has been toppled. Additionally, $U_\alpha(x)$ represents the number of times site x appears in the sequence α. When toppling site x, we update the configuration ω and the odometer reading U by defining the toppling operation acting on (ω, u) as:

$$\Phi_x(\omega, U) = (\mathsf{Instr}_x(U(x) + 1)(\omega), U + \delta_x),$$

where δ_x is a function on $\mathbb{Z}$ that takes the value 1 at x and 0 elsewhere. We say that Φ_x is *legal* if site x is unstable. A toppling sequence Φ_α is a *sequence of legal topplings* if, for each $1 \leq i \leq l$, the toppling Φ_{x_i} is legal after performing the preceding topplings $\Phi_{x_{i-1}} \cdots \Phi_{x_2}\Phi_{x_1}$. We call α *stabilizing* if, after performing all topplings in the sequence α, no unstable sites remain.

The abelian property is an essential feature of the ARW, enabling us to utilize the toppling procedure described in Section 3.3, where we stabilize the system one interval at a time. We formalize this property in the following lemma:

Lemma 3.3 (Abelian Property [19])**.** *Let α and β be two legal sequences of topplings of a configuration ω. Suppose $U_\alpha(x) = U_\beta(x)$ for all $x \in \mathbb{Z}$. Then their final configurations are identical, that is,*

$$\Phi_\alpha(\omega) = \Phi_\beta(\omega).$$

We will also make use of a particle-wise representation of ARW. In this representation, each particle has its own stack of i.i.d. instructions. The stacks for each particle are also independent. By the abelian property, at any time we can choose to move any active particle. We do this by executing the next unused instruction for that active particle. If this particle-wise process terminates for a given initial configuration, then it will have the same final distribution on the location of particles and odometers as the site-wise representation has.

3.3 Good Configurations

3.3.1 Toppling Procedure

By the abelian property (Lemma 3.3), the stabilization of the system is independent of the order in which sites are toppled. Therefore, we can perform the following inductive partial stabilization procedure, starting from the configuration ω^* with one active particle at the origin:

1. Given ρ, choose $\delta < (\rho - \rho_*)/3$ and a sufficiently large integer $k > 200/\delta^2$.

2. For each positive integer n, consider the interval $I_n = [-k^n, k^n]$.

3. For each n, proceed inductively as follows:

   (a) Evolve the system within the interval I_n, freezing active particles at the boundaries of I_n, until no active particles remain in the interior of I_n.

   (b) If active particles are present at the boundaries of I_n, extend the interval to I_{n+1} and repeat the process.

Throughout the remainder of this chapter, when we refer to the odometer, we are considering the odometer associated with the stabilization step on the interval I_{n+1}. Specifically, the stabilization begins with the configuration ω_{initial} and concludes with the configuration ω_{final} on I_{n+1}. Also, the nth step stabilization refers to our toppling procedure for stabilizing the interval I_n. Let **`Positive Odometer`** denote the subset of $(-k^{n+1}, k^{n+1})$ where the odometer is positive during this stabilization step.

3.3.2 Good Configurations

Our goal is to show that, for a 'typical' configuration in our stabilization process, the probability that the system stabilizes without active particles at the boundaries of I_{n+1} is very small, given that it satisfied the previous stabilization step within the interval

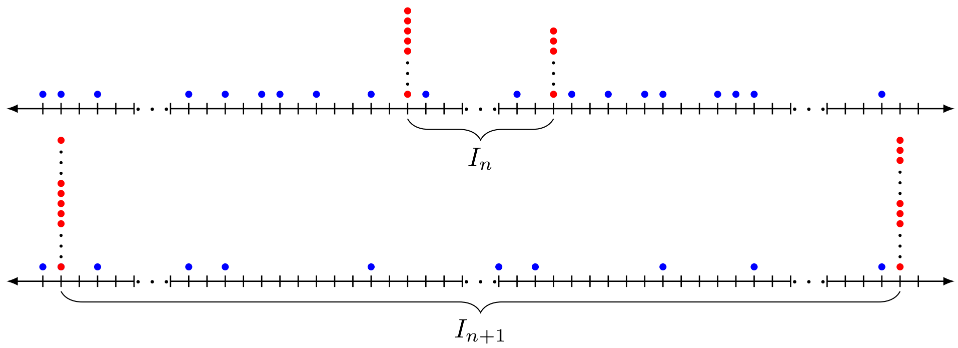

Figure 3.1: Example of the configurations ω_{initial} (top) and ω_{final} (bottom) during the stabilization of the interval I_{n+1} according to the toppling procedure. Red circles represent active particles, while blue circles represent sleeping particles.

I_n. However, there exists a class of configurations in which no active particles remain at the boundaries after stabilization. To address this, we introduce the concept of a *good* configuration. Let $\gamma < \delta/10$ be chosen such that both $k\gamma$ and $1/\gamma$ are integers.

Definition 3.4. A configuration ω is *good* for a positive integer n if ω satisfies the following criteria for the nth stabilization step:

Criterion 1 The interval I_n contains at least $2k^n(\rho_* + 2\delta)$ particles.

Criterion 2 Every subinterval of length γk^n within $I_n \setminus \{-k^n, k^n\}$ contains at most $\gamma k^n(\rho_* + \delta)$ particles.

3.3.3 Prior Results

To estimate the probability that the system fails to satisfy the criteria for a good configuration at step $n + 1$, given that it satisfies the criteria at step n, we use the following recent theorem.

Theorem 3.5 (Theorem 8.4 of [9])**.** *Let n be a positive integer. Consider activated random walk on $\mathbb{Z}$ with sleep rate $\lambda > 0$. Let σ be an initial configuration with no* *sleeping particles on the interval* $[a, a+n-1]$ *for some integer* a*. Let* X_n *be the number of particles left sleeping in the stabilization of* σ *on* $[a, a+n-1]$*. For any* $\delta > 0$*,*

$$\mathbb{P}_\sigma[X_n \geq (\rho_* + \delta)n] \leq Ce^{-cn},$$

where C *and* c *are positive constants depending only on* λ *and* δ*.*

In particular, we will use the following variant of the theorem.

Corollary 3.6. *Let* n *be a positive integer. Consider activated random walk on* $\mathbb{Z}$ *with sleep rate* $\lambda > 0$*. Let* σ *be an initial configuration with at least one active particle on the interval* $[a, a+n-1]$ *for some integer* a*.*

Let b *be an integer and* $\gamma \in (0, 1]$ *such that* $I_{b,\gamma} = [b, b+n\gamma-1] \cap \mathbb{Z} \subseteq I$*. Let* U *be a stable odometer on* $I_{b,\gamma}$ *that satisfies* $U(b) = u_0$ *and* $U(x) > 0$ *for every site* $x \in I_{b,\gamma}$ *containing a particle in* σ*. Define* $Y_{I_{b,\gamma}}(u_0)$ *to be the number of sleeping particles remaining in* $I_{b,\gamma}$ *after applying* U*. Then, for any* $\rho > \rho_*(\lambda)$*,*

$$\mathbb{P}_\sigma\left[Y_{I_{b,\gamma}}(u_0) \geq \rho n\gamma\right] \leq Ce^{-cn}$$

for some constants C *and* c*, depending only on* λ *and* ρ*.*

Proof. This is a restatement of equation (81) in the proof of Theorem 8.4 in [9] (plus an application of the union bound). □

3.3.4 Transition Probabilities

Let G_n denote the set of all good configurations on the interval I_n at step n. We define the event

$$D_{n+1} = \text{the resulting configuration on } I_{n+1} \text{ is not good or the odometer is somewhere equal to zero.}$$

Our goal is to analyze the probability of the event D_{n+1} conditioned on configurations $\tilde{\omega} \in G_n$. With a slight abuse of notation for $\tilde{\omega} \in G_n$ we write $\mathbb{P}_{\tilde{\omega}}$ to indicate the randomness of the distribution of particles in I_{n+1} at the end of the stabilization of I_{n+1}. To achieve this, we partition D_{n+1} into four distinct cases, denoted D^1_{n+1}, D^2_{n+1}, D^3_{n+1}, and D^4_{n+1}, based on the following considerations:

1. **Odometer Considerations:**

   - **Case** D^1_{n+1}: The sum of the odometers in I_{n+1} exceeds $k^{3.5(n+1)}$.
   - **Case** D^2_{n+1}: The sum of the odometer does not exceed $k^{3.5(n+1)}$, and the odometer is somewhere equal to zero.

2. **Criterion Considerations:** Identify which of the two criteria the configuration fails to satisfy.

   - **Case** D^3_{n+1}: The sum of the odometer does not exceed $k^{3.5(n+1)}$, and the configuration fails Criterion 1.
   - **Case** D^4_{n+1}: The sum of the odometer does not exceed $k^{3.5(n+1)}$, the odometer is everywhere positive, and the configuration fails Criterion 2.

This decomposition allows us to systematically analyze and bound the probability of D_{n+1} by considering each of these four cases separately.

Lemma 3.7.

$$D_{n+1} \subset \bigcup_{i=1}^{4} D^i_{n+1}.$$

Proof. Consider the case where the sum of the odometer readings in I_{n+1} exceeds $k^{3.5(n+1)}$. This scenario is captured by D^1_{n+1}. Suppose, instead, that the sum does not exceed $k^{3.5(n+1)}$. In this situation, either the odometer is zero at some location in I_{n+1}, leading to D^2_{n+1}, or the odometer is positive everywhere but the configuration is not good. In the latter case, the configuration must fail either condition Criterion 1, which implies D^3_{n+1}, or condition Criterion 2, which implies D^4_{n+1}. $\square$

Lemma 3.8. *For positive constants C and c depending on λ, we have*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}}\left(D^1_{n+1}\right) \le Ce^{-ck^{0.5(n+1)}}.$$

Proof. We use a particle-wise representation of ARW. There are at most $2k^{n+1}+1 \le 3k^{n+1}$ particles in I_{n+1}. If the sum of the odometers in I_{n+1} exceeds $k^{3.5(n+1)}$, then there must exist at least one particle that executes at least $(1/3)k^{2.5(n+1)}$ instructions (either steps or sleep instructions).

Each time a particle takes k^{2n+2} steps, it has a probability of at least $c'(\lambda) > 0$ of exiting I_{n+1}. Therefore, the probability that a particle takes more than $k^{2.5(n+1)}$ steps without reaching the boundaries of I_{n+1} decays exponentially in $k^{0.5(n+1)}$.

Applying the union bound over all particles in I_{n+1}, we obtain the desired result. □

Lemma 3.9. *There exist constants C and c such that*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}}\left(D^3_{n+1}\right) \le Ce^{-ck^{n+1}}.$$

Proof. Let X be the number of particles in I_{n+1}. Initially, each site contains one sleeping particle, following a Bernoulli distribution with mean ρ. In our toppling procedure at step $n+1$, we ensure that no particle escapes from I_{n+1}. Thus, the expectation of X is $\rho(2k^{n+1}+1)$.

We consider the event that X is less than $2k^{n+1}(\rho_*+2\delta)$. Since $\rho > \rho_* + 3\delta$, this value is less than the expected value by a fixed proportion of δ. By the Chernoff bound, the probability that X deviates below its mean by a fixed fraction decays exponentially in k^{n+1}. □

Lemma 3.10. *There exist positive constants C and c such that*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}}\left(D^4_{n+1}\right) \le Ce^{-ck^{n+1}}.$$

Proof. For D^4_{n+1} to occur, there must exist a subinterval of length γk^{n+1} with a particle density exceeding $\rho_* + \delta$ after the $(n+1)$th stabilization. There are at most $2k^{n+1}$

choices of subinterval and $k^{3.5(n+1)}$ choices of odometer at the leftmost edge of the subinterval. By Corollary 3.6, the probability of such an event for each of these choices decays exponentially in k^{n+1}. Applying the union bound over all possible subintervals and choices of the odometer we obtain the desired result. $\square$

3.4 Completing the Proof

We will break up the event D^2_{n+1} into pieces depending on the location of the leftmost and/or rightmost site with a zero odometer.

1. F^1_{n+1}: D^2_{n+1} occurs and there is no zero odometer $\geq k^n$,
2. F^2_{n+1}: D^2_{n+1} occurs and there is no zero odometer $\leq -k^n$ and
3. F^3_{n+1}: D^2_{n+1} occurs and there are zero odometers both $\leq -k^n$ and $\geq k^n$.

In the case that F^3_{n+1} occurs we define z to be the smallest value with positive odometer and z' to be the largest value with positive odometer. Then we can partition

$$F^3_{n+1} = \bigcup_{z',z} F^{3,z',z}_{n+1}.$$

Lemma 3.11.

$$D^2_{n+1} \subset F^1_{n+1} \cup F^2_{n+1} \cup \left(\bigcup_{z,z'} F^{3,z,z'}_{n+1} \right).$$

Proof. If D^2_{n+1} occurs and F^1_{n+1} and F^2_{n+1} do not occur then F^3_{n+1} must occur. Then there are z and z' such that $F^{3,z,z'}_{n+1}$ occurs. $\square$

To bound the probabilities of these events, recall that for a configuration ω, $\omega(i)$ represents the number of particles at site i. Moreover, we define that if there is a sleeping particle at site i, then $\omega(i) = 1$. Using this description, our next definition will characterize the environment in $I_{n+1} \setminus I_n$ before stabilizing.

Definition 3.12. We say that $\omega_{\text{initial}} \in \{0, 1, 2, \dots\}^{\{-k^{n+1},\dots,k^{n+1}\}}$ is **`plentiful`** if, for every $j \in \{1, 2, \dots, (k-1)/\gamma\}$, the following inequalities hold:

$$\sum_{i=k^n+(j-1)\gamma k^n+1}^{k^n+j\gamma k^n} \omega_{\text{initial}}(i) \geq \gamma k^n(\rho_* + 2\delta)$$

and

$$\sum_{i=k^n+(j-1)\gamma k^n+1}^{k^n+j\gamma k^n} \omega_{\text{initial}}(-i) \geq \gamma k^n(\rho_* + 2\delta).$$

The next definition describes the configuration that we expect to see after stabilization in the region **`Positive Odometer`**.

Definition 3.13. We say that $\omega_{\text{final}} \in \{0, 1, 2, \dots\}^{\{-k^{n+1},\dots,k^{n+1}\}}$ is **`sparse`** in $R \subset \mathbb{Z}$ if, for every integer j such that $[j+1, j+\gamma k^n] \subseteq R$, the following inequality holds:

$$\sum_{i=j+1}^{j+\gamma k^n} \omega_{\text{final}}(i) \leq \gamma k^n(\rho_* + \delta).$$

Next we bound the probabilities that the initial and final configurations are atypical.

Lemma 3.14. *There exist positive constants C and c such that*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}}\left(\omega_{initial} \textit{ is not } \texttt{plentiful}\right) \leq Ce^{-ck^n}.$$

Proof. In each interval, the total number of particles is the sum of γk^n independent Bernoulli random variables with mean ρ, where $\rho > \rho_* + 3\delta$. The probability that the total number of particles in an interval is less than $\gamma k^n(\rho_* + 2\delta)$ corresponds to a deviation below the expected value by at least $\gamma k^n(\rho - \rho_* - 2\delta) \geq \gamma k^n \delta$. Applying the Chernoff bound, the probability of such a deviation is exponentially small in k^n. Since there are $2(k-1)/\gamma$ intervals, the union bound implies that the probability that ω_{initial} is not **`plentiful`** is $2(k-1)/\gamma$ times an exponentially decreasing term. $\square$

Lemma 3.15. *There exist positive constants C and c such that*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}}\left(\omega_{final} \textit{ is not } \texttt{sparse} \textit{ in } \texttt{Positive Odometer}\right) \leq Ce^{-ck^n}.$$

Proof. This follows by the same argument used in Lemma 3.10. □

Lemma 3.16. *Let $\omega_{initial},\ \omega_{final} \in \{0,1,2,\dots\}^{\{-k^{n+1},\dots,k^{n+1}\}}$ be the initial and final configurations of the stabilization in I_{n+1}. If*

1. *there exist $z \le k^n$ such that $z \notin$* **`Positive Odometer`**,
2. *$\omega_{initial}$ is* **`plentiful`** *and*
3. *ω_{final} is* **`sparse`** *in $[z, k^{n+1})$,*

then

$$\sum_{i=z}^{k^{n+1}} i\ (\omega_{final}(i) - \omega_{initial}(i)) > (\delta/10)(k^{n+1})^2. \tag{3.1}$$

Proof. We will construct a pairing between the particles in ω_{initial} and ω_{final} that satisfies

1. at least $(\delta/2)k^{n+1}$ particles in ω_{initial} are paired with a particle in ω_{final} that is more than $(1/3)k^{n+1}$ to its right.
2. no particle in ω_{final} is paired with a particle in ω_{initial} that is more than $(3/\delta)k^n$ to its left,
3. at most $(3/\delta)k^n$ particles in ω_{final} is paired with a particle in ω_{initial} that is more than γk^n to its left,

These conditions imply that the sum on the left hand side of (3.1) is at least

$$\begin{aligned}(\delta/2)k^{n+1}(1/3)k^{n+1} - ((3/\delta)k^n)^2 - (\gamma k^n)(2k^{n+1}) &\ge (k^{2n})(k^2\delta/6 - 9/\delta^2 - 2\gamma k)\\ &\ge (k^{2n})(k^2\delta/10).\end{aligned}$$

The last inequality holds for sufficiently large k. Thus (3.1) holds.

Since ω_{initial} is **`plentiful`** and ω_{final} is **`sparse`** in $[z, k^{n+1}]$, in $[-k^n, k^n + (2/\delta)k^n]$ there are more particles in ω_{initial} than in ω_{final}. Thus we can pair every particle in this region in ω_{final} to a particle in the same region in ω_{initial}. This interval has length at most $(3/\delta)k^n$ and it has at most $(3/\delta)k^n$ particles in ω_{final}.

In the region from

$$[k^n + (2/\delta)k^n, k^{n+1}/2] \supset [k^{n+1}/10, k^{n+1}/2]$$

there are at least $\gamma k^{n+1}/3$ particles in ω_{initial} that are paired with a particle at k^{n+1}. Thus we can create the desired pairing. □

Lemma 3.17. *Let $U_{n+1}(x)$ be the odometer at x in the stabilization at stage $n+1$. Let* `Tilt` *be the event that*

1. $\sum_{x=-k^{n+1}}^{k^{n+1}} U_{n+1}(x) \leq k^{3.5(n+1)}$ *and*

2. *there exists* $z \notin$ `Positive Odometer` *with*

$$\sum_{j=z}^{k^{n+1}} j(\omega_{\mathit{final}}(j) - \omega_{\mathit{initial}}(j)) > (\delta/10)k^{2n+2}.$$

There exists C and c

$$\sup_{\tilde{\omega} \in G_n} \mathbb{P}_{\tilde{\omega}}(\texttt{Tilt}) \leq Ce^{-ck^{.5n}}$$

Proof. If $z \notin$ `Positive Odometer` then all of the particles that start to the right of z remain to the right of z throughout the stabilization process. Choose an order of updating the particles and call the configurations

$$\omega_0 = \omega_{\text{initial}}, \omega_1, \ldots, \omega_J = \omega_{\text{final}}.$$

Then

$$X(j) = \sum_{x=z}^{k^{n+1}} x(\omega_j(x) - \omega_{j-1}(x))$$

for $j = 1, 2, \ldots, J$ is a martingale where the step size is bounded by 1. Also note that

$$J \leq \sum_{x=-k^{n+1}}^{k^{n+1}} U_{n+1}(x) \leq k^{3.5(n+1)}.$$

By the Azuma-Hoeffding inequality we have that

$$\mathbb{P}(\texttt{Tilt}) \leq Ce^{-ck^{.5n}}.$$

□

Lemma 3.18. *There exists C and c such that*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}} F^1_{n+1} \leq Ce^{-ck^n} \qquad \textit{and} \qquad \sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}} F^2_{n+1} \leq Ce^{-ck^n}.$$

Proof. Suppose that F^1_{n+1} occurs and ω_{initial} is `plentiful`. And let z be the largest site with zero odometer. Then we have that z is less than k^n.

For the first inequality by Lemma 3.16 these events only happen if ω_{initial} is not `plentiful`, ω_{final} is not `sparse` to the right of z' or (3.1) holds (the center of mass changes a lot.) The probabilities of these events are bounded by Lemmas 3.14, 3.15 and 3.17.

The second inequality follows by symmetry. □

Lemma 3.19. *There exists C and c such that*

$$\sup_{\tilde{\omega}\in G_n} \mathbb{P}_{\tilde{\omega}} \left(\bigcup_{z,z'} F^{3,z,z'}_{n+1} \right) \leq Ce^{-ck^{n+1}}.$$

Proof. First we deal with the case that the set of positive odometers is connected. As $\tilde{\omega} \in G_n$ if ω_{initial} is `plentiful` then for every $z < k^n$ and $z' > k^n$ there are at least

$$(z' - z)(\rho_* + 2\delta) - 2\gamma k^n$$

particles in $[z, z']$. If $F^{3,z,z'}_{n+1}$ occurs then all of these particles are in $[z, z']$ after stabilization. If ω_{final} is `sparse` in (z, z') then there are at most

$$\begin{aligned}
(z' - z)(\rho_* + \delta) + 2\gamma k^n &= (z' - z)(\rho_* + 2\delta) - \delta(z' - z) + 2\gamma k^n \\
&\leq (z' - z)(\rho_* + 2\delta) - 2\gamma k^n - \delta(2k^n) + 4\gamma k^n \\
&< (z' - z)(\rho_* + 2\delta) - 2\gamma k^n
\end{aligned}$$

particles in $[z, z']$. Thus $F_{n+1}^{3,z,z'}$ only happens if ω_{initial} is not `plentiful` or ω_{final} is not sparse in $[z, z']$.

If the set of positive odometers is not connected then the set **`Positive Odometer`** is of the form $[z, z_1] \cup [z_2, z']$ for some z_1 and z_2. A very similar computation to the one above again shows that either ω_{initial} is not `plentiful` or ω_{final} is not sparse in **`Positive Odometer`**. Combining Lemmas 3.14 and 3.15 and summing up over all choices of z, z', z_1 and z_2 completes the proof. □

Lemma 3.20. *There exist C and c such that*

$$\sup_{\tilde{\omega} \in G_n} \mathbb{P}_{\tilde{\omega}} \left(D_{n+1}^2 \right) \leq C e^{-ck^{n^{0.5}}}.$$

Proof. By Lemmas 3.11, 3.18, and 3.19 and the union bound.

□

Lemma 3.21. *There exist positive constants C and c such that*

$$\sup_{\tilde{\omega} \in G_n} \mathbb{P}_{\tilde{\omega}} \left(D_{n+1} \right) \leq C e^{-ck^{n^{0.5}}}.$$

Proof. By Lemmas 3.8, 3.20, 3.9, and 3.10, the probabilities of the events D_{n+1}^1, D_{n+1}^2, D_{n+1}^3, and D_{n+1}^4 decay exponentially in k^{n+1} or $k^{0.5(n+1)}$. Therefore, by Lemma 3.7 there exist positive constants C and c such that

$$\sup_{\tilde{\omega} \in G_n} \mathbb{P}_{\tilde{\omega}} \left(D_{n+1} \right) \leq C e^{-ck^{n^{0.5}}}.$$

□

Proof of Upper Bound of Theorem 3.2. From Lemma 3.21 and the Borel-Cantelli lemma we get that the probability that D_n does not occur for any n is positive. By the definition of D_n if D_n does not occur then every particle in I_n moves at least once in the nth stabilization stage. Thus if D_n does not occur for any n then every particle moves infinitely often. Thus the system does not stabilize. □

Chapter 4

LOCAL DENSITY OF ACTIVATED RANDOM WALK ON $\mathbb{Z}$

This chapter is based on joint work with Christopher Hoffman and Jacob Richey. It studies the local one-site density in the point-source model and identifies the limiting single-site occupation probability in the bulk with the critical density.

4.1 Introduction

Many complex, real-world systems share a common pattern: they accumulate energy gradually and release it in sudden, cascading bursts, which are statistically scale-free. This phenomenon arises in earthquake dynamics, where stress accumulates progressively along geological fault lines before dissipating abruptly through seismic events of varying magnitudes, from negligible tremors to catastrophic ruptures. Similarly, financial markets display this behavior through price fluctuations in stocks and commodities. These observations led Bak, Tang, and Wiesenfeld in 1987 to introduce the unifying notion of *self-organized criticality* (SOC) [2]. The defining characteristic of SOC is that the system is autonomously driven toward a critical state without fine-tuning of parameters or dependence on the initial condition. The prevalence of this mechanism across both artificial and natural phenomena underscores its importance as a universal framework for understanding the behavior of complex systems.

Since then, many mathematical toy models have been introduced as candidates to capture and rigorously define SOC. One such model is *activated random walk* (ARW). ARW is an interacting particle system on $\mathbb{Z}^d$ whose particles are either active or sleeping. Active particles perform independent continuous-time simple random walks at rate 1 and attempt to fall asleep at rate $\lambda \in (0, \infty)$; an attempt succeeds only if the particle is alone at its site. Sleeping particles do not move and are instantaneously reactivated when an active particle arrives. Starting from finitely many particles, the dynamics stabilize almost surely in finite time.

ARW shares many common features with its spiritual predecessor, the abelian sandpile model, specifically its abelian property. From a modeling and technical perspective, ARW has two important advantages over the abelian sandpile model: first, the fact that individual particles perform random walks and become sleeping through independent mechanisms; and second, that ARW is expected to exhibit SOC in a robust way, while the abelian sandpile lacks some of the universality properties associated with SOC [6, 12].

There are several natural notions of criticality for ARW, including the fixed-energy, driven-dissipative, cycle, and point-source formulations. Hoffman, Johnson, and Junge proved that these four definitions agree in $d = 1$ [9]. We write $\rho_c \in (0, 1)$ for this common critical density. We next recall the point-source and driven-dissipative critical densities in the formulation used in [9].

Point-source model. Start ARW on $\mathbb{Z}$ from the initial condition $n\delta_0$, consisting of n active particles at the origin and no other particles, and let $\eta^{(n)} = \mathrm{Stab}(n\delta_0) \in \{0, \mathfrak{s}\}^{\mathbb{Z}}$ be the final stabilized configuration (we define the stabilization operator Stab in Section 4.2). Let $L_n = L_n(\eta^{(n)})$ be the number of sites in the smallest integer interval containing all sleeping particles in $\eta^{(n)}$, and define the point-source critical density

$$\rho_{\mathrm{PS}} \;:=\; \lim_{n\to\infty} \mathbb{E}\left[\frac{n}{L_n}\right],$$

whenever the limit exists.

Driven-dissipative model. Fix $m \in \mathbb{N}$ and consider ARW on the finite interval $[\![0, m]\!] := \{0, 1, \dots, m\}$ with sinks at -1 and $m + 1$. Consider the Markov chain on stable configurations in which, at each step, one adds a single active particle and then stabilizes until only sleeping particles remain in $[\![0, m]\!]$. The active particle may be added uniformly at random in $[\![0,m]\!]$ or at a fixed site; this choice does not affect the stationary distribution of the chain [13]. Let π_m^{DD} be the stationary distribution. If $\eta \sim \pi_m^{\mathrm{DD}}$, let

$$N_m(\eta) \ := \ \sum_{x=0}^{m} \mathbf{1}_{\{\eta(x)=\mathfrak{s}\}}$$

be the number of sleeping particles in $[\![0,m]\!]$, and define

$$\rho_{\mathrm{DD}} \ := \ \lim_{m\to\infty} \frac{1}{m+1}\, \mathbb{E}_{\pi_m^{\mathrm{DD}}}[N_m],$$

whenever the limit exists.

By [9], these limits exist and coincide, $\rho_{\mathrm{PS}} = \rho_{\mathrm{DD}} = \rho_c$. We use the point-source formulation in the statement of our main theorem, while the driven-dissipative viewpoint enters in Section 4.4.

So far, little progress has been made on a key aspect of ARW: the existence of a critical stationary distribution in the infinite-volume setting which exhibits the hallmarks of SOC. Depending on the choice of context, this critical measure arises in a few different possible forms, but the same critical object, which can be thought of as a probability measure on particle configurations, should be visible in each one. The limit measure is expected to exhibit the statistical features which constitute a definition of SOC, namely multi-scale power laws for correlation decays and avalanche sizes, and hyperuniformity, which is a reduction in the variance of the number of particles in a large region compared to independent randomness.

One way to define the critical measure, which is relevant for our main result, is presented by Levine and Silvestri in [15, Section 2.3]. There, it is conjectured to arise as the microscopic limit of ARW started from a large *point-source* initial condition. Namely, for each n let $\mathbb{P}_n$ denote the law (on the space of sleeping particle configurations $\{0,\mathfrak{s}\}^{\mathbb{Z}^d}$) of the final sleeping configuration of ARW on $\mathbb{Z}^d$ started from n active particles at the origin. It is conjectured in [15] that $\mathbb{P}_n$ converges locally, in the sense that for every finite set $V \subset \mathbb{Z}^d$ the restrictions $\mathbb{P}_n|_V$ converge as $n \to \infty$, and that the resulting limiting measure α is translation invariant and has a well-defined particle density $\rho_c(d, \lambda)$. Alternatively, the measures π_m^{DD} should converge locally to the same limit α as $m \to \infty$.

For the abelian sandpile model on $\mathbb{Z}^d$ with $d \geq 2$, an infinite-volume limit of the stationary measure exists [1]. Local probabilities under this limiting measure can sometimes be computed explicitly, especially in two dimensions where the burning bijection relates sandpiles to spanning trees. For example, exact height probabilities on $\mathbb{Z}^2$ (that is, the probability that the stable configuration has a given number of particles at the origin) are available in special cases [18]. In high dimensions, the single-site height distribution admits a Poisson-type asymptotic as $d \to \infty$ [10], and related computations have been carried out on self-similar graphs including the Sierpiński gasket [8]. These spanning tree techniques are not available for ARW, so different methods are needed.

Our main result identifies the single-site occupation probability of any subsequential local limit measure for point-source ARW. Concretely, for any site i not too far from the source, the probability that i contains a sleeping particle in the final configuration is approximately ρ_c.

Theorem 4.1. *Fix $\varepsilon > 0$. For any sequence of sites $i(n) \in \mathbb{Z}$ with*

$$|i(n)| \leq \frac{1}{2}\Big(1 - \varepsilon\Big)n,$$

we have

$$\mathbb{P}_n\Big(\eta(i(n)) = \mathfrak{s}\Big) = \rho_c + o(1) \quad \text{as } n \to \infty,$$

where $\eta \sim \mathbb{P}_n$ is the final sleeping configuration of point-source ARW started from $n\delta_0$.

As a consequence of our methods, specifically those in Section 4.3, we obtain that the limiting critical measure, if it exists, is shift invariant.

Proposition 4.2. *Let α be a probability measure on $\{0, \mathfrak{s}\}^{\mathbb{Z}}$ which is a subsequential limit of the point-source laws $\mathbb{P}_n$. Then α is shift invariant.*

To summarize, we find that for every fixed site $i \in \mathbb{Z}$ the one-site marginals $\mathbb{P}_n|_{\{i\}}$ converge and the limiting probability that i contains a sleeping particle is ρ_c. In particular, if the full microscopic limit measure α exists, it must be shift-invariant with average particle density ρ_c. These results confirm part of Conjectures 3–5 of [15].

4.1.1 Outline and Overview

For $i \in \mathbb{Z}$, let $S_i := \{\eta^{(n)}(i) = \mathfrak{s}\}$ be the event that site i contains a sleeping particle in the final stabilized configuration. Equivalently, if $\eta \sim \mathbb{P}_n$, then $S_i = \{\eta(i) = \mathfrak{s}\}$ and

$$\mathbb{P}_n\Big(\eta(i) = \mathfrak{s}\Big) = \mathbb{P}(S_i).$$

The idea of the proof is to show that (1) the function $i \mapsto \mathbb{P}_n(\eta(i) = \mathfrak{s})$ is relatively flat, and (2) that over any block of size $\omega(1)$ as $n \to \infty$, the particle density is close to ρ_c. We prove (1) in Section 4.3 using a coupling between two instances of internal diffusion limited aggregation (IDLA), which can be thought of as ARW with infinite sleep rate. (2) is proved in Section 4.4 by an application of the machinery of [9], in combination with a semi-artificial toppling sequence involving IDLA on the $\omega(1)$-sized block. The artificial topplings allow us to give an upper bound on the probability that the particle density in the block is atypical. These two uses of IDLA are distinct and different from those in the existing ARW literature. Finally, in Section 4.5 we combine (1) and (2) with elementary inequalities to obtain Theorem 4.1.

4.2 ARW Setup

We follow the same conventions as most recent work on ARW by using the so-called site-wise (Diaconis–Fulton) construction of the process. Namely, to every site we attach an infinite stack of instructions, and evolve the system by applying the instructions to the particles that arrive at that site. Thanks to the abelian property (Lemma 4.3), as long as the resulting configuration is stable (i.e. contains only empty sites and sleeping particles), the order in which sites are toppled is irrelevant. We give an abbreviated version of the setup, with particular focus on the 'flattening' toppling sequences that will be relevant for our proofs. We refer the reader to [19] for more details.

State space and configurations. We take $\mathbb{N} = \{0, 1, 2, \dots\}$. Let $\mathbb{N} \cup \{\mathfrak{s}\}$ be ordered by $0 < \mathfrak{s} < 1 < 2 < \cdots$. A *configuration* is $\eta \in \{0, \mathfrak{s}, 1, 2, \dots\}^{\mathbb{Z}}$, with $\eta(x)$ the state at $x \in \mathbb{Z}$: $\eta(x) = 0$ means no particle, $\eta(x) = \mathfrak{s}$ means one sleeping particle, and $\eta(x) = n \geq 1$ means n active particles. When an active particle arrives at a site containing a sleeping particle, it wakes it, so $\mathfrak{s} + 1 = 2$. We set $|\mathfrak{s}| = 1$ and write $|\eta(x)|$ for the number of particles at x. A site x is *stable* if $\eta(x) < 1$ and *unstable* otherwise.

Instruction stacks and their execution. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be the product space of instruction stacks, under which $(\mathsf{Instr}_x(k))$ are i.i.d. with the following distribution

$$\mathsf{Instr}_x(k) = \begin{cases} \texttt{Left} & \text{with prob. } \frac{1}{2(1+\lambda)}, \\ \texttt{Right} & \text{with prob. } \frac{1}{2(1+\lambda)}, \\ \texttt{Sleep} & \text{with prob. } \frac{\lambda}{1+\lambda}. \end{cases}$$

A `Left` (resp. `Right`) instruction removes 1 from $\eta(x)$ and adds 1 to $\eta(x-1)$ (resp. $\eta(x+1)$). A `Sleep` instruction at a site with exactly one active particle changes $1 \mapsto \mathfrak{s}$, while it has no effect if there are $k \geq 2$ active particles. We write $\mathbb{E}$ for expectation with respect to $\mathbb{P}$.

Topplings and odometer. One should think of each stack $(\mathsf{Instr}_x(k))_{k \in \mathbb{N}^+}$ as an infinite roll of train tickets; then particles arrive one at a time at the train station, and are given the next ticket in the roll to read and execute. We call each such execution a *toppling* at a given site. Given a finite sequence of sites $\kappa = (x_1, \dots, x_\ell)$, we may execute the *toppling sequence* which topples those sites in order. The *odometer* of a toppling sequence records how many times each site has been toppled:

$$U_\kappa : \mathbb{Z} \to \mathbb{N}, \qquad U_\kappa(x) = (\#\text{ of topplings at } x \text{ in } \kappa).$$

If some instructions have already been executed, and (η, U) is the current state and partially-executed toppling sequence odometer, the next state and odometer after toppling site x is

$$\Phi_x(\eta, U) = \Big(\mathsf{Instr}_x(U(x)+1)(\eta),\, U + \delta_x\Big),$$

where δ_x is 1 at x and 0 elsewhere. The move Φ_x is *legal* for (η, U) when x is unstable for η. For a toppling sequence $\kappa = (x_1, \ldots, x_\ell)$, write

$$(\Phi_\kappa(\eta), U_\kappa) \ := \ \Phi_{x_\ell} \circ \cdots \circ \Phi_{x_1}(\eta, 0),$$

so that $\Phi_\kappa(\eta)$ denotes the resulting configuration. We call a toppling sequence $\kappa = (x_1, \ldots, x_\ell)$ *legal* for an initial configuration η if for each $j = 1, \ldots, \ell$, the move Φ_{x_j} is legal for the state obtained after performing the first $j-1$ topplings of κ starting from $(\eta, 0)$. We call κ *stabilizing* for η if the resulting configuration $\Phi_\kappa(\eta)$ is stable (i.e., every site is stable). The abelian property shows that all legal stabilizing schemes are equivalent:

Lemma 4.3 (Abelian Property [19]). *Let κ and β be legal, stabilizing toppling sequences for a configuration η. Then $U_\kappa(x) = U_\beta(x)$ for every $x \in \mathbb{Z}$, and in particular,*

$$\Phi_\kappa(\eta) = \Phi_\beta(\eta).$$

Fix a realization of the instruction stacks $(\mathsf{Instr}_x(k))$. For a configuration η which admits a finite legal stabilizing toppling sequence, define the *stabilization* of η by

$$\mathrm{Stab}(\eta) \ := \ \Phi_\kappa(\eta),$$

where κ is any legal stabilizing toppling sequence for η. By Lemma 4.3, $\mathrm{Stab}(\eta)$ does not depend on the choice of κ. We similarly write

$$u(\eta) \ := \ U_\kappa$$

for the associated (well-defined) odometer function. For a finite set V, we write $\mathrm{Stab}_V(\eta)$ for the stabilization of the configuration η when legal topplings are allowed only at sites in V, and sites in V^{c} are treated as sinks (equivalently, particles that move outside V are removed (ignored)). We write $u_V(\eta, x)$ for the associated odometer, that is, the number of topplings performed at site $x \in V$ during the stabilization of η in V (and set $u_V(\eta, x) = 0$ for $x \notin V$). We reserve $\mathrm{Stab}(\cdot)$ (without a subscript) for the global stabilization map defined earlier; finite-volume stabilizations always carry the subscript V.

4.2.1 Point-source ARW

Fix a positive integer n, and consider the initial condition $n\delta_0$ consisting of n active particles at the origin. Let

$$\eta^{(n)} \ := \ \mathrm{Stab}(n\delta_0) \in \{0, \mathfrak{s}\}^{\mathbb{Z}}$$

be the resulting stabilized configuration. All randomness in this chapter comes from the i.i.d. instruction stacks, and we work on the stack space $(\Omega, \mathcal{F}, \mathbb{P})$ introduced earlier. For each site $i \in \mathbb{Z}$, define the event

$$S_i \ := \ \{\eta^{(n)}(i) = \mathfrak{s}\} \in \mathcal{F},$$

that site i contains a sleeping particle after stabilization. Its probability is $\mathbb{P}(S_i)$. When it is convenient to view $\eta^{(n)}$ as a random element of $\{0, \mathfrak{s}\}^{\mathbb{Z}}$, we denote by

$$\mathbb{P}_n \ := \ \mathbb{P} \circ (\eta^{(n)})^{-1}$$

the induced law on final sleeping configurations.

4.2.2 IDLA Toppling Sequence

We will use a specific type of toppling sequence, which is legal and stabilizing, in Sections 4.3 and 4.4. The toppling sequence works in two phases: a *flattening* phase, where particles perform IDLA steps, and then a second phase, where particles can be toppled in any legal order until the configuration is stable. For the sake of definiteness, the flattening phase evolves in stages by always toppling the leftmost site x where

$|\eta(x)| \geq 2$, so this phase ends when only sites with $|\eta(x)| \leq 1$ remain. Note that during this phase, any sleep instruction encountered is effectively voided, since the particle which executed the sleep is at a site with at least two particles. Thus, when started from the point-source initial condition $n\delta_0$ on $\mathbb{Z}$, the flattening phase is equivalent to performing IDLA (aka ARW with infinite sleep rate), except that it results in an interval of length n with each site containing a single active particle (whereas IDLA would leave a single sleeping particle at each site). In particular, by Lemma 4.3 for ARW with $\lambda = \infty$, all legal choices of flattening procedure yield the same intermediate configuration ω_{IDLA} obtained at the end of the flattening phase.

4.3 Coupling two one-dimensional IDLA clusters

We now prove a coupling result for two IDLAs. (Recall that IDLA is just ARW with $\lambda = \infty$, i.e. where particles perform independent simple random walks, and fall asleep instantly whenever they are alone.) Let $\mathcal{C}_n^{(s)}$ denote the IDLA cluster obtained by releasing n particles from the site $s \in \{0, 1\}$ and performing IDLA until all particles are asleep. We will construct an explicit coupling of the two clusters $\mathcal{C}_n^{(0)}$ and $\mathcal{C}_n^{(1)}$ so that they coincide with high probability as $n \to \infty$.

Let $K_n^{(s)}$ denote the number of occupied sites strictly to the right of s in $\mathcal{C}_n^{(s)}$ after stabilization. Note that the random variable $K_n^{(s)}$ completely determines $\mathcal{C}_n^{(s)}$ (for each $s \in \{0, 1\}$, respectively) because the IDLA cluster is always a connected interval in $\mathbb{Z}$. It follows that the two clusters $\mathcal{C}_n^{(0)}$ and $\mathcal{C}_n^{(1)}$ coincide if and only if

$$K_n^{(0)} = K_n^{(1)} + 1.$$

Thus, to show the desired coupling, it is equivalent to construct a coupling of the laws $K_n^{(0)}$ and $K_n^{(1)}$ for which $K_n^{(0)} = K_n^{(1)} + 1$ holds with high probability.

Let D_n denote the number of descents of a uniformly random permutation $\sigma \in S_n$, that is

$$D_n = \#\{1 \leq i \leq n-1 : \sigma(i) > \sigma(i+1)\}.$$

The following result of Mittelstaedt gives an exact description of the law of $K_n^{(s)}$ via the Eulerian distribution of descents D_n.

Theorem 4.4 (Theorem 1 of [16])**.** *For each source $s \in \{0,1\}$,*

$$K_n^{(s)} \overset{d}{=} D_n, \qquad P_n(k) := \frac{\langle n, k\rangle}{n!}, \quad 0 \le k \le n-1,$$

where $\langle n, k\rangle$ is the kth Eulerian number.

In particular, $\Pr(D_n = k) = P_n(k) = \langle n, k\rangle / n!$ for $0 \le k \le n-1$. We extend P_n to a probability mass function on $\mathbb{Z}$ by setting $P_n(k) = 0$ for $k \notin \{0, 1, \ldots, n-1\}$, and define its right shift

$$Q_n(k) := P_n(k-1), \qquad k \in \mathbb{Z}.$$

Recall that $\mathcal{C}_n^{(0)} = \mathcal{C}_n^{(1)}$ if and only if $K_n^{(0)} = K_n^{(1)} + 1$. By maximal coupling, the optimal success probability for coupling a random variable with law P_n to one with law Q_n is $1 - \mathrm{TV}(P_n, Q_n)$, where

$$\mathrm{TV}(P_n, Q_n) := \frac{1}{2} \sum_{k \in \mathbb{Z}} \Big| P_n(k) - Q_n(k) \Big|$$

is the total variation distance. Since $K_n^{(0)} \sim P_n$ and $K_n^{(1)} + 1 \sim Q_n$, there exists a coupling of $K_n^{(0)}$ and $K_n^{(1)}$ such that $K_n^{(0)} = K_n^{(1)} + 1$ with probability at least $1 - \mathrm{TV}(P_n, Q_n)$. Therefore it suffices to bound $\mathrm{TV}(P_n, Q_n)$. We use the following standard fact about the Eulerian distribution: the Eulerian numbers are log-concave in k [23], and hence the sequence $P_n(k)$ is unimodal in k. From this, a short telescoping argument gives an exact identity.

Lemma 4.5. *If $(P(k))_{k \in \mathbb{Z}}$ is unimodal and $Q(k) = P(k-1)$, then*

$$\mathrm{TV}(P, Q) = \max_k P(k).$$

Proof. Write $a_k = P(k) - Q(k)$. Since $\sum_k a_k = 0$ and the sign of a_k is nonnegative up to the mode and negative thereafter,

$$\mathrm{TV}(P, Q) = \frac{1}{2} \sum_k |a_k| = \sum_k (a_k)_+ = \max_t \sum_{k \le t} a_k = \max_t P(t),$$

because $\sum_{k\le t} a_k = P(t)$ by telescoping. □

Applying Lemma 4.5 to P_n and Q_n yields

$$\mathrm{TV}(P_n, Q_n) = \max_k P_n(k).$$

Thus the coupling problem reduces to bounding the maximal point mass of D_n.

Let $F_n(x) := \sum_{k\le\lfloor x\rfloor} P_n(k)$. Let Φ_n denote the CDF of $\mathcal{N}(\mu_n, \sigma_n^2)$ with

$$\mu_n = \frac{n-1}{2}, \quad \sigma_n^2 = \frac{n+1}{12}.$$

By Lemma 4.5, $\mathrm{TV}(P_n, Q_n) = \max_k P_n(k)$, so it remains to control the maximal atom of the Eulerian distribution. We obtain an $O(n^{-1/2})$ bound on $\max_k P_n(k)$ from a uniform normal approximation due to Özdemir.

Theorem 4.6 (Theorem 1.1 of [17])**.**

$$\|F_n - \Phi_n\|_\infty \;:=\; \sup_{x\in\mathbb{R}} |F_n(x) - \Phi_n(x)| \;\le\; C\, n^{-1/2}.$$

Lemma 4.7.

$$\max_k P_n(k) \le \left(\sqrt{\frac{6}{\pi}} + 2C\right) n^{-1/2}.$$

Proof. Since $P_n(k) = F_n(k) - F_n(k-1)$, we obtain

$$P_n(k) \le (\Phi_n(k) - \Phi_n(k-1)) + 2\|F_n - \Phi_n\|_\infty \le \max_{t\in\mathbb{R}} (\Phi_n(t) - \Phi_n(t-1)) + \frac{2C}{\sqrt{n}}.$$

And

$$\max_{t\in\mathbb{R}} (\Phi_n(t) - \Phi_n(t-1)) \le \frac{1}{\sqrt{2\pi}\sigma_n} = \sqrt{\frac{6}{\pi}} \cdot \frac{1}{\sqrt{n+1}} \le \sqrt{\frac{6}{\pi}} \cdot \frac{1}{\sqrt{n}}.$$

Therefore

$$\max_k P_n(k) \le \left(\sqrt{\frac{6}{\pi}} + 2C\right) n^{-1/2}.$$

□

Combining Lemmas 4.5 and 4.7, and following the IDLA toppling sequence as laid out in Section 4.2.2, shows that the probabilities $\mathbb{P}_n(\eta(i) = \mathfrak{s})$ and $\mathbb{P}_n(\eta(i+1) = \mathfrak{s})$ are close:

Lemma 4.8. *There exists $c > 0$ such that for all $n \in \mathbb{N}$ and all $i \in \mathbb{Z}$,*

$$\Big|\mathbb{P}_n(\eta(i) = \mathfrak{s}) - \mathbb{P}_n(\eta(i+1) = \mathfrak{s})\Big| \le c\, n^{-1/2}.$$

Proof. Applying Lemmas 4.5 and 4.7 gives

$$\mathrm{TV}(P_n, Q_n) = \max_k P_n(k) \le cn^{-1/2}$$

for a global constant $c > 0$. By maximal coupling, there exists a coupling of $\mathcal{C}_n^{(0)}$ and $\mathcal{C}_n^{(1)}$ such that $\mathcal{C}_n^{(0)} = \mathcal{C}_n^{(1)}$ with probability at least $1 - \mathrm{TV}(P_n, Q_n)$. This shows that we can couple $\mathcal{C}_n^{(s)}$, $s \in \{0, 1\}$, so that they are equal with probability at least $1 - cn^{-1/2}$.

Now consider two instances of point-source ARW on $\mathbb{Z}$ with n particles, one having source 0 and the other source 1. We couple these two instances by applying the flattening procedure described in Section 4.2.2 simultaneously to both instances. Specifically, for the IDLA phase, we use the coupling guaranteed by the previous paragraph. If the resulting IDLA clusters $\mathcal{C}_n^{(s)}$, $s \in \{0, 1\}$ are the same, then we couple the remaining (unused) portions of the instruction stacks to be identical and complete stabilization using these coupled stacks; otherwise, stabilize the two instances independently. In the former case, the final configurations are identical by construction. The result follows. $\square$

As a corollary, we obtain that any subsequential local limit of the measures $\mathbb{P}_n$ must be shift invariant.

Proof of Proposition 4.2. Let $(\theta\eta)(x) = \eta(x-1)$ denote the shift operator on configurations, and for a measure μ write $\theta\mu$ for its pushforward under θ. Let $\mathbb{P}_n^{(1)}$ be the point-source law on $\{0, \mathfrak{s}\}^{\mathbb{Z}}$ for n particles started at 1; by translation invariance of the stacks, $\mathbb{P}_n^{(1)} = \theta\,\mathbb{P}_n$. The coupling constructed in the proof of Lemma 4.8 gives $\mathrm{TV}(\mathbb{P}_n, \mathbb{P}_n^{(1)}) \le cn^{-1/2}$, hence $\mathrm{TV}(\mathbb{P}_n, \theta\,\mathbb{P}_n) \le cn^{-1/2}$. Let $n_k \to \infty$ be a subsequence along which $\mathbb{P}_n \to \alpha$. For any cylinder event E,

$$|\mathbb{P}_n(E) - \mathbb{P}_n(\theta^{-1}E)| \le \mathrm{TV}(\mathbb{P}_n, \theta\,\mathbb{P}_n) \le cn^{-1/2},$$

so evaluating at $n = n_k$ and letting $k \to \infty$ yields $\alpha(E) = \alpha(\theta^{-1}E)$. Thus α is shift invariant. $\square$

4.4 Block averages near the source

In this section we show that the average particle density over any block of growing size in the bulk is asymptotically equal to the critical density ρ_c. Fix $\varepsilon, \gamma \in (0,1)$. For each n, we choose a site $i = i(n) \in \mathbb{Z}$ satisfying

$$|i(n)| \leq \frac{1}{2}(1-\varepsilon)\, n. \tag{4.1}$$

We then consider the block

$$I_n = [\![i - \lfloor n^\gamma \rfloor,\, i]\!].$$

The bulk condition (4.1) is chosen so that, with high probability, the one–dimensional IDLA cluster produced by n particles at the origin contains the entire block I_n and leaves exactly one particle at each site of I_n (see Lemma 4.10 below).

We evolve the point-source system using the toppling procedure described in Section 4.2.2. We now briefly recall the procedure and add some auxiliary notation.

Phase 1 (IDLA flattening). We follow the *flattening* toppling procedure described in Section 4.2.2. This produces a configuration ω_{IDLA}.

Phase 2 (ARW stabilization). Starting from ω_{IDLA} we resume the full ARW dynamics on $\mathbb{Z}$, by performing any legal, stabilizing toppling sequence.

During Phase 2 we define the boundary fluxes

$$a^* := \#\{\texttt{Right} \text{ instructions used at the site } i - \lfloor n^\gamma \rfloor - 1 \text{ during Phase 2}\},$$

$$b^* := \#\{\texttt{Left} \text{ instructions used at the site } i + 1 \text{ during Phase 2}\}.$$

In other words, a^* and b^* are the total numbers of jumps entering I_n from the left and from the right, respectively, during Phase 2. Let D be the number of particles in I_n in the final stabilized configuration of the full ARW on $\mathbb{Z}$.

We will use the following auxiliary driven-dissipative process which is coupled to the evolution in phase 2. For any $a, b \geq 0$, we consider the following finite-volume system associated with I_n: run ARW on the interval I_n with sinks at $i - \lfloor n^\gamma \rfloor - 1$ and $i + 1$, started from the restriction of ω_{IDLA} to I_n with a additional active particles placed at $i - \lfloor n^\gamma \rfloor$ and b additional active particles placed at i. We use the same instruction stacks at sites in I_n as in the full system. Let $D_{a,b}$ denote the number of particles in I_n when this finite-volume system stabilizes. Our first goal is to relate D to D_{a^*,b^*}, so that the block density can be analyzed via a driven-dissipative system.

Lemma 4.9. *Fix n and consider the two–phase construction above. Then, for every realization of the instruction stacks, we have*

$$D = D_{a^*,b^*}.$$

Proof. Write $I_n = [L, R]$ where $L = i - \lfloor n^\gamma \rfloor$ and $R = i$. We now work with the instruction stacks as they stand at the beginning of Phase 2, after Phase 1 has already consumed some instructions.

Throughout this proof, for any configuration η on I_n, $\mathrm{Stab}_{I_n}(\eta)$ denotes the stabilized configuration obtained by starting from η and applying any legal, stabilizing toppling sequence that topples only sites in I_n (equivalently, $L - 1$ and $R + 1$ act as sinks). By the abelian property (Lemma 4.3) restricted to I_n, $\mathrm{Stab}_{I_n}(\eta)$ is well defined and independent of the choice of legal stabilizing sequence in I_n. For $x \in I_n$, define the addition operator

$$A_x(\eta) := \mathrm{Stab}_{I_n}(\eta + \delta_x).$$

It follows from Lemma 4.3 that for all $x, y \in I_n$ and all configurations η on I_n,

$$A_x A_y = A_y A_x \quad \text{and} \quad A_x(\mathrm{Stab}_{I_n}(\eta)) = A_x(\eta), \tag{4.2}$$

and more generally that for any $a, b \geq 0$,

$$A_L^a A_R^b(\eta) = \mathrm{Stab}_{I_n}\Big(\eta + a\delta_L + b\delta_R\Big). \tag{4.3}$$

Now consider Phase 2 of the full system on $\mathbb{Z}$, started from ω_{IDLA}. We follow the following I_n-first toppling procedure: at each step

- if there is an unstable site in I_n, topple one such site;
- otherwise, topple an unstable site in $\mathbb{Z} \setminus I_n$.

This defines a legal stabilizing toppling sequence on $\mathbb{Z}$, so by Lemma 4.3 it produces the same final configuration and odometer as any other legal stabilizing sequence.

Let $(\eta_t)_{t\geq 0}$ be the sequence of configurations during this I_n-first stabilization in Phase 2. Let

$$\tau_1 < \tau_2 < \cdots < \tau_M$$

be the times t at which a particle jumps into I_n from outside. By definition of a^* and b^*,

$$M = a^* + b^*.$$

Define configurations ζ_j on I_n as follows. Let ζ_0 be the restriction to I_n of the first configuration η_t for which every site in I_n is stable. Hence,

$$\zeta_0 = \mathrm{Stab}_{I_n}(\omega_{\mathrm{IDLA}}|_{I_n}).$$

For $j \geq 1$, at time τ_j a particle enters I_n at some boundary site $s_j \in \{L, R\}$. Between the last time at which I_n was stable (whose restriction is ζ_{j-1} by definition) and the instant τ_j, we topple only outside I_n, so the configuration on I_n just before τ_j is still ζ_{j-1}. And immediately after the jump at time τ_j the configuration on I_n is $\zeta_{j-1} + \delta_{s_j}$. After time τ_j we again topple only sites in I_n until every site in I_n is stable; let ζ_j be the resulting configuration on I_n. Then,

$$\zeta_j = \mathrm{Stab}_{I_n}(\zeta_{j-1} + \delta_{s_j}) = A_{s_j}(\zeta_{j-1}), \qquad j \geq 1.$$

Iterating, we obtain

$$\zeta_M = A_{s_M} \cdots A_{s_1}(\zeta_0). \tag{4.4}$$

Among the entry sites $s_1, \ldots, s_M$, the site L appears exactly a^* times and the site R appears exactly b^* times, so by the commutativity in (4.2),

$$A_{s_M} \cdots A_{s_1} = A_L^{a^*} A_R^{b^*}.$$

Combining this with (4.4), the identity $\zeta_0 = \mathrm{Stab}_{I_n}(\omega_{\mathrm{IDLA}}|_{I_n})$, and the relation $A_x(\mathrm{Stab}_{I_n}(\eta)) = A_x(\eta)$ from (4.2), we obtain

$$\zeta_M = A_{s_M} \cdots A_{s_1}(\zeta_0) = A_L^{a^*} A_R^{b^*}\Big(\zeta_0\Big) = A_L^{a^*} A_R^{b^*}\Big(\mathrm{Stab}_{I_n}(\omega_{\mathrm{IDLA}}|_{I_n})\Big) = A_L^{a^*} A_R^{b^*}\Big(\omega_{\mathrm{IDLA}}|_{I_n}\Big).$$

By construction of the I_n–first procedure, ζ_M is exactly the restriction to I_n of the final configuration of the full ARW on $\mathbb{Z}$ at the end of Phase 2, since after time τ_M there are no further jumps into I_n and we never topple inside I_n again once it is stable. Therefore D, the number of particles in I_n in the final stabilized configuration of the full system, is exactly the number of particles in ζ_M.

On the other hand, consider the finite-volume ARW on I_n with sinks at $L-1$ and $R+1$, started from the configuration $\omega_{\mathrm{IDLA}}|_{I_n} + a^*\delta_L + b^*\delta_R$, using the same instruction stacks at sites in I_n as at the beginning of Phase 2. By definition of Stab_{I_n} and by the identity (4.3), the final stabilized configuration of this finite system is

$$\mathrm{Stab}_{I_n}\Big(\omega_{\mathrm{IDLA}}|_{I_n} + a^*\delta_L + b^*\delta_R\Big) = A_L^{a^*} A_R^{b^*}\Big(\omega_{\mathrm{IDLA}}|_{I_n}\Big) = \zeta_M.$$

Thus the number of particles in I_n in this finite-volume system, which is D_{a^*,b^*} by definition, coincides with D. $\square$

We now show that the average particle density in I_n converges to ρ_c.

Lemma 4.10. *For any $\varepsilon, \gamma \in (0,1)$ and integer i satisfying (4.1), we have*

$$\sum_{j=i-\lfloor n^\gamma \rfloor}^{i} \mathbb{P}_n\Big(\eta(j) = \mathfrak{s}\Big) = (\rho_c + o(1))\, n^\gamma \qquad \text{as } n \to \infty. \tag{4.5}$$

Proof. In the final configuration each site contains at most one sleeping particle. So by linearity of expectation

$$\mathbb{E}[D] = \sum_{j \in I_n} \mathbb{P}_n\Big(\eta(j) = \mathfrak{s}\Big).$$

Therefore it suffices to show $\mathbb{E}[D] = (\rho_c + o(1))\, n^\gamma$. We first perform Phase 1 as outlined at the beginning of this section. Let A_1 be the event that, after this phase, there is not exactly one active particle at each site of I_n. By Theorem 1 of [16] we can deduce that there exist constants $c_1, C_1 > 0$ such that

$$\mathbb{P}(A_1) \le C_1 e^{-c_1 n^{.5}}. \tag{4.6}$$

On A_1^c the IDLA configuration ω_{IDLA} has exactly one active particle at each $j \in I_n$.

Next, we run Phase 2. By Lemma 4.9 we have $D = D_{a^*,b^*}$. For each fixed $a, b \ge 0$, consider the finite-volume ARW on I_n with sinks at $i - \lfloor n^\gamma \rfloor - 1$ and $i + 1$. In this system we start from the restriction of ω_{IDLA} to I_n (where each site $j \in I_n$ contains one active particle), then add a additional active particles at the left boundary site $i - \lfloor n^\gamma \rfloor$ and b additional active particles at the right boundary site i, and then stabilize. Theorem 2.1 of [13] implies that the law of the resulting stabilized configuration is the stationary distribution of the driven-dissipative ARW on I_n with these sinks and boundary driving. Fix $\delta > 0$, and for integers a and b with $a, b \ge 0$, define the deviation event

$$A_{2,a,b} := \{D_{a,b} \notin (n^\gamma \rho_c\, (1 - \delta/2)\, ,\, n^\gamma \rho_c\, (1 + \delta/2))\}\, .$$

By Proposition 8.5 and 8.6 of [9], we know there exist constants $c_2, C_2 > 0$ such that for all $a, b \ge 0$ and all sufficiently large n,

$$\mathbb{P}(A_{2,a,b}) \le C_2 e^{-c_2 n^\gamma}.$$

Next we control the size of the boundary fluxes. By the same argument as Lemma 3.5 of [4] there exist constants $c_3, C_3 > 0$ such that

$$\mathbb{P}(a^* > n^5 \text{ or } b^* > n^5) \le C_3 e^{-c_3 n}.$$

We now define a good event. Let A be the event that

1. A_1^c occurs,

2. $\left(\bigcup_{a,b\in[\![0,n^5]\!]} A_{2,a,b}\right)^c$ occurs, and

3. a^* and b^* are between 0 and n^5.

By a union bound with the fact that there are at most n^{10} pairs (a,b), we know that there exist constants $c, C, \beta > 0$ such that

$$\mathbb{P}(A^c) \leq Cn^{10}e^{-cn^\beta}.$$

Also note that $0 \leq D \leq n$ and on A

$$D = D_{a^*,b^*} \in (n^\gamma \rho_c\,(1-\delta/2)\,,\, n^\gamma \rho_c\,(1+\delta/2))\,.$$

Now we are ready to bound $\mathbb{E}[D]$. The lower bound is

$$\mathbb{E}[D] \geq \mathbb{E}[D\mathbb{1}_A] \geq \mathbb{P}(A)\, n^\gamma \rho_c\,(1-\delta/2)\,.$$

The upper bound is

$$\mathbb{E}[D] \leq \mathbb{P}(A)\, n^\gamma \rho_c\,(1+\delta/2) + \mathbb{P}(A^c)\, n \leq n^\gamma \rho_c\,(1+\delta)\,.$$

Combining these two equations above we get that

$$\mathbb{E}[D] = n^\gamma \rho_c\,(1+o(1))$$

as desired. □

4.5 Completing the Proof

Proof of Theorem 4.1. By Lemma 4.8, for any $j \in \mathbb{Z}$ we have

$$\Big|\mathbb{P}_n(\eta(j) = \mathfrak{s}) - \mathbb{P}_n(\eta(j+1) = \mathfrak{s})\Big| \leq c\,n^{-1/2}.$$

By the triangle inequality this implies that for all $i, j \in \mathbb{Z}$,

$$\Big|\mathbb{P}_n(\eta(i) = \mathfrak{s}) - \mathbb{P}_n(\eta(i-j) = \mathfrak{s})\Big| \leq c\,j\,n^{-1/2}. \tag{4.7}$$

Summing over $j \in [0, \lfloor n^\gamma \rfloor] \cap \mathbb{Z}$ and using (4.7), we obtain

$$\left| \sum_{j=i-\lfloor n^\gamma \rfloor}^{i} \mathbb{P}_n(\eta(j) = \mathfrak{s}) - \lfloor n^\gamma \rfloor \mathbb{P}_n(\eta(i) = \mathfrak{s}) \right| \le \sum_{j=0}^{\lfloor n^\gamma \rfloor} c\, j\, n^{-1/2} \le c'\, n^{2\gamma - 1/2}.$$

Now take $i = i(n)$ as in the statement of the theorem. By Lemma 4.10 with the same ε and γ, we have

$$\sum_{j=i-\lfloor n^\gamma \rfloor}^{i} \mathbb{P}_n(\eta(j) = \mathfrak{s}) = (\rho_c + o(1))\, n^\gamma.$$

Combining this with the previous inequality and dividing by n^γ gives

$$\left| \mathbb{P}_n(\eta(i) = \mathfrak{s}) - \rho_c \right| \le c\, n^{\gamma - 1/2} + o(1),$$

which tends to 0 as $n \to \infty$ by taking any $0 < \gamma < 1/2$. This proves the claim. $\square$

Remark. Note that the source of the $o(1)$ error in Theorem 4.1 is the proof of Lemma 4.10, and ultimately Lemmas 8.5 and 8.6 of [9], where δ depends only implicitly on n. Thus, our method does not give a quantitative estimate for the error term.

4.6 Future Directions

Our results verify the one–site part of Conjecture 3 of [15] in the one–dimensional point–source model, conditional on the existence of the full microscopic limit: if the point–source measures converge locally, then the limiting measure is shift invariant and has density ρ_c. A natural next step would be to extend our methods to multi–site events, for example to show that the joint law of the sleeping indicator field $(\mathbf{1}_{S_i})_{i \in \mathbb{Z}}$ converges on finite windows, and to relate its covariance structure to the hyperuniformity conjecture in [15].

Two further directions concern changing either the geometry or the driving mechanism. First, it is very plausible that an analogue of Theorem 4.1 should hold in two and higher dimensions for the point–source model on $\mathbb{Z}^d$. Our proof suggests that this would follow from an IDLA statement about the harmonic measure on the boundary of a large ball: namely, that the hitting distribution is essentially unchanged when the source is moved from the origin to a neighbor of the origin. Second, one expects a local density theorem of the same form to hold for the driven–dissipative ARW on large finite intervals. Our attempted coupling breaks down in this setting because particles can fall into the sinks at different times in the two coupled systems, so new ideas seem necessary. Establishing the analogues of Theorem 4.1 and Corollary 4.2 in these settings would give further evidence for universality of the limit measure.

Chapter 5
CONCLUDING REMARKS AND FUTURE DIRECTIONS

The two main results of this dissertation examine activated random walk from complementary sides of the universality problem. Chapter 3 studies how little initial activity is needed to see the supercritical phase. Chapter 4 studies what local density is seen deep inside a large point-source aggregate. In both cases, the answer is governed by the same critical density that appears in the one-dimensional density conjecture.

The first result shows that if the Bernoulli sleeping-particle density is above critical, then a single active particle at the origin can trigger indefinite activity with positive probability. This confirms that the supercritical phase is robust in a strong sense: it does not require the whole initial configuration to be active. The second result shows that the point-source aggregate has the correct one-site density throughout a macroscopic bulk window. It also implies shift invariance for subsequential microscopic limits of the point-source laws.

These results leave several natural problems open.

5.1 *Multi-site local limits*

The local-density theorem identifies one-site marginals. A natural next step is to prove convergence of finite-dimensional marginals for the point-source laws. This would address the full microscopic limit conjecture for the one-dimensional point-source aggregate. Such a result would need estimates on correlations between sleeping indicators at multiple sites, not only control of their average density.

5.2 Hyperuniformity and correlations

Levine and Silvestri's universality program predicts strong negative correlations and hyperuniformity in critical activated random walk states [15]. The point-source local-density theorem is compatible with this picture but does not estimate variances of particle counts. Proving sub-volume-order variance bounds for large windows remains a central challenge.

5.3 Higher-dimensional analogues

The methods in Chapter 4 use one-dimensional IDLA structure in an essential way. A higher-dimensional analogue would require a replacement for the one-dimensional coupling of IDLA clusters started from neighboring sources. One possible direction is to control how harmonic measure on a large IDLA boundary changes when the source is moved by one lattice step.

5.4 Other routes to the critical state

Another direction is to compare local limits obtained from the point-source, driven-dissipative, fixed-energy, and wake-chain formulations. The equality of critical densities in one dimension is now known, but equality of the limiting local measures remains largely open. Establishing such equivalence would turn density-level universality into full distributional universality.

BIBLIOGRAPHY

[1] Siva R. Athreya and Antal A. Járai. Infinite volume limit for the stationary distribution of abelian sandpile models. *Communications in Mathematical Physics*, 249(1):197–213, 2004. doi: 10.1007/s00220-004-1080-0.

[2] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59(4):381–384, 1987. doi: 10.1103/PhysRevLett.59.381.

[3] Ofer Biham, Erel Milshtein, and Ofer Malcai. Evidence for universality within the classes of deterministic and stochastic sandpile models. *Physical Review E*, 63(6):061309, 2001. doi: 10.1103/PhysRevE.63.061309.

[4] Madeline Brown, Christopher Hoffman, and Hyojeong Son. Activated random walks on $\mathbb{Z}$ with critical particle density. arXiv:2411.07609 [math.PR], 2024. doi: 10.48550/arXiv.2411.07609.

[5] Deepak Dhar. Some results and a conjecture for Manna's stochastic sandpile model. *Physica A: Statistical Mechanics and its Applications*, 270(1–2):69–81, 1999. doi: 10.1016/S0378-4371(99)00149-1.

[6] Anne Fey, Lionel Levine, and David B. Wilson. Driving sandpiles to criticality and beyond. *Physical Review Letters*, 104:145703, 2010. doi: 10.1103/PhysRevLett.104.145703.

[7] Nicolas Forien. Macroscopic flow out of a segment for activated random walks in dimension 1. *ALEA, Latin American Journal of Probability and Mathematical Statistics*, 22:557–577, 2025. doi: 10.30757/ALEA.v22-21.

[8] Nico Heizmann, Robin Kaiser, and Ecaterina Sava-Huss. Average height for abelian sandpiles and the looping constant on sierpiński graphs. *Computational and Applied Mathematics*, 44:227, 2025. doi: 10.1007/s40314-025-03139-5.

[9] Christopher Hoffman, Tobias Johnson, and Matthew Junge. The density conjecture for activated random walk. arXiv:2406.01731 [math.PR], 2024. doi: 10.48550/arXiv.2406.01731.

[10] Antal A. Járai and Minwei Sun. Asymptotic height distribution in high-dimensional sandpiles. *Journal of Theoretical Probability*, 34:349–362, 2021. doi: 10.1007/s10959-019-00962-5.

[11] Hang-Hyun Jo and Hyeong-Chai Jeong. Comment on "driving sandpiles to criticality and beyond". *Physical Review Letters*, 105(1):019601, 2010. doi: 10.1103/PhysRevLett.105.019601.

[12] Lionel Levine. Threshold state and a conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle. *Communications in Mathematical Physics*, 335(2):1003–1017, 2015. doi: 10.1007/s00220-014-2216-5.

[13] Lionel Levine and Feng Liang. Exact sampling and fast mixing of activated random walk. *Electronic Journal of Probability*, 29:Paper No. 184, 20 pp., 2024. doi: 10.1214/24-EJP1220.

[14] Lionel Levine and Yuval Peres. Laplacian growth, sandpiles, and scaling limits. *Bulletin of the American Mathematical Society*, 54(3):355–382, 2017. doi: 10.1090/bull/1573.

[15] Lionel Levine and Vittoria Silvestri. Universality conjectures for activated random walk. *Probability Surveys*, 21:1–27, 2024. doi: 10.1214/24-PS25.

[16] Kiana Mittelstaedt. A stochastic approach to Eulerian numbers. *The American Mathematical Monthly*, 127(7):618–628, 2020. doi: 10.1080/00029890.2020.1757359.

[17] Alperen Y. Özdemir. Martingales and descent statistics. *Advances in Applied Mathematics*, 140:102395, 2022. doi: 10.1016/j.aam.2022.102395.

[18] V. B. Priezzhev. Structure of two-dimensional sandpile. I. height probabilities. *Journal of Statistical Physics*, 74:955–979, 1994. doi: 10.1007/BF02188212.

[19] Leonardo T. Rolla. Activated random walks on $\mathbb{Z}^d$. *Probability Surveys*, 17: 478–544, 2020. doi: 10.1214/19-PS339.

[20] Leonardo T. Rolla, Vladas Sidoravicius, and Olivier Zindy. Universality and sharpness in activated random walks. *Annales Henri Poincaré*, 20(6):1823–1835, 2019. doi: 10.1007/s00023-019-00797-0.

[21] Robert Hough and Hyojeong Son. Cut-off for sandpiles on tiling graphs. *The Annals of Probability*, 49(2):671–731, 2021. doi: 10.1214/20-AOP1458.

[22] Robert Hough and Hyojeong Son. The spectrum of the abelian sandpile model. *Mathematics of Computation*, 90(327):441–469, 2021. doi: 10.1090/mcom/3565.

[23] Richard P. Stanley. Log-concave and unimodal sequences in algebra, combinatorics, and geometry. *Annals of the New York Academy of Sciences*, 576:500–535, 1989. doi: 10.1111/j.1749-6632.1989.tb16434.x.