EXACT SAMPLING AND FAST MIXING OF ACTIVATED RANDOM WALK

LIONEL LEVINE AND FENG LIANG

Abstract. Activated Random Walk (ARW) is an interacting particle system on the d-dimensional lattice $\mathbb{Z}^d$. On a finite subset $V \subset \mathbb{Z}^d$ it defines a Markov chain on $\{0,1\}^V$. We prove that when V is a Euclidean ball intersected with $\mathbb{Z}^d$, the mixing time of the ARW Markov chain is at most $1+o(1)$ times the volume of the ball. The proof uses an exact sampling algorithm for the stationary distribution, a coupling with internal DLA, and an upper bound on the time when internal DLA fills the entire ball. We conjecture cutoff at time ζ times the volume of the ball, where $\zeta < 1$ is the limiting density of the stationary state.

Dedicated to the memory of Vladas Sidoravicius.

Contents

Date: October 24, 2021.

2010 *Mathematics Subject Classification.* 37A25 60J10 82C22 82C24 82C26 .

Key words and phrases. Abelian property, activated random walk, coupling, cutoff, exact sampling, internal DLA, mixing time, self-organized criticality, stochastic abelian network, strong stationary time.

LL was partially supported by a Simons Fellowship, IAS Von Neumann Felowship, and NSF grant DMS-1455272.

1. Introduction: Activated Random Walk

A key feature of many complex systems is the release of stress sudden bursts. An example is the pressure between continental plates, released in earthquakes. Bak, Tang and Wiesenfeld called these "self-organized critical" (SOC) systems, and proposed a mathematial model of them, the abelian sandpile [5, 7]. But the abelian sandpile is non-universal: Even in the limit of large system size, its behavior depends delicately on the underlying graph [20] and on the initial condition [8, 15].

One of the best candidates for a *universal* model of SOC is Activated Random Walk (ARW) [25]. This is an interacting particle system with two species, **active** and **sleeping**. Active particles perform random walks and fall asleep at a fixed rate λ. Sleeping particles do not move, but become active when an active particle encounters them. To make explicit the connection to SOC, sleeping particles represent stress in the system, and a single active particle can cause a burst of activity by waking up many sleeping particles.

So far, one universality result has been proved for ARW: Rolla, Sidoravicius, and Zindy [26] show that there is a critical mean $\zeta_c = \zeta_c(\lambda, d)$ such that for any translation-invariant and ergodic configuration s of active particles in $\mathbb{Z}^d$ with mean ζ

$$\mathrm{P}(s \text{ stabilizes}) = \begin{cases} 1, & \zeta < \zeta_c \\ 0, & \zeta > \zeta_c. \end{cases} \tag{1}$$

Still missing is a rigorous connection between these infinite ARW systems and their finite counterparts. For instance, we expect that the ARW stationary distribution μ_V on a finite subset $V \subset \mathbb{Z}^d$ has an infinite-volume limit μ, and that its mean equals ζ_c. We also expect the microstructure of finite ARW clusters (such as the cluster of sleeping particles formed by stabilizing n chips at the origin in $\mathbb{Z}^d$) to converge to μ as $n \to \infty$. These conjectures are detailed in the forthcoming paper [23].

Recent work has succeeded in showing that ζ_c is strictly between 0 and 1 on many transitive graphs [28], including $\mathbb{Z}^d$ for $d \geq 3$ [29] and $d = 1$ [10]; that ζ_c is continuous and strictly increasing in the sleep rate [30], and tends to zero as the square root of the sleep rate [3]; and that ARW at sufficiently high density takes exponential time to stabilize on a cycle [4]. Despite all this progress, very little is known about the behavior of Activated Random Walk at criticality.

In this paper we examine ARW from a different perspective, by driving a finite ARW system to a stationary state. We give an exact sampling algorithm for the stationary state, and upper bound its mixing time (the time it takes to reach the stationary state from an arbitrary intitial state). **Fast mixing** is evidence for universality in that the system forgets its initial state quickly.

To see how mixing relates to universality, we can contrast ARW to the *non-universal* abelian sandpile model. In contrast to (1), the abelian sandpile has an interval of critical means [9], and the problem of whether a sandpile on $\mathbb{Z}^d$ stabilizes almost surely is not even known to be decidable [19]. The root cause of this non-universality is slow mixing: For example, the sandpile mixing time on both the ball $B(0, n) \cap \mathbb{Z}^d$ and on the torus $\mathbb{Z}_n^d$ is of order $n^d \log n$ [11, 12]. This extra log factor is responsible for the non-universality of the sandpile threshold state [18]. Our upper bounds (Corollaries 5 and 6), show that the extra log factor is *not* present for ARW, providing further evidence for universality.

1.1. **The ARW process.** Let P be the transition matrix of a discrete time Markov chain on a finite state space $V \cup \{z\}$. Here z is an absorbing state ($P(z,z) = 1$) called the "sink". Assume that every $v_0 \in V$ can access the sink, in that there exists a path $v_0, v_1, \dots, v_k = z$ such that $P(v_{i-1}, v_i) > 0$ for all $i = 1, \dots, k$.

Given P and a vector $\lambda = (\lambda_v)_{v \in V}$ with each $\lambda_v \in [0, \infty]$, we define the **ARW process** for integer $t \geq 0$ by

$$\sigma_t = \mathtt{S}[\sigma_{t-1} + \delta_{u_t}]. \tag{2}$$

In words, the state σ_t at each discrete time step is obtained from the previous state σ_{t-1} by adding a single active particle at u_t and then stabilizing. Here

- σ_t takes values in the hypercube $\{0, \mathtt{s}\}^V$. The symbol $\mathtt{s}$ stands for "sleeping" and will be explained below.
- $u_1, u_2 \dots \in V$ is a (possibly random) sequence of vertices.
- $\delta_v(w) = 1$ if $w = v$ and 0 otherwise;
- $\mathtt{S}$ is the **stabilization operator for activated random walk** with sleep rate λ and base chain P, which we now define.

Following [24], consider the total ordering on $\mathbb{N} \cup \{\mathtt{s}\}$

$$0 < \mathtt{s} < 1 < 2 < \dots.$$

Extend addition to a commutative operation on $\mathbb{N} \cup \{\mathtt{s}\}$ by declaring

$$0 + \mathtt{s} = \mathtt{s}$$

and

$$n + \mathtt{s} = n + 1$$

for all $n \neq 0$. In particular, $s + s = s + 1 = 1 + s = 2$. Note that if σ takes values in $\{0, \mathtt{s}\}$, then $\sigma + \delta_v$ takes values in $\{0, \mathtt{s}, 1, 2\}$.

An **ARW configuration** is a map

$$\sigma : V \to \mathbb{N} \cup \{\mathtt{s}\}.$$

If $\sigma(v) = n \geq 1$ then we say there are n active particles at v; if $\sigma(v) = \mathtt{s}$ then we say there is one sleeping particle at v; and if $\sigma(v) = 0$ then we say there are no particles at v.

A configuration taking values in $\{0, \mathtt{s}\}$ is called a **sleeping configuration**. The stabilization operator $\mathtt{S}$ takes an arbitrary configuration σ as input, and outputs a sleeping configuration. If σ takes values in $\{0, \mathtt{s}\}$, then we define $\mathtt{S}[\sigma] = \sigma$. Otherwise, we obtain $\mathtt{S}[\sigma]$ by a sequence of **firings** $\mathtt{F}_v$ of vertices with at least one active particle. Firing vertex v is defined in two cases, depending whether there are at least two active particles ($\sigma(v) \geq 2$) or only one ($\sigma(v) = 1$).

- Suppose $\sigma(v) \geq 2$. To fire v, move one particle from v to a random vertex drawn from $P(v, \cdot)$. Formally,

$$\mathtt{F}_v[\sigma] = \sigma - \delta_v + \delta_w \text{ with probability } P(v, w) \text{ for each } w \in V \cup \{z\}.$$

- Suppose $\sigma(v) = 1$. To fire v, put the particle at v to sleep with probability

$$q_v := \frac{\lambda_v}{1+\lambda_v}.$$

Otherwise move the particle from v to a random neighbor drawn from $P(v, \cdot)$. Formally,

$$\mathtt{F}_v[\sigma] = \begin{cases} \sigma - \delta_v + \mathtt{s}\delta_v & \text{with probability } q_v \\ \sigma - \delta_v + \delta_w & \text{with probability } (1-q_v)P(v,w) \end{cases}$$

for each $w \in V$.

The domain of σ is V, not $V \cup \{z\}$; so in the case $w = z$ the term δ_z is zero. This case represents a particle falling into the sink, where it is removed from the system.

We make three remarks:

1. Any ARW configuration σ reaches a sleeping configuration after some (random, but almost surely finite) number of firings of active vertices.

   This follows from our assumption in the very beginning, that V is finite and every vertex can access the sink: If any particle is still active, then try to move it along a path to the sink. If it falls asleep or strays from the chosen path, then pick another active particle and try again. There is a positive number ϵ such that each such trial has probability at least ϵ of depositing a particle in the sink. Since the number of particles is finite and cannot increase, we reach a sleeping configuration after a finite number of firings.

2. We define the **stabilization** $\mathtt{S}[\sigma]$ as the final, sleeping configuration of particles. A crucial fact is that the stabilization does not depend on the order of firings. This **abelian property**, which is proved for ARW in [25], also holds for a more general class of particle systems, the abelian networks [6].

3. The case $\lambda_v = \infty$ for all v is of special interest. It is called internal DLA (**IDLA**). Note that in this case $q_v = 1$, so that each site v "absorbs" the first particle it receives. So this process has a simple description: Each active particle moves according to the Markov chain P, until reaching an unoccupied site or the sink, where it remains forever. A fundamental quantity associated to IDLA is the time when V becomes full:

$$T_{\text{full}} = \min\{t \,:\, \sigma_t(v) > 0 \text{ for all } v \in V\}.$$

1.2. **Plan of the paper.** Our main goal is to upper bound the mixing time of the ARW process. We will first give a method for exact sampling from its stationary distribution (Theorem 1) and then show that the time T_{full} for IDLA to fill V is a **strong stationary time** for the ARW process (Theorem 2). To upper bound the mixing time, it therefore suffices to upper bound T_{full}. Despite the exponential size of its state space $\{0, \mathtt{s}\}^V$, we will prove that the mixing time of the ARW process is not much larger than $\#V$ (Theorem 3).

These three theorems are proved in Section 2 for a general base chain P. Then in Section 3, we examine the case that P is simple random walk on a Euclidean ball intersected with the d-dimensional integer lattice $\mathbb{Z}^d$, with sink at the boundary of the ball. We prove that with high probability

$$T_{\text{full}} \le \#V + (\#V)^\alpha$$

for some $\alpha < 1$ (Theorem 4).

After a brief discussion of lower bounds, we conclude with two conjectures in Section 4.

1.3. **Instructions; Resampling; Abelian Property.** The purpose of this section is to spell out the meaning of the phrase "order of firings" in the abelian property, so that we can give careful proofs of our main results.

Quench the randomness of ARW into a collection of **instructions** $(\rho_{n,v})_{n\in\mathbb{N},\, v\in V}$. The instruction $\rho_{n,v}$ dictates what will happen the nth time v is fired: a particle at v either tries to fall asleep with probability $\lambda_v/(1+\lambda_v)$, or it steps to w with probability $P(v,w)/(1+\lambda_v)$. We assume that each sequence $(\rho_{n,v})_{n\in\mathbb{N}}$ is independent and identically distributed (i.i.d.), and that all $\rho_{n,v}$ are independent.

Fix an ARW configuration σ, and let $(v) = (v_1, v_2, \dots, v_m)$ be a sequence of vertices to be fired in order. We say that (v) is a **legal execution** for σ if $\mathtt{F}_{v_{k-1}} \dots \mathtt{F}_{v_1}[\sigma](v_k) \ge 1$ for all $k = 1, \dots, m$. A legal execution is called **complete** for σ if $\mathtt{F}_{v_m} \cdots \mathtt{F}_{v_1}[\sigma]$ is a sleeping configuration. The **odometer** of an execution (v) is the function

$$f(w) = \#\{1 \le k \le m \,:\, v_k = w\},$$

which counts how many times each vertex w is fired. If (v) is any legal execution for σ and (v') is any complete execution for σ, then the odometer of (v) is less than or equal to the odometer of (v') (see [25] or [6, Lemma 3.4]). This inequality holds pointwise, both in V and in the quenched instructions. In particular, any two legal complete executions for σ have the same odometer. Therefore they use the same subset of the quenched instructions, so they result in the same sleeping configuration $\mathtt{S}[\sigma]$. Occasionally we will write this as $\mathtt{S}_\rho[\sigma]$ to make explicit the dependence on the instructions ρ.

Implicit in the definition of the ARW process is that new independent instructions are used to stabilize at each time step. But if we wish to use the previous paragraph, then the randomness for the entire chain should be expressed in terms of a *single* collection of instructions $\rho = (\rho_{v,j})_{v\in V,\, j\in\mathbb{N}}$. Our first order of business is to check that doing so does not change the distribution of $(\sigma_t)_{t\in\mathbb{N}}$. For this purpose we will use a lemma from [22].

For $f : V \to \mathbb{N}$, write $\mathcal{F}_f$ for the σ-field generated by the instructions $\rho_f := (\rho_{v,n})_{v\in V, n\le f(v)}$ (the "past"), and write $\rho^f := (\rho_{v,\, k+1+f(v)})_{v\in V,\, k\in\mathbb{N}}$ (the "future").

Lemma 1. (Strong Markov Property For Quenched Instructions, [22, Proposition 4]) *Let $F : V \to \mathbb{N}$ be a random function satisfying $\{F = f\} \in \mathcal{F}_f$ for all $f : V \to \mathbb{N}$. Then ρ^F has the same distribution as ρ, and ρ^F is independent of ρ_F.*

Now using a single collection of instructions ρ, let $F_t(v)$ be the number of instructions used at v during the first t time steps of the ARW process, let $\rho_t = \rho_{F_t}$ and let $\rho^t = \rho^{F_t}$; formally, for each $t \ge 1$ we define these inductively by

$$\sigma_t := \mathtt{S}_{\rho^{t-1}}[\sigma_{t-1} + \delta_{u_t}] \tag{3}$$

and $F_t := F_{t-1} + G_t$, where G_t is the odometer for stabilizing the right side.

Lemma 2. (Resampling Future Instructions) *Assume the driving sequence u is independent of the instructions ρ.*

Let $\widetilde{\rho}_1, \widetilde{\rho}_2, \ldots$ be independent families of instructions with the same distribution as ρ, and independent of u. Then $(u_t, F_t, \sigma_t)_{t\in\mathbb{N}}$ has the same distribution as $(u_t, \widetilde{F}_t, \widetilde{\sigma}_t)_{t\in\mathbb{N}}$, where

$$\widetilde{\sigma}_t = \mathtt{S}_{\widetilde{\rho}_t}[\widetilde{\sigma}_{t-1} + \delta_{u_t}]$$

and $\widetilde{F}_t - \widetilde{F}_{t-1}$ is the odometer for stabilizing the right side.

Proof. Fix functions $f_1 \le \cdots \le f_t$ and ARW configurations $\tau_1, \ldots, \tau_t$. Let A_t and $\widetilde{A}_t$ be the events $\{F_s = f_s, \sigma_s = \tau_s,\ s = 1, \ldots, t\}$ and $\{\widetilde{F}_s = f_s, \widetilde{\sigma}_s = \tau_s,\ s = 1, \ldots, t\}$ respectively.

Write P_u for the law of u, and P for the law of $(\rho, \tilde{\rho}_1, \tilde{\rho}_2, \ldots)$. For any *fixed* driving sequence $u_1, \ldots, u_t$, writing $\xi_s = \tau_{s-1} + \delta_{u_s}$ we have

$$\begin{aligned}
\mathrm{P}(A_t) = \prod_{s=1}^{t} \mathrm{P}(A_s | A_{s-1}) &= \prod_{s=1}^{t} \mathrm{P}(\mathtt{S}_{\rho^{s-1}}[\xi_s] = \tau_s, F_s = f_s | A_{s-1}) \\
&= \prod_{s=1}^{t} \mathrm{P}(\mathtt{S}_{\rho^{s-1}}[\xi_s] = \tau_s, F_s = f_s) \\
&= \prod_{s=1}^{t} \mathrm{P}(\mathtt{S}_{\widetilde{\rho}_s}[\xi_s] = \tau_s, \widetilde{F}_s = f_s) \\
&= \prod_{s=1}^{t} \mathrm{P}(\mathtt{S}_{\widetilde{\rho}_s}[\xi_s] = \tau_s, \widetilde{F}_s = f_s | \widetilde{A}_{s-1}) \\
&= \prod_{s=1}^{t} \mathrm{P}(\widetilde{A}_s | \widetilde{A}_{s-1}) = \mathrm{P}(\widetilde{A}_t).
\end{aligned}$$

In the second line we have used that the event A_{s-1} depends only on the past ρ_{s-1}, which is independent of the future ρ^{s-1} by Lemma 1. In the third line we have used that ρ^{s-1} has the same distribution as $\widetilde{\rho}_s$, again by Lemma 1. In the fourth line we have used that $\widetilde{A}_{s-1}$ depends only on the instructions $\tilde{\rho}_r$ for $r \le s-1$, which are independent of $\tilde{\rho}_s$ by hypothesis.

Now let $B = \{u_1 = v_1, \ldots, u_t = v_t\}$. Since u is assumed independent of ρ and $\widetilde{\rho}$, the proof is finished by multiplying by 1_B, taking P_u of both sides, and applying Fubini's theorem:

$$\mathbb{P}(A_t \cap B) = \mathrm{P}_u(\mathrm{P}(A_t) 1_B) = \mathrm{P}_u(\mathrm{P}(\widetilde{A}_t) 1_B) = \mathbb{P}(\widetilde{A_t} \cap B). \qquad \square.$$

The sequence $u = (u_t)_{t\ge 1}$ is called the **driving sequence**. As a consequence of Lemma 2, if u is an i.i.d. sequence then the ARW process (3) is a time-homogeneous Markov chain. For general u, the ARW process is not a Markov chain, but we will see that some techniques from the theory of Markov chains, such as the use of a strong stationary time to bound the mixing time, can still be applied. The reason we are interested in general driving is for future applications when V is a subset of a larger system V', and the driving comes from particles entering V as a result of stabilizing $V' \backslash V$.

In what follows we write $\mathtt{S}^F := \mathtt{S}_{\rho^F}$.

Lemma 3. (Abelian Property) *Let ϕ_1, ϕ_2 be ARW configurations, and let F be the odometer for stabilizing ϕ_1. Then*

$$\mathtt{S}[\phi_1 + \phi_2] = \mathtt{S}^F[\mathtt{S}[\phi_1] + \phi_2]. \tag{4}$$

Proof. Let (v) be a legal complete execution for ϕ_1 with instructions ρ, and let (w) be a legal complete execution for $\mathtt{S}[\phi_1] + \phi_2$ with instructions ρ^F. Then the concatenaion (v, w) is a legal complete execution for $\phi_1 + \phi_2$ with instructions ρ. □

Our main use of the abelian property will be to stabilize the driving particles all at once, instead of one at a time:

$$\sigma_t = \mathtt{S}[\sigma_0 + \phi_t]$$

where

$$\phi_t = \delta_{u_1} + \ldots + \delta_{u_t}. \tag{5}$$

This will allow us to couple the ARW and IDLA processes.

1.4. **Coupling ARW and IDLA.** Recall that ARW with infinite sleep rate is called IDLA. We write $\mathtt{S}_\infty$ for IDLA stabilization *without allowing any particles to fall asleep.* In other words, to perform $\mathtt{S}_\infty$, we let each active particle perform P-walk until reaching an unoccupied site or the sink z. In particular, if ϕ is an all active configuration, then $\mathtt{S}_\infty[\phi]$ is all active.

Every legal IDLA execution is also legal for ARW. One way to stabilize an ARW configuration is therefore to perform IDLA first, and then complete the ARW stabilization:

$$\mathtt{S}[\phi] = \mathtt{S}^G[\mathtt{S}_\infty[\phi]] \tag{6}$$

where G is the odometer for IDLA. In particular, with ϕ_t given by (5), we have a coupling between the IDLA process

$$\eta_t = \mathtt{S}_\infty[\sigma_0 + \phi_t]$$

and the ARW process

$$\sigma_t = \mathtt{S}[\sigma_0 + \phi_t] = \mathtt{S}^{G_t}[\eta_t] \tag{7}$$

where G_t is the odometer for IDLA-stabilizing $\sigma_0 + \phi_t$.

This coupling was used by Shellef [27] to prove nonfixation of certain infinite ARW systems. We will use it to bound the mixing time of the ARW process.

2. Main Results

Now we are ready to prove our exact sampling theorem for the ARW process $\sigma_t = \mathtt{S}[\sigma_{t-1} + \delta_{u_t}]$ with driving sequence $u = (u_t)_{t \in \mathbb{N}}$. We make no assumption that u is i.i.d. or even Markovian, but we will always assume that u is independent of the quenched instructions.

2.1. **Exact sampling.** Write $\bar{\mathtt{S}} := \mathtt{S}_{\bar\rho}$, where $\bar\rho$ is an independent copy of the instructions ρ.

Theorem 1. (Exact sampling from the ARW stationary distribution) *Let* $\sigma_0 = \bar{\mathtt{S}}[1_V]$. *Then for any driving sequence* u *and all* $t \geq 1$,

$$\sigma_t \stackrel{d}{=} \sigma_0.$$

Proof. For any ARW configuration ϕ, consider stabilizing $1_V + \phi$ in two ways: If we first move the extra particles ϕ, then they cannot fall asleep (as every $v \in V$ contains an active particle) so they all perform P-walk until reaching the sink. We can then stabilize 1_V. The second way is to stabilize 1_V, then add the extra particles ϕ, and stabilize again. Using (6) and (4),

$$\bar{\mathtt{S}}^G[1_V] = \bar{\mathtt{S}}[1_V + \phi] = \bar{\mathtt{S}}^H[\bar{\mathtt{S}}[1_V] + \phi]. \tag{8}$$

where G is the odometer for IDLA-stabilizing $1_V + \phi$ to 1_V, and H is the odometer for ARW-stabilizing 1_V. These equalities hold pointwise in $\bar\rho$.

Now take $\phi = \phi_t = \delta_{u_1} + \ldots + \delta_{u_t}$. By the Strong Markov Property, in equation (8) the left side $\stackrel{d}{=} \sigma_0$, and the right side $\stackrel{d}{=} \mathtt{S}[\sigma_0 + \phi_t] = \sigma_t$. $\square$

Theorem 1 identifies a stationary distribution for the ARW process. Next we give a sufficient condition for the stationary distribution to be unique. For $A \subset V$ we say that A can **access** all of V if for every $v \in V$ there exist $a \in A$ and $j \in \mathbb{N}$ such that $P^j(a, v) > 0$. We say that the driving sequence $u = (u_t)_{t\in\mathbb{N}}$ is **thorough** if the set

$$A := \{a \in V \,:\, \mathbb{P}(u_t = a \textit{ infinitely often}) = 1\}$$

can access all of V. In particular, if the base chain P is irreducible, then any nonempty set can access all of V, so *any* driving sequence is thorough.

Lemma 4. *Let* $\eta_t = \mathtt{S}_\infty[\phi_t]$ *be the IDLA-stabilization of* $\phi_t = \delta_{u_1} + \ldots + \delta_{u_t}$. *If* u *is thorough, then* $\mathbb{P}(\eta_t = 1_V$ *eventually*$) = 1$.

Proof. Let $A_t = \{v \in V \,:\, \eta_t(v) = 1\}$. If $A_t = V$, then $A_s = V$ for all $s \geq t$. Otherwise, since u is thorough, it happens infinitely often that P-walk started at u_{t+1} and stopped on exiting A_t has a positive probability to exit in $V \setminus A_t$, in which case A_{t+1} is strictly larger than A_t. Hence $\mathbb{P}(A_t = V \text{ eventually}) = 1$. $\square$

Let

$$\mathcal{R} := \left\{ \sigma \in \{0, \mathtt{s}\}^V \;\middle|\; \begin{array}{l} \sigma(v) = 0 \text{ for all } v \text{ such that } \lambda_v = 0, \text{ and} \\ \sigma(w) = \mathtt{s} \text{ for all } w \text{ such that } \lambda_w = \infty \end{array} \right\}.$$

Lemma 5. (Recurrent ARW Configurations) *If the driving sequence is thorough, then*

- *An ARW configuration* σ *is recurrent if and only if* $\sigma \in \mathcal{R}$; *and*
- $\mathcal{R}$ *is the unique communicating class of recurrent configurations.*

Proof. We first check that if $\sigma_0 \in \mathcal{R}$, then $\sigma_t \in \mathcal{R}$ for all t. For each vertex v with $\lambda_v = 0$, since $\sigma_0(v) = 0$ and no particle will ever fall asleep at v, we have $\sigma_t(v) = 0$ for all t. For each vertex v with $\lambda_v = \infty$, since $\sigma_0(v) = \mathtt{s}$ and the last particle left at v will always fall asleep there, we have $\sigma_t(v) = \mathtt{s}$ for all t.

To finish the proof, we now show that if the driving sequence is thorough, then every ARW configuration σ_0 can access every $\tau \in \mathcal{R}$.

By Lemma 4 there exists T such that $\eta_T = 1_V$, so $\sigma_0 + \phi_T$ has a legal IDLA execution to 1_V. Now starting from 1_V, for each site v such that $\tau(v) = 0$, let the particle at v perform P-walk to the sink. Each of these walks has positive probability to reach the sink before the particle falls asleep (here we use that $\tau(v) = \mathtt{s}$ for all v such that $\lambda_v = \infty$, so all such v are already occupied). Then let all remaining particles fall asleep immediately. This last step succeeds with probability $\prod \frac{\lambda_v}{1+\lambda_v}$, where the product is over all v such that $\tau(v) = \mathtt{s}$ (here we use that $\tau(v) = 0$ for all v such that $\lambda_v = 0$, so the product is > 0). If any step fails, then repeat the whole procedure from the beginning. □

In the case of i.i.d. driving, the ARW process is a Markov chain, so uniqueness of the stationary distribution follows immediately from Lemma 5. The next lemma shows uniqueness for more general driving. We write P for the law of the instructions, P_u for the law of the driving sequence, and $\mathbb{P} = \mathrm{P}_u \times \mathrm{P}$ for their joint law.

Lemma 6. *If the driving sequence u is thorough, then the ARW process has a unique stationary distribution, and the stationary distribution does not depend on u.*

Proof. By Theorem 1 the configuration $\bar{\mathtt{S}}[1_V]$ is stationary and does not depend on u.

To show uniqueness, let μ be any stationary distribution, and let $\sigma_0 \sim \mu$. By stationarity of μ, and the coupling (7), we have for all t and all ARW configurations ξ

$$\mu(\xi) = \mathbb{P}(\sigma_t = \xi) = \mathbb{P}(\mathtt{S}^{G_t}[\eta_t] = \xi).$$

Now for a fixed driving sequence u, by the Strong Markov Property, the future instructions ρ^{G_t} have the same distribution as $\bar{\rho}$. Since η_t depends only on the past instructions ρ_{G_t}, we have (pointwise in u)

$$\begin{aligned}
\mathrm{P}(\mathtt{S}^{G_t}[\eta_t] = \xi) &= \mathrm{P}(\bar{\mathtt{S}}[\eta_t] = \xi) \\
&\geq \mathrm{P}(\bar{\mathtt{S}}[1_V] = \xi, \eta_t = 1_V) \\
&= \pi(\xi)\mathrm{P}(\eta_t = 1_V)
\end{aligned}$$

where π is the distribution of $\bar{\mathtt{S}}[1_V]$. Taking P_u of both sides,

$$\mu(\xi) \geq \pi(\xi)\mathbb{P}(\eta_t = 1_V).$$

By Lemma 4 the right side converges to $\pi(\xi)$ as $t \to \infty$. Since both μ and π sum to 1 we conclude that $\mu = \pi$. □

In the case the stationary distribution of the ARW process is unique, we denote it by $\pi = \pi_{\lambda,P}$. We make a few remarks.

(1) Theorem 1 gives a reasonably fast **sampling algorithm** for $\pi_{\lambda,P}$: The time to stabilize $\mathtt{S}_{\lambda,P}[1_V]$ is upper bounded by the time to stabilize $\mathtt{S}_{0,P}[1_V]$, which is simply the time for all particles to reach the sink z. Writing T_{vz} for the

time for a P-walker started at v to hit z, the time to generate a sample from $\pi_{\lambda,P}$ is therefore at most the sum of hitting times

$$\sum_{v \in V} T_{vz}.$$

(2) The case when the driving sequence is constant, $u_t = v$ for all $t \in \mathbb{N}$, is already interesting. The ARW process $(\sigma_t)_{t\in\mathbb{N}}$ depends on the choice of v, but its stationary distribution does not. One way to see this directly is to define an operator $\mathtt{A}_v$ that adds one chip at v and then stabilizes. This $\mathtt{A}_v$ is a stochastic matrix of size $2^{\#V}$. Then $\mathtt{A}_v\mathtt{A}_w = \mathtt{A}_w\mathtt{A}_v$ by Lemma 3. The stationary distribution π is a left eigenvector of both $\mathtt{A}_v$ and $\mathtt{A}_w$.

(3) Despite the fast sampling algorithm, many properties of the stationary distribution $\pi_{\lambda,P}$ remain mysterious. For example, in the special case that P is simple random walk on a path $\{0, 1, \ldots, L\}$ with sink $z = L$, experiments indicate that $\pi_{\lambda,P}$ is **hyperuniform** in that the variance of the number of sleeping particles grows sublinearly with L. A number of other conjectures about $\pi_{\lambda,P}$ will be discussed in the forthcoming paper [23].

2.2. **Strong stationary time.** Our next goal is to show that the time T_{full} for IDLA to fill V is a strong stationary time for the ARW process. In words, the ARW process is exactly stationary at time T_{full} and all later times.

Theorem 2. (Strong stationary time) *Let $\mathbb{P} = \mathbb{P}_{\sigma_0,u,\lambda,P}$ be the law of the ARW process $(\sigma_t)_{t\in\mathbb{N}}$ with initial state σ_0, thorough driving sequence $u = (u_t)_{t\in\mathbb{N}}$, sleep rate vector λ, and base chain P on state space V. For all ARW configurations $\sigma_0, \xi \in \{0, \mathtt{s}\}^V$, and all $t \in \mathbb{N}$, we have*

$$\mathbb{P}(\sigma_t = \xi \,|\, T_{\text{full}} \le t) = \pi(\xi) \tag{9}$$

where $\pi = \pi_{\lambda,P}$ is the unique stationary distribution of the ARW process.

Proof. We will use the coupling (7) between the IDLA process η_t and the ARW process $\sigma_t = \mathtt{S}^{G_t}[\eta_t]$. For each fixed driving sequence u, the event

$$\{T_{\text{full}} \le t\} = \{\eta_t = 1_V\}$$

depends only on the past instructions ρ_{G_t}, which are independent of the future instructions ρ^{G_t} by the Strong Markov Property. So we have (pointwise in u)

$$\begin{aligned} \mathrm{P}(\sigma_t = \xi, T_{\text{full}} \le t) &= \mathrm{P}(\mathtt{S}^{G_t}[1_V] = \xi, T_{\text{full}} \le t) \\ &= \mathrm{P}(\mathtt{S}^{G_t}[1_V] = \xi)\mathrm{P}(T_{\text{full}} \le t). \end{aligned}$$

Note that G_t depends on u, but the future instructions ρ^{G_t} can be replaced with new independent instructions $\bar{\rho}$ by the Strong Markov Property, so $\mathrm{P}(\mathtt{S}^{G_t}[1_V] = \xi) = \mathrm{P}(\bar{\mathtt{S}}[1_V] = \xi)$ does not depend on u, and equals $\pi(\xi)$ by Theorem 1. So

$$\mathrm{P}(\sigma_t = \xi, T_{\text{full}} \le t) = \pi(\xi)\mathrm{P}(T_{\text{full}} \le t).$$

Recalling $\mathbb{P} = P_u \times \mathrm{P}$, we obtain (9) by taking P_u of both sides. □

2.3. **Upper bounds on mixing time.** Fix an initial configuration σ_0 and thorough driving sequence u. Writing μ_t for the resulting distribution of the ARW process σ_t at time t, for $\epsilon > 0$ let

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) = \min \left\{ t\colon \max_{\sigma_0} ||\mu_t - \pi||_{TV} \le \epsilon \right\}$$

where $||\cdot||_{TV}$ denotes the total variation distance between proability measures. Let

$$t_{\text{full}}(\texttt{IDLA}, u, \epsilon) = \min\{t : \mathbb{P}_{0,u,\infty,P}(T_{\text{full}} > t) \le \epsilon\}$$

be the first time that IDLA, started from the empty initial configuration, fills V with probability $\ge 1 - \epsilon$.

Theorem 3. (Upper bounds on mixing) *For any thorough driving sequence u, any sleep rate vector λ, any base chain P, and any $\epsilon > 0$,*

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) \le t_{\text{full}}(\texttt{IDLA}, u, \epsilon). \tag{10}$$

If the driving sequence $(u_t)_{t\in\mathbb{N}}$ is independent with the uniform distribution on V, then

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) \le \#V \log \#V + \log(1/\epsilon)\#V. \tag{11}$$

Finally, if the driving sequence u is a permutation of V, then the ARW process is exactly stationary at time $\#V$, so

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) \le \#V.$$

Proof. For $t \ge t_{\text{full}}(\texttt{IDLA}, u, \epsilon)$ we have by Theorem 2

$$\mu_t(\xi) \ge \mathbb{P}(\sigma_t = \xi, t \ge T_{\text{full}}) = \pi(\xi)\mathbb{P}(t \ge T_{\text{full}}) \ge (1 - \epsilon)\pi(\xi).$$

Summing over ARW configurations ξ for which $\pi(\xi) > \mu_t(\xi)$ yields

$$||\pi - \mu_t||_{TV} = \sum_\xi (\pi(\xi) - \mu_t(\xi))_+ \le \epsilon \sum_\xi \pi(\xi) \le \epsilon$$

which proves (10).

The inequality (11) follows from a standard coupon collector bound; see, for example, [17, Prop. 2.4], which implies that for $t > \#V \log \#V + \log(1/\epsilon)\#V$ we have

$$\mathbb{P}(\phi_t \ge 1_V) \ge 1 - \epsilon.$$

On the event $\phi_t \ge 1_V$, letting all extra particles perform P-walk until reaching the sink yields a legal execution from $\sigma_0 + \phi_t$ to 1_V, so the total variation distance between the laws of $\texttt{S}[\sigma_0 + \phi_t]$ and $\texttt{S}[1_V]$ is at most ϵ.

Finally, if the driving sequence is a permutation of V, then for $t = \#V$ we have $\phi_{\#V} = 1_V$, so $\sigma_t = \texttt{S}[\sigma_0 + 1_V] \overset{d}{=} \texttt{S}[1_V]$ is exactly stationary by Theorem 1. □

In the next section we will upper bound the right side of (10) when V is a discrete ball in $\mathbb{Z}^d$.

3. Bounds for the fill time of IDLA

For $r > 0$ let $B_r = B(0,r) \cap \mathbb{Z}^d = \{x \in \mathbb{Z}^d : |x| < r\}$ be the Euclidean ball of radius r intersected with $\mathbb{Z}^d$, viewed as a graph with nearest-neighbor adjacencies. Here $|x| := (x_1^2 + \cdots + x_d^2)^{1/2}$ denotes the Euclidean norm. Collapse the boundary

$$\partial B_r = \{y \in \mathbb{Z}^d \setminus B_r \colon |y - x| = 1 \text{ for some } x \in B_r\}$$

to a sink vertex.

We consider IDLA driven by simple random walk on B_r, in two different scenarios: **central driving** in which all particles start at 0, and **uniform driving** in which each particle starts at an indpendent random location in B_r.

Theorem 4. (Upper bound for the fill time of IDLA)

Let T_{full} be the time for IDLA with either central or uniform driving to fill B_r, and let $N = \#B_r$.

- *In dimension $d = 1$, for any $\alpha > \frac{1}{2}$ there is a constant $R = R(\alpha)$ such that for all $r \geq R$*

$$\mathrm{P}\{T_{\text{full}} > N + N^\alpha\} \leq \exp\{-c_1 r^{\alpha - \frac{1}{2}}\}.$$

- *In dimension $d \geq 2$, let $\alpha = 1 - \frac{1}{3d}$. Then for all sufficiently large r*

$$\mathrm{P}\{T_{\text{full}} > N + N^\alpha\} \leq \exp\{-c_2 r^{1/4}\}.$$

These two bounds are proved in Sections 3.2 and 3.3, respectively. The exponent $\frac{1}{2}$ is optimal for $d = 1$, but $1 - \frac{1}{3d}$ is not optimal for $d \geq 2$. Using methods of [1, 2, 13, 14], it can be improved to $1 - \frac{1}{d} + \delta$, at the cost of reducing $r^{1/4}$ on the right side to r^c for $c = c(d, \delta) > 0$; but we do not pursue this variation. The c_1 and c_2 above are absolute constants; the proof will show that $c_1 = \frac{1}{41}$ and $c_2 = \frac{1}{5}$ suffice.

Combining Theorem 4 with the bound (10), we obtain the following upper bounds on the mixing time of the ARW process:

Corollary 5. (Upper bound for ARW mixing on the ball)

Let u be either the central or uniform driving sequence on the ball B_r, let λ be any sleep rate vector, let P be the simple random walk on B_r, and let $N = \#B_r$. Then for any $\epsilon > 0$, we have for sufficiently large r,

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) \leq N + N^{1 - \frac{1}{3d}}.$$

An interesting question (see Conjecture 10) is whether the ARW process achieves cutoff in total variation at an earlier time ζN for some $\zeta < 1$.

By covering the torus $\mathbb{Z}_n^d$ with Euclidean balls, we obtain the following corollary, proved in Section 3.4.

Corollary 6. (Upper bound for ARW mixing on the torus.)

Let u be the uniform driving sequence on the discrete torus $\mathbb{Z}_n^d \setminus \{z\}$ with sink at z. Let λ be any sleep rate vector, let P be the simple random walk on $\mathbb{Z}_n^d$, and let $N = n^d$. Then for any $\epsilon > 0$ we have for sufficiently large n

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) \leq N + d^{1/2} N^{1 - \frac{1}{3d}}.$$

3.1. **Concentration inequalities.** To prepare for the proof of Theorem 4 we recall three concentration inequalities.

1. **Azuma-Hoeffding inequality.** Let S_t be a martingale with bounded differences $|S_t - S_{t-1}| \le b_t$. Then

$$\mathrm{P}(S_t - S_0 \ge s) \le \exp(-\frac{2s^2}{\sum_{i=1}^t b_i^2}). \tag{12}$$

2. **Bernstein inequality.** Let $X_1, \dots, X_t$ be independent mean zero random variables with $|X_i| \le 1$. Then

$$\mathrm{P}(|X_1 + \ldots + X_t| \ge s) \le 2\exp(-\frac{s^2}{2(\sum_{i=1}^t \mathrm{E}X_i^2 + \frac{1}{3}s)}).$$

   We will apply this inequality in the case $X_i = Y_i - \mathrm{E}Y_i$ where the Y_i are independent Bernoulli random variables. Writing $S = Y_1 + \ldots + Y_t$ and $\mu = \mathrm{E}S$, we obtain for $s \le \mu$

$$\mathrm{P}(|S - \mu| \ge s) \le 2\exp(-\frac{s^2}{3\mu}). \tag{13}$$

3. **Time to exit a ball.** Consider a simple random walk in $\mathbb{Z}^d$ starting at any point in the ball B_r. Let T be the first exit time of the walk from B_r. Then for sufficiently large t

$$\mathrm{P}(T \ge t) \le \exp(-\frac{t}{3(r+1)^2}). \tag{14}$$

   This follows from the fact that $\mathrm{E}T \le (r+1)^2$ regardless of where the random walk starts. By Markov's inequality and the strong Markov property, $\mathrm{P}(T \ge (k+1)e(r+1)^2 \,|\, T \ge ke(r+1)^2) \le \frac{1}{e}$ for all $k \in \mathbb{N}$. Therefore $\mathrm{P}(T \ge t) \le (\frac{1}{e})^{\lfloor t/e(r+1)^2 \rfloor}$, which implies (14) for sufficiently large t.

3.2. **Upper bound in dimension 1.** Consider IDLA with $2r+n$ particles in the interval $(-r, r)$, with particles killed if they reach an endpoint r or $-r$. By the abelian property, we may assume all particles are present at the beginning instead of being added one at a time. We stabilize IDLA in discrete time steps where at each time step, one particle moves either left by 1 or right by 1 with probability 1/2 each. For definiteness, we always move the leftmost active particle (recall that a particle is active in IDLA if and only if there is at least one other particle located at the same site). We keep track of the quantity

$$S_t = \sum_{i=1}^{2r+n} x_{i,t}$$

where $x_{i,t}$ is the location of the ith particle after t time steps. This S_t is a martingale with $|S_t - S_{t-1}| \le 2$; it measures the total left-right "imbalance" of the particles at time t.

Lemma 7. *Fix $\epsilon > 0$ and $r^{\frac{1}{2}+\epsilon} \le n \le r$. For IDLA with $2r+n$ particles in $(-r,r)$ with either central or uniform driving, there is a constant $R = R(\epsilon)$ such that for all $r \ge R$*

$$\mathrm{P}\{\textit{No particles reach } r\} \le \exp(-\frac{n}{10r^{1/2}}). \tag{15}$$

Proof. The total number of particles is $2r+n$, and at most one particle can settle at each site in $(-r,r)$. So on the event $\mathcal{B} := \{\text{No particles reach } r\}$, at least n particles exit at $-r$. Each exiting particle contributes $-r$ to the total imbalance S_T, where T is the time of stabilization of IDLA. If the interval $(-r,r)$ is completely full at time T, then the total contribution of the particles inside $(-r,r)$ to S_T is zero; moreover, every unoccupied site in $(-r,r)$ results in an additional particle exiting at $-r$, which can only make S_T smaller. Hence

$$\mathcal{B} \subset \{S_T \le -nr\}.$$

Now for any $t \ge 0$,

$$\mathrm{P}(\mathcal{B}) \le \mathrm{P}\{T > t\} + \mathrm{P}\{T \le t, \mathcal{B}\} \tag{16}$$

and since $S_t = S_T$ for all $t \ge T$,

$$\begin{aligned}\{T \le t, \mathcal{B}\} &\subset \{T \le t,\, S_T \le -nr\}\\ &= \{T \le t,\, S_t \le -nr\}\\ &\subset \{S_t - S_0 \le \frac{-nr}{2}\} \cup \{S_0 < -\frac{nr}{2}\}.\end{aligned} \tag{17}$$

We now make our choice of $t = nr^{5/2}$. By Azuma-Hoeffding (12),

$$\mathrm{P}\{S_t - S_0 \le \frac{-nr}{2}\} \le \exp(-\frac{2(nr/2)^2}{4t}) \le \exp(-\frac{n}{8r^{1/2}}) \tag{18}$$

for sufficiently large r. In the last inequality we have used that $n \ge r^{\frac{1}{2}+\epsilon}$.

With central driving, $S_0 = 0$. With uniform driving, S_0 is a sum of $2r+n \le 3r$ independent random variables with the uniform distribution on $(-r,r)$, so by Azuma-Hoeffding

$$\mathrm{P}(S_0 < -\frac{nr}{2}) \le \exp(-\frac{2(\frac{1}{2}nr)^2}{(3r)r^2}) \le \frac{1}{3}\exp(-\frac{n^2}{6r}) \tag{19}$$

for sufficiently large r. Combining this with (17) and (18) yields $\mathrm{P}(T \le t, \mathcal{B}) \le 2\exp(-\frac{n}{8r^{1/2}})$.

Finally, to bound the first term of (16), since $r \le n$,

$$\mathrm{P}\{T > t\} \le \sum_{i=1}^{2r+n} \mathrm{P}\{T_i > \frac{t}{3r}\}$$

where T_i is the total number of steps of taken by the ith particle during IDLA. By the simple random walk estimate (14), the right side is at most $3r\exp(-\frac{t}{9r(r+1)^2}) \le \frac{1}{2}\exp(-\frac{n}{10r^{1/2}})$ for sufficiently large r. Here we again use the lower bound $n \ge r^{1/2+\epsilon}$. ∎

Now we are ready to prove Theorem 4 in dimension 1.

Proof. On the event that IDLA with central driving does not fill B_r, either $-r$ or r receives no particles. By Lemma 7 and symmetry, this event has probability $\leq 2\exp(-\frac{n}{10r^{1/2}})$, which completes the proof in the case of central driving.

Now consider uniform driving. On the event that IDLA with uniform driving does not fill B_r, there is some $X \in B_r$ not visited by any particles, so all particles leaving $(-r, X)$ must exit to the left and all particles leaving (X, r) must exit to the right. For fixed $x \in (-r, r)$ let I be the larger of the two intervals $(-r, x)$ and (x, r). Let N_1 be the number of particles starting in I. Then N_1 has the binomial$(2r + n, p)$ distribution where $p = \#I/(2r)$. Since $\#I \geq r$ we have $EN_1 \geq \#I + \frac{n}{2}$. So by Azuma-Hoeffding,

$$\mathrm{P}\{N_1 < \#I + \frac{n}{4}\} \leq \exp(-\frac{2(n/4)^2}{2r+n}) \leq \exp(-\frac{n^2}{24r}).$$

On the complementary event $\{N_1 \geq \#I + \frac{n}{4}\}$, the probability that all N_1 particles exit I on one side is by Lemma 7 at most

$$\exp(-\frac{n/4}{(10\#I)^{1/2}}) \leq \exp(-\frac{n}{40r^{1/2}}).$$

Taking a union bound over $x \in (-r, r)$, the probability that IDLA with $2r+n$ particles does not fill $(-r, r)$ is at most $\exp(-\frac{n}{41r^{1/2}})$ for sufficiently large r. □

3.3. **Upper bound in higher dimensions.** To prove Theorem 4 in dimensions $d \geq 2$ we will use the method of Lawler, Bramson, and Griffeath [16]. In their shape theorem for IDLA, driving is from the origin and there is no sink. Here we adapt their method to uniform driving with sink.

Fix $d \geq 2$ and let $G_r(y, z)$ be the expected number of visits to z by simple random walk started at y before exiting the ball $B_r = \{z \in \mathbb{Z}^d : |z| < r\}$. We recall that G_r is symmetric in y and z, and for all $z \in B_r$

$$G_r(z, z) \leq c_1 \log r \tag{20}$$

for a constant c_1 depending only on d. In the proof of Theorem 4, we will use the following lower bound.

Lemma 8. *There is a constant $c_2 > 0$ depending only on d, such that for all $z \in B_r$*

$$\sum_{y \in B_r} \mathrm{P}_y(\tau_z < \tau_r) \geq c_2 \frac{r}{\log r}.$$

Proof. Recall that $G_r(z, z)\mathrm{P}_y(\tau_z < \tau_r) = G_r(y, z) = G_r(z, y)$; the first equality follows from the strong Markov property by noting that if the walk visits z before exiting B_r, then the number of visits to z before exiting B_r has a geometric distribution with mean $G_r(z, z)$. Now by (20),

$$c_1 \log r \sum_{y \in B_r} \mathrm{P}_y(\tau_z < \tau_r) \geq \sum_{y \in B_r} G_r(z, y).$$

The right side equals the expected time $\mathrm{E}_z\tau_r$ for simple random walk started at z to exit B_r. As a function of z, this expected time has discrete Laplacian -1 and vanishes outside B_r, so

$$\mathrm{E}_z\tau_r \geq r^2 - |z|^2 \geq r(r-|z|).$$

This completes the proof for all $z \in B_{r-1}$. For $z \in B_r - B_{r-1}$ note that simple random walk started at z has a constant probability of hitting B_{r-1} before exiting B_r, so $\mathrm{E}_z\tau_r$ is at least a constant times r. $\square$

Proof of Theorem 4 in dimensions $d \geq 2$. We consider first the case of uniform driving. Perform IDLA starting with $N + N^\alpha$ particles at independent uniform locations in the ball B_r, where $N = \#B_r$. Denoting by A_r the resulting random subset of B_r where particles stabilize, we must show that $\mathrm{P}(A_r \neq B_r) \leq \exp(-\frac{1}{4}r^{1/4})$ for sufficiently large r.

We modify the proof of the inner bound in [16] to account for killing at ∂B_r and uniform driving. For $z \in B_r$, denote by $E_z = \{z \notin A_r\}$ the event that no particle visits z during IDLA. By a union bound over z, it suffices to show that for sufficiently large r

$$\mathrm{P}(E_z) < 4\exp(-\frac{1}{4}r^{1/4}) \quad \text{for all } z \in B_r. \tag{21}$$

Fix an arbitrary ordering of the particles, and define

τ_z^i = time of first visit to z by the ith particle in simple random walk;

τ_r^i = time of first exit of B_r by the ith particle in simple random walk;

σ_i = stopping time of ith particle in the IDLA stabilization process.

$$M = \sum_{i=1}^{N+N^\alpha} \mathbb{1}_{\{\tau_z^i < \tau_r^i\}}; \tag{22}$$

$$L = \sum_{i=1}^{N+N^\alpha} \mathbb{1}_{\{\sigma_i < \tau_z^i < \tau_r^i\}};$$

$$\tilde{L} = \sum_{y \in B_r} \mathbb{1}_{\{\tau_z^y < \tau_r^y\}}.$$

Here τ_z^y is first hitting time of z for a simple random walk started at y; and τ_r^y is the first exit time of B_r for a simple random walk started at y.

Now we have for any $a \in \mathbb{R}$,

$$\begin{aligned}\mathrm{P}(E_z) = \mathrm{P}(M - L = 0) &\leq \mathrm{P}(M \leq a) + \mathrm{P}(L \geq a) \\ &\leq \mathrm{P}(M \leq a) + \mathrm{P}(\tilde{L} \geq a).\end{aligned} \tag{23}$$

The last inequality follows from the observation that after IDLA stabilization, each vertex can be occupied by at most one particle, so $\tilde{L} \geq L$.

Next we will show $\mathrm{E}M$ is substantially larger than $\mathrm{E}\tilde{L}$. Since $N + N^\alpha$ particles are dropped uniformly in B_r, and $N = \#B_r$, we have:

$$\mathrm{E}M = \frac{N + N^\alpha}{N} \sum_{y \in B_r} \mathrm{P}(\tau_z^y < \tau_r^y) = (1 + N^{\alpha-1})\mu. \tag{24}$$

where $\mu = \mathrm{E}\tilde{L}$. By Lemma 8 we have $\mu \geq c_2 \frac{r}{\log r}$ for all $z \in B_r$. Now we make our choice of exponent: $\alpha = 1 - \frac{1}{3d}$, so that for sufficiently large r

$$N^{\alpha-1} \geq c_3 r^{d(\alpha-1)} = c_3 r^{-\frac{1}{3}} \geq c_3 \mu^{-\frac{1}{3}-\epsilon}$$

for any $\epsilon > 0$ and $r \geq R(\epsilon)$, where the constant $c_3 > 0$ depends only on d. This implies

$$\mathrm{E}M - \mathrm{E}\tilde{L} \geq c_4 \mu^{\frac{2}{3}-\epsilon}.$$

Taking $a = (\mathrm{E}\tilde{L} + \mathrm{E}M)/2$ we have

$$\mathrm{P}(\tilde{L} \geq a) = \mathrm{P}(\tilde{L} - \mu \geq \frac{1}{2} c_4 \mu^{\frac{2}{3}-\epsilon})$$

By Bernstein's inequality (13), the right side is at most $2\exp(-\frac{1}{4}\mu^{\frac{1}{3}-2\epsilon})$.

Likewise, since $\mu \leq \mathrm{E}M \leq 2\mu$, we have by Bernstein's inequality

$$\mathrm{P}(M \leq a) \leq \mathrm{P}(M - \mathrm{E}M \leq -\frac{1}{4} c_4 (\mathrm{E}M)^{\frac{2}{3}-2\epsilon}) \leq 2\exp(-\frac{1}{4}\mu^{1/3-2\epsilon}).$$

We conclude from (23) that

$$\mathrm{P}(E_z) \leq 4\exp(-\frac{1}{4} r^{1/4})$$

which completes the proof in the case of uniform driving.

Now we adapt the proof to handle the case of central driving. Note driving enters the proof only in equation (24). In the case of central driving, we have instead

$$EM = (N + N^\alpha)\mathrm{P}(\tau_0^z \leq \tau_0^r).$$

To complete the proof in this case, it suffices to show

$$(N + \frac{1}{2}N^\alpha) G_r(0,z) \geq \sum_{y \in B_r} G_r(y,z) \quad \text{for all } z \in B_r. \tag{25}$$

This inequality differs in two respects from [16, Lemma 3], in which the N^α term is absent but z is restricted to the smaller ball $B_{(1-\epsilon)r}$. To prove (25), we use the **divisible sandpile** of [21]. Let $N_0 = N + C_0 N^{1-\frac{1}{d}}$ and

$$f(z) := N_0 G_r(0,z) - \sum_{y \in B_r} G_r(y,z).$$

This f has discrete Laplacian $1 - N_0\delta_0$ in B_r and vanishes on ∂B_r. The divisible sandpile in $\mathbb{Z}^d$ started with point mass $N_0\delta_0$ fills B_r. Writing u for the divisible sandpile odometer, u and f have the same discrete Laplacian in B_r; moreover $u \leq C_2$ on ∂B_r. (This is the last displayed equation in the proof of [21, Theorem 3.3].) By the maximum principle, $f - u \geq -C_2$ in B_r. Since $G_r(0,z) \geq C_3$ in B_r, it follows that

$$N_1 G_r(0,z) \geq \sum_{y \in B_r} G_r(y,z) \quad \text{for all } z \in B_r$$

where $N_1 = N_0 + C_4$. $\square$

3.4. **Upper bound for the torus.**

Proof of Corollary 6. The case $d = 1$ is immediate from Theorem 4.

For $d \geq 2$, let n be large enough so that the bounds in Theorem 4 hold for $r = n/4$. We cover the torus $\mathbb{Z}_n^d \setminus \{z\}$ by Euclidean balls of radius of r, while leaving the sink z uncovered. A simple but inefficient way to do this, which suffices for our purpose, is to take *all* balls $B(x, r)$ for $x \in \mathbb{Z}_n^d$ which do not contain z.

Now let B be one of the covering balls. By the abelian property, if IDLA with sink at ∂B fills B, then IDLA with sink at z also fills B.

One can check that for all $d \geq 2$ we have $(\#B/n^d)^{1/3d} \geq 1.05 d^{-1/2}$ for sufficiently large n. When $d = 2$, this bound follows from the fact that $\#B \geq 3(n/4)^2$ for sufficiently large n (since $\pi > 3$). For general d, the bound follows from the formula for the volume of the d-dimensional ball, along with the estimates $k^k/e^{k-1} \leq k! \leq (k+1)^{k+1}/e^k$.

Let $N = n^d$ and $\alpha = 1 - \frac{1}{3d}$. After dropping $t = N + d^{1/2}N^\alpha$ particles uniformly at random in $\mathbb{Z}_n^d \setminus \{z\}$, the number of particles starting in B is a sum of t independent Bernoulli random variables of mean $\#B/N$. This sum has mean $\geq (\#B) + 1.05(\#B)^\alpha$, so by Bernstein's inequality (13), the probability that B starts with less than $(\#B) + (\#B)^\alpha$ particles is at most $2\exp(-c(\#B)^{2\alpha-1})$, where $c > 0$ depends only on d. For sufficiently large n, on the event that B starts with at least $(\#B) + (\#B)^\alpha$ particles, the probability that IDLA does not fill B is at most $\exp(-c_2(n/4)^{1/4})$, by Theorem 4.

By a union bound over the covering balls, the probability that IDLA does not fill $\mathbb{Z}_n^d \setminus \{z\}$ is at most $n^d[2\exp(-c(\#B)^{2\alpha-1}) + \exp(-c_2(n/4)^{1/4})]$. Taking n large enough so that this probability is $< \epsilon$, we obtain from (10)

$$t_{\text{mix}}(\texttt{ARW}, u, \epsilon) \leq t_{\text{full}}(\texttt{IDLA}, u, \epsilon) \leq t. \qquad \square$$

3.5. **Lower bounds and wheel example.** In this section we state some lower bounds for the fill time of IDLA. The proofs are straightforward, so we indicate only the main idea.

The first lower bound shows that the exponent α in Theorem 4 cannot be improved to less than $\frac{1}{2}$ in dimension $d = 1$ or $1 - \frac{1}{d}$ in dimensions $d \geq 2$.

Proposition 7. *For $d \geq 1$, let T_{full} be the time for IDLA to fill the ball $B_r \subset \mathbb{Z}^d$, with sink at $\mathbb{Z}^d \setminus B_r$. Let $N = \#B_r$ and let $\beta = \max\{\frac{1}{2}, 1 - \frac{1}{d}\}$. The following holds for any driving sequence u satisfying $u_t \in B_{r-2}$ for all t, and also for the uniform driving sequence on B_r: For all $b > 0$ there exists $c > 0$ such that for all sufficiently large r*

$$\mathrm{P}(T_{\text{full}} > N + bN^\beta) > c.$$

The idea of the proof is to split the IDLA stabilization into two stages: In stage one, stabilize all particles starting inside B_{r-2}, stopping them when they hit ∂B_{r-2}; and in stage two, finish the stablization procedure. Let M be the number of particles resting at ∂B_{r-2} at the end of stage one. If M is large ($\geq CN^\beta$), then with nonvanishing probability, at least $2bN^\beta$ particles will exit B_r in stage two. If M is small ($< CN^\beta$), then with nonvanishing probability, $B_r \setminus B_{r-2}$ will not fill up in stage two.

Next we observe that for general graphs, T_{full} is not always upper bounded by $(1+o(1))\#V$. The wheel graph provides an example.

Proposition 8. *Let W_n be the graph formed by adding a sink vertex z to the cycle $\mathbb{Z}_n$, with an edge between each vertex of $\mathbb{Z}_n$ and z. For IDLA with uniform driving on W_n, there is a constant $b > 0$ such that*

$$\mathrm{P}(T_{\text{full}} > \frac{bn\log n}{\log\log n}) \to 1 \tag{26}$$

as $n \to \infty$.

The idea of the proof is that all particles move only a short distance ($\leq C\log n$) before falling into the sink. With high probability there are many tiny intervals (of length $c\log\log n$) that receive no chips from the uniform driving, and for at least one of these empty intervals, no chips reach the center of the interval during IDLA.

4. Conjectures

We conclude by stating two conjectures.

Conjecture 9. (Time for IDLA to fill a transitive graph)

Let V be a transitive graph with one vertex designated as sink. Then

$$\frac{T_{\text{full}}}{\#V} \to 1 \qquad \textit{in probability as } \#V \to \infty.$$

Conjecture 10. (Cutoff for ARW at the stationary density)

Let u be the uniform driving sequence on $B_r = B(0,r) \cap \mathbb{Z}^d$. Let $0 < \lambda < \infty$ be any constant sleep rate, and let P be the simple random walk on B_r with sink at $\mathbb{Z}^d \setminus B_r$. There exists a constant $\zeta = \zeta(\lambda, d) < 1$ such that

(1) For any $\epsilon > 0$,

$$\frac{t_{\text{mix}}(\mathtt{ARW}, u, \epsilon)}{\#B_r} \to \zeta \qquad \textit{as } r \to \infty.$$

(2) Writing $|\mathtt{S}[1_{B_r}]|$ for the number of particles in the ARW stationary state on B_r, we have

$$\frac{|\mathtt{S}[1_{B_r}]|}{\#B_r} \to \zeta \qquad \textit{in probability as } r \to \infty.$$

(3) $\zeta = \zeta_c$, the critical density for ARW stabilization in $\mathbb{Z}^d$.

We remark that the inequality $\zeta_c < 1$ has been proved in dimensions $d \geq 3$ by Taggi [29] and in dimension 1 by Hoffman, Richey, and Rolla [10]. It remains open in dimension 2.

Acknowledgements

The first author thanks Christopher Hoffman, Leonardo Rolla, and Vladas Sidoravicius for inspiring conversations. The second author thanks his parents for their long time supporting him.

References

[1] Amine Asselah and Alexandre Gaudillière (2013). From logarithmic to subdiffusive polynomial fluctuations for internal DLA and related growth models. *Annals of Probability* 41.3A: 1115–1159.

[2] Amine Asselah and Alexandre Gaudillière (2013). Sublogarithmic fluctuations for internal DLA. *Annals of Probability* 41.3A:1160–1179.

[3] Amine Asselah, Leonardo T. Rolla, and Bruno Schapira (2019). Diffusive bounds for the critical density of activated random walks. arXiv:1907.12694

[4] Riddhipratim Basu, Shirshendu Ganguly, Christopher Hoffman, and Jacob Richey (2019). Activated random walk on a cycle. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques* 55(3):1258–1277.

[5] Per Bak, Chao Tang and Kurt Wiesenfeld (1987). Self-organized criticality: an explanation of the $1/f$ noise, *Physical Review Letters* 59(4):381–384.

[6] Benjamin Bond and Lionel Levine (2016). Abelian networks I. Foundations and examples. *SIAM Journal on Discrete Mathematics* 30:856–874.

[7] Deepak Dhar (1990). Self-organized critical state of sandpile automaton models. *Physical Review Letters* 64:1613–1616.

[8] Anne Fey, Lionel Levine and David B. Wilson (2010). Driving sandpiles to criticality and beyond. *Physical Review Letters* **104**:145703

[9] Anne Fey, Ronald Meester, and Frank Redig (2009). Stabilizability and percolation in the infinite volume sandpile model, *Annals of Probability* 37(2):654–675.

[10] Christopher Hoffman, Jacob Richey, and Leonardo T. Rolla (2020). Active Phase for Activated Random Walk on $\mathbb{Z}$, https://arxiv.org/abs/2009.09491

[11] Bob Hough, Daniel C. Jerison, and Lionel Levine (2019), Sandpiles on the square lattice. *Communications in Mathematical Physics* 367:33–87.

[12] Bob Hough and Hyojeong Son (2021), Cut-off for sandpiles on tiling graphs. *Annals of Probability* 49.2: 671–731.

[13] David Jerison, Lionel Levine, and Scott Sheffield (2012), Logarithmic fluctuations for internal DLA. *Journal of the American Mathematical Society*, 25:272–301.

[14] David Jerison, Lionel Levine, and Scott Sheffield (2013), Internal DLA in higher dimensions. *Electronic Journal of Probability* 18(98):1–14.

[15] Hang-Hyun Jo and Hyeong-Chai Jeong (2010). Comment on "Driving sandpiles to criticality and beyond". *Physical Review Letters* 105.1: 019601.

[16] Gregory F. Lawler, Maury Bramson, and David Griffeath (1992). Internal Diffusion Limited Aggregation, *Annals of Probability* 20(4):2117–2140.

[17] David A. Levin and Yuval Peres (2017). Markov chains and mixing times, 2nd ed., Vol. 107, American Mathematical Society.

[18] Lionel Levine (2015). Threshold state and a conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle, *Communications in Mathematical Physics*, 335(2):1003–1017.

[19] Lionel Levine, Mathav Murugan, Yuval Peres, and Baris E. Ugurcan (2016). The divisible sandpile at critical density. *Annales Henri Poincaré* 17(7):1677–1711.

[20] Lionel Levine, Wesley Pegden, and Charles K. Smart (2017). The Apollonian structure of integer superharmonic matrices. *Annals of Math* 186:1–67.

[21] Lionel Levine and Yuval Peres (2009). Strong spherical asymptotics for rotor-router aggregation and the divisible sandpile. *Potential Analysis* 30(1), 1–27.

[22] Lionel Levine and Vittoria Silvestri (2020). How far do activated random walkers spread from a single source? arXiv:2011.02535

[23] Lionel Levine and Vittoria Silvestri (2021). Universality conjectures for Activated Random Walk. In preparation.

[24] Leonardo T. Rolla (2020). Activated Random Walks on $\mathbb{Z}^d$. *Probability Surveys* 17:478–544.

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

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

[27] Eric Shellef, Nonfixation for activated random walks (2009). arXiv:0910.3338

[28] Alexandre Stauffer and Lorenzo Taggi (2018). Critical density of activated random walks on transitive graphs. *Annals of Probability* 46(4):2190–2220.

[29] Lorenzo Taggi (2019). Active phase for activated random walks on $\mathbb{Z}^d$, $d \geq 3$, with density less than one and arbitrary sleeping rate. *Annales de l'Institut Henri Poincaré Probabilités et Statistiques* 55(3):1751–1764.

[30] Lorenzo Taggi (2020). Essential enhancements in Abelian networks: continuity and uniform strict monotonicity. arXiv:2003.00932

Lionel Levine, Department of Mathematics, Cornell University, Ithaca, NY 14853. https://pi.math.cornell.edu/~levine

Feng Liang, Department of Mathematics, Cornell University, Ithaca, NY 14853.