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

CHRISTOPHER HOFFMAN, JACOB RICHEY, AND HYOJEONG SON

Abstract. We study the one-dimensional activated random walk (ARW) on $\mathbb{Z}$ started with a point source of n particles at the origin. Let ρ_* denote the critical density. We prove that, uniformly throughout a linear-size bulk window around the source, the probability that a site in the stabilized configuration contains a sleeping particle converges to ρ_* as n grows.

1. Introduction

Many complex systems share a common pattern: they accumulate stress or energy gradually, then release it in sudden bursts. This pattern manifests 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 systems display this behavior through price fluctuations in stocks and commodities that conform to power-law distributions. These and other observations point to a fundamental organizing principle that Bak, Tang, and Wiesenfeld formalized in 1987 as *self-organized criticality* (SOC) [1]. The defining characteristic of SOC is that systems autonomously migrate toward critical states through their internal dynamics, requiring neither precise parameter calibration nor external forcing. The prevalence of this mechanism across both engineered and natural phenomena underscores its importance as a unifying framework for understanding complex system behavior.

After SOC was introduced, researchers pursued a central question: which mathematical models could universally represent this critical behavior? Among the candidates that emerged from this research program, the *deterministic sandpile model* has received considerable attention. The model operates on a graph structure where vertices accumulate nonnegative quantities of chips. When the chip count at a vertex reaches or surpasses its degree, that vertex undergoes a toppling event, transferring one chip to each adjacent vertex. These transfers can initiate cascading sequences of topplings throughout the system. The resulting dynamics generate complex fractal patterns [14], demonstrating the model's capacity to produce nontrivial emergent structures. Nevertheless, the strictly deterministic mechanics of the model impose constraints that prevent the manifestation of certain critical phenomena [6, 7, 10–12]. A genuinely universal model must demonstrate that its large-scale properties remain invariant across different microscopic specifications, ensuring that the system's essential characteristics persist despite perturbations to its underlying details.

To address these limitations, researchers developed the *stochastic sandpile model*, a probabilistic variant of the deterministic framework. The fundamental distinction lies in the toppling mechanism. Rather than distributing exactly one chip to every adjacent vertex during a toppling event, this stochastic variant disperses a predetermined quantity of particles among neighboring vertices selected randomly according to a specified probability distribution. Recent investigations suggest that this probabilistic approach successfully captures universality properties [2,4,5,9], marking a significant advancement toward a robust universal model. However, the stochastic framework poses analytical challenges. Pairwise correlations emerge between particle trajectories, substantially increasing the complexity of characterizing the system's dynamics.

The *Activated Random Walk* (ARW) model represents a promising recent development in the search for a universal SOC framework. The model establishes a continuous-time interacting particle system on $\mathbb{Z}$ with an initialization procedure that assigns particle counts to each site through independent and identically distributed sampling from an ergodic distribution characterized by mean density ρ_*. The model distinguishes between two particle states: active and sleeping configurations. Particles in the active state execute symmetric random walks with unit rate while simultaneously transitioning to the sleeping state at rate $\lambda \in (0, \infty)$. Conversely, sleeping particles maintain their positions without movement until contact with an active particle triggers awakening, with the reactivation occurring immediately when additional particles occupy the same location. The analytical tractability of the ARW model stems from the independence structure governing particle behavior, specifically the fact that individual particles conduct random walks and enter dormant states through independent mechanisms. This independence property distinguishes the ARW model from the stochastic sandpile framework and has facilitated substantial progress in theoretical understanding, leading to noteworthy analytical results.

A significant theoretical advance emerged from recent work by the first author in collaboration with Junge and Johnson [8], establishing that critical densities are not only well-defined but also identical across multiple ARW variants for any $\lambda > 0$. The analysis encompassed four distinct formulations: the driven-dissipative model, the point-source model, the fixed-energy model on $\mathbb{Z}$, and the fixed-energy model on the cycle. The demonstration that each of these frameworks possesses a critical density and that these values coincide, denoted collectively as ρ_*, provides rigorous validation of the *density conjecture*. This unification across different model specifications strengthens the case for the ARW framework as a viable universal SOC model, particularly because the existence of a well-characterized critical threshold ρ_* represents a hallmark of genuine criticality.

We consider point-source ARW on $\mathbb{Z}$ with n particles initially placed at the origin. After stabilization, let

$$S_i = \{\text{site } i \text{ contains a sleeping particle}\}.$$

Our goal is to quantify the local density of sleepers in the *bulk* near the source. We state our main theorem.

Theorem 1.1. *Fix $\varepsilon > 0$. For the point-source ARW on $\mathbb{Z}$ with n particles at the origin, and for any site*

$$i \in [0, \frac{1}{2}(\rho_* - \epsilon)n] \cap \mathbb{Z},$$

we have

$$\mathbb{P}(S_i) = \rho_* + o(1) \quad \text{as } n \to \infty.$$

We compare two point-source ARWs with sources at 0 and 1. We first run an internal diffusion limited aggregation (IDLA) "flattening" phase. These moves are legal for ARW, so exposing them first does not change the final ARW outcome. Lemma 3.5 couples the two IDLAs so that, with high probability, the resulting configurations agree. We then evolve both systems using the same randomness. On the event that the flattened configurations match up to a shift, this coupling forces the single-site sleeper probabilities at neighboring sites to be nearly equal, hence the sleeper marginal is almost constant across short blocks in the bulk. Finally, Lemma 4.1 identifies the average sleeper probability over such a short block with the critical density ρ_*. Combining "near constancy on blocks" with the block-average identification yields Theorem 1.1.

2. Site-wise Construction of ARW

To establish the upper bound we work with a site-wise (Diaconis–Fulton) construction of the ARW: to every site we attach an infinite stack of instructions. Thanks to the abelian Property (Lemma 2.1), the order in which sites are toppled is irrelevant once the number of topplings at each site is fixed; this lets us reason locally about moves at specified sites.

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

Instruction stacks and their execution. Each $x\in\mathbb{Z}$ carries an infinite stack $(\mathsf{Instr}_x(k))_{k\in\mathbb{N}^+}$ of i.i.d. instructions, independent across x and k, with

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

Unused instructions at x are executed at rate $\mathbb{1}_{\omega(x)\neq\mathfrak{s}}\,|\omega(x)|(1+\lambda)$. A `Left` (resp. `Right`) instruction removes 1 from $\omega(x)$ and adds 1 to $\omega(x-1)$ (resp. $\omega(x+1)$). A `Sleep` instruction at a site with exactly one active particle changes $1\mapsto\mathfrak{s}$, while it has no effect if there are $n\geq 2$ active particles.

Topplings and odometer. Toppling a site x means executing its first unused instruction. For a finite sequence of sites $\alpha=(x_1,\dots,x_\ell)$ we topple in that order. The *odometer*

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

records how many times each site has been toppled; for a sequence α we write $U_\alpha(x)$ for the number of occurrences of x in α. If (ω,U) is the current state, the toppling map at x is

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

where δ_x is 1 at x and 0 elsewhere. The move Φ_x is *legal* when x is unstable. A sequence Φ_α is a *sequence of legal topplings* if each Φ_{x_i} is legal at the moment it is performed; it is *stabilizing* if, after Φ_α, no site is unstable.

The abelian property permits an interval-by-interval stabilization scheme. We state it in the form we use:

Lemma 2.1 (Abelian Property [16]). *Let α and β be legal toppling sequences for a configuration ω. If $U_\alpha(x) = U_\beta(x)$ for every $x \in \mathbb{Z}$, then the resulting configurations coincide:*

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

Particle-wise viewpoint. We also use an equivalent particle-wise construction: each particle carries its own i.i.d. instruction stack, independent across particles. At any time one may move any active particle by consuming its next unused instruction; by the abelian property, the choice of which active particle to move is immaterial. Whenever the particle-wise process terminates from a given initial configuration, the joint law of the final particle locations and odometer agrees with that of the site-wise construction.

3. Coupling two one-dimensional IDLA clusters

We begin with coupling statement for IDLA. Let $\mathcal{C}_n^{(s)}$ be the IDLA cluster obtained by releasing n particles from the site $s \in \{0,1\}$. Let $K_n^{(s)}$ denote the number of occupied sites strictly to the right of s in $\mathcal{C}_n^{(s)}$ after stabilization.

Observe that the two clusters $\mathcal{C}_n^{(0)}$ and $\mathcal{C}_n^{(1)}$ coincide after stabilizing n particles if and only if

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

Thus it suffices to construct a coupling of $(K_n^{(0)}, K_n^{(1)})$ 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

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

The next result Mittelstaedt identifies the law of $K_n^{(s)}$ with the Eulerian distribution of descents D_n of a uniformly random permutation.

Theorem 3.1 (Theorem 1 of [15]). *For each source $s \in \{0,1\}$,*

$$K_n^{(s)} \overset{d}{=} D_n, \quad \mathbb{P}[K_n^{(s)} = k] = \frac{\langle n,k \rangle}{n!}, \quad 0 \le k \le n-1,$$

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

Set

$$P_n(k) = \mathbb{P}[D_n = k] \quad \text{and} \quad Q_n(k) = P_n(k-1) \quad (k \in \mathbb{Z}),$$

so Q_n is the right shift of P_n. Note that

$$\mathbb{P}(\mathcal{C}_n^{(0)} = \mathcal{C}_n^{(1)}) = \sup \mathbb{P}(X = Y)$$

where the supremum is over all couplings of $X \sim P$ and $Y \sim Q_n$. The maximal value of $\mathbb{P}(X = Y)$ is $1 - \mathrm{TV}(P_n, Q_n)$, where

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

is total variation distance. Therefore we bound $\mathrm{TV}(P_n, Q_n)$. We use standard fact about Eulerian distribution. The Eulerian numbers are log-concave in k [17], hence the sequence $P_n(k)$ is unimodal in k. From this, a short telescoping argument gives an exact identity.

Lemma 3.2. *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 3.2 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 $F_n(x) = \mathbb{P}(D_n \le x)$ and let Φ_n be the normal distribution with

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

Theorem 3.3 (Theorem 1.1 of [18])**.**

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

Lemma 3.4.

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

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

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

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

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

□

Lemma 3.5. *There exist $c, \gamma > 0$ and a coupling of two instances of IDLA on $\mathbb{Z}$, one with particles arriving at site 0, and the other with particles arriving at site 1, such that for all n, the resulting stable configurations are identical with probability at least $1 - cn^{-\gamma}$.*

Proof. Note that

$$\mathbb{P}(\mathcal{C}_n^{(0)} = \mathcal{C}_n^{(1)}) = \sup \mathbb{P}(K_n^{(0)} = 1 + K_n^{(1)}) = 1 - \mathrm{TV}(P_n, Q_n).$$

By Lemma 3.2 and 3.4,

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

for some universal c, which proves the claim.

$\square$

4. Block averages near the source

Lemma 4.1. *For any $\epsilon, \alpha > 0$ and any $i \in [0, (1-\epsilon)n/2] \cap \mathbb{Z}$,*

$$\sum_{j=i-\lfloor n^\alpha \rfloor}^{i} \mathbb{P}(S_j) = (\rho_c + o(1))n^\alpha. \tag{1}$$

Proof. Let D be the number of particles that stabilize between $i - \lfloor n^\alpha \rfloor$ and i, inclusive of both endpoints. First note that by linearity of expectation the left hand side of (1) is the expected value of D.

To bound $\mathbb{E}(D)$ we first perform IDLA freezing particles at $i - \lfloor n^\alpha \rfloor - 1$ and $i + 1$. Let A_1 be the event that IDLA does not leave one active particle at every site between $i - \lfloor n^\alpha \rfloor$ and i. From Theorem 1 of [15] we can deduce that

$$\mathbb{P}(A_1) < Ce^{-cn^{.5}}. \tag{2}$$

Then we stabilize. We let a^* be the number of particles that moved from $i - \lfloor n^\alpha \rfloor - 1$ to $i - \lfloor n^\alpha \rfloor$ and we let b^* be the number of particles that moved from i to $i+1$.

Now we momentarily leave behind our stabilization and consider a new family of starting configurations based on the IDLA process. For any $a, b \geq 0$ we consider a stabilization of a new configuration formed by adding some particles to the result of the IDLA process. We add a particles at $i - \lfloor n^\alpha \rfloor$ and b particles at i to the IDLA configuration and stabilize with sinks at $i - \lfloor n^\alpha \rfloor - 1$ and $i + 1$. (Note that the particles that were at $i - \lfloor n^\alpha \rfloor$ and $i + 1$ after IDLA do nothing in this process.) We can couple all of these new stabilizations with particles added and sinks with the stabilization of our original activated random walk by using the same set of instructions (after the IDLA has been performed) at every site.

For any a and b the distribution after this process is the stationary distribution on the driven dissipative system by Theorem 2.1 of [13]. Fix any $\delta > 0$. Let $A_{2,a,b}$ be the event that the number of particles that stabilize between $i - \lfloor n^\alpha \rfloor$ and i is not in

$$((n^\alpha)\rho_c(1 - \delta/2), (n^\alpha)\rho_c(1 + \delta/2)).$$

Let A_3 be the event that

(1) A_1^C occurs
(2) $(\cup_{a,b \in [0,n^5]} A_{2,a,b})^C$ occurs and
(3) a^* and b^* are between 0 and n^5.

Notice by the abelian property that D is the same as the number of particles left between $i - \lfloor n^\alpha \rfloor$ and i when we added a^* particles to $i - \lfloor n^\alpha \rfloor$ and b^* particles to i (with sinks at $i - \lfloor n^\alpha \rfloor$ and $i + 1$). So if A_3 occurs we have

$$D \in \Bigg((n^\alpha)\rho_c(1 - \delta/2), (n^\alpha)\rho_c(1 + \delta/2) \Bigg).$$

By the same argument as Lemma 3.5 of [3] the probability that either a^* or b^* is greater than n^5 is exponentially small in n. By Propositions 8.5 and 8.6 of [8] for any a and b

$$\mathbb{P}(A_{2,a,b}^C)$$

is also exponentially small in n^α. Thus combining (2) with the two bounds above we have

$$\mathbb{P}(A_3) \geq 1 - Cn^{10}e^{-cn^\beta}$$

for some β, C and c.

Finally we note that $0 \leq D \leq n$. Now we are ready to bound $\mathbb{E}(D)$.

$$\mathbb{E}(D) \geq \mathbb{P}(A_3^C)(n^\alpha)\rho_c(1-\delta/2) \geq (n^\alpha)\rho_c(1-\delta).$$

The upper bound is

$$\mathbb{E}(D) \leq \mathbb{P}(A_3)(n^\alpha)\rho_c(1+\delta/2) + \mathbb{P}(A_3^C)n \leq (n^\alpha)\rho_c(1+\delta).$$

Combining these two equations above we get that

$$\mathbb{E}(D) = (n^\alpha)\rho_c(1+o(1))$$

as desired. □

5. Completing the Proof

Proof of Theorem 1.1. Consider two instances of point-source ARW on $\mathbb{Z}$, both with n particles, one having source 0 and the other source 1. We may evolve the two ARW instances by first performing IDLA until the stack is fully flat (no more than one particle per site) – these topplings are always legal. We couple the dynamics in this phase using the coupling given by Lemma 3.5. When the IDLA is complete, if the configurations are not identical, we declare failure; otherwise, we couple by performing exactly the same instructions for both systems. It follows, since both systems have the same marginal distribution up to a shift by one, that

$$|\mathbb{P}(S_j) - \mathbb{P}(S_{j+1})| \leq cn^{-\gamma}. \tag{3}$$

By the triangle inequality, for $j = 1, 2, \ldots$, and taking i as in the statement of the Theorem,

$$|\mathbb{P}(S_i) - \mathbb{P}(S_{i-j})| \leq (j)(cn^{-\gamma}). \tag{4}$$

Summing over $j \in [0, \lfloor n^\alpha \rfloor] \cap \mathbb{Z}$ and using 4,

$$\left| \sum_{j=i-\lfloor n^\alpha \rfloor}^{i} \mathbb{P}(S_j) - \lfloor n^\alpha \rfloor \mathbb{P}(S_i) \right| \leq \sum_{j=0}^{\lfloor n^\alpha \rfloor} (j)(cn^{-\gamma}) \leq c' n^{2\alpha-\gamma}. \tag{5}$$

Combining with Lemma 4.1 (for any $\alpha < \gamma$) and dividing by n^α, we obtain

$$|\mathbb{P}(S_i) - \rho_c| \leq c' n^{\alpha-\gamma} = o(1). \tag{6}$$

□

References

[1] P. Bak, C. Tang, and K. Wiesenfeld, *Self-organized criticality: An explanation of the* $1/f$ *noise*, Phys. Rev. Lett. **59** (1987), 381.

[2] Ofer Biham, Erel Milshtein, and Ofer Malcai, *Evidence for universality within the classes of deterministic and stochastic sandpile models*, Phys. Rev. E **63** (2001May), 061309.

[3] Madeline Brown, Christopher Hoffman, and Hyojeong Son, *Activated random walks on F with critical particle density*, 2024.

[4] Concetta Campailla and Nicolas Forien, *Stochastic sandpile model: exact sampling and complete graph*, 2025.

[5] Concetta Campailla, Nicolas Forien, and Lorenzo Taggi, *The critical density of the stochastic sandpile model*, 2024.

[6] Deepak Dhar, *The abelian sandpile and related models*, Physica A: Statistical Mechanics and its applications **263** (1999), no. 1-4, 4–25.

[7] Anne Fey, Lionel Levine, and David B. Wilson, *Approach to criticality in sandpiles*, Phys. Rev. E **82** (2010), 031121.

[8] Christopher Hoffman, Tobias Johnson, and Matthew Junge, *The density conjecture for activated random walk*, 2024.

[9] ———, *A proof of self-organized criticality in a sandpile*, 2024.

[10] Robert D. Hough and Hyojeong Son, *Cut-off for sandpiles on tiling graphs*, Ann. Probab. **49** (2021), no. 2, 671–731.

[11] ———, *The spectrum of the abelian sandpile model*, Math. Comp. **90** (2021), no. 327, 441–469.

[12] Hang-Hyun Jo and Hyeong-Chai Jeong, *Comment on "driving sandpiles to criticality and beyond"*, Phys. Rev. Lett. **105** (2010), 019601.

[13] Lionel Levine and Feng Liang, *Exact sampling and fast mixing of activated random walk*, arXiv preprint arXiv:2110.14008 (2021).

[14] Lionel Levine and Yuval Peres, *Laplacian growth, sandpiles and scaling limits*, arXiv preprint arXiv:1611.00411 (2016).

[15] Kiana Mittelstaedt, *A stochastic approach to eulerian numbers*, arXiv:1902.03195 [math.CO] (2019).

[16] Leonardo T Rolla, *Activated random walks on* $\mathbb{Z}^d$, Probability Surveys **17** (2020), 478–544.

[17] Richard P. Stanley, *Log-concave and unimodal sequences in algebra, combinatorics, and geometry a*, Annals of the New York Academy of Sciences **576** (1989).

[18] Alperen Y. Özdemir, *Martingales and descent statistics*, Advances in Applied Mathematics **140** (2022), 102395.

Department of Mathematics, University of Washington, Seattle, WA 98195
Email address: `hoffman@math.washington.edu`

Alfréd Rényi Institute of Mathematics, Budapest, Hungary
Email address: `jrichey@renyi.hu`

Department of Mathematics, University of Washington, Seattle, WA 98195
Email address: `hjson@math.washington.edu`