

Activated random walk on a cycle

Riddhipratim Basu^a, Shirshendu Ganguly^b, Christopher Hoffman^c and
Jacob Richey^c

^aInternational Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Bangalore, India. E-mail: rbasu@icts.res.in

^bDepartments of Statistics and Mathematics, UC Berkeley, CA, USA. E-mail: sganguly@berkeley.edu

^cDepartment of Mathematics, University of Washington, Seattle, USA. E-mail: hoffman@math.washington.edu; jfrichey@math.washington.edu

Received 21 February 2018; revised 17 May 2018; accepted 24 June 2018

Abstract. We consider Activated Random Walk (ARW), a particle system with mass conservation, on the cycle $\mathbb{Z}/n\mathbb{Z}$. One starts with a mass density $\mu > 0$ of initially active particles, each of which performs a simple symmetric random walk at rate one and falls asleep at rate $\lambda > 0$. Sleepy particles become active on coming in contact with other active particles. There have been several recent results concerning fixation/non-fixation of the ARW dynamics on infinite systems depending on the parameters μ and λ . On the finite graph $\mathbb{Z}/n\mathbb{Z}$, unless there are more than n particles, the process fixates (reaches an absorbing state) almost surely in finite time. In a first rigorous result for a finite system, establishing well known beliefs in the statistical physics literature, we show that the number of steps the process takes to fixate is linear in n (up to poly-logarithmic terms), when the density is sufficiently low compared to the sleep rate, and exponential in n when the sleep rate is sufficiently small compared to the density, reflecting the fixation/non-fixation phase transition in the corresponding infinite system as established in (*Invent. Math.* **188** (2012) 127–150).

Résumé. Nous considérons la marche aléatoire activée (*Activated Random Walk*, ARW), un système de particules avec conservation de masse sur le cycle $\mathbb{Z}/n\mathbb{Z}$. Partant d'un état initial avec une densité $\mu > 0$ de particules actives, chacune d'entre elles évolue selon une marche simple symétrique à taux 1, et s'endort à taux $\lambda > 0$. Les particules endormies deviennent actives lorsqu'elles entrent en contact avec d'autres particules actives. Plusieurs résultats récents se sont penchés sur la fixation ou la non-fixation de la dynamique en volume infini, en fonction des paramètres μ et λ . Sur le graphe fini $\mathbb{Z}/n\mathbb{Z}$, à moins qu'il y ait plus de n particules, le processus se fixe (en atteignant un état absorbant) presque sûrement en temps fini. Nous établissons un premier résultat rigoureux sur ces systèmes finis, confirmant des prédictions bien connues de la littérature de physique statistique, en montrant que le nombre d'étapes avant fixation est linéaire en n (à des termes poly-logarithmiques près) lorsque la densité est suffisamment petite par rapport au taux d'endormissement, et exponentielle en n lorsque le taux d'endormissement est suffisamment petit par rapport à la densité, ce qui reflète la transition de phase entre fixation et non-fixation établie dans (*Invent. Math.* **188** (2012) 127–150) pour le système infini.

MSC: 60K35; 82C20; 82C22

Keywords: Activated random walk; Diaconis–Fulton representation; Abelian property; Self-organized criticality

1. Introduction

Consider the following interacting particle system on a one dimensional lattice. Given a configuration of particles, initially all active, the dynamics, which conserves the particles, proceeds as follows. Each active particle independently does a simple symmetric random walk at rate one in continuous time and falls asleep at rate $\lambda > 0$. Each sleepy particle is awakened when an active particle occupies the same site. This model, known as Activated Random Walk (ARW), has attracted interest in non-equilibrium statistical mechanics as well as probability literature in recent years in connection with studying fixed energy sandpile models [5,7,8,22,28–30]. The motivation of studying this model is

two-fold. ARW can be regarded as a special case of driven diffusive epidemic process introduced by Spitzer in 1970s, and studied later in [13–16]. ARW was also introduced in the physics literature as a more mathematically tractable approximation of the Stochastic Sandpile Model (SSM), and is one of the paradigm examples of the widely studied phenomenon of self-organized criticality (SOC) [4,6,20,21].

ARW is believed to manifest self-organized criticality when run in a finite volume with carefully controlled driven diffusive dynamics. However, the rigorous study of ARW has so far been mostly restricted to the case of infinite volume limit where the counterpart of SOC is known as Absorbing state Phase Transition (APT) (although some recent results have called into question the exact relationship between these two notions [10,11,17]). Absorbing state Phase Transition was rigorously established for ARW on \mathbb{Z} a few years ago in the fundamental work of Rolla and Sidoravicius [22]. Let us briefly explain their result. Consider ARW started with initial configuration of particles coming from a product measure with density μ ; denote this process by $\text{ARW}(\mu, \lambda)$. One would expect that for a fixed λ , if μ is very small, then all the particles will eventually fall asleep, whereas for large μ the activity would go on forever. Indeed, in [22] it was shown that for each $\lambda > 0$, there exists $\mu_c(\lambda) \in [\frac{\lambda}{\lambda+1}, 1]$ such that for $\mu < \mu_c$ the process $\text{ARW}(\mu, \lambda)$ on \mathbb{Z} fixates (i.e., the total number of jumps at origin is finite) almost surely, and for $\mu > \mu_c$ the process remains almost surely active forever. Observe that it is easy to understand heuristically why $\mu_c \leq 1$. If $\mu > 1$, there are “more particles than sites” and hence not all particles can eventually fall asleep [1,22,23]. Complementing the results of [22], the first three authors showed in [2] that for any fixed $\mu > 0$, the process almost surely does not fixate if λ is sufficiently small; thus showing $\mu_c \rightarrow 0$ as $\lambda \rightarrow 0$, and answering a question from [8,22]. Subsequently, a statement similar to the latter was proven for transient Euclidean lattices in [25], which also analyzed ARW dynamics on transitive graphs where the random walk is ballistic. However, to rigorously establish the critical or near-critical behaviour in these models seems far out of reach of the current mathematical techniques.

The results in [2,22] are in the setting of ARW on the infinite lattice \mathbb{Z} . Indeed, there has been a flurry of recent mathematically rigorous results on ARW following the breakthrough work [22], but most of them have been in the context of Euclidean lattices or other infinite graphs [1,23–27]. From the point of view of understanding self-organized criticality, it is interesting to study this model on a finite lattice, with say a periodic boundary condition. On a finite graph, if the total number of particles is more than the number of vertices, then the process will continue for ever. If the total number of particles is at most the number of vertices, this is an absorbing Markov chain, so all the particles will almost surely fall asleep after a finite time. One would expect the absorbing state phase transition to be manifested in the finite process as a phase transition for the absorption time. For many finite systems of these type, it is generally believed that absorption time has three different scalings with the system size, polynomial (with different exponents) for the sub-critical and critical systems, and exponential for the super-critical system. Indeed a version of the above statement in the set up of [22] and [2] respectively are the main results of this paper. In physics literature there are many non-rigorous and numerical results about the critical and near-critical scaling of this and related quantities for SSM and its many variants (see, e.g., [19] and references therein). However, as with the infinite system, rigorous analysis of the critical scaling behaviour remains a challenging problem.

1.1. Main results

Consider an n -cycle $\mathbb{Z}/n\mathbb{Z}$ with nearest neighbour edges. Fix $\lambda > 0$ and $\mu \in (0, 1)$. Consider the initial configuration with independent $\text{Ber}(\mu)$ many particles at each site. (We will denote the product Bernoulli measure by \mathbb{P}^μ). Consider ARW started with this configuration with sleep rate λ ; denote this process by $\text{ARW}(\mu, \lambda)$. As mentioned before, this is an absorbing Markov chain and hence the process reaches the absorbing state of all sleepy particles (the set of all such configuration will be henceforth called the cemetery set and written Δ) after a finite time almost surely. Let $T_n(\mu, \lambda)$ denote the total number of attempts by any active particle to either jump or to try to sleep. Note that the continuous time ARW dynamics can be coupled naturally to the following discrete time dynamics: at every positive integer time, pick an active particle uniformly at random. With probability $\frac{1}{2(1+\lambda)}$ each, the particle jumps to one of the neighbouring sites, and with probability $\frac{\lambda}{1+\lambda}$ it tries to fall asleep. It is easy to see that under the natural coupling, $T_n(\mu, \lambda)$ is the absorption time of the latter dynamics.

Our first result shows that if the particle density μ is sufficiently small compared to λ then $T_n(\mu, \lambda)$ is linear up to poly-logarithmic correction factors.

Theorem 1. Consider ARW(μ, λ) on $\mathbb{Z}/n\mathbb{Z}$. For any $\lambda > 0$ and $\mu < \frac{\lambda}{1+\lambda}$, there exist positive constants C_0, b depending on μ and λ , such that for all large enough n ,

$$\mathbb{P}(T_n(\mu, \lambda) > C_0 n \log^2 n) \leq \frac{1}{n^b}.$$

Observe that $\mu < \frac{\lambda}{1+\lambda}$ is precisely the regime in which [22] showed fixation on \mathbb{Z} . It is also easy to observe that the bound is tight up to the polylogarithmic factor. To see this, observe that the expected total number of jumps a particle takes before trying to fall asleep is $\frac{1+\lambda}{\lambda}$. Thus, if the number of particles is linear in n , the total number of jumps is also at least linear in n .

In the heavily super-critical regime, that is, when λ is sufficiently small compared to μ , we have the following complementary result (This result should be compared to [2, Theorem 1].).

Theorem 2. For any $0 < \mu < 1$, there exists $\lambda_0 > 0$ and $c > 0$ such that for any $\lambda < \lambda_0$, and all large enough n ,

$$\mathbb{P}(T_n(\mu, \lambda) < e^{cn}) < e^{-cn}.$$

Theorem 1 relies heavily on the uniformity of the locations of the particles in the initial configuration (note however, that the arguments in this paper do not depend on the specific nature of the Bernoulli distribution of the initial configuration). In fact, it can be shown that $T_n(\mu, \lambda)$ is at least of order n^3 when all the particles start at the origin. An exponential upper bound for $T_n(\mu, \lambda)$ is also relatively easy to establish. See Remarks 3.13 and 4.6 for further elaboration.

We list below the new contributions in this paper and relations to existing results: Although the linear to exponential phase transition for absorption time is widely expected in the statistical physics literature, to the best of our knowledge this is the first rigorous result establishing such a transition for some variant of fixed energy sandpile models. We rely crucially on the recent progress [2,22] in understanding ARW on the infinite line. However one needs certain new ideas to deal with the finite case. Following the argument of [22], the main obstacle in showing fast fixation for low particle density is the wrapping around issue, that is, to make sure the particles do not wrap around the cycle, and wake up already settled particles. We get around this by a block argument and a two-sided variant of the stabilizing algorithm in [22]. In the process of attaining the quantitatively optimal result Theorem 1, we encounter a particle system similar to internal erosion (see [17]).

For the slow fixation part, that is, Theorem 2, we first recall that the argument from [2] essentially tells us that when we stabilize a certain density of particles on $\mathbb{Z}/n\mathbb{Z}$ until they hit 0, for small enough sleep rate only a small fraction of the particles fall asleep. The key observation in this paper is that the above step can be applied iteratively for exponentially many rounds. A naive application of the iteration scheme only allows the number of steps to be logarithmic in n since a constant fraction of particles fall asleep in every round. However the finiteness of the environment allows us to recycle particles which fell asleep in earlier rounds once they get woken up in later rounds. To this end we strengthen the argument in [2] by showing that with exponentially small failure probability all the particles that were asleep at the start of the round get woken before the end of the round and not too many particles fall asleep. Thus we can sustain the process for exponentially many steps.

In the next section we elaborate on the above ideas further.

1.2. Sketch of the proofs

A crucial property of many interacting particle systems that serve as models of distributed networks is the Abelian property. Informally it means that the final outcome of a certain probabilistic experiment does not depend on the order in which operations at different sites are performed.

In the context of ARW, one exploits the Abelian property via the Diaconis–Fulton representation, which is roughly the following (see Section 2.1 for a more formal description). At every site in $\mathbb{Z}/n\mathbb{Z}$, we have a sequence of i.i.d. instructions (referred to as the “stack of instructions”) to either jump to one of the two neighbours, with probability $\frac{1}{2(1+\lambda)}$ each, or to fall asleep, with probability $\frac{\lambda}{1+\lambda}$. Given these stacks, one way to run the process is: as long as there is some active particle, pick an arbitrary site $x \in \mathbb{Z}/n\mathbb{Z}$ with at least one active particle, and use the first unused

instruction from the stack to topple the site. That is, if the instruction is a jump instruction, then the particle jumps to a neighbouring site accordingly, and otherwise tries to fall asleep.

The Abelian property then states that the final configuration of particles after every particle has fallen asleep does not depend on the order in which the sites were toppled. Thus in this language $T_n(\mu, \lambda)$ is the total number of instructions across all the stacks used until the end of the toppling process.

We also rely on the following monotonicity property of the ARW dynamics: given a set of stacks of instructions, while toppling sites, if we ignore any sleep instruction, then the total number of topplings required to reach Δ can only increase. Here ‘ignoring a sleep instruction’ means the configuration does not turn an active particle to a sleepy particle, even though the instruction is a sleep instruction. This monotonicity is certainly heuristically plausible and, indeed, is a well-known consequence of the Abelian property. (See Lemma 2.3 and the discussion preceding it for formal definitions.)

1.2.1. *Sketch of the proof of Theorem 1*

Given the above two properties, to prove Theorem 1 we will provide a toppling scheme which will end with a configuration in the cemetery set Δ . Our toppling procedure will ignore certain sleep instructions and hence by the monotonicity property, the total number of instructions used in the actual process in the Diaconis–Fulton representation is upper bounded by the number of instructions used in our scheme. The basic idea is to break the cycle $\mathbb{Z}/n\mathbb{Z}$ in to sub-intervals I_1, I_2, \dots of size $c_0 \log n$ for some constant c_0 . Our toppling scheme is then a combination of toppling schemes, one for each of the sub-intervals. The toppling scheme for I_i is designed to stabilize particles inside I_i for $i = 1, 2, \dots, \frac{n}{c_0 \log n}$. The toppling scheme in each of the intervals is a variant of the trap-setting procedure appearing in [22]. We prove that with very small failure probability (exponential in the size of the interval) the toppling procedure in the interval succeeds to stabilize everything. A union bound over all the intervals show that with high probability the procedure succeeds simultaneously for all the intervals, and hence stabilizes the system. Recall that the failure probability for each interval is exponential in the size of the interval which forces us to choose the size of the intervals to be logarithmic (in n) as otherwise the union bound over polynomially (in n) many such intervals will fail; this also explains the logarithmic correction term in the statement of Theorem 1.

We now briefly describe our toppling scheme (see Figure 1). It consists of broadly two parts.

- (1) Phase 1: Given the initial configuration, the first step is to gather particles which are initially located uniformly over $\mathbb{Z}/n\mathbb{Z}$, to a set of points we call **Sources**. We will take this set to be $\{\frac{c_0}{2} \log n, \frac{3c_0}{2} \log n, \dots\}$ for some carefully chosen value of c_0 , depending on the parameters μ and λ . Thus we first ignore all the sleep instructions and allow the particles to do independent random walks till they hit an element of **Sources**. Large deviation estimates imply that with high probability the number of particles at each **Source** at the end of this process is roughly $c_0 \mu \log n$. Recall that we are using the monotonicity property mentioned above, and hence we can ignore certain sleep instructions.
- (2) Phase 2: The intervals I_1, I_2, \dots inside which we run our toppling scheme are of size $c_0 \log n$, and centred at the sources. The proof proceeds by showing that there is a toppling procedure which carefully ignores certain sleep instructions allowing the particles to fall asleep only at certain well chosen ‘traps’ which prevents interaction with other particles. The remainder of the proof then shows that the above scheme succeeds with high probability.

Even though the trap setting scheme is inspired from [22], the argument in the latter was particularly tailored to the ARW process on the infinite line and does not directly work on the cycle: this is because, on the cycle, each particle has a chance of order $\frac{1}{n}$ to wrap around the cycle before falling into the trap and thereby waking up the already asleep particles. To circumvent this we introduce a two sided version, which, relying on estimates for random walk on the interval, can be shown to work in our setting.

1.2.2. *Sketch of the proof of Theorem 2*

For the proof of Theorem 2 we rely on the non-fixation result from [2]. The technical core of that paper (see the proof of [2, Lemma 18]) was to establish the following non-fixation phenomenon. Consider a sufficiently large interval $[0, r]$ with at least μr many active particles, and stabilize the particle system inside the interior of the interval $[0, r]$, that is, the particles are stopped upon hitting $\{0, r\}$. Then, given $\varepsilon > 0$, for any mass density μ , and for all small enough λ , at the end of the stabilization procedure, the number of particles accumulating at $\{0, r\}$ is at least $(1 - \varepsilon)$ fraction of the total number, with failure probability exponentially small in r . This was used in [2] to show infinite activity on the




Fig. 1. (a) The toppling scheme for fixation in the sub-critical regime: in the first step of the stabilization scheme, we ignore sleep instructions to get every particle to a nearby source vertex. Particles at a particular source are then stabilized inside an interval of length $c_0 \log n$ using a trap-setting procedure. (b) The toppling scheme for non-fixation in the supercritical regime: 0 and $n/2$ are the north and south ‘poles’ of the cycle. We run several rounds of the following stabilization loop where we first try to topple all particles located away from 0 or $n/2$ until they fall asleep or hit $\{0, n/2\}$. Then we try to topple particles starting from 0 until they hit $n/2$ or fall asleep, and afterwards do the same at $n/2$, namely topple all particles starting from $n/2$ until they hit 0 or fall asleep. These three actions are repeated until all particles are asleep: our proof shows that this loop can be sustained for exponentially many steps (in n) with high probability.

line, by considering a growing sequence of r and then using the above statement to show that particles from arbitrary far away would hit the origin, thus implying non-fixation. For a finite system, we cannot rely on an argument which uses particles arbitrarily far away. Instead we use the following ‘recycling particles’ approach: represent the cycle $\mathbb{Z}/n\mathbb{Z}$ as the interval $[-n/2, n/2]$ with endpoints identified. We run the following rounds of particle stabilization:

- (1) We first treat the ‘poles’ 0 and $n/2$ as boundary points, and topple particles not at those sites until they fall asleep or land at one of the sites 0 or $n/2$. By the arguments in [2], a constant fraction of all particles will make it to either 0 or $n/2$ with exponentially high probability.
- (2) We topple particles that ended up at 0 in the previous stage until they fall asleep or hit $n/2$. The particles that did not start at 0 do not move, but they can be woken up by other active particles during this stage. Using ideas from [2, Lemma 18], we show that most particles will make it to the boundary before falling asleep with exponentially high probability.
- (3) Just as the previous step, we only topple particles that started at $n/2$ at the end of the previous stage, running the dynamics until all such particles are asleep or at 0.

Since all these steps keep most particles awake with exponentially high probability, we can run the loop exponentially many times. Moreover, since each loop takes at least one stack instruction to run, the Abelian property implies that $T_n(\mu, \lambda)$ is at least exponential in n .

2. Abelian property for activated random walk

For brevity, we will denote $\mathbb{Z}/n\mathbb{Z}$ by \mathcal{C}_n . We follow [22] in formally describing the set up of ARW. To avoid unnecessary notational overhead we describe the bare minimum of the formalism necessary. We always work with ARW on \mathcal{C}_n for some fixed but large n , and the notation is adapted accordingly. In particular, for the remainder of this section, addition and subtraction will be considered modulo n whenever appropriate and we shall not mention that explicitly every time.

For any time $t \geq 0$ and location $x \in \mathcal{C}_n$, $\eta_t(x)$ denotes the state of the system at location x at time t . We write $\eta_t(x) = \rho$ if there is one sleepy particle at $x \in \mathcal{C}_n$ at time t . If there is not a sleepy particle present we let $\eta_t(x) \in \mathbb{N}$ denote the number of particles at $x \in \mathcal{C}_n$ at time t . Then $\eta_t = \{\eta_t(x)\}_{x \in \mathcal{C}_n}$ denotes the state of the system at time t .

We shall use two operations (called topplings) on the space of configurations. For $x \in \mathbb{Z}$ and $y = x \pm 1$, let $\tau_{x,y}(\eta)$ denote the configuration obtained by moving one of the active particles from x to y . This operation will be called **illegal** (for the configuration) if there are no active particles at x and the system remains unchanged. Let $\tau_{x,\rho}(\eta)$

denote the configuration obtained from η by making the solitary particle at x fall asleep. Moreover if x has more than one active particle, the sleep instruction has no effect, so $\tau_{x,\rho}(\eta) = \eta$. Again if there are no particles at x , this instruction is called illegal and the system is not changed.

Now we can formally define ARW as a finite state space continuous time Markov chain with transitions $\eta \rightarrow \tau_{x,y}\eta$ at rate $A(\eta_t(x))\frac{1}{2}\mathbf{1}_{y=x\pm 1}$, and $\eta \rightarrow \tau_{x,\rho}\eta$ at rate $\lambda A(\eta_t(x))$ where $A(\eta(x))$ denote the number of active particles at site x in configuration η . Let \mathbb{P}^ν denote the law of the process started from an initial configuration distributed according to ν .

2.1. Diaconis–Fulton representation

We will now describe the Diaconis–Fulton representation of the ARW dynamics which will be convenient for our purposes. For an extensive discussion of the Diaconis–Fulton representation of ARW dynamics, the Abelian property and its consequences, see [22]. For completeness we recall the relevant results from [22], suitably adapted to the setting of a finite cycle.

The Diaconis–Fulton representation [3,9] maps the ARW process to sequence of instructions attached to the sites. The advantage of this representation is the Abelian property, which allows one to disregard the order in which different steps were performed in certain settings. We start by introducing a series of notations. Recall the operations $\tau_{x,y}$ and $\tau_{x,\rho}$ from above. Now consider the following array of random variables:

$$\mathbf{\xi} = \begin{matrix} \dots & \xi_{(-2,1)} & \xi_{(-1,1)} & \xi_{(0,1)} & \xi_{(1,1)} & \xi_{(2,1)} & \dots \\ \dots & \xi_{(-2,2)} & \xi_{(-1,2)} & \xi_{(0,2)} & \xi_{(1,2)} & \xi_{(2,2)} & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \end{matrix}, \tag{2.1}$$

where $\xi_{(x,j)}$ are independent for any $x \in \mathcal{C}_n$ and $j \in \mathbb{N}$ and moreover,

$$\xi_{(x,j)} = \begin{cases} \tau_{x,x-1} & \text{with probability } \frac{1}{2(\lambda+1)}, \\ \tau_{x,x+1} & \text{with probability } \frac{1}{2(\lambda+1)}, \\ \tau_{x,\rho} & \text{with probability } \frac{\lambda}{\lambda+1}. \end{cases} \tag{2.2}$$

We will now show that using these instructions one can define a discrete time version of the ARW process. In fact we can define many such versions. But they will all have the same configuration when it is finally stabilized and the same set of instructions that have been implemented.

We call the $\xi_{(x,j)}$ ’s instructions at the site x and the underlying product measure \mathcal{P} .

Given a configuration η at each discrete time step t , one can choose (arbitrarily) an unstable site x and use the first unused element from the stack $\xi_{(x,\cdot)}$ and use it to perform the transition to a configuration η' at time step $(t + 1)$. As mentioned in Section 1.2, we call such an operation “toppling” at site x . We keep track of the number of topplings at every site. Let h be the configuration after applying $h(x)$ many topplings at each site $x \in \mathcal{C}_n$. Let us denote

$$h := (h(x) : x \in \mathbb{Z}/n\mathbb{Z}) \tag{2.3}$$

which we will call the odometer function. Let $\Phi_x(\eta)$ denote the configuration obtained by toppling the site x next, that is, we apply the instruction $\xi_{x,h(x)+1}$, and also increase h at x by one (h at other sites does not change). We say Φ_x is legal for η if x is unstable in η . For any sequence $\alpha = (x_1, x_2, \dots, x_k)$ we define the sequence of topplings at x_1 , followed by x_2 and so on through until x_k by Φ_α , that is, $\Phi_\alpha = \Phi_{x_k} \dots \Phi_{x_1}$. We now say that α is a **legal sequence** for initial configuration η if Φ_{x_i} is legal for $\Phi_{x_{i-1}} \dots \Phi_{x_1}(\eta)$ for all $i = 1, \dots, k$. We abuse notation a little to denote by h_α the odometer function after performing the sequence of toppling given by α , that is, for any $x \in \mathcal{C}_n$,

$$h_\alpha(x) = \sum_{i=1}^k \mathbf{1}(\alpha_i = x). \tag{2.4}$$

Given the above preparation we can now formally state the Abelian property, which says that given two sequence of legal topplings that result in the same odometer function (see (2.4)), the final configuration is the same in both the cases, that is, the order in which topplings are performed does not matter.

Lemma 2.1 (Abelian Property [22, Lemma 2]). *Given any two legal sequence of topplings α and α' such that $h_\alpha = h_{\alpha'}$, then*

$$\Phi_\alpha(\eta) = \Phi_{\alpha'}(\eta).$$

The next lemma is a consequence of the Abelian property. It shows that any legal sequence of topplings must occur in any stabilizing sequence (i.e., a legal sequence which leads to all stable sites).

Lemma 2.2 (Least Action Principle [22, Lemma 1]). *Let α, α' be two legal sequences of topplings such α stabilizes η , then $h_{\alpha'} \leq h_\alpha$, that is, all the topplings in α' are also needed in α .*

Observe that Lemma 2.2 immediately implies that any two stabilizing sequence must lead to the same odometer function, and in turn by Lemma 2.1 this implies that the final configuration after stabilization is also independent of the stabilizing sequence. This will imply that for any stabilizing sequence α , for an initial configuration of product $\text{Ber}(\mu)$ many particles, we have

$$T_n(\mu, \lambda) = \sum_{x \in \mathcal{C}_n} h_\alpha(x).$$

This will be the fundamental tool used in our proofs. Finally we need another lemma to compare the stabilizing sequences which formalizes the intuitively plausible statement: for a stabilizing sequence where we ignore some of the sleep instructions, the total number of jumps is larger than if we hadn't ignored those sleep instructions. Formally we need to introduce a new notation to define precisely the meaning of ignoring a sleep instruction. Recall the stack of instructions \mathcal{I} from (2.1) and the action of the instructions $\tau_{x,\rho}$ and $\tau_{x,x\pm 1}$ on the particle configuration η . Let us introduce a null instruction \mathfrak{n} which acts on a configuration to create no change, that is, $\mathfrak{n}\eta = \eta$ for all η . Now given a stack of instructions, $\mathcal{I} = \{\xi_{(\cdot,\cdot)}\}$, as in (2.1), let $\mathcal{I}' = \{\xi'_{(\cdot,\cdot)}\}$, be another stack, with the property that for each (x, j) either, $\xi_{(x,j)} = \xi'_{(x,j)}$, or $\xi_{(x,j)} = \tau_{x,\rho}$, and $\xi'_{(x,j)} = \mathfrak{n}$. Thus informally \mathcal{I}' is any set of toppling instructions obtained from \mathcal{I} by ignoring certain sleep instructions.

Lemma 2.3 ([22, Lemma 5]). *Given any $\mathcal{I}, \mathcal{I}'$ as above, and any initial configuration η , let α and α' be two legal toppling sequences stabilizing η , using instructions from \mathcal{I} and \mathcal{I}' , respectively. Let $h_\alpha(\cdot)$ and $h_{\alpha'}(\cdot)$ be the respective odometer functions. Then $h_{\alpha'}(x) \geq h_\alpha(x)$ for every $x \in \mathcal{C}_n$.*

Often our argument will be based on running the ARW dynamics on certain sub-intervals of \mathcal{C}_n , and hence to be completely formal one needs to introduce stacks corresponding to such intervals. However to avoid introducing additional notation, we will identify the latter in the natural way with the corresponding subset of stacks of \mathcal{I} .

3. Fast fixation under low density

We prove Theorem 1 in this section. By the discussion at the end of last section and Lemma 2.3, we shall provide an algorithm for toppling which will ignore some sleep instruction and which stabilizes all sites in \mathcal{C}_n within $O(n \log^2 n)$ many topplings with high probability. (Note that the probability here is over both the random initial configuration and also the random stack of instructions.) We shall formalize the sketch provided in Section 1.2 to build the toppling procedure.

Let η be a initial configuration of particles on \mathcal{C}_n distributed according to law \mathbb{P}^μ . Also fix a realization \mathcal{I} of the stack of instructions. For the remainder of this section, we shall always talk about toppling the configuration η sequentially using instructions from \mathcal{I} .

3.1. Random walk estimates

As explained before, the first phase will be to topple any unstable site that is not a **Source** and ignore any sleep instructions encountered in the process. So at the end of this phase of the toppling all the particles will be herded at

the **Source** vertices. Let $\eta^{(1)}$ denote the configuration at the end of this phase, which is supported on **Source** vertices. Because we ignore the sleep instructions, this procedure is the same as letting all the particles in η be independent simple symmetric random walks stopped at hitting one of the sources. Thus we shall need a couple of basic random walk estimates to estimate the distribution $\eta^{(1)}$, as well as the total number of jumps to reach that configuration.

Recall that the i th source is located at the vertex $z_i := (i - \frac{1}{2})c_0 \log n$. Let $\eta_i^{(1)}$ denote the number of particles $\eta^{(1)}$ has in this vertex. Also recall that to reduce notation we assumed that n is an integer multiple of $c_0 \log n$, which itself is an even integer. Let $K = \frac{n}{c_0 \log n}$ be the number of sources. In the general case we can take all the intervals to be $\lfloor c_0 \log n \rfloor$ except possibly one which has length between $\lfloor c_0 \log n \rfloor$ and $2\lfloor c_0 \log n \rfloor$.

The following lemma is our first random walk estimate.

Lemma 3.1. *For each $\varepsilon > 0$, there exists $a > 0$ such that for all large enough c_0 and n ,*

$$\mathbb{P}\left(\sup_{1 \leq i \leq K} |\eta_i^{(1)} - \mu c_0 \log n| \geq \varepsilon c_0 \log n\right) \leq e^{-ac_0 \log n}.$$

Proof. For this proof we shall forget about \mathcal{I} and the toppling procedure, and treat $\eta^{(1)}$ as the configuration obtained from letting all the particles of η perform independent simple symmetric random walks stopped at hitting any of the source vertices (thus particles initially located at source vertices do not move at all). We first recall a standard concentration inequality for sums of independent but not necessarily identically distributed Bernoulli variables that we will use later in the proof. Let X_1, \dots, X_k be independent Bernoulli variables with means p_1, \dots, p_k respectively. Let $v = p_1 + \dots + p_k$. Then

$$\mathbb{P}\left(\left|\sum_i^k X_i - v\right| \geq \delta v\right) \leq e^{-\frac{\delta^2 v^2}{k}}. \tag{3.1}$$

Let us consider the first source z_1 . Clearly, any particle that ended up at z_1 in $\eta^{(1)}$ must have been located at some $j \in V_1 = (-\frac{c_0}{2} \log n, \frac{3c_0}{2} \log n)$. Let Z_j denote the indicator of the event that there was a particle at j in η and that ended up at z_1 in $\eta^{(1)}$. Clearly

$$\eta_i^{(1)} = \sum_{j \in V_1} Z_j.$$

Observe that a standard Gambler's ruin calculation yields that the probability that a random walk started at $j \in V_1$ would reach z_1 before reaching either z_0 or z_2 is $g(j) = 1 - \frac{|j - \frac{c_0}{2} \log n|}{c_0 \log n}$. It follows that Z_j 's are independent Bernoulli variables with mean $\mu g(j)$. Observe that

$$\left|\sum_{j \in V_1} g(j) - c_0 \log n\right| \leq 1,$$

and hence using (3.1) we get that for each $\varepsilon > 0$, and for all n sufficiently large

$$\mathbb{P}\left(|\eta_1^{(1)} - \mu c_0 \log n| \geq \varepsilon c_0 \log n\right) \leq e^{-2ac_0 \log n} \tag{3.2}$$

for some a depending on μ, ε (but not on c_0). By the rotational symmetry of \mathcal{C}_n and of the law of the initial configuration \mathbb{P}^μ , we have the same bound for all $\eta_i^{(1)}$ for all $1 \leq i \leq K$. Now since the number of source vertices is less than n , using (3.2) and a union bound over all source vertices, we get

$$\mathbb{P}\left(\sup_{1 \leq i \leq K} |\eta_i^{(1)} - \mu c_0 \log n| \geq \varepsilon c_0 \log n\right) \leq n e^{-2ac_0 \log n} \leq e^{-ac_0 \log n},$$

where we have chosen c_0 sufficiently large (depending on μ, ε) so that the last inequality holds. This completes the proof of the lemma. \square

Recall the basic setting of toppling sites using instructions from stack \mathcal{J} . The next lemma will prove that the total number of instructions explored until the end of phase one is at most order $n \log^2 n$ with high probability.

Lemma 3.2. *Let $T^{(1)}(\mu, \lambda)$ denote the total number of instructions that have been explored until the end of phase one. Then there exists $C_2, \theta > 0$ such that*

$$\mathbb{P}(T^{(1)}(\mu, \lambda) > C_2 n \log^2 n) \leq n^{-\theta}.$$

For this step we shall need a concentration result for sums of geometric random variables. Although the result we need at this step is pretty standard, we shall need a more complicated variant later on, and we also need a concentration for some of exponential random variables for a later part of the argument. For convenience we quote, at this point, the following result from [12] which shall cover all our needs (see Theorems 2.1 and 5.1 there).

Lemma 3.3. *The following concentration results hold:*

- (i) *Fix $p \in (0, 1)$, and let Y_1, Y_2, \dots be i.i.d. geometric random variables with parameter p , so $\mathbb{E}Y_1 = 1/p$. Then for any $\delta > 0$ and any $M \in \mathbb{N}$,*

$$\mathbb{P}\left(\left|\sum_{i=1}^M Y_i - M/p\right| \geq \delta \frac{M}{p}\right) \leq 2 \exp\left(-(\delta - \log(1 + \delta)) \frac{M}{p}\right).$$

- (ii) *Suppose Z_1, Z_2, \dots are independent exponential random variables with means a_1, a_2, \dots , and set $a_* = \inf_i a_i$, $\kappa = \sum_{1 \leq i \leq M} a_i$. Then for any $\delta > 0$ and $M \in \mathbb{N}$,*

$$\mathbb{P}\left(\left|\sum_{i=1}^M Z_i - \kappa\right| \geq \delta \kappa\right) \leq 2 \exp(-a_* \kappa (\delta - \log(1 + \delta))).$$

For the proof of Lemma 3.2 we need the following result.

Lemma 3.4. *Let n be an integer multiple of r and consider fixed locations at distance r on C_n (without loss of generality take them to be multiples of r). Let k independent identically distributed lazy symmetric random walks started from arbitrary locations on C_n and stopped on hitting the nearest integer multiple of r . Let T_i be the total number of steps taken by the i th walk (including the lazy steps). Then there exists $C > 0$ such that*

$$\mathbb{P}\left(\sum_{i=1}^k T_i \geq Cr^2k\right) \leq e^{-Ck},$$

where C depends on the laziness parameter (i.e., the probability of not jumping).

Proof. Standard simple random walk estimates show that if τ is the hitting time of $\{0, r\}$ for a lazy simple random walk on \mathbb{Z} (with laziness p), then for each $x \in \llbracket 0, r \rrbracket$, we have $\mathbb{P}_x(\tau < Cr^2) \geq \frac{1}{2}$ where \mathbb{P}_x , denotes the probability measure for the random walk started at x and C depends only on the laziness parameter. It follows then that for any arbitrary location of the k particles, each T_i is stochastically dominated by r^2G where G is a geometric random variable with mean 2. The statement now follows from part (i) of Lemma 3.3. □

We are now ready to prove Lemma 3.2.

Proof of Lemma 3.2. First notice that the total number of instructions used by the particles ignoring the sleep instructions is the same as the total number of steps taken when the particles do independent lazy symmetric random walks on C_n , where the laziness (probability of not jumping) is $\frac{\lambda}{1+\lambda}$ (exactly the probability that an instruction is a sleep instruction). Now as an easy consequence of (3.1), for any $\varepsilon > 0$, the total number of particles in η is in $(\mu - \varepsilon, \mu + \varepsilon)n$

with probability at least $1 - e^{-cn}$ for some $c = c(\varepsilon, \mu) > 0$. We can therefore condition on the number of particles being $m \in (\mu - \varepsilon, \mu + \varepsilon)n$. Let T^* denote the total number of steps taken by these particles until the end of phase one. Using Lemma 3.4 with $r = c_0 \log n$, we get that $\mathbb{P}(T^* > Cn \log^2 n) \leq e^{-cn}$. \square

3.2. Phase two: Stabilizing from the sources

We describe the second phase of our toppling scheme now. Recall that we start with the configuration $\eta^{(1)}$ that is supported on the **Source** vertices. Throughout the section we shall assume that $\eta^{(1)}$ satisfies the high probability event described in Lemma 3.1, where ε will be chosen sufficiently small depending on μ and λ later. Recall the intervals $I_i = [(i - 1)c_0 \log n, ic_0 \log n]$ for $1 \leq i \leq K$. Observe that the i th source z_i is the midpoint of the interval I_i .

As mentioned before, for each $1 \leq i \leq K$, we start stabilizing particles at z_i sequentially, in any arbitrary manner of toppling, until one of the particles hit the boundary of I_i , in which case we term the process a failure. On the contrary, we denote by \mathcal{S}_i the event that all the $\eta_i^{(1)}$ particles all fall asleep before hitting the boundary of I_i : call this event **Success at source** z_i . Clearly on the event that \mathcal{S}_i occurs for all i , the system stabilizes. The main step is to show that \mathcal{S}_i occurs with high enough probability so that one can take a union bound over all intervals I_i . Because of the underlying symmetry, we state the following result for a generic interval $[-\frac{r}{2}, \frac{r}{2}]$, where we assume r is even to avoid rounding issues.

Proposition 3.5. *Consider ARW, with sleep rate λ , on $[-\frac{r}{2}, \frac{r}{2}]$, started with m particles at the origin. Let $\varepsilon > 0$ be such that $\mu + 2\varepsilon < \frac{\lambda}{1+\lambda}$, and let $m \leq (\mu + \varepsilon)r$. Let \mathcal{S} denote the event that all the particles fall asleep before any particle hits $\{-\frac{r}{2}, \frac{r}{2}\}$. Then there exists $c > 0$, such that for all r sufficiently large we have $\mathbb{P}(\mathcal{S}^c) \leq e^{-cr}$.*

Proposition 3.5 is the most technically complicated result that goes into the proof of Theorem 1, and the proof will be spread over the next two subsections. Before delving into this proof, we want to complete the remaining steps in the argument proving Theorem 1. First we need to prove the following easy lemma.

Lemma 3.6. *Consider ARW with sleep rate λ , started from $\eta^{(1)}$. Let $T^{(2)} = T^{(2)}(\mu, \lambda)$ denote the total number of steps taken by all the particles until stabilization. Then there exists $C_3 > 0$, and $\theta > 0$, such that on the event $\bigcap \mathcal{S}_i$, we have*

$$\mathbb{P}(T^{(2)} \geq C_3 n \log^2 n) \leq n^{-\theta}$$

for all n sufficiently large.

Proof. Observe that arguing as in the proof of Lemma 3.4, taking $r = c_0 \log n$, on \mathcal{S}_i , the number of steps taken by each particle started from z_i is dominated by $Cc_0^2 \log^2 n G$ where G is a geometric random variable with mean 2 and C depends on λ . The rest of the proof is identical to that of Lemma 3.4 and its application in the proof of Lemma 3.2. We skip the details. \square

We can now complete the proof of Theorem 1.

Proof of Theorem 1. Fix $\mu < \frac{\lambda}{1+\lambda}$ and recall our two phase stabilization procedure. Recall $T^{(1)}$ from Lemma 3.2 and $T^{(2)}$ from Lemma 3.6, and the stack of instructions \mathcal{I} from (2.1). Note that our toppling scheme produces a stack of instructions \mathcal{I}^* obtained from \mathcal{I} , where the sleep instructions which are ignored by our toppling scheme are replaced by n instructions (see the definitions before Lemma 2.3). Moreover, given the stack \mathcal{I}^* , by the Abelian property (Lemma 2.1), on the event $\bigcap \mathcal{S}_i$, the total number of instructions needed to stabilize is $T^{(1)} + T^{(2)}$, since our toppling scheme uses exactly those many instructions from \mathcal{I}^* . Finally, using Lemma 2.3, the total number of instructions used to stabilize, for the stack \mathcal{I} , is upper bounded by $T^{(1)} + T^{(2)}$. Thus it will suffice to show that

$$\mathbb{P}(T^{(1)} + T^{(2)} > C_0 n \log^2 n) \leq n^{-b}. \tag{3.3}$$

We first fix $\varepsilon > 0$ so that the hypothesis of Proposition 3.5 is satisfied. Then fix c_0, C_2, a so that the conclusions of Lemma 3.1 (with the same ε) and Lemma 3.2 hold. Thus, the event

$$\mathcal{A} = \left\{ \sup_{1 \leq i \leq K} |\eta_i^{(1)} - \mu c_0 \log n| \leq \varepsilon c_0 \log n \right\} \cap \{T^{(1)}(\mu, \lambda) < C_2 n \log^2 n\}$$

occurs with probability at least $1 - n^{-\theta}$ for some $\theta > 0$.

On \mathcal{A} , by definition, for each $1 \leq i \leq K$, $\eta_i^{(1)}$ satisfies the hypothesis of Proposition 3.5, and hence applying the latter with $r = c_0 \log n$ and a union bound we get

$$\mathbb{P}\left(\bigcap \mathcal{S}_i\right) \geq 1 - n e^{-hc_0 \log n} \geq 1 - n^{-\theta}$$

for some $h > 0$, and by choosing c_0 sufficiently large the final inequality holds for some $\theta > 0$. We can now infer (3.3) for all sufficiently large n , from Lemma 3.6, choosing C_0 sufficiently large compared to C_2 and C_3 and choosing b sufficiently small. □

3.3. Setting the traps

It remains to prove Proposition 3.5. For this proof we shall work with the Diaconis–Fulton representation of ARW on $[-\frac{r}{2}, \frac{r}{2}]$ using the stack of instructions \mathcal{J} as explained at the end of Section 2.

For the remainder of this section, we shall be in the set-up of Proposition 3.5. Also without loss of generality we shall assume that the total number of particles at the origin is $m = \mu r$, and we run ARW with sleep rate λ where $\mu < \frac{\lambda}{1+\lambda}$. Using Abelian property (Lemma 2.1), our goal will be to provide a toppling procedure (with some possibly ignored sleep instructions) that, when it succeeds, will lead to a stable configuration before any of the particles reach $\{-\frac{r}{2}, \frac{r}{2}\}$. Our job will be finished once we show that the algorithm succeeds with high probability. Next we describe in detail the steps of our algorithm, which employs a variant of the algorithm in [22, Section 5], along with a more complicated trap setting procedure in the finite setting. We elaborate on the differences from [22] and their necessity later, but first we describe the procedure.

3.3.1. Exploration, and locating the traps

Recall that the number of particles at the origin is m . Let us enumerate the particles y_1, y_2, \dots, y_m . The algorithm consists of applying a settling procedure to each particle. This procedure explores \mathcal{J} until it identifies a suitable trap for the particle. The exploration follows the path that the particle would perform if we always toppled the site it occupies, and stops when the trap has been chosen. In the absence of a suitable trap, we declare the algorithm to have failed. We set two initial barriers $a_0 = -\frac{r}{2}, b_0 = \frac{r}{2}$. We will recursively define barrier processes $a_0 < a_1 < \dots$ and $b_0 > b_1 > \dots$ that are functions of \mathcal{J} . Having defined a_i and b_i , we define a_{i+1} and b_{i+1} as follows:

Topple the particle y_{i+1} using the previously unexplored instructions in \mathcal{J} until it hits either a_i or b_i , that is, we always topple the site where this particle is currently located. At this stage we ignore all the sleep instructions. Eventually the particle hits either a_i or b_i , call the site hit q_i . Let us suppose $q_i = a_i$ and the exploration process hits q_i at step τ_i . We set $b_{i+1} = b_i$ and explore the accessed sites backwards from a_i , until we reach zero. If we reach a site v where the second to last instruction accessed was a sleep instruction (notice that the last instruction must have been a step towards a_i) that was ignored, then we set $a_{i+1} = v$, and call $a_{i+1} = \mathbf{Trap}_{i+1}$. See Figure 2. If no such site exists we declare the procedure to have failed. Observe that provided we can successfully set the barriers then they are moving towards the origin from both sides. We declare the trap setting scheme a success if $a_m < 0 < b_m$. Note that this is a sufficient condition for us to be able to set barriers for all the m particles.

In [22], the argument is based on considering the infinite half line and hence in that setting, it suffices to only consider a single barrier process $a_0 < a_1 < \dots$. In contrast, in our setting, we want the particles to not exit the interval $[-\frac{r}{2}, \frac{r}{2}]$, and hence do not want the situation where a certain random walk hits the barrier process $a_0 < a_1 < \dots$, after a long excursion outside the interval of our interest. This creates the necessity to have another barrier process, $\dots < b_1 < b_0$, which the particle would hit instead on such a journey, preventing its exit from the interval.




Fig. 2. This figure is similar to the one appearing in [22]. The blue, red and green paths are the exploration trajectories for three consecutive particles labelled y_k, y_{k+1}, y_{k+2} respectively, starting at the origin. The \leftarrow, \rightarrow denote jump instructions whereas X denotes a sleep instruction. The first particle is stopped on hitting the barrier a_{k-1} and hence it advances to a_k which is the closest site where the second last instruction was a sleep instruction ignored. The particle now falls asleep at a_k instead of exploring the path beyond a_k . As shown in Lemma 3.9, $a_k - a_{k-1}$ is dominated by a Geometric variable of mean $\frac{1+\lambda}{\lambda}$. However the barrier b_{k-1} stays as it is and is renamed b_k . The second particle hits barrier b_k which now advances to b_{k+1} where the particle uses the previously ignored sleep instruction to fall asleep, whereas a_k is renamed a_{k+1} . Thus the trap setting scheme proceeds to find traps for each individual particle to fall asleep.

3.3.2. Running the dynamics

Let us suppose that the realization of \mathcal{I} is such that the trap setting procedure is a success. On this event, let us give a toppling scheme which will utilize the traps to stabilize all the particles before they hit a_0 or b_0 . We topple the particles sequentially. Assuming the particles up to y_{i-1} have been settled, we start the toppling of the particle y_i , ignoring all the sleep instructions until the particle hits **Trap** _{i} for the last time before hitting a_{i-1} or b_{i-1} . Observe that because the instructions used were never used in the in exploration process of the previous particles, the path of this particle is the same as the exploration path, up until it hits **Trap** _{i} for the last time. We let the sleep instruction that was the second to last one accessed at **Trap** _{i} be executed and this settles the particle y_i at **Trap** _{i} . The key thing to notice here is that all the subsequent instructions accessed by the exploration process (but not in the actual dynamics) are located outside (a_i, b_i) , hence the future exploration processes will never try to access them by definition. This implies that the procedure can be continued with all the particles settled at their respective traps if the trap setting procedure succeeds, and additionally the consecutive exploration paths are independent of each other. This fact will be crucial for us when we try to estimate the growth of the barrier processes.

We summarize the upshot of the toppling procedure and the discussion above in the following lemma.

Lemma 3.7. *In the set-up of Proposition 3.5, suppose the trap setting procedure described above succeeds. Then S occurs.*

Proof. The proof is a straightforward consequence of Lemmas 2.3 and 2.1. □

Using Lemma 3.7, the proof of Proposition 3.5 will now be complete, if we show that the probability that the trap setting procedure fails is exponentially small in r . As mentioned before, our toppling scheme succeeds if $a_m < 0 < b_m$ where m is the total number of particles initially at 0. The proof of Proposition 3.5 follows from the next lemma.

Lemma 3.8. *Let $m = \beta r$ be the total number of particles at the origin, and consider the trap setting procedure described above. For $\beta < \frac{\lambda}{1+\lambda}$, there exists $c = c(\beta, \lambda) > 0$ such that for all sufficiently large r we have*

$$\mathbb{P}(a_m < 0 < b_m) \geq 1 - e^{-cr}.$$

Let us first provide a brief outline of our argument. Observe that at each stage i , exactly one of the barriers advances towards the origin. The probability of this being the left one or the right one is equal by symmetry at step 1 and hence equal to $\frac{1}{2}$. The proof now involves showing that it remains close to $\frac{1}{2}$ throughout. Also we show that the distance a barrier moves at each step has mean $\frac{1+\lambda}{\lambda}$. So the total distance covered by the barriers after βr many moves is approximately $\frac{1}{2}\beta r \frac{1+\lambda}{\lambda}$, which is smaller than $\frac{r}{2}$ because of the assumption on β and λ . Since the initial location of the barriers were at $-\frac{r}{2}$ and $\frac{r}{2}$, the above implies that none of the barriers cross the origin.

We now make the above argument formal. The first lemma we need is the following. A similar observation was already present in [22].

Lemma 3.9. *At the i th stage, the distance of \mathbf{Trap}_i from the barrier q_{i-1} hit by the i th exploration process is dominated by a geometric random variable with mean $1 + \frac{1}{\lambda}$ independent of everything else.*

Proof. Without loss of generality we assume that $q_{i-1} = a_{i-1}$. Recall from (2.2) that each instruction $\xi_{(x,j)}$ in \mathcal{I} is a sleep instruction with probability $\frac{\lambda}{1+\lambda}$ and jump instruction otherwise, independent of everything else. Thus, conditioning on the i th exploration path, the number of sleep instructions ignored at any site $x \in \mathcal{C}_n$ between successive jumps at x are i.i.d. random variables distributed as $\text{Geom}(\frac{1}{1+\lambda}) - 1$ (we adopt the standard notation of denote a Geometric random variable with mean p^{-1} by $\text{Geom}(p)$). In particular the number of sleep instructions between successive jumps is zero with probability $\frac{1}{1+\lambda}$. Thus at any site, the probability that there was a sleep instruction ignored before the last jump instruction is $\frac{\lambda}{1+\lambda}$. Thus $\mathbf{Trap}_i - a_{i-1}$ is dominated by $\text{Geom}(\frac{1+\lambda}{\lambda})$ variable, independent of everything else. Note that this is not a distributional identity, as $\mathbf{Trap}_i - a_{i-1}$ is bounded above by $r - a_{i-1}$. \square

The next step is to show that roughly half of the particles hit the barriers on either side. For $i = 1, 2, \dots, m$, let U_i denote the indicator that the i -th particle exploration process hits the left barrier first, that is, $q_{i-1} = a_i$. Also let $V_i = 1 - U_i$. We have the following lemma.

Lemma 3.10. *In the above set-up, for each $\delta > 0$, there exists $c = c(\delta) > 0$ such that*

$$\mathbb{P}\left(\sum_{i=1}^m U_i \geq \left(\frac{1}{2} + \delta\right)m\right) \leq e^{-cr}; \quad \mathbb{P}\left(\sum_{i=1}^m V_i \geq \left(\frac{1}{2} + \delta\right)m\right) \leq e^{-cr}.$$

The proof of Lemma 3.10 is involved and requires a somewhat complicated coupling to a different process; we postpone it to the next subsection. Using this, however, the proof of Lemma 3.8 is almost immediate and we complete that part of the argument now.

Proof of Lemma 3.8. Let X_1, X_2, \dots and Y_1, Y_2, \dots be two independent sequences of i.i.d. $\text{Geom}(\frac{\lambda}{1+\lambda})$ variables independent of the sequences $\{U_i\}_{i=1}^m$ defined above. It follows from Lemma 3.9 that a_m is stochastically dominated by $-\frac{r}{2} + \sum_{i=1}^m X_i U_i$, and similarly b_m stochastically dominates $\frac{r}{2} - \sum_{i=1}^m Y_i V_i$. Now clearly, using Lemma 3.3 it follows that for all $\delta, \delta' > 0$ there exists $c = (c, c') > 0$ such that

$$\mathbb{P}\left(\sum_{i=1}^{(\frac{1}{2}+\delta)m} X_i \geq (1 + \delta') \frac{1+\lambda}{\lambda} \left(\frac{1}{2} + \delta\right)m\right) \leq e^{-cr}.$$

Choosing δ and δ' sufficiently small so that $(1 + \delta') \frac{1+\lambda}{\lambda} (\frac{1}{2} + \delta)\beta < \frac{1}{2}$ (this is possible since $\beta < \frac{\lambda}{1+\lambda}$ and $m = \beta r$) it follows using Lemma 3.10 that $\mathbb{P}(a_m < 0) \geq 1 - e^{-cr}$. By symmetry an identical bound holds for $\mathbb{P}(0 < b_m)$ and we are done by taking a union bound. \square

3.4. Coupling with an internal erosion process

It only remains to prove Lemma 3.10. As alluded to before to this end we shall use a coupling to a process called internal erosion (see [18] for a nice exposition on the latter). Let X_1, X_2, \dots and Y_1, Y_2, \dots be two independent

sequences of i.i.d. $\text{Geom}(\frac{\lambda}{1+\lambda})$ variables. Let $S_i = \sum_{j=1}^i X_j$ and $T_i = \sum_{j=1}^i Y_j$ denote the sequence of partial sums. Let τ_1 (resp. τ_2) denote the largest positive integer i such that S_i (resp. T_i) is less than $\frac{r}{2}$. Now let $\{Z_i\}_{1 \leq i < \tau_1}$ (resp. $\{W_i\}_{1 \leq i < \tau_2}$) be a sequence of independent exponential random variables with means $f(i)$ (resp. $g(i)$) where $f(i) = \frac{r}{2} - S_i$ (resp. $g(i) = \frac{r}{2} - T_i$). Consider the two following continuous time counting processes:

$$N^{(1)}(t) = \sup \left\{ n \geq 0 : \sum_{i=1}^n Z_i \leq t \right\}; \quad N^{(2)}(t) = \sup \left\{ n \geq 0 : \sum_{i=1}^n W_i \leq t \right\}.$$

Set also $N(t) = N^{(1)}(t) + N^{(2)}(t)$. We shall crucially use the following connection of the above process with the trap setting procedure described in the previous subsection.

Lemma 3.11. *Consider the barrier processes $\{a_i\}$, $\{b_i\}$ and the internal erosion process described above. There is a coupling between the two processes satisfying the following: for all $t \geq 0$ such that $N^{(1)}(t) < \tau_1$ and $N^{(2)}(t) < \tau_2$, one has $a_{N(t)} = -\frac{r}{2} + \sum_{i=1}^{N^{(1)}(t)} X_i$ and $b_{N(t)} = \frac{r}{2} - \sum_{i=1}^{N^{(2)}(t)} Y_i$. Moreover, $N^1(t)$ is the number of times the barrier $a_0 < a_1 < \dots$ is hit among the first $N(t)$ particles.*

Proof. The proof is a consequence of the memoryless property of Exponential variables. Recall from the proof of Lemma 3.9 that the consecutive non-zero increments of the process $a_0 \leq a_1 \leq \dots$, are distributed as X_1, X_2, \dots , truncated at certain values which are functions of both the barrier processes. However note that while neither barrier process has reached zero, the issue of truncation does not arise. And hence we can couple the increments of the $\{a_i\}_{i \geq 1}$ exactly to the process $\{X_i\}_{i \geq 1}$ for the first τ_1 increments. Similarly the decrements of the process $\{b_i\}_{i \geq 1}$ can be coupled exactly to the process $\{Y_i\}_{i \geq 1}$ for the first τ_2 decrements (see Figure 2 for an illustration).

Thus to finish the proof of the lemma, we have to argue that the probability of the j th particle hitting the barrier a_{j-1} instead of b_{j-1} is the same as the process $N^1(t)$ increasing before $N^2(t)$ when $N(t) = j - 1$ for any j such that both $N^1(t) < \tau_1$ and $N^2(t) < \tau_2$. Note that the probability of the $(N(t) + 1)$ th particle hitting the barrier $a_{N(t)}$ as opposed to $b_{N(t)}$ has probability

$$\frac{b_{N(t)}}{b_{N(t)} - a_{N(t)}} = \frac{g(N^{(2)}(t))}{f(N^{(1)}(t)) + g(N^{(2)}(t))}. \quad (3.4)$$

Note that given the filtration up to time t , N_t^1 increases before $N^2(t)$ if and only if

$$\sum_{i=1}^{N^1(t)+1} Z_i - t \leq \sum_{i=1}^{N^1(t)+1} W_i - t.$$

Now given the filtration up to time t , using the memoryless property, it follows that

$$\sum_{i=1}^{N^1(t)+1} Z_i - t \text{ is distributed as } Z_{N^1(t)+1},$$

and similarly $\sum_{i=1}^{N^2(t)+1} W_i - t$ is distributed as $W_{N^2(t)+1}$. Thus using the fact that

$$\mathbb{P}(Z_{N^1(t)+1} < W_{N^2(t)+1}) = \frac{g(N^{(2)}(t))}{f(N^{(1)}(t)) + g(N^{(2)}(t))},$$

the proof is complete using (3.4). □

Using Lemma 3.11 to prove Lemma 3.10, it suffices to prove the following lemma. Recall that $m = \beta r$ is the total number of particles.

Lemma 3.12. Let $\beta < \frac{\lambda}{1+\lambda}$ and let $\delta > 0$ be fixed (and sufficiently small as a function of $\frac{\lambda}{1+\lambda} - \beta$). Let $M_1 = (\frac{1}{2} - \delta)\beta r$, $M_2 = (\frac{1}{2} + \delta)\beta r$. Let \mathcal{E} denote the event that there exists t such that $\{N^{(1)}(t) \geq M_2, N^{(2)}(t) \leq M_1\}$ or $\{N^{(2)}(t) \geq M_2, N^{(1)}(t) \leq M_1\}$. Then there exists $c > 0$ such that $\mathbb{P}(\mathcal{E}) \leq e^{-cr}$.

Proof. First observe that, by Lemma 3.3, $\max(S_{M_2}, T_{M_2}) \leq \frac{r}{2}$ (since δ is sufficiently small) with exponentially small failure probability and hence we have $\min(\tau_1, \tau_2) \geq M_2$ with exponentially small failure probability. Thus we can safely restrict our analysis to the latter event. Observe next that it suffices to prove that with exponentially (in r) small failure probability, we have

$$\sum_{i=1}^{M_2} Z_i > \sum_{i=1}^{M_1} W_i; \quad \sum_{i=1}^{M_2} W_i > \sum_{i=1}^{M_1} Z_i. \tag{3.5}$$

Conditional on the sequences S_i and T_i , the concentration estimates in Lemma 3.3 imply that the terms $A_1 := \sum_{i=1}^{M_1} Z_i$, $A_2 := \sum_{i=1}^{M_1} W_i$, $A_3 := \sum_{i=1}^{M_2} Z_i$, $A_4 := \sum_{i=1}^{M_2} W_i$ are all concentrated near their means p_1, p_2, p_3, p_4 , with exponentially small failure probability. The proof is then essentially completed by comparing the means. Note that the means are themselves random (functions of S_i and T_i) and hence the last detail is to show that the means themselves are concentrated.

Formally we first observe that $p_1 = M_1 \frac{r}{2} - \sum_{i=1}^{M_1} i X_{M_1-i+1}$ and similar expressions holds for p_2, p_3 and p_4 . Note that $\mathbb{E}(p_1) = M_1 \frac{r}{2} - \frac{1+\lambda}{\lambda} \frac{M_1^2}{2} + O(M)$. There are several ways to prove concentration of p_1 and below we sketch a way to use Lemma 3.3 to achieve this. Note that the latter only allows for sums of geometric variables, whereas we have a linear combination of them. Since we can afford to be rather crude with our estimates, we use the following decomposition

$$p_1 = M_1 \frac{r}{2} - \sum_{i=1}^{M_1} \sum_{j=1}^i X_j.$$

Thus by union bound, after applying Lemma 3.3 to each of the terms of form $\sum_{j=1}^i X_j$, it follows that: for all ε_1 small enough, there exists c depending on all the parameters except r , such that

$$\mathbb{P}(|p_1 - \mathbb{E}(p_1)| \geq \varepsilon_1 r^2) < e^{-cr}.$$

Similar analysis allows us to conclude similar bounds as above for p_2, p_3, p_4 . By choice of M_1 and M_2 , note that there exists ε_1 such that $\mathbb{E}(p_3) - \mathbb{E}(p_1) \geq 4\varepsilon_1 r^2$ and similarly $\mathbb{E}(p_4) - \mathbb{E}(p_2) \geq 4\varepsilon_1 r^2$. Thus we see that with probability at least $1 - e^{-cr}$, the sequences S_i, T_i are such that

$$p_3 - p_1 > 2\varepsilon_1 r^2 \quad \text{and} \quad p_4 - p_2 > 2\varepsilon_1 r^2.$$

Moreover, conditioned on the above events, for $j = 1, 2, 3, 4$, Lemma 3.3 implies the following concentration estimates:

$$\mathbb{P}\left(|A_j - p_j| \geq \frac{\varepsilon_1}{2} r^2\right) < e^{-cr}. \tag{3.6}$$

Thus combining the above inequalities and union bound the lemma follows. □

We are now ready to complete the proof of Lemma 3.10.

Proof of Lemma 3.10. We shall use the coupling described in Lemma 3.11. Let M_1, M_2 be as in Lemma 3.12. Now by the coupling discussed above and Lemma 3.12, it follows that with exponential (in r) failure probability, M_1 particles hit both barriers before M_2 particles hit any barrier.

Since the total number of particles is $m \leq M_1 + M_2$ it follows that, with exponentially high probability neither $\sum_{i=1}^m U_i$ nor $\sum_{i=1}^m V_i$ can exceed M_2 . Since $\beta < \frac{\lambda}{1+\lambda}$, we can safely ignore the exponentially (in r) unlikely event that $M_2 \geq \min(\tau_1, \tau_2)$ and hence assume that coupling in Lemma 3.11 does not fail. □

Remark 3.13. It can be shown that $T_n(\mu, \lambda)$ will be of order n^3 , if all the particles (approximately μn) were initially located at the same site and hence Theorem 1 relies heavily on the location of the particles in the initial configuration being uniform. To see this, note that by the above discussion regarding topplings, when a linear in n , say αn , number of particles start at the origin, then to stabilize, due to lack of space, at least $\frac{\alpha n}{2}$ particles must move outside an interval of size $\frac{\alpha n}{2}$ centred at the origin. Since a random walk path takes time $\Theta(n^2)$ on average to exit such an interval, the observation follows.

4. Slow fixation for low sleep rate

In this section we prove Theorem 2. That is, we prove that for any $\mu > 0$ and sufficiently small sleep rate λ , $\text{ARW}(\mu, \lambda)$ on \mathcal{C}_n takes at least exponentially many steps before reaching the absorbing state with failure probability exponentially small in n .

4.1. The stabilization loop

We shall now describe the toppling scheme outlined in Section 1.2 in more detail. Let $\mu \in (0, 1)$ be fixed and λ be sufficiently small. By Lemma 2.1 it suffices to exhibit a sequence of legal topplings with exponentially many steps. We shall show that our procedure satisfies this property with exponentially high probability if λ is sufficiently small.

While running this scheme, particles will switch between two different states, which we call states X and Y . Particles in state X follow normal ARW dynamics among themselves as described in Section 2.1, and can wake up sleeping Y -particles. Y -particles, on the other hand, do not move, and have no effect on the states of any other particles. Thus a state of the system during this toppling scheme consists of all the particles, each in state X or state Y , and each asleep or awake. When a site is toppled, only X -particles at that site follow the corresponding stack instructions from \mathcal{I} (see Section 2.1). So Y -particles only undergo the transition from sleepy to active when an active X -particle reaches the same site.

As described in Section 1.2, our toppling procedure will run multiple rounds of what we call stabilization loop. Formally, starting from a particle configuration η – that is, the values $\eta_i(x)$ for $x \in \mathcal{C}_n$ – stabilizing the system in a subset \mathcal{D} of sites in \mathcal{C}_n means choosing some particles to be in state X and the rest to be in state Y . Then running the particle dynamics described above, only toppling sites inside \mathcal{D} , until all X -particles are asleep in \mathcal{D} or are outside \mathcal{D} . Now the obvious strong Markov property of the above dynamics makes the different stabilization rounds conditionally independent which would be crucial in our calculations.

It will be convenient to identify \mathcal{C}_n with the interval $[-r, r]$ with $-r$ and r identified, that is, assume $n = 2r$ for integer r . We shall denote the origin by $\mathbf{0}$ and the identified vertex $r = -r$ will be denoted by \mathbf{r} . The stabilization steps we run will alternate between taking $\mathcal{D} = \mathcal{C}_n \setminus \{\mathbf{0}, \mathbf{r}\}$, $\mathcal{D} = \mathcal{C}_n \setminus \{\mathbf{0}\}$ and $\mathcal{D} = \mathcal{C}_n \setminus \{\mathbf{r}\}$.

1. **Stabilization Step A:** Stabilize all the particles in $\mathcal{C}_n \setminus \{\mathbf{0}, \mathbf{r}\}$. That is, treat particles in $\mathcal{C}_n \setminus \{\mathbf{0}, \mathbf{r}\}$ as X -particles, stopped upon hitting $\{\mathbf{0}, \mathbf{r}\}$, and all other particles as Y -particles. So at the end of this procedure all active particles will be at $\mathbf{0}$ or \mathbf{r} .
2. **Stabilization Step B:** Reset the X and Y labels: the particles initially at $\mathbf{0}$ become X -particles, and all other particles become Y -particles. Then stabilize all X -particles with particles stopped at \mathbf{r} . With the identification of \mathcal{C}_n with $[-r, r]$, this step is the same as stabilizing the ARW dynamics in the interior of $[-r, r]$ where the initial particle configuration is supported at the center of the interval, a special case of the more general process analyzed later in Lemma 4.5 using results from [2].
3. **Stabilization Step C:** This is identical to the **Stabilization Step B** above with the roles of $\mathbf{0}$ and \mathbf{r} interchanged.

The algorithm receives an initial particle configuration η on \mathcal{C}_n drawn from \mathbb{P}^μ as an input. Then we perform the **Stabilization Loop**, which is **Stabilization Step A**, followed by **Stabilization Step B** and **Stabilization Step C**. We repeat the **Stabilization Loop** until all of the particles are asleep.

We now state the main lemma about the **Stabilization Loop** and use it to prove Theorem 2.

Lemma 4.1. Fix $\mu \in (0, 1)$, $\varepsilon \in (0, 2\mu/5)$ and any particle configuration η with at least $(\mu - \varepsilon)n$ particles, and suppose at least $\mu n/2$ particles are active in η . Let $\tilde{\eta}$ denote the configuration after we have performed the **Stabilization**

Loop. Then for $\lambda = \lambda(\mu) > 0$ sufficiently small, there exists $c > 0$ such that

$$\mathbb{P}(\tilde{\eta} \text{ has less than } \mu n/2 \text{ active particles}) < e^{-cn}.$$

Using this lemma we now prove Theorem 2.

Proof of Theorem 2. By the Abelian property it suffices to demonstrate that with high probability there is a toppling algorithm that does not terminate before exponentially many steps are executed. By a Chernoff bound for any μ and ε the probability that there are at least $(\mu - \varepsilon)n$ particles is exponentially close to one. In the initial stage at least $\mu n/2$ of the particles are awake. By Lemma 4.1, the number of consecutive rounds that the **Stabilization Loop** is performed is at least $e^{\frac{\varepsilon}{2}r}$ with probability at least $1 - e^{-\frac{\varepsilon}{2}r}$ before all the particles are asleep. Each time the **Stabilization Loop** is performed, there must be at least one jump or sleep instruction occurring in it. As $n = 2r$ this completes the proof of Theorem 2. \square

To prove Lemma 4.1 we rely on two results that are proved by adapting the analysis in [2]. Our first goal is to show that with exponentially high probability after performing **Stabilization Step A** there are at least $\mu n/4$ active particles. By definition these active particles are all at $\mathbf{0}$ or \mathbf{r} .

Lemma 4.2. Fix $\mu \in (0, 1)$, $\varepsilon \in (0, 2\mu/5)$ and any particle configuration η with at least $(\mu - \varepsilon)n$ particles of which at least $\mu n/2$ are active.

Let $\tilde{\eta}^A$ denote the configuration after we have performed the **Stabilization Step A**. Then for $\lambda = \lambda(\mu) > 0$ sufficiently small, there exists $c > 0$ such that

$$\mathbb{P}(\tilde{\eta}^A \text{ has less than } \mu n/4 \text{ active particles}) < e^{-cn}.$$

Thus after performing **Stabilization Step A** there are likely to be either at least $\mu n/8$ active particles at $\mathbf{0}$ or at least $\mu n/8$ active particles at \mathbf{r} . Suppose there at least $\mu n/8$ active particles at $\mathbf{0}$. Our next result says that with high probability after running **Stabilization Step B** at least 90% of active particles that were at $\mathbf{0}$ are now active particles at \mathbf{r} . Also all Y particles are now active.

Lemma 4.3. Fix $\mu > 0$ and any particle configuration η with at least $A \geq \mu n/8$ active particles at $\mathbf{0}$.

Let $\tilde{\eta}^B$ denote the configuration after we have performed **Stabilization Step B**. Then for $\lambda = \lambda(\mu) > 0$ sufficiently small, there exists $c > 0$ such that

$$\mathbb{P}(\tilde{\eta}^B \text{ has all } Y \text{ particles active and at least } 0.9A \text{ active particles at } \mathbf{r}) > 1 - e^{-cn}.$$

We postpone the proofs of Lemma 4.2 and Lemma 4.3 for now and complete the proof of Lemma 4.1.

Proof of Lemma 4.1. After performing **Stabilization Step A**, by Lemma 4.2 there are at least $\mu n/4$ active particles at either $\mathbf{0}$ or \mathbf{r} except exponentially small failure probability. On this high probability event, after performing **Stabilization Step B** then except for exponentially small failure probability there are at least $\mu n/8$ active particles at \mathbf{r} by Lemma 4.3. (Note that this is true no matter how the active particles were split among $\mathbf{0}$ and \mathbf{r} at the start of **Stabilization Step B**.) On the event that, the high probability event happens at both of these steps, when we start **Stabilization Step C** there are at least $\mu n/8$ active particles at \mathbf{r} . Thus we can apply Lemma 4.3, and it follows that, again except for exponentially small failure probability, at the end of **Stabilization Step C** all Y particles are active and at least 90% of the X particles are active. Since, by hypothesis, the total number of particles is at least $(\mu - \varepsilon)n$ and $\varepsilon < 2\mu/5$, easy algebra shows that on the event that none of the three steps resulted in the failure events of exponentially small probability, the loop results in a particle configuration with at least $\mu n/2$ active particles. The lemma now follows by taking a union bound over the three failure events. (Note that implicitly we are using the obvious strong Markov property of the above dynamics.) \square

4.2. Proving Lemmas 4.2 and 4.3

Now we show how to derive Lemmas 4.2 and 4.3 from [2]. The following lemma will be the key in both arguments.

Lemma 4.4. Fix $\delta_0 \in (0, 1)$ and $c_0 > 0$. Consider ARW with sleep rate λ on the interval $[-r, r]$ starting from an initial configuration η with at least $\delta_0 r$ active particles. Let S denote the number of sleepy particles in $(-r, r)$ after stabilizing. There exists $C = C(\delta_0, c_0)$ and $\lambda_0 = \lambda_0(\delta_0, c_0) > 0$ such that for all $\lambda < \lambda_0$

$$\mathbb{P}(S \geq c_0 r) \leq e^{-Cr}.$$

for all r sufficiently large.

Proof. This lemma comes from calculations contained in [2]. Recall the odometer function from (2.3) and let $h(0)$ denote the odometer at the origin at the end of the stabilization process and let $\mathcal{E} = \{h(0) \leq r^6\}$. We break up $\{S \geq c_0 r\}$ in two parts

$$\{S \geq c_0 r\} \subset (\{S \geq c_0 r\} \cap \mathcal{E}) \cup \mathcal{E}^c.$$

First we use random walk estimates to prove $\mathbb{P}(\mathcal{E}^c) \leq e^{-Cr}$. This is the same as the argument presented in [2] (Lemma 32) but we provide the short proof for completeness. To start, note that the probability that a lazy random walk started arbitrarily inside $[-r, r]$ does not hit $\{-r, r\}$ in Kr^2 steps is at most $e^{-Kc'}$, where c' depends on the laziness parameter. Now there are at most $2r$ particles, each of which moves along an independent, $\frac{\lambda}{1+\lambda}$ -lazy, random walk trajectory. So the probability that any of these particles take more than $0.5r^5$ steps before hitting $\{-r, r\}$ is exponentially small. A union bound then implies that the sum of the number of steps taken by all the particles before reaching $\{-r, r\}$ is exponentially unlikely to be more than r^6 and hence we get

$$\mathbb{P}(\mathcal{E}) = \mathbb{P}(h(0) \leq r^6) \geq 1 - e^{-Cr}.$$

Fix c_1 such that $c_0/4 > c_1 > 0$. Equation (6.21) in [2] shows for λ sufficiently small (and r sufficiently large)

$$\mathbb{E}((e^{S_1} + e^{S_2})\mathbf{1}_{\mathcal{E}}) \leq e^{c_1 r},$$

where S_1 and S_2 denote the number of sleepy particles in $(-r, 0]$ and $[0, r)$ respectively. Thus by Markov's inequality,

$$\mathbb{P}(\{S \geq c_0 r\} \cap \mathcal{E}) \leq \frac{\mathbb{E}((e^{S_1} + e^{S_2})\mathbf{1}_{\mathcal{E}})}{e^{\frac{c_0}{2}r}} \leq \frac{2e^{c_1 r}}{e^{\frac{c_0}{2}r}} \leq e^{-Cr}$$

for some $C > 0$, since $S > c_0 r$ implies either S_1 or S_2 is at least $\frac{c_0 r}{2}$. □

Proof of Lemma 4.2. If there are at most $\mu n/4$ active particles after **Stabilization Step A** then at least $\mu n/4$ particles must have fallen asleep in this step. Thus either

- (1) there were initially at least $\mu n/5$ active particles on $(0, r)$ in η and there were at least $\mu n/20$ sleepy particles on $(0, r)$ in $\tilde{\eta}^A$ or
- (2) there were initially at least at least $\mu n/5$ active particles on $(-r, 0)$ in η and there were at least $\mu n/20$ sleepy particles on $(-r, 0)$ in $\tilde{\eta}^A$.

By Lemma 4.4 both of those events have exponentially small probability. Thus the lemma follows from an union bound over the two cases. □

For the proof of Lemma 4.3 we shall use the following lemma which is immediate from Lemma 4.4 by taking c_0 sufficiently small and we omit the proof.

Lemma 4.5. For each $\delta_0 > 0$, the following holds for λ sufficiently small. Consider stabilizing any particle configuration η supported on $[-r, r]$, that is, particles hitting $\{-r, r\}$ are ignored. Call the stabilized system η' . If η has A many active particles with $A \geq \delta_0 r$ then

$$\mathbb{P}(\text{the total number of particles supported on } \{-r, r\} \text{ in } \eta' \text{ is at least } 0.9A) \geq 1 - e^{-cr}.$$

Proof of Lemma 4.3. For this lemma we have to show the two events have exponentially small failure probability: (i) $\tilde{\eta}^B$ has at least $0.9A$ active particles at \mathbf{r} , and (ii) all Y particles are active. The first part is immediate from Lemma 4.5. For (ii), observe the following.

Each X -particle, whenever it moves, follows an independent random walk trajectory that is stopped at hitting $\{-r, r\}$ (the number of steps in this trajectory that is realized has a complicated dependent structure). By symmetry each of these trajectories are equally likely to end at $-r$ and r . Since $A \geq \mu n/8$ and using a standard Chernoff bound, it follows that with failure probability exponentially small in n , no more than $0.6A$ of these trajectories end at r (also at $-r$). Since by the first part, we know that at least $0.9A$ of these particles follows their trajectories to hit $\{-r, r\}$, this implies that except for exponentially small failure probability, both r and $-r$ are hit by at least $0.3A$ many X -particles, that is, \mathbf{r} is hit from both positive and negative side. This implies that every Y particle is hit by an X particle, and hence all Y particles are active in $\tilde{\eta}^B$. \square

We finish with a remark on lower bounding $T_n(\mu, \lambda)$.

Remark 4.6. An exponential upper bound for $T_n(\mu, \lambda)$ is relatively easy to establish. Note that starting from any configuration, ignoring sleep instructions, one can get each particle to a different location, using only polynomial in n many instructions with probability at least $1 - e^{-cn}$, since for a random walk on \mathcal{C}_n , the hitting time for any point is a sub-exponential variable at scale n^2 . Once the particles are all located at different sites, with probability at least $(\frac{\lambda}{1+\lambda})^n$, all of their next instructions are sleep instructions which causes the system to stabilize. Thus starting from any configuration, the probability of stabilizing in polynomially many steps is at least $(\frac{\lambda}{1+\lambda})^n$, which implies an exponential upper bound on the fixation time, by running independent trials of the above argument until one trial does succeed to stabilize the system.

Acknowledgements

The authors thank Lionel Levine, Vladas Sidoravicius and Lorenzo Taggi for useful conversations, and Leonardo Rolla for comments on a previous version of the paper. Research of R. Basu is partially supported by an ICTS-Simons Junior Faculty Fellowship and a Ramanujan Fellowship (SB/S2/RJN-097/2017) from the Govt. of India. S. Ganguly is supported by a Miller Research Fellowship. C. Hoffman is supported by a Simons Fellowship and NSF Grant DMS-1712701.

References

- [1] G. Amir and O. Gurel-Gurevich. On fixation of activated random walks. *Electron. Commun. Probab.* **12** (2010) 119–123. [MR2643591](#)
- [2] R. Basu, S. Ganguly and C. Hoffman. Non-fixation of symmetric activated random walk on the line for small sleep rate. *Comm. Math. Phys.* To appear, 2018. Preprint. Available at [arXiv:1508.05677](#).
- [3] P. Diaconis and W. Fulton. A growth model, a game, an algebra, Lagrange inversion, and characteristic classes. *Rend. Semin. Mat. Univ. Politec. Torino* **49** (1) (1991) 95–119. [MR1218674](#)
- [4] R. Dickman. Nonequilibrium phase transitions in epidemics and sandpiles. *Phys. A* **306** (2002) 90–97. [MR1933220](#)
- [5] R. Dickman, M. Alava, M. A. Muñoz, J. Peltola, A. Vespignani and S. Zapperi. Critical behaviour of a one-dimensional fixed-energy stochastic sandpile. *Phys. Rev. E* **64** (2001), 56104.
- [6] R. Dickman, M. A. Muñoz, A. Vespignani and S. Zapperi. Paths to self-organized criticality. *Braz. J. Phys.* **30** (2000) 27–41.
- [7] R. Dickman, A. Vespignani and S. Zapperi. Self-organized criticality as an absorbing-state phase transition. *Phys. Rev. E* **57** (1998) 5095–5105. [MR1628210](#)
- [8] R. Dickman, L. T. Rolla and V. Sidoravicius. Activated random walkers: Facts, conjectures and challenges. *J. Stat. Phys.* **138** (1–3) (2010) 126–142. [MR2594894](#)
- [9] K. Eriksson. Chip-firing games on mutating graphs. *SIAM J. Discrete Math.* **9** (1) (1996) 118–128. [MR1375419](#)

- [10] A. Fey, L. Levine and D. B. Wilson. Driving sandpiles to criticality and beyond. *Phys. Rev. Lett.* **104** (14) (2010) Article ID 145703.
- [11] B. Hough, D. Jerison and L. Levine. Sandpiles on the square lattice. Preprint, 2017. Available at [arXiv:1703.00827](https://arxiv.org/abs/1703.00827).
- [12] S. Janson. Tail bounds for sums of geometric and exponential random variables. *Statist. Probab. Lett.* **135** (2018) 1–6. [MR3758253](#)
- [13] H. Kesten and V. Sidoravicius. Branching random walk with catalysts. *Electron. J. Probab.* **8** (2003) 1–51. [MR1961167](#)
- [14] H. Kesten and V. Sidoravicius. The spread of a rumor or infection in a moving population. *Ann. Probab.* **33** (2005) 2402–2462. [MR2184100](#)
- [15] H. Kesten and V. Sidoravicius. A phase transition in a model for the spread of an infection. *Illinois J. Math.* **50** (2006) 547–634. [MR2247840](#)
- [16] H. Kesten and V. Sidoravicius. A shape theorem for the spread of an infection. *Ann. of Math. (2)* **167** (2008) 701–766. [MR2415386](#)
- [17] L. Levine. Threshold state and a conjecture of Poghosyan, Poghosyan, Priezzhev and Ruelle. *Comm. Math. Phys.* **335** (2) (2015) 1003–1017. [MR3316648](#)
- [18] L. Levine and Y. Peres. Internal erosion and the exponent $3/4$. Available at <http://www.math.cornell.edu/~levine/erosion.pdf>.
- [19] S. Lübeck. Universal scaling behavior of non-equilibrium phase transitions. *Internat. J. Modern Phys. B* **18** (31–32) (2004) 3977–4118.
- [20] S. S. Manna. Large-scale simulation of avalanche cluster distribution in sand pile model. *J. Stat. Phys.* **59** (1990) 509–521.
- [21] S. S. Manna. Two-state model of self-organized criticality. *J. Phys. A: Math. Gen.* **24** (1991) L363–L369.
- [22] L. T. Rolla and V. Sidoravicius. Absorbing-state phase transition for driven-dissipative stochastic dynamics on \mathbb{Z} . *Invent. Math.* **188** (1) (2012) 127–150. [MR2897694](#)
- [23] E. Shellef. Nonfixation for activated random walks. *ALEA Lat. Am. J. Probab. Math. Stat.* **7** (2010) 137–149. [MR2651824](#)
- [24] V. Sidoravicius and A. Teixeira. Absorbing-state transition for stochastic sandpiles and activated random walks. *Electron. J. Probab.* **22** (33) (2017) Article ID 22. [MR3646059](#)
- [25] A. Stauffer and L. Taggi. Critical density of activated random walks on transitive graphs. *Ann. Probab.* **46** (2018) 2190–2220. [MR3813989](#)
- [26] L. T. Rolla and L. Tournier. Non-fixation for biased activated random walks. *Ann. Inst. Henri Poincaré Probab. Stat.* **54** (2) (2018) 938–951. [MR3795072](#)
- [27] L. Taggi. Absorbing-state phase transition in biased activated random walk. *Electron. J. Probab.* **21** (2016) Article ID 13. [MR3485355](#)
- [28] A. Vespignani, R. Dickman, M. A. Muñoz and S. Zapperi. Driving, conservation, and absorbing states in sandpiles. *Phys. Rev. Lett.* **81** (1998) 5676–5679.
- [29] A. Vespignani, R. Dickman, M. A. Muñoz and S. Zapperi. Absorbing-state phase transitions in fixed-energy sandpiles. *Phys. Rev. E* **62** (2000) 4564–4582.
- [30] R. Vidigal and R. Dickman. Asymptotic behavior of the order parameter in a stochastic sandpile. *J. Stat. Phys.* **118** (1) (2005) 1–25. [MR2122547](#)